Table of Contents

1. Eyeglass PowerScale Edition Admin Guide..... 8
1.1. What's New with DR Product Releases..... 10
1.2. Introduction to this Guide..... 12
1.3. Eyeglass Ports Requirements , Scalability Limits and Phone Home Requirements..... 14
1.4. License Management with Eyeglass..... 25
1.5. Eyeglass Deployment Options..... 27
1.6. Eyeglass Backup Options..... 31
1.7. Eyeglass Jobs..... 33
1.8. Configuration Replication..... 50
1.9. Advanced Quota Replication..... 57
1.10. Manage Eyeglass Jobs..... 60
1.11. Failover Readiness Validations DR Dashboard..... 73
1.12. RPO Reporting and Trending Feature Guide..... 122
1.13. Role Based Access Controls RBAC..... 130
1.14. Configure Email, Twitter, Slack, Webhooks for Notifications of Eyeglass Monitoring Events..... 131
1.15. ECA (Eyeglass Clustered Agent) CLI Commands..... 140
1.16. DR Design Guides with Eyeglass..... 148
1.17. Eyeglass CLI Commands..... 149
1.18. Eyeglass Appliance Time Synchronization Best Practice..... 218
1.19. Eyeglass Automatic Updates for Recommended Packages..... 221
1.20. Update Eyeglass Appliance Network Settings..... 225
1.21. Appliance Security Updates and Eyeglass Updates with HTTP Proxy..... 227
1.22. Diagnostic Tools for Dark Site Support..... 229
1.23. Role Based Access Control And Authentication Guide..... 230
1.23.1. Overview..... 231
1.23.2. RBAC Requirements - Read Me First..... 237
1.23.3. RBAC Quick Start Steps..... 241
1.23.4. Simple Setup AD Group based RBAC..... 242
1.23.5. How to create a new Eyeglass Role..... 248
1.23.6. How to Login with RBAC..... 250
1.24. Eyeglass Alarm forwarding Guide - Syslog and Legacy SNMP..... 252
1.25. How to Setup Email alarms with Exchange Server..... 276
1.26. How to Change PowerScale IP Address in Eyeglass..... 281
1.27. Pre Post Failover Scripting Guide..... 284
1.28. Eyeglass Backup and Restore..... 307
1.29. Eyeglass API guide..... 312
1.29.1. How to build a Web server solution with the Eyeglass API..... 371
1.30. How to convert VMware Eyeglass appliance and Migrate to Microsoft Azure..... 385

1.31. TLS Certificate Procedures for Eyeglass....399
1.32. Configuring Eyeglass with dual NIC environments....411
1.33. Changing Eyeglass UI Behavior....413
1.34. Custom email routing by application or alarm subject contents....416
2. Eyeglass Ransomware Admin Guide....438
2.1. What's New....439
2.2. Introduction to this Guide....446
2.3. Planning and Design....450
2.3.1. How to determine threat response settings to meet your Company's Risk Profile....458
2.3.2. Eyeglass User Lockout Active Directory Planning....464
2.3.3. Ransomware Audit Events Required for all Deployments....467
2.3.4. Well Known Ransomware File Extension Whitelist....468
2.3.5. Security Event Descriptions....473
2.3.6. NFS Lockout Feature....480
2.3.7. Planning New application Workloads Best Practice....481
2.4. Everything about Detection, Configuration and Tuning Security Event....484
2.4.1. Ransomware - Threat Detection Settings Summary Explanation....485
2.4.2. Tuning Ransomware Defender Detections for Your Environment....490
2.4.3. Ransomware - Security Event Explanations....493
2.4.4. Best Practice for Tuning Ransomware Defender....499
2.5. How to Configure, Tune and View Ransomware Defender Threat Detection settings and Responses....504
2.5.1. How to Configure Monitor Mode and Ignored Lists....515
2.5.2. How To Manage False Positives and Learning Mode....529
2.5.3. Banned and Allowed File Type Configuration....537
2.5.4. Rapid Machine to Machine Malware Spreading Attack Defense....546
2.5.5. How to Manage Threat Detectors - Advanced Consult Support....548
2.5.6. How to Configure Snapshot Modes (Critical Path and SMB share snapshots) and Snapshot Quotas....551
2.6. How to respond to Security Events for Warning, Major or Critical Events....554
2.6.1. If Warning:....559
2.6.2. If Major:....560
2.6.3. If Critical:....562
2.6.4. Event State Descriptions....563
2.7. Operational Procedures For Common Tasks....573
2.7.1. How to login and Manage Ransomware Defender....575
2.7.2. Security Guard - Automated Security Testing....576
2.7.3. How To Configure HoneyPot file Tripwire....590
2.7.4. How to Flag a detection as False Positive....595
2.7.5. How to Enable or Disable Enforcement Mode....597

2.7.6. How to unlock a User that was locked out....600
2.7.7. Ransomware Defender Security Event Workflow for Warning Severity Detections....602
2.7.8. How to Recover Data after a Ransomware Attack....604
2.7.9. How to Enable Learning Mode and Monitor Learning mode Results.... 609
2.8. AirGap 2.0 Guide.... 614
2.9. How to Configure a Dell ECS and Data Protection Use Cases....688
3. Eyeglass Easy Auditor admin guide....694
3.1. What's New.... 696
3.2. Introduction to Easy Auditor Guide.... 704
3.3. Key Features.... 706
3.4. PowerScale Recommended Audit Event Configuration....711
3.5. Feature Limits.... 713
3.6. How to get started with Auditing....714
3.7. Who Audits the Auditor?.... 716
3.8. Planning, Design and How to Use Easy Auditor to Audit.... 717
3.9. How to Use Excel for advanced filtering of CSV Reports....791
3.10. How to Backup and Restore an Audit Database....795
3.11. Backup and DR for Audit Database with SyncIQ to a Remote Cluster....799
3.12. How to check Analytics database size....803
3.13. How to Use Easy Auditor for Typical Audit Use Cases....804
3.14. Audit Message Workflows.... 810
3.15. Advanced Configuration.... 822
3.16. Excluded Audit Events....828
3.17. How to Configure Syslog Forwarding of Formatted Audit messages to an External Syslog Server829
3.18. Data Retention of Audit Data and Archive.... 845
3.19. How to Purge or Archive PowerScale Audit logs....848
3.20. Bulk Ingest old Audit Data from PowerScale to Easy Auditor....851
4. LiveOps - Continuous Operations Admin Guide....857
4.1. Overview....858
4.1.1. LiveOps Continuous Operations Key Features.... 859
4.1.2. What’s New.... 860
4.2. LiveOPS Continuous Operations Dashboard.... 861
4.2.1. What do the columns Mean?.... 862
4.3. Overview - LiveOPS Snapshot Sync.... 863
4.3.1. How it Works.... 864
4.3.2. How to Enable....865
4.4. Overview - LiveOPS Dedupe Sync.... 867
4.4.1. How it Works.... 868
4.4.2. How to Enable....869

4.5. Overview - LiveOPS DR Test mode.... 870
4.5.1. DR Test Mode Architecture Diagram.... 871
4.5.2. Operating View.... 873
4.5.3. Prerequisites.... 875
4.5.4. How to Configure DR Test Mode.... 876
4.5.5. Isolated DR Test Mode High level Guide.... 883
4.6. How to Enable DR Test Mode.... 886
4.7. How to Disable DR Test Mode.... 892
4.8. How DR Test mode Jobs are displayed in Eyeglass UI.... 893
4.8.1. Jobs UI DR Test mode:.... 894
4.8.2. DR Assistant for DR Test mode.... 895
4.8.3. DR Dashboard for DR Test mode:.... 896
4.9. Advanced DR Test mode Configurations.... 898
4.9.1. Reasons for multiple DR test mode policies:.... 899
4.9.2. Procedures to DR Test different data sets independently.... 900
4.9.3. DR Test Mode States.... 903
5. Eyeglass Clustered Agent Admin Guide.... 904
5.1. What's New.... 905
5.2. ECA Cluster Topology Deployment Options.... 907
5.3. High Availability and Resilience.... 910
5.4. Eyeglass Active Directory User SID resolution considerations.... 912
5.5. How to monitor remote ECA clusters from Eyeglass.... 915
5.6. ECA Cluster Monitoring Tools.... 918
5.7. ECA Cluster Operational Procedures.... 920
5.8. ECA CLI Command Guide.... 928
5.8.1. How to run commands across all nodes.... 930
5.9. ECA Cluster Disaster Recovery and Failover Supported Configurations.... 931
5.10. How to Change Performance with VMware.... 943
5.11. Dark Sites - How to Health check Eyeglass and ECA clusters when opening a support case.... 948
5.11.1. Troubleshooting ECA Configuration Issues.... 960
6. Eyeglass Cluster Storage Monitor Admin Guide.... 962
6.1. Introduction to this Guide and What's New.... 963
6.2. Active Directory Managed Quotas Overview.... 966
6.3. How to Manage Storage by the Share or Export.... 976
6.4. Operations - Cluster Bulk Quota Management Features.... 978
6.5. Configuration - Cluster Storage Reports.... 981
6.6. Configuration - IGLS CLI commands to configure Cluster storage monitor features.... 983
6.7. Unlock My files Help Desk Application.... 984
7. RPO Trending and Reporting Guide.... 992

7.1. Overview....993
7.2. Getting Started.... 996
7.3. RPO Calculations:....1000
7.4. RPO Summary & Compliance Email Report.... 1004
7.5. Example Report:.... 1008
7.6. The Charts....1010
7.7. Charts Uses:.... 1014
7.8. Advanced Settings....1019
7.9. RPO CSV Reports....1024
8. Data And Config Migration Admin Guide....1027
8.1. What's New.... 1028
8.2. Overview....1030
8.3. Typical Use Cases:.... 1031
8.4. Supported Clusters.... 1032
8.5. Planning Migrations between Access zones.... 1033
8.6. Use Case #1 - System to Other Access zone Same cluster.... 1034
8.7. Use Case #2 - System to Other Access zone Remote cluster....1035
8.8. Use Case #3 - Merge Access Zones Configuration....1036
8.9. Use Case #4 - Overlapping Access Zones Configuration....1037
8.10. Prerequisites to Use the Migration Feature....1041
8.11. How to create a Data and Config Migration Job.... 1042
8.12. How To Re-apply Default SMB Share ACL Post Migration - only if Enable write access was disabled.... 1052
8.12.1. Add back the original Microsoft Default ACL's....1053
8.12.2. Check ACL's after changes to migrated parent folder.... 1054
8.12.3. Delete extras ACL's.... 1056
8.12.4. Check ACL Delete.... 1057
8.12.5. Connect to Smartconnect Name to test mount and write access to the data.... 1058
8.13. Planning Timeouts for Migration jobs....1064
8.14. Known Limitations.... 1065
8.15. End to End Data Migration Steps to Move Data/Config and Users to new Access Zone.... 1066
8.16. Successful Migration Job View - Example.... 1070
9. Eyeglass Search & Recover Admin Guide....1074
9.1. What's New with Search & Recover....1076
9.2. Why Eyeglass Search & Recover is the only solution for Scale Out NAS Content Indexing.... 1078
9.3. Index Update Intervals, Index Granularity, Index Ignored words, Indexing feature options.... 1079
9.4. Product Requirements, Cluster Sizing and Tested Scaling Limits.... 1083
9.5. Use Cases....1086
9.6. Search & Recover Known Limitations.... 1091
9.7. Search & Recover Solutions Guides.... 1092

9.7.1. Solution Guide - Script Download Examples..... 1099
9.7.2. Eyeglass Search & Recover - Archive Solution with Dell EMC ECS from Linux Host.....1108
9.7.3. Eyeglass Search & Recover - Archive Solution with Dell EMC ECS from Windows Host..... 1115
9.7.4. How to OCR Image Data with Search & Recover.....1125
9.7.5. Writeable Snapshots with File Clones Automation.....1133
9.7.6. Legal Hold Solution..... 1141
9.7.7. Smart Lock Data Retention , Lock status, and Expiry Reporting and Delete Solution.....1147
9.8. Deployment Diagrams, Firewall Ports , VM Requirements and Supported OneFS Releases.....1154
9.9. Understanding Search Security Results and User Login sessions.....1156
9.10. Search & Recover Cluster Configuration Steps.....1164
9.11. Cluster Administration - UI Access and Security Configuration.....1212
9.12. Search & Recover Cluster Operations..... 1215
9.13. User Search Guide..... 1243
9.14. Administrator Search Guide.....1253
9.15. File System Analytics with Quick Reports - Guide.....1257
9.16. Advanced Searching Syntax and Script Automation Editor..... 1285
9.17. DR Failover Considerations..... 1297
9.18. Management Diagnostic Tools.....1299
9.19. Monthly Index Backup Solution Guide.....1301
9.20. Trouble Shooting Search Results.....1306
9.21. Advanced Cluster Configuration..... 1309
9.22. Cloud Pool Reporting..... 1314
9.23. Eyeglass Search GraphQL API.....1319
9.23.1. Scripted API Searches with Search & Recover GraphQL API..... 1336
9.24. Content Classification Feature Guide.....1346
9.25. ShowBack and ChargeBack Guide.....1408
10. Eyeglass Golden Copy Admin Guide..... 1417
10.1. Golden Copy Overview..... 1418
10.2. Limitations, Requirements and Supported Target List..... 1425
10.3. Networking Deployment Overview, Cloud Provider IP Ranges, Proxy Configuration.....1435
10.4. Understanding Copy Performance Configurations.....1441
10.5. Cloud Provider Cost Considerations.....1444
10.6. Supported Data Security & Storage Class Life Cycle Options..... 1448
10.7. Appliance Security, Authentication & Hardening.....1450
10.8. Golden Copy Configuration Steps..... 1462
10.9. Golden Copy Back Bundle & Adv License Configuration Steps..... 1546
10.10. Golden Copy GUI - Beta..... 1564
10.11. Golden Copy VM Operations..... 1571
10.12. S3 Storage Bucket Configurations Options , Operations and Settings.....1574

10.13. Trouble Shooting File Copies....1586
10.14. Golden Copy Solutions Guides....1587
10.14.1. DR Solution for Azure....1589
10.14.2. Application Disaster Recovery to Amazon S3 and FSX Windows Service....1602
10.14.3. Search & Recover & Golden Copy Archiving Solution....1612
10.14.4. Immutable Storage for Azure Blob Storage....1622
10.14.5. Data LifeCycle - Cost Reduce with Archive to the Cloud....1626
10.14.6. Object Backup and Basic DR Solution....1630
10.14.7. Bulk Loading Data with Azure Data Box....1636
10.14.8. Bulk Recall of Data from AWS and Azure....1645
10.14.9. Bulk Loading Data with AWS Snowball....1647
11. Eyeglass Performance Auditor Admin Guide....1653
11.1. Performance Auditor Overview....1654
11.2. Performance Auditor Requirements....1656
11.3. How to Use Performance Auditor to Root Cause....1658
11.4. Performance Auditor Advanced Configuration....1688
12. Eyeglass AnyCopy Admin Guide....1690
12.1. What's New with AnyCopy....1691
12.2. Installation of AnyCopy....1694
12.3. Planning Guide for AnyCopy....1698
12.4. AnyCopy Admin & User Guide....1702

1. Eyeglass PowerScale Edition Admin Guide

Home Top

- What's New with DR Product Releases
- Introduction to this Guide
- Eyeglass Ports Requirements , Scalability Limits and Phone Home Requirements
- License Management with Eyeglass
- Eyeglass Deployment Options
- Eyeglass Backup Options
- Eyeglass Jobs
- Configuration Replication
- Advanced Quota Replication
- Manage Eyeglass Jobs
- Failover Readiness Validations DR Dashboard
- RPO Reporting and Trending Feature Guide
- Role Based Access Controls RBAC
- Configure Email, Twitter, Slack, Webhooks for Notifications of Eyeglass Monitoring Events
- ECA (Eyeglass Clustered Agent) CLI Commands
- DR Design Guides with Eyeglass
- Eyeglass CLI Commands
- Eyeglass Appliance Time Synchronization Best Practice
- Eyeglass Automatic Updates for Recommended Packages
- Update Eyeglass Appliance Network Settings

- Appliance Security Updates and Eyeglass Updates with HTTP Proxy
- Diagnostic Tools for Dark Site Support
- Role Based Access Control And Authentication Guide
- Eyeglass Alarm forwarding Guide - Syslog and Legacy SNMP
- How to Setup Email alarms with Exchange Server
- How to Change PowerScale IP Address in Eyeglass
- Pre Post Failover Scripting Guide
- Eyeglass Backup and Restore
- Eyeglass API guide
- How to convert VMware Eyeglass appliance and Migrate to Microsoft Azure
- TLS Certificate Procedures for Eyeglass
- Configuring Eyeglass with dual NIC environments
- Changing Eyeglass UI Behavior
- Custom email routing by application or alarm subject contents

© Superna LLC

1.1. What's New with DR Product Releases

Home Top

Release

1. 2.5.5 - Full support for 8.2 Onefs including default directory quota failover. All Eyeglass functionality is supported with this release.
    a. SyncIQ encryption can be used and does not impact Eyeglass Failover
2. 2.5.6 - DR Release with many new feature listed here
3. 2.5.7 - SyncIQ Monitor - Audits SyncIQ data sync by adding test data with timestamps to the source cluster and compares when synciq runs, and verifies timestamps on the target cluster. This provides the highest level of data off site backup confidence in your off site data. The job is disabled by default and requires configuration.
4. 2.5.7 update 1 - Major security updates with built in secure virtual network between Eyeglass and ECA VM's. ECA management VM GUI access now secured by authenticated proxy login over https, with no direct access, Eyeglass removes 4 open ports 2011, 2012, 2013 and 2014 are no longer required. Removal of additional components for web sockets.
    a. Full details available here.
5. 2.5.8 Adds Onefs 9.2.1.x > release full configuration backup to be created daily and protected by Eyeglass backup. If phone home is enabled this allows an automated fully automated cluster backup of key configuration information. HTTP, NDMP, NFS, SMB, Quota's, S3 and snapshot configuration data is

protected see below. This is enabled automatically and runs daily using the scheduled job shown below.

a. igls admin schedules | grep -C 2 ClusterConfigurationBackup

b. NOTE: Requires new permission of the eyeglass user, see documentation.

c.

```
ee152:/opt/superna/sca/data/config_backup/CL9200A/CL9200A-20211012210000 # ls -l
total 84
-rw-r--r-- 1 sca users   114 Oct 12 17:00 backup_readme.json
-rw-r--r-- 1 sca users   330 Oct 12 17:00 http_CL9200A-20211012210000.json
-rw-r--r-- 1 sca users   596 Oct 12 17:00 ndmp_CL9200A-20211012210000.json
-rw-r--r-- 1 sca users 18814 Oct 12 17:00 nfs_CL9200A-20211012210000.json
-rw-r--r-- 1 sca users  4240 Oct 12 17:00 quota_CL9200A-20211012210000.json
-rw-r--r-- 1 sca users  1629 Oct 12 17:00 s3_CL9200A-20211012210000.json
-rw-r--r-- 1 sca users 34625 Oct 12 17:00 smb_CL9200A-20211012210000.json
-rw-r--r-- 1 sca users   517 Oct 12 17:00 snapshot_CL9200A-20211012210000.json
```

© Superna LLC

1.2. Introduction to this Guide

Home Top

- Overview
- Software Release Updates
- Help Getting Started with Eyeglass

Introduction to this Guide

The purpose of this guide is to assist you in configuring and and administering your Eyeglass PowerScale Edition.

Overview

Eyeglass addresses an HA DR automation and orchestration platform for PowerScale clusters. As file based data becomes mission critical to business, a DR application focused on automating and synchronizing configuration data is needed to access data shares.

A solution that ensures all data (SyncIQ replication) and configuration data is also replicated for a DR event is only one aspect of a well designed high availability solution. Many factors need to be considered to failover a file based application.

Software Release Updates

The initial releases of the Eyeglass product focused on sync of configuration data, cluster and SyncIQ policy automation failover, and reporting automation of configuration and DR readiness. Superna is working continuously to improve Eyeglass PowerScale Edition with enhancements and feature updates. The What's Coming link will provide you with detail on past, current, and future enhancements.

Help Getting Started with Eyeglass

To get started you should understand the failover options and which one makes most sense for your environment. The Read me First section of the Eyeglass PowerScale Edition Documentation that covers Eyeglass DR Planning and Implementation will help to decide which option is best.

Superna also offers design and implementation services that can assist with this decision.

The next most important part of DR is planning, planning, planning and more planning. To assist with this planning process, we provide a checklist (Failover Planning Guide and checklist) that covers all the key steps that we expect customers to execute before attempting a failover with Eyeglass (or failover in general). Our support team will refer you to this document, if you ask for assistance in failover, to understand which of these preparation steps are completed. This document is designed to ensure your success. We suggest you review and adapt it to your business processes.

© Superna LLC

1.3. Eyeglass Ports Requirements , Scalability Limits and Phone Home Requirements

Home Top

- Eyeglass Ports Requirements
- Eyeglass Support and Phonehome Whitelist URL's
  - PC Browser Upload and support site usage URL whitelist for full access to (support.superna.net)
  - Download software and download license keys from support.superna.net
  - Phone Home Remote Monitoring Test Internet URL Steps
- Phone Home Message Flow
- Eyeglass Proxy Login Message Flow between Eyeglass VM and Isilon/PowerScale
- Eyeglass Scalability Limits and Appliance Memory Minimum Requirements

Eyeglass Ports Requirements

Port	Protocol	Direction	Eyeglass Release	Product Requires the ports open	Function
Operat					It is customer

ing Syste m Open Suse 15.1					responsibility to patch the operating system and allow Internet repository access for automatic patching. The OS is not covered by the support agreement.
DNS port 53 UDP	DNS	Eyeglass --> DNS server Eyeglass --> GroupNet DNS configured on all clusters		All	Functional DNS is a requirement for multiple validations needed for failover and Failover Readiness
Phone Home Monito ring TLS 443	TCP TLS 1.2	Eyeglass appliance --> Internet		All	DR Monitoring service remote monitoring OR phone home remote log upload for support and health checks for Ransomware Defender, Easy Auditor and Performance Auditor products. See phone home Internet ports that need to be opened below.
NTP 123	UDP	Eyeglass appliance --> NTP server in your environment		All	Time sync should use same NTP as the clusters. Should always disable vmware host VM time sync option.
SMTP 25	TCP	Eyeglass appliance --> Mail server in your environment		All	Email of alarms from Eyeglass to your mail server
OS Repo for Securit	TCP	Eyeglass appliance --> open suse mirror		All	1. URL to allow security updates http://dow nload.opensuse.or

y patches 80 http		repositories			g 2. NOTE: Security patches come directly from open suse and requires the appliance to have access to download the patches and apply on the weekly schedule.
HTTPS over port 8080	TCP TLS 1.2	Eyeglass appliance → Isilon/PowerScale cluster		All	REST API is authenticated using the service account created here. Authenticatio uses Isilon session authentiation method.
SSH port 22	AES	Eyeglass appliance → Isilon/Power Scale cluster		All	SSH access for some CLI commands
NFS TCP, UDP 111, 2049 (in some environments UDP 300))	TCP & UDP	Eyeglass appliance → Isilon/Power Scale Cluster		Ransomware Defender, Easy Auditor, Performance Auditor	Audit Data Ingestion
Syslog for ECA clusters to send logs to Eyegla	TCP	ECA clusters → Eyeglass Appliance		Ransomware Defender, Easy Auditor, Perfor	Syslog (non standard port) used to send ECA cluster VM logs to Eyeglass for support logs (enabled in Eyeglass can be disabled, if no ECA deployed).

ss using port 5514				mance Auditor	
HTTPS 443	TCP TLS 1.2 AES - unsigned certificate	admin pc browser → appliance		All	Secures client to browser access.
target port 80 → destination random TCP source port on the browser	only used to redirect to 443, can be blocked if needed	admin pc browser → appliance		All (**optional**)	If connection on ip address port 80 is made a http 301,302 redirect is returned on port 80 to switch the browser to https and url https:/x.x.x.x/eyeglass. **NOTE: No services run on port 80 and this is only used to redirect to port 443 HTTPS**
https 2011 websocket	TCP TLS 1.2 AES	admin pc browser → appliance	> 2.5.7 update 1 not required	DR	Websocket for real-time appliance to browser updates (redirected to 2012).
2012 TLS websocket	TCP TLS 1.2 AES	admin pc browser → appliance	> 2.5.7 update 1 not required	DR	Websocket for real-time appliance to browser updates (redirected to 2012).
2013	TCP	admin	> 2.5.7	Easy	Websocket for Easy

TLS websocket	TLS 1.2 AES	pc browser → appliance	update 1 not required	Auditor	Auditor wiretap feature (only required if this product is installed).
2014 TLS websocket (new Performance Auditor)	TCP TLS 1.2 AES	admin pc browser → appliance	> 2.5.7 update 1 not required	Performance Auditor	Websocket for Performance Auditor application (only required if this product is licensed and installed).
SSH 22	TCP AES	admin pc workstation → appliance		All	secure shell access.
Proxy login SMB 2 (only) 445	TCP	appliance → Isilon/PowerScale	> 2.5.7 SMB3 supported with encryption	All	Used to authenticate to AD through Isilon/PowerScale using standard Microsoft SMB authentication request for Role based login proxy interface.
SMB Security Guard Ransomware Defender SMB TCP 445 SMB2 only	TCP	appliance → Isilon/PowerScale	> 2.5.7 SMB3 supported with encryption	Ransomware Defender	Used by Ransomware Defender (if licensed) to simulate ransomware attack automation.
SMB Robo	TCP	appliance → Isilon/Power	> 2.5.7 SMB3	Easy Audito	Used by Easy Auditor to test audit health for audit

Audit Easy Auditor SMB TCP 445 SMB2 only		Scale	supported with encryption	r	data ingestion and database health automation (if licensed).
Dual DNS Delegation	UDP	appliance port 53 UDP DNS --> Groupnet(x) DNS servers		DR	This is new in 2.5.6 or later and requires Eyeglass to be able to access the Groupnet DNS servers to validate Dual DNS delegation is configured correctly. The OS DNS is not used since the DNS that must be configured correctly is used by Isilon/PowerScale itself.

Eyeglass Support and Phonehome Whitelist URL's

PC Browser Upload and support site usage URL whitelist for full access to

(support.superna.net)

1. https://*.zopim.com **(your pc browser --> Internet, Internet --> your pc browser)**
2. https://licenses.supernaeyeglass.com **(your pc browser --> Internet, Internet --> your pc browser)**
3. **https://support.superna.net (your pc browser --> Internet, Internet --> your pc browser)**
4. https://supernahelp.zendesk.com **(your pc browser --> Internet, Internet --> your pc browser)**
5. https://cloudapps.supernaeyeglass.com **(your pc browser --> Internet)**

Download software and download license keys from support.superna.net

1. https://software.supernaeyeglass.com **(your pc browser <-- Internet)**

2. https://licenses.supernaeyeglass.com **(your pc browser --> Internet, Internet --> your pc browser)**

Phone Home Remote Monitoring Test Internet URL Steps

Overview - faster more efficient support, enables proactive response without your involvement)
How to test firewall port access to required URL's

1. SSH to Eyeglass appliance as admin user
2. type admin password (default: 3y3gl4ss)
3. Execute below command to test get command:
   `wget https://na-static-phonehome.supernaeyeglass.com`
4. Execute below command to test post command:
   `curl -X POST -k http://na-static-phonehome.supernaeyeglass.com`
5. Send us the output of Step #3 and #4.
6. Done.
7. The Monitoring service requires the following URL's allowed
   a. https://cloudapps.supernaeyeglass.com (appliance to internet)
   b. Note: Superna has made the IP change for this URL therefore please whitelist IP address 35.244.217.10
   c. https://na-static-phonehome.supernaeyeglass.com (appliance to internet) IP address 35.207.34.234

Phone Home Message Flow

1. After Phone Home is enabled the following message flow is described below.
   a. NOTE: No inbound firewall rules need to be open from the Internett are required .

b. NOTE: No data is collected other than support logs that can be created in the About Icon Backup tab and is uploaded to support through the support if Phone Home is disabled.

c. NOTE: No remote control functionality is possible with the phone home feature

d. NOTE: No PHI data of any kind is collected in support data

e. NOTE: if configured a proxy device can be used to reach the Internet.

f. NOTE: url or IP based firewall information is provided on this page

2. Eyeglass Appliance, Golden Copy or Search & Recover product appliance phone home actions

   a. After initial enable of Phone home from the About Icon a registration request is sent via HTTPS POST to https://na-static-phonehome.supernaeyeglass.com that includes the appliance appliance ID, Appliance version information.

   b. Twice per 24 hours a hear beat is sent via an HTTPS POST to https://na-static-phonehome.supernaeyeglass.com that updates the portal that the appliance is still alive and running.

   c. Once per 5 minutes (random time within the 5 minutes) the appliances will send HTTPS GET request to the https://na-static-phonehome.supernaeyeglass.com url to see if remote log upload request has been requested.

i. If no request to upload logs, no actions are taken and phone home sleeps until next log upload request 5 minute window.

ii. If a request to upload logs is returned from the HTTPS GET request, the appliance will generate support logs zip file and then HTTPS POST the zip file to this url https://cloudapps.supernaeyeglass.com

Eyeglass Proxy Login Message Flow between Eyeglass VM and Isilon/PowerScale

1. Eyeglass browser https --> Eyeglass VM
2. Eyeglass VM --> SMB2 standard Microsoft authentication request sent to Isilon/PowerScale ip address used to add cluster to Eyeglass.
3. Isilon/PowerScale sends authentication request to AD to validate password.
4. Eyeglass --> sends rest api to Isilon/PowerScale requesting AD group membership for User X from login request.
5. Isilon/PowerScale returns Authentication request to Eyeglass VM.
6. Isilon/PowerScale returns list of AD groups the user is a member of in AD.
7. Eyeglass compares AD groups to Role based Access configuration to determine permissions in Eyeglass and displays Icons based on this security evaluation process.
8. User desktop loads based on role configured.

Eyeglass Scalability Limits and Appliance Memory Minimum Requirements

Scaling Limit Area	Tested Scaling Limits	Notes
Number of Managed Clusters 1 appliance	22	contact Support for RAM requirement
SyncIQ Policies All clusters	> 100 --> 64GB > 200 --> 84 GB	
Access Zones	> 10 --> 32GB > 30 --> 64GB > 50 --> 84 GB	Requires 32G-84G RAM
Failover job limitations	100 policies selected in a single failover	Requires 64G RAM
total object count (shares + exports + quotas)	< 5 000 --> 16 GB 5000 --> 10 000 32GB to 40GB > 10 000 --> 64 GB > 20 000 --> 84GB	
Clusters added to the appliance	4 --> 32GB 4 - 8 clusters --> 64GB > 10 64GB - 84 GB	
Performance Auditor	Review the requirements above and if none apply, Eyeglass requires minimum 32 GB RAM when Performance Auditor is licensed	Minimum 32GB

Concurrent administrators 3 or more	Number of current logged in administrators to Eyeglass GUI using RBAC or not using RBAC	Add 8 GB ram to above requirements

© Superna LLC

1.4. License Management with Eyeglass

Home Top

License Management with Eyeglass

Licensing

Eyeglass license keys are auto applied for all products.

DR Module Key Assignment

This auto assign function will apply DR module cluster keys to the largest node count cluster first, and remove node count from the total. Then apply node based keys from a pool of node keys to Clusters added to Eyeglass. This is completed until all node keys and cluster keys are assigned.

Ransomware Defender and Easy Auditor Key Assignment

Each writable cluster requires an agent license and agent maintenance. Licenses are automatically assigned to clusters. A cluster agent license can failover to the DR cluster automatically. If more clusters are managed than agent licenses than the incorrect cluster could be assigned a license. The license manager icon on the desktop allows selecting which cluster to license. The license manager can be used to select which cluster receives which agent license key.

Auto licensed devices detect the cluster that has SyncIQ policies that are writable to apply the license. If the device selected is not the correct device, then keys must be assigned using the licensed devices procedure below.

How to Manually Assign Licenses to Clusters for Products

1. Open License Manager Icon.
2. Select License Devices tab.
3. Select Unlicensed option for a product license that says Auto Licensed to remove the license from the cluster for that product.
4. Select User Licensed to select this device to receive a license.
5. NOTE: Another cluster with Auto Licensed may need to be set to Unlicensed before being able to apply User License to another device.
6. This will pin the license to the cluster with User Licensed option.

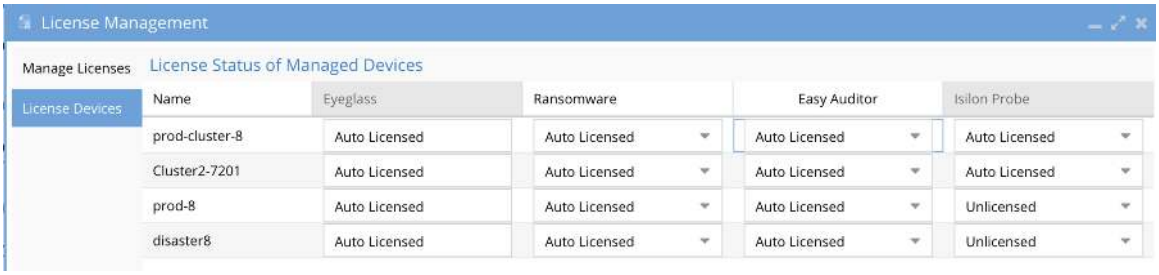

A system alarm will be issued in case insufficient licenses exist and more writable clusters are detected during inventory functions.

© Superna LLC

1.5. Eyeglass Deployment Options

Eyeglass Deployment Options

Several options exist that allow customers to deploy Eyeglass in several configurations.

Warm Standby

This procedure is to protect the production Eyeglass appliance. Using the Eyeglass Automated 7 Day Backup feature and SyncIQ, the procedure will sync the backup archives to a running Warm Standby Eyeglass appliance in a 2nd cluster.

Notes on this solution:

1. This process also means only one Eyeglass appliance has clusters added.
2. This only requires one set of license keys (license keys are appliance specific).
3. Requires SyncIQ enabled clusters to protect the archives.
4. Production Eyeglass syncs Export used for DR of Eyeglass configuration data.

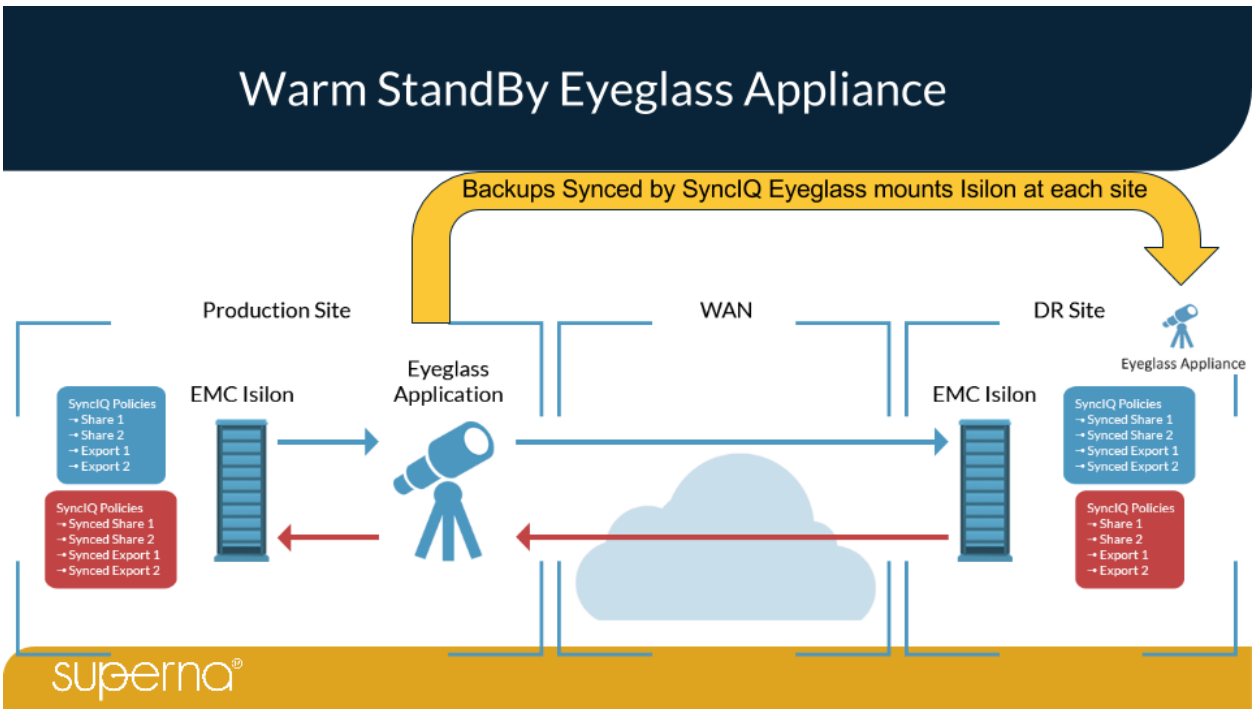

The following link provides detailed information on how to set up the Warm Standby Backup Procedure (Warm Standby)

3rd Site Witness (Supported)

In order to have all DR orchestration and information separate from the DR site it's best practice, when possible, to use a 3rd site to monitor your clusters with Eyeglass that is IP connected to Site A and Site B as shown below.

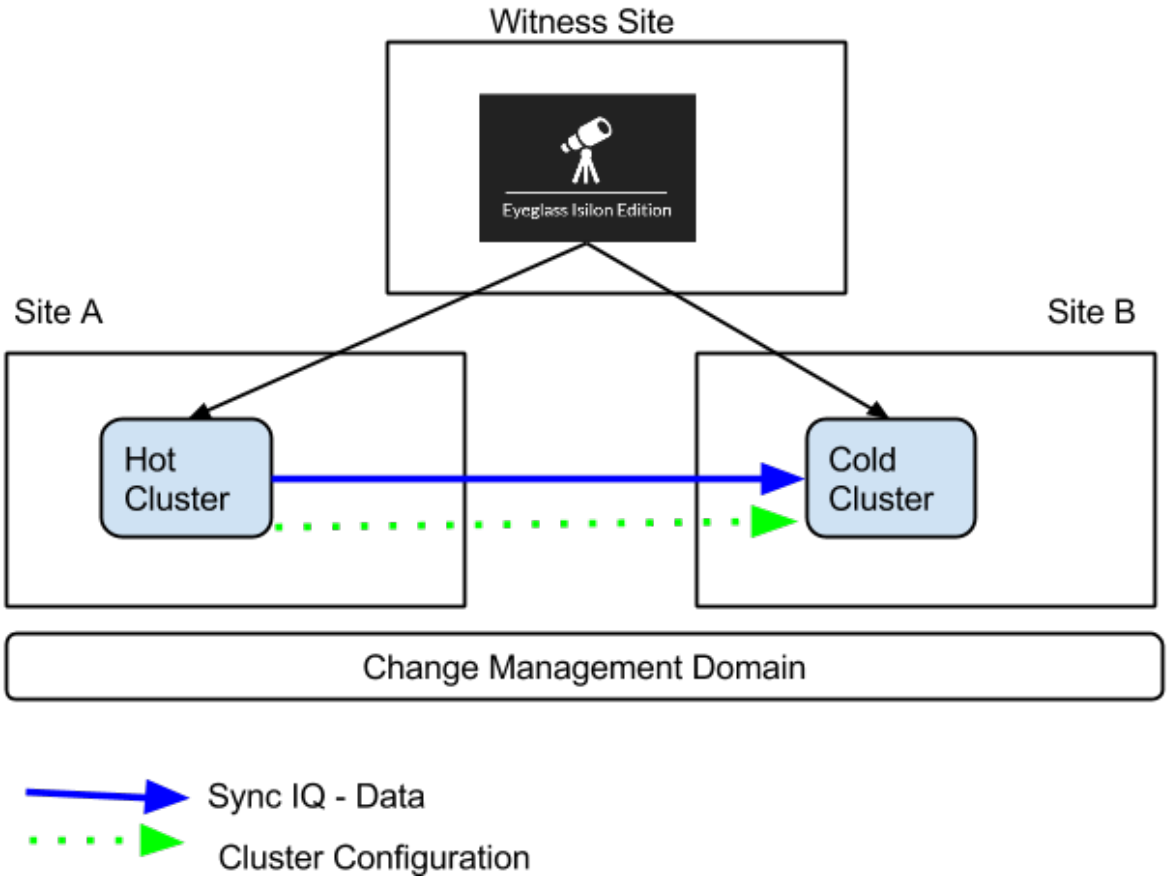

DR Site Witness (Supported)

In this configuration the witness function resides at the Cold or DR site, so that configuration and orchestration data is available in a DR event.

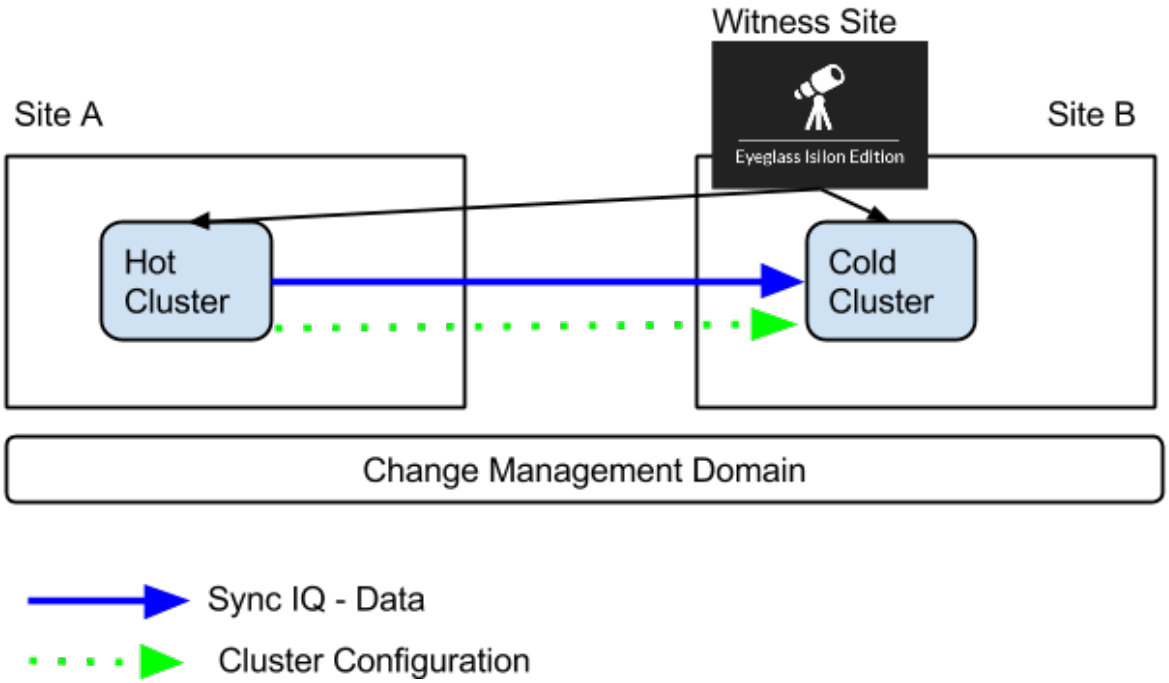
Witness Site
Eyeglass Isilon Edition
Site A
Site B
Hot Cluster
Cold Cluster
Change Management Domain
Sync IQ - Data
Cluster Configuration

© Superna LLC

1.6. Eyeglass Backup Options

Home Top

Eyeglass Backup Options

Eyeglass Warm Standby Deployment

The Warm Standby Eyeglass deployment protects the production Eyeglass appliance by using a 2nd unlicensed Eyeglass appliance and a backup and restore procedure for Eyeglass configuration. Please refer to How to setup Warm Standby Backup Procedure for more information.

Supported Replication Topologies for 2 Copies of Data

Multi protocol NFS and SMB Access Zone failover supports 2 copies of data protection, with Active / Passive configuration. An Access Zone cannot span clusters due to DNS delegation and simplified failover with Eyeglass that is fully automated.

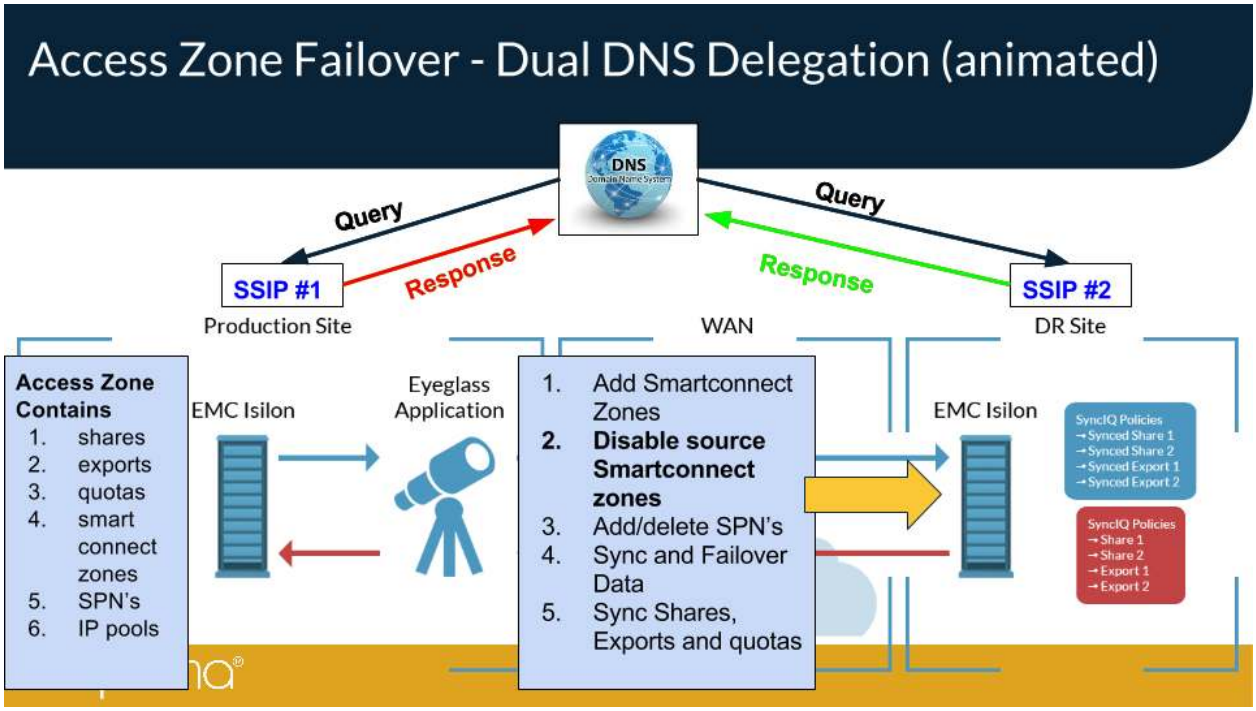

Eyeglass DR LiveOPS Features

LiveOPS are a set of features designed to allow continuous operations with no impact to production DR or cluster operations. The first feature, available in this new family of features, is DR Test mode. This allows customers to execute DR testing on 3rd copies of data that also includes configuration data synced from the production cluster into a test Access Zone.

For detailed instructions on configuration, operation, and requirements of LiveOPS features consult the Live Ops - Continuous Operations Admin Guide

© Superna LLC

1.7. Eyeglass Jobs

Home Top

- Understanding how Eyeglass Jobs Work
- Job Types
  - Configuration Replication Jobs Automatically Created by Eyeglass
    - 0) Type Unconfigured
  - How to Set Type on Unconfigured Jobs
  - New Unconfigured Alarm
    - 1) Share, Export, Alias replication (Type: AUTO)
- Job Creation:
    - 2) Access Zone Replication (Type: ZONES)
    - 3) Quota replication (Type: QUOTA)
    - 4) Configuration Replication: Snapshot Schedules (Type: FILESYSTEM)
  - Configuration Replication Jobs not Auto Created by Eyeglass
    - 1) CUSTOM (Type: CUSTOM)
  - 2) Configuration Replication: DFS Mode (Type: AUTODFS)
  - 3) Configuration Replication: Skip Share, Export, Alias replication Mode (Type: AUTOSKIPCONFIG)
  - 4) Disaster Recovery Testing: (Type: AUTOMATIC)
  - 5) Runbook Robot: (Type: RUNBOOKROBOT and AUTOMATIC)
  - 6) Failover Readiness Jobs (Type: AUTOMATIC)

Understanding how Eyeglass Jobs Work

Jobs are the basis for all automation in Eyeglass and job types dictate the type of automation that will be performed. Currently the job types supported are the following:

- **Unconfigured (new in 2.5.6 or later releases) - new policy detected that needs a type set.**
- Configuration Replication (shares, exports, permissions, nfs alias).
- DFS mode enabled config replication policies (shares only).
- Skip Share, Export, Alias replication.
- Quota Replication (quotas - all types).
- Snapshot schedule and Dedupe settings.
- Access Zone replication.

Failover related job types:

- Failover jobs (SyncIQ policy or Access Zone or DFS).
- Created by DR Assistant.
    - Failover Readiness.
- Runbook Robot jobs.
- Disaster Recovery Testing - LiveOPS DR test mode jobs.
- Access Zone data migration

- Created by Data Migration Icon

Further to the above jobs, the following modes can be set on jobs:

1. **Automatic** - These are built by Eyeglass after auto detection of SyncIQ policies which are then used to detect which shares, exports, nfs alias and quotas should be replicated. They are automatically created.
2. **DFS Mode** - These are set by the administrator and will be used when configuring DFS mode failover. This job type will rename SMB shares on the read-only cluster automatically. You can switch a job from DFS mode back to sync mode.
3. **AutoSkipConfig Mode** - These are set by the administrator and will be used for the case where you do not want SMB shares and NFS exports to be synced to the target cluster, but still want to be able to failover the data. You can switch a job from AutoSkipConfig back to sync mode.
4. **Custom**:
    a. These are user created and allow detection of shares, exports, quotas in a path in the file system, to be detected and added to a job, that will replicate the configuration to a target cluster without needing a SyncIQ policy to exist.
    b. The data and path must exist on the target cluster

Job Types

Configuration Replication Jobs Automatically Created by Eyeglass

0) Type Unconfigured

Purpose: New in 2.5.6 and later releases, all new detected policies will by shown in a new section called Unconfigured. These jobs must have a type set of Auto or DFS and then enabled before any sync actions will occur.

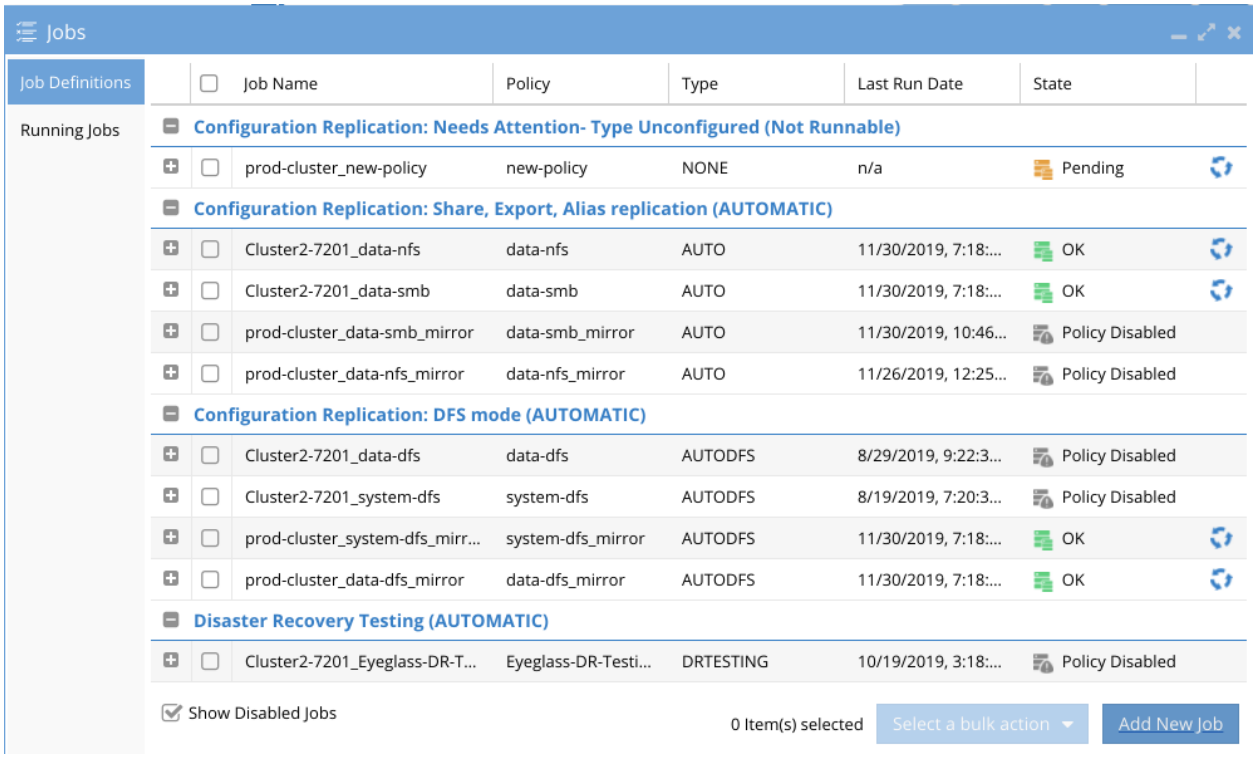

A new alarm is raised for each new policy detected that requires configuration applied before the policy enters into production DR status.

How to Set Type on Unconfigured Jobs

1. Open the Jobs icon and select an Unconfigured Job
2. Click Bulk Actions menu

a. 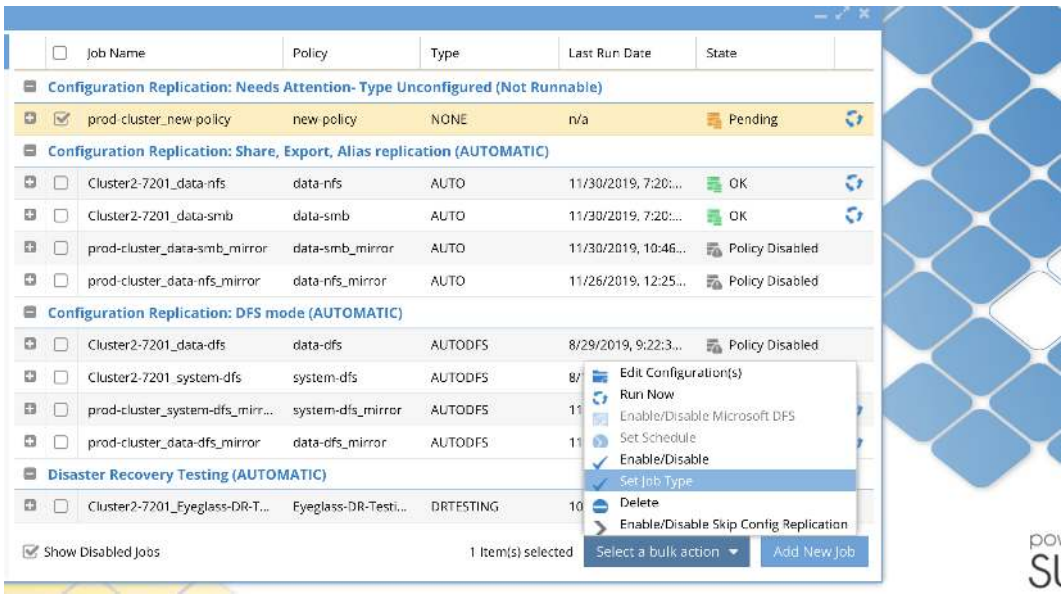

b. Then select AUTO or AUTODFS or SKIPCONFIG

c. 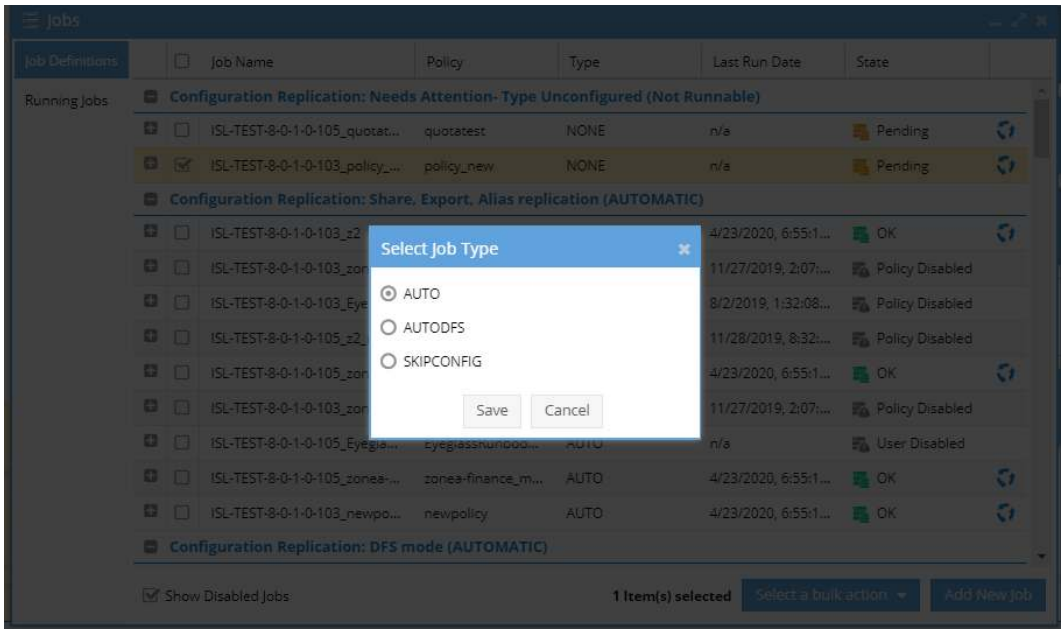

3. Then select the policy in the AUTO or AUTODFS or SKIPCONFIG section and choose Bulk Actions menu and enable it to put the policy into production sync state. If left in User Disabled state the policy is not synced and not available for DR in DR assistant.

New Unconfigured Alarm

When a new Synciq policy is created steps must be taken in Eyeglass to set the type and enable it for DR readiness. A new alarm is raised for any synciq policy that is detected during the inventory process that runs every 5 minutes by default. The new policy alarm names the cluster where the policy was detected.

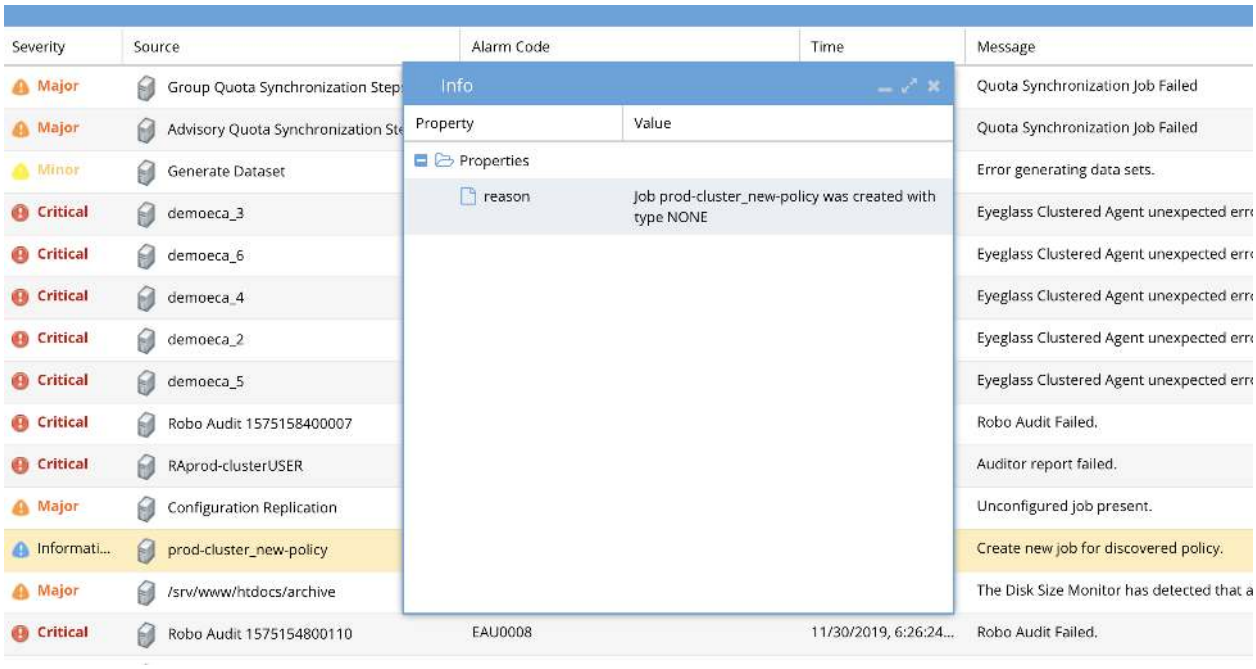

1) Share, Export, Alias replication (Type: AUTO)

Purpose:

- Identify shares / exports / nfs alias that are related to SyncIQ Policies detected based on SyncIQ policy source path.
- Synchronize these configuration items so that they exist on both clusters:
    - associated shares & their configuration
    - associated exports & their configuration
    - associated nfs alias & their configuration

Eyeglass Configuration Replication Job Name convention: < PowerScale Cluster name >_< SyncIQ Policy name >

Job Creation:

The Share, Export, Alias replication jobs are auto created by Eyeglass after auto detection of SyncIQ policies which are then used to auto-detect which shares, exports, nfs alias, Access Zones and quotas should be replicated by Eyeglass.

Schedule: All Eyeglass share/export/alias configuration replication Jobs execute on a 5 minute schedule.

Initialstate: User Disabled (does not run), Must be enabled by the administrator in the Jobs icon

2) Access Zone Replication (Type: ZONES)

Note: This replication occurs when the associated Zone is NOT the System Zone.

Note: Recommendation to leave disabled unless directed by support.

Note: Access zones are case sensitive when syncing.

Purpose:

- Identify Access Zone that is related to SyncIQ Policies detected based on SyncIQ policy source path.

- Synchronize Access Zone so that it exists on both clusters:
- Associated Zone & it's configuration, example user access mappings are synced.

Note: Deleted Zones on a source cluster are not deleted on the target cluster.

Eyeglass Zone Job Name convention: < PowerScale Cluster name >_< SyncIQ Policy name >-< ZONES >

Job Creation:

The Access Zone replication jobs are auto created by Eyeglass based on the SyncIQ Policy of the same name.

Schedule: Zone replication jobs execute on a 5 minute schedule.

Initialstate: User Disabled (does not run).

3) Quota replication (Type: QUOTA)

Background:

Quota jobs that are created based on SyncIQ autodetection are placed in a pending state. This state prevents quotas policies, that are collected and shown in the Inventory tree, from being applied to target clusters paths protected by SyncIQ policies. This is a best practice due to some scenarios that result in errors when quotas are applied to a target cluster file system.

The scenarios to apply quota policies are below.

- In a failover event, the quota job can be selected and "Run Now" option used AFTER the target cluster file system is writable as a

result of SyncIQ failover. This is run automatically under normal conditions. Cluster migrations is another use case where applying to target without a delete on source is desirable.

- Custom jobs can replicate quota policies on a schedule for a path selected in the job. The quotas are applied successfully only when the target file system path on the target cluster already exists. Any new quota created under the selected job path, will be detected and replicated only if the target path also exists.

Note: When you run a QUOTA Job associated with an AUTO share/export configuration replication Job, the Job is based on the Eyeglass current inventory view. If you have made a change in OneFS to a quota, the Eyeglass Inventory Task must have run (runs on a 5 minute schedule) prior to running the QUOTA job in order for the change to be applied on the target.

Warning review quota failover limitations:

This section should be reviewed when planning quota failover solutions.

Review Dell EMC Quota EMC KB - https://support.emc.com/kb/88602

Some combinations of quota domain settings and SyncIQ source and target settings are incompatible. There are several scenarios where this might occur:

- Multiple quota domains span **SyncIQ target subtrees**. SyncIQ translates source operations such as file or directory moves (mv) into similar actions on the target cluster. **Moving files or directories across quota domains is not supported, and the syncs will fail.**

- Multiple quota domains span **SyncIQ source subtrees**. If the source cluster will be used for failback, the failback operation could error if multiple quota domains span SyncIQ source subtrees.
- If the source cluster will not be used for failback, multiple quota domains can exist on the source.
- Quota domains exist in directories other than the top-level directory of the SyncIQ policy locations. If a SmartQuotas quota domain overlaps with a SyncIQ policy domain, and failback is desired, then the quotas created must exist only on the top-level directory of the SyncIQ policy source and target locations.
    - A QuotaScan job is still running when sync job starts. The quota scan identifies statistics about the files in a quota domain. If the quota scan does not finish identifying all the files that belong in the quota domain before the sync job starts, the sync will fail.
    - Nested subdirectories receiving new files below an applied quota path at the target side. As part of the transfer process, SyncIQ first creates files and directories in a temporary directory in the target path, and then later moves (renames) them into the final destination. If the final destination has a quota domain configured, this will run into the quota limitation of not being able to move directories into and out of quota domains.

Purpose:

- Identify Quotas that are related to SyncIQ Policies detected based on SyncIQ policy source path.
- Synchronize Quotas so that they exist on both clusters:

- Associated quotas & their configuration.

Eyeglass Quota Job Name convention: <PowerScale Cluster name>_<SyncIQ Policy name>_quotas

Job Creation:

The Quota Failover jobs are auto created by Eyeglass based on the SyncIQ Policy of the same name.

Schedule: Auto-created quota configuration replication Jobs **do not** run automatically. **They are run on-demand as part of a failover.**

Initialstate: User Disabled (does not run)

4) Configuration Replication: Snapshot Schedules (Type: FILESYSTEM)

Purpose:

Sync Snapshot schedules found on SyncIQ paths. Syncs the schedule as per SyncIQ policy paths defined. Will also read dedup path settings (scan and actual) for SyncIQ Policies that match only, and apply the path (corrected by policy path) to target cluster. Can be disabled independently of snapshots using igls command.

Job Creation:

Automatically build, user disabled by default must be enabled.

Schedule: The FILESYSTEM Jobs are run automatically on a 5 minute schedule.

Initialstate: The FILESYSTEM Jobs are disabled and must be enabled, execute at configuration replication scheduled defined.

Note: Snapshot schedule config sync does not overwrite existing Snapshot schedules on the target cluster which have a different Snapshot schedule name.

Configuration Replication Jobs not Auto Created by Eyeglass

1) CUSTOM (Type: CUSTOM)

Purpose:

Use when config data is not protected by SyncIQ policies. Scan path to find config data and replicates to target cluster and path defined in the job. Once created this Job will:

- Identify shares / exports / nfs alias that are related to the Custom Job path.
- Synchronize these configuration items so that they exist on both clusters:
- Associated shares & their configuration.
- Associated exports & their configuration.
- Associated nfs alias & their configuration.

Job Creation:

If you see Job Type “CUSTOM” it means that this is a share/export/nfs alias/quota configuration replication Job that was created manually in the Eyeglass web page. A CUSTOM job that was not created based on a SyncIQ Policy and a SyncIQ Policy is not required or allowed.

Note: Multiple Eyeglass configuration replication jobs where paths overlap is not supported. (i.e. A Custom Job path cannot overlap with

another Custom Job path or an “AUTO” Configuration Replication Job).

Note: Eyeglass custom job where path is the parent of another job is not supported.

Schedule:

- The CUSTOM Jobs are run automatically on a 5 minute schedule.
- QUOTA Jobs associated with a CUSTOM share/export configuration replication Job are run automatically on same replication schedule as the associated CUSTOM share/export replication Job.

Initialstate: User Disabled (does not run)

2) Configuration Replication: DFS Mode (Type: AUTODFS)

Purpose:

Please refer to Eyeglass SyncIQ Failover and Failback with Microsoft DFS for more information.

Job Creation:

DFS Mode is enabled manually from the Eyeglass web page. Please refer to Eyeglass SyncIQ Failover and Failback with Microsoft DFS for more information.

Schedule: The AUTODFS Jobs are run automatically on a 5 minute schedule.

Initialstate: The AUTODFS Jobs when enabled will have the same state as the AUTO Job that it came from.

3) Configuration Replication: Skip Share, Export, Alias replication Mode (Type: AUTOSKIPCONFIG)

Purpose:

Use this configuration mode for the case where the Shares, Exports or NFS Aliases purposely have different properties and permissions on source on target cluster and need to remain unique. In this case the Job remains enabled in the Jobs window but does not perform and configuration sync operations during Configuration Replication. However, because the Job is still enabled the associated SyncIQ policy can still be failed over thus executing other failover steps such as allow writes and resync prep.

Job Creation:

Skip Mode is enabled manually from the Eyeglass web page: Jobs -> select Job -> Select a bulk action -> Enable/Disable Skip Config Replication

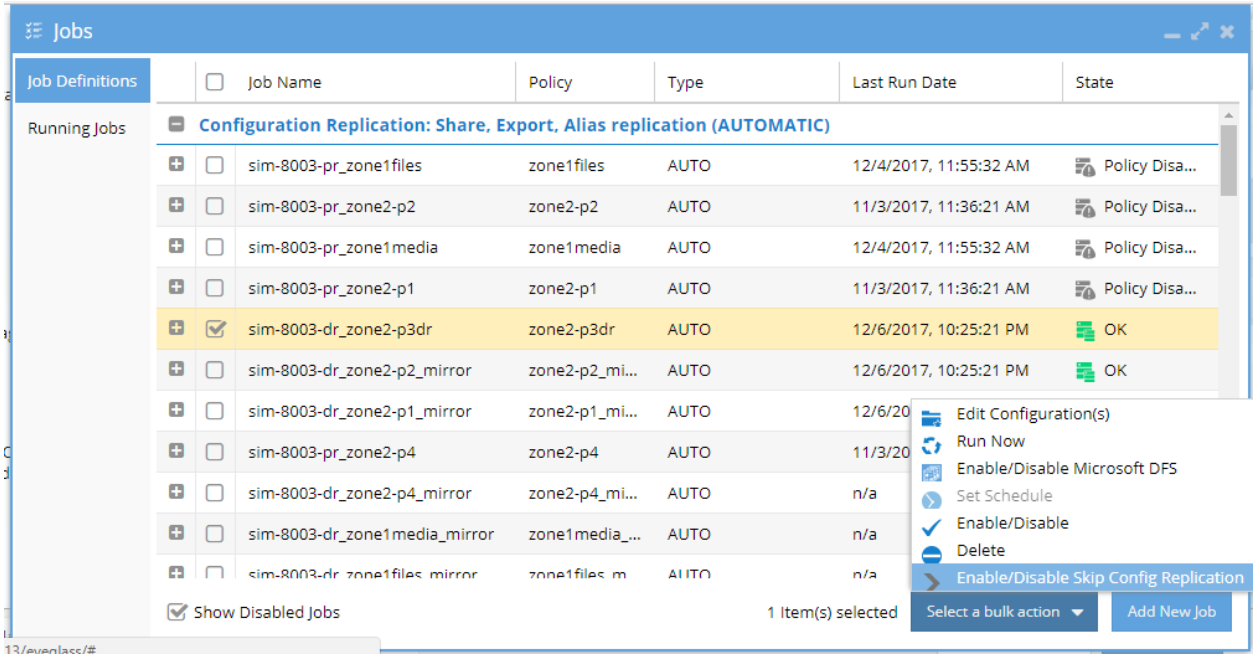

Schedule: The AUTOSKIPCONFIG Jobs are run automatically on a 5 minute schedule.

In the Running Jobs view the Job will appear in the execution list but will not list any Shares or Exports. The Info confirms that the configuration data is not synced.

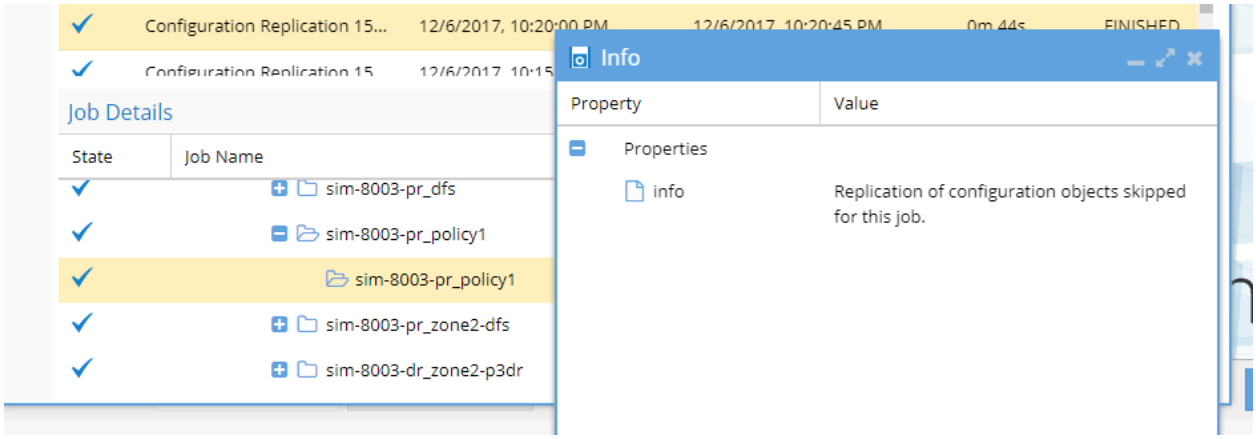

Initialstate: The AUTOSKIPCONFIG Jobs when enabled will have the same state as the AUTO Job that it came from.

4) Disaster Recovery Testing: (Type: AUTOMATIC)

Purpose:

3RD Copy LiveOPS feature, sync's config from prod to DR test access zone if enabled.

Job Creation:

The mode is auto, and is automatically built **if** Eyeglass detects DR test mode policy, user disabled by default must be enabled.

Schedule: The **AUTOMATIC** Jobs are run automatically on a 5 minute schedule.

Initialstate: The **AUTOMATIC** Jobs are disabled and must be enabled, execute at configuration replication scheduled defined.

5) Runbook Robot: (Type: RUNBOOKROBOT and AUTOMATIC)

Purpose:

Continuous DR feature to failover and back daily to DR readiness

Job Creation:

Mode is auto since it's automatically built **if** Eyeglass detects robot access zone or policy name, and user disabled by default must be enabled. Igls command to set the schedule

Schedule: The **RUNBOOKROBOT** Jobs are run automatically on a 24 hour schedule.

Initialstate: The **RUNBOOKROBOT** Jobs are disabled and must be enabled, execute at configuration replication scheduled defined.

6) Failover Readiness Jobs (Type: AUTOMATIC)

Purpose:

Analyzes Access Zones against failover readiness criteria (data, config, SPN, network SmartConnect mapping). Updates the DR dashboard with Access Zone readiness and criteria readiness status. Collects subnet and pool data to display failover mapping on the Access Zone Readiness panel on the DR Dashboard. Failover Readiness Jobs are created between replicating cluster pairs, one job for each direction.

Note: Involved in failover operations, it is used to update DR assistant on readiness, update DR dashboard for users to correct errors, alarms Zone Readiness status to alert administrator.

Job Creation:

Mode is auto and is built automatically by Eyeglass and set to user disabled. The job can be enabled to analyze access zone readiness, if access zone failover is not planned it can remain disabled.

Failover Readiness Job Name convention: PowerScale Cluster name_PowerScale Cluster name

Schedule: Execute on a 15 minutes schedule.

Initialstate: Disabled

© Superna LLC

1.8. Configuration Replication

Home Top

- Prerequisite for Configuration Replication
- Replicating Shares with $ for Security on Target Clusters
- How does Eyeglass Determine Uniqueness - the Details
- Uniqueness and Equality between Shares
- Uniqueness and Equality between Exports
- Uniqueness and Equality between Zones
- What Configuration Properties are Replicated by Eyeglass?
- Shares
- Exports
- Quotas
- Zones
- Snapshots
- Dedupe
- SyncIQ Policies

Replication

Prerequisite for Configuration Replication

In order for Eyeglass to be able to execute the configuration replication for the shares/exports/nfs alias/quotas/zones that have been detected for the SyncIQ policies, the following is required:

- Directories associated with the shares/exports/quotas/zones must exist on the target PowerScale cluster as defined in the SyncIQ policy.
- Access Zones associated with the shares/exports/quotas must exist on the target PowerScale cluster.

Replicating Shares with $ for Security on Target Clusters

This section describes how configuration replication can be used to hide shares on DR and unhide shares after failover. This is to hide the data on the DR cluster automatically.

Use Cases:

1. DFS mode for SMB failover with shares renamed on failover and hidden on DR cluster
2. Access Zone failover with shares Synced and hidden on DR cluster

Configuration:

1. Enable a policy for DFS mode from Jobs icon

Note: This will work for Access Zone failover as well.

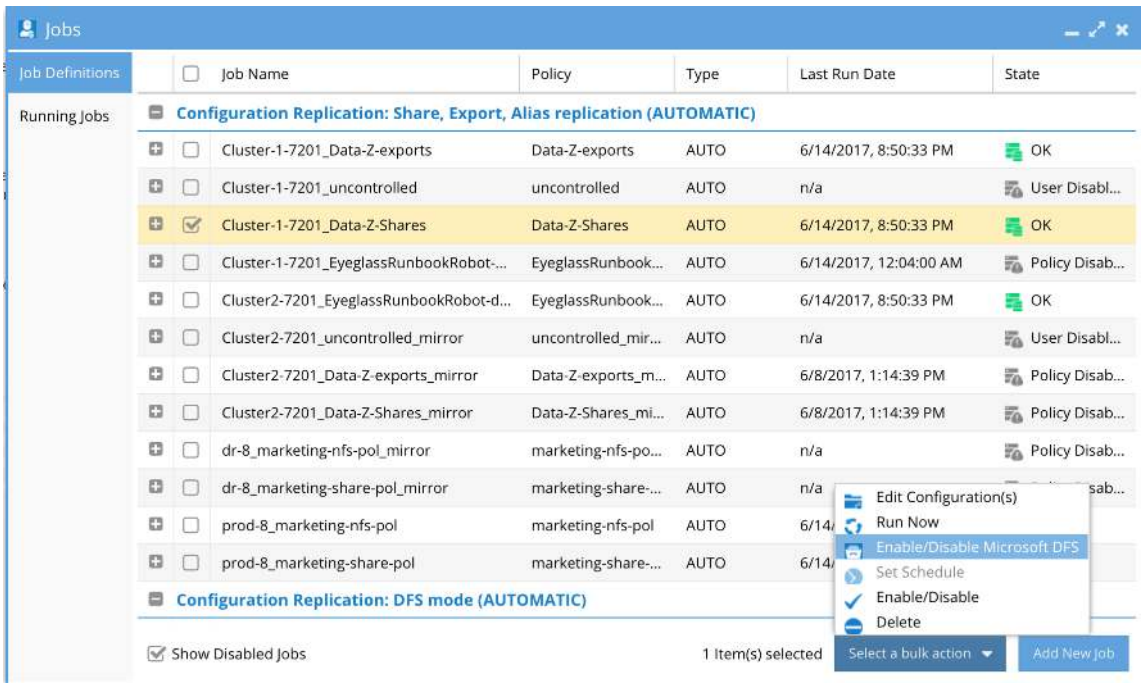

2. Now change the dfs config suffix.

1. This can be enabled by editing /opt/superna/sca/data/system.xml file. Change the tag to include $ as shown below:

<dfssharesuffix>$</dfssharesuffix>

NOTE: on an existing installation the old DFS renamed share will need to be manually deleted.

3. After failover the share names will flip and hide it on the Source cluster.

How does Eyeglass Determine Uniqueness - the Details

Uniqueness and Equality between Shares

The combination of share name and zone name is the unique value that we use to determine if one share equals another. This applies to records on the same cluster, as well as comparing shares between two PowerScale clusters.

Uniqueness and Equality between Exports

Uniqueness is determined by the client list and the path to sync between Access Zones on replicating clusters. Same path but different client list (all client lists) will be treated as different exports to sync within a given Access Zone.

Uniqueness and Equality between Zones

Zone name is used to determine if one zone equals another. This applies to records on the same cluster, as well as comparing zones between two PowerScale clusters.

What Configuration Properties are Replicated by Eyeglass?

Shares

For Shares all configuration properties are replicated with the following exceptions:

- Local users & their permissions.
- filesystem users & their permissions.

! SID for local and filesystem users are replicated but not mapped to users on the target. This means AD at the target cluster should be setup to resolve SID's using the same AD forest.

Note: Shares with variable expansion now sync correctly to the target cluster example %u share names.

Note: Locally created uses or groups on PowerScale are created and generate a cluster unique UID or GID. This means users with the same name on two clusters have different UID and GID. This is normal Linux machine behavior. Sync local users or groups have no value, since userX on cluster A will not be the same userX synced on cluster B, and therefore will not have access to any data secured to userX on cluster A. The best practice is use LDAP or Kerberos for user security for exports or Active Directory for unified security database.

Exports

For exports, all configuration information is replicated to target with the following **exceptions**:

- export id: Exports on source and target will not have the same export id.
- Dual path exports **must** have both paths fall under the same SyncIQ policy.
- FQDN clients listed **must** be resolvable by the target cluster. If this is not possible Eyeglass full sync mode **must** be used to override API validation and force create exports. See ilgs CLI command documentation section in this guide.

Quotas

For quotas, all configuration information is replicated to target with the following **exceptions**:

- Usage Accounting - include Snapshot Data (best practice to not replicate this setting with SyncIQ)
- Usage Accounting - include Data-Protection Overhead (best practice to not replicate this setting with SyncIQ)

- Quota definitions that use user or group quota from local or filesystem are unique across clusters even if the same name is used. This means local/filesystem users and well known accounts should not be used with the Quota Replication feature. Active Directory should be used with user and group quota and custom or automatic quota replication jobs.

Zones

The following configuration properties are replicated for Zones other than the System zone:

- Access Zone Name
- Zone base directory
- Authentication Providers (name only)
- User Mapping rules (PowerScale 7.1.x.x only)

! Authentication providers / Local providers must exist on the target to resolve the SID' used in share permissions and the filesystem ACL's.

Snapshots

- Snapshot schedule properties (only for paths that match SyncIQ policies cluster wide) with the following exceptions:
- next_run (auto-populated by OneFS)
- next_snapshot (auto-populated by OneFS)

Dedupe

- Dedupe path for job processing (only for paths that match SyncIQ policies cluster wide)

- Dedupe assessment paths (only for paths that match SyncIQ policies cluster wide)

SyncIQ Policies

- On Failover only the SyncIQ schedule is re-applied to newly created mirror policies
- File pattern filters set on SyncIQ policies are **NOT** synced on failover, these patterns can result in unprotected data during failover and failback. Failover and failback work flows require customer testing. All file filter use cases are untested without support or documentation.

© Superna LLC

1.9. Advanced Quota Replication

Home Top

Customers with large numbers of quotas can leverage new features to better manage quota protection, failover and inventory functions. Large quota count above 5000 quotas should consider the following new features.

1. **Quota Inventory Collection job** - Move quota inventory to new scheduled inventory task to speed up share/export change detection job, and allow failover configuration sync function to be faster without the need to collect quota's. NOTE: Failover of many quotas will still take a long time. Best Practice is to use DR assistant option to skip quota failover. Open a case with support to use command line option to failover quota's after a failover that skipped quotas during the failover.
    a. **How to Enable quota inventory collection**
        i. igls admin schedules set --id QuotaInventoryCollection_2_5_3 --enabled true
        ii. check the schedule igls admin schedules
        iii. **NOTE: Default quota collection inventory schedule is every 12 hours. Change the schedule to collect quotas more often if required.**
2. **Pre Sync Quotas** - **Warning: many quotas on DR PowerScale WILL slow down replication performance of SyncIQ. This should be understood and factored into the decision to enable this feature.** This option allows quota's to be synced on the normal

Configuration replication job that runs every 5 minutes. This will treat quotas like shares and exports and will create new quotas or update existing quota's every 5 minutes and sync the quota to other DR cluster. This pre-stages the quota before failover and allows the quota scan job to complete before failover. This options is only supported on Onefs 8 clusters. This can overcome quota failover issues when quota scan job causes SyncIQ steps to fail during a failover. It will also help address failover and failback scheduled with 1 or 2 days between failover and back.

a. **How to Enable presync Quotas**

   i. Defaults to disabled

   ii. ssh to appliance as admin

   iii. vim /opt/superna/sca/data/system.xml

   iv. find the tag <quota_presync_on_interval>false</quota_presync_on_interval> in the <process> section

   v. change false to **true**

   vi. save the file

   vii. restart the sca

   viii. sudo -s

   ix. enter admin password

   x. run command systemctl restart sca

   xi. Done.

xii. NOTE: Once this is enabled, quota's will be synced during normal configuration sync task. If quota inventory schedule is set, quota's will be synced during quota collection task AND configuration sync task.

xiii. NOTE: Quota will be Created, Updated if changed, AND deleted from the target cluster with this setting enabled.

© Superna LLC

1.10. Manage Eyeglass Jobs

Home Top

- Overview
- View Job Definition and Status
- Edit Configuration
- Delete
- Enable/Disable
- Enable/Disable Skip Configuration
- Newly Detected SyncIQ policies Job Type Unconfigured
  - How to set job type for unconfigured jobs
- Enable/Disable Microsoft DFS
- Run Now Action next to Job Name
- Run Now Menu Item Bulk Actions
- Running Jobs
- Troubleshooting Eyeglass Job
- What are the Eyeglass System Alarms?
- Renaming SyncIQ Policy & Eyeglass Jobs

Overview

The Jobs Feature is used to keep track of Eyeglass tasks. Two Job views are provided :

- **Job Definitions**: Use the Job Definition view to see details for and manage Eyeglass configuration replications Jobs.
- **Running Jobs**: Use the Running Jobs view to see the statue of active tasks.

View Job Definition and Status

Job Definitions Window :

Job Name : It displays name as was defined during job creation.

This column displays a + symbol along with it which when expanded provides more details about the job such as :

> Enabled/Disabled : The Job can be enabled or disabled from the Jobs window itself. When it is disabled, this field displays "USERDISABLED" State, If it is enabled, it displays "ENABLED" . For Auto Type Jobs, If the Policy is disabled from OneFS, it displays "POLICY DISABLED".

Job Type : It display two Job types : "AUTO" and "MANUAL".

- **Auto:** For Jobs automatically detected by Eyeglass when a policy is created in OneFS and the Job is scheduled from there.
- **Manual:** For Jobs that were manually created from the Jobs window.

Source : It displays PowerScale cluster name configured as Source for a replication Job.

Source Path: The source path is used to provide a point in the file system for Eyeglass to autodetect shares, exports, and quotas, to discover and include in the job. A job can be edited, using Edit Configuration(s) to remove a detected configuration that you don't want included in the job.

Destination: It displays PowerScale cluster name configured as Destination for replication Job.

Last Success: This field displays date when the job was last successfully run. If the Job was never run successfully, this field is empty.

Last Run Date: It displays the date when the job was last run regardless of its last state was failed or success.

State: It displays the current state of Job.

- **"OK"** if the Last Run date and Last Success date matches.
- **"ERROR"** if the Last Run date and Last Success date doesn't match.
- **"SYNC IQ RUNNING"** if the associated SyncIQ Job is in running state

Edit Configuration

To view the list of associated shares/exports/nfs alias/quotas for an Eyeglass Job, select Jobs.

To see the shares/exports/nfs alias associated with your Eyeglass configuration replication Job, select the checkbox for your Job and then select the **Select a bulk action** button. The **Edit Configuration(s)** option will take you to the Jobs - Edit Configuration(s) View window where you can expand the tree to see which shares and exports and nfs alias are associated with your job.

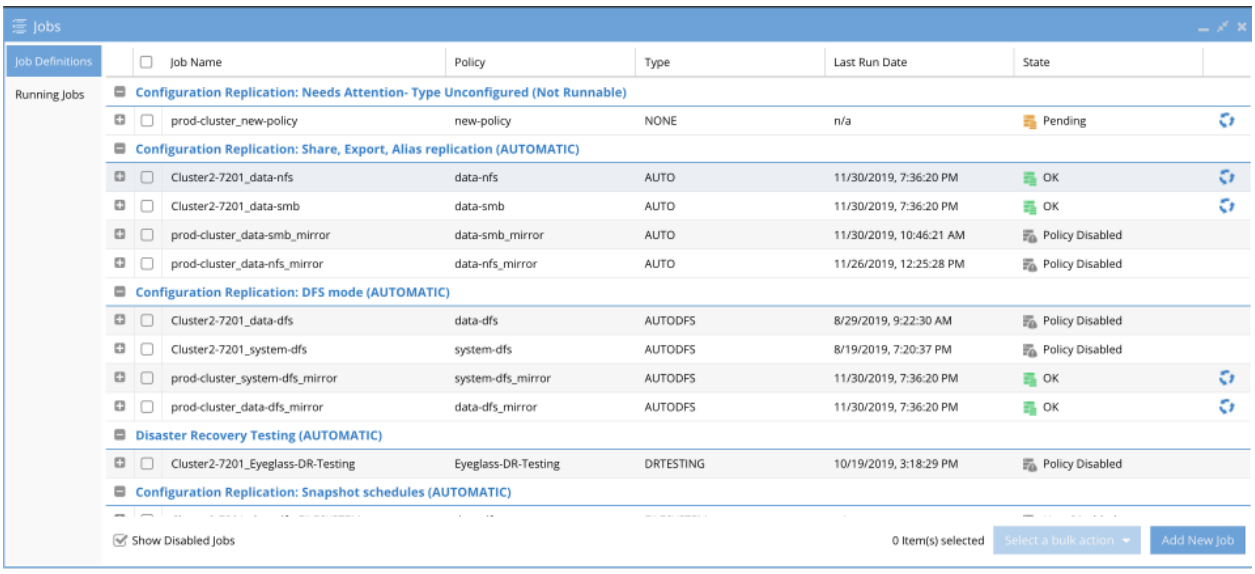

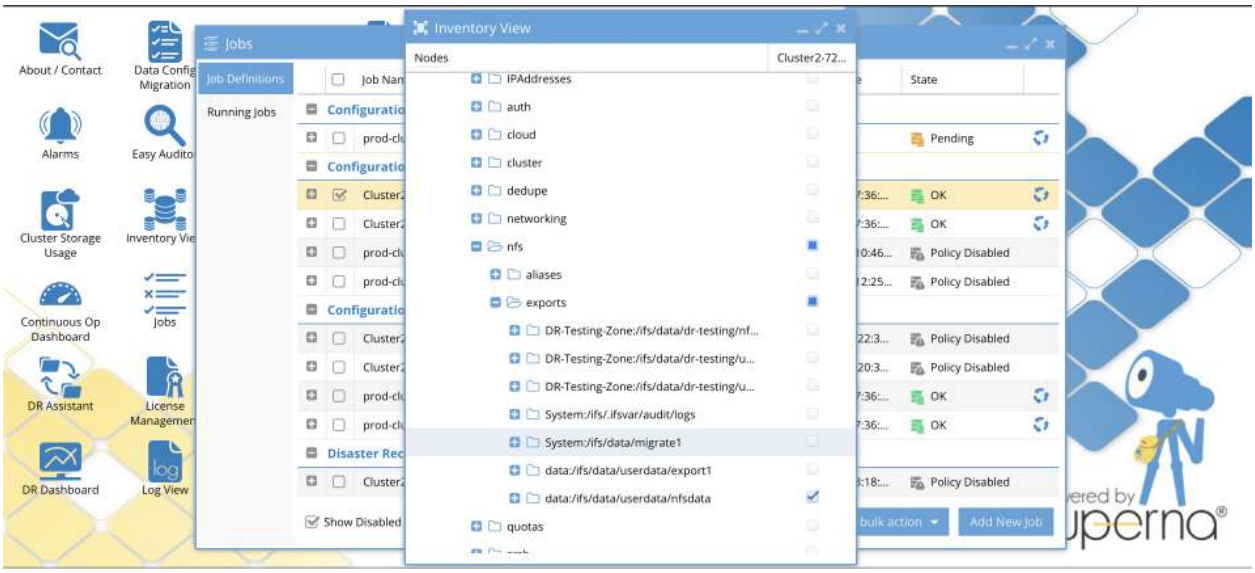

Delete

To delete manually created jobs and the jobs that are not currently running.

Enable/Disable

Able to enable/disable all type of jobs except Eyeglass Failover:Runbook Robot (AUTOMATIC) (TYPE = RUNBOOK).

Note: If disabled no actions will be taken on the policy and it will not be eligible for failover.

Enable/Disable Skip Configuration

Set this mode to skip syncing any configuration data on this policy but allow the policy to be eligible for failover. This mode would be used when 2 different sets of hosts will be used at production and DR site for NFS exports and the client lists will be different. This allows configuration to be managed manually on these policies.

Newly Detected SyncIQ policies Job Type Unconfigured

New in 2.5.6 and later releases, this allows an Unconfigured Job type to be set to Auto, AutoDFS, or AutoSkipConfig and then enabled for production use. All new policies detected will appear as unconfigured and require administrator to take action to enter production use.

An alarm will be raised when a new syncIQ policy is detected indicating that an action is required to set the DR job type for this new policy. This will generate an alarm every 5 minutes until the job type has been configured. If the policy is not in production, then set the job type and then set to user disabled.

How to set job type for unconfigured jobs

1. Login to Eyeglass
2. Open the Jobs icon
3. Select the unconfigured policy
4. Select bulk actions menu Set Type
5. Set the type
6. **NOTE: Once the job type is set, it will not be possible to set back to unconfigured status.**
7. 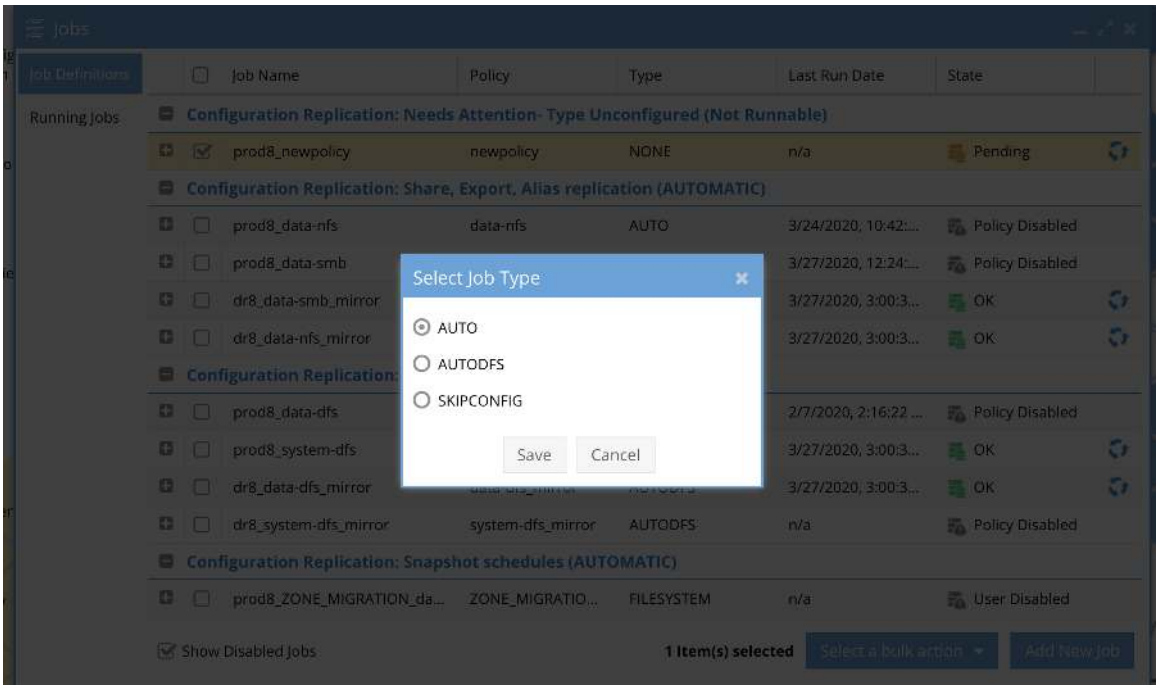

Enable/Disable Microsoft DFS

To enable/disable DFS mode for Configuration Replication: Share, Export, Alias replication type of jobs. Please refer to [Microsoft DFS Mode Failover Guide](#).

Run Now Action next to Job Name

New in 2.5.6 and later releases a run now icon next to jobs that are in a runnable state. This simplifies running configuration jobs or failover readiness jobs.

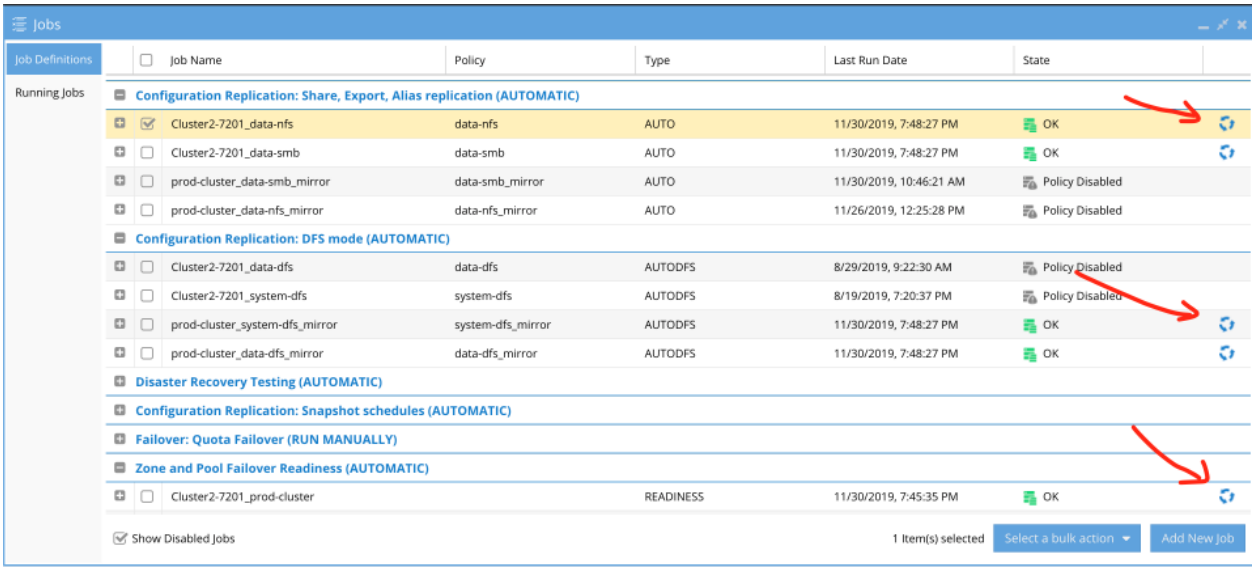

Run Now Menu Item Bulk Actions

Able to run Eyeglass Configuration Replication Jobs on demand.

1. Open the Jobs window.
2. Select an Eyeglass Configuration Replication Job.
3. Select a Job.
4. Select a bulk action.

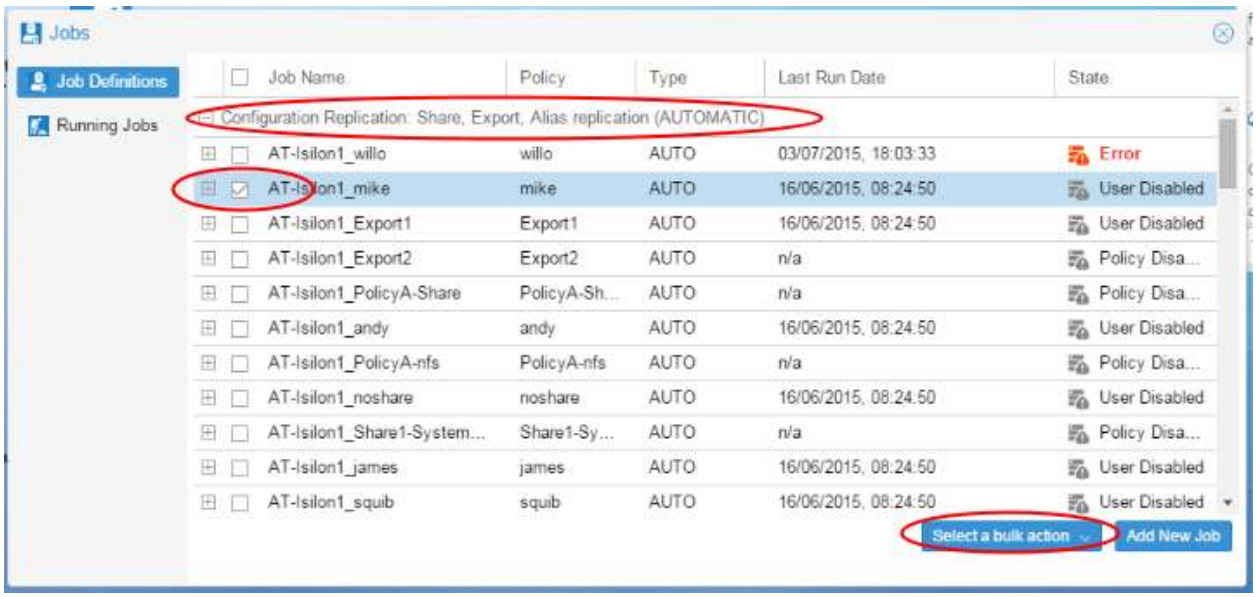

4. Elect Run Now. This will initiate the Eyeglass Configuration Replication task which will run all Eyeglass Configuration Replication Jobs.

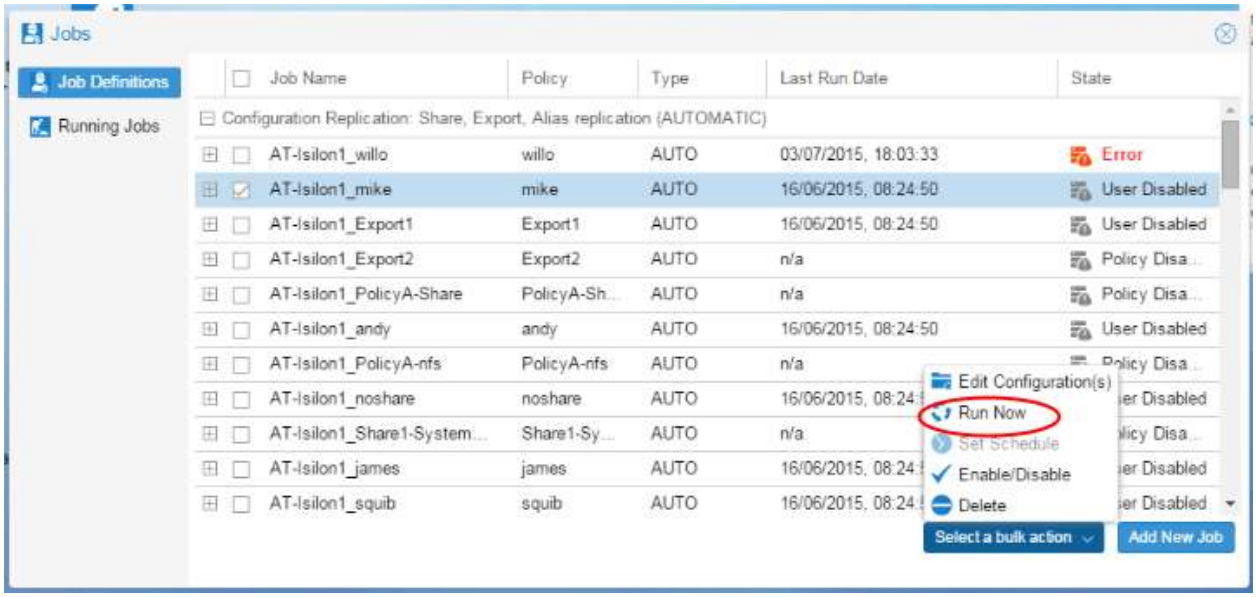

5. You can check the progress of the Configuration Replication Job in the Running Jobs view of the Jobs window.

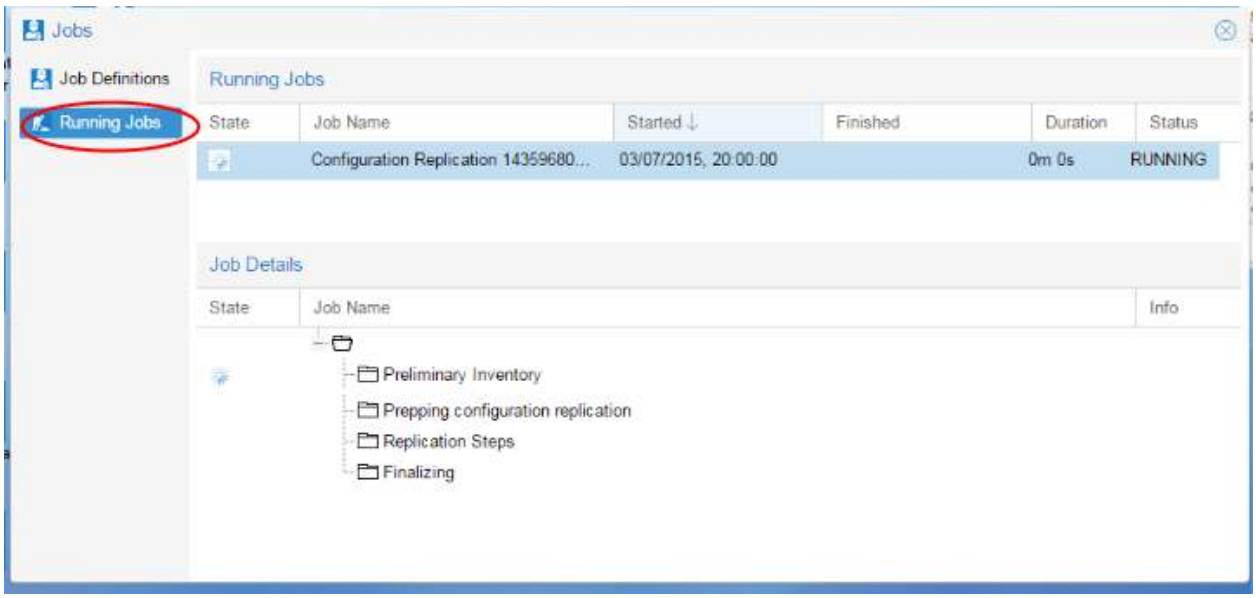

Running Jobs

This category displays the currently Running Jobs and its details such as :

Job Name : displays the name of Job

State : displays the state of Job : Finished or Running

Started : displays time when it was Started

Finished : displays time when it was Finished.

To view a Job in progress select **Running Jobs** from the **Jobs** window. In the bottom **Job Details** pane you will see the steps associated with the Job and the status for each step.

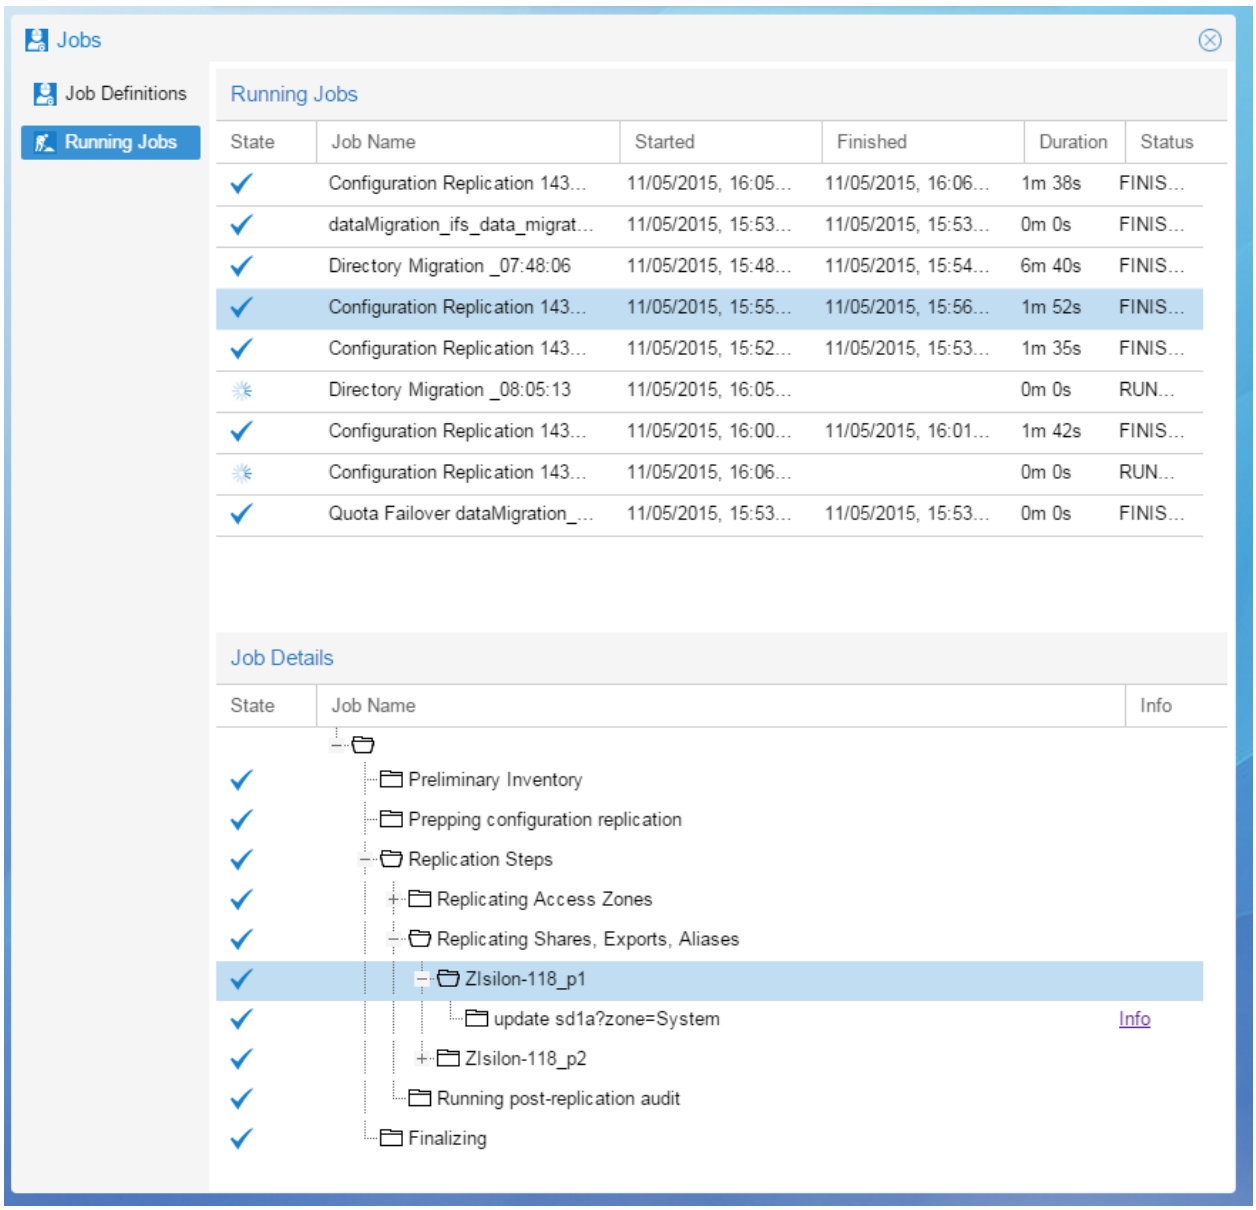

Troubleshooting Eyeglass Job

When a problem occurs, for a Job in progress the Running Jobs view will indicate at which step a problem occurred with a link to the information available for the error.

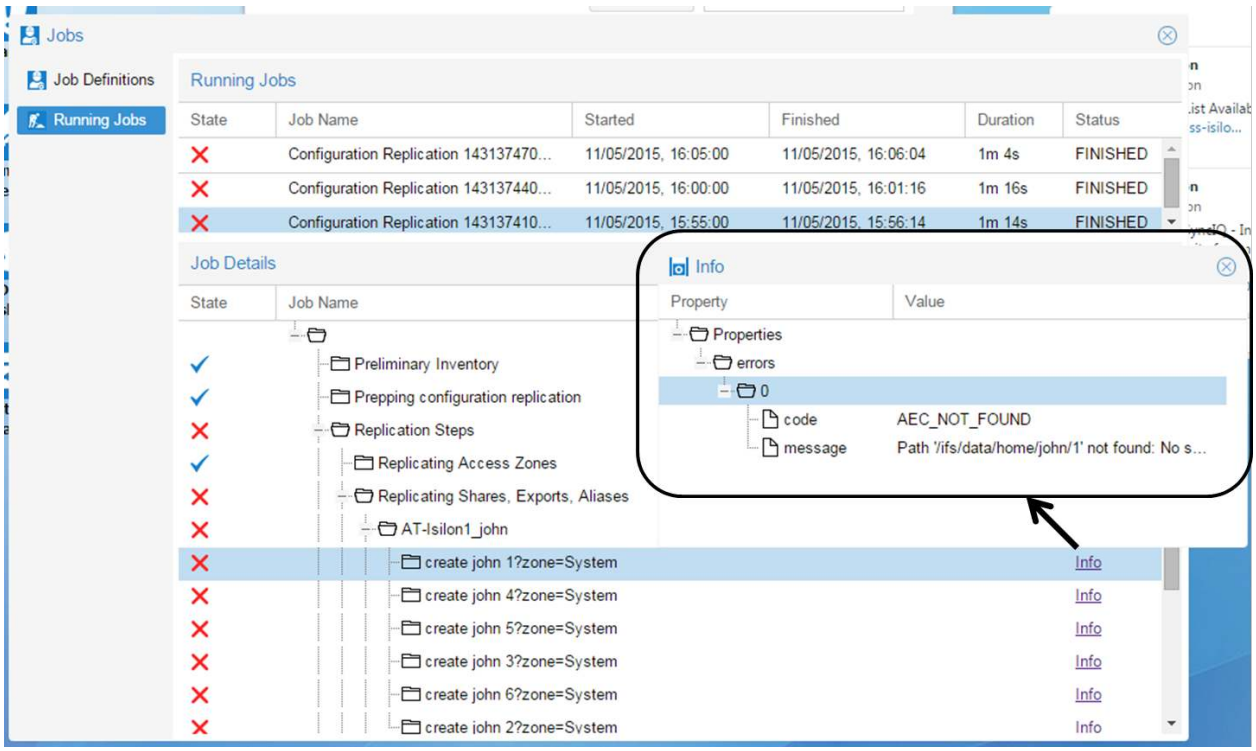

For a completed Job, the Eyeglass system alarm related to the failed Job may contain the extra information you need to troubleshoot the problem

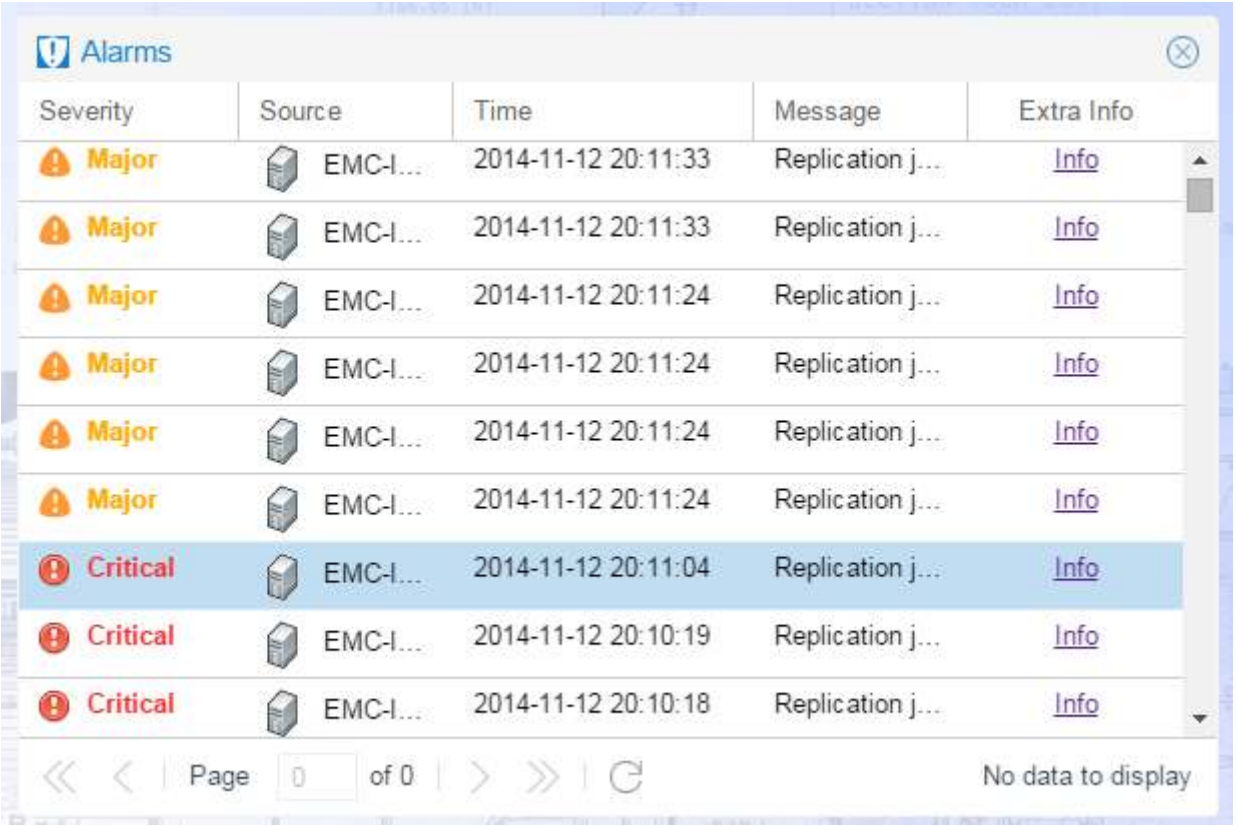

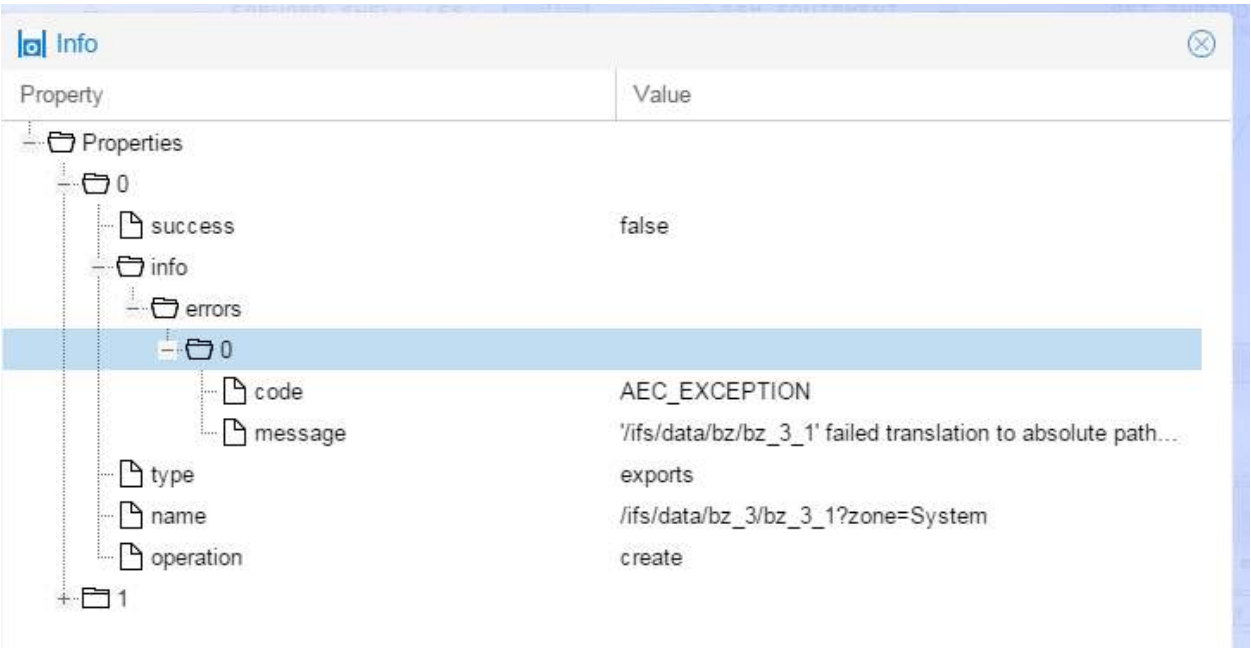

What are the Eyeglass System Alarms?

Refer to Eyeglass PowerScale Alarm codes for the list of system alarms.

Renaming SyncIQ Policy & Eyeglass Jobs

When a SyncIQ Policy is renamed, Eyeglass considers it to be a new SyncIQ Policy and therefore the following happens:

1. Eyeglass Job(s) related to the original SyncIQ Policy Name are deleted.
2. Eyeglass Job(s) for the new SyncIQ Policy Name are created but will be **set to Unconfigured state (2.5.6 or later release) and will need to be set to auto or autodfs and enabled.**

As a new Job, the Eyeglass Configuration Replication Jobs will have a state depending on the Eyeglass appliance default “Initialstate” configuration (see section igls adv initialstate in this Admin Guide). By default the Jobs will be disabled. Any other customizations to the Job such as DFS Mode must also be re-applied. For a SyncIQ Job that has been renamed, follow these steps in Eyeglass:

1. Login to the Eyeglass web page.
2. Open the Jobs window.
3. For any Jobs that were in DFS Mode previously, re-set to DFS Mode.
4. For any Jobs where sync of a share or export, and had been customized to be skipped, from the Jobs page, use Edit Configuration to deselect those objects from the Job.

In addition to the Job setup, RPO Reporting for the original SyncIQ Policy name will not be linked to RPO reporting for the new SyncIQ Policy name.

© Superna LLC

1.11. Failover Readiness Validations DR Dashboard

Home Top

- Monitor DR Readiness
- SyncIQ Audit Monitor - Verify Data is Replicating
  - Overview
  - Requirements
  - How it Works
  - Configuration
- Policy Readiness and DFS Readiness
  - Policy/DFS Readiness - OneFS SyncIQ Readiness
- New Validations in Release 2.0 and later
  - Previous Failed DFS failover share prefix
  - Policy Source Nodes Restriction
  - SyncIQ Policy Status Validations
  - Quota Domain Validation
  - SyncIQ File Pattern Validation
  - Corrupt Failover Snapshots
  - Duplicate SyncIQ Local Targets
  - Target Writes Disabled
  - Policy/DFS Readiness - Eyeglass Configuration Replication Readiness
  - Policy / DFS Readiness - Zone Configuration Replication Readiness

- Zone and IP Pool Readiness
    - New 2.5.6 or later releases
    - Zone and IP Pool Readiness - Zone Configuration Replication Readiness
    - Zone and IP Pool Readiness - Target Cluster Reachability
    - Zone and IP Pool Readiness - Date-Time Validation
    - Zone and IP Pool Path Validation
    - Access Zone and IP Pool FQDN Alias Validation
    - Access Zone and IP Pool - SPN (Service Principal Name) Active Directory Delegation Validation
        - Prerequisite: 2.5.6 or later
        - Documentation to Correct warnings
        - SPN AD Delegation Example
        - SPN test Failure Conditions and Error Messages
    - Access Zone and IP Pool - DNS Dual Delegation Validation
        - Prerequisite: 2.5.6 or later
        - Documentation to correct Warnings
        - UnSupported Configurations
        - Example Dual DNS validation
        - Warnings and Statuses for all Possible Dual Delegation Validations
            - Pool with ignore alias igls-ignore-
            - Zone or pool with correct settings:
            - SmartConnect Zone name/alias unknown:

    - One NS record for SmartConnect Zone name/alias:
    - Detected three NS records for SmartConnect Zone name/alias:
    - One of the NS record is incorrect for SmartConnect Zone name/alias:
    - There is no NS records for SmartConnect Zone name/alias:
    - Two NS records point to same IP addresses for SmartConnect Zone name/alias:
    - SSIP of source cluster (in cluster) is incorrect:
- All Failover type Domain Mark Validation
  - Prerequisite: 2.5.6 or later
  - Documentation to correct Warnings
- IP Pool Failover Readiness
  - Pool Validations
    - Pool Readiness Validation
    - Un-mapped policy validation Overview
    - Overall pool validation status
- How To Configure Advanced DNS Delegation Modes Required for Certain Environments
  - What's new
  - How to change the values
  - How the tags work

- How to Configure Advanced SPN delay mode for Active Directory Delegation
- How to Disable AD or DNS Delegation Validations - Advanced Option
- Network Visualization

Monitor DR Readiness

To view and manage your DR readiness, use the DR Dashboard window:

Policy Readiness: Provides a view of your readiness for SyncIQ Policy Failover.

Zone Readiness: Provides a view of your Access Zone Failover Readiness.

IP Pool Readiness: Provides a view of your IP pool failover Readdiness.

DFS Readiness: Provides a view of your readiness for a SyncIQ Policy DFS Mode Failover.

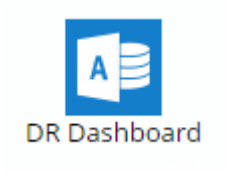

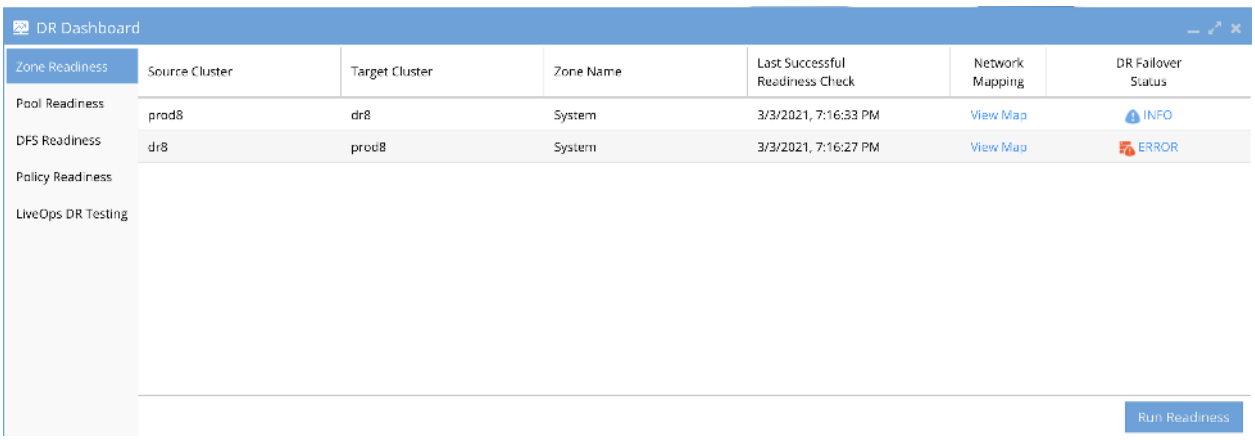

SyncIQ Audit Monitor - Verify Data is Replicating

Overview

Audits SyncIQ data sync by adding test data with timestamps to the source cluster and compares when synciq runs, and verifies timestamps on the target cluster. This provides the highest level of confidence in your off site data. The job is disabled by default and requires configuration.

Requirements

1. 2.5.7 or later

How it Works

1. A test file is created in the hidden directory with timestamp using the cluster file api
2. SyncIQ policies are monitored when they run successfully

3. After SyncIQ runs the remote DR cluster is checked with the file api to read the file and check the time stamp matches the source cluster.
4. This validation must pass, if it fails an alarm is raised to flag a data sync issue has been detected.
5. If the validation passes the job completes without any alerts being raised.
6. Successful Job monitor on a single SyncIQ policy configured for monitoring.

a. 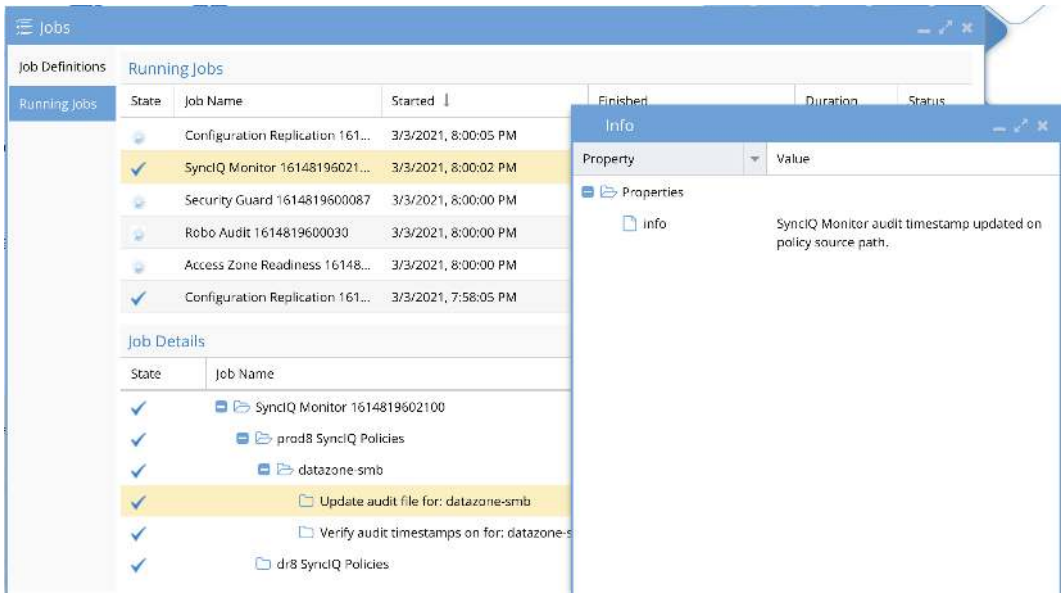

Configuration

To monitor a policies data integrity the following step need to be completed on each synciq policy. You can configure on any policy by following these steps. Only the policies that require monitoring need to follow these steps. All other steps are automated after completing these steps on a policy.

1. Login to the source cluster as the root

2. Create a hidden folder at the base of the synciq policy path example if the policy path is /ifs/data/userdata/smbdata
    a. mkdir -p /ifs/data/userdata/smbdata/.iglssynciqmonitor (not the . is required)
    b. chown eyeglass:wheel /ifs/data/userdata/smbdata/.iglssynciqmonitor (allows the eyeglass service account to create a test file in this folder)
3. Verify the synnciq monitor job is enabled
    a. ssh to eyeglass as admin
    b. sudo -s
    c. igls admin schedules set --id SyncMonitorTask --enabled true
    d. igls admin schedules (verify the job is enabled and default schedule is setup for hourly)
4. If the monitor detects an issue with data sync an alarm will be raised
    a. 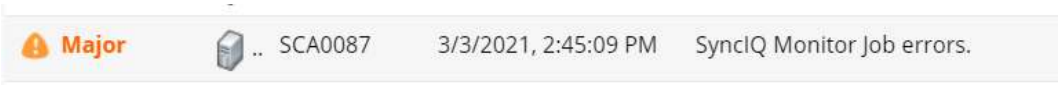

5.

Policy Readiness and DFS Readiness

The DR Failover Status on the Policy Readiness or DFS Readiness tabs provides you with a quick and easy way to assess your Disaster Recovery status for a SyncIQ Policy Failover (non-DFS mode) or SyncIQ Policy Failover (DFS mode).

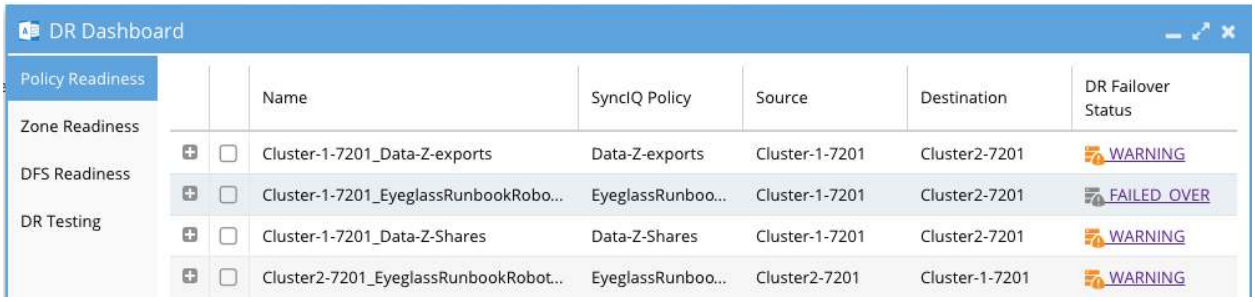

Each row contains the following summary information:

Column	Description	Notes
Name	Name of the Eyeglass configuration Replication Job	Eyeglass configuration replication Job created automatically for each SyncIQ Policy detected with exactly the same name + prefixed with PowerScale Cluster name. Quota Jobs are suffixed with “quotas”.
SyncIQ Policy	Name of the SyncIQ Policy associated with the Eyeglass Job	
Source	The PowerScale cluster that is the source configured in the SyncIQ Policy	Eyeglass Job will have same source
Destination	The PowerScale cluster that is the target configured in the SyncIQ Policy	Eyeglass Job will have same target
DR Failover Status	A status calculated by Eyeglass based on the failover validation criteria for the Policy or DFS Failover mode	This column displays the overall status. Select the link to see individual status for each validation criteria. **New option shows “Failed Over” for any policy that is read-only status in SyncIQ.**

Click on the DR Failover Status link to see the details for each of the failover validation criteria for a selected Job.

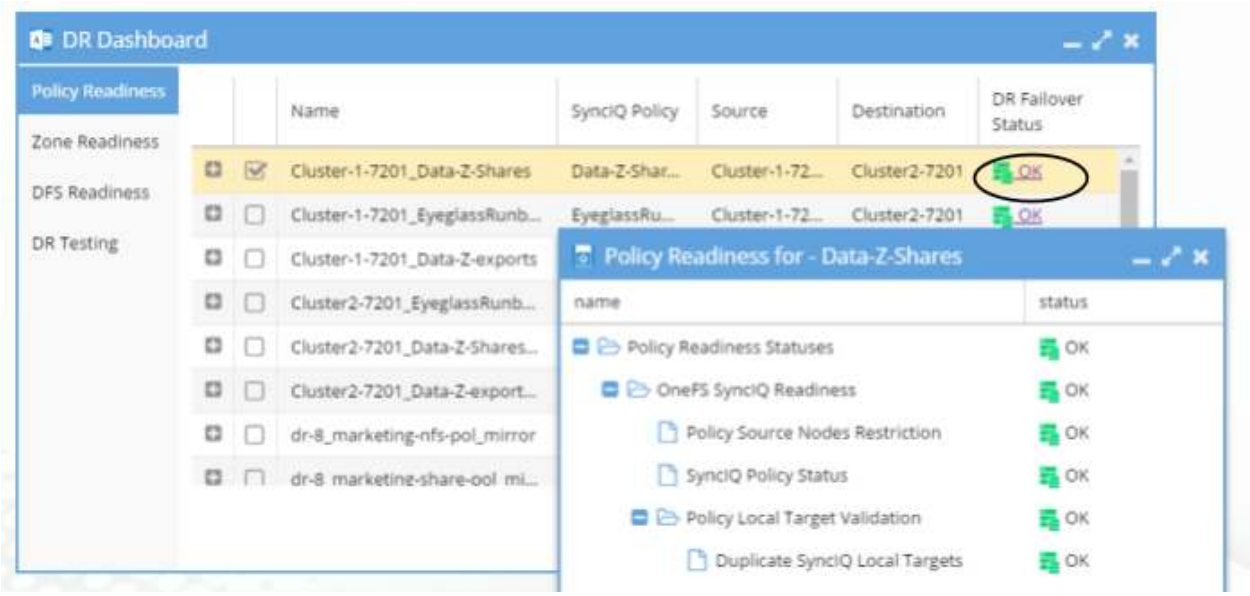

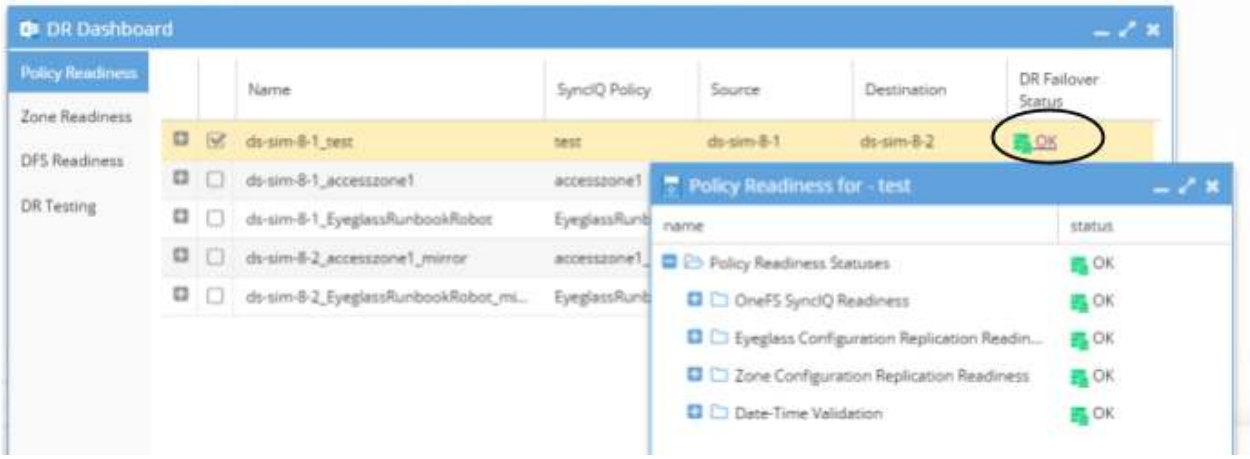

The DR Failover Status will be one of the following:

Status	Failover Impact
OK	Able to Failover
WARNING	Warning state does NOT block failover. **IMPORTANT: While failover is not blocked, the issue(s) causing this Warning may cause the failover to fail. Recommendation is to understand and resolve these issues prior to failover.**
ERROR	Error state BLOCKS failover.
DISABLED	Disabled state DOES block failover.
FAILED OVER	Failed Over state DOES block failover.

The DR Failover Status is based on Status for each of the following areas for Policy and DFS Failover:

OneFS SyncIQ Readiness	Have the SyncIQ policies in the Access Zone been successfully run and are they in a supported configuration?
Eyeglass Configuration Replication Readiness	Have the Eyeglass Configuration Replication jobs in the Access Zone been successfully run to sync configuration data for all policies that are members of the Access Zone?
Zone Configuration Replication Readiness	Has the Eyeglass Zone Configuration Replication job been successfully run to create target cluster Access Zones that don't already exist for configuration sync complete?
Date-Time Validation	Are the date-time differences between nodes and between Eyeglass and the clusters within an acceptable range that will not affect SyncIQ operations?

Additional information for Policy / DFS Readiness criteria is provided in the following sections.

Policy/DFS Readiness - OneFS SyncIQ Readiness

The OneFS SyncIQ Readiness criteria is used to validate that the SyncIQ policy has been successfully run and that it is in a supported configuration. For each SyncIQ Policy the following checks are done:

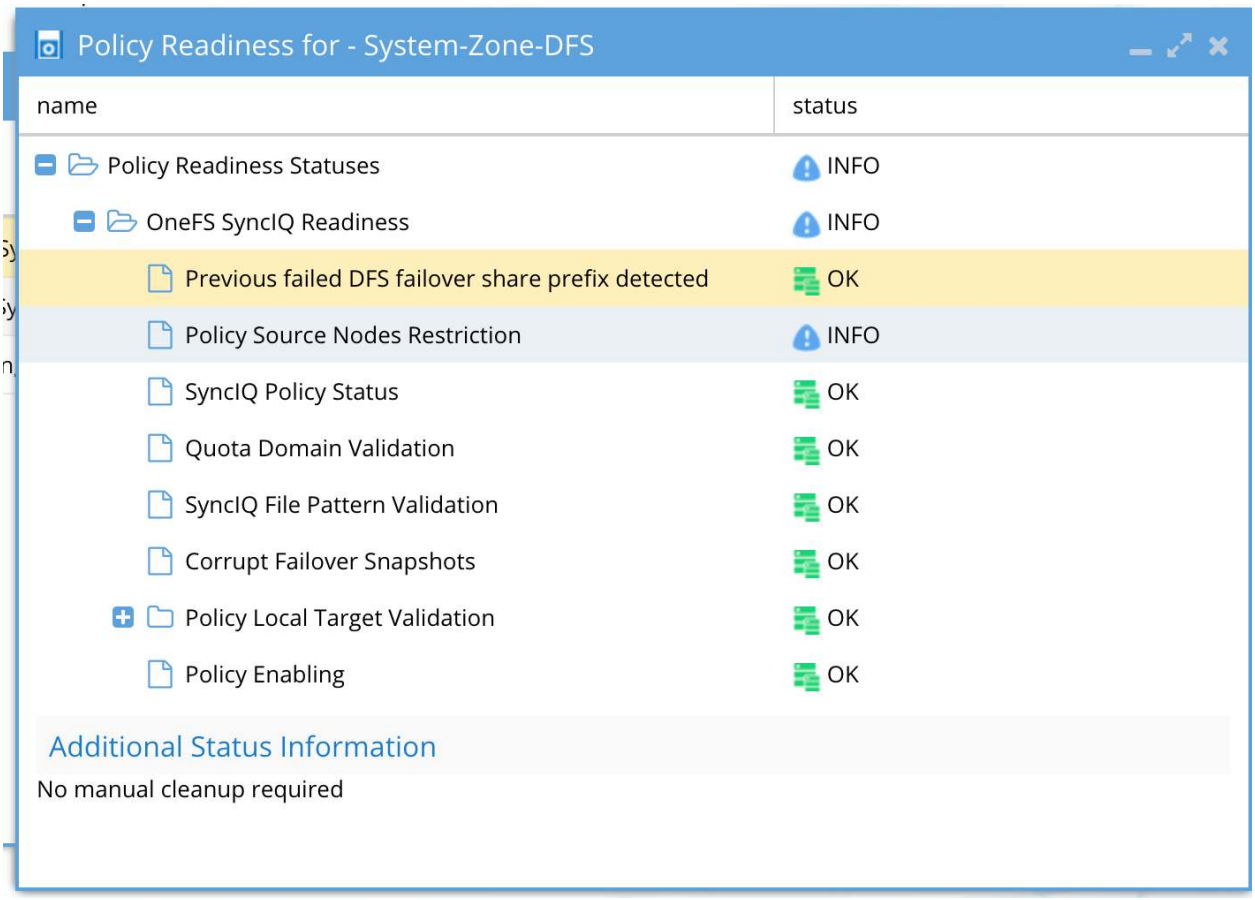

New Validations in Release 2.0 and later

SyncIQ Policy:	Notes:
Previous Failed DFS failover share prefix	This detects any prefixed dfs shares on the active cluster. The active cluster should not have any prefixed share on the SyncIQ policy. This validation will indicate if prefixed shares are detected that should be cleaned up prior to any failover.
Domain Mark Validation	Domain Mark validation applies to all failover types and is described here. **DR Status is Warning when the validation fails. Failover should not be started with this validation warning.**
Policy Source	Validate PowerScale best practices that recommend that SyncIQ Policies utilize the Restrict Source Nodes option to control which nodes replicate between clusters.

Nodes Restriction	
SyncIQ Policy Status Validations	DR Status is “OK” when all of the conditions below are met: • Your SyncIQ Policy is enabled. • Your SyncIQ Policy last state was finished or needs attention. DR Status is “Warning” when the condition below is met: • SyncIQ Policy has a last state that was not successful. • SyncIQ Policy has a last state that was paused or canceled. • SyncIQ Policy does not not have a last state (has never been run). • SyncIQ Policy has Excluded Directories and/or Included Directories configured. **IMPORTANT: SyncIQ Policy in Warning state MAY NOT be able to be run by Eyeglass assisted failover depending on it’s current status. If not run, you will incur data loss during failover.** • Example 1: SyncIQ Policy has an error state. If it cannot be run from the OneFS, it will also not be able to run from Eyeglass. • Example 2: SyncIQ Policy is paused. Eyeglass failover cannot RESUME a paused SyncIQ Policy - this must be resumed from OneFS **You must investigate these errors and understand their impact to your failover solution.** DR Status is “Disabled” when either of the conditions below are met: • Eyeglass configuration replication Job is user disabled. • Or the SyncIQ policy in OneFs is disabled.
Quota Domain	Detects a quota with needs scanning flag set. This flag will fail SyncIQ steps (run policy, Make Writable, and Resync prep) for any policy with a quota that has not been scanned and is missing a quota domain. **Failover should not be started**

Validation	**with this validation warning. Quota scan job should be run manually or verify if a quota scan job is in progress.**
SyncIQ File Pattern Validation	SyncIQ policies with file patterns set cannot be failed back and any files that do not match the file pattern will be read-only after failover. This file pattern is not failed over by Resync prep to mirror policies and Eyeglass does not support copying file access patterns to mirror policies. This setting should not be used for DR purposes. **This validation will show warning for any policy with a file pattern set. The file pattern should be removed from the policy to clear the warning.**
Corrupt Failover Snapshots	Validate that Target Cluster does not have an existing SIQ-<policyID>-restore-new or SIQ-<policyID>-restore-latest snapshot from previous failovers/synciq jobs for the Policy.
Policy Local Target Validation : • Duplicate SyncIQ Local Targets	Validate that there is only 1 Local Target per SyncIQ policy.
Policy Local Target Validation: • Target Writes Disabled	Validate that the target folder of SyncIQ policy has writes disabled.
Policy Enabling	Validate that the SyncIQ Policy is enabled in OneFS. If Disabled overrall DR Status is Disabled.

Policy/DFS Readiness - Eyeglass Configuration Replication Readiness

The Eyeglass Configuration Replication Readiness criteria is used to validate that the Eyeglass Configuration Replication job related to the SyncIQ Policy has been successfully run to sync the related configuration data. For the Eyeglass Configuration Replication Job the following check is done:

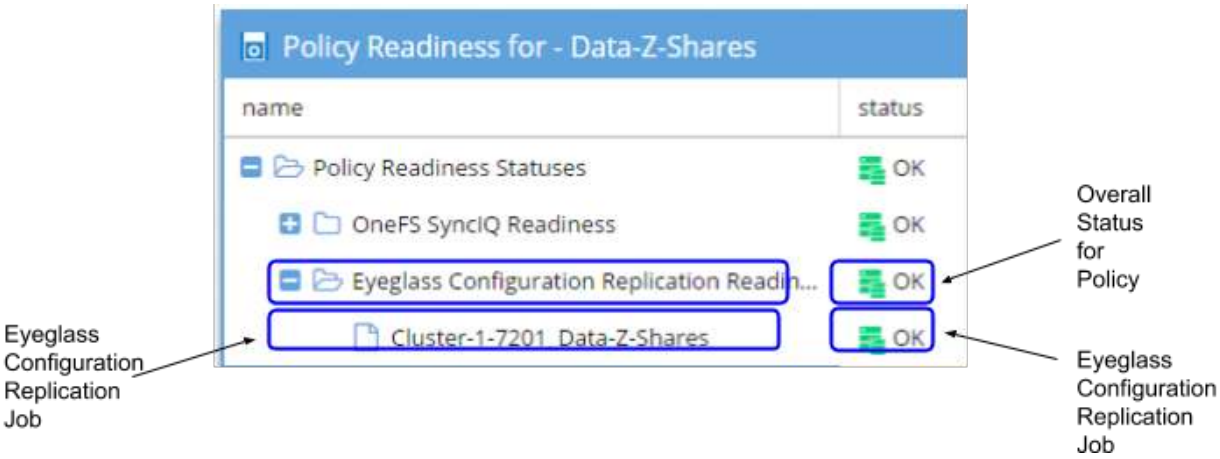

<Eyeglass Configuration Replication Job Name>	Validate that the Eyeglass Configuration Replication Job status is not in error state. DR Status is “OK” when: • Your Eyeglass configuration replication Job is enabled. • Your Eyeglass configuration replication Job Last Run and Last Success timestamp are identical. • Your Eyeglass configuration replication Job Audit Status is OK. DR Status is “Warning” when either of conditions below are met: Eyeglass configuration replication Job is new and has not been run yet.Eyeglass configuration replication Job has been successfully executed but not audited. Eyeglass configuration replication Job Last Run was not successful.Eyeglass configuration replication Job Audit was not successful. **DR Status is Warning when the validation fails. Failover should not be started with this validation warning.** DR Status is “Disabled” when either of the conditions below are met: • Eyeglass configuration replication Job is user disabled • OR the SyncIQ policy in OneFs is disabled

Policy / DFS Readiness - Zone Configuration Replication Readiness

The Zone Configuration Replication Readiness criteria is used to validate that the Zone Configuration Replication job related to the

SyncIQ Policy has successfully run. This is in order to create a new target cluster Access Zone for configuration sync completion.

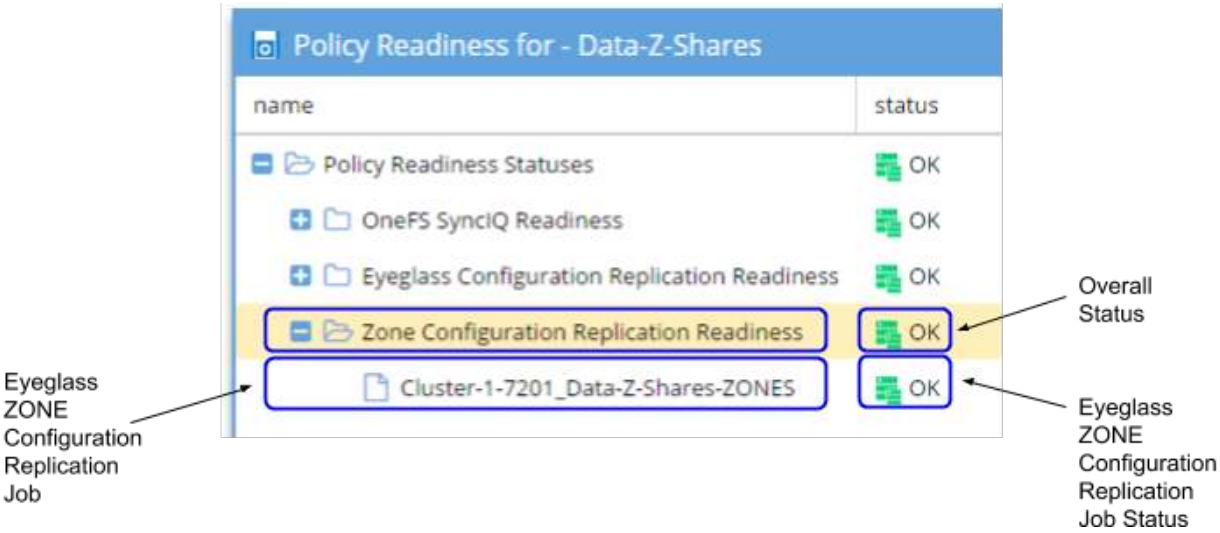

Policy/DFS Readiness - Date-Time Validation

The Date-Time Validation is used to validate that the time difference between the cluster nodes, and between clusters and Eyeglass, are within an acceptable range that will not affect SyncIQ operations. SyncIQ commands, like re-sync prep, can fail if the time between cluster nodes is greater than the time between the Eyeglass VM and the cluster with regards to latency of issuing the API call (a scenario where a node returns a timestamp for a step status message is earlier than the beginning of the Re-sync prep request). API calls can return different completion times between clusters. Differences here can cause re-sync prep failover commands to fail, if the difference between Eyeglass and the source cluster is greater than the time it takes for a resync prep command to complete.

NOTE: This condition under which timing differences cause resync failover commands to fail is rare and hard to detect manually. In release 1.8 Eyeglass can detect this condition. It's best practice to use

NTP on clusters and the Eyeglass appliance. This allows failover logs and the new feature in release 1.8 or later to collect cluster SyncIQ reports for each step and append to the failover log. This will make debugging multi step multi cluster failover simpler. This process will require time to be synced.

For each Cluster the following checks are done:

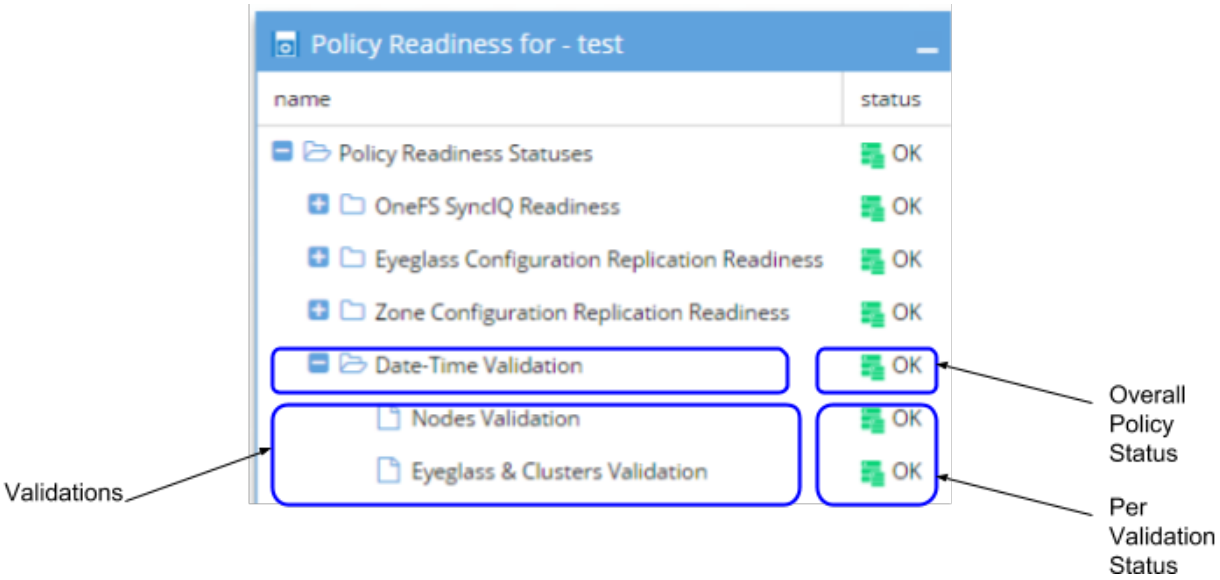

Date-Time Validation	Notes:
Nodes Validation	Validation that the maximum time difference between the nodes of a cluster is less than the time required for the cluster node time request made by Eyeglass to complete. DR Status Validation is Warning if fails. Typically you may proceed with this Warning.
Eyeglass & Clusters Validation	**NOTE:** Only applicable if Nodes Validation is OK. Validation that the earliest node time for a cluster and the Eyeglass appliance time are less than the time required for the cluster node time request made by Eyeglass to complete plus a default additional skew factor (default 1s). Executed if Nodes Validation is OK. DR Status Validation is Warning if fails. Typically you may proceed with this Warning.

DR Dashboard Job Details

Each Policy or DFS Job can be expanded in the DR Dashboard Policy Readiness or DFS Readiness view to see Job Details:

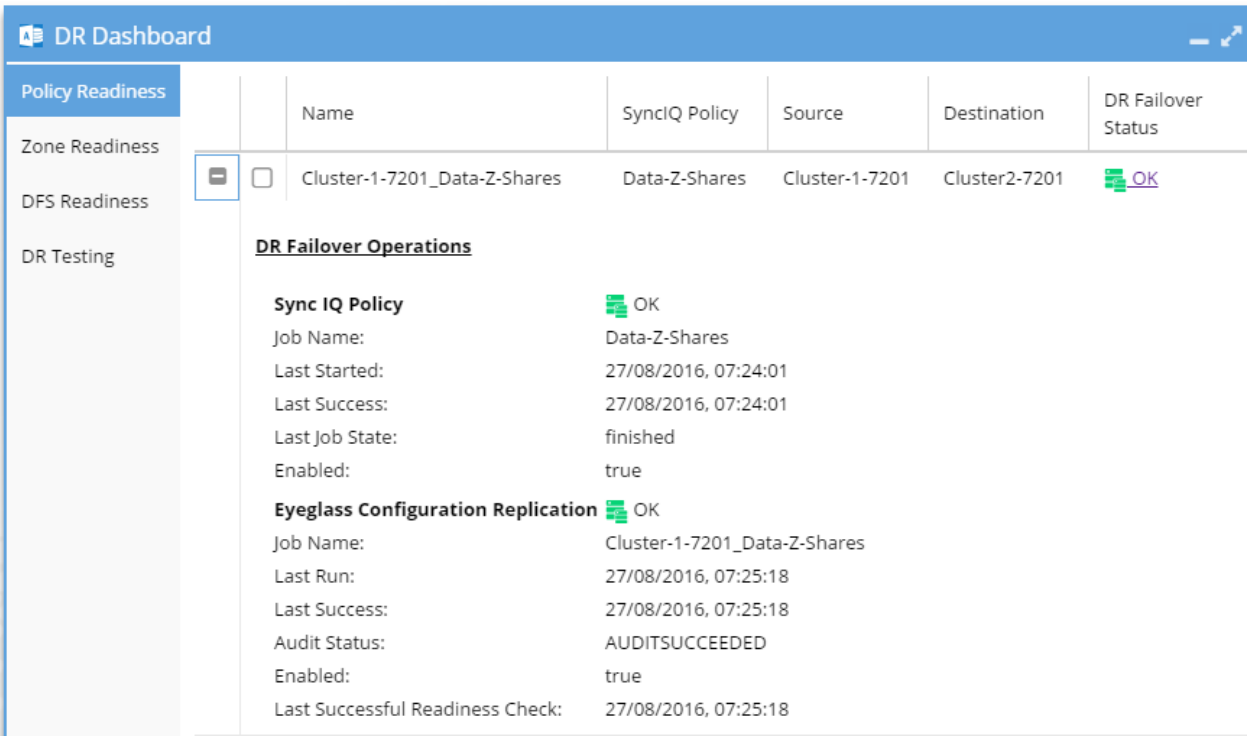

DR Dashboard Configuration Replication Job Details

Column	Description	Notes
SyncIQ Policy	All information in the SyncIQ Policy details section comes from the PowerScale Cluster itself.	If this section is empty then the Job is a custom Eyeglass job not associated with a SyncIQ policy.
Job Name	Name of the SyncIQ Policy.	Same name on the PowerScale cluster.
Last Started	Date/time when the last SyncIQ Policy job was started.	Information retrieved from SyncIQ Policy details on the PowerScale Cluster.
Last Success	Date/time when SyncIQ Policy last ran successfully on the PowerScale.	Information retrieved from SyncIQ Policy details on the PowerScale Cluster.
Last Job State	Status for Last SyncIQ Policy Job.	Information retrieved from SyncIQ Policy details on the PowerScale Cluster used to determine Overall DR Status.
Enabled	Indicates whether the PowerScale SyncIQ policy is enabled.	
Eyeglass Configuration Replication	All information in the Eyeglass Configuration Replication details section comes from Eyeglass.	
Job Name	Name of the Eyeglass Configuration Replication Job.	Eyeglass Configuration Replication Job created automatically for each SyncIQ

		Policy detected with exactly the same name and prefixed with PowerScale Cluster name. Quota Jobs also suffixed with “quotas”.
Last Run	Date/time when the last Eyeglass Configuration Replication Job was started.	Information used to determine Overall DR Status.
Last Success	Date/time when Eyeglass Configuration Replication Job was last run successfully.	Information used to determine Overall DR Status.
Audit Status	Indicates status of the Eyeglass Configuration Replication Job Audit.	After the Eyeglass Configuration Replication Job has completed, Eyeglass performs an audit to compare source and destination configuration and ensures that replicated configurations are identical.
Enabled	Indicates whether the Eyeglass Configuration Replication Job is enabled.	
Last Successful Readiness Check	Date/time when Eyeglass last successfully ran the Readiness Check Job.	

Zone and IP Pool Readiness

The Zone Readiness DR Failover Status provides you with a quick and easy way to assess your DR Status for an Access Zone Failover. The Zone Readiness check is performed for both directions of a replicating PowerScale cluster pair to ensure that you have your status, not only for failover but also for failback.

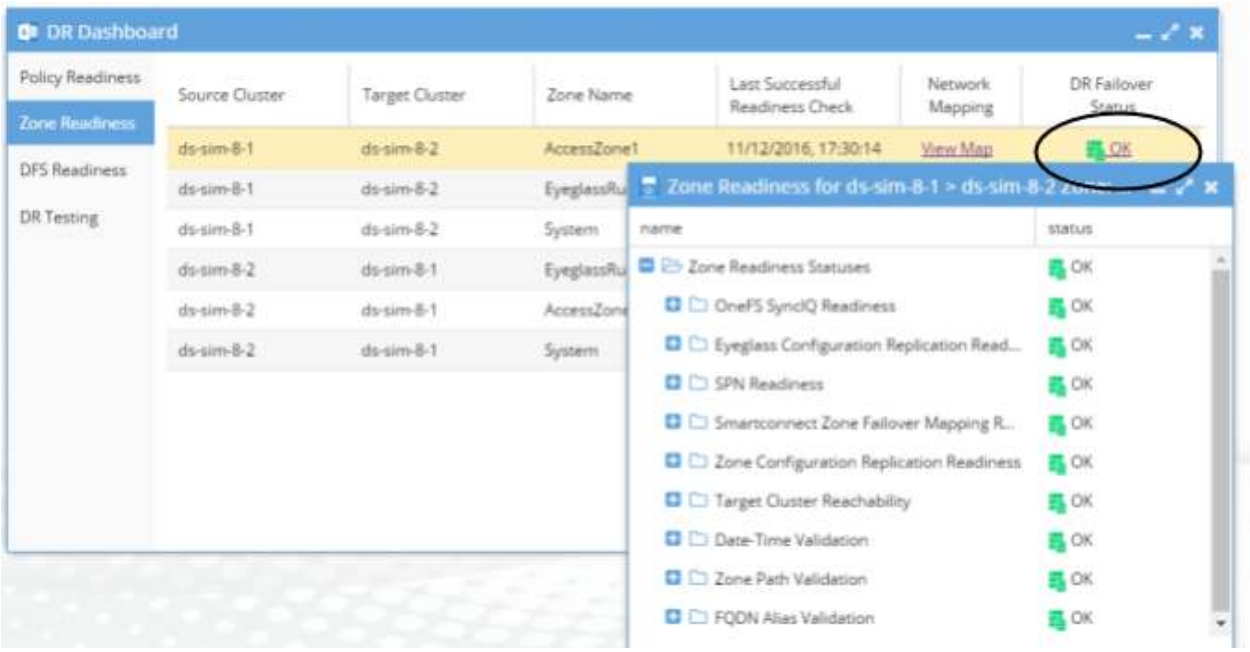

The Zone Readiness Status will be one of the following:

Status	Description
OK	All Required and Recommended conditions that are validated by Eyeglass software have been met.
WARNING	One or more Recommended conditions that are validated by Eyeglass software have not been met. Warning state does **NOT** block failover. Review the Access Zone Failover Guide Recommendations to determine impact for recommendations that have not been met.
ERROR	One or more of the Required conditions that are validated by Eyeglass software have not been met. Error state **DOES** block failover. Review the Access Zone Failover Guide Requirements for failover to determine resolution for these error conditions.
FAILED OVER	This Access Zone on this cluster has been failed over. You will be blocked from initiating failover for this Access Zone on this Cluster.

IMPORTANT: Not all conditions are validated by Eyeglass software. Please refer to the Eyeglass Access Zone Failover Guide for complete list of requirements and recommendations.

Notes:

1. For the case where the Target cluster pool that has the Eyeglass hint mapping for failover does not have a SmartConnect Zone defined:

1. On Failover the Access Zone will be in Warning state due to SPN inconsistencies.

2. On Failback it will not have the FAILED OVER status displayed.

2. For the case where there is no Eyeglass Configuration Replication Job enabled in an Access Zone there will be no entry in the Zone Readiness table for that Access Zone.

3. Until Configuration Replication runs, Policy Readiness for a policy in the Access Zone will be in Error.

The DR Failover Status is based on Status for each of the following areas for Access Zone Failover:

OneFS SyncIQ Readiness	Have the SyncIQ policies in the Access Zone been successfully run and are they in a supported configuration.
Eyeglass Configuration Replication Readiness	Have the Eyeglass Configuration Replication jobs in the Access Zone been successfully run to sync configuration data for all policies that are members of the Access Zone.
SPN Readiness	Is Active Directory delegation completed for cluster machine accounts to detect missing SPNs and remediate existing and newly created SmartConnect Zones as short and long SPN's created for cluster Active Directory machine accounts.
SmartConnect Zone Failover Mapping Readiness	Validation that confirms if all IP pools in the Access Zone have an Eyeglass hint (SmartConnect alias using igls syntax). Each SmartConnect Zone name associated to the IP pools (and any SmartConnect aliases) must be mapped to a target cluster IP pool prior to any failover. This ensures all SmartConnect names used to access source cluster data will failover to a target cluster IP pool. It is best practice and a requirement to create IP pools in matched pairs on source

	and destination cluster.
SmartConnect/IP Pool Readiness	SmartConnect/IP Pool Failover Readiness provides the status of whether the IP pool is ready for the failover, or has been already failed over. It also verifies each IP pool has a SmartConnect name applied. This validation will be used for IP pool based failover in addition to Access Zone failover where all pools must have a SmartConnect name defined.
Zone Configuration Replication Readiness	Has the Eyeglass Zone Configuration Replication job been successfully run to create target cluster Access Zones that don't already exist for configuration sync complete.
Target Cluster Reachability	Is Eyeglass able to connect to the Failover Target Cluster using API
Date-Time Validation	Are the date-time differences between nodes and between Eyeglass and the clusters, within an acceptable range that will not affect SyncIQ operations.
Zone Path Validation	Zone Path Validation provides the status of whether Access Zones have colliding paths. Status of OK indicates that the Access Zone paths have no conflicts. Status of ERROR indicates that this Access Zone collides with another Access Zone's path.
FQDN Alias Validation	If a cluster was added to Eyeglass with FQDN SmartConnect name for management, this SmartConnect zone must have an igls-ignore hint applied to avoid a failover impacting Eyeglass access. An ERROR means no igls-hint was found on the IP pool for the SmartConnect zone used for cluster management. OK means igls-ignore hint was found.

IMPORTANT

By default the Failover Readiness job which populates this information is disabled. Instructions to enable this Job can be found in the Eyeglass PowerScale Edition Administration Guide.

Note: If there are no Eyeglass Configuration Replication Jobs enabled there is no Failover Readiness Job.

Preparation and planning instructions for Zone Readiness can be found in the Access Zone Failover Guide:

Requirements for Eyeglass Assisted Access Zone Failover

Unsupported Data Replication Topology

Recommended for Eyeglass Assisted Access Zone Failover

Preparing your Clusters for Eyeglass Assisted Access Zone Failover

Additional information for Zone Readiness criteria is provided in the following sections.

Zone and IP Pool Readiness - OneFS SyncIQ Readiness

The OneFS SyncIQ Readiness criteria is used to validate that the SyncIQ policies in the Access Zone has been successfully run and that they are in a supported configuration. You will find one entry per SyncIQ Policy in the Access Zone. For each SyncIQ Policy the following checks are performed:

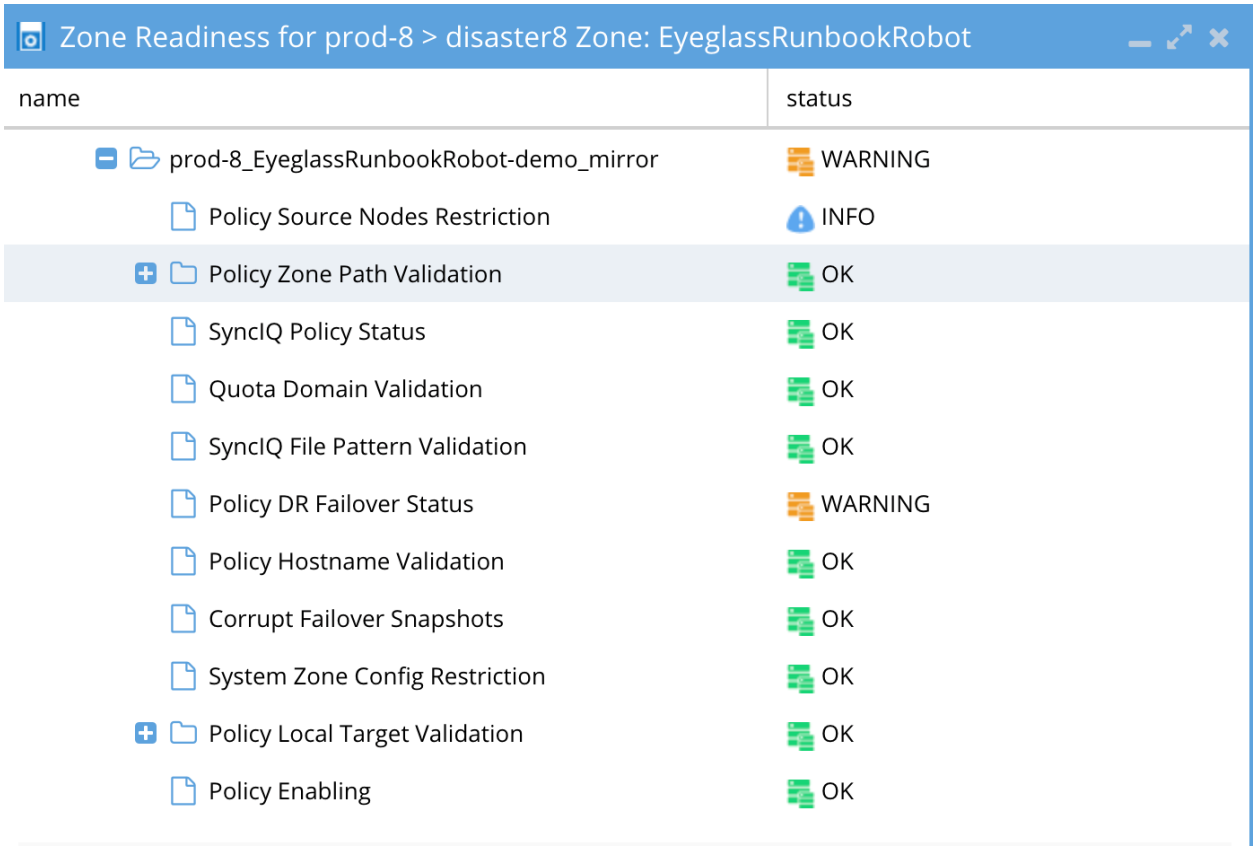

New Validations in Release 2.0 and later

SyncIQ Policy Check	Notes:
Policy Hot/Hot Validation	For Hot-Hot (Active-Active data) replication topology, validate that there is a dedicated Access Zone for each replication direction.

Policy Zone Path Check • Policy Source Path Check • Policy Target Path Check	Validate that the SyncIQ Policy(s) source root and target directories are at or below the Access Zone Base Directory.
Policy Source Nodes Restriction	Validate PowerScale best practices that recommend that SyncIQ Policies utilize the Restrict Source Nodes option to control which nodes replicate between clusters.
Policy Hostname Validation	Validate the the SyncIQ Policy target host hostname is associated with a subnet pool that is NOT going to be failed over.
Corrupt Failover Snapshots	Validate that Target Cluster does not have an existing SIQ-<policyID>-restore-new or SIQ-<policyID>-restore-latest snapshot from previous failovers/synciq jobs for the Policy.
System Zone Config Restriction	Validate that all shares, exports and alias have been created in the Access Zone that is being failed over. It is not supported to have shares, exports and alias with a path that is outside (higher in the file system) than the Access Zone base path.
Policy Enabling	Validate that the SyncIQ Policy is enabled in OneFS.
Quota Domain Validation	Detects a quota with needs scanning flag set. This flag will fail SyncIQ steps (run policy, Make Writable, and Resync prep) for any policy with a quota that has not been scanned and is missing a quota domain. **Failover should not be started with this validation warning. Quota scan job should be run manually or verify if a quota scan job is in progress.**
SyncIQ File Pattern Validation	SyncIQ policies with file patterns set cannot be failed back and any files that do not match the file pattern will be read-only after failover. This file pattern is not failed over by Re-sync prep to mirror policies and Eyeglass does not support copying file access patterns to mirror policies. This setting should not be used for DR purposes. **This validation will show warning for any policy with a file pattern set. The file pattern should be removed from the policy to clear the warning.**
Policy Status	Validate that the SyncIQ Policy is not in error state in OneFS.
Policy Local Target Validation : • Duplicate SyncIQ Local Targets	Validate that there is only 1 Local Target per SyncIQ policy.
• Policy Local Target Validation: • Target Writes	Validate that the target folder of SyncIQ policy has writes disabled.

Disabled	

Zone and IP Pool Readiness - Eyeglass Configuration Replication Readiness

The Eyeglass Configuration Replication Readiness criteria is used to validate that the Eyeglass Configuration Replication jobs in the Access Zone have been successfully run, to sync configuration data for all policies members of the Access Zone. For each Eyeglass Configuration Replication Job in the Access Zone, the following check is performed:

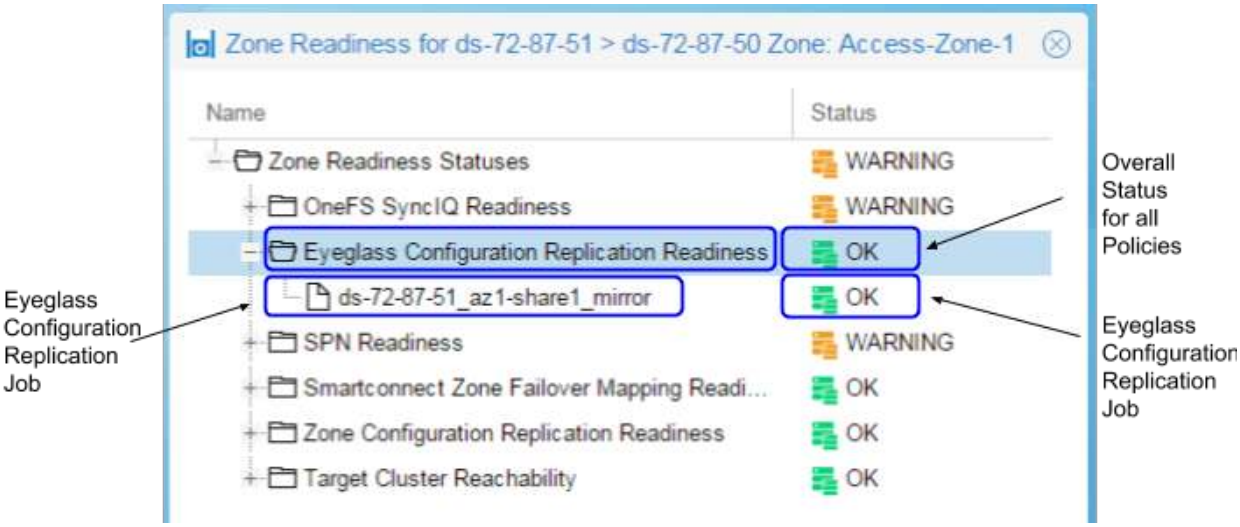

<Eyeglass Configuration Replication Job Name>	Validate that the Eyeglass Configuration Replication Job status is not in the ERROR state.

Note: With both enabled and disabled Eyeglass Configuration Jobs in the Access Zone, the Eyeglass Configuration Replication Readiness validation will only display status for the Enabled jobs.

Zone and IP Pool Readiness - SPN Readiness

Prerequisites: some new validations require 2.5.6 or later for enhanced SPN management and failover.

The SPN Readiness criteria is used to:

1. Detect missing SPNs and insert them into AD based on PowerScale's list of missing SPN's. Requires AD Delegation step to be completed to support auto insert feature.
2. Remediate existing and newly created SmartConnect Zones as short and long SPN's created for each cluster Active Directory machine account.
3. (2.5.6 or later releases) Checks the Case of the SPN in Active Directory versus the Smartconnect zone name case. NOTE: SPN's are case sensitive and must match the case of the cluster SmartConnect name or alias. For example HOST\Data.example.com and HOST\data.example.com are different and must match for correct kerberos authentication, and failover requires the case to match the PowerScale configuration. This validation will detect incorrect case.
4. (2.5.6 or later releases) Checks if syntax is correct in AD (i.e host\xxxx is lower case and not the correct syntax). The service class must be upper case HOST\xxxx (i.e "host\xxxx" is invalid) for failover and authentication.
5. (2.5.6 or later releases) Supports additional service classes for custom SPN insert into ad as well as failover support. The following SPN's are supported NFS, HDFS, WEB and any other custom SPN required for failover and automatic insertion. See Access Zone Failover configuration guide on how to enable custom SPN's.

This check is done for each domain for which each cluster is a member.

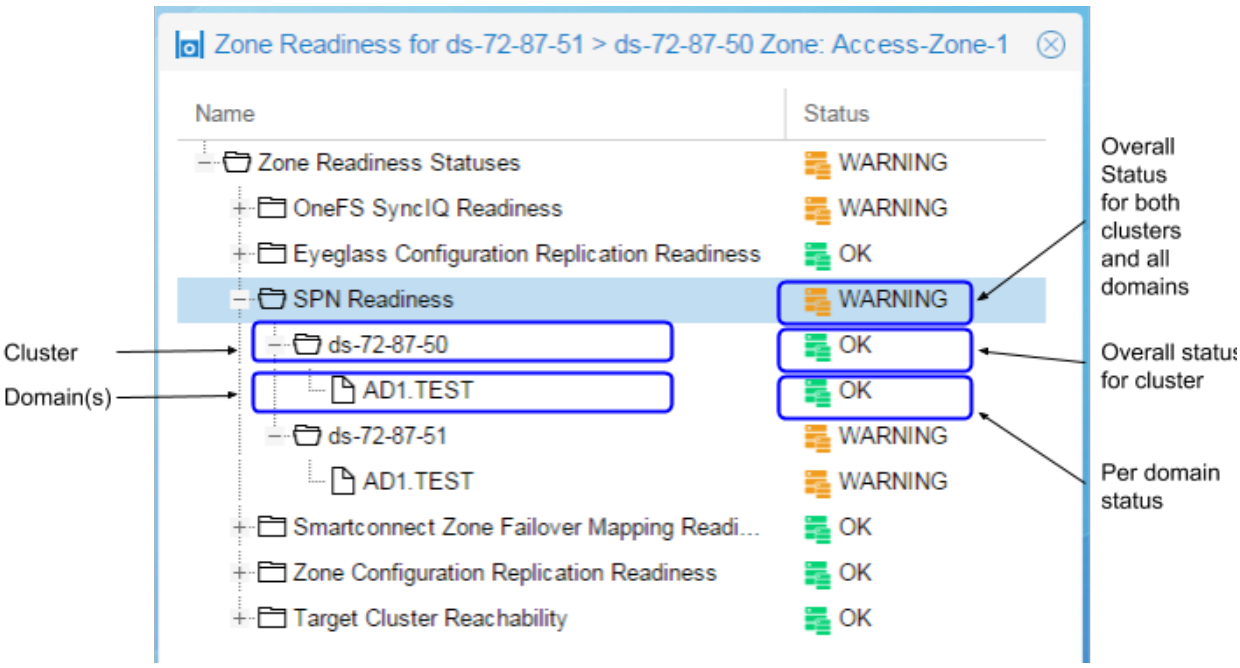

Note: For the case where the PowerScale Cluster is not joined to Active Directory, the SPN Readiness will show the following:

- For OneFS 7.2 the SPN Readiness check is displayed with message “Cannot determine SPNs”.
- For OneFS 8 the SPN Readiness check is not displayed in the Zone Readiness window.

New 2.5.6 or later releases

Each SPN that should be associated with the AD cluster computer object, based on the SmartConnect names or aliases, is displayed and shown as green "OK" if it matches case and SPN syntax. It will show warning for each SPN detected with incorrect case or syntax issue.

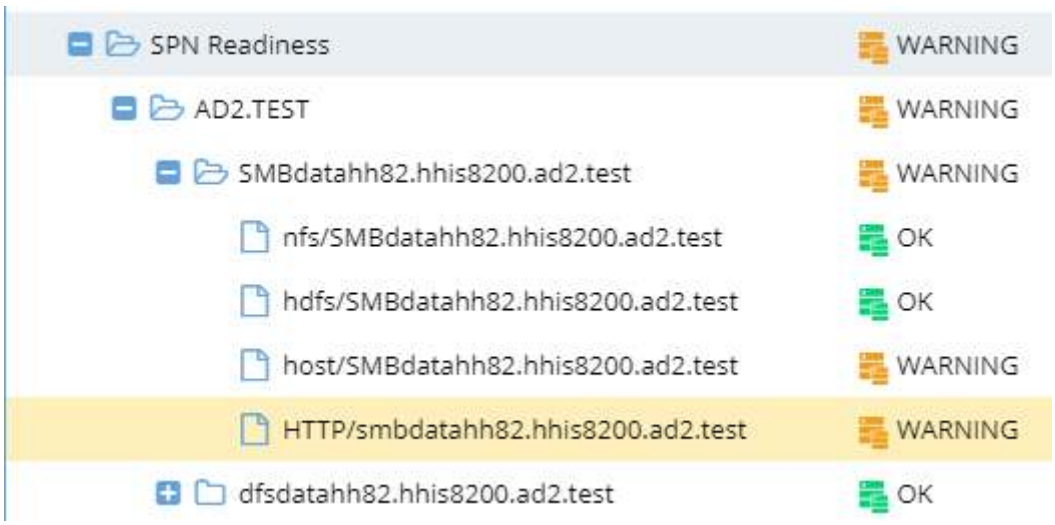

Example error shown when selecting a warning SPN entry will display in the bottom part of the validation UI.

1. **Warning for the lowercase SPN host/xxxxx**
    a. Not valid SPN. Service class in the SPN definition should be HOST
2. **Warning for the wrong case in the smartconnect name or alias**
    a. Not valid SPN. SPN entries are case sensitive and should match the case used on the SmartConnet name.

Zone and IP Pool Readiness - SmartConnect Zone Failover Mapping Readiness

The SmartConnect Zone Failover Mapping Readiness criteria is used to validate that the SmartConnect Zone alias hints have been created

between source and target cluster subnet IP pools. This check is done for each subnet:pool in the Access Zone.

For details on configuring the SmartConnect Zone Failover Mapping Hints, please refer to the [Eyeglass Access Zone Failover Guide](#).

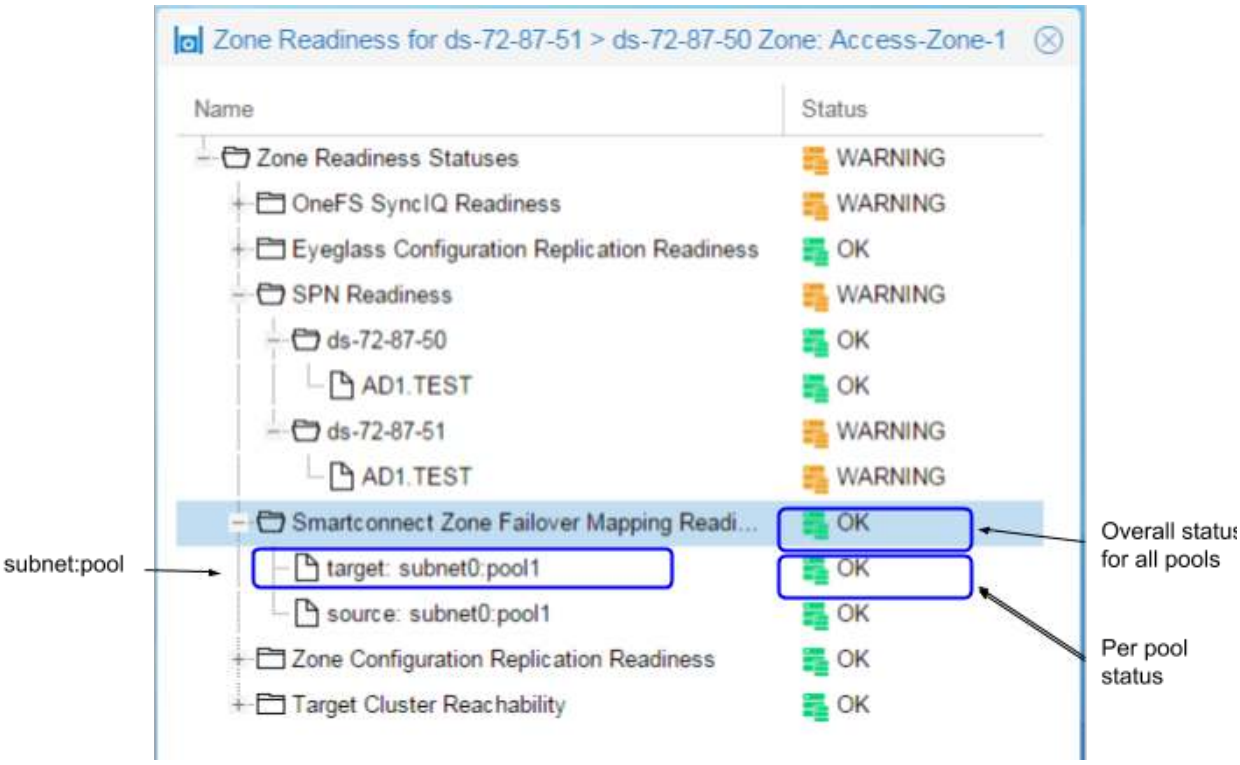

Use the Zone Readiness View Mapping feature to display pools in the Access Zone, and how they have been mapped using the SmartConnect Zone Alias hints.

Zone and IP Pool Readiness - View Mapping

Use the Zone Readiness View Mapping link to display the subnet:pool mappings with configured hints for the Access Zone.

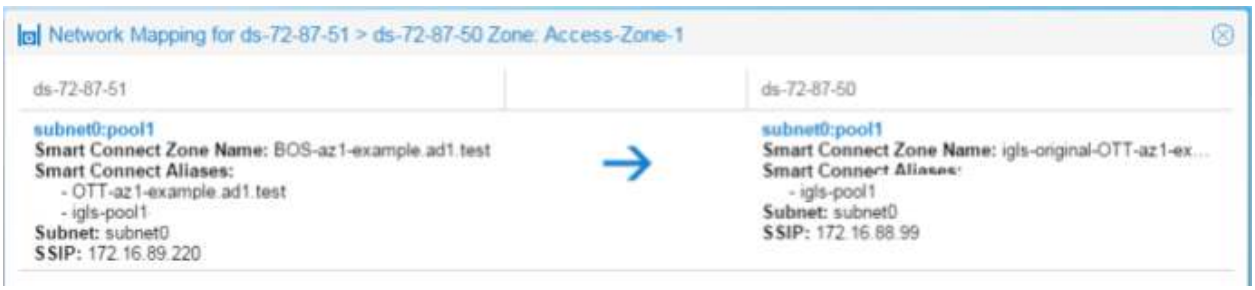

Zone and IP Pool Readiness - Zone Configuration Replication Readiness

The Zone Configuration Replication Readiness criteria is used to validate that the Zone Configuration Replication jobs in the Access Zone have been successfully run to create target cluster Access Zones that don't already exist for configuration sync complete.

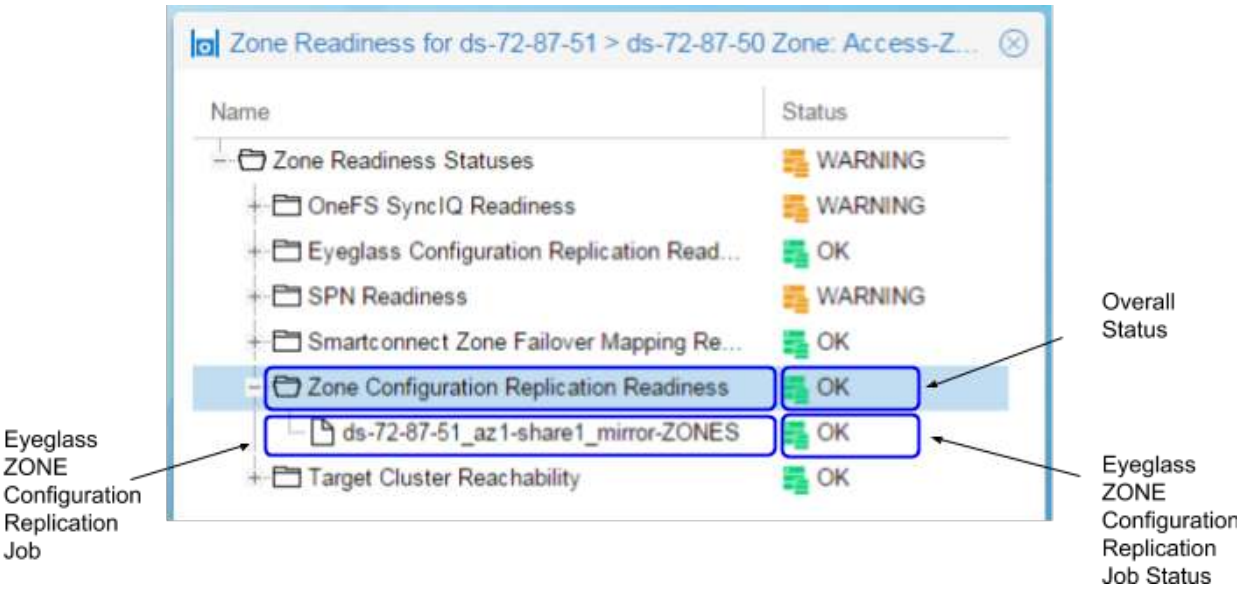

Zone and IP Pool Readiness - Target Cluster Reachability

The Target Cluster Reachability criteria is used to validate that Eyeglass is able to connect to the Failover Target Cluster using Onefs API.

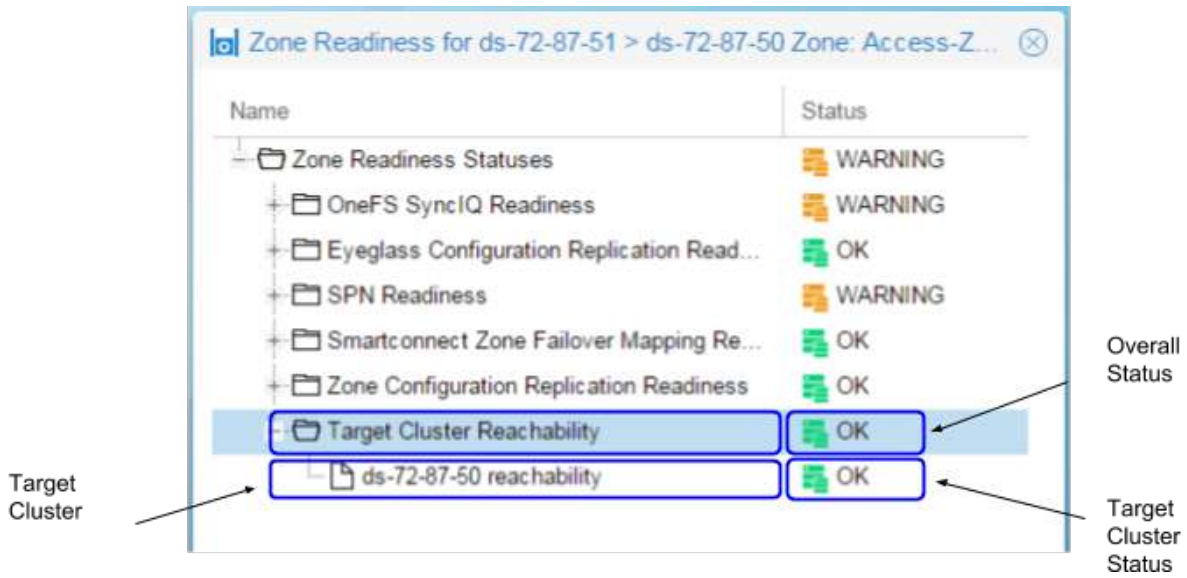

Zone and IP Pool Readiness - Date-Time Validation

The Date-Time Validation is used to validate that the time difference between the cluster nodes, and between clusters and Eyeglass are within an acceptable range that will not affect SyncIQ operations. SyncIQ commands like re-sync prep can fail if the time between cluster nodes is greater than the time between the Eyeglass VM and the cluster with regards to latency of issuing the API call (a scenario where a node returns a timestamp for a step status message is earlier than the beginning of the Re-sync prep request). API calls can return different completion times between clusters. Differences here can cause re-sync prep failover commands to fail, if the difference between Eyeglass and the source cluster is greater than the time it takes for a resync prep command to complete.

NOTE: This condition under which timing differences cause resync failover commands to fail is rare and hard to detect manually. In release 1.8 Eyeglass can detect this condition. It's best practice to use NTP on clusters and the Eyeglass appliance. This allows failover logs and the new feature in release 1.8 or later to collect cluster SyncIQ

reports for each step and append to the failover log. This will make debugging multi step multi cluster failover simpler. This process will require time to be synced.

For each Cluster the following checks are done:

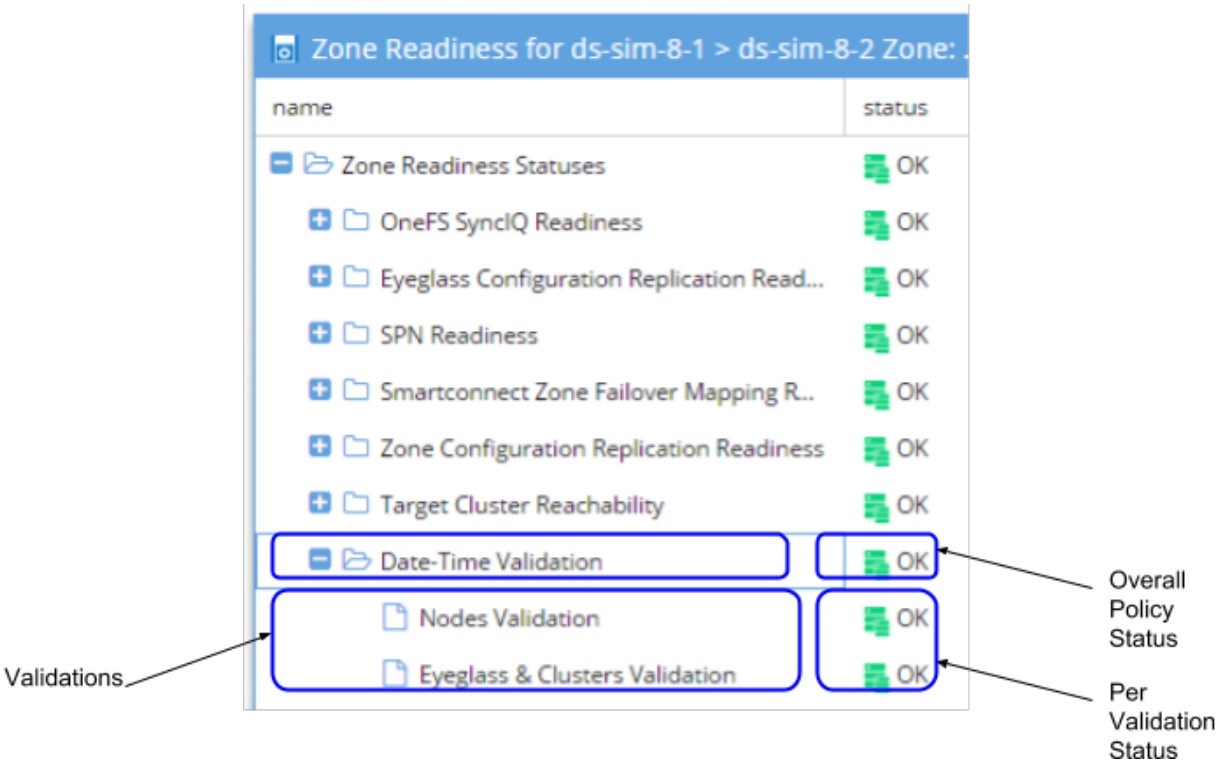

Date-Time Validation	Notes:
Nodes Validation	Validation that the maximum time difference between the nodes of a cluster is less than the time required for the cluster node time request made by Eyeglass to complete.
Eyeglass & Clusters Validation	Validation that the earliest node time for a cluster and the Eyeglass appliance time are less than the time required for the cluster node time request made by Eyeglass to complete plus a default additional skew factor (default 1s). Executed if Nodes Validation is OK.

Zone and IP Pool Path Validation

Zone Path Validation provides the status of Access Zones. Status of **OK** indicates that the Access Zone paths have no conflicts. Status

of **ERROR** indicates that this Access Zone collides with another Access Zone's path.

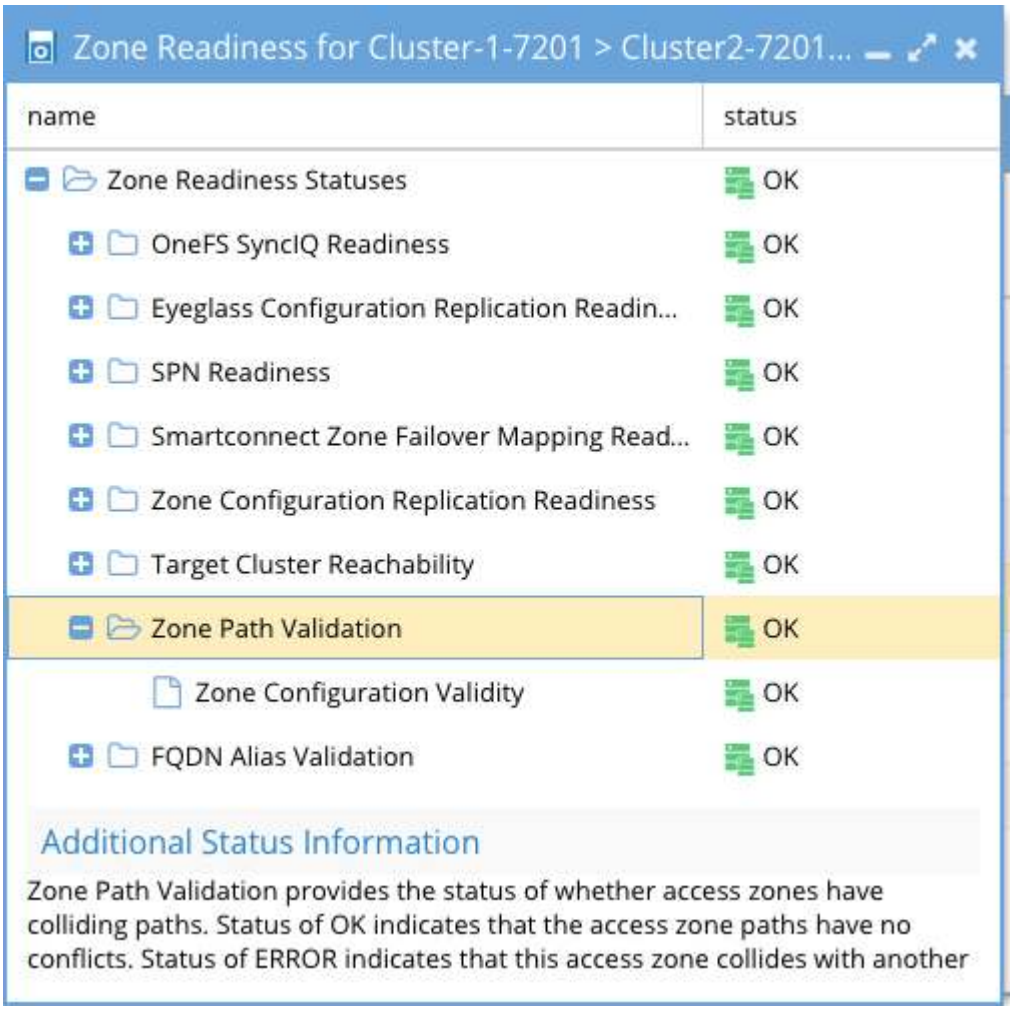

Access Zone and IP Pool FQDN Alias Validation

If cluster was added to Eyeglass with FQDN SmartConnect name for management, this SmartConnect Zone must have an igls-ignore hint applied to avoid a failover impacting Eyeglass access. **Error** indicates that no igls-hint was found on the IP pool for the SmartConnect Zone used for cluster management. **OK** indicates that igls-ignore hint was found.

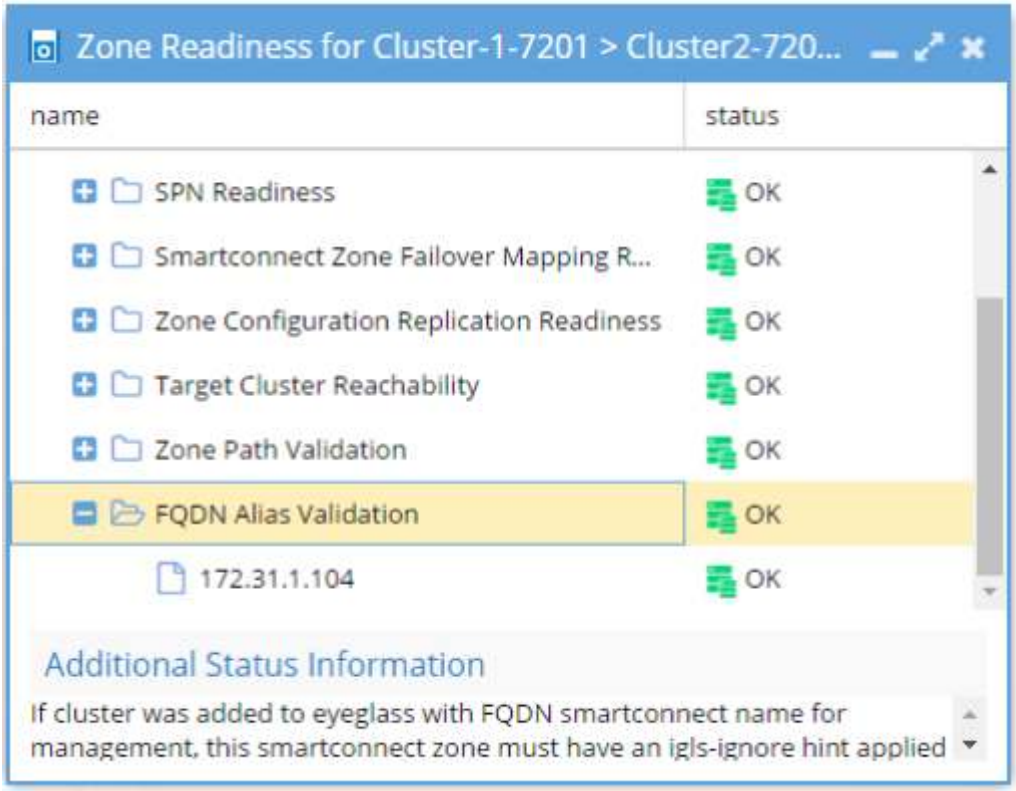

Access Zone and IP Pool - SPN (Service Principal Name) Active Directory Delegation Validation

Prerequisite: 2.5.6 or later

Documentation to Correct warnings

The SPN AD Delegation Validation will test SPN create and delete along with cross delegation between computer objects in AD. This test ensures the AD delegation step is completed. SPN's are used to authenticate SMB clients and SPN's are the same as SmartConnect Zone names and aliases with a host\smartconnect name syntax. These must be present on the cluster AD computer object that is currently the writable cluster to authenticate users. AD domain controllers only allow a computer object to validate passwords for a

service if the SPN list on the AD computer object lists the SPN for the request. Client machines send the SPN with the password for validation. For example "\\data.example.com\shareA" the SPN is "HOST\data.example.com" and must be registered against the cluster computer object in AD in the SPN property list.

This test will create and delete test SPN's against the clusters AD object and the cross delegation to verify AD delegation is configured for failover. The cross SPN delegation allows one cluster to edit the opposite clusters SPN property list using AD permissions. This is required for a real DR event where the cluster that owns the SPN's is not operational, and requires the opposite target cluster to be able to modify and take ownership of the SPN properties. This validation automatically tests all required permissions and will indicate which test failed to allow an administrator to easily identify which Delegation permissions is missing.

See below how to disable this validation.

SPN AD Delegation Example

In the example below you can see the create and delete tests against self, the computer object matching the cluster the test was executed against, and cross test which is the opposite cluster computer object. All tests must pass to be ready for failover. If a test fails you can use the test name to determine which AD permission is missing.

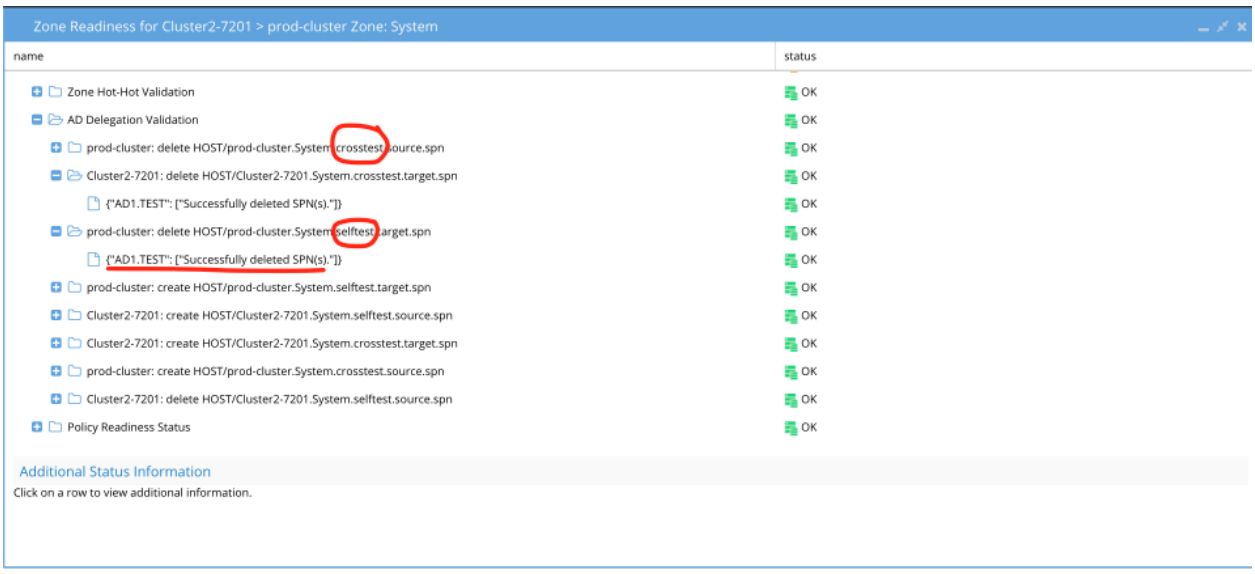

SPN test Failure Conditions and Error Messages

Use this table to identify which test failed and the error message to determine the root cause.

Cause	Impact	Error message	Missing on both
Missing sudoer permission for **isi_classic auth ads*** on one cluster. image: On PowerScale	Alarm(Major), Readiness: AD Delegation ERROR for affected cluster. 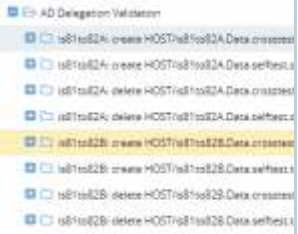	{"AD2.TEST": ["Error processing SPNS: Sorry, user eyeglass is not allowed to execute '/usr/bin/isi_classic auth ads spn create -s HOST/is81to82A.Data.crosstest.target.spn --machinecreds --account IS81TO82B$ -D AD2.TEST' as root on is81to82A-1."]}	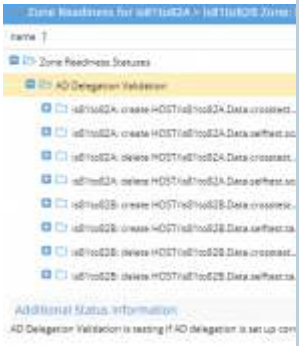
Missing AD Delegation for **SELF** test for one cluster. ☑ Read servicePrincip ☑ Write servicePrincip image: On AD	Alarm(Major) SPN processing (either checking or re Readiness: AD Delegation ERROR for self tests of the affected cluster.	Create: {"AD2.TEST": ["Error processing SPNS: LdapError: Failed to modify attribute 'servicePrincipalName' [50:Insufficient access]"]} Delete:{"AD2.TEST": []}	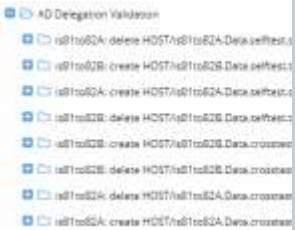

	AD Delegation Validation		
Missing AD Delegation for **CROSS** test for one cluster. ☑ Read servicePrincip ☑ Write servicePrincip image: On AD Note: Cross Test warning for clusterA check the AD cross delegation for clusterB whether clusterB has configured cross delegation permission for clusterA	Alarm(Major) SPN processing (either checking or re Readiness: AD Delegation ERROR for cross tests of the affected cluster. AD Delegation Validation	Create: {"AD2.TEST": ["Error processing SPNS: LdapError: Failed to modify attribute 'servicePrincipalName' [50:Insufficient access]"]} Delete: {"AD2.TEST": []}	AD Delegation Validation

Access Zone and IP Pool - DNS Dual Delegation Validation

Prerequisite: 2.5.6 or later

Documentation to correct Warnings

UnSupported Configurations

1. If Infoblox is configured using forwarders and not dual name servers this validation will not work. Forwarders is vendor specific configuration and not standards based DNS. Nslookup cannot remotely validate dual forwarding configuration. Recommendation is to use name servers with Infoblox versus dual forwards. This validation will need to be

disabled if Infoblox dual forwarding is configured. Open a case with support.

2. NOTE: Eyeglass will validate the groupNet DNS servers directly and will not use the OS DNS configured on Eyeglass, this is because the DNS servers that must be configured are the ones used by PowerScale itself. This requires the Eyeglass VM to have port 53 udp access to the groupnet DNS servers. If this is not possible, the validation must be disabled. See the section be

This validation will automatically validate each SmartConnect name and alias on pools has two name servers configured, and that the ip address returned is a subnet service ip that is servicing the IP pool by a subnet with the correct SSIP. If any of these tests fail it means failover will not be able to auto the DNS step leaving it a manual step. This will detect misconfiguration or missing dual DNS delegation before planned failovers. If A records are used in the delegation, the DNS name returned as the name server will have a reverse lookup done to validate the IP is a subnet service ip. This validation will also make sure the cluster pairs configured for failover have the correct SSIP on both clusters and can be found in inventory for the 2 clusters.

Example Dual DNS validation

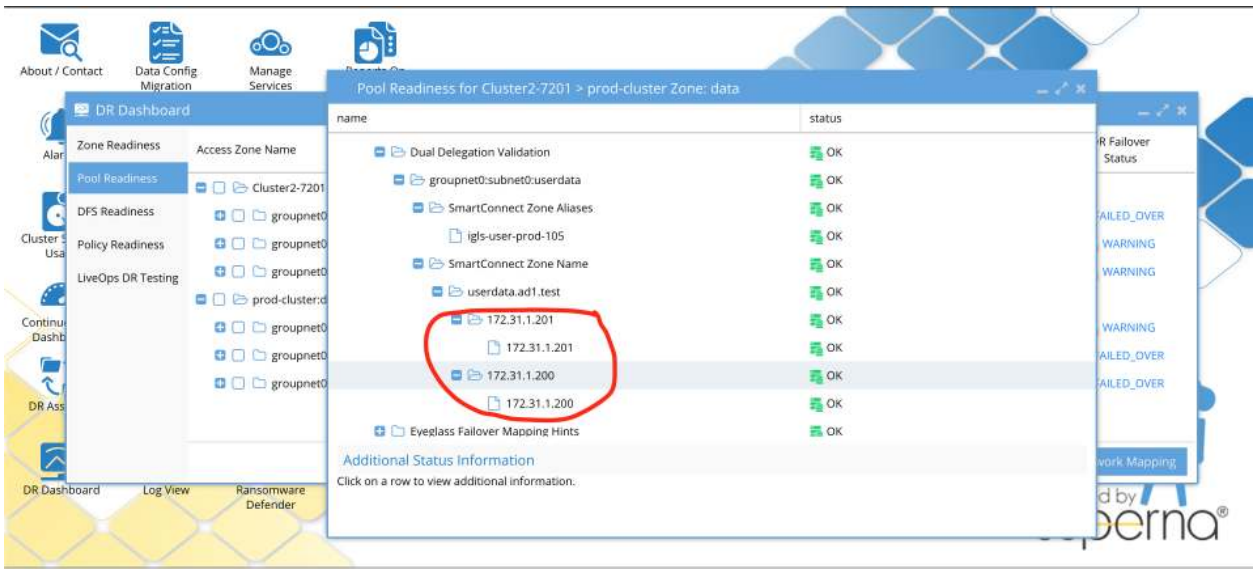

Warnings and Statuses for all Possible Dual Delegation Validations

Dual delegation validation cases and corresponding status information for SmartConnect zone name and SmartConnect zone alias:

Pool with ignore alias igls-ignore-

In Additional status Information shows: 'Pool has igls-ignore hint. Dual Delegation validation skipped.'

Zone or pool with correct settings:

In Additional status Information shows: 'IP address detected for cluster XXXX.'

SmartConnect Zone name/alias unknown:

WARNING

In Additional status Information shows: 'Could not resolve SmartConnect Zone name to a valid IP address.'

One NS record for SmartConnect Zone name/alias:

WARNING
In Additional status Information shows: 'DNS query did return only one IP address for this SmartConnect Zone name. There should be two IP addresses for the valid dual delegation setting.'

Detected three NS records for SmartConnect Zone name/alias:

WARNING
In Additional status Information shows: 'Dual delegation for this zone is not set up correctly. SmartConnect Zone name revolves to more than 2 IP addresses.'

One of the NS record is incorrect for SmartConnect Zone name/alias:

WARNING
In Additional status Information shows: 'The IP address does not reference valid cluster.'

There is no NS records for SmartConnect Zone name/alias:

WARNING
In Additional status Information shows: 'Could not resolve SmartConnect Zone name to a valid IP address.'

Two NS records point to same IP addresses for SmartConnect Zone name/alias:

WARNING
In Additional status Information shows: 'SmartConnect name server delegations are not dual delegated. Both name resolve to the same IP address.'

SSIP of source cluster (in cluster) is incorrect:

WARNING
In Additional status Information shows: 'The IP address does not reference valid cluster.'

All Failover type Domain Mark Validation

Prerequisite: 2.5.6 or later

Documentation to correct Warnings

1. **Consult Dell Documentation to run domain mark job**

This will validate that source and target cluster have a valid domain mark for accelerated failback. This will raise a warning if either cluster source or target is missing a domain mark. This is an important validation to ensure failovers do not take a long time to wait for domain mark to run during resync prep step.

The domain mark needs to be run manually on the cluster with Onefs job UI if the validation fails. This should always be completed prior to any planned failover. All policies are checked for this validation.

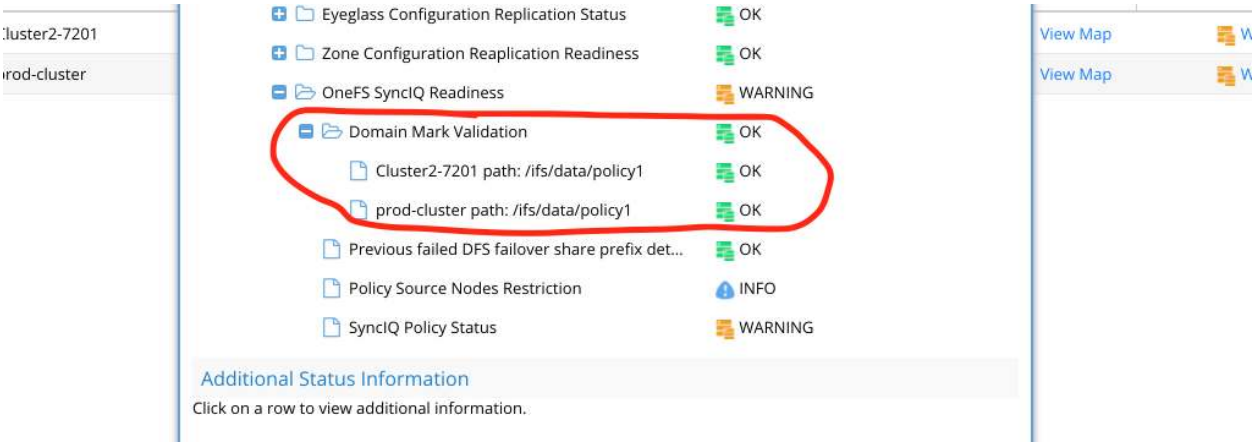

IP Pool Failover Readiness

This interface shows each Access zone and the IP pool defined within the access zone. Expanding each pool will show the SyncIQ policies that are mapped to the pool.

1. The Access zone column shows cluster:zone name.
2. Pool mapping will show the pool to pool igls-hints that map a pool on source cluster to target and allows viewing the mapping
3. Target cluster this pool will failover too
4. Last Successful Readiness Check is the day and time that Failover readiness Assessed this pool readiness
5. Map policy to pool allows mapping a policy or more than one policy to a pool and allows viewing the mapping for all pools in the access zone.
6. DR Failover Status shows the highest severity state for all validations OR it will show failed over status if the pool has been failed over.

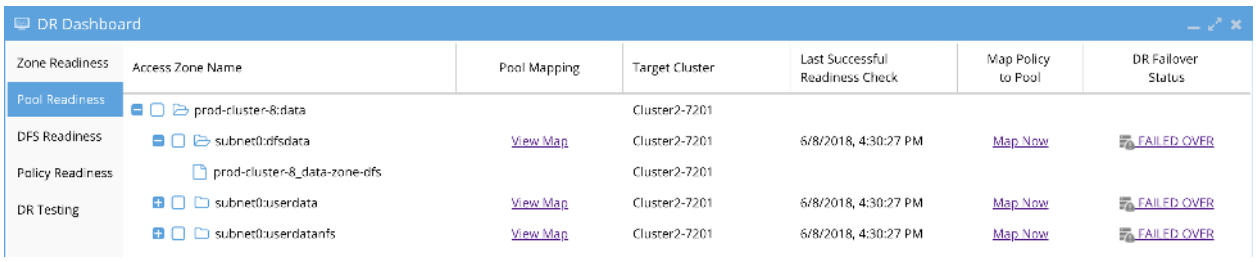

Pool Validations

The pool validations are the same as Access Zone readiness with the key difference being they only apply to the pool itself and not the whole zone. This allows a single pool to be viewed and ready for failover independently.

Pool Readiness Validation

The Pool Readiness validation that is unique to IP pool failover is the un-mapped policy smartconnect/ip pool status. and the overall pool readiness that summarizes the pools status.

Un-mapped policy validation Overview

1. **This verifies that all synciq policies in the zone have been mapped to a pool.**
2. **A pool may have more than one SyncIQ policy mapped.**
3. **A syncIQ policy may NOT be mapped to more than one pool.**

4. **Any Synclq policy not Mapped using the Dr Dashboard IP pool mapping interface, will raise this error message and WILL block failover for all pools in the access zone until corrected.**

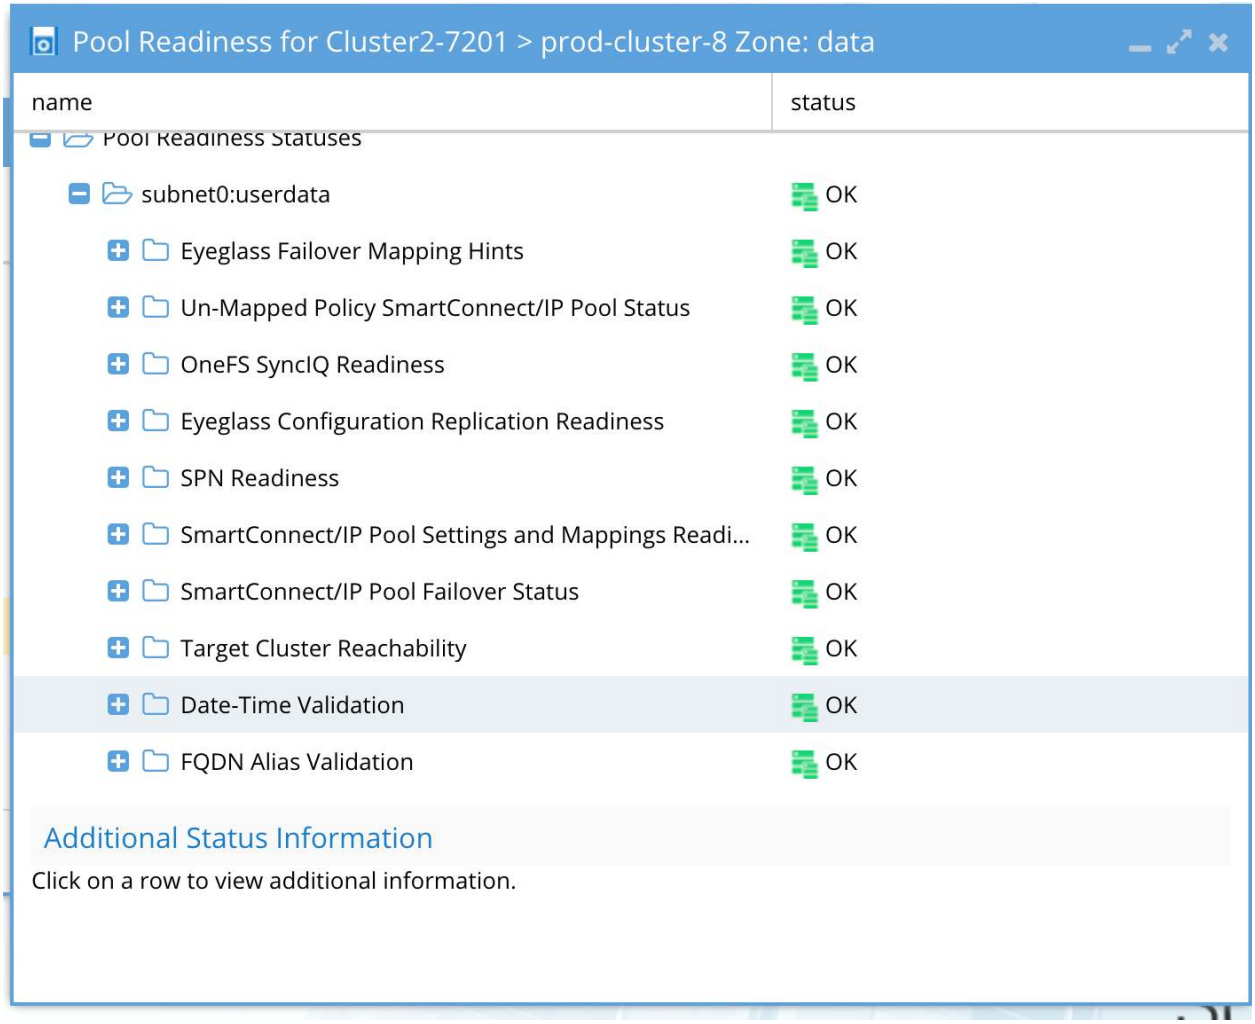

Overall pool validation status

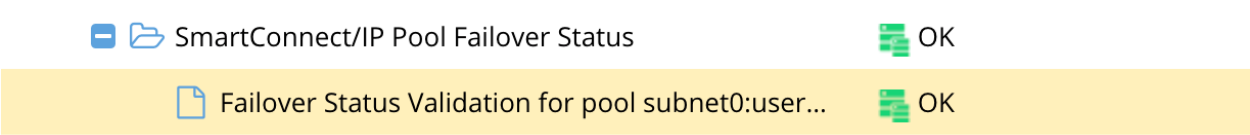

How To Configure Advanced DNS Delegation Modes Required for Certain Environments

What's new

1. **2.5.7 adds the ability to detect DNS dual delegations above the smartconnect zone name level in the DNS name space. example smartconnect data.example.com is the smartconnect name but the delegation is done at the example.com level in DNS. Now the validation will attempt to locate the delegation above the smartconnect level to locate the dual delegation. If this validation scans all DNS names and finds no dual name server entries it will result in a failed validation.**
    a. **additionnal logging shows all steps in the validation located here /opt/superna/sca/logs/readiness.log**

These settings allows control over DNS query servers and recursion options needed for some environments.
When to use the options:

1. If Eyeglass has no access to reach the groupnet DNS due to firewall then enable local OS DNS option below
2. If your DNS is Bluecat, bind or Infoblox we recommend disabling recursive lookups and use the option below to disable recursive lookups.
3. Combine both values if both scenarios apply

How to change the values

1. ssh to eyeglass as admin
2. sudo -s (enter admin password)
3. nano /opt/superna/sca/data/system.xml
4. Find the readinessvalidation tag and add the tag below with the settings below to change the setting required for your situation, only add the tag if you want to change to the settings below.
5. control+x to save and exit
6. systemctl restart sca (for the changes to take effect)
7. done

New system.xml parameters in the <readinessvalidation> section with the following defaults: **NOTE: The typo in the tag is needed in 2.5.6 release.**

<dualdelegation_use_eyglass_dns>true</dualdelegation_use_eyglass_dns>

<dualdelegation_recurse>**false**</dualdelegation_recurse>

How the tags work

Both are independent settings, so with these there are 4 possible states. When doing dual dns delegation validation, implement the following behavior based on the settings below:

1. if dualdelegation_use_eyglass_dns is true, instead of taking the DNS server from the isilon's groupnet, get the DNS server from

the local OS, and issue the dual dns delegation validation request(s) to that DNS server instead.

   a. If it's false, continue to issue requests to the isilon's dns server on the groupnet **(default mode)**.

2. if dualdelegation_recurse is false, turn off recursion on the DNS query.

   a. **if true DNS query will use recursive lookups and this is the default mode.**

How to Configure Advanced SPN delay mode for Active Directory Delegation

1. This advanced mode can be used to add a delay between the SPN create and delete test used during the validation. If AD domain controllers do not execute the create and delete fast enough this can fail the valiation test. This will add a delay in seconds between the commands. Default is no delay.
2. ssh to eyeglass as admin
3. sudo -s (enter admin password)
4. nano /opt/superna/sca/data/system.xml
5. Insert the readinessvalidation tags add the value below and change the 0 to 5 seconds.
6. Then control + X to save and exit the file.
7. systemctl restart sca (for the changes to take effect)
8. <readinessvalidation> <spnCreateDeleteWait>5</spnCreateDeleteWait> </readinessvalidation>

How to Disable AD or DNS Delegation Validations - Advanced Option

In some cases it may be required to disable these validations in some environments. If Infoblox is used with forwarding the DNS dual delegation will need to be disabled. In other cases, SPN's are not required and AD delegation is not completed. This option is global and will disable these validations from executing on all access zone's and pools.

1. **Login to Eyeglass as admin using ssh**
2. Disable DNS validation
    a. igls adv readinessvalidation set --dualdelegation=false
3. Disable SPN test Validation
    a. igls adv readinessvalidation set --spnsdelegation=false

Network Visualization

A new way to view PowerScale clusters, DR status, and jump to the DR Dashboard. The Network Visualization feature allows you to visualize DR and cluster replication. This first release offers the first in several enhancements aimed at visualizing data, data flows, and storage across one or more PowerScale clusters.

This view indicates which clusters are replicating to each other and direction. For each cluster, failover readiness status for all failover types is summarized. Any failover readiness error will show as a red arrow from the source to target cluster (failover direction). Warnings are displayed with an orange arrow. In the case where there are not errors or warnings the arrow will be green for active replication direction (failover direction) and Grey for the failed over (inactive) direction. This simplifies monitoring many clusters replicating.

To view Network Visualization:

1. Open the Icon.

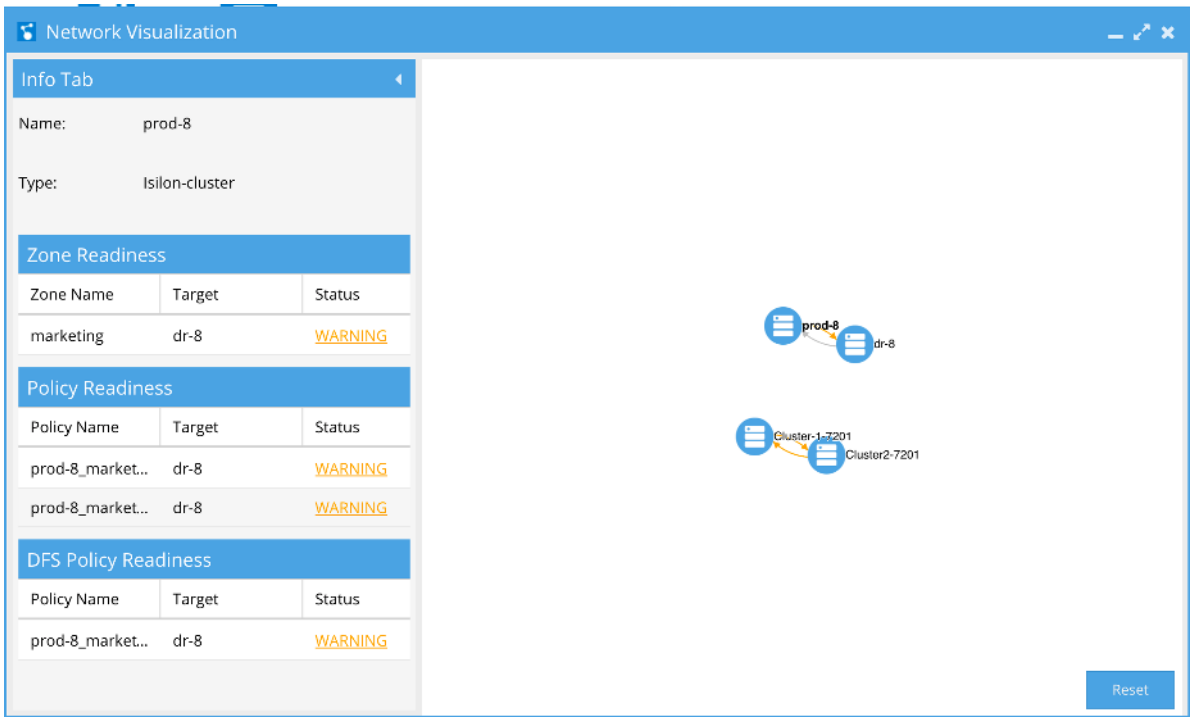

2. Zoom in or out to navigate depth of view.
3. Click and hold to drag the visual view objects.
4. Click a cluster to get view of active Sync Data on the cluster viewed by Failover mode and status.

5. Click on hyperlink to hump to DR Dashboard Directly from the Network Visualization window.

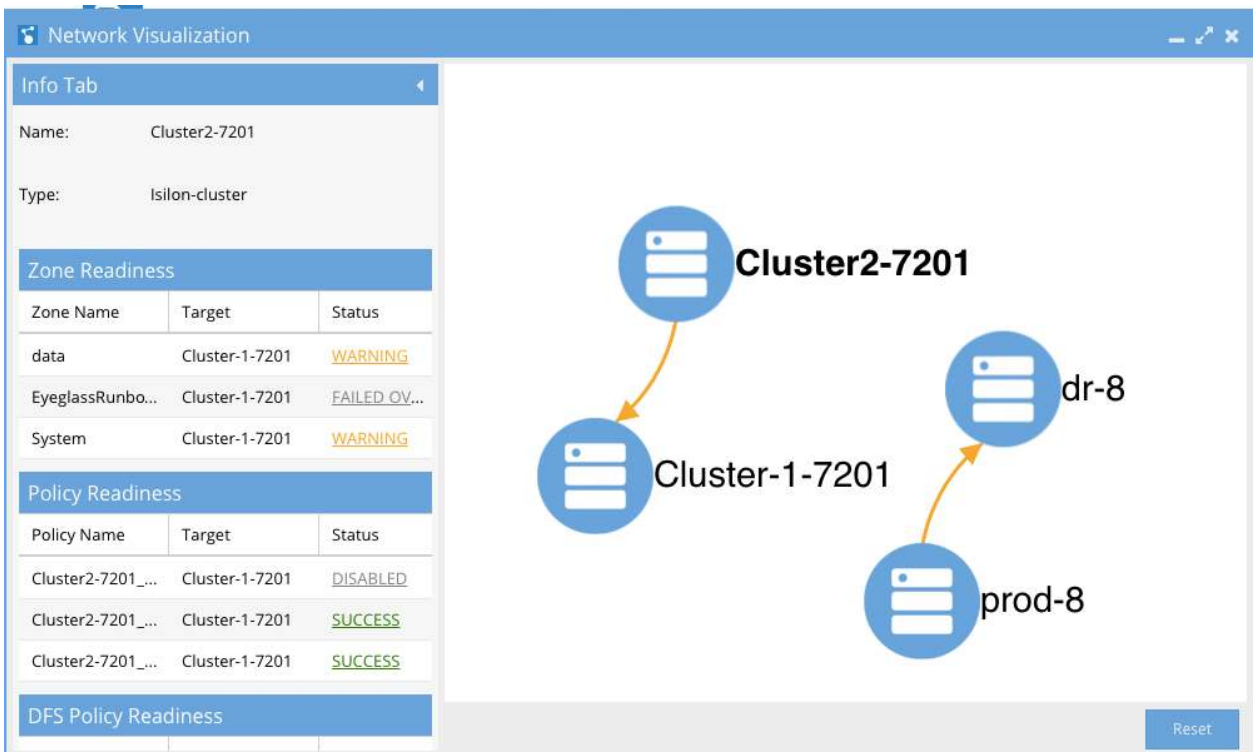

© Superna LLC

1.12. RPO Reporting and Trending Feature Guide

Home Top

RPO Reporting and Trending Feature Guide

See Eyeglass PowerScale Edition Recovery Point Objective Trending and Reporting

End User Offline Detection

The Eyeglass web page detects when the browser does not have Internet access and does not attempt to connect to Twitter or the Superna support (superna.help) site. In this case the Twitter icon on the Eyeglass web page is removed and the Help button is inactive.

Eyeglass Reporting

Eyeglass offers these reports:

1. Cluster reports that cover configuration>
2. RPO reports that include SyncIQ performance, and backup details on SyncIQ performance, for backup monitoring of policies used for backup only versus DR.

See RPO guide for details .

Cluster Usage reporting (requires Cluster Storage Monitor License keys).

See Cluster Storage Monitor admin guide for details .

Reports are now available from the Reports on Demand Icon.

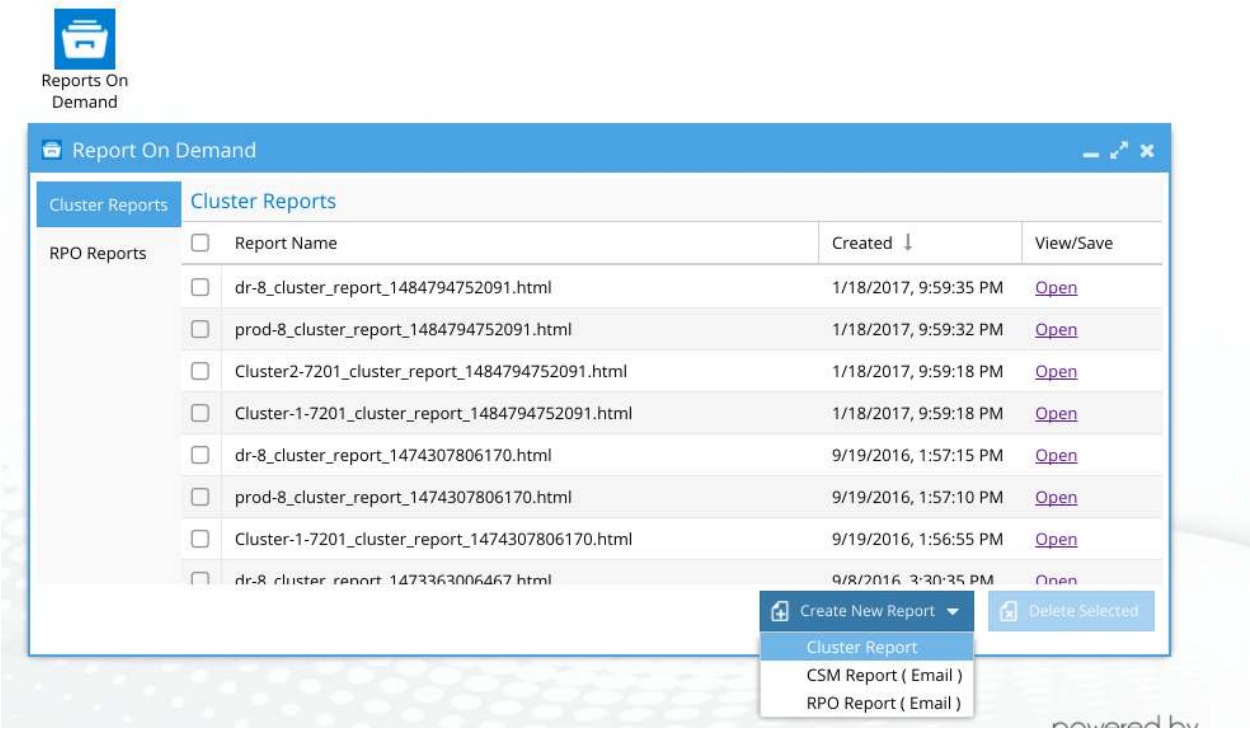

Eyeglass Cluster Config Reports

Eyeglass collects key cluster settings for the PowerScale clusters that it is managing, and makes them available in an HTML report format by email or on demand. Use the contents of this report to automate, and keep up to date, the documentation of cluster configuration for DR and your DR run book. This information is also valuable for IT planning, cluster expansion, networking and security

What's New in Cluster Reports

New in 1.9 cluster reports includes:

1. DNS settings for OneFS 8 now added to cluster reports.
2. Pools/Subnets in all groupnets shows up (with Access Zones) in OneFS 8 cluster reports. Also added zone name to pools in OneFS 7 reports.

Cluster Configuration Email Report

By default, the full cluster HTML report containing key cluster settings is emailed daily, at midnight, to all users who have been configured for email notification.

Cluster Configuration Report On-Demand

To generate a Cluster Configuration Report on-demand:

1. Login to the Eyeglass web page.
2. Open the Reports on Demand window.

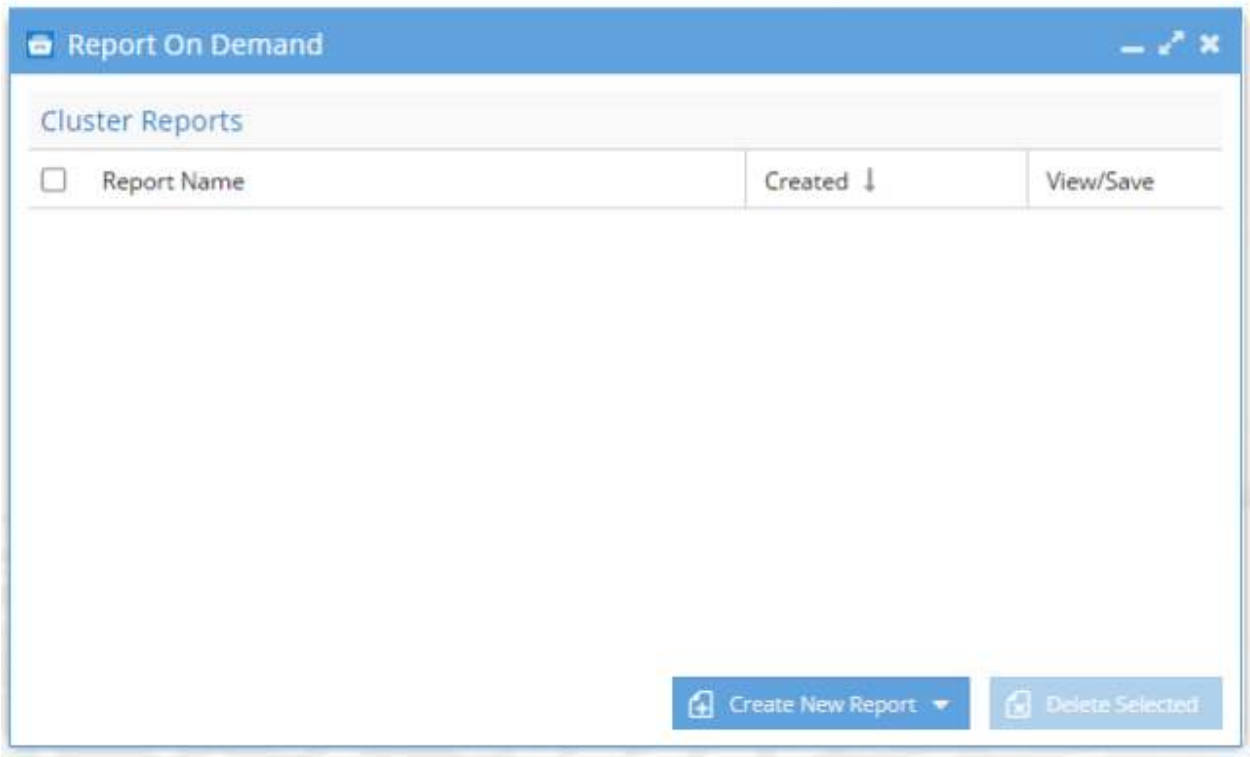

3. Select Create New Report/Cluster Report.

4. A confirmation dialog is displayed once the request has been submitted.

5. To monitor the progress of the report generation from the Running Jobs window, you can select View Running Jobs in the confirmation dialog.

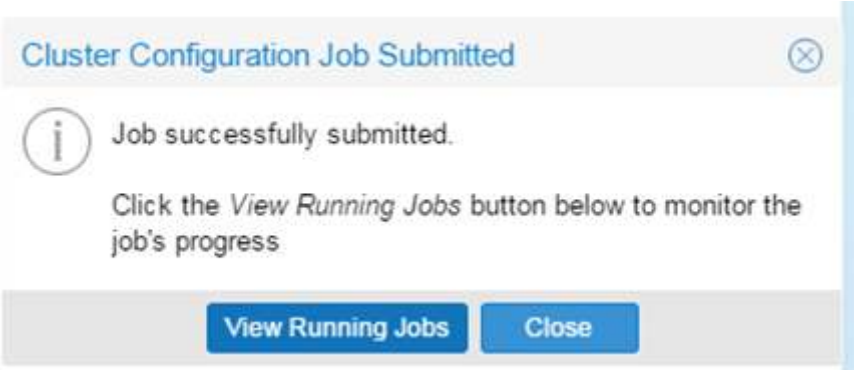

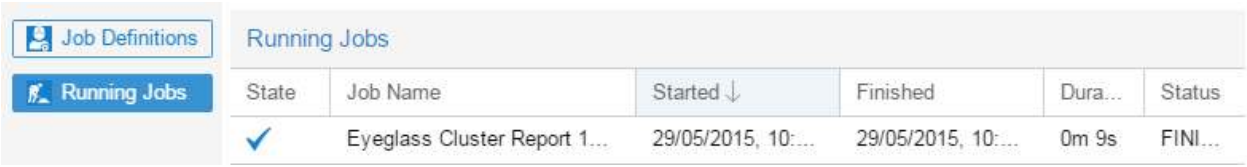

6. When the report has been generated it will appear in the Cluster Reports list. Report Name format includes the name of the cluster that the report is about <cluster name>_cluster_report_<timestamp>.html.

7. Select the Open link to view the report.

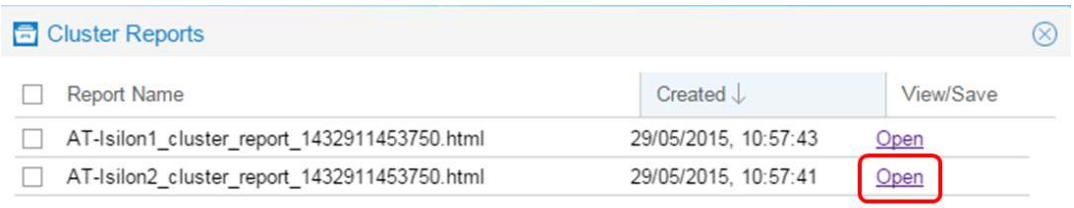

8. The report opens in a new Cluster Report Viewer window with a table of contents linked to each section of the report.

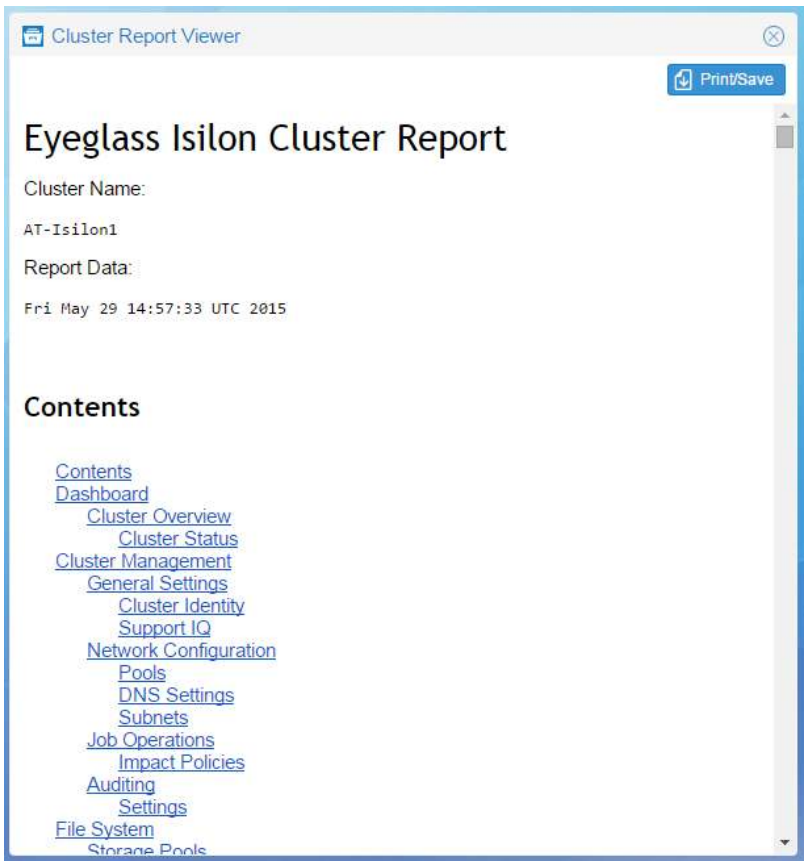

The report also displays changes in configuration - items added in green and items removed in red. Example below is adding a share permission:

```
name: policy1-group1
ntfs_acl_support: true
oplocks: true
path: /ifs/data/policy1/group1
permissions:

    permission: full
    permission_type: allow
    trustee:
      id: SID:S-1-5-21-1825440792-1775492485-428706412-1807
      name: AD01\\dorothyuser
      type: use
```

#added-diff-389

```
    permission: full
    permission_type: allow
    trustee:
      id: SID:S-1-1-0
      name: Everyone
      type: wellknown
```

9. To save the report as PDF, select the Print/Save button in the Cluster Report Viewer window and then select the Save as PDF option.

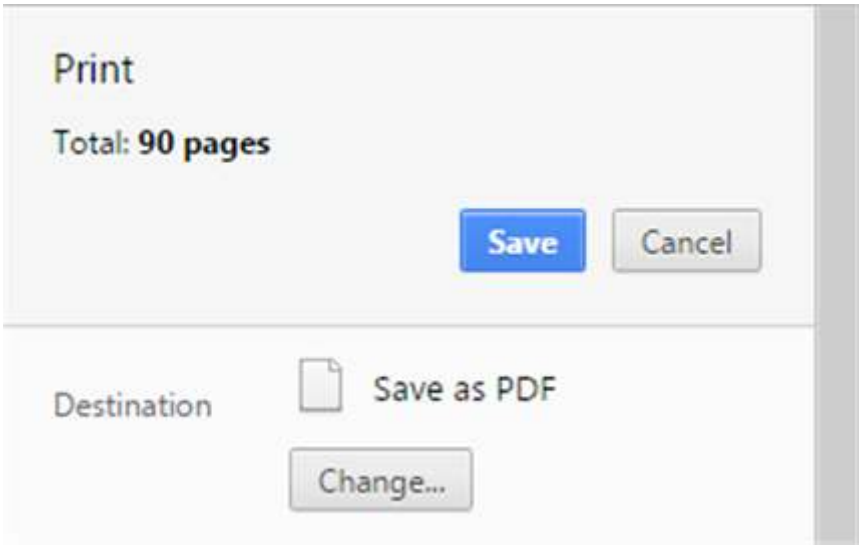

To remove old reports:

1. Use the check box to select reports for deletion.

2. Click Delete Selected.

Cluster Networking for Failover and Cluster Reports

Cluster networking is now added to cluster reports and can also be viewed in the zone inventory tree (see examples below). This information is key to assisting with Access Zone failover. Each Access Zone maintains a mapping to subnet pools and configuration data.

See the Failover Design Guide (Eyeglass Failover Design Guide) for instructions on creating hints to allow Eyeglass to failover networking with SmartConnect Zone migration, SPN updates, and pool mapping. This information is also needed to assist with Runbook Robot configuration within an Access Zone, and will allow customers to verify complete failover and failback operations between sites using Eyeglass automation.

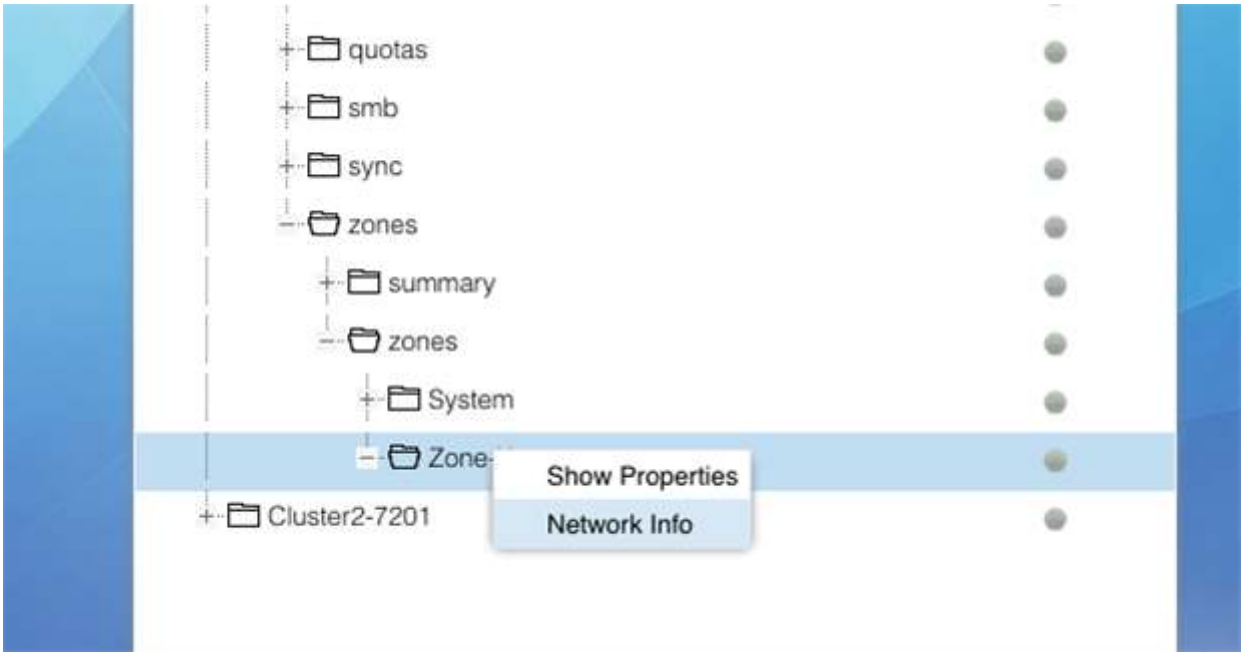

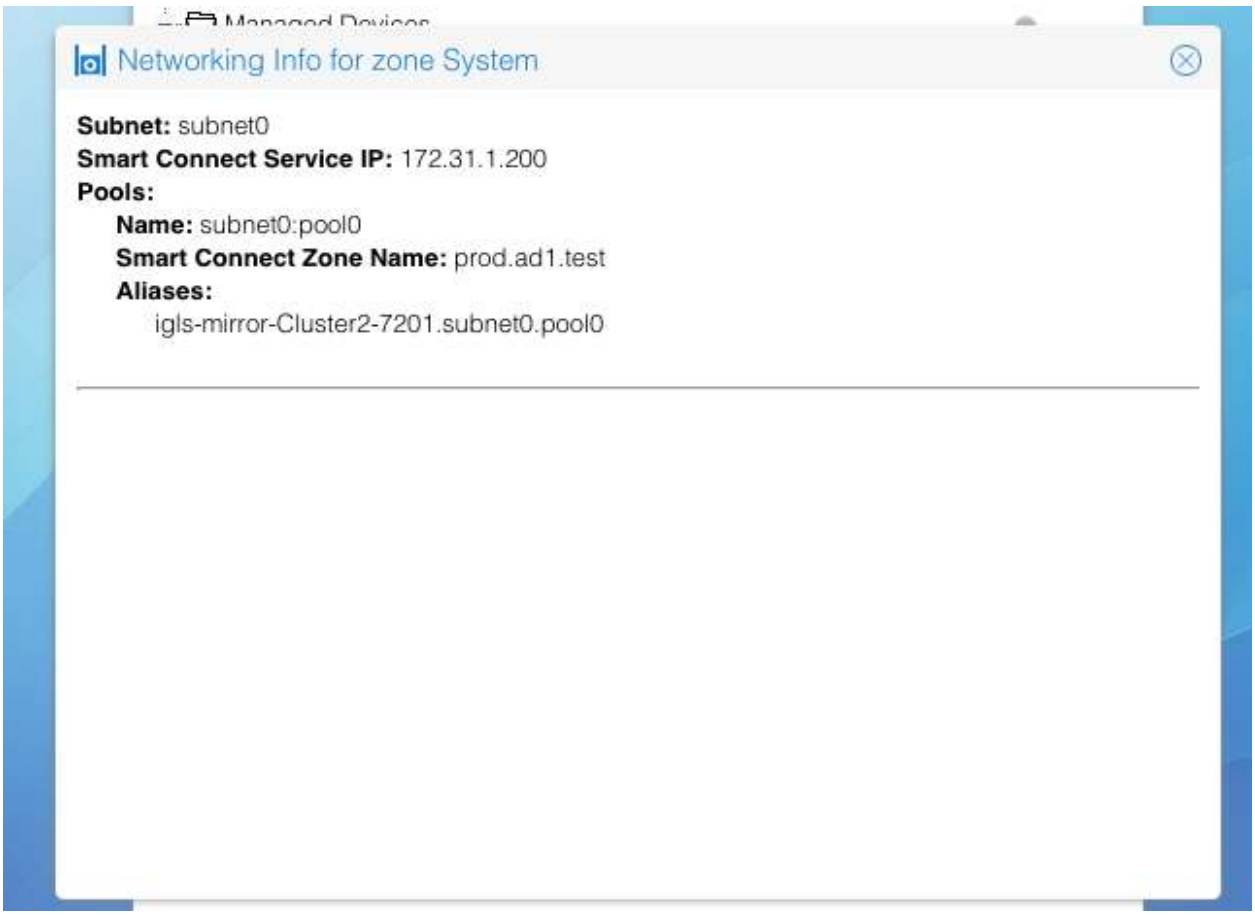
Networking Info for zone System
Subnet: subnet0
Smart Connect Service IP: 172.31.1.200
Pools:
Name: subnet0:pool0
Smart Connect Zone Name: prod.ad1.test
Aliases:
igls-mirror-Cluster2-7201.subnet0.pool0

© Superna LLC

1.13. Role Based Access Controls RBAC

Home Top

Role Based Access Controls RBAC

This new feature is found under the User Roles icon on the desktop. It allows user or groups to be mapped to Eyeglass roles. This feature can be used to create various user roles within Eyeglass from read-only access, to backup, to DR failover roles.

The complete guide on how to configure can be found in Role Based Access Control And Authentication.

© Superna LLC

1.14. Configure Email, Twitter, Slack, Webhooks for Notifications of Eyeglass Monitoring Events

Home Top

- Notification Options
- Limitations of Slack, Webook, Twitter Notifications
- How to configure Slack channel PowerScale events AND Eyeglass events
- Webhooks
- How to Configure Private Tweets of Alarms

Notification Options

For advanced alarm filtering and forwarding over SNMP or Syslog see the Eyeglass Alarm forwarding guide.

For advanced alarm custom routing configurations see this guide.

The Steps to configure Email Notification are found in the Eyeglass PowerScale Edition Quick Start Guide for Eyeglass Installation.

Additional information for Microsoft Exchange can be found in the tech note How to Setup Email alarms with Exchange.

Limitations of Slack, Webook, Twitter Notifications

1. All of these options require transparent NAT to reach the Internet services. No support for OS proxy is available for this notification options.

How to configure Slack channel PowerScale events AND Eyeglass events

1. Goto https://superna.slack.com/apps/manage/custom-integrations (for your domain) as admin.
2. Select Incoming WebHooks.
3. Configure as per screen shot pick the channel to receive alarms.
4. Cut and paste Webhook url for configuration in Eyeglass.

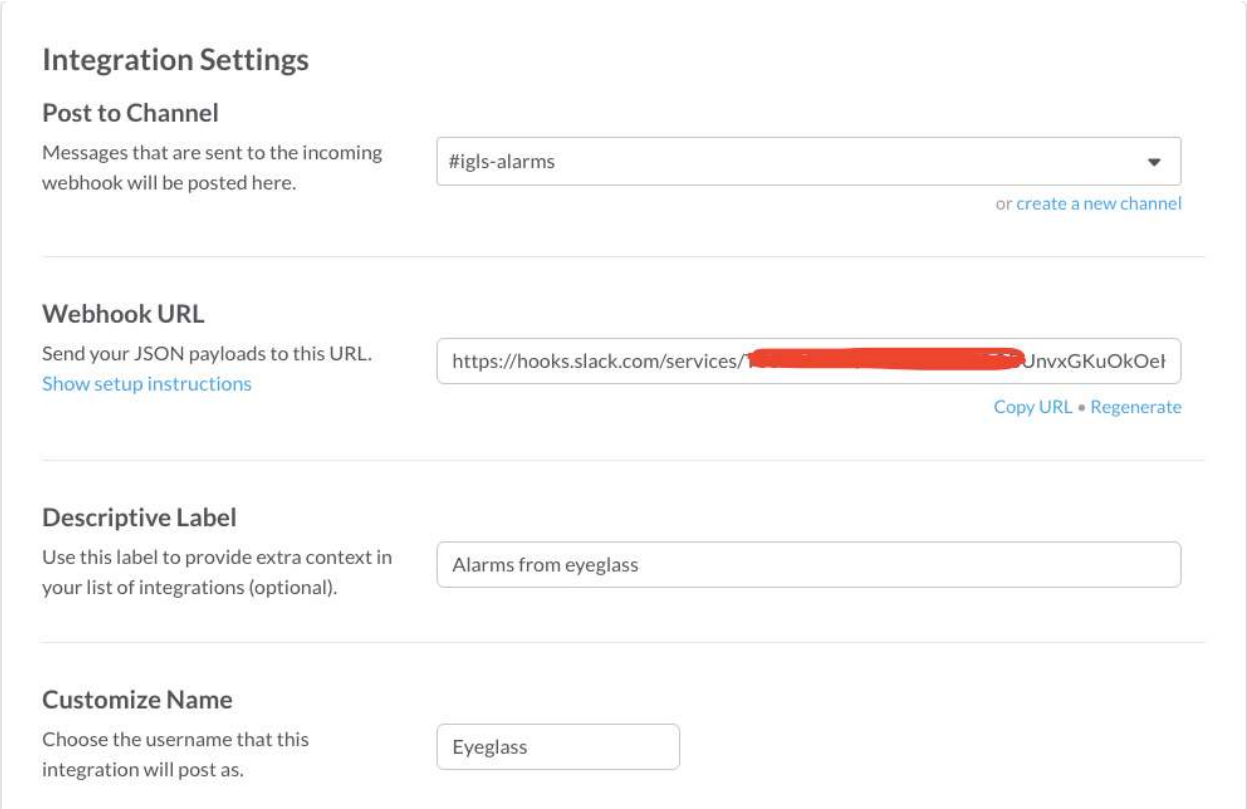

5. Login to Eyeglass.

6. Open notification center.
7. Select Slack tab.
8. Set the Webhook url created above.
9. Set the level alarms that will be sent to the channel.

Note: This includes DR events from Eyeglass and PowerScale events that are polled from the cluster.

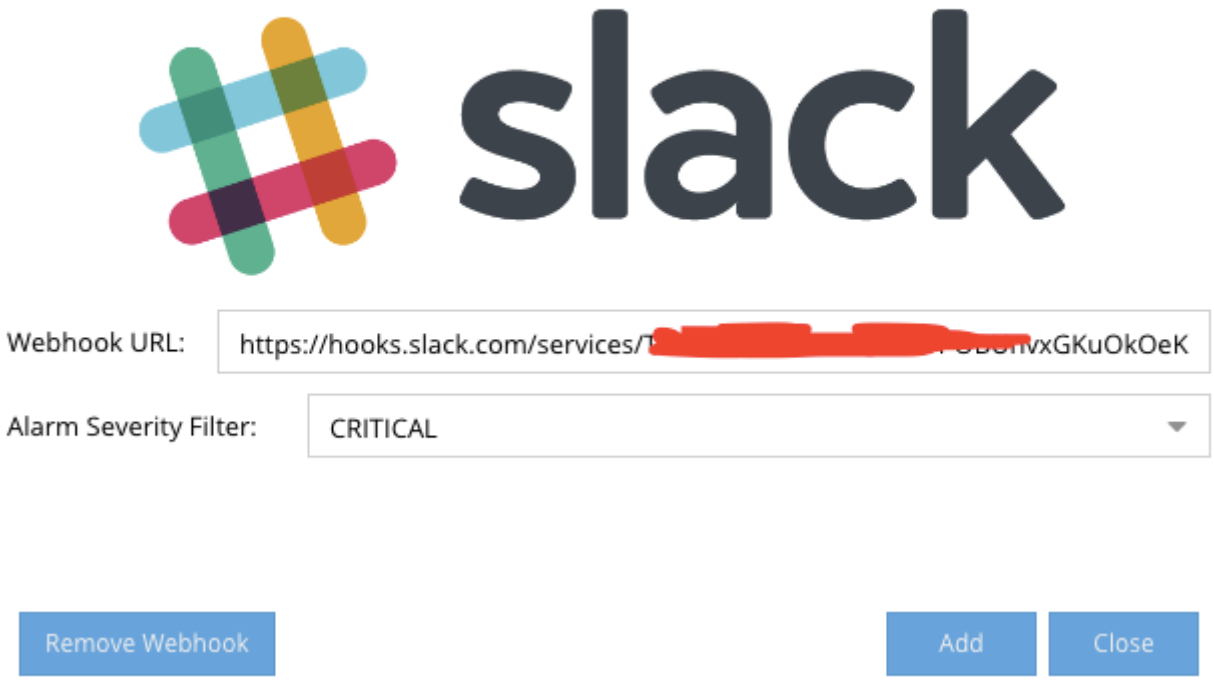

Webhooks

Webhooks uses http post commands to send events to any system that supports this interface.

1. Open notification center.
2. Click Webhooks tab.

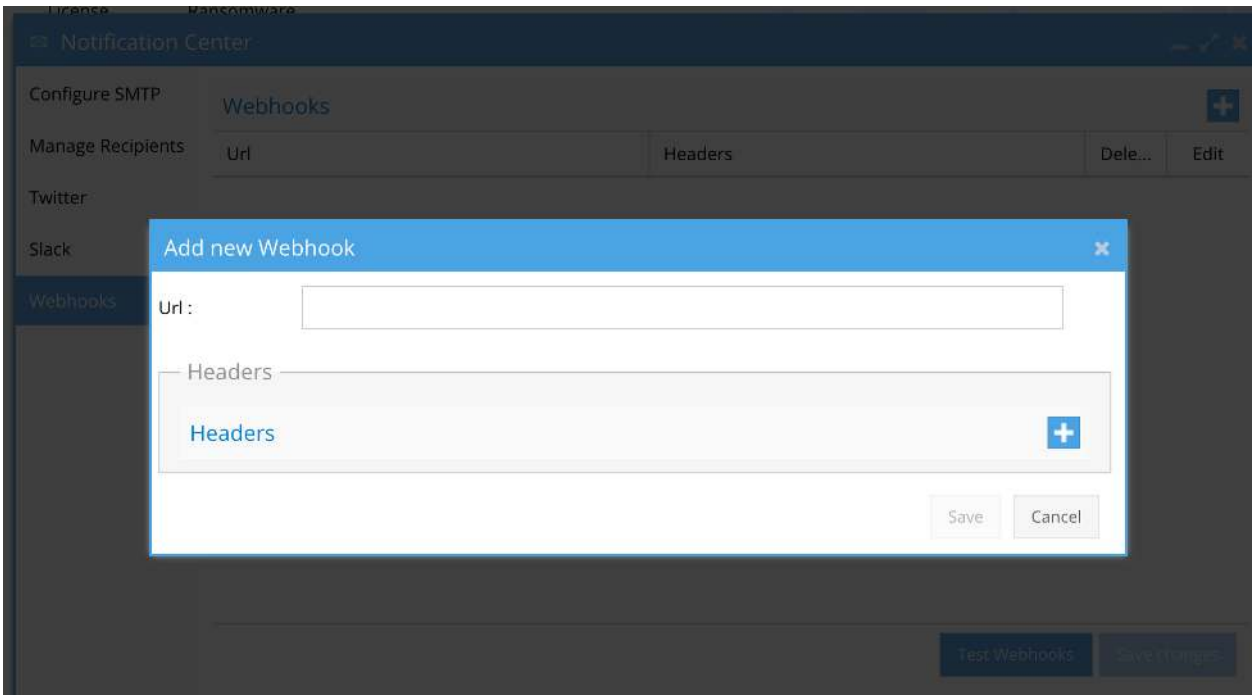

3. Enter url of the target system to receive the post command.
4. Click plus headers to enter any key value pairs needed to send to the upstream
   system

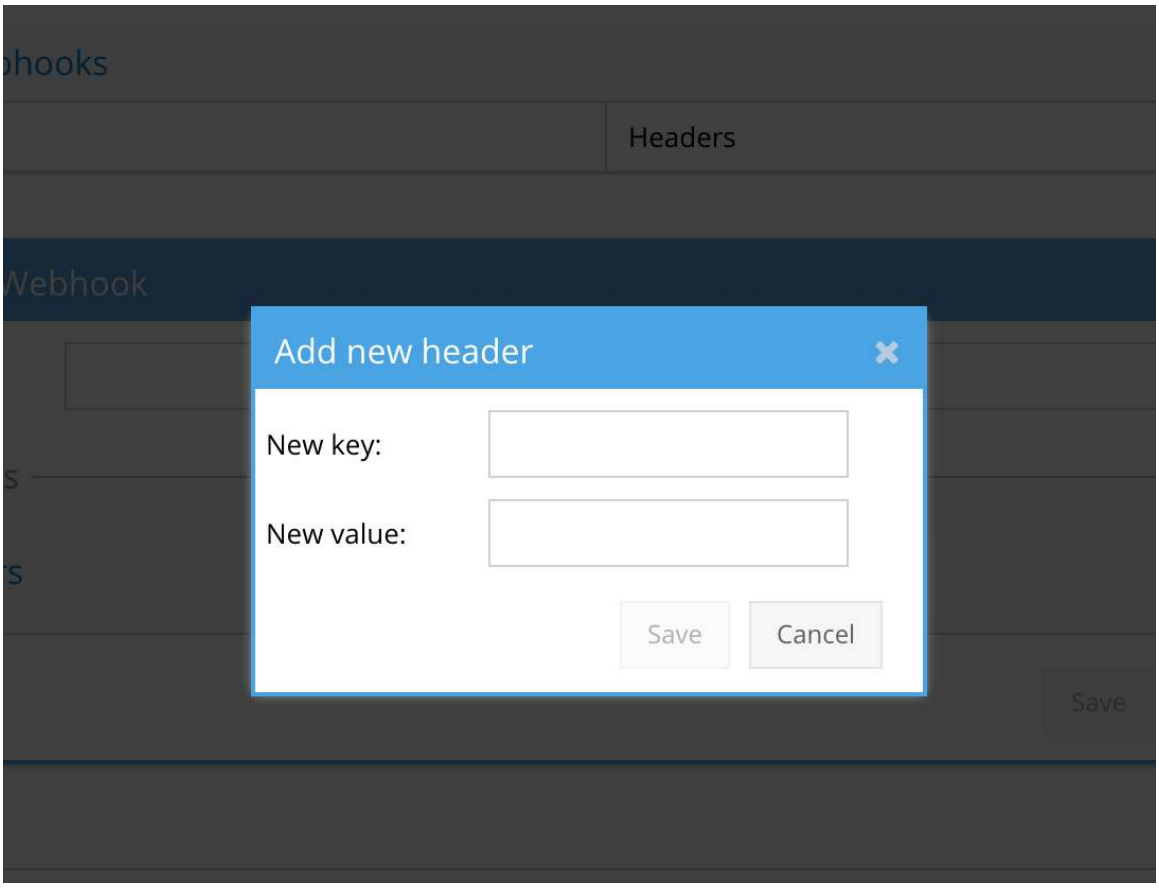

5. Once completed test with the test button and verify with the upstream system the test even was sent.
6. Now all alarms Eyeglass and PowerScale will be sent as Webhooks.

How to Configure Private Tweets of Alarms

Twitter App Setup:

1. https://twitter.com/signup and sign up for an account. Please note, this account is the account responsible for tweeting alarms. (Best to create an eyeglass specific account from scratch).

*You don't need to complete the 6-step introduction tutorial Twitter starts.

1. Once created check your inbox for the twitter verification email and verify it.
2. With the account verified, in order to use twitter apps the account must have a mobile phone number associated and verified.

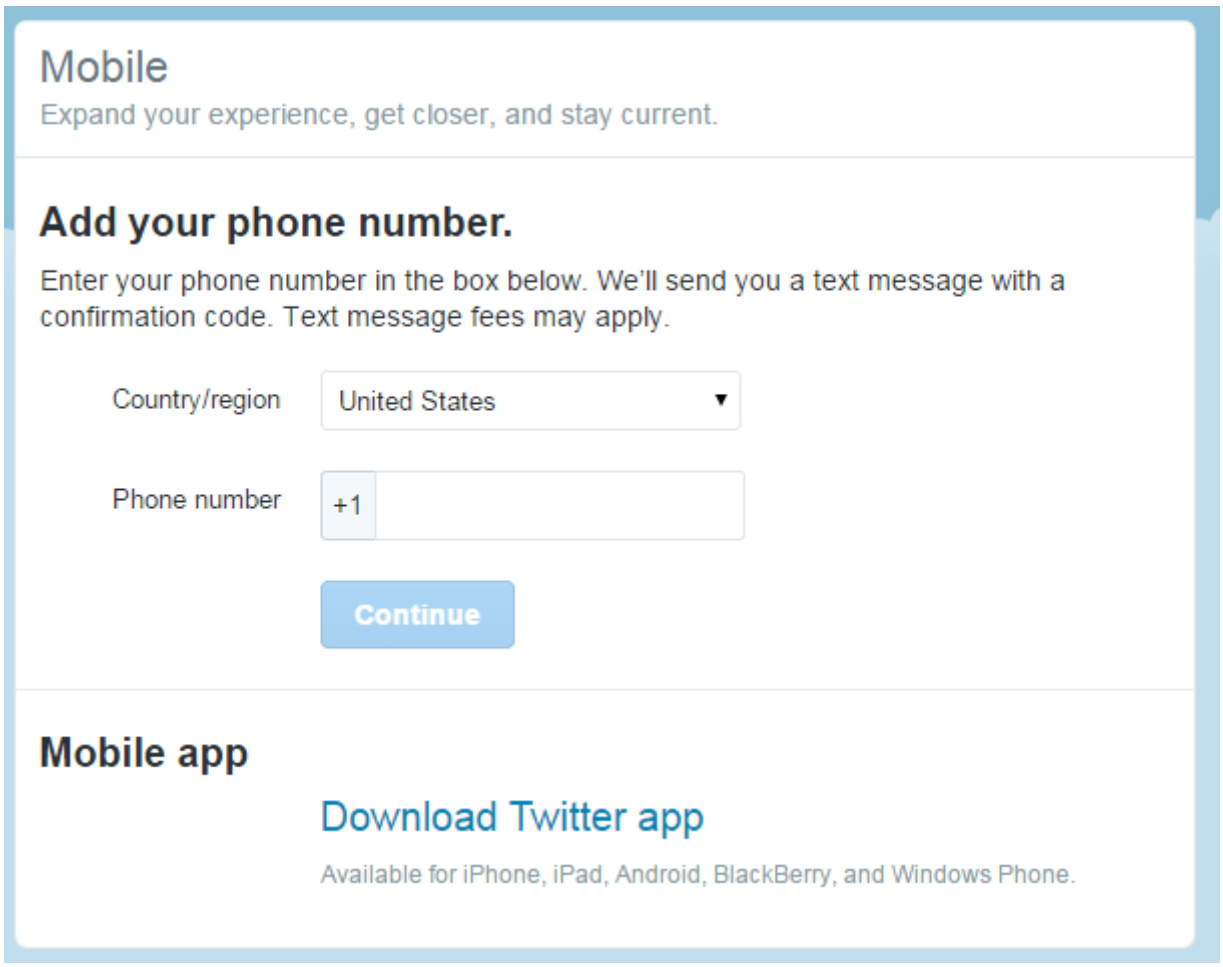

3. Go to https://apps.twitter.com/
4. Create New App

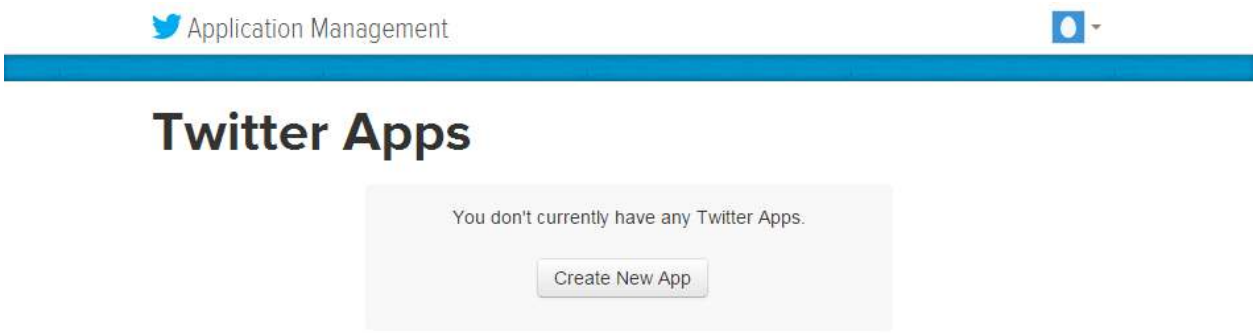

5. Fill in Name, Description and Website and “Create your Twitter application”.

6. With the application created, click the “Permissions” Tab

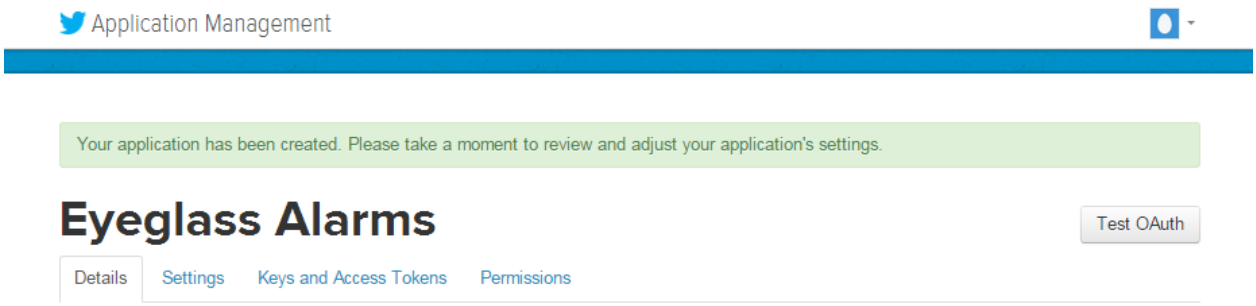

7. Change access to “Read and Write” and Update Settings.

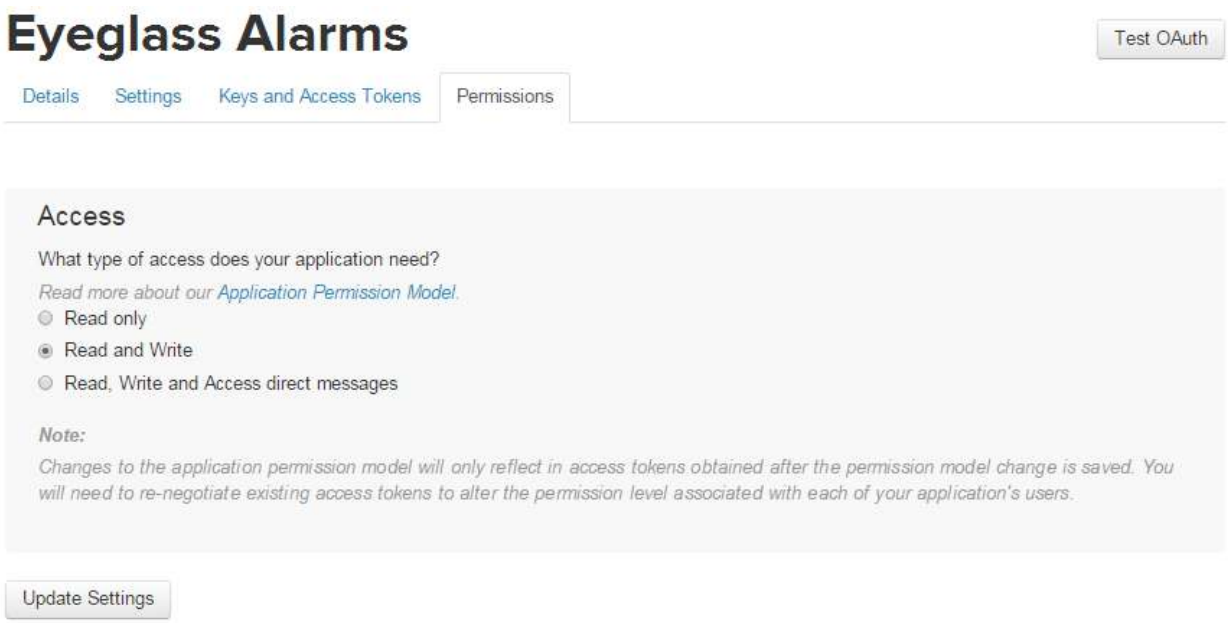

8. Click the Keys and Access Tokens tab, at the bottom of the page click “Create my access token”.

**If you wish to have the alarm tweets private (only to those who your appliance Twitter Account allows to follow) you can select the checkbox called “Protect my tweets” here https://twitter.com/settings/security (If selected, only those you approve will receive your Tweets. Your future Tweets will not be

available publicly. Tweets posted previously may still be publicly visible in some places)

Eyeglass Twitter Setup:

1. Eyeglass Main menu > Notification Center > Twitter
2. Copy the Keys and Tokens from the Twitter Keys and Access Tokens page into the appropriate fields on the Eyeglass UI.
3. Select the Alarm Severity Filter: Critical only tweets critical alarms, Major tweets Major and Critical alarms, etc.
4. Submit
5. Once submitted the account will now tweet out alarms when raised in Eyeglass to the Twitter account previously created.
6. To get real time notifications, we recommend following the recently created appliance on Twitter with your personal twitter accounts. You can find the follow button here: Eyeglass Main menu > Notification Center > Twitter > [Follow “name”] or by clicking the name of the eyeglasses twitter account hyperlinked in both the splash page and the Add/Edit/Delete Twitter Account” page.
7. Here’s a tutorial on setting up notifications with iPhone (same for most smart phones). http://www.wikihow.com/Get-Push-Notifications-for-a-Users-Tweets-on-Twitter-for-iPhone

NOTE: Twitter will apply it's own logic to prevent "re" tweet of what it considers to be the same message. This may prevent tweet of alarms that are raised and cleared and raised again within the period of time that Twitter is monitoring for re-tweets.

© Superna LLC

1.15. ECA (Eyeglass Clustered Agent) CLI Commands

Home Top

- Kafka commands
- Start ECA cluster
- Stop ECA cluster
- Check ECA container logs
- Login to ECA container
- ECA Cluster-wide log
- Run ECA related script cluster-wide
- Restart ECA containers cluster-wide
- Take down single ECA node containers
- Run components install eca
- Get a list of ECA containers and container stats
- Stop a specific ECA container
- Remove a container
- Restart a container
- Start a container
- Delete and re-add a ECA container or a list of containers
- Check ECA cluster-wide service status
- Combine ECA cluster-wide commands
- ECA disk usage check
- ECA cluster-wide disk I/O check
- Useful ECA commands

- Check hbase health for inconsistencies
- Check fastanalysis container file for extension
- Run ECA container in debug mode
- Schedule cron to restart containers

Kafka commands

ecactl zk --help (commands for zookeeper)

ecactl zk shell (opens shell)

ls /superna/eca/turboaudit (will list sub folders to get status)

ls /superna/eca/turboaudit/ecanodes

ls /superna/eca/turboaudit/auditfolders (list the cluster folders and select a folder name with /owner and use the get command to find the assigned ECA node of this audit folder on the PowerScale cluster)

Example:

get /superna/eca/turboaudit/auditfolders/00505698f0793f8bbb56fc176e2f7b6e204c_node001/owner

ecactl kafka topics --help

This will list the available flags for kafka.

ecactl kafka topics --describe

Which returns info about all kafka topics. If these commands do not error out, this can be considered a successful test.

ecactl cluster services list (lists each node and the services running on that node in a distributed cluster with different services running on each node.)

Basic ECA commands

Start ECA cluster

```
ecactl cluster up
```

Stop ECA cluster

```
ecactl cluster down
```

Check ECA container logs

```
ecactl logs <CONTAINER>
```

You can tail the container logs

```
ecactl logs --follow <CONTAINER>
```

You can also tail 'n' number line from ECA container logs

```
ecactl logs --follow --tail 100 <CONTAINER>
```

e.g.: This will output last 100n lines log

Mix command with "grep"

```
ecactl logs --follow --tail 100 <CONTAINER>  | grep -i '<KEYWORD>'
```

Mix multiple container logs with "egrep"

```
ecactl logs <CONTAINER> <CONTAINER>  | egrep -i '<KEYWORD>|<KEYWORD>'
```

Login to ECA container

```
ecactl containers exec <CONTAINER> /bin/bash
```

ECA Cluster-wide log

```
ecactl cluster exec'ecactl logs --follow --tail <CONTAINER>'
```

Run ECA related script cluster-wide

```
ecactl cluster exec '/opt/superna/eca/scripts/<SCRIPT_NAME>.sh'
```

Restart ECA containers cluster-wide

```
ecactl cluster exec 'ecactl containers restart <CONTAINER>'
```

Take down single ECA node containers

Login to ECA node X (Where ‘X’ represents the node ID)

```
ecactl containers down
```

To bring up ECA containers again, login to ECA master node 1

```
ecactl cluster up
```

Run components install eca

needed before first cluster up (2.5.4)

```
ecactl components install eca
```

Get a list of ECA containers and container stats

```
ecactl containers ps
ecactl cluster status
```

Stop a specific ECA container

```
ecactl containers stop <CONTAINER>
```

Remove a container

```
ecactl containers rm -f <CONTAINER>
```

Restart a container

```
ecactl containers restart <CONTAINER>
```

Start a container

```
ecactl containers up -d <CONTAINER>
```

Delete and re-add a ECA container or a list of containers

NEW 2.5.5 + Command:

```
ecactl cluster services restart --container <CONTAINER>
ecactl cluster services restart --all --container <CONTAINER>    ← Restart container on ALL nodes
```

Under the hood:

```
ecactl containers stop <CONT> && ecactl containers rm -f <CONT> && ecactl containers up -d <CONT>
```

Check ECA cluster-wide service status

```
ecactl cluster exec "systemctl status autofs"
```

Combine ECA cluster-wide commands

```
ecactl cluster exec '<COMMAND>' && ecactl cluster exec '<COMMAND>'
```

e.g.: ecactl cluster exec 'docker system df -v' && ecactl cluster exec 'docker stats -a --no-stream'

ECA disk usage check

```
ecactl cluster exec "df -h"
```

Mix command with “grep”

```
ecactl cluster exec "df -h | grep -i zk*"
```

e.g.: shows zk-ram disk mount information

ECA cluster-wide disk I/O check

```
ecactl cluster exec 'iostat -xyz'
```

Useful ECA commands

Run same command across ECA nodes

```
ecactl cluster exec "<command>"
ecactl cluster exec 'top -n 1 | grep -i "kib"'
```

```
ecaadmin@eca-1:~> ecactl cluster exec 'top -n 1 | grep -i "kib"'
KiB Mem:  16419312 total,  6267316 used, 10151996 free,       12 buffers
KiB Swap:  2103292 total,   631608 used,  1471684 free.  4940940 cached Mem
Connection to 172.25.3.173 closed.
KiB Mem:  16419312 total,  8285048 used,  8134264 free,      552 buffers
KiB Swap:  2103292 total,   918196 used,  1185096 free.  5628848 cached Mem
Connection to 172.25.3.174 closed.
KiB Mem:  16419312 total,  6821768 used,  9597544 free,      756 buffers
KiB Swap:  2103292 total,   620272 used,  1483020 free.  5115308 cached Mem
Connection to 172.25.3.175 closed.
```

e.g.: ecactl cluster exec 'sudo mount -a’

e.g.: ecactl cluster exec 'sudo umount -l /opt/superna/mnt/audit/_GUID_/_clusterName_’

e.g.: ecactl cluster exec 'sudo systemctl mask autofs’

e.g.: ecactl cluster exec 'sudo systemctl restart docker’

e.g.: ecactl cluster exec 'docker network ls | grep -i node*'

e.g.: ecactl cluster exec '/opt/superna/eca/scripts/delete_network.sh'

Check hbase health for inconsistencies

Hbase health check and repair covered in detail HERE

```
ecactl containers exec hbase-master /bin/bash
bin/hbase hbck
```

Check fastanalysis container file for extension

```
ecactl containers exec fastanalysis /bin/bash
cat ransomwareFilters.json
```

Run ECA container in debug mode

```
ecactl containers stop <container_name>
ecactl containers rm -f <container_name>
ecactl containers up -d <container_name> --debug
ecactl logs --follow <container_name>
```

Schedule cron to restart containers

1. ecactl cluster exec "sudo -E ecactl components restart-cron set fastanalysis,evtarchive,turboaudit 0 0,6,12,18 \'*\' \'*\' \'*\'"
    a. This command will set cron on three containers using the cron string shown above. **The above cron string is the default that should be used. Note: the outer double " is required for the ecactl cluster exec command**
    b. password for each node will be required

2. sudo -E ecactl components restart-cron set <container>[,<container>[...]] <cron interval>

3. sudo -E ecactl components restart-cron remove <container>[,<container>[...]] (this removes the cron setting for 1 or more containers)
4. NOTE: can be executed on a single node

© Superna LLC

1.16. DR Design Guides with Eyeglass

DR Design Guides with Eyeglass

The Steps to configure, plan and implement DR operations with Eyeglass is located here in a dedicated chapter on DR options including Run Book Robot implementation and design.

1. Eyeglass Start Here first
2. Eyeglass Failover Design Guide

© Superna LLC

1.17. Eyeglass CLI Commands

Home Top

- Security CLI Commands
  - Change cluster service account password CLI Command
- Appliance Management
  - Disk Management Monitoring
  - Alarm Database Table Compression Enabled or Disabled
- Easy Auditor CLI commands
  - Robot Audit
  - Event result percent for Robot Audit
  - Where did My Folder Go Query Event Limit
- Easy Auditor And Ransomware Defender common CLI commands
  - Easy Auditor and Ransomware Queue Reset to process only recent events versus backlog in the queue
- Ransomware CLI commands
  - Convert Ignored list to Monitor Only Settings list (> 2.5.7)
  - Lockout Root SID with SMB disable
  - Default Snapshot Expiry
  - Security Guard Delay Detection
  - NFS Event Processing and lockout
  - Extensions Whitelist
  - False Positive Override per user
  - Restore a locked out User when Security Event in Error State

- Ransomware Defender Banned file Version Commands
- Ransomware Defender AirGap CLI commands
  - Temporary Maintenance Commands Airgap Enterprise
  - Temporary Maintenance Command Airgap Basic
- igls adv failovermode
- igls adv failovertimeout
- igls adv full sync
- igls adv runbookrobot
- igls admin ignoreunresolvablehosts
- igls admin health
- igls admin appid
- igls admin version
- igls alarm active
- igls alarm all
- igls alarm settings
- igls appliance upgrade
- igls appliance restore
- igls appliance rediscover

- Advanced CLI Commands
  - igls adv adserver
    - example command
  - igls adv initialstate
  - igls adv PolicyFailover

    - Igls adv rundedupe
    - igls (disable new validations)
    - igls admin schedules
- Eyeglass Reports
    - The cluster diff report is now run from an igls command. If the the cluster has a large configuration, this report can run for hours. The daily report will not difference configurations and will only send the basic report.
    - Update Schedule
- Configuration Replication
    - Enable/Disable
    - Update Schedule
- Runbook Robot Schedule Interval
    - Enable/Disable
    - Update Schedule
- Failover Readiness for Access Zones and IP Pools
    - Enable/Disable
    - Update Schedule
- Runbook Robot Mount Export Enable Disable
- Advanced Commands
    - igls adv requesttimeout
    - igls adv spndelay
- Role Based Authentication Authentication CLI Commands
- Cluster Storage Monitor CLI commands

- Cluster Storage Reports Schedule CLI Commands
  - igls admin schedules (lists schedules)
- Quota collection Job Schedule for Large Quota Clusters Scheduling CLI Commands
  - How to Enable Dedicated Quota Inventory Collection Job and Quota pre-sync
- Advanced Quota Failover and Inventory collection CLI commands
- Cluster Storage Monitor automation quota commands (Cluster Storage Monitor Feature)
- Active Directory Group based Quota Management (Cluster Storage Monitor Feature)
- High level requirements to use this Feature
- How to configure AD group quota scheduled job (Required)
- How to start an onDemand AD Quota scan of AD and Quota creation
- User quotas or group quota templates
- igls csm tier help
- igls csm template help
- igls csm tier
- igls csm template
- igls csm template update
- Home Share AD Managed quota Configuration example
- Group Share AD managed quota Configuration example

    - igls adv adgroupmode
- RPO Reporting CLI Commands
- CSM Reporting CLI Commands
    - igls adv runreports --report_type=csm
- Advanced Commands Use if directed by support
    - Memory watch dog on Eyeglass
    - Database insertion validation
    - AD User/group to SID or SID to user/group (Easy Auditor, Ransomware Defender)

Eyeglass CLI Commands

The following Eyeglass CLI commands are available and can be executed directly from the Eyeglass shell or any ssh session to the appliance.

Security CLI Commands

This section covers security related cli commands

Change cluster service account password CLI Command

1. (2.5.6 or later required) igls adv changepwd --cluster T-A8200 --password 1 --restart true
    a. **NOTE: --restart true is required for the password to take effect, and will restart the SCA process to have the password change take effect. After using this command, login to the UI, open the jobs icon and verify configuration sync jobs are completing successfully. If restart true is not used, the SCA will need to be restarted at a later time using the command sudo systemctl restart sca.**
    b. **The Eyeglass service account must be used. AD accounts are not supported and not best practice since it reduces availability of the system with dependency on AD DC's**
    c. **igls adv changepwd help**

Appliance Management

Disk Management Monitoring

1. This command and configuration file allows increasing the support backup disk monitor alarm from the default of 800 MB to a higher value. The cli command below re-loads the new monitoring threshold.
2. Edit this file to change the 2nd rule to a higher value in MB /opt/superna/sca/conf/DiskSpaceMonitorConfig.xml save the file, and run the command below, to allow storing more backups before the alarm will be triggered.

3. If you get this alarm you can delete old backup files located in /srv/www/htdocs/archive
4. igls adv reloaddiskspacemonitorrules (this command will update the monitor to use the new limits configured in)

Alarm Database Table Compression Enabled or Disabled

1. igls adv managealarmdatacompression (this command checks if alarm data compression is abled to store in the database, contact support before attempting to change this)

Easy Auditor CLI commands

These commands are used for Easy Auditor configuration changes.

When a cluster has had auditing enabled a history of audit logs are stored on the cluster. This CLI command can be used to ingest old audit messages for searching. It can also be used to ingest data while ECA cluster was down or unable to reach a cluster to process audit messages.

igls rawsignals bulkLoadTAEvents --file=/opt/superna/sca/tmp/bulkLoadTAConfig.json

A json file is used to specify the cluster, the node names and the compressed files on PowerScale that should be ingested.

Sample file exists on the appliance to edit /opt/superna/sca/tmp/bulkLoadTAConfig.json

Example json file below

[{ "cluster_name": "sourcein8", "cluster_guid": "0050569f9a9f4d819b58261e950907a632ad", "node": [{ "node_id": "node001", "audit_files": ["00000000.gz", "00000001.gz", "00000002.gz", "00000003.gz"] }, { "node_id": "node002", "audit_files": ["00000000.gz", "00000001.gz", "00000002.gz", "00000003.gz"] }] }]

igls admin eaCsvArchivePath show

Use this to show the current location that csv reports are stored or change this location.

The location can be changed to an NFS mount on the Eyeglass appliance, to allow centralized reports to be stored automatically as they are generated by users. This keeps a secondary copy of all reports or searches executed with Easy Auditor.

Use:
igls admin eaCsvArchivePath set --value=<path where csv files should be saved> (to set the path.)
igls admin eaCsvArchivePath (will show you where they are currently saved. There should be no default, so initially they are not saved anywhere, but the step in the report job will be marked as successful if no path is set.)

Robot Audit

This feature performs continuous auditing by creating user events as an SMB connected user. The events are created, ingested and stored in the database. The Robot audit process runs reports, and counts file and directory events, and logs success or failure. This offers the highest level of confidence that audit data is being processed and stored. The audit lag is the time from when an event is created to when the data is searchable. This sets the time lag value to a value that avoids robot audit failures for a particular environment where event rate may require an increased value. Value is in minutes.

igls easyauditor roboaudit (shows current value)

igls easyauditor roboaudit set --eventlag=15 (sets new value to 15 minutes)

igls easyauditor roboaudit set --runpathreport=true (this command disables the path search that can take variable time to complete, and will exclude this from running as part of the test automation)

Event result percent for Robot Audit

Use only directed by support

igls easyauditor roboaudit set --reportpercent=60

Where did My Folder Go Query Event Limit

This command will set the limit of the number of events returned with a Where did my folder go search. NOTE: The feature has been tested to return 5000 events, over this limit may overwhelm the browsers ablity to display the data.

igls easyauditor folderquerylimit set --limit=2500

Easy Auditor And Ransomware Defender common CLI commands

Easy Auditor and Ransomware Queue Reset to process only recent events versus backlog in the queue

Easy Auditor has real time triggers, if a misconfigured trigger is configured, a back log of detections will end up in the queue to be processed. This will take a long time to process when it was setup in error. This command also applies to Ransomware Defender if a large number of detections occur. Each product has a separate real-time queue for processing, with Ransomware Defender processing taking priority.

If a large back log occurs from an Easy Auditor, or user activity flagged by Ransomware Defender, the following command can be used to skip to the end of the queue, to effectively ignore all previous detections and start processing from the end of the queue.

igls adv eventTriggers set --operation=reset --topic=rsw (will reset the processing on the Ransomware Queue)

igls adv eventTriggers set --operation=reset --topic=ea (will reset the processing on the Real time Active audit triggers Queue for Easy Auditor)

Ransomware CLI commands

Convert Ignored list to Monitor Only Settings list (> 2.5.7)

This command will convert all ignored list entries for path, user or ip address to the monitor list in 2.5.7 or later releases. This is the preferred method to protect data without a lockout applied. The same matching behavior works with the monitor mode list entries.

This command will convert all previous entries on the ignored list to monitor only list

igls rsw convertignoredlist

NOTE: All existing Ignored List entries are deleted and moved to Monitor Mode

Lockout Root SID with SMB disable

Use this command to enable SMB shutdown if root PowerScale user is detected with Ransomware behavior. Cluster wide shutdown of all SMB IO. **NOTE: Root user should never be used to access data since it can access all shares regardless of permissions**

igls admin lockroot --lock_root

General Ransomware Settings

Use this command to see general Ransomware Defender settings. **NOTE: some settings are managed in the GUI.**

igls rsw generalsettings

Sample output :


```
{
    "snapshot_expiry_hours": 48,
    "escalate": false,
    "critical_on": true,
    "monitor_only": false,
    "snapshotOn": "WARNING",
    "lock_root": false,
    "root_sids": [
        "S-1-1-1-0",
        "S-1-22-1-0"
    ]
}
```


Default Snapshot Expiry

Use this command to set the expiry default time on pro-active snapshots.

```
igls rsw generalsettings set --snapshot_expiry_hours 72
```

Security Guard Delay Detection

Use this command to change security timer to delay failure message when audit events are behind on the cluster.

igls rsw securityguardsettings help

demo2:/opt/superna/sca/conf # igls rsw securityguardsettings help

show(default):

Provides the following options:

1. Set security guard wait for event timer in seconds.

2. Set security guard restore timer in seconds.

set --<option>=<value>

Valid Options: sg_waitforevent_timer_seconds and sg_restore_timer_seconds

Example - set event timer :

igls rsw securityguardsettings set --sg_waitforevent_timer_seconds=600

Example - set restore permissions timer :

igls rsw securityguardsettings set --sg_restore_timer_seconds=600

NFS Event Processing and lockout

This section shows how to enable NFS IO processing and enable lockout for Ransomware events. This single command will process NFS IO, and apply lockouts based on thresholds configured in the GUI. The lockout function will remove the client ip address from the export definition to lockout the NFS host. It will not check host names

to map them to IP addresses in this release. Consult documentation for Lockout configuration.

To set the value from default of disabled to enabled:

igls rsw nfsevents set --enabled=true

To check the value that is currently set for NFS lockout:

igls rsw nfsevents

Extensions Whitelist

The file extension list tracks over 2000 well known extensions used in Ransomware incidents. Sometimes these are valid extensions in customer environments or applications. These CLI commands can be used to whitelist extensions by adding them to an allowed list.

igls rsw allowedfiles add --extensions='*.ext1' (add an extension to the list)
igls rsw allowedfiles (list all allowed files)
igls rsw allowedfiles remove --extensions='*.ext1' (remove an extension from the list)

False Positive Override per user

Contact support to enter values. Commands are provided as a reference, but support should be involved to provide values to modify user threat level settings. These command use SID to add or delete. The side cache file can be used to find the user name for a given SID. To read this file login to Eyeglass VM via ssh and run this command "cat /opt/superna/sca/data/ad_principal_cache.json" to see the SID and user names.

- Add an override for a user
    - igls rsw RSWUserOverride post --user=S-1-5-21-826284354-1834749432-1846952604-4825 --tdid=03 --parameter=X --multiplier=16.008001000000004
- Delete an override for a user
    - igls rsw RSWUserOverride delete --sid=S-1-5-21-826284354-1834749432-1846952604-4825 --tdid=03 --parameter=X --multiplier=16.008001000000004
- View the current overrides set by flag as false positive
    - igls rsw RSWUserOverride get

all command line inputs and outputs will use a SID. Resolve these to and from the SID format using the cache file above or AD tools. All entries in the file should be by SID. "[SID] : [Threat detector ID] : [Parameter ID] : [Multiplier]"

Restore a locked out User when Security Event in Error State

This command would be used only when a locked out user security event has an error due to cluster reachability, or some shares were not unlocked. This command will re-attempt the recovery of a user's share access. NOTE: do not use this command if you assigned users directly to shares versus using AD groups.

igls rsw RSWRestoreAccess set --user=<value> (where value is domain\user, note use uppercase domain name and quotes)

example: igls rsw RSWRestoreAccess set --user="TESTDOMAIN\usera"

A new igls CLI is defined to restore all access to all shares for a user. The cli command is "igls rsw RSWRestoreAccess set --user=<value>"

This CLI scans the Eyeglass database, finds all shares that have a "deny" for the user based on AD group membership only.
The CLI builds and executes the restore job for all of these shares.
This CLI is only used by Superna support in the case where there was a lockout that reported no shares successfully locked out, but where deny flags actually were applied on the PowerScale.

Test Step:

Have an Error state where no shares are saved in the database (i.e. "none lockedout" under "Shares" column in Ransomware Defender table).
Apply the command.
Expected Results:

The user is restored accessed to all shares he was assigned to.
Note:
Restoring access to a user who happened to be in a group added to this share. And not for a user who is directly added to this share.

Ransomware Defender Banned file Version Commands

1. These commands require 2.5.7 update 1 release
2. These commands allow apply a new version of the banned file lists from the Eyeglass Ransomware Defender or compares two versions for differences. The mode can be set to auto switch to latest as well.
    a. Note The Eyeglass VM requires Internet access to download and apply new versions of the banned file list.
3. igls rsw filefiltersettings - **list current settings**
4. igls rsw filefiltersettings --diff=<version1,version2> - **Run this command to see the differences between the two versions, it will**

show new extensions added to the file and extensions removed from the list.

5. igls rsw filefiltersettings --version=<version> - **select an available version to switch Eyeglass to the new version of the banned file**
6. igls rsw filefiltersettings set --mode=Latest - **Sets Eyeglass Ransomware Defender to check for new versions and automatically switch to the new version.**
7. igls rsw filefiltersettings set --mode=Fixed --version=<version> - **Sets Eyeglass Ransomware Defender to use a specific version of the file**

Ransomware Defender AirGap CLI commands

igls airgap disable (quick disable all policies, no syncing will occur, this is a maintenance state)

igls airgap enable (quick enable all policies)

If set to disable send major alarm ("airgap policy sync state disabled by administrator")

Temporary Maintenance Commands Airgap Enterprise

These commands only work if the variable for remote maintenance is enabled on the vault agent. It is disabled by default.

1. igls airgap vaultaccessrequest --interval=x (x is minutes)
    a. Sets a request in eyeglass that will be picked up by the vault agent vm. The vault agent checks in every 2 hours for requests. If the request is set the airgap will open for x minutes for maintenance and will auto close after x minutes
2. igls airgap vaultaccessview

   a. This command lists any requests that have been set and the interval requested

3. igls airgap vaultaccesscancel
   a. This command will cancel any pending requests that have been set.

Temporary Maintenance Command Airgap Basic

It is also possible to open an airgap for maintenance but specify a time period to leave the network open before it automatically closes. Use the connect and disconnect commands. The timeout ensures the network is not left open by accident exposing data to the network.

igls airgap connect help

Usage:
igls airgap connect --job=<job-name> --timeout=<duration>[m|h]
igls airgap connect --policy=<policy> --source=<source> --timeout=<duration>[m|h]

Enable network connection to the target device of an airgap policy with a timeout that will auto close the network connection once the timeout is reached.

Required arguments:

--job=<job-name> Use config settings of given AirGap job to enable network connection to it's target device.

--policy=<policy> Use config settings of AirGap <policy> to enable network connection to it's target device.

--source=<source> Name of the source cluster for the given AirGap policy.

--timeout=<duration>[m|h] Duration of time to keep the connection to the given AirGap policy's target device. You can follow the duration by the = suffix 'm' to indicate minutes or 'h' to indicate hours. If no suffix is given, duration defaults to minutes.

igls airgap disconnect help

Usage:

igls airgap disconnect --job=<job-name> OR use
igls airgap disconnect --policy=<policy> --source=<source>

Disable network connection to the target device of an airgap policy, without waiting for the time out period.

Required arguments:

--job=<job-name> Use config settings of given AirGap job to disable network connection to it's target device.

--policy=<policy> Use config settings of AirGap <policy> to disable network connection to it's target device.

--source=<source> Name of the source cluster for the given AirGap policy.

igls adv failovermode

This command is for a large number of policy failovers. It changes the default behavior of sequential make writable and resync prep commands to allow up to 10 parallel make write commands, or resync prep commands to be issued to the cluster at once. If a job on the cluster finishes, another is sent with the goal of keeping 10 jobs always running on the cluster until failover is complete.

High Speed Failover - Parallel Failover Flag Notes :

1. Allows make write step and resync prep to run in parallel with up to 10 threads, ensures that 10 policies are submitted to be processed at all times.
2. Testing has shown these steps for large quantity policy failover can improve failover times 3x to 4x.
3. Risk of a policy failure increases, and new flag will NOT stop the failover in progress. The process will continue to issue api calls to submit all SyncIQ policies in the failover job until all have been submitted. This runs the risk of more complex recovery if more than one policy fails to complete its step (Allow Writes OR resync Prep) .

igls adv failoverSettings set --parallel=true **(defaults to true as of 2.5.5 release),** not recommended to disable this contact support.

igls adv failovertimeout

Display the per step timeout for failover tasks. Advanced setting. Default 45 minutes. For very large policies (see [Eyeglass and PowerScale DR Best Practices]), igls adv failovertimeout can be increased to suggest value of 180 minutes.

igls adv failovertimeout get (returns current value)

igls adv failovertimeout set --minutes 180 (sets)

igls adv full sync

This advanced option should be enabled only after consulting with Superna support first. It overcomes a scenario where NFS exports are created with FQDN for client lists, and the FQDN values are NOT resolvable by the DR or target cluster. This scenario happens when DHCP leases expire DNS resolution, OR if FQDN values do not resolve any longer, and it's not possible to clean up this condition. OneFS 8 API behavior denies the creation of the exports with unresolved FQDN client list entries, and requires the force flag to override cluster rules on export creation. The force create override flag is disabled by default in Eyeglass to avoid conditions where duplicate exports are created.

Behavior

This sync mode will delete all shares and exports found on the target cluster that DO NOT exist on the source. This creates a full sync. The

default option in Eyeglass will leave any shares or exports found that do not exist on the source. With this option enabled, all extra config will be deleted to make an exact copy on the DR/target cluster.

igls adv fullsync set --fullsync=<true/false>

Default is false

igls adv runbookrobot

Allows a mode where the export auto create and update is disabled, and can be manually created on the Robot policy path, set the export settings with Eyeglass appliance ip address as root client, and other settings can be enabled manually. Each robot run will no longer create or update the export.

Default is true.

igls adv runbookrobot set --createExport=false

igls admin ignoreunresolvablehosts

This command can be used to enable or disable config sync of exports to allow client lists with unresolvable DNS or Netgroup entries. It is best practice to allow the DR cluster to resolve host names, or data will not be mountable after a failover.

The default setting is disabled and will raise and configuration sync error when attempting to create an export on the DR cluster when the DR cluster cannot resolve the client list host name or Netgroup.

igls admin ignoreunresolvablehosts set --value=true (use this command to allow unresolvable hosts on exports to sync)

igls admin ignoreunresolvablehosts set --value=false (use this command to disable it)

igls admin ignoreunresolvablehosts (use this command to see current value)

igls admin health

Display the overall health status of the Eyeglass appliance.

```
~> igls admin health

{

 "success": true

}
```

igls admin appid

Display the appliance id of the Eyeglass appliance.

```
~> igls admin appid


{

 "applianceCode": [

     ""

 ]

}
```

igls admin version

Display the Eyeglass component versions.

```
~> igls admin version
[
 {
    "release": [
        "38"
    ],
    "version": [
        "1.3"
    ],
    "name": [
        "eyeglass_ui"
    ]
 },
 {
    "release": [
        "34"
    ],
```

```
    "version": [
        "1.3"
    ],
    "name": [
        "eyeglass_rest"
    ]
},
{
    "release": [
        "64"
    ],
    "version": [
        "1.3"
    ],
    "name": [
        "eyeglass_sca"
    ]
}
]
```

igls alarm active

Retrieve the current active alarm list.

~> igls alarm active

{

"sync_key": "Share3-SystemZone",

"code": "SCA0002",

"severity": "Critical",

"timestamp": 1430350806854,

"source": "Share3-SystemZone",

"message": "Found a replication job where either the sou

rce or destination is not a managed network element.",

"extra_data": "{\"info\":\"The replication job for polic

y 'Share3-SystemZone' cannot be created because the target host

cannot be identified.\"}"

}

Note: To view this list incrementally, you can use the command:

~> igls alarm active | more

igls alarm all

Display the total alarms received in “results”.

```
~> igls alarm all
{
 "rows": [],
 "alarmsPerPage": "50",
 "results": "889"
}
```

igls alarm settings

This new command allows controlling the severity of any alarm and can be used to disable an alarm completely. Use with caution.

igls alarm settings help

Allows to set the following options for alarms:
1. Enable or disable raising alarm. Setting "raise" to false disables raising alarm.
2. Enable or disable alarm email notification. Setting "email" to false disables alarm email.
3. Setting alarm severity.
set --code=<AlarmCode> --raise=[false|true] --email=[false|true] --severity=[informational|warning|critical|major|minor|fatal]

--raise= false (disables the alarm)

--severity= (sets the severity of the alarm to the value entered here)

--email= (sends alarm true or false, if false it will display in the gui but no email or other method of alarm notification will be executed)

List of alarm codes can be found here.

igls appliance upgrade

Use for on line upgrade of the appliance software.

Usage: igls app upgrade [OPTIONS]

Download Eyeglass installer - update Eyeglass appliance

Options:

--url TEXT URL of an Eyeglass installer (optional)

--help Show this message and exit.

igls appliance restore

Restore Eyeglass data and configuration from Eyeglass Archive.

Note: must be logged in as admin or root user.

~> igls app restore

Release 2.5.5 or below

1. Usage: igls app restore [OPTIONS]
2. --anyrelease (this allows a version mismatch between the backup file appliance version and the target appliance version, this will skip restore of many items, check the upgrade guide for details. Retains licenses, and clusters and passwords)

Release 2.5.6 or later

1. igls app restore /srv/www/htdocs/archive (pass in a path and all available files are scanned to list the most recent backup found, and presents the file name with yes no option to proceed)
2. Or igls app restore /srv/www/htdocs/archive/backupfile.zip (This option allows full path to the file you want to restore)

igls appliance rediscover

This command should be used when directed by support. It will rebuild the Eyeglass database and preserve job status in the jobs icon with release 1.8 or later. The "igls appliance rediscover command" will prompt yes to continue. **NOTE: It will preserve the quota request, data recovery databases.**

Upon completion refresh the UI login screen. Go to running jobs to see initial discovery job is running to repopulate cluster information in the database inventory icon.

Once completed the job definition screen will show the jobs in previous state and show as pending . The jobs will run again on next scheduled interval or, you can force them to run with the "run now" option.

igls appliance report

(diagnostic log parsing tool run command) This command is for dark or secure sites where on site log analysis is required. The report summarizes all api, ssh and other errors, config sync analysis, failover analysis of each attempt and success or failure.

1. Run command: igls appliance report.
2. Wait for the report to complete.
3. See logs report on: https://<eyeglass IP address>/report/ .

Please refer to document: Eyeglass Backup and Restore .

Advanced CLI Commands

igls adv adserver

This command is used to build a user to SID cache information used by Eyeglass, Ransomware Defender, Storage cluster monitor and Easy Auditor. This avoids API lookups for user and AD information. In very large AD environments with 10 000 of thousands of users and groups it is more efficient to collect this information directly from AD domain controllers. This is also more reliable method to collect this information. This command can be used for an AD provider and configure a user to collect this information using LDAP from the domain controller.

igls ad adserver help -- displays help of the command

list(default): it shows all saved AD servers for which the user cache will be built directly from the server and not from the PowerScale set:Allows adding AD servers to the list. Each entry will have the following parameters:

--server=<value> (the AD server name or LDAP provider-it has to match the value from the PowerScale)
--domain=<value> (domain name),
--basedn=<value> (the distinguished name from where the server will search for users),
--logindn=<value> (distinguished name for the connecting user),
--password=<value>
and optional parameters which are:
--loginhost=<value> i(p or hostname of AD domain controller)
--ssl={true|false} (default is false),
--port=<value> (default is 389 for SSL off, 636 otherwise).

delete: deletes the entry specified by:
--server=<value> or
--all=true (deletes all entries of the list).

Examples:

--server=AD2.TEST (note this is the value shown in the PowerScale GUI for auth provider name

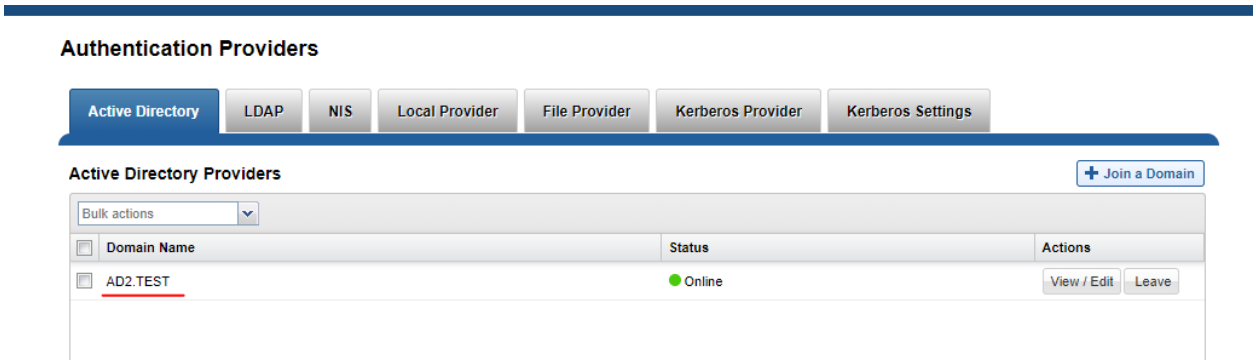

--baseDn=DC=ad2,DC=test (Distinguished to the start the search for users and groups)

--logindn=CN=Administrator,CN=Users,DC=ad2,DC=test (Distinguished to the user, used to authenticate to AD, can be normal user account)

--domain=AD01 (netbios name of the domain)

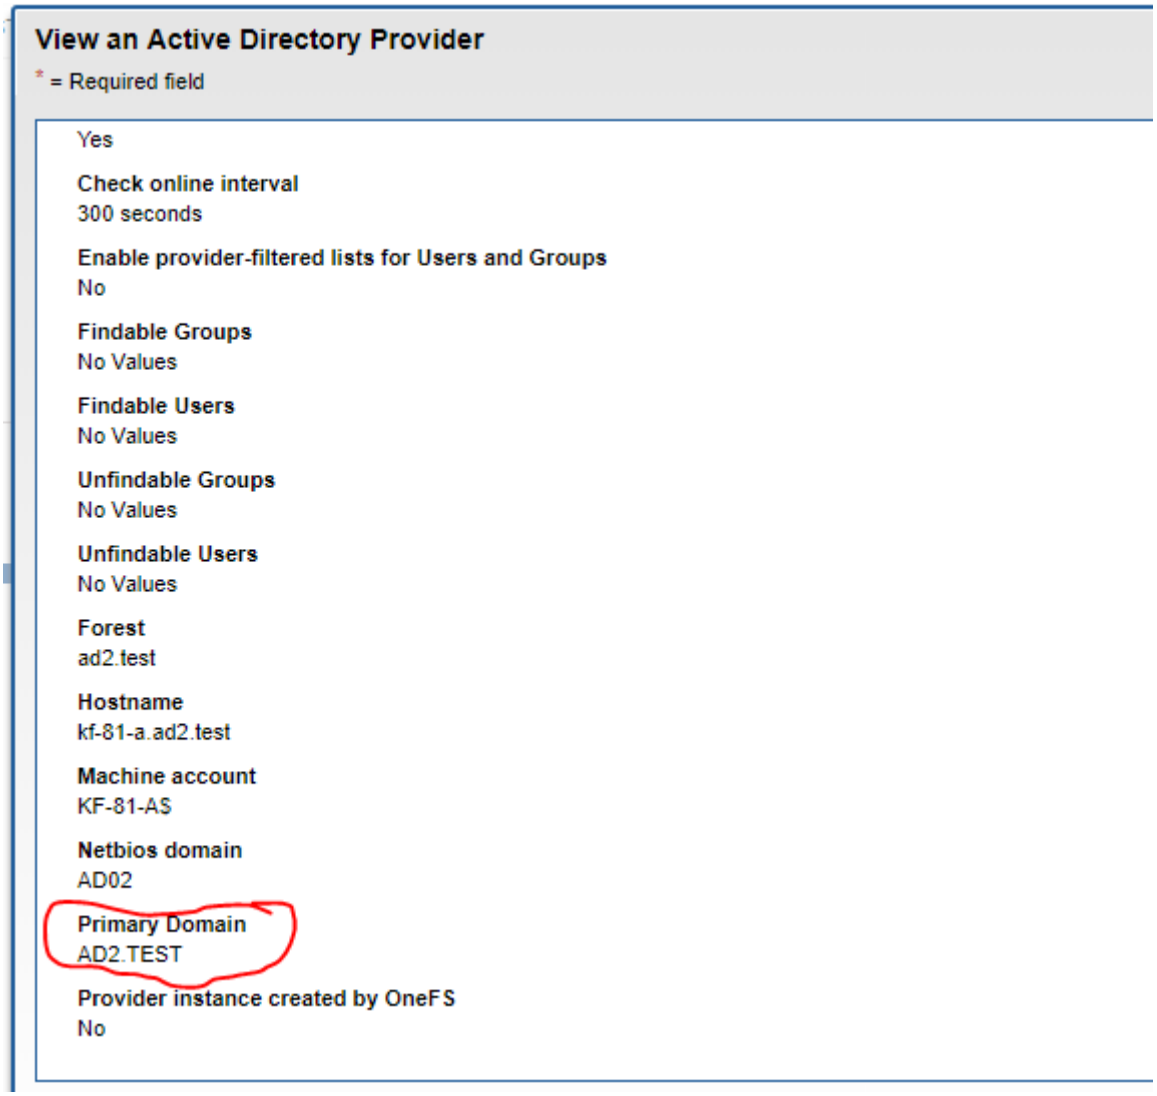

--loginhost= ip address of a domain controller

--ssl=false

--port=389

example command

1. Add AD configuration - see example below change yellow values:

a. igls adv adserver set **--server=AD1.TEST** --basedn=DC=ad1,DC=test --logindn=CN=Administrator,CN=Users,DC=ad1,DC=test **--domain**=AD01 --loginhost=172.16.80.6 **--ssl**=false **--port**=389 --password=3y3gl4ss!

b. Example when users are stored in a different OU:

i. igls adv adserver set --server=RNSM04-05.SUPERNA.NET --basedn=OU=t11543,DC=rnsm04-05,DC=superna,DC=net --logindn=CN=Administrator,CN=Users,DC=rnsm04-05,DC=superna,DC=net --domain=RNSM04-05 --loginhost=172.22.4.155 --ssl=false --port=389 --password=3y3gl4ss!

2. Delete AD configuration:

a. igls adv adserver delete --server=AD1.test

b. igls adv adserver delete --all

3. List all AD configurations:

a. igls adv adserver list

igls adv initialstate

Display and update initial state when Eyeglass creates a new Job. This command supports changing the initial state for the following Eyeglass Job types: ZONES, AUTO, CUSTOM, QUOTAS.

~> igls adv initialstate help

show(default):

Displays the initial states for new jobs.

set --<type>=<state>:

sets a job type to have a specific initial state.

Valid states are: enabled, disabled.

Valid types are: ZONES, AUTO, CUSTOM, QUOTA

Default: is shown below

Examples:

```
~> igls adv initialstate show
{
  "ZONES": "USERDISABLED",
  "AUTO": "ENABLED",
  "QUOTA": "ENABLED",
  "CUSTOM": "ENABLED"
~> igls adv initialstate set --custom=disabled
```

```
{
 "success": true
}
```

igls adv PolicyFailover

Enable and disable Eyeglass Configuration Replication task during SyncIQ Policy Failover.

```
~> igls adv PolicyFailover set --disablereplication=<state>
```

Valid states are: true, false

Examples:

Disable Eyeglass Configuration Replication task during SyncIQ Policy Failover:

```
~> igls adv PolicyFailover set --disablereplication=true
{
 "success": true
}
```

Enable Eyeglass Configuration Replication task during SyncIQ Policy Failover:

```
~> igls adv PolicyFailover set --disablereplication=false
{
 "success": true
```

}

Igls adv rundedupe

Disable dedupe setting process while allowing LiveOPS snapshot jobs to execute. All clusters global command.

igls adv rundedupe set --rundedupe=true/false (default true)

igls (disable new validations)

Prerequisite: 2.5.6 or later

This command will disable SPN AD delegation, dual dns delegation validations if they are not able to execute in a specific environment or of SPN's or dual DNS is not configured.

To display settings:

igls adv readinessvalidation

NOTE: if nothing is returned then default settings are in effect

To modify settings:

igls adv readinessvalidation set --[spnsdelegation|dualdelegation]=[true|false]

Example

igls adv readinessvalidation set --spnsdelegation=false

igls adv readinessvalidation set --dualdelegation=false

igls adv readinessvalidation set --spnsdelegation=false --dualdelegation=false

igls admin schedules

Display and update schedule for Eyeglass tasks. This command supports enabling, disabling and updating the schedule for the following tasks: Configuration Replication, Eyeglass Reports, Zone Readiness, Runbook Robot.

```
~> igls admin schedules list

[

  {

    "interval": "*/1 * * * *",

    "enabled": true,

    "id": "EventAuditProgress",

    "label": "Event Audit Progress Monitoring"
```

NOTE: Used with Easy Auditor to check audit lag

```
  },

  {

    "interval": "0 0 * * *",
```

"enabled": true,

"id": "InventoryReport",

"label": "Eyeglass Reports" **NOTE: Used with DR product configuration reports**

},

{

"interval": "*/15 * * * *",

"enabled": false,

"id": "PrintInventoryToSyslog",

"label": "Print Inventory to Syslog" **NOTE: Used with DR product**

},

{

"interval": "0 0 * * *",

"enabled": true,

"id": "QuotaRequestsReport",

"label": "Quota Requests Report" **NOTE: Used with cluster Storage Monitor product**

},

{

"interval": "*/1 * * * *",

"enabled": true,

"id": "RSWEventsMonitor",

"label": "Ransomware Events Monitoring" **NOTE: Used with Ransomware Defender**

},

{

"interval": "*/1 * * * *",

"enabled": true,

"id": "RSWHbaseScan",

"label": "Ransomware Hbase Scanning" **NOTE: Used with Ransomware Defender to check DB health at an interval**

},

{

"interval": "*/15 * * * *",

"enabled": true,

"id": "Readiness",

"label": "Zone Readiness" **NOTE: Used with Dr product assess one and pool readines**

},

{

"interval": "0 0 * * *",

"enabled": true,

"id": "RecoveryShareCleanUp",

"label": "Recovery Share Clean Up" **NOTE: Used with cluster storage monitor data recovery share deletion check**

},

{

"interval": "*/5 * * * *",

"enabled": true,

"id": "Replication",

"label": "Configuration Replication" **NOTE: Used with DR product to sync config**

},

{

"interval": "0 0 * * *",

"enabled": true,

"id": "RunbookRobot",

"label": "Runbook Robot" **NOTE: Used with DR product to ruh continuous DR feature**

},

{

"interval": "0 * * * *",

"enabled": true,

"id": "SecurityGuard",

"label": "Security Guard" **NOTE: Used with Ransomware Defender to test end to end detection**

},

{

"interval": "*/1 * * * *",

"enabled": true,

"id": "ServicesScan",

"label": "Services Scanning"

},

{

"interval": "0 0 * * *",

"enabled": true,

"id": "StorageMonitorReport",

"label": "Storage Monitor Report" **NOTE: Used with Cluster storage monitor report**

}

]

Eyeglass Reports

The cluster diff report is now run from an igls command. If the the cluster has a large configuration, this report can run for hours. The daily report will not difference configurations and will only send the basic report.

To execute an on demand difference report from today's cluster report to yesterday's use this cli command.

igls adv diffclusterreport

Enable/Disable

To enable/disable the schedule for the Eyeglass Reports, use this command:

igls admin schedules set --id InventoryReport --enabled <true|false>

Examples:

```
~> igls admin schedules set --id InventoryReport --enabled false
{
 "success": true
}
```

~> igls admin schedules set --id InventoryReport --enabled true

```
{
 "success": true
}
```

Update Schedule

To change the schedule for the Eyeglass Reports use this command. Valid intervals for reporting are: 1M, 2M, 3M, 4M, 5M, 6M, 10M, 15M, 20M, 30M, 1H, 2H, 3H, 4H, 6H, 8H, 12H, 1D, 7D, 31D.

~> igls admin schedules set --id InventoryReport --interval <interval>

Example:

~> igls admin schedules set --id InventoryReport --interval 7D

```
{
 "success": true
}
```

Configuration Replication

Enable/Disable

To enable/disable the schedule for Configuration Replication, use this command:

igls admin schedules set --id Replication --enabled <true|false>

Examples:

~> igls admin schedules set --id Replication --enabled false

```
{
 "success": true
}
```

~> igls admin schedules set --id Replication --enabled true

```
{
 "success": true
}
```

Update Schedule

To change the schedule for Configuration Replication use this command. Valid intervals for replication are: 1M, 2M, 3M, 4M, 5M, 6M, 10M, 15M, 20M, 30M, 1H, 2H, 3H, 4H, 6H, 8H, 12H, 1D, 7D, 31D.

igls admin schedules set --id Replication --interval <interval>

Example:

~> igls admin schedules set --id Replication --interval 10M

```
{
 "success": true
```

}

Runbook Robot Schedule Interval

Use this CLI command to change the interval from once per day.

Enable/Disable

To enable/disable the schedule for Runbook Robot, use this command:

igls admin schedules set --id RunbookRobot --enabled <true|false>

Examples:

```
~> igls admin schedules set --id RunbookRobot --enabled false

{
 "success": true
}
~> igls admin schedules set --id RunbookRobot --enabled true

{
 "success": true
}
```

Update Schedule

To change the schedule for RunbookRobot use this command. Valid intervals for Configuration Replication are: 1M, 2M, 3M, 4M, 5M, 6M, 10M, 15M, 20M, 30M, 1H, 2H, 3H, 4H, 6H, 8H, 12H, 1D, 7D, 31D.

igls admin schedules set --id RunbookRobot --interval <interval>

Example:

```
~> igls admin schedules set --id RunbookRobot --interval 10M

{
 "success": true
}
```

Failover Readiness for Access Zones and IP Pools

Enable/Disable

To enable/disable the schedule for the Zone Readiness job, use this command:

igls admin schedules set --id Readiness --enabled <true|false>

Examples:

```
~> igls admin schedules set --id Readiness --enabled false

{
 "success": true
```

```
}

~> igls admin schedules set --id Readiness --enabled true


{

 "success": true

}
```

Update Schedule

To change the schedule for the Zone Readiness job use this command. Valid intervals for reporting are: 1M, 2M, 3M, 4M, 5M, 6M, 10M, 15M, 20M, 30M, 1H, 2H, 3H, 4H, 6H, 8H, 12H, 1D, 7D, 31D.

```
~> igls admin schedules set --id Readiness --interval <interval>
```

Example:

```
~> igls admin schedules set --id Readiness --interval 2H


{
```

Runbook Robot Mount Export Enable Disable

Default is enabled to mount the cluster and create the test file:

igls adv runbookrobot show (show current value)

igls adv runbookrobot set --mount=true (default)

igls adv runbookrobot set --mount=false

Advanced Commands

igls adv requesttimeout

Description: Sets rest API timeout when cluster or wan responses take longer to return, this value can be increased.

igls adv requesttimeout (displays the timeout value)

igls adv requesttimeout set --inventory <time> (sets the timeout value to <time>)

Example:

igls adv requesttimeout set --inventory 300

igls adv spndelay

Description: used to increase the delay between SPN failover commands, that require domain controller to replicate the delete before the add spn can succeed. Release 1.8.3 removes the need for this command, by pinning spn failover commands to a single node and domain controller.

igls adv spndelay (displays the current setting)

igls adv spndelay set --seconds=<seconds> (set a delay between delete and create SPN during failover)

Example:

igls adv spndelay set --seconds=10

Role Based Authentication Authentication CLI Commands

When proxy login is used to login using pass through authentication to AD an ip address is needed to send an SMB authentication request for the users AD user id and password. This is done using the SMB protocol to a share on PowerScale using the SmartConnect FQDN to a share discovered in the zone. In order to provide a list of SmartConnect FQDN's to use to test the users password, use the CLI commands below. The authentication will use the ordered list below until the password test succeeds to allow the user to login.

During the process the AD groups are also retrieved to map the user to a role defined in the Users icon. If the user matches an AD group role, or a user added directly to a role, then the icons or permissions assigned to the role will display on the Eyeglass desktop.

These commands are used to add SmartConnect FQDN to the Access Zone where the users should authenticate to the AD provider assigned to the Access Zone.

1. For checking fqdn list: igls admin auth
2. For adding a new fqdn: igls admin auth add --fqdn <name>
3. For changing a fqdn: igls admin auth modify --fromfqdn <name> --tofqdn <newName>
4. For deleting a fqdn: igls admin auth delete --fqdn <name>
5. For deleting all fqdn's: igls admin auth delete --all true

Cluster Storage Monitor CLI commands

Cluster Storage Reports Schedule CLI Commands

See setting and getting schedules in the CLI section above, to specify job id, get and set functions, and enable and disable actions on scheduled reports.

igls admin schedules (lists schedules)

Used to enable or disable the daily report for cluster disk usage and quotas:


```
{
    "interval": "0 0 * * *",
    "enabled": true,
    "id": "StorageMonitorReport",
    "label": "Storage Monitor Report"
  }
```


Quota collection Job Schedule for Large Quota Clusters Scheduling CLI Commands

Use this CLI command to set a different collection interval for clusters with > 1000 quotas to avoid long configuration sync jobs that detect share, exports for DR syncing. It is best practice to always enable quota job to remove the quota collection from normal configuration sync. This will also remove quota collection from the inventory job on during SCA restart, normal configuration sync jobs and failover inventory step.

NOTE: 2.5.7 update 2 this job will be enabled by default.

How to Enable Dedicated Quota Inventory Collection Job and Quota pre-sync

1. igls admin schedules set --id QuotaInventoryCollection_2_5_3 --enabled true (this enables the schedule and default is every 4 hours)
    a. Note: igls admin schedules command will not display QuotaInventoryCollection when it is not enabled
2. Then restart Eyeglass sca service following steps below in order for change to take effect
    a. SSH to Eyeglass appliance
    b. Type: sudo su - (to elevate to root - enter admin user password)
    c. Type: systemctl restart sca
    d. Type: systemctl status sca (to verify sca service active and running after the restart)
3. Use igls admin schedules to change the default from twice per day at hour 6 and hour 18 to an alternate schedule. **Example - to change to 10 minutes use this for testing only.**

a. igls admin schedules set --id QuotaInventoryCollection_2_5_3 --interval 10M

b. You should see result success.

4. You should now see the config job running with one of 2 options showing.

   a. The default config sync job for shares and exports will show -- without quotas

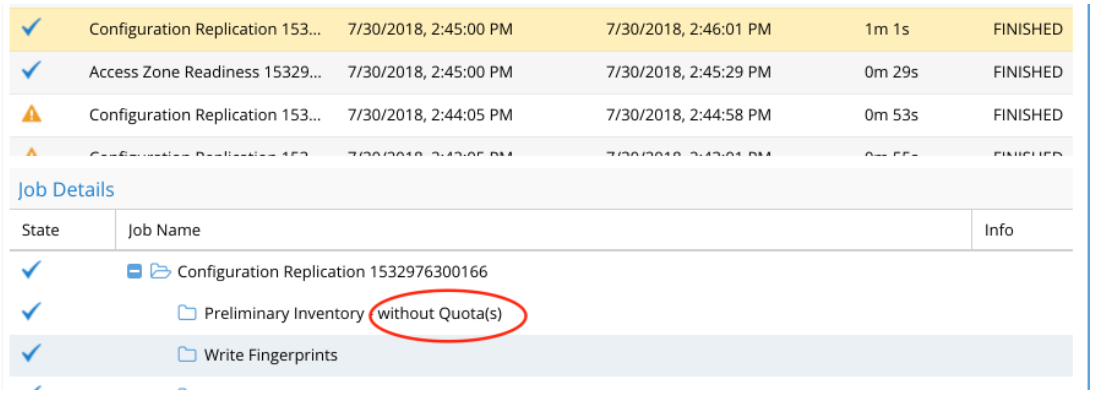

Advanced Quota Failover and Inventory collection CLI commands

Enable pre-sync of Quotas on Onefs 8.x clusters (DR licensed feature)

Quotas can be synced with a special configuration sync jobs to sync quotas at the same time shares and exports are synced.

NOTE: If pre-Sync is enabled, quota inventory must be enabled as well, follow steps above to enable the dedicated quota inventory collection job.

NOTE: pre-sync of quotas will impact synciq performance for replication of data, the more quotas on the target that exist, the slower synciq will replicate. Use with caution.

To check the pre-sync status

1. igls adv quotas

   Quota Advisory Sync Enabled: true

   Quota Advisory Sync Delete Mode: ENABLED

   Pre-Sync Quota On Interval: false

2. To enable pre sync of quotas requires the separate configuration job to be enabled as per above
   a. igls adv quotas set --quotapresync=true

1. Check that it is now set
   a. igls adv quotas

3. Verify that pre-sync shows enabled.

4. **Now enable quota dedicated quota collection job following the steps in the previous section. Pre sync is only supported using the dedicated quota collection inventory collection job.If you do not enable the separate quota collection pre sync will not function by design.**
   a. **Note the schedule to pre-sync quotas will be the same as the quota inventory collection schedule.**

Cluster Storage Monitor automation quota commands (Cluster Storage Monitor Feature)

This section covers auto advisory quota creation, and quota templates for AD managed quotas.

Note: Requires Storage cluster Monitor license

1. igls adv quotas help
2. Igls adv quotas (see current values).
3. Igls adv quotas set --quotasync=true (this enables the feature, false to disable).
4. Igls adv quotas set --quotasyncdelete=true **(defaults disabled, valid values are enabled/disabled/advanced).**

Active Directory Group based Quota Management (Cluster Storage Monitor Feature)

In order to use this AD group based quota management feature and new job type needs to be enabled , that runs on a default schedule of once per day. This group will evaluate the AD to group membership of users, and auto apply quota templates configured via the CLI.

High level requirements to use this Feature

1. New AD Group discovery job must be enabled first to retrieve AD groups and users (instructions below)
2. Separate quota inventory job must be enabled see here for instructions. (optional change schedule) and defaults to every 4 hours. **NOTE: This means quotas will be only collected if any changes every 4 hours.**

3. Create a Storage tier to label quotas (commands and examples below)
4. Create a template (commands and examples below)
5. Create ad groups, add users to groups
    a. **NOTE: Domain Users group cannot be used. A new AD group is required if the goal is a default domain wide quota. To add all domain users to a new AD group easily execute this command on a domain controller. Replace the object name for your domain group name.**
    b. **dsquery user -limit 0 | dsmod group "CN=newgroupname,CN=Users,DC=test ,DC=superna,DC=net" -addmbr**
6. Add AD groups to a SMB shares
7. Run the on demand AD group job to evaluate users groups and quota to create
8. To trouble shoot look at file on Eyeglass to debug
    a. cat /opt/superna/sca/logs/csm.log

NOTE: all references to AD domain should use uppercase characters

How to configure AD group quota scheduled job (Required)

This task will get AD group membership to evaluate which quotas should be applied. This is the task that monitors AD group changes to determine when to apply new quota or upgrade quota to a new tier.

1. Releases < 2.5.7

   a. igls admin schedules set --id ADGroupThresholds_2_5_5 --enabled true (this enables the schedule set to be every 2 hours).
      i. Note when ADGroupThresholds_2_5_5 is not enabled it will not be displayed using the igls admin schedules command
2. Releases > 2.5.7
   a. igls admin schedules set --id ADGroupThresholds_2_5_7 --enabled=false
3. Then restart Eyeglass sca service following steps below in order for change to take effect:
   a. SSH to Eyeglass appliance
   b. Type: sudo su - (to elevate to root - enter admin user password)
   c. Type: systemctl restart sca
   d. Type: systemctl status sca (to verify sca service active and running after the restart)
4. Verify the schedule is enabled:
   a. igls admin schedules
5. To change the schedule to check user to group AD membership more often than every 2 hours (the default schedule), make the following change with the IGLS command.
   a. Example for 10 minutes (this would be for testing only):
   b. igls admin schedules set --id ADGroupThresholds_2_5_7 --interval 10M

c. verify with "igls admin schedules"

How to start an onDemand AD Quota scan of AD and Quota creation

Use this command to start the scan job. This should be used for testing purposes. This will evaluate shares with template AD groups applied and determine of any quotas need to b created. Th monitor the evaluation tail -f /opt/superna/logs/csm.log

igls adv ADGroupThresholds

User quotas or group quota templates

Default mode is to create user quotas on templates. A parameter can be added to a template to change the quota type to be a group quota, and apply the group quota on a share where the template AD group has been applied to the permissions list. The command parameter example is shown below on the add template example using the [--quotasmode=[group|<default user>]] option on the create command.

igls csm tier help

Use this command to list storage tiers that have been created to group templates by tier.

list(default):
Lists all the tiers with their list of templates.

add: adds a new tier to the config file.
--id=<value>

delete: removes the tier from the config file.
--id=<value>

igls csm template help

Templates define a quota (hard, soft , accounting) and and AD group is assigned to a template. Templates are assigned to a tier (a label to group templates). The Tier must exist first before assigning a template to a tier.

list [--tier=<value>] [--name=<value>]: by default, lists all the templates from the config file and the tiers they belong to.

add: adds a new quota in the config file.
--tier=<value>
--name=<value>
[--hard=<value>]
[--hardunit=[PB|TB|GB|MB|KB|B|<default GB>]
[--soft=<value>]
[--softunit=[PB|TB|GB|MB|KB|B|<default GB>]
[--softgrace=<value>]
[--softgraceunit=[month|weeks|days|hours|minutes|<default hours>]
[--advisory=<value>]
[--advisoryunit=[PB|TB|GB|MB|KB|B|<default GB>]
[--quotasmode=[group|<default user>]]

update: updates an existing entry in the config file. (it takes the same arguments as "add")

delete: removes group from the config file.
--tier=<value>
--name=<value>

igls csm tier

Use this command to list the details of all tiers and the assigned template details.

igls csm template

Sample output shows tier assigned and AD group of the template.

Group quotas:

AD01\schema admins
tiers: silver

AD01\ai_testgroup
tiers: gold,silver

AD01\bronze
tiers: bronze

AD01\domain admins

tiers: gold

AD01\domain users
tiers: gold

List a specific template

igls csm template --name "AD01\testgroup"

List the template with the tier command to get details of the quota template

igls csm template --tier gold --name "AD01\gold"

List all templates in the gold Tier

igls csm template --tier gold

igls csm template add

Add new AD group template for user quota mode

Example:

igls csm template add --tier=bronze --name="AD01\bronze" --soft=200 --softunit=GB --softgrace=1 --softgraceunit=hours

Add new AD group template for group quota mode

Example:

igls csm template add --tier=bronze --name="AD01\silver" --soft=100 --softunit=GB --softgrace=1 --softgraceunit=hours **--quotasmode=group**

igls csm template update

Update an existing template

Example:

igls csm template update --tier=bronze --name="AD01\bronze" --soft=10 --softunit=GB --softgrace=1 --softgraceunit=days

igls csm template delete

Example:

igls csm template delete --tier=bronze --name="AD01\bronze" --quotasmode=user

Home Share AD Managed quota Configuration example

1. This guide will create 3 tiers bronze, silver and gold templates to offer 3 levels of quotas on the home folder.
2. Create a tiers
    a. igls csm tier add --id="bronze"
    b. igls csm tier add --id="silver"
    c. igls csm tier add --id="gold"
3. Create the templates and reference the tier name used above and assign a unique AD group that will be used to determine the users that will receive the quota.
    a. igls csm template add --tier=**bronze** --name=**"AD01\bronze"** --soft=200 --softunit=GB --softgrace=1 --softgraceunit=hours

b. igls csm template add --tier=**silver** --name=**"AD01\silver"** --soft=400 --softunit=GB --softgrace=1 --softgraceunit=hours

c. igls csm template add --tier=**gold** --name=**"AD01\gold"** --soft=800 --softunit=GB --softgrace=1 --softgraceunit=hours

4. Apply the Bronze , Silver and Gold AD groups (in this example the domain is AD01) to a share with full control (or read/write permissions) and move to the bottom of the share permission list.

   a. **NOTE: This group is not for assigning permissions to users, and is only used to indicate where quotas should be applied. This is why it should be moved to the end of the share list. Security groups should be higher on the share list.**

   b. This could be on a home directory path and will allow multiple templates to manage different tiers of quota limits for different users on the same path.

5. Run the quota scan job on demand to apply quotas and tail the CSM log to see what actions are taken

   a. tail -f /opt/superna/sca/logs/csm.log

   b. igls adv ADGroupThresholds

6. **NOTE: Make sure to enable the Quota Inventory quota job schedule to evaluate AD group membership of users and update quotas based on your configured templates. CLI steps are posted above in this guide.**

7. Done.

Group Share AD managed quota Configuration example

NOTE: The default mode for group quota templates (adgroupmode false), will create a group quota for any groups that are members of the template AD group, and found to be assigned to an SMB share. The Best Practice is not to use nested groups, and simply apply the template AD Group to SMB Shares where you want Group quotas created.

1. Create a tier: "igls csm tier add --id="bronze-group"
2. Create the template and name the tier: "igls csm template add --tier=bronze-group --name="AD01\bronze-group" --soft=200 --softunit=GB --softgrace=1 --softgraceunit=hours **--quotasmode=group** " (notice the group mode is set now)
3. Apply the AD group bronze-group. To a share with full control (or read/write permissions) and move to the bottom of the share permission list. **NOTE: This group is not for assigning permissions to users and is only used to indicate where quotas should be applied. This is why it should be moved to the end of the share list. Security groups should be higher on the share list.**

    a. This could be on a group share and will allow multiple templates to manage different tiers of group quotas on the same share path.

4. Run the quota scan job on demand to apply quotas and tail the CSM log to see what actions are taken

a. tail -f /opt/superna/sca/logs/csm.log

b. igls adv ADGroupThresholds

5. **NOTE: Make sure to enable the AD quota job schedule to evaluate AD group membership of users and update quotas based on your configured templates. CLI steps are posted above in this guide.**
6. Done.

igls adv adgroupmode

NOTE: Use this with caution, it will apply quotas based on user share access. Default is disabled. Do not change this setting unless directed by support.

Use this command to set the AD group mode used in quota templates. The **default is user mode** (which means this setting will show false). When enabled (shows true) it means the AD group named in the template will get the list of users and all the users groups, then all shares detected as accessible to each user listed in the groups (based on their AD groups assigned to shares) will have a user quota created on those shares based on the template definition. **NOTE: This can create a lot of quotas and limits the user on each share they have access to based on the template definition.**

igls adv adgroupmode (show current setting) Display the working mode for group quotas sync, user (default mode) or group mode. set --enabled=[true|false]

igls adv adgroupmode set --enabled=true (enables group mode)

igls adv adgroupmide set --enabled=false (disable group mode)

RPO Reporting CLI Commands

This section contains Eyeglass CLI commands related to the RPO Reporting feature.

igls adv runreports --report_type=rpo

Use this command to manually generate the SyncIQ Job Report and have it emailed. The time that the command is run is the starting time for the report and associated calculations. Each command example below are the type options.

igls adv runreports --report_type=rpo

igls adv skipscreenshots

Use this command to enable or disable RPO chart screenshots in the RPO Report.

To disable screenshots:

igls adv skipscreenshots set --skip=true

To enable screenshots:

igls adv skipscreenshots set --skip=false

CSM Reporting CLI Commands

This section contains Eyeglass CLI commands related to the CSM Reporting feature.

igls adv runreports --report_type=csm

Use this command to manually generate CSM Report and have it emailed.

igls adv runreports --report_type=csm

Advanced Commands Use if directed by support

HBASE Query Commands

igls hbase rowkeyscangenerator --cluster=<GUID> --path=<path> [--starttime=<DateTime>] [--endtime=<DateTime>] [--user=<user>] [--protocol={SMB | NFS}] [--operation={keys | data}] [--explain=<value>] [--dir=<dir>]

The mandatory parameters are the cluster GUID and the path information. Start and end time are optional and will be set automatically to the current day and one day back if missing. If start and end date-time information is provided, it needs to be done in the format dd-MM-yyyy HH:mm:ss. There are no default values for the user or protocol (the above shows the two options acceptable for protocol information).

The parameter operation identifies the type of request being asked for: keys represents the default request type and documents the list of start and end row keys based on the provided information. The resulting row keys will be documented into a file. Note that in this case, the addition of the optional parameter "explain" (the value is not

relevant), will break up the row key into its components and document the result as a table in the result file.

The optional parameter dir allows the caller to specify to where the result file is to be written (by default, this is /tmp). The generated file format name is ScanKeys_<CREATION_DATE>.log.

Memory watch dog on Eyeglass

get help

igls adv memorywatchdog help

show(default):
Allows to set and retrieve current values related to the forced garbage collection parameters.

set [--forcegc={true | false}] [--forcegcthreshold=<integer>] to set values or retrieve the current data.

get settings

igls adv memorywatchdog

Set memory threshold for memory watch dog to free up unused memory after crossing a threshold.

igls adv memorywatchdog set --forcegcthreshold=77 (sets GC to run over 77% memory used and writes debug log when this threshold has been crossed)

message sample

"Threshold has been reached, requesting GC to be executed if the GC is actually requested."

Database insertion validation

Default true, this removes orphaned records from db to avoid insertion errors. Do not use without direction from support.

igls adv verifydata [set] [--verify={true | false}]

AD User/group to SID or SID to user/group (Easy Auditor, Ransomware Defender)

These commands can be useful with Easy Auditor or Ransomware defender products to resolve user to SID, group to sid or the reverse.

1. User command
    a. igls adv resolve --user 'RNSM03\rwtest1' (NOTE: use single quotes around the user name and enter the domain in upper case)
    b. igls adv resolve --user "user@domain.com'

c. igls adv resolve --user S-1-5-21-51043000-931826463-941209176-1140

2. Group command

a. igls adv resolve --group 'RNSM03\rwtesting' (NOTE: use single quotes around the user name and enter the domain in upper case)

b. igls adv resolve --group SID:S-1-5-21-51043000-931826463-941209176-1112 (must add SID: before the sid)

© Superna LLC

1.18. Eyeglass Appliance Time Synchronization Best Practice

Eyeglass Appliance Time Synchronization Best Practice

The appliance is a VMware appliance which allows the ESX host and vCenter NTP time source to flow through to the appliance. In most environments, the virtual environment NTP source is distributed to each host. Each VM on that host receives a good time source via BIOS time set. The PowerScale cluster should be set to the same time source as the VMware environment to ensure time of events for DR logs are synced to the same source. If the cluster is using a different time source then the procedure below can be used to add NTP to the Eyeglass appliance.

1. SSH to appliance as admin then sudo -s to root .
2. type **'yast'** .
3. Select **Network Services** and then **NTP Configuration** .

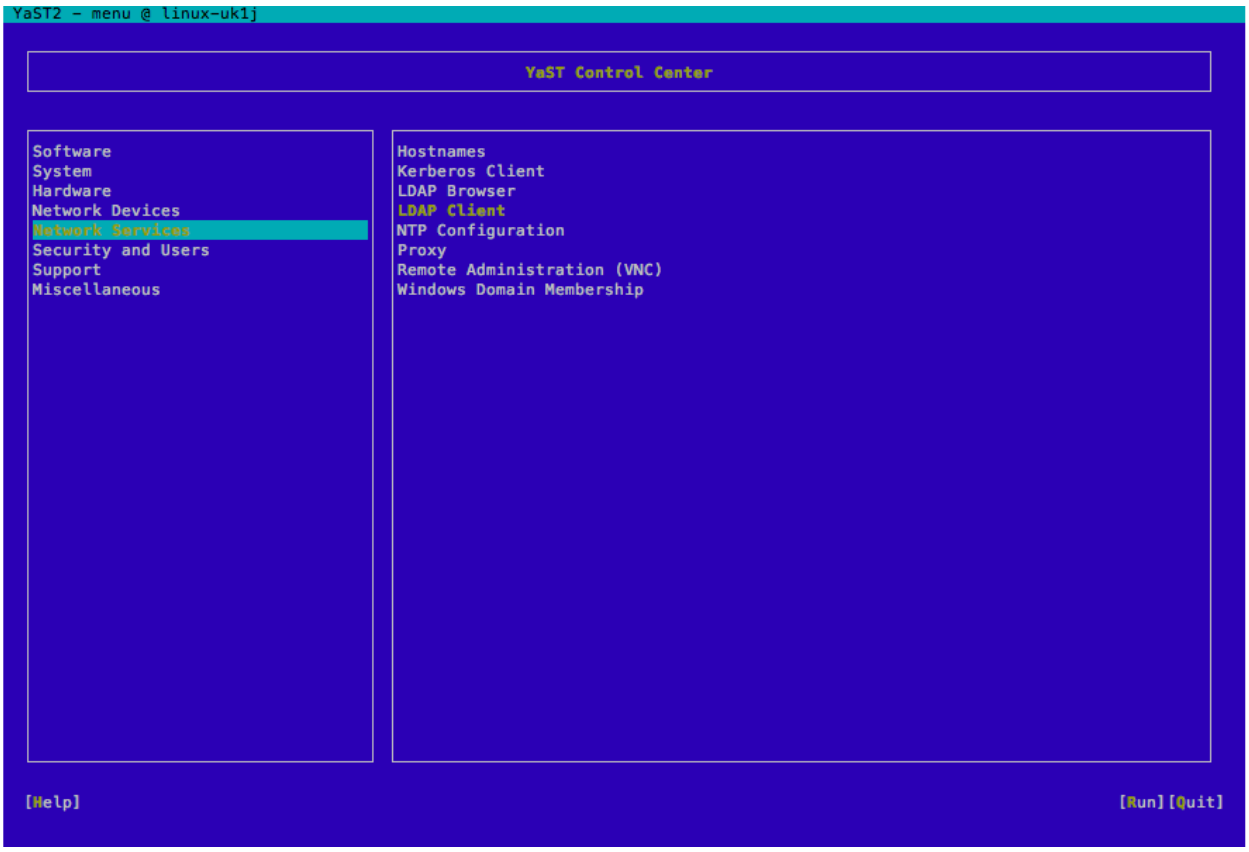

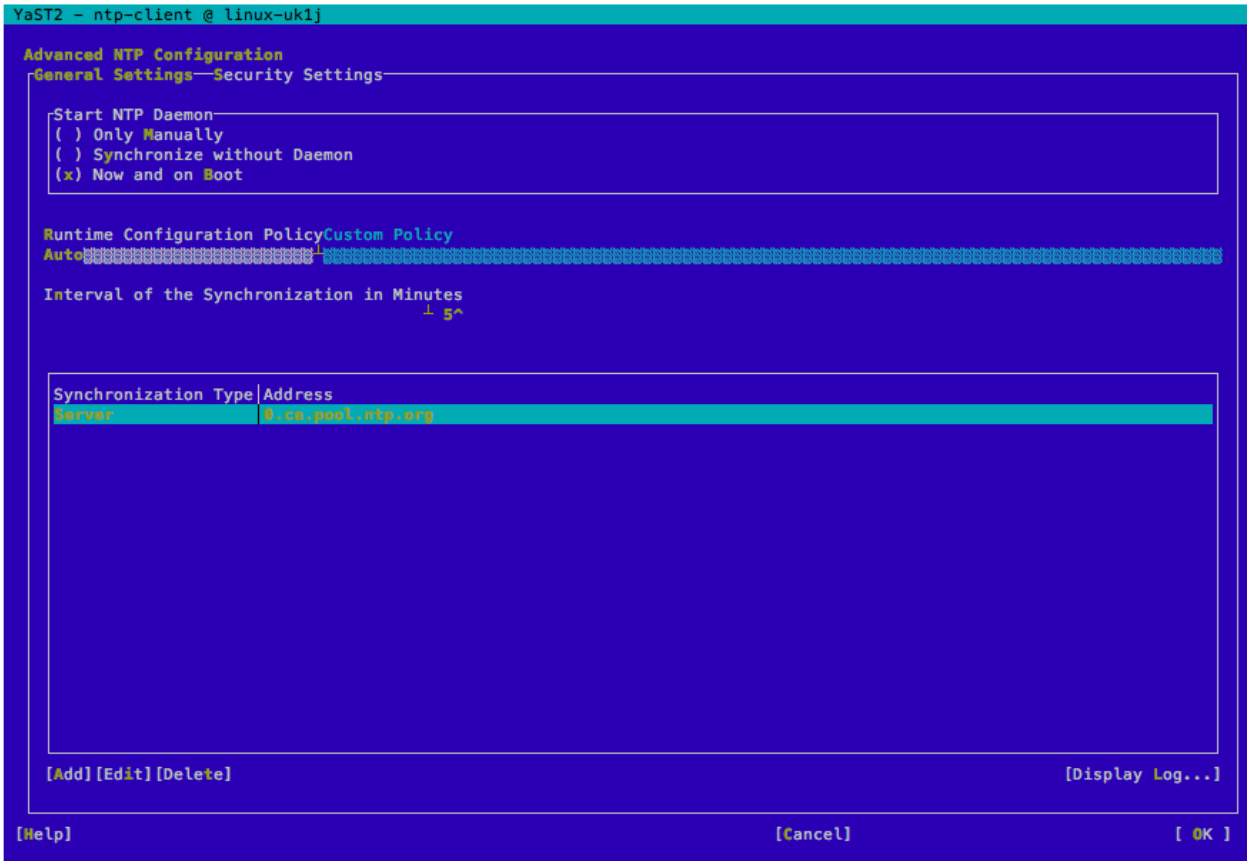

4. Using tab to **Start NTP Daemon** and select **Now and on Boot**.

5. Tab to **Add** to add a server source.

6. Select **Add Server** using tab to add a new server.

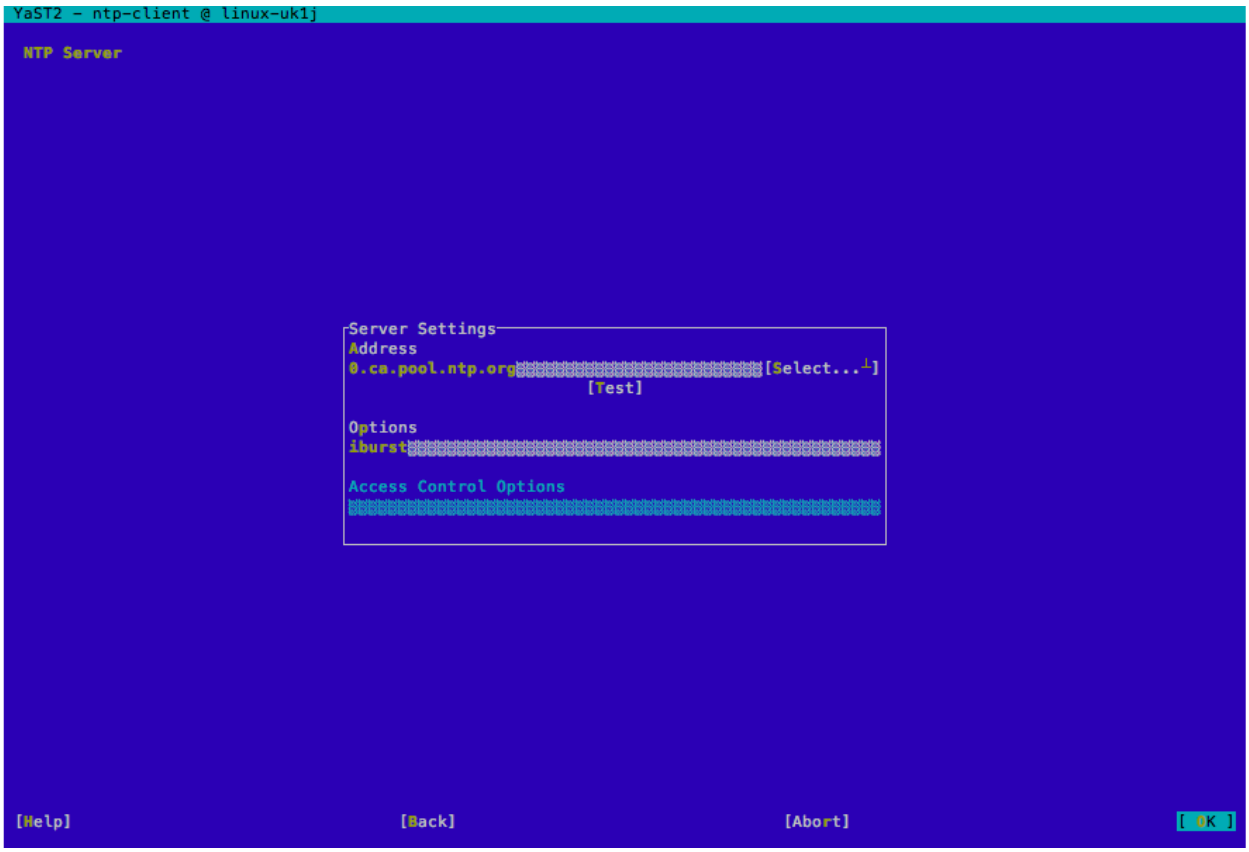

7. In the add server dialog box enter **ip or FQDN or hostname** (**Note:** requires DNS to be setup correctly on the appliance) of the NTP source.

8. Select the **Test** button to verify reachability and protocol connection for time responds correctly.

9. Select **OK** using tab

10. Done.

© Superna LLC

1.19. Eyeglass Automatic Updates for Recommended Packages

Eyeglass Automatic Updates for Recommended Packages

Eyeglass automatic recommended Suse updates are applied weekly only if the appliance has a Internet connection. Manual RPM updates are required for OS updates for customers without Internet access. The RPM's will need to be retrieved by customers from Internet hosted Repositories.

Changing repository URL on the appliance is possible for a customer that rehost OS patches internally. Use Zypper AR URL "name of repo" to add a repo or, consult openSuse documentation.

The appliance is setup by default to check weekly for recommended packages and automatically apply them. To change this setting:

1. ssh to the Eyeglass appliance.
2. Login as admin and sudo su to root or login as root.
3. Type '**yast**'. The YaST2 menu opens with Software selected by default.

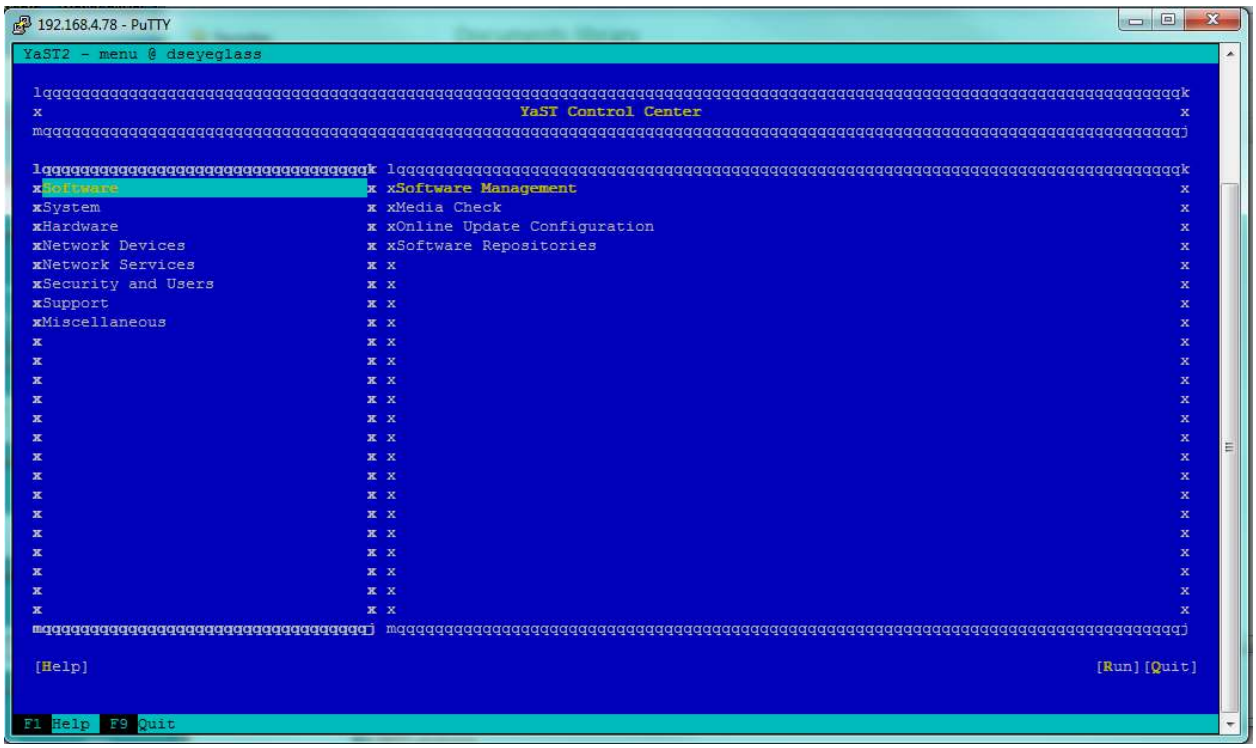

4. Use the right arrow key to move to the right hand menu and then the down arrow to highlight **Online Update Configuration**.

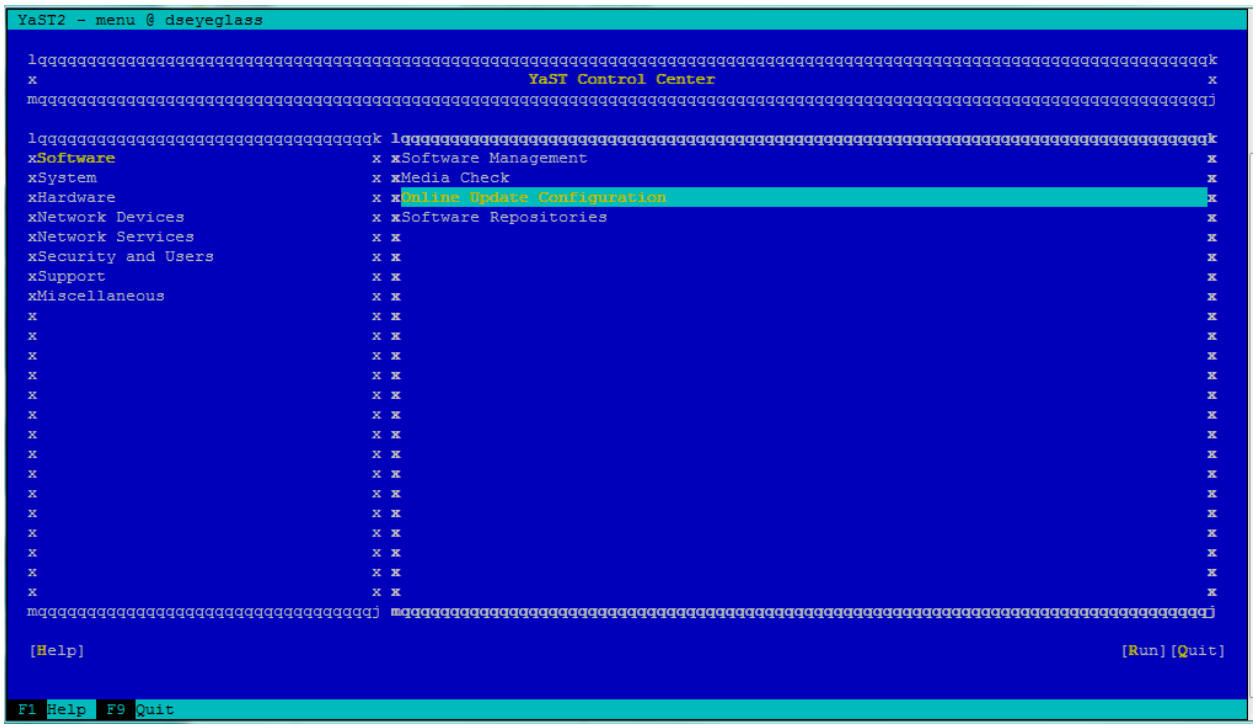

5. Hit **Enter** to select Online Update Configuration. The Online Update Configuration window opens.

6. On the screen below leave skip interactive patches enabled, to ensure required RPM versions Eyeglass requires are not overwritten. If changing this value contact support first to get instructions. **NOTE: Always create a VMware snapshot before updating the OS to avoid issues that require a rollback option. NOTE: license keys cannot be reset if a Operating System is corrupted , support cannot reset license keys. If unsure contact support first before unchecking this option.**

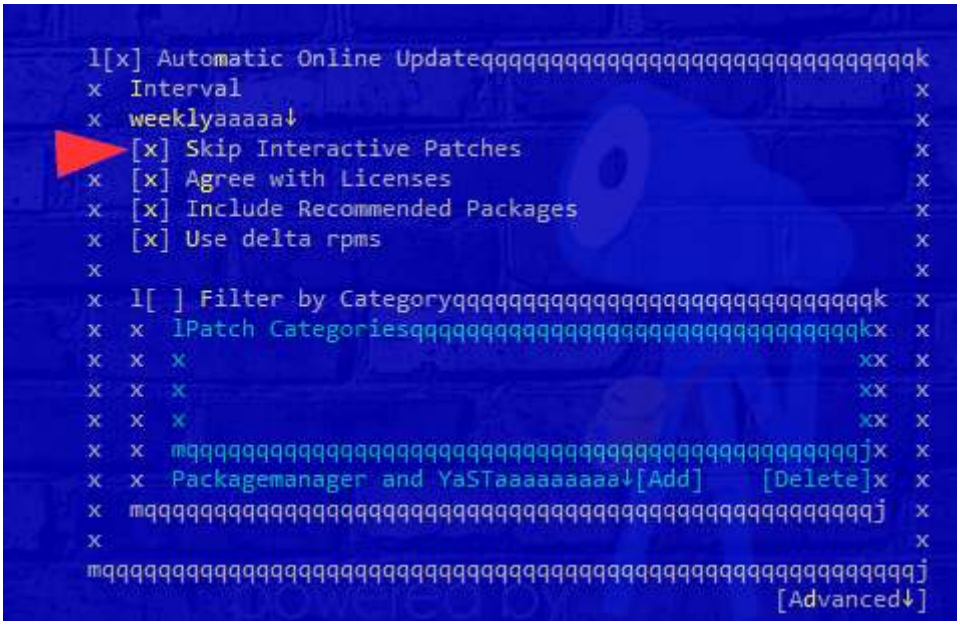

7. To **change the Interval**, Tab until the Interval is highlighted. Then use the arrow key to see Interval options. Use the arrow key again to highlight the interval you would like and Enter to select.

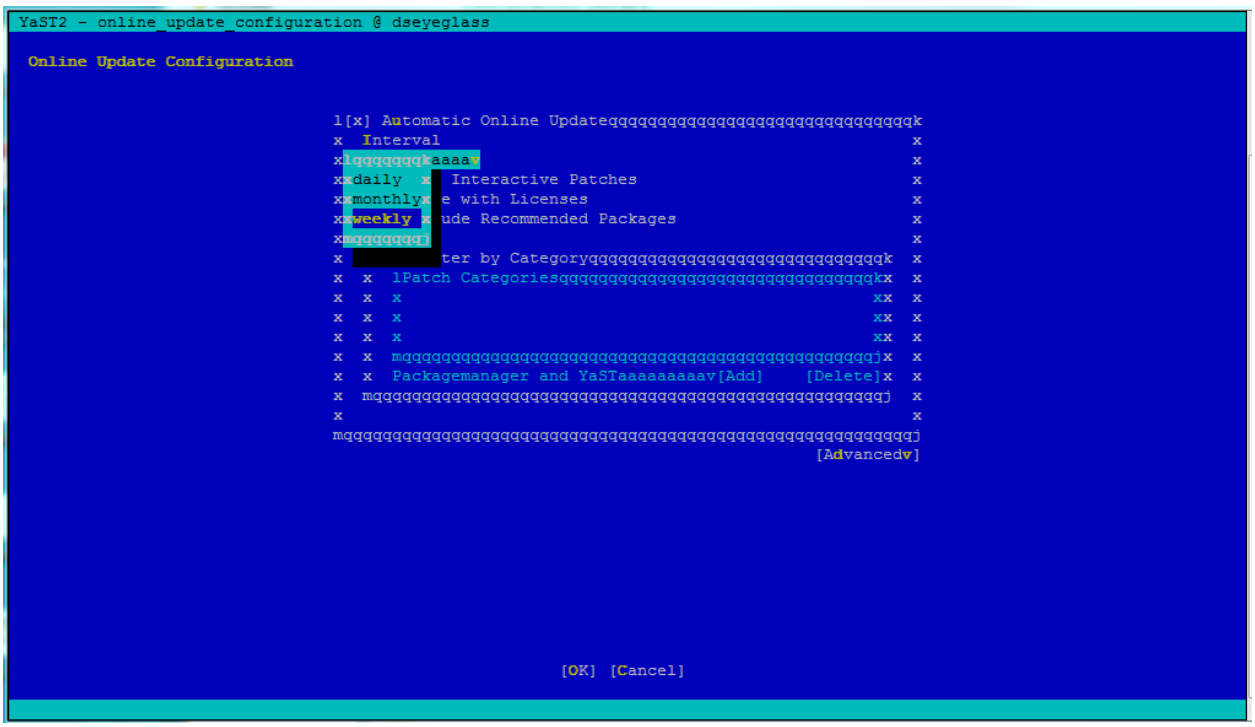

8. To enable/disable Automatic Updates, Tab until Automatic Online Update is highlighted. Then Enter to select/deselect this option.

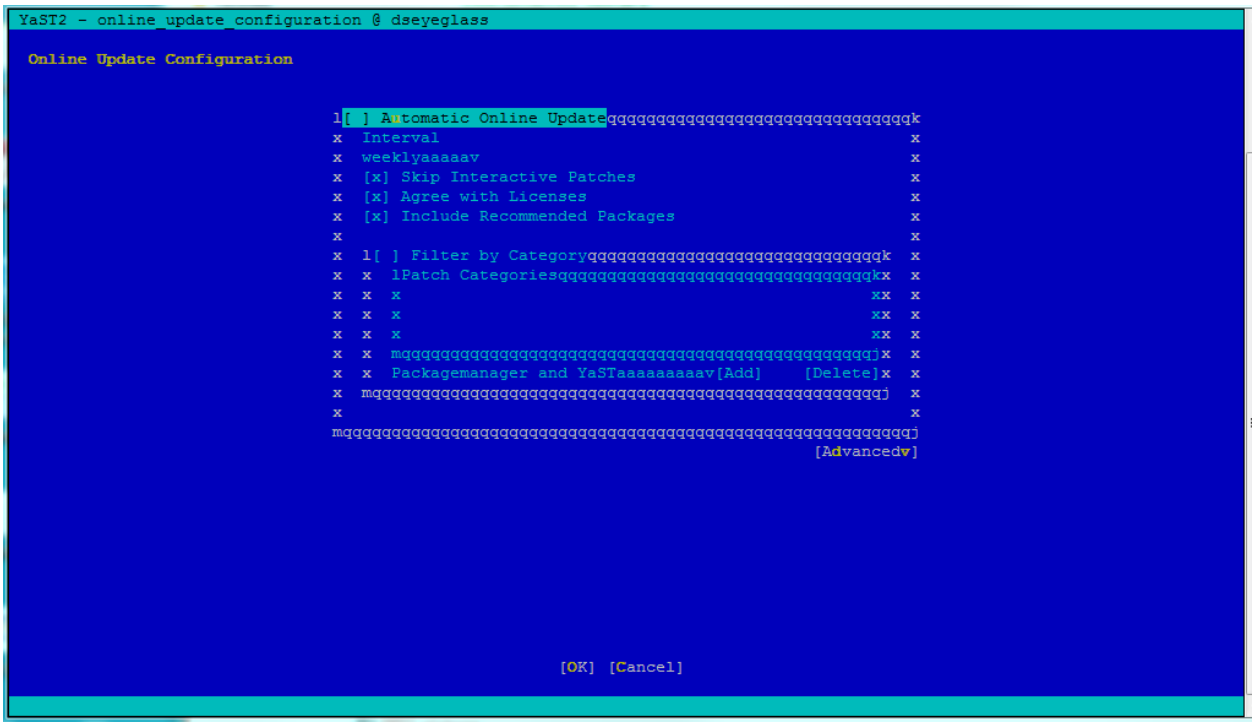

9. Tab to OK and Enter to save your changes.

© Superna LLC

1.20. Update Eyeglass Appliance Network Settings

Home Top

Update Eyeglass Appliance Network Settings

1. ssh to the Eyeglass appliance (if access this way is available) or, console to the Eyeglass Appliance Virtual Machine from vSphere.
2. Login as admin and sudo su to root or login as root.
3. Type '**yast**'. The YaST2 menu opens with Software selected by default.
4. Use the down arrow key and navigate to **Network Devices**. Then use the right arrow key to move to the right hand menu followed by the down arrow key to navigate to **Network Settings**.
5. Use the **Enter** key to select Network Settings.

To Change the Eyeglass Appliance IP Address:

1. Start at the **Network Settings** window.
2. Use the Tab key to highlight the **Edit option on the Network Settings window** and then **Enter**.
3. In the **Network Card Setup** window, use the Tab key to navigate to the field you want to change and make the required update.
4. Once you have made all of the required changes, use the Tab key to navigate to the **Next Option and Enter**. This will return you to the Network Settings window.
5. If no further updates are required use the Tab key to **navigate to OK and Enter** to save your changes. If further updates are required, follow the steps in the appropriate section.

To Change the Eyeglass Appliance DNS Settings:

1. Start at the **Network Settings** window.

2. Use the right arrow key to highlight the **Hostname/DNS option**.
3. Use the Tab key to navigate to the field that needs to be updated and make the required change.
4. Then **Tab to OK** to complete.

To Change the Eyeglass Appliance Routing Settings:

1. Start at the Network Settings window.
2. Use the right arrow key to highlight the **Routing Option**.
3. Use the Tab key to navigate to the field that needs to be updated and make the required change.
4. Then **Tab to OK** to complete.

Eyeglass Root Password

If it is required to have the root password to the appliance follow the procedure below.

1. Login as admin using ssh.
2. Then execute 'sudo -s'.
3. Then '**passwd**'.
4. Enter new and re-type new password.

© Superna LLC

1.21. Appliance Security Updates and Eyeglass Updates with HTTP Proxy

Appliance Security Updates and Eyeglass Updates with HTTP Proxy

Many customers use proxy access to the Internet, and require configuration of HTTP proxy on devices that need access to the Internet on port 80.

Eyeglass updates are hosted for online updates, and Operating System updates are also reachable with HTTP for security and package adds with zypper.

Steps to configure HTTP proxy for Eyeglass:

1. Login via SSH and admin password.
2. sudo -s (switch to root).
3. yast (run yast).
4. Navigate to the screen shots below to enter proxy configuration details:
    a. URL.
    b. User id and password.
    c. NOTE: You can also encode the user and password into the URL and may be required for some proxy devices
        i. **example url http://username:password@proxyaddress:port**

Note: This is untested with various proxy software solutions and not supported under maintenance contracts.

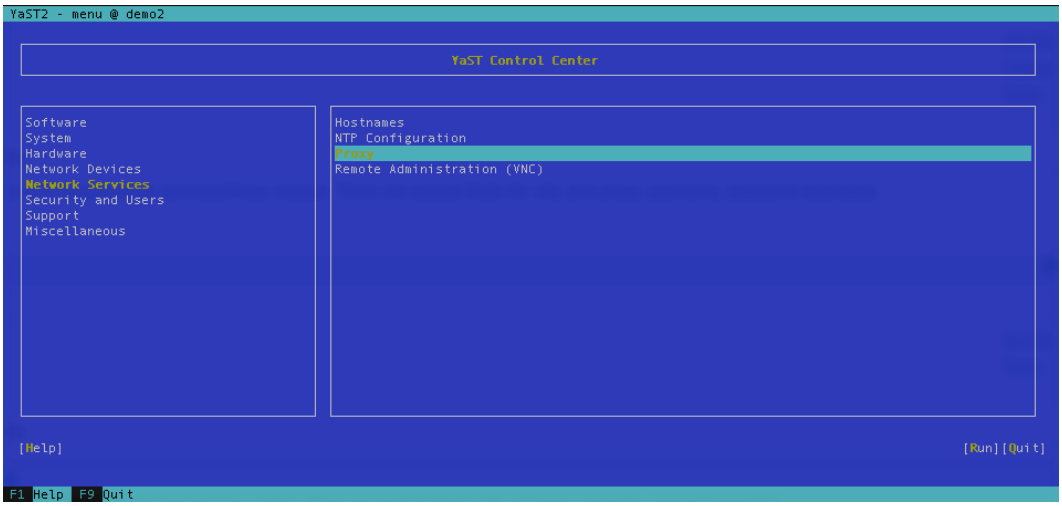

© Superna LLC

1.22. Diagnostic Tools for Dark Site Support

Diagnostic Tools for Dark Site Support

This feature allows diagnostic tool to process logs on the appliance and summarize errors from config sync, time to sync, failover jobs. This tool is the same log analysis tool used by support when logs are uploaded. This tool anonymizes the data in the logs and provides summary analysis of errors, alarms, and DR readiness.

Customers can schedule training with support to read the output when opening a case.

To execute the log analysis:

1. SSH to appliance as admin user.
2. Run command: **igls app report**
3. Wait for the report to complete.
4. See logs report on: https://<eyeglass IP address>/report/ .
5. Alternate location is /srv/www/htdocs/report with file name index.html can be downloaded via scp to review and open in a browser.

© Superna LLC

1.23. Role Based Access Control And Authentication Guide

Home Top

- Overview
- RBAC Requirements - Read Me First
- RBAC Quick Start Steps
- Simple Setup AD Group based RBAC
- How to create a new Eyeglass Role
- How to Login with RBAC

© Superna LLC

1.23.1. Overview

Home Top

- What's New
- Key Features:
- Built in Roles and user accounts:
- Use Cases for Custom Roles:
- Local Eyeglass OS user accounts:

What's New

1. Release 2.5.6 20258 or later has simplified adding users and groups to roles by validating the SID or GID in the GUI before saving the role. Using the SID and GID now allows AD group names to include spaced or some special characters.
    a. The user name to login can use any syntax with any case of the login without any special requirements example domain\username or username@domain can be used without regard for the AD UPN value of the account.
    b. The upgrade to this release will convert the RBAC file to include SID and GUID automatically without any user actions to switch to this new role mapping implementation.

Key Features:

Role Based Access Control (RBAC) for Eyeglass allows any role combination to be created based on Eyeglass desktop icons. Custom

roles can be created to meet any requirement of access to Eyeglass features:

1. Default admin user has all privileges
2. Default read-only role can see all icons
3. Create roles and assign icons of functionality
4. Map to user or group in AD, local Isilon/PowerScale users and groups.
5. All authentication is done through PowerScale API to an authentication provider, and SMB AD password validation to access zone SmartConnect FQDN's

Built in Roles and user accounts:

Eyeglass ships with built in roles and users as follows:

1. admin
    a. Has all permissions for all products
    b. Default password 3y3gl4ss
2. rwdefend
    a. Assigned the builtin role Ransomware Defender with ability to manage and monitor Ransomware Defender product
    b. Default password 3y3gl4ss

c. 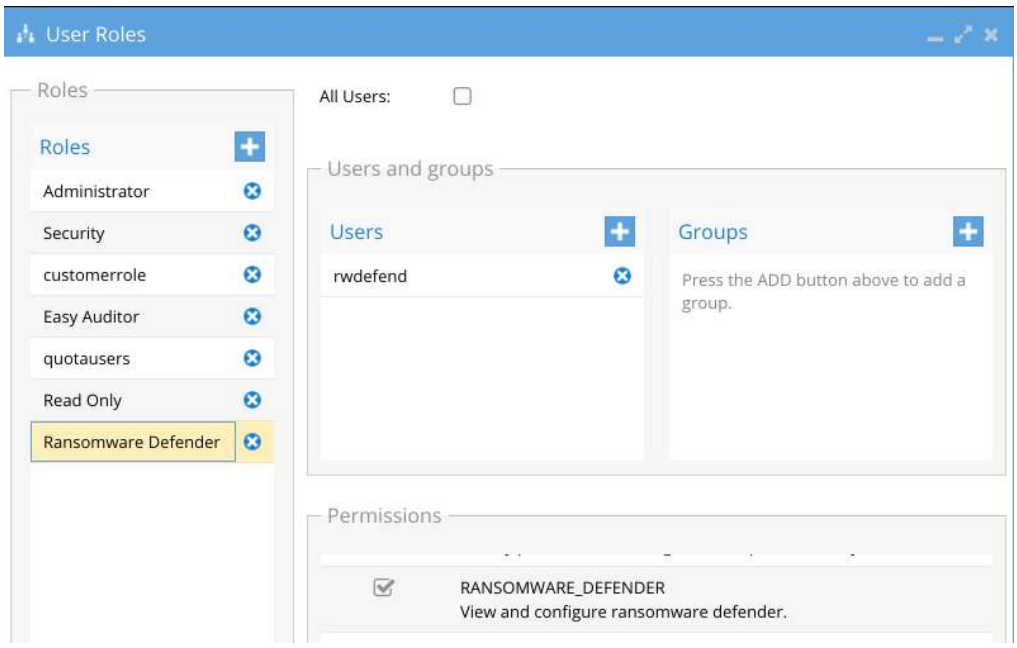

3. auditor

   a. User has read and modify permissions within the Easy Auditor application

   b. Default password 3y3gl4ss

   c. Assigned the Auditor builtin group role

   d. 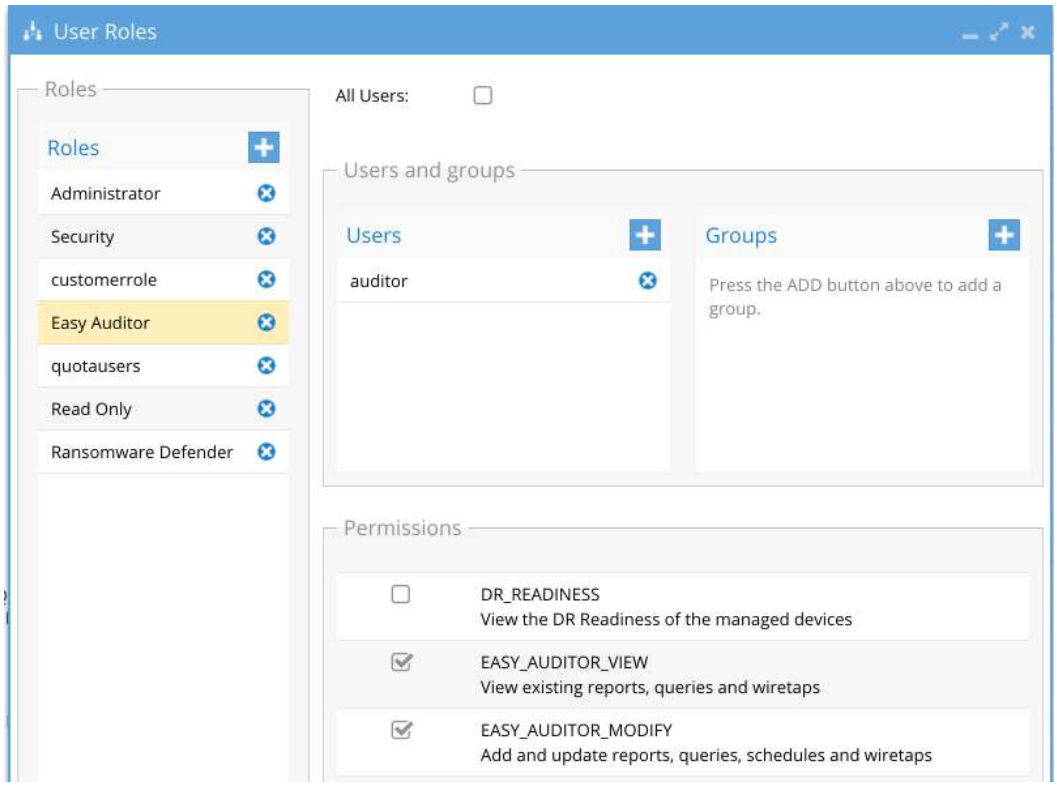

e. This also includes the manage remote services icon to see Eyeglass clustered agent status

f. 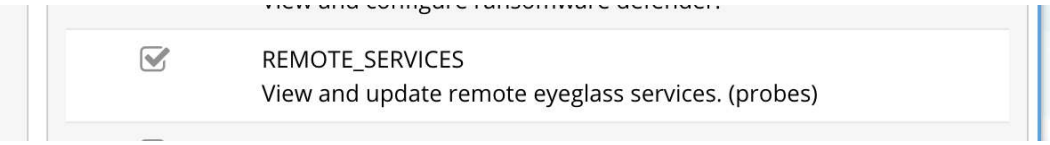

Use Cases for Custom Roles:

1. Monitoring only -readonly role
2. Departmental login for DR readiness view
3. Security for Ransomware monitoring (Ransomware Defender)
4. Unlock my files Help desk (Cluster Storage Monitor license)
    a. Use this permission to create a Help desk role for unlock my files portal

b. 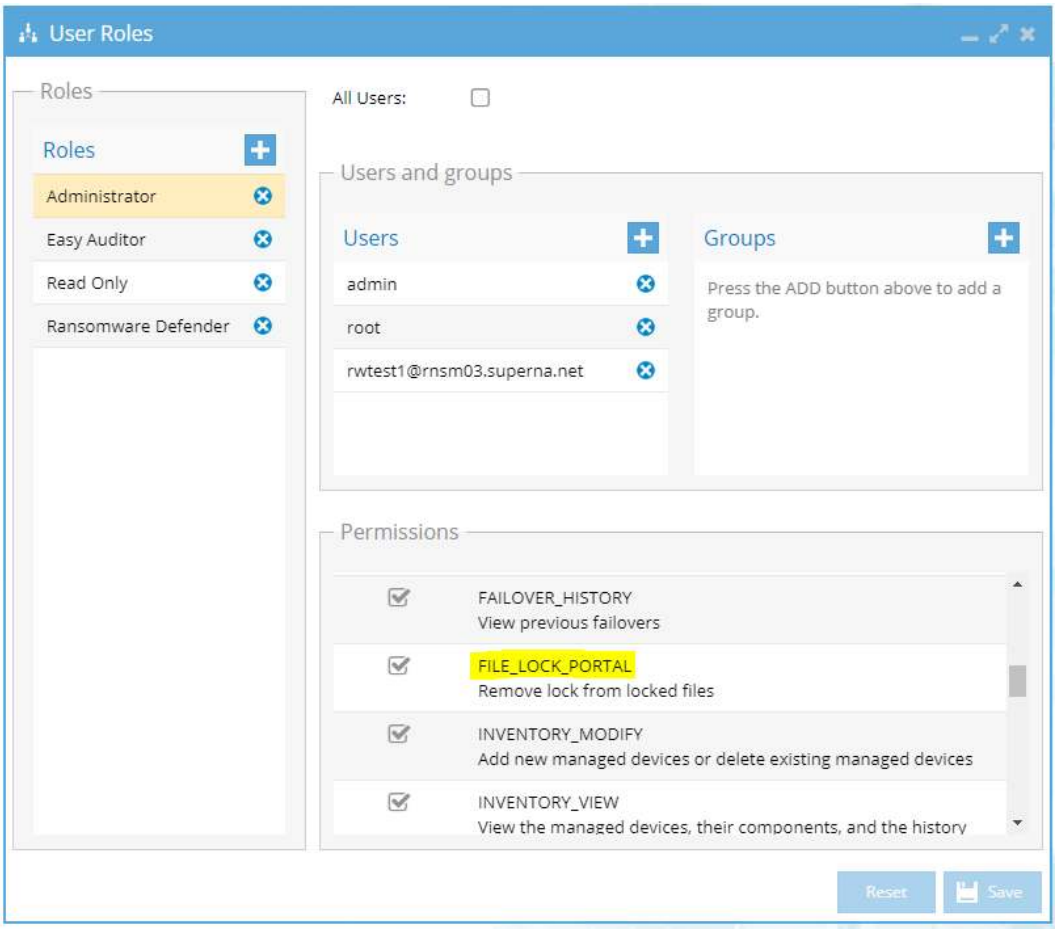

5. Auditing for file audit (Easy Auditor)
6. Failover only administration functions (i.e can not add new clusters)
7. Logging and monitoring only
8. Storage monitoring only (no DR functions)
9. Cluster reporting only (no DR functions or storage monitoring)
10. Centralized security to match PowerScale Role based Access groups to include DR functions with SyncIQ

Eyeglass Authentication Options

Local Eyeglass OS or Proxy Login are the two types of login Eyeglass supports. Local Login uses a user account created in the OS. Proxy Login options use PowerScale as the authentication provider, and proxies the user id and password to Ision for authentication validation and group membership. Custom roles can be created in Eyeglass that can use local or proxy login for access control.

The following sections describe how to configure and use Local Eyeglass OS or Proxy Login.

Local Eyeglass OS user accounts:

When entering local users we recommend the syntax UNIX_USERS to represent the domain. Example - UNIX_USERS\read (this is a convention to ensure its easy to know where this user will exist for authentication)

How to create new local user on the appliance:

1. Ssh admin@x.x.x.x
2. Sudo -s (enter admin pwd)
3. useradd <user name>
4. passwd <user name> (to set a password)

© Superna LLC

1.23.2. RBAC Requirements - Read Me First

Home Top

- Overview
- New in 2.5.6 20258 or later Releases
- General Requirements:
- Use Case: Applying AD Groups to an Eyeglass Role - Requirements below
  - AD group and AD domain Syntax Rules 2.5.6 20258 or later
- Use Case: Applying a user directly to an Eyeglass Role - Requirements below

Overview

Follow the guidelines in this section when adding groups or users to roles.

New in 2.5.6 20258 or later Releases

1. Adding users or groups will now validate the user or group can be resolved when saving the role, if the user or group can not be resolved to a SID or GUI from Active Directory an error will be displayed and the user or group will not be saved to the role.

1. General Requirements:

a. The proxy authentication requires the **system zone** to have an AD authentication provider added to allow for the password to be validated and AD group membership retrieved from AD.

b. SMB protocol port 445 open between Eyeglass VM and the cluster

c. **Trusted Domains** - AD Domain that is not directly added to PowerScale as an authentication provider can be used when adding users or groups . The trusted domains must trust the AD domain added to the system zone.

d. SMB2 protocol for AD authentication of users with an SMB share in the system zone.

e. System Zone authentication is the only supported proxy login and requires an AD provider in the system zone.

f. **SMB protocol must be enabled in system zone**.

   i. Login will attempt to validate password on all clusters added to Eyeglass using SMB and system zone authentication requests over SMB.

2. Use Case: Applying AD Groups to an Eyeglass Role - Requirements below

a. Verify the user account you are logging has AD groups displayed from PowerScale CLI. If no groups are shown, open an EMC SR.

   i. isi auth users view dfs1@ad1.test --show-groups **(The output MUST show Additional Groups: assigned to the**

user, the group used in the Eyeglass role must be listed in the output from this command.

b. AD group and AD domain Syntax Rules 2.5.6 20258 or later

   i. Upgrade to 2.5.6 2058 or later

   ii. Group name **cannot** have special characters other than dash or underscore or space

3. Use Case: Applying a user directly to an Eyeglass Role - Requirements below

a. The user you add must be the UPN format (user principal name) **user@domain name.**

   i. A user added to a role must use the UPN defined in AD (see screenshot example). You can verify the UPN with this Isilon CLI command **isi auth users view dfs1@ad1.test --show-groups**

ii. 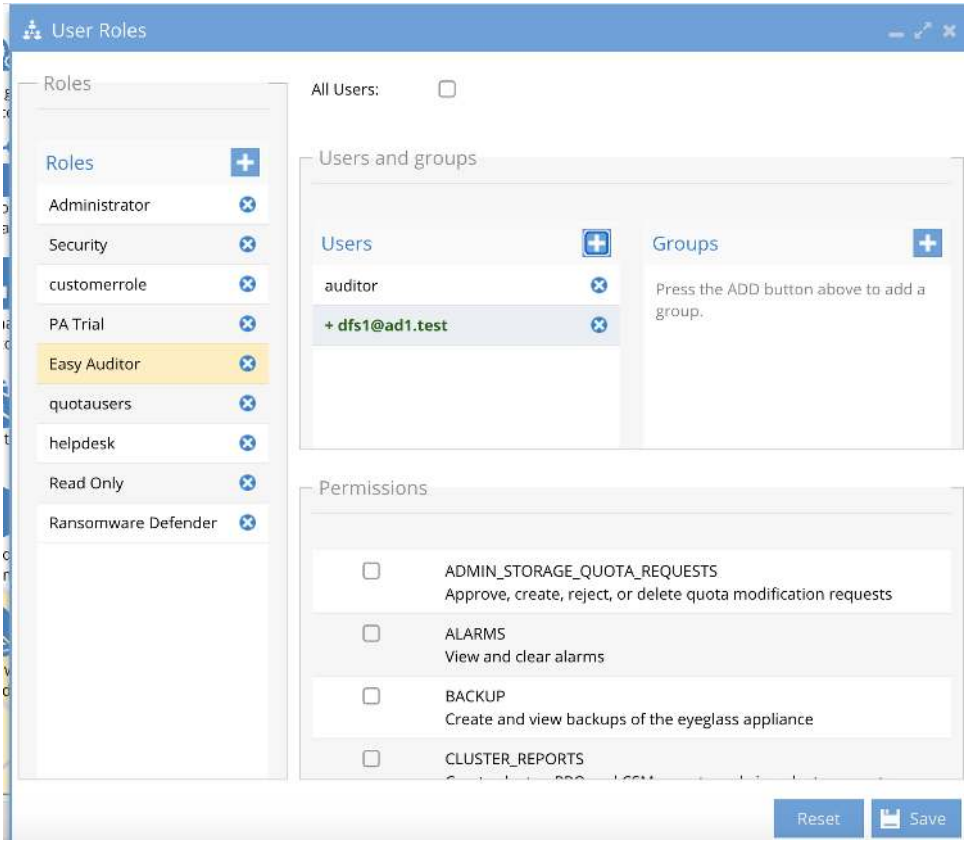

© Superna LLC

1.23.3. RBAC Quick Start Steps

Home Top

Follow these steps in order below to get role based access configured.

1. Review Active Directory Groups and user role requirements before configuring.
    a. Follow the Quick AD Group Solution Example
2. Create a new Role
3. Login and Authentication Steps to test role based permissions

© Superna LLC

1.23.4. Simple Setup AD Group based RBAC

Home Top

- Before You Start
- Configuration Steps

Before You Start

Read AD Group Name Requirements before starting Click Here.

Configuration Steps

1. Steps should be followed exactly
    a. **NOTE: In this example the domain name is AD02 and must be upper case, follow screenshots as a reference. We suggest using our exact group names to create your first RBAC role.**
2. Create an ad group named **superna-admins** (make sure it is all lower case, create as a global security group)

3. 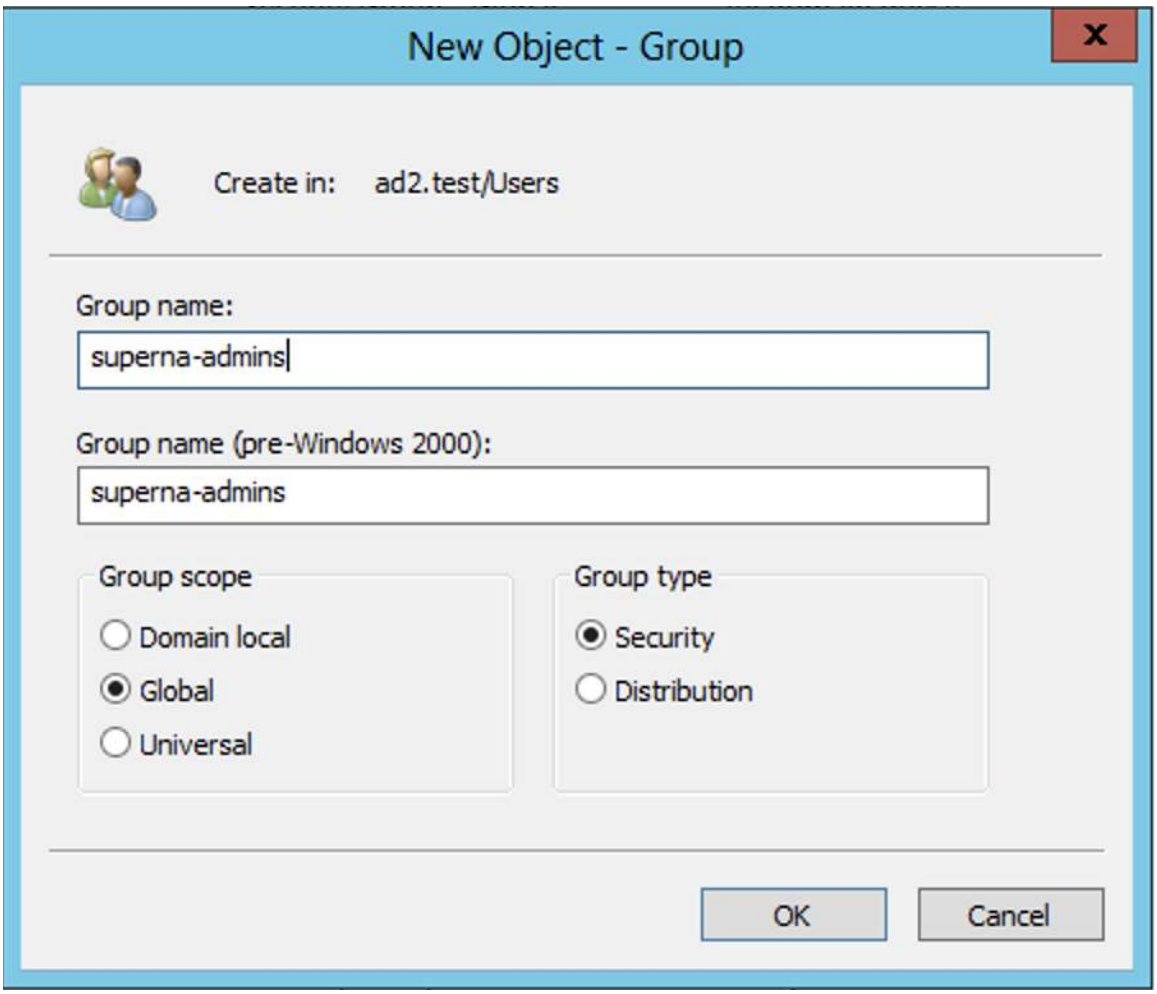

4. Add your user account to the group using Users and Computers Snapin console.
5. Create role in the Eyeglass User Roles Icon after loggin into Eyeglass as admin user.

6. 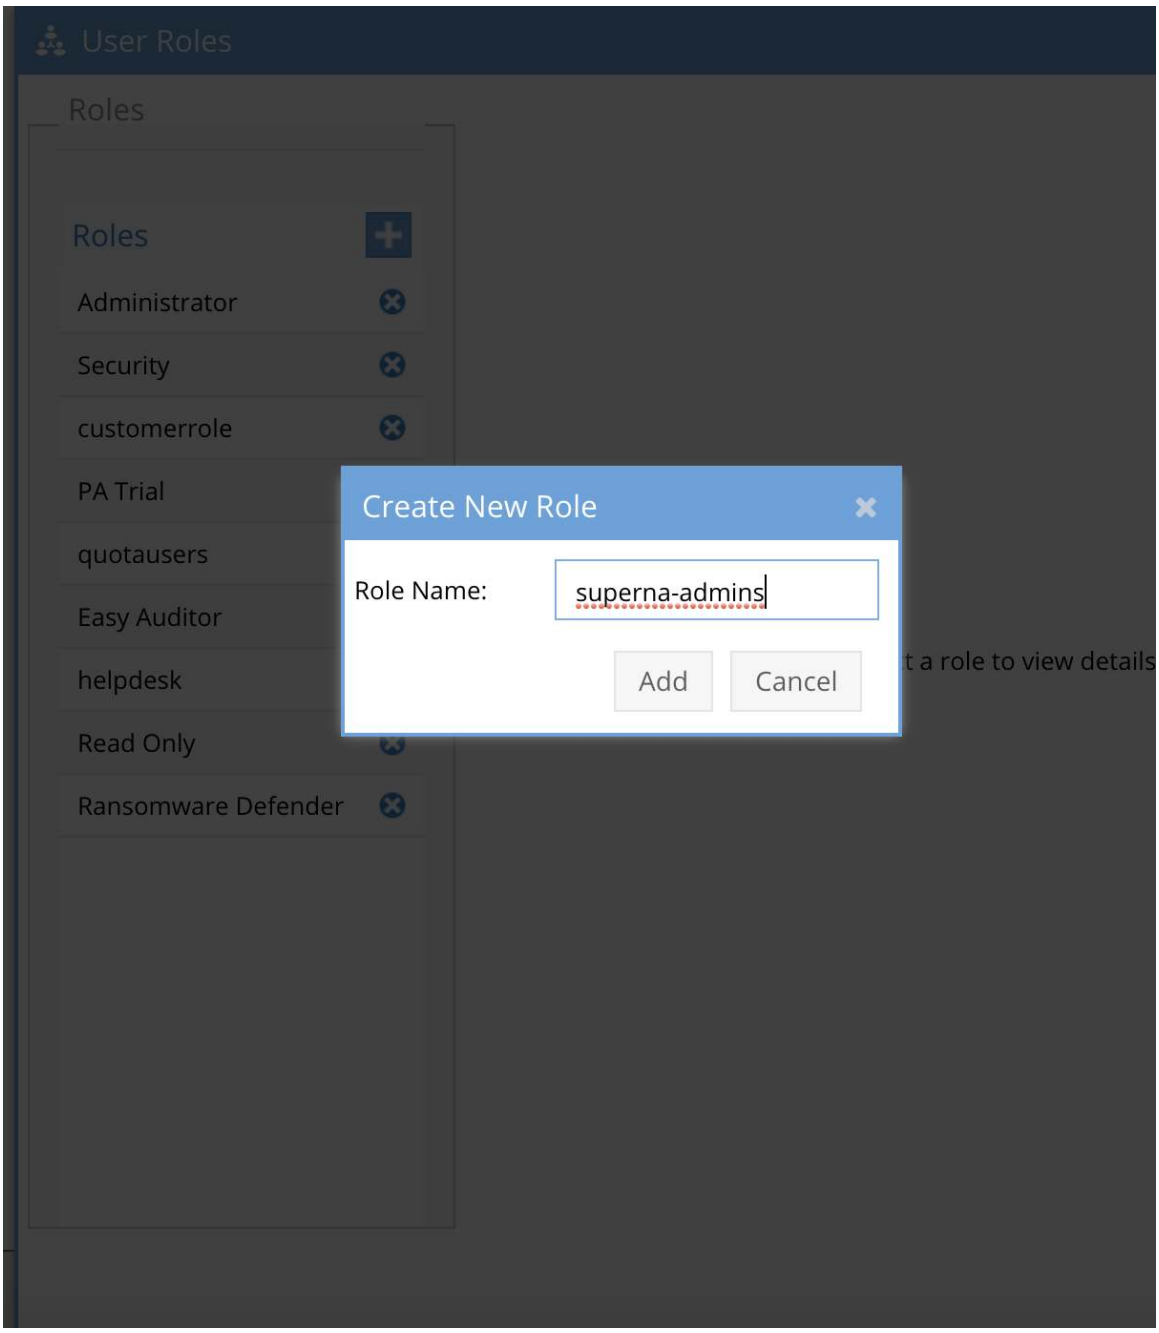

7. Add the AD group with **Upper case domain name (mandatory upper case)** and **lower case group name (mandatory lower case)**, see the diagram below as a reference.

8. 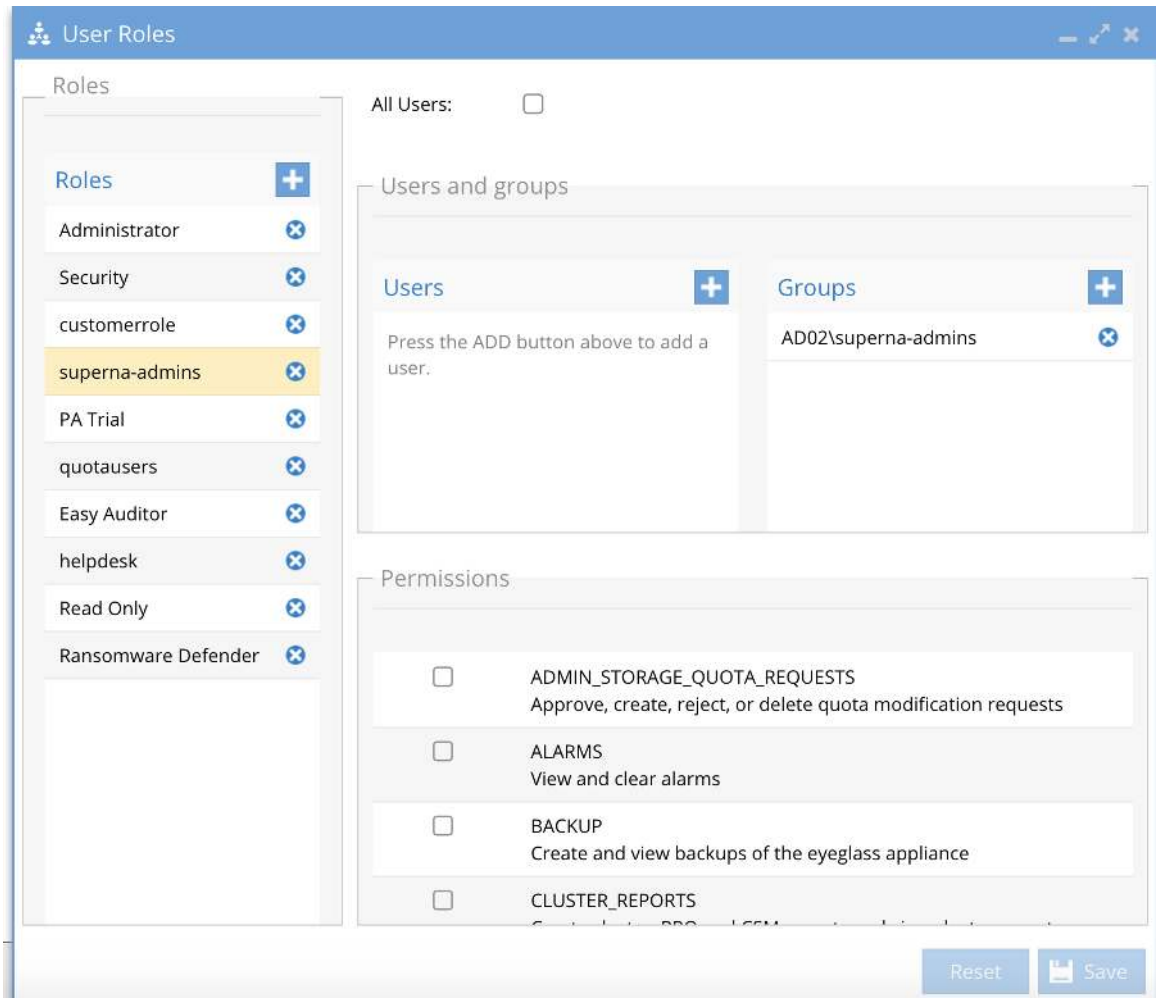

9. Select check boxes for the roles functions or icons that should be assigned to the role, **make sure to click save**.

10. Verify your cluster is ready for RBAC

    a. The user view command will list AD group.

    b. Run this command on your cluster **isi auth users view --user=user@domain --show-groups**

    c. The AD group created above must be listed in the output of the isi command.

        i. **WARNING: If AD the group is not listed in the output the RBAC role will not work.**

d. The ISI command above will also include the **DNS domain** property of the user and the **SAM Account Name**

   i. The login name will be **<sam account name>@<DNS Domain name>**

   ii. Example below:

      1. DNS Domain name = AD2.TEST
      2. SAM Account Name = demo1
      3. The additional groups must show the AD group **AD02\superna-admins**.
      4. The user name to enter to the proxy login would be **demo1@ad2.test**

e.

```
Isilon OneFS v8.2.1.0
prod8-1# isi auth users view --user=demo1@ad2.test --show-groups
                    Name: AD02\demo1
                      DN: CN=demo1,OU=demo systems,DC=ad2,DC=test
              DNS Domain: ad2.test
                  Domain: AD02
                Provider: lsa-activedirectory-provider:AD2.TEST
        Sam Account Name: demo1
                     UID: 1000482
                     SID: S-1-5-21-201832566-3187353407-2829991710-1834
                 Enabled: Yes
                 Expired: No
                  Expiry: -
                  Locked: No
                   Email: -
                   GECOS: demo1
           Generated GID: Yes
           Generated UID: Yes
           Generated UPN: No
           Primary Group
                      ID: GID:1000140
                    Name: AD02\domain users
        Home Directory: /ifs/home/AD02/demo1
       Max Password Age: -
       Password Expired: No
        Password Expiry: -
      Password Last Set: 2020-08-02T11:47:37
       Password Expires: No
                  Shell: /bin/zsh
                    UPN: demo1@AD2.TEST
User Can Change Password: Yes
     Additional Groups: AD02\superna-admins
                        AD02\bronze
prod8-1#
```

11. Login with **<sam account name>@<DNS Domain name>**

12. Follow the How to Login guide located here.

© Superna LLC

1.23.5. How to create a new Eyeglass Role

Home Top

How to create a new Eyeglass Role

1. Eyeglass Main menu select User Roles

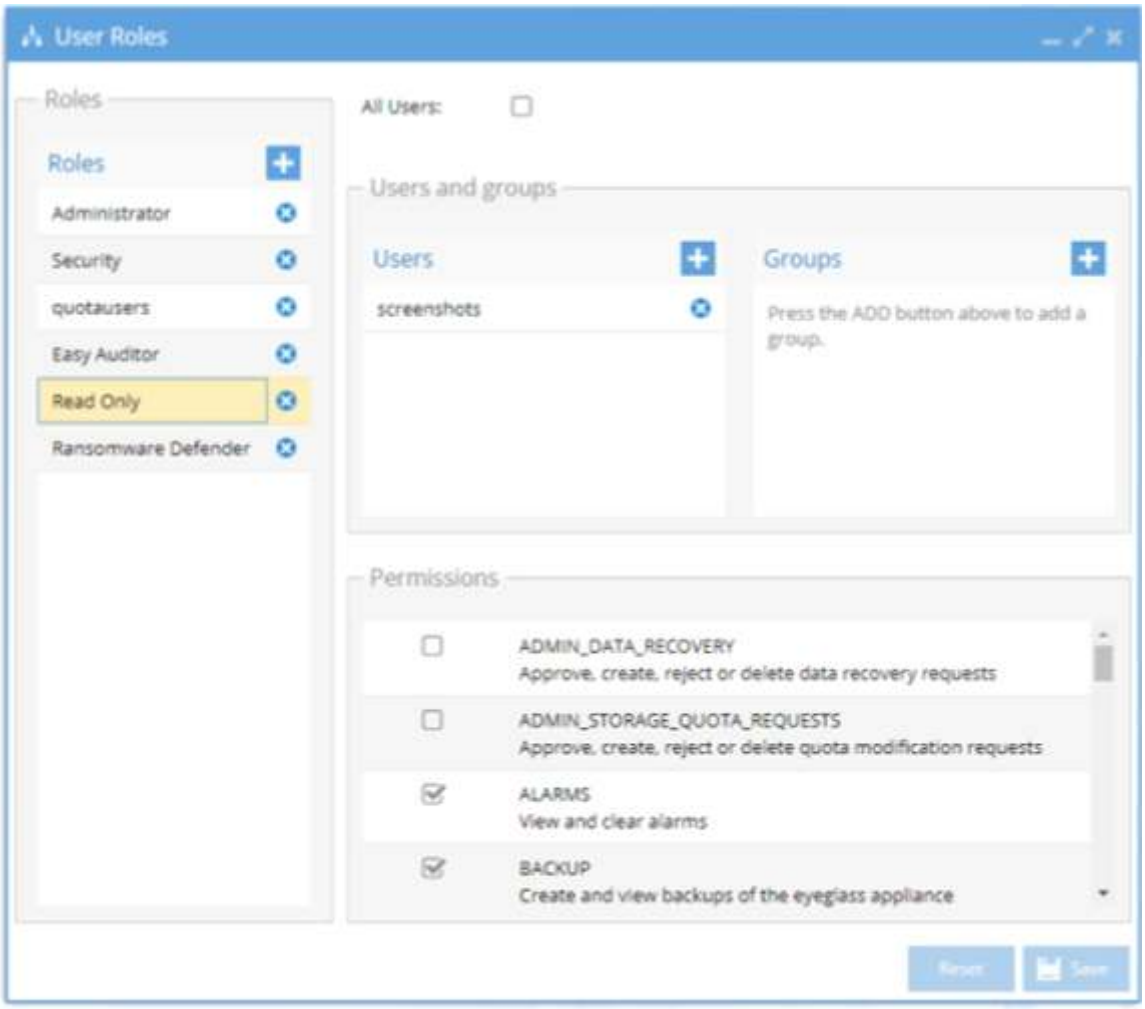

2. + next to Roles
3. Enter a Name
4. Select check boxes next to the permission to assign

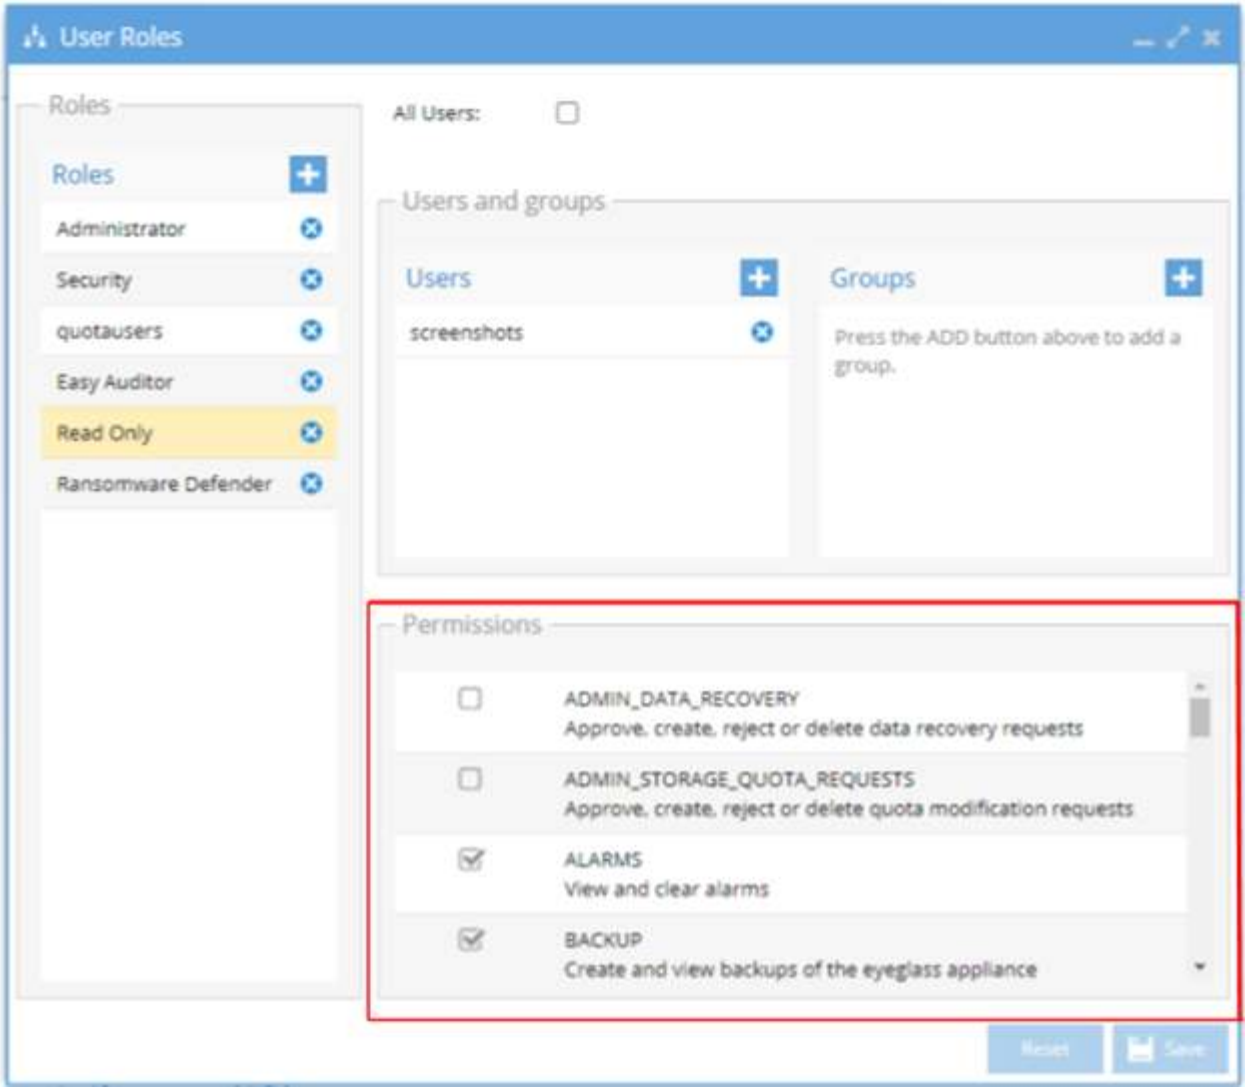

5. To Add Active Directory Groups click here
6. To Add Active Directory Users directly to the role click here.
7. Click Save
8. After configuration is complete login following Proxy login.

© Superna LLC

1.23.6. How to Login with RBAC

Home Top

How to Login with RBAC accounts on PowerScale or Active Directory Auth Provider Proxy login

1. Click the Proxy Authentication first. See screenshot below.

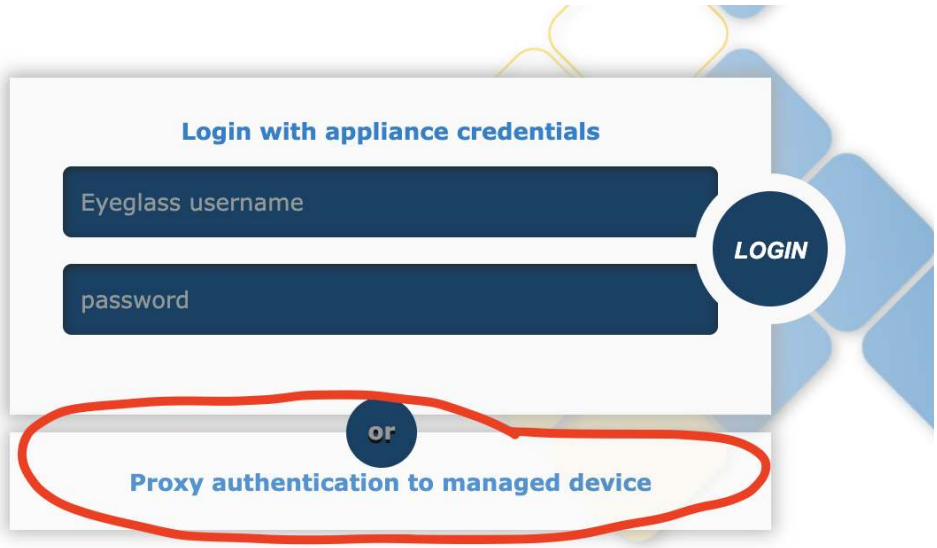

   a.

2. For users to authenticate they will need to switch to Proxy authentication to managed device login. The login should look like the screen shot below.

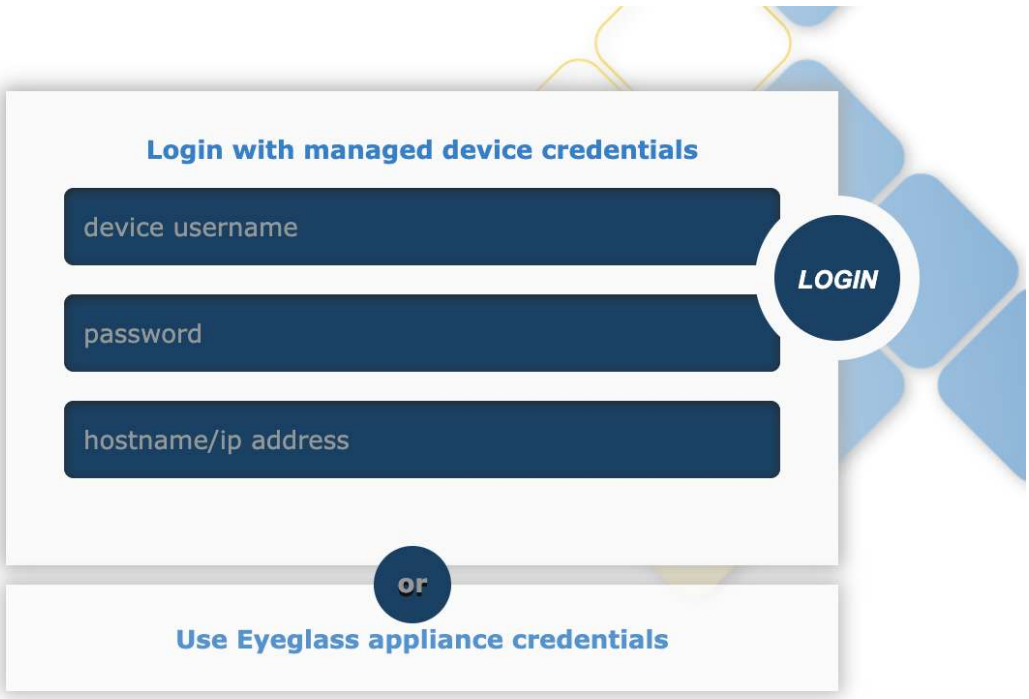

2. For example to login with an Active Directory domain user proxy authentication:
    a. Enter username in this format **<sam account name>@<DNS Domain name>**
    b. NOTE: Leave the ip address field blank, the IP used to add the cluster, and by default the System zone Authentication provider will be used.
3. For example to login with PowerScale local user proxy authentication:
    a. Enter username in this format **<username>@<cluster name>**

© Superna LLC

1.24. Eyeglass Alarm forwarding Guide - Syslog and Legacy SNMP

Home Top

- Overview:
  - Limitations
  - Deprecation Notice
  - Supported Alarms
  - Requirements:
- Configuration of SYSLOG Forwarding - > 2.5.7
- Configuration of SYSLOG Forwarding - Legacy Deprecated < 2.5.7
- How to Filter and Forward alarms
  - How to forward by alarm Severity
  - How to Forward by Alarm code (Recommended and Supported Method)
  - How to Filter by Application
    - How to Forward Ransomware Defender User Lockout and Restored Alarms Except for Security Guard Alarms
    - This example forwards all Ransomware Defender and Easy Auditor alarms
    - This example forwards all Ransomware Defender alarms Except for Security Guard alarms
- How to Integrate Ransomware Defender Events with a SIEM
  - Syslog format examples to be used for Parsing with a Syslog server -- > 2.5.7

    - How to search the Eyeglass appliance logs for examples of syslog alarm formatting
  - Example alarm formats
- How to Trouble shoot SYSLOG Forwarding
    - How to use packet capture to see syslog messages sent to your target syslog server
- Configuration of SNMP Forwarding - Legacy
  - How to send a test SNMP Trap
  - How to configure SNMP alarm Forwarding - Legacy
  - < 2.5.7 Configuration Example
  - > 2.5.7 Configuration Example
  - Example of SNMP Messages received from Eyeglass
    - SNMP Messages for Replication Jobs status
    - SNMP Messages for Policy Readiness
    - SNMP Messages for Zone Readiness
    - SNMP Messages for ALARM
    - SNMP Message for Overall DR Status
    - SNMP Message for Failover
    - SNMP Message for Ransomware Events

Overview:

Alarms can be forwarded over syslog or SNMP. We recommend syslog as SNMP is a legacy protocol with less capabilities. This guide explains how to filter and then forward alarms matching certain criteria. This guide

contains the most common examples of how to filter by application or by severity of the alarms.

NOTE: Only configure syslog forwarding OR snmp but not both. If you require both contact support.

Limitations

1. **Syslog and SNMP are limited in what information can be sent to these protocols. Email alerts will contain more information not available over Syslog and SNMP due to protocol limits.**
2. The intent of SNMP or Syslog forwarding is to make basic alarm type and severity, and detailed alarm data available in the GUI or via email that supports more text and attachments. This alarm solution provides notification of an alarm, the application that generated it, and the severity.
3. **Only the documented forwarding solutions below are supported.**

Deprecation Notice

1. **NOTE: /var/log/messages will be deprecated and no longer supported in an up coming release as a log to use for forwarding alarms. In a release 2.5.7 syslog will not be used and dedicated alarms log with syslog-ng.**

Supported Alarms

1. Eyeglass alarms listed here

Requirements:

1. Eyeglass OVF version 2.5.6 or greater. Open suse 15.1 or later. Upgrade to the latest OVF if required with the guide here.
2. Place the **SUPERNA-EYEGLASS-MIB** file onto your SNMP trap management station. It will be located here on the appliance /opt/pygls/lib/python3.6/site-packages/pygls/mibs

Configuration of SYSLOG Forwarding - > 2.5.7

1. The new alarm architecture will use a dedicated log that will roll over and provide alarm history external from the database and alarm history available in the GUI. Release 2.5.7 or later is required.
2. **README FIRST: If upgrading from a previous release or had configured < 2.5.7 syslog forwarding. You must replace the entire configuration with the example below.**
3. Review all the filter examples to match your requirement and replace the filter with one of the following scenarios.
    a. How to Filter by Severity, Application, or Alarm codes (recommended)
4. Log location
    a. /opt/superna/sca/logs/igls_alarms.log
5. Configure Syslog Forwarding
    a. Ssh to the appliance as admin user
    b. Sudo -s
    c. Enter admin password
    d. nano /etc/syslog-ng/conf.d/superna.conf
    e. The example below is going to forward specific ransomware defender events.
        i. Paste the text below into the file and change the text as follows:
        ii. replace x.x.x.x with the ip address of the syslog server ip address you want to forward messages

```
filter  f_superna {

        message("RSW0002") or
message("RSW0011") ;
                                    };
source igls_src {
file("/opt/superna/sca/logs/igls_alarms.log"); };
destination logserver { udp("x.x.x.x" port(514)); };

log {
    source(igls_src);
    filter(f_superna);
    destination(logserver);
};
```

1. After making changes syslog must be restarted to have the changes take effect
    a. systemctl restart syslog-ng
2. Check that its running
    a. systemctl status syslog-ng
3. It should show active running state
4. done

Configuration of SYSLOG Forwarding - Legacy Deprecated < 2.5.7

1. Ssh to the appliance as admin user
2. Sudo -s
3. Enter admin password
4. nano /etc/syslog-ng/conf.d/superna.conf (use vim on 42.3 OS)
5. Paste this text into the file and change the text as follows:

a. replace x.x.x.x with the ip address of the syslog server ip address you want to forward messages

```
filter f_superna {
    message("Severity:CRITICAL") or
message("Severity:WARNING") ;
};
destination logserver { udp("x.x.x.x" port(514)); };

log {
    source(src);
    source(chroots);
    filter(f_superna);
    destination(logserver);
};
```

1. After making changes syslog must be restarted to have the changes take effect
    a. systemctl restart syslog-ng
2. Check that its running
    a. systemctl status syslog-ng
3. It should show **active running state**
4. **done**

How to Filter and Forward alarms

This section provides examples of how to filter for alarms to forward to syslog.

How to forward by alarm Severity

To combine multiple Alarm severities or combine message strings see example below:
.

```
filter f_superna {
   message("Severity:CRITICAL") or message("Severity:MAJOR") ;
```

```
};
```

How to Forward by Alarm code (Recommended and Supported Method)

Use this filter example to the best option to forward exactly the alarms you need using the alarm code guide. All possible alarms are listed and provides the best option to simplify forwarding exactly the alarms you need to external systems. Get the Alarm codes and use them in the filters.

```
filter f_superna {
   message("RSW0002") or message("RSW0011") ;
};
```

How to Filter by Application

Each Eyeglass application has an alarm code to easily forward alarms based on the prefix.

1. Ransomware Defender prefix - RSW
2. Easy Auditor Prefix - EAU
3. DR - SCA

How to Forward Ransomware Defender User Lockout and Restored Alarms Except for Security Guard Alarms

In the example below 2 commonly used Ransomware Defender alarms are needed.

1. User locked out is RSW0002
2. User access restored is RSW0011

3. NOTE: Replace the yellow highlight with the security guard service account that you have configured.

```
filter f_superna {
   (message("RSW0002") or message("RSW0011"))  and not message("igls-sg") ;
};
```

This example forwards all Ransomware Defender and Easy Auditor alarms

```
filter f_superna {
   message("RSW") or message("EAU") ;
};
```

This example forwards all Ransomware Defender alarms Except for Security Guard alarms

NOTE: Replace the yellow highlight with the security guard service account that you have configured.

```
filter f_superna {
        (message("RSW")) and not message("igls-sg");
```

How to Integrate Ransomware Defender Events with a SIEM

1. The syslog message alarms generated by Ransomware Defender when a user is detect with Ransomware an alarm includes details with user ID, ip address and a subset of some of the files that were detected. The ip address can be used in a SIEM trigger to find the Ethernet port of the IP address and disable the port. See the example message format below.
2. Use the yellow highlighted sections below to build your parsing and trigger to capture the user name and PC IP address. Using

this information build a trigger in your SIEM to take action on the Ethernet port the PC is connected.

3. [DEBUG] IGLS_ALARMS:168 - Eyeglass, , Event: 2021-02-26 19:28:23.916, AID:AD02\sgdemo, Port:Nil, Type:null, EntityType:, Extra Data:{"clientIps":"172.31.1.65","info":"Successfully locked out user AD02\\sgdemo"}, Description:Locked user access.172.31.1.65, NSA, Severity:CRITICAL, Impact:false, Category:RSW0002

Syslog format examples to be used for Parsing with a Syslog server -- > 2.5.7

How to search the Eyeglass appliance logs for examples of syslog alarm formatting

1. Login to eyeglass vm over ssh as admin
2. cat /opt/superna/sca/logs/igls_alarms.log

Example alarm formats

[DEBUG] IGLS_ALARMS:168 - Eyeglass, , Event: 2021-02-26 20:15:23.983, AID:\ifs\data\dfsdata\dlp\, Port:Nil, Type:null, EntityType:, Extra Data:{"reason":"There is no smart quota for /ifs/data/dfsdata/dlp/ limited by a Data Loss Prevention threat detector. no limit is enforced."}, Description:There is no smart quota for a path limited by a Data Loss Prevention threat detector , NSA, Severity:MAJOR, Impact:false, Category:EAU0005

[DEBUG] IGLS_ALARMS:168 - Eyeglass, , Event: 2021-02-26 19:28:14.496, AID:AD02\sgdemo, Port:Nil, Type:null, EntityType:, Extra Data:{"clientIps":"172.31.1.65","event severity":"CRITICAL","user name":"AD02\\sgdemo","affected files":"\\\\prod8\\System\\ifs\\igls-securityguard\\igls-securityguard-test-file-1614385692201.iglsrswtest","affected Isilon clusters":["prod8"],"detectors":"THREAT_DETECTOR_06","number of affected files":"1","info":"Lockout required."}, Description:Ransomware event received. Event severity: CRITICAL, user: AD02\sgdemo172.31.1.65, NSA, **Severity:CRITICAL**, Impact:false, **Category:RSW0001**

[DEBUG] IGLS_ALARMS:168 - Eyeglass, , Event: 2021-02-26 19:28:23.916, AID:AD02\sgdemo, Port:Nil, Type:null, EntityType:, Extra Data:{"clientIps":"172.31.1.65","info":"Successfully locked out user AD02\\sgdemo"}, Description:Locked user access.172.31.1.65, NSA, **Severity:CRITICAL**, Impact:false, **Category:RSW0002**

How to Trouble shoot SYSLOG Forwarding

1. enable verbose logging
2. ssh to eyeglass as admin
3. sudo -s (enter admin password)
4. syslog-ng-ctl verbose --set=on
5. Check the statistics of the forwarding to the **logserver** label (this is the name assigned to the destination in all the examples)
6. syslog-ng-ctl stats | grep **logserver**

    a. If the counters are not incrementing or show zeros it means nothing has matched your filter and nothing was forwarded to your destination

7. Reset the stats to zero to test forwarding again to help trouble sheet the processed counter incrementing
    a. syslog-ng-ctl stats --reset

How to use packet capture to see syslog messages sent to your target syslog server

1. Use this command to monitor any udp syslog messages sent based on matching alarms
2. Login as admin
3. sudo -s (enter admin password)
4. Replace x.x.x.x in the command below with the ip address of your syslog server configured in the above settings file /etc/syslog-ng/config.d/f_superna.conf. This command will **NOT** display any data until a packet is sent to your syslog server based on the matching logic configured in your filter. Leave the command prompt running and continue to the next step.
    a. tcpdump -nnAs0 -i eth0 udp port 514 -v | grep -A 2 "x.x.x.x"
5. Open new ssh session as the admin user leaving the other session running
6. run the random test alarm command, this command will create a random alarm **(NOTE: The random alarm may not match your filter logic, adjust your filter logic to match on severity using the**

example above following all steps to edit the file and restart syslog-ng)

7. Run this command below several times until you see a packet appear in the first ssh session that is packet capturing all packets sent to your syslog server. This will help troubleshoot your filter and allow monitoring in realtime for any packets that are sent.
    a. **igls test AlarmTest**
    b. You may also run this command to verify any matches processed by Syslog-ng filter logic
    c. syslog-ng-ctl stats | grep **logserver**
8. Repeat the test command and stats command to verify your forwarding is working. Check your syslog server to verify the messages are appearing after you have verified the stats and packet capture show successful packets are sent.
9. Done.

Configuration of SNMP Forwarding - Legacy

1. `$ exec bash -l` (to reload your Bash session to pick up new environment settings)
2. `$ sudo -E pygls-snmptrap --setup` (to add the required entries to the syslog-ng configuration, and to configure the SNMP settings, you can re-run this command to change settings or edit this file /opt/superna/sca/conf/snmptraps.ini)
3. We need to specify the following

Server Address	IP Address of the SNMP Receiver
Port	Port number (Default 162)

SNMP Engine ID	SNMP Engine ID for SNMPv3
SNMP Version	Default 2c
Community String	Default public

Example:

```
Server Address: 172.22.22.29
Port: 162
SNMP Engine ID:
SNMP Version: 2c
Community String: public
```

4. To customize what is sent to SNMP trap destination follow instructions below for filtering based on alarm content
5. **The default configuration will forward all alarms to the SNMP destination.**
    1.

How to send a test SNMP Trap

`$ pygls-snmptrap --test` (to test sending snmp message to snmp receiver - verify this test message is received on SNMP server)
SNMPv2-MIB::snmpTrapOID.0 = SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = Superna Eyeglass Syslog-NG SNMP Notification Test Message

1. NOTE: by default the log filter will send all messages as traps. This will be a lot of traps messages.
2. It is recommended to replace the default with a specific filter of alarm severity. See next section below.

How to configure SNMP alarm Forwarding - Legacy

This explains how to select log message text to forward to SNMP. This can be used to send only INFO, Warning or Critical events. This can also be used to send specific events example Ransomware events or DR events. The default configuration will forward all alarms to the SNMP destination.

1. Ssh to the appliance as admin user
2. Sudo -s
3. Enter admin password
4. nano /etc/syslog-ng/conf.d/superna-snmp.conf
5. Edit this section below and change the text as follows to add or delete message strings to the f_superna_snmp filter section.

See example of alarm severity forwarding below. Adding additional strings allows application alarms to be forwarded.

< 2.5.7 Configuration Example

```
filter f_superna_snmp {
    message("Severity:CRITICAL") or
message("Severity:WARNING") ;
};

destination superna_snmp {
program(
"/usr/local/bin/pygls-snmptrap"
flush_lines(1)
flags(no_multi_line)
template("$ISODATE $HOST EYEGLASS
$MSGHDR$MSG\n")
);
};

log {
source(src);
source(chroots);
filter(f_superna_snmp);
destination(superna_snmp);
};
```

> 2.5.7 Configuration Example

```
filter f_superna_snmp {
    message("Severity:CRITICAL") or
message("Severity:WARNING") ;
};
```

```
source igls_src { file("/opt/superna/sca/logs/igls_alarms.log");
};
destination superna_snmp {
program(
"/usr/local/bin/pygls-snmptrap"
flush_lines(1)
flags(no_multi_line)
template("$ISODATE $HOST EYEGLASS
$MSGHDR$MSG\n")
);
};

log {
source(igls_src);
filter(f_superna_snmp);
destination(superna_snmp);
};
```

1. Save and Exit the file **(press the letter i to insert text, when done type : then type wq + enter key)**
2. Disable SNMP Mark Heartbeat
    a. Modify syslog-ng config file
    b. nano /etc/syslog-ng/syslog-ng.conf
    c. Add mark-freq(0); inside options { } clause.
    d. Example string: options { chain_hostnames(off); flush_lines(0); perm(0640); stats_freq(3600); threaded(yes); **mark-freq(0);** }
    e. Save and Exit the file (press the letter i to insert text, when done type : then type wq + enter key)

3. Now restart logging service
    a. systemctl restart syslog-ng
4. To verify the file was edited correctly and make sure syslog-ng is running
    a. systemctl status -l syslog-ng
5. done.

Example of SNMP Messages received from Eyeglass

SNMP Messages for Replication Jobs status

8/21/2017 3:55:26 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 = SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T03:55:26-04:00 rnsm04-igls-03 EYEGLASS bash[19595]: 2017-08-21 03:55:26,634 [pool-97-thread-2] DEBUG MAIN ReplicationTask:lambdarun982 [246] - ReplicationTask is done. 0 0 7619067 2
8/21/2017 3:55:21 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 = SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T03:55:21-04:00 rnsm04-igls-03 EYEGLASS bash[19595]: 2017-08-21 03:55:21,753 [pool-96-thread-1] DEBUG MAIN ReplicationTask:lambdarun980 [217] - Fetching post-configuration inventory. 0 0 7618578 2
8/21/2017 3:55:21 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 = SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T03:55:21-04:00 rnsm04-igls-03 EYEGLASS bash[19595]: 2017-08-21 03:55:21,753 [pool-96-thread-1] DEBUG MAIN ReplicationTask:lambdarun980 [214] - Unblocking deletes from the database 0 0 7618578 2
8/21/2017 3:55:20 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 = SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T03:55:20-04:00 rnsm04-igls-03 EYEGLASS bash[19595]: 2017-08-21 03:55:20,985 [pool-97-thread-1] DEBUG MAIN ReplicationTask:lambdarun979 [179] - Writing replication xml file. 0 0 7618502 2
8/21/2017 3:55:20 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 = SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T03:55:20-04:00 rnsm04-igls-03 EYEGLASS bash[19595]: 2017-08-21 03:55:20,968 [pool-97-thread-2] DEBUG MAIN ReplicationTask:lambdarun977 [124] - Writing fingerprints 0 0 7618499 2

8/21/2017 3:54:59 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 = SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T03:55:00-04:00 rnsm04-igls-03 EYEGLASS bash[19595]: 2017-08-21 03:55:00,021 [pool-97-thread-1] DEBUG MAIN ReplicationTask:lambdarun976 [109] - Fetching current inventory before running replication 0 0 7616408 2
8/21/2017 3:54:59 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 = SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T03:55:00-04:00 rnsm04-igls-03 EYEGLASS bash[19595]: 2017-08-21 03:55:00,017 [pool-97-thread-1] DEBUG MAIN ReplicationTask:lambdarun976 [104] - Clearing deleted items cache 0 0 7616408 2
8/21/2017 3:54:59 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 = SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T03:55:00-04:00 rnsm04-igls-03 EYEGLASS bash[19595]: 2017-08-21 03:55:00,003 [cron4j-task-10] INFO MAIN ReplicationTask:run [90] - Starting ReplicationTask 0 0 7616404 2

SNMP Messages for Policy Readiness

8/21/2017 3:41:09 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 = SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T03:41:09-04:00 rnsm04-igls-03 EYEGLASS bash[19595]: 2017-08-21 03:41:09,012 [pool-68-thread-1] DEBUG MAIN PolicyReadinessValidation:doPolicyValidation [194] - Policy readiness validation completed successfully 0 0 7533303 2
8/21/2017 3:41:00 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 = SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T03:41:00-04:00 rnsm04-igls-03 EYEGLASS bash[19595]: at com.superna.nde.jobengine.readiness.policyreadiness.PolicyReadinessValidation.doPolicyValidation(PolicyReadinessValidation.java:138) 0 0 7532456 2

SNMP Messages for Zone Readiness

8/21/2017 3:45:17 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 = SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T03:45:17-04:00 rnsm04-igls-03 EYEGLASS bash[19595]: 2017-08-21 03:45:17,296 [pool-75-thread-1] DEBUG MAIN ReadinessJobResultHandler:handleResult [64] - JOB rnsm04-c03_rnsm04-c04: Status: {"state":"FINISHED","jobStatus":"OK","started":1503301507126,"finished":1503301507532,"duration":406,"name":"AccessZoneValidation rnsm04-c03_rnsm04-c04","info":"Access Zone Validation","children":[],"modified":1503301507532} 0 0 7558132 2
8/21/2017 3:45:08 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 = SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T03:45:07-04:00 rnsm04-igls-03 EYEGLASS bash[19595]: 2017-08-21 03:45:07,480 [pool-80-thread-1] DEBUG MAIN AccessZoneValidation:doAccessZoneValidation [213] - { 0 0 7557218 2
8/21/2017 3:45:07 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 = SNMPv2-SMI::enterprises.50412.0.1

SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T03:45:07-04:00 rnsm04-igls-03 EYEGLASS bash[19595]: at com.superna.nde.jobengine.readiness.zonereadiness.operations.AccessZoneValidation.call(AccessZoneValidation.java:53) 0 0 7557171 2
8/21/2017 3:45:07 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 = SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T03:45:07-04:00 rnsm04-igls-03 EYEGLASS bash[19595]: at com.superna.nde.jobengine.readiness.zonereadiness.operations.AccessZoneValidation.call(AccessZoneValidation.java:70) 0 0 7557171 2
8/21/2017 3:45:07 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 = SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T03:45:07-04:00 rnsm04-igls-03 EYEGLASS bash[19595]: at com.superna.nde.jobengine.readiness.zonereadiness.operations.AccessZoneValidation.doAccessZoneValidation(AccessZoneValidation.java:315) 0 0 7557170 2
8/21/2017 3:45:07 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 = SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T03:45:07-04:00 rnsm04-igls-03 EYEGLASS bash[19595]: at com.superna.nde.jobengine.readiness.zonereadiness.operations.AccessZoneValidation.collectConfigReplication(AccessZoneValidation.java:1127) 0 0 7557170 2
8/21/2017 3:45:07 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 = SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T03:45:07-04:00 rnsm04-igls-03 EYEGLASS bash[19595]: at com.superna.nde.jobengine.readiness.zonereadiness.operations.AccessZoneValidation$$Lambda$428/501745496.apply(Unknown Source) 0 0 7557167 2
8/21/2017 3:45:07 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 = SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T03:45:07-04:00 rnsm04-igls-03 EYEGLASS bash[19595]: at com.superna.nde.jobengine.readiness.zonereadiness.operations.AccessZoneValidation.lambda$collectConfigReplication$743(AccessZoneValidation.java:1127) 0 0 7557167 2
8/21/2017 3:45:07 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 = SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T03:45:07-04:00 rnsm04-igls-03 EYEGLASS bash[19595]: 2017-08-21 03:45:07,226 [pool-28-thread-2] DEBUG MAIN AccessZoneValidation:doAccessZoneValidation [213] - { 0 0 7557154 2

SNMP Messages for ALARM

8/21/2017 3:57:00 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 = SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T03:57:00-04:00 rnsm04-igls-03 EYEGLASS bash[19595]: 2017-08-21 03:57:00,035 [cron4j-task-8] DEBUG MAIN AlarmDataManager:executeSave [2817] - >> Keys: Sync-Key: 'rnsm04-03', Severity: 'MAJOR', Description: 'ECA Service unreachable to scan for events' 0 0 7628414 2

8/21/2017 3:57:00 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 = SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T03:57:00-04:00 rnsm04-igls-03 EYEGLASS bash[19595]: 2017-08-21 03:57:00,034 [cron4j-task-8] INFO MAIN AlarmDataManager:executeSave [2815] - Sending alarm from '' to DB 0 0 7628413 2
8/21/2017 3:57:00 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 = SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T03:57:00-04:00 rnsm04-igls-03 EYEGLASS bash[19595]: 2017-08-21 03:57:00,028 [cron4j-task-8] DEBUG MAIN AlarmDataManager:executeSave [2817] - >> Keys: Sync-Key: '172.22.4.109', Severity: 'MAJOR', Description: 'ECA Node inactive or in error state' 0 0 7628412 2
8/21/2017 3:57:00 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 = SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T03:57:00-04:00 rnsm04-igls-03 EYEGLASS bash[19595]: 2017-08-21 03:57:00,028 [cron4j-task-8] INFO MAIN AlarmDataManager:executeSave [2815] - Sending alarm from '' to DB 0 0 7628411 2
8/21/2017 3:57:00 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 = SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T03:57:00-04:00 rnsm04-igls-03 EYEGLASS bash[19595]: 2017-08-21 03:57:00,025 [cron4j-task-8] DEBUG MAIN AlarmDataManager:executeSave [2817] - >> Keys: Sync-Key: '172.22.4.108', Severity: 'MAJOR', Description: 'ECA Node inactive or in error state' 0 0 7628411 2
8/21/2017 3:57:00 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 = SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T03:57:00-04:00 rnsm04-igls-03 EYEGLASS bash[19595]: 2017-08-21 03:57:00,025 [cron4j-task-8] INFO MAIN AlarmDataManager:executeSave [2815] - Sending alarm from '' to DB 0 0 7628411 2
8/21/2017 3:56:59 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 = SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T03:57:00-04:00 rnsm04-igls-03 EYEGLASS bash[19595]: 2017-08-21 03:57:00,019 [cron4j-task-8] DEBUG MAIN AlarmDataManager:executeSave [2817] - >> Keys: Sync-Key: '172.22.4.107', Severity: 'MAJOR', Description: 'ECA Node inactive or in error state' 0 0 7628408 2

SNMP Message for Overall DR Status

8/21/2017 3:47:12 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 = SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T03:47:12-04:00 rnsm04-igls-03 EYEGLASS bash[19595]: 2017-08-21 03:47:12,283 [pool-4-thread-33] DEBUG MAIN Policies:getAllPolicies [56] - [{"policy_name":"InsightIQ-NFSDS","policy_enabled":true,"policy_last_success":1497605568000,"policy_last_run":1497605568000,"policy_last_status":"finished","policy_status":"SUCCESS","overall_dr_status":"WARNING","job_status":"SUCCESS","job_name":"rnsm04-c03_InsightIQ-NFSDS","job_last_run":1503301518896,"job_last_success":1503301518896,"job_source":"rnsm04-c03","job_destination":"rnsm04-c04","job_enabled":true,"job_has_policy":true,"audit_status":"AUDITSUCCEEDED","policy_read

iness_last_success":1503301524918},{"policy_name":"z01-smb01-synciq","policy_enabled":true,"policy_last_success":1498033910000,"policy_last_run":1498033910000,"policy_last_status":"finished","policy_status":"SUCCESS","overall_dr_status":"WARNING","job_status":"SUCCESS","job_name":"rnsm04-c03_z01-smb01-synciq","job_last_run":1503301518900,"job_last_success":1503301518900,"job_source":"rnsm04-c03","job_destination":"rnsm04-c04","job_enabled":true,"job_has_policy":true,"audit_status":"AUDITSUCCEEDED","policy_readiness_last_success":1503301524923},{"policy_name":"z01-smb01-synciq_mirror","policy_enabled":false,"policy_last_success":1498033810000,"policy_last_run":1498033905000,"policy_last_status":"finished","policy_status":"DISABLED","overall_dr_status":"FAILED_OVER","job_status":"DISABLED","job_name":"rnsm04-c04_z01-smb01-synciq_mirror","job_last_run":1498033840357,"job_last_success":1498033840357,"job_source":"rnsm04-c04","job_destination":"rnsm04-c03","job_enabled":false,"job_has_policy":true,"audit_status":"AUDITSUCCEEDED","policy_readiness_last_success":1503301528213}] 0 0 7569630 2

SNMP Message for Failover

8/21/2017 5:49:55 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 = SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T05:49:54-04:00 rnsm04-igls-03 EYEGLASS bash[19595]: 2017-08-21 05:49:54,241 [pool-284-thread-1] INFO com.superna.nde.jobengine.failover.operations.AddReportsToLogs AddReportsToLogs:lambda$appendReportsToLog$617 [69] - { 0 0 8305934 2
8/21/2017 5:49:55 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 = SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T05:49:54-04:00 rnsm04-igls-03 EYEGLASS bash[19595]: 2017-08-21 05:49:54,238 [pool-284-thread-1] INFO com.superna.nde.jobengine.failover.operations.AddReportsToLogs AddReportsToLogs:lambda$appendReportsToLog$617 [70] - ********** ** 0 0 8305934 2
8/21/2017 5:49:55 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 = SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T05:49:54-04:00 rnsm04-igls-03 EYEGLASS bash[19595]: 2017-08-21 05:49:54,238 [pool-284-thread-1] INFO com.superna.nde.jobengine.failover.operations.AddReportsToLogs AddReportsToLogs:lambda$appendReportsToLog$617 [69] - { 0 0 8305907 2
8/21/2017 5:49:55 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 = SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T05:49:54-04:00 rnsm04-igls-03 EYEGLASS bash[19595]: 2017-08-21 05:49:54,236 [pool-284-thread-1] INFO com.superna.nde.jobengine.failover.operations.AddReportsToLogs AddReportsToLogs:lambda$appendReportsToLog$617 [70] - ********** ** 0 0 8305907 2
8/21/2017 5:49:54 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 = SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T05:49:54-04:00 rnsm04-igls-03 EYEGLASS bash[19595]: 2017-08-21 05:49:54,236 [pool-284-thread-1] INFO

com.superna.nde.jobengine.failover.operations.AddReportsToLogs
AddReportsToLogs:lambda$appendReportsToLog$617 [69] - { 0 0 8305880 2
8/21/2017 5:49:54 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 =
SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T05:49:54-04:00 rnsm04-igls-03 EYEGLASS
bash[19595]: 2017-08-21 05:49:54,234 [pool-284-thread-1] INFO
com.superna.nde.jobengine.failover.operations.AddReportsToLogs
AddReportsToLogs:lambda$appendReportsToLog$617 [70] -
****
** 0 0 8305879 2
8/21/2017 5:49:54 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 =
SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T05:49:54-04:00 localhost EYEGLASS
[INFO] SYSLOG:154 - Eyeglass, , Event: 2017-08-21 05:49:54.253, AID:rnsm04-c03_Policy
Failover 2017-08-21_05-47-05, Port:Nil, Type:null, EntityType:, Extra
Data:{"Status":"Success","Finished":1503308994249,"Started":1503308826347,"URL":"https://1
72.22.4.89/failover_logs/Policy_Failover__rnsm04-c03__2017-08-21_05-47-
05__SUCCESS/Policy_Failover__rnsm04-c03__2017-08-21_05-47-05__SUCCESS.json"},
Description:Failover Succeeded , NSA, Severity:INFORMATIONAL, Impact:false,
Category:SCA0040 0 0 8305846 2
8/21/2017 5:49:54 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 =
SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T05:49:54-04:00 rnsm04-igls-03 EYEGLASS
bash[19595]: 2017-08-21 05:49:54,234 [pool-284-thread-1] INFO
com.superna.nde.jobengine.failover.operations.AddReportsToLogs
AddReportsToLogs:lambda$appendReportsToLog$617 [69] - { 0 0 8305837 2
8/21/2017 5:49:54 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 =
SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T05:49:54-04:00 rnsm04-igls-03 EYEGLASS
bash[19595]: 2017-08-21 05:49:54,229 [pool-284-thread-1] INFO
com.superna.nde.jobengine.failover.operations.AddReportsToLogs
AddReportsToLogs:appendReportsToLog [66] -
****
** 0 0 8305837 2
8/21/2017 5:49:54 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 =
SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T05:49:54-04:00 rnsm04-igls-03 EYEGLASS
bash[19595]: 2017-08-21 05:49:54,228 [pool-284-thread-1] INFO
com.superna.nde.jobengine.failover.operations.AddReportsToLogs
AddReportsToLogs:appendReportsToLog [65] - SyncIQ Reports For Policy: z01-smb01-
synciq 0 0 8305836 2
8/21/2017 5:49:38 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 =
SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T05:49:38-04:00 rnsm04-igls-03 EYEGLASS
bash[19595]: 2017-08-21 05:49:38,740 [pool-298-thread-1] DEBUG MAIN
QuotaJobFactory:runPrepJob [77] - Is controlled failover? true 0 0 8304276 2
8/21/2017 5:49:38 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 =
SNMPv2-SMI::enterprises.50412.0.1

SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T05:49:38-04:00 rnsm04-igls-03 EYEGLASS bash[19595]: 2017-08-21 05:49:38,470 [pool-281-thread-1] DEBUG MAIN QuotaJobFactory:runPrepJob [77] - Is controlled failover? true 0 0 8304249 2
8/21/2017 5:47:28 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 = SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T05:47:28-04:00 rnsm04-igls-03 EYEGLASS bash[19595]: 2017-08-21 05:47:28,151 [pool-273-thread-1] DEBUG MAIN RunConfigurationReplication:handleReplication [48] - Starting replication during failover. 0 0 8291218 2
8/21/2017 5:47:06 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 = SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T05:47:06-04:00 rnsm04-igls-03 EYEGLASS bash[19595]: 2017-08-21 05:47:06,520 [pool-273-thread-1] DEBUG MAIN FailoverStep:call [132] - DONE Wait for other failover jobs to complete 0 0 8289054 2
8/21/2017 5:47:06 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 = SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T05:47:06-04:00 rnsm04-igls-03 EYEGLASS bash[19595]: 2017-08-21 05:47:06,516 [pool-273-thread-1] DEBUG MAIN FailoverStep:call [118] - Starting Wait for other failover jobs to complete 0 0 8289053 2
8/21/2017 5:47:05 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 = SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T05:47:05-04:00 rnsm04-igls-03 EYEGLASS bash[19595]: 2017-08-21 05:47:05,574 [pool-4-thread-120] INFO MAIN PolicyFailoverJobFactory:createJob [83] - in policy failover 0 0 8288959 2

SNMP Message for Ransomware Events

8/21/2017 6:36:25 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 = SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T06:36:25-04:00 rnsm04-igls-03 EYEGLASS bash[19595]: 2017-08-21 06:36:25,923 [pool-4-thread-146] DEBUG MAIN HandleRdaEvent:post [24] - Received ECA ransomware notification for sid S-1-5-21-4205747320-2446522354-1604720750-11190 0 0 8584996 2
8/21/2017 6:36:25 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 = SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T06:36:25-04:00 rnsm04-igls-03 EYEGLASS bash[19595]: 2017-08-21 06:36:25,923 [pool-4-thread-146] DEBUG MAIN RequestDispatcher:getPlugin [180] - retrieving plugin: com.superna.scaapi.plugins.ransomware.HandleRdaEvent 0 0 8584995 2
8/21/2017 6:36:25 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 = SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T06:36:25-04:00 rnsm04-igls-03 EYEGLASS bash[19595]: 2017-08-21 06:36:25,901 [pool-4-thread-146] DEBUG MAIN HandleRdaEvent:post [24] - Received ECA ransomware notification for sid S-1-5-21-4205747320-2446522354-1604720750-11190 0 0 8584993 2
8/21/2017 6:36:25 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 = SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T06:36:25-04:00 rnsm04-igls-03 EYEGLASS bash[19595]: 2017-08-21 06:36:25,901 [pool-4-thread-146] DEBUG MAIN

RequestDispatcher:getPlugin [180] - retrieving plugin:
com.superna.scaapi.plugins.ransomware.HandleRdaEvent 0 0 8584993 2
8/21/2017 6:36:21 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 =
SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T06:36:21-04:00 rnsm04-igls-03 EYEGLASS
bash[19595]: 2017-08-21 06:36:21,151 [Thread-31] INFO SYSLOG AlarmHandlerTask:run [154]
- Eyeglass, , Event: 2017-08-21 06:36:21.149, AID:RNSM04\rnsm04-t32, Port:Nil, Type:null,
EntityType:, Extra Data:{"severity":"WARNING","user name":"RNSM04\\rnsm04-
t32","files":["\\\\rnsm04-c03\\zone01\\ifs\\data\\zone01\\z01-
smb01\\Data01\\dtest3.txt","\\\\rnsm04-c03\\zone01\\ifs\\data\\zone01\\z01-
smb01\\Data01\\ctest4.txt","\\\\rnsm04-c03\\zone01\\ifs\\data\\zone01\\z01-
smb01\\Data01\\dtest1.txt"],"explanation":"New ransomware event created","sid":"S-1-5-21-
4205747320-2446522354-1604720750-11190"}, Description:Ransomware signal received. ,
NSA, Severity:CRITICAL, Impact:false, Category:SCA0061 0 0 8584517 2
8/21/2017 6:36:21 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 =
SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T06:36:21-04:00 localhost EYEGLASS
[INFO] SYSLOG:154 - Eyeglass, , Event: 2017-08-21 06:36:21.149, AID:RNSM04\rnsm04-t32,
Port:Nil, Type:null, EntityType:, Extra Data:{"severity":"WARNING","user
name":"RNSM04\\rnsm04-t32","files":["\\\\rnsm04-c03\\zone01\\ifs\\data\\zone01\\z01-
smb01\\Data01\\dtest3.txt","\\\\rnsm04-c03\\zone01\\ifs\\data\\zone01\\z01-
smb01\\Data01\\ctest4.txt","\\\\rnsm04-c03\\zone01\\ifs\\data\\zone01\\z01-
smb01\\Data01\\dtest1.txt"],"explanation":"New ransomware event created","sid":"S-1-5-21-
4205747320-2446522354-1604720750-11190"}, Description:Ransomware signal received. ,
NSA, Severity:CRITICAL, Impact:false, Category:SCA0061 0 0 8584517 2
8/21/2017 6:36:20 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 =
SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T06:36:20-04:00 rnsm04-igls-03 EYEGLASS
bash[19595]: 2017-08-21 06:36:20,301 [pool-4-thread-146] DEBUG MAIN
HandleRdaEvent:post [24] - Received ECA ransomware notification for sid S-1-5-21-
4205747320-2446522354-1604720750-11190 0 0 8584437 2
8/21/2017 6:36:20 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 =
SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T06:36:20-04:00 rnsm04-igls-03 EYEGLASS
bash[19595]: 2017-08-21 06:36:20,300 [pool-4-thread-146] DEBUG MAIN
RequestDispatcher:getPlugin [180] - retrieving plugin:
com.superna.scaapi.plugins.ransomware.HandleRdaEvent 0 0 8584437 2
8/21/2017 6:36:20 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 =
SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T06:36:20-04:00 rnsm04-igls-03 EYEGLASS
bash[19595]: 2017-08-21 06:36:20,282 [pool-4-thread-146] DEBUG MAIN
HandleRdaEvent:post [24] - Received ECA ransomware notification for sid S-1-5-21-
4205747320-2446522354-1604720750-11190 0 0 8584434 2
8/21/2017 6:36:20 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 =
SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T06:36:20-04:00 rnsm04-igls-03 EYEGLASS
bash[19595]: 2017-08-21 06:36:20,280 [pool-4-thread-146] DEBUG MAIN
RequestDispatcher:getPlugin [180] - retrieving plugin:
com.superna.scaapi.plugins.ransomware.HandleRdaEvent 0 0 8584434 2

8/21/2017 6:36:20 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 = SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T06:36:20-04:00 rnsm04-igls-03 EYEGLASS bash[19595]: 2017-08-21 06:36:20,265 [pool-4-thread-146] DEBUG MAIN HandleRdaEvent:post [24] - Received ECA ransomware notification for sid S-1-5-21-4205747320-2446522354-1604720750-11190 0 0 8584429 2
8/21/2017 6:36:20 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 = SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T06:36:20-04:00 rnsm04-igls-03 EYEGLASS bash[19595]: 2017-08-21 06:36:20,264 [pool-4-thread-146] DEBUG MAIN RequestDispatcher:getPlugin [180] - retrieving plugin: com.superna.scaapi.plugins.ransomware.HandleRdaEvent 0 0 8584429 2
8/21/2017 6:36:20 AM 172.22.4.89 SNMPv2-MIB::snmpTrapOID.0 = SNMPv2-SMI::enterprises.50412.0.1
SNMPv2-SMI::enterprises.50412.1.1 = 2017-08-21T06:36:20-04:00 rnsm04-igls-03 EYEGLASS bash[19595]: 2017-08-21 06:36:20,183 [pool-4-thread-146] DEBUG MAIN HandleRdaEvent:post [24] - Received ECA ransomware notification for sid S-1-5-21-4205747320-2446522354-1604720750-11190 0 0 8584421 2

© Superna LLC

1.25. How to Setup Email alarms with Exchange Server

Home Top

- Overview
- Typical Solution Configuration
- Steps for Connecting Eyeglass to Exchange Server TLS Disabled Anonymous SMTP with narrow Scope
- Use this procedure to reduce the scope of the relay to the Eyeglass appliance

Overview

The following guide is intended to assist customers with setup of Microsoft Exchange Server with Eyeglass email alarms, reports (change reports, alarm reports, Recovery Objective Reports) and phone home support. This is not intended to be a comprehensive guide on Exchange server setup but provides the minimum required setup needed to get email relay functioning.

Typical Solution Configuration

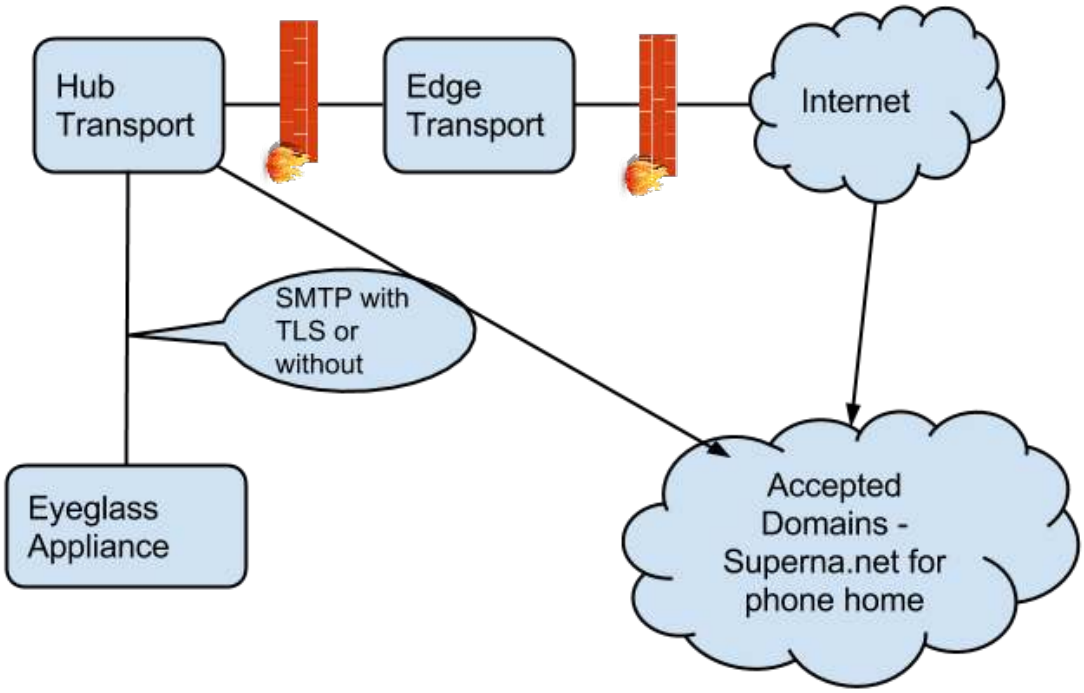

Steps for Connecting Eyeglass to Exchange Server TLS Disabled Anonymous SMTP with narrow Scope

Use this procedure to reduce the scope of the relay to the Eyeglass appliance.

1. Login to Eyeglass and open Notification Center from the start menu:
2. 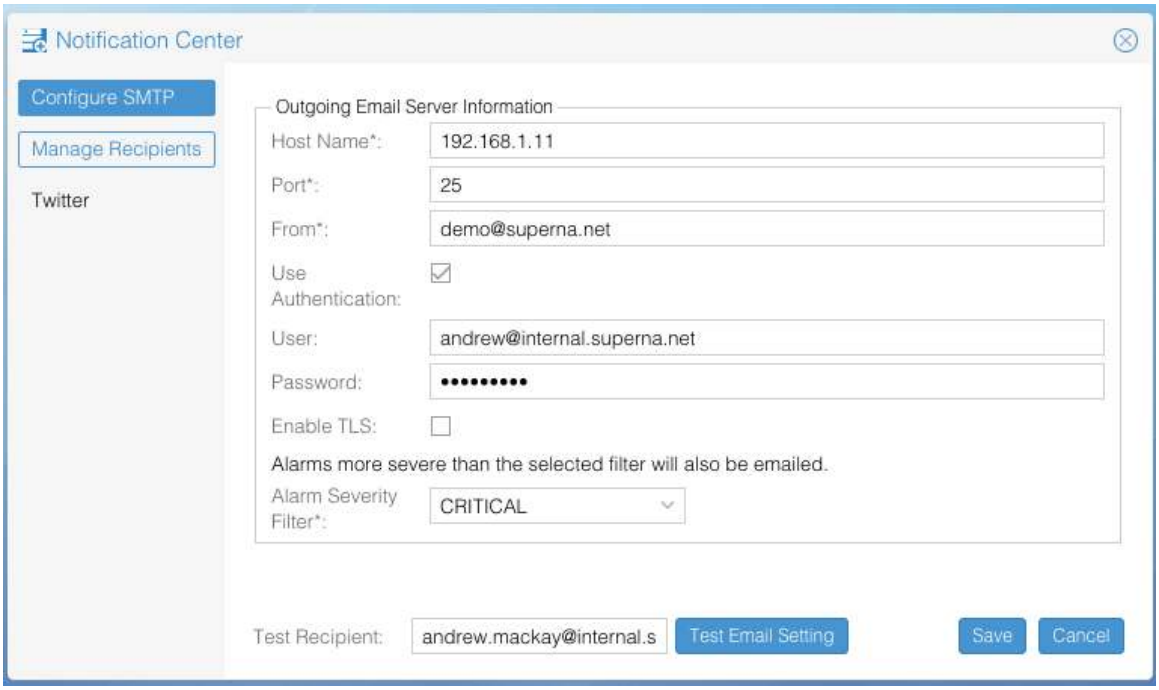

3. Add an email address to send alarms using SMTP information to Exchange Server.
    4. Enter the ip address of an **Exchange FrontEnd Transport server**
    5. Disable the TLS option on the Receiver Connector on Exchange. The Screenshots below shows a Receiver Connector configured with TLS disabled, and the check boxes to allow external systems to relay SMTP through Exchange with a narrow scope set to the Eyeglass source IP address as being allowed to relay through exchange.
    6. Screenshot below is with:
        1. Anonymous disabled and requires authentication to Exchange server
        2. TLS Disabled

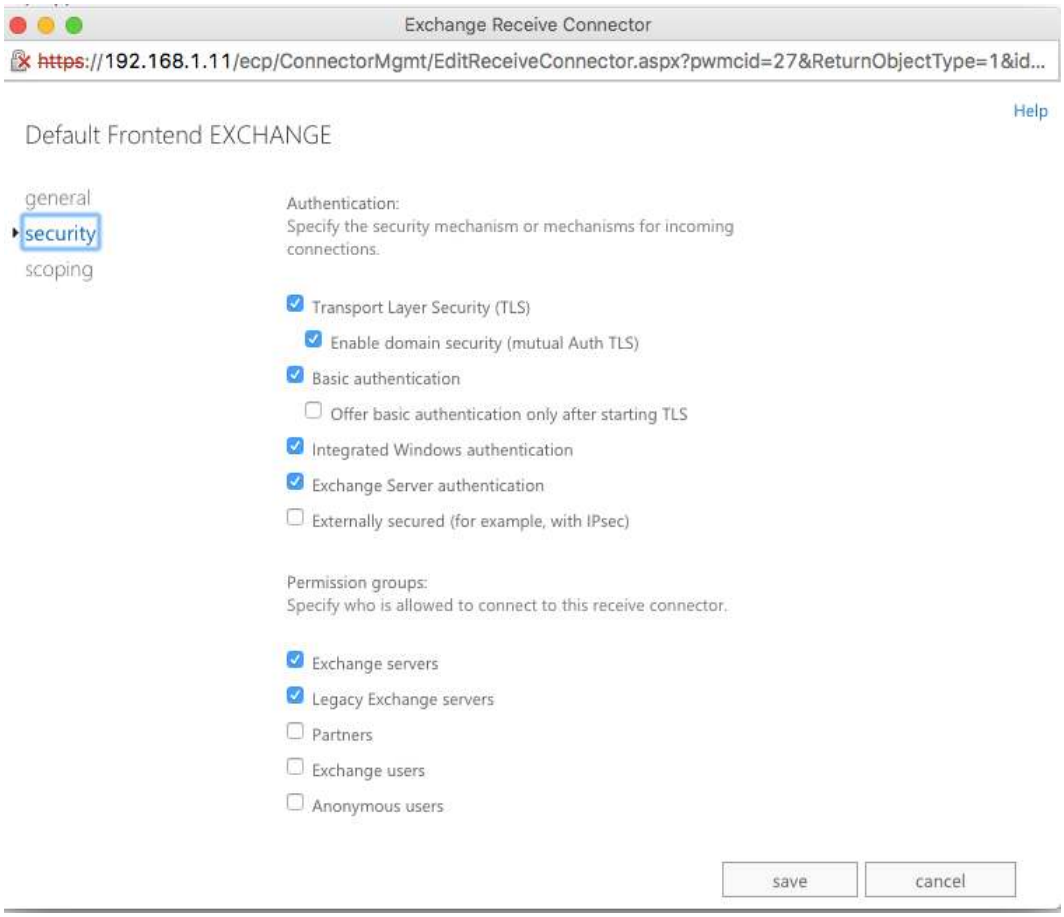

1.
2. Screenshot below is with:
    1. Enabled and Eyeglass configuration
    2. TLS Disabled

3. 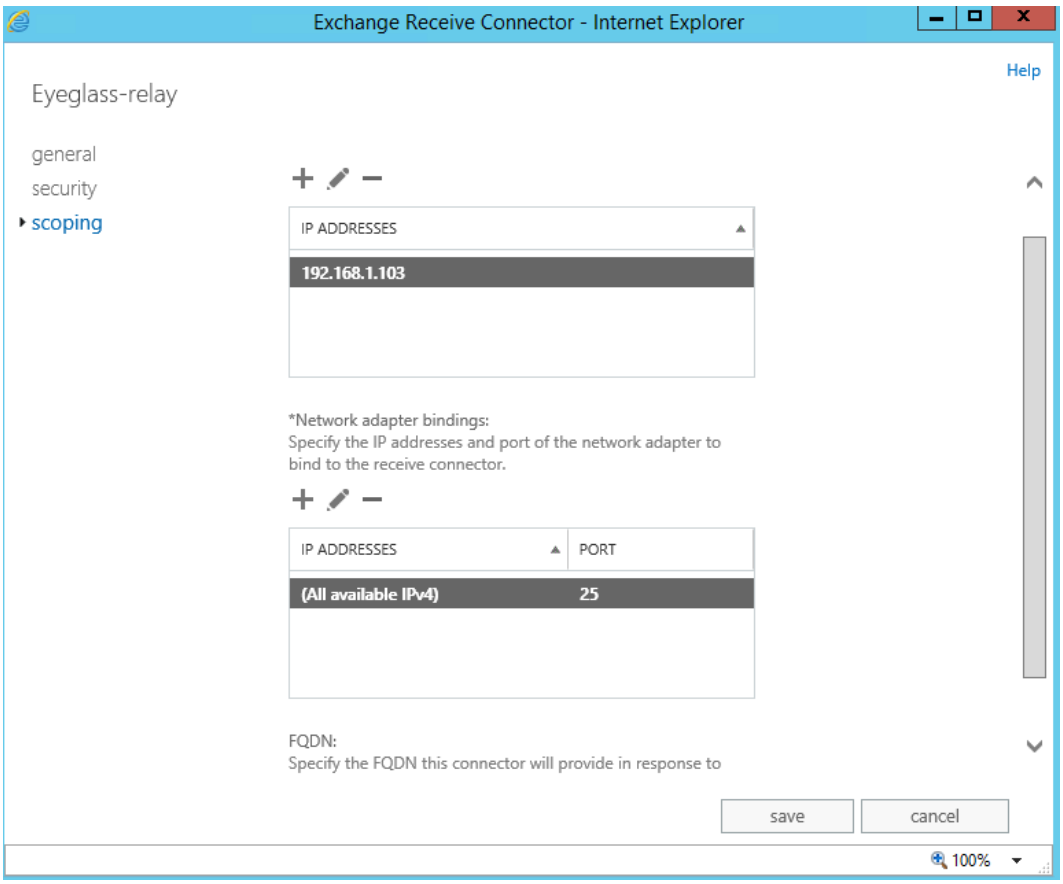

4. **NOTE: When using anonymous no authentication is required to send email so scope of Eyeglass ip address should be added to restrict this to Eyeglass only**

5. **NOTE: The port must be set to a unique port not already in use on the exchange server Example port 256.**

6. **NOTE: if exchange server logs this error** `certificate_unknown(46),` it indicates TLS parameter failure and TLS should be disabled as the certificate is unknown. Action Disable TLS on Eyeglass and retest.

7. **Note: The scope is set to the Eyeglass ip address for this connection policy using port 25 default but will need to change the port if the default Exchange Receive Connector is on this same machine, as a port conflict will occur on the host.**

8. Click the test button. You should receive an email as per below example.

9. 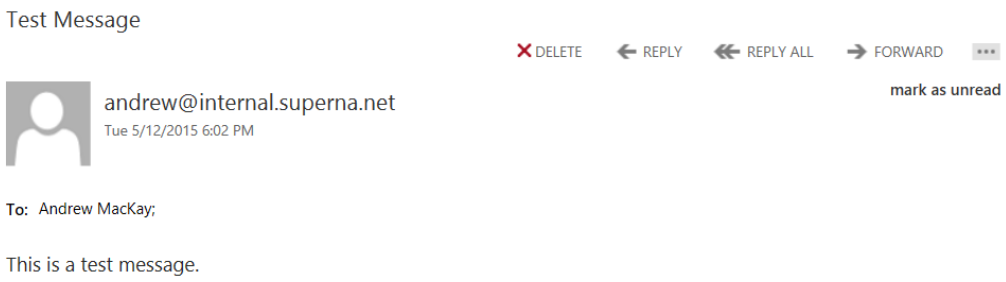

10. If a connection error occurs view the debug information to get details on the failure.
11. Verify the email is received if no error is returned.

Steps for Connecting Eyeglass to Exchange Server TLS Enabled Authenticated and Basic Authentication

1. Configure the Exchange server Receive Connector on the Transport Hub to use TLS and Basic authentication, with email and password based login for mail relay, as per Screenshot below.

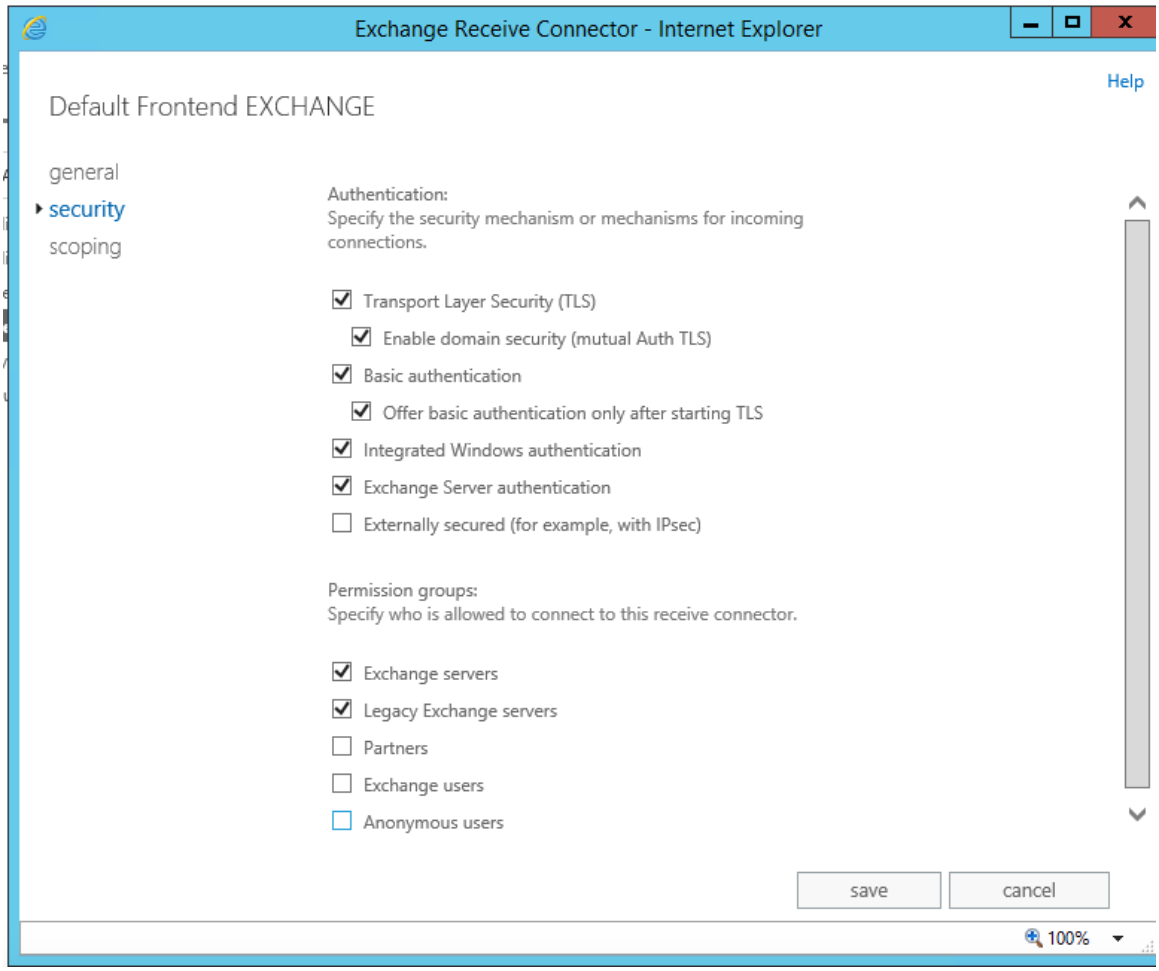

2.
3. **NOTE: The default Receive connector, enables Anonymous Users which is often disabled on Exchange installations.**

© Superna LLC

1.26. How to Change PowerScale IP Address in Eyeglass

Home Top

Edit PowerScale IP Address in Eyeglass

- Overview of Steps
- Procedure

Overview of Steps

IMPORTANT!
During this procedure, there will be an Eyeglass service interruption. Any configuration items added / updated / deleted on the source will not be synchronized to the target until the procedure is completed.

1. Prepare Eyeglass for the IP address change by disabling configuration replication.
2. Change the IP Address on the PowerScale Cluster itself.
3. Update Eyeglass for the new IP Address information.
4. Prerequisite: Networking between Eyeglass and the new IP address must be setup.
5. Enable Eyeglass configuration replication.

Procedure

5. Prepare Eyeglass for IP address change by disabling configuration replication.

- Log into Eyeglass web page.
- Click **Jobs** to open the Jobs window.
- **First disable all configuration replication jobs**

- Select all configuration replication jobs. You can use the checkbox at the top of the **Job Name** column.
- Scroll down the Job list and confirm that all Jobs have the State **User Disabled**.
- **Note**: If a Job is already in Policy Disabled state because the related SyncIQ policy is disabled, it will remain in this state. Eyeglass also does not run configuration replication for Jobs that are Policy Disabled so this state is OK to proceed with the Edit IP address procedure.
- Select **Select a bulk action**.
- Select **Enable/Disable**.

Disable replication task using Eyeglass shell

- Open the **Jobs** window and select the **Running Jobs** menu. If there is a Configuration Replication Job with Status **RUNNING**, wait for this job to complete. When it is completed, the Status will be **FINISHED**.
- Once Configuration Replication Job is **FINISHED**, click **Eyeglass Main Menu.**
- Enter the following command: igls admin schedules set --id Replication --enabled false
- Open the **Jobs** window again and select the **Running Jobs** menu. Again, confirm that there is no Configuration Replication Job with Status **RUNNING**. If there is, wait for this job to complete. When it is completed, the Status will be **FINISHED**.

Change the PowerScale Cluster IP address.

- This procedure is outside Eyeglass.
- Update Eyeglass with new IP address information.
- **Prerequisite: Networking between Eyeglass and the new IP address must be setup.**
  - Login to the Eyeglass web page.
  - Click **Inventory View**

- Right Click your desired cluster
- Click **Edit**
- Enter the new IP address in the **SmartConnect Service IP** field and **Submit**.
- **IMPORTANT!** Once you Submit, you cannot Edit the IP address again without having run the Eyeglass Replication Task at least once.

1.
2.

Enable Eyeglass configuration replication

- Enable one configuration replication Job.
  - Select one configuration replication Job.
  - Select: **Select a bulk action,** Select: **Enable/Disable**
  - The Job State updated to last known state
- Wait for next Replication Task to begin (within 5 minutes will start).
  - Check the status of the Configuration Replication job from the **Jobs** / **Running Jobs** window.
  - The Job should run without any error related to unknown source or target.
  - If no errors, enable remaining Configuration replication jobs.
- Procedure complete

© Superna LLC

1.27. Pre Post Failover Scripting Guide

Home Top

- What's new
- Script Engine Overview
- Typical Script Use Cases
- Script Engine Admin Procedures
- Script Engine CURL Tips
- Script Engine Understanding Remote Execution to Hosts
  - Video: How to Remount Exports Automation
  - ssh Passwordless Login to Remote Linux Hosts
  - How to Install Windows SSH Server for remote Powershell execution from Eyeglass with SSH
    - Enable SSH on Windows Server
    - Test the SSH Server with Powershell:
  - How to setup keyless ssh login for Eyeglass Script engine with Windows SSH and Powershell
  - Access Zone Example Bash Script using ssh Keys to Remotely Execute a Command
  - Remote Host Script Example:
- Eyeglass Access Zone Failover and Failback Script Example
- Adding Additional Script Language Support to the Appliance
- Script Run Time Variables

- Sample Execution Rules & Overall Failover Job Status Impact
- Sample Scripts in the Library
  - NodeJS - Example Script
  - NodeJS - Example Output
  - Python - Example Script
  - Python - Example Output
  - Bash - Example Script
  - Bash - Example Output
  - Consolidated Post Failover & Failback Script (Node.JS)
  - Consolidated Post Failover & Failback Script (Node.JS) I.E Multiple Zone Access

What's new

1. New script variables in 2.5.6 update 2 for handling failovver conditions in pre, post or unified scripts that need to know the status of the failover or errors to handle applicastion specific logic better. See the variable definition section for details.

Script Engine Overview

The new script engine feature provides an icon on the desktop that provides the following functions:

1. Save scripts.
2. Create new scripts.
3. Activate and deactivate scripts for failover.
4. Test scripts in simulated failover scenario.
5. Debug and edit scripts.
6. Scripts support in Bash, Nodejs, and Python languages.
7. Run time variables provide access to failover meta data, to complete simplified scripts that leverage variable replacement to handle many different failover scenarios.
8. **Pre-failover** - shutdown applications, unmount.
9. **Unified scripts** - handles either failover or failback with a single script, that can handle logic for either operation.
10. **Post-failover scripts** - runs when target cluster is writeable unmount, remount or mount only logic and application start up.

Typical Script Use Cases

Many failover scenarios depend on extra steps performed on devices, software, and infrastructure external to the NAS cluster. This tasks can now be automated with output captured and appended to Eyeglass failover logic logs.

- NFS host mount and remount automation.
- Application bring up and down logic to start applications post failover.
- Send alerts or emails.
- Run API commands on 3rd party equipment. (i.e. load balancer, switch, router or firewall).
- Shutdown an application.
- IP Load balancing solution and storage layer failover for web tier and storage tier dependencies.

The screenshot below shows the editor, admin console for script editing activation and testing:

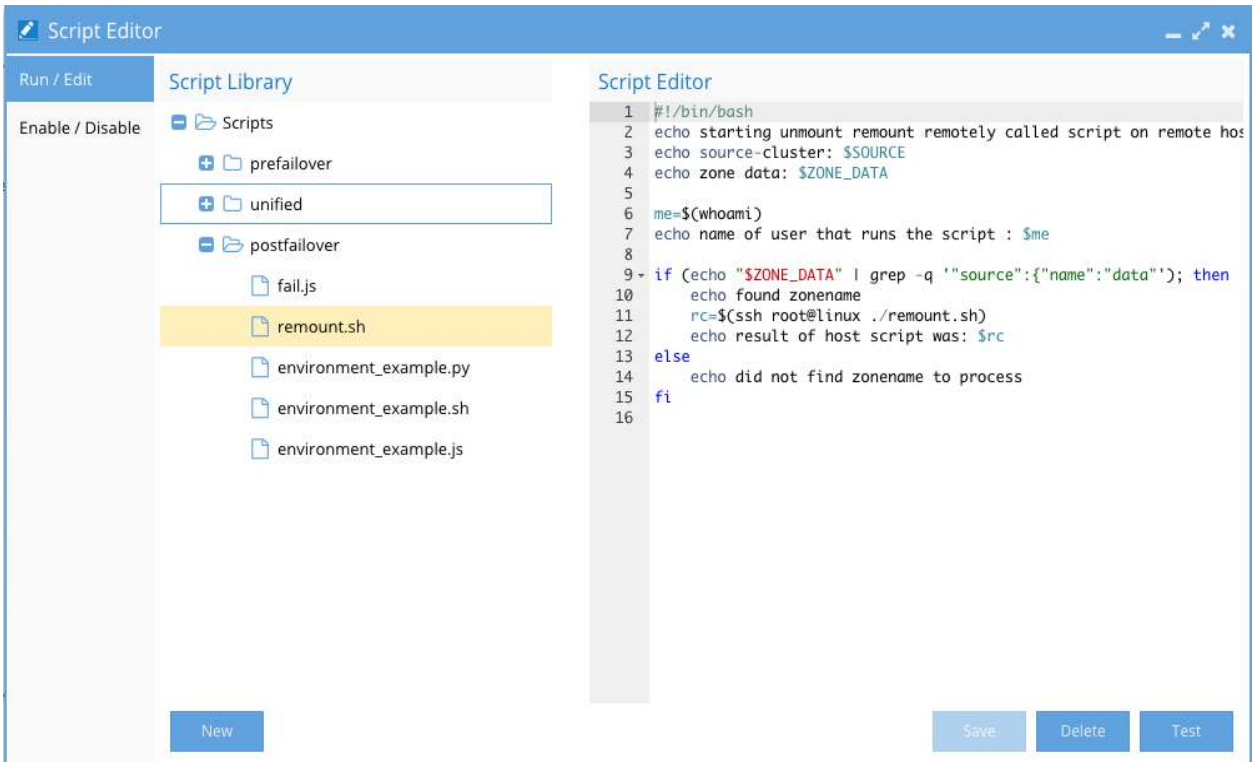

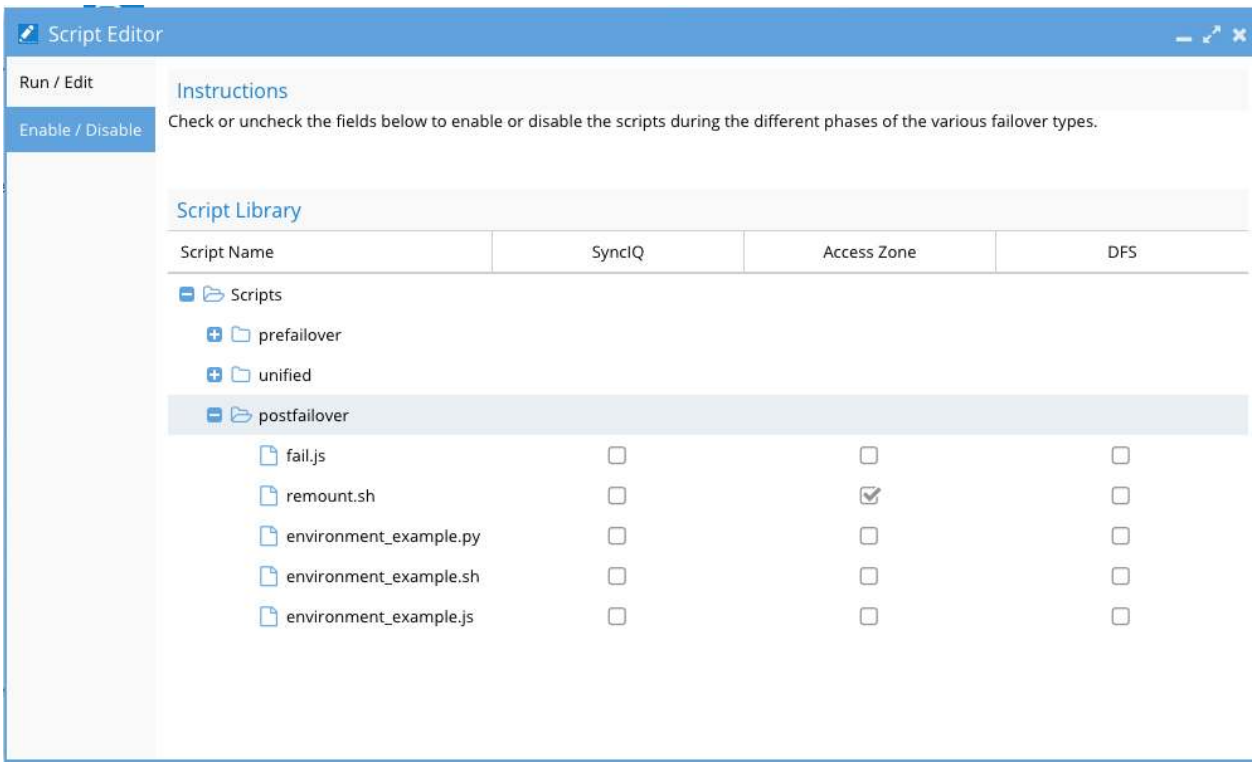

Script Engine Admin Procedures

Script Library - store many scripts in the library, but only activate some using the enable/disable menu, to enable or disable a script for one or more failover mode:

1. Add new scripts or delete existing scripts from the library.

2. Select a script to edit.
3. Numbered lines are for easier script editing and debugging.
4. Select a script click Test to see how it performs.
    a. You must select a failover mode and a SyncIQ policy (it won't be failed over). This is done to pass variables that might be used for this failover into the script test to allow easier script that can handle multiple policies.
    b. Decide if the script should be prefailover, unified or post failover script mode, and place it in the correct folder, to ensure it executes at the right location in the failover logic.
    c. Enable or disable scripts for each failover mode as required.
    d. NOTE: All scripts that are enabled, run for all failovers using the selected mode, so logic needs to handle each policy, Access Zone option or ensure logic does not run when not required.

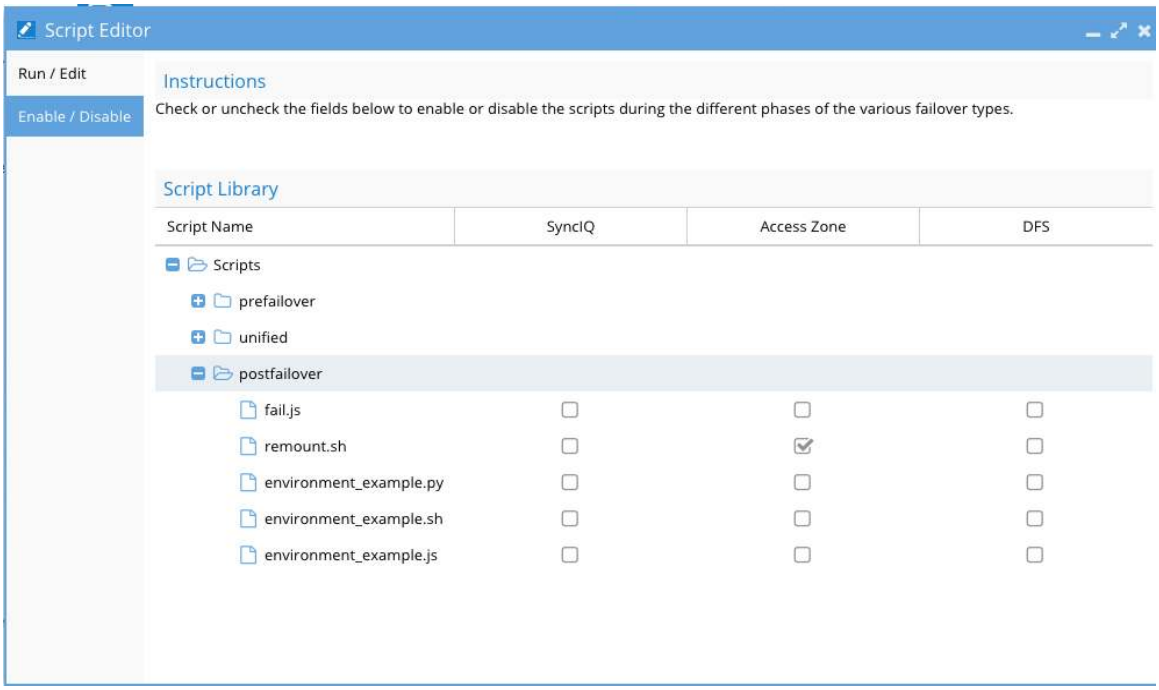

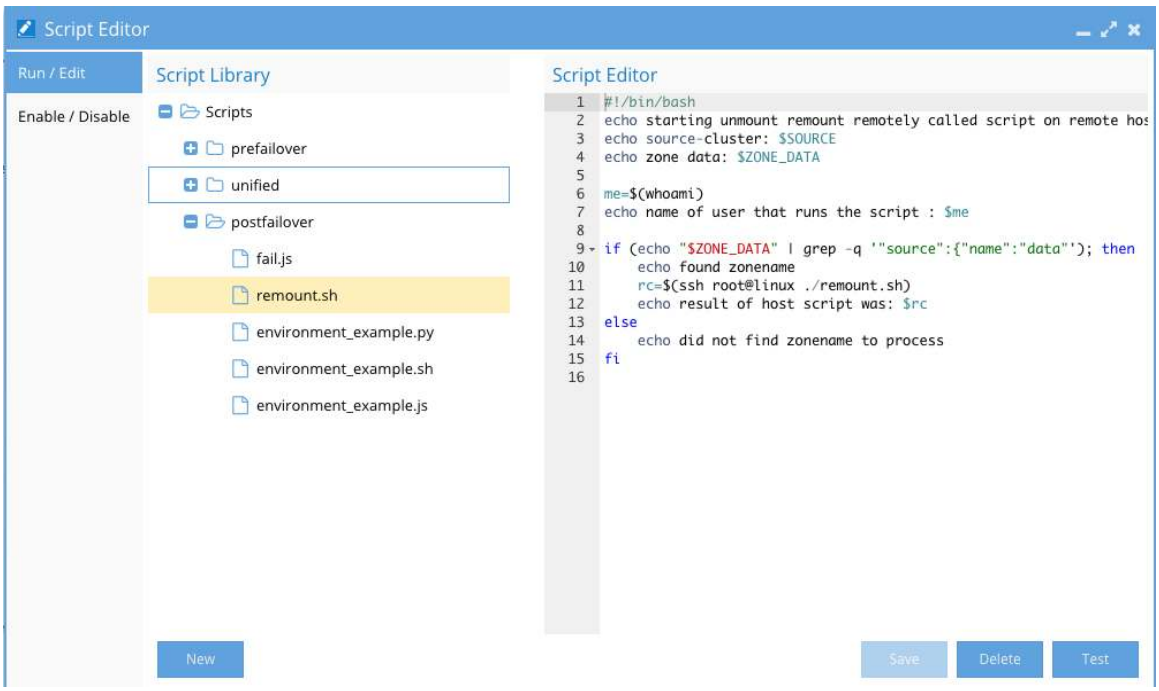

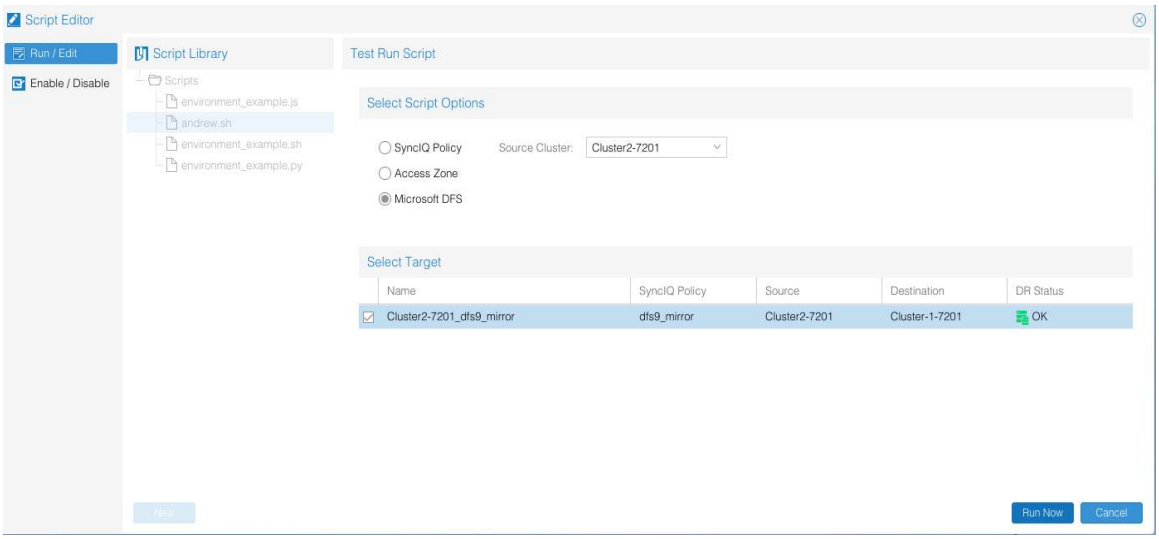

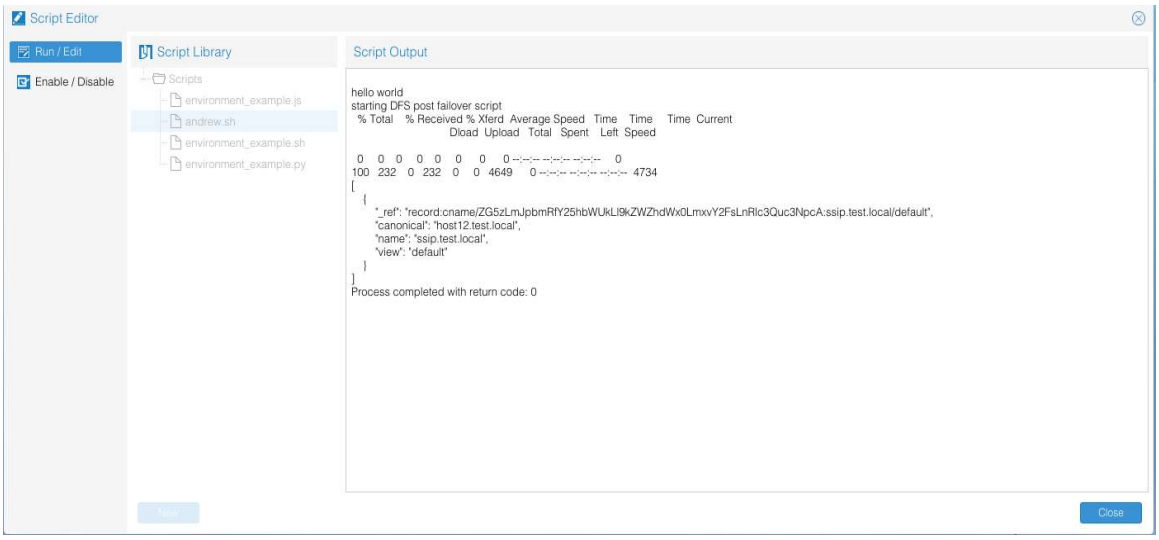

Script Engine CURL Tips

If you are using CURL as the method to automate with the API, **be aware** of the following:

- Curl -k will be required since the API is using a self signed certificate. This is not added to the CURL command with the CURL builder interface, and should be added to your CURL command.
- To avoid 411 response for content length, or body error from API server, add -d "" to the CURL command generated by the CURL builder interface of the API explorer.

Curl

```
curl -X POST --header 'Content-Type: application/json' --header 'Accept: application/json' --header 'api_key: igls-1cb6
```

Request URL

Script Engine Understanding Remote Execution to Hosts

Video: How to Remount Exports Automation

ssh Passwordless Login to Remote Linux Hosts

A common use case is running a script locally on Eyeglass to take an action on a remote Linux host, or running a command on a remote host to complete failover. When using ssh from a supported language, you can use Bash and .ssh keys to avoid passwords with follow these steps.

Note: Bash scripts run as the sca user, when they execute, this user also owns Eyeglass files and processes.

- Ssh to Eyeglass appliance.
- Sudo -s (to switch to root).
- Enter admin password.
- Cd /opt/superna (this is the home directory for the sca user used by Eyeglass processes).
- Create directory /opt/superna/.ssh/id_rsa.
- Type 'ssh-keygen -t rsa' do not set a password and accept all default prompts but enter a path of /opt/superna/.ssh/id_rsa.

- Now set ownership on files for remote execution all scripts run as the sca user:
- Su sca.
- Ssh user@remotehost (creates known_hosts file for target host, answer yes to accept ssh ID).
- Exit (you are now root again).
- Cd /opt/superna/.ssh.
- Chown sca *.
- Chgrp users *.
- ssh User@remotehost mkdir -p .ssh (User is the user that has access to the script that must execute, remotehost is dns or host name of remote linuxhost) - this will create .ssh if it does not already exist.
- cat /opt/superna/.ssh/id_rsa.pub | ssh User@remotehost 'cat >> .ssh/authorized_keys' (this places pub ssh keys into the remote users .ssh authorized keys file to allow passwordless login from a script).
- Enter password for User on remotehost.
- Test SCA remote ssh:
- Su sca .
- Ssh user@remotehost (if no pwd requested the setup is complete) .
- Done.

How to Install Windows SSH Server for remote Powershell execution from Eyeglass with SSH

Enable SSH on Windows Server

Tested on:
Windows server 2012 r2

Procedure:

1. Open Windows PowerShell
2. Change directory to the folder that we want to put the downloaded OpenSSH file
   Example: PS C:\Users\Administrator> cd C:\myfolder
1. We need to download the .zip file from the Github repository using the Invoke-WebRequest command. Change security protocol to TLS1.2 or TLS1.3 using the following command:
   PS C:\myfolder> [Net.ServicePointManager]::SecurityProtocol = [Net.SecurityProtocolType]::Tls12
2. Then download the binary using the Invoke-WebRequest:
   PS C:\myfolder> Invoke-WebRequest -Uri "https://github.com/PowerShell/Win32-OpenSSH/releases/download/v7.7.2.0p1-Beta/OpenSSH-Win64.zip" -OutFile "powershell.zip"
5. On a fresh installation, Windows 2012 R2 does not have the Expand-Archive command, so we will use .NET directly. Add-Type loads a .dll with the necessary .net functions in our current session. then [io.compression.zipfile] is a reference to that loaded .dll and ::ExtractToDirectory is the way to call a function from that dll :
   PS C:\myfolder> Add-Type -assembly "system.io.compression.filesystem"
6. Unzip the file
   PS C:\myfolder> [io.compression.zipfile]::ExtractToDirectory('C:\myfolder\powershell.zip','C:\myfolder')
7. Change into the directory that has been unzipped and launch the installation:
   PS C:\myfolder> cd .\OpenSSH-Win64
   PS C:\myfolder\OpenSSH-Win64> .\install-sshd.ps1
   [SC] SetServiceObjectSecurity SUCCESS
   [SC] ChangeServiceConfig2 SUCCESS
   [SC] ChangeServiceConfig2 SUCCESS
   sshd and ssh-agent services successfully installed
8. The following command will show the status of the SSHD service:
   PS C:\myfolder\OpenSSH-Win64> get-service | findstr ssh
   Stopped ssh-agent OpenSSH Authentication Agent
   Stopped sshd OpenSSH SSH Server

9. Start the service:
   PS C:\myfolder\OpenSSH-Win64> Start-Service sshd
   PS C:\myfolder\OpenSSH-Win64> Start-Service ssh-agent
10. If need to set the service to start automatically:
    PS C:\myfolder\OpenSSH-Win64> Set-Service -Name sshd -StartupType "Automatic"
    PS C:\myfolder\OpenSSH-Win64> Set-Service -Name ssh-agent -StartupType "Automatic"
11. If need to open firewall port for SSH:
    PS C:\myfolder\OpenSSH-Win64> netsh advfirewall firewall add rule name=SSHPort dir=in action=allow protocol=TCP localport=22

Test the SSH Server with Powershell:

1. SSH to Eyeglass as admin user
2. Run command: ssh *<user>@<windows-server-ip>*
3. Confirm the session establishment and provide the user password
4. Once logged in to the windows server via SSH, then run powershell command
5. Example:powershell.exe get-content env:computername

How to setup keyless ssh login for Eyeglass Script engine with Windows SSH and Powershell

1. ssh to Eyeglass as admin user
2. sudo su - (enter the password)
3. cd /opt/superna
4. mkdir .ssh
5. Run command : ssh-keygen -t rsa (enter the path of /opt/superna/.ssh/id_rsa and do not set a passphrase, accept all default prompts)
6. Change ownership of this .ssh directory: chown -R sca:users .ssh
7. su sca
8. ssh <user>@<windows-server-ip> (creates known_hosts file for target host, answer yes to accept ssh ID)
9. Exit ssh
10. Back to root: exit
11. ssh <user>@<windows-server-ip> mkdir -p .ssh (this will create .ssh if it does not already exist on windows server).
12. Exit ssh from windows server
13. cat /opt/superna/.ssh/id_rsa.pub (to get the content of the id_rsa.pub file)

14. Copy the content of id_rsa.pub file to .ssh/authorized_keys file on windows server (this places pub ssh keys into the remote users .ssh authorized keys file to allow passwordless login from a script). We can use text editor to copy the content of id_rsa.pub file to the authorized_keys file.
15. Test the SSH: su sca
16. ssh <user>@<windows-server-ip> (if no password requested the setup is complete)
17. Done

Access Zone Example Bash Script using ssh Keys to Remotely Execute a Command

The Access Zone Based failover preserves SmartConnect Zone names across failovers, which only requires an unmount and remount of the same FQDN SmartConnect Zone name. This means the failover logic can be used for failover or failback since it's' the same operation.
This sample solution uses a script that is unique on each host with the same name. Example; remount.sh placed in the user home directory used with the ssh remote execution rsa pub key. (See steps above to set up Eyeglass for ssh passwordless login to remote hosts).

Remote Host Script Example:

Script name **remount.sh,** placed in the home directory of the user account setup, for ssh login automation from the Eyeglass appliance:

```
#remount script
echo "remounting filesystem post failover"
umount -fl /mnt/data
mount -a
mount | grep "/mnt/data"
```

Eyeglass Access Zone Failover and Failback Script Example

```
#!/bin/bash
#
#  Script: remount.sh
# Purpose: unmount and remount from /etc/fstab persistent mounts post failover script, depends on remote script on remote host to execute
# Location: Eyeglass post failover scripting paste into script engine and enable for Access zone based failovers
```

#
echo starting unmount remount remotely called script on remote hosts
echo source-cluster: $SOURCE
echo zone data: $ZONE_DATA
me=$(whoami)
echo name of user that runs the script : $me
 # The variables set during failover include many variables and attributes of the Access Zone selected for failover that can be used to grep and select when # to apply failover logic. This can be used to group which hosts to automate failover based on the Access Zone selected for failover. The example below:
 # can be expanded to be used with per SyncIQ policy names using the same grep solution and variables shown in the Eyeglass echo example scripts.
 # This string "source":{"name":"xxxx" replace the xxxx with the name of the zone you want to failover, hint you can use test feature in script engine to
 # run the sample scripts with your clusters and access zones to see which string to grep for test see bolded section used below to select and access zone
 # zone data:
*#{"source":{**"name":"data"**,"subnets":[{"name":"subnet0","smartConnectServiceIp":"172.31.1.201","pools":[{"name":"subnet0:dfsdata","ranges":"172.31#.1.113-172.31.1.113","smartConnectZoneName":"dfsdata-dr.ad1.test","smartConnectAliases":["igls-ignore"]},{"name":"subnet0:userdata","ranges":"17#2.31.1.111-172.31.1.111","smartConnectZoneName":"userdata.ad1.test","smartConnectAliases":["igls-user-prod"]}]}],"target":{"name":"data","subnet#s":[{"name":"subnet0","smartConnectServiceIp":"172.31.1.200","pools":[{"name":"subnet0:dfsdata","ranges":"172.31.1.112-172.31.1.112","smartConn#ectZoneName":"dfsdata.ad1.test","smartConnectAliases":["igls-ignore"]},{"name":"subnet0:userdata","ranges":"172.31.1.110-172.31.1.110","smartCon#nectZoneName":"igls-original-userdata.ad1.test","smartConnectAliases":["igls-user-prod"]}]}]},"poolMap":[{"sourcePool":{"name":"subnet0:userdata","#ranges":"172.31.1.111-172.31.1.111","smartConnectZoneName":"userdata.ad1.test","smartConnectAliases":["igls-user-prod"]},"targetPool":{"name":"s#ubnet0:userdata","ranges":"172.31.1.110-172.31.1.110","smartConnectZoneName":"igls-original-userdata.ad1.test","smartConnectAliases":["igls-user-#prod"]}}]}*

if (echo "$ZONE_DATA" | grep -q '"source":{"name":"**data**"'); then

```
  echo found zonename
# remotely execute the remount.sh script on the remote host (NOTE: requires ssh pub keys from eyeglass on the remote host)
  rc=$(ssh root@linux ./remount.sh)
# remote script runs and returns output, can be output below to be captured in the failover log
  echo result of host script was: $rc
else
  echo did not find zonename to process
fi
```

Adding Additional Script Language Support to the Appliance

1. Nodejs can be added by: (https://nodejs.org/en/docs/)
    - ssh as root to the appliance .
    - then run **zypper install npm.**
    - answer yes (requires internet access).
    - bash - pre-installed in the OS (2.7.8).
    - Python - pre-installed in the OS (2.7.8).

Script Run Time Variables

The following variables can be used to pass in values to a script, to handle various policies or scenario's, using substitution of the values:

1. **FAILOVER_RAN** is true or false failover steps completed intended for single script that handles pre or post failover with a single script and needs to know if failover has happened or not.
    a. example pre script would stop application, unmount file system, post script logic would wait until failover has exercuted before mounting and start up the application.
    b. Only valid in a post failover script
    c. New in 2.5.6 update 2

2. **FAILOVER_STATUS** is "OK" "WARNING" or "ERROR" if your script logic needs to handle scenario's where the failover was not error free then use this variable.
    a. Only valid in a post failover script
    b. New in 2.5.6 update 2
3. **FAILOVER_SUCCESS** is true or false - this is a simple true or false without providing the specific error code the previous variable offers. This can be used to trigger logic after failover is successful versus a failure. This could be used to skip application steps on a failure to avoid mount issues or application startup problems.
    a. Only valid in a post failover script
    b. New in 2.5.6 update 2
4. **source** - Represents the name of the source cluster of the SyncIQ policy.
5. **target** - Represents the name of the target cluster of the SyncIQ policy.
6. **policy** - Used to return metadata about the policy itself see example output below (NodeJS - Example Output).
7. **failover_type** - SyncIQ, DFS or Access Zone.
8. **zone_data** - zone data about the Access Zone that can be used example SmartConnect Zone list and zone alias for DNS updates

Sample Execution Rules & Overall Failover Job Status Impact

1. They run after all Eyeglass automation.

2. One or more scripts can be enabled per failover type and both will execute in series during failover.
3. Return code provided by the script should return 0 to indicate the script had no errors and completed successfully.
4. Return code > 0 indicates an error.
5. Return codes can be set to any value and number and meaning in the script, Eyeglass takes no actions based on specific return codes.
6. Return codes are logged in the failover log for post failover review and debugging.
7. Script output is captured in the failover log. It is best practice to use the echo command to output script execution so that it's included in the Failover log.
8. If running two or more scripts each script should have discrete function to complete, AND should not have any dependency on other scripts. No ability to have IF script return code of X then 2nd script do Y exists.
9. Put host side script automation into its own script.
10. Put DNS automation logic into it’s own script.
11. Put application specific logic into it’s own script.
12. If any scripting logic needs dependant logic then a single script should be used for all functions.
13. Return code > 0 will failover the overall job status.

Sample Scripts in the Library

NodeJS - Example Script

```
#!/usr/bin/env node
console.log("these are the environment variables");
console.log("source", process.env.SOURCE);
console.log("target", process.env.TARGET);
console.log("type", process.env.FAILOVER_TYPE);
console.log("zone data", process.env.ZONE_DATA);
console.log("policy data", process.env.POLICY_DATA);
```

NodeJS - Example Output

```
These are the environment variables:
source {"pass":"password!","port":8080,"ip":"172.31.1.105","name":"Cluster2-7201","guid":"005056ba72edf6450c552312a728d3a22a23","user":"admin"}
```

target {"pass":"password!","port":8080,"ip":"172.31.1.104","name":"Cluster-1-7201","guid":"005056ba34580f410c55fd077989478a3821","user":"admin"}
type SYNCIQ
zone data
policy data
[{"name":"dfs9_mirror","targetIp":"172.31.1.104","targetHostname":"172.31.1.104","sourcePath":"/ifs/data/policy1","targetPath":"/ifs/data/policy1","enabled":true,"shares":[],"exports":[],"zones":[],"lastJobStatus":"running","lastSuccess":"null","lastStarted":"1446812101","schedule":"every 1 days every 5 minutes between 12:00 AM and 11:59 PM","sourceExludePaths":[],"sourceIncludePaths":[]}]
Process completed with return code: 0

Python - Example Script

```
#!/usr/bin/env python
import os
print "these are the environment variables"
print os.environ['SOURCE']
print os.environ['TARGET']
print os.environ['FAILOVER_TYPE']
print os.environ['ZONE_DATA']
print os.environ['POLICY_DATA']
```

Python - Example Output

these are the environment variables
{"pass":"password!","port":8080,"ip":"172.31.1.105","name":"Cluster2-7201","guid":"005056ba72edf6450c552312a728d3a22a23","user":"admin"}
{"pass":"password!","port":8080,"ip":"172.31.1.104","name":"Cluster-1-7201","guid":"005056ba34580f410c55fd077989478a3821","user":"admin"}
SYNCIQ
[{"name":"dfs9_mirror","targetIp":"172.31.1.104","targetHostname":"172.31.1.104","sourcePath":"/ifs/data/policy1","targetPath":"/ifs/data/policy1","enabled":true,"shares":[],"exports":[],"zones":[],"lastJobStatus":"finished","lastSuccess":"1446812701","lastStarted":"1446812701","schedule":"every 1 days every 5 minutes between 12:00 AM and 11:59 PM","sourceExludePaths":[],"sourceIncludePaths":[]}]
Process completed with return code: 0

Bash - Example Script

```
#!/bin/bash
```

```
echo these are the environment variables
echo source: $SOURCE
echo target: $TARGET
echo failover type: $FAILOVER_TYPE
echo zone data: $ZONE_DATA
echo policy data: $POLICY_DATA
```

Bash - Example Output

```
these are the environment variables
source: {"pass":"password!","port":8080,"ip":"172.31.1.105","name":"Cluster2-
7201","guid":"005056ba72edf6450c552312a728d3a22a23","user":"admin"}
target: {"pass":"password!","port":8080,"ip":"172.31.1.104","name":"Cluster-1-
7201","guid":"005056ba34580f410c55fd077989478a3821","user":"admin"}
failover type: SYNCIQ
zone data:
policy data:
[{"name":"dfs9_mirror","targetIp":"172.31.1.104","targetHostname":"172.31.1.104","sour
cePath":"/ifs/data/policy1","targetPath":"/ifs/data/policy1","enabled":true,"shares":[],"exp
orts":[],"zones":[],"lastJobStatus":"running","lastSuccess":"null","lastStarted":"14468127
01","schedule":"every 1 days every 5 minutes between 12:00 AM and 11:59
PM","sourceExludePaths":[],"sourceIncludePaths":[]}]
Process completed with return code: 0
```

Consolidated Post Failover & Failback Script (Node.JS)

```
#!/usr/bin/env node
var exec = require('child_process').exec;
var child;
var mycmd = 'echo
41d7297b7c79651bb94dcf676538f9b3b5ed6e8ed25e04c6ee38d14269e022cc | sudo -
S su root -c "sh /opt/superna/sca/failover.sh"';
var mycmd2 = 'echo
41d7297b7c79651bb94dcf676538f9b3b5ed6e8ed25e04c6ee38d14269e022cc | sudo -
S su root -c "sh /opt/superna/sca/failback.sh"';
// refresh name resolution
if (process.env.SOURCE.indexOf('cluster20') !== -1)
{
console.log("Failover");
child = exec(mycmd, function (error, stdout, stderr)
```

```
{
console.log('result output: ' + stdout);
console.log('result errors: ' + stderr);
}
);
}
else
{
   console.log("Failback");
   child = exec(mycmd2, function (error, stdout, stderr)
{
console.log('result output: ' + stdout);
console.log('result errors: ' + stderr);
}
);
}
var node_ssh = require('node-ssh');
var ssh = new node_ssh();
var cmd1 = "ls -l /proc/*/cwd | grep /mnt/z01-nfs01 | awk '{print $9}' | grep -o '[0-9]*' |
xargs kill -s 9",
   cmd2 = 'umount -fl /mnt/z01-nfs01',
   cmd3 = 'mount -t nfs -o vers=3 cluster20-z01.ad1.test:/ifs/data/zone01/z01-nfs01
/mnt/z01-nfs01',
   host = '172.16.81.161',
   user = 'root',
   pass = 'GoSuperna!';
console.log('Executing command: ' + cmd1);
console.log('Executing command: ' + cmd2);
console.log('Executing command: ' + cmd3);
console.log(' On host: ' + host);
ssh.connect({
  host: host,
  username: user,
  password: pass
}).then(function() {
   ssh.execCommand( cmd1, {
       stream: 'both'
   }).then(function(result) {
```

```
        console.log('result output: ' + result.stdout);
        console.log('result errors: ' + result.stderr);
}).then(function() {
   ssh.execCommand( cmd2, {
        stream: 'both'
   }).then(function(result) {
        console.log('result output: ' + result.stdout);
        console.log('result errors: ' + result.stderr);
   }).then(function() {
   ssh.execCommand( cmd3, {
        stream: 'both'
   }).then(function(result) {
        console.log('result output: ' + result.stdout);
        console.log('result errors: ' + result.stderr);
        console.log('ssh operation complete');
        ssh.end();
});
});
   });
});
```

Consolidated Post Failover & Failback Script (Node.JS) I.E Multiple Zone Access

The following node.js script is the example for handling multiple zone access:

```
// 1st Part refresh name resolution
var exec = require('child_process').exec;
var child;
var mycmd1 = 'echo
41d7297b7c79651bb94dcf676538f9b3b5ed6e8ed25e04c6ee38d14269e022cc | sudo -
S su root -c "sh /opt/superna/sca/failover-z01.sh"';
var mycmd2 = 'echo
41d7297b7c79651bb94dcf676538f9b3b5ed6e8ed25e04c6ee38d14269e022cc | sudo -
S su root -c "sh /opt/superna/sca/failover-z03.sh"';
var mycmd3 = 'echo
41d7297b7c79651bb94dcf676538f9b3b5ed6e8ed25e04c6ee38d14269e022cc | sudo -
S su root -c "sh /opt/superna/sca/failback-z01.sh"';
```

```
var mycmd4 = 'echo
41d7297b7c79651bb94dcf676538f9b3b5ed6e8ed25e04c6ee38d14269e022cc | sudo -
S su root -c "sh /opt/superna/sca/failback-z03.sh"';
if (process.env.SOURCE.indexOf('cluster20') !== -1)
{
  console.log("Failover");
  if (process.env.ZONE_DATA.indexOf('zone01') !== -1)
  {
  child = exec(mycmd1, function (error, stdout, stderr)
  {
  console.log('result output: ' + stdout);
  console.log('result errors: ' + stderr);
  }
  );
  }
  if (process.env.ZONE_DATA.indexOf('zone03') !== -1)
  {
  child = exec(mycmd2, function (error, stdout, stderr)
  {
  console.log('result output: ' + stdout);
  console.log('result errors: ' + stderr);
  }
  );
  }
}
if (process.env.SOURCE.indexOf('cluster21') !== -1)
{
  console.log("Failback");
  if (process.env.ZONE_DATA.indexOf('zone01') !== -1)
  {
  child = exec(mycmd3, function (error, stdout, stderr)
  {
  console.log('result output: ' + stdout);
  console.log('result errors: ' + stderr);
  }
  );
  }
  if (process.env.ZONE_DATA.indexOf('zone03') !== -1)
```

```
  {
  child = exec(mycmd4, function (error, stdout, stderr)
  {
  console.log('result output: ' + stdout);
  console.log('result errors: ' + stderr);
  }
  );
  }
}
// 2nd Part refresh mount
var node_ssh = require('node-ssh');
var ssh = new node_ssh();
var cmd1 = "ls -l /proc/*/cwd | grep /mnt/z01-nfs01 | awk '{print $9}' | grep -o '[0-9]*' |
xargs kill -s 9",
   cmd2 = 'umount -fl /mnt/z01-nfs01',
   cmd3 = 'mount -t nfs -o vers=3 cluster20-z01.ad1.test:/ifs/data/zone01/z01-nfs01
/mnt/z01-nfs01',
   cmd4 = "ls -l /proc/*/cwd | grep /mnt/z03-nfs01 | awk '{print $9}' | grep -o '[0-9]*' | xargs
kill -s 9",
   cmd5 = 'umount -fl /mnt/z03-nfs01',
   cmd6 = 'mount -t nfs -o vers=4 cluster20-z03.ad1.test:/ifs/data/zone03/z03-nfs01
/mnt/z03-nfs01',
   host = '172.16.81.161',
   user = 'root',
   pass = 'GoSuperna!';
console.log(' On host: ' + host);
if (process.env.ZONE_DATA.indexOf('zone01') !== -1)
  {
   console.log('Executing command: ' + cmd1);
   console.log('Executing command: ' + cmd2);
   console.log('Executing command: ' + cmd3);
   ssh.connect({
       host: host,
       username: user,
       password: pass
   }).then(function() {
   ssh.execCommand( cmd1, {
       stream: 'both'
```

```
  }).then(function(result) {
      console.log('result output: ' + result.stdout);
      console.log('result errors: ' + result.stderr);
  }).then(function() {
  ssh.execCommand( cmd2, {
      stream: 'both'
  }).then(function(result) {
      console.log('result output: ' + result.stdout);
      console.log('result errors: ' + result.stderr);

  }).then(function() {
  ssh.execCommand( cmd3, {
      stream: 'both'
  }).then(function(result) {
      console.log('result output: ' + result.stdout);
      console.log('result errors: ' + result.stderr);
      console.log('ssh operation complete');
      ssh.end();
  });
});
  });
  });
}
if (process.env.ZONE_DATA.indexOf('zone03') !== -1)
  {
  console.log('Executing command: ' + cmd4);
  console.log('Executing command: ' + cmd5);
  console.log('Executing command: ' + cmd6);
  ssh.connect({
      host: host,
      username: user,
      password: pass
  }).then(function() {
  ssh.execCommand( cmd4, {
      stream: 'both'
  }).then(function(result) {
      console.log('result output: ' + result.stdout);
      console.log('result errors: ' + result.stderr);
```

```
  }).then(function() {
  ssh.execCommand( cmd5, {
      stream: 'both'
  }).then(function(result) {
      console.log('result output: ' + result.stdout);
      console.log('result errors: ' + result.stderr);

  }).then(function() {
  ssh.execCommand( cmd6, {
      stream: 'both'
  }).then(function(result) {
      console.log('result output: ' + result.stdout);
      console.log('result errors: ' + result.stderr);
      console.log('ssh operation complete');
      ssh.end();
  });
});
  });
  });
}
```

Superna Eyeglass API Guide

It is now possible to failover using external applications such as VMware SRM or a script called from an application, web page or CURL command. The API Guide covers API Explorer to automatically build CURL commands that allows a single command failover over a policy or entire Access Zone. This also allows script engine logic to run if enabled at the end of failover.

The API and example VMware integration for failover is explained in the Superna Eyeglass API guide.

© Superna LLC

1.28. Eyeglass Backup and Restore

Home Top

Eyeglass Backup and Restore

Technical Note

- Abstract:
- Backup Procedure:
- Restore Procedure:

Abstract:

This technical note provides the Eyeglass backup/restore process that can be used to move to a new appliance and restore license keys that were applied to the previous appliance.

Eyeglass Backup and Restore

Note: Eyeglass backup/restore process that can be used to move to a new appliance and restore license keys that were applied to the previous appliance. This is the only supported method to move to a new appliance. **If a new appliance is deployed and the previous appliance was not backed up, sales@superna.net must validate your licenses. Sales is available Monday to Friday 9-5 EDT Customers should backup appliances before deleting them from vCenter! NOTE: Support is unable to assist with license keys and they do not have access any license key system to assist. They will mark case as solved and hand off to sales to follow up with you.**

NOTE: Daily automated backup with 7 days of backups is stored here /opt/superna/var/backup these backups can be copied from the appliacne for a restore or remount this path to external storage on PowerScale with NFS store a copy of the backup external to the appliance.

Backup

Backup Procedure:

To backup your Eyeglass appliance data, configuration and Licenses:

1. Login to the Eyeglass web page.
2. Open the About Eyeglass window.
3. Select the Backup menu.
4. 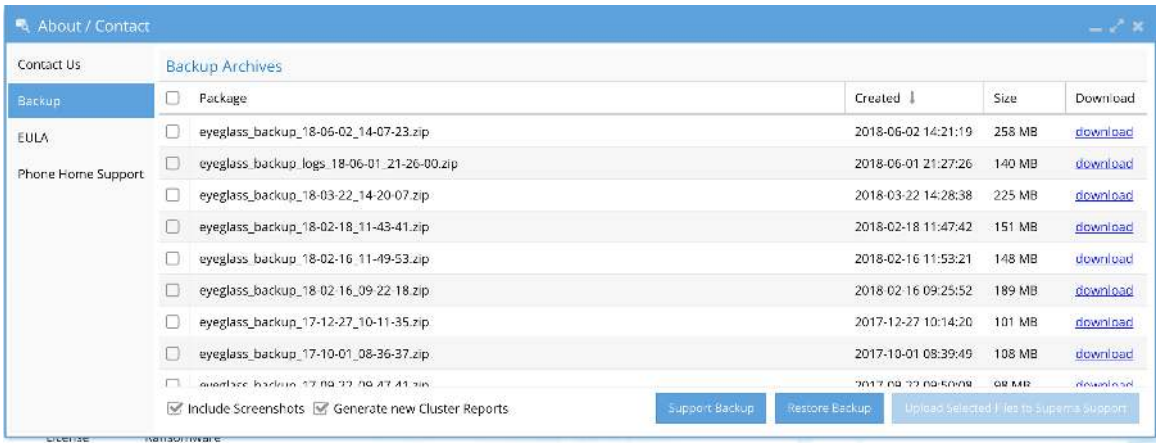

5. In the Backup window select the **Create New Restore Backup** button. **NOTE: Only Restore Backup should be used so that all private keys are restored for SSL. The Support backup does not include any private keys needed for SSL certificates.**
6. While the Backup is in progress you will see the “Creating backup archive” message.
7. 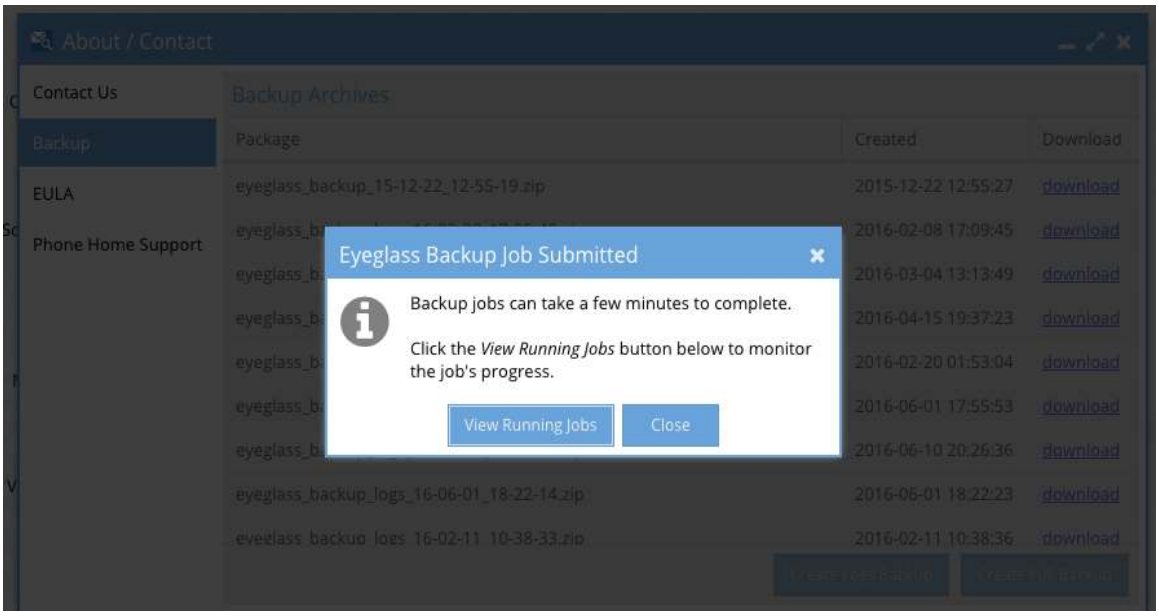

8. 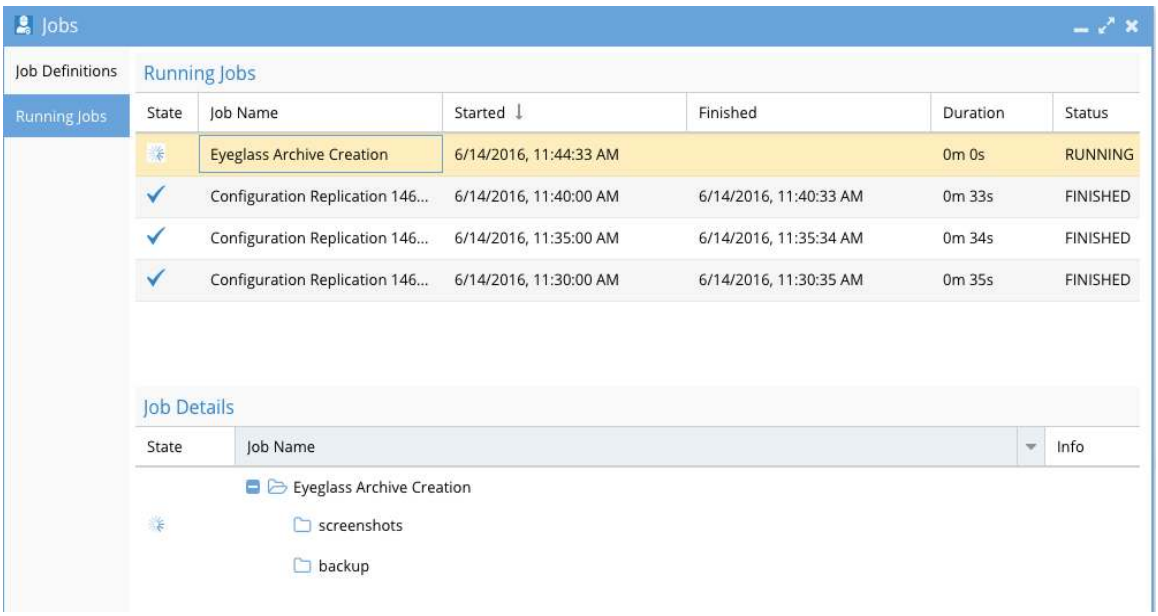

9. Once completed you will see the archive in the list.
10. 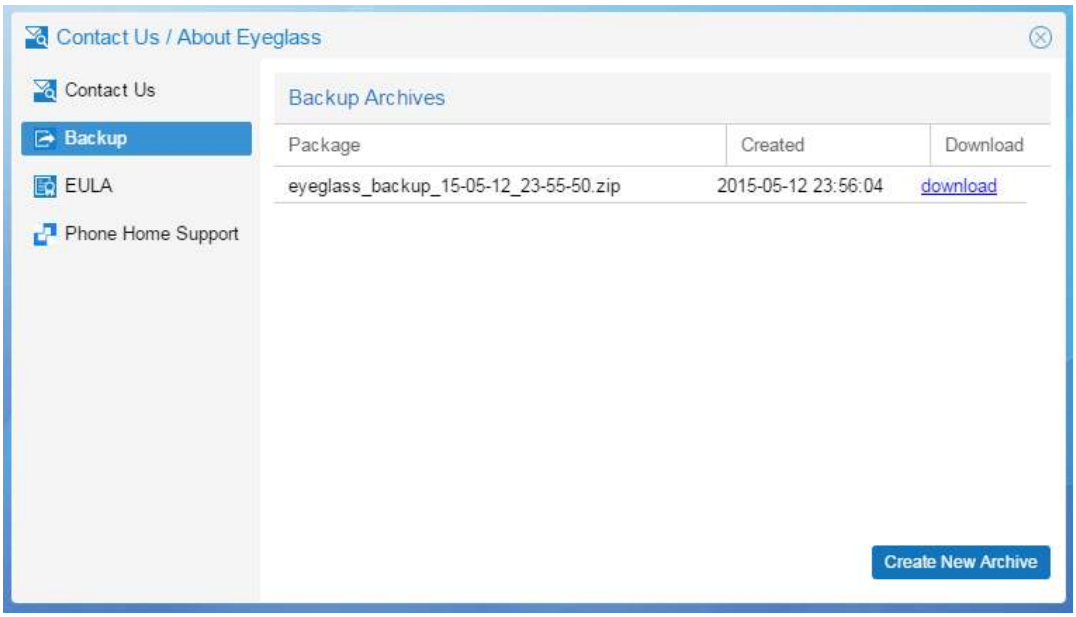

11. Select the download link to download the archive to your local computer.
12. Store the archive in a secure location.

Restore

IMPORTANT!

1. Eyeglass version at the time the backup was taken must match the Eyeglass version to which the backup is being restored

2. There will be an Eyeglass service interruption while the restore is in progress

Restore Procedure:

To restore Eyeglass data and configuration from an Eyeglass Restore Backup Archive file:

1. Take an Eyeglass full backup from your old Eyeglass appliance.
2. Download the full backup locally, and then copy the zip file backup, using scp or winscp to the newly deployed Eyeglass Appliance. It could be placed in /tmp for example **(or use one of the daily backups located here on the previous appliance /opt/superna/var/backup).**
3. Power off the old Eyeglass appliance. It is not supported to have multiple Eyeglass appliances managing the same clusters.
4. SSH to new Eyeglass appliance and login as admin (default password 3y3gl4ss). Issue “sudo su -” to enter in root mode (default password 3y3gl4ss).
5. From the command line execute one of the restore commands depending on the scenario:
    a. **Scenario #1 - Different version restore between the backup file and the new appliance - limited restore of settings.**
        i. igls app restore /tmp/<eyeglass_backup.xxxx.zip> --anyrelease **(This flag allows different releases between the backup and the new appliance. Note this will only restore the clusters login credentials and license keys, all other history is lost with this flag. This option should only be used if you are prepared to re-configure the new appliance jobs and schedules.)**
    b. **Scenario #2 - Same version restore between the backup file and the new appliance - full support for settings restore.**
        i. igls app restore /tmp/<eyeglass_backup.xxxx.zip> **(NOTE: must use the same version of Eyeglass**

backup and new appliance version must match to restore all information, see upgrade guide for details on what is restored)

c. **Scenario #3 - Restore settings to the same appliance - full support for settings restore.**
    i. igls app restore /opt/superna/var/backup **(this will display the most recent backup found on the internal daily backup path, and ask to confirm to use the most recent file to restore)**
    ii. Requires 2.5.6 or later
d. Replacing /tmp/<eyeglass_backup.xxxx.zip> with the name of the Eyeglass Archive file always including full path.
e. You will be prompted to continue. Enter “y” to continue.
f. Monitor the restore process for any errors.

6. Login to the new Eyeglass appliance and check if:
    a. Licenses have been added.
    b. Clusters have been added.
7. If any release flag is used, all job types will be user disabled, and will need to be configured along with task schedules configured on the old appliance.
1.

c.

d.

e.

f.

g.

© Superna LLC

1.29. Eyeglass API guide

Home Top

- Overview
- What's New with Eyeglass API
- API Documentation
- Architecture
- Use Cases
- How to video Overview of Eyeglass API
- How to Configure API Tokens
- Creating Tokens
- Deleting Tokens
- Launch API Explorer
- Token Authentication
  - Usage
- Full Version 2 API Routes Available
- job (v2) : Starting and stopping jobs.
  - GET /v2/jobs/failover
  - POST /v2/jobs/failover
  - POST /v2/jobs/failover/drtest
  - POST /v2/jobs/failover/rehearsal
  - DELETE /v2/jobs/failover/{id}
  - GET /v2/jobs/failover/{id}
  - GET /v2/jobs/failover/{id}/log
  - GET /v2/jobs/readiness
  - POST /v2/jobs/readiness (New in 2.5.7)
  - GET /v2/jobs/readiness/{id}
  - GET /v2/jobs/replication (New in 2.5.7)
  - POST /v2/jobs/replication (New in 2.5.7)

    - GET /v2/jobs/replication/{id} (New in 2.5.7)
- Full Version 1 API routes available
- alarms : Retrieve information about alarms
    - GET /alarms/active
    - GET /alarms/historical
- healthcheck
    - GET /healthcheck
- job : Starting and stopping failover jobs.
    - GET /jobs
    - POST /jobs
    - GET /jobs/{id}
    - GET /jobs/{id}/log
- node : Retrieve information about managed devices, access zones, and SyncIQ policies.
    - GET /nodes
    - GET /nodes/{id}
    - GET /nodes/{id}/policies
    - GET /nodes/{id}/policies/{name}
    - GET /nodes/{id}/pools
    - GET /nodes/{id}/pools/{name}
    - GET /nodes/{id}/zones
    - GET /nodes/{id}/zones/{name}
- POST jobs to Create Failovers
    - POST /jobs
- Create DR LiveOPS Failover Jobs
    - POST /jobs/drtest
- Create DR Rehearsal Failover Jobs (available in a Future Release)
    - POST /jobs/rehearsal

- GET /healthcheck
- GET /Alarms
- GET /Alarmhistory
- GET /jobs
  - Responses
- POST /jobs
- GET /jobs/{id}
- DELETE /jobs/{id}
- GET /nodes
- GET /nodes/{id}
- How to retrieve health check data and all validations for Access Zones, pools, DFS policies and Synciq policies
- GET /nodes/{id}/policies
- GET /nodes/{id}/pools
- GET /nodes/{id}/pools/{name}
- GET /nodes/{id}/zones
- ErrorModel
- Job
- Node
- Policy
- Zone
- Inline_response
- Job_failoverTarget
- Example 1: CMDB Integration API
- SyncIQ protection of shares and exports in the XML output Description
- Quota output in the XML file Description
- Asset Management output in the XML file Description
- How to Integrate with a CMDB

- Sample CMDB XML file for 4 clusters
- Example 2: VMWare SRM Integrated Failover Example
- Example 3: Initiate an Access Zone failover
  - Failover Setting Parameters - Eyeglass REST API
  - Eyeglass REST API - Failover Parameters Rules
  - Step 1: Get all of the PowerScale Clusters provisioned in superna eyeglass:
  - Step 2: Using the ID from step 1, get all of the access zones on the cluster:
  - Step 3: Initiate a failover
  - Step 4: Monitor the failover
- Example 4: Get all currently running failover jobs
  - Step 1: use the ?state=running query to filter by running jobs:
- Example 5: Get a historical record of all previous failover jobs
  - Step 1: use the ?state=finished query to filter by completed jobs:
- Ransomware Defender Airgap API's
  - GET /v2/jobs/airgap
  - POST /v2/jobs/airgap
  - POST /v2/jobs/airgap/status
- Ransomware Defender Smart Airgap API
  - GET /v1/ransomware/rswevents (this is for Ransomware Defender only)
  - GET /v1/securityevents (this is the smart airgap api that allows a filter of all or Ransomware Dender or Easy Auditor active alerts only events)
- Known Issues
  - T783 Rest API delete job error

- **T15042 Rest API policy readiness is missing output for Target Reachability check**
- **T10935 Pool failover "failovertarget" field must be "zone id"**
- **T15623 Pool failover API does not support multiple pool selection**
- **T15624 Failover API does not block controlled failover when source cluster unreachable**
- **T17428 API Policy Readiness returns incorrect Access Zone**

- **Known Limitations**
  - **T18079 Change Eyeglass Configuration Replication Job Disable/Enable API must be done at same time as Job Type Change**

Overview

Customers requiring end to end integrated failover of compute layer and dependant storage used by customer or external 3rd party applications like VMware SRM, can use Superna Eyeglass REST API to programmatically interact with Superna Eyeglass.

What's New with Eyeglass API

1. 2.5.7
    a. REST API for Eyeglass Automation
        i. versioned API support
        ii. New API's

1. Run a configuration job on demand and get the status
2. Run a DR Readiness job and get the status
3. Set newly discovered synciq Policy job type (auto, auto DFS or skip config)
4. Enable or disable Config jobs

2. In 2.5.6 or later the following api updates have been added:
    a. LiveOps DR Test mode enable and disable REST API has been added
    b. IP pool failover has been added
    c. All DR Assistant options in the GUI have been added to the API
    d. Eyeglass appliance health check API has been added
    e. DR Dashboard Access Zone, IP Pool, DFS mode and SyncIQ mode readiness information is now available in the API. Use the foreadiness = true flag on policy, pool and zone API calls to return full DR validations for each failover target.
    f. Retrieve the failover log using the rest API see API here.
    g. API to create DR Rehearsal jobs will be available in a future release.
    h. 2.5.6 patch 1
        i. Issues addressed, target cluster environment variable in script editor for DFS mode scripts now populates with correct values.

ii. REST API will accept a comma separated list of failover target ID's, allowing policy and DFS failovers to group multiple targets into a single failover job, which would be used for bulk fast failover of multiple targets.

iii. 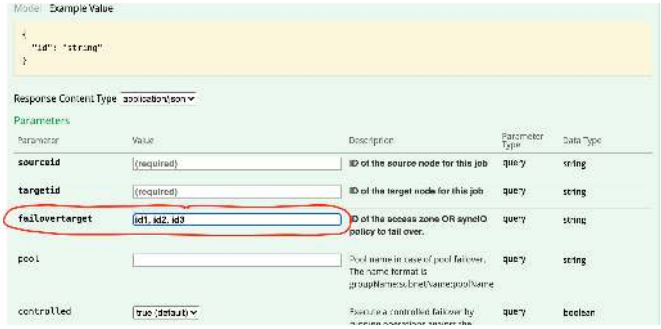

API Documentation

This information can be browsed with API Explorer using the url https://<x.x.x.x of eyeglass>/sera/docs/

Architecture

The API is REST based and Eyeglass provides an API explorer to explore the policies, nodes, jobs, needed to build failover scripts. The API explorer has a CURL builder feature. This feature allows users to browse the API select policies, or Access Zones and clusters, for failover automation, and have the CURL command built dynamically for simple cut and paste into a script file to execute failover.
The API uses token generation to identify and authenticate an application using the the failover API. This information is logged to trace failover requests back to an authorized script or user in the failover logs.

Use Cases

1. Polling for alarms in Eyeglass using alarm REST interface (new in 2.0 or later). See Get alarms and alarm history interfaces.
2. This can be used with CURL or REST to return active alarms Eyeglass only, not PowerScale alarms.
3. Build dashboards that monitor the readiness of policies, Access Zones or entire clusters. The API can be used in any development language that supports REST to build monitoring dashboards for end to end applications..
4. Request failover and monitor the success of the failover.
5. Build custom failover of one or more policy.
6. Integrate with post failover scripting solution that executes Eyeglass script engine post failover, and call this logic with the same API for failover. This enables application specific scripting to be initiated from the API, and all execution runs on the Eyeglass appliance.

How to video Overview of Eyeglass API

How to Configure API Tokens

Tokens are used for authentication of REST commands. The name parameter is used to identify the application using the token. Example: VMware SRM could have a token called “SRM”. Revoking a token also disables an application's access to Eyeglass.
Reference this Screenshot

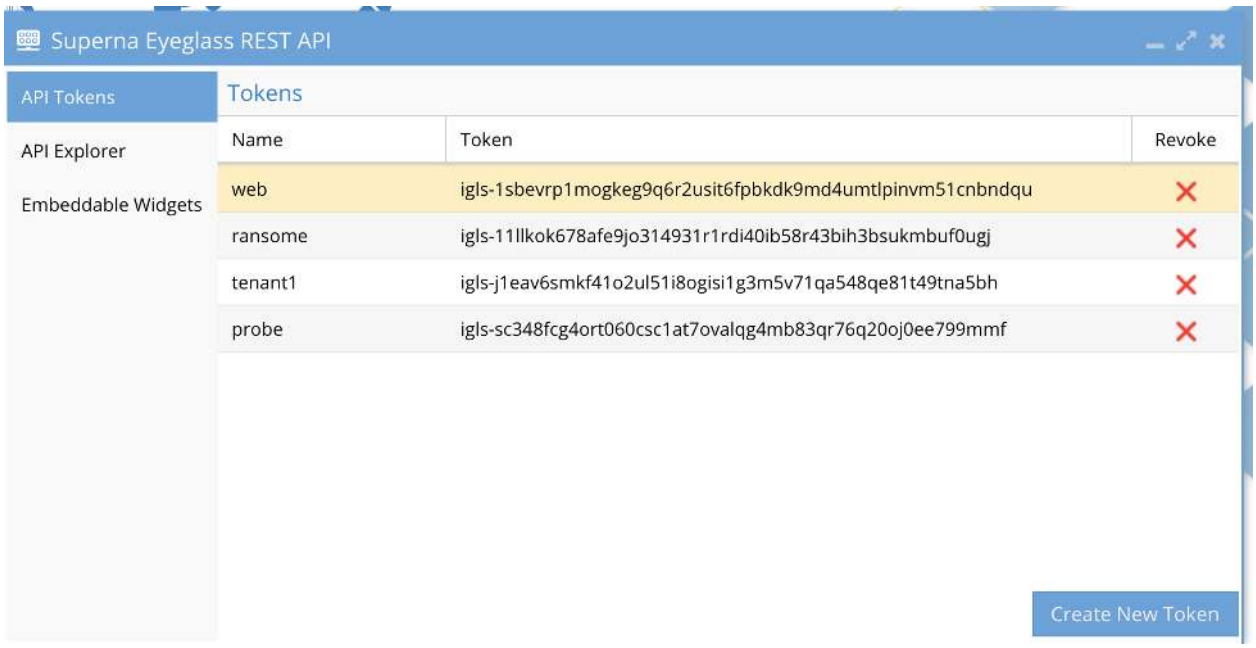

Creating Tokens

1. Open the Superna Eyeglass REST API icon on the main menu (Requires Enterprise license keys)
2. Click on Create New Token button and enter a name to reference the token

Deleting Tokens

1. Open the Superna Eyeglass REST API icon on the main menu (Requires Enterprise license keys)
2. Click the Revoke Red X to remove the token and block access to Eyeglass

Launch API Explorer

1. Click API Explorer tab and launch
2. The interface requires an API key to live view configuration API information to build CURL API

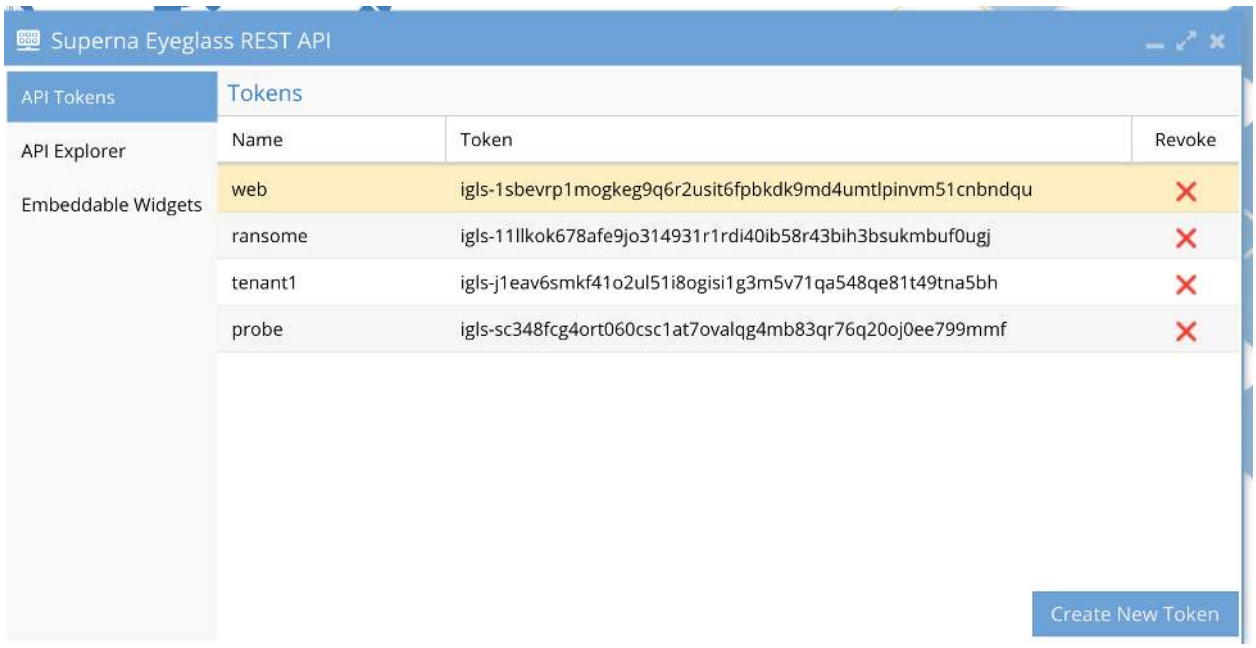

Watch this how to video to build a CURL command builder to failover Video

Token Authentication

Use the token that was created with CURL or development languages to authentication the application to Eyeglass. The API Explorer can be used to build curl command and shows syntax or how to contract an API call.

API Token Authentication in header
Include an http header named api_key with an authorized API token to authenticate to the Superna Eyeglass REST API.

Usage

The following curl command with the API token named igls-abc123 demonstrates authentication to the Superna Eyeglass REST API.

- curl --header "api_key: igls-abc123" --header "accept: application/json" https://eyeglass.example/sera/v1/jobs

Full Version 2 API Routes Available

Version 2 API provides new features described in What's New in this guide. Below is a list of all routes available, the new routes are listed with (New in 2.5.7)

- job (v2) : Starting and stopping jobs.
  - Show/Hide
  - List Operations
  - Expand Operations
    - GET /v2/jobs/failover
      - Get failover jobs
    - POST /v2/jobs/failover
      - Create a new failover job
    - POST /v2/jobs/failover/drtest
      - Enter/Exit DR test mode
    - POST /v2/jobs/failover/rehearsal
      - Create a new rehearsal job
    - DELETE /v2/jobs/failover/{id}

- GET /v2/jobs/failover/{id}
  - Retreive a failover job by ID
  -
- GET /v2/jobs/failover/{id}/log
  - Retrieve the logfile for a running or finished failover job
  -
- GET /v2/jobs/readiness
  - View all recent readiness jobs
  -
- POST /v2/jobs/readiness (New in 2.5.7)
  - Run zone readiness job
  -
- GET /v2/jobs/readiness/{id}
  - Retrieves a specific recently run readiness job, if it exists
  -
- GET /v2/jobs/replication (New in 2.5.7)
  - Get all recent replication jobs
  -
- POST /v2/jobs/replication (New in 2.5.7)
  - Run a configuration replication job
  -
- GET /v2/jobs/replication/{id} (New in 2.5.7)
  - Retrieves a specific replication job, if it was run recently

Full Version 1 API routes available

- alarms : Retrieve information about alarms
  - Show/Hide
  - List Operations
  - Expand Operations
    - GET /alarms/active
      - Get all active alarms
    - GET /alarms/historical
      - Get all historical alarms
- healthcheck
  - Show/Hide
  - List Operations
  - Expand Operations
    - GET /healthcheck
      - Get latest health-check timestamp
- job : Starting and stopping failover jobs.
  - Show/Hide

- List Operations
- Expand Operations
  -
- GET/jobs
  - Get jobs
  -
- POST/jobs
  - Create a new job
  -
    -
  -
- GET/jobs/{id}
  - Retrieve a job by ID
  -
- GET/jobs/{id}/log
  - Retrieve the logfile for a running or finished job

node : Retrieve information about managed devices, access zones, and SyncIQ policies.

- Show/Hide
- List Operations
- Expand Operations
  -
- GET/nodes
  - Get all nodes
  -

- GET/nodes/{id}
    - Find nodes by ID
-
- GET/nodes/{id}/policies
    - Find policies for a node
-
- GET/nodes/{id}/policies/{name}
    - Find policy by name
-
- GET/nodes/{id}/pools
    - Find pools for a node
-
- GET/nodes/{id}/pools/{name}
    - Find pool by name
-
- GET/nodes/{id}/zones
    - Find zones for a node
-
- GET/nodes/{id}/zones/{name}
    - Find zone name for a node

POST jobs to Create Failovers

NOTE: 2.5.6 patch 1 build will allow multi select of SyncIQ policy or DFS mode policies with a comma separated list in the source target field. See example below:

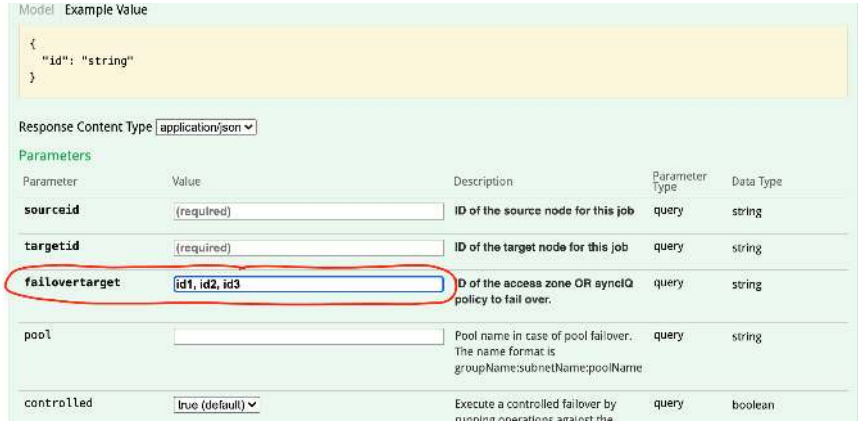

POST/jobs

- Create a new job

Implementation Notes

Launch a new job in Eyeglass.

Response Class (Status 201)

create job response

- Model
- Example Value

```
{   "id": "string" }
```

Response Content Type application/json

Parameters

Parameter	Value	Description	Parameter Type	Data Type
sourceid		ID of the source node for this job	query	string
targetid		ID of the target node for this job	query	string
failovertarget		ID of the access zone OR syncIQ policy to fail over.	query	string
pool		Pool name in case of pool failover. The	query	string

		name format is groupName:subnetName:poolName		
controlled	true (default) false	Execute a controlled failover by running operations against the source cluster as well as the target	query	boolean
datasync	true (default) false	Run the final incremental data sync before failover	query	boolean
configsync	true false (default)	Run a configuration sync before failover	query	boolean
resyncprep	true (default) false	Run resync prep on the source cluster to create the mirror policies	query	boolean
disablemirror	true false (default)	Disable mirror policies created on the failover target	query	boolean
quotasync	true (default) false	Run quota jobs to failover quotas to target	query	boolean
blockonwarnings	true (default) false	**Block failover on warnings**	query	boolean
rollbackrenameshares	true (default) false	Rollback renamed shares on failure	query	boolean
smbdataintegrity	true false (default)	SMB data integrity failover	query	boolean

Response Messages

HTTP Status Code	Reason	Response Model	Headers
default	error payload	Model Example Value `{ "code": 0, "message": "string"`	

Create DR LiveOPS Failover Jobs

POST/jobs/drtest

- Enter/Exit DR test mode

Implementation Notes

Enter/Exit DR test mode for a given policy

Response Class (Status 201)

create job response

- Model
- Example Value

```
{   "id": "string" }
```

Response Content Type application/json

Parameters

Parameter	Value	Description	Parameter Type	Data Type
enable	true (default) false	**True = Make target writable (Enter DR test mode). False = Make target read-only (Exit DR test mode)**	query	boolean
configsync	true false (default)	Run a configuration while DR test job	query	boolean
datasync	true (default) false	Run policy while DR test job	query	boolean
policy		DR testing policy id (as retrieved with /nodes/{id}/policies GET)	query	string

Response Messages

HTTP Status Code	Reason	Response Model	Headers
default	error payload	Model Example Value `{ "code": 0, "message": "string" }`	

Create DR Rehearsal Failover Jobs (available in a Future Release)

POST/jobs/rehearsal

- Create a new rehearsal job

Implementation Notes

Launch a new rehearsal job in Eyeglass.

Response Class (Status 201)

create rehearsal job response

- Model
- Example Value

```
{   "id": "string" }
```

Response Content Type application/json

Parameters

Parameter	Value	Description	Parameter Type	Data Type
enable	true (default) false	**Enable or disable (=false) the rehearsal mode**	query	boolean
sourceid		ID of the source node for this job	query	string
targetid		ID of the target node for this job	query	string
failovertarget		ID of the access zone OR DFS syncIQ policy to fail over.	query	string
pool		Pool name in case of pool failover. The name format is groupName:subnetName:poolName	query	string

Response Messages

HTTP Status Code	Reason	Response Model	Headers
default	error payload	Model Example Value	

```
{   "code": 0,   "message": "string" }
```

Methods

GET /healthcheck

Implementation Notes

Returns latest health-check timestamp from Superna Eyeglass.

https://igls/sera/v1/healthcheck

GET /Alarms

Use this example and replace API token to retrieve a list of active alarms.
This API call consumes the following media types via the Content-Type request header:

- application/json

Returns

- application/json

Query parameters for both routes;

since - epoch time of the earliest alarm to return.
until - epoch time of the latest alarm to return.
limit - the maximum number of alarms to fetch.
for example, we might query:

/sera/v1/alarms/historical?since=1499189000&until=1499190943&limit=50

CURL Example

```
curl  -X GET --header 'accept:application/json --header 'api_key:
igls-11llkok678afe9jo314931r1rdi40ib58r43bih3bsukmbuf0ugj'
'https://x.x.x.x/sera/v1/alarms/active?limit=20' -k
```

GET /Alarmhistory

Use this example and replace API token to retrieve a list of active historical alarms
This API call consumes the following media types via the Content-Type request header:

- application/json

Query parameters for both routes;

since - epoch time of the earliest alarm to return.
until - epoch time of the latest alarm to return.
limit - the maximum number of alarms to fetch.
for example, we might query:
CURL Example

```
curl  -X GET --header 'accept:application/json --header 'api_key:
igls-11llkok678afe9jo314931r1rdi40ib58r43bih3bsukmbuf0ugj'
'https://x.x.x.x/sera/v1/alarms/historical?limit=20' -k
```

GET /jobs

Get jobs (jobsGet)
Returns jobs from Superna Eyeglass.

Consumes

This API call consumes the following media types via the Content-Type request header:

- application/json

Query parameters

state (optional)
Query Parameter — filter running or complete jobs [all, running, finished]
success (optional)

Query Parameter – filter jobs by result success [true, false]

Return type

array[Job]

Produces

This API call produces the following media types according to the Accept request header; the media type will be conveyed by the Content-Type response header.

- application/json

Responses

200

jobs response

0

error payload

POST /jobs

Create a new job (jobsPost)

Launch a new job in Eyeglass.

Consumes

This API call consumes the following media types via the Content-Type request header:

- application/json

Query parameters

sourceid (required)

Query Parameter – ID of the source node for this job

targetid (required)

Query Parameter – ID of the target node for this job

failovertarget (required)

Query Parameter – ID of the access zone OR syncIQ policy to fail over.

Return type

inline_response_201

Produces

This API call produces the following media types according to the Accept request header; the media type will be conveyed by the Content-Type response header.

- application/json

Responses

201

create job response

0

error payload

GET /jobs/{id}

Retrieve a job by ID (jobsIdGet)
Retrieve a job by id

Path parameters

id (required)
Path Parameter — ID of the job to retrieve

Consumes

This API call consumes the following media types via the Content-Type request header:

- application/json

Return type

Job

Produces

This API call produces the following media types according to the Accept request header; the media type will be conveyed by the Content-Type response header.

- application/json

Responses

200

jobs response

0

error payload

DELETE /jobs/{id}

cancels a running job (jobsIdDelete)
Cancels a running job

Path parameters

id (required)
Path Parameter – ID of the running job to cancel

Consumes

This API call consumes the following media types via the Content-Type request header:

- application/json

Return type

Job

Produces

This API call produces the following media types according to the Accept request header; the media type will be conveyed by the Content-Type response header.

- application/json

Responses

200

jobs response

0

error payload

GET /nodes

Get all nodes (nodesGet)
Returns all Superna Eyeglass Managed Devices.

Consumes

This API call consumes the following media types via the Content-Type request header:

- application/json

Return type

array[Node]

Produces

This API call produces the following media types according to the Accept request header; the media type will be conveyed by the Content-Type response header.

- application/json

Responses

200

nodes response

0

error payload

GET /nodes/{id}

Find nodes by ID (nodesIdGet)
Returns the Superna Eyeglass Managed Devices based on ID

Path parameters

id (required)
Path Parameter – ID of the node to retrieve

Consumes

This API call consumes the following media types via the Content-Type request header:

- application/json

Return type

Node

Produces

This API call produces the following media types according to the Accept request header; the media type will be conveyed by the Content-Type response header.

- application/json

Responses

200

nodes response

0

error payload

How to retrieve health check data and all validations for Access Zones, pools, DFS policies and Synciq policies

Use the get pool, zone and policies node functions with the foreadiness flag set to true to return all validations provided by Eyeglass to detect failover readiness status.

See below API's and flag for readiness data

GET /nodes/{id}/policies

Find policies for a node (nodesIdPoliciesGet)
Returns the syncIQ policies for this node

forediness = true will return failover readiness validations to check the health of specific policy.

Parameters

Parameter	Value	Description	Parameter Type	Data Type
id		ID of the node to retrieve	path	string
foReadiness	true false (default)	**Retrieve also failover readiness status details**	query	boolean

Response Messages

HTTP Status Code	Reason	Response Model	Headers
default	error payload	Model Example Value `{ "code": 0, "message": "string" }`	

Path parameters

id (required)
Path Parameter — ID of the node to retrieve

Consumes

This API call consumes the following media types via the Content-Type request header:

- application/json

Return type

array[Policy]

Produces

This API call produces the following media types according to the Accept request header; the media type will be conveyed by the Content-Type response header.

- application/json

Responses

200

policies response

0

error payload

GET /nodes/{id}/pools

Find pool for a node (nodesIdpoolsGet)

Returns the pools for this node

foreadiness = true will return failover readiness validations to check the health of specific pool.

parameters

Parameter	Value	Description	Parameter Type	Data Type
id		ID of the node to retrieve	path	string
foReadiness	true false (default)	**Retrieve also failover readiness status details**	query	boolean

Response Messages

HTTP Status Code	Reason	Response Model	Headers
default	error payload	Model Example Value `{ "code": 0, "message": "string" }`	

GET /nodes/{id}/pools/{name}

Implementation Notes

Returns the pool by its name on a given node

GET /nodes/{id}/zones

Find zones for a node (nodesIdZonesGet)
Returns the access zones for this node
forediness = true will return failover readiness validations to check the health of specific zone.

Parameters

Parameter	Value	Description	Parameter Type	Data Type
id		ID of the node to retrieve	path	string
foReadiness	true false (default)	**Retrieve also failover readiness status details**	query	boolean

Response Messages

HTTP Status Code	Reason	Response Model	Headers
default	error payload	Model Example Value `{ "code": 0, "message": "string" }`	

Path parameters

id (required)
Path Parameter — ID of the node to retrieve

Consumes

This API call consumes the following media types via the Content-Type request header:

- application/json

Return type

array[Zone]

Produces

This API call produces the following media types according to the Accept request header; the media type will be conveyed by the Content-Type response header.

- application/json

Responses

200

zones response

0

error payload

ErrorModel

code
Integer
message
String

Job

failoverTarget
Job_failoverTarget
finished
Long The end time of the job
id
String
jobType
String The type of the job
Enum:
zone_failover
policy_failover
name
String Human readable name for this job
sourceNode
Node The source node for this job.
started
Long The start time of the job
success
Boolean True if there were no errors
targetNode

Node The target node for this job

Node

id
String
ip
String primary IP address of the node
name
String Unique name of the node

Policy

failoverReadiness
String Enumeration describing the failover readiness of the policy.
Enum:
ok
warning
error
id
String
name
String SyncIQ Policy Name
target
Node Node that is the target for this syncIQ policy
zone
Zone Access zone for this policy

Zone

failoverReadiness *String* Enumeration describing the failover readiness of the zone.
Enum:
ok
warning
error
id
String
name
String Access Zone Name

Inline_response

id
String

Job_failoverTarget

zone
Zone The access zone being failed over.
policies
array[Policy] The policies being failed over

Examples

Example 1: CMDB Integration API

CMDB Integration with ServiceNow or other CMDB systems that can use “http get” and “XML input files”.
CMDB integration requires key data about the cluster and configuration to be maintained in the CMDB to map services, service status or service resources. This is done for asset management, fault monitoring or common service availability with links between IT components that build a service. For example home directories requires PowerScale storage and Active directory.
The Eyeglass solution solves asset management and service status integration for shares and exports.
Supports CI objects (cluster, nodes, shares, exports, quotas)

1. CI attributes
2. cluster health, name, version, type, revision
3. node disk usage
4. share name, path, access zone, DR status (Active, DR, unprotected)
5. export path, access zone, DR status (Active, DR, unprotected)
6. quota path, type, usage, (hard, soft, advisory) limits

To access the file type https://x.x.x.x/servicenow/servicenow.xml to see the file.

This file is updated every 5 minutes with the latest information from all managed clusters. The XML file contains all of the information outlined above with any newly detected configuration added to the file.

SyncIQ protection of shares and exports in the XML output Description

The service component of shares and exports works as follows with data protection status indicated in the object

1. Assumption is anything with a share or export needs to be audited for SyncIQ data protection
2. Eyeglass discovers all shares and exports and sets attribute of the share in XML to a state.
    1. **unprotected** - no SyncIQ policy
    2. **Active** (means writable copy of the data as of now)
    3. **backup** (means DR copy not writable)

Quota output in the XML file Description

Quotas are also output to the XML file with type, path and usage in bytes. The usage is updated every 5 minutes to the file to keep the CMDB updated with quota usage information.

Asset Management output in the XML file Description

The cluster node count serial number output can be used for asset management updates. This information is dynamically updated, if cluster nodes are added the file will include the new nodes. If the software is updated on the cluster, the new version of the software will be included in this file.
This information can be used to maintain CMDB asset records for PowerScale clusters that is current and up to date.

How to Integrate with a CMDB

The file path is tested with ServiceNow that has the http get/XML file feature. This is available in most CMDB products and integration and transformation maps are outside the scope of support of Eyeglass. Professional Services can be purchased for assistance to integrate

with a CMDB. Integration typically requires getting the input file on a schedule and mapping the XML fields to a CI object.
This screenshot shows how integration can be done with ServiceNow.

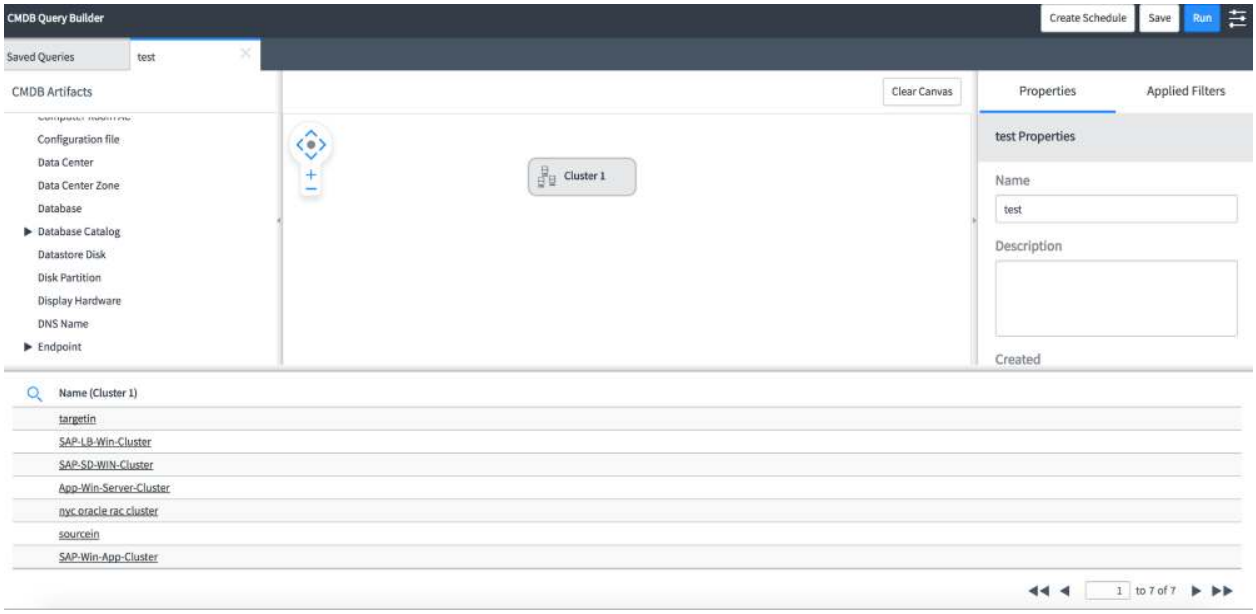

Sample CMDB XML file for 4 clusters


```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<elements>
   <element>
      <cluster>
         <name>Cluster-1-7201</name>
         <guid>005056ba34580f410c55fd077989478a3821</guid>
         <local_serial>V5052427352</local_serial>
         <onefs_version>
            <build>B_7_2_1_014(RELEASE)</build>
            <release>v7.2.1.0</release>
            <type>PowerScale OneFS</type>
            <version>PowerScale OneFS v7.2.1.0
B_7_2_1_014(RELEASE): 0x70201500000000E:Mon Jun 22
20:05:42 GMT 2015    root@sea-build7-
03:/b/mnt/obj/b/mnt/src/sys/IQ.amd64.release   clang version 3.3
(tags/RELEASE_33/final)</version>
            <revision>5049675513279283 34</revision>
         </onefs_version>
      </cluster>
      <nodes>
```

```
    <node>
      <health>OK</health>
      <serial_number>V5052427352</serial_number>
      <HDD_size>35GB</HDD_size>
      <HDD_used>13GB</HDD_used>
      <HDD_avail>23GB</HDD_avail>
      <VHS>n/a</VHS>
    </node>
  </nodes>
  <shares>
    <share>
      <name>SMB2</name>
      <path>/ifs/data/policy1</path>
      <zone>System</zone>
      <status>ACTIVE</status>
    </share>
    <share>
      <name>spacetest</name>
      <path>/ifs/data/policy1/space</path>
      <zone>System</zone>
      <status>ACTIVE</status>
    </share>
    <share>
      <name>igls-dfs-dfs1</name>
      <path>/ifs/data/userdata/dfs1</path>
      <zone>data</zone>
      <status>BACKUP</status>
    </share>
    <share>
      <name>igls-dfs-roam</name>
      <path>/ifs/data/userdata/dfs1/roam</path>
      <zone>data</zone>
      <status>BACKUP</status>
    </share>
    <share>
      <name>share1</name>
      <path>/ifs/data/userdata/share1</path>
```

```
        <zone>data</zone>
        <status>ACTIVE</status>
      </share>
    </shares>
    <exports>
      <export>
        <paths>/ifs</paths>
        <zone>System</zone>
        <status>UNPROTECTED</status>
      </export>
      <export>
        <paths>/ifs/data/robot</paths>
        <zone>EyeglassRunbookRobot-AccessZone</zone>
        <status>ACTIVE</status>
      </export>
      <export>
        <paths>/ifs/data/userdata/export1/andrew</paths>
        <zone>data</zone>
        <status>ACTIVE</status>
      </export>
      <export>
        <paths>/ifs/data/userdata/export1</paths>
        <zone>data</zone>
        <status>ACTIVE</status>
      </export>
    </exports>
    <aliases/>
    <quotas>
      <quota>
        <path>/ifs/data/userdata</path>
        <type>directory</type>
        <used>5299221577</used>
        <hardLimit>17008070491</hardLimit>
      </quota>
      <quota>
        <path>/ifs/data/userdata/export1/andrew</path>
        <type>directory</type>
```

```
    <used>12</used>
  </quota>
  <quota>
    <path>/ifs/data/policy1</path>
    <type>directory</type>
    <used>1972155571</used>
    <hardLimit>21474836480</hardLimit>
  </quota>
  <quota>
    <path>/ifs/data/policy1/space</path>
    <type>directory</type>
    <used>851444622</used>
  </quota>
  <quota>
    <path>/ifs/data/userdata/export1</path>
    <type>directory</type>
    <used>124216376</used>
    <advisoryLimit>1048576000</advisoryLimit>
  </quota>
  <quota>
    <path>/ifs/data/userdata/share1</path>
    <type>directory</type>
    <used>218</used>
  </quota>
  <quota>
    <path>/ifs/data/policy1/departmentB</path>
    <type>directory</type>
    <used>12</used>
    <hardLimit>26843545600</hardLimit>
  </quota>
  <quota>
    <path>/ifs/data/robot</path>
    <type>directory</type>
    <used>24833</used>
  </quota>
  <quota>
    <path>/ifs/data/policy1/departmentA</path>
```

```
            <type>directory</type>
            <used>12</used>
            <hardLimit>26843545600</hardLimit>
          </quota>
        </quotas>
      </element>
      <element>
        <cluster>
          <name>Cluster2-7201</name>
          <guid>005056ba72edf6450c552312a728d3a22a23</guid>
          <local_serial>V5052443373</local_serial>
          <onefs_version>
            <build>B_7_2_1_014(RELEASE)</build>
            <release>v7.2.1.0</release>
            <type>PowerScale OneFS</type>
            <version>PowerScale OneFS v7.2.1.0
B_7_2_1_014(RELEASE): 0x70201500000000E:Mon Jun 22
20:05:42 GMT 2015    root@sea-build7-
03:/b/mnt/obj/b/mnt/src/sys/IQ.amd64.release   clang version 3.3
(tags/RELEASE_33/final)</version>
            <revision>504967551327928334</revision>
          </onefs_version>
        </cluster>
        <nodes>
          <node>
            <health>-A--</health>
            <serial_number>V5052443373</serial_number>
            <HDD_size>35GB</HDD_size>
            <HDD_used>11GB</HDD_used>
            <HDD_avail>25GB</HDD_avail>
            <VHS>n/a</VHS>
          </node>
        </nodes>
        <shares>
          <share>
            <name>igls-dfs-SMB2</name>
            <path>/ifs/data/policy1</path>
```

```
    <zone>System</zone>
    <status>BACKUP</status>
  </share>
  <share>
    <name>igls-dfs-spacetest</name>
    <path>/ifs/data/policy1/space</path>
    <zone>System</zone>
    <status>BACKUP</status>
  </share>
  <share>
    <name>migrate1</name>
    <path>/ifs/data/migrate1</path>
    <zone>System</zone>
    <status>UNPROTECTED</status>
  </share>
  <share>
    <name>dfs1</name>
    <path>/ifs/data/userdata/dfs1</path>
    <zone>data</zone>
    <status>ACTIVE</status>
  </share>
  <share>
    <name>roam</name>
    <path>/ifs/data/userdata/dfs1/roam</path>
    <zone>data</zone>
    <status>ACTIVE</status>
  </share>
  <share>
    <name>share1</name>
    <path>/ifs/data/userdata/share1</path>
    <zone>data</zone>
    <status>BACKUP</status>
  </share>
  <share>
    <name>share2</name>
    <path>/ifs/data/userdata/share2</path>
    <zone>data</zone>
```

```
        <status>ACTIVE</status>
    </share>
    <share>
        <name>dfs1</name>
        <path>/ifs/data/dr-testing/dfs1</path>
        <zone>DR-Testing-Zone</zone>
        <status>BACKUP</status>
    </share>
    <share>
        <name>share1</name>
        <path>/ifs/data/dr-testing/share1</path>
        <zone>DR-Testing-Zone</zone>
        <status>BACKUP</status>
    </share>
    <share>
        <name>share2</name>
        <path>/ifs/data/dr-testing/share2</path>
        <zone>DR-Testing-Zone</zone>
        <status>BACKUP</status>
    </share>
</shares>
<exports>
    <export>
        <paths>/ifs</paths>
        <zone>System</zone>
        <status>UNPROTECTED</status>
    </export>
    <export>
        <paths>/ifs/data/migrate1</paths>
        <zone>System</zone>
        <status>UNPROTECTED</status>
    </export>
    <export>
        <paths>/ifs/data/robot</paths>
        <zone>EyeglassRunbookRobot-AccessZone</zone>
        <status>BACKUP</status>
    </export>
```

```
    <export>
        <paths>/ifs/data/userdata/export1/andrew</paths>
        <zone>data</zone>
        <status>BACKUP</status>
    </export>
    <export>
        <paths>/ifs/data/userdata/export1</paths>
        <zone>data</zone>
        <status>BACKUP</status>
    </export>
    <export>
        <paths>/ifs/data/dr-testing/export1/andrew</paths>
        <zone>DR-Testing-Zone</zone>
        <status>BACKUP</status>
    </export>
    <export>
        <paths>/ifs/data/dr-testing/export1</paths>
        <zone>DR-Testing-Zone</zone>
        <status>BACKUP</status>
    </export>
</exports>
<aliases/>
<quotas>
    <quota>
        <path>/ifs/data/userdata/dfs1</path>
        <type>directory</type>
        <used>5175108664</used>
    </quota>
    <quota>
        <path>/ifs/data/userdata/dfs1/roam</path>
        <type>directory</type>
        <used>2251508</used>
    </quota>
    <quota>
        <path>/ifs/data/migrate1</path>
        <type>directory</type>
        <used>769442173</used>
```

```
        </quota>
        <quota>
            <path>/ifs/data/userdata/share2</path>
            <type>directory</type>
            <used>12</used>
        </quota>
    </quotas>
</element>
<element>
    <cluster>
        <name>prod-8</name>
        <guid>005056ba67371492dd56f106ca5e3ff16028</guid>
        <local_serial>SV200-004EIJ-B96U</local_serial>
        <onefs_version>
            <build>B_8_0_1_007(RELEASE)</build>
            <release>v8.0.1.0</release>
            <type>PowerScale OneFS</type>
            <version>PowerScale OneFS v8.0.1.0
B_8_0_1_007(RELEASE): 0x800015000000007:Thu Sep  8 06:34:05
PDT 2016    root@sea-build10-
02:/b/mnt/obj/b/mnt/src/sys/IQ.amd64.release   FreeBSD clang
version 3.3 (tags/RELEASE_33/final 183502) 20130610</version>
            <revision>576462195412434951</revision>
        </onefs_version>
    </cluster>
    <nodes>
        <node>
            <health>OK</health>
            <serial_number>SV200-004EIJ-B96U</serial_number>
            <HDD_size>18.1GB</HDD_size>
            <HDD_used>2.6GB</HDD_used>
            <HDD_avail>15.5GB</HDD_avail>
            <VHS>n/a</VHS>
        </node>
    </nodes>
    <shares>
        <share>
```

```
        <name>migrate2</name>
        <path>/ifs/data/migrate2</path>
        <zone>System</zone>
        <status>UNPROTECTED</status>
    </share>
    <share>
        <name>dfs1</name>
        <path>/ifs/data/marketing/dfs1</path>
        <zone>marketing</zone>
        <status>ACTIVE</status>
    </share>
    <share>
        <name>share1</name>
        <path>/ifs/data/marketing/shares</path>
        <zone>marketing</zone>
        <status>ACTIVE</status>
    </share>
</shares>
<exports>
    <export>
        <paths>/ifs</paths>
        <zone>System</zone>
        <status>UNPROTECTED</status>
    </export>
    <export>
        <paths>/ifs/data/migrate2</paths>
        <zone>System</zone>
        <status>UNPROTECTED</status>
    </export>
    <export>
        <paths>/ifs/data/marketing/nfs/export1</paths>
        <zone>marketing</zone>
        <status>ACTIVE</status>
    </export>
</exports>
<aliases/>
<quotas>
```

```
        <quota>
          <path>/ifs/data/marketing/nfs/export1</path>
          <type>directory</type>
          <used>0</used>
        </quota>
        <quota>
          <path>/ifs/data/marketing/shares</path>
          <type>directory</type>
          <used>12</used>
        </quota>
        <quota>
          <path>/ifs/data/marketing/dfs1</path>
          <type>directory</type>
          <used>58599</used>
        </quota>
        <quota>
          <path>/ifs/data/migrate2</path>
          <type>directory</type>
          <used>769417411</used>
        </quota>
      </quotas>
    </element>
    <element>
      <cluster>
        <name>dr-8</name>
        <guid>005056ba657091badd564b1487f19066d641</guid>
        <local_serial>SV200-004EIJ-AQR3</local_serial>
        <onefs_version>
          <build>B_8_0_1_007(RELEASE)</build>
          <release>v8.0.1.0</release>
          <type>PowerScale OneFS</type>
          <version>PowerScale OneFS v8.0.1.0
B_8_0_1_007(RELEASE): 0x800015000000007:Thu Sep  8 06:34:05
PDT 2016    root@sea-build10-
02:/b/mnt/obj/b/mnt/src/sys/IQ.amd64.release   FreeBSD clang
version 3.3 (tags/RELEASE_33/final 183502) 20130610</version>
          <revision>576462195412434951</revision>
```

```
    </onefs_version>
  </cluster>
  <nodes>
    <node>
      <health>OK</health>
      <serial_number>SV200-004EIJ-AQR3</serial_number>
      <HDD_size>18.1GB</HDD_size>
      <HDD_used>2.9GB</HDD_used>
      <HDD_avail>15.2GB</HDD_avail>
      <VHS>n/a</VHS>
    </node>
  </nodes>
  <shares>
    <share>
      <name>migrate1</name>
      <path>/ifs/data/dr/migrate</path>
      <zone>System</zone>
      <status>UNPROTECTED</status>
    </share>
    <share>
      <name>igls-dfs-dfs1</name>
      <path>/ifs/data/marketing/dfs1</path>
      <zone>marketing</zone>
      <status>BACKUP</status>
    </share>
    <share>
      <name>share1</name>
      <path>/ifs/data/marketing/shares</path>
      <zone>marketing</zone>
      <status>BACKUP</status>
    </share>
  </shares>
  <exports>
    <export>
      <paths>/ifs</paths>
      <zone>System</zone>
      <status>UNPROTECTED</status>
```

```
        </export>
        <export>
            <paths>/ifs/data/dr/migrate</paths>
            <zone>System</zone>
            <status>UNPROTECTED</status>
        </export>
        <export>
            <paths>/ifs/data/marketing/nfs/export1</paths>
            <zone>marketing</zone>
            <status>BACKUP</status>
        </export>
    </exports>
    <aliases/>
    <quotas>
        <quota>
            <path>/ifs/data/dr/migrate</path>
            <type>directory</type>
            <used>167747631</used>
        </quota>
    </quotas>
  </element>
</elements>
```


Example 2: VMWare SRM Integrated Failover Example

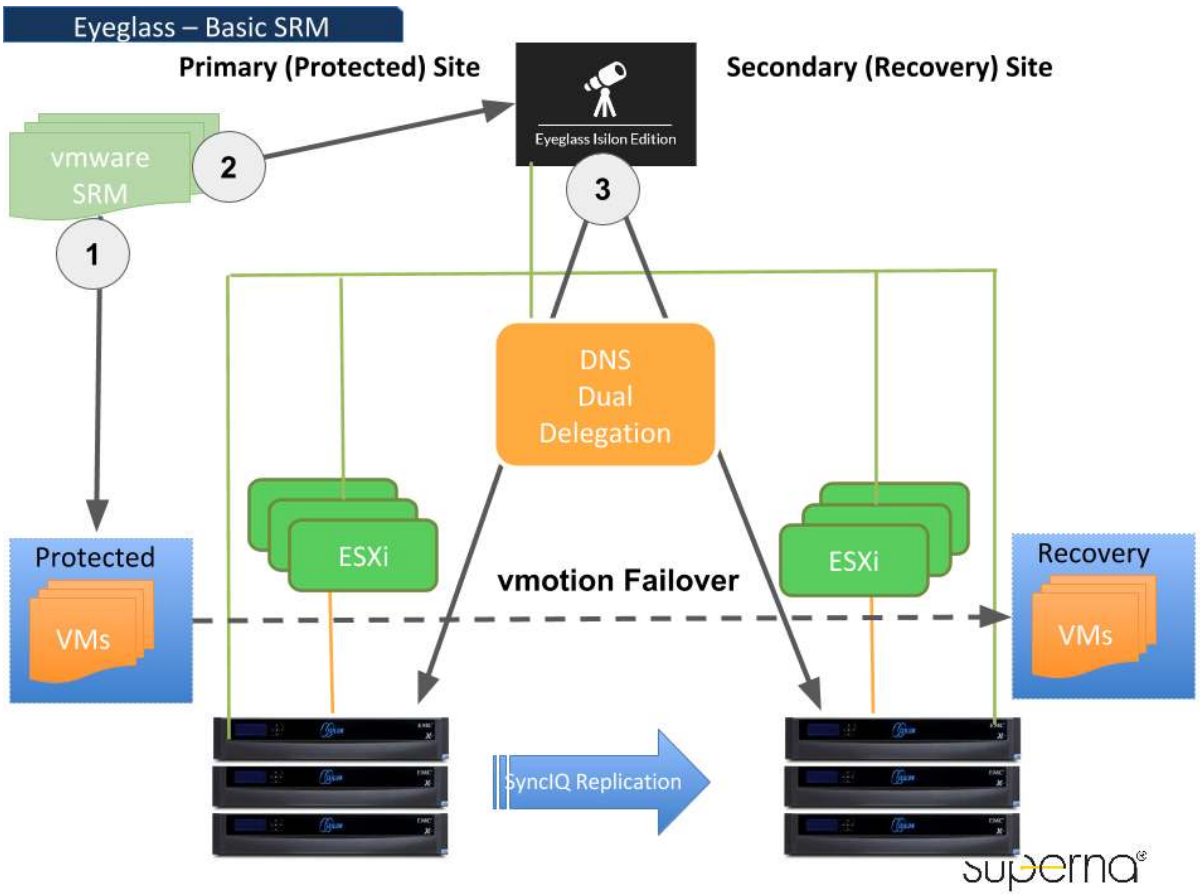

1. VMWare SRM starts site failover of VM's
2. External API calls Eyeglass REST API using curl command (created with CURL builder Eyeglass API explorer interface).
   VMWare SRM Creating Custom Recovery Steps
   1. Add command to the recovery plan: http://pubs.vmware.com/srm-61/index.jsp#com.vmware.srm.admin.doc/GUID-BABE0457-EB6F-4650-BB8B-01300ACAFF2F.html
   2. Create Message Prompts or Command Steps for Individual Virtual Machines: (http://pubs.vmware.com/srm-61/index.jsp#com.vmware.srm.admin.doc/GUID-45EE6522-3659-437F-B5AF-E9510AAA2CC8.html)
   3. On the Recovery Properties tab, click Pre-Power On Steps.
      1. Command on SRM Server Runs a command on Site Recovery Manager Server. This option is available for both pre-power on steps and post-power on steps

2. Complete rest of Step configuration

4. You must start the Windows command shell using its full path on the local host. **For example**, to run a script located in c:\alarmscript.bat, use the following command line:

   c:\windows\system32\cmd.exe /c c:\alarmscript.bat

5. You must install batch files and commands on the Site Recovery Manager Server at the recovery site.

6. Batch files or commands producing output that contains characters with ASCII values greater than 127 must use UTF-8 encoding. Site Recovery Manager records only the final 4KB of script output in log files and in the recovery history. Scripts that produce more output should redirect the output to a file rather than sending it to the standard output to be logged.

   Note: Batch files and commands must finish within 600 seconds, otherwise, the recovery plan terminates with an error. (See the following VMWare reference Change Recovery Settings). Therefore, during the configuration of SRM you may receive a timeout after 600 seconds . You can increase or decrease this value by editing the SRM configuration file (vmware-dr.xml). Look for the following section:

   <calloutCommandLineTimeout>600</calloutCommandLineTimeout>

7. Change value to the appropriate value.

1. Eyeglass API processes request from CURL command with timeout set and starts any failover mode supported by Eyeglass and requested by the CURL builder setup for the failover.

Example 3: Initiate an Access Zone failover

Failover Setting Parameters - Eyeglass REST API

The following failover setting parameters are available through the Eyeglass REST API curl command:

Parameter	Description
controlled	Execute a controlled failover by running operations against the source cluster as well as the target
datasync	Run the final incremental data sync before failover
configsync	Run a configuration sync before failover Should bd false for Release 2.5.6 and above.
resyncprep	Run resync prep on the source cluster to create the mirror policies
disablemirror	Disable mirror policies created on the failover target
quotasync	Run quota jobs to failover quotas to target
blockonwarnings	Blocks failover from starting if DR failover status is Warning when true. All Warnings in DR Readiness will block a failover and must be reviewed before disabling this option. Best Practice: Verify with support.superna.net warning validation prior to setting to false

rollbackrenameshares	This applies only to DFS mode failover and should be left enabled to automatically rollback SMB share rename step if required due to error in failover step.
smbdataintegrity	This mode disconnects any active SMB sessions prior to failover and ensures that no new sessions can be established on the failover source. It applies a deny read permission to the Everyone user to each share. Note any share with run as root bypasses all security and cannot be locked out from from a share.

Eyeglass REST API - Failover Parameters Rules

1. Failover Job is only allowed through this Eyeglass REST API when the Readiness shows OK, INFO or WARNING. Status. When the Readiness shows ERROR the failover job will be blocked and curl command will return this code:

```
{
"code": 500,
"message": "{\"code\":500,\"message\":\"Zone zone01 is not eligible for
failover\"}"
}
```

   We need to rectify the highlighted issues and then re-verify the Readiness does not show ERROR status before we can execute the Failover job

3. For uncontrolled Failover (controlled=false), the following parameters also need to be set as false: datasync, configsync, resyncprep, and disablemirror. Otherwise the curl command will return this output to highlight that requirement

```
{
"code": 500,
"message": "{\"code\":500,\"message\":\"For an uncontrolled failover,
datasync, configsync, resyncprep and disablemirror must all be false\"}"
```

```
}
```

4. For failover with parameter **resyncprep=false**, also need to leave the disablemirror parameter as false. If disablemirror is set as true while trying to run the failover job with resyncprep=false, the curl return this output.

```
{
"code": 500,
"message": "{\"code\":500,\"message\":\"Cannot include the disablemirror flag when reyncprep is false.\"}"
}
```

Step 1: Get all of the PowerScale Clusters provisioned in superna eyeglass:

Request:

curl -k -H "api_key: igls-4e81gu94mc3opgf7uuhdtrf6oo0a0arfcuajra65l834l8p53j" https://192.168.10.10/sera/v1/nodes

Response:

```
[
 {
   "id": "Kyle-8-A_00505698937a1b73bb5698242b10b5fe9a97",
   "ip": "172.16.86.238",
   "name": "Kyle-8-A"
 },
 {
   "id": "Kyle-8-B_00505698f0793f8bbb56fc176e2f7b6e204c",
   "ip": "172.16.86.248",
   "name": "Kyle-8-B"
 }
]
```

Step 2: Using the ID from step 1, get all of the access zones on the cluster:

Request:

```
curl -k -H "api_key: igls-
p9cbjc3bjkgbo3ceph9t29jholrtvuc08p17ri7na73eal5g9nv"
http://localhost:8089/sera/v1/nodes/Kyle-8-
A_00505698937a1b73bb5698242b10b5fe9a97/zones
```

Response:


```
[
 {
   "failoverReadiness": "warning",
   "id": "zone-Kyle-8-A_kylezone",
   "name": "kylezone"
 }
]
```


Step 3: Initiate a failover

Using the id of the zone and the IDs of the source and target clusters, post a new job to the jobs route to initiate a failover:

Request

```
curl -X POST -k -H "api_key: igls-
p9cbjc3bjkgbo3ceph9t29jholrtvuc08p17ri7na73eal5g9nv" -H
"Content-type: application/json"
"http://localhost:8089/sera/v1/jobs?sourceid=Kyle-8-
A_00505698937a1b73bb5698242b10b5fe9a97&targetid=Kyle-8-
B_00505698f0793f8bbb56fc176e2f7b6e204c&failovertarget=zone-
Kyle-8-A_kylezone"
```

Response


```
{
 "id": "job-1457385733807-630400755"
}
```


Step 4: Monitor the failover

Using the ID of the job, monitor the status of the failover. When the job is complete, there will be a “finished” property in the output and a “success” property indicating whether the job was a success or a failover.

Request

```
curl -k -H "api_key: igls-
p9cbjc3bjkgbo3ceph9t29jholrtvuc08p17ri7na73eal5g9nv"
http://localhost:8089/sera/v1/jobs/job-1457385733807-630400755
```

Response


```
{
 "failoverTarget": {
  "zone": {
   "failoverReadiness": "error",
   "id": "zone-Kyle-8-A_kylezone",
   "name": "kylezone"
  }
 },
 "finished": 1457385819964,
 "id": "job-1457385733807-630400755",
 "jobType": "zone_failover",
 "name": "Access_Zone_Failover__Kyle-8-A__2016-03-07_16-22-
13",
 "sourceNode": {
  "id": "Kyle-8-A_00505698937a1b73bb5698242b10b5fe9a97",
  "ip": "172.16.86.238",
  "name": "Kyle-8-A"
 },
 "started": 1457385733813,
 "success": true,
 "targetNode": {
  "id": "Kyle-8-B_00505698f0793f8bbb56fc176e2f7b6e204c",
```

"ip": "172.16.86.248",
"name": "Kyle-8-B"
}
}

Example 4: Get all currently running failover jobs

Step 1: use the ?state=running query to filter by running jobs:

Output from this route will not have “success” or “finished” parameters in the payload.

Request:

```
curl -k -H "api_key: igls-
p9cbjc3bjkgbo3ceph9t29jholrtvuc08p17ri7na73eal5g9nv"
http://localhost:8089/sera/v1/jobs\?state\=running
```

Response:

{
"failoverTarget": {
"zone": {
"failoverReadiness": "warning",
"id": "zone-Kyle-8-A_kylezone",
"name": "kylezone"
}
},
"id": "job-1457385733807-630400755",
"jobType": "zone_failover",
"name": "Access_Zone_Failover__Kyle-8-A__2016-03-07_16-22-
13",
"sourceNode": {
"id": "Kyle-8-A_00505698937a1b73bb5698242b10b5fe9a97",
"ip": "172.16.86.238",
"name": "Kyle-8-A"
},

"started": 1457385733813,
"targetNode": {
"id": "Kyle-8-B_00505698f0793f8bbb56fc176e2f7b6e204c",
"ip": "172.16.86.248",
"name": "Kyle-8-B"
}
}

Example 5: Get a historical record of all previous failover jobs

Step 1: use the ?state=finished query to filter by completed jobs:

Output from this route will not have “success” or “finished” parameters in the payload.

Request:

```
curl -k -H "api_key: igls-
p9cbjc3bjkgbo3ceph9t29jholrtvuc08p17ri7na73eal5g9nv"
http://localhost:8089/sera/v1/jobs?state=finished
```

Response:

{
"failoverTarget": {
"zone": {
"failoverReadiness": "warning",
"id": "zone-Kyle-8-A_kylezone",
"name": "kylezone"
}
},
"id": "job-1457385733807-630400755",
"finished": 1457385819964,
"jobType": "zone_failover",

"name": "Access_Zone_Failover__Kyle-8-A__2016-03-07_16-22-13",
"sourceNode": {
"id": "Kyle-8-A_00505698937a1b73bb5698242b10b5fe9a97",
"ip": "172.16.86.238",
"name": "Kyle-8-A"
},
"started": 1457385733813,
"success": true,
"targetNode": {
"id": "Kyle-8-B_00505698f0793f8bbb56fc176e2f7b6e204c",
"ip": "172.16.86.248",
"name": "Kyle-8-B"
}
}

Ransomware Defender Airgap API's

These API's require 2.5.8 release. They allow Airgap basic job automation to start an Airgap job, monitor the job status and list Airgap jobs that are configured. Requires Release 2.5.8 or later

These API's can be used to list airgap policy jobs, run and airgap job on demand and monitor it's status. To get the details of the usage use the API explorer GUI on usage and parameters needed for each API.

-
  - GET /v2/jobs/airgap
    - Get all airgap jobs
-
  - POST /v2/jobs/airgap

- Start an airgap job

-

- POST /v2/jobs/airgap/status
  - post airgap job status

Ransomware Defender Smart Airgap API

This api allows external applications to get data security or threat level of production file or object data. External applications can use this api to get status for decision making or actions based on an active threat. 2 API's exist with one focusing on Ransomware Defender only and the other exposing the smart Airgap API for external applications to use.

Requirements:

1. Release 2.5.8 or later

GET /v1/ransomware/rswevents (this is for Ransomware Defender only)

GET /v1/securityevents (this is the smart airgap api that allows a filter of all or Ransomware Dender or Easy Auditor active alerts only events)

Known Issues

T783 Rest API delete job error

Description: Deleting a job from the API Explorer requires an empty JSON object to be specified in the body.
Workaround: Enter ‘{}’ (without quotes) into the body field of the DELETE /jobs section of the API explorer to cancel a job.

T15042 Rest API policy readiness is missing output for Target Reachability check

Description: The SyncIQ policy readiness retrieved using REST API is missing the output for the Target Reachability check. If the Target Reachability validation fails the overall failover status is correctly shown as ERROR and failover cannot be initiated.
Workaround: To assess target reachability:

- Target reachability alarms related to Inventory or Configuration Replication would have been sent.
- From the Eyeglass web interface, Eyeglass/PowerScale reachability can be viewed from the Continuous Operations dashboard.
- All failover readiness criteria can be viewed from the Eyeglass web interface DR Dashboard.

T10935 Pool failover "failovertarget" field must be "zone id"

Description: The "failovertarget" field must be "zone id" even though description indicates "ID of the access zone OR syncIQ policy to failover" .
Workaround: Enter "zone id" for "failovertarget" when initiating pool failover.

T15623 Pool failover API does not support multiple pool selection

Description: From Eyeglass DR Assistant a Pool failover can be initiated for multiple pools but this is not supported from the API.
Workaround: Run concurrent failover for multiple pools.

T15624 Failover API does not block controlled failover when source cluster unreachable

Failover API does not validate source cluster reachability and will allow a controlled failover to start even if source cluster unreachable. Controlled failover in this case is expected to fail as it will attempt steps against the source cluster. When source cluster is not reachable uncontrolled failover should be used.
Workaround: Use manual process to verify source cluter reachability and initiate the appropriate controlled or uncontrolled failover.

T17428 API Policy Readiness returns incorrect Access Zone

Failover API to retrieve Policy Readiness information returns the incorrect Access Zone for environments with multiple Access Zones.
Workaround: None required. Access Zone does not affect Policy Failover and Access Zone Readiness and Failover correctly assign policy to correct Access Zone.

Known Limitations

T18079 Change Eyeglass Configuration Replication Job Disable/Enable API must be done at same time as Job Type Change

The API to change Eyeglass Configuration Replication Job state for Disable/Enable must be done together with the job state change in order to succeed.

© Superna LLC

1.29.1. How to build a Web server solution with the Eyeglass API

Home Top

- How to use this Guide
- Overview
- Common use cases:
- Requirements
- Components
- Diagram
- Failover Gateway
- Quickstart
- Configuration
- Example
- Nginx example config
- Usage

How to use this Guide

This example walks through how to build a secure solution using the Eyeglass api. The API is design to allow customers to build business logic to automate failover or present data to business units or end customers in a managed service.

NOTE: The solution documented here is an example only as a reference to get started on building solutions with the API. The material

on this page is "as is", and not covered under the support contract. It has no warranty updates and is free to use , change modify under the MIT license.

Overview

The objective of this solution is to provide an interface to securely initiate failover or failback and provide status of failover remotely without needing the Eyeglass GUI. This example will show how to build a proxy solution and will show how curl can be used but a well design solution would embed the api in a web page to simplify the user experience. Developing a webpage is out side the scope of this example and for simplicity curl will be used as an example to access the API proxy.

Common use cases:

1. multi tenant within a single company with business units or development teams that needed to test failover and failback unassisted by Storage admin
2. Hosted PowerScale and Eyeglass managed service solution to multiple end customers that are not the same company

Requirements

1. Security is a key requirement in all uses cases and should be factored into a well architected design. This guide will provide high level architecture on deploying this solution.
2. Separation of customers logically and physically (out of scope for this guide)
3. Separation of management infrastructure from end users or companies. This is out of scope for this guide and should be designed with networking considerations along with authentication to any service that enables failover of data. (out of scope for this guide)

Components

1. Eyeglass
2. PowerScale pair with synciq
3. Webserver - Nginx used in this example
4. Nodejs installed with webserver

Diagram

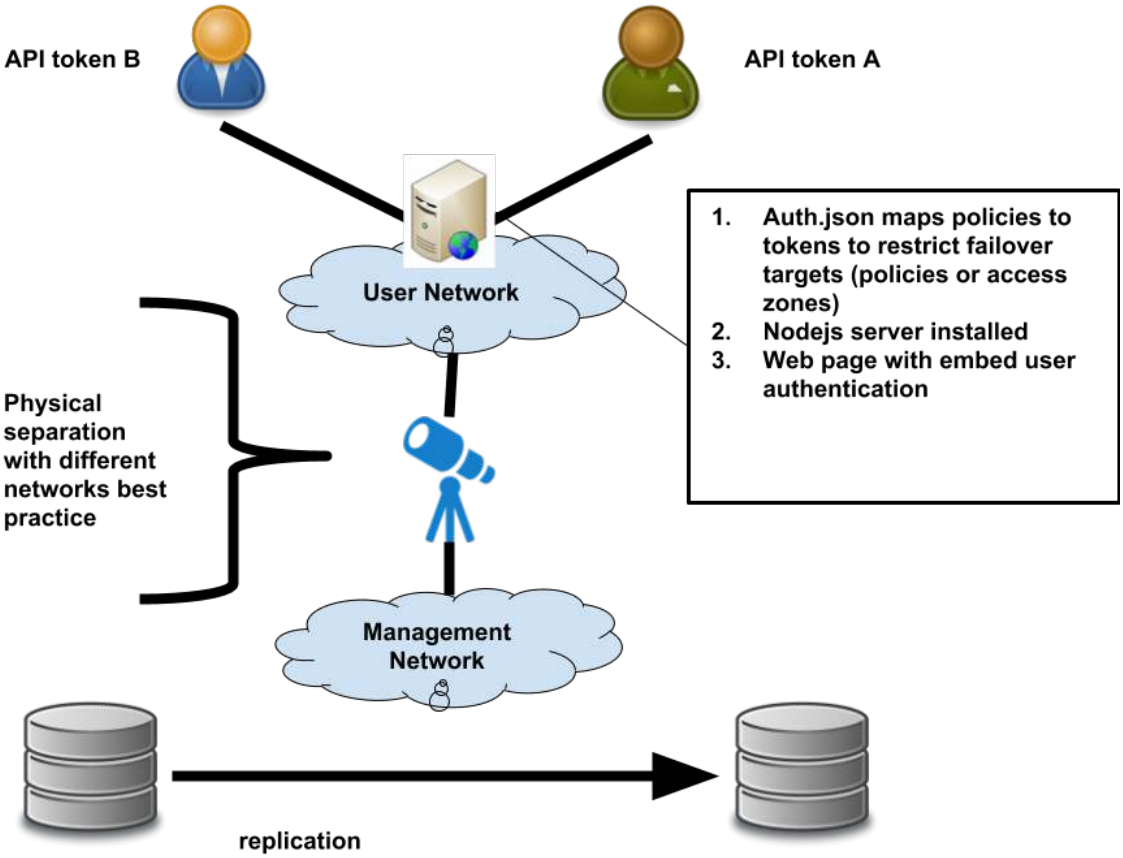

Failover Gateway

The failover gateway is a small nodejs application that can sit in front of the Superna Eyeglass REST API (SERA) and provide token authentication to failover a subset of SyncIQ policies and/or Access Zones.

Quickstart

- Install nodejs 8 or higher.
- Clone this repository to a directory on the web server (failover gateway example): `git clone https://bitbucket.org:superna/failover_example`
- Install dependencies: `npm install`
- Configure the `auth.json` file (see below)

- start the server: `npm start`
- NPM listens on port 8080

Configuration

All configuration is done through the `auth.json` file in the root of the project. The values to be filled in can be determined by using the Superna Eyeglass REST API Explorer:

1. Login to your eyeglass appliance.
2. From the main menu select "Eyeglass REST API"
3. Generate a REST API token, and copy to the clipboard.
4. Select the API Explorer tab, and click the Launch API Explorer button
5. In the api_key box at the top right, paste your copied API key and click Explore
6. Use the `/nodes` route to get the IDs of the PowerScale Clusters provisioned in eyeglass.
7. use the `/nodes/<id>/policies` and `/nodes/<id>/zones` routes to get the policies and access zones for a given cluster.
8. use the information to build the `auth.json` file in the root of the failover gateway project downloaded above, a sample file is included in the download.
9. Deploy a web server and install a node server, run the node example code from above. This is example is failover gateway example
10. Configure the web server for https following web server documentation

11. sample nginx configuration is part of the git download to redirect http to the node js server code listening on port 8080
12. License file is included in the git download
13. Review diagram for secure deployment example
14. Follow example below to create the mapping from a token created in eyeglass to a policy or access zone using the api browser tool in Eyeglass. See the Eyeglass API guide


```
{
"eyeglassIp": "192.168.1.140",
"apiKeys": {
"key-1": [
{
"targetId": "target-PowerScale-id",
"sourceId": "source-PowerScale-id",
"failoverTarget": "policy-name"
},{
"targetId": "target-PowerScale-id",
"sourceId": "source-PowerScale-id",
"failoverTarget": "policy-name_mirror"
},{
"targetId": "target-PowerScale-id",
"sourceId": "source-PowerScale-id",
"failoverTarget": "zone-name-1"
}
],
"key-2": [
```

{
"targetId": "target-PowerScale-id",
"sourceId": "source-PowerScale-id",
"failoverTarget": "policy-name-2"
}
]
}
}

Where in the above, the user that is provided with the `key-1` api key has authorization to failover only the policies `policy_name` and `policy_name_mirror`, as well as the access zone `zone-name-1`. The user provided with `key-2` can only failover `policy-name-2` and nothing else.

Example

Getting all PowerScales:

GET /nodes

[
{
"id": "ISL-TEST-8-1-0-0-144_0050569ff3955440815c200978332081a600",
"ip": "192.168.1.244",
"name": "ISL-TEST-8-1-0-0-144"
},
{

"id": "ISL-TEST-8-1-0-0-
146_0050569f65335842815c970993ede1cfa604",
"ip": "192.168.1.246",
"name": "ISL-TEST-8-1-0-0-146"
}
]

Getting all policies for each of the two PowerScales returned by the previous call:

GET /nodes/ISL-TEST-8-1-0-0-
144_0050569ff3955440815c200978332081a600/policies

[
{
"failoverReadiness": "ok",
"id": "policy-ISL-TEST-8-1-0-0-144_dfs",
"name": "dfs",
"target": {
"id": "ISL-TEST-8-1-0-0-
146_0050569f65335842815c970993ede1cfa604",
"ip": "192.168.1.246",
"name": "ISL-TEST-8-1-0-0-146"
},
"zone": {
"failoverReadiness": "error",
"id": "zone-ISL-TEST-8-1-0-0-144_System",
"name": "System"

}
}
]

GET /nodes/ISL-TEST-8-1-0-0-146_0050569f65335842815c970993ede1cfa604/policies

[
{
"failoverReadiness": "warning",
"id": "policy-ISL-TEST-8-1-0-0-146_dfs_mirror",
"name": "dfs_mirror",
"target": {
"id": "ISL-TEST-8-1-0-0-144_0050569ff3955440815c200978332081a600",
"ip": "192.168.1.244",
"name": "ISL-TEST-8-1-0-0-144"
},
"zone": {
"failoverReadiness": "error",
"id": "zone-ISL-TEST-8-1-0-0-146_System",
"name": "System"
}
}
]

Constructing an auth.json file to allow failover and failback of only these policies:

{
"eyeglassIp": "192.168.1.140",
"apiKeys": {
"igls-1uhrns90b9v4s9cggtt6vs5dlrjj38hkeapdc4ro1sdjnu0bcqtm": [
{
"targetId": "ISL-TEST-8-1-0-0-
146_0050569f65335842815c970993ede1cfa604",
"sourceId": "ISL-TEST-8-1-0-0-
144_0050569ff3955440815c200978332081a600",
"failoverTarget": "policy-ISL-TEST-8-1-0-0-144_dfs"
},{
"sourceId": "ISL-TEST-8-1-0-0-
146_0050569f65335842815c970993ede1cfa604",
"targetId": "ISL-TEST-8-1-0-0-
144_0050569ff3955440815c200978332081a600",
"failoverTarget": "policy-ISL-TEST-8-1-0-0-146_dfs_mirror"
}
]
}
}

Nginx example config

This is a sample configuration file to redirect port 80 to the node js server listening on 8080.

NOTE: this is an example contact your web server admin to configure for your environment and setup https.

```
#user  nobody;

worker_processes  1;



#error_log  logs/error.log;

#error_log  logs/error.log  notice;

#error_log  logs/error.log  info;



#pid        logs/nginx.pid;




events {

    worker_connections  1024;

}




http {

    include       mime.types;

    default_type  application/octet-stream;



    #log_format  main  '$remote_addr - $remote_user [$time_local] "$request" '
```

```
#                  '$status $body_bytes_sent "$http_referer" '
#                  '"$http_user_agent" "$http_x_forwarded_for"';

#access_log  logs/access.log  main;

sendfile        on;
#tcp_nopush     on;

#keepalive_timeout  0;
keepalive_timeout  65;

#gzip  on;

server {
    listen       80;
    server_name  localhost;

    #charset koi8-r;

    #access_log  logs/host.access.log  main;
```

```
location / {
            proxy_read_timeout   600s;
            proxy_pass http://localhost:8080/;
        }

        #error_page  404              /404.html;

        # redirect server error pages to the static page /50x.html
        #
        error_page   500 502 503 504  /50x.html;
        location = /50x.html {
            root   html;
        }
}
```

Usage

In order to issue calls to the server, the user has to be provided with the API token and the name of the failoverTargets they are authorized to operate on (as listed in `auth.json`). Once this is provided, end users can use curl to launch failovers:

```
curl -X POST http://<failover_gateway_ip>:8080/failover?apiKey=<api_key>&failoverTarget=<failoverTarget>
```

when launching failovers, the call above returns a json object with a jobId:

```
{"id":"job-1552504373143-255009881"}
```

together with the api key, the jobId can be used to query the status of running failovers:

```
curl http://<failover_gateway_ip>:8080/status?apiKey=<api_key>&jobId=<job_id>
```

© Superna LLC

1.30. How to convert VMware Eyeglass appliance and Migrate to Microsoft Azure

Home Top

How to convert VMware Eyeglass appliance and Migrate to Microsoft Azure

- Limitations of 3rd party products within this solution Example
- Overview
- Pre Setup Steps and Tools for upload to Azure
- Steps to Convert Eyeglass VMware appliance to VHD and automatically upload to Azure
- How to convert VHD to Managed Disk for Use in Azure
- How to Connect to Eyeglass and configure firewall ports on Eyeglass VM
- Optional - How to copy an existing VHD directly to a resource group storage blob

Limitations of 3rd party products within this solution Example

- References to PowerShell and Microsoft tools is provided as reference only and not included in product support.
- Issues with exact PowerShell examples or Azure procedures require a support case with Microsoft
- Reference links are provided "as is" with no support

- **NOTE: this document is not covered by support contract and is provided "as is" with 3rd party product expertise required**

Overview

This solutions guide assists with examples of how to move a VMware Eyeglass appliance into Azure to use Azure as a 3rd site. This solution example provides examples that must be changed to apply to your environment. It is expected that Azure expertise and PowerShell expertise is known to complete these steps.

Pre Setup Steps and Tools for upload to Azure

1. Install certificate remote management software:
    a. Download here https://azure.microsoft.com/en-gb/resources/samples/resource-manager-powershell-certificate-authentication/
    b. Install certificate authentication package for VM convert tool
2. Open PowerShell as administrator on PC OS that has PowerShell installed
3. Run these commands:
    a. Import-Module AzureRM.Resources
    b. Set-ExecutionPolicy RemoteSigned
4. Create Azure management certificate to allow remote tools to connect to Azure:

a. Copy and paste below to PowerShell prompt (NOTE: the cert file will be created in your temp folder in your profile C:\Users\username\AppData\Local\Temp>)

b. Import-Module PKI ## This will import public key infrastructure modules into Powershell

New-SelfSignedCertificate -DnsName "MVMC" -CertStoreLocation "Cert:\CurrentUser\My" ## Creates private cert in the current user's Personal certificate store

Export-Certificate -Cert (Get-ChildItem Cert:\CurrentUser\My\ -DnsName MVMC) -FilePath $env:TEMP\MVMC.cer ## Exports public cert created above using New-SelfSignedCertificate

Import-Certificate -FilePath $env:TEMP\MVMC.cer -CertStoreLocation Cert:\CurrentUser\Root ## Imports the cert into the current user's Trusted Root Certification Authorities store

c. Reference this link to if you need additional information http://adriank.org/generating-azure-management-certificate-microsoft-virtual-machine-converter-3-1/

d. Upload the certificate file .cer to Azure

a. 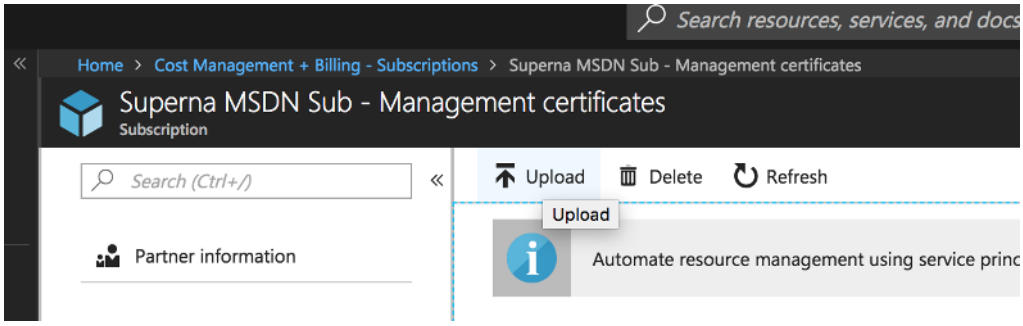

b. Record the Azure subscription id

c. Notice the hash value needed and record this value

a. 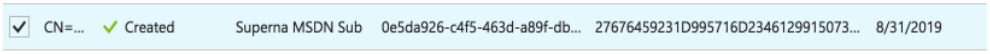

d. Download VMware to azure conversion tool and install it https://www.microsoft.com/download/details.aspx?id=42497

e. You will need azure subscription id and management key hash to authenticate to Azure to use this tool

Steps to Convert Eyeglass VMware appliance to VHD and automatically upload to Azure

1. You will need the following to complete these steps:
   a. Azure subscription id
   b. Management certificate hash
   c. Ip address for vCenter where Eyeglass VM is running
   d. **IMPORTANT: Configure Eyeglass VM for DHCP before converting**
2. Connect to the VM Console from VMware vCenter management interface:
   a. Login as admin user
   b. Type sudo -s
   c. Enter admin password

d. Type yast

e. Go to networking section edit Ethernet settings and switch to DHCP

i. 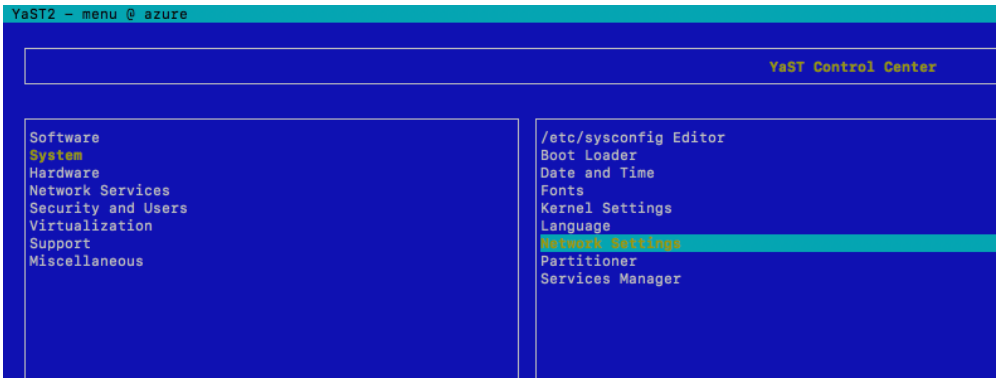

ii. Edit the Ethernet adapter (use tab key)

iii. Set to DHCP

iv. 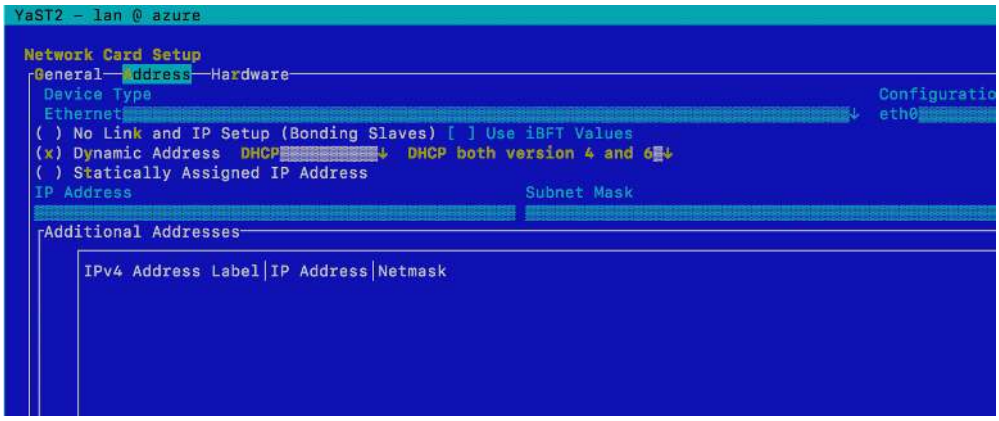

v. Save each screen and exit yast

f. Shutdown the Eyeglass VM

i. Type shutdown

ii. Verify its shutdown in vCenter

3.

1. **Open Microsoft Virtual Machine Converter**
2. **NOTE: Some VMDK's may require updates depending on the VMware version used (references are information only , support with Microsoft is required Reference)**

3. Additional tools for editing VMDK if required. Reference.

4. 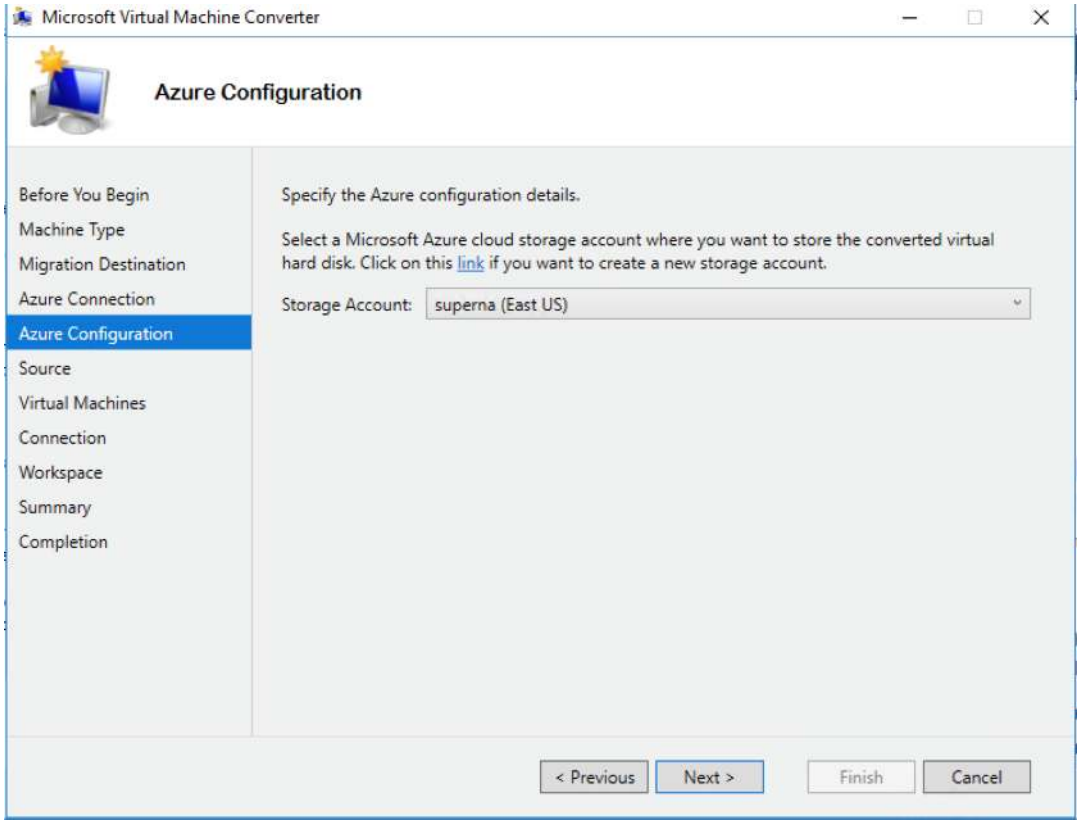

5. 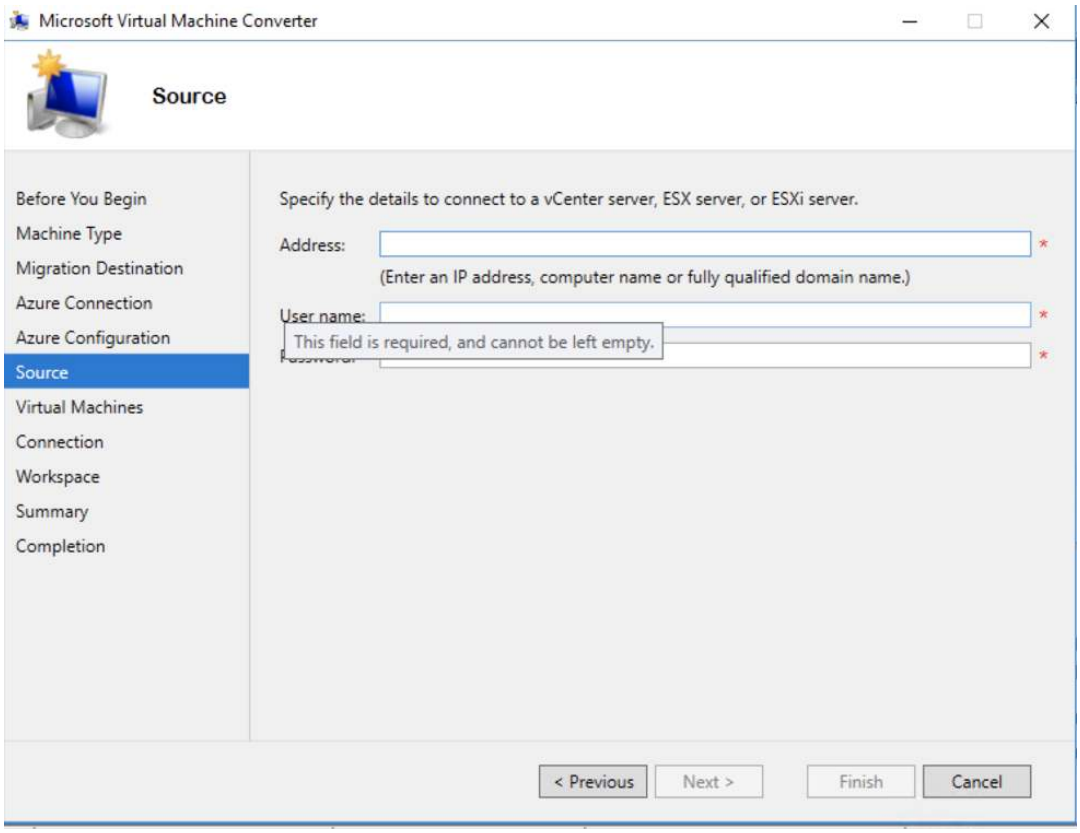

6. 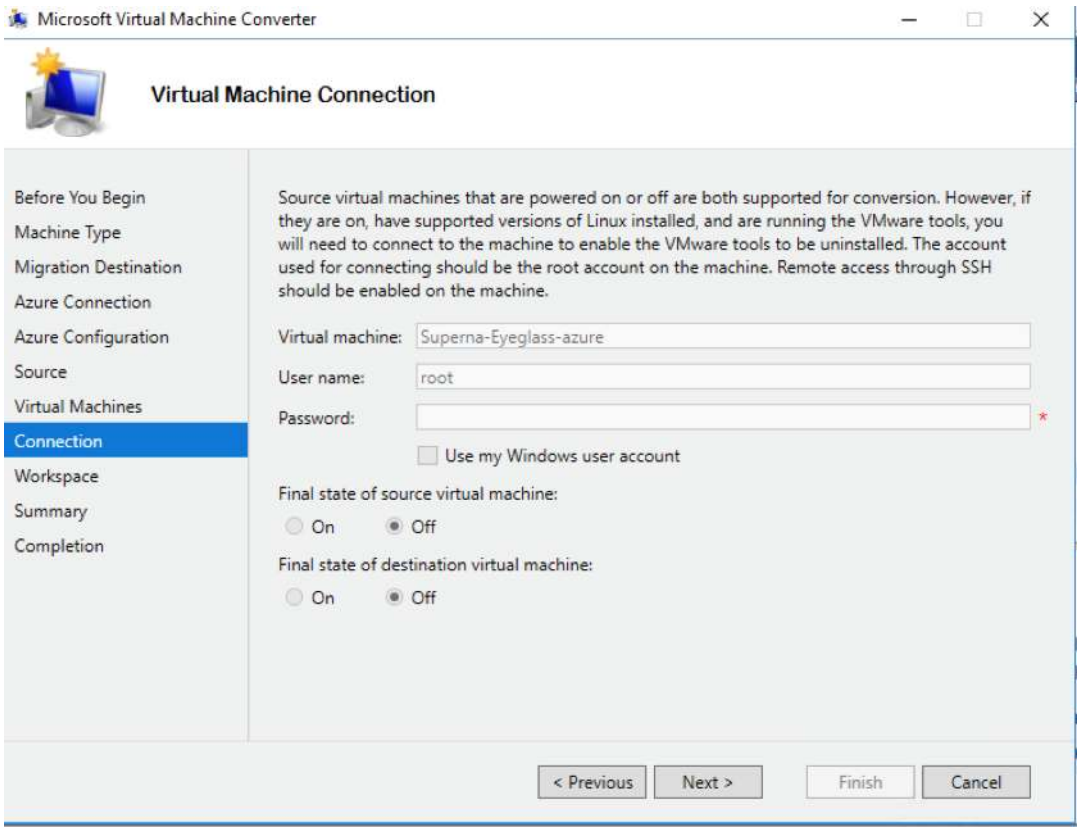

7. When selecting a location to convert ensure you have the disk space for the entire disk size of 80Gb.

8. 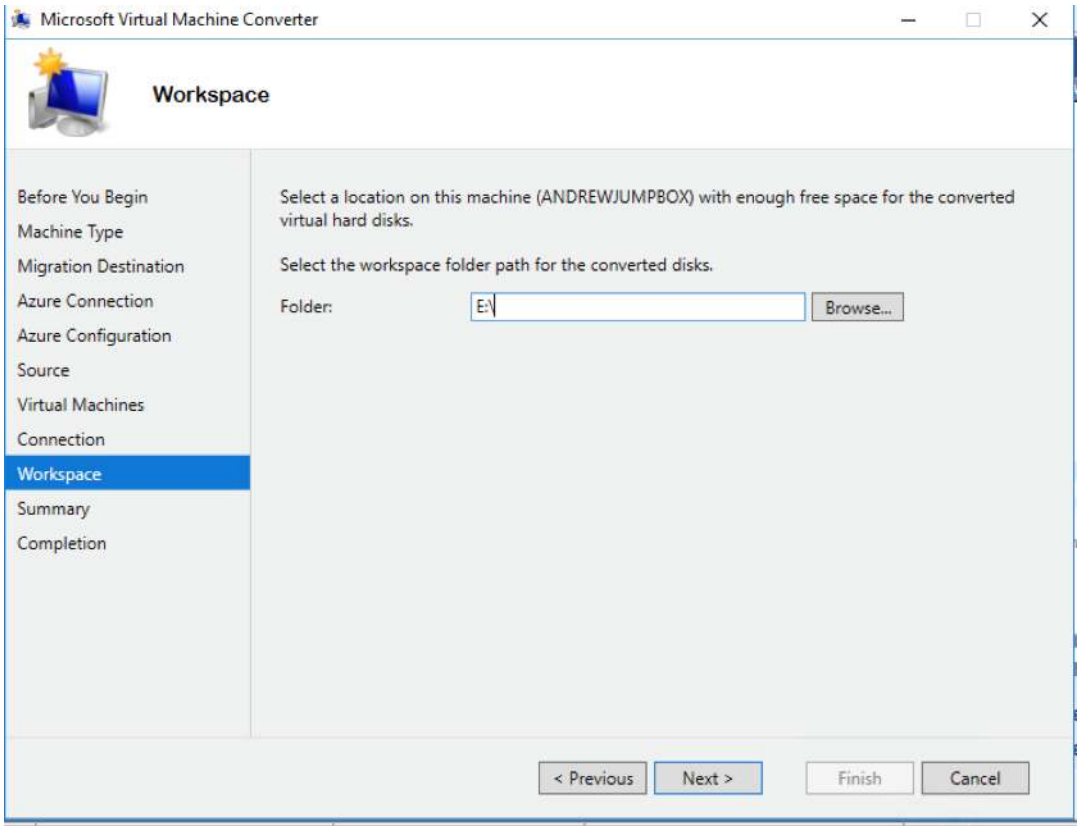

9.

10. 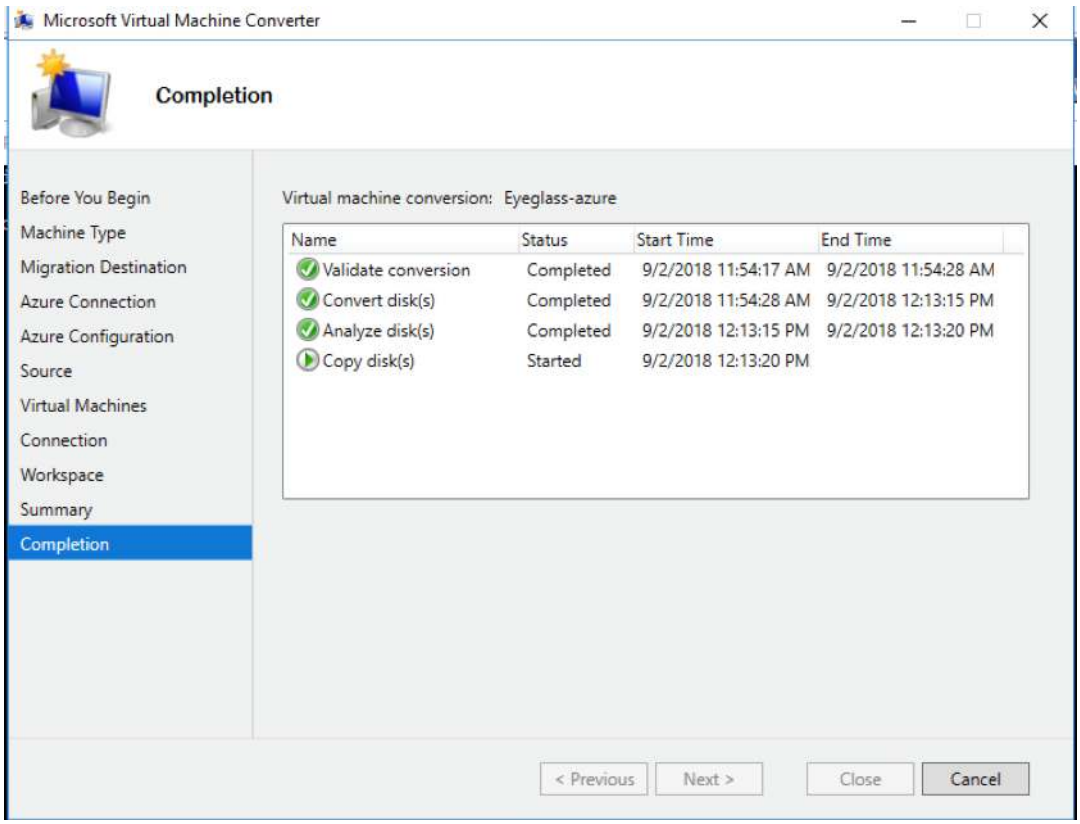

11. Wait for the last step copy Disks to complete. This step copies the VHD to a storage blob in Azure. **NOTE: This will take a long time to copy to Azure no progress is shown and can take an hour or more.**

12. **Done**.

How to convert VHD to Managed Disk for Use in Azure

1. After the previous step completes find the VHD in the storage blob **(note the VHD will be in a storage bucket named after the VM in VMware)**

2. 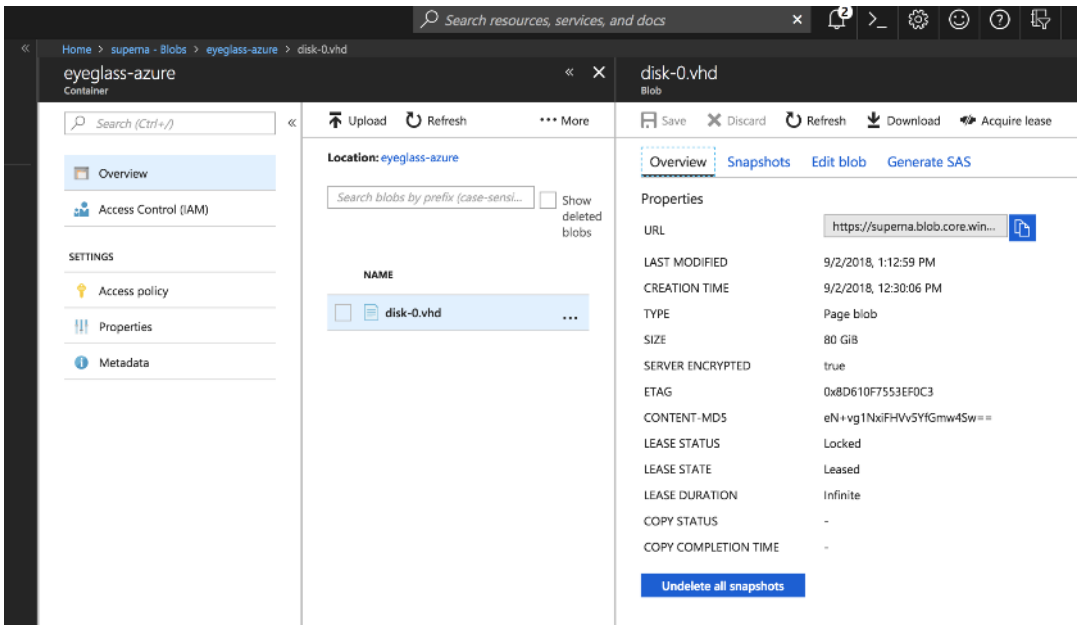

3. Copy the url of the VHD object from the overview screen to use with the command below.

4. Now execute the PowerShell commands to convert the VHD to managed disk in Azure. Copy and paste each line to PowerShell.

   a. Replace yellow highlights with options that make sense for your Azure environment (You will require Azure login account to complete this step) Reference link

   b. Connect-AzureRmAccount

   c. Install-Module AzureRM -RequiredVersion 6.0.0

   d. New-AzureRmResourceGroup -Location "East US" -Name RG01

   e. New-AzureRmDisk -DiskName eyeglass2 -Disk (New-AzureRmDiskConfig -AccountType Standard_LRS -Location 'East US' -CreateOption Import -SourceUri https://xxx.blob.core.windows.net/eyeglass-azure-2/disk-0.vhd) -ResourceGroupName 'RG01'

5. Now create the vm from Azure console (reference commands link)
6. Login to the Azure portal (https://portal.azure.com/)
   a. Open the Cloud Shell from the menu top right shows Azure CLI.
   b. 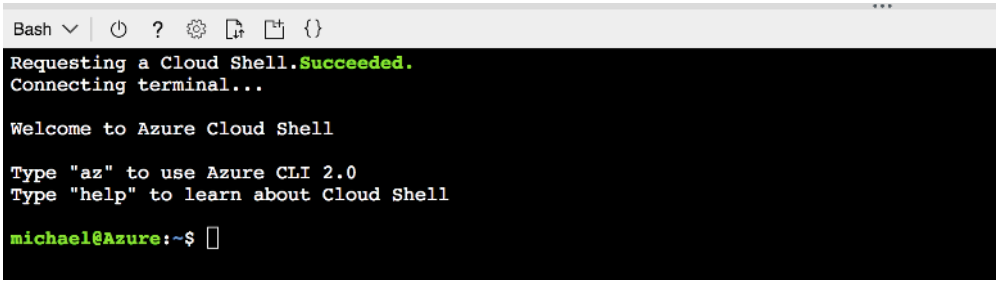

   c. Paste this command into the shell (command reference) NOTE: VM Size must be supported configuration shown in blue highlight)
   d. NOTE: Ensure you use the same resource group, location and disk name from the previous steps, yellow highlighted. The green highlight is the name of the VM once created you can change this name to anything you like.
   e. az vm create --resource-group RG01 --location eastus --name eyeglass-vm2 --os-type linux --attach-os-disk eyeglass2 --size standard_D4s_v
   f. 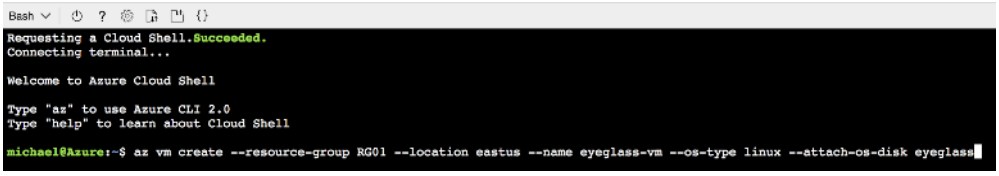

   g. Command completes with output showing the id of the VM that was created. This step can take 1-2 minutes to complete.
   h. Done create VM from managed disk steps.

1.

How to Connect to Eyeglass and configure firewall ports on Eyeglass VM

1. **NOTE: The previous step also starts the VM. Verify it is in a running status in the Azure Portal. DHCP option will pickup the private ip address from Azure, and map a public IP address to this VM, after running the command above**.
2. ssh is enabled by default test access to the VM by getting the public IP that was assigned on VM creation
3. Click the connect button and get the ssh ip address from the pop up window
4. 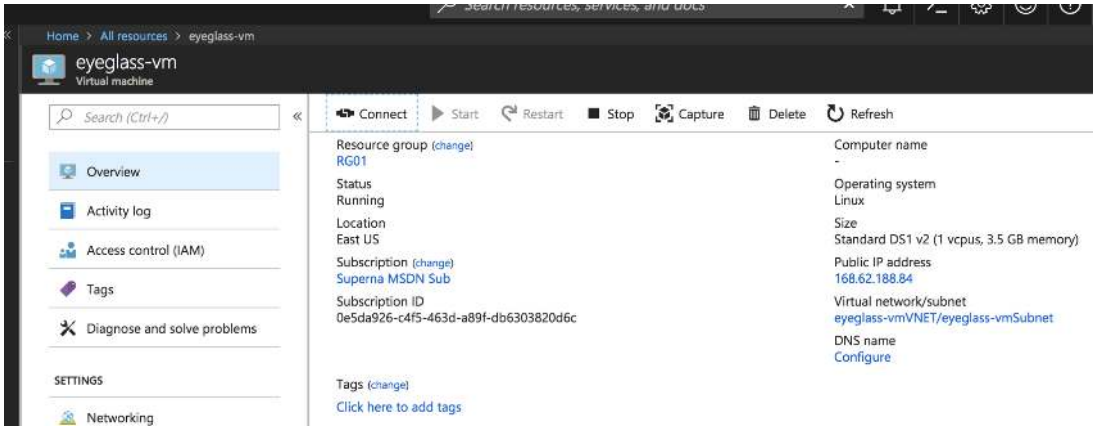

5. ssh admin@x.x.x.x where x.x.x.x is the public ip address
6. If successful ssh connection, login as the admin user to Eyeglass. This confirms the VM booted successfully and DHCP was assigned a private IP in Azure.
7. Now configure inbound GUI access over https port 443 TCP
    a. Open the Network Security Group associated to the Eyeglass VM.

b. 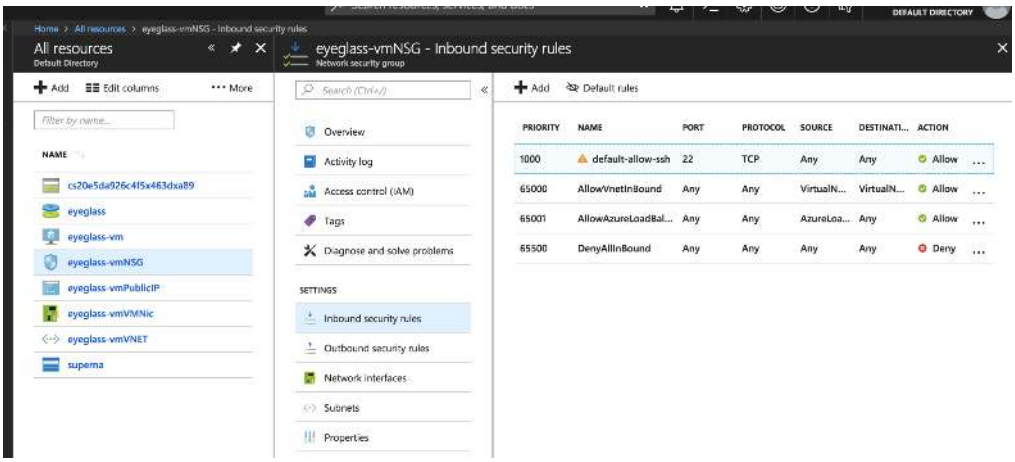

c. Add the inbound firewall rule as shown below (or customize for your security requirements) NOTE: Firewall example below is for testing only. You must factor in your security requirements. This example rule should be deleted after testing.

d. 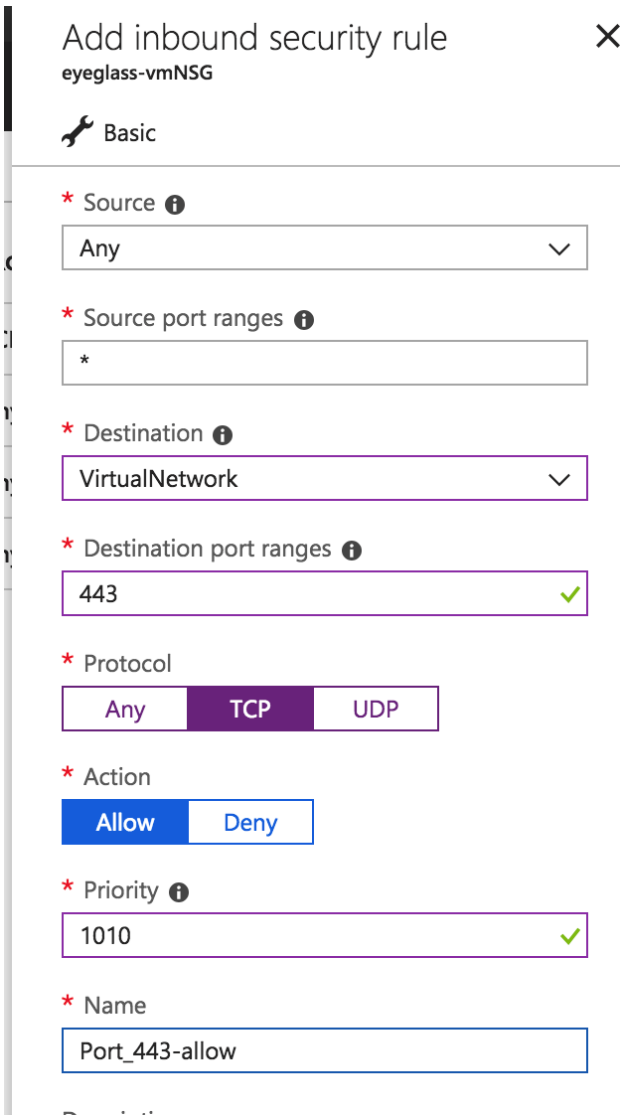

e. Test access to the web GUI

f. https://x.x.x.x the webUI should load successfully, where x.x.x.x is the public ip address

g. Now follow the firewall port configuration is required to allow Eyeglass to manage PowerScale and allow Eyeglass administrators to access the WebUI.

   i. This setup requires all ports and ip addresses required for Eyeglass are allowed. All ports and directions are outlined below.

ii. The scope to complete this in Azure is outside the scope of this guide and not covered by product support. Consult with an Asure system administrator.

h. **See the firewall ports guide here.**

i. **Done.**

Optional - How to copy an existing VHD directly to a resource group storage blob

```
Add-AzureRmVhd -ResourceGroupName ResourceGroup -
Destination https://xxx.blob.core.windows.net/vms `
-LocalFilePath "C:\Users\Public\Documents\Virtual hard disks\myVHD.vhd"
```

© Superna LLC

1.31. TLS Certificate Procedures for Eyeglass

Home Top

- How to replace self signed certificate on Eyeglass Appliance GUI - Quick Replace
  - Details:
  - Prerequisites:
  - Configuration Steps:
- How to Create Certificate Authority Root Cert on the Eyeglass appliance to Sign a Cert Request for the Eyeglass Appliance
- How to create a certificate Request in the Eyeglass Appliance for signing by an External Certificate Authority Server
- How to Install a signed certificate in an Eyeglass Appliance.
- How to Sign a Cert Request and Export a Certificate with Microsoft CA Server

How to replace self signed certificate on Eyeglass Appliance GUI - Quick Replace

Details:

The following procedure can be used to generate a new self signed certificate and apply it on the Eyeglass appliance.

Prerequisites:

NOTE: This only replaces the 443 main cert, if you want to replace the cert used for websockets and the WebUI self signed cert, follow the instructions here for an external CA signing process.

Configuration Steps:

Note: There will be an Eyeglass service interruption when performing this procedure.

1. SSH to the Eyeglass as admin
2. Default password is 3y3gl4ss
3. sudo su (to root)
4. Default password is 3y3gl4ss
5. systemctl stop sca
6. systemctl stop lighttpd
7. mv /opt/superna/sca/.secure/ssl.pem /tmp/ssl.pem.old
8. /opt/superna/bin/create_ssl_keys.sh /opt/superna/sca/.secure/ssl
9. chown sca:users /opt/superna/sca/.secure/*
10. systemctl start sca
11. systemctl start lighttpd
12. Done.

How to Create Certificate Authority Root Cert on the Eyeglass appliance to Sign a Cert Request for the Eyeglass Appliance

1. This procedure can be used if you do not have an external CA within your organization, and need to sign a Cert to change the certificate on Eyeglass without needing use an external CA. These steps will create a CA Root key, and CA Root cert on the appliance, and create a CA signing cert to be used for signing requests for the appliance.
2. ssh to Eyeglass as admin user
3. sudo -s (enter admin password)
4. mkdir -p /opt/ca
5. cd /opt/ca
6. **Create Root CA Key for Signing other Certificates**
    - openssl genrsa -passout pass:foobar -out rootCA.key 2048 (change the yellow highlight passphrase and store this value for use to sign certificates in the future. This is required to gain access to the CA key for signing)
    - Now we have a private root key(rootCA.key), and a root CA(rootCA.pem). If you want all the clients/PC/browsers to accept your authorized certificate, you need to put your root CA in their local trusted stores (e.g. OS's trusted certificates repositories)
7. **Self-sign the CA's signing certificate**
    a. openssl req -x509 -new -nodes -key rootCA.key -days 3650 -out rootCA.pem
        - **You will be prompted to enter information about the Root CA certificate and you should enter information**

for country, province, city etc. specific to your organization

- This CA Root Cert will be valid for 10 years , change yellow highlighted value for a different expiry.

8. Create the appliance Certificate Request and Sign it with the Root CA Certificate
   a. Create the private key
      - openssl genrsa -out eyeglass.key 2048
   b. Create the Certificate Request
      - openssl req -new -key eyeglass.key -out eyeglass.csr
      - NOTE: You will be required to enter information about your environment country, city, company, email, optionally enter a passphrase to protect the request.
   c. Sign the Request with the root CA certificate key and signing certificate created in the steps above
      - openssl x509 -req -in eyeglass.csr -CA rootCA.pem -CAkey rootCA.key -CAcreateserial -out eyeglass.cer -days 365
      - NOTE: Change days to allow extend how long the cert is valid
      - check the cert
        - openssl x509 -in eyeglass.cer -text -noout
   d. Now follow the instructions in this guide to install the certificate into the appliance.

How to create a certificate Request in the Eyeglass Appliance for signing by an External Certificate Authority Server

1. **NOTE: Use this procedure when you have an External CA server to sign certificates for your organization.**

2. First create a configuration file inside /tmp directory. You can named it "iglscert.cnf" in Eyeglass Appliance. Below is an example: **NOTE: the FQDN of the appliance should be used for the CN = property in the cnf file. NOTE: The alt DNS section should be setup to match the FQDN of the appliance and use * to wildcard the host name. Look at the yellow section below.** It is also possible to add the ip address here if you want to access by IP address use IP.1 = x.x.x.x syntax in the CNF file.

```
[ req ]
default_bits = 2048
prompt = no
encrypt_key = no
default_md = sha256
distinguished_name = dn
req_extensions = v3_req

[ dn ]
CN = iglscert.superna.local
emailAddress=support-team@superna.net
O = SUPERNA
```

OU = Support Team

L = Ottawa

ST = Ontario

C = CA

[v3_req]

subjectAltName = @alt_names

[alt_names]

DNS.1 = iglscert.superna.local

DNS.2 = *.superna.local

3. Now, create a CSR (Certificate Signing Request) file and a server key file in /tmp directory using the following command in Eyeglass Appliance: **NOTE: The path to the private .key file will be needed when installing the signed certificate in the next section.**
    - **openssl req -new -config /tmp/iglscert.cnf -keyout /tmp/iglscert.key -out /tmp/iglscert.csr**
4. Use the following command to verify the certificate information:
    - **openssl req -text -noout -verify -in /tmp/iglscert.csr**
5. Take the verified CSR file to your Windows Server CA or other CA and get it signed [Signed certificate must be in Base-64-encoded X.509 format] and **have the CER extension**. Once you have the file signed, copy it back to Eyeglass Appliance using WinSCP, and install it using the steps below.

- See the section at the end on how to use Microsoft Certificate Authority to sign a CSR request and download and exported CER formated x.509 certificate. **This is an example only consult vendor documentation for signing certs with your organizations CA server.**

6. After receiving the CER format certificate signed by your CA server follow the instructions here to install a signed certificate.

How to Install a signed certificate in an Eyeglass Appliance.

1. Get your **certificate in .cer format** to complete this procedure. Follow the steps here to create a certificate request.

2. Locate the private key (.key file in the tmp directory on Eyeglass from the steps completed above) and certificate (.cer file from the CA used to sign the CSR request). The file should have a private X.509 key and certificate signed by a trusted certificate authority. **It must be X.509 Certificate CER format.**

    - **Example: iglscert.key (created from CSR steps above) and iglscert.cer (exported signed certificate file) from Microsoft CA or 3rd party CA.**

3. Login to Eyeglass

    - ssh as admin user

- sudo -s (enter the admin password)
- Copy the certificate file to Eyeglass using a tool like WinSCP tool

4. **Follow these instructions if you used the Eyeglass Root CA procedures to Generate the Self Signed Cert for the Appliance**
    - type **whoami (make sure you are the root user before proceding)**
    - scacli replace-certificate --privateKey=/opt/ca/eyeglass.key --certificate=/opt/ca/eyeglass.cer
    - **Skip to step #6**

5. **SKIP If External CA was not Used: Follow these Steps if you had the Certificate Request Signed by an External CA Server, otherwise**
    - Strip the key file and convert it to PEM format by executing below command **(note the .key is the private key created from the create CSR request completed in the steps above and stored in /tmp/iglscert.key path when creating the CSR request and private key)**
    - **openssl rsa -in /tmp/iglscert.key -out /tmp/iglscert.pem**
    - Now replace the certificate with existing Eyeglass cert:
    - **scacli replace-certificate --privateKey=/tmp/iglscert.pem --certificate=/tmp/iglscert.cer**

6. Browse the Eyeglass certificate directory:

- **cd /opt/superna/sca/.secure**

7. Move the existing .pem file:
   - **mv ssl.pem ssl.pem.orig**
8. Concatenate the new key file information into a single private key + certificate needed for lighttpd web server:
   - **cat ssl.pem.orig ssl > ssl.pem**
9. change file ownership:
   - **chown sca.users /opt/superna/sca/.secure/***
10. Restart Lighttpd and sca service:
    - **systemctl restart lighttpd sca**
11. Now, login to Eyeglass Web UI and use the FQDN to access and verify the Certificate in a Browser. Use any browsers view certificate feature to verify the correct certificate and expiry dates and showing correctly. **NOTE: the FQDN used to access Eyeglass should be the value set in the CSR request fields under [alt_names])**
12. done

How to Sign a Cert Request and Export a Certificate with Microsoft CA Server

Follow this procedure if you are using a Microsoft CA for signing certificates.

1. ssh to Eyeglass as admin user
2. sudo -s (enter admin password)
3. cat /tmp/iglscert.csr
4. Copy all of the text in this file from the ssh session, and create a file on the PC that has access to the Microsoft Certificate Authority administration GUI. Name the file "iglscert.req", paste the contents of the iglscert.csr to this file and save it.
5. Right click the CA server name and Select **Submit New Request:**

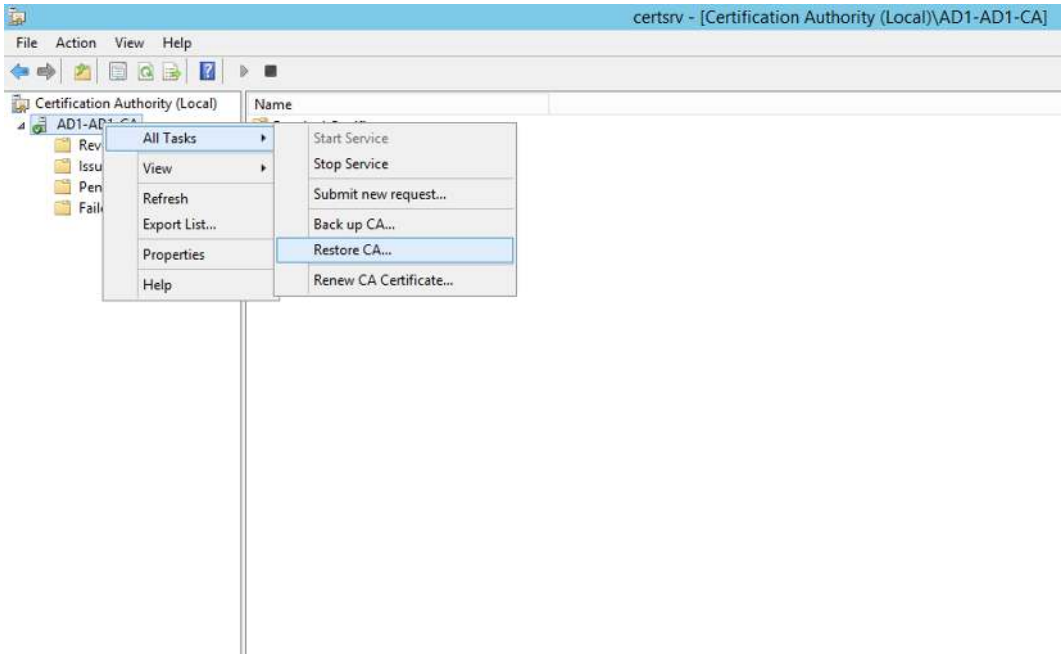

   -
   - Browse to the file you created iglscert.req and submit the request.
6. Click on the pending folder for the Certificate Authority.
7. 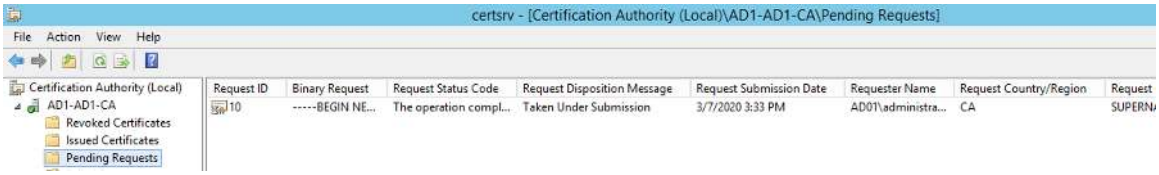

8. Right Click the pending request All Tasks and click the Issue option.

9. 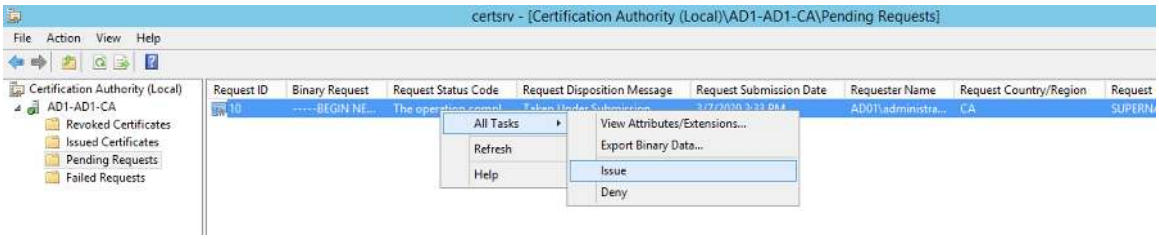

10. Now click on the Issued Certificates folder. Find the Issued Certificate and double click the Cert to display the cert, then select the **Details tab.**

11. 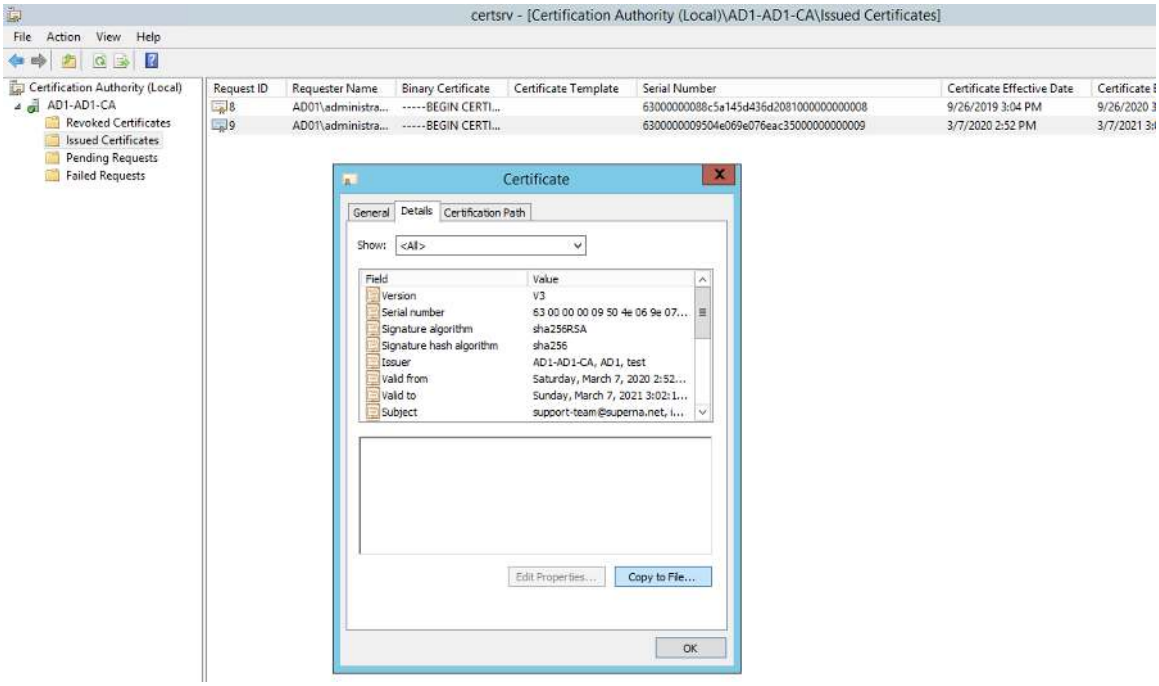

12. Click the Copy to File option to open the export Cert Wizard. Select the Base 64 x.509 CER format option

13.

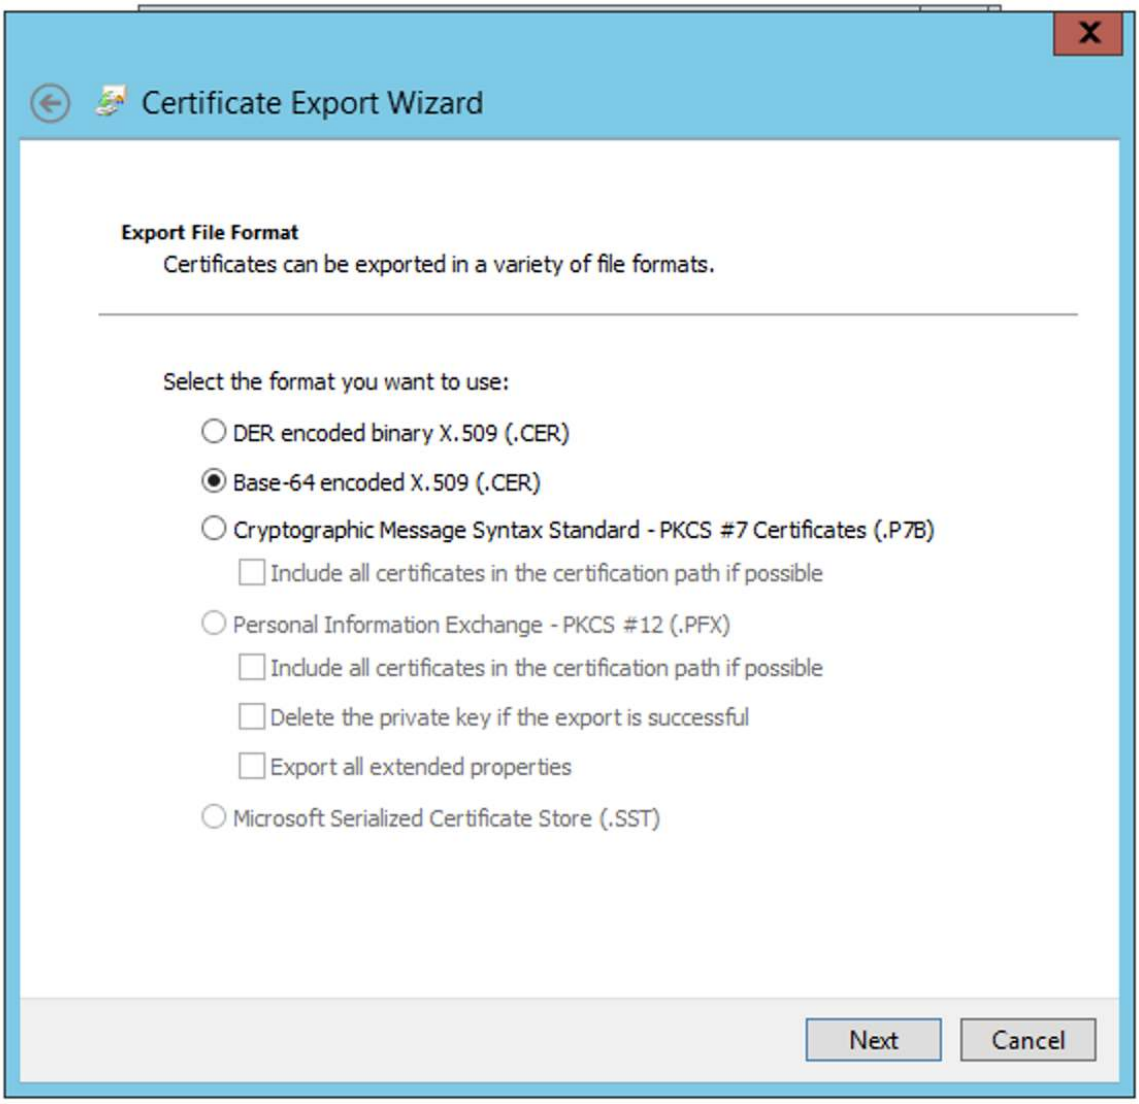

14. Save the CER file iglscert.cer. Then follow how to install a signed cert instructions above.

© Superna LLC

1.32. Configuring Eyeglass with dual NIC environments

Home Top

Configuring Eyeglass with dual NIC environments

This solution can be used for customers that do not allow VM's to access both clusters from a single network IP interface due to security or network design. The solution below allows Eyeglass to be modified to support dual NIC solution and static routes to reach remote clusters.

Prerequisites:

1. SSH to Eyeglass as admin, sudo -s
2. shutdown
3. Follow VMware steps with vCenter vSphere to add a 2nd nic to Eyeglass, use e1000 nic type
4. Attach to 2nd virtual switch
5. start up the VM

Note: Eyeglass must be able to login via PAPI API and SSH over these ip networks.

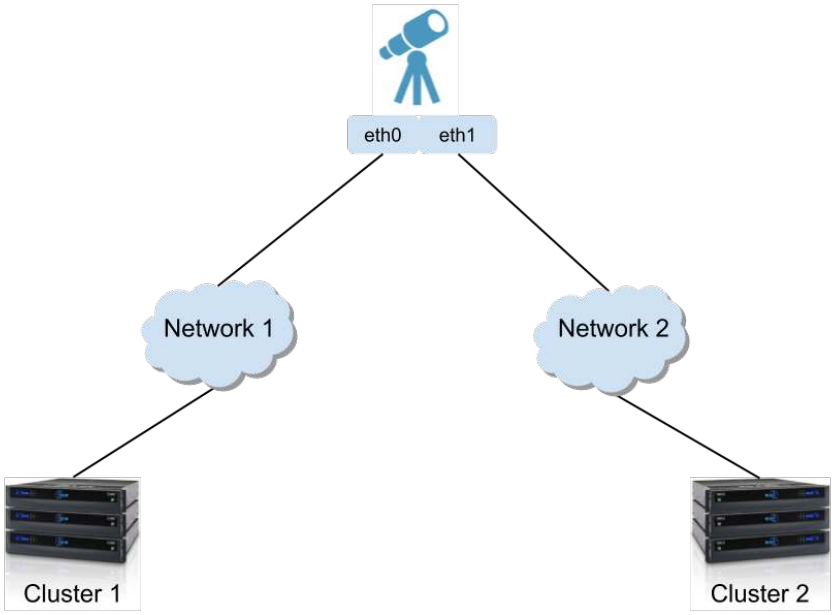

1. SSH to Eyeglass appliance login as admin then sudo -s to become root user.
2. Assume a single Eyeglass with two NICs, one configured for the address 10.105.16.100; the other for 10.105.17.101. To added the other route please follow the next steps to add your route .
3. Type the command yast2 , then tab to reach network settings > , then tab to reach Routing.
4. On Routing please tab to reach Add button and hit enter .
5. Add your destination IP , Subnet Mask (hint : Use \24 instead of 255.255.255.0 or your own subnet mask in the same format otherwise you'll have an error) then tab to reach Ok, then Hit enter .
6. Now can see the routes on the table. (note: These are persistent routes).
7. Change the cluster network setting as needed for your environment. Note: this simple network has the cluster on the same subnet as the Eyeglass NIC's. Typically a router would be the next top address in the static routes for one cluster, and the default route 0.0.0.0 would be used to reach the other cluster.

© Superna LLC

1.33. Changing Eyeglass UI Behavior

Home Top

- Overview
- How to Disable the Twitter Software Notification Feed on the Login Page
- How to extend the UI API timeout for long running commands example Unlock My files

Overview

The Eyeglass UI behavior has some capabilities to change its behavior. The sections below cover these modification steps.

How to Disable the Twitter Software Notification Feed on the Login Page

1. Requires 2.5.6 or later
2. Login to Eyeglass using ssh as admin
3. sudo -s (enter root password)
4. nano /srv/www/htdocs/eyeglass/js/eyeglass_globals.js

5. Locate the tag twitter_feed_enabled:!0 and change to twitter_feed_enabled:0 (by removing the ! character before the 0)
6. press ctrl key + w and type twitter to locate the correct section.
7. Use arrow keys to the ! character and delete this character
8. Save the file ctrl key + x
9. Type Y to save the file
10. Refresh the UI browser and the twitter feed will disappear.
11. **NOTE: This modification will not be persistent over upgrades**

How to extend the UI API timeout for long running commands example Unlock My files

1. When Eyeglass manages a large number of objects, some of the GUI windows or long running search for files with unlock my files with the Cluster Storage Monitor product. The timeout needs to be changed
2. Default value is set to 45 seconds and if you need to change these parameters editing the file Eyeglasss_globals.js as follows:
3. Login to Eyeglass via ssh
4. vim /srv/www/htdocs/eyeglass/js/eyeglass_globals.js
5. Type the / character to search and type 45 , then press enter so they cursor stays at the current location, the section should show this text (**{ajax_get_timeout_seconds:45**)

6. Type x key twice to delete the 45
7. Type i (to start an insert operation)
8. Type 90 (for 90 seconds)
9. Press Esc key to exit insert mode
10. Type :wq (colon then wq to write and save)
11. Now refresh the UI page and the timeout will now be 90 seconds for a long running request.
12.

© Superna LLC

1.34. Custom email routing by application or alarm subject contents

Home Top

- Overview
  - Limitations and Unsupported Configurations
- Requirements
- How to switch to Eyeglass mail routing based on postfix OS SMTP Relay
  - Eyeglass - Switch to use postfix SMTP local OS mail relay service no Authentication
  - Adding Filtering and Forwarding Rules to Postfix
  - Example 1: Send All Ransomware Alarms to a specific email (person or group email) AND Discarding Or Redirecting Security Guard daily Self Test Emails:
  - Example 2: Send Only alarms when a user is locked out by Ransomware Defender to a specific email (person or group email) but will not be sent to the Eyeglass admin defined to receive all alarms
  - Example 3: Send Easy Auditor Reports to a specific user or group email , note only this user or group will receive report emails
  - Example 4: Drop all Easy Auditor Report emails
  - Example 5: Send Easy Auditor Triggers to a specific email OR send a specific search report to a specific user

    - Mandatory Step - Activate Filtering Rules for subject or body of email rules
- How to edit Filter rules
- How to Test your new filter or forwarding rule
- How to convert previous Opensuse OS postfix main.cf to Opensuse 15.3 or later Postfix format
- Advanced Postfix Configuration For SMTP Authentication and TLS Configuration
    - How to Debug Postfix mail relay Issues

Overview

This solution guide explains how to configure custom email routing of specific alarms or notifications within Eyeglass. This is most commonly used for Easy Auditor or Ransomware Defender to route email notifications to a specific email or group distribution email. This also ensures other system level alarms are not sent to these emails. The steps below explain how to setup postfix email routing options. It is also possible to drop emails silently but the alarm is still visible in the GUI Alarms Icon.

Limitations and Unsupported Configurations

1. When a rule matches your criteria it will exit and no longer match any other rules listed in the configuration files.

2. Redirection - Multiple emails on a rule or multiple rows with different emails is unsupported. Redirecting emails should always use a group email.
3. Any example not listed below is not supported.

Requirements

1. Opensuse 15.3 OS Latest OS is recommended .

How to switch to Eyeglass mail routing based on postfix OS SMTP Relay

1. Most Eyeglass deployments use Notification center to enter SMTP details of your mail server. These steps will switch to a local SMTP engine in the operating system.
    a. **Requirements**:
        i. This example assumes you are using anonymous non authenticated SMTP over port 25.
        ii. The advanced section below covers authentication + TLS configuration
    b. **Steps**
        i. **Setup Eyeglass OS SMTP to Send mail to your mail server**
            1. ssh to Eyeglass as admin
            2. sudo -s (enter admin password)
            3. Edit the postfix setting: nano /etc/postfix/main.cf

4. control+w to search for the word **relayhost , to locate the correct instance without a comment. To find the line that is not commented (no # at the front of the line). You will need to press control+w [enter], repeat this 8 times to find the very last occurrence of relayhost that does not have the # comment in-front of the line.**
5. Edit the relayhost parameter (NOTE: leave the square brackets are required as per example below)
    a. relayhost = [**DNS or ip of your SMTP mail server**]:25 (leave the square brackets)
6. Add this value to the bottom line in the file.
    a. This will allow redirected emails to show the new email in the To field for emails that are redirected.
    b. **enable_original_recipient = no**
7. control+x answer **y** and enter to save and exit
8. Restart postfix service
    a. systemctl restart postfix
    b. Checked that postfix service is running: **systemctl status postfix**

ii. Eyeglass - Switch to use postfix SMTP local OS mail relay service no Authentication

1. From Eyeglass UI => Notification Center => Configure SMTP => Outgoing Email Server Information
2. Host Name change to: **localhost**
3. Port should stay set to: 25
4. From email can stay the same or change: [any e-mail address] e.g. eyeglass@<your domain>
5. Specify test e-mail recipient and click TEST E-mail Setting. Expect to receive test e-mail successfully.
6. Do not proceed until this step is successful with a test email being received.
7. For Advanced email configuration with authentication and TLS cipher control.
8. Done - No further steps are required unless email redirect rules are needed.

iii. Adding Filtering and Forwarding Rules to Postfix

1. This requires adding rules for filtering and forwarding emails based on the subject of the alert email or the body. These are provided as common examples below.
2. login as admin user over ssh
3. sudo -s (enter admin password to become root user)

4. **Preparing Configuration files for Content Filtering Rules (Mandatory Step)**
    a. Run these commands to enable content filtering on both subject of emails and the body of emails.
        i. **touch /etc/postfix/body_checks**
        ii. **touch /etc/postfix/header_checks**
        iii. **postconf -e "body_checks = regexp:/etc/postfix/body_checks"**
        iv. **postconf -e "header_checks = regexp:/etc/postfix/header_checks"**
5. **How to Edit the Rules files for email header_checks and for email body_checks**
    a. This section explains which files to edit depending on how you want to filter or forward alarm or report emails. See specific examples in the sections below that you can use for specific scenarios.
    b. Type the command below to edit email subject filter file
        i. **nano /etc/postfix/header_checks** (Use this file for email subject line filtering)
        ii. **nano /etc/postfix/body_checks** (only edit this file if you need email content filtering rules)

c. **Requirements for email Filtering Rules**

i. Under line 1 hit enter a few times to make space to add filters. **Note: you can add more than one filter with different emails and conditions**

ii. **Each rule must be on its own line.**

iii. **Rules are processed in the order listed in the file and the first match will exit the filtering logic.**

iv. **The syntax uses regex syntax when creating filters and actions. A full list is defined here on the postfix man pages.**

iv. Example 1: Send All Ransomware Alarms to a specific email (person or group email) AND Discarding Or Redirecting Security Guard daily Self Test Emails:

1. The example will send all Ransomware alarms to a specific email regardless of the configured Notification center alarm configuration and allow discarding or redirecting Security Guard test emails. This example requires body content matching.

a. Prerequisite: Switch Eyeglass to use post fix and add your mail server to the relayhost property. See instructions above.

b. NOTE: this will match the string in bold and will send to a different email address.

c. This also means that alarms for this product will not be sent to the Eyeglass administrator.

d. Security Guard emails will be discarded or forwarded

e. Example assumes Security Guard Cluster user name is "igls-sg," change this value depending on the name of the Security Guard user configured in Ransomware Defender.

2. Discard Security Guard and Forward all other Ransomware Alarms to an specific email

   a. nano /etc/postfix/body_checks (place each value below on it's own line the order matters to correctly discard first)

      i. /igls-sg/ DISCARD

      ii. /Ransomware Defender/ REDIRECT customer_email@domain.com

   b. Continue to activate step

3. **OR to Forward Security Guard emails to email A and forward All other Ransomware Alarms to Email B**
    a. **nano /etc/postfix/body_checks (place each value below on it's own line)**
        i. /igls-sg/ REDIRECT customer_email_A@domain.com

        ii. /Ransomware Defender/ REDIRECT customer_email_B@domain.com
4. control+x answer y and the enter key to save and exit
5. Continue to activate step

v. Example 2: Send Only alarms when a user is locked out by Ransomware Defender to a specific email (person or group email) but will not be sent to the Eyeglass admin defined to receive all alarms

1. **Prerequisite**: Switch Eyeglass to use post fix and add your mail server to the relayhost property. See instructions above.
2. nano /etc/postfix/header_checks (this file for email subject line filtering)

3. /^Subject: .*Locked/ REDIRECT xxxx@domain.name
4. control+x answer y and the enter key to save and exit
5. Continue to activate step

vi. Example 3: Send Easy Auditor Reports to a specific user or group email , note only this user or group will receive report emails

1. **Prerequisite**: Switch Eyeglass to use post fix and add your mail server to the relayhost property. See instructions above.
2. nano /etc/postfix/header_checks (this file for email subject line filtering)
3. /^Subject: Easy Auditor Report/ REDIRECT **xxxx@domain.name**
4. control+x answer y and the enter key to save and exit

vii. Example 4: Drop all Easy Auditor Report emails

1. **Prerequisite**: Switch Eyeglass to use post fix and add your mail server to the relayhost property. See instructions above.
2. nano /etc/postfix/header_checks (this file for email subject line filtering)

3. /^Subject: Easy Auditor Report/ DISCARD
4. control+x answer y and the enter key to save and exit
5. Continue to activate step

viii. Example 5: Send Easy Auditor Triggers to a specific email OR send a specific search report to a specific user

1. **Prerequisite**: Switch Eyeglass to use post fix and add your mail server to the relayhost property. See instructions above.
2. If you want to send a custom trigger or a saved report to a specific user email or group email, you first need the saved report name or the trigger name.
3. Get trigger names. We recommend using the word trigger in all triggers and then a name after to make matching alerts easier. Alerts will include the trigger name in the body of email.

a. 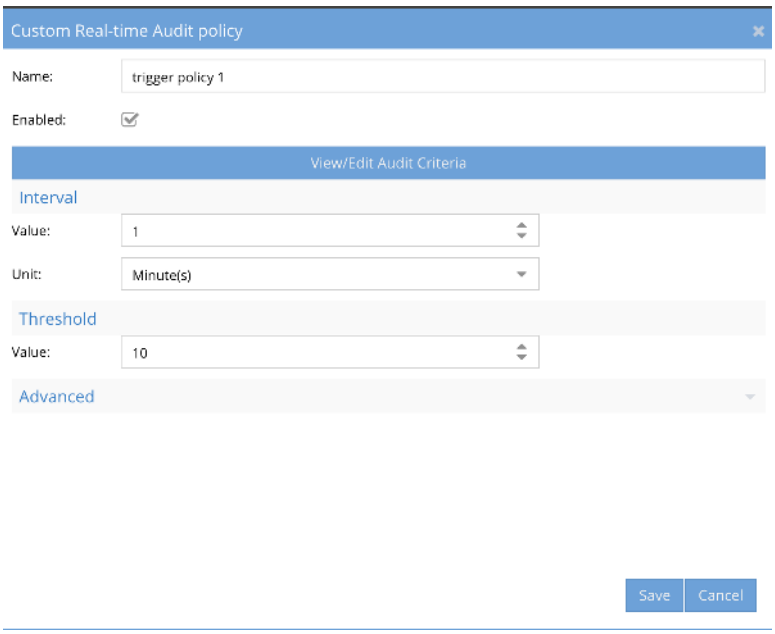

b.

4. For a content filter you must add the rule to the body_check file. See 3 different examples below.
5. nano /etc/postfix/body_checks (only edit the file if you need content filtering rules, see examples below)

a. /trigger/ REDIRECT xxxx@domain.name **(sends all triggers with the name trigger in the name)**

b. /trigger policy 1/ REDIRECT xxxx@domain.name **(sends all triggers with a specific name of a trigger in this example the trigger name is "trigger policy 1")**

c. /departmentXReport/ REDIRECT xxxx@domain.name **(sends all Easy Auditor report results with a saved report run manually or on a schedule with a name of "departmentXReport")**

d. 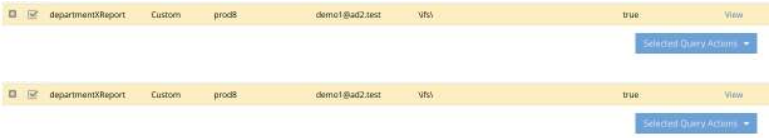

e. Continue to activate step

ix. Mandatory Step - Activate Filtering Rules for subject or body of email rules

1. Reload the rule set to take effect and restart postfix process

   a. **postfix reload**

   b. **systemctl restart postfix**

   c. **done**

How to edit Filter rules

1. Each time you edit the filter files
    a. nano /etc/postfix/body_checks (email body rules)
        i. OR
        ii. nano /etc/postfix/header_checks (email header rules)
    b. **You must restart postfix before any changes will take effect.**
    c. Activate filter rules

How to Test your new filter or forwarding rule

1. The examples that use Easy Auditor reports or Ransomware Defender can be tested using Security Guard and Robot Audit run now option to trigger a new alarm.
2. To verify your rule worked you can tail the mail routing log to see the rule rewrite the original email recipient with new email address.
3. tail -f -n 100 /var/log/mail (this command will monitor the mail log during the test)
4. or you can search the log
5. grep orig_to /var/log/mail* (this will locate all entries that the redirect rule triggered and shows the original email and new email used)
6. Example output shows to original email and the new to email as per below.
    a.
```
Feb 26 07:56:52 anycopy postfix/smtpd[53186]: disconnect from localhost[127.0.0.1] ehlo=1 mail=1 rcpt=1 data=1 quit=1 commands=5
Feb 26 07:56:52 anycopy postfix/qmgr[52534]: 658EE3812A: from=<demo@superna.net>, size=1737, nrcpt=1 (queue active)
Feb 26 07:56:52 anycopy postfix/smtp[53188]: Untrusted TLS connection established to 192.168.1.250[192.168.1.250]:587: TLSv1 with cipher AES128-SHA (128/128 bits)
Feb 26 07:56:52 anycopy postfix/smtp[53188]: 658EE3812A: to=<[redacted]gmail.com>, orig_to=<[redacted]@superna.net>, relay=192.168.1.250[192.168.1.250]:587, delay=0.08, d
lays=0.02/0/0.01/0.05, dsn=2.6.0, status=sent (250 2.6.0 <1048304108.2931161434421241 5.JavaMail.sca@anycopy> Queued mail for delivery)
Feb 26 07:56:52 anycopy postfix/qmgr[52534]: 658EE3812A: removed
```
7. Done

How to convert previous Opensuse OS postfix main.cf to Opensuse 15.3 or later Postfix format

1. This procedure is for backup and restore of an appliance running 15.1 or 15.2 to a 15.3 or later OS that had postfix configured. Release 2.5.8 backup includes the main.cf, body and header check files. **NOTE: You may need to get these file from your old appliance if the old appliance release is < 2.5.8**
2. Login to the new appliance and follow the steps below.
    a. ssh admin
    b. sudo -s
    c. unzip the Eyeglass backup used to restore the new appliance configuration on a Window PC , locate the postfix backup files and open main.cf in notepad and locate the relayhost line that includes your mail server.

i. 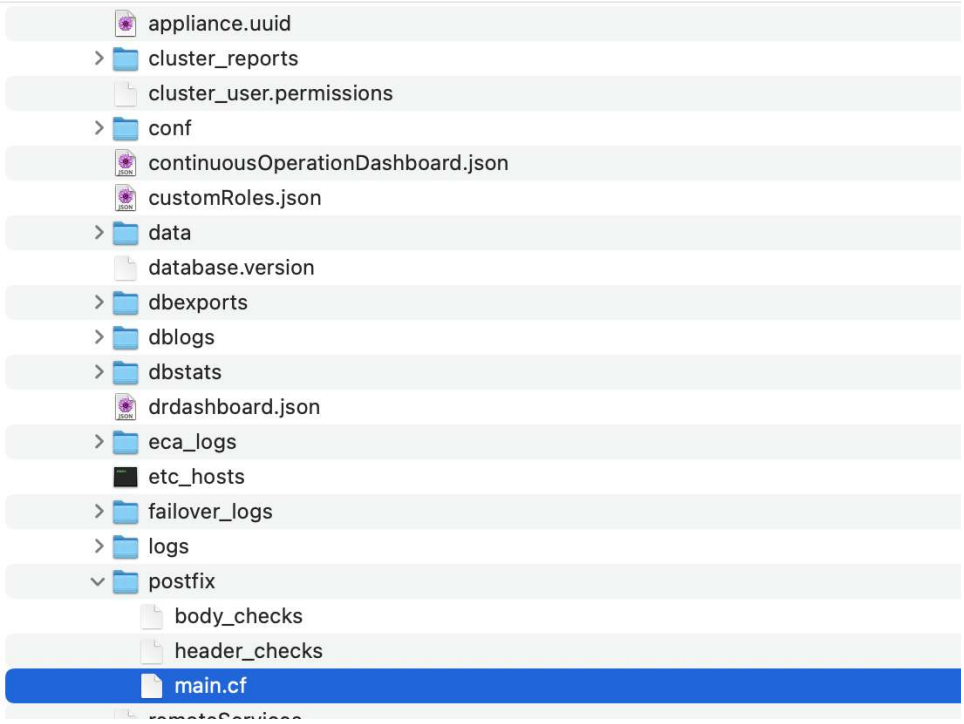

d. Edit the postfix file on the new appliance

   i. **nano /etc/postfix/main.cf**

   ii. control + w enter relayhost press enter, you will need to search 8 times to find the the line without # comment

   iii. Fix the relayhost field with the values from the main.cf opened in notepad from the backup file.

      1. Add the value from the backup to the relayhost = line in the file opened from the ssh session with nano editor

      2. relayhost = [**DNS or ip of your SMTP mail server**]:25 (leave the square brackets)

e. Re-apply body and header checks prepare new appliance

   i. Open the header_checks and body_checks files from the backup zip file in notepad

ii. From ssh session to new appliance

1. **touch /etc/postfix/body_checks**
2. **touch /etc/postfix/header_checks**
3. **postconf -e "body_checks = regexp:/etc/postfix/body_checks"**
4. **postconf -e "header_checks = regexp:/etc/postfix/header_checks"**

f. Re Apply previous forwarding filtering rules

i. **nano /etc/postfix/body_checks** (paste body_checks from notepad to this file)

ii. control + x answer yes to save and exit

iii. **nano /etc/postfix/header_checks** (paste header_checks from notepad to this file)

iv. control + x answer yes to save and exit

g. Activate the configuration

i. **postfix reload**

h. Test email in Eyeglass and test your rules

Advanced Postfix Configuration For SMTP Authentication and TLS Configuration

1. This section allows configuration of authentication to your mail system and control of TLS options.
2. **NOTE: Authentication + TLS is the only supported configuration.**

3. **NOTE: This assumes you have configured the relay and basic settings in the section above for switching to postfix MTA.**
4. ssh as admin user
5. switch to root
    a. sudo -s (enter admin password)
6. Fix the host file to provide a host name for localhost used by postfix mail relay
    a. nano /etc/hosts file, add the Eyeglass hostname fqdn in the localhost section. example:

    127.0.0.1 localhost igls01 igls01.ad1.test
7. Edit the postfix setting for Authentication.
    a. nano /etc/postfix/main.cf
    b. Verify the **relayhost** parameter is set already following steps in the basic configuration section above and change the port to use the TLS port. control-w type relayhost [enter] and repeat 8 times to find the correct entry.
        i. relayhost = [x.x.x.x]:587 (leave the square brackets)
        ii. Replace x.x.x.x with your mail relay host IP or FQDN
    c. Configure Authentication
        i. edit the following parameters for TLS
        ii. nano /etc/postfix/main.cf
            1. press control+w enter sasl [enter] to locate the section for authentication configuration. Set the settings as per below
        iii. **smtp_sasl_auth_enable = yes**

iv. **smtp_sasl_security_options = noanonymous**

v. **smtp_sasl_password_maps = lmdb:/etc/postfix/sasl_passwd**

vi. **smtpd_sasl_auth_enable = no**

vii. Create the password file /etc/postfix/sasl_passwd replace the yellow values to match your environment before creating the file.

1. nano /etc/postfix/sasl_passwd (enter the line below and change the yellow highlights to match your environment.
2. [x.x.x.x]:587 user:password
   a. x.x.x.x is replaced with the ip address or host name of your mail relay host
   b. This assumes port 587 is used on your mail relay host which is the default TLS SMTP port. Change the port to match your mail server.
   c. User is the user name for authentication. Example user@domain.com
   d. Password is the password for this service account user.
3. Convert /etc/postfix/sasl_passwd into a format that Postfix can read and remove clear text password file:
   a. postmap lmdb:**/etc/postfix/sasl_passwd**

b. This creates a file **/etc/postfix/sasl_passwd.lmdb**

c. You can now delete or remove the clear text password in the /etc/postfix/sasl_passwd

d. **rm /etc/postfix/sasl_passwd** (removes clear text password file, since it is not used by postfix. The DB file is used to read the user and password.

4. Secure the /etc/postfix/sasl_passwd.lmdb file.

   a. The file must be owned by root, and no one else should have read access to that file

      i. chown root:root /etc/postfix/sasl_passwd.lmdb

      ii. chmod 600 /etc/postfix/sasl_passwd.lmdb

5. Restart postfix service

   a. systemctl restart postfix

d. TLS Configuration for Secure SMTP

   i. nano /etc/postfix/main.cf

   ii. Control+w type tls [enter]

   iii. Locate these sections and configure them to match

      1. smtp_use_tls = yes

      2. smtp_tls_loglevel = 1

   iv. Done

e. This configures authenticated TLS SMTP relay. Continue with the steps below to debug and test the mail is correctly forwarding from Eyeglass to your SMTP mail system.

How to Debug Postfix mail relay Issues

1. Verify the mail is correctly being forwarded with authentication by tailing the mail log.
    a. tail -f /var/log/mail
2. Open the Notification Center Icon in Eyeglass from the main menu options, and send a test email to verify authentication and TLS connections are successful.
3. Example of a successful TLS SMTP message relay from Eyeglass to external SMTP host with TLS 1.2 enabled. the 250 OK code indicates successful delivery. The SMTP error or TLS errors will be visible in this log for debugging.
    a.
```
Jan  9 10:20:05 anycopy postfix/smtpd[2341]: connect from localhost[127.0.0.1]
Jan  9 10:20:05 anycopy postfix/smtpd[2341]: 0FCD52696C: client=localhost[127.0.0.1]
Jan  9 10:20:05 anycopy postfix/cleanup[2344]: 0FCD52696C: message-id=<175090405 6.3361610205605032.JavaMail.sca@anycopy>
Jan  9 10:20:05 anycopy postfix/smtpd[2341]: disconnect from localhost[127.0.0.1] ehlo=1 mail=1 rcpt=1 data=1 quit=1 commands=5
Jan  9 10:20:05 anycopy postfix/qmgr[1340]: 0FCD52696C: from=<demo@superna.net>, size=1781, nrcpt=1 (queue active)
Jan  9 10:20:05 anycopy postfix/smtp[2345]: connect to smtp.gmail.com[2607:f8b0:4023:1404::6d]:587: Network is unreachable
Jan  9 10:20:05 anycopy postfix/smtp[2345]: Untrusted TLS connection established to smtp.gmail.com[142.250.123.109]:587: TLSv1.2 with cipher ECDHE-ECDSA-CHACHA20-POLY1305 (256/256 bi
ts)
Jan  9 10:20:07 anycopy postfix/smtp[2345]: 0FCD52696C: to=<[redacted]@superna.net>, relay=smtp.gmail.com[142.250.123.109]:587, delay=2.6, delays=0.04/0.01/0.63/1.9, dsn=2.0.0, st
atus=sent (250 2.0.0 OK  1610205607 a8sm8914432ilc.20 - gsmtp)
Jan  9 10:20:07 anycopy postfix/qmgr[1340]: 0FCD52696C: removed
```

1. Check for any queued mail that failed
    a. sudo mailq
    b. any queued mail that failed will be listed with a reason code
    c. To reattempt delivery of the mail use the command below
    d. sudo postfix flush
2. To provide support with configuration of your postfix configuration use this command

a. postconf -n

© Superna LLC

2. Eyeglass Ransomware Admin Guide

Home Top

- What's New
- Introduction to this Guide
- Planning and Design
- Everything about Detection, Configuration and Tuning Security Event
- How to Configure, Tune and View Ransomware Defender Threat Detection settings and Responses
- How to respond to Security Events for Warning, Major or Critical Events
- Operational Procedures For Common Tasks
- AirGap 2.0 Guide
- How to Configure a Dell ECS and Data Protection Use Cases

© Superna LLC

2.1. What's New

Ransomware Defender for ECS

1. This will require additional license keys and will allow event data from ECS to be processed by the ECA cluster to offer real time object data protection with alerting and lockout.

2.5.7

1. **Major Release**
2. AirGap 2.0 - A complete solution to protect your data with protractive behavior monitoring of the source data access combined with Smart AirGap technology to manage SyncIQ policy replication to a 3rd AirGap Isilon.
3. **Smart AirGap** is unique solution to Ransomware Defender that suspends copy operations when an active threat to your data is detected. Unlike other solutions that will copy encrypted data to the offline copy.
4. Ransomeware Defender manages the AirGapped Isilon in-band over the replication network ensuring your isolated Isilon is never exposed on your network.
5. Supported Protocols
    a. SMB

b. NFS

c. S3 (The AD user mapped to the S3 keys will be locked out and the S3 IO is not blocked and no snapshot is created on the S3 path where the detection)

6. Automated AirGap Management ensures the AirGap is open and closed automatically before and after SyncIQ block level incremental copies complete. Fastest AirGap solution allows your 3rd copy to be an hour behind production. Not days like other solutions.

7. Virtual AirGap manages the network to ensure your data is offline and not accessible over the network when no data sync's are in progress.

8. New Behavior detections expands behavior analysis combined with honeypot and managed banned list of 2500+ extensions provides the highest level of data protection.

9. Support for Authenticated User SMB Share permissions will now lock on shares that grant access to users using this well known AD group.

10. **Major Feature Updates**

    a. **Learning Mode**. Automates the process of monitoring user behavior and apply settings needed to adjust settings needed. This will manage user behaviors and extension based detections from the banned list of files.

    b. **Monitor mode by user, path or IP address.** Removes the need to whitelist and allows monitor mode applied to a path, IP address or an AD user name. This retains detection, and snapshots without any lockout. This provides new method that will replace whitelisting in most cases.

c. Updated threat detector settings for user behavior detection - new detection vector

d. Banned file list versioning

    i. Multiple file versions allows transitioning to a new file version with latest extensions or roll back to a previous version

e. Banned file hosted in a new location compatible with phone home URL's

    i. Eyeglass deployments that use phone home will now be able to leverage phone home url to retrieve the banned list to simplify firewall and url white listing.

f. Allowed File Extension List Redesigned to File Filter Feature

    i. The Banned file list is now managed get by Eyeglass and not the ECA. This means proxy and phone home will allow retrieving the updated dated file list from the Internet.

    ii. Now all banned files are displayed with a searchable interface. Each file can be enabled, disabled or monitor mode status.

    iii. Ability to add custom file extensions is supported.

    iv. CLI command to convert whitelist entries to new monitor mode settings.

g. **Dual Vector Warning detection** - A new behavioral detection option looks for different behaviors within the Warning severity.

This new option will add one additional pattern of suspicious user activity that is designed to ignore spikes in user detection signals and provides a new analysis vector on user IO behavior to generate warnings.

2.5.5

1. New architecture - allows deployments without HDFS Access Zone requirement for standalone deployments.
2. UI updates allows viewing of:
    a. Flag as false positive user settings.
    b. White listed file extensions from the master known extension list.
3. Archived Events allow action menu to flag as false positive after an event has been archived.
4. Security Guard feature now cleans up files on the igls-honeypot SMB share created during the simulated attack. This clean runs on each scheduled, or on demand execution of the Security Guard feature.

2.5.3

1. **Honey pot File traps**
    - Detection at the folder level, allows files to be placed in specific folder locations, as detection of any type of Ransomware behavior attack that combines file access to Honeypot trap files that Defender uses to track Ransomware at the folder level, and does not depend on a specific file IO pattern for detection.

- Uses immediate lock out logic when this detection trap is tripped.
- Administrators can create this trap on any folder in the file system as needed.

2. **Roaming Profile Support**
    a. Roaming profiles on PowerScale shares writes files using a common Ransomware IO pattern trigger a lockout.
    b. New Relative path whitelist support allows only the directories of the profile to be added to the whitelist, and still protect data in the users profile. Example: whitelist /ifs/data/roamingprofilessharepath/*/Appdata This will ignore all user Appdata (the profile path) in each users home directory on a share that stores all users home directories.

1.9.2 Has new supportability enhancements, a feature to disable real-time critical lockout action, and use only time delayed response for security events. Full feature description in this release is available here.

1.9.3 Offers auto snapshot feature to protect paths and shares when any Ransomware has affected a user workstation. All shares the user has access to have a snapshot applied with a 48 hour expiry. This is enabled or disabled, with default enabled for all detection severities. Requires SnapshotIQ license on the cluster. Full feature description in this release is available here.

1.9.5

- ECA cluster now uses fluentd to collect logs and send to Eyeglass over syslog on port 5514 udp, cluster startup enhanced to debug HDFS configuration and provide validation errors.

- IGLS commands expanded for settings.
- Ability to set snapshot expiry default from 48 hours to another value with IGLS command.
- Ability to set security guard event timer to wait for events that are delayed by PowerScale forwarding rates. IGLS command.
- Ability to set security guard restore permissions timer to ensure restore permissions action has time to complete. IGLS command.
- NFS host lockout supported (enabled with IGLS command, disabled by default). This feature will remove the IP address from the client list(s) and re-save the export definition.

NOTE: DNS to IP resolution to will not be done for client lists that use FQDN. The feature requires client list to use ip address to successfully lockout.

- Default disabled since this can lead to stale mounts for NFS hosts.

NOTE: Admin guide has all IGLS commands for all products.

1.9.6

- Security guard delay IGLS command to delay how long Security guard waits for events to appear from the simulated attack. This solution accommodates variation in CEE forwarding rates from the cluster to the ECA cluster.

2.0

- Mark as false positive on security events allows AI teaching feature of user behaviors with per user behavior learning.

- File extension whitelist feature to allow well known bad file extension to be ignored if used in your organization (IGLS commands).
- Built in role and user for managing Ransomware Defender.

© Superna LLC

2.2. Introduction to this Guide

Introduction to this Guide

Overview

This guide covers the configuration, setup, and monitoring of Ransomware Defender. The solution is deployed with a 3 VM cluster, that processes PowerScale & ECS audit files, with an active active design for maximum availability to survive hardware or software failures.

The active defense solution monitors for user behaviors that are malicious, consistent with Ransomware encryption techniques of customer files. Network-attached SMB mounts on user workstations exposes PowerScale critical data. It can also protect object storage on Dell ECS.

Three levels of detection are possible Warning, Major, and Critical with automated defense options increasing with detection levels.

Abbreviations:

- **ECA:** Ransomware Defender Application - the entire Ransomware defender stack that runs in a separate VM outside of Eyeglass.
- **ECA:** Eyeglass Clustered Agent

Prerequisites, Requirements, and Feature Limitations:

Prerequisite:

Read this first:

It's assumed that all workstations and other entry points for Ransomware are running current virus malware software. Ransomware Defender is a second line of defense product.

ECA Installation Requirement:

The Superna Eyeglass Clustered Agent (ECA) vAPP used by Ransomware Defender therefore a Single host for ECA VM's OR multiple hosts for a high level of availability is required (See the Eyeglass Clustered Agent vAPP Install Guide.)

Licensing Requirement:

Registered PowerScale or Dell ECS clusters licensed for Eyeglass DR qualify to be licensed with Ransomware Defender. Each writable cluster requires an agent license and agent maintenance. The license manager is used to assign licenses.

A system alarm will be issued in the case of insufficient licenses and more writable clusters are detected in the audit event messages.

Note: A cluster can be monitored by the ECA without a license when it's the cold or DR cluster.

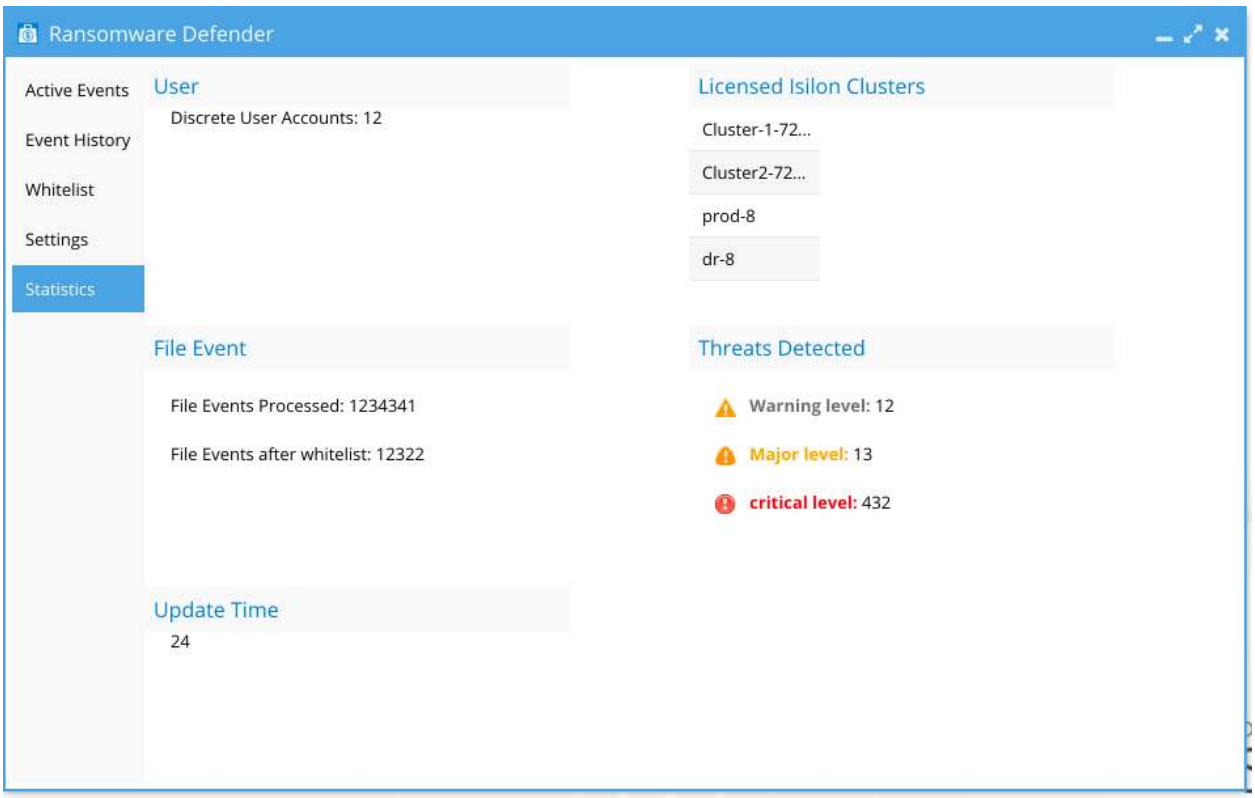

Additional Requirements:

1. Eyeglass VM installed.
2. Cluster discovery licenses (per node or per cluster) that need to be managed by the Eyeglass instance.
3. Ransomware feature license and a Ransomware agent license for all writable clusters protected by Ransomware Defender.
4. CPU limits applied to ECA cluster object in vCenter.
5. Hardware recommendation (see install guide).

Feature limitations:

1. SMB shares created with variable expansion will only support %U for snapshot creation.
2. NFS lockout is supported but disabled by default unless IGLS to enable.

3. NFS lockout requires license lists to use ip address to correctly lockout.
4. Object storage and all buckets

© Superna LLC

2.3. Planning and Design

Home Top

- Overview
- Automated Ransomware Defense Actions
- Securing Root user on PowerScale
- Why file extension filtering feature on SMB shares is a security risk
- External Data Security Expectations for A Secure Environment

Overview

The Ransomware Defender solution for PowerScale requires existing Eyeglass DR cluster licenses for each PowerScale cluster plus, an Eyeglass clustered agent license.

The Eyeglass Ransomware Defender solution is intended to be a last line of defense for critical NAS data stored on PowerScale. A best practice defense should include virus software on laptops and workstations, along with email gateway or IDS network solutions.

For Dell ECS, Ransomware Defender can monitor object access with behavior based detection and offers alerting and lockout of authenticated users.

The intended use case for Ransomware Defender assumes malware has circumvented all existing defenses, leaving critical NAS and object data exposed to attack.

The diagram below shows the traditional approach to security with perimeter defenses, with the primary purpose of ensuring malware never enters your network.

The Eyeglass solution builds a new security perimeter inside your network with active defense to threats.

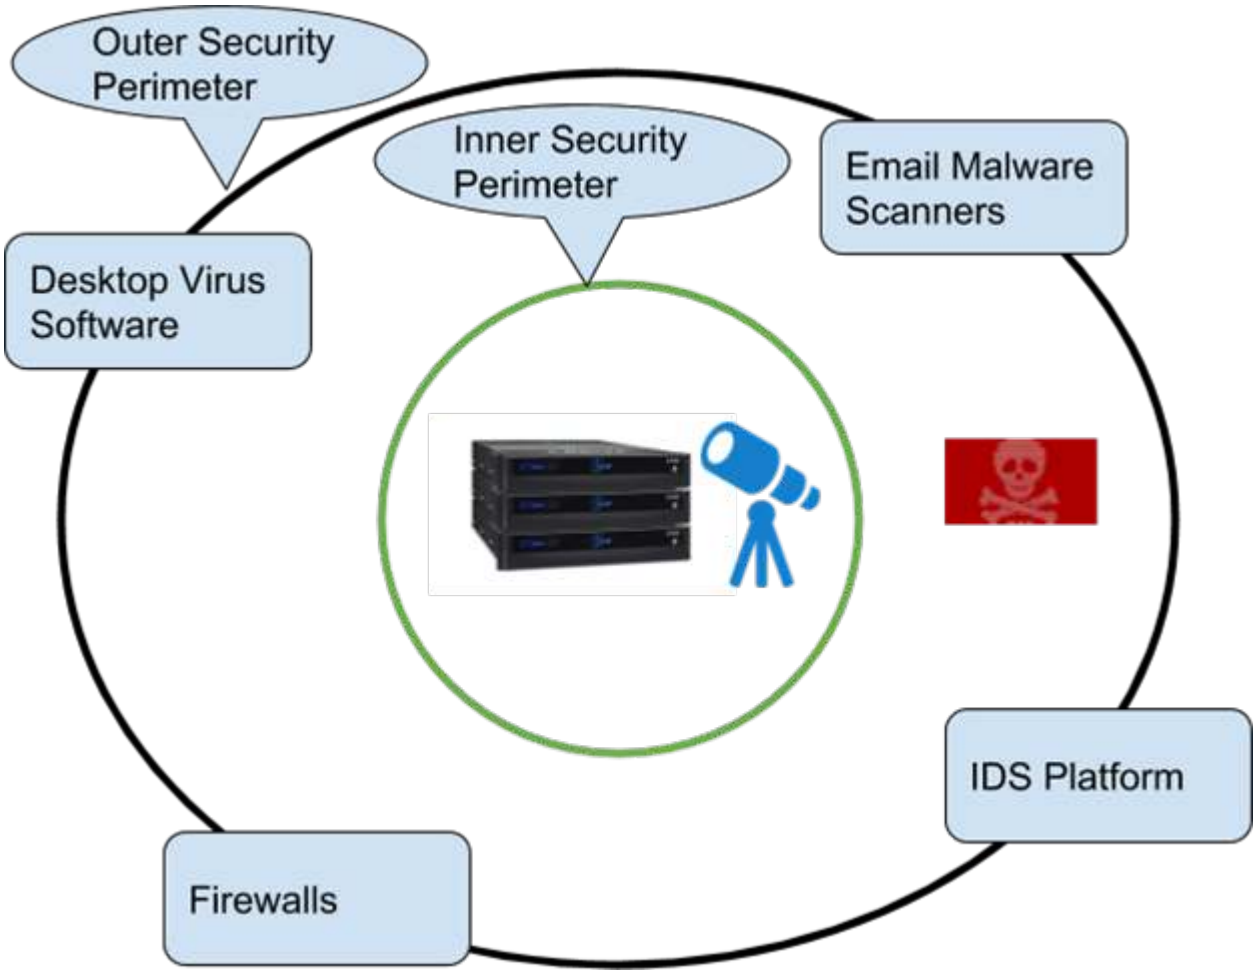

Automated Ransomware Defense Actions

1. **Warning -** Send an alert (email, snmp, syslog), no enforcement taken.

2. **Timed lockout (Major)** - a time-delayed lockout of the user account that triggered the security event. The lockout can be stopped before the timer expires.
3. **Immediate lockout (Critical)** - User lockout begins in real-time once a critical event is detected.
4. **Snapshot** the file system (all severities) using SnapshotIQ on all affected share paths. (Isilon/Powerscale only)

Securing Root user on PowerScale

The root user should never be used to access data on PowerScale. The reason this user is a high-security risk is that root has access to all shares even if access has not been granted to the root user. This security risk could allow a compromised machine, using the root user, to access data and could encrypt all data on the cluster.

Eyeglass Ransomware Defender offers a mode configured through IGLS CLI to disable SMB protocol on the PowerScale clusters managed by Eyeglass. This will ensure if a ransomware event is detected the compromised machine does not destroy all data on all clusters. **See** Eyeglass CLI command for Ransomware **.**

The root user can NOT be locked out with a deny permission, which is why SMB protocol disable is the only way to protect data.

NOTE: IF YOU USE RUN AS ROOT ON SHARES YOU ARE EXPOSING DATA TO VERY HIGH SECURITY RISK SINCE NO LOCKOUT WILL BE POSSIBLE. THIS IS BECAUSE THE USER SID, THAT IS SENT WHEN AN AD USER ACCESSES DATA WITH RUN AS ROOT

ENABLED, IS THE ROOT USER NOT THE ACTUAL AD USER.

We recommend to NOT use run as root on shares for the reason above, AND it fails all security audits of PowerScale in all industry standards (PCI, HIPPA, FedRAMP, ITSG, etc…). Remove run as root option on all shares.

The default setting is to disable the SMB Automated response on a cluster if the root user SID has tripped a threat detector. To enable this mode **VERIFY** no run as root user shares exist.

This can be done using the Eyeglass cluster configuration report

1. Login to Eyeglass
2. Open Reports on Demand Icon
3. Select Create New Report

4. Wait until the report is finished by viewing running jobs

5. Select Open/Print option for the finished report from Reports on Demand, after running jobs shows the report creation is completed.
6. Click Cancel on Print option (if using Chrome).
7. Control-F to search the page option
8. Search for “run_as_root”
9. If any Shares are found with this option set see below.
10. DO NOT ENABLE LOCK ROOT FEATURE.

Why file extension filtering feature on SMB shares is a security risk

1. This Isilon feature changes the audit events created which means Ransomware Defender will not be able to "see" the active ransomware attack and no alert will be sent since isilon suppressed the audit data.
    a. This means you will blind to an active Ransomware attack based on the blocked extension added to the SMB shares.
    b. NOTE Ransomware Defenders primary detection vector is user behavior
    c. In 2.5.7 or later you can customize the banned file list and is the recommended method to add banned file extensions since each extension can be enabled (protection active), disabled (added but no active monitoring) or monitor mode

that will alert but will not lock out users. These 3 states provides more options than the blocking only.

External Data Security Expectations for A Secure Environment

1. Ransomware Defender is a component of an overall security solution that must include the following best practices in order to correctly deploy a security solution.
    a. A data security plan should include multiple layers of security including end point protection and a backup system to recover data. Ransomware Defender is not intended to replace other security solutions or backups of your data.
    b. Backup data should be stored off line so that is is not connected to a network.
    c. Offline copy of data using a cyber vault or airgap.
2. The specification and operational management of this product Requires:
    a. May not detect or prevent any or all malicious code or that use of the licensed program and related updates or upgrades will keep company's network or computer systems free from viruses or other malicious or unwanted content or safe from intrusions or other security breaches

b. Product usage assumes end point protection Anti-virus software is in place on all operating systems, devices, computers.
c. All computers with operating systems are patched regularly and all zero day patches are applied immediately.
d. SPAM filters are implemented for phishing attacks
e. All CVE's are acted upon with patches and remediation applied
f. All firewalls, security devices are running current versions and configured correctly to protect networks
g. The compute infrastructure is maintained and provides minimum product requirements for cpu, memory, disk latency.
h. The end users and IT are trained to respond to a Ransomware attack and have a run book to respond to an incident.
i. End users are trained regularly for phishing attacks and social attacks intended to compromise computers with Malware/Ransomware
j. All product alerts are acting on in a timely manner in the infrastructure and patched weekly for all critical security patches.
k. Security guard product feature is monitored daily for proper product functionality
l. Honeypot feature is implemented fully on all SMB shares and NFS exports

© Superna LLC

2.3.1. How to determine threat response settings to meet your Company's Risk Profile

Home Top

- Ransomware Defender Protection Modes
- How to determine threat response settings to meet your Company's Risk Profile
- Threat Response Settings
  - Automated Threat Responses Settings
  - Recommended Threat Response Setting for Low Risk tolerance
  - Recommended Threat Response Settings for Medium Risk tolerance
  - Recommended Threat Response Settings for Medium-High Risk tolerance

Ransomware Defender Protection Modes

The 3 different modes all protect the file system and each mode determines how you should respond to alarms. The target events per day should be < 1 per day and see the process below to achieve this with learning mode.

Production Modes that protect the File System

Mode	File System Protection for production	User Lockout	File, User and IP Tracking	Snapshots	Operational steps
Learning	✓	No	Yes	Yes	React to alarms, lockout user from GUI
Monitor	✓	No	Yes	Yes	React to alarms, lockout user from GUI
Enforcement	✓	Yes	Yes	Yes	React to alarms, unlock users from GUI

1. Target for learning mode is < 1 user event per day.
2. To reduce events further increase minor, major , critical threshold by adding 25 to each value, wait a day and repeat until no events received per day.

superna®

Manage, Protect, and Secure Unstructured Data at Scale

How to determine threat response settings to meet your Company's Risk Profile

The Ransomware Defender product has several options to tune the detection and response to a Ransomware attack. The more sensitive the detection the more likely a false positive can occur. Threat response options are outlined below with business impact considerations for each option. This section should be reviewed to determine how to configure the product in your environment.

Risk tolerance and business impact need to be assessed to determine the best settings for your environment. The section below outlines the recommendations for each threat detection level.

Threat Level Severity	Action	Snapshot Data Protection and Recovery	Business Impact

		Enabled (all Shares a user can access has snapshot applied)	
Warning	No action taken. Email alert is sent	X	No impact on applications or user access to data. Snapshot is applied to protect the file system.
Major	Timed lockout of user. Email alert is sent	X	Business applications or servers write data that are not added to the ignore list, can be locked out. **Impact:** application downtime until restore access completed. **Recommendation**: add to ignore list.
Critical	Immediate lockout of user. Email alert is sent	X	**Impact:** application downtime until restore access completed. No wait time from detection to lockout for administrators to determine action. **Recommendation**: add to ignore list or Disable critical actions.

Threat Response Settings

Automated Threat Responses Settings

1. Critical Severity - Lockout of user account - is immediate
2. Major Severity - A delayed lockout Grace Period is set (user account lockout delayed by X minutes)
3. Auto Snapshot of the file system at share path - on detection of ANY severity

Recommended Threat Response Setting for Low Risk tolerance

Monitor Only Mode enabled - Email Alerts

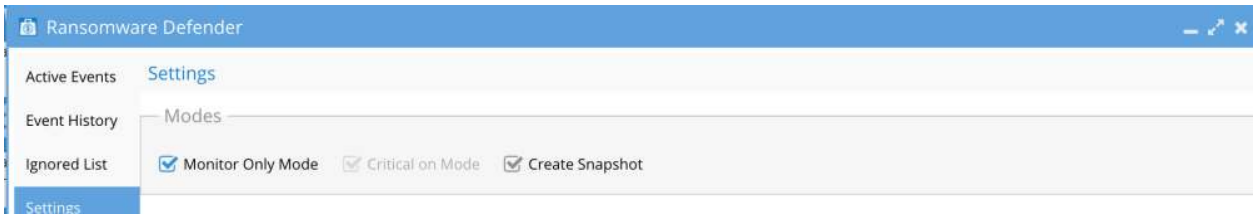

Recommended Threat Response Settings for Medium Risk tolerance

NOTE: In this configuration files can be encrypted up to the Grace Period value, but a snapshot has protected the file system at the point of detection allowing for accelerated recovery of files. The security event lists all affected files to build a recovery list of files.

1. "Critical on Mode" uncheck to disable immediate lockouts
2. Set Major delayed lockout timer (Grace Period) to a value that allows an administrator to reach and determine if lockout should occur (In the Screenshot below the "Grace Period" is set to 60 Minutes)
3. "Create Snapshot" Mode enabled

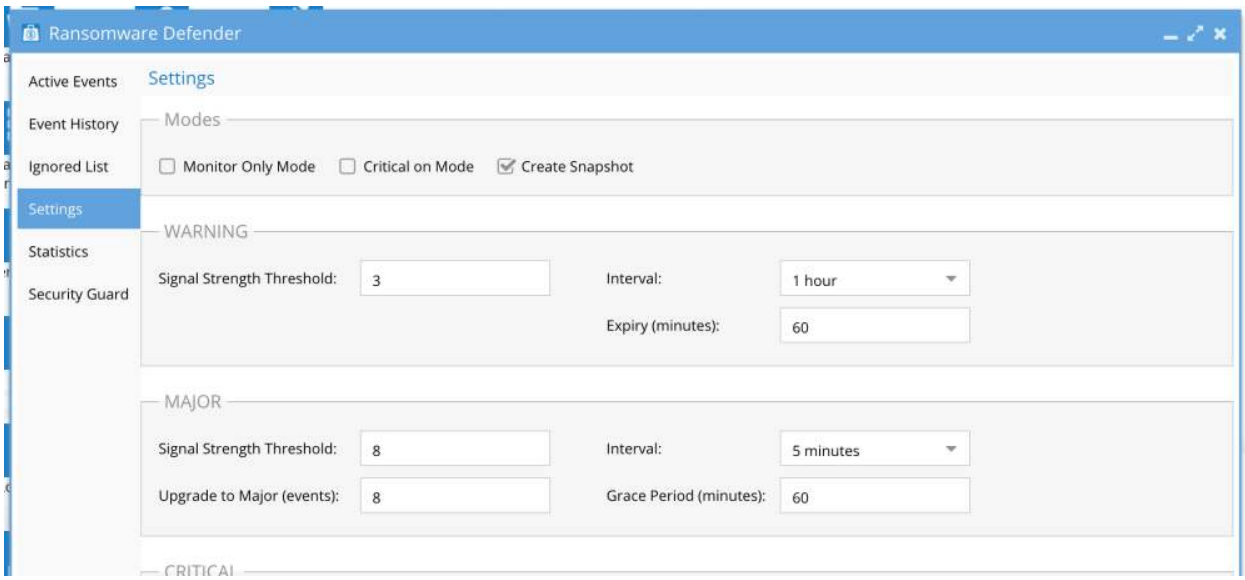

Recommended Threat Response Settings for Medium-High Risk tolerance

NOTE: In this configuration files users are locked out immediately, the risk of false-positive with a lockout is higher.

1. “Critical on Mode” checked to enable immediate lockouts.
2. Set Major delayed lockout timer “Grace Period” to a value that allows an administrator to reach and determine if lockout should occur. (In the Screenshot below the “Grace Period” is set to 60 Minutes)
3. “Create Snapshot” mode enabled.

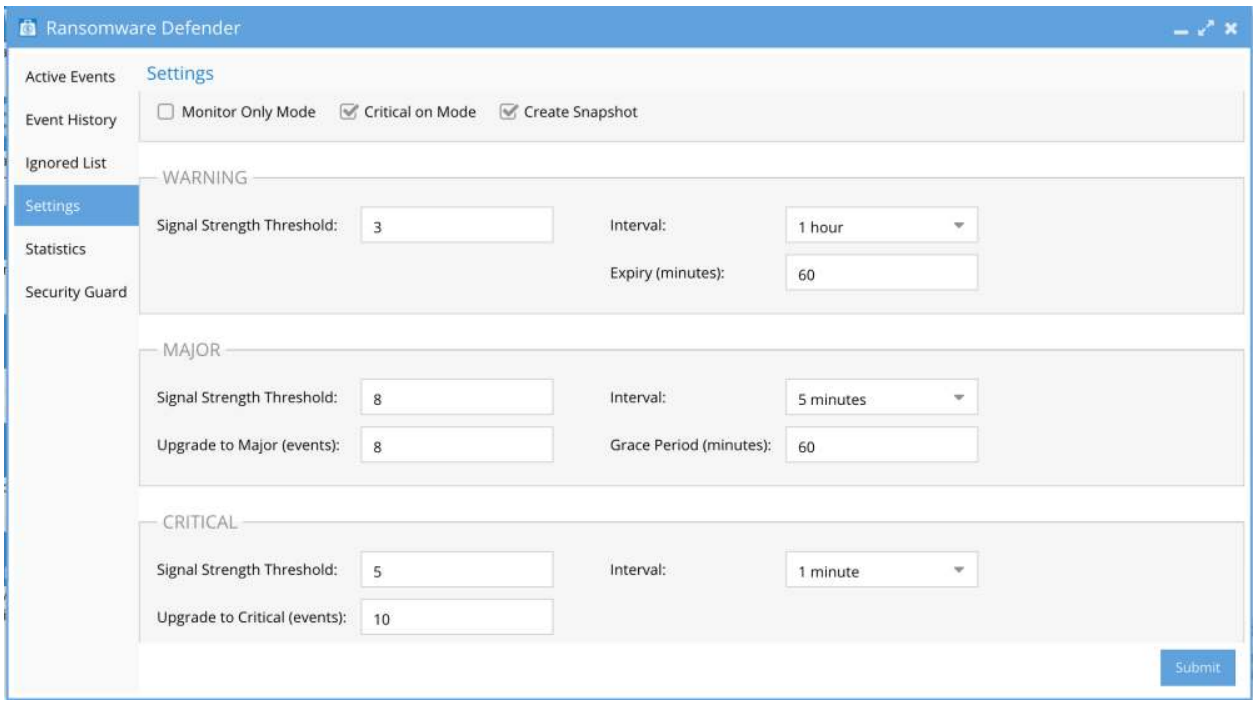
Ransomware Defender
Active Events
Event History
Ignored List
Settings
Statistics
Security Guard
Settings
Monitor Only Mode
Critical on Mode
Create Snapshot
WARNING
Signal Strength Threshold:
3
Interval:
1 hour
Expiry (minutes):
60
MAJOR
Signal Strength Threshold:
8
Interval:
5 minutes
Upgrade to Major (events):
8
Grace Period (minutes):
60
CRITICAL
Signal Strength Threshold:
5
Interval:
1 minute
Upgrade to Critical (events):
10
Submit

© Superna LLC

2.3.2. Eyeglass User Lockout Active Directory Planning

Home Top

- Eyeglass User Lockout Active Directory Planning
- Scenario #1
- Scenario #2

Eyeglass User Lockout Active Directory Planning

The lockout process identifies all shares the user has access permissions, based on searching all shares, in all Access Zones, on all clusters managed by Eyeglass. This list of shares will have a real-time deny permission added to the share for the affected user.

A special case is handled for the “Everyone” well-known group, which how it operates in multi-domain Active Directory configurations should be understood.

Two scenarios can exist with AD domains on PowerScale clusters.

Scenario #1

- The first is parent and child AD domains that are members of the same forest, and a trust relationship exists.

Scenario #2

- The second scenario covers two domains that are not members of the same forest, and no trust relationship exists between the domains

The “Everyone” well-known group, if applied to a share in each scenario, is shown below, and lockout permission is applied regardless of which domain the user is located. This is required since Eyeglass has no way to know if the domains trust each other. This solution ensures all “Everyone” shares are locked out, which is more secure than skipping some shares.

Reference the diagram below.

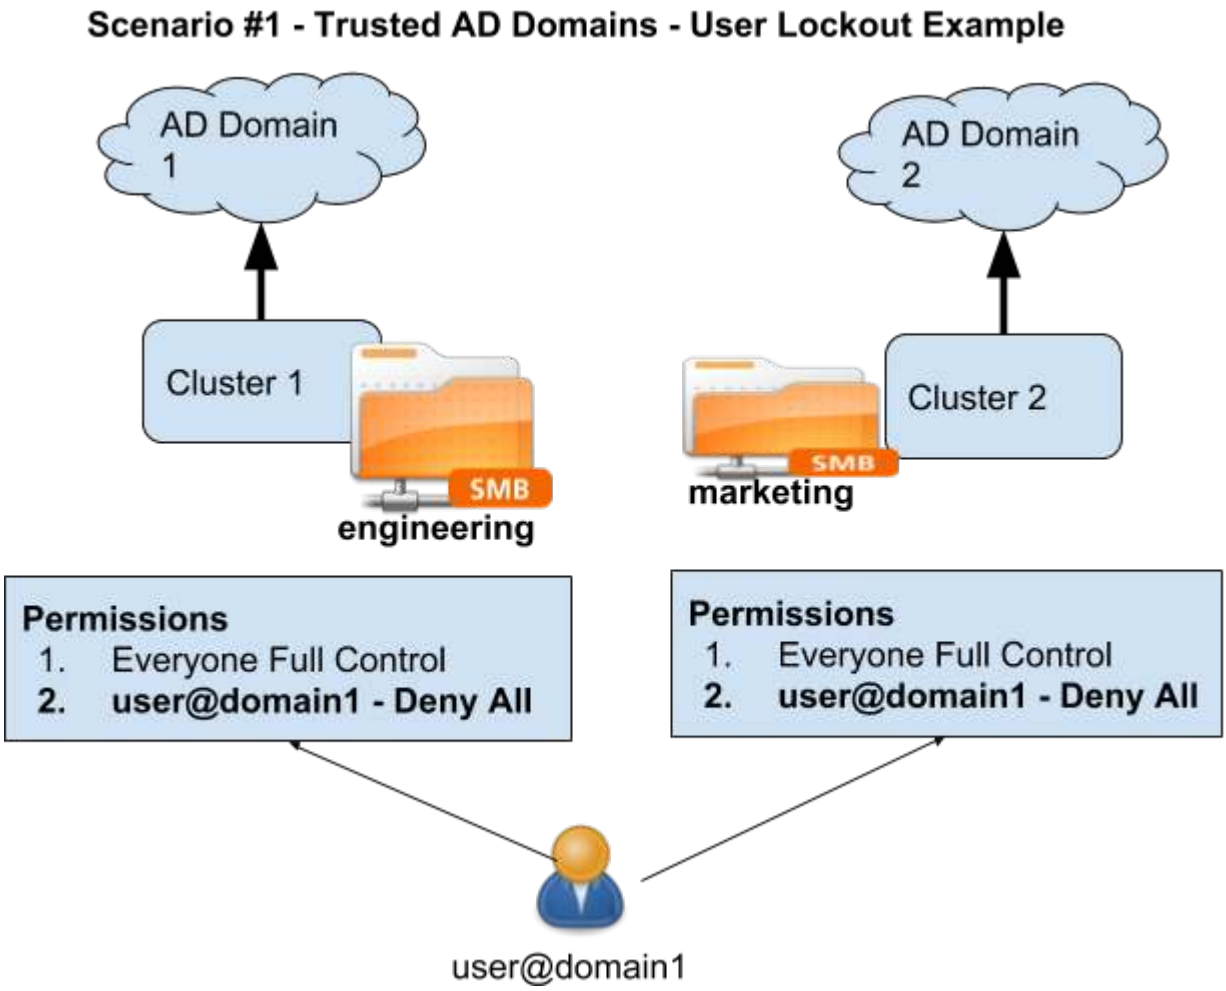

Scenario #2 - Untrusted AD Domains - User Lockout Example

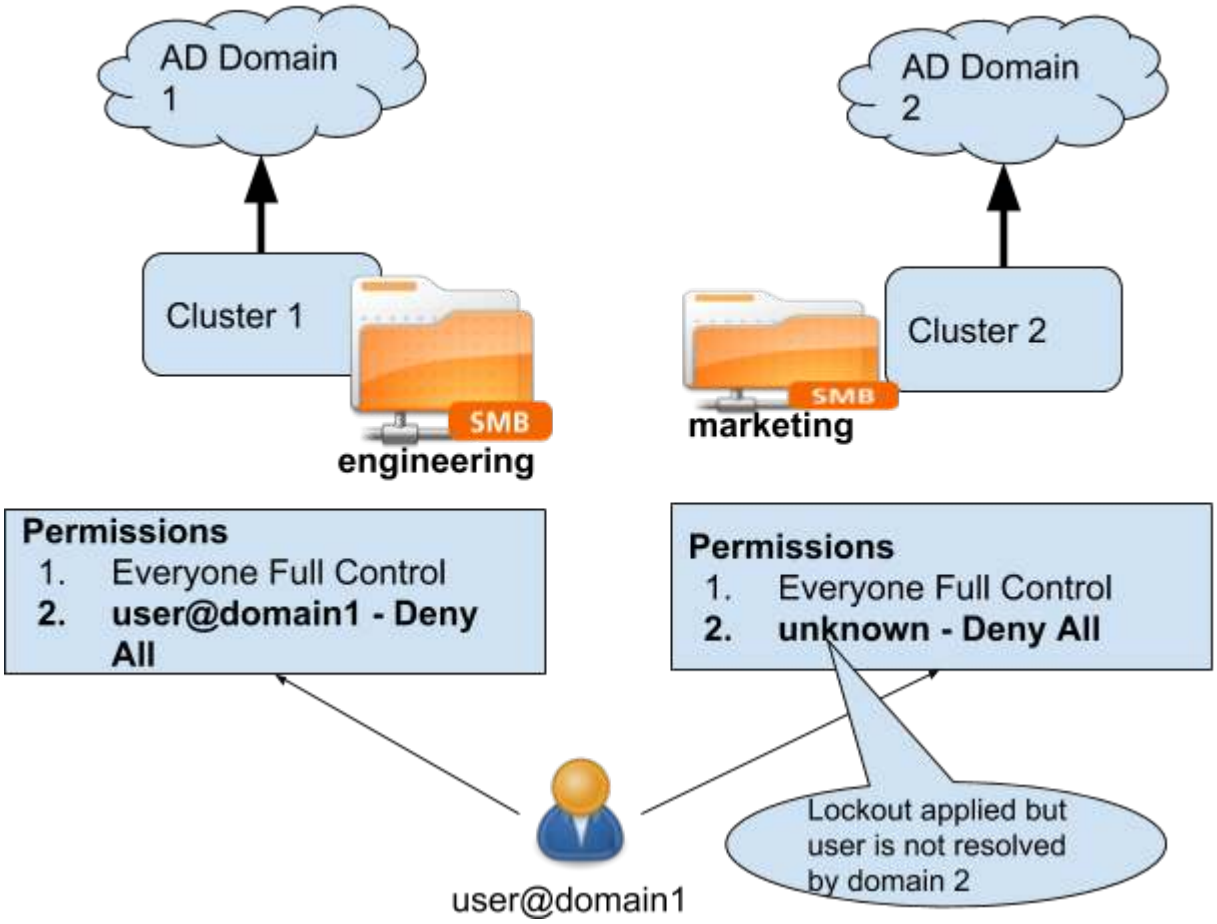

© Superna LLC

2.3.3. Ransomware Audit Events Required for all Deployments

Home Top

- Overview
- Audit Success:
- How to check your settings

Overview

This section covers the required audit events for a supported Ransomware Defender deployment, this applies to all OneFS releases.

Audit Success:

1. close | create | delete | read | rename| write

How to check your settings

1. SSH to your Isilon cluster
2. Run the following command: isi audit settings global view
3. Run the following command: isi audit setting view --zone=<name of access zone>
4. Repeat for each access zone

© Superna LLC

2.3.4. Well Known Ransomware File Extension Whitelist

Home Top

- Well known Ransomware File Extension List 2.5.7 >
  - Versioned Banned files List
  - How to switch banned file url to latest, default or a newer file
- Well Known Ransomware File Extension List < 2.5.6

Well known Ransomware File Extension List 2.5.7 >

New in 2.5.7 or later releases is versioned Banned list. This new feature allows switching between current banned list and new file lists that are published or selecting latest option to always use the latest list. If the new list adds extensions that are in use in your environment a lockout could occur. This new feature allows controlled switching to a new file and the ability to see what new extensions have been added to the list in the new file.

Phone home integration now moves the dynamic sync of the banned files from the ECA VM's to the Eyeglass VM and supports the same phone home URL's that have been whitelisted. This simplifies access to the file without needing to add any new firewall or proxy rules to use this new feature.

Versioned Banned files List

This feature allows a version of the banned file list to be selected, auto selected, differenced to allow migration from one version to another using controlled commands below. The files will appear in the File Filter tab in the Ransomware Defender GUI.

1. **Requirements**
    a. Requires 2.5.7 update 1
2. This section will provide a list of versioned files with a link to a file containing all the new extensions add to the version of the file.
3. File Versions commands
    a. igls rsw filefiltersettings (list current version and settings)
    b. igls rsw filefiltersettings --version=<version> (select a version to be active)
    c. igls rsw filefiltersettings --diff=<version1,version2> (show the changes between one version and another to know what new file extensions have been added)
    d. igls rsw filefiltersettings set --mode=Latest (always pickup and use the latest version available online)
    e. igls rsw filefiltersettings set --mode=Fixed --version=<version> (select a fixed version as a static version of the banned list to use)

How to switch banned file url to latest, default or a newer file

1. **NOTE: 2.5.7 will default to the same file used in 2.5.6 release but stored in the new location url cloudapps.supernaeyeglass.com**

2. Login to the eyeglass vm as admin
3. sudo -s (enter admin password)
4. cd /opt/superna/sca/data
5. nano system.xml
6. locate the tag called <rsware>
7. paste one of the following tags under the tag above
    a. default file used from 2.5.6 Release
        i. <rsw_threat_file_url>https://cloudapps.supernaeyeglass.com/supernaRansomwareFilters.json</rsw_threat_file_url>
    b. Latest Version URL will pick up the latest version. This could introduce new extensions that trigger a lockout
        i. <rsw_threat_file_url>https://cloudapps.supernaeyeglass.com/latest/supernaRansomwareFilters.json</rsw_threat_file_url>
    c. Versioned file allows standardizing on a specific file version to control which file is used on your appliance.
        i. <rsw_threat_file_url>https://cloudapps.supernaeyeglass.com/<YYYYMMDD>/supernaRansomwareFilters.json</rsw_threat_file_url>
8. Save the file with control+x and answer yes to save the file.
9. Restart the SCA for the changes to take effect and download the new file to be cached on the appliance for processing. This file is then automatically pushed to the ECA nodes to use the version of the file selected.

10. systemctl restart sca

11. Login to Ransomware Defender and search for files on the versioned file list of new extensions to verify the show up in the GUI. If they do not show up, it means Eyeglass does not have access to reach the Internet URL.

12.

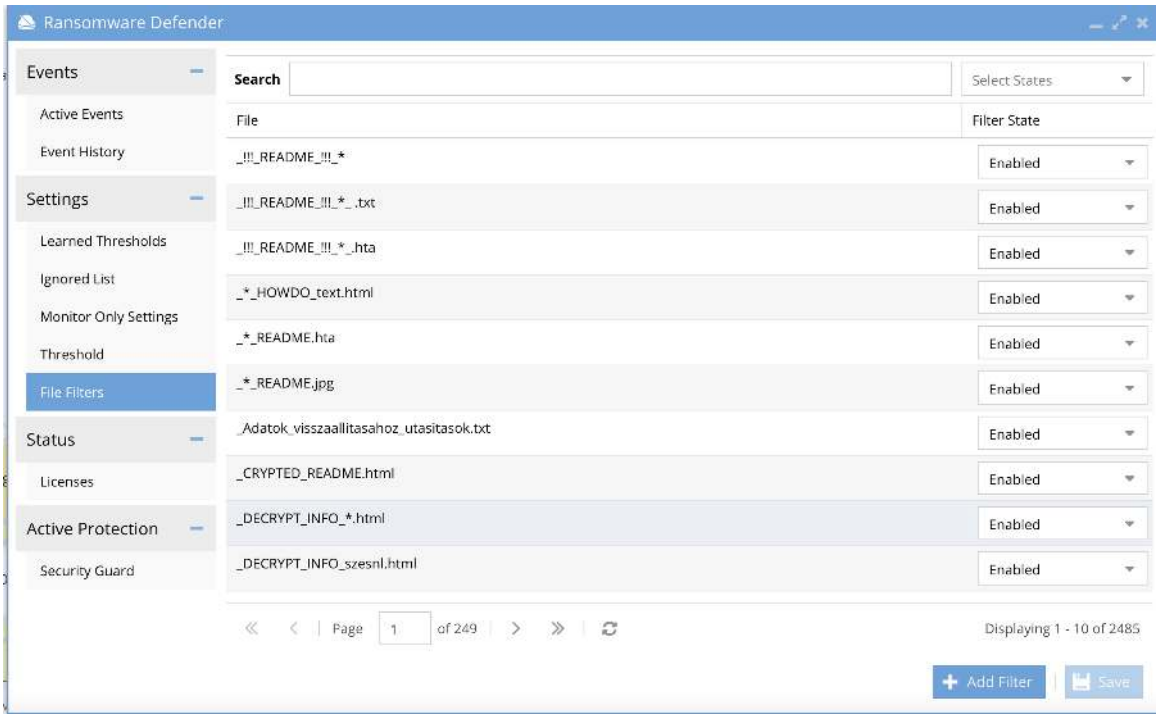

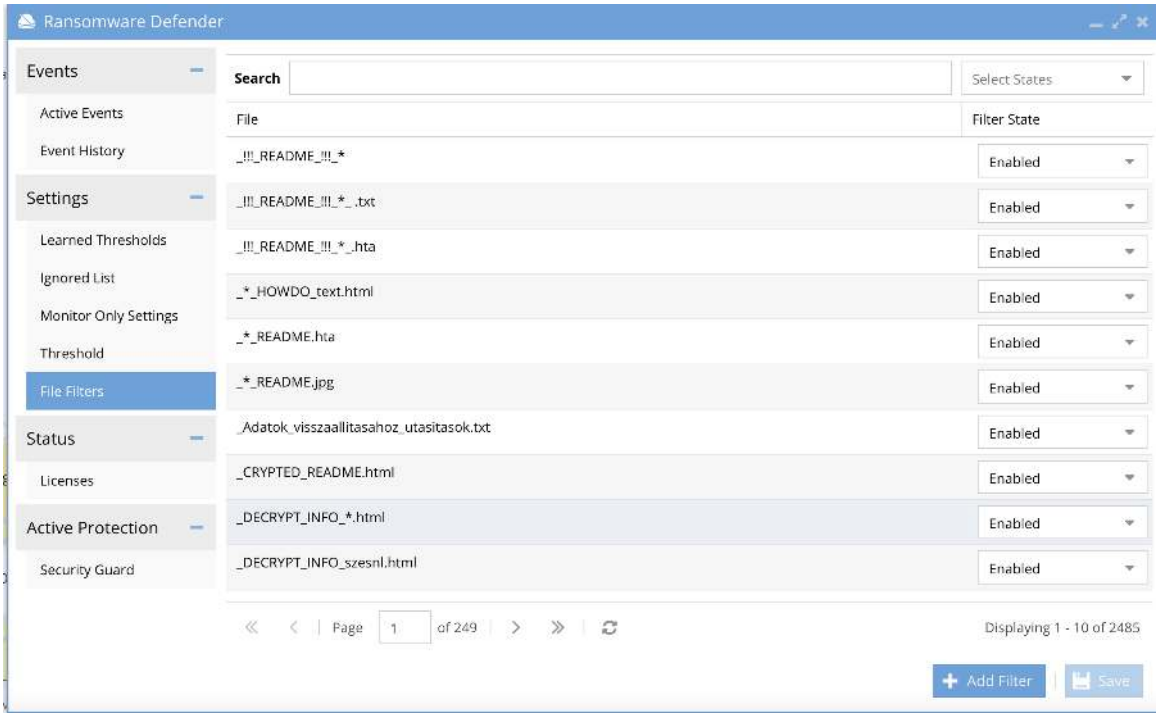

13. done

Well Known Ransomware File Extension List < 2.5.6

Ransomware Defender maintains a dynamic list of well known bad file extensions that are suspicious. This list is over 1000 extensions. It is common for some applications or enterprises to use a file extension on this list. This feature allows whitelisting the extension in use that will trigger security detections. **NOTE: The file listed below is deprecated as of release 2.5.6.**

The whitelist is maintained with igls commands in the admin guide igls section. The command allows adding, listing and removing extensions from the list.

NOTE: The exact extension syntax to use must match this file exactly as found in this document. You can search this document with a browser https://storage.googleapis.com/rwdefender.superna.net/supernaRansomwareFilters.json

Example below. See the guide for all commands **note single quotes.**

igls rsw allowedfiles add --extensions='*.ext1'

© Superna LLC

2.3.5. Security Event Descriptions

Home Top

- Detected User Security Event Descriptions
- Security Event State Descriptions
- Security Event Possible Action Descriptions

Detected User Security Event Descriptions

Once a user security event appears in Active Events, the following table outlines the column definitions and descriptions of each state of the security event.

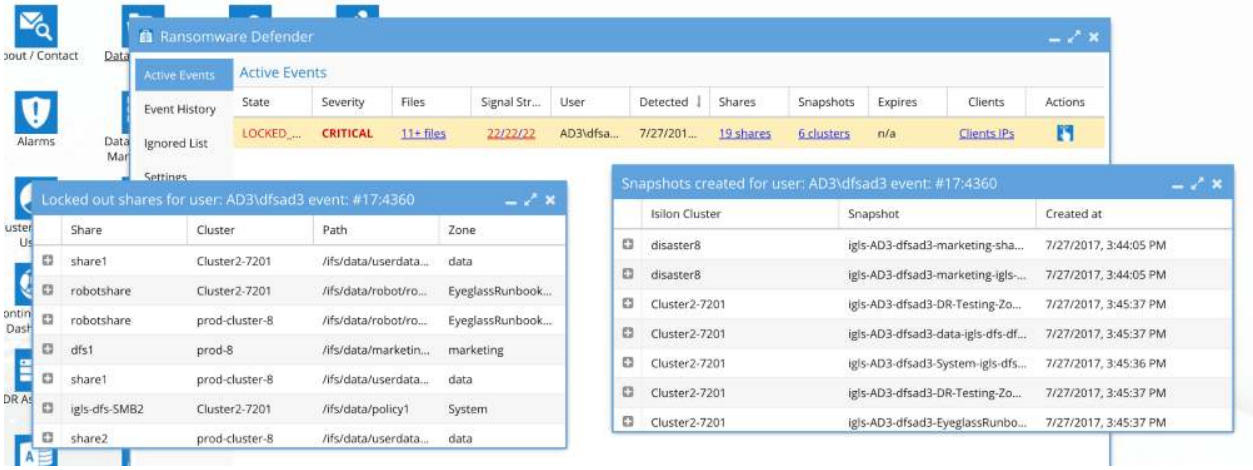

Column Name	Description
State	**Warning** - Threat rate threshold crossed. **Delayed Lockout** - Major Threat rate threshold crossed.

	Locked Out - Critical Threat rate threshold crossed.
Severity	**Warning** - Threat detector peak rate threshold for this event was crossed. **Major** - Threat detector peak rate threshold for this event was crossed. **Critical** - Threat detector peak rate threshold for this event was crossed.

Files	A count of files that tripped the threat detectors for this event. Click to browse the file system path to see the location on the disk that the user was accessing.
	
	• Two tabs are shown:
	  ○ One is a list of files that the user was accessing within the last hour since the event was detected (All Files).
	  ○ Affected Files is a list of files that tripped the threat detectors.
	• All files should be inspected to verify integrity

Signal Strengths	**Each number from left to right is warning peak/ major peak / critical peak threat rate file count**. This indicates the highest count seen for each severity configured in the settings tab. The metric is a count per minute. The higher the number for each severity indicates a higher security risk detected for the user behavior. It indicates more files were involved in the threat detection security event. When comparing two different security events, higher numbers indicates more files tripped the threat detector.
User	The domain and user account of the affected user or NFS ip and UI will be shown
Detected	Date and time representing the beginning of the security event. This event will stay until it is auto-archived or is updated as resolved, or unresolved

	status.
Protected Shares	Lists the cluster, share name, and access one of a share that had a lockout applied. Expanding will display the deny permission and existing ACL applied to the share.
Snapshots	Lists the snapshot name, time, and path that was protected by data-protection and recovery snapshot.
Expires	This will show the time remaining before auto-archive as unresolved is applied to the event. The auto-archive feature will only apply to events detected as warning, and will monitor the event for this time period before archiving the event as unresolved. OR If a timed lockout is active, the time remaining until a lockout will occur.
Clients	This has a pop-up link to list the source ip address of the client machine the user was logged into when the signal event was detected. This assists in finding the client on routers and switches in the environment. **Multiple ip's can be listed for a client if they are logged into more than one machine.**
Actions	Click to bring up the security event history of the event, all previous actions taken and menu to select available actions depending on the state of the security event.

Security Event State Descriptions

A Ransomware event in Eyeglass can be in one of the following states:

State	Description

WARNING	New Ransomware events with a WARNING severity initially have a WARNING state.
DELAYED_LOCKOUT	New Ransomware events with a MAJOR severity initially have a DELAYED_LOCKOUT state. This implies that the user has not yet been locked out, but will be if the event is not acknowledged.
LOCKED_OUT	New Ransomware events with a CRITICAL severity initially have a LOCKED_OUT state. MAJOR severity events that are not acknowledged before the grace period elapses also have a LOCKED_OUT state. WARNING severity events have a LOCKED_OUT state if the Administrator explicitly locks out the user.
ACKNOWLEDGED	A WARNING severity event can be acknowledged to indicate that the admin has seen the event and is monitoring the situation. MAJOR severity events change to ACKNOWLEDGE when the admin intervenes before the grace period has elapsed. CRITICAL severity events can never be ACKNOWLEDGED.
ACCESS_RESTORED	An event is in RESOLVED state when the Administrator has restored access to a locked-out user.
SELF_RECOVERY	An event is in SELF_RECOVERY state when the Administrator has initiated a workflow for the user to recover the affected files. See the Data Recovery section in this guide.
RECOVERED	An event is in RECOVERED state when the user file recovery process is complete.

	RECOVERED state events are not listed in the Active Events tab on Eyeglass. They are listed in the Event History tab.
UNRESOLVED	An event is in UNRESOLVED state when the Administrator has archived the event, but not explicitly restored access to the user. UNRECOVERED state events will are not listed in the Active Events tab on Eyeglass. They are listed in the Event History tab.
ERROR	An event is in ERROR state when Eyeglass has attempted to initiate an action on the Administrator's behalf, but that action has failed.

Security Event Possible Action Descriptions

The following actions are available to the Administrator at different stages of the Ransomware event lifecycle. The *Required States* column lists the state that the event must be in, for the action to be available. Whenever an action is submitted, a new record is added to the event's history.

Action	Required States	Result
Comment	ANY	Adds a comment to the event history.
Acknowledge	WARNING	Changes the event to ACKNOWLEDGED state.
Stop Lockout Timer	DELAYED_LOCKOUT	Changes the event to ACKNOWLEDGED state. Disables any countdown for the grace period on MAJOR severity events.

Lockout	WARNING, DELAYED_LOCKOUT	Initiates the procedure on Eyeglass to revoke access to the user's shares. Changes the event to the LOCKED_OUT state.
Restore User Access	LOCKED_OUT	Initiates the procedure on Eyeglass to restore access to any shares where access was revoked in the lockout step. Changes the event to ACCESS_RESTORED state.
Initiate Self Recovery	ACKNOWLEDGED, ACCESS_RESTORED	Launches the Eyeglass workflow to allow the user to recover all files associated with this event. This procedure will put the event into the RECOVERED state when it is complete. Events in the RECOVERED state. See the Data Recovery section in this guide.
Mark as recovered	ACKNOWLEDGED, ACCESS_RESTORED, SELF_RECOVERY	Allows the admin to manually mark an event as having been recovered. This can happen if the administrator manually restores files, or the user decides that they do not need the encrypted files.
Archive as Unresolved	WARNING, ACKNOWLEDGED,	The administrator can archive an event in nearly any state. The event gets put into event history and is no longer shown on the

	LOCKED_OUT, ACCESS_RESTORED, SELF_RECOVERY, ERROR	active events screen.
Create Snapshot	Manually apply a snapshot to shares in the security event	Run this action if the auto snapshot was disabled. It allows manual apply of snapshots to shares.
Delete Snapshot	Manually delete snapshots applied to share path security events.	Run this action if snapshots were applied and you want to manually delete BEFORE the auto expiry set on the snapshot.

© Superna LLC

2.3.6. NFS Lockout Feature

NFS Lockout Feature

This is now supported and enabled with IGLS command. The default is disabled. Once enabled, NFS source IP in the audit message is used to find exports that list this on a client list. The IP is removed from the export and re-saved. This will lockout the NFS host mount.

NOTE: This can cause stale mount issue on the hosts.

See Eyeglass Ransomware CLI section in this guide for configuration

© Superna LLC

2.3.7. Planning New application Workloads Best Practice

Home Top

- Overview
- When to use this process?
- Requirements
- Best Practice

Overview

When adding new applications that will mount cluster storage or new work flows, it is a good idea to monitor the IO of these new workloads using Monitored path, user or IP to allow learning mode to monitor the behavior of this new workload. This allows learning mode (Release 2.5.7 or later) to monitor this work load and customize settings for this workload.

When to use this process?

If you are currently in enforcement mode with user lockout mode enabled, then you can follow this process when adding a new application or workload to your environment.

Requirements

1. Release 2.5.7 or greater

Best Practice

1. Set up Monitoring
    a. This will monitor , detect, raise alerts and create snapshots to protect the workload but it will not lockout the application. Guide here.
        i. If your application uses a service account to access the storage with an AD user, configure a monitor only setting under the settings menu of the Ransomware Defender icon. Enter the user domain\username and save the configuration.
        ii. If the application uses NFS, then add a host ip entry on the monitor only settings and save the configuration.
2. Enable Learning mode
    a. This mode requires monitor mode enabled and then learning mode. See the guide here.
3. Wait several days with the workload running or until no warnings appear in the active events window of Ransomware Defender. Review the Flag as False positive tab and locate the service account or IP host from above steps.
    a. If no entry is found then the workload behavior is not going to be locked out. You can chose to leave the monitor only setting for this user or ip host OR you can remove the

monitor only setting and disable learning mode and monitor mode to return to enforcement mode.

b. If there is an entry for this user or ip host, then monitor for another day to verify the multiplier value stays the same. This means that the settings are correct to avoid any detections and lockouts for this workload.

© Superna LLC

2.4. Everything about Detection, Configuration and Tuning Security Event

Home Top

Everything you need to learn about security events, detections, configuration and understanding Security events in Ransomware Defender.

- Ransomware - Threat Detection Settings Summary Explanation
- Tuning Ransomware Defender Detections for Your Environment
- Ransomware - Security Event Explanations
- Best Practice for Tuning Ransomware Defender

© Superna LLC

2.4.1. Ransomware - Threat Detection Settings Summary Explanation

Home Top

- Ransomware - Threat Detection Summary Explanation
- Dual Vector Warning Detection in 2.5.7 or later
- Ransomware Security Signal Events and Detection Overview
  - Eyeglass Active Responses to Threats

Ransomware - Threat Detection Summary Explanation

Ransomware Defender uses threat detectors to create detection signals. Four detection vectors execute in real-time 1) User behavior detection 2) Honey pot file (trip wire) detection 3) banned file extension and 4) If signals are being tripped for multiple users, Ransomware Defender will escalate the severity of all security events more quickly than just considering the individual users in isolation.

A security event will be raised when the criteria for Warning, Major or Critical have been met. The criteria used to evaluate the severity of the security event requires:

1. The cumulative total number of signals for a single user has crossed the **signal strength threshold values** defined for Warning or Major or Critical within a sliding window evaluated over the time Intervals configured for Warning, Major or Critical. See image below for recommended settings. **(NOTE: These settings should not be changed without consulting support)**

2. NOTE: Each new signal received will trigger a calculation to determine if it is a Warning, Major or Critical. A security event will match only the highest severity after all three calculations are completed.
3. NOTE: A security event can be upgraded to the next highest severity based on the cumulative signal count crossing the next highest threshold. The corresponding user action will apply Warning alert only, Major timed lockout, Critical immediate lockout, snapshots are used for all Severities.

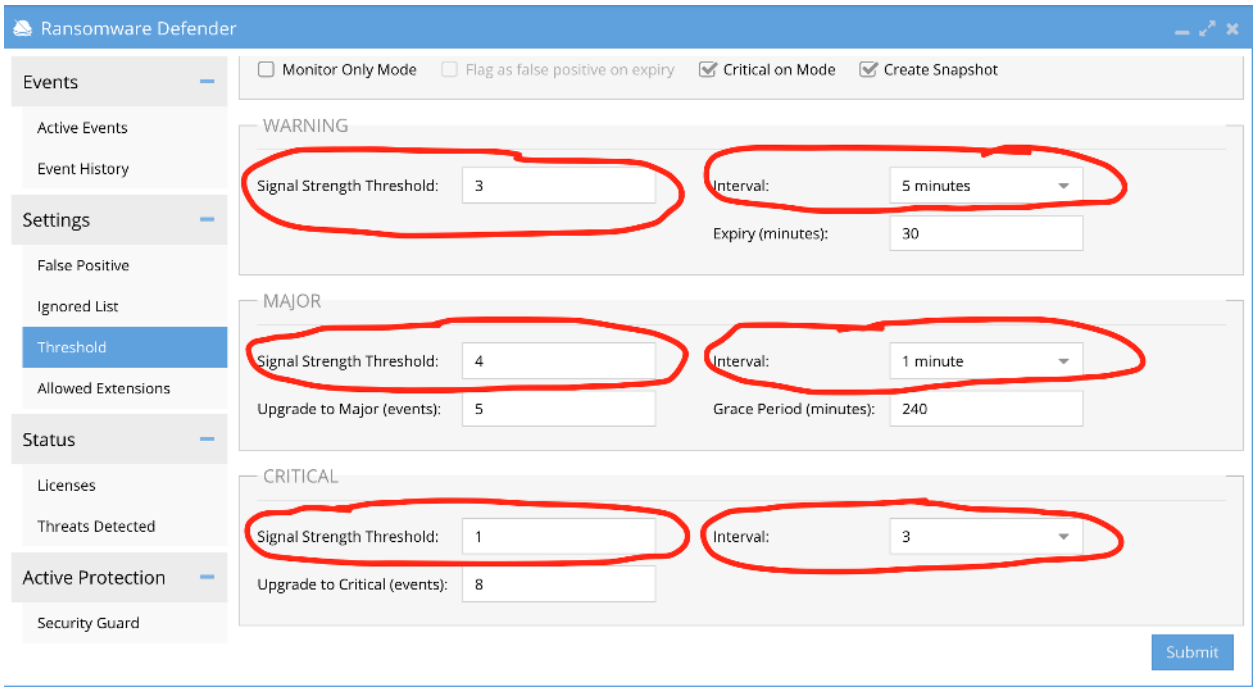

Dual Vector Warning Detection in 2.5.7 or later

A new behavioral detection option looks for different behaviors within the Warning severity. This new option will add one additional pattern

of suspicious user activity that is designed to ignore spikes in user detection signals and provides a new analysis vector on user IO behavior to generate warnings. This allows analysis of signals against a single and dual vector detection function. One vector may not trip a warning but the dual vector logic can detect and raise a warning.

This feature also allows customization to add N dual vector detection settings by clicking the add button to add a new dual vector setting. The product defaults to a single and dual vector setting. **NOTE: Warnings trigger proactive snapshots on all shares accessible to the user.**

In the screenshot example below

1. Single Vector Warning is default from < 2.5.7 release setting. 80 signals or more in a 5 minute time window will trigger a warning but it does not matter when the 80 signals appear in the 5 minute window. This is a single vector detection.
2. Dual Vector Warning is 30 signals in a 30 minute window but the 15 signals must persist over a time period > than 10 minutes from the first signal timestamp to the last signal timestamp. This 2nd vector will only raise a warning if both conditions are true. This second vector operates within the first window, in this example 30 minutes.
    a. Additional Dual vector triggers can be added

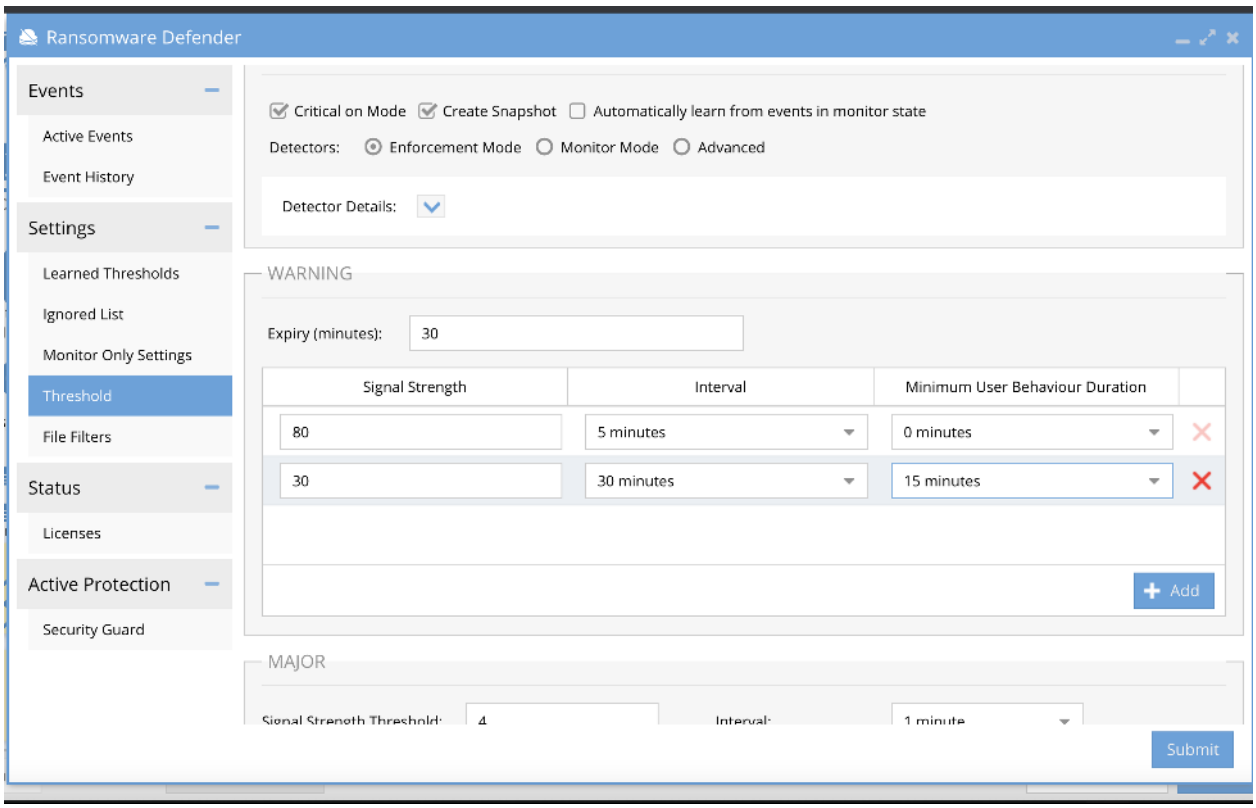

Ransomware Security Signal Events and Detection Overview

Ransomware Defender is a per-user monitoring solution that operates at PowerScale and Dell ECS Scale. This means each user's file activity is monitored individually for user behaviors that trigger threat detection patterns. This builds a zero day solution to identify patterns of IO, that are detected and weighted, without needing definition file-based detection.

The weight is called "**signal strength**" and determines how Eyeglass will respond to the threat.

Three threat levels are defined:

1. **Warning** - No action taken only alarm email sent to the administrator

2. **Major** - Timed lock of the user account in minutes from the event
3. **Critical** - Immediate lockout of the user account

Eyeglass Active Responses to Threats

1. Lockout action means deny permission on all shares the user has access across all managed PowerScale clusters (not just the cluster where the event was detected)
2. (Isilon/Powerscale only) Create snapshots if suspicious events are seen, and snapshot all shares a user has access. This feature is enabled or disabled and applies to Warning, Major, and Critical event types.
    a. It protects share paths with uniquely named snapshot, one per share detected for the user and defaults to 48 hour expiry
    b. In a multi-user infection scenario, this can protect the 2nd, 3rd, etc.. user's data on group shares that were snapshoted by user one infection. This offers maximum data protection.
    c. IGLS command is available to change the default expiry on snapshots. See Eyeglass Ransomware defender CLI in this guide.

© Superna LLC

2.4.2. Tuning Ransomware Defender Detections for Your Environment

Home Top

- Case 1 Release < 2.5.6 - If your Current Ransomware Defender Threshold settings are not at or above Warning 80, Major 83, Critical 100
- Case 2 - Release 2.5.7 or later - If your Current Ransomware Defender Threshold settings are not at or above Single Vector Warning 80 Dual Vector Warning 5 signals 30 minutes , Major 83, Critical 100

Case 1 Release < 2.5.6 - If your Current Ransomware Defender Threshold settings are not at or above Warning 80, Major 83, Critical 100

Make the changes below.

1. Open Ransomware Defender Icon, click on thresholds tab.

a. 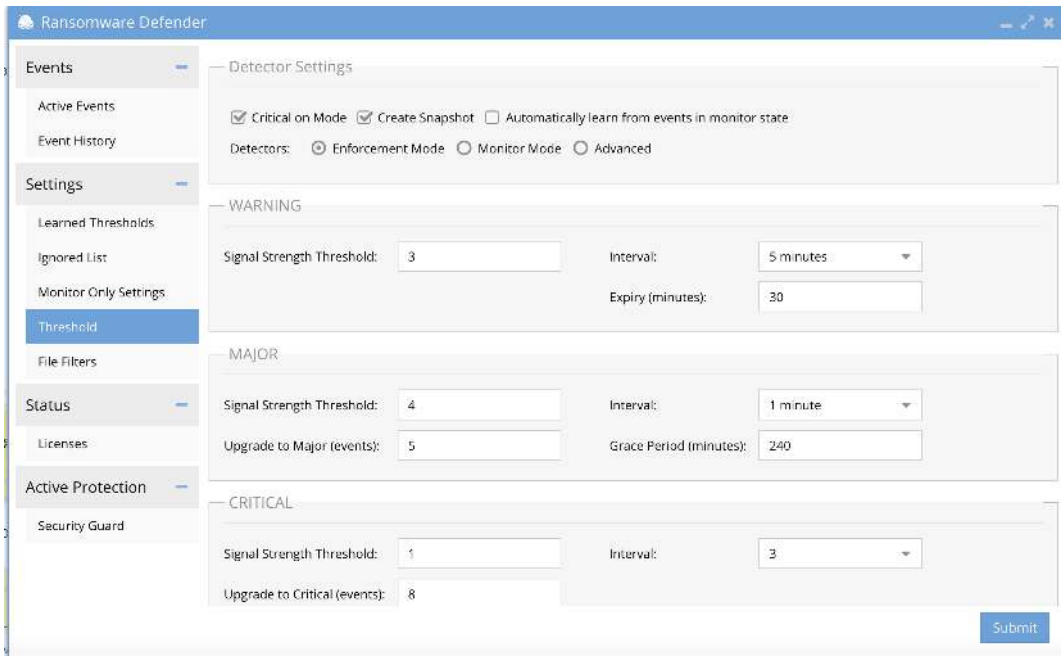

2. Change the signal strength threshold for "Warning" to equal 80
3. Change the signal strength threshold for "Major" to equal 83
4. Change the signal strength threshold for "Critical" to equal 100
5. Click "Submit" to save the changes

Case 2 - Release 2.5.7 or later - If your Current Ransomware Defender Threshold settings are not at or above Single Vector Warning 80 Dual Vector Warning 5 signals 30 minutes , Major 83, Critical 100

Make the changes below.

1. Open Ransomware Defender Icon, click on threshold

2. 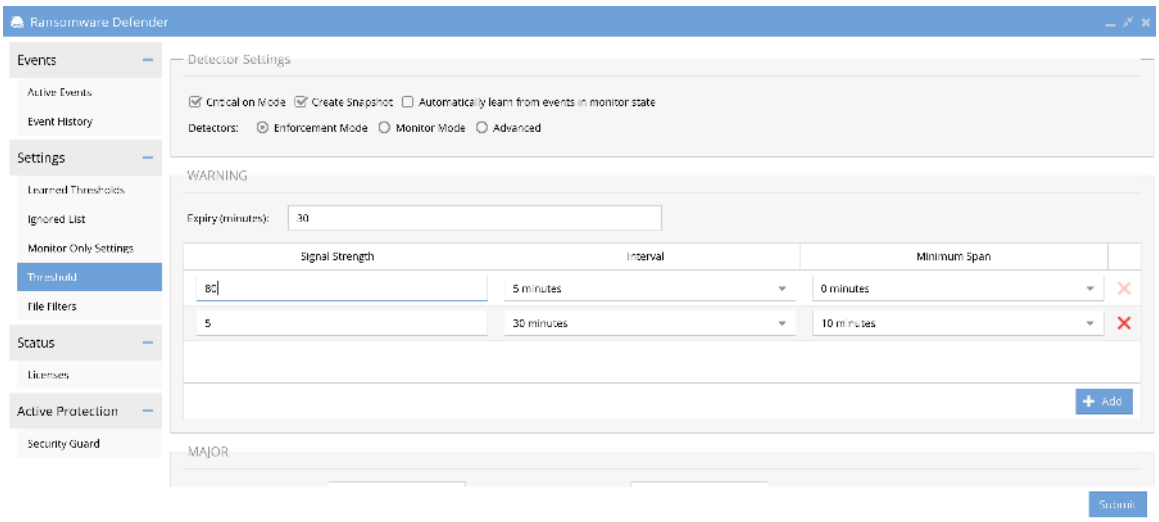

3. Change the signal strength threshold for "Warning" to equal 80
4. Change the signal strength threshold for dual vector "Warning" to equal 20 signals, 30 minute interval and 10 minute 2nd vector window.
5. Change the signal strength threshold for "Major" to equal 83
6. Change the signal strength threshold for "Critical" to equal 100
7. Click "Submit" to save the changes

© Superna LLC

2.4.3. Ransomware - Security Event Explanations

Home Top

- Security Event Column Definitions
- Shares and Snapshots columns in Ransomware Active Events
- Files and CSV Download Column Ransomware Active Events
- Threat Detection Signal Strengths in Ransomware Active Events
  - Signal Strength Window In Ransomware Active Events

Security Event Column Definitions

- **File Event:** a discrete event published by PowerScale's event stream based on a user action, for example: open file, close file, write or read to file.
- **User**: The user or NFS IP of the locked out account
- **Shares** : List of shares that were locked out
- **Snapshots**: List of snapshots of SMB share paths taken during the lockout
- **Threat Detectors:** Logic used by Eyeglass Ransomware Defender to determine if a group of File Events is potentially associated with a Ransomware attack. There are multiple independent threat detectors used by Eyeglass Ransomware Defender during analysis that are assessed in parallel.
- **Signal:** Occurrence of one or more File Events that have been flagged by one or more threat detectors as a potential Ransomware Event.

- **Signal Strength:** For a given Signal, the number of threat detectors that were triggered. A higher Signal Strength has a higher probability of being a Ransomware Event.
- **Ransomware Event:** A collection of signals whose combined Signal Strengths exceeds the user-set threshold in Eyeglass.
- **Files**: The list of files associated to the detection for the specified user.
- **Lock Out**: The date and time of the security event
- **Client IP**: The ip address of the users PC that is infected.
- **Actions:** The history of all steps taken during the detection is listed with time stamps. All share lockout or snapshot create tracking is logged along with any failed API calls to lockout or create snapshots. All future actions stay with the event history, example restore user, comments, flag as false positive etc... If an event is archived the history will stay with the event in the Event History tab.

Shares and Snapshots columns in Ransomware Active Events

1. The shares column lists shares that were locked out for this specific user listed in the User column. The shares are detected based on AD group membership of the detected user and locks out on all shares on all clusters managed by the Eyeglass.
2. The snapshots column lists the snapshots that we triggered during the lockout. The snapshots are taken on the SMB share paths of all shares on all clusters as detected as accessible to the user. The snapshots default with an expiry of 48 hours.

Protected shares for user: AD3\dfsad3 event: #17:4030

Share	Cluster	Zone

Permissions

permission: READ **permission type:** DENY

trustee

id: SID:S-1-5-21-2947135865-3844123249-188779117-1129

name: AD3\dfsad3, **type:** USER

permission: FULL **permission type:** ALLOW

trustee

id: SID:S-1-1-0

name: Everyone, **type:** WELLKNOWN

Ransomware Defender

Events
- Active Events
- Event History

Settings
- False Positive
- Ignored List
- Monitor Only Settings
- Threshold
- Allowed Extensions

Status
- Licenses
- Threats Detected

Active Events

State	Severity	Files	Signal ...	User	Detected	Shares	Snaps...	Expires	Clients	Actions	Locked...
LOCKE...	CRITIC...	41+ files	0/8/8	AD01\...	8/3/20...	none l...	none c...	n/a	Clients...		8/3/20...
LOCKE...	CRITIC...	3+ files	0/3/3	AD01\...	8/3/20...	9 shares	2 clust...	n/a	Clients...		8/3/20...

Snapshots created for user: AD01\dfs1 event: #17:4319

Isilon Cluster	Snapshot	Created at
prod8	igls-AD01-dfs1-data-share1-17_4319-...	8/3/2020, 6:14:26 PM
prod8	igls-AD01-dfs1-data-dfs-17_4319-159...	8/3/2020, 6:14:37 PM
prod8	igls-AD01-dfs1-System-smb2-17_431...	8/3/2020, 6:14:25 PM
prod8	igls-AD01-dfs1-System-veeam-17_43...	8/3/2020, 6:14:33 PM
dr8	igls-AD01-dfs1-System-igls-dfs-smb2-...	8/3/2020, 6:14:27 PM
dr8	igls-AD01-dfs1-data-igls-dfs-dfs-17_4...	8/3/2020, 6:14:29 PM
dr8	igls-AD01-dfs1-System-dfs-17_4319-1...	8/3/2020, 6:14:25 PM

Files and CSV Download Column Ransomware Active Events

1. The files section allows browsing of files that triggered the detection. A CSV download option provides a historical file list that contains all files from all detections for the specific user. This provides a history for all files that triggered events for each

user when they have had multiple detections. The files will rollover and maintain a current file and then historical files will be listed with a data and time stamp.

2. **NOTE:** The CSV rows will list signals , not files in each row. The files column cell for each row in the csv will list all files related to the signal. If using Excel double click the cell to see the list of files for that signal. **Remember a signal is a pattern of user behavior that affects multiple files.**

    a. The CSV will store the first 1000 files associated to a detection, there can be more files related to a detection that will not be stored in the CSV. The CSV could get very large to store all files so only the first 1000 are stored.

    b. In most cases, the CSV will store all affected files, it is best to use Easy Auditor to run a user report to get a complete list of files touched by a user.

    c. 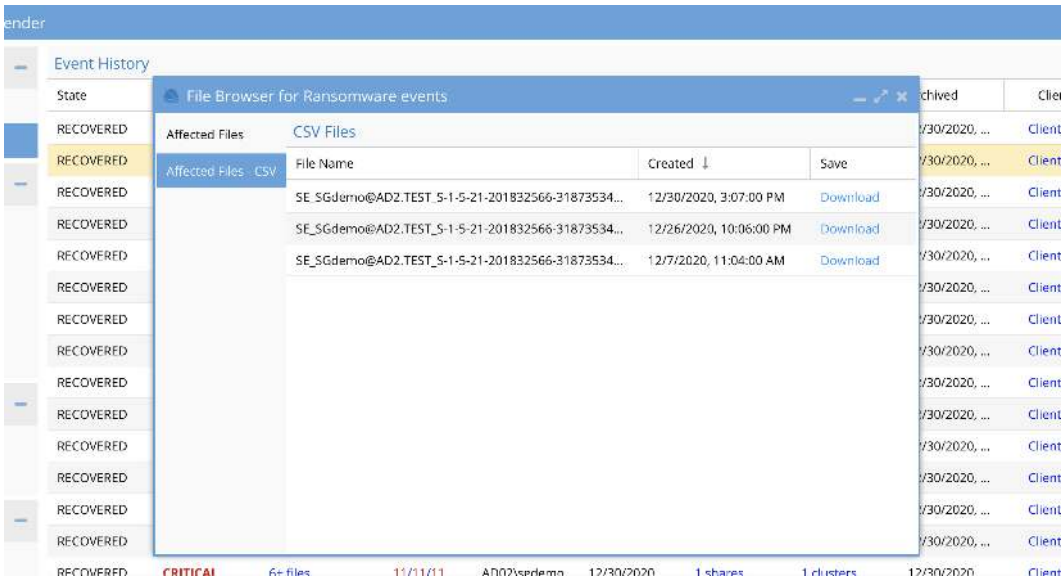

Threat Detection Signal Strengths in Ransomware Active Events

Threat detection Signal Strength is a measure of the severity of a user's File Event behavior. The higher the count the higher the severity of the detection.

The Signal Strength is displayed in the Eyeglass Ransomware Defender window Active Events or Event history tab.

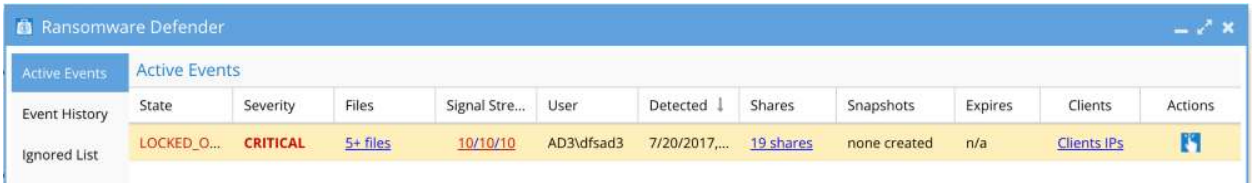

The numbers represent the peak of Warning, Major, and Critical signal strengths that were recorded in the entire lifetime of the Ransomware event. In addition, a list of shares by cluster and access zone is listed that have lockouts applied.

Signal Strength Window In Ransomware Active Events

When clicking on the Signal Strength in Eyeglass, you can see the threat detectors that contributed to the event. Note that this is not broken down by severity, and represents the total of the Threat Detector types that were tripped throughout the lifetime of the event for this user security event. **These are not documented and are for support only.**

Signal Strengths

Signal Strength Breakdown (90)

Type	Value
THREAT_DETECTOR_04	14
THREAT_DETECTOR_07	16
THREAT_DETECTOR_06	60

Note that the sum of these values can be greater than the peak signal strengths described above since it's possible that the lifetime of the Ransomware event is greater than the interval for the thresholds.

© Superna LLC

2.4.4. Best Practice for Tuning Ransomware Defender

Home Top

- Overview
- Best Practise

Overview

These sections cover guidelines on how to decide if whitelist of a detection or flag as false positive is the best option.

Definitions:

1. **Snapshot Critical path (Release 2.5.7.2 >)** - New feature for customers using everyone full control on SMB shares to disable share snapshots per SMB share when ACL security is used. This check box enables or disables share level snapshots . This feature is designed to be used with critical path check box and entering paths to create snapshots each time a user detection occurs only the critical paths will have snapshots created.
    a. NOTE: Critical path and share snapshots can both be enabled. This allows application data ie. a critical path to be snapshotted when a user detection occurs. NOTE: snapshots are only created on SMB shares the user has access to via AD groups. This allows other critical paths to get a snapshot created even when the user does not have access to this critical path.

b. 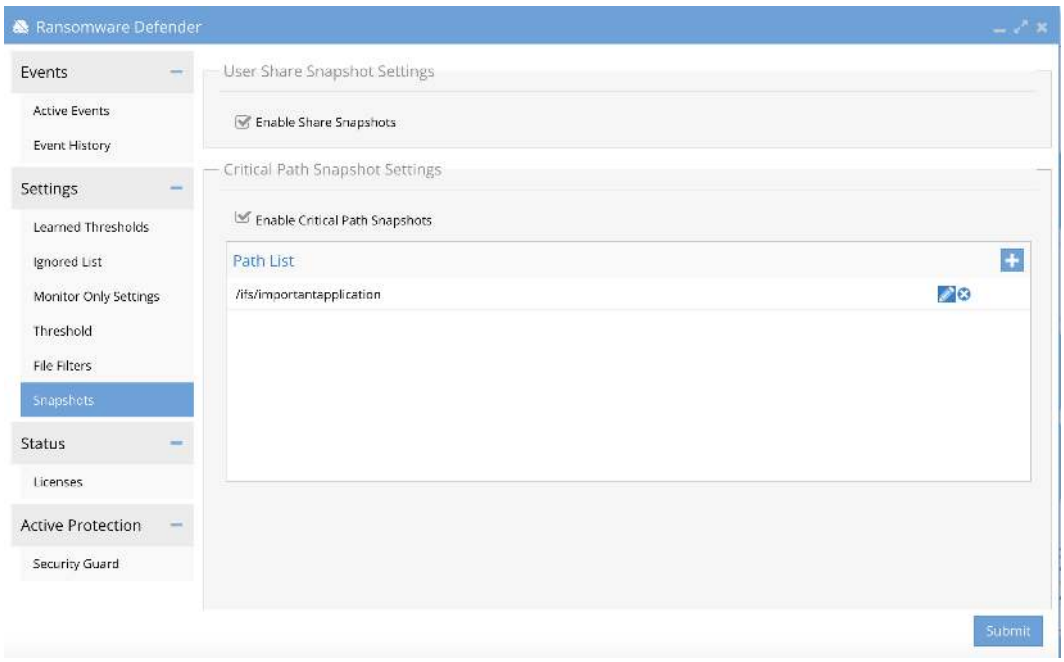

2. **Snapshot Quota Limit (Release 2.5.7.2 >)** - If everyone full control SMB shares are used or many detections occur in a short period of time a lot of snapshots can be created. This feature will allow setting a limit to the number of snapshots that can be created. Once snapshots expiry, it will allow more snapshots to be created up to the quota limit. The quota can be changed to a different value. The default is 1000 snapshots.
3. **Monitor Mode Lists (Release 2.5.7 >)** - Using monitor mode lists allows path , user or IP address detection and alerting with lockout actions. This should be done for any application where lockout would not be the desired action but alerting and snapshots are needed to protect the data created by the application. Learn how to configure Monitor mode lists here.
    a. **Whitelist option** - Using the whitelist option for path, user or ip address will ignore any detections for any of these configured white list entries. This should be used for data that does not require any monitoring or protection and

raising alerts would generate too many false positives. Learn how to configure whitelist here.

4. **Flag As False Positive Option** - This applies custom detection rules for the user listed in an event. This option will not ignore future detections by the same user and will allow all detection methods to detect, lockout and snapshot. This should always be used to tune settings for any detection. Learn how to manage flag as false positive here.
5. **Allowed Extensions Option** - This removes a file extension from the banned list of well known extensions. See the list of extensions published here. This option means the file extension used by the user listed in the event is on the banned list. You should review the extension and confirm this is a legitimate extension used by an application in your environment. Once you have confirmed this is a safe extension you and remove it from the banned list following the CLI guide example here. Learn how to manage banned file types here.
    a. Once a file extension is on the allowed list all future detections by Any user will no longer trigger a detection.
    b. **NOTE: The files with an allowed extension are still protected by user behavior detection, this does not reduce protection of the file system data. It will only disable extension detection for this application.**

Best Practise

1. Always apply flag as false positive as the first action on any new event after you have reviewed the detection and determined this is not actual Ransomware, then wait to see if any future detections occur and repeat the flag as false positive steps again.
    a. If you have flagged a user detection event 3 separate times as flag as false positive, then we recommend to monitor mode list this user, see next section.
    b. See guide here on how to flag as false positive.
    c. NOTE: Always use flag as false positive for user event detections avoid a whitelist for actual end users.
2. **Monitor only settings lists** should be used for server based applications versus users. Server applications that modify the file system in a manor that is detected as Ransomware should have the server based application data added to the monitor mode list.
    a. Use a user based monitor list for the service account used by the application Or use a source IP of the server IP address. This offers the best option to allow the data in the file system to be protected but assumes the server itself will be monitored for any Ransomware behavior in the file system.
    b. See the guide on how to configure monitor mode list here.
3. **Ignore list** - Only use this list for data that does not require any protection. Temporary data that can be recreated that is tripping detections should be added to the ignore list. Use a path based ignore list at the highest level of the file system.

a. See the guide on how to configure ignore lists here.

© Superna LLC

2.5. How to Configure, Tune and View Ransomware Defender Threat Detection settings and Responses

Home Top

- Overview
- Severity Threat Level Severity Definitions and Responses
- Automated Configuration with Learning Mode
  - Deployment Overview with Learning Mode
- Threat Detector Threshold Manual Configuration and Definitions
  - How to change Threat Detection Settings
- How to Enable “Monitor Only Mode” to Disable User lockout Actions
- How to Enable/Disable Critical Event Severity Detection
- How to Disable Snapshot Action for all Detections

Sub Topics

- How to Configure Monitor Mode and Ignored Lists
- How To Manage False Positives and Learning Mode
- Banned and Allowed File Type Configuration
- Rapid Machine to Machine Malware Spreading Attack Defense
- How to Manage Threat Detectors - Advanced Consult Support
- How to Configure Snapshot Modes (Critical Path and SMB share snapshots) and Snapshot Quotas

Overview

The detection of a Ransomware event will be contained strictly to the ECA nodes. Ransomware Defender will be responsible for taking action against the users access to cluster data and notifying administrators. This section identifies the behaviors that the Ransomware Defender takes when the ECA identifies a threat and how to configure settings that align to your company security policies or risk tolerance.

This section explains how to understand the settings and the impact of making changes to the detection settings and is provided as a reference. **The recommended settings are defined in the Tuning Ransomware Defender Detections for Your Environment and changes to these settings should be done by consulting with support.**

Severity Threat Level Severity Definitions and Responses

There are three Signal Strength Threshold levels defined, and Ransomware Defender will take different actions for each:

Threat Level	Ransomware Defender Action
WARNING	Eyeglass sends an email to notify any subscribed administrator(s) of the threat but takes **no** direct action.
MAJOR	Eyeglass begins a "**delayed lockout**" procedure. Notify the administrator(s) that a threat has been detected, and the user will be locked out after X minutes, unless the admin logs in and explicitly cancels the action. This grace period is configurable in Eyeglass settings.
CRITICAL	The user lockout is **immediate**, and the

	administrator(s) are notified.

Automated Configuration with Learning Mode

This new mode requires 2.5.7 or later release and automates all aspects of customizing settings based a mode that learns about applications, and user behaviors used in each environment that Ransomware Defender is deployed.

Deployment Overview with Learning Mode

1. This new mode will change the deployment process to enable monitor mode and then enable learning mode after deployment.
2. After deployment monitor mode and learning mode are enabled by default.
3. Wait 2-3 days after deployment to review settings that have been auto applied by Learning mode. **Learning mode should be left enabled for a normal work flow time period to allow it to "see" all the work flows in your environment.**
4. Login to Ransomware Defender --> settings section
    a. View the learned Thresholds tab and review all users listed.
        i. If any of the listed users are application server service accounts. Then create a monitor only mode entry for this service account. This will disable lockouts but still alert and snapshot. See the guide here how to add monitor mode user account. Repeat for each service account.

b. View the File Filters tab

i. Filter the list based on the monitor status. This status is set by learning mode when a user detection is based on a banned file extension. These detections will set the extension to monitor state (alert, snapshot no lockout)

1.

ii. Review the list of application extensions to verify each extension listed is a legitimate application used in your environment.

iii. For applications that are legitimately using an extension on the banned list, change the state to disabled (this allows the file extension use and will be ignored for future detections). Make sure to click save.

iv. Leaving an extension in monitor state means it will still trigger detections and snapshots but no user lockout will occur. **Leave this state if you are unsure about this extension in your environment.**

v. To switch the extension from monitor state to enforcement mode change the state to enabled and

click save. **Make this change to enabled state when you want this extension to trigger a user lockout.**

vi. For application extensions that do not appear to be legitimate in your environment you should investigate the application extension to verify it is not a Ransomware detection.

vii. If you have custom extensions that you want to add, use the add button. **NOTE: Use with caution, this will lockout users if in enabled status.**

5. Exit Learning mode after the above steps are completed. How to disable learning mode.

Threat Detector Threshold Manual Configuration and Definitions

Ransomware Defender allows the administrator to configure the thresholds to determine what actions are taken based on crossing the thresholds. For a detailed explanation on how detection works see this section.

- Different thresholds are available for the WARNING, MAJOR, and CRITICAL severities of a detection.
- The MAJOR severity also allows the specification of the Grace Period (the time between event detection and lockout). Timed Lockout can be stopped with action menu on an active event.

The figure below shows the settings UI.

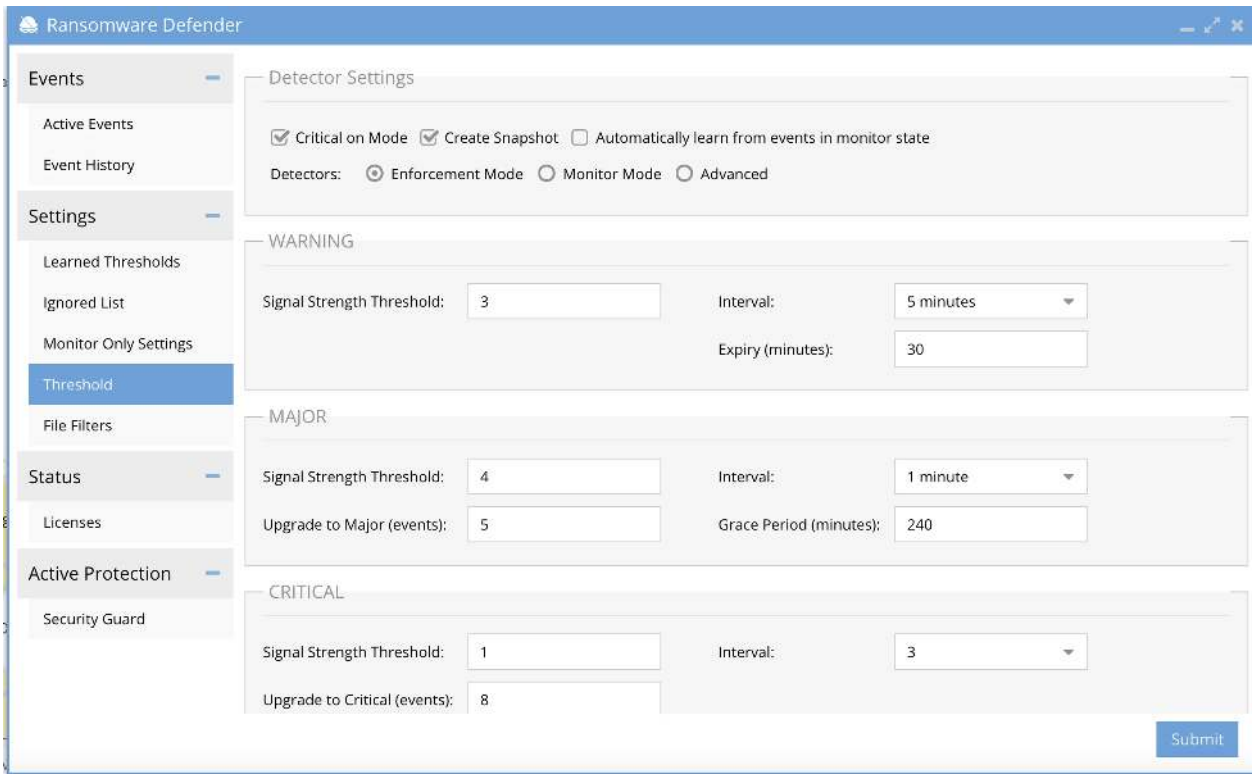

How to change Threat Detection Settings

1. Open Ransomware Icon
2. Click the Settings tab --> thresholds change the thresholds and click submit to save the settings.

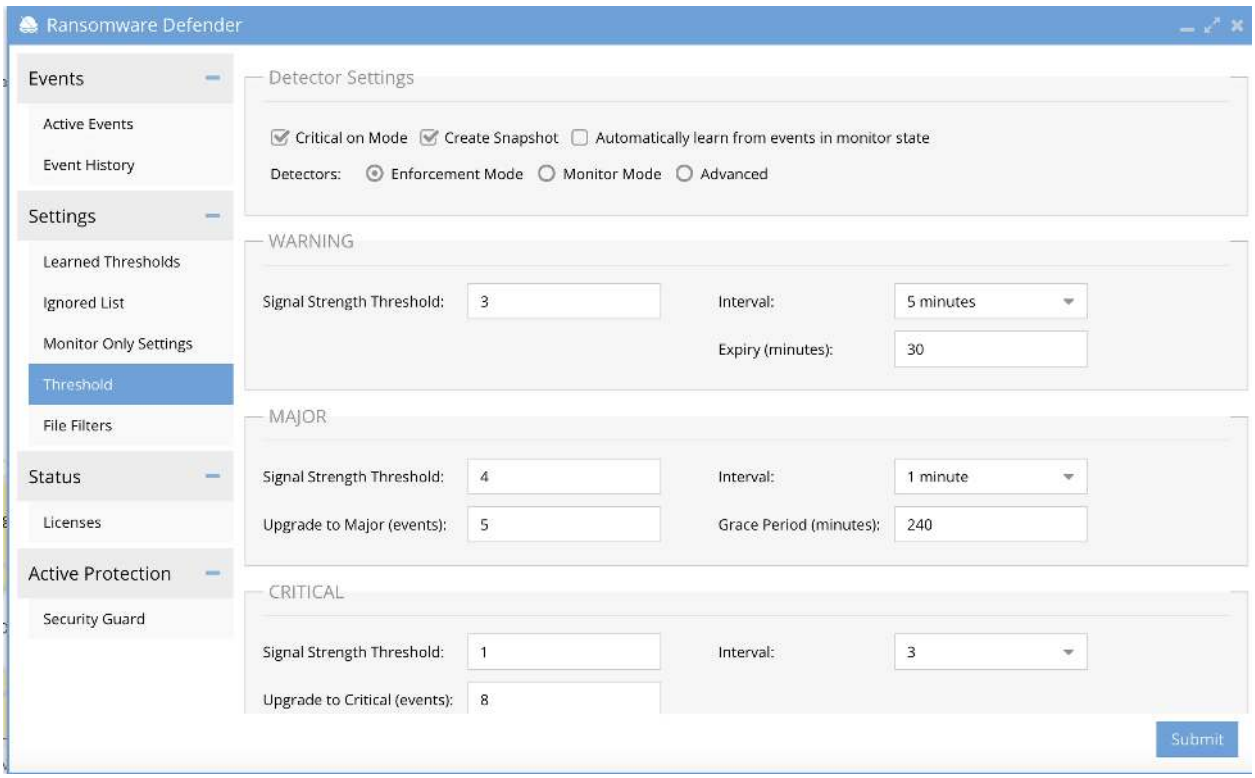

3. Upgrade to Major (events)

This setting will advance all warning events from Warning to Major. This assumes that many warnings in a short period of time is suspicious and promotes all Warning events to a Major severity. This accelerates the response if a large number of users enter the Warning detection level *even though the Major Signal Strength Threshold has not been crossed.*

For example, based on the above settings, if there were 8 Active Events at the Warning threshold for different users, those events would be advanced to Major severity, even though the Signal Strength for any of those events had not met the 8 Signals in 5 minutes Major Threshold configuration.

Upgrade to Critical (events)

This setting will advance an event from Major to Critical for the case where the number of users configured here has an active Major event even though the Critical Signal Strength Threshold has not been crossed. This assumes that many Major detections is suspicious and promoting the severity to lockout the users sooner as a proactive response to protect file system data.

For example, based on the above settings, if there were 10 Active Events at the Major threshold for different users, those events would be advanced to Critical severity e*ven though the Signal Strength for any of those events had not met the 5 Signals in 1 minute Critical Threshold configuration*.

4. The lower the Signal Strength Threshold the more sensitive the detection will become. Changing to a larger number can avoid false positives depending on IO patterns within your PowerScale environment.
5. The Grace Period (minutes) sets how long a Major security event detection will wait before locking out the user named in the security event.
6. **Best practice:** This should be set to a value that ensures an administrator can review the event and determine if lockout should occur or be canceled. It is the response to review an event before the lockout occurs.
7. **NOTE: Recommended to consult support before making any changes.**

How to Enable "Monitor Only Mode" to Disable User lockout Actions

Monitor mode is used after installation to disable any actions for Major and Critical events and to baseline the environment. It can also be used to quickly disable user actions if too many false positives are detected. All detections will continue to alert administrators but no action will be taken to lockout. Snapshots will still be applied during detections.

See steps to enable monitor mode and enter enforcement mode.

How to Enable/Disable Critical Event Severity Detection

This option will disable immediate lockout action and will only use the major timed lockout option. This is recommended if risk tolerance for a lockout on users should be reviewed by an administrator, using the timed out lockout feature on Major severity detection.

NOTE: For this setting to take effect all active events must be archived as resolved or unresolved before changing the mode.

1. Open Ransomware Defender window.
2. Under Settings Thresholds
3. Select the checkbox “Critical on Mode” so it is unchecked and click submit to save.
4. This will disable the automatic lockout feature if the signal threshold is reached for Critical and means only delayed lockout Major severity will be the highest level of user action response.
5. Done.

6. 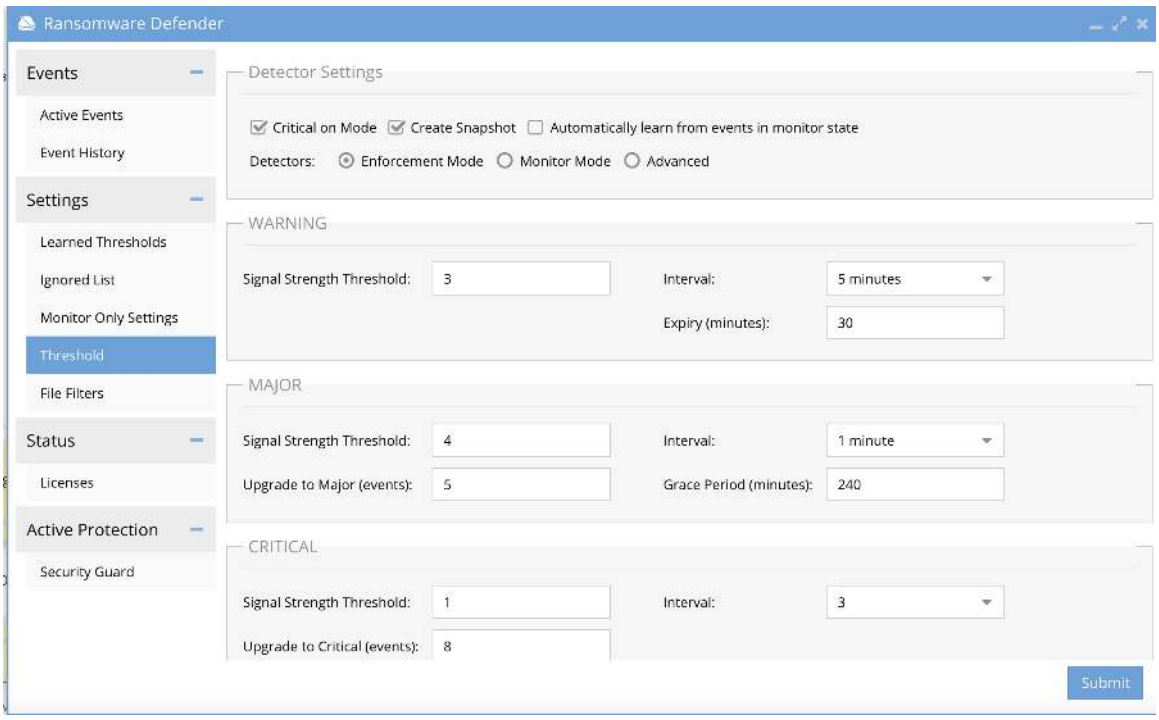

How to Disable Snapshot Action for all Detections

This is not recommended for normal operations. Snapshots are critical part of the solution for data recovery. Consult with support before disabling.

1. Open Ransomware Defender icon
2. Under settings, thresholds window
3. Uncheck Create Snapshot

4. 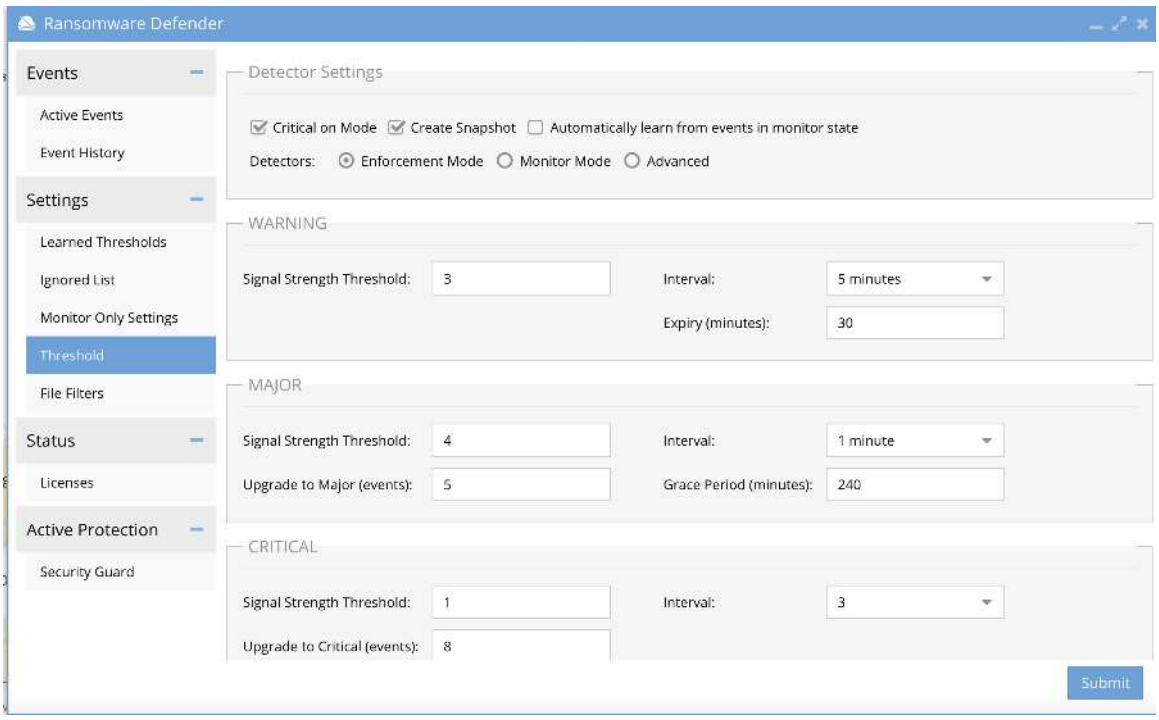

5. Click Submit button to save

© Superna LLC

2.5.1. How to Configure Monitor Mode and Ignored Lists

Home Top

- Monitor Mode List Overview
  - How to Configure Monitor Mode List
  - How to Enable Monitor Mode for an ECS Cluster
- Ignored Whitelists Overview
  - How to convert whitelists to monitor mode lists
  - Ignore List Configuration Procedures for Whitelisting path, user or source ip address
- Partial Path Matching Whitelists and Monitor mode lists and Wildcard Path Match Lists
  - ECS Object Pattern Matching for bucket and path
  - How to apply a partial path whitelist for roaming profiles

Monitor Mode List Overview

This option requires 2.5.7 or later release. This provides new option that is preferred over a whitelist that will ignore all detections for user, path or IP. This option allows a monitor mode to be applied to a user, path or source IP and has the same rules to add entries as whitelisting.

Protection offered by this option allows applications, service accounts to be monitored for detections with alerting and snapshots but no

lockout will occur. This is preferable to ignore detections for applications or service accounts.

See how to convert ignore list to monitor mode list below.

How to Configure Monitor Mode List

1. Open Ransomware Defender window.

   Select Settings --> Monitor Only Settings

2. 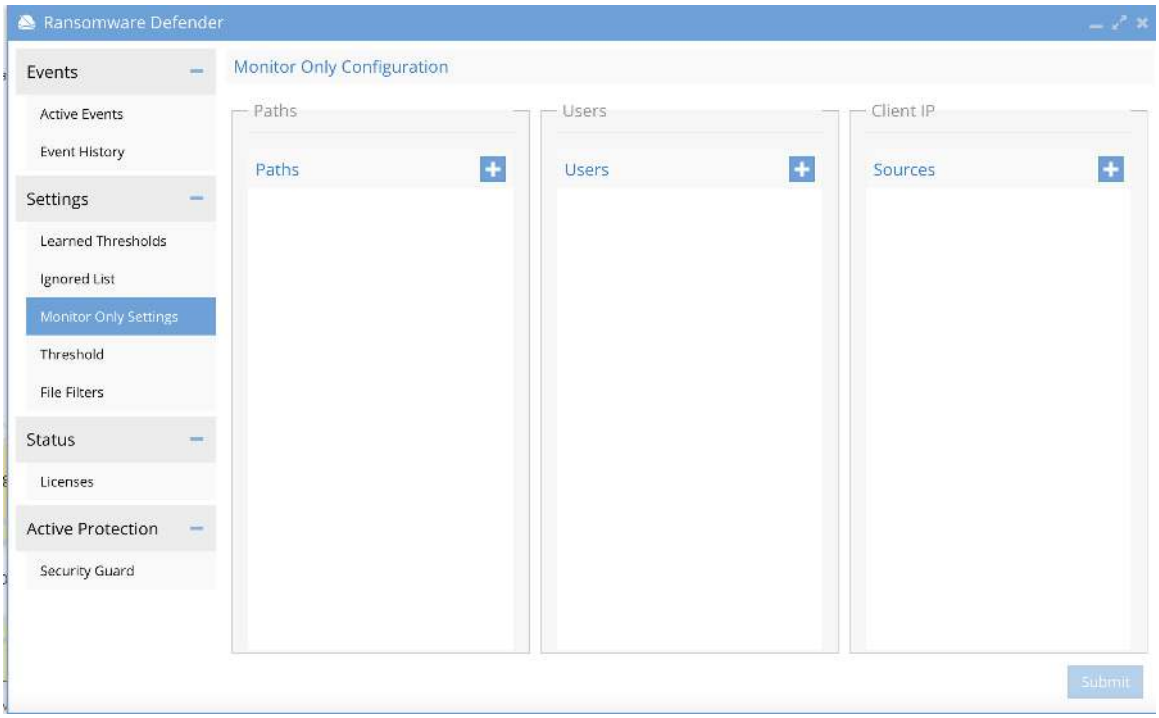

3. Use the + to add a path, AD service account or source IP for servers and click submit.

   a. Select Isilon or ECS and fill in the information

b.

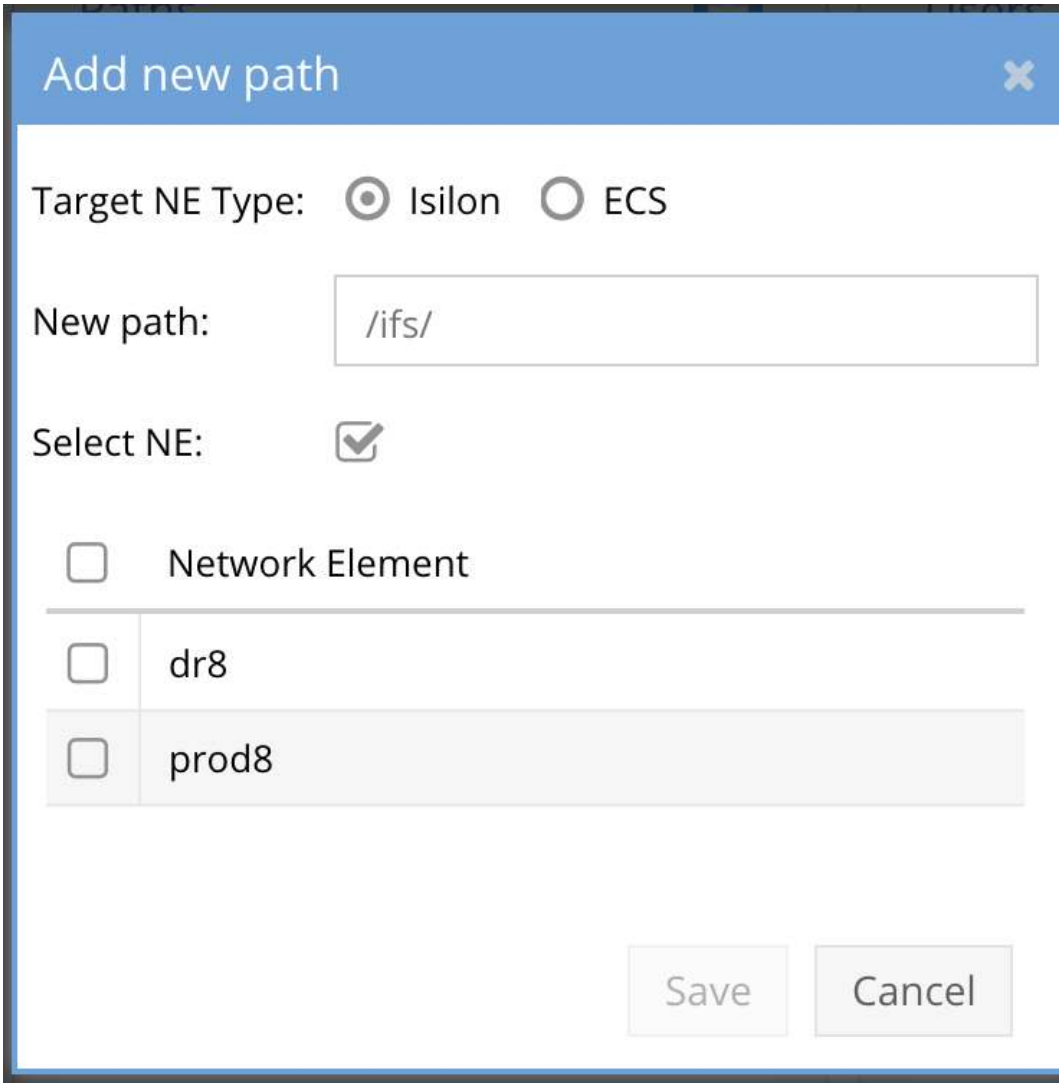
Add new path
Target NE Type:
Isilon
ECS
New path:
/ifs/
Select NE:
Network Element
dr8
prod8
Save
Cancel

c.

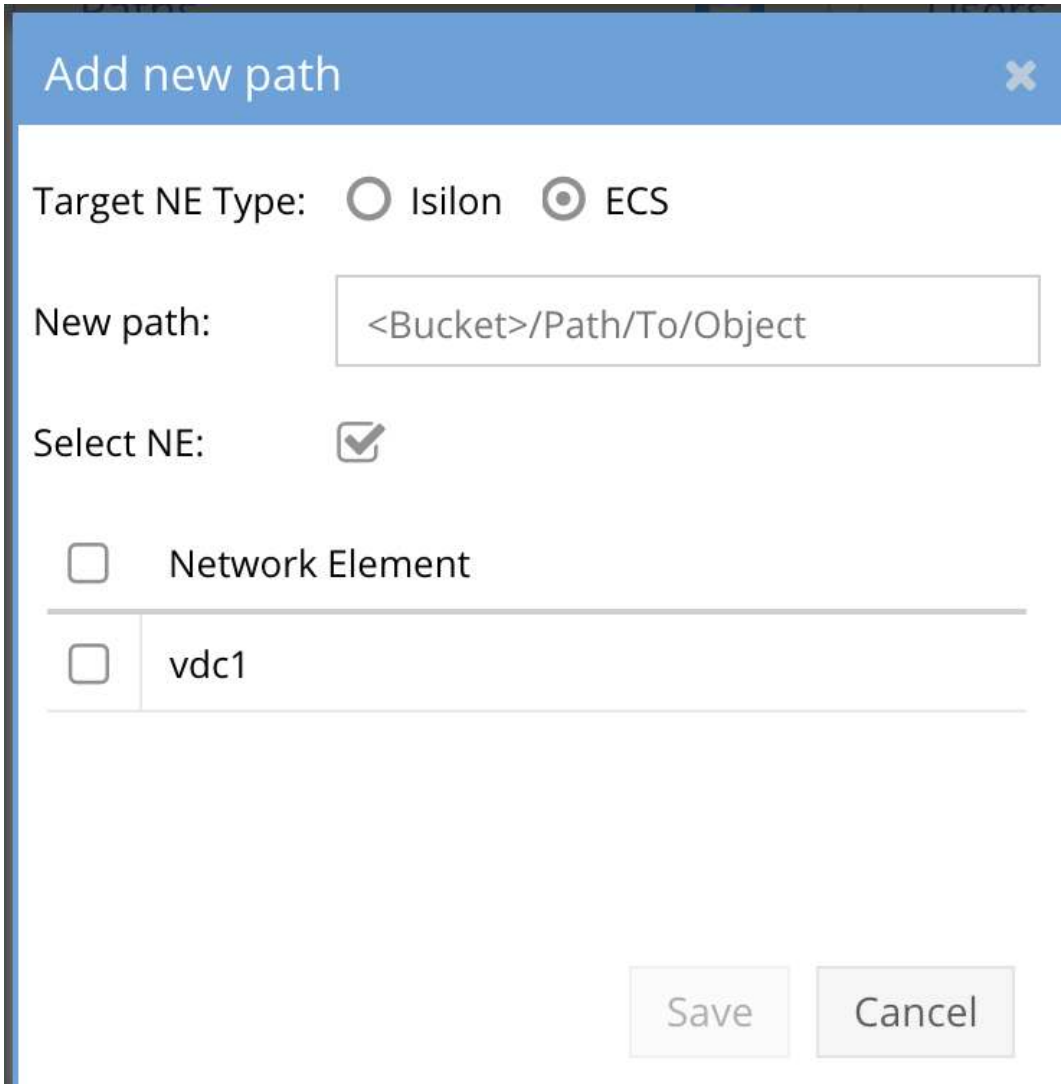

d. For ECS enter bucket name and path to object example bucket1/path1/path2/** to match all objects under path2.

4. The same rules defined for syntax of Whitelists is the same for Monitor mode lists. See pattern matching examples below.

How to Enable Monitor Mode for an ECS Cluster

1. The ECS monitor mode option is set using a wild card match for all buckets all objects. Open the monitor settings tab. Click to add a path entry and configure as per the screenshot.

2. 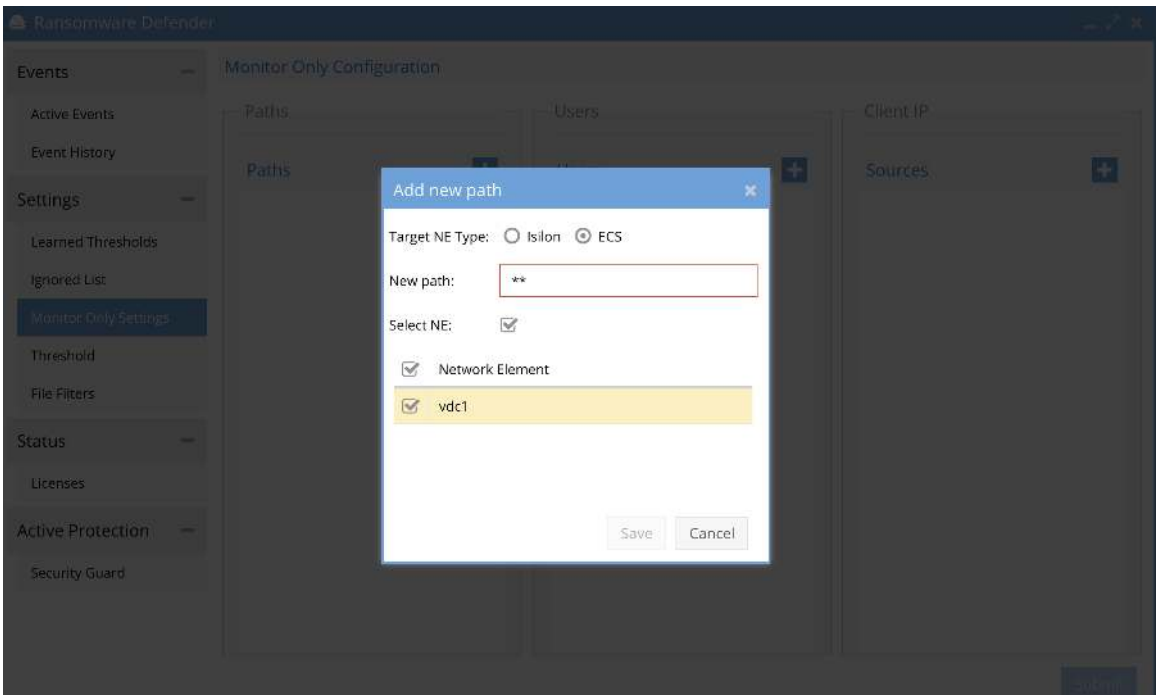

Ignored Whitelists Overview

Ransomware Defender allows the administrator to specify paths, users, and client or server IP address to exclude from Ransomware processing. No detections will be processed once a whitelist is applied based on the matching criteria.

New in 2.5.7 is Monitor only option for Users, Paths, our source IP's.

Best Practise:

This option should always be used before whitelisting since monitor only mode will provide data protection with snapshots and alerting without a lockout.

How to convert whitelists to monitor mode lists

1. This is now the recommended option for all customers upgrading to 2.5.7 to convert the whitelist to the monitor mode list. This offers increased data protection without lockout. Whitelist should be used for less critical data or data that can easily be recreated.
2. Two methods exist to convert the cli command here.
3. GUI option available on the ignore list GUI.
4. See the GUI option at the bottom of the screen.
5. 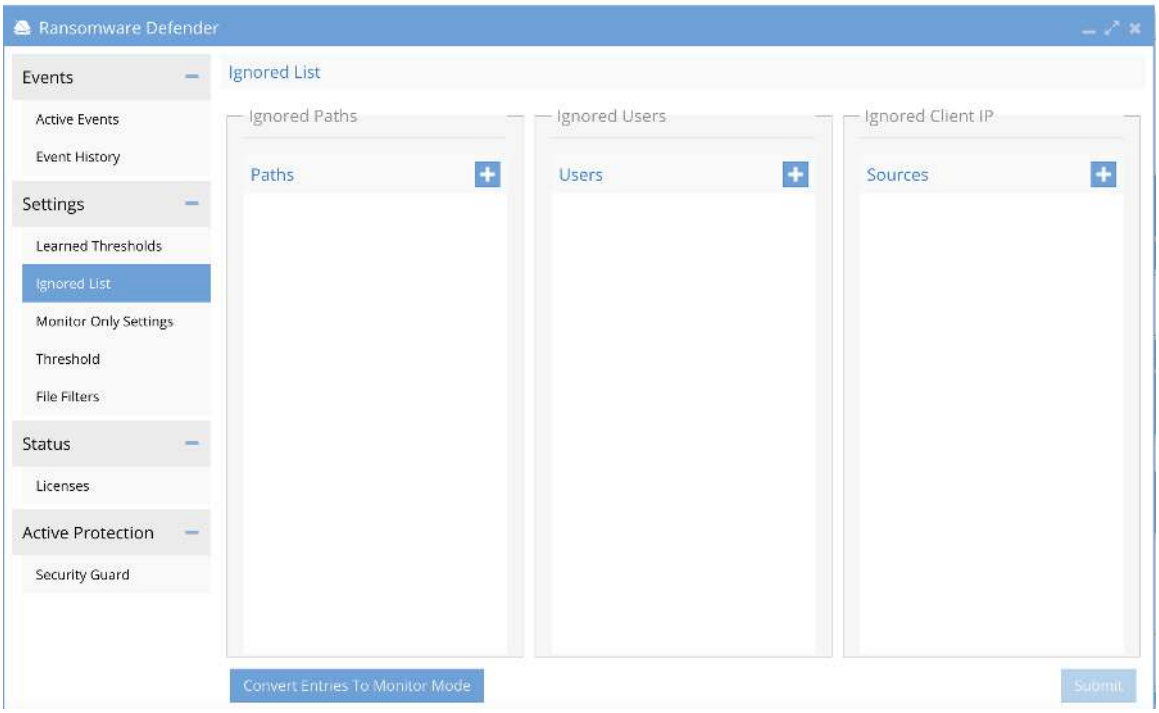

6. You will need to confirm the conversion process that will remove all ignore list settings and move them to the monitor mode list.

7. 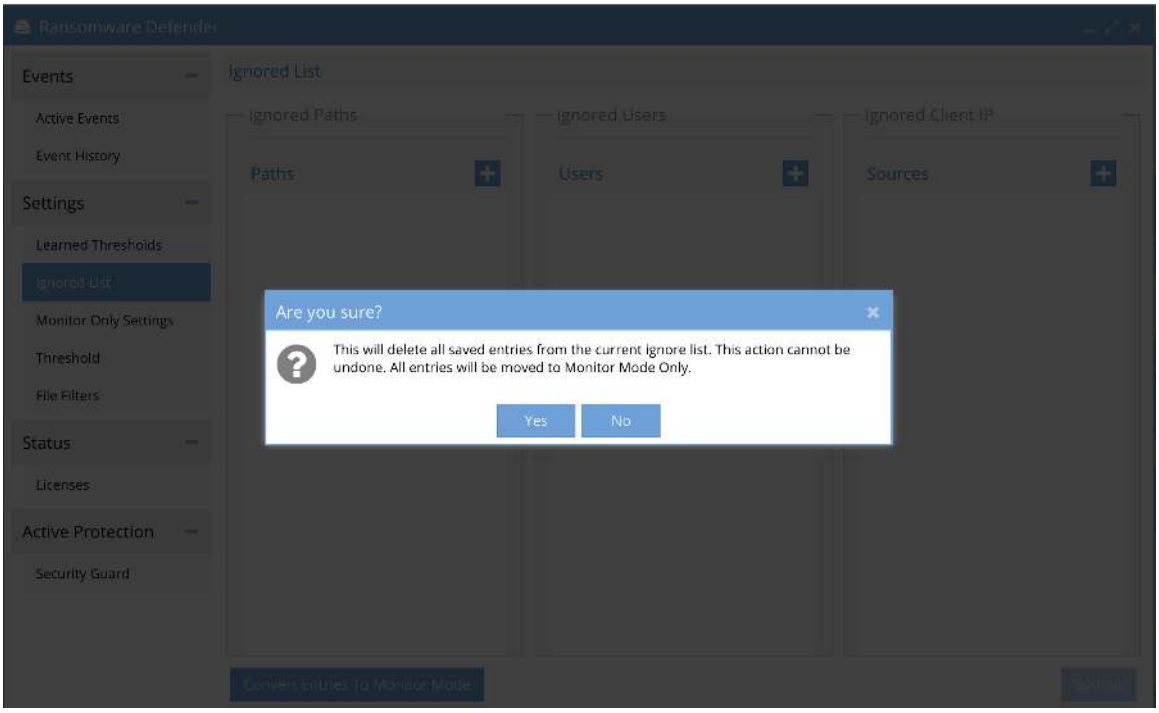

Ignore List Configuration Procedures for Whitelisting path, user or source ip address

Follow the steps below to add ignore list of paths, uses, or server/client source IP. This will skip processing all events based on the settings applied in this section. NOTE: Changes take effect immediately after saved . **Consult support for recommended settings. NOTE: Use monitor mode lists first before using a whitelist that ignores detections.**

1. Open Ransomware Defender window.
2. Select Settings --> Ignored List tab

3. 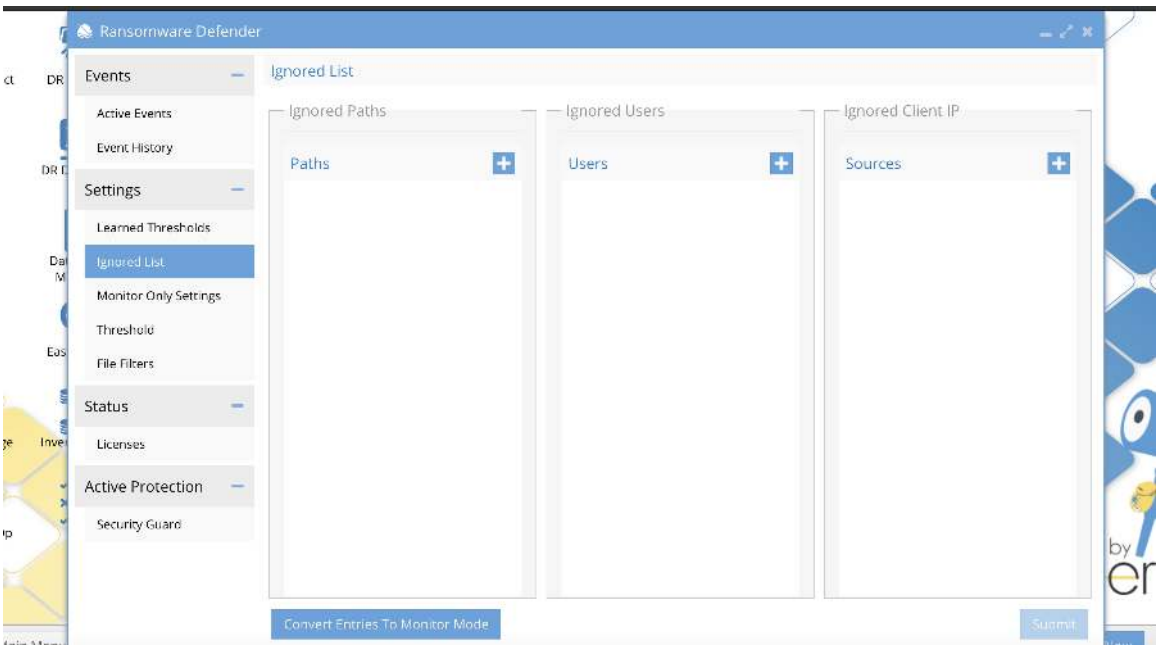

4. Enter a path, AD user domain\userid, or server or client IP address and save.

a. 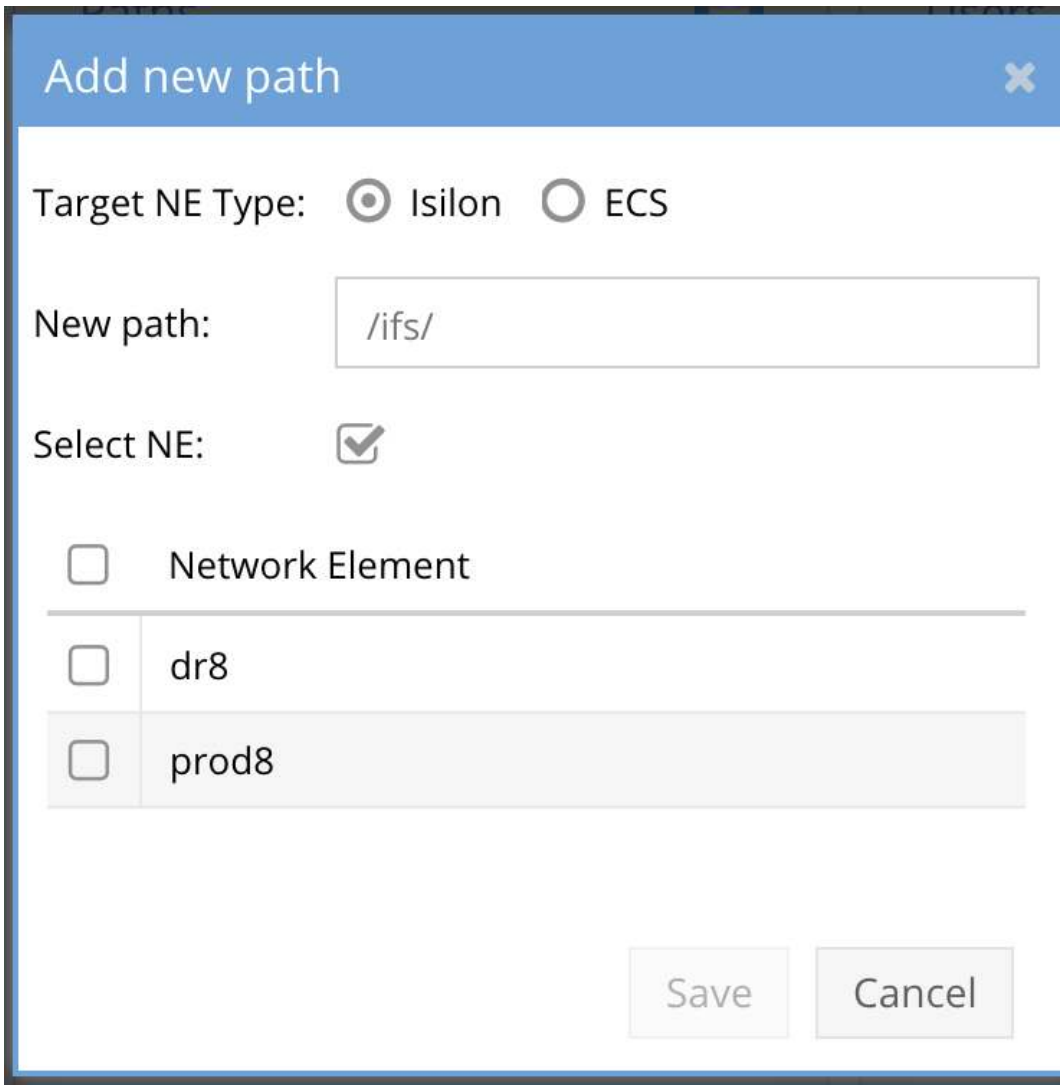
Add new path
Target NE Type:
Isilon
ECS
New path:
/ifs/
Select NE:
Network Element
dr8
prod8
Save
Cancel

b. 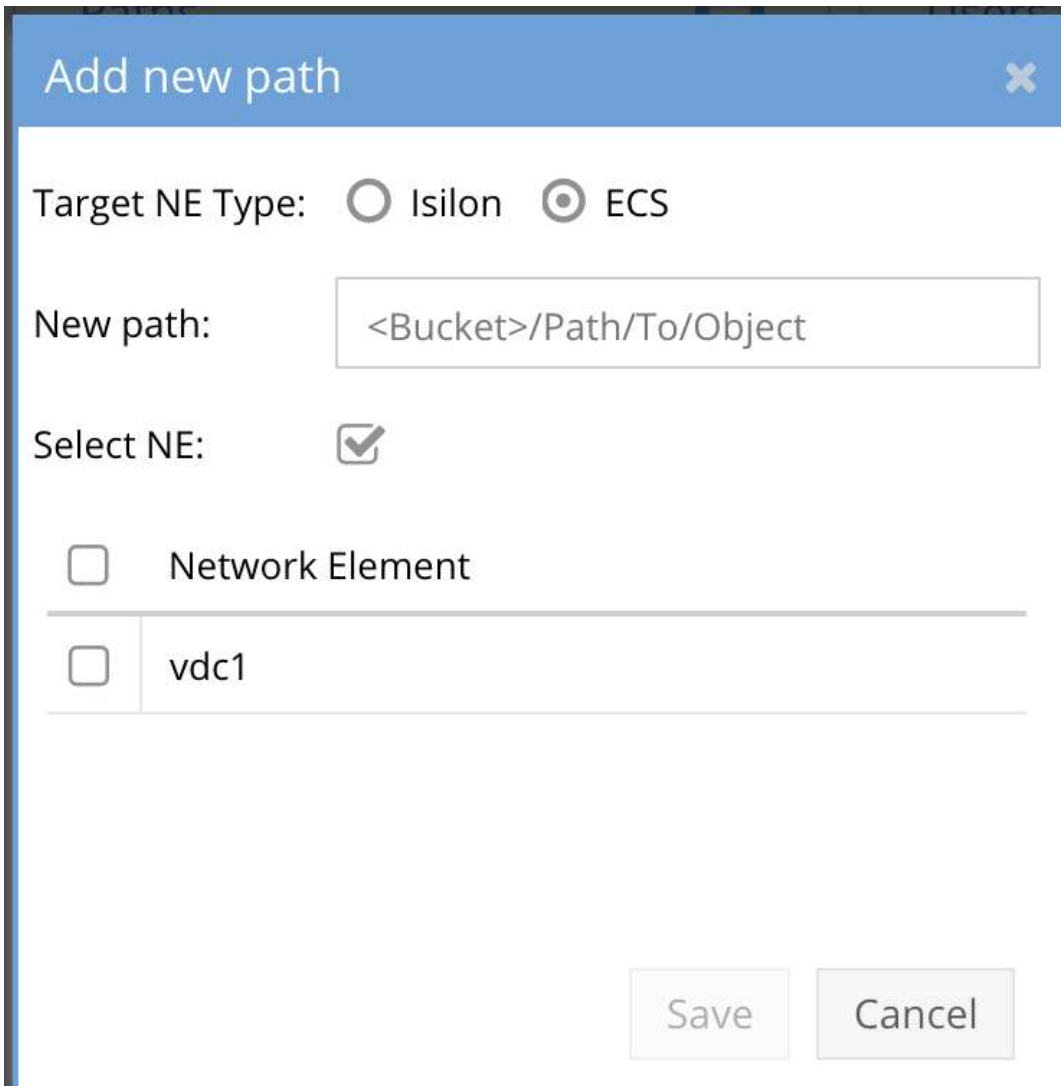

5. NOTE: all 3 options can be used
6. New Paths (full path is required example /ifs/data/xxx)
7. Active Directory Users (domain\userid or user@domainname)
8. ECS users should be enter by name
9. Client IP is the IP of a client or server.
10. **NOTE: each ignore column is an OR, meaning if ANY of the listed ignore values is found in an audit message it will be dropped before processing. The first matched ignore list will drop the audit event.**

Partial Path Matching Whitelists and Monitor mode lists and Wildcard Path Match Lists

Partial path matches for different uses cases, where a partial match is required. This option works for paths or objects paths with ECS objects.

Example, use case is roaming profiles stored on an SMB share. The profile is saved each time the user logs on and off the desktop session. Windows writes data to the roaming profile in a similar IO pattern to Ransomware using a read and write with a hashed file name. This can trip a false positive lockout.

How path and file matching works with whitelists

1. The * character is in place of a directory in a path. Example: /ifs/data/home/*/.
2. If using * at the end of a path, it becomes a file name wild card. Example: /ifs/data/home/usera/* is any file in the usera folder.
3. To wild card multiple subdirectories below "usera" folder this syntax is required. Example: /ifs/data/usera/** This will whitelist all subfolder below "usera" folder.

The portion of the profile where these files are stored for roaming profiles is below the **AppData** portion of the path.

Example of a partial path from a roaming profile: The .tmp file in this example has a file name structure similar to Ransomware variants, and can trip a threat detector and lockout the user.

AppData\Roaming\Microsoft\Windows\Start Menu\Programs\System Tools\{7347b4c0-96d8-45e3-abe2-d7ffde9840a4}.tmp

Example false positive from a roaming profile detection on a SMB share

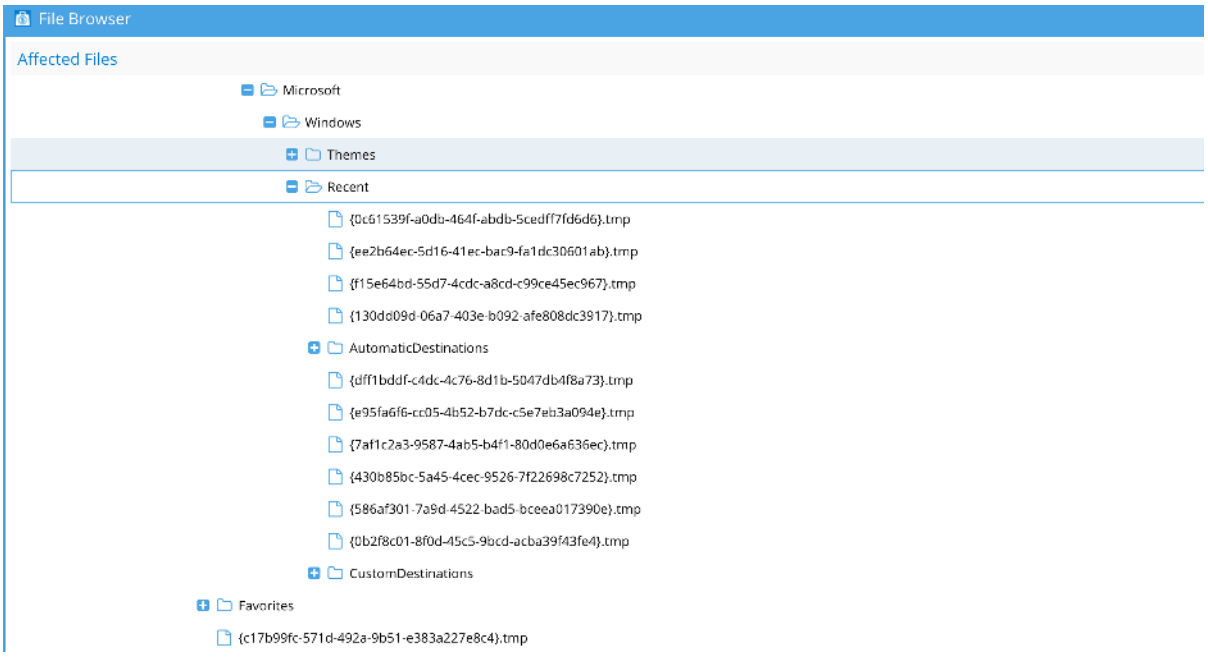

ECS Object Pattern Matching for bucket and path

1. Match the whole ECS cluster using the pattern:
    a. **
2. Match a specific path in any bucket using the pattern:
    a. */path/to/object
3. Match all direct children of a path using the pattern:
    a. <bucket>/path/*
4. Match all descendants of a path using the pattern:
    a. <bucket>/path/**

5. Match exact object:
    a. <bucket>/path/to/object

How to apply a partial path whitelist for roaming profiles

1. Open Ransomware Defender Icon, click on Ignored List tab under settings.
2. A Path-based entry will be used. Click the plus sign and select a cluster from the list.
3. Enter a path using the information below, and click submit to save once done.
4. Example: full path to a share storing roaming profiles for many users.
5. If your unc path to store profiles was this: \\ad1.test\FOdemo\Corpdata\%USERNAME%, and this was an PowerScale path of /ifs/data/corpdata/:
    a. Enter a whitelist path of /ifs/data/corpdata/*/AppData/**
6. This will replace the first * with the user name and the second * with any path in the roaming profile, and ignore roaming profile updates to these paths for all users storing roaming profiles on corpdata home directory path

© Superna LLC

2.5.2. How To Manage False Positives and Learning Mode

Home Top

- How to Teach Ransomware Defender about false positives - Learning Mode
- How to manually flag an security Event as False Positive
- How to View or Delete a Flag as False positive user setting
- How to manually configure per user Threat level settings with IGLS CLI

How to Teach Ransomware Defender about false positives - Learning Mode

1. **Requires:** Release 2.5.7 or later
2. **NOTE: When learning mode is enabled and learning is active a lot of snapshots can be created. Monitor the snapshot usage on your cluster. Snapshots are created with 48 hour expiry by default and will clean up within 2 days.**
3. **Learning Modes**
    a. **Full Learning Mode -** This mode applies to all security events detected an no lockouts will occur and all security events will be used for learning.
    b. **Monitor mode list Learning Mode -** This mode allows both enforcement and learning of monitor mode list entries. In this mode all security events that do **Not** match a

monitor mode learning mode list entry will be enforced and lockouts can occur based on thresholds. For events that match an entry on the monitor mode lists learning will be applied.

i. Use Case: Service accounts or new application work loads can be added to the monitor mode list by path, user or server IP address to allow learning mode to automatically configure settings for this workload.

4. **Full Learning Mode**

a. Enable Monitor mode (settings tab --> Thresholds) to allow user behaviors to be detected without actions taken to lockout.

b. Now enable Learning mode from the Thresholds screen once **monitor mode** is enabled Settings --> Threshold --> click "**Automatically learn from events in monitor state**". Click submit to save.

c. 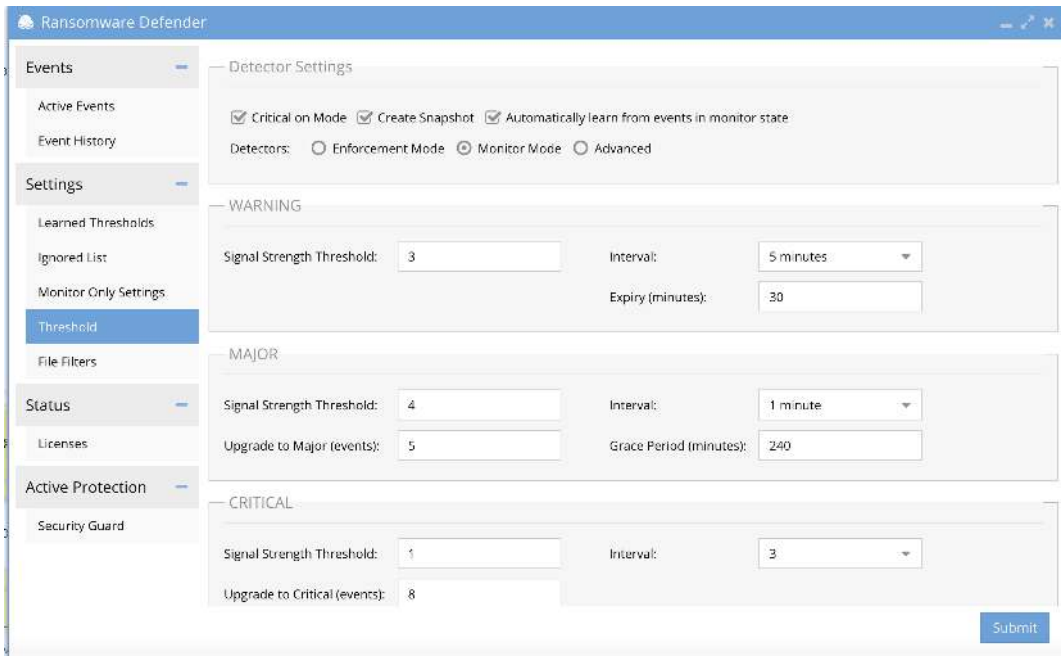

5. **Monitor mode list Learning Mode**

   a. Enable Learning mode from the Thresholds screen once **monitor mode** is enabled Settings --> Threshold --> click "**Automatically learn from events in monitor state**". Click submit to save. Example screenshot below.

   b. 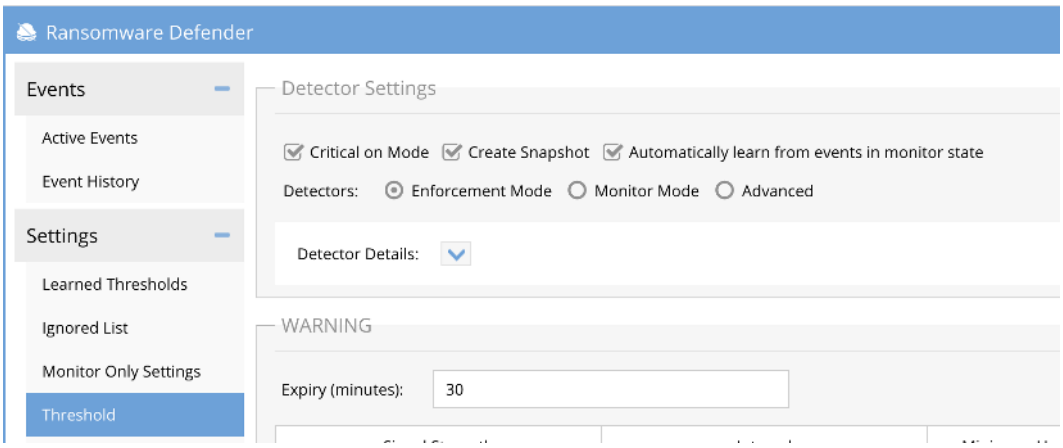

6. Leave this enabled for 2-3 business days and monitor the customized user behavior settings on the Learned Thresholds tab.

   a. This is where Learning mode will place customized settings. It will also set file extension detections on the File Filter tab into a disabled state so this file extension will not be detected as Ransomware.

b. 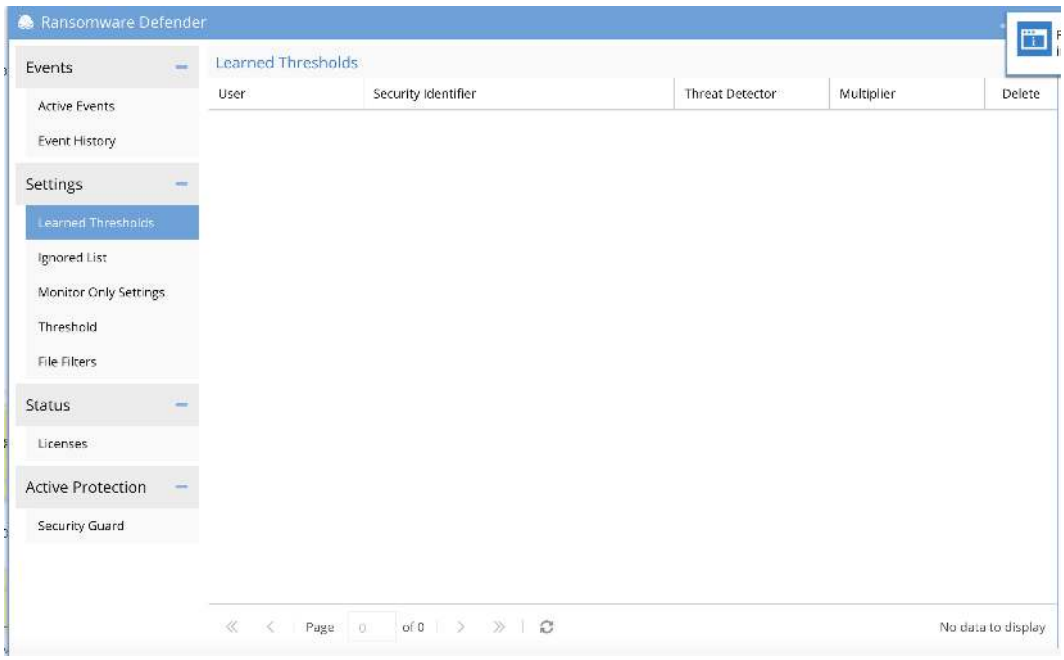

7. **The process to disable Learning Mode and then enter Enforcement Mode.**
    a. Review user settings on the **Learned Thresholds tab** to approve the list of users or NFS hosts or delete entries as needed. Consult with support or accept the learned behaviors.
    b. Review the File Filter list extensions that are disabled status, these extensions have been placed on the Allowed list and will not trigger a detection.
        i. Use the filter option to locate all the disabled file extensions by entering Disabled in the filter box.

ii. 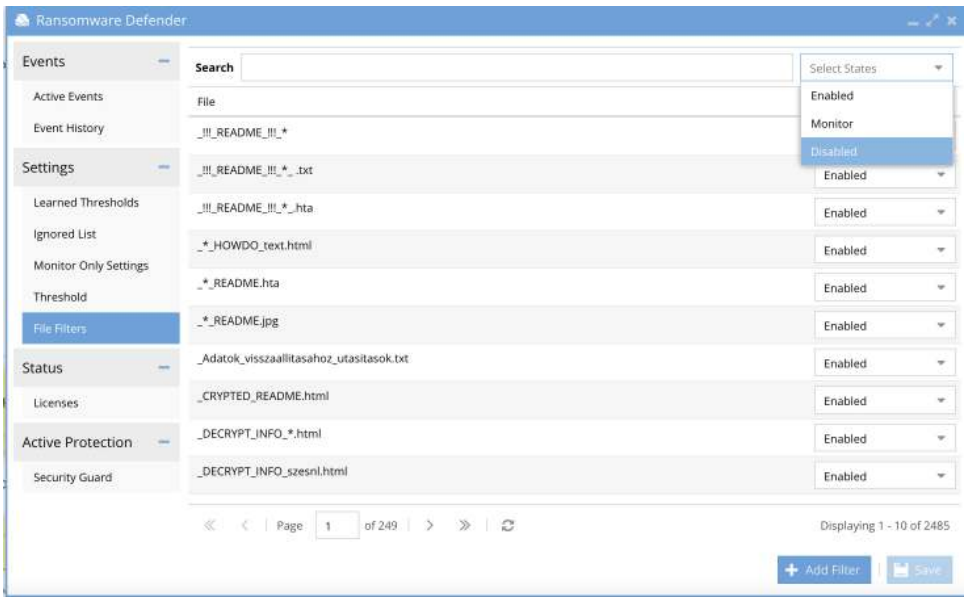

iii. Review all the extensions that were detected and disabled. If they are acceptable no action needed.

iv. To change the setting on the extension to enable enforcement and detection of this file extension, you may also chose monitor mode on the file extension to allow detection, snapshot but no lockout for this file extension.

1. 3 possible modes for each file extension enabled (full enforcement), disabled (ignored), monitor mode (detect, alert, snapshot and no lockout)

c. Disable Learning mode once the file settings are confirmed from the Settings-->Threshold tab and click submit to save. This only disables learning mode and remains in Monitor mode.

d. To enter enforce mode mode disable monitor mode from the Settings-->Threshold tab and click submit to enter enforcement mode.

How to manually flag an security Event as False Positive

1. Open the actions menu and select false-positive action and submit.

2. 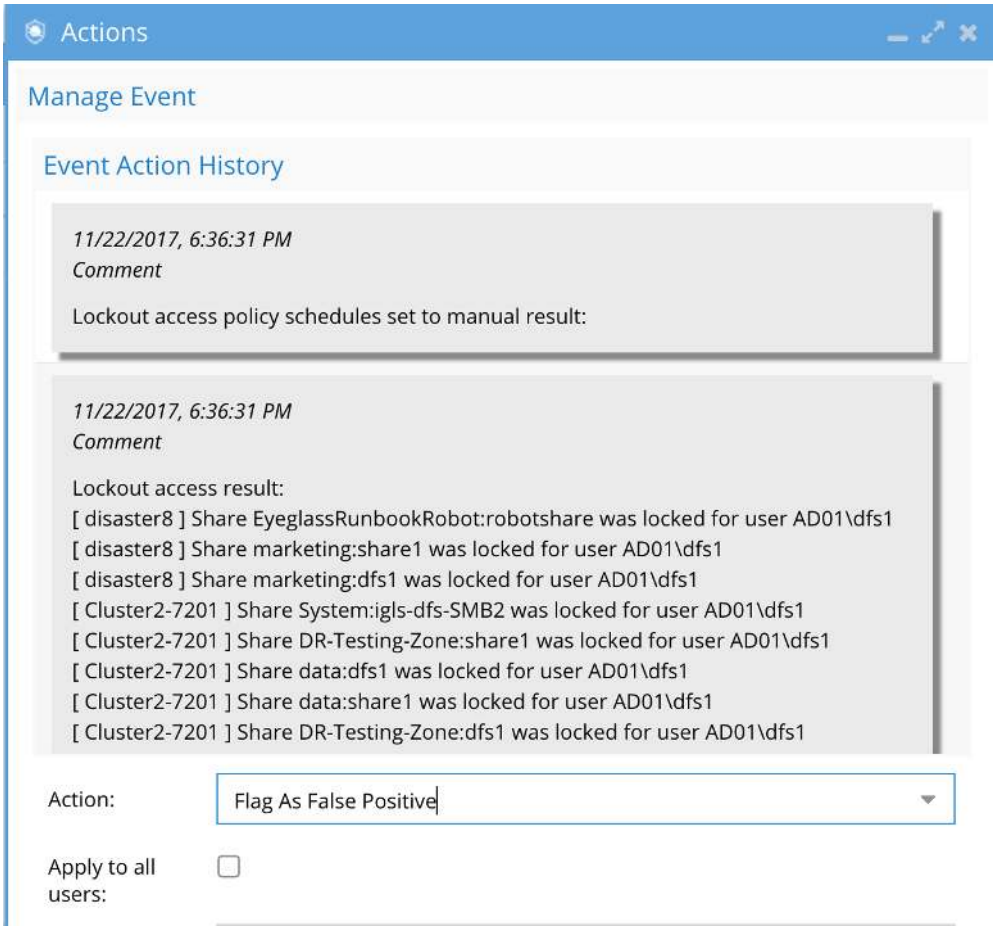

3. This will update the settings for this user. This change is real-time and will take effect immediately.

4. To view the settings for any custom user settings or flag as false-positive user settings click on the False Positive tab under the

settings

menu.

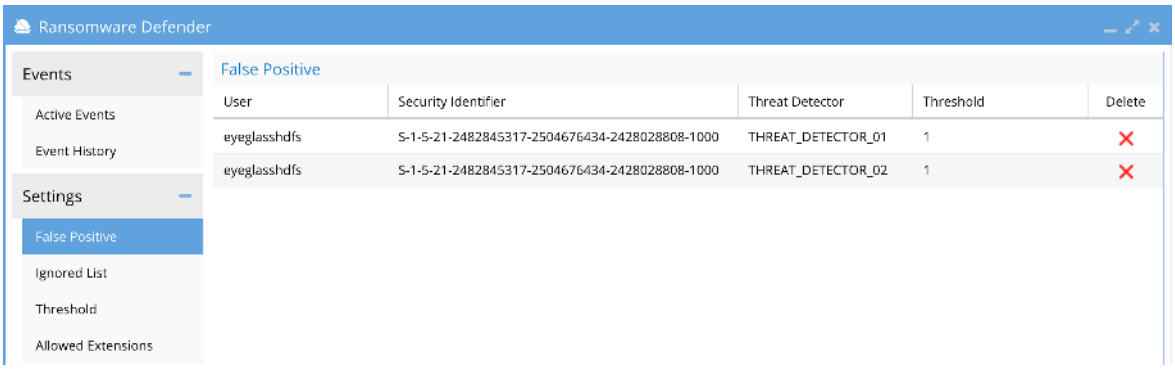

How to View or Delete a Flag as False positive user setting

If you accidentally flagged as false positive or want to undo a user override setting. Follow these steps.

1. Open the Ransomware Defender Icon
2. Click on Settings-->False Positives tab
3. Find the user setting in the list and click the delete red X to remove this setting.

a.

4. **NOTE: This change will take effect immediately and the ECA will be updated with the new settings for new events that are processed.**

How to manually configure per user Threat level settings with IGLS CLI

IGLS CLI commands exist to add and delete per user threat level override settings without waiting for a security event to teach.

Enter commands to create unique settings per user. This avoids the need for whitelisting users and can customize the settings per user. These settings are downloaded to the ECA cluster and processed in real-time once set as events flow through the cluster.

See Admin guide for complete documentation on the CLI commands

© Superna LLC

2.5.3. Banned and Allowed File Type Configuration

Home Top

- How to view a security event with a Threat Detector 7 Banned File detection
- How to View Allowed files Types (< 2.5.7)
  - How to view files on the Allowed Extensions list (< 2.5.7)
  - How to Ignore a File Extension using the Ignore whitelist (< 2.5.7)
- How to add Custom File extensions (2.5.7 or greater)
- How to Manage Banned File Extensions with Enforcement Modes (2.5.7 >)

How to view a security event with a Threat Detector 7 Banned File detection

Ransomware Defender has a dynamic list of 2000 or more known file extensions that are associated with Ransomware. This list is updated and imported from the Internet when the file is changed. Some environments use files on the banned extension list and may show up as a detection flagged as Threat Detector 7. You can view the Threat Detector type from the active events window by clicking the Threat detectors to view the detection type. When a file on the banned file

extension list is detected you may need to allow this file type in your environment by adding it to the whitelist. Follow the steps below to view which file type was matched to the banned file list.

1. Login to Eyeglass and open the Ransomware Defender Icon
2. Click Active Events
3. Click on the Threat detectors column of the active event you want to check.
4. If you see Threat Detector 7 listed continue below to identify which extension was found on the detection.
5. Click the actions menu
6. Scroll to the bottom of the event history and scroll up to locate the matching file type rule that trip the Threat Detector 7 banned file extension detection. See the screenshot example below.

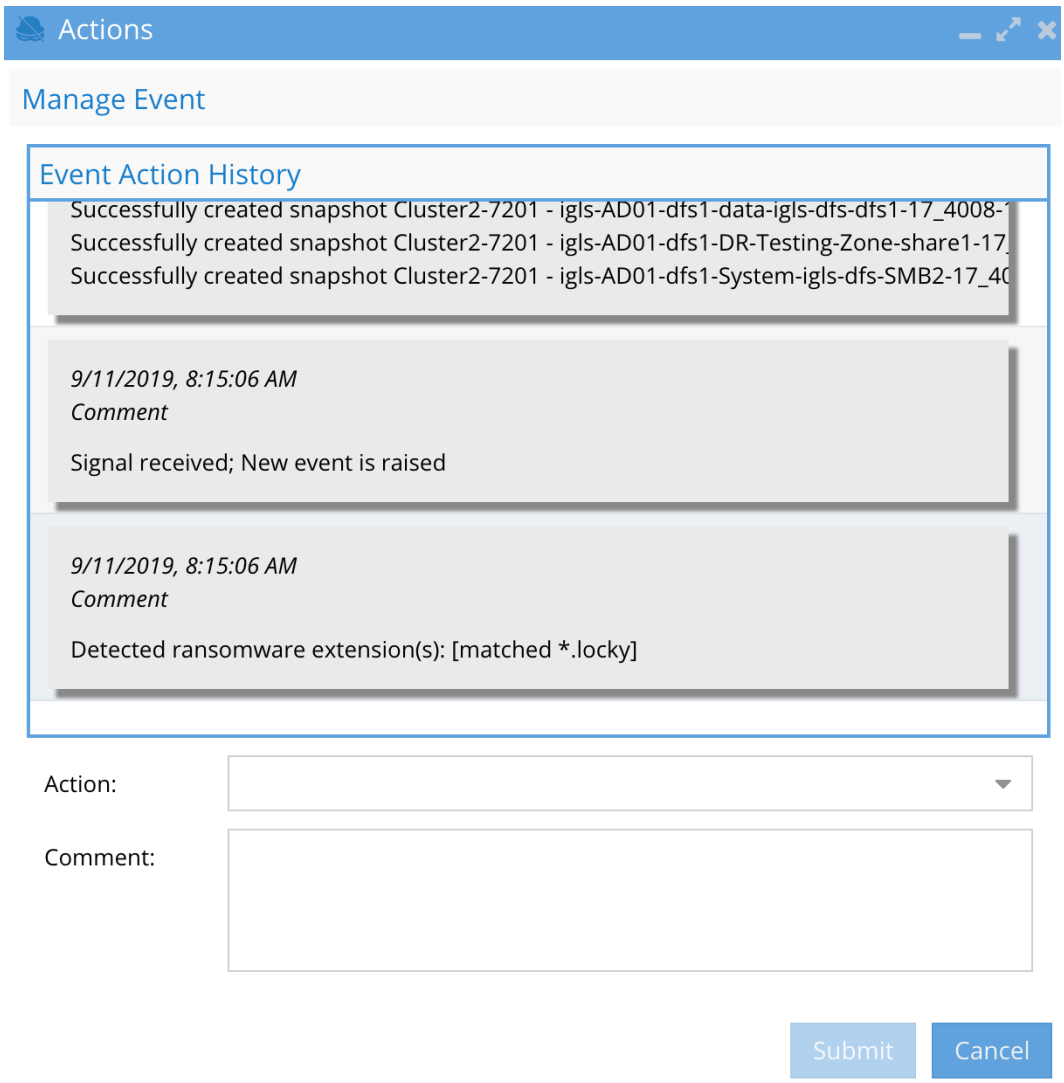

a.

b. In this example the file extension match was a file with an extension of .locky, this is not an extension that should ever be whitelisted and used for example only.

7. If the file extension is a legitimate file type used in your environment you will need to add this file to the allowed extensions list.

   a. Release 2.5.7 > Flag the event as false positive will place the file extension into the File Filters list that will set the extension to disabled status.

   b. Release < 2.5.7 igls rsw allowedfiles add --extensions='*.ext1' (note .ext1 is the extension found in

the Action history, change the file extension to match your detected file type)

8. See the next section to view the files on the allowed list.

How to View Allowed files Types (< 2.5.7)

1. Click on the Settings Allowed extensions tab
2. This will show any CLI enabled whitelisted file extensions or flag as false positive.
    a. A well-known list of Ransomware extensions is managed by Ransomware Defender and this list can sometimes conflict with files used within your environment. CLI commands place files on the allow list and this can be viewed in the GUI.
3. 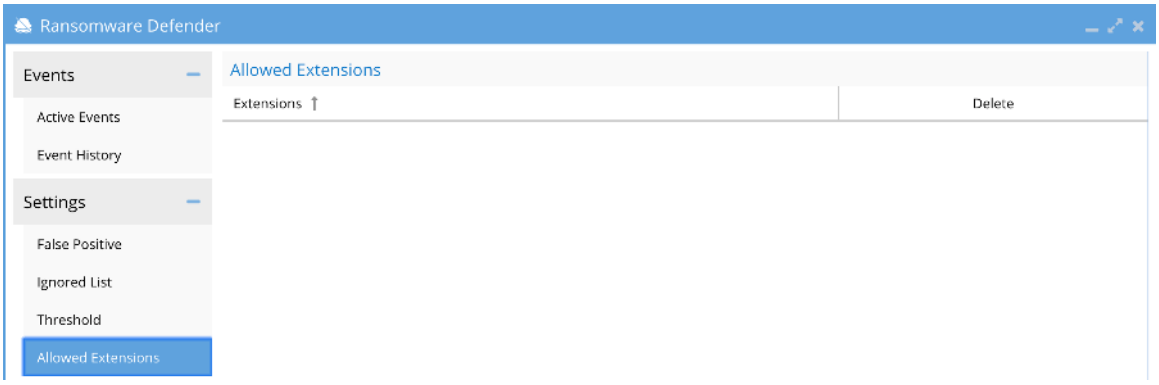

How to view files on the Allowed Extensions list (< 2.5.7)

This tab shows files that have been removed from the banned files master list, using the CLI commands to remove files from the master banned list or the flag as false positive option (release 2.5.7 or later). This interface also allows files on the allow list to be removed by deleting the entry in the table.

1. Open Ransomware Defender, Settings, Allowed Extensions
2. View the extensions that are on the allowed list, these files will not trip a detector if processed for users. These extensions are on the well-known master list of banned extensions. This allows an override to customize these extensions in your environment.
3. The delete button will remove the extension from the list and will be processed as a banned file once deleted. This change will take effect immediately for any new events processed.
4. 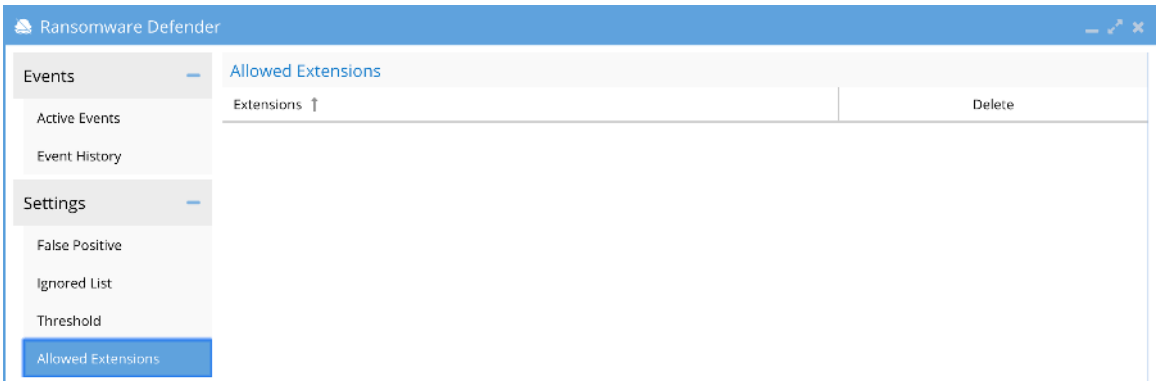

How to Ignore a File Extension using the Ignore whitelist (< 2.5.7)

This option allows using a new whitelist option to remove a file type from processing anywhere in the file system.

1. Open Ransomware Defender Icon, click on the Ignored List tab under settings.
2. A Path-based entry will be used. Click the plus sign, and select a cluster from the list.

3. Enter a path using the information below, and click submit to save once done.
4. Enter a whitelist of file type with extension .tmp /ifs/**/*.tmp
    a. This will whitelist any .tmp anywhere in the file system.

How to add Custom File extensions (2.5.7 or greater)

1. Open Ransomware Defender icon settings-->File Filters tab
2. Click the add filter button to enter a custom file extension. Enter the extension with *.xxx and set the mode to enabled (lockout), Disabled (no detection or lockout) or Monitor mode (detect, snapshot, no lockout). Click Add.
3. 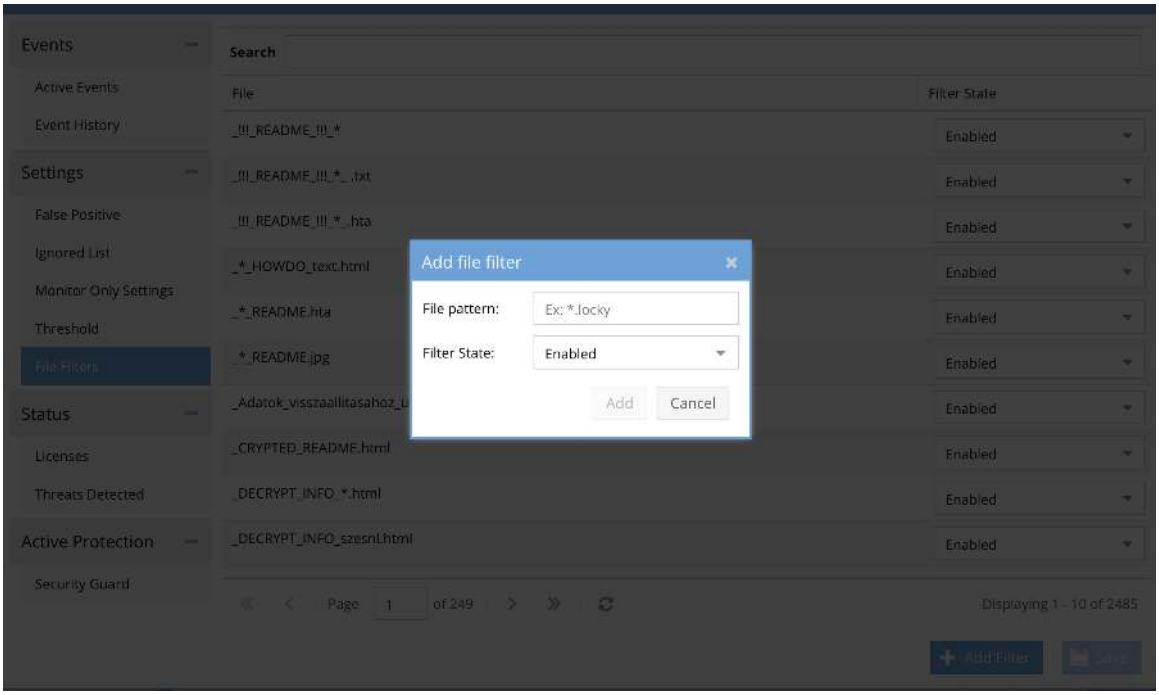

How to Manage Banned File Extensions with Enforcement Modes (2.5.7 >)

1. New in 2.5.7 the ability to search the banned file list and and set the mode on each extension to enabled, disabled or monitor mode.
2. **Modes**: enabled (lockout), Disabled (no detection or lockout) or Monitor mode (detect, snapshot, no lockout)
3. Filter based on builtin , custom extensions or both.
4. Filter file extensions by state using the filter box or filter based on builtin extensions or custom extensions.

a. 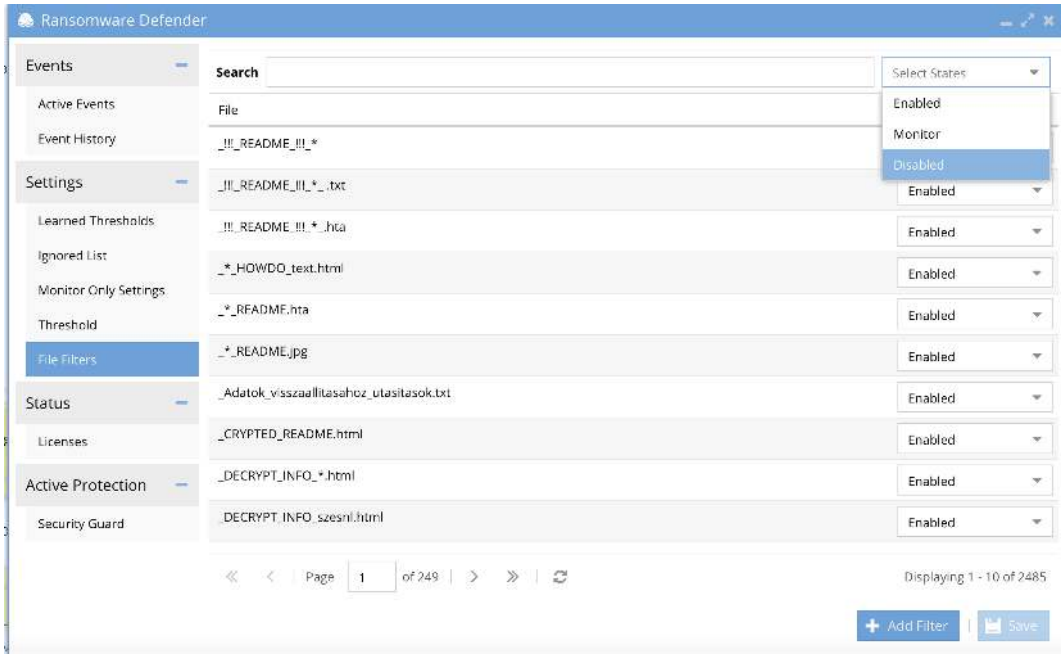

b. 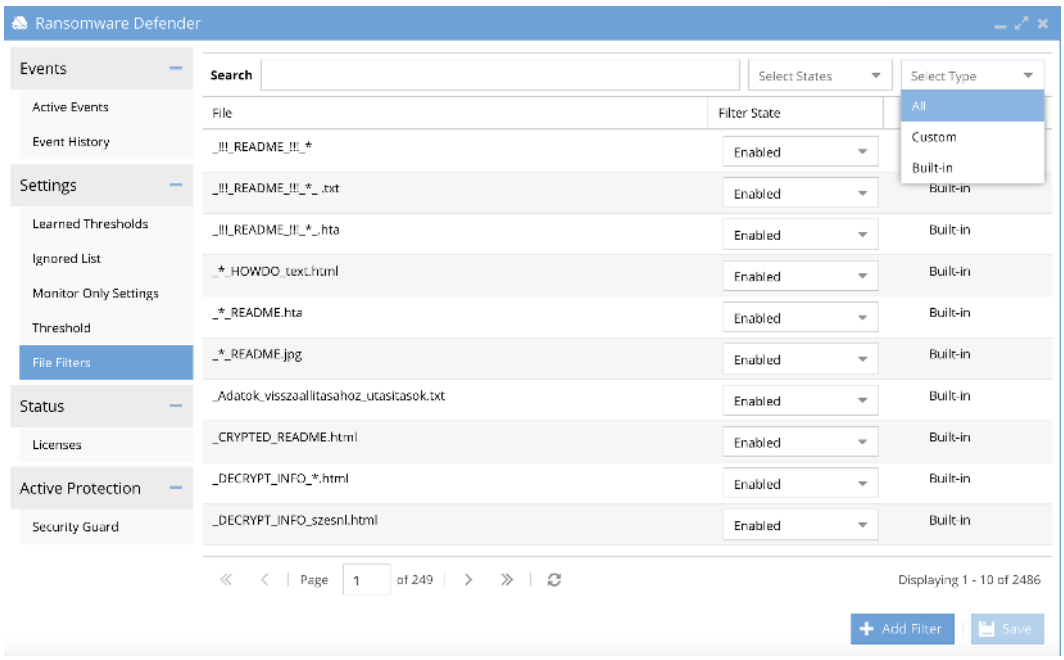

5. Search for the extensions by typing letters for the extension, set the mode and then click the save button. NOTE: Changes will take effect immediately.

6. 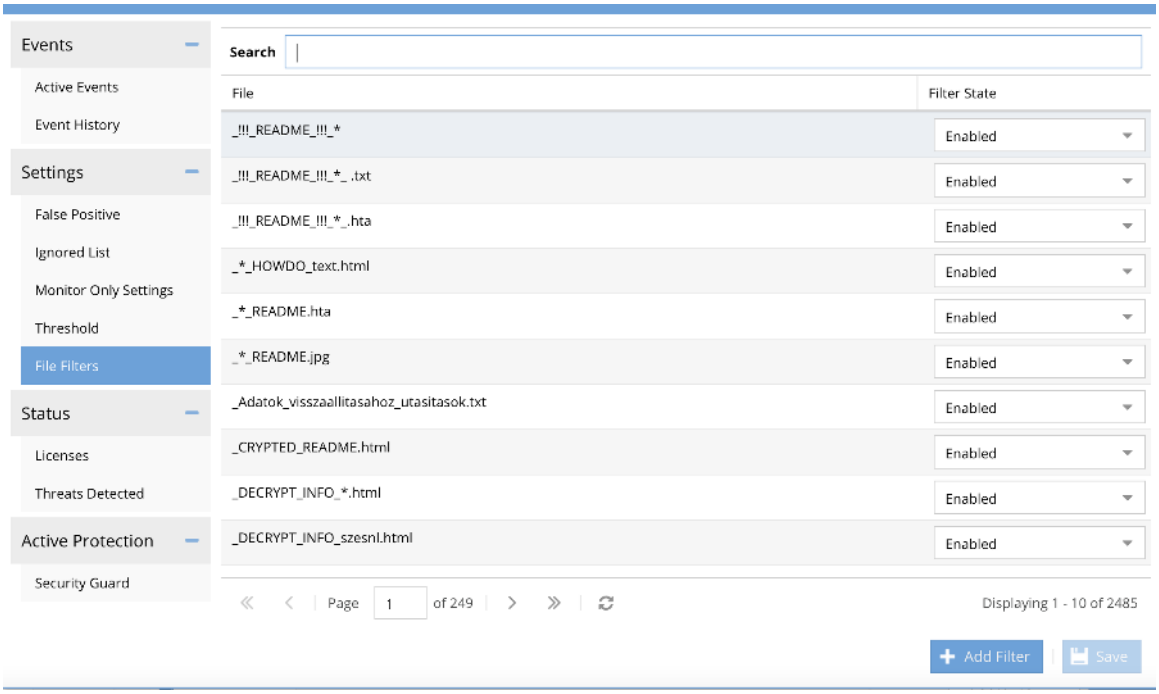

7. Search for an extension example locky to easily find an extension

8. 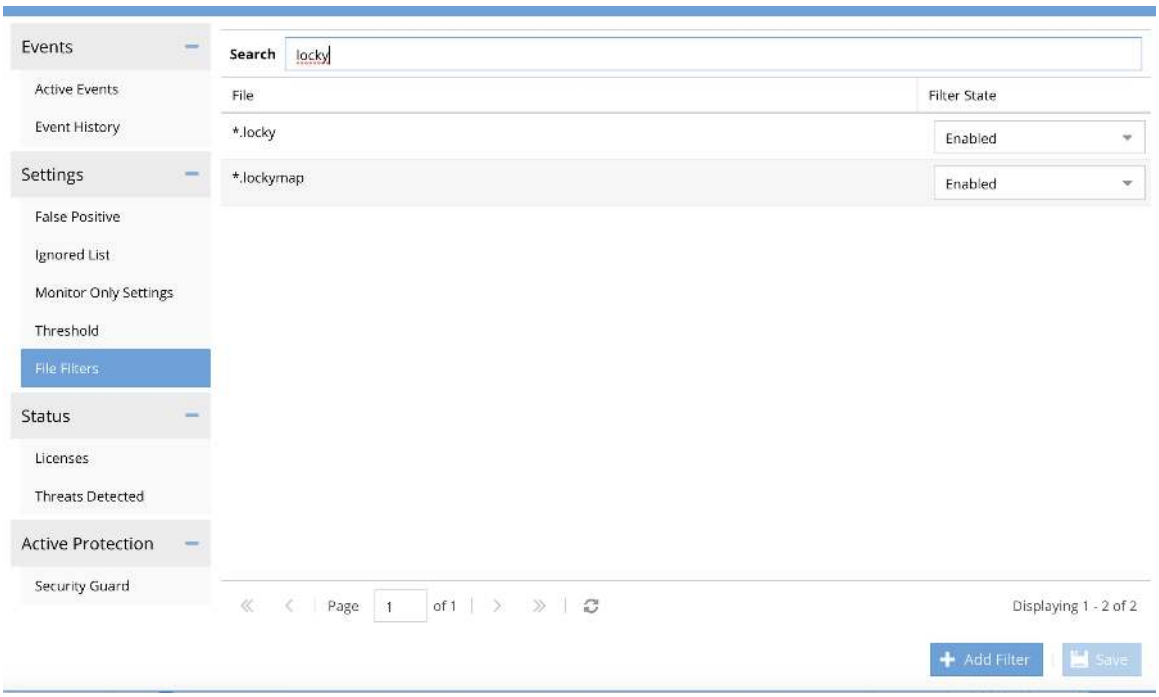

© Superna LLC

2.5.4. Rapid Machine to Machine Malware Spreading Attack Defense

Home Top

Overview

Ransomware Defender can use multiple cluster detections, to elevate the automated response due to the severity of the detection, and number of concurrent security events. Refer to the diagram below:

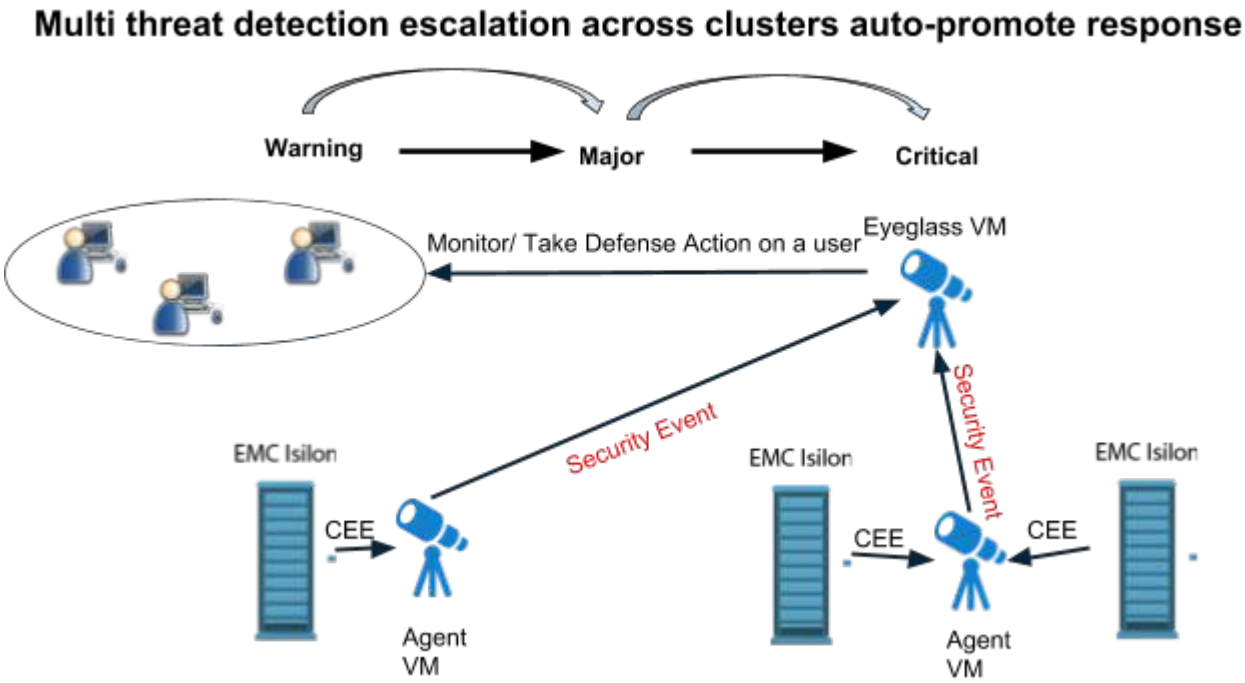

Rapid Machine to Machine Malware Attack Auto Response Escalation Configuration

This feature is designed to protect against a multi-user scenario where malware affects many machines in a short period of time, and when malware is spreading from machine to machine. The goal in this scenario is to escalate the response automatically based on the number of concurrent events. The example below walks through how

warning → major → critical response escalation will occur based on settings.

Best Practice: Set the Warning to **Major** to a higher number e.g. 5, and **Major** to **Critical** to half of the warning e.g. 8.

1. **Major** and **Upgrade to Critical** events are set to upgrade the severity to this level when a lower severity detection event matches or exceeds the number entered.
2. **Example (A)** if **Upgrade to Major (events)** is set to 8 this means if 8 separate Warning events are detected the response will be auto-upgraded to **Major** and timed lockout will be started. (see screenshot below)
3. **Example (B)** if **Upgrade to Critical (events)** is set to10 this means if 10 separate **Major** events are detected the response will be auto-upgraded to **Critical** and immediate lockout will be activated. (See screenshot below)

© Superna LLC

2.5.5. How to Manage Threat Detectors - Advanced Consult Support

Home Top

- Overview
- How to Disable all threat Detectors
- How to Customize each Threat Detector State

Overview

This is new advanced option in 2.5.7 or later releases. The Ransomware Defender solution is built on many threat detectors that detect user behavior, honeypot files and banned file extensions. This is an option that should never be changed without support advising when to use this option. Changing these settings can disable detection without understanding the impact of the changes.

How to Disable all threat Detectors

1. This option will disable all threat detectors and will disable all detections for all severities. **This option should never be used unless you need to temporarily need to disable detections.**
2. Login to Ransomware Defender
3. On the Settings --> Threshold tab
4. Click **Monitor All** and then the submit button. **NOTE: This disable all detections until you change back to Enable All Option and click submit button again.**

5. 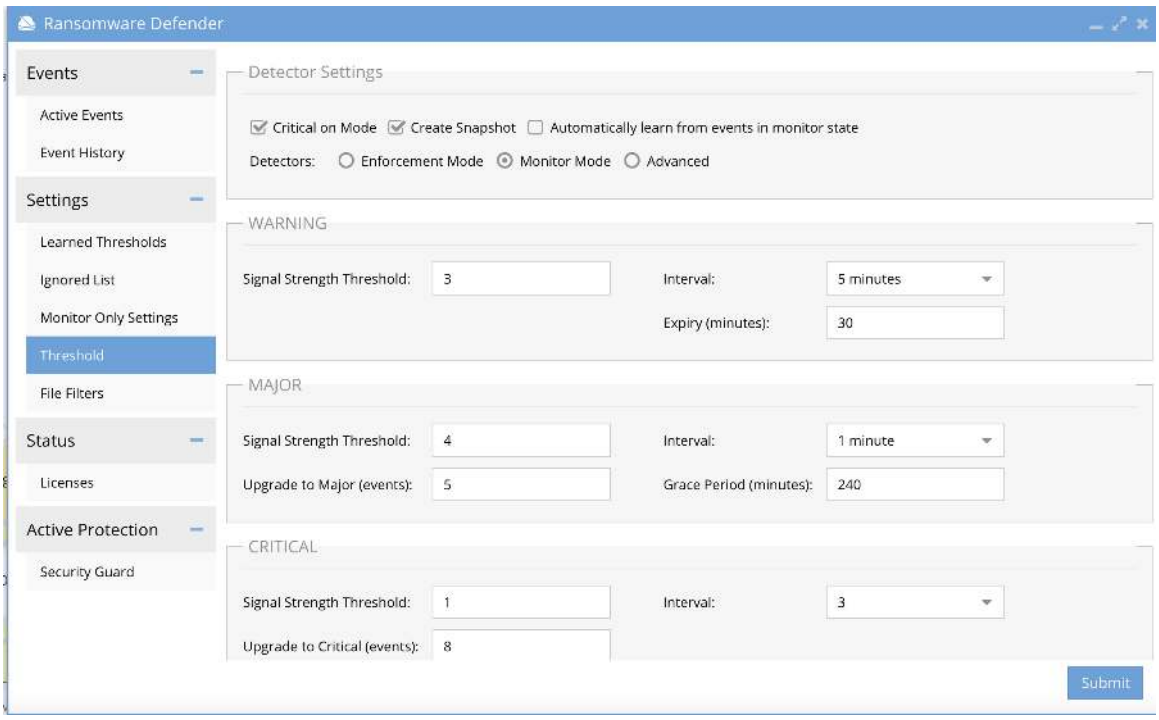

How to Customize each Threat Detector State

1. The customize option allows individual threat detectors to be disabled and switched into Monitor only mode. Monitor only mode allows the detector to alert, snapshot but it will not lockout. If the threat detector is disabled it will no longer protect the file system. **NOTE: These settings should not be changed without consulting with support, any misconfiguration can disable protection of your file system.**
2. **See the screenshot below on how to switch the mode of each threat detector.**

3. 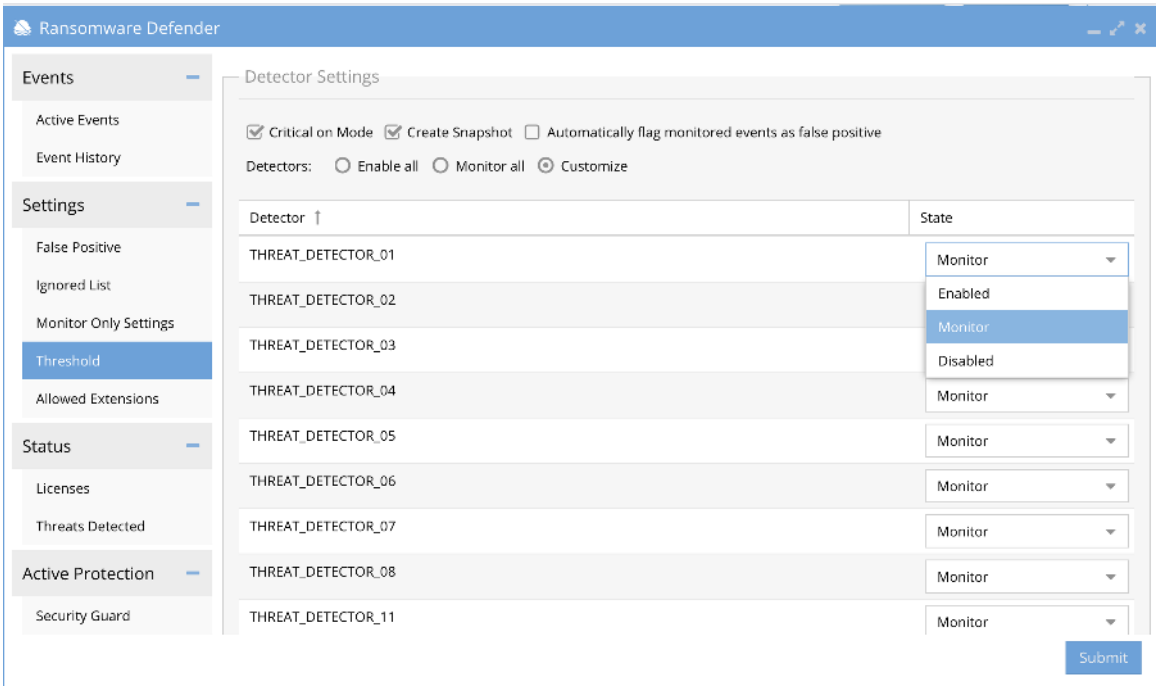

© Superna LLC

2.5.6. How to Configure Snapshot Modes (Critical Path and SMB share snapshots) and Snapshot Quotas

Home Top

- Overview
- Requirements
- Considerations
- Configuration Snapshot Modes and Snapshot Quota

Overview

This feature allows SMB user share snapshots to be disabled. This would be used when ACL security is used and most shares use everyone full control permissions allowing all users to have access to all shares. In this configuration, a lot of snapshots can be created for a single user detection. The other use case allows targeting snapshots on specific critical paths in the file system when any user detection occurs and disabling SMB share level snapshots. In addition snapshot quota allows specifying a limit on the number of snapshots that Ransomware Defender can create.

Requirements

1. Release 2.5.7.2 or later

Considerations

1. If user share snapshot mode is disabled , this also disables the event action menu create snapshot manual action on events.

Configuration Snapshot Modes and Snapshot Quota

1. Click the tab Snapshots
    a. unClick Enable Share Snapshots to disable snapshots applied to user SMB shares detected by AD group permissions
    b. Click Enable Critical Path Snapshots and then the + sign to add path to the list of paths that will have a snapshot applied on each and ever detection event. The snapshot will be created even if the user does not have access to this path. Use this option to protect application data or any critical data on the cluster. Add paths as needed.
    c. Change the snapshot quota value to a higher or lower number to set the limit. Once the snapshot limit is reached no more snapshots will be created until snapshots expiry which allows snapshots to be taken again up to the limit. Default snapshot expiry is 48 hours.
    d. NOTE: both share snapshot and critical path can be enabled independently.
    e. NOTE: Always click submit after making changes.

f. 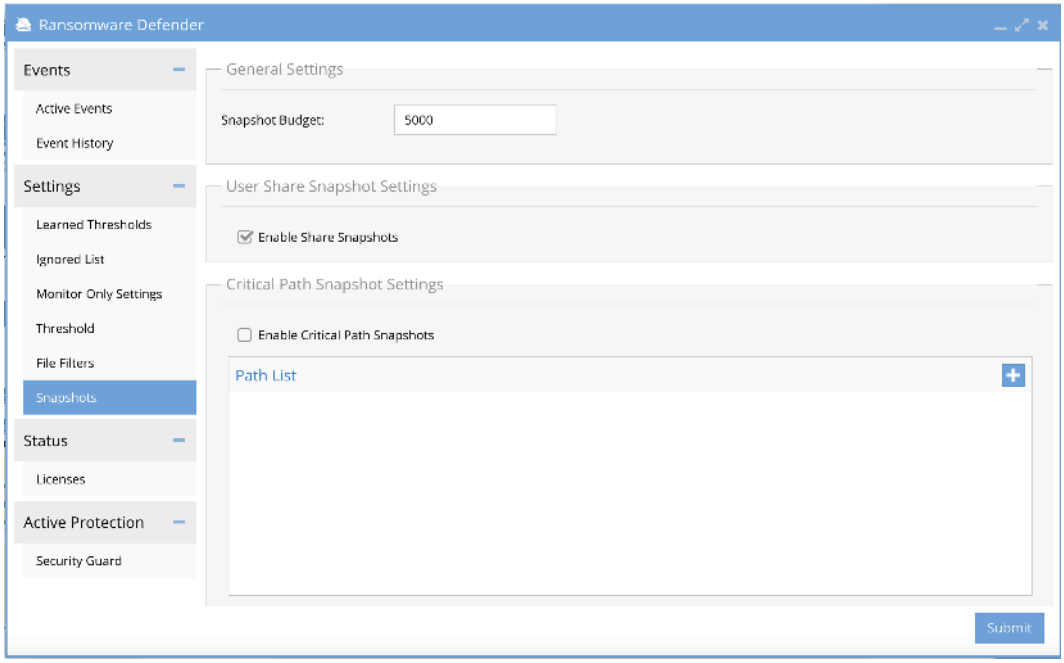

© Superna LLC

2.6. How to respond to Security Events for Warning, Major or Critical Events

Home Top

- Read Me First
- Security Event Triage Process to Collect Key Information for Your Security Incident Response Process

Read Me First

On detection of a possible ransomware event we can provide you with a relative severity of the detection and then provide with you information to make a decision such as files affected, user associated with event and client IP. This information should be used in conjunction with your security incident response process and security tools such as virus scanning and inspecting users computer to reach a conclusion whether or not it is a ransomware event. **A detection event should be evaluated based on the signal strength column and the severity column on an Event as additional information for you security response process.**

The information below should be used when following the process in this guide, get the signal strength and severity of the event and then review ranges below to determine the relative severity of the **User Behavior. Your Ransomware Defender settings determine when a user will be locked out based on the signal strength value and settings**

for Severity. If you are in Monitor Mode, no lockout will occur and manual lockout is required from the actions Menu.

Severity	Files	Signal Streng...	User	Detected ↓	Shares	Snapshots	Archived	Clients	Actions
CRITICAL	2+ files	2/2/2	AD02\sgdemo	8/3/2020, 3:0...	1 shares	1 clusters	8/3/2020, 3:1...	Clients IPs	

Quick Assessment

1. Use Severity and signal strength values to assess the event
2. **Severity** level is determined by the settings in your environment
   a. **Warning** - is considered **low** probability of infection
   b. **Major** - an **increased** probability of infection or a **medium** User Behavior
   c. **Critical** - is a **strong** Ransomware user behavior and an indication of a possible infection. A decision to lockout the user should be determined along with reviewing the files affected
3. **Signal Strength** - For any Severity level a higher signal strength count indicates the User Behavior is a continuous pattern. A continuous repeating detection pattern is consistent with Ransomware activity and should be used as an input to a response action based on the security incident response process.

Security Event Triage Process to Collect Key Information for Your Security Incident Response Process

When a security event has been detected, the steps to review and take actions should be followed exactly as listed below and then follow your internal security incicdent response process.

1. Open Ransomware Defender Active Events tab.
2. Review the files that tripped the detector and to download the CSV file associated to the event. Review these files.
    a. Each event shows the number of files associated with the active or archived security event. **NOTE: The list of files shown here are the files that tripped the detector and were stored with the event. A limit of 100 files is stored in the Eyeglass database, that populates the browser tree. See screen shot below**
    b.

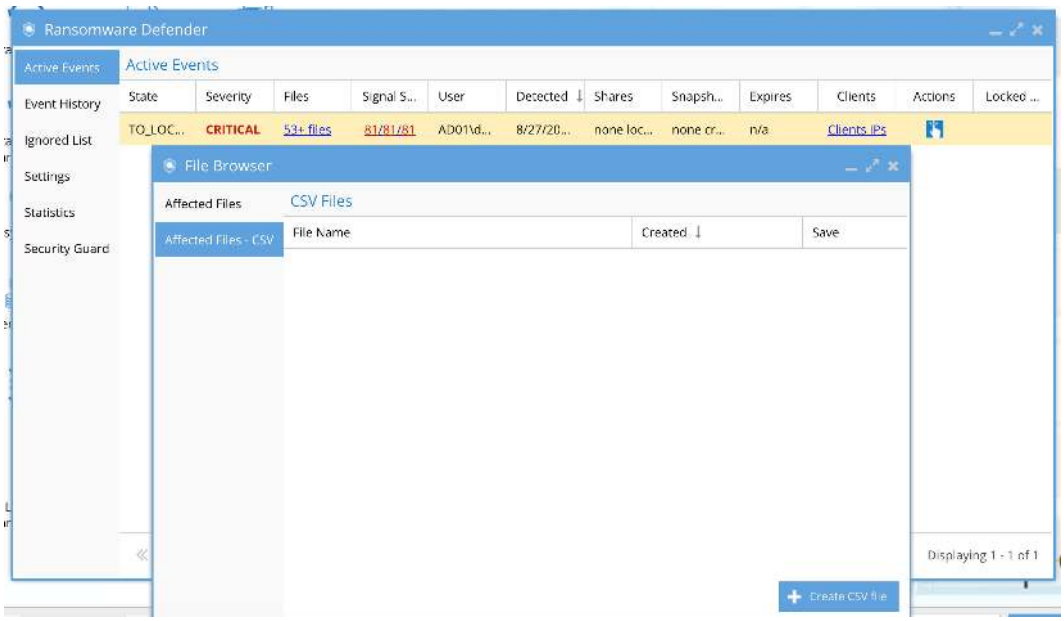

    c. The CSV download of affected files will list all sections for this user for all time, the date stamp in the file should be used to indicate when the detection occurred, the file name. NOTE: It is possible to have more files associated with the event than the CSV shows, if the event was active for more

than 1 hour. The list of files in the CSV represents only files deemed to have tripped a detector, and NOT files accessed or saved in this time period that did NOT trip a detector.

3. Review the list the following elements of the active event to review what occurred from the lockout:
    a. Click on the "Shares" link to review the shares with a lockout applied.
    b. Click on the "Snapshots" link to review the snapshots that were applied to shares.
    c. Click on the "Clients IP's" link to review the source IP of the subnet of the infected computer.
    d. Click the "Actions" menu to review the time stamps of each action applied to this event including share lockout, snapshots. The Action menu also provides a menu of actions that will be used based on the Warning, Major and Critical event severity, to make a decision on next steps.
    e. Sample event screenshot
        i. 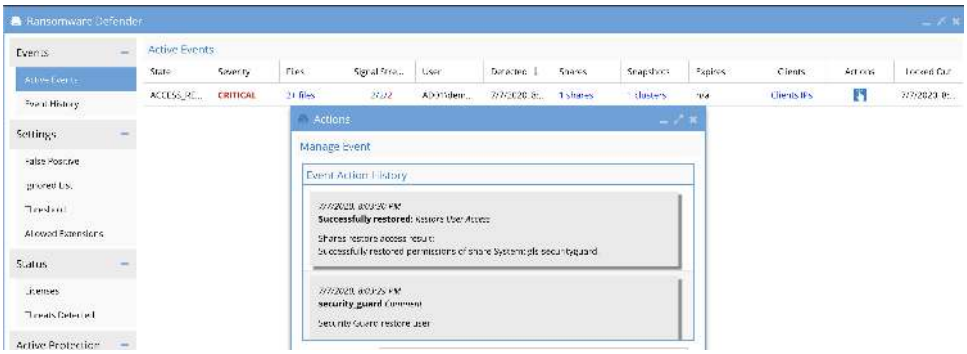

4. Review the severity (Warning, Major and Critical) response steps in the topics below. To evaluate which files tripped the security

detectors use the steps in this section to understand how the files download and browse window assists.

- If Warning:
- If Major:
- If Critical:
- Event State Descriptions

© Superna LLC

2.6.1. If Warning:

Home Top

If Warning:

1. Review the list of files affected for user account IP address by selecting the link in the Files column.
2. The file list view shows files that triggered the security event. The last hour of files accessed by the user is shown, and should be reviewed for possible compromise or data recovery.

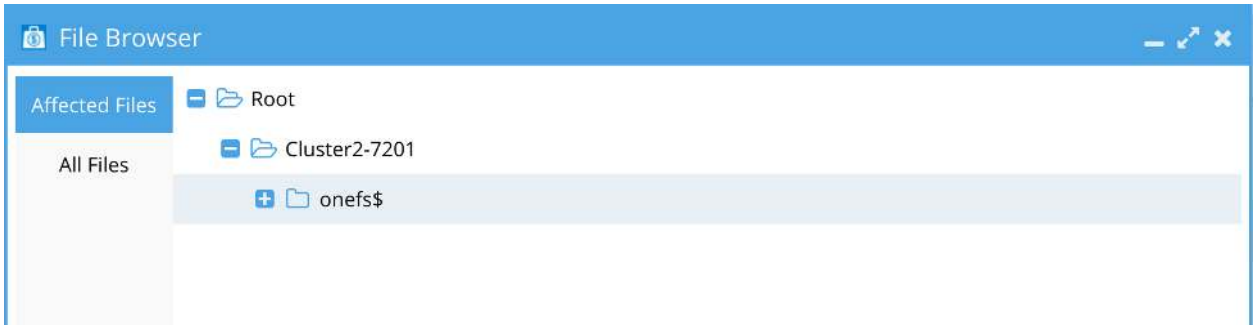

3. If the affected files are the result of normal file operations and not a malicious event, the event can be marked as resolved with the actions menu.(See Security Event Action State Descriptions section below).
4. If you need to flag as false-positive see instructions here on Flag As False Positive.
5. Security event closed and moved to the Event History tab.

© Superna LLC

2.6.2. If Major:

Home Top

If Major:

1. Review affected files, user name, and IP address to locate user in AD and your organization.
2. Review time to lockout timer in the Active Events tab which is the time until the lockout will be issued.
    a. If you determine this is a false alarm by contacting the user along with an assessment of the affected files, use the Action Menu to Stop the Lockout timer and then mark security event as **Resolved** (See Security Event Action State Descriptions section below).
3. If you determine it is a malicious security event, you can accelerate the lockout timer by using the Action menu to select **Lockout Now**. (See Security Event Action State Descriptions section below).
4. **Recovery**: Re-image machine or other recovery procedures that your policies require. Determine which files are to be recovered on PowerScale by selecting the files option on the security event. From this screen, you can download a CSV file of trigger files AND files from the last 1 hour of activity.
5. **Restore User Access:** Take this step after it has been determined it is safe to restore access to the user. The actions menu can be used to remove the user account lockout for all cluster shares to which that user had access. Using the Actions menu restore user access. (See Security Event Action State Descriptions section below). Click here for instructions.

6. The security event will now be in **Restored** state and can be archived to the Event History tab. Using the actions menu, submit a Mark As Resolved action. (See Security Event Action State Descriptions section below).
7. If you need to flag as false-positive see instructions here on Flag As False Positive.
8. Done.

© Superna LLC

2.6.3. If Critical:

Home Top

If Critical:

1. The security event will have a lockout applied immediately since it is a critical detection.
2. **Recovery**: Reimage machine or other recovery procedures that your policies require. Determine which files to be recovered on the PowerScale by selecting the files option on the security event. From this screen, you can download a CSV file of trigger files AND files from the last 1 hour of activity.
3. **Restore User Access:** After it has been determined it is safe to restore access to the user. The actions menu can be used to remove the user account lockout for all cluster shares to which that user had access. Using the Actions menu restore user access. (See Security Event Action State Descriptions section below). **Click here for instructions.**
4. The security event will now be in Restored state and can be archived to the Event History tab. Using the actions menu submit a **Mark As Recovered** action (See Security Event Action State Descriptions section below).
5. If you need to flag as false-positive see instructions here on Flag As False Positive.
6. Done.

© Superna LLC

2.6.4. Event State Descriptions

Home Top

Security Event Action State Descriptions

Once a user security event appears in the Active Events tab the following operations are possible by clicking the **Actions icon.** Each state has several possible actions. The table below describes the options available for each state of a security event.

State of Event	Possible Actions
Warning State 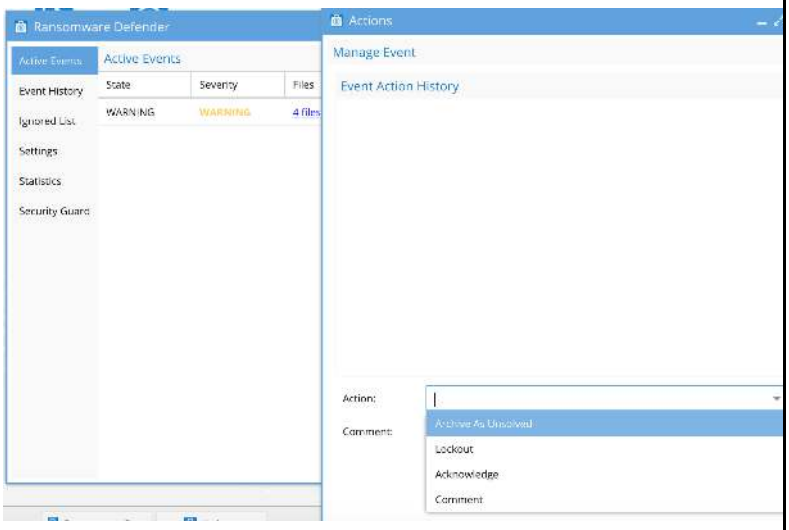 	1. **Comment on the event to update the security response or assessment of the event. Can be viewed by other administrators that review the security event history.** a. **Archive as Unsolved -** Moves event to the History tab. b. **Lockout -** From the Access Restored state it's possible to re-lockout the user again from the action menu. This applies deny permission to

	all shares stored within the lockout event. c. **Acknowledged State** - An administrator has acknowledged this event but has not marked as resolved. In this state the user is **not** locked out or in timed lockout states. d. **Create Snapshot -** Manual snapshot created on all share paths in the security event. e. **Delete Snapshot** - Manual snapshot deleted on all share paths in the security event.
Locked out User State (Critical Severity Threat Detection)	1. **Comment** on the event to update the security response or assessment of the event. Can be viewed by other administrators that review the security

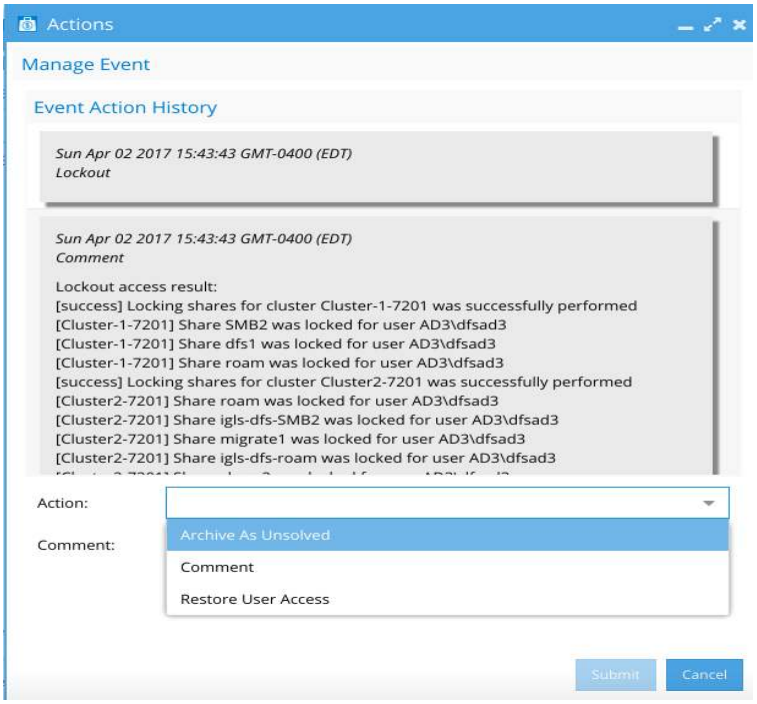

event history.

2. **Restore User Access - This will reverse the lockout and grant access to the shares that were locked out. Review the lockout details for a full list of shares and clusters that lockout was applied.**

1. **Once Restore User Access is launched, this will start a restore access job (running jobs window) and real-time restore access to the share that was last locked out.**
2. Verify the user has access
3. Verify a cluster share to confirm that the restore access was successful

3. **Archive as Unsolved - Leaves the lockout applied and moves the event to the History tab. Not recommended unless the user access is permanently revoked.**
4. **Create Snapshot -** Manual snapshot created on all share

	paths in the security event. 5. **Delete Snapshot -** Manual snapshot deleted on all share paths in the security event.
Access Restored State 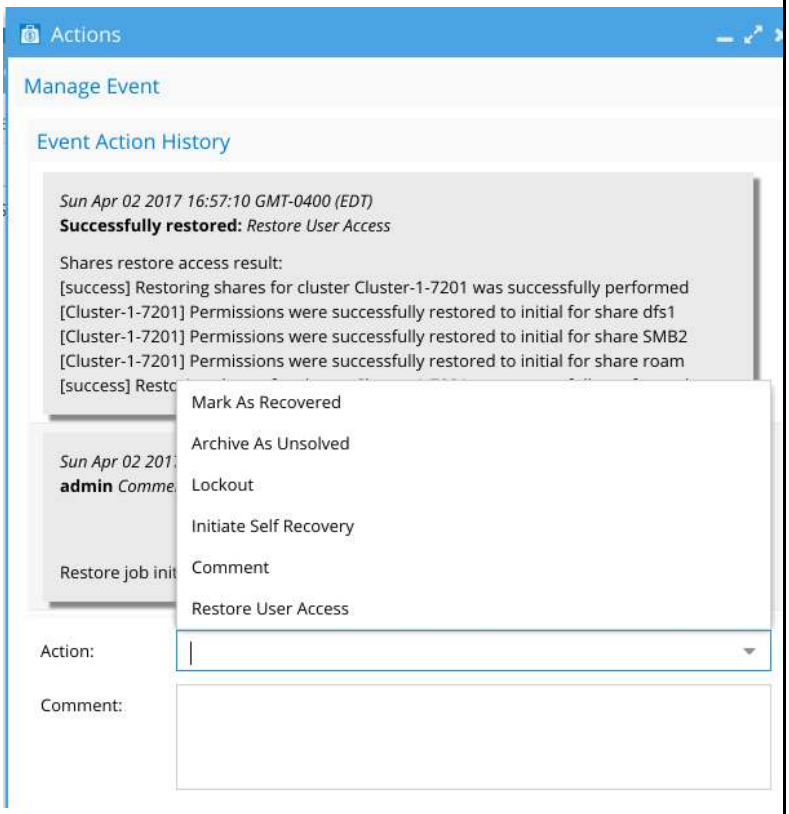 	1. **Mark as Recovered - This option allows archiving the security event to the history tab.** 2. **Lockout - From the Access Restored state it's possible to re-lockout the user again from the action menu. This applies deny permission to all shares stored within the lockout event.** 3. **Initiate Self Recovery - This option will only function if the Cluster Storage Monitor add-on is purchased. It integrates with the Backup Recovery User portal to create secured shares to snapshots and DR data that allow the user to recover data from snapshots. The temporary shares**

	will have a 2 day lifetime by default, after which they will be deleted. The shares are secured only to the user involved in the lockout. The data recovery request will require approval in the Data Recovery Manager Icon. See the Data Recovery section in this guide. (If licensed) 4. Comment - on the event to update the security response or assessment of the event. Can be viewed by other administrators that review the security event history. 5. Restore User Access - (Allows to re-run this job in the event a share or update failed) This will reverse the lockout and grant access to the shares that were locked out. Review the lockout details for a full list of shares and clusters that lockout was applied. 1. Once Restore User Access is launched, this will start a

	restore access job (running jobs window) and real-time restore access to the share last that was locked out. 2. Verify that the user has access 3. Verify a cluster share to confirm that restore access was successful 6. **Archive as Unsolved** - Leaves the lockout applied and moves the event to the History tab. **Note: recommended unless user access is permanently revoked.** 7. **Create Snapshot -** Manual snapshot created on all share paths in the security event. 8. **Delete Snapshot** - Manual snapshot deleted on all share paths in the security event.
Delayed Lockout state	1. **Lockout** - From the Access Restored state it's possible to re-lockout the user again from the action menu. This applies deny permission to all shares stored

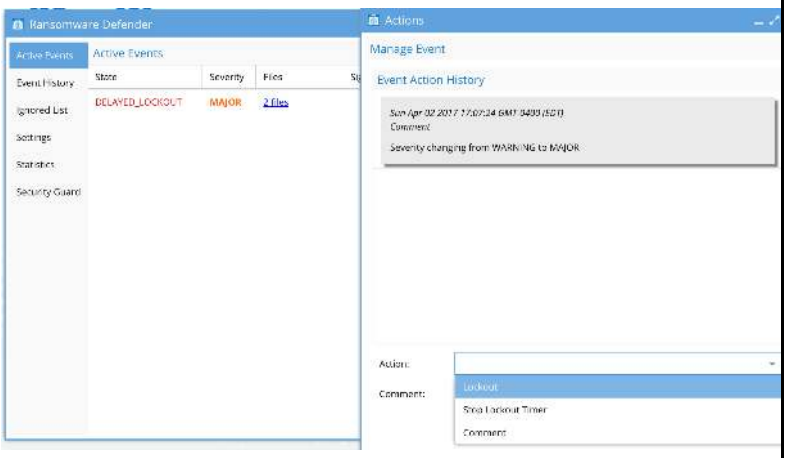	within the lockout event. 2. **Stop Lockout Timer-** This option can be used to stop the timed lockout. This would be used when the investigation determines the user account should not be locked out. 3. The status changes to **Acknowledged a nd the lockout will stop.** 4. **Comment** on the event to update the security response or assessment of the event. Can be viewed by other administrators that review the security event history. 5. **Create Snapshot -** Manual snapshot created on all share paths in the security event. 6. **Delete Snapshot** - Manual snapshot deleted on all share paths in the security event.
Acknowledged State	1. **Comment on the event to update the security response or assessment of the event. Can be**

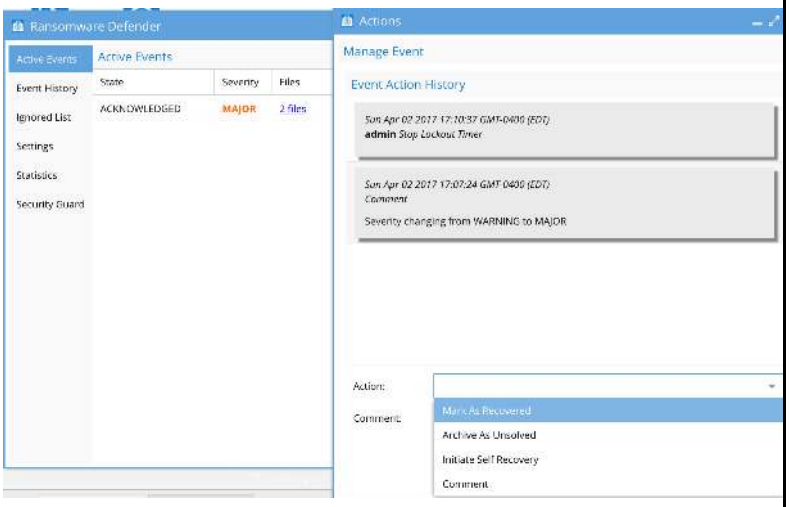	**viewed by other administrators that review the security event history.** 2. **Archive as Unsolved** - Leaves the lockout applied and moves the event to the History tab. **Note: recommended unless user access is permanently revoked.** 3. **Mark as Recovered -** This option allows archiving the security event to the history tab. 4. **Create Snapshot -** Manual snapshot created on all share paths in the security event. 5. **Delete Snapshot** - Manual snapshot deleted on all share paths in the security event.
Archived Event on Event History	1. **Comment** on the event to update the security response or assessment of the event. Can be viewed by other administrators that review the security event history. 2. **Create Snapshot -** Manual snapshot

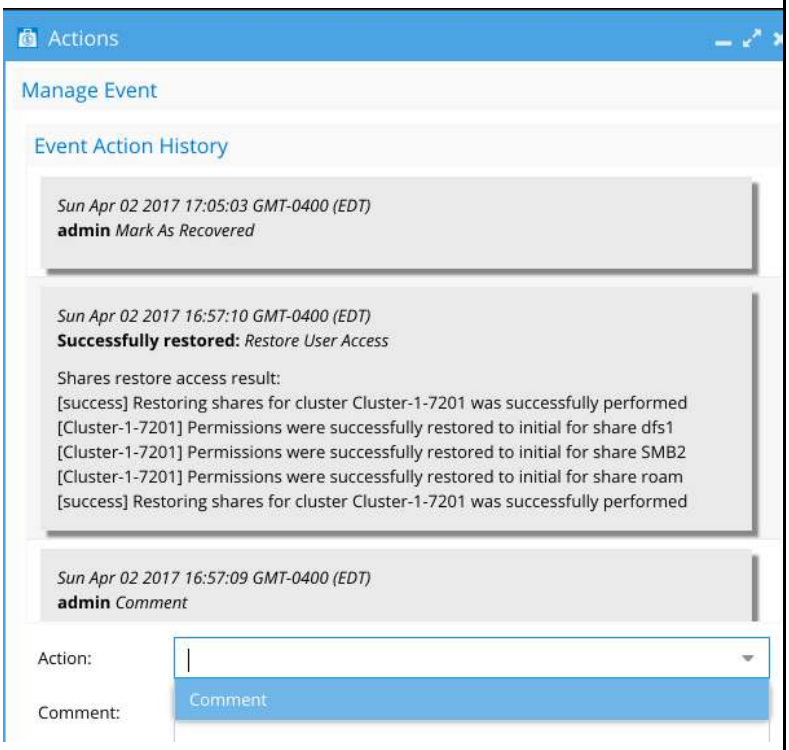

created on all share paths in the security event.

3. **Delete Snapshot** - Manual snapshot deleted on all share paths in the security event.

Error State

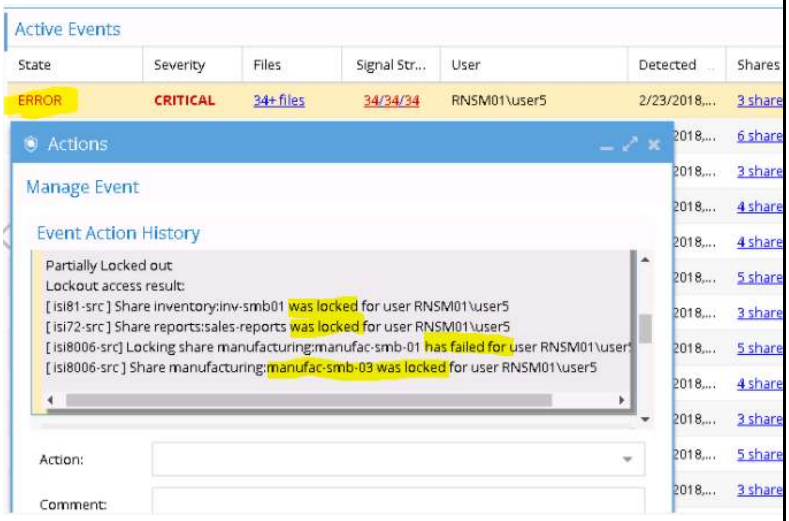

1. **Open the Action to see the Event Action History** Here you will see which shares had an issue in Lockout or Restore and reason
2. **Lockout -** If **Lockout** Error is related to an AEC_CONFLICT, then select the Lockout action again to re-attempt to complete the Lockout.
3. **Restore** - If **Restore** Error is related to an AEC_CONFLICT, then select the Restore action again to re-attempt to

	complete the Restore.

© Superna LLC

2.7. Operational Procedures For Common Tasks

Home Top

- Overview

Overview

This section covers daily operational actions that would be used for detections that occur during normal operations when no Ransomware attack is present in the environment.

- How to login and Manage Ransomware Defender
- Security Guard - Automated Security Testing
- How To Configure HoneyPot file Tripwire
- How to Flag a detection as False Positive
- How to Enable or Disable Enforcement Mode
- How to unlock a User that was locked out
- Ransomware Defender Security Event Workflow for Warning Severity Detections
- How to Recover Data after a Ransomware Attack
- How to Enable Learning Mode and Monitor Learning mode Results

© Superna LLC

2.7.1. How to login and Manage Ransomware Defender

Home Top

How to login and Manage Ransomware Defender

A built-in role and user account exist to separate the management of Ransomware settings and event monitoring.

1. Assigned the built-in role Ransomware Defender with the ability to manage and monitor Ransomware Defender product
2. Default password 3y3gl4ss
3. 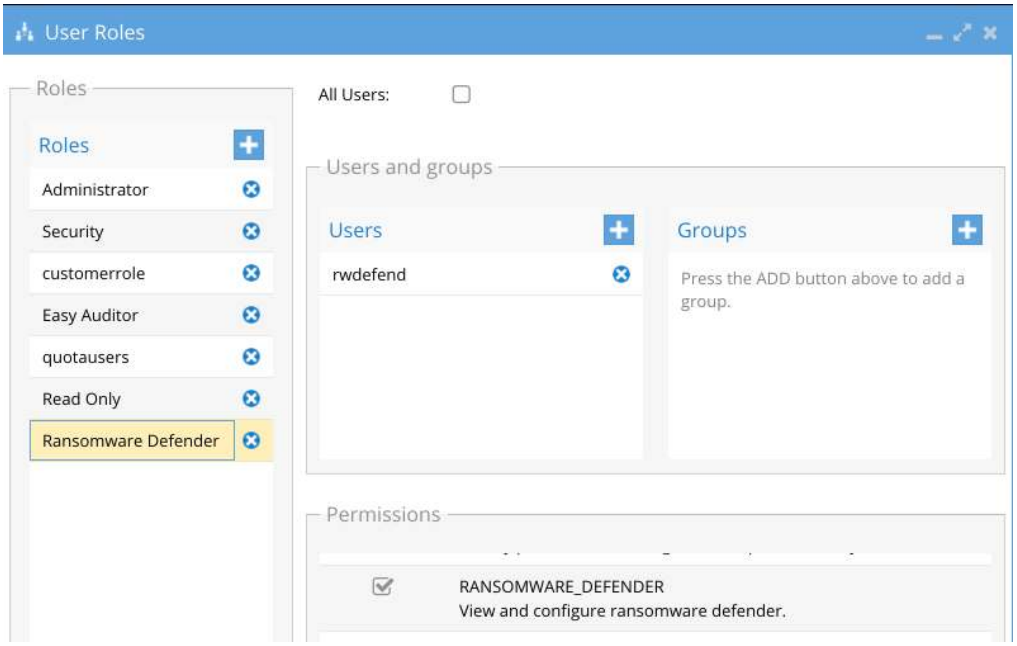

How to login

1. Login to Eyeglass appliance and enter either admin or rwdefend user with default password
2. Click on Ransomware Defender Icon

© Superna LLC

2.7.2. Security Guard - Automated Security Testing

Home Top

- Overview
- Simulated Attack
- Prerequisites (Isilon Powerscale and ECS)
  - Security Guard Lockout Behavior (Isilon Powerscale)
- Configuration (Isilon PowerScale)
  - How to Run on Demand Security Guard Penetration test (Isilon Powerscale)
  - How to Review Security Guard Penetration test history and logs (Isilon Powerscale)
- How to test Ransomware Defender with your own Custom File Extension (Isilon Powerscale)
- Configuration (ECS)
  - Security Guard Lockout Behavior (ECS)
  - How to Run on Demand Security Guard Penetration test (ECS)
- Advanced Configuration Security Guard CLI Commands (Isilon Powerscale)

Overview

Ransomware Defender monitors cluster IO for suspicious user behavior. Under normal day to day conditions, no actions are required since alerts are sent in the event of a Warning, Major or Critical security event.

The Security Guard feature simulates a Ransomware attack on a daily basis to validate all components are functioning, including alerting and lockout of user sessions. Once configured administrators get daily updates that Ransomware Defender is actively monitoring and responding to Ransomware events.

This offers you the highest level of confidence that your environment is ready in the event a malicious virus is inside your network and finds shares to attack data.

The feature will create a “honeypot share with name igls-securityguard” in the System Zone of each cluster managed by a Ransomware agent license key. The feature can simulate an attack on demand, or on a scheduled interval.

Simulated Attack

1. Creates share automatically secured to the service account.
2. Share name igls-securityguard. (Isilon)
3. ECS Bucket name xxx (ECS)
4. Cleans up old files from the last execution.
5. Creates test files using a well-known extension to trigger a simulated attack response from Ransomware Defender Clustered agent.
6. Verifies the user lockout occurs by checking that files cannot be written to the share.
7. Initiates the recovery of the user and verifies access to the share again.

8. Reports success and failure per step.
9. Emails administrator results.

Prerequisites (Isilon Powerscale and ECS)

1. Service Account Test User (Isilon Powerscale)
    a. A local PowerScale user created in the system zone local provider **example igls-securityguard**
    b. **Use an Active Directory service user only if multiple clusters are licensed for security guard. Best practise for a single cluster is a local account.**
    c. System Zone must be enabled in the audit configuration on the PowerScale cluster.
2. Security guard Service Account Test User (ECS)
    a. Create a bucket object user
    b. 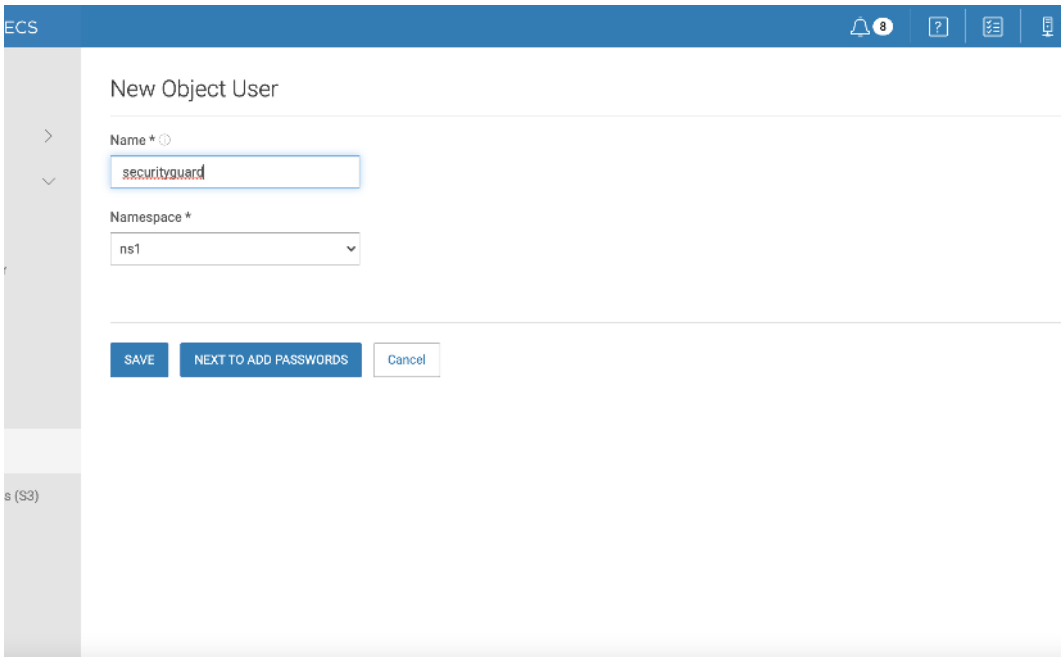

    c. Click next to add passwords
    d. Use the Generate and add keys button

e. Record the secret key to enter into the security guard configuration.

3. Repeat the steps above to create the **Bucket version service account user** that is used to enable bucket versioning to protect buckets that are under an attack. Bucket versioning will protect objects using version feature on the ECS cluster.

   a. Create the bucket version user eyeglassversions

   b. save the secret key to update to the ECS cluster.

   c. Open the Inventory Icon

   d. Right click the ECS cluster

   e. click add to and fill in the name space, user and secret key and click Submit to save. See the bucket version configuration inn the ECS section of the guide.

   f. 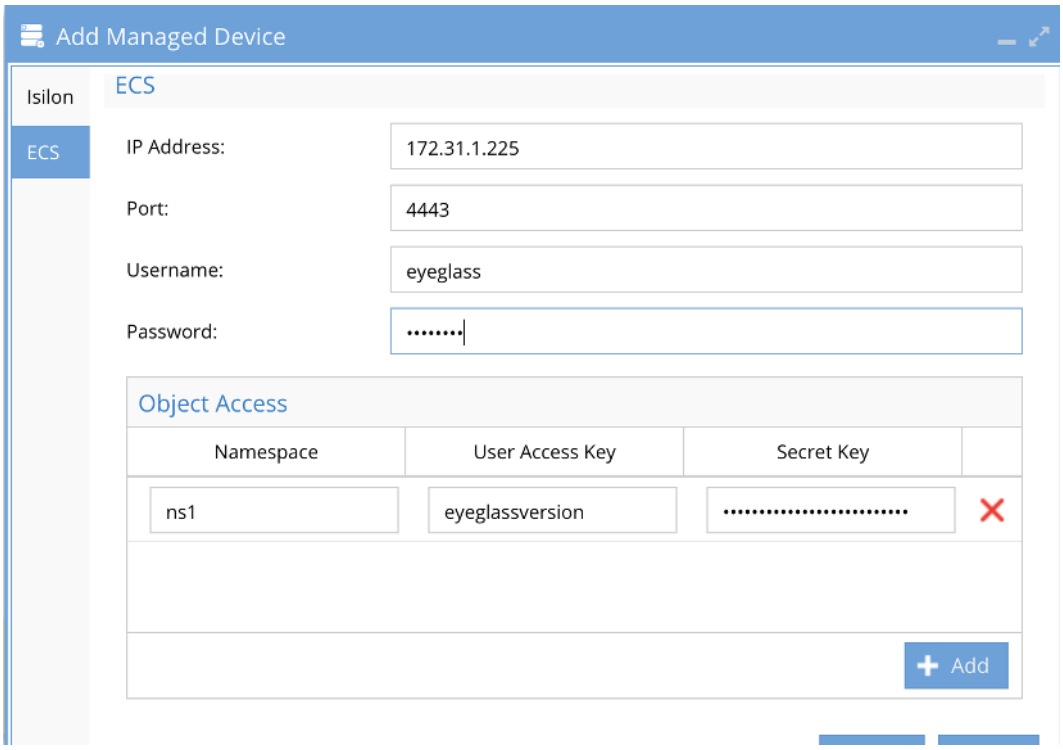

Security Guard Lockout Behavior (Isilon Powerscale)

1. The user does not need to be added to any shares. The Security Guard will create its own share in System Zone called **igls-securityguard** , and add the service account user to the share.
2. If you add the service account user to other shares, only the igls-securityguard share will have files written during the execution of a simulated attack.
3. Additional shares that have the service account add to the share permissions **WILL** have the service account access locked out during simulated attacks.

Configuration (Isilon PowerScale)

1. Open the Ransomware Defender window on the desktop and select the Security Guard.

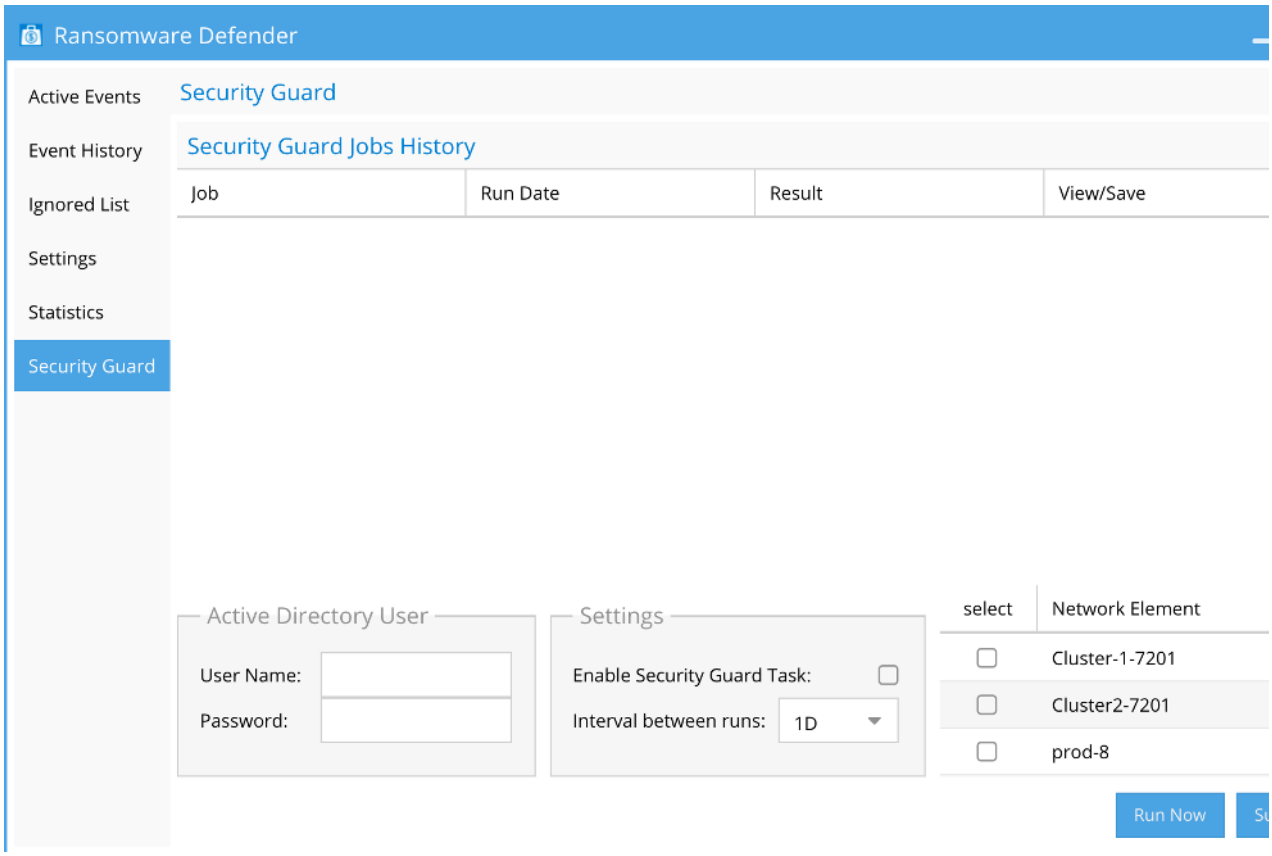

2. For local PowerScale user enter username@clustername

   a. NOTE: for multi cluster set up with AD user enter the user with user@domain.com (replace with your AD domain)

1. **Settings:**
   a. Enable Security Guard Tasks.
   b. Interval Between Runs - Set interval to schedule simulated attacks.
4. Select the checkbox of each cluster to simulate the attack.
5. Submit - Saves settings.
6. Run Now - Tests Security guard on demand.

How to Run on Demand Security Guard Penetration test (Isilon Powerscale)

1. Open the Ransomware Defender window (see screenshot below).
2. Select Security Guard tab.
3. Select each licensed cluster to test.
4. Select Run Now (see screenshot below).

Ransomware Defender

Active Events
Event History
Ignored List
Settings
Statistics
Security Guard

Security Guard

Security Guard Jobs History

Job	Run Date	Result	View/Save
Security_Guard	2017-03-31_21-32-22	SUCCESS	Open

Active Directory User

User Name: demosg@ad1.test

Password: ••••••••

Settings

Enable Security Guard Task: ☑

Interval between runs: 1D

select	Network Element
☑	Cluster-1-7201
☐	Cluster2-7201
☐	prod-8

Run Now Submit

5. Open Jobs window.
6. Running Jobs tab to monitor progress (see screenshot below).

Jobs

Job Definitions
Running Jobs

Running Jobs

State	Job Name	Started	Finished	Duration	Status
	Configuration Replication 149...	3/31/2017, 5:32:22 PM		0m 0s	RUNNING
	Security Guard 1490995942029	3/31/2017, 5:32:22 PM		0m 0s	RUNNING
✓	Configuration Replication 149...	3/31/2017, 5:30:00 PM	3/31/2017, 5:31:10 PM	1m 10s	FINISHED
✓	Access Zone Readiness 14909...	3/31/2017, 5:30:00 PM	3/31/2017, 5:30:10 PM	0m 10s	FINISHED
✓	Configuration Replication 149...	3/31/2017, 5:25:00 PM	3/31/2017, 5:26:12 PM	1m 12s	FINISHED
✓	Restoring multicluster shares....	3/31/2017, 5:24:30 PM	3/31/2017, 5:24:31 PM	0m 0s	FINISHED

Job Details

State	Job Name	Info
	Security Guard 1490995942029	
	Ransomware simulation on Cluster-1-7201	
✓	Check NE state	
✓	Check or create igls-honeypot share	
	Run Configuration Replication Now	
	Create test files in igls-honeypot share	

How to Review Security Guard Penetration test history and logs (Isilon Powerscale)

1. Open the Ransomware Defender window.
2. Select Security Guard tab.
3. Select each licensed cluster to test
4. Select Run Now (see screenshot below).

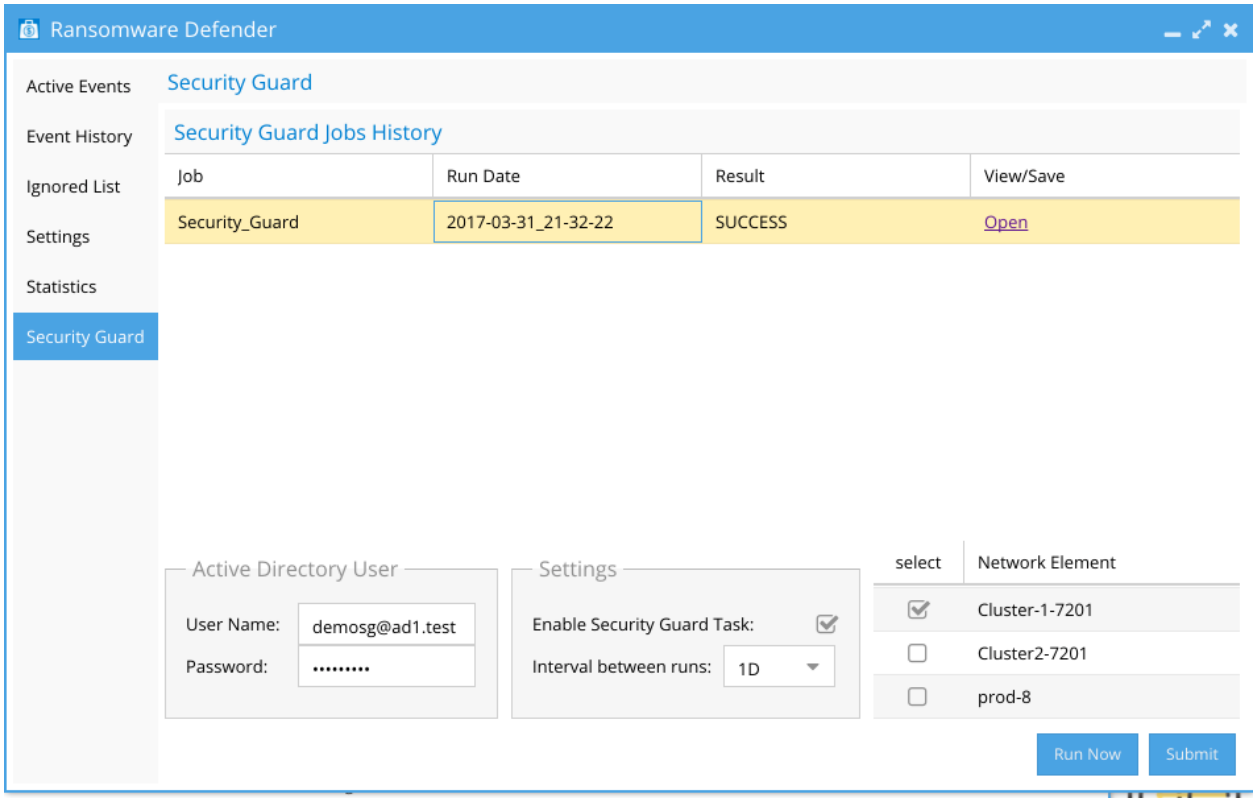

5. Click Open link to review results.

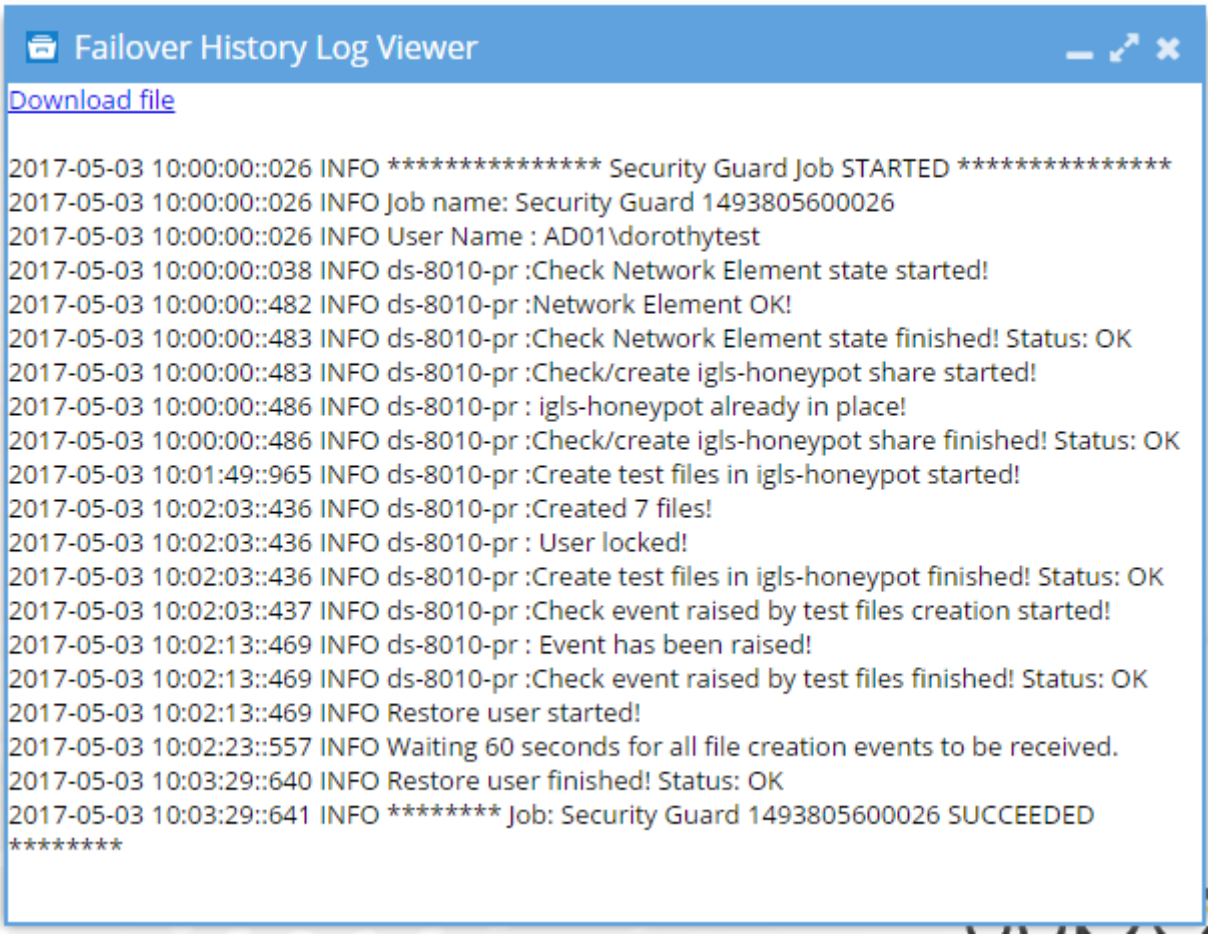

How to test Ransomware Defender with your own Custom File Extension (Isilon Powerscale)

1. Use this feature to test with your own file extension to allow testing complete user lockout and recovery.
2. **Requirements**:
    a. 2.5.7 or later release
3. **Configuration**
    a. Open the Ransomware Defender Icon
    b. Click File Filters tab

c. Click Add file extension button

d. Add a customer file extension that is unique for testing and not used in your environment. (This is important step). Select the Enable option to add the extension.

e. Now check your critical threshold values on the **Thresholds** tab record this value.

4. **How to test**

a. Mount a smart connect name and share in an access zone with auditing enabled. example \\fqdn\smb-share-name

b. Create more files with your custom file extension than the Critical Threshold value to trigger a lockout

c. You can now test creating files with this extension to trigger a lockout action and restore workflow. This can also be used to test alarm creation for integration with SEIM tools.

Configuration (ECS)

1. Requirements

a. Release 2.5.8 or later

b. ECS added to inventory

c. securityguard object user service account is created

2. Open the Security Guard tab and scroll to the bottom

3. Fill in the fields for ECS cluster, name space for the bucket and object user and secret key

a. 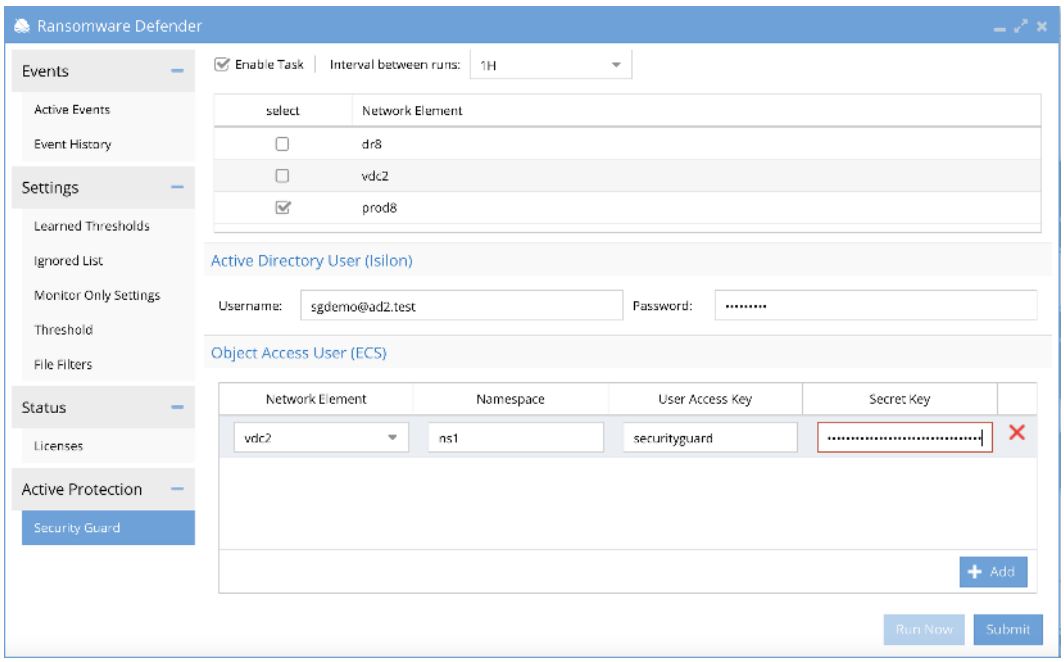

4. Click Submit

5. Now select the ECS cluster in the job settings area and click Enable Task and set the security guard interval to 24 hours and click the submit button to save settings.

a. 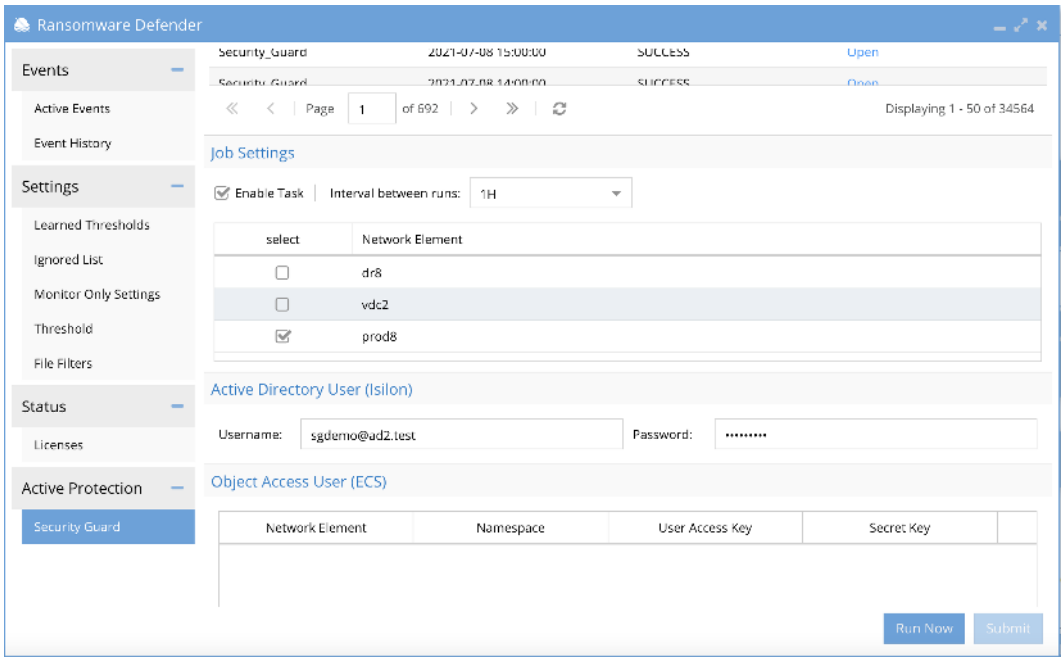

Security Guard Lockout Behavior (ECS)

1. The bucket is created automatically and the object user is added to the bucket for testing
2. The object user should not be used for any other purpose and the user should not be assigned to any other buckets

How to Run on Demand Security Guard Penetration test (ECS)

1. Open the Ransomware Defender window (see screenshot below).
2. Select Security Guard tab.
3. Select the licensed ECS cluster to test in the Job settings section.
4. Select Run Now (see screenshot below).
5. 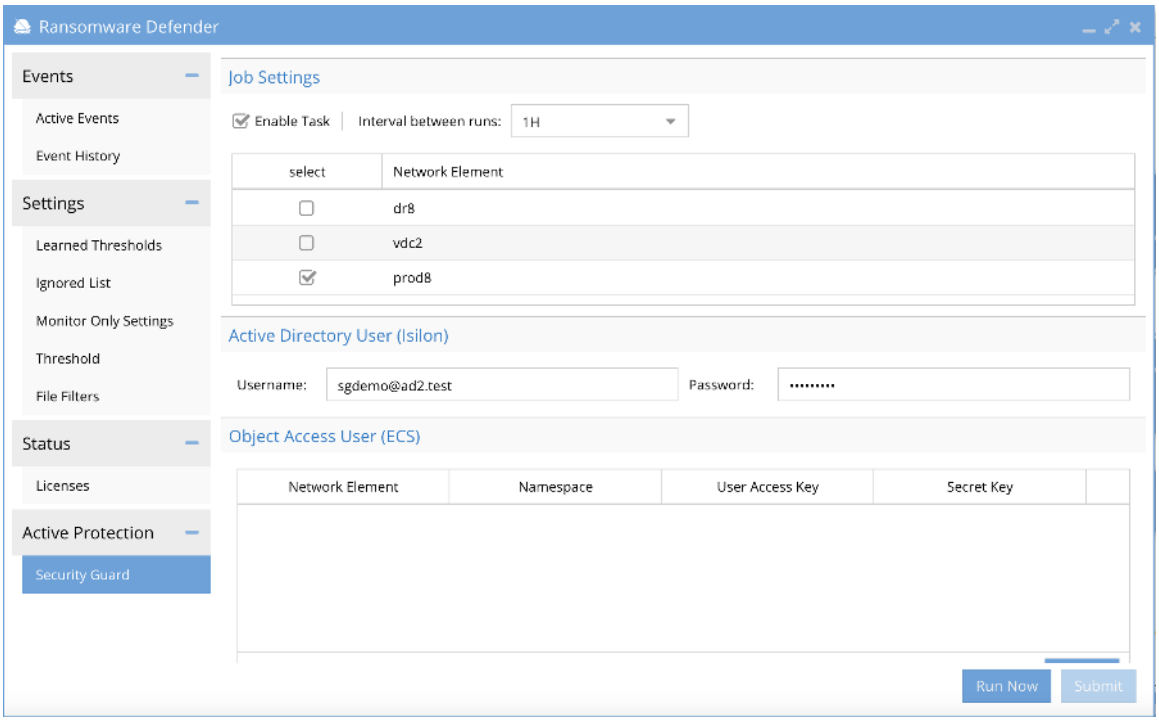

5. Open Jobs icon window

7. View the Running Jobs tab to monitor progress (see screenshot below).
    g.

Advanced Configuration Security Guard CLI Commands (Isilon Powerscale)

In some environments, audit events are delayed before they are sent to the ECA for processing. The security feature writes 100 files, one per second. If the detection of events does not occur before this 100 seconds, the Security Guard will fail the test.

The second phase of Security Guard will restore user permissions and test write access again to the share. This can also have a timer applied to extend the time between the lockout and restore step, to allow authentication and share settings to replicate to the cluster.

These advanced settings can be configured from the CLI to check the timers and set new higher values.

Consult the Ransomware CLI guide.

© Superna LLC

2.7.3. How To Configure HoneyPot file Tripwire

Home Top

- How to Configure Honeypot traps feature
  - Overview
  - Requirements
  - Configuration
  - Files at Base of SMB Share Configuration
  - Sub Folder configuration
- How change the file name used for Honeypot files

How to Configure Honeypot traps feature

This feature allows a new type of detection based on honeypot files placed anywhere in the file system. This would be used where very sensitive data exists and will allow faster detection times for these locations in the file system. They can be placed in as many locations as needed.

1. Can detect slow attack variants of Ransomware, or non-standard IO patterns.
2. Detects file access to encrypt the file itself, and allows immediate critical lockout response.
3. Can detect ransomware even when user behavior does not detect the initial attack pattern.

4. Can reduce the number of files encrypted with any IO that touches these honeypot files, to reduce file system damage.

Overview

1. **Placed at the root of SMB or NFS mounts since this is the first place Ransomware can locate files since drive letters are mounted to the base of the SMB share.**
2. **Uses files as bait for Ransomware, and detects atypical IO access patterns to any of the files in the folder along with many IO access patterns, to find variants that do not use a pattern previously seen before.**
3. **Needs only 1 Signal to raise a Security Event. With Monitor Mode OFF / Critical Mode OFF one signal will place the event into the Major - DELAYED Lockout threshold. With Monitor Mode OFF / Critical Mode ON one signal will place the event into Critical threshold for immediate lockout.**

Requirements

1. A minimum of 3 files should exist in the honeypot at the base of the share, to trip the detector.
2. Each share that needs protection needs the files created
3. Create honey pot files at the base of the share and in a subfolder for maximum protection.

Configuration

Place the following files anywhere in the file system following the procedure below. **From Windows Client repeat these steps on each SMB share that requires Honeypot files configured.**

Best Practice: Always create files at the base of the share and in at least 1 subfolder.

NOTE: The file pattern to match by default is *igls-honeypot-*

Files at Base of SMB Share Configuration

1. Mount the SMB share where you want the honeypot files: smb01 (example share) to drive letter z:
2. Using cmd command prompt cd to this mount point. e.g. Z:\
3. Create 3 files under the SMB share with these exact names :
    a. igls-honeypot-1
    b. igls-honeypot-2
    c. igls-honeypot-3

Sub Folder configuration

1. Mount the SMB share where you want the honeypot files: smb01 (example share smb01 is /ifs/data/smb01) to drive letter z:
2. Using cmd command prompt cd to this mount point. e.g. Z:\
3. Created a folder name: igls-honeypot in this smb01 share (will be created with path: /ifs/data/smb01/igls-honeypot)
4. Create 3 files under the igls-honeypot subfolder:
    a. igls-honeypot-1 - Path: /ifs/data/smb01/igls-honeypot/igls-honeypot-1

b. igls-honeypot-2 - Path: /ifs/data/smb01/igls-honeypot/igls-honeypot-2

c. igls-honeypot-3 - Path: /ifs/data/smb01/igls-honeypot/igls-honeypot-3

How change the file name used for Honeypot files

1. Use this procedure to change the name of the honeypot file names needed for detection.
2. Procedure:
    a. add the following to docker-compose.overrides.yml:
    b. nano /opt/superna/eca/docker-compose.overrides.yml
    c. Paste the below text into the file and make sure the spaces are respected exactl as shown below. Each indent is 2 spaces

```
version: '2.4'
services:
  fastanalysis:
    environment:
      - ECA_RWD_HONEYPOT_DIR_PATTERN
```

1. add the following to eca-env-common.conf: (note any file that contains the string below will trip the detector this string can be changed from this example)
2. nano /opt/superna/eca/eca-env-common.conf
    a. export ECA_RWD_HONEYPOT_DIR_PATTERN="file.docx"
3. In this example you would create files named 1-file.docx, 2-file.docx etc..

4. Follow the steps above to create files at the base of a share and in a subfolder under the share.
5. This will require a cluster down and up after the edits are complete
    a. ecactl cluster down
    b. then
    c. ecactl cluster up

© Superna LLC

2.7.4. How to Flag a detection as False Positive

Home Top

Overview

These steps will flag the event as a false positive for the user that was detected. **NOTE: This will not whitelist the event or user, it will increase the user behavior settings for this user to avoid detection of this user rate of file behavior in the future.**

1. Open Ransomware Defender Icon
2. Click Active events tab
3. Click button in Actions column of the active event, Select the Flag as False Positive action menu

4. 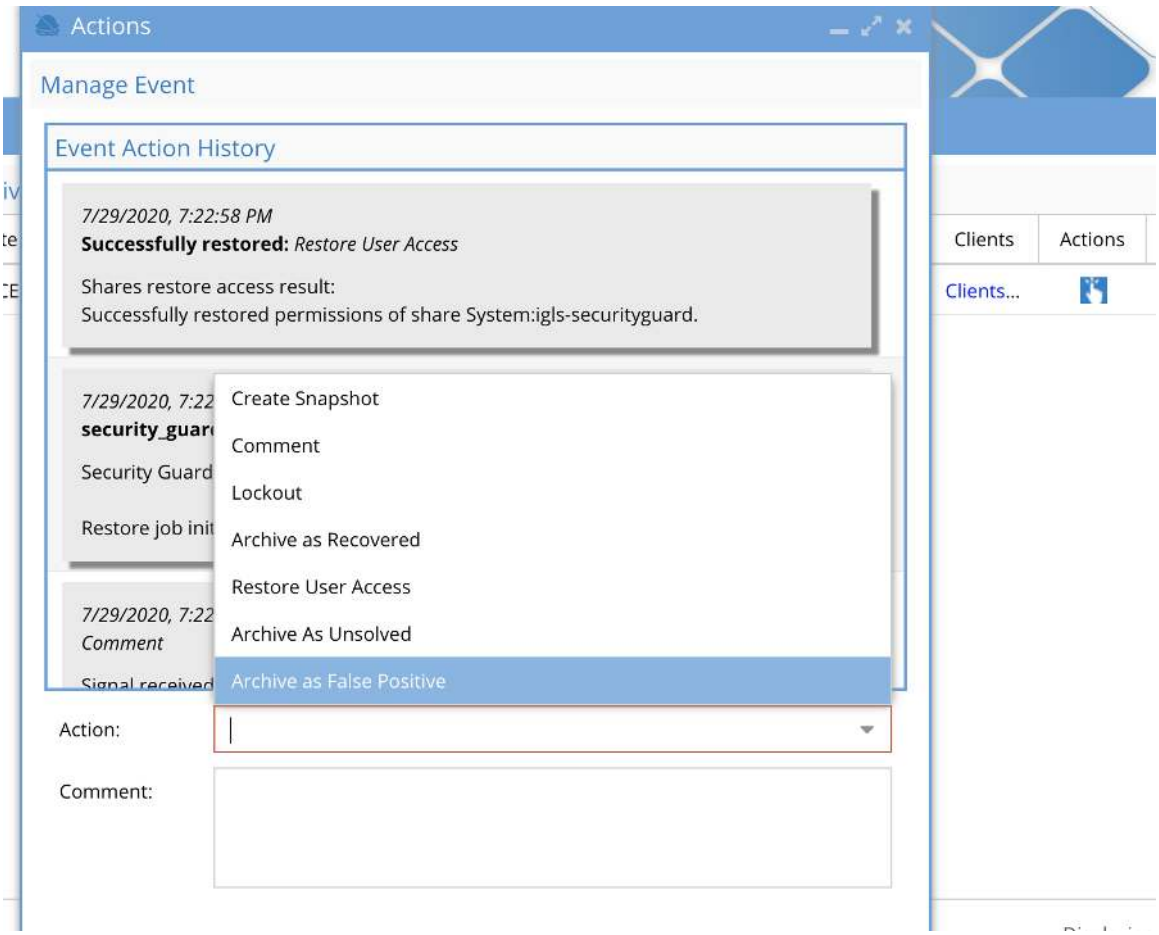

5. Click Submit
6. Verify event is archived to the Event History
7. Verify the false positive was registered
8.

9. Done

© Superna LLC

2.7.5. How to Enable or Disable Enforcement Mode

Home Top

- Overview
- How to Enable Monitor Mode
- How to Enable Enforcement Mode

Overview

This procedure is how to enable or disable enforcement mode. Monitor mode enabled will detect, alert and snapshot data. Disabled Monitor mode will detect, lockout, alert and snapshot. See steps below to enable or disable enforcement mode.

How to Enable Monitor Mode

1. Login to Eyeglass
2. Open Ransomware Defender Icon
3. Click on Settings --> Thresholds and **select** Monitor mode and click submit

4. 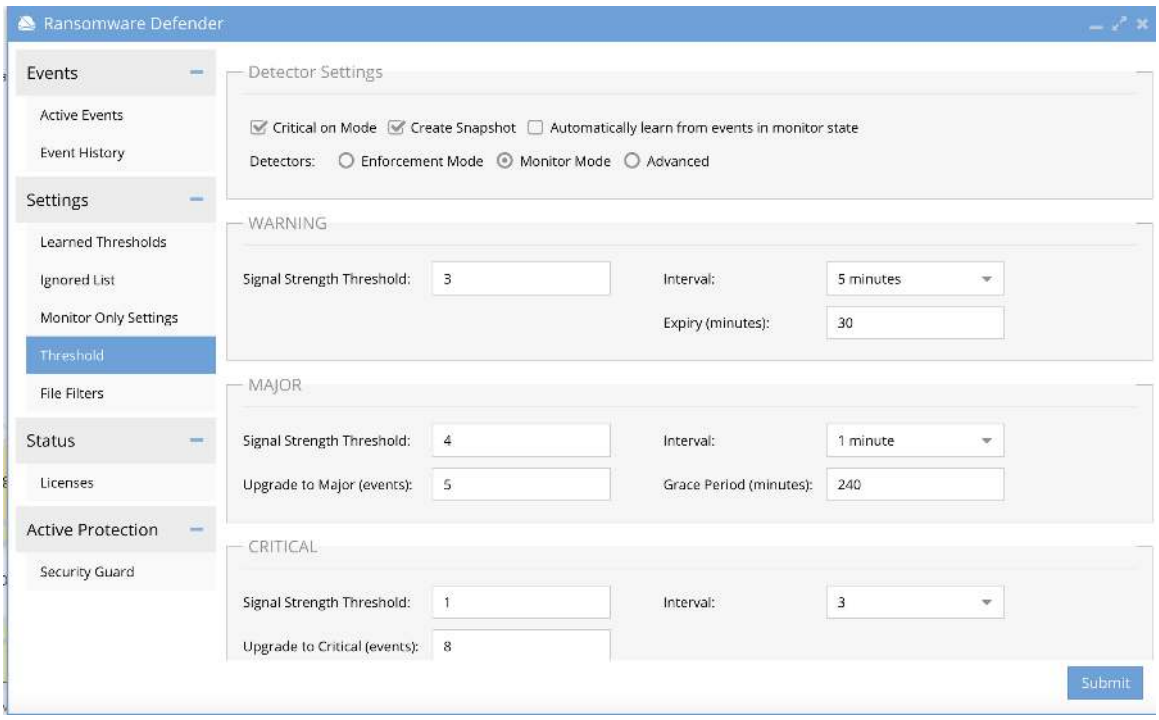

5. Done

How to Enable Enforcement Mode

1. Login to Eyeglass
2. Open Ransomware Defender Icon
3. Click on Settings --> Threshold tab and **select** Enforcement mode and click submit

4. 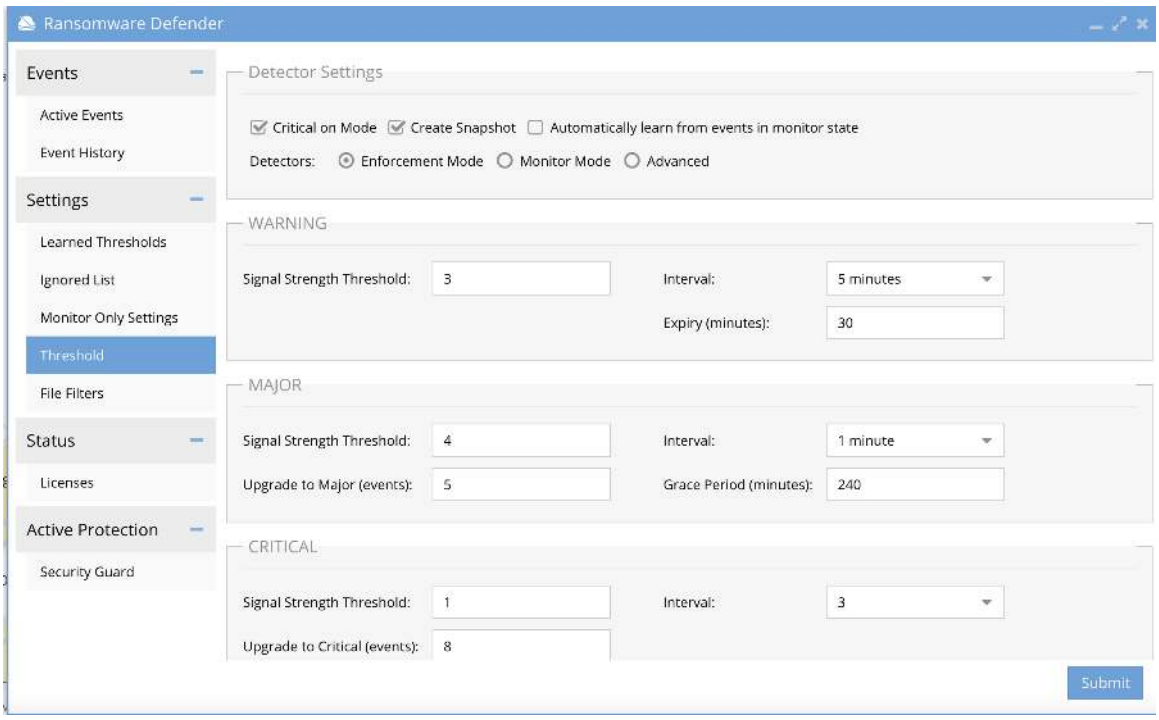

5. Done

© Superna LLC

2.7.6. How to unlock a User that was locked out

Home Top

Overview

These steps will unlock a user and remove the lockout on SMB shares that were locked out.

1. Open Ransomware Defender Icon
2. Click Active events tab
3. Click button in Actions column of the active event, select the menu item Restore User Access

4. 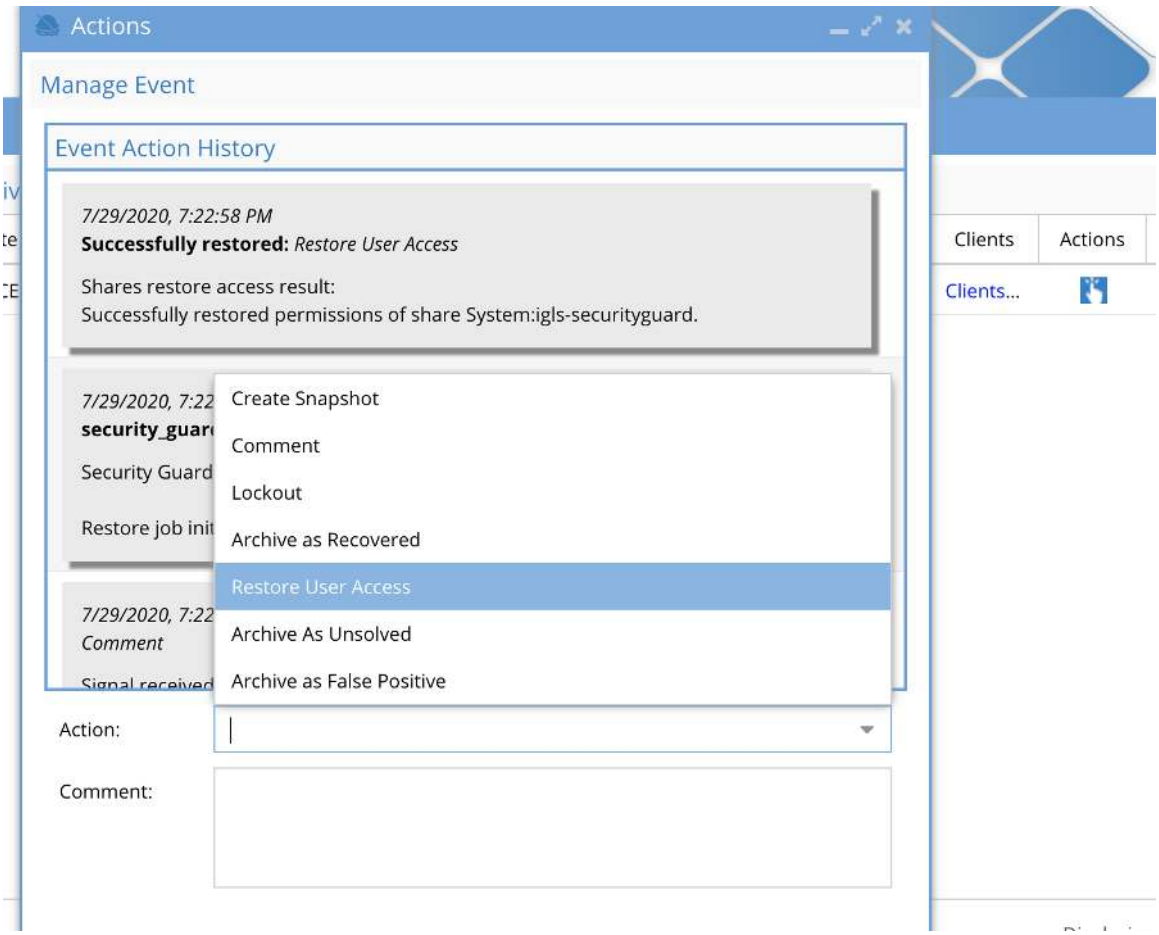

5. Click Submit
6. Verify state changes to Access Restored State
7. Review Actions column menu log to review SMB shares that were processed. They are listed in the log for the event.

© Superna LLC

2.7.7. Ransomware Defender Security Event Workflow for Warning Severity Detections

Ransomware Defender Security Event Workflow for Warning Severity Detections

Under normal working state, it will be normal to see some user behaviors detected as warnings in the active events window. These events will stay in active monitoring state for a period of time (settable in the settings tab). To continue to monitor this user behavior for new threat detectors and rates of detection, promote the event to Major or Critical.

If the user's activity continues to fire threat detectors at or below the Warning rate, the security event will remain in Active monitoring state and will not be Auto Archived. If the signal threshold cross Major or Critical thresholds then actions will be applied to the user behavior. If the severity of the detection stays at the warning level, the event will be auto archived to history. See the steps below to change the default 30 minute expiry settings.

Auto Archive Warning Security Events

This feature simplifies the monitoring of low-grade security events. Warning security events will stay active, as long as new threat detectors for this user continue to be detected, during the auto-archive timeout period. This feature will auto-archive of the event if no new threat detectors fire for this user's security event. The expires column can be used to monitor which events will auto-archive in X minutes from the Active Events window.

How to change the default Auto Archive Expired Warning Events

Use this procedure to set the time period a warning event will stay visible in the active event window before it's archived to the history tab. A longer time period allows for tracking a user's behavior for a longer time period.

1. Open the Ransomware Defender window.
2. Select the Settings tab --> Thresholds
3. Change the auto-archive timeout from the default of 3 minutes to another value in minutes. See screenshot below (Ransomware Defender> Settings > Thresholds> Warning > **Expiry (minutes)** box).
4. 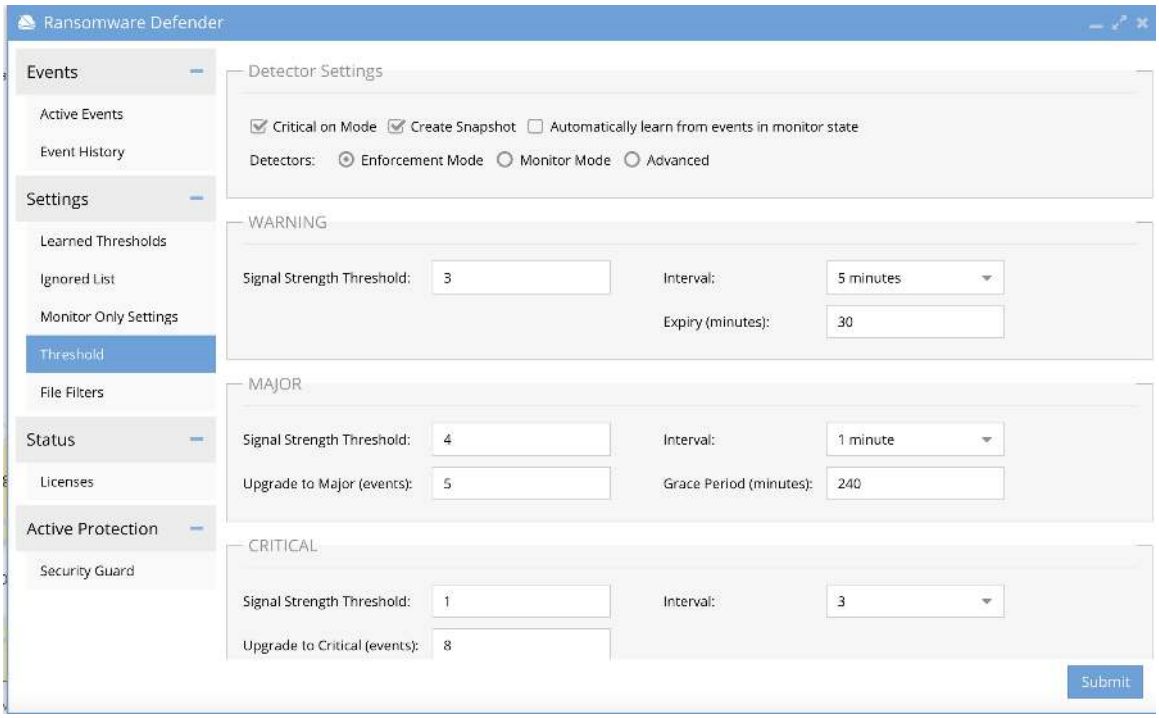

5. Click submit button
6. Note: In the diagram above in the Warnings section the **auto-archive timeout** has been **changed to 30 minutes** in the **Expiry** Box

© Superna LLC

2.7.8. How to Recover Data after a Ransomware Attack

Home Top

- Overview
- When to start data recovery after a Ransomware Attack
- How to recover data after a Ransomware Attack

Overview

The steps in this topic cover the process to recover files using Ransomware Defender CSV files.

When to start data recovery after a Ransomware Attack

Many steps should be completed before starting any data recovery steps.

1. Inspect all of the IT infrastructure to verify impact and no remaining threat exists in the environment. Example Active Directory, DNS, application servers, desktop PC's and any other system required for normal IT operations.
2. The chief security officer or similar role within your Enterprise should declare a data recovery start phase. This phase may not start for many days depending on how long the security audit of the safety of the infrastructure takes to complete.

3. Until this phase has been declared by senior security management within your company no data recovery should be attempted. The risk of data being attacked again from a persistent active threat will increase your recovery phase.
4. Ransomware Defender users in lockout state should remain in this state until the recovery phase is completed and the infected PC's or VM's have been remediated.

How to recover data after a Ransomware Attack

1. To begin the data recovery phase start to build a list of snapshots with creation time stamps in a document.
2. Login to Ransomware Defender open the active events tab and open the snapshots list for each locked out user and record the date and time stamps of each SMB share. Example below.

a.

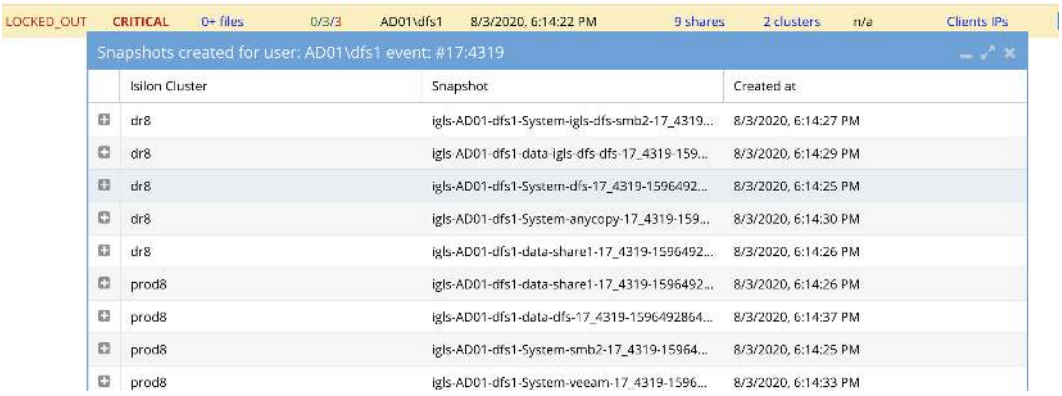

LOCKED_OUT	CRITICAL	0+ files	0/3/3	AD01\dfs1	8/3/2020, 6:14:22 PM	9 shares	2 clusters	n/a	Clients IPs

Snapshots created for user: AD01\dfs1 event: #17:4319

Isilon Cluster	Snapshot	Created at
dr8	igls-AD01-dfs1-System-igls-dfs-smb2-17_4319...	8/3/2020, 6:14:27 PM
dr8	igls-AD01-dfs1-data-igls-dfs-dfs-17_4319-159...	8/3/2020, 6:14:29 PM
dr8	igls-AD01-dfs1-System-dfs-17_4319-1596492...	8/3/2020, 6:14:25 PM
dr8	igls-AD01-dfs1-System-anycopy-17_4319-159...	8/3/2020, 6:14:30 PM
dr8	igls-AD01-dfs1-data-share1-17_4319-1596492...	8/3/2020, 6:14:26 PM
prod8	igls-AD01-dfs1-data-share1-17_4319-1596492...	8/3/2020, 6:14:26 PM
prod8	igls-AD01-dfs1-data-dfs-17_4319-1596492864...	8/3/2020, 6:14:37 PM
prod8	igls-AD01-dfs1-System-smb2-17_4319-15964...	8/3/2020, 6:14:25 PM
prod8	igls-AD01-dfs1-System-veeam-17_4319-1596...	8/3/2020, 6:14:33 PM

3. For each locked out user download the most recent CSV file listed

a. 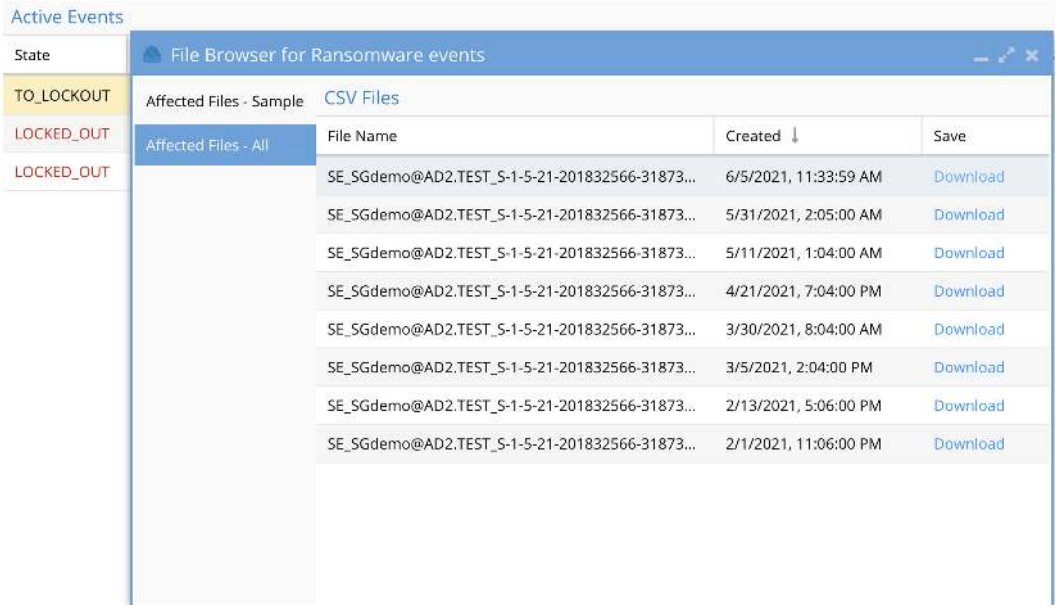

4. The CSV files contain the first 1000 files of the users activity during the attack. To get a more precise list of files including files the user touched prior to the detection an Easy Auditor user report can be run using /ifs/ path and last 24 hours for the search. Example below.

a. 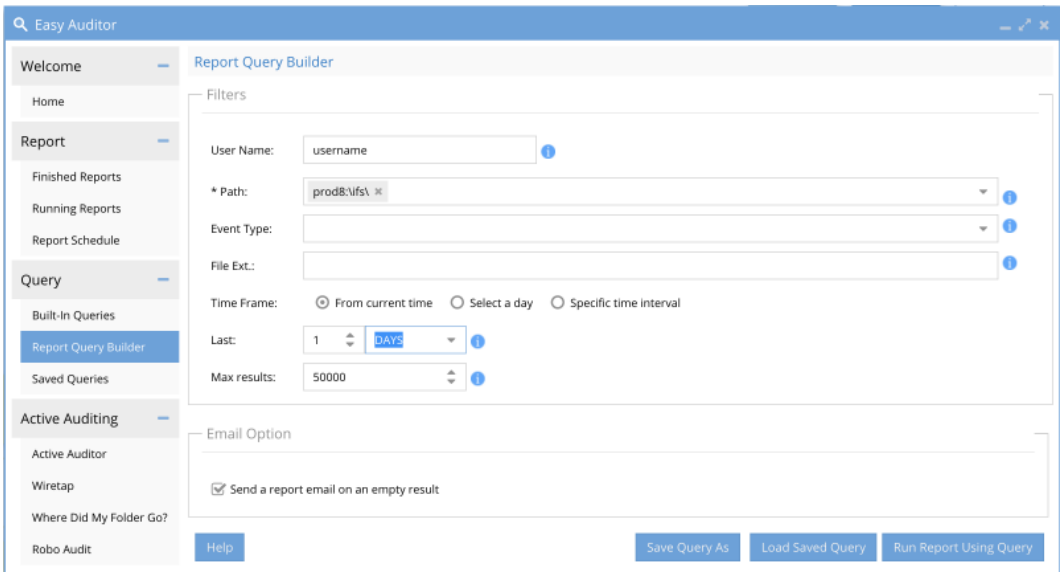

5. Using the CSV files and Easy Auditor reports review the absolute path of the files and user the snapshots taken from the first user that was detected. This user will have the oldest detection time in the active events window. See example below.

a. 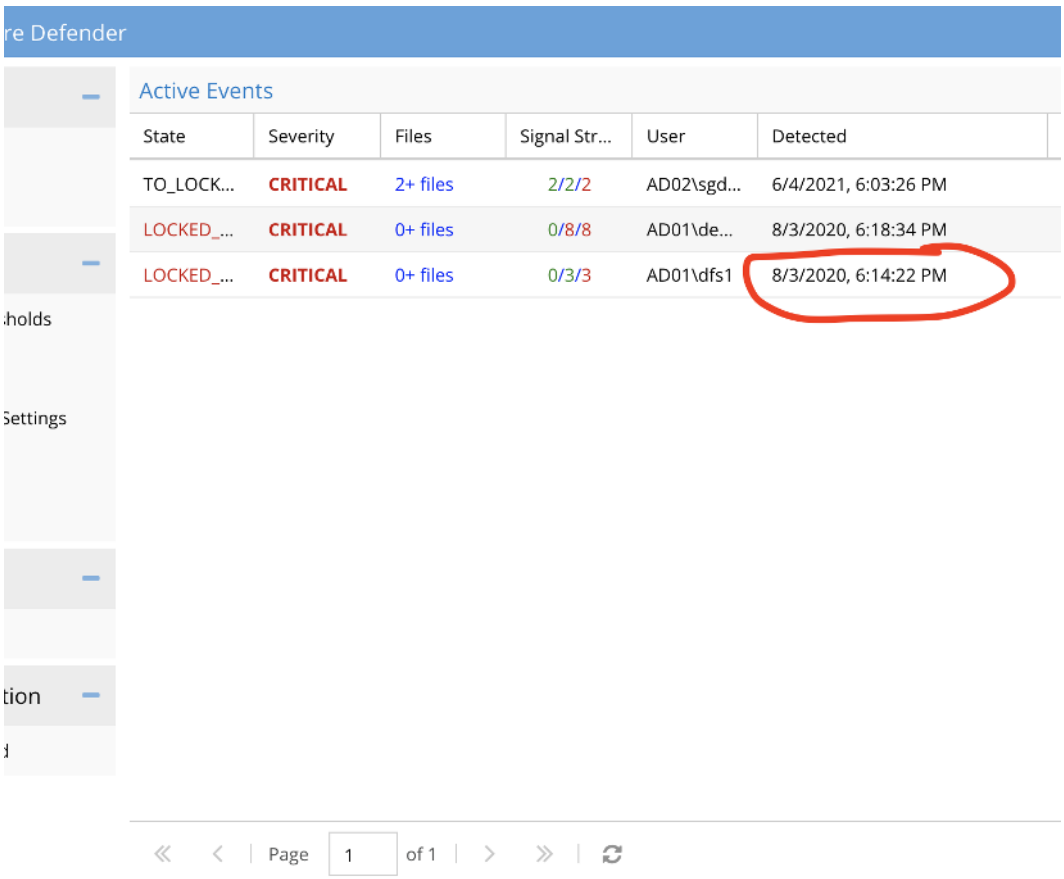

6. Using the snapshot list from this procedure browse to the snapshots listed for user one and use this snapshot to restore the files in the CSV by dragging the files from the snapshot back into the file system. Repeat these steps for each file in the CSV or Easy Auditor CSV report for each user.
    a. NOTE: A visual inspection of the file system where data is being restored should be done during this process. You may delete any encrypted files that are found in the file system or store data for analysis later.
        i. For follow analysis an administrator only SMB share can be created for a post mortem and encrypted files can be moved to this SMB share. These files should be reviewed by security personnel before deletion.

b. NOTE: You may find ransomware notes or possible other strange or unidentifiable file types , these files should be moved to the port mortem SMB share for further analysis by security personnel.

c. NOTE: After completing the post mortem and data recovery the encrypted/compromised files should be deleted.

7. **Unlocking Locked Users -** This should require approval from the CSO or similar senior management. Follow the procedures here.

8. Done

© Superna LLC

2.7.9. How to Enable Learning Mode and Monitor Learning mode Results

Home Top

How to Teach Ransomware Defender about false positives - Learning Mode

1. **Requires:** Release 2.5.7 or later
2. **NOTE: When learning mode is enabled and learning is active a lot of snapshots can be created. Monitor the snapshot usage on your cluster. Snapshots are created with 48 hour expiry by default and will clean up within 2 days.**
3. **Learning Modes**
    a. **Full Learning Mode -** This mode applies to all security events detected an no lockouts will occur and all security events will be used for learning.
    b. **Monitor mode list Learning Mode -** This mode allows both enforcement and learning of monitor mode list entries. In this mode all security events that do Not match a monitor mode learning mode list entry will be enforced and lockouts can occur based on thresholds. For events that match an entry on the monitor mode lists learning will be applied.
        i. Use Case: Service accounts or new application work loads can be added to the monitor mode list by path, user or server IP address to allow learning mode to automatically configure settings for this workload.
4. **Full Learning Mode**

a. Enable Monitor mode (settings tab --> Thresholds) to allow user behaviors to be detected without actions taken to lockout.

b. Now enable Learning mode from the Thresholds screen once **monitor mode** is enabled Settings --> Threshold --> click "**Automatically learn from events in monitor state**". Click submit to save.

c. 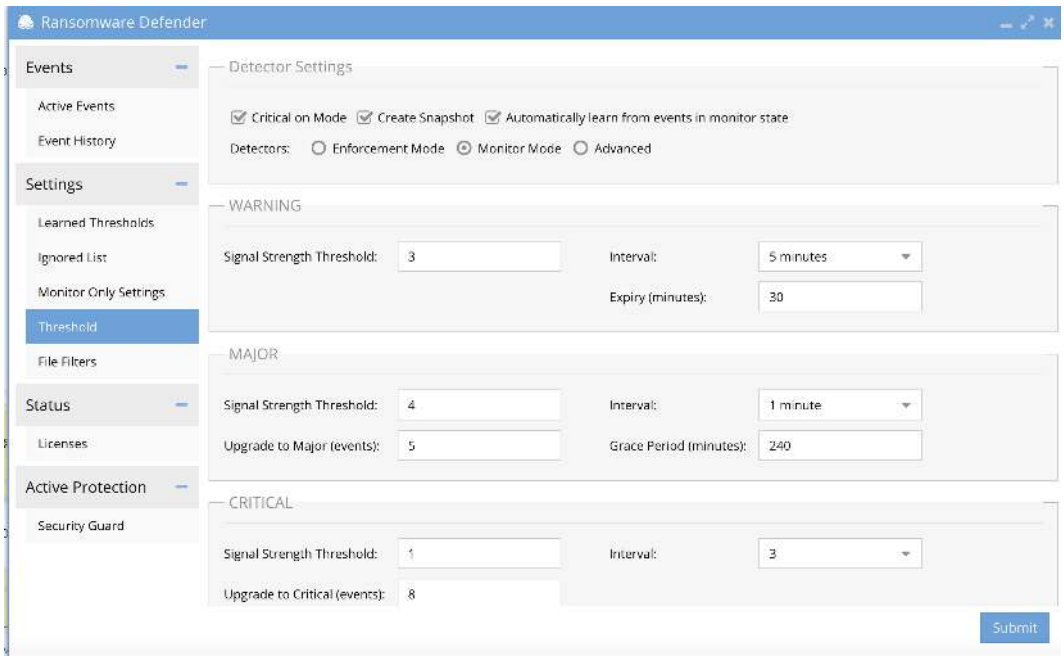

5. **Monitor mode list Learning Mode**

a. Enable Learning mode from the Thresholds screen once **monitor mode** is enabled Settings --> Threshold --> click "**Automatically learn from events in monitor state**". Click submit to save. Example screenshot below.

b. 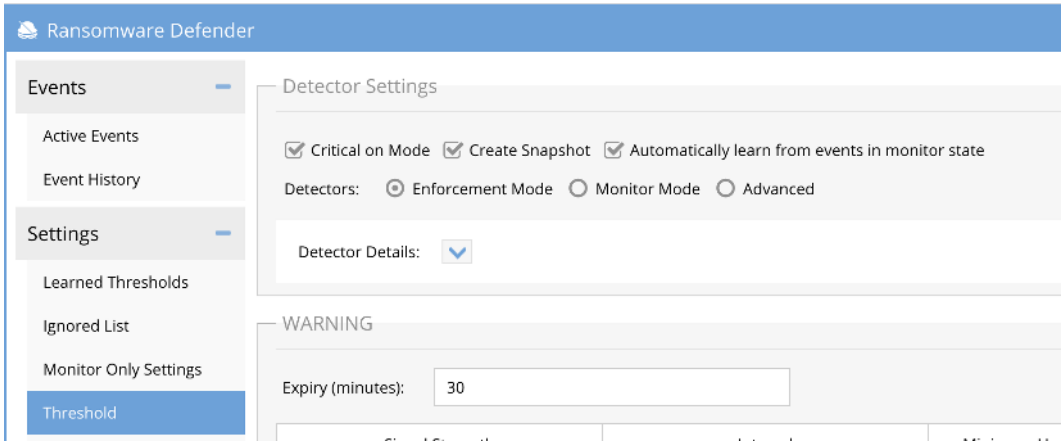

6. Leave this enabled for 2-3 business days and monitor the customized user behavior settings on the Learned Thresholds tab.

   a. This is where Learning mode will place customized settings. It will also set file extension detections on the File Filter tab into a disabled state so this file extension will not be detected as Ransomware.

   b. 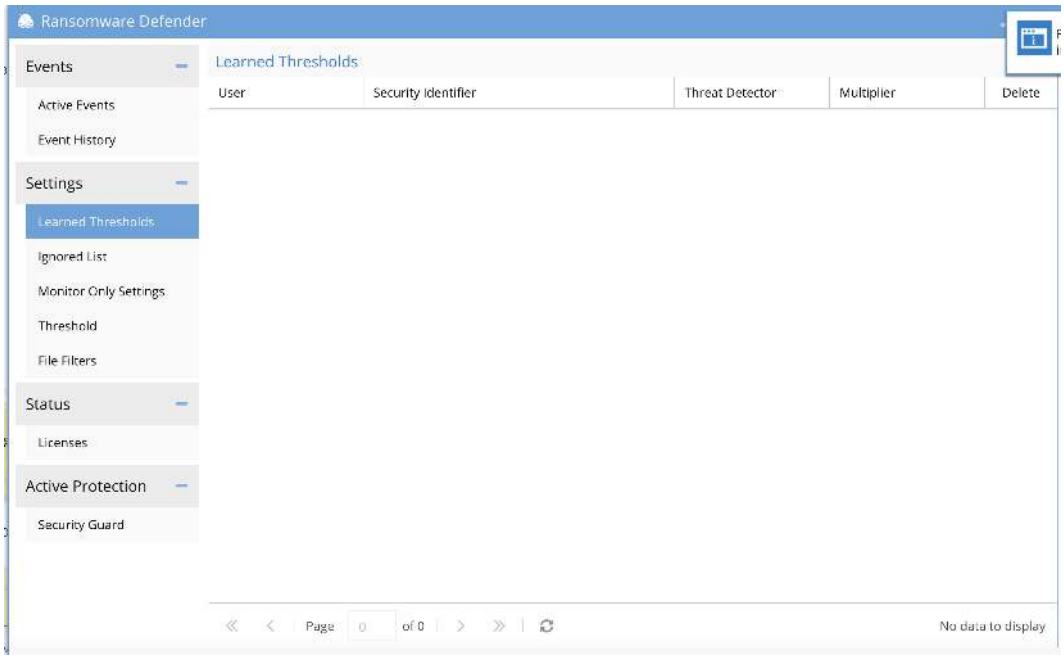

7. **The process to disable Learning Mode and then enter Enforcement Mode.**

a. Review user settings on the **Learned Thresholds tab** to approve the list of users or NFS hosts or delete entries as needed. Consult with support or accept the learned behaviors.

b. Review the File Filter list extensions that are disabled status, these extensions have been placed on the Allowed list and will not trigger a detection.

   i. Use the filter option to locate all the disabled file extensions by entering Disabled in the filter box.

   ii. 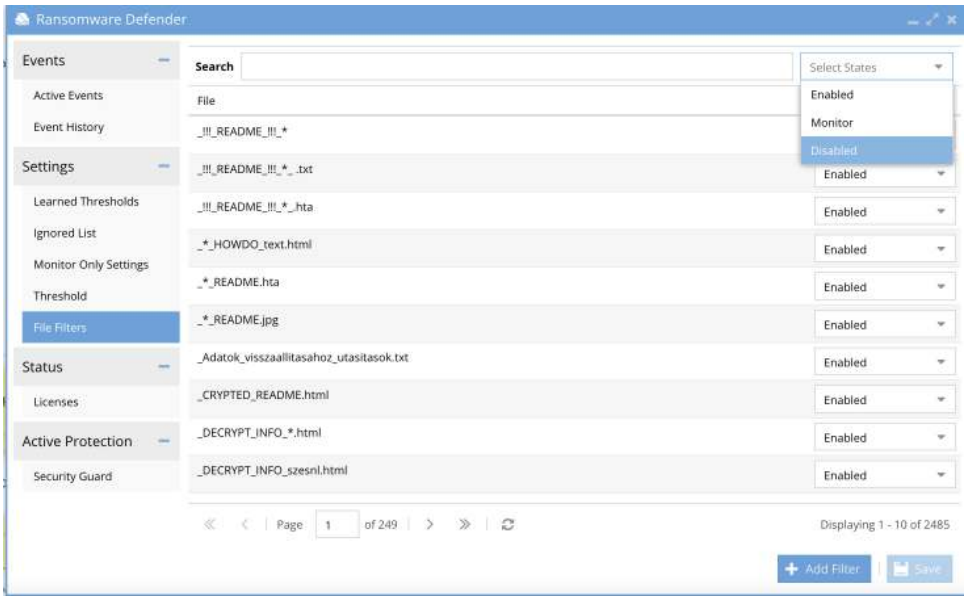

   iii. Review all the extensions that were detected and disabled. If they are acceptable no action needed.

   iv. To change the setting on the extension to enable enforcement and detection of this file extension, you may also chose monitor mode on the file extension to allow detection, snapshot but no lockout for this file extension.

      1. 3 possible modes for each file extension enabled (full enforcement), disabled (ignored),

monitor mode (detect, alert, snapshot and no lockout)

c. Disable Learning mode once the file settings are confirmed from the Settings-->Threshold tab and click submit to save. This only disables learning mode and remains in Monitor mode.

d. To enter enforce mode mode disable monitor mode from the Settings-->Threshold tab and click submit to enter enforcement mode.

© Superna LLC

2.8. AirGap 2.0 Guide

Home Top

- Overview
  - Golden Copy File to Object Integration
  - Easy Auditor Integration Enables Custom Vault Open Criteria
- Key Features
- Solution Summary diagram
- AirGap Location Independent Solution
- How protect high change rate and low change rate data
- FAQ - Buyers Guide to Cyber Vaults
- Requirements & Prerequisites
  - Firewall Vault Network
  - Additional Requirements for Enterprise AirGap Licensed Deployments
- High Level Configuration Steps
- Security Configuration of Components
  - Eyeglass VM Security
  - Vault Cluster Configuration
  - Role Based Management of the AirGap Feature
- Detailed Deployment Diagrams
  - Network Considerations for Layer 2 or Layer 3 Vault Network
    - Overview

    - Layer 2 or Layer 3 - Fan-In cluster protection (Enterprise Airgap License)
  - Virtual AirGap Mode - Layer 3 Vault Network (Basic Airgap License)
  - Inside the Vault mode Deployment (Enterprise AirGap License Required)
    - Overview
    - Inside the Vault Physical Topology - Layer 3 example
    - How to Enable Inside the Vault Agent VM (Enterprise AirGap)
  - Operations of Vault Data Replication
  - Data Flow Example for Data Replication with Enterprise Airgap
  - Vault Management Data Flow Example for Enterprise Airgap
- Configuration Steps for AirGap Setup
  - Overview video
  - How to setup syncIQ policies for AirGap
  - How to Configure AirGap policies and setup Virtual AirGap (Basic Airgap License)
  - How to test an Airgap Policy job
  - How Alarms from the vault Isilon are Viewed and Forwarded
  - How to Expand the Airgap Sync Job Timeout and the Airgap job prefix name
- Operational Procedures for AirGap Management
- How to stop AirGap Replication in an Emergency

- How to monitor replication AirGap policy success failure
- How to Monitor AirGap Replication Reports
  - How to enable or disable the Airgap daily summary report or change the schedule
- How to pause all AirGap policies to complete Vault cluster maintenance
- How to Pause the AirGap policies for maintenance with a timed auto close of the AirGap Network

- How to Configure Enterprise AirGap Ransomware Defender Enterprise Airgap Agent
  - Overview
  - Topology and Communications
  - Prerequisites
  - Configuration Steps
- Operational Procedures Enterprise Airgap
  - How To reach outside the vault through the vault cluster it is possible to open and close the vault with cli commands
  - How to Open the Airgap for maintenance from Ransomware Defender CLI
  - How to list running jobs
  - How to run an Airgap job from the Vault agent VM
  - How to monitor a running airgap job
  - How to check the remaining time of a maintenance window request on the vault agent

- How to configure Vault cluster Log Gather Automation for Hardware Support
- Advanced Vault Agent and Airgap Eyeglass Configurations
  - Scheduled check for new or changed Airgap policies
  - How to change the name of the Airgap policies.
- Security
  - Airgap Audit log
- Recovery Scenarios
  - Considerations
  - Partial Vault Data Recovery Scenario
  - Complete Vault Data Recovery Scenario
  - DR Vault Data Access Scenario - Rapid Recovery

Overview

Superna offers several products to protect data from Ransomware or unauthorized access. We recommend the Ransomware Defender product as the primary tool to protect data since it covers all requirements for detection, prevention, and recovery. Ransomware Defender includes an AirGap 2.0 solution that provides the only solution on the market that integrates user behavior detection of the protected data to suspend data updates to the secure vault copy until administrators take action on the alarms. Complete role based administration solution with split roles for user behavior monitoring and separate AirGap administration.

Golden Copy File to Object Integration

Expanding on file based projection solution for Isilon is the introduction of Golden Copy Advanced that can copy files to Objects off site locations example Amazon S3 or Azure to allow. This integration allows Golden copy to get real-time updates on source cluster data threats from Ransomware Defender user behavior monitoring and suspend syncing to the S3 targets.

Easy Auditor Integration Enables Custom Vault Open Criteria

1. Customers that own the Easy Auditor platform can extend the vault auto close replication criteria using Easy Auditor active auditing. This extends the security of protecting the vault data by using builtin triggers for DLP , Mass delete or custom triggers to control vault replication.
    a. This enables security teams to apply user aware, network aware policies that will stop replication for any active events in Easy Auditor.
    b. This provides a powerful capability to customize when the vault replication should occur depending on the data that is being protected. No other solution offers fully customizable real time triggers to control vault replication.
    c. The guide on active auditing can be referenced here.

Key Features

1. Unified data protection at the source and Automated AirGap in a single product.

2. Integrated user behavior detection "Smart AirGap" ensures data sync is suspended when suspicious user behavior is detected on the source production cluster.
3. Split roles between AirGap management and Vault cluster access and passwords and day to day monitoring of data protection of the source production cluster.
4. Fastest incremental always sync solution on the market powered by SyncIQ, keeps the AirGap open the least amount of time to sync the block changes of individual files into the Vault.
5. Lowest Cost storage with longest retention of Vault data with SnapshotIQ that only stores the block differences of the changes made to the file system in production. This translates to the longest retention of your versioned data with the lowest cost.
6. The only solution that offers < 2 hour repaid recovery of any quantity of data. Protect Peta bytes of data and bring it online < 2 hours provides unparalleled recovery speed to handle the worst case data recovery scenarios.
    a. Eyeglass DR mirrors shares, exports and quotas to the vault cluster to ensure your rapid recovery solution has the exact data access security as production. **No other solution considers configuration data as a critical component of a Cyber Vault.**
7. Easy Auditor extends the solution to include full file auditing solution to monitor suspicious user behavior covering Data Loss Prevention, Mass Delete detection and custom triggers along with historical searching of audit data.
8. Full User auditing within the Unified Desktop

9. Future API integration with Golden Copy to allow Golden Copy copy/sync control when active threat to source cluster is detected. For more information on license requirements for Golden copy click here.
10. Integration with Easy Auditor will block replication to the Vault cluster if Mass delete, DLP or other custom triggers have active events.
11. Inside vault Smart Airgap solution provides a VM inside the vault to automate the Airgap with the same functionality as the virtual Airgap mode.

Solution Summary diagram

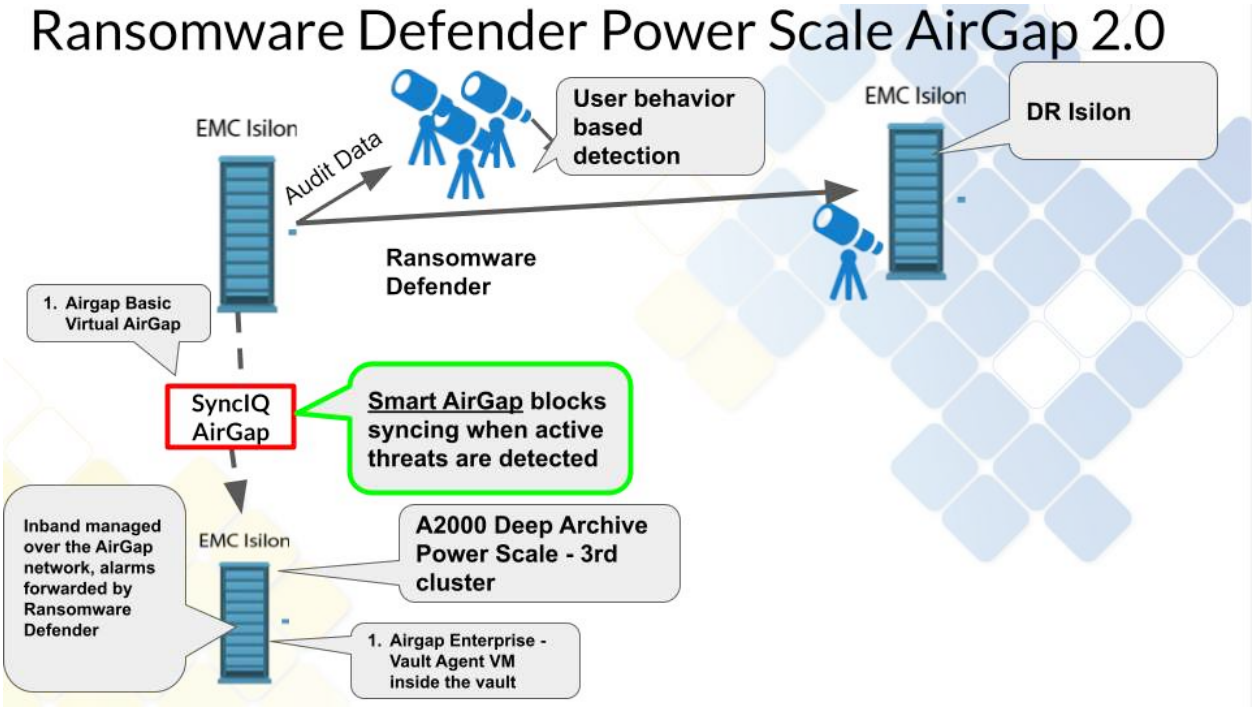

AirGap Location Independent Solution

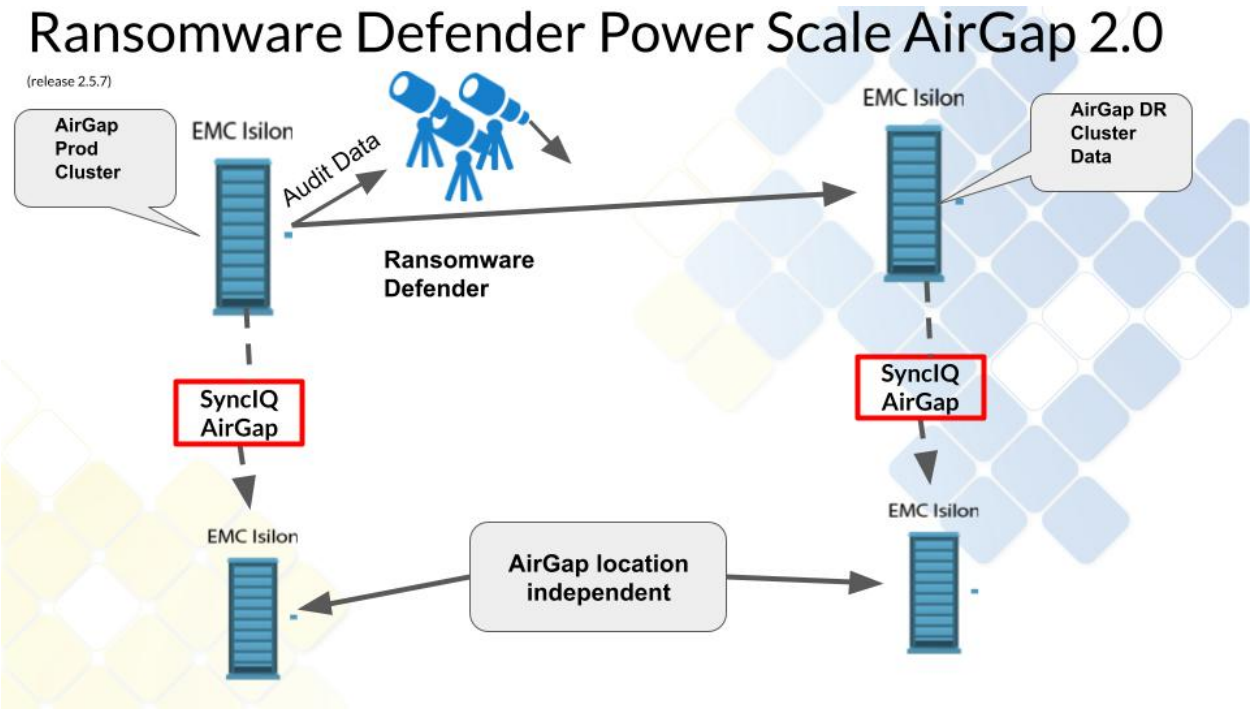

How protect high change rate and low change rate data

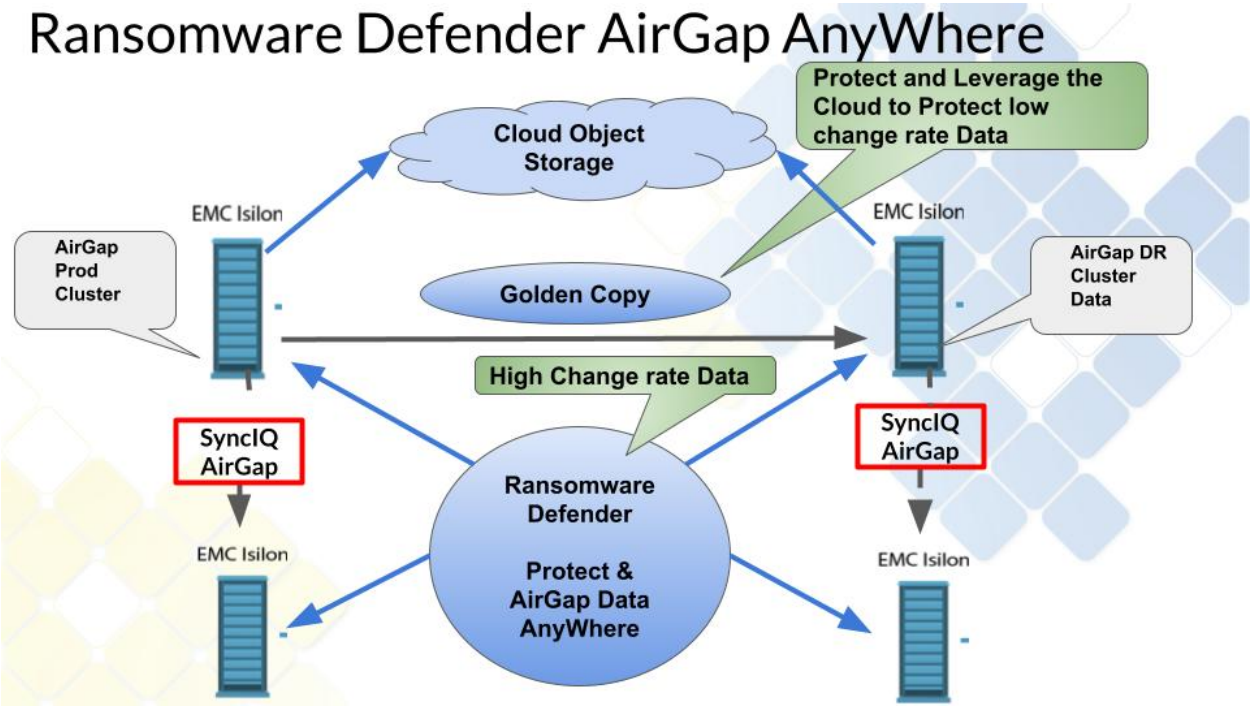

FAQ - Buyers Guide to Cyber Vaults

1. How will the data be protected in the vault?
    a. **The 3rd PowerScale is the vault since SyncIQ locks the replicated data in a read-only state, this data cannot be deleted or modified even by the root user on the cluster.**
2. Can I use my DR cluster as the target?
    a. We recommend a 3rd cluster because DR replication should be as fast as possible ie 5 minutes to protect against site failure. A 3rd Cyber vault copy should be slower example 24 hours to allow detection time of data issues that would require the 3rd copy. This buffer time is important to ensure compromised data is not replicated into the vault.
        i. The 2nd reason is an AirGap means fire-walling and securing the 3rd copy device. The DR should be available and reachable for a DR event and blocking all access to the DR site would compromise a DR solution effectiveness and readiness. The firewall would still require Eyeglass access to the DR cluster so this solution is not as secure as a dedicated Vault cluster device.
        ii. In summary, a 3rd copy AirGap copy of data and DR have opposite requirements that prevent using the DR cluster as the AirGap device due to sync interval and security impact of fire-walling your DR site.

3. Can the Airgap be opened and closed from the inside of the cyber vault?
    a. Yes Ransomware Defender supports 2 modes of operation with an outside the vault automation and an inside the vault automation that requires additional resources inside the cyber vault.
4. Is the data immutable? What is the granularity?
    a. Yes, the data is immutable and cannot be modified regardless of the permission applied to the data
    b. The entire source cluster, or only certain paths, can be replicated to the vaulted PowerScale. This is a matter of creating more SyncIQ policies to sync specific paths of data to the vault.
5. How will the AirGap be managed? Internal to vault or external? What are the requirements?
    a. The recommendation is the vault PowerScale is located next to the source cluster since this is not a DR copy of the data. Off-site remote is supported using SyncIQ which allows the vault PowerScale to be located anywhere. This will add AirGap networking requirements to ensure no IP routing to the remotely located cluster, and will not allow any access other than from the source production PowerScale. This is a networking implementation with firewalls but is fully supported.
        i. Considerations for location of the vault onsite versus remote site means you are trading off rapid recovery

with the onsite device versus remote that offers a longer recovery of data to an onsite device.

ii. Recommendation is always use an on site solution since DR is designed to protect against site failures and a Cyber vault is protecting data from an on site threat.

6. How will a customer get into the vault to make future modifications?
    a. Inline SSH access to through the production cluster. This will allow access to the vault PowerScale over SSH with a command that can open the AirGap for maintenance, and will automatically close the AirGap after X minutes.
    b. Optionally a physical management interface can be connected to the vault PowerScale for GUI access. This would require physical access to connect and disconnect the management interface. The recommendation **is to use the inline SSH access since the PowerScale CLI provides all management commands.**
7. What type of maintenance should be expected?
    a. A feature will collect alarm data via the AirGap when a copy is running, and forward the data through the Normal Eyeglass alarm feature (email, syslog snmp etc..). This allows full alarm event monitoring of the vault PowerScale even when it is fully disconnected most of the time. Any serious alarms detected would then require remediation, or CLI access to debug the issue on the Vault PowerScale.

b. Benefits: **inline AirGap vault aware alarm monitoring avoids most scheduled maintenance on the vault PowerScale until a critical alarm requires action.**

8. What is needed in the vault to perform analytics or reporting?
    a. All reporting and analytics are done by Ransomware Defender from the source cluster. All data copy jobs are monitored, trended, and reported over 24 hours, 30 days, and 60 days. A report is emailed daily that covers success, failure, throughput, average AirGap open time, and more along with a CSV with raw per path and policy reporting.
    b. **Benefit: Fully automated monitoring and reporting.**
9. How will data be made available outside the vault in a recovery scenario?
    a. The vault PowerScale can have data exposed for 3 different scenarios and offers flexibility that no other vault backup solution offers since PowerScale can serve data over SMB, NFS as well as act a vault.
        i. **Access to a subset of the vaulted and immutable data**. Connect the management interface, create an SMB share on the data path that is required for recovery. NOTE: The share data will be read-only regardless of the permissions since it is locked by SyncIQ.
        ii. **Access or recover all the data.** SyncIQ steps can be executed to reverse replicate the data from the vault PowerScale back to the production cluster and will use SyncIQ speed and performance advantage to restore data and ACL permissions to the source

cluster. This would be used when the volume of data that is needed to be recovered is a very large % of the data.

iii. **Emergency Operations mode.** This mode highlights the advantages of a vault PowerScale solution. This scenario turns the vault PowerScale into the production cluster to serve data directly from the vault without a recovery phase. This allows getting operational as fast as possible to operate the business. High-level steps for Emergency Operations Mode are outlined below.

1. Connect management interface of the Vault PowerScale to the network
2. Eyeglass DR can copy the shares, exports, quotas to the vault PowerScale (requires the Vault PowerScale to be added to Eyeglass, delete DR cluster first and add Vault PowerScale)
3. Pre-staged IP pools, SmartConnect names from the production cluster needed to get access to some of the data that is urgently needed. Connect interfaces to the network. Update DNS to point at the new vault cluster with NS recored edit.
4. Execute a failover from Eyeglass to the Vault cluster and sync all shares, exports and quotas to re-secure the vault data.

5. Start accessing the PowerScale data.
6. Done.

10. How will testing of the data recovery be done?

a. In vault:

i. Via ssh from source cluster - Use Open AirGap command to ssh to the vault cluster and then use scp command to copy the files to the production cluster.

b. External:

i. On a quarterly basis - Connect the management interface on the vault cluster to the network, create an SMB test share mount, and test read the protected data. **NOTE:** no write access will be allowed in this test mode.

ii. Delete the share and disconnect the management interface to complete the test.

iii. **Benefit: vault replication is still active during a test with the source cluster still in full production mode. No downtime is needed for this test.**

11. Who is creating the recovery runbooks?

a. Everything is documented in this guide to operate and test the vault data, and most management is automated with no day to day tasks needed. The professional service also offers customers assistance in the design, implementation and operations. See here.

b. No Run books are needed since most tasks are automated. Pre-staged rapid recovery option is recommended.

12. Is there any dependence on NTP or other services?

    a. No, the cluster time can free run from its own clock.

13. Are there any additional hardware/software components that may be recommended or required to make the overall solution work?

    a. 2 AirGap Options are available:

        i. Virtual AirGap - Requires a layer 2/3 switch to route between the source PowerScale and the Vault PowerScale. All other requirements are within Ransomware Defender to manage static routes to reach the Vault PowerScale.

14. Is there any visibility in production that a copy is being sent to a vault?

    a. Yes, full reporting 24 hour, 30 day and 60 day reporting on all copied data with success, failure, and throughput metrics along with AirGap average open time. This is the time the Gap is open and should be minimized at all times. The solution reports on this daily, or an on-demand report can be created.

Requirements & Prerequisites

1. License Requirements

a. Ransomware Defender license for each source cluster that is protected by the vault.

b. **Airgap Enterprise** - Agent VM license for inside the vault automation.

   i. No DR license for the Vault cluster is required.

   ii. Future roadmap to include Smartlock automation feature inside the vault.

c. **AirGap Basic** - This is builtin to the Ransomware Defender license. This will only support virtual AirGap with outside the vault automation.

   i. Vault cluster DR license with maintenance is required for a supported AirGap.

2. A PowerScale cluster, any make model with OneFS release matching source cluster, sized for the data set and change rate and retention of data required. Sizing can be done with assistance from the sales team.

3. Dedicated Airgap Ethernet switch

   a. Minimum of 4 x 10G ports for synciq port connections.

   b. **NOTE: It is not recommended to use the front end ethernet switches to connect to the vault cluster using a vlan. Physical connections offers the best practice network separation and reduces the attack surface.**

4. **OR Firewall - Enhanced network option for the vault networking**

   a. **This option provides additional control of ports and data flow into the and out of the vault cluster. This provides enhanced security to reduce the potential attack surface and provides**

logging. The inside vault switch can be a firewall to lock down this network.

5. **Powerscale Production Node Connections**

    a. **Best Practise with High Availability**: At least 2 nodes and 1 interface per node on the production cluster and 2 nodes on the Vault cluster connected to the AirGap Ethernet switch.

    b. **Next Best option without High Availability**: One node and 1 interface connected from production cluster to the Airgap Ethernet network switch.

    c. **NOTE: if the production cluster has no available ports the choices above offer lower port count for physical separation connections to the AirGap network. It is also possible to add nodes to the production cluster that are dedicated to connect to the AirGap network.**

Firewall Vault Network

1. Production cluster --> to the vault cluster with vault network firewall deployment

    a. ssh (maintenance only)

    b. 8080 https TCP --> **optional** used for vault cluster hardware alarm collection and free space reporting only during data syncing

    c. synciq ports TCP ports 5666, 5667, 2097, 2098, 3147 and 3148

Additional Requirements for Enterprise AirGap Licensed Deployments

1. Hardware Recommendation:
    a. A dual socket server with 128 GB of RAM, 1 and 10G Ethernet interface options, 2 - 4T B of local flash storage.
        i. Hardware should be future proofed to allow additional VM's to run for cyber security protection solutions and Windows desktop(s) for administrators that has key tools installed to allow a guaranteed clean, secure OS desktop to be used for used recovery operations or upgrades to the Isilon hardware, firmware and software.
2. Ransomware Defender VM Agent
    a. Vmware ESX host server that will run a single Ransomware Defender VM with 16G ram 130G disk and 4 x vcpu. This ESX host only needs to run this one VM but can be used to run other applications inside the vault.
3. Networking - Basic Airgap option more secure option below
    a. Ethernet switch to connect the VMware ESX host to the management ports on the vault cluster. The 1G Ethernet interfaces can be used to connect to the ESX host using the system zone management interfaces.
    b. Allows device expansion in the vault for future equipment
4. OR Firewall - Enhanced network option for the vault networking

a. **This option provides additional control of ports and data flow into the and out of the vault cluster. This provides enhanced security to reduce the potential attack surface and provides logging. The inside vault switch can be a firewall to lock down this network.**

High Level Configuration Steps

1. Install 3rd PowerScale at the same location as the cluster with data to be protected.
    a. **Best practice deployment:**
        i. **Use a bastion host (VM connected to private IP vault management network) and complete all configuration of the vault cluster through this bastion host. This avoids connecting the vault cluster to the corporate network during commissioning steps. The vault cluster should never be exposed to the network directly.**
    b. **NOTE: The airgap can be located at the DR location using the DR copy as the source of the data to copy to the vault.**
    c. **Vault PowerScale Requirements on Deployment :**
        i. **Cyber Recovery RunBook:** As much of the pre-configuration, labeling, Ethernet port planning (VLAN's,) cabling and logic configuration as possible should be completed at deployment time to speed up recovery scenarios described in this guide.
            1. The configuration steps completed should be documented in a **Cyber Recovery RunBook**.

This will be used along with this guide when executing a cyber recovery scenario. This guide documents the high-level steps needed to complete recovery. These high-level steps should be turned into detail specific steps for your environment and added to the **Cyber Recovery Runbook**.

ii. Management System zone access network configured but should be disconnected physically after installation.

iii. **Vault Cluster High level hardening**

1. Note: The Advanced service will provide detailed hardening of the vault cluster. This service scope is out lined here. The information below is not the complete solution and only identifies high level steps.
2. Delete all default shares and NFS exports on the cluster.
3. Stage and plan physical ports, or VLANs required for the Vault cluster nodes, to be connected to the production network in the event a rapid recovery scenario is required. These cables should be physically in place but not connected, with labeling applied to each cables Ethernet port connection. The node Ethernet interfaces should be the minimal configuration

needed to serve data for production IP pools and Access Zones.

4. Stop the SMB and NFS services.
5. Add NTP server (even though it will not be reachable). Used for Rapid Recovery.
6. Add DNS servers (even though they will not be reachable). Used for Rapid Recovery.

iv. For additional hardening consult Dell documentation on how to apply additional changes for hardening.

1. **NOTE: The Airgap Design And Implementation Service includes hardening of the vault cluster based on Dell documentation.**

v. **Production Powerscale IP Address space and IP pool for replication with SyncIQ and management IP pool**

1. Use the default groupnet and subnet (enable vlan tagging on the subnet).
2. A syncIQ pool and management pool will be created in this subnet
3. Review the layer 3 vs layer 2 vault network pros and cons.

vi. **Production Powerscale Management IP pool**

1. Create a new IP pool in the new private IP subnet, and configure at least 2 nodes to join the management IP pool and set the pool mode to dynamic for HA IP address failover.

2. NOTE: Make sure the IP pool is set to System Access zone
3. NOTE: Vault cluster does not require a management IP pool facing the vault replication network

vii. **Vault and Production PowerScale SyncIQ replication IP pool**

1. Create a new IP pool in the default subnet, and configure at least 2 nodes to join the replication SyncIQ IP pool (static IP pool), for HA replication access to the vault PowerScale from the source PowerScale.
2. **NOTE: Make sure the IP pool is set to System Access zone**

viii. **Vault cluster Inside Airgap**

1. **If the inside Airgap solution is used an ethernet switch , ESX host and VM deployment are required for inside the Airgap automation.**
2. **Ransomware Defender vault VM requires the eyeglass minimum permissions configured on the vault cluster. See guide.**
3. **The vault cluster is added to the VM using the eyeglass service account.**
4. **SSH tunnel to a production cluster to allow communications with Ransomware Defender**

from within the vault. See the Enterprise license Vault Agent configuration in this guide.

2. **Source PowerScale**
    a. Create a new IP pool called "**Vault Replication**", and add at least 2 nodes, and 2 interfaces to this pool. No SmartConnect name is required for this pool. The pool must be in the system zone.
3. **Layer 3 Vault Replication switch**
    a. The network between the source PowerScale and the vault PowerScale will require a layer 3 device between the clusters. The Interface on the source PowerScale IP pool will have a static route with a next-hop of the layer 3 vault switch added, to reach the private subnet created on the vault PowerScale.
    b. **NOTE: This does not need to be a managed device and should be a statically configured routing device. It can be a larger switch using VLAN routing, but this exposes the potential for misconfiguration and allowing routing into the vault network. This is a business cost decision as a VLAN routing configuration can also be used.**
    c. **Best Practice: Use a dedicated switch with physical separation from production networks, and do not enable management of this switch, or leave the management port of the switch disconnected.**
4. **Ransomware Defender and Eyeglass steps**
    a. Vault policies are created on the source PowerScale with the a policy name prefixed with rw-airgap-xxxx where xxxx

can be any text to describe the policy, more than one policy can be created if required.

i. **SyncIQ Policy details:**

1. The source path should NOT be a path that is used as a DR replication path on your production cluster.

    a. Example: DR policy /ifs/data/zone1, the AirGap policy can use a policy path above or below this DR policy source path. i.e. above would be: /ifs/data, or below: /ifs/data/zone1/somepath.

    b. **Reason:** In a full re-sync recover from the Vault PowerScale, DR cluster mirror policies will cause an overlapping SyncIQ condition with 2 clusters trying to write data into the same path on the production cluster. This will block the Vault cluster policies from running successfully in a full recovery scenario.

    c. **Solution:** Avoid the overlapping condition by using non-overlapping paths when creating the AirGap policies. This may mean creating more policies to replicate all the data and to avoid the overlap with DR. This is the best option to avoid several manual steps in a full re-sync recover

scenario and will make recovery simpler, faster, and less complicated.

2. Schedule = set to manual (note Ransomware Defender manages the policy)
3. **Target host** - is the IP address of an IP on the vault PowerScale replication pool.
4. Restrict at source option enabled, and select the vault replication Pool created above to force replication traffic to use the vault pool node interfaces. This is the same pool that will have the static route applied for virtual Airgap mode.
5. Create the policy with the same target path used on the source path

b. Configure the policy replication schedule on the Eyeglass appliance, the recommended schedule is daily at midnight.

c. **Virtual Airgap mode**

   i. Add the static route to the Eyeglass Ransomware AirGap GUI.

   ii. The static route will be the next hop of the layer 3 vault switch, and target network will be the private network subnet created on the vault PowerScale replication pool.

d. **Inside the vault Airgap mode with Vault Agent VM**

   i. This requires the Vault Agent VM to be deployed on a dedicated ESX host that is secured inside the vault with the vault cluster.

ii. Configure inside the vault agent vm to connect to the vault cluster with minimum permissions user

iii. Add management IP pool to source cluster mapping information (see guide for more details)

iv. Verify Ransomware Defender reachability with test command to verify Airgap interfaces can be opened closed and remote API calls to the Ransomware Defender are functional.

e. Reporting requires no steps other than configuring email on Eyeglass to receive the daily Airgap sync reports.

f. Vault PowerScale alarm monitoring requires a service account user on the vault cluster. This simplifies management and monitoring of the Vault PowerScale if any hardware faults are detected. Alarm collection is completed during replication windows when the network is open. This means alarms will only be collected once a day if the replication schedule is daily.

Security Configuration of Components

The sections below outline additional security configuration that should be implemented when deploying the AirGap feature.

Eyeglass VM Security

1. Implement the hardening guidelines and password complexity and password management using this guide.

   a. Implement the fail to ban feature to auto ban and firewall failed login attempts to eyeglass using the hardening guide above.
   b. Configure 2 factor SSH on Eyeglass, ECA and the Vault Agent VM following this guide.
2. Firewall ssh and https access to the Eyeglass VM to management network jump box (administration VM) that has 2 factor authentication.
3. Using the ECA firewall requirements restrict ECA ports to only be authorized between Eyeglass and ECA and allow ssh access to the ECA from a network jump-box only. See firewall guide here.
4. Configure Role for AirGap management and configuration separately from Ransomware Defender management. see next section.
5. Full user UI access and configuration auditing covered in this guide.

Vault Cluster Configuration

1. NOTE: The AirGap Design and implementation service covers more hardening, the items below are the minimum changes that should be applied to the cluster.
2. Use only local accounts and no AD provider. This simplifies security infrastructure and ensure a lower attack surface to the device itself.

3. Enable Configuration Auditing to track all changes made to the cluster configuration.
4. Disable all non essential services
    a. SMB, NFS
    b. Delete default SMB share and NFS export
5. Disable all built in users accounts except for the root user
    a. The password should be a random password 20 characters or longer with upper case, lower case, numbers and at least 1 special character
    b. This password should be created and managed by the senior management within the security team and should not be shared with anyone outside of the security group.
6. Create the Eyeglass service account with minimum permissions for vault alarm collection, see minimum permissions guide.

Role Based Management of the AirGap Feature

1. The AirGap feature is added to the Ransomware Defender builtin role
2. This allows the AirGap management to be separate from day to day Ransomware Defender management. See example of dedicated role option that can be removed from the Ransomware Defender role and added to a custom role.

3. **Recommendation:** CSO or senior security management personnel should be assigned this role. The personnel with this role should be separate from the Ransomware Defender personnel.

4. 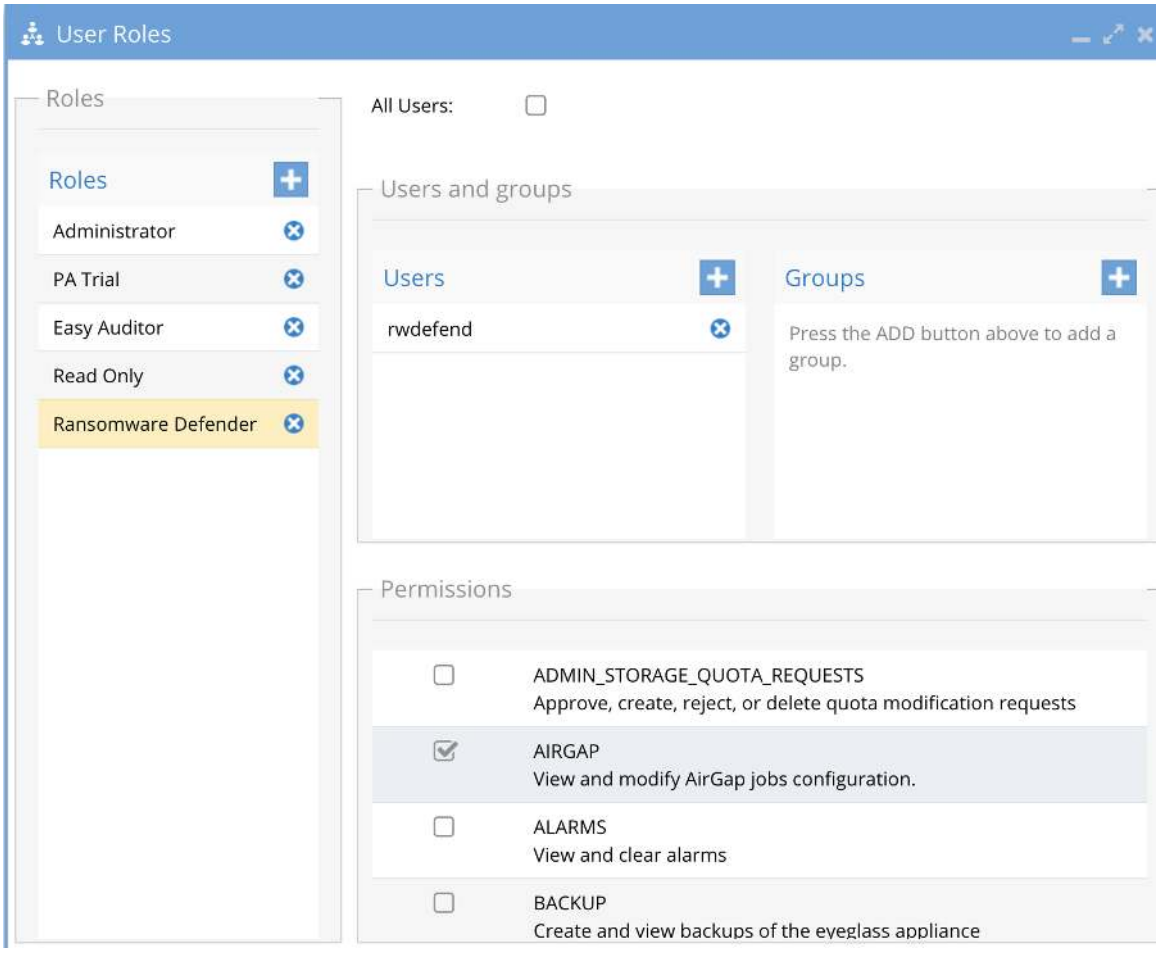

Detailed Deployment Diagrams

Network Considerations for Layer 2 or Layer 3 Vault Network

Overview

The vault network itself can be designed using layer 2 or layer 3 between the prod and vault clusters. The choices and best practices are as follows.

Layer 2 or Layer 3 - Fan-In cluster protection (Enterprise Airgap License)

1. **Highest Secure networking option** - Inside the vault VM agent with Enterprise Airgap license.
2. The Vault network the connects each protected cluster to the vault can use layer 2 flat vlan or a layer 3 network with routing between the protected clusters and the vault cluster.
    a. Example below is a layer 3 network example

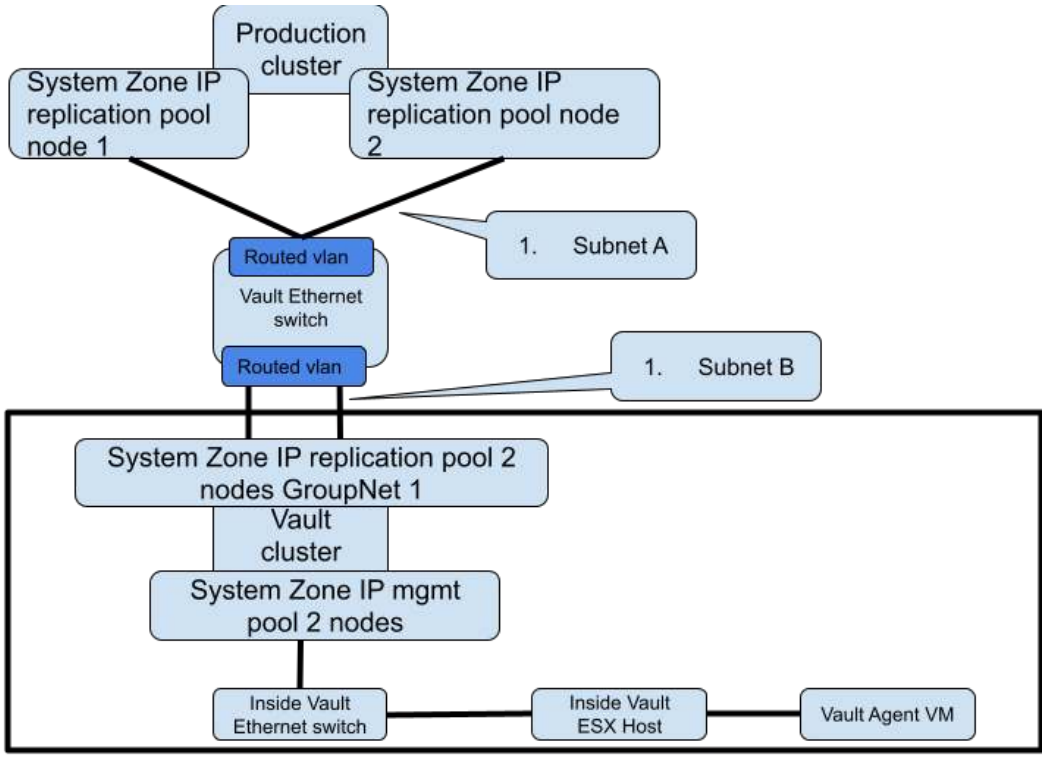

    b.

    c. Example below is layer 2 vault network

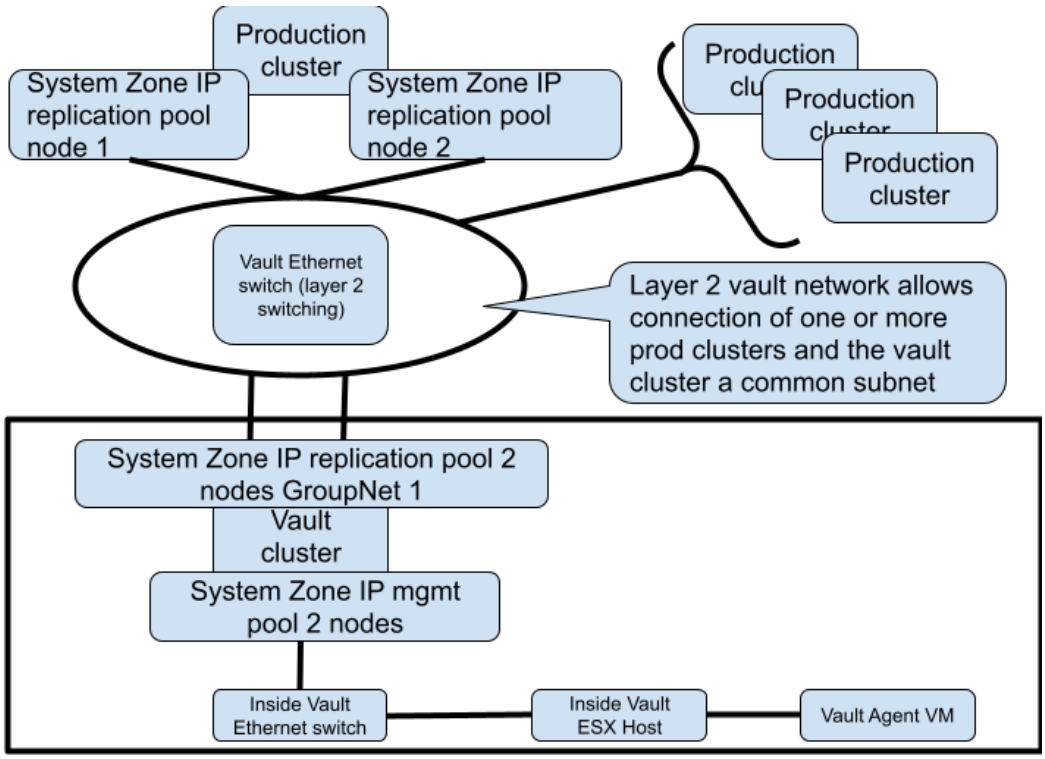

d.

3. **Best Practise:**

   a. For multiple source production clusters, a single vault network subnet allows all clusters to attach to a single layer 2 network and a single vlan can be used between all the clusters.

   b. A layer 3 vault network allows a firewall to be used between the source protected clusters and the vault cluster to add additional traffic firewall rules between the clusters.

1. Virtual AirGap Mode - Layer 3 Vault Network (Basic Airgap License)

a. The diagram below shows the networking required for source and vault PowerScale clusters, and how the vault switch and ip static routes should be configured for initial setup and configuration. The static route added on the source PowerScale will be added to the Ransomware Defender configuration to open and close the virtual AirGap.

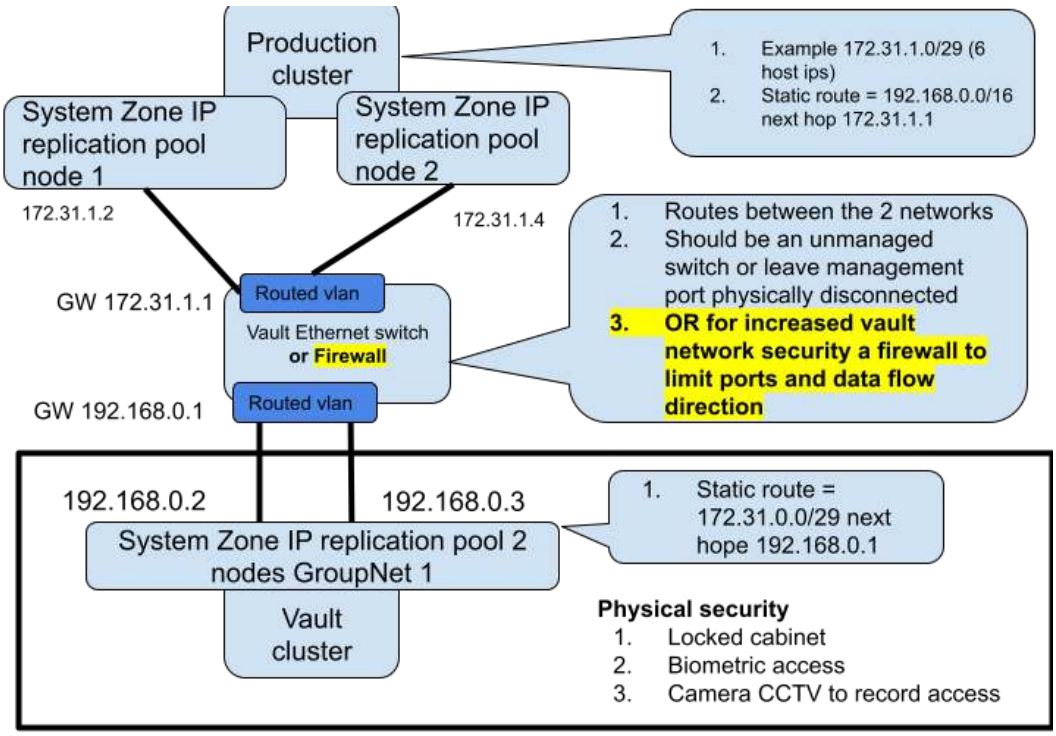

b.

2. Inside the Vault mode Deployment (Enterprise AirGap License Required)

Overview

This provides an alternate mode of operation with an inside the vault host and VM that opens and closes the vault from within the vault. This requires the Enterprise Airgap

license key. This is done by removing the replication interfaces from the IP pool which removes the IP address from the interfaces. This mode places a VM inside the vault and disables the IP stack that connects the vault cluster to any outside network. This mode offers Smart Airgap feature and all the same automation and enables this through a hardened Linux OS that autonomously manages the Airgap and verifies if is safe to replicate data. The slides below show how this mode is deployed.

Inside the Vault Physical Topology - Layer 3 example

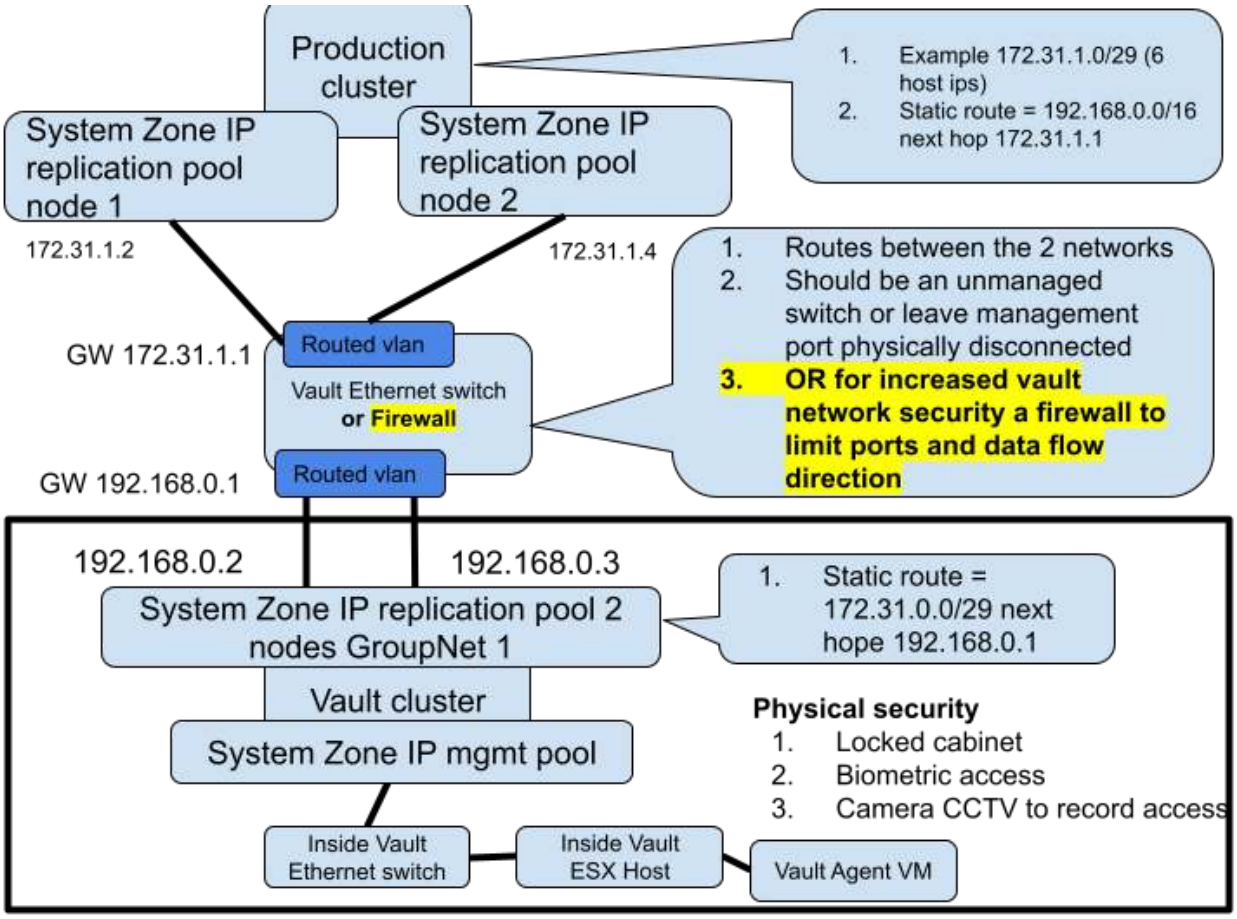

How to Enable Inside the Vault Agent VM (Enterprise AirGap)

1. Requirements:
    a. Airgap Enterprise agent VM license key is installed during deployment to enable this mode.
2. To enable Airgap policies to be managed by the secure hardened inside the Vault VM agent the Airgap administrator must switch from Virtual Airgap mode defaults to inside the vault mode.
3. The inside the vault VM agent will collect schedules configured in the Airgap UI and import them during initial setup.
    a. Once activated all vault open and close operations are managed by the Vault VM agent vm securing the vault access and shutting the vault networking if Ransomware defender or Easy Auditor triggers have alarms raised.
    b. **Smart Airgap** - Active alarms will cause the vault vm agent to shut the network down without replicating data.
    c. **NOTE: a static route will still need to be added, a fake route can be used that has no relevance to replication network. This requirement will be removed in a future release. example route 192.168.1.0/24 next hope 192.168.1.1 (NOTE: this assumes you are not using 192.168.x.x ip ranges)**
4. Open the Airgap Icon
5. Click on Settings
6. Enable "Managed by vault agent" check box and click save

7. 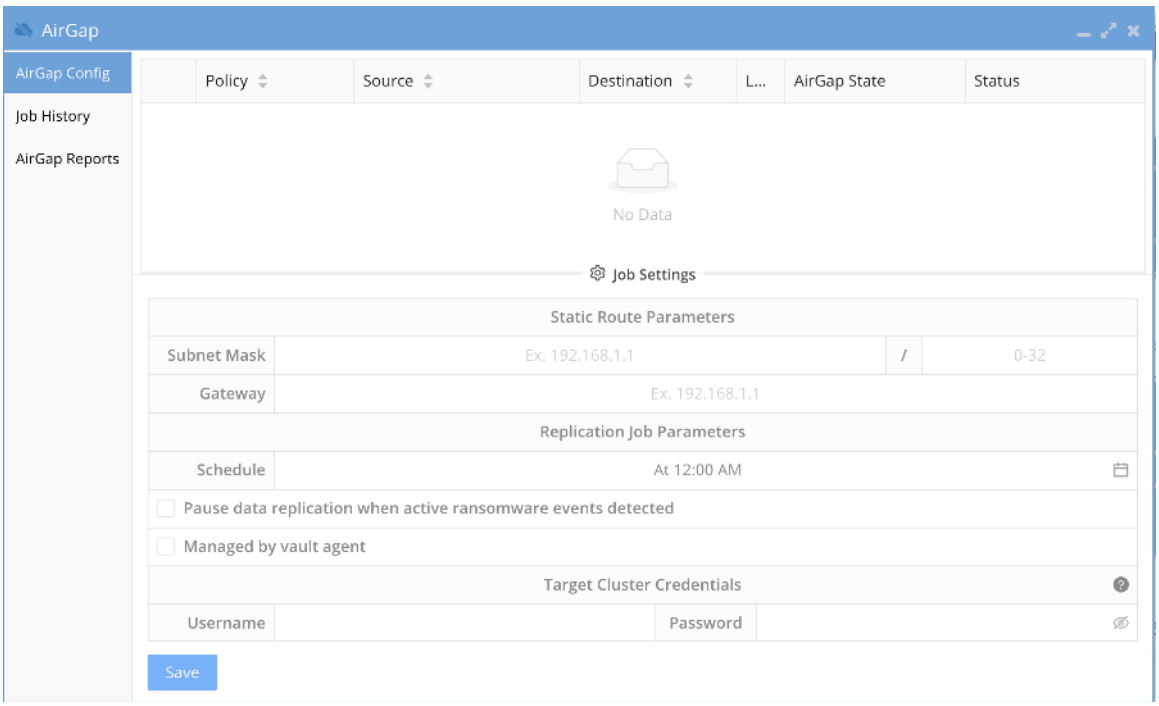

8. Done

Operations of Vault Data Replication

1. When using the inside the vault VM the CLI commands to force open the vault for maintenance are not supported and physical console access is required to gain access to the vault cluster or ESX host and VM. A physical keyboard mouse inside the locked cabinet will be required. This is a more secure operating mode.
2. 2 modes exist on the vault agent that allow a 2 hour maintenance heartbeat API request from the vault agent VM. This heartbeat API checks for a request for a maintenance access window using Airgap CLI command on the Eyeglass VM.

a. This is defaulted to disabled but can be enabled to check for a request every 2 hours to open the vault for maintenance for a timed window in minutes.

b. When the request is for 60 minutes the vault will close automatically after 60 minutes.

3. Ransomware Defender Smart Airgap API reachability failure is a fail safe for the vault. If the Eyeglass Smart Airgap API end point for safe replication cannot be reached the inside vault agent VM will fail safe and will close the vault on any failures keeping the data safe inside the vault.

Data Flow Example for Data Replication with Enterprise Airgap

1. 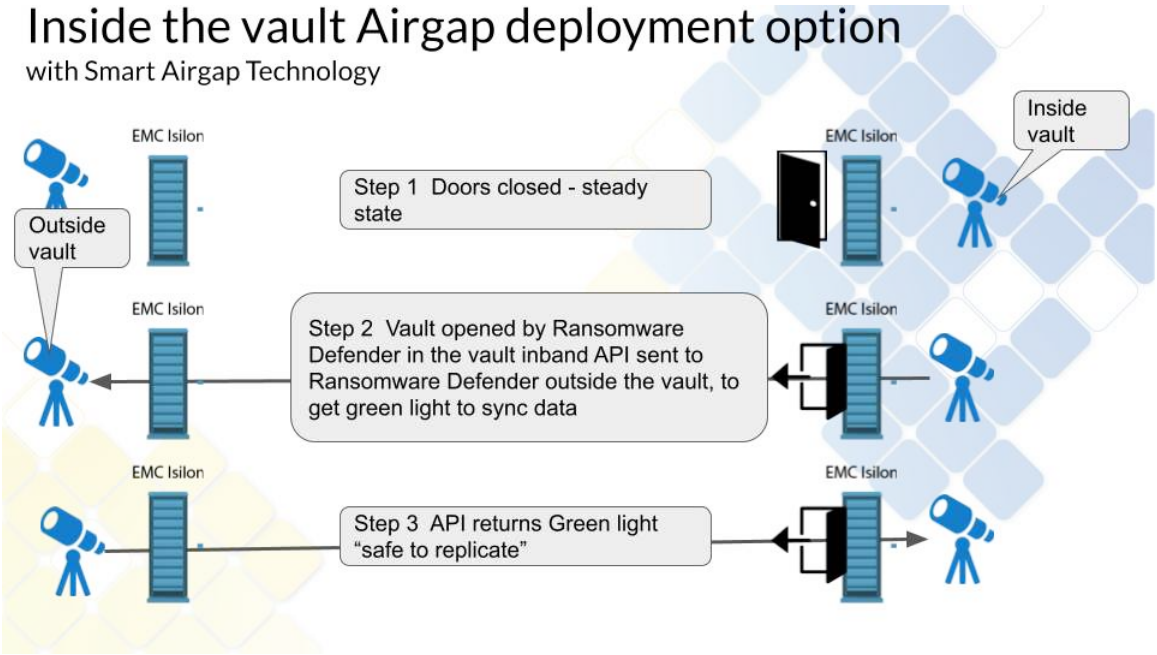

2. 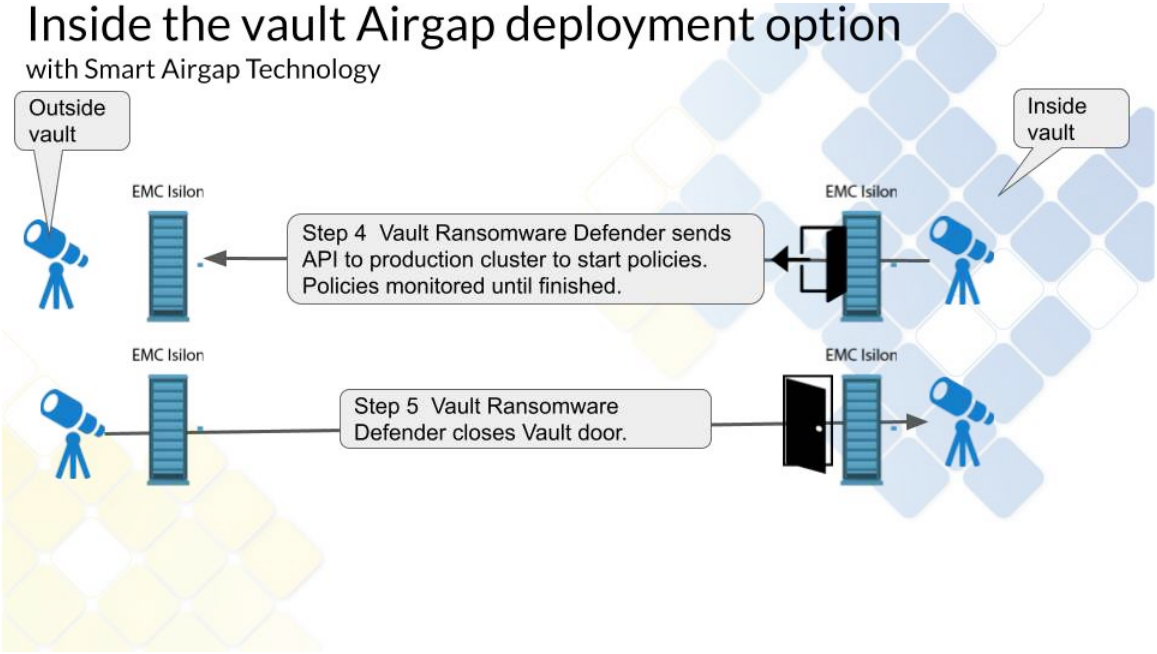

Vault Management Data Flow Example for Enterprise Airgap

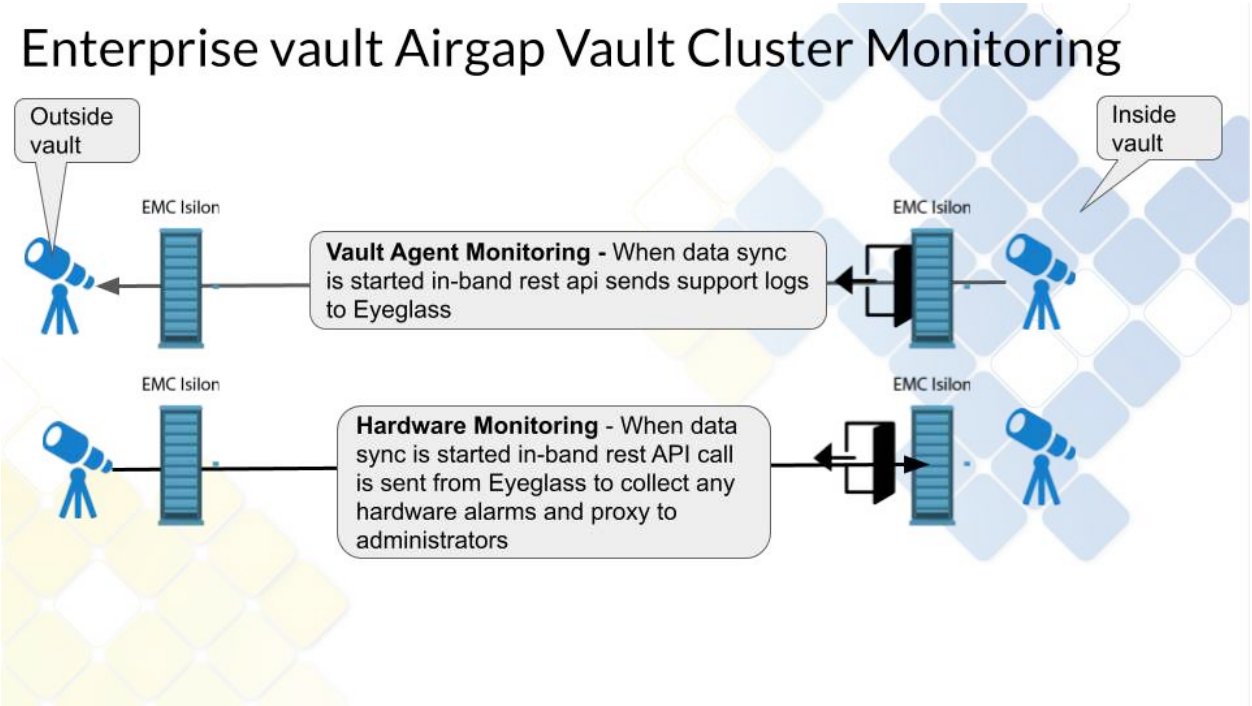

Configuration Steps for AirGap Setup

Overview video

How to setup syncIQ policies for AirGap

1. The AirGap policies are created on the Isilon and use restrict at source pool created in the physical configuration outlined in this guide. This ensures vault cluster replication traffic will use the correct nodes and physical interfaces. This also ensures Virtual AirGap to control the static route on this replication IP pool.
2. Select the source path based on your data protection requirements that select the data that should be protected in the vault. **NOTE: multiple policies can be created to protect different paths and change the replication schedule for each policy within the AirGap management GUI.**
3. The Name of the policy must use the following naming to be treated as an AirGap policy
    a. rw-airgap-xxxx where xxxx is unique part of the policy name. See Advanced settings to change the policy name prefix.
4. **Synciq Policy Property requirements**
    a. sync mode
    b. no schedule set leave at manual
    c. **Mandatory** - restrict at source pool set to the AirGap pool for synciq replication. This is required for for Basic Airgap to add the static route to the correct pool
5. **Data Retention**

a. This is an important consideration to provide maximum protection and options to recover data in a worst case data recovery scenario.

b. Longer SyncIQ data Retention will require more space with longer retention. Data change rates will determine how many days of retention.

c. When creating the policy enable Target Snapshots mode and set the retention in days. See example below.

d. 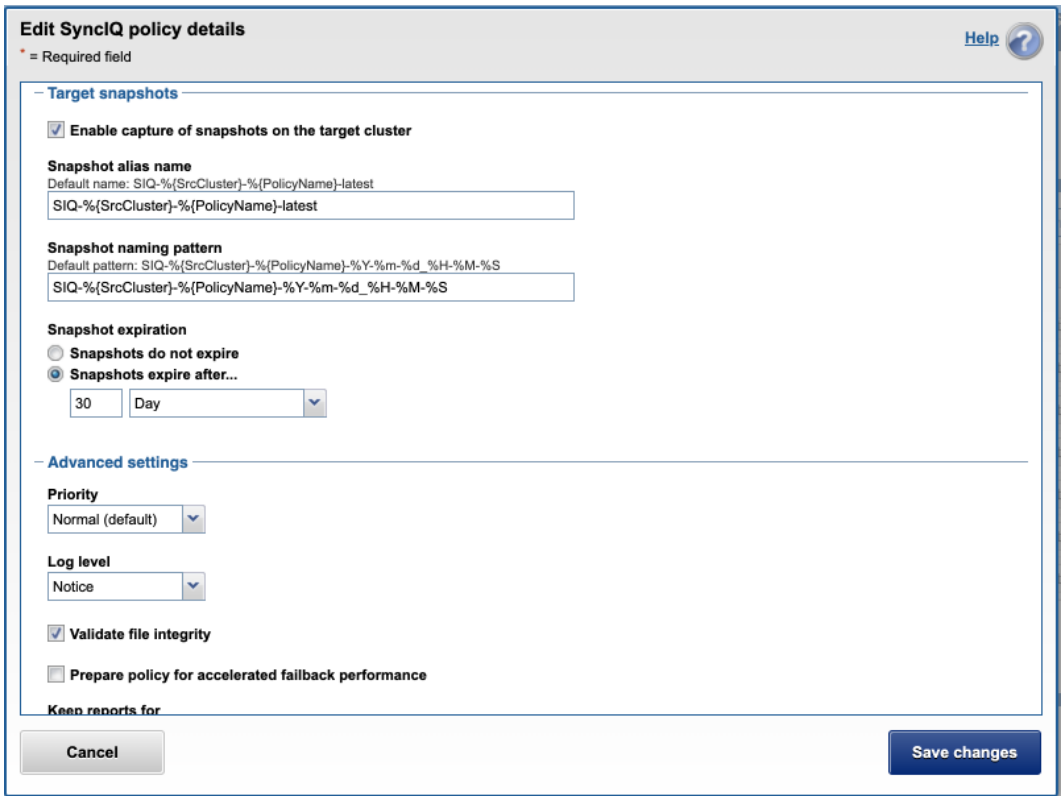

6. New policies will appear in the AirGap icon AirGap Config tab.

   a. Configuration replication inventory defaults to 5 minutes to detect new SyncIQ policies

7. Locate the policy in the Un-configured section of the jobs icon. The policy must be run once before it will move to the Airgap section of the jobs window.

a. 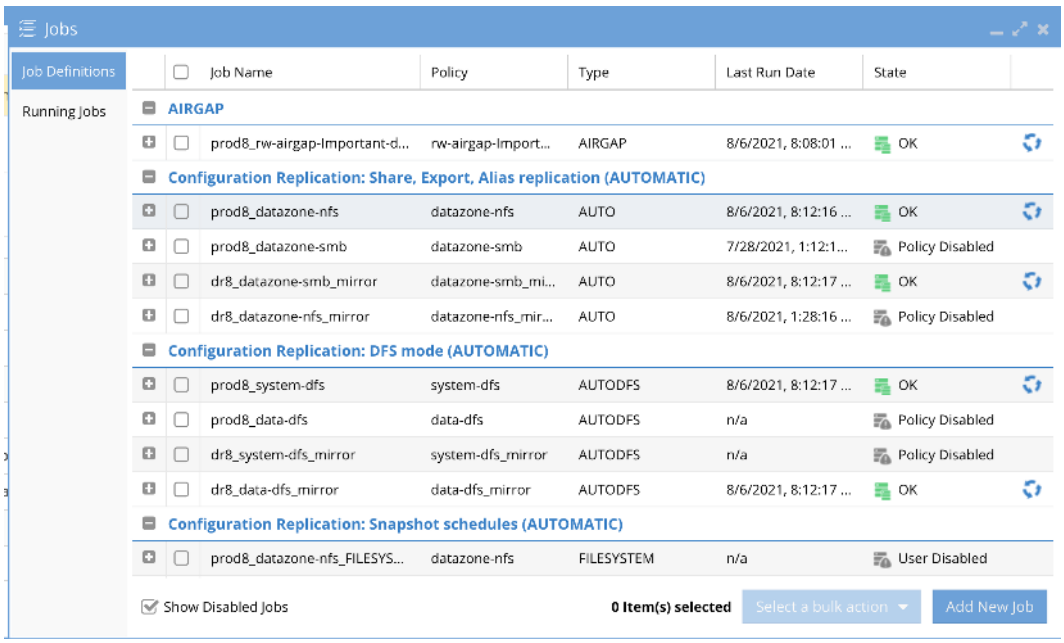

8. Verify the new policy appears in the AirGap icon

a. 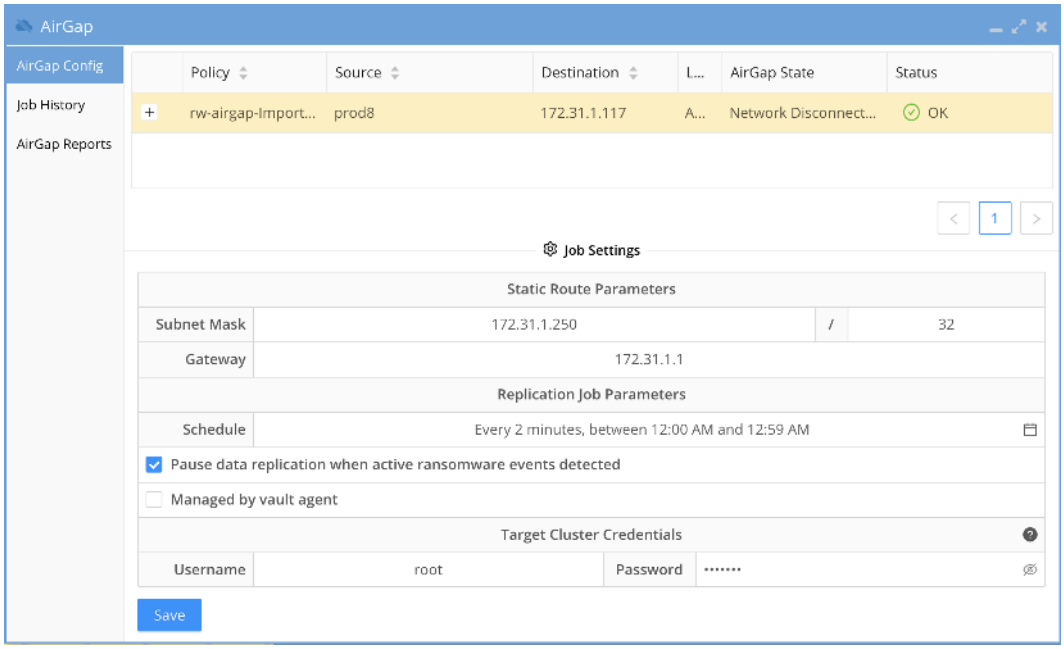

9. done.

How to Configure AirGap policies and setup Virtual AirGap (Basic Airgap License)

1. The SyncIQ AirGap policy should be configured as per the above section. Open the AirGap Icon to verify you can see the policy. NOTE: The schedule is not set and the policy is not managed in this state.

2. 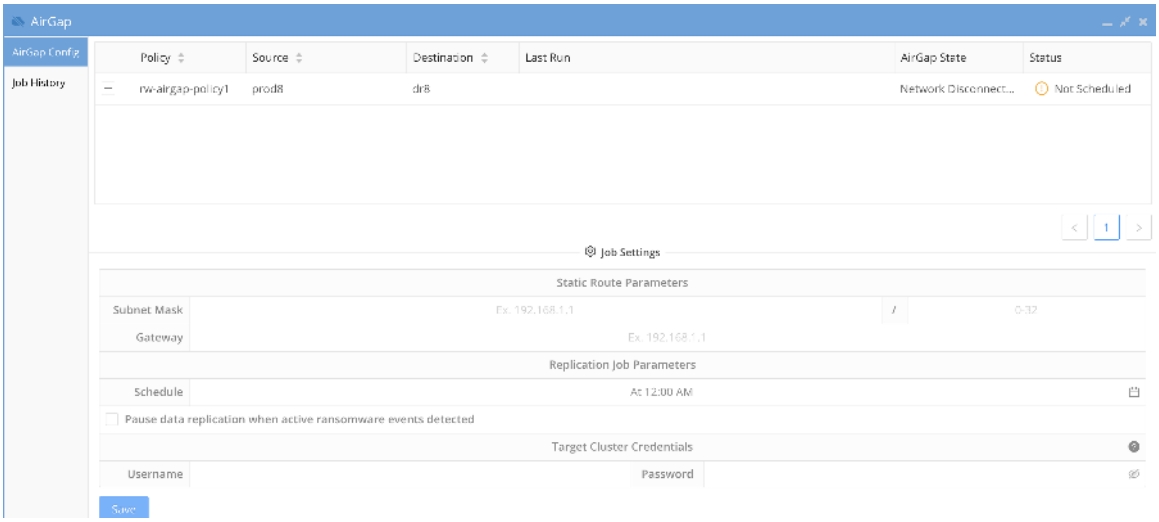

3. Enter the Virtual AirGap subnet range , network mask bits and next hop gateway.
    a. subnet = the network that the vault Isilon cluster IP pool that is configured for SyncIQ replication.
    b. The subnet mask bits to apply to the subnet entered. example 24 bits for a 255.255.255.0 subnet
    c. The next hop gateway IP address will be the IP address of a router between the production Isilon IP pool for AirGap and the Vault Isilon. Refer to the diagrams above on how to network the clusters together on a private network that is only reachable by the production cluster via the IP pool configured for the AirGap. See example vault cluster subnet of 192.168.0.0/24 and next hope of 192.168.1.1

d. **NOTE: You must enter a valid subnet that starts with the broadcast address for the subnet example 192.168.1.0/24 is the start of the subnet. An invalid subnet would be 192.168.1.1/24 since this does not include the broadcast address.**

e. 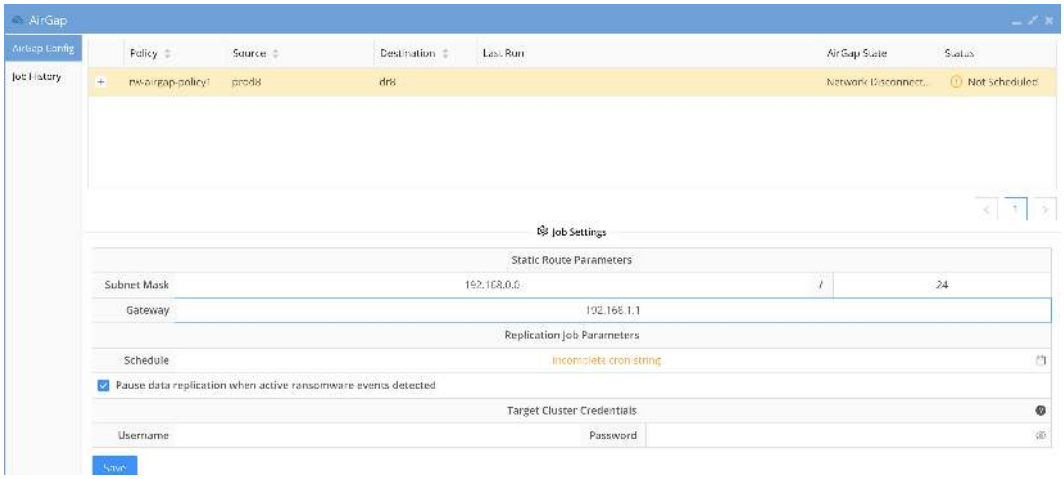

f. Next configure the schedule by clicking the calendar icon and completing the scheduling.

g. 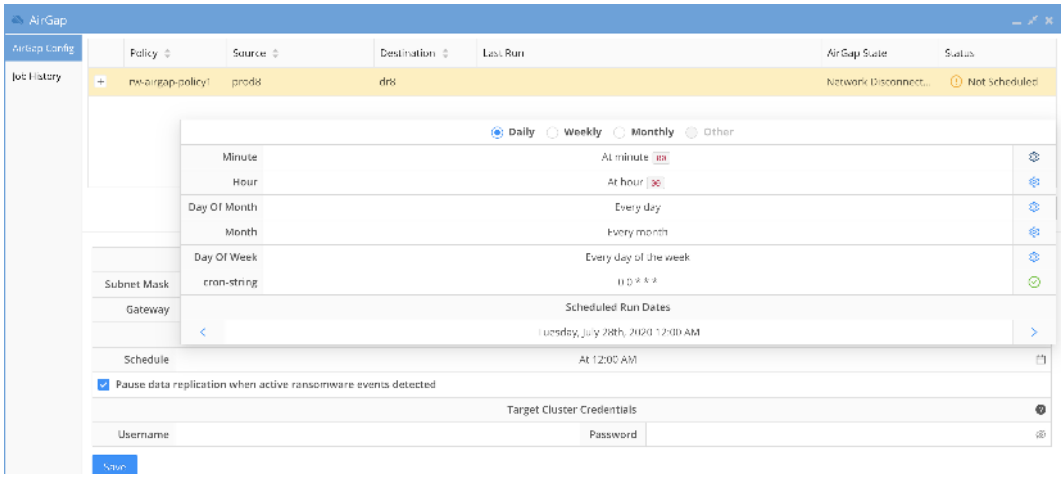

h. The defaults radio button at the top if the windows allows simply setup for daily , weekly, monthly

i. Select Other to have a custom schedule and complete all fields to complete the custom schedule.

j. **Recommendation: Always enable pause data replication when active ransomware events detected. This is the Intelligent data protection option that overcomes limitations on other backup based cyber vaults that allow encrypted comprised data to be copied into the vault.**

   i. Note the check box "**Pause data replication when active Ransomware events detected**" This enables Smart AirGap mode that will monitor user behaviors for any activity that could be considered Ransomware this includes warning, major or critical detections.

   ii. If these alarms are not cleared or managed as resolved in Ransomware Defender Icon the copy schedule will be skipped until an administrator makes addressed the alarms.

   iii. If Easy Auditor is installed all Active Auditor trigger active alarms will also block replication to the vault and must be cleared to allow replication.

      1. DLP, Mass Delete or custom triggers all block vault replication.

      2. Suggested Configuration to enable a honeypot trigger to monitor snooping of open SMB shares. See the guide.

k. **Best Practices: This option should always be enabled to offer the highest protection level of your data and ensures no copies are stopped until an administrator makes a decision on the events. When the events are cleared by an administrator AirGap will resume copies on the next**

schedule incremental update schedule. Consult support if you plan to disable this check box. If disabled the schedule will run regardless of what alerts are present in Ransomware Defender.

l. Target cluster Credentials

m.

n. The user and password should be the service account created above section for vault cluster configuration. This service account is a minimum privilege user to collect alarm data only.

   i. These credentials are used to retrieve alarms from the vault Isilon in-band while the AirGap is open and proxies alarms on the Vault cluster to administrators to monitor physical hardware issues that may occur.

   ii. This ensures an automated solution that is lights out at all times to secure the vault data.

o. Then submit the save button and the state should now change to show next schedule replication and AirGap state

p. 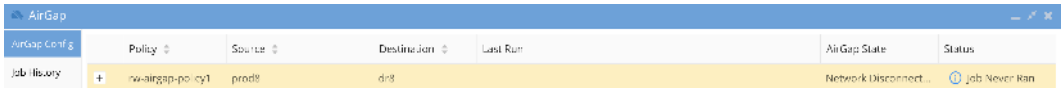

q. The AirGap policy is now in production mode.

How to test an Airgap Policy job

1. Open the Jobs icon and click the run now icon to start the job and then monitor the job from job history tab of the airgap Icon.
2. 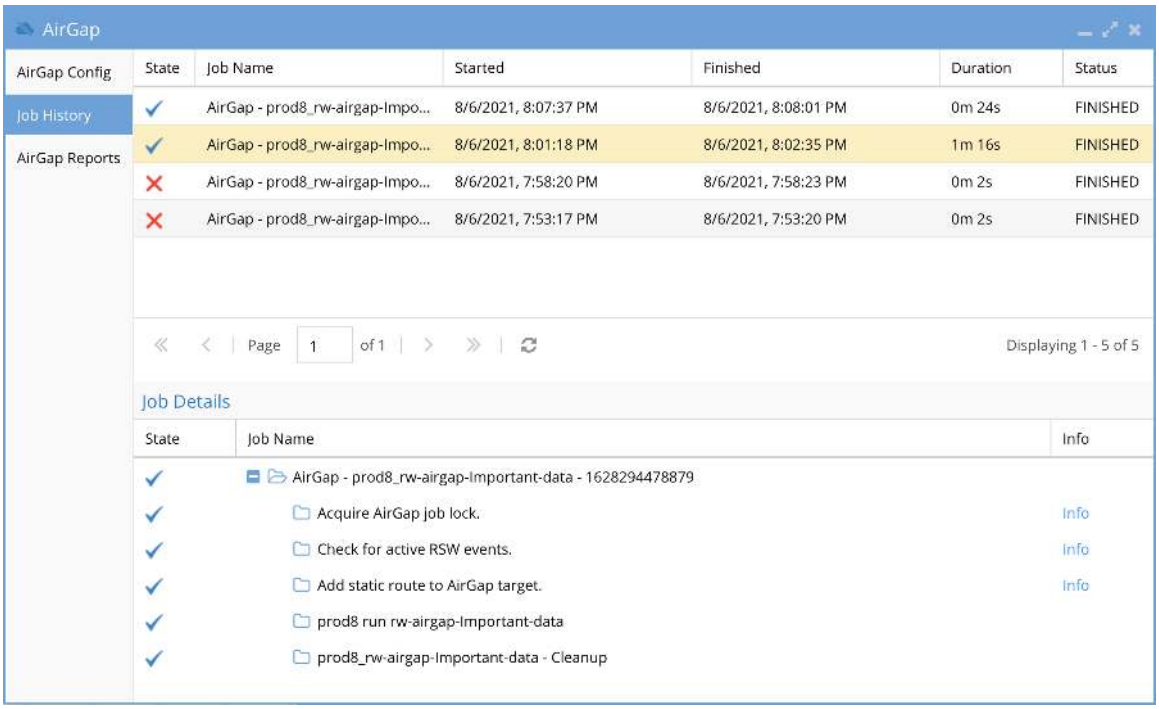

How Alarms from the vault Isilon are Viewed and Forwarded

1. Configuring the Target cluster credentials allows remote alarm collection during incremental AirGap copies using the in-band replication network to collect alarms.
2. Alarms are forwarded through email only and will not display in the Active Alarms icon that is reserved for Eyeglass alarms only. The history alarm will display on the Managed Cluster Alerts tab of the Alarms Icon.
3. Sample email proxy alarm

Alarm Report 2020-07-21 10:34:10 EDT - [AIRGAP-PROXY] One or more drives (location(s) Bay 7, Bay 8, Bay 9, Bay 10, Bay 11, Bay 12, Bay 13, Bay 14, Bay 15 / type(s) HDD, HDD, HDD, HDD, HDD, HDD, HDD, HDD, HDD) are not healthy. Alarm Severity: CRITICAL Inbox ×

scheng@superna.net
to me

10:34 AM (6 minutes ago)

Alarm Report 2020-07-21 10:34:10 EDT

Source	IP Address	Alarm Code	Sub-System	Time Raised	Description
1	172.25.28.200	AIRGAP_PROXY		2020-07-13 11:28:16 EDT	[AIRGAP-PROXY] One or more drives (location(s) Bay 7, Bay 8, Bay 9, Bay 10, Bay 11, Bay 12, Bay 13, Bay 14, Bay 15 / type(s) HDD, HDD, HDD, HDD, HDD, HDD, HDD, HDD, HDD) are not healthy.

More information about this alarm

Alarm Extra Info 2020-07-21 10:34:10 EDT

{"AirGap target IP":"172.25.28.200","AirGap target name":"SC-9000C","AirGap target version":"v9.0.0.0","AirGap target GUID":"0050569f3945f9780c5f320b0b77caad2640"}

4.
5. To route alarms to a specific email address use the Eyeglass custom email routing guide here.
6. Example Tab in Alarms
   a. 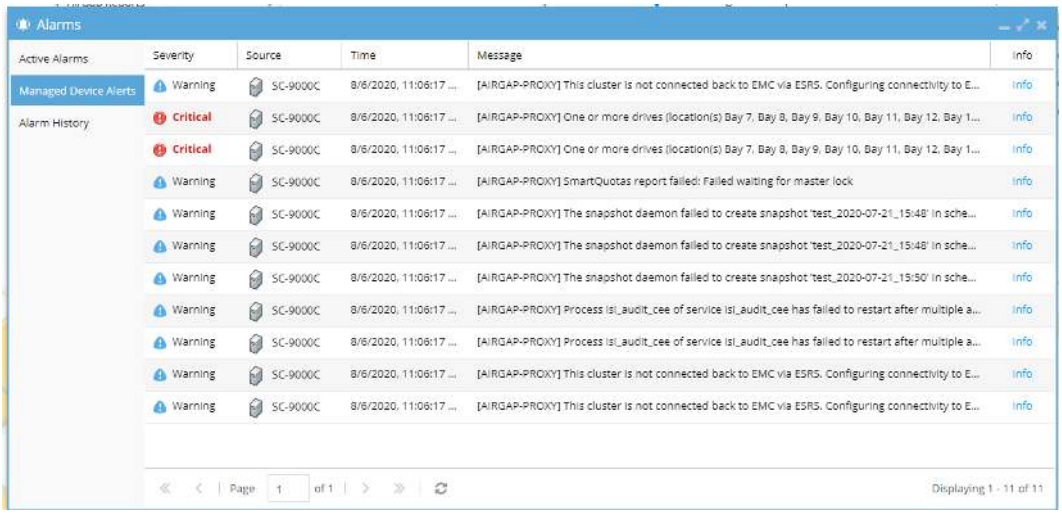

How to Expand the Airgap Sync Job Timeout and the Airgap job prefix name

1. The default timeout is 240 minutes or 4 hours and will fail a sync job that takes longer. This only applies to incremental syncs. These steps can also be used to change the default prefix that is used to match the airgap synciq policy as an airgap policy.
2. To change this timeout value
   a. On the eyeglass vm login as admin

b. nano /opt/superna/sca/data/system.xml

c. Add an airgap section with tags as per below and change the policy prefix value and or the timeout value in minutes.

d. Save the file with control+x and answer yes to save and exit.

e. <airgap>

   <policyPrefix>rw-airgap-</policyPrefix>

   <logsMaxAgeInDays>7</logsMaxAgeInDays>

   <airgapJobTimeout>240</airgapJobTimeout>

   </airgap>

Operational Procedures for AirGap Management

1. After the initial configuration, running the AirGap policies manually will start the large first full sync of the data. This can be done from the Onefs GUI SyncIQ tab.
2. Monitor the initial data sync phase, and then enable AirGap on Ransomware Defender to take over the sync schedule and manage the AirGap replication automatically.
3. **Day to Day Administration**

   a. The Vault PowerScale is monitored in-band by Ransomware Defender to collect alarms. This allows administrators to monitor the vault PowerScale without needing to expose the vault PowerScale to the external

network. When the AirGap is open to sync data, the in-band management is done over SSH from the production PowerScale to the vault PowerScale.

How to stop AirGap Replication in an Emergency

1. If you suspect you IT environment has been compromised in any way it is important to shutdown the AirGap permanently to protect the AirGap copy of the data.
2. See the AirGap CLI command in the CLI guide here.
    a. ssh to Eyeglass as the admin user and issue this command to disable and isolate the AirGap data.
    b. igls AirGap disable

How to monitor replication AirGap policy success failure

1. Login to Eyeglass and open the AirGap icon and click on the Jobs History tab to review the history of the replication jobs
2. 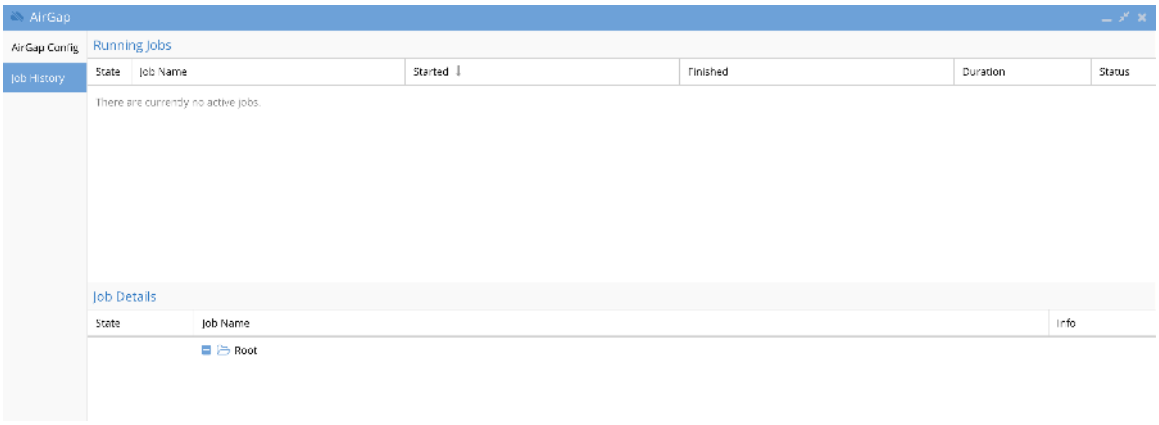

How to Monitor AirGap Replication Reports

1. The SyncIQ jobs are managed and reported on by a dedicated AirGap report. Setting up report notification in notification center with an email recipient set to reports , will receive the AirGap replication report. Consult Eyeglass admin guide on how to configure email and recipients.

How to enable or disable the Airgap daily summary report or change the schedule

1. igls admin schedules list **(to check the current schedule)**
2. igls admin schedules set --id AirGapReportsTask --interval 7D **(to change schedule to every week)**
3. igls admin schedules set --id AirGapReportsTask --enabled false **(to disable the report)**

How to pause all AirGap policies to complete Vault cluster maintenance

1. This mode should be used to complete network or vault cluster maintenance and stops policy replication
2. See the AirGap CLI command in the CLI guide here.

How to Pause the AirGap policies for maintenance with a timed auto close of the AirGap Network

1. This option uses the igls AirGap connect, and disconnect command and operate separately on a per policy basis and sets a timer to keep the AirGap network open for X minutes or hours.

This ensures the AirGap network is not left open accidentally and automatically closes the AirGap network after the timer expires.

2. See the AirGap CLI command in the CLI guide here.

How to Configure Enterprise AirGap Ransomware Defender Enterprise Airgap Agent

Overview

This section covers how to configure the vault agent VM on the ESX host that is deployed inside the secure vault. This VM manages the vault cluster and orchestrates all replication from protected source clusters. The vault agent uses a secure SSH tunnel from the vault cluster to a source protected cluster to reach Ransomware Defender VM to send secure messages to orchestrate replication tasks, upload logs, download new policies or protected clusters and updates to schedules configured in the AirGap UI.

Topology and Communications

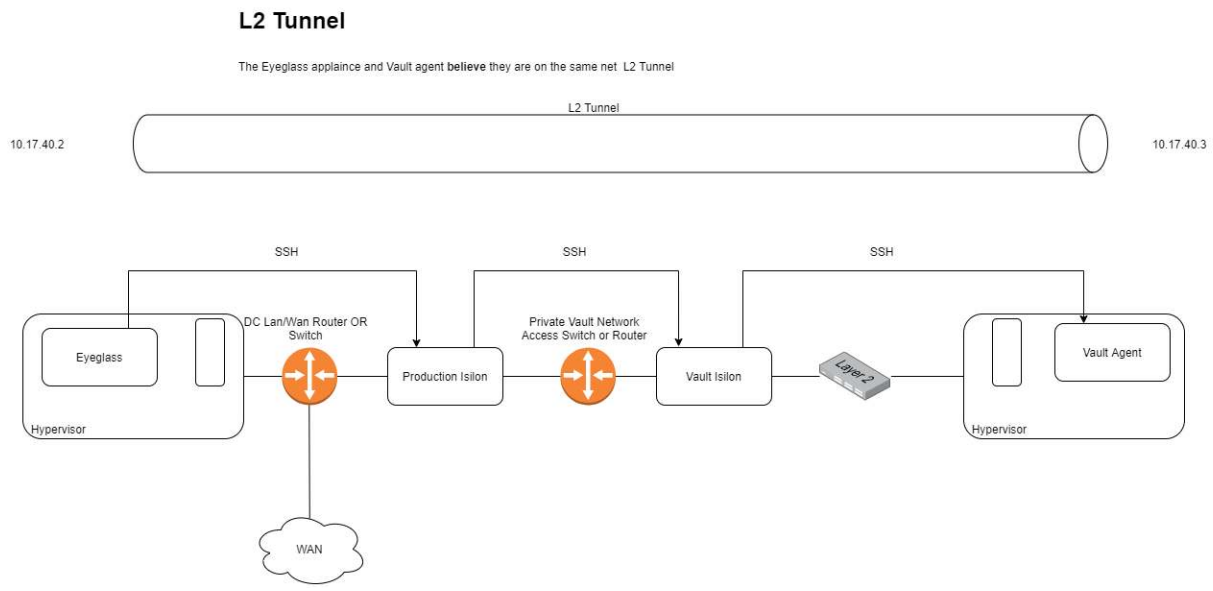

Prerequisites

1. Enterprise Airgap license
2. ECA single VM deployed
    a. After first boot run this command
    b. ovf set-value -f mode=vault-agent
    c. How to startup the software
    d. ecactl cluster up

Configuration Steps

1. Install the Eyeglass Vault Agent (EVA) license in Eyeglass
    a. Login to Eyeglass, open the license manager Icon and click upload new license zip file. This license is required to enable the managed by vault check box in the Airgap icon.

b. 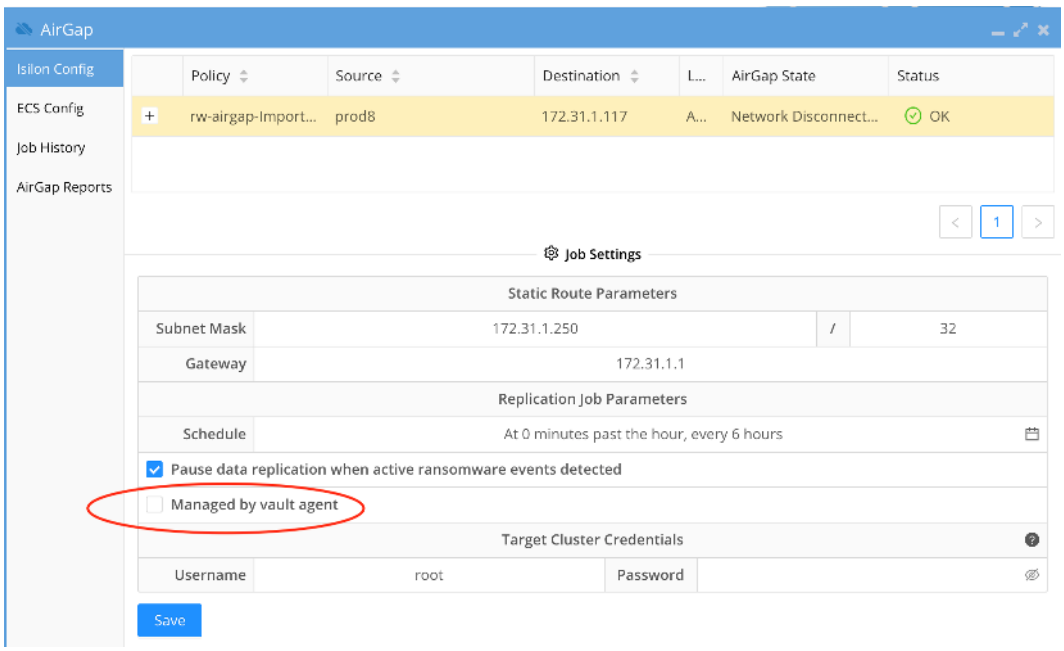

2. Configure Keyless ssh from Vault cluster to each protected source cluster to allow an ssh tunnel to be created for communications between the Vault Agent and the Eyeglass VM.
    a. NOTE: The minimum permissions user eyeglass should be created on all clusters that are protected or the vault cluster. The minimum permissions guide lists the ISI commands to verify have been applied to the eyeglass user.
    b. Login to the **protected** cluster that will be used for the SSH tunnel as the **eyeglass** user.
        i. mkdir .ssh
        ii. procedure done
    c. Login to the **vault** cluster over ssh as the **eyeglass** user
        i. create an ssh key pair
            1. run this command: ssh-keygen -t rsa
            2. Hit Enter for default path

3. Hit Enter for passphrase
4. An ssh key pair should be created in /ifs/home/eyeglass/.ssh
5. Copy public key (id_rsa.pub) to the primary protected cluster that will be used for Eyeglass communications.
    a. scp /ifs/home/eyeglass/.ssh/id_rsa.pub eyeglass@x.x.x.x:/ifs/home/eyeglass/.ssh/id_rsa.pub

ii. Complete protected cluster keyless SSH configuration

1. ssh as eyeglass user to the production cluster
2. cd .ssh
3. cat id_rsa.pub >> authorized_keys
4. chmod 600 authorized_keys
5. done

iii. Test keyless SSH from vault to production cluster

1. ssh to the vault cluster as eyeglass user
2. ssh again to the production cluster
3. If no password is requested then keyless ssh was successful, if a password prompt is presented it means a step was missed and review all steps above were completed.

iv. done.

3. Add the vault cluster IP and ip address for the SSH tunnel to this cluster. (Requires license key applied)

a. ecactl isilons add --vaulthost x.x.x.x --user eyeglass --vaultPoolName groupnet0.subnet0.synciq --vaultsynciqexternalInterface 1:ext-1, 2:ext-1, 3:ext-1, 4:ext-1

b. **--vaulthost** - isilon management IP address in system zone x.x.x.x is the system pool IP address on the inside of the vault.

c. **--vaultPoolName** - the IP pool on the vault cluster used to receive synciq data from a protected cluster

d. --**user** - service account created on the vault cluster for the vault agent VM

e. --**vaultsynciqexternalInterface** - This is the list of interfaces in the synciq pool. enter the node and interface name in a comma separated list example 1:ext-1, 2:ext-1, 3:ext-1, 4:ext-1

4. Add a protected cluster to the EVA VM to create the secure tunnel to reach Eyeglass Ransomware Defender

   a. NOTE: This cluster must be reachable over the vault cluster synciq IP pool interfaces configured in the vault cluster add command.

   b. NOTE: This cluster will be used for all communications via a secure SSH tunnel from the vault cluster and will be used to send api calls to Ransomeware Defender.

   c. ecactl isilons add --protectedhost x.x.x.x --user eyeglass

5. Test API communications over the secure tunnel to a protected cluster

a. NOTE: This tests the ssh secure tunnel from the vault cluster to the named protected cluster and issues a test api to Ransomware Defender to verify end to end communications.
b. ecactl isilons list (to get the cluster name from the add command)
c. ecactl airgap check --prod <protected cluster name>
d. Example output
e. Opening vault connection..

   Command succeeded

   Running command on vault.. whoami; hostname

   eyeglass

   Prod-cluster

   Running command on prod.. whoami; hostname

   eyeglass

   Prod-cluster

   Running command on eyeglass.. ' 'https://172.25.49.15/sera/v1/healthcheck' -k

   "\"Wed Aug 18 07:16:00 EDT 2021\""

   Closing vault connection

   Command succeeded

   DONE!

6. Modify a cluster configuration

a. ecactl isilons modify --name ISL-EASEE-8-2-1-0-172-25-47-73 --vaultsynciqexternalInterface 1:ext-1, 2:ext-1, 3:ext-1, 4:ext-1, 5:ext-1

7. List clusters of type vault and protected

   a. ecactl isilons list

8. Remove a cluster

   a. ecactl isilons remove --name ISL-EASEE-8-2-1-0-172-25-47-75

9. List Synciq Jobs between the vault and a remote cluster

   a. ecactl airgap list

10. Retrieves the airgap policies configured in Ransomware Defender, retrieves the schedules configured for each policy and saves this information locally to run on a schedule. NOTE: New airgap policies and schedule changes are checked each time the vault opens to run a job. A secondary schedule can be configured to check for new configuration or schedule changes independently of scheduled airgap sync schedules.

    a. ecactl airgap schedules

11. Run an Airgap job on demand to test an incremental sync of data into the vault.

    a. NOTE: Use the ecactl airgap schedules command to get the names of remote cluster synciq policies configured within the Airgap Icon in Eyeglass. This will retrieve from all Ransomware Defender managed clusters.

    b. ecactl airgap runjob --job ISL-EASEE-8-2-1-0-172-25-47-73_rw-airgap-test3

i. --job - this is the name of the policy returned from the list command above.

Operational Procedures Enterprise Airgap

How To reach outside the vault through the vault cluster it is possible to open and close the vault with cli commands

1. ecactl vault open - this will open the vault interfaces from the IP pool
2. ecactl vault close - this will close and remove the vault interfaces from the IP pool

How to Open the Airgap for maintenance from Ransomware Defender CLI

1. **NOTE: This command is default disabled on the vault agent and must be explicitly enabled. This is for doing maintenance or temporary access to the vault and then disabling this feature after work is completed. The vault agent will Not open the vault be default even if this command is used.**
2. **In order to enable this feature on the vault agent.**
    a. **nano /opt/superna/eca/eca-env-common.conf**
    b. **add this variable**
    c. **export EYEGLASS_OPEN_VAULT_ENABLED=true**
    d. **save the file control+x**

   e. **ecactl cluster down**
   f. **ecactl cluster up**

3. igls airgap vaultaccessrequest --interval=x (request to open the airgap for x minutes, after this time the airgap will auto close, the heartbeat check for pending requests is every 2 hours by default and will open the airgap for x minutes only once the pending request has been seen)
4. igls airgap vaultaccessview (view pending requests to open)
5. igls airgap vaultaccesscancel (cancel a pending request)

How to list running jobs

1. ecactl jobs running

How to run an Airgap job from the Vault agent VM

1. ecactl airgap startjob --job <job name> (use ecactl airgap list to get job names)

How to monitor a running airgap job

1. ecactl jobs view --follow --id job-1630432432546-879575052 (replace with job name)

How to check the remaining time of a maintenance window request on the vault agent

1. Use this command if you have requested a timed maintenance window from the eyeglass vm. This command runs on the vault

agent. **NOTE: The maintenance window time will survive and vault agent upgrade or cluster restart.**

2. ecactl airgap checkopen

   a.
```
ecaadmin@vanew-1:~> ecactl airgap checkopen

Time remaining for vault close: 13 minutes, 3 seconds.
```

How to configure Vault cluster Log Gather Automation for Hardware Support

1. These steps enable automation to allow automation to collect a log gather and place this on the production cluster to allow Dell Support the ability to verify the health of the vault cluster, if any alarms are proxied by the vault agent through Eyeglass managed devices.
2. **Requirements**:
   a. 2.5.8 build 228 or later
3. **Configuration**:
   a. ecactl airgap startisilonloggather -> to start the job now
   b. ecactl airgap isilonloggather -> to read the schedule
   c. **Recommended for all deployments**
      i. ecactl airgap isilonloggather --setschedule "0 0 * * *" => to set the schedule for the job every day at midnight
4. Logging will output location of the log gather gz file

a. Starting the vault gather job. Will find it under the production cluster /ifs/data/home/eyeglass/IsilonLogs-<vault_cluster_name>.tgz
b. The file can be provided to Dell Support to monitor or investigate vault hardware cluster issues.

5. How to change production cluster path for the log gather in eca-env-common.conf
   a. export EVA_VAULT_LOG_GATHER_PATH_ON_PROD=/ifs/xxxx (change xxxx to path to store the log gather in a different location)

Advanced Vault Agent and Airgap Eyeglass Configurations

Scheduled check for new or changed Airgap policies

1. Configuration check schedule can be changed using a variable in the conf file. This scheduled task will open the vault and download new policies or schedule changes and then close the vault. The vault will only be open for a few seconds.
   a. on the vault agent VM
   b. nano /opt/superna/eca/eca-env-common.conf
   c. add this line with a cron string interval. This is a 5 minute example.

d. export TASKMASTER_AIRGAP_SCHEDULING_CRON="*/5 * * * *"

e. control + x to save

f. ecactl containers restart taskmaster

How to change the name of the Airgap policies.

1. Use this procedure to change the name of the airgap policy prefix
2. Login to eyeglass vm over ssh
3. nano /opt/superna/data/system.xml
4. add the tag below in a new section
5. <airgap><policyPrefix>rw-airgap-</policyPrefix></airgap>
6. Edit the yellow value to a string that will prefix all airgap policies.
7. AIRGAP jobs are listed in Jobs window and AirGap Config'

Security

Airgap Audit log

1. The Eyeglass VM audits the reachability of the vault cluster and logs this information to a dedicated log.
2. This log can be found on Eyeglass below, the check is run every 5 minutes.
3. cat /opt/superna/sca/logs/AirgapAudit.log

Recovery Scenarios

Depending on the recovery requirements, the following three scenarios describe the high-level steps to gain access to the vault data, that are possible with the AirGap 2.0 solution.

Considerations

1. This procedures in this section should never be started until the all clear has been declared. This would require all supporting infrastructure (AD, DNS etc..) with no threat present in the environment. CSO or similar role in your organization should not request this procedure start until the environment has been cleaned of all active threats. The risk that this final copy is comprised due to active threat still present in the environment.

Partial Vault Data Recovery Scenario

1. In this scenario, you require access to some of the data in the vault, due to an issue with the production data affecting a subset of the data protected by the vault. The PowerScale vault makes this very simple.
    a. **Method #1- Windows Explorer + SMB Share**

i. Physically connect the PowerScale management interface to the vault Ethernet switch management VLAN or port.

ii. Connect to the PowerScale WebUI, and login.

iii. Enable the SMB protocol on the cluster. Protocols tab, enable SMB protocol.

iv. Create an SMB share on the path that stored the data you require. Since there is no AD provider on the vault cluster, a local user will be needed to authenticate to the SMB share. The admin user can be used for this authentication.

v. **NOTE: The data is locked by SyncIQ and immutable while connected to the network.**

vi. Copy the data from the recovery SMB share using a Windows PC connected to the production cluster folder location. Using Windows Explorer, copy the files or folders to the production cluster.

vii. Once the data restore is completed:

1. Delete the SMB Share.
2. Disable the SMB protocol.
3. Disconnect the Management port Ethernet cable.
4. Done.

b. **Method #2 - SCP in-band copy**

i. This method is best when a path of data needs to be restored and may require a long copy process. **NOTE:**

SSH needs to be enabled on the vault cluster this may have been disabled with hardening steps. This will prevent this method from being used.

ii. SSH to the production cluster.

iii. Open the AirGap - Use the Ransomware Defender CLI command to open the AigGap with the timed open command. This command will open the AirGap for X minutes and will automatically close when the timer expires.

iv. Remotely copy data from the Vault PowerScale to the production PowerScale using the SCP command. Use the secure in-band replication network to copy data.

v. This is syntax of the command that needs to be customized for your environment:

   1. scp -rp <**user**>@**x.x.x.x**:/ifs/data/**yyy**/* /ifs/data/**ccc**
   2. The user should be the admin.
   3. x.x.x.x is the ip address of the remote vault PowerScale replication pool IP address (pick any IP address in the pool).
   4. yyy - is the path on the remote vault PowerScale that contains the data to be restored.
   5. ccc - is the location on the production PowerScale where the data should be copied.

6. NOTE: -rp means recursive copy and -p preserves the date stamps on the files when copied.

vi. NOTE: The timed AirGap open command should be set long enough to ensure the copy completes. Monitor the SCP command progress until it completes.

vii. The AirGap will auto close after the timer expires.

viii. Done.

Complete Vault Data Recovery Scenario

1. This scenario covers recovery of all the data on a production cluster protected by the AirGap SyncIQ policies. This would be a worst-case scenario where the data in the Vault is determined to be the best copy of the data to be used for recovery.
2. This scenario assumes that the production PowerScale itself is in a usable state, and simply needs data recovered to get back into an operational state.
3. **NOTE: This is a recovery of last resort. This option should only be considered, if after a full evaluation of the data state is completed, it is determined the production data and its snapshots offer no restore option. It is strongly suggested that you open a case with support, and open a Dell SR for a joint meeting to make sure this procedure should be used.**

4. **Procedures to reverse Replicate Vault PowerScale data back to Production:**
    a. Run the igls CLI command to disable AirGap SyncIQ policies from replicating.
    b. Run the igls CLI command to open the AirGap and open it for hours. (Estimate the time needed to complete reverse replication. It is best to use a large number of hours to ensure the copy can complete without the AirGap closing during the copy process. Example: 1000 hours used on the CLI command)
    c. Issue the SyncIQ CLI command for resync prep on the production cluster AirGap policies (repeat for each policy if you have more than one AirGap policy configured). Consult Dell PowerScale documentation on the command syntax. Verify the command completes successfully using the view jobs CLI command for SyncIQ.
    d. SSH to the vault PowerScale using an IP address of the remote vault PowerScale replication IP pool.
    e. Issue the SyncIQ command to list the SyncIQ policies. You should now see a <AirGap policy name>_mirror policy created by the Resync prep process (if you do not see any mirror policies, open a Dell SR for assistance with SyncIQ)
    f. The <AirGap policy name>_mirror can now be used to replicate the data in the vault back to the production cluster. NOTE: This assumes the production PowerScale is fully operational and can serve the data as required, once the data is re-synced from the Vault PowerScale back to

the production PowerScale. (Repeat this step for each AirGap policy)

i. **NOTE Data Impacting Steps: Make sure you have confirmed all data protected by the AirGap policy needs to be restored. Their is no way back after starting a resync from the vault PowerScale to the production cluster, and resync Prep will make the data on the production cluster read-only blocking all IO to the data.**

ii. Run the <AirGap policy name>_mirror to reverse replicate the data from the vault PowerScale to the production cluster. Use the cluster ISI cli view synciq command to monitor the progress of the copy job.

iii. Once the SyncIQ job(s) complete exit the SSH session with Vault PowerScale cluster. You should now have an SSH session on the production cluster, verify by looking at the command prompt cluster name before you continue.

iv. Data SyncIQ Allow writes step is required to allow the data to be accessible to users, and applications since it is locked by SyncIQ.

    1. Issue the SyncIQ isi or OneFS GUI command to allow writes on the AirGap SyncIQ policy path(s).

        a. Repeat this step on all AirGap policies on the production cluster.

        b. This will mark the data as writable.

c. If SMB shares are in place the data is now accessible and usable by end-users and applications.

d. See Dell Documentation if you are unfamiliar with these commands or steps.

v. **Recovery of all fault data is now complete. Note: To reconfigure the AirGap policies to re-protect production data, open a Support case to get the best method to re-configure the AigGap policies. It might be best to re-sync a full copy into the vault.**

DR Vault Data Access Scenario - Rapid Recovery

1. Unlike backup strategies that use a backup and restore workflow, the Ransomware Defender AirGap 2.0 PowerScale Vault solution stands as the only rapid recovery solution. The PowerScale Vault can vault and lock data in an immutable state and restore data with replication. It is the only AirGap solution that can avoid data recovery completely by using the PowerScale as the file serving device.

   a. **It is recommended to build a RunBook for the steps below based on your production cluster. A well-documented recovery process would allow full data recovery and user/application access in less than 2 hours. NOTE: It is possible to pre-configure most of the steps below to save time during a recovery.**

   b. **Benefits of Rapid Recovery**

i. Eliminates the data copy step by converting the vault into the file severing device. This is the fastest possible recovery option available to customers, that need to get operational in the shortest possible time frame.

ii. If a DR license key is purchased Rapid Recovery allows syncing production shares, exports and quotas.

iii. Integrates file serving and vaulting in a single device.

iv. Supports partial recovery scenario with in-band or out of band recovery options.

v. Supports immutable locked data.

vi. Supports data integrity during copy operations.

vii. Supports versioning of data that provides multiple recovery time periods.

2. **How to perform a Rapid Recovery**

   a. **Vault Cluster Physical Steps**

      i. Connect Ethernet cables from each node to the production network, on the correct ports or VLAN's needed to reach production networks used by the production cluster.

   b. Configuration Data Sync Steps Eyeglass:

      i. SSH to Eyeglass and run the open AirGap command for 1 hour.

      ii. Physically connect the vault PowerScale management interface to the network.

iii. Login to the production protected cluster and disable the airgap SyncIQ policies

iv. In Eyeglass, add the vault PowerScale using the Add Network Element menu option (note this requires a DR license for the vault).

    1. Enter the management IP address, and supply the eyeglass user name and password. Wait for inventory to complete by viewing the Inventory icon.

v. Open Data & Configuration Migration Icon in Eyeglass

    1. Configure a migration policy that matches the source path for each airgap policy and enter the target path used on the airgap policies if different from the source cluster.
    2. Complete the access zone and cluster selection to select the source production cluster and access zone and target vault cluster and access zone. The access zone is where the configuration data is stored that should be copied to the vault.
    3. IMPORTANT: Click the configuration only check box.

4. 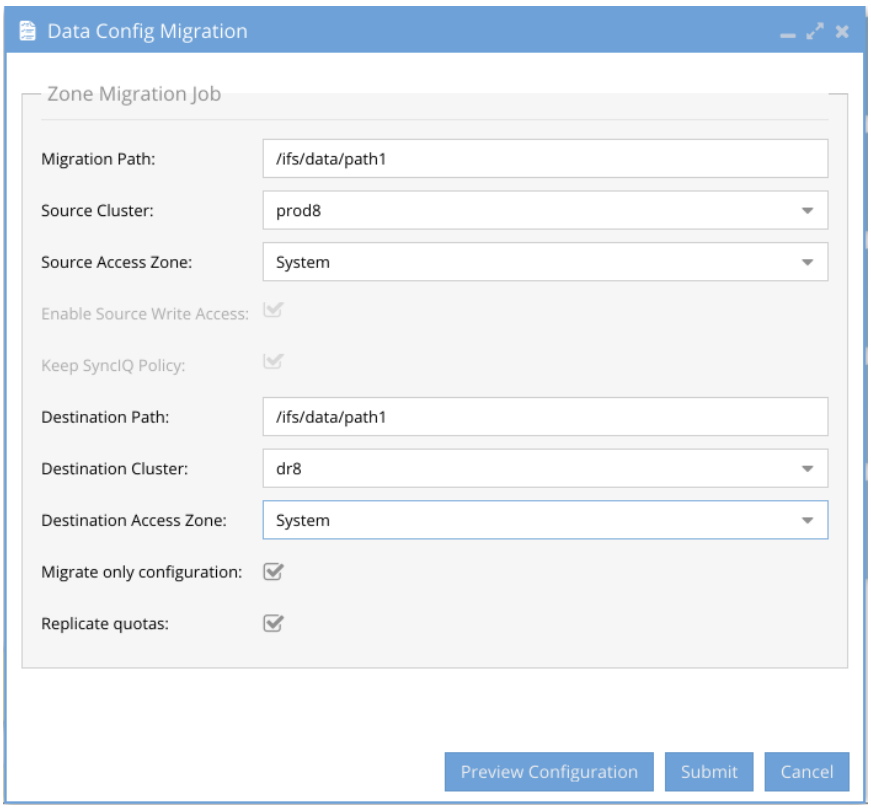

5. Once ready to copy the configuration, Click the Submit button to start the copy function.

c. **Data Write Access Steps**

i. Data on the Vault PowerScale is locked by SyncIQ and will need to be unlocked to provide access to users and applications.

ii. Login to the OneFS management IP address connected in the previous steps. Navigate to Data Protection SyncIQ menu, select the local targets tab, and select the More button and allow writes option on each of the AirGap policy paths. This will allow the locked copy to become the writable copy of the data.

d. **Authentication Providers, Access Zones, IP Pools**

i. Verify all the Access Zones were created on the Access Zone tab in Onefs GUI

ii. Create all IP pools that existed on the production cluster

    1. Assign the SmartConnect names to each pool that existed on the production cluster.
    2. Assign nodes to each pool as required.
    3. Configure new IP ranges or re-use production cluster ip ranges (**NOTE: this will require removing interfaces on the production cluster IP pools, so that IP addresses can be re-used on another cluster without IP address conflicts in the network**). Set other settings on the pool as needed.
    4. Edit each IP pool and assign the correct Access Zone to each IP pool and save the pool.

iii. **DNS** - Update DNS SmartConnect name delegations to point at the vault PowerScale subnet service IP.

iv. **Active Directory** - To speed up the SPN recovery steps, the fastest method is to delete the production cluster AD computer object that will remove all SPN's (Service Principal Names) from AD and the global Catalog. This is required to allow the Vault cluster to register all the SmartConnect name SPN's to its own AD computer object.

v. Active Directory AD Providers - Add the AD provider(s) required for your AD configuration that

was used on the production cluster. AD providers cannot be added in advance since SPN registration conflict would occur with the production cluster. This step must be done at this phase of recovery to simplify SPN bulk registration in AD.

1. Add each AD provider with AD administrator credentials. This will create the computer object and register all SmartConnect names currently configured on the IP pools.
2. Add each Authentication provider to each Access Zone that exists on the vault cluster and set the order with AD provider listed first in the list for each Access Zone.

e. **Test Data Access**

i. At this phase of the recovery, sufficient cluster configuration has been completed to start testing SMB and NFS data access over IP and SmartConnect names.

ii. Attempt IP mount of an SMB share first before testing SmartConnect name access. Then test NFS mount via IP and then SmartConnect names. It is likely debug efforts will be required before going into production.

iii. All teams should be involved AD administrators, DNS administrators, Network Administrators, NAS administrators.

f. **Recovery Complete**

i. This configuration is temporary until a production cluster can be assessed for production use. The Vault PowerScale is expected to operate at lower throughput levels and is designed to provide critical application recovery while planning the full recovery of production systems.

© Superna LLC

2.9. How to Configure a Dell ECS and Data Protection Use Cases

Home Top

- Overview
  - Video Overview
- How to add Licenses
- How to add ECS to Eyeglass Inventory
- Data Protection Use Cases
  - Overview
  - Backup Data on ECS
  - Long Term Legal Hold Data
  - Dell ECS Geo Drive
  - Archive Data

Overview

This topic covers how to add an ECS cluster to Eyeglass inventory and license it for Ransomware Defender.

Video Overview

How to add Licenses

1. Login to eyeglass as admin

2. Open the License manager icon
3. NOTE: Follow steps here to retrieve the license key with the order token.
4. Upload the zip file ECS node licenses and accept the license terms
5. Upload the zip file for the Ransomware Defender for ECS agent license and accept the license terms
6. Open the License manager icon after the refresh
7. Click the Licensed Devices tab
8. Under the Ransomware Defender column, locate the ECS(s) and switch the license type to user licensed and click Submit
9.

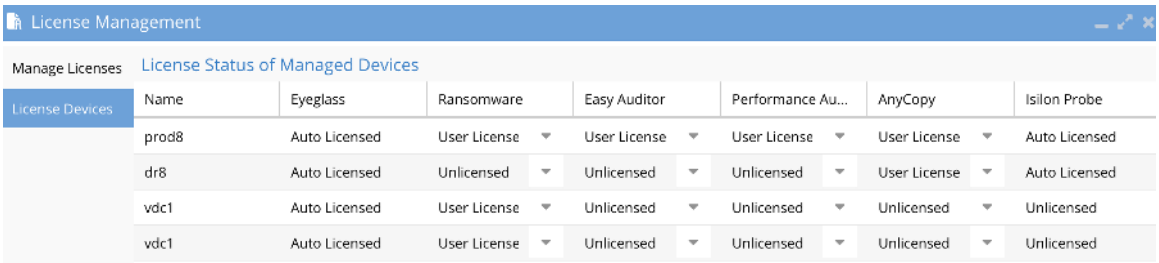

Name	Eyeglass	Ransomware	Easy Auditor	Performance Au...	AnyCopy	Isilon Probe
prod8	Auto Licensed	User License	User License	User License	User License	Auto Licensed
dr8	Auto Licensed	Unlicensed	Unlicensed	Unlicensed	User License	Auto Licensed
vdc1	Auto Licensed	User License	Unlicensed	Unlicensed	Unlicensed	Unlicensed
vdc1	Auto Licensed	User License	Unlicensed	Unlicensed	Unlicensed	Unlicensed

How to add ECS to Eyeglass Inventory

1. Before You begin: Follow the minimum permissions service account creation in the minimum permissions topic here.
2. Login to eyeglass as admin
3. Click the Eyeglass Main menu , bottom left of console
4. Click Add Managed Device

5. Select the ECS tab
6. Add the ip address of the ECS management IP, the eyeglass service account user and password
7. 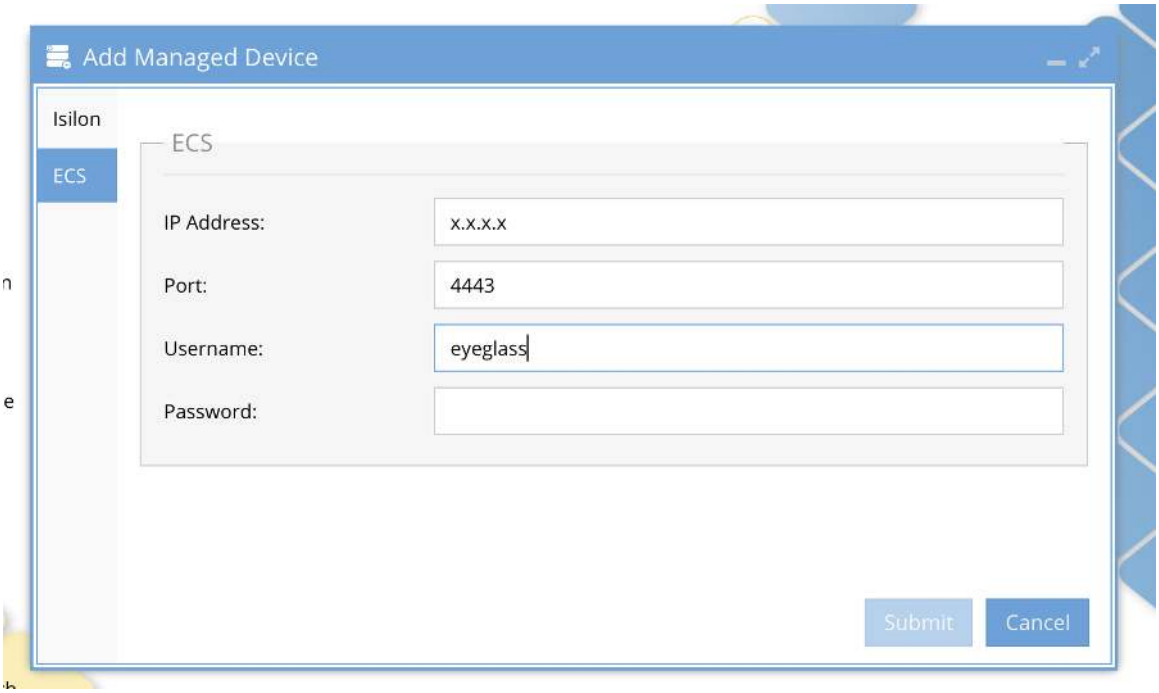

8. Click Submit
9. Answer yes or no to add an additional device.
10. Inventory collection job will appear in the Job icon on the running jobs tab.
11. Open the inventory icon to verify that ECS inventory completed successfully. The inventory can be expanded to show the nodes, name space and storage buckets.

12.

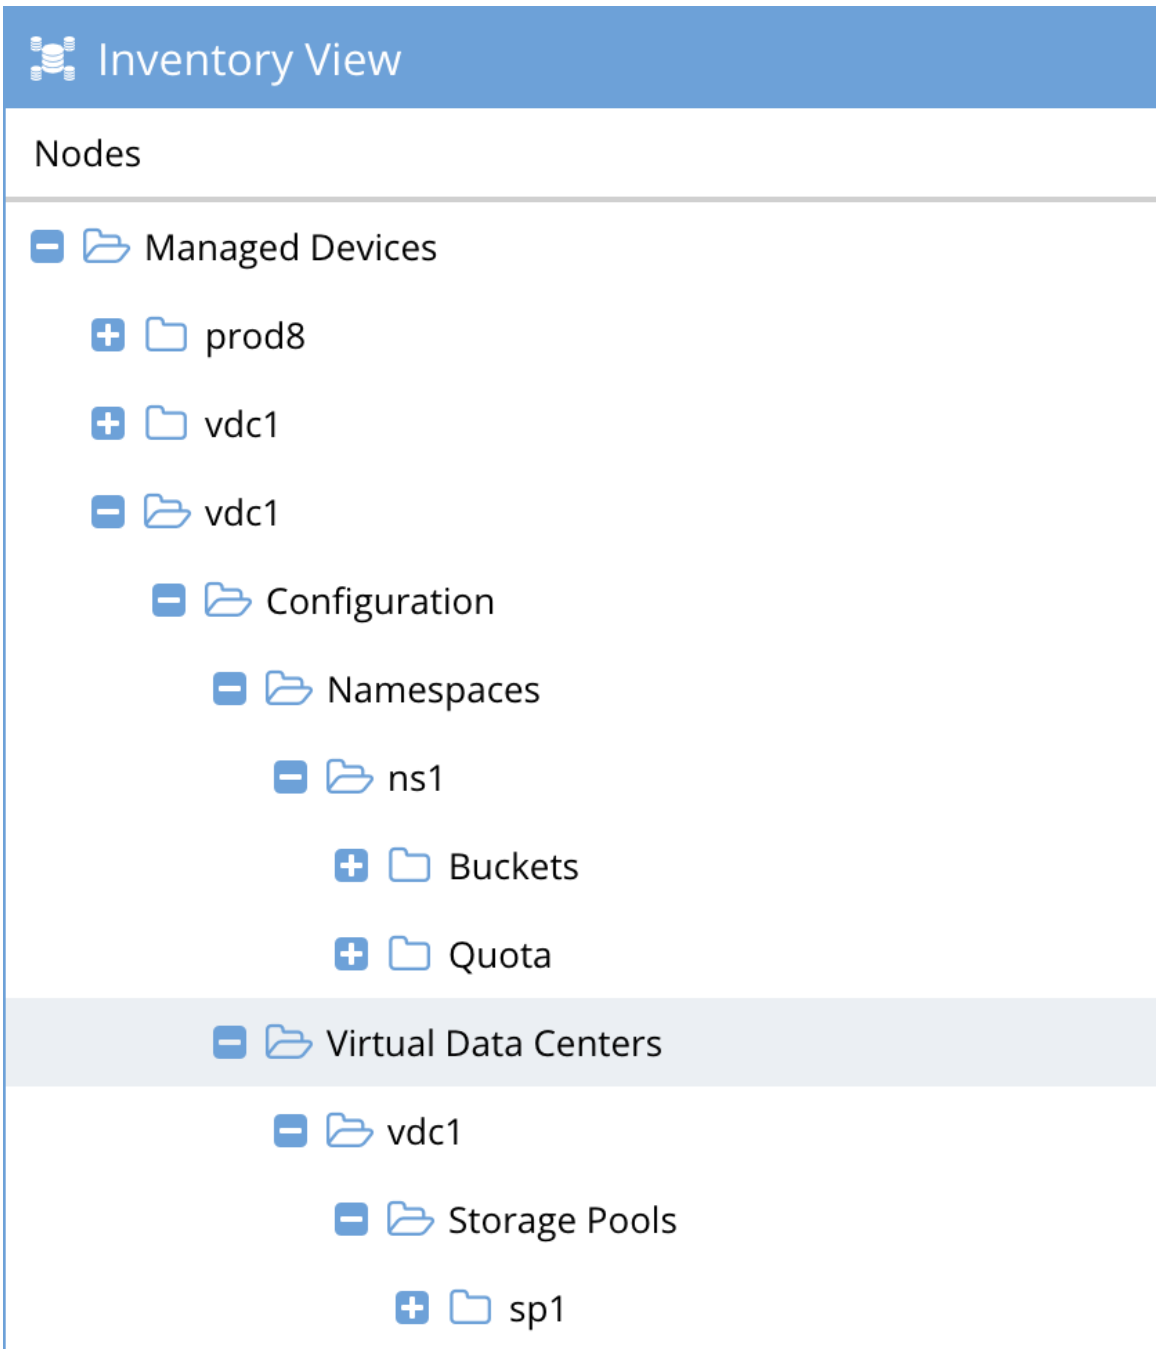

Data Protection Use Cases

Overview

These are common use cases that expose data to Ransomware risks. The first layer of protection is alerting administrators that suspicious activity may be affecting the data described in these use cases. The Ransomware Defender for ECS solution offers both alerting and early warning in

addition to mitigation protection with user account lockout feature that stops any encryption on a per user basis.

Backup Data on ECS

1. Backup applications that store backup data on ECS need the data protected. Backup data and applications are targets for malicious actors that want to take out backup data before an attack.

Long Term Legal Hold Data

1. This is another data type that is commonly stored on ECS for long term storage of legal hold data. This data is required for legal compliance and would also be a target to inflict maximum damage by a malicious actor.

Dell ECS Geo Drive

1. Geo Drive allows syncing local PC data to ECS as a backup solution backed by object storage. This solution enables long term archive for PC data with a locally mounted Windows driver letter.
    a. Ransomware will be able to attack Geodrive data through the Windows OS mounted drive.

Archive Data

1. Long term storage of enterprise archive data is commonly stored on ECS object storage. This data is exposed by server applications that write this data and potentially exposes PB of data protected by a single user and secret key

© Superna LLC

3. Eyeglass Easy Auditor admin guide

Home Top

- What's New
- Introduction to Easy Auditor Guide
- Key Features
- PowerScale Recommended Audit Event Configuration
- Feature Limits
- How to get started with Auditing
- Who Audits the Auditor?
- Planning, Design and How to Use Easy Auditor to Audit
- How to Use Excel for advanced filtering of CSV Reports
- How to Backup and Restore an Audit Database
- Backup and DR for Audit Database with SyncIQ to a Remote Cluster
- How to check Analytics database size
- How to Use Easy Auditor for Typical Audit Use Cases
- Audit Message Workflows
- Advanced Configuration
- Excluded Audit Events
- How to Configure Syslog Forwarding of Formatted Audit messages to an External Syslog Server
- Data Retention of Audit Data and Archive
- How to Purge or Archive PowerScale Audit logs
- Bulk Ingest old Audit Data from PowerScale to Easy Auditor

© Superna LLC

3.1. What's New

Home Top

What's New

For a full list see the feature list page here.

Easy Auditor Enhancements:

1. Quick Scan Path Search - New architecture to accelerate results for path searching when no user is specified. The user search is already indexed in a way to easily find all events by a user. The new search index will offer the same search speed for a path search. (patch release coming soon)

2. AI Analytics of user behavior - Analysis of the Auditor database can determine the optimal Ransomware Defender settings to best protect data and avoid false positives. (patch release coming soon)

3. WireTap provides filtering, folder browsing and event filtering. Complete update with advanced filtering options full screen UI. Realtime IO monitoring of users, paths , folder trees, or single folder. Allows debugging performance issues.

4. Real-time Syslog Forwarding - Allow the ECA cluster to forward formatted syslog message to 3rd parties example SIEM tools, event filtering for user, path, event type with regex filters

5. Where did my folder go? It will now track directory deletes in a fast cache lookup, and copy and paste results to Excel

6. HDFS protocol auditing - Supported now with current release
7. Builtin reports have been enhanced for performance and provide partial results while they execute
8. Optimized active audit triggers offers more performance at higher event rates to real-time DLP and Mass delete triggers
9. Active Auditor - Realtime Audit Triggers - Automate security, "No MORE Report Reading"
    a. If this happens and OR that happens send an alert, triggers do not use the database and process event data with stream based analytics.
    b. Predictive Analytics - Each custom trigger created evaluates event data over 1 minute intervals and every 5 minutes a prediction computation runs to provide more accuracy to your security policies getting triggered.
    c. Combine path, user and event types into a customized real-time audit policy that continuously monitors events and fires a trigger when the condition is met.
    d. Geofencing by user or path - Network Aware Security - Real-time triggers can use the source ip of hosts or even entire subnets. This allows a whole new security layer that can alert when access to storage is from authorized subnets or detect remote access from VPN or Wifi Guest networks
        i. Combined with user, path, file action, file name and more options powerful Geofence polices can be

created to secure your data with network aware policies.

e. This allows and event to be sent via email or configure syslog forwarding to a SIEM.

f. The only customizable real-time audit solution with no lag auditing for PowerScale.

10. S3 Object data access reporting with Easy Auditor for Onefs 9.x releases.

a. Supports reporting on data access based S3 protocol access to the cluster

b. Supported Features

i. Reporting based on query builder

ii. Wiretap

iii. Where did my Folder go? delete of folders, file only

Ransomware Defender Enhancements:

1. No HDFS needed!!!! We have redesigned Ransomware Defender to no longer needed HDFS. Easier to install with fewer dependancies

2. New GUI for flag as false positive to view users that have been flagged and reset the a user to factor default detection settings

3. Allow file list add UI for whitelisting files on the dynamic extension list

4. SIEM Integration - audit data real-time syslog forwarding

Security Enhancements:

1. IGLS cli command to automate changing eyeglass service account password and restarting the process to take affect. Useful for customers with a lot of clusters and regular password change policy can now automate this task.

Cluster Storage Monitor:

1. Large AD direct collection over LDAP will support direct connect to AD to collect users and groups. Supports 1 million or greater object collection in < 2 minutes.
2. This new collection method will be shared by all products that need this information example Ransomware, defender, Cluster Storage Monitor all need user to SID resolution and user to group information.

Eyeglass Cluster Agent:

1. New distributed model allows remote sites to be managed by Mini-Eca a single VM to collect audit data and forward it centrally for process, analysis, storage and searching. This is designed for customers that have distributed clusters and want centralized security and ransomware defense of all clusters.
2. New model can support PowerScaleSD clusters used at edge locations and offer centralized services.

1. **New in 2.5.3**

a. Auto Save Report to an NFS mount igls command to change location of saves and remount to PowerScale for centralized report storage

b. Streamed Result Reporting - Long running queries will now return data for every 50,000 records and the CSV file will allow partial download of results while the query continues to execute.

c. Reports now support 1,000,000 records for CSV download

d. Active Auditor - real-time policy based auditing

   i. Mass Delete detection

   ii. Data Loss Prevention detection

   iii. Actions- alarm, lockout, snapshot

e. What's Happening now? A new way to audit based on stream processing technology that builds an indexed in memory view or active paths with IO visible in a new file system audit viewer that is visible before being written to the database.

   i. Avoids searching for audit events with the last 48 hours. Auto refreshed based on current event stream

   ii. Allows filtering based on time or event type

   iii. 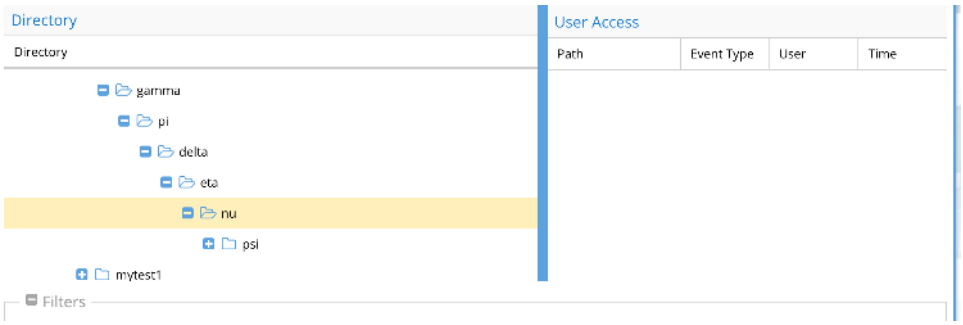

f. New ECA Alarm detection for audit event ingestion issues

g. New ECA Alarm for failure to write to Analytics Database
h. New ingest IGS CLI select a date range of gz PowerScale archived audit events.
    i. Ingest missing data
    ii. Ingest data on disk before Easy Auditor installation
    iii. Avoids and detects duplicate events during ingestion process
i. Load Balance processes on 6 node ECA clusters
j. Historical search logs UI archives all query logs to the PowerScale over HDFS with UI to download or navigate logs
k. Support for 1 Million events in CSV reports
l. Support for continuous results feature that allows retrieval of partial report data while its running 50 000 events at a time. Cancel a report search if the data required is already returned
m. Support for NFS User ID in reports for NFS audit events plus source IP of the NFS client in reports
n. **New Builtin reports**
    i. **HIPAA Report and Security report** for Login , logoff an failed logins by AD user report. Each time a user authenticates to PowerScale or is logged off a netbios session an audit record will be saved for a new Builtin report

ii. **Employee Exit Report** - captures all user activity for 30 days for sending to HR as a record upon an employee leaving your company

o. **HA Ingestion with TuruboAudit**

i. if an ECA node goes down Turboaudit will move ingestion log processing to another ECA node. When the ECA comes online turboaudit will balance the work load between ECA nodes again.

ii. checkpointing log position is written into cluster controller and allows another ECA to process at the same place in the log file that was last processed

iii. Load balancing node audit ingestion between ECA nodes. Dynamic load balancing audit log files between ECA nodes.

p. **RobotAudit -** This feature performs continuous auditing by creating user events as an SMB connected user. The events are created , ingested and stored in the database. The Robot audit process runs reports and counts file and directory events and logs success or failure. This offers the highest level of confidence that audit data is being processed and stored. The audit lag is the time from when an event is created to when the data is searchable.

2. **Existing Release**

a. Quick searching for audit events with filtering and data range

b. Running reports with csv and summary HTML reporting with download and email

c. Scheduled reporting of searches to find specific audit events
d. Wiretap real-time event monitoring by user or path
e. Where did my folder go search interface for directory renames
f. Scalable storage with HDFS and HBASE billions of audit records stored in compressed search able format
g. Native PowerScale storage for Analytics database protected by snapshotIQ and syncIQ lowers the cost of storing and protecting audit data
h. Real-time event processing with automatic triggered responses
i. Integrated with Eyeglass DR and Ransomware Defender into unified single pane of glass
j. Role based management and login with centralized AD or PowerScale user account database

© Superna LLC

3.2. Introduction to Easy Auditor Guide

Home Top

Introduction to this Guide

Overview

This guide covers configuration, setup and monitoring of a Easy Auditor install. The solution is deployed with a 6 VM cluster that process PowerScale audit files with an active active design for maximum availability to survive hardware or software failures. The primary storage of audit data is on PowerScale leveraging HDFS and inline compression to reduce storage costs. Embedded HA and load balanced processing servers simplifies deployment and offers high availability monitoring.

The active solution monitors user behaviours with both real time and report based auditing. The product assists with securing data, identifying performance issues, preventing data loss , audit data for compliance and identify excessive permissions..

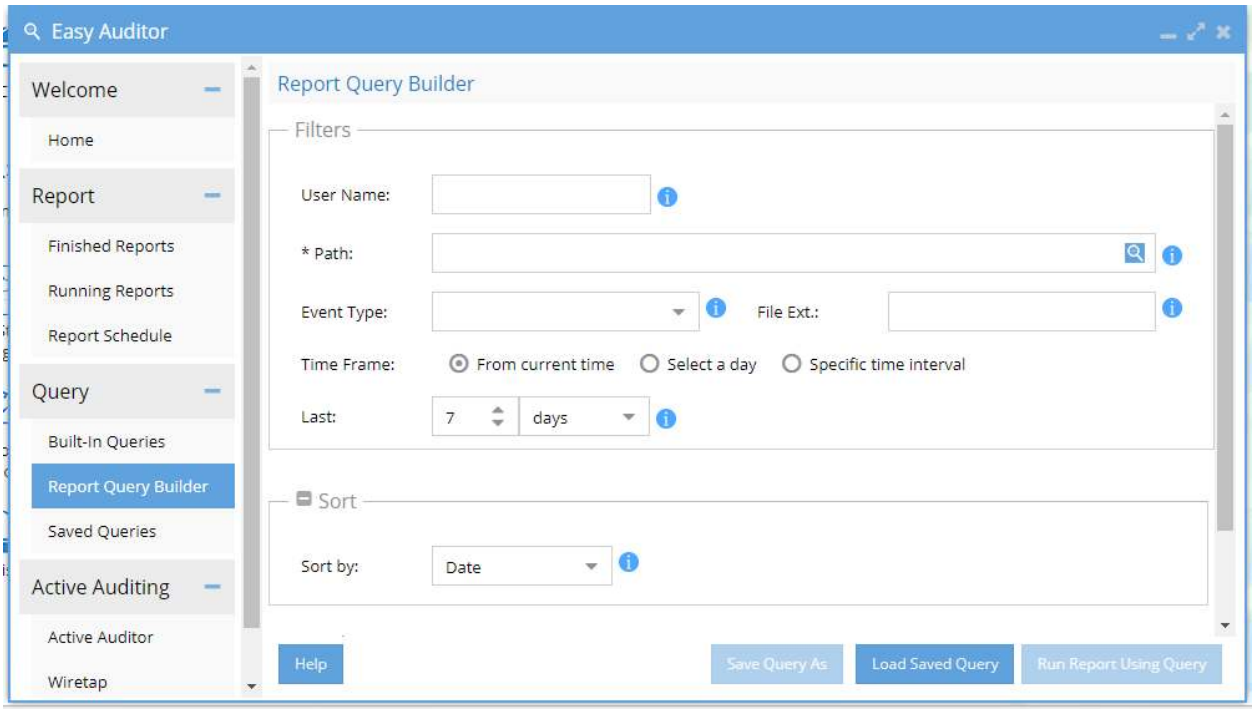

© Superna LLC

3.3. Key Features

Key Features

Active Auditing - Real time audit features

1. **Custom Real time triggers - define your own policies with many fields user, path, file , extension , source ip and more with and or logic to detect any condition in real time and send notifications.**
2. **Mass delete** - many files deleted by a user within a timed period
3. **Data loss prevention** - detect a user reading data from a path, as a percent of the data on the path.
4. **Continuous Search Results -** All searches now return 50 000 records during the search up to a maximum of 1, 000, 000 records. This allows a user to view results as they are found while the search continues to look for more records. This streams results to the report tab and allows the CSV to be downloaded during a search to see if any events found are of interest. If the results found during a search are the only results needed, then the search can be canceled from running reports or allowed to complete.
5. **Configurable actions**:
    a. Alerts (email, syslog, SNMP)
    b. Filesystem snapshot of affected path
    c. User SMB share lockout
6. WireTap - real time file system audit data decoded
    a. Wiretap a path or a user

b. Raw decode of audit events decoded to open files, file actions by user.

c. 2.5.5 new filter option for events, local path or sub folders option. Full screen mode allows viewing more events. Streams based processing queues events for display.

7. Use cases:

a. Performance of file activity by user or application

b. Application IO profiling

c. File locking

d. Group share activity monitoring

e. Real time or historical playback of audit data

8. Where did my folder go? - quickly find renamed folders by users in group shares

a. Search by user, path and date range

b. Identify directory renames by user with old path and new path shown to make reverting data a simple process

c. 2.5.5 adds directory deletes and UI filter to see folder moves or deletes or both.

9. Robot Audit - Automates event creation and report validation on scheduled basis to ensure healthy audit system is maintained on a regular basis. Updated to provide self test feature. Recommended for all customers to enable.

10. Supported Protocol auditing

a. SMB

b. NFS

   c. S3 object (Onefs 9.x)
   d. HDFS (only vendor that can support HDFS auditing)

Report Query Builder

1. GUI Search of audit data by cluster, user , path, date range , file action and file type
2. Save queries for later use or scheduling to run on an interval
3. Schedule queries to run on a interval to email when the query is satisfied.

Report

1. Pre-built reports for stale user access and excessive permissions
2. Top users (create and delete file actions) by file count
3. Top users by quantity of data written
4. Scheduled or on demand reports and queries
5. Reporting
   a. Search by path, user, event type
   b. Analytics reports (top writers by GB, file count create, delete)
   c. Security reports (stale access to data, user share access, login log off reports)
   d. Employee Exit report
6. 2.5.5 New threat detector AI reports use historical data to automatically build recommended settings for Ransomware Defender detection settings. Simple click and apply to activate AI based recommendations.

Role based Login

1. Use the built in Auditor role
2. User is auditor
3. Default password is 3y3gl4ss
4. Or create a new role to separate security, auditing and DR roles with AD group based roles customized to your needs with the user roles icon. See RBAC guide for details.

Scalability

1. Stores audit data on PowerScale
2. Leverage SnapshotIQ, SyncIQ to protected audit data
3. Tier audit data with pools
4. Compresses Audit data approximately 10:1
5. Leverage scale out nas with HDFS on PowerScale and IP pools to expand Disk IO performance
6. Leverage Eyeglass architecture to scale out compute with 6 and 9 node query clusters for scaling to the largest customer sites.

Availability

1. NFS audit data ingestion avoids the cost of CEE servers. Performs with real-time event processing versus stored and forward used by CEE.
2. HDFS + PowerScale and Easy Auditor allows **billions** of rows of audit data to be stored. No aging, pruning is necessary to reduce size of the audit database, providing lossless audit data storage.

Where to get Professional Services

1. To get assistance with auditing configuration and design professional services can be quoted by emailing sales@superna.net
2. Review Audit Service description

© Superna LLC

3.4. PowerScale Recommended Audit Event Configuration

Home Top

Overview

This section covers the recommended audit events that should be configured for Easy Auditor that will provide the best balance of security versus load on the cluster.

OneFS 8.2

1. Audit Success:
    a. close_file_modified, close_file_unmodified, create_directory, create_file, delete_directory, delete_file, get_security_directory, get_security_file, logoff, logon, open_file_noaccess, open_file_read, open_file_write, read_file, rename_directory, rename_file, set_security_directory, set_security_file, write_file

OneFS < 8.2

1. Audit Success:
    a. close | create | delete | get_security | logoff | logon | read | rename | set_security | write

© Superna LLC

3.5. Feature Limits

Feature Limits

Function	Tested Limit	Comments
Concurrent quick searches	10	
Concurrent reports	1	any other report jobs will be queued until a report slot is available to run on the ECA cluster
Enterprise Compliance Mode Clusters	Supported	Requires compadmin user to add cluster to Eyeglass
Concurrent active wiretap monitoring sessions	2	Rate limits are applied on the event rate displayed in the UI
Defined Wiretap configurations (user or path)	2	Rate limits are applied on the event rate displayed in the UI
Active Audit - Mass Delete policies	2	Resources may need to increase to support policies. RAM increase and GHZ
Active Audit - DLP policies	2	Resources may need to increase to support policies. RAM increase and GHZ
Custom Active Audit triggers	25	Tested limit of concurrent triggers
Number of records returned from a query	1,000,000	
Longest Running single Search job	20 hours	Timeout for jobs consuming search resources are capped at 20 hours. This is to avoid a long running search consuming too many resources. An advanced configuration to enable concurrent searches is possible. See guide here on the process to enable concurrent searches.

© Superna LLC

3.6. How to get started with Auditing

Home Top

How to get started with Auditing

Users with auditing permissions are described below. It should should be determined who should have access to the Easy Auditor icon. Best practise is separation of duties with a dedicated auditor administrator.

1. Admin user does **not** have audit permissions by default but the administrator role in Users Roles icon can be used to add the auditor permissions to the admin user.

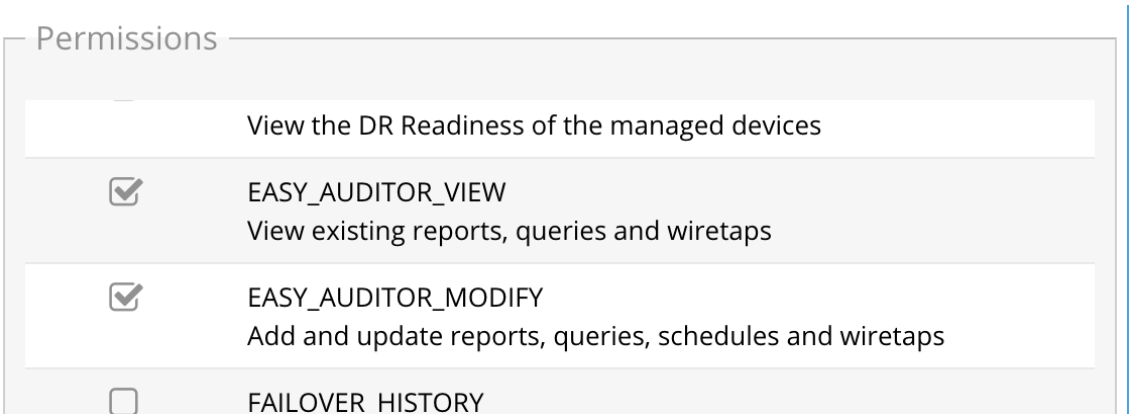

1.

2. Auditor user is a new user id that allows separation of duties from DR and non security audit functions in eyeglass.

1. **Separation of duties** is required for compliance to most industries regulations example HIPPA, PIC
2. Login as the **auditor** user with default password **3y3gl4ss**
3. Change the auditor password (steps to change the password)

4. 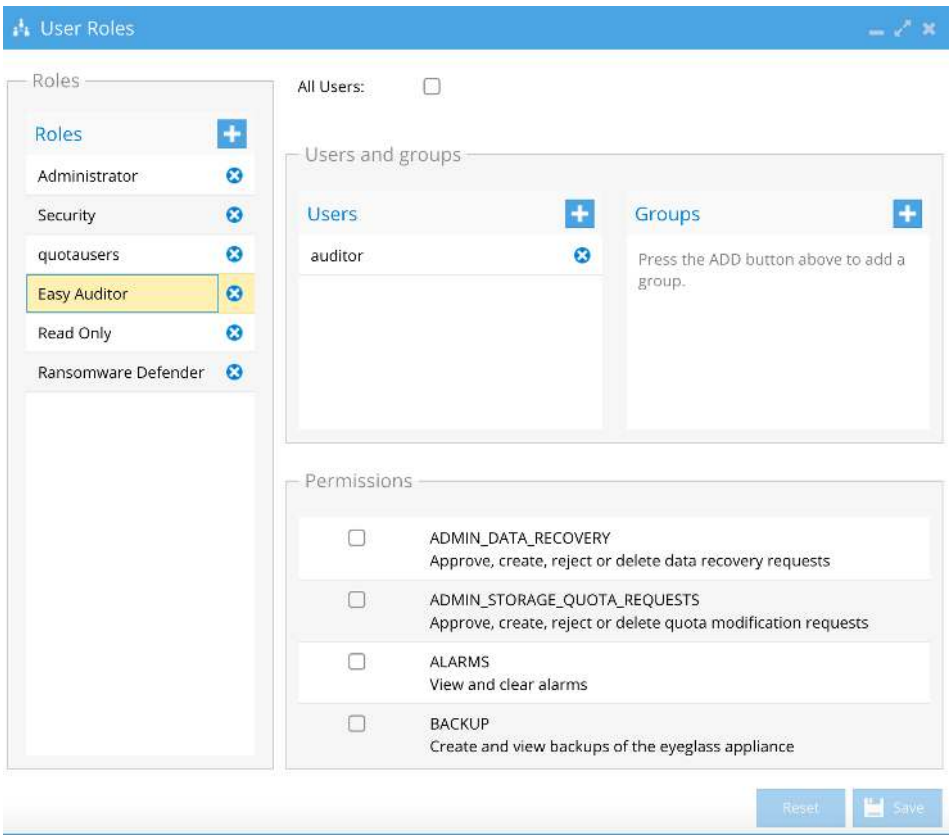

3. Configure email reporting for Audit administrators

1. Login to Eyeglass as admin user
2. Open notification center
3. Select recipients tab
4. Add new email receipt and select auditor only reports from the list.

5. 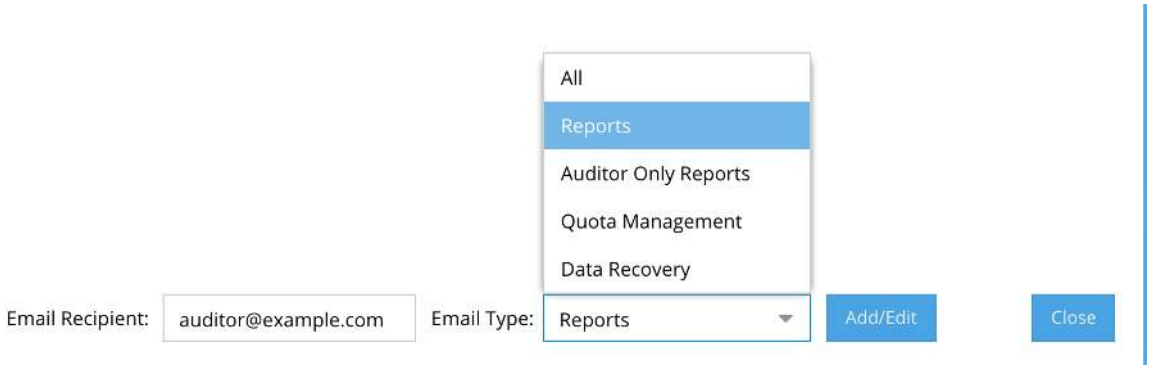

© Superna LLC

3.7. Who Audits the Auditor?

Who Audits the Auditor?

Requirements

1. Release > 2.5.7

Eyeglass will log all login activity and major actions taken within the Easy Auditor UI and all UI's . The Eyeglass audit log is stored in the file system and is included in eyeglass appliance backup zip files.

1. View the appliance hardening guide for steps to review user actions.

© Superna LLC

3.8. Planning, Design and How to Use Easy Auditor to Audit

Home Top

- Overview
- Read Me First
- Planning
- Daily Operating Procedures
- Deployment Topology Recommendation:
- Query Performance Configuration
- Expanding ECA cluster Size for performance
- Reduce the Load on PowerScale Cluster by Disabling Audit Events
- Audit Use cases
- How to Configure What PowerScale will audit
- How To Enable User logon and logoff support for the built in Logon monitor Compliance Report
- How to Configure and Operate Easy Auditor
  - Report Query Builder
    - How to search for Audit Events
    - How to Search for audit data for a specific file
  - Built In Reports
    - Data Access Report
    - Top User Create files Report

- Top User Deleted files Report
- Stale User Access Report
- Access User Report
- Login Monitor Report (compliance)
- Employee Exit Report
- Count
- AI User Behavior Analytics Reports - Detects and Applies Ransomware Defender settings (Beta Feature)
- Best Practise
- How to run the AI reports (beta)
- How to view the AI report User settings applied to Ransomware Defender Configuration

- Saved Queries Tab
  - How to delete a saved search
  - How to Load a Query and Search
  - How to Load and Schedule a Query
  - How to delete a saved Query
  - How to delete run a Query as a report

- Running Reports Tab
  - Running Reports Monitoring
    - How to Cancel a running report
    - Best Practice:
    - How to Monitor Progress of a Running Search

- Finished Reports
  - Filter Reports tab
- Report Schedule Tab
  - How to delete a saved Scheduled Query
  - How to load a Query and set a Schedule
  - How to edit a saved Scheduled Query
- Active Auditor Tab
  - Active Auditing Tab Overview
    - Overview
  - How to Configure Custom Real time Audit Policies
    - Overview:
    - How to Manage customer Real time Audit policies
    - How to Configure real time audit policies
    - How the UI elements work for real time audit policy rules
  - UI Examples
    - Multiple individual rules using OR between rules
    - Multiple individual rules using AND between rules
    - A group of rules using AND between rules
    - A group of rules using OR between rules
    - A group plus comparison to an individual rule using AND between the group rules and an AND between the group and single rule
  - Auditor Active Trigger Use Case Examples

    - HoneyPot Detection
  - How to configure Mass Delete protection
    - Overview:
  - How to configure Data Loss Prevention
    - Overview:
  - How to manage Active Auditor Events
  - Actions available to Manage Active Auditor Events
  - How to Archive an active auditor event
  - How to Display active auditor historical events
- WireTap
  - How to Configure Wiretap
- Where Did My Folder go? Tab
  - Overview
  - Limitations
  - Use Case
  - How to Use Where did my folder go?
    - Overview
- RoboAudit
  - How Use RobotAudit
  - How to Monitor RobotAudit jobs
    - Open the Jobs Icon
    - Click running jobs

    - Select the RobotAudit job and expand to see the steps that are being executed.
  - How to verify RobotAudit test Results
- Auditor Role Based Access
- How to view debug logs for Searches

Overview

The Easy Auditor solution for PowerScale requires existing Eyeglass DR cluster licenses for each PowerScale cluster plus an Eyeglass clustered agent license.

The Eyeglass Easy Auditor solution allows customers to set policies for both real-time policies and scheduled searches to alert on file actions. Static reports can also be scheduled for more complex longer running analytics.

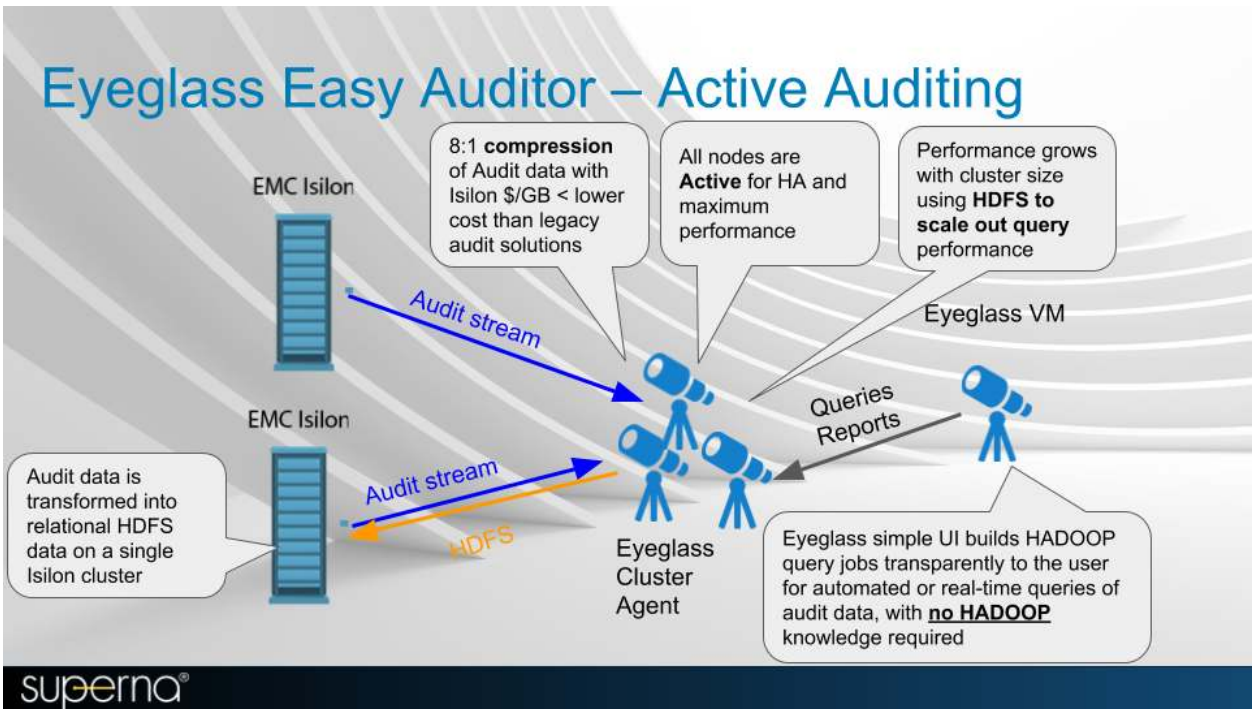

Read Me First

1. Complexity of auditing a multi protocol file system requires understanding of NFS and SMB.
    a. NFS - is a case sensitive protcol that means remove mounts and operations will generate case sensitive file and directory audit events. These audit events are sent to Easy Auditor to process with the case preserved. This means locating data based on the path and file name case will be honored as expected.
    b. SMB - Is not a case senstive protocol. NTFS view of file systems will honor case in the display

of Explorer but operations over SMB can open, delete files using the wrong case and the action will still suceed. This also means the case of the file operation sent to Isilon will be sent to the audit records and can result in a mismatch between the actual file system case and the audit event recorded. This is how Isilon works when creating audit events. Reach this technical advisory of how mouting shares with a subdirectory mount can generate audit data that does not match the actual file system. Tech note.

i. If a user renames a path below an SMB share mount audit data saved in the database will still have the old path case in the database. This means searching for audit data within Easy Auditor will attempt to search the current file system path as entered into the GUI but will not find any records since the database stores the previous case path.

ii. Work Around: Always begin a search at the base path of the SMB share. This will ensure that all records will be located, even if users

change the path on folders and files under the SMB share path.

Planning

Several decisions are required to configure auditing for 1 or more clusters with a single ECA cluster.

1. Number of clusters to assign to a single ECA cluster
2. Number of audit events per second (number of active smb or nfs connections is used to size event rate)
3. WAN link speed to send audit events over the wan to a centralized ECA cluster. See mini ECA architecture option for distributed auditing.
4. Query performance

Daily Operating Procedures

1. It is expected that alarms are monitored daily and acted upon. Some alarms indicate network connection issues to managed clusters. The impact of not monitoring alarms is missed audit data ingestion. This is daily task.
    a. If audit data is present in the database as a result of network issues in your environment, re-ingestion will be required. See guide here. NOTE: This is a slow background process to ingest old audit data stored on the cluster.

2. Robo Audit feature should always be configured and monitored to verify normal audit data ingestion, and storage in the database. This feature also tests searching on a daily basis.

Deployment Topology Recommendation:

1. Centralize the ECA cluster when possible and use the WAN link to send audit events. Audit data is xml over NFS and tolerates latency well. See the installation guide for bandwidth guidelines.

Query Performance Configuration

Two key factors of the Auditor Analytics database is write performance to PowerScale over HDFS for storing audit event streams and read performance for queries. Additional advanced configuration can be completed after contacting support.

To expand a cluster from 6, 9 or 12 VM's follow the steps below.

Event Rate Total per ECA Cluster	Number of ECA VM's	Comments
10000 or greater events per second (run this command isi statistics query current --nodes=all --stats=node.disk.xfers.rate.sum) and send results support	6 to 9 to 12 ECA VM cluster	Nodes 2-12 only run containers for read and writing data and analysis. These VM's can have memory lowered, contact support for assistance.

Expanding ECA cluster Size for performance

See Install guide expanding cluster size to increase to 6 or 9 or 12 nodes. Increasing cluster size will increase event rate processing and analysis job report speed.

Reduce the Load on PowerScale Cluster by Disabling Audit Events

PowerScale can disable some event types to reduce audit work load.

1. New in 8.2 release each audit event can be enabled or disabled. Use this command **isi audit settings modify --help. Open a case with support to get a recommendation of which events can be disabled.**
    a. --audit-success (close | close_directory | close_file | close_file_modified | close_file_unmodified | create | create_directory | create_file | delete |

    delete_directory | delete_file | get_security | get_security_directory | get_security_file | logoff | logon | open | open_directory | open_file |

    open_file_noaccess | open_file_read | open_file_write | read | read_file | rename | rename_directory | rename_file | set_security | set_security_directory |

    set_security_file | tree_connect | write | write_file | all) | --clear-audit-success | --add-audit-success (close | close_directory | close_file |

    close_file_modified | close_file_unmodified | create | create_directory | create_file | delete | delete_directory | delete_file | get_security |

    get_security_directory | get_security_file | logoff | logon | open | open_directory | open_file | open_file_noaccess | open_file_read | open_file_write | read

| read_file | rename | rename_directory | rename_file | set_security | set_security_directory | set_security_file | tree_connect | write | write_file | all) |

--remove-audit-success (close | close_directory | close_file | close_file_modified | close_file_unmodified | create | create_directory | create_file | delete |

delete_directory | delete_file | get_security | get_security_directory | get_security_file | logoff | logon | open | open_directory | open_file |

open_file_noaccess | open_file_read | open_file_write | read | read_file | rename | rename_directory | rename_file | set_security | set_security_directory |

set_security_file | tree_connect | write | write_file | all)]

2. In previous OneFS releases there is less control over events that are enabled or disabled
    a. isi audit settings modify --help
    b. --add-audit-success (close | create | delete | get_security | logoff | logon | read | rename | set_security | tree_connect | write | all) |
    c. Open a support case to get recommendations on events to disable.

Audit Use cases

The following use cases can be addressed by Easy Auditor

1. Find file deletes in the file system using searches
2. Configure a scheduled query to find deletes or other file actions and get alerts real time
3. Quickly find renamed directories using “Where did my folder go?” And revert the files to the previous location
4. Monitor secure shares for users copying data from the share
5. Report on user activity for compliance with HIPPA , PCI
6. Identify the top users for creates and deletes
7. Performance monitoring paths in the file system to profile application workflows or users
8. Track user activity for security audits
9. Find insider threats with advanced search
10. Store long term audit data for compliance reporting
11. Identify excessive permissions to data to assist with remove access to users that do not require access

How to Configure What PowerScale will audit

These ISI commands should be used to validate and change what events PowerScale will audit per access zone.

List audit settings on an access zone

isi audit settings view --zone=data (data is the zone name)

Default should look like this example

Audit Failure: create, delete, rename, set_security, close

Audit Success: create, delete, rename, set_security, close

How To Enable User logon and logoff support for the built in Logon monitor Compliance Report

Run these commands per access zone (example zone name below is data) to enable logon and logoff auditing

isi audit settings modify --zone=data --audit-success=**logon,logoff**,create,delete,rename,set_security,close

isi audit settings modify --zone=data --audit-failure=**logon,logoff,**create,delete,rename,set_security,close

How to get list a of audit configuration options

isi audit settings modify --zone=data --help (zone name data is example only)

Verify the settings with list command

Audit Failure: close, create, delete, logoff, logon, rename, set_security

Audit Success: close, create, delete, logoff, logon, rename, set_security

How to Configure and Operate Easy Auditor

Use this section on how the configure and use Easy Auditor to search audit data.

Report Query Builder

Use this tab to search by user(s), path(s) , file extension ,file action and date range.

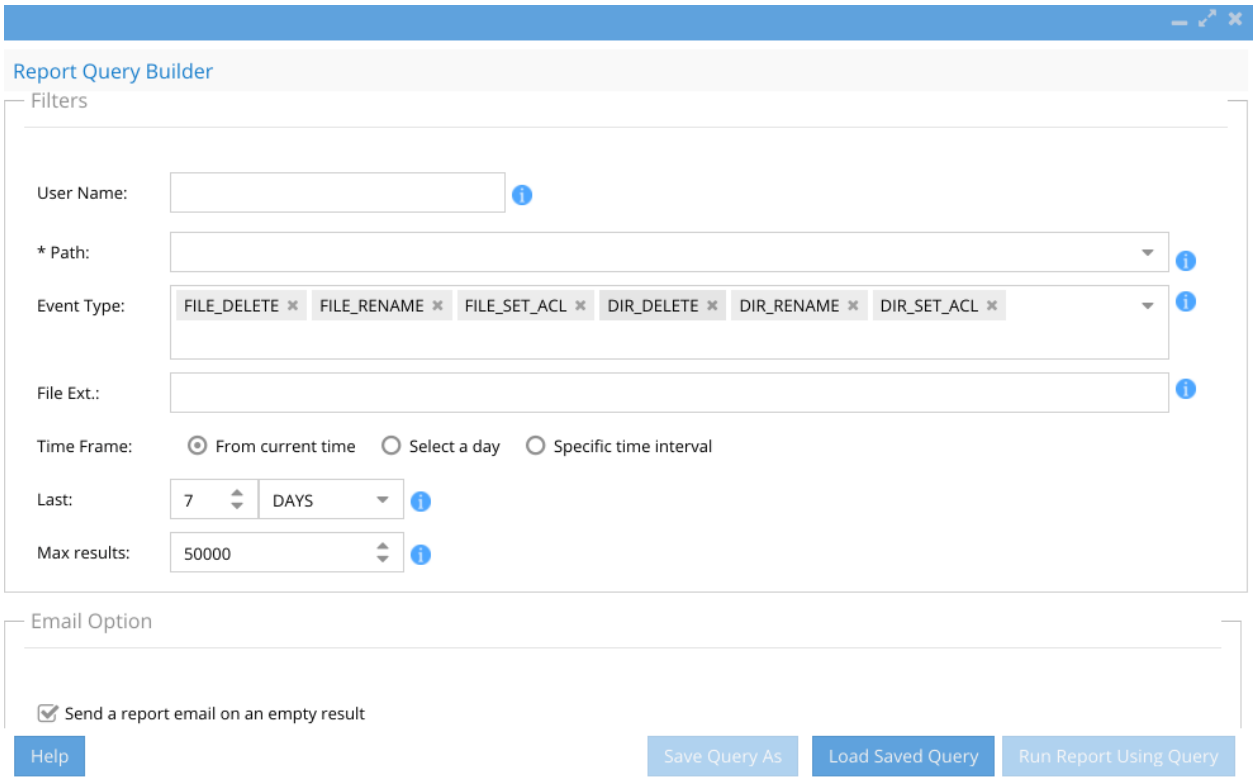

How Queries are Executed

How to search for Audit Events

1. The search event types defaults to the most common searches, click the drop down to see all available event types. Reducing the scope of the event types will speed up search times.
2. Search by entering data into the GUI fields
    a. Blank for the user field means any user. to Search for specific user enter user@domain or domain\user (domain must be upper case). **NOTE: This will resolve an active directory user to a SID and will support Windows SMB protocol user auditing.**

b. **The path field must be entered to select a cluster and path to audit. To speed up search times select a path as close to the audit events you are interested in finding. example don't enter /ifs this will increase search times.**

c. Extension field will match files with a specific extension to limit to file types of interest and this can reduce search times by ignore file types you do not need results. Enter only the extension example docx or xlsx or pdf **(NOTE: do not enter the period)**

d. **Time Range:**

   i. **Enter days in the past**

   ii. **a spefiic day**

   iii. **or custom date and time range**

   iv. NOTE: searching over a shorter time window will decrease search times. Always narrow the time window on a search and use multiple searches if the first search does not provide the audit data you are looking for.

   v. **Max results default is 50 000 records, the search will finish once the max results has been crossed. Best Practice is to narrow the search criteria to avoid returning too much data in a search. The maximum results supported is 1 Million events, this will generate a very large CSV and take more time to complete the search.**

e. How to Search for audit data for a specific file

i. If you know the name of the file, then follow GIF instructions to enter the file name to the search path NOTE: file name search is case sensitive.

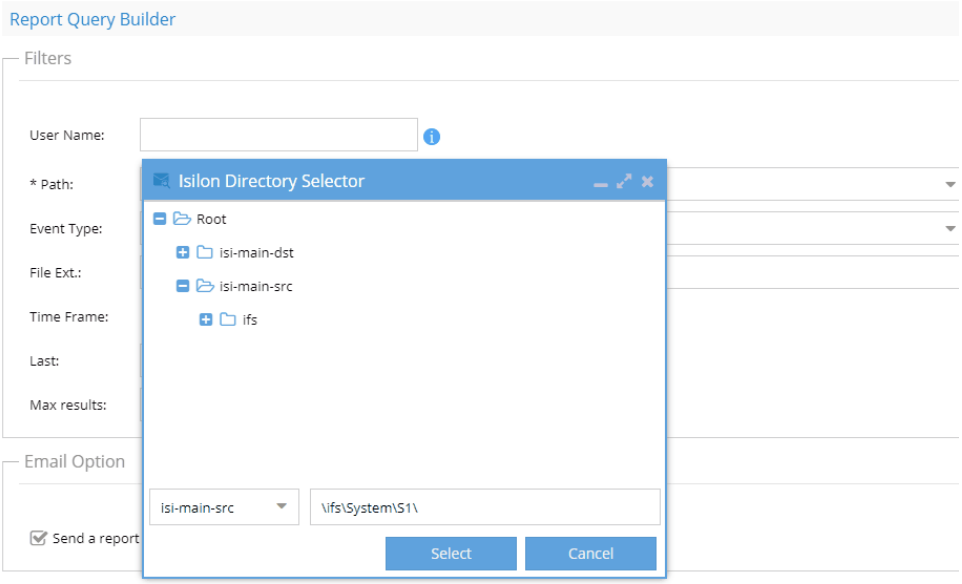

ii.

3. **NOTE: Email option if enabled will send an email even if the query returns no data. Disable this check box to skip sending an email if no data is returned. This option is best used for scheduled queries.**
4. Use the save option to save the query to saved query tab to reload at a later date.
5. Use the Run Report as Query
6. Execute a search with the search button
    a. Sort
    b. Use Sort to return the results in the preview screen according to these settings
    c. **NOTE: Search supports up to 1,000,000 records. Saved queries can be run as a report to return more data as csv download files.**

d. **NOTE: All Searches use Continouse mode that will update the CSV on report tab each time 50 000 records are found. This allows results to be viewed during the execution of search for faster preview or results review. The CSV can be repeatedly downloaded during a long search as required. If the results are found, the search can be canceled from the running reports tab.**

Built In Reports

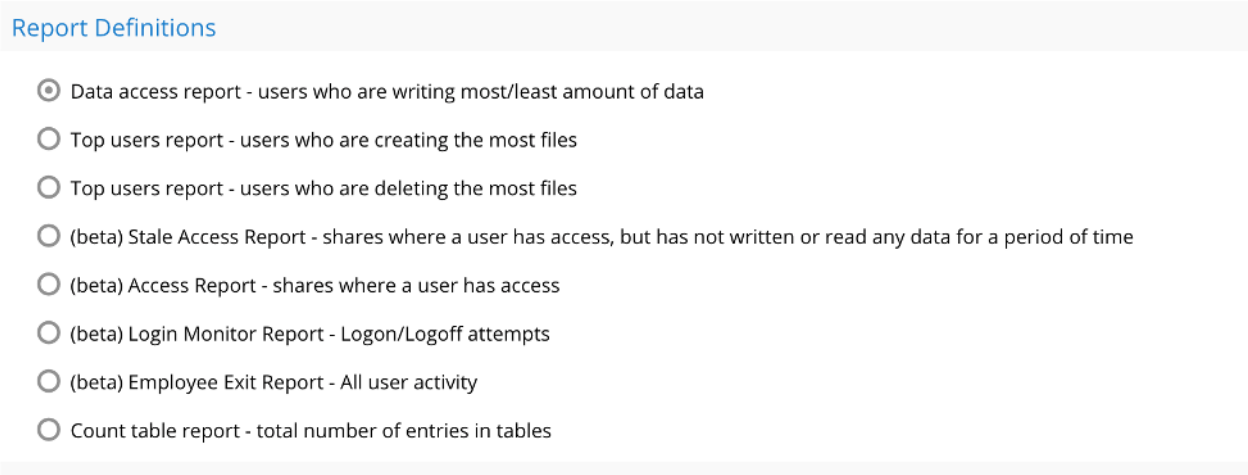

Data Access Report

This report will be used to determine which users write the most data in GB to the cluster. Use the parameters to determine the report output.

Recommendation: Use this report to determine highest load users and help determine node count and smartconnect design. This report can also assist with tiering policies for paths and users that are consuming the most or the least of the clusters resources.

NOTE: This report is not intended to be used to calculate cluster usage. It samples data over a time period to provide a relative usage sample of users over that time period who wrote the most data to the cluster. The report samples the database and may not read all data in the time period selected.

Warning: As of release 2.5.5 the event used for this report is disabled, it is a very high volume event type when data is written which creates very large audit database. This event type has no value for auditing data and is now disabled and will not be ingested or stored in the audit database. In order to use this report. A cluster filter must be changed to allow processing and storing of this event type. Contact support to have the filter removed to allow this report to be available. In many environments the rate of this event is low enough to leave the event filter disabled.

Report Parameters

1. Highest or lowest x %
2. Cluster
3. Time period to analyze

Top User Create files Report

This report will be used to determine the users that create the most files regardless of the size of the files. Use the parameters to determine the report output.

Recommendation: Use this report to determine highest file creation users and help determine node count and smartconnect design. This

report can also assist with tiering policies for paths and users that are consuming the most or the least of cluster resources.

Report Parameters

1. Highest or lowest x %
2. Cluster
3. Time period to analyze

Top User Deleted files Report

This report will be used to determine the users that delete the most files regardless of the size of the files. Use the parameters to determine the report output.

Recommendation: Use this report to determine highest file deletion by users. This list of users can be used to track who is deleting content on the file system. This can be used to help determine node count and smartconnect design. This report can also assist with tiering policies for paths and users that are deleting content that could be tired.

Report Parameters

1. Highest or lowest x %
2. Cluster
3. Time period to analyze

Stale User Access Report

This report will be used to build a list of users that have accessed data using SMB shares and calculate the last read or write of each share

they have access to mount based on AD group membership. This report will list all users that can mount shares and whether they have accessed data during the reporting time period.

Recommendation: Use this report to determine which users may not require access to SMB shares based on access patterns. This is a security report that can be shared with departments that manage SMB share access.

NOTE: This is a long running report that can take hours to complete on a large database. Large user count Active Directory domains can also cause the report to run longer to analyse data access.

Report Parameters

1. Cluster
2. Time period to analyze

Access User Report

This report will be used to build a list of users and SMB shares and map out user to share access to determine excessive permissions or validate existing share access that may not be inline with the desired security policies. The report generates a list of shares and a list of active directory users and groups that have access SMB share. The AD groups can be expanded to a list of users for NAS administrators that do not have access to Active Directory.

Recommendation: Use this report to determine which users may not require access to SMB shares based AD group membership.. This is a

security report that can be shared with departments that manage SMB share access.

NOTE: Large user count Active Directory domains can also cause the report to run longer to analyse data access.

Report Parameters

1. Cluster

Login Monitor Report (compliance)

This report satisifies HIPAA and PCI Requirements to track login to systems storing compliance data. NOTE: This builtin report has moved to the Query builder interface as of release 2.5.6 patch release. This simplifies searching and speeds up the reporting process.

1. **Prerequisites**
    a. **NOTE: logon and logoff events are disable on PowerScale by default follow these steps to enable. Replace the zone name for the zone you want to audit**
        i. isi audit settings modify --add-audit-success logon --zone System
        ii. isi audit settings modify --add-audit-success logoff --zone System
        iii. isi audit settings modify --add-audit-failure logoff --zone System
        iv. isi audit settings modify --add-audit-failure logon --zone System
2. How to run a Logon Report

a. Open the Query build tab and enter a search based on the screenshot below. This will retrieve all the logon and logoff events for a day. The path field can be left at /ifs as it is not used in the search. Select a day or a time range as required for your search. **NOTE: Searching over more days will take longer to complete the search.**

i. 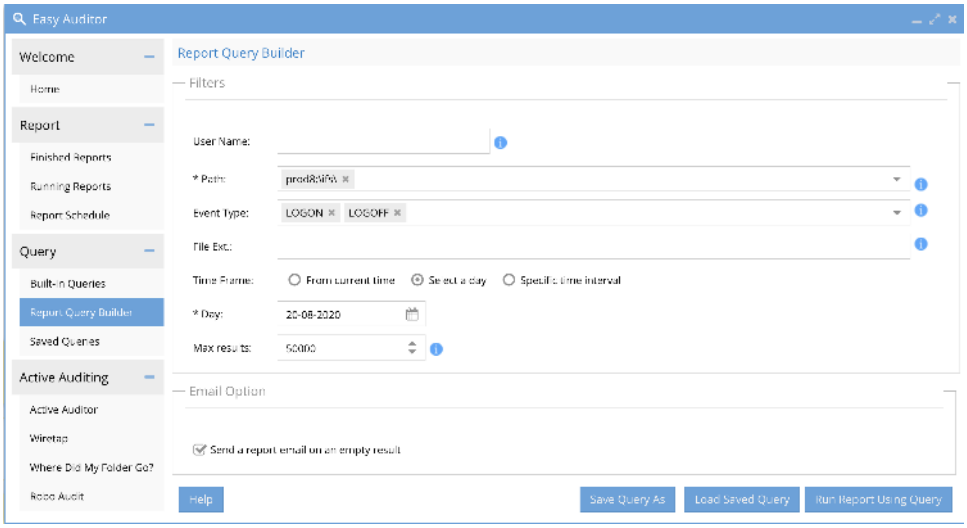

ii. Sample report in the Finished reports tab

iii. 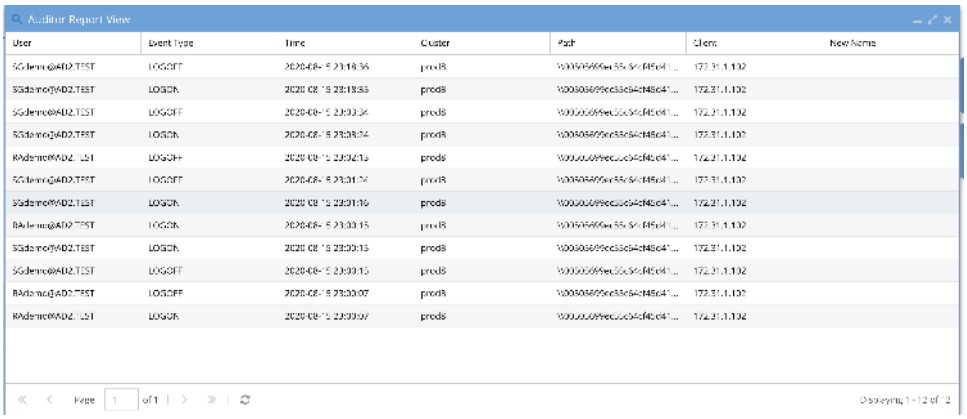

Employee Exit Report

This report would be run prior to an employee leaving your company and provides a view of all files the user accessed in any way over the

last 30 days. This is typically done as part of an HR process. The report will show user activity by day over the last 30 days.

1. Select the employee exit report
2. 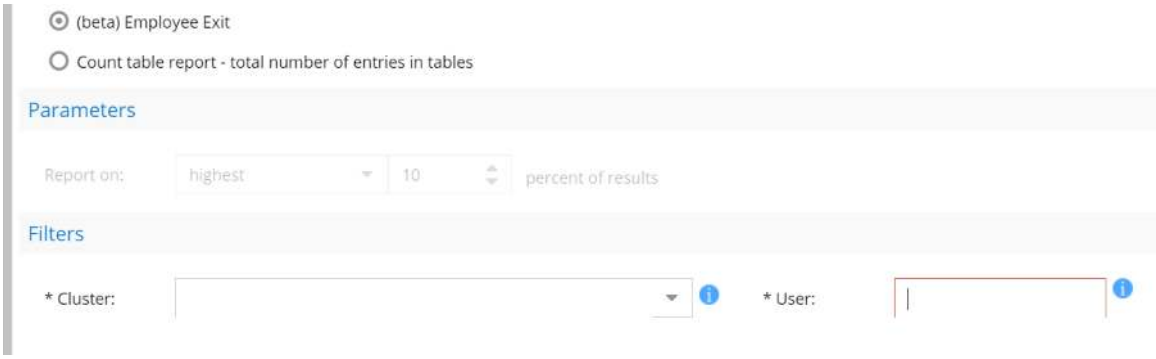

3. Enter the userid user@domainname or domain\userid
    a. if the user SID cannot be resolved an error will be returned.
4. Select the cluster
5. Save and name the report
6. Execute the report
7. View results on the report tab
8. Download CSV to send to HR
9. Example email report
10. superna®

Easy Auditor Report for Query"bz1" 2018-05-26 11:42:49 EDT

Cluster: Isi805-bz-1

Auditor Employee Exit Report

Spark Application Id : app-20180526153945-0006

Isi805-bz-1	admin	S-1-22-1-10	Tue May 22 22:00:42 EDT 2018
Daily Usage: Wednesday, May 23, 2018	Total Events: 382	First event at: Tue May 22 22:00:42 EDT 2018	Last event at: Wed May 23 19:34:40 EDT 2018
Event Type	Client IP	Time	Zone:{Directory,File}
FILE_CLOSE	10.100.252.46	19:34:40	System:/ifs/data/shares/s2/c.Petya
FILE_OPEN_NOACCESS	10.100.252.46	19:34:40	System:/ifs/data/shares/s2/c.Petya
FILE_CLOSE	10.100.252.46	19:34:40	System:/ifs/data/shares/s2/c.Petya

Count

This report will be recommended to run by support to count the rows in each system table in the analytics database. On a large database this job can take a very long time to complete. We recommend using DB size procedure in this guide.

NOTE: This report can run for hours on a large database and generate IO to the cluster. Run in off peak hours example overnight or weekends.

Report Parameters

1. None

AI User Behavior Analytics Reports - Detects and Applies Ransomware Defender settings (Beta Feature)

These reports analyzes user behavior in the auditor database and builds recommended settings customized to IO pattern found in the audit database. This simplifies configuration of Ransomware Defender settings based on historical data.

The output of the report shows the settings that will be applied to the Ransomware Defender settings. These settings modify the threat detectors on the ECA cluster with an override setting. Each Threat detector has a report that will recommend the settings based on historical data in the Audit database.

Best Practise

The threat detector AI report defaults are design to apply the best settings based on actual user activity. This feature will apply detector settings based on 90% of the user IO sampled with defaults and then suggest the remaining 10% of the top users sampled with per user overrides that apply settings for these top 10% of the users. This balanced approach will reduce detections for most user IO and then apply specific settings for the to 10%.

The report defaults are shown below

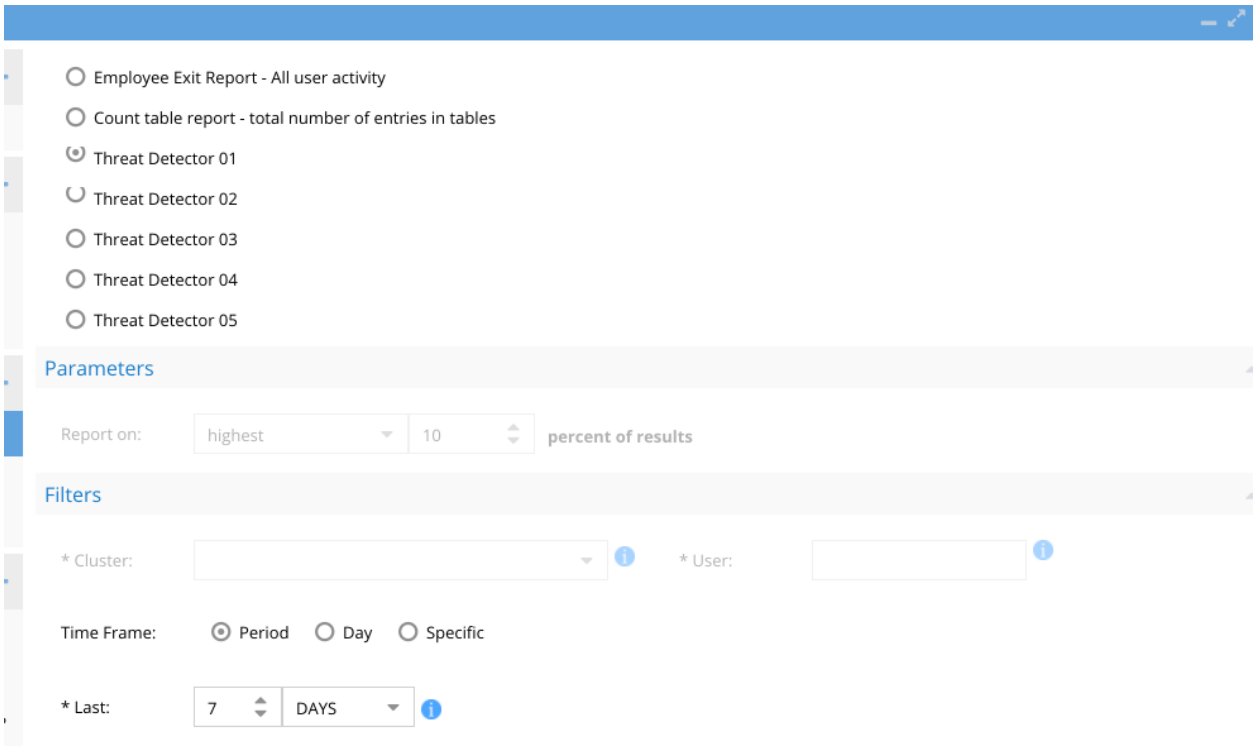

1. Leave top 10% of users at the default setting
2. select the cluster to analyze
3. leave the user input blank to analyze all users versus a single user
4. Leave the default of 7 days
5. Save and run the job definition

How to run the AI reports (beta)

1. From the Builtin reports window Select Threat Dector AI report 1 or 2 and enter the number of days to analyze. The default of 7 days is recommended.
2. Save the report name to submit the report
3. Monitor the report job from the running jobs window
4. When the report finishes click the view button to review and apply the settings.

See examples of Threat Detector 1 and 2 reports. The Save and Apply button will apply the settings to the ECA cluster settings. The threat detector section identifies the global settings that should cover 90% of the user behaviors. The bottom section lists the user sids that represent the top 10% of the user behaviours that would trip the detectors and sets values that will apply to each of the users listed to avoid detections.

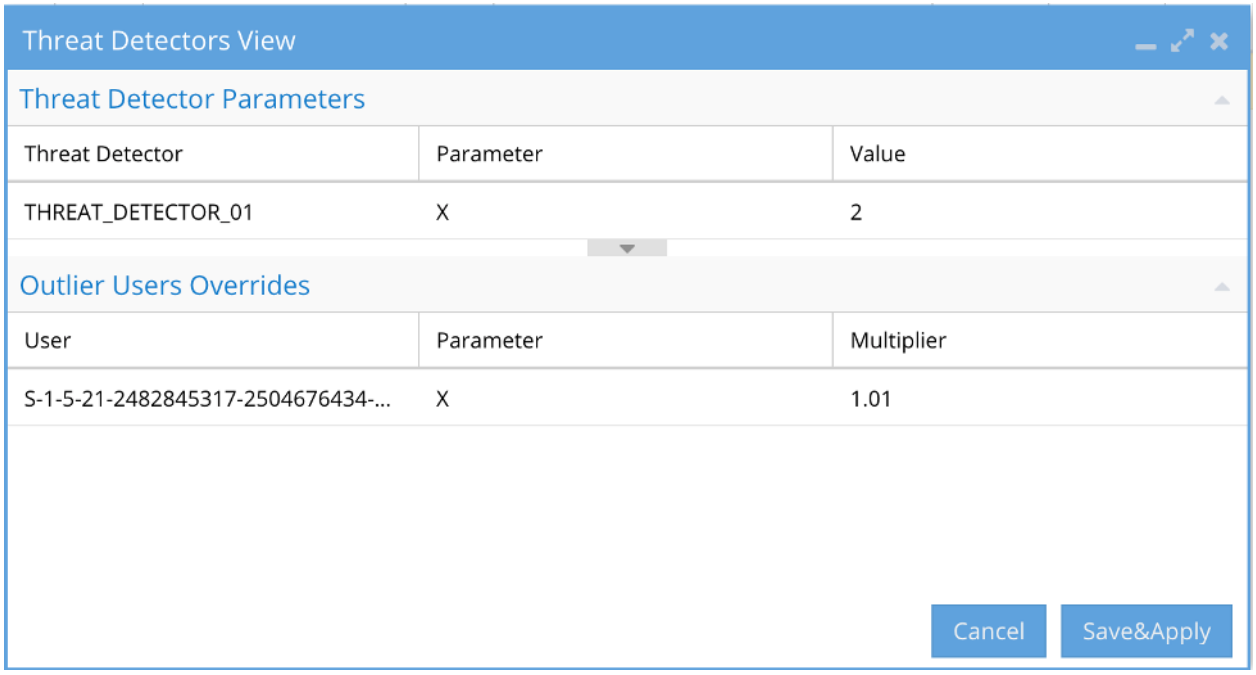

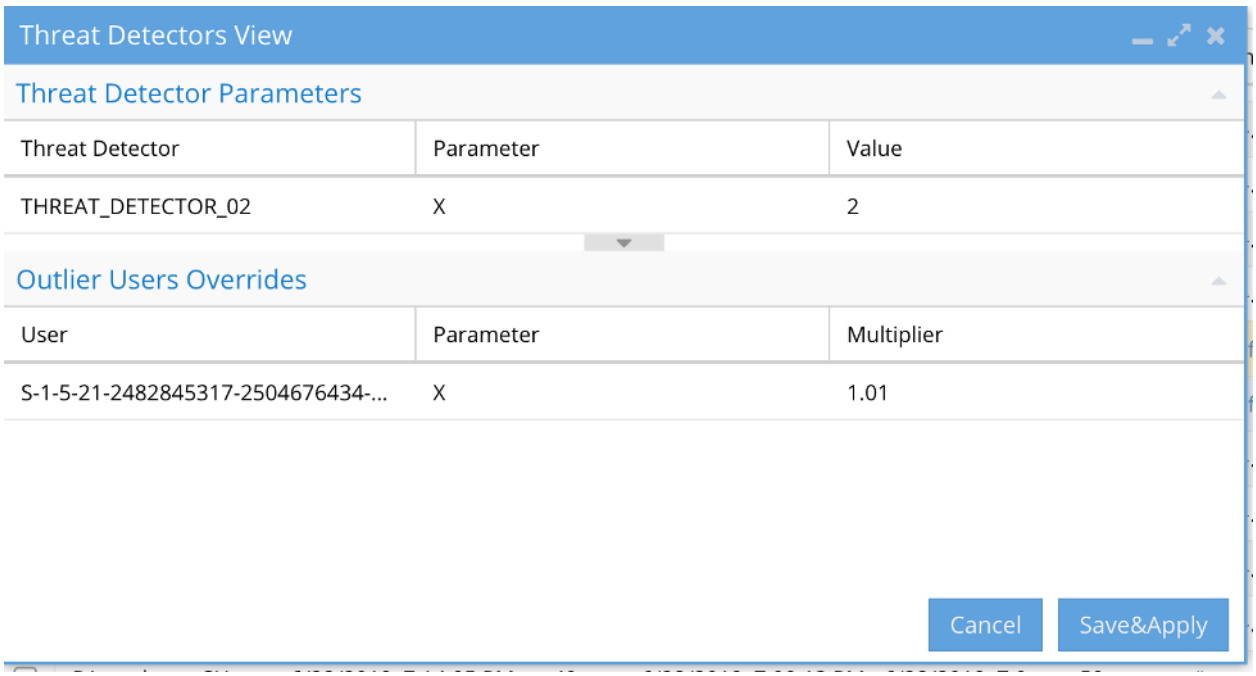

How to view the AI report User settings applied to Ransomware Defender Configuration

1. Open the Ransomware Defender icon
2. Under settings click the False Positive tab
3. example below of the top 10% user override settings applied the the AI reports
4. 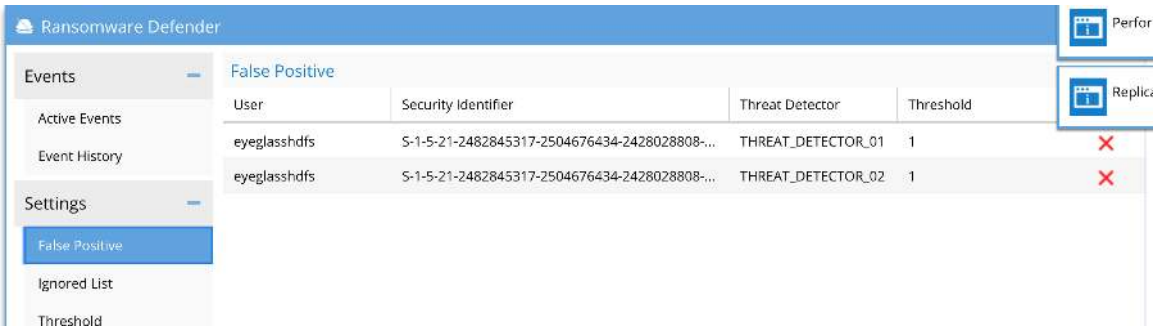

Saved Queries Tab

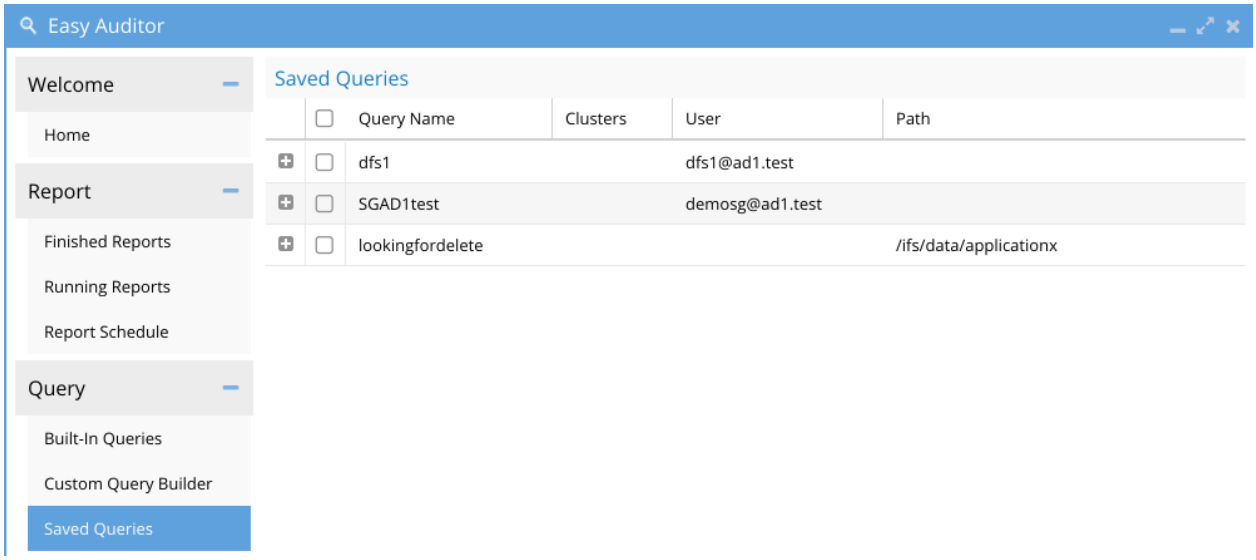

How to load a saved search

1. Use the saved query tab and select a query and click load query
2. Execute the search with search button

How to delete a saved search

1. Select one or more saved queries with a check box
2. Click delete

How to Load a Query and Search

1. Select checkbox of the query
2. Click Load Query
3. Then click search from the search UI

How to Load and Schedule a Query

1. Select the checkbox of the query

2. Click the Load Scheduled Query
3. On Schedule tab set the interval the query should run and other schedule settings

How to delete a saved Query

1. Select the checkbox of the query
2. Click the Delete Query

How to delete run a Query as a report

Use this option to return all rows and generate a CSV of the results versus preview

1. Select the checkbox of the query
2. Click Run Report
3. Then go to the running reports tab to monitor completion
4. Results are available on the Report history tab

Running Reports Tab

Shows all active running report jobs and details of the running report along with duration and status. Use this tab to monitor a running job and the duration of the running report. Click the link to see the finished report.

Running Reports Monitoring

How to Cancel a running report

1. Click the cancel link to cancel a running job

2. Once a job is finished a clickable link is displayed to take you to the results on the Finished Report tab.

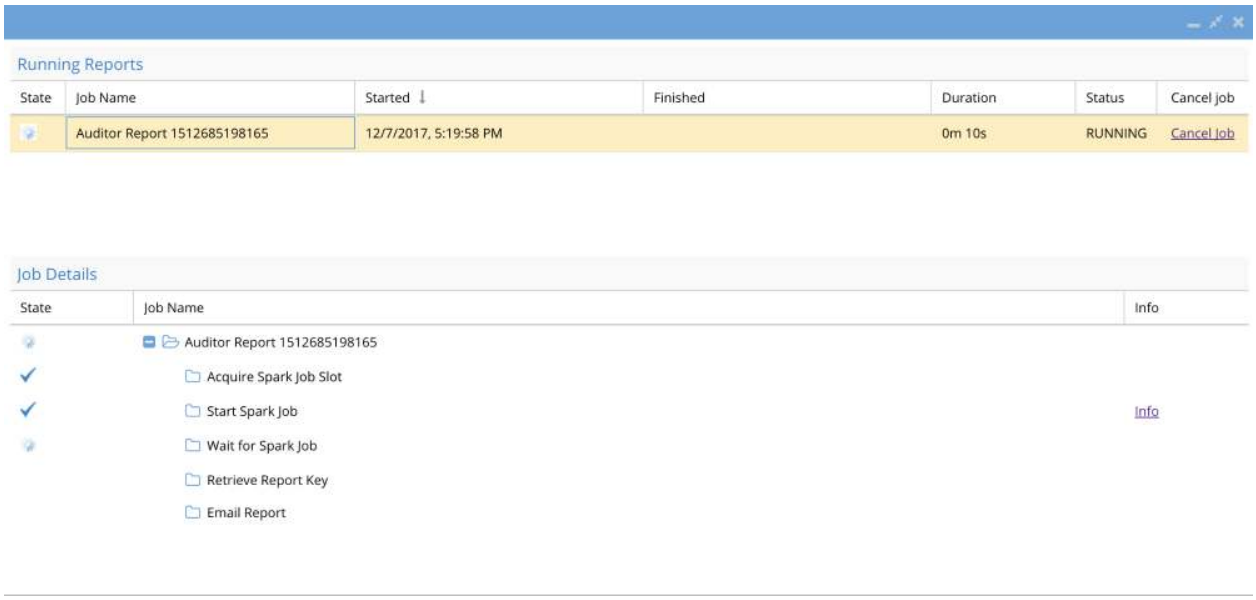

Best Practice:

1. Large reports with a lot of data will take longer to complete, use this screen to determine if changing the query should be done to speed up the report for long running reports.
2. Cancel or view reports that are completed or running from this screen.

How to Monitor Progress of a Running Search

When a search is running you can monitor its percent completion to know how far along the search has progressed and how many file events have been found during the search. You can also see partial results of a running search on the Reports tab with View preview and Download CSV while a search is running.

1. Click on the Running Status in the Running Reports tab to view the real time statistics of the search including % completion,

found events and any failed searches in a portion of the database.

2. 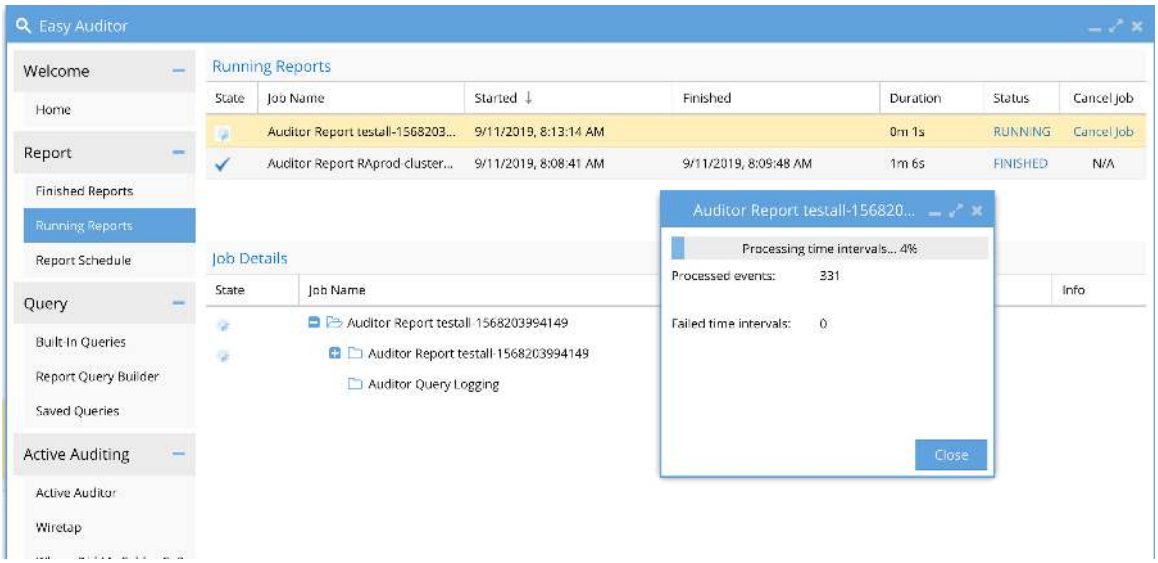

Finished Reports

This tab stores all reports and results for download

The saved reports shows completion time, duration, date range covered, job ID, number of rows of results and a download link to the report. A debug link to the search query details is now available for support purposes.

For long running queries the CSV file can be downloaded to get partial data returned from the query, using the continous search mode.

Finished Reports

Report Id: | Run Type: MANUAL | Status: ALL | Search

Report ID	Status	Completion Time	Durat...	Data Time Range	Event...	Download Spark Logfile	Run Ty...	Viewer	Download
usagerepo...	SUCC...	5/9/2018, 3:56:55 PM	18.75s		1	"app-20180509195638-0000"	MANUAL	View	Download
dfs1last30	SUCC...	5/8/2018, 7:58:58 PM	1.82s	4/9/2018, 11:53:22 AM - 5/8/2...	5531	undefined	MANUAL	View	Download
quicksearch	SUCC...	5/8/2018, 4:33:58 PM	23.77s	5/8/2018, 8:00:01 AM - 5/8/20...	50000	undefined	MANUAL	View	Download
nfsuserdfs	SUCC...	4/21/2018, 10:19:21 AM	10.48s	4/21/2018, 10:15:38 AM - 4/21...	16	"driver-20180421141907-0000"	MANUAL	View	Download
stale	SUCC...	3/3/2018, 7:58:29 PM	5.07s	2/24/2018, 7:57:59 PM - 3/3/2...		"driver-20180304005801-0001"	MANUAL	View	Download
stale	SUCC...	3/1/2018, 7:28:15 AM	5.89s	2/22/2018, 7:27:36 AM - 3/1/2...		"driver-20180301122739-0000"	MANUAL	View	Download
sg	SUCC...	2/19/2018, 9:23:55 AM	1.04s	2/12/2018, 10:02:49 AM - 2/19...	2929	undefined	MANUAL	View	Download

Filter Reports tab

Use the Search button and dialog box to filter the list of display reports.

1.
2. Select all to search and filter any Report ID, Run type (Manual or scheduled), Status (success or failed)
3. Example to find all reports that begin with User . Enter **User** in the search box.
4. Example to find all the scheduled reports enter **scheduled**
5. **NOTE: default filtering will NOT show scheduled reports and only shows manually executed reports, to see scheduled reports filter**

the report. This was done incase scheduled searches had 100's of reports that would make it hard to find manual searches.

Report Schedule Tab

Use this tab to view , load and delete scheduled queries.

NOTE: See tested feature limits for scheduled search limits

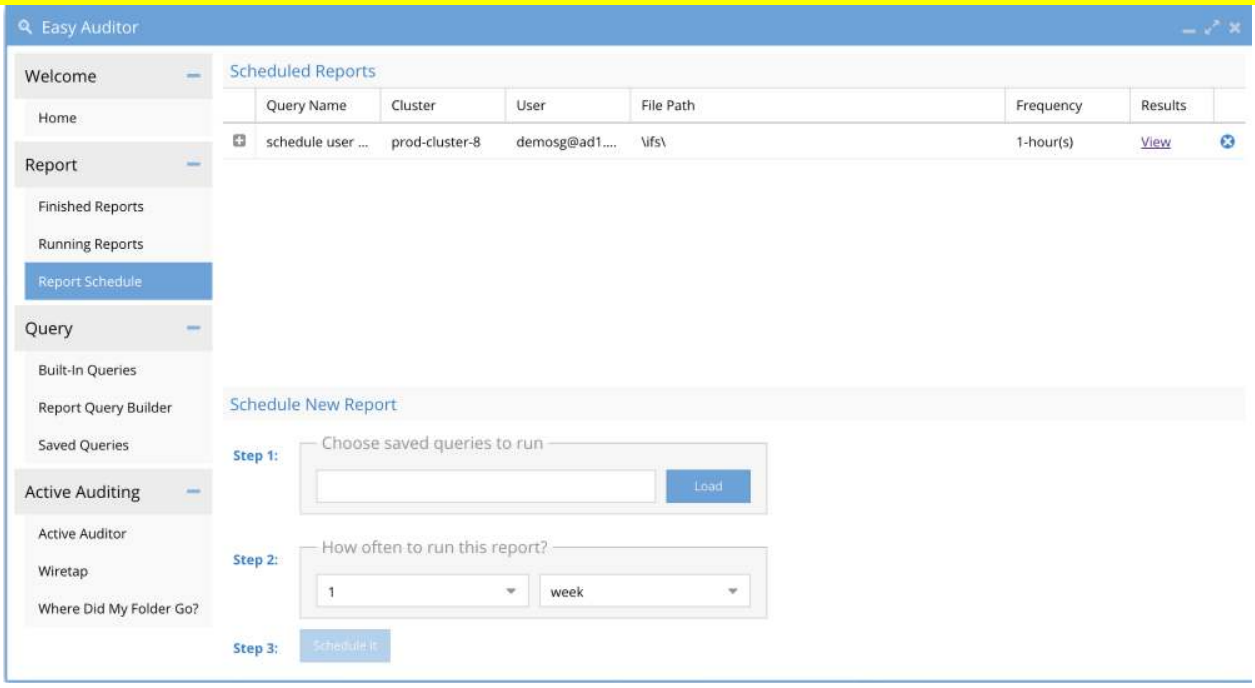

How to delete a saved Scheduled Query

1. Select the Schedule checkbox
2. Click Delete Schedule

How to load a Query and set a Schedule

1. Click the load saved Query
2. Select a query from the list
3. Set the schedule

4. Click Schedule to save

How to edit a saved Scheduled Query

1. Select the schedule from the list with the checkbox
2. Change the schedule
3. Click the schedule button

Active Auditor Tab

This tab is for configuring real-time audit features, to secure and protect data from delete and data loss on secure shares. Any active triggers are displayed on this page with relevant information for the active audit trigger.

Active Auditing Tab Overview

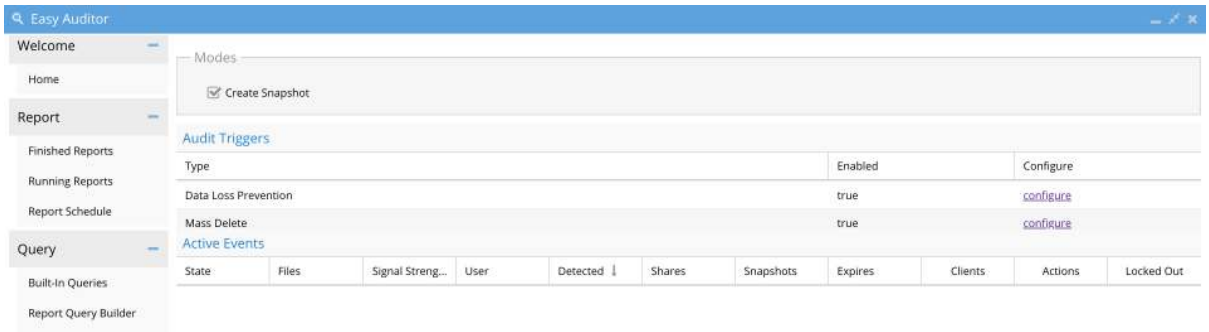

Overview

The policies on this tab are a policy that the ECA cluster will execute in real-time as events are processed.

The feature allows for per user monitoring of file deletes or data copies upto threshold over a period of time

The main tab indicates if the audit feature is enabled or disabled.

NOTE: These triggers are supported for SMB users only. Not supported for NFS client connections

How to Configure Custom Real time Audit Policies

Overview:

This feature elminates most of the reasons to run reports and automates security of almost any scenario. This feature allows administrators to build simple or complex policies that identify patterns, users, actions or even network access to the file system. This solution is designed to elminate running reports in a reactive security mode and enables for the first time a proactive security solution for auditing.

This feature also includes predictive threshold monitoring that takes the threshold crossing settings and uses rate prediction to provide more accurate detections.

How to Manage customer Real time Audit policies

1. **Select the Active Auditor Configure option for Custom real-time policies**

a. 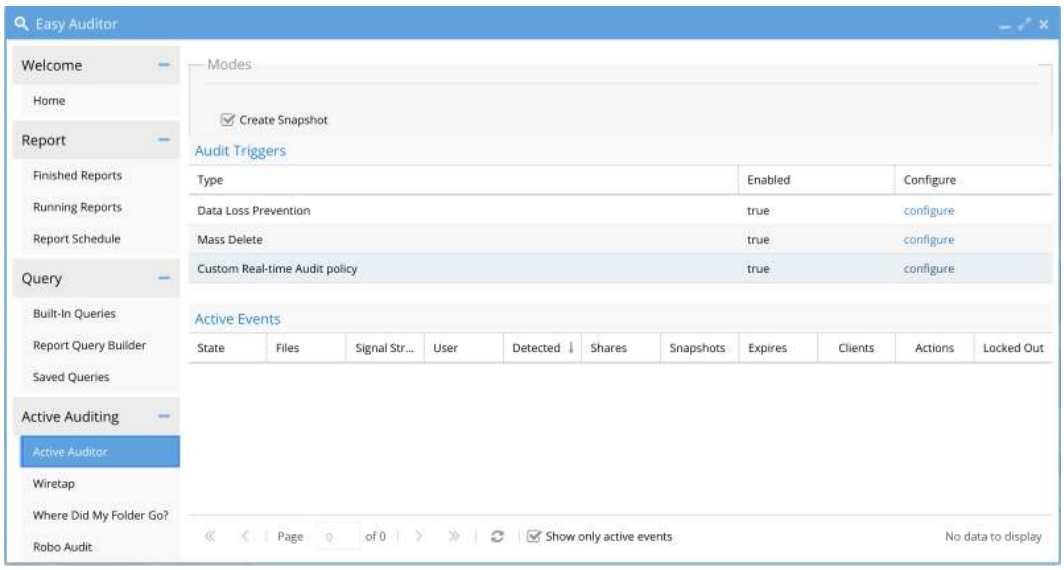

2. The image below shows all configured triggers. The enable check box enables all trigger processing. A trigger can be edited or deleted.

3. 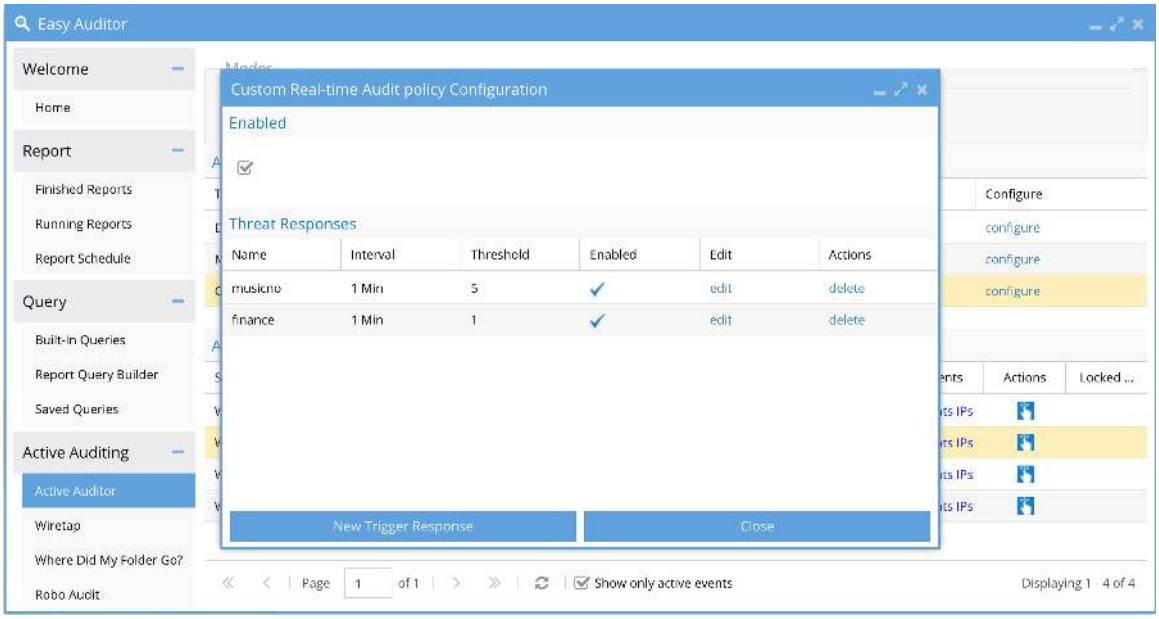

4. To Create a new trigger click the New Trigger Response button.

How to Configure real time audit policies

1. Click Configure in the Custom Audit policies screen
2. Click New Trigger Response button

3. Enter a unique policy name. Note: This name will be used in email alerts and syslog messages. **NOTE: Only characters and numbers, no spaces or special characters in the name.**

   a. 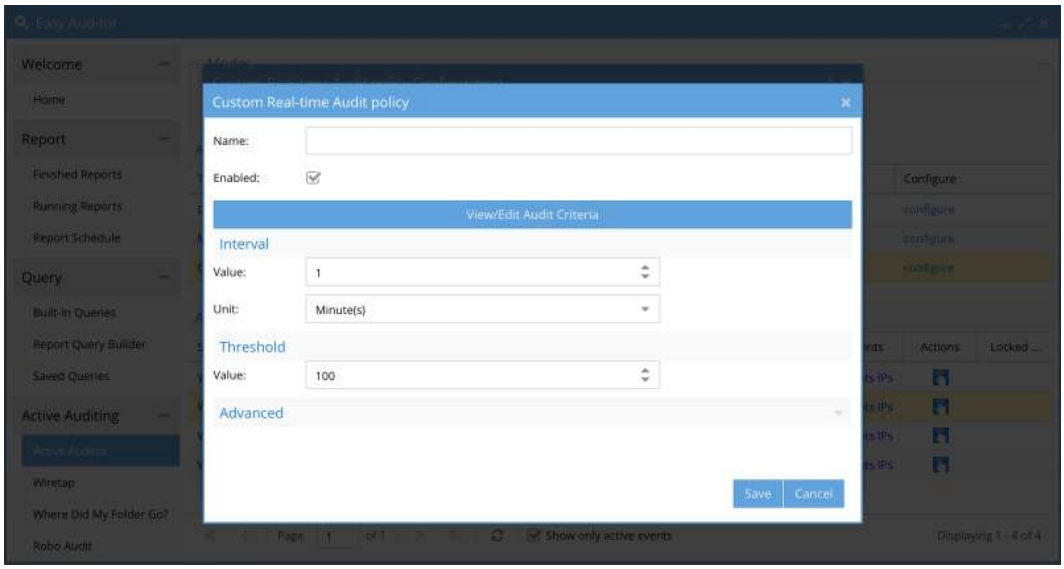

4. The enable check is defaults to enabled. A trigger can be disabled temporarily if required.

5. The **Unit section** is the period of time that the policy definition will monitor events to determine if the condition defined was met. This is settable in minutes from 1 to 60 minutes.

6. The **Threshold** section is the number of times the condition must be met within the Unit in minutes to trip the policy. Until this is true no alert will be sent. This setting can be tuned to ensure the policy does not trip to easily or change the Unit section to monitor over a longer period of time. Valid values is any number greater than 0.

   a. **NOTE: This threshold will be managed by predictive rate monitoring feature that calculates rate of detection to determine if the policy will be true at the unit value set in minutes. This will provide more accurate triggers without needing to worry about a specific rate for detection. The**

active auditor trigger window wil indicate if the detection was predicted or actual.

7. Optional - Advanced configuration of policy evaluation
    a. 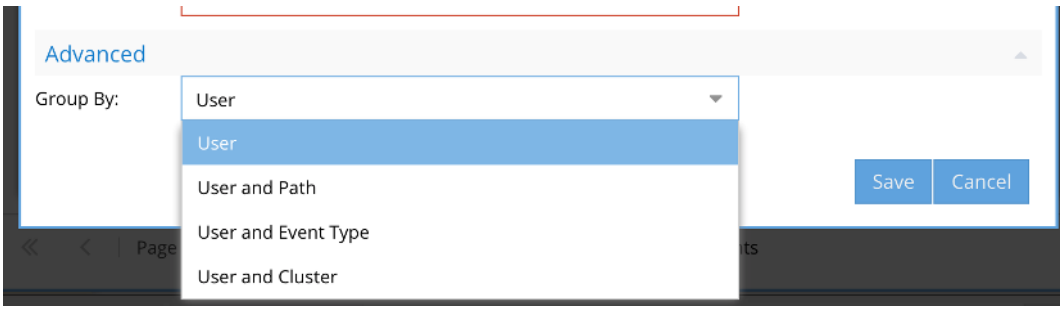

    b. The default is group by user which means the trigger processing evaluates the policy per user.
        i. The user and path option means the trigger processing evaluates the policy per user AND A path meaning a user and a single path must met the criteria configured in the policy.
        ii. User and cluster would evaluate the trigger policy and trip only when the user and a cluster match the criteria, unlike the default setting where events for the user are evaluated from events from any cluster.
8. Real time Policy query definition
    a. Click the button View/edit audit criteria
    b. Audit Event criteria Definitions

i. 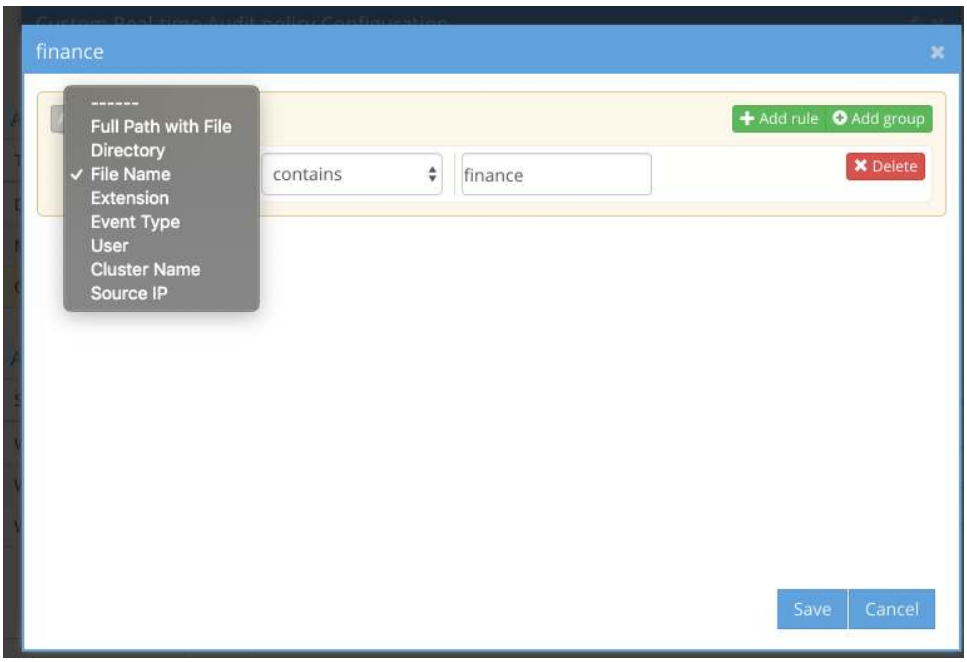

ii. **Full path with file - enter a full path including the file name to monitor an exact file example /ifs/data/zone1/marketing/marketing-plan.docx (Supported comparison options equal, not equal, contains, doesn't contain, ends with, doesn't end with, begins with, doesn't begin with)**

iii. **Directory** - the path where the event will occur example /ifs/data/zone1/marketing **(Supported comparison options equal, not equal, contains, doesn't contain, ends with, doesn't end with, begins with, doesn't begin with)**

iv. **File name -** the name of the file without regard for the path of the file example marketing-plan.docx **(Supported comparison options equal, not equal, contains, doesn't contain, ends with, doesn't end with, begins with, doesn't begin with)**

v. **Extension -** just the extension of the file where the event will occur example pdf or docx (Supported comparision equals, does not equal)

vi. **Event Type - Create,** Read, Write, Delete, Rename (Supported comparison equals)

vii. **User -** the AD user (syntax MUST be DOMAIN\user upper case domain name) to monitor a single user (Supported comparision equals, does not equal)

viii. **Cluster -** the cluster name the event was generated from. This should only be used with multiple clusters that are managed by Easy Auditor (Supported comparision equals, does not equal)

ix. **Source IP -** The source ip or subnet range the event originated. example 192.168.1.80/32 for a single host or 192.168.1.0/24 to identify any host in the 255.255.255.0 subnet. (Supported comparision in, not in)

c. **Using multiple conditions with AND and OR**

i. The simplest method to build criteria is using and or operands. For example The file name AND user X, This directory OR user X

d. Using the Grouping feature for policy criteria

i. This allows grouping several criteria together and selecting AND or OR for the group of criteria. example file name AND path MUST be seen during

the unit value in minutes, OR means any of the criteria in the group need to have been seen within the Unit in minutes

e. Complex combinations of individual criteria and groups. It is possible to combine a group plus individual criteria using the AND and OR operands within the group and between the group and individual criteria.

f. How the UI elements work for real time audit policy rules

   i. The black selection indicates that AND or OR has been selected for rule criteria

   ii. A Group has its own AND or OR operand selector for the group inside the group

   iii. Delete a group or rule with the delete button next to the element you want to delete from the policy

   iv. The gray bar on the left side of the UI indicates a seperator between rules or groups and rules. This will indicate how the AND or OR will be evaluated .

   v. The + rule and + group icons are used to add a group or rule to the policy

vi. 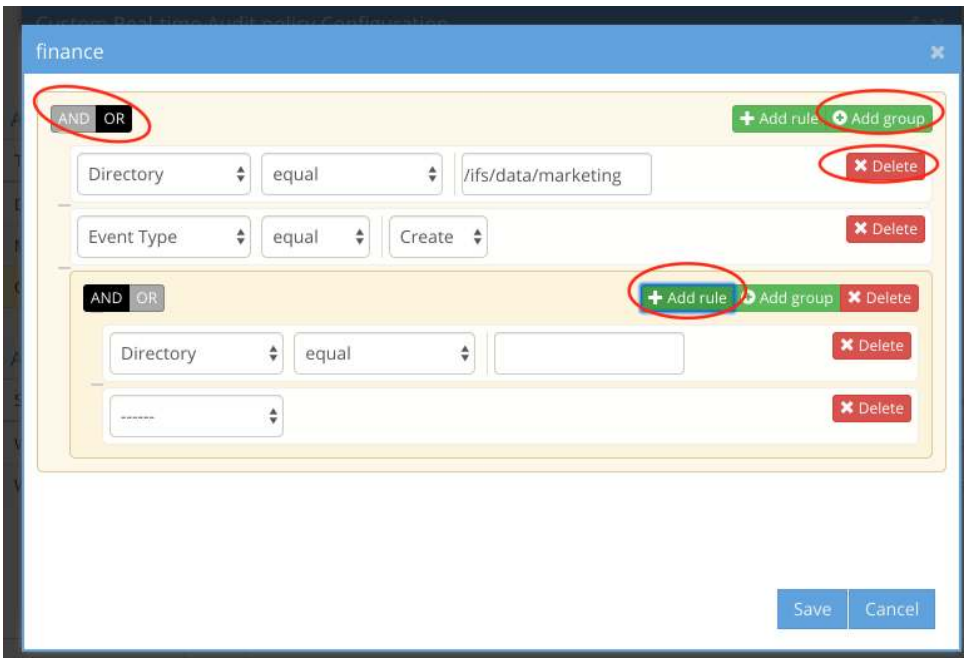

9. UI Examples

10.

a. Multiple individual rules using OR between rules

b. 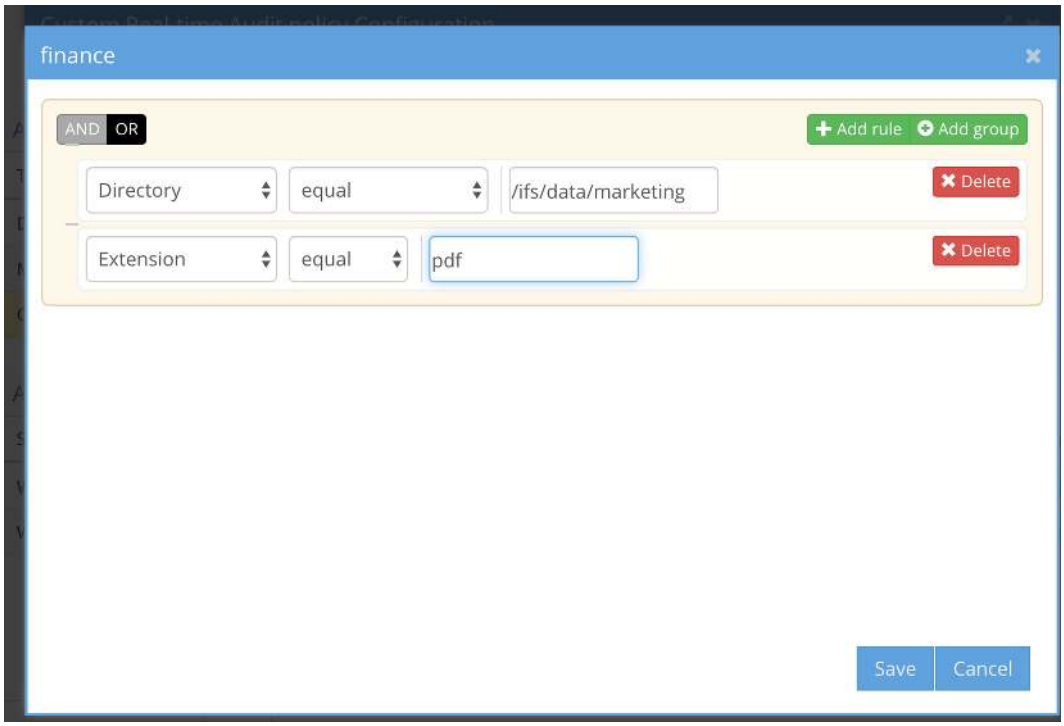

c. Multiple individual rules using AND between rules

i. 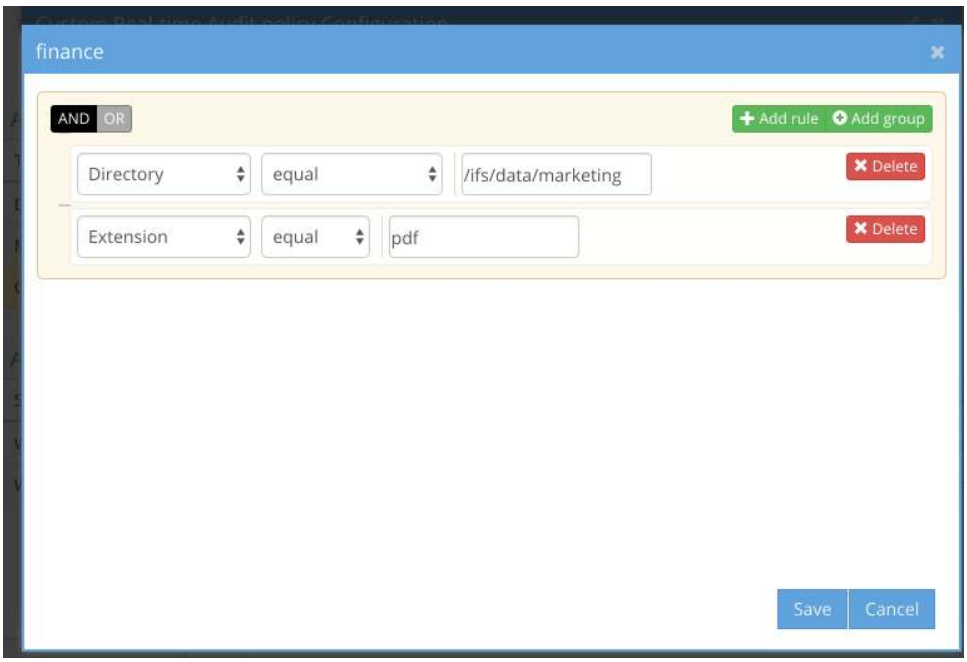

d. A group of rules using AND between rules

i. 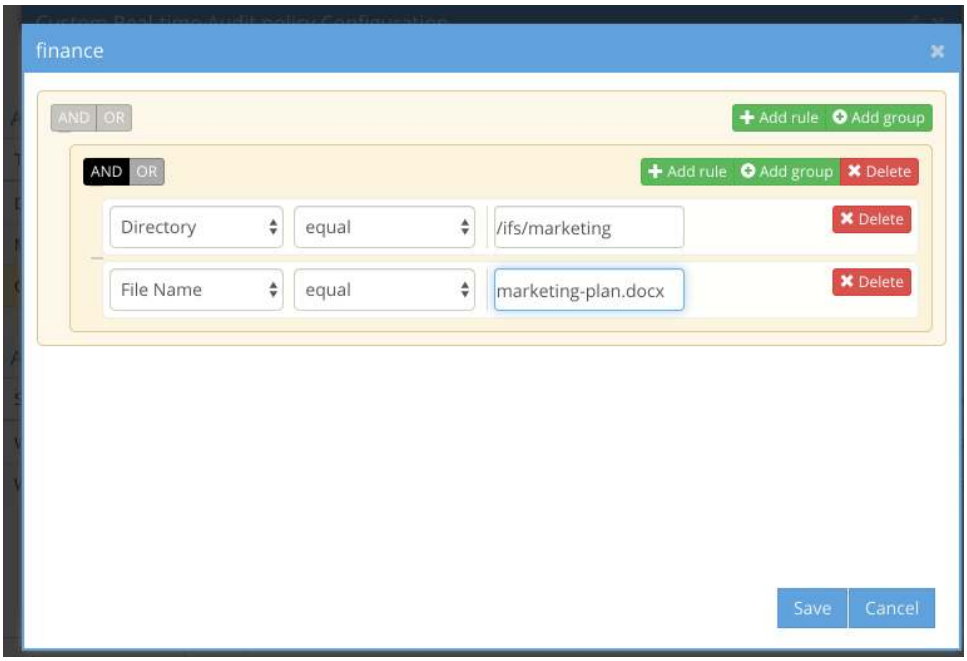

e. A group of rules using OR between rules

i. 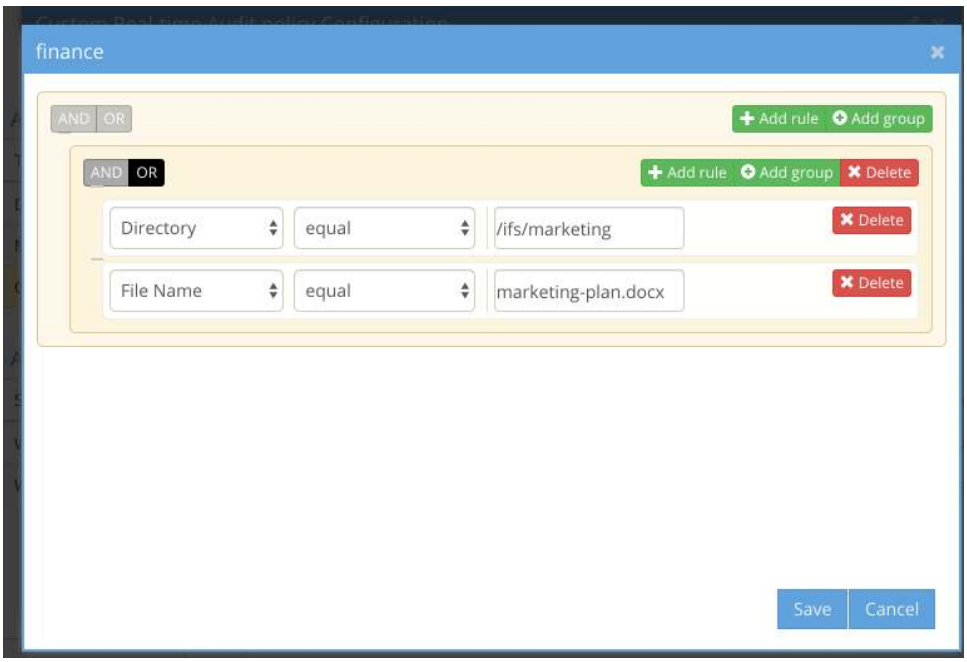

f. A group plus comparison to an individual rule using AND between the group rules and an AND between the group and single rule

i. 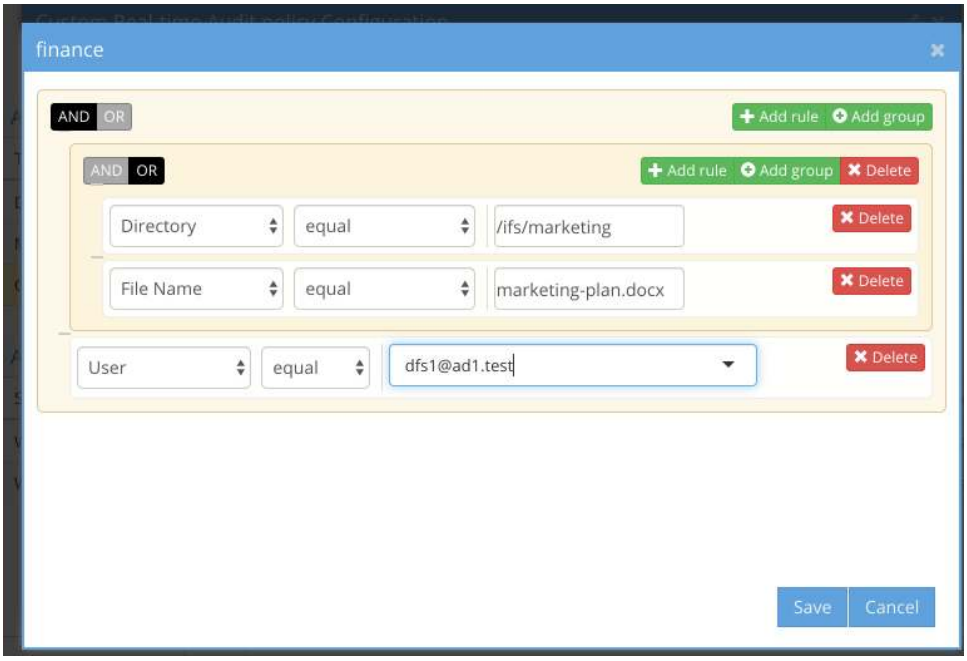

11. Save the definition of the policy to activate the rule. NOTE: the policy will be sent to the ECA immediately for processing events

Auditor Active Trigger Use Case Examples

This section covers how to use Active triggers to solve specific security monitoring use cases.

HoneyPot Detection

1. This provides a method to create a honeypot SMB share to detect an insider threat looking for open shares to read data.
2. Create an SMB share with everyone full control with a name that may attract someone looking around the network for open sharers. Add fake file names that look like sensitive data.
    a. example SMB share named - finance on path /ifs/data/honeypot/finance
3. Open Easy Auditor Active Triggers section and add a custom trigger as per screenshot. Then click the view edit button.

a. 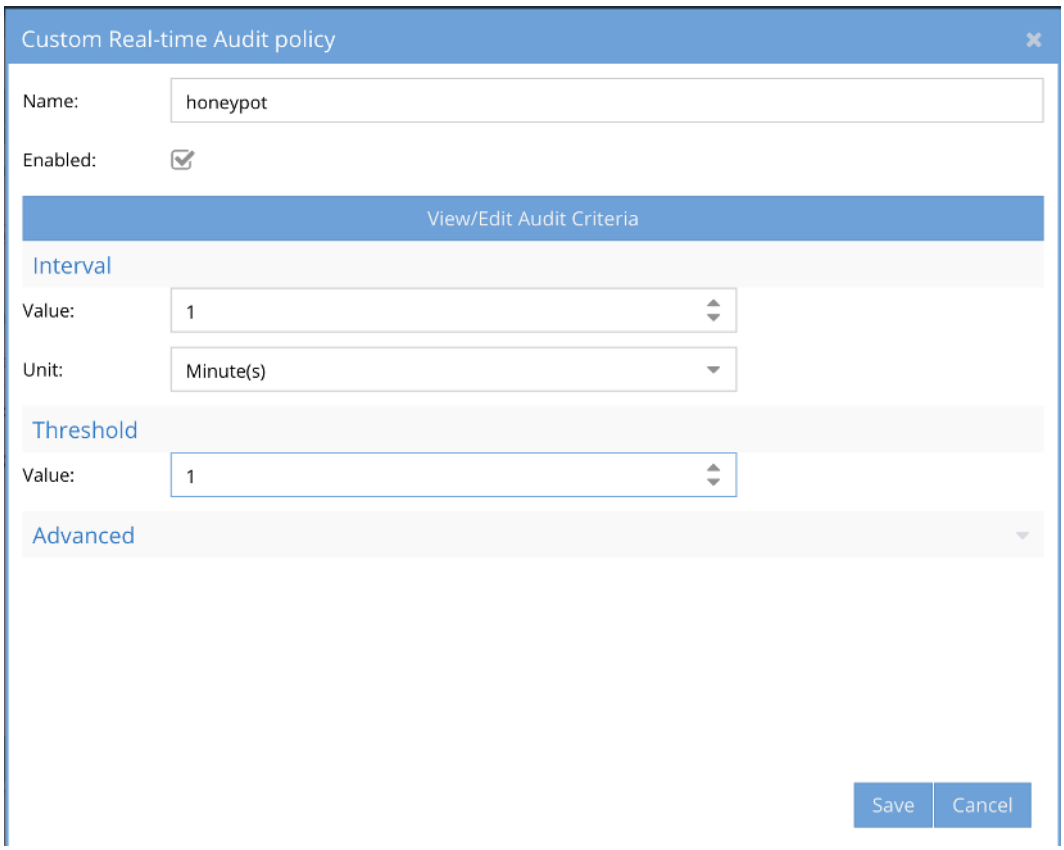

b. Use the rule editor to add the path to the trigger in this example /ifs/data/honeypot/finance and click save.

c. 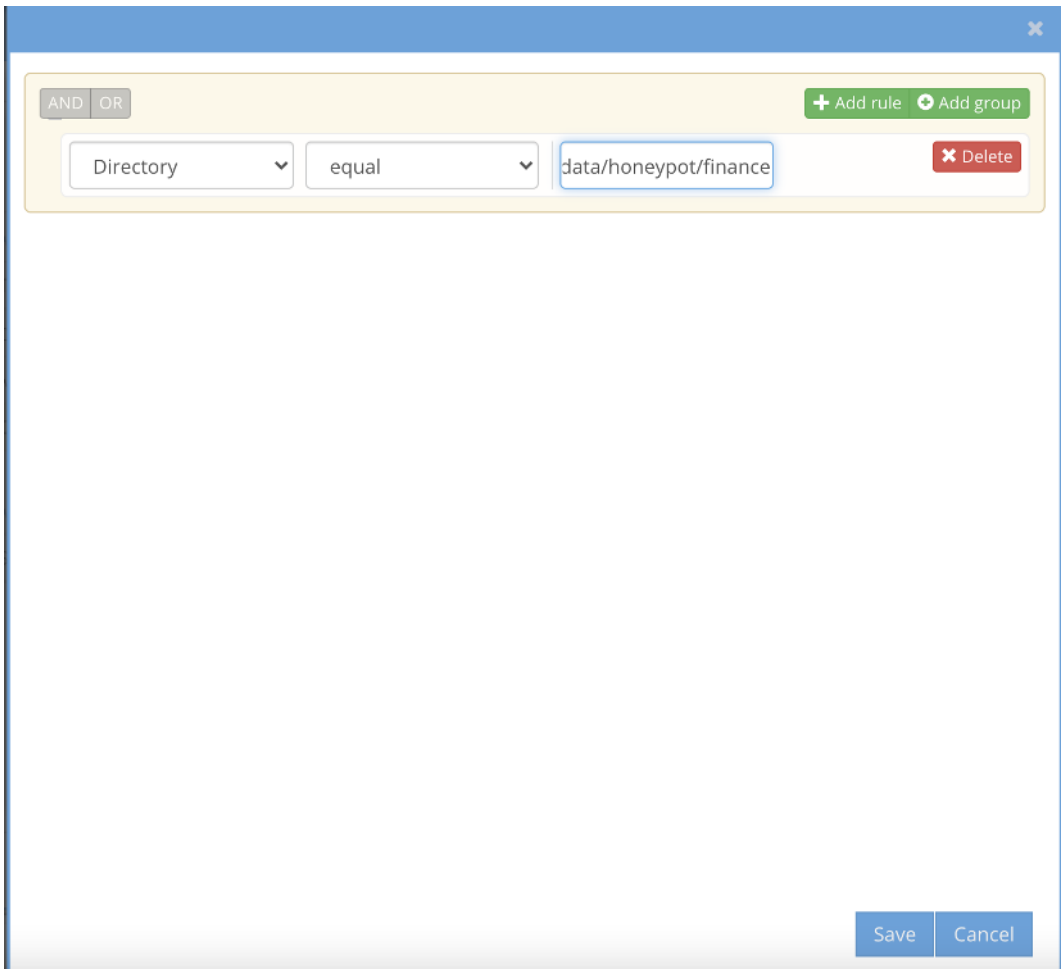

d. Now click Save to save the trigger definition.

e. Test the trigger by accessing files on the share. A detection will look like below.

f. 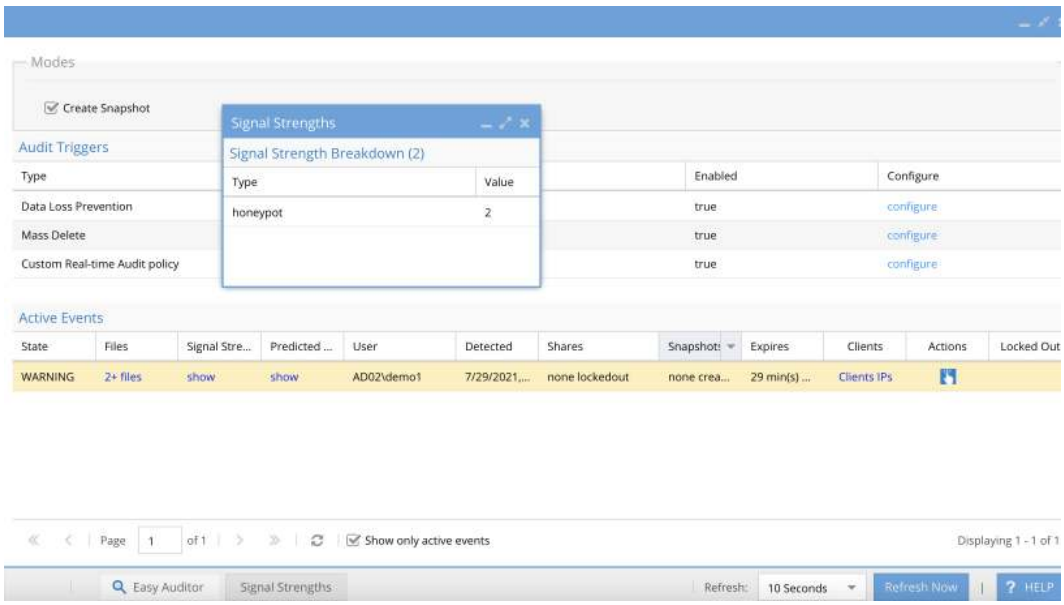

How to configure Mass Delete protection

Overview:

Monitors users deleting files on any share or export upto administrator defined threshold over administrator defined time period. The feature counts deletes by user using the ECA cluster real-time detectors and raises an alert when the policy is violated.

- Provide visibility into delete actions on the file system before user cases are opened regarding deleted files.
- Simplifies recovery of accidental deletes from auto applied snapshots
- Provides security monitoring of deleted fails in real-time

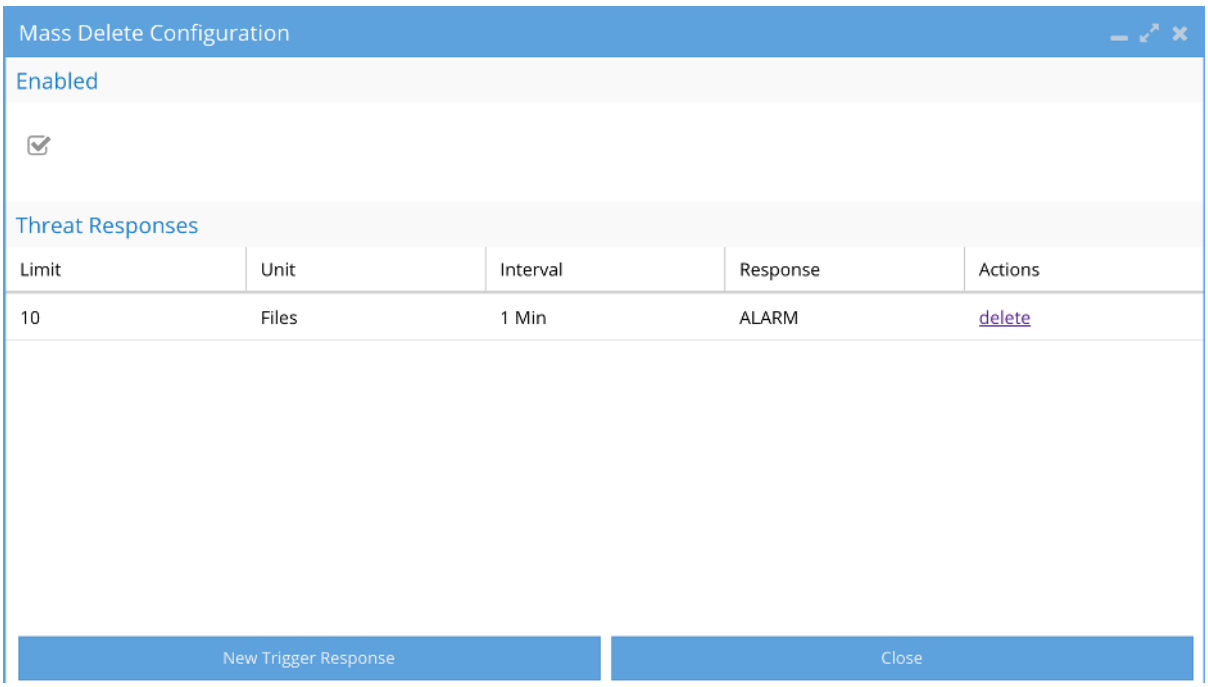

1. Click Configure
2. Enable the audit feature
3. Click **New Response** to create policy to set one or more responses to crossing a file delete threshold per user.
4. 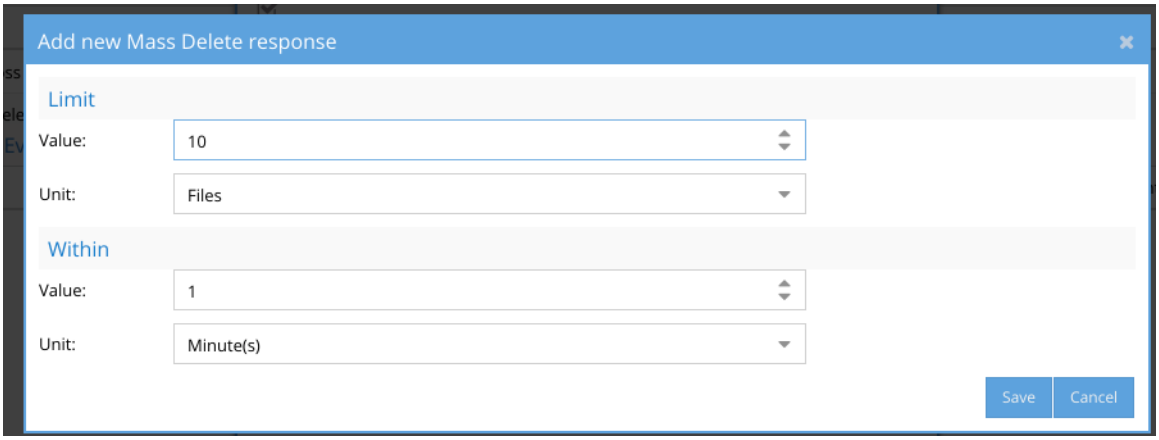

5. Threshold for number of files to be deleted. **NOTE minimum value is 1000 within 1 minute, the time period cannot be changed in this release**

6. **Best Practice:** 1000 files to start and adjust higher if too many notifications are sent.
7. Time period of which the deletes should occur within.
8. **Best Practice:** The rate at which deletes occur does not change the severity of the delete. The goal is to set at a rate that detects many deletes in a short period of time. This value can be adjusted up or down depending on the number of notifications that are sent.
9. Possible Response Actions:
    a. Email alert - Sends an alert of the user, path and crossing of the policy criteria
    b. Snapshot the path(s) being deleted to provide a restore point. The snapshot has a time to live of 48 hours be default.
10. **Possible Actions on an Active event**
    a. **User Lockout** If selected the user that trips this trigger will have SMB access denied on the first share above the path being monitored.
11. Best Practice: Enable both email alert and snapshot response for mass delete.
12. Click close

How to configure Data Loss Prevention

Overview:

Monitors users copying files on any secure share or path. This assists with real-time monitoring of secured data from bulk copy operations that are not authorized or indication of potential data loss scenario. The feature will monitor the capacity of the file system path, using an auto applied accounting quota and allows administrator to set % of the data any user can read from the path before the audit trigger is detected.

- Automatically monitors secure data access by users
- Alerts administrators of access by user with date, time, ip address of the access
- Protects against bulk copy of secure data
- Provides visibility into user data access
- Proactive security measure to simply auditing of secure data.
- Secures sensitive data from insider threats

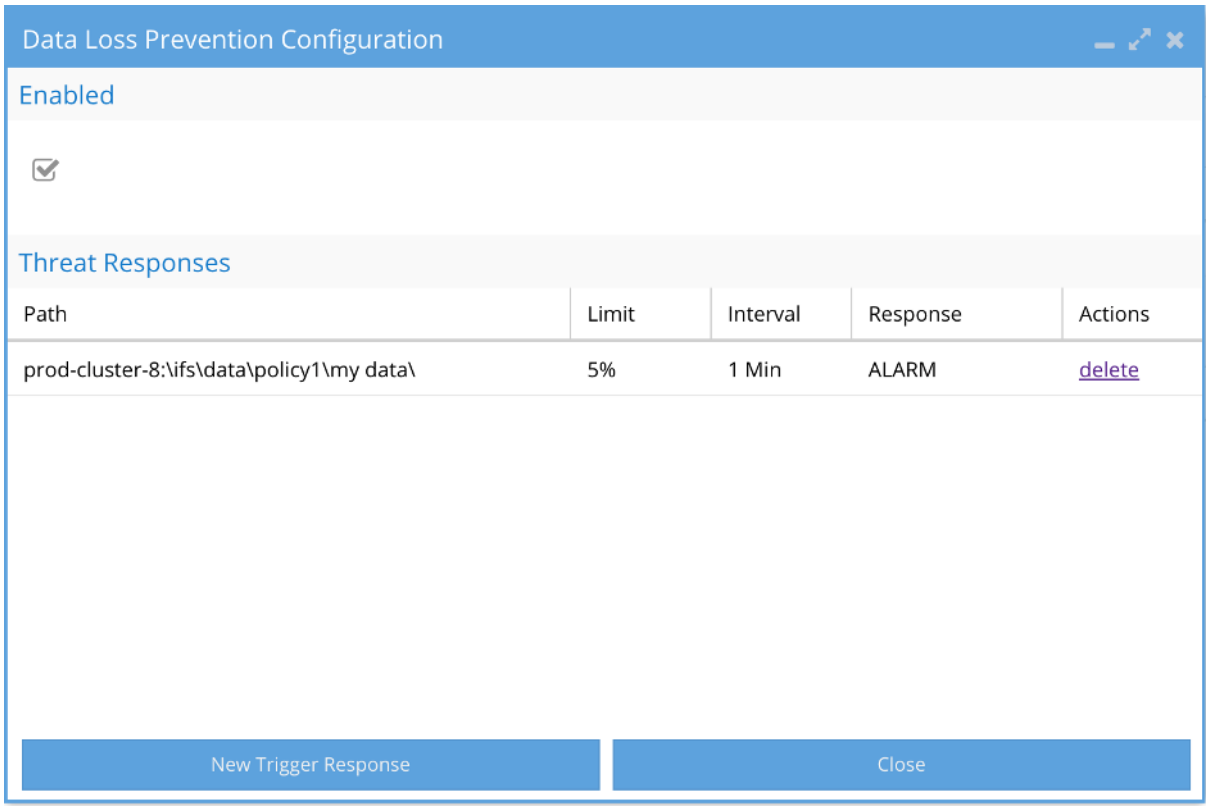

1. Click Configure button on main active auditor screen
2. Enable the audit feature
3. Click **New Response** to Create policy to set per user response actions when crossing data copy % threshold of the total data on the file system path.
4. 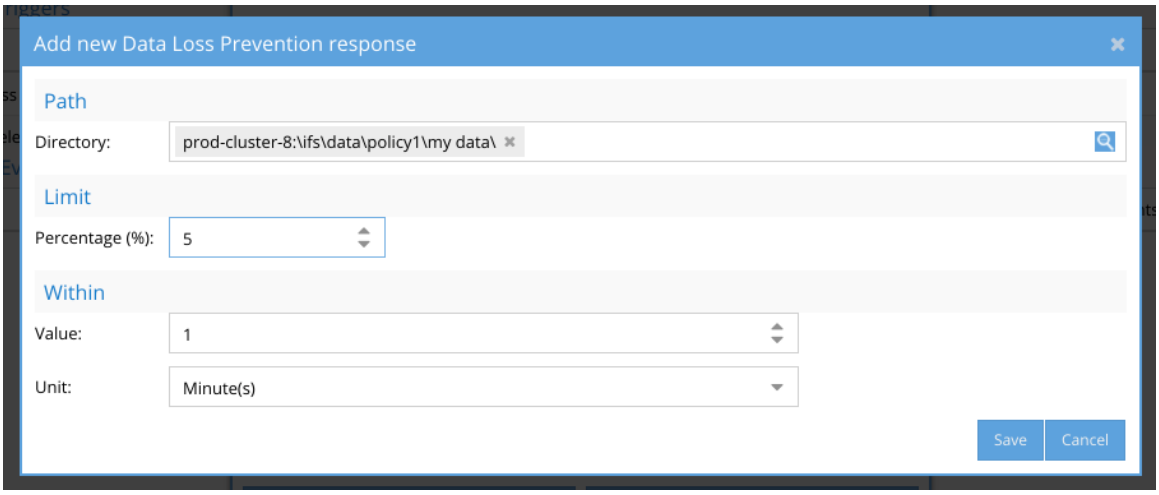

5. Enter the path in the file system to monitor. Best Practise: Enter a path equal to a share path being monitored.
6. Threshold **%** -user the slider bar to select
7. **Best Practice:** 10% is a good starting point to catch copy actions by users. **NOTE: in this relese minutes is not selectable and defaults to 1 minute**

3. **Time period** over which the copy will cross the threshold.

1. **Best Practice:** The rate at which copying occurs can affect the trigger detection. The goal is to set at a time period low enough to ensure the threshold is crossed. This value can be adjusted up or down depending on results of trigger testing.

2. **Possible Response Actions:**

    a. **Email alert** - Sends an alert of the user, path and crossing of the policy criteria

b. Snapshot the path(s) being deleted to provide a restore point. The snapshot has a time to live of 48 hours be default.

3. **Possible Actions on an Active event**

   a. **User Lockout** If selected the user that trips this trigger will have SMB access denied on the first share above the path being monitored.

   b. Best Practice: Enable both email alert and snapshot response for mass delete.

4. Click close

How to manage Active Auditor Events

When Mass delete or Data Loss prevention events are detected the following options are available per incident. When a user has an entry in the active events any triggered event will be added to the same active event. An active event will stay active until acknowledge (action menu) and can be mark as recovered status (action menu).

An active event has the following information:

Active Events

State	Files	Signal Streng...	User	Detected ↓	Shares	Snapshots	Expires	Clients	Actions	Locked Out
WARNING	57+ files	1/1/1	AD01\dfs1	5/6/2018, 9:5...	none locked...	4 clusters	59 min(s) 46 ...	Clients IPs		

1. State of the security event
2. List of files that tripped the detector
3. Signal Strength lists the active audit detectors that have been detected for this user ID. In this example both detectors have been tripped for this user.

a. Signal Strengths

Signal Strength Breakdown (2)

Type	Value
Auditor - Mass Delete	1
Auditor - Data Loss Prevention	1

4. User id column is the AD user name
5. Detection date and time
6. Share lockout list (if action menu is used to lockout, the list of shares will be listed here)
7. Snapshot names (if snapshot check box is enabled, Mass Delete and Data loss Prevention triggers will automatically snapshot shares the user has access to based on AD group membership
8. Expiry minutes - How long before the timer on a detection will trigger secondary actions. No secondary actions are defined in this release.
9. Client ip address - This is the source ip of the machine the user was logged into.
10. Action Menu provides a list of actions for this security event.
11. Lockout column indicates if the user account has had lockout applied. that will deny the user access to all shares the have access to.
    a. If the action menu lockout action is applied to the user id. The event is updated. See sampe image showing lockoout

date and time along with a list of shares that have had the lockout applied.

b. 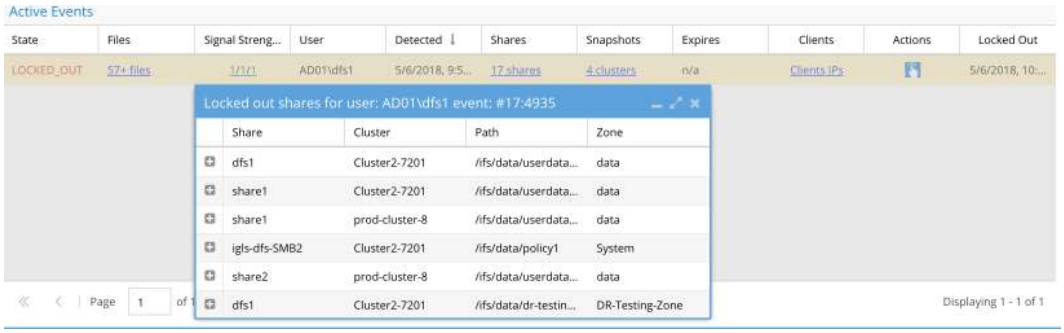

c. The state of the event is changed to Locked Out

d. To remove the lockout use the action menu. "Restore User Access"

e.

Actions available to Manage Active Auditor Events

An event has several possible actions that can be applied, lockout user, create snapshot, acknowledge, comment and flag as false postive

1. Click the action button for an active event.

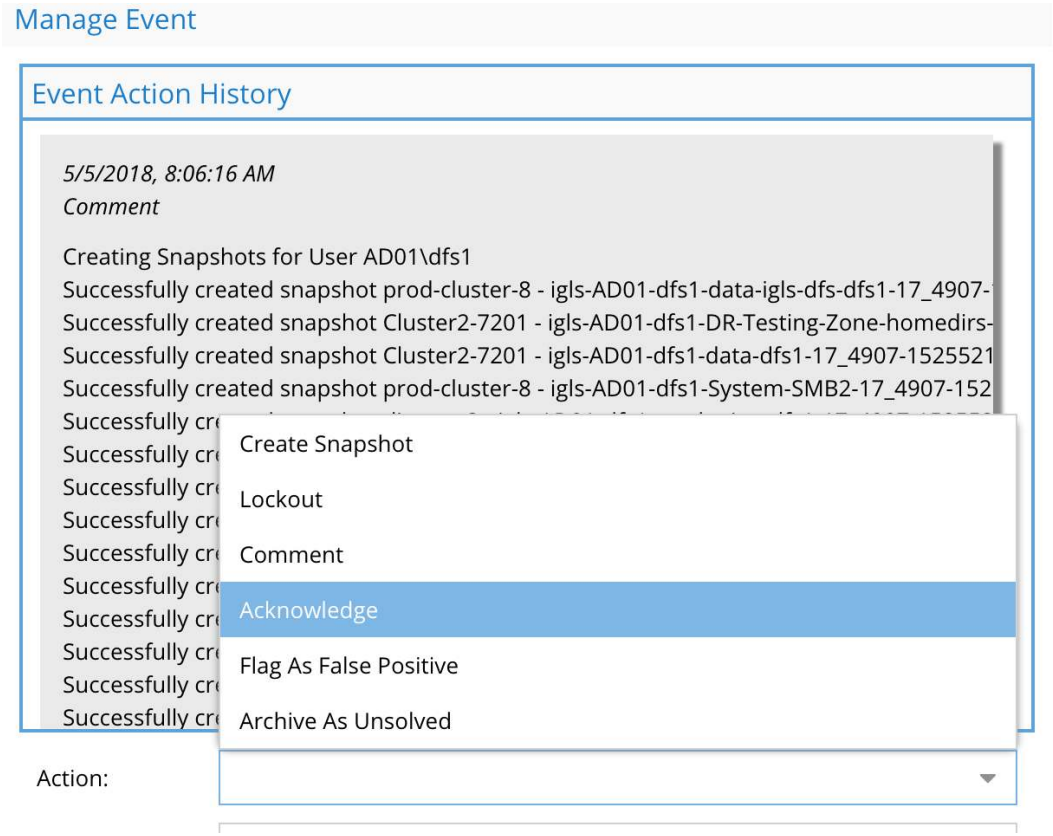

a.

b. Comment - add comment to the event for other administrators to see and store with the event when archived

c. Create Snapshot - only required if auto snapshot is disabled

d. Lockout - apply deny read to All shares the use has access to mount

e. Acknowledge - set status to show the event has been reviewed. Other administrators will know this event has been looked at.

f. Archive as unsolved - No follow up or issue was caused by the event after reviewting the details of the event. The event can be archived using this option

g. Flag as false postive - In cases where the default trigger settings detect an event with too low a threshold , this option sets an override for the specific user id to not trip detection and sets a custom threshold for this single user. Consider changing the detector settings if too many false positives.

How to Archive an active auditor event

To archive an event after all follow up has been completed follow these steps.

1. Determine if any action is required and apply as required. See above action options described above.
2. To clear an event from the active list and move to the historical list.
3. Use the action menu and apply the **acknowledge** event action.
    a. Or use the action menu and appy the **Archive As Unsolved.** (this would be used if no follow up is done on the event and you want to archive without any investigation)
4. Then apply the Action **Mark as Recovered**
5. The event will disappear from the active list. See below how to see historical events.

How to Display active auditor historical events

1. Click the active auditor tab
2. Uncheck the "Show active events"

3.

4. Done

WireTap

This feature is a tool to allow security admins and auditors to easily look at a sequence of events by one or more user actions in the file system to build a complete picture of what happened.

NOTE: Due to event rates and system load only 2 Wiretap filters can be created at a time. If a 3rd wiretap is required one of the existing Wiretaps must be deleted first.

NOTE: Wiretap will only output 25 events per second even if the path or user event rate is higher. The sampled events should provide enough to view the file actions without overloading the UI at rate that cannot be viewed in real time. This also means that some events will be dropped to meet the event rate, this is expected behavior to allow visualization of events as they happen.

NOTE: If the event rate is high on the monitored path the browser requires sufficient RAM and cpu to process the audit data stream. If

the browser seems slow or unresponsive you do not have a fast enough pc. You should filter the events to display fewer events.

The feature relies on the real-time processing of audit events managed by the ECA cluster to decode and stream audit data to Eyeglass wiretap UI.

- The feature allows wiretapping a path in the file system to monitor multiple users accessing data in a specific path in the file system or monitor a user anywhere on the file system.
- A folder can be monitored or all sub folders, file or directory events can be added to a filter to reduce information shown in the UI. Adding a user name will filter all data by user or leave it blank for any user IO on the path.
- NOTE: For high rate paths the data is rate limited to 25 events per second and some events maybe dropped. The browseer cannot recieve data at higher rates and the scroll rate is also too fast to see the events.
- NOTE: by defining a wiretap session the ECA cluster is monitoring for events that match the wiretap. The events are only forwarded to Eyeglass when the wiretap watch window is open.

How to Configure Wiretap

1. 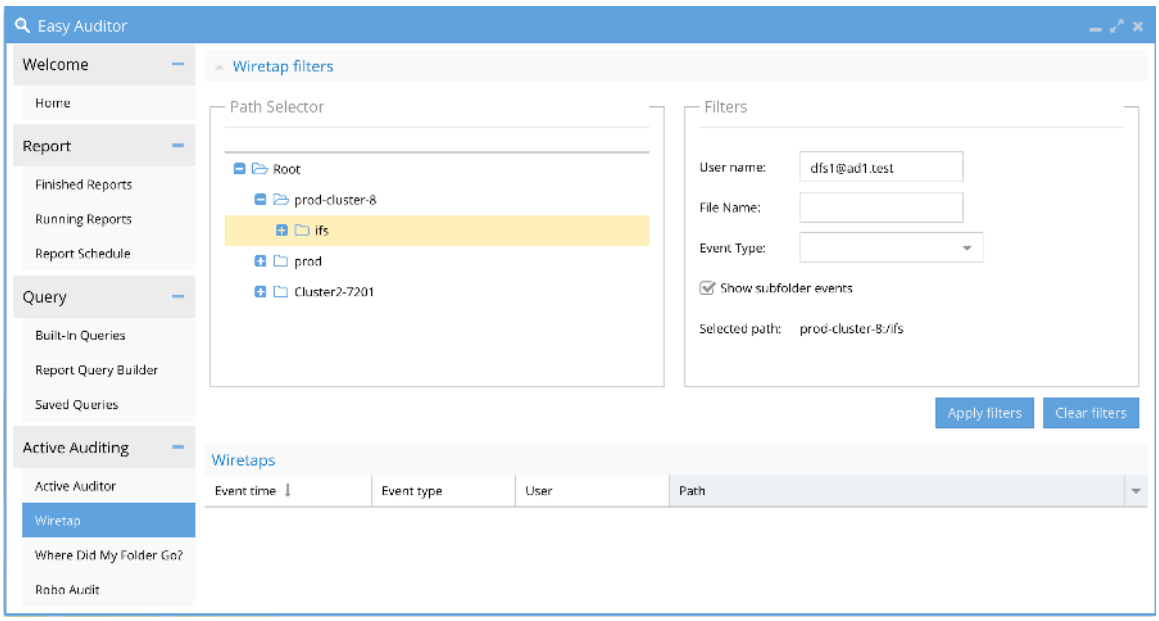

2. Browse the cluster paths to select a path
3. Click sub folders events or disable this to reduce the event rate
4. Enter a user id or leave blank to see events from all users
5. Enter a file name to filter IO on a specific file in the folder
6. Click Apply Filters. This will trigger a message to the ECA to filter events for Wiretap. It may take several seconds for events to appear in the UI.
7. Click clear Filters to remove all settings from the ECA.
8. To Hide the filter UI and get more screen space to see events click the upper arrow to hide or show the filter UI.

9. 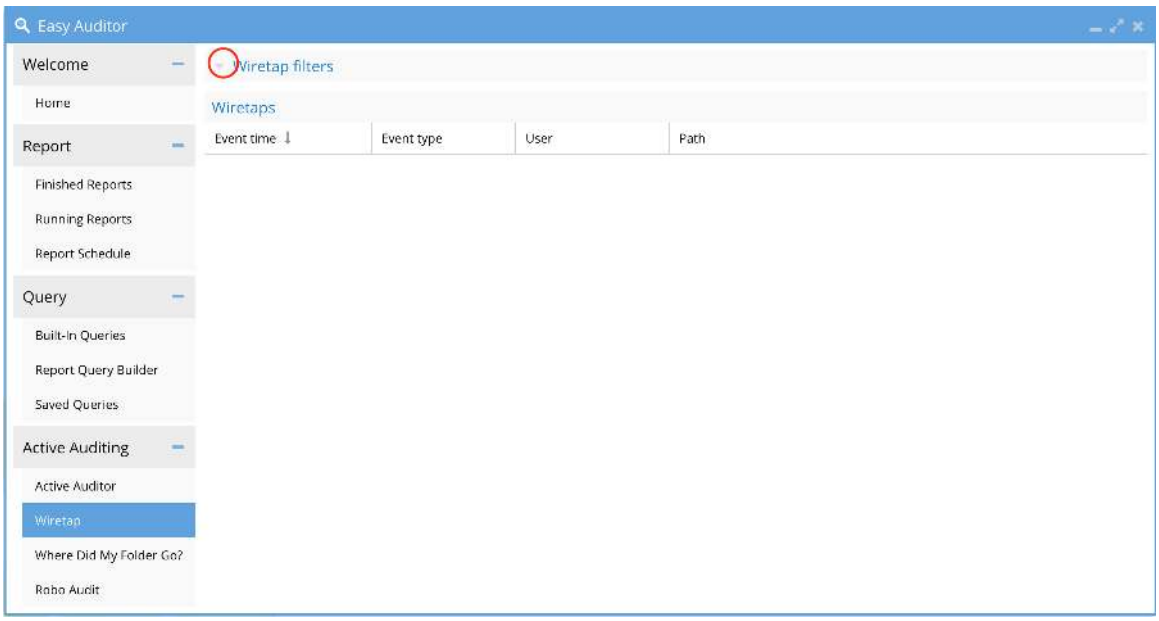

10. WireTap UI with events example

11. 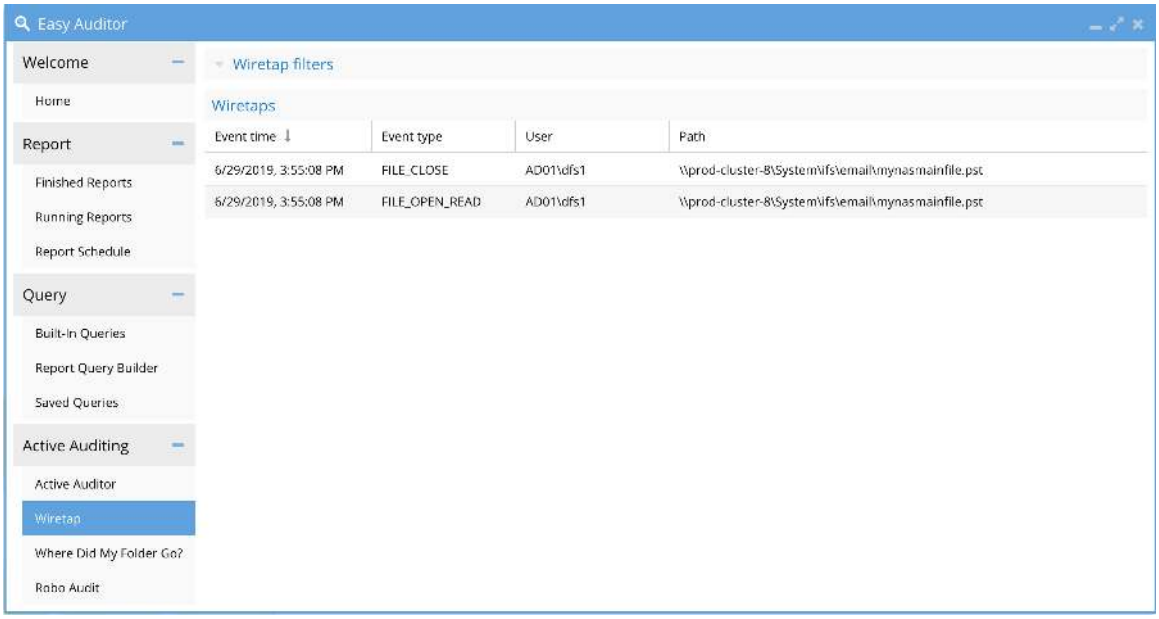

Where Did My Folder go? Tab

Overview

This feature assists with a common issue with folders moved from user drag and drop actions on NAS shares, often resulting in the user

and/or other users unable to locate the files. This turns into a help desk case to locate "missing data", consuming support staff time and effort to locate the missing files. **New in 2.5.5 directory deletes will be reported as well. New in 2.5.6 patch build adds File delete support which is a common use case when single files are deleted.**

This solution accelerates and simplifies addressing missing data requests. The Role based access feature in Eyeglass has a read only role that can be assigned to the help desk to lookup user folder names to assist with locating data.

Limitations

1. 2.5.6 builds after 84 will only return 5000 entries by default and a CLI command will allow increasing this number of results returned. This limit accelerates the search performance.
    a. **NOTE: If the search results indicate Results 5000 (Limit 5000) in the GUI (see example below), you MUST increase the limit and run the search again.**
        i. ssh to Eyeglass vm as admin user and run this command
        ii. **igls easyauditor folderquerylimit set --limit=10000**
2. Copy to clipboard is tested to 5000 entries. The copy to clipboard may not copy all rows when greater than 5000 responses are returned. The solution is too narrow the time period searched or be more specific with the path to be searched.

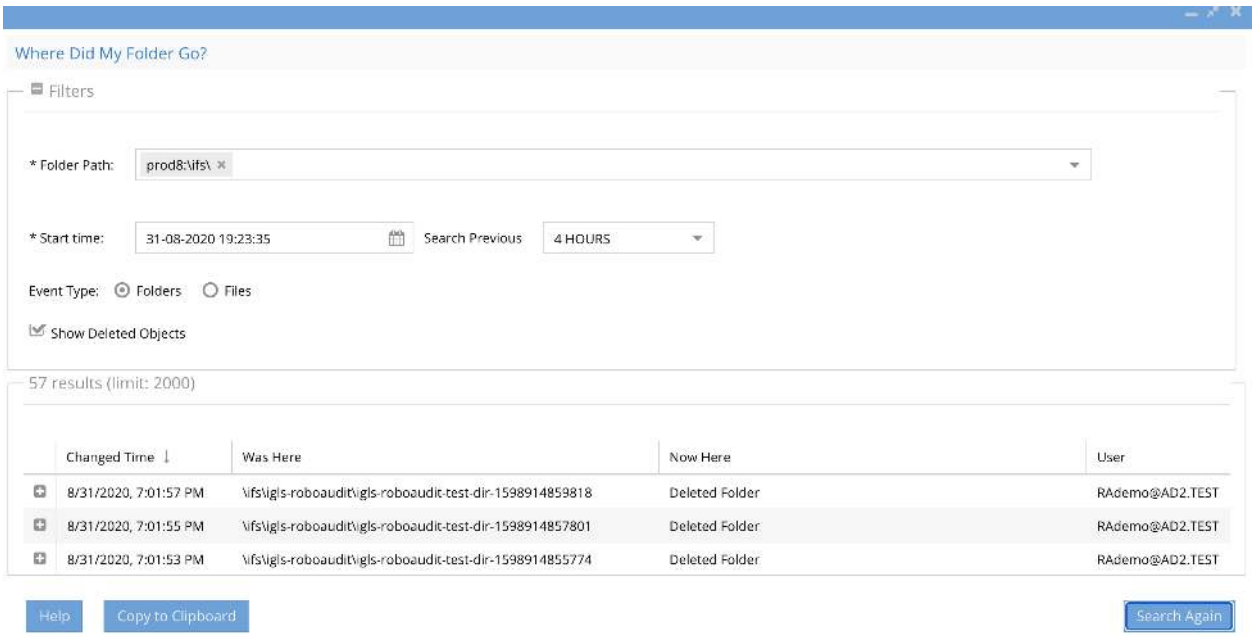

Use Case

1. A user drags and drops a folder on a share
2. A user directory rename
3. A file is deleted (2.5.6 after build 84))
4. A file is renamed (2.5.6 after build 84)
5. A user directory delete (2.5.5 >)
6. NOTE: Does not capture drag and drop between SMB shares, since this is a copy event
7. NOTE: Does not capture cut and paste of a direction on an SMB share or between SMB Shares since this operation triggers a create and delete operation. This type of action can found using queries and reports to look for file create and delete events.

How to Use Where did my folder go?

Overview

This feature helps locate drag and drop folders in the file system. It does not detect renamed files. New in 2.5.5, Directory deletes will be tracked in a fast cache lookup in addition to directory renames. A UI selector will allow filter results to show both Directory or file renames and deletes. A quick an easy copy to clipboard button allows pasting into Excel to sort the user or folder you are looking for more easily. New in 2.5.6 patch build is Deleted or renamed file support with more controls on the time range to search.

1. Use the directory selector to pick the cluster and the path (mandatory). Path must begin with \ifs\ or select a path location closer to the suspected path to reduce the results returned and locate the event faster.
2. Select the files or folders option to search one or the other. **NOTE: 2.5.6 builds after 84 build will store directory and file deletes and renames by default.**
    a. **NOTE: Searching for folders or files will return renames and deletes at the same time. The check box hides deleted folders or files and does not execute a new search. If the results do not show you the file or folder you should expand the time range or select a path closer to where the delete or rename occured.**

3. Select a day for the search (mandatory). NOTE: This feature requires search day by day or hour by hour range to ensure fast search results when their is a lot of data to search.
    a. In 2.5.6 build after 84 you can now select search previous 1, 2, 4, 6 or 24 hours to narrow down the amount of data returned from a search using the drop down selector. To specify the start time of your search you can use new time control to change the start time of the search. When selecting a previous day, the search will automatically start from the current time on your pc.
    b. 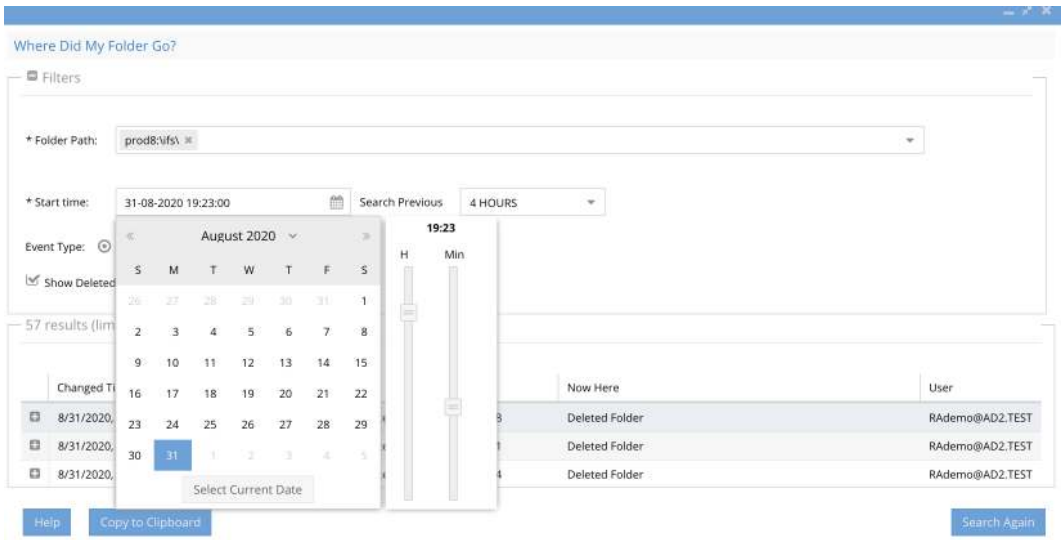

4. **Check the Show Deleted objects to filter to show Directory and file Deletes in addition to Renames. The default will display only rename events for files and folders.**
5. The results display
    a. The time of when the directory was moved
    b. The user id that executed the move or delete operation
    c. The original folder path **before** the move

d. The new location of the folder path **after** the move (this will be blank and show a delete for a directory delete)

e. The column Now Here will display Deleted folder or Deleted File or the path of a moved folder

f. See new results counter showing how many results were returned and the current limit default of 5000 results. In order to increase the number of results returned per search see the Eyeglass CLI guide here.

g. 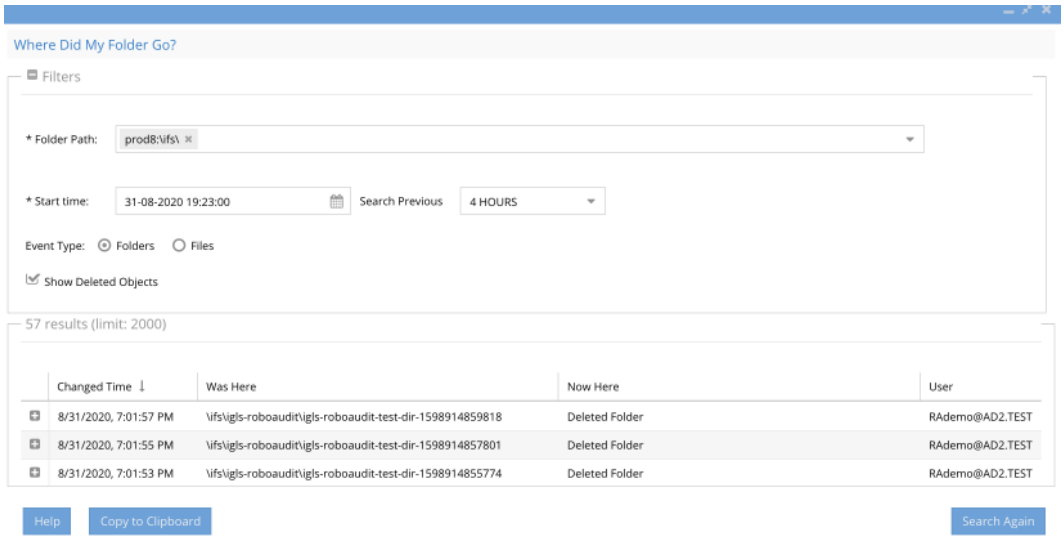

6. Using the search results identify the path and user that executed the folder move operation to move the data back to the previous location

7. **See screenshot below and the new option in > than 2.5.5 or later to copy results to the clipboard and paste into Excel.**

a. 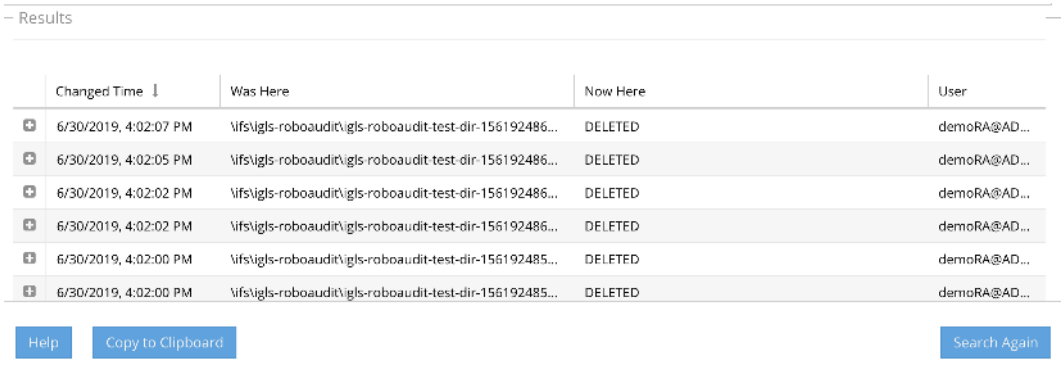

8. Done

RoboAudit

This feature performs continuous auditing by creating user events as an SMB connected user. The events are created , ingested and stored in the database. The Robot audit process runs reports and counts file and directory events and logs success or failure. This offers the highest level of confidence that audit data is being processed and stored. The audit lag is the time from when an event is created to when the data is searchable.

How Use RobotAudit

1. Open Easy Auditor Icon
2. Click the RoboAudit tab in the active auditor menu

a. 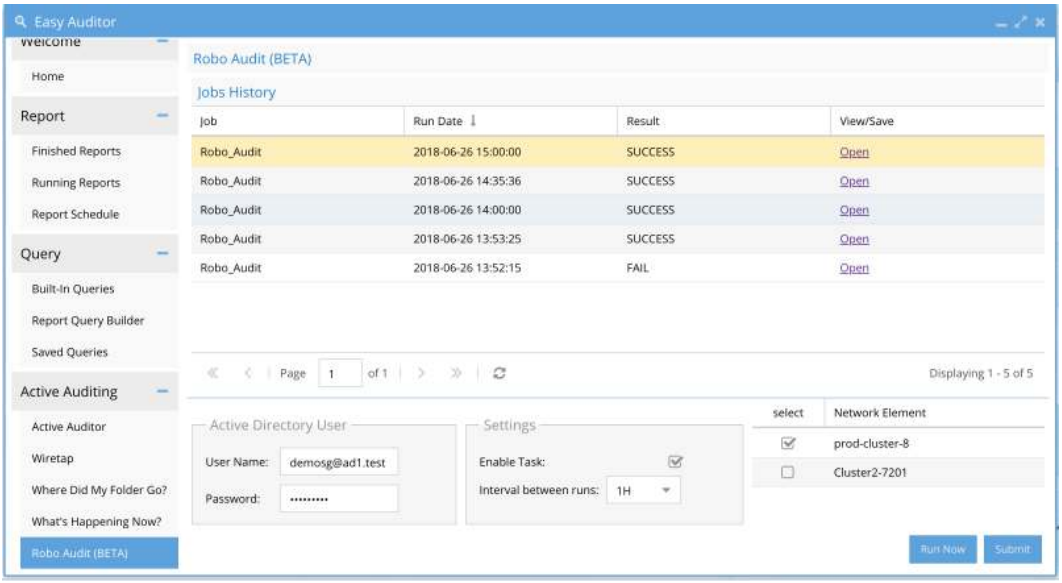

b.

c. **User Service Account Create an a local PowerScale user in the system zone local authentication provider OR c**reate an AD service account. This account should be a normal AD user and should not be the same as the Ransomware Defender Security guard user.

d. For local PowerScale user account enter user@clustername OR ror Active Directory user enter the user id user@domain or domain\user syntax with password. **NOTE: if running Ransomware defender user a different service account user to avoid a lockout on the RobotAudit service account user.**

e. Enable the task check box and select an interval default is 1D or once per day. The default should be used unless you need to debug something and can increase to once per hour.

f. Select the check box next to each cluster to test audit message flow.

g. click submit button to save settings. NOTE: The userid and password authentication will be tested at this point and an error returned if it fails.

h. Now click Run Now button to start the first execution.

How to Monitor RobotAudit jobs

1. Open the Jobs Icon
2. Click running jobs
3. Select the RobotAudit job and expand to see the steps that are being executed.
   a. 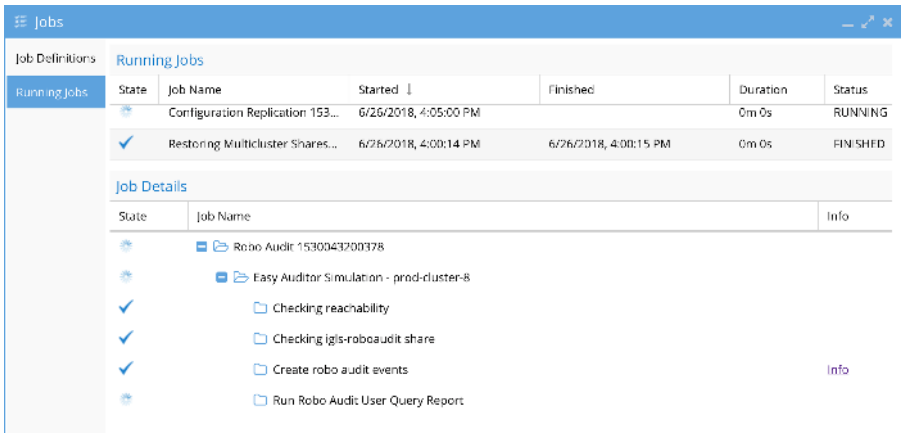

   b. done

How to verify RobotAudit test Results

1. Click on the RobotAudit tab in Easy Auditor Icon
2. Click open link to view the log file and all steps completed or any errors
   a. Example succesful log

b. Robo Audit Log Viewer

Download file

```
2018-06-26 15:00:00::061 INFO *************** Robo Audit Job STARTED ***************
2018-06-26 15:00:00::062 INFO Job Name: Robo Audit 1530039600061
2018-06-26 15:00:00::062 INFO User Name : demosg@ad1.test
2018-06-26 15:00:00::062 INFO Run Path Report: false
2018-06-26 15:00:00::062 INFO ******************************************************
2018-06-26 15:00:00::164 INFO prod-cluster-8 : Checking reachability - STARTED
2018-06-26 15:00:00::180 INFO prod-cluster-8 : is reachable
2018-06-26 15:00:00::180 INFO prod-cluster-8 : Step: "Checking reachability" Result: SUCCESS. Checking reachability -
FINISHED - Status: OK
2018-06-26 15:00:00::181 INFO prod-cluster-8 : Checking igls-roboaudit share - STARTED
2018-06-26 15:00:00::592 INFO prod-cluster-8 : igls-roboaudit already in place!
2018-06-26 15:00:00::593 INFO prod-cluster-8 : Step: "Checking igls-roboaudit share" Result: SUCCESS. Checking igls-
roboaudit share - FINISHED - Status: OK
2018-06-26 15:00:00::593 INFO prod-cluster-8 : Create robo audit events - STARTED
2018-06-26 15:00:32::811 INFO prod-cluster-8 : Created & Deleted - 10 directories, 10 files
2018-06-26 15:00:32::811 INFO prod-cluster-8 : Step: "Create robo audit events" Result: SUCCESS. Create robo audit
events - FINISHED - Status: OK
2018-06-26 15:00:32::812 INFO prod-cluster-8 : Run Robo Audit User Query Report - STARTED
2018-06-26 15:00:32::813 INFO Waiting for 300 seconds
2018-06-26 15:05:32::813 INFO prod-cluster-8 : Initiating Robo Audit USER Query Report
2018-06-26 15:06:40::399 INFO prod-cluster-8 : Run Robo Audit User Query Report 1530039932813 completed
successfully
2018-06-26 15:06:40::399 INFO prod-cluster-8 : Verifying robo audit USER query report
2018-06-26 15:06:40::407 INFO prod-cluster-8 : USER query report - EventType: DIR_CREATE RetrievedCount: 10
ExpectedCount: 10
2018-06-26 15:06:40::407 INFO prod-cluster-8 : USER query report - EventType: DIR_DELETE RetrievedCount: 10
ExpectedCount: 10
2018-06-26 15:06:40::407 INFO prod-cluster-8 : USER query report - EventType: FILE_CREATE RetrievedCount: 10
ExpectedCount: 10
2018-06-26 15:06:40::407 INFO prod-cluster-8 : USER query report - EventType: FILE_DELETE RetrievedCount: 10
ExpectedCount: 10
2018-06-26 15:06:40::407 INFO prod-cluster-8 : Step: "Run Robo Audit User Query Report" Result: SUCCESS. Run Robo
Audit User Query Report - FINISHED - Status: OK
2018-06-26 15:06:40::407 INFO ******** Job: Robo Audit 1530039600061 SUCCEEDED ********
```

c. Example failed log where the SMB share did not exist to create test files

Robo Audit Log Viewer

Download file

```
2018-06-26 13:52:15::094 INFO *************** Robo Audit Job STARTED ***************
2018-06-26 13:52:15::094 INFO Job Name: Robo Audit 1530035535092
2018-06-26 13:52:15::094 INFO User Name : demosg@ad1.test
2018-06-26 13:52:15::095 INFO Run Path Report: false
2018-06-26 13:52:15::095 INFO ******************************************************
2018-06-26 13:52:15::104 INFO prod-cluster-8 : Checking reachability - STARTED
2018-06-26 13:52:15::104 INFO prod-cluster-8 : is reachable
2018-06-26 13:52:15::104 INFO prod-cluster-8 : Step: "Checking reachability" Result: SUCCESS.
Checking reachability - FINISHED - Status: OK
2018-06-26 13:52:15::105 INFO prod-cluster-8 : Checking igls-roboaudit share - STARTED
2018-06-26 13:52:16::000 INFO prod-cluster-8 : igls-roboaudit created/updated!
2018-06-26 13:52:16::000 INFO prod-cluster-8 : Step: "Checking igls-roboaudit share" Result:
SUCCESS. Checking igls-roboaudit share - FINISHED - Status: OK
2018-06-26 13:52:16::000 INFO prod-cluster-8 : Create robo audit events - STARTED
2018-06-26 13:52:16::034 ERROR prod-cluster-8 : Create robo audit events - FAILED. Reason -
STATUS_BAD_NETWORK_NAME(3221225676/3221225676): Could not connect to
\\172.31.1.104\igls-roboaudit
2018-06-26 13:52:16::034 ERROR prod-cluster-8 : Step: "Create robo audit events" Result: FAILURE:
Create robo audit events - FINISHED - Status: ERROR
2018-06-26 13:52:16::035 ERROR ********* Job: Robo Audit 1530035535092 FAILED *********
```

d.

e. If RoboAudit indicates failure a case support should be opened with support.

Auditor Role Based Access

Easy auditor has 2 permissions the user roles icon that are automatically assigned to the auditor user. See below.

To create additional access to read only or read/write permissions consult the RBAC guide and use the permissions below.

RBAC guide

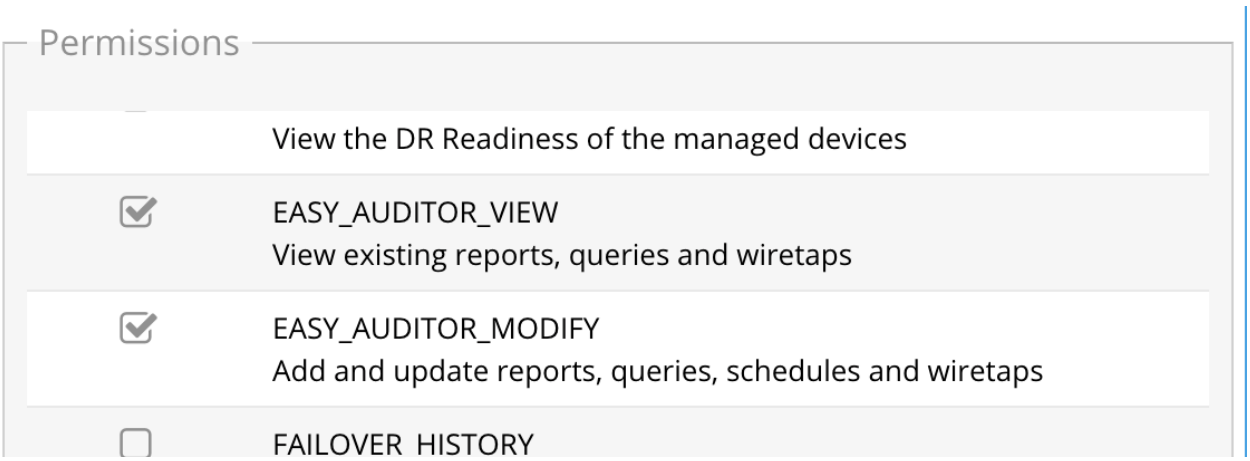

How to Centralize CSV saved reports on a NAS Export

Use this option to mount an NFS epxort on PowerScale to auto store reports centrally as a second copy.

1. Create a mount to an PowerScale and secure to all ECA node ip addresses with write access
2. Create a mount point on the Eyeglass appliance and create a mount point to the PowerScale export path.
3. eyeglasss-61:/opt/superna/sca/data # mkdir EA_reports

   eyeglasss-61:/opt/superna/sca/data # chown sca:users EA_reports

   eyeglasss-61:/opt/superna/sca/data # chmod 755 EA_reports

   eyeglasss-61:~ # igls admin eaCsvArchivePath set --value=/opt/superna/sca/data/EA_reports

   {

   "success": true

}

eyeglasss-61:~ # igls admin eaCsvArchivePath

/opt/superna/sca/data/EA_reports

4. In this example the mount point path is /opt/superna/sca/data/EA_reports
5. sudo -s (enter admin password)
6. type 'yast'
7. From the menu select Network Services --> then NFS client
8. Using Tab button select to add option
9. Complete inputs with FQDN to remote host, local path created in the step above
10. 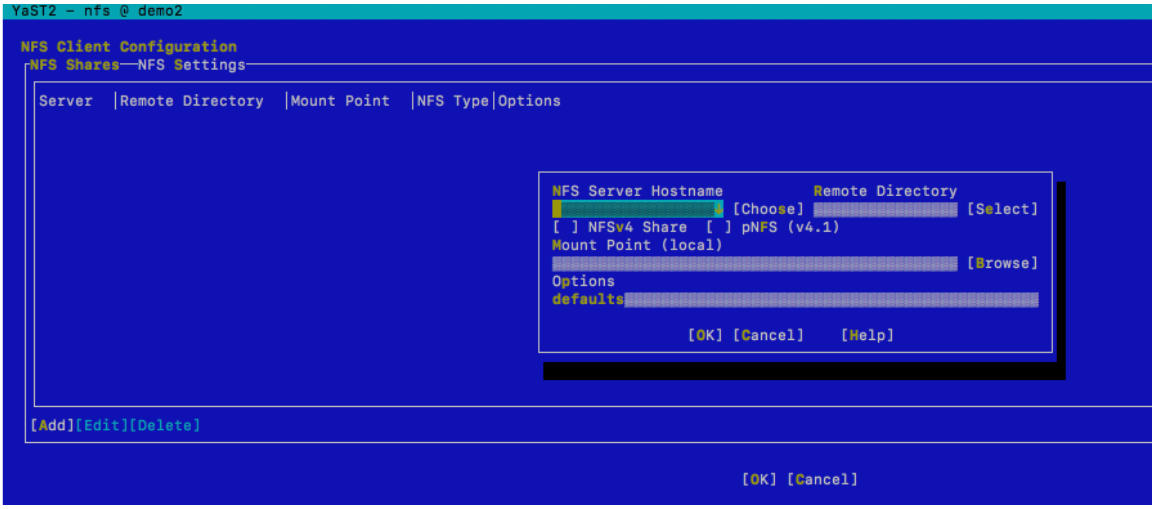

11. done.

How to view debug logs for Searches

The debug logs for each search executed are stored in HDFS and be viewed from the spark-history server webUI.

1. Login to eyeglass desktop
2. open managed services
3. expand node 1 and find spark history url to access the logs
4. 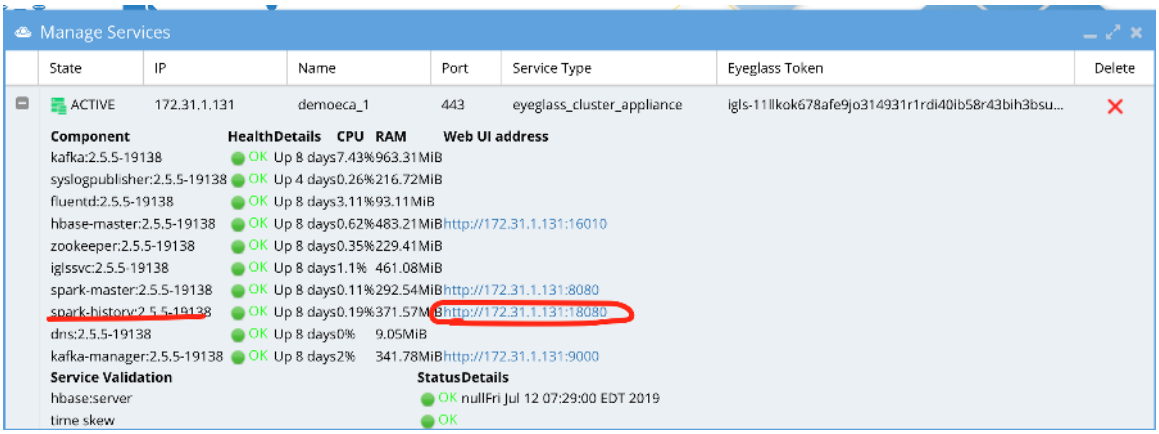

5. 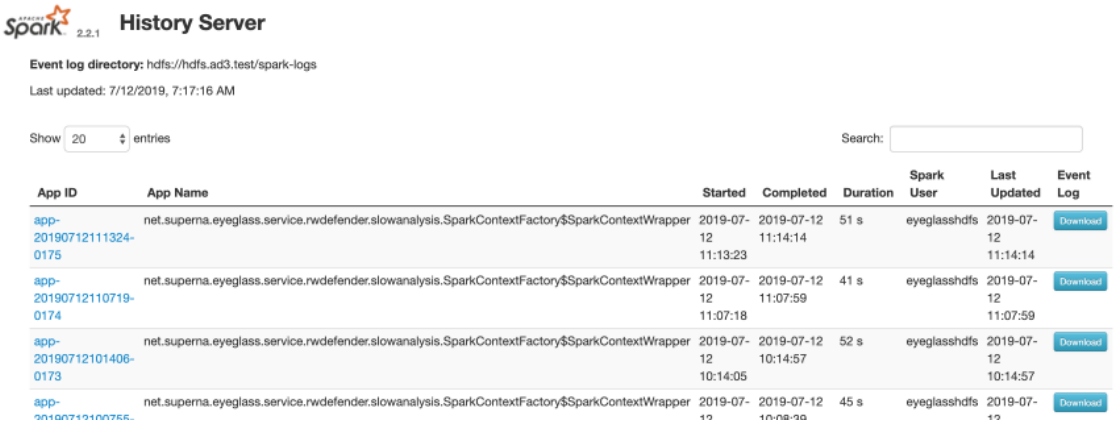

© Superna LLC

3.9. How to Use Excel for advanced filtering of CSV Reports

Home Top

How to Use Excel for advanced filtering of CSV Reports

This section covers advanced filtering options with Easy Auditor CSV exports using Microsoft Excel. Typical date and time filtering and combinations of columns are easily managed with Excel column heading filters. The following sections walk through how to apply date and time filters to filter event data more precisely.

on date and time can use this date format

1. Enable column heading row 1 data filter
2. **Time Filtering** enter a custom date and time column format as per below to allow advanced time based filtering. example **yyyy-mm-dd h:mm:ss AM/PM**

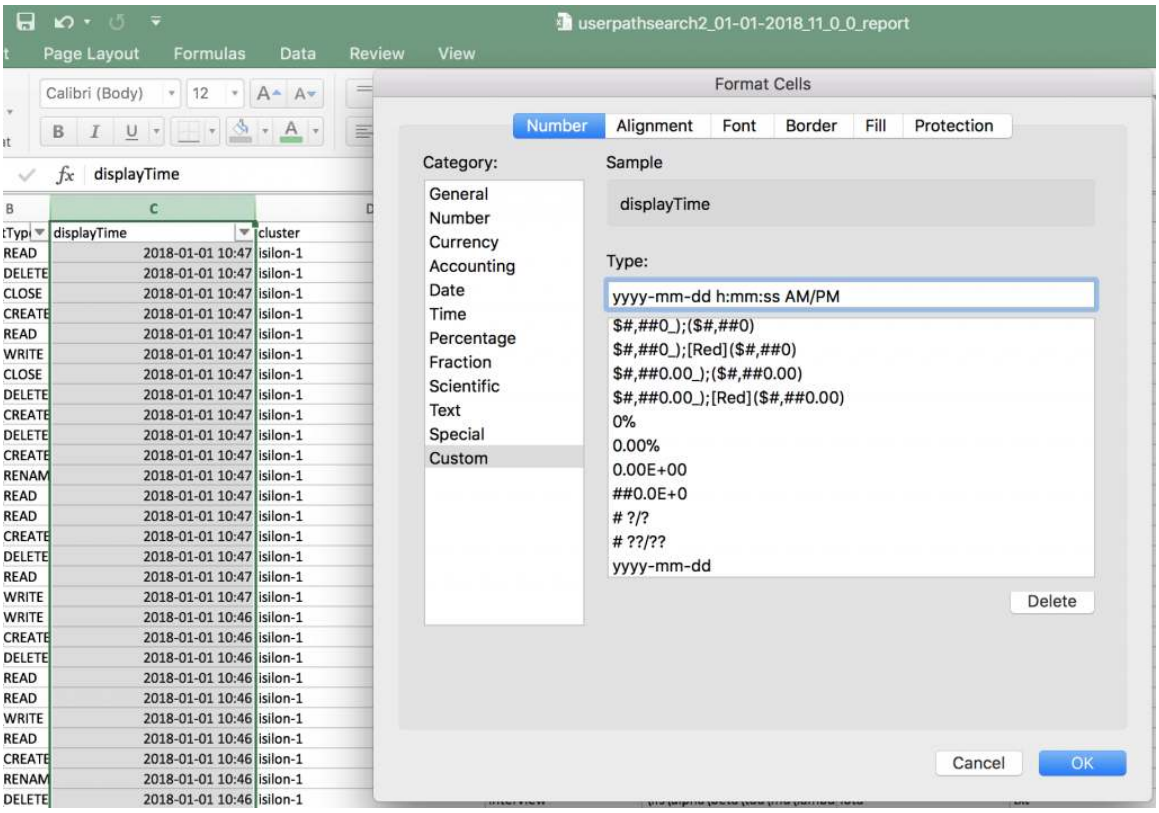

1. Advanced filters can be entered with seconds granularity using before , after or between filters. See example below.

2. 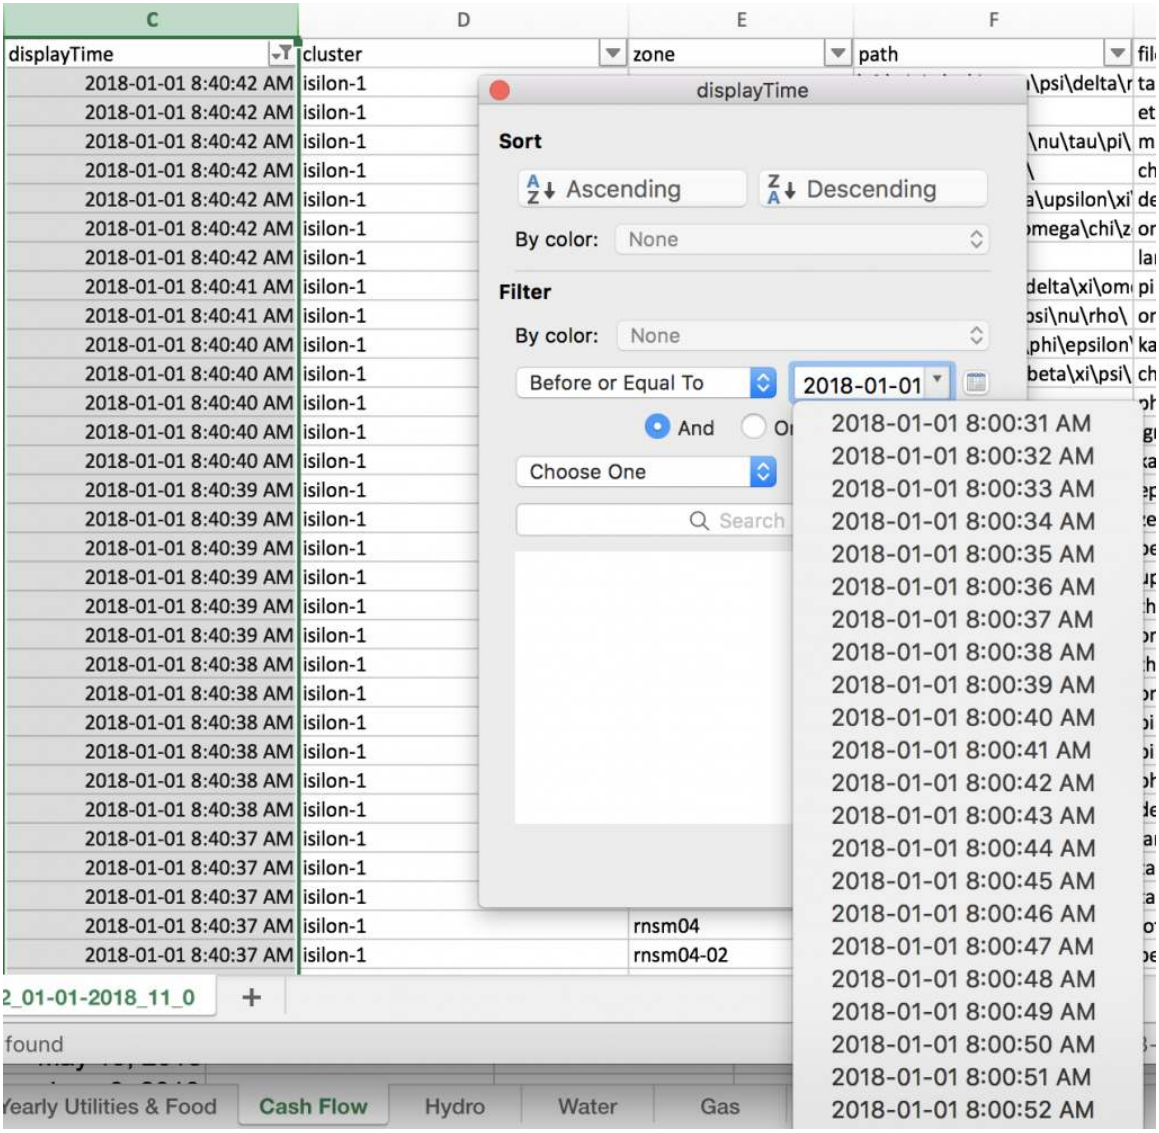

3. File Extension filtering

1. After enabling filtering on the column heading

2. The file extension column can be used to quickly find all files with an extension

3. 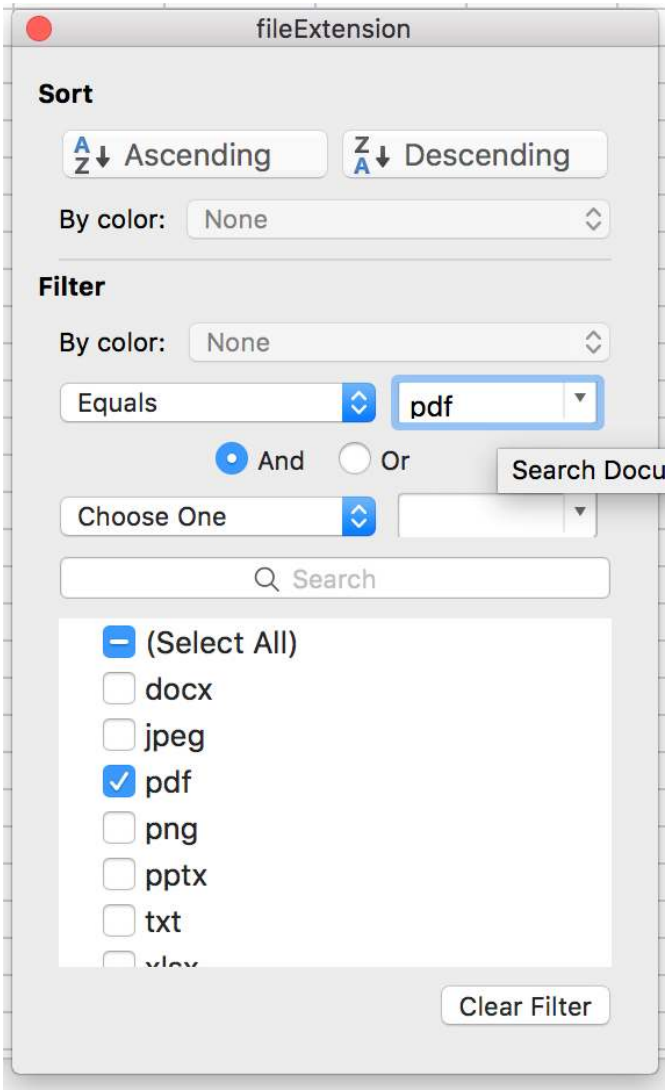

4. Path or partial path search
1. Entering contains in the path filter allows matching any partial match to a path

2. 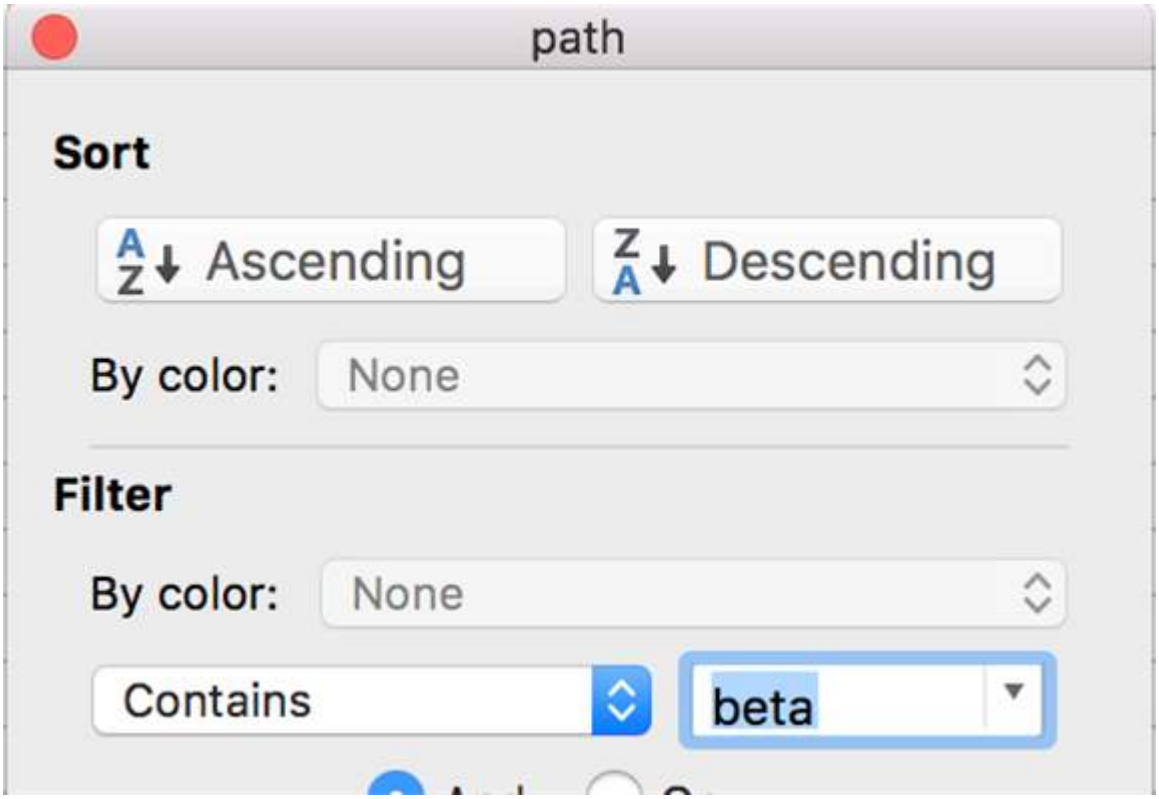

5. **Combined filters -** All column headings can be filtered to allow complex filter to find time range, extension, partial path match and filter on access zone.

1. 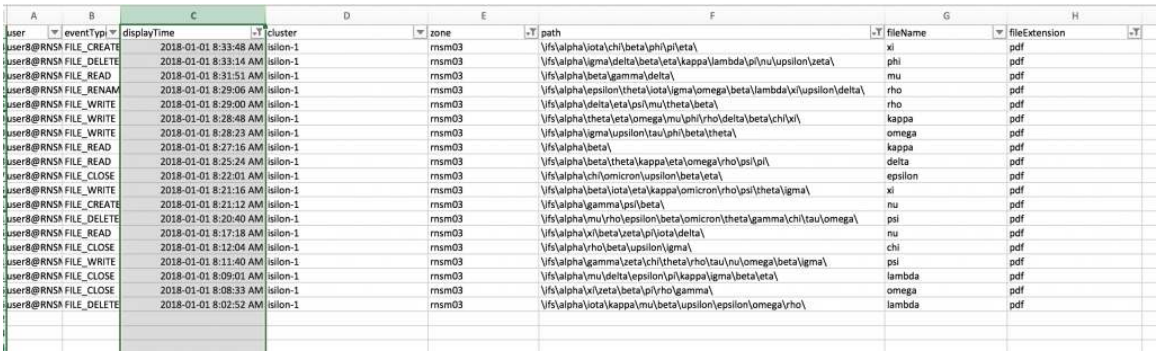

user	eventTypi	displayTime	cluster	zone	path	fileName	fileExtension
user8@RNSA	FILE_CREATE	2018-01-01 8:33:48 AM	isilon-1	rnsm03	\fs\alpha\iota\chi\beta\phi\pi\eta\	xi	pdf
user8@RNSA	FILE_DELETE	2018-01-01 8:33:14 AM	isilon-1	rnsm03	\fs\alpha\igma\delta\beta\eta\kappa\lambda\pi\nu\upsilon\zeta\	phi	pdf
user8@RNSA	FILE_READ	2018-01-01 8:31:51 AM	isilon-1	rnsm03	\fs\alpha\beta\gamma\delta\	mu	pdf
user8@RNSA	FILE_RENAM	2018-01-01 8:29:06 AM	isilon-1	rnsm03	\fs\alpha\epsilon\theta\iota\igma\omega\beta\lambda\xi\upsilon\delta\	rho	pdf
user8@RNSA	FILE_WRITE	2018-01-01 8:29:00 AM	isilon-1	rnsm03	\fs\alpha\delta\eta\psi\mu\theta\beta\	rho	pdf
user8@RNSA	FILE_WRITE	2018-01-01 8:28:48 AM	isilon-1	rnsm03	\fs\alpha\theta\eta\omega\mu\phi\rho\delta\beta\chi\xi\	kappa	pdf
user8@RNSA	FILE_WRITE	2018-01-01 8:28:23 AM	isilon-1	rnsm03	\fs\alpha\igma\upsilon\tau\phi\beta\theta\	omega	pdf
user8@RNSA	FILE_READ	2018-01-01 8:27:16 AM	isilon-1	rnsm03	\fs\alpha\beta\	kappa	pdf
user8@RNSA	FILE_READ	2018-01-01 8:25:24 AM	isilon-1	rnsm03	\fs\alpha\beta\theta\kappa\eta\omega\rho\psi\pi\	delta	pdf
user8@RNSA	FILE_CLOSE	2018-01-01 8:22:01 AM	isilon-1	rnsm03	\fs\alpha\chi\omicron\upsilon\beta\eta\	epsilon	pdf
user8@RNSA	FILE_WRITE	2018-01-01 8:21:16 AM	isilon-1	rnsm03	\fs\alpha\beta\iota\eta\kappa\omicron\rho\psi\theta\igma\	xi	pdf
user8@RNSA	FILE_CREATE	2018-01-01 8:21:12 AM	isilon-1	rnsm03	\fs\alpha\gamma\psi\beta\	nu	pdf
user8@RNSA	FILE_DELETE	2018-01-01 8:20:40 AM	isilon-1	rnsm03	\fs\alpha\mu\rho\epsilon\beta\omicron\theta\gamma\chi\tau\omega\	psi	pdf
user8@RNSA	FILE_READ	2018-01-01 8:17:18 AM	isilon-1	rnsm03	\fs\alpha\xi\beta\eta\pi\iota\delta\	nu	pdf
user8@RNSA	FILE_CLOSE	2018-01-01 8:12:04 AM	isilon-1	rnsm03	\fs\alpha\rho\beta\upsilon\igma\	chi	pdf
user8@RNSA	FILE_WRITE	2018-01-01 8:11:40 AM	isilon-1	rnsm03	\fs\alpha\gamma\zeta\chi\theta\rho\tau\nu\omega\beta\igma\	psi	pdf
user8@RNSA	FILE_CLOSE	2018-01-01 8:09:01 AM	isilon-1	rnsm03	\fs\alpha\mu\delta\epsilon\pi\kappa\igma\beta\eta\	lambda	pdf
user8@RNSA	FILE_CLOSE	2018-01-01 8:08:33 AM	isilon-1	rnsm03	\fs\alpha\xi\zeta\beta\pi\rho\gamma\	omega	pdf
user8@RNSA	FILE_DELETE	2018-01-01 8:02:52 AM	isilon-1	rnsm03	\fs\alpha\iota\kappa\mu\beta\upsilon\epsilon\omega\rho\	lambda	pdf

© Superna LLC

3.10. How to Backup and Restore an Audit Database

Home Top

How to Backup and Restore an Audit Database

A key advantage to Easy Auditor architecture is using PowerScale native features to protect the audit data. The following sections explains how to backup and restore the Analytics database.

Backup the Audit Database with SnapshotIQ

1. Create a scheduled snapshot of the HDFS root directory that contains the Audit Database directory with PowerScale SnapshotIQ . Example:
2. Recommended schedule daily snapshot at noon 7 days a week, with 30 day retention

Access zone basepath for audit database is **/ifs/data/igls/analyticsdb**

HDFS root directory: **/ifs/data/igls/analyticsdb/eca**

Audit Database directory: **/ifs/data/igls/analyticsdb/eca**

3. If creating a manual snapshot by using the PowerScale GUI, do not leave the snapshot name blank.
1. A default snapshot name will be applied automatically (e.g. “Snapshot: 2017Nov09, 10:59 PM"). That name format is not supported for ECA cluster due to special character support with HDFS.

2. That name format will prevent the ECA cluster to be brought up. Provide a normal name for the snapshot. Avoid to use name with the ":" character.
4. If creating a scheduled snapshot, also avoid to use the name with the ":" character (e.g. ScheduleName_Duration_%Y-%m-%d_%H:%M). That name format is not supported for ECA cluster. That name format will prevent the ECA cluster to be brought up. Provide a name pattern without ":" character for the snapshot.

Note: Please refer to PowerScale documentation for creating a snapshot, including to create a SnapRevert domain

Restore the Audit Database with SnapshotIQ

1. ssh to ECA master node (node 1). Login as ecaadmin
2. Run command: **ecactl cluster down**.
3. Wait until nodes are down
4. PowerScale command:
    a. **NOTE: You need to run snaprevert domain mark job first if not already done. See screenshot.**

b. 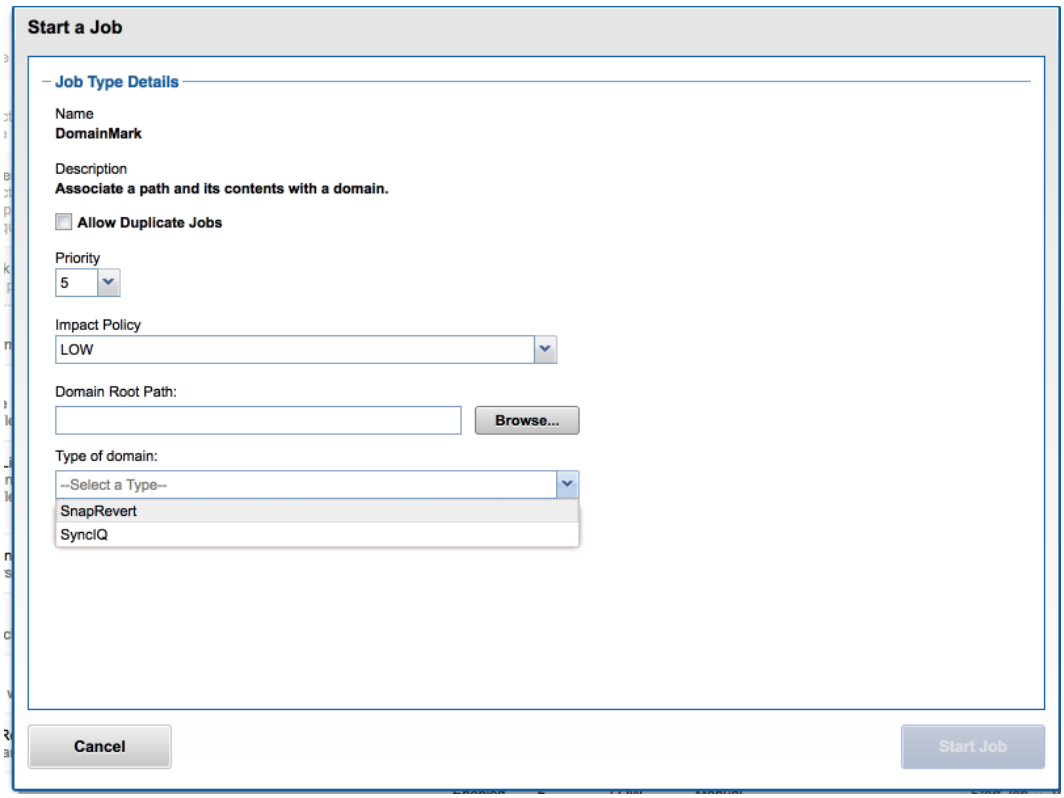

c. **isi job jobs start snaprevert --snapid** ***xxxx*** (verify the correct snapid of the snapshot to revert to)

5. To verify the snapshot revert job status, PowerScale command to list running jobs: **isi job jobs list**
6. Once the snapshot revert job has completed
7. After ECA Cluster VMs are up, then bring up ECA Cluster
8. ssh to ECA master node (node 1)
9. Login as ecaadmin
10. Run command: **ecactl cluster up**.
11. NOTE: During cluster up uncommitted transactions are replayed to the database, this can be seen from the HBASE Region server GUI logs http://x.x.x.x:16030 this can take longer to startup the cluster
12. Sample below

13. Verify that ECA Cluster is up and audit database status return no error. Command: **ecactl db shell**
14. Until the status appears like above, HBASE is not fully operational.
15. Done

2017-11-10 08:59:52,729 WARN [main] util.NativeCodeLoader: Unable to load native-hadoop library for your platform... using builtin-java classes where applicable

HBase Shell; enter 'help<RETURN>' for list of supported commands.

Type "exit<RETURN>" to leave the HBase Shell

Version 1.2.6, rUnknown, Mon May 29 02:25:32 CDT 2017

hbase(main):001:0> status

1 active master, 2 backup masters, 3 servers, 0 dead, 2.6667 average load

© Superna LLC

3.11. Backup and DR for Audit Database with SyncIQ to a Remote Cluster

Home Top

- Overview
- How to Use SyncIQ to replicate and protect the Audit Database to a remote PowerScale Cluster
- How to Use the replicated copy of the Audit Database and Failover to the Warm Standby ECA Cluster

Overview

This solution backups the Audit Database to a remote cluster to provide a remote backup and DR copy at the same time.

How to Use SyncIQ to replicate and protect the Audit Database to a remote PowerScale Cluster

1. **Prerequisites**
    a. Follow the ECA failover guide prerequisites and prepare the DR cluster before following this guide. Ensure all steps are completed in the prerequisites section of this guide.
2. Create a SyncIQ Policy on the production cluster to replicate the audit database to a directory under the HDFS root directory with a replication schedule to the target DR cluster. Example:
    a. HDFS root directory: **/ifs/data/igls/analyticsdb/**

b. SyncIQ Policy Source Path: **/ifs/data/igls/analyticsdb/eca1/**

c. SyncIQ Policy Target Path on the remote cluster: **/ifs/data/igls/analyticsdb/eca1/**

d. Recommended policy Schedule: once a day at noon, 7 days a week

e. Complete the policy configuration name, description and target host property of the policy.

f. Run the policy after it is created to copy the database

g. Verify the policy completes successfully.

3. Done

How to Use the replicated copy of the Audit Database and Failover to the Warm Standby ECA Cluster

This procedure assumes an ECA cluster will be deployed at the remote site to use the database copy and monitor the DR cluster after failover. See the ECA disaster recovery guide that covers how to configure a Warm standby ECA. All of the DR target cluster prerequisites and Scenario #2 Warm Standby steps are assumed to be completed. This section covers the steps to use the replicated copy of the database using the Warm Standby ECA cluster.

Procedure to Mount the Audit Database with the Warm Standby ECA:

1. Use Eyeglass SyncIQ policy failover option with DR Assistant to failover the Audit Database policy. This will automate all steps for the SyncIQ policy and configure reverse replication. See DR assistant guide here.

a. After the failover is successful, the DR cluster copy of the audit database will be writeable and reverse replication to re-protect the audit database.

2. Bring up the Warm standby ECA Cluster at the DR site

   a. ssh to ECA master node (node 1)

   b. Login as ecaadmin

   c. Run command:

      i. **ecactl cluster up**

      ii. NOTE: During cluster up uncommitted transactions are replayed to the database, this can be seen from the HBASE Region server GUI logs http://x.x.x.x:16030 (x.x.x.x is node 1 of the ECA). This will delay the cluster up process while the database replays transactions.

      iii. 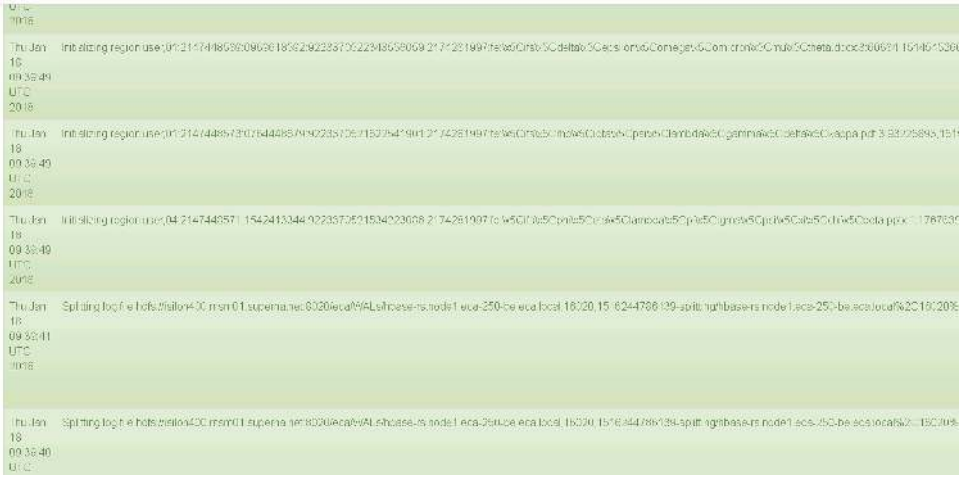

   d. Verify that ECA Cluster is up and audit database status returns no error.

      i. Command: **ecactl db shell**

      ii. **Type" status [enter]**

iii.

iv. **no error messages should be returned.**

e.

© Superna LLC

3.12. How to check Analytics database size

How to check Analytics database size

1. Login to the cluster that stores the audit database using ssh
2. Change to the root of the audit database access zone (default is /ifs/analyticsdb)
3. du -h
4. This will sum the files in the database
5. **Best Practise:** Create an advisory quota on /ifs/analyticsdb to monitor space usage.

© Superna LLC

3.13. How to Use Easy Auditor for Typical Audit Use Cases

How to Use Easy Auditor for Typical Audit Use Cases

This section walks through typical audit use cases and assists with suggested features to address the audit requirements.

- Urgent Request to re-act to a security event
- Application Performance Issue for NAS share or export
- Compliance Reporting
- User Behavior Audit
- Triggers for Network Aware Monitoring
  - example 1
  - example 2
  - example 3
  - example 4

Urgent Request to re-act to a security event

This type of request has urgency and can be a data leak, user termination or information delete request that needs answers fast. This requires instanst access to navigate the file system where activity needs to be reviewed in muliple folders since its not 100% clear what you are looking for. The new feature in Easy Auditor is targetted at this use case.

Options to audit this use case

1. **WireTap? - This feature allows browsing the file system and instantly see all file activity by user and path. This UI allows advanced filtering of specific events, IO by a single user or even a file name. This tool can assist with live security incident in the file system since this tool is viewing live audit data based on the filters configured. This speeds up the investigation work. It also avoids searching the database since all data is streamed to the GUI .**
    a. **Only possible with real-time event processing platform like Easy Auditor**

User Reports of missing files in a share path

Options to audit this use case

1. **Where did my folder go? Browse to the path /ifs and search for all directory rename (move) or deletes with a single click. See if delete or rename events are the root cause of the issue. Simple copy to clipboard to sort in Excel if you get a lot of results. Purpose built index for this common every day issue.**
2. **Scheduled Query**: Create a search with advanced search tab and enter the cluster and path in the file system to monitor. Save the query and then use the schedule tab to run every hour to alert you on any deletes in that path
3. Same as above but enter a file extension as well to narrow the delete query and schedule every hour
2. **WireTap:** Create a wiretap session to monitor the path in real-time if the delete is a recurring issue on a path. The wiretap can monitor a path if it unknown who deleted the file(s). If its a specific user issue, wiretap the user to monitor user activity while they execute a sequence to reproduce the delete issue.

Application Performance Issue for NAS share or export

Users raise issue about performance of an application or data access. This can be caused by file locking or temp file creation on the NAS share versus local disk or poor application workflow accessing network shares/exports.

Options to audit this use case

1. **Wiretap:** Create a wiretap session for the user or path with performance issue. Monitor while asking end users to re-attempt the application operations. Path based wiretap is best when multiple users raise performance issue on a share. Create use based wiretap when an application performance issue for single users.

Compliance Reporting

This report uses the logon and log off audits to report on user acccess to storage. It can also report on failed logon attempts. HIPAA, PCI and many other industry regulations require an inscope device must be able to report on user authenticated access attempts. This report will meet this requirement. It is also required to know who has access to data for tracability.

Options to audit this use case

1. Builtin Reports
    a. Login Monitor report - shows which users logged in and logged off the array, including failed login attempts
    b. **Stale access Report** - shows which users accessed data they have permissions to see over a time period. Users that do not access data can be considered for removal.

Security best practice states least access model should be followed.

c. Access Report - Shows a list of users that can mount SMB shares. The user groups are expanded on the SMB level security to build a full uaer list that can access a share. Used to send to departements to verify data access privileges.

Excessive Permissions Analysis

The excessive permissions report assists with identifying users with access to data that is no longer being accessed. This report can help with compliance and securing access to data. The report analysis users that have accessed shares and resolves their share access from AD group membership and lists users with access to shares but no actual file activity within the report range.

This list of users are candidates to have group membership reduced to narrow access to data.

Options to audit this use case

1. **Builtin Excessive Permissions Report:** Open the Report History to open the report.

User Behavior Audit

Random user audits or suspicious file access auditing is a common requirement in security departments. Easy Auditor provides several tools to perform proactive audits of file access.

Options to audit this use case

1. **Wiretap:** Create a wiretap session with per user option. The session can be actively monitored or saved and run a report to build a report of all file access since the creation of the wiretap session.
2. **Search:** Build a search based on user id and a date range , that will return all file access on all shares within the data range. In the preview screen of the search select run report.

Triggers for Network Aware Monitoring

1. Active Auditing and triggers allow pre-built logic and custom rules to triger on any type of proactive monitoring needed to secure data.
    a. **Data loss prevention** trigger can monitor users doing a bulk copy of sensitive data. Recommened on a subset of your data example financial or other sensitive data. Configure the trigger to monitor the % of data on the path that is normal usage. example 5-10%. Experiment with a % to avoid false positive scenarios
    b. **Mass Delete protection trigger** - Use this pre built trigger to monitor paths for high rate deletes by users to get visibility to deletes and user behaviours. Monitor triggers to verify if users have odd work flows that are impacting the cluster with high rate deletes and copies.
    c. **Custom triggers** - This option is very powerful option to create rules with and or logic using audit data fields and thresholds and time windows.
        i. example 1

1. monitor a path of data based on the source ip subnet of hosts touching the data. Use this option when application servers on a subnet are the only authorized machines to access the data and indentify any user trying to access the data directly with subnet aware triggers

ii. example 2

1. identify banned file types with a simple extension based trigger example mp3 extension filter can locate users touching, creating, or accessing banned file types

iii. example 3

1. monitor a users behavior with a user based trigger based on deletes. If a user is suspected of deleting data in a group share, configure a user monitor trigger for deletes and get notifications any time the user deletes data

iv. example 4

1. Monitor all access to centralized storage from VPN or wifi using subnet aware policy to identify who and what is being access from external subnets.

© Superna LLC

3.14. Audit Message Workflows

Audit Message Workflows

This section shows expected audit messages for typical file action work flows to assist with auditing applications and user file access.

The Turbo audit workflows cover tested file actions with Turbo Audit enabled. This is the default configuration.

- Audit Message Workflows with Turbo Audit - SMB
  - SMB (Turbo Audit): Create a File
  - SMB (Turbo Audit): Rename a File
  - SMB (Turbo Audit): Write to a File
  - SMB (Turbo Audit): Delete a File
  - SMB (Turbo Audit): Create a Folder
  - SMB (Turbo Audit): Delete a Folder
  - SMB (Turbo Audit): Rename a Folder
  - SMB (Turbo Audit): Set ACL of a file
  - SMB (Turbo Audit): Set ACL of a Directory
  - SMB (Turbo Audit): User with Read-Only ACL open a File
- Audit Message Workflows with Turbo Audit - NFS
  - NFS (Turbo Audit): Create a File
  - NFS (Turbo Audit): Rename a File
  - NFS (Turbo Audit): Write to a File
  - NFS (Turbo Audit): Delete a File

- NFS (Turbo Audit): Create a Folder
- NFS (Turbo Audit): Delete a Folder
- NFS (Turbo Audit): Rename a Folder
- NFS (Turbo Audit): Modify ACL of a file
- NFS (Turbo Audit): Modify ACL of a Directory

Audit Message Workflows with Turbo Audit - SMB

WorkFlow Description	File Audit messages Expected
SMB (Turbo Audit): Create a File	S_FILE_CREATE ..\parent_dir\desktop.ini FILE_CLOSE_MODIFIED ..\parent_dir\dir_name\file_name FILE_WRITE ..\parent_dir\dir_name\file_name DIR_CLOSE ..\parent_dir\ DIR_OPEN ..\parent_dir\ FILE_CREATE ..\parent_dir\dir_name\file_name
SMB (Turbo Audit): Rename a File	S_FILE_CREATE ..\parent_dir\desktop.ini DIR_CLOSE ..\parent_dir\ FILE_CLOSE ..\parent_dir\dir_name\file_name FILE_RENAME ..\parent_dir\dir_name\file_name DIR_OPEN ..\parent_dir\ FILE_CLOSE ..\parent_dir\dir_name\file_name
SMB (Turbo	FILE_CLOSE_MODIFIED ..\parent_dir\dir_name\file_name

Audit): Write to a File	FILE_WRITE ..\parent_dir\dir_name\file_name DIR_CLOSE ..\parent_dir\ DIR_OPEN ..\parent_dir\ FILE_READ ..\parent_dir\dir_name\file_name FILE_OPEN_WRITE ..\parent_dir\dir_name\file_name FILE_CLOSE ..\parent_dir\dir_name\file_name DIR_CLOSE ..\parent_dir\ DIR_OPEN ..\parent_dir\ FILE_READ ..\parent_dir\dir_name\file_name FILE_OPEN_WRITE ..\parent_dir\dir_name\file_name
SMB (Turbo Audit): Delete a File	S_FILE_CREATE ..\parent_dir\desktop.ini FILE_CLOSE ..\parent_dir\dir_name\file_name FILE_DELETE ..\parent_dir\dir_name\file_name DIR_CLOSE ..\parent_dir\ DIR_OPEN ..\parent_dir\
SMB (Turbo Audit): Create a Folder	S_FILE_CREATE ..\parent_dir\desktop.ini DIR_CLOSE ..\parent_dir\dir_name\new_dir_name DIR_CREATE ..\parent_dir\dir_name\new_dir_name
SMB (Turbo Audit): Delete a	S_FILE_CREATE ..\parent_dir\desktop.ini DIR_CLOSE ..\parent_dir\dir_name\current_dir_name

Folder	DIR_DELETE ..\parent_dir\dir_name\current_dir_name DIR_OPEN ..\parent_dir\dir_name\current_dir_name
SMB (Turbo Audit): Rename a Folder	S_FILE_CREATE ..\parent_dir\desktop.ini DIR_CLOSE ..\parent_dir\dir_name\current_dir_name DIR_RENAME ..\parent_dir\dir_name\current_dir_name DIR_OPEN ..\parent_dir\dir_name\current_dir_name DIR_CLOSE ..\parent_dir\dir_name DIR_OPEN ..\parent_dir\dir_name DIR_CLOSE ..\parent_dir\dir_name\current_dir_name DIR_OPEN ..\parent_dir\dir_name\current_dir_name
SMB (Turbo Audit): Set ACL of a file	FILE_CLOSE ..\parent_dir\dir_name\file_name FILE_SET_ACL ..\parent_dir\dir_name\file_name DIR_CLOSE ..\parent_dir\dir_name DIR_CLOSE ..\parent_dir\dir_name DIR_OPEN ..\parent_dir\dir_name DIR_CLOSE ..\parent_dir\dir_name DIR_OPEN ..\parent_dir\dir_name DIR_CLOSE ..\parent_dir\dir_name DIR_OPEN ..\parent_dir\dir_name DIR_CLOSE ..\parent_dir\dir_name DIR_OPEN ..\parent_dir\dir_name DIR_OPEN ..\parent_dir\dir_name DIR_CLOSE ..\parent_dir\dir_name

SMB (Turbo Audit): Set ACL of a Directory	DIR_CLOSE ..\parent_dir\dir_name\current_dir_name DIR_OPEN ..\parent_dir\dir_name\current_dir_name DIR_CLOSE ..\parent_dir\dir_name\current_dir_name DIR_OPEN ..\parent_dir\dir_name\current_dir_name DIR_CLOSE ..\parent_dir\dir_name\current_dir_name DIR_CLOSE ..\parent_dir\dir_name\current_dir_name DIR_OPEN ..\parent_dir\dir_name\current_dir_name DIR_SET_ACL ..\parent_dir\dir_name\current_dir_name DIR_OPEN ..\parent_dir\dir_name\current_dir_name DIR_CLOSE ..\parent_dir\dir_name\current_dir_name DIR_CLOSE ..\parent_dir\dir_name DIR_OPEN ..\parent_dir\dir_name DIR_OPEN ..\parent_dir\dir_name\current_dir_name DIR_CLOSE ..\parent_dir\dir_name\current_dir_name DIR_OPEN ..\parent_dir\dir_name\current_dir_name DIR_CLOSE ..\parent_dir\dir_name DIR_OPEN ..\parent_dir\dir_name DIR_CLOSE ..\parent_dir\dir_name\current_dir_name DIR_OPEN

	..\parent_dir\dir_name\current_dir_name DIR_CLOSE ..\parent_dir\dir_name\current_dir_name DIR_OPEN ..\parent_dir\dir_name\current_dir_name DIR_CLOSE ..\parent_dir\dir_name DIR_OPEN ..\parent_dir\dir_name S_FILE_CREATE ..\parent_dir\desktop.ini DIR_CLOSE ..\parent_dir\dir_name\current_dir_name DIR_CREATE ..\parent_dir\dir_name\current_dir_name
SMB (Turbo Audit): User with Read-Only ACL open a File	FILE_CLOSE ..\parent_dir\dir_name\file_name DIR_CLOSE ..\parent_dir\dir_name DIR_OPEN ..\parent_dir\dir_name FILE_CLOSE ..\parent_dir\dir_name\file_name FILE_OPEN_NOACCESS ..\parent_dir\dir_name\file_name FILE_CLOSE ..\parent_dir\dir_name\file_name FILE_OPEN_NOACCESS ..\parent_dir\dir_name\file_name FILE_CLOSE ..\parent_dir\dir_name\file_name FILE_OPEN_NOACCESS ..\parent_dir\dir_name\file_name DIR_OPEN ..\parent_dir\dir_name DIR_CLOSE ..\parent_dir\dir_name DIR_CLOSE ..\parent_dir\dir_name DIR_OPEN ..\parent_dir\dir_name DIR_CLOSE ..\parent_dir\dir_name

	FILE_OPEN_NOACCESS ..\parent_dir\dir_name\file_name FILE_CLOSE ..\parent_dir\dir_name\file_name DIR_OPEN ..\parent_dir\dir_name FILE_CLOSE ..\parent_dir\dir_name\file_name FILE_OPEN_NOACCESS ..\parent_dir\dir_name\file_name FILE_CLOSE ..\parent_dir\dir_name\file_name FILE_OPEN_NOACCESS ..\parent_dir\dir_name\file_name FILE_OPEN_READ ..\parent_dir\dir_name\file_name DIR_CLOSE ..\parent_dir\dir_name DIR_OPEN ..\parent_dir\dir_name S_FILE_CREATE ..\parent_dir\dir_name\desktop.ini

Audit Message Workflows with Turbo Audit - NFS

WorkFlow Description	File Audit messages Expected
NFS (Turbo Audit): Create a File	DIR_CLOSE ..\parent_dir\dir_name FILE_CLOSE_MODIFIED ..\parent_dir\dir_name\file_name FILE_WRITE ..\parent_dir\dir_name\file_name FILE_CLOSE ..\parent_dir\dir_name\file_name

	FILE_CREATE ..\parent_dir\dir_name\file_name DIR_OPEN ..\parent_dir\dir_name DIR_CLOSE ..\parent_dir\dir_name DIR_OPEN ..\parent_dir\dir_name DIR_CLOSE ..\parent_dir\dir_name DIR_OPEN ..\parent_dir\dir_name DIR_CLOSE ..\parent_dir\dir_name DIR_OPEN ..\parent_dir\dir_name
NFS (Turbo Audit): Rename a File	DIR_CLOSE ..\parent_dir\dir_name DIR_CLOSE ..\parent_dir\dir_name FILE_RENAME ..\parent_dir\dir_name\file_name DIR_OPEN ..\parent_dir\dir_name DIR_OPEN ..\parent_dir\dir_name DIR_CLOSE ..\parent_dir\dir_name DIR_CLOSE ..\parent_dir\dir_name DIR_OPEN ..\parent_dir\dir_name DIR_OPEN ..\parent_dir\dir_name
NFS (Turbo Audit): Write to a File	DIR_CLOSE ..\parent_dir\dir_name FILE_DELETE ..\parent_dir\dir_name\file_name.swp FILE_DELETE ..\parent_dir\dir_name\file_name~ FILE_CLOSE ..\parent_dir\dir_name\file_name FILE_CLOSE_MODIFIED ..\parent_dir\dir_name\file_name FILE_SET_ACL ..\parent_dir\dir_name\file_name FILE_WRITE

	..\parent_dir\dir_name\file_name FILE_CLOSE ..\parent_dir\dir_name\file_name FILE_CREATE ..\parent_dir\dir_name\file_name DIR_CLOSE ..\parent_dir\dir_name FILE_RENAME ..\parent_dir\dir_name\file_name DIR_OPEN ..\parent_dir\dir_name FILE_DELETE ..\parent_dir\dir_name\tempfile FILE_SET_ACL ..\parent_dir\dir_name\tempfile FILE_CLOSE ..\parent_dir\dir_name\tempfile FILE_SET_ACL ..\parent_dir\dir_name\tempfile FILE_CLOSE ..\parent_dir\dir_name\tempfile FILE_CLOSE ..\parent_dir\dir_name\tempfile FILE_CREATE ..\parent_dir\dir_name\tempfile FILE_CLOSE_MODIFIED ..\parent_dir\dir_name\file_name.swp FILE_CLOSE ..\parent_dir\dir_name\file_name.swp FILE_SET_ACL ..\parent_dir\dir_name\file_name.swp FILE_WRITE ..\parent_dir\dir_name\file_name.swp FILE_CLOSE ..\parent_dir\dir_name\file_name.swp FILE_SET_ACL ..\parent_dir\dir_name\file_name.swp FILE_CLOSE ..\parent_dir\dir_name\file_name.swp FILE_CREATE

	..\parent_dir\dir_name\file_name.swp FILE_DELETE ..\parent_dir\dir_name\file_name.swp FILE_DELETE ..\parent_dir\dir_name\file_name.swx FILE_CLOSE ..\parent_dir\dir_name\file_name.swx FILE_SET_ACL ..\parent_dir\dir_name\file_name.swx FILE_CLOSE ..\parent_dir\dir_name\file_name.swx FILE_CREATE ..\parent_dir\dir_name\file_name.swx FILE_CLOSE ..\parent_dir\dir_name\file_name.swp FILE_SET_ACL ..\parent_dir\dir_name\file_name.swp FILE_CLOSE ..\parent_dir\dir_name\file_name.swp FILE_CREATE ..\parent_dir\dir_name\file_name.swp DIR_CLOSE ..\parent_dir\dir_name DIR_OPEN ..\parent_dir\dir_name DIR_OPEN ..\parent_dir\dir_name FILE_CLOSE ..\parent_dir\dir_name\file_name FILE_CLOSE ..\parent_dir\dir_name\file_name DIR_CLOSE ..\parent_dir\dir_name DIR_OPEN ..\parent_dir\dir_name
NFS (Turbo Audit): Delete a	DIR_CLOSE ..\parent_dir\dir_name FILE_DELETE ..\parent_dir\dir_name\file_name

File	DIR_CLOSE ..\parent_dir\dir_name DIR_OPEN ..\parent_dir\dir_name DIR_OPEN ..\parent_dir\dir_name DIR_CLOSE ..\parent_dir\dir_name DIR_OPEN ..\parent_dir\dir_name
NFS (Turbo Audit): Create a Folder	DIR_CLOSE ..\parent_dir\dir_name\new_dir_name DIR_CLOSE ..\parent_dir\dir_name DIR_CREATE ..\parent_dir\dir_name\new_dir_name DIR_OPEN ..\parent_dir\dir_name DIR_OPEN ..\parent_dir\dir_name DIR_CLOSE ..\parent_dir\dir_name
NFS (Turbo Audit): Delete a Folder	DIR_CLOSE ..\parent_dir\dir_name DIR_CLOSE ..\parent_dir\dir_name\current_dir_name DIR_DELETE ..\parent_dir\dir_name\current_dir_name DIR_CLOSE ..\parent_dir\dir_name DIR_OPEN ..\parent_dir\dir_name DIR_OPEN ..\parent_dir\dir_name\current_dir_name DIR_OPEN ..\parent_dir\dir_name\current_dir_name DIR_CLOSE ..\parent_dir\dir_name\current_dir_name DIR_OPEN ..\parent_dir\dir_name DIR_CLOSE ..\parent_dir\dir_name DIR_OPEN ..\parent_dir\dir_name
NFS (Turbo	DIR_CLOSE ..\parent_dir\dir_name DIR_CLOSE ..\parent_dir\dir_name

Audit): Rename a Folder	DIR_RENAME ..\parent_dir\dir_name\current_dir_name DIR_OPEN ..\parent_dir\dir_name DIR_CLOSE ..\parent_dir\dir_name DIR_OPEN ..\parent_dir\dir_name DIR_OPEN ..\parent_dir\dir_name DIR_CLOSE ..\parent_dir\dir_name DIR_OPEN ..\parent_dir\dir_name
NFS (Turbo Audit): Modify ACL of a file	DIR_CLOSE ..\parent_dir\dir_name FILE_CLOSE ..\parent_dir\dir_name\file_name FILE_SET_ACL ..\parent_dir\dir_name\file_name FILE_CLOSE ..\parent_dir\dir_name\file_name DIR_OPEN ..\parent_dir\dir_name
NFS (Turbo Audit): Modify ACL of a Directory	DIR_CLOSE ..\parent_dir\dir_name\current_dir_name DIR_CLOSE ..\parent_dir\dir_name DIR_SET_ACL ..\parent_dir\dir_name\current_dir_name DIR_OPEN ..\parent_dir\dir_name\current_dir_name DIR_OPEN ..\parent_dir\dir_name

© Superna LLC

3.15. Advanced Configuration

Home Top

- Save all Reports Centrally to PowerScale Filesystem
- Filter-Out Event Messages - Turbo Audit
- How to change the file and directory events supported with Where did my folder go? Feature
- How to increase search concurrency and search performance

Save all Reports Centrally to PowerScale Filesystem

1. This feature allows all reports to be saved as CSV to an nfs mount created on the Eyeglass appliance
2. igls admin eaCsvArchivePath set --value=/opt/superna/sca/data/EA_reports
3. Now create an NFS export secured to the ip address of the eyeglass appliance
4. mkdir /opt/superna/sca/data/EA_reports
5. chown sca:users /opt/superna/sca/data/EA_reports
6. chmod 755 /opt/superna/sca/data/EA_reports
7. Create an NFS mount to the path EA_Reports
8. example **below to mount the report mount**
9. You will now need to perform the following mount command.

a. SSH into your Eyeglass appliance and gain root access. sudo -s (then enter the eyeglass password)
b. Execute the following: vim /etc/fstab
c. Add the command line below and replace it as indicated:
d. <Source-cluster-IP>:/<path_of_export> /opt/superna/sca/data/EA_reports nfs rw 0 0
e. **Replacing**:
    i. <path_of_export> by path of the NFS Export configured on your Source cluster.
    ii. Then execute the previous command using:
    iii. mount -a
    iv. Navigate to the export to make sure it is writeable
    v. cd /opt/superna/sca/data/EA_reports
    vi. ls > test.file
    vii. This command should succeed if write access is working.

Filter-Out Event Messages - Turbo Audit

Event Messages can be filtered out from the Audit Event processing to reduce the storage usage as well as the rate of processing events.

1. To configure the filter, add the following line in the **/opt/superna/eca/eca-env-common.conf** file
2. export BYPASSED_EVENT_TYPES=***<list of Events to be filter - comma separated>***
3. Default Events Filtered:

4. To filter-out DIR_SET_ACL,DIR_OPEN,DIR_CLOSE,DIR_SET_SEC events, add this line in the **/opt/superna/eca/eca-env-common.conf** file
5. export BYPASSED_EVENT_TYPES=DIR_SET_ACL,DIR_OPEN,DIR_CLOSE,DIR_SET_SEC
6. Verify that the Turbo Audit mode is also enabled
    a. export USE_TURBOAUDIT=true
    b. The supported list of events that can be specified in the Filter:
    c. FILE_OPEN_NOACCESS
    d. FILE_OPEN_READ
    e. FILE_OPEN_WRITE
    f. FILE_CREATE
    g. FILE_RENAME
    h. FILE_DELETE
    i. FILE_CLOSE
    j. FILE_CLOSE_MODIFIED
    k. FILE_SET_ACL
    l. FILE_READ
    m. FILE_WRITE
    n. DIR_CREATE
    o. DIR_RENAME
    p. DIR_DELETE

q. DIR_SET_ACL

r. DIR_OPEN

s. DIR_CLOSE

How to change the file and directory events supported with Where did my folder go? Feature

1. 2.5.6 builds after 84 default to the following file and directory events delete,rename. These events are stored in the Where did my folder go index. If you need to remove high volume event types, typically file delete events follow these steps.
2. login to node 1 of the ECA cluster as ecaadmin
3. nano /opt/superna/eca/eca-env-common.conf
4. add a line as per example below and remove any events from the list that you no longer want stored in the Where did my folder go index. **NOTE: This does not remove the events from the main index and the events will still be searchable with the query builder searches.**
5. export EVTARCHIVE_EVENT_TYPES=DIR_RENAME,DIR_DELETE,FILE_DELETE,FILE_RENAME
6. Save the file with control + X and answer yes to save
7. Restart the cluster with ecactl cluster down and then ecactl cluster up to have the changes take effect.

How to increase search concurrency and search performance

1. In order to increase search performance follow these steps to increase resources. Search performance is directly related to resources.
    a. **Scenario 1** - Allow concurrent search jobs to run in parallel up to a limit of 5
    b. **Scenario 2** - Long running searches that scan many days for many different event types, can exhaust memory. We recommend to use the specific time interval to search a week at a time or increase resources to allow for along running search.
    c. Increase the number of Isilon nodes in the HDFS pool from 3 to 6 or 9 nodes to increase load balancing of IO requests.
    d. (**Required**) Increase RAM on ECA nodes 2-N to 32 GB, open a case to have memory allocation changed for the search engine spark-workers and spark master containers. This is required to allocate the memory to searching.
    e. (**Required**) Increase CPU cores from 4 to 8 per VM on nodes 2-N. This is required for concurrent jobs.
    f. (optional) Increase the number of ECA nodes and expand the cluster size, if using 6 VM's increase to 9 VM's Deploy another ECA ova following the installation guide. Open a case to get the cluster expanded. This increases the

database and the search engine and has the highest impact on search performance.

2. **Contact Support once resources have been assigned to the VM's**

© Superna LLC

3.16. Excluded Audit Events

Home Top

The following events are excluded from ingestion, processing and storage. These event types are not material to auditing the filesystem and often appear in high rates and consume space in the database. If you believe you want to store these event types then open a support case.

Excluded Events

1. **FILE_OPEN_NOACCESS**
2. **FILE_CLOSE**
3. DIR_OPEN
4. DIR_CLOSE

© Superna LLC

3.17. How to Configure Syslog Forwarding of Formatted Audit messages to an External Syslog Server

Home Top

- Overview
- How to configure syslog forwarding
  - Example syslog message format sent by the ECA
  - How to view Syslog forwarding Statistics
  - How to debug syslog forwarding when you syslog server does not receive messages
- How to Configure event filtering before forwarding
  - Example filter for a path called /ifs/data/smb01/test123
  - Example of how to exclude audit records from /ifs/.ifsvar
- Syslog Configuration Forwarding Parameters

Overview

This configuration is optional and only needed when events should be forwarded another logging system link Splunk or other logging tools. The ECA can run an additional dock container that consumes events and formats for syslog forwarding. This container can run on all nodes and allows for parallel forwarding of events.

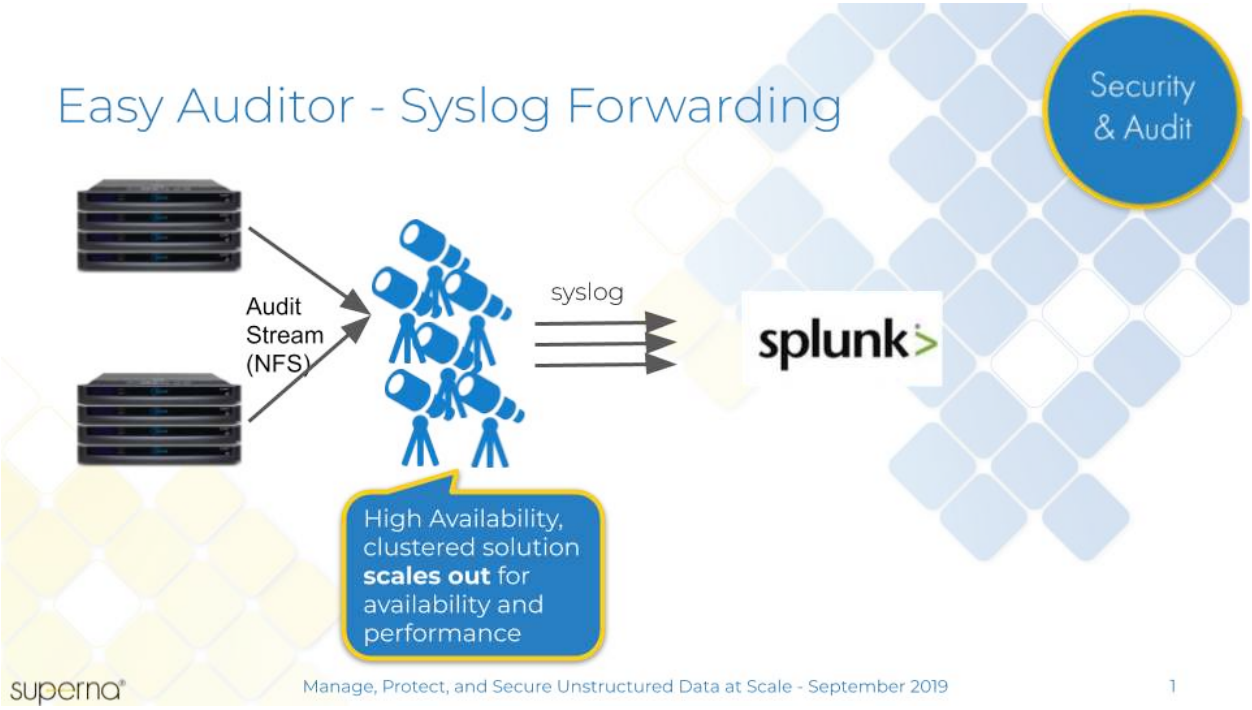

How to configure syslog forwarding

1. Login to each node that you want to enable syslog forwarding over ssh as ecaadmin. NOTE: each node needs the file edited to configure the forwarding. The instructions assume all nodes except node 1 will forward to syslog server.
2. vim /opt/superna/eca/conf/syslogpublisher/log4j2.xml
3. Add the ip address of the syslog server and the UDP port for your syslog server. **NOTE: You must edit the default port in the file 5140 and change this to the port used by your syslog server. NOTE: Default syslog servers use port 514**

a. 172.25.8.30 - PuTTY

```
<?xml version="1.0" encoding="UTF-8"?>
<Configuration>
  <Appenders>
    <Console name="STDOUT" target="SYSTEM_OUT">
      <PatternLayout pattern="%highlight{%d %C{1}:%L %-5level: %msg%n%throwable}"/
>
    </Console>
    <Syslog name="SupernaSyslog" format="RFC5424" facility="LOCAL0"
            host="172.25.8.25" port="5140" protocol="UDP" appName="ECA"
            messageId="AuditLogs" id="Event" connectTimeoutMillis="10000"
            newLine="true" mdcId="mdc" includeMDC="true" enterpriseNumber="18060">
     </Syslog>
  </Appenders>
  <Loggers>
    <Root level="ALL">
     <AppenderRef ref="STDOUT"/>
    </Root>
    <Logger name="org.apache.log4j.xml" level="info"/>
    <Logger name="SYSLOG" level="ALL">
      <AppenderRef ref="SupernaSyslog"/>
    </Logger>
  </Loggers>
</Configuration>
~
~
~
~
~
~
~
~
~
~
~
~
~
~
~
~
~
"log4j2.xml" 22L, 804C                                    8,72          All
```

4. Save the file
5. The docker container does not start by default to start this container, Add an entry to the docker overrides file to start the container on nodes 2-N.
   a. vim /opt/superna/eca/docker-compose.overrides.yml
   b. add a section as per image below make sure to space the entries exactly as shown

c.

6. save the file
7. To update all nodes with the new syslog configuration run the following command:
    a. ecactl cluster push-config
8. To start the container now and follow the steps below:
    a. Now create the container on all nodes
        i. ecactl cluster exec "ecactl containers up -d syslogpublisher"
        ii. NOTE: This will start the container
        iii. This will start the container on node 1 and should be removed in production environments
        iv. On node 1 run this commands to stop and remove the container, Answer yes to the prompt to remove.
            1. ecactl containers stop syslogpublisher
            2. ecactl containers rm syslogpublisher
9. Verify your syslog server is now receiving events sample syslog format below.
10. To monitor the forwarding function and events received and sent use this command to monitor the syslog container on one eca node (note all ECA nodes are forwarding events).

a. ecactl logs --tail 200 --follow syslogpublisher (this command will show stats every minute for events received by the container and sent to your syslog server).

11. done

Example syslog message format sent by the ECA

1. 2019-07-07T20:49:22.328Z syslogpublisher.node1.demoeca.eca.local ECA 1 AuditLogs - {"eventCode":"0x8","path":"\\\\00505699a9f1aecd965b770a3472e43955d2\\System\\ifs\\spark-logs\\.f6b02da1-6f42-43a3-b02a-ae87564c255b","protocol":"HDFS","server":"node 172.31.1.131 07/07 16:49:21.964
2. 2019-07-07T20:49:22.332Z syslogpublisher.node1.demoeca.eca.local ECA 1 AuditLogs - {"eventCode":"0x40","path":"\\\\00505699a9f1aecd965b770a3472e43955d2\\System\\ifs\\spark-logs\\.f6b02da1-6f42-43a3-b02a-ae87564c255b","protocol":"HDFS","server":"nod 172.31.1.131 07/07 16:49:21.964
3. 2019-07-07T20:49:22.333Z syslogpublisher.node1.demoeca.eca.local ECA 1 AuditLogs - {"eventCode":"0x8000","path":"\\\\00505699a9f1aecd965b770a3472e43955d2\\System\\ifs\\spark-logs\\.f6b02da1-6f42-43a3-b02a-ae87564c255b","protocol":"HDFS","server":"n 172.31.1.131 07/07 16:49:21.964

4. 2019-07-07T20:49:22.334Z syslogpublisher.node1.demoeca.eca.local ECA 1 AuditLogs - {"eventCode":"0x40000","path":"\\\\00505699a9f1aecd965b770a3472e43955d2\\System\\ifs\\spark-logs\\.f6b02da1-6f42-43a3-b02a-ae87564c255b","protocol":"HDFS","server":" 172.31.1.131 07/07 16:49:21.964
5. 2019-07-07T20:49:22.334Z syslogpublisher.node1.demoeca.eca.local ECA 1 AuditLogs - {"eventCode":"0x40","path":"\\\\00505699a9f1aecd965b770a3472e43955d2\\System\\ifs\\spark-logs\\.f6b02da1-6f42-43a3-b02a-ae87564c255b","protocol":"HDFS","server":"nod 172.31.1.131 07/07 16:49:21.964
6. 2019-07-07T20:49:22.335Z syslogpublisher.node1.demoeca.eca.local ECA 1 AuditLogs - {"eventCode":"0x40","path":"\\\\00505699a9f1aecd965b770a3472e43955d2\\System\\ifs\\spark-logs\\.f6b02da1-6f42-43a3-b02a-ae87564c255b","protocol":"HDFS","server":"nod 172.31.1.131 07/07 16:49:21.964
7. 2019-07-07T20:49:22.335Z syslogpublisher.node1.demoeca.eca.local ECA 1 AuditLogs - {"eventCode":"0x20","path":"\\\\00505699a9f1aecd965b770a3472e43955d2\\System\\ifs\\spark-logs\\.f6b02da1-6f42-43a3-b02a-ae87564c255b","protocol":"HDFS","server":"nod 172.31.1.131 07/07 16:49:21.964
8. 2019-07-07T20:59:40.816Z syslogpublisher.node1.demoeca.eca.local ECA 1 AuditLogs - {"eventCode":"0x40","path":"\\\\0050569960fcd70161594d21dd2

2a3c10cbe\\System\\ifs\\data\\policy1\\search\\cow.txt","protocol":"SMB2","server":"node001","clientIP":"1 172.31.1.131 07/07 16:59:40.430

9. 2019-07-07T20:59:40.817Z syslogpublisher.node1.demoeca.eca.local ECA 1 AuditLogs - {"eventCode":"0x40","path":"\\\\0050569960fcd70161594d21dd22a3c10cbe\\System\\ifs\\data\\policy1\\search\\cow.txt","protocol":"SMB2","server":"node001","clientIP":"1 172.31.1.131 07/07 16:59:40.446
10. 2019-07-07T20:59:40.817Z syslogpublisher.node1.demoeca.eca.local ECA 1 AuditLogs - {"eventCode":"0x2","path":"\\\\0050569960fcd70161594d21dd22a3c10cbe\\System\\ifs\\data\\policy1\\search\\cow.txt","protocol":"SMB2","server":"node001","clientIP":"17 172.31.1.131 07/07 16:59:40.446
11. 2019-07-07T20:59:40.821Z syslogpublisher.node1.demoeca.eca.local ECA 1 AuditLogs - {"eventCode":"0x40","path":"\\\\0050569960fcd70161594d21dd22a3c10cbe\\System\\ifs\\data\\policy1\\search\\cow.txt","protocol":"SMB2","server":"node001","clientIP":"1 172.31.1.131 07/07 16:59:40.446
12. 2019-07-07T20:59:40.823Z syslogpublisher.node1.demoeca.eca.local ECA 1 AuditLogs - {"eventCode":"0x40","path":"\\\\0050569960fcd70161594d21dd22a3c10cbe\\System\\ifs\\data\\policy1\\search\\cow.txt","protocol":"SMB2","server":"node001","clientIP":"1 172.31.1.131 07/07 16:59:40.446

13. 2019-07-07T20:59:40.829Z syslogpublisher.node1.demoeca.eca.local ECA 1 AuditLogs - {"eventCode":"0x40","path":"\\\\0050569960fcd70161594d21dd22a3c10cbe\\System\\ifs\\data\\policy1\\search\\cow.txt","protocol":"SMB2","server":"node001","clientIP":"1 172.31.1.131 07/07 16:59:40.446
14. 2019-07-07T20:59:40.831Z syslogpublisher.node1.demoeca.eca.local ECA 1 AuditLogs - {"eventCode":"0x40","path":"\\\\0050569960fcd70161594d21dd22a3c10cbe\\System\\ifs\\data\\policy1\\search\\cow.txt","protocol":"SMB2","server":"node001","clientIP":"1 172.31.1.131 07/07 16:59:40.461

How to view Syslog forwarding Statistics

1. NOTE: No log exists to see events within the ECA. The forwarding feature uses programmatic access to an internal message bus that is not exposed to viewable.
2. To view statistics of each ECA nodes forwarding function run this command on node 1 of the eca cluster after logging as the ecaadmin user.
3. ecactl cluster exec "ecactl logs --tail 20 syslogpublisher"
4. Each node will output the events it received for forwarding and the number of files it sent to the configured syslog server
5. example
    a. In the example below the Sent events shows the total all time and the rate per second over the last minute. Stats are updated each minute on each node. The example below

shows a rate of 794 audit events forwarded per second over the last minute. The **last event ts** date is the date of the time stamp indicating when the event itself was created on the cluster. This gives you an idea of how current relative to the current time versus the forwarding functions progress.

b. syslogpublisher | 2020-09-21 12:58:28,269 AnalysisModule:146 INFO : Events Sent: | total 15040319 | **rate 794.77** | **last event ts**: Mon Sep 21 12:58:28 UTC 2020

How to debug syslog forwarding when you syslog server does not receive messages

1. These steps assume you have check firewalls and verified this is not the issue and that the correct forwarding port has been used on the ECA configuration.
2. login to eca node 2, 3, 4, 5 etc.. since each node forwards syslog messages. You should use the stats command in the above section to determine which ECA node is showing sent audit events. Then use tcpdump on that node to capture UDP messages.
3. sudo -s (to become root user)
4. zypper in tcpdump (requires Internet connection to ECA)
5. Monitor UDP syslog on the ECA node
    a. tcpdump -i eth0 udp port 514 (this command will display all UDP packets on port 514 to the console)

6. The messages will look like this once a packet is captured with the destination host name or IP shown. Yellow highlight in the example.
    a. 09:54:43.379985 IP 172.31.1.135.45750 > syslog.internal.superna.net.syslog: SYSLOG local0.info, length: 664

How to Configure event filtering before forwarding

Use this configuration to select events from a specific path or below and forward only these events to the syslog server. This avoids a large volume of syslog data being sent when only a subset is needed. This same concept can be used to pattern match on a SID, event type in the raw syslog message. You will need to experiment with the pattern match for specific events.

The example below covers path based matching, to match against other fields setup forwarding first to syslog, then review the fields in the event messages to build matching filters for other fields such as user or event type.

1. Review the syntax below and edit the log4j2.xml file to add your filter
    a. **Example to match all syslog events for the path /ifs/data/smb01/test123**
    b. **login to eca node 1 as ecaadmin over ssh**
    c. **vim /opt/superna/eca/conf/syslogpublisher/log4j2.xml**

d. **Insert a line using yellow highlighted example below and adjust the filter for your matching criteria.**
e. **Save the file**
f. **:wq**
g. **done**
h. **Push the configruation to all nodes**
i. **ecactl cluster push-config**
j. **Restart the syslog publishing container on all ECA nodes to reload the configuration**
k. ecactl cluster services restart --container sysylogpublisher --all
l. **done**

Example filter for a path called /ifs/data/smb01/test123

After making a change to this file you must

1. ecactl cluster push-config
2. ecactl cluster services restart --container sysylogpublisher --all

```
<?xml version="1.0" encoding="UTF-8"?>
<Configuration>
<Appenders>
<Console name="STDOUT" target="SYSTEM_OUT">
<PatternLayout pattern="%highlight{%d %C{1}:%L %-5level:
%msg%n%throwable}"/>
</Console>
```

```
<Syslog name="SupernaSyslog" format="RFC5424"
facility="LOCAL0"
host="172.22.4.19" port="514" protocol="UDP" appName="ECA"
messageId="AuditLogs" id="Event" connectTimeoutMillis="10000"
newLine="true" mdcId="mdc" includeMDC="true"
enterpriseNumber="18060">
<RegexFilter regex=".*ifs.*data.*smb01.*test123.*" useRawMsg="true"
onMatch="ACCEPT" onMismatch="DENY"/>
</Syslog>
</Appenders>
<Loggers>
<Root level="ALL">
<AppenderRef ref="STDOUT"/>
</Root>
<Logger name="org.apache.log4j.xml" level="info"/>
<Logger name="SYSLOG" level="ALL">
<AppenderRef ref="SupernaSyslog"/>
</Logger>
</Loggers>
</Configuration>
```

Example of how to exclude audit records from /ifs/.ifsvar

After making a change to this file you must

1. ecactl cluster push-config
2. ecactl cluster services restart --container sysylogpublisher --all

```
<?xml version="1.0" encoding="UTF-8"?>
```

```
<Configuration>
<Appenders>
<Console name="STDOUT" target="SYSTEM_OUT">
<PatternLayout pattern="%highlight{%d %C{1}:%L %-5level:
%msg%n%throwable}"/>
</Console>
<Syslog name="SupernaSyslog" format="RFC5424"
facility="LOCAL0"
host="172.22.4.19" port="514" protocol="UDP" appName="ECA"
messageId="AuditLogs" id="Event" connectTimeoutMillis="10000"
newLine="true" mdcId="mdc" includeMDC="true"
enterpriseNumber="18060">
<RegexFilter regex="^((?!.*ifsvar.*).)*$" useRawMsg="true"
onMatch="ACCEPT" onMismatch="DENY"/>
</Syslog>
</Appenders>
<Loggers>
<Root level="ALL">
<AppenderRef ref="STDOUT"/>
</Root>
<Logger name="org.apache.log4j.xml" level="info"/>
<Logger name="SYSLOG" level="ALL">
<AppenderRef ref="SupernaSyslog"/>
</Logger>
</Loggers>
</Configuration>
```

Syslog Configuration Forwarding Parameters

Adavanced options for forwarding.

Parameters:

host - The name of the host to connect to.

port - The port to connect to on the target host.

protocolStr - The Protocol to use.

sslConfiguration - TODO

connectTimeoutMillis - the connect timeout in milliseconds.

reconnectDelayMillis - The interval in which failed writes should be retried.

immediateFail - True if the write should fail if no socket is immediately available.

name - The name of the Appender.

immediateFlush - "true" if data should be flushed on each write.

ignoreExceptions - If "true" (default) exceptions encountered when appending events are logged; otherwise they are propagated to the caller.

facility - The Facility is used to try to classify the message.

id - The default structured data id to use when formatting according to RFC 5424.

enterpriseNumber - The IANA enterprise number.

includeMdc - Indicates whether data from the ThreadContextMap will be included in the RFC 5424 Syslog record. Defaults to "true:.

mdcId - The id to use for the MDC Structured Data Element.

mdcPrefix - The prefix to add to MDC key names.
eventPrefix - The prefix to add to event key names.
newLine - If true, a newline will be appended to the end of the syslog record. The default is false.
escapeNL - String that should be used to replace newlines within the message text.
appName - The value to use as the APP-NAME in the RFC 5424 syslog record.
msgId - The default value to be used in the MSGID field of RFC 5424 syslog records.
excludes - A comma separated list of mdc keys that should be excluded from the LogEvent.
includes - A comma separated list of mdc keys that should be included in the FlumeEvent.
required - A comma separated list of mdc keys that must be present in the MDC.
format - If set to "RFC5424" the data will be formatted in accordance with RFC 5424. Otherwise, it will be formatted as a BSD Syslog record.
filter - A Filter to determine if the event should be handled by this Appender.
configuration - The Configuration.
charset - The character set to use when converting the syslog String to a byte array.
exceptionPattern - The converter pattern to use for formatting exceptions.
loggerFields - The logger fields
advertise - Whether to advertise

© Superna LLC

3.18. Data Retention of Audit Data and Archive

Home Top

- Overview
- Actions

Overview

Two types of retention are supported, online and long term archive.

1. Online means audit data that is searchable in the index. Operational requests for auditing typically do not exceed 30 days but we support up-to 18 months online searchable depending on the audit rate. This is for optimal performance and maintaining a manageable database size. Contact support to get automatic data retention applied to the database. Typical values are 6 months or 1 year online searchable.
2. Database is the secondary version of audit data and for all long term retention requirements raw audit in GZ format should always be stored for long term retention as the data is shareable in a format needed for auditors and is compatible with OneFS tools. The database is not in a format that auditors can use in a sharable format.
3. Database management tasks and size of DB require online searching to prune data older than 18 months maximum.

4. The ECA VM deployment is designed for online search and VM count may need to be increased to manage a large database up to 18 months of data.
5. Long term storage of audit data
    a. PowerScale audit data must be purged as the raw audit data is stored on the PowerScale in GZ format and is never deleted. For long term storage of audit data depending on business need this format should be stored in an archive location in GZ format. We recommend purging these GZ files twice per year. See EMC SR requirement and steps documented here.
    b. **NOTE: Preserve the folder hierarchy of the audit data with nodes and created and modified date stamps when archiving this data. The only method to ingest GZ data is based on the data stamp of the GZ files that determines the date range covered by the GZ files for ingestion.**
    c. The GZ files impact ingestion audit performance if 1000's of files are left on PowerScale and the purge process allows them to be removed and archived at the same time. **This is the other requirement to remove old GZ data to reduce NFS audit data bandwidth ingestion.**

Actions

1. Open a support case to get the database retention set to 6m , 1 year or 18 months
2. Review PowerScale GZ purge procedure
3. Identify long term archive location for GZ audit data

© Superna LLC

3.19. How to Purge or Archive PowerScale Audit logs

Home Top

How to purge or archive Audit PowerScale logs

PowerScale stores audit messages in archived compressed files and does not have an automatic purge process. These steps should be used to correctly remove old GZ files and ensure audit protocol is operating normally after the purge process on all nodes in the cluster.

CAUTION!

This procedure will stop capturing audit events on the cluster during the time auditing is disabled. **NOTE: We recommend you open an EMC SR with PowerScale steps, Superna support cannot support this procedure or trouble shoot steps on the cluster releated to this procedure.**

IMPORTANT!

This procedure must be performed using the "root" account on the cluster. **Please consult EMC for an updated procedure.**

1. Stop the ECA cluster
    a. ssh eca master node as ecaadmin
    b. ecactl cluster down
2. Run the following commands to turn off audit logging
    OneFS 8.0.0 and later
    1. isi audit settings global modify --protocol-auditing-enabled=no
    2. isi audit settings global modify --config-auditing-enabled=no (**only if enabled before)**
3. Run the following commands to stop the isi_audit_d, isi_audit_cee and isi_audit_syslog processes from automatically restarting:
    a. isi services -a isi_audit_d ignore
    b. isi services -a isi_audit_cee ignore
    c. isi services -a isi_audit_syslog ignore
4. Run the following commands to end the isi_audit_d and isi_audit_cee processes:
    a. isi_for_array 'pkill isi_audit_d'

   b. isi_for_array 'pkill isi_audit_cee'
   c. isi_for_array 'pkill isi_audit_syslog'
5. Run the following command to ensure that no isi_audit processes are running on the cluster:
   a. isi_for_array pgrep -l isi_audit
6. Run the following commands to change directory to the audit directory.
   a. cd /ifs/.ifsvar/audit
7. Run the following command to backup the audit directory and allow for the files to be recreated:
   a. mv /ifs/.ifsvar/audit /ifs/.ifsvar/audit.bak
   b. Archive the moved audit data to long term storage to retain a permanent copy of the source data or decide if you want to delete this data by consulting with your compliance department requirements.
8. Run the following commands to inform the Master Control Program (MCP) to resume monitoring the audit daemons. MCP automatically restarts the audit daemons and reconstructs the audit directory on each node when the isi_audit_d process is running.
   a. isi services -a isi_audit_d monitor
   b. isi services -a isi_audit_cee monitor
   c. isi services -a isi_audit_syslog monitor
9. Run the following command to check that audit processes have restarted:
   a. isi_for_array -s pgrep -l isi_audit
10. Run the following command to verify that audit data was removed and reconstructed:
   a. find /ifs/.ifsvar/audit
11. Run the following command to re-enable audit logging:
   a. **OneFS 7.1.0 - 7.2.1:**
      i. isi audit settings modify --protocol-auditing-enabled=Yes
      ii. isi audit settings modify --config-auditing-enabled=Yes (**only if enabled before)**
   b. **OneFS 8.0.0 and later**
      i. isi audit settings global modify --protocol-auditing-enabled=Yes
      ii. isi audit settings global modify --config-auditing-enabled=Yes (**only if enabled before)**
12. Run the following command to verify log files are being populated after audit processes have restarted:
   a. Reset audit log to current day and time
      i. isi audit settings global modify --cee-log-time "Protocol@2017-11-21 04:13:00" (**use a current date and time**)
      ii. isi_audit_viewer -t protocol
         1. Verify output from this command returns correctly last logged event.

13. On ECA master node **as ecaadmin user**
    a. ecactl cluster up
14. Login to eyeglass and verify the Managed Services icon shows active and green ECA nodes. NOTE: heartbeats take 2-5 minutes before the ECA cluster is completely up
15. If running Ransomware Defender run the Security Guard feature to test that audit messages are being processed correctly
16. End procedure

© Superna LLC

3.20. Bulk Ingest old Audit Data from PowerScale to Easy Auditor

Home Top

- Overview
  - Limitations
  - PowerScale Steps
  - Eyeglass steps
- How to View Progress of Bulk Audit Data Ingestion
  - Start up Queue Monitor Process
  - View Ingestion Jobs
  - View Event Ingestion progress

Overview

Use these instructions to re-ingest audit data from PowerScale's audit directory into Easy Auditor's index.

Limitations

1. Not supported for ingesting days, weeks months or years of data. Only supported for targetted specific date of an event. Support will not assist in bulk ingestion. Support of any use case other than targetted day is not supported under the support contract.
2. No possible method to predict time to ingest audit data.

3. IMPORTANT: Maximum number of files that can be added to json file to be run at any time = 20. ANY HIGHER NUMBER IS NOT SUPPORTED. Initial testing should be done with only 1 file. Use cron to run at non-busy time.
4. NOTE: Submitting more than one json file will queue the jobs and only 1 ingestion job will execute at a time.
5. NOTE: bulk ingestion is a background task and processing of current audit data has higher priority. There is no way to change this priority and no way to predict completion times since active audit data will take priority.

PowerScale Steps

1. SSH to the PowerScale cluster you intend to re-ingest audit logs from
2. Navigate to the directory you intend to re-ingest audit logs from. This directory is at the bottom of the below path (choose node8 as it was the most recent but yours will vary):
   `cd /ifs/.ifsvar/audit/logs/node008/protocol`
3. List the contents. This will assist determining which audit data based on dates/times to re-ingest. The audit logs are listed as .gz files.
   `ls -lT`

```
isilon-1-3# ls -lT
total 3592736
-rw-------    1 root  wheel   49153531 Sep 12 04:57:18 2018 00000000.gz
-rw-------    1 root  wheel   47456977 Sep 13 22:00:13 2018 00000001.gz
-rw-------    1 root  wheel   46336731 Sep 15 06:09:49 2018 00000002.gz
-rw-------    1 root  wheel   48135225 Sep 16 23:40:21 2018 00000003.gz
-rw-------    1 root  wheel   47521049 Sep 18 22:19:48 2018 00000004.gz
-rw-------    1 root  wheel   47191156 Sep 20 11:50:55 2018 00000005.gz
-rw-------    1 root  wheel   47560030 Sep 21 18:25:22 2018 00000006.gz
```

Eyeglass steps

1. SSH to Eyeglass CLI as `admin`

2. Navigate to

   `cd /home/admin`

3. Create a file

   `touch bulkingest.json`

4. Open the file in vim editor

   `vim bulkingest.json`

5. Copy paste content below (if ingesting from a single node), substituting in the following for your own:

```
[{
 "cluster_name": "YOUR_PowerScale_CLUSTER_NAME",
 "cluster_guid": "YOUR_PowerScale_CLUSTER_GUID",
 "node": [{
     "node_id": "node008",
     "audit_files": ["node_audit_file.gz",
"node_audit_file.gz"]
   }
 ]
}]
```

6. Save the file

   `:wq!`

```
7.If you wish to ingest from multiple nodes, please
  use the below code

  EXAMPLE:

  [{
   "cluster_name": "YOUR_PowerScale_CLUSTER_NAME",
   "cluster_guid": "YOUR_PowerScale_CLUSTER_GUID",
```

```
 "node": [{
     "node_id": "node008",
     "audit_files": ["node_audit_file.gz",
"node_audit_file.gz"]
   },
   {
     "node_id": "node003",
     "audit_files": ["node_audit_file.gz",
"node_audit_file.gz", "node_audit_file.gz",
"node_audit_file.gz"]
   }
 ]
}]
```

8. Save the file

   `:wq!`

9. Execute the bulkingest.json file (must be absolute path, does not matter where the file is located):

   ```
   igls rswsignals bulkLoadTAEvents --file=/home/admin/bulkingest.json
   ```

   NOTE: Depending on how large a period of time is being ingested, it can take some time to complete.

How to View Progress of Bulk Audit Data Ingestion

1. NOTE: Each Isilon node has historical audit data and each file is 1GB in size compressed, each node may have multiple files to ingest for a given day. The higher the audit rate the more GZ files that will need to be ingested. This can be a slow process for multiple days of historical audit data.

2. Start up Queue Monitor Process

   a. Login to node 1 of the ECA cluster and run this command

   b. ecactl containers up -d kafkahq

3. View Ingestion Jobs

   a. Open the Eyeglass **Jobs Icon running jobs tab** to view the bulk ingestion task to verify it is running, each CLI command will start an audit job to process the files in the json configuration file

      i. RUNNING Jobs Screen

      ii. You need to wait for this to finish before submitting more files. Currently processing

      iii. **Wait for Spark Job** step must show a blue checkmark to indicate it is finished processing this job.

         1. Spinning symbol means it is in progress.

      iv. 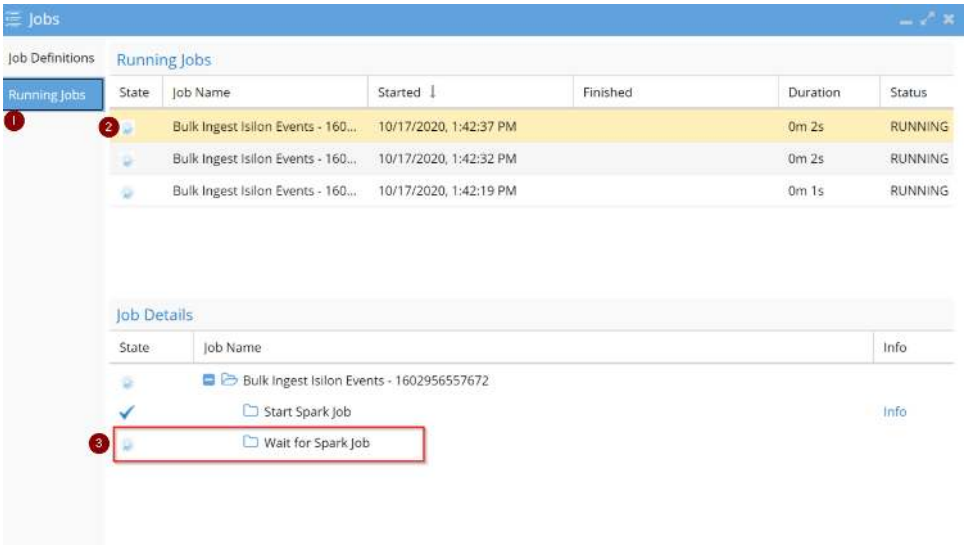

4. View Event Ingestion progress

a. **Using a browser http://x.x.x.x:9000 (x.x.x.x is node 1 of the ECA cluster IP address)**

b. **Topics** - in our case we are tracking specific bulkingestion topic

c. **Bulkingestion** - The topic used to track current progress on ingestion task

d. **Lag** - this can go UP/Down depending on the ingestion speed. A value of 0 means the ingestion job is finished and no more files events are being processed for the current active job. Check running Jobs icon to verify the active job shows finished. Any queued Jobs will start automatically.

e. **Count** - this field should always show an increase as new events are processed from all jobs. Once you add more json files to the queue this value will increase as events are ingested. This field will always increase each time ingestion jobs are run.

f.

© Superna LLC

4. LiveOps - Continuous Operations Admin Guide

- Overview
- LiveOPS Continuous Operations Dashboard
- Overview - LiveOPS Snapshot Sync
- Overview - LiveOPS Dedupe Sync
- Overview - LiveOPS DR Test mode
- How to Enable DR Test Mode
- How to Disable DR Test Mode
- How DR Test mode Jobs are displayed in Eyeglass UI
- Advanced DR Test mode Configurations

© Superna LLC

4.1. Overview

Overview

This Admin Guide is for the LiveOps features being introduced in Eyeglass releases. LiveOps allows customers to move from DR solutions to "**Continuous Operations**".

1. Continuous Operations allows storage administrators to allow online operations during production hours without risk.
2. Ensures all daily data protection and storage management functions are synced to the replication pair cluster.
3. Allows a storage admin to failover some or all data to a cluster pair, and to operate the business with all storage policies for storage management in Sync, data in sync, configuration in sync and snapshot and dedupe settings.

© Superna LLC

4.1.1. LiveOps Continuous Operations Key Features

LiveOps Continuous Operations Key Features

The LiveOps Continuous Operations features introduce a dashboard on the DR Dashboard that provides Continuous Operations Readiness, in the same way DR Readiness is calculated and displayed for clusters under Eyeglass management.

1. DR Test Mode
2. Snapshot settings Sync
3. Dedupe settings sync

© Superna LLC

4.1.2. What's New

What's New

Release 2.5.6

1. Multiple DR test mode policies can be enabled at the same time.
2. Config sync task can be disabled when executing, in order to save time on execution.
3. Rest API support to stop and start with option to turn off config job during execution.

© Superna LLC

4.2. LiveOPS Continuous Operations Dashboard

LiveOPS Continuous Operations Dashboard

Similar to the DR dashboard that provides status on DR Readiness, errors for Access Zones and clusters, the LiveOPS icon on the Eyeglass Desktop provides a single pane of glass to see cluster and policy sync Status for Snapshots and dedupe settings on clusters.

You can also view cluster reachability from the Eyeglass appliance to display all managed clusters. This dashboard also shows cluster release OneFS release and the effective API version in use between Eyeglass and the cluster.

© Superna LLC

4.2.1. What do the columns Mean?

Home Top

What do the columns Mean?

1. **Cluster Name** column lists cluster managed by Eyeglass.
2. **Cluster reachability** column indicates if Eyeglass can login to the cluster (tested every minute).
3. **Cluster Version** indicates the detected version of OneFS the cluster is running.
4. **Effective Cluster Version** means Eyeglass is in mixed API mode and uses the lowest cluster version API, you can see below a OneFS 8 cluster is operating at API 7.x. This means only objects or attributes supported in OneFS 7.x will sync to 8.x clusters.
5. **Continuous OPs Status** Summarizes each policy and cluster wide status, indicated snapshots and dedupe settings are in sync and audited between the cluster pairs that are replicating.

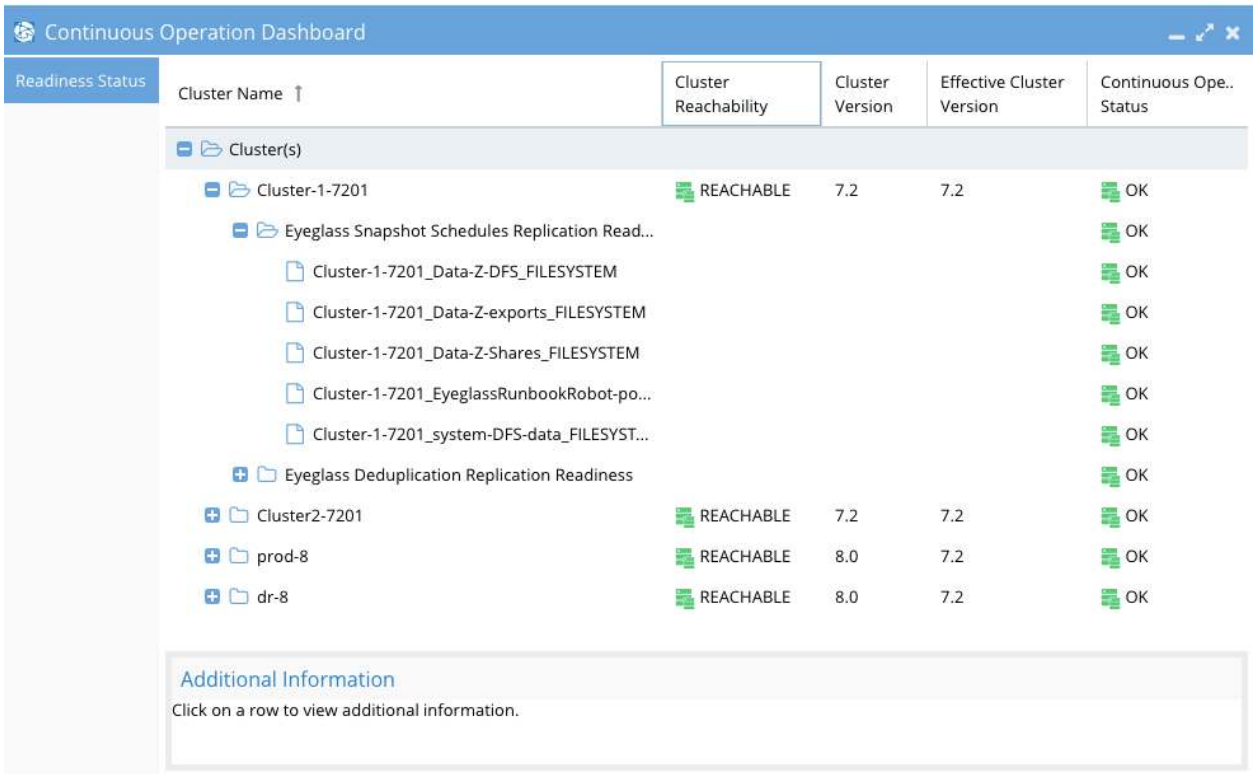

© Superna LLC

4.3. Overview - LiveOPS Snapshot Sync

Overview - LiveOPS Snapshot Sync

Replicating pairs of clusters now have a new job type in the Jobs window that will scan for Snapshot policies at or under SyncIQ policies, and sync them to the same or different path based on SyncIQ target path. This process scans for changes and syncs the changes to the schedule or other settings.

© Superna LLC

4.3.1. How it Works

How it Works

1. Any changes to the Snapshot settings will be synced to the peer cluster of the policy including path changes.
2. If the policy is failed over, then the target cluster owns the Snapshot settings and any changes or new Snapshots will be synced back to the source cluster.
3. Any other options on the settings tab are not synced.

© Superna LLC

4.3.2. How to Enable

Home Top

How to Enable

1. Open Jobs icon

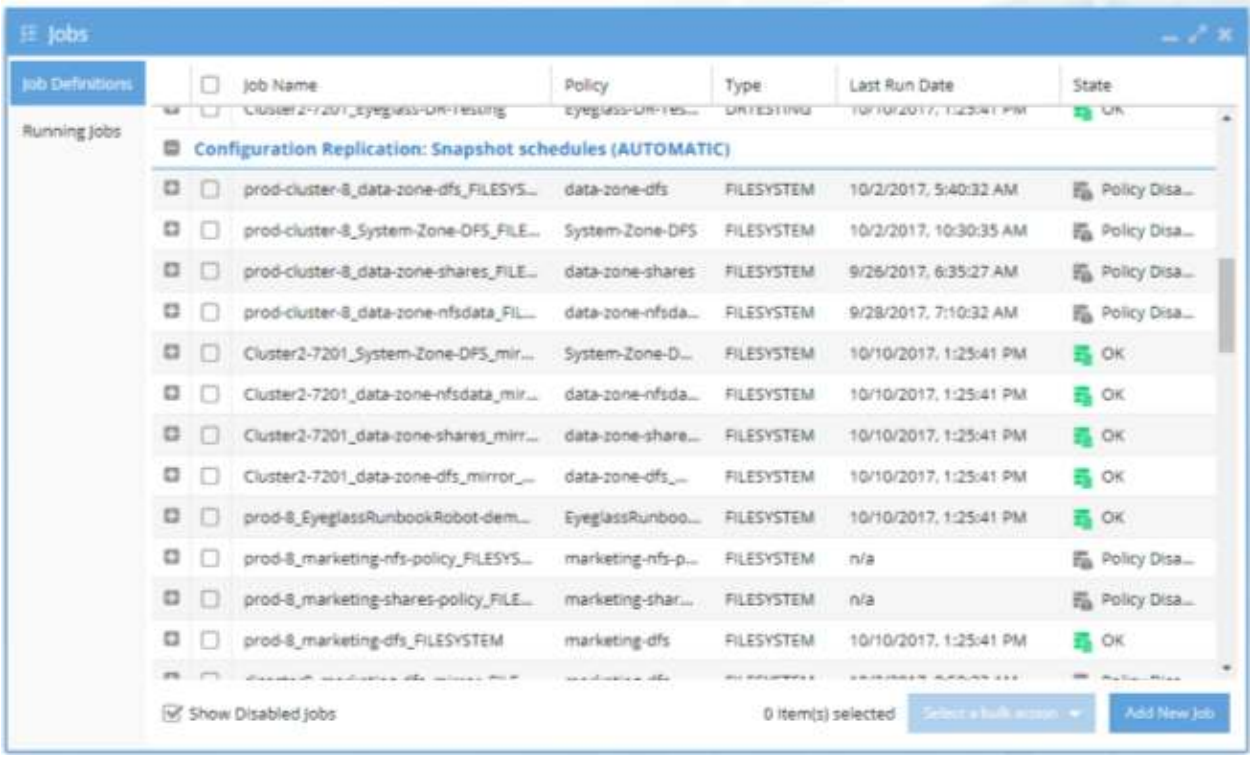

2. Enable the Snapshot Sync jobs
3. Select one or wait until Configuration sync runs again.
4. The Running Jobs window shows which policies have snapshot updates that are synced to the peer cluster on the same path specified by the SyncIQ policy

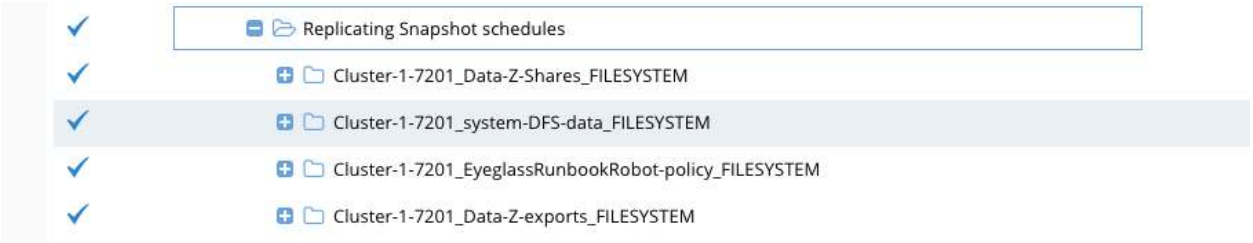

5. Verify by checking the snapshot schedules tab on OneFS UI

© Superna LLC

4.4. Overview - LiveOPS Dedupe Sync

Overview - LiveOPS Dedupe Sync

Replicating dedupe settings between replicating clusters allows dedupe to process the data on the DR cluster and achieve the same disk space savings, so that post failover, the cluster usage matches the source cluster usage. This ensures normal operations on the target cluster without any time delay for dedupe jobs to reduce the data before normal usage levels are achieved.

© Superna LLC

4.4.1. How it Works

Home Top

How it Works

1. File system paths and assessment paths are added to the source cluster AND match a SyncIQ policy path that will sync to the target cluster of the SyncIQ policies managed by Eyeglass.

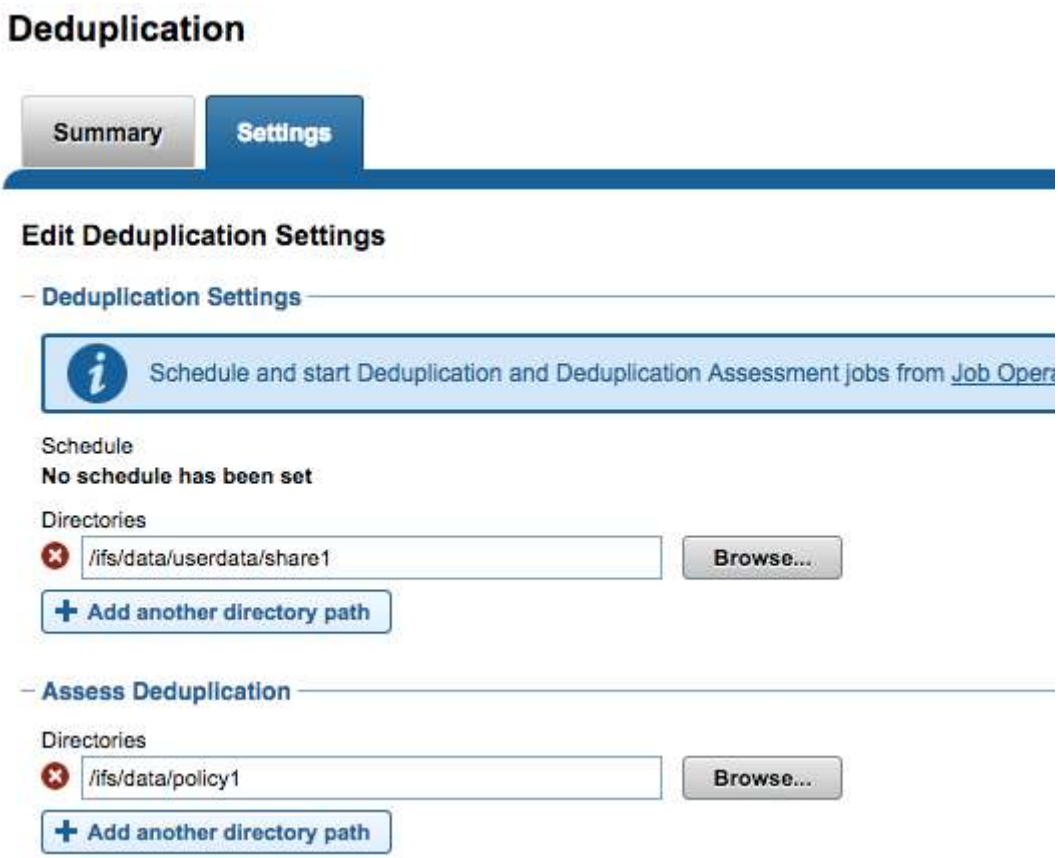

2. The dedupe scheduled job is not synced and must be setup or changed on the target cluster manually.
3. Licenses for dedupe should exist on the target cluster.

© Superna LLC

4.4.2. How to Enable

Home Top

How to Enable

1. Dedupe paths are **auto** synced with normal configuration replication, no setup required, it's automatic.
2. Dedupe paths added to the target cluster are only paths that match a SyncIQ policy path or below in the file system.
3. Dedupe settings for path are synced once for all policies not per policy

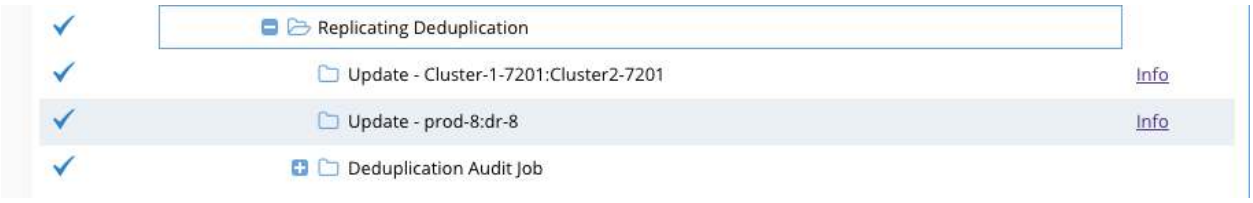

4. Verify paths are added to the target cluster

© Superna LLC

4.5. Overview - LiveOPS DR Test mode

Overview - LiveOPS DR Test mode

Eyeglass has introduced LiveOps features aimed at zero downtime, with the first feature offering DR testing that allows IT use cases to be executed without incurring downtime or impact to production systems.

1. Test DR procedures during normal business hours with full copy of production data. Storage administrators can validate application data integrity post failover, test end to end application procedures under various simulated conditions.
2. Upgrade testing of applications in a sandbox that allows writeable copy of production data.
3. Execute planned monthly or quarterly DR tests.
4. Mirror shares, exports to test access to production data in a sandbox.
5. Allows isolated testing using a different SmartConnect name to allow testing without overlap with production

The key capability introduced with LiveOps DR Testing is the ability to avoid impacting DR readiness for failover with SyncIQ replication fully operational and Eyeglass configuration sync running between production and DR clusters. Production to DR replication and failover abilities is never disrupted during DR Test mode.

© Superna LLC

4.5.1. DR Test Mode Architecture Diagram

Home Top

Refer to these diagrams to understand two different deployment scenario's with DR test mode for testing only **Isolated** and **Integrated** option where production AD and DNS are used for DR testing.

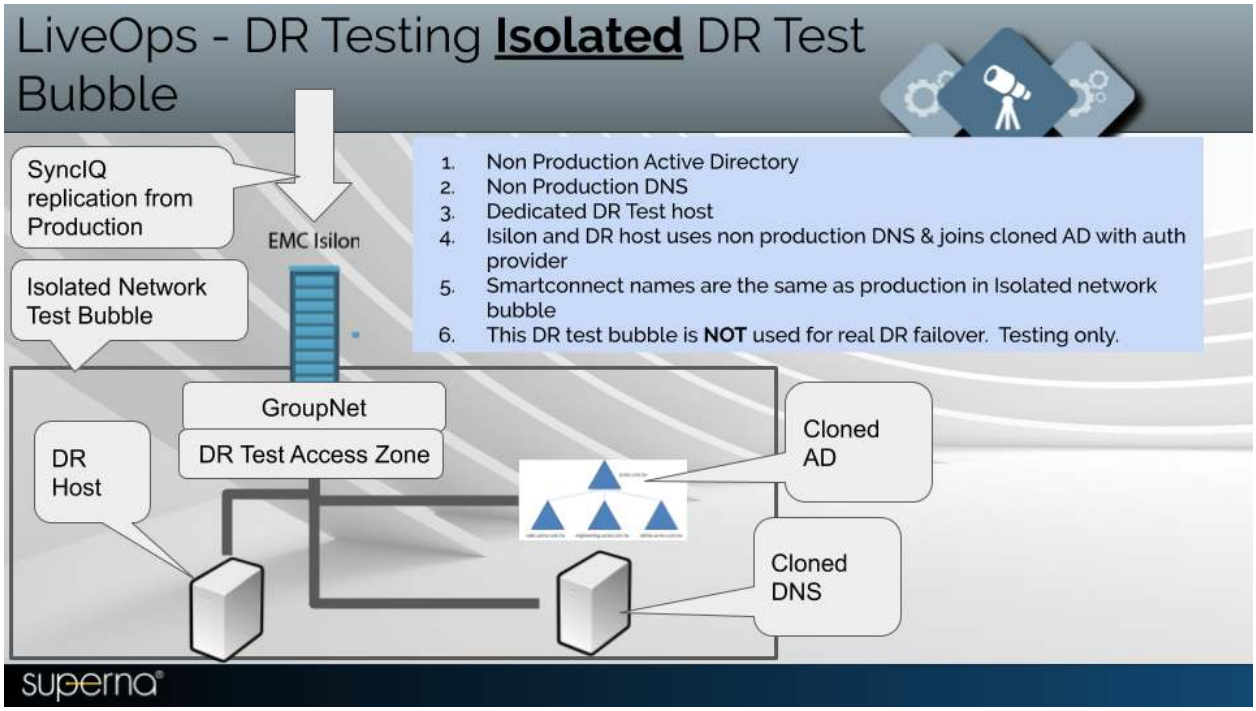

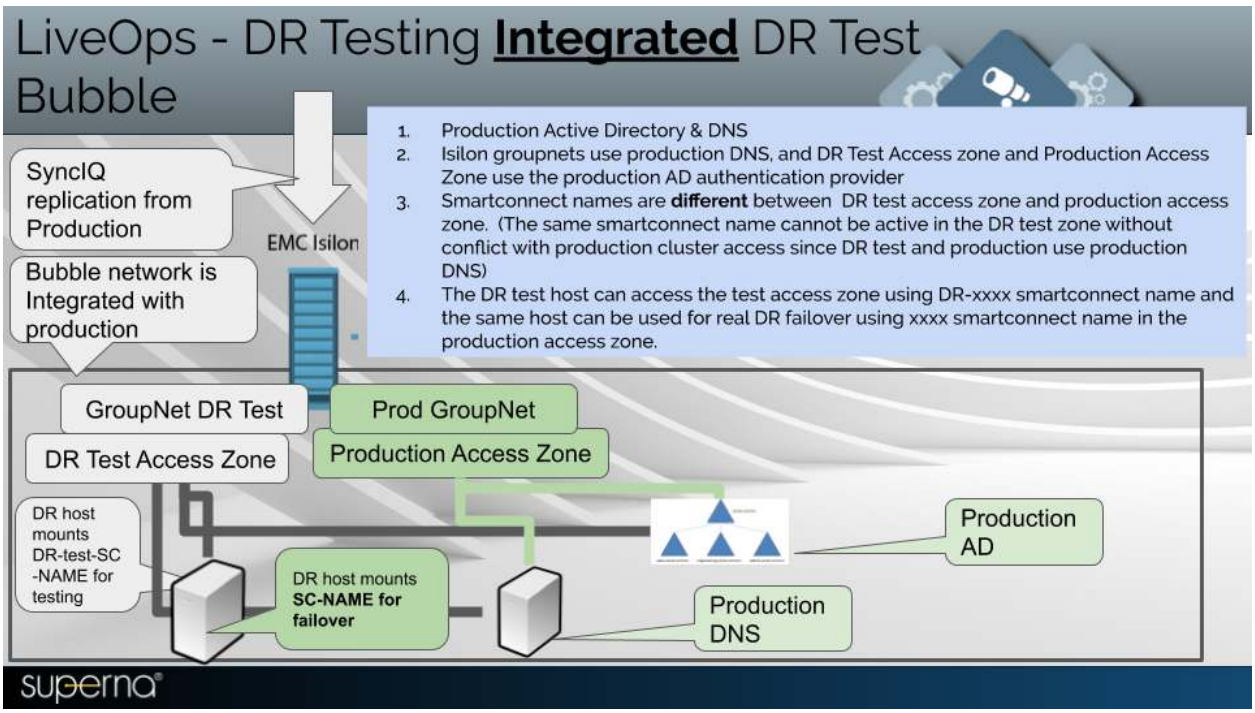
LiveOps - DR Testing Integrated DR Test Bubble
1. Production Active Directory & DNS
2. Isilon groupnets use production DNS, and DR Test Access zone and Production Access Zone use the production AD authentication provider
3. Smartconnect names are different between DR test access zone and production access zone. (The same smartconnect name cannot be active in the DR test zone without conflict with production cluster access since DR test and production use production DNS)
4. The DR test host can access the test access zone using DR-xxxx smartconnect name and the same host can be used for real DR failover using xxxx smartconnect name in the production access zone.
SyncIQ replication from Production
EMC Isilon
Bubble network is Integrated with production
GroupNet DR Test
Prod GroupNet
DR Test Access Zone
Production Access Zone
DR host mounts DR-test-SC -NAME for testing
DR host mounts SC-NAME for failover
Production AD
Production DNS
superna

© Superna LLC

4.5.2. Operating View

Home Top

Operating View

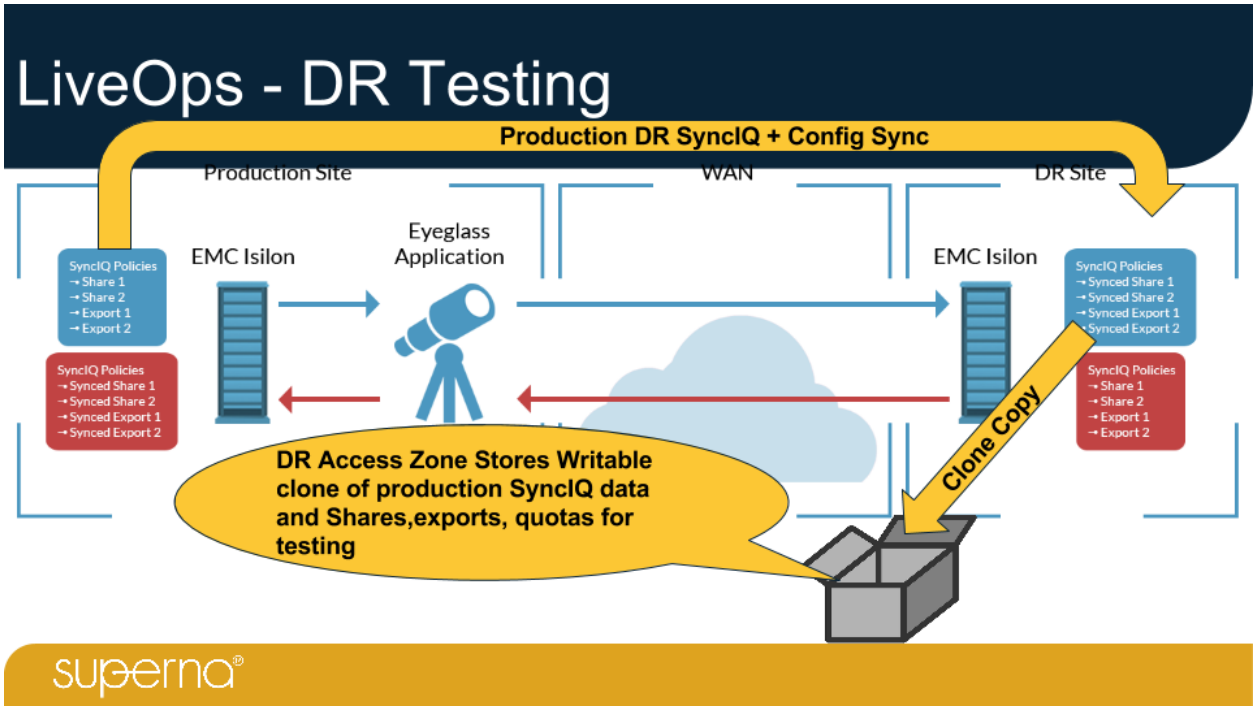

3rd copy of data is read-only until ready for DR Testing. Data and configuration is in sync with production in a sandbox Access Zone that isolates data and access using an Access Zone with dedicated IP pool to for DR testing.

In Eyeglass it's easy to see the status of the 3rd copy from the DR Dashboard.

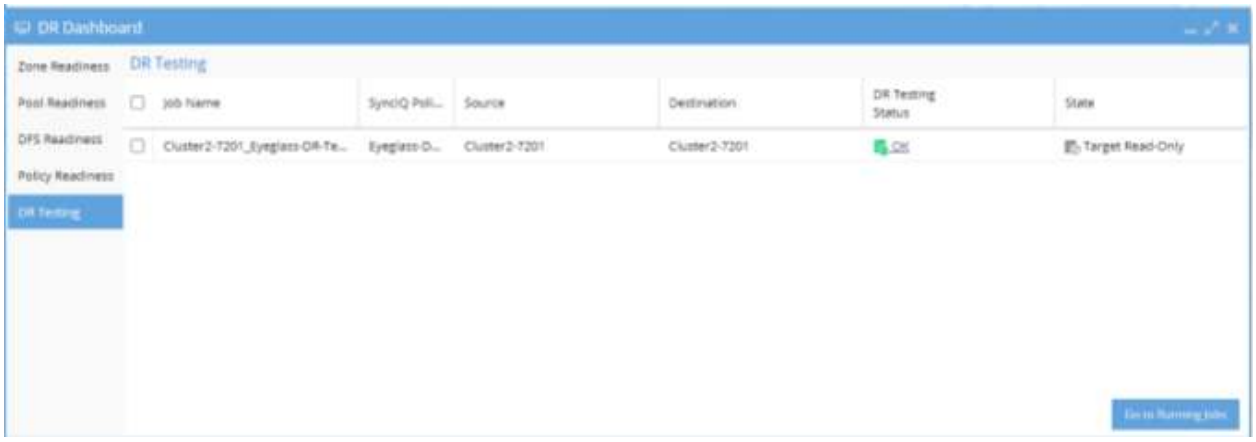

© Superna LLC

4.5.3. Prerequisites

Prerequisites

1. Parallel execution requires release 2.5.6 or later
2. API execution requires 2.5.6 or later
3. skipping data and config sync options requires 2.5.6 or later
4. Free Disk space on the Target DR cluster equal to the size of the test data set
5. Superna Eyeglass Release latest release
6. DR test policy with daily sync interval, to select DR test data and maintain a staggered copy of dr data
7. Supported PowerScale OneFS Release as per the Superna Eyeglass Feature Release Compatibility matrix found in the Release Notes.
8. SyncIQ policies used on the same cluster should use the SSIP or a pool IP address vs localhost to allow Eyeglass to manage the job correctly.

© Superna LLC

4.5.4. How to Configure DR Test Mode

Best Practice

1. ***For large quantities of data, it is best to break up the tests into smaller tests to avoid long sync times. This can be done with multiple access zones***
2. Always select data under DR path used by production policies to avoid overlapping schedules, this also also DR test mode policies to run more often and keep the DR test mode copy in sync.
3. Only Execute one DR Test mode job at a time and wait for it to complete
4. Use all nodes in the cluster when creating the SyncIQ policy and use Subnet service IP for the target host property of the policy.

Important! *This solution requires enough disk space on the target cluster to store the DR copy of the data, and enough space to contain the "Test data set" which can be all the DR data or a subset based on the policy path created.*

If the cluster has an PowerScale dedupe license key, we recommend adding DR folder and the DR test folder to the dedupe paths configuration. This will reclaim nearly 100% of the disk space used to create the 3rd copy on the DR cluster, since the data is 100% duplication of the source path.

In the example below /ifs/data/userdata is the target folder for the production cluster to copy DR data with SyncIQ. This data is copied

into /ifs/data/dr-testing and configuring the dedupe policy and running it will be able to reclaim 100% of the disk usage for DR testing.

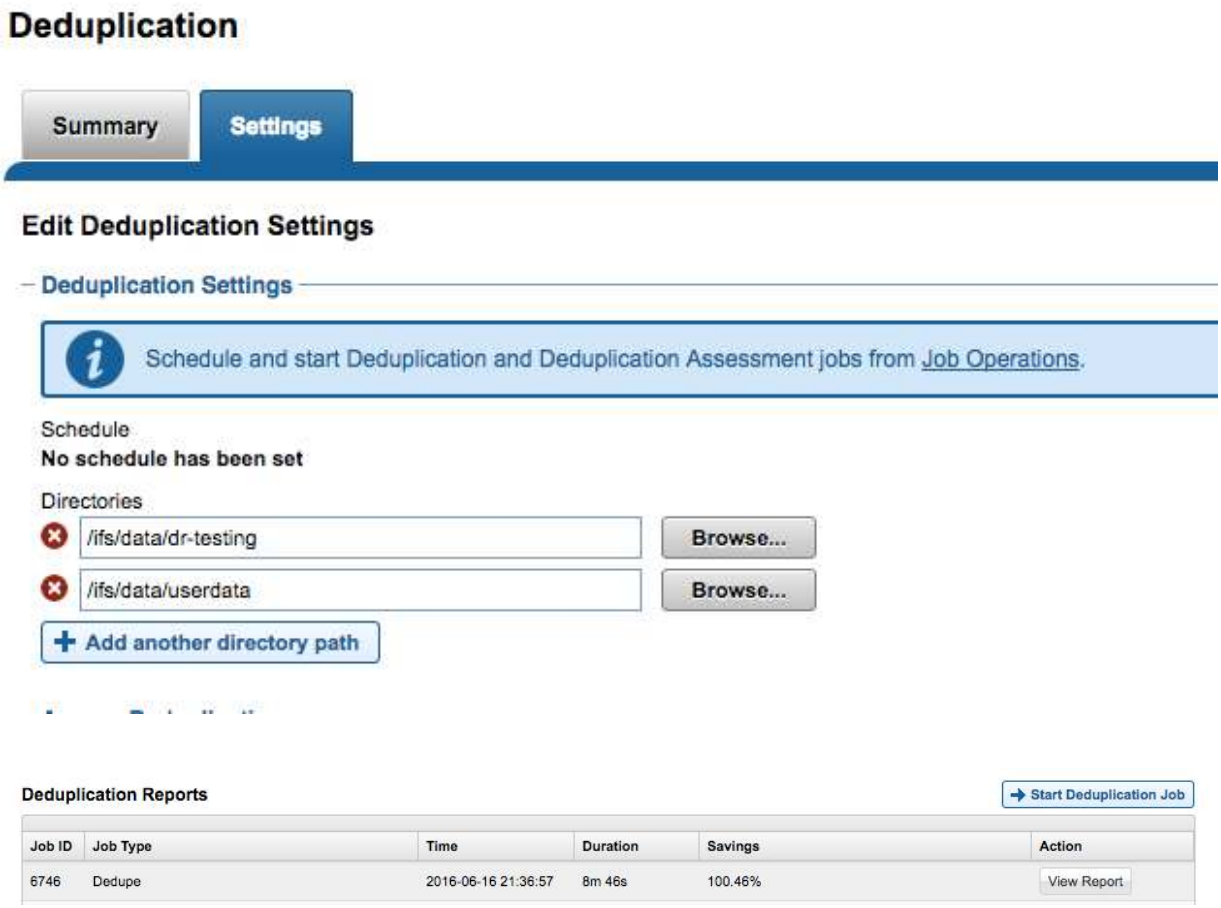

1. Production to DR clusters SyncIQ synchronizes data as follows:
2. Example: SyncIQ policy on Prod /ifs/data/corpdata.
3. Destination path DR /ifs/data/dr/corpdata.
4. Eyeglass Sync's configuration data between Prod and DR clusters (shares, exports, aliases).
5. **Note: Quotas are not used with DR test mode since the SyncIQ policies are one way failover. Quotas will not exist on the DR clusters read-only path and no quotas will be detected in DR test mode. If quotas are required for DR testing, they will need to be created manually.**

3. Create an Access Zone that will be used to test access to production data. Suggested base path "**/ifs/data/dr-testing"** (example only) on the DR cluster. The name of the Access Zone can be any name but it should have a good name like this "**DR-Testing-Zone**".
4. The Access Zone will need an IP pool associated to a IP pool and SmartConnect zone name and should be setup for DR testing applications on a dedicated network.
5. The Access Zone MUST be different than the Access Zone used for production data.
6. The Access Zone MUST use the same Authentication Provider as the production Access Zone.
7. On the DR Cluster a SyncIQ policy is created using prefix of "**Eyeglass-DR-Testing**" and uses the DR copy of the data as a source path.
    a. Note: name can include hyphen and other text to add a description.
8. SyncIQ Job for DR Testing (see advanced section for multi policy options)
    a. SyncIQ policy source path for DR cluster "/ifs/data/dr/corpdata".
    b. Destination "/ifs/data/dr-testing". This path must match the Access Zone base path created above. A sub folder can be used as well, the key requirement is the target path **MUST** fall under the target Access Zone base path that will be used for DR Testing.
    c. **Note: If you change the target path after the DR Testing Zone Access Zone exists, Eyeglass Zone Replication will**

not be able to update the path for this existing DR Testing Access Zone as per PowerScale default behaviour.

d. Destination cluster is the DR cluster (same cluster policy):

e. Use an IP address of a subnet service IP on the DR cluster or a pool ip address, leave the default option to use all nodes in the cluster to speed up the copy and sync process. Do not use a value of local host, since Eyeglass needs to validate the target cluster is managed by Eyeglass.

f. Manually run the policy after creating it (first copy can take a long time to run depending on the amount of data and will consume cluster resources, schedule this in off peak times).

g. Setup a schedule on this policy to an interval that matches your testing requirements to maintain sync with production data. NOTE: information below about overlapping policies.

h. IMPORTANT: For case where Eyeglass-DR -Testing SyncIQ Policy Source Path is the same as the production SyncIQ Policy Target Path:

i. Stagger the schedule such that the Eyeglass-DR-Testing policy Jobs do NOT start while the production SyncIQ policy is in a running state. The Eyeglass-DR-Testing policy, which starts running at the same time as the production SyncIQ policy is already running, will result in Sync Job failure. This means DR sync policy schedule should be at least 12 hours or 24 hours to maintain a near copy of production data. NOTE: running DR test mode policy on a short schedule using none overlapping

schedules configuration is the best option to avoid long sync times when activating DR test mode. In addition, DR test mode uses failover timeout value for all steps which defaults to 45 minutes. If you are not following best practices above ,this value will need to be changes. See the CLI guide in the admin manual for the CLI command to change the failover timeout value.

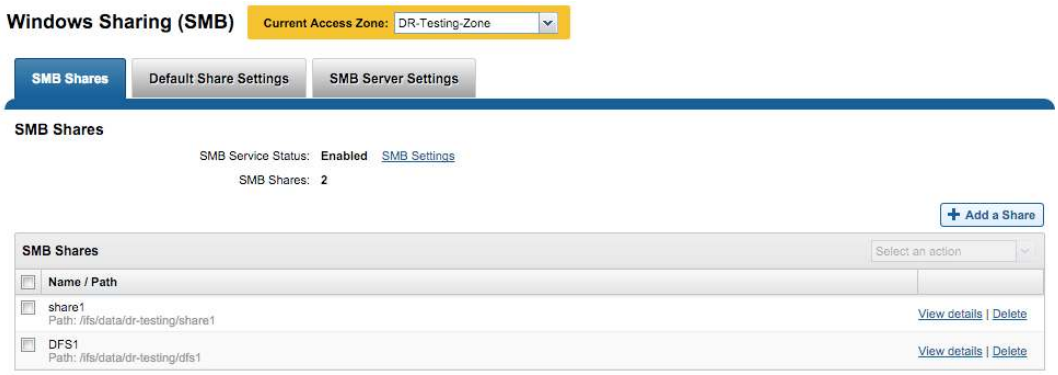

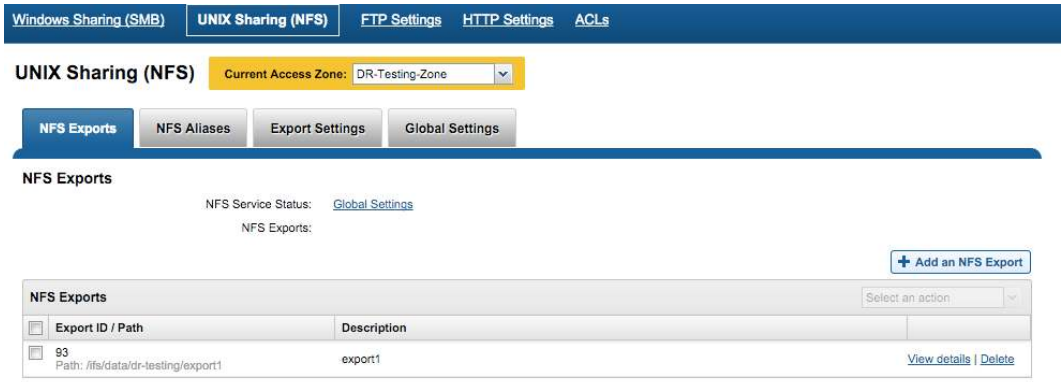

i. This is not an issue if the Eyeglass-DR-Testing policy is a sub-folder of the production SyncIQ Policy Target Path.

ii. Example:

iii.

Production SyncIQ Policy Target Path	Eyeglass-DR-Testing SyncIQ Policy Source Path	Schedule Overlap Allowed

/ifs/data/dr/corpdata	/ifs/data/dr/corpdata	No
/ifs/data/dr/corpdata	/ifs/data/dr/corpdata/app1	Yes

iv. **Note: Make sure to run the policy once, after creating it, to ensure the target path folder structure has been created.**

9. Check the Eyeglass Jobs Window after config sync has run successfully to detect the new policy **AND enable** the policy since it will be in user-disabled state by default. It will appear in a new section called "Disaster Recovery Testing"

a. 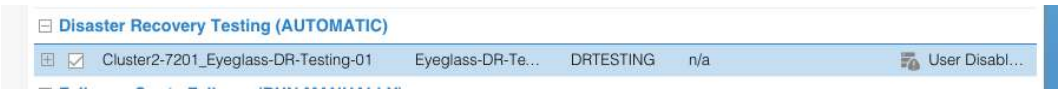

b. Run the job now by selecting and using Bulk actions run Now.

c. Use running jobs tab to verify it completed successfully.

d. Now check the cluster Access Zone SMB and export screens on OneFS to verify the configuration data from the production Access Zone has been created in the DR Testing access zone.

e. The example below shows the DR test Access Zone and path has configuration data created with the correct path for the DR testing Access Zone used in this example.

i. 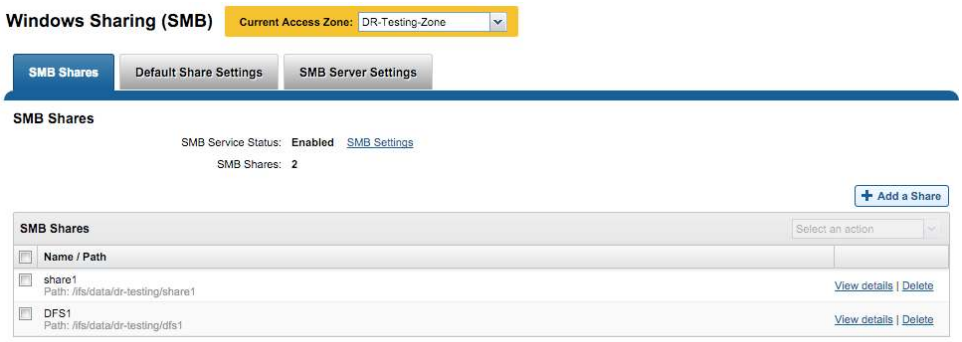

ii. 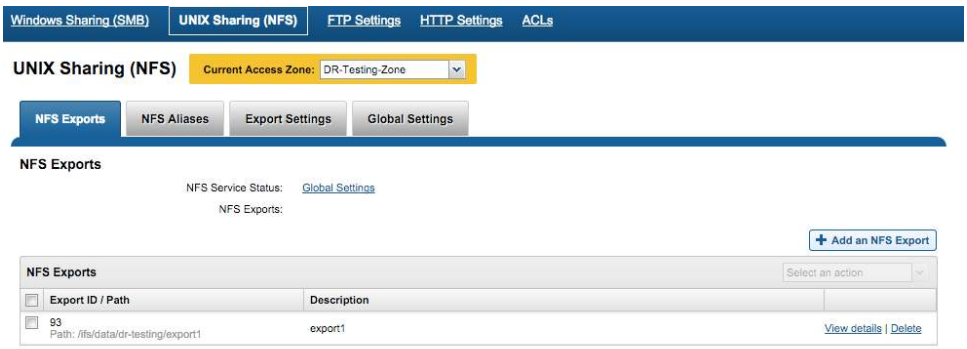

© Superna LLC

4.5.5. Isolated DR Test Mode High level Guide

- Overview
  - Prerequisites
- LiveOps DR Test mode - Isolated Network
  - Configuration for Isolated Network

Overview

This guide is not intended to be a step by step guide on PowerScale configuration. Experience on PowerScale networking is required to fully configure an isolated network bubble with firewalls, vmware hosts, AD and DNS. This guide provides specific steps about how to use PowerScale 8.2 and later feature for 2nd connection to the same AD domain to create an Isolated test environment that can use the Live OPS DR test mode feature to recreate a production clone of data, shares, exports, AD, hosts for testing DR scenario's without impact to production.

This guide only covers the AD provider and groupnet steps and high level steps needed in VMware for AD and DNS only. Consult with subject matter experts within your organization for assistance. **NOTE: Support is not able to assist with external device configuration within your infrastructure**

Prerequisites

1. Onefs 8.2
2. Eyeglass any supported version

LiveOps DR Test mode - Isolated Network

LiveOps DR Test mode can be configured on an isolated network with a non-production AD and DNS (create clone from production AD and DNS).
As the cloned AD will have the same domain as the production AD, we need to add this cloned AD to DR PowerScale Cluster on an isolated groupnet with new parameters supported in OneFS 8.2 or newer for multi-instance Active Directory: new instance provider name and new machine account.

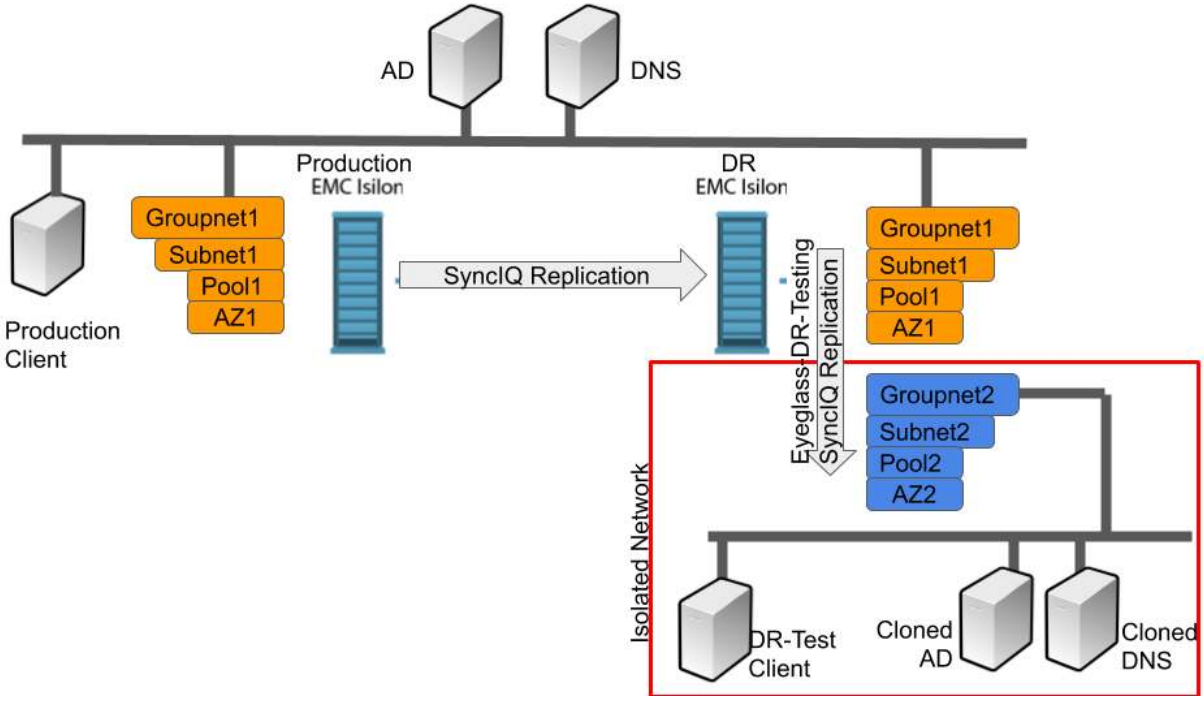

Configuration for Isolated Network

1. Ensure that DR PowerScale cluster has network interface that connect to the isolated network subnet
2. Create VM Clone of Production AD and DNS VMs (use vSphere VM Clone)
3. Update cloned AD and DNS:
    4. IP Address according the the IP range of the isolated network subnet
    5. Remove DR PowerScale Cluster machine account from this cloned AD (example: **isidr**)
6. Create new **Groupnet** on DR PowerScale Cluster (Example: **groupnet2**) and configure the DNS server to point to the cloned DNS IP address
7. Create new **Subnet** for this new Groupnet (Example: **subnet2**) and configure the netmask, SSIP and other network settings (VLAN) according to the isolated network subnet
8. Create new **Access Zone** in this new Groupnet (Example: **AZ2**) and set the groupnet to groupnet2.

9. Create new **IP Pool** in the new Subnet (Example: **pool2**) with the required interface and IP range in the isolated network and specify the access zone as AZ2 and configure the smartconnect zone name for this pool (Example: **dr-test.domain.com**)
10. Create a new A record for new isolated subnet SSIP IP address in Cloned DNS
11. Create new DNS delegation in Cloned DNS server for this new smartconnect zone name that point to that SSIP IP address
12. Add new AD Authentication Provider in DR PowerScale Cluster by joining the cloned AD. This requires Multi-Instance Active Directory feature in OneFS 8.2.0 or newer:
    13. Same domain name as the production domain name (Example: **domain.com**)
    14. Specify new provider instance name (Example: **AD-Cloned**)
    15. Specify new machine account (Example: **isidr-c**)
    16. Specify the groupnet as **groupnet2**
    17. User and password same as the production user and password to join AD

Example:

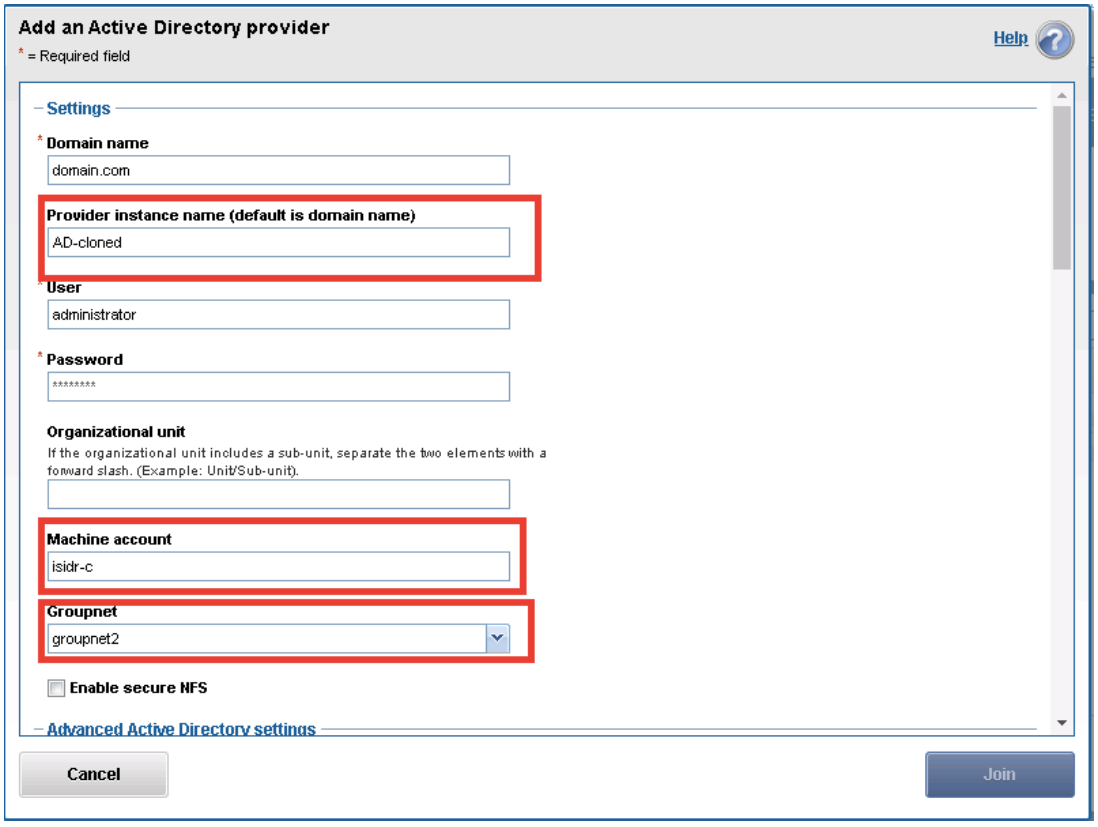

18. Verify that both production AD and cloned AD (with new provider instance name) are shown with status online.
19. Add this new AD provider (**AD-Cloned**) to the isolated access zone (**AZ2**) as the AD provider
20. Next step is to configure LiveOps DR test as per this LiveOP DR Test configuration document
21. Follow this procedure to enable DR Test Mode. Once DR Test mode is enabled, test data access from DR-Test client

© Superna LLC

4.6. How to Enable DR Test Mode

Home Top

- New Features
- How to Enable DR Test Mode
  - Open the DR assistant (< 2.5.6 releases)
  - Open DR Assistant (2.5.6 or later Releases)
- To Exit DR Test mode (All Releases)

New Features

In 2.5.6 option exists to skip config sync during enable in the GUI and the API. The api now supports enable and disable action.

Concurrent execution of 3 policies at a time has been tested and requires the configuration sync option to be disabled in the GUI or API to run concurrent DR Tests.

How to Enable DR Test Mode

Once DR Test mode is enabled from the DR Assistant icon, the Access Zone data in the DR Test Access Zone becomes writable and available for testing. The feature will make a final sync to ensure a recent copy is available for testing. The configuration is also synced to ensure it's a close mirror of the production configuration.

1. Eyeglass will detect the DR test policies based on the prefix name of “Eyeglass-DR-Testing” prefix, and place this policy on the DR Assistant DR Testing tab.
2.

3. The configuration data is also synced to the Access Zone created above during normal configuration sync jobs. See the section example in the Jobs window:
   a. 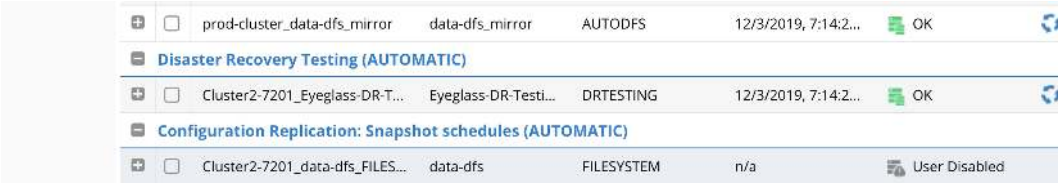

4. Configuration Synced:
   a. The configuration data is now synced from the DR cluster SyncIQ path used as the source (which is a mirror of the production Access Zone path configuration data).
   b. All configuration data that matches the SOURCE path used on the DR testing policy, is now created in the new Access Zone detected as the path of the DR Testing policy.
   c. If a subset of the DR data is needed for testing, then a subset of the data and config can be synced by building a policy with different source paths.
5. To enable write access with DR test mode read the notes below:
   a. **IMPORTANT: For the case where Eyeglass-DR -Testing SyncIQ Policy Source Path is the same as the production SyncIQ Policy Target Path:**

i. DO NOT enable DR Test Mode while the production SyncIQ policy is in a running state, OR allow this policy to run during a DR test. We recommend daily schedule or set to manual on the DR test procedure day. The Eyeglass-DR-Testing policy which starts running at the same time as the production SyncIQ policy is already running will result in Sync Job failure.

b. NOTE: < 2.5.6 releases configuration sync is started for each DR test policy execution, this cannot be disabled and a single DR test mode policy should be enabled at a time.

   i. NOTE: < 2.5.6 releases If you have more than one DR test mode policy only enable one at a time. Concurrent DR Test mode is not supported.

c. NOTE: for 2.5.6 or later releases a GUI option exists to turn off data sync and config sync steps before starting a DR test mode policy. Up to 3 policies have been tested to execute at the same time only if the config sync option is unchecked.

d. NOTE: 2.5.6 or later supports API enable and disable of DR test mode policies to be used with scripts or any REST API application. See Eyeglass API documentation here.

6. Open the DR assistant (< 2.5.6 releases)

   a. Select LiveOPS DR Testing Tab.

   b. Select the policy you want to enable for testing.

   c. Select a bulk action -> Make Target Writable.

d. Goto Running Jobs window to monitor completion of the DR test enabled job. A Job for “Enable DR test mode” is created that can be viewed from the running jobs window.

e. Once data and config is fully in Synced:

   i. DR Test mode is enabled.

   ii. DR Testing can now begin.

7. Open DR Assistant (2.5.6 or later Releases)

   a. Select LiveOPS DR Testing Tab.

   b. Select 1 or more policies (up to 3 maximum).

   c. Select Bulk Action --> Make Target Writable.

   d. NOTE: Config run option is disabled by default. Data sync can also be disabled to speed up the allow writes step.

   e. NOTE: If executing multiple concurrent DR tests on different policies config sync must be disabled, it is not supported to run concurrent DR tests without unchecking config sync option. This can also be done with the REST API from a curl script as well. See api guide.

   f.

g. 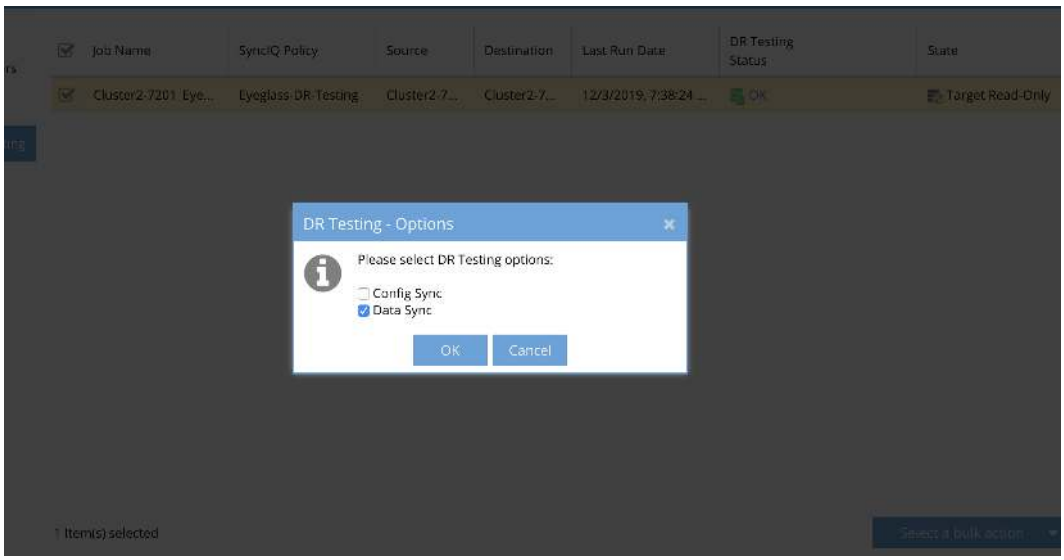

h. Monitor the DR Dashboard until it shows completed, or click running jobs to view each step.

i. 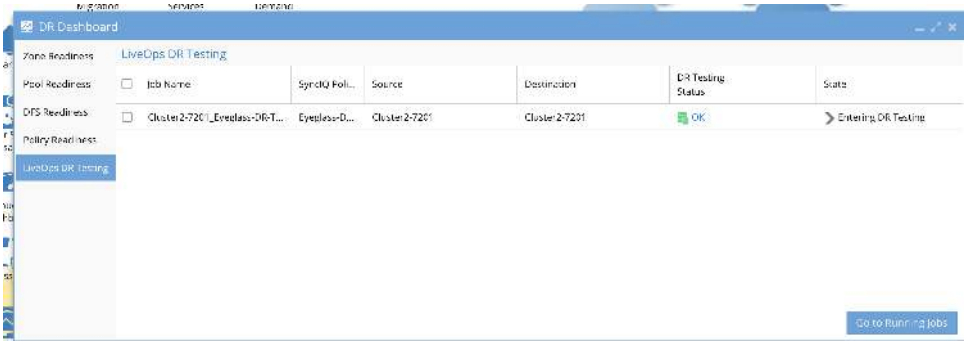

i. Once it has completed a popup window will appear show DR Test mode success message and how long it took to complete. Click on the DR Test mode job to view each step in the running jobs window.

j. 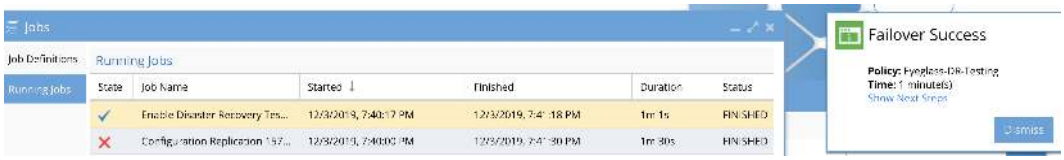

k. Once completed the DR Dashboard will show Write Enabled.

l.

m. Done

To Exit DR Test mode (All Releases)

1. See next section here

© Superna LLC

4.7. How to Disable DR Test Mode

Home Top

How to Disable DR Test Mode

1. Open DR Assistant --> Select LiveOPS DR Testing Tab.
2. Select policies that show Target Write Enabled .
3. Select the policy --> Bulk Action --> Make Target Read Only.
4. 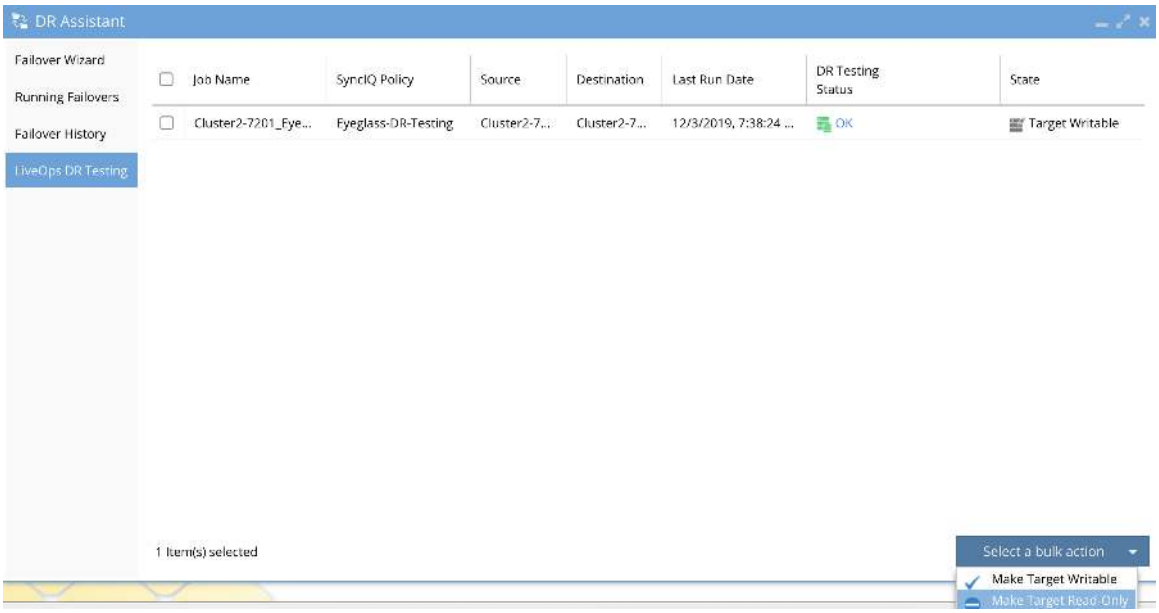

© Superna LLC

4.8. How DR Test mode Jobs are displayed in Eyeglass UI

Home Top

- Jobs UI DR Test mode:
- DR Assistant for DR Test mode
- DR Dashboard for DR Test mode:

© Superna LLC

4.8.1. Jobs UI DR Test mode:

Home Top

Jobs UI:

1. Job from production cluster is displayed as job type: “Configuration Replication” .
2. Job from DR cluster that is prefixed with “Eyeglass-DR-Testing” is displayed as job type: “Disaster Recovery Testing (AUTOMATIC)”.
3. DR type Job can be run manually or on schedule with other configuration replication jobs.
4. User can put the job in USERDISABLED or ENABLED state, while the job itself can put it in POLICY DISABLED after running Enabling DR Test.
5. 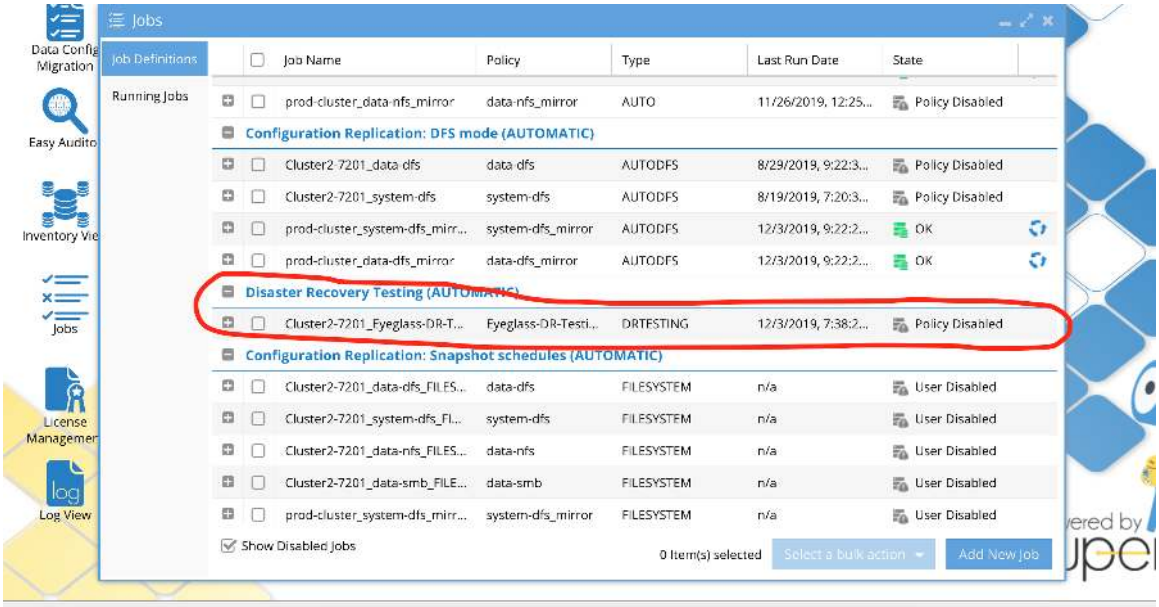

© Superna LLC

4.8.2. DR Assistant for DR Test mode

Home Top

DR Assistant:

All jobs that are prefixed with "Eyeglass-DR-Testing" are displayed in DR Assistant. This is the UI from where we can Enable/Disable DR Test Mode.

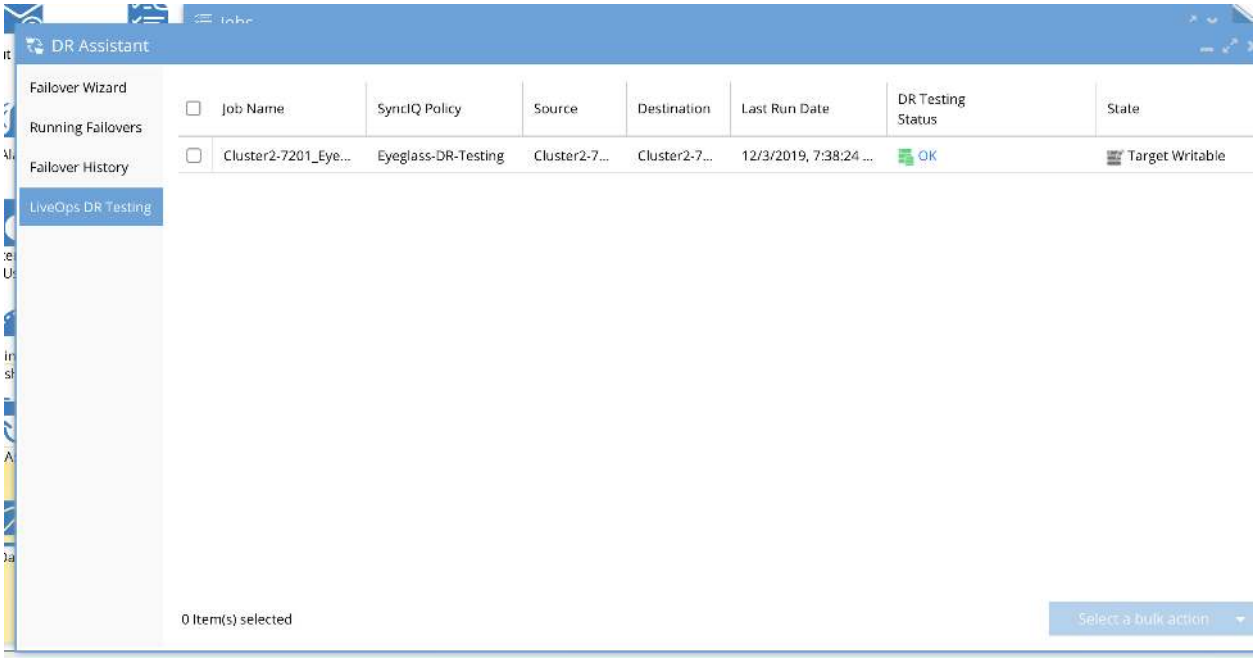

© Superna LLC

4.8.3. DR Dashboard for DR Test mode:

Home Top

DR Dashboard:

From DR Dashboard, we can see the state of DR type Jobs. Different states that the Jobs can be in are :

1. **Entering DR Testing :** When Enable DR test mode is initiated from DR Assistant and job is still running, its state is "*Entering DR Testing*".
2. **Target Writable :** When DR Testing job finishes its "*Entering DR Testing*" phase, it will end in the "*Target Writable*" status.
3. **Exiting DR Testing :** When Disable DR test mode is initiated from DR Assistant and job is still running, its state is "*Exiting DR Testing*".
4. **Target Read-Only :** When DR Testing job finishes its "Exiting DR Testing" phase, it will end in "*Target Read-Only*" state.

Note : "Entering DR Testing" and "**Exiting DR Testing**" **are transitory states**, they exist only during the execution phase of a DR Testing job. **The stable states of the DR Testing** jobs are:

- "**Target Writable**", which involves disabling the policy corresponding to the current job and allowing writes on target destination.
- "**Target Read-Only**", where the policy attached to the current job is re-enabled and writes on the target destination are disallowed.

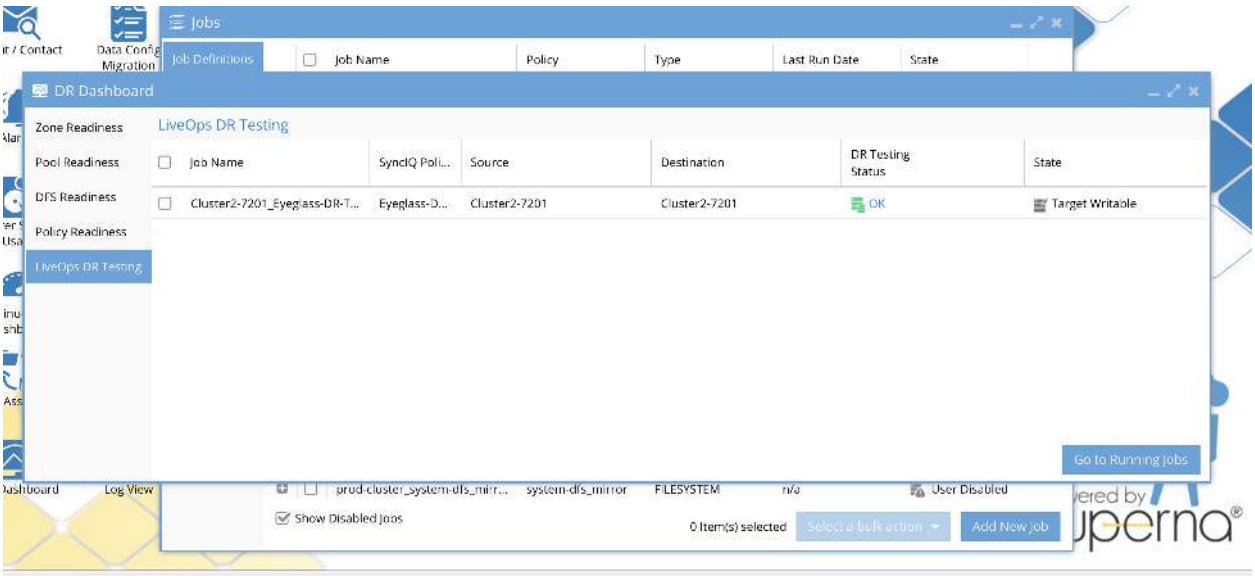
Jobs
Job Definitions
Job Name
Policy
Type
Last Run Date
State
DR Dashboard
Zone Readiness
Pool Readiness
DFS Readiness
Policy Readiness
LiveOps DR Testing
LiveOps DR Testing
Job Name
SyncIQ Poli...
Source
Destination
DR Testing Status
State
Cluster2-7201_Eyeglass-DR-T...
Eyeglass-D...
Cluster2-7201
Cluster2-7201
OK
Target Writable
Go to Running Jobs
Show Disabled Jobs
0 Item(s) selected
Add New Job

© Superna LLC

4.9. Advanced DR Test mode Configurations

Advanced DR Test mode Configurations

In order to control which applications are tested, multiple policies can be configured, selecting different source paths and directing to different target paths.

- Reasons for multiple DR test mode policies:
- Procedures to DR Test different data sets independently
- DR Test Mode States

© Superna LLC

4.9.1. Reasons for multiple DR test mode policies:

Reasons for multiple DR test mode policies:

1. Too much data to make a full copy, multiple policies allows targeting shares or exports or both for application specific testing.
2. Different groups testing different data need to execute DR testing at different times.
3. Application upgrade testing only requires a subset of the overall DR data to test.

© Superna LLC

4.9.2. Procedures to DR Test different data sets independently

Home Top

- Copy into 3 separate DR testing Access Zones with 3 separate policies
- Copy data set into 1 DR Testing Access Zone From 3 separate policies

Copy into 3 separate DR testing Access Zones with 3 separate policies

1. Create a policy with the required policy prefix name and name for application test scenario (example add -name of app or -shares).
2. NOTE: ensure the path includes shares and exports required for testing):
3. The target path can be any path depth below the base Access Zone path and allows moving the data to a different path than exists in production (example /ifs/data/dr-testing/applications/shares/application1).

Note: Each policy target path must be a different Access Zone.

3. Run policies to sync data, and set sync schedule as normal on SyncIQ policies to match your sync requirements, and to maintain a full copy .
4. Enable the DR Test mode policies in the Eyeglass jobs window after they have been discovered.
5. Run the job with bulk actions run now option.

6. Verify the configuration data is created in the DR Test Access Zone.
7. Open DR Assistant, DR Testing tab.
8. Select one or more policies to enable for DR Test mode in a writeable file system.

DR Testing

Job Name	SyncIQ Pol...	Source	Destination	State
Cluster2-7201_Eyeglass-DR-Testing-expo...	Eyeglass-D...	Cluster2-7201	Cluster2-7201	Target Writable
Cluster2-7201_Eyeglass-DR-Testing-shares	Eyeglass-D...	Cluster2-7201	Cluster2-7201	Target Writable

Copy data set into 1 DR Testing Access Zone From 3 separate policies

4. Create 3 separate policies with the required policy prefix name and a name for application test scenario. For example: Eyeglass-DR-Testing-1, Eyeglass-DR-Testing-2 and Eyeglass-DR-Testing-3.
5. Create each policy to select a source path that matches the target application within its own Access Zone base path. Each policy source path should be within an individual Access Zone. (**Note:** ensure the path includes shares and exports required for testing).
6. Create target path inside the target test Access Zone created for DR Testing.
7. The target path can be any path depth below the base Access Zone path and allows moving the data to a different path than exists in production (example /ifs/data/dr-testing/applications/shares/application1).

   a. Note: Each policy target path must be the same Access Zone.

8. Run policies to sync data and set sync schedule as normal on SyncIQ policies to match your sync requirements and to maintain a full copy up to date.
9. Enable the DR Test mode policies in the Eyeglass jobs window after they have been discovered.
10. Run the job with bulk actions Run Now option.
11. Verify the configuration data is created in the DR test Access Zone.
12. Open DR Assistant, DR Testing tab.
13. Select one or more policies to enable for DR Test mode in a writeable file system.

© Superna LLC

4.9.3. DR Test Mode States

DR Test Mode States

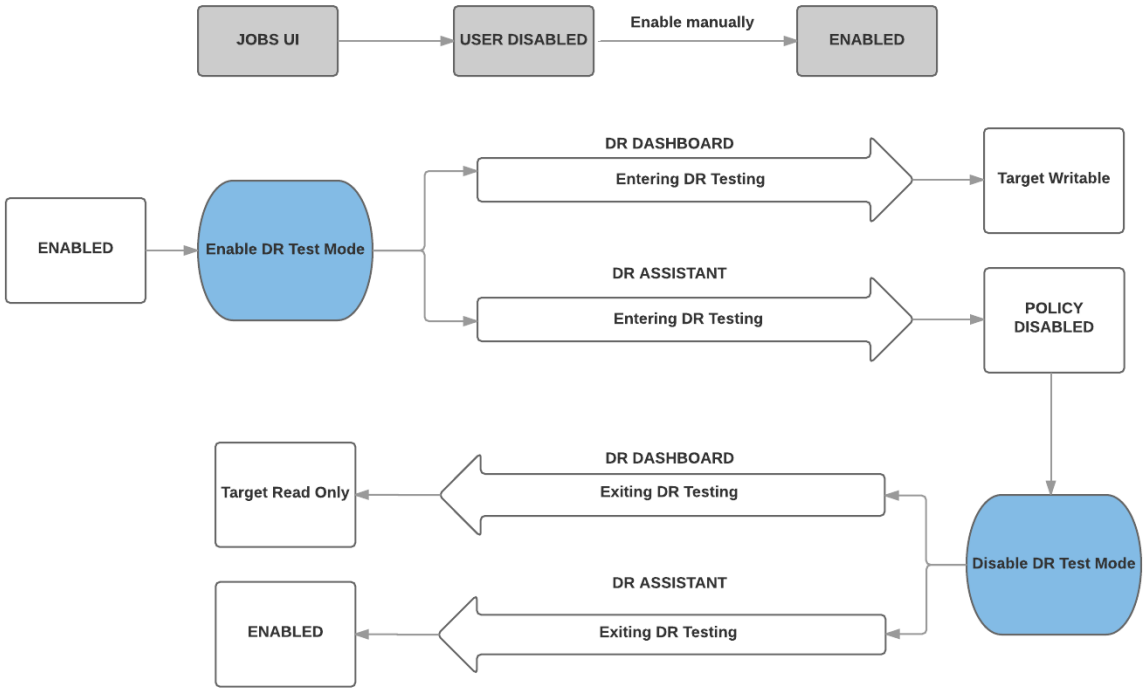

Copyright Superna LLC 2017

© Superna LLC

5. Eyeglass Clustered Agent Admin Guide

Home Top

- What's New
- ECA Cluster Topology Deployment Options
- High Availability and Resilience
- Eyeglass Active Directory User SID resolution considerations
- How to monitor remote ECA clusters from Eyeglass
- ECA Cluster Monitoring Tools
- ECA Cluster Operational Procedures
- ECA CLI Command Guide
- ECA Cluster Disaster Recovery and Failover Supported Configurations
- How to Change Performance with VMware
- Dark Sites - How to Health check Eyeglass and ECA clusters when opening a support case

© Superna LLC

5.1. What's New

Home Top

What's New

This covers Ransomware Defender and Easy Auditor ECA deployment configurations and operational maintenance.

2.5.7

1. opensuse 15.2
2. automatic firewall
3. authenticated management tools

2.5.5

1. Disk space monitoring
2. dual OS and data disks
3. Enhanced alarm monitoring
4. opensuse 15.1
5. Mini ECA for distributed deployment options
6. cluster up checks NFS mounts
7. AutoNFS mount feature for centralized mount control

2.5.3

1. **HA Turbo Audit ingestion allows ECA node to check point location in audit file and failover processing to other nodes automatically and resume processing at the same location in the audit log**

2. **Load Balancing on ECA node rejoining the cluster will allocate log file processing evenly across the cluster**
3. **Ingestion of existing archived compressed PowerScale audit logs**
    a. allows missing data to be re-ingested
    b. allows ingestion of old data prior to Easy Auditor installation
4. Master processes are now moved to nodes 4-6 and ability to start more cluster wide services for processing audit triggers or ransomware analysis for high through put environments

© Superna LLC

5.2. ECA Cluster Topology Deployment Options

Eyeglass Clustered Agent Deployment Options

It's best practice to place ECA clusters near PowerScale clusters to reduce latency between the cluster and the ECA instance that is processing the audit log.

The ECA is cluster can be deployed centrally in 2 configurations.

1. Centralized with NFS over the WAN - in this configuration a central cluster has access over a WAN link to mount the remote cluster audit data over NFS. This is recommended in metro WAN deployments
2. Centralized with remote mini ECA cluster - This option is longer latency links between clusters with centralized processing and uses a single VM or 2 for HA at remote sites to collect audit data over a local NFS mount and forwards events to the centralized ECA cluster for processing. This is recommended if the WAN is slow or latency is high > 10 MS RTT.

Supported Topology #1

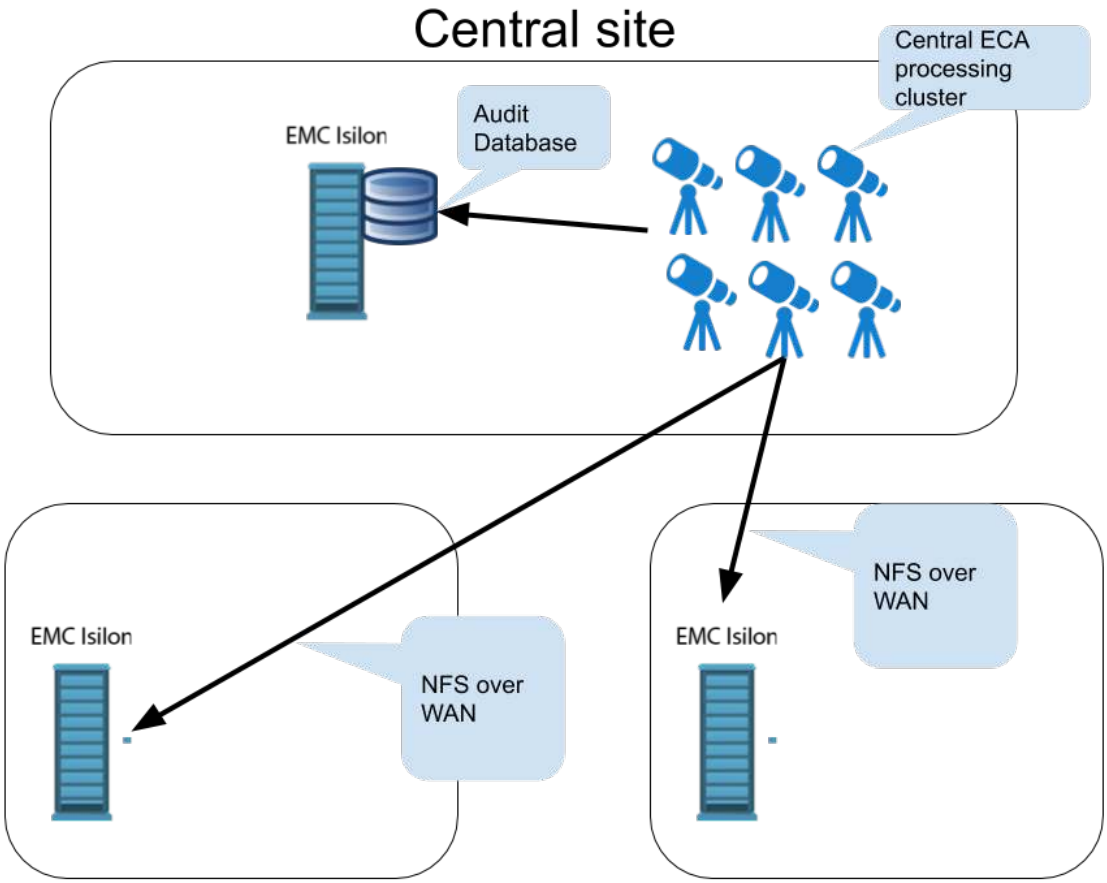

Supported Topology #2

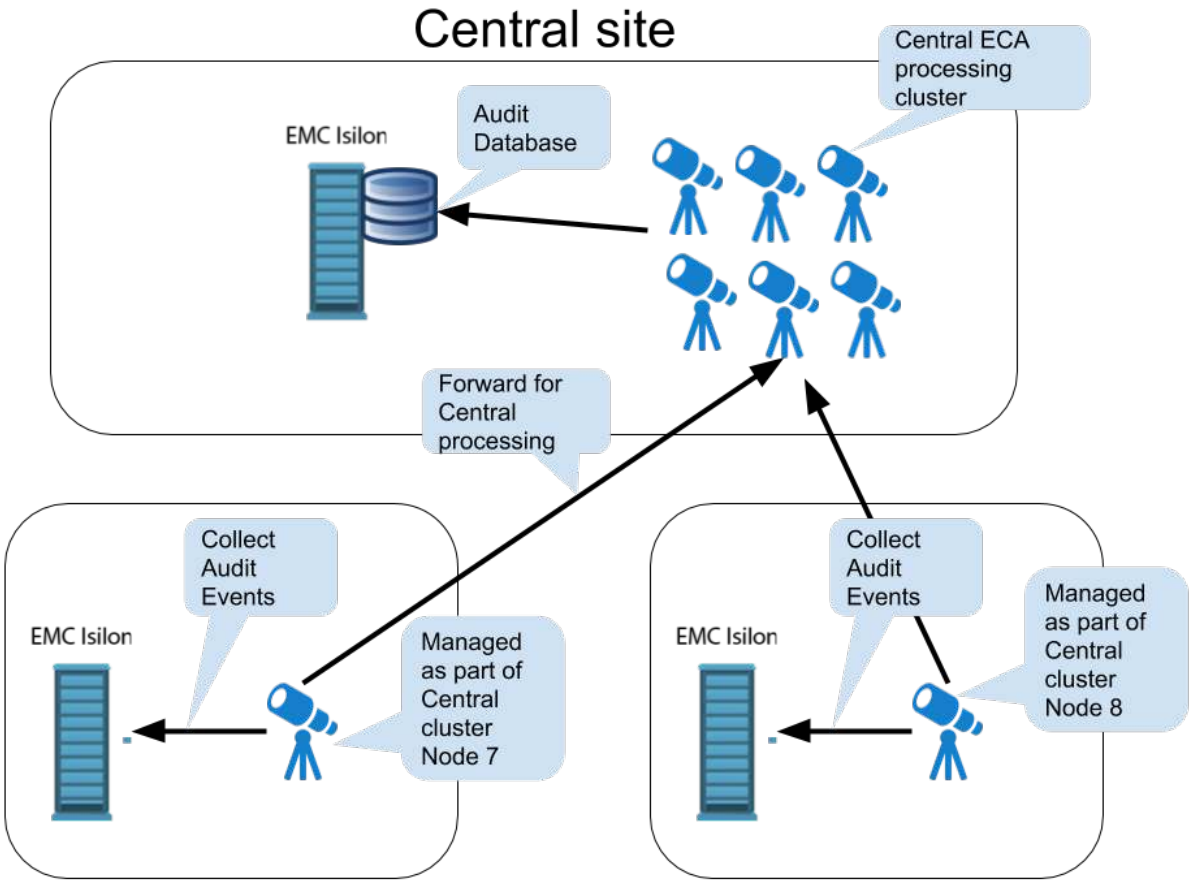
Central site
Central ECA processing cluster
EMC Isilon
Audit Database
Forward for Central processing
Collect Audit Events
EMC Isilon
Managed as part of Central cluster Node 7
Collect Audit Events
EMC Isilon
Managed as part of Central cluster Node 8

© Superna LLC

5.3. High Availability and Resilience

High Availability and Resilience

The ECA cluster is an active active design that offers Matrix Processing of events. This design uses dedicated docker containers that perform a specific function on each node. The solution allows for multiple container failures within a node and between nodes.

The solution allows a distribution of event processing at the functional container level on any of the nodes in the cluster. This allows greater than a single point of failover within a node and between nodes. This ensures processing contains under most common conditions with greater than 2x HA level of redundancy.

Cluster Operational Requirements

The platform is a robust high performance event processing cluster for threat and audit detection capabilities. The cluster will remain operational as long as 2 of the 3 nodes are running and can reach the HDFS cluster database.

Architectural Data flow of audit events through the Eyeglass Clustered Agent

How the ECA processes incoming events, should be understood when debugging

1. The ECA cluster is an a**ctive active active** solution which means all nodes process and analyze audit data from the cluster.

2. The cluster load balances audit messages to each node in the cluster.
3. Each user in AD hashed and assigned to one node in the cluster so single user behavior patterns can be processed by a single node in the cluster.
4. If a node goes down another node takes over the active directory user processing for the failed node .

© Superna LLC

5.4. Eyeglass Active Directory User SID resolution considerations

Active Directory Planning

Your SID to friendly name resolution uses PowerScale Authentication providers to resolve a SID for all products that use the ECA.

Eyeglass User Lockout Active Directory Planning Ransomare Defender

The lockout process identifies all shares the user has access permissions based on searching all shares in all access zones on all clusters managed by Eyeglass. This list of shares will have a real-time deny permission added to the share for the affected user.

A special case is handled for the “Everyone” well known group which should be understood how it operates in multi-domain Active Directory configurations.

Two scenarios can exist with AD domains on PowerScale clusters.

Scenario #1:

- The first is parent and child AD domains that are members of the same forest and a trust relationship exists.

Scenario #2:

- The second scenario covers two domains that are not members of the same forest and no trust relationship exists between the domains

The “Everyone” well known group if applied to a share in each scenario is shown below and a lockout permission applied regardless of which domain the user is located. This is required since Eyeglass has no way to know if the domains trust each other or not. This solution ensures all everyone shares are locked out, which is more secure than skipping some shares.

Reference the diagram below.

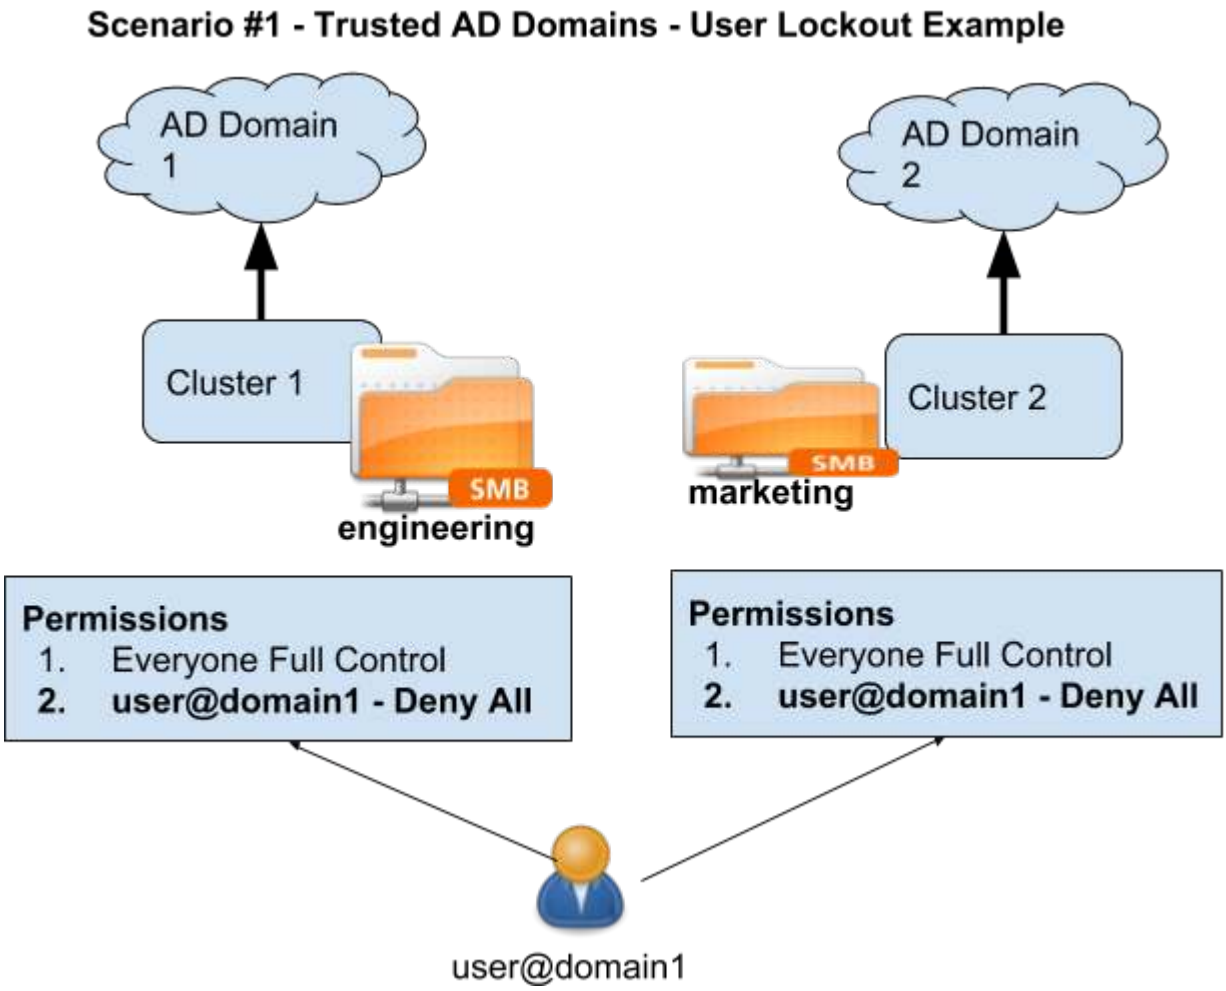

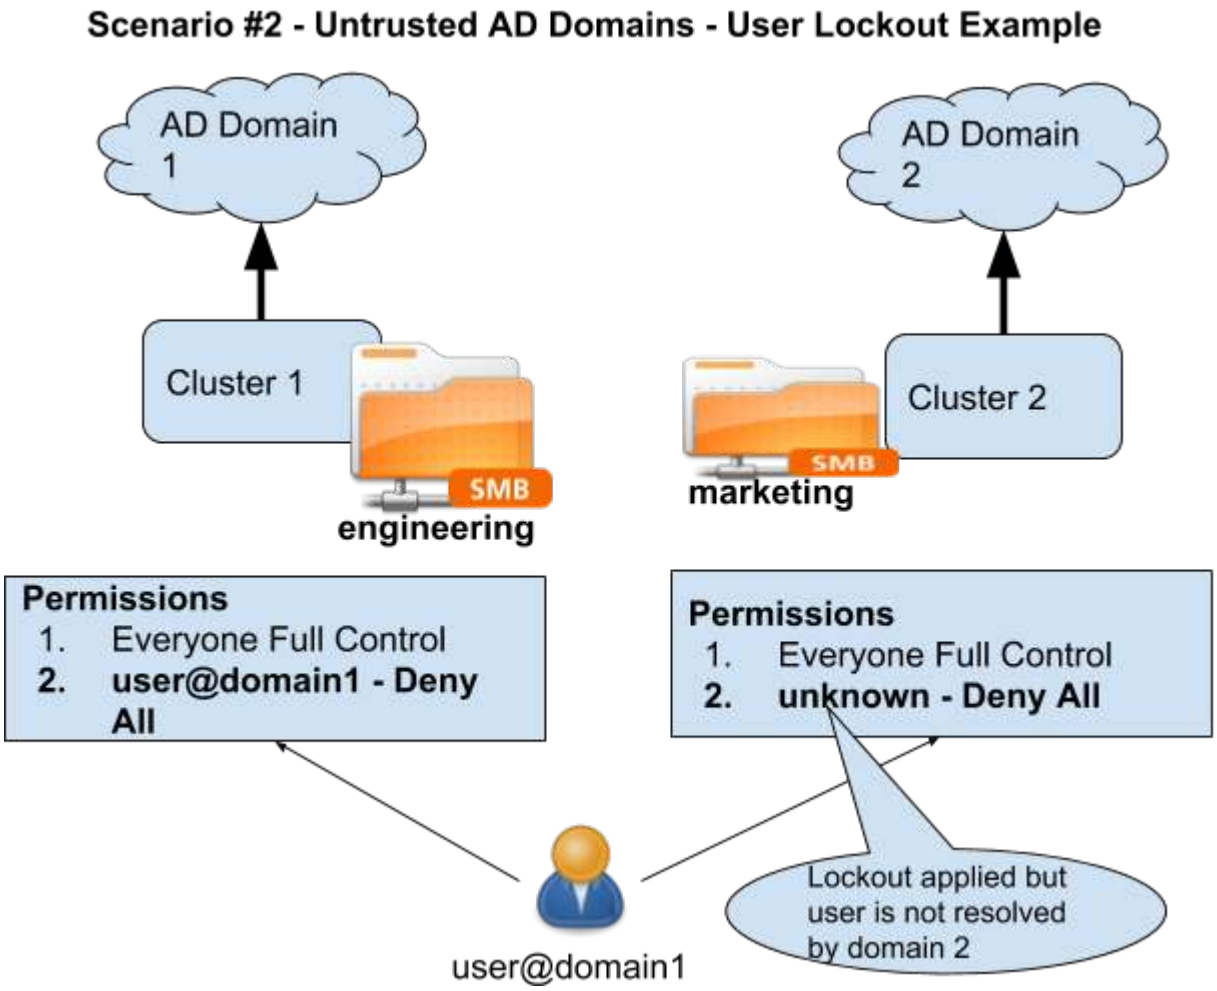
Scenario #2 - Untrusted AD Domains - User Lockout Example
AD Domain 1
AD Domain 2
Cluster 1
Cluster 2
SMB
engineering
SMB
marketing
Permissions
1. Everyone Full Control
2. user@domain1 - Deny All
Permissions
1. Everyone Full Control
2. unknown - Deny All
Lockout applied but user is not resolved by domain 2
user@domain1

© Superna LLC

5.5. How to monitor remote ECA clusters from Eyeglass

Home Top

- Remote Service Authentication and Protocol
  - Service Registration Monitoring in Eyeglass
  - Service States
  - Health States

Remote Service Authentication and Protocol

Eyeglass can communicate with multiple Ransomware Defender or Easy Auditor endpoints. Each endpoint must have a unique API token, generated by the Superna Eyeglass REST API window:

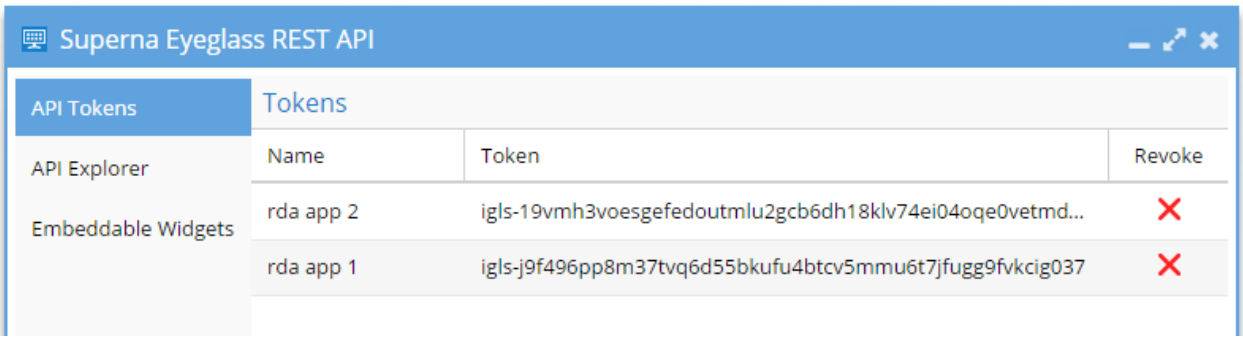

Once a token has been generated for a specific ECA, it can be used in that ECA's startup command for authentication, along with the location of Eyeglass.

Communication with the ECA is bidirectional at the start (ECA -> Eyeglass for security events). Eyeglass will query the analytics database and test database access on regular interval.

The ECA should:

1. Heartbeat
2. Notify Eyeglass of any detected threats
3. Periodically send statistics on processed events.
4. Periodically poll for updated Ransomware definitions, thresholds, and Ignore list settings.

Service Registration Monitoring in Eyeglass

Eyeglass icon "Manage Services" displays all registered ECA's and CA UIM probes operating remotely from the Eyeglass appliance. The screenshot below shows 3 ECA nodes registered and the health of each process running inside the node.

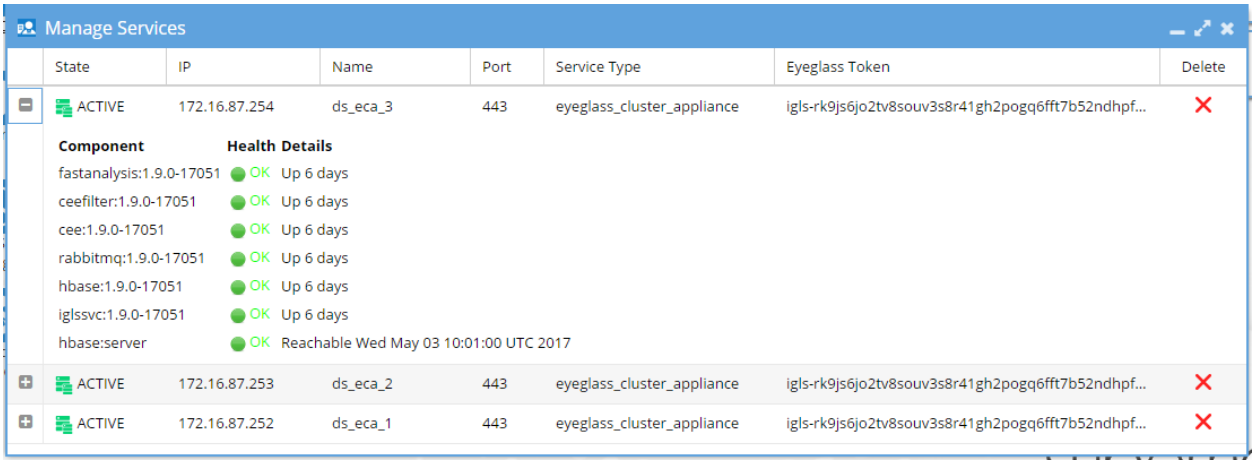

Service States

1. Active: Has checked in with heartbeat
2. In-Active: Has failed to heartbeat, no longer processing

Health States

1. Up - running and up time in days
2. Down - not running

The Delete icon per service registration should not be used unless directed by support. This will remove the registration from the remote service.

© Superna LLC

5.6. ECA Cluster Monitoring Tools

ECA Cluster Monitoring Tools

The ECA is built on docker containers with each container providing a different function to process messages. If both Ransomware And Auditor are licensed on an ECA cluster the following additional UI's exist for monitoring HDFS write performance, spark job analysis jobs.

Note x.x.x.x is any node in the cluster, the master function only runs active on one node and can move from one node to another based on voting on startup.

HBASE Database Monitoring UI

1. http://x.x.x.x/hbase-master (master controls all region servers and provides overall cluster configuration. One node is elected the master but all ECA nodes can become HBASE Master that oversees each Region server.
2. http://x.x.x.x/hbase-rs (Region server controls a portion of the audit database and controls writes to this portion of the database)

Spark Monitoring UI

1. http://x.x.x.x/spark-master - Spark master shows all running queries across the cluster, each node runs
2. http://x.x.x.x:spark-worker - Spark worker shows running queries on each node in the cluster and the number of tasks to complete along with job logs and errors
3. http://x.x.x.x/spark-history - Spark History server provides interface to see all searches done across the cluster and all logs related to these searches for debugging. The searches are encoded by ID visible in Eyeglass reports.

Spark History Server Report logs Monitoring UI

1. http://x.x.x.x/spark-history - Spark queries have logs stored on the PowerScale for each query for diagnostics and reviewing post query execution. This UI can be used to download query logs based on the Auditor report driver ID number. This will also be available for download in the Eyeglass report table as well.

Kafka Monitoring UI

2. http://x.x.x.x:kafkahq/ - Kafka monitoring of message event processing, queues depths, event rates. On ECA node one only open ui click add cluster enter **zookeeper:2181** for the cluster to have auto detection find the correct Kafka broker to add. Note only 1 UI is required to monitor all brokers. This process handles events for fast analysis and writing to the Analytics DB

© Superna LLC

5.7. ECA Cluster Operational Procedures

Home Top

- Eyeglass Cluster Maintenance Operations
  - Cluster OS shutdown or restart
  - Cluster Startup
- ECA Cluster Node IP address Change
- Change ECA Management tool Authentication password
- Single ECA Node Restart or Host crash Affect 1 or more ECA nodes
- Eyeglass ECA Cluster Monitoring Operations
  - Checking ECA database Status:
  - Check overall Cluster Status
  - Check Container stats memory, cpu on an ECA node

Eyeglass Cluster Maintenance Operations

Note: Restart of the OS will not auto start up the cluster post boot. Follow steps in this section for cluster OS shutdown, restart and boot process.

Cluster OS shutdown or restart

1. To correctly shutdown the cluster
2. Login as ecaadmin via ssh on the master node (Node 1)
3. ecactl cluster down (wait until all nodes are down)

4. Now shutdown the OS nodes from ssh login to each node
5. ssh to each node
6. Type sudo -s (enter admin password)
7. Type shutdown

Cluster Startup

1. ssh to the master node (node 1)
2. Login as ecaadmin user
3. ecactl cluster up
4. Verify boot messages shows user tables exist and signal table exists (this step verifies connection to analytics database over HDFS on startup)
5. Verify cluster is up
6. ecactl cluster status (verify containers and table exist in the output)
7. Done.

ECA Cluster Node IP address Change

To correctly change the cluster node ip addresses:

1. Login as ecaadmin via ssh on the master node (Node 1)
2. ecactl cluster down (wait until completely down)
3. Sudo to root
    a. sudo -s (enter admin password)
    b. Type yast

c. Navigate to networking to change the IP address on the interface)

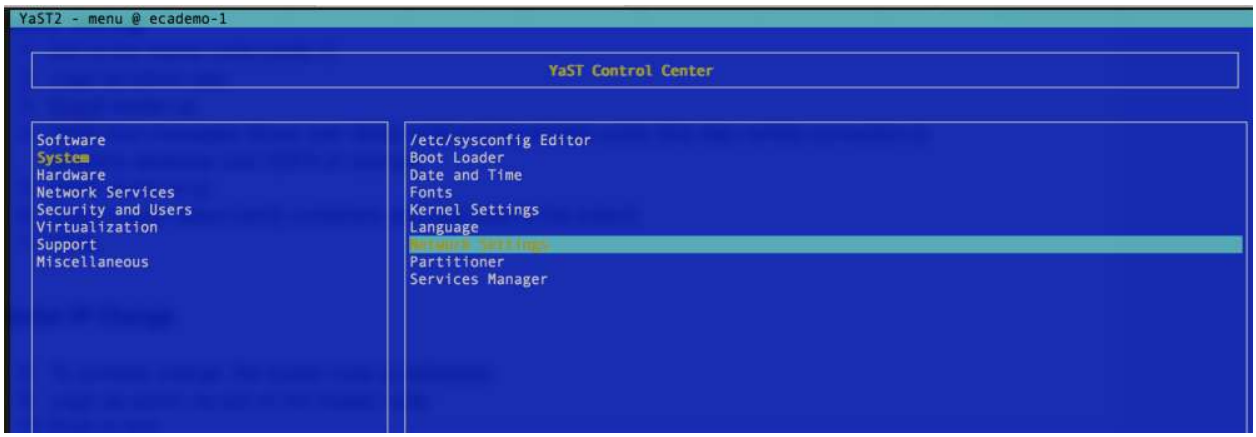

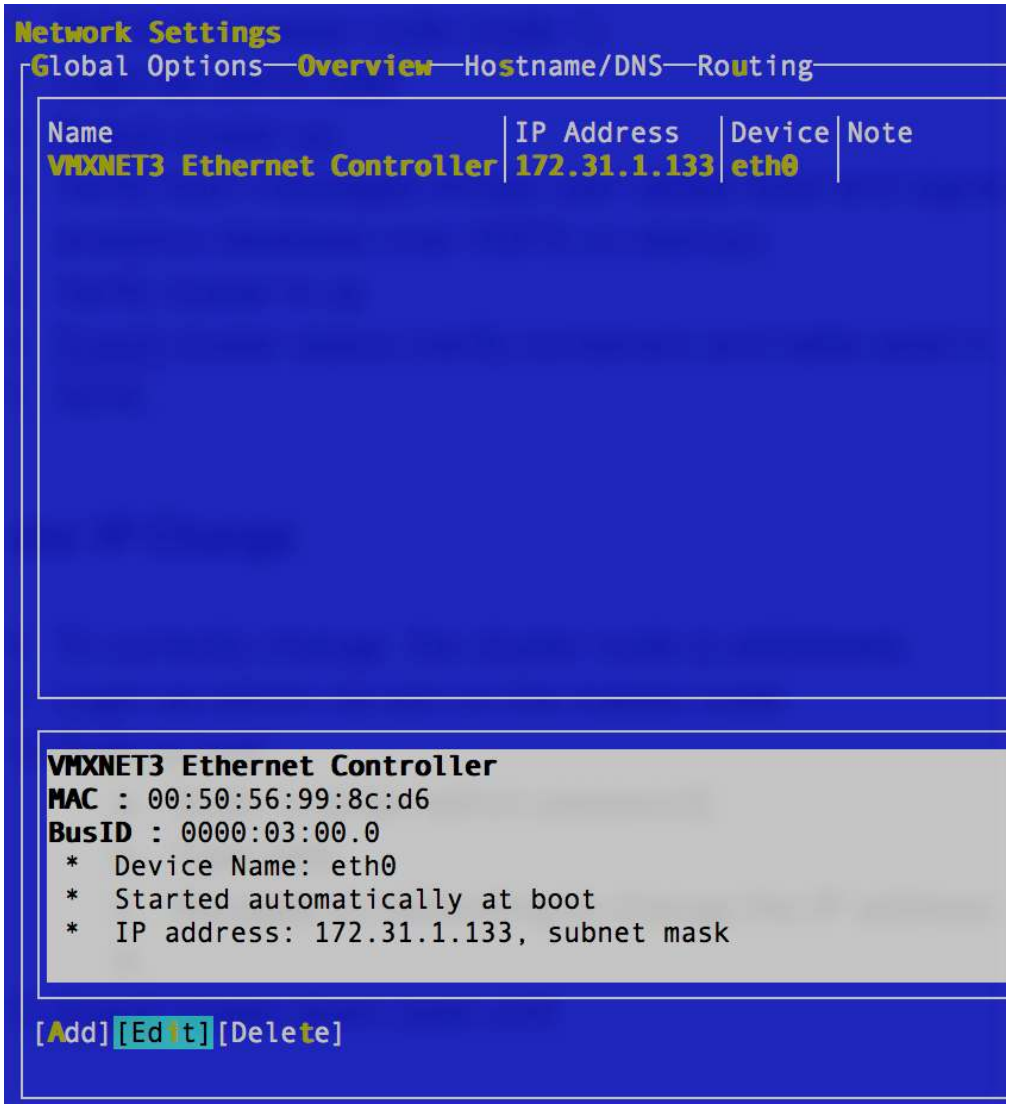

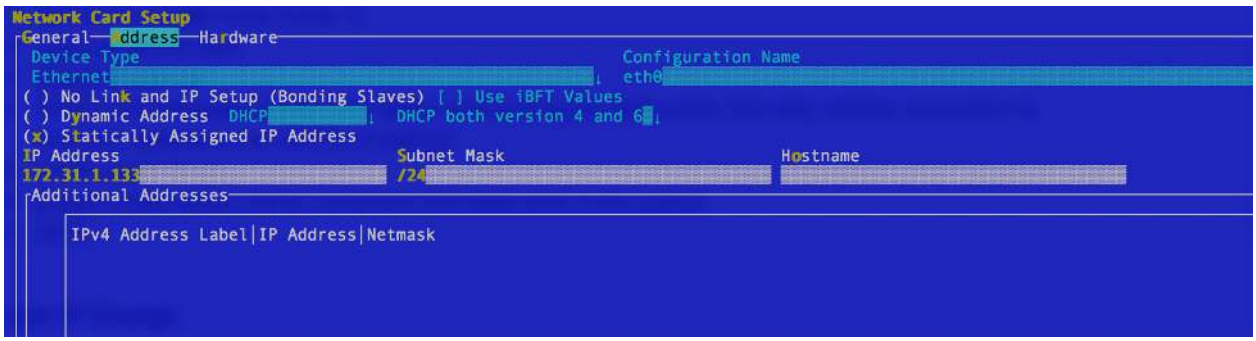

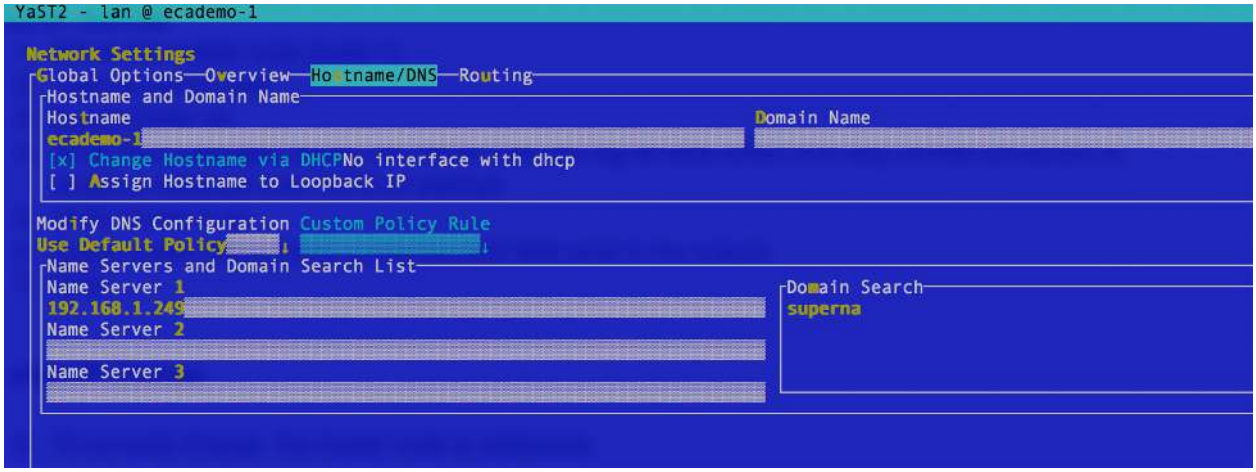

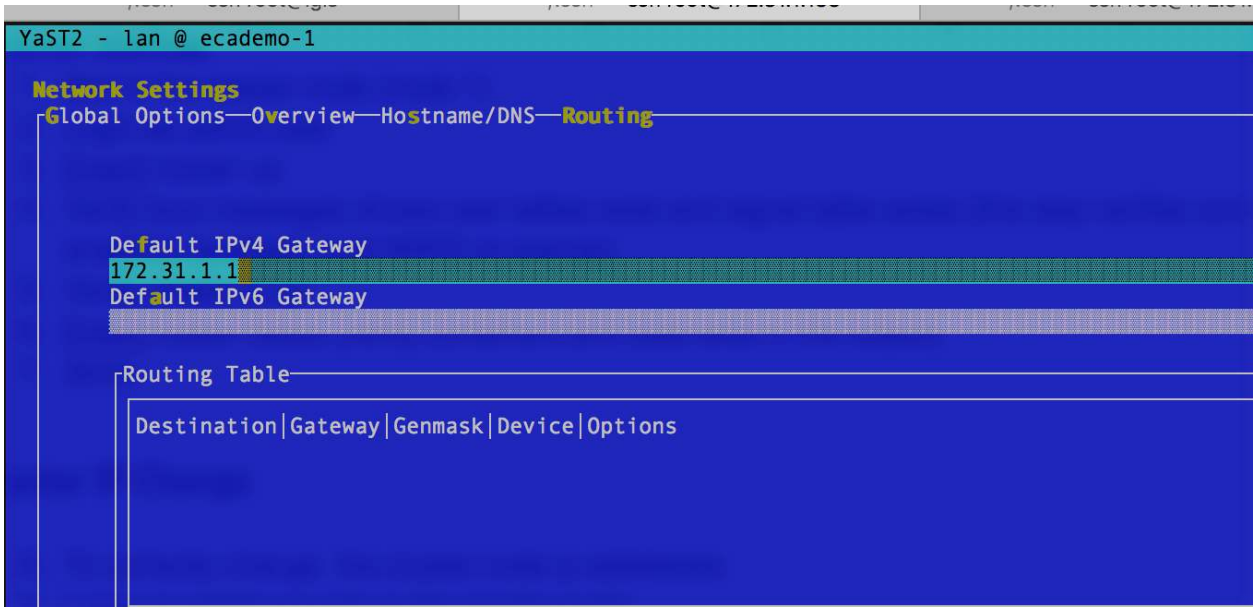

4. Each screenshot shows ip, dns, router settings

8. Save and exit yast

9. Repeat on all nodes in the cluster

4. Once completed changes verify network connectivity with ping and DNS nslookup

8. Edit with ‘ nano /opt/superna/eca/eca-env-common.conf ’ on the master node (Node 1)
9. Edit the ip addresses of each node to match new new settings
10. export ECA_LOCATION_NODE_1=x.x.x.x
11. export ECA_LOCATION_NODE_2=x.x.x.x2
12. export ECA_LOCATION_NODE_3=x.x.x.x3
13. Control X to exit and save
14. Modify the Isilon NFS mount permissions on all clusters managed by the ECA instance. Replace the IP's to include all ECA node ip addresses. Example below shows 3 IP, check your cluster node count to update the command below to match your deployment.
    - isi nfs exports modify --id 3 -f --add-root-clients="x.x.x.x, y.y.y.y, z.z.z.z"
15. Start cluster up
16. From master node (Node 1)
17. ecactl cluster up (verify boot messages look as expected)
18. Eyeglass /etc/hosts file validation
19. Once the ECA cluster is up
20. Login to Eyeglass as admin via ssh
21. Type cat /etc/hosts
22. Verify the new ip address assigned to the ECA cluster is present in the hosts file.
23. If it is not correct edit the hosts file and correct the IP addresses for each node.

24. Login to Eyeglass and open the Manage Services window. Verify active ECA nodes are detected as Active and Green.
25. You should see the old ip addresses and inactive ECA nodes with the old ip addresses , click the red X next to each to delete these entries from the managed services icon.
26. Done

Change ECA Management tool Authentication password

1. Release 2.5.7 and later now protects all management tools on the ECA cluster with a user name and password over a HTTPS login page. This includes hbase, kafka, spark UI's that are accessible from the Managed Services icon in the Eyeglass GUI.
    a. The login to this UI is ecaadmin and default password is 3y3gl4ss
2. Login to node 1 over ssh as ecaadmin user and run the command below
    a. NOTE: Replace <password> with the password
3. ecactl cluster exec "htpasswd -b /opt/superna/eca/conf/nginx/.htpasswd ecaadmin <password>"
4. done. The new password is active immediately on all nodes.

Single ECA Node Restart or Host crash Affect 1 or more ECA nodes

Use this procedure when restarting one ECA node, which under normal conditions should not be done unless directed by support. The other use case is when a host running an ECA VM is restarted for maintenance and a node will leave the cluster and needs to rejoin.

1. On the master node
2. Login via ssh as ecaadmin
3. Type command : ecactl cluster refresh (this command will re-integrate this node back into the cluster and check access to database tables on all nodes)
4. Verify output
5. Now type: ecactl db shell
6. type : status
7. Verify no dead servers are listed
8. If no dead servers
9. Login to Eyeglass GUI, check Managed Services and verify all nodes are green.
10. Cluster node integration procedure completed.

Eyeglass ECA Cluster Monitoring Operations

Checking ECA database Status:

1. ecactl db shell [enter]
2. status [enter]

Check overall Cluster Status

1. ecactl cluster status

Check Container stats memory, cpu on an ECA node

1. ecactl stats (auto refreshes)

© Superna LLC

5.8. ECA CLI Command Guide

Home Top

ECA CLI Command Guide

The following table outlines each CLI command and purpose,

For a complete list of commands see the admin guide here.

CLI Command	Function
ecactl cluster <command>	**UP** - Bring up the cluster across all nodes, **Down** - Bring down the cluster across all nodes **Status** - gets status of all processes and connection to HDFS database **Refresh** - Use this command on an ECA node when it was restarted and needs to rejoin an existing cluster.
ecactl cluster down --hard	This should only be used if the normal down process hangs or a container will not shutdown correctly. This command should only be used unless directed by support. This will shutdown the database without a clean shutdown process.
ecactl cluster down -bg	The background flag is a rapid down for faster cluster down executed in parallel for upgrades. Requires 2.5.7 or later

ecactl stats	Returns container memory and cpu statistics to view health of all containers
ecactl containers ps	List all running containers
ecactl logs --follow iglssvc (other services are rmq, fastanalysis)	Tail the Eyeglass agent on a node, used for debugging.
ecactl cluster exec <command> Example ecactl cluster exec ecactl containers ps	Use this command to execute the commands arguments across all ECA nodes Run from ECA master node Use only with single commands (not multiple concatenated commands)

© Superna LLC

5.8.1. How to run commands across all nodes

How to run commands across all nodes

For example, to get the hostname of each node, do:

- ecactl cluster exec hostname

To get the running containers on each node, do:

- ecactl cluster exec ecactl containers ps

To restart the turboaudit container on all three nodes, do:

- ecactl cluster exec "ecactl containers stop turboaudit && ecactl containers rm -r turboaudit && ecactl containers up -d turboaudit"

This will work with sudo commands too, but you'll be prompted for the `ecaadmin` password each time. For example:

- ecactl cluster exec sudo systemctl restart docker

© Superna LLC

5.9. ECA Cluster Disaster Recovery and Failover Supported Configurations

Home Top

- Overview
- Unsupported ECA HA Configurations
- Prerequisites for both Scenarios for the DR Cluster and DR Site
- Scenario #1 - Production Writeable cluster fails over to DR site ECA cluster stays at the production site (Longer RTO)
  - Post Failover Reconfiguration Steps
- Scenario #2 - Production Writeable cluster fails over to DR site and Production site ECA cluster will failover to warm standby ECA cluster at the DR Site (Lower RTO)
  - Prerequisites to Complete Before a Failover
  - Post Failover Reconfiguration Steps

Overview

This section covers how the ECA cluster can be configured for failover to another site or for migrating the ECA cluster to another cluster when the current Isilon will be decommissioned or after a failover and the failover cluster is now active. This guide applies to Ransomware Defender, Easy Auditor and Performance Auditor.

Unsupported ECA HA Configurations

1. It is not supported to have 2 different ECA clusters using the same Eyeglass VM.
2. It is not supported to stretch a ECA between 2 data centers. An ECA cluster is not designed to be stretched and handle node failures. It is designed as a local cluster with load balancing and local node failover. Warm Standby is the only supported HA solution.

Prerequisites for both Scenarios for the DR Cluster and DR Site

1. **Easy Auditor checklist** - All steps below are found in the installation guide. **The check list is a summary of key steps that must be completed using the installation guide. All Steps MUST be completed before ECA cluster can manage the target cluster.**
    a. (Mandatory) If Warm Standby ECA cluster is used it must have the same VM count as the production site ECA cluster. This option is only used when ECA cluster failover to the 2nd site is required.
    b. (Mandatory) The Easy Auditor Database has been synced to the DR site into the correct access zone - The guide here can be followed to protect the audit database with SyncIQ.
    c. (Mandatory) The IP pool is created in the Eyeglass access zone with at least 3 nodes in the pool to receive HDFS IO

from the ECA and the smartconnect name is created for HDFS access.

d. (Mandatory) Smartconnect DNS Delegation is created for HDFS access zone in your DNS infrastructure following Dell documentation.

e. (Mandatory) Onefs HDFS license is applied to the failover target cluster for the Easy Auditor Database.

f. (Mandatory) All firewall ports are open between the ECA cluster and the failover target cluster

g. (Mandatory) A System zone NFS export is created with the IP addresses of the warm standby ECA cluster IP addresses at the failover site for the ECA mount to ingest audit data.

h. (Mandatory) Smartconnect DNS Delegation is created for NFS access in the system zone in your DNS infrastructure following Dell documentation.

i. (Mandatory) ECA will require a mount path with the target cluster name (case sensitive) and cluster GUID (Cluster Management > General GUID is visible on this page)

j. (Mandatory) Prepare DR cluster for HDDFS database following steps below. The installation guide has detailed steps. The steps below are high level steps that need to be completed.

k. Create the local Hadoop user for the HDFS database (**eyeglasshdfs**) in the System access zone of the DR PowerScale cluster as per Preparation of Analytics Database Cluster documentation, specify the same UID as

the production PowerScale cluster's hdfs eyeglasshdfs user. Example:

l. isi auth users create --name=eyeglasshdfs --provider=local --enabled=yes --password-expires=no --zone=system --uid=*eyeglasshdfs_uid_on_production_PowerScale*

m. Set the permissions on the HDFS folder using root user on the DR cluster and use the commands below

n. mkdir -p /ifs/data/igls/analyticsdb/eca1/

o. chown -R eyeglasshdfs:'Isilon Users' /ifs/data/igls/analyticsdb/eca1/

p. chmod -R 755 /ifs/data/igls/analyticsdb/eca1/

Scenario #1 - Production Writeable cluster fails over to DR site ECA cluster stays at the production site (Longer RTO)

In this scenario the active cluster is failed over to the DR cluster and the ECA needs to be configured to monitor the DR cluster. This procedure takes time to complete versus Scenario #2 warm standby ECA that offers faster RTO. **NOTE: Assistance can be scheduled with services, reconfiguration is not under the support contract and must be scheduled with services that operates M-F based on resource availability.**

Post Failover Reconfiguration Steps

1. Before failover, audit data is ingested from the production site cluster and the ECA VM's are located at the production site.
2. After failover to the DR Site, the DR cluster needs to have auditing enabled and configure the NFS export configured with a

smartconnect name delegation configured to allow the ECA cluster in the Production site to mount the DR cluster audit folder. This will allow audit data to be ingested by the ECA cluster.

3. Eyeglass license must be assigned to the DR site cluster after failover. The license manager UI in eyeglass is used to assign the license to the DR site cluster.
    a. Set the production site cluster to unlicensed and click submit.
    b. Then set the DR site cluster to user licensed and click submit.

4. 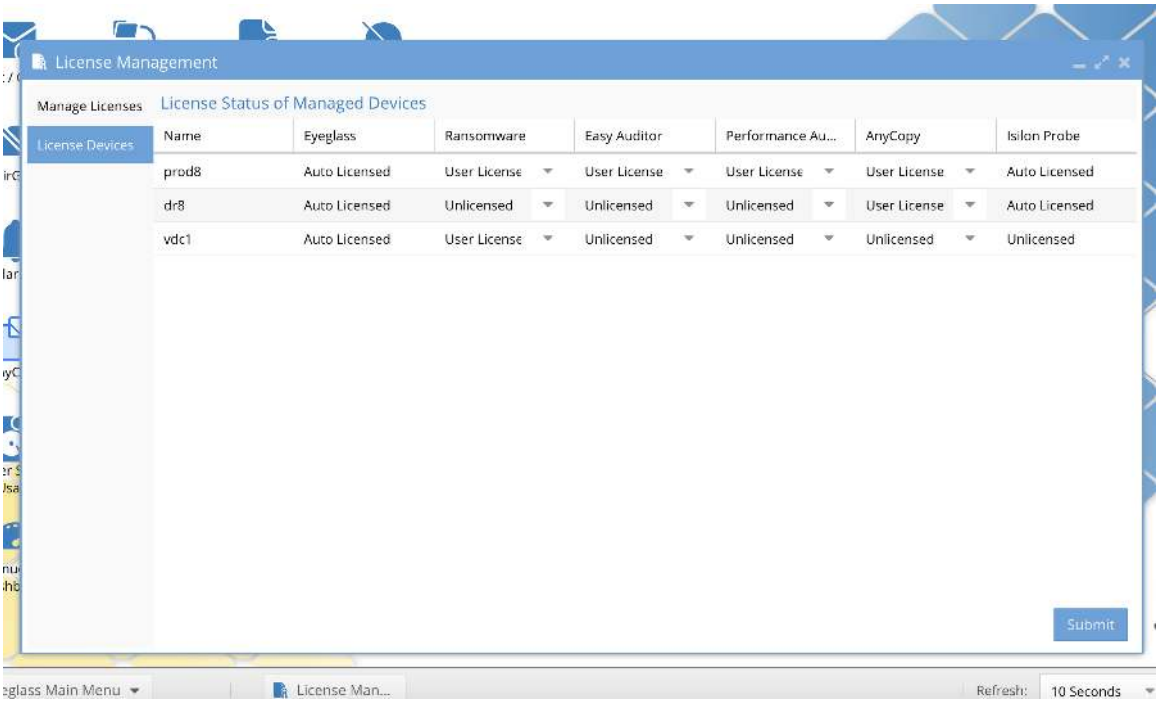

5. **Steps to complete (Follow the Installation guide for detailed steps)**
    a. Get cluster name (case sensitive) and GUID to be used in later steps.

b. Enable Auditing on the DR cluster Onefs GUI Cluster Management --> Auditing --> enable protocol Auditing and add all access zones to the list and save the configuration.

c. Create target cluster NFS export for audit data ingestion.

d. Create smartconnect name to mount the export in system zone, setup DNS delegation for smart connect name to the IP pool to be used for NFS mount in system zone.

e. Create mount path on the ECA cluster using target cluster name and GUID. See guide for steps to create the mount point on all ECA nodes using information collected from Step #1.

f. Edit auto.nfs file on ECA node 1 (/opt/superna/eca/data/audit-nfs/auto.nfs) with new smartconnect name on the DR cluster in system zone and enter mount point for the NFS mount using the path from the step above. Comment with the # character the previous mount entry to prevent mounting the previous cluster. Save the file. (example /opt/superna/mnt/audit/00505699ec55c64cf45d411e36ac285fff13/prod8 -fstype=nfs,nfsvers=3,ro 172.31.1.104:/ifs/.ifsvar/audit/logs)

g. Sync the Configuration to all ECA nodes

    i. ecactl cluster push-config

h. unmount the audit NFS export from the previous active cluster on the ECA cluster.

    i. ecactl cluster exec "sudo umount -a -t autofs (note the password for ecaadmin will be requested for each node)

i. remount the new cluster NFS audit export
    i. ecactl cluster exec "sudo systemctl restart autofs
    ii. Verify the mount was successful on each node, the command below should show the mount on each node.
    iii. ecactl cluster exec mount | grep "ifsvar"

j. Restart services
    i. ecactl cluster services restart --container turboaudit --all

k. Complete license steps documented above if not already completed

l. Reconfigure Security guard and roboaudit features to self test the new cluster. A new local account needs to be created on the new target cluster and edit the the configuration to switch to the target cluster. This requires editing the user and changing the cluster that the automation will run against. The license step must be completed before you can select the target cluster.
    i. See security guard guide here, and roboaudit guide here.
    ii. NOTE: SMB port must be open between eyeglass and the new target cluster.

iii. Run security guard and roboaudit and monitor the job from the Jobs Icon --> running jobs tab.

m. Done

Scenario #2 - Production Writeable cluster fails over to DR site and Production site ECA cluster will failover to warm standby ECA cluster at the DR Site (Lower RTO)

In this scenario the ECA cluster site is impacted and the Warm standby ECA cluster will become the active ECA cluster. In this scenario it is assumed the Eyeglass VM is already located at the DR site as the best practice or the warm standby Eyeglass VM has been activated at the DR site.

Prerequisites to Complete Before a Failover

1. Deploy a Warm Standby Eyeglass vm following this guide. The active Eyeglass VM is located at the DR site, if it is not the active appliance follow the guide to make the Eyeglass VM at the DR site the active appliance.
2. Deploy a 2nd warm standby ECA cluster at the DR site and configure it to use the Eyeglass VM at the DR site. Complete all steps to setup the DR cluster to be protected by the ECA following the installation guide here to install the Warm standby ECA cluster. The points below are used to customize the ECA

installation for the DR cluster. Use the preparation information in this guide to collect the information needed for the steps below.

a. **NOTE: The audit data NFS mount point will use the DR cluster name and cluster GUID. the cluster name and GUID will be added to the warm standby ECA mount points.**

b. **NOTE: The eca configuration file /opt/superna/eca/eca-env-common.conf references the DR site Eyeglass VM (export EYEGLASS_LOCATION= and the API token export EYEGLASS_API_TOKEN) edit this file and add the warm standby VM ip address. Use the same API token value used on the production ECA and get this from the eca-env-common.conf file.**

c. **NOTE: Verify the the hdfs database URL is set correctly for the DR cluster eyeglass access zone and smartconnect name.**

    i. **Make sure the URL matches the production ECA cluster path at the end of the URL see yellow highlight example below to verify. If this path does not match on the DR cluster the ECA will not be able to locate the database files during cluster up command.**

    ii. **edit /op/superna/eca/eca-env-common.conf locate the line below and make sure the DR HDFS smartconnect name is added and path is the same as the production ECA.**

    iii. export ISILON_HDFS_ROOT='hdfs://dr_hdfs_sc_zone_name:8020/eca1'

d. Audit Database DR protection guide should be followed to Sync the Audit database to the DR cluster. This is required to switch the HDFS database to the DR cluster after failover. Follow the Guide here to configure Database sync to the DR cluster

e. In Warm standby mode the ECA cluster should be down (ecactl cluster down) until the warm standby cluster is needed.

f. **Mandatory** - Ensure the warm standby VM's (Eyeglass and ECA) are always running the same release version as deployed at the production site.

Post Failover Reconfiguration Steps

1. After failover to the DR cluster, shutdown the **production** site eca cluster
   a. ecactl cluster down (login to ECA node 1)
2. Follow the steps in the Easy Auditor database protection guide to make the DR cluster database writeable. See the guide here.
   a. NOTE: Make sure to run the synciq policy manually to sync all changes to the database when the ECA cluster is shutdown.
   b. The guide steps above will bring up the cluster and validate the database is healthy
   c. Do not proceed until the ECA and database is fully operational.
   d. **Login to eyeglass and verify the Managed Services Icon shows all ECA VM's are green and no warnings.**

3. Eyeglass license must be assigned to the DR site cluster after failover. The license manager UI in eyeglass is used to assign the license to the DR site cluster.
    a. Set the production site cluster to unlicensed and click submit.
    b. Then set the DR site cluster to user licensed and click submit.
    c. 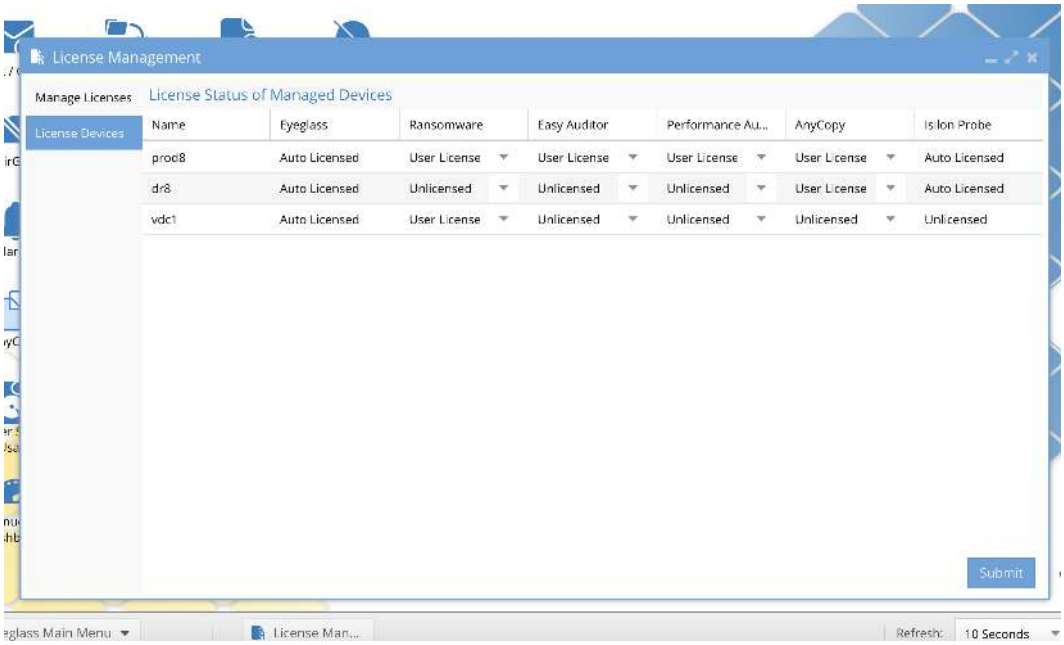

4. Reconfigure Security guard and roboaudit features to self test the new cluster. A new local account needs to be created on the new target cluster and edit the the configuration to switch to the target cluster. This requires editing the user and changing the cluster that the automation will run against. The license step must be completed before you can select the target cluster.
    a. See security guard guide here, and roboaudit guide here.
    b. NOTE: SMB port must be open between eyeglass and the new target cluster.

c. Run security guard and roboaudit and monitor the job from the Jobs Icon --> running jobs tab.

5. Done

© Superna LLC

5.10. How to Change Performance with VMware

Home Top

- Eyeglass ECA Performance Tuning
  - vCenter ECA OVA CPU Performance Monitoring
  - vCenter OVA CPU limit Increase Procedure
  - How to check total ESX host MHZ Capacity

Eyeglass ECA Performance Tuning

The ECA cluster is mostly CPU intensive operation.

1. If the average CPU utilization of the ECA cluster as measured from vCenter and averages 75% or greater, it is recommended to increase the CPU limit applied by default on the ECA cluster.

Default ECA OVA cluster reservation provides 12000 MHZ shared across all the VM's.

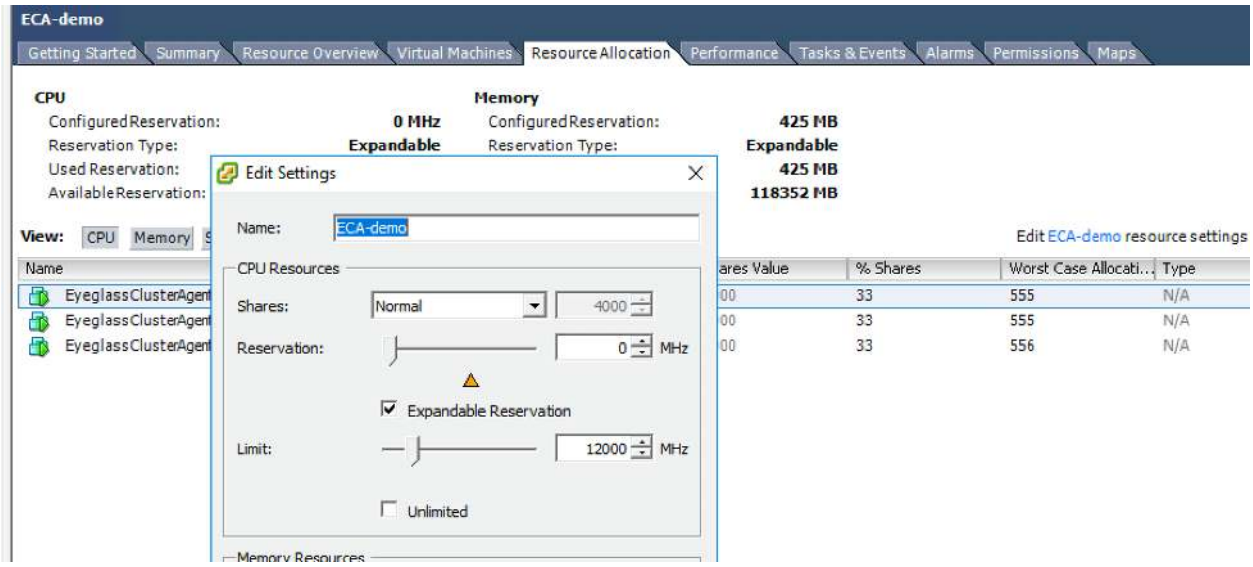

vCenter ECA OVA CPU Performance Monitoring

1. To determine if the ECA MHz limit should be increased.
2. Using vCenter select the OVA cluster Performance tab

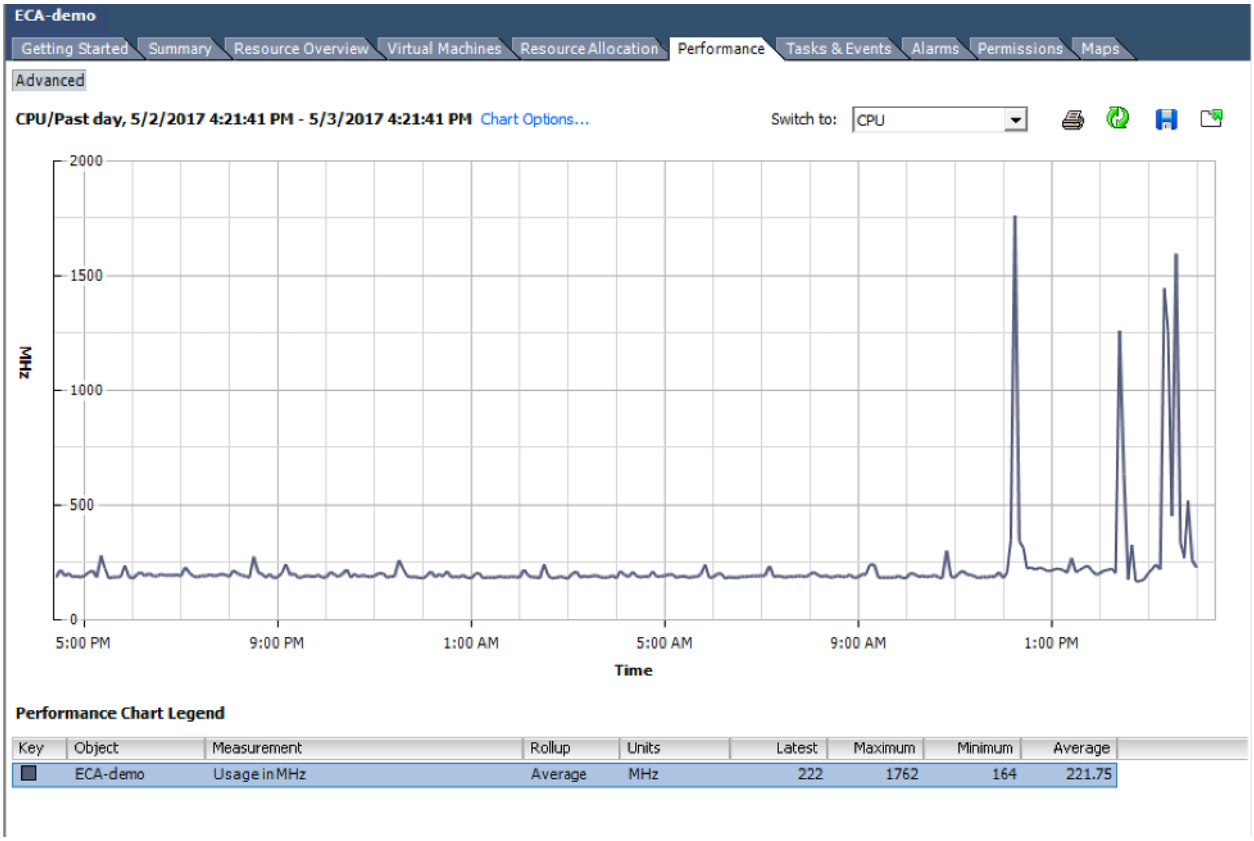

3. As shown above the average Mhz usage is 221 well below the 12000 limit. No change would be required until the average cpu MHZ shows 9000 MHz or greater. The screenshot shows spikes in CPU but the average cpu is the statistic to use.
4. To increase the limit follow procedure below.

vCenter OVA CPU limit Increase Procedure

1. If it's determined an increase is required to new value, it is recommended to increase by 25% and monitor again. Example 12000 * 25% = 3000 MHz
2. Select the Resource Allocation tab on the OVA

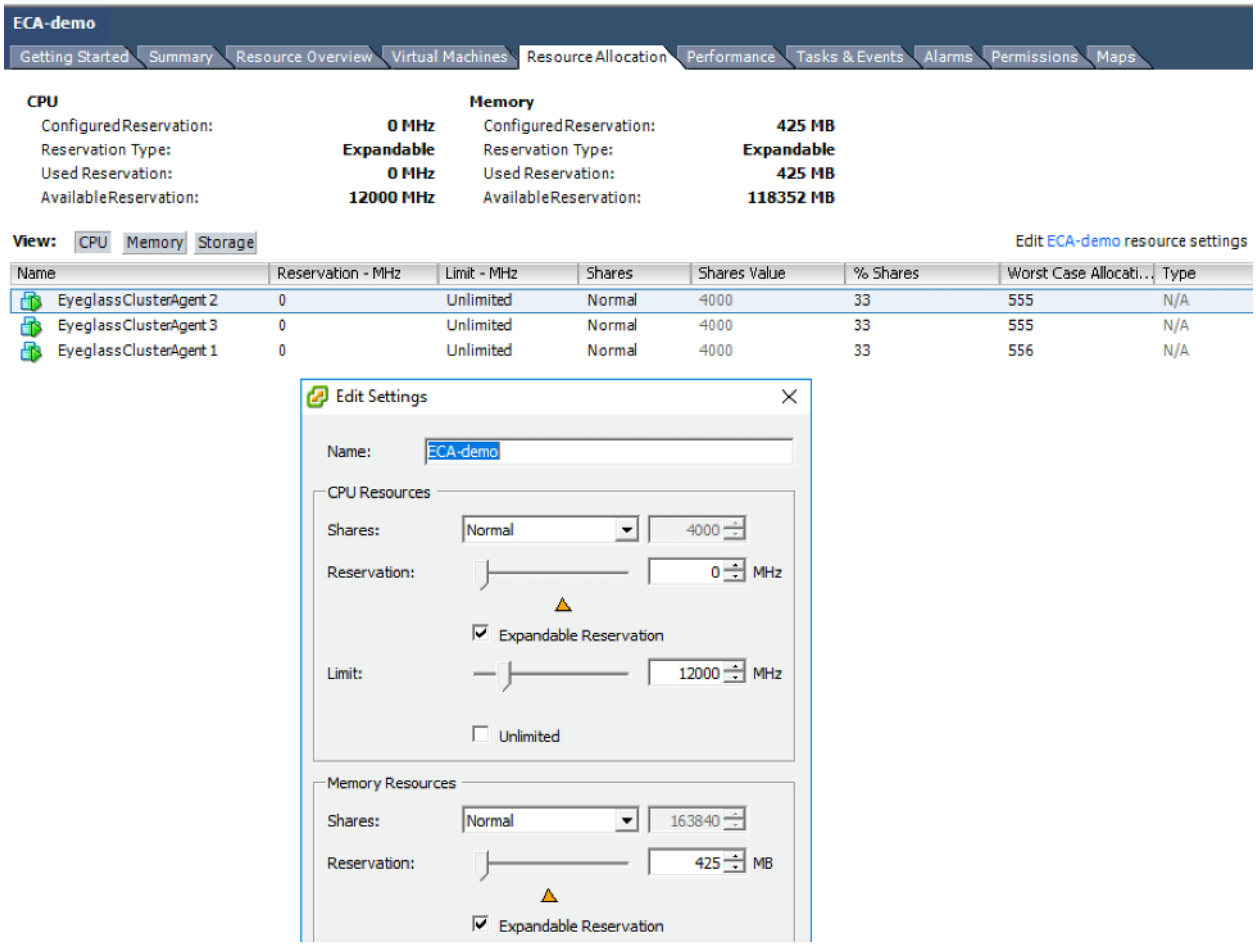

3. Change the Limit value from 12000 to 15000 to increase by 25% and click ok.

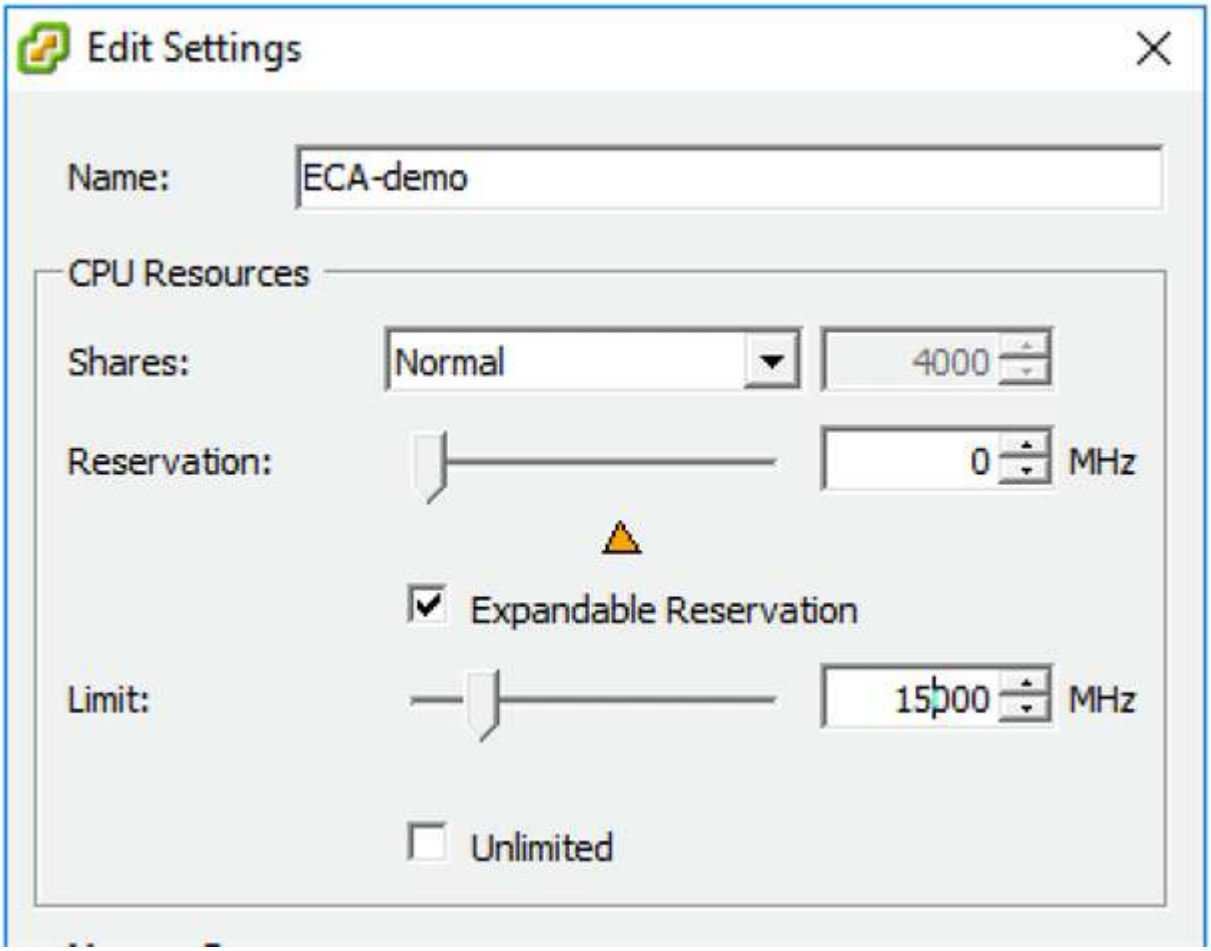

4. Click ok to apply the settings.

How to check total ESX host MHZ Capacity

1. Get the ESX host Summary tab CPU capacity

2. 2.699 GHZ * 1000 = 2699 MHZ per core * 16 = 43,184 MHz of total capacity.
3. This host example is using 6136 MHz of the total 43,184 capacity so there is plenty of unused CPU capacity available on this host.

© Superna LLC

5.11. Dark Sites - How to Health check Eyeglass and ECA clusters when opening a support case

Home Top

- Overview
- How to run the the Log Parser Report
- Steps for more detailed Health Check of components (optional)
- ECA Specific Support data in the Support Report
- ECA Health check for Dark-Sites
  - Hbase-master Database Health Check
  - Kafka-manager Health Check
  - Spark-master Health Check

Overview

This process is for dark sites to generate the log parsing output that provides a summary of information to support. The parsing report can be reviewed for information be for sending to support.

1. Step 1 - run the eyeglass log parsing report, review and send as attachment to support case
2. Step 2 - if requested more detailed steps to verify component health.

How to run the the Log Parser Report

To determine ECA health from log parser output, first run Eyeglass log parser report.

1. Login to Eyeglass CLI as admin
2. Run the following command
   igls app report
3. Wait for it to finish can take 15 minutes to run.
4. Then open your Google Chrome browser [preferred] and go to the following link
   https://<eyeglass IP address>/report/
5. Check the log parser and provide information to Eyeglass support team as requested.
6. Attach the report to the support case
7. NOTE: this report has minimal information and can be reviewed by security for any items in the report.
8. Done.

Steps for more detailed Health Check of components (optional)

Use these steps if support requests additional details after sending the parsing report to support.

Use the Log Parser table of content to browse to sections

Navigation		
Search:		Clear
1	Access Zones And Policies Mapping	No data
2	Active Alarms	Data exists
3	Active Ransomware Defender Events	No data
4	Audit	No data
5	CPU Usage	No data
6	Cluster Information	No data
7	Cluster User Permissions	Data exists
8	Commands Executed To Gather OrientDB Info	Data exists
9	Consumer leaving a topic - from Kafka broker logs	Data exists
10	Continuous Operation Dashboard	Data exists
11	Cronjob count of profiler.sh	Data exists
12	DFS Readiness	No data
13	DR Testing Readiness	No data
14	ECA Disk Usage	Data exists
15	ECA Nodes: GET/POST Failures	Data exists
16	ECA Nodes: cluster up/down, HBASE up/down	Data exists
17	ECA POST requests: heartbeat	Data exists
18	ECA POST requests: notifications	Data exists
19	ECA var/log/messages: out of memory occurrences	No data
20	ECA: Error posting heartbeat to eyeglass	Data exists
21	ECA: Security Guard Events	No data
22	ECA: Threat Detector Overloading	Data exists
23	Event Rates	Data exists
24	Evt Archive - Event Rates	Data exists
25	Existence Of Auth Providers	No data
26	Failovers	Data exists
27	HBase Garbage Collection Errors	Data exists
28	Installed Licenses	Data exists
29	Log Parser Recent Upgrades	Data exists
30	Memory Usage	No data
31	More Than 1 Access Zone Per Base Path	Data exists

1. Check Active Alarm for any ECA related issue

2. Check Remote Services and ensure the Active section states true

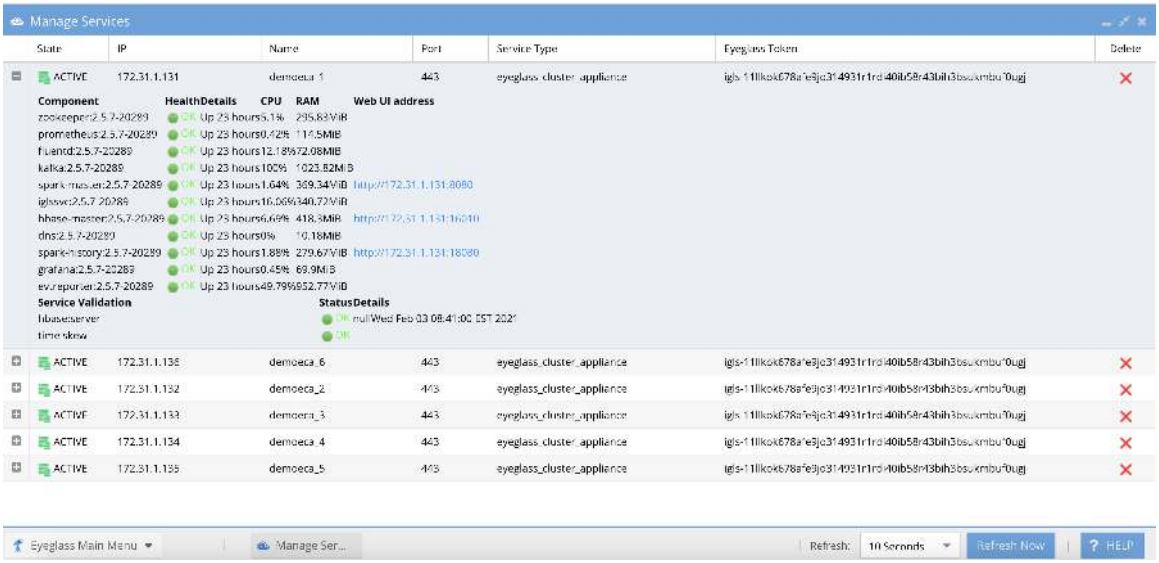

3. Everything under Health section in RemoteServices - Running Containers should be GREEN and OK. When you open support ticket related to ECA containers in Error state, agents will ask for current container state. Most common ones are turboaudit, fastanalysis, evtarchive, spark-worker etc

   Support team may ask you to check Remote Services - Running Containers, remote Services - Validation section for ECA troubleshooting

   If any ECA node is not in true state, that node is NOT in healthy state.

4. Check Remote Services - Validation to determine any time skew and/or issue with HBase-scanning

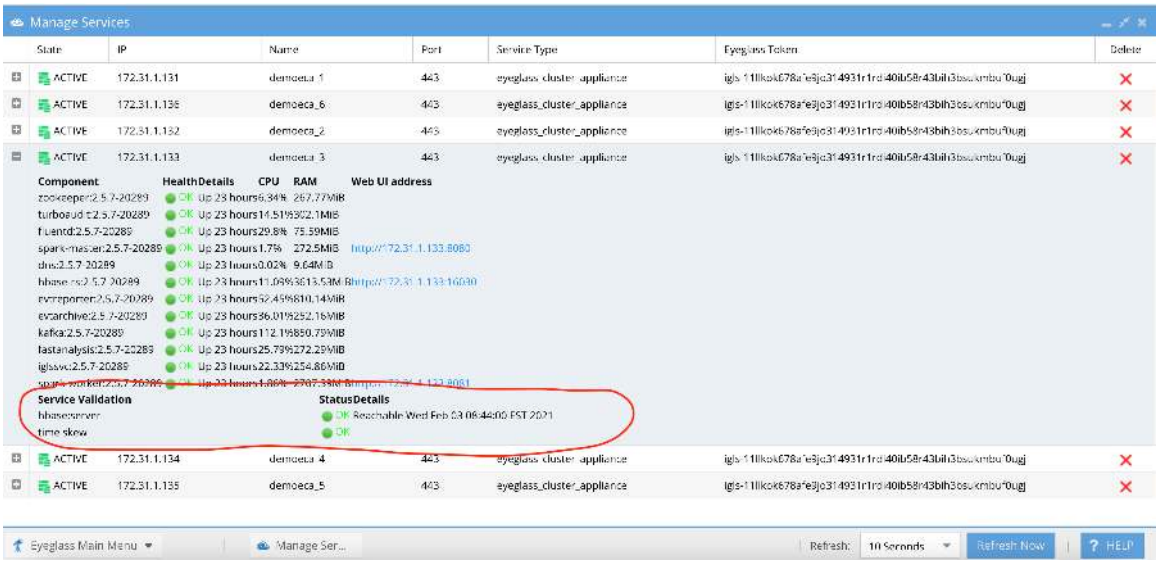

In the picture above, hbase:server Validation is in OK state. Time skew is OK.

5. Check Security Guard for most recent events and provide details. ERROR state indicates an unhealthy ECA environment.

a. From Table of content click on Security Guard Events

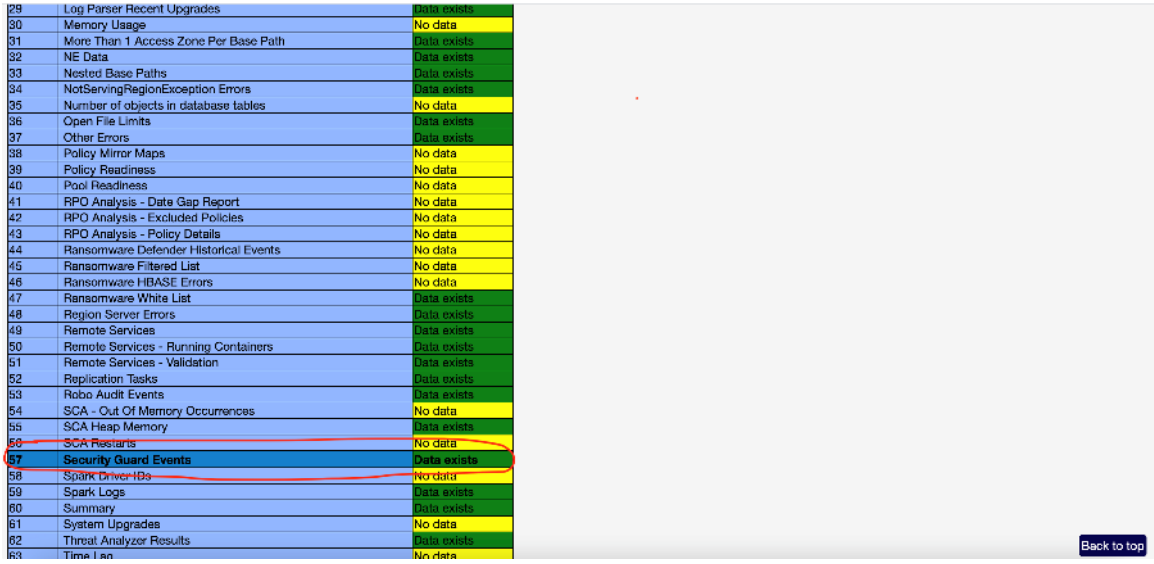

b. Check for most recent events and provide details

Security Guard Events	
Date	Result
Wed Feb 03 07:00:00 EST 2021	OK
Wed Feb 03 06:00:00 EST 2021	OK
Wed Feb 03 05:00:00 EST 2021	OK
Wed Feb 03 04:00:00 EST 2021	OK
Wed Feb 03 03:00:00 EST 2021	OK
Wed Feb 03 02:00:00 EST 2021	OK
Wed Feb 03 01:00:00 EST 2021	OK
Wed Feb 03 00:00:00 EST 2021	OK
Tue Feb 02 23:00:00 EST 2021	OK
Tue Feb 02 22:00:00 EST 2021	OK
Tue Feb 02 21:00:00 EST 2021	OK
Tue Feb 02 20:00:00 EST 2021	OK
Tue Feb 02 19:00:00 EST 2021	OK
Tue Feb 02 18:00:00 EST 2021	OK
Tue Feb 02 17:00:00 EST 2021	OK
Tue Feb 02 16:00:00 EST 2021	OK
Tue Feb 02 15:00:00 EST 2021	OK
Tue Feb 02 14:00:00 EST 2021	OK
Tue Feb 02 13:00:00 EST 2021	OK
Tue Feb 02 12:00:00 EST 2021	OK
Tue Feb 02 11:00:00 EST 2021	OK
Tue Feb 02 10:00:00 EST 2021	OK
Tue Feb 02 09:00:00 EST 2021	ERROR
Tue Feb 02 08:00:00 EST 2021	OK
Tue Feb 02 07:00:00 EST 2021	OK
Tue Feb 02 06:00:00 EST 2021	OK

c. If they are in ERROR state, support team will ask you to open Eyeglass Web UI and provide log details to support team.

- Eyeglass Web UI → Ransomware Defender widget → Security Guard → Open failed log

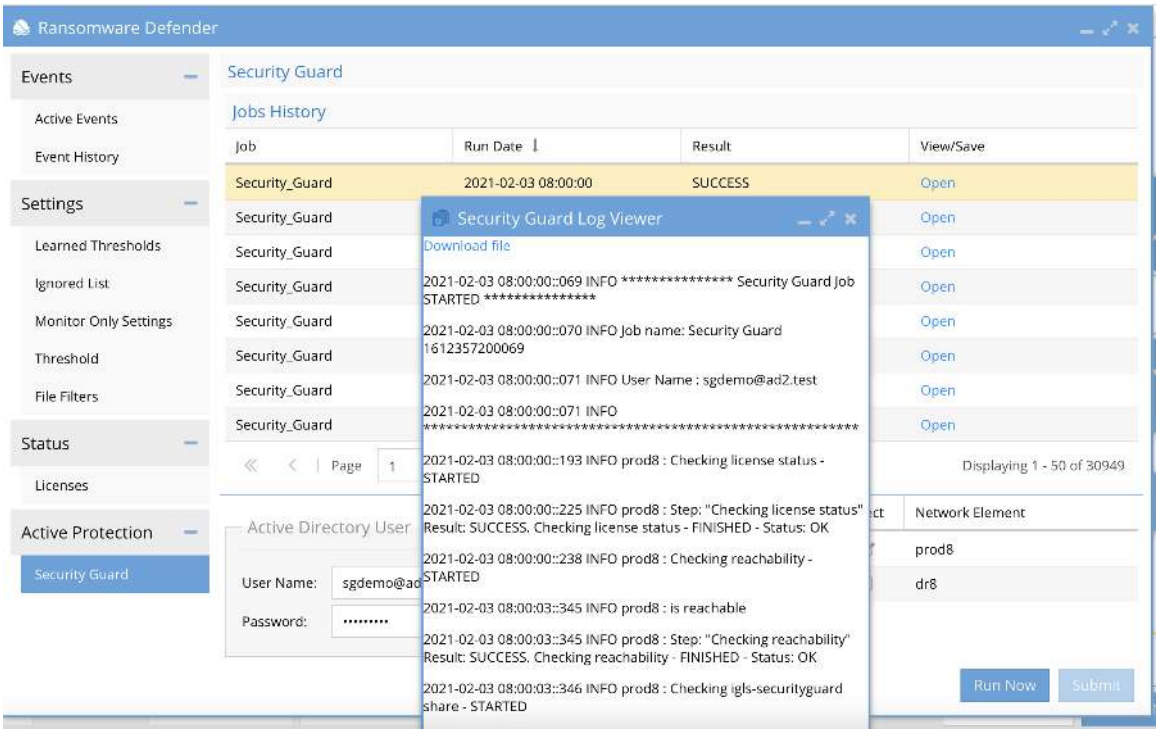

6. **Check turboaudit health and event rates by click the Event Rates/Turbo Audit table of content menu Item number 23 Event rates**

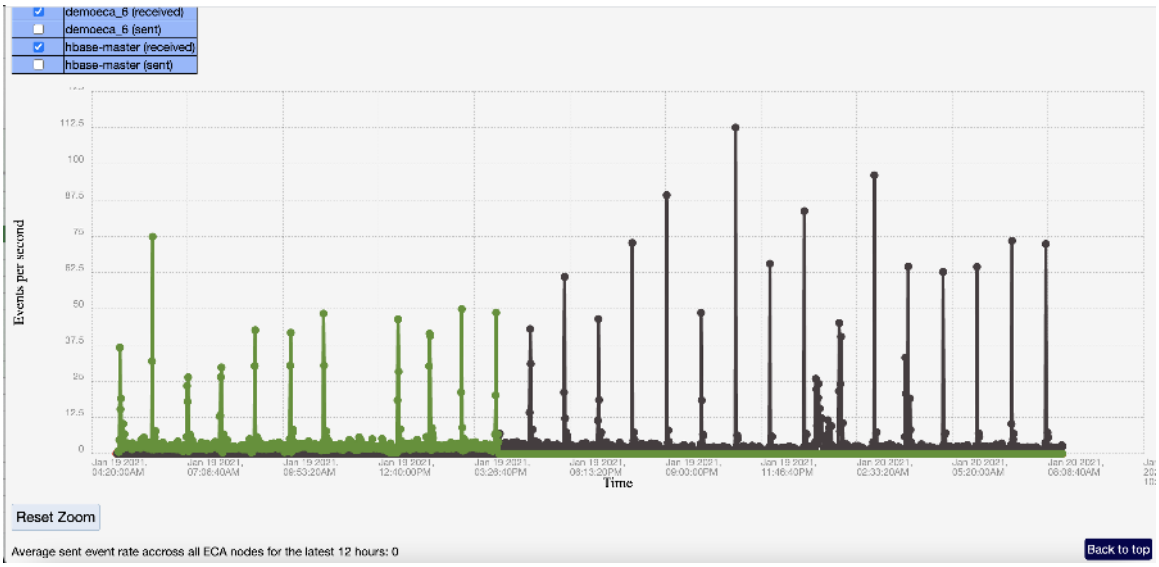

NOTE: Sent and receive Rate for nodes 2-6 ECAs (check sent

and received independently).
NOTE: For “Received” MAKE SURE above 0 EVTS all the way to the right side of graph. IF ALL ZERO this is an issue. Flag for case follow up

7. **Check evtarchive container health and event rates by click the Evt Archive Event Rates table of content** menu item number 24 "Evt Archive - Event Rates"

8. 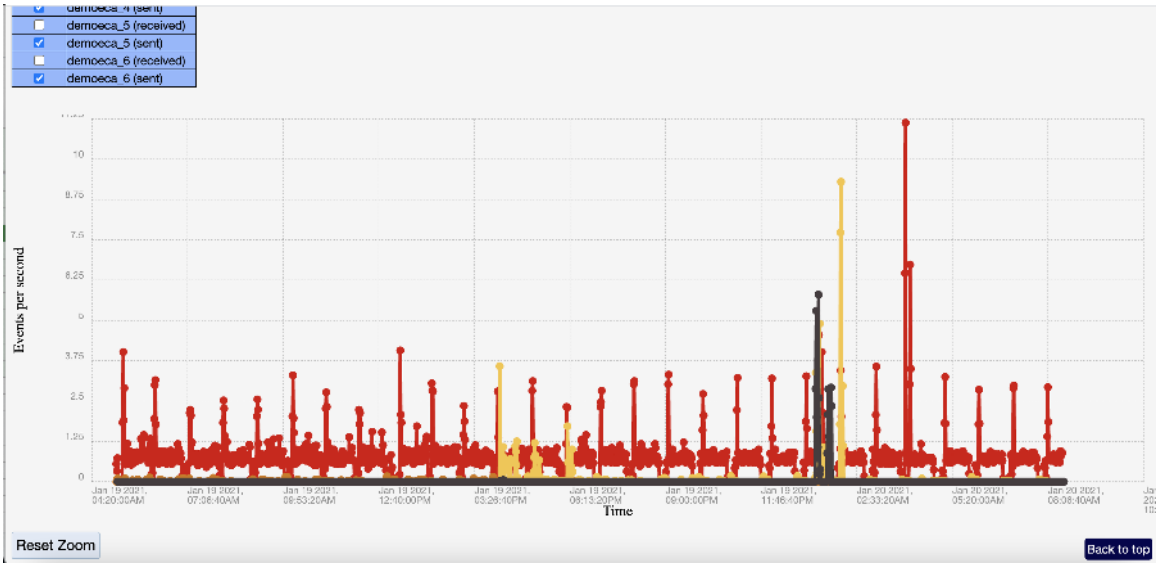

NOTE: MAKE SURE above 0 EVTS all the way to the right side of graph. IF ALL ZERO this is an issue. Flag for case follow up

9. Check ECA node disk space

    f. You may be asked to check disk space as well as CPU/MEMORY usage for ECA nodes. Select the ECA disk usage TOC to check disk space.

ECA Specific Support data in the Support Report

1. First check Manage Services in Eyeglass. If you notice a red/yellow warning, open the widget
2. If containers are showing INACTIVE note down the node # and names of the containers
3. Login to ECA node # using SSH session [Putty]
4. First thing to check is Eyeglass version and summary. Support team will ask for:
    a. Eyeglass version,
    b. Total RAM,
    c. OS version etc
5. If you are running EOL version of Eyeglass, support team will request you to upgrade Eyeglass appliance.
6. If Eyeglass Replication jobs are in ERROR state, click to expand and then hover over to read the ERROR message
    a. 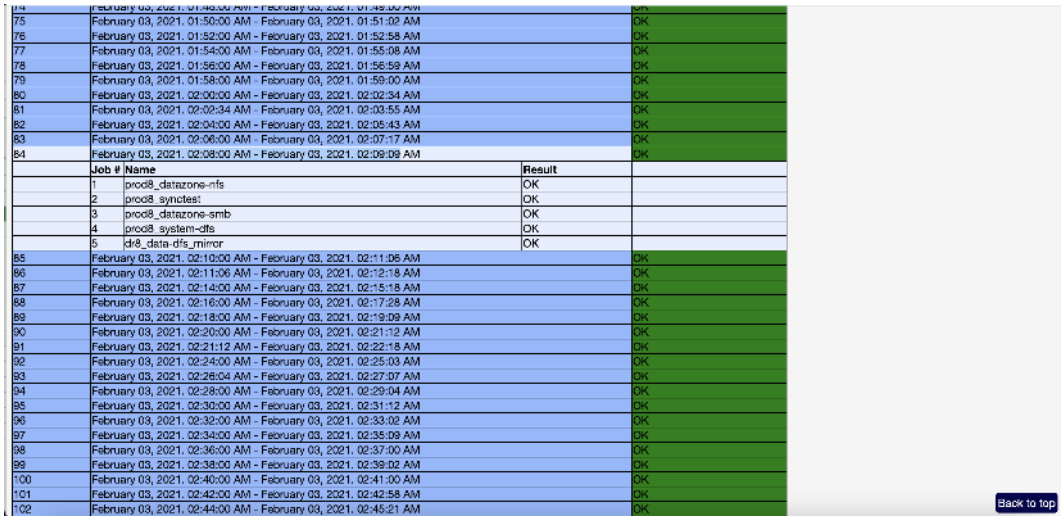

75	February 03, 2021. 01:50:00 AM - February 03, 2021. 01:51:02 AM		OK
76	February 03, 2021. 01:52:00 AM - February 03, 2021. 01:52:58 AM		OK
77	February 03, 2021. 01:54:00 AM - February 03, 2021. 01:55:08 AM		OK
78	February 03, 2021. 01:56:00 AM - February 03, 2021. 01:56:59 AM		OK
79	February 03, 2021. 01:58:00 AM - February 03, 2021. 01:59:00 AM		OK
80	February 03, 2021. 02:00:00 AM - February 03, 2021. 02:02:34 AM		OK
81	February 03, 2021. 02:02:34 AM - February 03, 2021. 02:03:55 AM		OK
82	February 03, 2021. 02:04:00 AM - February 03, 2021. 02:05:43 AM		OK
83	February 03, 2021. 02:06:00 AM - February 03, 2021. 02:07:17 AM		OK
84	February 03, 2021. 02:08:00 AM - February 03, 2021. 02:09:09 AM		OK
Job #	Name	Result	
1	prod8_datazone-nfs	OK	
2	prod8_synctest	OK	
3	prod8_datazone-smb	OK	
4	prod8_system-dfs	OK	
5	dr8_data-dfs_mirror	OK	
85	February 03, 2021. 02:10:00 AM - February 03, 2021. 02:11:06 AM		OK
86	February 03, 2021. 02:11:06 AM - February 03, 2021. 02:12:18 AM		OK
87	February 03, 2021. 02:14:00 AM - February 03, 2021. 02:15:18 AM		OK
88	February 03, 2021. 02:16:00 AM - February 03, 2021. 02:17:28 AM		OK
89	February 03, 2021. 02:18:00 AM - February 03, 2021. 02:19:09 AM		OK
90	February 03, 2021. 02:20:00 AM - February 03, 2021. 02:21:12 AM		OK
91	February 03, 2021. 02:21:12 AM - February 03, 2021. 02:22:18 AM		OK
92	February 03, 2021. 02:24:00 AM - February 03, 2021. 02:25:03 AM		OK
93	February 03, 2021. 02:26:04 AM - February 03, 2021. 02:27:07 AM		OK
94	February 03, 2021. 02:28:00 AM - February 03, 2021. 02:29:04 AM		OK
95	February 03, 2021. 02:30:00 AM - February 03, 2021. 02:31:12 AM		OK
96	February 03, 2021. 02:32:00 AM - February 03, 2021. 02:33:02 AM		OK
97	February 03, 2021. 02:34:00 AM - February 03, 2021. 02:35:09 AM		OK
98	February 03, 2021. 02:36:00 AM - February 03, 2021. 02:37:00 AM		OK
99	February 03, 2021. 02:38:00 AM - February 03, 2021. 02:39:02 AM		OK
100	February 03, 2021. 02:40:00 AM - February 03, 2021. 02:41:00 AM		OK
101	February 03, 2021. 02:42:00 AM - February 03, 2021. 02:42:58 AM		OK
102	February 03, 2021. 02:44:00 AM - February 03, 2021. 02:45:21 AM		OK

Back to top

7. Provide the error msg to Eyeglass support team
8. Support team will ask for additional information such as Active Alarm. Scroll down the page to locate the alarms. Use menu item #2 for active alarms

ECA Health check for Dark-Sites

When Eyeglass support team is asking for health check status please perform the following steps and provide output to agents:

1. Login to Eyeglass Appliance Web UI
2. Check for error/warning in Manage Services.
3. If you have received an Eyeglass alarm - collect the information and provide to support
4. Open Eyeglass Manage Services widget and check containers state
5. ECA containers provide Web UI access to collect information. You can locate the Web UI link from Manage Services window

Hbase-master Database Health Check

Support team may ask for hbase-master server status. To collect hbase-master statistics,

a) Browse to http://<ECA_NODE1_IP>:16010

b) Take Screenshot:

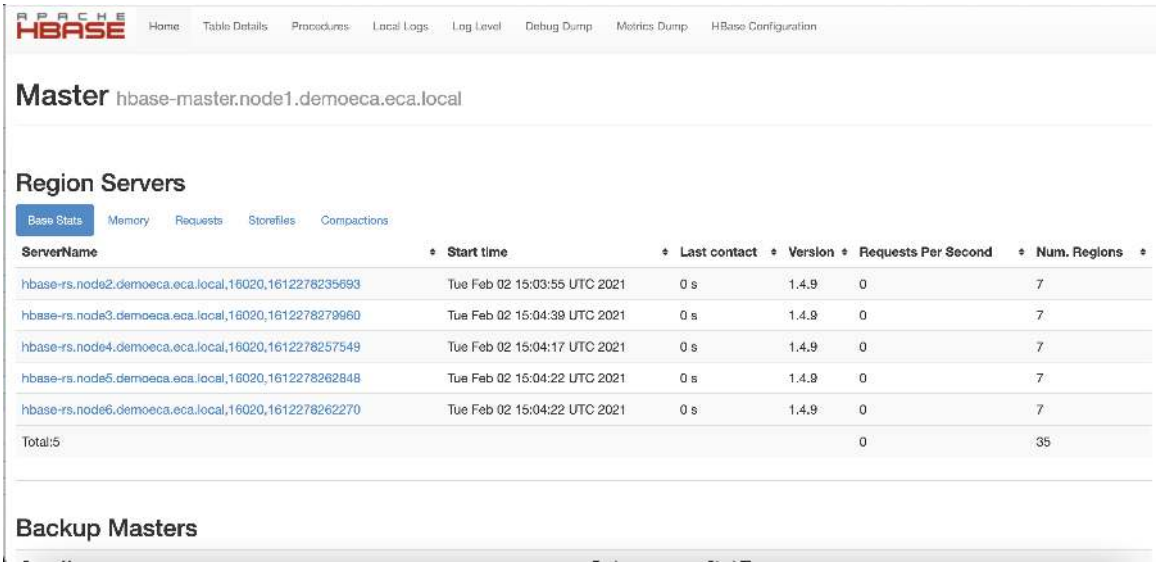

c) Check Hbase tables. Look for Offline/Failed regions take screenshot

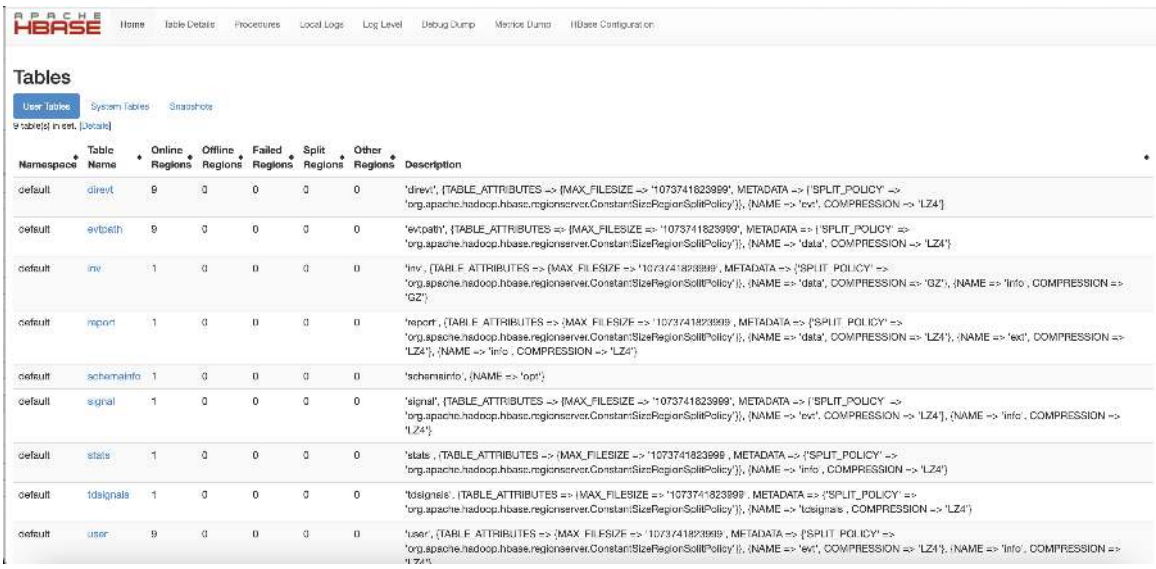

d) Collect the hbase-master server info and provide to support

Kafka-manager Health Check

Support team may ask for kafka-manager status. To collect kafka-manager statistics,

a) Browse to http://<ECA_NODE1_IP>/kafkahq

b) If kafka-manager is NOT configured, set it up before using it from node 1 of the eca

ecactl containers up -d kafkahq

ecactl containers start kafkahq

Take screen shot and add to the support case

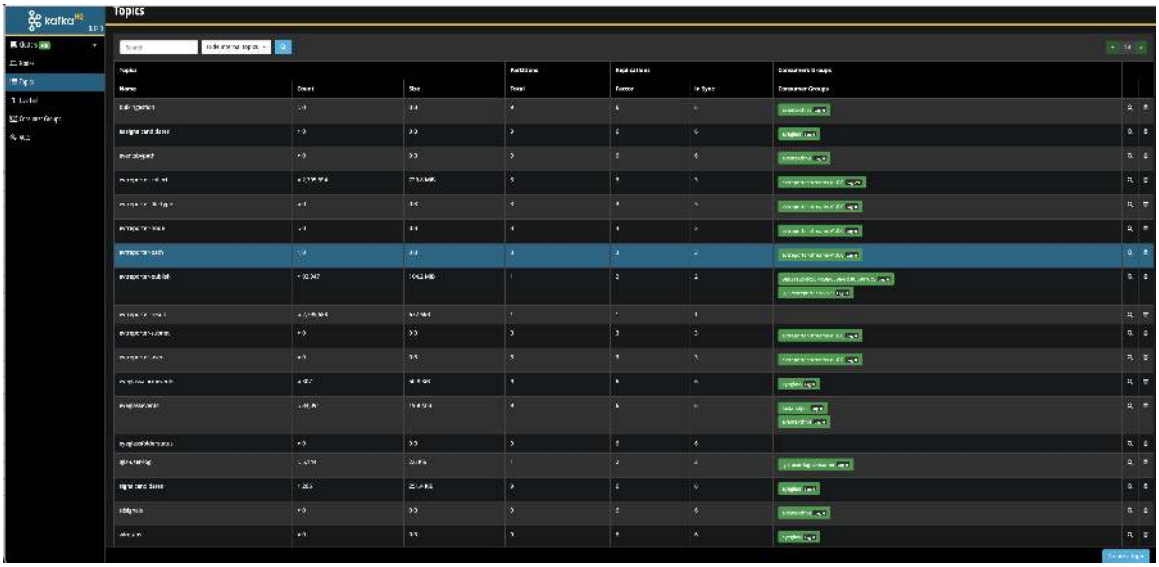

Spark-master Health Check

When Easy Auditor reports having issue, support team may ask you to check spark-master GUI. For this reason, you need to browse to spark-master Web UI

1. Take a screenshot as per below.

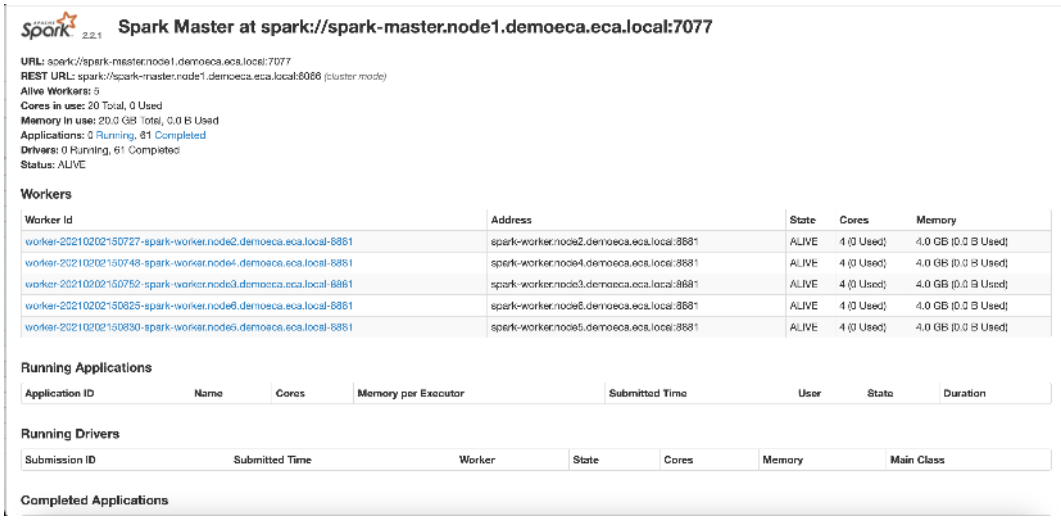

2.

© Superna LLC

5.11.1. Troubleshooting ECA Configuration Issues

Home Top

Troubleshooting ECA Event Processing

This section covers how to troubleshoot cluster event ingestion

How to configure monitoring

1. Login to eca
2. ecactl containers up -d kafkahq
3. Then goto https://x.x.x.x/kafkahq (x.x.x.x is node 1 of the eca cluster)
4. enter the ecaadmin user and password (default is 3y3gl4ss on the ECA)
5. Screenshot the topics list and provide to support. The shows processing and the lag column with a positive number indicates the backlog in event processing.
6. 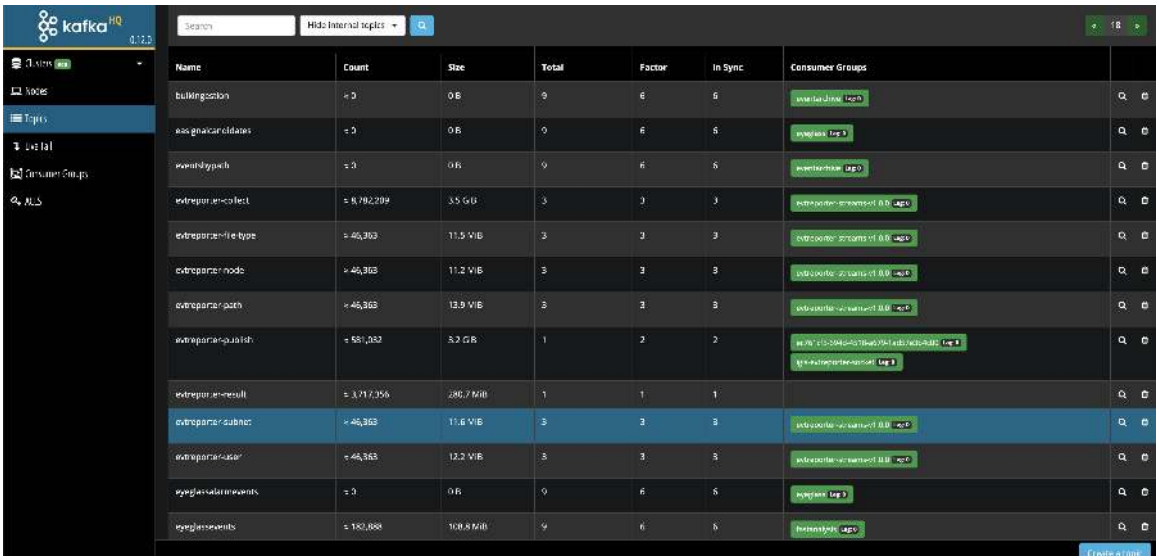

7. done

© Superna LLC

6. Eyeglass Cluster Storage Monitor Admin Guide

Home Top

- Introduction to this Guide and What's New
- Active Directory Managed Quotas Overview
- How to Manage Storage by the Share or Export
- Operations - Cluster Bulk Quota Management Features
- Configuration - Cluster Storage Reports
- Configuration - IGLS CLI commands to configure Cluster storage monitor features
- Unlock My files Help Desk Application

© Superna LLC

6.1. Introduction to this Guide and What's New

Home Top

- Overview
- What's New
- Overview Video's
- License Requirements:

Overview

As data grows, and clusters are deployed in remote locations the cost of administration and management of these complex systems is growing. Better tools are required to automate display of summary usage as well as produce quick searches across multiple clusters for quota usage. Cluster Storage Monitor product is focused on storage consumption, storage tier usage, cluster health, quota usage and managing locked files. The Cluster Storage Monitor will reduce administration cost by:

- Simplifying storage reporting
- Providing health check at a glance,
- Removing manual steps from quota administration and allowing AD management of quota's
- Providing a help desk tool to manage locked and open files without requiring permissions to the storage management interface
- Cluster Storage Monitor features can be used for quota storage chargeback with reporting for all quota types.

What's New

See what new features are coming with each new release here

1. 2.5.5 - adds Direct LDAP to AD to collect users and groups for large AD environments. This caches users and groups needed for AD managed groups and collect directly from AD on a schedule. The CLI guide explains how to configure this feature. Recommended for large user and group AD forests example > 10 000 users and groups.
2. Active Directory Group Quota Management feature allows managing user quotas using AD groups
3. Pre Sync Quotas (DR feature) sync quotas to DR cluster continuously (See DR Design guide for details)
4. 2.5.4 Unlock my files! feature allows a help desk to be delegated permissions to find open files and break the lock.

Overview Video's

License Requirements:

1. Existing cluster licenses for DR Configuration Replication, each cluster will be enabled .
2. License key activates all licensed clusters, maintenance purchased separately.
3. Trial key limits:

a. Storage by share/export tab will only display shares or exports with a Name, or with a path that includes pattern of igls-quota-sharename or /ifs/data/somepath/igls-quota-export1.

b. The Product license key is global and activates all clusters under management with valid cluster license and removes the share or export name limitation used only for lab testing or trial of the feature.

© Superna LLC

6.2. Active Directory Managed Quotas Overview

Home Top

- Overview
- Overview diagram
  - Quota Types
  - Active Directory Groups used to target where to apply quotas or How to apply Quotas
  - AD Group membership and quota creation Schedules
- Use Cases:
  - User Home Directory Automatic Quota Assignment - User Quota Mode
  - Migrating from Linked Quotas to AD managed quotas
- How to Configure Group Quota Templates
- How to Monitor Cluster Storage Monitor AD managed Quota Evaluation Results

Overview

To simplify quota administration this feature allows Active Directory groups to be used to detect quotas that should be automatically applied to users or directories based on group membership or apply group quotas on a share.

Overview diagram

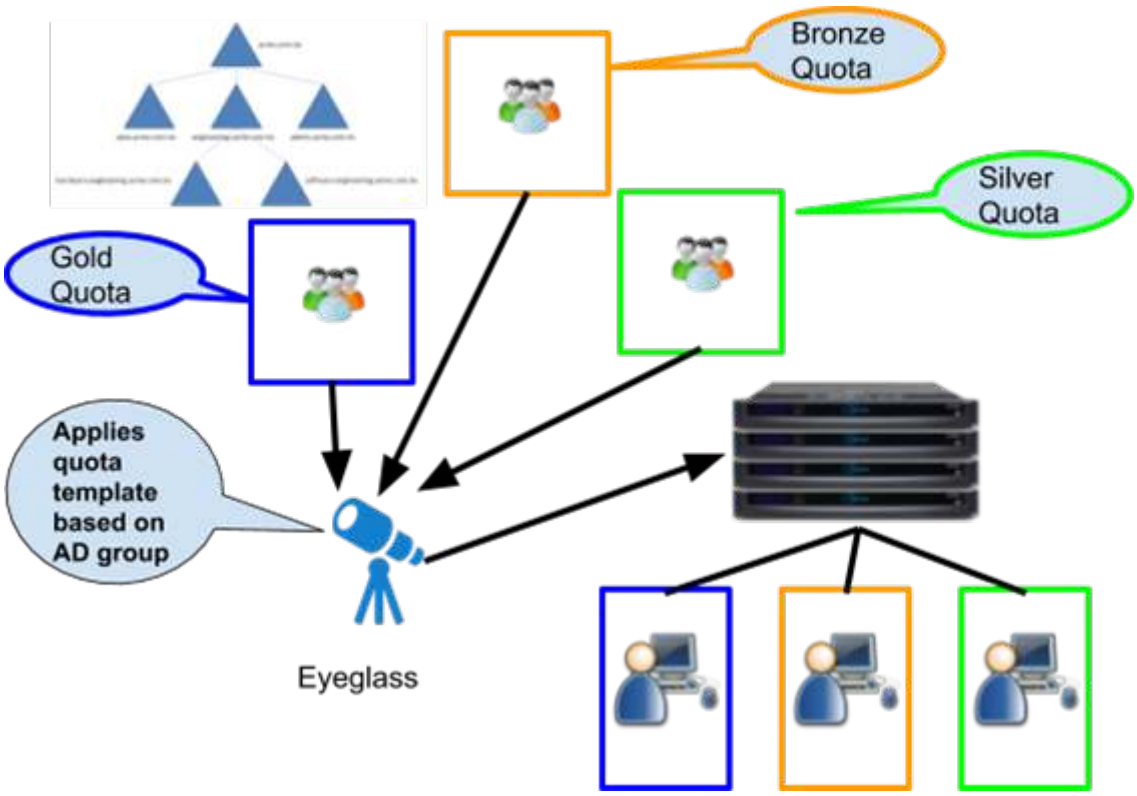

This capability allows quota templates to be created in Eyeglass based on Gold, Silver and Bronze labels or a custom label name and define hard, soft or advisory quotas. No direct connection to AD is required with all user group AD membership retrieved using the PowerScale API's.

Quota Types

1. User Quotas
2. Group Quotas
3. **Quota Enforcement type:** Soft, Hard, Advisory
4. Directory quotas are included in reporting but cannot be managed by AD group feature.

Active Directory Groups used to target where to apply quotas or How to apply Quotas

Eyeglass quota templates can be configured to create automatic detection and apply quotas. This feature is configured with Eyeglass CLI to create

the templates. The running jobs icon will show the detection job and apply quota job running on its schedule. The schedule is managed by the CLI commands (see the CLI guide here). Once configured Eyeglass will run a special quota job that detects AD group membership and reviews quotas applied to clusters and updates user quotas to add, or update quota to match AD group membership based on the template definitions.

NOTE: Deleting quotas is disabled and requires manual delete to avoid accidental AD group change causing delete of quotas. Options exist to handle quota deletion, see CLI guide on options to handle different scenario's . The default settings do not delete quotas.

NOTE: If a user has quota increased by being a member of a template AD group, and then added to a template AD group that would lower the quota. Eyeglass will not lower the quota. This ensures that the user does not end up with blocked writes from a reduced quota. It also allows a higher quota to be manually changed on the cluster without risk of Eyeglass changing the quota to a lower value.

The admin guide covers all the IGLS cli command to configure the features. This guide covers configuration, planning and some example cli commands.

AD Group Modes on Templates

A template has an AD group mode that determines how the AD group will be viewed when evaluating when to create a quota for a user.

1. **Quota Template User mode Enabled** - an AD groups user membership are used to create an PowerScale user quota on all shares that have the AD group assigned to share permissions

2. **Quota Template Group mode Enabled** - An PowerScale Group quota will be applied on shares that have the AD group assigned to share permissions

AD Group membership and quota creation Schedules

AD Group membership task once enabled (see IGLS commands in the CLI guide), will evaluate all AD groups created in quota templates to determine users that should have new quota created or updates to existing quotas. The default schedule is every 12 hours. This can be changed.

The quota updates are done on a different schedule and is determined by how quota inventory is configured.

1. Default quota inventory collection occurs during normal configuration replication jobs that run every 5 minutes by default. This means that quotas will be created based on the last execution of the AD group membership task approximately 5 minutes after the AD group task completes. **Recommendation: for < 1000 quotas use default configuration**
2. If quota inventory schedule has been configured for large quota collection, this task runs collection once per day by default. This means quota creation step for AD managed quotas will now follow the quota collection schedule as well. **Recommendation: For faster quota creation after AD group membership changes, align the quota inventory (default once per day) and the AD group task (default 12 hours) to new values. If not change to the defaults is done, then quota updates will occur once per day.**

AD Groups for Security Versus Quota assignment Best Practice.

AD groups used for securing access to shares can be used in quota templates. This is when all users that have access to the share should have a quota applied. NOTE: if some users that have access to the share based on an AD security group require a different quota value, this can be accomplished with a second quota template in Eyeglass.

Note: the highest quota value will be applied for any given user if a conflict exists between the quota templates when more than one AD group tier template matches a given user.

Recommendation: Use AD security groups already present on shares when possible to simplify management of quotas

AD group for security and for quota AD group can be different if only a subset of the users require a quota. This is possible by creating the new quota template AD group and assigning to the share. The share permission should be equal to the security group and best practice is to place the AD group at end of the share list so it is evaluated last.

Recommendation: Create second AD group for quota only detection and auto creation when a share requires only a subset of users to have a quota applied. You can also use the security group as the default quota setting and then create a second higher quota limit AD group for those users that require a higher quota limit from the default.

Use Cases:

User Home Directory Automatic Quota Assignment - User Quota Mode

1. Create a template and enable the AD group **User mode** to apply a User quota to the members of the Active Directory Group named in the template.
2. Assign the AD group to the share permissions list of one or more shares (you will need to apply everyone full control or read\write Share level permissions or at a minimum the same security access as the users security AD group applied to the share). **Best Practice**: Move the share permission to the bottom of the share permission list so it is evaluated last.
3. Supports shares with %U variable expansion feature on PowerScale or normal share names.
4. Eyeglass will retrieve the user group membership on an scheduled interval (see IGLS command for changing this default schedule 'igls admin schedule') and will create or update user quota's on all shares that have the Active Directory group applied to share permissions from the template created on step 1.

Migrating from Linked Quotas to AD managed quotas

PowerScale user quota's allows an everyone feature that auto creates a user quota for all ad users in a domain under a path and links this user quota to a parent quota allowing simply edits to all linked quotas to a new value.

The problem is all users get the same quota assigned and unlinking the quota is the only way to override a quota for a specific user or group of users.

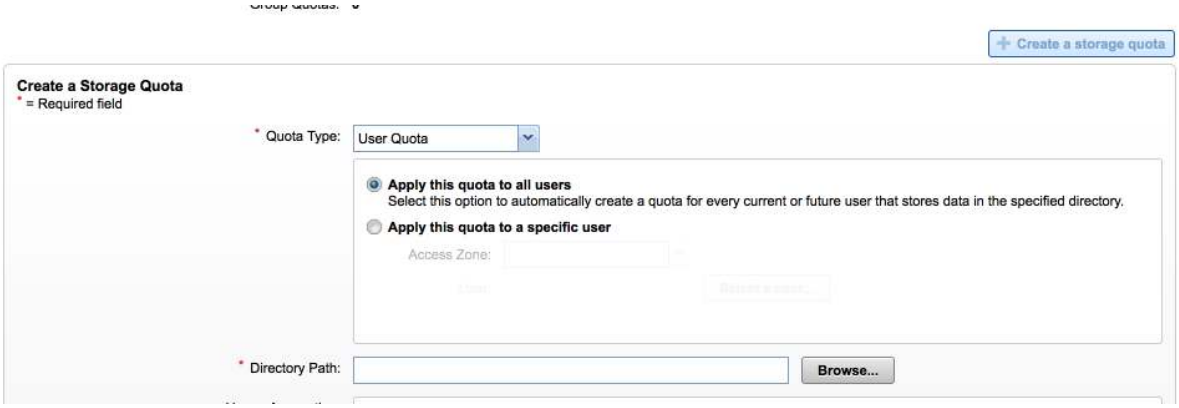

AD managed eyeglass quotas allows AD groups to offer different quotas to users on the same path or across multiple paths and even across clusters.

These steps allow migrating from a parent quota with links to child user quotas to Superna Eyeglass AD managed quotas.

1. Create a template in Eyeglass to define a quota definition
2. Create more than one template and assign different AD users to each group.
3. Apply the AD groups to the share permissions list of the share where you want the user quotas applied.
4. Optional apply the AD group to other shares on the same cluster or different clusters managed by Eyeglass.
5. Eyeglass will now verify the AD groups and start to apply quotas, if a user quota already exists linked to a parent quota that was already in use. Eyeglass will unlink the quota and apply the quota definition defined in Eyeglass.

6. This provides a seamless transition to AD managed quotas and can be staged with only a subset of the users by creating a group with limited number of users in the AD group.
7. Both AD managed and parent linked quotas on the same path can co-exist.

Group Share Automatic Quota Assignment - Group Quota

1. Create a template and enable the AD group Group mode. This will locate all the shares with the AD group specified in the template and apply a group quota using the same AD group.
2. The user to group task does not require users in the group and the quota apply task can apply the group quota on the next scheduled update. This can be controlled using the igls admin schedule command).
3. **NOTE: If using a single group share and then using ACL's on subfolders to secure the group space. Then create a <share name>$ to hide the share and set access to read only or Deny. This share will not be used by users to connect to the group space. The AD template group applied to this share is a marker for Eyeglass to apply the group quota**

4.

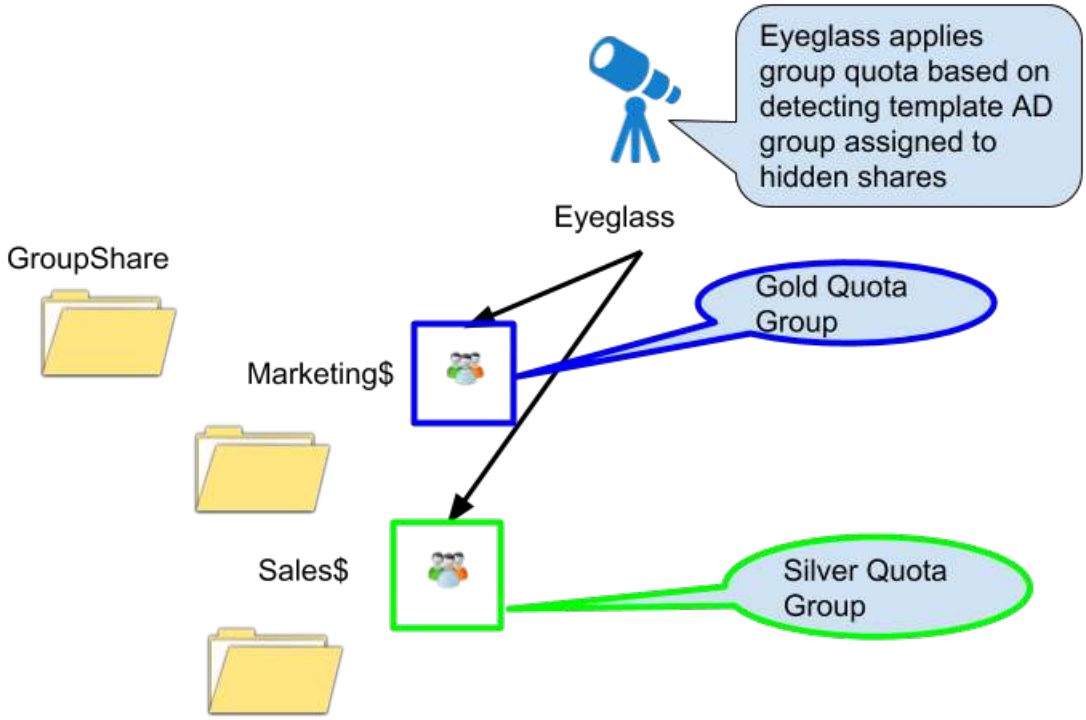

How to Configure Group Quota Templates

1. All CLI commands required to create templates, modify and delete templates and changing schedules of the two tasks required for this feature is located here.

How to Monitor Cluster Storage Monitor AD managed Quota Evaluation Results

1. Login to Eyeglass over ssh
2. tail -f /op/superna/sca/logs/csm.log

3. This log will show the evaluation of AD groups and missing quotas and if quota updates ares skipped.

© Superna LLC

6.3. How to Manage Storage by the Share or Export

Home Top

- Automatic Share & Export Advisory Quota Mode
  - Use Cases
  - How to Enable Auto Quota Mode

Automatic Share & Export Advisory Quota Mode

This completely automates management of storage by share or export by detecting any new share or export without an advisor quota and creating it automatically. This ensures the quota reports are 100% up to date, tracks all shares and exports. Requires no administration or process to enable quotas. The daily CSV report will show the auto created quotas.

Use Cases

1. This can also be useful with Quota portal when not all shares have quotas applied.
2. Defaults to disabled.
3. Automatically removes quota of share or export is deleted, path changed.

How to Enable Auto Quota Mode

For full cli documentation

1. Ssh as admin user to appliance or use web shell from main menu.

2. igls adv quotas help.
3. igls adv quotas (see current values).
4. igls adv quotas set --quotasync=true (this enables the feature, false to disable).
5. igls adv quotas set set --quotasyncdelete=enabled (**defaults disabled, valid values are enabled/disabled/advanced**).
    a. **Disabled** means no quotas will be deleted after share or export delete, they must be deleted manually.
    b. **Enabled** When a share is deleted, this mode **deletes all quotas in that path** unless another share exists in the same path **NOTE: do not enable quota sync delete unless you are sure you can. This mode will detect a share being deleted and delete all quotas at this path and below in the file system.**
    c. **(recommended mode) Advanced** When a share is deleted, this mode **deletes only Eyeglass created quotas in that path** unless another share exists in the same path.

© Superna LLC

6.4. Operations - Cluster Bulk Quota Management Features

Home Top

- Cluster Bulk Quota Management Features
- Quota Advanced Search
- Bulk Quota Changes

Cluster Bulk Quota Management Features

Managing quotas in environments with thousands or tens of thousands can create a huge administrative effort to track and manage. The Advanced Search feature allows searching using various criteria across one or more clusters

Quota Advanced Search

1. Login to Eyeglass and open the Quota Search Icon.
2. Click Advanced Search and combine fields to find quotas to produce a search result.

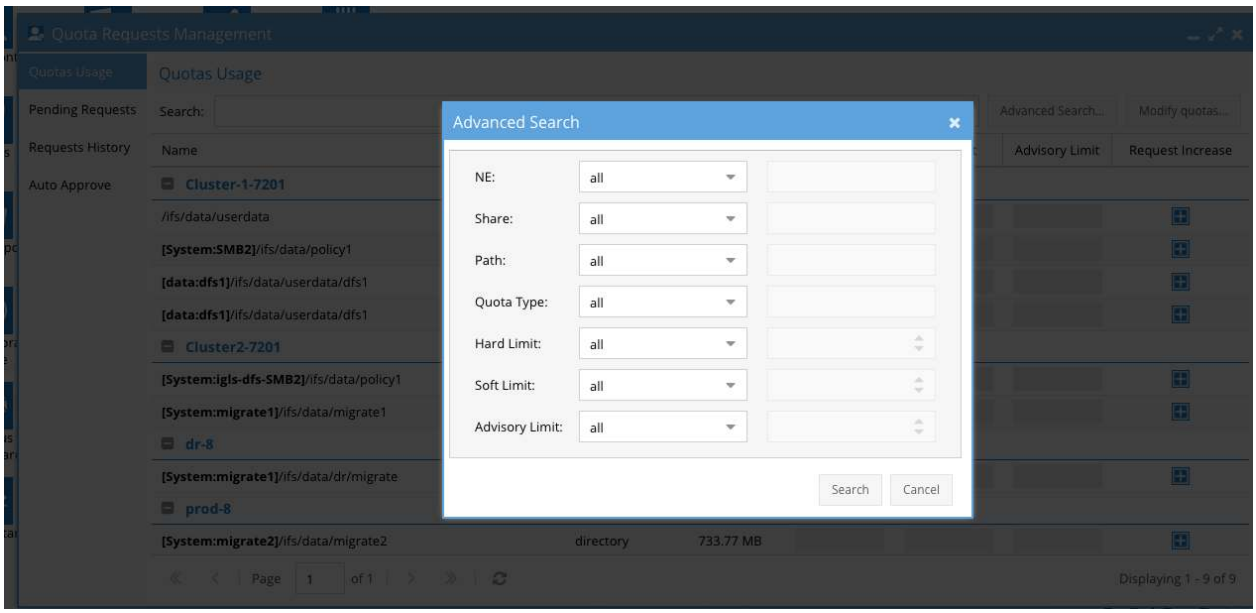

4. Review the quota details, edit or make a request on behalf of the user.

Bulk Quota Changes

1. Login to Eyeglass and open the Quota Search Icon.
2. Click Advanced Search and combine fields to find quotas to produce a search result.
3. Now click Modify Quotas button. This will allow changes based on the search results.

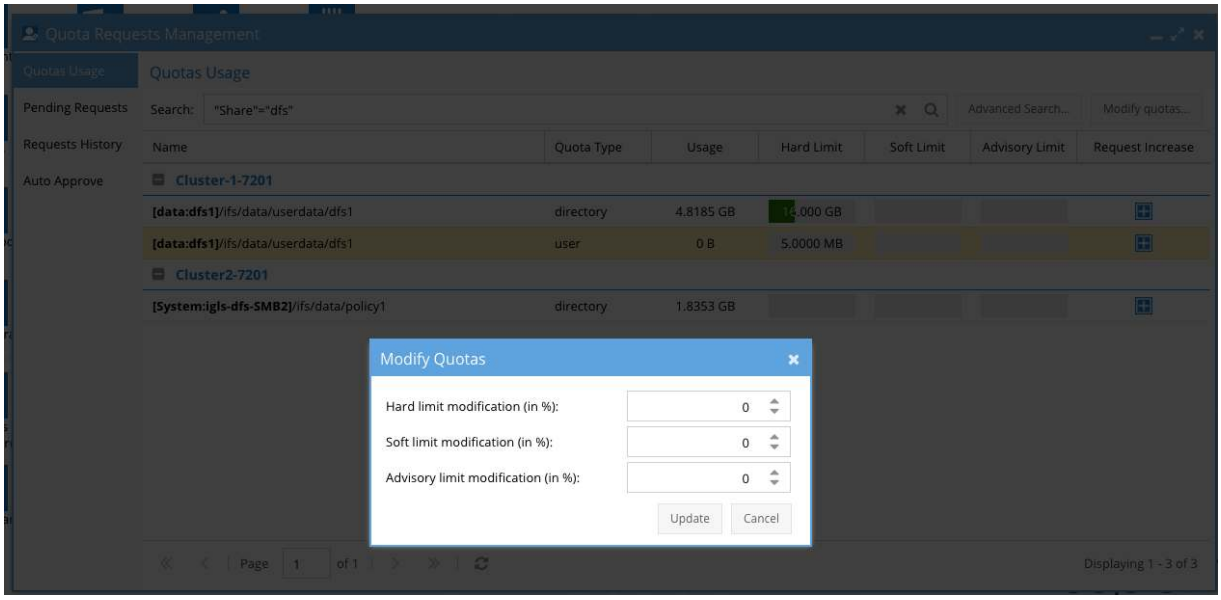

2. Make the changes as % up or down.
    b. **IMPORTANT: Any change made to a parent linked quota will be applied to all child linked quotas**
4. Select update and a quota update job will be submitted for the next configuration replication cycle to apply the changes.
5. This is a very powerful feature that works across clusters. **Note: no undo function exists, changes are permanently applied.**
6. Consult the running jobs window to verify successful quota updates. Verify from OneFS quota UI.

© Superna LLC

6.5. Configuration - Cluster Storage Reports

Home Top

- Cluster Storage Usage Reports
- Cluster Storage Report Includes
- How to Run Cluster Storage Report On Demand
- How to Change Cluster Storage Report Schedule

Cluster Storage Usage Reports

This feature sends automated reports via email. The reports can be used for chargeback and each report provides different data. The reports present cluster, node, share/export and quota usage data.

Cluster Storage Report Includes

1. 3 CSV files provide: Cluster Usage, storage pool usage , Hardware health and quota usage summary.
2. A quota summary CSV file is attached
    a. The CSV includes a sum of all hard, soft and advisor quotas and % of available disk space, and overhead options o the quota.
    b. All quota types will be listed in the report on all clusters added to Eyeglass.
    c. 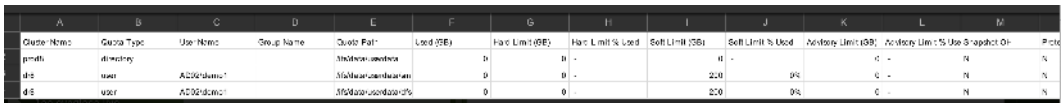

	A	B	C	D	E	F	G	H	I	J	K	L	M	
	Cluster Name	Quota Type	User Name	Group Name	Quota Path	Used (GB)	Hard Limit (GB)	Hard Limit % Used	Soft Limit (GB)	Soft Limit % Used	Advisory Limit (GB)	Advisory Limit % Use	Snapshot On	Fsto
	pmdb	directory			/fs/dbvar/usedvice	0	0	-	0	-	0	-	N	N
	dfs	user	AD2\Admin1		/fs/dbvar/usedbvar/am	0	0	-	200	0%	0	-	N	N
	dfs	user	AD2\Admin1		/fs/dbvar/usedbvar/dfs	0	0	-	200	0%	0	-	N	N

How to Run Cluster Storage Report On Demand

1. Open Reports on Demand.

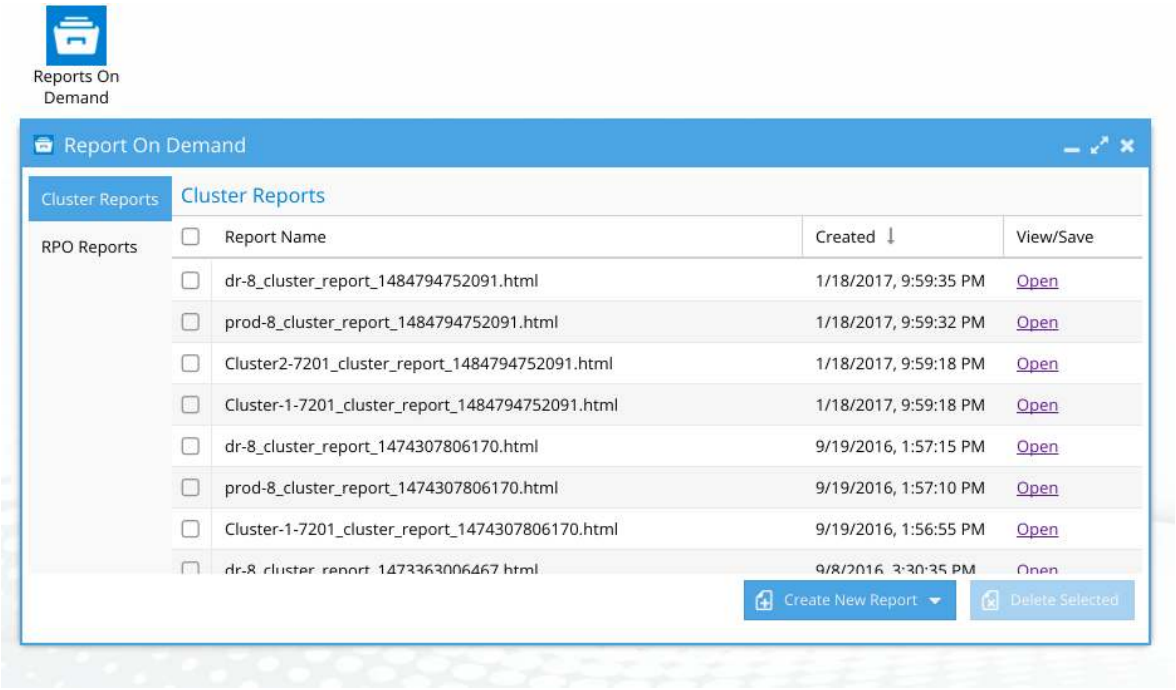

2. Select Create New Report button
3. Select CSM report option to run the report and email the results.

How to Change Cluster Storage Report Schedule

1. Using igls command below, default is daily report and runs on a daily schedule.

1. "interval": "0 0 * * *",
2. "enabled": true,
3. "id": "StorageMonitorReport",
4. "label": "Storage Monitor Report"

2. Show current schedules
    a. igls admin schedule.
3. How set new schedule see Eyeglass Administration Guide .

© Superna LLC

6.6. Configuration - IGLS CLI commands to configure Cluster storage monitor features

IGLS CLI commands to configure Cluster storage monitor features

CLI Guide

© Superna LLC

6.7. Unlock My files Help Desk Application

Home Top

- What's new
- Read Me First
  - Please read Notes below to understand more about locked files
- Known Issues
- The Use Case
- Requirements
- How to add Unlock My Files Icon permission to the admin user role
- How to delegate Unlock My Files to the Help Desk
- How to find Locked files and Break locks

What's new

1. 2.5.6 update 2 includes a completely new tool for unlocking files.
2. Partial results feature will search each access zone and return results from each access zone as they are received.
3. A cluster selection requires 1 and only 1 cluster to be selected for searching for open and locked files. This limits the time need to complete the search.

4. Search progress bar indicates the search is still executing and shows which access zones have pending results.

Read Me First

Users leave files open and locked. This requires finding the PowerScale node with the open file and issuing a command to break the lock. This is a time consuming process. This feature is designed to allow a help desk to securely find and break locks on files without any permissions on PowerScale.

Please read Notes below to understand more about locked files

1. **This feature only displays files that are locked and opened with write permissions. Applications that open for read only do not place locks on files. example notepad, wordpad do not lock files**
2. **NOTE: Very important information about open files, not all applications use persistent file handles and many applications do not lock files ever. If a file does not show up in the results it likely means the application does not use persistent file handles and does not lock files. This is application dependant.**
3. **NOTE: This feature is not intended to list or report on all open and locked files. It is intended to search for a specific open file.**
4. **NOTE: Only 100 files are returned per search**

Known Issues

1. See Release note for releases before 2.5.6 update build < = 200158
2. Service account user used to add the cluster to eyeglass must support sudo command over ssh

The Use Case

1. Unlock my files! is a new desktop icon with a dedicated permission to allow delegation to a help desk to unlock files for end users.
2. Search for a file based on name or partial string match
3. Find the file in the list
4. Click break lock button
5. Done.

Requirements

1. Requires Onefs 8.x or later.
2. Upgrade to 2.5.6 update 2 build > 200158 , major enhancements in this build. Follow the upgrade guide to get upgraded.
    a. Special character handling in searches
    b. Partial results returned per access zone scan allows faster results to be returned while other access zones are being scanned for open files
    c. Per cluster search allowing selection of a specific cluster to search
    d. Help text on how to search efficiently is now in the GUI
3. MANDATORY STEPS BEFORE USE See guide with sudoer permissions changed required

How to add Unlock My Files Icon permission to the admin user role

1. Login as admin
2. Open User Roles icon
3. find the file lock portal permission and enable it.
4. Save admin role
5. 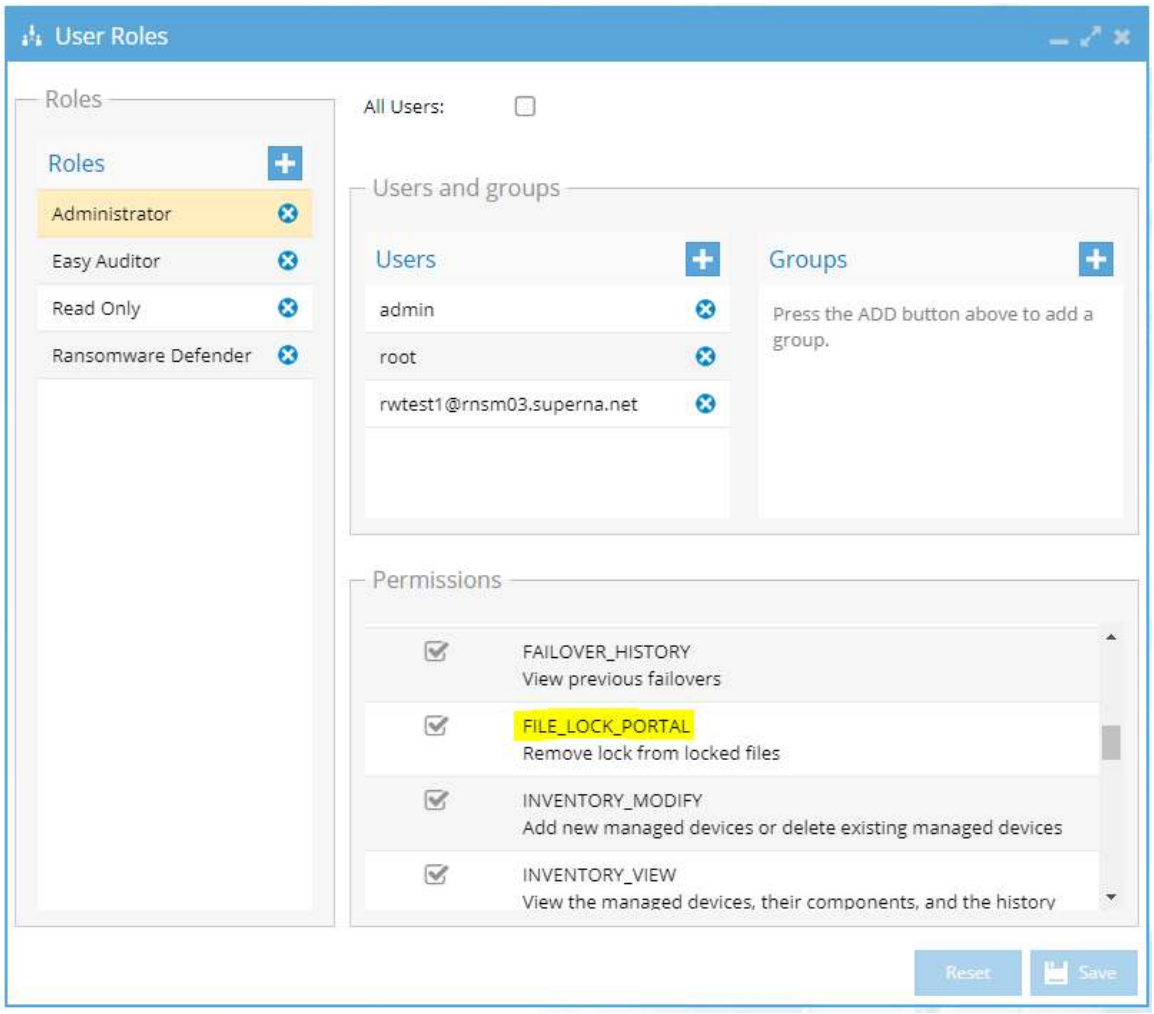

How to delegate Unlock My Files to the Help Desk

1. Please read the Full RBAC Guide for trouble shooting and configuring RBAC.
2. The RBAC feature is required to provide AD login and delegation to a Help desk to unlock files.
3. **Prerequisites**
    a. Create an AD group in Active Directory for this role
        i. **NOTE: group name cannot have spaces or special characters, the group name should be all lower case.**
4. Open the User Roles icon
    a. Create new role in User roles called "Help desk" with the + sign
    b. Assign the AD group to this role
        i. **NOTE: syntax of AD group MUST be DOMAIN\lowercasegroupname (where domain name is uppercase, and ad group name is lower case)**
    c. Add the **file lock portal** permission to this role and save the role.

d. 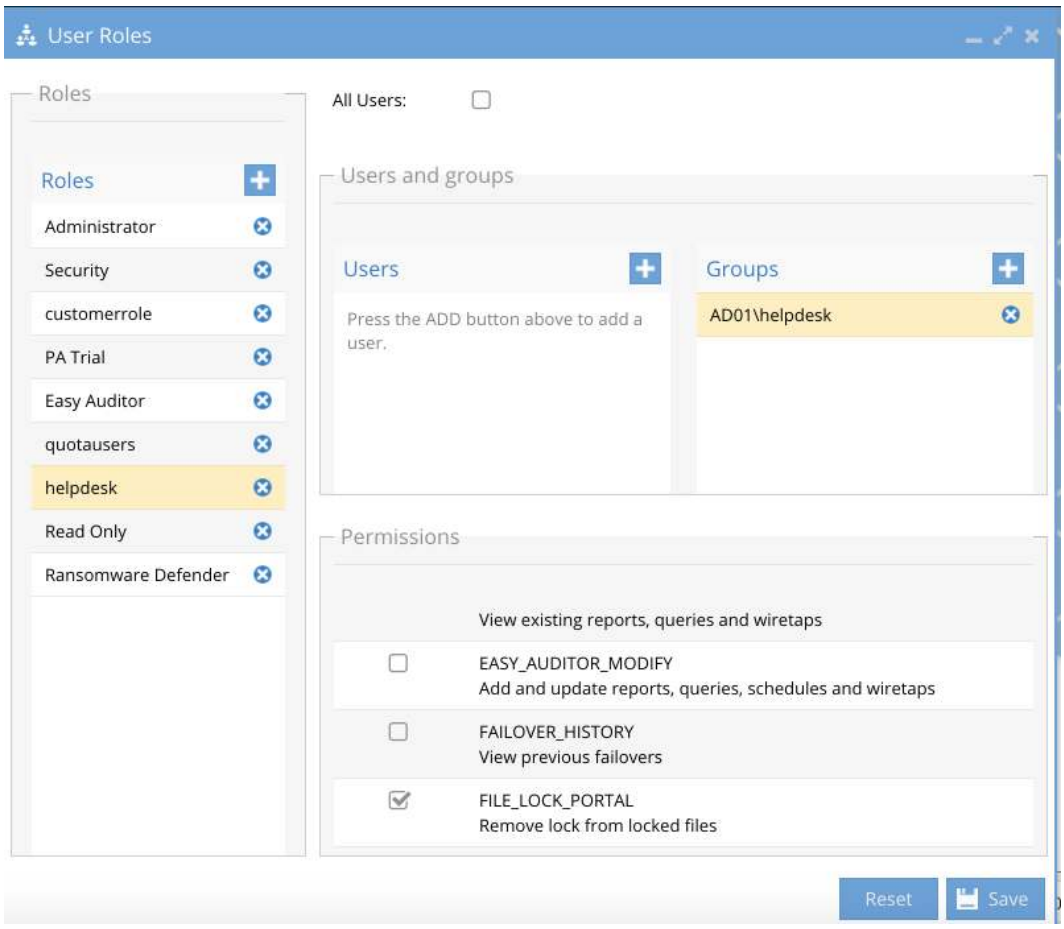

5. Click Save

How to find Locked files and Break locks

1. Type the name of the file or partial match of the file.
2. **NOTE: do not enter access zone name, do not enter special characters for the path , only enter text for the path, or exact file name. Example do not enter /ifs/data/marketing enter marketing to find open files in folders named marketing.**
3. **See the example below searching for a file with db in the extension. The red circle indicates how many Access zones exist and which zones have searches completed. Partial results per access zone will be returned the user interface.**

4. 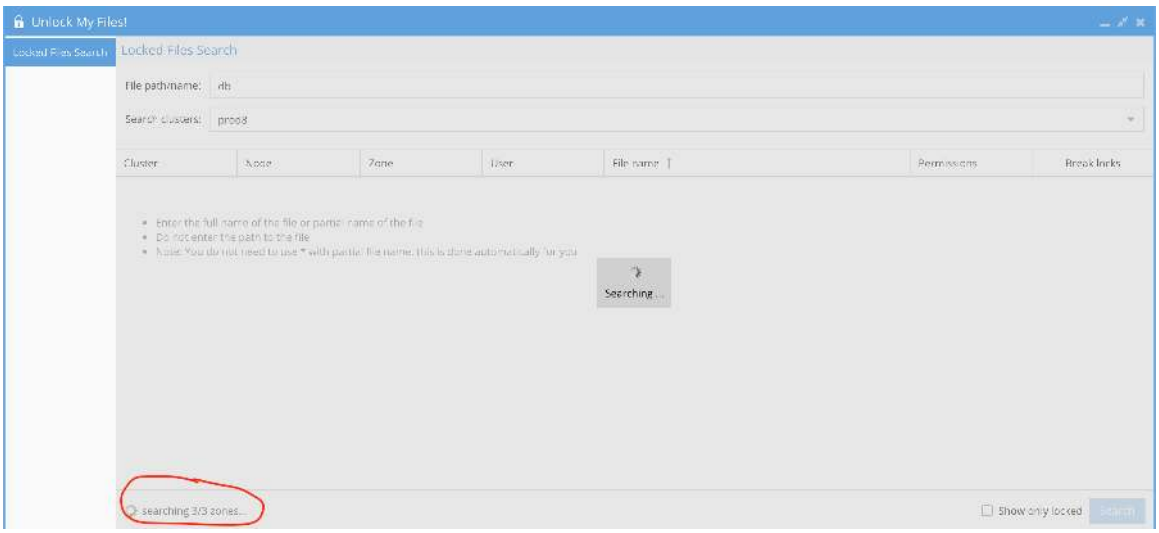

5. Once all results for all access zones are completed the GUI will indicate a completed status. See example below.

   a. 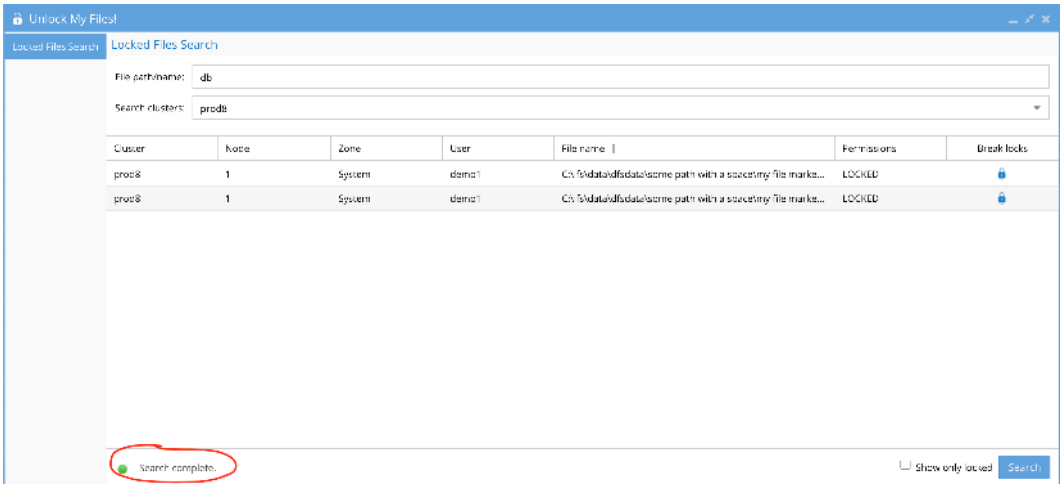

6. Best Practice:

   a. Only enter the name of the file, or text of the directory the file is located.

   b. Using 2.5.6 update 2 and later. Files with spaces can be entered exactly as the file name looks with spaces. No need for quotes to search. See example.

i.

c. **example path is /ifs/data/groups/marketing/project123.docx search for** project123.docx **or enter marketing to search for files open in the marketing folder**

7. Results will return all matching files and directories

   a. **NOTE: Only 100 files are returned, if the file is not listed a more accurate file match search is required.**

   b. **NOTE: The locked column will shows files with a SMB file lock applied the client application. If it does not display locked status it means the application did not open the file with a lock request.**

8. **Use the show only locked files filter to view only files with lock requests on the files**

9. Click the break lock icon next to the file. Answer yes to the break lock when prompted.

   a. **NOTE: Breaking a lock on an application can lose data if the application has not committed the data to the file.**

10. Verify with user that the file is unlocked and attempt to re-open the file.

© Superna LLC

7. RPO Trending and Reporting Guide

Home Top

- Overview
- Getting Started
- RPO Calculations:
- RPO Summary & Compliance Email Report
- Example Report:
- The Charts
- Charts Uses:
- Advanced Settings
- RPO CSV Reports

© Superna LLC

7.1. Overview

Overview

This guide provides information on Eyeglass PowerScale Edition Recovery Point Objective (RPO) Trending and Reporting. How to set up RPO Trending and Reporting, reporting calculations, definitions, charts and outputs, and advanced settings.

RPO Monitoring What's New

As of Release 1.9, Eyeglass now includes:

- RPO Reporting and backup monitoring report now track failed SyncIQ jobs and the report includes 30 day and last 24 hour view out per policy and per cluster of all failed SyncIQ jobs that started but finished with an error code.
- This allows the report to track the percentage of completed SyncIQ jobs when used in a backup monitoring solution.

As of Release 1.8, Eyeglass now includes:

- On-demand report generation CLI and GUI option.
- Total GB transferred in the reporting period as well as Avg GB transferred per Job in the reporting period.
- Advanced Settings for Report Screenshot Enable/Disable, Report Time Range, Transfer Rate Troubleshooting Threshold, Interval Troubleshooting Threshold.

As of Release 1.6, Eyeglass now includes:

- All new features are automatically enabled after the upgrade:
- Per SyncIQ RPO reporting in the reports.
- CSV file with per cluster and per policy data included for including in Excel or other reporting tools.
- Automatic 30 Day rolling average per policy data loss in minutes graph PNG files attached to the email report for simple review and inclusion into reports or PowerPoint.

The solution now allows simple reporting for business units where a SyncIQ policy or more than one policy is specific to a business unit that requires SLA on the DR service of the files they consume on a cluster.

The new section to this guide has been added with examples of the output of the Per Cluster and Per Policy CSV and images that are automatically generated now every night.

Recovery Point Objective Key Features

With the Eyeglass PowerScale Edition Recovery Point Objective (RPO) Trending and Reporting feature you will have the data to answer these questions:

- Am I meeting my RPO business target on my PowerScale cluster?

The Eyeglass Solution: A daily email with your RPO business target compliance by cluster over the last 24 hours and the last 30 days in a simple and easy to read summary report.

- Are my SyncIQ Policies performing as expected?

The Eyeglass Solution: A daily email with SyncIQ policies by cluster where job interval or transfer rate does not meet historical performance over the last 30 days.

- How do I optimize my SyncIQ schedule to lower my RPO to match my WAN bandwidth and data change rate?

The Eyeglass Solution: Tune your SyncIQ Replication based on tracking your data change rate with deep-dive graphing that allows the time of day, week, month trending to assist with lowering your replication schedule or increasing it to achieve shorter replication cycles with SyncIQ.

- How far back can I recover my business data?

The Eyeglass Solution: Graph and trend your business's recovery point graphed as the age of data in minutes in the past per SyncIQ policy.

- How much bandwidth is used and when is it used?

The Eyeglass Solution: Graph and trend GB data transferred per SyncIQ Policy to help with WAN bandwidth planning and quality of service at the network lawyer to ensure your critical business data has adequate bandwidth to meet your RPO objectives.

© Superna LLC

7.2. Getting Started

Home Top

Getting Started

Follow these 3 easy steps to get RPO Trending and Reporting up and running on your Eyeglass appliance:

1. Check if RPO Trending and Reporting License is installed.
2. Setup RPO target by Cluster.
3. Setup Email Notification for daily reports to be enabled.

Check if RPO Trending and Reporting License is Installed

Eyeglass PowerScale Edition RPO Trending and Reporting requires a separate feature license. Open the Manage Licenses window to check your licenses. If you see license type “PowerScale RPO Reporting” you are licensed for this feature.

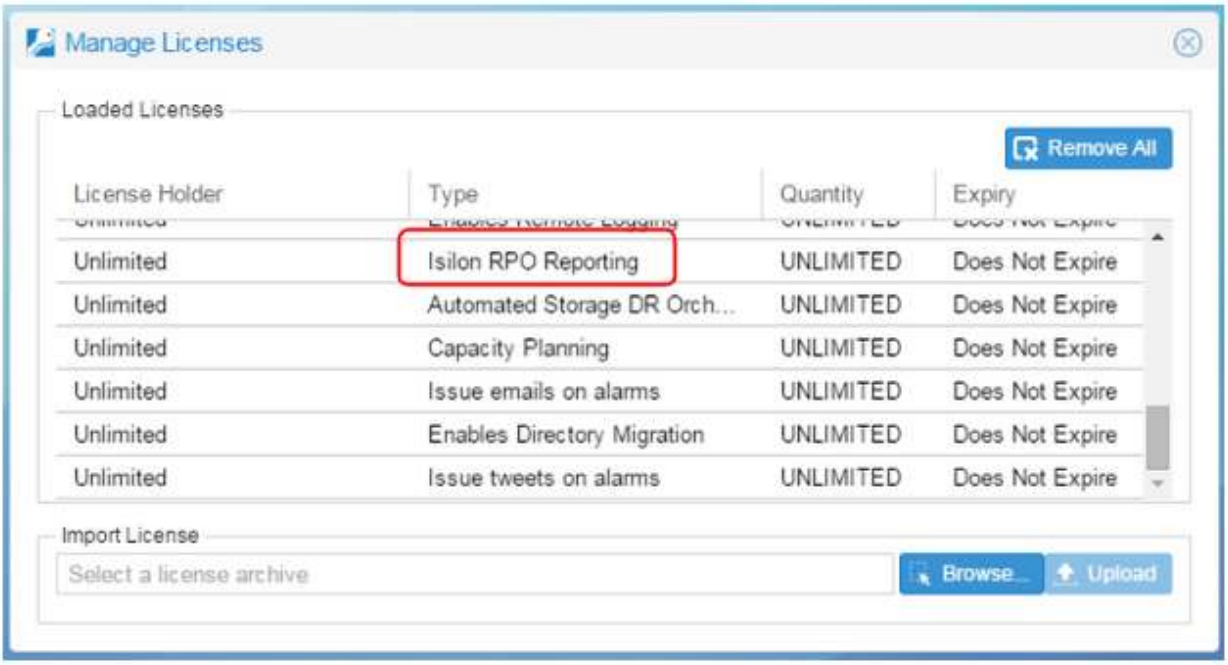

If you do not have the PowerScale RPO Reporting license, please contact your Eyeglass sales representative.

Setup RPO Target by Cluster

To do the analysis, Eyeglass requires you to enter your RPO target by cluster in minutes. This is the target RPO that your company would like to achieve. Eyeglass will calculate the actual RPO achieved, and provide a comparison to the target entered here in daily emailed reports.

Note: This will be an average RPO for the entire cluster.

For a new cluster, the RPO is entered on the Add PowerScale Cluster window in the Maximum RPO Value field. Enter the RPO target in minutes.

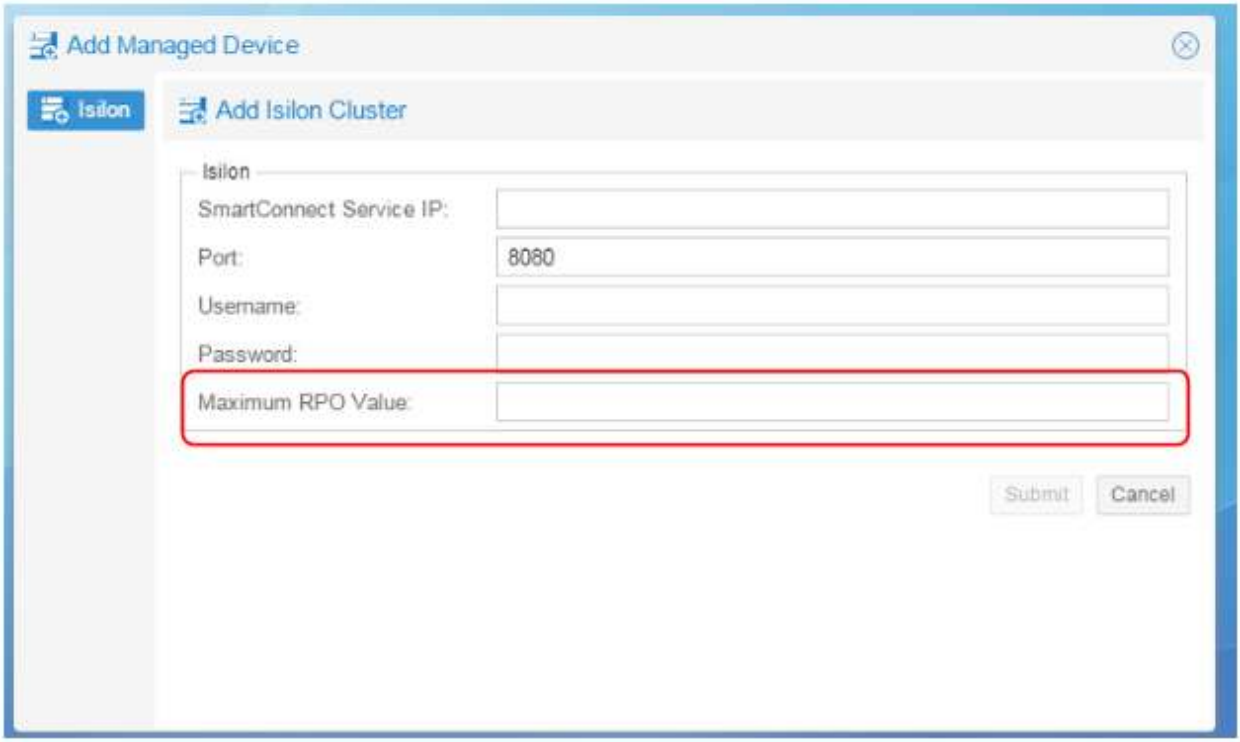

To update the RPO value for an existing Cluster, open the Inventory View window, right-click on the cluster you want to update, and select Edit.

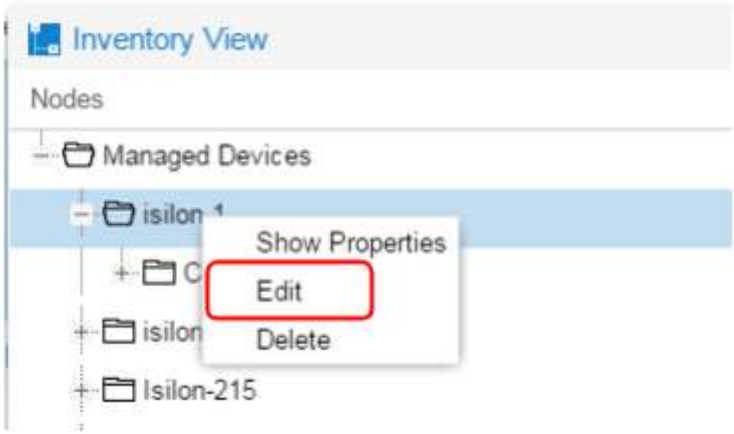

In the Edit PowerScale Cluster window that opens enter the new RPO target in minutes in the Maximum RPO Value field and then Submit to save your changes.

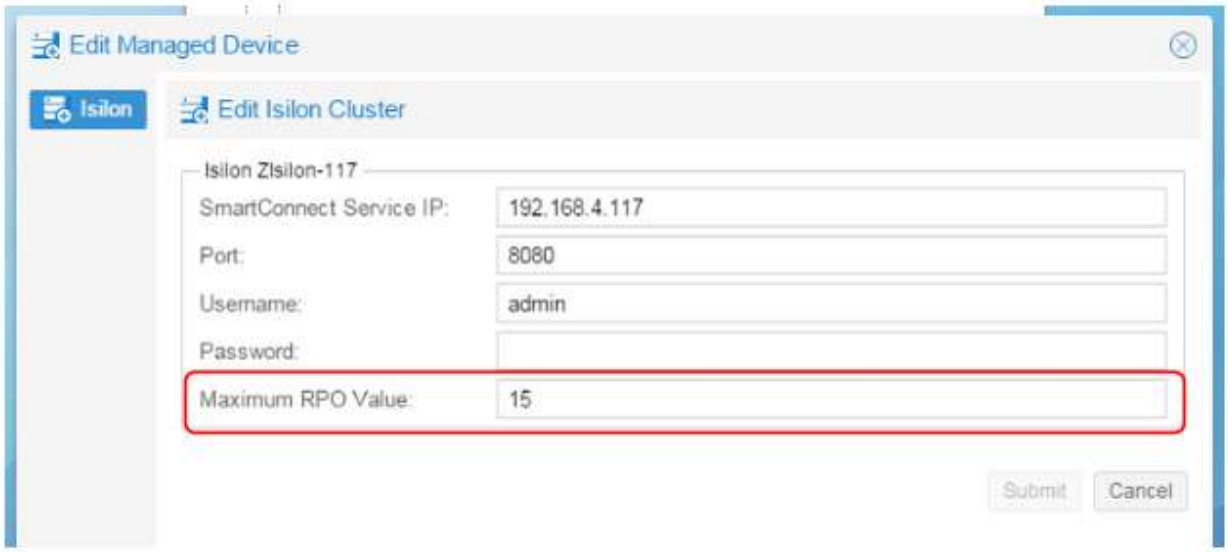

Setup Email Notification

To receive the daily RPO compliance email, your Eyeglass appliance must have an email server configured and email recipient email

addresses configured. Please refer to the Eyeglass PowerScale Quickstart Guide for details on setting up Email Notification.

© Superna LLC

7.3. RPO Calculations:

Home Top

- General Notes
- Maximum Age of Unreplicated Data
- Job Duration
- Amount of Replicated Data
- Recovery Point Analysis
- Average Data Transfer Rate in Mb per second
- SyncIQ Jobs Troubleshooting

General Notes

Data for RPO calculations are based on the SyncIQ Job Reports that OneFS generates each time a SyncIQ Policy Job is run. These reports are collected as follows:

- When a PowerScale cluster is provisioned in Eyeglass, the last 10 reports for each SyncIQ Policy are collected.
- Once every 5 minutes, the last 10 reports for each SyncIQ Policy are collected.
- Reports for failed jobs are not collected and are not included in the statistics.
- For a canceled SyncIQ Job, the RPO calculation starts when the canceled Job was started and ends when the next successful Job for that SyncIQ Policy is completed.

- If a SyncIQ Policy is deleted in OneFS, Eyeglass summary statistics include reports for the deleted policy until they are no longer relevant for the reported time frame.

Maximum Age of Unreplicated Data

The last completion time of a SyncIQ policy does not represent your recovery point. The “age” plot assumes the change rate of the data was the same as the last SyncIQ policy report and, this data is in flight but not yet successfully replicated.

Example: Last successful completion of a SyncIQ policy was 10 minutes ago and, the job took 5 minutes to replicate.

For this example, the Maximum Age of Unreplicated Data is calculated as the 10 minutes since the Job was last successfully completed + the 5 minutes it would take to replicate the data (assumed to be the same as last time a Job was successfully executed). So, in this example, a DR event would result in 15 minutes of data being lost.

Job Duration

The Job Duration reported in Eyeglass is the Duration that is reported in the OneFS SyncIQ Report for a policy Job.

Amount of Replicated Data

The Amount of Replicated Data reported in Eyeglass is the Total Data that is reported in the OneFS SyncIQ Report for a Policy Job.

Recovery Point Analysis

Using the target RPO and average SyncIQ Job duration, data transfer rate and data change rate for a cluster, Eyeglass will calculate a recommended SyncIQ Job interval to reduce the number of Jobs that violate the target RPO.

Average Data Transfer Rate in Mb per second

The Average Data Transfer Rate reported in Eyeglass is the Total Data that is reported in the OneFS SyncIQ Report for the Cluster Policy Jobs divided by the total duration of all Jobs for the reported time period.

SyncIQ Jobs Troubleshooting

SyncIQ Policy Interval and Transfer rate are assessed against the average SyncIQ Policy interval and transfer rate over the last 30 days.

Note: Interval is calculated as the difference between the start time of 2 consecutive SyncIQ Jobs. Interval may not be the same as the schedule, for example, if a Job takes longer to run than its schedule.

The following anomalies are reported in the daily email report by policy:

- The number of jobs where the transfer rate is at least 50% below the average transfer rate (the measured difference factor - ie 50% - is configurable - please refer to the Advanced Settings section of this document for details).
- The number of SyncIQ Policy jobs that did not run in the last 24 hours, that should have been run based on the 30-day average interval.

- The number of jobs where the interval is greater than double the average interval (the measured difference factor - ie 2x or double - is configurable - please refer to the Advanced Settings section of this document for details).

© Superna LLC

7.4. RPO Summary & Compliance Email Report

RPO Summary & Compliance Email Report

Prerequisites:

- RPO target configured for each cluster managed by Eyeglass.
- Eyeglass email notification configured.
- PowerScale RPO Feature License.

How to Generate the RPO on Demand and Specify Time Period

Follow these steps to generate a report on demand for the current reporting period that is 24 hours or change the data range that the report covers and have the report emailed.

Generate Report from the CLI

1. Follow the Eyeglass Administration Guide to execute report generation from CLI commands.
2. See RPO Reporting CLI commands in the Eyeglass Administration Guide.

Generate on UI and Change the Reporting Period

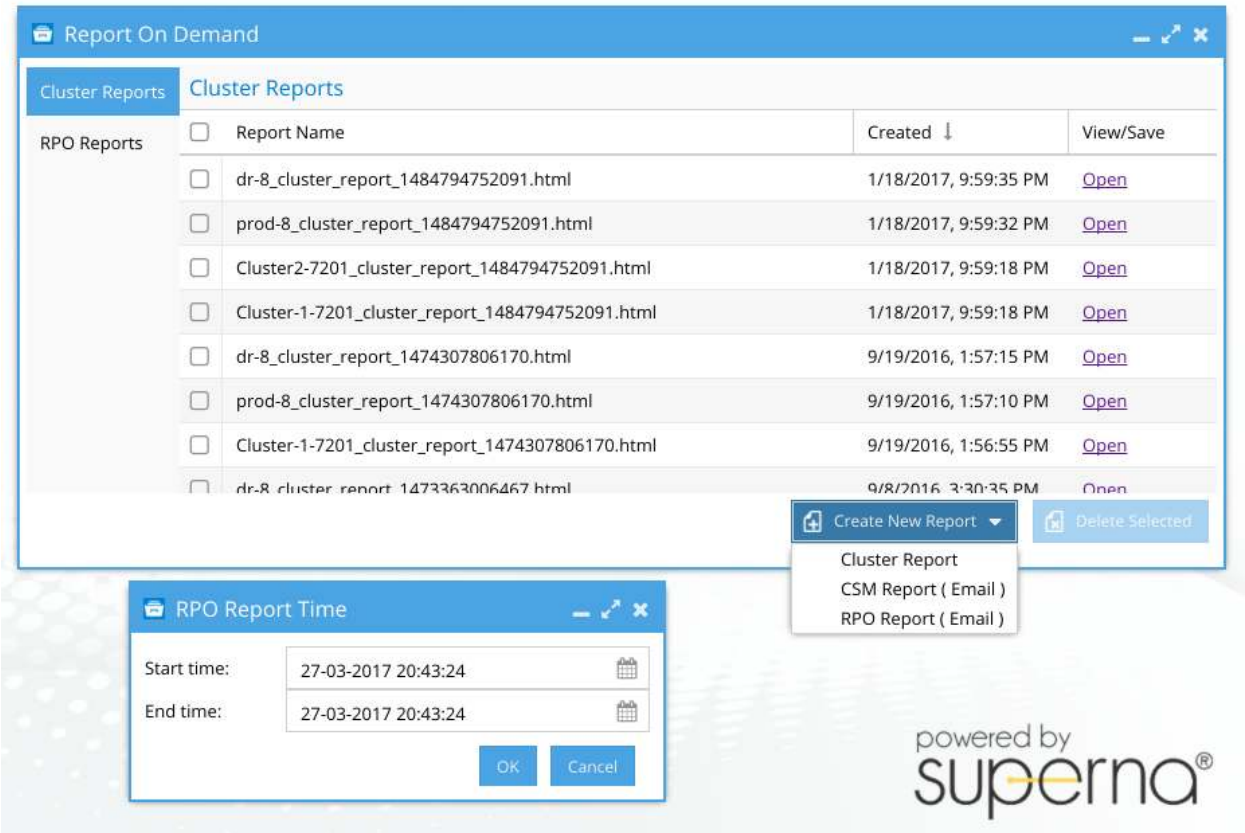

Details

The RPO Summary and Compliance Email Report is sent out each day at midnight. It contains the following information for each cluster managed by Eyeglass:

Field	Description
RPO	The RPO target in minutes as entered into Eyeglass for the cluster.
Number of Jobs	The total number of SyncIQ Policy Jobs that ran over the previous 24 hours and the previous 30 days by cluster.
Total Amount of Replicated Data in GB	The total amount of replicated data in GB per cluster. This statistic is provided based on the total number of SyncIQ Jobs that ran over the previous 24 hours and the previous 30 days for each cluster managed by Eyeglass.

Percentage of Jobs Violating RPO	Percentage of SyncIQ Jobs where the maximum age of the un-replicated data has been calculated to be greater than the RPO target for the cluster. This statistic is calculated based on the total number of SyncIQ Jobs that ran over the previous 24 hours and the previous 30 days for each cluster managed by Eyeglass.
Average Amount of Replicated Data in GB	The average amount of replicated data in GB per SyncIQ Job. This statistic is provided based on the total number of SyncIQ Jobs that ran over the previous 24 hours and the previous 30 days for each cluster managed by Eyeglass.
Average Job Duration in minutes	Average Job Duration in minutes. This statistic is provided based on the total number of SyncIQ Jobs that ran over the previous 24 hours and the previous 30 days for each cluster managed by Eyeglass.
Average Data Transfer Rate in Mb per second	Average Data Transfer Rate in Mb per second. This statistic is provided based on the total number of SyncIQ Jobs that ran over the previous 24 hours and the previous 30 days for each cluster managed by Eyeglass.
Recovery Point Analysis - Diagnostics	Diagnostics provides a recommendation for SyncIQ Policy schedule such that the RPO target for the cluster can be better met.
SyncIQ Job Troubleshooting	Summary of any SyncIQ Policies where the Job interval or transfer rate falls below historical average over the last 30 days. This may indicate a networking issue or cluster resource on the cluster that is impacting replication performance. Adding more nodes to the SyncIQ pool for replication or worker threads should be increased. Jobs that failed to run in the last 24 hours (no job report) are also captured on this list and should be investigated

	on the cluster to see why the job failed to run.

© Superna LLC

7.5. Example Report:

Home Top

Example Report:

SyncIQ Jobs Report 2016-12-16 00:00:00 UTC

Number of Jobs

Cluster	last 24 hours	last 30 days
ds-sim-8-1	12	23
ds-sim-8-2	0	21

Total Amount of Replicated Data in GB

Cluster	last 24 hours	last 30 days
ds-sim-8-1	less than 0.01	less than 0.01
ds-sim-8-2	no data	less than 0.01

Percentage of Jobs Violating RPO

Cluster	RPO	last 24 hours	last 30 days
ds-sim-8-1	3	100	43
ds-sim-8-2	5	no data	23

Average Amount of Replicated Data in GB

Cluster	last 24 hours	last 30 days
ds-sim-8-1	0.00	0.00
ds-sim-8-2	no data	0.00

Average Job Duration in minutes

Cluster	last 24 hours	last 30 days
ds-sim-8-1	0.18	0.13
ds-sim-8-2	no data	0.14

Average Data Transfer Rate in Mb per second

Cluster	last 24 hours	last 30 days
ds-sim-8-1	0.00	0.00
ds-sim-8-2	no data	0.00

Recovery Point Analysis

Cluster	Diagnostics
ds-sim-8-2	
ds-sim-8-1	Policy accesszone1: setting to run job at the interval 2 minutes per 24 hour period will lower the RPO violation rate.

Refer to RPO Trending and Reporting in the Eyeglass PowerScale Edition documentation for an explanation of the calculation of diagnostics.

SyncIQ Jobs Troubleshooting

ds-sim-8-2

Policy Name	Job ID	Detected Problems
EyeglassRunbookRobot_mirror		No SyncIQ jobs have been run for the last 24 hours. The average 30 days interval between jobs is 320.14minutes.

ds-sim-8-1

Policy Name	Job ID	Detected Problems
test		Found 43 jobs that have the transfer rate lower than the policy average rate 193.92..

Refer to RPO Trending and Reporting in the Eyeglass PowerScale Edition documentation detail on SyncIQ Jobs Troubleshooting

** average transfer rate in Mb/s

© Superna LLC

7.6. The Charts

The Charts

Prerequisites

- RPO target configured for each cluster managed by Eyeglass.
- PowerScale RPO Feature License.

SyncIQ Job RPO Chart

The Eyeglass Job Duration Chart plots the Maximum Age of Un-replicated Data per SyncIQ Job successfully run in minutes over time, and on the same chart, the Job Duration for each SyncIQ Job run. Up to 6 SyncIQ Policies can be selected to be graphed on the same chart.

To generate the Data Transfer Chart:

1. Login to the Eyeglass web page.
2. Open the **DR Dashboard.**
3. Select the checkbox for the policy of interest.
4. For a multi-Job chart, up to 5 additional policies can be selected.
5. Select the **Generate SyncIQ Job Charts** button.
6. Select the **From** and **To** date and time in the Report Time Range Setting window.
7. Select the **Launch SynciQ Job RPO Chart** button.
8. The **SyncIQ Job RPO Chart** opens.

- Maximum age of Unreplicated Data plotted in minutes corresponding to each SyncIQ Job successfully executed (diamond marker).
- Job Duration plotted in minutes corresponding to each SyncIQ Job successfully executed (circle marker).
- Each policy plotted in a different color.
- Mouse over a data point to see the details.

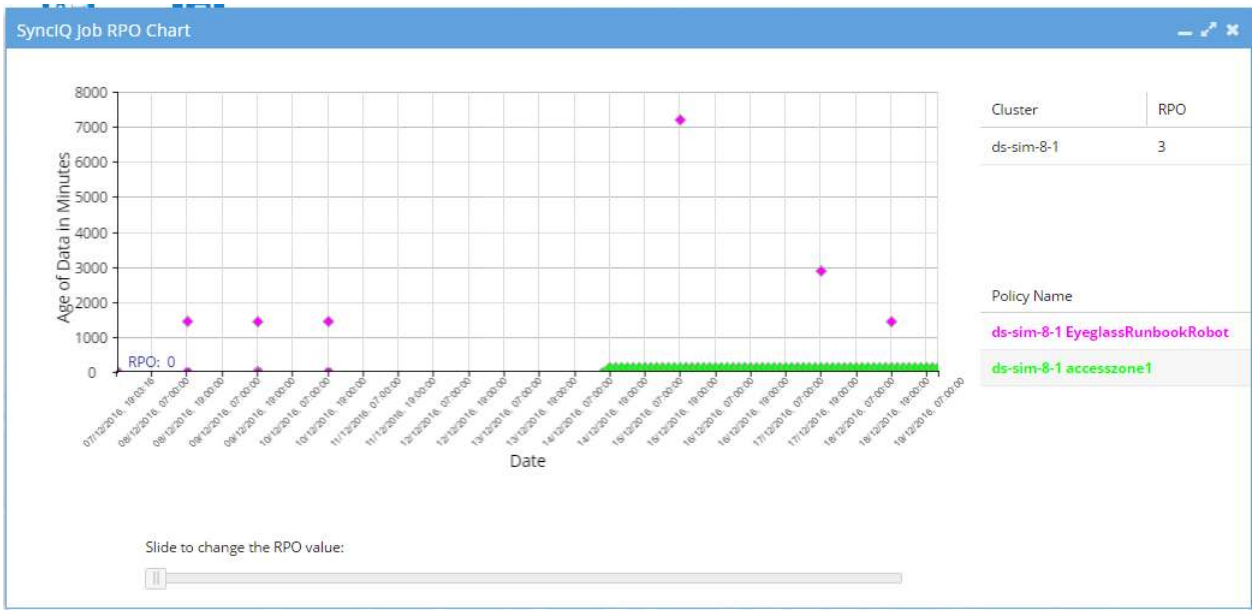

9. Select a cluster on the right-hand table to add the RPO target for that cluster to the graph. Points above the RPO target have not met the target and are displayed in red. You can use the RPO slider at the bottom of the chart to see the effect of changing RPO target.

Data Transfer Chart

The Eyeglass Data Transfer Chart plots the total amount of data replicated per SyncIQ Job successfully run in GB over time. Up to 6 SyncIQ Policies can be selected to be graphed on the same chart.

To generate the Data Transfer Chart:

1. Login to the Eyeglass web page.
2. Open the **DR Dashboard**.
3. Select the checkbox for the policy of interest.
4. For a multi-Job chart, up to 5 additional policies can be selected.
5. Select the **Generate SyncIQ Job Charts** button.
6. Select the **From** and **To** date and time in the Report Time Range Setting window.
7. Select the **Launch SyncIQ Job Data Transfer Chart** button.
8. The **SyncIQ Job Data Transfer** chart opens.

- Total data transferred is plotted in GB for each job executed for the selected policies in the selected time period.
- Each policy plotted in a different color.
- Mouse over a data point to see the details.

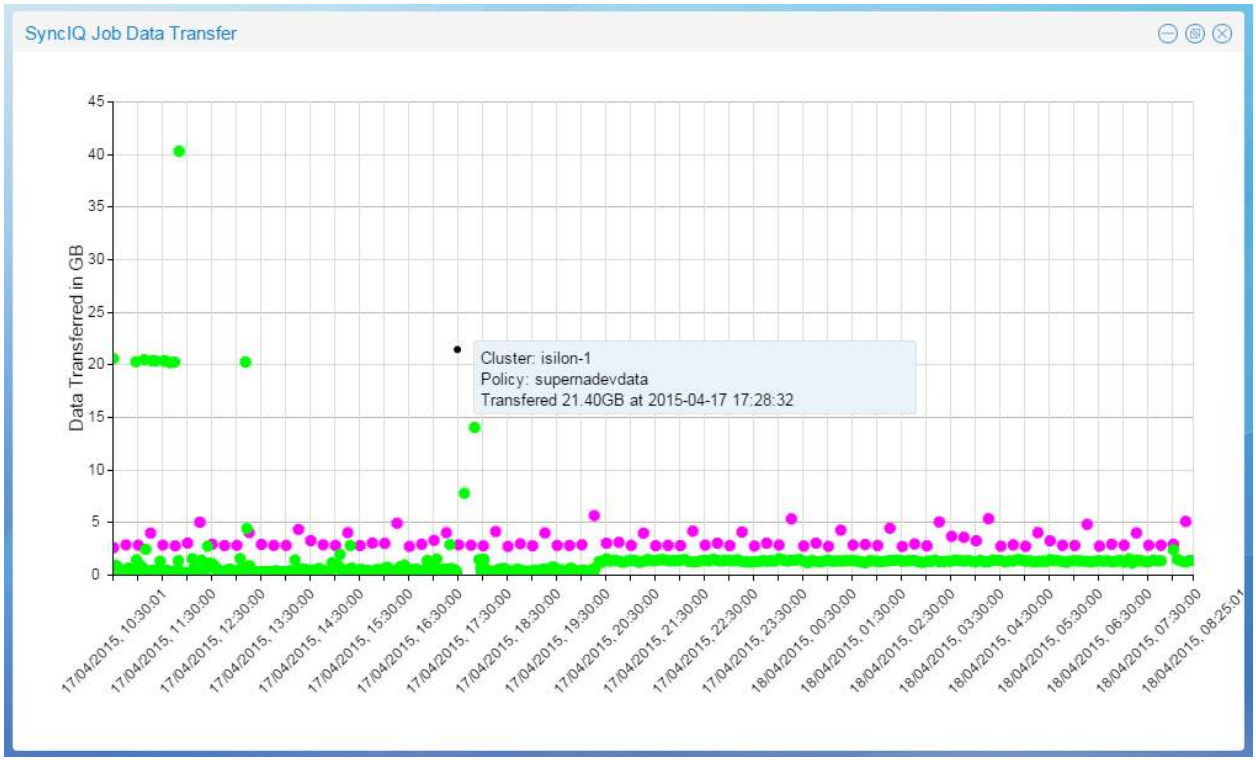

© Superna LLC

7.7. Charts Uses:

Charts Uses:

Am I Meeting my RPO Business Target on my PowerScale Cluster?

Use the Eyeglass daily RPO compliance email to quickly and easily gain a view of whether or not you are meeting your business targets. For each cluster that Eyeglass is managing, you will receive an email comparing the last 24 hours to the last 30 days. To get the details you can use the Eyeglass graphs to find out why RPO times might be increasing outside your targets.

The summary report provides the percent of SyncIQ jobs that failed to meet the objectives, and which provides a quick summary of your targets per cluster.

Are my SyncIQ Policies Performing as Expected?

Use the Eyeglass daily email to quickly and easily gain a view of whether your SyncIQ Policies are performing as expected. The SyncIQ Jobs Troubleshooting section will highlight by cluster:

- Which SyncIQ Policies did not run that should have run in the last 24 hours.
- Which SyncIQ Policies had jobs where the transfer rate was lower than the policy average rate over the last 30 days.
- Which SyncIQ Policies had jobs where the interval is greater than the policy average interval over the last 30 days.

How far back can I Recover my Business Data?

Use the Eyeglass Job Duration Chart to analyze the maximum age of your data per SyncIQ policy over a selected time period. The maximum age number is calculated based on the last time the data was successfully replicated to the remote cluster, plus the time it would take to replicate the same amount of data.

This chart can be customized:

- Start and end date and time.
- 1 to 6 SyncIQ policies on the same graph.
- Moving the RPO scale to show which days and time the SyncIQ data "Age" exceeded your cluster target. This is shown as red dots above the dotted line.
- To ensure 100% of all change rate data is below your target, and to understand your worst-case RPO value within the time period, slide the RPO slider to the right until no Red plots exist. This value is now the worst-case data age exposure.

How Much Bandwidth is Used and When is it Used?

Use the Eyeglass Job Data Transfer Chart to analyze the maximum age of your data per SyncIQ policy over a selected time period. This chart can be customized:

- Start and end date and time.
- 1 to 6 SyncIQ policies on the same graph.

How to Use the Average Data Transfer Rate Analysis?

This provides an overall view of the average WAN rate, in Mbps, that is required to maintain the RPO in the reports. If the goal is to lower

the RPO, this number can be used to provide the WAN or network team input on the current network load required to meet current RPO levels.

Network QOS or SyncIQ threads per node can be verified to increase network throughput, and use this report summary to track improvements in WAN throughput.

Per SyncIQ and Cluster Wide Graph and Data Examples

Example #1 of 30 day Per SyncIQ graph

Note: File name is the name of the SyncIQ policy.png

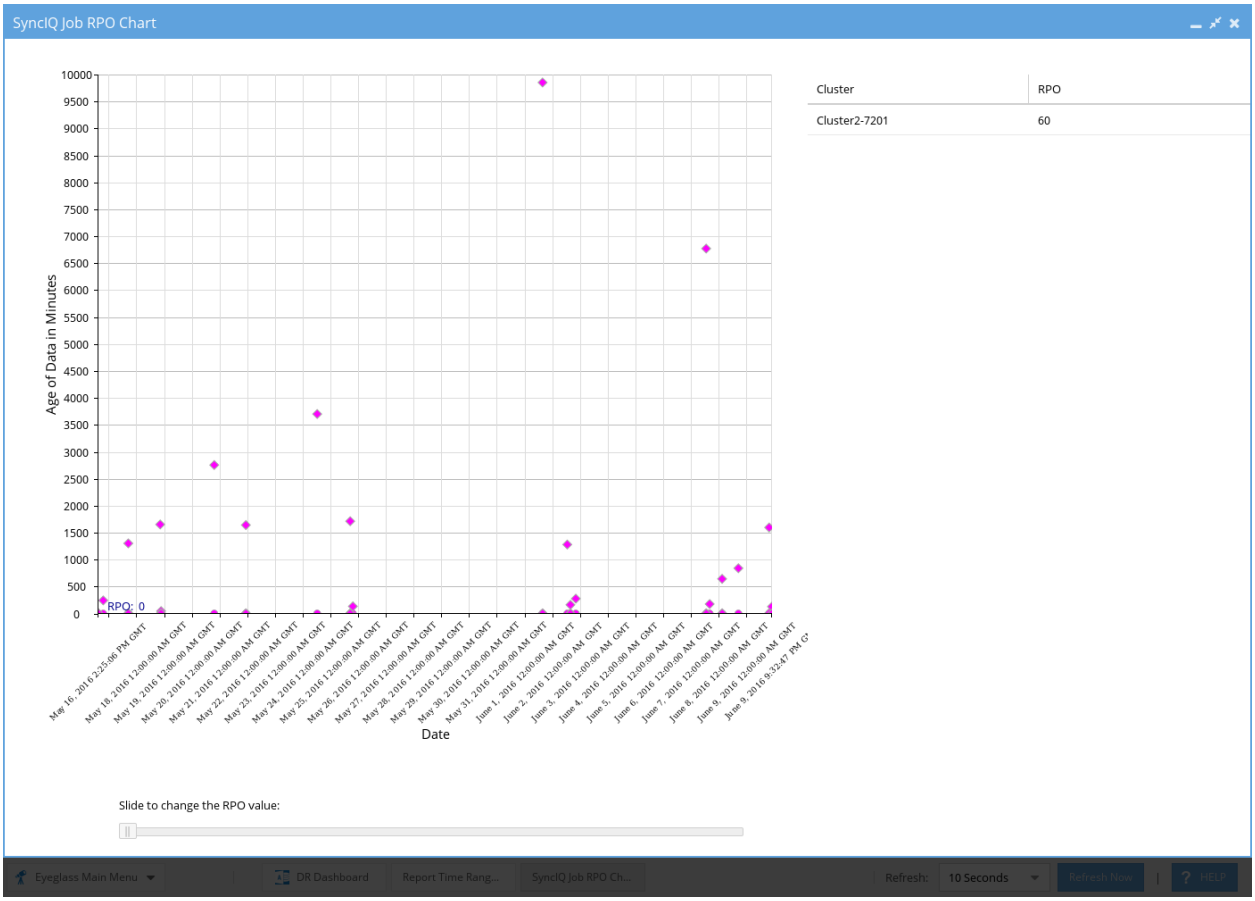

Example #2 of 30 day Per SyncIQ graph

Note: File name is the name of the SyncIQ policy.png

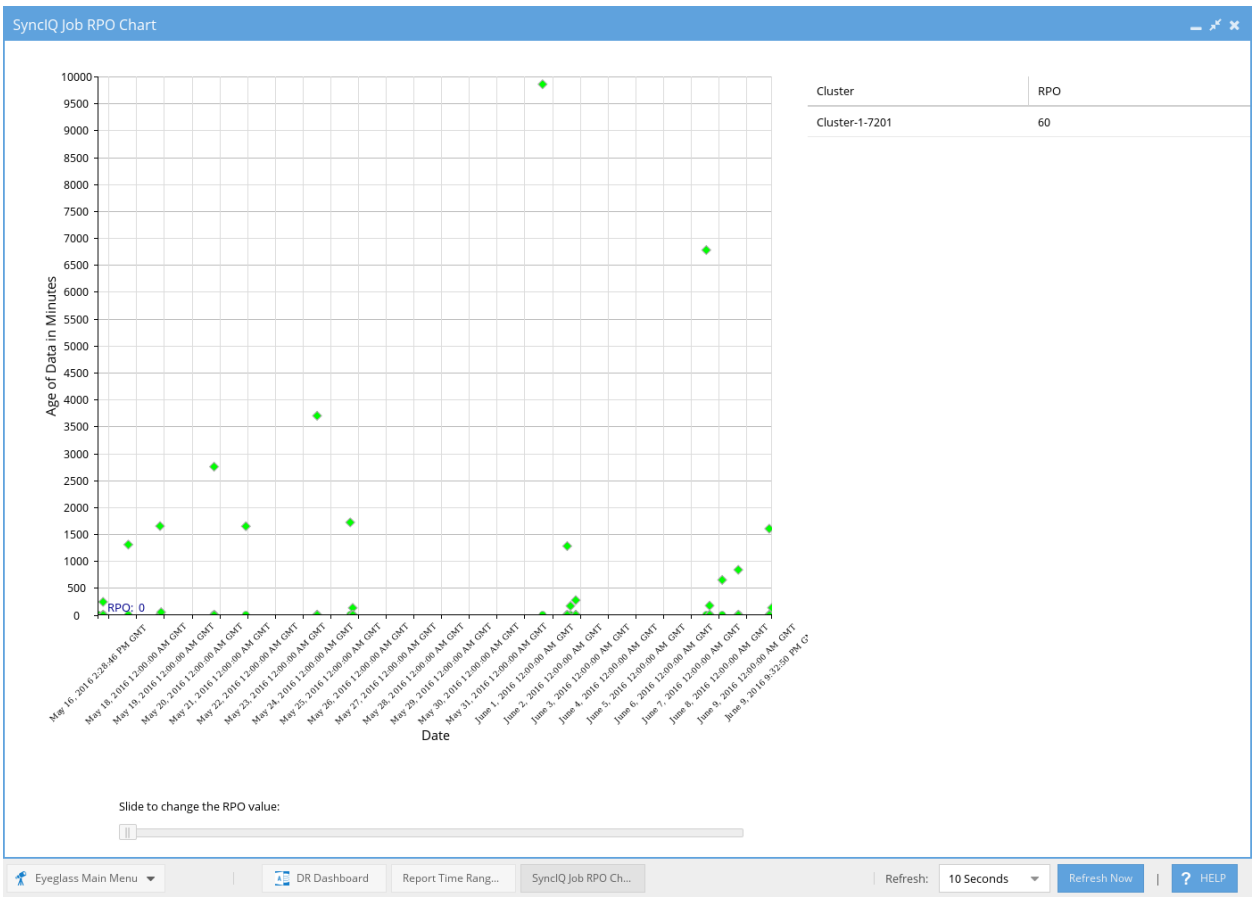

Example Cluster Wide RPO CSV file data

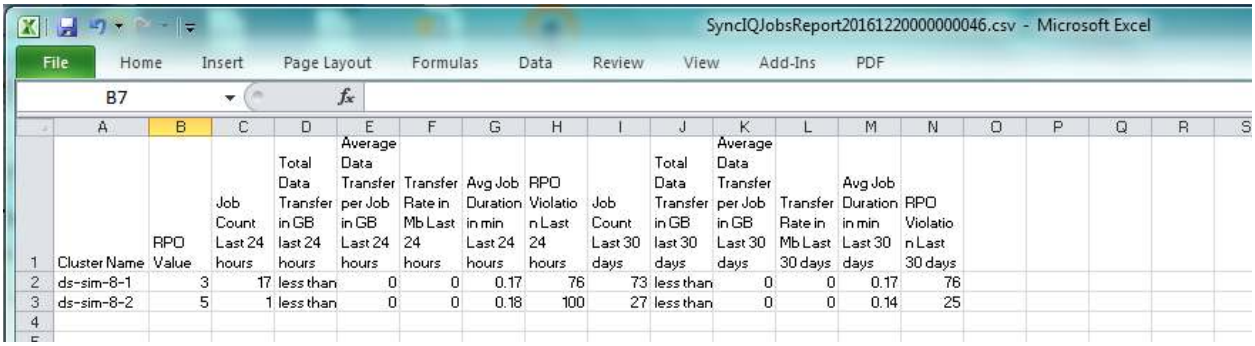

SyncIQJobsReport20161220000000046.csv - Microsoft Excel

Cluster Name	RPO Value	Job Count Last 24 hours	Total Data Transfer in GB last 24 hours	Average Data Transfer per Job in GB Last 24 hours	Transfer Rate in Mb Last 24 hours	Avg Job Duration in min Last 24 hours	RPO Violation Last 24 hours	Job Count Last 30 days	Total Data Transfer in GB last 30 days	Average Data Transfer per Job in GB Last 30 days	Transfer Rate in Mb Last 30 days	Avg Job Duration in min Last 30 days	RPO Violation Last 30 days
ds-sim-8-1	3	17	less than	0	0	0.17	76	73	less than	0	0	0.17	76
ds-sim-8-2	5	1	less than	0	0	0.18	100	27	less than	0	0	0.14	25

Example Per SyncIQ RPO CSV file data

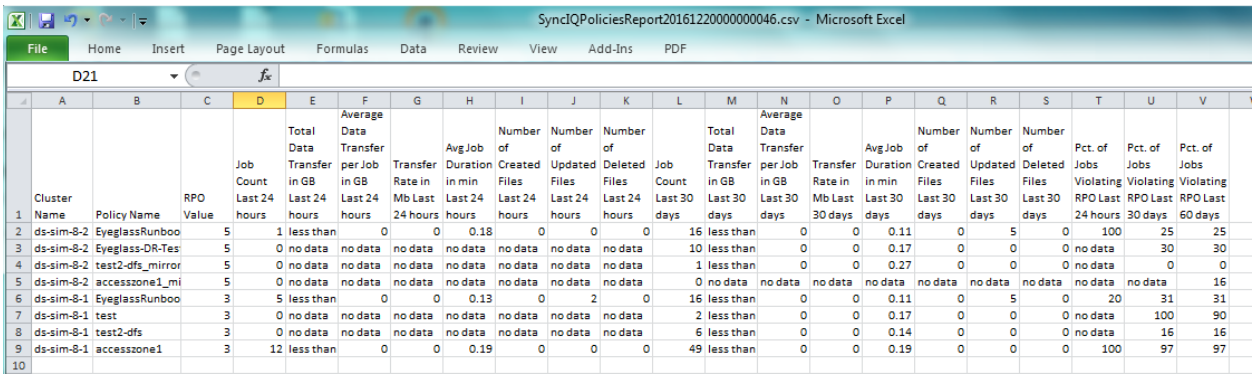

SyncIQPoliciesReport20161220000000046.csv - Microsoft Excel

Cluster Name	Policy Name	RPO Value	Job Count Last 24 hours	Total Data Transfer in GB Last 24 hours	Average Data Transfer per Job in GB Last 24 hours	Transfer Rate in Mb Last 24 hours	Avg Job Duration in min Last 24 hours	Number of Created Files Last 24 hours	Number of Updated Files Last 24 hours	Number of Deleted Files Last 24 hours	Job Count Last 30 days	Total Data Transfer in GB Last 30 days	Average Data Transfer per Job in GB Last 30 days	Transfer Rate in Mb Last 30 days	Avg Job Duration in min Last 30 days	Number of Created Files Last 30 days	Number of Updated Files Last 30 days	Number of Deleted Files Last 30 days	Pct. of Jobs Violating RPO Last 24 hours	Pct. of Jobs Violating RPO Last 30 days	Pct. of Jobs Violating RPO Last 60 days
ds-sim-8-2	EyeglassRunboo	5	1	less than	0	0	0.18	0	0	0	16	less than	0	0	0.11	0	5	0	100	25	25
ds-sim-8-2	Eyeglass-DR-Tes	5	0	no data	no data	no data	no data	no data	no data	no data	10	less than	0	0	0.17	0	0	0	no data	30	30
ds-sim-8-2	test2-dfs_mirror	5	0	no data	no data	no data	no data	no data	no data	no data	1	less than	0	0	0.27	0	0	0	no data	0	0
ds-sim-8-2	accesszone1_mi	5	0	no data	no data	no data	no data	no data	no data	no data	0	no data	no data	no data	no data	no data	no data	no data	no data	no data	16
ds-sim-8-1	EyeglassRunboo	3	5	less than	0	0	0.13	0	2	0	16	less than	0	0	0.11	0	5	0	20	31	31
ds-sim-8-1	test	3	0	no data	no data	no data	no data	no data	no data	no data	2	less than	0	0	0.17	0	0	0	no data	100	90
ds-sim-8-1	test2-dfs	3	0	no data	no data	no data	no data	no data	no data	no data	6	less than	0	0	0.14	0	0	0	no data	16	16
ds-sim-8-1	accesszone1	3	12	less than	0	0	0.19	0	0	0	49	less than	0	0	0.19	0	0	0	100	97	97

© Superna LLC

7.8. Advanced Settings

Advanced Settings

Run SyncIQ Job Report On-Demand

To run the SyncIQ Job Report On-Demand:

1. ssh to the Eyeglass appliance.
2. Login as the admin user.
3. Enter the command:

igls adv runreports

The time that the command is run is the starting time for the report and associated calculations.

Disable Screenshots for SyncIQ Job Report

To disable screenshots for the SyncIQ Job Report:

1. ssh to the Eyeglass appliance.
2. Login as the admin user.
3. Enter the command:

igls adv skipscreenshots set --skip=true

(To enable screenshots: igls adv skipscreenshots set --skip=false)

Modify SyncIQ Job Report Schedule

Standard Schedule

To change the SyncIQ Job Report Schedule to a standard schedule (1M 2M 3M 4M 5M 6M 10M 15M 20M 30M 1H 2H 3H 4H 6H 8H 12H 1D 7D 31D):

Note: Default schedule is 1D (once every 24 hours at midnight)

1. ssh to the Eyeglass appliance.
2. Login as the admin user.
3. Enter the command:

igls admin schedules set --id InventoryReport --interval <interval>

Example:

igls admin schedules set --id InventoryReport --interval 12H

Custom Schedule

To change the SyncIQ Job Report Schedule to a custom schedule (for example, to change the schedule to run at 09:00 every day).

1. ssh to the Eyeglass appliance.
2. sudo su to root user (default admin password 3y3gl4ss).
3. cd /opt/superna/sca/data
4. Make a backup of the file we are going to edit:

cp sync.xml sync.xml.bak

5. vi sync.xml
6. Update the line that starts with the tag <InventoryReport so that cron string is correct for the interval you would like the report to run at. Example below to run the report daily at 09:00:

<InventoryReport IsConfigurable="true" Label="Eyeglass Reports">0 9 * * *</InventoryReport>

7. Save your changes.
8. Restart the Eyeglass sca service:

systemctl restart sca

Modify SyncIQ Job Troubleshooting Thresholds

SyncIQ Job Troubleshooting Thresholds are the factors used when comparing SyncIQ Job Report data to the 30-day average data to determine whether a Troubleshooting notification should be posted. There are 2 thresholds configured:

1. Transfer Rate Threshold

The 24 hour Transfer Rate Troubleshooting notice is posted when the 24 hour Transfer rate is less than the 30 day Average Transfer Rate / Transfer Rate Threshold. By default, the Transfer Rate Threshold is 2. Thus SyncIQ Job Troubleshooting notice is posted when the 24 hour Transfer Rate is less than 50% of the average Transfer Rate over the last 30 days.

2. Interval Threshold

The 24-hour Interval Troubleshooting notice is posted when the 24-hour Interval is greater than the 30-day Average Interval * Interval Threshold. By default, the Interval Threshold is 2. Thus SyncIQ Job Troubleshooting notice is posted when the 24 hour Interval is more than double the average Interval over the last 30 days.

To Change the Troubleshooting Thresholds:

1. ssh to the Eyeglass appliance.

2. sudo su to root user (default admin password 3y3gl4ss).
3. cd /opt/superna/sca/data
4. Make a backup of the file we are going to edit:

cp system.xml system.xml.bak

5. vi system.xml
6. To update the Transfer Rate Threshold, edit the line that starts with the tag <**transferatethld**>.

Example below to change to 3 - meaning Troubleshooting message only posted when Transfer Rate is less than ⅓ of the average 30 day Transfer Rate:

<transferatethld>3</transferatethld>

7. To update the Interval Threshold, edit the line that starts with the tag <**intervalthld**>.

Example below to change to 5 - meaning Troubleshooting message only posted when Interval is greater than 5 times the average 30-day Interval:

<intervalthld>5</intervalthld>

8. Save your changes.
9. Restart the Eyeglass sca service:

systemctl restart sca

Modify SyncIQ Job Report Time Range

SyncIQ Job Job Report Time Range is the number of hours SyncIQ Job Report data is analyzed from the time the report is run. By default the last

24 hours from the time the report was run are analyzed. To customize the Report Time Range:

1. ssh to the Eyeglass appliance.
2. sudo su to root user (default admin password 3y3gl4ss).
3. cd /opt/superna/sca/data
4. Make a backup of the file we are going to edit:

cp system.xml system.xml.bak

5. vi system.xml
6. Edit the line that starts with the tag <**reporttimerange**> .

Example below to change the Report Time Range to 12 hours (value in hours):

<reporttimerange>24</reporttimerange>

7. Save your changes.
8. Restart the Eyeglass sca service:

systemctl restart sca

© Superna LLC

7.9. RPO CSV Reports

Home Top

- Overview
- Example RPO CSV Report
- Example Policy CSV Report

Overview

The CSV reports are attached to an email sent daily on a schedule. The jobs report is the same information as the html report but in CSV format. The 2nd policy CSV provides policy level view of 24 hour, 30 day, 60 day statistics. PCT means percent and is calculated based on the RPO value entered in the inventory cluster UI. See example below of where this is entered.

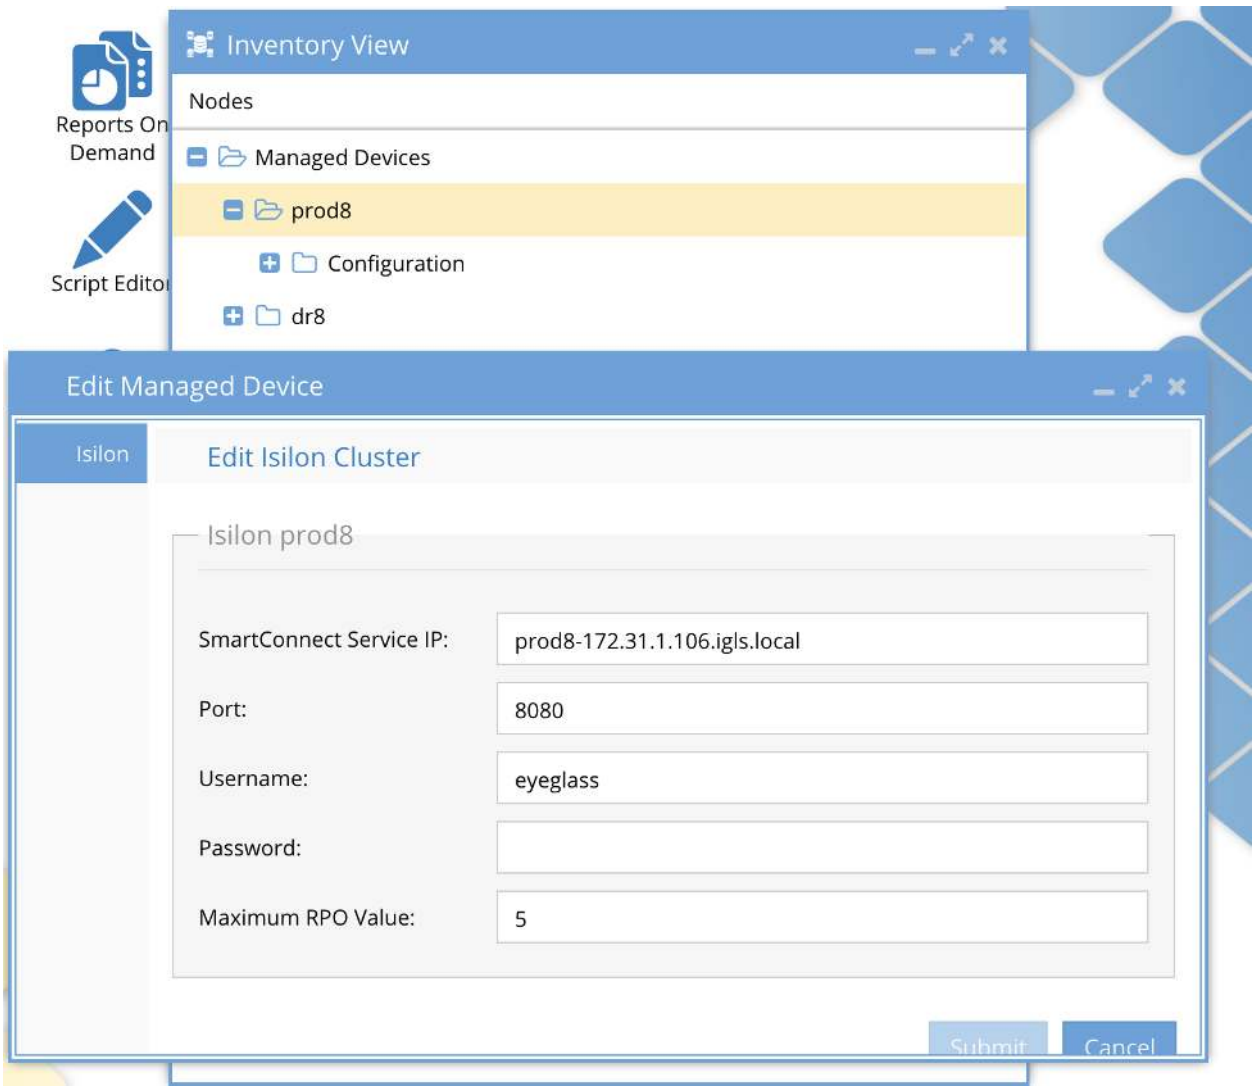

Example RPO CSV Report

Cluster Name	RPO Value	Total Job Count during 24 hours	Finished Job Count during 24 hours	Failed Job Count during 24 hours	Total Data Transfer in GB last 24 hours	Average Data Transfer per Job in GB Last 24 hours	Transfer Rate in Mb Last 24 hours	Avg Job Duration in min Last 24 hours	RPO Violation Last 24 hours	Total Job Count Last 30 days	Finished Job Count Last 30 days	Failed Job Count Last 30 days	Total Data Transfer in GB last 30 days	Average Data Transfer per Job in GB Last 30 days	Transfer Rate in Mb Last 30 days	Avg Job Duration in min Last 30 days	RPO Violation Last 30 days
prod8	5	1	1	0	less than 0.0	0	less than 0.0	4.47	100	49	49	0	0.42	0.01	0.44	2.71	75
dr8	5	0	no data	no data	no data	no data	no data	no data	no data	23	23	0	less than 0.0	0	less than 0.0	0.23	30

Example Policy CSV Report

The RPO value can be edited and calculations added to the CSV to view compliance on a per policy basis.

Cluster Name	Policy Name	RPO Value	Total Job Count Last 24 hours	Finished Job Count during 24 hours	Failed Job Count during 24 hours	Total Data Transfer in GB Last 24 hours	Average Data Transfer per Job in GB Last 24 hours	Transfer Rate in Mb Last 24 hours	Avg Job Duration in min Last 24 hours	Number of Created Files Last 24 hours	Number of Updated Files Last 24 hours	Number of Deleted Files Last 24 hours	Total Job Count Last 30 days	Finished Job Count Last 30 days	Failed Job Count Last 30 days	Total Data Transfer in GB Last 30 days	Average Data Transfer per Job in GB Last 30 days	Transfer Rate in Mb Last 30 days	Avg Job Duration in min Last 30 days	Number of Created Files Last 30 days	Num Upd Files 30 d
dr8	data-dfs_mi	5	0	no data	no data	no data	no data	no data	no data	no data	no data	no data	0	no data	no data	no data	no data	no data	no data	no data	no d
dr8	system-dfs_	5	0	no data	no data	no data	no data	no data	no data	no data	no data	no data	0	no data	no data	no data	no data	no data	no data	no data	no d
dr8	data-nfs_mi	5	0	no data	no data	no data	no data	no data	no data	no data	no data	no data	12	12	0	less than 0.0	0	less than 0.0	0.23	0	
dr8	data-smb_m	5	0	no data	no data	no data	no data	no data	no data	no data	no data	no data	11	11	0	less than 0.0	0	less than 0.0	0.23	0	

U	V	W	X	Y	Z
Number of Created Files Last 30 days	Number of Updated Files Last 30 days	Number of Deleted Files Last 30 days	Pct. of Jobs Violating RPO Last 24 hours	Pct. of Jobs Violating RPO Last 30 days	Pct. of Jobs Violating RPO Last 60 days
no data	no data	no data	no data	no data	100
no data	no data	no data	no data	no data	no data
0	0	0	no data	33	29

© Superna LLC

8. Data And Config Migration Admin Guide

Home Top

- What's New
- Overview
- Typical Use Cases:
- Supported Clusters
- Planning Migrations between Access zones
- Use Case #1 - System to Other Access zone Same cluster
- Use Case #2 - System to Other Access zone Remote cluster
- Use Case #3 - Merge Access Zones Configuration
- Use Case #4 - Overlapping Access Zones Configuration
- Prerequisites to Use the Migration Feature
- How to create a Data and Config Migration Job
- How To Re-apply Default SMB Share ACL Post Migration - only if Enable write access was disabled
- Planning Timeouts for Migration jobs
- Known Limitations
- End to End Data Migration Steps to Move Data/Config and Users to new Access Zone
- Successful Migration Job View - Example

© Superna LLC

8.1. What's New

What's New

1. New in 2.5.6
    a. New Dedicated Icon "Data and Config Migration"
    b. Ability to copy only configuration data from one zone to another without need a synciq policy to exist. This can be used by entering any source path and any target path with auto access zone detection. NOTE: The paths must exist. This new job type can stay and run during normal configuration task to keep config data in sync during a migration. New check box: "Migrate only configuration"
    c. Ability to auto detect existing syncIQ policies on the source and target path and use them to sync configuration data. This allows the copy policy to be setup and migration of configuration can use the existing policies in place.
    d. Migration SyncIQ policies can now be selected and used in SyncIQ mode failover in the DR Assistant. This allows the cut over to be done in Eyeglass using a one way failover. This is only supported between 2 clusters and not between access zones on a single cluster.
2. New in 1.9 Access Zone Migration now allows the SyncIQ policy created to persist after the initial copy and Config sync phase.
3. This allows for phased cutover and incremental sync of data before the final cutover to new access zone or cluster for migrated data.

4. The Policy will appear in the Jobs windows to support incremental config sync changes as well as data sync.

© Superna LLC

8.2. Overview

Overview

Data migration is not just a one time operation. It's a continuous operation to move data between clusters, and within access zones on clusters and from one access zone to another clusters access zone. This feature assists with moving data and configuration data (shares, exports, quotas, nfs aliases) with the data and updating the path and access zone on the target path.

© Superna LLC

8.3. Typical Use Cases:

Typical Use Cases:

1. Split an access zone into two for failover granularity reasons
2. Move an application to its own access zone for security
3. Split data and application load between clusters
4. Move data + configuration data to new access zone in the DR cluster for testing
5. Move data + configuration data to new access zone to achieve active active clusters
6. Migrate data + configuration from several remote cluster to a central cluster with into the same access zone (Fan in) or separate access zones)

© Superna LLC

8.4. Supported Clusters

Supported Clusters

1. PowerScale all models
2. PowerScaleSD
3. See Release notes for feature matrix support and OneFS supported releases

© Superna LLC

8.5. Planning Migrations between Access zones

Planning Migrations between Access zones

Various options exist to move data between access zones and clusters with this feature. This allows moving data and configuration data in various configurations and some planning is required. When planning data migration review the source and destination paths and access zones you plan to move data from and too based on the rules below.

1. **Source path** - The source access zone is selected when configurations is submitted based on the path matching an access zone base path
2. **Target path -** The target access zone same or different cluster access zone is auto detected based on the path matching a base path access zone.

IMPORTANT:

TARGET ACCESS ZONE MUST USE SAME AUTHENTICATION PROVIDERS AS SOURCE ACCESS ZONE as Eyeglass will not be able to translate User and Group SID between AD providers.

© Superna LLC

8.6. Use Case #1 - System to Other Access zone Same cluster

Home Top

Use Case #1 - System to Other Access zone Same cluster

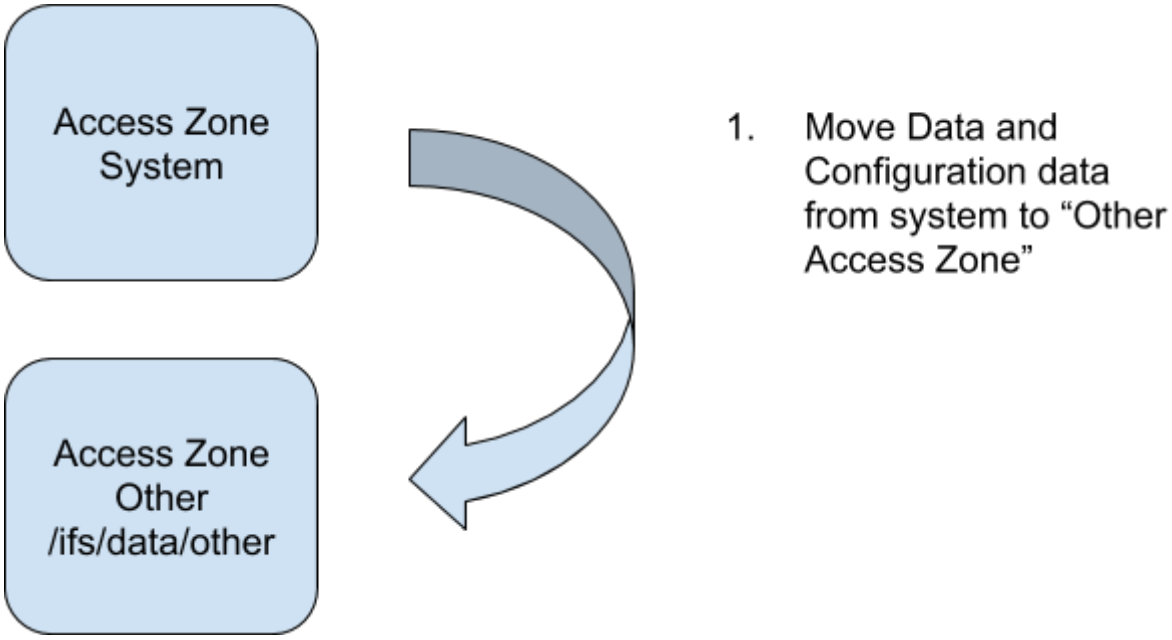

© Superna LLC

8.7. Use Case #2 - System to Other Access zone Remote cluster

Use Case #2 - System to Other Access zone Remote cluster

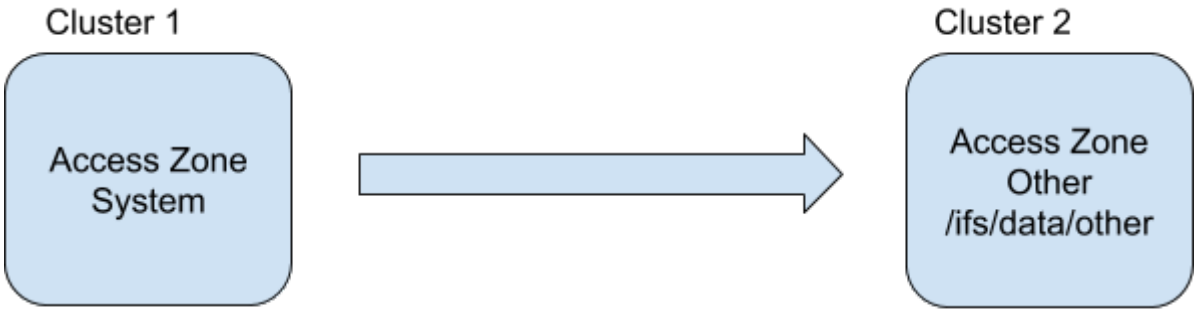

1. Move Data and Configuration data from system to “Other Access Zone” to a remote cluster

© Superna LLC

8.8. Use Case #3 - Merge Access Zones Configuration

Use Case #3 - Merge Access Zones Configuration

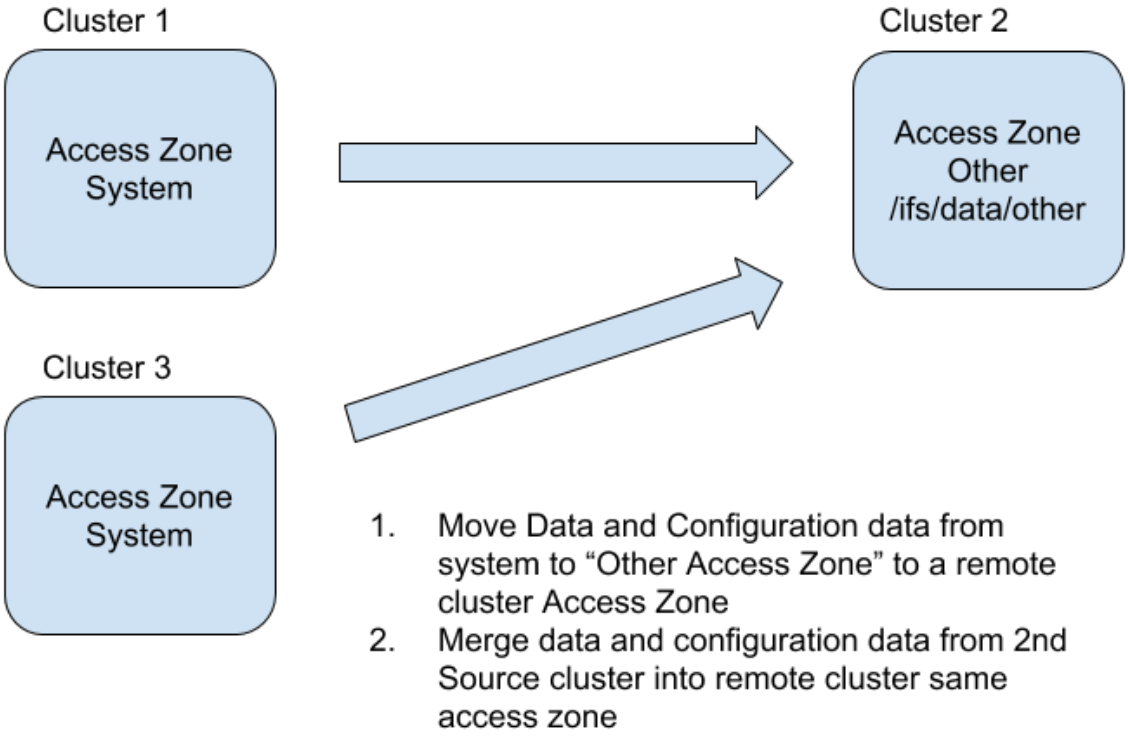

© Superna LLC

8.9. Use Case #4 - Overlapping Access Zones Configuration

Use Case #4 - Overlapping Access Zones Configuration

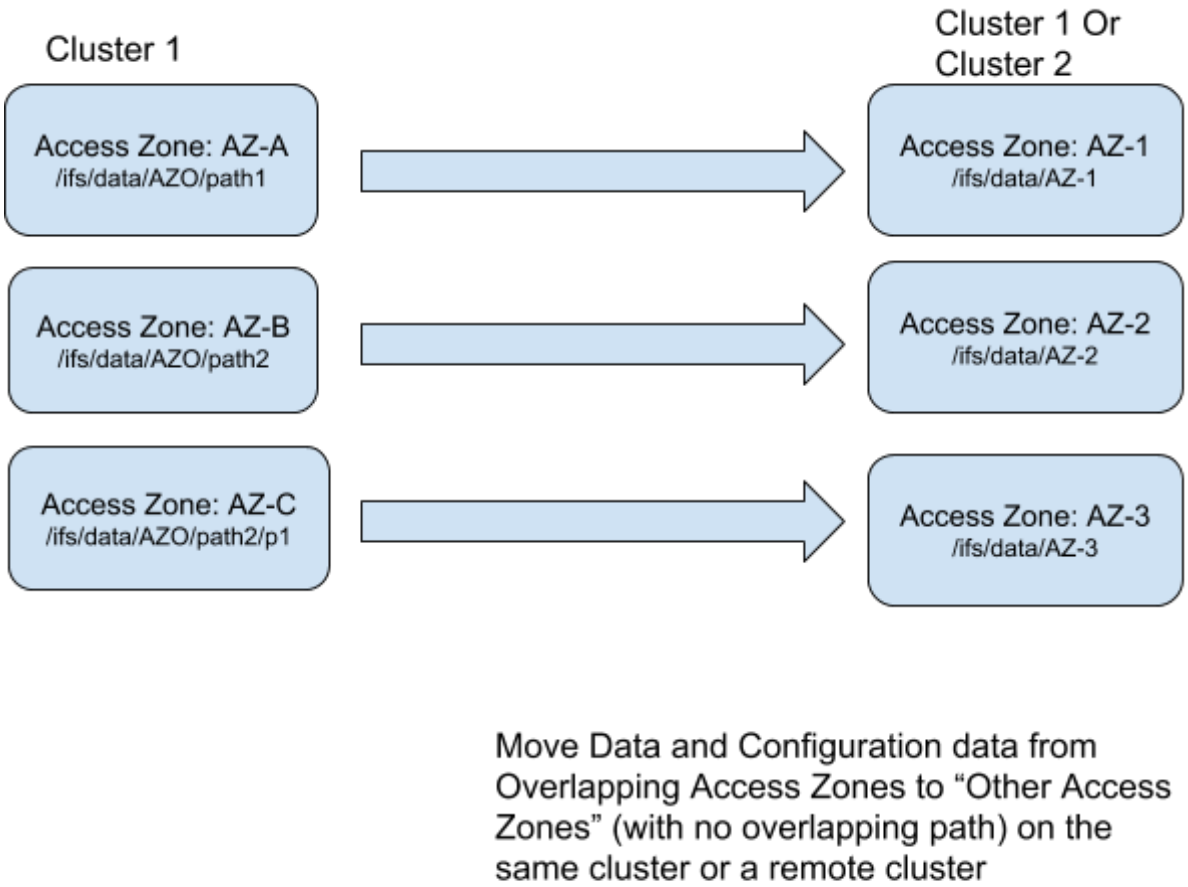

Example:

4 Access Zones with overlapping paths:

Access Zones

Zones

Zones

+ Create an access zone

Select an action

Zone	Path	Groupnet	Actions
zone-sales	/ifs/zone-sales	groupnet0	View / Edit Delete
region-saleszone	/ifs/zone-sales/region-saleszone	groupnet0	View / Edit Delete
marketing-saleszone	/ifs/zone-sales/marketing-saleszone	groupnet0	View / Edit Delete
clientA-datasales	/ifs/zone-sales/region-saleszone/clientA-datasales	groupnet0	View / Edit Delete

In each access zone there are configuration objects as shares , exports, nfs aliases and quotas.

The goal is to migrate the overlapping Access Zones to NEW Access Zone path that does not overlap on same or different cluster.

On Eyeglass, select job window, click add new job, select migration job tab: Select each source access zone path and zone name as source for migration, you can keep synciq policy option for incremental sync. **Note: The initial SyncIQ policy is created with copy option. It must be changed to incremental sync and a schedule applied to maintain sync between source and destination paths.**

Add New Job

Replication Job

Migration Job

Zone Migration Job

Migration Path:	/ifs/zone-sales/marketing-saleszone
Source Cluster:	isi8-source
Source Access Zone:	marketing-saleszone
Enable Source Write Access:	☑
Keep SyncIQ Policy:	☑
Destination Path:	/ifs/AZ-3/marketing
Destination Cluster:	isi8-source

Preview Configuration Submit Cancel

Check if the migration job has finished successfully.

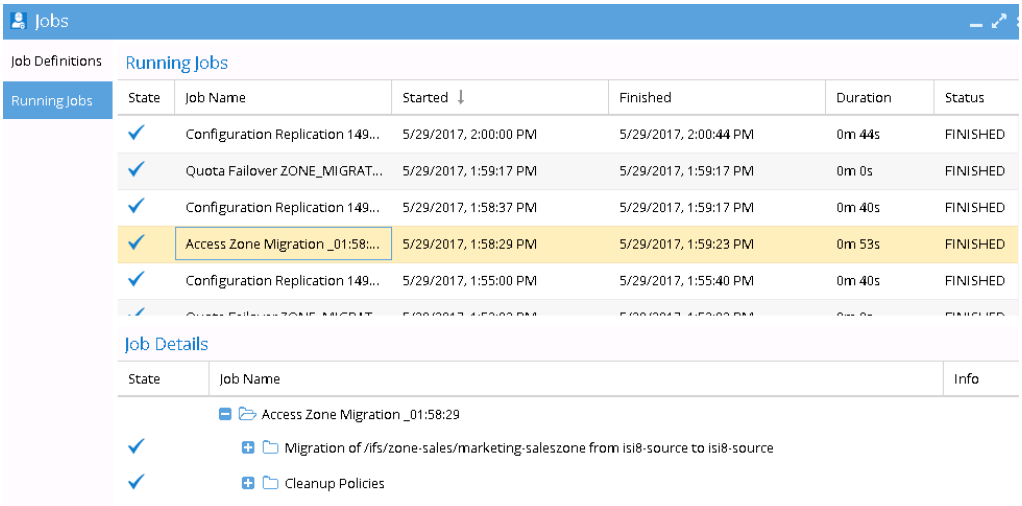

Check the configuration objects on target Access Zone:

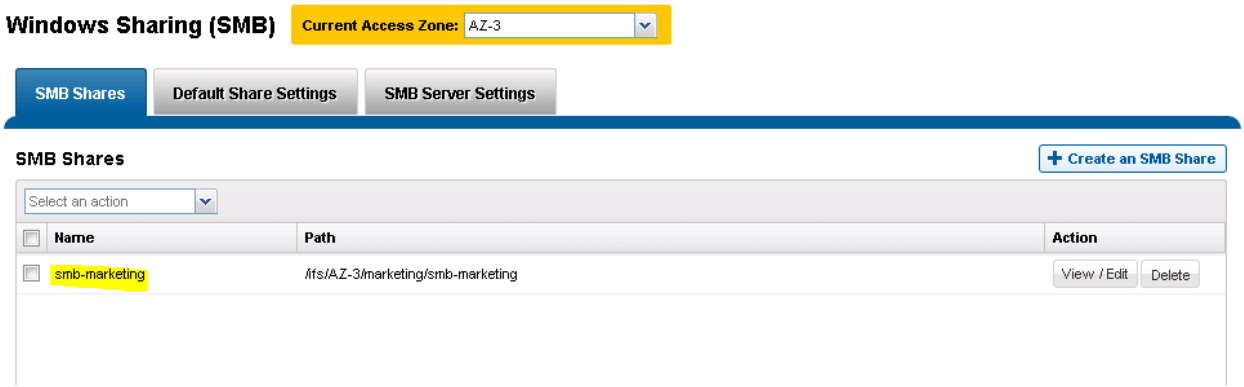

Note: The overlapping path will cause the quotas that are migrated to be duplicated on the new Access Zone paths.

Quotas & Usage

Define quota display | Select an action

Quota Type	Quota Path	Usage	Hard Limit		Soft Limit		Advisory Limit		
			Limit	% Used	Limit	% Used	Limit	% Used	
Directory	/ifs/AZ-1/sales-zone/marketing-saleszone	96 B	--	--	10 GB	< 1%	--	--	View details | Delete
Directory	/ifs/AZ-1/sales-zone/region-saleszone	219 B	--	--	15 GB	< 1%	--	--	View details | Delete
Directory	/ifs/AZ-1/sales-zone/region-saleszone/clientA-datasales	92 B	--	--	7 GB	< 1%	--	--	View details | Delete
Directory	/ifs/AZ-1/sales-zone/test-01	136 B	--	--	5 GB	< 1%	--	--	View details | Delete
Directory	/ifs/AZ-2/region-saleszone	219 B	--	--	15 GB	< 1%	--	--	View details | Delete
Directory	/ifs/AZ-2/region-saleszone/clientA-datasales	92 B	--	--	7 GB	< 1%	--	--	View details | Delete
Directory	/ifs/AZ-3/marketing	96 B	--	--	10 GB	< 1%	--	--	View details | Delete
Directory	/ifs/AZ-4/clinetA-datasales	92 B	--	--	7 GB	< 1%	--	--	View details | Delete
Directory	/ifs/zone-sales/marketing-saleszone	96 B	--	--	10 GB	< 1%	--	--	View details | Delete
Directory	/ifs/zone-sales/region-saleszone	219 B	--	--	15 GB	< 1%	--	--	View details | Delete
Directory	/ifs/zone-sales/region-saleszone/clientA-datasales	92 B	--	--	7 GB	< 1%	--	--	View details | Delete
Directory	/ifs/zone-sales/test-01	136 B	--	--	5 GB	< 1%	--	--	View details | Delete

Usage accounting includes data-protection overhead

© Superna LLC

8.10. Prerequisites to Use the Migration Feature

Prerequisites to Use the Migration Feature

The Eyeglass appliance must have the initial state for Quota Jobs (type QUOTA) set to Enabled to run an Access Migration Job. By default these are Disabled. To Enable them follow these steps:

1. ssh to the Eyeglass appliance and login as admin user.
2. Enter the following CLI commands:

igls adv initialstate set --quota=enabled

3. Enter the following CLI command to check settings:

igls adv initialstate show

4. Check that you see following in the list

"QUOTA": "ENABLED",

© Superna LLC

8.11. How to create a Data and Config Migration Job

Home Top

- Read Me First
- How to start a Migration Job

Read Me First

1. **2.5.6 or later**
    a. **Existing SyncIQ policies can be detected and used to copy configuration data from a source path to a target path on the same or different cluster, this will avoid Eyeglass creating a new policy to sync data. This allows more flexability when to start the data copy for a migration.**
    b. Configuration only option will allow any source to any target path (same or different clusters) to copy shares, exports and quotas and skip any data policy creation steps. Simple check box in the UI to enable configuration only migration job. This will create a new job in the job icon that will persist to sync data from one path to another using the normal configuration job schedule. This job can be deleted once it is no longer required. **NOTE: target path must exist to create the configuration data. No syncIQ policy needs to exist for this feature to work.**
        i. **Not Supported with MultiHop Configuration Replication**: **For same cluster configuration copy on**

overlapping paths, only the source cluster that replicates writeable data can be selected as the source path for a copy job. A DR cluster cannot be used as the source path for a configuration only migration.

c. **New dedicated Icon on the desktop to access the migration features.**

2. **IMPORTANT NOTE:** Since Data copy phase can take hours to complete, the steps in a migration need a timeout and it's hard coded to 15000 minutes or approx 10 days. Any migration of data longer than this will fail.
3. Contact Support before using this feature to ensure your use case is supported.

How to start a Migration Job

1. Open Data Config Migration Icon on the desktop (release 2.5.6 and higher previous releases Jobs Data and Config Migration Icon on the desktop).

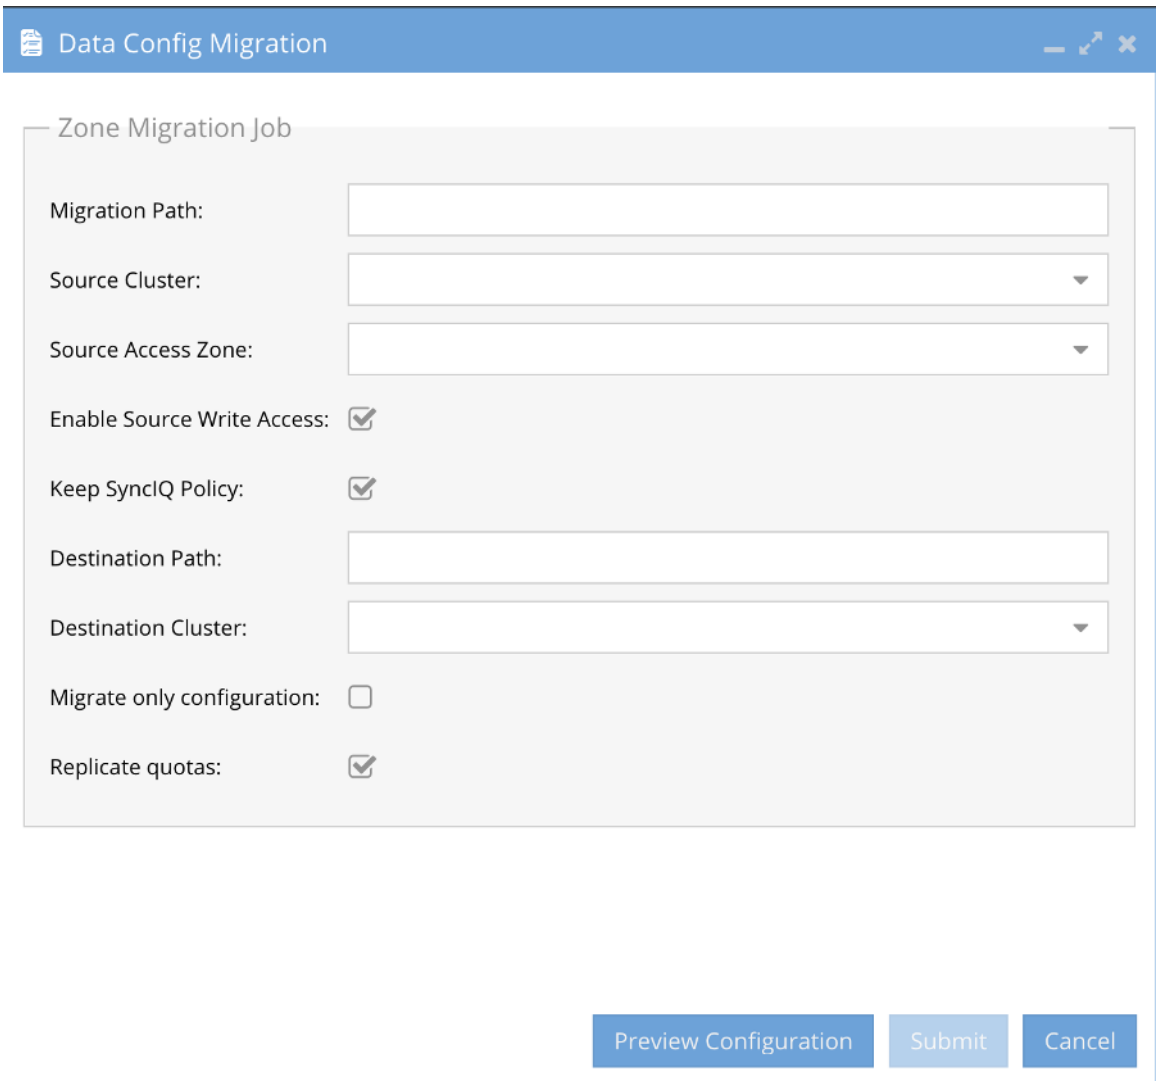

2.

3. Enter source path to migrate on the source cluster
    a. Note: All Configuration data shares, exports, NFS aliases, Quotas must exist at path or below be included in the migration
    b. Enter the source cluster in the drop down (**only Managed clusters in Eyeglass are listed**)
    c. Source Access Zone is a drop down list of all listed access zones detected. Select the zone where the configuration data exists. **NOTE: The reason this is not detected is some clusters allowed overlapping access zone base path which means a source path can have configuration from one or MORE access zones.**

d. **if you want to block user access on the source cluster path, you must de-select "Enable Source Write Access" . Advanced option contact support before using this feature.**

   i. NOTE: if blocking user access to the source cluster path and source cluster path is protected by an existing SyncIQ policy. Then you MUST De-Select "**Enable source write access**" option when the source path is protected by SyncIQ policy. **(WARNING: De-Selected means a the lock policy will block writes to the source folder during migration, the migration job will fail since the path is under a SyncIQ domain. Also note that post migration default top level ACL's will need to be reapplied see section in this guide)**

   ii. This option when de-selected will lock the source path and deny all IO regardless of share or export access settings. The migrated data will inherit the locking SyncIQ policy ACL's on the parent folder when the migration is done.

e. The Target path entered into the migration job will have modified permissions that will need to be restored **BEFORE** users can access the data in the new location. This provides a 2nd level of data locking before the new data is in production.

   i. See Detailed Steps below to restore the ACL settings on the parent folder.

f. **Configuration sync only check box (default disabled) (2.5.6 or later release)** - Use this if no synciq policy exists from source to target path and the data is already synced, this option will create a special job in the Jobs configuration icon to read all the configuration data from the source path and create it on the target path and will create the

configuration data in the access zone that was detected on the target path.

i. NOTE: In 2.5.6 build 84 or eariler this feature is a mirror mode and will remove any overlapping configuration data found on the target path or below. This will be changed in a patch release to default to merge mode and will leave all target configuration data as is and copy only new configuration to the target path.

ii. NOTE: the path must exist on the target location entered relative to the source path configuration objects. This means the entire folder structure must exist on the target path entered. The Share and NFS create will not create a folder if it does not exist.

iii. NOTE: After the configuration migration job has completed. The job should be deleted from the jobs window by selecting the job check box on the left --> bulk actions and delete option. This config job should not be left in place.

iv. 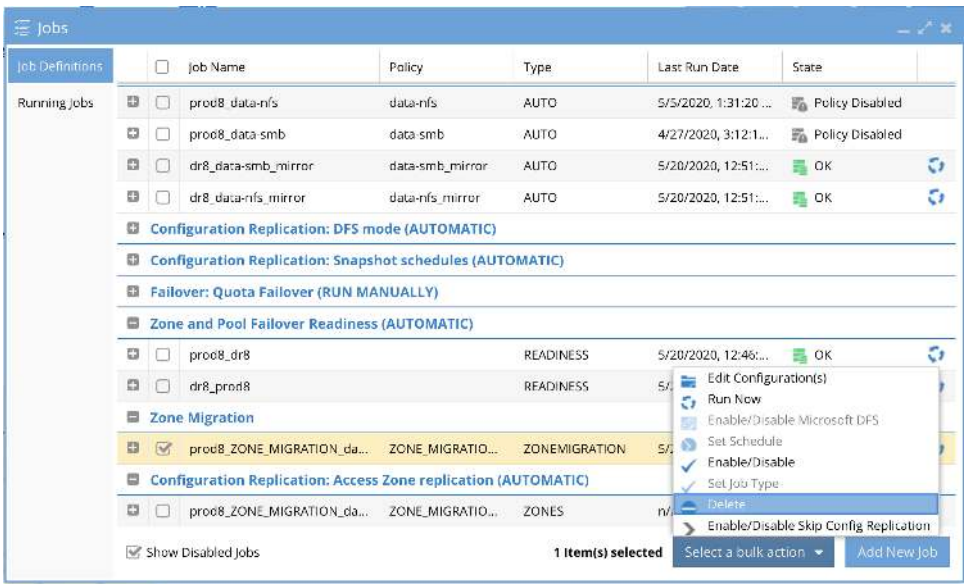

g. **Keep SyncIQ Policy should be checked.** Do not uncheck without asking support. This feature is used for incremental data sync after initial sync of data and allows scheduled cut over process to ensure all data and configuration is synced before schedule cut over.

   i. Check this box to leave the SyncIQ policy after the policy runs. Leaving the policy allows multi runs from OneFS UI or set a schedule to keep the target path updated before the scheduled cut over day.

h. The Configuration data is only synced one time on first job run using Copy mode on the SyncIQ policy.

i. 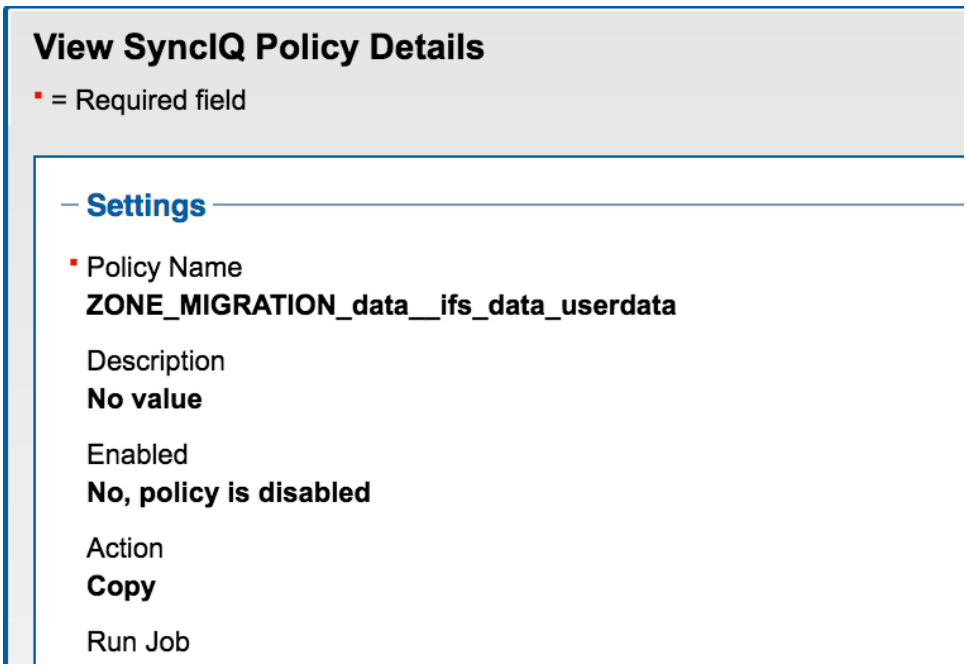

i. Data can by synced incrementally by setting syncIQ policy Sync mode and setting a schedule on the migration policy created by Eyeglass or running it manually.

i. 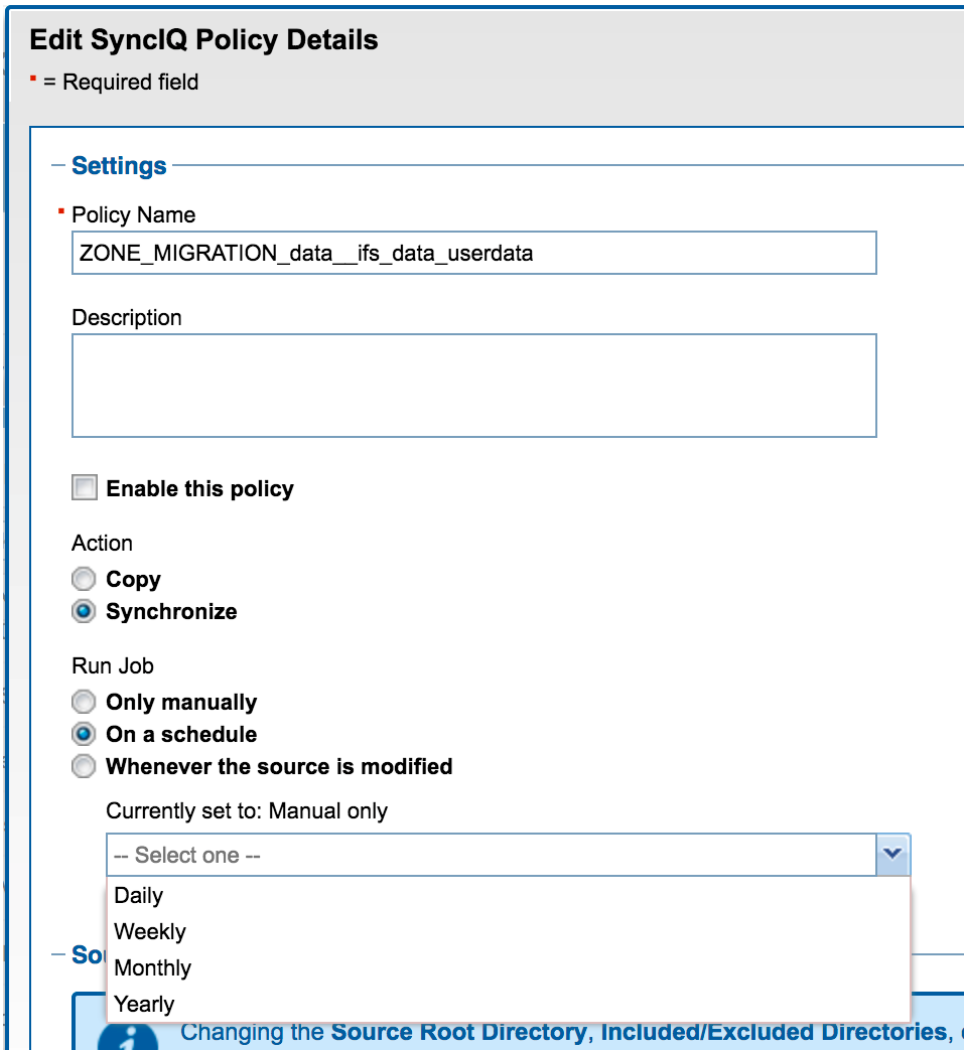

j. The policy will appear in the Jobs windows under ZoneMigration section and can be used to incrementally sync configuration data that has changed.

i. 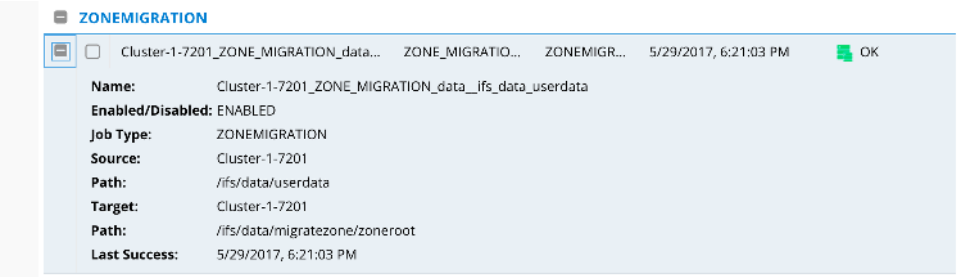

ii. The target Access Zone will be auto detected based on path matching of the Access Zone Base path (on local or remote cluster migrations)

1. Note: This can be changed and configuration path will be updated on shares,exports,quotas and aliases during the migration.
2. Note: Path can be on the same cluster or remote cluster (SyncIQ policy will copy data to the target cluster and it must be IP reachable by the source cluster)
3. Note: Path cannot be the target of an existing SyncIQ policy as it will be in read-only state which will block migration. Must be writable location on target cluster.
4. Note: No path or data can exist on the target cluster. The target path is checked if it exists. If it exists the migration will not continue. An empty target path is required.

k. Enter destination cluster from the drop down (must be managed cluster in Eyeglass with DR license key)

l. Select the Preview option to verify which shares, exports and quotas, aliases were discovered for migration and validate this is expected.

m. **Auto detect Existing SyncIQ (New in 2.5.6 or later releases)** - If a SynIQ policy is detected on the source cluster migration path and correct target cluster path then no migration policy will be created and the existing syncIQ policy will be used to detect configuration data to sync. This allows the Data and Config migration to be used after a data sync copy is already in progress.

i. **NOTE: No administrator action is required, Eyeglass will automatically detect if SyncIQ policy needs to be created for the migration.**

n. Click the Submit button to start the migration

o. Monitor from the Running Jobs tab of the Jobs Icon.

p. **(2.5.6 or later releases)** When ready to cut over the data to the new cluster or access zone, the DNS update for subnet service ip and SPN changes are manual steps but the data can be failed over with DR Assistant using SyncIQ policy failover mode.

© Superna LLC

8.12. How To Re-apply Default SMB Share ACL Post Migration - only if Enable write access was disabled

How To Re-apply Default SMB Share ACL Post Migration - only if Enable write access was disabled

Use this procedure **ONLY** if you deselected allow “Enable write access” on the migration job. This migration option will apply ACL’s on the target that block all access. The ACL’s are applied only at the top level folder path. Any ACL’s that have been applied to child paths of the parent migration path, will retain the ACL’s post migration.

The top level parent modified ACL blocks access to the data even if the share or export level permission allows write access or full control. It is the combination of Share/export and ACL’s that allow write access to the data.

© Superna LLC

8.12.1. Add back the original Microsoft Default ACL's

Add back the original Microsoft Default ACL's

1. M8000A-1# chmod +a# 0 group Administrators allow dir_gen_all,object_inherit,container_inherit group1
2. M8000A-1# chmod +a# 1 creator_owner allow dir_gen_all,object_inherit,container_inherit,inherit_only group1
3. M8000A-1# chmod +a# 2 everyone allow dir_gen_read,dir_gen_execute group1
4. M8000A-1# chmod +a# 3 group Users allow dir_gen_read,dir_gen_execute,object_inherit,container_inherit group1
5. M8000A-1# chmod +a# 4 group Users allow std_synchronize,add_file,add_subdir,container_inherit group1

© Superna LLC

8.12.2. Check ACL’s after changes to migrated parent folder

Check ACL’s after changes to migrated parent folder

1. M8000A-1# ls -lze
2. total 2
3. drwxrwxr-x + 4 root wheel 99 Nov 17 20:42 group1
4. OWNER: user:root
5. GROUP: group:wheel
6. 0: group:Administrators allow dir_gen_all,object_inherit,container_inherit
7. 1: creator_owner allow dir_gen_all,object_inherit,container_inherit,inherit_only
8. 2: everyone allow dir_gen_read,dir_gen_execute
9. 3: group:Users allow dir_gen_read,dir_gen_execute,object_inherit,container_inherit
10. 4: group:Users allow std_synchronize,add_file,add_subdir,container_inherit
11. 5: user:root allow dir_gen_read,dir_gen_write,dir_gen_execute,std_write_dac,delete_child

12. 6: group:wheel allow dir_gen_read,dir_gen_write,dir_gen_execute,delete_child
13. 7: everyone allow dir_gen_read,dir_gen_execute

© Superna LLC

8.12.3. Delete extras ACL’s

Delete extras ACL’s

1. M8000A-1# chmod -a# 7 group1
2. M8000A-1# chmod -a# 6 group1
3. M8000A-1# chmod -a# 5 group1

© Superna LLC

8.12.4. Check ACL Delete

Check ACL Delete

1. M8000A-1# ls -lze
2. total 2
3. drwxrwxr-x + 4 root wheel 99 Nov 17 20:42 group1
4. OWNER: user:root
5. GROUP: group:wheel
6. 0: group:Administrators allow dir_gen_all,object_inherit,container_inherit
7. 1: creator_owner allow dir_gen_all,object_inherit,container_inherit,inherit_only
8. 2: everyone allow dir_gen_read,dir_gen_execute
9. 3: group:Users allow dir_gen_read,dir_gen_execute,object_inherit,container_inherit
10. 4: group:Users allow std_synchronize,add_file,add_subdir,container_inherit

© Superna LLC

8.12.5. Connect to Smartconnect Name to test mount and write access to the data

Home Top

- Connect to Smartconnect Name to test mount and write access to the data
  - Review the default ACL's applied to Shares by OneFS

Connect to Smartconnect Name to test mount and write access to the data

1. This step verifies the ACL's and SPN smartconnect name mount succeeds.
2. Using AD account that has permissions to the share moun the FQDN of the smartconnect name of the new location of the data.
3. Test write access to the share
4. If successful the ACL's applied was correctly completed
5. Done.

Review the default ACL's applied to Shares by OneFS

1. Review the default ACL's on shares created with Default Microsoft ACL's
2. Create share (directory does not exist)

3. 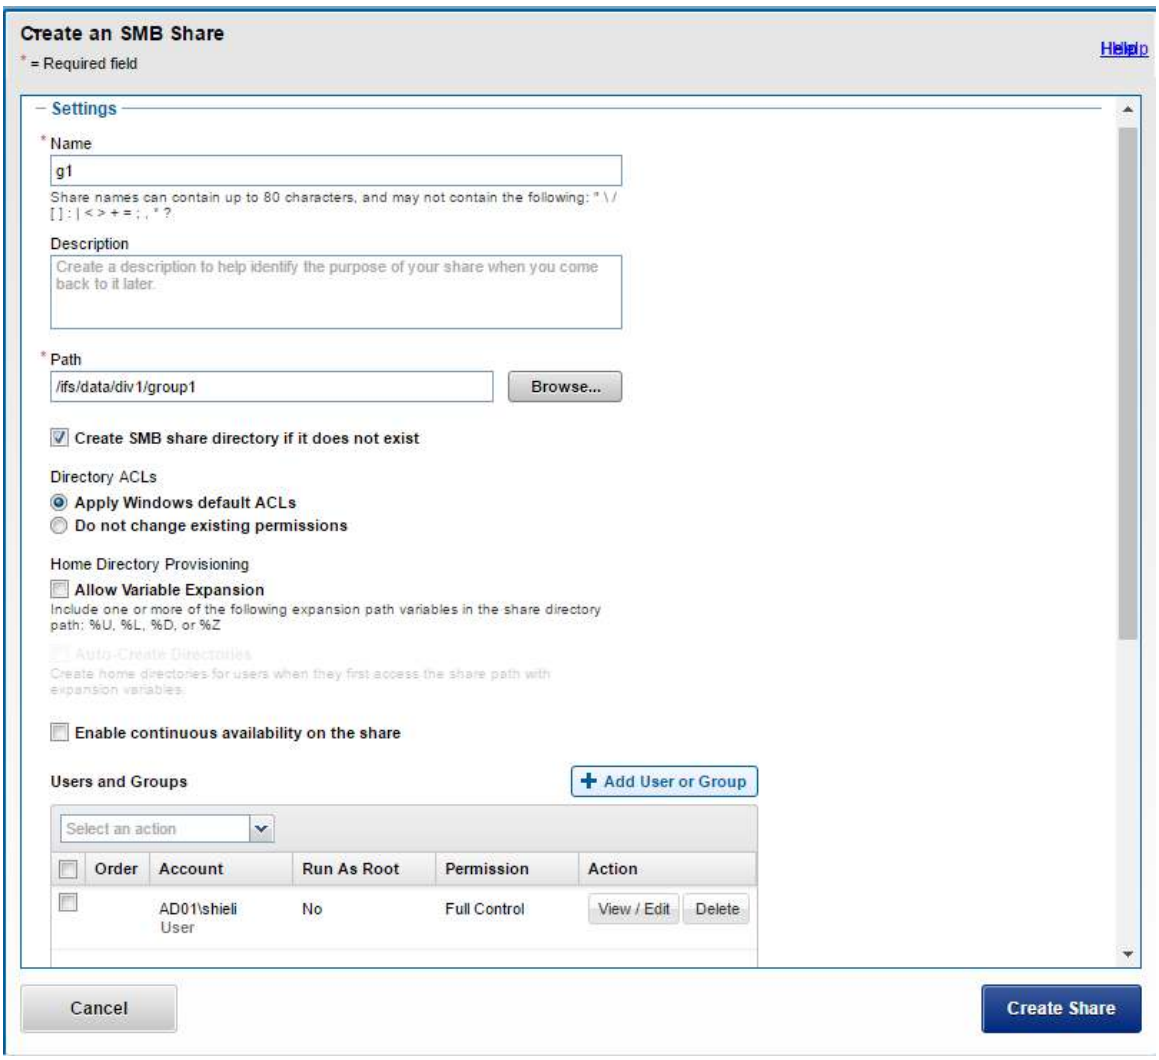

4. Folder permissions after folder is automatically created

5. 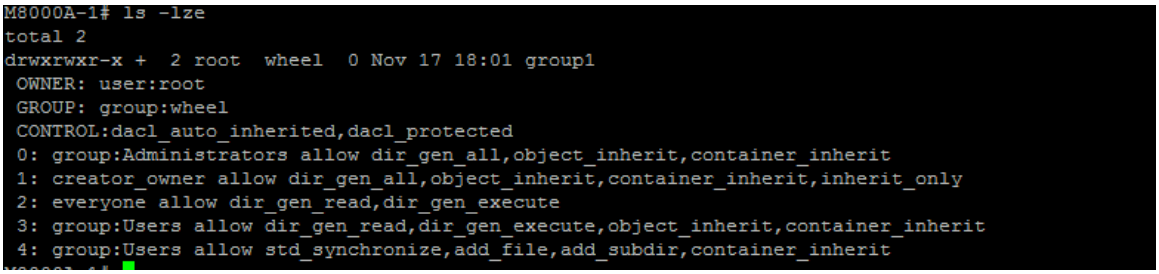

```
M8000A-1# ls -lze
total 2
drwxrwxr-x +  2 root  wheel  0 Nov 17 18:01 group1
 OWNER: user:root
 GROUP: group:wheel
 CONTROL:dacl_auto_inherited,dacl_protected
 0: group:Administrators allow dir_gen_all,object_inherit,container_inherit
 1: creator_owner allow dir_gen_all,object_inherit,container_inherit,inherit_only
 2: everyone allow dir_gen_read,dir_gen_execute
 3: group:Users allow dir_gen_read,dir_gen_execute,object_inherit,container_inherit
 4: group:Users allow std_synchronize,add_file,add_subdir,container_inherit
```

6. Share security settings:

7. 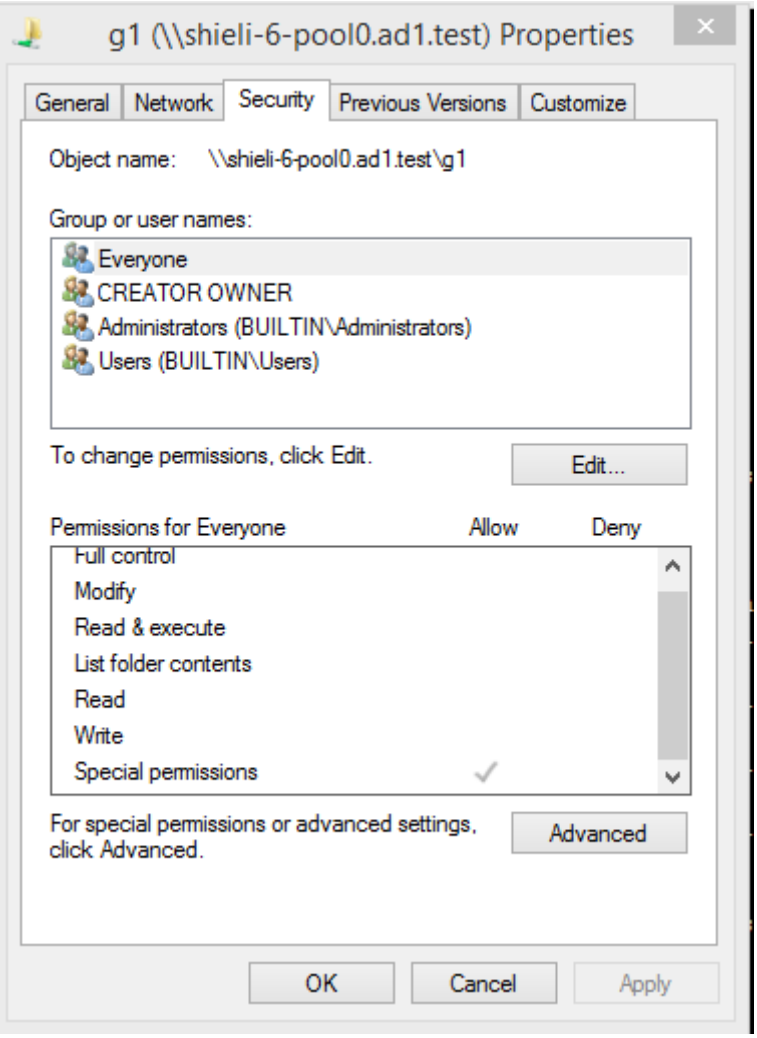
g1 (\\shieli-6-pool0.ad1.test) Properties
General
Network
Security
Previous Versions
Customize
Object name: \\shieli-6-pool0.ad1.test\g1
Group or user names:
Everyone
CREATOR OWNER
Administrators (BUILTIN\Administrators)
Users (BUILTIN\Users)
To change permissions, click Edit.
Edit...
Permissions for Everyone
Allow
Deny
Full control
Modify
Read & execute
List folder contents
Read
Write
Special permissions
For special permissions or advanced settings, click Advanced.
Advanced
OK
Cancel
Apply

8. 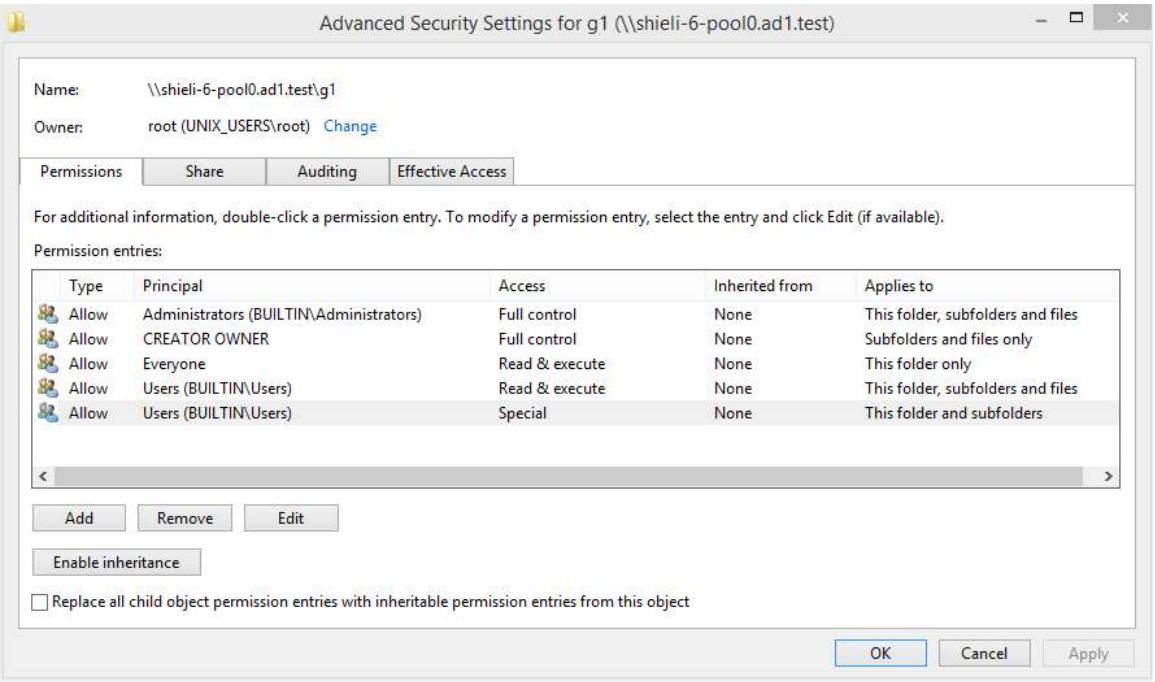
Advanced Security Settings for g1 (\\shieli-6-pool0.ad1.test)
Name: \\shieli-6-pool0.ad1.test\g1
Owner: root (UNIX_USERS\root) Change
Permissions
Share
Auditing
Effective Access
For additional information, double-click a permission entry. To modify a permission entry, select the entry and click Edit (if available).
Permission entries:
Type Principal Access Inherited from Applies to
Allow Administrators (BUILTIN\Administrators) Full control None This folder, subfolders and files
Allow CREATOR OWNER Full control None Subfolders and files only
Allow Everyone Read & execute None This folder only
Allow Users (BUILTIN\Users) Read & execute None This folder, subfolders and files
Allow Users (BUILTIN\Users) Special None This folder and subfolders
Add
Remove
Edit
Enable inheritance
Replace all child object permission entries with inheritable permission entries from this object
OK
Cancel
Apply

9.

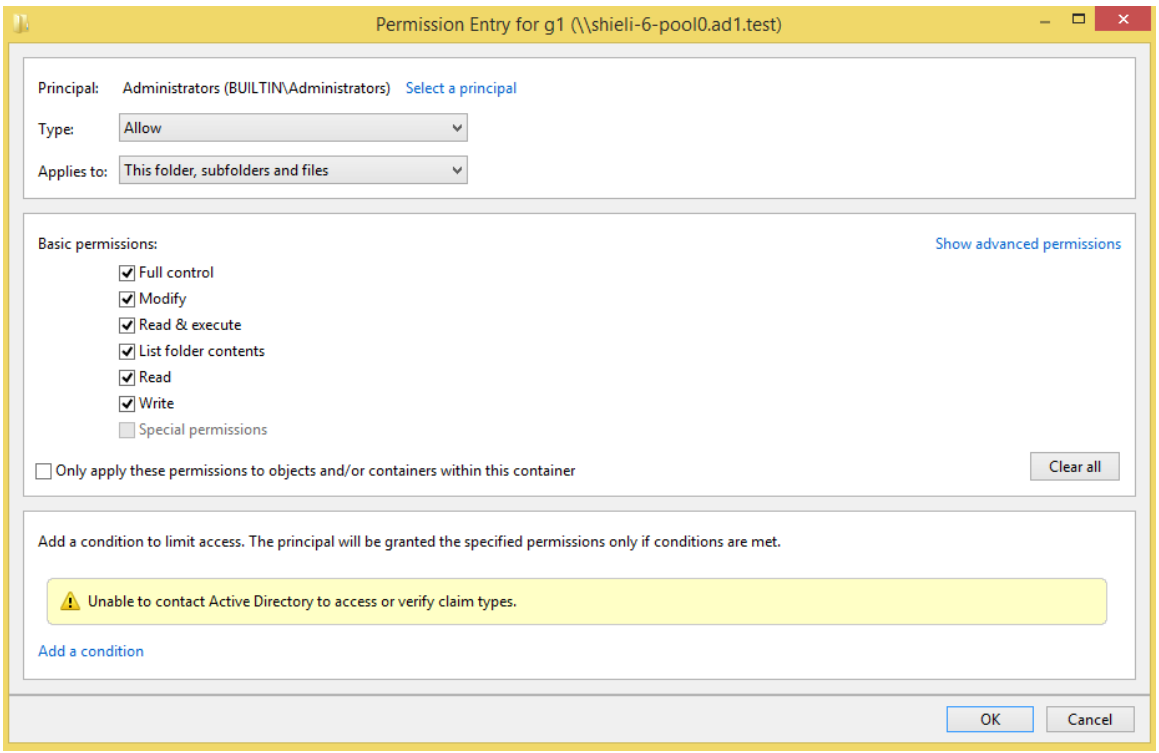
Permission Entry for g1 (\\shieli-6-pool0.ad1.test)
Principal: Administrators (BUILTIN\Administrators) Select a principal
Type: Allow
Applies to: This folder, subfolders and files
Basic permissions:
Show advanced permissions
Full control
Modify
Read & execute
List folder contents
Read
Write
Special permissions
Only apply these permissions to objects and/or containers within this container
Clear all
Add a condition to limit access. The principal will be granted the specified permissions only if conditions are met.
Unable to contact Active Directory to access or verify claim types.
Add a condition
OK
Cancel

10.

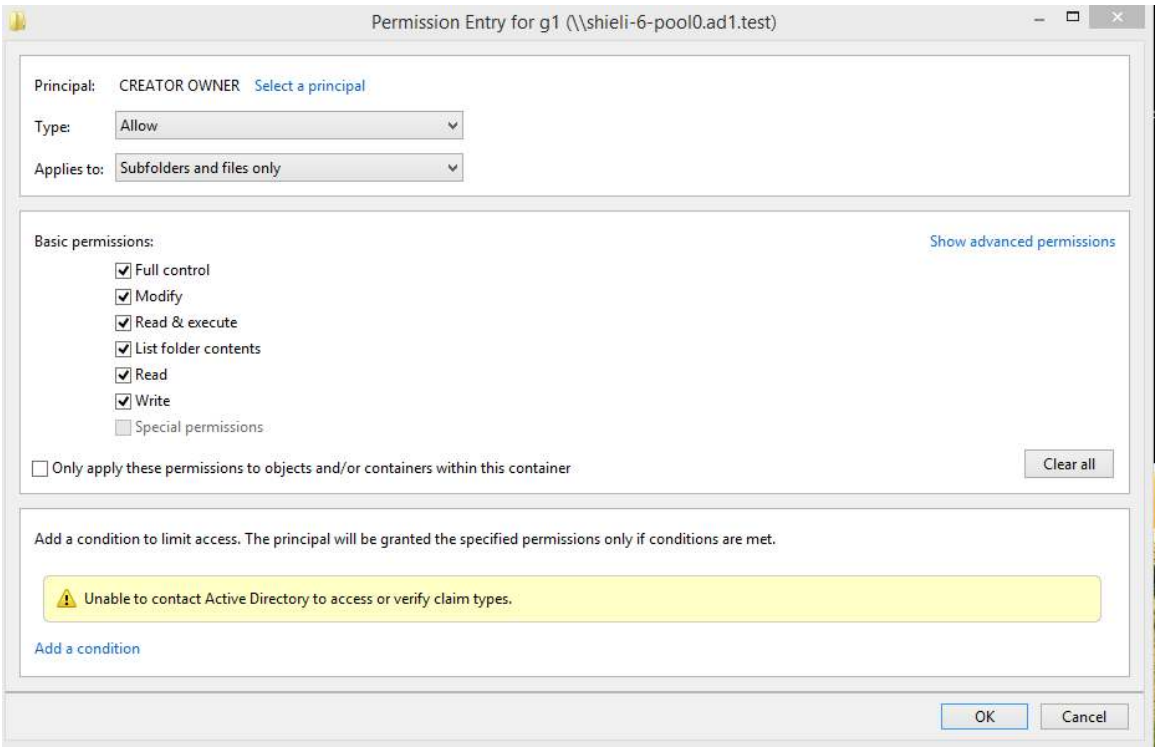
Permission Entry for g1 (\\shieli-6-pool0.ad1.test)
Principal: CREATOR OWNER Select a principal
Type: Allow
Applies to: Subfolders and files only
Basic permissions:
Show advanced permissions
Full control
Modify
Read & execute
List folder contents
Read
Write
Special permissions
Only apply these permissions to objects and/or containers within this container
Clear all
Add a condition to limit access. The principal will be granted the specified permissions only if conditions are met.
Unable to contact Active Directory to access or verify claim types.
Add a condition
OK
Cancel

11.

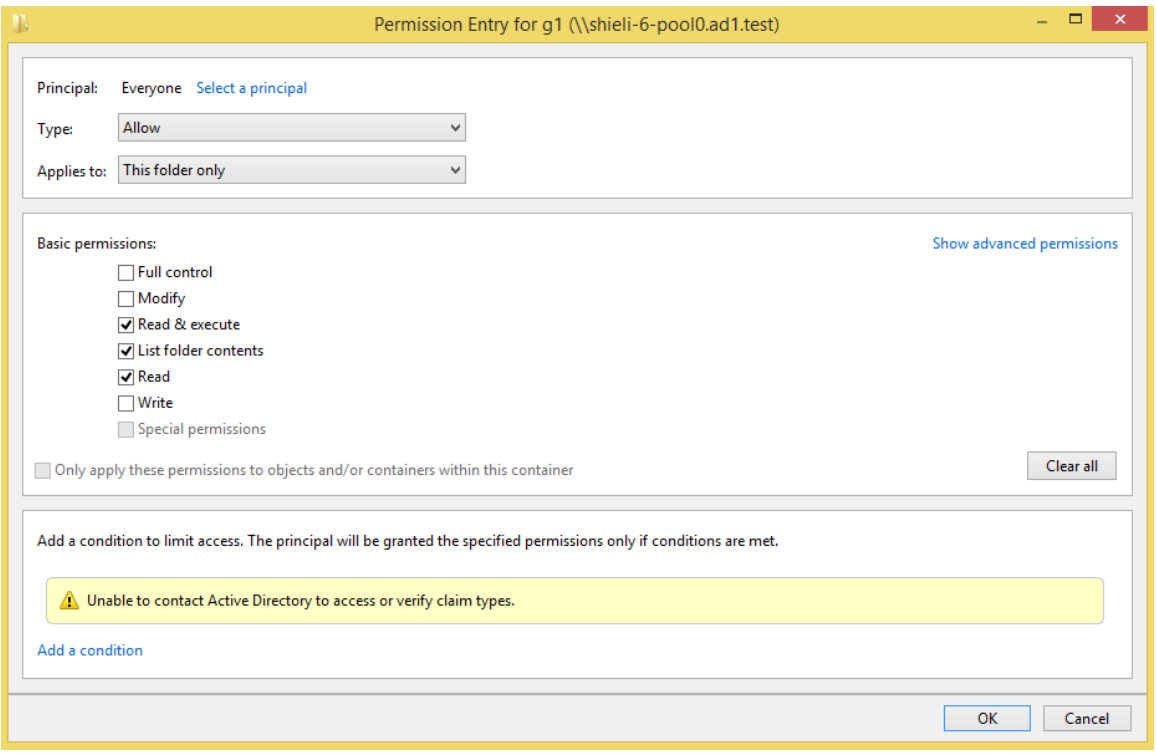
Permission Entry for g1 (\\shieli-6-pool0.ad1.test)
Principal: Everyone Select a principal
Type: Allow
Applies to: This folder only
Basic permissions:
Show advanced permissions
Full control
Modify
Read & execute
List folder contents
Read
Write
Special permissions
Only apply these permissions to objects and/or containers within this container
Clear all
Add a condition to limit access. The principal will be granted the specified permissions only if conditions are met.
Unable to contact Active Directory to access or verify claim types.
Add a condition
OK
Cancel

12.

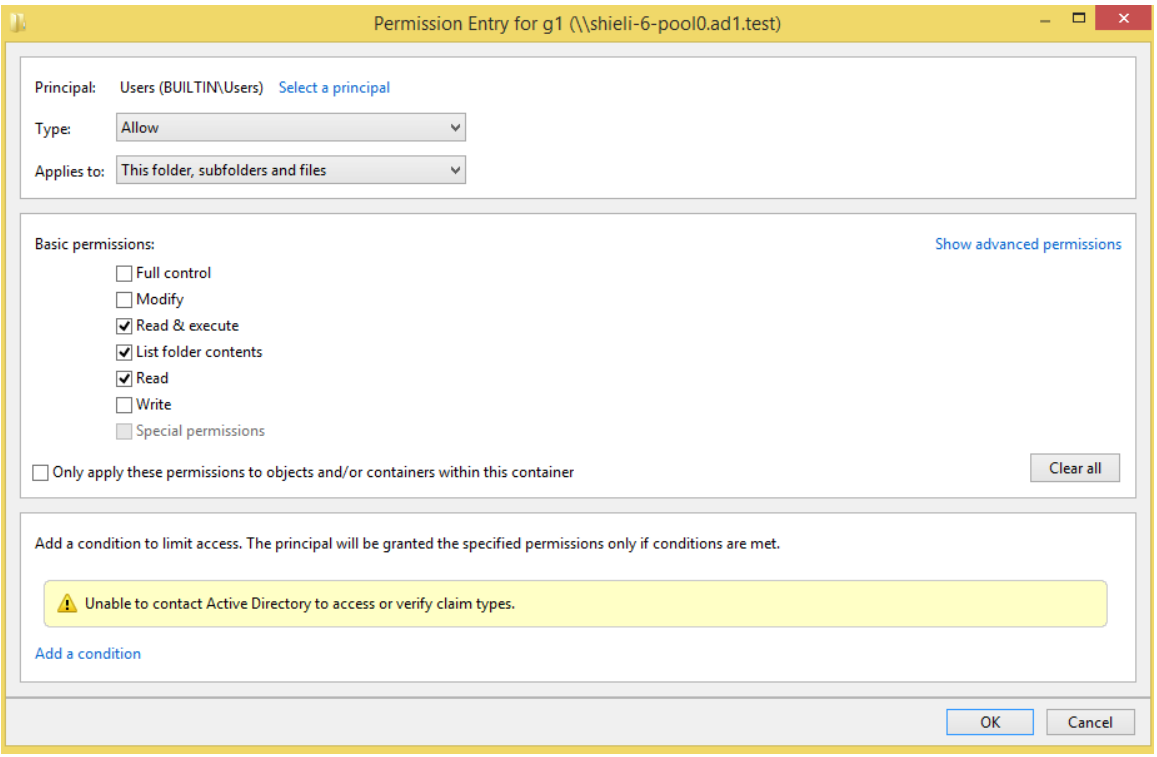
Permission Entry for g1 (\\shieli-6-pool0.ad1.test)
Principal: Users (BUILTIN\Users) Select a principal
Type: Allow
Applies to: This folder, subfolders and files
Basic permissions:
Show advanced permissions
Full control
Modify
Read & execute
List folder contents
Read
Write
Special permissions
Only apply these permissions to objects and/or containers within this container
Clear all
Add a condition to limit access. The principal will be granted the specified permissions only if conditions are met.
Unable to contact Active Directory to access or verify claim types.
Add a condition
OK
Cancel

13.

14.

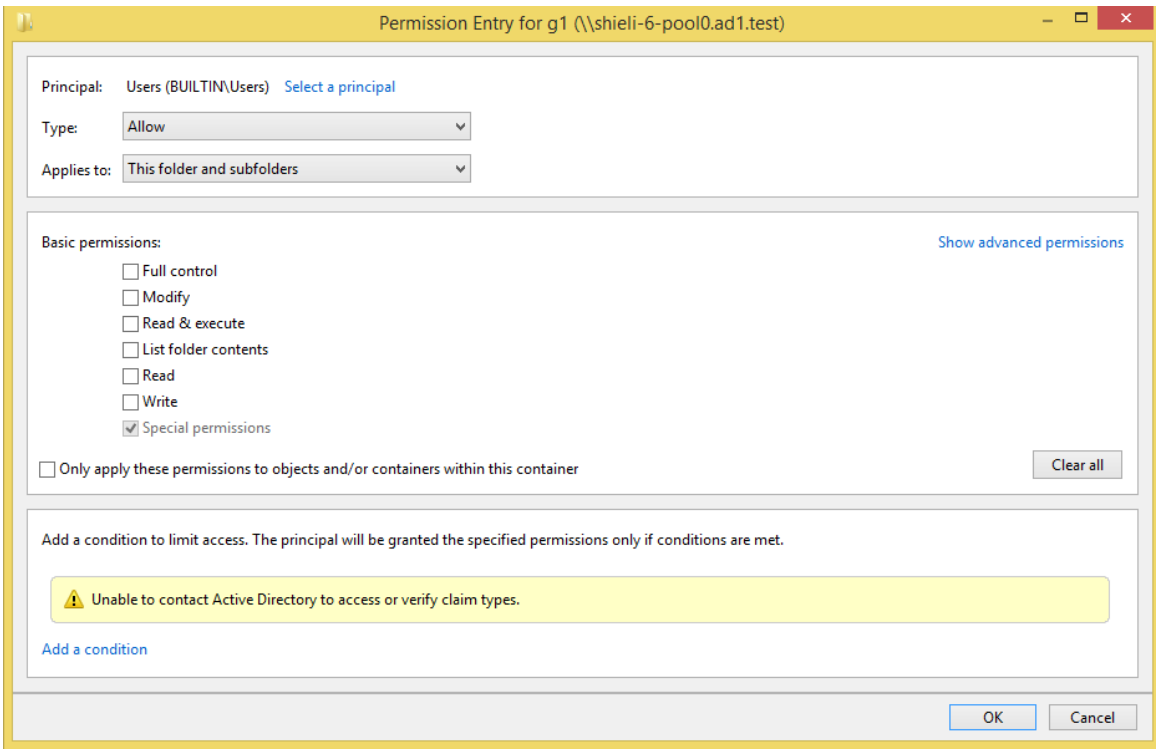

© Superna LLC

8.13. Planning Timeouts for Migration jobs

Planning Timeouts for Migration jobs

These can help debug timeouts that may occur for long running jobs. All Default timers below should be ok for most migrations.

- Run policy: 15000 minutes (10 days)
- Wait for migration of config to Complete: 50 sec
- Wait for locking policy to Complete: 50 sec
- Wait for opened files to be closed: 300 sec (this will delay the start of the job and will fail if force flag not enabled)
- Cleanup migration policy: 75 sec (only if keep policy unchecked)
- Cleanup locking policy : 75 sec (only if allow source access is unchecked)

© Superna LLC

8.14. Known Limitations

Known Limitations

- Cannot migrate igls-dfs shares. In this case writeable copy of data with non-prefixed shares must be migrated.

© Superna LLC

8.15. End to End Data Migration Steps to Move Data/Config and Users to new Access Zone

Home Top

End to End Data Migration Steps to Move Data/Config and Users to new Access Zone

#	Steps Outline	Details of Test Setup
1	**Does not delete original shares/exports** **Use the same smartconnect zone to access data post migration**	**WITH SYNCIQ POLICY** Same cluster access zone migration testzone1 -> testzone2 Partial data covered by syncIQ policy div1 Share g1: /ifs/testzone/div1/group1 Share g1_1: /ifs/testzone1/div1/group1/group1_1 Share g2:/ifs/testzone1/div1/group2 Share g3:/ifs/testzone1/div1/group3 Migrate Share g1: /ifs/testzone1/div1/group1 to /ifs/testzone2/div1/group1 Share g1_1: /ifs/testzone1/div1/group1/group1_1
2	**Ne**: setup new access zone with AD provider	testzone2 setup
3	**External**: Refrain users from writing data being migrated	Method to be determined by customer
4	**Eyeglass**:	Eyeglass job should be on all the time E.g. not all data covered by a policy are

	Do not disable Eyeglass job	migrated
5	**Eyeglass**: Access zone migration from testzone1 to testzone2	Note: must select 'Enable Source Write Access' if source path is covered by a policy to proceed Data replicated Share/export/quota's created for new Access Zone on source
6	**PowerScale**: associate new Access Zone to IP pool	For both source and target clusters setup IP pool in the new access zone
6A	**PowerScale:** Setup schedule to incrementally sync data on the zonemigration policy.	This is created by Eyeglass.
6B	**Eyeglass:** Run incremental config sync from Jobs window using the Zone migration policy name created by the migration job.	Note: data sync on new directory paths that are created for config data must already exist.
6C	The steps 6A and 6B should be repeated up until the final day of the cut over to the new access zone. This step should be done before moving the smartconnect zone name from the old IP pool to the New IP pool created in the steps below.	
6D	**Schedule Maintenance Window:** 1. Repeat step 6C 2. Rename or delete source access zone smartconnect zone IP pool name. 3. Create smartconect zone name on new access zone Pool on the new access zone. 4. Verify DNS resolves correctly to the new IP pool using nslookup FQDN of smartconnect name. 5. Verify from OneFS shares and exports on target access zone name exist as expected	

6E	1. On Day of cut over use DR Assistant using SyncIQ policy failover mode. 2. Failover data with DR Assistant	NOTE: Requires 2.5.6 or later release to select Migration policies in DR Assisstant
7	**User**: able to access to new share using the same SmartConnect Zone name	Note: having problem with connection is not updated after associated to new AZ. Seems to be a Windows problem, Shows correct shares in zone only after reboot.
8	**Old shares**: not accessible using the same SmartConnect Zone name since old SmartConnect Zone name has been renamed	Users left behind who were using that SmartConnect Zone name have no access and now need a new one - requires new pool and new SPN for new SmartConnect Zone name
9	**PowerScale**: Reprotect data to DR cluster	Create new synciq policy **Create new policy with new path** **Run new policy**
10	**PowerScale**: Create new corresponding access zone on target cluster	
11	**Eyeglass**: Enable Config Replication Job	**Run new job - shares/exports are replicated to target cluster** **Note: original job is still running - protecting shares left behind** **If select 'Delete Source Configuration', original shares are deleted after access zone migration, old shares on target cluster are deleted after config replication.**
12	**Eyeglass**: Readiness job	Re-run readiness
13	**Eyeglass**: mirror policy	Re-run failover

© Superna LLC

8.16. Successful Migration Job View - Example

Successful Migration Job View - Example From Running Jobs Tab of the Jobs Icon

- Successful Migration Job View - Example From Running Jobs Tab of the Jobs Icon
  - Migration Job Setting - Example 1
    - Successful Migration Job View - Example 1
  - Migration Job Setting - Example 2
    - Successful Migration Job View - Example 2
  - Migration Job Setting - Example 3
    - Successful Migration Job View - Example 3
  - Migration Job Setting - Example 4 - Config Only Migration
    - Successful Migration Job View - Example 4

Migration Job Setting - Example 1

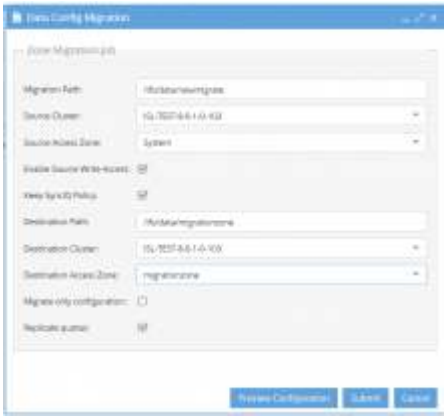

Successful Migration Job View - Example 1

Migration Job Setting - Example 2

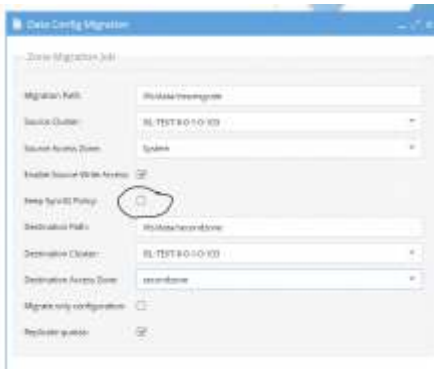

Successful Migration Job View - Example 2

Migration Job Setting - Example 3

Successful Migration Job View - Example 3

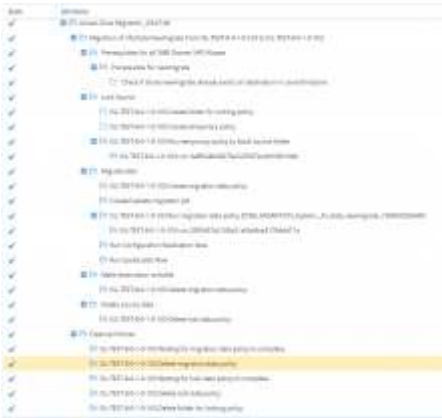

Migration Job Setting - Example 4 - Config Only Migration

Successful Migration Job View - Example 4

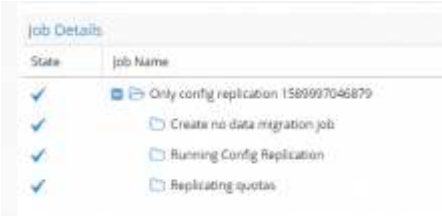

© Superna LLC

9. Eyeglass Search & Recover Admin Guide

Eyeglass Search & Recover Admin Guide

- Overview
- Software Release Updates
- Installation

Introduction to this Guide

The purpose of this guide is to assist you in configuring and administering Eyeglass Search and Recover.

Overview

Eyeglass Search and Recover provides a solution to simplify:

- Finding data based on content
- E-discovery
- Capacity management
- Recovering files from snapshots.

The product is designed to scale with native PowerScale array incremental indexing.

Software Release Updates

The initial release of the Eyeglass Search & Recover product is focused on user, and administrator use cases for content and file meta data searching. The What's Coming link will provide you with detail on past, current and future enhancements.

Installation

Installation guide can be found here for ECA clusters. Find the section that covers the Search & Recover cluster deployment.

© Superna LLC

9.1. What's New with Search & Recover

Home Top

Search & Recover What's New

1. See the new features of Release 1.1.1 here https://www.supernaeyeglass.com/feature-descriptions.
2. Release 1.1.5 will add:
    a. Overview Video
    b. Show Back & Charge Back feature
        i. Scheduled searches from all Search UI's and Quick Reports allows custom search and email results daily, weekly, monthly. Add $ to you search by file pool and send charge back reports to department or project owners
        ii. Schedule any search for security , monitoring usage, data growth
    c. **File Pool Aware search UI** and File pool aware Quick Reports enabled file pool awareness to all existing Quick Reports.
    d. **Classification Engine -** Full custom tagging feature to read the contents of files with regex pattern matching mapped to a custom tag. The Classification tags are automatically populated in all Search UI's and Quick reports allowing security, compliance or custom tagging of data using simple tags. The feature does not require any more disk

space than meta data indexing since the contents of the documents are not full text indexed.

e. **New Quick Report** - Whats in the File Pool? Automatic file pool reporting imports file pool policies directly into the Search interface.

f. Cluster selection in Search UI and Quick Reports. Allows search installation with multiple clusters managed by a single installation to select which cluster search in UI.

g. Snapshot only mode for indexing snapshots only with increased scalability of the number of snapshots that can be monitored for changes.

h. Integration with Golden copy to index copied files with search result CSV files can be imported into Golden Copy to archive data in the search results.

i. BETA - S3 indexing (metadata only) allows indexing object names and properties into the index with the ability to search through objects added through S3 bucket scanning.

j. BETA - File system capacity browser

k. Phone Home Support

© Superna LLC

9.2. Why Eyeglass Search & Recover is the only solution for Scale Out NAS Content Indexing

Home Top

How Eyeglass Search Solves content ingestion challenges

1. **Preserve Last Accessed Attribute**:

   Ingest & index data from automated PowerScale Snapshots preserves the last accessed attribute on files.

2. **Incremental Indexing**: Snapshot comparison allows native PowerScale change file detection feature to provide a list of changed files. The benefits are:

   - Incremental file metadata and content reindexing.
   - Massive reduction in IOPS required to maintain a current index.
   - Ensures files are an exact point time when indexed.
   - Avoids open file issues when indexing files on a live file system.

3. **Security**: Mounting a read only locked snapshot avoids any live file system permissions changes to index the data, closing this security gap in other solutions.

© Superna LLC

9.3. Index Update Intervals, Index Granularity, Index Ignored words, Indexing feature options

Home Top

- Indexing Features
- Index Update Interval
  - Committing Indexed Documents to Disk Interval
  - Index Search Results Interval
- Index Ignored Words
- Unsupported Characters in Folder names , Files and content

Indexing Features

1. **Multiple clusters** are supported by a single Search & Recover cluster with Per Path ingestion capabilities.
2. **Alternate Data Streams**
   a. Only the main data stream will be indexed, any alternate data streams will not be indexed with content
3. **Per Path Ingestion Capabilities:**
   - File system metadata only (default mode).
   - Full content + file system metadata.
4. **Snapshot Aware mode** ingests files in snapshots for historical searching and file version support in the search results:
   - **Snapshot Aware mode** will index files in a snapshot. This mode will *incrementally* index each new snapshot discovered and retain the history of files within the snapshot.

5. **Include or exclude Override** Allows paths, plus wild card or file types, to be skipped or added to the index.

Example 1: A path configured for metadata only can add an **include option** for a specific path, or file type to be full content indexed

Example 2: A path configured for full content can add an **exclude option** of a path or file types to skip processing time on low value data. This will improve performance and reduce the index size.

Index Update Interval

As full and incremental indexing jobs are running and indexing documents, it is important to understand when documents are committed to disk and when indexed content is available for searching.

Committing Indexed Documents to Disk Interval

As documents are being indexed, results are logged in a transaction log. A hard commit is completed that applies all transaction log entries to the master index when either of the following criteria have been met:

- 20 000 documents have been submitted to the index - **OR**
- 5 minutes has passed since the last hard commit to the index

Whichever threshold is hit first triggers a disk write to the master index from the transaction logs.

NOTE: These background processes do not mean the documents will appear in search results. This process ensures data integrity of

the index and ensures on cluster crash that the boot process does not require a lengthy log replay of uncommitted documents to the index.

Index Search Results Interval

Indexed content becomes available for searching as described below:

1. A background process updates the search results cache allowing newly added documents to appear in search results. This operation is an expensive cluster wide process and will be updated based on either of these criteria:

- If **10 Million documents** are added to the index, the result cache will be updated to include these new documents in results - **OR**
- If **30 minutes** have passed since document was indexed

Whichever threshold is hit first will trigger new documents to be available in the index for searching.

Index Ignored Words

When documents are indexed stop words are skipped, these are words that fill the index space but offer low value in searching. For example in the English language stop words include: the, their, a, an etc.. Eyeglass Search has stop words configured for English, German, Spanish, Portuguese and French as defined in the stop words file below.

stop words

Unsupported Characters in Folder names , Files and content

The following characters are currently unsupported if they exist in a file name or directory. If a directory has a special character listed below all child subfolders will not have any files indexed.

1. { } < > ' / [ ] ( ) : " \ # $
2. Content special characters that cannot be used in search syntax: + - && || ! ( ) { } [ ] ^ " ~ * ? : \ / #

© Superna LLC

9.4. Product Requirements, Cluster Sizing and Tested Scaling Limits

Home Top

- Sizing the Search Cluster
- Tested Limits

Sizing the Search Cluster

File Count	Full Content Index or Meta data only Index	ECA Cluster Node Count	Sustained Disk Throughput over 1 minute	Average IO Disk Latency (iostat -xyz)	Disk size
> 250 Million (files content indexed) * Or number of clusters added to search appliance great than 1	Full Content ****	7 ECA nodes with 20 minimum GB RAM per VM **	200 MB/s Read 100 MB/s Write	Avg Read < 10 ms Avg write < 10 ms ***	430 GB x 7 VM's (starting size and will require more disk space to be added over time)
< 250 Million (files content indexed) * Or number of clusters added to search appliance great than	Full Content ****	4 ECA nodes with 20 minimum RAM per VM **	200 MB/s Read 100 MB/s Write	Avg Read < 10 ms Avg write < 10 ms ***	430 GB x 4 VM's (starting size and will require more disk space to be added over time)

1					
> 1 Billion * or number of clusters greater than 1	Meta data	7 ECA nodes nodes with 20G RAM per VM**	200 MB/s Read 100 MB/s Write	Avg Read < 10 ms Avg write < 10 ms ***	430 GB x 7 VM's
< 1 Billion * Or number of clusters added to search appliance great than 1	Meta data	4 ECA nodes nodes with 20G RAM per VM**	200 MB/s Read 100 MB/s Write	Avg Read < 10 ms Avg write < 10 ms ***	430 GB x 4 VM's

*** Note this is starting disk size, content indexing and metadata indexing will require adding more storage as more files are indexed. If disk space utilization reaches 70% all indexing will stop automatically. The above are only starting disk space requirements. Additional RAM is required as per below**

**** Content indexing or high file count or quick reports that run against a high file count requires more RAM per VM 20G is the minimum and file count will determine total ram per node. Heap usage per node must be below 75% used. If Heap usage rises additional RAM will be required for each node. Heap usage can be viewed on node one https://x.x.x.x/solr then select solr cloud**

***** Indexing rate is directly dependent on read latency first and then write latency 2nd. If read latency is above specified values, indexing rate will drop as a factor of the read latency to the disks in the VM.**

****** Content indexing - It is expected to add disk space to the index as different content types, and numbers of files that require content ingestion will increase the index size. This is normal, expected, and supported to add disk space online during indexing operations. Content type and volume of content indexing differences does not allow prediction of disk space required. Rough estimates for planning purposes should assume 10%-20% of the original data size. File formats vary greatly in how much text vs formatting or images a file contains.**

Tested Limits

The following are tested limits. These do not represent actual limits. These numbers will be updated with future releases.

Scaling Limit Item	Tested Value
Number of files or directories meta data only index (requires 7-9 VM configuration)	12 billion
Maximum file size for content Indexing	500 MB
Number of files in a directory	1 Million
Number of subdirectories in a path	1 Million
Number of paths added for indexing	25
Number of clusters added to single appliance for indexing	2
Number Snapshots to Monitor (1.1.5)	25

© Superna LLC

9.5. Use Cases

Home Top

- High level Use Cases
- How to decide which indexing method meets your Use case

High level Use Cases

Search & Recover can solve many use cases with powerful search capabilities and incremental file system change ingestion.

1. **Show Back & Charge Back (1.1.5)**
    a. Show usage by file pool / storage tier
    b. Charge back with scheduled reports showing the cost of data per tier and sum total by user by path or any other combination
2. **Data classification system (1.1.5)**
    a. Custom tagging for compliance PHI, PCI, GDPR or any other compliance tagging or custom tagging solution. Tags are applied using regex matching and full content parsing of over 1000 file types.
    b. Search UI is data classification tag enabled for all reporting searches, including show back and charge back reporting by Data classification.
3. **S3 object indexing (1.1.5)**
    a. Ability to walk an S3 bucket and index the objects in the bucket for file and object unified searching.
4. **End user File Search Portal**

   a. Enables searching on file name, file system attributes, file type attributes, File content, file metadata.
   b. Google Simple search interface.
5. **Administrator E-Discovery Tool**
   a. Enables searching on email, employee id, zip files , pst file support and many more document types to assist with E-Discovery on file based data.
6. **Administrator Search and Script**
   a. Search for files - generate a script that can operate using SmartConnect+share name or /ifs path on the cluster.
   b. Scripts can zip, move, copy, run ISI commands (example worm lock, cloudpool) even archive data easily, or any other action Storage admin needs.
7. **Capacity Management**
   a. Share capacity is growing but who is responsible for the growth?
   b. Search by path summarizes total size owned by all users that own files under the search path.
   c. Use Case #1 - Total Space consumed by user: Run this search monthly, download CSV, sort by largest file total to identify users that are consuming the most disk space month over month.
   d. Use Case #2 - Spaced consumed over a time period: Time period analysis is possible by searching with file created, modified dates within a date range to identify who created the most content within the time period.

8. **End user File Recovery Portal (Release1.1.2 or later)**
    a. Enables indexing files in PowerScale snapshots.
    b. Users can find all versions of a file based on name, path & contents of files.
    c. Users see SmartConnect + share UNC path directly to a file in a snapshot for easy access.

How to decide which indexing method meets your Use case

1. **Metadata only indexing on a path** - Fastest Ingestion mode.
    a. **Supports:** File system User searching/snapshot searching of files , paths, owners, file size, date stamps (created, modified, last accessed)
    b. **Use Cases:**
        i. Users and administrators can find files with wild card searches of any of the fields above. Download research results and Download custom CSV for scripting actions against the file system.
        ii. Capacity Planning reports to find who consumed space on path.
        iii. Stale data searches using last modified or last access date searches.
        iv. File type profiling of the file system data.

2. **Full content indexing on a path** - Metadata ingests in stage 1, full content is slower using 2nd stage ingestion queue.
    a. **Supports:**
        i. Content search of files from the supported file types listed above.
        ii. **NOTE: All searches will return data from all indexed paths based on the users permissions to shares of the returned data. See Secure Search section above. The key words will be used to match against metadata and content depending on how the folder ingestion was configured.**
    b. **Use Cases:**
        i. All of the metadata uses cases plus below.
        ii. **Full Content Ingestion only Use Cases** :
            1. **Contents of files searching for users to locate files faster.**
            2. **Locate duplicate data based on content.**
            3. **Compliance Searches for HIPAA, PCI, eDiscovery.**
            4. **Forensic ssecurity searches associated to an investigation.**

© Superna LLC

9.6. Search & Recover Known Limitations

Home Top

Known Limitations

1. The full path to a file cannot exceed 1012 bytes due to REST API limitations in OneFS. This means data at folder levels beyond this cannot be indexed.
2. SMB Share authentication requires Active Directory provider on Powerscale. LDAP is not supported.
3. Scheduled reports that returns a list of files is limited to 50 000 or less

© Superna LLC

9.7. Search & Recover Solutions Guides

- User Search
- File Automation for PowerScale Administrators
- Capacity Management
- Smartlock Data Reporting
- How to Search for Worm Locked Files Future Expiry
- Snapshot & Replicate Data Protection for PowerScale(FUTURE)

User Search

PowerScale customers can extract the value of data by allowing full text search to find files, duplicates, and collaborate more easily.

Solution: Eyeglass Search & Recover user self serve search portal:

- Find files by the file contents not by file name - this is how people find information on the Internet.

- Find old untouched files with simple 30, 60, 90 day search.
- Find “My files” based on file ownership for logged in user.
- Find information you need created by your colleagues. Collaborate and re-use information that already exists in your organization.
    - Find all high resolution images (meta data search) example all images.

File Automation for PowerScale Administrators

PowerScale administrators need to automate tasks on the file system. Identifying files with scripts and walking the file system is slow and complicated.

Solution: Eyeglass Search & Recover product provides powerful searching based on content, age of file, accessed time, path etc., and assists with script generation to simplify data management tasks.

1. Run ACL change command to a group of files or verify ACL's
2. Zip a group of files (in any folder based on a search criteria) - Example: compliance search
3. Copy, move files (from any path in the file system)
4. Run ISI commands on a group of files - Example: smartlock, cloudpool,
5. Create shadow copy of a set of files for writable snapshot capabilities

Capacity Management

Share capacity is growing but who is responsible for the growth?

Solution: Search by path or time period and summarize total size owned by all users that posses files under the search path.

Use Case #1 - Total Space consumed by user: Run path search monthly, download CSV, sort by largest file total to identify users that are consuming the most disk space month over month.

Use Case #2 - Space consumed over a time period: Time period analysis is possible by searching with file created, modified dates within a date range to identify who created the most content within the time period

Smartlock Data Reporting

Use this procedure to search and find worm data, and generate a report of file lock status using script download and ISI commands.

1. Search for the worm lock data on a path:

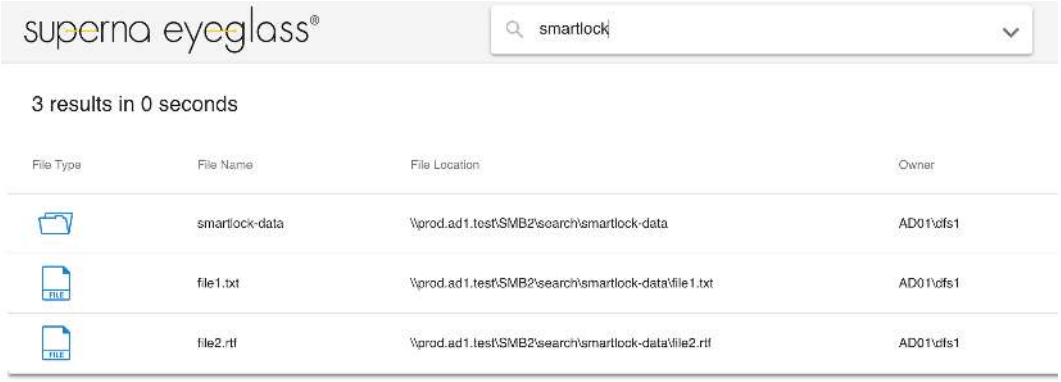

   a.

2. Click the script download icon, and add the ISI commands to query the list of files and redirect the results to a file:

   a. 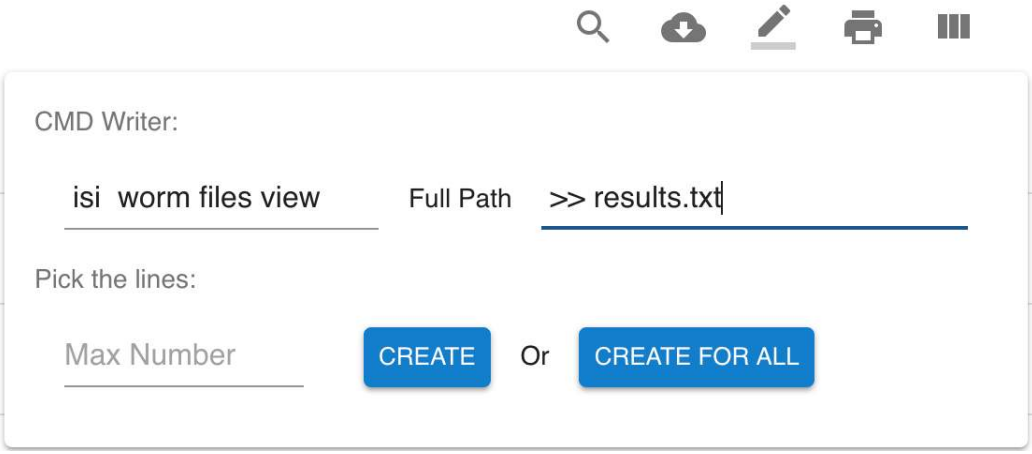

3. See output script example below that will be used to copy and paste or copy the file to the PowerScale for execution:

   a.

```
#!/bin/sh
# Solr Query Summary:
# -Content: smartlock

 isi worm files view "/ifs/data/policy1/search/smartlock-data"   >> results.txt
 isi worm files view "/ifs/data/policy1/search/smartlock-data/file1.txt"   >> results.txt
 isi worm files view "/ifs/data/policy1/search/smartlock-data/file2.rtf"   >> results.txt
```

4. See results output example below from the script:
   a.
```
prod-cluster-8-1# cat results.txt
WORM Domains
ID    Root Path
---------------------------------------------
65553 /ifs/data/policy1/search/smartlock-data
WORM Domains
ID    Root Path
---------------------------------------------
65553 /ifs/data/policy1/search/smartlock-data

WORM State: COMMITTED
   Expires: 2018-11-24T10:36:07
WORM Domains
ID    Root Path
---------------------------------------------
65553 /ifs/data/policy1/search/smartlock-data

WORM State: COMMITTED
   Expires: 2018-11-24T10:36:16
```
5. Done

How to Search for Worm Locked Files Future Expiry

Use this procedure to find locked files and identify all the files that have an expiry on a specific day, or use a date range search to find files on a specific day, month or year. This would be useful to plan or report on files that will be deletable in the future. The script download can be used to automate the delete or archive functions for expired files on a given day.

1. Smartlock sets the Last Accessed Date on files into the future based on the Smartlock domain settings that are configured for files that are committed.

a. 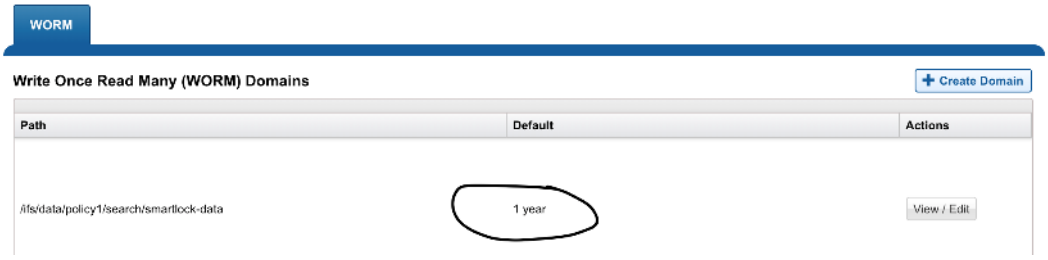

2. See example below Windows file attributes showing the file is locked until 2019:

a. 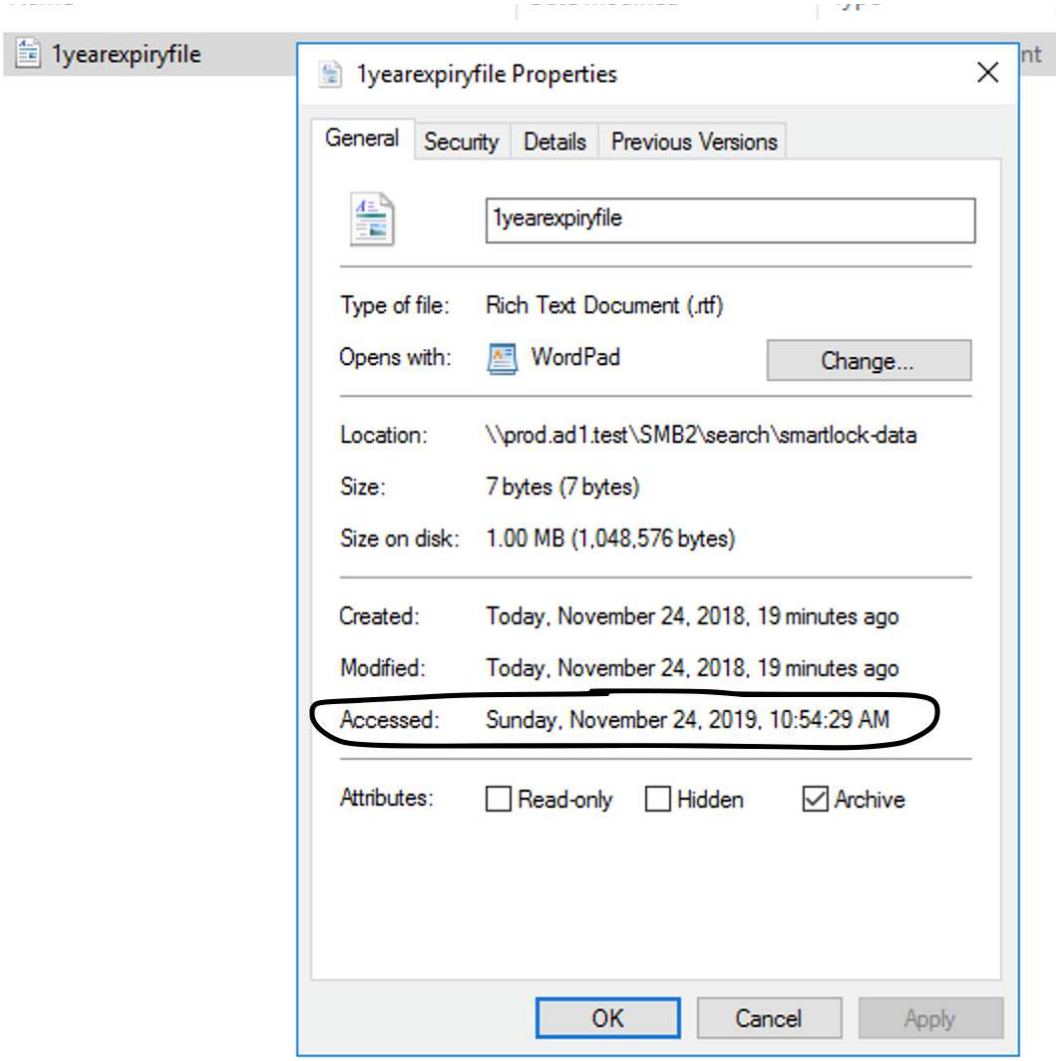

3. Build an advanced Search to find files with expiry in the future using last accessed date query:

a. 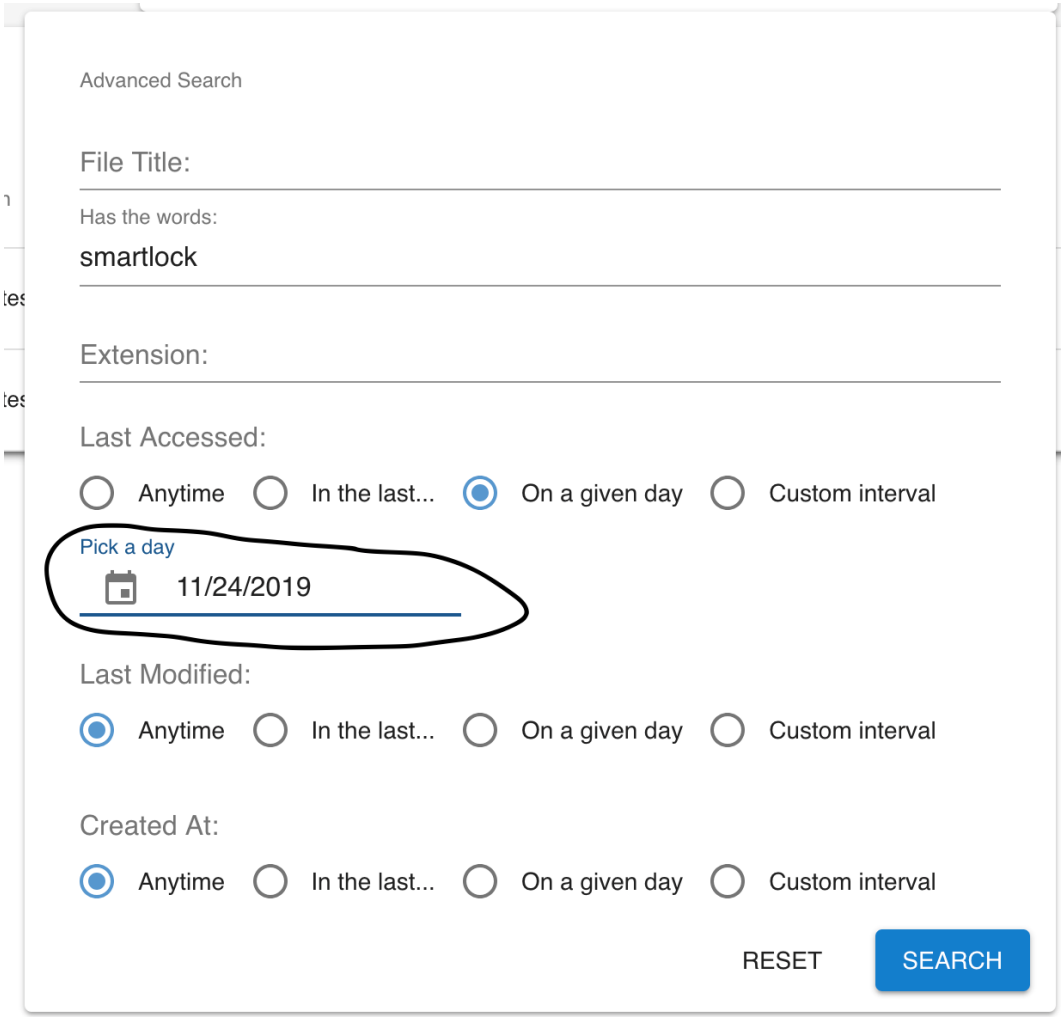

b.

TIP: To search for a date range in the future, select **Custom Interval** and set the End time first (to a date in the future) and then set the Start time. Start time then can also be in the future but before the End time.

Done.

Snapshot & Replicate Data Protection for PowerScale (FUTURE)

PowerScale customers using Snapshot and replicate data protection strategy need a solution to find files in Snapshots to complete the solution.

Solution: Eyeglass Search & Recover user self serve search and recover portal.
Indexes Snapshot data and monitors paths for new or expired snapshots

1. Incrementally indexes files and content on PROD and DR cluster snapshots.
2. User Self Serve Recovery portal finds files in Snapshots and provides a UNC path for users to directly access data in any Snapshot.
3. Historical view of files in expired Snapshots is retained.
4. File version list returned for a file found in multiple Snapshots.

© Superna LLC

9.7.1. Solution Guide - Script Download Examples

Home Top

- Script Example Tools
- Solution to retrieve list of directories with Not Inherited ACL Flag:
  - Eyeglass Search & Recover - Creating Script
  - Configure the script content:
- Solution to retrieve list of directories with “Everyone” ACL:
  - Eyeglass Search & Recover - Creating Script
  - Configure the script content:
- Solution to find Deleted AD user files
- How to move or copy search results to a staging area on the Cluster

Script Example Tools

- Tested with Powershell on Windows Server 2012 R2 (PowerShell 4.0)

Solution to retrieve list of directories with Not Inherited ACL Flag:

This solution will find directories that have the inheritance disabled option in the file system blocking ACL's from flowing down the file system. This can cause unexpected access to files or block access to parts of the file system. This solution will help identify where this flag is disabled.

Eyeglass Search & Recover - Creating Script

1. Login to Eyeglass Search & Recover
2. Enter type:directory as the search keyword

3. Eyeglass Search & Recover will return the list of directories
4. Click CMD Writer Icon

5. Configure the script content:
    1. In First CMD section, enter: "**Get-Acl -Path**"
    2. Select "File Location"
    3. In Second CMD section, enter: "**| select path,owner -ExpandProperty access | where { !$_.IsInherited } | export-csv -path c:\scripts\output2.csv -Append**"
    4. In Script format, select "Plain"
    5. Check the Surround File location in quotes. This will allow a path with spaces to be handled correctly in the script.
    6. Click CREATE For All button
    7. Rename the downloaded file with a .ps1 extension

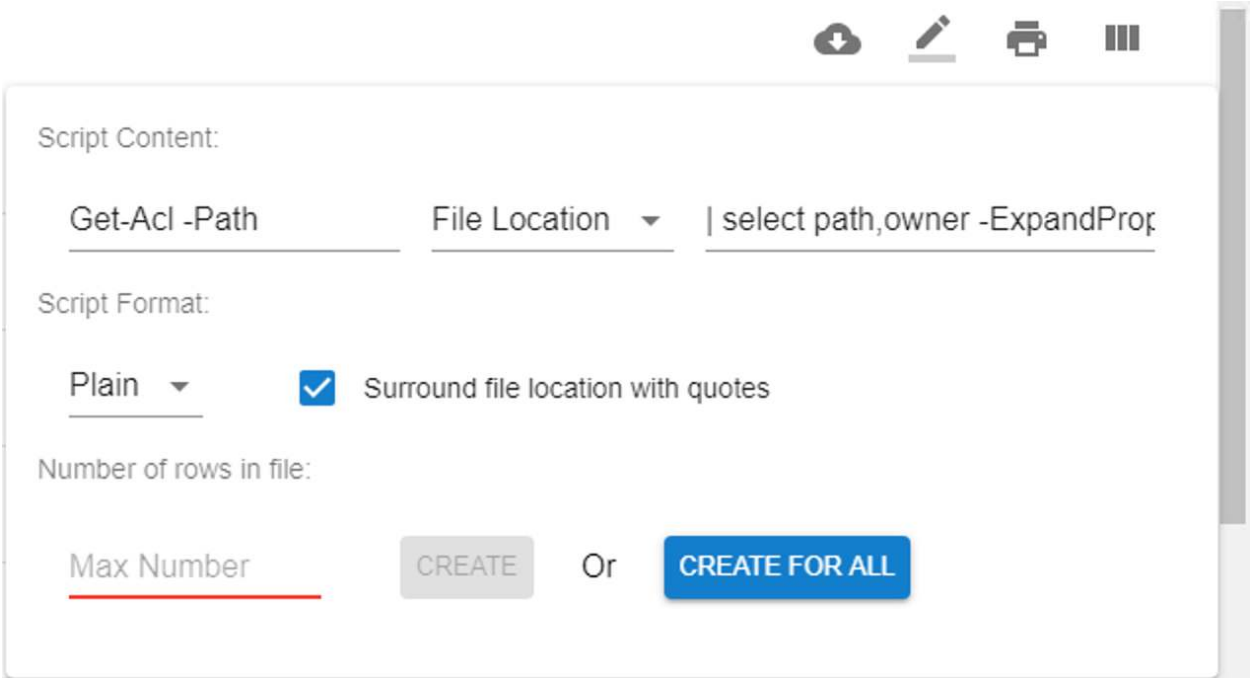

6. Save the modified powershell script
7. Run the script and the output2.csv will content the list of directories with Not Inherited flag
8. See example output below

1. 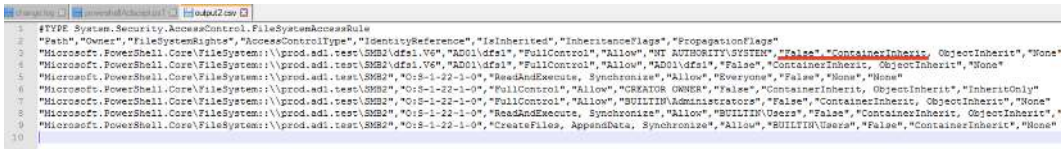

Solution to retrieve list of directories with “Everyone” ACL:

This solution will help find ACL entries that list Everyone, this ACL entry may over expose data and should be used to locate and address potentially over exposed data.

Eyeglass Search & Recover - Creating Script

1. Login to Eyeglass Search & Recover
2. Enter type:directory as the search keyword

3. Eyeglass Search & Recover will return the list of directories

4. Click CMD Writer Icon

5. Configure the script content:

    8. In First CMD section, enter: “**Get-Acl -Path**”
    9. Select “File Location”
    10. In Second CMD section, enter: “**| select path,owner -ExpandProperty access | where {$_.IdentityReference -like "*Everyone*"} | export-csv -path c:\scripts\output2.csv -Append**”
    11. Check the Surround File location in quotes. This will allow a path with spaces to be handled correctly in the script.
    12. Click CREATE For All button
    13. Rename the downloaded file with a .ps1 extension

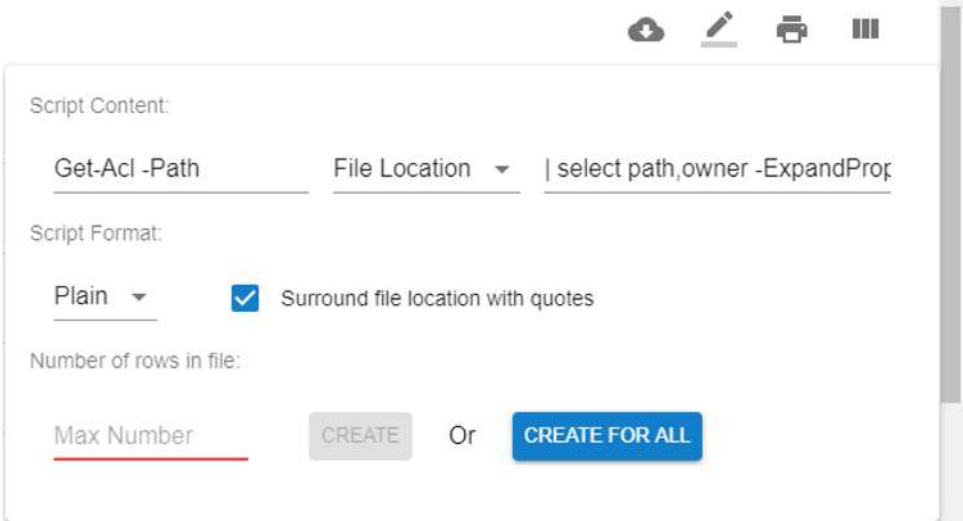

6. A powershell script file will be generated and downloaded
7. Save the modified powershell script
8. Run the script and the output2.csv will content the list of directories “Everyone” ACL

Solution to find Deleted AD user files

When a user is deleted in AD the file ownership on disk will display a UID that will no longer resolve to a friendly name.

1. This requires clear the PowerScale user cache before running the script. This should be done in off peak hours since the cluster will need to resolve users to names once the cache is flushed.
   a. **isi auth user flush**
2. Search for files in a directory administrator login so that file path entry screen is shown (note a search administrator is added using the CLI commands)
   a. 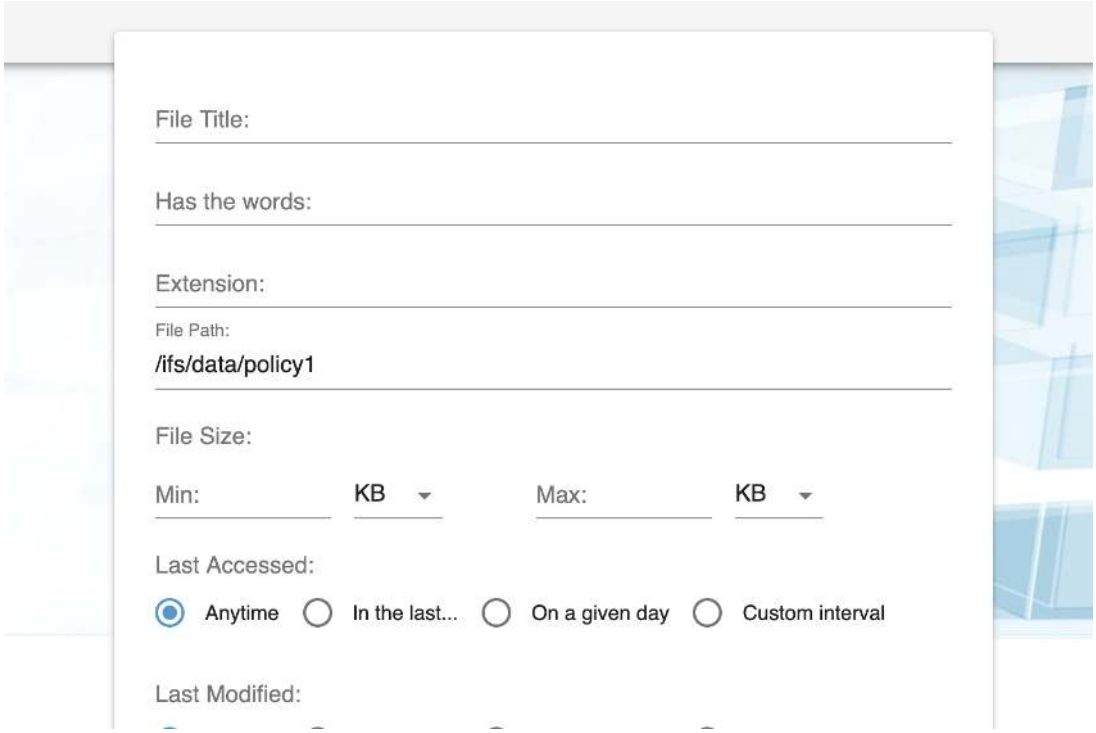

   b. **This will return all files at this path and below (note this can create a very large result) You may need to edit the script file to delete rows for directories you do not need to search)**
3. Click the script editor icon above the search results
4. Enter **find** into the first script content field
5. Then enter **-ls | awk '{print $5"\t\t"$12}' | grep "10" (note this will find all UID's that start with 10)**

a. 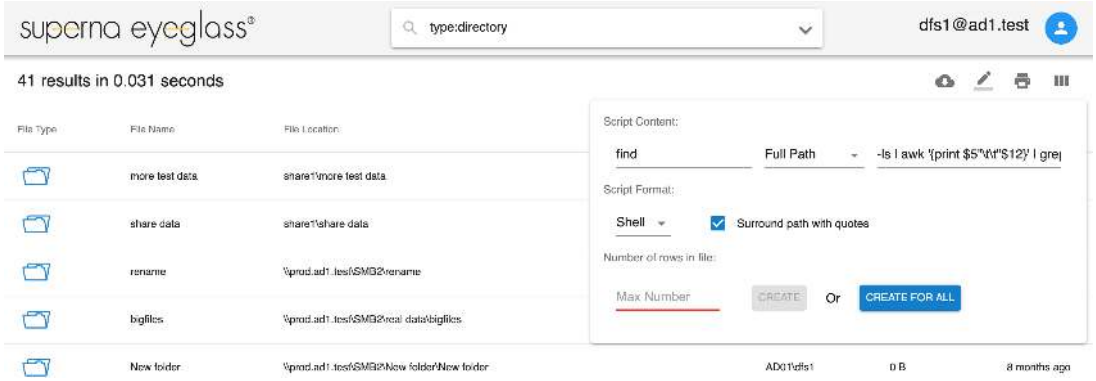

6. Now download by clicking the create for all button
7. Copy this shell script to the cluster and make the file executable with chmod 777 file name
8. Now run the file.
9. **NOTE: this may take some time to run on a lot of files only file names that have a UID in the file system will be output to the console as the script runs. You can modify this script to log results to a file as required.**
10. **Done any file with a UID is not resolved by the cluster and is likely a deleted user in Active Directory.**

How to move or copy search results to a staging area on the Cluster

1. This solution allows searching for files and then copy or move them to a staging area. This could be used for legal discovery of data or archive the data. The key is being able to retain the folder structure of the files in the search results, along with permissions, ownership of the files. The archive requirement can use Superna Eyeglass Golden Copy to sync the staging folder to S3 storage

for long term archive and protect the file system metadata of the files using S3 properties..

2. Execute Seach with the advanced search interface using any of the available search options example find all files with last accessed > 6 months or last modified > 1 year as an example.

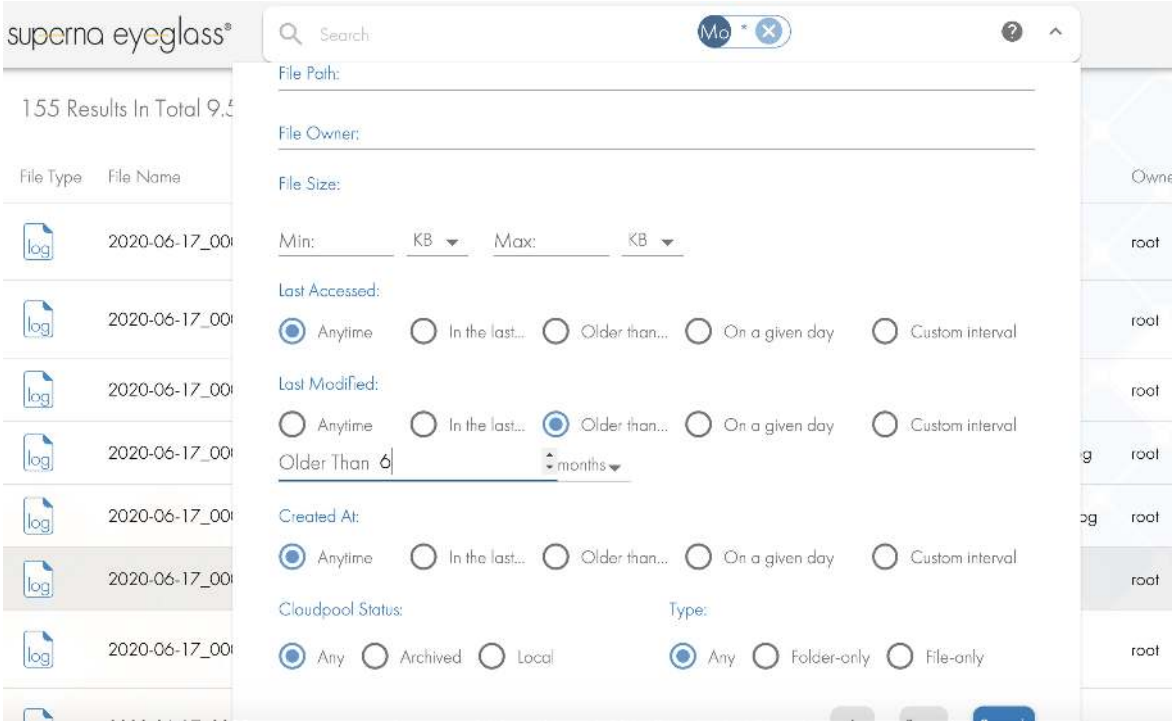

3.

4. Now click the command writer creator

   a. If you want to "**Move**" the search result files to the staging path in this example /ifs/data/staging (the staging directory must already exist) then paste the following into the first box rsync -axuvR --delete-after --progress and then paste /ifs/data/staging into the second box. Now click the "**create for all**" button that will create a bash script for all search results and download the file to your PC.

      i. Change to "**Copy**" The --delete-after flag treats the copy like a move and this parameter is removed it will be a

copy. To copy versus move paste the following into the first box rsync -axuvR --progress

b. 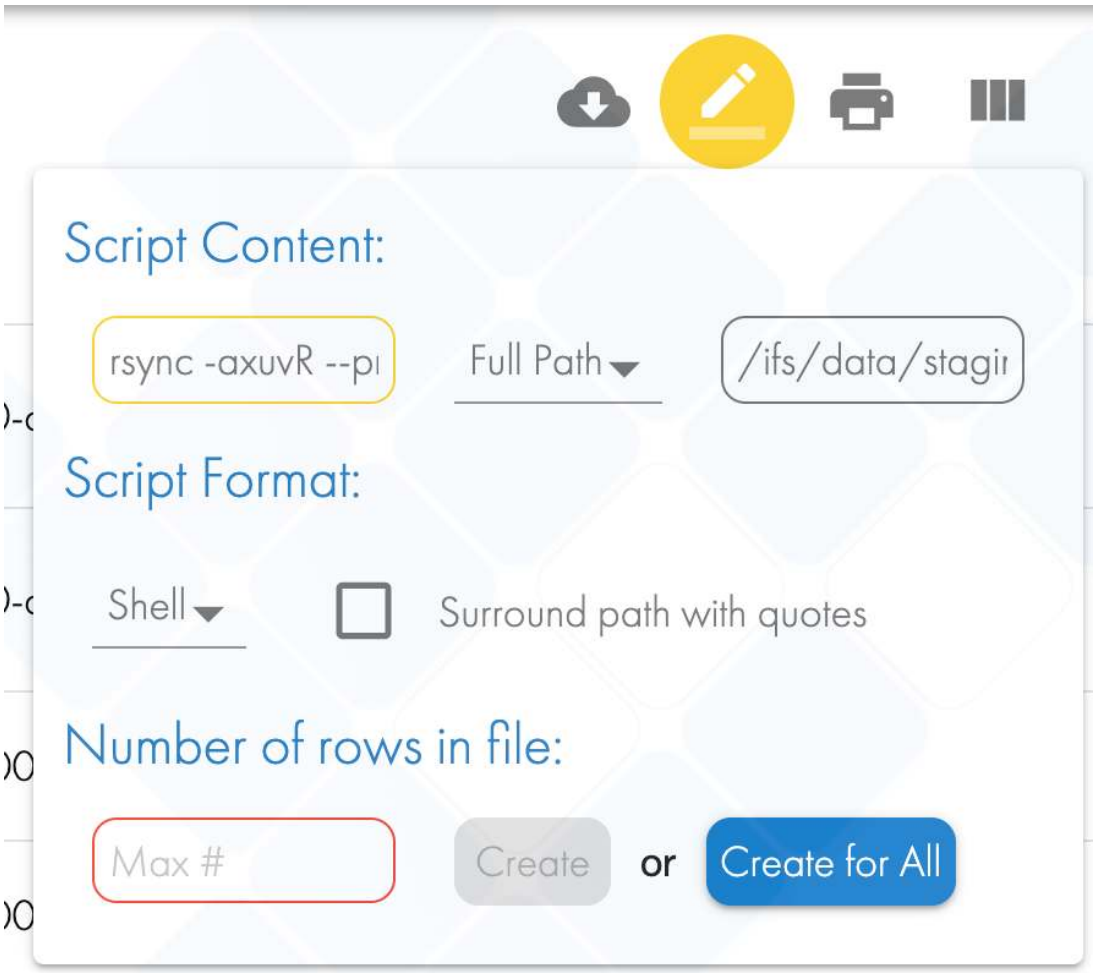

5. Edit the .sh file that is downloaded to verify all the files in the file that will be processed.
6. Copy the .sh file to the cluster as root user. use SCP to copy the file or another method.
7. ssh to the cluster as root user where the script is located
8. chmod 777 script name
9. ./<script name>
10. Verify script executes and then verify the target location has the files.

11. ls /opt/data/staging
12. If using Golden Copy create a sync folder definition with an S3 target to sync all new files copied to the staging path on a schedule.
13. Done

© Superna LLC

9.7.2. Eyeglass Search & Recover - Archive Solution with Dell EMC ECS from Linux Host

Home Top

Eyeglass Search & Recover - Archive Solution with Dell EMC ECS

- Overview
- Solution Test Environment
- ECS Configuration
- S3curl installation
- Create PowerScale NFS mount
- Mount PowerScale NFS to ECA-1
- Create Scripts from Eyeglass Search & Recover to push to ECS and remove from PowerScale

Overview

This solution allows simple scripted upload of data into ECA using Search's powerful searching and script creation feature to simplify the process of bulk copy or move operations into object storage. The tested solution uses NFS to read the files and send over the S3 protocol to ECS storage bucket for archive.

Solution Test Environment

- Eyeglass Search & Recover
- Dell EMC PowerScale (Source)
- Dell EMC ECS v3.3 (Archive)
- s3curl

ECS Configuration

1. Create a new test user: test1

(Manage → Users → New Object Users)

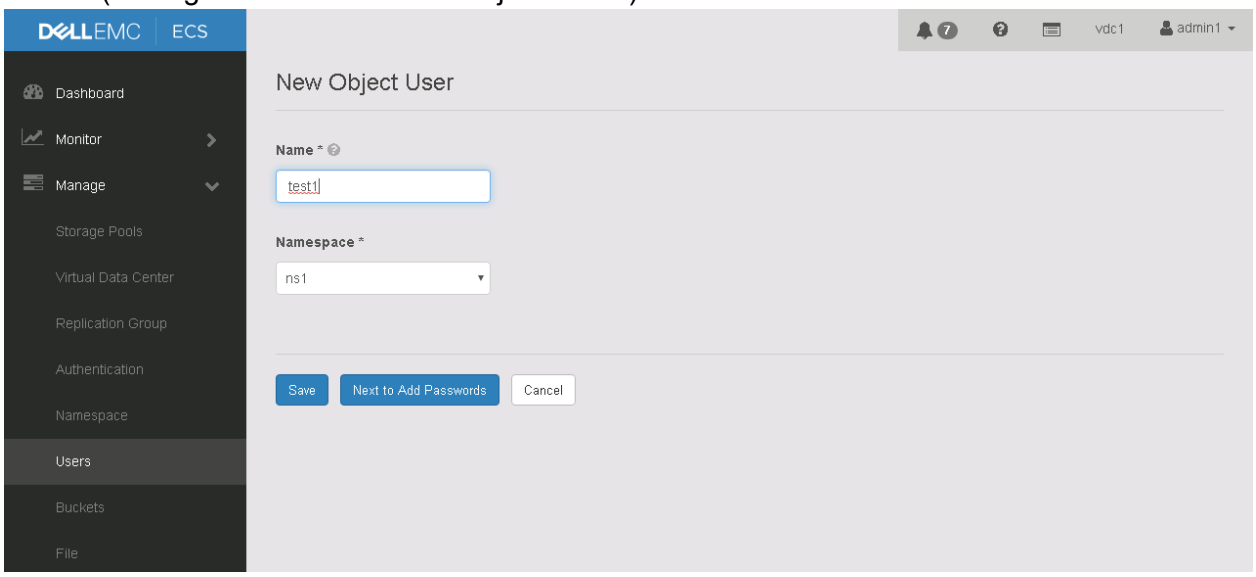

2. Next to Add password → Generate & Add Secret Key (S3)

   Show Secret Key to get the key

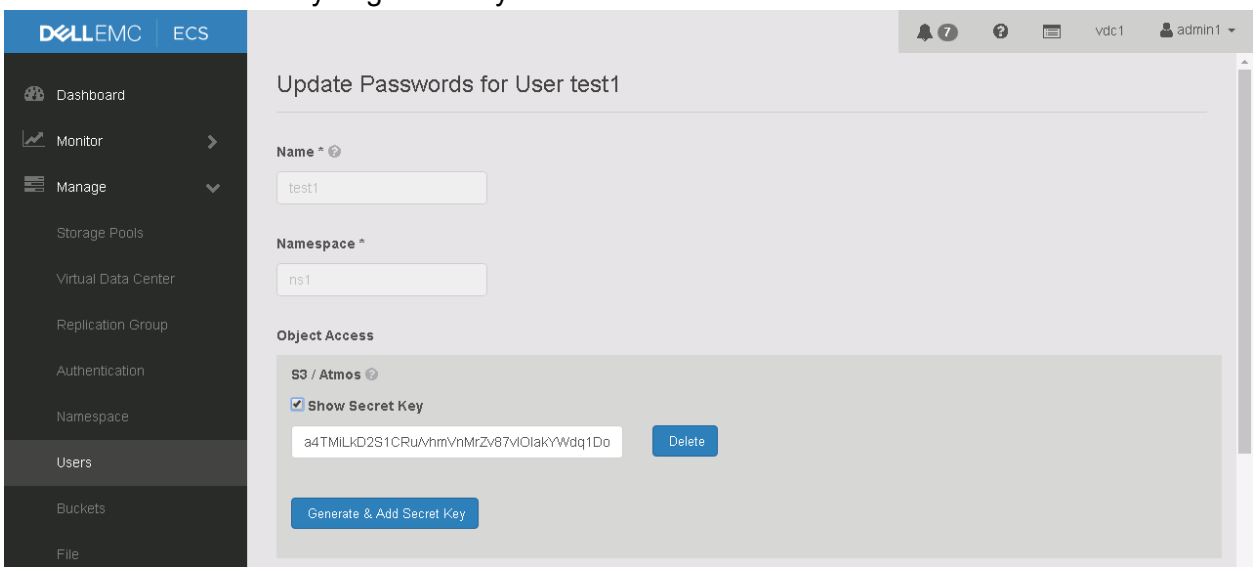

3. Create Bucket

   Manage → Buckets → New Bucket
   Set the Bucket Name, e.g. Bucket2
   And set the owner → test1 user

S3curl installation

1. ssh to ECA node 1
2. sudo su -
3. cd /tmp

4. git clone https://github.com/EMCECS/s3curl.git
5. Set the profile and credentials file, and change the mode of the file

```
cd $HOME
touch ~/.s3curl
chmod 600 ~/.s3curl
vi ~/.s3curl
```

Add these lines into the profile file - specify the userid and secret key

```
%awsSecretAccessKeys = (
 # ECS account
 ecsid => {
 id => 'test1',
 key => 'a4TMiLkD2S1CRu/vhmVnMrZv87vlOlakYWdq1Do7',
 },
);
push @endpoints , (
    '172.22.4.25',
);
```

Create PowerScale NFS mount

Create PowerScale NFS mount for the folder's contents will be archived to ECS.
Configure the ECA-1 IP address as the root client.
Example of the path: /ifs/data/search2/folder101

Mount PowerScale NFS to ECA-1

This mount is used to read the data from PowerScale that will be archived. The Search cluster node 1 can be used for the read and archive funtion

1. ssh to ECA-1
2. sudo su -
3. cd /mnt
4. mkdir folder101
5. Mount. Example

   mount -t nfs -o vers=3 172.22.4.166:/ifs/data/search2/folder101 /mnt/folder101

Create Scripts from Eyeglass Search & Recover to push to ECS and remove from PowerScale

The search used to locate files can be any search including content aware searches or age based search using last accessed or last modified data stamps on files or any combination.

1. Login to Eyeglass Search & Recover
2. Search the folder that the contents will be archive. Example: folder01
3. Click the CMD Writer icon

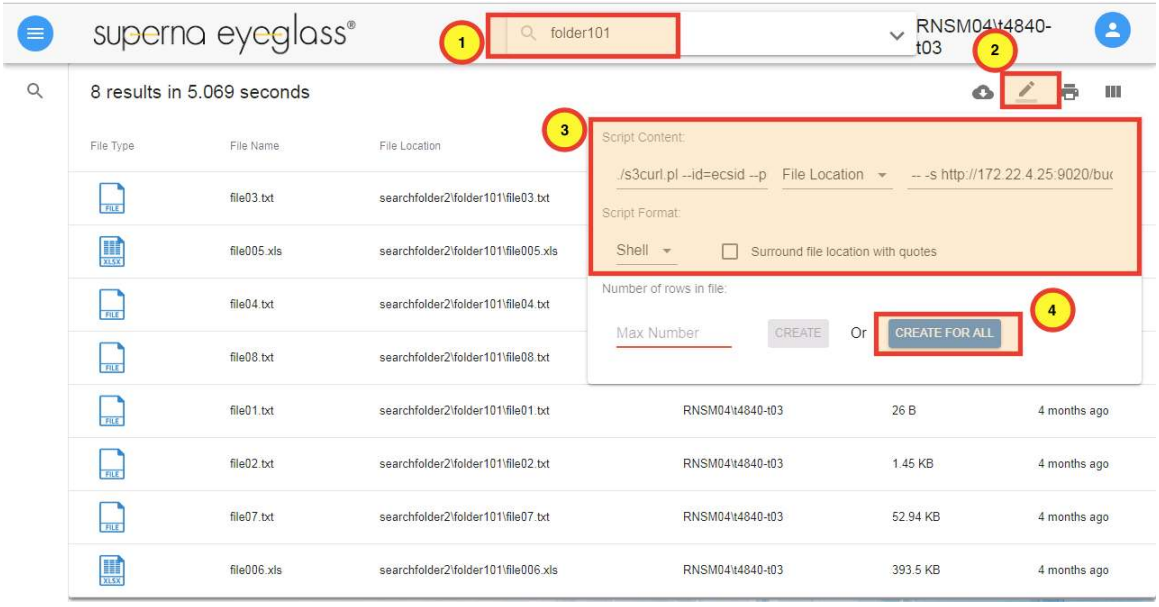

4. Define the Script content as follow:

cmd	Select	cmd
./s3curl.pl --id=ecsid --put=	Full Path	-- -s http://***ecs-node-ip***:9020/***bucket-name/***

Example:

cmd	Select	cmd
./s3curl.pl --id=ecsid --put=	Full Path	-- -s http://172.22.4.25:9020/bucket2/

5. Click “Create For All” button

6. Open and Modify the script use text editor tools to have the correct naming for the file location from the mount directory and the file name that will be placed in ECS bucket

Example:

Before modification :

```
#!/bin/sh
# Solr Query Summary:
# -Content: folder101
./s3curl.pl --id=ecsid --put= /ifs/data/search2/folder101/file03.txt  -- -s http://172.22.4.25:9020/bucket2/
./s3curl.pl --id=ecsid --put= /ifs/data/search2/folder101/file005.xls  -- -s http://172.22.4.25:9020/bucket2/
./s3curl.pl --id=ecsid --put= /ifs/data/search2/folder101/file04.txt  -- -s http://172.22.4.25:9020/bucket2/
./s3curl.pl --id=ecsid --put= /ifs/data/search2/folder101/file08.txt  -- -s http://172.22.4.25:9020/bucket2/
./s3curl.pl --id=ecsid --put= /ifs/data/search2/folder101/file01.txt  -- -s http://172.22.4.25:9020/bucket2/
./s3curl.pl --id=ecsid --put= /ifs/data/search2/folder101/file02.txt  -- -s http://172.22.4.25:9020/bucket2/
./s3curl.pl --id=ecsid --put= /ifs/data/search2/folder101/file07.txt  -- -s http://172.22.4.25:9020/bucket2/
./s3curl.pl --id=ecsid --put= /ifs/data/search2/folder101/file006.xls  -- -s http://172.22.4.25:9020/bucket2/
```

After modification:

```
#!/bin/sh
# Solr Query Summary:
# -Content: folder101
./s3curl.pl --id=ecsid --put=/mnt/folder101/file03.txt  -- -s http://172.22.4.25:9020/bucket2/file03.txt
./s3curl.pl --id=ecsid --put=/mnt/folder101/file005.xls  -- -s http://172.22.4.25:9020/bucket2/file005.xls
./s3curl.pl --id=ecsid --put=/mnt/folder101/file04.txt  -- -s http://172.22.4.25:9020/bucket2/file04.txt
./s3curl.pl --id=ecsid --put=/mnt/folder101/file08.txt  -- -s http://172.22.4.25:9020/bucket2/file08.txt
./s3curl.pl --id=ecsid --put=/mnt/folder101/file01.txt  -- -s http://172.22.4.25:9020/bucket2/file01.txt
./s3curl.pl --id=ecsid --put=/mnt/folder101/file02.txt  -- -s http://172.22.4.25:9020/bucket2/file02.txt
./s3curl.pl --id=ecsid --put=/mnt/folder101/file07.txt  -- -s http://172.22.4.25:9020/bucket2/file07.txt
./s3curl.pl --id=ecsid --put=/mnt/folder101/file006.xls  -- -s http://172.22.4.25:9020/bucket2/file006.xls
```

7. Save the modified file and copy to the ECA-1 node under the same path of the s3curl command (example: /tmp/s3curl/archivescript1.sh)
8. Change mode to executable

   Chmod +x archivescript1.sh
9. Execute the script

   ./archivescript1.sh
10. To verify the content of the bucket, after file copy has been completed, run this command:

    ./s3curl.pl --id=ecsid -- -s http://172.22.4.25:9020/bucket2 | xmllint --format -

    Example of the output:


```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<ListBucketResult xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Name>bucket2</Name>
  <Prefix/>
  <Marker/>
  <MaxKeys>1000</MaxKeys>
  <IsTruncated>false</IsTruncated>
  <ServerSideEncryptionEnabled>false</ServerSideEncryptionEnabled>
  <Contents>
    <Key>file005.xls</Key>
    <LastModified>2019-06-14T10:04:27.802Z</LastModified>
    <ETag>"3ce27fddbbf9a902dd7bbace9ffd5a1b"</ETag>
    <Size>247296</Size>
    <StorageClass>STANDARD</StorageClass>
    <Owner>
      <ID>test1</ID>
      <DisplayName>test1</DisplayName>
    </Owner>
  </Contents>
  <Contents>
    <Key>file01.txt</Key>
    <LastModified>2019-06-21T07:54:37.736Z</LastModified>
    <ETag>"cbc2d26585417e941c806eaf563a685b"</ETag>
    <Size>26</Size>
    <StorageClass>STANDARD</StorageClass>
    <Owner>
      <ID>test1</ID>
      <DisplayName>test1</DisplayName>
    </Owner>
  </Contents>
  <Contents>
    <Key>file02.txt</Key>
    <LastModified>2019-06-17T06:29:50.966Z</LastModified>
    <ETag>"a7bfc78c55cf0a6bffa3cc5f898aa050"</ETag>
    <Size>1481</Size>
    <StorageClass>STANDARD</StorageClass>
    <Owner>
      <ID>test1</ID>
      <DisplayName>test1</DisplayName>
    </Owner>
  </Contents>
  <Contents>
    <Key>file03.txt</Key>
    <LastModified>2019-06-21T08:39:17.591Z</LastModified>
    <ETag>"0bb572b08b954ee1b9047e603f2326b2"</ETag>
    <Size>69</Size>
    <StorageClass>STANDARD</StorageClass>
    <Owner>
      <ID>test1</ID>
      <DisplayName>test1</DisplayName>
    </Owner>
```

```
  </Contents>
  <Contents>
    <Key>file04.txt</Key>
    <LastModified>2019-06-21T08:39:18.558Z</LastModified>
    <ETag>"ca9f024b6bae43f9c05a0357f645c0c2"</ETag>
    <Size>2059</Size>
    <StorageClass>STANDARD</StorageClass>
    <Owner>
      <ID>test1</ID>
      <DisplayName>test1</DisplayName>
    </Owner>
  </Contents>
  <Contents>
    <Key>myfile.txt</Key>
    <LastModified>2019-06-14T09:38:30.256Z</LastModified>
    <ETag>"327922bccbb2249a049608ebda97197c"</ETag>
    <Size>11</Size>
    <StorageClass>STANDARD</StorageClass>
    <Owner>
      <ID>test1</ID>
      <DisplayName>test1</DisplayName>
    </Owner>
  </Contents>
</ListBucketResult>
```


11. Verify that the files are copied to ECS bucket

12. Create another script to remove the files from PowerScale NFS folder

13. Repeat step#3, and define the command as follows:

cmd	Select	cmd
rm -f	Full Path	

14. Click “Create For All” button

15. Open and Modify the script use text editor tools to have the correct naming for the file location from the mount directory

Example:

Before modification :

```
#!/bin/sh
# Solr Query Summary:
# -Content: folder101
rm -f  /ifs/data/search2/folder101/file03.txt
rm -f  /ifs/data/search2/folder101/file005.xls
rm -f  /ifs/data/search2/folder101/file04.txt
rm -f  /ifs/data/search2/folder101/file08.txt
rm -f  /ifs/data/search2/folder101/file01.txt
rm -f  /ifs/data/search2/folder101/file02.txt
rm -f  /ifs/data/search2/folder101/file07.txt
rm -f  /ifs/data/search2/folder101/file006.xls
```

After modification:

```
#!/bin/sh
# Solr Query Summary:
# -Content: folder101
rm -f  /mnt/folder101/file03.txt
rm -f  /mnt/folder101/file005.xls
rm -f  /mnt/folder101/file04.txt
rm -f  /mnt/folder101/file08.txt
rm -f  /mnt/folder101/file01.txt
rm -f  /mnt/folder101/file02.txt
```

```
rm -f  /mnt/folder101/file07.txt
rm -f  /mnt/folder101/file006.xls
```

16. Copy the script to ECA node 1, change to executable and run it to remove file from PowerScale.
17. Verify that file is no longer there.

© Superna LLC

9.7.3. Eyeglass Search & Recover - Archive Solution with Dell EMC ECS from Windows Host

Home Top

Eyeglass Search & Recover - Scripted Archive Solution with Dell EMC ECS (s3cmd and SMB)

- Overview
- Solution Test Environment
- ECS Configuration
- s3cmd installation
- Create PowerScale SMB Share
- Create Scripts from Eyeglass Search & Recover for Push to ECS and Delete From PowerScale

Overview

This solution allows simple scripted upload of data into ECA using Search's powerful searching and script creation feature to simplify the process of bulk copy or move operations into object storage. The tested solution uses SMB to read the files and send over the S3 protocol to ECS storage bucket for archive.

Solution Test Environment

- Eyeglass Search & Recover
- Dell EMC PowerScale (Source)
- Dell EMC ECS v3.3 (Archive)
- s3cmd (Windows platform)

ECS Configuration

1. Create a new test user: test1

   (Manage → Users → New Object Users)

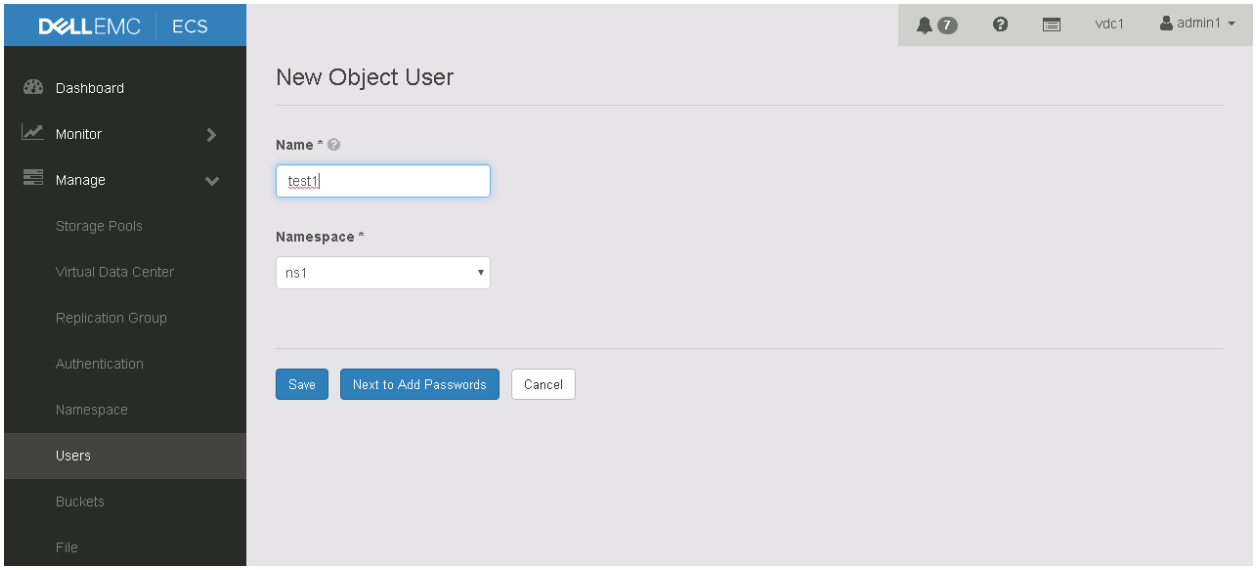

2. Next to Add password → Generate & Add Secret Key (S3)

   Show Secret Key to get the key

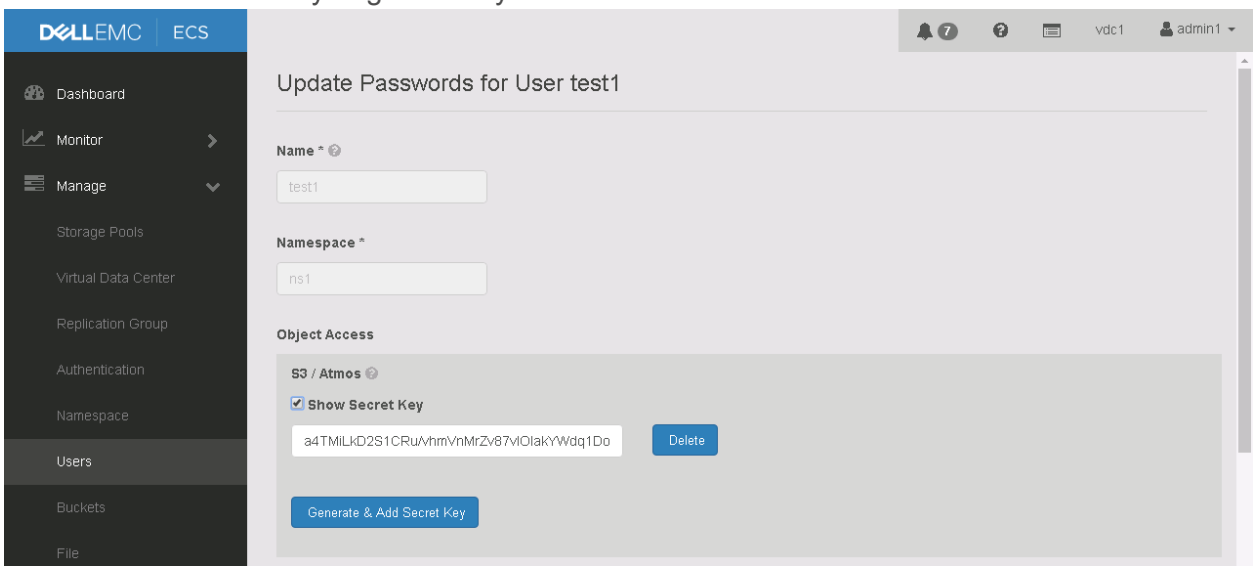

3. Create Bucket

   Manage → Buckets → New Bucket
   Set the Bucket Name, e.g. Bucket2
   And set the owner → test1 user

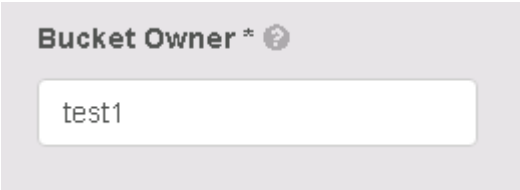

s3cmd installation

1. On a Windows machine, install python (python 2.7 or higher version)
2. Set the environment variable for the path to this python installation directory (e.g. c:\Python27)
3. Download s3cmd from http://s3tools.org/download
4. Extract to c:\s3cmd directory (e.g. c:\s3cmd-2.0.2). Change directory to this directory cd c:\s3cmd-2.0.2
5. Configure s3cmd by using this command:
6. `python s3cmd --configure`
7. (*it may prompt a warning about missing python-dateutil module. To install this python-dateutil module, change directory to python script directory (e.g cd c:\Python27\Scripts\ and then run* `pip install python-dateutil` *command to install that module. Once installed, re-run the s3cmd configuration command from c:\s3cmd-2.0.2 directory*)
8. Specify the following settings:
9. `Access Key: test1` (t***he S3 user that was created before***)
10. `Secret Key: a4TMiLkD2S1CRu/vhmVnMrZv87vlOlakYWdq1Do7` (***secret key for that s3 user***)
11. `Default Region: US` (***accept default***)
12. `S3 Endpoint: ecs25.ad1.test:9020` (***Set the endpoint - DNS name /IP address of our ECS node***)
13. `DNS-style bucket+hostname:port template for accessing a bucket:`
14. `ecs25.ad1.test:9020` (***set the bucket setting***)

15. Save the setting

16. Example:

17. c:\s3cmd-2.0.2>python s3cmd --configure

18. ERROR: Option --preserve is not yet supported on MS Windows platform. Assuming --no-preserve.

19. ERROR: Option --progress is not yet supported on MS Windows platform. Assuming --no-progress.

20. Enter new values or accept defaults in brackets with Enter.

21. Refer to user manual for detailed description of all options.

22. Access key and Secret key are your identifiers for Amazon S3. Leave them empty for using the env variables.

23. Access Key: test1

24. Secret Key: a4TMiLkD2S1CRu/vhmVnMrZv87vlOlakYWdq1Do7

25. Default Region [US]:

26. Use "s3.amazonaws.com" for S3 Endpoint and not modify it to the target Amazon S3.

27. S3 Endpoint [s3.amazonaws.com]: ecs25.ad1.test:9020

28. Use "%(bucket)s.s3.amazonaws.com" to the target Amazon S3. "%(bucket)s" and "%(location)s" vars can be used

29. if the target S3 system supports dns based buckets.

30. DNS-style bucket+hostname:port template for accessing a bucket [%(bucket)s.s3.amazonaws.com]: ecs25.ad1.test:9020

31. Encryption password is used to protect your files from reading

32. by unauthorized persons while in transfer to S3

33. Encryption password:

34. Path to GPG program:

```
When using secure HTTPS protocol all communication with Amazon S3
servers is protected from 3rd party eavesdropping. This method is
slower than plain HTTP, and can only be proxied with Python 2.7 or newer
Use HTTPS protocol [Yes]: No
On some networks all internet access must go through a HTTP proxy.
Try setting it here if you can't connect to S3 directly
HTTP Proxy server name:
New settings:
  Access Key: test1
  Secret Key: a4TMiLkD2S1CRu/vhmVnMrZv87vlOlakYWdq1Do7
  Default Region: US
  S3 Endpoint: ecs25.ad1.test:9020
  DNS-style bucket+hostname:port template for accessing a bucket: ecs25.ad1.test:9020
  Encryption password:
  Path to GPG program: None
  Use HTTPS protocol: False
  HTTP Proxy server name:
  HTTP Proxy server port: 0
Test access with supplied credentials? [Y/n] Y
Please wait, attempting to list all buckets...
Success. Your access key and secret key worked fine :-)
Now verifying that encryption works...
Not configured. Never mind.
```

58. `Save settings? [y/N] y`

59. `Configuration saved to 'C:\Users\administrator\AppData\Roaming\s3cmd.ini'`

60. Verify s3cmd, use this command to list bucket that belongs to that user

61. `c:\s3cmd-2.0.2>python c:\s3cmd-2.0.2\s3cmd ls`

62. `2019-06-14 09:33  s3://bucket2`

Create PowerScale SMB Share

1. Create PowerScale SMB sharefor the folder's contents will be archived to ECS.
2. Example of the path: /ifs/data/search2/
3. And assign the user that will run the s3cmd command to have the permission to read (for upload) and for deleting file after copy to cloud, require read and write permission

Create Scripts from Eyeglass Search & Recover for Push to ECS and Delete From PowerScale

1. The search used to locate files can be any search including content aware searches or age based search using last accessed or last modified data stamps on files or any combination.
   a. Login to Eyeglass Search & Recover
   b. Search the folder that the contents will be archived. Example: folder1
   c. Click the CMD Writer icon
   d. 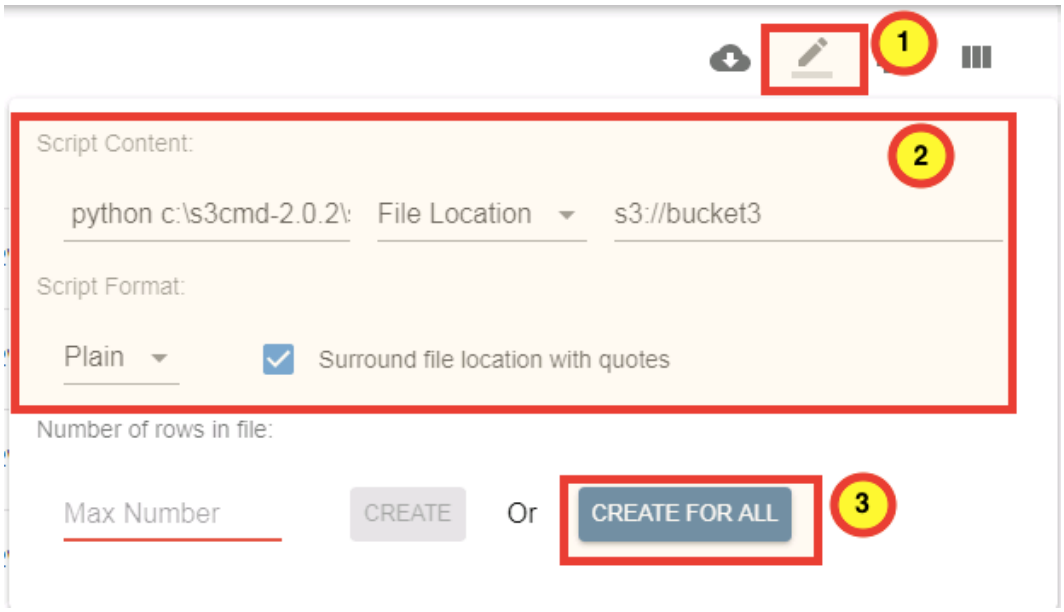

e. Define the Script content as follows:

f.

cmd	Select	cmd	Script Format	Surround file location with quotes
python c:\s3cmd-2.0.2\s3cmd put	File Location	s3:// *<bucket-name>*	Plain	Checked
cmd	**Select**	**cmd**	**Script Format**	**Surround file location with quotes**
python c:\s3cmd-2.0.2\s3cmd put	File Location	s3://bucket2	Plain	Checked

g. Example:

h. Click “Create For All” button

i. Open the script use text editor tools and save it as a batch file

   i. Example:

   ii. # -Content: folder1

   iii. python c:\s3cmd-2.0.2\s3cmd put "\\rnsm04-c08.ad1.test\searchfolder2\folder1\10.txt" s3://bucket2

   iv. python c:\s3cmd-2.0.2\s3cmd put "\\rnsm04-c08.ad1.test\searchfolder2\folder1\1.txt" s3://bucket2

   v. python c:\s3cmd-2.0.2\s3cmd put "\\rnsm04-c08.ad1.test\searchfolder2\folder1\2.txt" s3://bucket2

   vi. python c:\s3cmd-2.0.2\s3cmd put "\\rnsm04-c08.ad1.test\searchfolder2\folder1\3.txt" s3://bucket2

   vii. python c:\s3cmd-2.0.2\s3cmd put "\\rnsm04-c08.ad1.test\searchfolder2\folder1\9.txt" s3://bucket2

   viii. python c:\s3cmd-2.0.2\s3cmd put "\\rnsm04-c08.ad1.test\searchfolder2\folder1\5.txt" s3://bucket2

ix. python c:\s3cmd-2.0.2\s3cmd put "\\rnsm04-c08.ad1.test\searchfolder2\folder1\8.txt" s3://bucket2

x. python c:\s3cmd-2.0.2\s3cmd put "\\rnsm04-c08.ad1.test\searchfolder2\folder1\4.txt" s3://bucket2

xi. python c:\s3cmd-2.0.2\s3cmd put "\\rnsm04-c08.ad1.test\searchfolder2\folder1\6.txt" s3://bucket2

xii. python c:\s3cmd-2.0.2\s3cmd put "\\rnsm04-c08.ad1.test\searchfolder2\folder1\7.txt" s3://bucket2

xiii. Example of the output:

1. c:\s3cmd-2.0.2>python c:\s3cmd-2.0.2\s3cmd put "\\rnsm04-c08.ad1.test\searchfolder2\folder1\3.txt" s3://bucket2
2. WARNING: Module python-magic is not available. Guessing MIME types based on file extensions.
3. upload: '\\rnsm04-c08.ad1.test\searchfolder2\folder1\3.txt' -> 's3://bucket2/3.txt' (15 bytes in 0.6 seconds, 23.40 B/s) [1 of 1]
4. c:\s3cmd-2.0.2>python c:\s3cmd-2.0.2\s3cmd put "\\rnsm04-c08.ad1.test\searchfolder2\folder1\4.txt" s3://bucket2
5. WARNING: Module python-magic is not available. Guessing MIME types based on file extensions.
6. upload: '\\rnsm04-c08.ad1.test\searchfolder2\folder1\4.txt' -> 's3://bucket2/4.txt' (15 bytes in 1.2 seconds, 12.47 B/s) [1 of 1]

j. To verify the content of the bucket, after file copy has been completed, run this command:

i. Python c:\s3cmd-2.0.2\s3cmd ls s3:\\bucket2

ii. Example of the output:

iii. c:\s3cmd-2.0.2>python c:\s3cmd-2.0.2\s3cmd ls s3://bucket2

iv. 2019-06-27 06:34 15 s3://bucket2/1.txt

v. 2019-06-27 07:34 15 s3://bucket2/2.txt

vi. 2019-06-27 07:54 15 s3://bucket2/3.txt

vii. 2019-06-27 07:54 15 s3://bucket2/4.txt

k. Verify that the files are copied to ECS bucket

(Example:. dir \\rnsm04-c08.ad1.test\searchfolder2\folder1)

i. Create another script to remove the files from PowerScale SMB share

ii. Repeat step#3, and define the command as follows:

iii.

cmd	**Select**	**cmd**	**Script Format**	**Surround file location with quotes**
del	File Location		Plain	Checked

iv. Click “Create For All” button

v. Open the script use text editor tools and save it as a batch file

1. Example:
2. # -Content: folder1
3. del "\\rnsm04-c08.ad1.test\searchfolder2\folder1\10.txt"
4. del "\\rnsm04-c08.ad1.test\searchfolder2\folder1\1.txt"
5. del "\\rnsm04-c08.ad1.test\searchfolder2\folder1\2.txt"

6. del "\\rnsm04-c08.ad1.test\searchfolder2\folder1\9.txt"
7. del "\\rnsm04-c08.ad1.test\searchfolder2\folder1\3.txt"
8. del "\\rnsm04-c08.ad1.test\searchfolder2\folder1\5.txt"

vi. Run that batch file script to remove files from PowerScale

vii. Verify that files are no longer on PowerScale

© Superna LLC

9.7.4. How to OCR Image Data with Search & Recover

Home Top

- Overview:
- Requirements:
- Summary:
- Search for OCR Input Data for Processing with Command Builder:
  - Create NFS mount on ECA node 1 for image processing
  - Copy Batch Script to ECA Node for processing Image Data

Overview:

This guide helps walk through how bulk OCR of image data can be done with the Search & Recover command builder and then index the results for searching. This solution depends on an open source OCR library that is available to be installed on the Search & Recover appliance.

Requirements:

Install OCR libraries on ECA node 1 as follows:

1. ssh to node 1 as ecaadmin.
2. sudo -s (enter ecaadmin).
3. zypper install tesseract-ocr (requires Internet connection).
4. Or manual download https://software.opensuse.org/download.html?project=Publishing&package=tesseract-ocr .

Summary:

- This script example shows how command builder can help generate a script file to automate the OCR detection of images based on the Tessact Open ource OCR library, installed on the Search & Recover appliance.
- For larger quantities professional services should be purchased to assist with scripting a parallel solution to multi thread process image data. The example in this guide is quick start guide on how easy OCR solutions can be built with Search & Recover.
- Once the text files are content ingested search results will return the image file name with a txt extension to allow navigation to the folder containing the image.

Search for OCR Input Data for Processing with Command Builder:

1. Using the Search & Recover GUI, locate the OCR data by using any type of search to list the files. It is common to store all OCR scanned data under a single path. This example assumes this is the case.

2. Using the FIle Path option enter the path to the OCR data (i.e. /ifs/data/dfsdata/search/ocr) and add image file extensions (i.e. tif jpg) to a list with spaces to the Extension input box. Click the check box for files only.

3. 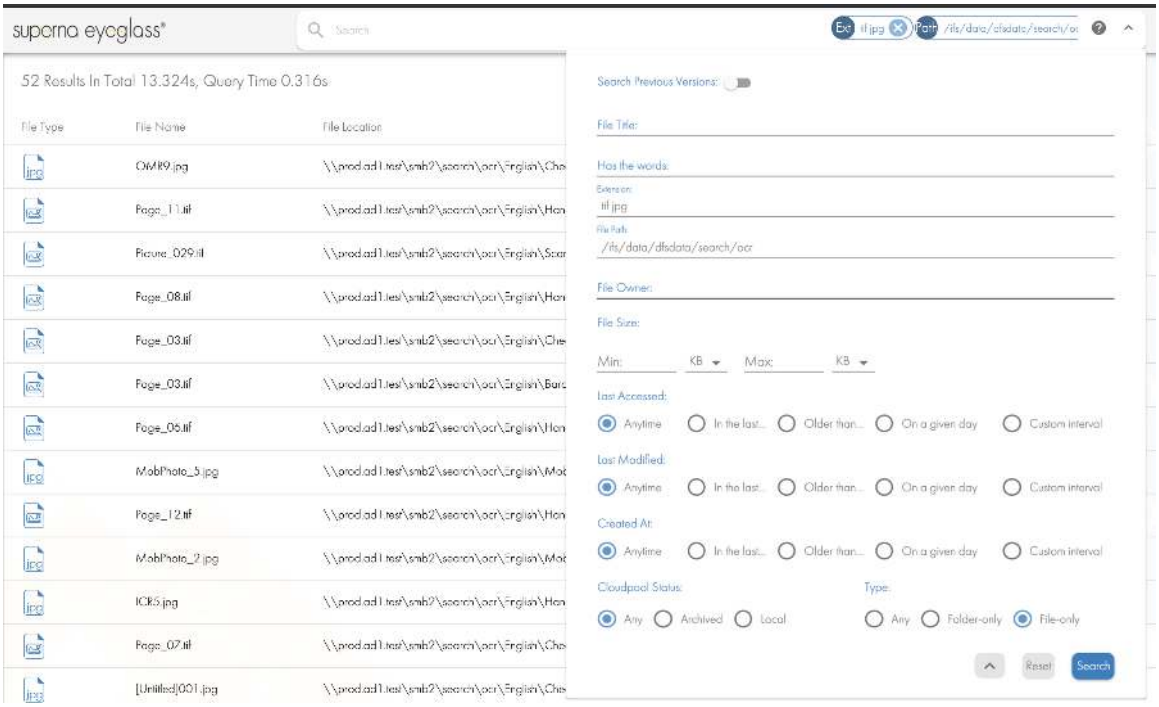

4. Per the screenshot above, this will locate all files with images under the path entered and only list those files in the results.

5. Using the command builder icon generate the file list and enter the OCR command "tesseract" into the first dialog box.

a. 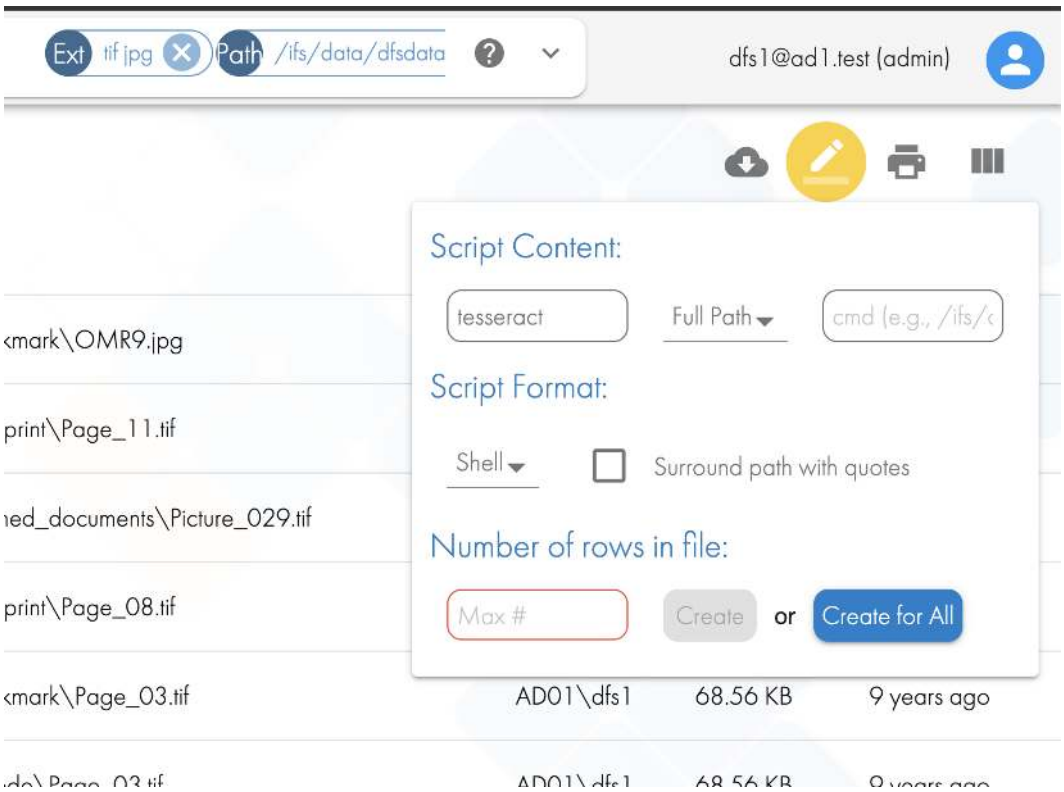

6. Excel is an easy tool to modify the script file to specify the output file name and path. Open the file in Excel and import as CSV using **space** as the separator. See example below. **NOTE: you may need to fix file names with spaces in the path or file name.**

a.

	A	B	C	D	E	F	G	H
2	#	Solr	Query	Summary:				
3	#	#NAME?	Ext.:	tif	jpg			
4	#	#NAME?	Path.:	/ifs/data/dfsdata/search/ocr				
5								
6								
7	tesseract	/ifs/data/dfsdata/search/ocr/English/Scanned_documents/t1.tif						
8	tesseract	/ifs/data/dfsdata/search/ocr/English/Barcode/Barcode_3.jpg						
9	tesseract	/ifs/data/dfsdata/search/ocr/English/Handprint/ICR6.tif						
10	tesseract	/ifs/data/dfsdata/search/ocr/English/Scanned_documents/Picture_025.tif						
11	tesseract	/ifs/data/dfsdata/search/ocr/English/Checkmark/Page_06.tif						
12	tesseract	/ifs/data/dfsdata/search/ocr/English/Handprint/Page_03.tif						
13	tesseract	/ifs/data/dfsdata/search/ocr/English/Barcode/barcode3.tif						
14	tesseract	/ifs/data/dfsdata/search/ocr/English/Handprint/Page_02.tif						
15	tesseract	/ifs/data/dfsdata/search/ocr/English/Barcode/barcode1.jpg						
16	tesseract	/ifs/data/dfsdata/search/ocr/English/Business_cards/doc10002.tif						
17	tesseract	/ifs/data/dfsdata/search/ocr/English/Handprint/Page_09.tif						
18	tesseract	/ifs/data/dfsdata/search/ocr/English/Handprint/Page_04.tif						
19	tesseract	/ifs/data/dfsdata/search/ocr/English/Handprint/Page_07.tif						
20	tesseract	/ifs/data/dfsdata/search/ocr/English/Mobile_Photos/MobPhoto_4.jpg						
21	tesseract	/ifs/data/dfsdata/search/ocr/English/Scanned_documents/New Image.jpg						
22	tesseract	/ifs/data/dfsdata/search/ocr/English/Handprint/Page_13.tif						
23	tesseract	/ifs/data/dfsdata/search/ocr/English/Checkmark/OMR10.tif						
24	tesseract	/ifs/data/dfsdata/search/ocr/English/Handprint/ICR3.jpg						
25	tesseract	/ifs/data/dfsdata/search/ocr/English/Barcode/Page_06.tif						
26	tesseract	/ifs/data/dfsdata/search/ocr/English/Checkmark/OMR5.tif						
27	tesseract	/ifs/data/dfsdata/search/ocr/English/Mobile_Photos/IMG_0122.jpg						
28	tesseract	/ifs/data/dfsdata/search/ocr/English/Business_cards/doc10003.tif						
29	tesseract	/ifs/data/dfsdata/search/ocr/English/Checkmark/OMR7.jpg						
30	tesseract	/ifs/data/dfsdata/search/ocr/English/Checkmark/OMR9.jpg						
31	tesseract	/ifs/data/dfsdata/search/ocr/English/Handprint/Page_08.tif						
32	tesseract	/ifs/data/dfsdata/search/ocr/English/Checkmark/Page_03.tif						
33	tesseract	/ifs/data/dfsdata/search/ocr/English/Barcode/Page_03.tif						
34	tesseract	/ifs/data/dfsdata/search/ocr/English/Handprint/Page_06.tif						

b. Now copy column B files to Column C and it should look like this image below. Save the file as .sh text file.

c.

	A	B	C	D	E	F	G	H
1	#!/bin/sh							
2	#	Solr	Query	Summary:				
3	#	#NAME?	Ext.:	tif	jpg			
4	#	#NAME?	Path.:	/ifs/data/dfsdata/search/ocr				
5								
6								
7	tesseract	/ifs/data/dfsdata/search/ocr/English/Scanned_documents/t1.tif	/ifs/data/c	fsdata/search/ocr/English/Scanned_documents/t1.tif				
8	tesseract	/ifs/data/dfsdata/search/ocr/English/Barcode/Barcode_3.jpg	/ifs/data/c	fsdata/search/ocr/English/Barcode/Barcode_3.jpg				
9	tesseract	/ifs/data/dfsdata/search/ocr/English/Handprint/ICR6.tif	/ifs/data/c	fsdata/search/ocr/English/Handprint/ICR6.tif				
10	tesseract	/ifs/data/dfsdata/search/ocr/English/Scanned_documents/Picture_025.tif	/ifs/data/c	fsdata/search/ocr/English/Scanned_documents/Picture_025.tif				
11	tesseract	/ifs/data/dfsdata/search/ocr/English/Checkmark/Page_06.tif	/ifs/data/c	fsdata/search/ocr/English/Checkmark/Page_06.tif				
12	tesseract	/ifs/data/dfsdata/search/ocr/English/Handprint/Page_03.tif	/ifs/data/c	fsdata/search/ocr/English/Handprint/Page_03.tif				
13	tesseract	/ifs/data/dfsdata/search/ocr/English/Barcode/barcode3.tif	/ifs/data/c	fsdata/search/ocr/English/Barcode/barcode3.tif				
14	tesseract	/ifs/data/dfsdata/search/ocr/English/Handprint/Page_02.tif	/ifs/data/c	fsdata/search/ocr/English/Handprint/Page_02.tif				
15	tesseract	/ifs/data/dfsdata/search/ocr/English/Barcode/barcode1.jpg	/ifs/data/c	fsdata/search/ocr/English/Barcode/barcode1.jpg				
16	tesseract	/ifs/data/dfsdata/search/ocr/English/Business_cards/doc10002.tif	/ifs/data/c	fsdata/search/ocr/English/Business_cards/doc10002.tif				
17	tesseract	/ifs/data/dfsdata/search/ocr/English/Handprint/Page_09.tif	/ifs/data/c	fsdata/search/ocr/English/Handprint/Page_09.tif				
18	tesseract	/ifs/data/dfsdata/search/ocr/English/Handprint/Page_04.tif	/ifs/data/c	fsdata/search/ocr/English/Handprint/Page_04.tif				
19	tesseract	/ifs/data/dfsdata/search/ocr/English/Handprint/Page_07.tif	/ifs/data/c	fsdata/search/ocr/English/Handprint/Page_07.tif				
20	tesseract	/ifs/data/dfsdata/search/ocr/English/Mobile_Photos/MobPhoto_4.jpg	/ifs/data/c	fsdata/search/ocr/English/Mobile_Photos/MobPhoto_4.jpg				
21	tesseract	/ifs/data/dfsdata/search/ocr/English/Scanned_documents/New Image.jpg	/ifs/data/c	fsdata/search/ocr/English/Scanned_documents/New Image.jpg				
22	tesseract	/ifs/data/dfsdata/search/ocr/English/Handprint/Page_13.tif	/ifs/data/c	fsdata/search/ocr/English/Handprint/Page_13.tif				
23	tesseract	/ifs/data/dfsdata/search/ocr/English/Checkmark/OMR10.tif	/ifs/data/c	fsdata/search/ocr/English/Checkmark/OMR10.tif				
24	tesseract	/ifs/data/dfsdata/search/ocr/English/Handprint/ICR3.jpg	/ifs/data/c	fsdata/search/ocr/English/Handprint/ICR3.jpg				
25	tesseract	/ifs/data/dfsdata/search/ocr/English/Barcode/Page_06.tif	/ifs/data/c	fsdata/search/ocr/English/Barcode/Page_06.tif				
26	tesseract	/ifs/data/dfsdata/search/ocr/English/Checkmark/OMR5.tif	/ifs/data/c	fsdata/search/ocr/English/Checkmark/OMR5.tif				
27	tesseract	/ifs/data/dfsdata/search/ocr/English/Mobile_Photos/IMG_0122.jpg	/ifs/data/c	fsdata/search/ocr/English/Mobile_Photos/IMG_0122.jpg				
28	tesseract	/ifs/data/dfsdata/search/ocr/English/Business_cards/doc10003.tif	/ifs/data/c	fsdata/search/ocr/English/Business_cards/doc10003.tif				
29	tesseract	/ifs/data/dfsdata/search/ocr/English/Checkmark/OMR7.jpg	/ifs/data/c	fsdata/search/ocr/English/Checkmark/OMR7.jpg				
30	tesseract	/ifs/data/dfsdata/search/ocr/English/Checkmark/OMR9.jpg	/ifs/data/c	fsdata/search/ocr/English/Checkmark/OMR9.jpg				
31	tesseract	/ifs/data/dfsdata/search/ocr/English/Handprint/Page_08.tif	/ifs/data/c	fsdata/search/ocr/English/Handprint/Page_08.tif				
32	tesseract	/ifs/data/dfsdata/search/ocr/English/Checkmark/Page_03.tif	/ifs/data/c	fsdata/search/ocr/English/Checkmark/Page_03.tif				
33	tesseract	/ifs/data/dfsdata/search/ocr/English/Barcode/Page_03.tif	/ifs/data/c	fsdata/search/ocr/English/Barcode/Page_03.tif				

d. You may need to save as CSV and then use a text editor to search and replace the comma for a space. You can also remove the comments at the top of the file.

7. Create NFS mount on ECA node 1 for image processing

8. An NFS mount is needed on the cluster to allow Search & Recover to OCR the images, and create a .txt version of the file. This mount will need root mount options. See steps below to create the NFS export on /ifs.

9. The screenshot shows the NFS root client export on path /ifs/ :

a. 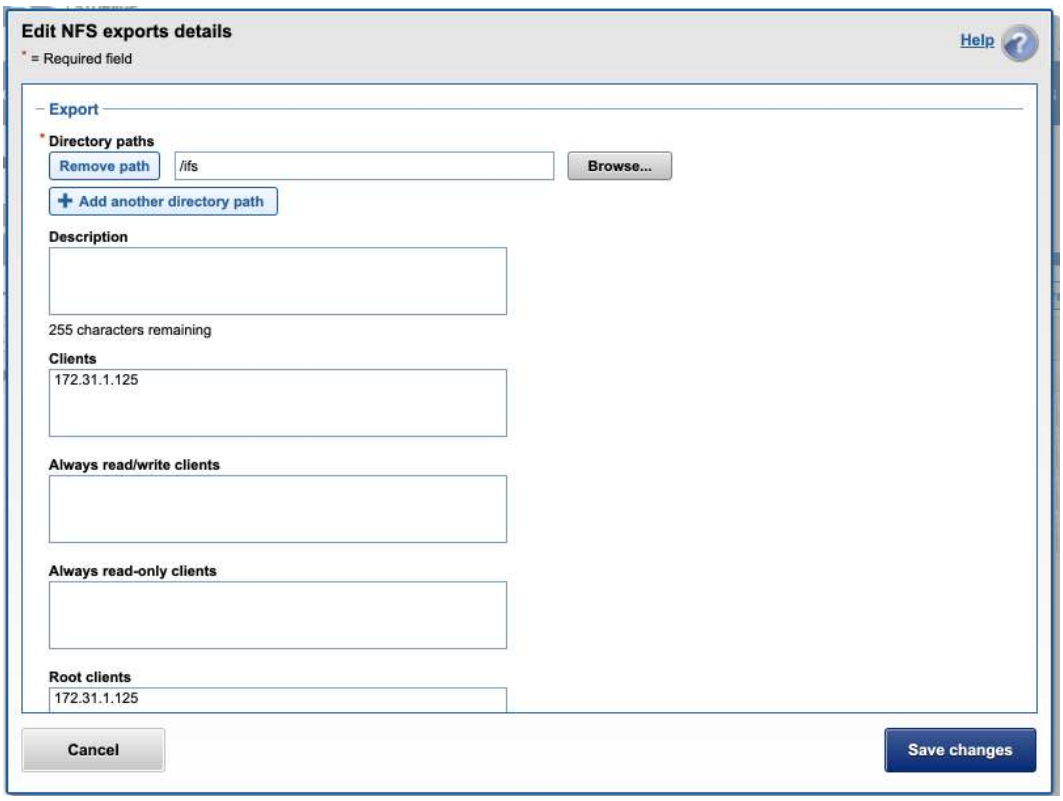

10. Now create the mount point on Node 1 of the Search & Recover appliance.
11. ssh to ecaadmin node 1 as ecaadmin.
12. sudo -s (enter ecaadmin password to become root user).
13. mkdir -p /ifs .
14. Mount the cluster with this command (**NOTE: use /etc/fstab for a persistent mount point to handle reboots**)
    a. mount 172.31.1.104:/ifs /ifs (**Note: use SmartConnect name vs ip address used in the example**)
15. Verify by typing "mount".
16. Verify with ls /ifs to make sure you see files and directories returned.

17. Copy Batch Script to ECA Node for processing Image Data

18. Copy script file to Search & Recover node 1 with scp or winscp tool using ecaadmin user the file will be copied to /home/ecaadmin.
19. **Example with scp from the command line:** scp ocr.sh ecaadmin@172.31.1.125:/home/ecaadmin/ocr.sh .
20. Change permissions:
    a. chmod 777 /home/ecaadmin/ocr.sh
21. Execute the OCR conversion (NOTE: This can take a long time to complete, potentially hours).
22. cd /home/ecaadmin/ .
23. ./ocr.sh &> results.txt .
24. Monitor progress with this command:
    a. tail -f /home/ecaadmin/results.txt .
25. Once your script finishes you will have a file matching the same file name with .txt added to the file. The .txt file will contain the text extracted from the image file.
26. Search & Recover incremental will detect the new .txt files and index the content of the files if content ingestion is enabled on the OCR path.

© Superna LLC

9.7.5. Writeable Snapshots with File Clones Automation

Home Top

- Overview
- How to create writeable Snapshot of data for testing
- How to Verify the Cloned Data Consumes no additional space
- Browse the Writeable Clone data for testing

Overview

A very common requirement is to test on a large data set and making full copies is very slow and consumes a lot of space. This solution guide will walk through how to clone a directory tree of data to make use of File Clone feature in OneFS to create writeable copies of data BUT without consuming space of the original data set. This is done with the Search & Recover Command builder tool, Excel and a text editor. The cloned data only consumes space based on the changes to the files at the block level. This is a copy on first write solution.

How to create writeable Snapshot of data for testing

1. This example will use a source path of /ifs/data/dfsdata/search

2. 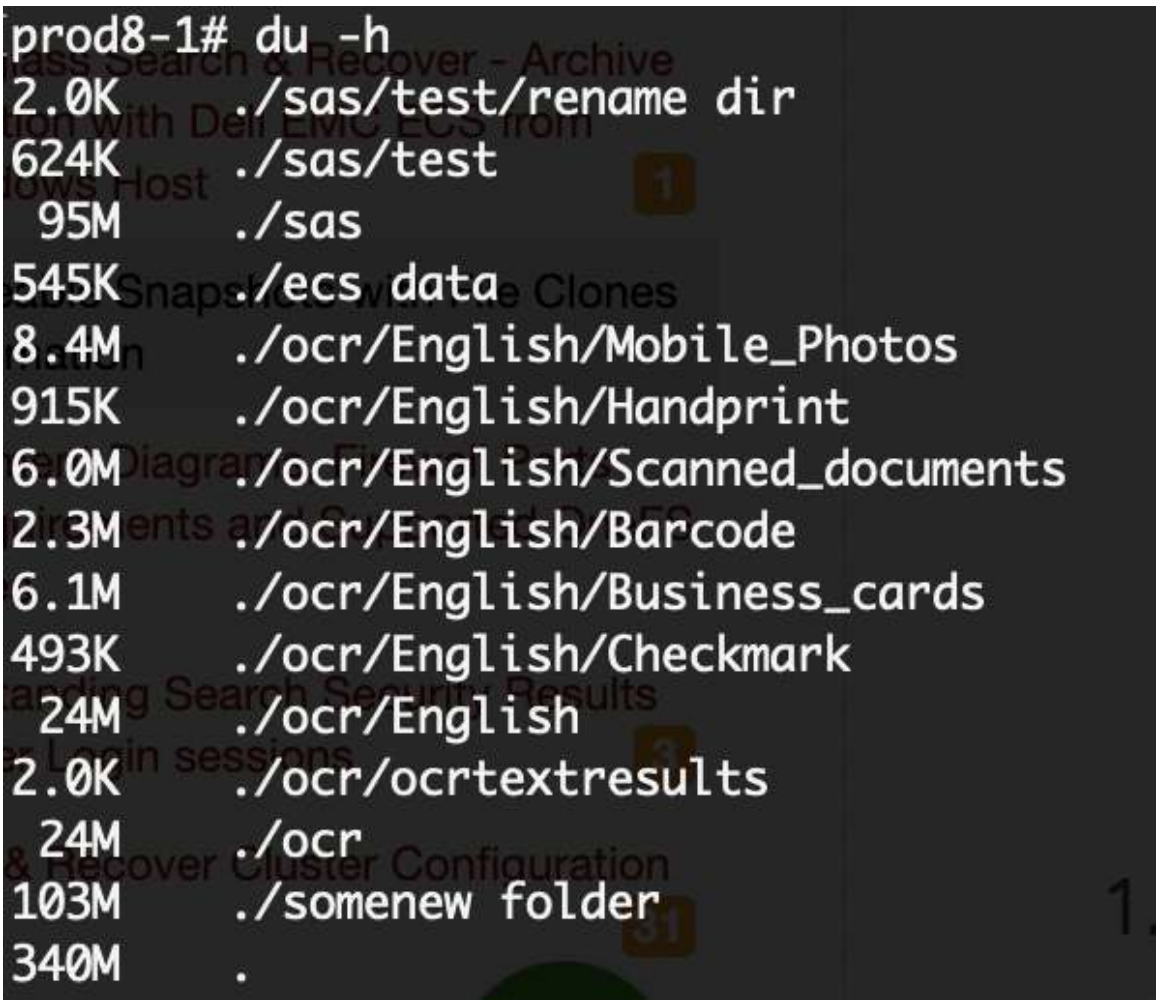

```
prod8-1# du -h
2.0K    ./sas/test/rename dir
624K    ./sas/test
 95M    ./sas
545K    ./ecs data
8.4M    ./ocr/English/Mobile_Photos
915K    ./ocr/English/Handprint
6.0M    ./ocr/English/Scanned_documents
2.3M    ./ocr/English/Barcode
6.1M    ./ocr/English/Business_cards
493K    ./ocr/English/Checkmark
 24M    ./ocr/English
2.0K    ./ocr/ocrtextresults
 24M    ./ocr
103M    ./somenew folder
340M    .
```

3. NOTE: This folder path has 340 MB of data
4. Create the target folder structure where the cloned data will reside
    a. ssh to the cluster as root or admin
    b. cd /ifs/data/dfsdata/search (change the path from this example)
    c. rsync -av -f"+ */" -f"- *" . /ifs/data/dfsdata/clonedsearch/ (change target path to the location where you plan to create the writeable clones)
5. In the Search & Recover GUI advanced window enter the fully qualified path into the File path field where the data is located. In

this example /ifs/data/dfsdata/search. Then select the files only option in the advanced search UI and then execute the search.

a. 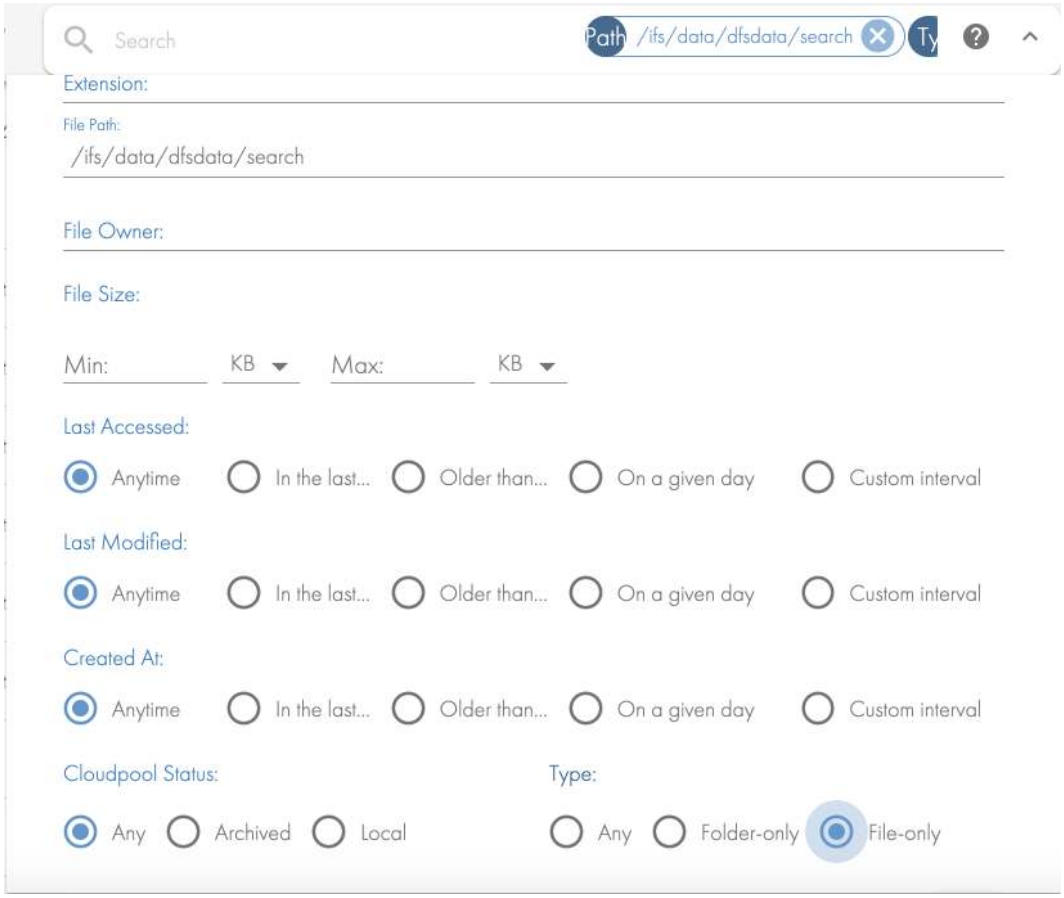

6. Click the Command build icon and enter this into the first box "cp -c", click the check box to add double quotes, switch to plain text option and enter nothing into the second box and click create all

a. 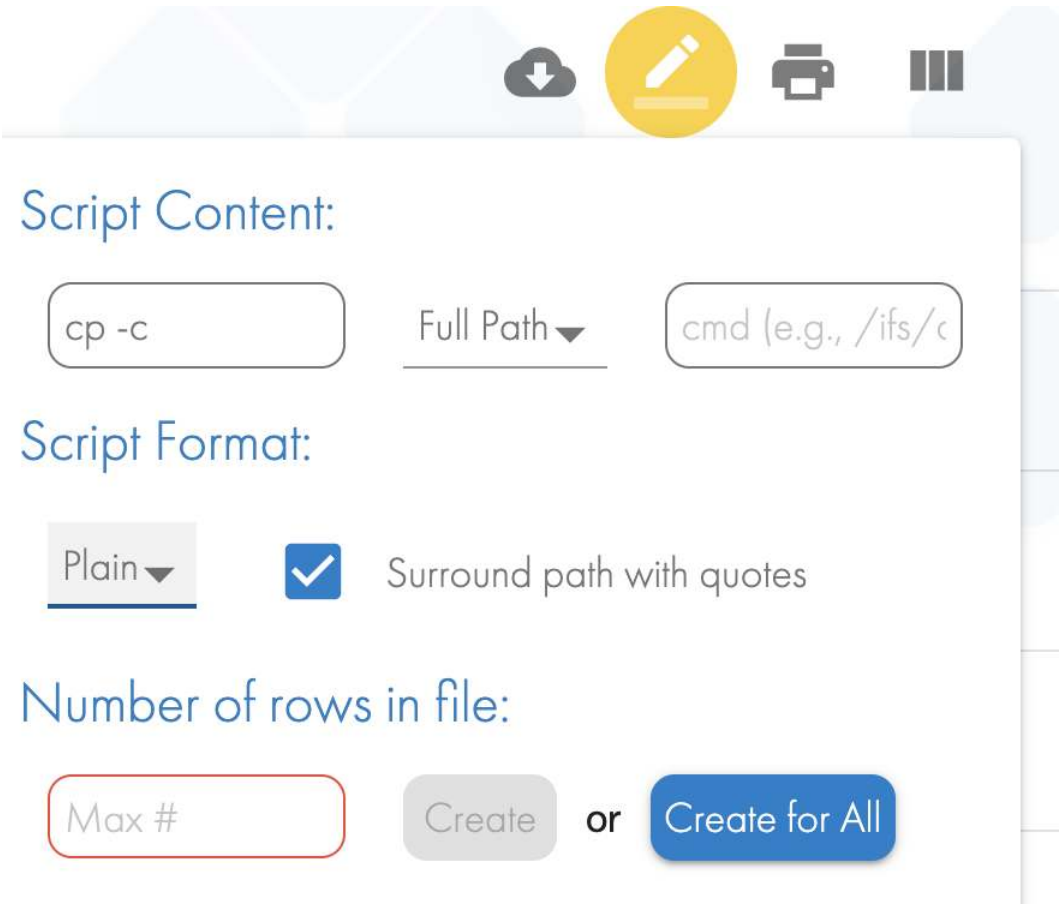

7. Excel is an easy tool to modify the script file to specify the output file name and path.

   a. Import from data file default settings

   b. 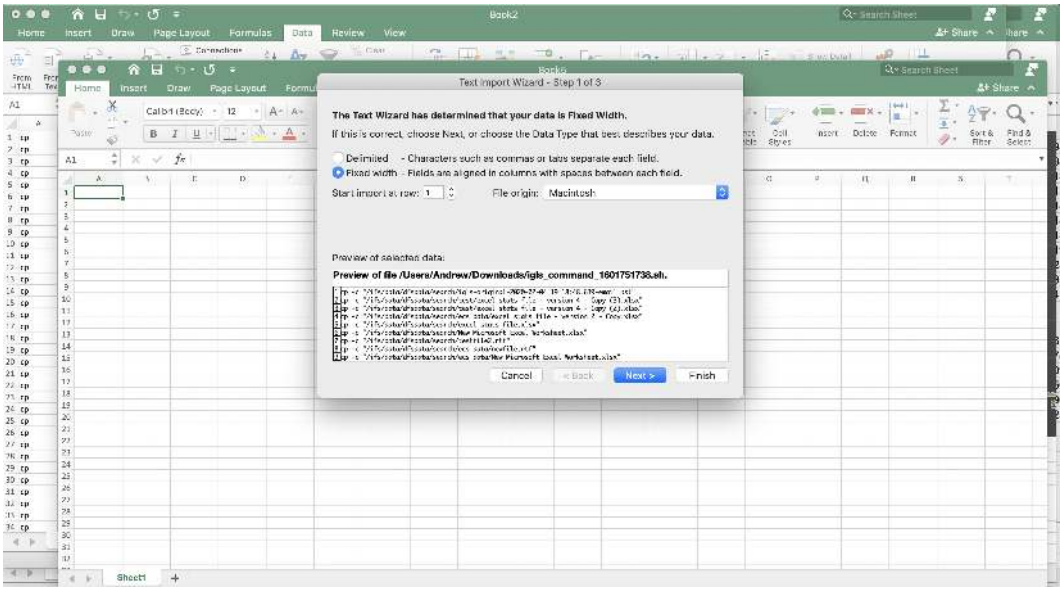

c. Delete the first few header rows of the script.

d. Open a new blank Excel file

   i. Then copy the column with the source files column and past into the new blank Excel file.

   ii. Using the Replace feature in Excel, we will select the target path column in the new file, change the Replace to Search Columns, enter the source path to replace and the target path to replace. The target path is the path created earlier where all the cloned files will be copied.

   iii. In this example the source path was /ifs/data/dfsdata/search/ but we will replace this with the target path of /ifs/data/dfsdata/clonedsearch/. Enter the fields and replace all button.

   iv. 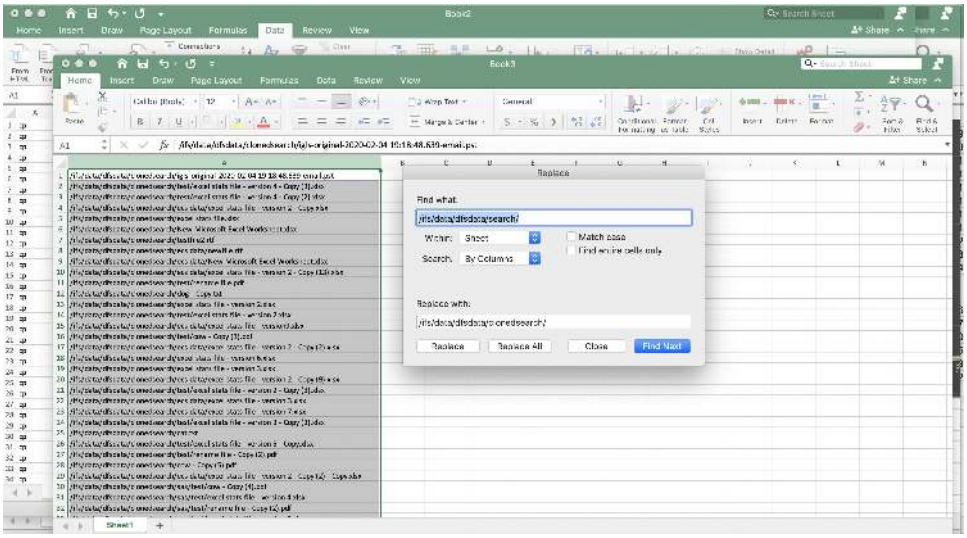

e. Copy the target path column and paste into the main script file

8. Now Save the file as CSV using Save As

9. You will need to save as CSV and then use a text editor to search and replace the triple quote """ with " and replace comma's with a space. This is required to encode files and folders with spaces in " for the script to run correctly.
10. Rename the file to clone.sh (shell script) and scp the script to the Isilon cluster.
11. chmod 777 clone.sh
12. bash clone.sh (this will run the script to clone and this can take a long time depending on the size of the data set). Monitor the target path for files appearing in the folder tree to track progress of the copy script. NOTE: This can run for a long time and is best to leave this script running on terminal within a VM to allow you to disconnect from session and leave the script running.

How to Verify the Cloned Data Consumes no additional space

1. On an original file in the source path run this command. NOTE: Change directory or enter full path to the file.
2. isi get -D mediumfile.zip (source file is 147 MB) **Notice the physical blocks value of 18019**

3. 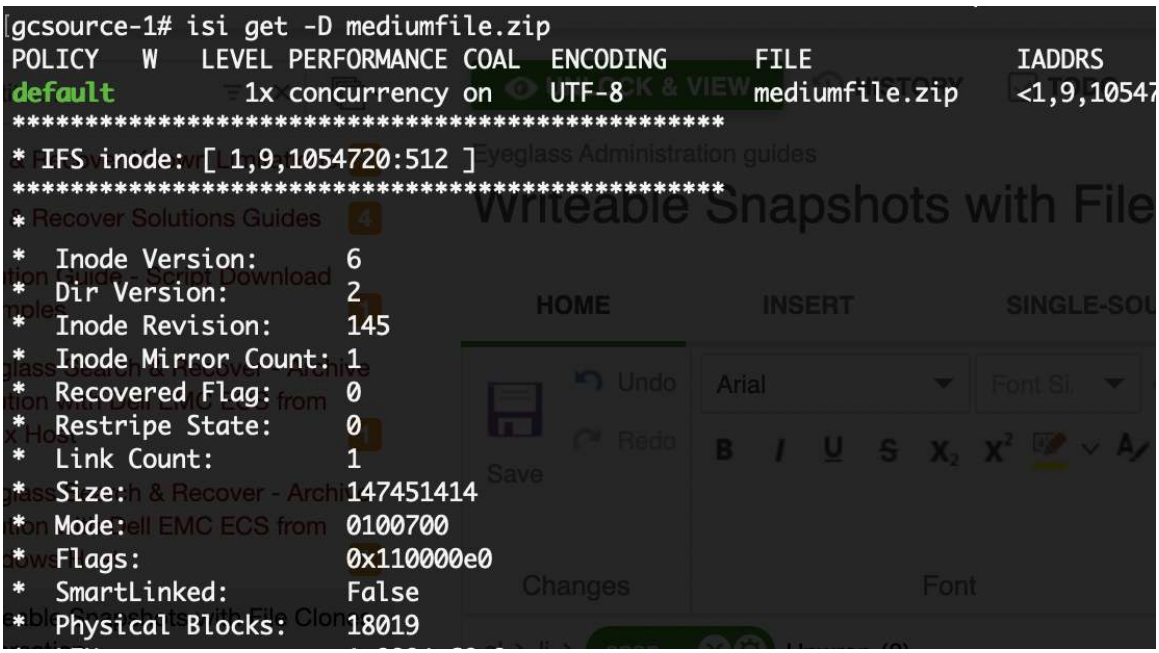

4. Now run the command on the cloned file with isi get -D cmediumfile.zip **Note: the file only consumes 5 blocks versus 18019 blocks.**

5. 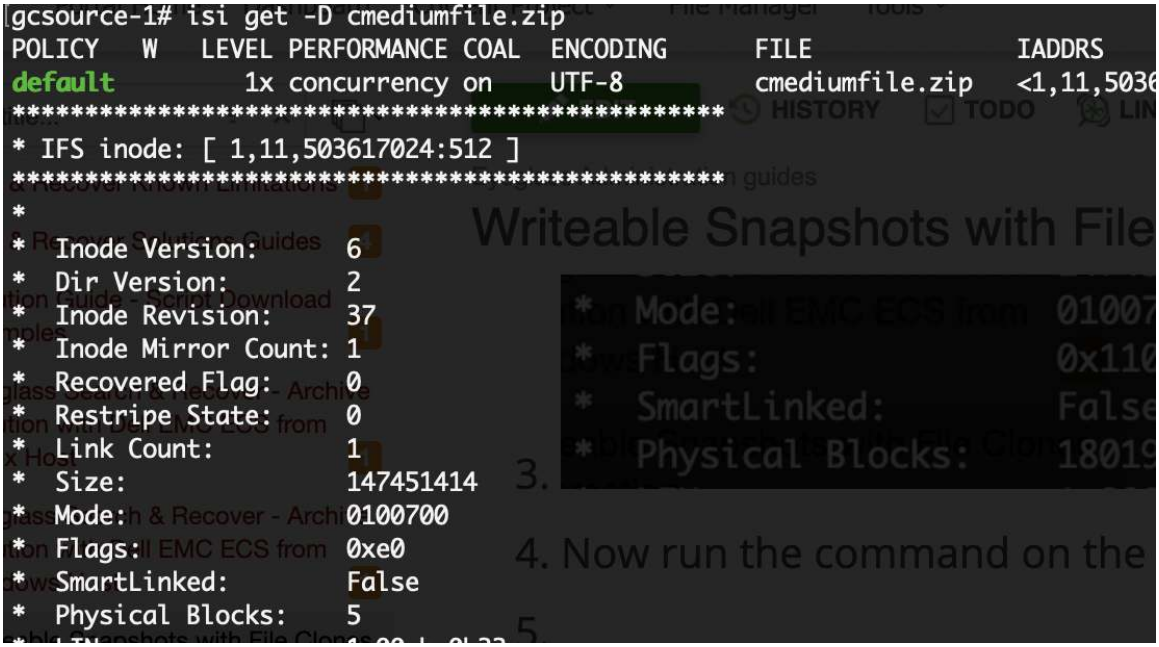

6. done

Browse the Writeable Clone data for testing

1. Create an SMB share to present the cloned data to users or developers for testing. Any writes to the data will consume space equal to the modified blocks changed within each cloned file.

2. 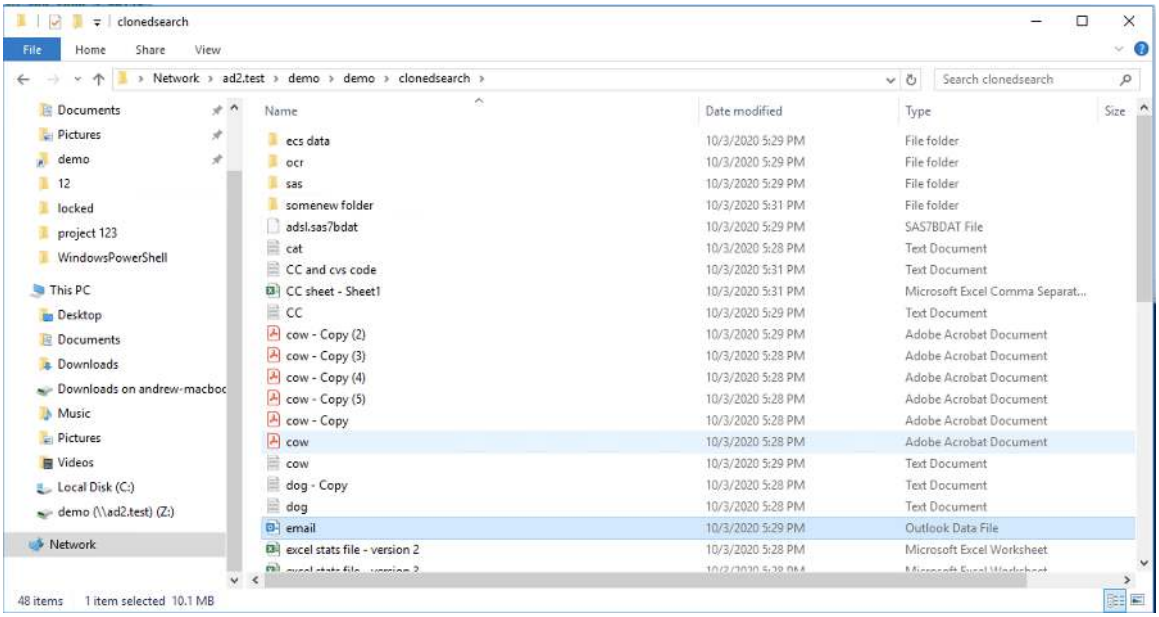

© Superna LLC

9.7.6. Legal Hold Solution

Home Top

- Overview
- Create the Legal hold Project Space
- Create SmartLock policies on the Legal Hold Project Space
- How to Create Legal Search Users in Search & Recover
- How Legal Discovery Users Provide File Lists for Project Retention
- How the administrator copies legal hold files to project space and shares them with Legal Discovery Users
- How to Configure Indexing on the legal hold path

Overview

The legal discovery workflow allows Search & Recover to be used by legal departments to locate file data of interest using content aware search index. Key criteria for this solution is allowing legal to search data without the need to have permissions to the data itself. Once file data of interest is located, it is required to make a copy of this data and comply with legal hold requirements. Legal hold requirements means the file are locked and cannot be modified regardless of the permissions set on the files, and the hold should have a retention period to protect this data for the duration of the legal hold project.

This guide will explain how to configure a legal hold solution with Search & Recover and Dell Isilon Smartlock feature that provides locking and retention policies.

Create the Legal hold Project Space

1. Create a file system location for for smartlock policies to be applied. Example /ifs/legalhold
2. This folder will hold legal hold project data with projects stored under this path example /ifs/legahold/project1 would store all files identified by legal as requiring legal hold for project1

Create SmartLock policies on the Legal Hold Project Space

1. Requirements: Smartlock License
2. This screenshot is only an example of how to configure smartlock retention on the folder. It will lock the file 10 minutes after creation and hold the files for 30 days.
    a. Adjust these settings based on legal hold requirements. NOTE: retention policies can also be created on a per project folder under the /ifs/legalhold project space, if different retention periods are required per project

3. 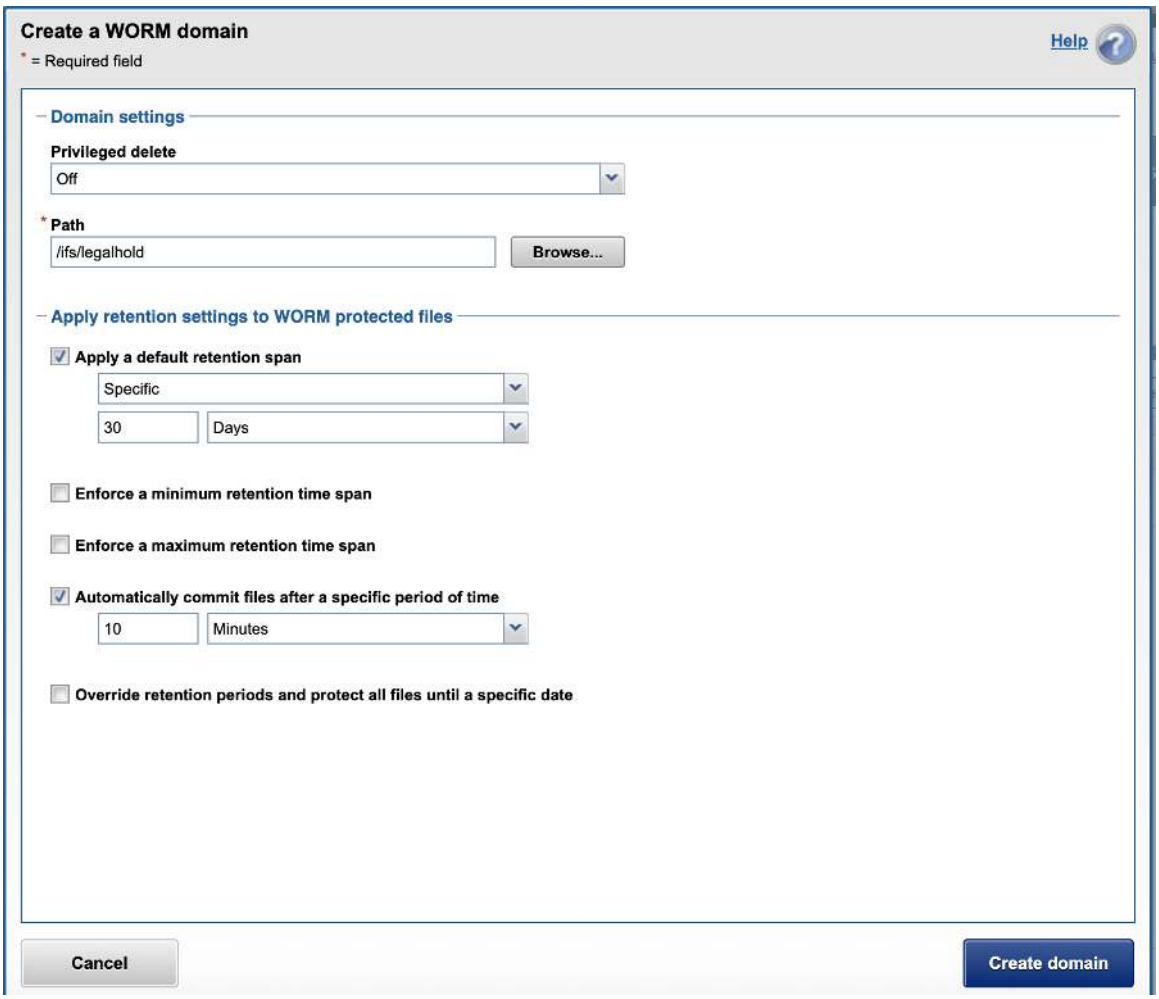

How to Create Legal Search Users in Search & Recover

1. Creating a search administrator with Active Directory users for each legal user that needs to search for files using the following command for each user.
2. Follow the steps on this link to add a path to an AD user. This grants the user permissions to return search results from this path and below. The default command provides access to

search contents of files. This command will override SMB or file system ACL security and return results from a search.

How Legal Discovery Users Provide File Lists for Project Retention

1. Users complete searches based on key words to locate legal discovery documents. The user guide provides provides tips on how to search.
2. The legal discovery user can export a list of files to be copied to the legal hold project folder provided by the administrator. Example path for a project /ifs/legalhold/project1 will be used in this example.
3. The user would click the command builder icon and paste **rsync -axuvR --progress** into box 1 and paste **/ifs/legalhold/project1** into box 2 and then click Create all button to save the export list. See example below.
4. This would be repeated for each list of files required for legal hold. The list of export files would be emailed to the administrator.

5. 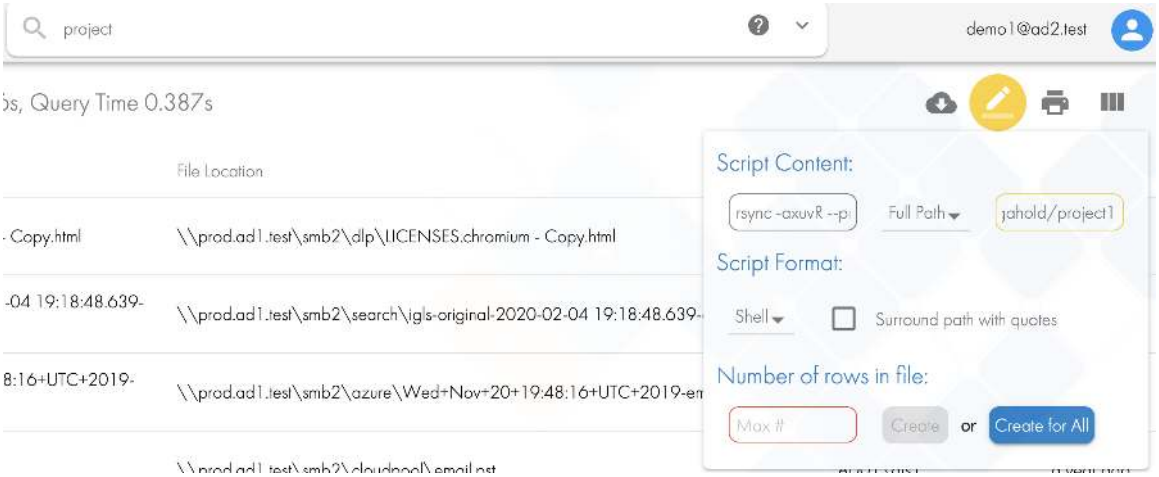

How the administrator copies legal hold files to project space and shares them with Legal Discovery Users

1. The export script files can be copied to the cluster keeping the .sh extension
2. Change permissions to execute with chmod 777 filename.sh
3. Then run the files to copy the files to project1 folder. Note the file location relative to it's original location will be retained.
4. After 10 minutes the files will be locked by Smartlock retention policy set on the folder.
5. An SMB Share can be created on the /ifs/legalhold/project1 path to share with the legal discovery users to open and review documents. The documents will be locked and read-only until the retention period expires.

How to Configure Indexing on the legal hold path

1. Adding the legal hold path to the index will allow the data copied to the legal hold project space will allow full content searches of the locked project data.
2. See the configuration steps to add the folder and start indexing on the folder in the configuration topic in this guide. Guide is here.

© Superna LLC

9.7.7. Smart Lock Data Retention , Lock status, and Expiry Reporting and Delete Solution

Home Top

- Solution Overview
  - Use Case #1 - Compliance Report on Locked Data Storage Requirement for a Future Year
  - Use Case #2 - Locate Expired Retention Data for Scripted Deletion

Solution Overview

Customers using Isilon/Powerscale Wormlock feature face challenges on reporting on the lock status of files and locating files that will expiry in the future based on the lock retention. Once files expire after the retention period, many customers want to locate and delete this data to save space. This solution guide explains how this can be managed within Search & Recover.

Use Case #1 - Compliance Report on Locked Data Storage Requirement for a Future Year

1. Locking data for retention is only half the solution. Reporting on what is locked, how long it's locked, when it was locked are critical requirements to manage retention data.

2. This will be available in release 1.1.5 with a dedicated quick report. To report on data that needs to be stored based on retention requires a search to search file all files on path and sum up the file size and then subtract data that is expired based on the last accessed time stamp of the data within the time period of the report. It will report on future years automatically and show storage requirements after removing expired data from the results.
3. A work around in the current release to report on data that will need to be stored in a future year.
    a. Create an advisory quota on the compliance data path you want to analyze and get the GB's stored on this path.
    b. Using the What's growing Old? Quick Report select the Last accessed for Search by and Group by year. Enter the path where the retention data is stored. See example below.

c. 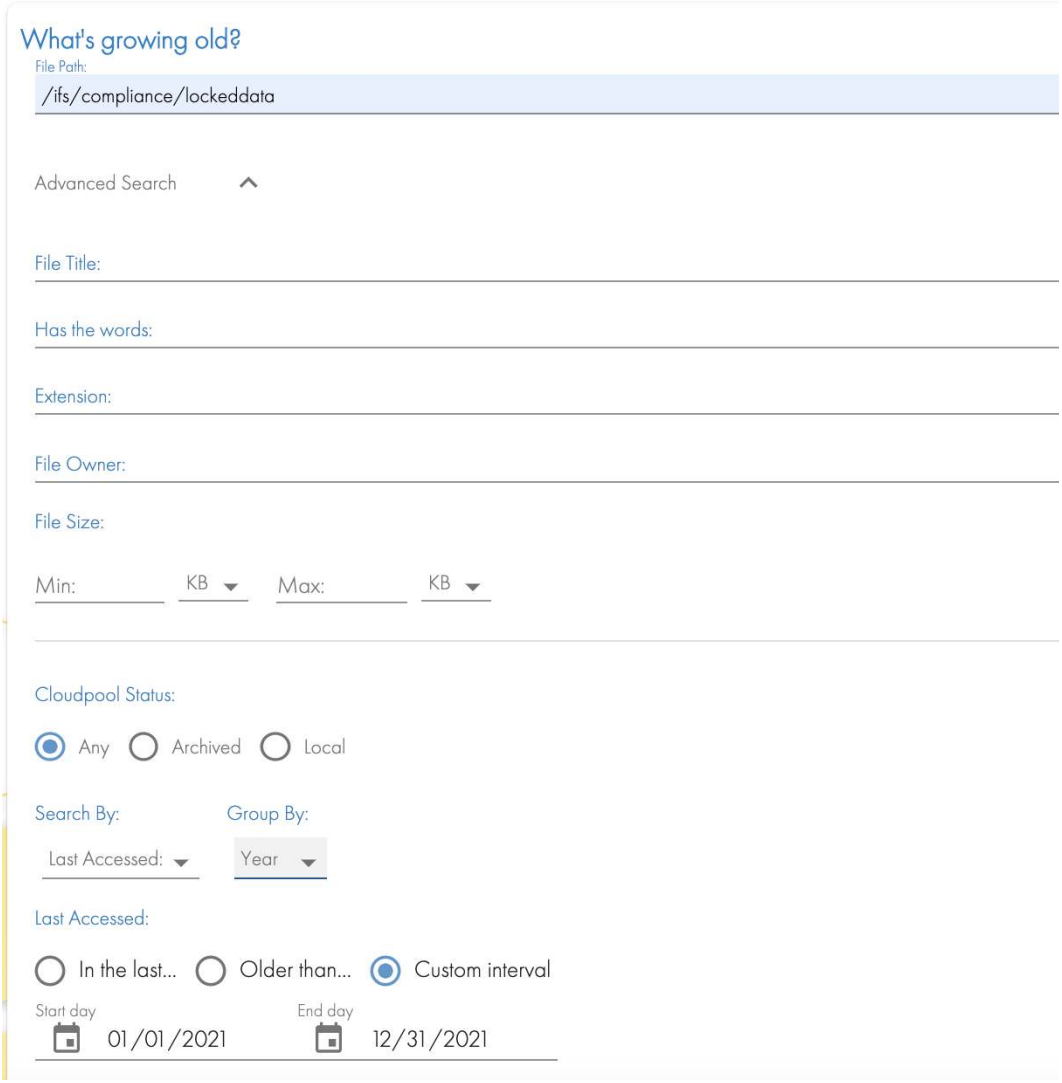

d. The results table below the graph will contain the sum of all the data that will expiry within the year entered in the date range.

e. Now estimate the creation rate of data within your compliance folder that you are analyzing.

   i. Using What's Growing old? Quick report and change search by to Created and group by month. Select a time period to sample the new compliance created for retention. The resulting table will show new created

data by month and can be used to estimate the creation rate of data in future months.

f. To calculate the storage requirement for that year entered above for analysis:
    i. take the quota GB value (this is all data stored)
    ii. add the estimated new data data per month
    iii. Then subtract the Quick report value for expired data.
    iv. The resulting answer is the net new storage requirement for that year. This assumes the expired data has been deleted following the steps in the next section. It is possible your storage requirements go down if more expired data exists versus newly created data.
g. Repeat these steps for future years to build a trend of how the long term storage requirements for compliance data.

Use Case #2 - Locate Expired Retention Data for Scripted Deletion

1. Expired retention data can be deleted to save space. Locating this data and automating the deletion is a key requirement for customers with retention data.

2. Open quick reports tab and use the What's growing old ? Quick report interface.
    a. Enter the path where the retention data is stored
    b. 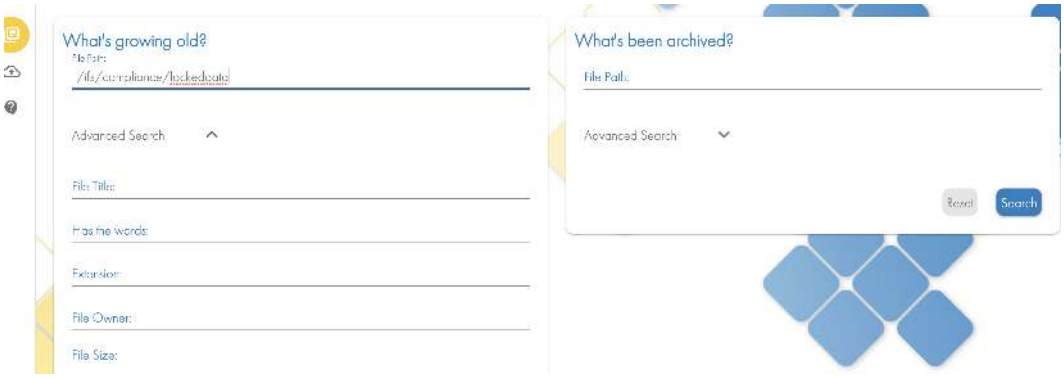

    c. Change the search by to use the last accessed file attribute
    d. Using the "in the last" option select month or year. This will use today's date and search into the past to locate all files with a last accessed time stamp and report on the data using the group by option to sort the data into months or years.
    e. All the data in the results will be data that can be deleted since it matches a date in the past which means the expiry date of the data allows this data to be deleted. See example input for a search using the last month and group by month options.
        i.

1. 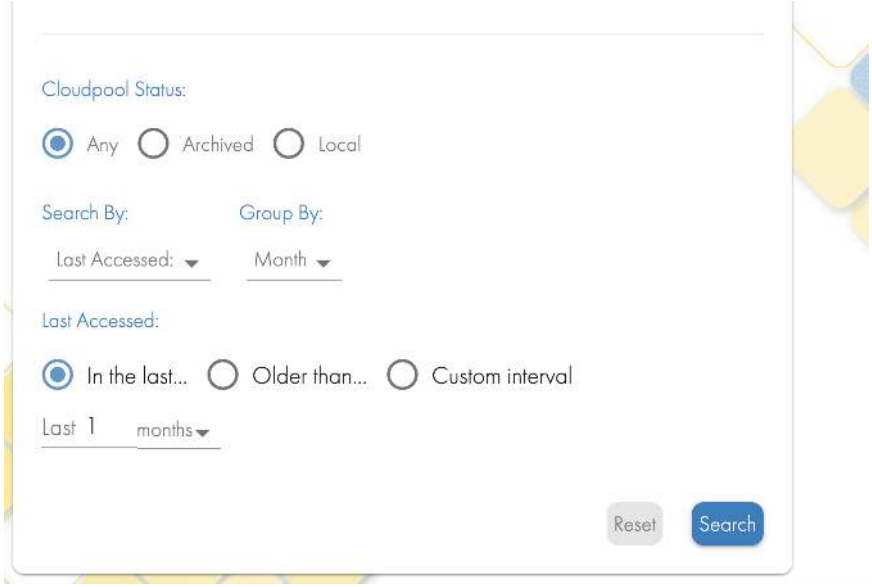

   ii. To report on more expired data use in the last year.

f. Once the report completes you can use the command builder tool to generate a bash script to delete these files.

g. Using the table results below the graph select all table rows and select Show me the files option to list the files that match the search.

h. 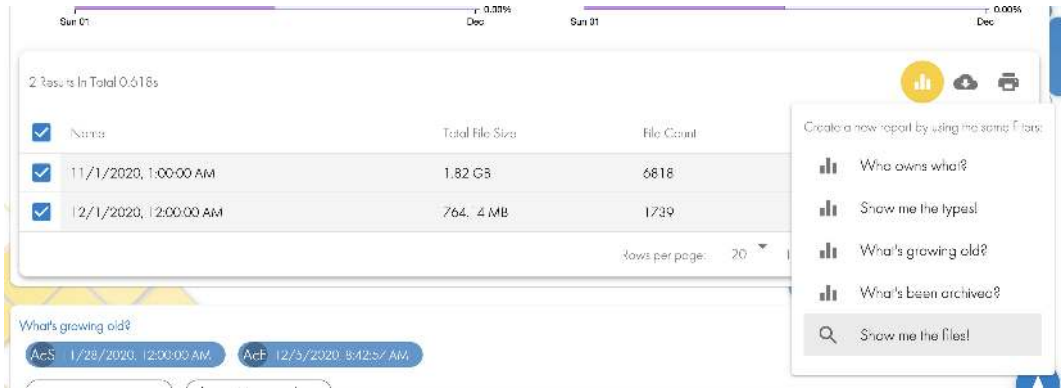

i. Use the screenshot example below to generate a bash script to delete the file and log the delete to a file

j.

k. Save the file to your cluster and change the extension to .sh, then change permissions of the file to allow it to be executed with chmod 777 filename.sh

l. Login as root user on your cluster to execute this script.

m. ./filename.sh

n. Save the deletelog.txt for long term compliance reporting of what expired data has been deleted.

o. Done.

© Superna LLC

9.8. Deployment Diagrams, Firewall Ports , VM Requirements and Supported OneFS Releases

Home Top

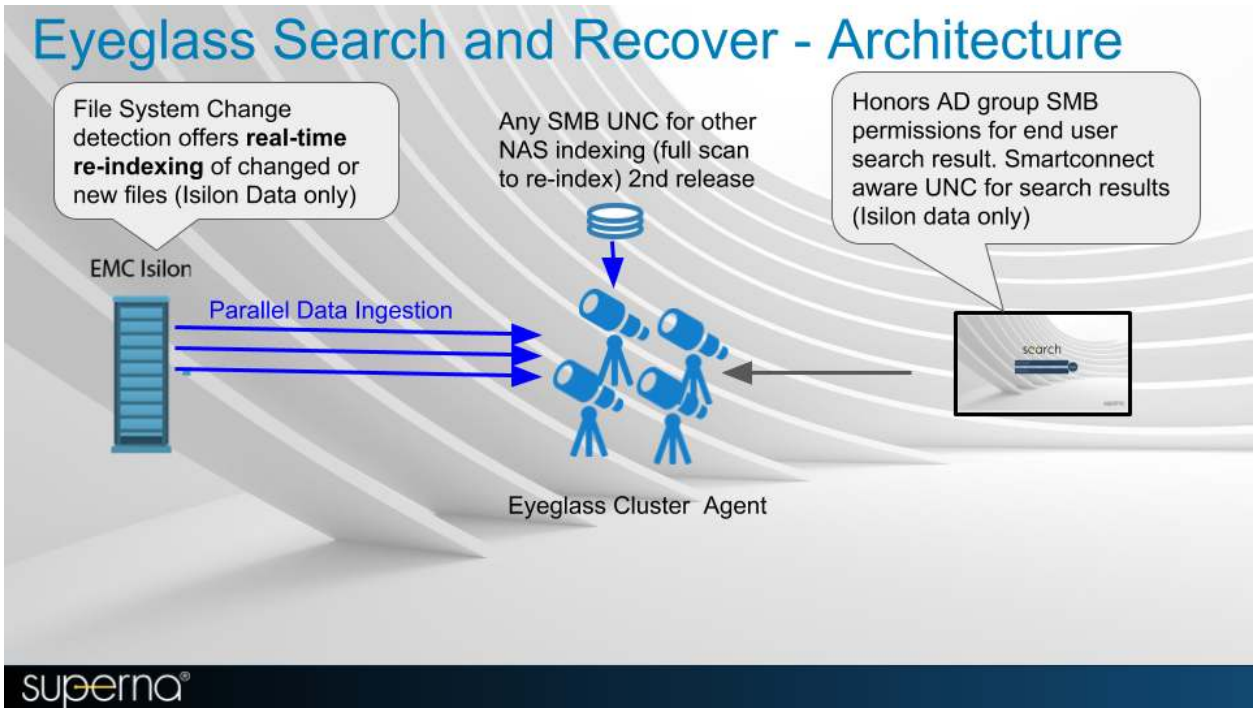

The Eyeglass Search & Recover Cluster

VM specifications:

1. VMware OVA.
2. **Minimum:** 4-6 x 16G RAM , 4 x vCPU, 400 G Hard Disk, 1Gig Ethernet Interface (See sizing section for number required).
3. Back Log file processing buffer and local index are stored locally.
    a. 300 Million for content indexing.
    b. 300 Million for metadata only indexing.
4. **Includes**: index Engine, ingestion engine, GUI, CLI.

5. Supported Browsers: Chrome, IE11.
6. OneFS Supported versions: 8.0.0.x, 8.0.1.x, 8.1.x.x, 8.2.x.x.
    a. 1.1.5 release to support One9.0, 9.1, 9.2

Firewall Ports

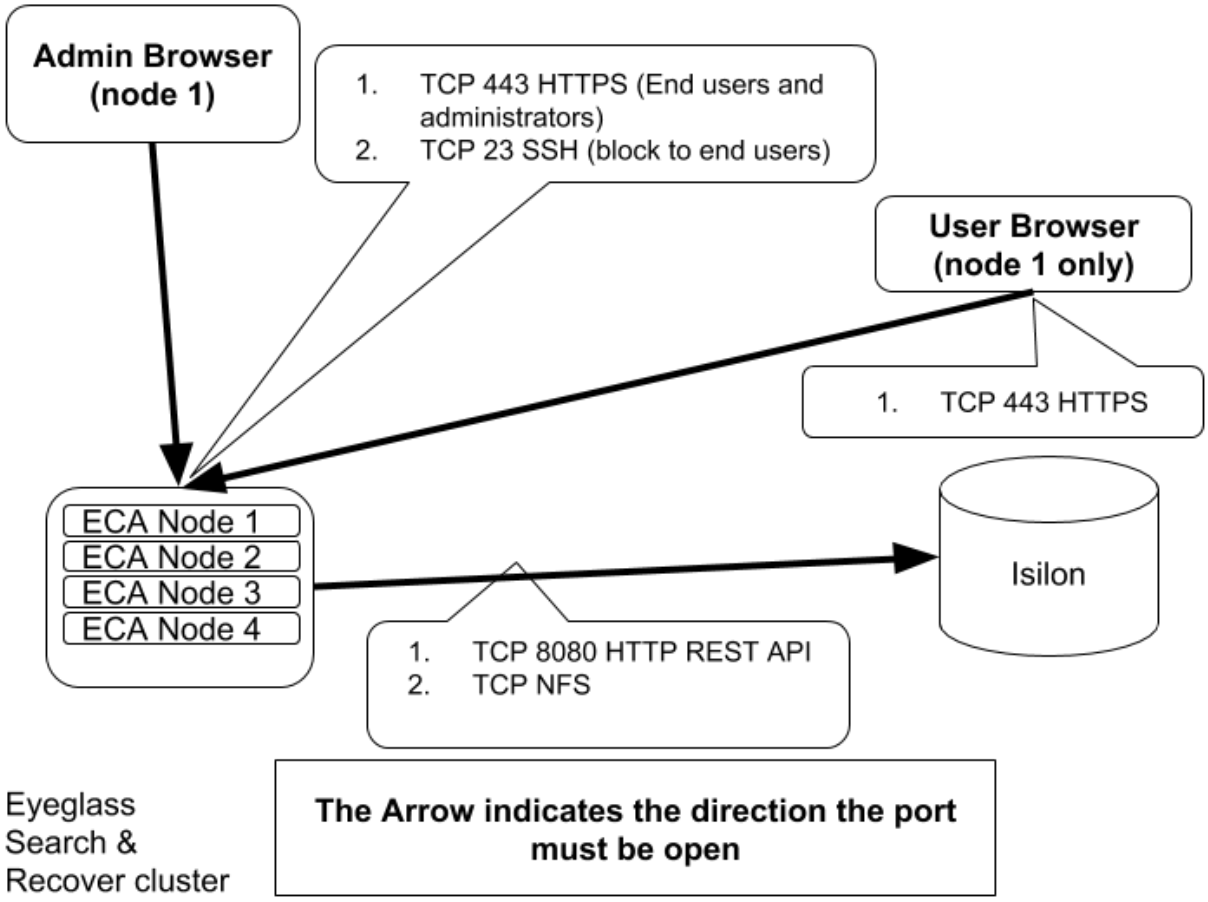

© Superna LLC

9.9. Understanding Search Security Results and User Login sessions

Home Top

- Overview
- 5 security modes to secure search results:
- User login Session Control

Overview

This section explains how the various modes of search security are managed, and how user logon sessions are controlled. These is an important considerations before providing users access to search for indexed content.

5 security modes to secure search results:

1. **SMB Share Security mode:**
    a. A user will login with AD User id and password that is verified by the PowerScale against Active Directory.
    b. Search & Recover cluster will retrieve the users AD groups and calculate which SMB shares the user has read or write access permissions.

c. All searches for the logged in user will **only** return results for data at, or below the path defined by the SMB shares that they have PowerScale access.

2. **File Owner Security Mode:**
    a. In this mode results are returned **only** if the user is listed as owner of the files, and share level permissions will **Not** be used to filter results returned to the user.
    b. **Use Case:** Home directories typically have 1 SMB share and many directories for each user under the SMB share path. A user owns the folder and all data within their home directory. Using File Owner Mode will ensure only data created or owned by the user will be returned in search results.
3. **Hybrid Security Mode:**
    a. Enabling SMB and File Owner mode at the same time will filter data using both modes. This means group share searches for User X will only return data to User X for files they own in the group share, and will NOT return data for files owned by other users in the group share.
4. **SHARE ACL Mode:**
    a. Requirements
        i. Requires 1.1.2 or later
        ii. Release 1.1.5 removes the need for root user proxy to read ACL's and uses backup and restore permission on the service account. Covered in the minimum permissions guide.

b. **Limitations**

i. This feature does not behave like a file server that dynamically evaluate all folder ACL's in a file system tree. This feature only statically scans folders at one level in the file system to cache the results.

ii. Does not support nested folder ACL evaluation. A single level in the file system below a share will be scanned , that depth is controlled with a global varible.

c. **Use Case**

i. **Group or project share space** - This assumes a share for groups or projects is used and all the folders under this share have ACL's applied to determine which folder users see based on AD group ACL's applied.

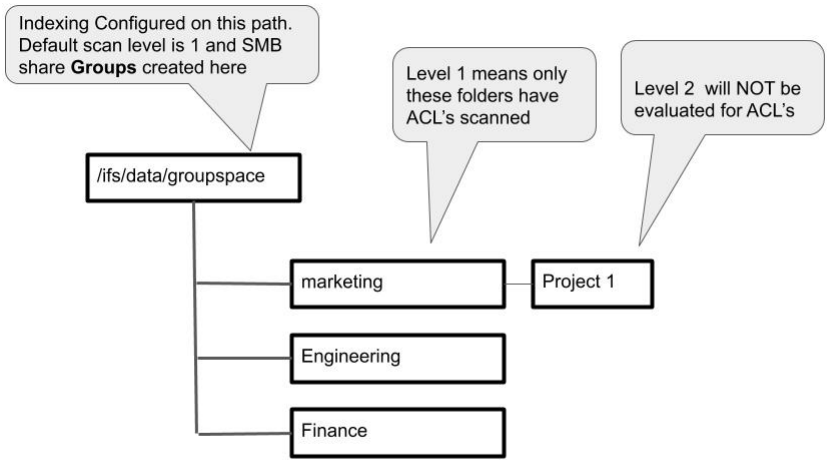

ii.

iii. Example: For each of the shares under an indexed folder path, a list of folders that are nested under each SMB share path that is scanned to a folder depth as controlled by the ECA_AUTH_ACL_DEPTH levels deep cluster wide setting, will be scanned and

AD groups and users will be evaluated to build a security view.

iv. **Example with 1 Share**:

1. If the share is on /ifs/groupshares and ECA_AUTH_ACL_DEPTH has a value of 1 (the default), will enumerate the directory children of /ifs/groupshares, for example [/ifs/groupshares/groupA, /ifs/groupshares/groupB]. This is repeated for each folder under the SMB share path.

v. **Example with multiple shares**:

1. Indexed folder is /ifs/data/projects , and 3 shares exist below SMB share projectA /ifs/data/projects/projectA, SMB share projectB /ifs/data/projects/projectB and projectC /ifs/data/projects/projectC and default index depth is 1.

    a. Each of these shares will be scanned 1 folder depth below the share path and ACL's collected to evaluate which folders users can see.

d. **Key Information About ACL Mode**

i. This mode will evaluate share permissions first and then look at ACL user and group security applied to the file system under each SMB share path that falls under a path configured for indexing.

ii. The depth in the file system below share path is controlled globally and defaults to level 1 below the indexed path will be scanned for ACLs. The ACL's are used to evaluate which paths a user is allowed to see in search results. The ACL's are evaluated and cached every 1 hour by default to detect changes to the file system ACL's.

iii. The user must have at least read permissions to a share and to a folder ACL to get results returned.

iv. Deny permissions are also factored into the users search results.

v. See the table below that shows expected search results depending on the ACL's applied to the SMB share, the ACL using user and group logic.

vi.

ID	SMB Access	Group ACL	User ACL	Expected Results	Test Results
1	Deny	<any>	<any>	User cannot login, or, if user has access through another share, user cannot see any search results on the ACL's share path	User can login, no results are returned when searching.
2	Allow	Deny	Deny	User can login but cannot see any search results on the	Login is successful, no search results are returned.

				share path	
3	Allow	Deny	Allow	User can login but cannot see any search results on the share path	Login is successful, no search results are returned.
4	Allow	Allow	Deny	User can login but cannot see any search results on the share path	Login is successful, no search results are returned. Login with different group user is successful, can perform search.
5	Allow	Allow	Allow	User can login and see results	All users from group can login and perform search.
6	Allow	None	None	User can login and see results	Login and search is successful with all group users.
7	Allow	None	Allow	User can login and see results	Login and search is successful.
8	Allow	None	Deny	User can login but cannot see any search results on the share path	Login is successful, no search results are returned. Login with any other user from group, no search results returned.
9	Allow	Allow	None	User can login and see results	Login and search is successful with all group users.

10	Allow	Deny	None	User can login but cannot see any search results on the share path	Login is successful with all users added on group, no search results are returned.

5. **Admin only Mode**:
    a. This mode is global and only allows users listed on the admin list to login. See Configuration section on how to configure.

The first 4 security modes above can be configured per indexed folder.

User login Session Control

Users that login to the UI will have the AD userid and password validated, and a web token session sent to the browser. This session token will survive browser restarts. It is a signed authentication token with an expiry set by the Search cluster. The default is 12 hours. This can be changed on the cluster. **NOTE: Any changes made to the cluster configuration will not be visible to end users until they logout and login again.**

1. ssh to the Search master node (node 1).
2. Edit the config file to change the value
    a. vim /opt/superna/eca/eca-env-common.conf
3. add the line below to set expiry time in minutes (x is minutes - replace with minutes)

    a. export SEARCHMW_SESSION_EXPIRATION_MINS=x
4. shut down the Search cluster
    a. ecactl cluster down
5. Then bring up the cluster
    a. ecactl cluster up

© Superna LLC

9.10. Search & Recover Cluster Configuration Steps

Home Top

- Quick Start Steps
- Search CLI Basics
- Adding and viewing License keys
  - Licensing CLI Commands
- Adding, Viewing Clusters
  - How to Change a PowerScale cluster's IP address , enable Snapshot Recovery root password or change service account user name
- Adding, Viewing, and Starting Full Indexing Jobs Section
  - Content Ingestion Overview
  - List of supported file types for full content ingestion and indexing
- How to Enable Security Mode for Search Results on Indexed Data
  - SHARE_ACL mode configuration:
- How to add a folder path to be Indexed
  - (Optional Advanced Configuration) Folder Ingestion processing of include, exclude patterns and metadata or full content overrides
    - Order of processing
  - How to Configure Common Use Cases to include or exclude a path or file type Best Practice

    - (Optional Advanced Configuration) Overview and Examples of include and exclude
- How to Index S3 Storage Buckets
- How to Index Snapshots for User Self Serve File Recovery
    - Requirements
    - Limitations:
    - Understanding Snapshot Index Modes
    - Use Cases:
    - How to Configure - Snapshot Monitor Mode Integrated with Folder Indexing
    - How to Configure - Snapshot Only Mode
        - Requirements:
        - Overview:
        - How to Configure Snapshot Only Mode:
- How to Enable File Global Recovery Modes to control snapshot Restore file Collisions
    - Requirements:
    - Configuration:
- How to start a full index or incremental job on a folder path
- How to Manage Scheduled Jobs (Global and Folder full and incremental)
    - Requirements:
    - Schedule Job Definitions

    - Commands to Manage Schedules (enable, disable, set schedule)
    - Example Full and incremental index job Schedule Configuration
- How to Monitor Index Job Status
- How to Monitor Ingestion with the stats command
- Running Inventory Scans and Viewing users and SMB Share Access
- How to Enable User Authentication to Data within Access Zones and Return Search Results with Smartconnect UNC's to Files
    - CLI commands to add Zone FQDN to Authentication and Search Results Display
- How to Enable Administrator Search Security Override
    - Compliance, File System Antlaytics Administrators
    - How to configure a List of Search administrators
- How to add Data Owner Search Administrators
    - Data Owner Admin Configuration
- How to Configure admin only login mode and block user login

Quick Start Steps

This quick setup guide provides exact steps to get up and running, and a link to learn more if needed.

Note:

- **All searchctl commands must be run as the ecaadmin user from Search node 1**

1. **License Keys:**
    a. Copy license zip file to Search node /home/ecaadmin directory and change permissions chmod 777.
    b. searchctl licenses add --path /home/ecaadmin/**<name of zip>**.zip.
    c. Learn more about - license key CLI commands .
2. **Add a cluster to Inventory:**
    a. searchctl isilons add --host **x.x.x.x** --user root **(x.x.x.x Subnet Service IP of the system zone DO NOT USE FQDN. PowerScale does not support Session cluster wide session authentication. Use root user for quick setup (Create eyeglass service account user for production use).**
    b. Learn more about - Add a cluster CLI commands.
3. **Add a folder to be indexed:**
    a. searchctl folders add --isilon **<name of cluster>** --folder /ifs/data/testsearch **(NOTE: The name of the cluster is returned from set #2, record the folder ID returned from this command, default add command is metadata only indexing).**
    b. Learn More About - Adding folders and index option CLI commands.
4. **Start Index Job for a folder that has been added:**
    a. searchctl folders index --id <folder ID> **(NOTE: Replace <folder ID> with ID from step #3. Example: only 3fe1c53bdaa2eedd).**

b. Learn more About - Starting index jobs for folders CLI commands.

5. **(Optional step if indexed data is not present in the System Access zone) Add SmartConnect UNC to be added to user search results for Access Zones with Indexed Data:**

   a. This feature allows users to see a UNC path with SMB shares inserted into the results to enable simple click to copy and open end ensures results in none system zone are displayed to users.

   b. searchctl settings zoneunc add --isilon prod-cluster-8 --zone **system** **--fqdn** <smartconnect name> **(NOTE: Repeat for each Access zone that has indexed data, SmartConnect name should be from each Access Zone added).**

   c. **Learn more about - Managing Search Dynamic UNC path CLI commands.**

6. **Monitor Indexing Job Progress**

   a. This command will show progress of files as they are indexed, with real-time updates every few seconds.

   b. searchctl folders stats --id <folder ID> **(NOTE: replace <folder ID> with ID from step #3, example only 3fe1c53bdaa2eedd).**

   c. Lean more about - Monitoring Index Job with Stats CLI commands.

   d. Learn more about - Index Update Intervals.

   e. **WARNING: Incremental indexing for changed files runs every hour after a folder is added, and runs on the hour.**

Full index scan starts after the index job is started, but files are "committed" to the index every 30 minutes, which means files will NOT be returned until 30 minutes after starting the index job. See below.

f. For testing use the commit command to force files that are indexed to be visible in the search results.

   i. searchctl solr commit.

7. Start Searching

   a. Open a browser to https://<ip of node 1>.

   b. Login as a user with **SMB share permissions** to the folder path added for indexing (Userid syntax user@domain.com) .

   c. **Type enter in the search bar this will return all files indexed so far.**

   d. Refer to the user search guide or administrator search guide and Advanced Searching Guide.

   e. NOTE You will need to monitor the index job stats to see if any files have been indexed BEFORE trying to search for files.

Search CLI Basics

The Search & Recover CLI uses ecactl CLI syntax, or a new shortcut to search only commands called searchctl. This limits what needs to be typed.

1. searchctl supports -h for help on a command.

2. seachctl supports -v for verbose and easier to read output for some commands.
3. searchctl supports **tab** completion.
4. searchctl supports **tab tab** to show list of available commands.
5. ecactl does not support tab or tab tab features.

Adding and viewing License keys

1. Licensing is per PowerScale node or per PowerScale cluster. The license allows a cluster to be added to the configuration for indexing.
    a. When a cluster is added to the configuration for indexing the node count is detected and reduced from node count licenses, when no more node licenses are available you will be able to add clusters to the configuration.
    b. If cluster based licensing is used, each cluster added to the configuration will reduce the count of clusters from the license count.
    c. **NOTE: License keys are locked to the cluster GUID when the license key is installed and a cluster is added . License keys cannot be moved to another cluster. License keys cannot be reset if the wrong cluster was added.**
    d. **NOTE: Unlicensed clusters will not be indexed, and a license error will display in the UI to end users.**
    e. **NOTE: Search results for unlicensed clusters in the index will not be returned in the results list.**

Licensing CLI Commands

1. **searchctl licenses add** (Uploads a new license zip file)
    a. searchctl licenses add --path <full path to zip file>.
    b. NOTE: license zip file permissions chmod 777 filename.zip.
2. **searchctl licenses uninstall** (Removes all of the licenses on the system).
3. **searchctl licenses list** (Lists the currently installed licenses).
4. searchtl licenses applications list (list all applications on a unified deployment with Golden Copy).

Adding, Viewing Clusters

1. To Add a cluster:
    a. NEW feature supports automatic load balancing and HA features for maximum performance to index the file system. This operates like SmartConnect but 100% supports session authentication and CRSF secured clusters.
        i. **Select an IP address in the system zone IP pool but do NOT use the SSIP.** This will turn on Load Balance mode and will:
            1. inventory all the IP's in the pool.

2. Use round robin API calls to each node to increase performance.
3. Supports failover if a node fails.

ii. **NOTE: Not recommended. but adding via SSIP will disable load balance mode and will only send api calls to the SSIP node.**

b. searchctl isilons add --host x.x.x.x (pool ip) --user yyy [--applications APPLICATIONS]

c. [--applications APPLICATIONS] is used on the unified deployment with Golden Copy. Use GC for Golden Copy cluster and SR for a Search & Recover cluster

d. **NOTE: xxxx is the a pool IP address of a management in the system zone.**

e. NOTE: yyy is the local user created on the cluster. This service account can be created by following minimum permissions documented in this guide.

2. To list clusters and license status:

   a. searchctl isilons list

3. To remove a cluster:

   a. **NOTE: Do NOT remove a cluster and try to add a different cluster, this will be blocked. Licenses are bound to the cluster when its added. Support will be unable to assist. Sales will be required to assist with the purchase of a new license for a 2nd cluster. You will be able to add the same cluster back to the configuration.**

b. **NOTE: The snapshots created to ingest content are not deleted and must be manually deleted from the cluster**

c. searchctl isilons remove --name (use searchctl isilons list to get the exact name of the cluster)

How to Change a PowerScale cluster's IP address , enable Snapshot Recovery root password or change service account user name

1. List the current ip address:
   a. searchctl isilons list.
2. How to change ip address for an PowerScale cluster in inventory:
   a. searchctl isilons modify --name <PowerScale_Name> --ip x.x.x.x --user --update-password (Get the PowerScale name from the list command, x.x.x.x is the ip address used to add the cluster, --user is the service account name normally eyeglassadminSR.
   b. Example to add password for root user and snapshot monitor feature:
      i. searchctl isilons modify --name xxxx --ip y.y.y.y {--root-pw} (hit enter and a prompt will ask for password, xxxx is the name of the cluster displayed from the list command and y.y.y.y is new ip address to use for REST API calls to the cluster).

ii. The --root-pw requires the root user password to be used with recovery portal with snapshot monitor mode. Only enter root user password if you plan to enable snapshot mode feature.

c. Example to change the user and password to connect to the cluster:

i. searchctl isilons modify --name SC-8120A --ip 172.25.27.32 --user (you will be prompted to enter the user name and then the password).

d. Example to change the service account password only:

i. searchctl isilons modify --name SC-8120A --ip 172.25.27.32 --update-password (you will be prompted to enter the new password)>

Adding, Viewing, and Starting Full Indexing Jobs Section

Content Ingestion Overview

Full content ingestion uses a 2 stage approach when processing a path configured for ingestion. The means files are added to a queue for content ingestion are first entered for metadata ingestion first and second stage for full content. This approach allows search results to appear to users based on path, name and extension quickly while allowing content ingestion to be processed by the full content ingestion queue.

A parallel process monitors incremental ingestion by detecting changes in the file system for any configured paths. The changed files are processed in a separate queue from full content ingestion processing and will process metadata and content ingestion at the same time. This allows file content that is changed to appear with full content, and allows users to find active content in the file system. These two queues operate until the full ingestion is completed and incremental ingestion runs continuously.

The objective of this solution is to allow active content to appear in the results faster than stale content.

List of supported file types for full content ingestion and indexing

This is the list of file types that support content ingestion in this release.

https://tika.apache.org/1.18/formats.html#Full_list_of_Supported_Formats

How to Enable Security Mode for Search Results on Indexed Data

This section is important to understand before adding folders to be indexed. If multiple modes are used on different folders the security of the results is processed for each folder and all results are returned from all indexed folders.

1. The flag --auth-type {SHARE_ACCESS,FILE_OWNER,SHARE_OWNER, SHARE_ACL} is used when adding a folder to be indexed, this flag is used when adding a folder path to be indexed. **NOTE: If the flag is not used the default is share access mode**
    a. **Share Access mode** - Means a users SMB share paths are used to restrict results to data that is at or below SMB share paths they have access to mount, AND the data is indexed with Share_Access mode at or below the share paths they have permissions.
    b. **File Owner mode** - Will only return results to the user on the folder or below if the user owns the file in the file system. **NOTE: Use this mode for the home directory folder.**
    c. **Share Owner Mode** - This mode combines share and file ownership filters on results. This should be used on group share paths if a group share is secured using ACL's in the file system, versus share level permissions. Combining the security mode on a path means the user MUST be a member of the share to see the results, AND must be owner on the file in the results.

2. SHARE_ACL mode configuration:
    a. This mode is designed for group share space with a share and ACL's applied to folders directly below the share.
        i. See detailed explanation of this security mode here.

ii. NOTE: Release 1.1.2 or later is required for this security mode

iii. NOTE: Release 1.1.5 removes the requirement for proxy root user to read ACL's and uses service account backup and restore role permissions on the cluster.

iv. This feature allows ACL's in the file system to determine if a user should see results from a given indexed path.

1. This feature will not evaluate all folder ACL's under a folder due to performance reasons.
2. A cluster wide setting controls how many sub folders below each SMB share path will be scanned for ACL's to build the user filters on search results.

b. **SHARE_ACL mode** defaults directory depth to 1 and means only 1 folder below the indexed folder will be scanned for ACL permissions to determine the users access. The cluster configuration to change this requires the following steps:

i. vim /opt/superna/eca/eca-env-defaults.conf

ii. find the export ECA_AUTH_ACL_DEPTH=1 (change to a depth value up to 10)

iii. save the file

iv. Then restart the cluster

v. ecactl cluster down

vi. followed by below to ensure the change takes effect

vii. ecactl cluster up

How to add a folder path to be Indexed

NOTE: Default mode is metadata only indexing. See below for an example of how to enable full content .

1. To add a folder to be indexed with **metadata ONLY:**
   a. **searchctl folders add --isilon <name of PowerScale> --folder /ifs/something**
   b. [--metadata-only] [--includes INCLUDES]

      [--excludes EXCLUDES]

      [--metaIncludes META_INCLUDES]

      [--fullIncludes FULL_INCLUDES]

      [--snapshotMode TYPE]

      [--auth-type {SHARE_ACCESS,FILE_OWNER,SHARE_OWNER,SHARE_ACL}]
2. To add a folder with **full content AND metadata:**
   a. **searchctl folders add --isilon <name of PowerScale> --folder /ifs/something --metadata-only false.**
   b. **NOTE: Name of PowerScale is the PowerScale cluster name of cluster added to Search & Recover.**
3. To list folders that are indexed (returns folder id used for other commands):

a. **searchctl folders list.**

b. **searchctl folder list --verbose (provides more details on the configuration of the folder configuration).**

4. To remove an indexed folder:

   a. **searchctl folders remove --id ID** (get the folder id with searchctl folders list).

   b. **NOTE: The snapshots created to ingest content are not deleted and must be manually deleted from the cluster.**

5. **(Advanced Option)** To modify an indexed folder, and change includes or excluded file types:

   a. **NOTE: Modify command will require all settings needed and will replace previous settings with the new settings. If adding extensions or paths for content indexing, all required paths or extensions need to be added when modifying a folder configuration.**

   b. searchctl folders modify --id ID (add new flag values below to update the folders settings).

   c. [--metadata-only {true or false} [--includes INCLUDES]

      [--excludes EXCLUDES]

      [--metaIncludes META_INCLUDES]

      [--fullIncludes FULL_INCLUDES]

      [--snapshotMode TYPE]

      [--auth-type {SHARE_ACCESS,FILE_OWNER,SHARE_OWNER,SHARE_ACL}]

(Optional Advanced Configuration) Folder Ingestion processing of include, exclude patterns and metadata or full content overrides

Content ingestion configuration allows for includes and excludes to overide default ingestion rules which will ingest all file types all paths under the configured path. In addition, a folder configured for metadata can only have an override to full content index paths, or even specific file types. The reverse is also supported on a full content ingestion folder to apply an override to metadata index certain paths or specific files. Uses cases below explain the use cases.

Order of processing

1. Includes patterns are processed first.
2. Then excludes patterns are processed 2nd.
3. Then folder override for metadata or full content is processed 3rd.

How to Configure Common Use Cases to include or exclude a path or file type Best Practice

1. **Home Directory or Group share space:**
    a. **Best Practice:** Index the home directory for metadata only and include the file types you want to index as full content. This reduces the index size to focus on high value content only.

b. **How to Configure Content Indexing by file extension:**

   i. This command will index contents of files matching the above extensions in the home directory.

      1. **searchctl folders modify --id <ID> --fullIncludes="*.ppt,*.docx,*.xls,*.pdf"**

   ii. This command will exclude all the roaming profile registry data in the home directory and full content index files by extension. This will also reduce low value content to be indexed and searchable.

      1. **searchctl folders modify --id <ID> --exclude "**/AppData/**" --fullIncludes="*.ppt,*.docx,*.xls,*.pdf"**

2. **Full Content Folder added with file types that cannot be indexed:**

   a. **Best Practice:** A directory path with a lot of image formats and some content types that can be indexed, should be optimized to avoid processing file types that do not have content to index.

   b. **Note: The folder was added for full content indexing**

   c. **How to Configure:**

      i. **searchctl folders modify --id <ID> --metaIncludes="*.png,*.jpeg,*.tiff"** (NOTE: To add new extensions, you must apply all previous and new to modify the folder)

      ii. **This command will skip an attempt to process these file types for content ingestion, and only process them for metadata on this ingestion folder.**

(Optional Advanced Configuration) Overview and Examples of include and exclude

For the searchctl folders [add|modify] commands, add new arguments:

A glob is a pattern match syntax to match files or folders using examples shown below.

Flag	Description
--include	File paths matching this glob will be included in the indexing operation. If not specified, all files will be included.
--exclude	File paths matching this glob will be excluded from indexation. This flag only applies to those files that are included by the --include flag. If not specified, no files will be excluded.
--metaIncludes	File paths matching this glob will be indexed with metadata only. This argument only applies to files that are included by the --include and --exclude flags. It will have no effect if applied to folders that have the --metadata-only flag set to true.
--fullIncludes	Only file paths matching this glob will be full content indexed. This argument only applies to files that are included by the --include and --excludeflags. It will have no effect if applied to folders that have the --metadata-only flag set to false.
--auth-type {SHARE_ACCESS,FILE_OWNER,SHARE_OWNER}]	Default security is Share level access results filtering

Examples:

Exclude everything in the user's appdata profile:

--exclude '/ifs/home/*/AppData/**'

Only index docx and pdf files, and exclude everything in a tmp directory:

--include '*.pdf,*.docx' --exclude '/ifs/data/home/tmp/**'

Only index docx, pdf and bmp files, and but treat bmp files as metadata only.
--include '*.pdf,*.docx,*.bmp' --metaIncludes '*.bmp'

Index all files except those in AppData, but only do full content for pdf and docx
--exclude ''/ifs/home/*/AppData/**' --fullIncludes '*.pdf,*.docx'

Index all files with full content, except for those with a .png suffix which should be metadata only:
--metadata-only=false --metaIncludes="*.png"

Index all files as metadata only, except for docx, which should be included for full content.
--metadata-only=true --fullIncludes="*.docx"

How to Index S3 Storage Buckets

This 1.1.5 feature allows S3 storage to be added as a target for indexing object names. Results will return the https url to the object.

1. login to node 1 as ecaadmin
2. nano /opt/superna/eca/eca-env-common.conf and add this variable **export ARCHIVEWORKER_ENABLE=true** and save the file with control+x answer Y to save
3. S3 commands
    a. searchctl s3 [-h] {add,list,index,remove}
    b. usage: searchctl s3 add [-h] --endpoint **ENDPOINT** --secretkey **SECRETKEY**

    [--accesskey ACCESSKEY] [--region **REGION**] --bucket

    BUCKET [--container **CONTAINER**] --cloudtype

    CLOUDTYPE

c. [--includes INCLUDES] - Use this to include or exclude object keys in the object store.

   [--excludes EXCLUDES]

d. Cloud types are: **aws, ecs, other**

e. Example command for ECS

   i. searchctl s3 add --accesskey username --secretkey Q5EY6abpUdbNRC7t --endpoint **https://172.25.24.53:9021** --bucket test --cloudtype ecs

f. Example command for AWS

   i. searchctl s3 add --accesskey AKIAIsdf45LN3GQ --secretkey AGV7tMIPOmlpSVsctyoqaP7k6Oxv --endpoint s3.ca-central-1.amazonaws.com --region ca-central-1 --bucket mybucketname --cloudtype aws

g. searchctl s3 list - Use to show configured s3 endpoints

h. searchctl s3 remove --id xx - Use list command to get the S3 ID to remove, where xx is the configuration id.

i. **searchctl s3 index --id xx** (**Use this command to start the S3 bucket walk and object index task, where xx is the configuration id of the s3 target**)

How to Index Snapshots for User Self Serve File Recovery

Requirements

1. Release 1.1.2 or later for snapshot monitor mode
2. **Release 1.1.4** adds snapshot only mode that provides recover portal for many snapshots for a recover solution. This is a separate mode from snapshot monitor that integrates with file system indexing. See this section.
3. Release 1.1.5 adds multiple snapshots under an indexed path and snapshots that match the indexed path, and snapshots schedules that overlap on a single path.

Limitations:

1. Limitations Snapshot monitor mode release 1.1.2:
    a. Limited to 1 snapshot monitor per indexed folder path in this release using snapshot monitor mode.
    b. Limited to paths below the indexed folders not at the same path as the indexed folder.
2. Limitations Snapshot folder only mode release 1.1.5:
    a. 25 folders added in snapshot only mode
    b. Snapshots at the same level as the indexed folder will be supported.
    c. Multiple snapshots under an indexed path will be supported.
3. **NOTE: Configure the proxy root user password to allow file recovery from snapshots. See modify PowerScale command to add root password that is used only when restoring files from snapshots. Release < 1.1.5 only CLI command is here.**

4. **NOTE: Snapshot monitor mode will start monitoring and differencing snapshots after the CLI command is applied. This means no files will be indexed until a new snapshot is created by the snapshot schedule before new files will be indexed in the snapshot path. To test this feature you will need to wait or increase the frequency of the snapshots to get changed, new files added to the snapshot.**
5. **NOTE: Expired snapshots will orphan files in the snapshot index in this release and will still be returned to users in search results but will fail to restore the file. A later release will add purge expired snapshot data from the Snapshot index which will remove expired snapshot data from the index automatically.**
6. **See the User Guide on user procedures to search and recover files.**

Understanding Snapshot Index Modes

1. **Snapshot Monitor Mode** integrates with indexed folder paths, and allows a snapshot below the index folder path to be monitored and included in incremental indexing of the existing snapshots scheduled on the PowerScale. This has the limits above with a single snapshot below the indexed folder path.
    a. **Use Case: Use this mode with a small quantity of snapshots < 3 and primary goal is file system indexing with some snapshots.**

2. **Snapshot Folder Only Mode** allows adding a folder and specify the snapshot only mode that will only index the snapshot data, and will not index the file system itself. This is the recovery only solution and multiple snapshot only folders can be added. See limitations above. This mode is recommended for backup administrator use cases that search snapshots for recovery requests.
    a. **Use Case: Use this mode to indexing snapshots only** This mode would be used when the file system indexing is not the primary objective, but recovery of snapshot data by backup administrators or end users is the objective.

Use Cases:

1. **User Self Service Restore** - Allow users to see versions of files that exist to restore from snapshots.
2. **Backup Admin** - A backup admin can find files for users in snapshots.
3. **Find deleted files in snapshots -** deleted files in the file system often exist in the snapshots. This allows users and administrators to recover files easily from snapshots when they are not present in the active file system.

How to Configure - Snapshot Monitor Mode Integrated with Folder Indexing

This mode is added to an indexed folder and will index metadata within snapshots at this path. The feature will monitor new snapshots that appear from a schedule and will difference the changes and index new files in the snapshot. This provides users or administrators with the ability to search snapshot metadata for files. **NOTE: The path must be added for indexing before the Monitor Mode can be enabled on a path at or below the indexed path configured in the appliance.**

1. User interface **will allow searching the snapshots in the advanced options window**.
2. CLI commands to configure monitoring snapshots below indexed folders that are already configured:
    a. Get the folder id of the target folder:
        i. searchctl folders list
    b. searchctl snapshotmonitor **add** --folderid <folderid> --path <snapshot path>.
        i. the folder id is the parent path folder id that is already configured for indexing, to list all folder id's: **searchctl folders list.**
    c. searchctl snapshotmonitor **remove** --folderid <folderid> --path <snapshot path> (snapshot path is the path entered that has snapshots configured in Onefs and added for monitoring of new snapshots)
    d. searchctl snapshotmonitor **list** --folderid <folderid>

How to Configure - Snapshot Only Mode

Requirements:

1.1.4 or later release

Overview:

This mode is used to only index snapshots on a cluster, to use Search & Recover as a recovery tool for backup administrators, or allow end users to search through all snapshots and perform self server file recovery. See next section on configuring global file collision settings.

How to Configure Snapshot Only Mode:

Two steps:

- Add the snapshot folder in snapshot only mode and run the baseline deep scan index on the snapshot folder
- Then you must enable snapshot monitor to index new snapshots created by the Cluster's scheduled snapshots

Step 1:

1. addsearchctl folders --isilon <cluster name> --folder <FOLDER_PATH> --snapshotMode "SNAPSHOT_ONLY" (the folder path is the path where a snapshot schedule is configured to protect data.
2. **How create the baseline Index of the snapshot folder**
    a. searchctl folders index --id <FOLDER_ID> --snapshot <SNAPSHOT_NAME>
    b. The snapshot name is the name of the snapshot in the OneFs gui or CLI and specifies which snapshot creates the

baseline index of the snapshot. **NOTE: It should be a recent snapshot not a snapshot from the past, because the snapshots between the baseline and new snapshots are NOT indexed.**

c. **NOTE:** The index folder command supports a flag to deep index the snapshot only folders, and create a baseline in the index that includes all files in the snapshot. This is a full scan of this path, but all data will be added to the snapshot index and searchable with the search previous versions option in the GUI. **See below: How to search for Snapshot only files.**

Step 2

1. searchctl snapshotmonitor add --folderid <folderid> --path <snapshot path>
2. The folder ID is the from Step 1, snapshot path is the same used in Step 1. This command will monitor for new snapshots that appear and index the differences between the baseline deep scan, executed in in Step 1, and the new snapshot that appears on this path based on the cluster schedule.

How to search for Snapshot only files

NOTE: This mode will allow searches that select the search backups flag in the user interface to locate files within snapshots. See screenshot.

How to Enable File Global Recovery Modes to control snapshot Restore file Collisions

This sets the user or administrator defaults for file recovery from snapshots feature. This restore from snapshot is a cluster side operation that copies the user selected file back into the file system, and secures the file to the user logged into the Search & Recover user interface. **NOTE: The owner of the file will be the Search & Recover service account EyeglassAdminSR.**

Requirements:

1. **Snapshot monitor mode must be enabled to index snapshots at or under a path configured for indexing.**
2. **Release 1.1.2 or later releases.**
3. **NOTE: Configure the proxy root user password to allow file recovery from snapshots. See modify PowerScale command to add root password that is used only when restoring files from snapshots. CLI command is here.**

Configuration:

1. **searchctl settings filerecovery mode {OVERWRITE,OVERWRITE_AND_BACKUP_ORIGINAL, NO_OVERWRITE_AND_RESTORE}**
2. Usage: ecactl search settings filerecovery mode --isilon HOST OVERWRITE,OVERWRITE_AND_BACKUP_ORIGINAL,NO_OVERWRITE_AND_RESTORE} (NOTE: HOST is the cluster name)
    a. **OVERWRITE - This will overwrite the file with the same name in the file system if it exists, or simply create the file in the file system from the snapshot.**
    b. **OVERWRITE_AND_BACKUP_ORIGINAL -** This will overwrite the file with the same name in the file system if it exists, and will backup the existing file in the file system with igls-original-<current_date>-<filename>.
    c. **NO_OVERWRITE_AND_RESTORE -** This mode will never overwrite the file if it exists in the file system, will not create the file if it does not exist in the file system, and will create a restore file as follows igls-restored-<date>-<filename>.

How to start a full index or incremental job on a folder path

Requirements:

1. Full index jobs supported with < 1.1.5
2. Incremental index on demand > 1.1.5

How to start a full index job on a path that has already been added:

1. searchctl folders index --id (where id is the folder id, list ID's with **searchctl folders list**)
    a. Option flags include:
        i. [--action {RESET,RECOVER}]
            1. **RESET flag** will restart the full index job from the top of the path entered (**NOTE: Use if directed by support only, this can not be undone and requires a complete re index of all data**)
                a. Example command syntact searchctl folders index --id 3fc3613a0fe814b8 --action RESET)
            2. **RECOVER flag** this flag will resume a full index from where it left off, if an index is interrupted during full ingestion. This could happen due to networking issues, REST API reachability to the cluster, or Host failure running for the Search cluster. **Note:** Node 1 is responsible for all index job management, if this host is restarted or fails, a RECOVER flag command will be required for all paths that never completed the full index.
                a. Example command syntax "searchctl folders index --id 3fc3613a0fe814b8 --action RECOVER)"

3. --incremental (requires 1.1.5 or later) This option allows running a snapshot based compare incremental job before the next scheduled incremental job.

2. Example: searchctl folders index --id 3fc3613a0fe814b8
    a. **NOTE: This will start a file and directory scan to index all files at /ifs/data and below.**
3. New commands have been added to allow targeted full re-ingestion of a single folder, or all folders below the target path.
4. **--subdir <path>** this is required to enter the path to rescan all files in the folder, but it will not walk any child paths found within this folder.
5. **--recursive** (optional, default is true) this is not required if the folder and children folders are expected to be full indexed. If only a single folder needs to be index this should be set to false.
    a. Examples:
        i. searchctl folders index --id xxxxxx --subdir /ifs/data/toindex/somesubfolder (index's this path and all children folders)
        ii. searchctl folders index --id xxxxxxxx --subdir /ifs/data/toindex/somesubfolder --recursive false (will only index the subdir folder)
6. --solrUpdate - (release 1.1.2) over time some types of file system actions can leave orphaned directories, for example renaming a directory can leave the old directory and path of old

files. This index option will fix the index and remove orphaned folders and files. This is a result of the PowerScale change list not supporting rename directory events.

7. --content (release 1.1.5 or later) This allows running an index job on a folder when the fullincludes flag was used to add additional file extensions for content indexing. This comment can be used with the --subdir command to specify where to start the scan. The index job will not tree walk the file system but will instead query the index for files that match the --fullincludes flag and place these files in the queue for content indexing. If a file is already content indexed and has not been updated the file will be skipped. This will speed up a content scan update on a large path of data when new file types are added.
    a. Any content indexing configuration on the folder will be used when searching the file system for files to be queued for content indexing assessment.

How to Manage Scheduled Jobs (Global and Folder full and incremental)

Requirements:

1. Release 1.1.5 or >

Schedule Job Definitions

1. INVENTORY - collects shares, acl's and user information for security - must be enabled

2. INCREMENTAL_INGESTION - enables incremental changelist scheduled to run against all defined folders. Default disabled
3. FULL_INGESTION - enables full index job on all folders, this will skip files already in the index with date stamp compare to the index. Default disabled
4. DAILY_REPORT_SCHEDULE - Sends daily reports at this time. Default enabled once per day
5. SOLR_HEALTH_WATCHDOG - Enables health check on the index process for support purposes.

Commands to Manage Schedules (enable, disable, set schedule)

1. Schedule modify syntax
    a. searchctl schedules modify [-h] --id ID (--schedule SCHEDULE | --disabled)
    b. SCHEDULE is a cron string with double quotes
    c. ID values can be listed with searchctl schedules list
2. List Schedules
    a. searchctl schedules list (list schedules)
3. disable a schedule
    a. searchctl schedules --ID xxxx --disable

Example Full and incremental index job Schedule Configuration

1. Enable incremental on all folders with 1 hour interval
    a. searchctl schedules modify --id INCREMENTAL_INGESTION --schedule "0 * * * *"
2. Enable full index job on all folders with 1 hour interval (note this will skip files that are already in the index automatically)
    a. searchctl schedules modify --id FULL_INGESTION --schedule "0 * * * *"

How to Monitor Index Job Status

NOTE: Execute commands on node 1 of the cluster.

1. searchctl jobs running.
    a. This command will show all running jobs full and incremental, and the current state of the job along with the date and time it started.
    b. job id folder id type started at

```
--------------------------  -------------  ----------------------------  ------------

job-1550880760575311032660  FullIngestion  Sat Feb 23 00:12:40 UTC 2019  SCANNING
```

2. searchctl jobs history.
    a. Use this command to see the start and stop times for previous full and incremental jobs, as well as the status of the job.

3. searchctl jobs view --id <job-xxxxxxxxxxxx> (Use this command to monitor the status on a running job).
4. OR searchctl jobs view --id <job-xxxxxxxxxxxx> --follow **(Use this to monitor an active running job progress through steps, with real-time updates).**
   a. Use this command to view details of the running see example below:

```
b. ecaadmin@demosearch-1:~> searchctl jobs view --id job-1550880760575311032660
```

```
Folder ID: 3fe4b6a5d4b3c899

FullIngestion (  Running ...   )
----Take snapshot of /ifs/data ( SUCCESS : 0.17 seconds )
----update snapshot alias ( SUCCESS : 0.45 seconds )
----Walking File System at /ifs/data (  Running ...   )
```

   c. See example of a completed job:

```
d. FullIngestion ( SUCCESS : 2 minutes, 13.12 seconds )
----Take snapshot of /ifs/data ( SUCCESS : 0.17 seconds )
----update snapshot alias ( SUCCESS : 0.45 seconds )
----Walking File System at /ifs/data ( SUCCESS : 2 minutes, 12.37 seconds )
```

----Collect settings (SUCCESS : 0.13 seconds)

Status: SUCCESS

How to Monitor Ingestion with the stats command

1. This command only shows stats that have values default. add --all to see all stats available.
2. searchctl folders stats --id <id of job> <--no-stream> **<--all>**
1. To get the job id of a folder index, run the index command **searchctl folders list**
2. "name": "PowerScale-1",

   "indexedFolders": [

   {

   "id": "",
3. searchctl folders stats --id **<folder ID here>** (optional flag --no-stream to get stats without auto refresh)
4. Sample stats
5. **NOTE: The rates columns is a rate per second average over the time period.**

Per Node stats command allows monitoring statistics for a single node or for all nodes

1. ecactl search stats view (--folder <folder_id> | --node <node_id>) [--all] [--no-stream]
2. If entering a node the stats will be specific to the nodes processing of indexed data.

```
Statistics for folder: 3fe3631c41a7e74a

name                                  total_alltime    total_min    total_hr
  total_day    rate_min    rate_hr    rate_day

------------------------------------  ---------------  -----------  ----------  -----------  --
--------  ---------  ----------

FULL/FILES_ACCEPTED                                0          0          0
        0          0          0          0

FULL/FILES_CONTENT_ERRORED                         0          0          0
        0          0          0          0

FULL/FILES_CONTENT_INDEXED                         0          0          0
        0          0          0          0

FULL/FILES_IGNORED                                 0          0          0
        0          0          0          0

FULL/FILES_METADATA_ERRORED                        0          0
        0          0          0          0          0

FULL/FILES_METADATA_INDEXED                        0          0          0
        0          0          0          0

FULL/FOLDERS_ACCEPTED                              0          0          0
        0          0          0          0

FULL/FOLDERS_IGNORED                               0          0          0
        0          0          0          0
```

```
FULL/FOLDERS_METADATA_ERRORED                0        0
        0        0        0        0        0

FULL/FOLDERS_METADATA_INDEXED                0        0
        0        0        0        0        0

INCREMENTAL/FILES_ACCEPTED                   2        0        0
        0        0        0        0

INCREMENTAL/FILES_CONTENT_ERRORED            0        0
        0        0        0        0        0

INCREMENTAL/FILES_CONTENT_INDEXED            0        0
        0        0        0        0        0

INCREMENTAL/FILES_IGNORED                    0        0        0
        0        0        0        0

INCREMENTAL/FILES_METADATA_ERRORED           1        0
        0        0        0        0        0

INCREMENTAL/FILES_METADATA_INDEXED           0        0
        0        0        0        0        0

INCREMENTAL/FOLDERS_ACCEPTED                 0        0
        0        0        0        0        0

INCREMENTAL/FOLDERS_IGNORED                  0        0
        0        0        0        0        0

INCREMENTAL/FOLDERS_METADATA_ERRORED         0
        0        0        0        0        0        0
```

```
INCREMENTAL/FOLDERS_METADATA_INDEXED              0
        0         0         0         0         0         0
```

Running Inventory Scans and Viewing users and SMB Share Access

1. To run inventory command and collect cluster information:
    a. searchctl isilons runinventory.
2. To display AD users collected from inventory:
    a. searchctl isilons list --users.
3. To display SMB shares collected from inventory:
    a. searchctl isilons list --shares.
4. To display details about a users SMB share path access. Use this command to identify the filters applied to search results for a given user. It will list the Access zone, the path and the cluster:
    a. searchctl users view --name **user@domain.com.**
    b. **searchctl users view --name 'DOMAIN\\user'** . (NOTE: the domain must be upper case and double backslash must be used to seperate the user from the domain) .
    c.

```
Attribute    Value
-----------  ------------------------------------------  -------------  -----------  ----
----------
Name:        AD01\dfs1
```

```
SID:        S-1-5-21-1825440792-1775492485-428706412-1157

DLLN:       AD01\dfs1

UPN:        dfs1@AD1.TEST

Shares:

            Path                                Share Name     Access Zone
 PowerScale

            ----                              ----------     -----------  ------

            /ifs/data/userdata/dfs1                    igls-dfs-dfs1  data
       prod-cluster-8

            /ifs/data/userdata/share2                  share2        data
       prod-cluster-8

            /ifs/data/userdata/share1                  share1        data
       prod-cluster-8

            /ifs/data/policy1                       SMB2          System
      prod-cluster-8
```

How to Enable User Authentication to Data within Access Zones and Return Search Results with Smartconnect UNC's to Files

This section is required to enable authentication to the WebUI for users, allows users to see Smartconnect UNC path to files for copying to the clipboard, and opening files from Windows Explorer or Mac

Finder. If this is not configured, users will see a full path to the file from /ifs which will not be accessible without a UNC path to the file.

Authentication Requirements for User Data in Access Zones

The configured FQDN per Access Zone setting is required for all Access Zones that will have users authenticating on the WebUI login page. Each user that logs in will have the userID and password checked against each Access Zone FQDN configured, to verify the users has access to data. The first Access Zone that validates the user credentials will exit the authentication process and proceed to identify all SMB shares in all Access Zones. Review the authentication data flow below.

NOTE: Each Access Zone used to authenticate users MUST have a configured FQDN entered in to the configuration, AND MUST have at least 1 SMB shared within the Access Zone to be used for authentication and password validation.

Authentication Data Flow

1. **Access Zone system - FQDN UNC authentication request with Userid and password against an SMB share discovered in System Zone.**
    a. **If successful exit and identify data access to SMB shares in ALL Access Zones.**
    b. **If authentication fails check next Access Zone UNC FQDN that was configured.**

2. **Access Zone Data1 - FQDN UNC authentication request with userID and password against an SMB share discovered in Data1 Zone.**
    a. **If successful exit and identify data access to SMB shares in ALL Access Zones.**
    b. **if authentication fails check next Access Zone.**
3. **Repeat until user is authenticated or denied access to search login page.**
4. **At the end of this process all SMB shares in ALL zones are used to filter login results to the user.**

CLI commands to add Zone FQDN to Authentication and Search Results Display

1. searchctl settings zoneunc add --isilon clusterA --zone ZoneA --fqdn mycluster.example.com
    a. This command will return all search results on ClusterA for all files in ZoneA with \\mycluster.example.com\<sharename>
2. searchctl settings zoneunc list
3. searchctl settings zoneunc remove --isilon clusterA --zone ZoneA

How to Enable Administrator Search Security Override

By default results are secured by SMB permissions, file ownership or both SMB access and then file ownership. File owner security will block results for administrator use cases. This feature is used to disable all security for administrator users to execute file system searches that are required to analyze the file system and use the automation script feature or eDiscovery use cases **even when their AD account does not have access to the data on the cluster**.

Compliance, File System Antlaytics Administrators

A list of AD user ID's are added to an approved list to override all security on search results. Local user accounts on the applicance are also supported when AD users are not required, and local login to the search UI is required.

Use cases for this feature

1. **Compliance officer content search**
2. **eDiscovery administrator**

How to configure a List of Search administrators

New in 1.1.1 build AD group support for administrator groups.

usage: ecactl search settings admins add [-h] [--name NAME] [--group GROUP] [--local]

1. searchctl settings admins [add | list | remove] searchctl settings admins add --name <userid> [--local]
    a. searchctl settings admins add --name username@domain.com
        i. To Use local user on the appliance without using AD accounts
            1. searchctl settings admins add --name **ecaadmin** --local **(local default OS account on the cluster)**
            2. NOTE: This uses local users and AD is always the best practice. A single default OS account already exists on the appliance that can be used.
            3. NOTE: Authentication is against the password in the /opt/superna/eca/conf/nginx/.htpasswd file instead of the PowerScale AD provider. Additional users must manually be added to the .htpasswd file AND via searchctl admins add.
            4. To add additional local users
                a. ssh ecaadmin@xxxx (ip of node 1)
                b. sudo -s
                c. useradd xxxx (xxxx is the name of the user)
                d. exit
                e. Set the web login password (Note OS login is not required) change the user name for xxxx and yyyy for the password.

i. ecactl cluster exec "htpasswd -b /opt/superna/eca/conf/nginx/.htpasswd xxxx yyyy"

f. Add the local admin with:

i. searchctl settings admins add --name xxxx --local

ii. To add AD group of administrators

1. searchctl settings admins add --group "groupname@domain.com" (use groupname case senstive with @ AD DNS name syntax)

b. searchtl settings admins remove --sid SID (the SID is required to remove a user)

c. To remove group admin use domain syntax with 3 \\\ to escape the \ character

i. searchctl settings admins remove --sid "AD01\\\mlrsw1"

How to add Data Owner Search Administrators

This feature allows a data search administator that has span of control over one or more paths, regardless of SMB or ACL permissions modes set on the indexed folder. The paths added to these data owners is added to the users existing security profile and allows them to search file metadata and content within documents on the allowed path.

For example if the users SMB share permissions grants them access to a subset of the data, based on AD groups and SMB share permissions, the paths added with this Data Owner admin will be added to the existing auto detected permissions and allow analytics and searching at this path and below.

This feature also allows the data owner to be restricted to metadata only. It will block content searching to protect data in the index and allow the data owner admin to report on data, but not identify by content. This is available in 1.1.1 or later releases.

Use cases for this feature

1. **Department admin for reporting on business unit data**
2. **Project administrator for reporting on project data**

NOTE: If the Data Owner Admin does not have access to the files in the search results or reports, they will not be able to open the files, and no additional file access is possible from the Search results.

NOTE: The PowerScale must be added to the appliance inventory and the path must be at or below an indexed folder configurd in the system.

Data Owner Admin Configuration

1. Add Data Owner Admin to a path:
    a. **searchctl adminaccesslist add --user <user@domain.com> --isilon <PowerScale Name> --path </ifs/path/to/folder> [--metadata-only] .**
2. Remove Data Owner Admin from a path:

   a. **searchctl adminaccesslist remove --user <user@domain.com> --isilon <PowerScale Name> --path </ifs/path/to/folder> .**

3. List Data Owner Administrator assigned paths:
   a. **searchctl adminaccesslist list --user <user@domain.com> --isilon <PowerScale Name> .**

How to Configure admin only login mode and block user login

This is for administrator only mode, where end users do not need to login to the UI, and allows an administrator listed on the admin list to login while all other users are blocked if not on the list. NOTE: This includes the local ecaadmin account.

1. Login to node 1 of the Search cluster as admin over ssh .
2. edit conf file and make the change below.
3. vim /opt/superna/eca/eca-env-common.conf .
4. Add this line **export SEARCHMW_ADMIN_MODE_ONLY=true .**
5. Save the file.
6. Shut down the cluster and start up again to take effect.
7. ecactl cluster down .
8. wait until down completes.
9. ecactcl cluster up .

10. Now only users listed on the admin list will be allowed to login to execute searches.

© Superna LLC

9.11. Cluster Administration - UI Access and Security Configuration

Home Top

Cluster Administration and Operations tools are available using special URL's available on node 1 ip address of the cluster. These WebUI's are secured with a password configured during deployment. This section also covers how to change passwords on admin tool UI and change the self signed cert on the webUI to a signed certificate.

- URL to Access Admin UI's
- How to add a signed cert to user login GUI

URL to Access Admin UI's

1. **Solr index engine** - allows index health status, document count, index size, error messages, advanced queries for administrators only.
    a. https://x.x.x.x/solr
2. **Kafka Message processing** - used to process file ingestions for full and incremental jobs. HA features and cluster wide view of processing of messages.
    a. https://x.x.x.x/kafka-manager
3. **How to Access**
    a. Access the URL enter the user "ecaadmin" and the password that was configured during deployment.

How to add a signed cert to user login GUI

Best Practice:

1. Access the WebUI from node 1 and create a DNS entry for node to create a FQDN to create a signed cert.
2. The objective is to install the signed cert for nginx ECA Node-1
3. Create A record in DNS name for ECA Node-1 and verify with nslookup.. Example eca1.domain.com
4. SSH to ECA Node-1 as ecaadmin
5. cd /opt/superna/eca/conf/nginx
6. Verify that the nginx.key is there with **ls -la**
7. Create csr with that key file
    a. **Command: openssl req -key nginx.key -new -out nginx.csr**
    b. **SCP the nginx.csr file for signing**
    c. **Or type cat nginx.csr and copy and paste the text to submit for signing.**
8. When it is asked about the Common Name: provide the **fqdn** of ECA Node-1 (the name registered in DNS e.g. search.domain.com)
9. **With that CSR certificate submit the request to Certificate Authority at your enterprise**
    a. **NOTE: These steps are CA specific, consult with your security team**
10. **Once received the signed certificate encoded in PEM format**
    a. scp (use WinSCP for WIndows) and copy this file to ECA-1 under /opt/superna/eca/conf/nginx with name nginx.crt

b. **NOTE: if not in PEM format, convert to PEM format or ask your Security team for pem format**

c. Replace existing nginx.crt certificate with this new signed CA certificate.

d. mv nginx.crt nginx.crt.bak (backup old file)

e. cp /**pathtonewfile**/nginx.crt
to /opt/superna/eca/conf/nginx/nginx.crt

11. **Restart nginx**

a. ecactl containers stop nginx

b. ecactl containers start nginx

12. **Or bring down and up the ECA cluster to push the config to all the other ECA nodes**

a. ecactl cluster down

b. ecactl cluster up

13. Verify the certificate when accessing the UI (e.g. https://FQDN)

© Superna LLC

9.12. Search & Recover Cluster Operations

Home Top

- Cluster Operations CLI commands
- How to Initialize Cluster Management Across all nodes
- How to List all Kafka topics and definitions
- How to Start and Stop the Cluster
- How to Change the IP address of an PowerScale Cluster
- How to change TLS security settings when connecting to clusters that do not support the highest security algorithms
- How to Enable and Use PhoneHome support
- How to Upgrade the cluster Online
- How to change the downloads and Admin tool UI password
- How to collect support logs and submit a support case
- Backing up and Restore the Cluster Configuration
- How to Configure Automated backups to store on External Storage and protect the appliance Configuration
- How to check for Alarms
- How to configure Alarm notification
  - Quick Start SMTP Configuration for Notifications :
  - Setup Syslog channel for Notifications:
  - Setup an SMTP or Syslog Channel for Notifications:
  - Create a Notification Group:
  - Manage Recipients for SMTP and Syslog channels:
  - Manage Notification Suppression Alarm Configurations:

- How to force commit indexed files to the Index
- Index Weekly Maintenance Task Schedule
- How to factory reset an installation (Warning this deletes the index)
- How to delete a path of data folders and files from the Index
- How to Expand the Capacity of the Index Path on each VM
- How to change the TLS Certificate for the WebUI with a new Self Signed Cert
- Content Indexing Performance Advanced Configuration Settings
  - Scale Out Content Indexing & Classification Only VM Deployment Option

Cluster Operations CLI commands

The following sections cover cluster operation commands.

How to Initialize Cluster Management Across all nodes

Use this command once after installation to setup ssh keys and SSL certificate for the web UI as well sync key files to all nodes.

1. ecactl components install eca .

How to List all Kafka topics and definitions

1. searchctl topics list
2. searchctl kafka describe --topic <topic_name>

How to Start and Stop the Cluster

1. **ecactl cluster up** - starts up cluster services on all nodes .
2. **ecactl cluster down** - stops all cluster services on all nodes.

How to Change the IP address of an PowerScale Cluster

1. Use this command to change the ip address of a cluster that has already been added to the Search & Recover appliance.
2. searchctl isilons modify --name NAME --ip IP **(name is the cluster name as reported by searchctl isilons list, IP is node ip in system zone in a an IP pool with dynamic mode enabled) .**

How to change TLS security settings when connecting to clusters that do not support the highest security algorithms

1. In some cases it is required to change the TLS security settings used to connect to clusters or other devices.

2. The java security file can be found here /opt/superna/eca/conf/java.
3. Edit this file with vim to make changes to the settings.
4. To apply the settings restart the isilongateway container.
5. ecactl cluster containers restart isilongateway.
6. done.

How to Enable and Use PhoneHome support

To enable remote log collection for faster support of cluster issues, enable the phonehome feature.

Login to node 1 and run this command to start the registration process, and listen for remote log upload requests.

1. ecactl phonehome now .
2. If the appliance has a valid license key and the following whitelisted urls below, the registration will succeed. If an Internet proxy is required then configure the proxy .
    a. Proxy configuration.
    b. sudo -s (enter admin password).
    c. Type yast.
    d. Enter proxy configuration under network services menu and proxy and save.
3. whitelist urls
    a. The Monitoring service requires the following URL's allowed:

    b. See Here for details on whitelist

4. How to Disable phonehome:
    a. ecactl phonehome stop
5. How to upload logs directly to support:
    1. ecactl phonehome logsupload

How to Upgrade the cluster Online

This command will check for new code online and download the code, and requires Internet access from the appliance to *.superna.net URL over HTTPS.

1. Each login to the cluster will check for new software upgrades.
2. The command to run the upgrade:
    a. ecactl cluster upgrade.
3. NOTE: This will download the new code, shutdown the cluster, upgrade the cluster code.
4. Verify if any upgrade was successful and look for errors.
5. Start up the cluster:
    a. ecactl cluster up
    b. verify no startup errors are visible.

How to change the downloads and Admin tool UI password

Installation requires the WebUI password to be set. To reset or change this password follow these steps:

1. Login to node 1 over ssh as ecaadmin user and run the command below
    a. NOTE: Replace <password> with the password
2. ecactl cluster exec "htpasswd -b /opt/superna/eca/conf/nginx/.htpasswd ecaadmin <password>"
3. done. The new password is active immediately on all nodes.

How to Display Cluster Diagnostic and Cluster Version

1. **ecactl cluster diagnostics -** use this command to display cluster wide health and diagnostic data (FUTURE)
2. **ecactl version -** (shows the current code version).
3. **ecactl zk -** (advanced support only commands for zookeeper).

How to collect support logs and submit a support case

1. **ecactl cluster loggather --tail 5000** - Use this command to collect logs and create support zip. --tail option specifies how many lines are collected for certain logs and will reduce file size. Suggested value 5000 for support.

2. To get the appliance ID required to upload support data:
    a. run the command **ecactl version .**
    b. Record the appliance ID value. This is required to upload support logs.
3. Download the support zip from:
    a. https://x.x.x.x/downloads/loggather/
4. Then follow normal support backup upload instructions located here.

Backing up and Restore the Cluster Configuration

To protect the configuration of the search cluster, the backup should be created and and stored off the appliance.

Backup command

1. **ecactl cluster backup .**
2. Results are stored in /opt/superna/var/search/downloads/archive/ .
3. To access the downloads page open a browser and enter https://x.x.x.x/downloads/ of node 1 ip address. This will require the administrator WebUI user and password set up during installation. **User name is ecaadmin and password used during setup.**

a.

```
172.16.82.176 - PuTTY
ecaadmin@ineca-1:~> ecactl cluster backup

Initiating configuration backup...

Connecting to node 172.16.82.178....
Creating tmp directory /tmp/20181123160919/ineca3...
Copying environment configuration file...
Copying docker overrides file...
Copying common configurations...
Copying nginx configurations...
Copying zookeeper configurations...
Copying solr configurations...
Copying keystore configurations...
Sync configuration backup...
Removing tmp directory /tmp/20181123160919/ineca3....
*******************************************

Connecting to node 172.16.82.177....
Creating tmp directory /tmp/20181123160919/ineca2...
Copying environment configuration file...
Copying docker overrides file...
Copying common configurations...
Copying nginx configurations...
Copying zookeeper configurations...
Copying solr configurations...
Copying License configurations...
Copying keystore configurations...
Sync configuration backup...
Removing tmp directory /tmp/20181123160919/ineca2....
*******************************************

Connecting to node 172.16.82.176....
Creating tmp directory /tmp/20181123160919/ineca1...
Copying environment configuration file...
Copying docker overrides file...
Copying common configurations...
Copying nginx configurations...
Copying zookeeper configurations...
Copying solr configurations...
Copying License configurations...
Copying keystore configurations...
Sync configuration backup...
Removing tmp directory /tmp/20181123160919/ineca1....

Creating zip file....
Cleaning tmp directories...
*******************************************

Configuration backup successful

ecaadmin@ineca-1:~>
```

b.

```
ecaadmin@search-1:/opt/superna/var/search/downloads/archive> ll
total 30652
-rw-r--r-- 1 ecaadmin ecaadmin 31385740 Dec  8 09:01 search_config_backup_20181208090051.zip
ecaadmin@search-1:/opt/superna/var/search/downloads/archive>
```

Restore Command:

1. NOTE: The cluster must be shut down first to restore. ecactl cluster down
2. ecactl cluster restore --path /opt/superna/var/search/downloads/archive/<name of backup file>.zip
    a.

```
ecaadmin@ineca-1:~> ecactl cluster restore --path /opt/superna/var/archive/search_config_backup_20181123160919.zip

Initiating restore...

Extracting file /opt/superna/var/archive/search_config_backup_20181123160919.zip...

Restoring configuration on node 172.16.82.176....
sending incremental file list
docker-compose.overrides.yml
eca-env-common.conf
conf/common/ThreatLevels.json
conf/common/keystore.jceks
conf/common/licensedNe.json
conf/common/queuebfs.json
conf/common/schedule.json
conf/common/sync.xml
conf/common/whitelist.json
conf/common/overrides/.gitignore
conf/nginx/eca.conf
conf/nginx/nginx.crt
conf/nginx/nginx.key

sent 1,011 bytes  received 419 bytes  2,860.00 bytes/sec
total size is 11,603  speedup is 8.11
sending incremental file list
authorized_keys
id_rsa
id_rsa.pub
known_hosts

sent 301 bytes  received 158 bytes  306.00 bytes/sec
total size is 6,368  speedup is 13.87
sending incremental file list
solr.xml
zoo.cfg
configsets/_default/conf/managed-schema
configsets/_default/conf/params.json
configsets/_default/conf/protwords.txt
configsets/_default/conf/solrconfig.xml
configsets/_default/conf/stopwords.txt
configsets/_default/conf/synonyms.txt
configsets/_default/conf/lang/contractions_ca.txt
configsets/_default/conf/lang/contractions_fr.txt
configsets/_default/conf/lang/contractions_ga.txt
configsets/_default/conf/lang/contractions_it.txt
configsets/_default/conf/lang/hyphenations_ga.txt
configsets/_default/conf/lang/stemdict_nl.txt
configsets/_default/conf/lang/stoptags_ja.txt
configsets/_default/conf/lang/stopwords_ar.txt
configsets/_default/conf/lang/stopwords_bg.txt
configsets/_default/conf/lang/stopwords_ca.txt
configsets/_default/conf/lang/stopwords_cz.txt
configsets/_default/conf/lang/stopwords_da.txt
configsets/_default/conf/lang/stopwords_de.txt
configsets/_default/conf/lang/stopwords_el.txt
configsets/_default/conf/lang/stopwords_en.txt
configsets/_default/conf/lang/stopwords_es.txt
configsets/_default/conf/lang/stopwords_eu.txt
configsets/_default/conf/lang/stopwords_fa.txt
configsets/_default/conf/lang/stopwords_fi.txt
configsets/_default/conf/lang/stopwords_fr.txt
configsets/_default/conf/lang/stopwords_ga.txt
configsets/_default/conf/lang/stopwords_gl.txt
configsets/_default/conf/lang/stopwords_hi.txt
configsets/_default/conf/lang/stopwords_hu.txt
configsets/_default/conf/lang/stopwords_hy.txt
configsets/_default/conf/lang/stopwords_id.txt
```

b.

```
172.16.82.176 - PuTTY
configsets/_default/conf/solrconfig.xml
configsets/_default/conf/stopwords.txt
configsets/_default/conf/synonyms.txt
configsets/_default/conf/lang/contractions_ca.txt
configsets/_default/conf/lang/contractions_fr.txt
configsets/_default/conf/lang/contractions_ga.txt
configsets/_default/conf/lang/contractions_it.txt
configsets/_default/conf/lang/hyphenations_ga.txt
configsets/_default/conf/lang/stemdict_nl.txt
configsets/_default/conf/lang/stoptags_ja.txt
configsets/_default/conf/lang/stopwords_ar.txt
configsets/_default/conf/lang/stopwords_bg.txt
configsets/_default/conf/lang/stopwords_ca.txt
configsets/_default/conf/lang/stopwords_cz.txt
configsets/_default/conf/lang/stopwords_da.txt
configsets/_default/conf/lang/stopwords_de.txt
configsets/_default/conf/lang/stopwords_el.txt
configsets/_default/conf/lang/stopwords_en.txt
configsets/_default/conf/lang/stopwords_es.txt
configsets/_default/conf/lang/stopwords_eu.txt
configsets/_default/conf/lang/stopwords_fa.txt
configsets/_default/conf/lang/stopwords_fi.txt
configsets/_default/conf/lang/stopwords_fr.txt
configsets/_default/conf/lang/stopwords_ga.txt
configsets/_default/conf/lang/stopwords_gl.txt
configsets/_default/conf/lang/stopwords_hi.txt
configsets/_default/conf/lang/stopwords_hu.txt
configsets/_default/conf/lang/stopwords_hy.txt
configsets/_default/conf/lang/stopwords_id.txt
configsets/_default/conf/lang/stopwords_it.txt
configsets/_default/conf/lang/stopwords_ja.txt
configsets/_default/conf/lang/stopwords_lv.txt
configsets/_default/conf/lang/stopwords_nl.txt
configsets/_default/conf/lang/stopwords_no.txt
configsets/_default/conf/lang/stopwords_pt.txt
configsets/_default/conf/lang/stopwords_ro.txt
configsets/_default/conf/lang/stopwords_ru.txt
configsets/_default/conf/lang/stopwords_sv.txt
configsets/_default/conf/lang/stopwords_th.txt
configsets/_default/conf/lang/stopwords_tr.txt
configsets/_default/conf/lang/userdict_ja.txt
igls_shard2_replica_n4/core.properties
igls_shard3_replica_n10/core.properties
iglshistory_shard2_replica_n4/core.properties
iglshistory_shard3_replica_n10/core.properties

sent 4,350 bytes  received 2,968 bytes  14,636.00 bytes/sec
total size is 212,702  speedup is 29.07
sending incremental file list
var/myid
var/version-2/acceptedEpoch
var/version-2/currentEpoch
var/version-2/log.100000020
var/version-2/log.100000311
var/version-2/snapshot.0
var/version-2/snapshot.10000001f

sent 66,079 bytes  received 114,915 bytes  72,397.60 bytes/sec
total size is 134,221,356  speedup is 741.58
*******************************************

Cleaning tmp directories...

Successfully restored backup
```

How to Configure Automated backups to store on External Storage and protect the appliance Configuration

1. The appliance will create a daily backup and keep 7 backups without any configuration.
2. The auto backups are stored here: /opt/superna/var/search/downloads/archive/auto .
3. Create an NFS export on Isilon to store these backups. Example: /ifs/goldencopy/dailybackup with read/write permissions
4. Edit /etc/fstab to mount this Isilon export and save the daily backup external to the appliance. Example: fstab entry located below can be customized for your environment.
5. **Example fstab entry**
    a. FQDN:/ifs/goldencopy/dailybackup nfs nfsvers=3 0 0 (where FQDN is the DNS name to connect to the cluster) save the file .
    b. mount -a (will mount the exports in fstab)
    c. mount (to check the mount was successful)

How to check for Alarms

This command is used to check overall system events:

1. **searchctl notifications list**
2. The data returned includes the error code , severity of the event, the time of the event, subsystem involved and a description. Depending on the event type additional fields may be populated as well.

How to configure Alarm notification

The alarm system supports channels with smtp and syslog as available choices. The alarm system supports groups of recipients that can be added to a channel. A suppression feature allows rate limiting events or disabling them.

usage: searchctl notifications [-h]

{list,channels,groups,recipients,suppkeys}

optional arguments:

-h, --help show this help message and exit

operations:

{list,channels,groups,recipients,suppkeys}

list Lists notifications

channels Commands for adding/removing/viewing notification

channels

groups Commands for adding/removing/viewing notification

groups

recipients Commands for adding/removing/viewing notification

recipients

suppkeys Commands for adding/removing/viewing suppression keys

suppkeys

Quick Start SMTP Configuration for Notifications :

1. Create a notification channel.
2. Test the notification channel.
3. Create notification channel group.
4. Add recipients to the channel group.
5. **Create a notification group and add recipients :**
    a. Examples:
        i. Anonymous SMTP example:
            1. searchctl notifications channel addsmtp --host x.x.x.x --port 25 --sender fromemail@domain.com --tlsEnabled false
        ii. Authenticated connection example:
            1. searchctl notifications channel addsmtp --host x.x.x.x --port 25 --sender fromemail@domain.com --tlsEnabled false --user supernacorp (NOTE this will prompt for password, --user is only required authenticated SMTP connection)

6. **Test the channel Configuration with a test email:**
    a. Test the configuration (**replace yellow highlighted with values from your environment**).
    b. searchctl notifications channels testsmtp --id A5tNdkgIT --recipient fromemail@domain.com (--id is the channel id found from searchctl notifications channels list
    c. The response will return diagnostic information of the test send message.
7. **Add Notification Group (optional step):**
    a. **Optional step - a default group exists called default that is set to INFO level. Only create a group if you want to set the severity filter to a higher level.**
    b. searchctl notifications groups add --name **alerts** --severity **MAJOR(The severity option will send alarms and reports with a severity of INFO or higher, a value of MAJOR will only send MAJOR and CRITICAL alarms)**
8. **Add Recipients to the Channel group:**
    a. Get the channel ID using this command for the next step: **searchctl notifications channel list**
    b. searchctl notifications recipients addsmtp --group **default** --channel xxxxxxx --to targetemail@example.com
    c. Modify the yellow highlights for your environment (assumes default group is used) .

9. Done.

Setup Syslog channel for Notifications:

1. Create the notification channel:
    a. searchctl notifications channels addsyslog --logger SYSLOG .
2. Configure and test a Syslog channel:
    a. edit /opt/superna/var/notification/log4j2.xml .

```
b.<Syslog name="SupernaSyslog" format="RFC5424"
  facility="LOCAL0"          host="172.25.8.25"
  port="5140" protocol="UDP" appName="EyeglassSearch"
  messageId="SystemEvent" id="Event"
  connectTimeoutMillis="10000"          newLine="true"
  mdcId="mdc" includeMDC="true"
  enterpriseNumber="18060"> </Syslog>
```

    c. Locate the host and change the ip address to the ip of the syslog server you want to send events.
    d. searchctl notifications channels testsyslog --id MMhjYPd0e (get the channel id by running searchctl notifications channels list) .
    e. This will send a test syslog entry to the syslog to verify communications.
    f. Optionally add a suppression key see below.
    g. Done.

Setup an SMTP or Syslog Channel for Notifications:

1. **Configuration** :

a. searchctl notifications channel addsmtp --host x.x.x.x --port 25 --sender fromemail@domain.com --tlsEnabled false --user supernacorp .
    i. This example assumes user is required for authentication and prompt for password is required to complete the command.
b. searchctl notifications channel addsyslog --logger SYSLOG (enter SYSLOG for logger) .

2. **List :**
    a. searchctl notifications channel list --verbose .
3. Remove configuration:
    a. searchctl notifications channels remove --id xxxxxx (use list command for id) .

Create a Notification Group:

A notification group is a channel plus alarm severity and receipts. This allows different emails to receive different severity alerts

1. Create Notification group and assign receipts:
    a. searchctl notifications groups add --name INFO --severity INFO (**name is any name you want to assign, severity is the minimum alarm level that will be sent on this channel** {CRITICAL,MAJOR,MINOR,WARNING,INFO,TRACE}).
    b. Assign recipients to the SMTP channel group (see recipients commands below):

i. searchctl notifications recipients addsmtp --group INFO --channel xxxxxxx --to targetemail@example.com (note channel id can be found listing all the configured channels).

ii. Done - emails will now be sent.

2. Group commands:

   a. searchctl notifications groups add .

   b. searchctl notifications groups list (list existing).

   c. searchctl notifications groups remove --name xxxx (name of the group).

Manage Recipients for SMTP and Syslog channels:

Use these commands to add recipients for channel groups.

1. Assign recipients to a channel and group:

   a. searchctl notifications recipients addsmtp --group NAME --channel xxxx --to email@example.com (add group name, channel ID see above for list channels to get the ID , email target)

   b. searchctl notifications recipients addsyslog --group NAME --channel xxxx **(add group name, channel ID see above for list channels to get the ID)**

2. **Manage recipients commands**

a. **searchctl notifications recipients list --group xxxxx** --verbose **(xxxxx is the group name)**
b. searchctl notifications recipients remove --id --group (ID is the channel ID and group is group ID, use the list command to get both ID's)

Manage Notification Suppression Alarm Configurations:

Use this optional configuration to rate limit a specific alarm code for a specific recipient group to only receive a certain alarm once per hour, day or never. This can be useful for notifications that are reoccurring and need to be rate limited.

1. Add a suppression key:
    a. searchctl notifications suppkeys add --group GROUP --code CODE --frequency {ONE_PER_HOUR,ONE_PER_DAY,NEVER}
        i. group is the group name, code is the alarm code, frequency is used to ensure the rate of the alarm never exceeds the values of once per hour, once per day or never).
    b. searchctl notifications suppkeys list --group name (enter a group name to list the group with an assigned suppression configuration).
    c. searchctl notifications suppkeys remove --group GROUP --id ID (group and supresskey id, use the

list command for the group to get the suppression key ID) . This will remove the suppression key configuration from the group.

How to force commit indexed files to the Index

The files that are processed by a full index or incremental scan job are not committed to the index until a threshold in time or number of files is crossed. For testing purposes or when using the path level re ingestion cli command the commit command will force the processed files into the index. **NOTE: This is an expensive disk operation and should not be used frequently unless required.**

1. searchctl solr commit

Index Weekly Maintenance Task Schedule

Every Saturday a task runs on each indexed folder to remove files and directories that have been moved. This process compares the index to what PowerScale is showing on a folder by folder basis, and issues deletes to the index based on orphaned folders in the index. This process runs at midnight and can take several hours to complete.

The schedule can be changed setting a cron string in the /opt/superna/eca/eca-env-default.conf file and locate this variable export FULL_INGESTION_TASK_CRON="0 0 * * 6"

A cli command allows running this job on demand as well. See the CLI guide for folders here.

How to factory reset an installation (Warning this deletes the index)

Do not execute this command unless directed by support. This will delete all configuration settings including the index.

1. ecactl cluster down .
2. ecactl cluster factory-reset .
    a. Enter the ecaadmin password and wait for the commands to finish.
3. To restart with a fresh configuration start up the cluster to create an empty index:
    a. ecactl cluster up .

How to delete a path of data folders and files from the Index

Sometimes it may be desired to remove data in bulk from the index. This procedure allows a delete from the index using a path based solution that deletes all files and folders.

1. Login to the Solr index .
2. https://x.x.x.x/solr (x.x.x.x is node 1 of the Search & Recover appliance) .
3. Enter the ecaadmin user name and password set during installation .
4. On the left hand menu select the collection option and pick igls .
5. Then select the documents menu option .
6. Complete the screen as shown and enter the path with the exact characters as shown in the screenshot and click submit the document. **Note: The delete can take hours to take affect in search results. The delete will delete all files that begin with the path that was entered.**
    a. Text to copy and paste note the slashes are required as shown: **<delete><query>"path:\/ifs\/data\/somepath"</query></delete>**

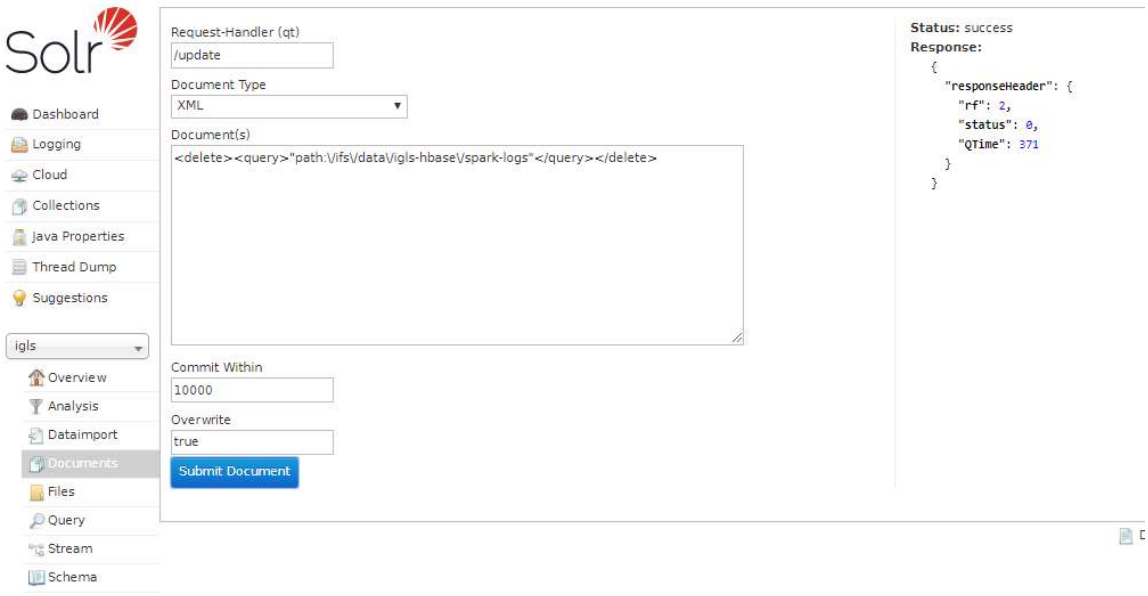

7.
8. Now login to Search & Recover node 1 as ecaadmin.
9. Commit index changes with this command

a. **searchctl solr commit**

How to Expand the Capacity of the Index Path on each VM

If the index disk gets up to 80%, space must be added. Space is consumed if the number of files consuming metadata space uses up available space, or if content indexing is enabled. Content indexing space consumption is expected and will require expanding the disks on all virtual machines 2 through N, where N is the last VM in the cluster.

Requirements:

Access to VMware vCenter must be available to add disk space to the VM.

Procedure:

1. Create a new virtual disk of at least **200GB**
2. Collect Virtual Device Node i.e.: This example uses `SCSI (0:2)` where `0` is the host node ID

3. 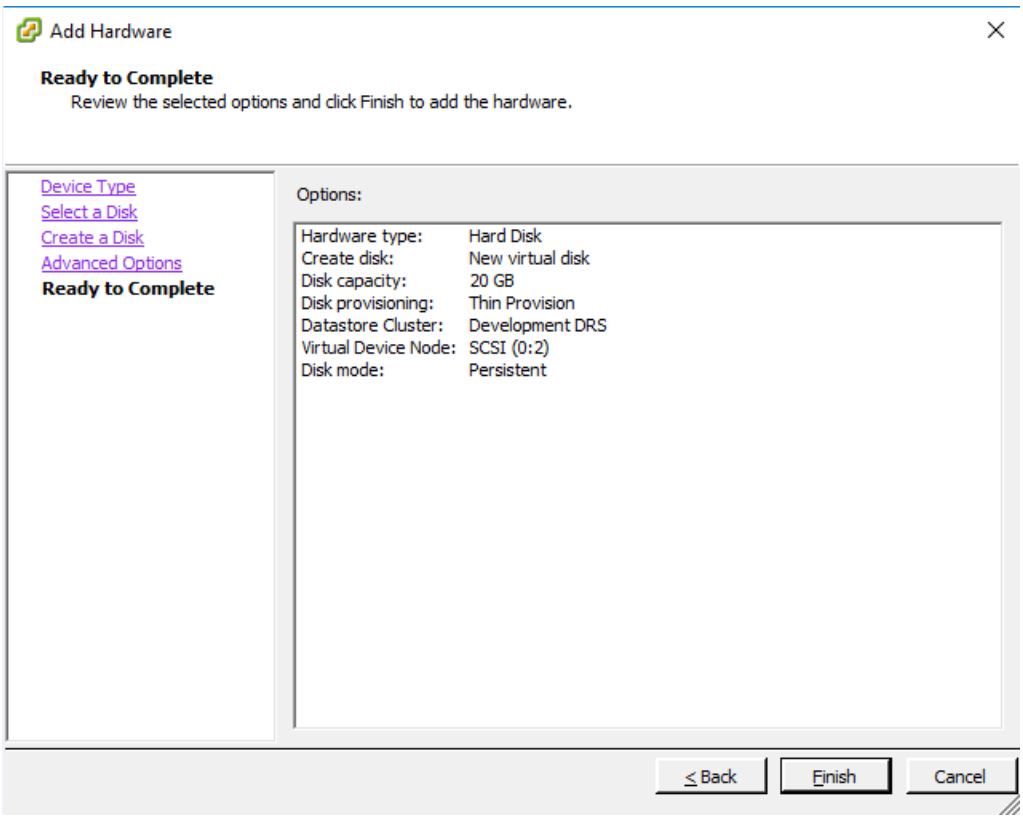

4. Login to Search & Recover node 2 through N as ecaadmin .
5. Switch to root with sudo -s (enter ecaadmin password) .
6. Scan for the new disk in OS. **(host0 should be the default if you do not find a disk contact support).**
    a. echo "- - -" > /sys/class/scsi_host/host0/scan
7. Check for new disk:
    a. fdisk -l
    b. See image below:

c. 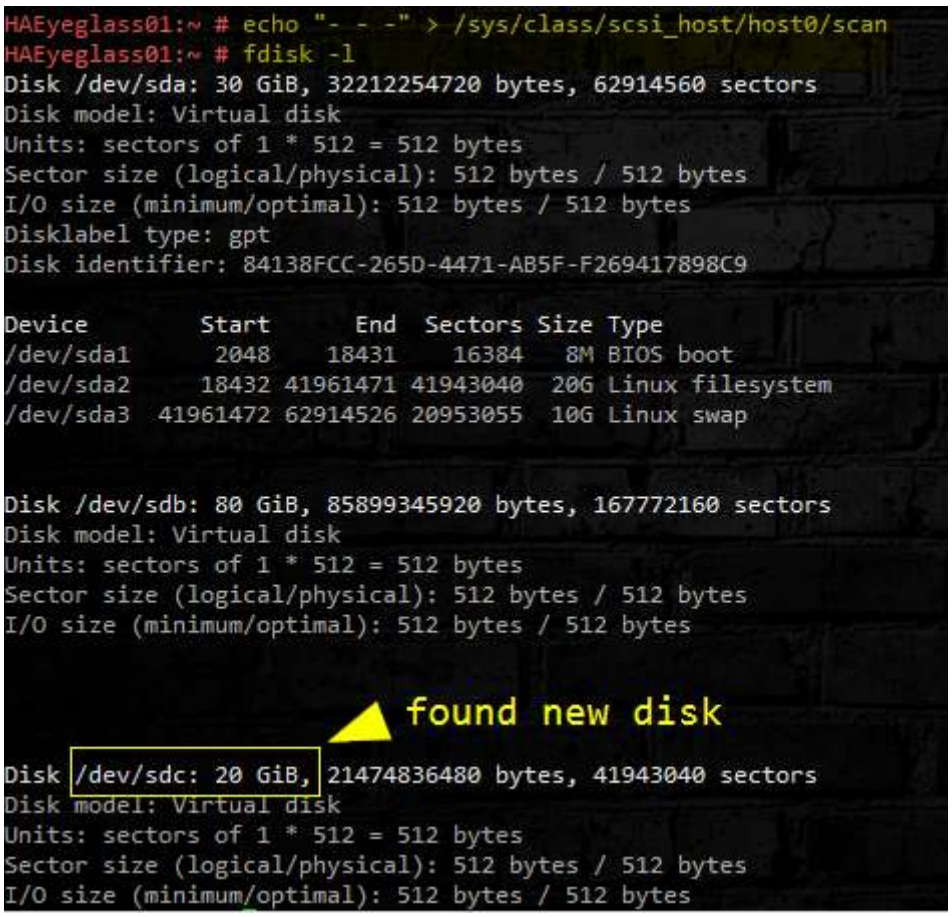

8. Add the newly added disk to the /opt/data btrfs filesystem. In our example it is `**/dev/sdc**` but if you have added disk before you will need to verify the device name from the fdisk -l output. **Run the command below to add the new disk to the existing file system.**
    a. **btrfs device add -f /dev/sdc /opt/data**
9. **At this point the metadata is only stored on the first disk, to distribute (balance) it across the devices run the following command:**
    a. **btrfs filesystem balance /opt/data &**
    b. **NOTE: command will run in background [DO NOT REBOOT]**
10. **Check if the File System balancing job is running with command below.**

a. **jobs**

11. **Check for the newly added disk. Look for `/opt/data` disk size:**

1. **df -h**
2. **Verify the new disk has expanded the capacity of the** /opt/data file system.
3.

```
HAEyeglass01:~ # df -h
Filesystem      Size  Used Avail Use% Mounted on
devtmpfs        7.9G  4.0K  7.9G   1% /dev
tmpfs           7.9G  166M  7.7G   3% /run
tmpfs           7.9G     0  7.9G   0% /sys/fs/cgroup
/dev/sda2        40G  4.4G   35G  11% /
/dev/sdb         90G  763M   87G   1% /opt/data
/dev/sda2        40G  4.4G   35G  11% /var
```

Repeat the steps above on each VM 2 to N, where N is the last VM in the Search & Recover cluster.

Done.

How to change the TLS Certificate for the WebUI with a new Self Signed Cert

1. Access the WebUI from node 1 and create a DNS entry for node to create a FQDN to create a signed cert.
2. The objective is to install the signed cert for nginx ECA Node-1
3. Create A record in DNS name for ECA Node-1 and verify with nslookup.. Example search.domain.com
4. SSH to ECA Node-1 as ecaadmin
5. cd /opt/superna/eca/conf/nginx
6. Verify that the nginx.key is there with ls -la

7. Create csr with that key file. Then create self-signed cert to replace nginx cert.
    a. cp -p nginx.crt nginx.cert.bak
    b. openssl req -key nginx.key -new -out nginx.csr
    c. openssl x509 -req -sha256 -days 365 -in nginx.csr -signkey nginx.key -out nginx.crt
8. When it is asked about the Common Name: provide the fqdn of ECA Node-1 (the name registered in DNS e.g. search.domain.com)
9. Restart nginx
    a. echo y | ecactl cluster push-config
    b. ecactl containers stop nginx
    c. ecactl containers rm -f nginx
    d. ecactl containers start nginx
10. Verify the certificate when accessing the UI (e.g. https://FQDN)

Content Indexing Performance Advanced Configuration Settings

The following settings require a cluster restart to take effect. **These should not be changed unless confirmed by support.**

These variables should be set on node 1 of the cluster by editing /opt/superna/eca/eca-env-common.conf .

1. export INGESTION_WORKER_PARALLEL_LIMIT=5
    a. **Number of threads to submit documents per ingestionworker container, higher number increases concurrent work processing but will also increase CPU utilization of the indexing container.**
2. export INGESTION_WORKER_MAX_FILE_SIZE_MB=100
    a. **Sets the max file size in MB to index, files over this limit will be skipped during indexing, larger file sizes will consume more cpu utilization and reduce indexing performance.**
3. export INGESTION_WORKER_MIN_FILE_SIZE_MB=0
    a. **Sets the min file size to index default is all files regardless of size, this may consume too much size in the index with many small files that have no content.**
4. export INGESTION_WORKER_FULL_PRIORITY=10
    a. Default ingestionworker container will round robin between indexing incremental changes to the file system and full scan indexing. Increasing the priority number for full will change the ratio to prefer more indexing of full scan files versus changed files.
5. export INGESTION_WORKER_INCREMENTAL_PRIORITY=10
    a. See above. Changing this to a higher value than the full indexing priority will mean more incremental file changes are processed versus full scan files.
6. export INVENTORY_TASK_CRON="* * * * *"
    a. The inventory task collects users share permissions and cluster inventory needed to index.

7. export INCREMENTAL_INDEX_TASK_CRON="* * * * *"
    a. This determines how often file system paths are monitored for changes for incremental re-indexing .

Scale Out Content Indexing & Classification Only VM Deployment Option

1. **Overview:** This option can expand a Search & Recover clusters content ingestion workers with a small foot print VM using only 4G RAM and a small disk without an index to focus entirely on content ingestion analysis.
2. Additional VM's can be added one by one to focus only content indexing tasks. This will also be used in feature 1.1.5 content classification feature that does not store the indexed document in the index saving space on disk.
3. This requires deployment of the Search & Recover single VM content only OVF, configuration by support to join this VM to your existing cluster.
4. Download the Search & Recover zip file and deploy the OVF for content only (download instructions for all products https://www.supernaeyeglass.com/downloads
5. Once the VM is deployed with an IP address and all firewall rules are in place. Contact Support to join this to your existing cluster.

© Superna LLC

9.13. User Search Guide

Home Top

- How to Login and Search for Documents
- How to Recover old versions of Files (Release 1.1.2)
- How to Use the Advanced Search Options
- How to Use Dynamic Document Tagging
  - Use Cases for this feature

How to Login and Search for Documents

How Login and Search Results are managed by PowerScale Security on SMB Shares

Users and administrators must login with Active Directory user id, using DOMAIN\userid or user@domain.com syntax. The authentication request is proxied through the PowerScale to Active Directory to validate the password. The users AD groups are employed to identify all SMB shares that the user has read or write permission. This process of SMB share validation will be limited to PowerScale clusters that are have at least 1 indexed path.

NOTE: All Searches will filter results to return ONLY file or directories that fall at or below SMB shares the user has AD permissions to see based on PowerScale SMB security settings.

1. At the login screen enter Active Directory user id and password (**NOTE: syntax requires upper case domain name DOMAIN\userid or userid@example.com**)

2. 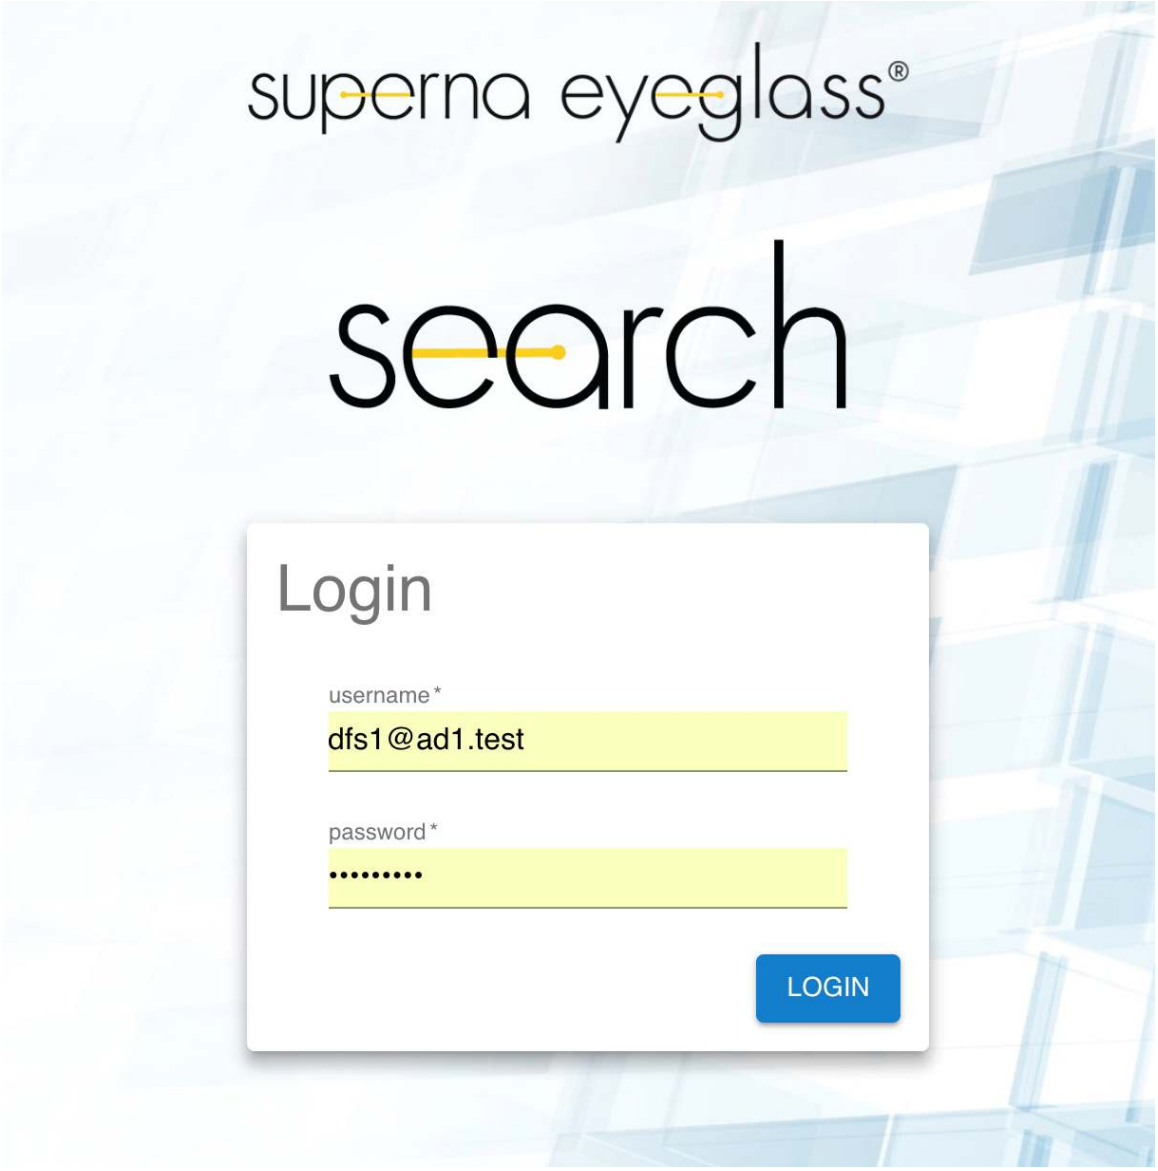

3. Enter a basic search by entering key words into the search box.

4. Note: by default the key words will search file names, directory names and contents of files. (**NOTE only folders that are configured for content indexing will return results based on the contents of the file**).

5.

6. You can copy and paste the location of a file in the search results using right click mouse action.

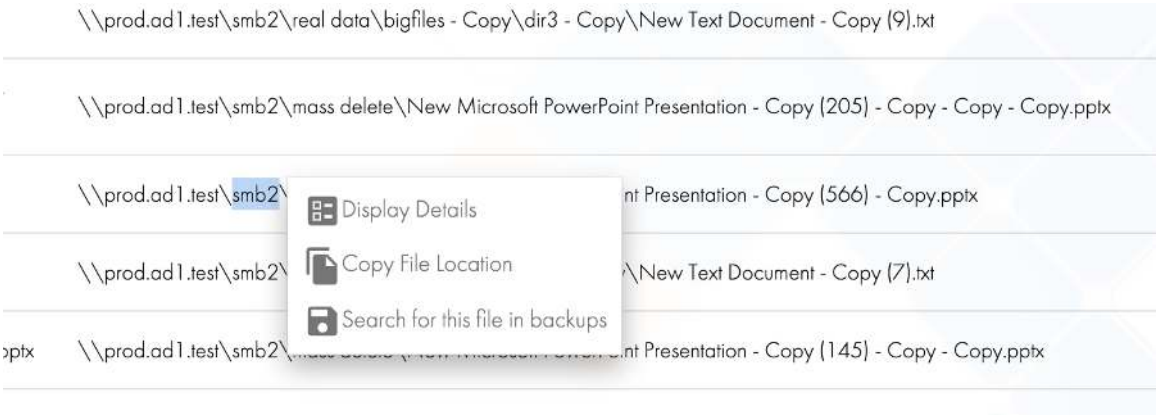

7.

8. Then paste into Microsoft Explorer to open the file from the search results.

9. You can also download the search results into an Excel file or CSV file format. Click the cloud download button, select the file type download option and use "Download All" to get all files added to the report, or enter the number of rows to add to the results. **NOTE: Limit on the browser download is 100 000 files from search results.**

10.

11. When done searching click logout button.

How to Recover old versions of Files (Release 1.1.2)

If the administrator has enabled snapshot monitor mode , the file versions can be displayed and used to recover files protected by the storage device. This allows old versions of file versions, and file location copy and paste, to access old file versions.

Steps to see file versions (if available)

1. Complete a search.
2. Right Click on any file in the search results to search for existing file versions in snapshots that protect and store multiple versions of a file. The Option shown below "Search for this file in backups".
3. 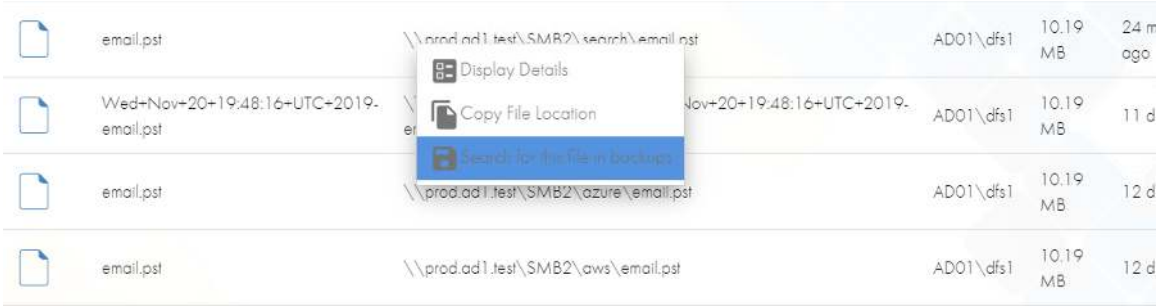

4. The Search results will list all versions of that file that exist in the backups. You can see the age of the file in the last modified column. You can select a version of the file to restore to the file system by Right Clicking the file version to display the "Recover File" option. See Screenshot below:

5. 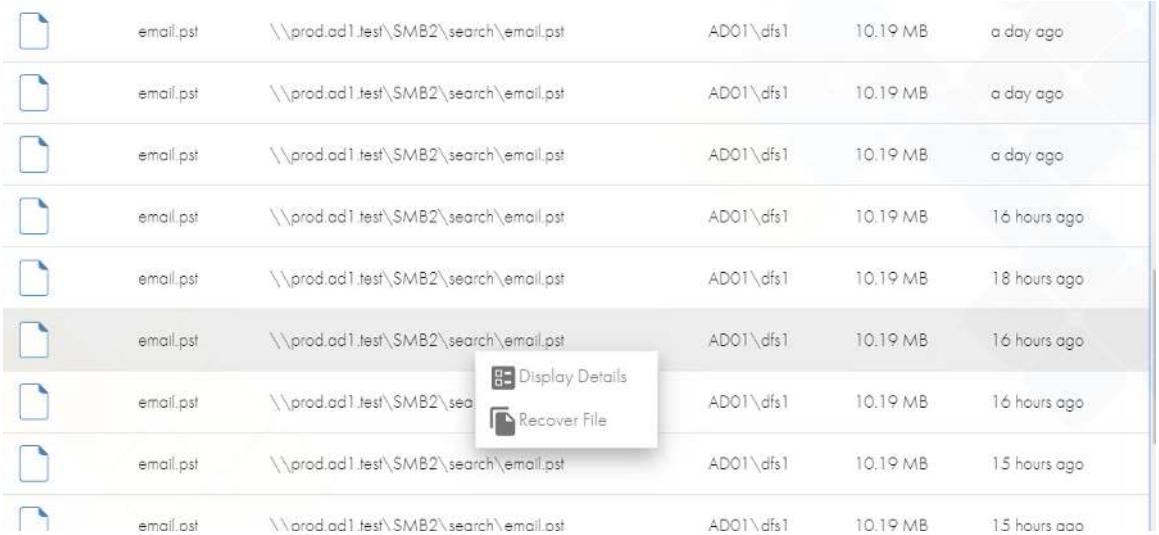

6. Select the "Recover File" option to display information about when the file was backed up and where this file will be restored.

7. 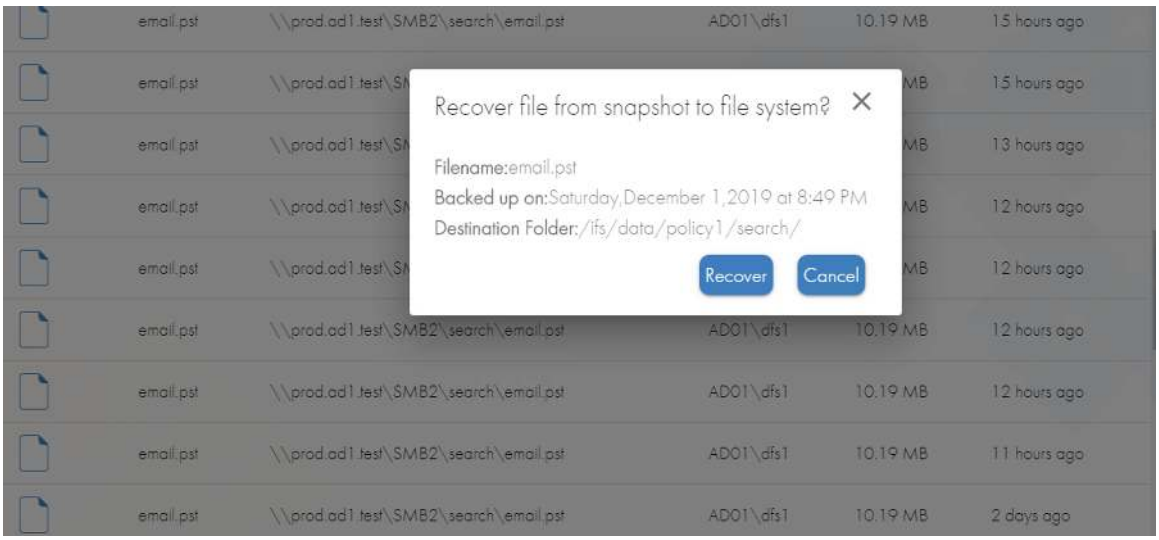

8. Select the "Recover" option to restore the file or "Cancel" to select another file version.

9. The restore progress is shown

10. 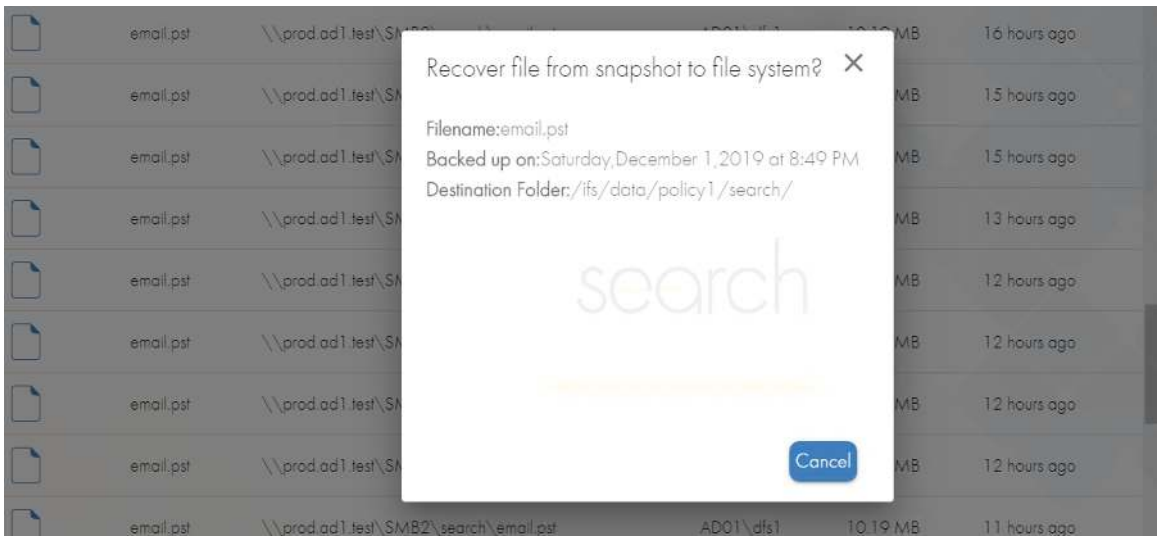

11. The results of the restore operation is displayed with the following information and options:

    a. Summary of the file name, backup date and location where the file will be restored.

    b. Success of the recover operation .

    c. A Copy to Clipboard option for the restored file, and file name that identifies the file name as igls-restored -date of the file version and original file name .

    d. A Copy to Clipboard option for the original file name in the file system that was overwritten with the restored version.

    e. A Copy to Clipboard option for the folder that contains the restored files.

f. 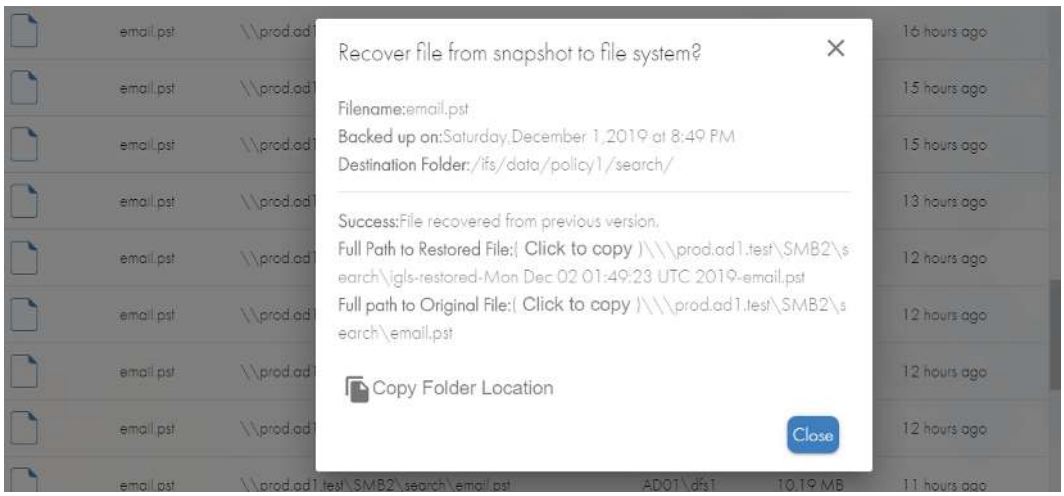

12. If the file is not available in the backup snapshot location an error will be returned. See example below:

a. 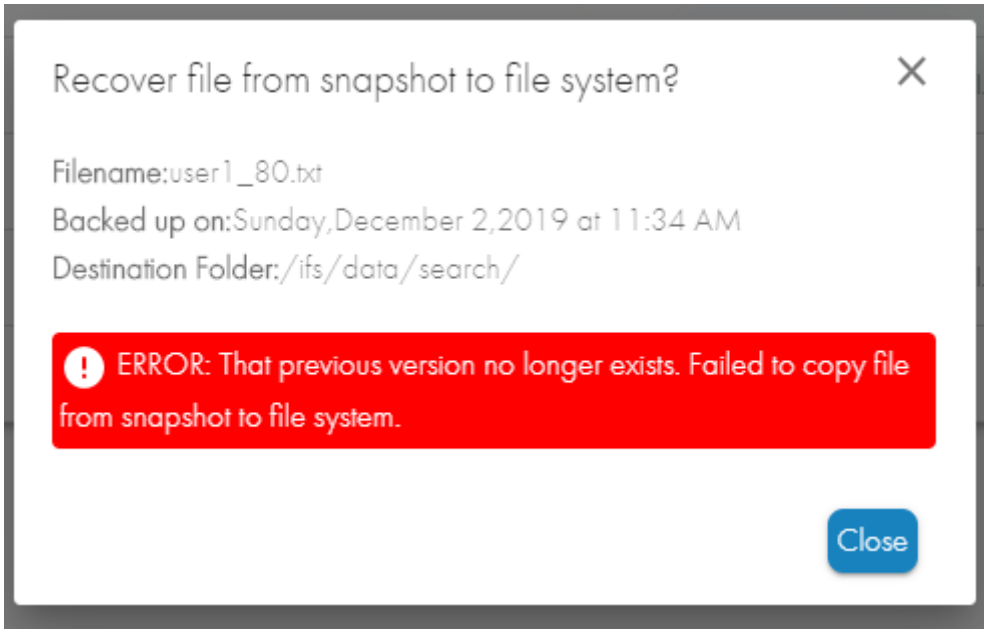

13. Copy to Clipboard a path to a file or folder to paste into Windows Explorer to open the file or folder. Click Close to return to the backup search results.

14. NOTE: The administrator has configured restore overwrite settings and protecting the original file during restore operations.

How to Use the Advanced Search Options

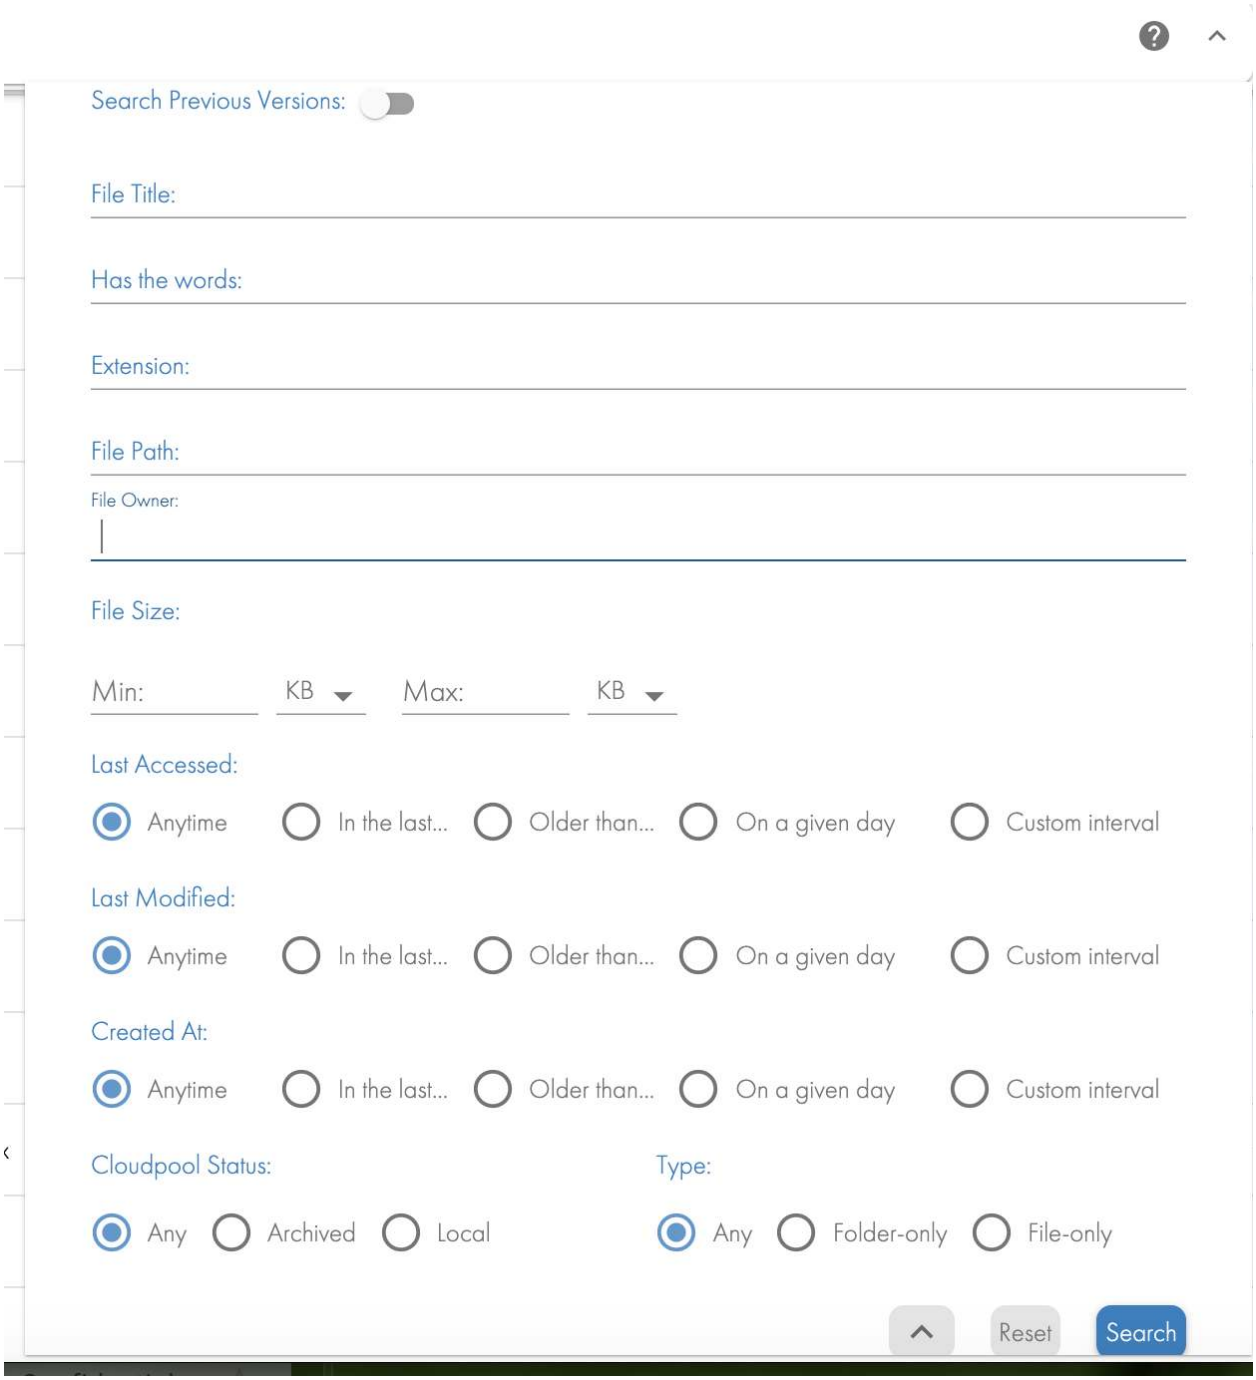

1. File title: only searches file names
2. Has the Words: A content search of words and phrases in supported document types for content indexing. Your administrator must enable content indexing first.

3. File Extension: the extension of the file example pdf . You can provide multiple extension types using a space character between extensions (Example: **docx pdf**). **NOTE: The period or dot is not required.**
4. File Owner: This is the owner of the files. NOTE: the domain must be entered upper case and the slash must be escaped as follows. AD01\\username (user id is case sensitive).
5. File Size: Min and Max range of file sizes.
6. Last Accessed, Last Modified, Created Date stamps on files:
    a. default is anytime.
    b. In the last hours , days, months , years (to search using a simple range based on today's date).
    c. Older than hours, days, months , years (to search using a simple range based on a day in the passed from today's date) .
    d. On a specific day.
    e. Custom - this is a start day and time and an end day and time.
7. Cloud pool Status: If PowerScale cloud pools feature is in use, this allows locating and reporting on stubbed or archived files. The Default is return results for any type of file or change to search only the file system or archived data.
8. Type: This defaults to searching for files and folders names but you can specify folders only or files only to narrow your search results.

How to Use Dynamic Document Tagging

This feature uses content aware indexing to group and locate documents based on tags within documents.

Use Cases for this feature

1. Group project data based on project ID.
2. locate employee data by employee ID.
3. Allow users to tag documents with document ID's to locate with simply searches.
4. Create document ID's and leverage wild cards * ? to easily build file lists from multiple projects.
5. Classify documents for security with confidential, or other security classifiers.

Watch the video showing examples of how to use Dynamic project tags:

© Superna LLC

9.14. Administrator Search Guide

Home Top

- Overview
- Security and Administrator Searches
- Download Search Results to CSV or Script downloads
  - Limitations on GUI Download of CSV
  - How to Download Large Search or cmdwriter Files

Overview

This section covers administrator specific use cases. The user search guide UI is the same for administrators, and all advanced search features are the same with the considerations below.

Security and Administrator Searches

Administrators have the same security requirements for results. This means smb share permissions will filter results based on access of the AD account logged in. To see more of the file system, the administrator will need to create shares to grant access for searching data for e-Discovery or file system automation. If ACL search mode is enabled, this will block results if the administrator does not own the files which is always the case. See the configuration guide on how to enable Admin bypass to return all files even if they are not owned by the administrator.

How to search within a specific path and below for advanced searches

Administrators often need to reduce the scope of a search to a specific path in the file system. This option is only presented to users defined in the admin overrides list. See the configuration commands to add users to this list. Any user that logins on this list will see a path entry that allows a full path to be entered to bound the scope to this path or below.

1. Advanced Search + admin only Path entry field. This will be enhanced in future releases to allow file system browsing. Enter a path **without** the trailing / to indicate a path.

Example: /ifs/data/somepath (this is ok). ~~/ifs/data/something/~~ (this is not ok)

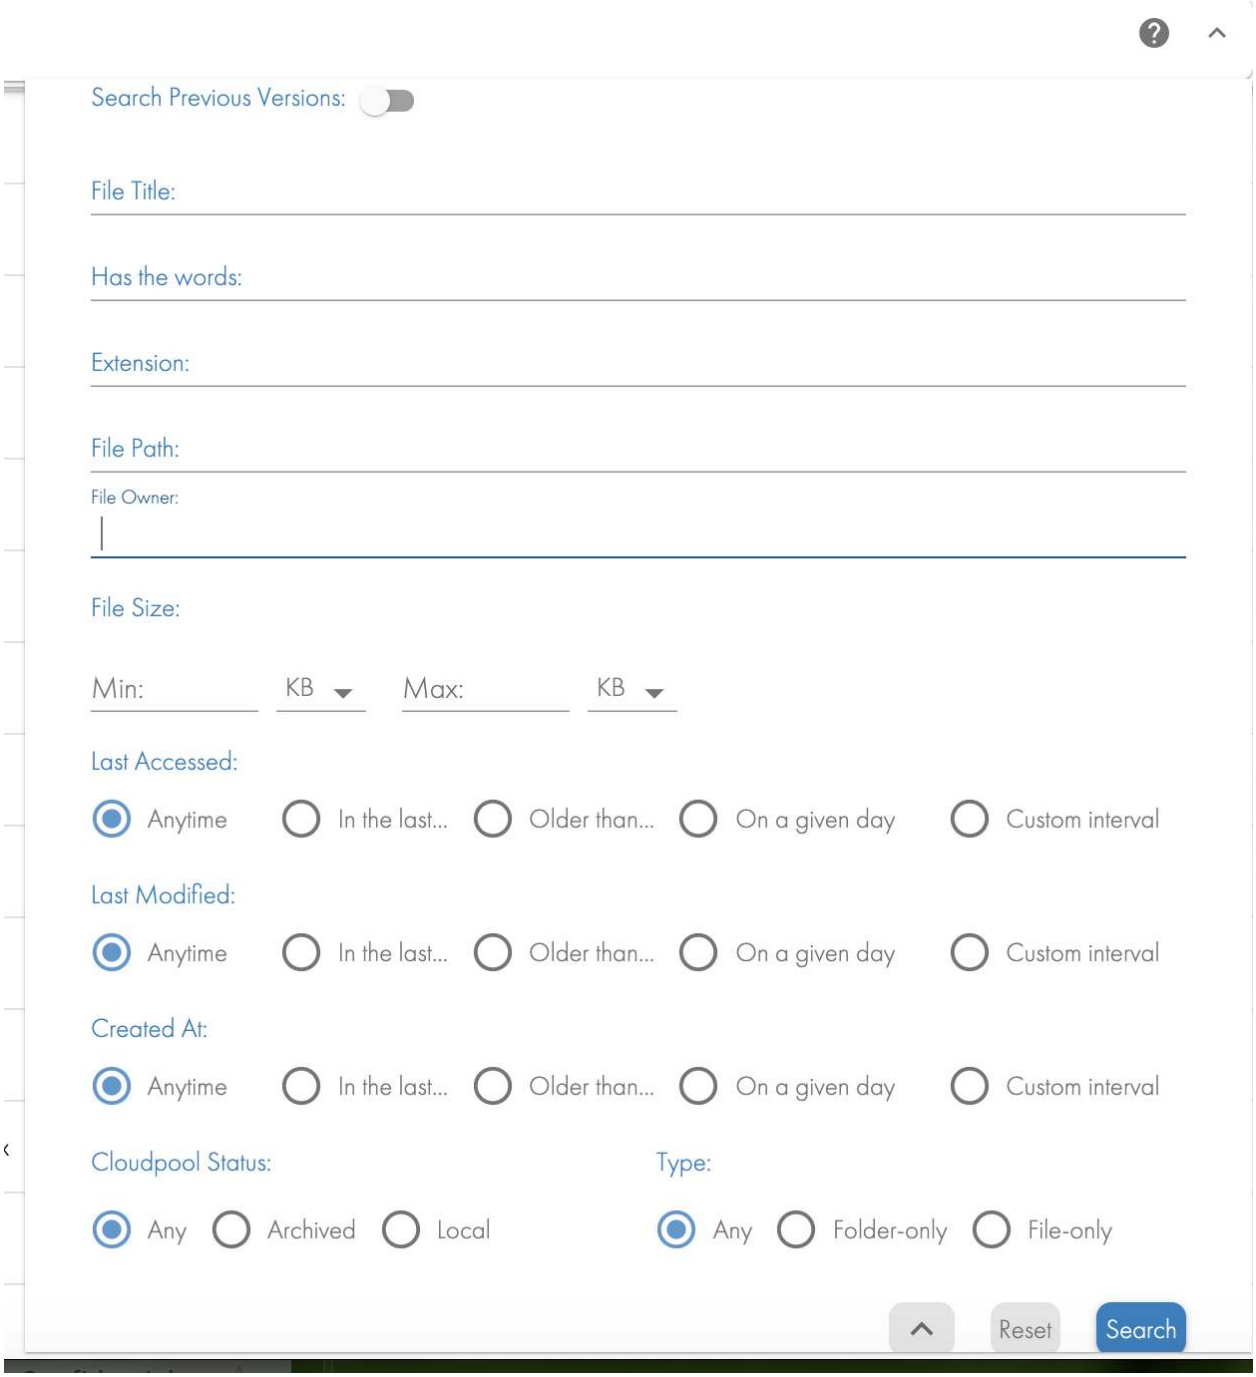

Download Search Results to CSV or Script downloads

See the user guide for how to download search results to a CSV in the GUI. Administrators have use cases to download to csv for several different use cases including data analytics, capacity management, security, file system action automation scripts. These use cases

require very large csv script downloads. The appliance backup location is also stored in the download url.

Limitations on GUI Download of CSV

1. csv download of 100,000 results or less can be done on the UI and directly downloaded to your pc.
2. csv files larger than 100 000 files will automatically be stored on the download link. See the instructions below to access downloads.

How to Download Large Search or cmdwriter Files

1. Using a browser and the webUI user and password. Default user is ecaadmin and the password is set during installation for the Webui and download page access.
    a. **Note: This page is secured for administrators and not intended for end users.**
2. https://x.x.x.x/downloads (enter the user id and password).
3. Under this directory you will find a **CSV** folder where all CSV search results will be stored for all users that requested large results save operations.
4. Under the **cmdwriter** folder any administrator saved scripts will be stored in this folder.

© Superna LLC

9.15. File System Analytics with Quick Reports - Guide

Home Top

- Overview
  - Common Tasks
  - How to Enter Paths to narrow results
  - How to Target Data with Advanced Options for Quick Reports
  - Who Owns what?
    - Use Cases
  - Show me the File types
    - Use Cases
  - What's growing old?
  - What's Been Archived?
  - Who used the space?
  - How to read and use the What's Growing Old? Reports
    - Use Cases
  - Data Classification Quick Report
  - What's in the File Pool?
  - When can I delete my worm lock data?
  - How Fast is My Data Growing?
  - How to drill into Quick Reports to identify the files behind the results
- How to Schedule Quick Reports and email results

- How to edit or delete Scheduled Reports
- File System Interactive Browser - Beta
  - Overview

Overview

Quick reports provide various summaries of file system data. Quick reports are end user aware which means the logged in user sees reports relevant to their security access to data. Each quick report has advanced options to narrow the report to a path, extension, file size and content of the files. Content based searches provides unique ability to focus reports on content not just date stamps.

Quick reports return a graph or chart to simplify data analysis, along with a table format of the results that can be downloaded to a CSV file.

Quick reports are more than just reports, they allow selecting a subset of the data in the report, and drilling into the results to identify the exact files from the table results. This means you can locate the exact data easily, and more importantly the script automation feature is available to automate some task.

Common Tasks

1. Automate a file system action with the file results from Quick Reports. See Script examples that cover many common tasks covered in the Search & Recover Solutions Guides chapter.

2. Learn more about the files in the quick reports to manage the file system. See how to use the Drill in Actions menu option on search results later in this guide.

How to Enter Paths to narrow results

1. The path input fields must be absolute paths .
2. The syntax is /ifs/data/somepath .

How to Target Data with Advanced Options for Quick Reports

All quick reports support options to target data in the file system. The options below can be combined to build a targeted report on any criteria including the contents of the files. See the options to target data with advanced options.

1. Path .
2. file size or file size range .
3. file extension .
4. Content .
5. Date Stamps - Supports a range, specific day or In the last Days options:
    a. Created Date stamp .
    b. Modified date stamp .
    c. Last Accessed (if enabled on the file system).
6. File Owner:

a. **Note use this to bound a search by a specific AD user. Syntax must be upper case domain name with double slash. Example: "AD01\\userx", and the user is case sensitive match.**

7. Date Search Options:

   a. **Any Time** - Includes all possible dates in the index and does not narrow the search results.

   b. **In the Last** - Allows selecting a time period that starts with today and includes dates stamps in the last x days, minutes, or seconds in the passed using today's date.

   c. **Older than** - This option allows a search that starts in the past using today's date. The choice is older than x weeks, months or years. Example: Older than the last 6 months, will start the search 6 months in the past starting from today's date and will locate files from the beginning of time or the oldest file time stamp

   d. **On a given Day** - Selects a single day with a calendar pop up to select a day, can be any day in the past. Click the date icon and use the arrow to change the month. Click on the year at the top of the window to select the year.

i. 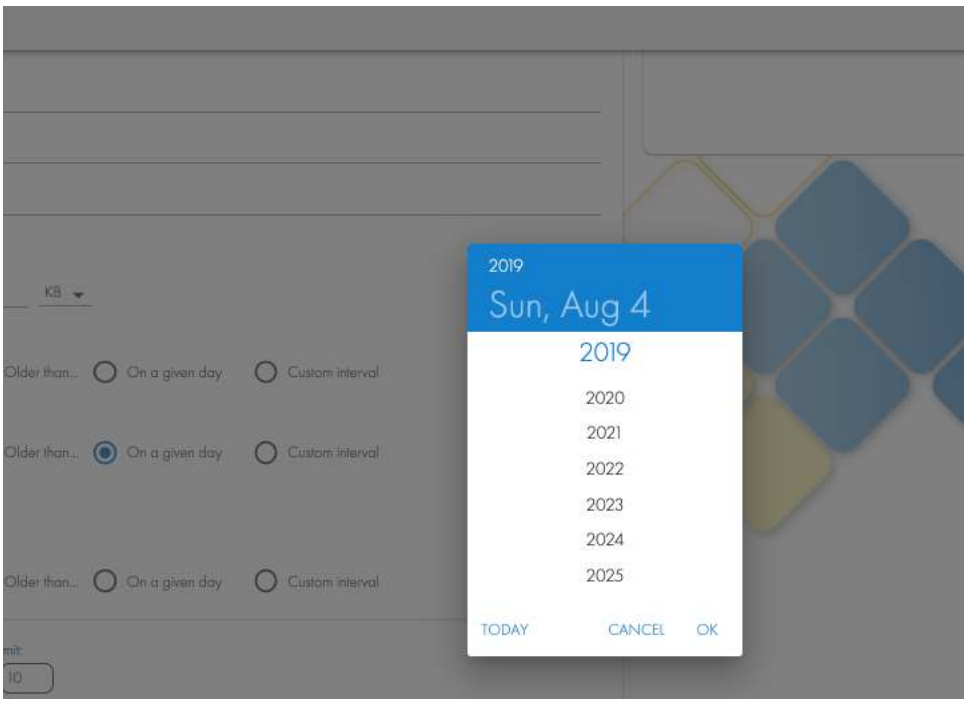

e. **Custom Interval** - Pick an exact date range with a start date and time and end date and time.

i. 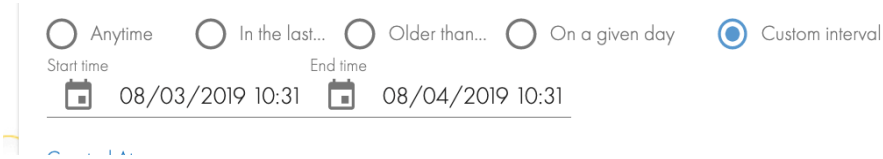

8. Sort by - returns the table data sorted by:
    a. Total file size sum by user.
    b. Total file count.
    c. Average file size.
    d. Largest file size.
9. Sort by Ascending or Descending - applies to the sort by option, not available on all reports.
10. Limit is the number or rows of data to return to the UI results screen:
    a. If a user report it will only return a sorted top 10 users.

b. if an extension report it will only return a sorted 10 file extensions.

Who Owns what?

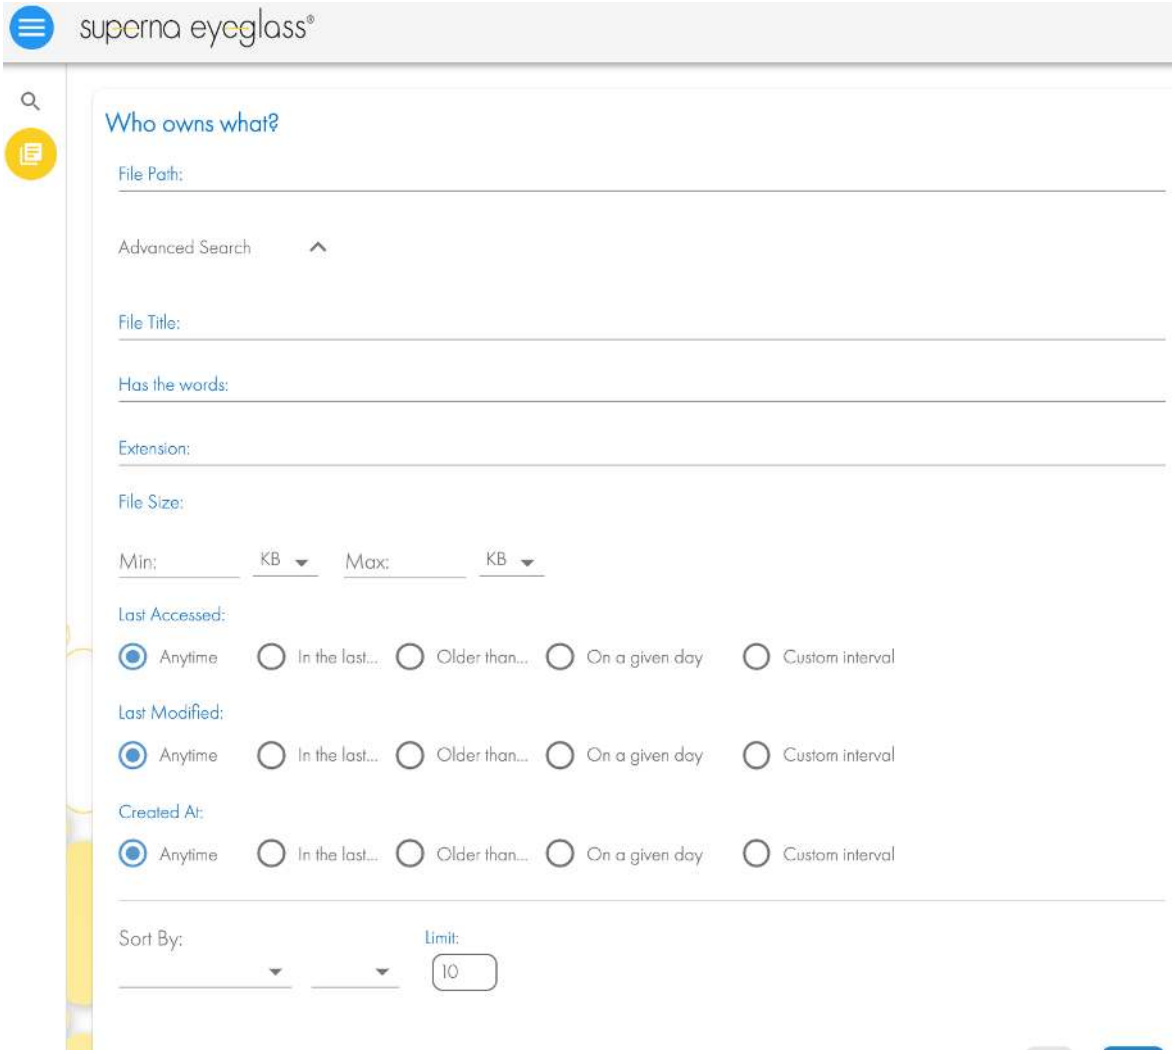

This report summarizes files the ownership of the files and sums the file sizes by user. The report will sum the file size, file count and average file size by user id, using the if ownership attribute of the files. This report can be narrowed by path to identify owner summary for a specific path, and all data under this path. **NOTE: the default only**

returns the top 10 results sorted highest to lowest. See limit advanced option to get more results returned.

Use Cases

1. Identify **project data owners** at a path and below, can be achieved by entering a path to the search
2. **Dynamic quota** - to report on usage by month, path or content using date range filters for a user
    a. Example: Add a path where the search should start, the results will sum all users that own files at or under this path. Find the user name in the results, and download the CSV report.
3. **Disk space Rapid Growth To determine "Who"** created data within a specific time period. Example: A month where usage on a path spiked unexpectedly.
    a. Enter the path.
    b. Enter a custom interval for Created time stamp and use the first day of the month and the ending date of the month.
    c. This will return all files created in that month summing file sizes by user . This will identify who is responsible for the data growth on that path for that month.

Show me the File types

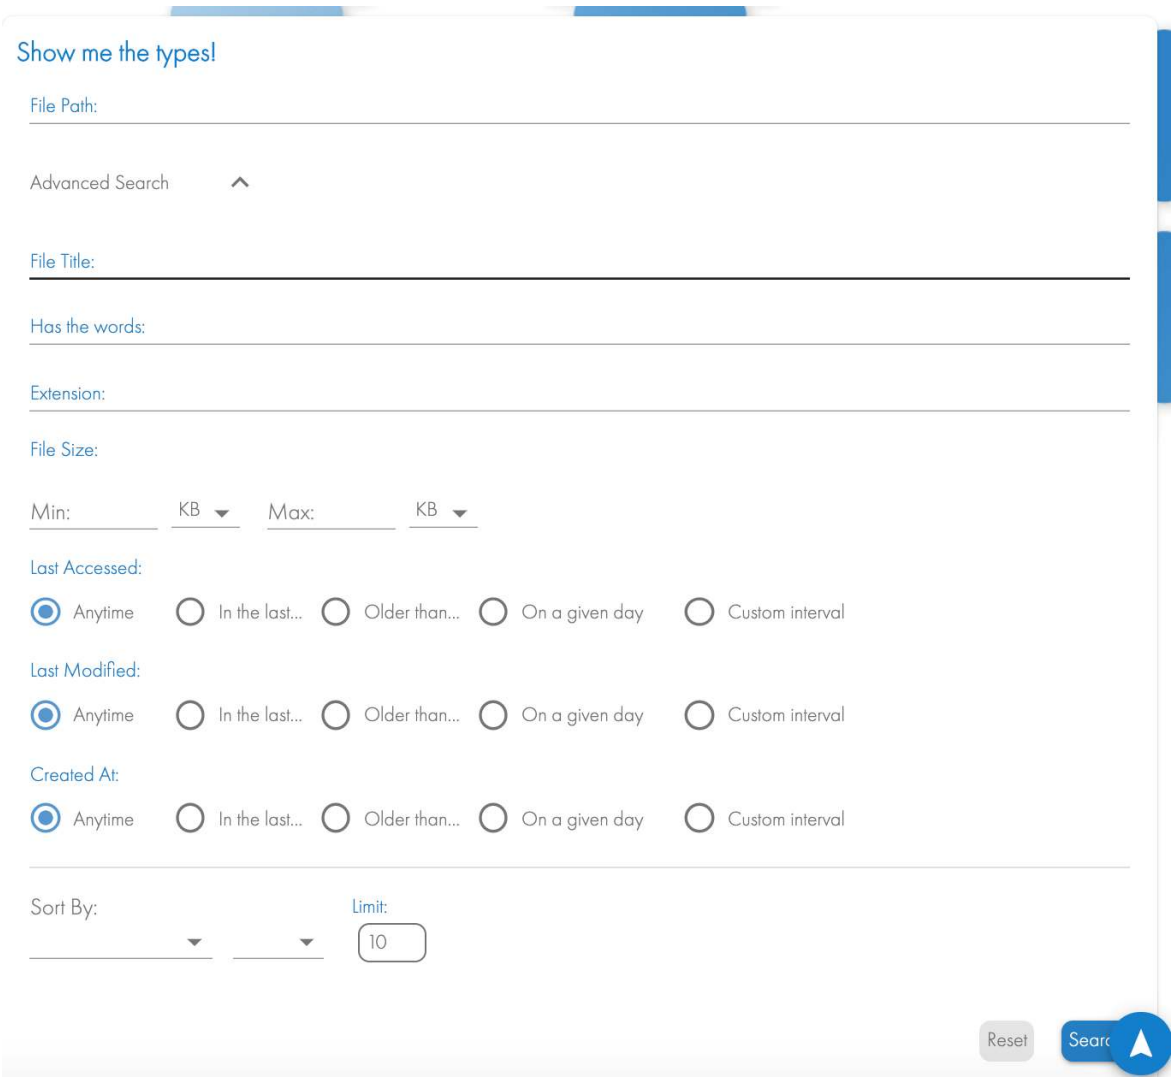

Use Cases

1. **Data Classification** - This report can help classify data to find file types that might be candidates for deletion. Example: .tmp, mp3, .avi or identify unknown file types on the cluster. This can be narrowed down by path, dates as mentioned above to focus the results to a specific location.
2. **File pool policy Design** - The file types and quantity of data by type can allow file pool policies to place low IO file types onto slower media. Examples may include pdf's that are commonly copies of other files in pdf format.

3. **Acceptable Use Security Report -** Some file types might be banned. For example avi, mpeg, mp3 or mp4 are examples of files that should not be stored on corporate storage.

What's growing old?

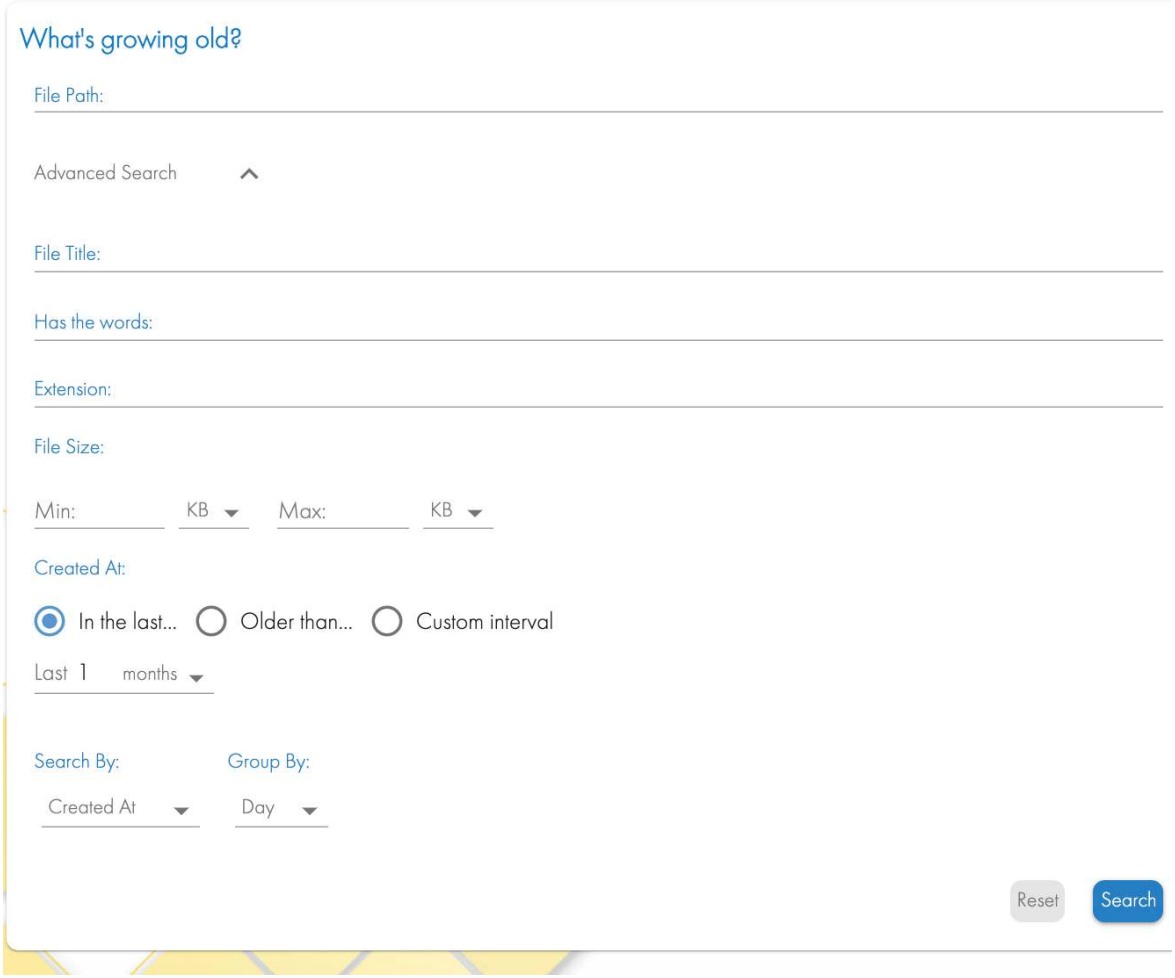

This report will assist with locating data based on age of the file or access time stamps (if enabled on the cluster). This will help identify data for archive or deletion based on age. This report can use creation date, modified date, or last accessed date stamps to locate data. In addition the advanced options narrow the search by path, content, and file size. The report can identify data in the last X days, month, years, or custom date ranges. The key search criteria is data older than x months. This will simplify locating data.

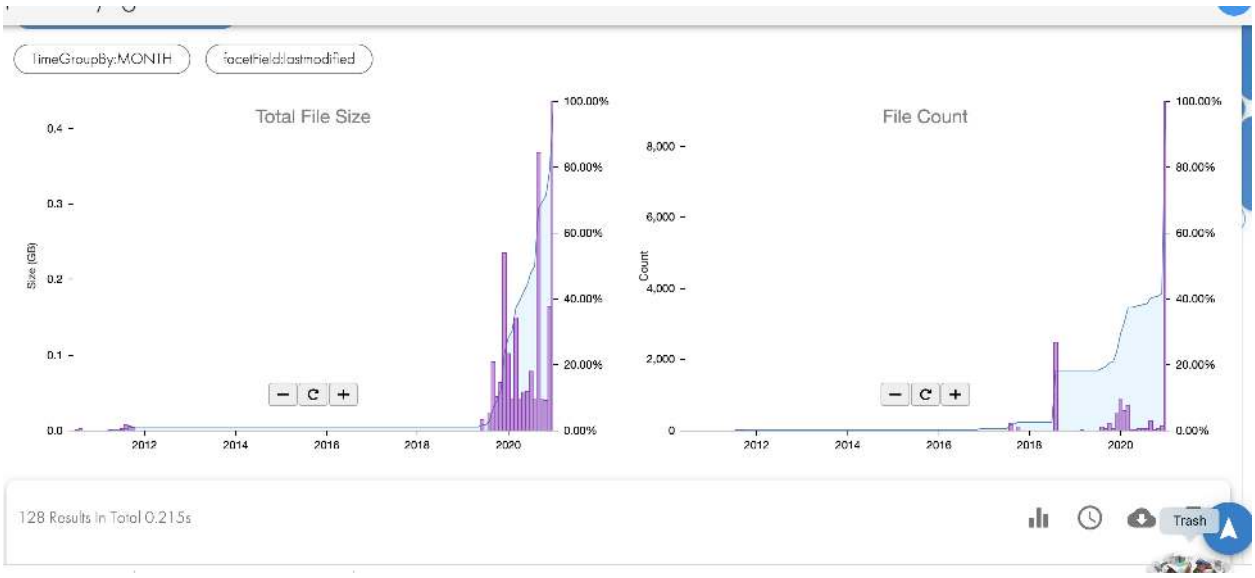

What's Been Archived?

Search & Recover is PowerScale cloudpool aware and indexes the stub status during indexing. This allows a report that shows users what % of the data they have access to is archived as a stub, versus in the file system. This report also supports advanced options to narrow the results down to a path. The screenshot below shows the advanced options that can be used to narrow the results including content aware.

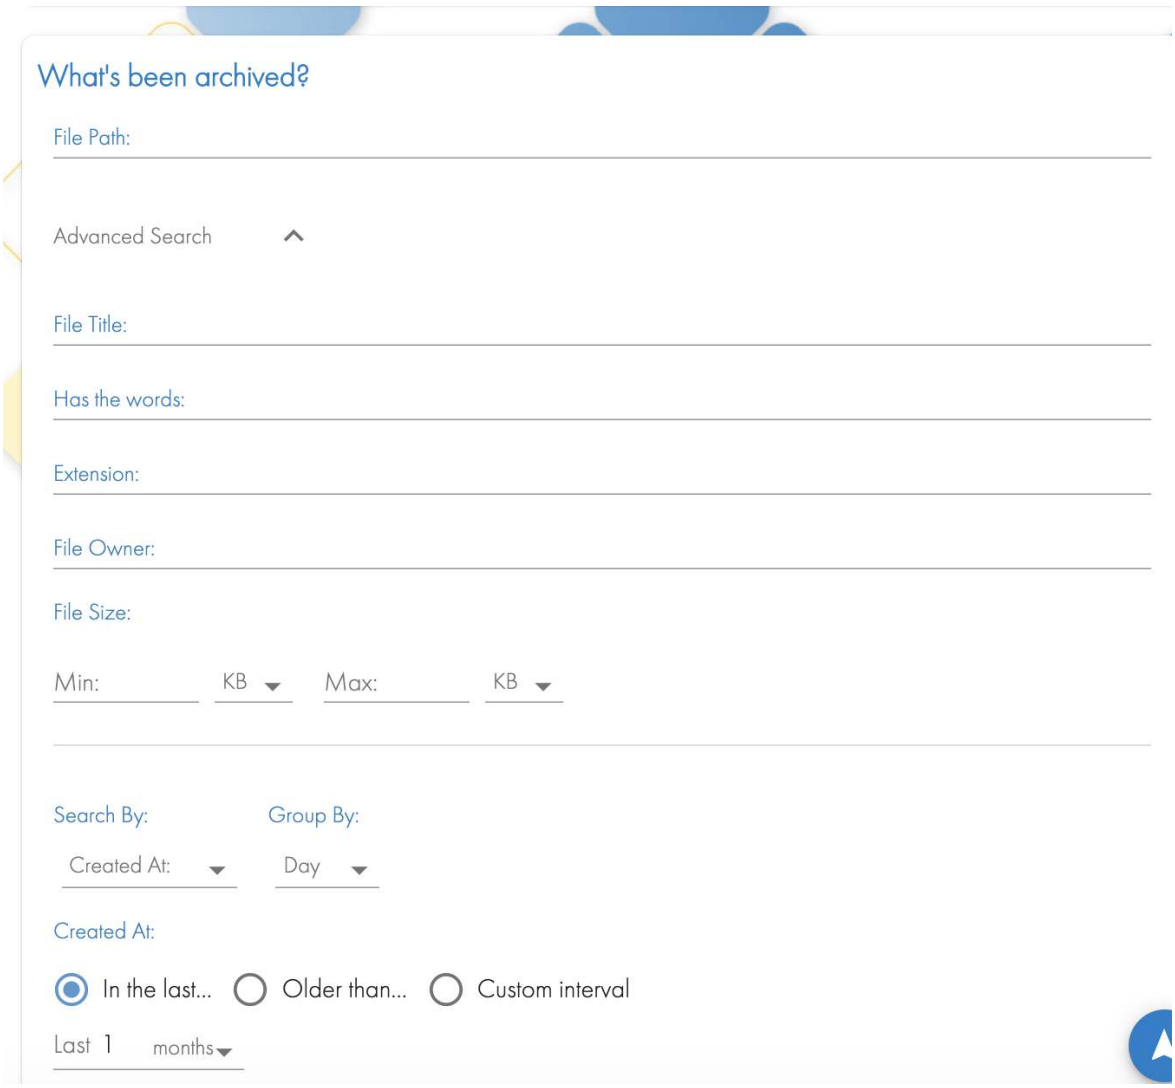

The image below shows results that indicates the % of archived versus file system data. The table below the graph shows the date range where data was found based on the file date stamp used to run the search, each column indicates file summary stats for that day or month depending on the report settings. Use the option to drill in to the results as needed using the icon on the report menu bar.

Who used the space?

A common issue is identifying repaid disk usage on a path in the file system, and identify who is responsible for this growth. Quotas can indicate the used space but not who is responsible for the growth over time.

This can be done using the What's Growing Old? Quick report and setup with a schedule in release 1.1.5.

Example below:

1. Identify who used space in a path over the last week.
2. Enter the path to the top level folder you want to analyze .
3. Use the Search by "Created At" date stamp and group by Day.
4. Created At section change to **In the last** to weeks and enter 1 week.
5. See screenshot below.

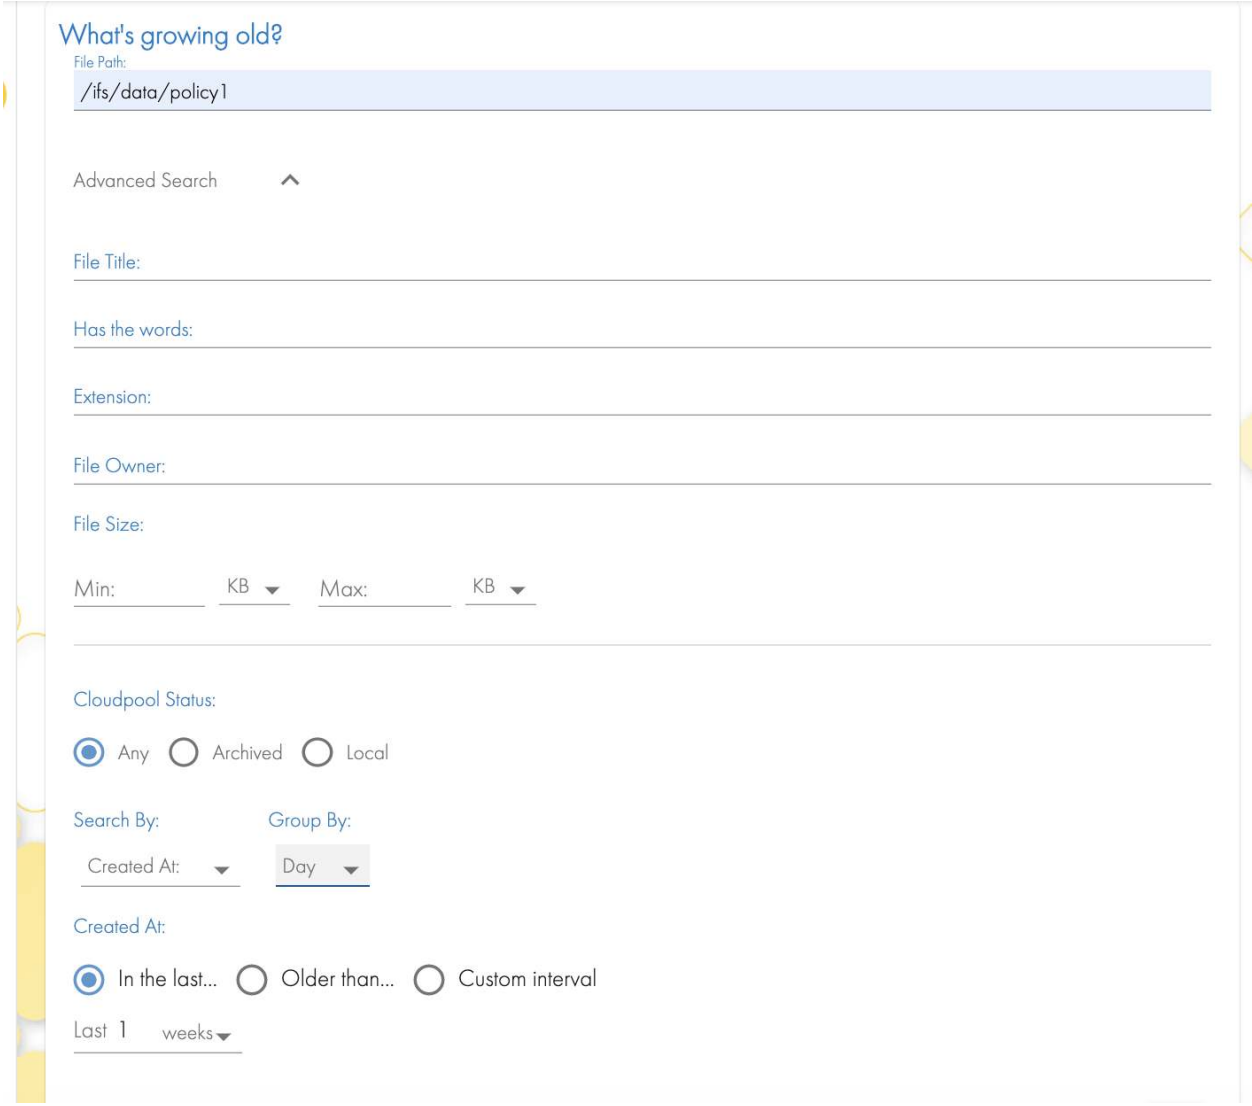

The search results shows the data created by day in the table and graphs.

Select all days in the table (check box at the top of the table) or select individual days that show a lot of data was created.

Use the scheduled report feature to schedule this quick report.

In the example below select the day with the most data created with the check box then select the drillin **"Who Owns that?"**

A pop up window will appear to further select what data you want returned.

Select "Total file Size" to get the sum of all data owned by users on this day. Set Ascending or Descending sort and the number of users to return in the search.

The results below shows the user responsible for the data created in the last week:

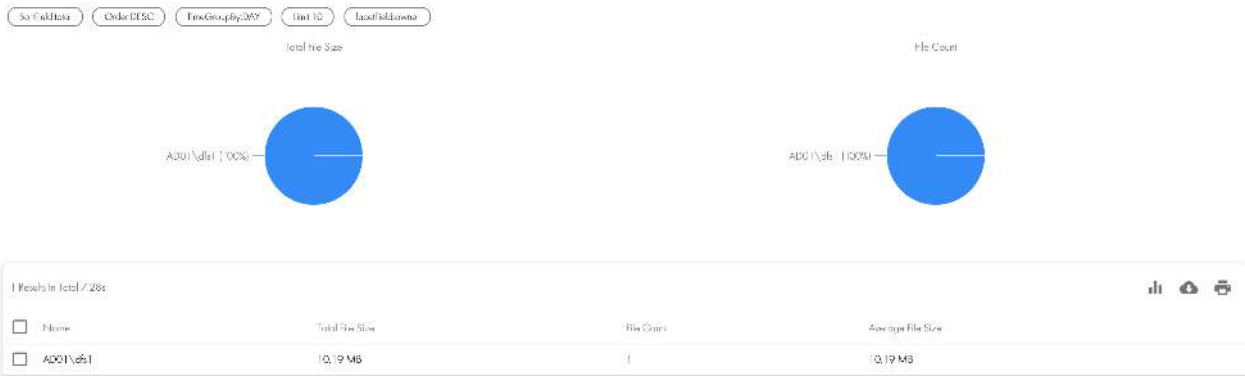

How to read and use the What's Growing Old? Reports

Mouse over the left graph and review the Context Box , that shows the following:

1. The File Size percentage value shows the percent of data this file date represents of the Total sum of all data returned by the search criteria. For example in the graph below the Context Box shows that 38.96% of the data was May of 2005 and older.
    a. As the mouse moves to the right the % will increase until you get to 100% of the data in the search results.
2. **Total File Size** shows the sum of all the files in GB that has a date stamp equal to or older than the date shown in the Context Box. This allows you to see how much data could be archived or deleted and what % this represents.
3. **Days Ago** - indicates how many days in the past using today's date to calculate this age.

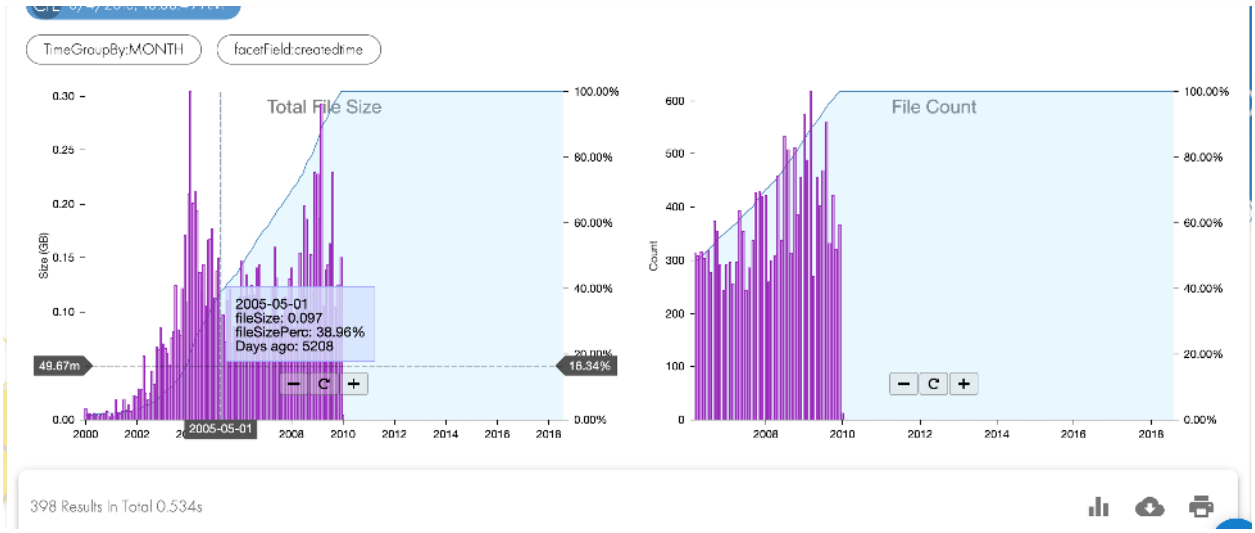

Use Cases

1. Locate data for **Archive, Deletion.**
    a. Identify storage space consumed by age of files (using created, modified or last accessed time stamps).
2. **File System Analytics -** Manage by path by age for departments or project data.
    a. File policy design can use the results to move files to archive based on age.
    b. Identify the quantity of data that will be moved to archive is easy. To identify the GB that will be moved to archive, mouse over the graph to determine the sum of the data in GB on a given vertical graph.
    c. Use the % option to determine your objective for data reduction.
    d. Use the Date of the final selected % for archive or delete.
    e. Edit the search using the custom interval option and a start date many years in the past (i.e. 1970 January 1st to locate all files). Enter the end date from Step #4 above.
    f. When the results are returned you can now use the "Show my the files" option to get a full list of all the files location in the files system and download a csv, or use the script option to assist with delete, archive or move to staging area before deleting. See next section.

Data Classification Quick Report

1. Once data classification feature is configured and tagging files have been indexed the report will show distribution of document

classifications configured for tagging. The classification report simplifies reporting on your classified data.

2. Requirements:
    a. Release 1.1.5 or later
    b. Configure Classification tags following the guide.
3. Select the classification box in any advanced search interface to select the tags you have configured for your data and use the and OR option to combine multiple classification tags to the report using AND OR logic for the search. The Advanced option allows additional options to narrow your search to a specific path or use the date range for the search.
4. 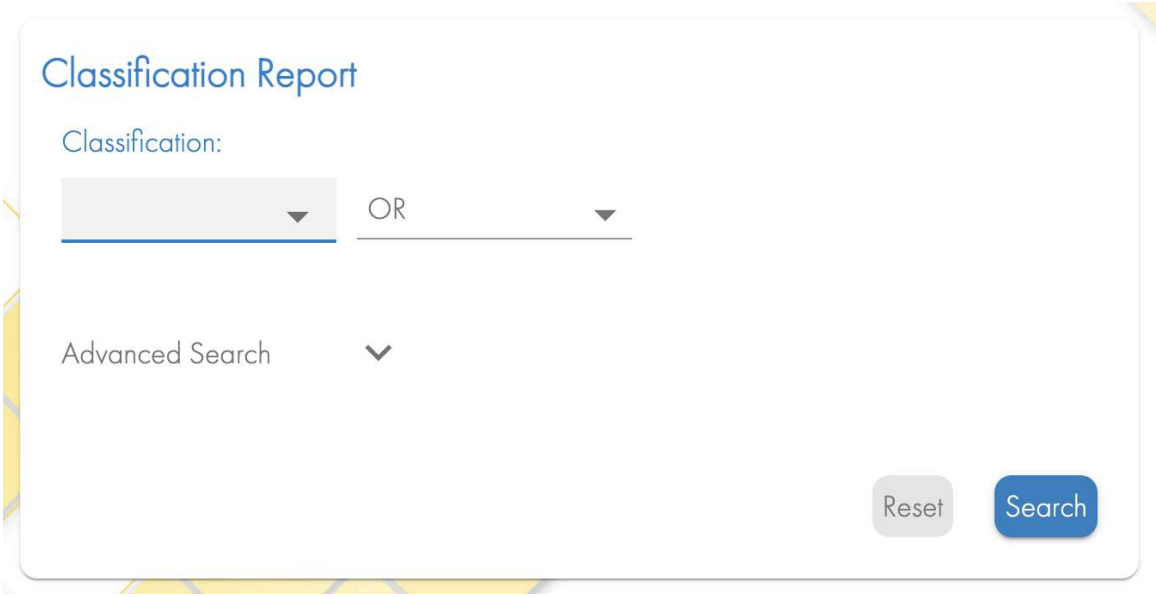

5. Using the Quick report to summarize all your classification tags by data quantity by tag or file count

a. 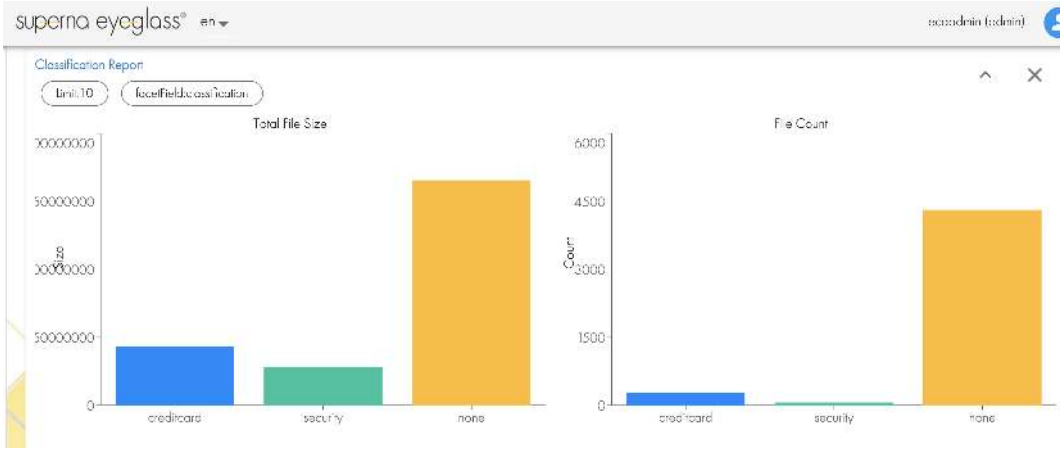

b. 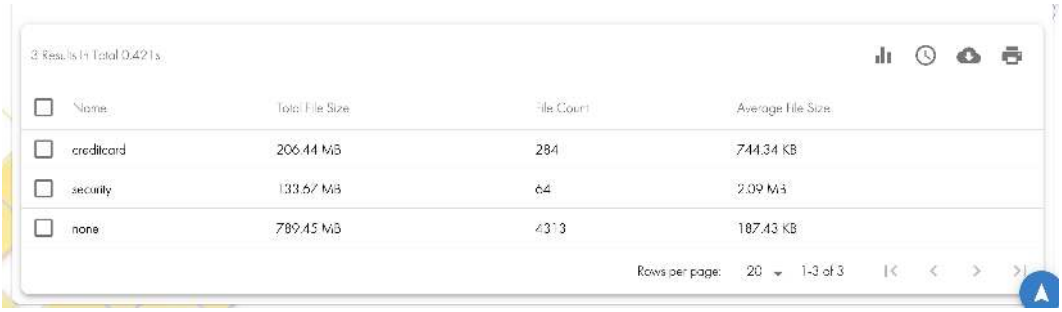

What's in the File Pool?

1. This quick report allows reporting on which files exist in a file pool example archive tier, fast tier or a cloudpool tier. This is typically needed for show back and charge back reporting.
2. Requirements: Release 1.1.5 or later
3. This quick report will report on 1 or more pools that are auto detected from the cluster and selected in the advanced area of the report UI.
4. See the example below. If you select multiple pools the repot area will show the sum of all data in each pool. you can also drill in to a pool to locate the files within that pool.

5. 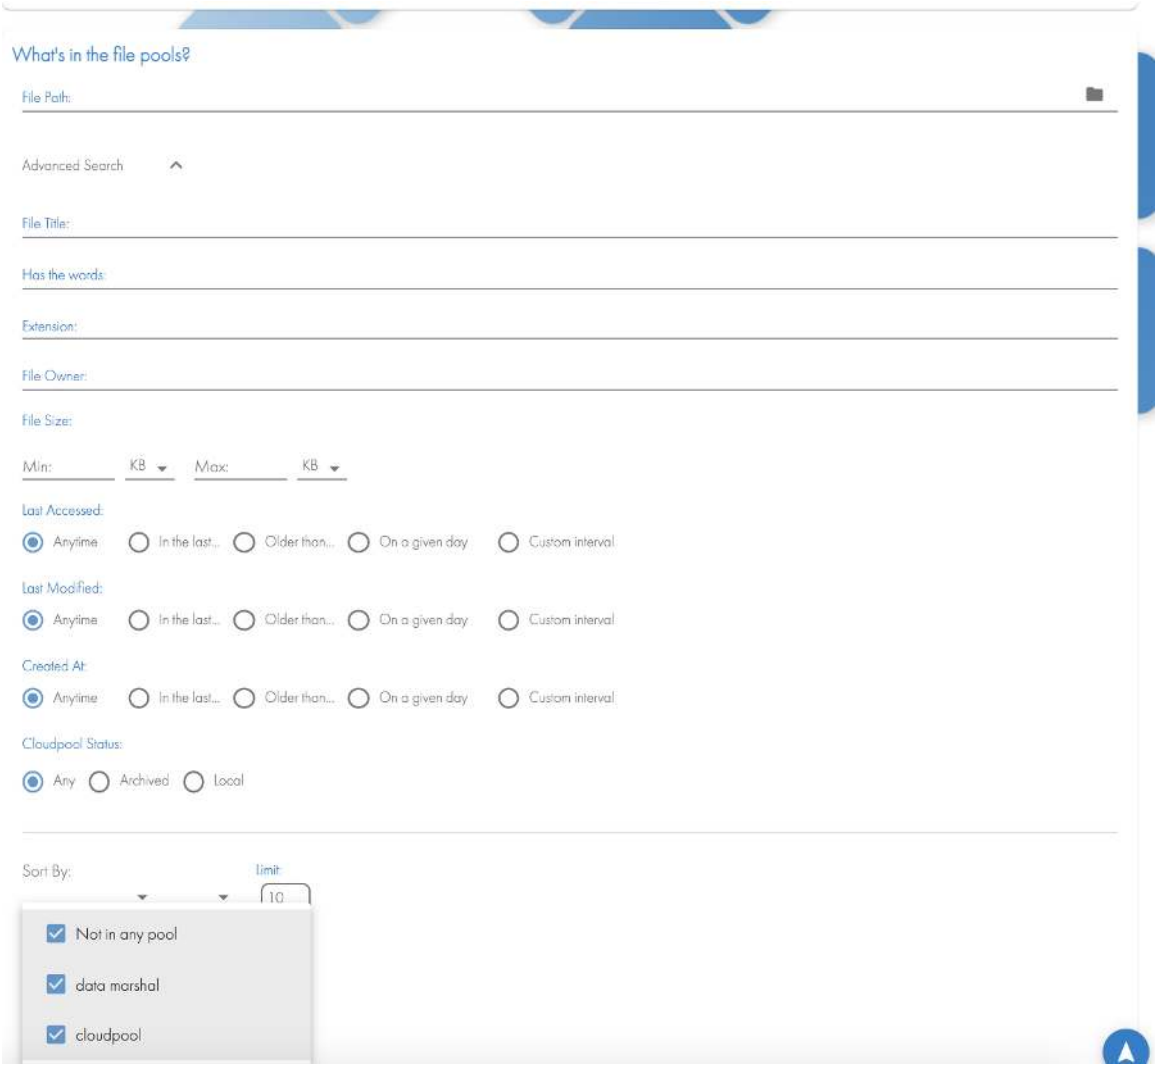

6. Example Report showing each File pool policy usage graph and table.

7.

When can I delete my worm lock data?

1. This report will assist administrators to locate and see the quantity of data into the future that will expiry lock status. This is data that can now be deleted since it has expired its retention period. This will assist administrators with reclaiming disk space and predict disk usage more accurately.
2. This report will also allow drill in to locate data in a given month or week that has expired and can be deleted with the script command builder feature to assist with deleting expired data regardless of its location within the Worm folder structure. This is done using the last accessed date stamp.
3. Requirements: Releasee 1.1.5 or later
4. Example Report

a. 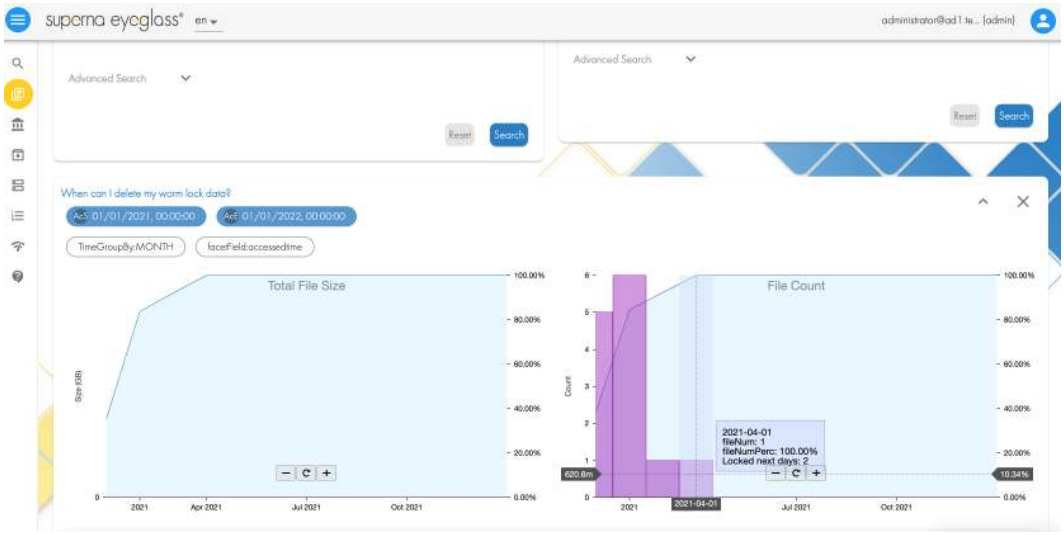

How Fast is My Data Growing?

1. This quick report helps forecast future data growth per month by sampling the previous x months of data. This report will show month by month future data growth using an extrapolation that uses the files created date stamps.
    a. Forecasts data created and files created
2. Requirements: Release 1.1.5 or later
3. The default will forecast data usage 1 year into the feature. The advanced options allow this to be changed to additional future years. It will sample the previous 12 months of created data to complete the forecast. The table view can be downloaded as CSV or Excel report with a break down per month showing data created and files created.
4. Options:
    a. Group by month or year
    b. Forecast period 1 year or multiple years
    c. Limit analysis to a path
    d. Additional search controls to bound your search are available in the advanced configuration.

5. 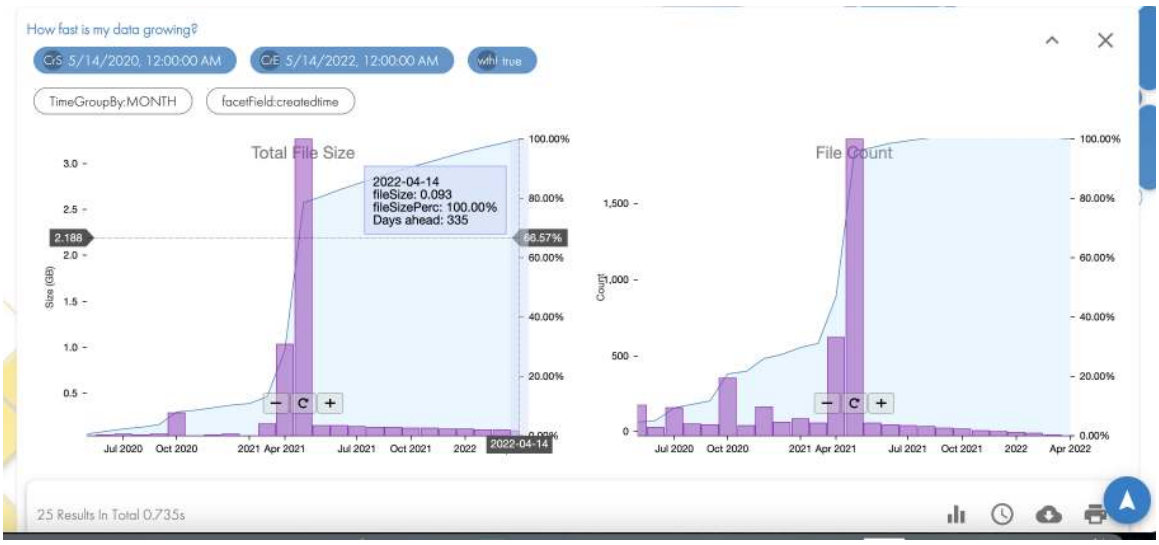

How to drill into Quick Reports to identify the files behind the results

After completing a quick report, a table of the results is shown directly below the graphs. This table allows selecting a day, or row of the results depending on the type of quick report desired. Then selecting the action menu:

1. Select data in the table by selecting check boxes next to the data you want to analyze further. You can select all or several different date ranges. **NOTE: The data in the table depends on the report setting and can be grouped by day, month or year.**
2. Then click the action menu icon.
3. Select show me the files to get a list of files that match the row selection of the quick report.

4.

How to Schedule Quick Reports and email results

Requirements

1. Configure email channel - see guide.
2. Release 1.1.5 or later
3. NOTE: Quick reports or file searches can be scheduled

Configuration

1. Click the clock icon above the quick report results

2. 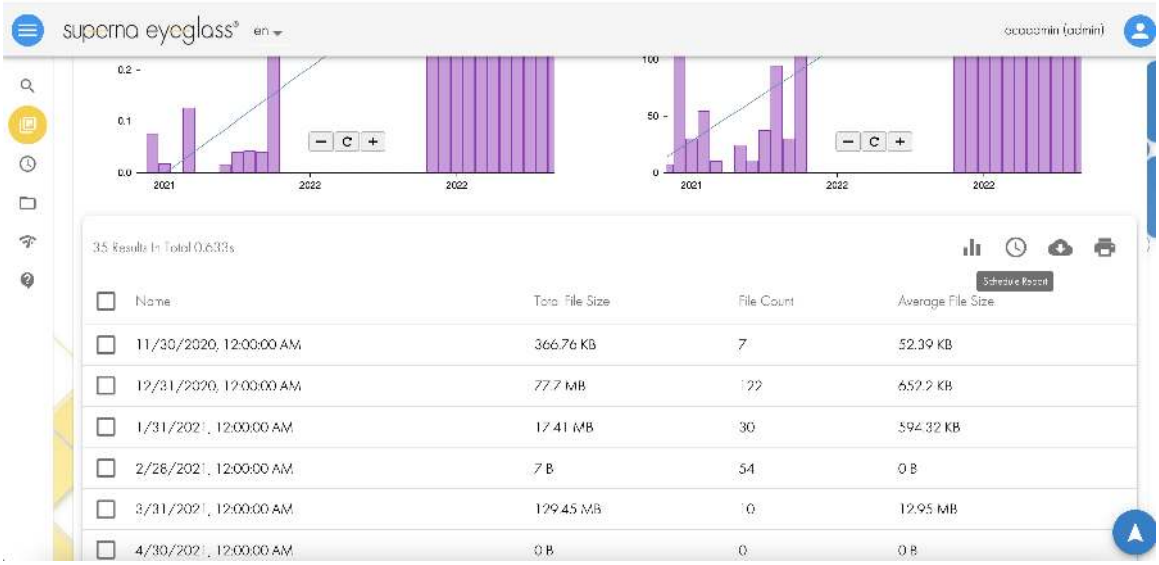

3. Fill in:
    a. name of the scheduled report
    b. format csv or excel
    c. interval to run the report
    d. email addresses for To and CC
    e. Subject of the email
4. Click test button to verify the email is received
5. Click submit to save

6. 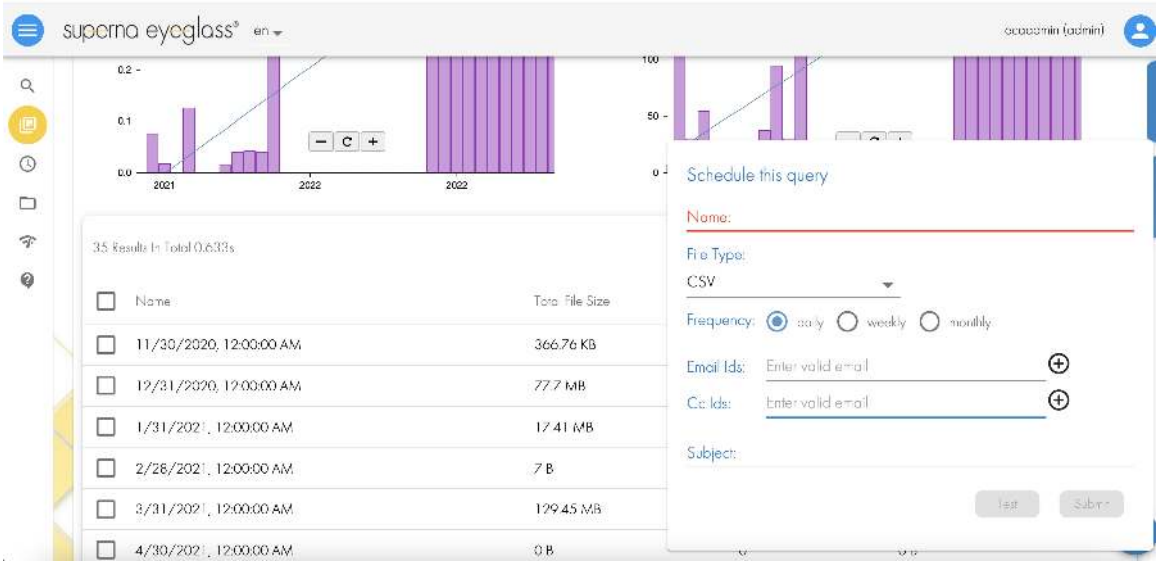

How to edit or delete Scheduled Reports

Requirements

1. Release 1.1.5 update 1

Configuration

1. Login as an admin user
2. click side bar clock icon
3. click pencil icon to edit and make changes
4. Selection check box to select and the trash icon will appear to delete

5. 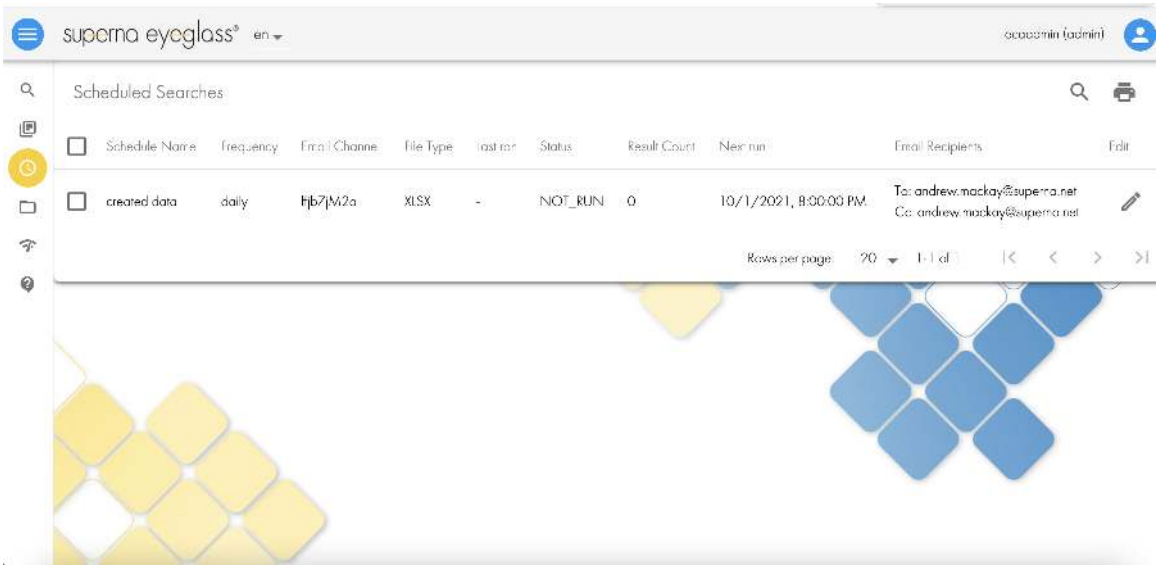

File System Interactive Browser - Beta

Overview

This new tool allows browsing the file system to locate folders that are consuming the majority of the space or the number files in the folder. Each time the tree is expanded index queries are executed and the folder tree is color coded to show the folder that consume 90% of the space of the parent folder total usage, 75% to 90% and less than 75%. This color code scheme allows identifying where in the file system space is consumed.

The length of the colored bar indicates the relative % of the usage of the parent and the color indicates the range of the %. In the example below the dfsdata folder and the Isilon_support folder both fall in the range 75% to 90 % of the parent folder /ifs/data and /ifs/data is > 90% of /ifs.

superna eyeglass
en
ecaadmin (admin)
prod8
ifs
.snapshot
veeam
test
home
igls-roboaudit
igls-securityguard
synctest
data
staging
userdata
dfsdata
mass delete
lockout
lowercasefolder
cloudpool
honeypot
Isilon_Support
Bytes
2.8 GB
0 B
0 B
0 B
0 B
0 B
0 B
0 B
2.8 GB
0 B
6.2 KB
1.1 GB
0 B
0 B
0 B
10.1 MB
7 B
1.7 GB
Files
5636
0
0
0
0
0
0
0
5636
0
181
4991
2496
0
1
10
2
462

© Superna LLC

9.16. Advanced Searching Syntax and Script Automation Editor

Home Top

- Use Cases
- How to Use the Script Automation Editor
  - How to download very large scripts
  - How to use the Script Editor
- Search Operators explained with examples
  - Language Detection codes
- How to search for directories only
- How to count the number of files at a path and below
- How to search the size field to find files with certain file size ranges
- How to use Wild Cards for single character or multiple characters
- How to search with AND to join terms together
- How to Search for email addresses and partial email addresses
- How to search for a phrase
- How to search for a term that MUST exist or MUST NOT exist in a document to narrow the search to a userid, employee

Use Cases

1. **User Search:** Finding email addresses in documents.
2. **Compliance:** Find credit numbers in a document.

3. **e-Discovery, GDPR:** Finding employee or customer ID with wild cards.
4. **e-Discovery, User Search, GDPR: S**single character and multi character searches with ? and *.
5. **e-Discovery, User Search, GDPR:** Finding documents with more than one mandatory term in the document example customer ID + some other term.
6. **File Recovery, User Search:** Find files with wild cards.
7. **Capacity Management:** Find all files greater than or less than a size or in a range of file sizes.

How to Use the Script Automation Editor

The Script Automation editor icon on the UI allows administrators or end users to generate a script download that can automate some tasks with the results of a search. The CSV and script download icons limit results to 100 000 rows, any search results over this limit will automatically run as a server side job to create the file that is available from the secure download URL.

How to download very large scripts

1. The server side download URL provides backup archive, and CSV script downloads.

2. Using a browser access the URL below (**NOTE: the secure access userid and password are required**). Default userid is ecaadmin. The password is set during installation.
    a. https://x.x.x.x/downloads (x.x.x.x is node 1 of the search cluster)

How to use the Script Editor

1. The window shown below allows users to select an absolute path in the script download (i.e. /ifs/data/....), which is useful when the script will be run directly on the PowerScale using ssh and a .sh script.
    a. The other option is to insert the UNC path in the downloaded script which enables a script to be run by users over SMB connections to the cluster. The UNC path is the SmartConnect address and SMB share to reach the file over the network. This allows powershell or batch cmd files to run automation against a result of files.
    b. 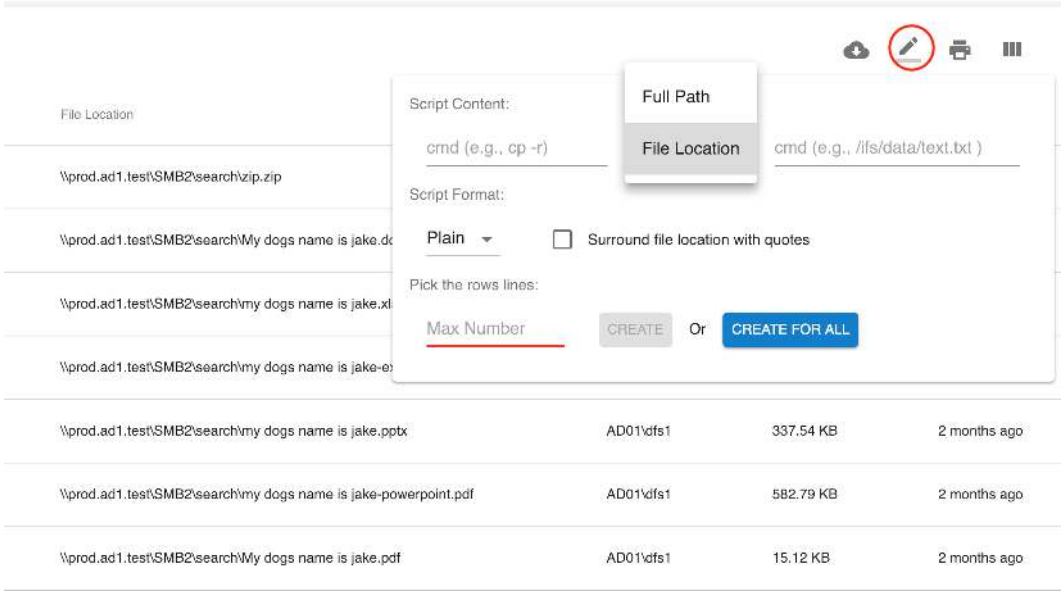

2. The other option available in the script download is plan file download with rows "Plain" option. This would be used when the script language needs to be changed or modified. Typically used when bash shell execution will not be used. The Shell option will add bash shell and download the file with a .sh extension so that the file can be copied to PowerScale to be run on the cluster itself. This would typically be used by PowerScale administrators.

   a. 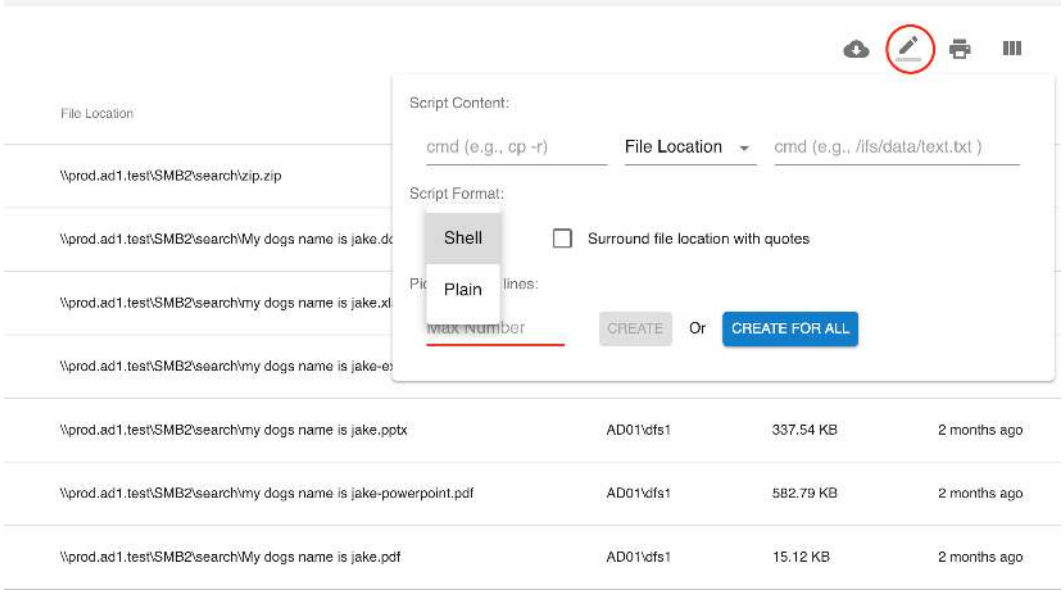

3. Enter the script action in the first dialog (i.e. cp, mv , del, rm etc..) depending on the script language you plan to use. Enter the output of the command which could use >> results.txt , or | grep "some string to look at".
   a. Examples:
      i. isi get -d
      ii. cp
      iii. mv

Search Operators explained with examples

1. + and - term operator
2. ? single character wild card
3. * multiple character wild card
4. fieldname: (specific schema search examples below):
    a. **filesize**: (file size in bytes).
    b. **extension**: (file extension).
    c. **language_s**: (detected language of the document).
    d. **clustername**: (cluster name).
    e. **ishidden**: (hidden file in the file system).
    f. **path**: (search the full path to a file using key words).
    g. **type**: (find files or directories valid values are **files** or **directory**).
    h. **owner**: (ad or Linux owner of the file in the form domain name\\userid. NOTE: the double slash is needed i.e. owner:AD01\\usera users).
    i. **group:** (group ownership of the file example group:AD01\\domain users).

Language Detection codes

Note: all languages are tested or supported. Provided as the language codes, these codes help searching the language metadata field.

Language Code	Language
af	Afrikaans

ar	Arabic
bg	Bulgarian
bn	Bengali
cs	Czech
da	Danish
de	German
el	Greek
en	English
es	Spanish
et	Estonian
fa	Persian
fi	Finnish
fr	French
gu	Gujarati
he	Hebrew
hi	Hindi
hr	Croatian
hu	Hungarian
id	Indonesian
it	Italian
ja	Japanese
kn	Kannada
ko	Korean
lt	Lithuanian
lv	Latvian
mk	Macedonian
ml	Malayalam
mr	Marathi
ne	Nepali
nl	Dutch
no	Norwegian
pa	Punjabi
pl	Polish
pt	Portuguese

ro	Romanian
ru	Russian
sk	Slovak
sl	Slovene
so	Somali
sq	Albanian
sv	Swedish
sw	Swahili
ta	Tamil
te	Telugu
th	Thai
tl	Tagalog
tr	Turkish
uk	Ukrainian
ur	Urdu
vi	Vietnamese
zh-cn	Simplified Chinese
zh-tw	Traditional Chinese

How to search for directories only

1. Use this search to find directories and below a specific path (NOTE: This requires the logged in user to be on the admin list to use the absolute path option shown below)
    a. In the search box type "type:directory" to include in the search filter. Enter the absolute path in the "File Path" field to constrain the search to this path and below. The results count the number of directories. Downloading results will provide a directory result csv.

b. 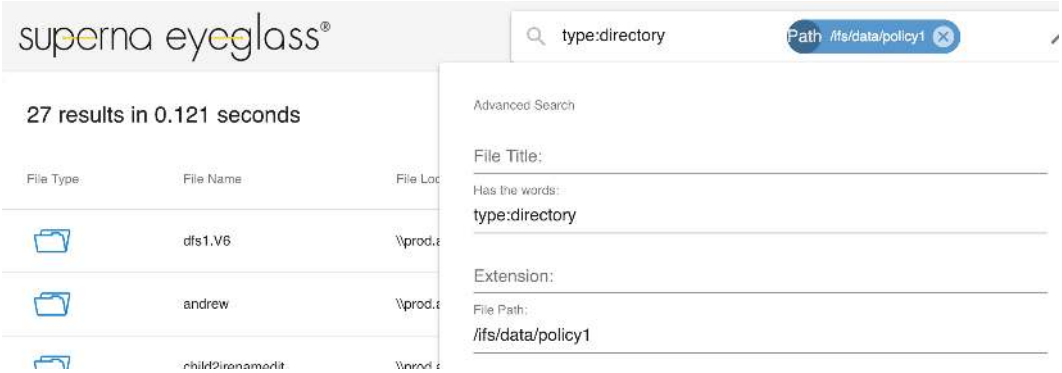

c. Or use the folders only option in the advanced search field.

d. 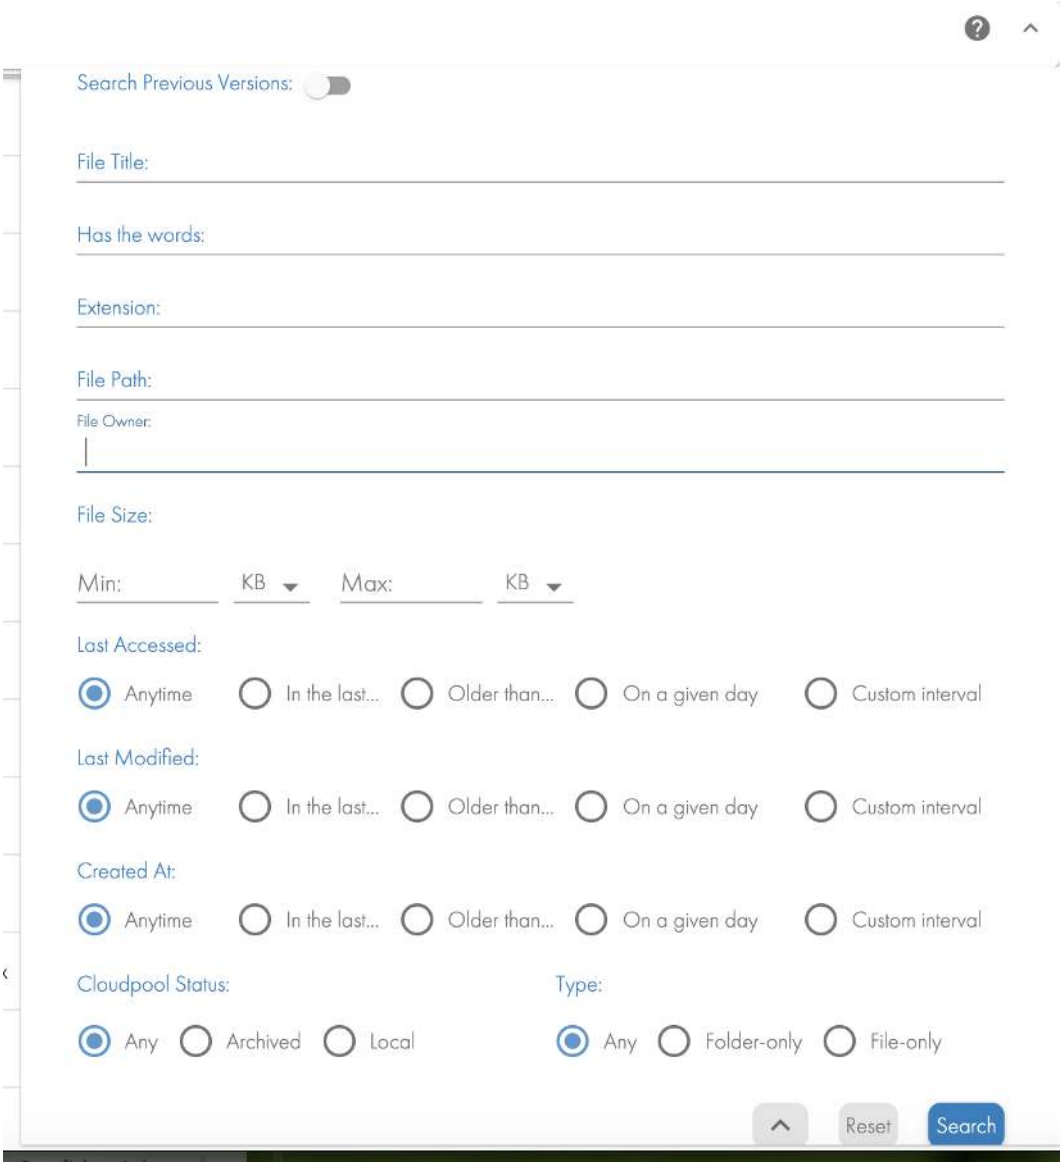

How to count the number of files at a path and below

1. Add your user id to the admin list (see the configuration guide for the cli command)
2. In the Advanced Search drop down, enter the path in using absolute path starting with "/ifs" and enter the starting point of the search.
    a. Enter no text in the Search box to match all files, the file result count equals the number of files at the path entered and below.
    b. 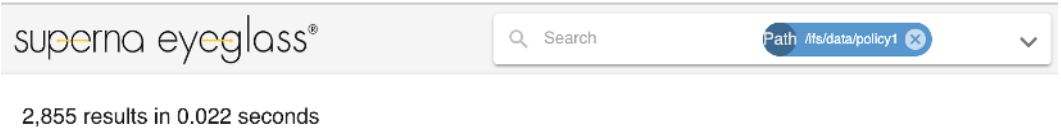

How to search the size field to find files with certain file size ranges

1. This search allows a specific field to be searched, and when combined with other search terms allows for very controlled search of data. These examples can be entered into the search bar versus using the advanced UI
    a. **Enter Search: filesize:[10 TO 100] AND extension:pdf** (this will find all the pdf files that are 10 - 100 bytes).
    b. **Enter Search: filesize:[100 TO *] AND extension:pdf** (this will find all the pdf files that are 100 bytes or greater).

c. **NOTE: TO must be upper case AND operator must be upper case. File sizes are stored as bytes and will require a GB or MB to byte conversion. Many online calculators can assist with this conversion.**

How to use Wild Cards for single character or multiple characters

1. In the search dialog box you can use ? to wild card a single character or use * to represent any characters. See examples below:
    a. Employee ID: 3 letters + 3 numbers example abc123 **Enter Search: abc1??** will find all employees that start with abc1.
    b. Credit card search for PCI: **Enter Search: (master card) 5524-????-????-???? will find all documents with this pattern begging with 5524.**
    c. Multi character wildcard: **Enter Search: abc***

How to search with AND to join terms together

1. **Enter search for 2 terms:** dog AND cat (use the AND upper case to join the terms that must be in the document).
2. This means both dog and cat must be in the document to match.

3. The default for terms is OR with no operator specified example. dog cat is the same as dog OR cat.

How to Search for email addresses and partial email addresses

1. Type the full email address into the search bar will find exact match in documents.
2. To search a partial email address use wild cards to find the addresses per the following examples:
    a. *** multi character** wild card use case: To find "test.user@superna.net" combine a phrase and operator to find documents:
        i. test* AND "@superna.net"
        ii. *.user AND "@superna.net"
    b. **? single character** wild card use case: To find "test.user@superna.net" combine a phrase and operator to find documents:
        i. t???.user AND "@superna.net"
        ii. test.us?? AND "@superna.net"

How to search for a phrase

1. Use double quotes.
2. "my dogs name" These words in this exact order must appear in the document.

How to search for a term that MUST exist or MUST NOT exist in a document to narrow the search to a userid, employee

There are many use cases for this type of search. For example a search to find documents that must have an employee ID.

1. **Enter Search for employee with userid of abc123 and add plus sign + to the term that must exist. Example "**+abc123 performance reviews" (this search will find all documents that have abc123 and then find all documents that have performance OR reviews as words in those documents containing abc123.
2. How to find documents that do not have a term.
    a. **Enter Search without a userid of abc123: -abc123 performance reviews** (You can combine this with other terms, for example find all documents that do not contain the user ID abc123 and may contain contain performance OR reviews terms)

© Superna LLC

9.17. DR Failover Considerations

Overview

The license keys are attached to a read write cluster:

1. License keys are bound to a cluster.
2. The DR cluster will require NFS exports created.
3. The DR cluster may have a different path for the same content found on the source cluster, since SyncIQ allows changing the target path on the DR cluster.
4. The DR cluster will require adding the folder paths to the configuration.
5. License keys cannot be moved to DR clusters without opening a request to sales, only DR copy of source cluster data will be considered for a license key move. The move is a permanent move of the license.
6. **Limitations:**
    a. The original indexed files from the production cluster will still be found in the index, new or changed content on the DR cluster will be ingested.
    b. This means the same document Prod and DR will be returned in results, since the file uniqueness is based on the absolute path from /ifs to the file.
    c. The results will show the prod cluster modified date as being older and the DR cluster version being newer.

d. Deleting a file on DR will NOT delete the reference to the file in the index from the Prod cluster, if the absolute path is different between prod and dr clusters.

e. File Delete when the path on prod and dr is the same the index will remove the file from the index.

f. Secure Access to the old cluster will not be possible for end users.

© Superna LLC

9.18. Management Diagnostic Tools

Home Top

The following tools are used to monitor ingestion of files and monitor the index and diagnose health issues of the search cluster process.

These tools have WebUI's secured through a webe proxy on node 1. The WebUI's are secured over https and requires a user and password to access. **NOTE: the default user is "ecaadmin" and the default password is "3y3gl4ss". This password is changed during installation. This password should be secured as the index management UI can by pass the AD security of the indexed files.**

File Index Message queue monitoring

1. https://x.x.x.x/kafka-manager/ **(on node 1 only)**
2. The monitoring of backlog of file ingestion for incremental or full ingestion of files is located on this URL :

https://x.x.x.x/kafka-manager/clusters/search/consumers

Index Health, Diagnostics and Management

1. https://x.x.x.x/solr
2. Cluster index health
   URL https://x.x.x.x/solr/#/~cloud?view=graph
3. Cluster index size per
   shard https://x.x.x.x/solr/#/~cloud?view=nodes

© Superna LLC

9.19. Monthly Index Backup Solution Guide

Home Top

To protect the index we recommend a monthly or bi-weekly back to NFS export on the PowerScale. This provides a recovery point of a large index stored within the Search & Recover cluster.

- BACKUP: Backup Index
- RESTORE: Restore Index

BACKUP: Backup Index

The BACKUP command will backup Search & Recover indexes and configurations for a specified Index. The BACKUP command takes one copy from each shard for the indexes. For configurations, it backs up the configSet that was associated with the collection and metadata. Use the following command to back up igls Search & Recover collection and associated configurations to PowerScale over NFS:

`/admin/collections?action=BACKUP&name=iglssearchbackup1&collection=igls&location=`/opt/superna/mnt/backup/&async=*task-id*

Procedure:

1. Create NFS export on PowerScale for this backup. Example: Create NFS export with path "/ifs/searchindexbackup", and configure to let Search & Recover nodes to have read and write permission to this NFS export by adding the Search and Recover ip to the read/write client list on the export.
    a. ssh to the cluster as root

i. mkdir -p /ifs/searchindexbackup

b. Create the user named "eyeglasshdfs" in the local system provider, no password is required when creating this user. This user will own the files on PowerScale.

c. Configure ownership of that NFS export path on PowerScale: chown -R eyeglasshdfs:"PowerScale Users" /ifs/searchindexbackup .

d. Change mode of this directory: chmod -R 777 /ifs/searchindexbackup .

2. **On each of the Search & Recover cluster nodes:**

a. Mount the NFS export to the mount point on each solr node. Replace yellow with SmartConnect name. Example: mount -t nfs -o nfsvers=3 <dns name of smartconnect>:/ifs/searchindexbackup /opt/superna/mnt/solr-backup .

b. Repeat these steps on each node starting at node 2 to 4 or 7 depending on the Search cluster size.

c. To ensure the NFS mount persists a reboot:

i. Complete these steps on nodes 2 - X (X is the last node in the cluster, depending on the size of your Search & Recover cluster) vim /etc/fstab .

ii. Replace yellow highlight with the correct values for your cluster. NOTE: the FQDN should be a SmartConnect name for a pool in the System Access Zone IP Pool SmartConnect.

iii. FQDN:/ifs/searchindexbackup /opt/superna/mnt/solr-backup nfs ro 0 0 .

iv. Save the file.

3. **Restart the cluster to allow new mount to be visible:**

a. SSH to node 1 of the cluster as ecaadmin.

b. ecactl cluster down (wait for this to finish).

c. ecactl cluster up .

d. Verify that the NFS mounted directory is in the mount list of solr container.

   i. ecactl containers exec solr mount .

4. **Execute the backup command**:

   a. location=/opt/superna/mnt/backup **(this is a the local location in the VM that is mounted to the PowerScale export).**

   b. Task-id = 1 (any integer can be used to monitor task (will be used to check the status with REQUESTSTATUS command).

   c. Login to node 2 using ssh and ecaadmin user account.

   d. Run this command:

      i. curl 'http://node2-IP:8983/solr/admin/collections?action=BACKUP&name=iglssearchbackup1&collection=igls&location=/opt/superna/mnt/backup/&async=*1*'

      ii. Then use this command to monitor progress:

         1. curl 'http://node2-IP:8983/solr/admin/collections?action=REQUESTSTATUS&requestid=*1*'

   e. Once that task has been completed, the action=REQUESTSTATUS will return the status of backup (success/failed).

   f. **Note for a large index this backup can take hours.**

   g. Once completed login via ssh to the PowerScale and verify the backup directory contains files.

   h. The size of the index backup will be smaller than the index size on the cluster.

RESTORE: Restore Index

The RESTORE command will create a Index with the specified name in the collection parameter. Use the following command to restore igls Search & Recover index and associated configurations:

`/admin/collections?action=RESTORE&name=iglssearchbackup1&collection=igls&location=`/opt/superna/mnt/backup&async=*task-id*

The target collection should not be present at the time the API is called, as Search & Recover will create this collection. In order to restore with the same collection name, we should delete the existing collection with DELETE Command.

Procedure:

1. Delete existing Index from the Collection screen in the GUI:

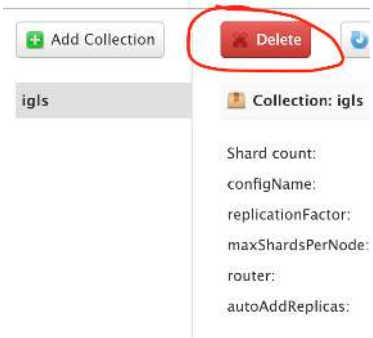

   a.

2. **Restore collection:**
   a. Login to node 2 of the cluster and execute this command:
   b. Example: "curl 'http://node2-ip:8983/solr/admin/collections?action=RESTORE&name=`iglssearchbackup1`&collection=igls&location=/opt/superna/mnt/backup/&async=*1*'"
   c. Check the status of the request task for that task-id:
      i. curl 'http://node1-IP:8983/solr/admin/collections?action=REQUESTSTATUS&requestid=*task-id*'
      ii. **NOTE: This process can take hours on a large restore**

iii. Once that task has been completed, the action=REQUESTSTATUS will return the status of backup (success/failed).

d. Use GUI to verify the collection after the restore that everything is green:

i. 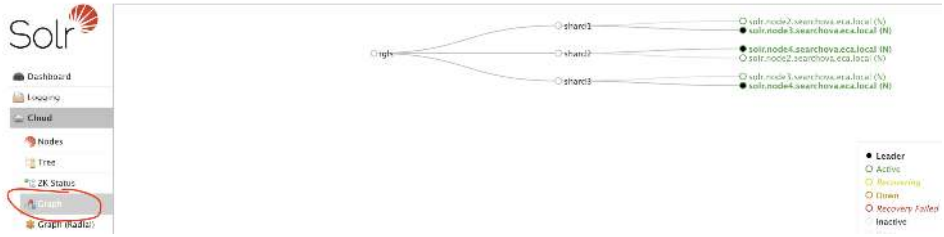

e. Done.

f. Verify ingestion tasks are functioning be creating new files and verify you can search for new files.

g. Use health check process to verify ingestion and stats command to see that files are being added to the index successfully. See Configuration section.

1.

© Superna LLC

9.20. Trouble Shooting Search Results

Home Top

- Did you run the index job on the folder after adding it?
- User can not login?
- Users can login but no search results returned?
- User can search for file names and paths but not based on the contents of the file?
- I cannot find any files when I search?

1. Did you run the index job on the folder after adding it?

a. If you are not sure, then go back to the quick start steps to review how to start the index to index the path.

2. User can not login?

a. The user AD login is tested against the cluster ip address used to add to Search & Recover. If no SMB shares exist in the System Zone, the login will fail. It will also fail if no auth provider can resolve the user in the System Zone.

b. If the user belongs to another Access Zone , then a UNC path must be added to Search & Recover for that Access Zone to allow authentication. See the guide on adding additional SmartConnect UNC paths for authentication in any Access Zones that have indexed data. Learn more

about - Managing Search Dynamic UNC path CLI commands

3. Users can login but no search results returned?

a. When a user logins their SMB shares are determined based on AD group membership. The list of SMB shares determines that paths in the file system that results will be returned to the user.

b. If a user has no SMB share path with indexed data no results will be returned.

c. Use this commands listed here to view the users SMB shares, paths and zones that were auto discovered based on AD group membership. This will show which paths to expect result if data has been indexed on the path or below the path. **(NOTE: the SMB share path, and all folders below this path are included in the search results, but only paths that have been indexed will be returned).**

4. User can search for file names and paths but not based on the contents of the file?

a. Default settings on folders is meta data only, a modify folder command is required to add the file extensions that should have the content indexed (i.e. *.pdf, *.docx).

b. See the modify folder command for examples on how to modify a folder configuration for file extension content indexing. **NOTE: a folder index job will need to be started again for existing files on disk that have already been**

indexed. Incremental indexing will start to index the file extensions on the next incremental job. Default incremental indexing is hourly.

5. I cannot find any files when I search?

a. Please review all of the above to make sure these have been done first.

b. **To monitor a running index job, and see if any errors are occurring during indexing, follow these steps:**

   i. **Start a folder index job. Instructions here.**

   ii. Monitor the running index job and look for errors. Instructions here.

© Superna LLC

9.21. Advanced Cluster Configuration

Home Top

The clusters configuration can be customized from defaults. The cluster can be shutdown, the main configuration file updated, and restart the cluster for the changes to take effect.

- Compressed File Content Ingestion Handling
- How to configure faster Content indexing with Parallel Content Ingestion per Worker
- How to Analyze Content Parsing Latency Distribution
- Advanced Cluster Configuration

Compressed File Content Ingestion Handling

1. Shutdown the cluster first.
    a. ecactl cluster down .
2. Edit the file "/opt/superna/eca/eca-env-common.conf ".
    a. vim /opt/superna/eca/eca-env-common.conf .
3. Add variable to set the size of compressed files that should be processed for full content indexing by decompressing the file, and then indexing the files inside the compressed file.
4. Add the line below to change the size of the compressed files:
    a. export INGESTION_WORKER_MAX_COMPRESSED_FILE_SIZE_MB=x (where x is a number in MB's, (default

value is 0.4 * INGESTION_WORKER_MAX_FILE_SIZE_MB)).

5. Add the line below to control the compressed file types that will be processed:
    a. export INGESTION_WORKER_COMPRESSED_FILE_EXTENSIONS=.zip,.tar.gz .
6. Save the file with changes :wq .
7. Start the cluster:
    a. ecactl cluster up .

How to configure faster Content indexing with Parallel Content Ingestion per Worker

1. Follow these step to enable more CPU consumption to indexing more files per working in parallel .
2. Shut down the cluster
    a. ecactl cluster down .
3. Edit the file "/opt/superna/eca/eca-env-common.conf"
    a. vim /opt/superna/eca/eca-env-common.conf
4. Add variable to set the size of compressed files that should be processed for full content indexing by decompressing the file, and then indexing the files inside the compressed file.
5. Add the line below to change the size of the compressed files:

   a. export INGESTION_WORKER_PARALLEL_LIMIT=x (where x is a number of files to process at a time, default is 5 files))

6. Save the file with changes :wq
7. Start the cluster
   a. ecactl cluster up

How to Analyze Content Parsing Latency Distribution

1. These steps will use a hidden attribute in the schema to return latency of parsing documents. This will indicate if content indexing is processing longer or shorter time based on the reading and parsing of text in various document types. Support can use this data to optimize indexing speed for content.
2. Login to solr UI
   a. https://x.x.x.x/solr (enter ecaadmin user and password)
   b. Click on the collection named IGLS on the left hand side
   c. Select the query meny option on the left hand side
   d. Fill in the query paramters as follows. See Screenshot below.
   e. In the q filed enter -0> **contentindexedat:[* TO *]**
   f. in the start , rows enter **0 and 0**
   g. In the Raw Query Parameters -- **facet.range=contentparsetimer&facet.range.start=0&facet.range.gap=100&facet.range.end=30000&stats=true&stats.field=contentparsetimer**

h. Click the Facet check box to enable it.

i. In the Facet Query enter -- **facet.range=contentparsetimer&facet.range.start=0&facet.range.gap=100&facet.range.end=30000&stats=true&stats.field=contentparsetimer**

j. **Click the Execute Query button**

k. **Results should display similar to the screen shot. Copy this text response completely and post it in to a support case for analysis.**

l.

m. 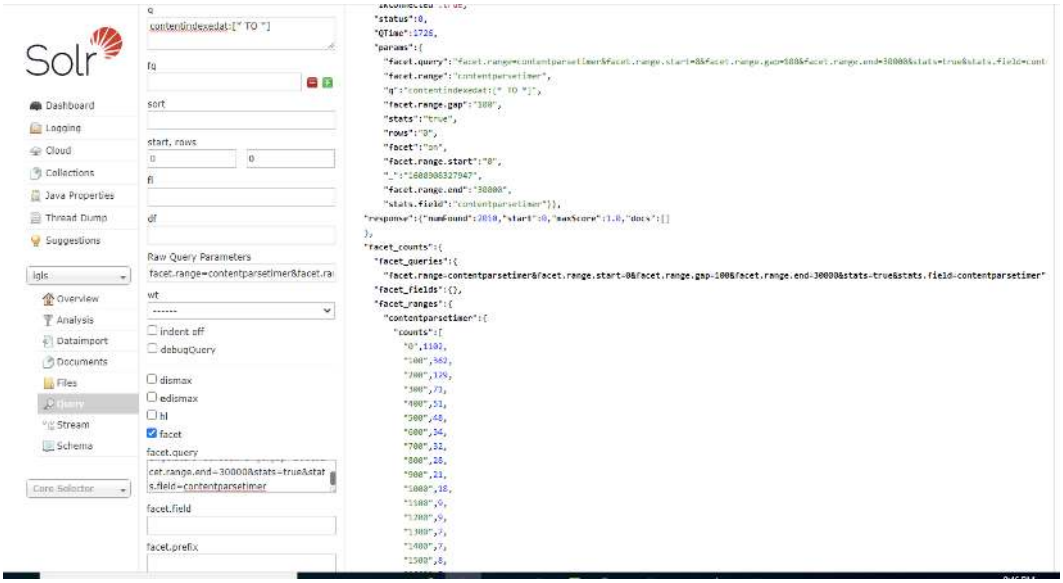

Advanced Cluster Configuration

1. **Suppress PowerScale changelist mode - Use only if directed by support**

   a. Search creates snapshots with a 5 day expiry. When we've disabled incremental ingestion, but are processing a long-running full ingestion, this snapshot can be deleted and

we're left without one on disk. To avoid the issue, add an optional setting to "eca-env-common.conf" to disable incremental that will stop processing after the snapshot has been created, but before the changelist is created. This will keep creating snaps on the PowerScale, but will not process any changelist data.

b. add to eca-env-common.conf and requires a restart of the taskmaster container on node 1 .

c. export SUPPRESS_CHANGELIST_TASK_CRON="0 0 * * *"

© Superna LLC

9.22. Cloud Pool Reporting

Home Top

- Overview
  - Use Cases
- Examples of how to Search for File Stubs
- How to report on GB's of Cloudpool stubs that are being accessed
- How To Configure Skip Indexing Stub Files with content Indexing
  - NOTES
  - Procedure

Overview

Search & Recover is the only product that fully supports Cloudpool reporting. This desribes how Cloudpool status can be used in the Search UI and Quick reports.

Use Cases

1. Index cloud pool data to track which files are stubs and which are not to make sure File Pool policies are correctly matching the expected files.
2. Quick report to show the total sum of all stub files including total count of files, and total GB of stub files currently in the cloud tier.

3. Advanced search ability to search file files using any supported criteria, including content with option to only return files that are stub files.

Examples of how to Search for File Stubs

1. Advanced Search UI Option allows searching for Stub files only.

   a. 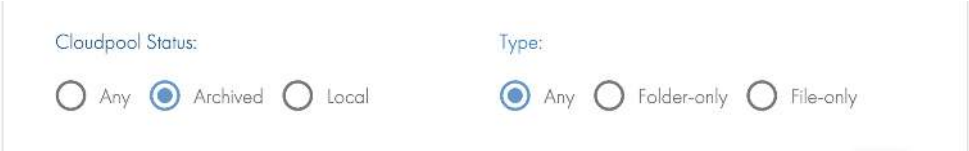

2. Quick Report Who Owns What?, What are the Types?, What's growing old? all have an advanced search option to select Cloudpool Status Archived when searching to narrow the results to only include stub files.

How to report on GB's of Cloudpool stubs that are being accessed

1. If Cloudpool stubs are being accessed frequently, it indicates the files are not good candidates to be tiered to slower S3 tier. These steps will identify Cloudpool files that are still being accessed and will allow reporting on the files along with the owner of the files.
2. Open the Quick reports Tab and configure the What's Growing Old? Quick report.
3. Select the Cloudpool status Archived

4. Select the Last Accessed option in the last 1 month

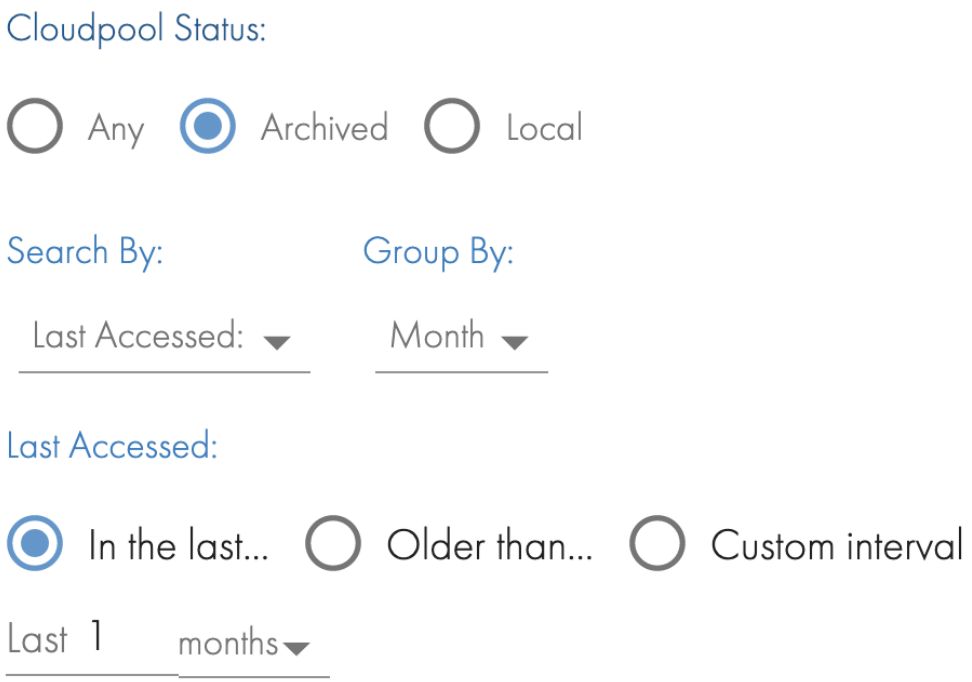

a.

5. The results will show the quantity of data in the last 30 days that has been accessed (requires the last accessed attribute to be enabled on the cluster).
6. The table below will show the quantity of data and selecting this with the check box and Action menu Show my the files will allow reporting on the files that are being accessed.
7. This list would be reviewed to see which users are accessing data and where in the the file system the files exist. This review can be used to revise the File Pool policy that is stubbing files.

How To Configure Skip Indexing Stub Files with content Indexing

NOTES

1. Once files are stubbed they can still be read by any user. This will inflate the portion of the file the user is accessing. Indexing can inflate files that are stored in the Cloud causing a lot of bandwidth from the Cloud repository to the PowerScale cluster.
2. In some cases it may not desirable to inflate files during indexing.
3. The following advanced option can be configured to override content indexing settings, if the file is detected as a stub file, to avoid inflating the file.
4. **Note: This option is applied to the cluster and requires a cluster down and up operation.**

Procedure

1. ssh to node 1 as ecaadmin.
2. Open the master config file "/opt/superna/eca/eca-evn-common.conf".
3. Add a variable export INFLATE_CLOUDPOOL_STUBS=true .
4. Save the file.
5. ecactl cluster down .
6. Wait until the cluster is down.
7. ecactl cluster up.
8. Done.

© Superna LLC

9.23. Eyeglass Search GraphQL API

Home Top

- Overview
- Summary
- Query Format and routes
  - Endpoint:
  - Parameter Encoding
  - Authentication
  - login
  - Query
    - Response
    - Example:
- Queries
  - fileInfo
    - Query
    - Response
    - Example 1: Run an empty query to get all of the file names and owners
    - Example 2: Paginate through *.txt files
  - getSummaryReport
  - Query
    - Response
    - Example 1: Get the counts and sizes of top two file extensions
    - Example 2: Groups all of the files on a path by last modified date

Overview

The API guide provides API usage guidelines for customers that want to automate searches and use the file results in an automation process. The API could also be used to integrate the search results in indexing into custom developed applications or web pages. The scope of this integration is outside the scope of the API documentation.

Summary

The Eyeglass Search GraphQL Api (ESGA) is an authenticated, remote interface that runs over http for querying an Eyeglass search appliance for files.

Query Format and routes

Queries to ESGA can be issued to any node in the search cluster. Queries run over https with the bulk of the query passed as URL parameters.

Endpoint:

All Queries must be issued to: https://ip.of.search.node/graphql

Parameter Encoding

GraphQL queries must be issued as a URL encoded value to the query http parameter:
https://ip.of.search.node/graphql?query=url_encoded_graphql_query
Examples of using curl to encode the query can be seen below.

Authentication

Authentication is achieved through the retrieval of a json web token. This token is to be included in the “Bearer” header of all future calls to ESGA.

login

Query

Schema		
`login(id: String!, pass: String!): LoginResult`		
Argument	**Type**	**Value**

id (required)	String	Username + domain of the user logging in, in the user@domain.com syntax. For local users, omit the domain.
pass (required)	String	Password of the user attempting to log in.

Response

Schema		

`type: LoginResult {`
`  user: User!`
`  token: String!`
`}`		

Field	**Type**	**Value**
user (non-null)	Object (`User`)	User Object Object for the logged in user.
token (non-null)	String	The JSON Web Token to be used in authorizing future ESGA calls.

`type: User {`
`  name: String!`
`  role: String!`
`}`		

Field	**Type**	**Value**
name (non-null)	String	The name of the logged in user, in DOMAIN\username format
role (non-null)	String	one of: USER or ADMIN

Example:

Login with the username: testuser@exampledomain.com, using password: NotReal!:

```
curl -s -G -k https://search.igls.com/graphql --data-urlencode 'query={
  login(id:"testuser@exampledomain.com", pass:"NotReal!") {
    user {
      name
          role
    }
    token
  }
}'
```

```
{
        "data": {
```

```
            "login": {
                "token":
"eyJhbGciOiJIUzI1NiJ9.eyJzdWIiOiJTSUQ6Uy0xLTUtMjEtMjAxODMyNTY2LTMxODczNTM0MDc
tMjgyOTk5MTcxMC0xNDQ4Iiwicm9sZSI6IlVTRVIiLCJleHAiOjE1NzAyODk3MTR9.0vs-
tnOgs0cOBSYB_SOuNHdmV7NT6YisTwuIbZFMkJE",
                "user": {
                    "name": "EXAMPLEDOMAIN\\testuser",
                    "role": "USER"
                }
            }
        }
}
```


For any authenticated endpoints, the value of `token` needs to be used in the `Authorization` header using the `Bearer` schema

```
Authorization: Bearer <token>
```

Queries

fileInfo

The main query to execute searches. Returns a list of file records that match the provided query.

Query

Schema

```
fileInfo(
  content: String,
  fileName: String,
  fileExt: String,
  lastModifiedStart: String,
  lastModifiedEnd: String,
  fileOwner: String,
  lastAccessStart: String,
  lastAccessEnd: String,
  creationStart: String,
  creationEnd: String,
  startCursor: Int,
  createNewQuery: String,
  filePath: String,
  clusterGuid: String,
  fileSizeMin: String,
  fileSizeMax: String,
  rowsPerQuery: Int,
  group: String,
  stub: String,
  historySearch: Boolean
```

): FileInfoList		
Argument	**Type**	**Value**
content	String	Search for words in files. Accepts simple words and phrases, or more complicated Lucene queries, of the form “dog AND (cat OR bat)”
fileName	String	Search for files with this exact filename, not including the path, but including the extension. Accepts wildcards.
fileExt	String	Search for files with this file extension. The extension is defined as the set of characters after the last dot (.) in the filename.
lastModifiedStart	String	filter by documents that have a last modified timestamp later in time than the value that is set here. Dates must be formatted with the `YYYY-MM-DDThh:mm:ssZ` format. For example, 03:14 in the morning of March 14, 2020 is written `2020-03-14T03:14:15Z`
lastModifiedEnd	String	filter by documents that have a last modified timestamp earlier in time than the value that is set here. See lastModifiedStart for syntax and example.
lastAccessStart	String	filter by documents that have a last accessed timestamp later in time than the value that is set here. atimes must be enabled on your PowerScale cluster. See lastModifiedStart for syntax and example.
lastAccessEnd	String	filter by documents that have a last accessed timestamp earlier in time than the value that is set here. atimes must be enabled on your PowerScale cluster. See lastModifiedStart for syntax and example.
creationStart	String	filter by documents that have a created timestamp later in time than the value that is set here. See lastModifiedStart for syntax and example.
creationEnd	String	filter by documents that have a created accessed timestamp earlier in time than the value that is set here. See lastModifiedStart for syntax and example.
fileOwner	String	return only documents that match this file owner. Owners must be specified in the DOMAIN\user format.
filePath	String	Search for files that begin with the specified path, starting with /ifs
clusterGuid	String	Search for files that exist on the cluster specified by this GUID.
fileSizeMin	String	In bytes, the minimum file size to return in the search results.
fileSizeMax	String	In bytes, the maximum file size to return in the search results.

group	String	Filter the search results to those that are owned by this particular group. Groups must be in DOMAIN\group syntax.
stub	String	True or False, limit the results to files that are SmartLinked files (for true), or files that are not SmartLinked files (for false).
historySearch	Boolean	Execute the search in the history collection instead of the main. Reports on previous versions of documents in snapshots.
startCursor	Int	Cursor to use to start the search. Provide this value in conjunction with rowsPerQuery to paginate the search results.
rowsPerQuery	Int	Maximum number of records to return. If the number of available results exceeds rowsPerQuery, use the cursor value to paginate future responses.

Response

Schema

```
type FileInfoList {
  pageInfo: PageInfo
  fileList: [FileInfo]
}
```

Field	Type	Value
pageInfo	Object (`PageInfo`)	Information on the query and on rows.
fileLIst	List of Objects (`FileInfo`)	A list of `FileInfo` objects

```
type PageInfo {
 totalNum: Long!
 startCursor: Int!
 hasNextPage: Boolean!
 qTime: Int
}
```

Field	Type	Value
totalNum (non-null)	Long	Total number of records that were matched by this query
startCursor (non-null)	Int	A cursor to use to resume this query
hasNextPage (non-null)	Boolean	True if this is not the final page of results

qTime	Int	The time (in milliseconds) that the query took to return.

```
type FileInfo {
  fileName: String
  clusterName: String
  path: String
  displayPath: String
  ownerSid: String
  ownerName: String
  creationTime: String
  lastWriteTime: String
  lastAccessTime: String
  fileSize: String
  extension: String
  type: String
  contentIndexedAt: String
  metadataIndexedAt: String
  group: String
  stub: String
  changedTime: String
  blockSize: String
  mode: String
  language_s: String
  snapShotName: String
  isHidden: String
}
```

fileName	String	The filename (without path) of this file.
clusterName	String	The name of the PowerScale cluster where this file resides
path	String	The absolute path to the file on the PowerScale cluster
displayPath	String	The windows-specific network UNC of the file as computed by the user's share access.
ownerName	String	The name of the user that owns this file
creationTime	String	Epoch seconds timestamp of when this file was created.
changedTime	String	Epoch seconds timestamp of the last time this file was modified
lastAccessTime	String	Epoch seconds timestamp of the last time this file was accessed. Requires atime tracking on the PowerScale to be enabled.
fileSize	String	In bytes, the size of this file

extension	String	The file extension. The extension is defined as the set of characters after the last dot (.) in the filename.
type	String	either CONTAINER for directory or OBJECT for a file.
contentIndexedAt	String	Epoch seconds of the server time when the content of this file was last indexed.
metadataIndexedAt	String	Epoch seconds timestamp of the server time when the metadata of this file was last indexed.
group	String	the group that owns the file.
stub	String	True if the object is a smartlink file that's stubbed to a cloudpool.
blockSize	String	The blocksize of the file
mode	String	The octal mode of the file, according to standard posix file permissions.
language_s	String	The detected language of the file's content.
isHidden	String	True if this represents a hidden file.

Example 1: Run an empty query to get all of the file names and owners

```
curl -s -G -k -H Authorization:"Bearer
eyJhbGciOiJIUzI1NiJ9.eyJzdWIiOiJTSUQ6Uy0xLTUtMjEtMjAxODMyNTY2LTMxODczNTM0MDct
MjgyOTk5MTcxMC0xNDQ4Iiwicm9sZSI6IlVTRVIiLCJleHAiOjE1NzAyODk3MTR9.0vs-
tnOgs0cOBSYB_SOuNHdmV7NT6YisTwuIbZFMkJE" https://search.igls.com/graphql --
data-urlencode 'query={
  fileInfo {
        pageInfo {
            totalNum
        }
        fileList {
            fileName
            ownerName
        }
  }
}'
```

```
{
        "data": {
            "fileInfo": {
                "fileList": [
                    {
                        "fileName": "kf-demo-u1",
                        "ownerName": "root"
```

```
                    },
                    {
                        "fileName": "First Installation.txt",
                        "ownerName": "AD02\\kf-demo-u1"
                    },
                    {
                        "fileName": "Post Installation.txt",
                        "ownerName": "root"
                    },
                    {
                        "fileName": "Post Touch.txt",
                        "ownerName": "root"
                    }
                ],
                "pageInfo": {
                    "totalNum": 4
                }
            }
        }
}
```


Example 2: Paginate through *.txt files

Below are two curl requests for search results, the second one using the page info returned by the first call. In this example we're limiting the number of results returned on each page to 1 for clarity.

```
ccurl -s -G -k -H Authorization:"Bearer
eyJhbGciOiJIUzI1NiJ9.eyJzdWIiOiJTSUQ6Uy0xLTUtMjEtMjAxODMyNTY2LTMxODczNTM0MDct
MjgyOTk5MTcxMC0xNDQ4Iiwicm9sZSI6IlVTRVIiLCJleHAiOjE1NzAyODk3MTR9.0vs-
tnOgs0cOBSYB_SOuNHdmV7NT6YisTwuIbZFMkJE" https://search.igls.com/graphql --
data-urlencode 'query={
  fileInfo(rowsPerQuery:1, fileExt:"txt") {
        pageInfo {
            totalNum
            startCursor
            hasNextPage
        }
        fileList {
            fileName
            ownerName
        }
  }
}'
```


```
{
        "data": {
            "fileInfo": {
```

```
                "fileList": [
                    {
                        "fileName": "First Installation.txt",
                        "ownerName": "AD02\\kf-demo-u1"
                    }
                ],
                "pageInfo": {
                    "hasNextPage": true,
                    "startCursor": 1,
                    "totalNum": 3
                }
            }
        }
}
```


The call to get the second page of results uses `startCursor:1` since that was returned by the first page.

```
curl -s -G -k -H Authorization:"Bearer
eyJhbGciOiJIUzI1NiJ9.eyJzdWIiOiJTSUQ6Uy0xLTUtMjEtMjAxODMyNTY2LTMxODczNTM0MDct
MjgyOTk5MTcxMC0xNDQ4Iiwicm9sZSI6IlVTRVIiLCJleHAiOjE1NzAyODk3MTR9.0vs-
tnOgs0cOBSYB_SOuNHdmV7NT6YisTwuIbZFMkJE" https://search.igls.com/graphql --
data-urlencode 'query={
  fileInfo(startCursor:1, rowsPerQuery:1, fileExt:"txt") {
        pageInfo {
            totalNum
            startCursor
            hasNextPage
        }
        fileList {
            fileName
            ownerName
        }
  }
}'
```


```
{
        "data": {
            "fileInfo": {
                "fileList": [
                    {
                        "fileName": "Post Installation.txt",
                        "ownerName": "root"
                    }
                ],
                "pageInfo": {
                    "hasNextPage": true,
                    "startCursor": 2,
                    "totalNum": 3
```

```
                }
            }
        }
}
```

getSummaryReport

The main query to execute roll up searches. This uses Solr's faceting engine to group and summarize fields. (See https://lucene.apache.org/solr/guide/8_2/faceting.html for how faceting works.) This endpoint provides a simple abstraction on a few of the fields to roll up the queries and return summaries of the results. Filters that are applied in the fileInfo queries are also applicable here.

Query

Schema

```
getSummaryReport(
  content: String,
  fileExt: String,
  filePath: String,
  fileName: String,
  fileSizeMin: String,
  fileSizeMax: String,
  creationStart: String,
  creationEnd: String,
  lastAccessStart: String,
  lastAccessEnd: String,
  lastModifiedStart: String,
  lastModifiedEnd: String,
  stub: String,
  facetSortField: String,
  facetSortOrder: FacetSortOrder,
  facetTimeGroupBy: FacetTimeGroupBy,
  facetLimit: Int,
  facetField: String!,
  clientTimeZone: String
): SummaryReportList
```

Argument	Type	Value
content	String	Search for words in files. Accepts simple words and phrases, or more complicated Lucene queries, of the form "dog AND (cat OR bat)"
fileName	String	Search for files with this exact filename, not including the path, but including the extension. Accepts wildcards.
fileExt	String	Search for files with this file extension. The extension is defined as the set of characters after the last dot (.) in the

		filename.
lastModifiedStart	String	filter by documents that have a last modified timestamp later in time than the value that is set here. Dates must be formatted with the `YYYY-MM-DDThh:mm:ssZ` format. For example, 03:14 in the morning of March 14, 2020 is written `2020-03-14T03:14:15Z`
lastModifiedEnd	String	filter by documents that have a last modified timestamp earlier in time than the value that is set here. See lastModifiedStart argument above for syntax and example.
lastAccessStart	String	filter by documents that have a last accessed timestamp later in time than the value that is set here. atimes must be enabled on your PowerScale cluster. See lastModifiedStart argument above argument above for syntax and example.
lastAccessEnd	String	filter by documents that have a last accessed timestamp earlier in time than the value that is set here. atimes must be enabled on your PowerScale cluster. See lastModifiedStart argument above for syntax and example.
creationStart	String	filter by documents that have a created timestamp later in time than the value that is set here. See lastModifiedStart argument above for syntax and example.
creationEnd	String	filter by documents that have a created accessed timestamp earlier in time than the value that is set here. See lastModifiedStart for syntax and example.
filePath	String	Search for files that begin with the specified path, starting with /ifs
fileSizeMin	String	In bytes, the minimum file size to return in the search results.
fileSizeMax	String	In bytes, the maximum file size to return in the search results.
stub	String	True or False, limit the results to files that are SmartLinked files (for true), or files that are not SmartLinked files (for false).
facetSortField	String	The result field to sort on. Can be one of: - total - count - average

facetSortOrder	Object(enum)	`enum FacetSortOrder {`
		` ASC`
		` DESC`
		`}`
		Sort results ascending (ASC) or descending (DESC)

facetTimeGroupBy	Object(enum)	`enum FacetTimeGroupBy {`
		`DAY`
		`WEEK`
		`MONTH`
		`}`
		For time range based facet queries, group the results by day, week, or month.

facetLimit	Int	The number of groups to return.
facetField	String (required)	The field to use for faceting. Supported fields are: - owner - group - extension - createdtime - lastmodified - filename - clustername - stub
clientTimeZone	String	The offset in hours of the current time zone. For Eastern Standard time, enter “-4”.

Response

Schema

```
type SummaryReportList {
  pageInfo: PageInfo
  summaryList: [SummaryReport]
}
```

Field	Type	Value
pageInfo	Object (`PageInfo`)	Information on the query and on rows.
summaryList	List of Objects (`SummaryReport`)	A list of `SummaryReport` objects.

```
type PageInfo {
 totalNum: Long!
 qTime: Int
}
```

Field	Type	Value
totalNum (non-null)	Long	Total number of records that were matched by this query
qTime	Int	The time (in milliseconds) that the query took to return.

```
type SummaryReport {
  name: String
  count: Long
  total: Long
  average: Float
}
```

name	String	The name of this unique record. Value depends on the value that was chosen for faceting.
count	Long	The number of file records matching this name.
total	String	The summed value (in bytes) of all file records in this bucket.
average	String	The average filesize of all file records in this bucket.

Example 1: Get the counts and sizes of top two file extensions

Gets all of the extensions on the system, groups by extension, and sums the total bytes. Reports the data by default in descending count order.

```
curl -k -L -s -G -H Authorization:"Bearer
eyJhbGciOiJIUzI1NiJ9.eyJzdWIiOiJTSUQ6Uy0xLTUtMjEtMTk0MzI5MTcwNi0yNjMzMTY1NTY0
LTIyOTIzNDIyNS0xMzE4IiwicM9sZSI6IkFETUlOIiwiZXhwIjoxNTcwNjQ3NjQ5fQ.jzusrMyQdw
v2vRLY2_8_ER2UunjwRYXoRwOAPpDAHrg" https://172.25.1.101/graphql --data-
urlencode 'query={
  getSummaryReport(facetField:"extension") {
    pageInfo {
          totalNum
        }
    summaryList {
          name
          count
          total
          average
        }
  }
}'
```

```
{
    "data": {
        "getSummaryReport": {
            "pageInfo": {
                "totalNum": 2
            },
            "summaryList": [
                {
                    "average": 229351.0,
```

```
                    "count": 34170,
                    "name": "ZIP",
                    "total": 7836955408
                },
                {
                    "average": 380176.0,
                    "count": 32483,
                    "name": "txt",
                    "total": 12349288443
                }
            ]
        }
    }
}
```


Example 2: Groups all of the files on a path by last modified date

Accept a given path as a filter query. Group all of the files by last modified date, with monthly bucket results. Limit the results to a year's worth of data by adding start and end modified dates to the query.

```
curl -GkLs -H Authorization:"Bearer
eyJhbGciOiJIUzI1NiJ9.eyJzdWIiOiJTSUQ6Uy0xLTUtMjEtMTk0MzI5MTcwNi0yNjMzMTY1NTY0
LTIyOTIzNDIyNS0xMzE4Iiwicm9sZSI6IkFETUlOIiwiZXhwIjoxNTcwNjQ3NjQ5fQ.jzusrMyQdw
v2vRLY2_8_ER2UunjwRYXoRwOAPpDAHrg" https://172.25.1.101/graphql --data-
urlencode 'query={
  getSummaryReport(
    facetField:"lastmodified"
    facetTimeGroupBy: MONTH
    lastModifiedStart: "1539027461000"
    lastModifiedEnd: "1570563482000"
    filePath: "/ifs/data"
  ) {
    pageInfo {
          totalNum
        }
    summaryList {
          name
          count
          total
          average
        }
  }
}'
```


```
{
    "data": {
        "getSummaryReport": {
```

```
            "pageInfo": {
                "totalNum": 10
            },
            "summaryList": [
                {
                    "average": 127.0,
                    "count": 30444,
                    "name": "Fri Feb 01 00:00:00 UTC 2019",
                    "total": 3878443
                },
                {
                    "average": 0.0,
                    "count": 0,
                    "name": "Fri Mar 01 00:00:00 UTC 2019",
                    "total": 0
                },
                {
                    "average": 0.0,
                    "count": 0,
                    "name": "Mon Apr 01 00:00:00 UTC 2019",
                    "total": 0
                },
                {
                    "average": 0.0,
                    "count": 0,
                    "name": "Wed May 01 00:00:00 UTC 2019",
                    "total": 0
                },
                {
                    "average": 0.0,
                    "count": 0,
                    "name": "Sat Jun 01 00:00:00 UTC 2019",
                    "total": 0
                },
                {
                    "average": 0.0,
                    "count": 0,
                    "name": "Mon Jul 01 00:00:00 UTC 2019",
                    "total": 0
                },
                {
                    "average": 0.0,
                    "count": 0,
                    "name": "Thu Aug 01 00:00:00 UTC 2019",
                    "total": 0
                },
                {
                    "average": 0.0,
```

```
                    "count": 0,
                    "name": "Sun Sep 01 00:00:00 UTC 2019",
                    "total": 0
                },
                {
                    "average": 0.0,
                    "count": 0,
                    "name": "Tue Oct 01 00:00:00 UTC 2019",
                    "total": 0
                },
                {
                    "average": 0.0,
                    "count": 0,
                    "name": "Fri Nov 01 00:00:00 UTC 2019",
                    "total": 0
                }
            ]
        }
    }
}
```


© Superna LLC

9.23.1. Scripted API Searches with Search & Recover GraphQL API

Home Top

- Overview
- Scheduled Content Search with Email Alerts Example
  - Description:
  - Script:
- Scheduled File Count Monitoring in a Folder with email alerts Example
  - Description:
  - Script:
- Monitor a Directory for files being created in the last hour
  - Description:
  - Script:
- Setup a script to run on a Schedule with Cron

Overview

The API can be used to automate searches and apply logic to the results. Several examples are available below.

1. Search for credit card or SSN on newly added files and send an alert email on any hits.

2. Monitor a folder for a threshold number of files and send an alert email if the threshold is crossed.

Scheduled Content Search with Email Alerts Example

Description:

- To search file with specific path that contain specific Social Security Number (SSN) (i.e: filePath:"/ifs/data/search3/folder1",content:"333-44-5555")
- If found that file, send e-mail by utilizing mailx function on Eyeglass Search node.

Script:

1. Login to any Eyeglass Search node as ecaadmin user.
2. Create a script (Example: /home/ecadmin/contentsearchtrigger.sh).

Bash Script:
==========

```
#!/bin/bash
## declare mail variables
##email subject
subject="Eyeglass Search Found file with requested SSN"
##sending mail as
from="eyeglassSR@exampledomain.com"
## sending mail to
to="admin1@exampledomain.com"
## send carbon copy to
```

```
also_to="admin2@exampledomain.com"

## token
token=$(curl -s -G -k https://<eyelgassSR-node1-
IPaddress>/graphql --data-urlencode 'query={
   login(id:"searchuser@exampledomain.com", pass:"NotReal!")
{
     token
   }
 }'| awk -F'"' '{print $8}')

## Search
found=$(curl -s -Gk -H Authorization:"Bearer $token"
https://<eyelgassSR-node1-IPaddress>/graphql --data-
urlencode 'query={
  fileInfo(filePath:"/ifs/data/search3/folder1",content:"333
-44-5555") {
        pageInfo {
            totalNum
        }
  }
}'| tr -dc '0-9')

## check if found
if [[ "$found" > 0  ]]; then

        ## send email if file with requested SSN found
        echo -e "Found\n\nthe requested Social Security
Number" | mailx -s "$subject" -r "$from" -c "$to" "$also_to"
fi

exit 0
```

3. Change mode to executable:
   chmod +x contentsearchtrigger.sh
4. Run the script:
   ./contentsearchtrigger.sh
5. Check e-mail.
6. Then follow setup on a schedule instructions.

Scheduled File Count Monitoring in a Folder with email alerts Example

Description:

This example will count the number of files in a folder (Example: Path:"/ifs/data/search3/folder1"). If the number of files is greater than threshold, send e-mail by utilizing mailx function on Eyeglass Search node.

Script:

1. Login to any Eyeglass Search node as ecaadmin user
2. Create a script (Example: /home/ecadmin/foldermonitor.sh)
3. Requires modifications for emails and password to authenticate to the index
4. Can use eccaadmin user
5. The depth property is how many folder levels the search should use. This example /ifs/data/path1/path2 is a depth 5 to limit the search to files in path2 folder, level 5 depth means the folders equal a depth and the file name is also included in the depth value. Adjust the script to use a depth value equal to the path used in the search. See bolded

yellow highlight in the script to set the depth value of the search.

Note: This is not a recursive search and will not count files below the path in the search.

Bash Script:

==========

```
#!/bin/bash
## declare mail variables
##email subject
subject="Eyeglass Search Found the number of files is greater than threshold"
##sending mail as
from="eyeglassSR@exampledomain.com"
## sending mail to
to="admin1@exampledomain.com"
## send carbon copy to
also_to="admin2@exampledomain.com"

## token
token=$(curl -s -G -k https://<eyelgassSR-node1-IPaddress>/graphql --data-urlencode 'query={
login(id:"searchuser@exampledomain.com", pass:"NotReal!") {
token
}
}'| awk -F"" '{print $8}')
## Search. Set the depth of the search where each folder equals 1 level of depth in the example 5 means 5 folders in the path
found=$(curl -s -Gk -H Authorization:"Bearer $token" https://<eyelgassSR-node1-IPaddress>/graphql --data-urlencode 'query={
```

```
fileInfo(filePath:"/ifs/data/search3/folder1",content:'depth:5'") {
pageInfo {
totalNum
}
}
}'| tr -dc '0-9')

## check if greater than threshold (Replace the <Threshold with the actual threshold number>
if [[ "$found" > <Threshold> ]]; then
## send email if found greater than threshold
echo -e "Found\n\nthe number of files greater than threshold" | mailx -s "$subject" -r "$from" -c "$to" "$also_to"
fi

exit 0
```

==========

1. Change mode to executable
2. chmod +x foldermonitor.sh
3. Run the script
4. ./foldermonitor.sh
5. Check e-mail
6. Then follow setup on a schedule instructions.

Monitor a Directory for files being created in the last hour

Description:

This solution can assist with application workflows that expect files to be created in a directory on a regular basis. If no new files are created it may indicate the application process has crashed or has an issue writing data to the directory. This solution can use an hourly search to check for new files created in the last hour and runs on a schedule, to monitor for files > 0 or some other threshold and send an email alert.

1. To count the number of files created for the last 1 hour in a folder (i.e.: Path:"/ifs/data/search3/folder1") .
2. If the number of files is zero, send e-mail by utilizing mailx function on Eyeglass Search node.
3. Time format: Epoch .

Script:

1. Login to any Eyeglass Search node as ecaadmin user .
2. Create a script (Example: /home/ecadmin/monitornewfileslasthour.sh).

Bash Script:
==========

```
#!/bin/bash
## declare mail variables
##email subject
subject="Eyeglass Search detect 0 file was created during last 1 hour"
##sending mail as
from="eyeglassSR@exampledomain.com"
## sending mail to
to="admin1@exampledomain.com"
## send carbon copy to
also_to="admin2@exampledomain.com"

## token
token=$(curl -s -G -k https://<eyelgassSR-node1-IPaddress>/graphql -
-data-urlencode 'query={
login(id:"searchuser@exampledomain.com", pass:"NotReal!") {
token
}
}'| awk -F"" '{print $8}')

## Stat Time and End Time for search last 1 hour
starttime=$(date +%s%N -d "1 hour ago" | cut -b1-13)
endtime=$(date +%s%N | cut -b1-13)

## Search. Set the folder and depth of the search
found=$(curl -s -Gk -H Authorization:"Bearer ${token}"
https://<eyelgassSR-node1-IPaddress>/graphql --data-urlencode
"query={
```

```
fileInfo(filePath:\"/ifs/data/search3/folder1\",content:\"depth:5\",creation
Start:\"${starttime}\", creationEnd:\"${endtime}\") {
pageInfo {
totalNum
}
}
}"| tr -dc '0-9')

## check if number of file equal to zero
if [[ "$found" = 0 ]]; then
## send email if 0 file was created during last 1 hour
echo -e "Found\n\n0 file was created during last 1 hour" | mailx -s
"$subject" -r "$from" -c "$to" "$also_to"
fi

exit 0
```

==========

1. Change mode to executable.
2. chmod +x monitornewfileslasthour.sh
3. Run the script:
4. ./monitornewfileslasthour.sh
5. Check e-mail.
6. Then follow setup on a schedule instructions.

Setup a script to run on a Schedule with Cron

1. Login to node 1 as ecaadmin use examples below for hourly, daily or use a different cron tab string. Find examples here.
2. Type crontab -e .
3. Press letter i (to insert).
4. Copy and paste the job information into the editor.
5. Then press esc key.
6. Then press :
7. Followed by wq (write and quit).

Example cron entries:

1. Runs hourly:
    a. 0 * * * * $HOME/myscript.sh (note $HOME = /home/ecaadmin path)
2. Runs Daily at midnight:
1. 0 0 * * * $HOME/myscript.sh

© Superna LLC

9.24. Content Classification Feature Guide

Home Top

- Overview of the Document Classification System
  - How Classification and Tagging Works
- Classification Use Cases
- How to enable Classification Tagging without Content Indexing
  - How to Configure a Classification Rule
    - How to list classification rules
    - How to add a new classification rule
    - How to modify an existing classification rule
    - How to remove a classification rule
    - How to View the rules file
    - How to test a file with your Classification rules before production Role out
- Common Compliance Standards Overview
- PII and Sensitive PII Defined
- Security Classification Schema For PII
- How to use RegEx Expressions to Locate Compliance Data
  - Regex Expression Testing Tools
  - Basic Regex Syntax for Most Scenarios
  - Tested RegEx Expressions
  - Identify Executable Files
- Files with URL's

- Document Security Classification
- Name
- Address
    - Street Address
    - North American Street Address
    - German Street Address
- State, Region or Province
    - State, Region or Province Key Words
    - French Province
    - Italian Region or Province
    - Spain
    - UK
    - Canadian Province
    - US State
- US Zip Code
    - Telephone Number Key Words
    - French Telephone Number
    - Italian Telephone Number
    - Spanish Telephone Number
    - UK Telephone Number
    - North American Telephone Number
- Social Insurance Number
    - Social Insurance or Social Security Key Word

        - French Social Insurance Number
        - Italian Codice Fiscale
        - Spanish Social Insurance
    - NIE (Número de Identificación de Extranjero) for non Spanish Citizens
        - DNI (Documento Nacional de Identidad) for Spanish citizens
        - UK NINO
        - Canadian Social Insurance Number
        - US Social Security Number
        - Date of Birth
- Email
- IBAN: Europe Only
- Bank Account Numbers
    - Germany
    - France
    - Italy
    - Spain
- Bank Account Numbers: Canada and US only
    - Account Numbers Canada and US
- Bank Account Number US (10 - 12 digits) and Canada (7- 12 digits)
    - US
    - Canada

- Bank Routing Numbers
  - US Routing Number
  - Canadian Routing Numbers
- Routing and Bank Account Numbers
  - US
  - Canada
- Credit Card Number
  - ACH Format Compliance Check
  - Credit Card CVV
  - Expiration Date
- Medical Data
  - Medical Record Number (MRN)
  - Weight
  - Height

Overview of the Document Classification System

This guide covers different classification types that can be located in documents and the ability to configure the classification feature to read documents and match to classification types that are configured and apply a tag to the custom classification attribute in the search schema. This allows simple easy searches and scheduled reporting on compliance data.

This allows aggregate reports and file lists for data that has been tagged by the classification system

Once the tagging configuration is in place, full scan index can tag existing data and incremental mode can detect new documents or modified documents and scan the documents for compliance data. This feature automates the whole process and reduces the complexity and time spent classifying data.

This feature can be used with custom tags for internal document classification for any vertical industry that has specific document types. This system is flexible to tag documents based on regex syntax matching that supports substitution and wild card for letters numbers in documents.

How Classification and Tagging Works

1. The classification uses regex to process all the text that is provided by the content indexing engine when a full or incremental folder scan runs with content ingestion enabled in classification mode.
2. The text of each document is decoded and sent to the classification engine.
3. The classification engine is configured with regex rules below to locate compliance data and apply a tag to the document in the index.
4. The tag name is specified in the classification rule.
5. A document can match more than one classification rule and will have more than one tag applied.
6. The text of the document is not indexed or stored in the index which reduces the space needed in the index. **NOTE: no key word searches will be possible unless full content indexing is enabled in addition to classification mode.**

7. Once a document is tagged you can use the classification option in advanced search or quick reports to narrow your search results to documents that match a specific tag.
8. See example below, the list of tags will be available with the option to combine them with AND/OR search logic with multiple tags and can multi select tags.
9. 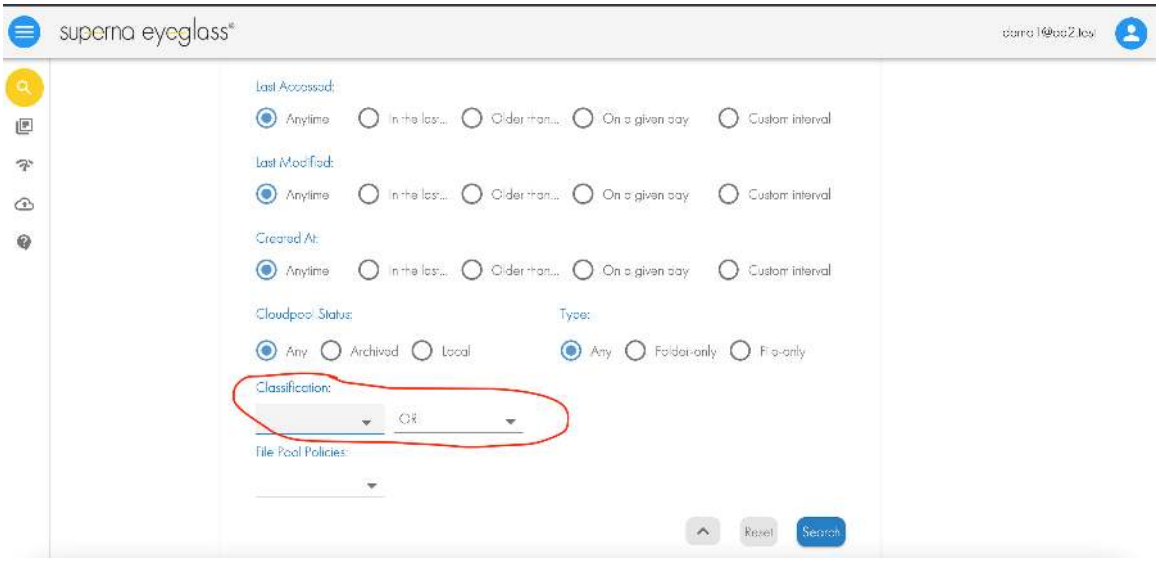

Classification Use Cases

1. New Search & Recover Deployment
    a. Index metadata and leverage content aware data classification without storing indexed content in the database. This option does not require expanding disk space to store classification tags. Disable full content index storage following steps here.
2. Existing Deployment with Content Indexing Enabled
    a. In the initial release of 1.1.5 tags can not be applied unless data is re-indexed and stored in the database. The presence of a classification tag configuration will trigger the re-index process to force a full re-index of content to be scanned for tags and the file will be re-added to the index even if the file on disk has not changed since it was last indexed. This will take time to re-

index. This requirement should be considered before using this feature.

How to enable Classification Tagging without Content Indexing

1. Use this option to enable classification when full content indexing is not required. This option has no additional disk space requirements other than the metadata index sizing guidelines.
2. Login to node 1
3. nano /opt/superna/eca/eca-env-common.conf
4. paste this to the file
5. export ADD_CONTENT_TO_INDEX=false
6. control+x to save
7. ecactl cluster down
8. ecactl cluster up
9. done

How to Configure a Classification Rule

1. Each classification rule is added to a single file and evaluated against the text of each file processed for content and classification processing.
2. Login to Search & Recover node 1 as ecaadmin
3. searchctl CLI commands will be added to edit this file and live update the ingestion engine to start using the new rules file. The below information is information to validate your rules file.
4. How to list classification rules

a. searchctl classification list

5. How to add a new classification rule

   a. searchctl classification add --name NAME --regex REGEX | [--regex-file /home/ecaadmin/regex.txt]

      i. --name - this is the name of the tag that will appear in the GUI when searching. Use single quotes at the beginning and end of the tag name.

      ii. --regex - this is the regex value used to match text for this tag. See examples in this guide or create your own regex matching string. Use single quotes at the beginning and end of the regex pattern.

      iii. (optional) --regex-file <path to file > This flag should be used for any complex regex with special characters that cannot be passed into the CLI command due to bash shell restrictions. The regex in the file can use any special characters needed.

   b. example

      i. searchctl classification add --name 'email' --regex '[\\w-\\.]+@([\\w-]+\\.)+[\\w-]{2,4}'

6. How to modify an existing classification rule

   a. searchctl classification modify --name NAME --regex REGEX

   b. Enter an existing tag name and new regex value to replace the matching logic.

7. How to remove a classification rule

    a. This command will start a job that searches the index for the files matching the tag and will place these files in a queue to re-index the files. This second re-index will no longer be tagged with the removed tag. This job can add a lot of files to be processed and will take time to remove them from the index. Once the command is executed the GUI will no longer display the tag for searching.

    b. searchctl classification remove --name NAME

        i. **NOTE: A job is launched to submit all matching files to be re-indexed, this will re-read the content of these files and run match rules for all remaining tags. This can generate read bandwidth to re-process these files and re-read all files to retag the files based on remaining data classification rules.**

8. How to View the rules file

    a. cat /opt/superna/eca/data/classifiers.json

```
{

"email":"[\\w-\\.]+@([\\w-]+\\.)+[\\w-]{2,4}",

"phone":"(\\+\\d{1,2}\\s)?\\(?\\d{3}\\)?[\\s.-]\\d{3}[\\s.-]\\d{4}",

"ipv4":"(([0-9]|[1-9][0-9]|1[0-9]{2}|2[0-4][0-9]|25[0-5])\\.){3}([0-9]|[1-9][0-9]|1[0-9]{2}|2[0-4][0-9]|25[0-5])"

}
```

How to test a file with your Classification rules before production Role out

1. This feature allow testing your classifiers easily with a single command
2. Example:
    a. searchctl classifications test --isilon <cluster name> --path /ifs/data/search/filetotest.txt
    b. Retrieves the file over the rest API.
    c. Tests the content of the file against all configured classifications
    d. Prints to the console which classifications matched and which string was found
    e. See example below that shows which tags matched and which strings in the file matched the regex
        i. 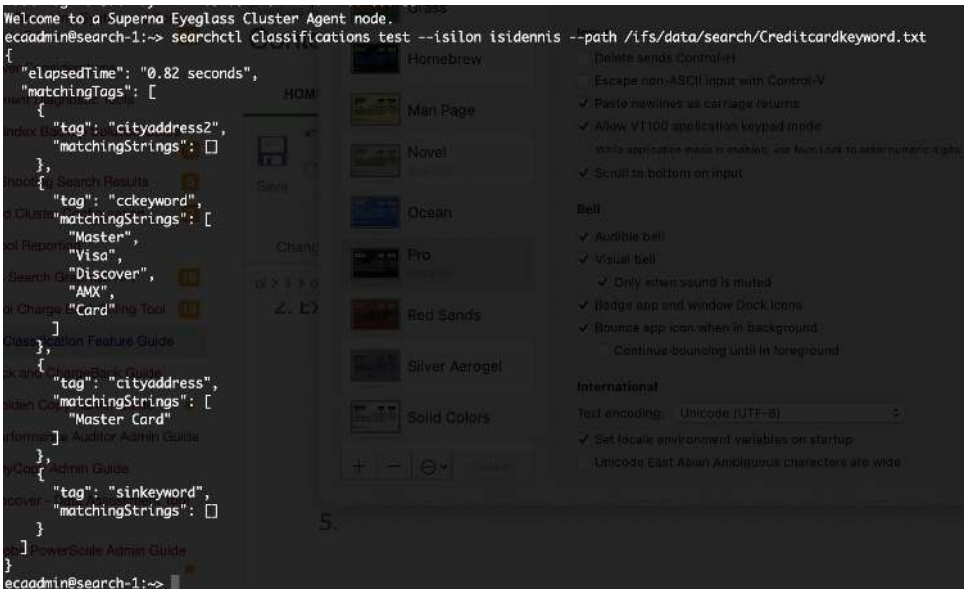

Common Compliance Standards Overview

Businesses are collecting and storing more and more individuals' PII. As a result, various jurisdictions have passed legislation to limit the use, distribution, and accessibility of PII, while allowing companies who need it to manage the data safely. GDPR in the European Union, HIPAA (Health Insurance Portability and Accountability Act) and PCI DSS (Payment Card Industry Data Security Standard) in the United States, along with state laws, other data breach laws, and other regulations control what defines PII.

The most common compliance standards covered below provides background on the key requirements to locating and protecting data that falls into one of these categories.

PCI DSS Data

PCI DSS provides standards for the processes and systems that merchants and vendors use to protect information. This information includes: Cardholder data such as the cardholder's name, the primary account number, and the card's expiration date and security code.

PHI Data:

PHI is any health information that can be tied to an individual, which under HIPAA means protected health information includes one or more of the following 18 identifiers. If these identifiers are removed the information is considered de-identified protected health information, which is not subject to the restrictions of the HIPAA Privacy Rule.

1. Names (Full or last name and initial)
2. All geographical identifiers smaller than a state, except for the initial three digits of a zip code if, according to the current publicly available data from the U.S. Bureau of the Census: the geographic unit formed by combining all zip codes with the same three initial digits contains more than 20,000 people; and the initial three digits of a zip code for all such geographic units containing 20,000 or fewer people is changed to 000
3. Dates (other than year) directly related to an individual
4. Phone Numbers
5. Fax numbers
6. Email addresses
7. Social Security numbers
8. Medical record numbers

9. Health insurance beneficiary numbers
10. Account numbers
11. Certificate/license numbers
12. Vehicle identifiers (including serial numbers and license plate numbers)
13. Device identifiers and serial numbers;
14. Web Uniform Resource Locators (URLs)
15. Internet Protocol (IP) address numbers
16. Biometric identifiers, including finger, retinal and voice prints
17. Full face photographic images and any comparable images
18. Any other unique identifying number, characteristic, or code except the unique code assigned by the investigator to code the data

GDPR Data:

The data subjects are identifiable if they can be directly or indirectly identified, especially by reference to an identifier such as a name, an identification number, location data, an online identifier or one of several special characteristics, which expresses the physical, physiological, genetic, mental, commercial, cultural or social identity of these natural persons. In practice, these also include all data which are or can be assigned to a person in any kind of way. For example, the telephone, credit card or personnel number of a person, account data, number plate, appearance, customer number or address are all personal data.

Classifying and Securing PII

PII and Sensitive PII Defined

PII: NIST 800-122 defines PII (Personally Identifiable Information) as any information that permits the identity of an individual to be directly or indirectly inferred, including any other information that is linked or linkable to that individual, regardless of whether the individual is a U.S. citizen, legal permanent resident, visitor to the U.S., or employee or contractor to a government agency or corporation.

Sensitive PII (SPII): According to the US Department of Homeland Security Handbook for Safeguarding Sensitive PII is Personally Identifiable Information, which if lost, compromised, or disclosed without authorization, could result in substantial harm, embarrassment, inconvenience, or unfairness to an individual. Sensitive PII includes:

- Social security numbers
- Bank account numbers
- Passport information
- Healthcare related information
- Medical insurance information
- Student information
- Credit and debit card numbers
- Drivers license and State ID information

Security Classification Schema For PII

In order to comply with the various restrictions in place covering the security of PII at least three levels of security classifications are recommended: Public, Private and Restricted.

1. **Public Data** is PII data that is in the "public domain", which includes data found in public records, public telephone directories, business directories, business cards, newspapers, social media platforms and websites. While anyone can access this data, access controls to prevent it's unauthorized use is recommended. If an individual does not want the data disclosed it should be classified as Private Data. If it's linkable to disassociated data such as name, address, social insurance, bank account, credit card and medical record data, it should be classified as Restricted Data.

2. **Private Data** is PII data an individual may not wish to disclose, such as date-of-birth, home address and phone number. Private data should be covered by a moderate level of protection. If it is linkable to disassociated data such as name, address, social insurance, bank account, credit card and medical record data. It should be classified as Restricted Data.

3. **Restricted Data** is sensitive PII (SPII), and is data which cannot obtain through legitimate means. Restricted data includes social insurance numbers, bank account, credit card

details, medical information, etc. Restricted data should be covered by the highest level of security controls.

How to use RegEx Expressions to Locate Compliance Data

The following are Regex expressions that will help to identify data as defined by PII, CPI, PHI and GDPR standards. All of these expressions have been tested using a mock database containing data for Germany, France, Italy, Spain, UK, Canada and US. The mock data Included:

1. Document Security Classification
2. Name
3. Address (Street, City, State or Province)
4. Postal Code
5. Telephone
6. Email
7. Date of Birth
8. Social Insurance Number
9. Bank Account (Account Number, Routing Number, IBAN Number)
10. Credit Card (number, CVV, Exp Date)
11. Medical Record Number
12. Blood Type, Height

Regex Expression Testing Tools

A regex is a string of text that allows you to create patterns that help match, locate, and manage text.
The https://regex101.com/ test site was used to test and validate all the following expressions:

Basic Regex Syntax for Most Scenarios

1. The following regex examples can be used for most matching scenarios. Use these examples to build your own rules.
    a. Case insensitive matching
        i. (?i)xxxxx - Using (?i) and then a string will match the string regardless of case
    b. Combining key words with OR logic
        i. xxxx | yyyy - Using the pipe command will match string xxxx OR yyyy
    c. Combining key words with AND logic
        i. xxxx AND yyyy - using upper case AND command will match string xxxx AND yyyy
    d. Numeric range matching
        i. Example 1 - matching a single number between 0 to 9
            1. [0-9]
        ii. Example 2 - Matching a 2 digit 10 to 99
            1. [1-9][0-9]
        iii. Example 1 - matching - 0 to 255
            1. [0-9]|[1-9][0-9]|1[0-9][0-9]|2[0-4][0-9]|25[0-5].
        iv. Examples of numerous number matching options
            1. **000..255:** ^([01][0-9][0-9]|2[0-4][0-9]|25[0-5])$
            2. **0 or 000..255:** ^([01]?[0-9]?[0-9]|2[0-4][0-9]|25[0-5])$
            3. **0 or 000..127:** ^(0?[0-9]?[0-9]|1[01][0-9]|12[0-7])$
            4. **0..999:** ^([0-9]|[1-9][0-9]|[1-9][0-9][0-9])$
            5. **000..999:** ^[0-9]{3}$

6. **0 or 000..999:** ^[0-9]{1,3}$
7. **1..999:** ^([1-9]|[1-9][0-9]|[1-9][0-9][0-9])$
8. **001..999:** ^(00[1-9]|0[1-9][0-9]|[1-9][0-9][0-9])$
9. **1 or 001..999:** ^(0{0,2}[1-9]|0?[1-9][0-9]|[1-9][0-9][0-9])$
10. **0 or 00..59:** ^[0-5]?[0-9]$
11. **0 or 000..366:** ^([012]?[0-9]?[0-9]|3[0-5][0-9]|36[0-6])$

Tested RegEx Expressions

The following Classification Tags and RegEx expressions have been tested using ssh prompt. They contain the necessary escape characters to work with bash prompt over ssh and can be copied and pasted to the ssh session to set up a classification rule. The data base used to test these expressions contained words in German, French, Italian, Spanish and English. German, French and Spanish languages contain accented (i.e é, ë, ê, è, ñ) and or special characters (i.e. German language ß). The following tested RegEx's do not include accommodation for accented or special characters. In a database search using these expressions data up to a accented character will be shown, the accented character and any letters following will not be displayed.

Identify Executable Files

Format: This Regular Expression will identify executable files in the format of <file name>.<File Extension>. The expression will identify the following executable file extensions:

BAT, BIN, CMD, COM, CPL, EXE, GADGET, INF1, INS, INX, ISU, JOB, JSE, LNK, MSC, MSI, MSP, MST, PAF, PIF, PS1, REG, RGS, SCR, SCT, SCT, SHB, SHS, U3P, VB, VBE, VBS, VBSCRIPT, WS, WSF, WSH, ACTION, APP, COMMAND, CSH, IPA, KSH, OSX, OUT, PRG, RUN, WORKFLOW

Identify Executable Files RegEx Expression: searchctl classification add --name 'executablefiles' --regex '(.*\.(BAT|BIN|CMD|COM|CPL|EXE|GADGET|INF1|INS|INX|ISU|JOB|JSE|LNK|MSC|MSI|MSP|MST|PAF|PIF|PS1|REG|RGS|SCR|SCT|SCT|SHB|SHS|U3P|VB|VBE|VBS|VBSCRIPT|WS|WSF|WSH|ACTION|APP|COMMAND|CSH|IPA|KSH|OSX|OUT|PRG|RUN|WORKFLOW))'

Tag: executablefiles

Key Words: Executable File

Keyword RegEx Expression: searchctl classifications add --name 'executablefilekeyword' --regex '((?i)(Executable|File|Files|Identify|Extension|Extensions))'

Phrase Example: Identify Executable File Extensions

Security: File extensions are not considered PII, PCI, PHI, or GDPR data

Files with URL's

Format: This Regular Expression can be used to find URL's in a document.

Search URL's RegEx Expression: searchctl classification add --name 'searchurl' --regex '(http|ftp|https):\/\/([\w_-]+(?:(?:\.[\w_-]+)+))([\w.,@?^=%&:\/~+#-]*[\w@?^=%&\/~+#-])'

Tag: searchurl

Key Words: Find, Search, URL, URL's

Keyword RegEx Expression: searchctl classifications add --name 'searchurlkeyword' --regex '((?i)(URL|URL's|Search|Identify|Find))'

Phrase Example: Find URL's, Identify URL

Security: URL's are not considered PII, PCI, PHI, or GDPR data

Document Security Classification

Format: string match case insensitive
Document Security Classification RegEx Expression: searchctl classifications add --name 'docclass' --regex '((?i)(Public|Private|Restricted|Confidential|Secret|Top|Unrestricted|Propritary|Sensitive|Unclassified|Tres secret defense|Secret defense|Confidentiel defense|Diffusion restreinte|STRENG GEHEIM|GEHEIM|VS-VERTRAULICH|VS-NUR FUR DEN DIENSTGEBRAUCH|Segretissimo|Segreto|Riservatissimo|Riservato|Secreto|Reservado|Confidencial|Difusion Limitada|OFFICIAL-SENSITIVE|OFFICIAL|Confidentiel|Protected|Protoge))'

Name

Format: First Last Name
Name RegEx Expression: searchctl classification add --name 'names' --regex '[a-zA-Z]+[\-"""\s]?[a-zA-Z]{1,40}'

Key Words: Name, Nombre, Nome, First Name, Last Name, Apellido, Cognome, Nom de famille, Familienname, Nachname

Name Keyword RegEx Expression: searchctl classification add --name 'nameskey' --regex '(Name|Nombre|Nome|First|Last|Middle|Zuerst|Letzte|Mitte|Milieu|D"""'abord|Durer|Primo|Prima|Medio|Media|Scorso|Scorsa|Ultimo|Ultima)'

Phrase Example: Customers Name, Patients Name, Kundenname, Nom du client, Nome del cliente, Nombre de las clientas
Security:

	Public Data	Private Data	Restricted Data
PII	If Public Domain Data, or stand alone (not	If data subject may not wish to	If linkable to disassociated data such as Social Insurance, bank

	likable to other data)	disclose	account, credit card data.
PCI			If linked with Credit Card data
PHI			Full or last name and initial
GDPR			If linkable to other data can lead to the identification of a particular person

Address

Street Address

Security:

	Public Data	**Private Data**	**Restricted Data**
PII	If stand alone or Public Domain Data (not likable to other data)	If data subject may not wish to disclose	If linkable to disassociated data such as Name, Social Insurance, bank account, credit card data.
PCI			If linked with Credit Card data
PHI			X
GDPR			If linkable to other data can lead to the identification of a particular person

North American Street Address

Format: Number followed by street name and designation i.e. 125 June Street
Street RegEx Expression: searchctl classification add --name 'nastreet' --regex '(?i)\d+(\s+\w+){1,}\s+(?:st(?:\.|reet)?|dr(?:\.|ive)?|pl(?:\.|ace)?|ave(?:\.|nue)?|rd|road|lane|ln.|drive|way|court|ct.|plaza|square|run|parkway|point|pike|square|driveway|trace|park|terrace|blvd)'

Key Words: Drive, drive, Street, street, St. Place, Lane, Ln., ln., Road, Rd., rd., Avenue, avenue, Ave., ave., Circle, Court, court, Ct., ct., Way, Plaza, Square, square, Run, run, Parway, parkway, Pkw., pkw., Point, point, Pike, pike, driveway, trace, Park, park, terrace, Terrace, blvd

North American Street Address Keyword Regex Expression: searchctl classification add --name 'nastreetkeyword' --regex '((?i)(Drive|drive|Street|street|Place|Lane|Ln.|ln.|Road|Rd.|rd.|Avenue|avenue|Ave.|ave.|Circle|Court|court|Ct.|ct.|Way|Ave|Plaza|Square|square|Run|run|Parway|parkway|Pkw.|pkw.|Point|point|Pike|pike|driveway|trace|Park|park|terrace|Terrace|blvd))'

German Street Address

Format: Street Name followed by Number i.e. Kieler Strasse 89
French Street Name RegEx Expression: searchctl classification add --name 'gstreet' --regex '([a-zA-Z]+\s*-?[a-zA-Z]+\s*-?[a-zA-Z]+\s*\d{1,3})'

Key Words: Weg, Straße, Allee, Mitte

German Street Address Keyword Regex Expression: searchctl classification add --name 'gstreetkeyword' --regex '(Weg|Strasze|Allee|Mitte)'

French Street Address

Format: Number followed by street designation (i.e. rue, quai) followed be name. i.e. 48 rue des Lacs
French Street Name RegEx Expression: searchctl classification add --name 'frstreet' --regex '((\d{1,3}),?\s\b(cours|rue|Rue|boulevard|Avenue|avenue|Quai|route)\s([a-zA-Z]{1,15})\s?([a-zA-Z]{1,15})\s?([a-zA-Z]{1,15})\s?([a-zA-Z]{1,15}))'

Key Words: cours, rue, Rue, boulevard, Avenue, avenue, Quai, route

French Street Address Keyword Regex Expression: searchctl classification add --name 'frstreetkeyword' --regex '((?i)(cours|rue|Rue|boulevard|Avenue|avenue|Quai|route))'

Italian Street Address

Format: Street designation (i.e. Via, corso) followed be name, followed by number i.e. Piazza Trieste e Trento 136
Italian Street Name RegEx Expression: searchctl classification add --name 'itstreet' --regex '(Via|Corso|Viale|Piazza|Largo) ([a-zA-Z]+?),? (\d+)'

Key Words: Via, Corso, Viale, Piazza, Largo

Italian Street Address Keyword Regex Expression: searchctl classification add --name 'itstreetkeyword' --regex '((?i)(Via|Corso|Viale|Piazza|Largo))'

Spanish Street Address

Format: Street name, followed by number i.e. Rio Segura 47
Spanish Street Name RegEx Expression: searchctl classification add --name 'spstreet' --regex '([a-zA-Z./]+\s*[a-zA-Z]+\s*[a-zA-Z]+\s*[a-zA-Z]+\s*\d{1,3})'

Key Words: Calle

Spanish Street Address Keyword Regex Expression: searchctl classification add --name 'spstreet' --regex '(Calle|calle)'

United Kingdom Street Address

Note: UK Street Expression is the same as the North American Expression
Format: Number followed by street name and designation i.e. 92 Great Western Road
UK Street RegEx Expression: searchctl classification add --name 'ukstreet' --regex '(\d{1,5}\s(\w+)?([a-zA-Z]+)(\.)?\s?([a-zA-Z]+)\s?\-?([a-zA-Z]+)\s?(Drive|drive|Street|street|Place|Lane|Ln.|ln.|Road|Rd.|rd.|Avenue|avenue|Ave.|ave.|Circle|Court|court|Ct.|ct.|Way|Ave|Plaza|Square|square|Run|run|Parway|parkway|Pkw.|pkw.|Point|point|Pike|pike|driveway|trace|Park|park|terrace|Terrace|blvd)?)'

Key Words: Drive, drive, Street, street, St. Place, Lane, Ln., ln., Road, Rd., rd., Avenue, avenue, Ave., ave., Circle, Court, court, Ct., ct., Way, Plaza, Square, square, Run, run, Parway, parkway, Pkw., pkw., Point, point, Pike, pike, driveway, trace, Park, park, terrace, Terrace, blvd

UK Street Address Keyword Regex Expression: searchctl classification add --name 'ukstreetkeyword' --regex '((?i)(Drive|drive|Street|street|Place|Lane|Ln.|ln.|Road|Rd.|rd.|Avenue|avenue|Ave.|ave.|Circle|Court|court|Ct.|ct.|Way|Ave|Plaza|Square|square|Run|run|Parway|parkway|Pkw.|pkw.|Point|point|Pike|pike|driveway|trace|Park|park|terrace|Terrace|blvd))'

Street Address Key Words (North America, Germany, France, Italy, Spain, UK)

Key Words: Drive, drive, Street, street, St. Place, Lane, Ln., ln., Road, Rd., rd., Avenue, avenue, Ave., ave., Circle, Court, court, Ct., ct., Way, Plaza, Square, square, Run, run, Parway, parkway, Pkw., pkw., Point, point, Pike, pike, driveway, trace, Park, park, terrace,

Terrace, blvd, Weg, Straße, Allee, Mitte, cours, rue, Rue, boulevard, Avenue, avenue, Quai, route, Via, Corso, Viale, Piazza, Largo, Calle, Address, Addresse, Indirizzo, Habla a, Dirección

Street Address Keyword RegEx Expression: searchctl classification add --name 'streetkeyword' --regex '((?i)(Drive|drive|Street|street|Place|Lane|Ln.|ln.|Road|Rd.|rd.|Avenue|avenue|Ave.|ave.|Circle|Court|court|Ct.|ct.|Way|Ave|Plaza|Square|square|Run|run|Parway|parkway|Pkw.|pkw.|Point|point|Pike|pike|driveway|trace|Park|park|terrace|Terrace|blvd|Weg|Strasze|Allee|Mitte|rue|Rue|cours|rue|boulevard|Quai|route|strada|Via|Corso|Viale|Piazza|Largo|Strada|Calle|calle|Address|Addresse|Indirizzo|Habla a|Direccion))'
Phrase Example: Street address, Adresse de rue, Dirección

City Name

Format: City Name
City RegEx Expression: searchctl classifications add --name 'city' --regex '(?:[A-Z][a-z.-]+[ ]?)+'

Key Words: City, Stadt, Ville, Città, Ciudad

City Name Keyword RegEx Expression: searchctl classifications add --name 'citykeyword' --regex '(City|Citta|Ciudad|Ville|Stadt).*$

Phrase Example: City address, Stadtadresse, Adresse de la ville, Indirizzo della città, Dirección de la ciudad
Security: If a stand alone element no security restrictions required, only if in conjunction with name and street address the following security restrictions apply:

	Public Data	Private Data	Restricted Data
PII	If stand alone or Public Domain Data (not likable to other data)	If data subject may not wish to disclose	If linkable to disassociated data such as Name, Street and State Address, Social Insurance, bank account, credit card data.
PCI	If stand alone or Public Domain Data (not likable to other data)		If linked with Credit Card data

PHI	If stand alone or Public Domain Data (not likable to other data)		X
GDPR	If stand alone or Public Domain Data (not likable to other data)		If linkable to other data can lead to the identification of a particular person

State, Region or Province

State, Region or Province Key Words

Key Words: State, Zustand,Province, Provincia,Region
Keyword RegEx Expression: searchctl classifications add --name 'srpkeyword' --regex '(State|Zustand|Province|Provincia|Region)'

German States

Format: 2 Letter abbreviation or name of German State
German State RegEx Expression: searchctl classifications add --name 'gstate' --regex '\b(?:BW|Baden-Wurttemberg|BY|Bavaria|Bayern|BE|Berlin|BB|Brandenburg|HB|Bremen|HH|Hamburg|HE|Hesse|Hessen|NI|Lower Saxony|MV|Mecklenburg-Vorpommern|NW|North Rhine-Westphalia|Nordrhein-Westfalen|RP|Rhineland-Palatinate|SL|Saarland|SN|Saxony|Sachsen|ST|Saxony-Anhalt|Sachsen-Anhalt|SH|Schleswig-Holstein|TH|Thuringia|Thuringen|Freistaat Bayern|Rheinland-Pfalz|Nordrhein-Westfalen|Niedersachsen|Zustand)\b'

Key Words: Zustand

State Region or Province Keyword RegEx Expression: searchctl classifications add --name 'srpkeyword' --regex '(State|Zustand|Province|Provincia|Region)'
Phrase Example: Zustand
Security: If a stand alone element no security restrictions required, only if in conjunction with name and street address the following security restrictions apply:

	Public Data	Private Data	Restricted Data
PII	If stand alone or Public Domain Data (not likable to other data)	If data subject may not wish to disclose	If linkable to disassociated data such as Name, address, Social Insurance, bank account, credit card data.

PCI	If stand alone or Public Domain Data (not likable to other data)		If linked with Credit Card data
PHI	If stand alone or Public Domain Data (not likable to other data)		X
GDPR	If stand alone or Public Domain Data (not likable to other data)		If linkable to other data can lead to the identification of a particular person

French Province

Format: 3 letter abbreviation or Province name.
French Province RegEx Expression: searchctl classifications add --name 'fstate' --regex '(Aquitaine|IDF|Ile-de-France|CBL|Centre|Basse-Normandie|Lorraine|Bretagne|Picardie|Languedoc-Roussillon|Rhone-Alpes|BRE|Brittany|NOR|Normandy|HDF|Hauts-de-France|Grand Est|Pays de la Loire|Centre-Val de Loire|BFC|Bourgogne-Franche-Comte|NAQ|Nouvelle-Aquitaine|ARA|Auvergne-Rhone-Alpes|OCC|Occitanie|Provence-Alpes-Cote d"""Azur|COR|Corsica)\b'

Key Words: Province

Sate Region or Province Keyword RegEx Expression: searchctl classifications add --name 'srpkeyword' --regex '(State|Zustand|Province|Provincia|Region)'

Phrase Example: Province
Security: If a stand alone element no security restrictions required, only if in conjunction with name and street address the following security restrictions apply:

	Public Data	**Private Data**	**Restricted Data**
PII	If stand alone or Public Domain Data (not likable to other data)	If data subject may not wish to disclose	If linkable to disassociated data such as Name, address, Social Insurance, bank account, credit card data.
PCI	If stand alone or Public Domain Data (not likable to other		If linked with Credit Card data

	data)		
PHI	If stand alone or Public Domain Data (not likable to other data)		X
GDPR	If stand alone or Public Domain Data (not likable to other data)		If linkable to other data can lead to the identification of a particular person

Italian Region or Province

Format: 2 Letter abbreviation or Region or Province Name (Note, Region is like a US State, and Province is like a US county)
Italian Region RegEx Expression: searchctl classifications add --name 'istate' --regex '\b(Provincia|ABR|Abruzzo|BAS|Basilicata|CAM|Campania|CAL|Calabria|CAM|Campania|EMI|Emilia-Romagna|FRI|Friuli-Venezia Giulia|LAZ|Lazio|LIG|Liguria|LOM|Lombardia|MAR|Marche|MOL |Molise|PIE |Piemonte|PUG |Puglia|SAR |Sardegna|SIC|Sicilia|Südtirol|TOS|Toscana|TRE|Trentino-Alto Adige|UMB|Umbria|VAL|Valle d"""'Aosta|VEN|Veneto|PD|Padova|VA|Varese|BS|Brescia|MI|Milano|PA|Palermo|AN|Ancona|CS|Cosenza|CN|Cuneo|VV|Vibo Valentia|FI|Firenze|OR|Oristano|TN|Trento|CH|Chieti|AP|Ascoli Piceno|VR|Verona|AT|Asti|MS|Massa|VC|Vercelli)\b'

Key Words: Provincia, Regione

State Region or Province Keyword RegEx Expression: searchctl classifications add --name 'srpkeyword' --regex '(State|Zustand|Province|Provincia|Region)'
Phrase Example: Provincia, Regione
Security: If a stand alone element no security restrictions required, only if in conjunction with name and street address the following security restrictions apply:

	Public Data	**Private Data**	**Restricted Data**
PII	If stand alone or Public Domain Data (not likable to other data)	If data subject may not wish to disclose	If linkable to disassociated data such as Name, address, Social Insurance, bank account, credit card data.
PCI	If stand alone or Public Domain Data (not likable to other data)		If linked with Credit Card data

PHI	If stand alone or Public Domain Data (not likable to other data)		X
GDPR	If stand alone or Public Domain Data (not likable to other data)		If linkable to other data can lead to the identification of a particular person

Spain

Format: Province Name or 1 or 2 letter abbreviation
Spanish Province RegEx Expression:searchctl classifications add --name 'spstate' --regex '\b(A Coruna|Araba|VI|Alava|Albacete|A|Alicante|Alacant|Almeria|O|Asturias|Avila|Badajoz|Balears|B|Barcelona|Bizkaia|Biscay|Burgos|Cacer|Cadiz|Cantabria|CS|Castellon|Castello|Ciudad Real|Cordoba|Cuenca|Gipuzkoa|Girona|Granada|Guadalajara|Huelva|HU|Huesca|Huca|J|Jaen|LO|La Rioja|Las Palmas|Leon|Lleida|Lugo|M|Madrid|MA|Malaga|Murcia|NA|Navarre|Nafarro|OR|Ourense|Palencia|Pontevedra|SA|Salamanca|Santa Cruz de Tenerife|Segovia|Sevilla|SE|Seville|Soria|Tarragona|Teruel|Toledo|V|Valencia|VA|Valladolid|Zamora|Zaragoza|Provincia)\b'

Key Words: Provincia

State Region or Province Keyword RegEx Expression: searchctl classifications add --name 'srpkeyword' --regex '(State|Zustand|Province|Provincia|Region)'
Phrase Example: Provincia
Security: If a stand alone element no security restrictions required, only if in conjunction with name and street address the following security restrictions apply:

	Public Data	Private Data	Restricted Data
PII	If stand alone or Public Domain Data (not likable to other data)	If data subject may not wish to disclose	If linkable to disassociated data such as Name, address, Social Insurance, bank account, credit card data.
PCI	If stand alone or		If linked with Credit Card data

	Public Domain Data (not likable to other data)		
PHI	If stand alone or Public Domain Data (not likable to other data)		X
GDPR	If stand alone or Public Domain Data (not likable to other data)		If linkable to other data can lead to the identification of a particular person

UK

Format: Name
RegEx Expression: searchctl classifications add --name 'ukcountry' --regex '\b(Great Britain|England|Scotland|Wales|Northern Ireland|Western Isles)\b'

Key Words: Country or Region

State Region or Province Keyword RegEx Expression: searchctl classifications add --name 'srpkeyword' --regex '(State|Zustand|Province|Provincia|Region)'
Phrase Example: Country or Region
Security: If a stand alone element no security restrictions required, only if in conjunction with name and street address the following security restrictions apply:

	Public Data	**Private Data**	**Restricted Data**
PII	If stand alone or Public Domain Data (not likable to other data)	If data subject may not wish to disclose	If linkable to disassociated data such as Name, address, Social Insurance, bank account, credit card data.
PCI	If stand alone or Public Domain Data (not likable to other		If linked with Credit Card data

	data)		
PHI	If stand alone or Public Domain Data (not likable to other data)		X
GDPR	If stand alone or Public Domain Data (not likable to other data)		If linkable to other data can lead to the identification of a particular person

Canadian Province

Format: Province Name or Abbreviation
Canadian Province RegEx Expression: searchctl classifications add --name 'caprov' --regex '\b(AB|ALB|Alta|Alberta|BC|CB|British Columbia|LB|Labrador|MB|Man|Manitoba|N[BLTSU]|Nfld|NF|Newfoundland|NWT|Northwest Territories|Nova Scotia|New Brunswick|Nunavut|ON|ONT|Ontario|PE|PEI|IPE|Prince Edward Island|QC|PC|QUE|QU|Quebec|SK|Sask|Saskatchewan|YT|Yukon|Yukon Territories)\b'

Key Words: Province

State Region or Province Keyword RegEx Expression: searchctl classifications add --name 'srpkeyword' --regex '(State|Zustand|Province|Provincia|Region)'
Phrase Example: Province
Security: If a stand alone element no security restrictions required, only if in conjunction with name and street address the following security restrictions apply:

	Public Data	**Private Data**	**Restricted Data**
PII	If stand alone or Public Domain Data (not likable to other data)	If data subject may not wish to disclose	If linkable to disassociated data such as Name, address, Social Insurance, bank account, credit card data.
PCI	If stand alone or Public Domain Data (not likable to other data)		If linked with Credit Card data

PHI	If stand alone or Public Domain Data (not likable to other data)		X
GDPR	If stand alone or Public Domain Data (not likable to other data)		GDPR does not apply to Canadian citizens living in the Canada

US State

Format: State Name or Abbreviation
US State RegEx Expression: searchctl classifications add --name 'usstate' --regex '(?:Ala(?:(?:bam|sk)a)|Arizona|Arkansas|California|Colorado|Connecticut|Delaware|District of Columbia|Florida|Georgia|Hawaii|Idaho|Illinois|Indiana|Iowa|Kansas|Kentucky|Louisiana|Maine|Maryland|Massachusetts|Michigan|Minnesota|Miss(?:(?:issipp|our)i)|Montana|Nebraska|Nevada|New (?:Hampshire|Jersey|Mexico|York)|North (?:(?:Carolin|Dakot)a)|Ohio|Oklahoma|Oregon|Pennsylvania|Rhode Island|South (?:(?:Carolin|Dakot)a)|Tennessee|Texas|Utah|Vermont|Virginia|Washington|West Virginia|Wisconsin|Wyoming|A[KLRZ]|C[AOT]|D[CE]|FL|GA|HI|I[ADLN]|K[SY]|LA|M[ADEINOST]|N[CDEHJMVY]|O[HKR]|PA|RI|S[CD]|T[NX]|UT|V[AT]|W[AIVY])'

Key Words: State

State Region or Province Keyword RegEx Expression: searchctl classifications add --name 'srpkeyword' --regex '(State|Zustand|Province|Provincia|Region)'
Phrase Example: State name
Security: If a stand alone element no security restrictions required, only if in conjunction with name and street address the following security restrictions apply:

	Public Data	Private Data	Restricted Data
PII	If stand alone or Public Domain Data (not likable to other data)	If data subject may not wish to disclose	If linkable to disassociated data such as Name, address, Social Insurance, bank account, credit card data.
PCI	If stand alone or		If linked with Credit Card data

	Public Domain Data (not likable to other data)		
PHI	If stand alone or Public Domain Data (not likable to other data)		X
GDPR	GDPR does not apply to US citizens living in the US		GDPR does not apply to US citizens living in the US

US Zip Code

Format: 5 digit code (xxxxx), or 5 plus 4 digit code (xxxxx-xxxx)
City RegEx Expressions:

searchctl classification add --name 'zip' --regex '([0-9]{5})'

searchctl classification add --name 'zip9' --regex '([0-9]{5}-[0-9]{4})'

Key Words: Zip Code, Postal Code

City Name Keyword RegEx Expression: searchctl classifications add --name 'zipkeyword' --regex '((?i)(Zip|Postal|Code))'

Phrase Example: City address, Stadtadresse, Adresse de la ville, Indirizzo della città, Dirección de la ciudad
Security: If a stand alone element no security restrictions required, only if in conjunction with name and street address the following security restrictions apply:

	Public Data	**Private Data**	**Restricted Data**
PII	If stand alone or Public Domain Data (not likable to other data)	If data subject may not wish to disclose	If linkable to disassociated data such as Name, Street and State Address, Social Insurance, bank account, credit card data.
PCI	If stand alone or Public Domain Data (not likable to other		If linked with Credit Card data

	data)		
PHI	If stand alone or Public Domain Data (not likable to other data)		X
GDPR	If stand alone or Public Domain Data (not likable to other data)		If linkable to other data can lead to the identification of a particular person

Telephone Number

Telephone Number Key Words

Key Words:
Phone,Telephone,Number,Numero,Número,Mobile,Telefonnummer,Numéro,Telefono,Telefon,Téléphone,Teléfono,Handynummer,Poratble,Cell,Cellular,Cellulare|móvil

Telephone Number Keyword RegEx Expression: searchctl classification add --name 'phonekeyword' --regex '((?i)(Phone|Telephone|Number|Mobile|Telefonnummer|Telefono|Telefon|Telephone|Telefono|Handynummer|Poratble|Cell|Cellular|Cellulare|movil|Numero|de|telephone|portable))'

German Telephone Numbers

Formats:
6 2 2 (XXXXXX XX XX)
5 2 2 2 (XXXXX XX XX XX)
4 2 2 2 (XXXX XX XX XX)
RegEx Expression: searchctl classification add --name 'germanphone' --regex '([+][0-9]{1,3}[.-])?([(]{1}[0-9]{1,6}[)])?([0-9 .-\/]{3,20})'

Key Words: Telefonnummer, Telephon, Handynummer

German Telephone Number Keyword RegEx Expression: searchctl classification add --name 'dephonekeyword' --regex '((?i)(Telefonnummer|Telefon|Handynummer))'
Phrase Example: Telefonnummer
Security: If a stand alone element no security restrictions required, only if in conjunction with name and street address the following security restrictions apply:

	Public Data	Private Data	Restricted Data
PII	If Public Domain Data	If data subject may not wish to disclose	If linkable to disassociated data such as Name, address, Social Insurance, bank account, credit card data.
PCI			If linked with Credit Card data
PHI			X
GDPR			X

French Telephone Number

Format: XX XX XX XX
RegEx Expression: searchctl classification add --name 'frenchphone' --regex '\d{10}|\+33\d{9}|\+33\s\d{1}\s\d{2}\s\d{2}\s\d{2}\s\d{2}|\d{2}\s\d{2}\s\d{2}\s\d{2}\s\d{2}'

Key Words: Numéro de téléphone, Numéro de portable

French Telephone Number Keyword RegEx Expression: searchctl classification add --name 'frphonekeyword' --regex '((?i)(Numero|de|telephone|portable))'

Phrase Example: Numéro de téléphone, Numéro de portable
Security: If a stand alone element no security restrictions required, only if in conjunction with name and street address the following security restrictions apply:

	Public Data	Private Data	Restricted Data
PII	If Public Domain Data (not likable to other data)	If data subject may not wish to disclose	If linkable to disassociated data such as Name, address, Social Insurance, bank account, credit card data.
PCI			If linked with Credit Card data
PHI			X
GDPR			X

Italian Telephone Number

Land Line
Format: XXXX XXXXXXX
RegEx Expression: searchctl classification add --name 'itphone' --regex '([0][1-9][0-9]{1,2}\s?[0-9]{6,11})'

Mobile Phone
Format:3xx xxxxxxxx or 3xxxxxxxxx
RegEx Expression: searchctl classification add --name 'itmobilephone' --regex '([0][1-9][0-9]{1,2}\s?[0-9]{6,11})'

Key Words: Numero di telefono, Numero di cellulare

ItalianTelephone Number Keyword RegEx Expression: searchctl classification add --name 'itphonekeyword' --regex '((?i)(Numero|telefono|de|cellulare))'

Phrase Example: Numero di telefono, Numero di cellulare
Security: If a stand alone element no security restrictions required, only if in conjunction with name and street address the following security restrictions apply:

	Public Data	Private Data	Restricted Data
PII	If stand alone or Public Domain Data (not likable to other data)	If data subject may not wish to disclose	If linkable to disassociated data such as Name, Social Insurance, bank account, credit card data.
PCI			If linked with Credit Card data
PHI			X
GDPR			X

Spanish Telephone Number

Format: XXX XXX XXX
RegEx Expression: searchctl classification add --name 'spanishphone' --regex '[6,7]\d{2}\s\d{3}\s\d{3}'

Key Words: Número de teléfono, Número de teléfono móvil

Spanish Telephone Number Keyword RegEx Expression: searchctl classification add --name 'spphonekeyword' --regex '((?i)(Numero|de|telefono|movil))'

Phrase Example: Número de teléfono
Security: If a stand alone element no security restrictions required, only if in conjunction with name and street address the following security restrictions apply:

	Public Data	Private Data	Restricted Data
PII	If stand alone or Public Domain Data (not likable to other data)	If data subject may not wish to disclose	If linkable to disassociated data such as Name, Social Insurance, bank account, credit card data.
PCI			If linked with Credit Card data
PHI			X
GDPR			X

UK Telephone Number

Format: XXX XXXX XXXX
RegEx Expression: searchctl classification add --name 'ukphone' --regex '\s*\(?(0|\+44)(\s*|-)\d{4}\)?(\s*|-)\d{3}(\s*|-)\d{3}\s*|(\s*\(?(0|\+44)(\s*|-)\d{3}\)?(\s*|-)\d{3}(\s*|-)\d{4}\s*)|(\s*\(?(0|\+44)(\s*|-)\d{2}\)?(\s*|-)\d{4}(\s*|-)\d{4}\s*)|(\s*(7|8)(\d{7}|\d{3}(\-|\s{1})\d{4})\s*)|(\s*\(?(0|\+44)(\s*|-)\d{3}\s\d{2}\)?(\s*|-)\d{4,5}\s*)'

Key Words: Telephone Number, Mobile Number, Cellular Number, Cell Number

UK Telephone Number Keyword RegEx Expression: searchctl classification add --name 'ukphonekeyword' --regex '((?i)(Telephone|Number|Mobile|Cellular|Cell))'

Phrase Example: Telephone Number, Mobile Number, Cellular Number, Cell Number
Security: If a stand alone element no security restrictions required, only if in conjunction with name and street address the following security restrictions apply:

	Public Data	Private Data	Restricted Data
PII	If stand alone or Public Domain Data (not likable to other data)	If data subject may not wish to disclose	If linkable to disassociated data such as Name, Social Insurance, bank account, credit card data.
PCI			If linked with Credit Card data
PHI			X
GDPR			X

North American Telephone Number

Format: XXX-XXX-XXXX
RegEx Expressions:

searchctl classifications add --name 'naphone' --regex '(\+\d{1,2}\s)?\(?\d{3}\)?[\s.-]\d{3}[\s.-]\d{4}'

searchctl classification add --name 'usaphone' --regex '(?:\+?1[-.]?)?\(?([0-9]{3})\)?[-.]?([0-9]{3})[-.]?([0-9]{4})'

Key Words: Telephone Number, Mobile Number, Cellular Number, Cell Number

North America Telephone Number Keyword RegEx Expression: searchctl classification add --name 'naphonekeyword' --regex '((?i)(Telephone|Number|Mobile|Cellular|Cell))'

Phrase Example: Telephone Number, Mobile Number, Cellular Number, Cell Number
Security: If a stand alone element no security restrictions required, only if in conjunction with name and street address the following security restrictions apply:

	Public Data	Private Data	Restricted Data
PII	If stand alone or Public Domain Data (not likable to other data)	If data subject may not wish to disclose	If linkable to disassociated data such as Name, Social Insurance, bank account, credit card data.
PCI			If linked with Credit Card data
PHI			X
GDPR			GDPR does not apply to Canadian or US citizens living in Canada or the US

North American Fax Number

Format: XXX-XXX-XXXX
RegEx Expressions:

searchctl classification add --name 'fax' --regex '\+?\d{0,3}[- (]{0,2}?\d{3}[-)]{0,2}?\d{3}[-)]{0,2}?\d{4}?'

Key Words: Fax, Fax Number

North America Telephone Number Keyword RegEx Expression: searchctl classification add --name 'faxkeyword' --regex '((?i)(Fax|Number))'

Phrase Example: Fax, Fax Number
Security: If a stand alone element no security restrictions required, only if in conjunction with name and street address the following security restrictions apply:

	Public Data	Private Data	Restricted Data
PII	If stand alone or Public Domain Data (not likable to other data)	If data subject may not wish to disclose	If linkable to disassociated data such as Name, Social Insurance, bank account, credit card data.
PCI			If linked with Credit Card data
PHI			X
GDPR			GDPR does not apply to Canadian or US citizens living in Canada or the US

Social Insurance Number

Social Insurance or Social Security Key Word

Key Words: Social Security, Pension Insurance, Rentenversicherung, SIN, Social Code, Sozialgesetzbuch (SGB), INSEE, Codice fiscale, NIE, Número de Identificación de Extranjero, DNI, Documento Nacional de Identidad, NINO, NI, No, Social Insurance, SIN, Social Security, SSN

Keyword RegEx Expression: searchctl classification add --name 'ssnkeyword' --regex '.*((?i)(Social|Insurance|Number|Security|SIN|SSN|Pension|INSEE|Numero|Securite|Social e|Codice fiscale|DNI|Documento|Nacional|Identidad|NIE|Identificacion|Extranjero|NI|National|Rentenversicherung|Sozialgesetzbuch|Deutsche|Code|German|System|Identification|Italino|Italian)).*'

German Social Insurance Number

Format: 12 123456 A 123
RegEx Expression: searchctl classification add --name 'germansin' --regex '([0-9]{2}\s[0-9]{6}\s[a-zA-Z]\s[0-9]{3})'

Format: 12 123456 A 12 1
RegEx Expression: searchctl classification add --name 'germansin2' --regex '([0-9]{2}\s[0-9]{6}\s[a-zA-Z]\s[0-9]{2}\s[0-9])'

Format: 12 12345678 A 121
RegEx Expression: searchctl classification add --name 'germansin3' --regex '([0-9]{2}\s[0-9]{8}\s[a-zA-Z]\s[0-9]{3})'

Format: 12 12345678 A 12 1
RegEx Expression: searchctl classification add --name 'germansin4' --regex '([0-9]{2}\s[0-9]{8}\s[a-zA-Z]\s[0-9]{2}\s[0-9])'

Key Words: Social Security, Pension Insurance, Rentenversicherung, SIN, Social Code, Sozialgesetzbuch (SGB)

Social Insurance/Security Keyword RegEx Expression: searchctl classification add --name 'desinkeyword' --regex '((?i)(Social|Security|Pension|Insurance|Rentenversicherung|SIN|Social|Code|Sozialgesetzbuch|SGB))'
Phrase Examples: “German Social Security Number”, “German Social Security System (Sozialversicherung)”, Deutsche Sozialversicherungsnummer
Security:

	Public Data	Private Data	Restricted Data
PII			If living in North America
PCI			X
PHI			X
GDPR			X

French Social Insurance Number

Format: 1234567891234 12
RegEx Expression: searchctl classification add --name 'frenchsin' --regex '([0-9]{13}\s[0-9]{2})'

Format: 1 12 12 12345 123 12

INSEE RegEx Expression: searchctl classification add --name 'frenchsin2' --regex '([1-9]\s[0-9]{2}\s[0-9]{2}\s[0-9]{5}\s[0-9]{3}\s[0-9]{2})'

Key Words: INSEE

French INSEE Keyword RegEx Expression: searchctl classification add --name 'frsinkeyword' --regex '((?i)(NIRPP|INSEE|National|Identification|Number|numero|d"""inscription|au|repertoire|de|securite|sociale))'
Phrase Examples: “National Identification Number” , “"numéro d'inscription au répertoire", "numéro de sécurité sociale"
Security:

	Public Data	Private Data	Restricted Data
PII			If living in North America
PCI			X
PHI			X
GDPR			X

Italian Codice Fiscale

Format: 16 characters alphanumeric code (i.e. MRT MTT86S28 F205Z)
RegEx Expression: searchctl classification add --name 'italiansin' --regex '([A-Za-z]{3}\s[A-Za-z]{3}\s[0-9|A-Z]{5}\s[0-9|A-Z]{5})'

Format: MRTMTT86S28F205Z (No spaces in number)
Codice Fiscale RegEx Expression: searchctl classifications add --name 'italiansin2' --regex '([A-Za-z]{6}[0-9lmnpqrstuvLMNPQRSTUV]{2}[abcdehlmprstABCDEHLMPRST]{1}[0-9lmnpqrstuvLMNPQRSTUV]{2}[A-Za-z]{1}[0-9lmnpqrstuvLMNPQRSTUV]{3}[A-Za-z]{1})|([0-9]{11})'

Key Words: Codice fiscale

Italian Codice Fiscale Keyword RegEx Expression: searchctl classification add --name 'itsinkeyword' --regex '((?i)(Codice|fiscale|italiano|Italian|fiscal|code))'
Phrase Examples: Italian fiscal code, Codice fiscale italiano

Security:

	Public Data	Private Data	Restricted Data
PII			If living in North America
PCI			X
PHI			X
GDPR			X

Spanish Social Insurance

NIE (Número de Identificación de Extranjero) for non Spanish Citizens

Format: K,L,X,Y,or Z, followed by 7 digits followed by a letter (i.e K1234567A)
NIE RegEx Expression: searchctl classification add --name 'spanishsin' --regex '((?i)([K-L]|[X-Z])\d{7}([A-Z]))'

Key Words: NIE, Número de Identificación de Extranjero

Spanish NIE Keyword RegEx Expression: searchctl classification add --name 'spsinkeyword' --regex '((?i)(NIE|Numero|de|Identificacion|Extranjero|DNI|Documento|Nacional|Identidad))'
Phrase Example: Número de Identificación de Extranjero
Security:

	Public Data	Private Data	Restricted Data
PII			If living in North America
PCI			X
PHI			X
GDPR			X

DNI (Documento Nacional de Identidad) for Spanish citizens

Format: The DNI format is 8 digits followed by a letter (example: 12345678A). The same number is used for one's driver's license.

There are two other classes of DNI for citizens of Spain:

- 'K' (example: K12345678A) for children under 14 who are Spanish residents.
- 'L' (example: L12345678A) for Spanish citizens who live outside Spain.

DNI RegEx Expression: searchctl classification add --name 'spanishsin2' --regex '[K-L]?)\d{8}[A-Z]'

Key Words: DNI, Documento Nacional de Identidad

Spanish DNI Keyword RegEx Expression: searchctl classification add --name 'spsinkeyword' --regex '((?i)(NIE|Numero|de|Identificación|Extranjero|DNI|Documento|Nacional|Identidad))'
Phrase Example: Documento Nacional de Identidad
Security:

	Public Data	Private Data	Restricted Data
PII			If living in North America
PCI			X
PHI			X
GDPR			X

UK NINO

Format: Two prefix letters, six digits and one suffix letter. (i.e. QQ123456C)
NINO RegEx Expressions:

searchctl classification add --name 'uknino' --regex '(?!BG|GB|NK|KN|TN|NT|ZZ)[A-CEGHJ-PR-TW-Z][A-CEGHJ-NPR-TW-Z](?:\s*\d{2}){3}\s*[A-D]'

searchctl classification add --name 'uknino2' --regex '[A-CEGHJ-PR-TW-Z]{1}[A-CEGHJ-NPR-TW-Z]{1}[0-9]{6}[A-DFM]{0,1}'

Key Words: NINO, NI, No.

UK NINO Keyword RegEx Expression: searchctl classification add --name 'ukninokeyword' --regex '((?i)(National|Insurance|number|NINO|NI|No.))'
Phrase Example: National Insurance number

Security:

	Public Data	Private Data	Restricted Data
PII			If living in North America
PCI			X
PHI			X
GDPR			X

Canadian Social Insurance Number

Format: 123-123-123 or 123 123 123
Social Insurance Number RegEx Expression: searchctl classification add --name 'canadiansin' --regex '(\d{3}-?\s?\d{3}-?\s?\d{3})'
Key Words: Social Insurance, SIN

Canadian Social Insurance Keyword RegEx Expression: searchctl classification add --name 'casinkeyword' --regex '((?i)(Social|Insurance|Number|SIN))'
Phrase Example: Social Insurance Number
Security:

	Public Data	Private Data	Restricted Data
PII			X
PCI			X
PHI			X
GDPR			GDPR does not apply to Canadian citizens living in Canada

US Social Security Number

Format: 123-12-123 or 123 12 1234
Social Security Number RegEx Expressions:

searchctl classification add --name 'usssn' --regex '(\d{3}-?\s?\d{2}-?\s?\d{4})'

searchctl classification add --name 'usassn' --regex '(?!(000|666|9))\d{3}-(?!00)\d{2}-(?!0000)\d{4}$|^(?!(000|666|9))\d{3}(?!00)\d{2}(?!0000)\d{4}'

Key Words: Social Security, SSN

US Social Security Keyword RegEx Expression: searchctl classification add --name 'usssnkeyword' --regex '((?i)(Social|Security|Number|SSN))'
Phrase Example: Social Security Number
Security:

	Public Data	Private Data	Restricted Data
PII			X
PCI			X
PHI			X
GDPR			GDPR does not apply to US citizens living in the US

Date of Birth

Format: m-d-yy or mm-dd-yy, or m-d-yyyy, or mm-dd-yyyy, or m/d/yy, or mm/dd/yy, or m/d/yyyy, or mm/dd/yyyy.
Date of Birth RegEx Expression: searchctl classification add --name 'dob' --regex '([0-1]?)([0-9])-?\/?[0-3]?[0-9]-?\/?\d{2,4}'

Format: m/d/yy, or mm/dd/yy, or m/d/yyyy, or mm/dd/yyyy.
Date of Birth RegEx Expression: searchctl classification add --name 'dob2' --regex '([0-1]?)([0-9])\/[0-3]?[0-9]\/\d{2,4}'

Format: m-d-yy or mm-dd-yy, or m-d-yyyy, or mm-dd-yyyy, .
Date of Birth RegEx Expression: searchctl classification add --name 'dob3' --regex '([0-1]?)([0-9])-[0-3]?[0-9]-\d{2,4}'

Format: mm/dd/yyyy, or mm-dd-yyyy, or mm.dd.yyyy, or m/d/yyyy, or m-d-yyyy or m.d.yyyy
Date of Birth Expression: searchctl classification add --name 'usadob' --regex '(?:(?:(?:0?[13578]|1[02])(\/|-|\.)31)\1|(?:(?:0?[1,3-9]|1[0-2])(\/|-|\.)(?:29|30)\2))(?:(?:1[6-9]|[2-9]\d)?\d{2})$|^(?:0?2(\/|-|\.)29\3(?:(?:(?:1[6-9]|[2-9]\d)?(?:0[48]|[2468][048]|[13579][26])|(?:(?:16|[2468][048]|[3579][26])00))))$|^(?:(?:0?[1-9])|(?:1[0-2]))(\/|-|\.)(?:0?[1-9]|1\d|2[0-8])\4(?:(?:1[6-9]|[2-9]\d)?\d{2})'

Key Words: DOB, Date, Data, Fecha, Birth, Geburt, naissance, nascita, nacimiento

DOB Keyword RegEx Expression: searchctl classification add --name 'dobkeyword' --regex '((?i)(Date|Birth|DOB|Geburtsdatum|Naissance|Nascita|Fecha|Nacimiento|Geburt|Data))'

Phrase Example: Date of Birth, Geburtsdatum, Date de naissance, Data di nascita, Fecha de nacimiento
Security: As a stand alone data element it is Public Data and not restricted, but with Social Security Number and Name it is restricted data.

	Public Data	Private Data	Restricted Data
PII	If stand alone or Public Domain Data (not likable to other data)		If associated with name
PCI			If associated with name
PHI			Not Applicable
GDPR			If associated with name

Email

Email RegEx Expressions:

searchctl classifications add --name 'email' --regex '[\w-\.]+@([\w-]+\.)+[\w-]{2,4}'

searchctl classification add --name 'email2' --regex '([a-zA-Z0-9_\-\.]+)@([a-zA-Z0-9_\-\.]+)\.([a-zA-Z]{2,5})'

Key Words: Email

Email Key Word RegEx Expression: searchctl classifications add --name 'emailkeyword' --regex '((?i)(Email|E-mail|Address|Adresse|Correo|electronico|Indirizzo\\e-mail|correo|Direccion|email))'
Key Phrase: Email, E-mail, Email Address, Correo electrónico, E-Mail-Addresse, Adresse e-mail, Indirizzo email, Dirección de correo electrónico

Security:

	Public Data	Private Data	Restricted Data
PII			X
PCI			X
PHI			X
GDPR			X

IBAN: Europe Only

Note IBAN is used in Europe, not Canada or US. The first two digits are for the country code (i.e "DE" for Germany) followed by a check digit comprising two characters and the national account number BBAN (basic bank account number), which is made up of the eight-digit bank sort code and the ten-digit account number

Germany, Italy, France, Spain, UK (all formats)
RegEx Expression: searchctl classifications add --name 'iban' --regex '(?:DE)\d{20}|(?:FR)\d{19}[A-Z]\d{5}|(IT)\d{2}[A-Z]\d{22}|(?:ES)\d{22}|(?:GB)\d{2}[A-Z]{4}\d{14}|(?:DE)\d{2}\s\d{4}\s\d{4}\s\d{4}\s\d{4}\s\d{2}|(?:FR)\d{2}\s\d{4}\s\d{4}\s\d{4}\s\d{4}\s\d{1}[A-Z]\d{2}\s\d{3}|(IT)\d{2}\s[A-Z]\d{3}\s\d{4}\s\d{4}\s\d{4}\s\d{4}\s\d{3}|(?:GB)\d{2}\s[A-Z]{4}\s\d{4}\s\d{4}\s\d{4}\s\d{2}|(?:ES)\d{2}\s\d{4}\s\d{4}\s\d{4}\s\d{4}\s\d{4}'

Note the following Key Words and Keyword Regex Expression apply to the Bank Acounts for Germany, France, Italy Spain and the UK:
Key Words: IBAN, BBAN, Bank Konto, Kontonummer, Bank, Numéro, compte, bancaire, Numero, conto, bancario, Número, cuenta, bancaria, Bank, Account, Number, Intrenational, IBAN

IBAN Keyword RegEx Expression: searchctl classifications add --name 'ibankeyword' --regex '((?i)(Bank|Account|Number|Numero|Nummer|Routing|Transit|Bankario|Financial|Institution|German|French|Italian|Spanish|Kontonummer|Compte|Bancaire|Conto|International|Basic|Cuenta|Banco|Konto|Bankaire|Internacional|Internazionale|Internationale|BANK|ACCOUNT|NUMBER|NUMERO|CUENTA|BANCARIA|IBAN|BBAN))'
Phrase Examples: International Bank Account Number, Basic Bank Account Number, Numero de cuenta de Banco Internacional, Numero di conto bancario internazionale, Numéro de compte bancaire international, Internationale Kontonummer

Security:

	Public Data	Private Data	Restricted Data
PII			X
PCI	When storing account details PCI does not apply; it only applies to payment cards		
PHI			NA
GDPR			X

Bank Account Numbers

Germany

Format: DE29 1349 3022 4806 3708 88
German Bank Account RegEx Expression: searchctl classifications add --name 'ibande' --regex '(?:DE)\d{2}\s\d{4}\s\d{4}\s\d{4}\s\d{4}\s\d{2}'

Format: DE12345678901234567890
German Bank Account RegEx Expression: searchctl classifications add --name 'ibande2' --regex '(?:DE)\d{20}'

Key Words: Bank Konto

German IBAN Keyword RegEx Expression: searchctl classifications add --name 'debakeyword' --regex '((?i)(Kontonummer|Bank|Konto|konto|IBAN))'
Phrase Examples: Kontonummer, Bank Konto
Security:

	Public Data	Private Data	Restricted Data
PII			X
PCI	When storing account details PCI does not apply; it only applies to payment cards.		
PHI			X
GDPR			X

France

Format: FR29 9169 9073 6047 1767 2P38 290
French Bank Account RegEx Expression: searchctl classifications add --name 'ibanfr' --regex '(?:FR)\d{2}\s\d{4}\s\d{4}\s\d{4}\s\d{4}\s\d{1}[A-Z]\d{2}\s\d{3}'

Format: FR1234567895487613478A12345
French Bank Account RegEx Expression: searchctl classifications add --name 'ibanfr2' --regex '(?:FR)\d{19}[A-Z]\d{5}'

Key Words: Numéro de compte bancaire

French IBAN Keyword RegEx Expression: searchctl classifications add --name 'frbakeyword' --regex '((?i)(Numero|compte|bancaire|INUMERO|DE|COMPTE|BANCAIRE|IBAN))'
Phrase Examples: Numéro de compte bancaire, Compte bancaire
Security:

	Public Data	Private Data	Restricted Data
PII			X
PCI	When storing account details PCI does not apply; it only applies to payment cards.		
PHI			X
GDPR			X

Italy

Format: IT89 R374 3328 2256 2103 0440 500
Italian Bank Account RegEx Expression: searchctl classifications add --name 'ibandit' --regex '(?:IT)\d{2}\s[A-Z]\d{3}\s\d{4}\s\d{4}\s\d{4}\s\d{4}\s\d{3}'

Format: IT12A1234567891234567891234
Italian Bank Account RegEx Expression: searchctl classifications add --name 'ibandit2' --regex '(IT)\d{2}[A-Z]\d{22}'

Key Words: Numero di conto bancario

Italian IBAN Keyword RegEx Expression: searchctl classifications add --name 'itbakeyword' --regex '((?i)(Numero|conto|bancario|di|DI|NUMERO|DI|CONTO|BANCARIO|IBAN))'

Phrase Examples: Numero di conto bancario, Conto bancario

Security:

	Public Data	Private Data	Restricted Data
PII			X
PCI	When storing account details PCI does not apply; it only applies to payment cards.		
PHI			X
GDPR			X

Spain

Format: ES00 4514 6356 8256 8068 7871
RegEx Expression: searchctl classifications add --name 'ibanesp' --regex '(?:ES)\d{2}\s\d{4}\s\d{4}\s\d{4}\s\d{4}\s\d{4}'

Format: ES1234567901234567890123
RegEx Expression: searchctl classifications add --name 'ibanes2' --regex '(?:ES)\d{22}'

Key Words: Número de cuenta bancaria

Spanish IBAN Keyword RegEx Expression: searchctl classifications add --name 'spbakeyword' --regex '((?i)(Numero|cuenta|bancaria|de|DE|NUMERO|CUENTA|BANCARIA|IBAN))'

Phrase Examples: Número de cuenta bancaria, Cuenta bancaria
Security:

	Public Data	Private Data	Restricted Data
PII			X
PCI	When storing account details PCI does not apply; it only applies to payment cards.		
PHI			X
GDPR			X

United Kingdom

Format: GB96 CZTD 5452 0531 6999 91
RegEx Expression: searchctl classifications add --name 'ibanduk' --regex '(?:GB)\d{2}\s[A-Z]{4}\s\d{4}\s\d{4}\s\d{4}\s\d{2}'

Format: GB12ABCD12345678912345
RegEx Expression: searchctl classifications add --name 'ibanduk2' --regex '(?:GB)\d{2}[A-Z]{4}\d{14}'

Key Words: Bank, Account, Number, Intrenational, IBAN

UK IBAN Keyword RegEx Expression: searchctl classifications add --name 'ukbakeyword' --regex '((?i)(Bank|Account|Number|Intrenational|BANK|ACCOUNT|NUMBER|IBAN))'

Phrase Examples: International Bank Account Number, IBAN Number
Security:

	Public Data	Private Data	Restricted Data
PII			X
PCI	When storing account details PCI does not apply; it only applies to payment cards.		
PHI			X
GDPR			X

Bank Account Numbers: Canada and US only

Account Numbers Canada and US

Canadian Account Numbers breakdown:
Financial Institution 3 digits
Bank Transit Number 5 digits
Account Number 7 - 12 digits
Total 15 -20 digits

US Account Numbers Breakdown:
US Routing Number 9 digits
Account Number 10 - 12 digits
Total 19 - 21 Digits

Bank Account Number US (10 - 12 digits) and Canada (7- 12 digits)

US

Format: 10 to 12 digits
US Bank Account RegEx Expression: searchctl classifications add --name 'usba' --regex '(\d{10,12})'

Key Words: Bank Account Number

Keyword RegEx Expression: searchctl classifications add --name 'nabakeyword' --regex '((?i)(Bank|Account|Number|BANK|ACCOUNT|NUMBER|Routing|ROUTING|Transit|Financial|Institution|Canadian|Canada|US))'
Phrase Examples: Bank Account, Bank Account Number
Security:

	Public Data	Private Data	Restricted Data
PII			X
PCI	When storing account details PCI does not apply; it only applies to payment cards.		
PHI			X
GDPR			X

Canada

Format: 7 to 12 digits
Canadian Bank Account RegEx Expression: searchctl classifications add --name 'caba' --regex '(\d{7,12})'

Key Words: Canadian Bank Account Number

Keyword RegEx Expression: searchctl classifications add --name 'cabakeyword' --regex '((?i)(Bank|Account|Number|BANK|ACCOUNT|NUMBER|Routing|Transit|Financial|Institution|Canadian|Canada))'
Phrase Examples: Bank Account, Bank Account Number
Security:

	Public Data	Private Data	Restricted Data
PII			X
PCI	When storing account details PCI does not apply; it only applies to payment cards.		
PHI			X
GDPR			X

Bank Routing Numbers

US Routing Number

Format: Nine Digit Bank Routing Transit Number complying with ABA Rules:
US Routing Number RegEx Expression: searchctl classifications add --name 'usbr' --regex '((0[0-9])|(1[0-2])|(2[1-9])|(3[0-2])|(6[1-9])|(7[0-2])|80)([0-9]{7})'

Key Words: Bank Routing Number

Keyword RegEx Expression: searchctl classifications add --name 'usbrkeyword' --regex '((?i)(Bank|Account|Number||BANK|ACCOUNT|NUMBER|ROUTING|Routing|Transit|Financial|Institution|US))'
Phrase Examples: US Bank Routing Number
Security: Public information

	Public Data	Private Data	Restricted Data
PII	If stand alone and is not likable to bank account data.		If linked with bank account data.
PCI	When storing account details PCI does not apply; it only applies to payment cards.		
PHI	If stand alone and is not likable to bank account data.		If linked with bank account data.
GDPR	If stand alone and is not likable to bank account data.		If linked with bank account data.

Canadian Routing Numbers

Canadian Routing numbers are composed of a 5 digit Bank Transit Number and a 3 digit Financial Institution number.

Canadian Routing Number RegEx Expression: searchctl classifications add --name 'cabr' --regex '(\d{5}-?\s?\d{3})'

Key Words: Routing Number, Financial Institution, Transit Number

Keyword RegEx Expression: searchctl classifications add --name 'cabrkeyword' --regex '((?i)(Bank|Account|Number||BANK|ACCOUNT|NUMBER|ROUTING|Routing|Transit|Financial|Institution|Canadian|Canada))'
Phrase Examples: Canadian Bank Routing Number, Financial Institution, Bank Transit Number
Security: Public information

	Public Data	Private Data	Restricted Data
PII	If stand alone and is not likable to bank account data.		If linked with bank account data.
PCI	When storing account details PCI does not apply; it only applies to payment cards.		
PHI	If stand alone and is not likable to bank account data.		If linked with bank account data.
GDPR	If stand alone and is not likable to bank account data.		If linked with bank account data.

Routing and Bank Account Numbers

US

Format: Nine Digit Bank Routing Transit Number plus 10 to 12 digit Bank Account Number
US Routing plus Bank Account RegEx Expression: searchctl classifications add --name 'usbrba' --regex '(((0[0-9])|(1[0-2])|(2[1-9])|(3[0-2])|(6[1-9])|(7[0-2])|80)([0-9]{7}))-?\s?(\d{10,12})'

Key Words: Bank Account Number

Keyword RegEx Expression: searchctl classifications add --name 'usbrbakeyword' --regex '((?i)(Bank|Account|Number||BANK|ACCOUNT|NUMBER|ROUTING|Routing|Transit|Financial|Institution|US))'

Phrase Examples: Bank Account Data
Security:

	Public Data	Private Data	Restricted Data
PII			X
PCI	When storing account details PCI does not apply; it only applies to payment cards.		
PHI			X
GDPR			X

Canada

Format: 5 digit Bank Transit Number, a 3 digit Financial Institution number and 7 to 12 digit Account Number
Canadian Routing plus Bank Account RegEx Expression: searchctl classifications add --name 'carba' --regex '(\d{5}-?\s?\d{3}-?\s?\d{7,12})'

Key Words: Bank Account Number

Keyword RegEx Expression: searchctl classifications add --name 'carbakeyword' --regex '((?i)(Bank|Account|Number||BANK|ACCOUNT|NUMBER|ROUTING|Routing|Transit|Financial|Institution|Canadian|Canada))'
Phrase Examples: Bank Account Data
Security:

	Public Data	Private Data	Restricted Data
PII			X
PCI	When storing account details PCI does not apply; it only applies to payment cards.		
PHI			X
GDPR			X

Credit Card Number

Regex Expression for Visa (Vpay), Visa Electron, AMX, MasterCard, Maestro, Diners Club Cart Blanche, Diners Club International, Diners Club US & Canada, Discover to match Length, IIN Ranges and Spacing Patterns shown in [Credit Card IIN Ranges & Spacing Patterns]:
searchctl classifications add --name 'creditcard' --regex '((2[2-7]\d{2}-?\s?\d{4}-?\s?\d{4}-?\s?\d{4})|(4\d{3}-?\s?\d{4}-?\s?\d{4}-?\s?\d{4})|(3[0-9]\d{2}-?\s?\d{6}-?\s?\d{4,5})|(5[0-8]\d{2}-?\s?\d{4,6}-?\s?\d{4,5}-?\s?\d{0,4}-?\s?\d{0,3})|(6[0-9]\d{2}-?\s?\d{4,6}-?\s?\d{4,5}-?\s?\d{0,4})-?\s?\d{0,3})'

Key Words: Credit Card, CCN

Keyword RegEx Expression: searchctl classifications add --name 'creditcardkeyword' --regex '((?i)(Credit|Card|Number|Master|Discover|American Express|Maestro|Visa|CC|Diners Club|International|Cart Blanche|Electron|AMX|Kreditkartennummer|Numero|Carte|credito|de|carta|tarjeta))'
Phrase Example: Credit Card Number, Kreditkartennummer, Numéro de Carte de Crédit, Numero di carta di credito, Número de tarjeta de crédito, Mastercard, VISA, Discover Card, American Express, AMX, Maestro, Cart Blanche, Diners Club, Diners Club International
Security:

	Public Data	Private Data	Restricted Data
PII			X
PCI			X
PHI			X
GDPR			X

ACH Format Compliance Check

Automated Clearing House (ACH) is an electronic network for financial transactions. These transactions include credit and debit transactions. ACH credit transfers include direct deposit payroll and vendor payments.Can use the following RegEx to pattern match numbers to check ACH format compliance, check for 1 to 17 alpha-numerics:
ACH Format Compliance Check Regex Expression: searchctl classifications add --name 'achcompliance' --regex '\w{1,17}'

Key Words: ACH, ACH Check

Keyword RegEx Expression: searchctl classifications add --name 'achkeyword' --regex '(ACH|Compliance|Check)'

Phrase Examples: ACH Compliance
Security:

	Public Data	Private Data	Restricted Data
PII			X
PCI			X
PHI			X
GDPR			X

Credit Card CVV

Format: xxx
Credit Card CVV RegEx Expression (3 Digit): searchctl classifications add --name 'cvv3' --regex '\d{3}'

Format: xx, or xxx, or xxxx
Credit Card CVV RegEx Expression (2 to 4 Digit): searchctl classifications add --name 'cvv' --regex '\d{2,4}'

Key Words: CVV, CVC

Keyword RegEx Expression: searchctl classifications add --name 'cvvkeyword' --regex '((?i)(CVV|CVC|Card|Karten|Verification|Number|Code|Verificacion|Verifiziernummer|Uberprufung|Verifica|Wert|Valeur|Valore|Valor|Le code de|Codigo|Codice|Numero|Nombre|Nummer))'
Phrase Example: Card Verification Value, Card Verification Code, Credit Card Verification, CVV Number, CVC Number, Credit Verification Value Number, Credit Card Verification Number, Karten Verifiziernummer, Valeur de vérification de la carte, Valore di verifica della carta, valor de verificación de la tarjeta, Kartenbestätigungscode, Le code de vérification de carte, Codice di verifica della carta, Código de verificación de la tarjeta
Security: As a stand alone data element there are no restrictions but, In conjunction with Credit Card Number and Expiration Date the following restrictions apply

	Public Data	Private Data	Restricted Data

PII			In conjunction with Credit Card Number and Expiration Date
PCI			In conjunction with Credit Card Number and Expiration Date
PHI			In conjunction with Credit Card Number and Expiration Date
GDPR			In conjunction with Credit Card Number and Expiration Date

Expiration Date

Format: mm/yy
Expiration Data RegEx: searchctl classifications add --name 'exp2' --regex '(?:0[1-9]|1[0-2])\/[0-9]{2}'

Format: mm/yy or mm/yyyy
Expiration Data RegEx: searchctl classifications add --name 'exp2or4' --regex '(?:0[1-9]|1[0-2])\/[0-9]{2,4}'

Key Words: Expiration Date, Haltbarkeitsdatum, Ablaufdatum der Kreditkarte, Date d'expiration, Data di scadenza, Fecha de caducidad

Keyword RegEx Expression: searchctl classifications add --name 'ccexpkeyword' --regex '((?i)(Credit|Card|Expiration|Date|Data|Fecha|Haltbarkeitsdatum|Ablaufdatum|Kreditkarte|scadenza|caducidad|Credito|Anerkennung|Credito|Tarjeta|Carta|Carte|Karte))'
Phrase Examples: Credit Card Expiration, Credit Card Expiration Date, Ablaufdatum der Kreditkarte, Date d'expiration de la carte de crédit, Data di scadenza della carta di credito, Fecha de vencimiento de la tarjeta de crédito
Security: As a stand alone data element there are no restrictions but, In conjunction with Credit Card Number and CVV the following restrictions apply.

	Public Data	Private Data	Restricted Data
PII			In conjunction with Credit Card Number and CVV
PCI			In conjunction with Credit Card Number and CVV
PHI			In conjunction with Credit Card Number and CVV

GDPR			In conjunction with Credit Card Number and CVV

Medical Data

Medical Record Number (MRN)

Format: xxx-x-xxxxx, or xxx x xxxxx, or xxxxxxxxx
Medical Record RegEx Expression: searchctl classifications add --name 'mrn' --regex '[0-9]{3}-?\s?[0-9]{1}-?\s?[0-9]{5}'

Format: 05xxxxxxx
Medical Record RegEx Expression: searchctl classification add --name 'usamrn' --regex '(05)\d{7}'

Key Words: MRN, Medical Record Number, Nummer der Krankenakte, Numéro de dossier médical, Numero di cartella clinica, Numero de historia clinica

Keyword RegEx Expression: searchctl classifications add --name 'mrnkeyword' --regex '((?i)(Medical|Record|Number|MRN|Nummer|Krankenakte|Numero|dossier|medical|cartella|clinica|historia|clinica|Geduldig|Patiente|Patient|Paziente|Paciente|Med Rec))'

Phrase Example: Patients Medical Record Number, Nummer der Patientenakte, Numéro de dossier médical du patient, Numero di cartella clinica del paziente, Número de historia clínica del paciente
Security:

	Public Data	Private Data	Restricted Data
PII			X
PCI	PCI Appplies to Credit Card Data only		
PHI			X
GDPR			X

National Provider Identifier (NPI)

Format: 10 digits XXXXXXXXXX

RegEx Expression: searchctl classification add --name 'npi' --regex '\d{10}'

Key Words: NPI,"National Provider Identifier"

Keyword RegEx Expression: searchctl classifications add --name 'npikeyword' --regex '((?i)(NPI|National|Provider|Identifier|ID))'

Phrase Example: NPI,"National Provider Identifier"
Security:

	Public Data	Private Data	Restricted Data
PII			If linked to name, medical record or other data that could link to an individual
PCI	PCI Appplies to Credit Card Data only		
PHI			X
GDPR			X

Financial Identification Number (FIN)

Format: 1200xxxxxx
RegEx Expression: searchctl classification add --name 'fin' --regex '(1200)\d{6}'

Key Words: FIN,"Financial Identification Number","Financial ID","Patient ID","Encounter ID","Encounter #"

Keyword RegEx Expression: searchctl classifications add --name 'finkeyword' --regex '((?i)(FIN|Financial|Identification|Number|ID|Patient|Encounter))'

Phrase Example: FIN,"Financial Identification Number","Financial ID","Patient ID","Encounter ID","Encounter #"
Security:

	Public Data	Private Data	Restricted Data
PII			If linked to name, medical record or other data that could link to an individual
PCI	PCI Appplies to Credit Card Data only		

PHI			X
GDPR			X

Medical License Number

Format: 2 letter prefix followed by 5 digits (i.e. ME.10003)
RegEx Expression: searchctl classification add --name 'mln' --regex '[A-Z]{2}\.\d{5}'

Key Words: License,"License #","Medical License Number"

Keyword RegEx Expression: searchctl classifications add --name 'mlnkeyword' --regex '((?i)(License|Medical|Number|MLN))'

Phrase Example: FIN,"Financial Identification Number","Financial ID","Patient ID","Encounter ID","Encounter #"
Security:

	Public Data	Private Data	Restricted Data
PII			If linked to name, medical record or other data that could link to an individual
PCI	PCI Appplies to Credit Card Data only		
PHI			X
GDPR			X

J Number

Format: J ollowed by 8 digits (Jxxxxxxxx)
RegEx Expression: searchctl classification add --name 'jnumber' --regex '[Jj][0-9]{8}"

Key Words: "J #","J Number"

Keyword RegEx Expression: searchctl classifications add --name 'Jkeyword' --regex '((?i)(J|Number))'

Phrase Example: "J #","J Number"
Security:

	Public Data	Private Data	Restricted Data
PII			If linked to name, medical record or other data that could link to an individual
PCI	PCI Appplies to Credit Card Data only		
PHI			X
GDPR			X

DEA Registration Number

Format: 2 letter prefix followed by 7 digits (i.e. MR0622476)
RegEx Expression: searchctl classification add --name 'dea' --regex '[A-HJ-NPR-UX]{1}[A-Z]{1}\d{7}'

Key Words: DEA,"DEA Registration","Drug Enforcement Administration"

Keyword RegEx Expression: searchctl classifications add --name 'deakeyword' --regex '((?i)(DEA|Registration|Drug|Enforcement|Administration|Number))'

Phrase Example: DEA,"DEA Registration","Drug Enforcement Administration"
Security:

	Public Data	Private Data	Restricted Data
PII			If linked to name, medical record or other data that could link to an individual
PCI	PCI Appplies to Credit Card Data only		
PHI			X
GDPR			X

Blood Type

Format: X(+/-)
Blood Type RegEx Expression: searchctl classifications add --name 'bloodtype' --regex '(A|B|AB|O)[+-]'

Key Words: Blood Type

Keyword RegEx Expression: searchctl classifications add --name 'btkeyword' --regex '((?i)(Blood Type|Blutgruppe|Groupe sanguin|Gruppo sanguigno|Tipo de sangre))'

Phrase Example: Blood Type, Blutgruppe, Groupe sanguin, Gruppo sanguigno, Tipo de sangre

Security:

	Public Data	Private Data	Restricted Data
PII			If linked to name, medical record or other data that could link to an individual
PCI	PCI Appplies to Credit Card Data only		
PHI			X
GDPR			X

Weight

Metric or English:

Weight RegEx Expression: searchctl classification add --name 'weight' --regex '(\d*\.?\d+)\s*(pounds?|lbs?|lb?|kgs?|kg)'

Key Words: Weight

Keyword RegEx Expression: searchctl classifications add --name 'weightkeyword' --regex '((?i)(Weight|Gewicht|Poids|Peso))'

Phrase Example: Weight, Gewicht, Poids, Peso,

Security:

	Public Data	Private Data	Restricted Data
PII			If linked to name, medical record or other data that could link to an individual
PCI	PCI Appplies to Credit Card Data only		

PHI			X
GDPR			X

Height

Format: cm or ft, in or ‘ “.
RegEx Expression: searchctl classification add --name 'height' --regex '(\d{1,3}\s?(ft)\s?\d{1,2}\s?(in)?)|(\d*\.?\d*\scm)'

Format: cm
RegEx Expression: searchctl classification add --name 'heightm' --regex '(\d*\.?\d*\scm)'

Format: ft, in or ', "
RegEx Expression: searchctl classification add --name 'heighti' --regex '(\d{1,3}'''''?\s?(ft)?\s?\d{1,2}\"?\s?(in)?)'

Note: For English height data, the characters ‘ and “ are format sensitive. For example
5'9" worked
while
5’6” did not work

Key Words: Height

Keyword RegEx Expression: searchctl classifications add --name 'heightkeyword' --regex '((?i)(Height|Hohe|Hohe|la taille|Taille|Altezza|Altura))'

Phrase Example: Height, Höhe, la taille, Altezza, Altura
Security:

	Public Data	Private Data	Restricted Data
PII			If linked to name, medical record or other data that could link to an individual
PCI	PCI Appplies to Credit Card Data only		

PHI			X
GDPR			X

© Superna LLC

9.25. ShowBack and ChargeBack Guide

Home Top

- Overview
  - Use Cases:
- Requirements
- How To Schedule a Search
- How to List Saved Searches and Delete Saved Searches CLI and Web GUI
- How to see scheduled search logs for debugging
- Example Report
- Customize Email Template
- Advanced - How to Configure the time Scheduled Reports are Emailed

Overview

The show back and chargeback feature is designed to automate reporting on department data, group shares or project data. The solution will be available in a new side menu tab called Show back & ChargeBack that will be visible to administrator users only. Users will not be able to see this new tab after they login. The most common scenario's that will be supported are listed below. It will support daily, weekly, monthly scheduling. File pool aware search results allow

reporting on each storage pool separately and the total on a specific path. Example archive file pool, cloudpool file pool, performance file pool. Showback will be usage only in GB's and ChargeBack will allow the data to be assigned a $ amount for each file pool or tier of data to show differential pricing for each tier.

Use Cases:

1. A scheduled email report on a specific file system path showing the sum of the data on the path.
2. A scheduled email report on a specific file system path showing the stale data total based on last modified or last accessed attribute
3. A scheduled email report on a specific file system path showing the file pool break down of stored data and the sum of all file pool usage.
4. A scheduled email report report on a specific file system path showing data a summary of the data owners.
5. A scheduled email report covering any of the scenario's above that shows usage and shows a $ amount associated to the data or pool the data is stored in.

Requirements

1. Release 1.1.5 or >
2. Only Administrator users configured from the CLI will have the option to create saved searches.

How To Schedule a Search

1. Login to the Search & Recover webUI as a user listed on the admin list
2. The Show back & Charge Back UI is the icon on the top right menu as shown below. It will appear after completing a search or a quick report.

a. 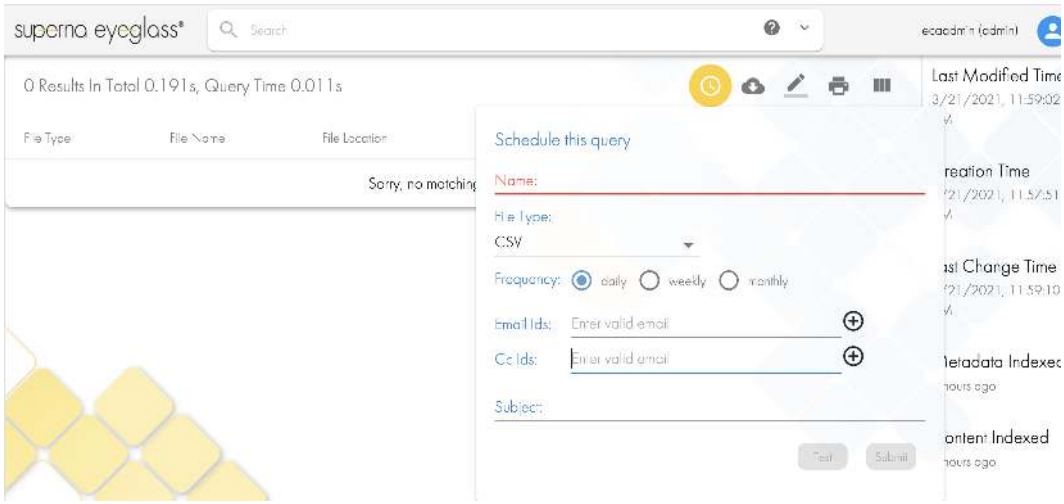

b. Quick Reports Example. The screenshot below is Who owns What? report

c. 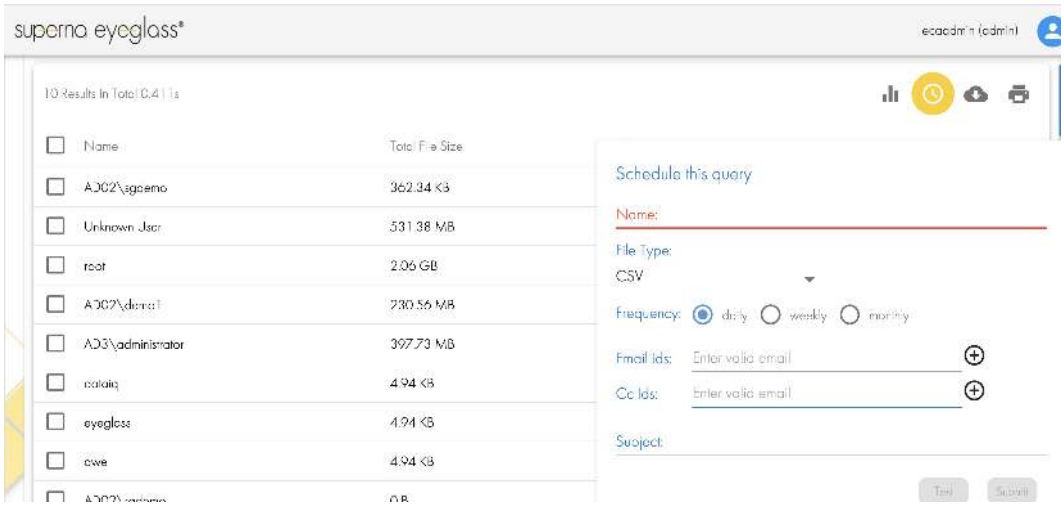

3. Configure the search parameters for the advanced Search UI and execute the search to display the Schedule icon.
4. Quick Reports also support saved searches, configure the search parameters , execute the search to display the results and then click the scheduled search icon and configure the fields as per the below definitions.
5. Fields to complete
    a. **Name** Give the saved report a good name as this will be the only way to identify the saved report from the CLI. No special characters, spaces are allowed.
    b. **File type** is CSV or Excel format for the attached report format
    c. **Frequency** Is daily, weekly or monthly
    d. **Email and CC** Email is the primary person to receive the report, CC is the carbon copy email (optional).
    e. **Subject** The subject of the email for this report
6. **Test button** - This will send an on demand email to validate the results and verify the report was received and verify the contents before saving it.
7. **Submit button** - This will save the report for scheduling.
8. NOTE: If click submit multiple times, this will save a duplicate scheduled search.
9. See steps below to list or remove a saved search.
10. Done

How to List Saved Searches and Delete Saved Searches CLI and Web GUI

1. The saved searches are visible from the CLI and WebGUI. The search criteria used are not displayed so the name of the saved search needs to be descriptive.
2. CLI
    a. List Saved Searches
        i. searchctl scheduledsearches list
        ii. The fields that are listed ID (unique ID for each search), email TO, CC email, Schedule, FileType, Email subject
    b. Remove a Saved Search
        i. searchctl scheduledsearches remove --id xxxxx
3. WebGUI
    a. The scheduled search gui is only available to ecaadmin or other admin users.
    b. The scheduled searches are listed with the configuration and shows last run time.
    c. The schedules can be edited and re-saved or select one or more schedules and use the trash icon to delete them.

d. 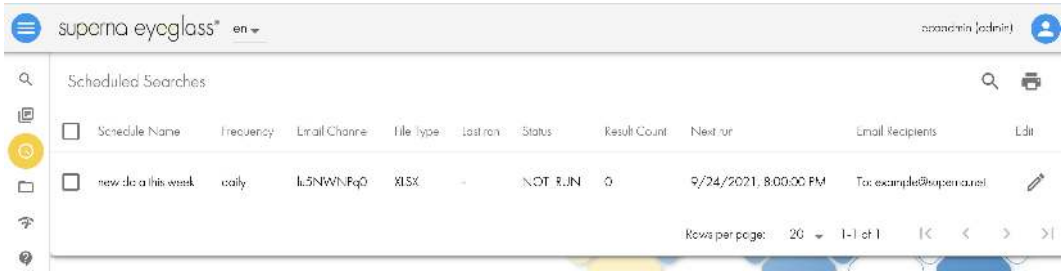

How to see scheduled search logs for debugging

1. Node 1 of the search appliance stores the scheduled search logs
2. login to node 1 as ecaadmin
3. cat /opt/superna/var/logs/searchmw/scheduledSearches.log
4. use this log to see issues with scheduled searches.

Example Report

1.

	A	B	C	D	E
1	**Name**	**Total File Size**	**File Cour**	**Average File Size**	
2	AD02\sgdemo	371040	11595	32	
3	Unknown User	557190137	5571	100016	
4	root	2215090238	1511	1465976	
5	AD02\demo1	241755810	352	686806	
6	AD3\administrator	417052795	32	13032899	
7	dataiq	5059	9	562	
8	eyeglass	5059	9	562	
9	qwe	5059	9	562	
10	AD02\rademo	0	4	0	
11	AD02\administrator	0	1	0	
12					

2. Cat Report for the week-report-Sun Mar 21 20:17:00 UTC 2021.csv

	A	B	C	D	E	F
1	File Type	File Name	File Location	Owner	File Size	Last Modified
2	file	cat.txt	/ifs/data/dfsdata/clon	root	67	Sat Oct 03 21:45:49
3	file	cat.txt	/ifs/data/dfsdata/sect	Unknown User	49	Sun Oct 27 12:45:04
4	file	cat.txt	/ifs/data/dfsdata/sear	Unknown User	56	Tue Dec 10 17:53:5
5	file	cat.txt	/ifs/data/dfsdata/aws/	Unknown User	49	Sun Oct 27 12:45:04
6	file	cat.txt	/ifs/data/dfsdata/azur	Unknown User	49	Sun Oct 27 12:45:04

3. Cat Report for the week 1 searchDemo ×

searchdemo@superna.net 4:17 PM (25 minutes ago)
to me

Please find attached search report.

Cat Report for the ...

Customize Email Template

The email template can be customized by editing an html template file on node 1 of the Search & Recover appliance. NOTE: Only 1 report template is supported.

1. Login to node 1 via ssh as ecaadmin user.
2. Edit the file to add additional text to the email body.
3. nano /opt/superna/var/search/downloads/searchreport/emailTemplate.html

4. Save the file after making edits
5. control+x answer yes to save the file.
6. New report emails will use this template when sending reports.

Advanced - How to Configure the time Scheduled Reports are Emailed

1. Changing the time of the daily, weekly and monthly scheduled reports can be customized to determine when searches are executed and emailed to users.
    a. The defaults are midnight of the day, Midnight on Sunday for weekly, and the first of the month at midnight.
2. login to eca node 1 as ecaadmin
3. nano /opt/superna/eca/eca-env-common.conf
4. Paste these variables to the file and adjust the cron string to change when the daily , weekly and monthly reports are emailed.
    a. export ECA_SEARCH_REPORT_SCHEDULE_DAILY="0 0 * * *"
    b. export ECA_SEARCH_REPORT_SCHEDULE_WEEKLY="0 0 * * 0"
    c. export ECA_SEARCH_REPORT_SCHEDULE_MONTHLY="0 0 1 * *"

© Superna LLC

10. Eyeglass Golden Copy Admin Guide

Home Top

- Golden Copy Overview
- Limitations, Requirements and Supported Target List
- Networking Deployment Overview, Cloud Provider IP Ranges, Proxy Configuration
- Understanding Copy Performance Configurations
- Cloud Provider Cost Considerations
- Supported Data Security & Storage Class Life Cycle Options
- Appliance Security, Authentication & Hardening
- Golden Copy Configuration Steps
- Golden Copy Back Bundle & Adv License Configuration Steps
- Golden Copy GUI - Beta
- Golden Copy VM Operations
- S3 Storage Bucket Configurations Options , Operations and Settings
- Trouble Shooting File Copies
- Golden Copy Solutions Guides

© Superna LLC

10.1. Golden Copy Overview

Home Top

- Overview
- What's New
- Deployment Diagram
- Architecture
- Licensing
- Target Use Cases:
  - Synced Backup Copy of Data with 1 version of a file
  - Synced Backup Copy of Data with N versions of a file
  - One Time copy of data for long term storage or archive
  - Isilon/PowerScale Snapshot copy and Archive
  - Full backup Mode
- Key Features

Overview

This product offers high speed file to object copy, and sync of files to native object store over the S3 protocol.

What's New

1. The latest releases enables
   a. Onefs 8.1.2 or later changelist API folder rename support to ensure renames are synced to the object store in the correct path
   b. 2 factor login over ssh

c. post job stats steps recomputes job view stats

d. Incremental sync jobs auto retry failed files

e. Coming soon

   i. QOS full vs incremental jobs - will prioritize one of the other. Global setting applies to all folders

Deployment Diagram

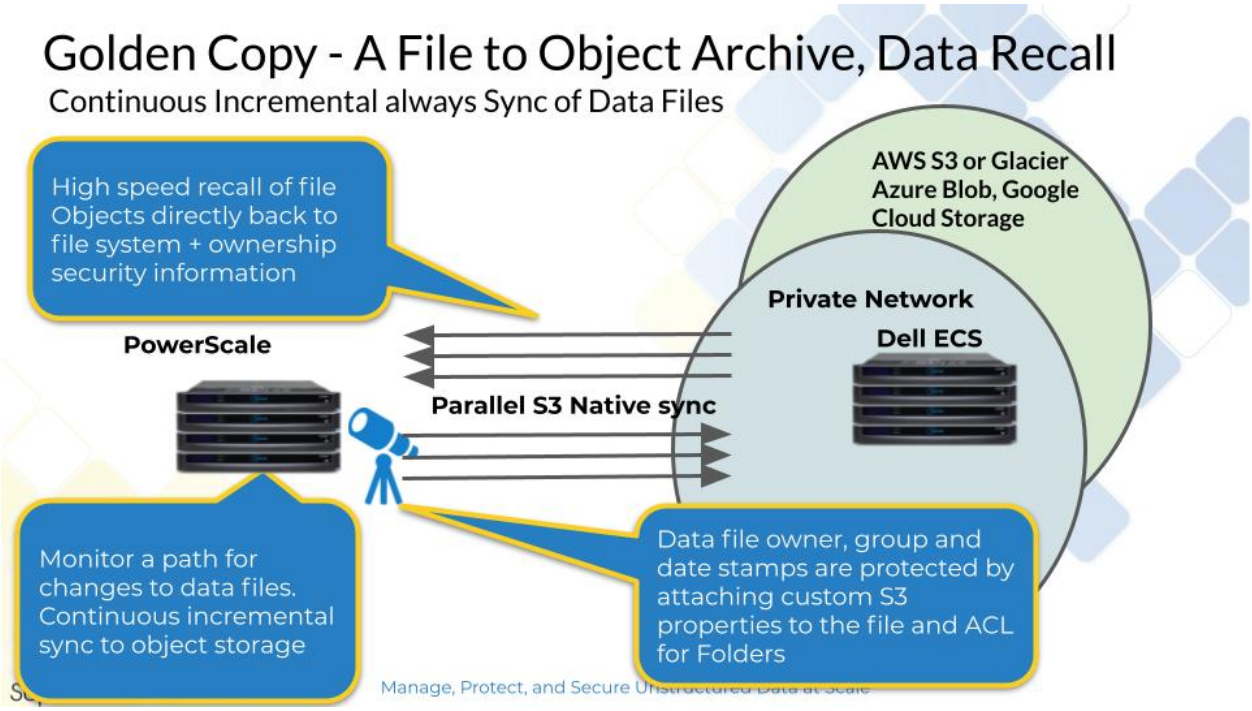

Architecture

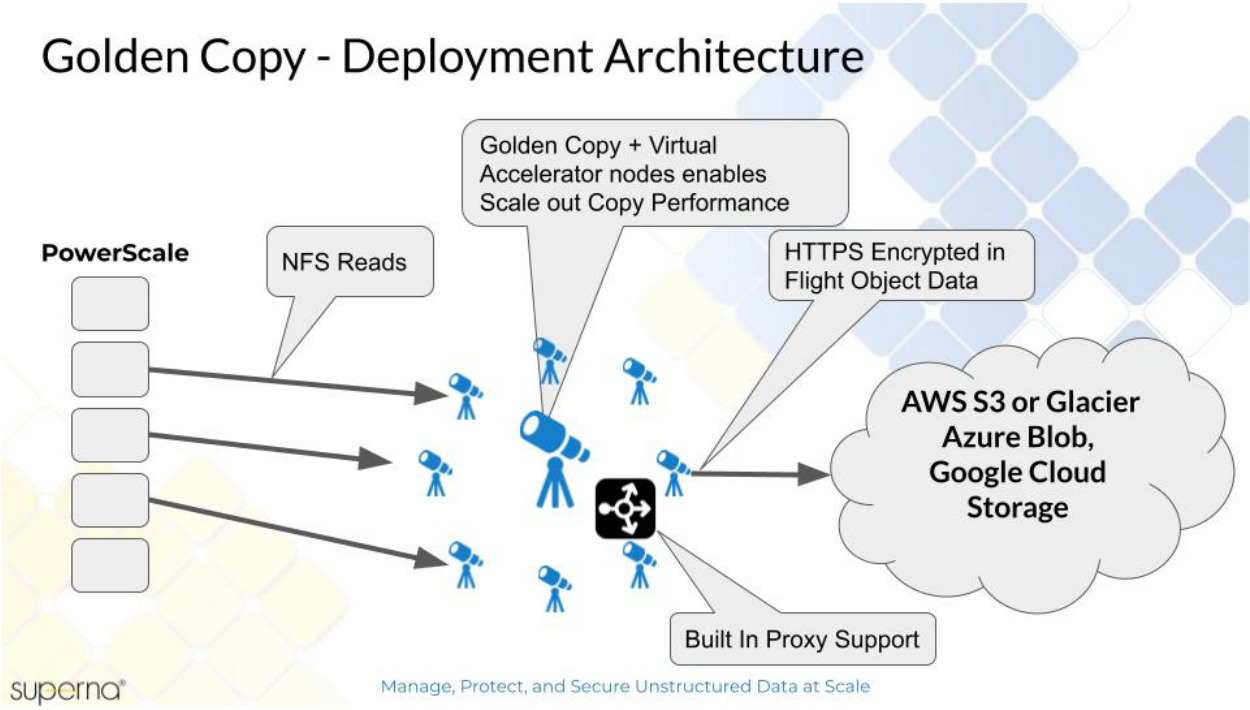

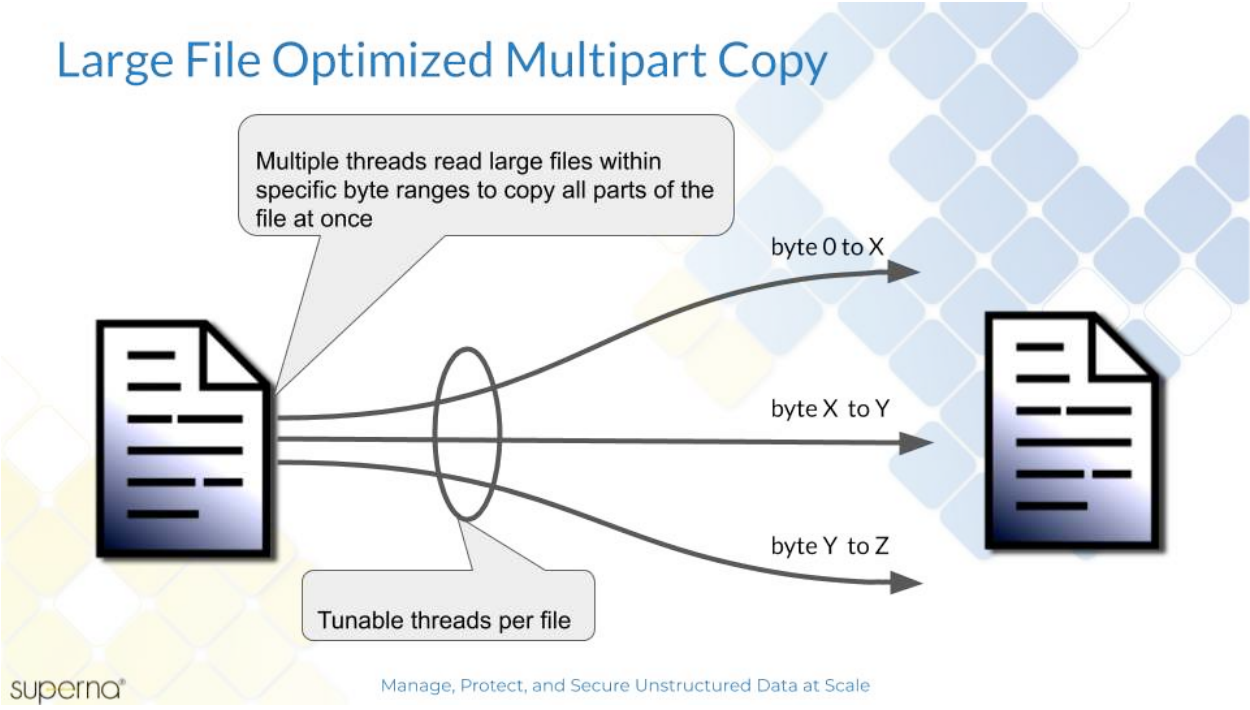

Licensing

1. The Golden Copy tool is a cluster based license with unlimited node count and unlimited copy size of data.

2. The license key is locked to a cluster and cannot be moved to a different cluster once the license has been applied to a cluster. An unlock license key can be purchased if this is required to move a license.

Target Use Cases:

Synced Backup Copy of Data with 1 version of a file

1. Requires: Base product license, recommended to use Backup bundle advanced license key for additional reporting and monitoring.
2. Only 1 copy or version of a file, file modifications will overwrite the previous version of the object with the new object
3. Higher probability of recalling data
4. Storage tier is near line and assumes data recall is important with shortest possible time to recall data
5. Deleted files are either synced as deleted objects OR turn off deletes and retain all deleted data inn the S3 storage
6. No expiry set on the S3 target, each version of a file is retained
7. Full copy followed by incremental schedule using Isilon change list api and snapshots for incremental always backup.

Synced Backup Copy of Data with N versions of a file

1. Requires: Backup Bundle or advanced license key
2. Same as the above use case but allows S3 bucket versioning to be enabled. Specify how many versions of a file should be available at any one time.

3. Example daily incremental and store 7 versions of a file provides 7 day restore window of data. The 8th day backup will automatically purge the oldest version of the file with a 7 version S3 bucket policy.
4. NOTE: Versions consume the file size of the file.
5. Object retention expiry would not be used in this use case.
6. Version aware restore/recall of data is supported with the backup bundle version of Golden Copy and allows an older than or newer than flag to pick which version of the day you want to recall. Example restore data and version from Wednesday would select the closest match to Wednesday backup on a per object basis.

One Time copy of data for long term storage or archive

1. Requires: The base product license
2. Full copy of a path to long term object storage
3. Low probability of recalling data
4. Optional
    a. object lock retention applied
    b. tiering objects to low cost storage tier for long term retention

Isilon/PowerScale Snapshot copy and Archive

1. Requires: Base product license
2. One time copy of a pre-existing snapshot for long term storage

3. Expiry of objects will not be used on the S3 bucket in this solution

Full backup Mode

1. Requires: the backup bundle or advanced license key
2. Specify the number of full copies of a folder, the full backup creates a folder on the S3 target for each full backup copy.
3. Object expiry would be set to the number of days each full backup should be retained before it is deleted.
4. NOTE: This solution requires N full copies of the folder space in the target S3 storage but offers a full backup of all data
5. Example
    a. folder X - 2 full monthly backups and each backup should be retained for 60 days, will ensure that 2 full backups exist at any one time. The oldest backup would be purged by the object expiry on the S3 bucket

Key Features

1. A PowerScale Integrated tool to sync a copy of a path(s) on PowerScale to an S3 storage bucket.
    a. Uses PowerScale snapshot change list to support fast incremental syncs.
2. Direct Restore - Restore Data from S3 to PowerScale and re-apply file meta data (owner, ACL's, group, data stamps). PowerScale file structure is restored directly from S3 storage on to the file system without an out of band copy.
    a. Redirected recall to a different cluster
    b. Recall to a different folder path
    c. Recall based on modified date range

3. Protects file metadata automatically with S3 metadata tags (ACL's, owner, group, data stamps etc).
4. MD5 checksum support for data integrity copies
5. File system to object store audit jobs
6. Bandwidth rate limiting per copy job
7. Sync mode or copy mode or both
8. **Delayed Deletes to S3 copy or no deletes on source propagated to the target storage bucket -** Recycle mode puts deleted files found during incremental sync and copies them to a special bucket to store deleted files. Setting a TTL on this bucket allows automatic purging on a differed schedule after files are deleted.

© Superna LLC

10.2. Limitations, Requirements and Supported Target List

Home Top

- Firewall Port Requirements:
- Product Limitations and S3 Target Requirements
- S3 Target Device Specific Feature Support
  - Dell ECS
  - Amazon Amazon S3
  - Microsoft Azure Blob
  - Google Cloud Storage
  - Cohesity
  - OpenIO
  - Ceph version 15 or later Octopus
  - MinIO https://min.io/
  - Scality Open Source
  - Wasabi
  - IBM Cloud Object Storage
  - BackBlaze S3
  - Cloudian

Firewall Port Requirements:

1. SSH and HTTPS 8080 REST API from Golden Copy to PowerScale IP pool in the System Zone.
2. S3 port 443 https to target S3 storage.
3. NFS access from Golden Copy and Virtual Accelerator Nodes to Systems zone IP pool
4. Full firewall port table and directions can be found here.

Product Limitations and S3 Target Requirements

1. **Data copied to the bucket using tools other than Golden Copy will be recalled but metadata will not be recalled for this data. The data must exist under the Golden Copy path in the bucket <cluster name>/ifs/xxxxx if the data is copied under this object path it will be recalled without metadata. (Requires 1.1.6 or later)**
2. **Symlink files are skipped and will not be copied**
3. **hardlink files are copied**
4. **Fast Incremental performance requires OneFs 8.2.1 or later releases. The change list API includes metadata and removes the requirement to query the file system for metadata of the files in the change list.**
    a. **Releases before 8.2.1 do not support change list + metadata, this requires additional query to the file system for each file in the change list this will impact performance of processing a large change list of files.**
    b. **Solution is to upgrade to 8.2.1 of Onefs**
5. **Copying paths longer than 1001 characters due to REST API limitation on characters per request. This means files at folders below this level will not be copied to the S3 bucket the path is calculated**

/ifs/.snapshot/.... path to file and file name. NOTE: Windows SMB path length limit is 260 characters.

6. Concurrent Running jobs - Releases above 1.1.4 21073 or later
    a. NOTE: More concurrent jobs does not mean increased throughput. Copy resources are shared across all running jobs.
    b. Number of folders configured for Incremental always scheduled concurrently is a maximum 80. It is Not supported to exceed this concurrent folder incremental limit, contact support.
    c. In order to exceed 80 folders with incremental it is possible to offset additional folders to start on different hours to stay within the 80 concurrent limit. example every 4 hours start 5 folder incremental and the next 4 hour interval start 5 different folders.
    d. Full copy concurrent testing limit of 80 jobs but only 10 jobs execute at a time with resources shared across all jobs that are running.
7. Report creation is limited to 15 M success files but can be extended by adding disk to the vm to support exporting copy job reports in json format to any file count required by allocating more disk space for logging successful file copies.
8. Concurrent running jobs display is limited to 10 active running jobs in the CLI any job submitted after 10 will not be displayed but it will be queued once one of the first 10 jobs completes, queued jobs will appear as running.
9. Largest Tested single file multipart upload is 1TB file. No architectural limit blocks larger files other than the S3 target max object size.

10. Directories with no files stored in them will not be copied to the object store as an empty directory. The directory will not be visible in the object store.

11. ADS or Alternate data stream files will not have ADS streams copied

12. Characters supported in file and directory names as per Amazon S3 specification listed here.

    a. Tested special character list for files and folders

        i. Safe Characters: ~ ` ! @ # $ % ^ & () _ + - = [] {} : ; ' , . \ : * ? " < > |

        ii. Windows does not support these characters in file or folder names \ / : * ? " < > |

        iii. Azure does not support: \ %

        iv. Language characters have been tested successfully for non ascii characters in file names or folders. Not all languages or characters have been tested.

    b. Not supported

        i. All other special characters not listed above

13. Connecting to S3 targets on the Internet requires Golden Copy IP addresses to be NATed to the Internet. This applies to AWS and Azure or any other Internet S3 target.

14. S3 Targets must support custom object properties for protecting meta data of the files and folders.

15. Isilon or Powerscale any folder can be copied including snapshots. In addition the DR cluster that is read-only due to synciq can also be used as a source of a folder copy.

16. **Recall Cluster Targets**

    a. Golden Copy Basic license 1 recall target configured at a one time.

    b. Golden Copy Advanced or Backup Bundle License recall target limit of 5 with the ability to select a target during a recall job creation

S3 Target Device Specific Feature Support

1. Dell ECS

    a. Compliance mode is supported

    b. Bucket Versions are supported

    c. Retention policies and versions at the bucket level in a combined policy is not supported by ECS

    d. Application Aware meta data content type encoding

    e. Checksum support

    f. Supports a load balance option with health check of ECS nodes + load balancing copies evenly to ECS nodes provided. Typically has higher performance than using a load balancer.

2. Amazon Amazon S3

    a. Snowball Edge

    b. AWS auth v4

    c. Transfer Acceleration Buckets. (planned future release)

    d. Client Side KMS keys - not supported - TLS inflight encryption for inflight security only.

e. Retention policies
f. Access and Security authentication only
g. Application Aware meta data content type encoding
h. Checksum support
i. Multipart upload and download
j. Object Lock support

3. Microsoft Azure Blob

a. Blob storage REST API
b. Data Box
c. Retention policies
d. checksum support
e. Multipart uploads
f. Does Not support multipart downloads
g. No content type encoding support binary/octet used for all files
h. Object lock supported release 1.1.6 or later.

4. Google Cloud Storage

a. Native SDK API
b. Checksum support
c. Multipart upload and download
d. all metadata options supported

5. Cohesity

a. See vendor documentation for versioning support, and object retention policy support.
b. NOTE: Does not support load balancing feature in Golden Copy due to no support for multi part copy of chunks load balanced across Cohesity nodes.
c. checksum support
d. Application Aware meta data content type encoding
e. Checksum support

6. OpenIO
    a. Versioning not tested
    b. Requires --meta-prefix when adding folders and value of oo- .
    c. checksum support
    d. Application Aware meta data content type encoding
    e. Checksum support
7. Ceph version 15 or later Octopus
    a. aws v4 signature only
    b. checksum support
    c. Application Aware meta data content type encoding
    d. Checksum support
8. MinIO https://min.io/
    a. aws v4 signature only

b. Requires --meta-prefix xxx (where xxx is a number to identify folder objects separately from object key paths)

c. cloud type other must be used

d. checksum support

e. Application Aware meta data content type encoding

f. Checksum support

9. Scality Open Source

a. Versioning not tested

b. cloud type other must be used

c. checksum supported

d. example add command

i. searchctl archivedfolders add --isilon gcsource --folder /ifs/archive --accesskey accessKey1 --secretkey verySecretKey1 --endpoint http://x.x.x.x:8000 --region us-east-1 --bucket bucket --cloudtype other

create a folder which will become the bucket name when created at the root of the target device.

e. Application Aware meta data content type encoding

f. Checksum support

10. Wasabi

a. Versioning tested

b. Checksum support

c. cloud type other

d. full permissions to the bucket are required

e. example command

i. searchctl archivedfolders add --isilon gcsource --folder /ifs/archive --accesskey --secretkey --endpoint https://s3.wasabisys.com --bucket gctest --cloudtype other

11. IBM Cloud Object Storage

a. Versioning tested

b. Checksum support

c. cloud type other

d. example command

i. searchctl archivedfolders add --isilon gcsource --folder /ifs/archive --accesskey --secretkey --endpoint https://<endpoint> --bucket gctest --cloudtype other

12. BackBlaze S3

a. Versioning tested

b. Checksum support

c. cloud type other

d. example command

i. searchctl archivedfolders add --isilon gcsource --folder /ifs/archive --accesskey --secretkey --endpoint https://<endpoint> --bucket gctest --cloudtype other

13. Cloudian

a. Versioning tested

b. Checksum support

c. Multi part upload

d. cloud type AWS

e. note: bucket url must resolve to the end point ip address using url access method. example **gctest.s3-us-east.internal.superna.net** must resolve in DNS on the Golden Copy VM's

f. Example Command

   i. searchctl archivedfolders add --isilon gcsource --folder /ifs/archive --accesskey --secretkey --endpoint https://s3-us-east.internal.superna.net --bucket **gctest** --region **us-east** --cloudtype aws

© Superna LLC

10.3. Networking Deployment Overview, Cloud Provider IP Ranges, Proxy Configuration

Home Top

- Overview
- PowerScale Networking Options
  - S3 Target IP ranges that must be reachable to the Golden Copy Virtual Nodes
  - Networking Options on Golden Copy
- Firewall Requirements
- Deployment Networking Diagram
- Proxy Configuration to Reach Internet Cloud Providers
  - Variable Definitions
  - Example of AWS Proxy configuration
- Disk Space Management
- Phone Home Proxy Configuration

Overview

This section provides information on how the Golden Copy VM must be configured for access to PowerScale, how files are copied to Internet S3 targets, and options that exist to configure.

PowerScale Networking Options

S3 Target IP ranges that must be reachable to the Golden Copy Virtual Nodes

If you control outbound connections, each Golden copy node will require a firewall rule providing access to the Storage services for Cloud providers.

1. AWS:
    a. If you need to secure outbound access to a range of ip addresses, you can follow AWS documents here:
        i. Range of IP's by service type https://ip-ranges.amazonaws.com/ip-ranges.json follow guide here to firewall ports needed for your region https://docs.aws.amazon.com/general/latest/gr/aws-ip-ranges.html
2. Azure:
    a. Download the IP range json file from here .
    b. Search the file for the word "AzureStorage". Locate your region to get a specific list of IP addresses that must be accessible to whitelist to the PowerScale nodes that will copy files.

Networking Options on Golden Copy

Rate limiting feature on each folder or global default can be configured within Golden Copy. See the Configuration section for details on adding a folder with a rate limit applied.

1. Default networking:
    a. Golden Copy VM's require access to ip addresses of PowerScale nodes within the System Zone over NFS and REST API port 8080.
    b. File copy traffic will use the default route table on Golden Copy virtual nodes to reach the Internet S3 targets.

Firewall Requirements

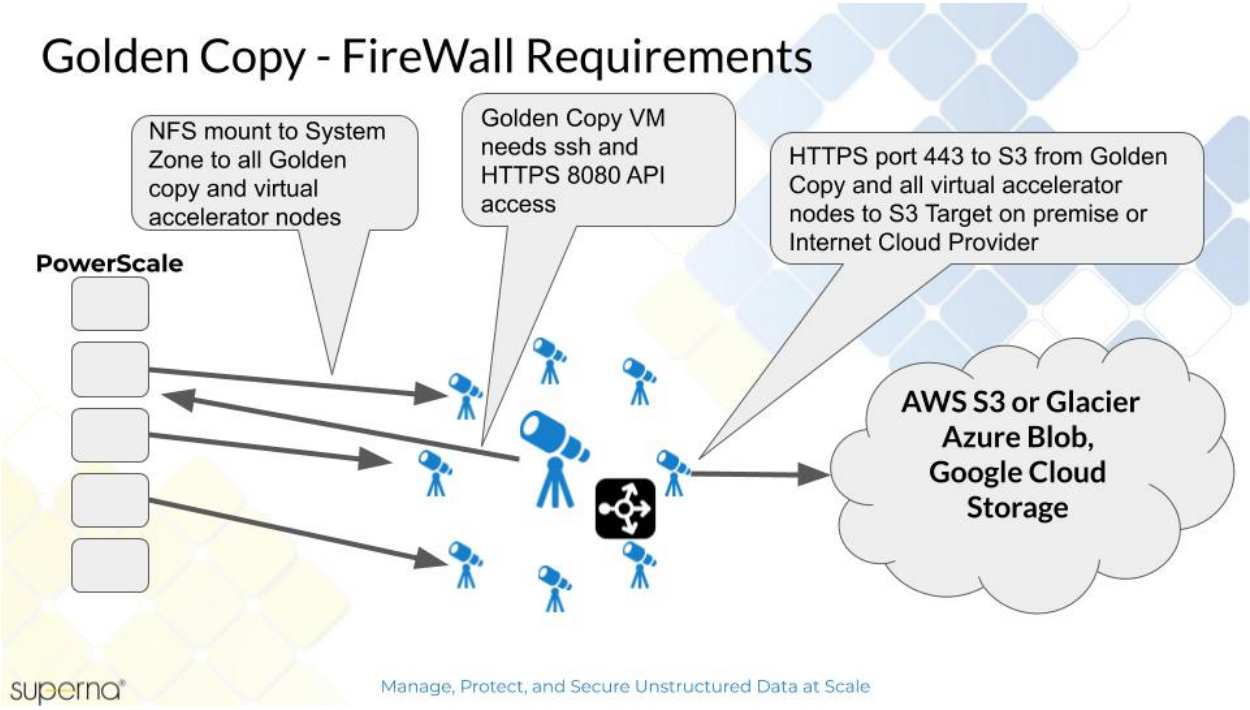

Deployment Networking Diagram

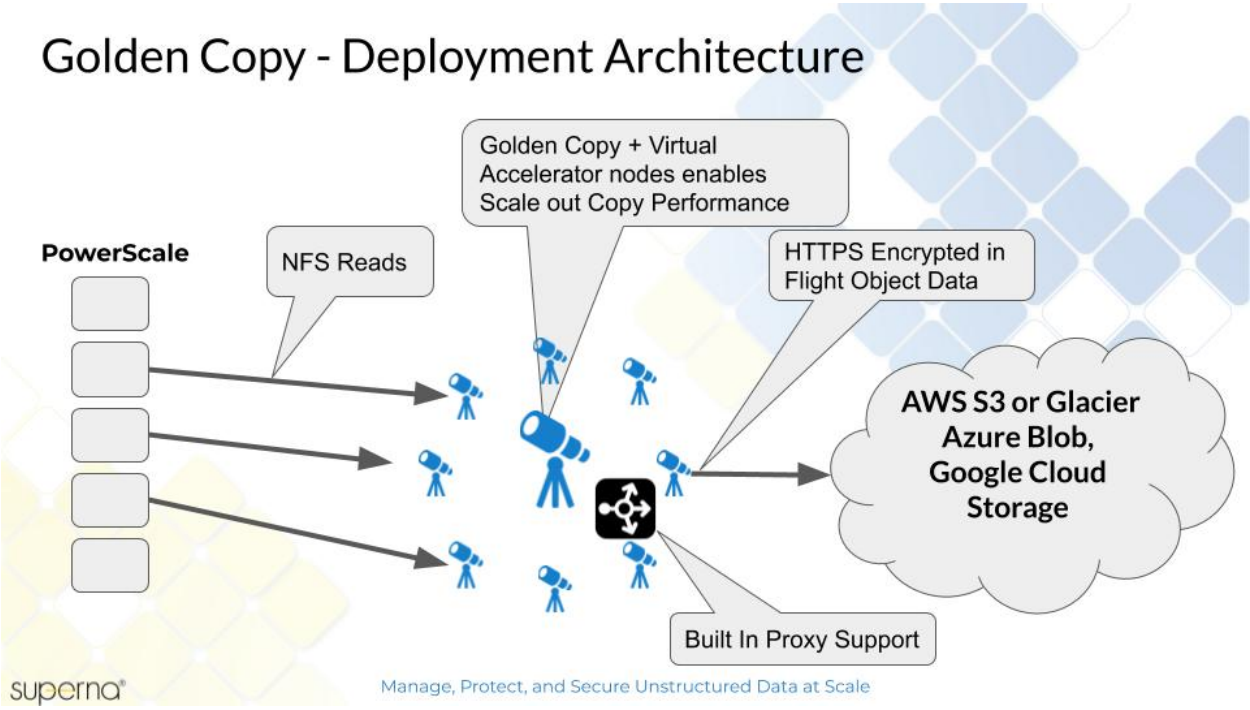

Proxy Configuration to Reach Internet Cloud Providers

1. If transparent source IP NAT is not available to reach Internet storage providers, Golden Copy supports proxy access through standard

proxy devices. Follow the steps below to enable proxy copy configuration.

   a. NOTE: if using AWS or any S3 compliant target devices use the AWS proxy configuration steps.

2. login to Golden Copy VM as ecaadmin over ssh
3. nano /opt/superna/eca/eca-env-common.conf
4. Add the variables defined below to configure a proxy target. Then restart the Golden copy software.
   a. example variable to add and then save the file with control+x key and answer yes
   b. export AWS_PROXY_HOST=x.x.x.x
   c. export AWS_PROXY_PORT=xxxx
   d. export AWS_PROXY_PROTOCOL=https
5. ecactl cluster down
6. ecactl cluster up

Variable Definitions

AWS_PROXY_HOST: Sets the proxy host the client will connect through.
AWS_PROXY_PORT: Sets the proxy port the client will connect through. Some proxies require a port example 8080
AWS_PROXY_PROTOCOL: Set the Protocol to use for connecting to the proxy. (HTTP or HTTPS)
AWS_PROXY_PASS: (optional if required by the proxy) Sets the proxy password to use when connecting through a proxy.

AWS_PROXY_USER: (optional if required by the proxy) Sets the proxy user name to use if connecting through a proxy.

Example of AWS Proxy configuration

```
export AWS_PROXY_HOST=172.25.38.42
export AWS_PROXY_PORT=3128
export AWS_PROXY_PROTOCOL=HTTP
export AWS_PROXY_USER=proxyuser
export AWS_PROXY_PASS=3y3gl4ss
```

Disk Space Management

1. The appliance comes with a 400G disk to store queued data to copy and reporting data on jobs. This space may need to be increased under the following conditions:
    a. The reporting records are unique per folder, folder counts over 25 may need to add disk space on all nodes. Suggested increments of 100GB per 25 folders.
    b. If rate limiting is applied to the WAN copy, the space consumed by queued file records can also increase age beyond the 400GB default.
    c. If exports of job summaries are run and stored on the appliance additional space may need to be added.
    d. If the error the rate is high, this will consume space until the issue is resolved.
    e. Multiple full archive jobs

Phone Home Proxy Configuration

1. To allow the phone home service to use a proxy, the operating system proxy needs to be used. Follow these steps here that also applies to Golden Copy VM.
2. Open Suse OS proxy configuration guide

© Superna LLC

10.4. Understanding Copy Performance Configurations

Home Top

- Overview
  - How to scale out and scale up copy performance
- Dell ECS Performance Copy Mode
  - How load balancing works

Overview

This section should be reviewed when deciding how to configure which Golden Copy for copying files, and which IP pool interfaces will be used to copy files. In addition, decisions on parallel copy influences the speed of copies and the bandwidth consumed for a copy job. Golden Copy also supports multipart uploads and downloads.

How to scale out and scale up copy performance

1. Deploy Virtual Accelerator nodes to increase copy performance. The configuration is 1 VM or 6 VM's. The main Golden Copy VM is the control VM and the remaining 5 VM's are only used for copying. All VM's must have an NFS mount to the source cluster.
   a. See the install guide to deploy VAN nodes here.
2. Parallel threads can be increased to use more threads on multi part uploads. The default is 10 threads.
   a. This can be changed following steps below.
   b. ssh to node 1 as ecaadmin

c. nano /opt/superna/eca/eca-env-common.conf

d. Add this tag to double the threads on large files. This should only be used if the majority of files are over 1GB in size.

e. export ARCHIVE_PARALLEL_THREAD_SDK=20

f. control+x , answer yes to save

g. ecactl cluster down

h. ecactl cluster up

Dell ECS Performance Copy Mode

1. Golden Copy supports a load balance option with health check of ECS nodes + load balancing copies evenly to ECS nodes provided.
2. This will maximize performance and eliminates the requirement to use an external load balcner that adds cost to the overall solution.

How load balancing works

1. Each Golden Copy node uses a S3 health check to each ECS node, this is performed every minute. If an ECS node is unreachable or fails the S3 health check it is removed from the pool of ECS nodes.
2. If the ECS nodes passes the S3 health check in the next 1 minute, it will be added back to the load balancing pool. The algorithm will ensure a balanced number of copies per ECS node using the parallel thread setting on the Golden copy vm.

© Superna LLC

10.5. Cloud Provider Cost Considerations

Home Top

- Overview
- Golden Copy Feature and Cost Impacts when using Cloud Storage Providers
- Cost Factors To Consider when using Cloud Storage
- Golden Copy Amazon Amazon S3 Calculator
- Cloud Provider On line Object Storage Cost Calculators

Overview

1. Several factors can affect the cost of storing data in Cloud provider storage services. This guide covers best practices for large data transfers and how to estimate monthly costs. Some Golden Copy features can incur charges when using Cloud Storage. These should be reviewed to understand potential fees incurred from the use of these features.

Golden Copy Feature and Cost Impacts when using Cloud Storage Providers

1. Cloud providers charge for api calls, uploads, downloads and tier transitions. Read the feature list below and how usage of these features can incur charges from your cloud provider.

a. **Archive job** - api costs, HTTP put API costs, storage costs per GB

b. **Recall job** - api costs for get objects and data costs from egress of

c. **Re-run Archive job** - api costs to check if a file exists in the object store

d. **Data Audit job** - api costs , get object egress data charges equal to requested GB to audit

e. **Incremental job** - api costs to post objects, storage costs

f. **Life cycle policies** in the cloud provider Storage bucket - The cost to transition objects from one tier to another can attract costs per object moved between tiers.

g. Recall data from archive tiers to active tier can attract recall costs.

Cost Factors To Consider when using Cloud Storage

1. Ingress data to a Cloud Provider typically does not have bandwidth charges. Ingress data is data copied into the storage provider.
2. API requests have a fee in most cases and each file equals 1 API call. Large files over 10MB are copied using a multi part upload for example a 100MB file is sent in 10 parts and would be 11 api calls. This can be calculated as follows: the quantity of your data that is over 10MB in file size divided by 10 can estimate the number of API calls to factor into the cost estimates. The Superna AWS calculator will do all these calculations for you to estimate your costs.
3. Egress network data charges are much higher than ingress data charges, egress data is any S3 object data downloaded from your Cloud Provider. Golden Copy recall jobs will incur egress bandwidth

charges. The quantity of data in the recall job should be reviewed before initiating a recall job to understand the cost impacts that may be charged our Cloud providers.

4. **Data Storage:** Use the information below to help estimate your long term storage costs. Sum up the data volume you plan to archive or copy to the Cloud provider, and estimate your data change rate over a month.
    a. NOTE: A GB of data on Isilon will be a GB of data in a Cloud provider service. No compression or modification is applied to the data during the copy process.
5. **Best Practice for Initial Data Loading:**
    a. Bulk data loading should always be used to speed up the transfer of a large file data sets. This option should be used for large data set files that average 10MB or more. Small file file systems do not consume a lot of bandwidth and will not be faster when using data transfer devices.
    b. Amazon Amazon S3 - Use the Snowball device and guide to bulk load data for large file file systems, see the guide here.
    c. Microsoft Azure Blob Storage - Use Data Box device and guide to bulk load data, see the guide here.

Golden Copy Amazon Amazon S3 Calculator

1. Superna offers an integrated Amazon Amazon S3 calculator that is designed to calculate costs based on how Golden Copy copies data. It requires information about your data and change range rate to summarize the costs for Standard, Glacier and Deep archive tiers in Amazon S3. This tool takes all the complexity out of figuring out your costs. This integrated tool uses Amazon pricing API and

provides accurate estimates to build your archive/backup business case.

2. Try the Eyeglass Golden Copy Amazon Amazon S3 Calculator now.

Cloud Provider On line Object Storage Cost Calculators

1. Amazon AWS - https://calculator.s3.amazonaws.com/index.html
2. Microsoft Azure - Azure pricing calculator

© Superna LLC

10.6. Supported Data Security & Storage Class Life Cycle Options

Home Top

- Overview
- Data Inflight Encryption
- Data at Rest Encryption
- Data Retention Policies
- Object Data Version Control

Overview

Data that is copied to objects can have several policies applied for security , retention and versioning. This section covers supported options.

Data Inflight Encryption

1. Adding endpoints using https ensures that inflight data will be encrypted using TLS security protocol.
2. The endpoint can use self signed or signed certificates. Certificate signing is external to Golden Copy, and no configuration is required to support signed or unsigned.

Data at Rest Encryption

1. S3 targets that support customer provided keys to encrypt data at rest would be configured on the S3 provider bucket level.
2. Consult S3 target documentation about default at data at rest encryption keys are applied without any configuration at the storage bucket level.
3. No support for object level encryption keys.

Data Retention Policies

1. Data retention for objects is configured at the bucket level using the S3 target administration policies.
2. Create different storage buckets to set different retention levels for copied objects.

Object Data Version Control

1. Version control is configured at the bucket level following S3 target documentation. Golden Copy supports versioning by updating existing objects with a newer version, **NOTE: No configuration is required within Golden Copy to use Versioning on your S3 target device**. If versioning is enabled on the storage bucket both versions will be available using S3 bucket browsing tools.
2. A future version of Golden Copy will support version aware recall feature that allows a specific version of data to be recalled based on a date range. Check documentation for the build and version number that support version aware recall. The command supports recall data older than or Newer thann x date and time.
    a. Azure object versioning configuration
        i. General Overview
        ii. How to Enable or disable Blob Versioning
    b. AWS object versioning configuration
        i. General Overview
        ii. How to Enable S3 Versioning
    c. All other supported S3 targets consult vendor documentation.

© Superna LLC

10.7. Appliance Security, Authentication & Hardening

Home Top

- Overview
- User Roles
- Authentication & Login
  - Overview
- Read Me First
- How to Enable Administrator RBAC for local OS Users or AD Users (Mandatory)
  - How to Configure Active Directory OS Login for Application Users
  - How to Configure Additional AD User Appliance Administrators
  - How to Configure Additional Local OS Appliance Administrators
  - How to create a New RBAC Operator Role for Monitoring Copy Jobs and Stats
- How to Assign CLI commands to a Role
- How to Create a Custom Role
- Managing Access to CLI, WebUI
  - How to Access the command line:
  - How to Access the WebUI:
  - Local OS administration and application cluster operations

    - How to add new local admin with Access to the WebUI and CLI
    - How to Add AD Users with Access to the WebUI:
- How to Configure password policies for local users and ban failed login attempts

Overview

This section covers login and authentication options for Golden Copy administrators.

User Roles

1. Golden Copy has an administrator role for cluster configuration and cluster operations (up, down, upgrades). The user is "ecaadmin".
2. If additional administrator users are created they are only able to configure copies and run copy jobs. Cluster operations are blocked.
3. This guide explains how to create new users, groups and create roles by assigning CLI commands to a group.

Authentication & Login

Overview

The Golden Copy vm supports OS login options and Application authentication options to allow configuration changes. The ecaadmin user is the appliance administrator that is used for managing the software operating environment changes to system configuration files and can also be used to start up and shutdown the application.

The options for OS login include locally created users and groups in the OS to control who can access CLI or Active Directory OS login can also be used.

Read Me First

1. **This feature is designed for 2 use cases**
    a. **Creating new administrators with all CLI command access with either local os login or remote AD login**
    b. **Creating a role for operators to monitor copy job progress and view stats but will unable to change configuration.**
    c. **Creating additional administrators**
    d. **NOTE: ecactl commands are restricted to the ecaadmin user that is the only appliance administrator**

How to Enable Administrator RBAC for local OS Users or AD Users (Mandatory)

1. Activate additional administrators with this command. Without this command no admin users will be allowed to run any commands other than the builtin ecaadmin appliance administrator. The builtin group for administrators is called ecactl.
    a. searchctl settings groups add --name ecactl --cmd '*'

How to Configure Active Directory OS Login for Application Users

1. **This requires joining the appliance to Active Directory.**
    a. **login as ecaadmin over ssh**
    b. **sudo -s (enter ecaadmin password)**
    c. **Type YAST**
    d. **Navigate to Windows Domain Membership**

e. 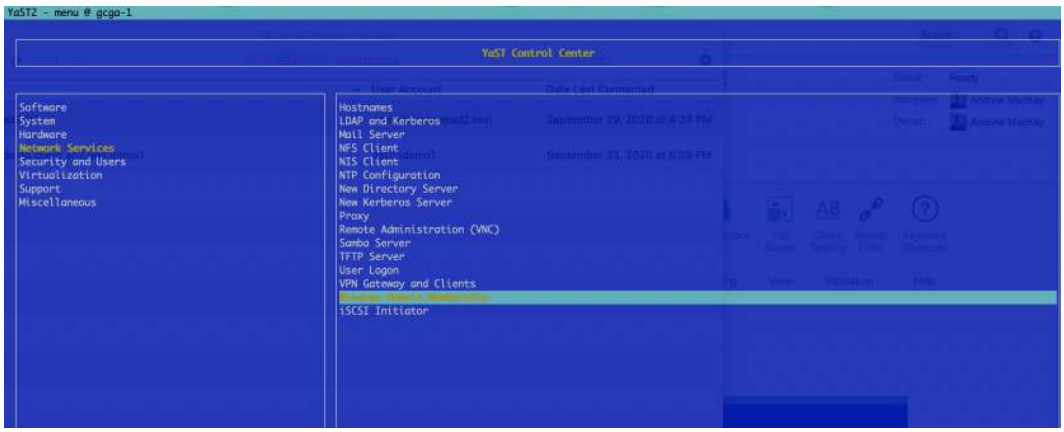

f. Enter your AD domain and select options as per the screenshot below

g.

h. Enter domain admin or an account that allows computers to be joined to the domain.

i. Once joined successfully, exit YAST using tab to select quit.

j. Test ssh login with AD

k. example login syntax for an AD domain called ad2.test with user demo1. NOTE: the double slash is required to escape slash character

   i. Linux ssh client login

      1. ssh ad2.test\\demo1@x.x.x.x OR

      2. ssh demo1@ad2.test@x.x.x.x

   ii. Putty windows tool

1. **User name:** demo1@ad2.test

l. **This allows OS login but blocks application CLI commands until the AD user is added to a role.**

m. **Proceed to** How to Configure Additional AD appliance Administrators .

How to Configure Additional AD User Appliance Administrators

1. **Follow these steps to add an AD user to the builtin ecactl local group that provides access to all searchctl commands.**
2. **NOTE: You must follow the steps to join the appliance to AD before using these procedures.**
3. **Login as ecaadmin user**
4. sudo /usr/sbin/usermod -a -G ecactl AD02\\xxxx (AD02 is the domain name in upper case and xxxx is the user name in AD to add to the group, NOTE the double slash is required as an escape character)
5. **Logout and test the user access**
6. example login syntax for an AD domain called ad2.test with user demo1. NOTE: the double slash is required to escape slash character
   a. Linux ssh client login
      i. ssh ad2.test\\demo1@x.x.x.x OR
      ii. ssh demo1@ad2.test@x.x.x.x
   b. Putty windows tool
      i. User name: demo1@ad2.test

7. test CLI commands

How to Configure Additional Local OS Appliance Administrators

1. This command will add a new OS local user that will be added to the local OS group **ecactl** providing access to full appliance cli.
2. Ssh ecaadmin@x.x.x.x
3. searchctl settings users add --name NAME (you will be prompted to enter the ecaadmin appliance administrator password)
4. Set the password for the new user
    a. sudo passwd NAME (where the user name is entered, you will be prompted to enter the eccaadmin user password, then you will be prompted to set the user password)
5. or searchctl settings users remove --name NAME (you will be prompted to confirm)
6. This new user will be able to login and run all commands required to manage the appliance, the user is added to the builtin administrator group called **ecactl**. **NOTE: activation of all commands must be completed to enable commands.**
    a. To view the users in this group **cat /etc/group**
7. **Login over ssh to test CLI commands**

How to create a New RBAC Operator Role for Monitoring Copy Jobs and Stats

1. Use this command to create a new group role to assign CLI commands to create a user role that defines an **Operator** role. This

user will be able to monitor copy jobs and statistics on folder configurations.

2. ssh to the appliance as ecaadmin
3. sudo /usr/sbin/groupadd **operator**
    a. (note: do not use spaces or special characters, when prompted for a password enter the ecaadmin password)
4. Add the new Group to the appliance with operator role commands
    a. searchctl settings groups add --name operator --cmd 'searchctl archivedfolders list' 'search jobs *' 'searchctl archivedfolders stats *'
5. Add an AD user to the new operator group **(NOTE: double slash is required to escape the \ character, use this AD syntax when adding the user)**
    a. sudo /usr/sbin/usermod -a -G operator AD02\\demo1
6. (optional) How to remove a user from a group **(note: double slash to escape the \ character is required for AD users)**
    a. sudo gpasswd -d AD02\\demo1 ecactl
7. **example login syntax for an AD domain called ad2.test with user demo1. NOTE: the double slash is required to escape slash character**
    a. Linux ssh client login
        i. ssh ad2.test\\demo1@x.x.x.x OR
        ii. ssh demo1@ad2.test@x.x.x.x
    b. Putty windows tool
        i. **User name:** demo1@ad2.test
8. Test CLI commands

How to Assign CLI commands to a Role

1. **Overview**
    a. **The commands below allow creation of customer user roles based on a local OS group that can contain local users or AD users. The CLI commands can be assigned to the group. A default group is created named ecactl that is reserved for appliance administration only. New groups can be created using the steps above to create additional roles, the Operator role is defined in this guide.**
2. **Usage to add and remove CLI commands from a group**
    a. searchctl settings groups add --name <group name> --cmd COMMANDS [COMMANDS ...]
    b. searchctl settings groups remove --name <group name> --cmd COMMANDS [COMMANDS ...]
3. **Examples of adding CLI commands to an existing group**
    a. Giving access to cli commands as an administrator and all subcommands.
        i. searchctl settings groups add --name ecactl --cmd '*'
    b. Giving access to only searchctl archivedfolders list
        i. searchctl settings groups add --name DemoGr1 --cmd 'searchctl archivedfolders list'
    c. To add multiple commands at once use a space between each command

i. searchctl settings groups add --name DemoGr1 --cmd 'searchctl archivedfolders *' 'searchctl jobs *'

How to Create a Custom Role

1. **Overview**
    a. This example will create a new role with a group called role1 and assumes AD users are used, in this example the user is AD02\demo1 user in the domain called AD02. The prerequisite requires the appliance is joined to Active Directory following steps in this guide.
    b. This role will provide monitoring and configuration commands to add or delete or modify folder configurations and start archive copy jobs
2. Login as ecaadmin user over ssh
3. Create the group
    a. sudo /usr/sbin/groupadd role1
4. Add a user to the group
    a. sudo /usr/sbin/usermod -a -G operator AD02\\demo1
5. Add CLI commands to the group called **role1**
    a. searchctl settings groups add --name **role1** --cmd 'searchctl archivedfolders *' 'searchctl jobs *' 'searchctl archivedfolders stats *'
6. Done

Managing Access to CLI, WebUI

How to Access the command line:

1. ssh to the appliance with a local user account.
    a. The default user is: "ecaadmin", and the default password is: "3y3gl4ss"
    b. **example login syntax for an AD domain called ad2.test with user demo1. NOTE: the double slash is required to escape slash character**
        i. Linux ssh client login
            1. ssh ad2.test\\demo1@x.x.x.x OR
            2. ssh demo1@ad2.test@x.x.x.x
        ii. Putty windows tool
            1. **User name:** demo1@ad2.test
    c. done

How to Access the WebUI:

1. https://x.x.x.x (this will present the login page and login with user@domain name and password. NOTE: The user must be added as a local admin first see below) .

Local OS administration and application cluster operations

1. cluster CLI commands, webUI and upgrade administrator user is "ecaadmin".
2. SSH login or webUI login.
3. Configuration and appliance cluster operations.
4. OS configuration changes requires sudo -s to become root user and this is granted to the ecaadmin appliance administrator only.

How to add new local admin with Access to the WebUI and CLI

1. Create new local OS administrator with the command **command in this section. This will add the user to the CLI admin role.**
2. This command will add the user to the webUI.
    a. searchctl settings admins add --name <OS user name> --local
3. Then set the password for the webUI login
    a. Installation requires the WebUI password to be set. To reset or change this password follow these steps:
    b. Login to node 1 over ssh as ecaadmin user and run the command below
    c. NOTE: Enter the password after the user name
    d. ecactl cluster exec "htpasswd -b /opt/superna/eca/conf/nginx/.htpasswd <user name> password"
4. done. The new password is active immediately on all nodes.

How to Add AD Users with Access to the WebUI:

1. Active Directory users can only access the WebUI, no CLI access will be granted.
    a. The important thing to note is that login will be proxied through to the isilon authentication provider. This means the user you're attempting to give access must have access to at least one SMB share on your isilon to test the password. TCP Port 445 must be open from node 1 to the Isilon
2. searchctl settings admins add --name <AD user name> (NOTE: syntax is user@domain name)

How to Configure password policies for local users and ban failed login attempts

The all product steps can be followed here for password complexity and life time policies as well as ban failed login attempts

Operations Guide - All products Hardening Guide

© Superna LLC

10.8. Golden Copy Configuration Steps

Home Top

- Supported S3 Protocol Targets
- How to prepare a Cluster for Golden Copy Management
- Quick Start Steps for Golden Copy Setup
  - Prerequisites:
  - License Keys
  - How to Add a cluster to Inventory
  - Archiving to a Dell ECS
  - Archiving to AWS
  - Archiving to Google Cloud Storage
  - Archiving to Azure Blob Storage
  - Archive to Cohesity
  - Archive to BackBlaze
- Most Common Every Day CLI Commands to Manage Archive jobs
- How to Test Data Copy and Start the copy Job for a folder that has been added to Golden Copy
  - How to test file copy target permissions before starting a copy job:
- How to Modify the Cluster service account or password
- How to add folders to be copied or synced
- Overview of Archive Job Types

    - Concurrent Job Prioritization Incremental vs Full
    - How to configure Incremental jobs to Sync Deleted Files to S3 Object Store
- How to start a Full or Incremental Archive Job
- How to Schedule Jobs on Folders (Full Archive, Incremental Archive or Archive Data Audit jobs)
    - Overview of Archive Job Types
    - Configuring Archive Job Schedules on Folders
        - Add an Incremental Schedule to a folder
        - Modify a Schedule on a Folder
        - Disable a Schedule on a folder
        - Add a Data Audit job Schedule to a folder
- Monitor, View running Jobs, show Job History, Show folder job history, Summarize Job stats, Monitor Progress, Auto Email Progress and Cancel a Copy Job
    - How to Enable Auto Progress Emails for copy Jobs
- How to Manage Folders (List , Modify and Remove)
    - list (list configured folders)
    - modify (change configuration of an existing folder)
    - How to remove an archive folder
    - How to Re-run Only Failed Files from a Copy Job
- How to Recall Data from Object Back to File
    - Overview:
    - Backup and Restore Use Cases

    - License Key Dependancies
  - Limitations:
  - Requirements:
  - Prerequisite NFS Export Configuration Steps
  - Logical Diagram of a Recall
  - Recall command syntax and options
  - How to recall object data to the Recall Staging Area
- How to Monitor Copy Job Performance and Job Summary logs and Error Logs
  - Stats Command for real time job throughput monitoring
  - Job View command monitors progress in MB and file count with % completion
  - How to Monitor Ethernet Interface Mbps during a Copy Job
  - How to view the Detailed json Copy Job Logs
  - How to view copy job errors on failed copies
- How to Manage File Copy Performance
  - How to Increase Copy Performance with concurrent file and large file thread count
  - How to Shape Bandwidth of Archive Jobs
    - Overview:
- How to Monitor Network bandwidth Usage from the appliance
- How to Copy a Snapshot to S3 Storage
- How to Configure Delayed Deletes with Sync Mode
  - Overview

    - Requirements:
    - How to Configure Delayed Delete Mode
- How to list and change System Scheduled Tasks
- How to Configure a Folder Alias to Handle: Cluster Decommission Use Case, Source Cluster Name Change, Switch to DR cluster as data source and Data full Copy
- Storage Target Configuration Examples
    - Dell ECS Bucket Creation Walk Through
    - Amazon AWS Bucket Creation Walk Through
        - How to setup minimum permissions in AWS for Golden Copy
            - S3 Policy Permissions Requirements
            - Quick Start Method
            - Complete Steps to Create User and Policy Following All Steps (skip if you used quick start above)
        - How To restrict S3 Bucket Access by Source Public IP of your Data Center
        - How to Enable and Use Accelerated Transfer Mode on a bucket
    - Google Cloud Storage Creation Walk Through
- Azure Blob Storage Creation Walk Through
    - How to create Azure Storage Account
        - How to create an Azure Blob Container in a Storage Account

- Cohesity Walk Through Example
  - How to Create the Storage View and Service Account User on a Cohesity Storage Array
- Appliance Global Settings Configuration
  - How to set Default Settings for Snapshot Expiry for all Folders
  - How to set File Checksum Global Settings for All Folders
- Advanced Configurations to Appliance Configuration

Supported S3 Protocol Targets

Review supported S3 Targets and limitations here.

How to prepare a Cluster for Golden Copy Management

1. **Mandatory Step - Create the service account for the Search & Recover/Golden copy products on the source cluster, and configure the sudo file on the source cluster.**
2. **Mandatory Step Time Sync:**
   a. **Time sync of the cluster and Golden Copy must have accurate time sync and be within 15 minutes of NTP synced time source. If this is not done then S3 targets will reject all uploads with an error message due to time skew too great.**
   b. **Setup NTP on your source Isilon/Powerscale cluster and Golden Copy to ensure accurate time.**

Quick Start Steps for Golden Copy Setup

This quick setup guide provides the steps to get configure a folder for archive copy or sync jobs and a link to learn more detailed CLI options.

Prerequisites:

1. All searchctl commands must be run as the ecaadmin user from Golden Copy node 1
2. Adding the PowerScale by IP address is recommended.
3. The eyeglassSR service account must be created following the minimum permissions guide for Golden Copy. The guide is here.
4. Default login is user ecaadmin and password 3y3gl4ss

License Keys

1. Copy license zip file to Search node /home/ecaadmin directory and change permissions chmod 777 .
2. searchctl licenses add --path /home/ecaadmin/**<name of zip>**.zip .
    a. Verify the license is installed.
3. searchctl licenses list
4. Get help on licenses
    a. searchctl licenses --help
5. searchctl licenses uninstall (Removes all of the licenses on the system).
6. searchctl licenses applications list (Lists all the application types that are licensed on the System, used on the unified Search & Recover and Golden Copy cluster).
7. searchctl isilons license --name NAME --applications APPLICATIONS - Use this command to assign a license key to a cluster. Example to assign the advanced license key

a. searchctl isilons license --name <cluster name> --applications GC

How to Add a cluster to Inventory

1. searchctl isilons add --host < **ip address of Isilon in system zone>** --user eyeglassSR --applications {GC, SR}
    a. **[Advanced backup bundle license key required]** [**--goldencopy-recall-only**] Use this option to add a cluster for a redirected recall job, this cluster type is available with the backup bundle or the upgrade to the advanced license key.
    b. [**--applications APPLICATIONS**] This is required parameter to assign the cluster to the search application, or the Golden copy application. For Golden Copy product enter **GC**
    c. **NOTE: Once a license is assigned to a cluster it is locked and cannot be removed. A unlock license key must be purchased to be able to re-assign a license.**

Archiving to a Dell ECS

1. searchctl archivedfolders add --isilon gcsource --folder /ifs/archive --accesskey CqX2--bcZCdRjLbOlZopzZUOsl8 --secretkey zZMDh3P7fcMD2GUvLNK5md7NVk --endpoint https://x.x.x.x:9021 --bucket <bucket name> --cloudtype ecs --endpoint-ips x.x.x.x,y.y.y.y,k.k.k.k
2. See Walk through on Setting up the S3 Bucket here.
    a. NOTE: --isilon is the name of the cluster not the ip address. Get the name using searchctl isilons list
    b. NOTE: change the yellow highlights with correct values.

c. Supports a load balance option with health check of ECS nodes + load balancing copies evenly to ECS nodes provided. Requires the --endpoint-ips flag and specify data ECS ip addresses.

d. NOTE: The following settings are needed in the /opt/superna/eca/eca-env-common.conf. The S3 heartbeat is enabled by default and will remove an ECS node from load balance table if it becomes unreachable. This setting enables round robin load balance algorithm. A future release to support alternate load balancing mode for even connections.

    i. ssh to node 1

    ii. nano /opt/superna/eca/eca-env-common.conf

    iii. paste the setting below

    iv. Control+x answer yes to save

    v. Cluster down up is required.

    vi. export ARCHIVE_ENDPOINTS_ROUND_ROBIN=true

e. NOTE: A service account user should be created.

f. NOTE: replace endpoint-IPS with data nodes on ECS that can receive copied data. Alternate syntax is x.x.x.x-y.y.y.y to use a range of IPS.

g. NOTE: A dedicated View should be created for S3 storage without any other protocols enabled.

h. NOTE: the access key and secret key can be found by logging into the console select Admin menu --> Access Management -->, click on your user id and record the access ID and secret ID. You need to login to the console as the user configured for authentication to get the S3 keys.

Archiving to AWS

1. searchctl archivedfolders add --isilon prod-cluster --folder /ifs/data/policy1/aws --accesskey AKIAIs3GQ --secretkey AGV7tMlPOmqaP7k6Oxv --endpoint s3.**ca-central-1**.amazonaws.com **--region ca-central-1** --bucket mybucketname --cloudtype aws
    a. NOTE: The region is mandatory field with Amazon S3 .
    b. NOTE: The endpoint must use the region encoded URL. In the example above the region is ca-central-1 and is used to create the end point URL.
    c. See how to configure an AWS Bucket

Archiving to Google Cloud Storage

1. **Get the Google authentication key following the Google Cloud Storage Walk Through Guide. Copy the key to the node 1 of Golden Copy using winscp and copy the file to /home/ecaadmin**
2. **searchctl archivedfolders add --isilon gcsource --folder /ifs/archive --secretkey /home/ecaadmin/ <service account gcs keyname.json> --bucket mybucketname --cloudtype gcs**
    a. **Change all yellow highlights to the correct value for your environment**

Archiving to Azure Blob Storage

1. searchctl archivedfolders add --isilon gcsource --folder /ifs/archive --secretkey NdDKoJffEs9UOzdSjTxUlaE9Xg== --endpoint blob.core.windows.net --container gc1 --accesskey **<storage account name>** --cloudtype azure

a. NOTE: The storage account name is used as the access key with Azure.

b. NOTE: Get the access keys from the Azure console.

c. NOTE: The container name must be added with the --container flag.

d. NOTE: cloudtype flag must be azure.

e. See how to configure a container in Azure.

Archive to Cohesity

1. searchctl archivedfolders add --isilon gcsource --folder /ifs/archive --accesskey CqX2--bcZCdRjLbOlZopzZUOsI8 --secretkey zZMDh3P7fcMD2GUvLNK5md7NVk --endpoint https://x.x.x.x:3000 --bucket <bucket name> --cloudtype other
2. See Walk through on Setting up the S3 Bucket here.
    a. NOTE: change the yellow highlights with correct values.
    b. NOTE: A service account user should be created.
    c. NOTE: A dedicated View should be created for S3 storage without any other protocols enabled.
    d. NOTE: the access key and secret key can be found by logging into the console select Admin menu --> Access Management -->, click on your user id and record the access ID and secret ID. You need to login to the console as the user configured for authentication to get the S3 keys.
    e. See how to configure S3 buckets with Cohesity.

Archive to BackBlaze

1. searchctl archivedfolders add --isilon gcsource --folder /ifs/archive --accesskey 000257a71c6ed0001 --secretkey K000xlgzxveQ --endpoint https://s3.us-west-000.backblazeb2.com --bucket gctest --cloudtype other
    a. NOTE: the URL may be different for your bucket, replace all highlighted settings

Most Common Every Day CLI Commands to Manage Archive jobs

1. These commands are most commonly used for day to day. Many advanced flags exist on commands that should be reviewed. The commands below are the minimum flags needed to complete the task.
2. Listing running jobs
    a. searchctl jobs running
3. Listing history of all previous jobs
    a. searchctl jobs history
4. Listing all configured folders and the folder ID's and paths
    a. searchctl archivedfolders list
5. How to run an archive job on a folder
    a. NOTE: This will not recopy everything, it will check the target if the file exists and will skip it and only copy new files missing from the target or update the target if a file changed. It will not detect deleted files. This requires an incremental job to be configured on the folder. Steps in this guide to configure.
    b. searchctl archivedfolders archive --id xxx (xxx is the folder ID found using the list folders command above)
6. Follow the progress of a running job

   a. searchctl jobs running (get list of running job ID's)
   b. searchctl jobs view --follow --id xxxxx (xxxx is the job ID name from the step above)
7. Get the processing rates of a folder that has a running job
   a. searchctl archivedfolders list (get the folder ID)
   b. searchctl stats --folder xxxx (folder ID of the path with a running job from the above step)

How to Test Data Copy and Start the copy Job for a folder that has been added to Golden Copy

How to test file copy target permissions before starting a copy job:

1. Before starting a copy job, it is best practice to test file copy permissions with the test feature.
2. searchctl archivedfolders test --id <folderID> (list folder ID searchctl archivedfolders list) **The command must pass all tests before starting an archive job. Resolve all issues.**
3. This command will complete the following validations:
   a. Isilon/PowerScale S3 connectivity test (port test)
   b. File upload to S3 target test
   c. Validation this file was copied to S3 target
   d. Deleting this file from S3 test

How to Modify the Cluster service account or password

1. searchctl isilons modify
    a. --name NAME (name of cluster to modify)
    b. [--ip IP] (change ip in system zone to another cluster node)
    c. [--user] (service account name change , should always use eyeglassSR unless directed by support)
    d. [--update-password] (service account password)

How to add folders to be copied or synced

1. searchctl archivedfolders [-h] {add,history,rerun,errors,archive,metadata,s3stat,recall,list,modify,remove,test,configure,getConfig,stats,export,audit,notifications}

   --isilon HOST --folder PATH [--accesskey ACCESSKEY] [--secretkey SECRETKEY] [--endpoint ENDPOINT] [--bucket BUCKET] [--region REGION]

   [--container CONTAINER] --cloudtype {aws,ecs,azure,other,gcs,blackhole} [--recyclebucket TRASHBUCKET] [--skip-s3-file-exists {True,False,true,false}]

   [--endpoint-ips ENDPOINT_IPS] [--meta-prefix METAPREFIX] [--includes INCLUDES] [--excludes EXCLUDES] [--incremental-schedule INCREMENTALCRON]

[--full-archive-schedule FULLCRON] [--archive-data-audit-schedule ARCHIVEDATAAUDIT] [--tier TIER] [--backup-num BACKUPNUM] [--cluster-name CLUSTERNAME]

a. add (add a folder for ECS, AWS, Azure)
b. --path (PowerScale path to copy or sync example /ifs/data/projectx)
c. --force (used to add AWS Snowball device and bypass connectivity checks to AWS, only use when following the Snowball guide)
d. [--tier] default is **standard (AWS), Cool (Azure)**
    i. Requires Golden Copy Advanced license or Backup Bundle license
    ii. **Requires 1.1.6**
    iii. **Azure**
        1. flag to specify the tier the API calls are sent to, this should match the container tier configuration options are Access tier for Azure e.g. **hot, cool, archive**) Without this flag the default is cold.
    iv. **AWS**
        1. specify AWS tier using (**standard** (default), **standard_IA, glacier, glacier_deeparchive**)
e. --cloudtype {**aws,ecs,azure,other, gcs, blackhole**} (type of target storage)
    i. the Other type can be used to add S3 storage targets for testing before it is qualified formally and some S3 targets will require other.

ii. other - AWS v4 S3 authentication signature

iii. otherv2 - AWS v2 S3 authentication signature

f. [--region **REGION**] (required for AWS)

g. --cluster-name - Allows creating an alias for the root folder name in the storage bucket. Use a string to replace the actual cluster name in the folder used to copy all the data under this root folder name. See the use cases for this feature here.

h. --bucket **BUCKET** (required for all storage targets, except Azure which uses the --container flag)

i. --container <container name> (Azure only requires the container flag)

j. --endpoint **ENDPOINT** (required for ecs to include URL and port used for storage buckets, for azure the URL needs to include the storage account name example blob.core.windows.net)

k. --secretkey **SECRETKEY** (required for all storage targets)

l. --accesskey **ACCESSKEY** (required for all storage targets for Azure this is the storage account name)

m. [--skip-s3-file-exists] {true, **false**} (this option defaults to false which means S3 is checked first to verify if the file to be copied already exists in the S3 bucket with the same last modified date stamp. If set to true this check will be skipped and all files will be copied to the S3 bucket in a full copy mode. This will overwrite all files in the storage bucket even if they already exist in the Storage bucket.

i. **Use this option to avoid AWS or Azure fees for issuing LIST, Get commands if you have a large number of files in the path**

n. [--recyclebucket TRASHBUCKET Name] - Enter the storage bucket name to store deleted files detected during sync mode operations. The deleted files will be copied to the recycle bucket where a TTL can be set on the bucket to keep deletes for a period of time before they are deleted permanently . **example --recyclebucket mytrashbucketname**

o. [--prefix xxx] This is an option allows Copy Mode to run but insert a prefix into the storage bucket path when the copy runs. This would be useful to make a copy into a new location in the storage bucket without updating the existing path.

   i. <bucket_root>/<cluster_name>/[optional prefix]/ifs/

   ii. example --prefix temp will insert temp into the path when you want a temporary copy in the S3 bucket.

p. [--endpoint-ips] (list of ip or range to load balance ECS nodes, only used for ECS targets or Other targets that support multiple endpoints). Check Limitations and Requirements for S3 Targets that support load balancing of multi part uploads.

   i. NOTE: replace endpoint-IPS with data nodes that can receive copied data. Alternate syntax is x.x.x.x-y.y.y.y to use a range of IPS.

q. [--meta-prefix METAPREFIX] - The default prefix is x-amz-meta- for storing meta data properties on objects. Some S3 targets require custom meta data in the http headers to be used. Example open-io requires oo- to be used. This flag

allows changing the meta data http header tag for S3 targets that require this.

r. Checksum control on upload is now managed globally. This will calculate a checksum to include in the headers to be validated by the target before confirming upload was successful. See link here to configure.

s. Scheduling jobs

   i. The syntax for the Cron for all jobs see scheduling examples here.

   ii. [--incremental-schedule INCREMENTALCRON]

   iii. [--full-archive-schedule FULLCRON]

   iv. [--archive-data-audit-schedule ARCHIVEDATAAUDIT] - compares the file system data to the target bucket and will ensure they are in sync with add new files, upload modified or delete files on the target. This is a long running job and should be scheduled at most weekly or monthly.

t. [--includes INCLUDES] [--excludes EXCLUDES] - These optional flags allow including files or folders with pattern matching to either include files and folders or exclude from the copy process. (**Release 1.1.4 or later**)

   i. --includes - File paths matching this glob will be included in the archiving operation. If not specified, all files will be included.

   ii. --excludes - File paths matching this glob will be excluded from archiving. This flag only applies to those files that

are included by the --include flag. If not specified, no files will be excluded.

iii. Examples:

1. Exclude everything in the user's appdata profile: --exclude '/ifs/home/*/AppData/**'

Only archive docx and pdf files, and exclude everything in a tmp directory: --include '*.pdf,*.docx' --exclude '/ifs/data/home/tmp/**'

Only archive docx, pdf and bmp files --include '*.pdf,*.docx,*.bmp'

Archive all files except those in AppData, but only do full content for pdf and docx --exclude '/ifs/home/*/AppData/**'

Overview of Archive Job Types

2 types of folder modes exist full archive and incremental.

1. **Full Archive:** The full archive will tree walk a folder identify which files already exist on the target object store and skip them or copy new files or update modified files. It will not delete files that exist on the target but do not exist on the source file system path.
    a. If you run a Full archive job multiple times it will only copy new files, modified files detected on the tree walk. This acts like an incremental and can locate missing data in the S3 target.

2. **Incremental Archive:** This mode is enabled with a schedule on the folder. This mode uses the change list and creates snapshots that are compared to detect created, modified and deleted files to copy to the target S3 object store. This mode will default to a mode that will note delete files from the object store if files are found to be deleted on the file system.
    a. See the procedure below to enable deleting files from S3 storage during incremental sync jobs.
    b. An alternate delete mode exists called delayed deletes. This mode allows a folder to be configured with a second storage bucket to hold deleted objects when they are deleted from the file system. This can be reviewed here.
    c. The default setting will sync deletes from the file system to the S3 target during incremental.

Concurrent Job Prioritization Incremental vs Full

1. When full and incremental jobs run at the same time a QOS setting can be a changed to prioritize one job type over the other.
2. The default setting will prioritize incremental jobs over full archive jobs.
    a. This means resources will used to archive incremental queues at the cost of copying any files from back log in the full archive jobs. The full archive job may stop completely or progress very slowly until the incremental back log is completely copied. This process will repeat each time an incremental job starts.
3. How to change the default to prioritize full archive jobs.
    a. Login to node 1 as ecaadmin

b. nano /opt/superna/eca/eca-env-common.conf
    i. add these variables to the file
    ii. export ARCHIVE_FULL_TOPIC_REGEX="archivecontent-*"
    iii. export ARCHIVE_INCREMENTAL_TOPIC_REGEX="nomatch"
c. save with control+x answer yes to save
d. Restart cluster for the change to take effect
    i. ecactl cluster down
    ii. ecactl cluster up

How to configure Incremental jobs to Sync Deleted Files to S3 Object Store

1. This is a global setting for all incremental jobs and will detect deleted files on the file system and will delete the object in the S3 target storage.
2. NOTE: The default wil sync deletes.
3. Login to Golden Copy as ecaadmin
    a. nano /opt/superna/eca/eca-env-common.conf
    b. add a line to this file by copying and pasting the line below
    c. **export ARCHIVE_INCREMENTAL_IGNORE_DELETES=false (this is the default change to true to leave deleted files as objects)**
4. save the file with control+x answer yes

5. Shutdown the cluster and start up again to take effect
    a. ecactl cluster down (wait until it finishes)
    b. ecactl cluster up
6. done

How to start a Full or Incremental Archive Job

1. searchctl archivedfolders **archive**
2. --force (used to add AWS Snowball device and bypass connectivity checks to AWS, only use when following the Snowball guide)
3. --id ID (folder id)
4. --uploads-files (requires 1.1.6 build > 21124) accepts a file with full path to files with one file per line (carriage return after each file). A text file with a path per line will be accepted and each file will have it's meta data queried and included with the copied object.
    a. NOTE: file path must be full path to file example /ifs/xxx path to file and is a case sensitive file path. The case must match the actual file system case for path and file name.
    b. NOTE: The search export file format has been deprecated for a flat file format with only file path.
    c. NOTE: Make sure only 1 file per line with a carriage return at the end of each line.
5. --incremental (requires 1.1.4 build > 21002) This option will run an on demand snapshot based changelist to detect created, modified and deleted files since the last incremental job ran and copy these changes to the S3 target configured.

6. --follow (Requires 1.1.4 build > 21002) the archive job will be started and will move directly to a monitor UI to view progress without needing to use searchctl jobs view.
7. --auto-rerun this will queue all failed copies into a new job to automatically retry all failed copies. Requires 1.1.4 or later
8. [--recursive {true,false}] (can be used to update a single path only or recursive update to copy all data under a path to the storage bucket, **this is optional and the default without using this flag is recursive copy of all data under the archived path entered**)
9. [--subdir SUBDIR] (**if a recopy of some data is required under an existing archived folder, or a new folder was added under an archived folder, this option allows entering a path under the archived path to copy only this subfolder, this can be combined with --recursive option if required.**)
10. [--s3update] **File System to S3 Audit Feature** compares S3 bucket to PowerScale path and fixes any differences between the folders and files and will delete files in the S3 bucket that no longer exist in the file system path.
    a. **NOTE: Use this to audit the file system path to the S3 bucket to remove files from S3 that do not exist in the file system. This can be used on folders configured for copy mode versus sync mode that will sync deleted files into the storage bucket.**
11. --snapshot <snapshotName> if the name of the snapshot is added with this option, then no new snapshot will be created for the archive path and the snapshot name path will be used as the source of the file copy. Normally, a snapshot is created when a folder archive job is started and data is copied from snapshot. This option allows existing snapshots to be used as the source of an archive job.

a. When to use this option:

i. This option allows an existing snapshot to be used. This option would be used to copy a previous snapshot of a path to S3 storage for long term storage and allow deletion of the snapshot to free up usable space on the PowerScale once the file copy completes.

ii. The path referenced by the snapshot name must match exactly the folder path added to Golden Copy

iii. The path copied in the S3 bucket will match the folder path added to Golden Copy

iv. NOTE: The full inflated size of the path the snapshot is protecting will be copied. Snapshot copies are not space efficient when copied from the PowerScale to external storage.

12. --skip-acl Requires 1.1.4 update 2 or later. This will skip ACL permissions encoding but will encode all other metadata. This would be used if it is desired to not copy security information into objects where it is visible.

How to Schedule Jobs on Folders (Full Archive, Incremental Archive or Archive Data Audit jobs)

Overview of Archive Job Types

1. **2 scheduled job types for backup**

   a. **--full-archive-schedule Use this flag to identify this job type. This will walk a file path copy all files and will check the S3 target if the file already exists and will skip files if needed while**

copying. The comparison is done using the last modified date stamp on a file.

b. --incremental-schedule Use this flag to identify this job type. This job will use the Isilon/Power Scale change list API to snapshot that folder path and offload the comparison to the cluster to return file system changes for incremental always syncing. The created, modified files will be synced to the S3 target.

   i. NOTE: The default behavior will NOT sync deleted files to the S3 target. This means deleted files will be detected but will remain on the S3 target. To change this default behavior see the advanced configuration here.

c. --archive-data-audit Use this flag to identify this job type. This job will audit the file system path and the S3 bucket to identify files that are missing in S3 but it will also identify files that no longer exist no the file system but also exist on the S3 target. This will delete objects when a file has been deleted from the file system. Recommendation: Run this job on demand or on a schedule to audit the S3 object store data.

Configuring Archive Job Schedules on Folders

1. This feature requires 1.1.4 build > 178
2. NOTE: times scheduled will be in GMT time zone.
3. The input is a cron string to create the interval for the schedule. See the cron web site for assistance in creating a cron string https://crontab.guru/examples.html
4. The (other parameters) is a place holder for any other options being added or modified in the examples below)

Add an Incremental Schedule to a folder

1. Run Every Hour
    a. searchctl archivedfolders add (other parameters) **--incremental** "0 * * * *"
2. Run Every 2 hours
    a. searchctl archivedfolders add (other parameters) --incremental "0 */2 * * *"
3. Run Every 6 hours
    a. searchctl archivedfolders add (other parameters) --incremental "0 */6 * * *"
4. Run Once a Day at midnight
    a. searchctl archivedfolders add (other parameters) --incremental "0 0 * * *"

Modify a Schedule on a Folder

1. This command uses a different syntax see below
2. Once per day examples of incremental and data audit job types
    a. searchctl archivedfolders modify --id <folder id> (other parameters)**--incremental** "0 0 * * 0"
    b. searchctl archivedfolders modify --id <folder id> (other parameters)**--archive-data-audit** "0 0 * * 0"

Disable a Schedule on a folder

1. **searchctl archivedfolders modify --id <folder-id>** --full-archive "NEVER" --incremental "NEVER"

Add a Data Audit job Schedule to a folder

This job type compares all data on the file system to the object store target and ensures missing data is added, deleted data is removed from the object store and updated files are updated in the object store. This is a data compare job that can be scheduled to run once a week or month on a folder. NOTE: This can take hours to run.

1. searchctl archivedfolders add (other parameters)--archive-data-audit "0 0 * * 0"
    a. Suggested schedule is weekly on Sunday

Monitor, View running Jobs, show Job History, Show folder job history, Summarize Job stats, Monitor Progress, Auto Email Progress and Cancel a Copy Job

1. These commands will show progress of files as they are copied, with real-time updates every few seconds.
    a. searchctl jobs view --id job-1574623981252553755794 --follow **(replace the job name with the name returned from the archive command).**
    b. searchctl archivedfolders stats --id <folder ID> **(NOTE: replace <folder ID> with ID from step #3, for example only: 3fe1c53bdaa2eedd).**
    c. **searchctl stats --folder xxxx**

2. Learn more about - Monitoring Copy Jobs, View Job History, Viewing folder job History:
    a. searchctl jobs running (**shows running jobs**).
    b. searchctl jobs view --id (id is returned from the command above, **monitors progress**).
    c. searchctl jobs history (shows the **history of previous jobs** including incremental sync jobs, inventory jobs).
        i. [--folderid] - allows filtering the history to a specific folder to list the jobs
        ii. [--type TYPE] - allows filtering to show a specific job type "Incremental", "Full Archive", "Inventory"
        iii. [--output OUTPUT] - Json format output of the history data
        iv. [--tail x] where x is last x jobs executed example --tail 100 (100 last jobs executed)
        v. Example command
            1. searchctl jobs history --folderid <folderid> - this command can very useful to list all the jobs that have run against a single folder. This allows viewing the history of all jobs on a single folder.
            2. searchctl jobs history --tail 100 (return the most recent 100 job executions)
    d. searchctl jobs cancel --id (with a job ID provided **a copy job can be canceled.** Note: it takes time for a cancel to take effect).
3. Show job , all folders and golden copy node Summary statistics

a. This command will provide stats of all jobs executed on all defined folders or all golden copy nodes or summary of all jobs.

b. This provides a global view statistics based on each of the view options. This command provides live updates to the CLI based on active jobs.

c. searchctl jobs summary [-h] [--no-stream] (--nodes | --folders | --jobs)

d. searchctl jobs summary --folders (shows the summary of all jobs stats for all folders)

i. 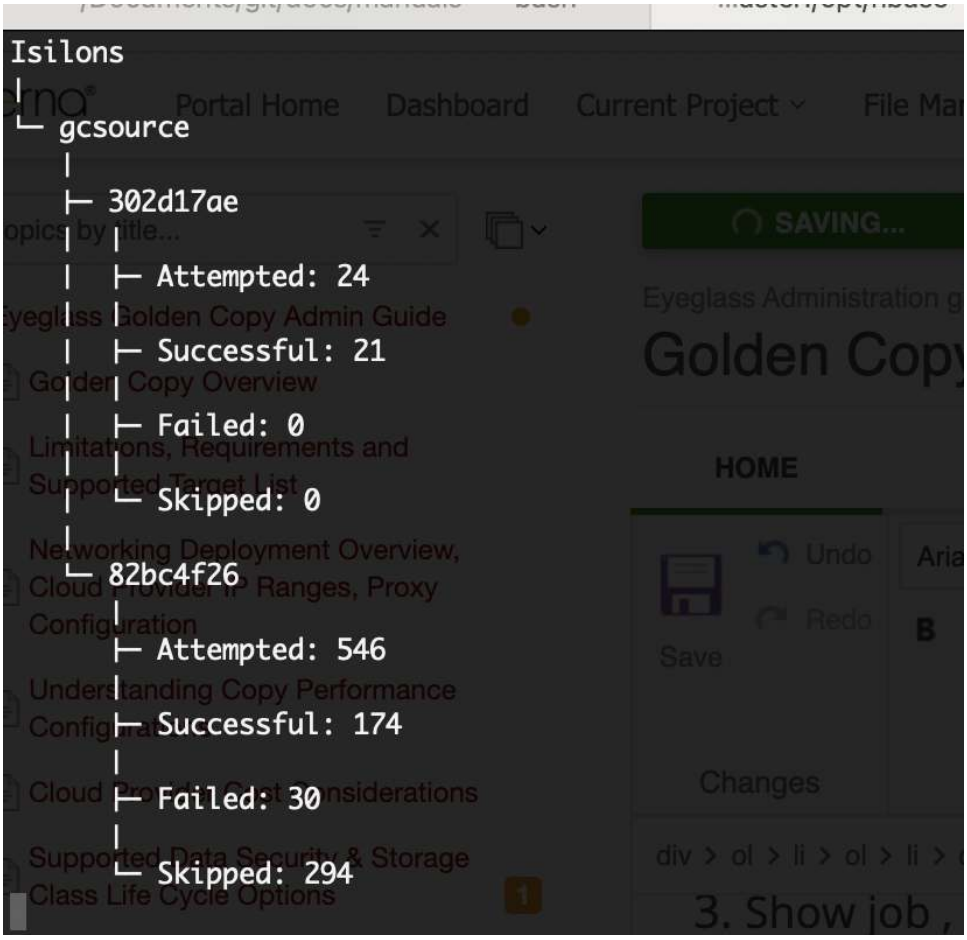

e. searchctl jobs summary --nodes (shows the summary of all jobs stats for all golden copy VM's)

i. 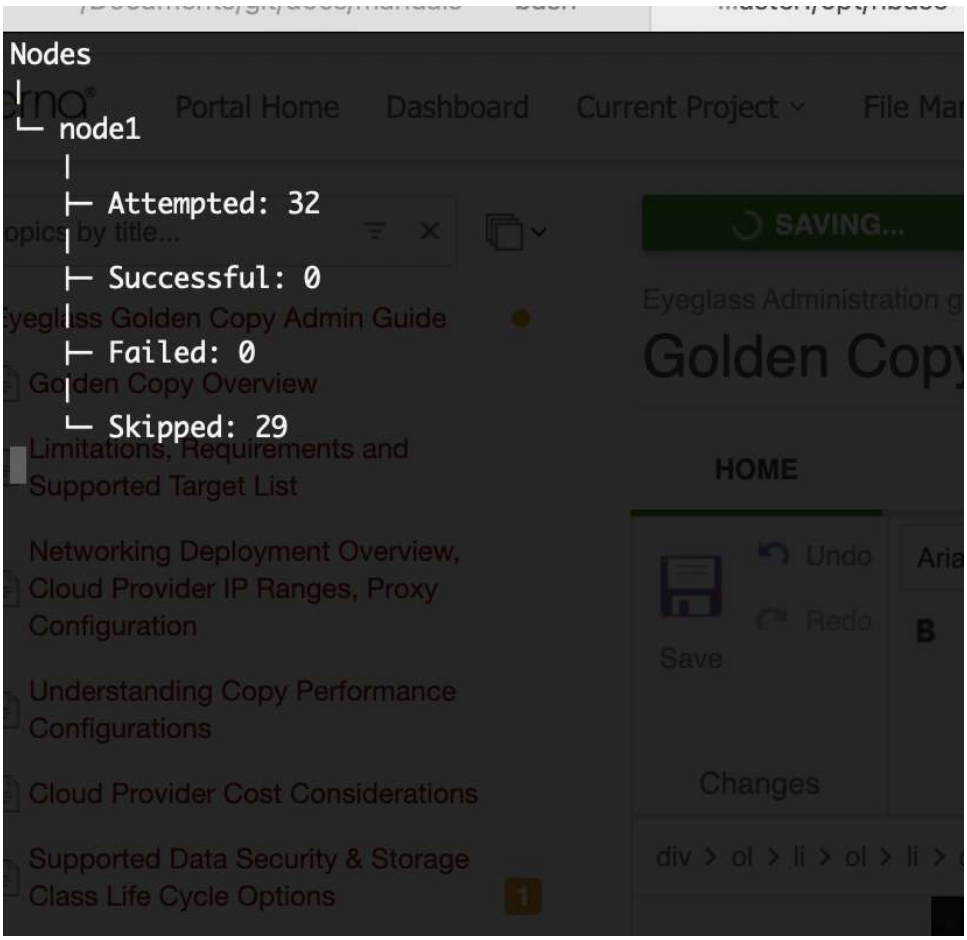

f. searchctl jobs summary --jobs (shows the summary of all job stats for all nodes and all folders

4. Export a completed job into an HTML report with steps here.

How to Enable Auto Progress Emails for copy Jobs

This feature allows administrators to monitor copy job progress from their email. Once enabled, running job progress will be emailed every 24 hours (this can be changed). This avoids the need to login to simply monitor progress. It can also assist with support since the emails can be attached to a support case. The feature allows a group email to be configured.

Prerequisite:

1. Enable a notification channel with SMTP see guide here.
    a. Search & Recover Cluster Operations
    b. Get the email channel groups with this command
        i. searchctl notifications groups list
2. searchctl archivedfolders notifications addgroup --isilon [host] --group [notification-group-name]
    a. Adds a group email to be notified with job progress summary reports
    b. Defaults to disabled status
    c. Enable the notification with:
        i. searchctl archivedfolders notifications modify --isilon <cluster name> --enabled true
3. searchctl archivedfolders notifications list
    a. lists all notifications that are configured
4. searchctl archivedfolders notifications removegroup --group [notification-group-name]
    a. removes the email group
5. searchctl archivedfolders notifications modify --isilon [host] --enabled [true or false] --groups [group-names-comma-separated]
    a. The modify command allows the option to disable the notifications without removing it.
6. The default interval to receive updates is 24 hours. To change this default.
    a. Use the schedules CLI to changes the frequency , this examples shows 5 minutes

b. searchctl schedules modify --id JOBS_SUMMARY --schedule "*/5 * * * *"

c. searchctl schedules list

How to Manage Folders (List , Modify and Remove)

1. List , modify and remove commands and examples

list (list configured folders)

1. searchctl archivedfolders list
2. searchctl archivedfolders list --verbose (adds all flags to the output)

modify (change configuration of an existing folder)

1. --id ID (folder id)
2. [--cloudtype {aws,ecs,azure, gcs, other,blackhole}] (type of target storage)
3. [--region REGION] (Required for AWS)
4. [--tier] specify storage tier (Advanced license or backup bundle required)
5. [--bucket BUCKET] (Required for all storage targets, except Azure)
6. --cluster-name - Allows creating an alias for the root folder name in the storage bucket. Use a string to replace the actual cluster name in the folder used to copy all the data under this root folder name. See the use cases for this feature here.
7. [--container CONTAINER] (Required for Azure only and should list container name).
8. [--endpoint ENDPOINT] (Required for Azure, AWS, ECS).

9. [--secretkey SECRETKEY] (Required for all storage targets).
10. [--accesskey ACCESSKEY] (Required for all storage targets, for ECS this is the user id name, for Azure this is the storage account name) .
11. [--skip-s3-file-exists {true, false}] (Optional: see explanation of this in above) .
12. [--recyclebucket TRASHBUCKET] (See create explanation in the "How to add folders to be copied or synced" section above) .
13. [--meta-prefix xxx] - (See add folder in the "How to add folders to be copied or synced" section for explanation).
14. [--endpoint-ips] (List of ip or range to load balance ECS target only).
15. [--includes INCLUDES] [--excludes EXCLUDES] - (see the "How to add folders to be copied or synced" section for detailed explanation and examples).
16. Scheduled jobs - See job schedule examples here.
    a. [--incremental-schedule INCREMENTALCRON] - incremental sync job
    b. [--full-archive-schedule FULLCRON] - full copy job
    c. [--archive-data-audit-schedule ARCHIVEDATAAUDIT] - data audit source and target comparison

How to remove an archive folder

1. searchctl archivedfolders **remove --id ID (folder id)**
2. [-h] get help

How to Re-run Only Failed Files from a Copy Job

1. This feature allows a retry of only the failed files listed in a copy job, versus running the entire copy job again. This is more efficient and faster process than running the entire job again.
2. First list the job history for an archive folder.
3. **searchctl archivedfolders history**
    a. Then select the job to re-run with a status of failed to get the job ID, then run the command below:
        i. **searchctl archivedfolders rerun --id**xxxxx (where xxxx is the job-xxxxxx)
    b. This will locate all the filed files in the last job and reattempt to copy these files in a new job. This will generate a new summary and progress report.

How to Recall Data from Object Back to File

Overview:

Once data is archived to objects, you may need to recall some or all the data, to the same or different cluster. This section covers how that is done from the CLI. All recalls job recall data to a staging recall path where the administrator would move the tree of data back into the main file system or share the data directly from the recall path.

The recall NFS mount is on /ifs/goldencopy/recall this path is where all recall jobs will recall data. The recalled data will be created as a fully qualified path /ifs/data/datarecalled under the mount path above.

Backup and Restore Use Cases

The backup and restore use case capabilities depends on which licenses are installed. The enhanced recall features require the Backup bundle or upgrade Advanced License key.

License Key Dependancies

1. Base Golden Copy License key allows recall of data only, recall data + metadata, recall to same cluster the data was copied from. It also allows the NFS recall mount to be changed to a different cluster. This is a manual process.
2. Golden Copy Advanced (Backup Bundle) key allows additional options: object version aware recall based on file metadata date range, redirected recall to an alternate cluster other than the source cluster, storage tier aware recall example archive tier recall before recall to on site storage.

Limitations:

1. Golden Copy base license recall to an unlicensed cluster will only recall data, without metadata.

Requirements:

1. 1.1.4 build 2006x Release or later
2. Some recall features require the Advanced license key as explained below. These commands will be identified as requiring the advanced license key.

Prerequisite NFS Export Configuration Steps

1. Make sure the NFS export is created on the clusters that are targets of a recall and and the mount is added to Golden Copy and all VAN nodes.

2. Basic Golden Copy licenses can switch the NFS mount using /etc/fstab entry to point at a different cluster to recall the data.
    a. Create NFS mount on the new target cluster. Review installation requirements here.
    b. Change the mount used for recall
        i. sudo -s (enter ecaadmin password)
        ii. edit the nano /etc/fstab
        iii. find the recall mount entry '<CLUSTER_NFS_FQDN>:/ifs/goldencopy/recall /opt/superna/mnt/recall/<GUID>/<NAME> nfs defaults,nfsvers=3 0 0
        iv. Replace <CLUSTER_NFS_FQDN> with the ip address or DNS smartconnect name of the cluster target
        v. control+x to save the file
        vi. Type exit (to return to ecaadmin session)
        vii. unmount the recall mount
            1. ecactl cluster exec "sudo umount /opt/superna/mnt/recall/<GUID>/<NAME>"
                a. enter the admin password when prompted on each node.
        viii. mount new cluster recall mount
            1. ecactl cluster exec "sudo mount -a"
                a. enter the admin password when prompted on each node.
        ix. The recall job will now recall data to the new cluster target.

3. Advanced or Backup Bundle Licensed appliances can use the add cluster target recall only option and create unique NFS recall mounts for each cluster to allow selecting the target cluster when building the recall job.

Logical Diagram of a Recall

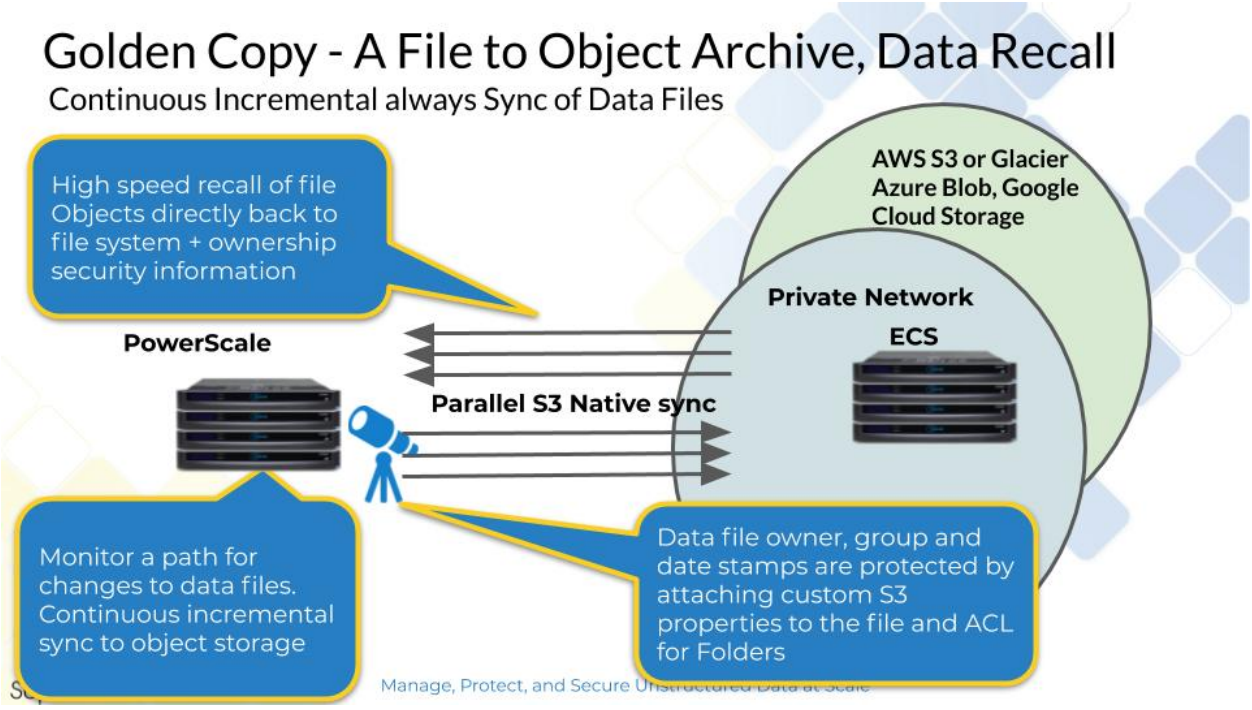

Recall command syntax and options

1. searchctl archivedfolders recall [-h] --id ID[--subdir SUBDIR]
    a. --ID - Is the folder id.
    b. [--subdir] - Allows entering a path below the folder path to select a subset of the data example archive folder is /ifs/archive and --subdir can be /ifs/archive/subdirectory.
    c. [--apply-metadata] If this flag is not used only data will be recalled and the files will not have owner, group or mode bits updated, folders will not have ACL permissions applied.

i. NOTE: Skipping metadata will recall data faster. The recall starts with data and then separately runs a 2nd job to apply metadata, this 2 stage approach allows for fast data recall. The metadata will be applied after the data is recalled completely. The metadata is placed in a queue and can be applied later using a special job that will run on demand. If this flag is omitted you can still run the 2nd job type to apply metadata later.

ii. See How to run a metadata job

d. **(Advanced License)** [--target-cluster] - This is the name of the target cluster that the recall will use to store the recalled data. The cluster must be added to Golden Copy first to use this flag.

e. **(Advanced License)** [--start-time STARTTIME] - This allows selection of data to recall based on the meta data stamps on the objects.

i. **This command is version aware and will check the version of objects to select the best match.**

ii. This is a date format for the beginning of the date range. example

start last modified time. Example: "2020-05-21T17:44:40-04:00"

f. **(Advanced License)** [--end-time ENDTIME] - This is used with the --start-time to specify the end date and time or the date range.

i. **This command is version aware and will check the version of objects to select the best match.**

ii. This uses the same date format. Example: "2020-05-21T17:44:40-04:00"

g. **(Advanced License)** [--timestamps-type {modified,created}] - This flag is used with the date range command to specify which time stamp to use the created time stamp or the modified time stamp stored with the object metadata.

    i. This command is version aware and will check the version of objects to select the best match.

2. Running the job will recall data and re-apply all possible meta data to the file system including ACL's folder and file owner, group and mode bits.

How to recall object data to the Recall Staging Area

1. This example will recall data to the staging areas located here /ifs/goldencopy/recall. There are flags to handle overwriting data that already exists on the destination path and applying metadata to files and folders or reapplying metadata.

    a. Metadata includes the following:

        i. Files - owner, group, mode bits

        ii. Folders - Owner, group ACL on the folder

    b. **Example command to recall only a specific path**

        i. searchctl archivedfolders recall --id xxxx --subdir /ifs/bigfile --apply-metadata

        ii. NOTE: Recall jobs will always overwrite data in the staging area if a previous recall job had already been executed.

    c. **Recall data and apply metadata to files and folders**

        i. searchctl archivedfolders recall --id <folderid> **--apply-metadata**

ii. **Note: Metadata will be applied after all the data has been recalled first**

d. **Recall metadata only or re-apply metadata to a previous recall job**

i. searchctl archivedfolders metadata --jobid <recall Job Id>

ii. **NOTE: This command requires the jobid of a previous recall job that has completed already. Use searchctl jobs history to get previous job ID's.**

How to Monitor Copy Job Performance and Job Summary logs and Error Logs

1. Stats Command for real time job throughput monitoring

a. The following command will monitor real time stats for an archive folder with bytes copied, files copied and error rates.

b. searchctl archivedfolders stats --id xxx (where the xxx is the folder id found from searchctl archivedfolders list)

c. OR use feature rich command **searchctl stats --folder xxx**

d. NOTE: The rate statistics columns are per second but average over the last minute, hour or day.

e. NOTE: Files Retry Pending, Bytes Retry Pending requires the --auto-rerun flag on archive jobs to queue failed files to be retried

at the end of the archive job. This is available in release 1.1.4 builds > 178.

f. NOTE: Rerun stats show files that were retried at the end of the archive job and will only display if --auto-rerun flag was used. This is available in release 1.1.4 builds > 178.

g. NOTE: Accepted stats is related to tree walking the folder that is configured for archiving and is based on REST API retrieval of files and folders.

h. NOTE: Full statistics are recorded during full archive jobs. Incremental stats will appear if incremental jobs are configured on the folder.

i. NOTE: Files and folders have separate statistics

j. NOTE: Metadata is (owner, group, mode bits, created, modified, accessed, ACL)

k.

```
Statistics for folder: 3fc44980bb524894
name                              total_alltime   total_min   total_hr    total_day   rate_min    rate_hr     rate_day
--------------------------------  --------------  ----------  ----------  ----------  ----------  ----------  ----------
FULL/BYTES_ARCHIVED               21,096,828      0           21096828    21096828    0           35164.9     35164.9
FULL/BYTES_ARCHIVE_ERRORED        450,887,680     0           450887680   450887680   0           751556      751556
FULL/BYTES_RETRY_PENDING          450,887,680     0           450887680   450887680   0           751582      751582
FULL/FILES_ACCEPTED               25              0           25          25          0           0.0416757   0.0416757
FULL/FILES_ARCHIVED               18              0           18          18          0           0.0300052   0.0300052
FULL/FILES_ARCHIVED_ERRORED       7               0           7           7           0           0.0116668   0.0116668
FULL/FILES_ARCHIVE_ACCPETED       31              0           31          31          0           0.0516722   0.0516722
FULL/FILES_RETRY_PENDING          7               0           7           7           0           0.0116676   0.0116676
FULL/FOLDERS_ACCEPTED             6               0           6           6           0           0.0100022   0.0100022
FULL/FOLDERS_ARCHIVED_METADATA    6               0           6           6           0           0.0100015   0.0100015
FULL/FOLDERS_ARCHIVE_ACCEPTED     6               0           6           6           0           0.0100002   0.0100002
RERUN/BYTES_ARCHIVE_ERRORED       450,887,680     0           450887680   450887680   0           751639      751639
RERUN/FILES_ARCHIVED_ERRORED      7               0           7           7           0           0.0116687   0.0116687
RERUN/FILES_ARCHIVE_RERUN_ACCEPTED 7              0           7           7           0           0.0116668   0.0116668
```

2. Job View command monitors progress in MB and file count with % completion

a. This command can check the progress of a running job, shows the MB of files queued for copy, the MB of files archived, and % completed. The same information is shown for file count queued, archived and %. Error rate is also shown

b. searchctl jobs view --id job-1591744261081-1928956557 --follow (use searchctl jobs running to get the job id)

c. New in builds > 1.1.4 178 build is the **Retry Pending** Stat that tracks files that had an error will be queued to be retried at the end of the archive job. This new stat is only active if the --auto-rerun flag is used on the archive job.

d.

```
ecaadmin@gc113-1:~> searchctl jobs view --id job-1602171295308-221348677 --follow

Folder ID: 3fc44980bb524894

Upload: 450.12MB accepted, 20.12MB archived, 0B skipped ( 4.47% complete )
Count: 31 accepted, 24 archived, 0 skipped ( 77.42% complete )
Errors: 31 attempted, 7 errored ( 22.58% error rate )
Retry Pending: 7

GoldenCopy Full Archive
----Take snapshot of /ifs/data/folder1 ( SUCCESS : 1.17 seconds )
----Create Distributed Kafka topic for REPORT for /ifs/data/folder1 ( SUCCESS : 0.19 seconds )
----Create Distributed Kafka topic for CONTENT for /ifs/data/folder1 ( SUCCESS : 0.05 seconds )
----Trigger before/after job scripts ( SUCCESS : 0.25 seconds )
----File System walk (GoldenCopy) distributedly ( SUCCESS : 0.03 seconds )
----Wait for job on archive-bfs-walk-3fc44980bb524894 ( SUCCESS : 2 minutes, 40.04 seconds )
----Trigger before/after job scripts ( SUCCESS : 0.35 seconds )
----Collect settings ( SUCCESS : 0.36 seconds )
----Pass the failed jobId to data ( SUCCESS : 0.00 seconds )
----Create rerun for Archive-1602171295308 (   Running ..    )
        ----Create Kafka topic for /ifs/data/folder1 ( SUCCESS : 0.01 seconds )
        ----Create Kafka topic for /ifs/data/folder1 ( SUCCESS : 0.01 seconds )
        ----Get failed archive events ( SUCCESS : 1.81 seconds )
        ----Wait for job on archivecontent-3fc44980bb524894 (   Running ..    )
```

How to Monitor Ethernet Interface Mbps during a Copy Job

1. This tool will be available in new OVF builds. In the current product, it must be installed and requires Internet access to repositories
2. ssh to the Golden Copy vm as ecaadmin
3. sudo -s (enter ecaadmin password)
4. zypper install nload (answer yes to install)
5. exit
6. type **nload**
7. Once the UI loads use the right arrow key until you see Eth0 displayed at the top left of the UI. This will now display TX and RX bandwidth current, average, min and max values with a graph showing relative bandwidth usage.

8. 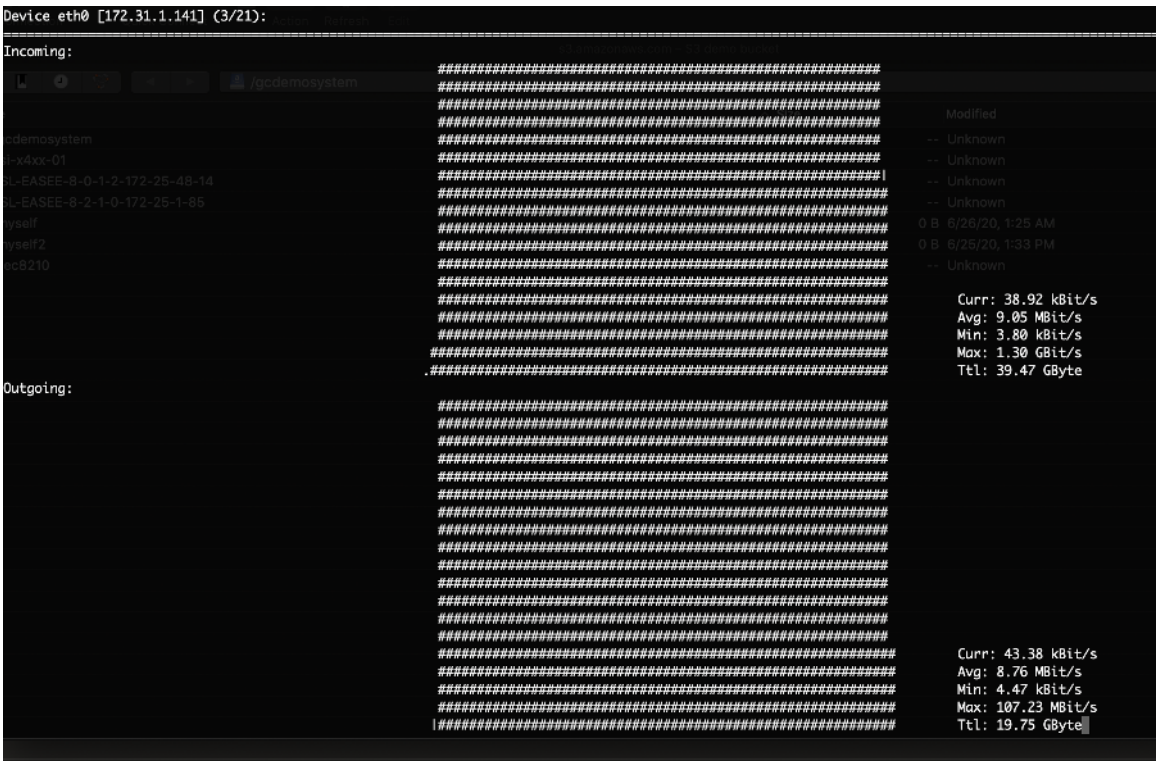

How to Generate the HTML Report After a Copy Job Completes

As of 1.1.4 or later it is required to run the export report command to produce the HTML report for a copy job. The recommended method to monitor a copy is using the CLI monitor command. **NOTE: This command will only execute when copy job is completed**.

1. Get the job id "**searchctl job history"** (Use the job id for the next command) .
2. searchctl archivedfolders export --jobid <job-id> --errors-only true
    a. Recommended to use the --errors-only flag to extract only errors and reasons for the error
    b. NOTE: The json files created can be very large since they will include each file copied, errored or skipped.

3. The report is queued to extract the report data that is stored in the archive report Web folder. Follow the steps below to view the exported HTML report.
4. Login to the report page with https://x.x.x.x/downloads/archivereport/<folder name>/Full/ .
5. Locate the <folder name>-<date and time>-summary-checkpoint.html and click on this file to view it.
6. 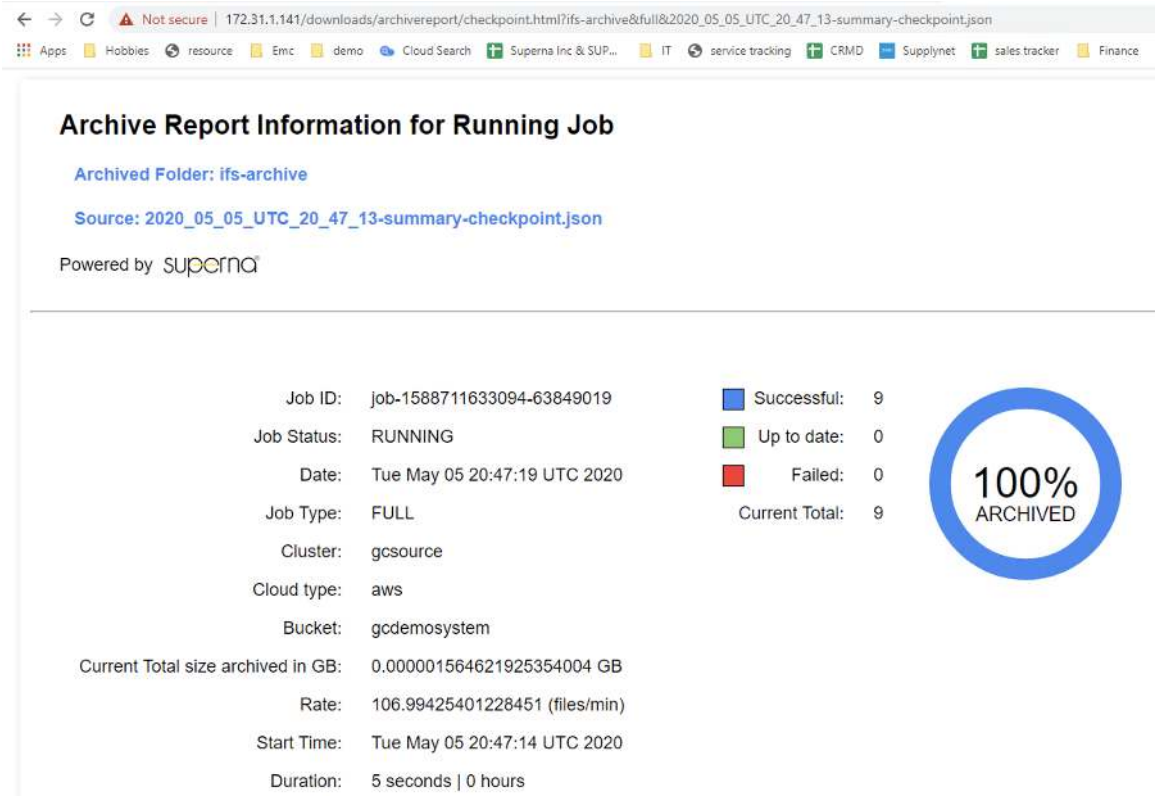

1. How to view the Detailed json Copy Job Logs
   a. The HTML export command also generates detailed json logs that store each file that was copied.
      i. searchctl archivedfolders export --jobid <job-id>
   b. This will start a job to generate the json files. These files can be very large and may take many minutes to create.

c. The jobs logs can be accessed from https://x.x.x.x/downloads/archivereport (where x.x.x.x is the ip address of the Golden Copy VM).

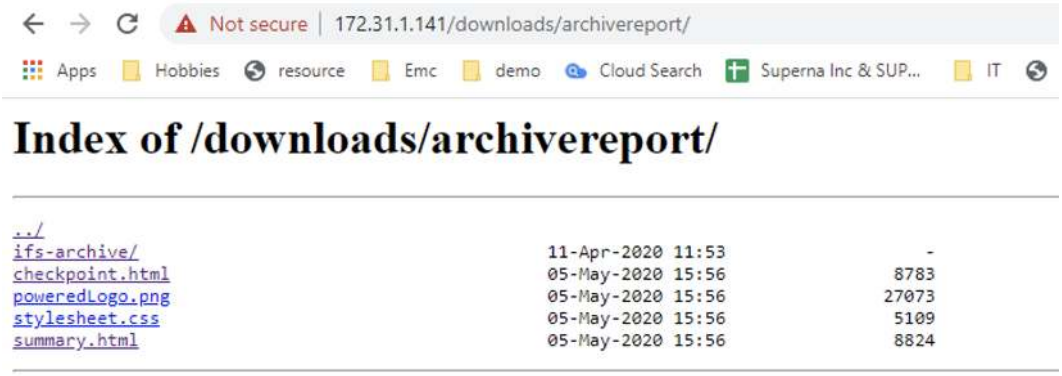

d.

e. The default login is: "ecaadmin", with password: "3y3gl4ss".

f. The log directory for each execution of a job report export command. It will create a folder based on the path that was archived example ifs-archive:

   i. The job logs will be contained in the full and incremental folders sorted by date.

      1. 3 files are created for each copy job. See example screen shot below.
      2. date-time.JSON (full log with all files that successfully copied, many numbered versions of this file will exist).
      3. -summary.html (report view that shows the report for the entire job).
      4. -summary.json (JSON version of the summary report used by the HTML file to display).
      5. -errors.json (if there are failed copy files this file extension will appear and stores all the files that failed to copy)

ii. 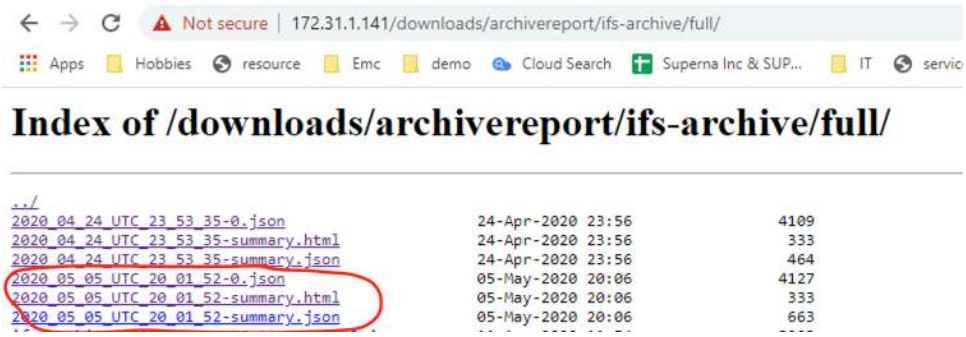

2. How to view copy job errors on failed copies

a. Use this command to find the reason for the failed copies. The job id can be found with **searchctl jobs history. NOTE: Requires 1.1.4 or later.**

b. searchctl archivedfolders errors <--Id JOBID> [--head | --tail | --at TIME] [--count N], where:

   i. JOBID is the ID of the job that was run

      --count N prints N records (default 10)

      --head (default) starts printing from the earliest detected error

      --tail prints up to the last detected error in the job

      --at TIME will print errors starting from the given time. Use the same time format as T15023

      --head, --tail, and --at are mutually exclusive.

   ii. Example command to quickly find that last 20 reasons files failed to copy.

      1. searchctl archivedfolders errors --Id xxxx --tail --count 20

How to Manage File Copy Performance

There are 3 methods that can be used to control performance of a copy job.

1. The first is the number of virtual accelerator nodes deployed to distribute copy tasks to more virtual machines. See the installation guide on how to deploy virtual accelerator nodes. The guide is here.
2. The 2nd parameter that controls how many files will be copied concurrently per Golden Copy VM or per Accelerator node.
3. The 3rd option applies to large file copies (greater than 10 MB), and allows increasing the number of threads that copy byte ranges of the file. A large value will increase bandwidth requirements.

Best Practice: Test a copy job and monitor performance statistics command to monitor files per second counter (searchctl archivedfolders stats --id <folder id>. This command shows bytes per second and files per second over the last minute. The higher the value indicates higher performance.

How to Increase Copy Performance with concurrent file and large file thread count

Golden Copy VM defaults to the values shown below for concurrent file copies and threads per large file. This value applies to each VM Golden Copy and Virtual Accelerator nodes. This can be changed globally, following the steps below:

1. ssh to Golden Copy VM as ecaadmin .
2. nano /opt/superna/eca/eca-env-common.conf
3. Add the line to this file and set the number of files to desired concurrent copy limit and the parallel thread count for large file copies. **NOTE: Consult with support , increasing these numbers may not increase performance in some scenarios.**

   a. export ARCHIVE_PARALLEL_THREAD_COUNT=100 (number of concurrent files per Golden Copy VM or Accelerator Nodes, increasing this number may not increase performance unless sufficient bandwidth is available)
   b. export ARCHIVE_PARALLEL_THREAD_SDK=10 (Number of separate threads used to copy a single large file, higher number will increase bandwidth utilization)
4. control key + x key to exit .
5. Answer yes to save the file .
6. ecactl cluster push-config
7. ecactl cluster services restart --container archiveworker --all
8. done

How to Shape Bandwidth of Archive Jobs

Overview:

1. The feature is using Traffic shaping and not rate limiting. Traffic shaping (also known as packet shaping) is bandwidth management technique that delays the flow of certain types of network packets in order to ensure network performance. In the case of Golden Copy the file copies are delayed to ensure that over time bandwidth would average out to the desired configuration.
   a. NOTE: Monitoring the interface will show the network usage will be above the set shaping value which is expected with traffic shaping. This is because the interface is 10Gbps and allows the data to leave the VM at a high rate for short bursts.

2. NOTE: Monitoring the interface will show the network usage will be above the set shaping value which is expected with traffic shaping. This is because the interface is 10Gbps and allows the data to leave the VM at a high rate for short bursts.
3. In Release Build 1.1.4 21002
    a. nano /opt/superna/eca/eca-env-common.conf
    b. export ARCHIVE_NETWORK_RATE_MB=xx (xx is MB per second value to all copy bandwidth)
    c. control+x answer yes to save
    d. ecactl cluster push-config
    e. ecactl cluster services restart archiveworker
    f. done

How to Monitor Network bandwidth Usage from the appliance

To monitor utilization of the Ethernet interface in real - time.

1. ssh to the appliance as ecaadmin
2. sudo -s (enter admin password)
3. zypper in nload (requires internet access to the appliance to install the package)
    a. answer yes to install
4. nload (this will display tx and rx bandwidth)

Because Golden Copy is logged in via ssh any snapshot can be used as a source path to copy data Therefore, copying a snapshot is no different than any other path on the PowerScale

NOTE: This will not support continuous sync mode since that depends on snapshot change tracking. This mode is a copy of a snapshot for a long term archive.

1. **Prerequisites**
    a. Add an archived folder with S3 target configuration. See below on path considerations when copying a snapshot.
    b. If the snapshot is higher up the file system than the folder archive path (i.e. a snapshot on /ifs/data and an archived folder on /ifs/data/toarchive) than the archived folder, then the archived folders' base path will be used for the file system walk of files to copy.
    c. If the snapshot is lower down the file system (i.e. a snapshot on /ifs/data/toarchive/somesubdir and an archived folder of `/ifs/data/toarchive'), then the snapshot path will be used as the root for the file system walk to copy files.
2. Command to specify a snapshot path to be copied:
    a. searchctl archivedfolders archive --id xxxx --snapshot <snapshotName> (where xxxx is the id of the archive folder ID and snapshot name is the name of the snapshot in the Onefs Gui).
3. Done.

How to Configure Delayed Deletes with Sync Mode

1. This feature allows protecting deleted files that need to be retained for a period of time. This feature uses a 2nd storage bucket to hold deleted files, and uses the time to live feature on a storage bucket to auto delete files created in the bucket after x days. This storage bucket life cycle policy is configured manually on the storage bucket following S3 target documentation.
    a. This provides a recovery location for deleted data in sync mode and allows a retention period in days to allow recovery using the life cycle management feature of S3 storage buckets.

2. Overview

a. 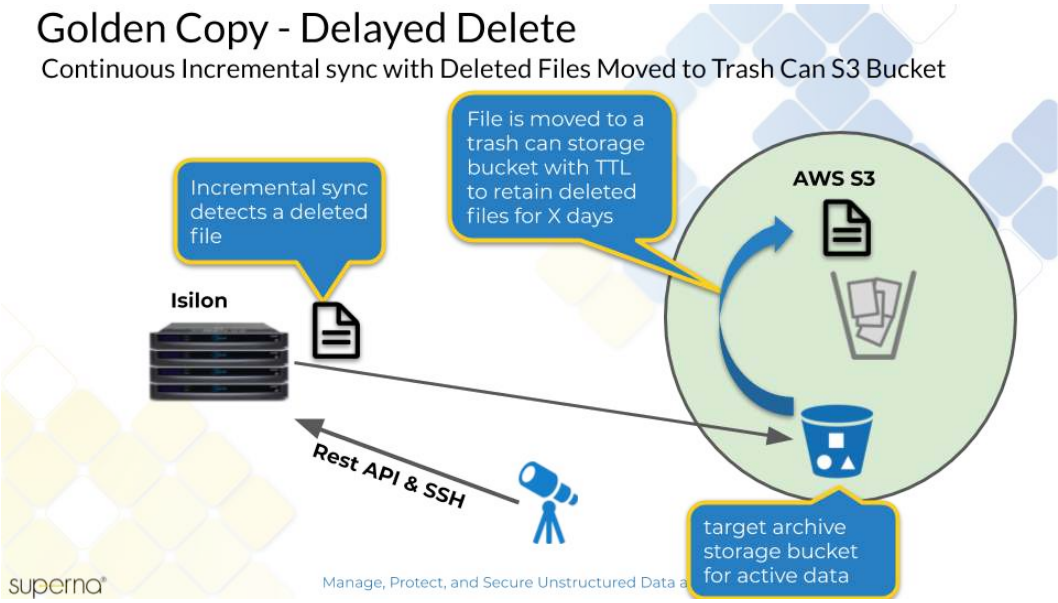

3. Requirements:

    a. The S3 storage provider must be the same for the target storage bucket and the trash can storage bucket
    b. The TTL expiry policy must be created on the trash can storage bucket using the S3 target device documentation. This value is generally set in days to retain a file before it is deleted automatically.

4. How to Configure Delayed Delete Mode

a. Add a folder and use the --recyclebucket option to specify the name of the S3 storage bucket that will act as the trash can for deleted files

b. Example only: searchctl archivedfolders add --isilon gcsource --folder /ifs/gcdeletetest --accesskey xxxxxxxxx --secretkey yyyyyyyyyy --endpoint s3.regionname.amazonaws.com --region region --bucket targetbucketname --cloudtype aws --recyclebucket name-of-trashcan-bucket

c. See Screen shots of a storage bucket source and target after a file is deleted on the source cluster path. NOTE: The incremental job runs to detect file changes to the path including deletes. When a file is detected as deleted and --recyclebucket option is used the file is moved to the trash can bucket and then deleted from the target S3 storage bucket configured on the archive path.

d.

i. 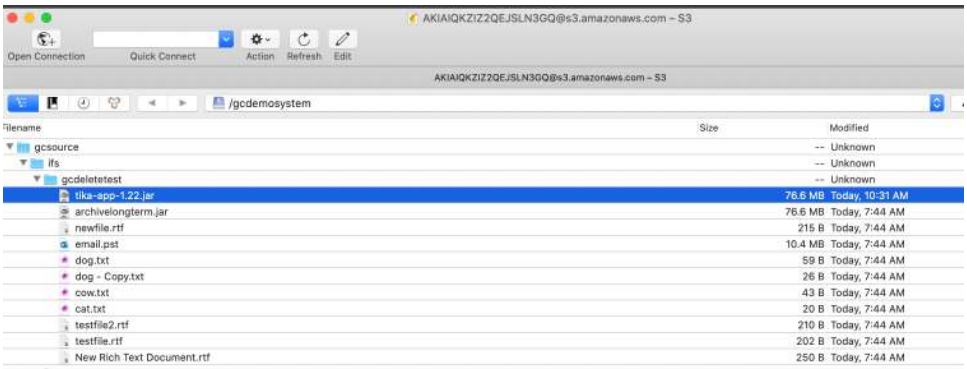

e. File is deleted from source path:

f. 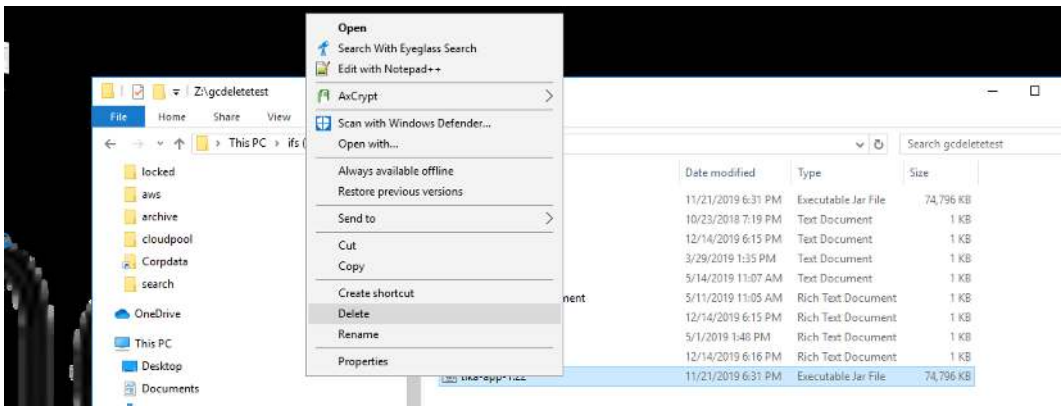

g. File in the trash can bucket showing the expiry property is set:

h. 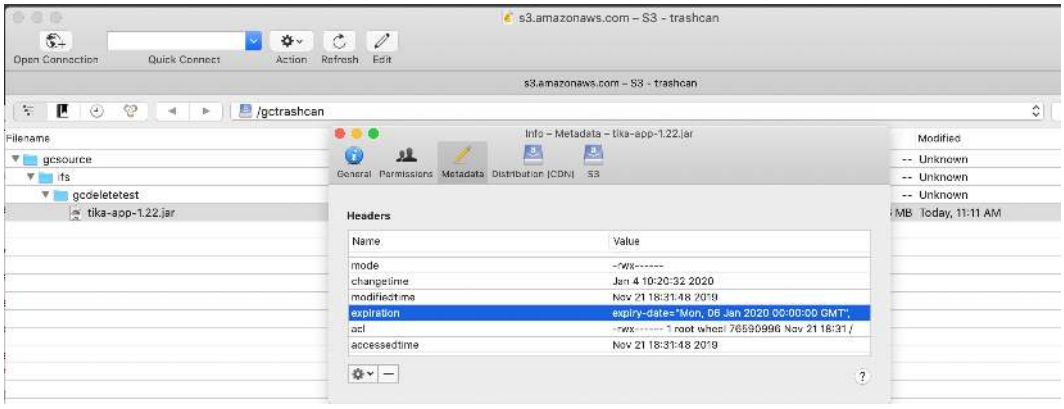

i. Done

How to list and change System Scheduled Tasks

Job Definitions

1. **INVENTORY -** This collects networking , user information from the cluster
2. **PERSIST_JOBS_HISTORY -** this job syncs job status and history to on disk backup copy
3. **UPDATE_SNAPSHOTS** - Used for Search & Recover 1.1.5 or later releases to make sure content indexing snapshot alias always points to a current snapshot and avoids content ingestion errors due to expired snapshots. Defaults to daily.

4. **SYNC_ARCHIVEDFOLDER_SCHEDULES -** This job polls the the Golden Copy configuration to update the master scheduling container. Consult with support.

1. To list all system scheduled tasks
    a. searchctl schedules list

How to Configure a Folder Alias to Handle: Cluster Decommission Use Case, Source Cluster Name Change, Switch to DR cluster as data source and Data full Copy

1. **Use Cases:** A Cluster is decommissioned, a clusters name is changed, switching to the DR cluster to copy data, or creating a new full copy data under a new the folder.
    a. Each of these use cases leaves the data in the bucket under the old cluster name folder at the root of the storage bucket. This solution allows creating an alias to override the folder name used to store data in the bucket . The objective is to allow a new cluster to use the older folder name for archive and recall jobs to gain access to data previously copied from another cluster.
        i. See screenshot below shows data is stored under a cluster named **gcsource**

ii. 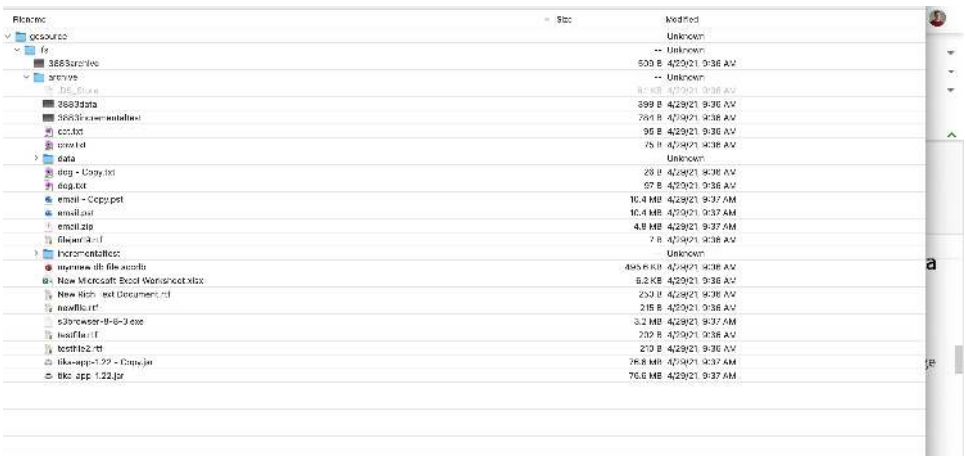

b. The solution is to copy data or recall data under the previous cluster folder name of "gcsource" but using data stored on a new cluster . An alias is created on the folder definition to copy data into the same tree structure in the storage bucket when the cluster name and the folder name in the bucket do not match.

   i. **Example 1**: old cluster name "gcsource" and new cluster is "gctarget". Use this command to change the folder name used to copy or recall data on a folder definition.

      1. A new folder definition is created to connect to the new cluster but adds an alias that references the old cluster name.

         a. searchctl archivedfolders add --isilon **gctarget** --folder path yyy (add authentication flags, bucket and endpoint flags) --cluster-name **gcsource**

         b. This folder definition is connecting to cluster **gctarget** but using a folder alias of **gcsource**, this will copy data into the storage bucket under the folder **gcsource** where the old data is stored.

ii. Example 2: Create a new full copy of data under a new folder alias.

1. In this example a folder definition is modified to add an alias for the cluster name, this will cause data to be copied under a new folder at the base of storage bucket.
2. searchctl archivedfolders modify --id xxxx --cluster-name **newcopy**
3. In this example an existing folder is modified to copy data under a new folder root named **newcopy**. The original data is left under the original folder named after the cluster. All copy and archive jobs will now use the new folder name.
4. NOTE: The old data will remain under the old folder name and will not be used for copy or recall jobs.

Storage Target Configuration Examples

Dell ECS Bucket Creation Walk Through

1. Login to ECS and open the Manage and then users tab.
2. Edit object_user1 .
3. Generate a new secret key :

a. 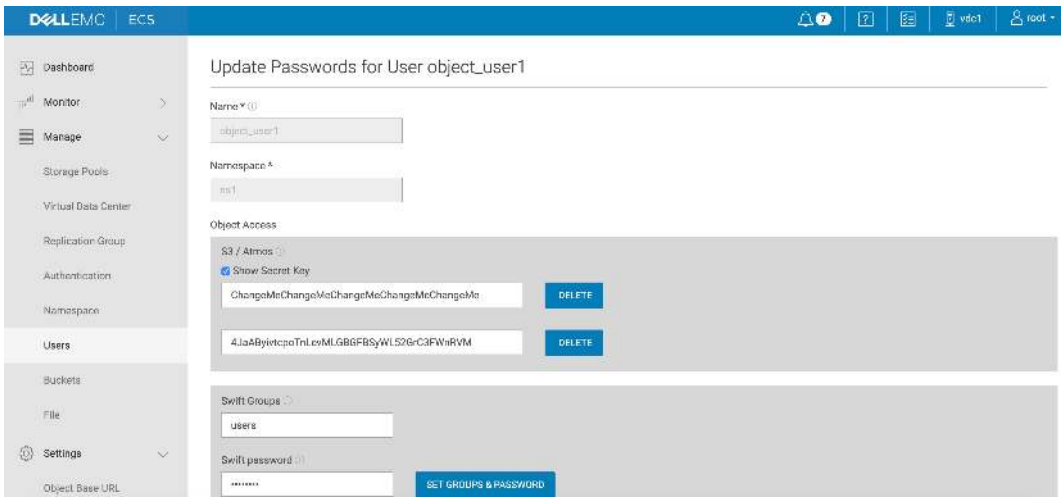

4. Now click on the buckets menu tab.

5. Create a new bucket and enter a name along with the the owner set to the object_user1 user. Click next .

a. 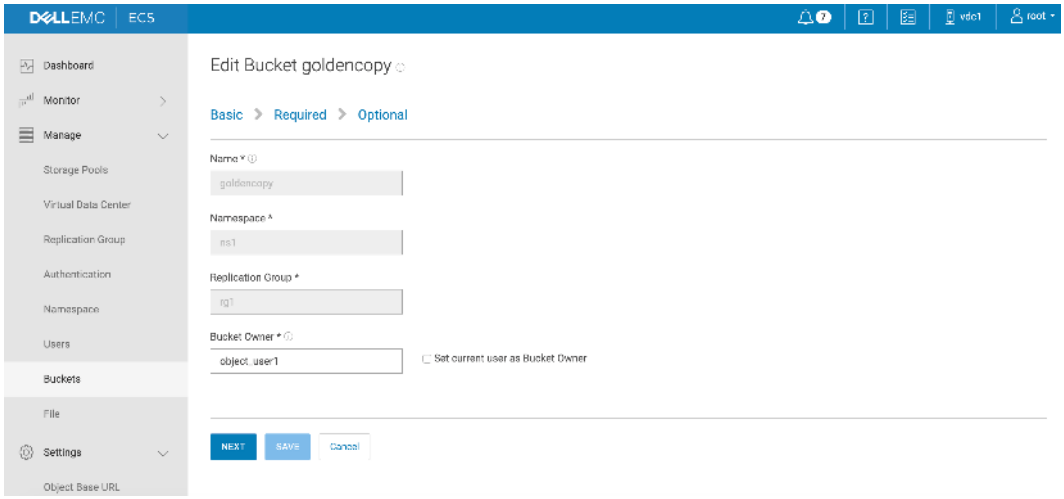

6. Enable meta data search and enter meta data tags that will be indexed.

   a. Click add, select user, and enter tags as per screen shot below with type set to string:

b. 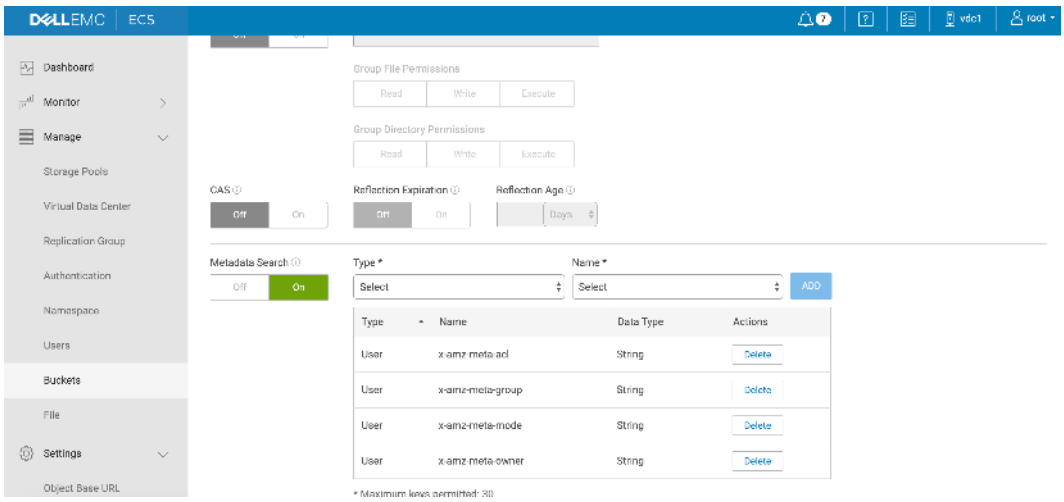

7. Done .

Amazon AWS Bucket Creation Walk Through

1. **Complete Steps to the storage bucket**

   a. Login to Amazon web console .

   b. Select S3 service .

   c. 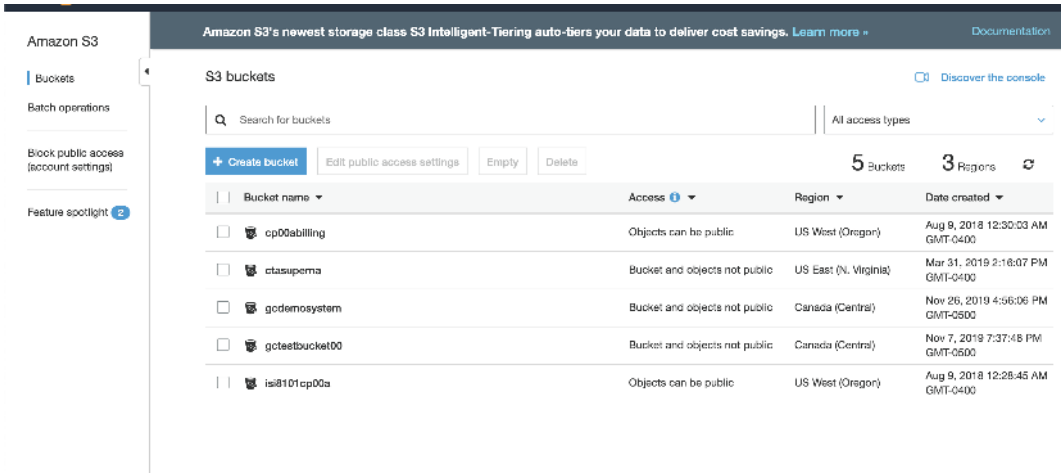

   d. Click Create bucket.

   e. Enter the bucket name, select a region for your bucket, and click create.

f. 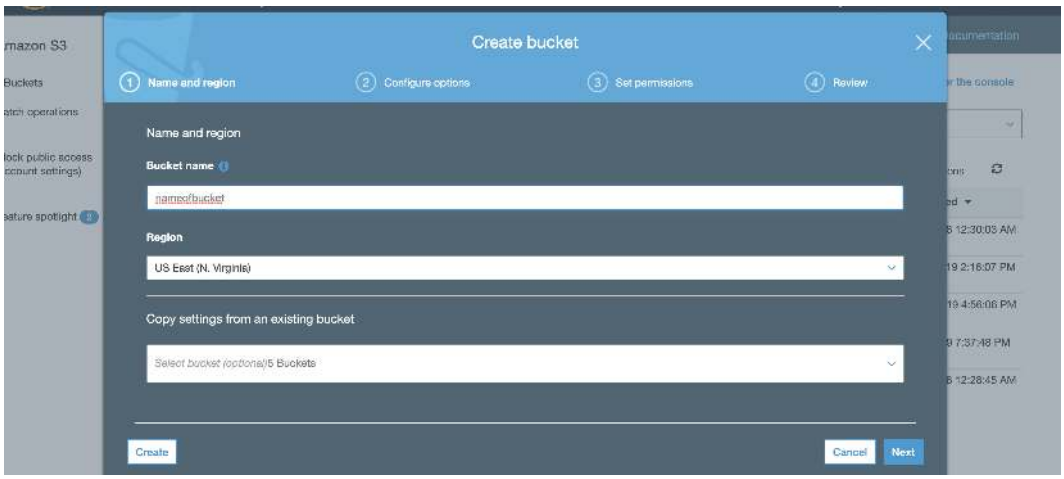

How to setup minimum permissions in AWS for Golden Copy

S3 Policy Permissions Requirements

1. **S3 permissions lists the following policy scope options, indicating optional and mandatory resource scope:**
    a. **<u>Access Point</u> (Optional) - This is used with AWS SDK, CLI or REST API. Access point create and usage is documented here https://docs.aws.amazon.com/AmazonS3/latest/dev/using-access-points.html If you use Access points you need to specify the access point ARN in the policy, and assign the access point to the bucket (see AWS documentation).**
        i. **Example: Access point url has the following syntax**
            1. **Access point ARNs use the format arn:aws:s3:*region*:*account-id*:accesspoint/*resource*.**
            2. **Example**: AWS URL that must be used when adding an archive folder s3-**accesspointname**.*Region*.amazonaws.com
        ii. **NOTE: Sample policy sets the Access point to * to all**

b. Bucket (Mandatory) - The sample policy file includes a sample bucketname that must be replaced with your bucket name. This resource scope is mandatory in a policy.

c. Jobs (Optional) - The jobs resource is not used by Golden Copy, is not required and the sample policy file sets this to * to all. Jobs is used to automate tasks against S3 buckets.

d. Objects (Mandatory) - Golden Copy requires access to all objects with permissions set in the sample policy file. No restricted access to objects should be applied and this is unsupported to block access to objects in a storage bucket dedicated for Golden copy. (AWS documentation)

    i. Sample policy sets this resource scope to * to allow access to all objects.

Quick Start Method

NOTE: We recommend following the guide to learn how to create permissions. AWS supports JSON format policies. The sample policy can be downloaded here and pasted into the JSON tab of the AWS console.

NOTE: You must edit this file and replace the sample storage bucket name gcdemosystem with the name of the storage bucket. Find this string in the file and change the bucket name for your environment "arn:aws:s3:::gcdemosystem"

1. Quick start Steps:

    a. Login to AWS console

    b. Goto IAM

c. Click Policies left side menu
d. Click Create policy
e. Click JSON tab
f. Copy and paste the fixed JSON example file into the policy create
g. Click Review Policy
h. Give the policy a name example Goldencopy
i. Click Create Policy
j. Now click Users on left menu
k. Click add user add name goldencopy
l. Click check box for Access type Programmatic access
m. Click Next for permissions
n. click Attach existing policies directly, search for the policy name you created above ex goldencopy. Select the check box to assign.
o. Click through the rest of the options to create the user and record the access key and secret key needed to add archive folders.
p. Done

Complete Steps to Create User and Policy Following All Steps (skip if you used quick start above)

1. Open the IAMS User screen in AWS:

2. 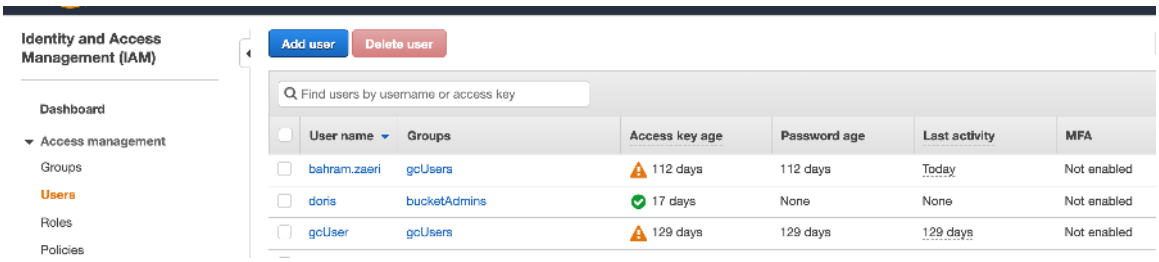

3. Create new user with User name "goldencopy" and select the programmatic check box.

   a. 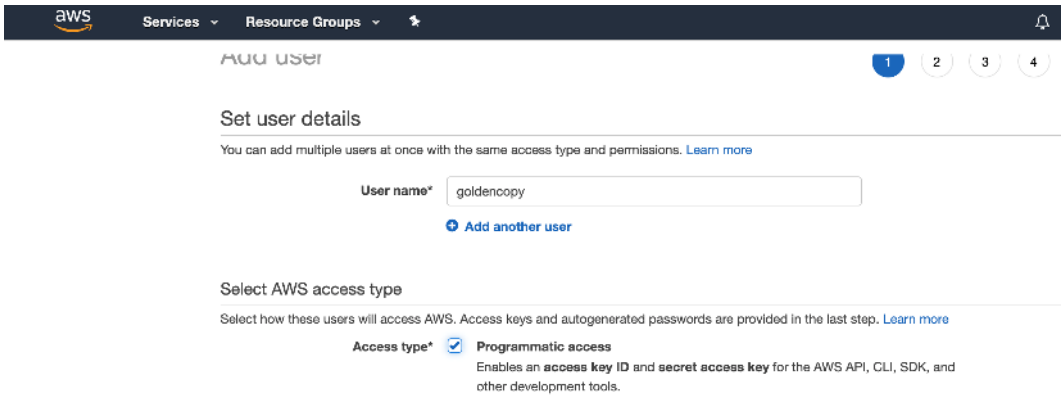

4. Click Next.

5. Change to "Attach existing policies directly" option.

6. 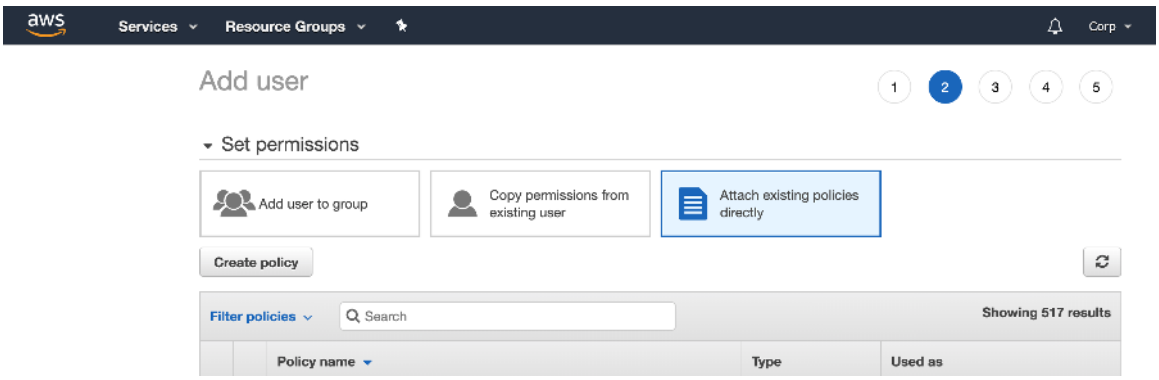

7. Click the Create Policy button.

   a. Click on "Service" and type "S3" and select this service

b. 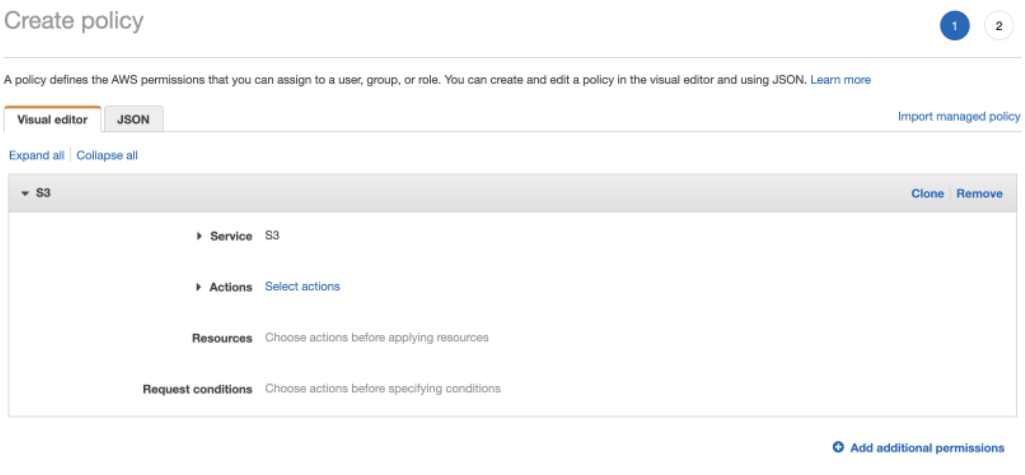

8. Now Click "Select Actions".

   a. Select the permissions as per the screen shot.

9. 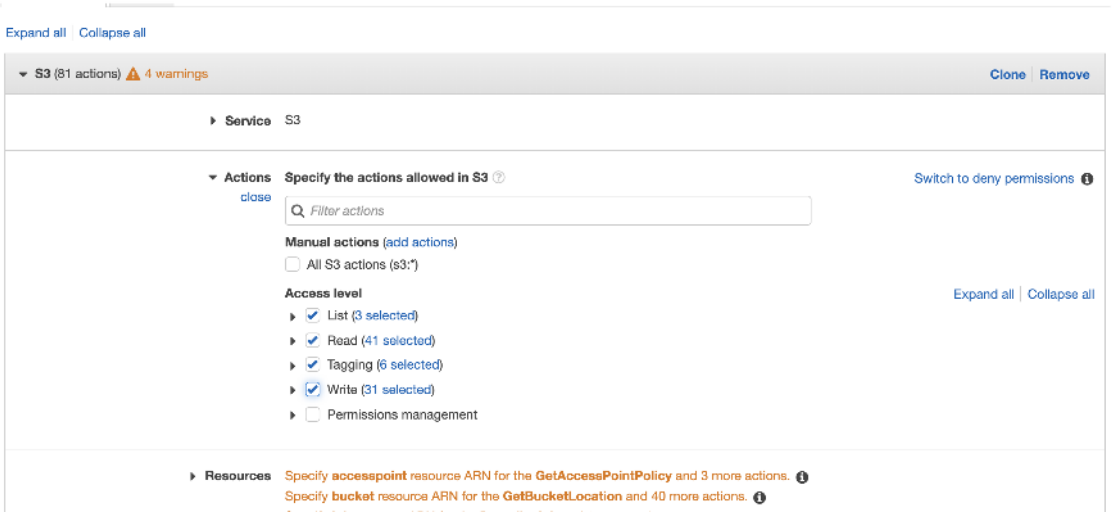

10. Now specify the bucket(s) created to store golden copy target data be adding each **ARN** to the policy. **Click on "Add Arn" next to "bucket" .**

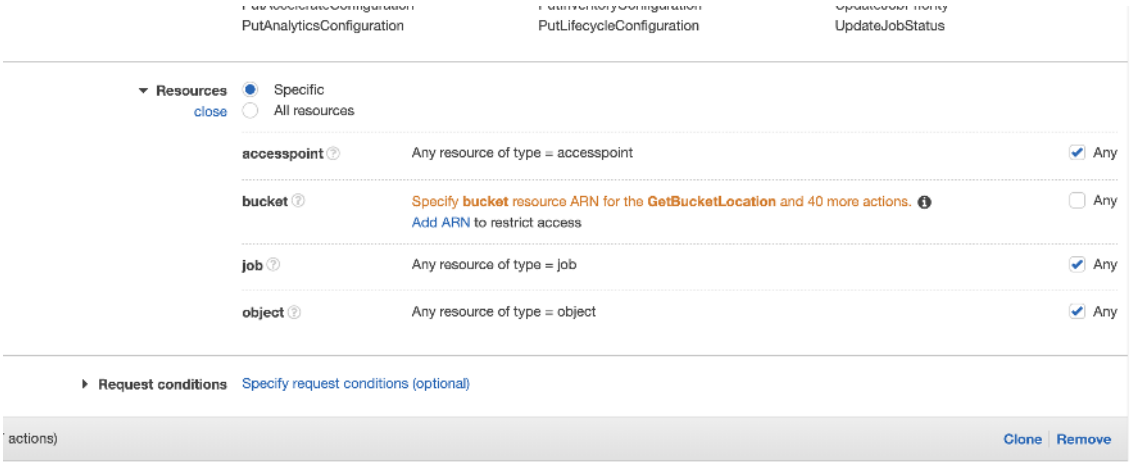

11.

12. The Add ARN window appears, enter the storage "Bucket name" created for Golden Copy (i.e. "gcdemosystem" as shown in the sample screenshot below).

13. 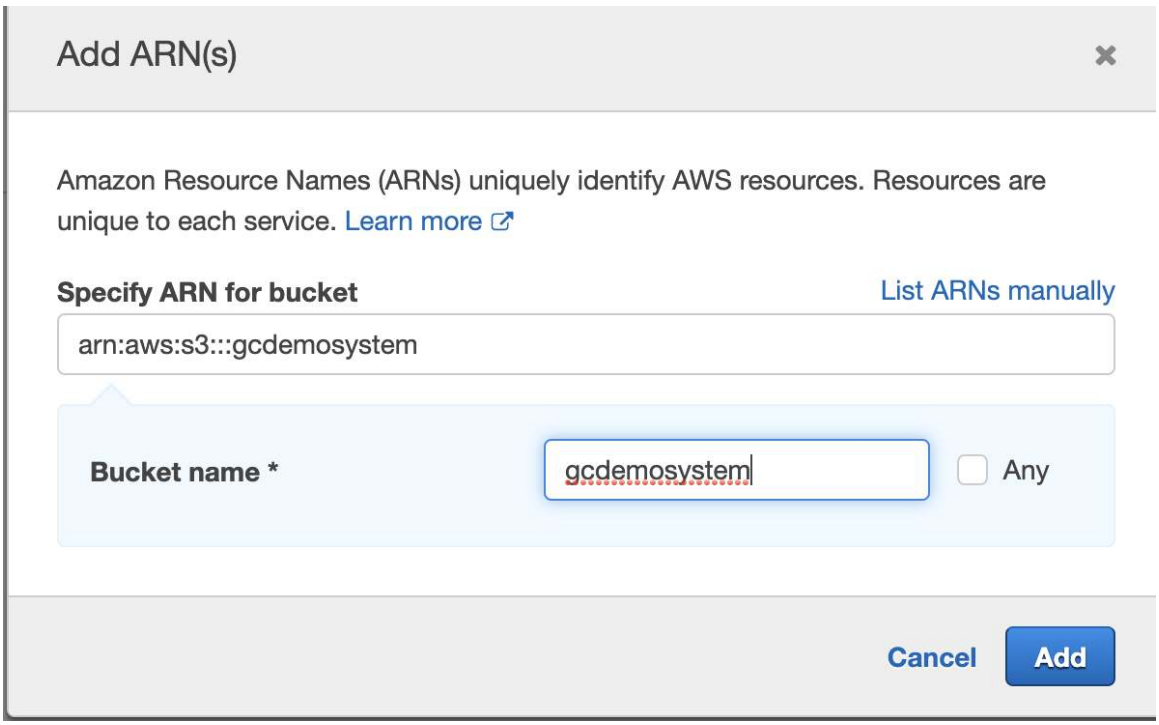

14. Optional - **Source IP Restriction** to the Bucket by Selecting the **"Request conditions"** option.

    a. **Enter a specific Public IP address or a range of addresses. This would be the public facing IP address used for any Internet access from the Data Center where the PowerScale is located. Example s to a specific IP address but a range can be added.**

b. 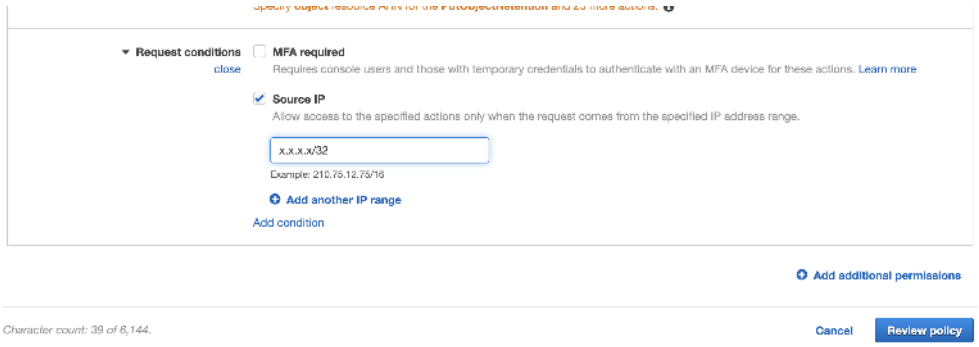

15. Now Click the "Review Policy" button bottom right of the UI (shown in the screenshot above).

16. Enter the Name of the policy, Description, and click "Create Policy".

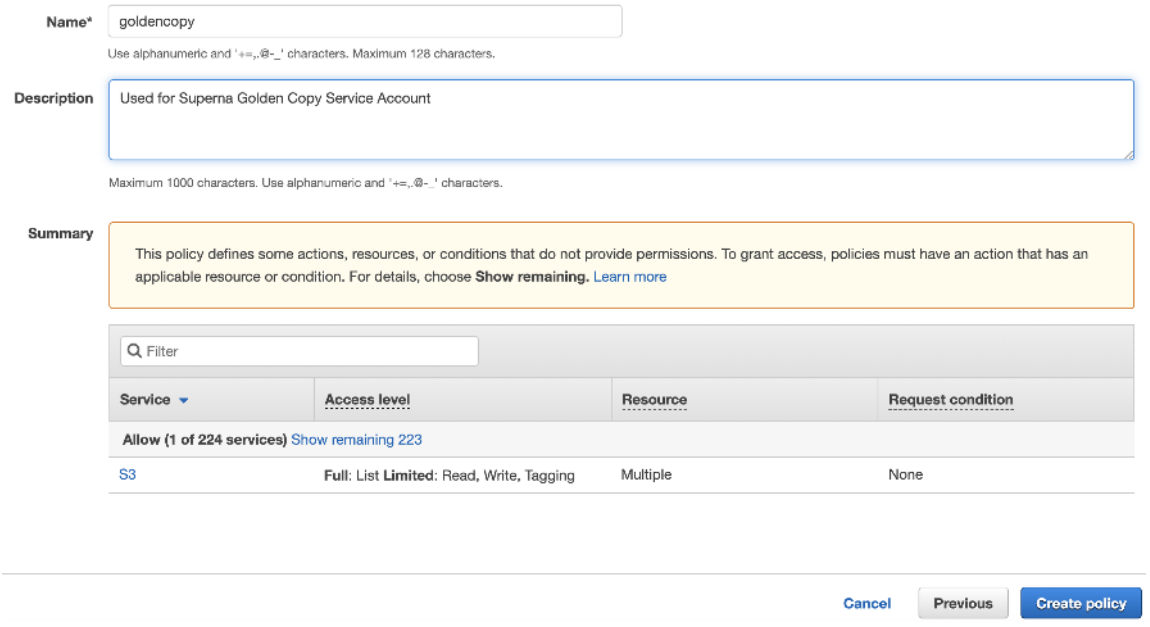

17.

18. Now return to the IAM Create User browser tab and click the Refresh Icon to reload new policies. Type "Goldencopy" into the "Filter policies" dialog box and select the Goldencopy policy you created above.

19. 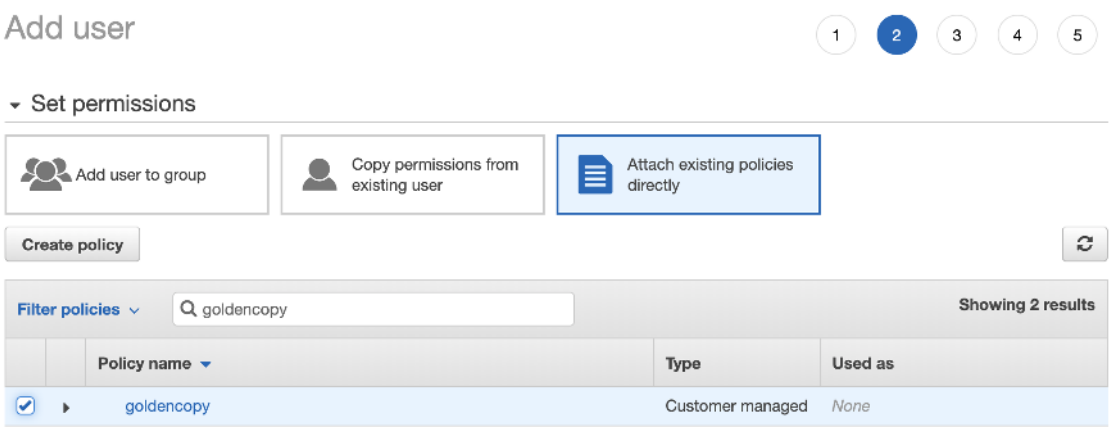

20. Click Next, and then click Next on the tags screen .

21. Click the "Create User" button .

22. On the final screen you need to record the Access ID and the secret key for the Goldencopy user. **Record this in a secure location for use when adding archive paths to Golden Copy.**

23. 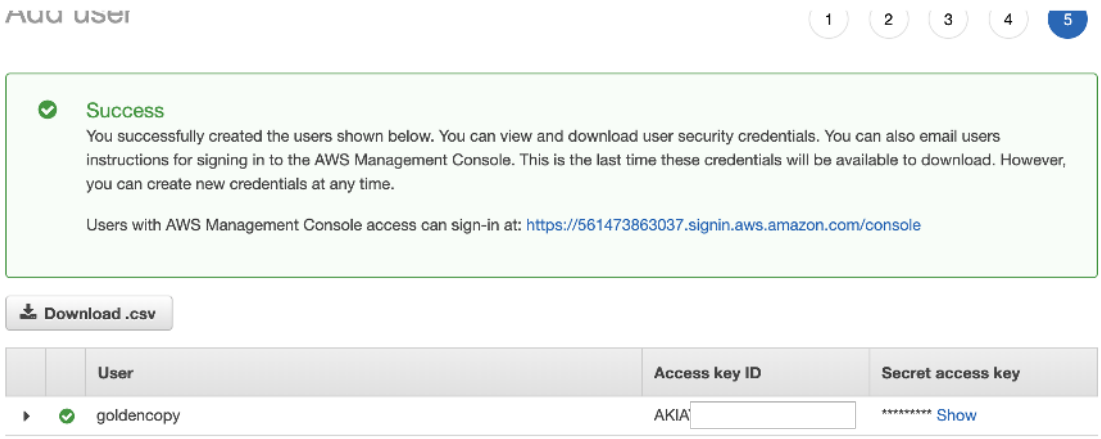

24. **NOTE: You will need your bucket name, region, access key and secret key to configure Amazon S3 target.**

25. Done.

How To restrict S3 Bucket Access by Source Public IP of your Data Center

1. NOTE The role in IAMS can create a single policy that restricts access to a list of IP addresses. Use 1 method to limit access. This provides a 2nd option to limit using a bucket policy.
2. NOTE: This assumes that a proxy or source IP NAT is in use and Cloud provider will only see your public ip address. If you have a range or pool of IP addresses than you need to include all IP addresses used by your NAT or proxy implementation.
3. An S3 bucket policy can also be used to allow access from a range of IP's or a specific ip address. Use this example to restrict access to your Data Center public IP address.
4. Get your public facing ip address that will be used the by Golden Copy or Virtual Accelerator Nodes.
    a. Method #1 - curl :
        i. Login to Golden Copy VM over ssh and run this command "curl ifconfig.io" .
        ii. This should return the IPv4 ip address configured for public Internet access to use with the policy.
    b. Method #2 - visit an ip locate website from a data center subnet :
        i. Google for "what is my ip address" to get the IP v4 ip address.
5. Replace the x.x.x.x with your ip address in the example policy below.
6. In the Amazon S3 service console click on the storage bucket configured for Golden Copy:

a. 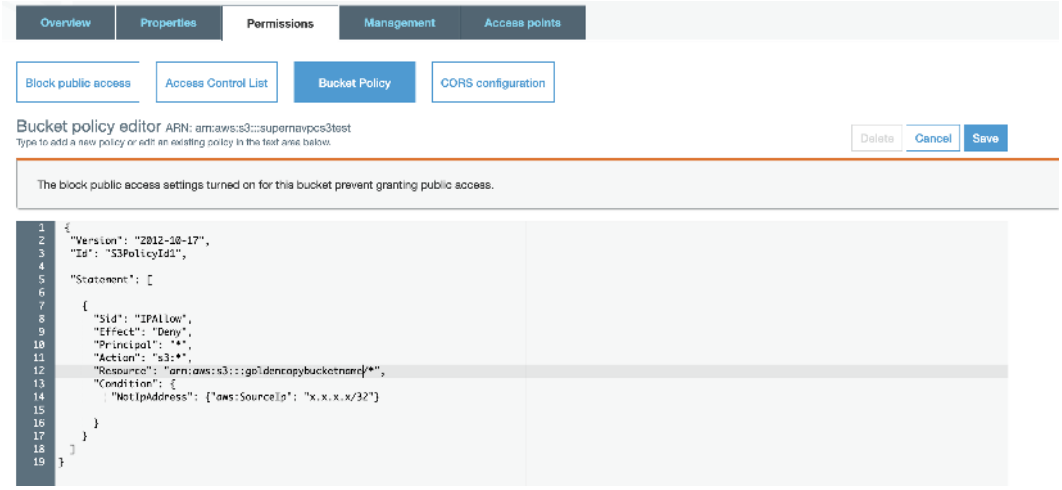

7. Replace x.x.x.x with your ip address, replace **goldencopybucketname** with your storage bucket name.
8. **NOTE: to get the Bucket ARN for the resource property. You can see this next to the Bucket Policy Editor in the screen shot above.**
9. Edit the policy text shown below and save to the bucket policy .
10. Done.

```
{
"Version": "2012-10-17",
"Id": "S3PolicyId1",

"Statement": [

{
"Sid": "IPAllow",
"Effect": "Deny",
"Principal": "*",
"Action": "s3:*",
"Resource": "arn:aws:s3:::goldencopybucketname/*",
"Condition": {
"NotIpAddress": {"aws:SourceIp": "x.x.x.x/32"}
}
```

```
}
]
}
```

How to Enable and Use Accelerated Transfer Mode on a bucket

1. **Requirements:**
    a. Release 1.1.6 or later
2. Why use this mode? (AWS documentation reference)
    a. You might want to use Transfer Acceleration on a bucket for various reasons, including the following:
        i. You have customers that upload to a centralized bucket from all over the world.
        ii. You transfer gigabytes to terabytes of data on a regular basis across continents.
        iii. You are unable to utilize all of your available bandwidth over the Internet when uploading to Amazon S3.
3. Enabling Transfer Acceleration on an Amazon S3 Bucket Guide Link.
4. How to Configure in Golden Copy
    a. When adding the folder the --aws-accelerated-mode true flag must be used.
    b. The endpoint must use the Accelerated endpoint FQDN (i.e bucketname.s3-accelerate.amazonaws.com) to access an acceleration-enabled bucket.

Google Cloud Storage Creation Walk Through

1. Requirements:

    a. Storage Buckets should use the Uniform Access Control method when setting up new storage bucket. See the guide here for details.

2. Login to the Google Cloud Platform console and open IAM to create a service account, click create and then Done.

3. 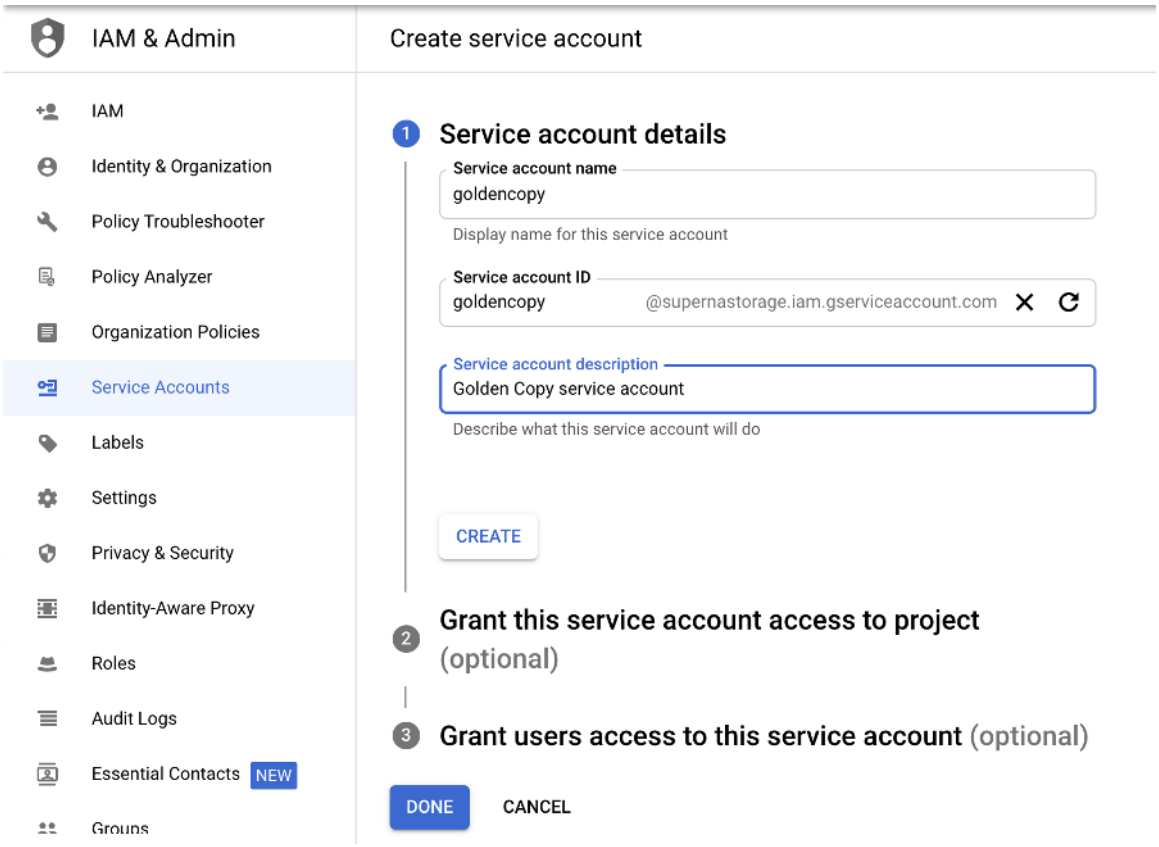

4. Click on the created service account to create a json key and click on Create new key select Json

5. 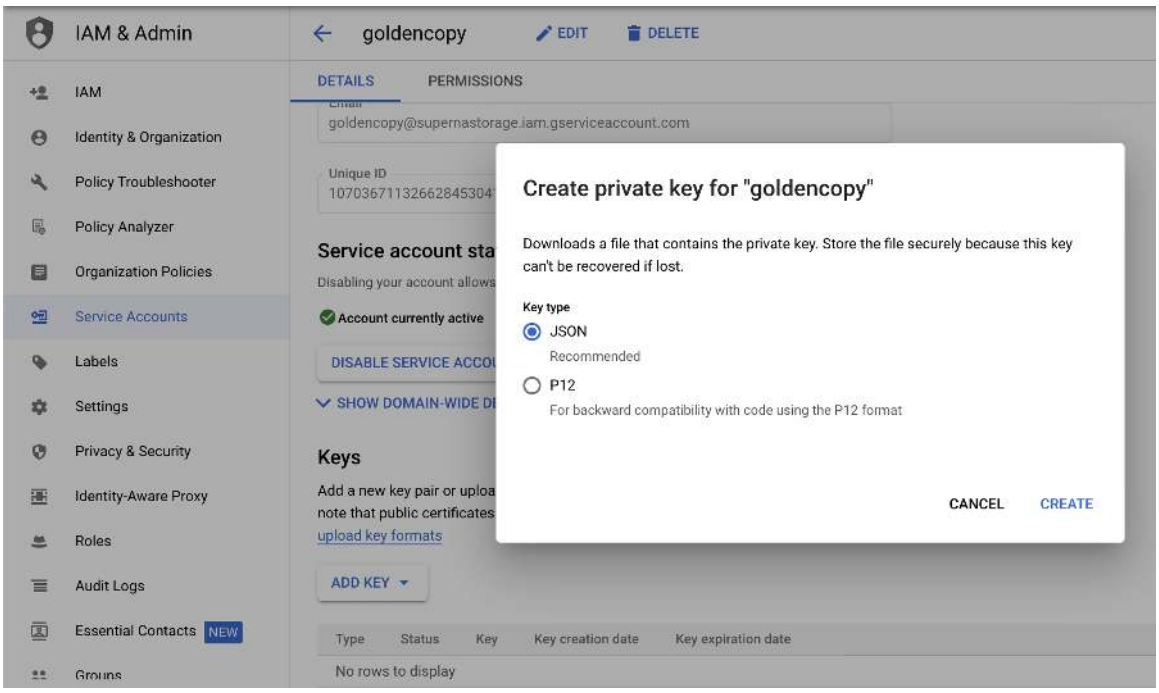

6. The file will be downloaded to your pc and used later.

7. 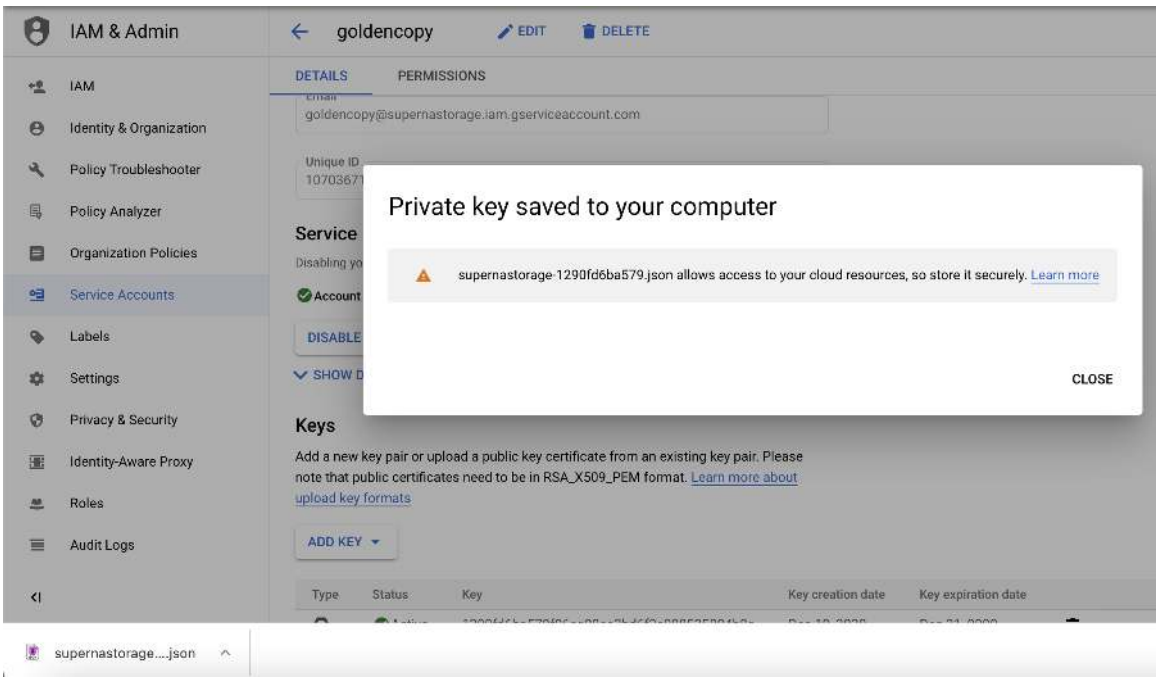

8. Locate the Storage bucket created to store Golden Copy data (steps not covered to create buckets, see Google Documentation) and select the permissions tab. Click Add. NOTE: You will need the email ID on the service account created in the step above.

9. 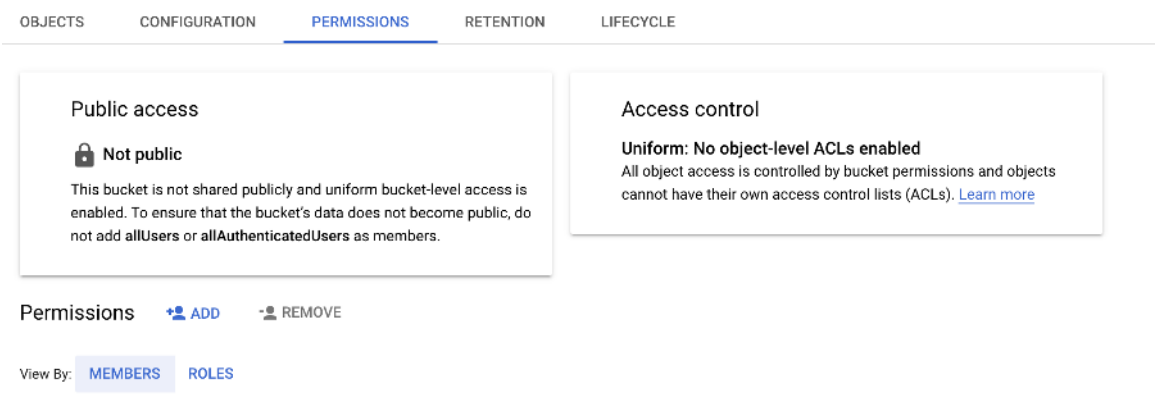

10. Use the Storage Object Creator role.

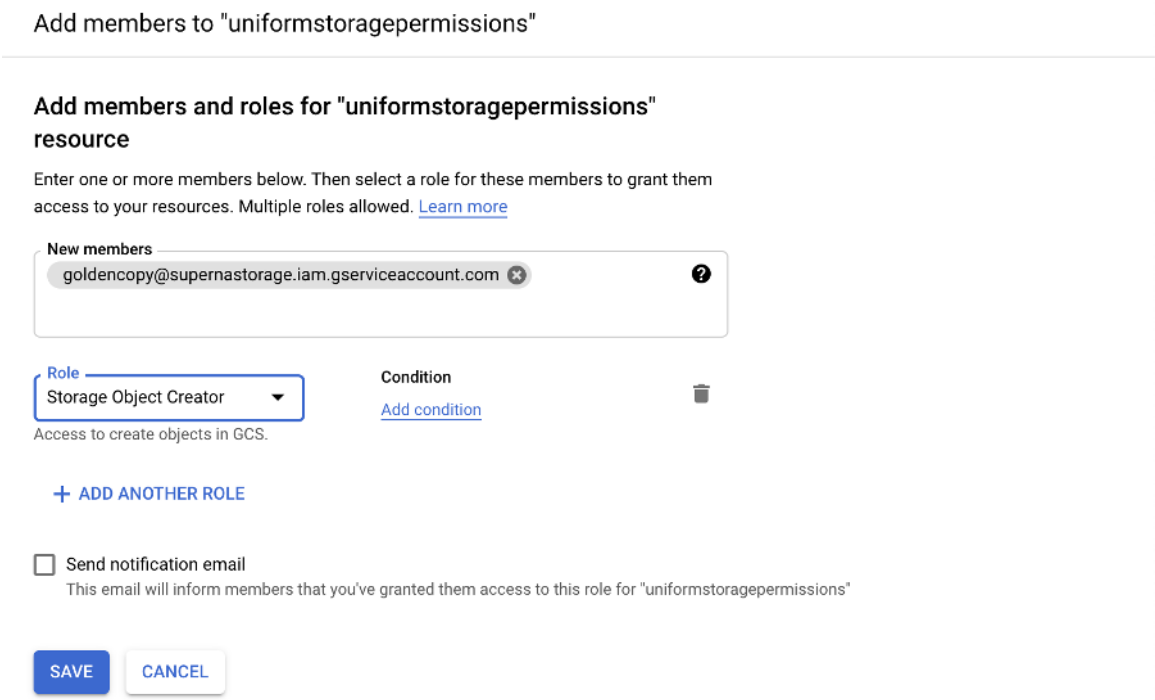

11.

12. Repeat these steps on each Storage Bucket that will be used with Golden Copy.

13. Copy the json authentication file to Golden Copy node 1 with winscp to the /home/ecaadmin path.

    a. Authenticate using the ecaadmin user.

b. This completes all the Google console steps and the remaining steps will cover how to add a Google Cloud Storage end point in the quick start section of this guide.

Azure Blob Storage Creation Walk Through

Azure Blob Storage is not compatible with S3 protocol and uses a proprietary REST API. The configuration is similar to S3 bucket concepts. A container is similar to a bucket and a blob is the same as an object.

How to create Azure Storage Account

This process creates a storage account if one does not already exist. The storage account contains a Container that will store data copied from Golden Copy.

1. Login to Azure console.
2. Select Home --> Storage accounts --> Create storage account --> + sign to create .
3. Assign or create a "Resource group" (i.e."goldencopy"), enter "Storage account name" (i.e. "goldencopy2"), fill in the remaining fields and click "Next: Networking".

4. Home > Storage accounts > Create storage account

Create storage account

your resources.

Subscription * Visual Studio Enterprise

Resource group * goldencopy

Create new

Instance details

The default deployment model is Resource Manager, which supports the latest Azure features. You may choose to deploy using the classic deployment model instead. Choose classic deployment model

Storage account name * goldencopy2

Location * (US) East US

Performance Standard Premium

Account kind StorageV2 (general purpose v2)

Replication Read-access geo-redundant storage (RA-GRS)

Access tier (default) Cool Hot

Review + create < Previous Next : Networking >

5. Select a Networking configuration, note authentication will be used. This section limits the networks or virtual networks this blob will be reachable. Configuring networking to allow your on site data center network to be able to reach the endpoints. Consult Azure Documentation on how to secure access from your on premise network to Azure.

6. 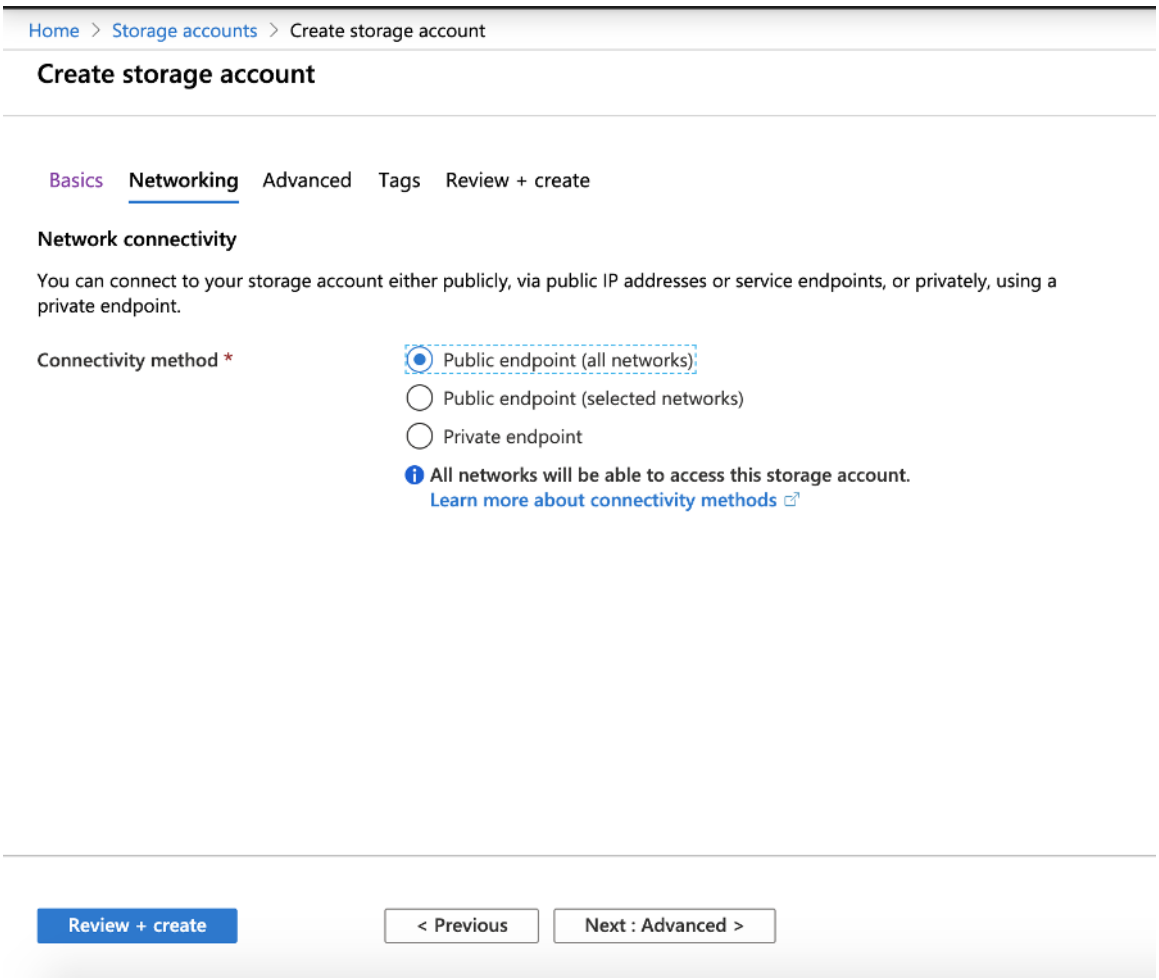

7. Click "Next: Advanced" and accept Default settings or change according to your requirements.

8. 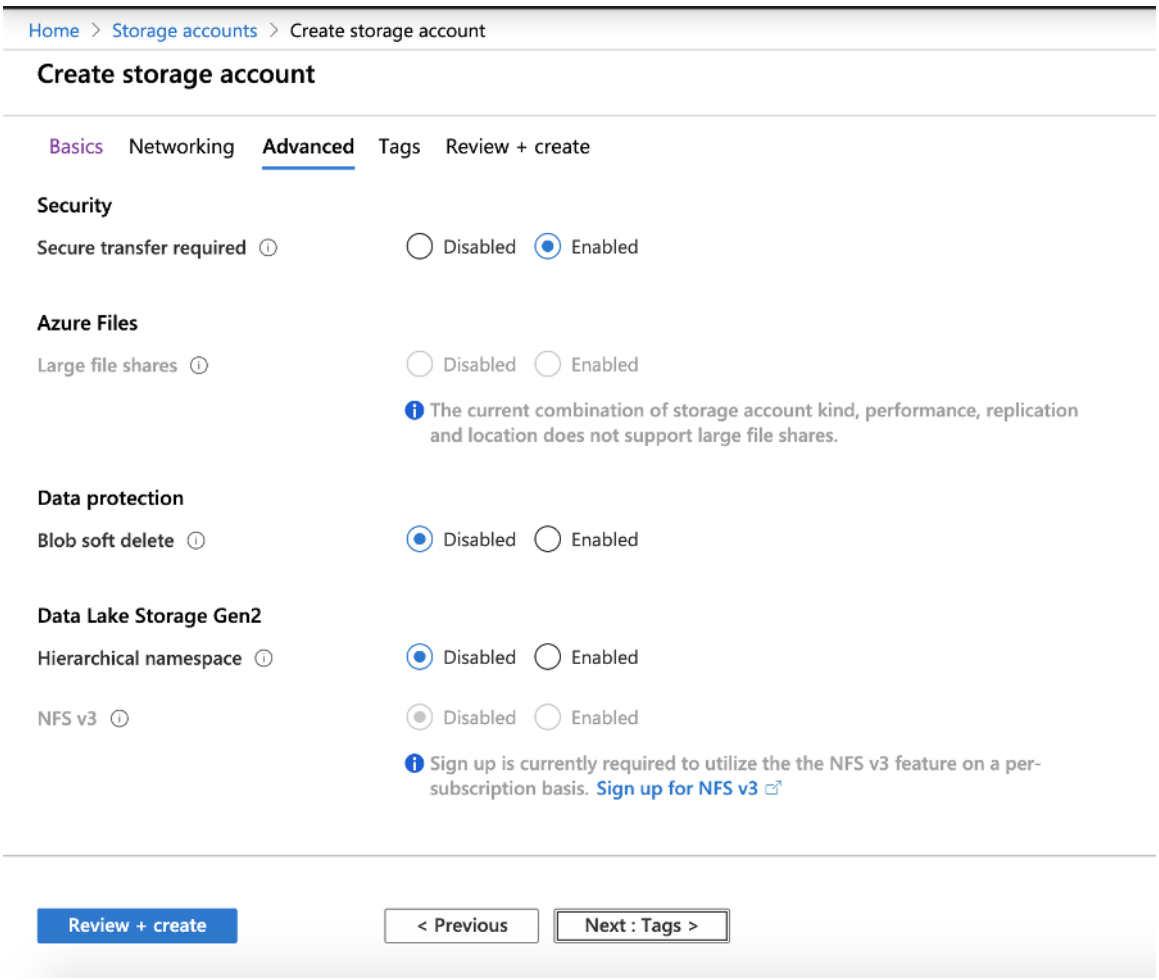

9. Click "Next: Tags".
10. Leave the Tags Screen at defaults click "Next".
11. On the Final screen review all settings and click create.
12. The completes the Storage Account Configuration. Continue to the next section to create a container.

How to create an Azure Blob Container in a Storage Account

1. Select the Storage Account in the Azure console.
2. On the menu on the left, under"Blob service" click on the "Containers" option.

3. 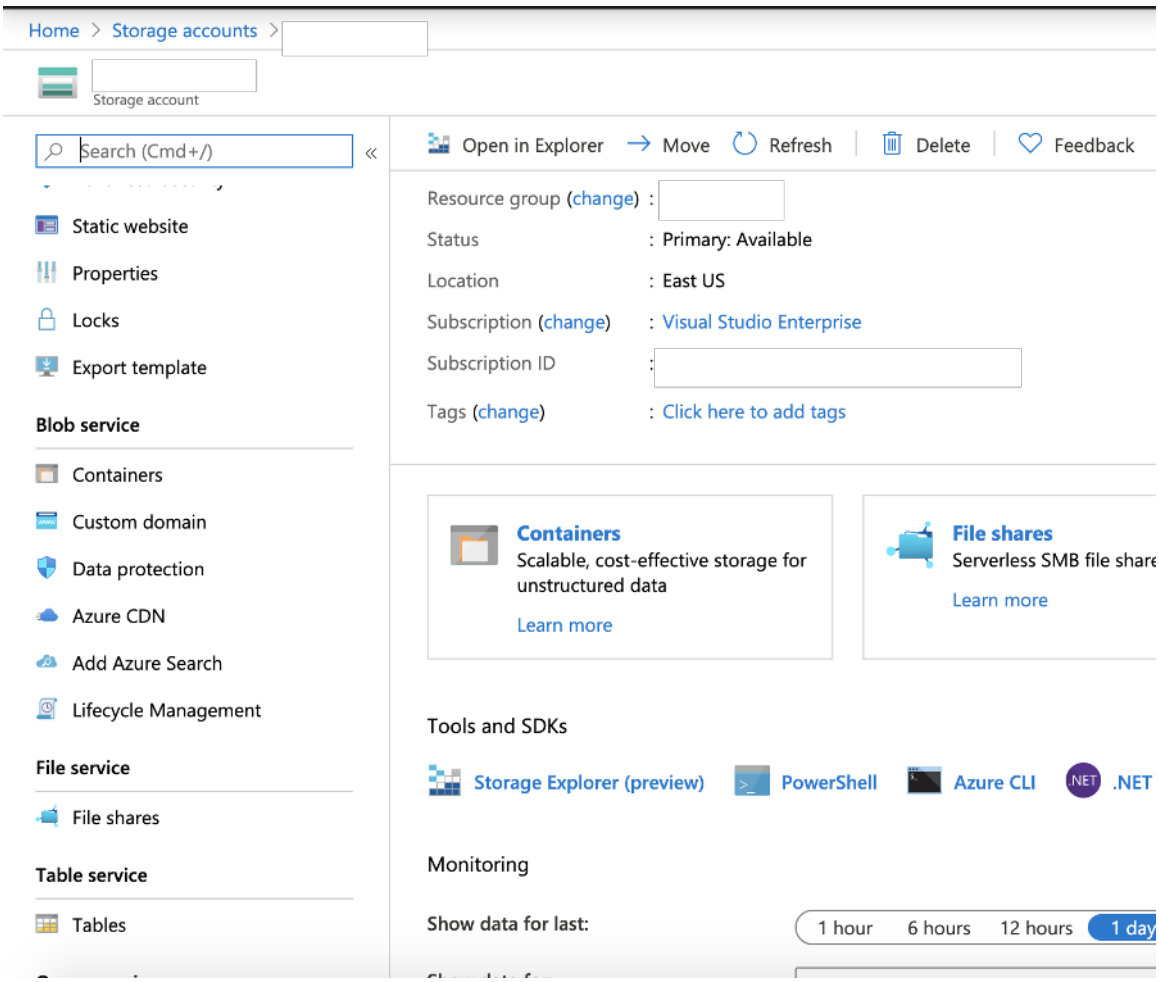

4. Click the + to create a "New container", under "Name" enter a unique name and leave the default setting as "Private no anonymous access". Click "Create".

5. 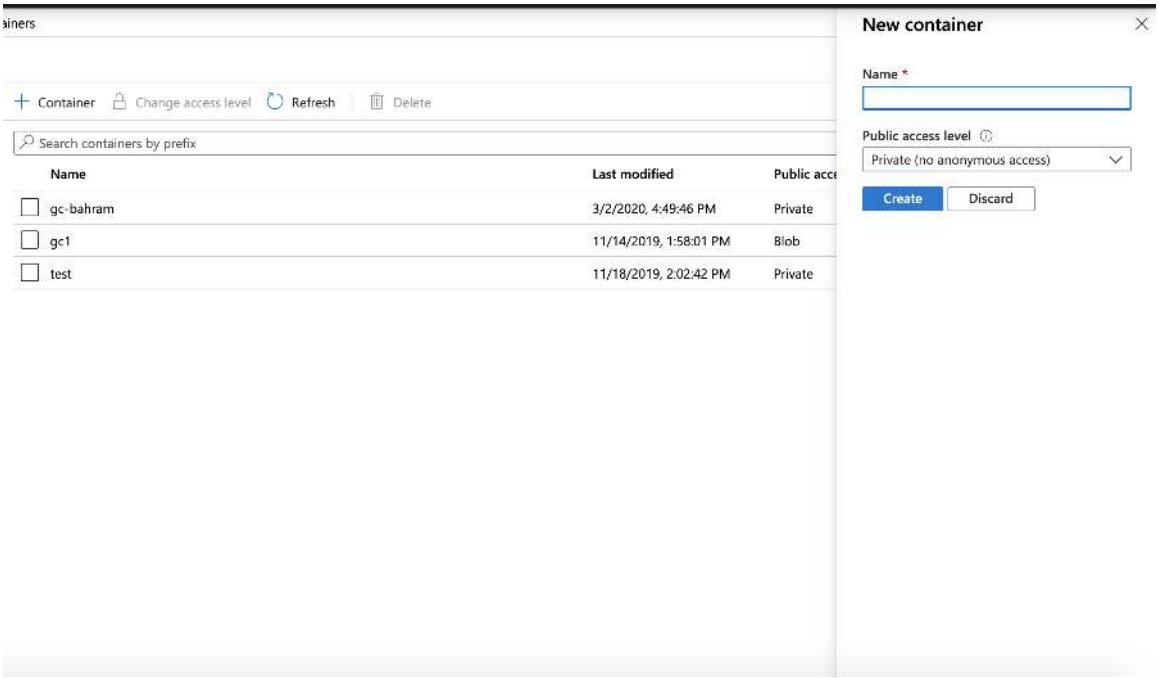

6. After creation click on the "Properties" tab and record the URL to the container for use when configuring Golden Copy folders.

7.

8. Authentication keys for Blob access can be found on the Storage account --> Settings --> Access Keys .

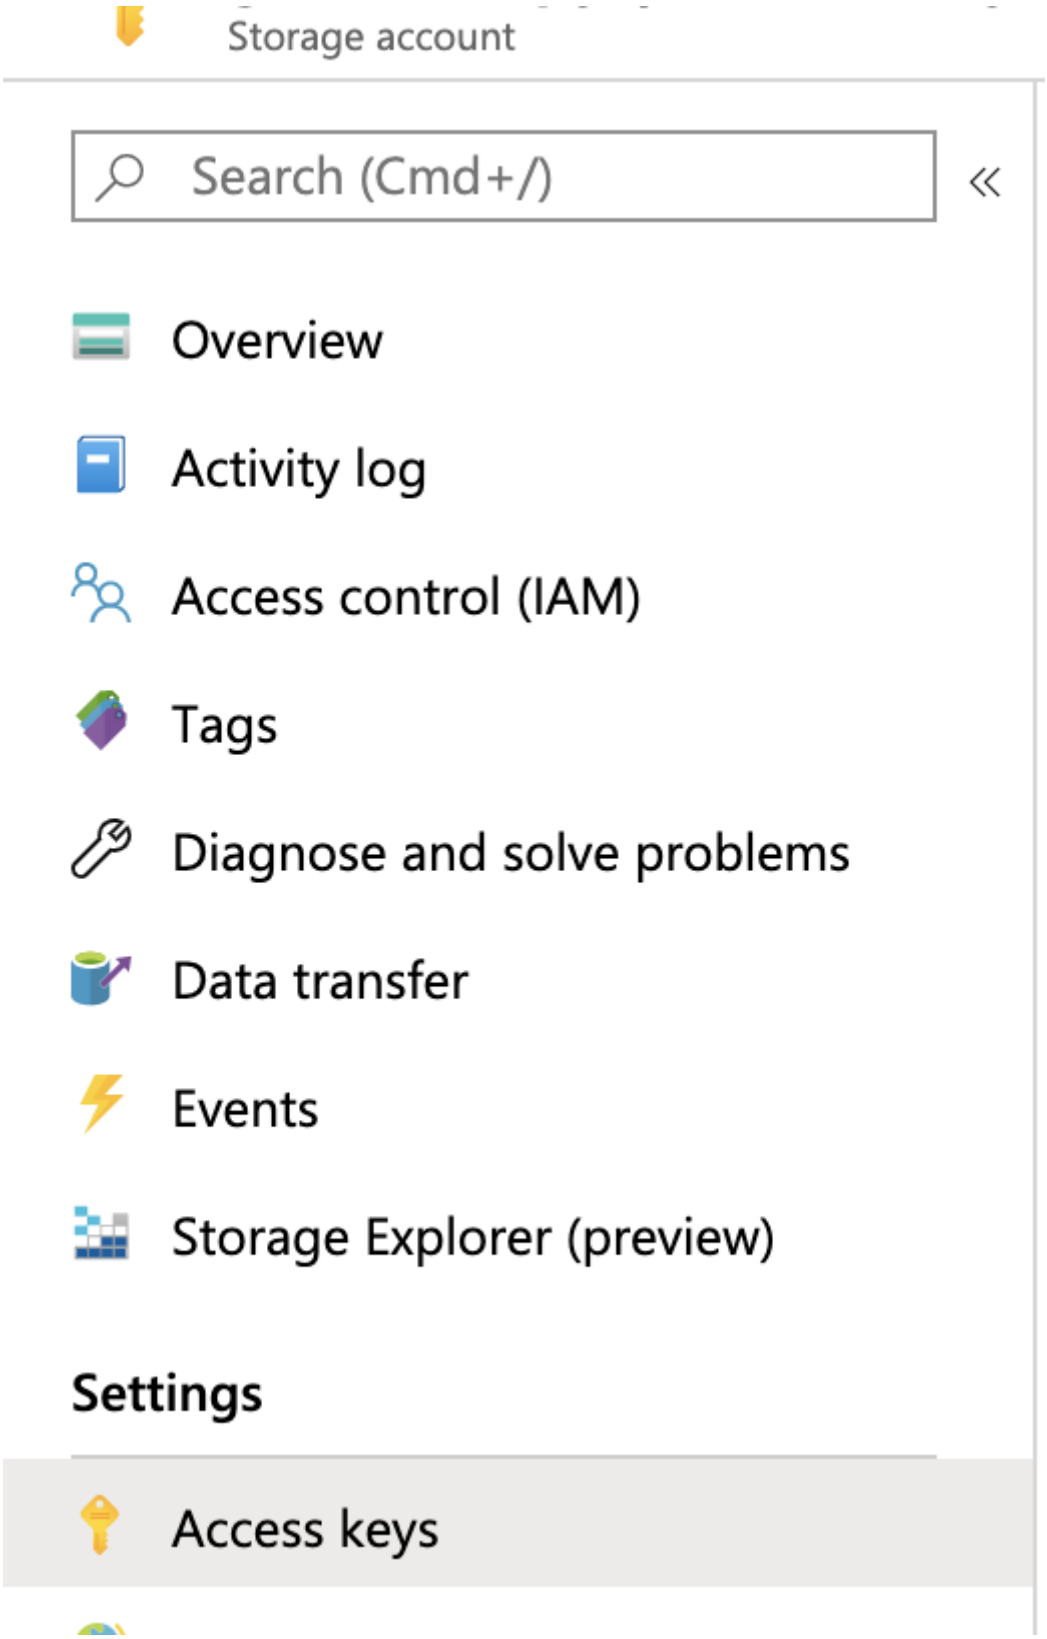

9.

10. On this screen record the Key 1 access key for use when configuring Golden Copy .

11. Done, you will need the following information from these steps

a. **Storage Account.**

b. **Container URL.**

c. **Access Key for the Storage Account.**

Cohesity Walk Through Example

How to Create the Storage View and Service Account User on a Cohesity Storage Array

1. Create Service Account for Authentication with Golden Copy:
    a. Under the Admin menu --> Access Management.
    b. Create a new "Service Account User" with local user option with the admin role assigned (i.e. "eyeglassSR").
    c. **NOTE: To ensure the storage view gets the bucket ACL's assigned, it is required to create the storage view logged in as the service account user created in this step.**
2. Login to the console as the new Service Account User (see note above).
3. A Storage View is required with only the S3 protocol enabled for write access over the S3 protocol.
4. The name of the view will generate an S3 bucket name. Use all lowercase and no special characters to create the Storage View name.
5. This Storage View name will be needed when adding the folder to Golden Copy folder configuration using the bucket parameter.
6. In the console use the Platform Menu --> View to create the new Storage view for S3

   a. Logged in as the new user, view the user in Admin menu --> Access Management --> click on the user and record the access ID and the secret ID.
   b. Use the values shown on the step above when adding the folder using the example above.
7. Done.

Appliance Global Settings Configuration

How to set Default Settings for Snapshot Expiry for all Folders

Folders have default settings that can be viewed and set using commands in this section.

1. **How to set the Global Default snapshot Expiry for long running copy jobs**
   a. searchctl archivedfolders configure --full-copy-snap-expiry x (where x is in days)
   b. **NOTE: Use this command to extend the source snapshot taken for a long running copy job that could take many days to complete. The default is 10 days unless changed.**
2. **How to view the Global Default settings on folders**
   a. **searchctl archivedfolders getConfig**
3. **NOTE: Each folder add command also allows a per folder override , see folder add command to set per folder.**

How to set File Checksum Global Settings for All Folders

1. This option enables software MD5 checksum and includes the checksum in the S3 headers to allow the AWS, ECS or other S3 targets to recompute the MD5 checksum before saving the object.

a. **Commands**

i. checksum property is controlled by **searchctl archivedfolders configure** -h

1. searchctl archivedfolders configure --checksum <ON | OFF | default>

a. value ON computes the checksum and adds a visible metadata property for files and folders. **NOTE: This requires a file to be read twice, once for the checksum and a 2nd time to upload it with the checksum visible on the metadata This option is preferred since independent data integrity is now possible for stored objects after they have been copied.**

b. value OFF does not compute the checksum and avoids any in-built default checksum computing

c. value DEFAULT is the default behavior for each cloud target

i. OFF for Azure and partial ON

ii. AWS/ S3 Targets of type other, which means it is calculated but no visible property on the object)

iii. Google Cloud - TBD

ii. checksum property can be viewed with

1. searchctl archivedfolders getConfig

Advanced Configurations to Appliance Configuration

The parameters allow advanced changes to the appliance and should only be changed when directed by support.

1. nano /opt/superna/eca/eca-env-common.conf **Note: Only make changes if advised by support!**
    a. export ARCHIVE_PARALLEL_MULTI_UPLOAD=True (this enables multi part upload for files over a certain size)
    b. export ARCHIVE_MAX_PART_SIZE= (10MB default for larger files , leave at default)
    c. export ARCHIVE_PARALLEL_THREAD_COUNT=10 (sets the number of concurrent files to copy across all PowerScale nodes in the pool)
    d. export ARCHIVE_S3_PREFIX="igls-" **OneFS 9 requirement** example: to replace xxx with string to prefix folder objects that store ACL information and avoids a collection with directory names on NAS devices example "igls-"
        i. Each folder object will have this prefix applied to avoid the collection with duplicate folder names. This flag will apply a prefix to folders that are uploaded that are stored to retain ACL's on folders. You will see for each folder uploaded, one folder stores files, and an additional object with the same name as the folder. This is where the ACL is stored for the folder that contains files.
        ii. When using the file system based S3 target storage example, OneFS 9 with S3 support, 2 folders with the

same name are not supported. This requires the folder object to have a different name. This prefix will be applied to folder objects when the target storage enforces unique object keys and object names. **Consult support before using this flag.**

e. export ARCHIVE_ENABLE_SDK=true (Enables Native SDK mode - default mode as of 1.1.4 > 178 build)

f. export ARCHIVE_PARALLEL_THREAD_COUNT=100 (number of concurrent files per Golden Copy VM or Accelerator Nodes, increasing this number may not increase performance unless sufficient bandwidth is available)

g. export ARCHIVE_PARALLEL_THREAD_SDK=10 (Number of separate threads used to copy a single large file, higher number will increase bandwidth utilization)

h. export ARCHIVE_SMART_INCREMENTAL=false (default value)

    i. Change the value to true to enable, this will need to be added to common conf file and cluster down and up to take effect.

    ii. This enables a fast incremental feature to skip collecting metadata that requires additional api calls to collect, when this is true metadata will only support owner, group (none text value, and sid/gid only), no mode bits will be collected.

    iii. **This is only recommended if a very high change rate of files is expected and the performance of fast incremental is required.**

i. export ARCHIVE_BLOCK_PARALLEL_JOBS=true (default true, Blocking Parallel on single VM Golden copy, multi VM deployments must be used to allow parallel jobs).
    i. Set to false to allow multi VM deployments to run multiple copy jobs.
j. export ARCHIVE_FULL_PARALLEL_JOBS_ALLOWED=x
    i. Default is 1 change to a value no greater than 30 and requires multi VM deployment for a supported configuration
k. export ARCHIVE_INCREMENTAL_PARALLEL_JOBS_ALLOWED=x
    i. Default is 1 change to a value no greater than 30 and requires multi VM deployment for a supported configuration
l. export ARCHIVE_TOTAL_JOBS_ALLOWED=60
    i. Default 2, for 1 Full and 1 Incremental, Maximum supported value is 60 which allows for 30 folders with full copy jobs and 30 incremental jobs

© Superna LLC

10.9. Golden Copy Back Bundle & Adv License Configuration Steps

Home Top

- Overview
- Golden Copy Advanced License And Backup Bundle Feature Configuration
  - Overview
    - How to Assign an Advanced License to a Cluster
    - Cloud Storage Tier Aware Copy and Sync (AWS, Azure)
    - Version Aware Recall / Restore
      - Overview
        - Use Case #1: Recall "hot data" first
        - Use Case #2: Recall "cold data" last
        - Use Case #3: Recall files with a specific created or modified time stamp
    - How to redirect recall object data to a different target cluster
      - Overview
    - Target Object Store Stats Job to Monitor the backup object store
      - Overview
      - How to Configure Full Backup Mode Overview
      - How to configure full backup mode CLI commands
      - How to Report on the object count and q uantity of data protected on a folder Definition

    - How to Run a Data Integrity Job
    - How to Enable Ransomware Defender Smart Airgap
- Golden Copy Pipeline Workflow License Overview
    - Overview
    - Requirements
    - Configuration Examples
- Automation with Golden Copy Rest API
    - Overview
    - How to use the Golden Copy API

Overview

This topic covers installations licensed with the backup bundle or advanced license key features. These features require the license key installed.

Golden Copy Advanced License And Backup Bundle Feature Configuration

Overview

These features require the Advanced license addon for Golden Copy to use the Backup use case features. The backup features make managing backup workflows easier monitor and automated reporting. New features

to protect data allow several new workflows. Data integrity features , API access for automation with external tools and new reporting options.

Requirements

1. Advanced Golden Copy license applied
2. Golden Copy Backup Bundle license

How to Assign an Advanced License to a Cluster

1. A Golden Copy Advanced license enables the features in this guide for a cluster managed by Golden Copy. Use the steps below.
2. searchctl isilons license --name <clusterName> --applications GCA
    a. NOTE: An available advanced license must exist for this command to succeed and the cluster also requires a base Golden Copy license before an Advanced license key can be applied.

Cloud Storage Tier Aware Copy and Sync (AWS, Azure)

1. This feature allows copying or syncing data directly into an archive tier for AWS S3 and Azure targets. Google Cloud storage requires creating the bucket with a storage class and does not support setting the tier of individual objects.
    a. This feature avoids life cycle policies to move objects to different tiers.

2. The add folder and modify folder CLI command allows specifying the tier that objects should be copied into.
    a. [--tier] default is **standard (AWS), Cool (Azure)**
        i. Requires Golden Copy Advanced license or Backup Bundle license
        ii. **Requires 1.1.6**
        iii. **Azure**
            1. flag to specify the tier the API calls are sent to, this should match the container tier configuration options are Access tier for Azure e.g. **hot, cool, archive**) Without this flag the default is cold.
        iv. **AWS**
            1. specify AWS tier using (**STANDARD** (default), **STANDARD_IA, GLACIER, DEEP_ARCHIVE, INTELLIGENT_TIERING, ONEZONE_IA, OUTPOSTS, REDUCED_REDUNDANCY**) Use upper case tier name.
            2. NOTE: Not all tier options are valid for all use cases. Consult AWS documentation.
3. Example command
    a. searchctl archivedfolders add --isilon gcsource --folder /ifs/archive --secretkey xxx --endpoint blob.core.windows.net --container gc1 --accesskey yyyy --cloudtype azure --tier STANDARD_IA

Version Aware Recall / Restore

Overview

1. Requires:
    a. Release 1.1.6
2. Full and incremental with S3 bucket versioning allows multiple versions of files to be protected using S3 policies configured on the target storage. The Storage bucket must have versioning enabled and the folder should be configured in sync mode, or have run a copy job multiple times to detect file changes and update objects with a new version. **NOTE: Storage bucket version configuration is external to Golden Copy consult your S3 storage device documentation.**
    a. This allows recall jobs to select files based on a date range using older than x date or newer than Y date. This allows selecting files based object creation date (the date the backup ran) using the older newer than flags on the recall job.
    b. **NOTE: The date range is evaluated against the object creation date of the object in the version history of an object. This date is when the object was backed up.**
    c. **NOTE: If you run multiple recall jobs with the same path the files in the recall staging area under /ifs/goldencopy/recall will be overwritten if they already exist.**
3. This feature also adds the ability to scan the metadata in the objects properties to recall files based on created or modified data stamps of the files that existed on the file system at the time they were backed up.
4. **The recall command adds to new options with the following date syntax. Use double quotes.**

a. --newer-than "<date and time>" (yyyy-mm-dd HH:MM:SS e.g 2020-09-14 14:01:00)

b. --older--than "<date and time>" (yyyy-mm-dd HH:MM:SS e.g 2020-09-14 14:01:00)

5. **Use Case #1: Recall "hot data" first**

   a. This solution is when a large recall of data is needed to be recalled /restored but you want the most recent data recalled first. This would use the newer than flag to select a date example 2 weeks in the past.

   b. Example

      i. searchctl archivedfolders recall --id 3fd5f459aab4f84e --subdir /ifs/xxx **--newer--than** "2020-09-14 14:01:00" --apply-metadata

6. **Use Case #2: Recall "cold data" last**

   a. This solution would be used to recall data after "hot data" is recalled since it has not been recently updated. This would use a recall job and the older than flag , using the example above using the same date 2 weeks in the past with the older than flag would start a recall job to locate and recall data unmodified that is at least 2 weeks or older.

   b. **Example**

      i. searchctl archivedfolders recall --id 3fd5f459aab4f84e --subdir /ifs/xxx **--older--than** "2020-09-14 14:01:00" --apply-metadata

7. **Use Case #3: Recall files with a specific created or modified time stamp**

    a. This use case allows scanning the metadata that Golden Copy encodes into the properties of the objects as criteria to select data to recall the data based on the created date stamp of the files or modified time time stamp.

    b. [--start-time STARTTIME] (yyyy-mm-dd HH:MM:SS e.g 2020-09-14 14:01:00)

    c. [--end-time ENDTIME] (yyyy-mm-dd HH:MM:SS e.g 2020-09-14 14:01:00)

    d. [--timestamps-type {modified, created}] The default is modified date stamp. The files are backed up with created and modified time stamps, this flag allows selecting which time stamp to use when evaluating the older than or newer than dates.

    e. Example to scan for files with a last modified date stamp between Sept 14, 2020 and Sept 30 2020 under the /ifs/xxx folder.

        i. searchctl archivedfolders recall --id 3fd5f459aab4f84e --subdir /ifs/xxx --start-time "2020-09-14 14:01:00" --end-time "2020-09-30 14:01:00" --apply-metadata

    f. Example to scan for files with a created date date stamp between Sept 14, 2020 and Sept 30 2020 under the /ifs/xxx folder.

        i. searchctl archivedfolders recall --id 3fd5f459aab4f84e --subdir /ifs/xxx --start-time "2020-09-14 14:01:00" --end-time "2020-09-30 14:01:00" **--timestamps-type created** --apply-metadata

How to redirect recall object data to a different target cluster

Overview

1. Use this option to add a cluster to Golden copy that does not require a license, since it will be used as a recall target only. This option requires the Advanced or Backup Bundle license key.
1. **Requires**
    a. Release 1.1.6
    b. **Advanced license key or backup bundle license**
2. This process requires adding the target cluster to Golden Copy, this cluster does not require a license when using the --goldencopy-recall-only flag.
3. searchctl isilons add --host IP **address** --user EyeglassSR [--isilon-ips x.x.x.x, y.y.y.y] --goldencopy-recall-only
4. Follow steps below:
    a. searchctl archivedfolders recall --id ID [--subdir SUBDIR] --target-cluster TARGET**CLUSTER** **--apply-metadata**
    b. **NOTE: The recall NFS mount must be created on the target cluster on the /ifs/goldencopy/recall path before a recall can be started.**
    c. Replace **TARGETCLUSTER** with the redirected cluster name added above to redirect the restore to the new target cluster.

Target Object Store Stats Job to Monitor the backup object store

Overview

This job type will scan all the objects protected by a folder definition and will provide a count of objects, the sum total of data along with the the age of the objects.

In addition it will summarize:

1. File count
2. Total data stored
3. oldest file found
4. newest file found

How to Configure Full Backup Mode Overview

Overview

This feature allows specifying the number of full backup copies to maintain. This feature must integrate with the target storage data retention feature that will automatically delete full backup copies after X days. The feature will be applied to a folder and accepts the number of copies to maintain. Each copy will be placed into a folder in the target storage with a date stamp with all data placed under this folder. The folder definition includes a scheduled interval to make each full copy example weekly or monthly. The data retention is configured on the target bucket to 2 x the retention setting to have a previous copy and current copy available for recall.

Example: full copies interval is weekly (every 7 days) with 2 full copies maintained at all times, the bucket retention would be set to 15 days so that 2 full weekly copies will be retained at all times. On day 15 the first copy will be deleted by the storage target.

How to configure full backup mode CLI commands

1. When adding a folder a new parameter --backup-num x is added to indicate how many full copies this folder is going to be retained. This flag will create a folder with a date stamp in the target storage device each time a full archive job is executed against this folder.
    a. Example --backup-num of 3 with a full archive schedule of weekly and S3 storage bucket retention set to 22 days (7 x 3 = 21 + 1 day). The folder will need a schedule applied using the full archive schedule to weekly. On the 22nd day the first backup taken will be deleted by the storage target. On the 21st day 3 full copies of the data exist in 3 different folders with date stamps when each full backup job was started.
2. The full archive schedule can be configured with the command below. This example is every 7 days on Sunday at midnight.
    a. searchctl archivedfolders add (other parameters) --full-archive-schedule "0 0 * * 0"
3. Make sure to set the object retention on the storage bucket used on the folder definition. This is what will handle the deletes of object data. Follow your vendors S3 documentation to set object retention features.

How to Report on the object count and quantity of data protected on a folder Definition

1. searchctl archivedfolders s3stat --id <folderID>
    a. NOTE: Usage charges for cloud provider storage will be assessed based on API list and get requests to objects. The job can be canceled at any time using searchctl jobs cancel command.

b. Use the searchctl jobs view --follow --id xxxxx (job ID) to view the results

2. Sample output

a. 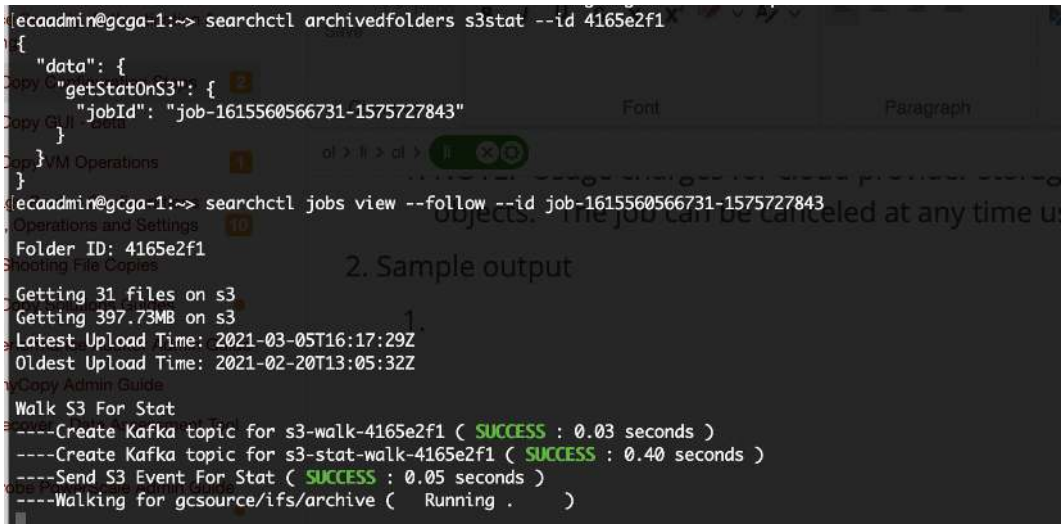

How to Run a Data Integrity Job

1. **Overview**

   a. This feature provides a data integrity audit of folders by leveraging the metadata checksum custom property. Random files are selected for audit on the target device , downloaded the file computes the checksum and compares to the checksum stored in the metadata. If any files fail the audit the job report will summarize failures and successful audit passing. This verifies your target device stored the data correctly.

b. 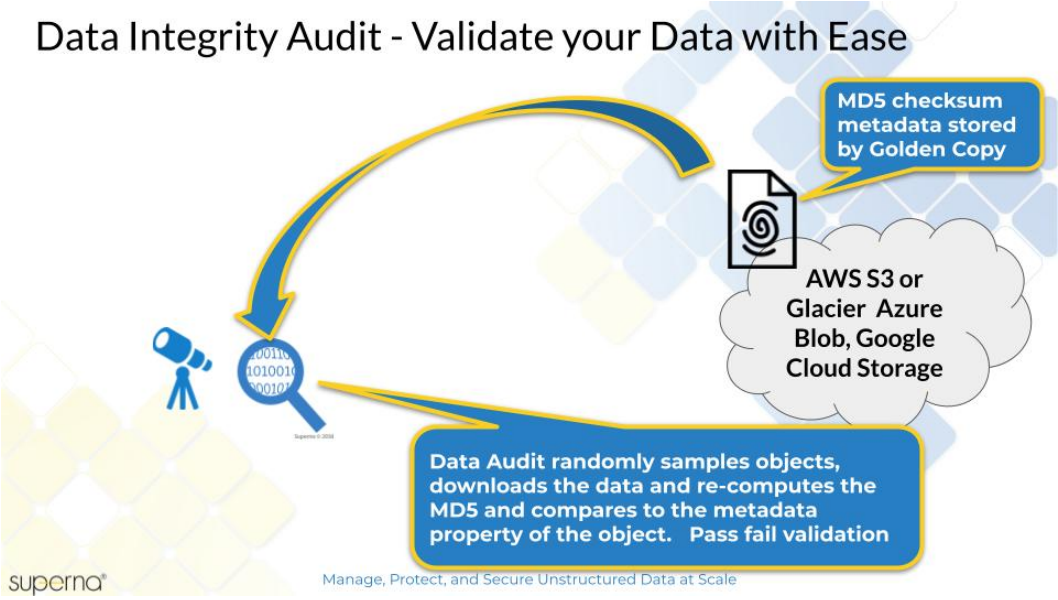

2. **Requirements**
    a. Release 1.1.6 update 2 or later
    b. **NOTE: If data did not have a checksum applied during a copy the job will return 100% error rate.**
    c. This feature will require data is copied with the --checksum ON global flag enabled. See the Global configuration settings.
        i. searchctl archivedfolders getConfig
        ii. searchctl archivedfolders configure -h
3. **How to run a data integrity job**
    a. searchctl archivedfolders audit --folderid yy --size x (note x is GB total of data that will be randomly audited, and yy is the folder ID that will be audited)
    b. Using the job id to monitor the success of the job
        i. searchctl jobs view --id xxxxx

ii.

How to Enable Ransomware Defender Smart Airgap

1. This feature integrates with Ransomware Defender to enable Smart Airgap. This blocks full or incremental jobs if an active alarm is raised in Ransomware Defender or Easy Auditor Active Auditor triggers.
2. login to node 1 of Golden Copy as ecaadmin
3. nano /opt/superna/eca/eca-env-common.conf
4. Add these variables and generate a new API token in eyeglass to authenticate Golden Copy API calls. This can be completed from Eyeglass GUI, main action menu, Eyeglass Rest API, API tokens and create a new new token named for Golden Copy.
5. Add these variables and enter the eyeglass ip address and api token.
6. export EYEGLASS_LOCATION=x.x.x.x
7. export EYEGLASS_API_TOKEN=yyyy
    a. copy api token created in eyeglass api menu to replace yyyy
8. export ARCHIVE_RSW_CHECK_THRESHOLD=WARNING

   a. **Options are WARNING, MAJOR, CRITICAL to determine the severity of alarm that will block the backup process. If set to warning then all severities will block, if set to Major then Warnings will be ignored, if set to critical then warning and major will be ignored**

9. export ARCHIVE_RSW_CHECK_INTERVAL_MINS=1
   a. **How often to poll eyeglass for ransomware events, recommended to use 1 minute**
10. export ARCHIVE_RSW_CHECK_ENABLED=TRUE
   a. **True / False to enable the functionality. True required to activate the feature.**
11. control + x to save the file
12. ecactl cluster down
13. ecactl cluster up

Golden Copy Pipeline Workflow License Overview

Overview

This feature license allows S3 to file workflows and S3 to S3 work flows in addition to the File to Object, Object to File workflow available in Golden Copy. The S3 to File direction require incremental detection feature that will leverage the date stamps on the file system set to match the object time stamp. This will allow incremental sync from S3 to file or S3 to S3.

Use Cases

1. Media workflows to pickup media contribution from a 3rd party from an Cloud S3 bucket and transfer the data to an on premise Powerscale for editing workflows
2. Media workflow to download S3 output from a rendering farm that produces output that is needed on premise for video editing workflows.
3. HPC cloud analysis for AI/ML that requires on premise data to be copied to an S3 Cloud bucket for analysis input to AI/ML that produces an output in a different bucket that needs to be copied back on premise.

These workflows are file to object and object to file with different source and destinations along with scheduled copies or incremental sync in both directions example on premise to cloud and cloud back on premise.

The solution is designed to allow scheduled incremental in both directions to pickup new or modified files only from the S3 bucket and copy to the cluster.

Requirements

1. Pipeline license key applied to Golden Copy

Configuration Examples

1. Use this command to add a folder to receive data from an S3 bucket
    a. searchctl archivedfolders add/modify --isilon HOST --folder PATH [--source-path SOURCEPATH] [--recall-schedule RECALLCRON]

i. --source-path "SOURCEPATH" (enter the path with double quotes ")

ii. the path on s3 that contains the data we want to copy. Defaults to /<clustername>/<foldername>.

iii. --recall-schedule RECALLCRON Recall job Cron expression e.g. "*/1 * * * *"

2. searchctl archivedfolders recall --id ID --source-path SOURCEPATH

   a. --id ID Id of archived folder

   b. --source-path SOURCEPATH the path on s3 that contains the data to copy.

3. Example to create a pipeline configuration from an S3 bucket and path to a file system path on the cluster

   a. searchctl archivedfolders add --folder /ifs/notUploadedByGC --isilon ofs3830 --source-path "/notUploadedByGC" --recall-schedule "*/30 * * *" --cloudtype aws --bucket gcsoak2 --secretkey <> --accesskey xxxxx

      i. --folder /ifs/notUploadedByGC path where the data will be copied to on the cluster

      ii. --source-path "/notUploadedByGC" - The S3 path to start the copy from in the bucket

      iii. --recall-schedule "*/30 * * *" - The schedule to scan the S3 bucket to copy new data or modified data found in the bucket. This example is scanning the S3 bucket every 30 minutes

b. After adding the folder you can run the job to scan the bucket and copy data
    i. searchctl archivedfolders recall --id <folderID> --sourcePath "/notUploadedByGC"

Automation with Golden Copy Rest API

Overview

The rest API can be used to automate copy jobs, monitoring of jobs to allow integration to application work flows that require data movement tasks from a file system to S3 or from S3 back to a file system. Examples include media work flows, machine learning and AI training where Cloud computing is used with data and the results are returned to on premise file systems.

How to use the Golden Copy API

© Superna LLC

10.10. Golden Copy GUI - Beta

Home Top

- Overview
- Login to Golden Copy GUI - Beta
- System Monitoring Health Dashboard Tab - GA
  - Monitor Archived Folder Statistics - GA
- Archived Folders Tab - Beta
- Isilon Management Tab - Beta
- Archive Jobs Tab - Beta

Overview

Covers basic functions of the GUI.

Login to Golden Copy GUI - Beta

1. Before login a local user must be added
2. searchctl settings admins add --name ecaadmin --local .
3. Then access https://x.x.x.x to login .
4. Main Screen:

a. 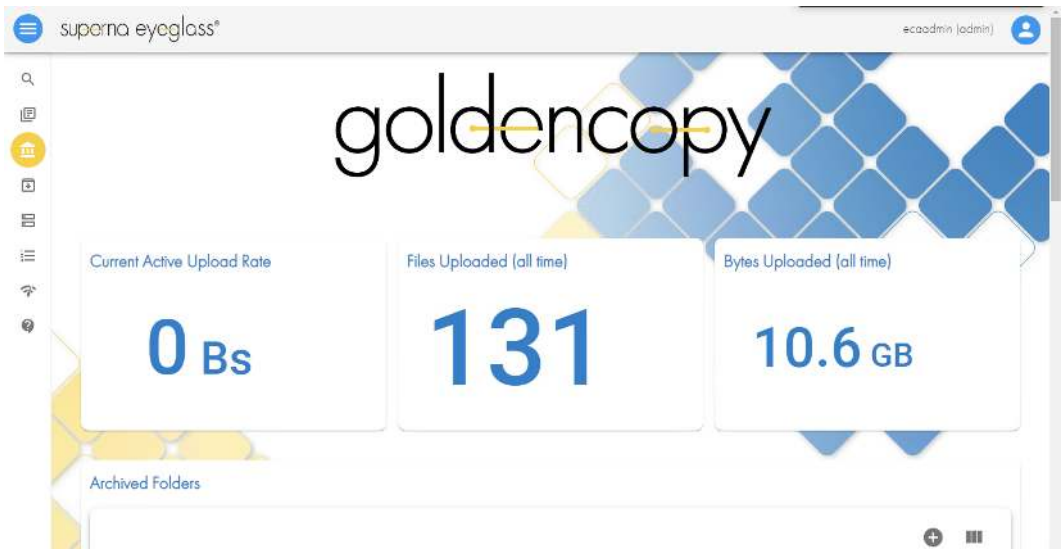

System Monitoring Health Dashboard Tab - GA

1. Appliance health metrics for all components are displayed here in real-time with history .

2. 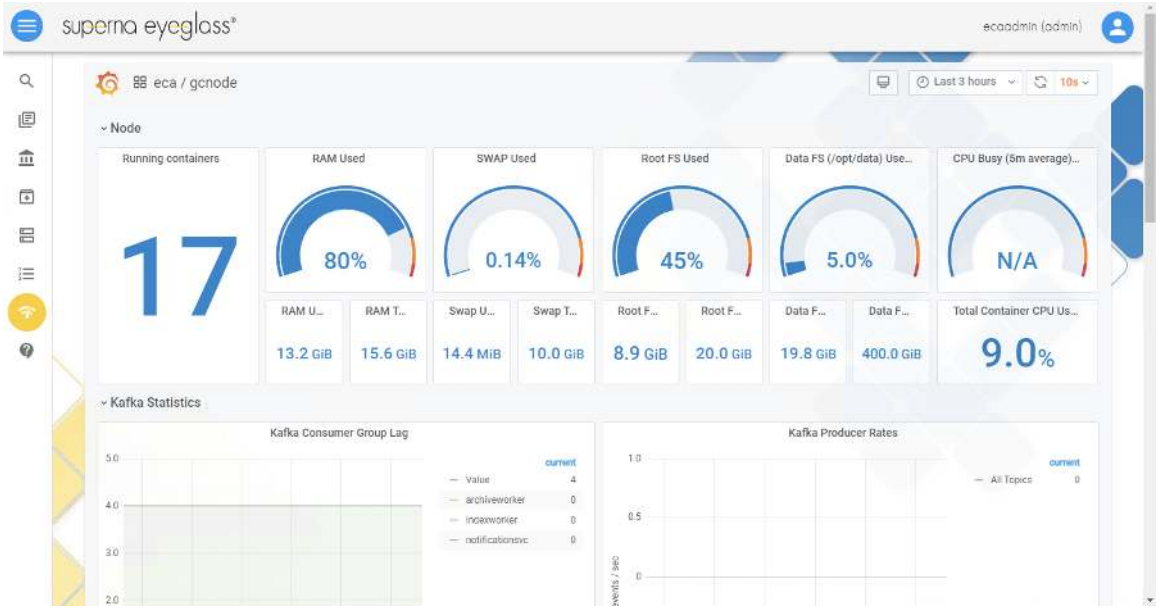

Monitor Archived Folder Statistics - GA

1. Click on the System Icon to load the default dashboard
2. Click on the ECA /Node text at the top right
3. Select gcfolder dashboard

a.

b. The gcfolder dashboard monitors stats on all folders and shows real-time updates of cumulative file count and files copied per second in the bottom graphs. Mouse over the graph to see files waiting to be uploaded, archived or errored.

c. 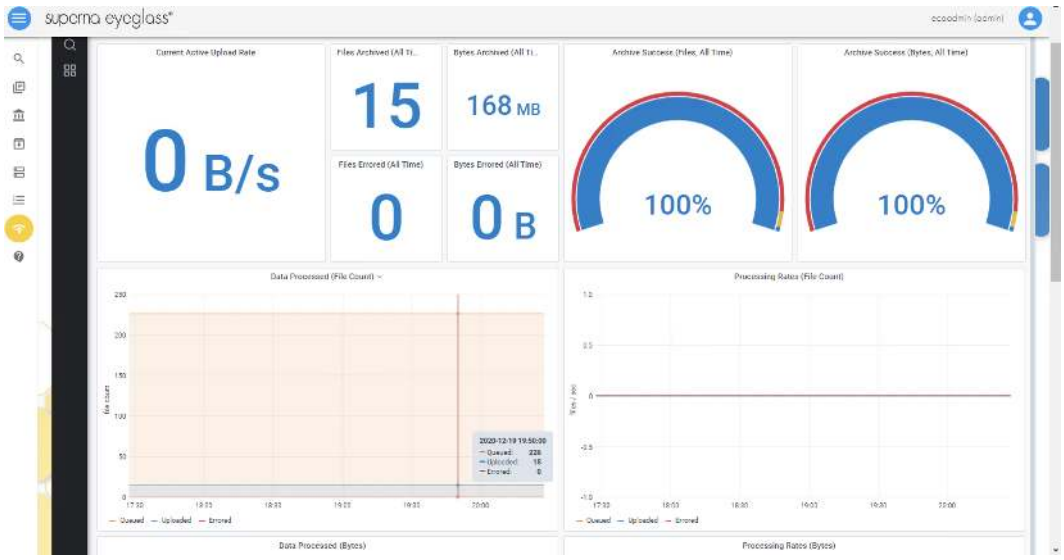

Archived Folders Tab - Beta

1. 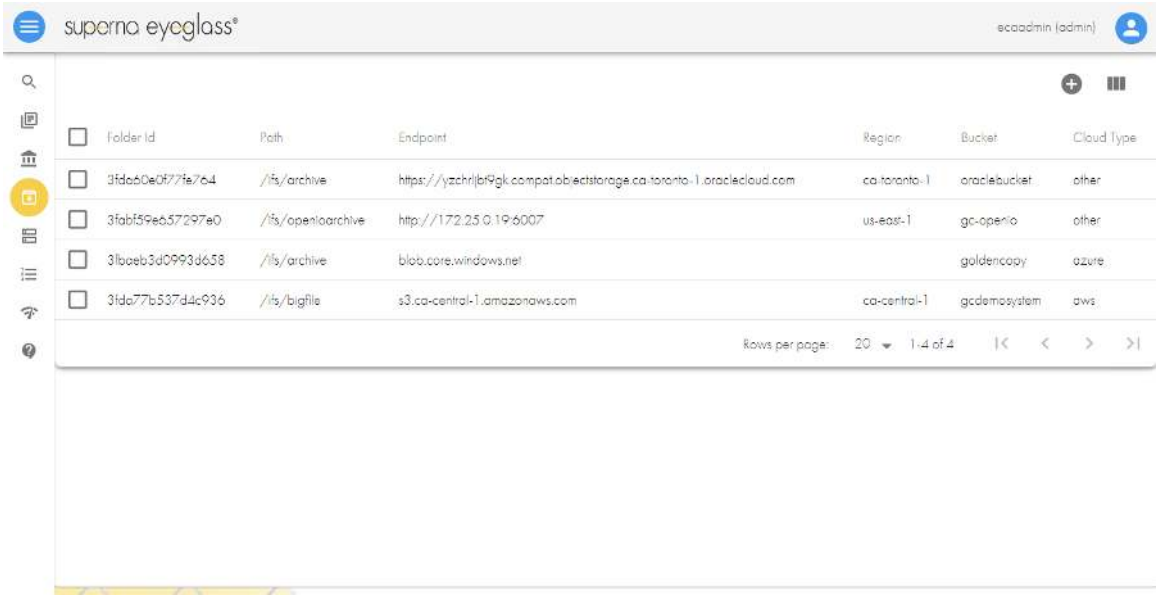

2. Lists all configured folders with option to add new columns to the display.

3. Click the plus sign to add new folder configurations.

4.

Isilon Management Tab - Beta

1. Add, delete Isilon clusters.

2. 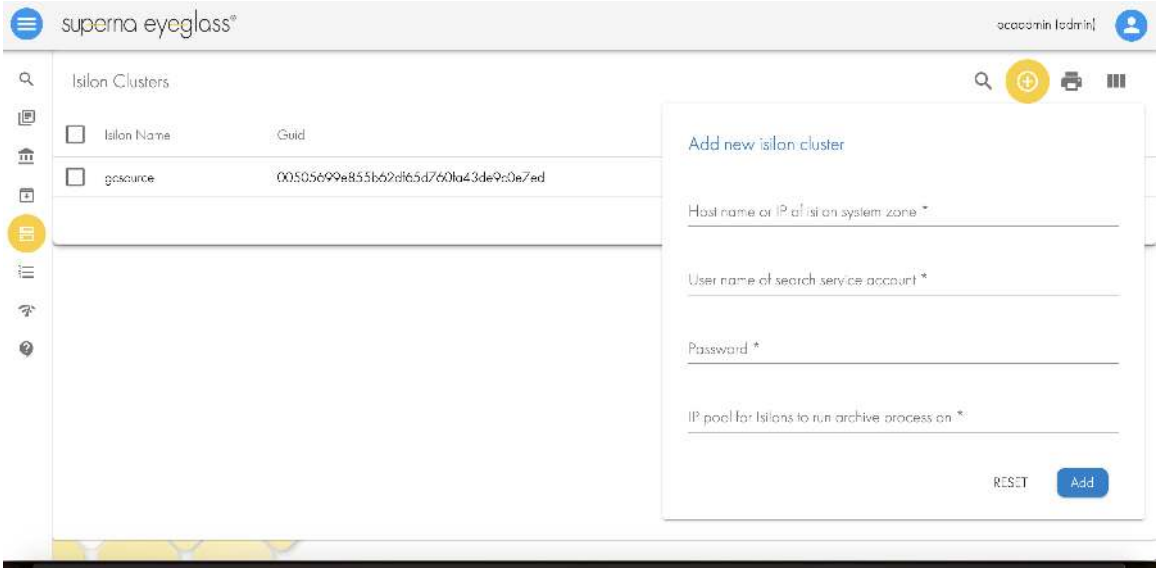

Archive Jobs Tab - Beta

1. View active copy jobs and history of previous jobs.

2. 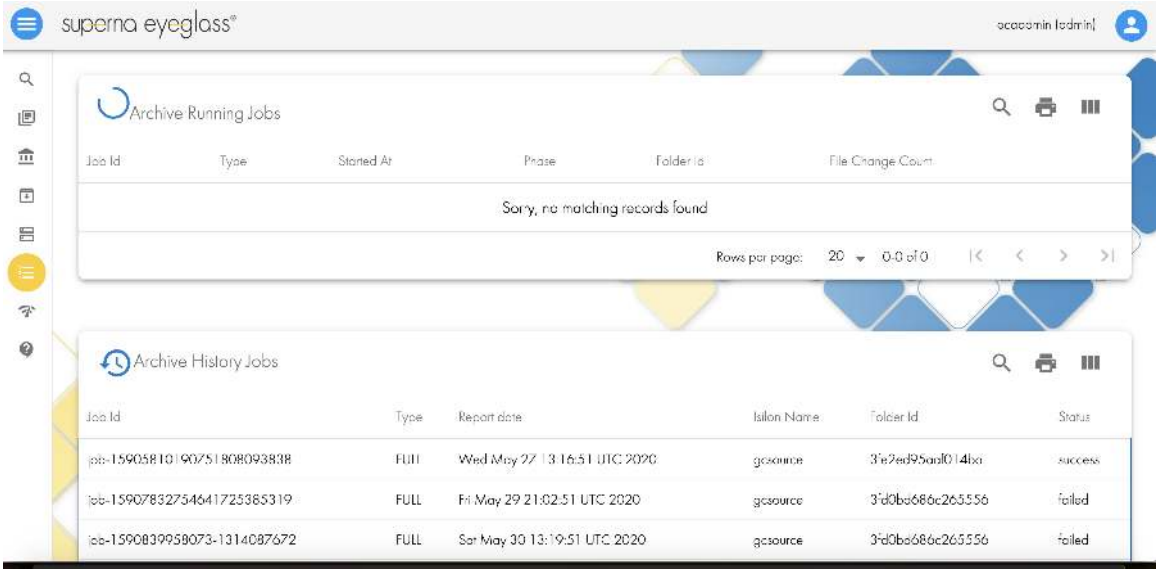

© Superna LLC

10.11. Golden Copy VM Operations

Home Top

- Cluster Operations CLI commands
- How to Start and Stop the Cluster
- How to Change the IP address of an PowerScale Cluster
- How to change TLS security settings when connecting to clusters that do not support the highest security algorithms
- How to Enable and Use PhoneHome support
- How to collect support logs and submit a support case
- Backing up and Restore the Cluster Configuration
- How to check for Alarms
- How to configure Alarm notification
  - Quick Start SMTP Configuration for Notifications
  - Setup Syslog channel for Notifications
  - Setup an SMTP or syslog Channel for Notifications
  - Create a Notification Group
  - Manage Recipients for SMTP and Syslog channels
  - Manage Notification Suppression Alarm Configurations
- How to Remove a Cluster Node from Active Copy Operations

Cluster Operations CLI commands

The following sections cover cluster operation commands.

How to Start and Stop the Cluster

1. **ecactl cluster down (use this to shutdown all services).**

2. **Use this script to startup all services:**
    a. **ecactl cluster up .**

How to Change the IP address of an PowerScale Cluster

How to change TLS security settings when connecting to clusters that do not support the highest security algorithms

How to Enable and Use PhoneHome support

How to collect support logs and submit a support case

Backing up and Restore the Cluster Configuration

How to check for Alarms

How to configure Alarm notification

Quick Start SMTP Configuration for Notifications

Setup Syslog channel for Notifications

Setup an SMTP or syslog Channel for Notifications

Create a Notification Group

Manage Recipients for SMTP and Syslog channels

Manage Notification Suppression Alarm Configurations

How to Remove a Cluster Node from Active Copy Operations

1. If a cluster node that was added for copy job operations is no longer required, or needs maintenance, or has performance issues it is possible to remove a node on the fly with active copy jobs. This will stop sending files to this node(s).
    a. searchctl isilons modify --name <name of cluster> --isilon-ips x.x.x.x y.y.y.y .(Enter the list of IP's that should be used and remove the IP of the

node you no longer want to receive copy job files from the list or range of IP's)

b. See options for list of ip's or a range in "Add a cluster to Inventory" in the Golden Copy Configuration Steps guide for more details.

© Superna LLC

10.12. S3 Storage Bucket Configurations Options , Operations and Settings

Home Top

- Overview
- How to Cancel a Running Archive Job
- How to Configure S3 Storage bucket retention, versioning to meet your File copy Archive Use case
  - Use Case: Data Retention with Scheduled Copies with or without bucket Versioning
  - Use Case: One Time Copy Data Archiving
- How to clean up orphaned multi part file uploads
  - Amazon S3 Procedure
- How to Configure Amazon AWS Glacier
  - How to apply single File storage class change for testing Glacier
  - How To recall a file from Glacier:
  - How to configure Glacier Lifecycle policies on a storage bucket
- How to Configure Storage Tier Lifecycle Policies with Azure Blob Storage

Overview

This section covers storage bucket optional configurations and operational procedures.

How to Cancel a Running Archive Job

1. First get a list of running jobs.
2. searchctl jobs running (this command will return job id's).
3. searchctl jobs cancel --id job-xxxxxxxx (enter the job id you want to cancel. NOTE: this will stop copying files, but files that are already copied will be left in the S3 storage up to the time the copy was canceled).

How to Configure S3 Storage bucket retention, versioning to meet your File copy Archive Use case

Use Case: Data Retention with Scheduled Copies with or without bucket Versioning

1. This copy mode will run on a schedule to copy all folders configured without --manual flag setting.
2. The full SmartCopy will check if the target S3 file exists and skip existing files that are the same version.
3. Files that are modified AND a previous version exists in the S3 storage will be updated in the S3 storage.
4. Object Retention:
    a. Best practice is to set the S3 retention policy at the bucket level so that all new objects automatically get retention set per object. Use different storage buckets with different retention policies or use versioning feature on the S3 storage target to set retention per version of the file. See vendor documentation on how to configure retention policies.
5. (Optional) Enable Bucket versioning to keep previous versions with an expiry set for each version in days. This will will allow modified files to be uploaded while preserving the old version of the file should it need to be restored.
    a. Example: For Amazon S3 enables each version of a file uploaded with 45 day expiry.

b. 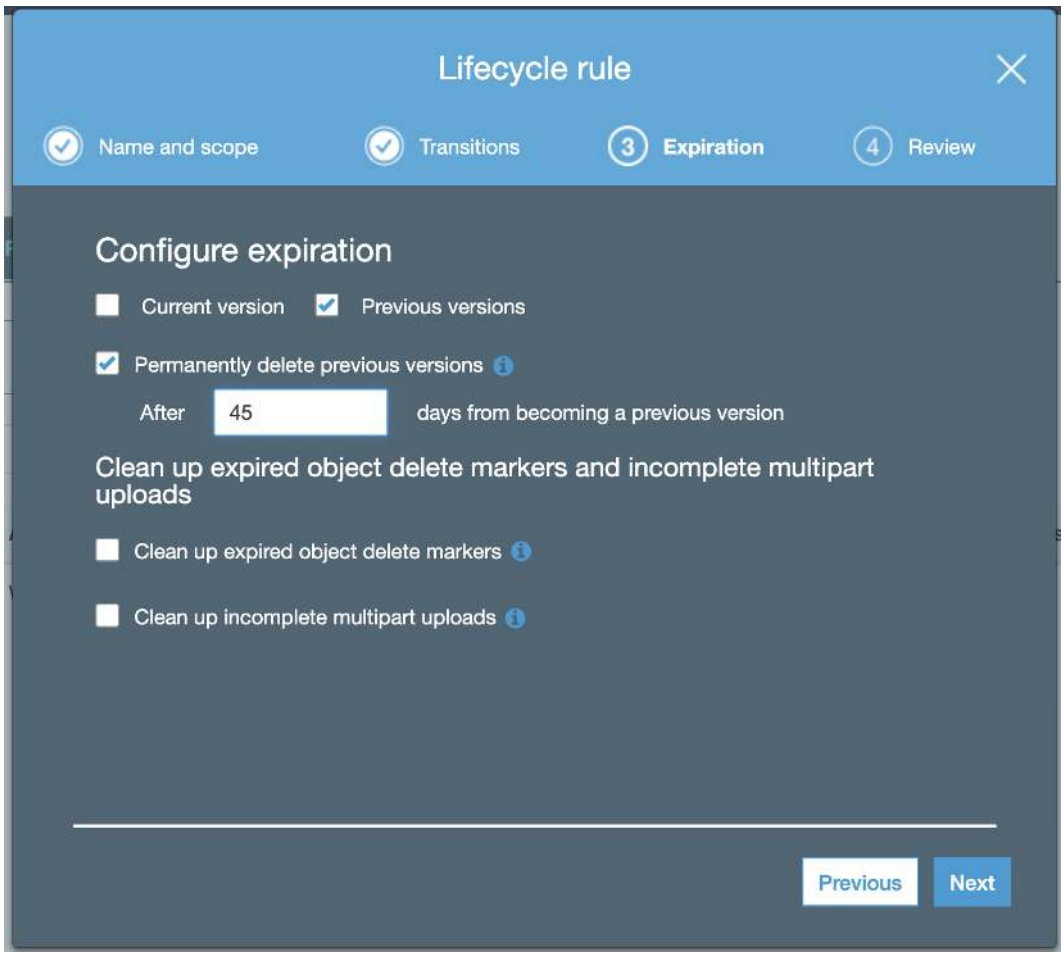

c. NOTE: If a file changes several times between Full Copy jobs only the last version modified prior to the copy job will be stored.

Use Case: One Time Copy Data Archiving

1. To clean up data and copy to S3 storage for long term archive and legal retention. This use case is a 2 step process.
2. Create a folder target configuration but add the --manual flag which excludes the folder from scheduled copies.
3. Run the archive job (see guide here on running archive jobs).
4. The copy job will run and review the copy results for success and failures. See guide here.
5. Once you have verified all all the data is copied with an S3 browser tool (i.e. Cyberduck, or S3 Browser), run the PowerScale Tree Delete command to submit a job to delete the data that was copied. Consult Dell Documentation on Tree Delete command.

6. **Recommendation:** Create a storage bucket for long term archive and set the Time to Live on the storage bucket to met your data retention requirements. Consult the S3 target device documentation on how to configure the TTL.
7. **Copy to new target path:**
    a. The run archive job has a flag to add a prefix with --prefix xxx where xxx is the prefix to add to the S3 path (see guide here) to the S3 storage bucket path which is useful to copy the data to a 2nd path in the S3 storage and not overwrite the existing path. This allows creating multiple folders of the same source path in the target S3 storage. This can be used for present the S3 data over file sharing protocols (i.e. ECS feature or direct access to the S3 data for application use).
8. NOTE: You can run a copy job multiple times using the CLI if data has changed on the source path and a full smart copy job will run again.

How to clean up orphaned multi part file uploads

Multi part upload can fail for various reasons and leaves orphaned incomplete files in the storage bucket that should be cleaned up since they are not a complete file. See the procedures below for each S3 provider. **NOTE: Golden Copy will issue an Abort Multi Part API command to instruct the target to delete partial uploads. The procedures below should still be enabled in case the S3 target does not complete the clean up or does not support the Abort Multi Part API command.**

Amazon S3 Procedure

1. Login to Amazon S3 console.
2. Click on the storage bucket name.
3. Click on Management tab.

4.

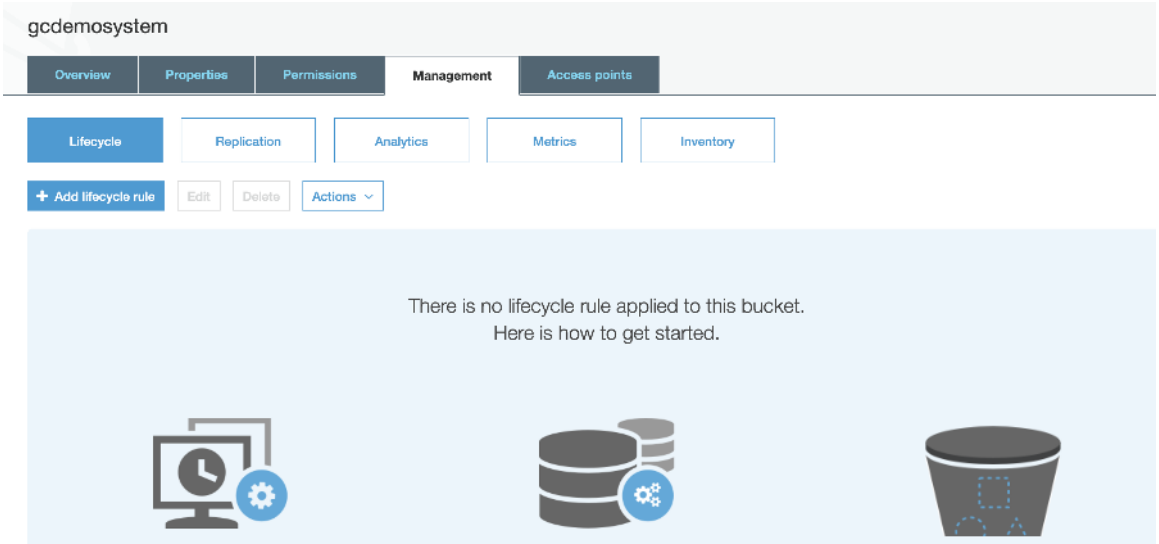

5. Click Add life cycle rule.

6. Name the rule incomplete uploads:

7.

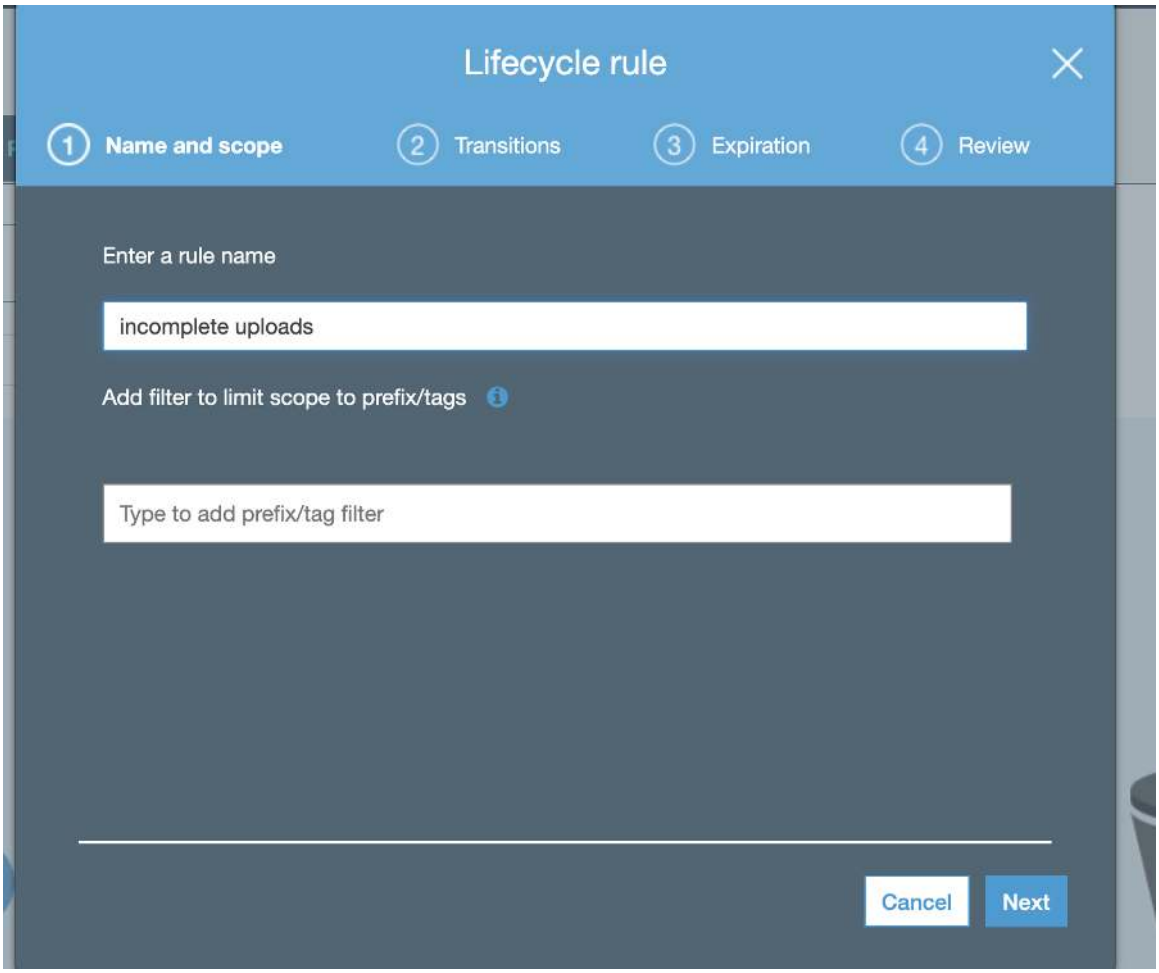

8. Leave defaults on this screen:

9.

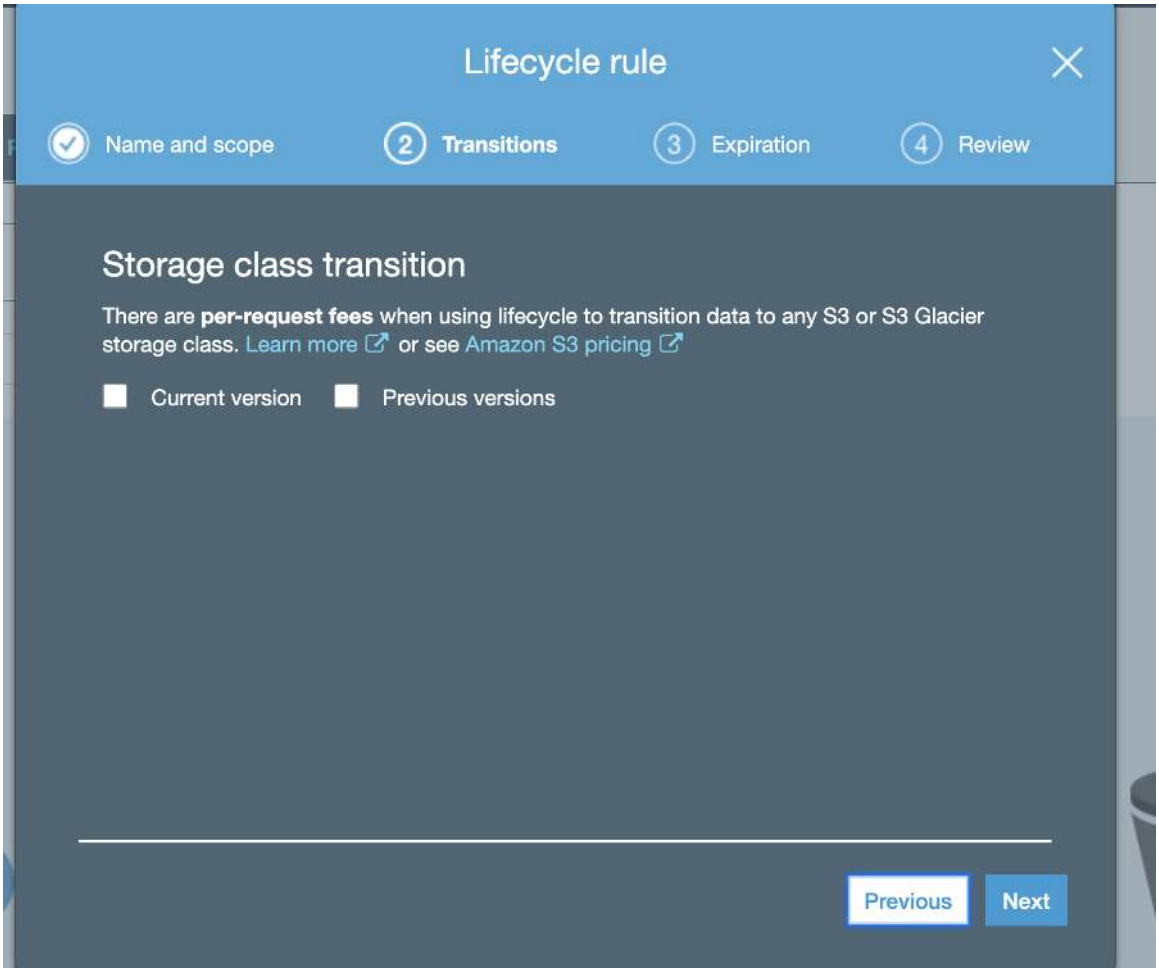

10. Configure as per the screenshot below:

11.

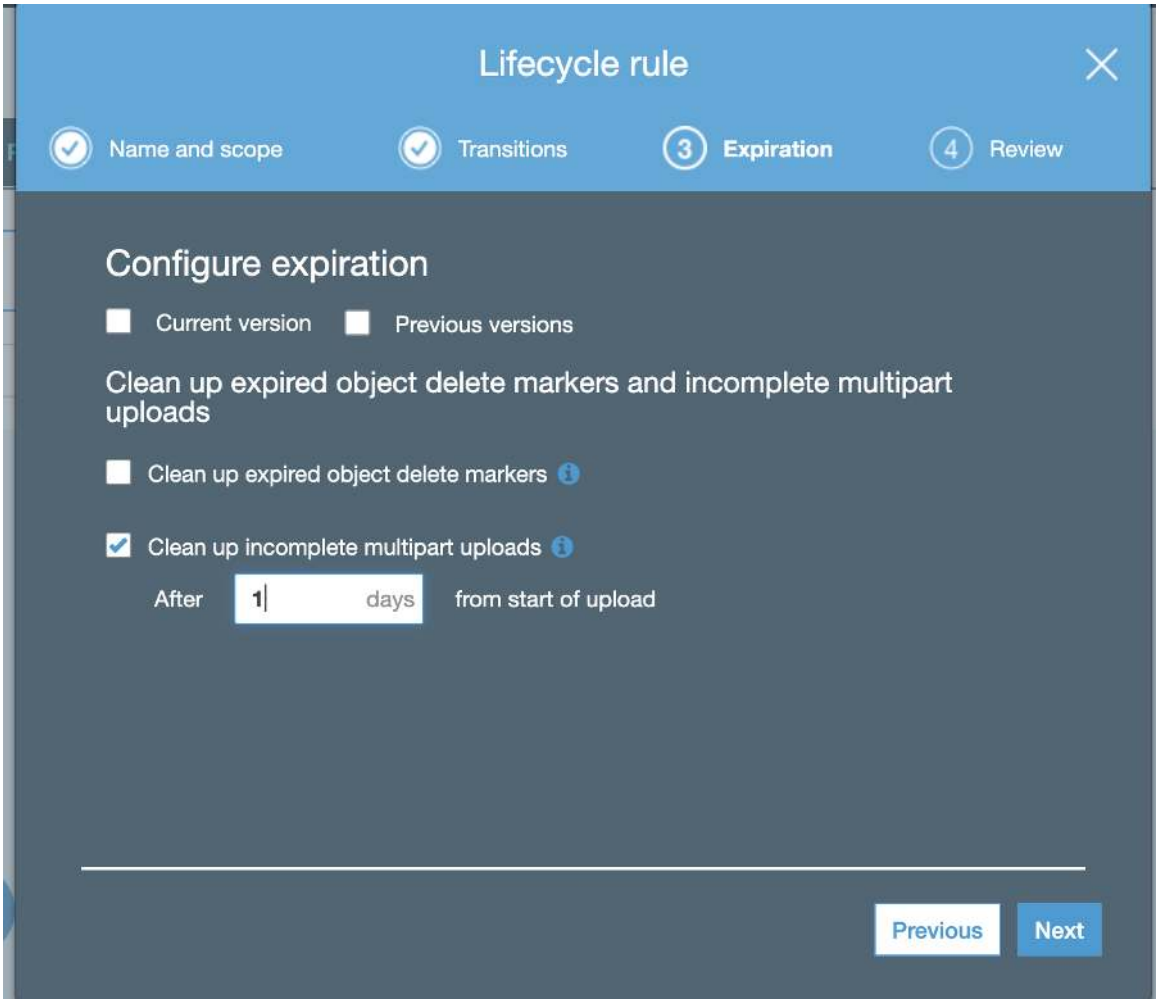

12. click "Next" and save the rule.

13. Done.

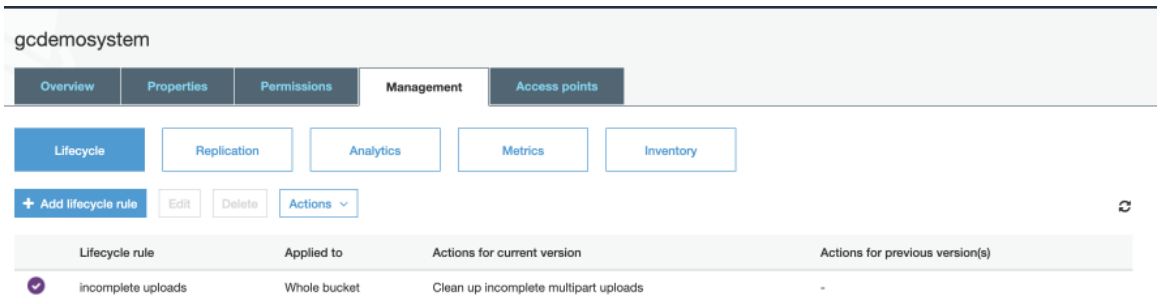

14.

1. Glacier is an S3 service tier that offers long term long cost archive. Once an object has been moved to Glacier it will no longer be accessible from an S3 browser or by Golden Copy for restore until the inflate process has been completed.
2. The steps below explain how to edit a single files storage class and how to create life cycle policy to move all uploaded data to Glacier.

How to apply single File storage class change for testing Glacier

1. Login to the Amazon S3 portal and click on a file in the storage bucket and select the "Properties" tab.
2. 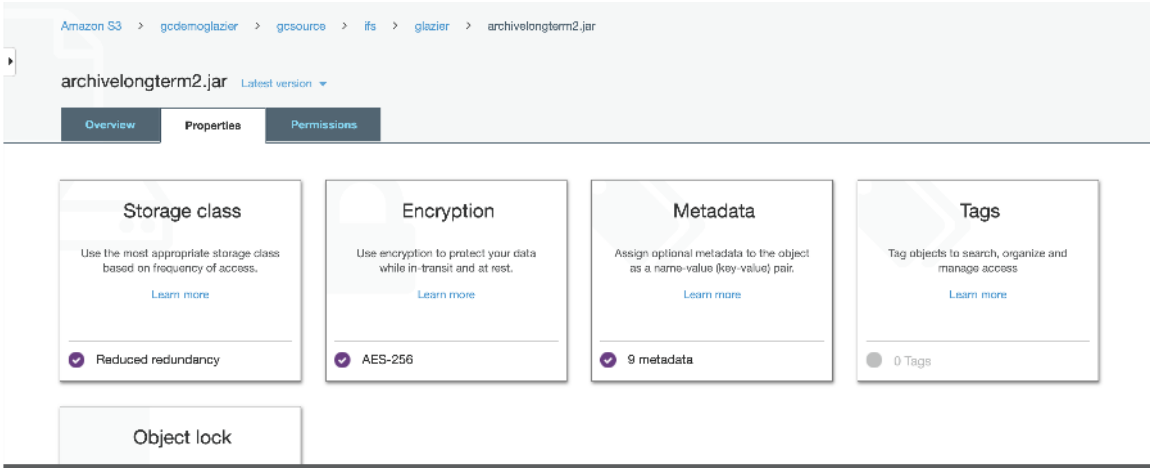

3. click on the "Storage class" option, select "Glacier" and save the change.
4. 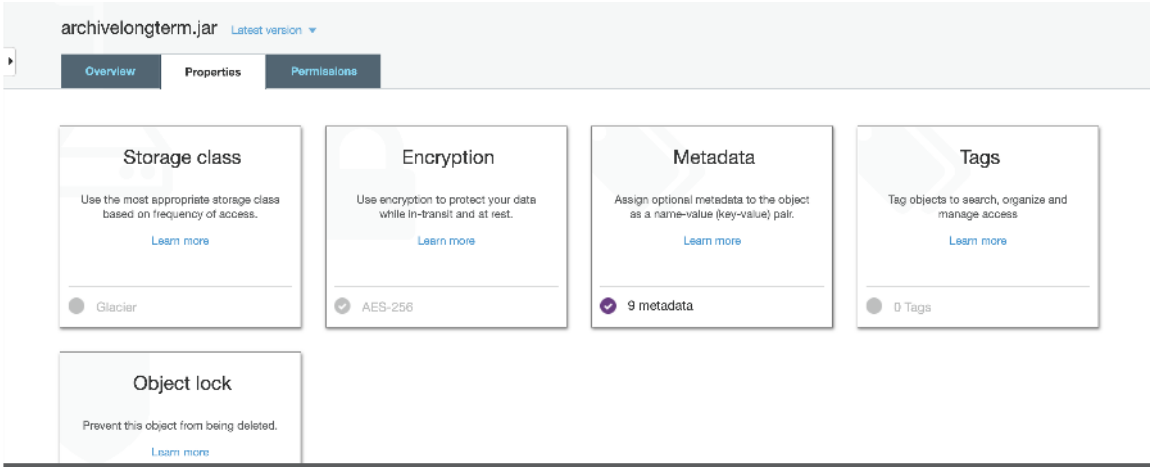

5. Note: after this change the file will no longer be downloadable from an S3 browser or Golden Copy.

6. 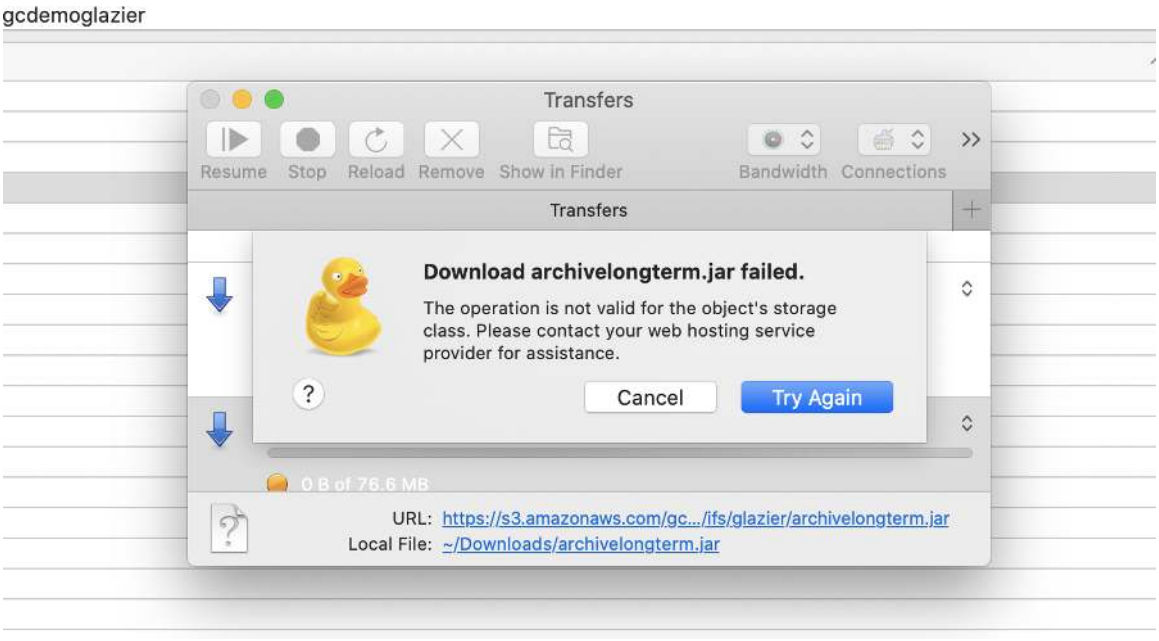

How To recall a file from Glacier:

1. Select the file in the Amazon S3 bucket and select the check box and click "Restore".

2. 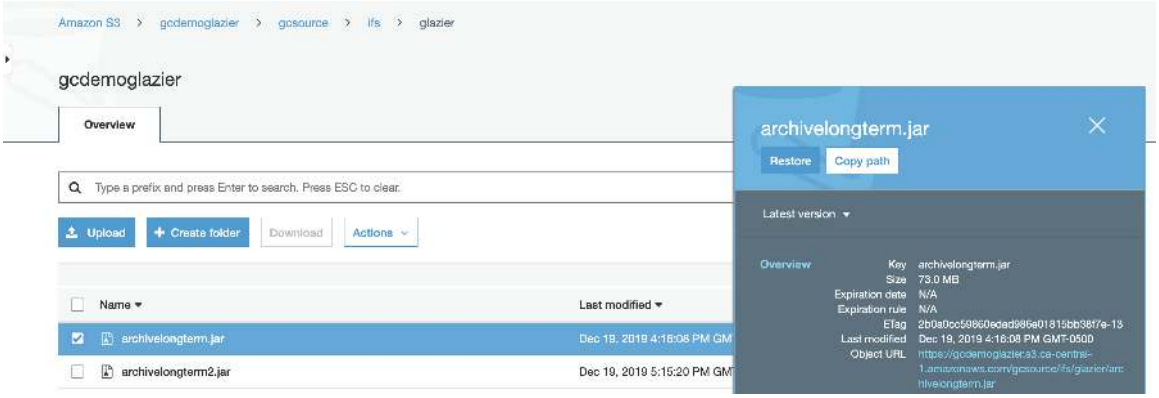

3.

   a. Note: Bulk operations or api restore is possible and you should consult AWS documentation.

4. Complete the restore, restore speed options and days the restored file will be available for access to submit the request.

5. Once the restore has been completed, the property on the file is updated to show how long the file can be downloaded.

6. 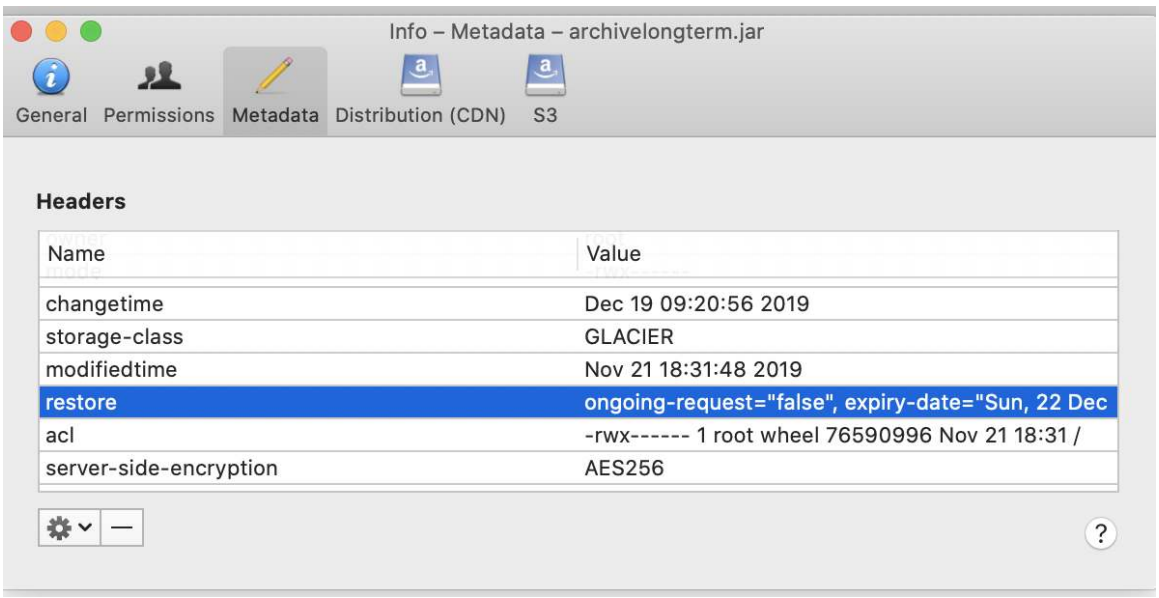

7. Done.

How to configure Glacier Lifecycle policies on a storage bucket

1. Select the bucket Management tab.
2. Click "Add Lifecyle Rule".
3. Enter a rule name.
4. 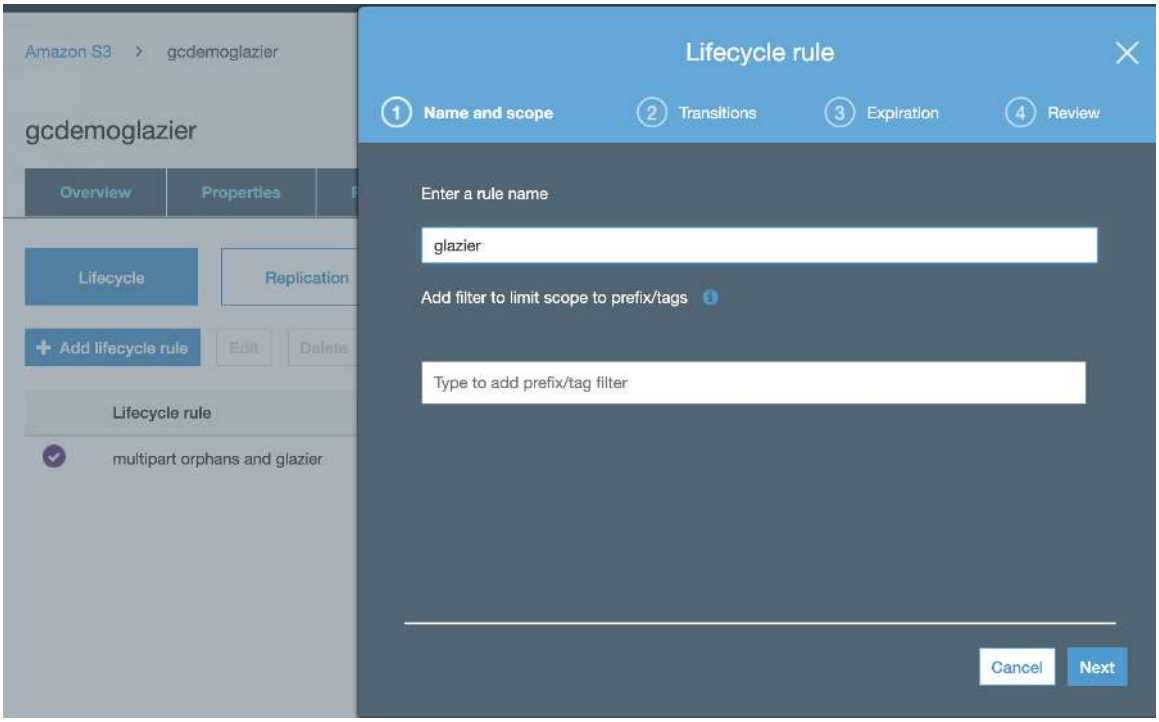

5. Select "Current version" or versions check box depending on your version setting on the bucket.

6. Click "Object creation" rule drop down and select Glacier

7. 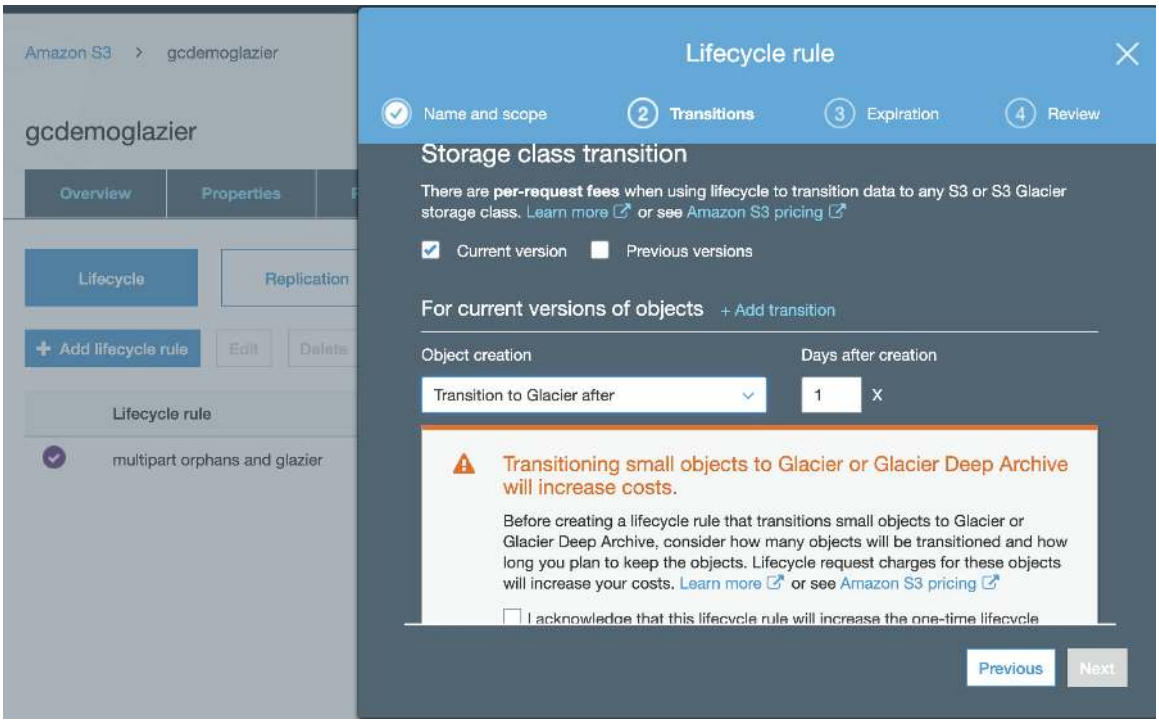

8. Click "Next".

9. Optional but recommended enable clean up incomplete mulipart uploads and set to 1 day

10. 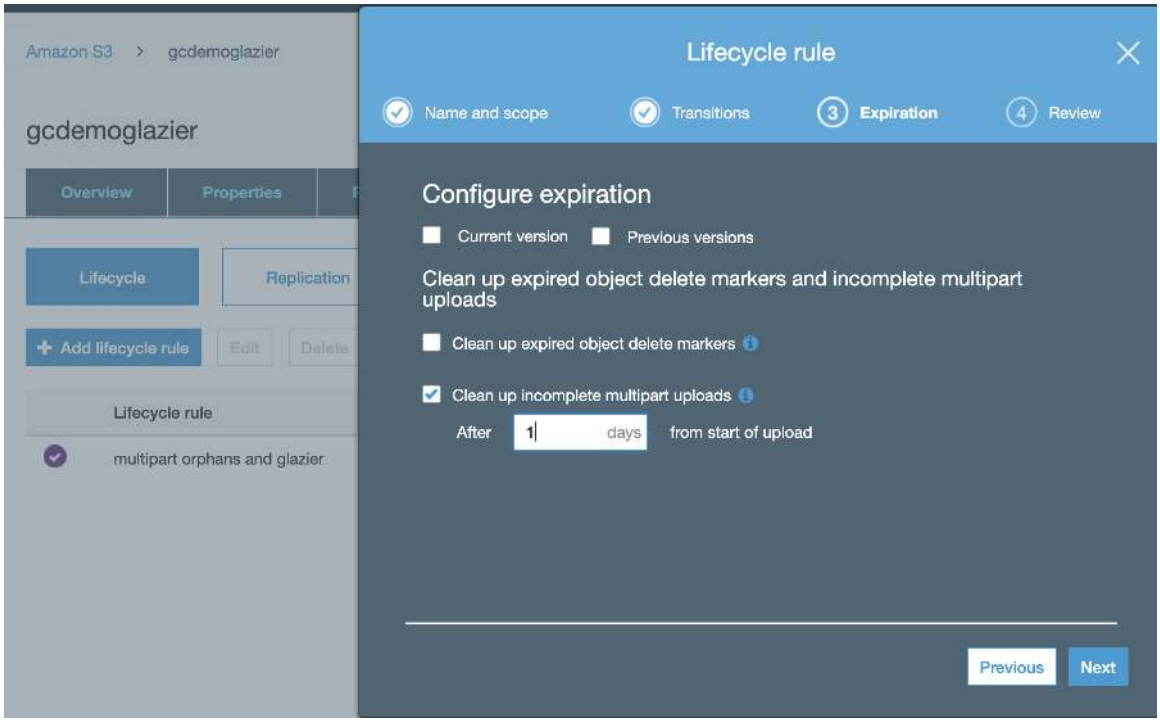

11. Click "Next" and save the policy.
12. Note all newly created objects will move to Glacier tier 1 day after creation.
13. Done.

How to Configure Storage Tier Lifecycle Policies with Azure Blob Storage

1. Centralized Azure lifecyle managment allows simple policies to move large numbers of objects with policies.
2. This simple guide shows how to create a storage policies to move objects based on last accessed or modified dates on the objects to determine when to move the object to a lower cost storage tier.
3. https://docs.microsoft.com/en-us/azure/storage/blobs/storage-lifecycle-management-concepts?tabs=azure-portal#azure-portal-list-view

© Superna LLC

10.13. Trouble Shooting File Copies

Home Top

- Logs to provide to support
- Get Errors detected during an Archive Job

Logs to provide to support

1. The error log records any file that could not be copied during a copy task.
    a. cat /opt/superna/var/logs/archiveworker/error.log
    b. Attach to support case.

Get Errors detected during an Archive Job

1. Extract errors from a copy job with this command
    a. searchctl archivedfolders errors --jobId xxxx --count 20 --tail (get the copy jobid with searchctl jobs running, or searchctl jobs history)

© Superna LLC

10.14. Golden Copy Solutions Guides

Home Top

Overview

The guides below are focused on specific use cases with Golden Copy.

- DR Solution for Azure
- Application Disaster Recovery to Amazon S3 and FSX Windows Service
- Search & Recover & Golden Copy Archiving Solution
- Immutable Storage for Azure Blob Storage
- Data LifeCycle - Cost Reduce with Archive to the Cloud
- Object Backup and Basic DR Solution
- Bulk Loading Data with Azure Data Box
- Bulk Recall of Data from AWS and Azure
- Bulk Loading Data with AWS Snowball

© Superna LLC

10.14.1. DR Solution for Azure

Home Top

- Overview
- Support Statement
- Requirements
- High Level Steps
  - Golden Copy Steps
  - Azure DR Application & Blob to File Share Configuration Steps
- DR Recovery of Blob Data to SMB shares in Azure Procedures
- How to connect remote Computers to the Azure Cloud Share
- How to Sync Changed File Data back to Blob Container for Recall

Overview

This solution explains how Golden Copy Blob sync mode can be used for application specific DR recovery in Azure. The solution uses blob storage and File Shares in Azure Storage accounts to allow recovery of one or more applications in Azure. This solution will cover the high level steps needed to get the data prepared for an application server running in Azure. There are many aspects to accessing and securing data within Azure that are not covered in this guide. Consult with Azure documentation on aspects of running and accessing

applications within Azure for your requirements. This solution documents the SMB share solution. For a more seamless integration Azure blob allows NFS mount of blob storage for native access to object data in place. See the guide here that describes NFS support for blob data.

Support Statement

1. This guide is provided "as is" and Azure steps are examples only, support contract does not include support for Azure. Customers must have expertise in Azure to complete steps not covered in this guide.

Requirements

1. Azure subscription.
2. PowerScale Onefs 8.1.x or 8.2.x.
3. Golden Copy Installed and deployed
4. All Azure resources in the example should be created within the same resource group.
5. Active Directory joined from on premise to Cloud instance of Active Directory hosted by Azure. This guide does not cover these steps, consult with Azure documentation on how to complete this configuration and secure access to the SMB shares in Azure.

High Level Steps

Golden Copy Steps

These steps explain how to protect a single application's data stored on PowerScale with the goal of providing a DR recover option for the application in Azure.

1. Create Azure Storage Account and Blob Container to store the application data.
    a. These steps are covered here.
    b. In this example the storage account created is **DRApp1,** the storage container name is **application1,** and then follow the steps to get the **secret key** needed for authentication**.**
2. Install Golden Copy and configure sync mode on an application folder to sync the path on the PowerScale where the application data is stored and is required for DR recovery.
    a. Azure example in the configuration guide to add a folder.
    b. Example for Azure to add folder command in Sync Mode (default mode when adding a folder). This example assumes the application data is all stored under the path /ifs/data/applicationdata on the PowerScale cluster named gcsource:
        i. searchctl archivedfolders add --PowerScale gcsource --folder /ifs/data/applicationdata --secretkey NdDKoJffEs9U9Xg== --endpoint blob.core.windows.net --container application1 --bucket **DRApp1** --cloudtype azure

c. Start the Sync job full sync to copy the application data to the blob storage in Azure.

    i. searchctl archivedfolders archive --id (**use searchctl archivedfolders list to get the folder ID for /ifs/data/applicationdata**)

3. Now move on to Azure DR steps.

Azure DR Application & Blob to File Share Configuration Steps

1. In the Azure portal the following tasks and resources need to exist:

    a. Windows server VM for Blob to file share conversion (size of the VM does not matter).

    b. Application 1 VM(s) required for the application to execute in Azure.

        i. **NOTE: The scope of this is outside this document's intended purpose. It may require multiple VM's, a resource familiar with Azure services and application server clone. Migration to Azure should be consulted.**

2. Azure Windows server blob conversion VM configuration:

    a. Verify blob container files are synced. You can view the blob files using the Azure portal. See example viewing the storage account and then the container contents (in this example the container name is "gc1")

i. 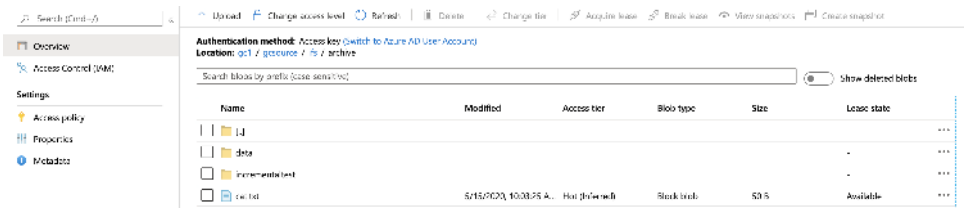

b. Create Storage Account File share by clicking on the storage account, then the File Shares and then click + File Share. In the "New file share" window enter a Name for the application share (i.e Example below "application1"):

i. 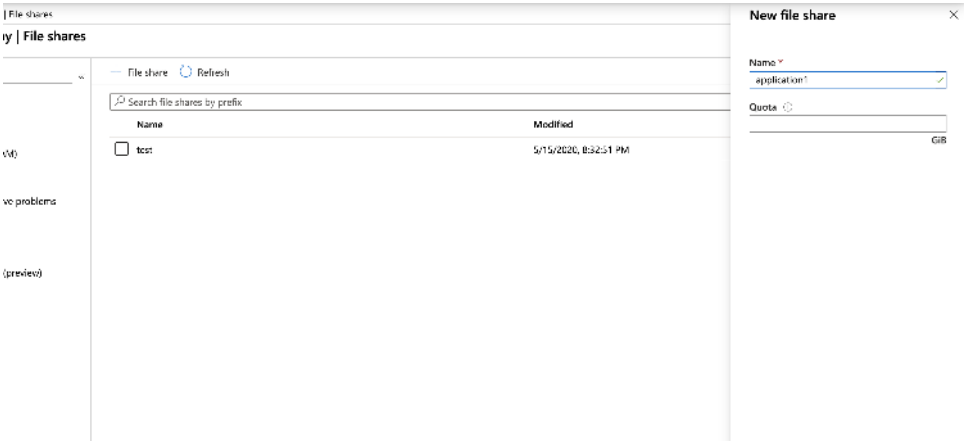

ii.

iii. Now copy the connect code for a Windows VM clicking on the 3 dots, and then "connect" option. Select the drive letter defaults:

iv. 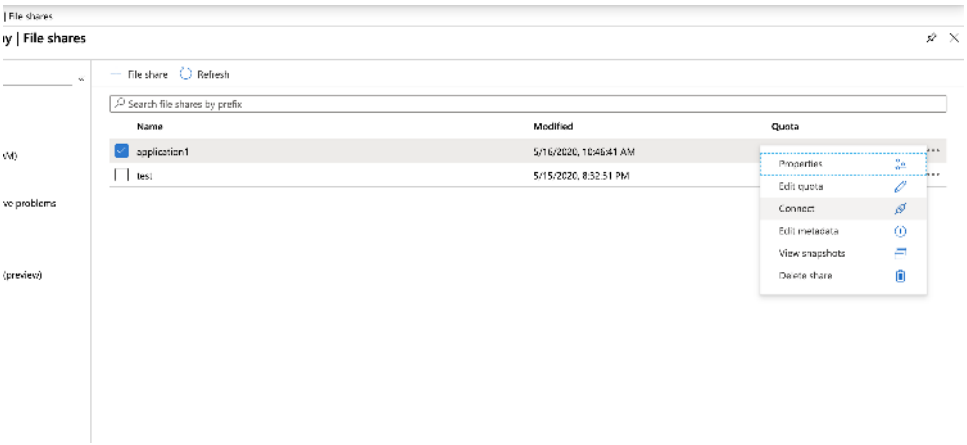

v. Copy the connect code for powershell

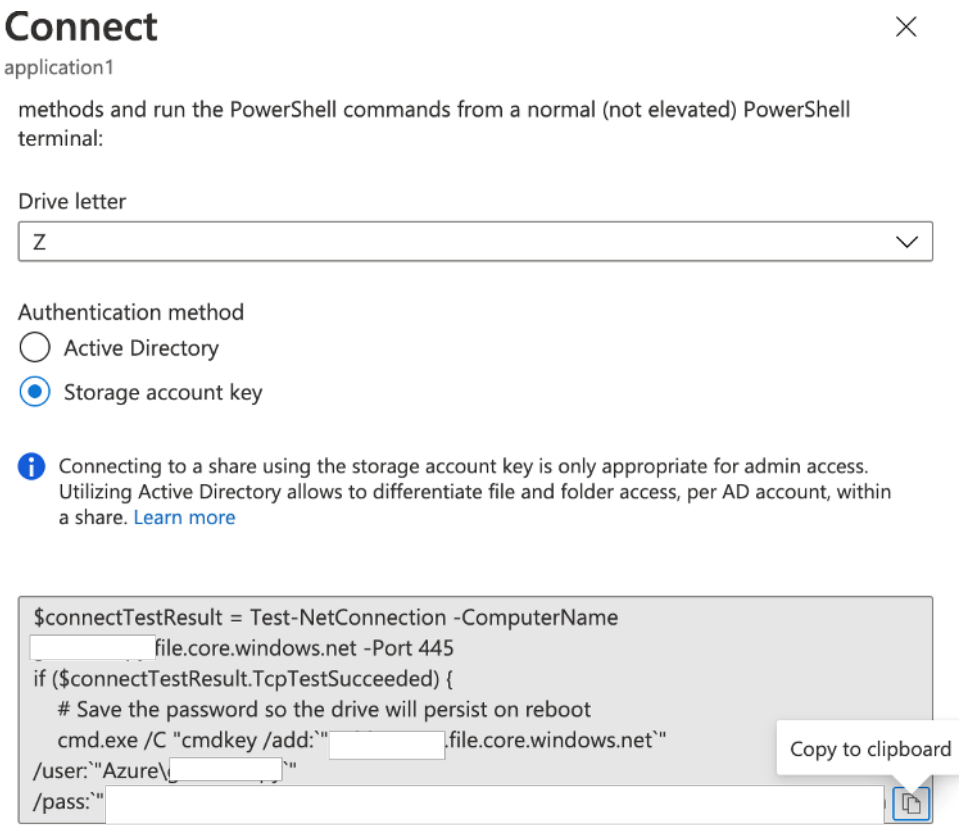

vi. This script will check to see if this storage account is accessible via TCP port 445, which is the

c. Login to the Windows Conversion Server VM created above. (NOTE the VM should be in the same resource group as the Storage Account).

   i. Open a PowerShell prompt and paste the connection code to mount the "application1" share.

   ii. Verify in Windows explorer the share can be opened and create a test file:

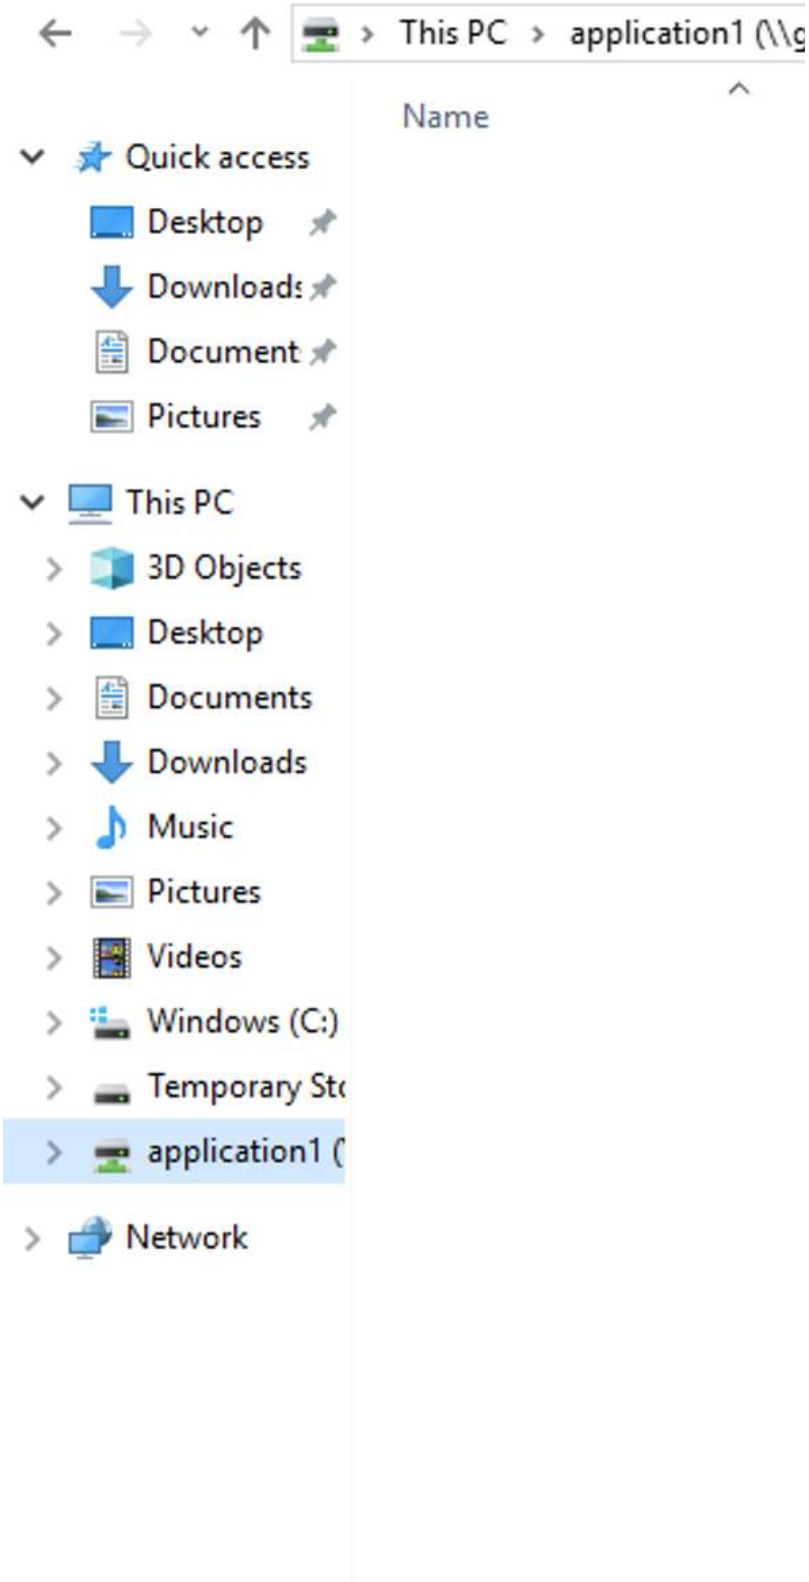

iii.

iv. Done.

d. Install Azure Storage Explorer tools into the Conversion VM:

i. The installer can be found here https://azure.microsoft.com/en-us/features/storage-explorer/ .

ii. Open Storage Explorer, click the User Icon and login to your Azure portal administrator account. This is the subscription option to provide access to all resources in your Azure account.

iii. You should now see all storage accounts created within your Azure account, and the top level of the files uploaded by Golden Copy. The cluster name is the first folder visible in the Blob storage container.

iv.

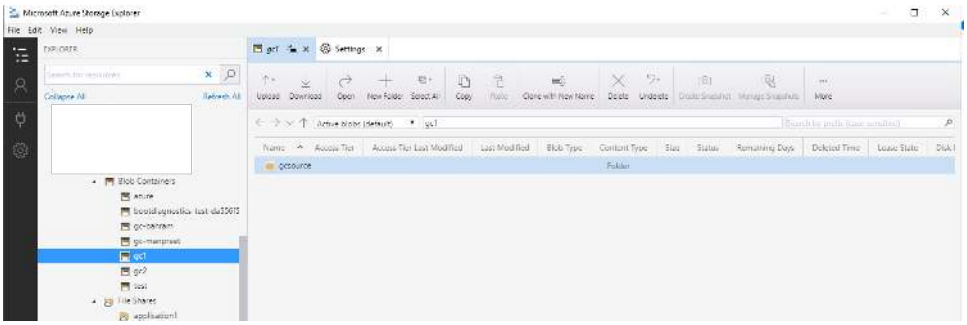

e. Infrastructure Configuration Completed, you are now ready for DR.

f. Done

DR Recovery of Blob Data to SMB shares in Azure Procedures

1. In a DR scenario, data must be copied from blob to the file share created above.

2. The process to move data from Blob container to File share can be completed with the Azure Storage Explorer tool. This step can take time to complete depending on the quantity of data.
3. Open Azure Storage Explorer.
    a. Browse to the storage account name and expand to select the container name. Navigate on the right hand pane to the "application1" folder data. Then select the "application1" folder and click the "Download: button. See example below:
    b. 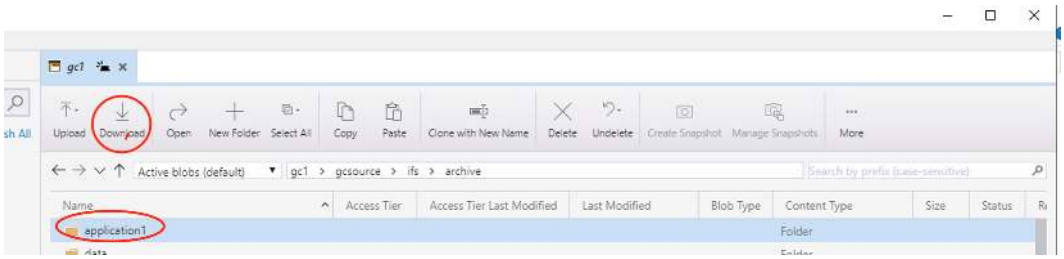

    c. Select the Mapped drive letter (Z: in this example) to copy the blob data to the File share "application1" created for the Application.
    d. The download job is queued and visible at the bottom of the UI and will show progress and completion of the copy:
        i. 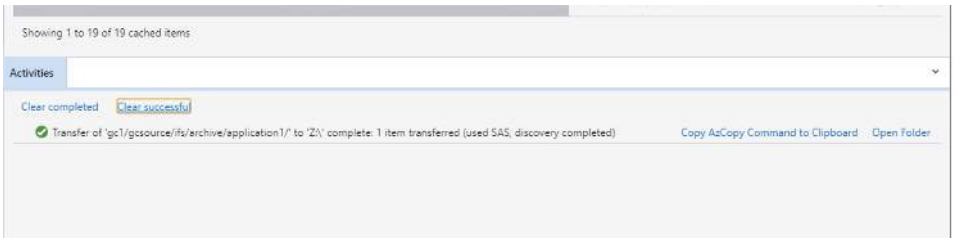

    e. Once the copy is completed, or during the long transfer time, view the Z: mapped share to view the files that are appearing in the File share.

i. 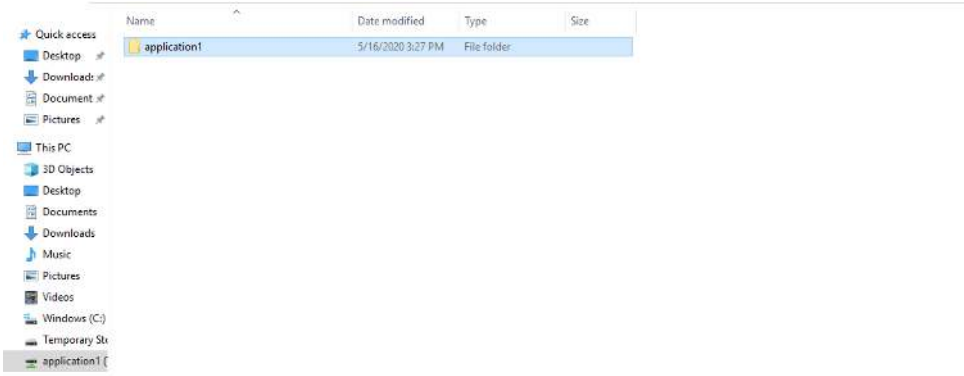

f. Now that the Storage is copied to storage account file share the storage recover steps are completed. Any application mounting and testing steps now require application VM's to mount the "application1" share to gain access to data in the File share.

g. This completes the example storage failover solution.

How to connect remote Computers to the Azure Cloud Share

1. An administrator can use the same powershell command to connect from a remote location to the Azure file share. This allows administrators quick easy access to the data remotely. The powershell connection uses SMB3 encryption so the data inflight is encrypted.
2. User Connections to Azure File Shares.
    a. This should be done using Integrated Active Directory authentication.

b. This guide provides setup, overview of how to integrate on premise AD to with Azure AD service to secure shares and ACL's on the shares.

c. Microsoft Azure guide on AD authentication

How to Sync Changed File Data back to Blob Container for Recall

1. Overview:

    a. Data that is added or changed after the DR event will need to have this data synced back to the Blob container before it can be synced back to Isilon using Golden Copy Recall.

2. This will be completed Using Storage Explorer Tool installed on the management VM in Azure.

3. Select the blob container name in storage explorer on the left had menu, then click the upload button. **NOTE: The administration VM must have a mapped drive letter to the File share in place.**

    a. Select the path on the drive letter mapped to the file share that needs to be synced back to the blob container. NOTE: partial data sync is possible by selecting only the data you need to sync back. See example below.

b. 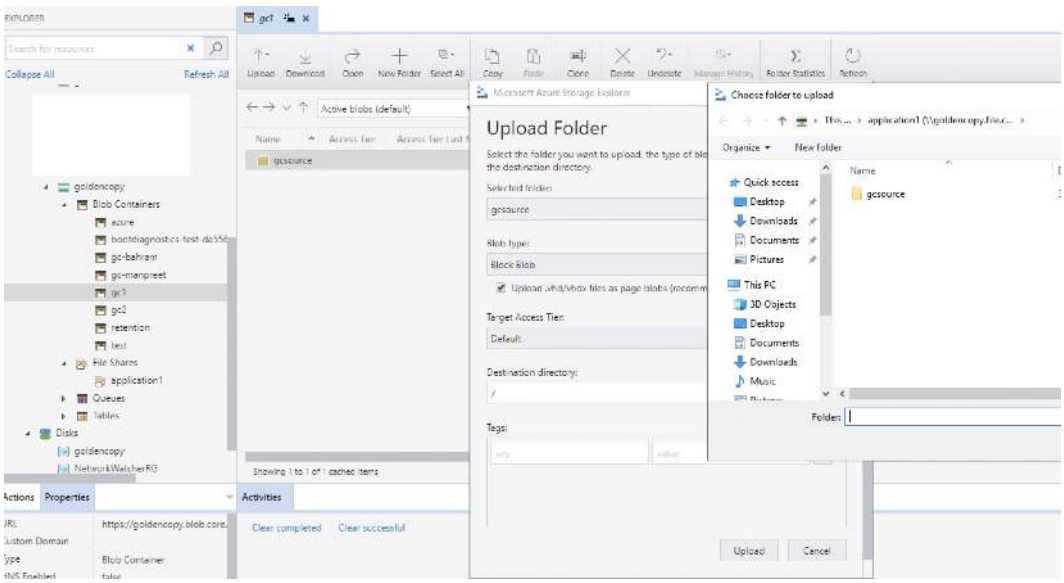

c. Accept the overwrite options to replace data on the target blob container.

d. 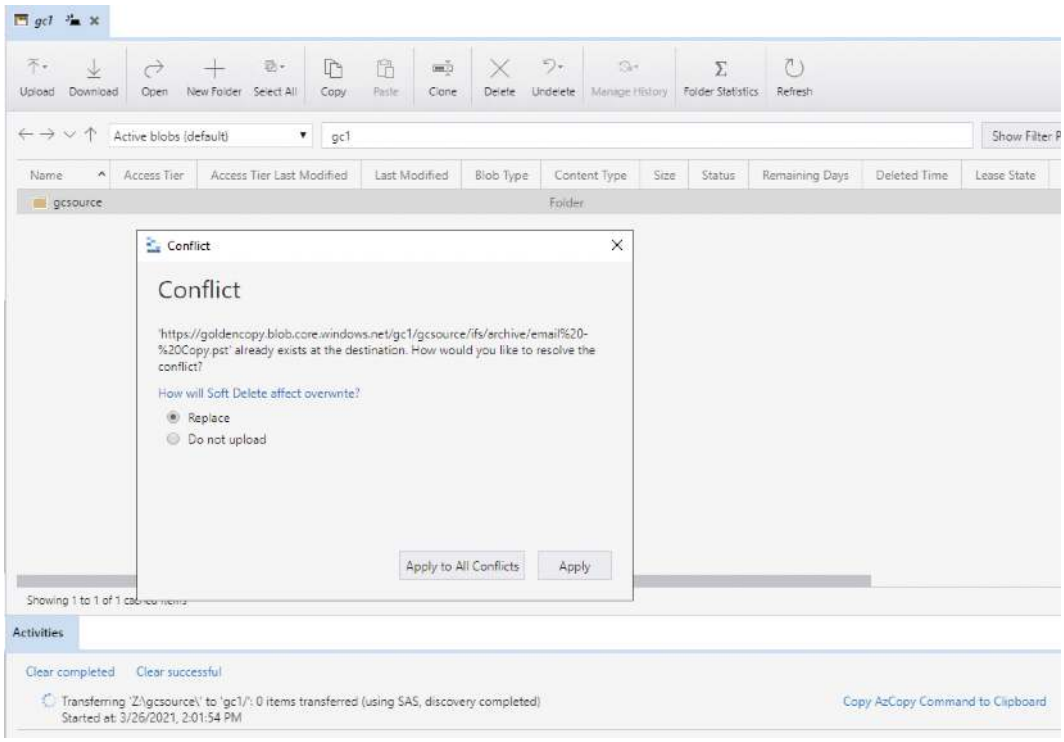

e. Wait for the copy to complete.

f. Any changed data from Azure will be copied back into the Blob storage. NOTE: overwriting data will remove the

metadata in the container which means metadata will no longer be restored by Golden copy.

4. To recall the data with Golden Copy see this guide link here to recall the container blob data back onto your cluster.

© Superna LLC

10.14.2. Application Disaster Recovery to Amazon S3 and FSX Windows Service

Home Top

- Overview
- Requirements
- High Level Steps
  - Golden Copy Steps
- AWS DR Application & S3 to FSx SMB Share Configuration Steps
  - Active Directory Deployment Options in AWS
  - Create FSx File system for the Application Server Data
  - Application Server AMI in EC2
  - AWS DataSync Service configuration
  - Create a DataSync Task to sync S3 to FSx Filesystem Share
    - Test DataSync Copy
  - Test Mount FSx Share from the ECA Application Server VM
- Application DR Recovery Procedures
  - How to Initiate DR Test or Actual DR of the application to AWS EC2 and FSx File System

Overview

This solution explains how Golden Copy Amazon S3 sync mode can be used for application specific DR recovery in AWS. The solution uses S3 storage and the FSx SMB Share Service in AWS to allow recovery of one or more applications in AWS. This solution will cover the high level steps needed to get the data prepared for an application server running in AWS. There are many aspects to accessing and securing data within AWS that are not covered in this guide. Consult with AWS documentation on aspects of running and accessing applications within AWS for your requirements.

Requirements

1. AWS S3 , FSx Configured Service
2. PowerScale Onefs 8.1.x or 8.2.x
3. Golden Copy
4. This example is only covering within a single availability zone

High Level Steps

1. Sync Isilon file system data to Amazon S3 for applications that need DR
2. Deploy EC2 application instances that will be used for DR. They can be left powered off.
3. Configure FSx SMB service in AWS to hold the SMB application data within AWS to be used by EC2 application VM
4. Configure Datasync service in AWS to sync S3 data to FSx SMB share, in the event of a disaster recovery scenario
5. Start EC2 application server to recovery your application in AWS.

Golden Copy Steps

These steps explain how to protect a single application's data stored on PowerScale with the goal of providing a DR recover option for the application in AWS.

1. Create Amazon S3 Bucket to store the application data.
    a. These steps are covered here.
    b. In this example the storage bucket created is **GCDEMOSYSTEM.**
    c. You will need the access and secret key to authenticate to the bucket.
2. Install Golden Copy and configure sync mode on an application folder to sync the path on the PowerScale where the application data is stored and is required for DR recovery.
    a. AWS example in the configuration guide to add a folder.
    b. Set an incremental schedule on the folder following guide here.
    c. Start the Sync job full sync to copy the application data to the blob storage in AWS.
        i. searchctl archivedfolders archive --id (**use searchctl archivedfolders list to get the folder ID for /ifs/data/applicationdata**)
    d. You can verify the data is visible in the S3 bucket from the AWS Console

i. 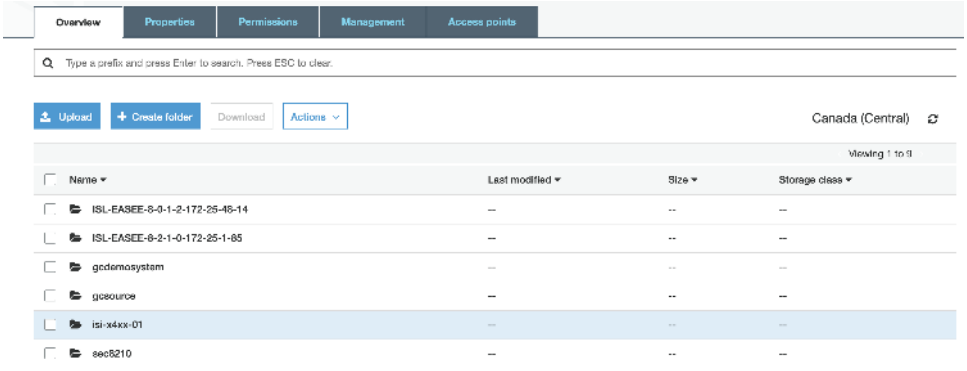

3. Now move on to AWS DR steps.

AWS DR Application & S3 to FSx

SMB Share Configuration Steps

1. Active Directory Deployment Options in AWS

    a. AWS Managed AD instance - guide here - This option requires the AWS AD to have a trust in place to the on premise Active Directory Domain to ensure on premise and cloud authentication is the same.

    b.

i. 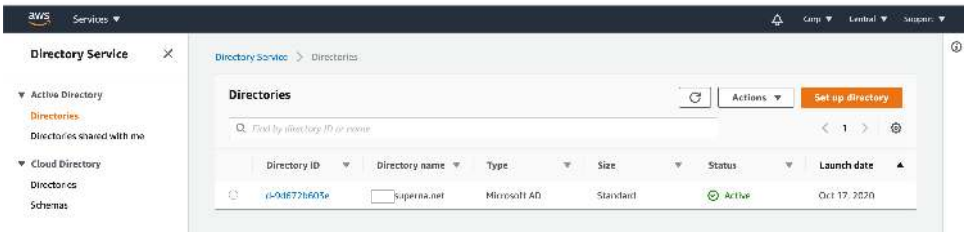

2. Create FSx File system for the Application Server Data

a. The FSx File System is a managed SMB server in AWS that can present SMB shares to applications servers. This File System can store data for multiple applications and have multiple SMB shares for each application. This example guide only presents a single SMB share for 1 application.

b. Make sure to create the Windows FSx File System large enough to store the application data. The FSx file system must also be connected to Active Directory. The disk can be extended to add space from the AWS console if required.

c. Follow the FSx guide to configure a new file system.

d. 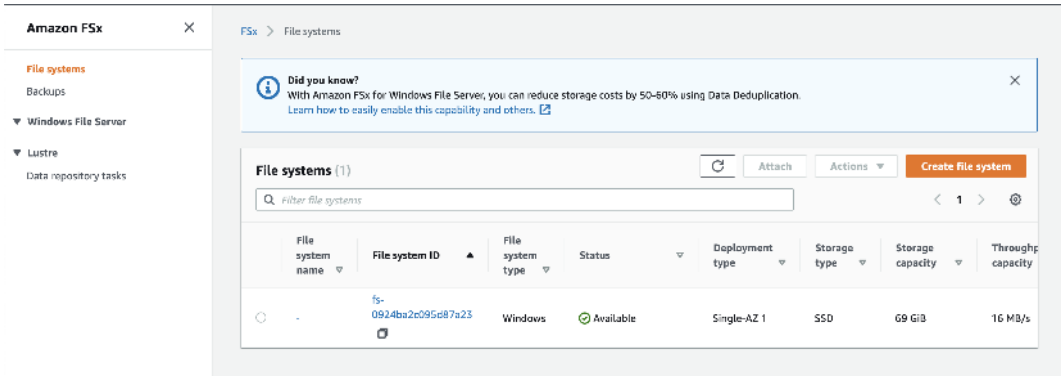

e. Retrieve the FSx File System DNS end point for mounting Shares from Application servers or Management Servers

   i. Login to the FSx management page in the AWS console and record the DNS endpoint, ip address and power shell endpoint needed to mount SMB shares within AWS VPC. See example below.

ii. 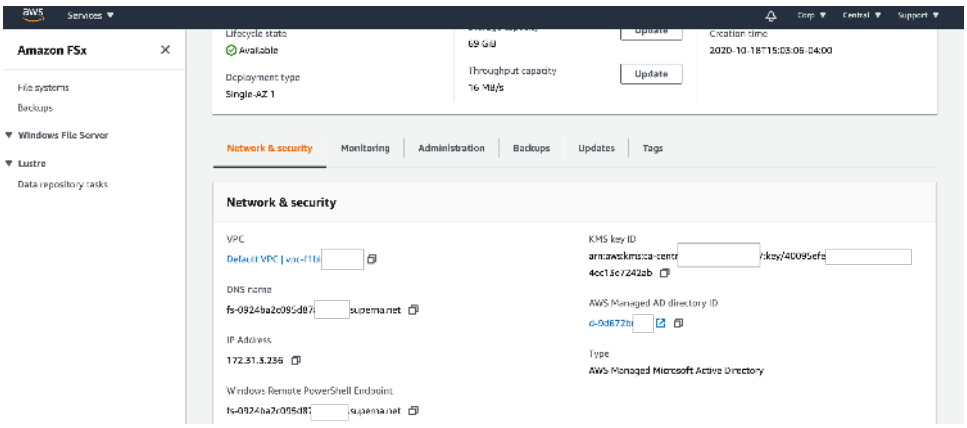

f. Create or Manage SMB shares in FSx File Systems

   i. To create or manage SMB shares in an FSX File System requires a management server in AWS to ensure access to manage the file system is available in a disaster.

   ii. Best Practice: Create a management server to connect to FSX File Systems.

   iii. Follow the guide to use power shell or Windows Management server GUI tool

   iv. RDP to the management Windows server logged in as domain admin user that has permissions to the File System FSX shares. Consult FSx documentation on the Managed AD domain group that allows management of FSx File Systems.

      1. open fsmgmt.msc tool to create or modify SMB shares and set permissions

      2. 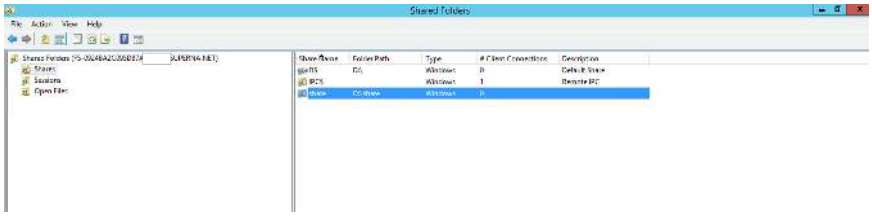

3. Application Server AMI in EC2

a. Application virtual machines to mount the SMB shares. Your application server should be pre-staged in AWS EC2 and joined to Active Directory so that correct SMB share permissions will allow the application to mount the data.

   i. **NOTE: The scope of this is outside this documents intended purpose, and may require multiple VM's and a resource familiar with AWS services and application server clone and migration to AWS should be consulted.**

4. AWS DataSync Service configuration

a. This service copies data from S3 AWS buckets and a specific path of data within the bucket to a destination FSx SMB share. This copy configuration can be preconfigured in advance of a DR event. The steps below cover some of the key steps needed to setup and configure the Datasync service. To learn more about DataSync see the documentation page here.

b. Deploy the Datasync Agent following this guide. The Datasync EC2 instance should be deployed in the same region as your storage bucket and FSx file system.

c. The next step involves creating an endpoint. This is documented here. Since the DataSync service is only copying data within AWS choose the VPC endpoint option.

d. Continue to activate the agent following this procedure

5. Create a DataSync Task to sync S3 to FSx Filesystem Share

   a. The documentation can be found here. A location is required for FSx and S3 bucket. Using the links below create a location in DataSync to authenticate to your FSX file share for your application, this requires an AD user with permissions to the FSx filesystem share.

   b. Creating a Location for Amazon FSx for Windows File Server

   c. Creating a Location for Amazon S3

   d. Create the task that defines the source (s3) to target FSX following steps here

6. Test DataSync Copy

   a. Perform a test copy to verify the task is setup correctly. Once the test is complete you can delete the data in the share. This requires a management VM running Windows to mount the share and verify the data was copied and also to delete the data.

   b. This is done by selecting the task and starting the job.

   c. 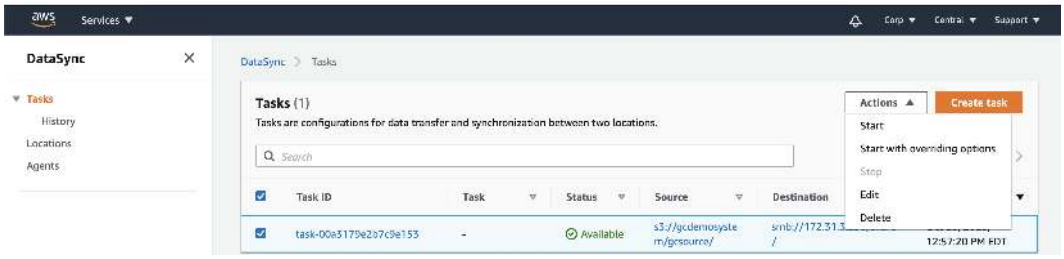

d. 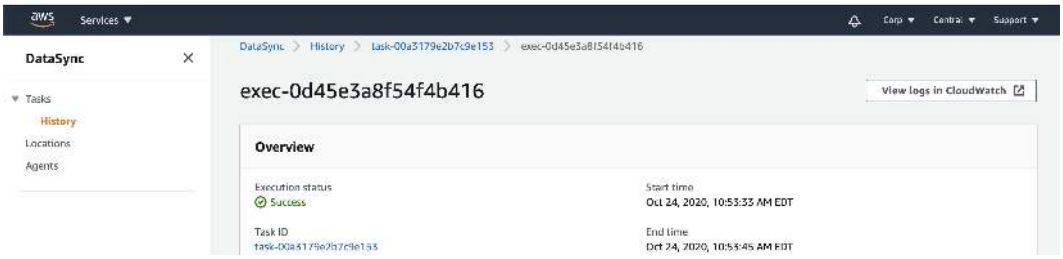

7. Test Mount FSx Share from the ECA Application Server VM

   a. From a the Windows application server in EC2 mount the SMB Share Configured for this application server to verify the DataSync Job has copied the data successfully. See example below.

   b.

      i. 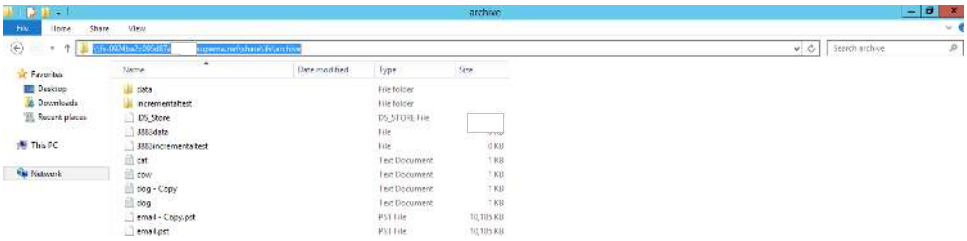

   c.

Application DR Recovery Procedures

1. **Overview:** The process would be used in the event of a disaster or if you plan to exercise a DR test. The high level steps below would be used to move an application to operate from within EC2 AWS VPC. These steps can take time to complete depending on the quantity of data in Amazon S3 that needs to be copied into the FSx file System.

2. How to Initiate DR Test or Actual DR of the application to AWS EC2 and FSx File System

a. Open AWS DataSync console
b. Start the application datasync task to sync data into the FSx Share location target. This process can take time depending on how much data needs to be copied.
c. Monitor the DataSync task's copy performance and completion status on the console.
d. Proceed to startup up the EC2 instance of the application server
e. Login to the application server over RDP and verify the SMB mount. Data should already appear in the mount as the copy progresses.
f. Startup the application software once the DataSync task has completed.
g. DR test completed.

© Superna LLC

10.14.3. Search & Recover & Golden Copy Archiving Solution

Home Top

- Solution Overview - How to use an Archive Staging Area With Search & Recover and Golden Copy
- Deprecation Notice
  - How to locate aging data for Archive With Search and Recover
  - How to move the files to the Archive Staging Area
  - How Recall data for Users
- Solution Overview - How to archive data in place with Search & Recover and Golden Copy
  - Requirements
  - How to locate aging data for Archive With Search and Recover
  - Copy Search results file to Golden Copy for Archive
  - Delete the Data Using Search & Recover
  - Recalling Archived Data

Solution Overview - How to use an Archive Staging Area With Search & Recover and Golden Copy

This work flow requires both products to use this work flow to provide an archive solution with S3 or Azure blob storage. This work flow would require users to request access to data that has been archived

and a recall job created to return data to the cluster. This solution is a bulk archive solution aimed at doing large clean up every few months.

Deprecation Notice

1. As of 1.1.4 the search & recover csv format will be deprecated with a 1.1.4 build to support a format of CSV import that uses only an absolute path of /ifs/... to a file to be copied. This will allow customer scripts or command builder feature in Search & Recover to generate a simple csv or flat file with only path to files to be copied by Golden copy.
2. Required 1.1.4 build > 21124

How to locate aging data for Archive With Search and Recover

1. Typically files age based on modified or last access are common criteria to locate data
2. Login to search
3. Open Quick reports and What's Growing old?

4. 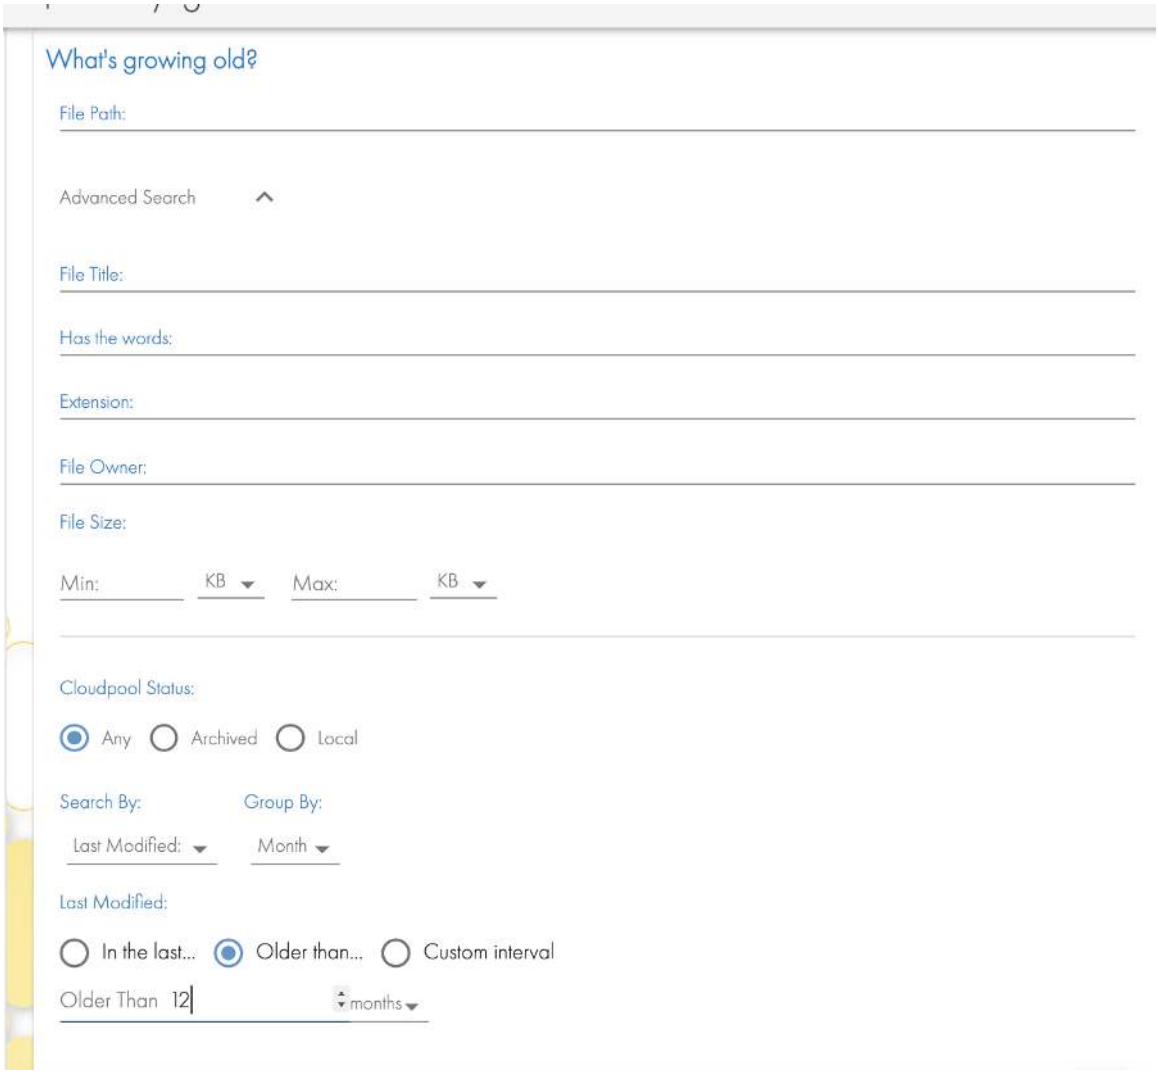

5. Move the results to Archive Staging Area.
    a. Create /ifs/archivestaging
    b. Use this Guide to move data from the search result list to archive staging area. Use the target path above to replace the examples used here
6. Repeat the above steps when you want to clean up data and move to an S3 bucket. The example script copy above has the option to move data to the staging area. This solution will also retain the original path of the file under the staging folder base path. example /ifs/archivestaging/ifs/data/moveddata

How to move the files to the Archive Staging Area

1. On the Golden Copy VM at the /ifs/archivestaging to an archived folder definition. See quick start examples here.
2. Configure the folder for copy mode
3. Run the copy job
4. Verify all data was successfully copied by viewing the job report or the running job.
    a. searchctl job view --id <job name > --follow
    b. searchctl archivedfolders export --id <job name >
5. Delete the data under /ifs/archivestaging

How Recall data for Users

1. If a user request to recall data comes is received you can view the S3 bucket with Any S3 browser tool to locate the path of the files.
2. Start a recall job on Golden Copy vm to recall the data back into the staging area. Example
    a. searchctl archivedfolders recall --id <folder id> --subdir /ifs/archivestaging/ifs/path-to-user-data
3. Once the recall job finishes you can create an SMB share on top of the recalled data to provide access to the recalled data.
4. Summary:
    a. This keeps the data in the original location that the user last had access to it

b. The staging area moves data into an area that is easily copied and managed seperate from the file system.
c. Users must request data back and this allows monitoring of data access requests.

Solution Overview - How to archive data in place with Search & Recover and Golden Copy

This solution keeps the data in place , meaning it does not use a staging area to copy or move data to the archive. Data can be recalled back using Golden Copy recall feature. NOTE: Recall will not place the data into the same location it was archived from. The data will be recalled to the staging area but the relative file path will be preserved.

Requirements

1. Search & Recover
2. Golden Copy
3. S3 Target Storage

How to locate aging data for Archive With Search and Recover

1. Typically files age based on modified or last access are common criteria to locate data
2. Login to search
3. Open Quick reports and What's Growing old?
4. 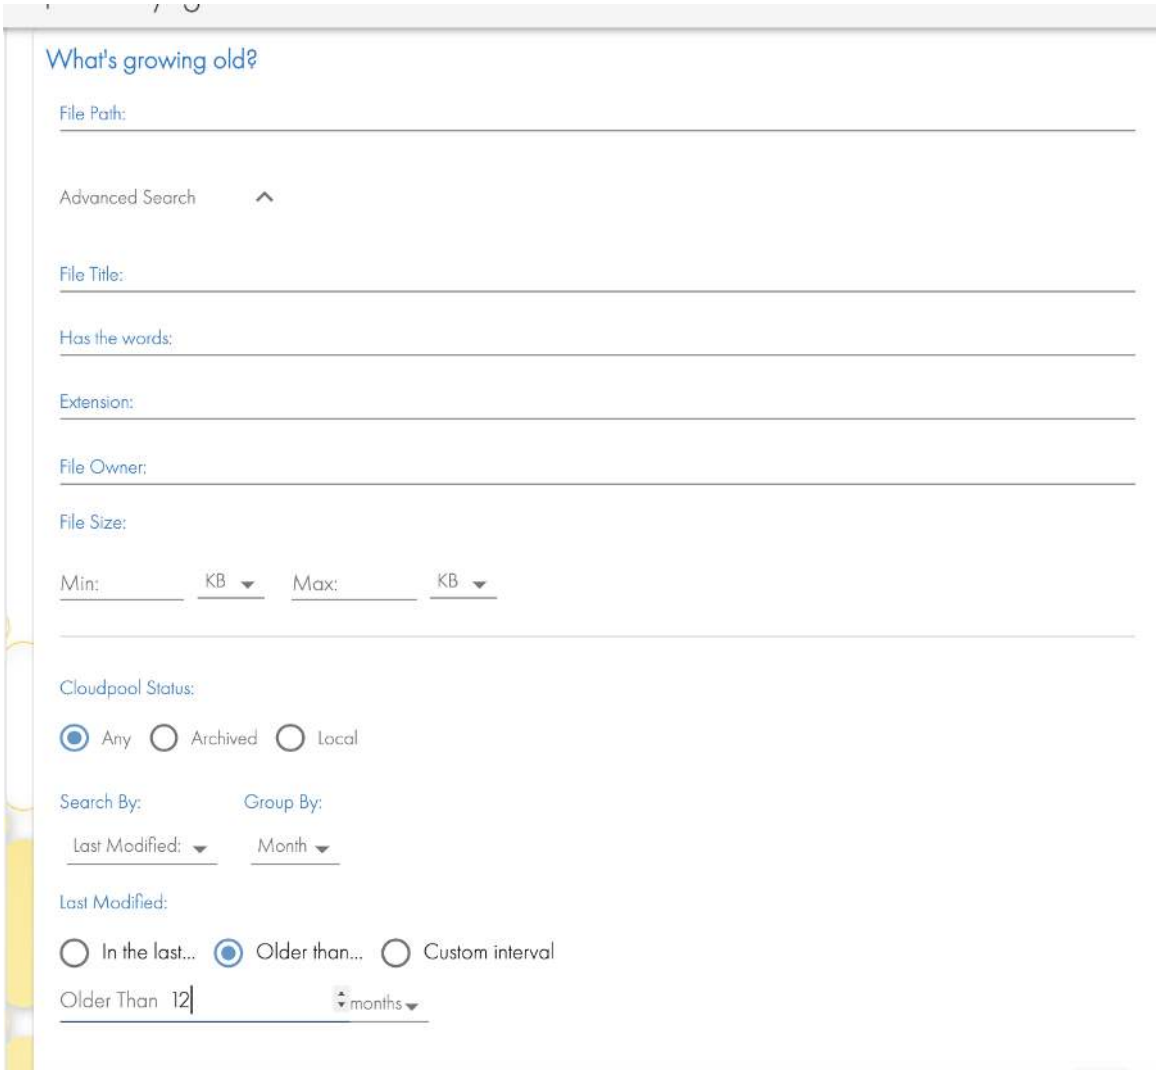

5. The resulting table will show the month and year with data that meets the age search criteria. NOTE: You may enter a path to narrow the search to a specific path in the file system as well as

using other search criteria in the Advanced window of the What's Growing Old? Quick Report.

6. Select the months or rows in the table and select the menu option "Show me the files"

   a. 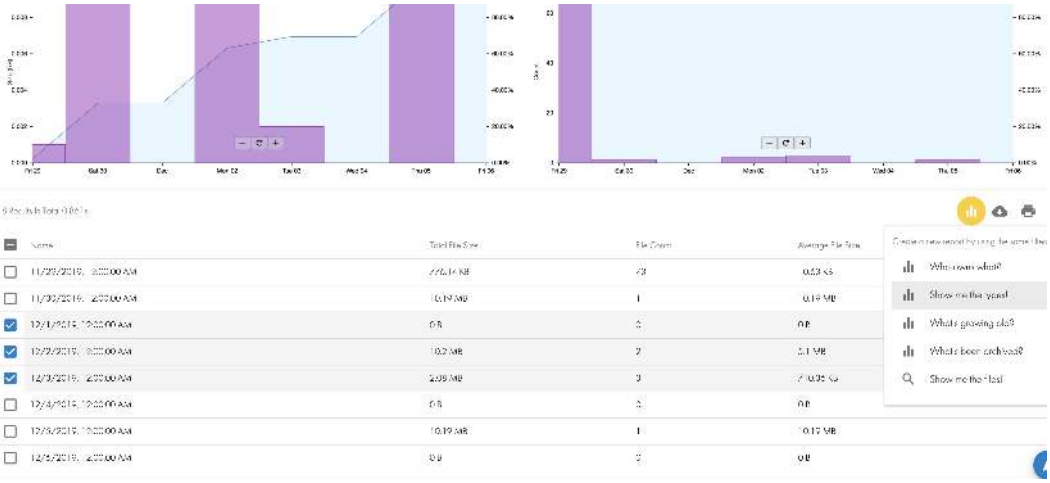

Copy Search results file to Golden Copy for Archive

1. From the files results list you will need to save the files to a CSV file that can be used with Golden Copy to archive the files in the search list. NOTE: You may edit the CSV to delete files that should not be archived.
2. Due to depreciation notice the CSV file must be edited and delete all columns except the file path /ifs/.... this new simpler format will be the only format supported. Command builder feature allows export a search result to a file with path to file one per row file with a carriage return after each row in the file.

3.

4. Using winscp to copy the saved csv file from Search to Golden Copy node 1 and save the file in /opt/superna/var/csvimport
5. Authenticate as the ecaadmin user.
6. Run an archive job and specify the csv file
    a. searchctl archivedfolders archive --id 3fc3795a731daef8 --uploads-file /opt/superna/var/csvimport/igls_search_1606924303.csv
7. Monitor the copy job progress as normal
8. Done

Delete the Data Using Search & Recover

1. **OPTIONAL** - If the objective is to delete the aged files from the file system follow this steps.
    a. Login as an admin user to get the command builder UI option in the UI when running the steps above.
    b. Use the script builder icon and enter **rm** ssh command into the box one and then log the delete to a file in the second box as per screenshot. This will create a bash script with each file in the file to delete the files in the results.

c. 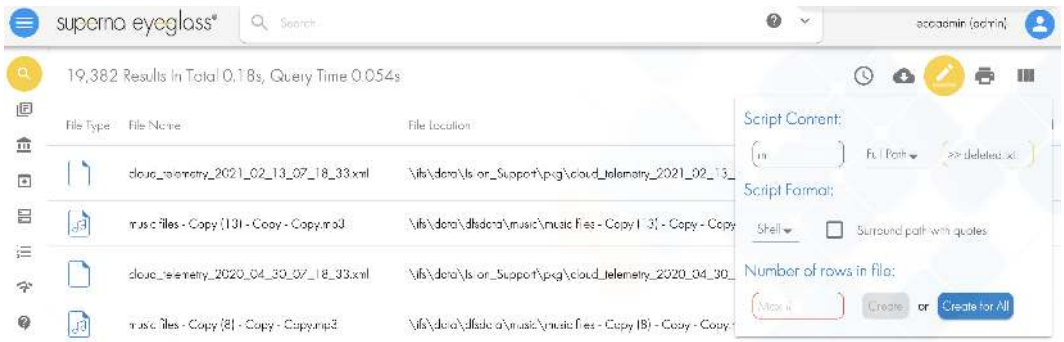

d. Copy this bash script to the Isilon, any location is ok.

   i. Make the file executable **chmod 777 filename.sh**

   ii. Make sure you are logged in as a user with permissions to delete files regardless of the ACL's example root.

   iii. Execute the script command store the log somewhere as a record of the deletion.

e. NOTE: The data copied will have the same path in the object store as existed on the cluster file system.

Recalling Archived Data

The data that was archived will not be required in most cases. If the recall of the data is required. The Golden copy recall feature can recall the archived data based on path only to the recall staging area. The data can be shared from this location or copied back into the original file location. The recall staging area retains the absolute file path when it is recalled.

See the recall instructions here.

© Superna LLC

10.14.4. Immutable Storage for Azure Blob Storage

Home Top

- Overview
- Configuring Blob Container Retention
- How to Configure Legal Hold

Overview

This solution guide explains common configurations for leveraging legal hold or Azure blob ummutable retention options.

Configuring Blob Container Retention

1. Review all documentation on Azure's retention features for blob storage here.
2. Login to Azure console and open the storage account container access policies tab and select retention

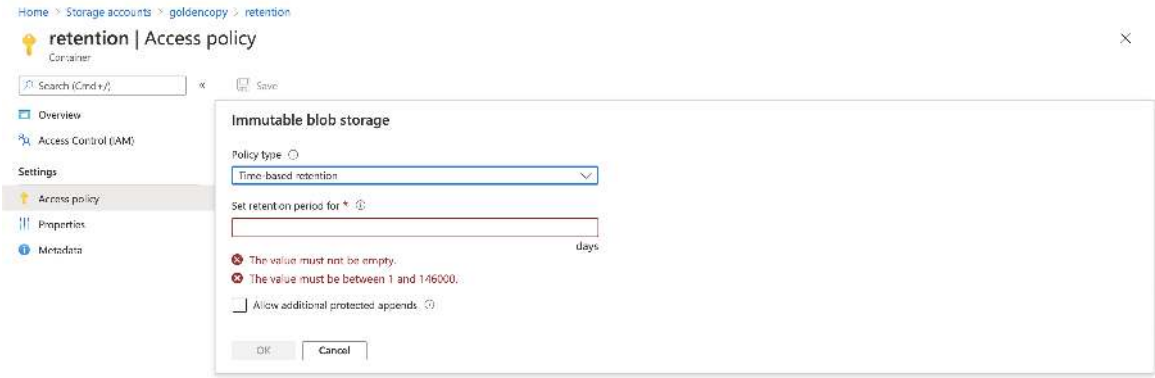

3.

4. enter the Retention in days and save the policy.
    a. NOTE: only 1 retention policy is allowed per container
5. The default mode allows changes to the policy and allows deleting the policy in a trial only mode. It is best practice to test copy data, configure incremental schedules on folders to verify everything is working as expected before locking the retention policy.
    a. NOTE: once a policy is locked the data cannot be deleted.
6. **Read all documentation and limitations before locking a policy since this cannot be undone and the retention of the data will be locked and paid for the duration of the lock.**

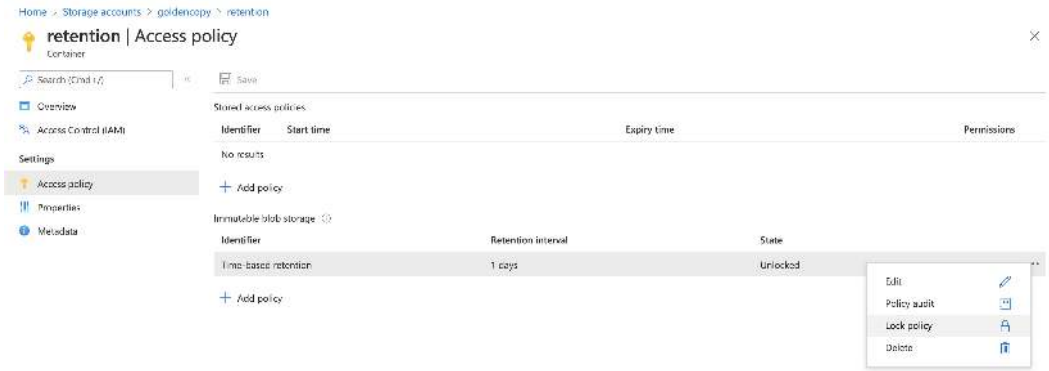

    a.

7. The retention settings are applied to the creation date of each object. Incremental mode on a folder will not be able to update objects that are modified on file system.
8. If you see errors in an archive copy job you can view the errors with the command below changing the job id.
    a. searchctl archivedfolders errors --id job-1605642959317-2077753906 --tail --count 25
    b. If the reason column shows **BlobImmutableDueToPolicy** It means a modified file on disk was copied and the object already existed and is in a locked state.

How to Configure Legal Hold

1. The difference between legal hold versus retention mode is that a legal hold can be applied and then removed at any time. Retention policies once applied and the policy locked cannot be undone and data cannot be deleted until the retention date has expired.
2. In the Azure console the legal hold can be applied along with a tag to describe the reason for the hold.

a.

3. To remove the legal hold after it is applied, click the edit option and then the trash option.

a.

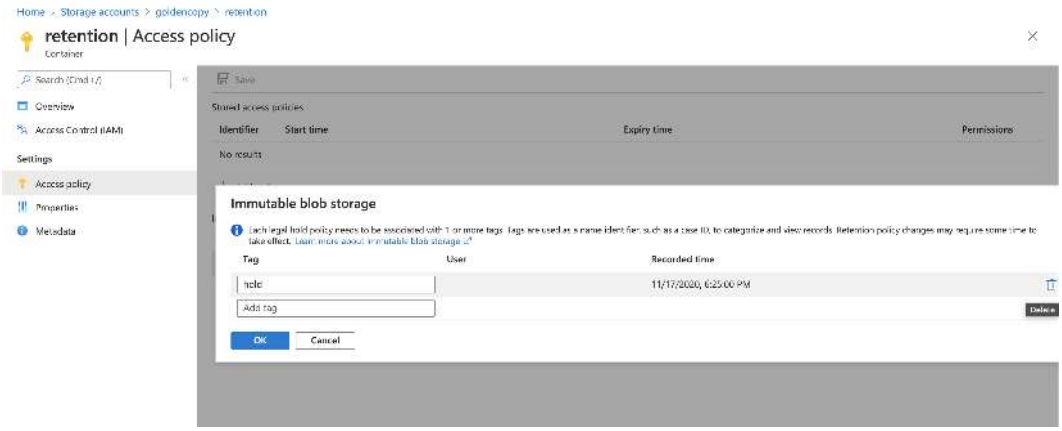

4. If you upload a file that already exists to a legal hold container you will see an error that legalholdpolicy using the errors command to display copy errors.

© Superna LLC

10.14.5. Data LifeCycle - Cost Reduce with Archive to the Cloud

Home Top

- Overview
- Prerequisites
- Use Case
- How to Export Search Results to Golden Copy
- How to Copy data with Search & Recover Export Files

Overview

Golden Copy can integrate with Search & Recover by allowing a search result saved to a CSV to be used as an input to Golden Copy to copy the data. This allows locating data using Search & Recovers powerful search engine that allows file size, age type owner as criteria. The data can be any where in the file system. Unlike Golden Copy that is path based copy or sync only copy.

Prerequisites

1. A Golden Copy folder definition with a cluster and path configured for copying. The folder definition needs to exist, it is not required to copy any data from this folder.

2. The path configured on the folder definition is the base path that must be used when searching for files with Search & Recover

Use Case

1. Identifying old data for archive purposes is a key use case. Search & Recover allows searching by date stamps using modified or last accessed and exporting search results to Golden Copy to copy data. See the solution 4 step process.

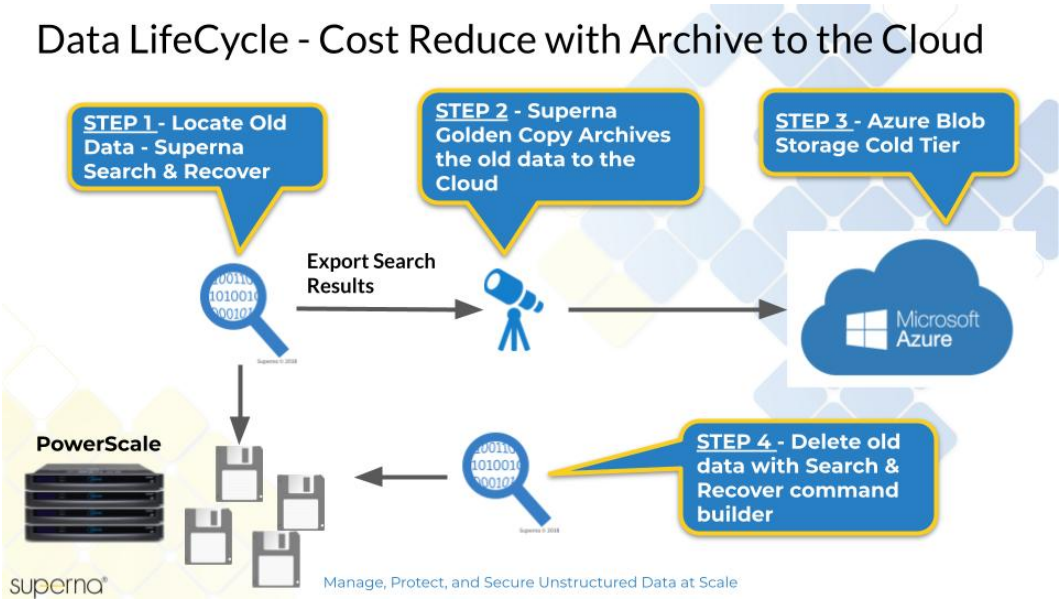

a.

How to Export Search Results to Golden Copy

1. Login to search and execute a search using any of the search options. Enter the path configured within Golden Copy where the search will begin. The search results must match the base path defined within Golden Copy folder definition.
2. Execute the search using any criteria needed age, data stamp size, owner group etc. NOTE: You can also use any of the Quick reports as well. Make sure to enter the path field in the

search interface to bound the search to be files under the Golden Copy folder definition path.

3. Then click the save to csv option from the results table. See example below. Use the Download All Button.

4. 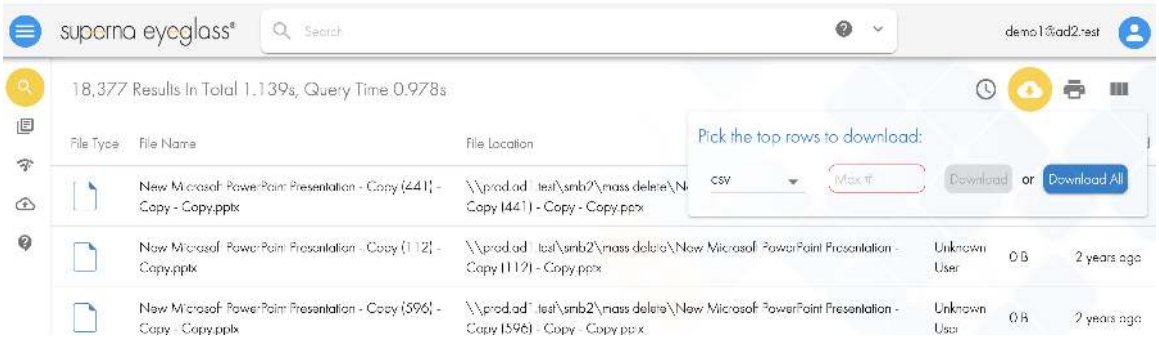

5. Edit the CSV and remove all columns except the path column and make sure each file is listed one per row with carriage return. The golden copy flat file input is 1 file per line, no comma needed.

6. Move to the next section.

How to Copy data with Search & Recover Export Files

1. Using winscp to copy the saved csv file from Search to Golden Copy node 1 and save the file in /home/ecaadmin Authenticate as the ecaadmin user.

2. Run an archive job and specify the csv file

   a. searchctl archivedfolders archive --id 3fc3795a731daef8 --uploads-files /home/ecaadmin/igls_search_1606924303.csv

3. Monitor the copy job progress as normal
4. Done

© Superna LLC

10.14.6. Object Backup and Basic DR Solution

Home Top

- Overview
- Solution Summary
- Configure Golden Copy in Sync Mode
- Install Cyberduck or S3 Browser for DR Access to Data
- DR Scenarios

Overview

This solution is not traditional DR with failover and fail back. This solution is aimed at single cluster configurations that want to leverage Cloud services as an off site copy with the ability to access the data if required. This solution does not handle security of the original data set that is copied from files to objects. This solution assumes accessing the data quickly meets the business requirement without needing to recall all the offsite data.

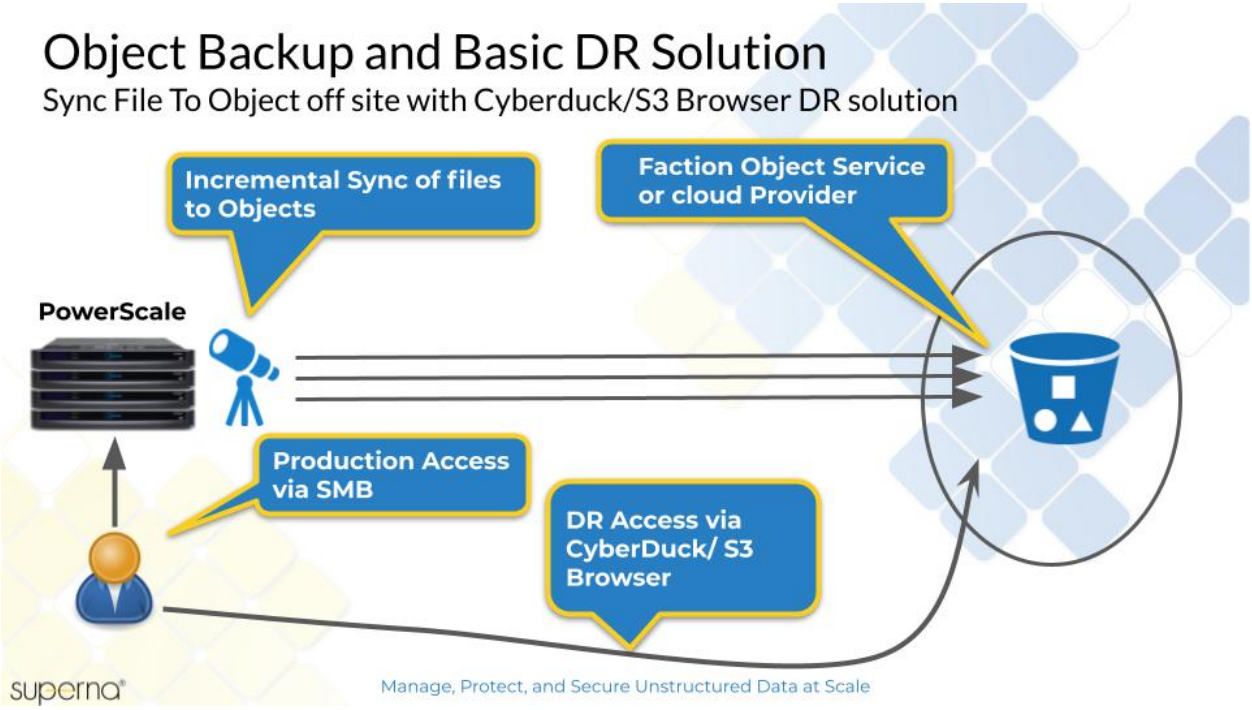

Solution Summary

1. < 100TB data set
2. Data Access is the key priority
3. Failover is not required (applications, or large number of end users)
4. Reduced speed of data access in a DR scenario
5. Existing SMB and ACL security is not required in a DR scenario
6. Dell Faction Object service, or Cloud vendor object storage equal to the size of the data set to protect off site.

Configure Golden Copy in Sync Mode

1. Configure Golden Copy for Sync mode with incremental sync mode. Follow this link here for details.
2. Start the copy full copy job to start the initial data sync.

3.

Install Cyberduck or S3 Browser for DR Access to Data

1. Install cyberduck (download) or S3 Browser (download) on a Windows PC to access DR data.
2. Configure access to the storage bucket used by Golden Copy
    a. Example for Dell ECS

b. 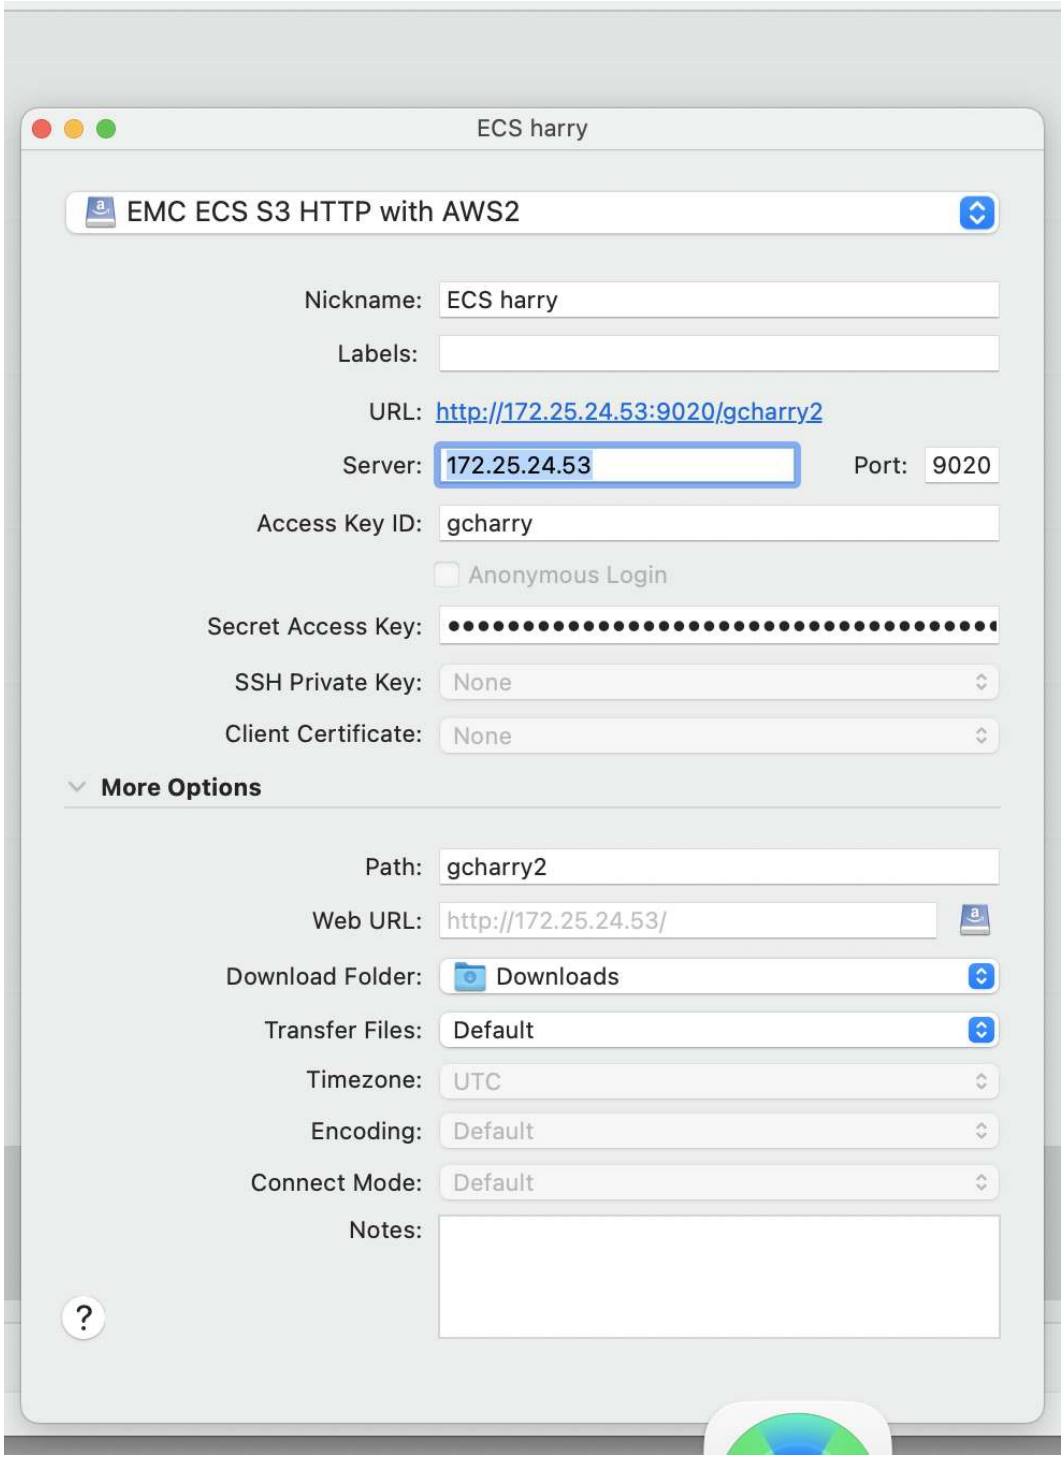

3. Test the connection to the storage bucket and download test files to verify access is working as expected.

a. 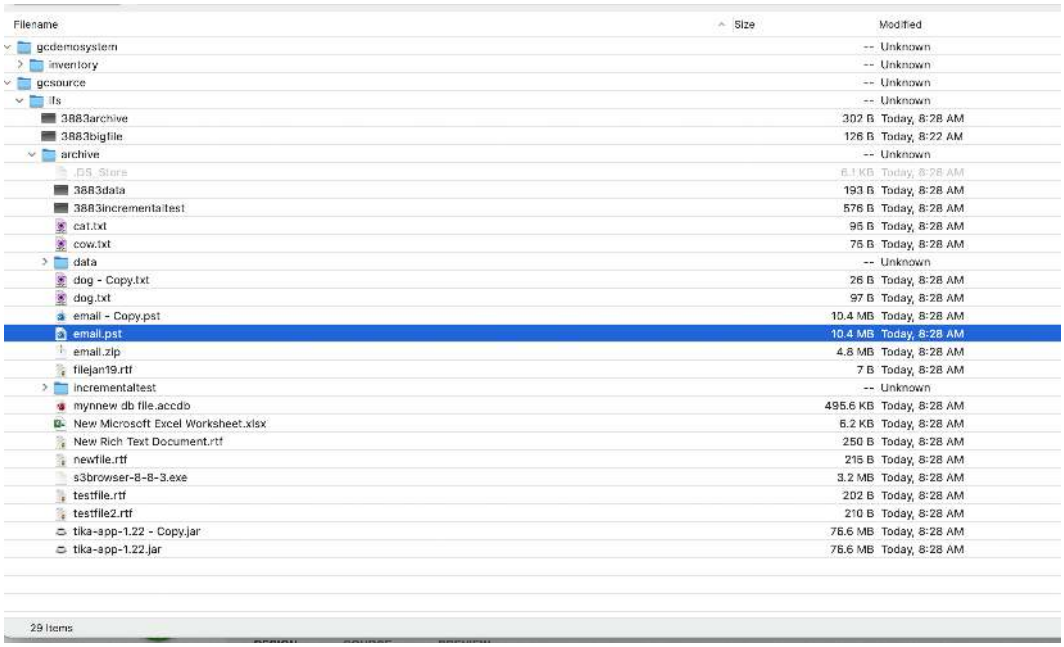

Filename	Size	Modified
gcdemosystem	--	Unknown
inventory	--	Unknown
gcsource	--	Unknown
lfs	--	Unknown
3883archive	302 B	Today, 8:28 AM
3883bigfile	126 B	Today, 8:22 AM
archive	--	Unknown
.DS_Store	6.1 KB	Today, 8:28 AM
3883data	193 B	Today, 8:28 AM
3883incrementaltest	576 B	Today, 8:28 AM
cat.txt	98 B	Today, 8:28 AM
cow.txt	76 B	Today, 8:28 AM
data	--	Unknown
dog - Copy.txt	26 B	Today, 8:28 AM
dog.txt	97 B	Today, 8:28 AM
email - Copy.pst	10.4 MB	Today, 8:28 AM
email.pst	10.4 MB	Today, 8:28 AM
email.zip	4.8 MB	Today, 8:28 AM
filejan19.rtf	7 B	Today, 8:28 AM
incrementaltest	--	Unknown
mynnew db file.accdb	495.6 KB	Today, 8:28 AM
New Microsoft Excel Worksheet.xlsx	6.2 KB	Today, 8:28 AM
New Rich Text Document.rtf	250 B	Today, 8:28 AM
newfile.rtf	215 B	Today, 8:28 AM
s3browser-8-8-3.exe	3.2 MB	Today, 8:28 AM
testfile.rtf	202 B	Today, 8:28 AM
testfile2.rtf	210 B	Today, 8:28 AM
tika-app-1.22 - Copy.jar	76.6 MB	Today, 8:28 AM
tika-app-1.22.jar	76.6 MB	Today, 8:28 AM

29 Items

4. This completes the configuration. This would be repeated on any host that needs access to the data in a DR scenario.

DR Scenarios

1. Source cluster is not accessible.
    a. Use the S3 mount tool data access described in this document.
2. Source cluster is accessible.
    a. Use the S3 mount tool data access described in this document.
    b. Start Golden Copy bulk recall from object storage back onto the source cluster using the temporary access method described in this document.
    c. Follow this link to recall steps in Golden Copy

© Superna LLC

10.14.7. Bulk Loading Data with Azure Data Box

Home Top

- Overview
- When to use
- Major steps to order Data Box , install, load data and return
- How to Load data onto a Azure Data Box with Golden Copy
- How to switch the folder to incremental sync mode

Overview

Azure Data Box is an on-side device used to bulk load data into Azure Blob Storage that is faster than using an Internet connection. This section provides a high level explanation of how to use Golden Copy to load file system data onto an Azure Data Box. This process ensures the file system metadata is maintained on the objects copied to the Data Box device. Once the data is imported to your storage account and container, Golden Copy can maintain an incremental update of the data over the Internet.

When to use

1. This option should only be used for very large files and when Internet bandwidth is limited to veery low rates. This option is not faster for small files or when 100 Mbps or more of Internet bandwidth is available. Consult with support before using this option. In most cases sending data directly over the Internet will always be faster.

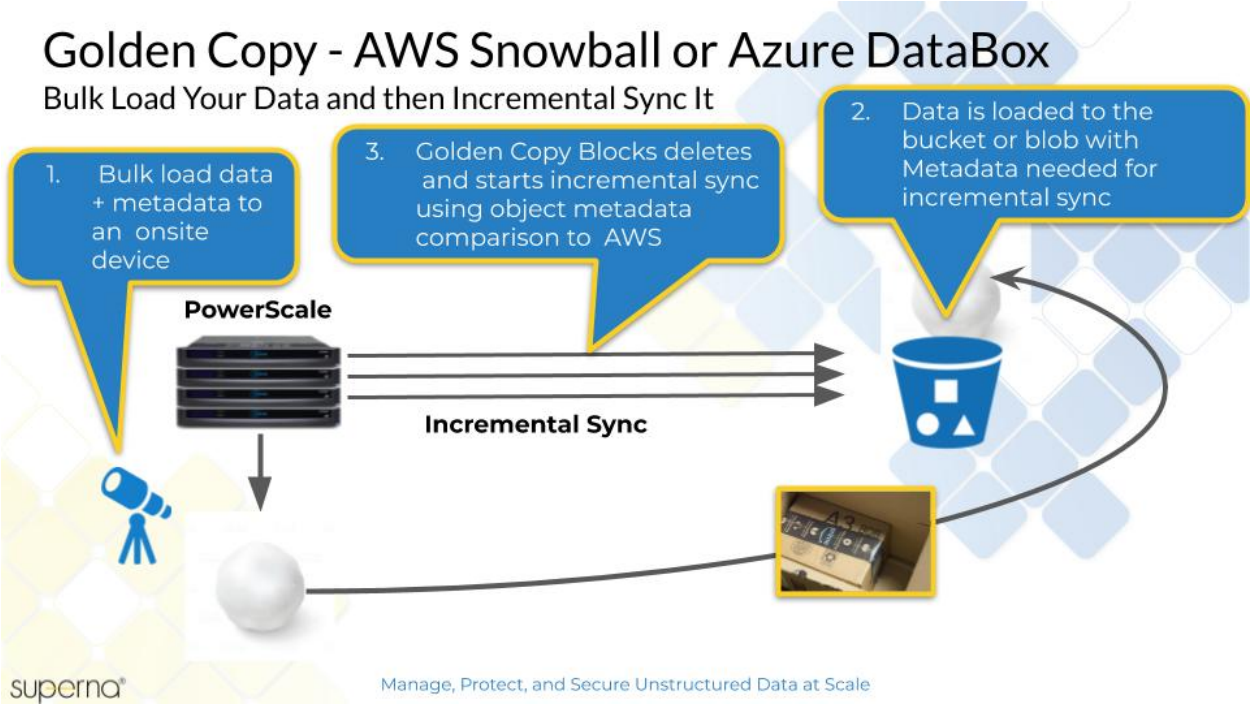

Major steps to order Data Box , install, load data and return

1. Follow steps here to order your Data Box making sure to have a **resource group**, **storage account created of type Blob Storage** and in a region.
2. Create Storage Account in Azure.
3. Create a container in the Storage Account to receive bulk load data.
4. Complete the steps on the quick start page:

a. Order.

b. Unpack.

c. Connect and unlock.

d. Copy Data Steps click here to follow the Golden Copy Steps.

e. Ship to Azure.

f. Verify your data.

g. Clean up resources .

How to Load data onto a Azure Data Box with Golden Copy

1. **Create Blob container**

    a. In the Azure console open the Storage Account used to order the data box, create a container name to store your PowerScale data, and set the Public Access to Private.

2. The Data Box should be on the network with an IP address following steps in the quick start guide.

    a. To get the blob storage API endpoint follow the steps on this link https://docs.microsoft.com/en-us/azure/databox/data-box-deploy-copy-data-via-rest#add-device-ip-address-and-blob-service-endpoint .

    b. Record the end point DNS name from the steps above on the Data Box Azure console page and record the Access ID.

i. 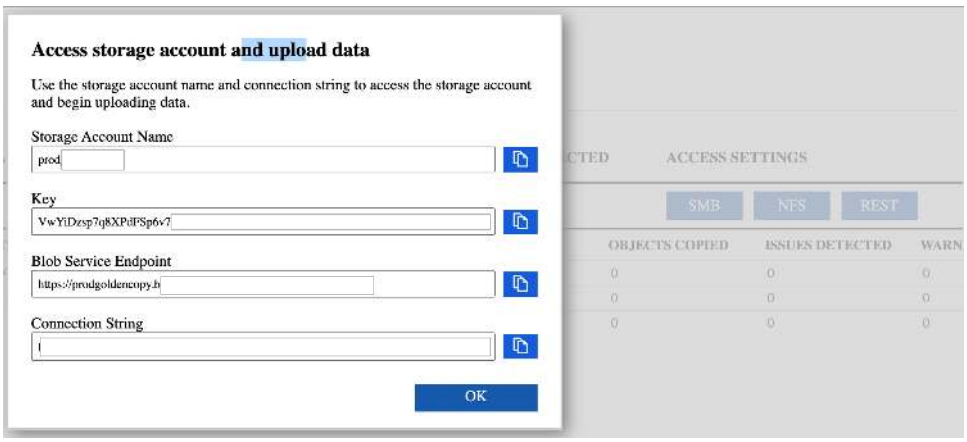

c. On a Windows PC download Azure Storage Explorer and test access to the Data Box by adding host entry to the Windows PC as per the quick start guide section here.

d. **Mandatory Step:**

   i. Create Azure container in the storage account that the Data Box was ordered with. The bulk load data will be stored in this container. It is required to create the same container on the Data Box to upload data.

   ii. Using Storage Explorer connect to the Data Box (use the Storage Account when ordering the Data Box) and create a container with the same name (case sensitive) that was created in Azure. See example below.

iii. 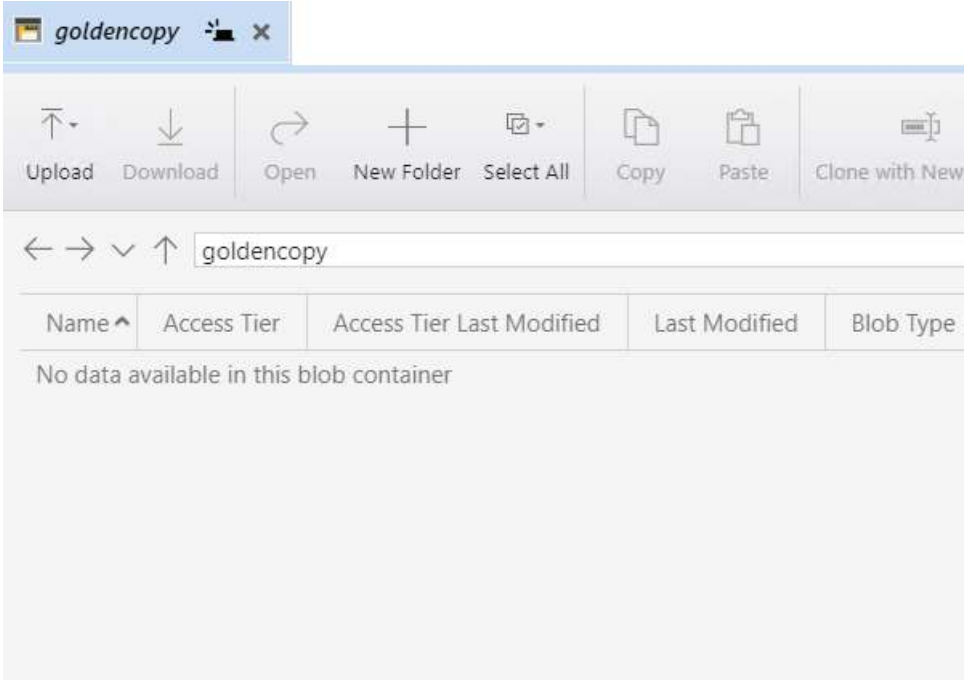

e. Once connection is verified and the container is created, continue to the next section

3. **Update DNS with Data box A record using matching blob endpoint**

   a. The quick start guide requires the endpoint to resolve to the on-side ip address of the Data Box. Add a domain record matching the endpoint listed on the Data Box configuration page. A Windows DNS example is shown below, where domains are added to the blob domain name and an A record added to complete the DNS entry lookup pointing at the on-side IP address of Data Box. A com domain needs to be added to the DNS server to build the service endpoint DNS entry. **The full fqdn will be similar to this .<storageaccountname>.blob.<serialnumber>.microsoftdatabox.com**

b. Create a **blob.<serialnumber>.microsoftdatabox.com** zone on the DNS server used by Isilon. Replace <serialnumber> with the value from your unique databox endpoint URL. Then create an A record in the next step.

c. Create the A record under the new zone from the step above. The A record will be <storage account name>.blob as the host A record and replace the ip address with the IP used on the databox on your LAN. See example below **(actual name was blanked out example only). The FQDN should show the exact same DNS entry as the REST endpoint shown on the Databox admin UI. Verify before saving.**

i. 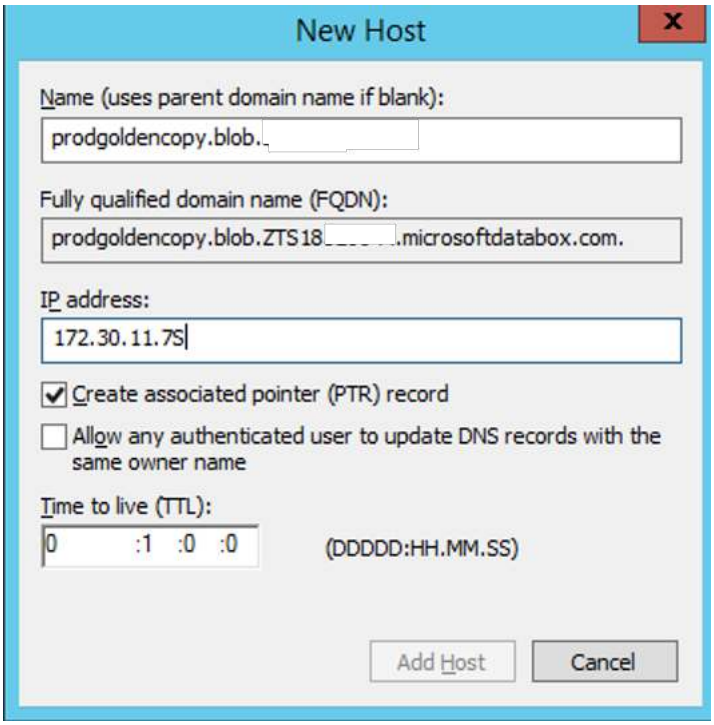

d. **Verify DNS Resolution**

i. Now login to the PowerScale cluster with ssh (root user) and verify the DNS name resolves. Ping the service endpoint and verify the DNS entry resolves correctly to the IP address of the Data Box.

ii. **Do not proceed until the endpoint DNS name resolves correctly.**

iii. If the ping does not return correct IP it means the DNS server used by the Groupnet 0 is not able to resolve the name. Double check the DNS configured on groupnet 0 System zone is configured correctly.

4. **Add the Data Box device to Golden Copy using the "searchctl archivefolder add" command.**

   a. searchctl archivedfolders add --isilon **<cluster name>** --folder /ifs/archive --secretkey **<secret key>** --endpoint **<databox endpoint DNS name>** --container <container name> --bucket **<storage account name>** --disable-incremental true --cloudtype azure

   b. NOTE: You specify the Storage Account name configured when the Data Box device was ordered.

   c. NOTE: Incremental Mode is disabled on the folder since the data is shipped back and takes time to load into the container. This can be switched on again with the steps below.

5. Start the archive job to copy data to the Data Box:

   a. Get the folder id from the add folder command using **searchctl archivedfolders list.**

   b. searchctl archivedfolders archive --id xxxxxxx .

   c. Monitor the archive copy job with:

      i. searchctl jobs running (get the job ID)

      ii. searchctl jobs view --id job ID

6. Wait for the Archive job to complete. This could be hours or days depending on the amount of data. Use the commands above to monitor progress.
7. Return Data Box to Azure. Follow the return and clean steps in the Azure guide (referenced above)
8. Remove DNS entry added to DNS for the Data Box device.
9. Once notification from Azure about data loaded into the container, login to Azure portal and open the container to verify data is present in the blob storage.
    a. Click on a file, and click properties to verify the custom metadata is visible on the object
10. Done

How to switch the folder to incremental sync mode

1. Once the data is loaded into the Azure container and verified from the Azure console, it's possible to switch the folder configuration from bulk load mode to incremental sync mode. The metadata encoded into the objects allows incremental sync mode.
2. The folder configuration used for bulk data load can be modified to use the Azure production endpoint.
3. Get the folder id:
    a. searchctl archivefolders list .
4. Modify the folder to use the Azure blob endpoint:

a. searchctl archivedfolders modify --id xxxx --endpoint **<storage account service endpoint DNS name>** --disable-incremental false .

b. NOTE: the endpoint is different for the Internet hosted blob container and must be changed. See Configuration guide for details.

5. Verify incremental schedule is configured and wait for the next incremental job to run.

   a. Verify the incremental job runs and that incremental copy was able to add new files, or delete objects for deleted files using the html report. The html report export command can be found in the cconfiguration guide.

© Superna LLC

10.14.8. Bulk Recall of Data from AWS and Azure

Home Top

- Overview
- When to use
- AWS Snowball Data Export
- Azure Data Box Data Export

Overview

In cases where large quantities of data are required to be recalled it's more cost effective to use either the AWS Snowball or Azure Data Box solution to avoid egress charges. The chosen solution should be used following the guides on how to request or export data from a storage bucket or container.

When to use

1. This option should only be used for very large files and when Internet bandwidth is limited to veery low rates. This option is not faster for small files or when 100 Mbps or more of Internet bandwidth is available. Consult with support before using this option. In most cases sending data directly over the Internet will always be faster.

AWS Snowball Data Export

1. This guide covers ordering an export Snowball device to be shipped to your location. https://docs.aws.amazon.com/snowball/latest/ug/create-export-job-steps.html
2. Use the Snowball client to copy data from your Snowball to a mounted export on the cluster to start your copy.
3. NOTE: This will not apply metadata during the recall of data from Snowball.
4. A release of Golden Copy will add support for Snowball recall + meta data copy for high speed data recovery, using the cluster nodes to retrieve data over S3 from the Snowball device.

Azure Data Box Data Export

1.

© Superna LLC

10.14.9. Bulk Loading Data with AWS Snowball

Home Top

- Overview
- When to use
- Major steps to order Snowball, install, load data and return
  - Create an Export Job
  - Receive the AWS Snowball device
  - Connect the AWS Snowball device to Your Local Network
  - Transfer Data
  - Follow Golden Copy Steps below
  - Return the Device
- How to Load data onto a Snowball Device with Golden Copy
  - Configure a profile
  - Unlock
  - Get Access key id
  - Get Secret access key

Overview

AWS Snowball is an on-site device used to bulk load data into Amazon S3 that is faster than using an Internet connection. This section provides high level explanation of how to use Golden Copy to

load file system data onto an AWS Snowball. This process ensures the file system metadata is maintained on the objects copied to the Snowball device. Once the data is imported to your AWS storage bucket, Golden Copy can maintain an incremental update of the data over the Internet.

When to use

1. This option should only be used for very large files and when Internet bandwidth is limited to veery low rates. This option is not faster for small files or when 100 Mbps or more of Internet bandwidth is available. Consult with support before using this option. In most cases sending data directly over the Internet will always be faster.

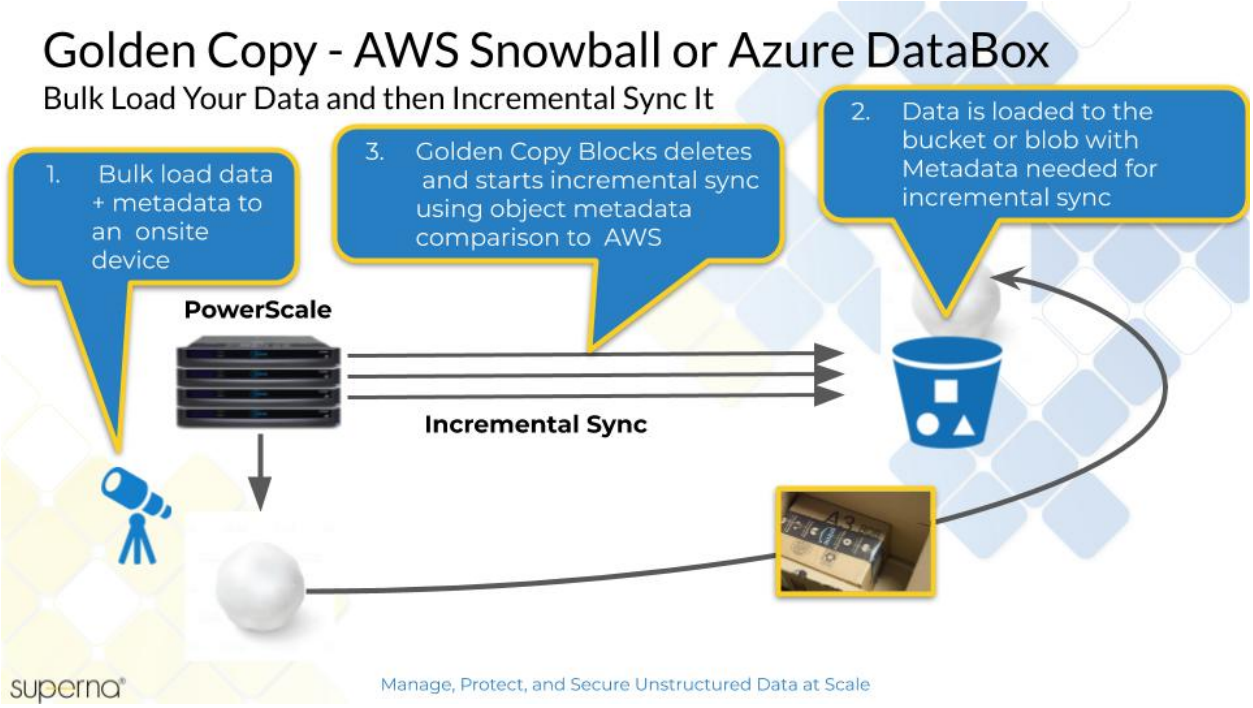

The AWS Snowball end to end process is explained here.

Major steps to order Snowball, install, load data and return

Topics

Create an Export Job

Receive the AWS Snowball device

Connect the AWS Snowball device to Your Local Network

Transfer Data

Follow Golden Copy Steps below

Return the Device

How to Load data onto a Snowball Device with Golden Copy

1. Start here with AWS Guide to create an import job, and have a Snowball device shipped to your location https://docs.aws.amazon.com/snowball/latest/ug/create-import-job-steps.html
2. Setup the Snowball device on your network low latency to the PowerScale cluster.
3. Get prepared to authenticate to the Snowball Edge:

a. Get manifest file and unlock code:

i. Login to AWS- go to the Snowball service, expand job, get credentials and unlock key:

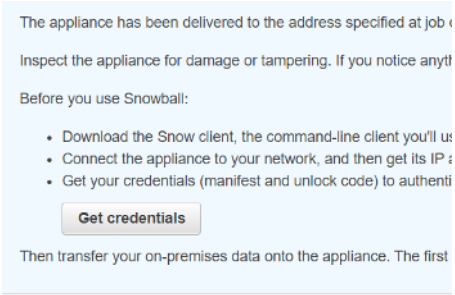

ii.

iii. Save the manifest file locally. This example uses: g:\snowball\ :

Configure a profile

This prevents you from having to enter manifest and unlock code in subsequent commands. Profile named gcsnowball in this case - could have been anything - gcsnowball chosen to match snowball in AWS.

```
C:\Program Files
(x86)\SnowballClient\bin>snowballEdge.b
at configure --profile gcsnowball
Configuration will stored at
C:\Users\dmanning\.aws\snowball\config\
snowball-edge.config
Snowball Edge Manifest Path:
g:\snowball\JIDa54cf5fd-9a0b-426e-8873-
a05c556e4c3b_manifest.bin
Unlock Code: xxxx-0b5bf-b4859-492aa-
yyyy
Default Endpoint: https://192.168.1.11 (IP of the
snowball device on your network)
```

Unlock

```
C:\Program Files
(x86)\SnowballClient\bin>snowballEdge.b
at unlock-device --endpoint
https://192.168.1.11 --manifestfile
g:\snowball\xxxxxxxxx-9a0b-426e-8873-
yyyyyy_manifest.bin --unlock-code xxxx-
0b5bf-b4859-492aa-yyyy
```

Get Access key id

```
C:\Program Files
(x86)\SnowballClient\bin>snowballEdge.bat list-
access-keys --profile gcsnowball
{
  "AccessKeyIds" : [
"AKIACEMGU3DCNBEEIYYSGGA5XPBW" ]
}
```

Get Secret access key

```
C:\Program Files
(x86)\SnowballClient\bin>snowballEdge.bat get-
secret-access-key --access-key-id
AKIACEMGU3DVEEIYYSGGA5XPBW --profile gcsnowball
```

iv. Save the AccessID and Secret key to add archive folder to Golden Copy.

v.

4. Add the Snowball device to Golden Copy using the "searchctl archivefolder add" command, and using the "--force" flag that is required to disable validations used with AWS.
    a. NOTE: You specify the bucket name configured when the Snowball device was ordered. example command
    b. NOTE: incremental mode is disabled on the folder since the data is shipped back and takes time to load into the

container. This can be switched on again with the steps below.

c. searchctl archivedfolders add --folder /ifs/data/archivedata --isilon **<cluster name>** --accesskey xxxxxx --secretkey yyyyyy --endpoint http://x.x.x.x:8080 --bucket <bucket name> --region snow --disable-incremental true --cloudtype other --force

d. The "Access key ID" and "Secret access key" are created using the create local user process documented AWS Snowball Edge create local user and summarized above. Replace the yellow highlighted sections to match your installation when you created the Snowball Edge order.

5. Start the archive job to copy data to the SnowBall Edge:

   a. Get the folder id from the add folder command using searchctl archivedfolders list:

      i. searchctl archivedfolders archive --id xxxxxxx --force (NOTE: Use the force flag to by pass AWS checks that do not apply to Snowball devices)

   b. Monitor the archive copy job with:

      i. searchctl jobs running (get the job name)

      ii. searchctl jobs view --id jobname

6. Archive job completes .

7. Return AWS Snowball device (see steps above).

8. Done.

© Superna LLC

11. Eyeglass Performance Auditor Admin Guide

Home Top

- Performance Auditor Overview
- Performance Auditor Requirements
- How to Use Performance Auditor to Root Cause
- Performance Auditor Advanced Configuration

© Superna LLC

11.1. Performance Auditor Overview

Home Top

- Overview

Overview

Performance issues can happen at any time, and Scale out NAS presents a challenge to quickly find the root cause. The business is being impacted every minute data is unavailable, and you are under pressure to find answers to resolve the issue quickly.

The solution requires identifying the user, the node, the path or subnet that's the source of the unexpected workload impacting business critical workflows.

The performance auditor tool simplefies the understanding of how the cluster's file system is being used, and a simple easy to use interface allows switching between user, path, node or subnet view of the workload.

One click root cause makes Performance auditor a key tool to identify the source of degradation with 1 second updates to all IO across the cluster. The data source for this unique view is audit data generated by the cluster. The Superna suite of products uses real-time audit data ingestion process and stream based processing based on Apache Kafka Streams technology.

Historical data is stored allowing reply , rewind and fastword through real time data to locate performance issues.

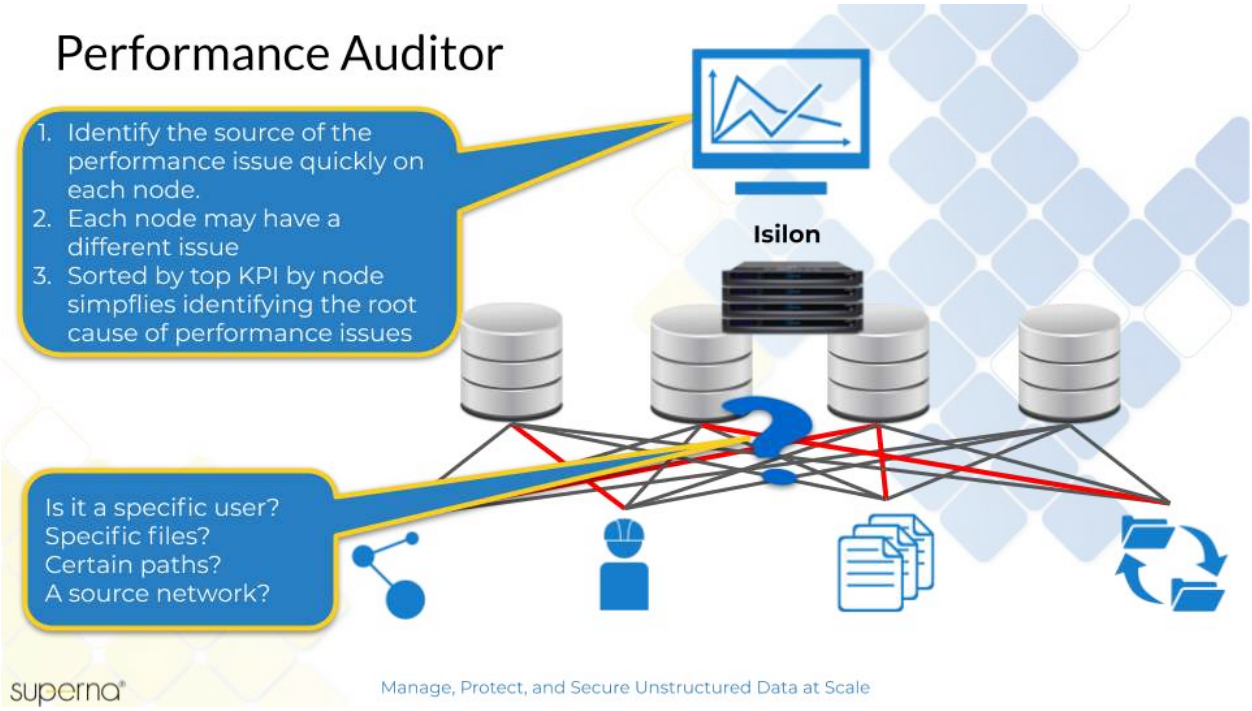
Performance Auditor
1. Identify the source of the performance issue quickly on each node.
2. Each node may have a different issue
3. Sorted by top KPI by node simpflies identifying the root cause of performance issues
Isilon
Is it a specific user?
Specific files?
Certain paths?
A source network?
superna
Manage, Protect, and Secure Unstructured Data at Scale

© Superna LLC

11.2. Performance Auditor Requirements

Requirements for Performance Auditor in Unified deployments

1. Auditing enabled on monitored clusters.
2. NFS mount for monitored clusters to ingest audit data.
3. Browser PC to Eyeglass VM must open port 2014 for UI updates.
4. Eyeglass clustered Agent cluster deployed with release 2.5.6 or later code level.
5. Audit Event > 6000 events per second (contact support to get this value from support logs), will require a memory increase on the ECA cluster nodes when running Ransomware Defender Easy Auditor:
    a. > 6000 events per second = 20GB memory per ECA VM.
    b. < 6000 events per second can stay at 16GB per ECA VM.
6. Large node count clusters require more ECA VM's and more memory
    a. PowerScale clusters over 20 nodes will require additional ECA VM's for processing and will require additional memory per ECA VM.

b. 30 nodes or more = 9 ECA VM's with 20GB RAM per VM when running unified Ransomware, Easy Auditor and Performance Auditor applications.

7. **Minimum Audit Events Required**

   a. For 8.2 or later clusters that allow configuration of individual audit events, the following audit success events can be set to reduce the audit work load on the cluster and enable only the audit events required for performance auditor:

      i. close_file_modified, open_file_read, open_file_write, read_file, write_file, read, write .

© Superna LLC

11.3. How to Use Performance Auditor to Root Cause

Home Top

- Whats New
- Important Information on Performance Auditing - Read Me First
- Functional Limitations
- RBAC Role Required to Use Performance Auditor
- Baseline Performance Monitoring Overview
  - How to use the Baseline Category Average
  - How to Trend cluster IO over time for the top cluster nodes
- How to Use the UI
  - How to Navigate the UI
  - How to View Historical Performance data with the Rewind Feature
    - How to view long term historical summary performance data graphs
  - How to switch the UI display from one cluster to another
  - Cluster Wide Performance View Shows the Top Users, Paths, File Types, and Subnets and Category Baseline Averages
  - User View shows the Top Users with Baseline Average and Nodes, Path, Subnet and File Types breakouts below
  - Path View shows the Top paths with Baseline Average and node, user and subnet breakouts below

    - Subnet View shows the Top subnets with Baseline Average and nodes, users, and paths breakouts below
    - File Types View Shows the Top application types with Baseline Average and nodes, user and subnet breakouts below
- How to change the throughput rate shown for all Views
- How to Switch to Application Requests Statistics
- How to trouble shoot by Pinning Users, files, nodes, subnets to Performance Views
    - How to pin a User using with Drag and drop
    - How to remove a pinned Record
    - How to add a pinned record for a user, file, subnet, node application type not displayed in the UI
- How to enable MB committed mode for long running file copy application use cases
    - How to Enable MB per minute committed mode

Whats New

1. **2.5.7**
    a. **Historical cluster node reads, writes and total IO is now available historically in a live graph.**
2. **2.5.6 update 2 includes Rewind**

   a. This feature stores real time category data for up to 14 days allowing administrators the ability to rewind and playback real time performance in the past.
   b. The feature allows date selection and skip forward and backwards through the historical data with a pause play capability.
3. **Unified release** - Eyeglass DR Edition, Ransomware Defender, Easy Auditor, Cluster Storage Monitor and Performance Auditor all share a common desktop, and the ECA cluster can process Ransomware Defender, Easy Auditor and Performance Auditor statistics.
4. This counts the number of discrete application commits to a file (reads or writes) are occurring per second. For example, a file copy commits all the bytes at the end of the file copy. An open file that has data saved to it on a regular basis will show several application requests per second.
5. This new statistic is visible per node, user, file or subnet and application extension.
6. The stat is useful to compare applications to each other when looking at performance issues. Applications that commit data more often will impact the network than applications that commit less often.
7. This statistic is application centric versus storage centric. For example, the IOP is a storage device performance counter where as the Application Requests is based directly on how the application is using the file system.

8. Category Baseline Averages - Each category display now calculates a running average performance stat that can be compared to all other top resource consumers to provide guidance if the current throughput is above or below the average. This allows a quick comparison to current real time performance to determine if you are above or below the today's baseline average.
    a. The baseline average is computed by averaging all IO for a given category during the last minute, so this provides relative context to the top consumers above the average over the same time period.

Important Information on Performance Auditing - Read Me First

1. Data is sampled in a sliding 60 second window and updated each second to the UI. Results and relationships are crated every second to update the UI.
2. The UI ONLY shows the top 5 resources in each category that consume resources nodes, files, users, subnets and applications. The pinning feature allows adding a resource not in the display of top resource consumers.
3. Files that are copied within a minute will show a rate over the entire minute even if the file copy takes 10 seconds, the rate will be averaged over 1 minute. This flattens out IO rates without spikes in performance as seen from a packet network capture tool.

a. **NOTE: Copying a file with Windows and then comparing to Performance Auditor is not a valid test. Performance Auditor is looking at committed data by the application layer. Applications can copy data but not commit the data to the file. This is a key difference to understand between counting packets and MB's saves to a file.**

4. Files copied or read that are large long IO's over 1 minute in length, will record the octets only once the client application commits the read as completed, or commits the write as done. Performance Auditor monitors file creates in a cache (1 Million files by default), and will watch for the audit event with octets read or written to compute the rate in MB's. **NOTE: This means that as a file is copying it will not be displayed in the UI until it commits data to the file. This happens at the end of the file copy process, at this time the UI will show the correct MB's rate this file copy consumed, but the actual file copy could have started many minutes early.**
5. Files that are opened for read or write will always show the correct rate averaged over the minute as reads and writes are issued by the application and committed by the SMB or NFS protocol each time an application saves or reads. This type of file IO will be displayed in the minute that the IO operation occurred.

Functional Limitations

1. SMB IO will display the AD user

2. NFS hosts will display IP address and UID in numerical format
3. HDFS IO will display local users defined on the cluster
4. Pinning feature only supports AD SMB user names

RBAC Role Required to Use Performance Auditor

1. Release 2.5.6 update 2 20258 build or later
2. The Performance Auditor icon will not display unless added to the administrator role or create a new role.
3. The administrator role can be edited to display the new icon if a valid license has been applied. After adding the role and saving, you must logout and login again.

a. 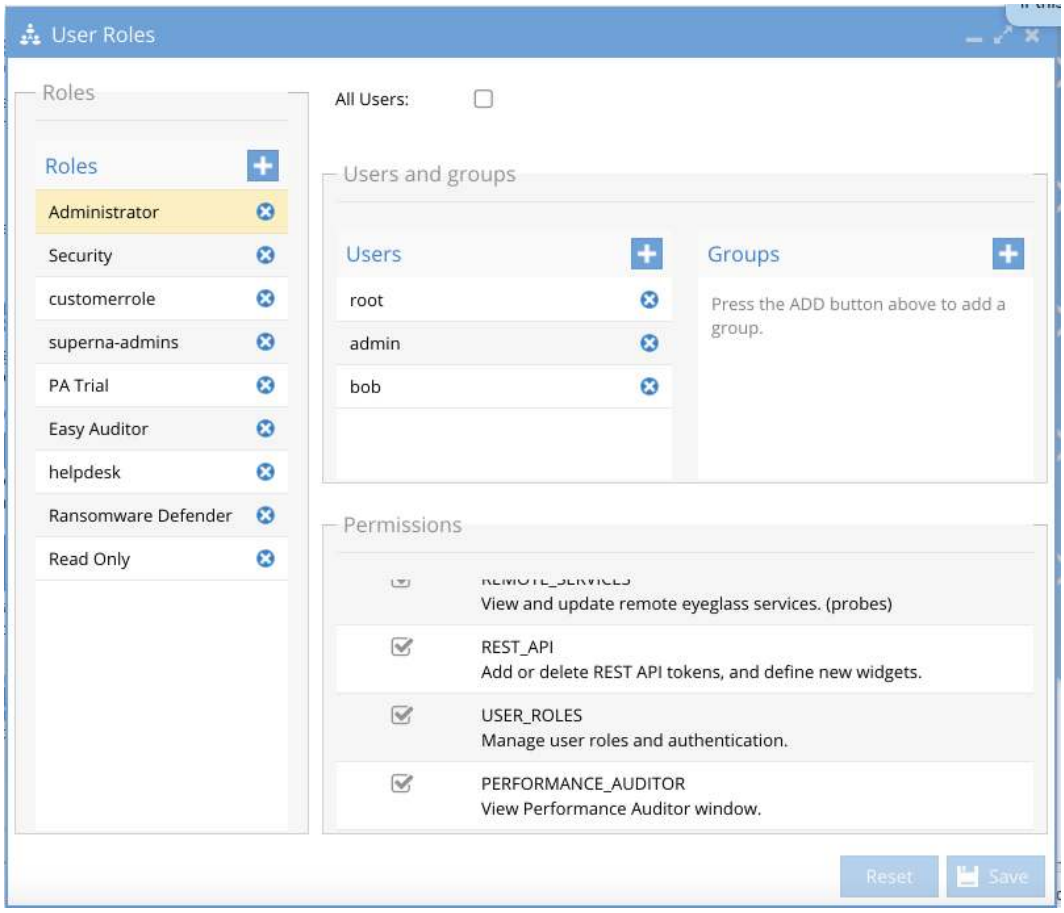

Baseline Performance Monitoring Overview

To understand performance it is important to know what normal performance looks like. This is possible with Performance Auditor's baseline Category Average feature. This feature will average all nodes, users, files, subnets and application extensions for read and write bandwidth to produce a cluster wide average. This average can be used to compare with the top 5 resources consumers in each category and see how high above the average the object is performing.

An average will not provide the best possible comparison, so Performance Auditor computes 2 positive and negative standard deviations above and below the average to provide a band that represents a more accurate average workload comparison.

How to use the Baseline Category Average

1. Click on the Average indicator in the GUI to show the Read, Write and Total break down for the category total average. This can be used to compare to the top consumers and will display on the top right main performance display. Click Anywhere in the UI to unselect the Average indicator display in the top right of the UI.
2. Now view the graph on the left to see the yellow shaded band that shows the full range of the average including the 2 standard deviations, this provides a more accurate view of the baseline average for the category.
3. NOTE: The average is calculated across the same 60 second window as the top consumers in each category, so the comparison is the same time period shown for the top user, nodes, files, subnets and applications.
4. See example average band below:

5. 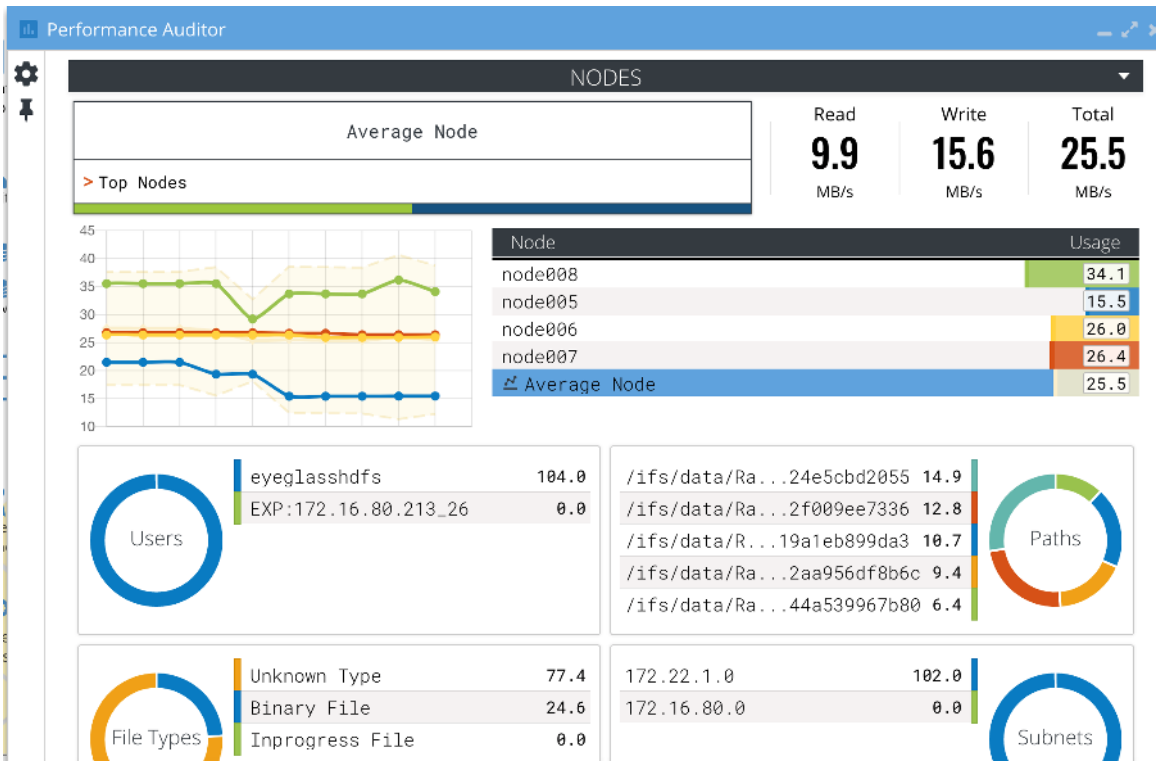

How to Trend cluster IO over time for the top cluster nodes

1. New in release 2.5.7 is the history mode that provides a time series graph of historical cluster performance data.
2. The history view can be accessed from the history link shown below.
    a.

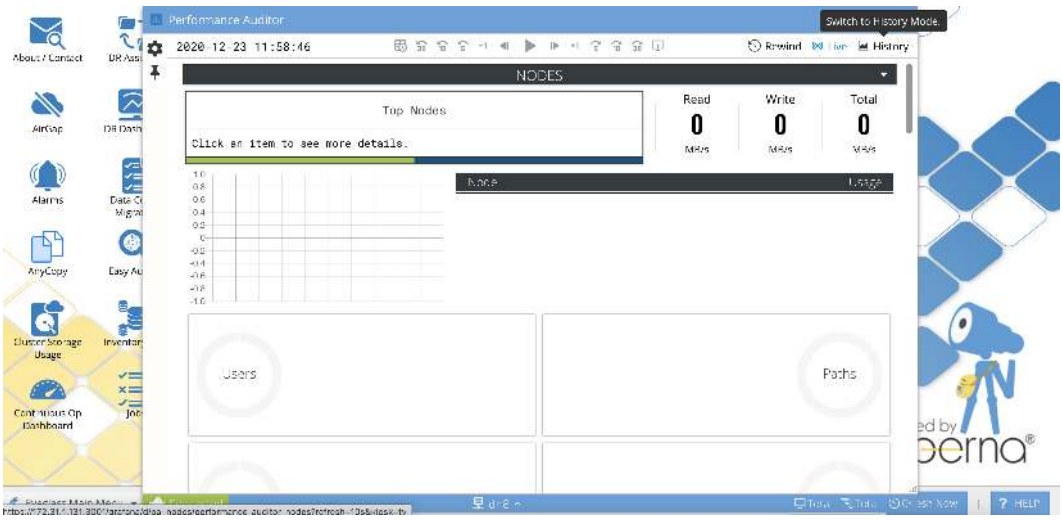

3. The graph below is an example of historical live data view

4.

How to Use the UI

The interface uses views to show the performance data from cluster wide view that shows all nodes, users, paths, and subnets consuming read and write bandwidth IO displayed in KB's,MB's or GB's. The rate selector can toggle between MB's or GB's and is found in the top right of the UI.

How to Navigate the UI

1. How to switch between Vies (Users, Nodes, Paths,Subnet, File Types) using the drop down arrow on the top right of the UI.

   a. 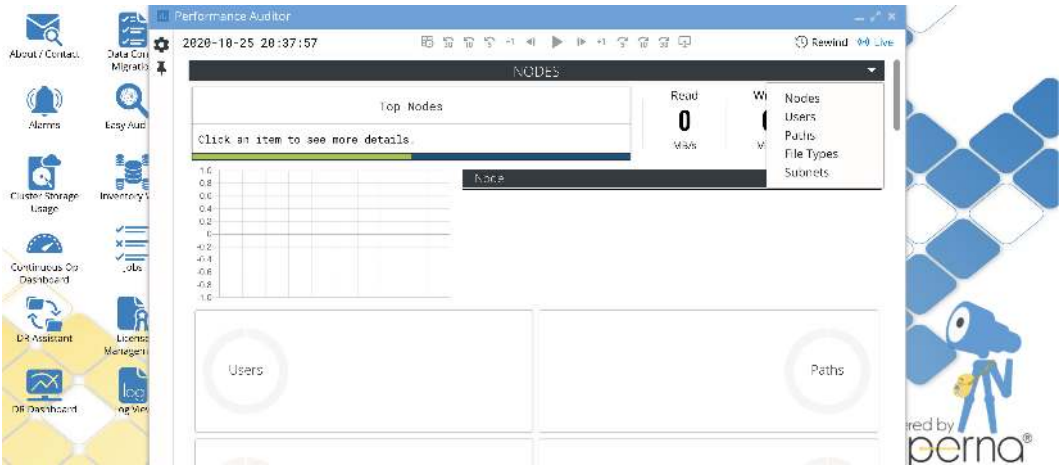

2. Select the rate in the settings icon.
3. In any view select an object in the main window to show the specifics for the selected object below. Example in the default view selecting a node will dynamically show users, paths , subnets and file types that are consuming resources on the selected node. This concept applies to all views.
4. By selecting an item in any view, anywhere in the UI will add the object to the Top box at the top of the UI (see example below). This allows viewing this object's performance with Reads and Writes bandwidth broken out including the sum of both for a Total.

a. 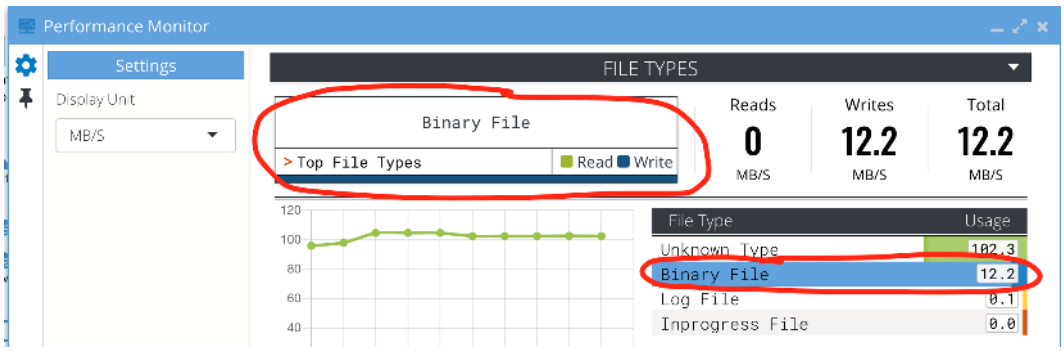

5. Performance Data indicator at the bottom of the Settings Window, shows "Connected" to the Performance collector and how long ago the "Last update" was revived by the UI. This tells you how current the data is that is displayed. Connected means the back end data source is providing performance data to the UI.

How to View Historical Performance data with the Rewind Feature

1. The historical data is stored for approximately 14 days before it is auto deleted. This time range can be extended and will consume more disk space on the ECA cluster nodes.
2. You can only view Rewind mode or Live mode at one time.
3. Click the Rewind button at the top right for Historical mode or Click Live mode to switch back to current date and time.
4. UI Controls
    a. The date and time of the data is shown at the top left of the display.

b. Play pause button to stop playback or start again. Once it plays it will playback 1 second updates at the day and time selected.

c. Select the Calendar button to pick a date and time. Once selected use the + - 1 minute, + - 5 minute, + - 10 minute and + - 30 minute button to fast forward or rewind through performance data to locate the issue.

d. Once you have selected a day and time to view, you can switch views to see different category data nodes, users, paths, subnets and applications.

5. 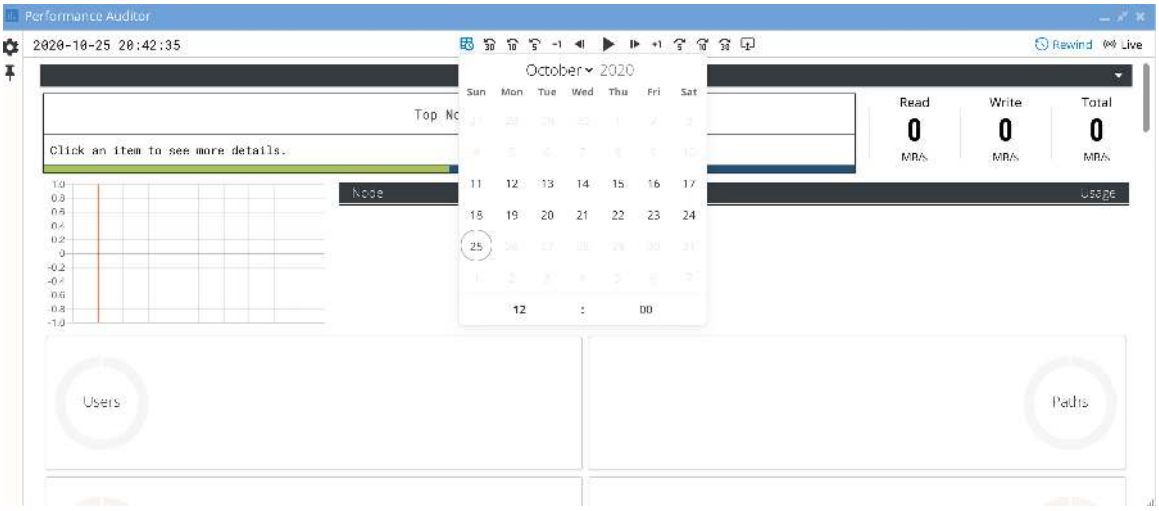

6. Share a Day and time with another administrator or save a book mark. This Share mode stores the time series index number that can be copied to clipboard and saved to import and rewind to this exact point in time in the historical data.

   a. This can be shared with other administrators to view this date time and time by using the import option. Paste the number and click import.

b. 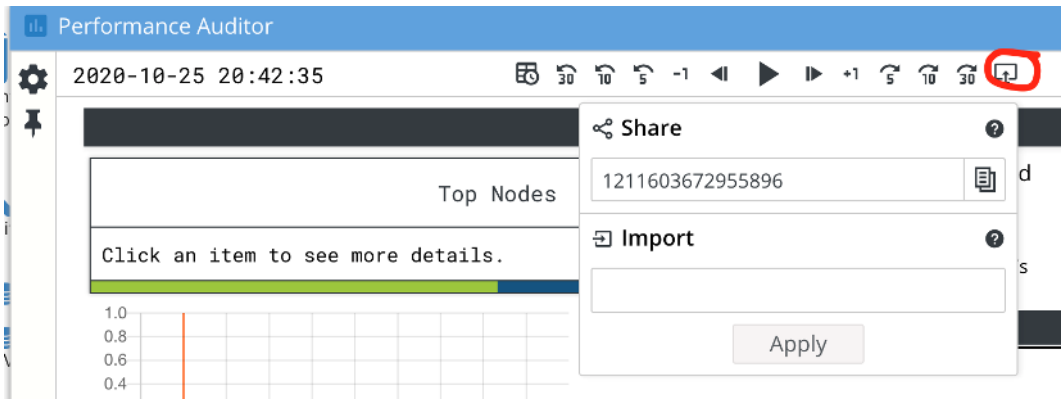

How to view long term historical summary performance data graphs

1. This data is stored in a time series database and allows graphing a subset of the performance data in an interactive graph.
2. Access this historical graphs with the History link top right of the UI. This will open a new browser tab.
3.

How to switch the UI display from one cluster to another

1. At the bottom of the display in the middle you will see the cluster name that the Performance UI is currently displaying. Click the cluster name to see a list of clusters sending statistics from the ECA cluster. **NOTE: NFS mount to the audit director is required before any performance data will display from a cluster.**

2. 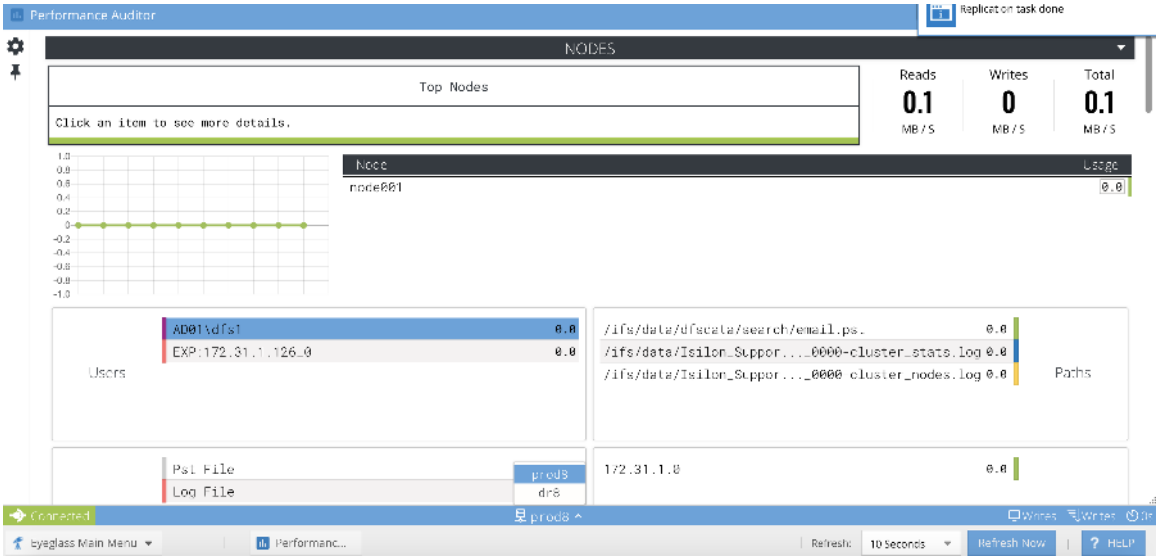

Cluster Wide Performance View Shows the Top Users, Paths, File Types, and Subnets and Category Baseline Averages

NOTE: The Average Category stat is an average for all items and not just the top 5 resource consumers displayed in the UI. Example the nodes average is all nodes in the cluster averaged over the last minute to allow comparison to the top nodes in this view.

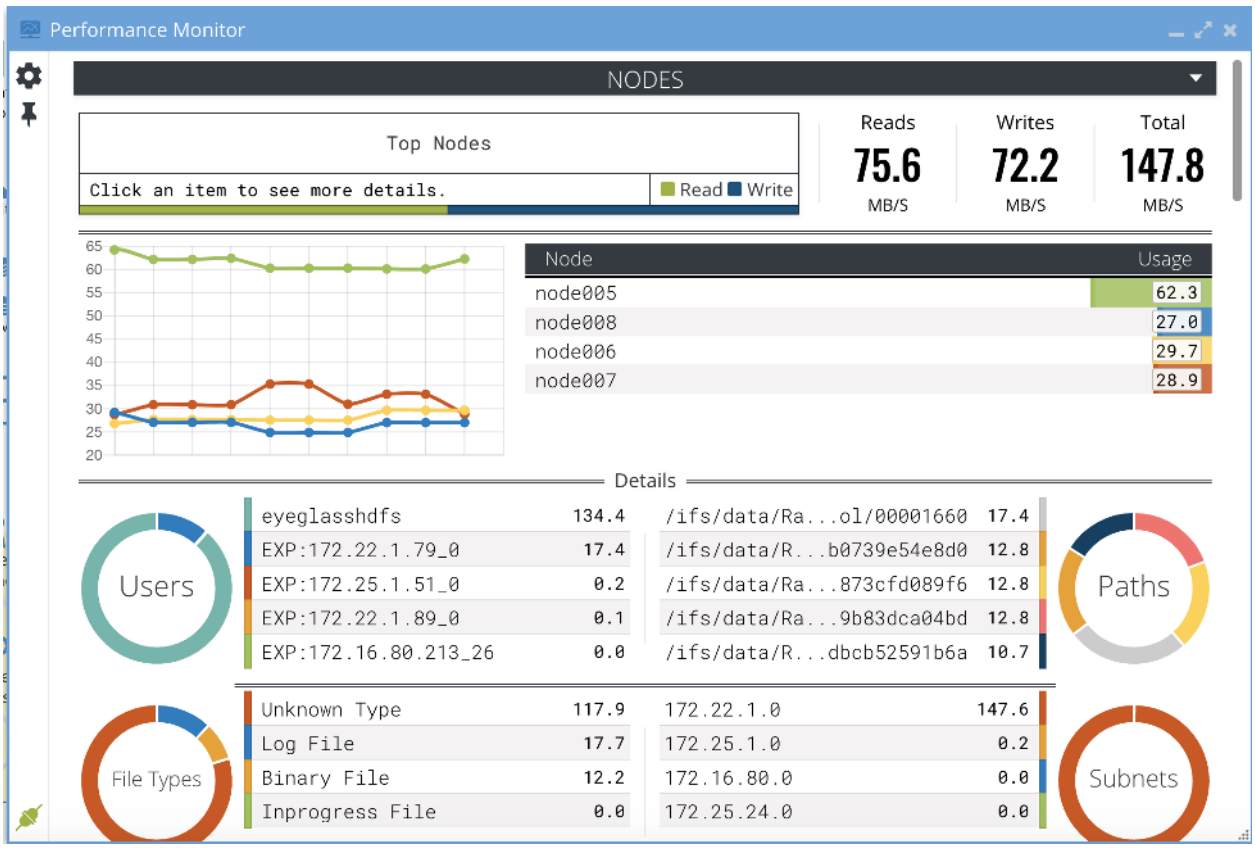

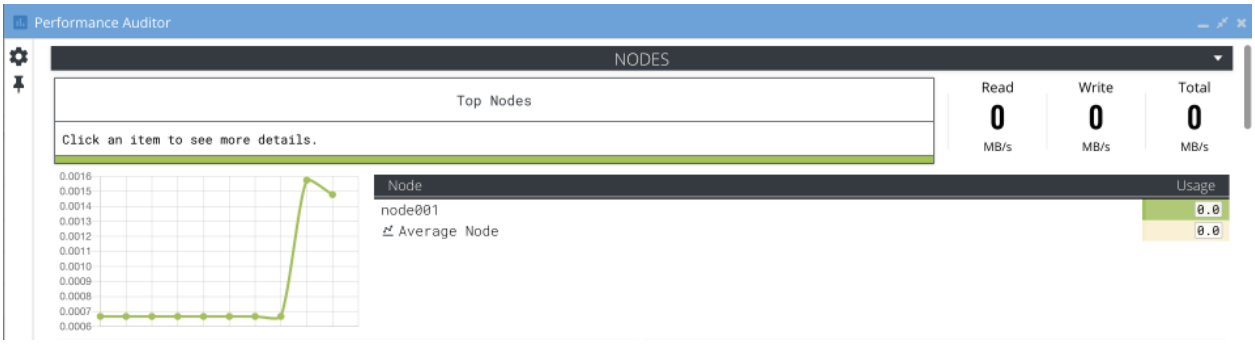

User View shows the Top Users with Baseline Average and Nodes, Path, Subnet and File Types breakouts below

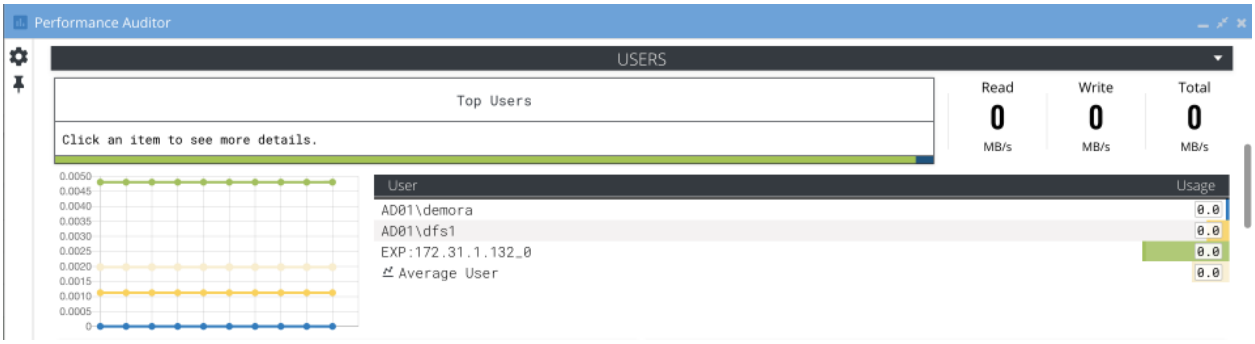

Path View shows the Top paths with Baseline Average and node, user and subnet breakouts below

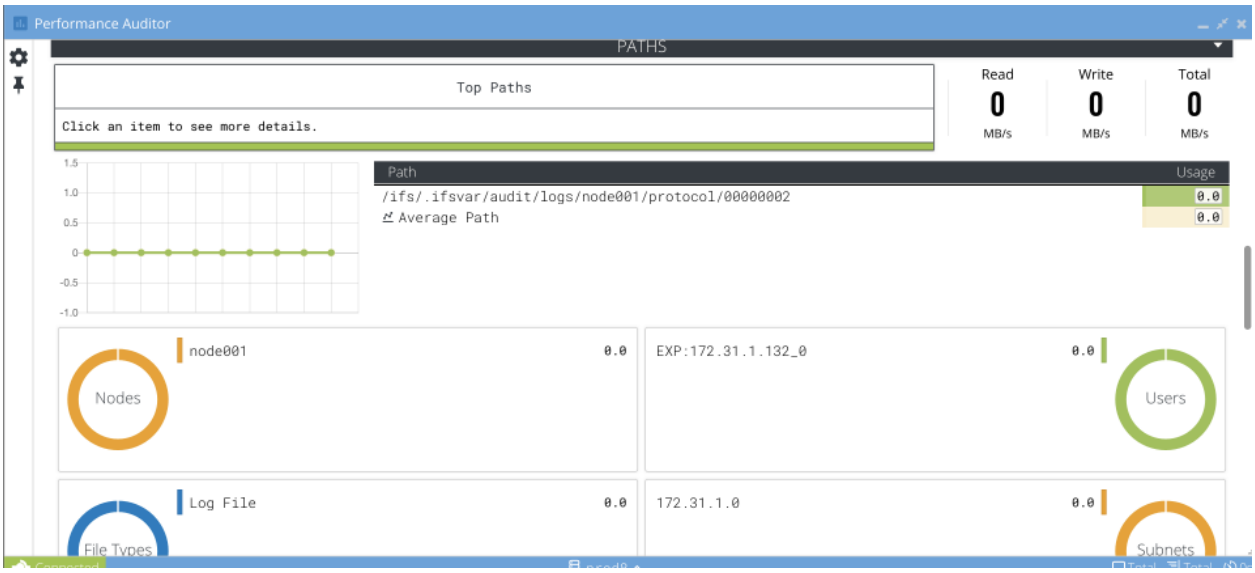

Subnet View shows the Top subnets with Baseline Average and nodes, users, and paths breakouts below

NOTE: Subnets shown are based on all source IP's of IO across the cluster nodes with a /24 subnet mask applied to show a break down by subnet. Performance Auditor has no way to know what the actual subnet mask is for your network. This view provides a relative break down of source IP's, use this view and combine this data with your knowledge of the actual subnet masks in your environment.

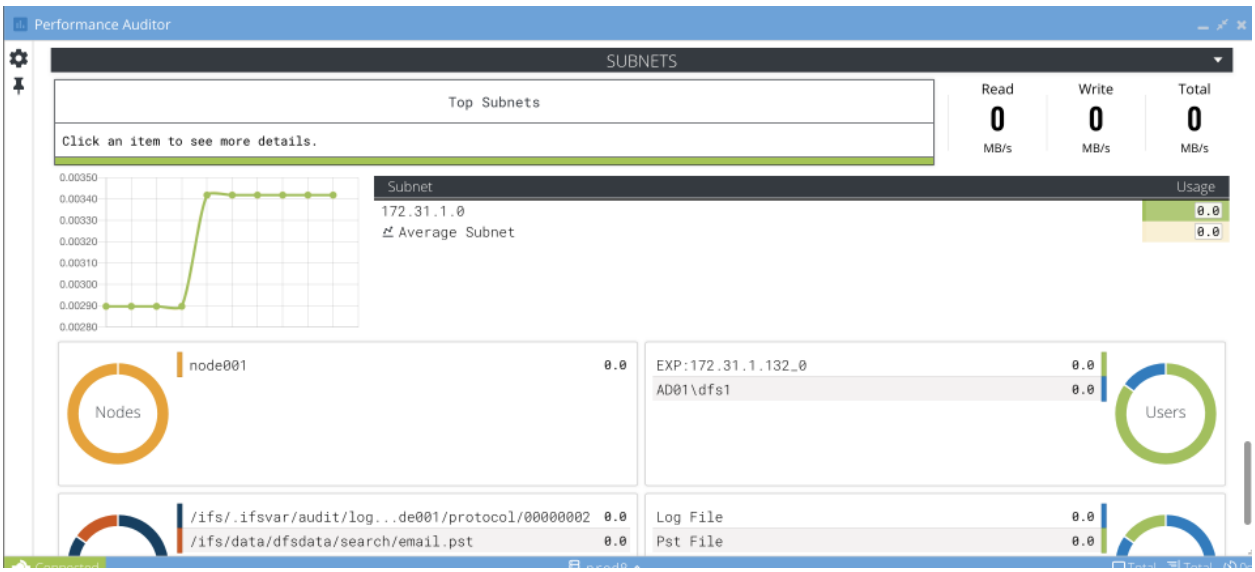

File Types View Shows the Top application types with Baseline Average and nodes, user and subnet breakouts below

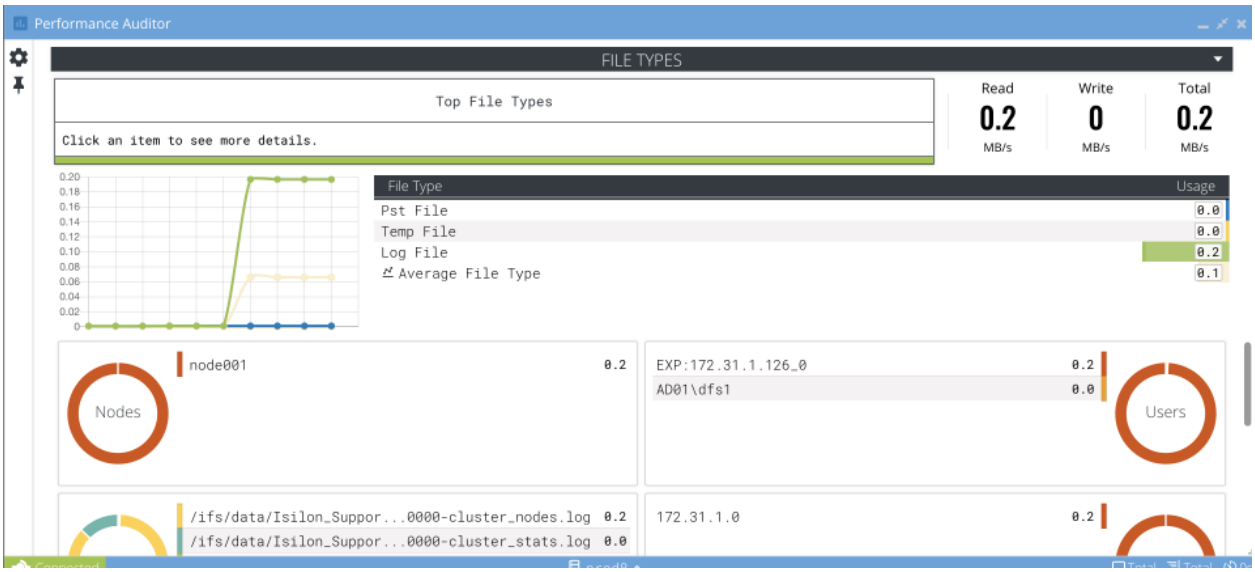

How to change the throughput rate shown for all Views

1. Click the Settings icon.

2. Sort each performance category using the sort by option.
3. Unclick the "Compute average rate over 1 minute" to switch to MB committed per minute versus the rate of throughput calculation.
4. Use the "Display Key" drop down to change the display rate for Reads, Writes or Total, and “Units” to select KB, MB or GB per second rate display in all UI's.
5. To switch the category display to show the Application Requests per second select All Requests under app Requests.
6. 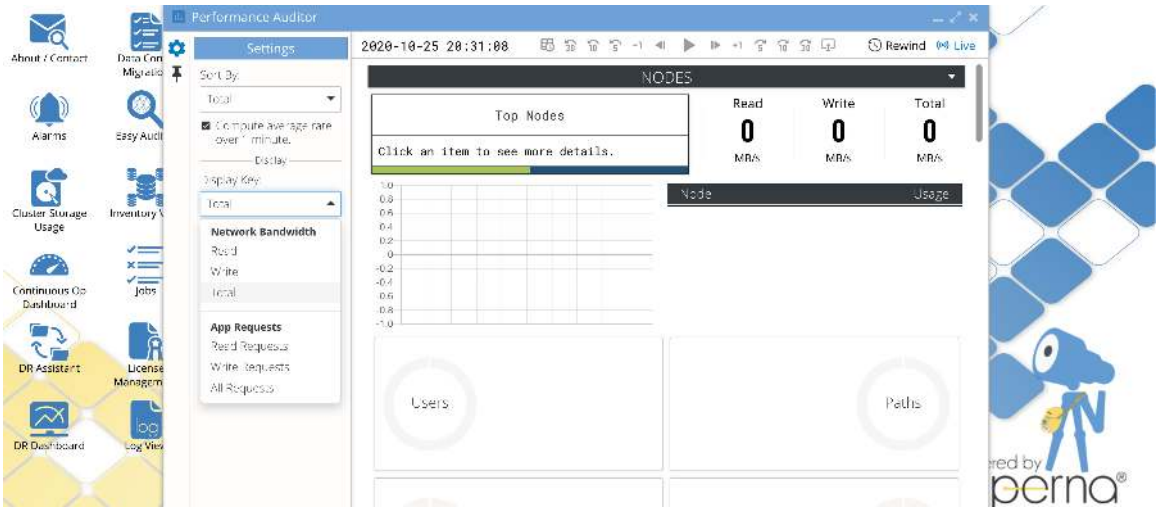

How to Switch to Application Requests Statistics

1. Click the Settings icon and select the "Display Key" drop down and select "Application Requests". All displays will now show Application Requests per second. This is similar to IOPS but is application specific IO read, write or total application requests per second over the last 1 minute sample window. This provides insight into the application IO patterns and how often the application commits data on a write or how much read ahead is being done by the application.

a. Example: A higher rate of application requests means smaller IO requests and less tolerance to network latency since more round trips are required with each application commit for a write or reading data many times per second vs caching data locally and larger reads.

2. Set the Display Key to read, writes or total to show that stat in each category view. Example to show Reads for all nodes set the Display Key to "Read" and now all nodes will show Reads for the value in the top 5 list for each area (nodes, paths, users, subnets, applications)

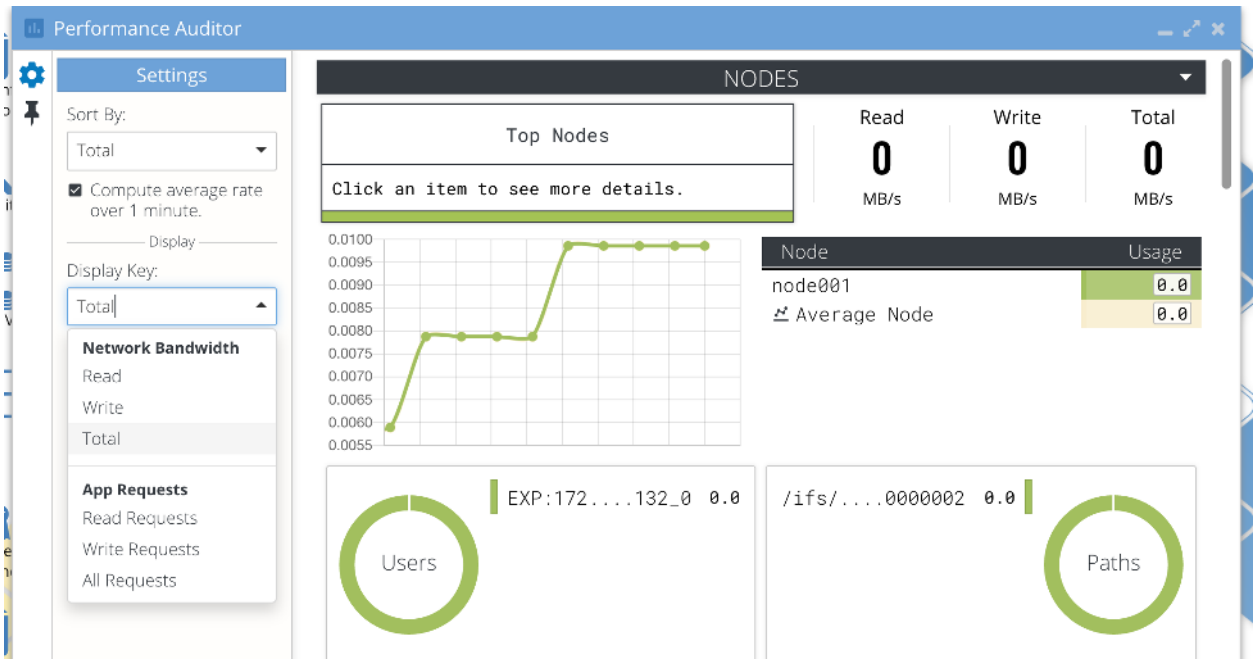

How to trouble shoot by Pinning Users, files, nodes, subnets to Performance Views

1. Use Case: A common issue is a complaint of performance issues by a user or application using NAS storage. The pinning feature is designed to allow focusing on a node, user, file or

subnet to the display how it compares with the top five resource consuming aspects of the cluster.

2. **NOTE: You can only pin 1 object in this release**
3. Example #1 User with performance issue is sharing a node with a backup application that is creating large files and using most of the node resources.
4. Example #2 a top 5 user is visible in the GUI and you want to continue to monitor this users file performance. Pinning this user to the display will keep this user's files and statistics in the display to compare to other users nodes subnets etc.. Without pinning the user may disappear from the top 5 list.

How to pin a User using with Drag and drop

1. Click the Pin on the left side of the UI:

   a. 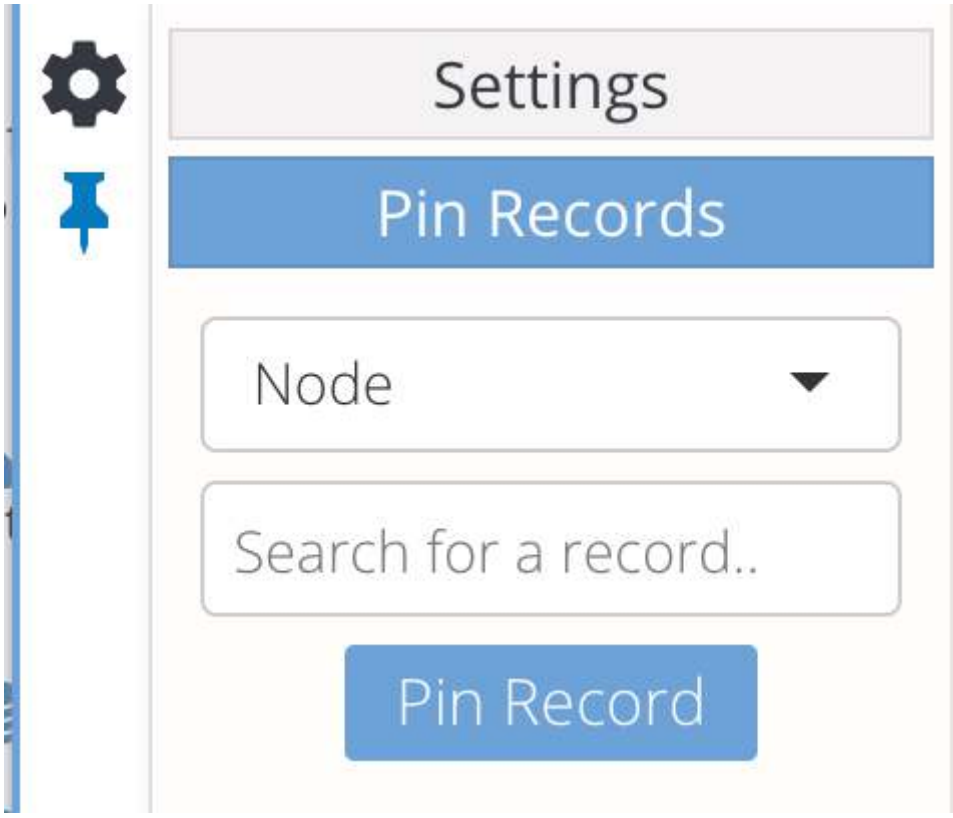

2. Select the User View:

a. 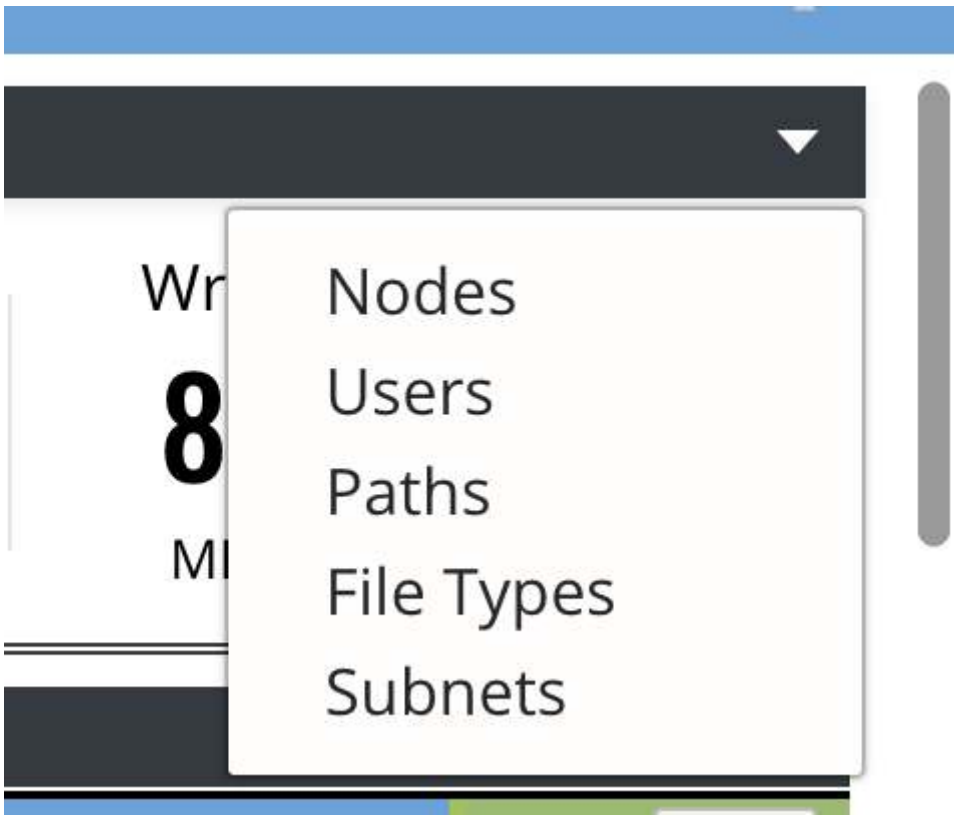

3. Drag the user from the main view over to the pin record area. It will display the pinned user:

a. 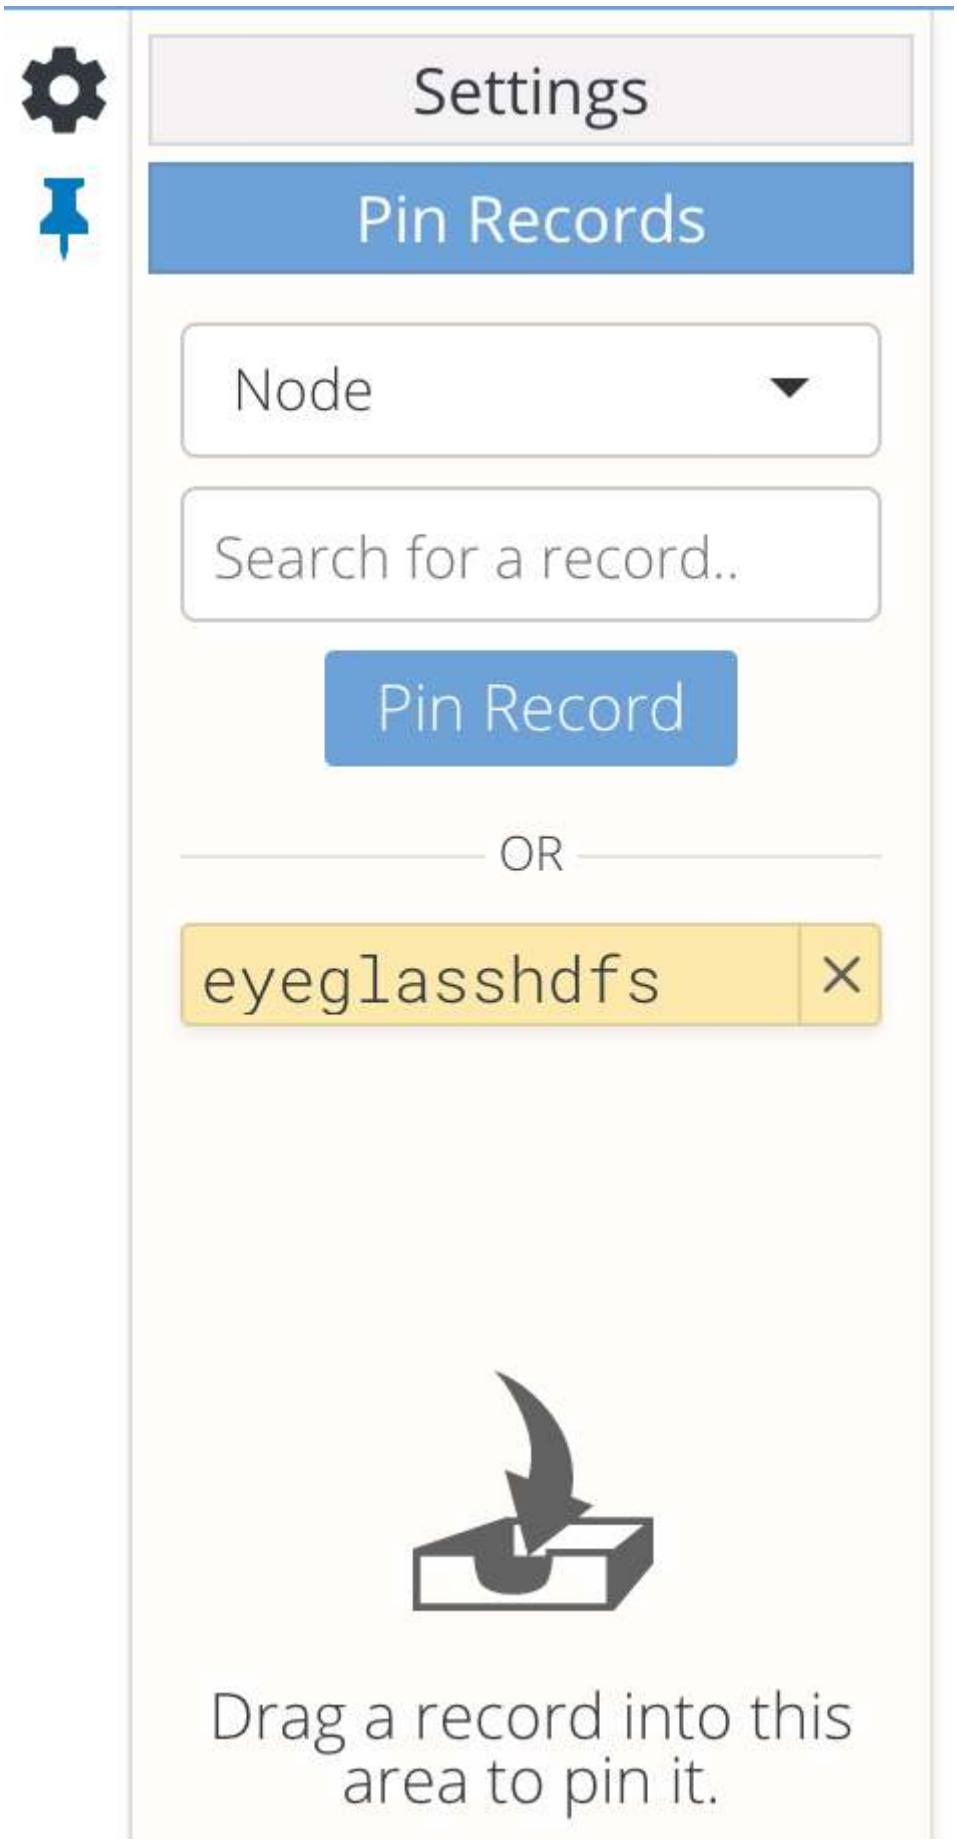

4. You can hide the pinned view by clicking on the pin icon.

5. NOTE: The pinned record will only be visible on the View for the record type. For example pinning a user will display on the user view, pinning a node will display on the nodes view. You can select any record type on any view to pin the record but the pinned record is only going to stay on the record type view.
6. Once pinned an icon indicates the record has been pinned which means it may not be a top resource consuming record but is pinned to the display for comparison to other records within the view. This allows you to monitor the record (user, node , file type,, subnet, file).

a. 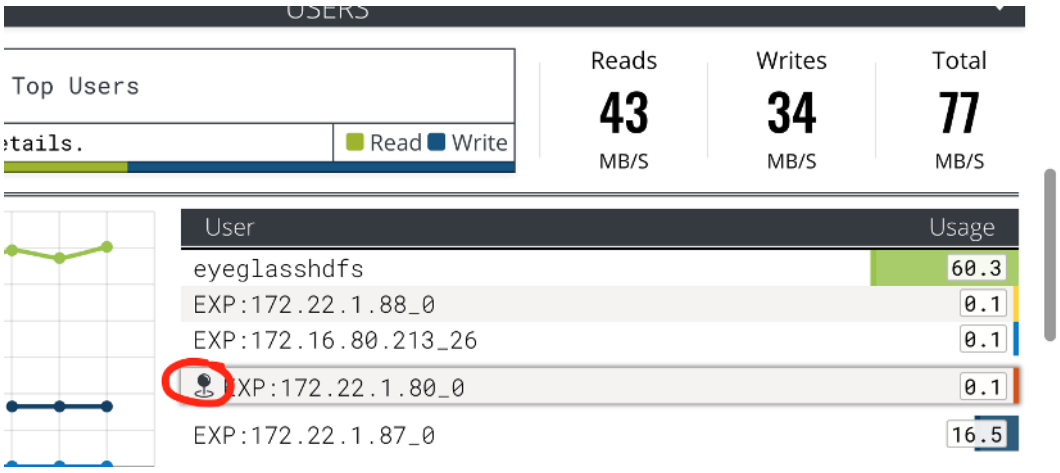

How to remove a pinned Record

1. Click the pinned icon on the left side of the UI.
2. Click the X next to the pinned record:

3. 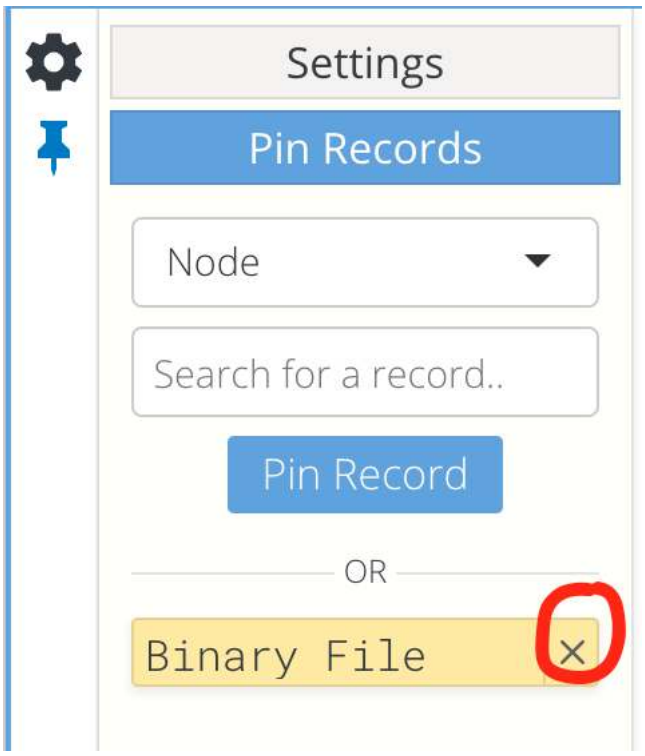

How to add a pinned record for a user, file, subnet, node application type not displayed in the UI

1. **User Pinning example:**
    a. Click the pinned icon.
    b. Select the record type selector drop down and select "User".
    c. Enter the user with "user@domain.com" syntax.
    d. If the user is resolved to a SID correctly it will be added to the display. If the user is not resolved you will receive an error message after a timeout "Error: Username 'xxxxx' could not be resolved."

e. 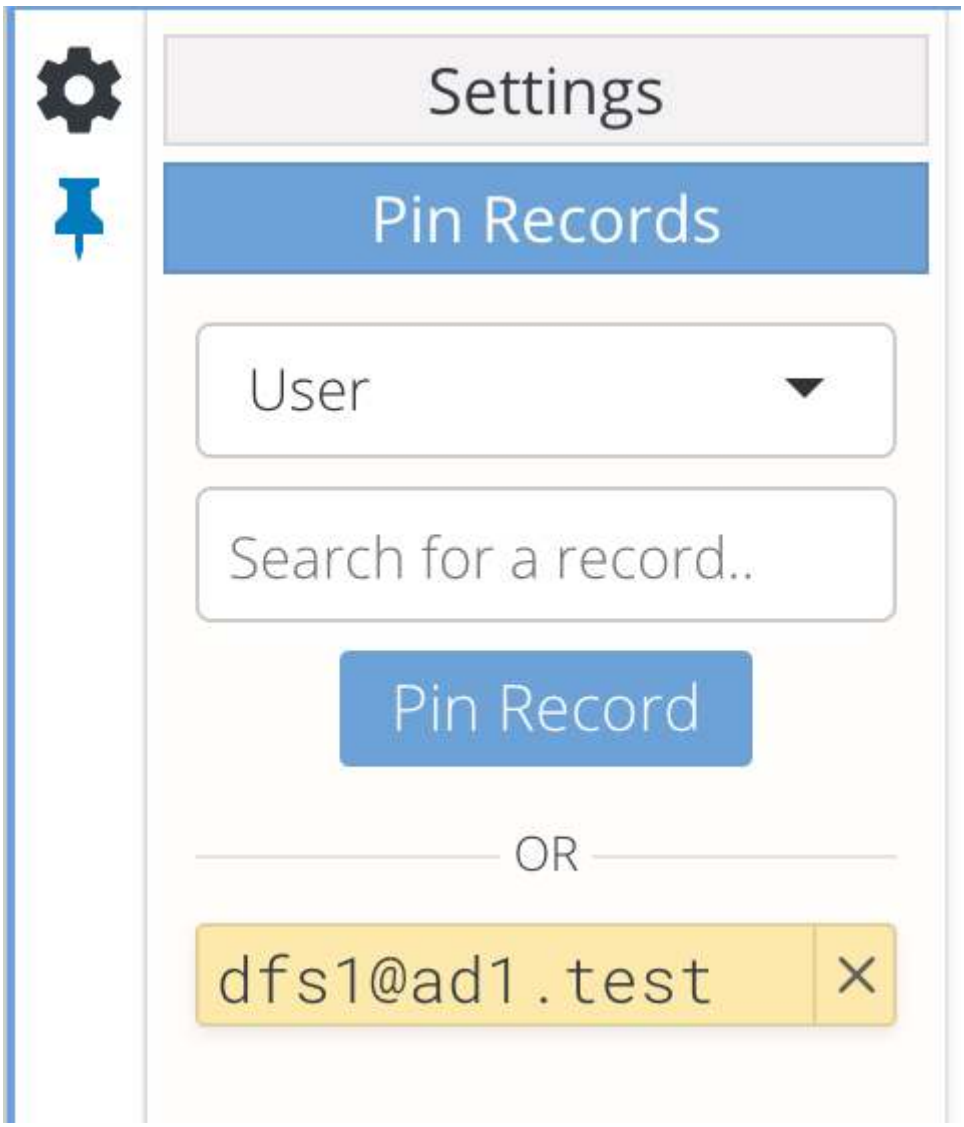

2. **Node Pinning Example**
    a. This option allows you to add a node, not listed in the top 5, to the display, and compare its usage to that of the top 5.
    b. Type the name of the node using the syntax shown similar to other nodes listed in the active view.

c. 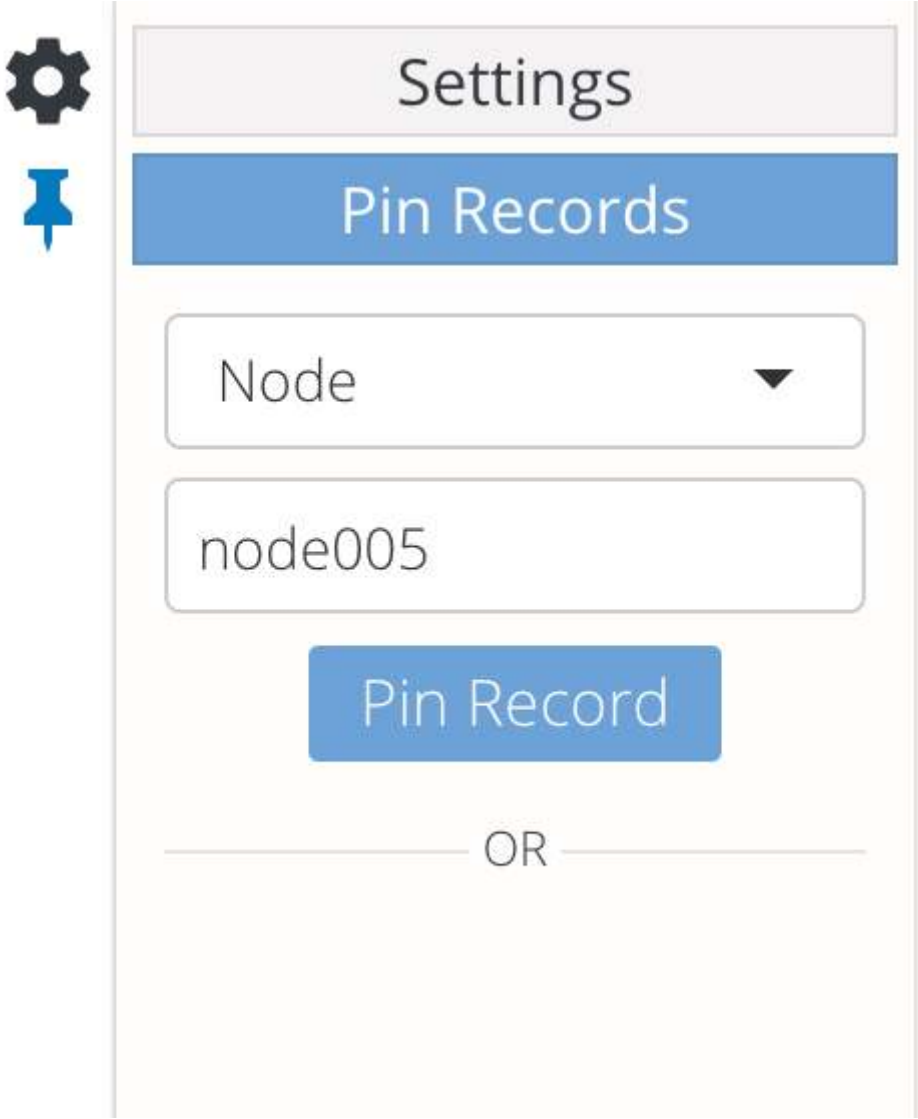

3. **Subnet Pinning Example**
    a. For Subnets use the drag and drop method or enter x.x.x.x and assume a slash 24 subnet or 255.255.255.0.
    b. Select the "Subnet" from the drop down and enter a subnet record example 172.25.24.0 this will grab all IO in the 172.25.24.0/24 subnet to display.
    c. **How to enter your own subnet: Simply enter the first 3 octets of your subnet and .0 at the end to collect all the hosts in that match 255.255.255.0 default subnet matching logic.**

d. 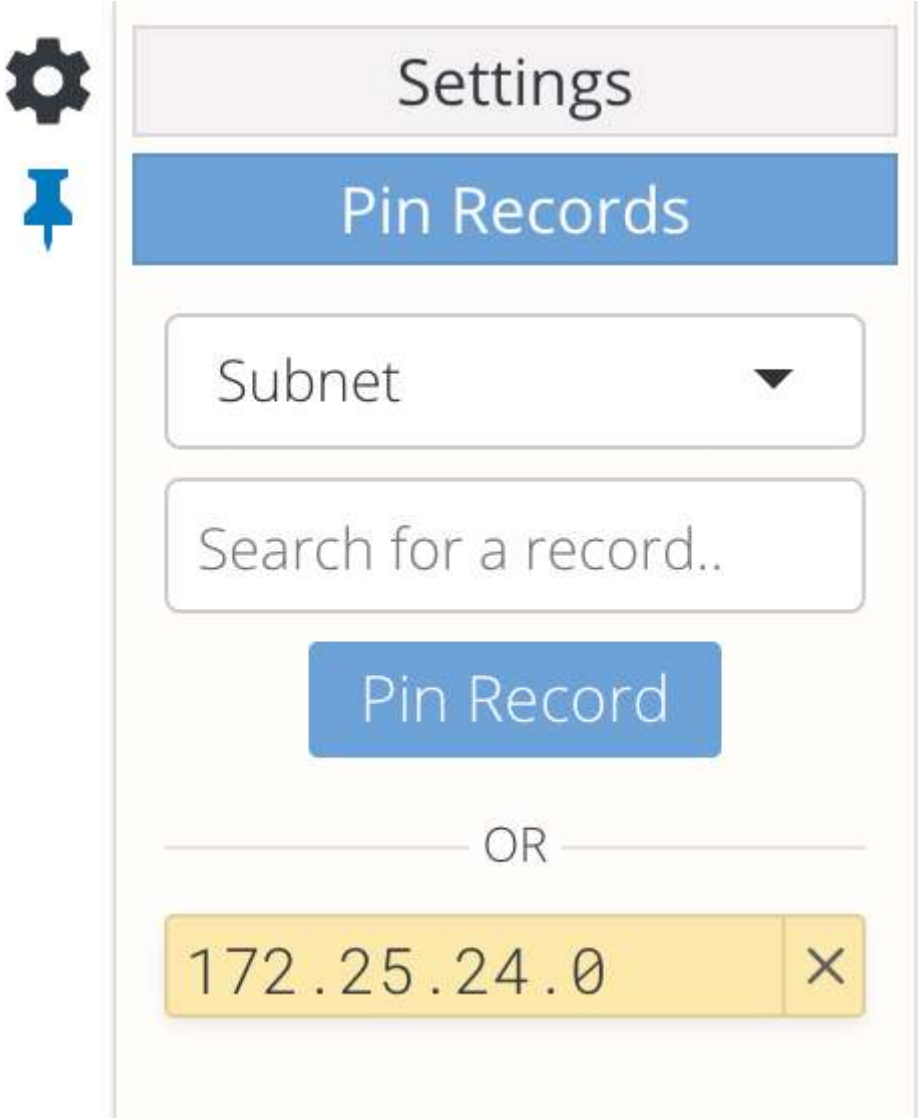

4. **Node Application Type Pinning Example**
    a. For application pinning select "File Type" option in the drop down menu.
    b. Enter the extension (i.e. pst will monitor all pst file IO across the cluster).

c. 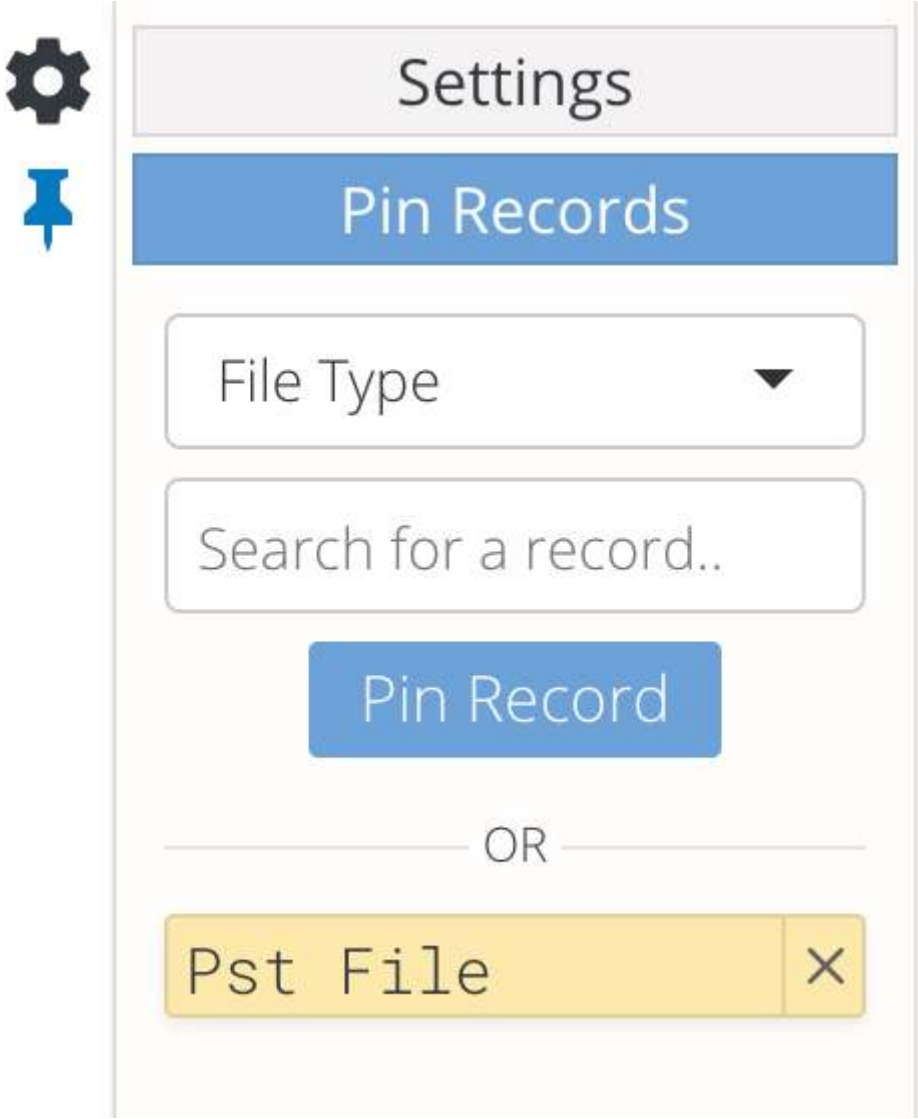

How to enable MB committed mode for long running file copy application use cases

1. This can be found in the "Settings" icon at the top left of the user display.
2. This feature switches from rate based measurement to commented data per 60 second window, which captures long running file copies or applications that creates very large files and take longer than a minute to create.

3. This will toggle all displays to show total MB of data committed per minute, which is another metric for load on nodes when applications commit data at the end of the create logic. This is true for many large file copies or similar new file create work loads like video rendering. This metric is not rate based but shows the total MB's committed by applications over the last 60 seconds .
4. The example below shows 31.6 Total MB were read in the last 60 second sample window. Each 1 second update is the total Sum of all files that were committed to storage.

How to Enable MB per minute committed mode

1. Click the Settings gear icon.
2. Uncheck the "Compute average rate over 1 minute" check box.
3. 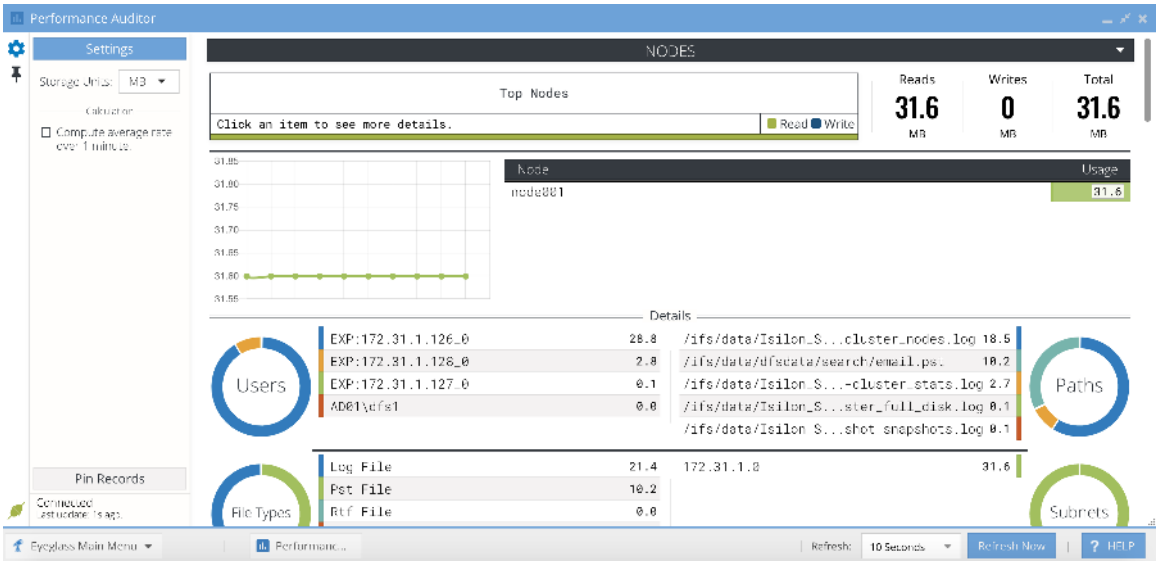

© Superna LLC

11.4. Performance Auditor Advanced Configuration

Home Top

- How to increase the file transaction monitoring window

The settings in this section allows changing how performance data is computed.

How to increase the file transaction monitoring window

1. By default file event data is monitored for a 5 minute window and 1 million open files. Contact support to recommend changes if you have long running IO workflows that last longer than 5 minutes to copy data.
2. ssh to node 1 of the ECA node 1 as ecaadmin
3. vim /opt/superna/eca/eca-env-common.conf
4. Add the variable:
    a. export EVT_REPORTER_CREATE_EVENT_RETENTION_MINS=x (where x minutes a value in minutes)
5. Save the file with :wq
6. ecactl cluster down

7. Then
8. ecactl cluster up
9. Done

© Superna LLC

12. Eyeglass AnyCopy Admin Guide

Home Top

- What's New with AnyCopy
- Installation of AnyCopy
- Planning Guide for AnyCopy
- AnyCopy Admin & User Guide

© Superna LLC

12.1. What's New with AnyCopy

Home Top

- Overview 2.5.7 Update 1
- Overview 2.5.7
- Use Cases

Overview 2.5.7 Update 1

1. **One to Many Copy jobs** - Allows content distribution from a single source cluster to many target paths on the same cluster or multiple clusters or any combination of path and cluster targets
2. **Scheduling Copy Jobs** - The ability to create a copy job and set an hourly, daily, weekly, monthly schedule on a job
3. **Retain SyncIQ policy Optio**n - This can now be disabled and still allow scheduling the job. This will create the synciq policy before each scheduled job.

Overview 2.5.7

1. The latest release offers power users storage side copy performance offload to Isilon clusters.
2. Maintains data security based on SMB share detection for end users to browse smartconnect SMB shares for data to copy between clusters.

3. RBAC with user and administrator roles to control access to reporting.
    a. Administrator role can see all copy job history
    b. User role can only see job history for jobs the created
4. Email notification list of copy jobs to update end users when a copy is finished, updates during the copy process to avoid the need to login to the UI to check on progress. Final email summary of copy performance and UNC path of the target of the data copy job.
5. slack channel support with webhook to publish copy updates to a slack channel
6. Real time monitoring of SyncIQ jobs with throughput and files progress tracking
7. Historical job run stores the last results of the copy for each execution
8. Copy audit logs stores target folder file summary of all files copied using Isilon change list API to scale tracking files using snapshot technology. SyncIQ job summary is also stored for diagnostics.
9. SyncIQ policies are created on demand and can be left behind to run over and over again or auto delete and recreate policies as needed.
10. Re run previous copy jobs

Use Cases

1. Power user offload copying large datasets powered by SyncIQ. Remove network traffic from user PC's and networks with same cluster or cluster to cluster transparent copy tool, leveraging the multi node synciq engine to move data at high speed
2. Medical research, Media and Entertainment work flows to enable researches and media data to movement of large data sets without IT assistance. Fast, secure and self serve tool to push data movement to end users while allowing centralized monitor, auditing and control of data movement.
3. Data Archive - Move data from production to archive clusters with file copy audit log for compliance and proof of what data was archived. Centralized copy tool tracks all file changes for each copy job using Isilon snapshots and changelist API feature. Track all diagnostics with SyncIQ job reports stored permanently as a compliance record of the copy jobs.

© Superna LLC

12.2. Installation of AnyCopy

Home Top

- Overview
- Prerequisites
- Installation Steps

Overview

AnyCopy is an addon product license to an Eyeglass deployment. It enables storage offload copies with role based access, audit trail, scheduling and integrated data security with smartconnect and SMB share aware user copy tool.

Prerequisites

1. Onefs 8.2.1.x for all clusters
2. Synciq and snapshotIQ License on all source and target clusters
3. Eyeglass DR manager license for all clusters to be licensed for AnyCopy.
4. Eyeglass user must be assigned the backup and restore permission assigned to the eyeglass admin role

a. 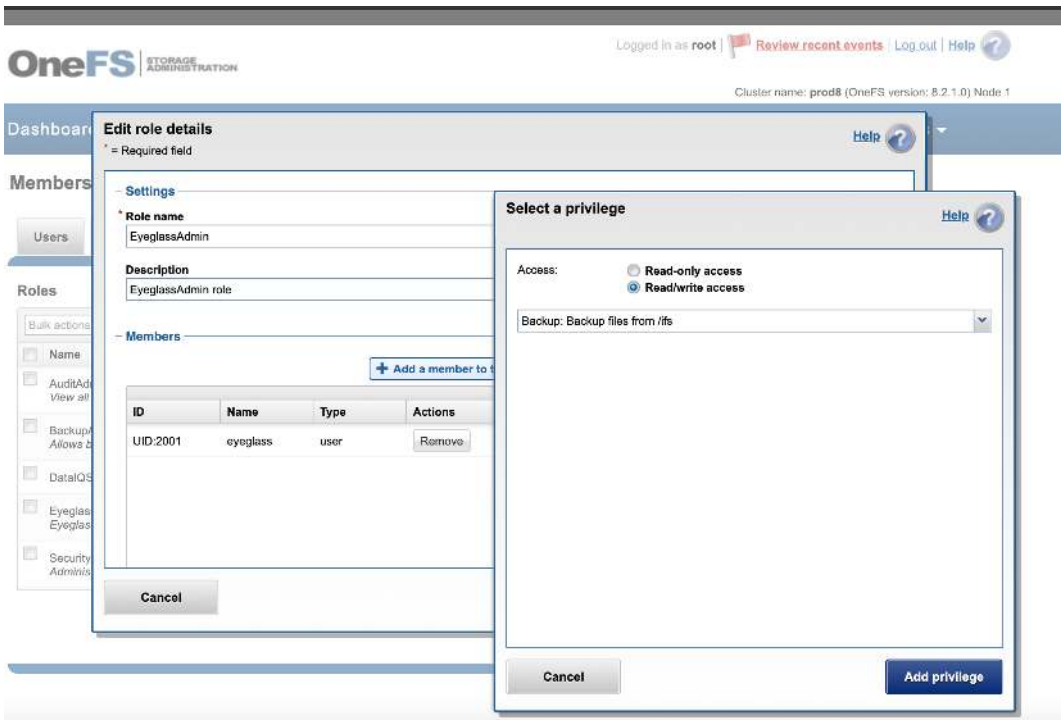

Installation Steps

1. Installation requires an Eyeglass VM deployed and clusters added following the normal see guide here.
2. Install the license key following steps here to download the key and install the key into eyeglass
3. Login to Eyeglass and assign the key using the License Manager Device tab and set the clusters to User Licensed

4. License Management

Manage Licenses | License Status of Managed Devices

License Devices

Name	Eyeglass	Ransomware	Easy Auditor	Performance Au...	AnyCopy	Isilon Probe
prod8	Auto Licensed	User License	User License	User License	User License	Auto Licensed
dr8	Auto Licensed	Unlicensed	Unlicensed	User License	Unlicensed	Auto Licensed

Submit

5. RBAC Roles

   a. 3 roles are defined View only , User and Administrator

   b. User Roles

Roles

Roles
- Administrator
- Security
- customerrole
- superna-admins
- PA Trial
- Easy Auditor
- helpdesk
- Ransomware Defender
- Read Only

All Users:

Users and groups

Users
- root
- admin
- bob

Groups

Press the ADD button above to add a group.

Permissions

View and modify AirGap jobs configuration.

ANYC_VIEW
Can view only data they have access via smb. Cannot create/edit/r...

ANYC_COPY
Can view his own anycopy jobs. Can create/edit/run anycopy jobs

ANYC_ADMIN
Can view anycopy jobs for all users. Can create/edit/run anycopy j...

Reset | Save

c. Create an RBAC role and assign roles to AD groups as needed or use builtin user admin for the admin function. All steps to create RBAC roles is listed here

6. Done

© Superna LLC

12.3. Planning Guide for AnyCopy

Home Top

- Overview
- Data Security for Copy Jobs

Overview

The AnyCopy tool allows delegation of high performance data coping between clusters or within a cluster using SMB share level security. The roles include a read only user to view copy jobs, a user role and and administrator role that can configure pre and post scripts and review all copy jobs created by users. It is recommended to start by creating RBAC roles to meet your requirements for monitoring copy jobs.

SyncIQ polices created by AnyCopy will be created in copy mode only and will be named as follows <name of copy job>_User_Copy_<unique number>. The default behavior for all copy jobs create the policy of the duration of the copy and then clean up and delete the policy afterwards. This is done to avoid leaving many policies for manual clean up. There is a GUI option to leave the policy after the job.

SyncIQ default behavior will use all nodes in the cluster when replicating data, this will require the correct firewall ports are open between any clusters that are licensed for AnyCopy. If bandwidth usage is a concern between clusters, the network bandwidth feature

on the cluster should be used to allocate bandwidth usage for AnyCopy.

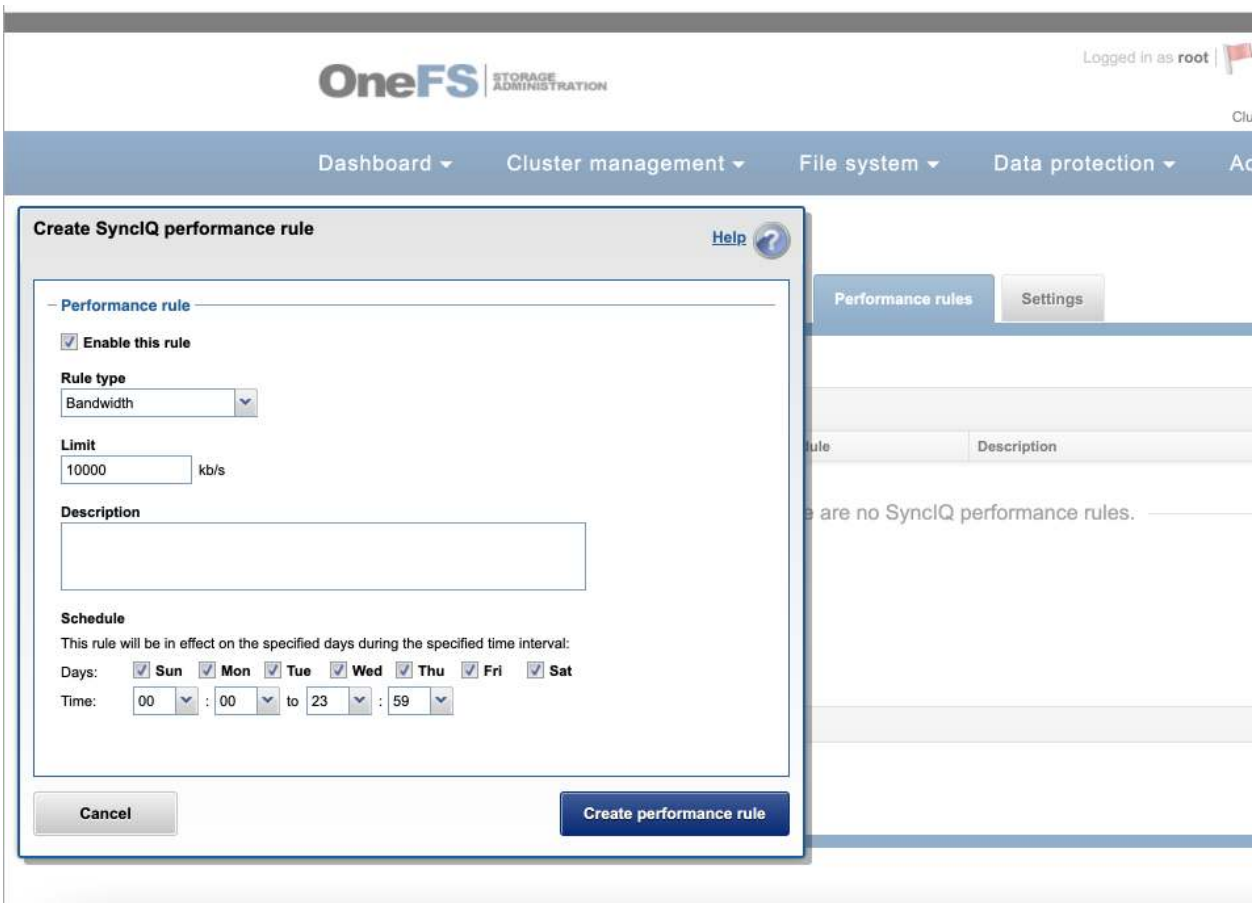

To change the default nodes used for SyncIQ on any policy that is created is possible using the SyncIQ settings tab to set the IP pool used by all synciq policies. This will allow restricting which nodes are used for AnyCopy policies that are created.

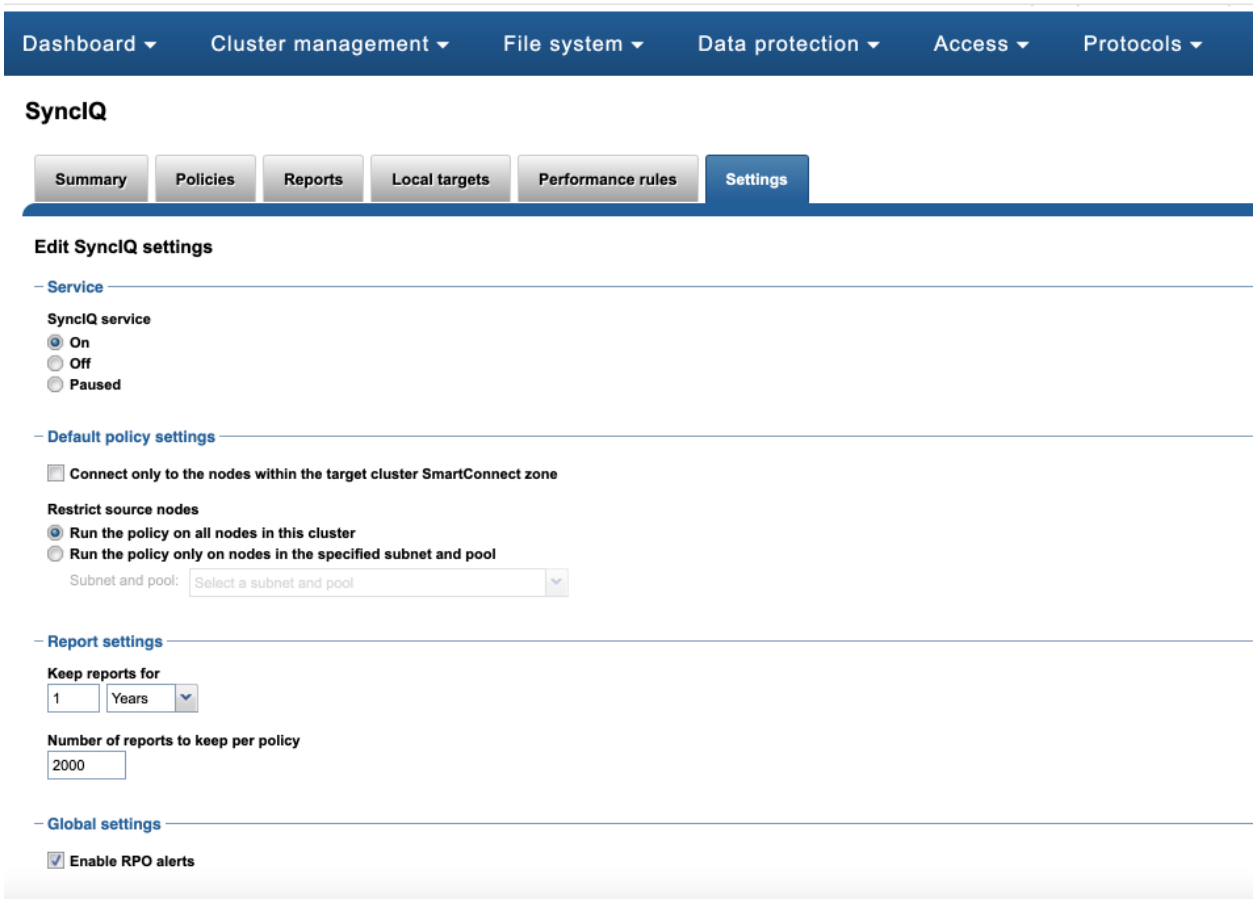

Data Security for Copy Jobs

The AnyCopy product is intended to maintain data security while allowing end users or administrators to start copy jobs using SyncIQ. This means user and administrator role within AnyCopy Eyeglass appliance must have SMB share access to any data on the source and target cluster to be allowed to configure a copy job. This maintains the data permissions already in place. This also means RBAC login is required to configure a copy job even if the user is logged in as administrator locally to the eyeglass VM. The copy job history will store the user name that created the copy job as an audit trail of the copy job creator.

If a user access to an SMB share is removed and they no longer have access to the SMB share, they will be blocked from running the job again from job history. Security is applied in all UI screens to validate access before allowing the job to run.

If everyone full control SMB share permissions are used, then the local admin user in Eyeglass will be allowed to create copy jobs. This is the only exception that otherwise require AD login via RBAC to Eyeglass to auto detect SMB groups and share level access. NOTE: NTFS ACLs are not used to restrict selection of folders for source or target paths.

© Superna LLC

12.4. AnyCopy Admin & User Guide

Home Top

- Overview
- Prerequisites
- Limitations
- Creating a Copy Job
- How to monitor a running Copy Job
- How to View Copy Job Definitions, Job History, Job edit, Pre and Post script Assignment
- How to Copy to Multiple targets - Content Distribution One to Many
- How to Edit a Copy Job Definition
- How to Schedule a Copy Job
- How to Review a Copy Job Definition
- How to Delete a Job from the Finished Jobs Tab
- Advanced Administrator Pre and Post Script UI
  - How to add a command to Pre or Post Operations with Path Variable Substitution
  - Pre and Post Script Command Examples
  - Solution Example post Script Remove Inheritance and file ACE's with Read only data
  - How to Extend the Copy Log Changelist job timeout
  - How to Test script logic
- POSIX Mode - How it Works

- Advanced POSIX Permissions Mode Configuration

Overview

This guide covers how to use AnyCopy to create copy jobs and review settings for copy jobs, copy job history, monitoring copy job progress, copy job audit logs.

Prerequisites

1. RBAC roles have been created for users and administrators of AnyCopy using an AD group and following the RBAC guide here.
2. Clusters have syncIQ and SnapshotIQ licenses
3. All Eyeglass DR & AnyCopy licenses have been assigned to the source and target clusters.
4. The Eyeglass admin role has been assigned the backup and restore permission for all clusters to be managed by AnyCopy.

Limitations

1. Pre script operations will fail the entire copy job if the pre script return code for the ssh command is > 0
2. Post script operations will complete the copy job even if the script return codes is in error
3. Copy Jobs require inventory to complete and will wait until cluster inventory is completed before starting steps.

4. The copy job file list audit log has no paging to view, for large file lists download the CSV and view in a text editor
5. The copy job file list is created using Isilon change list using snapshots before and after the copy job. If files are added to the target path before the change change list step runs, these files will appear in the the copy job report file listing even though they were not copied.
6. SMB share level permissions are used to determine source and target SMB path selection. NTFS ACL's are not used to restrict access to select a folder.
7. **File Collision handling** - If a copy job is run more than once, older files on the source path can overwrite a newer version of that file on the target. Deleted data on the source path, will not be deleted from the target, since SyncIQ is using copy mode. These limitations are SyncIQ with no work around.
8. See a full list of release note items here.

Creating a Copy Job

1. Login with your AD user name user@domain.com using the proxy login User interface. Refer to the RBAC guide link here.
2. Open the AnyCopy Icon
3. Mandatory Fields - Refer to the image below
    a. **Job name** (no special characters or spaces are allowed in the job name)
    b. **Source**:

i. open this input and select a cluster

ii. Expand a smartconnect name (DNS name used to mount the cluster)

iii. Expand an SMB share

iv. Expand the file system to select the source path that you want to copy all the data at this folder and below to the target folder.

v. **NOTE: SMB shares listed are based on the logged in users AD group membership and it will include shares with the well known Everyone permission.**

c. **Destination:**

i. Repeat the same steps as above to select the target cluster path. NOTE: You can select the same cluster as the source to copy data between 2 different SMB shares and offload the copy to the cluster.

d. **Notifications:**

i. For a long running copy you may enable email or slack channel updates by selecting the check box "Notify Progress" , enter how often updates should be sent and enter email addresses using a comma separated list of emails that should be notified about this copy job.

ii. NOTE: This is useful to notify the end user who is waiting for this copy to be completed to start a new workflow with the data.

iii. (**optional**) Slack - enter the webhook url to publish updates on progress to a slack channel.

e. **Advanced Section**

i. **Block Access to source data** - This will lock the file system before the copy starts and no users will be able to save data on the source path if this is used. It will apply a deny read permission on the SMB shares detected on the source path and below. Once the copy completes the lockout is removed from the SMB shares.

ii. Allow Scheduling - this will leave the synciq policy after the copy job finishes and allow setting up a schedule within AnyCopy to run this copy on a regular basis.

1. It also enables the option to use the on demand run now option in the Jobs History tab.

iii. Copy Share Security to Target - This will read the SMB share definitions found under thee source path and will create the same SMB share(s) on the target cluster and path with the same security options and permissions configured.

iv. Run Scripts before and after allows a scripted action to be taken on the source path or target path or both. These scripts must be setup by an administrator and should not be used unless directed by the AnyCopy administrator.

f. Starting the copy jobs by finishing the settings configuration and click **Start Copy Job.** See next section on how to monitor the copy job.

4. 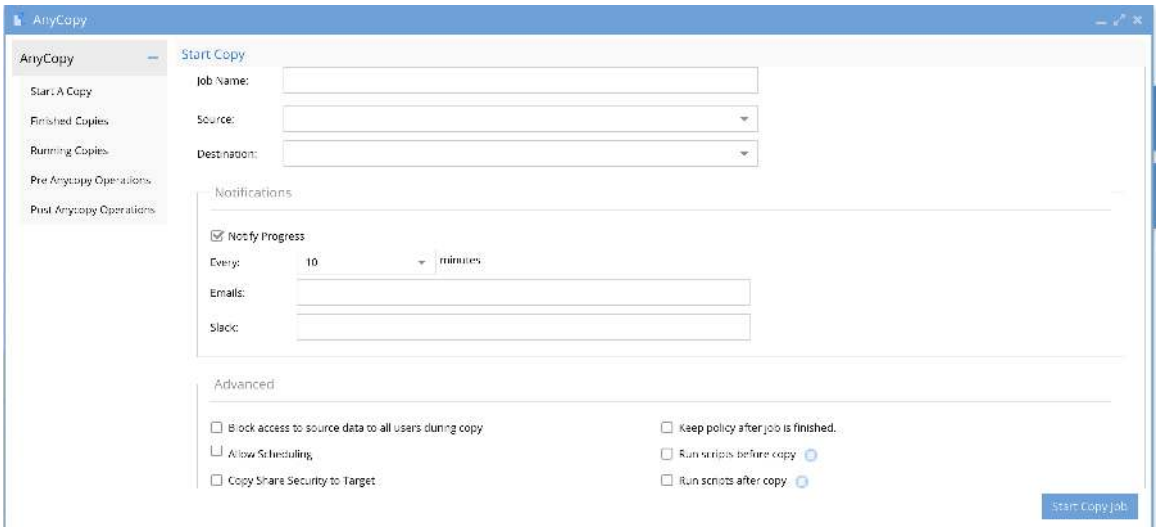

How to monitor a running Copy Job

1. Click no the running on the Running Copies tab to view running jobs.
2. Details about the job are displayed, status and copy details.
    a. Clicking on Copy Details will allow monitoring real time progress, throughput and number of files copied and GB's transferred.
    b. The lower part of the screen shows each step the copy job is executing depending on the options that were selected.

3. 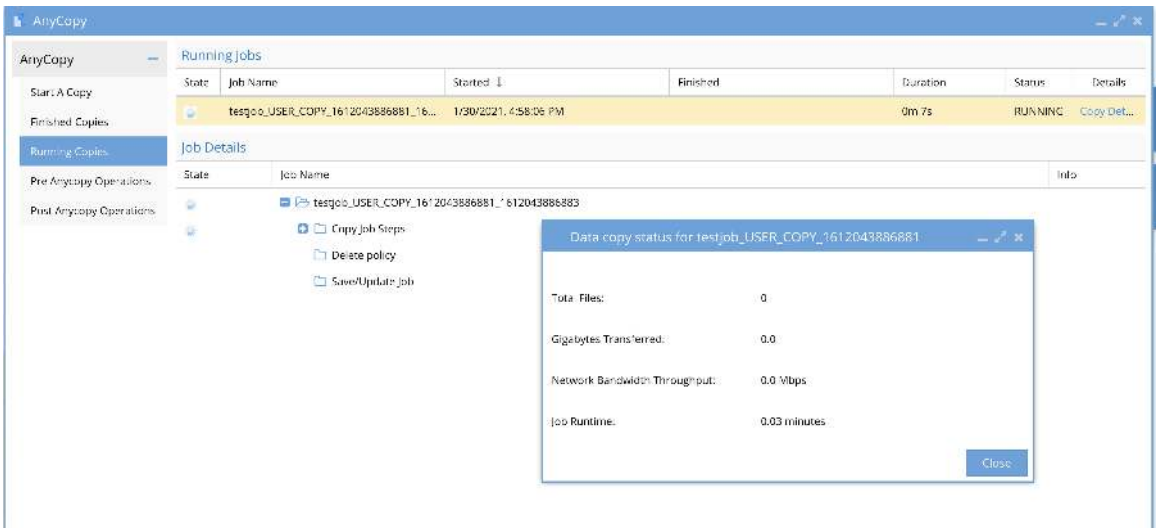

4. If email notifications was enabled you will get emails showing progress of the copy jobs.
    a. Progress emails will be sent on the interval selected in the copy job. A final summary email that shows the total files, throughput, job duration, job details and UNC path to the target cluster SMB share is sent as per the example below.

b. 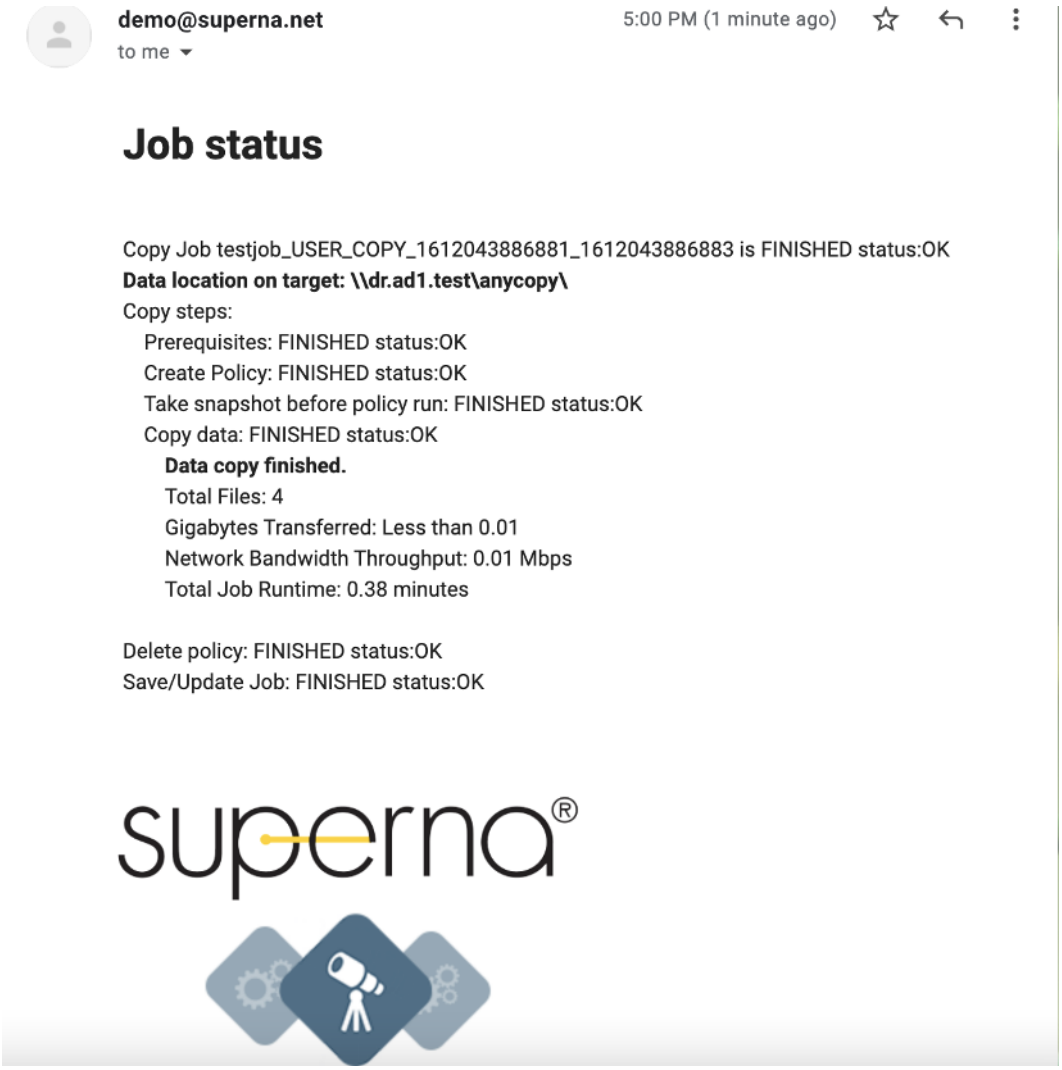

c.

How to View Copy Job Definitions, Job History, Job edit, Pre and Post script Assignment

1. Completed jobs and history of each job execution (scheduled or on demand run now) are accessible from the Finished Jobs tab. This page explains all the controls for jobs after they have run at least once.

2. 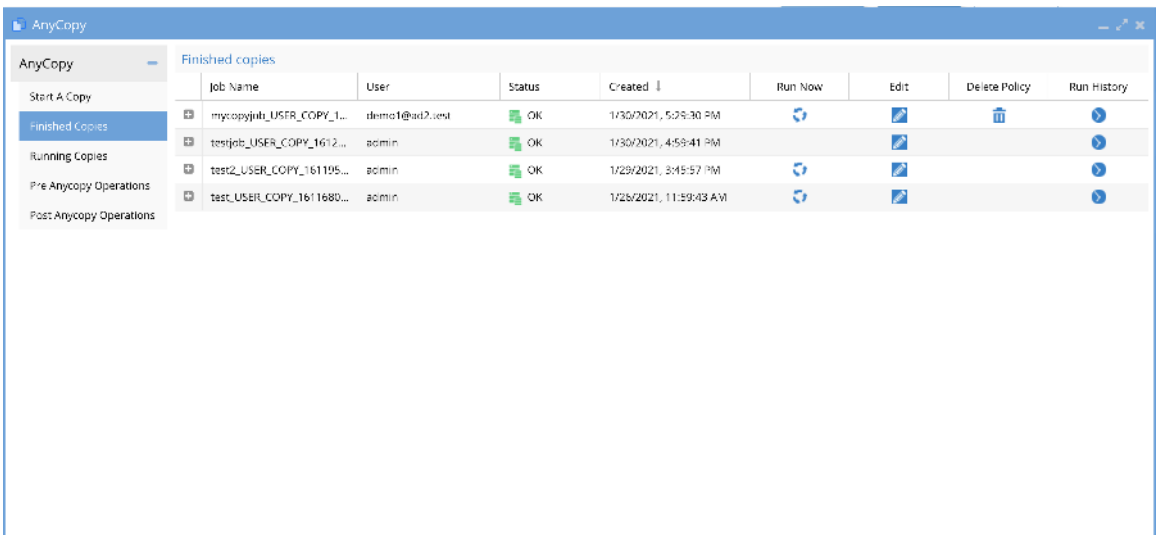

3. **Column Definitions**

   a. **Job name** - this column is the name of the job entered by the user

   b. **User** - This column lists the user name of the logged in user that created the job.

      i. NOTE: Users can only see jobs they created. Administrator role in AnyCopy can see all jobs created by all users as per the screenshot below.

   c. **Status** - This indicates if the copy job completed successfully without error and will show OK or a red error.

   d. **Created** - This is the date and time the copy job was created.

   e. **Run Now -** This column will show an icon only if the Schedule policy option was enabled, which will leave the policy on the source cluster after the copy finishes. Pressing the run now option will start the job again based on its defined settings.

i. Any job without the icon, indicates the synciq policy was deleted and the job cannot be run on demand.

ii. NOTE: If the job definition is edited to allow scheduling, then the run now icon will still show the icon and running the job will automatically recreate a new SyncIQ policy if it was previously deleted based on the job definition settings.

f. **Edit** - This opens a dialog box to allow editing the job definition and re-saving it, this is also where the job schedule is configured. The edit options are the same as the job create settings user interface. Make any needed changes and click Save job Configuration button.

i. On the edit page the option to select pre or post job operations check box(s) . Selecting either pre or post and the icon next to the option will provide selection of operation labels configured by your administrator.

ii. 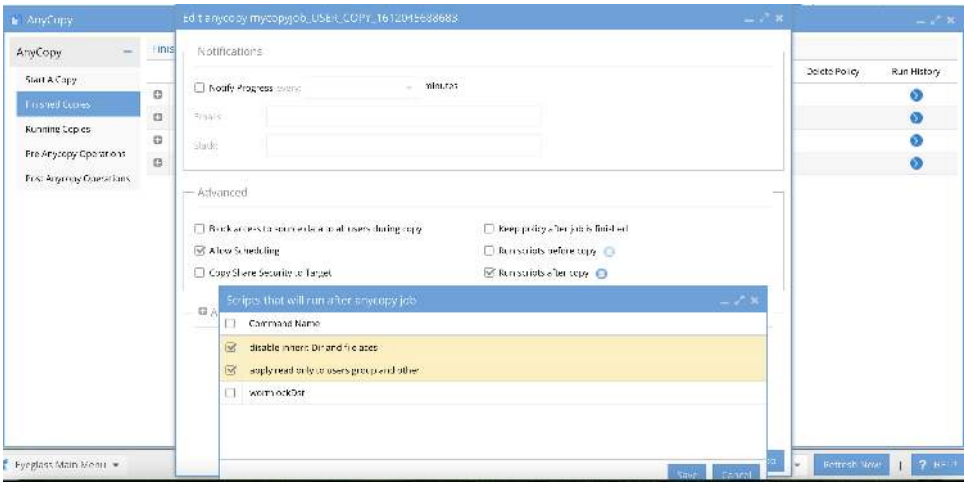

g. **Delete Policy -** This column will show a trash can icon only if the policy was left behind and still exists on the cluster.

i. Pressing this icon will immediately delete the synciq policy from the cluster and will display a confirmation box that this will be updated in inventory after the next schedule inventory job completes.

ii. After the next scheduled configuration inventory job runs in Eyeglass the trash can icon will disappear.

iii. Running the on demand option will recreate the SyncIQ policy again.

h. **Run History** - Each execution of a job via a schedule or on demand run now will create a run history with the date and time the job was run and 2 columns with each one providing access to logs.

i. 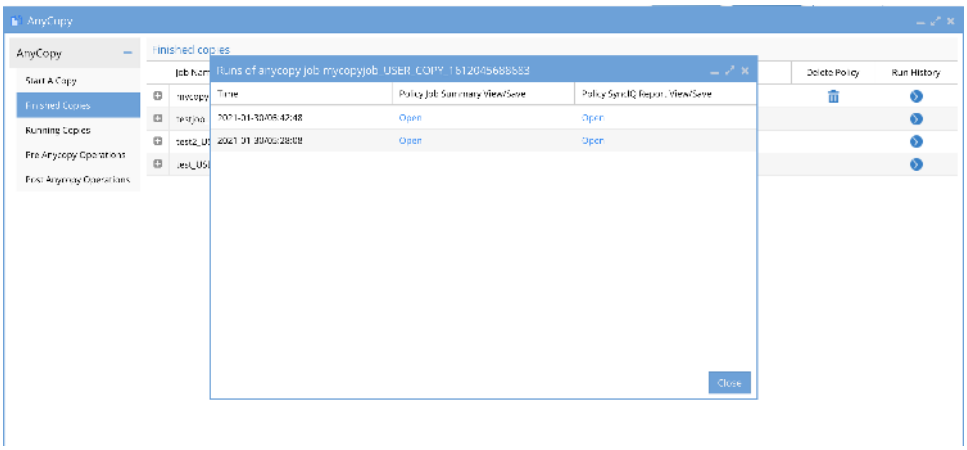

ii. **Time** - Shows the date and time of each job execution

iii. **Policy Job Summary** - will provide an audit log that lists all files copied to the target folder path.

1. This file list is created using the Isilon changelist feature, where a snapshot is taken on the target path before the copy is started. After the copy is finished a 2nd snapshot is taken and the cluster API provides the differences which is the file list.

2. 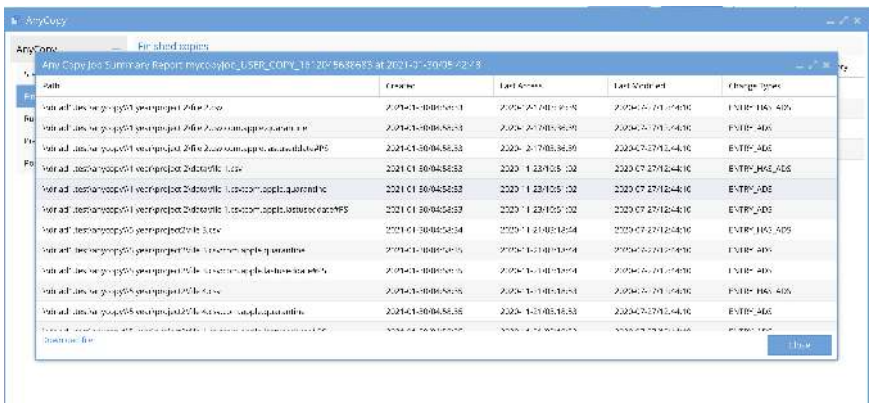

3. The file list is provided as UNC path to the files using \\smartconnect name\SMB share name\path to file.
4. **Date Stamps of the files** - created, modified, last accessed
5. **Action taken** - Entry ADS means file was added to the target folder.
6. Option to download this file to a CSV is available at the bottom of the window.

iv. **Policy SyncIQ Report** - This shows the diagnostic log captured from SyncIQ that is useful for support debugging or to align the SyncIQ job settings used on the SyncIQ policy. It can also be used with Dell support if errors occur.

1. The file also has a download to txt file option in the UI.

How to Copy to Multiple targets - Content Distribution One to Many

1. Requirements: 2.5.7 update 1

2. When creating a new job definition simply click in the target drop down arrow to add additional targets. Each target can be the same as the source cluster or a different cluster or any combination of local and remote targets.
    a. This allows copying data to multiple shares on the same cluster or to remote clusters or both.
3.

4. Fill in the remaining job definitions as normal including the option to schedule this copy job on an interval.

How to Edit a Copy Job Definition

1. NOTE: A copy job must run once before you can edit the configuration.
2. NOTE: An administrator can edit other user jobs. Users can only edit job definitions they created themselves.
3. Click the Finished Jobs tab
4. Click the edit icon for the job you want to edit the configuration.

5. 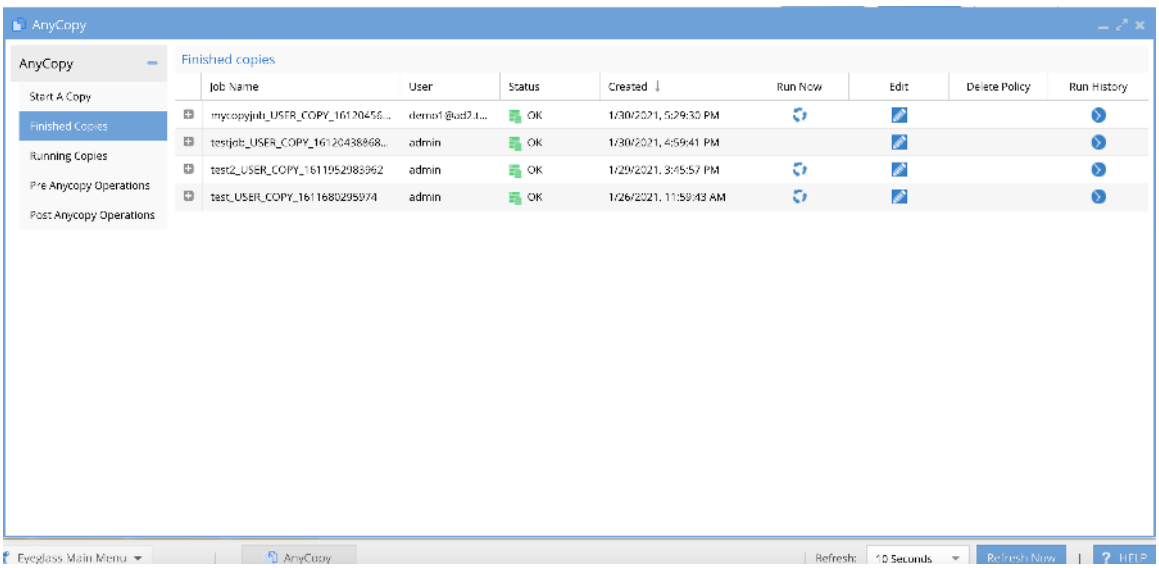

How to Schedule a Copy Job

1. Requires 2.5.7 update 1
2. NOTE: An administrator can edit other user jobs. Users can only edit job definitions they created themselves.
3. When creating a new job definition Click the Allow Scheduling checkbox and set the interval. You can also uncheck keep policy after job is finished. This will delete and create the synciq policy each time the job is scheduled. This is useful if the target directory contains data that you want to write to in-between copy jobs.

4. 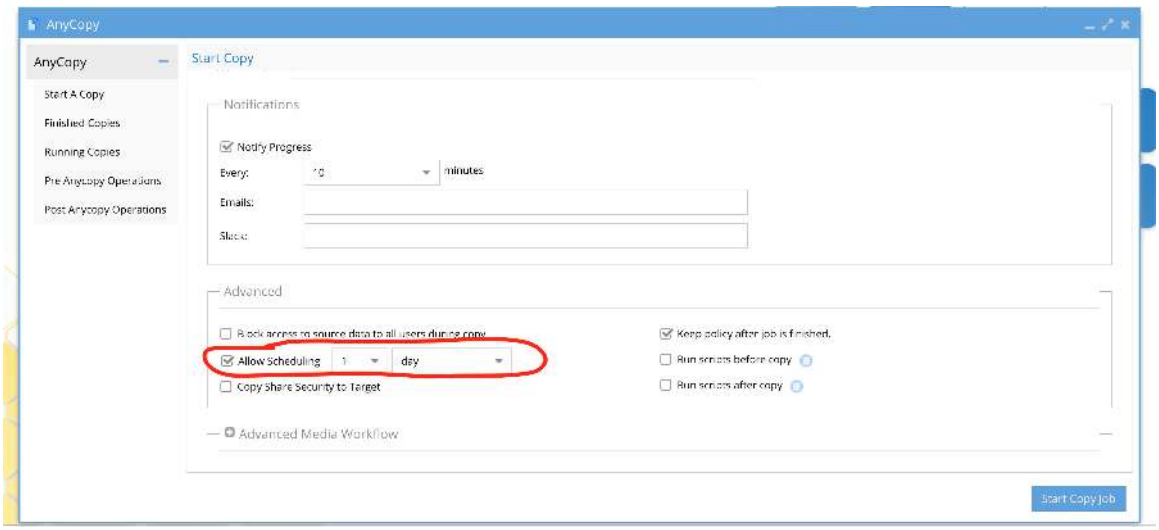

5. Complete the filling in the job details.
6. To edit the schedule at a later time, follow the How to edit a job definition in this guide.

How to Review a Copy Job Definition

1. From the Finished Jobs tab the list of jobs shows columns as described above. To review the configuration summary of a job definition. Click the + sign to expand the job details.

2. 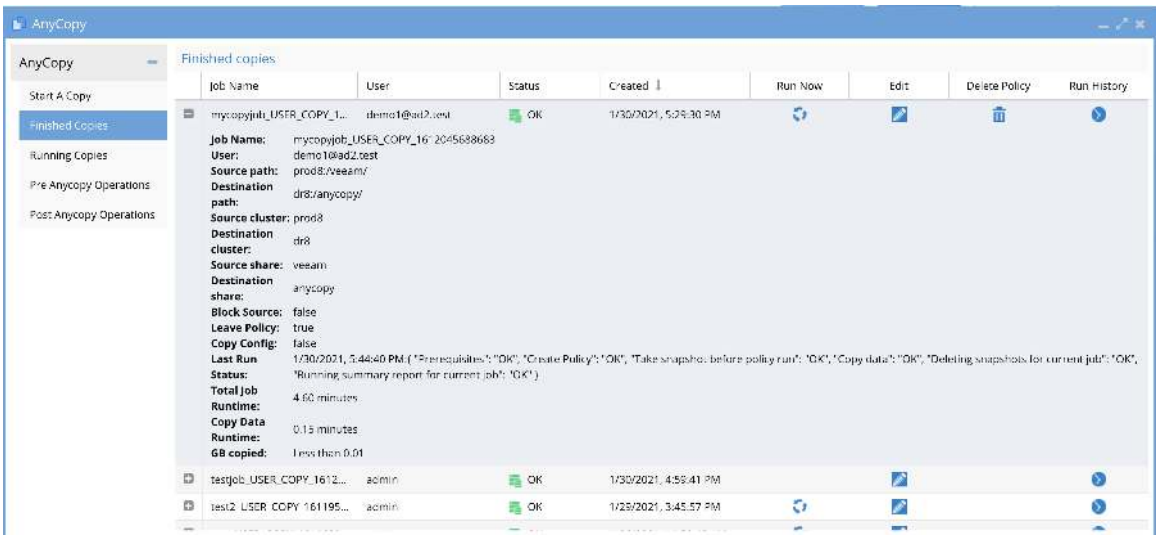

3. The options will match the Job configuration or edit screen and are summarized on this screen.
4. The screen also shows the "Last run" details and summarizes if the last run was successful, the job run time, and GB copied. This provides a quick way to see when the last job ran and the details without using the job history view.

How to Delete a Job from the Finished Jobs Tab

1. Open the AnyCopy Icon
2. Click the finished job tab
3. Use the Delete job trashcan next to the job you want to delete. NOTE: This cannot be undone.
4.

Advanced Administrator Pre and Post Script UI

1. **Overview**:
    a. Pre and post script options allow create commands to run on the source or target path or both. The commands can be ordered and configured as pre or post.

b. The pre and post commands are a library of functions that can be selected by users using a command label assigned to the commands.

    i. When users enable pre or post scripts they must also select the commands from the library to apply to their job. They will only see the command label and not the command itself in the UI.

    ii. If the user is instructed to use more than on command the order of execution will be managed by the administrator who orders the commands in the pre and post administration UI.

c. Only AnyCopy Administrators can create pre and post script operations for users to use

d. The pre and post commands execute on the Isilon over ssh only Linux commands that run on Isilon can be used. You can use ISI commands. The command should be single line command to execute all logic or issue multiple commands.

e. For more complex logic it is possible to create a scripts directory on the Isilon file system and reference a script name for more complex logic that requires a simpler way to maintain the logic within the bash script and simple reference this script name in AnyCopy UI.

2. **The pre AnyCopy Operations tab** - This is where new commands can be added for pre copy steps to execute.

3. **The Post AnyCopy Operations tab** - This is where new commands can be added for post copy steps to execute.

4. Path substitution using variables
    a. The following variables can be used to insert an absolute path into your pre and post commands
    b. $DSTPATH - destination absolute path example /ifs/data/projectA is the destination copy target and this path will be replaced if this variable is used
    c. $SRCPATH - source absolute path example /ifs/data/projectA is the source path on the source cluster configured in AnyCopy.

How to add a command to Pre or Post Operations with Path Variable Substitution

1. 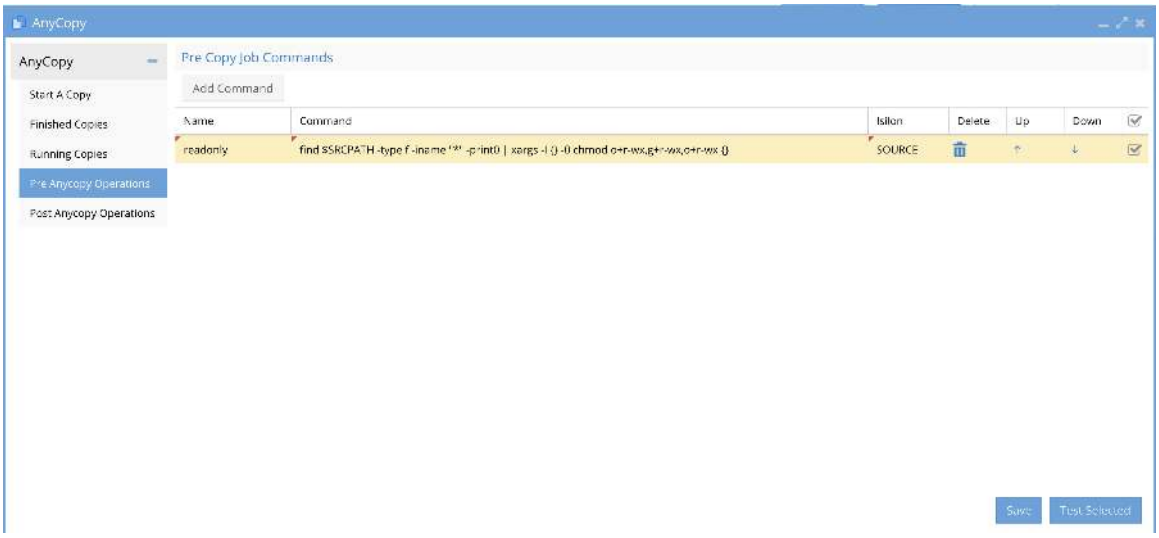

2. **Columns**:
    a. Name - this is a name or label on the command to provide an idea of what it does. no special characters or spaced can be used.

b. **Command** is a single line command that will be executed over ssh and should use recursive flags to process a change across all files and folders below the directory.

c. **Cluster** - The command can be run on the source or target cluster regardless of the pre or post setting. Select source or target for this command.

d. **Delete** - This option allows deleting a command from the pre or post list of commands

e. **Up/Down** - If multiple commands are used the up and down arrow orders the commands with the command listed at the top of the UI being executed first. Re-order the commands based on the correct order of processing needed by your script logic.

f. **Selection Check Box** - This check box allows testing your script logic by selecting the command and then clicking the **Test Selected** button.

Pre and Post Script Command Examples

These are examples to lock down data after a copy to remove all directory and file permissions and inheritance.

1. **Apply read only to users group and other**

    a. This command example will set the permissions on the source files to read only for Other and will try to set Owner,

Group to read only. NOTE: Directory inheritance may need to be disabled first. **Using chmod u+r,g+r,o+r**

b. Uses $DSTPATH variable

c. Recursive for files only

d. sleep 2; find **$DSTPATH** -type f -iname "*" -print0 | xargs -I {} -0 chmod u+r,g+r,o+r {}

e. Cluster target

f. Used as post script

2. **Remove inheritance and remove file ACES**

a. This command disables directory permissions from flowing down the directory tree so that permissions can be customized after the copy is completed. This examples uses 3 commands on 1 line with a semi colon to allow each to run using a single line. This can be used to combine steps into a single command entry

b. It applies a block of inheritance with +dacl_protected , removes directory ACL and synthetic auto inheritance

c. It removes all ACE entries from all files using **chmod -D**

d. Uses **$DSTPATH** variable

e. Recursive for directories only and recursive for files to remove all ACE entries using **chmod -c +dacl_protected** and **chmod -c -sacl_auto_inherited and chmod -c -dacl_auto_inherited**

f. Cluster is set to target

g. Used as post script

h. sleep 10; find **$DSTPATH** -type d -iname "*" -print0 | xargs -I {} -0 **chmod -c +dacl_protected** {};find **$DSTPATH** -type d -iname "*" -print0 | xargs -I {} -0 **chmod -c -dacl_auto_inherited** {};find **$DSTPATH** -type d -iname "*" -print0 | xargs -I {} -0 **chmod -c -sacl_auto_inherited** {} ; find **$DSTPATH** -iname "*" -print0 | xargs -I {} -0 **chmod -D** {}

Solution Example post Script Remove Inheritance and file ACE's with Read only data

1. See how the above examples were used to create a solution to remove all permissions and inheritances of file and folder permissions.

2. 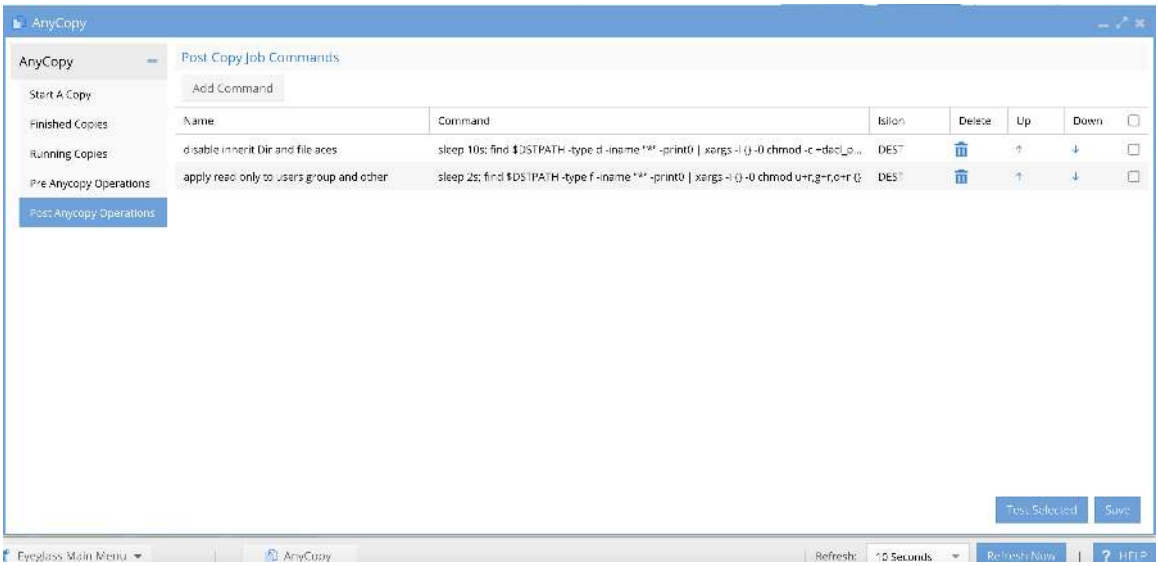

3. Configure the Copy job to select the 2 labels configured in the post script operations tab and ordered as shown.

4. 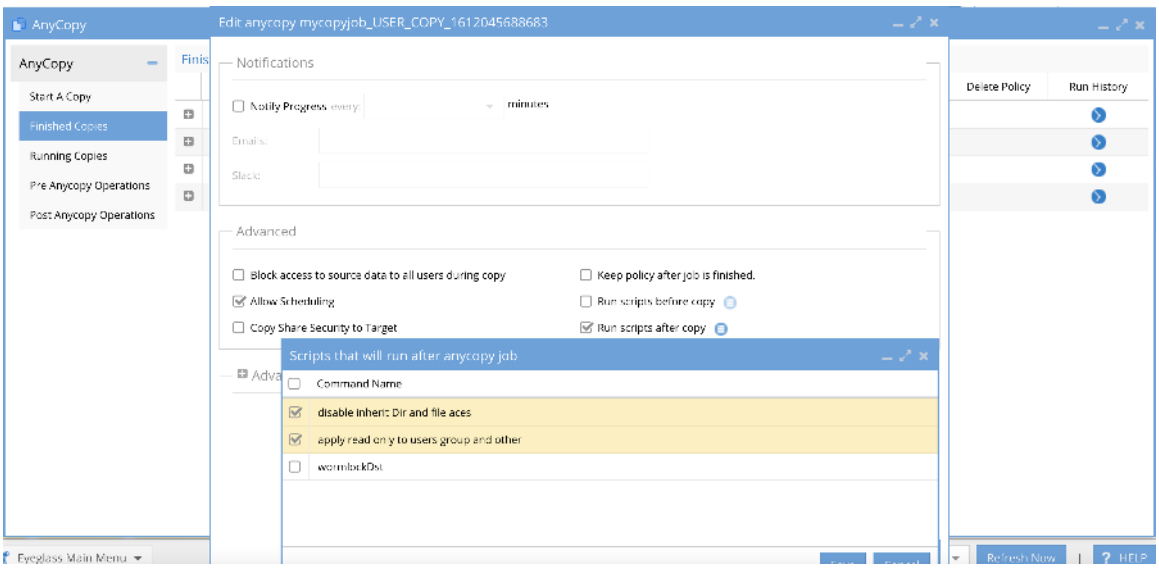

How to Extend the Copy Log Changelist job timeout

1. Open an SSH session to Eyeglass
2. Login as admin
3. Elevate to root using : sudo su (enter admin password when prompted)
4. Type the command below and press Enter
5. sudo su
6. Type the command below and press Enter
7. sudo vi /opt/superna/sca/data/system.xml
8. Once the file is opened, find/locate the tag below
9. <changelistTimeout>300</changelistTimeout>
10. Press i or "Insert" from your keyboard
11. Replace 300 with new time out value - suggest 1800 the tag should look like <changelistTimeout>1800</changelistTimeout>

12. If the file does not contain the tag add the following lines at the bottom of the file but before the </config> with new timeout value (example is 1800 s)
13. <anycopy><changelistTimeout>1800</changelistTimeout> </anycopy>
14. Save the file:
15. wq!
16. Restart SCA service
17. systemctl restart sca

How to Test script logic

1. Not available in current release

POSIX Mode - How it Works

1. **Overview:** If you are using NFS or AD authenticated users to access file systems with unified permissions using UID and GID locally generated by the cluster or stored in Active Directory RFC2703 UID and GID as properties of the user account can be used to present folders to the logged in user based on their UID/GID or the other permission on a folder. This allows posix permissions for owner, group or other to restrict what folders the user can see when selecting data to copy.
    a. **How Owner, group, other permissions determine what is browsable in the GUI.**

i. The source input box enforces owner , group and other Only if the logged in user has **Read permissions** to the folder. If they do not, the folder will not display.

ii. The target input box enforces owner , group and other Only if the logged in user has **Read & Write** permissions to the folder. If they do not, the folder will not display.

2. Example below shows 2 folders with folder1 secured to an AD user using UID and GID stored in AD, and folder2 is owned by root and wheel.

   a. 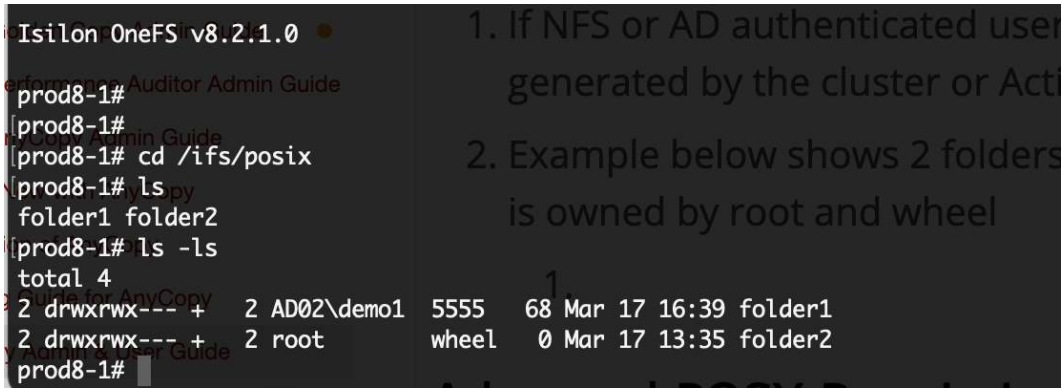

3. In this example the user demo1 logs into Eyeglass and has been assigned the AnyCopy user role to login. When the user tries to browser under an SMB share (NOTE: the user must have access to the SMB shares where posix data is stored.)

   a. After login Eyeglass will have the UID and GID assigned to the user in Active Directory.

4. The user will expand folders under the share and a REST API is used to check the owner and group of the folders, the user must be either owner or member of the group granting access to the folder. See screenshot below that only lists folder1.

5. 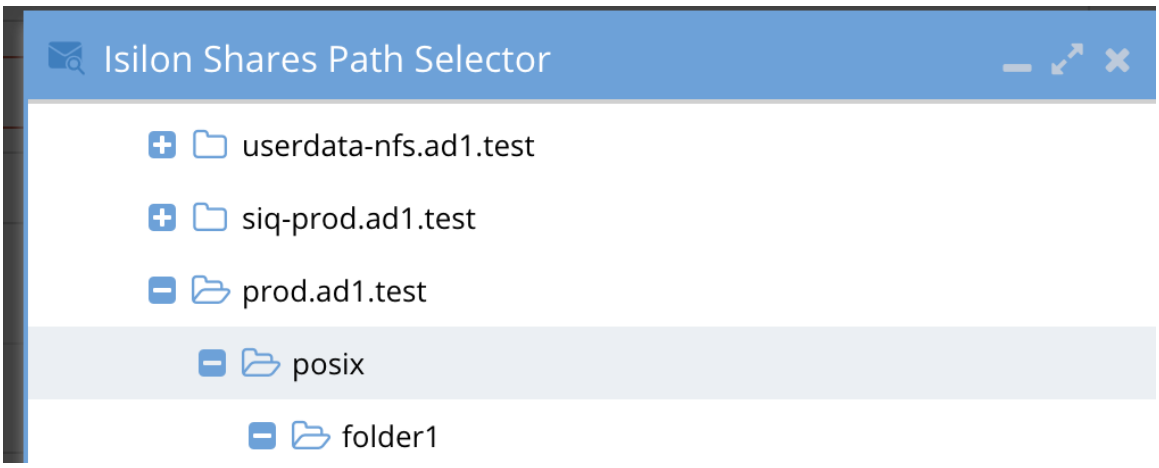

Advanced POSIX Permissions Mode Configuration

1. To enable this mode
2. ssh as admin
3. sudo -s
4. nano /opt/superna/sca/data/system.xml
5. Inside the tags <anycopy> and </anycopy> add this value
6. <usePosixPermissions>true</usePosixPermissions>
7. control+x answer yes to save
8. systemctl restart sca
9. done

© Superna LLC