Table of Contents

1. Eyeglass PowerScale Edition Quick Start Guide

Home Top

Introduction to this Guide

Use this document to get your new Eyeglass installation up and running fast with all the best options.
For planning DR and understanding design choices with Eyeglass use the Eyeglass Start Here First Guide

System Requirements

1. vSphere 6.0 ESX host or higher or hyper-v with vhdx appliance appliance requires
    a. **vcenter supported deployment clients**
        i. **vcenter 6.5 Flex or html5**
        ii. **vcenter 6.7 Flex client (vmware bug broke OVF with html webUI)**
        iii. **vcenter 7.0.1 Build: 17491160**

    b. **NOTE: All other tools are unsupported. Requires vapp property support.**

2. 4 vCPU
3. 16 GB RAM (RAM must be upgraded based on the scalability table here)
4. 30G OS partition plus 80 GB disk Total disk size in VMware 110G
5. Latency from admin PC to Eyeglass VM GUI < 15 ms ping times
6. AD Auth provider in System zone for RBAC and other SID to user resolution API requirements
7. Chrome Browser (Required), Browser must support Websockets, Internet Explorer is not supported.
    a. The browser must not disable 3rd party cookies required for authentication sessions and file downloads.
8. Eyeglass Port Requirements: Eyeglass-Ports-Requirements

IMPORTANT INFORMATION REGARDING ADDING CLUSTERS TO EYEGLASS READ-ME FIRST

1. **ONEFS 8.2 - This release disabled API and SSH access to the SSIP and introduces a range of SSIP addresses in a subnet. IF YOU ARE CURRENTLY USING AN SSIP THE BELOW STEPS ARE MANDATORY. NOTE: Requires Release 2.5.5 or later**
    a. **Upgrade to the latest release using this guide.**
    b. **Best Practise is to have an dynamic smartconnect IP pool in system zone for the IP used by Eyeglass to provide an HA connection to the cluster**
    c. **Login to Eyeglass open the Inventory Icon**

    d. Right click the 8.2 cluster and select the Edit option
    e. Enter any IP in the dynamic pool range
    f. Re-enter the password for the eyeglass service account
    g. Click submit
    h. Repeat for each 8.2 cluster in the inventory tree.
    i. Done.
    j. NOTE: Dynamic IP pools will fail the IP address to a new node if the node fails or is taken down for maintenance.
    k. NOTE: CSRF patch referenced below still applies to 8.2 clusters, that do not support session based authentication across the cluster and this blocks use for smartconnect for load balancing Eyeglass API's.

2. NOTE: For release 2.5.5 and later, you must use a node IP with dynamic IP allocation in subnet with pool in the System Access zone to add PowerScale clusters to Eyeglass (typically the management subnet). Using SmartConnect Zones is no longer supported due to PowerScale CSRF patch which disabled basic authentication and does not share session token between PowerScale nodes. For more details please refer to Technical Advisory #15 and Technical Advisory #17.

-

Supported OneFS releases

1. Please refer to the Release Notes for the Eyeglass PowerScale Edition version that you are installing.

Feature Release Compatibility

1. Please refer to the Release Notes for the Eyeglass PowerScale Edition version that you are installing.

Eyeglass Scalability Limits

1. Please refer to the Eyeglass Admin Guide Scalability limits.
2.

Eyeglass Firewall Port Requirements

Please refer to the firewall ports table.

Video Tutorial - Installing Eyeglass for PowerScale

The following link provides a video tutorial outlining how to install Eyeglass for PowerScale, add clusters and an overview of features.

New Eyeglass Installation

For a new Eyeglass installation, complete the following steps:

Download Eyeglass (Mandatory)

1. Download Eyeglass OVF, VHDX from Superna web site following instructions here Latest Appliance Download

2.

Deploy the Eyeglass Appliance (Mandatory)

Eyeglass is delivered in an OVF format for easy deployment in your vCenter environment. Deploy the OVF and then follow the wizard to setup networking for this Linux appliance. You will need to know:

1. subnet and network required so that appliance will have IP connectivity to the PowerScale clusters that it’s managing, and the users that are using it
2. IP address for the appliance
3. (Optional) SmartConnect Zone for management access to the cluster
4. Gateway
5. DNS server
6. NTP Server

IMPORTANT: If you are using hostname or FQDN for the target cluster in your SyncIQ policies or SmartConnect Zone for adding clusters to Eyeglass, the DNS information entered here must be able to resolve

back to a discovered cluster IP Address (should resolve to a SyncIQ SmartConnect Zone IP pool IP address), in order for Eyeglass to perform configuration replication. If the hostname cannot be resolved, Eyeglass will not create the associated configuration replication Job.

Steps to Deploy the OVF with vSphere Client (Mandatory)

OVF Deployment steps :
Step 1 : Download an OVF zip file from Latest Appliance Download.
Step 2 : Unzip the contents of the zip file from Step 1 onto a computer with vSphere web or Windows client installed.
Step 3 : Login to the vCenter with appropriate login credentials.
Step 4 : Single click on VMware vSphere client on the Desktop. Login with appropriate login credentials.
Step 5 : Once logged in to VMware client, you can see different Menus on the top left of the application. Next, go to the File menu and select Deploy OVF Template.
Step 6 : Browse to the location of OVF files you've downloaded and unzipped in step 1 and 2. Select OK and then Next.
Next, You will see the OVF template details. Verify the details and proceed by selecting Next. Notice download size to be under allocated disk size limit.
Step 7 : Choose a unique name for the virtual machine and select Inventory location for the deployed template. Once done, select Next.
Step 8 : Select the host/cluster where you want to run the deployed template and then Next.
Step 9 : Select the Resource pool within which you wish to deploy the template.
Step 10 : Select a destination storage for virtual machine files, select Next
Step 11 : Select Disk Format for the datastore you selected in previous step.
Step 12 : Enter the networking properties for the Eyeglass appliance VM in the OVF properties display. Replace with correct settings for your environment.

IMPORTANT: If you are using hostname for the target in your SyncIQ policies, the DNS information entered here must be able to resolve this host back to the Cluster IP Address in order for Eyeglass to perform configuration replication. If the hostname cannot be resolved, Eyeglass will not create the associated configuration replication Job.

Step 13 : When done, verify your settings and deploy the OVF

After deployment:

Step 1 : Power On the virtual machine.

1. The Eyeglass appliance is deployed with following default admin user password:
2. ssh to eyeglass vm as admin
    a. **sudo systemctl status superna-on-boot** (enter admin password and verify the first boot process completes)
    b. default login and password: **admin/3y3gl4ss**
3. Can also be used to login to the Eyeglass UI or SSH
4. **NOTE: It is highly recommended to reset the default password after the appliance is deployed.**

Setup Time zone and NTP (Mandatory)

1. Setup NTP server (published online list here)
2. Setup Timezone for log time alignment and SyncIQ operations.
3. Follow Animated GIF below to set using YAST
4. ssh as admin user,
5. sudo -s
6. Enter admin password
7. yast

8. 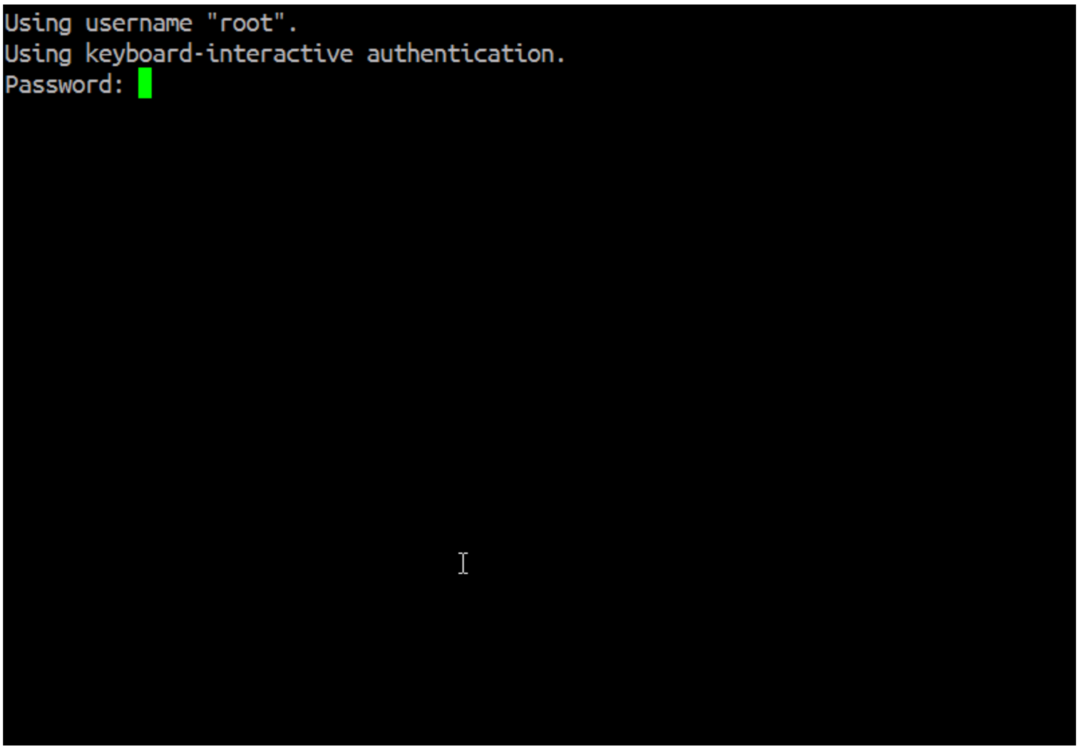

```
Using username "root".
Using keyboard-interactive authentication.
Password:
```

1.
    a.

Eyeglass Initial Configuration (Mandatory)

Your Eyeglass initial configuration steps are:

1. Login to the Eyeglass UI
2. Install License
3. Create Eyeglass service account first for each PowerScale cluster with Minimum Privileges (if not done configure Clusters in Eyeglass using root user)
4. Add Clusters

Login to the Eyeglass UI (Mandatory)

To login to the Eyeglass web UI, enter the following URL into your browser (Chrome preferred) replacing <Eyeglass IP address> with the real IP address assigned to the appliance:

1. https://<Eyeglass IP address>

2. You have 2 options for login authentication:
3. Login with appliance credentials - use the admin user and password configured on the appliance
4. Default user/password: admin / 3y3gl4ss

Install License (Mandatory)

1. Retrieve your Eyeglass License keys (instructions provided here).
2. Upload the license zip file provided to you by Superna:
3. **IMPORTANT: Do not unzip the license file. Upload the zip file.**
4. 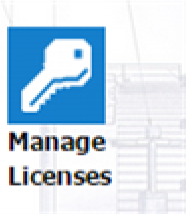

5. **IMPORTANT: You will be asked to accept the Eyeglass EULA and Phone Home after selecting the Upload button. License will not be loaded unless EULA is accepted.**
6. 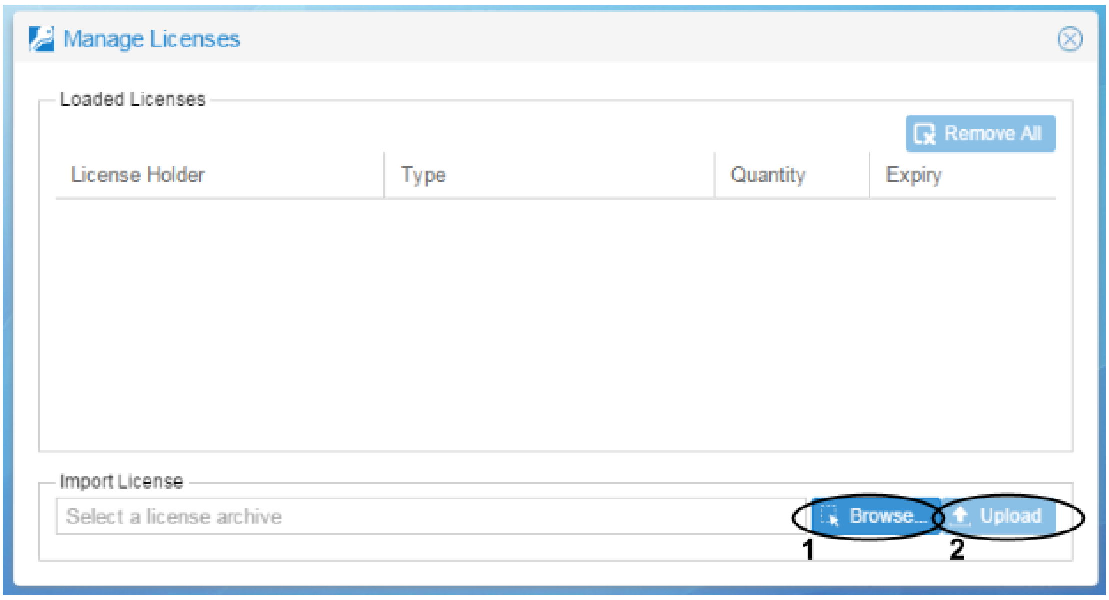

7. Done

Add PowerScale Clusters (Mandatory)

1. NOTE: No Auto Refresh Inventory View use the refresh button bottom right of the GUI

2. **This window does not auto refresh after adding a cluster. You must click the refresh button bottom right to verify when a cluster has finished discovery. This process can take 5-10 minutes typically.**
3. NOTE: Password special characters and length.
    a. These characters cannot be used [{ (any bracket open) , }]) (any bracket close), ~ (tilde), ` (back quote), \ (back slash), / (forward Slash), & , *, $ this is not a full list and more special characters may not work. Password length should be < = 20 characters.
4. NOTE: Cluster DNS Setup and Add Cluster to Inventory:
    a. **If discovery takes a very long time to complete (> 10 minutes), then to check to make sure that cluster configuration data can resolve external URL. Cloud pools use a URL to a storage bucket, and if this URL can not complete a DNS lookup to an IP address, then API calls that discover cloud pools will take too long to complete and will timeout the cluster discovery.   Make Sure all URL and DNS resolution is functioning on the cluster.**
5. IMPORTANT: After Discovery of a Cluster’s SyncIQ policies all Eyeglass configuration jobs are disabled automatically
    a. Configuration Replication Jobs for zones, shares, exports and NFS alias protected by SyncIQ Policy are automatically created and in the USERDISABLED state after successful provisioning in Eyeglass. Enabling these Jobs will be part of the installation steps.
6. IMPORTANT: Clusters on source target must be in the support feature matrix

   a. PowerScale cluster replication pairs must be running a supported OneFS version as documented in the System Requirements / Feature Release Compatibility matrix.

7. IMPORTANT: Before you add a Cluster to Eyeglass verify SyncIQ FQDN Name resolution

   a. **This step is important to allow Eyeglass to automatically build configuration replication jobs correctly. Eyeglass will resolve the FQDN of the SyncIQ policy and then compare the returned ip address to all PowerScale clusters added to the Eyeglass appliance. If no match is found, Config Sync jobs will fail to be added to the jobs window, until name resolution works correctly. A system alarm is also raised that indicates no matching clusters found for the SyncIQ policies on Cluster named X.**

8. Adding Clusters for Eyeglass

   a. PowerScale clusters must be added to Eyeglass using a node IP from an IP pool in the System Access Zone. Do not use the SSIP.
   b. **Create the eyeglass service account: To create an Eyeglass service account with minimum privileges follow the instructions provided in Eyeglass Service Account Minimum Privileges.**
   c. To verify SyncIQ target host FQDN:
      i. Login to Eyeglass via ssh
      ii. Validate that the FQDN of SyncIQ policy targets will resolve correctly on Eyeglass
      iii. nslookup <FQDN>
      iv. **NOTE: If it does not resolve validate DNS and make sure DNS can resolve this FQDN or Eyeglass will not auto detect SyncIQ replication relationships**
   d. From the Eyeglass UI add the PowerScale Clusters between which Eyeglass will be replicating the share and export configuration data

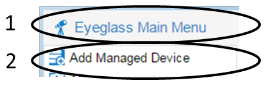

e. 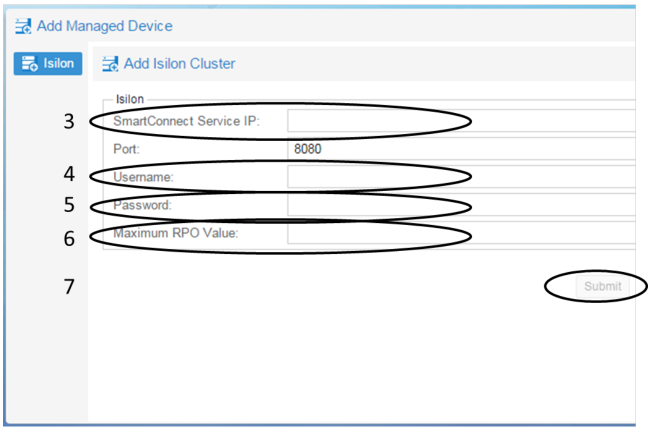

SmartConnect Service IP:

Port: 8080

Username:

Password:

Maximum RPO Value:

f. Note:

   i. If you get authentication failure when clicking submit. It can be one of these issues:

   ii. Bad password (make sure before looking at next causes)

   iii. SmartConnect Service must be IP address format. Also it must be an IP pool from System Access zone for PAPI API calls to be supported.

   iv. If your cluster is running original 7.2.x.x, 8.0.0.0, 8.0.0.1, 8.0.0.2 the TLS security protocols allowed weaker security algorithms and key sizes. Eyeglass 1.9 OVF and later has hardened security settings. In this case you may need to edit /opt/superna/java/jre1.8.0_05/lib/security/java.security and comment out the line “jdk.tls.disabledAlgorithms=MD5, SHA1, DSA, RSA keySize < 2048, SSLv2Hello, SSLv3, TLSv1, TLSv1.1”

      1. After editing this file an Eyeglass sca service restart is required
      2. systemctl restart sca

g. Maximum RPO Value is the Recovery Point Objective for the cluster in minutes. If you are using the RPO feature, this target is used during RPO analysis. More information about

Eyeglass RPO analysis can be found in Feature Overview - RPO Trending and Reporting.

9. After the cluster(s) are added successfully passing the authentication validation verify Inventory collection

10. Inventory Collection After Clusters are added

    a. Once the PowerScale is added, Eyeglass will automatically run an inventory task to discover the PowerScale components. When completed, the discovered inventory can be seen in the Inventory View.

    b. Click the Inventory Icon and verify the inventory completes as per below

        i. 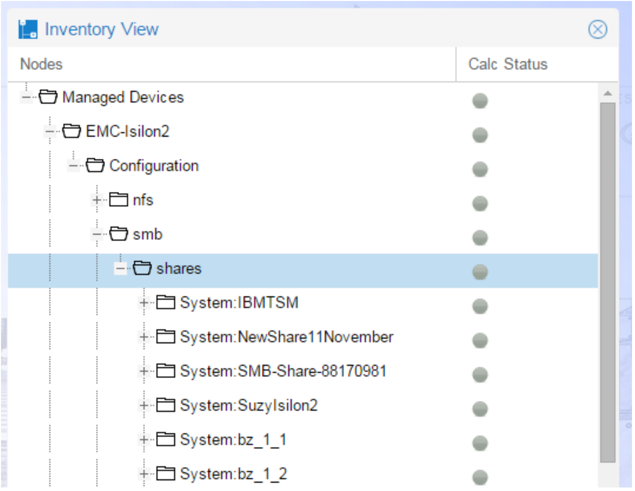

Inventory View

Nodes	Calc Status
Managed Devices	●
EMC-Isilon2	●
Configuration	●
nfs	●
smb	●
shares	●
System:IBMTSM	●
System:NewShare11November	●
System:SMB-Share-88170981	●
System:SuzyIsilon2	●
System:bz_1_1	●
System:bz_1_2	●

Enable Eyeglass Jobs (Mandatory)

1. Newly discovered SyncIQ policies will now appear as unconfigured in the Jobs Icon. You must set the job type as auto (mirrors configuration data between clusters) DFS (only used for DFS mounted SMB shares), and Skip config (special case when no configuration sync is required). **NOTE: Once you set the type they will be userdisabled and must be enbled for production use.**
2. Enable the jobs in the auto or DFS sections of the jobs icon with check box and bulk action menu to enable them. See section below on the steps.
3. **Optional:**

a. Enable File system job to sync snapshot schedules to DR
b. Zone job to sync access zones to DR (consult support before enabling not required in most cases)
c. Zone and Pool Failover Readiness job is disabled and only enabled if you plan to use IP Pool or Access Zone failover feature. Consult the Getting started guide to decide on the failover mode.
d. **NOTE: Do not enable or run quota jobs, these are managed automatically by the failover process. These jobs should not be enabled or run.**

Pre-requisite for Enabling Configuration Replication

1. If you have an Active - Active Replication Topology (for data), confirm that you do not have an unsupported share or NFS Alias environment described in the diagram below:

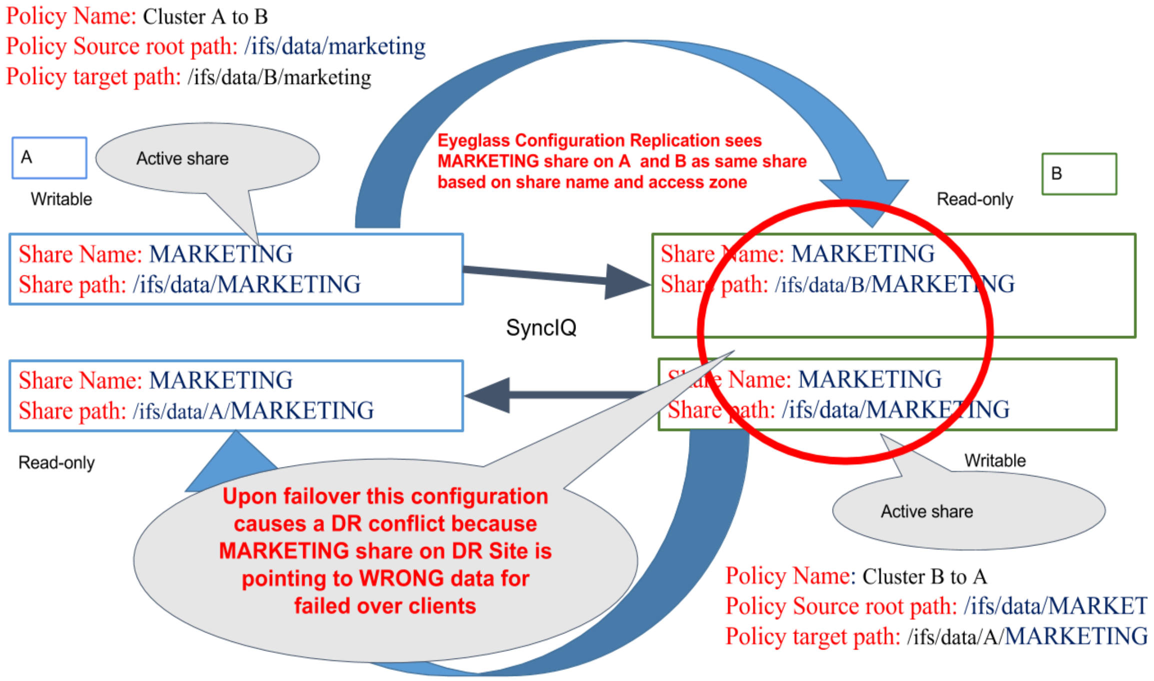
Policy Name: Cluster A to B
Policy Source root path: /ifs/data/marketing
Policy target path: /ifs/data/B/marketing
A
Active share
Writable
Eyeglass Configuration Replication sees MARKETING share on A and B as same share based on share name and access zone
B
Read-only
Share Name: MARKETING
Share path: /ifs/data/MARKETING
Share Name: MARKETING
Share path: /ifs/data/B/MARKETING
SyncIQ
Share Name: MARKETING
Share path: /ifs/data/A/MARKETING
Share Name: MARKETING
Share path: /ifs/data/MARKETING
Read-only
Writable
Active share
Upon failover this configuration causes a DR conflict because MARKETING share on DR Site is pointing to WRONG data for failed over clients
Policy Name: Cluster B to A
Policy Source root path: /ifs/data/MARKET
Policy target path: /ifs/data/A/MARKETING

2.

3. Review Eyeglass Admin Guide Jobs description to understand what the Configuration Replication Jobs will do

4. Review Eyeglass Admin Guide for Configuration Replication Pre-requisites

5. Review how Eyeglass determines uniqueness for configuration items and what properties are replicated.

Enable Jobs for Configuration Replication (Mandatory)

1. Next step is to enable your Share, Export, NFS Alias (AUTO) Jobs for Configuration Replication. This can be done on a Job by Job basis by following these steps:
2.

3. Select the Configuration Replication Job to be enabled.
4. 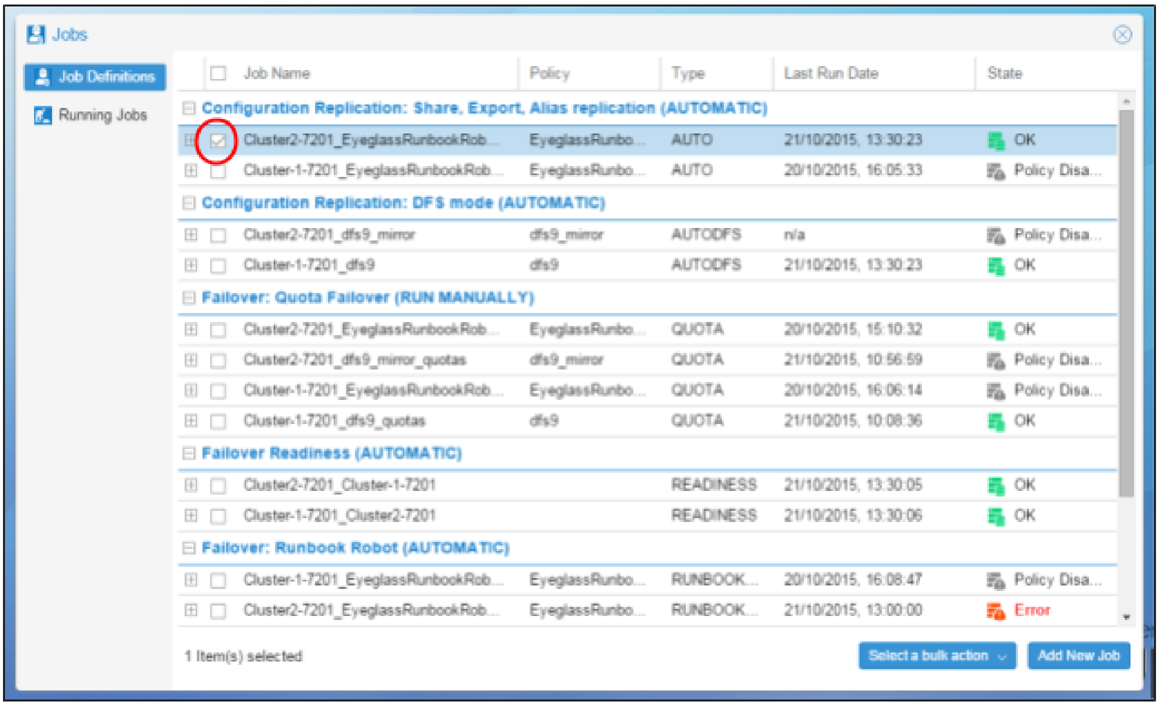

□	Job Name	Policy	Type	Last Run Date	State
	Configuration Replication: Share, Export, Alias replication (AUTOMATIC)				
☑	Cluster2-7201_EyeglassRunbookRob...	EyeglassRunbo...	AUTO	21/10/2015, 13:30:23	OK
□	Cluster-1-7201_EyeglassRunbookRob...	EyeglassRunbo...	AUTO	20/10/2015, 16:05:33	Policy Disa...
	Configuration Replication: DFS mode (AUTOMATIC)				
□	Cluster2-7201_dfs9_mirror	dfs9_mirror	AUTODFS	n/a	Policy Disa...
□	Cluster-1-7201_dfs9	dfs9	AUTODFS	21/10/2015, 13:30:23	OK
	Failover: Quota Failover (RUN MANUALLY)				
□	Cluster2-7201_EyeglassRunbookRob...	EyeglassRunbo...	QUOTA	20/10/2015, 15:10:32	OK
□	Cluster2-7201_dfs9_mirror_quotas	dfs9_mirror	QUOTA	21/10/2015, 10:56:59	Policy Disa...
□	Cluster-1-7201_EyeglassRunbookRob...	EyeglassRunbo...	QUOTA	20/10/2015, 16:06:14	Policy Disa...
□	Cluster-1-7201_dfs9_quotas	dfs9	QUOTA	21/10/2015, 10:08:36	OK
	Failover Readiness (AUTOMATIC)				
□	Cluster2-7201_Cluster-1-7201		READINESS	21/10/2015, 13:30:05	OK
□	Cluster-1-7201_Cluster2-7201		READINESS	21/10/2015, 13:30:06	OK
	Failover: Runbook Robot (AUTOMATIC)				
□	Cluster-1-7201_EyeglassRunbookRob...	EyeglassRunbo...	RUNBOOK...	20/10/2015, 16:08:47	Policy Disa...
□	Cluster2-7201_EyeglassRunbookRob...	EyeglassRunbo...	RUNBOOK...	21/10/2015, 13:00:00	Error

1 Item(s) selected

5. Select a bulk action and then select the Enable/Disable option.

6. 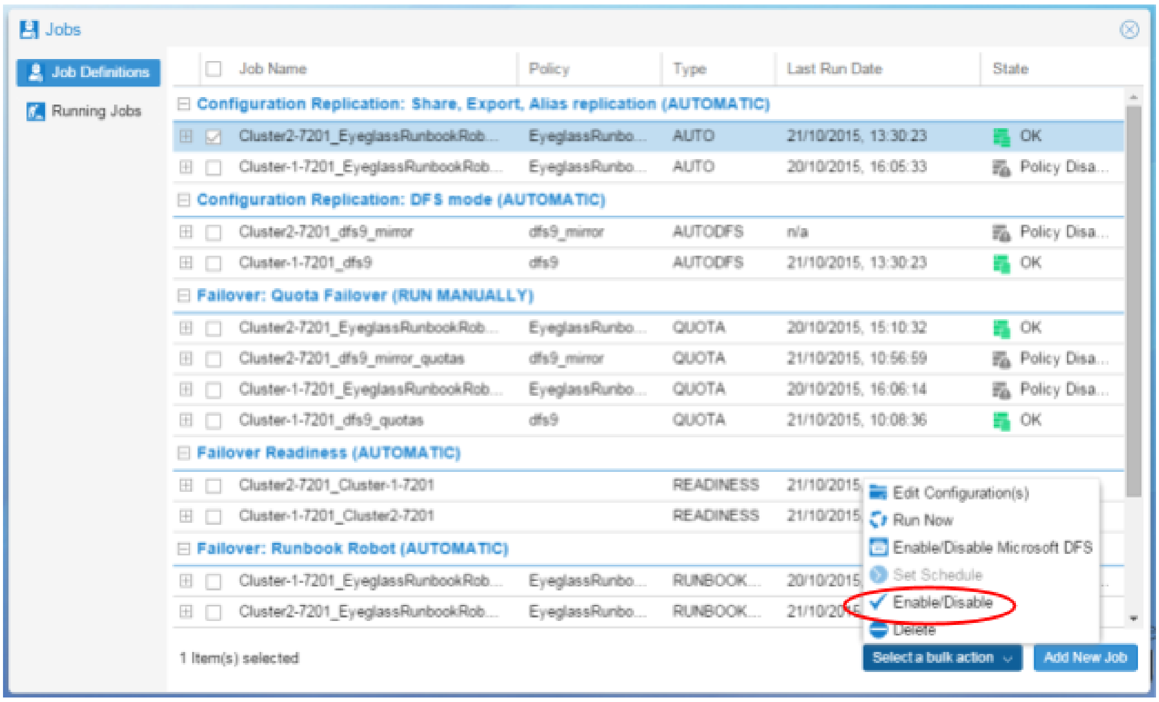

☐	Job Name	Policy	Type	Last Run Date	State
	Configuration Replication: Share, Export, Alias replication (AUTOMATIC)				
☑	Cluster2-7201_EyeglassRunbookRob...	EyeglassRunbo...	AUTO	21/10/2015, 13:30:23	OK
☐	Cluster-1-7201_EyeglassRunbookRob...	EyeglassRunbo...	AUTO	20/10/2015, 16:05:33	Policy Disa...
	Configuration Replication: DFS mode (AUTOMATIC)				
☐	Cluster2-7201_dfs9_mirror	dfs9_mirror	AUTODFS	n/a	Policy Disa...
☐	Cluster-1-7201_dfs9	dfs9	AUTODFS	21/10/2015, 13:30:23	OK
	Failover: Quota Failover (RUN MANUALLY)				
☐	Cluster2-7201_EyeglassRunbookRob...	EyeglassRunbo...	QUOTA	20/10/2015, 15:10:32	OK
☐	Cluster2-7201_dfs9_mirror_quotas	dfs9_mirror	QUOTA	21/10/2015, 10:56:59	Policy Disa...
☐	Cluster-1-7201_EyeglassRunbookRob...	EyeglassRunbo...	QUOTA	20/10/2015, 16:06:14	Policy Disa...
☐	Cluster-1-7201_dfs9_quotas	dfs9	QUOTA	21/10/2015, 10:08:36	OK
	Failover Readiness (AUTOMATIC)				
☐	Cluster2-7201_Cluster-1-7201		READINESS	21/10/2015	
☐	Cluster-1-7201_Cluster2-7201		READINESS	21/10/2015	
	Failover: Runbook Robot (AUTOMATIC)				
☐	Cluster-1-7201_EyeglassRunbookRob...	EyeglassRunbo...	RUNBOOK...	20/10/2015	...
☐	Cluster2-7201_EyeglassRunbookRob...	EyeglassRunbo...	RUNBOOK...	21/10/20	

1 Item(s) selected

7. On the next Configuration Replication cycle, the enabled Job will be run.

1.

Setup Eyeglass for Email Notification (Mandatory)

1. **Configure SMTP**
2. **Configure Email Recipients**

Configure SMTP (Mandatory)

1. Enter the information for your email server in the **Notification Center / Configure SMTP** tab.

2. 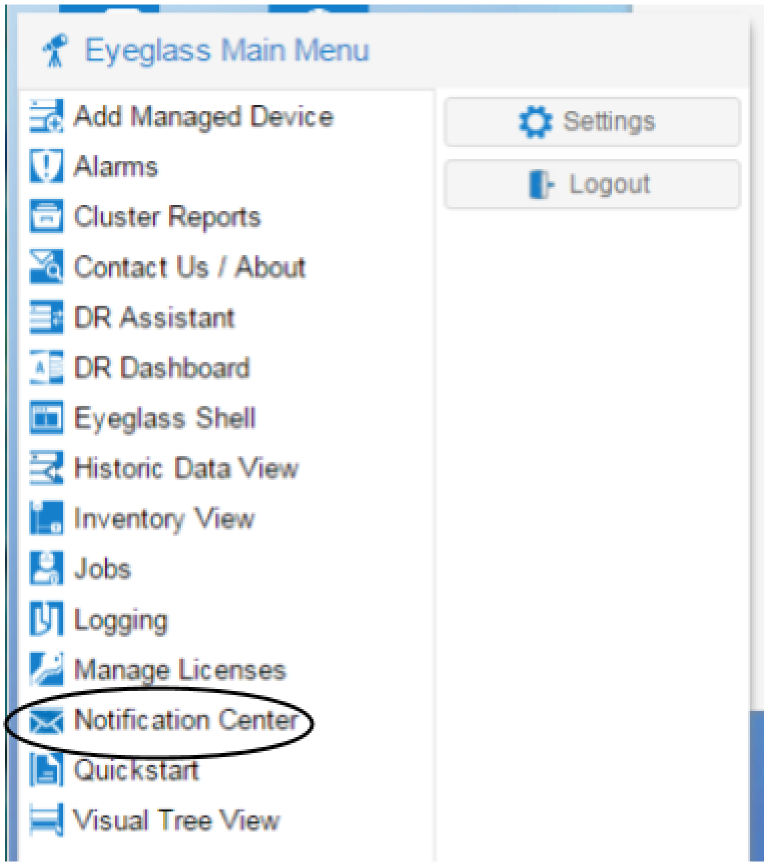

3. 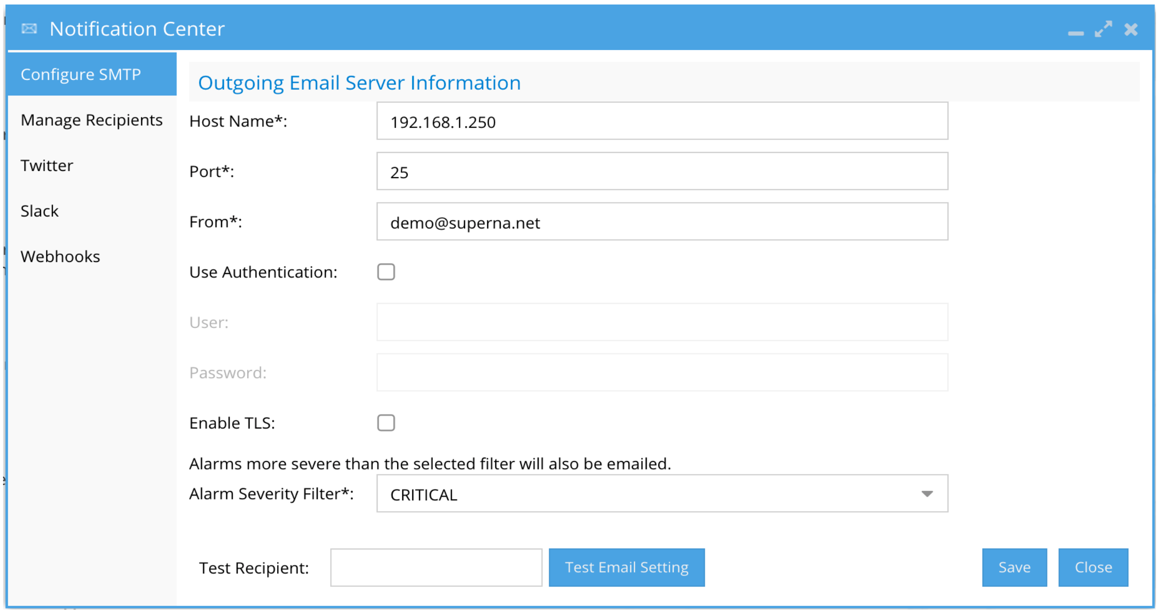

Host Name*: 192.168.1.250

Port*: 25

From*: demo@superna.net

Use Authentication: □

User:

Password:

Enable TLS: □

Alarms more severe than the selected filter will also be emailed.

Alarm Severity Filter*: CRITICAL

Test Recipient:

4.

5. Host name: Enter the host name for your email server

6. Port: Enter the port which should be used for sending email

7. From: Enter the email address of the sender of the email. Typically this is required to be a valid email address recognized by the email server.
8. Use Authentication: Select if email server requires an authenticated login
9. User: User or email address for authentication
10. Password: Password for authentication
11. Enable TLS: Select the Enable TLS check box if your email server expects TLS communication.
12. Alarm Severity Filter: Select level of alarms for which you would like to receive email.
13. Use the Test Email Setting button to check that the email server information added is correct. If an error occurs, you will get error codes from the SMTP connection. The "no error" response indicates successful connection. If an error is returned the debug response should be sent to support.superna.net.
14. Save your changes.

    -

Configure Email Recipients (Mandatory)

1. Enter the information for your email server in the **Notification Center / Manage Recipients** tab.
2. Email Recipient: Enter the email address that emails will be sent to.
3. Select the report type this user receipt
    a. All

b. Reports (RPO, cluster configuration)

c. Easy Auditor product (All reports and email notifications)

d. Cluster Storage Monitor product reports (quota usage)

e. Cluster Storage Monitor product Data recovery portal emails)

4. 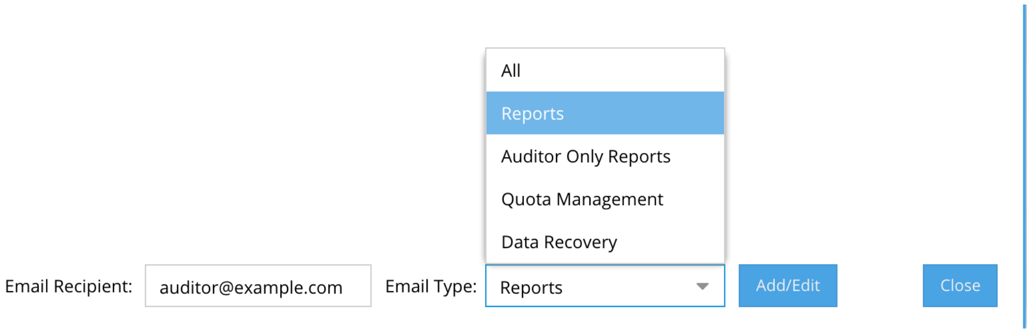

Email Recipient: auditor@example.com

Email Type: Reports

5. Select the Add button.

6. 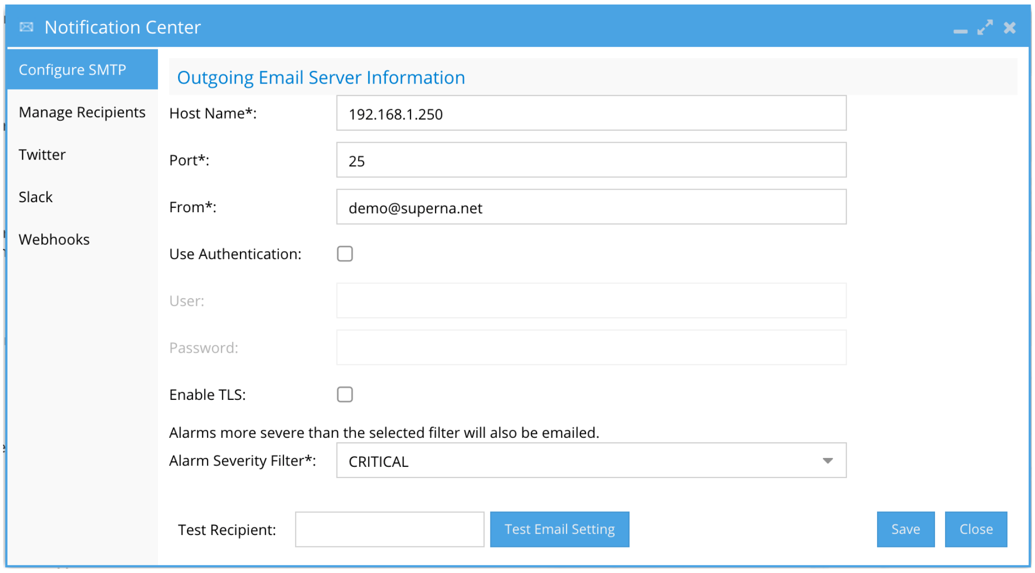

Notification Center

Outgoing Email Server Information

Host Name*: 192.168.1.250

Port*: 25

From*: demo@superna.net

Use Authentication: □

User:

Password:

Enable TLS: □

Alarms more severe than the selected filter will also be emailed.

Alarm Severity Filter*: CRITICAL

Test Recipient:

7. For other Notification center configuration options see the admin guide topic.

2.

Protecting the Eyeglass appliance (Optional)

1. See How to configure warm standby

How to Change Eyeglass Appliance Networking Configuration (yast network utility) after OVA deployment (Optional)

Note: if you need to update the configuration at any time, ssh to the appliance admin user and then sudo su - to root and use the yast2 command to open the wizard.

Step 1: sudo su - to root
Root user password is unique to the appliance and has no default.
Use the “sudo su” command to change to root user if desired.
Step 2: Type yast2 lan at the prompt.
Now follow steps to setup IP information on the appliance.

Notes on Using yast2

The yast2 interface is a generic text based user interface (TUI) that predates modern mouse and windows desktop environments. The downside of TUI’s are that they can be tricky to navigate if you’ve never used them before - the benefit of TUI’s is that they are widely compatible across platforms and do not require any graphical user interface to be deployed.

yast2 navigation Tips

Use the Tab (**TAB**)key to move from one field to the next.

Use Shift + Tab (**SHIFT TAB**) key combination to move backwards

Use the ↑ ← ↓ → (**ARROW**) keys to move around within a field

Use the Alt (**ALT**) key, along with the bold letter in the interface to select that specific field, tab or option

1. The Network Devices / Network Settings window is open.
2. In Network Settings screen open second tab: Hostname/DNS.
3. Type Hostname for your Eyeglass appliance
4. Type DNS server IP. In this example: 192.168.1.250

IMPORTANT: If you are using hostname for the target in your SyncIQ policies, the DNS information entered here must be able to resolve this host back to the Cluster IP Address in order for Eyeglass to perform configuration replication. If the hostname cannot be resolved, Eyeglass will not create the associated configuration replication Job.

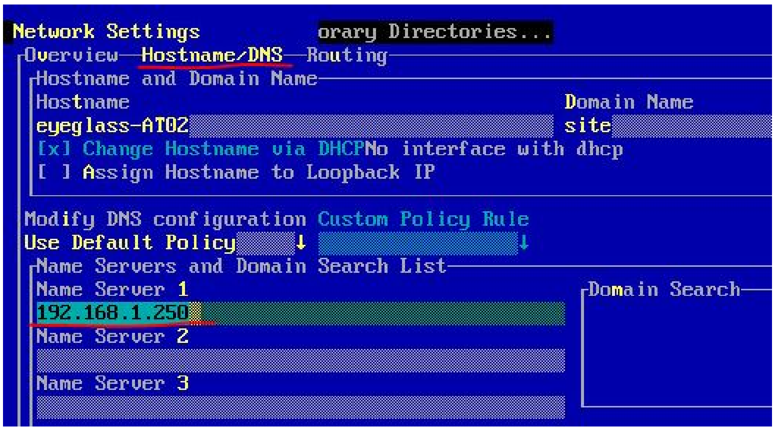

```
Network Settings             orary Directories...
Overview—Hostname/DNS—Routing
Hostname and Domain Name
Hostname                                      Domain Name
eyeglass-AT02                                 site
[x] Change Hostname via DHCPNo interface with dhcp
[ ] Assign Hostname to Loopback IP

Modify DNS configuration Custom Policy Rule
Use Default Policy    ↓                 ↓
Name Servers and Domain Search List
Name Server 1                                 Domain Search
192.168.1.250
Name Server 2

Name Server 3
```

9. In Network Settings screen open first tab: Overview. Choose Edit.

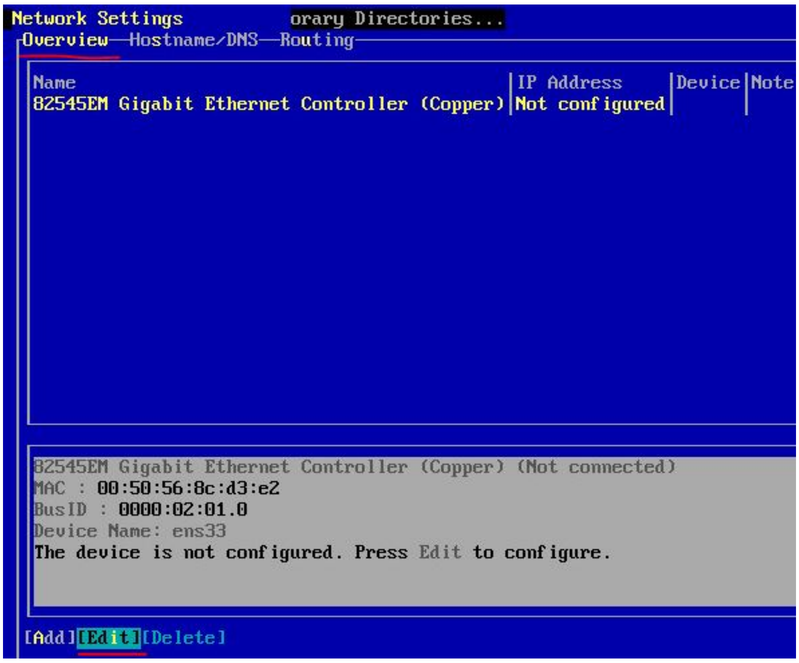

```
Network Settings            orary Directories...
Overview—Hostname/DNS—Routing

Name                                        |IP Address    |Device|Note
82545EM Gigabit Ethernet Controller (Copper)|Not configured|      |

82545EM Gigabit Ethernet Controller (Copper) (Not connected)
MAC : 00:50:56:8c:d3:e2
BusID : 0000:02:01.0
Device Name: ens33
The device is not configured. Press Edit to configure.

[Add][Edit][Delete]
```

1. In Network Card Setup screen, choose 2nd tab: Address. Type the **same** Hostname you entered when deploying the Eyeglass OVF..

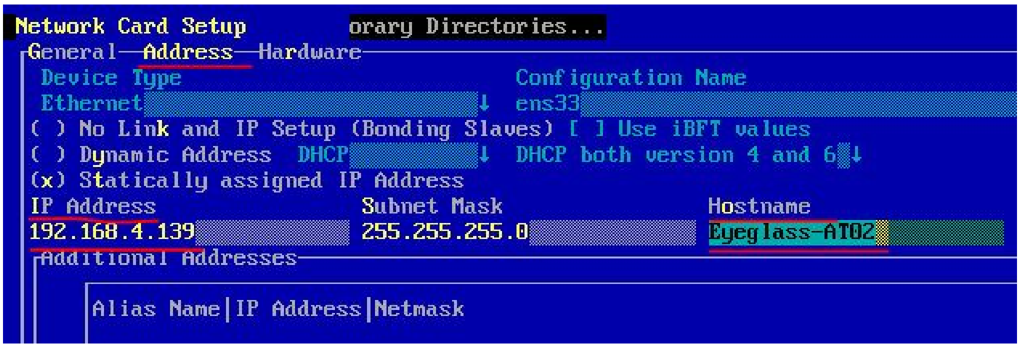

```
Network Card Setup          orary Directories...
┌General──Address──Hardware──────────────────────────────────────
│ Device Type                           Configuration Name
│ Ethernet                          ↓   ens33
│( ) No Link and IP Setup (Bonding Slaves) [ ] Use iBFT values
│( ) Dynamic Address  DHCP          ↓   DHCP both version 4 and 6 ↓
│(x) Statically assigned IP Address
│IP Address                Subnet Mask              Hostname
│192.168.4.139             255.255.255.0            Eyeglass-AT02
│┌Additional Addresses──────────────────────────────────────────
││
││    Alias Name│IP Address│Netmask
```

3. Next - OK - Next

1.1. Eyeglass Service Account Minimum Privileges

Home Top

Overview

Eyeglass communicates with PowerScale clusters to perform discovery and add/update/delete of share, export and quota configuration information. The minimum PowerScale cluster node user privileges required for Eyeglass/PowerScale connectivity to successfully perform configuration replication and support other Eyeglass features are:

NOTE: Any change to Eyeglass Service Account privileges requires an Eyeglass sca service restart to recognize the change (procedure below).

NOTE: AD or LDAP user is not supported, this lowers the system availability and adds dependency on AD/LDAP servers for API calls, local users on PowerScale have no dependence on AD/LDAP, in addition this generates too many authentication requests for API calls.

In addition to creation of the Eyeglass service account on the PowerScale cluster, the sudoer file on the cluster must be updated to allow the Eyeglass service account to execute OneFS CLI commands that require Elevated Permissions to run as root.

Step 1 - Creating the local PowerScale Eyeglass User - PowerScale Command Line For Eyeglass DR, Ransomware Defender, Easy Auditor and Storage Cluster Monitor

Use these permissions for all of the products above.

1. Follow create steps below
2. Then move to step 2 SUDO file configuration

To provision user and role from the PowerScale Cluster command line:

These commands below are executable by ssh as **root** on PowerScale and then right click : **Note: Service account set to password never expires.**

```
isi auth roles create --name EyeglassAdmin --description "EyeglassAdmin role"
isi auth users create eyeglass --enabled yes --password 3y3gl4ss
isi auth users modify eyeglass --password-expires no
isi auth roles modify EyeglassAdmin --add-user eyeglass
isi auth roles modify EyeglassAdmin --add-priv-ro ISI_PRIV_LOGIN_PAPI
isi auth roles modify EyeglassAdmin --add-priv ISI_PRIV_AUTH
isi auth roles modify EyeglassAdmin --add-priv ISI_PRIV_ROLE
isi auth roles modify EyeglassAdmin --add-priv ISI_PRIV_NFS
isi auth roles modify EyeglassAdmin --add-priv ISI_PRIV_SMB
isi auth roles modify EyeglassAdmin --add-priv ISI_PRIV_NETWORK
isi auth roles modify EyeglassAdmin --add-priv ISI_PRIV_QUOTA
isi auth roles modify EyeglassAdmin --add-priv-ro ISI_PRIV_LOGIN_SSH
isi auth roles modify EyeglassAdmin --add-priv ISI_PRIV_AUDIT
isi auth roles modify EyeglassAdmin --add-priv ISI_PRIV_SYNCIQ
isi auth roles modify EyeglassAdmin --add-priv-ro
 ISI_PRIV_NS_IFS_ACCESS
isi auth roles modify EyeglassAdmin --add-priv-ro
 ISI_PRIV_EVENT
isi auth roles modify EyeglassAdmin --add-priv-ro ISI_PRIV_HDFS
isi auth roles modify EyeglassAdmin --add-priv-ro
ISI_PRIV_REMOTE_SUPPORT
isi auth roles modify EyeglassAdmin --add-priv
ISI_PRIV_SNAPSHOT
isi auth roles modify EyeglassAdmin --add-priv-ro
ISI_PRIV_SMARTPOOLS
isi auth roles modify EyeglassAdmin --add-priv-ro
ISI_PRIV_WORM
isi auth roles modify EyeglassAdmin --add-priv-ro
ISI_PRIV_STATISTICS
isi auth roles modify EyeglassAdmin --add-priv
ISI_PRIV_JOB_ENGINE
isi auth roles modify EyeglassAdmin --add-priv-ro
ISI_PRIV_CLOUDPOOLS
isi auth roles modify EyeglassAdmin --add-priv-ro
ISI_PRIV_DEVICES
isi auth roles modify EyeglassAdmin --add-priv-ro
ISI_PRIV_FILE_FILTER
isi auth roles modify EyeglassAdmin --add-priv-ro
ISI_PRIV_HARDENING
isi auth roles modify EyeglassAdmin --add-priv-ro
ISI_PRIV_NDMP
isi auth roles modify EyeglassAdmin --add-priv-ro
ISI_PRIV_MONITORING
isi auth roles modify EyeglassAdmin --add-priv-ro
ISI_PRIV_ANTIVIRUS
```

```
isi auth roles modify EyeglassAdmin --add-priv-ro ISI_PRIV_FTP
isi auth roles modify EyeglassAdmin --add-priv-ro ISI_PRIV_HTTP
isi auth roles modify EyeglassAdmin --add-priv-ro ISI_PRIV_NTP
isi auth roles modify EyeglassAdmin --add-priv-ro
ISI_PRIV_SYS_UPGRADE
* new in 2.5.8 for config backup *
isi auth roles modify EyeglassAdmin --add-priv
ISI_PRIV_CONFIGURATION
```

Easy Auditor Additional Permissions Required:

If using Easy Auditor add this additional permission

```
isi auth roles modify EyeglassAdmin --add-priv-
ro ISI_PRIV_NS_TRAVERSE
```

Step 1A - Search & Recover and Golden Copy Product Service Account Local PowerScale User Account Creation:

1. Follow these instructions to create a dedicated Search & Recover/Golden Copy.

```
isi auth roles create --name EyeglassAdminSR --description
"Eyeglass Search & Recover role"
isi auth users create eyeglassSR --enabled yes --password
3y3gl4ss
isi auth users modify eyeglassSR --password-expires no
isi auth roles modify EyeglassAdminSR --add-user eyeglassSR
isi auth roles modify EyeglassAdminSR --add-priv-ro
ISI_PRIV_LOGIN_PAPI
```

```
isi auth roles modify EyeglassAdminSR --add-priv-ro ISI_PRIV_AUTH
isi auth roles modify EyeglassAdminSR --add-priv-ro ISI_PRIV_SMB
isi auth roles modify EyeglassAdminSR --add-priv ISI_PRIV_SNAPSHOT
isi auth roles modify EyeglassAdminSR --add-priv-ro ISI_PRIV_DEVICES
isi auth roles modify EyeglassAdminSR --add-priv-ro ISI_PRIV_NS_TRAVERSE
isi auth roles modify EyeglassAdminSR --add-priv-ro ISI_PRIV_NS_IFS_ACCESS
isi auth roles modify EyeglassAdminSR --add-priv ISI_PRIV_JOB_ENGINE
isi auth roles modify EyeglassAdminSR --add-priv-ro ISI_PRIV_NETWORK  (New As of May 2019)
```

****** Only required ACL security mode and file recovery portal feature **************

```
isi auth roles modify EyeglassAdminSR --add-priv-ro ISI_PRIV_IFS_BACKUP
isi auth roles modify EyeglassAdminSR --add-priv-ro ISI_PRIV_IFS_RESTORE
```

********* new as of 1.1.5 File pool reporting ***********

```
isi auth roles modify EyeglassAdminSR --add-priv-ro ISI_PRIV_CLOUDPOOLS
isi auth roles modify EyeglassAdminSR --add-priv-ro ISI_PRIV_SMARTPOOLS
```

```
*****  only add below for Golden Copy Product 1.1.6 and later requires these  *****
isi auth roles modify EyeglassAdminSR --add-priv-ro ISI_PRIV_LOGIN_SSH
isi auth roles modify EyeglassAdminSR --add-priv ISI_PRIV_NETWORK
isi auth roles modify EyeglassAdminSR --add-priv-ro ISI_PRIV_SMARTPOOLS
**********  Golden Copy ACL permissions collection and apply on recall *******************
isi auth roles modify EyeglassAdminSR --add-priv-ro ISI_PRIV_IFS_BACKUP
isi auth roles modify EyeglassAdminSR --add-priv-ro ISI_PRIV_IFS_RESTORE
```

Step 2 - SUDO Root Level Commands Needed for Eyeglass DR & Airgap , Cluster Storage Monitor Unlock My Files

In order to execute the some commands from the CLI that are not available in the PAPI for OneFS and require root-level (sudo) privileges for execution, this allows service accounts to run the command without having root access.

Apply the settings that apply to the applications you have purchased, apply to all clusters managed by Eyeglass products.

Steps to create sudo entries on PowerScale

1. Edit the sudoer file using the PowerScale **isi_visudo** command.
2. Sudo file opens in vi editor.

3. **NOTE: Add lines for the applications you need to enable, each product has different sudo requirements, review each section that applies below.**
4. Save your changes. (: then type wq!)
5. Repeat step for each cluster managed by Eyeglass.

SUDO file Updates Required for DR Edition & Airgap

Add the following lines DR product:

```
eyeglass ALL=(ALL) NOPASSWD: /usr/bin/isi_classic auth ads*

# (new for 2.5.6 DR validation)

eyeglass ALL=(ALL) NOPASSWD: /usr/bin/isi_classic domain info*
```

yellow highlighted entries should be deleted for 2.5.6 or later releases

```
eyeglass ALL=(ALL) NOPASSWD: /usr/bin/isi_classic networks* (only required if cluster is 7.x.x.x)
eyeglass ALL=(ALL) NOPASSWD: /usr/bin/isi_for_array isi status*
eyeglass ALL=(ALL) NOPASSWD: /usr/bin/isi status*
eyeglass ALL=(ALL) NOPASSWD: /usr/bin/isi_for_array date +%s
```

SUDO File Updates Required for Cluster Storage Monitor Unlock My Files Product

Add the following lines:

#(new Nov 12, 2019 for 2.5.5 latest release Cluster Storage Monitor UnLock My files

```
eyeglass ALL=(ALL) NOPASSWD: /usr/bin/isi_for_array -s isi_run -z ?* isi_classic smb file*
eyeglass ALL=(ALL) NOPASSWD: /usr/bin/isi_for_array isi_run -z ?* isi_classic smb file*
```

Airgap basic or Enterprise deployments on the protected clusters

```
# used for vault cluster alarm collection
eyeglass ALL=(ALL) NOPASSWD: /usr/bin/isi_for_array -n ? curl *
eyeglass ALL=(ALL) NOPASSWD: /usr/bin/isi_for_array -n ?? curl *
```

How to Restart Eyeglass services after making permissions changes:

1. SSH to Eyeglass appliance
2. Type: sudo su - (to elevate to root user - enter the admin user password)
3. Type: systemctl restart sca
4. Type: systemctl status sca (to verify sca service active and running after the restart)

How to use Eyeglass DR Edition with Compliance Mode Clusters

For clusters using compliance mode sudoer and root access is not permitted.
This means that clusters must be added to Eyeglass using the user below:
Compadmin

How to use Eyeglass DR Edition with Clusters using STIG Hardening

1. You can enable STIG OS hardening on the cluster using:
2. isi hardening apply --profile=STIG --report=true
3. done

How to Create Eyeglass Service Account on Dell ECS For Ransomware Defender

1. Login to the ECS as admin, expand the management menu, click users, select Management users tab, Click New User.
2. Complete the form as per below example. User name eyeglass with administrator role assigned and set a password.

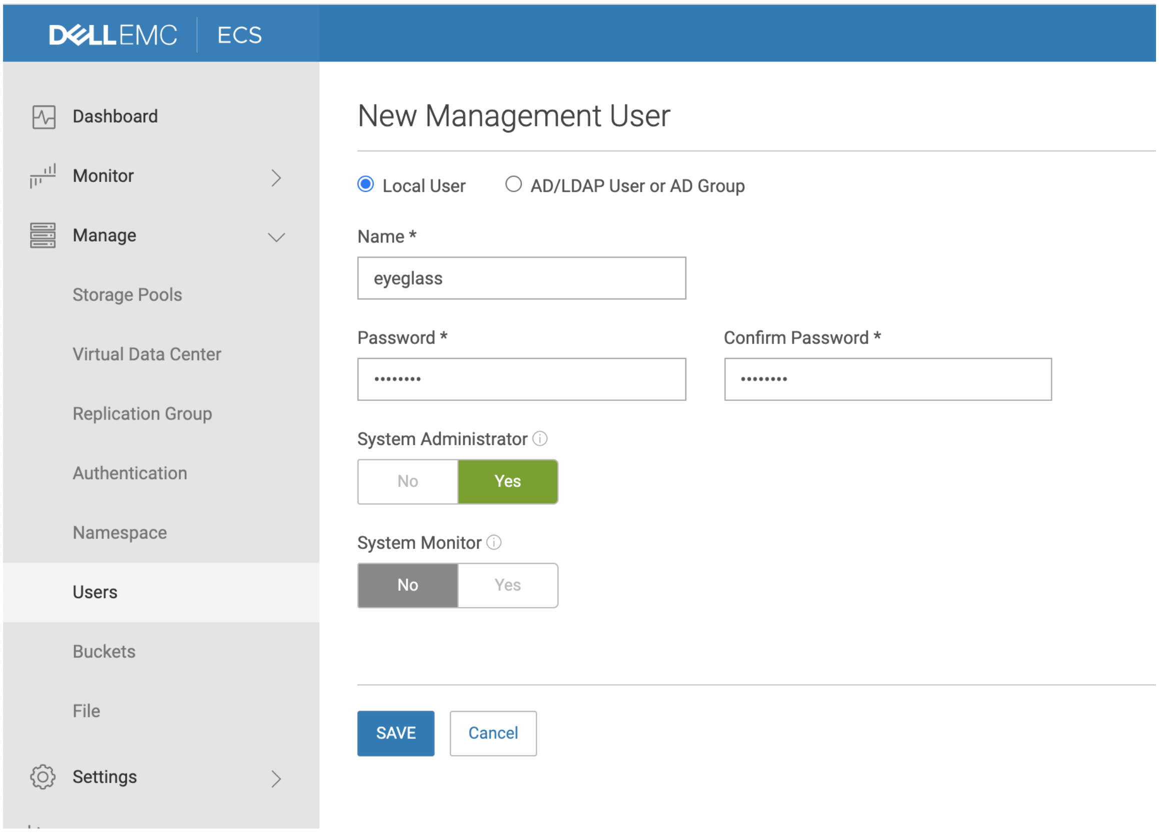

New Management User

● Local User ○ AD/LDAP User or AD Group

Name * eyeglass

Password * ••••••••

Confirm Password * ••••••••

System Administrator: No [✓ Yes]

System Monitor: [✓ No] Yes

3.

4.

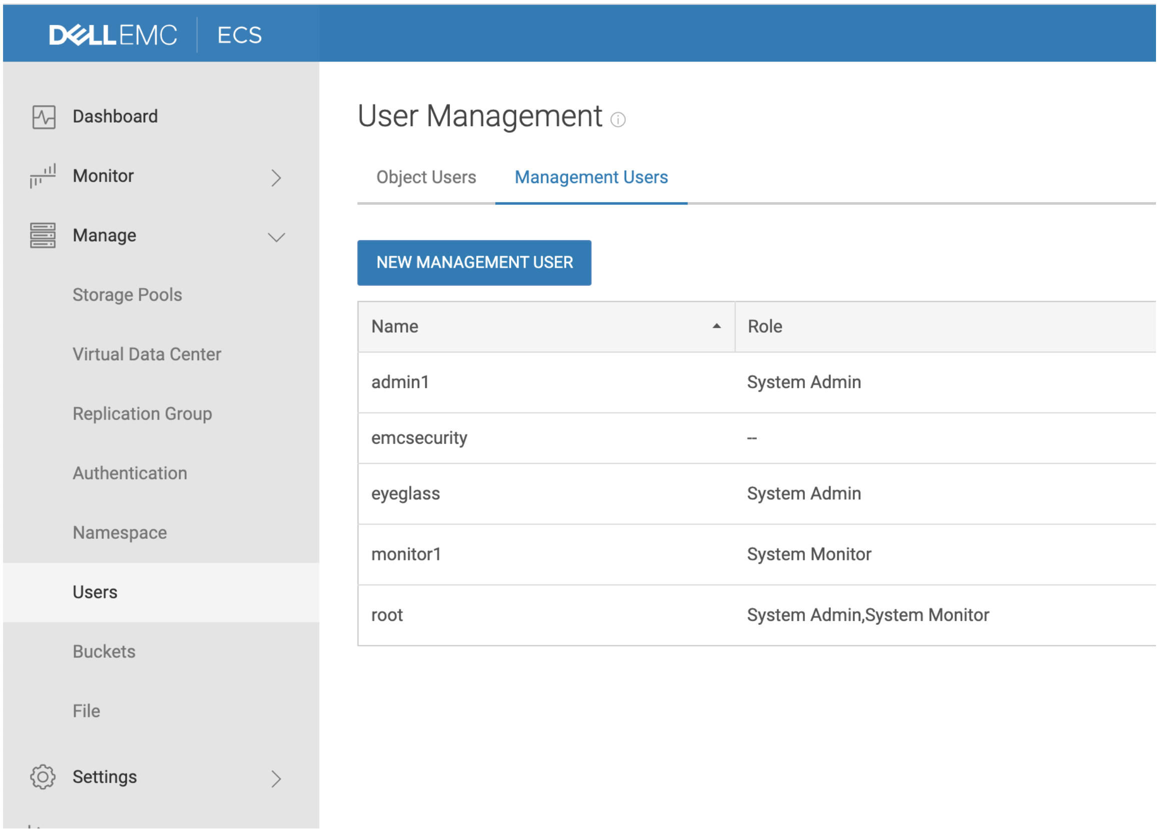

Name	Role
admin1	System Admin
emcsecurity	--
eyeglass	System Admin
monitor1	System Monitor
root	System Admin,System Monitor

How to create the Airgap Eyeglass Service Account for the Vault Cluster

Follow this steps on the vault cluster during vault cluster hardening. Replace the yellow with a strong password, that should only be known by the security officer or security personnel.

NOTE: The protected production cluster requires additional sudo permissions for vault alarm collection. See the eyeglass service account update section for the eyeglass

Eyeglass vault service account creation

```
isi auth roles create --name EyeglassAdminvault --description "EyeglassAdmin role Vault"
isi auth users create eyeglassvault --enabled yes --password <set a strong password here>
isi auth users modify eyeglassvault --password-expires no
isi auth roles modify EyeglassAdminvault --add-priv-ro ISI_PRIV_LOGIN_PAPI

isi auth roles modify EyeglassAdminvault --add-priv-ro ISI_PRIV_DEVICES

isi auth roles modify EyeglassAdminvault --add-priv-ro ISI_PRIV_NS_IFS_ACCESS

isi auth roles modify EyeglassAdminvault --add-priv ISI_PRIV_JOB_ENGINE

isi auth roles modify EyeglassAdminvault --add-priv ISI_PRIV_NETWORK

isi auth roles modify EyeglassAdminvault --add-priv-ro ISI_PRIV_LOGIN_SSH

isi auth roles modify EyeglassAdminvault --add-priv ISI_PRIV_SYNCIQ

isi auth roles modify EyeglassAdminvault --add-priv-ro ISI_PRIV_EVENT
```

Eyeglass service account Creation

```
isi auth roles create --name EyeglassAdmin --description
"EyeglassAdmin remote role"
isi auth users create eyeglass --enabled yes --password <set a
strong password here>
isi auth users modify eyeglass --password-expires no
isi auth roles modify EyeglassAdmin --add-priv-ro
ISI_PRIV_EVENT
```

Edit the sudo file to enable log gather support

1. Edit the sudoer file using the PowerScale **isi_visudo** command.
2. Sudo file opens in vi editor
3. eyeglassvault ALL=(ALL) NOPASSWD: /usr/bin/isi_gather_info
4. Save your changes. (: then type wq!)
5. Repeat step for each cluster managed by Eyeglass.

1.2. Eyeglass Warm Standby Direct Sync Guide

Home Top

Overview

This procedure is to protect the production Eyeglass appliance using a 2nd Eyeglass appliance. This solution offers sync of Eyeglass backup file to the 2nd appliance that can take over operations under the following conditions:

1. **Controlled failover** - switching the active appliance from one data center to the other
2. **Uncontrolled failover protection -** the 2nd appliance can be used for failover operations since it has a current near real time synced status and copy of all policies and configuration data (shares, exports and quotas) needed to complete a failover to the surviving cluster.
   a. **NOTE: If you declare an uncontrolled failover is required. You must ensure the active appliance is at the site where the data will be become active. The Warm standby appliance procedure should be used if the active appliance is at the site that had the disaster.**

The appliance should be co-located with the data that will be writeable.

Definitions:

1. **Active Appliance** - Responsible for syncing configuration data and is the primary appliance for all failover operations
2. **Active Sync Appliance** - 2nd appliance that has a backup synced from the active appliance

Operating Considerations

1. If you execute a planned controlled failover you should generate a backup from the active appliance after the backup completes. This backup reflects the current DR state after failover. The daily backup and sync process will generate a backup but you require a current view of the clusters DR status. **Always generate a backup after a planned failover.**
2. **Best Practise:**
    a. Enable phone home and request support to enable daily appliance backup to create an off site backup of your appliance. This backup is retained for 14 days before it is automatically deleted. This provides an off site backup of your DR state.

    b. After a planned failover push an off site backup using direct to support option.

        i. Select the Support backup you generated with the check box, select the "Upload selected file directly

Superna support". This will push an off site backup that can be requested back if needed.

c. OR If phone home is disabled and firewall or proxy does not allow the direct upload to support option, then download a newly created support backup from the Active appliance (after the failover), then upload to the support page following steps here. This backup is retained for 14 days before it is automatically deleted.

d. In addition to the above options, download the support backup file and store a copy at the opposite site from where the active appliance is located. Document this location internally to use with the recover options below if it is determined this backup is the most recent backup.

e. **NOTE: The daily sync backup enabled by direct sync is 24 hours old and is typically the backup required in most cases.**

Deploy 2nd Appliance & Prep for Sync

1. Follow the install guide to deploy a 2nd appliance. Guide link.

Configure Keyless SSH on Active Appliance

1. Login via ssh as the admin user
2. type: ssh-keygen (hit enter to all prompts)
3. ssh-copy-id -i /home/admin/.ssh/id_rsa.pub admin@x.x.x.x (where x.x.x.x is the ip address of the 2nd warm appliance, enter

yes to accept the ssh finger print, enter the remote warm standby appliance admin user password)

4. Test keyless ssh was successful
    a. ssh admin@x.x.x.x (where x.x.x.x is the ip address of the remote warm standby appliance)
    b. if no password prompt then everything worked
5. done

Configure Scheduled Cron Sync from Active to Warm Standby

1. Login to the **Warm StandbyAppliance** as admin user over ssh and set the permissions
    a. sudo mkdir -p /opt/superna/var/backup/
    b. sudo setfacl -m u:admin:rwx /opt/superna/var/backup/
2. Login to the **Active Appliance** to Test copy backup files
    a. Login as the admin user
    b. scp /opt/superna/var/backup/* admin@x.x.x.x:/opt/superna/var/backup/ (x.x.x.x is the ip of the warm standby appliance)
    c. if successful continue, if not debug the steps above.
3. Create the script on the **Active Appliance**
    a. Login as the admin user over ssh
    b. nano /home/admin/warmstandby.sh

    c. paste the string below into the file and replace x.x.x.x with ip of the **warm standby appliance**:

        i. scp /opt/superna/var/backup/* admin@**x.x.x.x**:/opt/superna/var/backup/

    d. Press control+x answer yes to save the file

    e. chmod 777 /home/admin/warmstandby.sh

4. Schedule the script to run daily at noon with cron on the **Active Appliance**

    a. Login to the **Active Appliance** over ssh as the admin user

    b. sudo -s (enter admin password)

    c. cd /etc/cron.d

    d. type this command:

        i. **echo "0 8 * * * admin /usr/bin/timeout 6h /home/admin/warmstandby.sh" > iglsstandby**

    e. Restart the cron service to pickup the new script

        i. systemctl restart cron

5. Verify backup copy script is running

    a. Wait until the next scheduled copy

    b. Login to the **warm standby appliance** as admin user and list the folder to verify files have the correct daily date stamp

    c. ls -ls /opt/superna/var/backup/

How to Restore the Warm Standby Appliance to become Active Procedures (Requires 2.5.6 release or later)

1. **You can restore a backup to your second Warm Standby Eyeglass appliance to make the standby the active appliance. SSH as the admin user into the warm standby appliance and execute the following steps:**
    a. ssh to as **admin** user on the **Warm Standby Eyeglass appliance**
    b. Run the restore command
        i. **NOTE that command must include full path to backup zip file or use path only and restore command will detect the most recent backup it discovers on the path based on the time stamp on the backup zip files.**
        ii. **Option #1 Auto detect the most recent file**
            1. igls app restore /opt/superna/var/backup/
                a. **This option will auto detect the most recent backup files and displays it for confirmation to use the suggested file, accept the file selection. You will be prompted to enter the admin password again to become the root user after entering the command.**
        iii. **Option #2 Use a specific file**

1. igls app restore /opt/superna/var/backup/<name_of_backup.zip>

   a. **NOTE: you will be prompted to enter the admin password again to become the root user after entering the command.**

c. You will be prompted with a confirmation of yes/no.

   i. **For Testing Only Answer No to exit the process.**

      1. This step should be used to test the restore procedure without actually running the restore.
      2. This will not restore the database and will not make the standby appliance the active appliannce.

d. For Production Active appliance switch to the Warm Standby Answer answer **Yes**

   i. Continue and monitor the command execution until it completes before trying to login to the GUI, it may take 15-20 seconds before you can login to the web UI.

   ii. **IMPORTANT: Original Eyeglass appliance should be powered off if switching to the Warm Standby Appliance. NEVER have two appliances operating against the same clusters. This can cause a conflict and is not supported. Verify your backup appliance is up and running after a restoration. Open a Browser to ip address of warm standby appliance.**

    e. **Once the restore process completes, you can login to the GUI to start a failover job following the documentation for executing a failover. see Guide here.**

How to Complete Appliance Switch Test Procedures

1. Install 2 appliances as per above steps in this guide
2. Shutdown the active appliance VM (power off)
3. Follow instructions to switch to the Warm standby Appliance
4. Verify the login to the GUI and open the jobs icon , and running jobs tab to verify normal configuration sync jobs are running, view the DR Dashboard after waiting at least 15 minutes to get a current view of DR Readiness.
5. Test completed
6. How to revert to the Active Appliance follow these steps:
    a. Login warm standby appliance
    b. Factory reset the warm standby appliance
        i. Login via ssh as admin user
        ii. run this command
            1. ./opt/superna/sbin/reset.sh
            2. This command will delete the database and remove all clusters added to the appliance and return to a default state.
    c. Power on the Active Appliance VM again

d. Login and verify jobs, and DR dashboard

e. Done

1.3. Legacy - Eyeglass Warm Standby with SyncIQ Guide

Home Top

Legacy Guide

Eyeglass Warm Standby Configuration and Restore Execution Procedures

What's New

1. Release 2.5.6 offers an auto detection of most recent backup to use which speeds up the restore process and decreases RTO by eliminating maual steps to select or list files to use for restore. This can remove minutes of time to recovery the 2nd appliance. The procedures to restore the 2nd appliance with auto detect mode should be possible to complete in less < 1 minute with single CLI command needed

Abstract:

This procedure is to protect the production Eyeglass appliance using the automated 7 day backup feature in Eyeglass and using SyncIQ to sync the backup archives to a 2nd cluster where the Warm standby Eyeglass appliance is running.

Notes on this solution:

- This process also means only **One** Eyeglass appliance has clusters added.
- This only requires one set of license keys (license keys are appliance specific)
- Requires SyncIQ enabled clusters to protect the archives
- Production Eyeglass sync's Export used for DR of Eyeglass configuration data

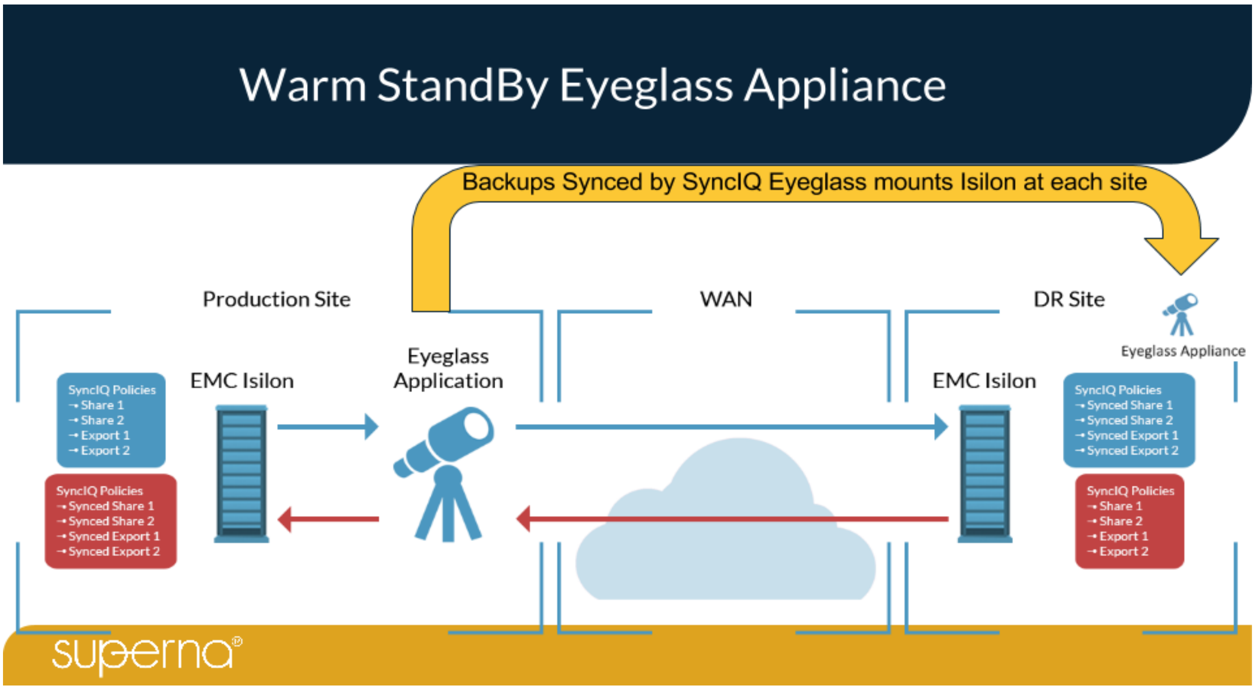

Procedure Overview below:

If you already have an second Eyeglass backup appliance, you may skip to Configuration Steps, otherwise start by installing a second OVF Appliance using the installagion guide.
You will need a second Eyeglass appliance set up. However, it will not manage any clusters and will be “empty”. This appliance is present for the sole purpose of restoring an Eyeglass backup archive to it, just in case a problem should arise in your original Eyeglass appliance.
Deploy the Eyeglass OVF and set up the network settings as usual. Confirm the installation went well by logging into your appliance and if all went well, stop there. **Do not** add any clusters, **licenses**, etc. to the appliance. **Everything** will be restored after restoring a backup archive.

Requirements

1. Standby appliance deployed at 2nd site
2. Standby appliance running the same release as the production appliance
3. All firewall ports opened to the clusters that need to be managed
4. DNS, NTP configuration is identical to the production appliance

Configuration

1. On your Source cluster, create an NFS export and change the following settings:
    a. Example path /ifs/eyeglasswarmstandby
    b. Under “clients” and “root clients”, type your 2 Eyeglass IPs, separated by a comma.
    c. Under “Permissions”, check “Mount access to subdirectories”.
    d. save the export
    e. Set permissions in the file system
        i. ssh to the PowerScale as root user
        ii. set permissions
            1. chmod 777 **/ifs/eyeglasswarmstandby**
    f. done
2. Create a SyncIQ policy pointing to the NFS export path while inputting the target cluster ip address or Smartconnect FQDN used for replication and desired target directory on the remote cluster example **/ifs/eyeglasswarmstandby**
    a. Run the syncIQ policy manually to create the remote directory.
    b. Set the schedule on the policy to 1 time per day
    c. Wait 5 minutes or start Eyeglass Configuration sync job to sync the export to the DR cluster.

   d. Ensure the export was synced by checking in your Eyeglass appliance the Configuration Replication Jobs is successfully completed and verifying that the export was created on the target cluster using Onefs GUI.

   e. **NOTE: Change the policy setting to "when the source is modified" if you plan to change the scheduled backup process to run more often than once per day. The instructions to increase the interval to run the backup process is documented below.**

   f. done

3. You will now need to perform the following mount command on the **production eyeglass appliance**.

   a. SSH into your Eyeglass appliance and switch to root user with sudo -s (enter admin password)

   b. Run this command and enter the admin password when prompted

      i. sudo chown sca:users /opt/superna/var/backup

      ii. sudo chmod 777 /opt/superna/var/backup

   c. **Execute the following:** vim /etc/fstab

   d. Add the command line below and replace it as indicated:

   e. <Source-cluster-IP>:/ifs/eyeglasswarmstandby
      /opt/superna/var/backup nfs   rw   0 0

   f. Then execute the command:

   g. mount -a

   h. verify no errors are returned from this command

        i. Type "mount" to verify the mount was successful. You should see the NFS export listed in the output of the mount command.

4. **Backups are created automatically by Eyeglass on a daily basis at 2:00 am. If you would like to test out that this procedure is working and backups are being written to the cluster, you may manually run the script via SSH anytime you'd like, rather than waiting 24 hours. To run the script, do the following:**
    a. SSH into your Eyeglass appliance and gain root access. (login as admin and type sudo -s to switch to root)
    b. Change directory to: /opt/superna/bin
    c. Execute the following: ./backup_last_seven_days.sh (this will create a backup)
        i. **NOTE: this schedule can be changed to a more refrequent schedule following these steps**
        ii. **NOTE: backup runs by default on daily schedule**
        iii. **How to Change the schedule**
            1. ssh to eyeglass as admin
            2. enter sudo -s (enter password for admin)
            3. nano /etc/cron.d/sca-backup
            4. change the cron string to a new value (Use this set to help with the cron string (https://crontab.guru/)
            5. save the file with ctrl key + x answer yes to save the file

   d. Verify that backup file has been created on the cluster mount path.

5. Now, you will need to SSH into your **warm standby** Eyeglass appliance and execute the following command to mount the NFS:

   a. SSH into your Eyeglass appliance and with **sudo -s** (enter admin password)

   b. **Run these commands**

      i. sudo chown sca:users /opt/superna/var/backup

      ii. sudo chmod 777 /opt/superna/var/backup

   c. Execute the following: nano /etc/fstab

      i. Add the command line below and replace it as indicated:

      ii. <Target-cluster-IP>: /ifs/eyeglasswarmstandby
       /opt/superna/var/backup nfs   rw   0 0

      iii. mount -a

      iv. Run the SyncIQ Policy on your Source Cluster to copy the backup files to the DR cluster.

         1. Ensure backup files are present from the Warm Standby appliance by listing the path

6. **Verify Warm Standy Appliance has visibility to replicated backup files are visible and have a current date stamp.**

   a. sudo -s (enter admin password)

   b. cd /opt/superna/var/backup

   c. Type ls (to list contents of the mount point from step above)

d. **Verify backup archive(s) exist from the source cluster, synced by SyncIQ to target cluster and mounted visible to the “Warm standby Eyeglass appliance”**

How to Restore the Warm Standby Appliance to become Active Procedures (Requires 2.5.6 release or later)

1. **You can restore a backup to your second Warm Standy Eyeglass appliance from the mounted path. SSH into the backup appliance and execute the following:**
   a. ssh to as **admin** user on the **Warm Standby Eyeglass appliance**
   b. Run the restore command - **NOTE that command must include full path to backup zip file or use path only and restore command will detect the most recent backup it discovers on the path.**
      i. **Option #1 Auto detect the most recent file**
         1. igls app restore /opt/superna/var/backup **(this option auto detect the most recent backup files and displays it for confirmation to use the suggested file, accept the file selection. You will be prompted to enter the admin password again to become root after entering the command)**
      ii. **Option #2 Use a specific file**
         1. igls app restore /opt/superna/var/backup/<name_of_backup.zip

> (NOTE: you will be prompted to enter the admin password again to become root after entering the command)

c. You will be given a confirmation (yes/no).

    i. **For Testing answer No to exit the process. This step should be used to test the restore procedure without actually running the restore.**

d. For Production appliance recovery answer **Yes** to continue and monitor the command execution until it completes before trying to login to the GUI, it may take 15-20 seconds before you can login.

    i. **IMPORTANT: Original Eyeglass appliance should be off or gone if restoring. NEVER have two appliances operating against the clusters under management. This can cause a conflict and is not supported. Verify your backup appliance is up and running after a restoration, by https to ip address of warm standby appliance.**

e. **Once the restore process completes, you can login to the GUI to start a failover job following the documentation for exeucting a failover.**

How to Restore the Warm Standby Appliance to become Active Procedures (< 2.5.6 releases)

You can restore a backup to your second Warm Standy Eyeglass appliance from the mounted path. SSH into the backup appliance and execute the following:

1. ssh to as **admin** user on the **Warm Standby Eyeglass appliance**
2. Run the restore command - **note that command must include full path to backup zip file.**
    a. use this command to get list of all available backup files ls /opt/superna/var/backup
    b. igls app restore /opt/superna/var/backup/<name_of_backup.zip> **(NOTE: you will be prompted to enter the admin password again to become root after entering the command)**
3. You will be given a confirmation (yes/no).
    a. **For Testing answer No to exit the process. This step should be used to test the restore procedure without actually running the restore.**
4. For Production appliance recovery answer **Yes** to continue and monitor the command execution until it completes before trying to login to the GUI, it may take 15-20 seconds before you can login.
    a. **IMPORTANT: Original Eyeglass appliance should be off or gone if restoring. NEVER have two appliances operating against the clusters under management. This can cause a conflict and is not supported. Verify your backup appliance is up and running after a restoration, by https to ip address of warm standby appliance.**

5. **Once the restore process completes, you can login to the GUI to start a failover job following the documentation for exeucting a failover.**

2. Eyeglass PowerScale Edition Upgrade Guide

Home Top

- Validate - Service account permissions, Eyeglass Job Status, Pool Mappings, Licenses and Cluster Inventory
- Validate - Ransomware Defender, Easy Auditor, Performance Auditor License Assignment
- Validate Ransomware Defender and Easy Auditor settings

Read Me First

New Validations. Expect to get warnings post upgrade. This is expected and is a key feature of this release to detect failover conditions that must be addressed to be ready for failover. If you get alarms it is good, so the issues can be fixed.

1. **SyncIQ Domain Mark for fast fail back** - checks both clusters to verify if SyncIQ domain mark exists and will raise warning if not found (NOTE: it will raise warning if you did not apply sudo update on step #1)
2. **SPN Delegation for Access zone and IP pool failover** - This check AD Delegation was completed to the cluster AD objects and the opposite cluster (cross test). If any test fails it will raise a warning with exactly which delegation permission is missing.
3. **DNS Dual Delegation -** This ensures automatic DNS resolution is in place before a failover. This validation will inspect your DNS configuration to determine if the Delegation is correct, A records resolve to subnet service ip's in the correct subnet. If you use Infoblox this validation must be disabled. It only supports standards based Name server delegations and not DNS forwarding.

Before Executing Upgrade Steps

1. The upgrade will disrupt Eyeglass services for less than 10 minutes
2. A VM level snapshot should be taken before upgrade to allow rollback to the previous version of the appliance

In-place Upgrades

Scenario #1 - Appliances running Older Open Suse OS to the latest release

This option allows customers to upgrade to the latest release without deploying a new OVF to get the latest operating system. **NOTE: OS version 42.3, 15.1 no longer receives security updates and is customers choice to stay on this OS of the appliance. NOTE: The OS is not covered by the support contract.**

1. To check the OS version
2. ssh as admin user to Eyeglass
3. type cat /etc/os-release
4. The OS version is displayed

Scenario #2 - Appliances running Open Suse 15.1 or 15.2 OS to the latest release

1. No special steps
2. Continue to upgrade instructions here.

In-place Upgrade - Installer Download and Upgrade Procedures

1. To complete an offline upgrade:
2. Login to support site with a registered support account https://support.superna.net
3. Scroll down on page after login to locate the software download validation form.

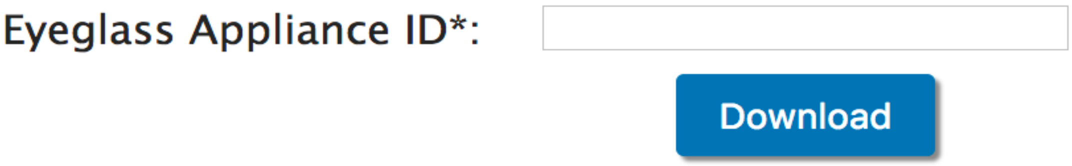

Eyeglass Appliance ID*:

4.
5. Get the appliance ID from the about window of the Eyeglass desktop

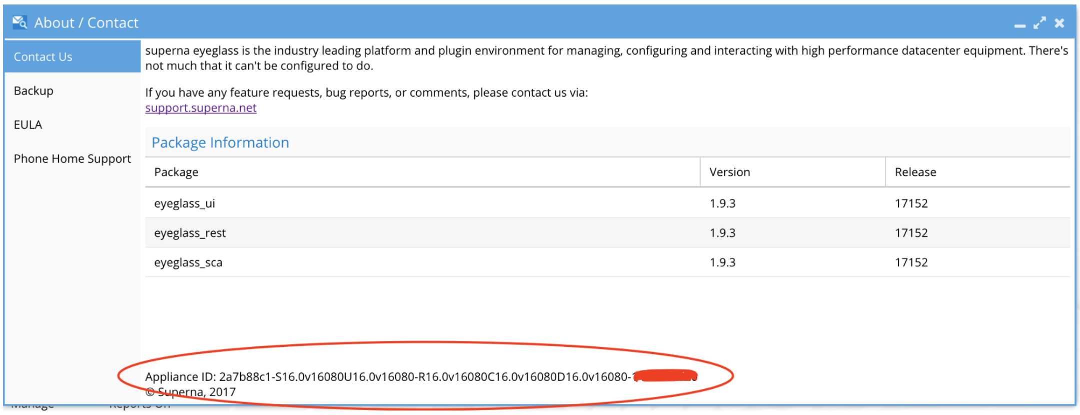

About / Contact

superna eyeglass is the industry leading platform and plugin environment for managing, configuring and interacting with high performance datacenter equipment. There's not much that it can't be configured to do.

If you have any feature requests, bug reports, or comments, please contact us via: support.superna.net

Package Information

Package	Version	Release
eyeglass_ui	1.9.3	17152
eyeglass_rest	1.9.3	17152
eyeglass_sca	1.9.3	17152

Appliance ID: 2a7b88c1-S16.0v16080U16.0v16080-R16.0v16080C16.0v16080D16.0v16080-[redacted]

© Superna, 2017

6.
7. Enter the appliance ID and click download button to retrieve the offline installer. NOTE: This command checks for an active support contract, and will only download software if support contract validation is successful.
8. Use winscp tool (google winscp download) to copy the offline package onto the appliance with the admin user and password.

9. ssh to the Eyeglass appliance and sudo to root command: **sudo su - (you will be prompted for the admin password again)**
10. Make the offline package executable: **chmod 755 <filename>**
11. Run the installer: **./<filename>**
12. You may be prompted for Phone Home Agreement if not previously set. Enter '**y**' or '**n**' to continue. Phone Home allows remote monitoring and faster support that allows remote log collection.
13. Once the update is completed, login to the Eyeglass web page.
14. **IMPORTANT**: Refresh any open Eyeglass window to ensure that you have latest changes.
15. Check the About Eyeglass window and verify version shows the version you downloaded. The full list of releases can be found **here**.
16. Complete
17. Check Post upgrade steps here

Upgrade to A New Appliance from an Old Appliance

1. Follow the steps in the next sections to complete the backup and restore process from an old appliance to a new appliance

Step 1 - Upgrade Path From Old Appliance Versions to Open Suse 15.x OS with the latest Release - Backup/Restore Method

All Appliance versions prior to latest version are using Open Suse OS versions that no longer have security patches available (13.1, 13.2 , 42.1, 42.3, 15.1). **Use this upgrade option to get upgraded to the latest Eyeglass release and get the latest Open Suse 15.x OS that includes automatic security patch updates. NOTE: If you are using an older version appliance backup file some settings are not retained depending on the backup file release version. The table in this document outlines settings that are migrated.**

1. Follow steps to download the new OVF here
2. Deploy new Eyeglass VM using the install guide as a reference.
3. **NOTE: The new appliance ip address can be different than the old appliance IP.**
4. Reference the table of settings that are migrated in the next section.
5. After the new appliance is deployed and you can login to the webUI and ssh then continue with the steps below.

Step 1a - Review Table of Migrated Settings

Eyeglass Configuration Item	**Source Appliance software version > 1.8.0**
Restoring local credentials for clusters	Yes
Restoring licenses keys	Yes
Adjusting licenses keys to latest format	Yes

Job Schedules	Yes
Job Initial state Setting (enabled, disabled)	No
custom settings with igls adv command.	Yes
Restore Notification Center settings1 Post restore Edit Notification Settings and set the	Yes
Restoring failover log history (if available)	Yes
Restoring custom RBAC roles (if available)	Yes
Restoring API tokens (if available)	Yes (as of 1.9.0)
Restoring Ransomware Defender security guard logs (if available)	Yes
Restoring cluster Configuration reports (if available)	Yes
Restoring Current Job state (enabled, disabled, DFS mode) (if available)	Yes
Alarm history	No
Old Backups Archives	No
Cluster Storage Monitor Data (if available)	No
RPO Generated Reports	No
RPO Report Data	No
Failover Scripts	Yes

Ransomware Defender Settings and History (if available) • Ransomware Defender History • Ransomware Defender ignored list settings • Ransomware Defender Statistics • Ransomware Defender Settings • Security Guard configuration2 2. schedule is restored but no other settings - these need to be re-added manually with user service account and password.	No

Step 1b - Review Historical Eyeglass Data & Settings that are Not Restored before continuing

1. All existing Eyeglass databases are removed, no backup is made.
2. **NOTE: This will delete databases and they will be rediscovered on startup. DO NOT USE this method if you have Cluster Storage Monitor or Ransomware Defender historical events or RPO Report data that you need to retain. Contact support if this applies to you scenario.**

Step 2 - Information to Record before Upgrading

1. Take a screenshot of the Eyeglass Jobs window prior to upgrade. This can be used as a reference to verify Job state and type. **Example auto type or dfs type.**
2. **(if IP pool mode configured)** Take a screenshot of the IP Pool failover policy to pool mappings

3. **(if Ransomware Defender)** Take a screenshot of Ransomware Defender
    a. Flag as False Positive (2.5.6 and lower) / Learned Thresholds (2.5.7 and higher)
    b. Ignored List
    c. Threshold window all settings
    d. Allowed Files (2.5.6 and lower) / File Filters (2.5.7 and higher)
    e. Monitor Only Settings (2.5.7 and higher)
4. **(if Easy Auditor)** Take a screenshot of Easy Auditor
    a. Active Auditor Triggers configured

Step 2a - Automated Appliance Configuration Import

1. Requirements:
    a. Must be running 2.5.7.1 or later new OVA appliance. **Check About Icon for the version. NOTE: Do not use this command if not running the correct version and use Step 3 below if running a version < 2.5.7.1**
    b. SSH access from new appliance to the old appliance
    c. New appliance is deployed on a new ip address
    d. 2.5.7.1
        i. ECA must be updated to point at the new eyeglass IP address
    e. 2.5.8

        i. This release will automatically update the ECA configuration and update the Eyeglass IP address and API token and push the change to all ECA nodes

    f. **NOTE: This option will not migrate custom threat file settings on ECA. This is not a common modification to ECA deployments and does not apply to most deployments.**

    g. **NOTE: This command can take 10 minutes to run.**

2. On the new appliance login as admin

    a. run import command: igls app pull-config --ip=<ip> --user=<user>

    b. IP = address of the old eyeglass appliance

    c. --user = admin

    d. You will be prompted for the password

3. The command will automate

    a. create a backup

    b. copy the backup to the new appliance

    c. apply the backup to the new appliance.

    d. Update the ECA configuration to use the new eyeglass IP address

        i. NOTE: after this command completes you will need to restart the ECA cluster to update the firewall configuration.

            1. login to ECA node 1 and run **ecactl cluster down**, followed by **ecactl cluster up**

    e. Shutdown the OS on the source appliance

f. done.

1. Once the steps complete -   Skip to - Check Post upgrade steps here

Step 3 - Restore Zip File (old appliance) and Restore to New Appliance Procedures

1. Take an Eyeglass **Restore backup** from your old Eyeglass appliance.
2. Download the **Restore backup** locally and then copy the zip file backup using scp or winscp to the newly deployed Eyeglass Appliance. It should be placed in /tmp folder .
3. **See Restore Backup button that is required versus support backup. The Restore backup includes SSL private keys, the support backup does not. This applies to Releases > 2.5.3**
4.

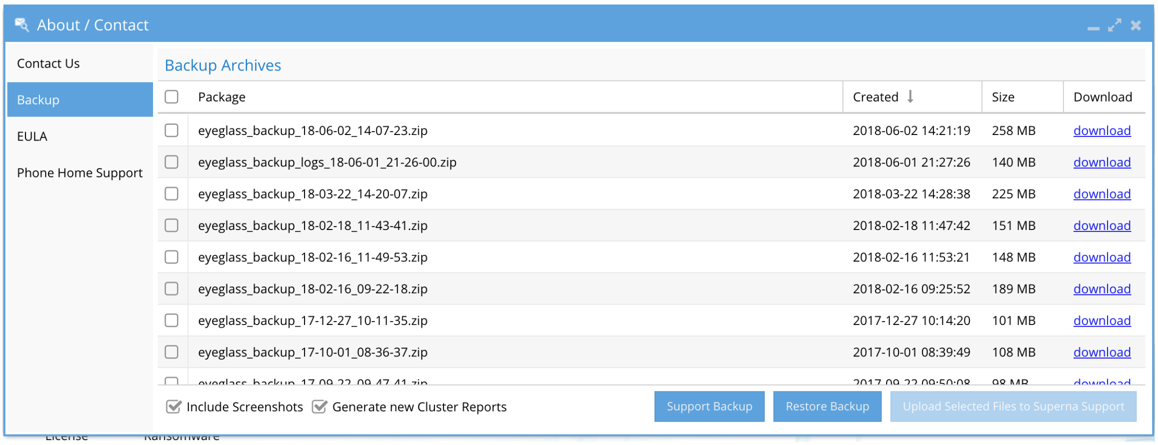

☐	Package	Created ↓	Size	Download
☐	eyeglass_backup_18-06-02_14-07-23.zip	2018-06-02 14:21:19	258 MB	download
☐	eyeglass_backup_logs_18-06-01_21-26-00.zip	2018-06-01 21:27:26	140 MB	download
☐	eyeglass_backup_18-03-22_14-20-07.zip	2018-03-22 14:28:38	225 MB	download
☐	eyeglass_backup_18-02-18_11-43-41.zip	2018-02-18 11:47:42	151 MB	download
☐	eyeglass_backup_18-02-16_11-49-53.zip	2018-02-16 11:53:21	148 MB	download
☐	eyeglass_backup_18-02-16_09-22-18.zip	2018-02-16 09:25:52	189 MB	download
☐	eyeglass_backup_17-12-27_10-11-35.zip	2017-12-27 10:14:20	101 MB	download
☐	eyeglass_backup_17-10-01_08-36-37.zip	2017-10-01 08:39:49	108 MB	download
☐	eyeglass_backup_17-09-22_09-47-41.zip	2017-09-22 09:50:08	98 MB	download

☑ Include Screenshots ☑ Generate new Cluster Reports

5. **Power off the old Eyeglass appliance.** It is not supported to have multiple Eyeglass appliances managing the same clusters.
6. SSH to **new** Eyeglass appliance and login as admin (default password 3y3gl4ss). Issue “sudo su -” to enter in root mode (default password 3y3gl4ss).
7. From the command line execute the command

a. **igls app restore /tmp/<eyeglass_backup.xxxx.zip> --anyrelease**

b. Replacing **/tmp/<eyeglass_backup.xxxx.zip>** with the name of the Eyeglass Archive file always including full path.

c. You will be prompted to continue. Enter “**y**” to continue.

d. For example:

   i. **igls app restore /tmp/eyeglass_backup_17-07-05_20-42-08.zip --anyrelease**

   ii. **Do you want to revert to the archive at /tmp/eyeglass_backup_17-07-05_20-42-08.zip? [y/N]: y**

e. Once the restore is complete continue below

f. **For 2.5.6 to 2.5.7 anyrelease restore** where Ransomware Defender or Easy Auditor products are used then these additional steps are required to restore custom Ransomware Defender/Easy Auditor Settings:

   i. If you have this file present on your ECA node 1, /opt/superna/eca/conf/common/overrides/ThreatLevels.json, copy it to the new Eyeglass appliance into the /opt/superna/sca/data directory and apply same owner and permission as other files in the folder.

      1. sudo to root

      2. copy the ThreatLevel.json file into the /opt/superna/sca/data directory

      3. chmod 644 /opt/superna/sca/data/ThreatLevels.json

4. chown sca:users /opt/superna/sca/data/ThreatLevels.json

ii. On the new Eyeglass appliance if this file /opt/superna/sca/data/rwdefender/RSWSettings.json exists delete it

1. rm /opt/superna/sca/data/rwdefender/RSWSettings.json

iii. Then download the matching 2.5.7 run file for Eyeglass upgrade and copy it to the 2.5.7 appliance and then run it to restore custom Ransomware Defender settings - see instructions here for In Place Upgrade - Installer Download and Upgrade Procedures.

8. Check Post upgrade steps here

    a.

Post-Upgrade Steps (All Upgrade Paths)

Validate - Service account permissions, Eyeglass Job Status, Pool Mappings, Licenses and Cluster Inventory

1. **Mandatory Step - Check minimum permissions sudo section in this document are all in place for your release. This will generate errors if permissions are not correct.   Use this guide to review sudo permissions.**
2. Login to the new Eyeglass appliance and check:
    a. **Open Jobs window and verify all jobs modes are set correctly** and appear in either config sync or DFS section. The

screenshot taken before should be used to check the jobs are in the correct mode.

   i. If the jobs are in the wrong mode please set the mode correctly with the bulk actions menu.

3. **(only If IP Pool failover Mode is configured otherwise skip)** use the screenshot taken above to verify the synciq pool mappings using the DR Dashboard mapping screen to verify they look correct.
4. Open **License Manage**r Icon and verify Licenses are visible.
5. Open **Inventory Icon** and verifyClusters are displayed.
6. Log in to the Eyeglass web page and open the Eyeglass Main Menu -> Notification Center and verify that the Alarm Severity Filter is correctly set
7. 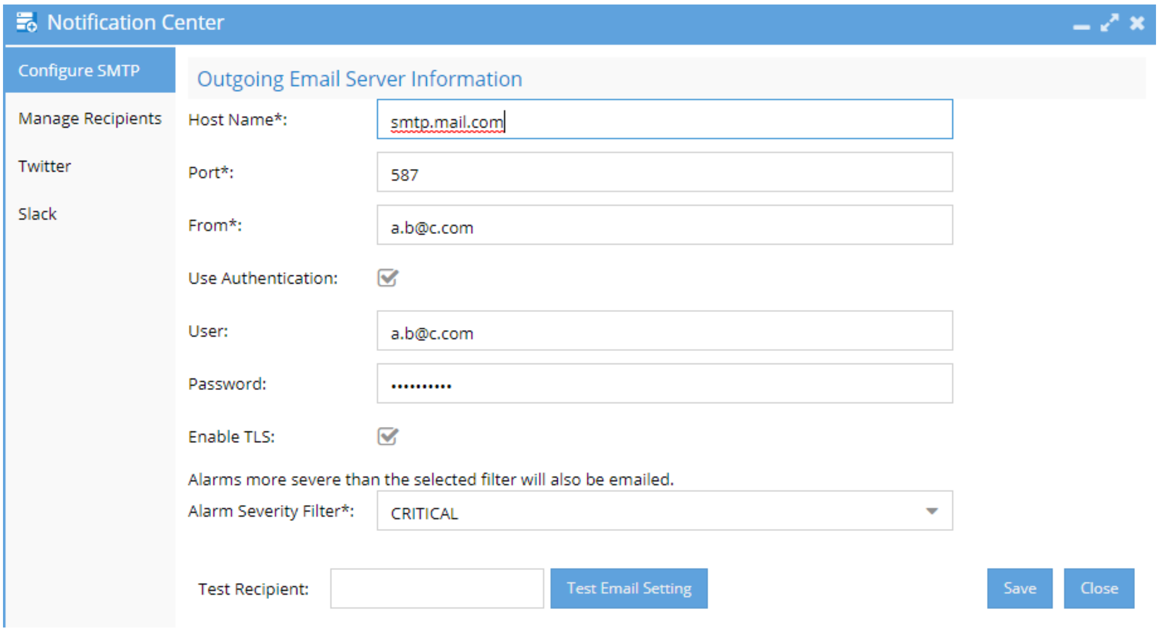

Host Name*: smtp.mail.com

Port*: 587

From*: a.b@c.com

Use Authentication: ☑

User: a.b@c.com

Password: ••••••••••

Enable TLS: ☑

Alarms more severe than the selected filter will also be emailed.

Alarm Severity Filter*: CRITICAL

Test Recipient:

8. And verify that the Email Recipients are correctly set with the correct Email Type.

9. 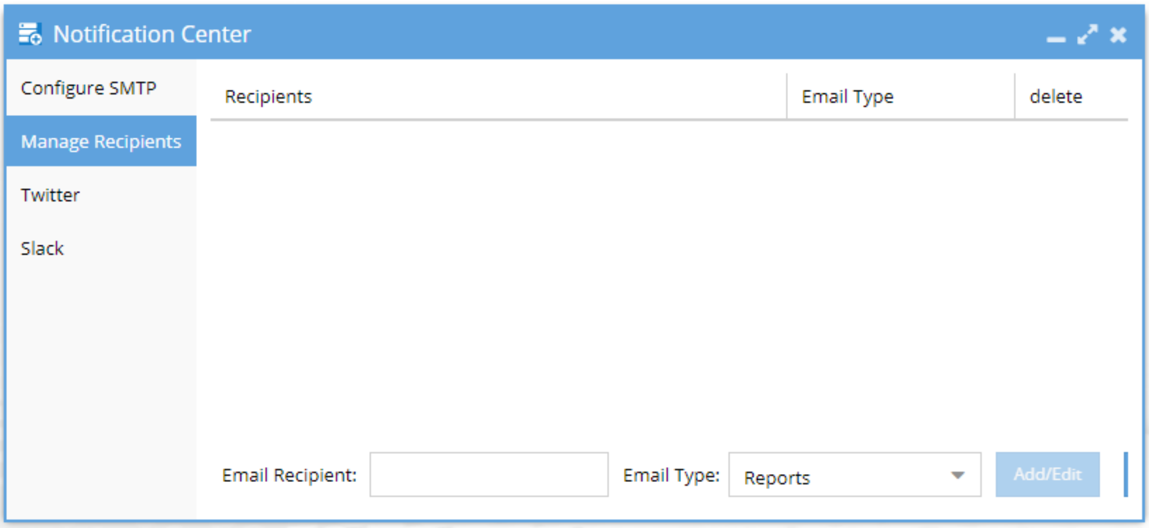

Email Type: Reports

10. done

Validate - Ransomware Defender, Easy Auditor, Performance Auditor License Assignment

1. **This step is mandatory to ensure licenses are assigned to the correct cluster. This release no longer supports auto assigned license mode to clusters**
2. Login to Eyeglass
3. Open **License Manager**
4. Click on **Licensed Devices tab**
   a. 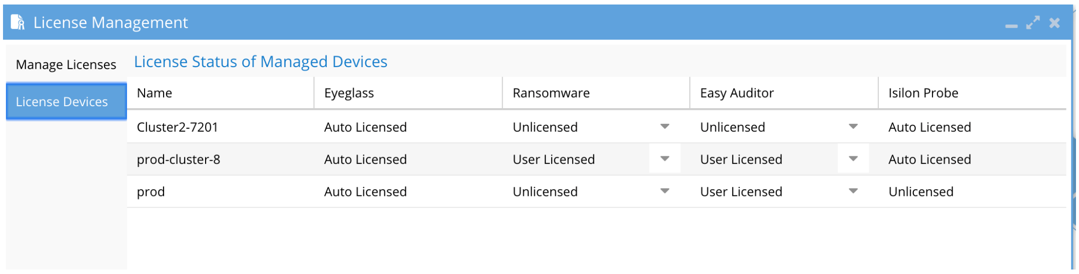

   License Status of Managed Devices

Name	Eyeglass	Ransomware	Easy Auditor	Isilon Probe
Cluster2-7201	Auto Licensed	Unlicensed	Unlicensed	Auto Licensed
prod-cluster-8	Auto Licensed	User Licensed	User Licensed	Auto Licensed
prod	Auto Licensed	Unlicensed	User Licensed	Unlicensed

5. **First STEP:** Set each cluster that should **NOT** be licensed to **Unlicensed status** using the drop down menu.

6. **2nd STEP:** Set each cluster listed to **User Licensed** for the product(s) that should be assigned to this cluster. Example the production writeable clusters should be set to **User Licensed** for Ransomware Defender or Easy Auditor.
7. **Click the submit button to save.**

Validate Ransomware Defender and Easy Auditor settings

Validate that Ransomware Defender and Easy Auditor settings preserved after upgrade:

1. **(if Ransomware Defender)** verify
    a. Flag as False Positive (2.5.6 and lower) / Learned Thresholds (2.5.7 and higher)
    b. Ignored List
    c. Threshold window all settings
    d. Allowed Files (2.5.6 and lower) / File Filters (2.5.7 and higher)
    e. Monitor Only Settings (2.5.7 and higher)
2. **(if Easy Auditor)** verify
    a. Active Auditor Triggers configured

3. Eyeglass VMware MarketPlace Deployment Guide

Home Top

- Overview
- Subscribe to the MarketPlace Eyeglass OVA
- How to Deploy Eyeglass from a Content Library

Overview

A new deployment option for on premise vCenter content library subscription option or VMware on AWS Cloud deployment. This guide will cover the content library solution with the VMware marketplace.

NOTE: Not all versions of vcenter content library are tested.   If you receive a deployment error due to compatibility issue with your version of vcenter, please download the OVA for direct deployment from https://support.superna.net and deploy the OVA following this guide.

Subscribe to the MarketPlace Eyeglass OVA

1. Click this link to go to the Eyeglass Marketplace listing.   You will need to login with your vmware account to be able to subscribe to the appliance.
2. Click subscribe button

a. 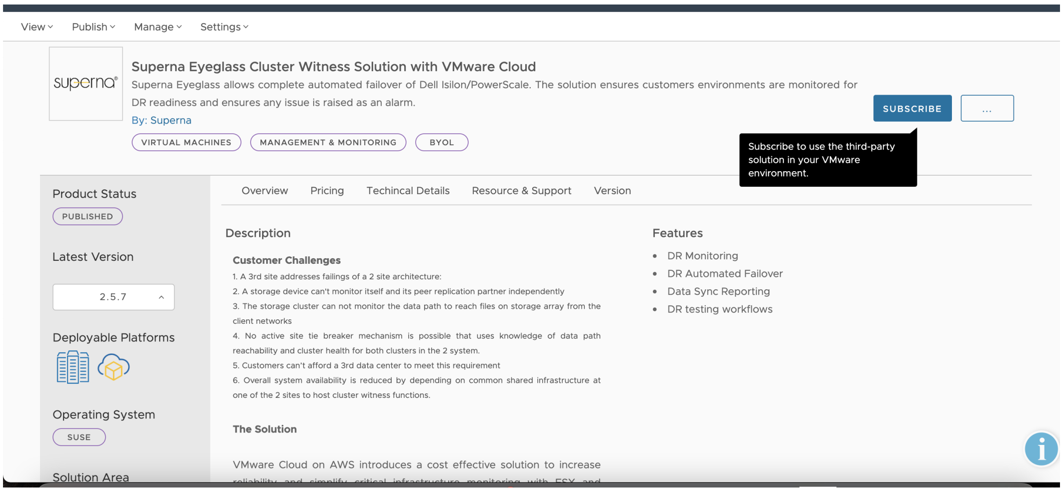

Superna Eyeglass Cluster Witness Solution with VMware Cloud

Superna Eyeglass allows complete automated failover of Dell Isilon/PowerScale. The solution ensures customers environments are monitored for DR readiness and ensures any issue is raised as an alarm.

By: Superna

VIRTUAL MACHINES | MANAGEMENT & MONITORING | BYOL

Subscribe to use the third-party solution in your VMware environment.

Product Status: PUBLISHED

Latest Version: 2.5.7

Deployable Platforms

Operating System: SUSE

Solution Area

Description

Customer Challenges

1. A 3rd site addresses failings of a 2 site architecture:
2. A storage device can't monitor itself and its peer replication partner independently
3. The storage cluster can not monitor the data path to reach files on storage array from the client networks
4. No active site tie breaker mechanism is possible that uses knowledge of data path reachability and cluster health for both clusters in the 2 system.
5. Customers can't afford a 3rd data center to meet this requirement
6. Overall system availability is reduced by depending on common shared infrastructure at one of the 2 sites to host cluster witness functions.

The Solution

VMware Cloud on AWS introduces a cost effective solution to increase

Features

- DR Monitoring
- DR Automated Failover
- Data Sync Reporting
- DR testing workflows

3. Select the on premise option and name the content library entry Eyeglass

a. 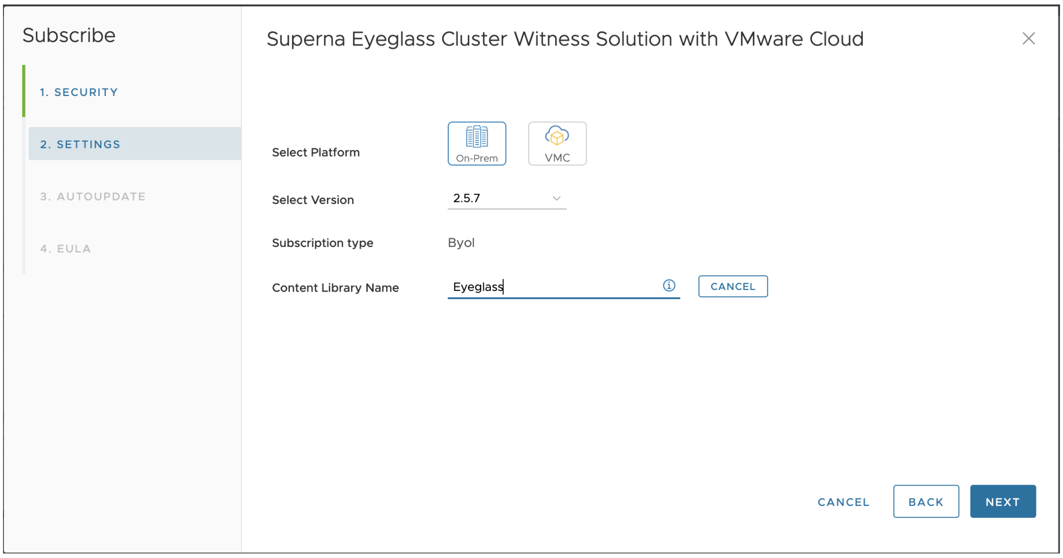

Superna Eyeglass Cluster Witness Solution with VMware Cloud

Select Platform: On-Prem

Select Version: 2.5.7

Subscription type: Byol

Content Library Name: Eyeglass

4. Enable auto update with 2 versions

a. 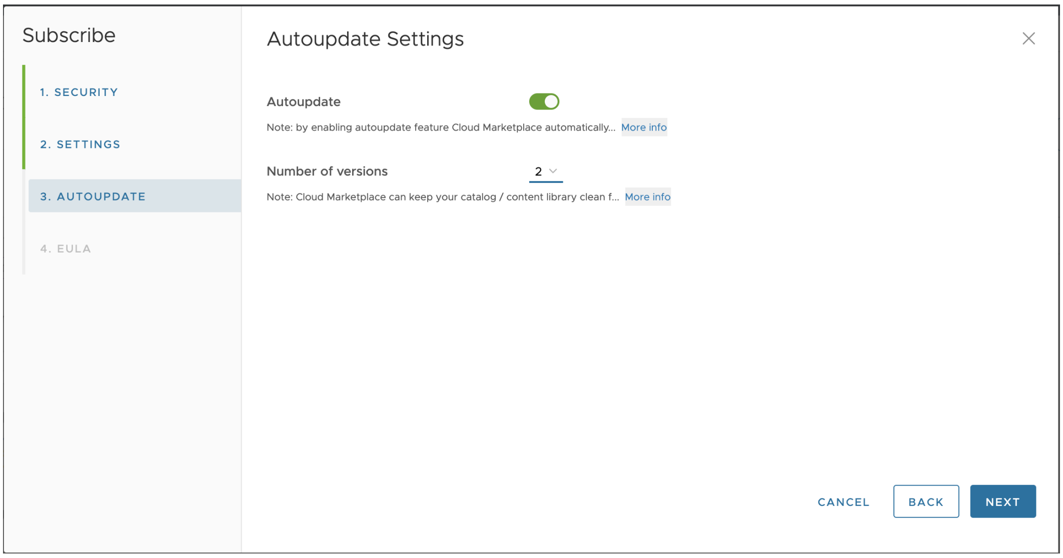

Subscribe

1. SECURITY
2. SETTINGS
3. AUTOUPDATE
4. EULA

Autoupdate Settings

Autoupdate

Note: by enabling autoupdate feature Cloud Marketplace automatically... More info

Number of versions 2

Note: Cloud Marketplace can keep your catalog / content library clean f... More info

5. Accept the EULA check box and click finish

6. Click the Check subscription status link

a. 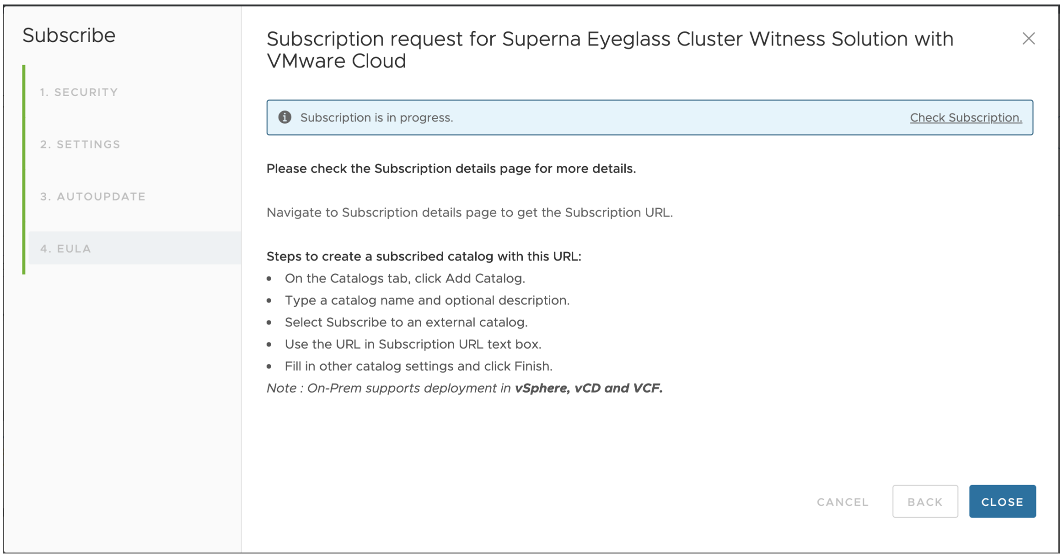

Subscribe

1. SECURITY
2. SETTINGS
3. AUTOUPDATE
4. EULA

Subscription request for Superna Eyeglass Cluster Witness Solution with VMware Cloud

Subscription is in progress. Check Subscription.

Please check the Subscription details page for more details.

Navigate to Subscription details page to get the Subscription URL.

Steps to create a subscribed catalog with this URL:

- On the Catalogs tab, click Add Catalog.
- Type a catalog name and optional description.
- Select Subscribe to an external catalog.
- Use the URL in Subscription URL text box.
- Fill in other catalog settings and click Finish.

Note : On-Prem supports deployment in ***vSphere, vCD and VCF.***

7. Scroll Down to get the subscription URL

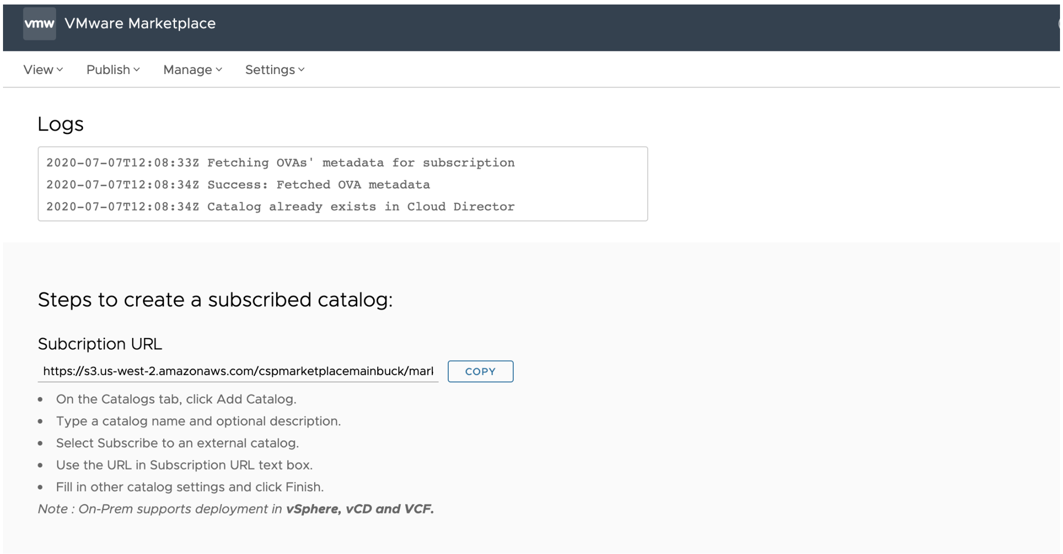

Logs

```
2020-07-07T12:08:33Z Fetching OVAs' metadata for subscription
2020-07-07T12:08:34Z Success: Fetched OVA metadata
2020-07-07T12:08:34Z Catalog already exists in Cloud Director
```

Steps to create a subscribed catalog:

Subcription URL

https://s3.us-west-2.amazonaws.com/cspmarketplacemainbuck/marl

- On the Catalogs tab, click Add Catalog.
- Type a catalog name and optional description.
- Select Subscribe to an external catalog.
- Use the URL in Subscription URL text box.
- Fill in other catalog settings and click Finish.

*Note : On-Prem supports deployment in **vSphere, vCD and VCF.***

a.

8. Copy the subscription URL

9. Login to vcenter and open the content library UI

10. Click + to add a new library

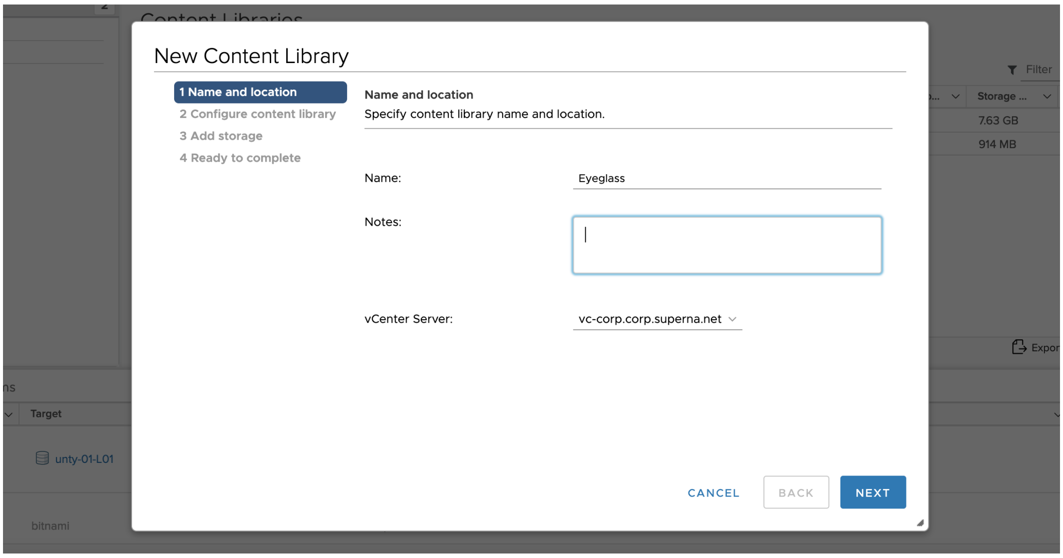

New Content Library

Name and location

Specify content library name and location.

Name: Eyeglass

Notes:

vCenter Server: vc-corp.corp.superna.net

a.

11. Click next and switch to subscription library and paste the url into the url input.

a. 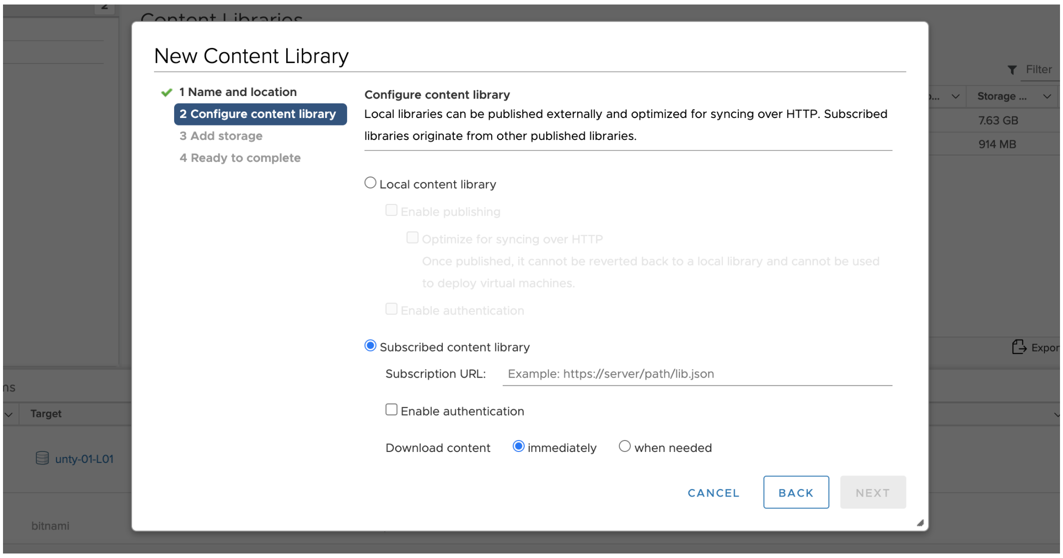

New Content Library

✓ 1 Name and location
2 Configure content library
3 Add storage
4 Ready to complete

Configure content library

Local libraries can be published externally and optimized for syncing over HTTP. Subscribed libraries originate from other published libraries.

○ Local content library

□ Enable publishing

□ Optimize for syncing over HTTP

Once published, it cannot be reverted back to a local library and cannot be used to deploy virtual machines.

□ Enable authentication

● Subscribed content library

Subscription URL: Example: https://server/path/lib.json

□ Enable authentication

Download content ● immediately ○ when needed

12. Select a data store to store the OVA image

a. 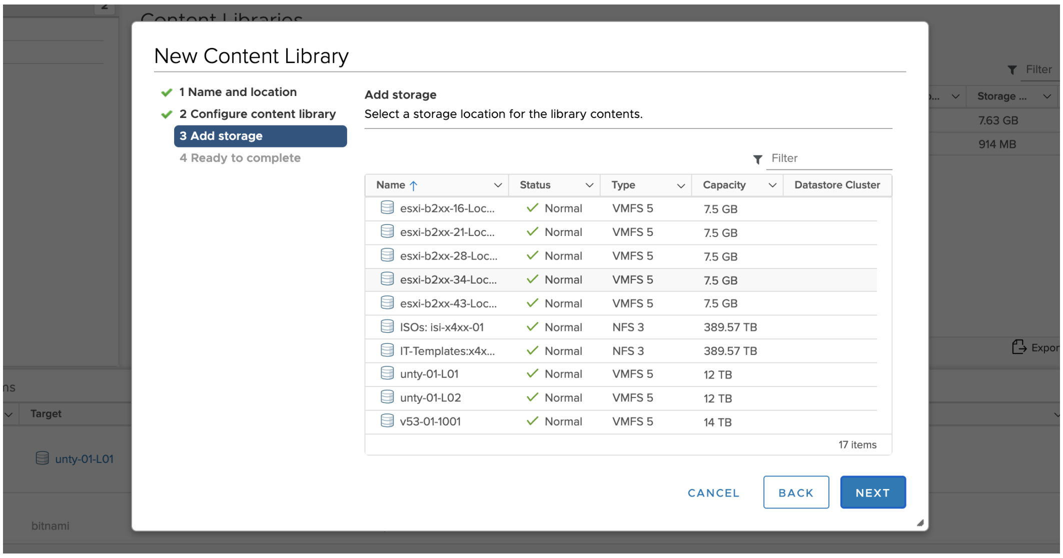

New Content Library

✓ 1 Name and location
✓ 2 Configure content library
3 Add storage
4 Ready to complete

Add storage

Select a storage location for the library contents.

Name ↑	Status	Type	Capacity	Datastore Cluster
esxi-b2xx-16-Loc...	✓ Normal	VMFS 5	7.5 GB	
esxi-b2xx-21-Loc...	✓ Normal	VMFS 5	7.5 GB	
esxi-b2xx-28-Loc...	✓ Normal	VMFS 5	7.5 GB	
esxi-b2xx-34-Loc...	✓ Normal	VMFS 5	7.5 GB	
esxi-b2xx-43-Loc...	✓ Normal	VMFS 5	7.5 GB	
ISOs: isi-x4xx-01	✓ Normal	NFS 3	389.57 TB	
IT-Templates:x4x...	✓ Normal	NFS 3	389.57 TB	
unty-01-L01	✓ Normal	VMFS 5	12 TB	
unty-01-L02	✓ Normal	VMFS 5	12 TB	
v53-01-1001	✓ Normal	VMFS 5	14 TB	

17 items

13. Click next and then finish to create the library

14. This will trigger a download and may take time to complete before the template is available for deployment.

15. Done.

How to Deploy Eyeglass from a Content Library

1. Login to vCenter
2. Open Content Library
3. Click on Eyeglass Content Library
4. Right Click the OVA template

a. 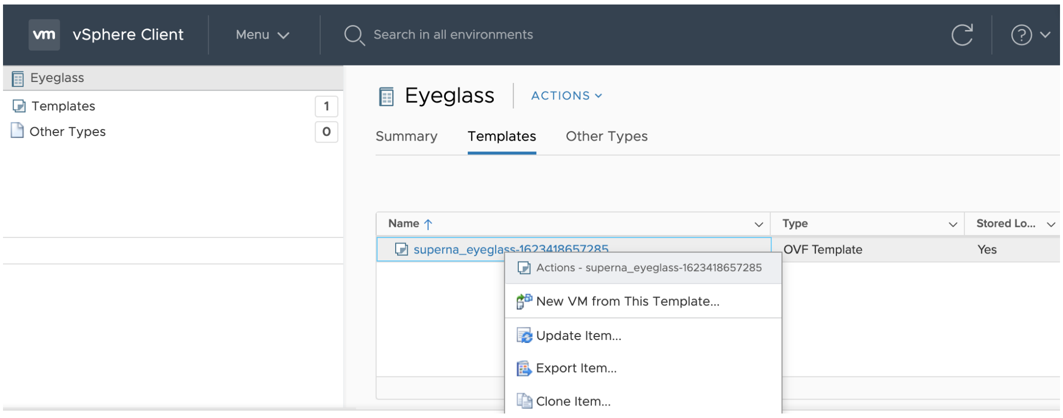

5. Select New VM from This Template Option
6. Follow the VM deployment Wizard and select options to deploy the Eyeglass VM.
7. On the VM Network tab make sure to **change the Destination Eyeglass Network to a network in your vmware environment**.

a. 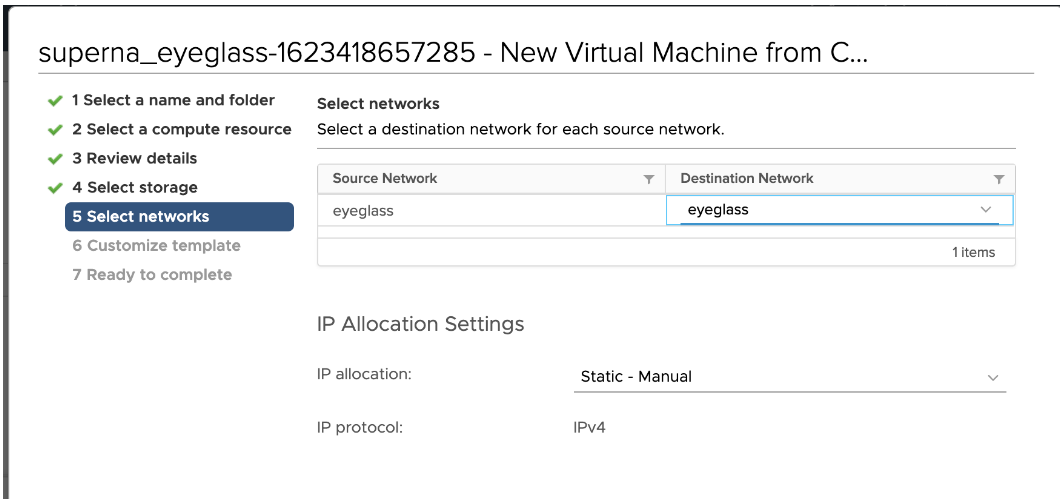

superna_eyeglass-1623418657285 - New Virtual Machine from C...

Select networks

Select a destination network for each source network.

Source Network	Destination Network
eyeglass	eyeglass

1 items

IP Allocation Settings

IP allocation: Static - Manual

IP protocol: IPv4

8. Data Disk Size must be left at 80GB. Do not change this value.

9. On the Customize Template tab --> Network section

   a. Edit the pre-populated IP address, network mask and gateway to match your network environment.

   b. Enter NTP IP address

   c. Enter a host name

   d. Enter DNS IP address (leave DNS search List blank)

   e. 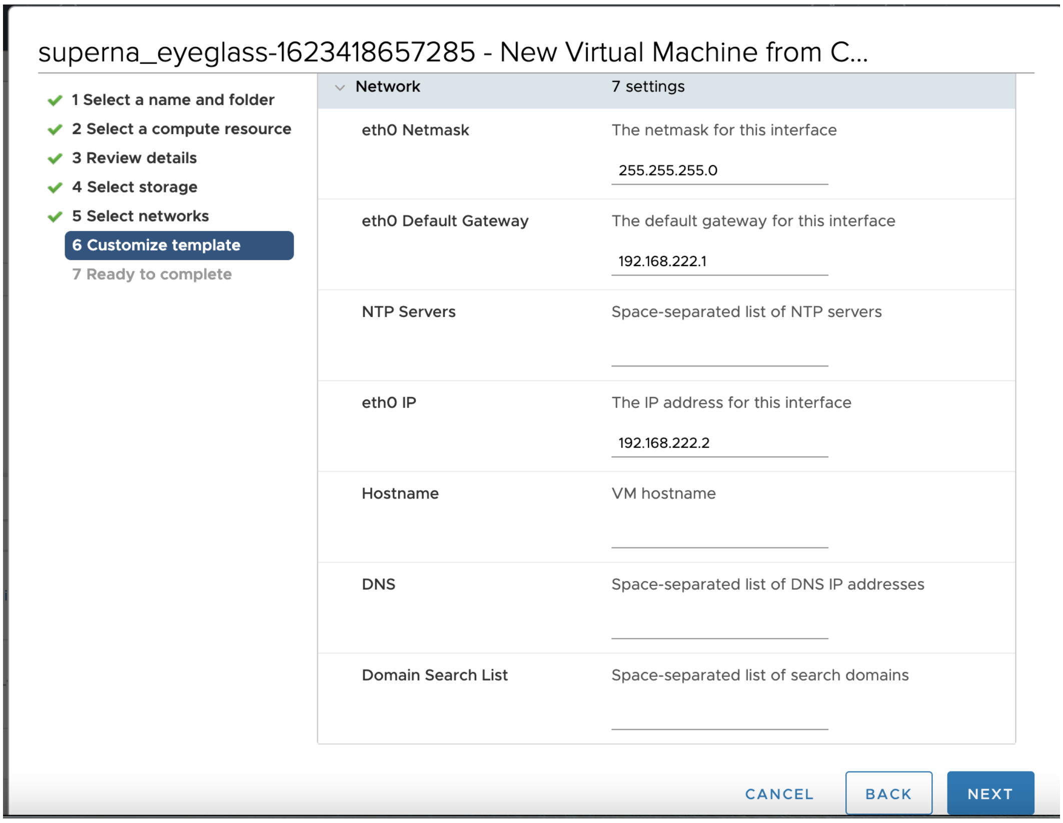

superna_eyeglass-1623418657285 - New Virtual Machine from C...

✓ 1 Select a name and folder
✓ 2 Select a compute resource
✓ 3 Review details
✓ 4 Select storage
✓ 5 Select networks
6 Customize template
7 Ready to complete

Network	7 settings
eth0 Netmask	The netmask for this interface 255.255.255.0
eth0 Default Gateway	The default gateway for this interface 192.168.222.1
NTP Servers	Space-separated list of NTP servers
eth0 IP	The IP address for this interface 192.168.222.2
Hostname	VM hostname
DNS	Space-separated list of DNS IP addresses
Domain Search List	Space-separated list of search domains

10. Complete the wizard to deploy the OVA appliance.

11. Complete the remaining installation tasks following the installation guide.

4. Eyeglass Search & Recover Installation and Upgrade Guide

Home Top

VMware ESX Host Compute Configuration and Cluster Sizing

See Sizing guide here.

Firewall Rules and Direction Table

NOTE: These rules apply to traffic into the VM and out of the VM. All ports must be open between VM's, private VLAN's and firewall between VM's is not supported.

Port	Direction	Function
Operating System Open Suse 15.3		It is customer responsibility to patch the operating system and allow Internet repository access for automatic patching. The OS is not covered by the support agreement.
443 HTTPS (TCP)	Browser → Search Cluster	UI Access, management tools (password protected)

22	Admin PC → Search Cluster	Management access to the CLI
8080 (HTTPS) over	Search Cluster → PowerScale	REST API Access
445 TCP SMB authentication	Search Cluster → PowerScale	SMB authentication of user name and password
NFS (optional for content)	Search Cluster → PowerScale	File Content Ingestion

Eyeglass Search & Recover Installation Prep Steps

The Search & Recover appliance is based on the ECA cluster architecture and has similar steps to complete installation.

Before you begin, collect the information below and verify all pre-requistes are completed:

1. See the sizing your cluster based on your environment to determine the number of nodes required. Guide Here.
2. Must have 4-7 IP addresses to assign during OVA deployment of the Search & Recover cluster depending on the size of the cluster (4 node or 7 node cluster)
    a. Must have DNS IP, router IP, subnet mask, NTP server IP
3. **Permissions for Service Account:** PowerScale REST API access with file traverse permissions to ingest files and directories by the Search & Recover cluster. See the minimum permissions guide

for a full list of permissions required for the eyeglass service account used by all ECA cluster products. Guide here.

4. (optional skip this section if Meta Data only indexing is needed)
    a. File Full Content data Indexing Requirements
    b. Must have DNS smartconnect name assigned to a management IP pool in the System zone for the NFS export used to ingest content from the snapshots folder.
    c. Get each Cluster GUID and name that will be indexed. Record these values for steps below.
        i. Login to the cluster OneFS GUI and open the Cluster Management --> General settings menu and record the cluster GUID and cluster name. Example below.
    d.

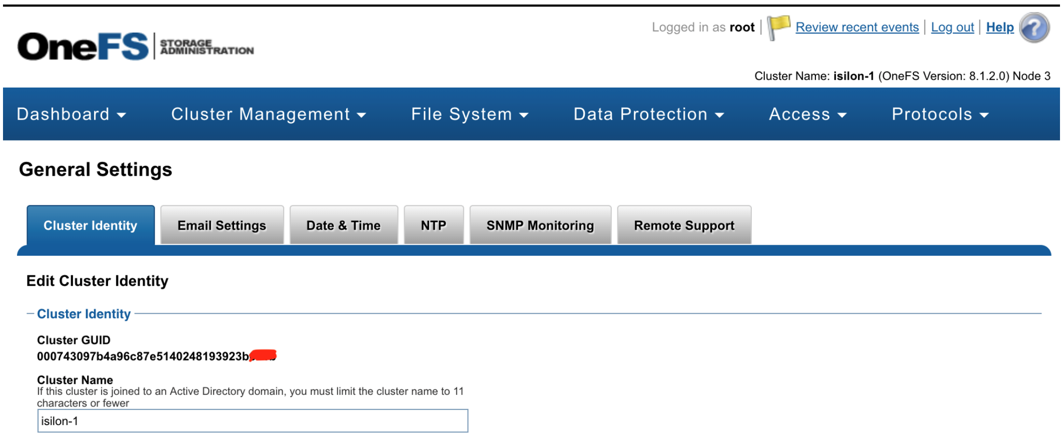

Cluster Name: isilon-1 (OneFS Version: 8.1.2.0) Node 3

General Settings

Edit Cluster Identity

Cluster Identity

Cluster GUID
000743097b4a96c87e5140248193923b[redacted]

Cluster Name
If this cluster is joined to an Active Directory domain, you must limit the cluster name to 11 characters or fewer
isilon-1

    e. Repeat for each cluster that will be licensed and indexed.
    f. Create an NFS export in the System Access Zone for full content on all clusters that will be indexed. See example below where the IP addresses entered are the Search & Recover cluster IP addresses. The export is created on the /ifs/.snapshot directory with root client list used and add all nodes except node 1.

g. 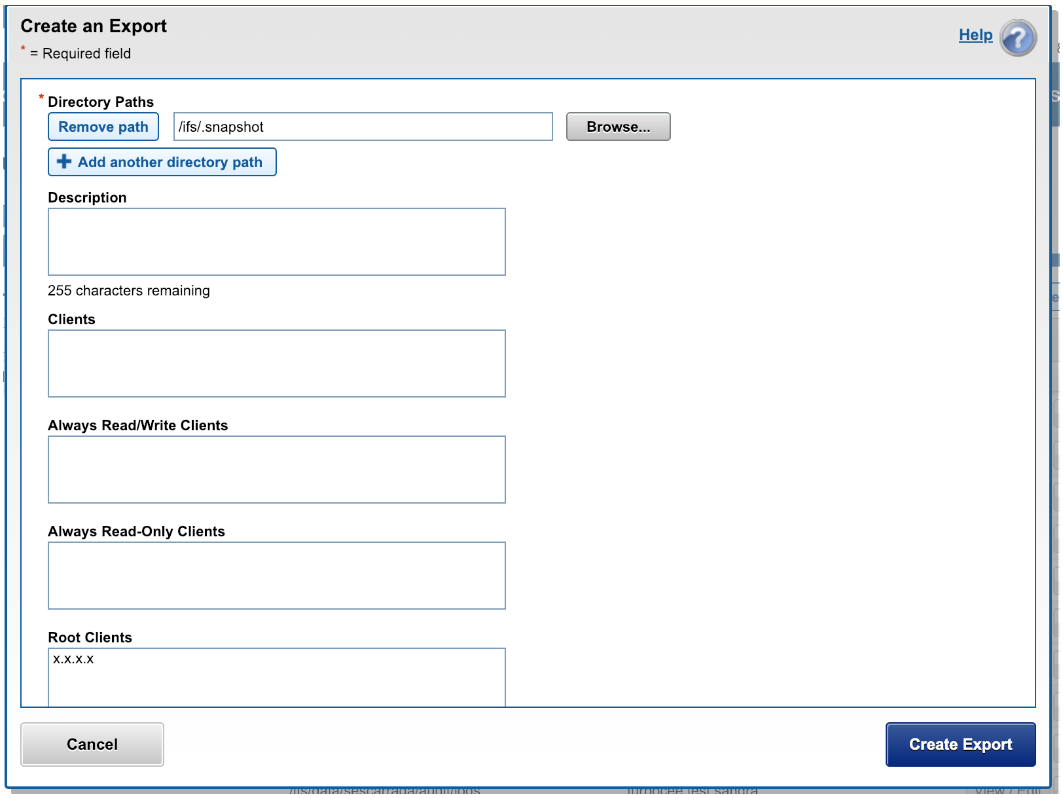

Create an Export

* = Required field

* Directory Paths

Remove path /ifs/.snapshot

Description

255 characters remaining

Clients

Always Read/Write Clients

Always Read-Only Clients

Root Clients

x.x.x.x

Search & Recover OVA Deployment and Cluster VM Configuration

1. Download the OVA following instructions here
2. **Deploy with vCenter**

   NOTE: vCenter 6.5 and 6.7 use FLASH or FLEX interface. HTML5 interface is not supported.

   a. Select the OVF file with the required node count based on cluster sizing pre-requisite step.

   b. Set the node ip addresses, gateway IP, DNS, NTP IP and index volume size during the OVA deployment. Use the disk size values from this sizing table.

c. **Note: Data disk size value MUST be left at the default.**

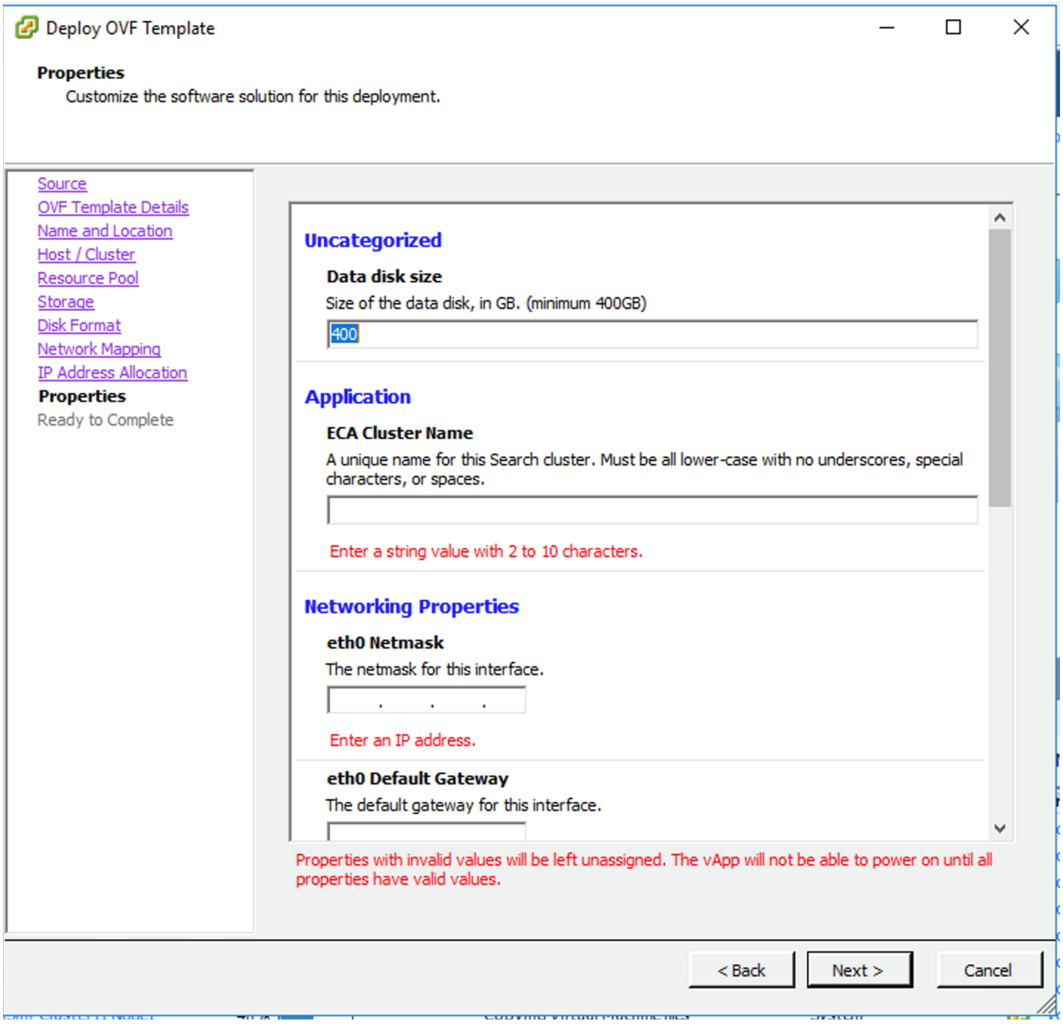

Deploy OVF Template

Properties
Customize the software solution for this deployment.

Uncategorized

Data disk size
Size of the data disk, in GB. (minimum 400GB)
400

Application

ECA Cluster Name
A unique name for this Search cluster. Must be all lower-case with no underscores, special characters, or spaces.

Enter a string value with 2 to 10 characters.

Networking Properties

eth0 Netmask
The netmask for this interface.

Enter an IP address.

eth0 Default Gateway
The default gateway for this interface.

Properties with invalid values will be left unassigned. The vApp will not be able to power on until all properties have valid values.

d. Set the **ECA cluster name** **(no special characters should be used, all lower case)**

e. Power on the OVA

f. SSH to the node 1 IP address

g. Login with user **ecaadmin** and default password **3y3gl4ss**

3. **Configure and test NTP (NOTE if Internet access is allowed for NTP protocol, then no configuration is needed, use this procedure to test access and or change NTP to an internal server IP)**

   a. **sudo -s**

b. **Enter ecaadmin password**

c. **type yast** **(NOTE navigation uses the Arrow keys and TAB key and shift TAB key to move between selections on the UI)**

d. **Select Network Services to enter NTP Configuration**

i. 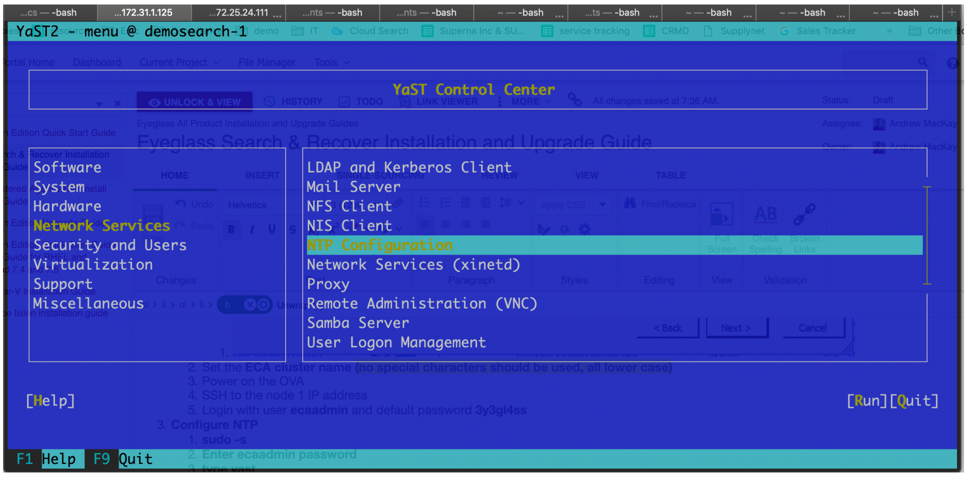

```
YaST2 - menu @ demosearch-1

                         YaST Control Center

Software               LDAP and Kerberos Client
System                 Mail Server
Hardware               NFS Client
Network Services       NIS Client
Security and Users     NTP Configuration
Virtualization         Network Services (xinetd)
Support                Proxy
Miscellaneous          Remote Administration (VNC)
                       Samba Server
                       User Logon Management

[Help]                                              [Run][Quit]

F1 Help  F9 Quit
```

e. Test Access to Internet NTP servers

i. TAB to the **Edit** button and press enter

ii. 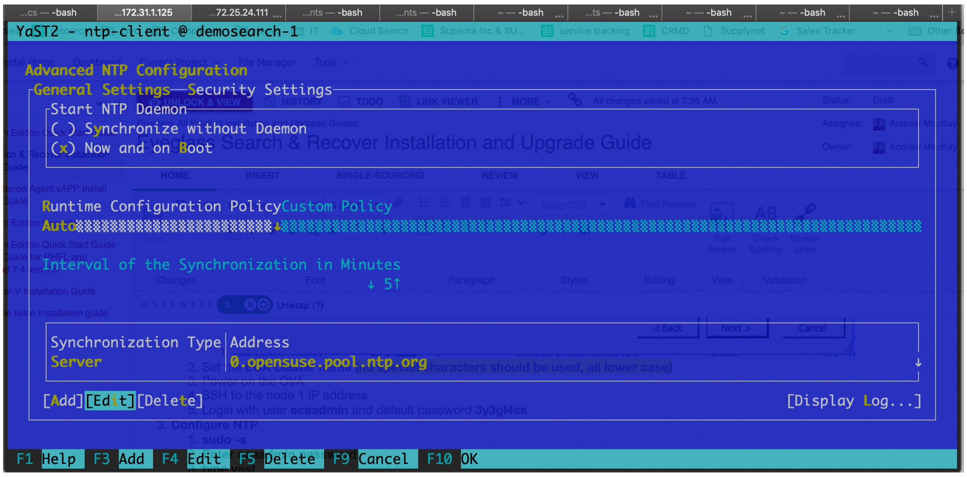

```
YaST2 - ntp-client @ demosearch-1
Advanced NTP Configuration
 General Settings  Security Settings
  Start NTP Daemon
  ( ) Synchronize without Daemon
  (x) Now and on Boot

 Runtime Configuration PolicyCustom Policy
 Auto
 Interval of the Synchronization in Minutes
                                  ↓ 5↑

  Synchronization Type|Address
  Server              |0.opensuse.pool.ntp.org

 [Add][Edit][Delete]                        [Display Log...]

F1 Help  F3 Add  F4 Edit  F5 Delete  F9 Cancel  F10 OK
```

iii. Now **Tab** to the **Test** button and press enter

iv. 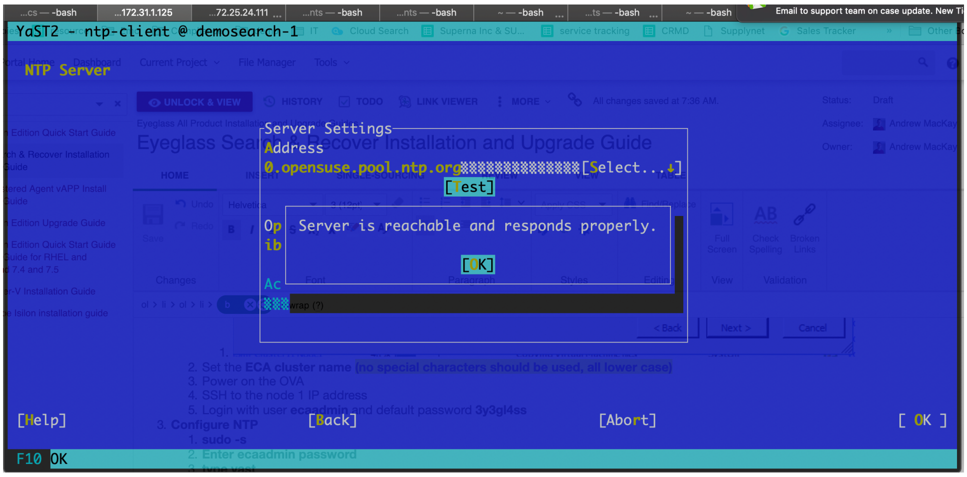

```
YaST2 - ntp-client @ demosearch-1
NTP Server

Server Settings
Address
0.opensuse.pool.ntp.org              [Select...↓]
                [Test]

   Server is reachable and responds properly.
                   [OK]

[Help]        [Back]          [Abort]         [ OK ]
F10 OK
```

v. If Internet access is enabled then test response will look like the image above.

vi. If it fails

1. Replace the NTP server entry with your Internal NTP server IP address
2. Use the **TAB** key to select the OK button, press enter
3. Press **F10** to exit the NTP configuration screen
4. Use the **TAB** key to Select the YAST **Exit** button.

vii. **Repeat NTP test and/or configuration on ALL Search nodes**

4. <u>**Configure Cluster Master Configuration file**</u>
5. <u>**IMPORTANT: Next step must be run as ecaadmin user. If still root user from previous steps exit before continuing.**</u>
6. <u>**IMPORTANT: Next step must be run from Search node 1**</u>
   a. Run the cluster setup command (sets up SSH between nodes, configures master file, set password for administration WebUI)

i. NOTE: This will setup ssh access between nodes and configure the cluster.

ii. Type y or yes when prompted to continue.

iii. You will be prompted to enter the default password to complete this setup process. The default password is 3y3gl4ss.

iv. You will prompted to login to each node during this one time process

v. Run setup command: ecactl components configure-nodes

vi. After the setup completes on all nodes verify the contents of the master configuration file

1. cat /opt/superna/eca/eca-env-common.conf
2. Verify the ip addresses are listed for each node
3. Verify the cluster name is correctly set in the file variable export **ECA_CLUSTER_ID=**

b. (Skip section if you only plan on indexing metadata) Configure NFS Export for Content Ingestion ONLY

i. Search and Recover uses PowerScale snapshots to ingest content. Follow the steps below to add the export to each of the VM's.

ii. **What you will need to complete this step on nodes 2 - x (where x is the last node IP in the cluster):**

1. Cluster GUID and cluster name for each cluster to be indexed

2. Cluster name as shown on top right corner after login to OneFS GUI

iii. **Change to Root user**

1. sudo -s
2. enter ecaadmin password 3y3gl4ss

iv. **Create local mount directory (repeat for each cluster)**

1. mkdir -p /opt/superna/mnt/search/GUID/clusternamehere/ (replace GUID and clusternamehere with correct values)
2. Use this command to run against all nodes, you will be prompted for ecaadmin password on each node.
   a. ecactl cluster exec "sudo mkdir -p /opt/superna/mnt/search/00505699937a5e1f5b5d8b2342c2c3fe9fd7/prod-cluster"

v. **(Skip this section if you only plan on indexing metadata) Configure automatic NFS mount Required for Full Text Indexing**

1. **Prerequisites**
   a. This will add a mount for content indexing to FSTAB on all nodes
   b. Build the mount command using cluster guid and cluster name replacing the yellow highlighted sections with correct values for

your cluster. **NOTE: This is only an example**

   c. You will need a smartconnect name to mount the snapshot folder on the cluster. The Smartconnect name should be a system zone IP pool

   d. Replace smartconnect FQDN and <> with a DNS smartconnect name

   e. Replace <GUID> with cluster GUID

   f. Replace <name> with the cluster name

2. **On each node in the cluster:**

   a. ssh to the node as ecaadmin

   b. sudo -s

   c. enter ecaadmin password

   d. echo '<CLUSTER_NFS_FQDN>:/ifs/.snapshot /opt/superna/mnt/search/<GUID>/<NAME> nfs defaults,nfsvers=3 0 0'| sudo tee -a /etc/fstab

   e. mount -a

   f. mount to verify the mount

   g. exit

   h. Login to next node via ssh

3. repeat steps on each node

4. **NOTE: The mount must exist BEFORE you start the cluster.**

7. Start up the cluster
8. IMPORTANT: Next step must be run as ecaadmin user. If still root user from previous steps exit before continuing.
9. IMPORTANT: Next step must be run from Search node 1
   ecactl cluster up
   a. Wait until all steps complete on all nodes verify no error messages are deplayed
   b. **Administration WebUI password**
      i. When prompted enter a secure strong password 8 characters or longer with upper case, lower case , numbers and special characters.   This password protects access to administration webUI.
10. **Get the Search appliance id and make a record of it - will be required to retrieve license**
    ecactl version
11. **Deployment done**
12. **NEXT Steps Search & Recover Cluster Logical Configuration**
    a. Configuration steps to add licenses, add clusters, add index folders is covered in the **Quick Start Steps** of the Search & Recover Admin guide.

How To Upgrade Search & Recover Cluster

1. NOTE: Always take a vmware level snapshot of all VM's before any upgrade to be able to roll back to a previous release.
2. NOTE: Contact support before any upgrade to get upgrade changes between releases for compatibility of the index.
3. NOTE: New permissions are required for the latest release for the service account used by Search Update the minimum permissions for the eyeglassSR user follow the guide here and apply all permissions again to ensure all permissions are correct.
4. Login to the support site https://support.superna.net and download the offline Search & Recover upgrade file
5. Using a tool like Winscp copy the upgrade file to node 1 using ecaadmin user to authenticate
6. shutdown the cluster
    a. Login to node 1 as ecaadmin over ssh
    b. run the command "ecactl cluster down"
    c. wait for the cluster to shutdown
    d. Update EyeglassAdminSR service account with an additional permission or verify the service account has the correct permissions. Verify all permissions listed here have been applied by running the command **isi auth roles view EyeglassAdminSR .** The networking previlege is now required.
    e. Modify the install file copied to the cluster node 1
        i. Assuming the file was copied to default location in /home/ecaadmin
        ii. cd /home/ecaadmin

iii. chmod 777 **<name of install file here>**

f. Run the installer

i. **./name of install file here**

ii. when prompted enter the ecaadmin password

iii. wait for all nodes to be upgraded to the new version

g. Start the cluster

i. ecactl cluster up

ii. wait until all nodes are started

h. done.

4.1. Search & Recover Hyper-V Deployment

Home Top

- Deploy HyperV Search & Recover Appliance

Deploy HyperV Search & Recover Appliance

1. Add 16GB RAM
2. Deploy 2nd vhdx disk for data [400GB]
    a. Configure /opt/superna/eca/eca-env-common.conf with ip addresses of all ECA nodes on node 1 master.
    b. 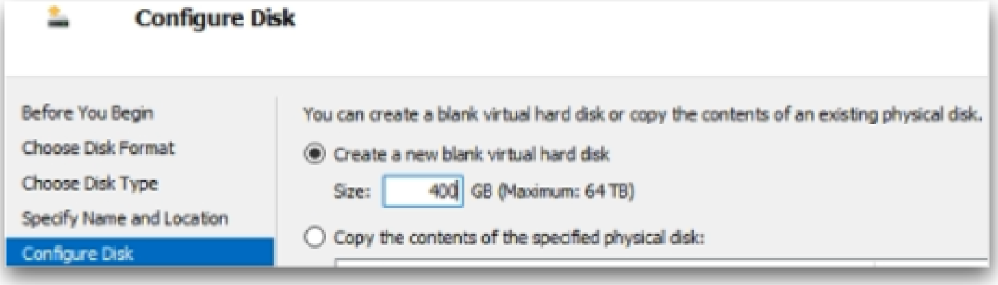
    Configure Disk
    Before You Begin
    Choose Disk Format
    Choose Disk Type
    Specify Name and Location
    Configure Disk
    You can create a blank virtual hard disk or copy the contents of an existing physical disk.
    ● Create a new blank virtual hard disk
    Size: 400 GB (Maximum: 64 TB)
    ○ Copy the contents of the specified physical disk:
    c. 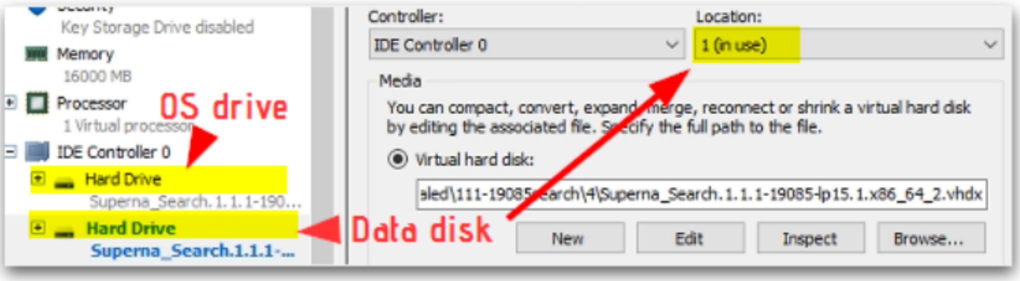
    Key Storage Drive disabled
    Memory
    16000 MB
    Processor
    1 Virtual processor
    IDE Controller 0
    Hard Drive
    Superna_Search.1.1.1-190...
    Hard Drive
    Superna_Search.1.1.1-...
    Controller: IDE Controller 0
    Location: 1 (in use)
    Media
    You can compact, convert, expand, merge, reconnect or shrink a virtual hard disk by editing the associated file. Specify the full path to the file.
    ● Virtual hard disk:
    aled\111-19085earch\4\Superna_Search.1.1.1-19085-lp15.1.x86_64_2.vhdx

3. Boot up and wait for 15-20 minutes to allow script to run
4. Login to node 1

5. Check the on-boot script log and run the command [Also do this on on all node]

    a. tail -4 /var/log/superna-on-boot.log

    b. 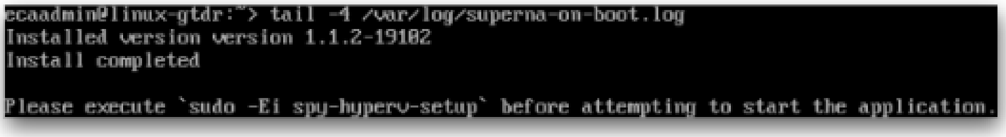

```
ecaadmin@linux-gtdr:~> tail -4 /var/log/superna-on-boot.log
Installed version version 1.1.2-19102
Install completed

Please execute `sudo -Ei spy-hyperv-setup` before attempting to start the application.
```

6. Run the command and enter `ecaadmin` passwd [3y3gl4ss] then configure network for the nodes

    a. sudo -Ei spy-hyperv-setup

    b. 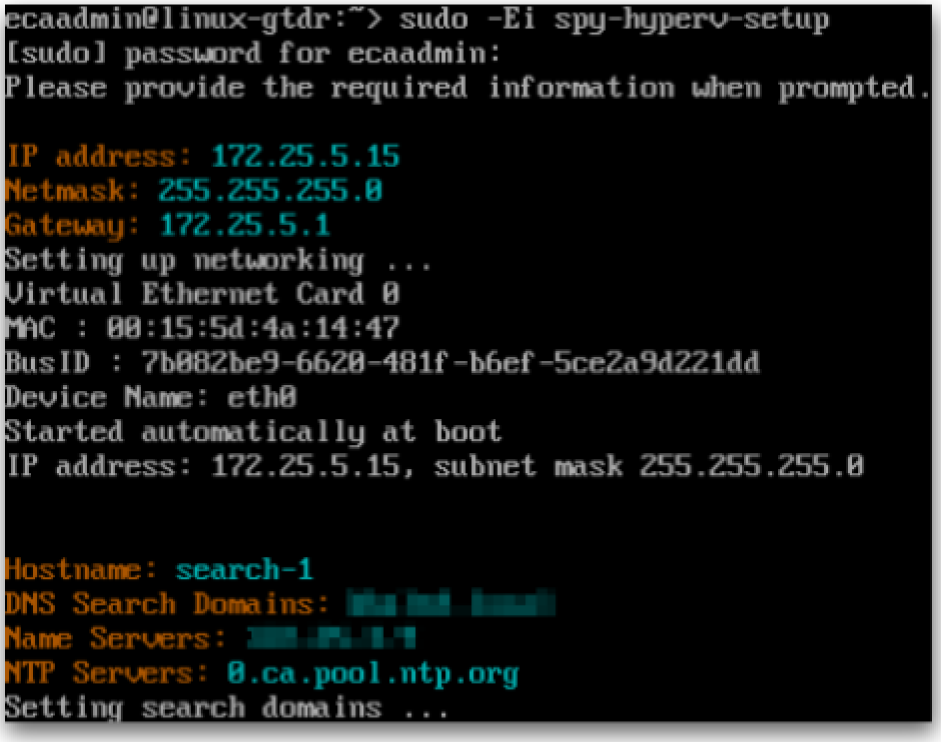

```
ecaadmin@linux-gtdr:~> sudo -Ei spy-hyperv-setup
[sudo] password for ecaadmin:
Please provide the required information when prompted.

IP address: 172.25.5.15
Netmask: 255.255.255.0
Gateway: 172.25.5.1
Setting up networking ...
Virtual Ethernet Card 0
MAC : 00:15:5d:4a:14:47
BusID : 7b082be9-6620-481f-b6ef-5ce2a9d221dd
Device Name: eth0
Started automatically at boot
IP address: 172.25.5.15, subnet mask 255.255.255.0


Hostname: search-1
DNS Search Domains: [redacted]
Name Servers: [redacted]
NTP Servers: 0.ca.pool.ntp.org
Setting search domains ...
```

7. [IMPORTANT: Run this step after 7-8] Node1 is Master node. The rest are child nodes. Enter `y` for node 1 ONLY. Do not press `y` until node 2-4 are marked as `n`.Go to the next step.

    a. 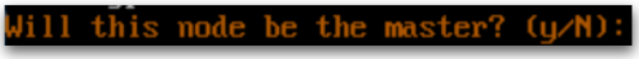

```
Will this node be the master? (y/N):
```

8. Repeat step 4-6 on node 2-4
9. Node 2-4, when prompted for `master` node, Enter `n`
10. Go to step 7 and press `y` to complete the master node setup
11. Enter cluster name and the child nodes IP [space separated]

    a.

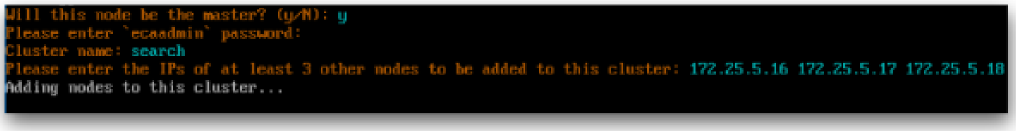

```
Will this node be the master? (y/N): y
Please enter `ecaadmin` password:
Cluster name: search
Please enter the IPs of at least 3 other nodes to be added to this cluster: 172.25.5.16 172.25.5.17 172.25.5.18
Adding nodes to this cluster...
```

12. Wait for it to finish then exit the nodes and log back in

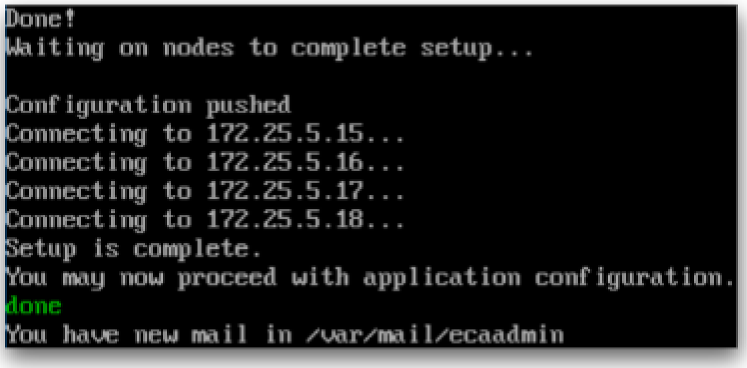

```
Done!
Waiting on nodes to complete setup...

Configuration pushed
Connecting to 172.25.5.15...
Connecting to 172.25.5.16...
Connecting to 172.25.5.17...
Connecting to 172.25.5.18...
Setup is complete.
You may now proceed with application configuration.
done
You have new mail in /var/mail/ecaadmin
```

    a.

13. Use putty to access node 1
14. Done
15. Configuration tasks should be completed following the Normal guide here.

4.2. How to Migrate Eyeglass Search Appliance Index from ECA OpenSuse 15.1 to OpenSuse 15.3 OS

Home Top

When to use this procedure

This procedure allow moving an index from an old appliance with 15.1 OS to a new appliance running opensuse 15.3. It will require VMware access to edit VM's and attach VMDK disk from the old appliance to the new appliance.

Pre-Migrate steps

1. Clone the existing VM to protect the index in case a roll back is required or backup the VM with vmware backup tools.
2. NOTE: No snapshots can exist on the VM for roll back since the VMDK disks cannot be moved to another VM if snapshots exist.
3. NOTE: if no backup has been completed and issues impact the index disk during migration the only method to revert requires a backup or clone of the VMDK. The other option is re-index the data using the new appliance vs backup or cloning the old appliance.

High Level steps

1. deploy Superna Search appliance
2. configure NFS mounts
3. Create a backup of old appliance
4. power on new vApp to make sure all nodes are up and reachable, power off
5. remove without delete hard disk 2 from old vApp node 1-N
6. remove/delete hard disk 2 on node1-4 of new vApp and add removed hard disks from old vApp to new vApp node 1-N
7. power on new vApp
8. copy zip to new vApp and restore configuration from backup

Detailed Steps

1. Create Search Cluster backup on old appliance:
   a. SSH to ECA node 1 as user: ecaadmin
   b. Type command: ecactl cluster backup
   c. Follow this procedure to retrieve this backup file. Find the section on how to retrieve the backup file on this link.
2. Create logs folder under **/opt/data/superna** on all ECA nodes
   a. cd /opt/data/superna
   b. mkdir logs
   c. chmod 775 logs
   d. Repeat the above steps on all ECA nodes
3. Bring down old ECA cluster
   a. SSH to ECA node 1 as user: ecaadmin

    b. Type command: ecactl cluster down

4. Using vCenter UI, power off the vApp
    a. Download new Eyeglass Search OVF based on OpenSuse 15.3 and deploy these vApp, as per documented install procedure
5. Configure the new Eyeglass Search vApp to have the same configuration as the old Eyeglass Search vApp. Assign the following when deploying new vApp
    a. Same ECA cluster name
    b. Same IP Addresses for ECA nodes
6. Once the deployment of new vApp has been completed, power on this new vApp and then
    a. SSH to this new ECA node 1 as user ecaadmin
    b. Type command: ecactl components configure-nodes
    c. Edit the Search Cluster backup zip file (Open this zip file by using zip tool utility e.g. 7-zip) to remove known_host file from each node folders in that backup zip file, under path /<node-x>/home/ecaadmin/.ssh/
    d. Copy the updated Search cluster backup zip file to this new ECA node 1. Use WinSCP
    e. Restore from the backup use command: **ecactl cluster restore --path <path-to-copied-backup-file>**
7. Once restore has completed, create local directory on ECA node 2 - last node for mounting PowerScale Snapshot NFS export (***Only If require Content Ingestion***)
    a. ssh to ECA node 1 as user ecaadmin
    b. sudo su -

c. mkdir -p /opt/superna/mnt/search/***<GUID-of-PowerScale-Cluster>***/***<Cluster-name>***

d. Repeat the above steps for ECA node 3 - last node

e. Modify file /etc/fstab on ECA node 2 - ECA last node.

f. Open the copied fstab file from old ECA node 2 and copy the mount to the PowerScale Snapshot folder setting line and insert this into the fstab file on new ECA node 2 - last node

g. On each node (Node 2 - last node) complete these steps:

h. ssh ecaadmin@x.x.x.x (ip of each eca node)

i. sudo -s (enter ecaadmin password when prompted)

j. nano /etc/fstab

k. paste mount line into the file control+x to save and exit

l. Test the mount in fstab on the node

m. NOTE: you should still be the root user from above steps

n. type command --> mount -a

o. if no mount error you should not see any output from this command

p. Check mount and type --> mount [enter]

q. Review the output to make sure the mount is visible

r. Repeat all steps above on ECA node 2 - last node

8. Upgrade to the latest code

   a. copy upgrade file to node 1

   b. chmod 777 <upgrade file>

   c. ./<upgrade file>

9. done

1. Vmware Steps to move the index to the new appliance

    1. Edit the setting of the new ECA VM node 1 (Warning: Do this on the new ECA vApp VMs, do not do this step on the old ECA vApp VMs), and remove Hard Disk 2 with option “Remove from virtual machine and delete files from disk”. Example:

    2. 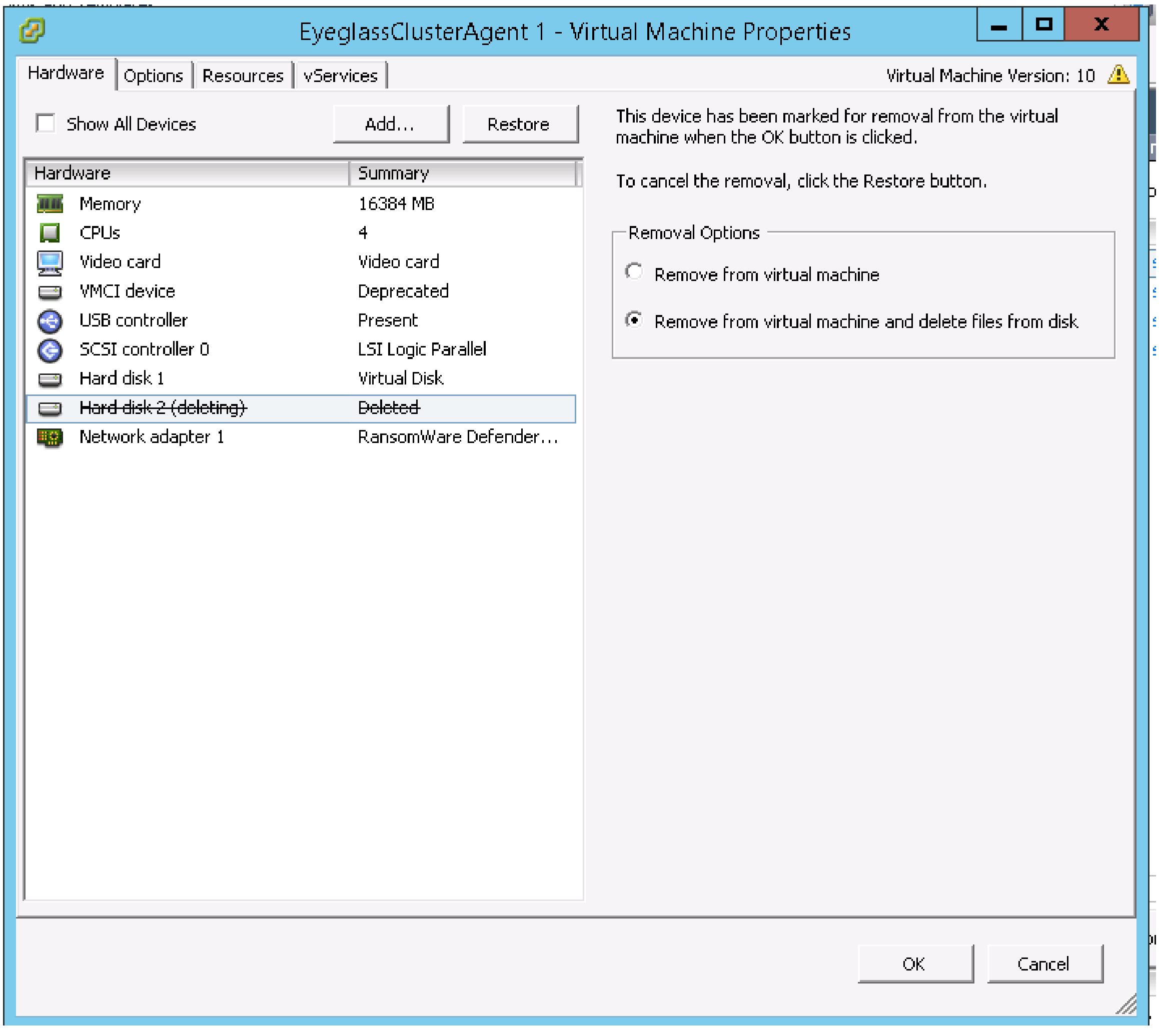

EyeglassClusterAgent 1 - Virtual Machine Properties

Virtual Machine Version: 10

□ Show All Devices

Hardware	Summary
Memory	16384 MB
CPUs	4
Video card	Video card
VMCI device	Deprecated
USB controller	Present
SCSI controller 0	LSI Logic Parallel
Hard disk 1	Virtual Disk
~~Hard disk 2 (deleting)~~	~~Deleted~~
Network adapter 1	RansomWare Defender...

This device has been marked for removal from the virtual machine when the OK button is clicked.

To cancel the removal, click the Restore button.

Removal Options

○ Remove from virtual machine

● Remove from virtual machine and delete files from disk

    3. Edit the setting of the old ECA VM node 1 (Now do this step on the old ECA vApp VMs) and record the Datastore Disk File location for Hard Disk 2 and then remove only from VM inventory (Warning: Only remove from VM, but do not delete files from disk). Chose: “Remove from virtual Machine”. Example:

4. 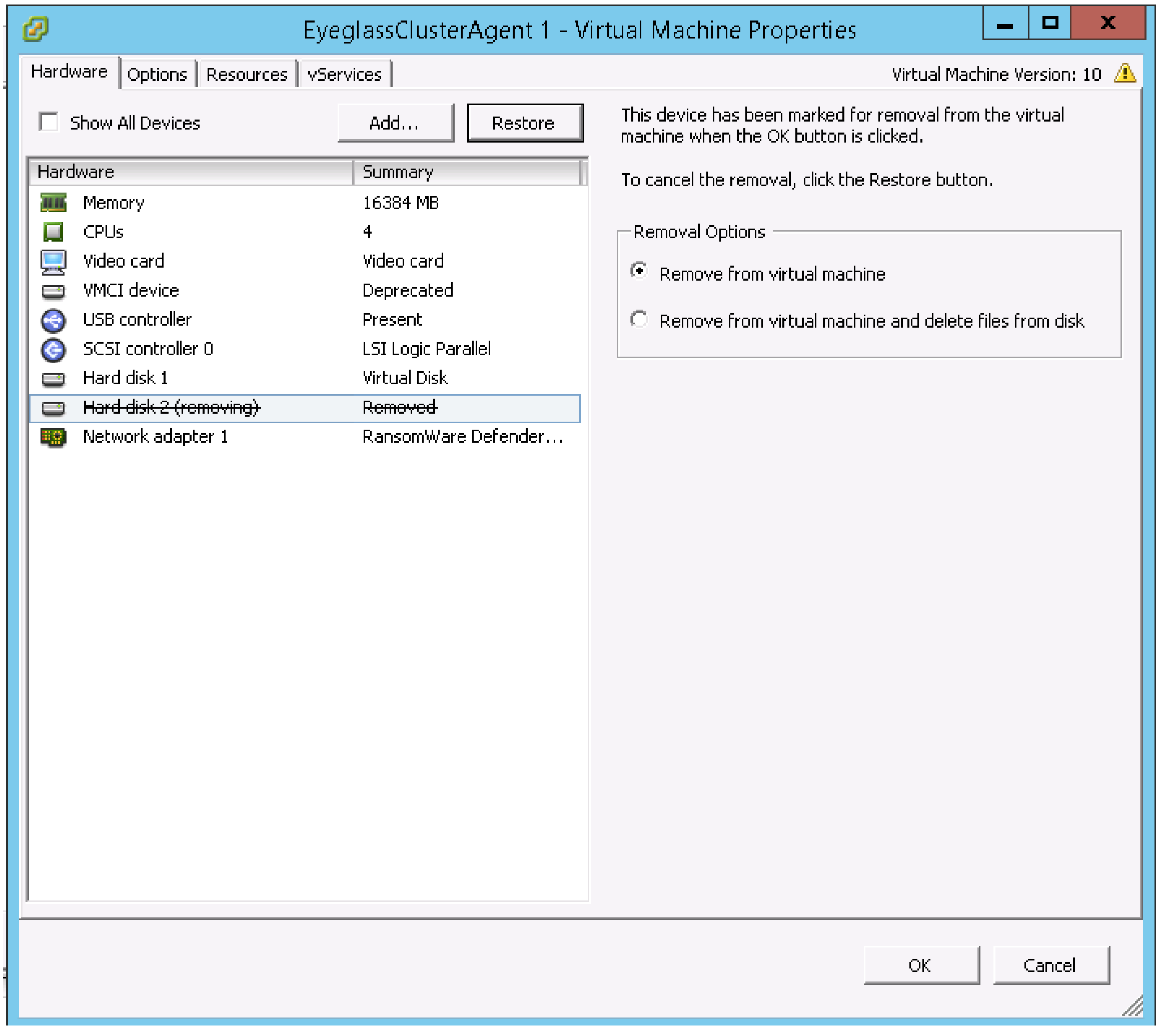

EyeglassClusterAgent 1 - Virtual Machine Properties

Hardware | Options | Resources | vServices — Virtual Machine Version: 10

□ Show All Devices — Add... — Restore

Hardware	Summary
Memory	16384 MB
CPUs	4
Video card	Video card
VMCI device	Deprecated
USB controller	Present
SCSI controller 0	LSI Logic Parallel
Hard disk 1	Virtual Disk
~~Hard disk 2 (removing)~~	~~Removed~~
Network adapter 1	RansomWare Defender...

This device has been marked for removal from the virtual machine when the OK button is clicked.

To cancel the removal, click the Restore button.

Removal Options

● Remove from virtual machine

○ Remove from virtual machine and delete files from disk

OK — Cancel

5. IMPORTANT STEP: Record the location of the VMDK disk on the datastore and record this path and for use in a later step. **NOTE: Record the full path in the data store to the disk for this VM, you will need this exact path to the vmdk disk to attach to the new appliance VM disk.**

   a. 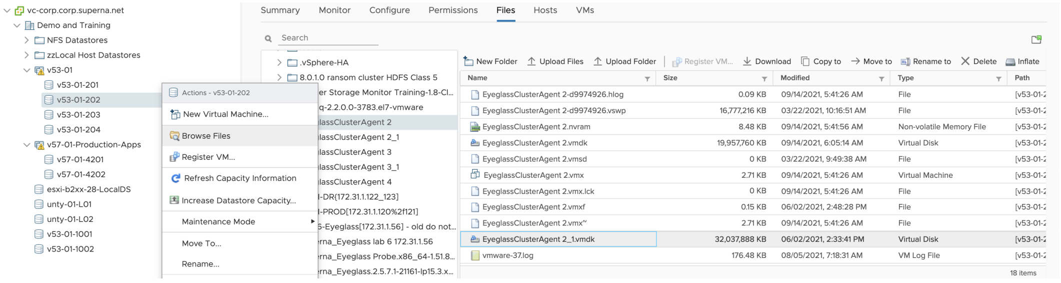

   - vc-corp.corp.superna.net
     - Demo and Training
       - NFS Datastores
       - zzLocal Host Datastores
       - v53-01
         - v53-01-201
         - v53-01-202
         - v53-01-203
         - v53-01-204
       - v57-01-Production-Apps
         - v57-01-4201
         - v57-01-4202
       - esxi-b2xx-28-LocalDS
       - unty-01-L01
       - unty-01-L02
       - v53-01-1001
       - v53-01-1002

   .vSphere-HA
   8.0.1.0 ransom cluster HDFS Class 5
   er Storage Monitor Training-1.8-Cl...
   q-2.2.0.0-3783.el7-vmware
   lassClusterAgent 2
   lassClusterAgent 2_1
   lassClusterAgent 3
   lassClusterAgent 3_1
   lassClusterAgent 4
   l-DR{172.31.1.122_123]
   l-PROD[172.31.1.120%2f121]
   6-Eyeglass[172.31.1.56] - old do not...
   erna_Eyeglass lab 6 172.31.1.56
   erna_Eyeglass Probe.x86_64-1.51.8...
   erna_Eyeglass.2.5.7.1-21161-lp15.3.x...

Name	Size	Modified	Type	Path
EyeglassClusterAgent 2-d9974926.hlog	0.09 KB	09/14/2021, 5:41:26 AM	File	[v53-01-2
EyeglassClusterAgent 2-d9974926.vswp	16,777,216 KB	03/22/2021, 10:16:51 AM	File	[v53-01-2
EyeglassClusterAgent 2.nvram	8.48 KB	09/14/2021, 5:41:56 AM	Non-volatile Memory File	[v53-01-2
EyeglassClusterAgent 2.vmdk	19,957,760 KB	09/14/2021, 6:05:14 AM	Virtual Disk	[v53-01-2
EyeglassClusterAgent 2.vmsd	0 KB	03/22/2021, 9:49:38 AM	File	[v53-01-2
EyeglassClusterAgent 2.vmx	2.71 KB	09/14/2021, 5:41:26 AM	Virtual Machine	[v53-01-2
EyeglassClusterAgent 2.vmx.lck	0 KB	09/14/2021, 5:41:26 AM	File	[v53-01-2
EyeglassClusterAgent 2.vmxf	0.15 KB	06/02/2021, 2:48:28 PM	File	[v53-01-2
EyeglassClusterAgent 2.vmx~	2.71 KB	09/14/2021, 5:41:26 AM	File	[v53-01-2
EyeglassClusterAgent 2_1.vmdk	32,037,888 KB	06/02/2021, 2:33:41 PM	Virtual Disk	[v53-01-2
vmware-37.log	176.48 KB	08/05/2021, 7:18:31 AM	VM Log File	[v53-01-2

18 items

6. Re-add that 2nd disk from old ECA VM 1 to the correspondent new ECA VM 1. NOTE: You must use the VMDK using the location in the

datastore from the step above, where the path to this 2nd disk was recorded. Click Add

7.

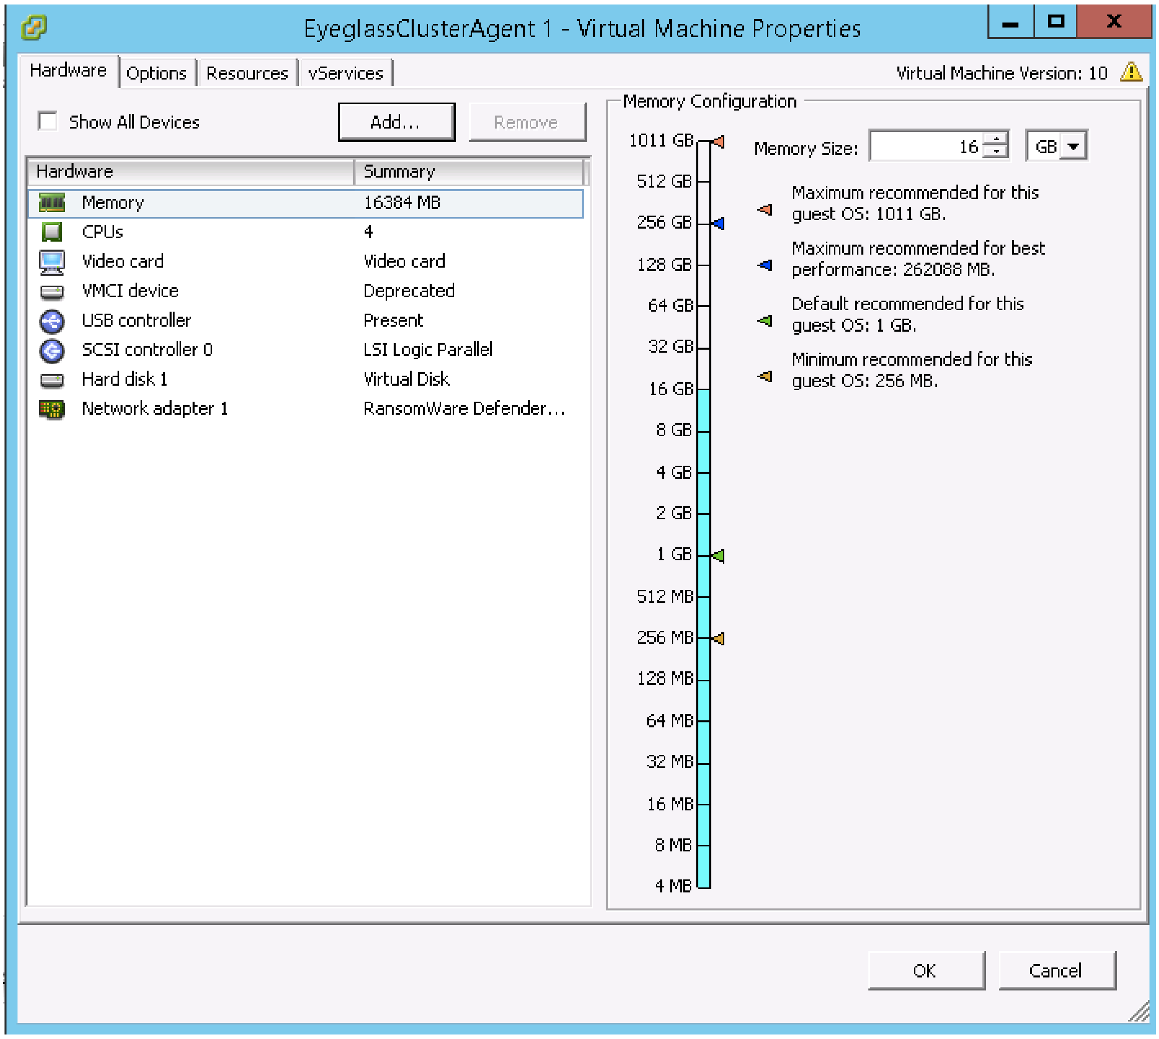

EyeglassClusterAgent 1 - Virtual Machine Properties

Hardware | Options | Resources | vServices

Virtual Machine Version: 10

□ Show All Devices

Hardware	Summary
Memory	16384 MB
CPUs	4
Video card	Video card
VMCI device	Deprecated
USB controller	Present
SCSI controller 0	LSI Logic Parallel
Hard disk 1	Virtual Disk
Network adapter 1	RansomWare Defender…

Memory Configuration

Memory Size: 16 GB

Maximum recommended for this guest OS: 1011 GB.

Maximum recommended for best performance: 262088 MB.

Default recommended for this guest OS: 1 GB.

Minimum recommended for this guest OS: 256 MB.

8.

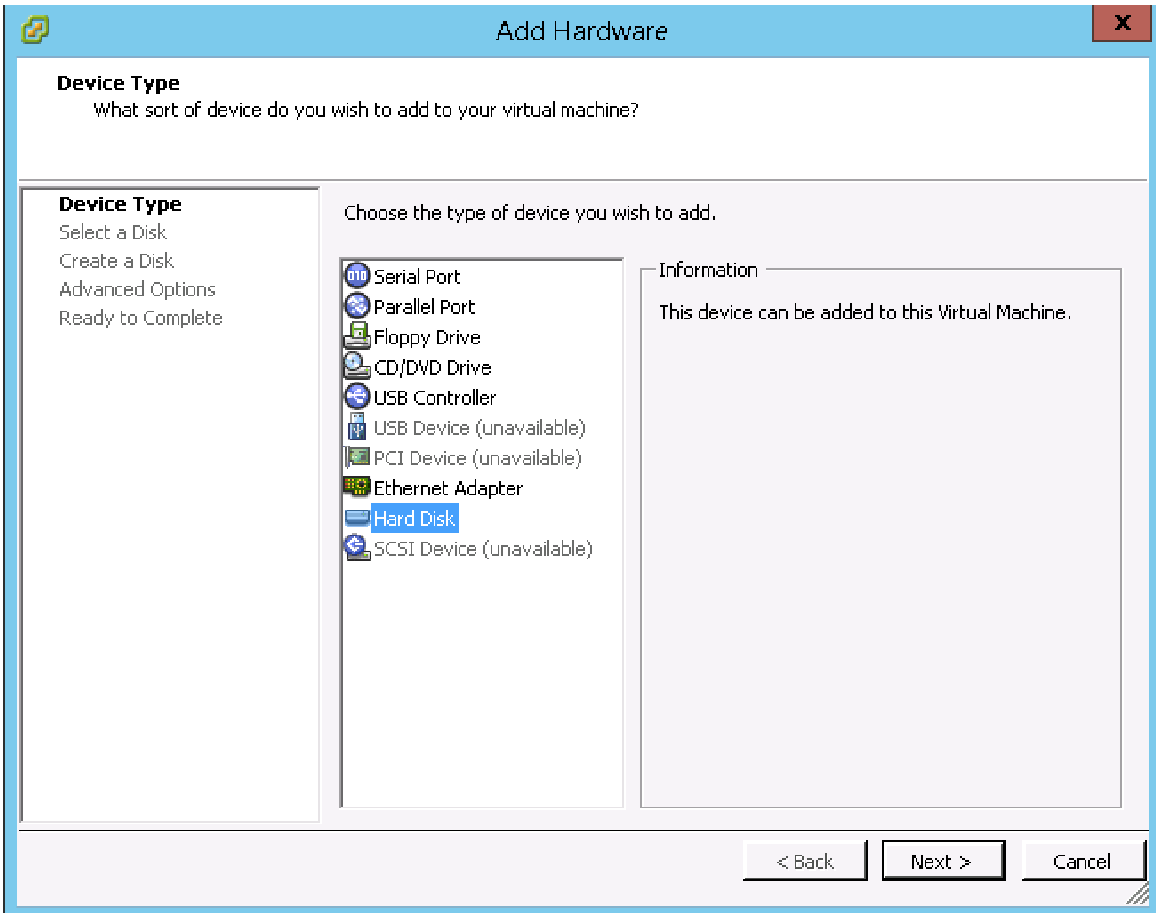

Add Hardware

Device Type

What sort of device do you wish to add to your virtual machine?

Choose the type of device you wish to add.

Information

This device can be added to this Virtual Machine.

9. Select “Use an existing virtual disk”

10.

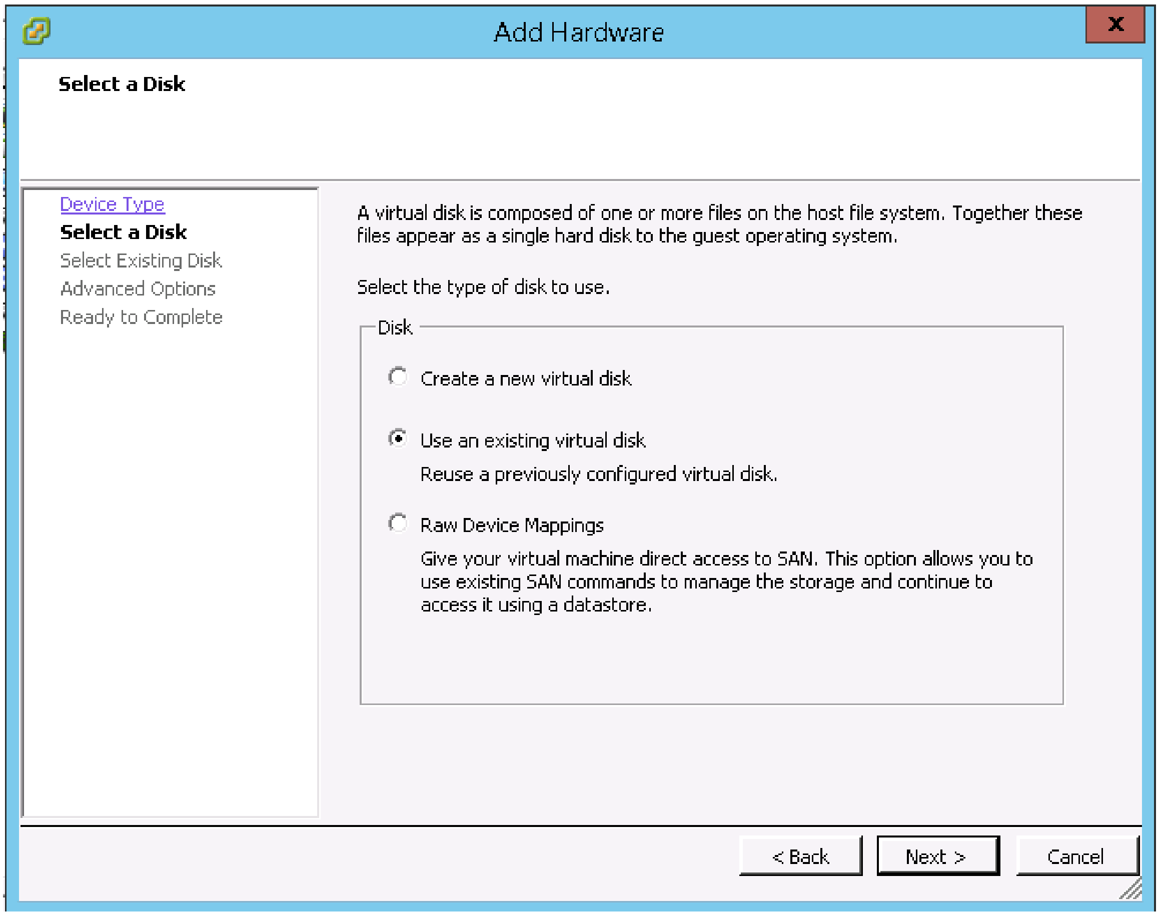

Add Hardware

Select a Disk

A virtual disk is composed of one or more files on the host file system. Together these files appear as a single hard disk to the guest operating system.

Select the type of disk to use.

Disk

○ Create a new virtual disk

● Use an existing virtual disk

Reuse a previously configured virtual disk.

○ Raw Device Mappings

Give your virtual machine direct access to SAN. This option allows you to use existing SAN commands to manage the storage and continue to access it using a datastore.

11. Specify the correct Disk File Path (Warning: Do not choose wrong disk) , **use the path to the datastore and folder and vmdk recorded from the step above.**

12.

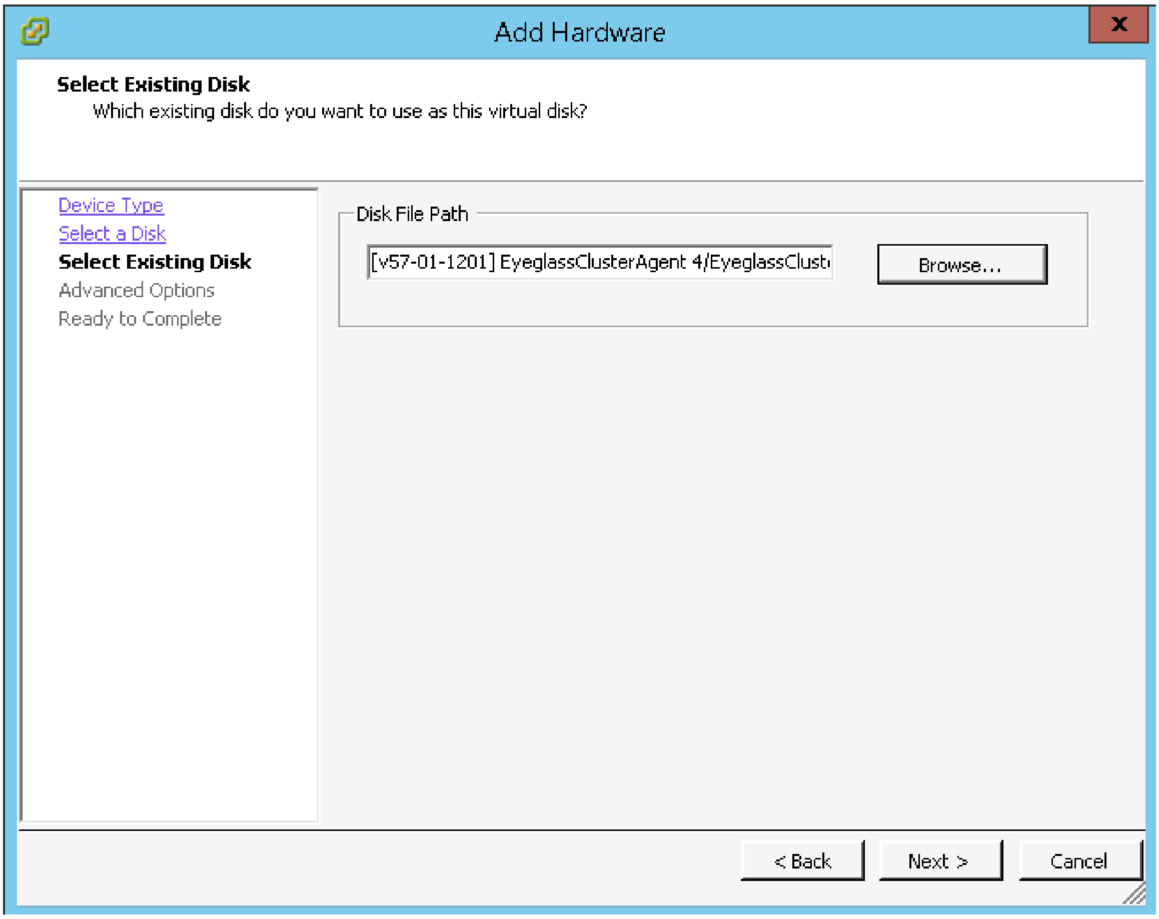

Add Hardware

Select Existing Disk

Which existing disk do you want to use as this virtual disk?

Disk File Path

[v57-01-1201] EyeglassClusterAgent 4/EyeglassClust

13. Accept the default advanced options and click “Next”

14.

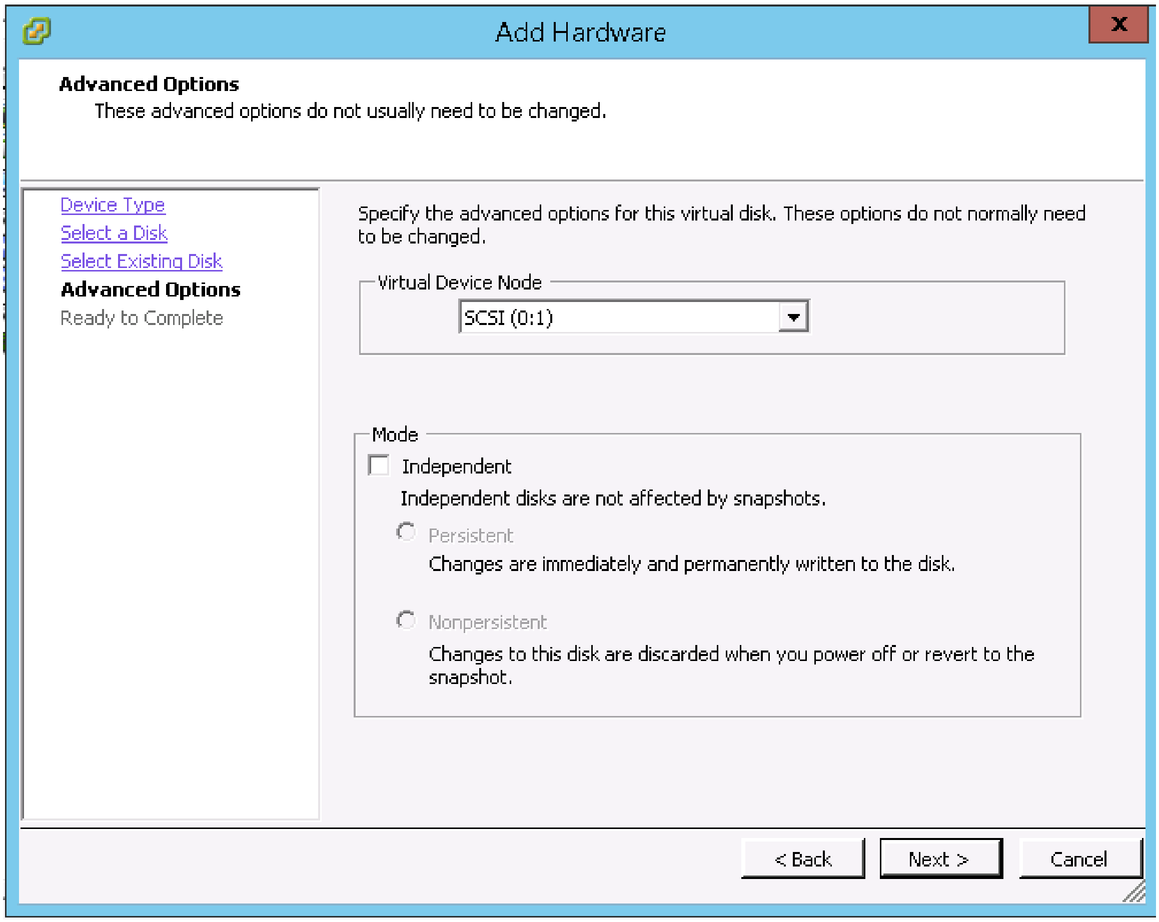

Add Hardware

Advanced Options

These advanced options do not usually need to be changed.

Specify the advanced options for this virtual disk. These options do not normally need to be changed.

Virtual Device Node

SCSI (0:1)

Mode

□ Independent

Independent disks are not affected by snapshots.

○ Persistent

Changes are immediately and permanently written to the disk.

○ Nonpersistent

Changes to this disk are discarded when you power off or revert to the snapshot.

15. Click “Finish”

16.

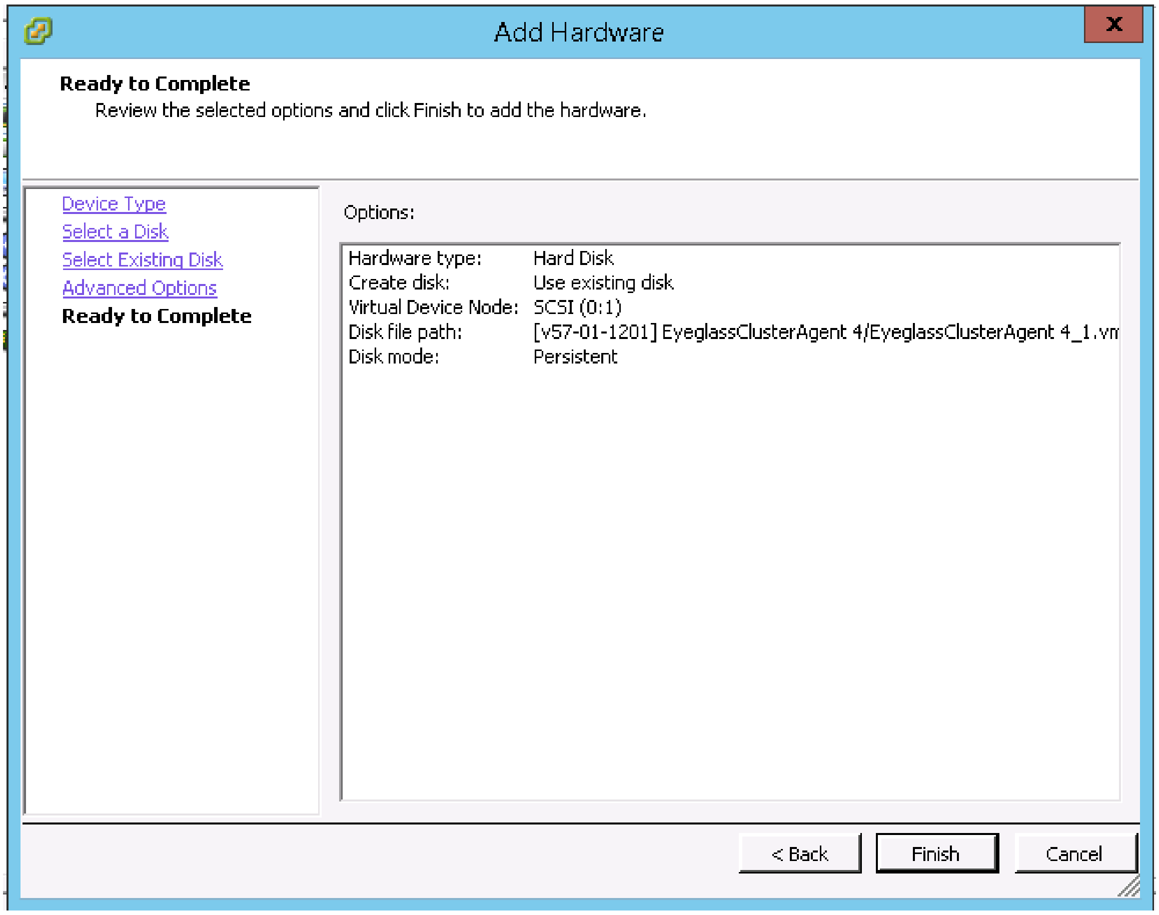

Ready to Complete

Review the selected options and click Finish to add the hardware.

Options:

Hardware type:	Hard Disk
Create disk:	Use existing disk
Virtual Device Node:	SCSI (0:1)
Disk file path:	[v57-01-1201] EyeglassClusterAgent 4/EyeglassClusterAgent 4_1.vm
Disk mode:	Persistent

17.

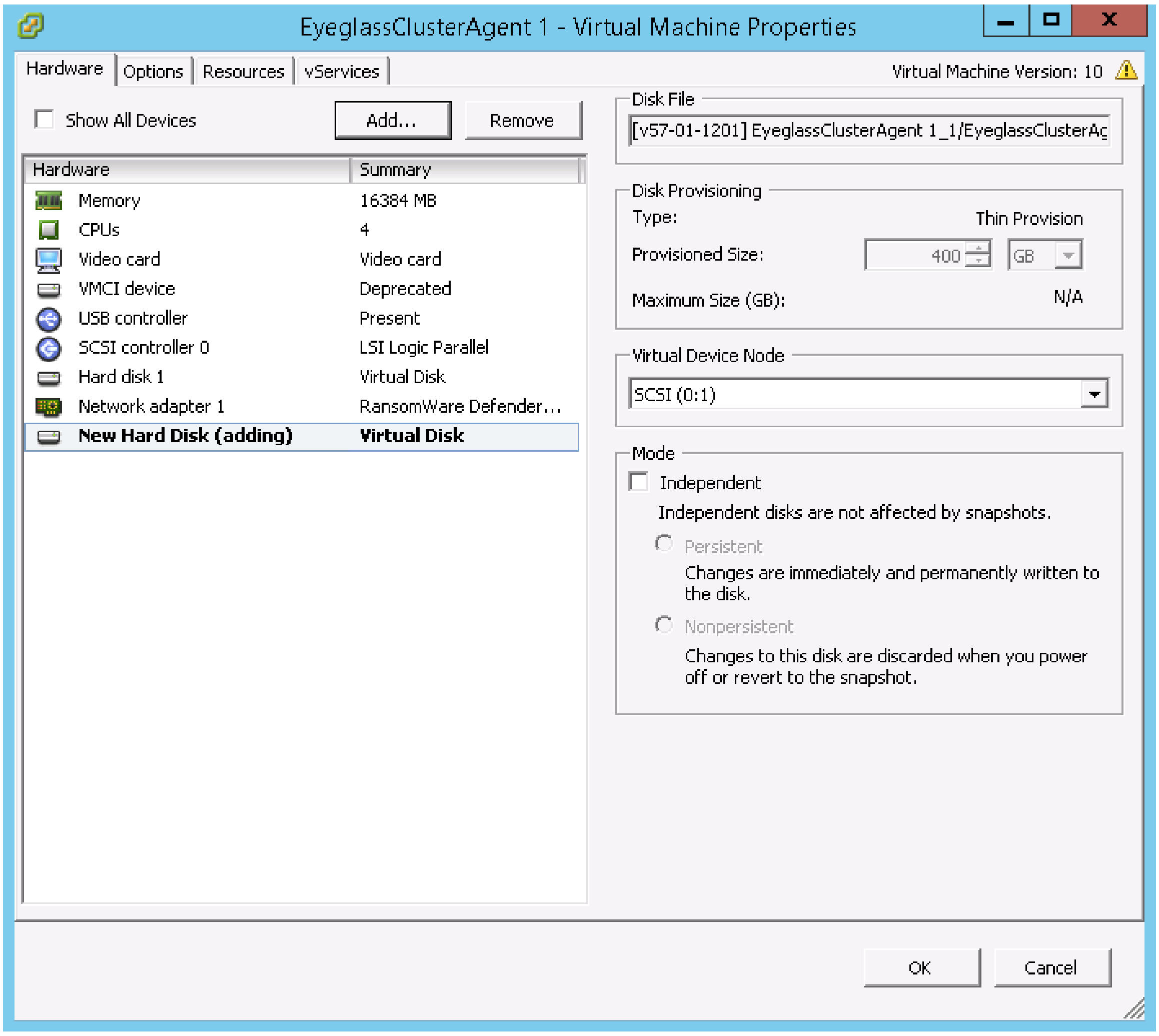

EyeglassClusterAgent 1 - Virtual Machine Properties

Virtual Machine Version: 10

□ Show All Devices

Hardware	Summary
Memory	16384 MB
CPUs	4
Video card	Video card
VMCI device	Deprecated
USB controller	Present
SCSI controller 0	LSI Logic Parallel
Hard disk 1	Virtual Disk
Network adapter 1	RansomWare Defender...
New Hard Disk (adding)	**Virtual Disk**

Disk File
[v57-01-1201] EyeglassClusterAgent 1_1/EyeglassClusterAg

Disk Provisioning
Type: Thin Provision
Provisioned Size: 400 GB
Maximum Size (GB): N/A

Virtual Device Node
SCSI (0:1)

Mode
□ Independent
Independent disks are not affected by snapshots.
○ Persistent
Changes are immediately and permanently written to the disk.
○ Nonpersistent
Changes to this disk are discarded when you power off or revert to the snapshot.

18. IMPORTANT STEP: Repeat the 2nd disk VMDK migration from the old appliance to the new appliance VM's for all remaining ECA VM's

19. Mandatory Step: Take a vmware level snapshot of all search nodes before proceeding to the next steps. This is the only way to roll back if any issues block the upgrade.

20. Done

Power on New appliance

1. SSH to ECA node 1 as ecaadmin

2. Ping each ip address in the cluster until each VM responds. **NOTE: Do not continue if you cannot ping each VM in the cluster.**
3. From ECA node 1: **ecactl cluster up**
4. Verify that new ECA can be brought up successfully
5. Verify Search license: searchctl licenses list
6. Verify registered PowerScale cluster: searchctl PowerScales list
7. Verify configured folder: searchctl folders list
8. Verify from Eyeglass Search UI https://<eca-node1-ip> that able to login and search the existing data.
9. Add new data and once the next incremental ingestion and commit has been completed, verify from Eyeglass Search UI.
10. Done

1.

5. Eyeglass Golden Copy Installation and Upgrade Guide

Home Top

Overview

Golden Copy VM contains all management and GUI functions and can copy data directly from a single VM. Additional Virtual accelerator nodes (VAN's) can be deployed to scale out the performance of the copy jobs.

VMware Requirements

1. vCenter 6.x, 6.5 and 7.0.1 Build: 17491160

VM Specifications for 3 Scaling Configurations

1. Small Configuration **Lab testing** - 1 x VM with 4x vCPU , 16G of ram, 400G hard disk
2. Small Configuration **Production Use** with default VM resources 1 x VM with 4x vCPU , 16G of ram, 400G hard disk
    a. Limit of 4 folder definitions
    b. > 4 folder definitions requires additional disk space to store file copy history for each folder. Additional 110 GB for 10 folders added
    c. NOTE: Multi VM deployments provide additional disk space on the VM cluster for storing file copy history
    d. disk latency read and write latency < 20 ms (test with command iotstat -xyz)
3. **Vertical Scaling high Performance Archiving** - 1 x VM with 12x vCPU, 32 G of ram, 600 G hard disk

a. Before power on, modify RAM and CPU to match above settings

b. > 4 folder definitions requires additional disk space to store file copy history for each folder. Additional 110 GB for 10 folders added

c. disk latency read and write latency < 20 ms (test with command iotstat -xyz)

d. modify the following file to expand the parallel file copies per VM

   i. nano /opt/superna/eca/eca-env-common.conf

   ii. Add a line

      1. export ARCHIVE_PARALLEL_THREAD_COUNT=400

   iii. control+x to save and exit

   iv. Change memory configuration (**note the the spacing must be Exactly as shown below**)

   control+x to save and exit

   v. nano /opt/superna/eca/docker-compose.overrides

   vi. version: '2.4'

```
services:
  indexworker:
    mem_limit: 8GB
    mem_reservation: 8GB
    memswap_limit: 8GB
```

```
archiveworker:
  mem_limit: 8GB
  mem_reservation: 8GB
  memswap_limit: 8GB

kafka:
  mem_limit: 4GB
  mem_reservation: 4GB
  memswap_limit: 4GB
```

4. Scale out Performance high performance and concurrent copy jobs - 6 x VM with 4x vCPU , 16G of ram, 400G hard disk
    a. > 4 folder definitions requires additional disk space to store file copy history for each folder. Additional 110 GB for 10 folders added
    b. disk latency read and write latency < 20 ms (test with command iotstat -xyz)

Cloud Storage Network Requirements

1. Direct NAT (private ip to public IP) network
2. Proxy configuration not currently supported

Firewall Rules and Direction Table

NOTE: These rules apply to traffic into the VM and out of the VM. All ports must be open between VM's, private VLAN's and firewall between VM's is not supported.

Port	Direction	Function
Operating System Open Suse 15.1		It is customer responsibility to patch the operating system and allow Internet repository access for automatic patching. The OS is not covered by the support agreement.
22	Admin PC → Golden Copy VM	Management access to the CLI
8080 (HTTPS) and 22 SSH	Golden Copy VM → PowerScale	REST API Access and SSH
NFS UDP/TCP port 111, TCP and UDP port 2049 UDP 300	Golden Copy VM → PowerScale Virtual Accelerator nodes → PowerScale	NFS mount in System Zone
port 9020 9021 for Dell ECS	Golden Copy VM and VAN VM's -> S3 (https 9021) (http 9020)	S3 protocol (https 9021, http 9020)
AWS	Golden Copy VM and VAN VM's -> S3 (https 443)	S3 protocol (https 443)
Azure Blob	Golden Copy VM and VAN VM's -> Azure Blob rest api https 443	Azure blob storage rest api

Firewall Diagram

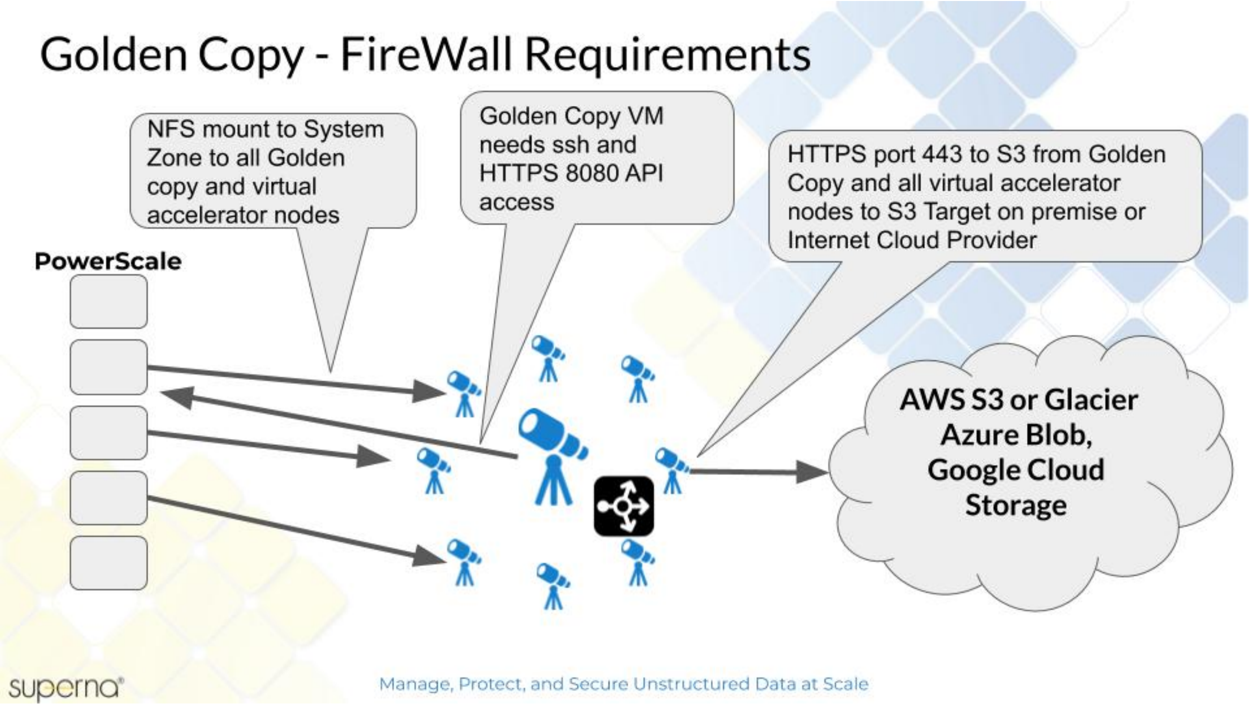

Golden Copy - FireWall Requirements

NFS mount to System Zone to all Golden copy and virtual accelerator nodes
Golden Copy VM needs ssh and HTTPS 8080 API access
HTTPS port 443 to S3 from Golden Copy and all virtual accelerator nodes to S3 Target on premise or Internet Cloud Provider
PowerScale
AWS S3 or Glacier Azure Blob, Google Cloud Storage
superna
Manage, Protect, and Secure Unstructured Data at Scale

Isilon/Power Scale Cluster NFS Mount Preparation Steps (Mandatory)

1. An IP pool created in the System access zone that with at least 3 Nodes as members. Must have DNS smartconnect name assigned to a management IP pool in the System zone for the NFS export used to read content from the snapshots folder and a 2nd NFS export for data recall.
2. Get each Cluster GUID and name that will be indexed. Record these values for steps below.
    a. Login to the cluster OneFS GUI and open the Cluster Management --> General settings menu and record the cluster GUID and cluster name. Example below.

3.

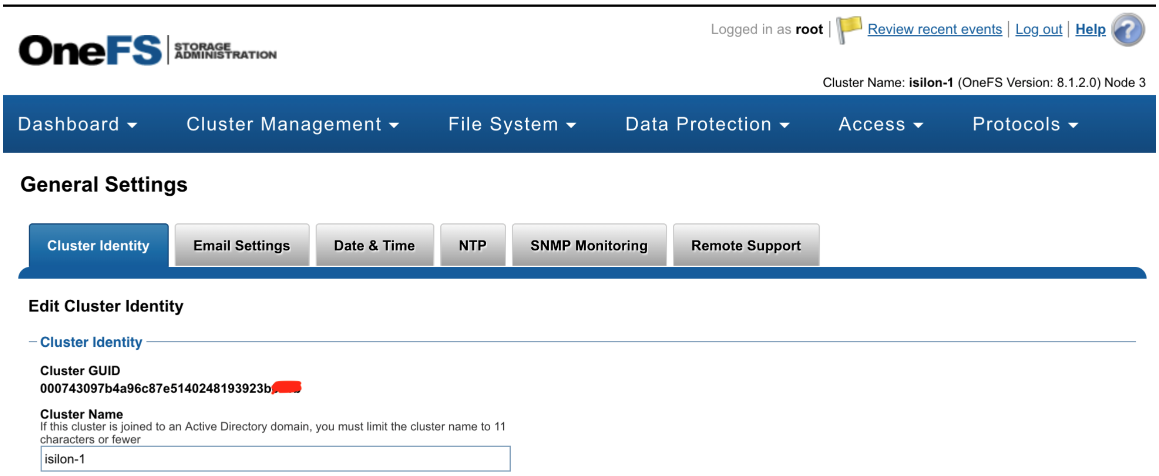

Cluster Name: **isilon-1** (OneFS Version: 8.1.2.0) Node 3

General Settings

Edit Cluster Identity

Cluster Identity

Cluster GUID
000743097b4a96c87e5140248193923b[redacted]

Cluster Name
If this cluster is joined to an Active Directory domain, you must limit the cluster name to 11 characters or fewer

isilon-1

4. Repeat for each cluster that will be licensed and used as a source cluster to copy data.

5. Create an NFS export in the System Access Zone for full content on all clusters that will be used as a source for archiving data. See example below where the IP addresses entered are the Golden Copy VM's. The export is created on the /ifs/.snapshot directory with root client list and clients list. Add Golden Copy and all Virtual Accelerator Node IP addresses.

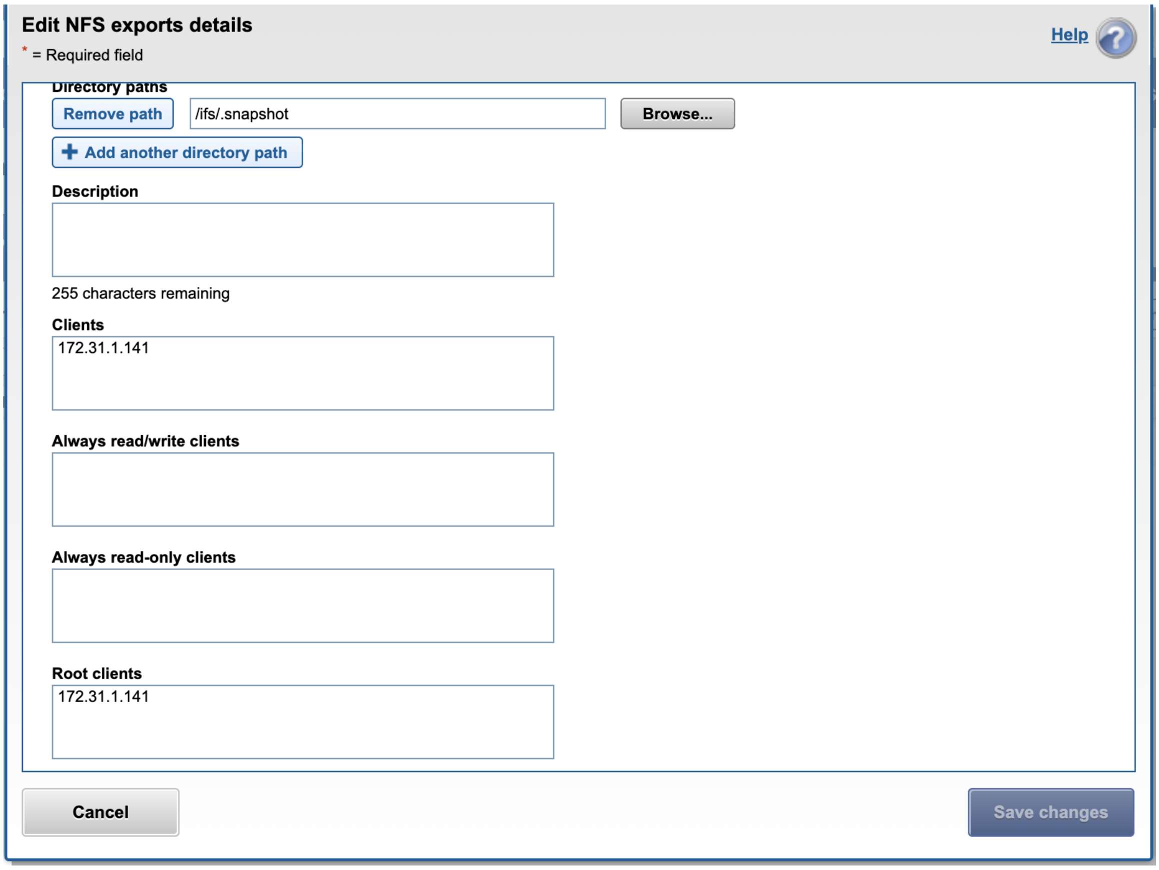

6. Edit NFS exports details

* = Required field

Directory paths: /ifs/.snapshot

Description:

255 characters remaining

Clients: 172.31.1.141

Always read/write clients:

Always read-only clients:

Root clients: 172.31.1.141

7. Create the recall NFS folder /ifs/goldencopy/recall using the cluster root user over ssh. Then create the export.

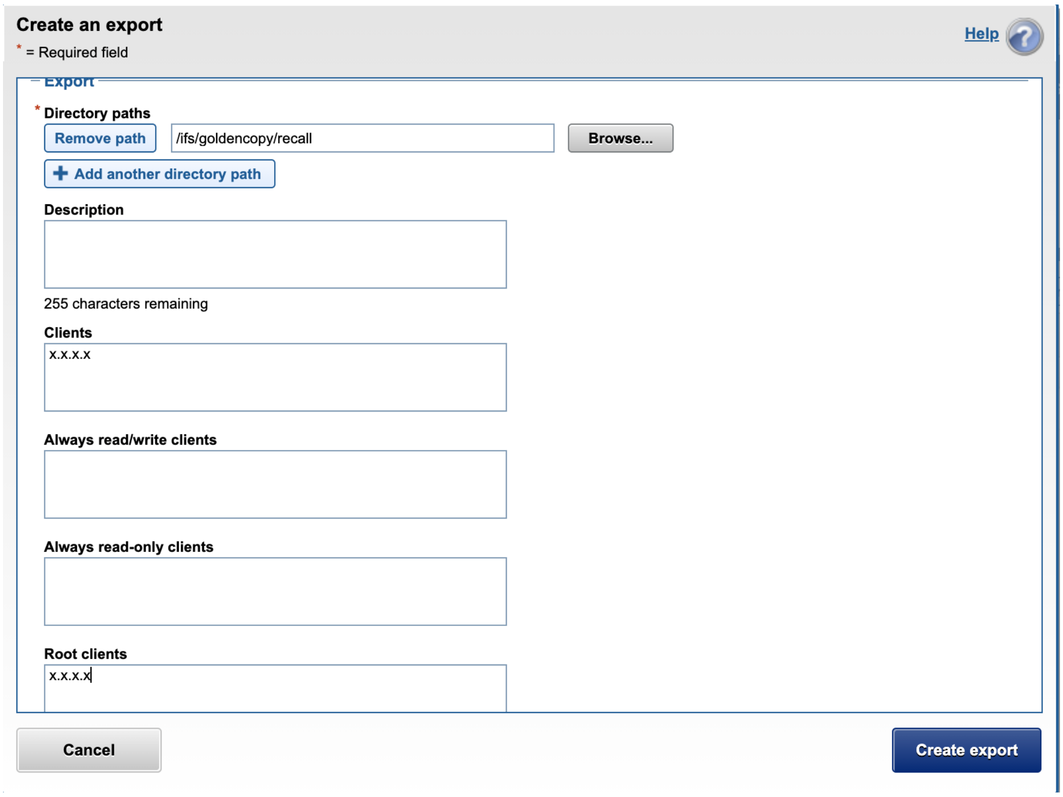

a. Create an export

* = Required field

* Directory paths: /ifs/goldencopy/recall

Description:

255 characters remaining

Clients: x.x.x.x

Always read/write clients:

Always read-only clients:

Root clients: x.x.x.x

8. Done

Eyeglass Golden Copy Service Account Preparation Steps (Mandatory)

The Golden Copy appliance is based on the ECA cluster architecture and has similar steps to complete installation.

Before you begin, collect the information below and verify all prerequisites are completed:

1. **Permissions for Service Account:** PowerScale REST API access with file traverse permissions to ingest files and directories by the Golden Copy VM. See the minimum permissions guide for a full list of permissions required for the eyeglassSR service account used by all ECA cluster products. Guide here.

Golden Copy OVA Deployment and Cluster VM Configuration (Mandatory)

1. Download the OVA following instructions here
2. **Deploy with vCenter**
   NOTE: vCenter 6.5 and 6.7 use FLASH or FLEX interface. HTML5 interface is not supported.
   a. Select the OVA file.

b. Set the node ip addresses, gateway IP, DNS, NTP IP

c.

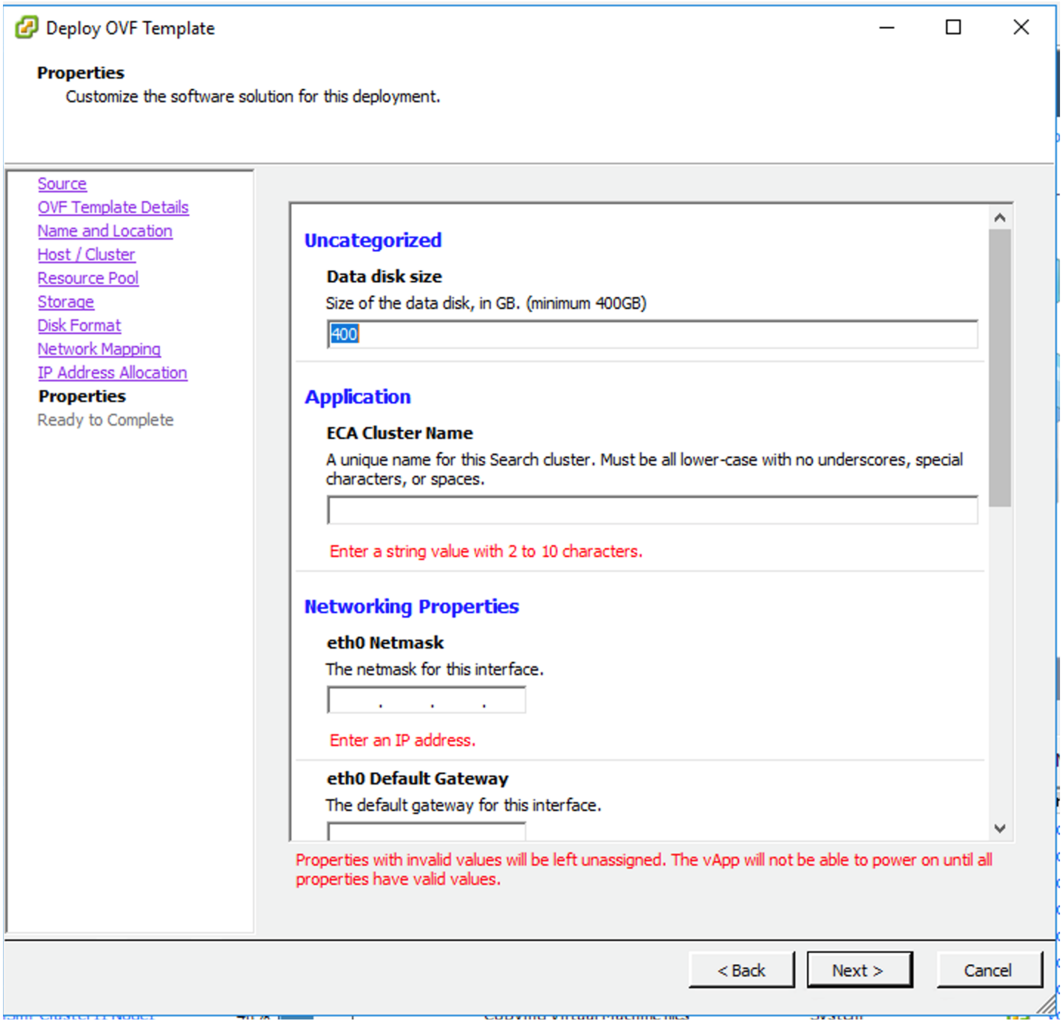

Deploy OVF Template

Properties

Customize the software solution for this deployment.

Uncategorized

Data disk size

Size of the data disk, in GB. (minimum 400GB)

400

Application

ECA Cluster Name

A unique name for this Search cluster. Must be all lower-case with no underscores, special characters, or spaces.

Enter a string value with 2 to 10 characters.

Networking Properties

eth0 Netmask

The netmask for this interface.

Enter an IP address.

eth0 Default Gateway

The default gateway for this interface.

Properties with invalid values will be left unassigned. The vApp will not be able to power on until all properties have valid values.

d. Set the **ECA cluster name** **(no special characters should be used, all lower case)**

e. NOTE: If using vertical scaling configuration edit the VM configuration with 12 vCPU and 32G of ram before power on.

f. Power on the OVA

g. SSH to the node 1 IP address

h. Login with user **ecaadmin** and default password **3y3gl4ss**

3. **Start up the cluster**

4. **ecactl cluster up**

5. **Get the appliance id and make a record of it - will be required to retrieve the license.**
    a. **ecactl version**
6. **Deployment done**
7. **NEXT Steps Golden Copy Cluster Logical Configuration**
8. Configuration steps to add licenses, add clusters, add archive folders is covered in the **Quick Start Steps** of the Golden Copy Admin guide.

How to Deploy Virtual Accelerator Nodes (VAN's) (Optional)

This node type is optional and allows distributed scale out copy performance. The Golden Copy VM can copy files without VAN VM's deployed.

NOTE: VAN deployment requires 6 VM's

1. Download the OVA following instructions here
2. **Deploy with vCenter**
   NOTE: vCenter 6.5 and 6.7 use FLASH or FLEX interface. HTML5 interface is not supported.
    a. Select the OVA file.
    b. Set the node ip addresses, gateway IP, DNS, NTP IP
3. Set the ECA cluster name (no special characters should be used, all lower case)
4. Repeat 6 times to deploy all 6 VM's

5. Power on the VM's
6. SSH to the Golden Copy VM node 1 (first ip address VM deployed)
7. Login with user ecaadmin and default password 3y3gl4ss
8. **Add each VM ip from node 1 using the command below:**
    a. **ecactl cluster add-node <ip_of_new_node>** (note all 6 VM's must be booted and pingable)
9. Upgrade each VM to the same release
    a. Download the upgrade file to each VAN vm and run the installer after making it executable with chmod 777 /home/ecaadmin/upgradefilename.run
    b. Run the upgrade
        i. ./home/ecaadmin/upgradefilename.run
    c. complete the upgrade on all VM's
10. ecactl cluster up (from node 1)
11. Verify boot process executes on all nodes in the cluster
12. This will now allow copy jobs to use additional VAN's to copy files.
13. Manage configuration from node 1 only.

Golden Copy and VAN VM node NFS Mount Configuration (Mandatory)

1. Golden Copy uses PowerScale snapshots to copy content. Follow the steps below to add the NFS export to each of the VM's that was created in the steps above. 2 NFS mounts are required, 1 for copying data and one for recalling data.
2. **You will need to complete this steps on all nodes**
    a. Cluster GUID and cluster name for each licensed cluster
    b. Cluster name as shown on top right corner after login to OneFS GUI
3. **Change to Root user**
    a. ssh to each VM as ecaadmin over ssh
    b. sudo -s
    c. enter ecaadmin password 3y3gl4ss
4. **Create local mount directory (repeat for each Isilon cluster)**
    a. mkdir -p /opt/superna/mnt/search/GUID/clusternamehere/ (replace GUID and clusternamehere with correct values)
    b. mkdir -p /opt/superna/mnt/recall/GUID/clusternamehere/
    c. **(Only if you have Virtual accelerator nodes, otherwise skip)** Use this command to run against all Golden Copy nodes, you will be prompted for ecaadmin password on each node.
        i. **NOTE: Must run from the Golden Copy VM and all VAN VM's must be added to the eca-env-common.conf file.**
        ii. **NOTE: example only.**

iii. ecactl cluster exec "sudo mkdir -p /opt/superna/mnt/search/00505699937a5e1f5b5d8b2342c2c3fe9fd7/clustername"

iv. ecactl cluster exec "sudo mkdir -p /opt/superna/mnt/recall/00505699937a5e1f5b5d8b2342c2c3fe9fd7/clustername"

5. **Configure automatic NFS mount After reboot**

   a. **Prerequisites**

      i. This will add a mount for content indexing to FSTAB on all nodes

      ii. Build the mount command using cluster guid and cluster name replacing the yellow highlighted sections with correct values for your cluster. **NOTE: This is only an example**

      iii. You will need a smartconnect name to mount the snapshot folder on the cluster. The Smartconnect name should be a system zone IP pool

      iv. Replace smartconnect FQDN and <> with a DNS smartconnect name

      v. Replace <GUID> with cluster GUID

      vi. Replace <name> with the cluster name

   b. **On each VM in the Golden Copy cluster:**

      i. ssh to the node as ecaadmin

      ii. sudo -s

      iii. enter ecaadmin password

iv. echo '<CLUSTER_NFS_FQDN>:/ifs/.snapshot /opt/superna/mnt/search/<GUID>/<NAME> nfs defaults,nfsvers=3 0 0'| sudo tee -a /etc/fstab

v. echo '<CLUSTER_NFS_FQDN>:/ifs/goldencopy/recall /opt/superna/mnt/recall/<GUID>/<NAME> nfs defaults,nfsvers=3 0 0'| sudo tee -a /etc/fstab

vi. mount -a

vii. mount to verify the mount

viii. exit

ix. Login to next node via ssh

c. repeat steps on each VM

6. done

How to Configure Multi Golden Copy VM Parallel Copy (Mandatory)

1. **Vertically scaled VM or multi Golden Copy VM deployments**
   a. The default deployment limits concurrent copies to 1 folder with a full or incremental job running. This must be changed for multi VM deployments to allow multiple folders to execute concurrent jobs (full or incremental).
2. **Single VM Limitations:**
   a. Single VM deployment is only supported with single folder concurrent job execution.
3. **Steps to enable Job Concurrency**

a. Login to VM node 1 as eccadmin

b. nano /opt/superna/eca/eca-env-common.conf

c. Copy and paste the following settings shown below. This enables full or incremental on up to 30 folders defined within Golden Copy across all clusters added to Golden Copy.

    i. Consult product supported limits of jobs in the admin guide.

```
# for blocking parallel jobs of any kind, true (enabled) by default
export ARCHIVE_BLOCK_PARALLEL_JOBS=false
# number of parallel full archive jobs allowed, works if `ARCHIVE_BLOCK_PARALLEL_JOBS` is disabled
export ARCHIVE_FULL_PARALLEL_JOBS_ALLOWED=30
# number of parallel incremental archive jobs allowed, works if `ARCHIVE_BLOCK_PARALLEL_JOBS` is disabled
export ARCHIVE_INCREMENTAL_PARALLEL_JOBS_ALLOWED=30
# total number of parallel jobs allowed, defaults to 1 FULL and 1 INCREMENTAL
export ARCHIVE_TOTAL_JOBS_ALLOWED=60
```

How To Upgrade Golden Copy Cluster

Offline Cluster No Internet Method

1. Login to the support site https://support.superna.net and download the offline Golden Copy upgrade file
2. Using a tool like Winscp copy the upgrade file to node 1 using ecaadmin user to authenticate
3. shutdown the cluster
    a. Login to node 1 as ecaadmin over ssh
    b. run the command "ecactl cluster down"
    c. wait for the cluster to shutdown
    d. Modify the install file copied to the cluster node 1
        i. Assuming the file was copied to default location in /home/ecaadmin
        ii. cd /home/ecaadmin
        iii. chmod 777 **<name of install file here>**
    e. Run the installer
        i. **./name of install file here**
        ii. when prompted enter the ecaadmin password
        iii. wait for all nodes to be upgraded to the new version
    f. Start the cluster
        i. ecactl cluster up
        ii. wait until all nodes are started
    g. done.

6. Eyeglass Clustered Agent vAPP Install and Upgrade Guide (Ransomware Defender, Easy Auditor, Performance Auditor)

Home Top

Abstract:

This Guide provides a step by step procedure for installing the Superna Eyeglass clustered agent vAPP used by Ransomware Defender, Easy Auditor and Performance Auditor. NOTE: Only follow steps in each section that names the product in the section

What's New

1. Syslog forwarding of ECA logs to eyeglass
2. Uses FluentD container for local logging and forwarding

3. Cluster Startup now checks HDFS configuration before starting and provides user feedback on validations
4. 3, 6 or 9 or upto 99 ECA node control and upgrade
5. Delayed startup option for containers
6. Statistics per container cli command
7. Kafka manager UI
8. **New 2.5.5 Ransomware Defender does not require HDFS or a smartconnect name pool for HDFS but if Easy Auditor is also installed than HDFS pool is still requried.**

Definitions

1. **ECA** - Eyeglass Clustered Agent - the entire stack that runs in a separate VM outside of Eyeglass that processes audit data

Deployment and Topology Overview

Deployment Diagram (Ransomware Defender, Easy Auditor, Performance Auditor)

This diagram shows a three node ECA cluster

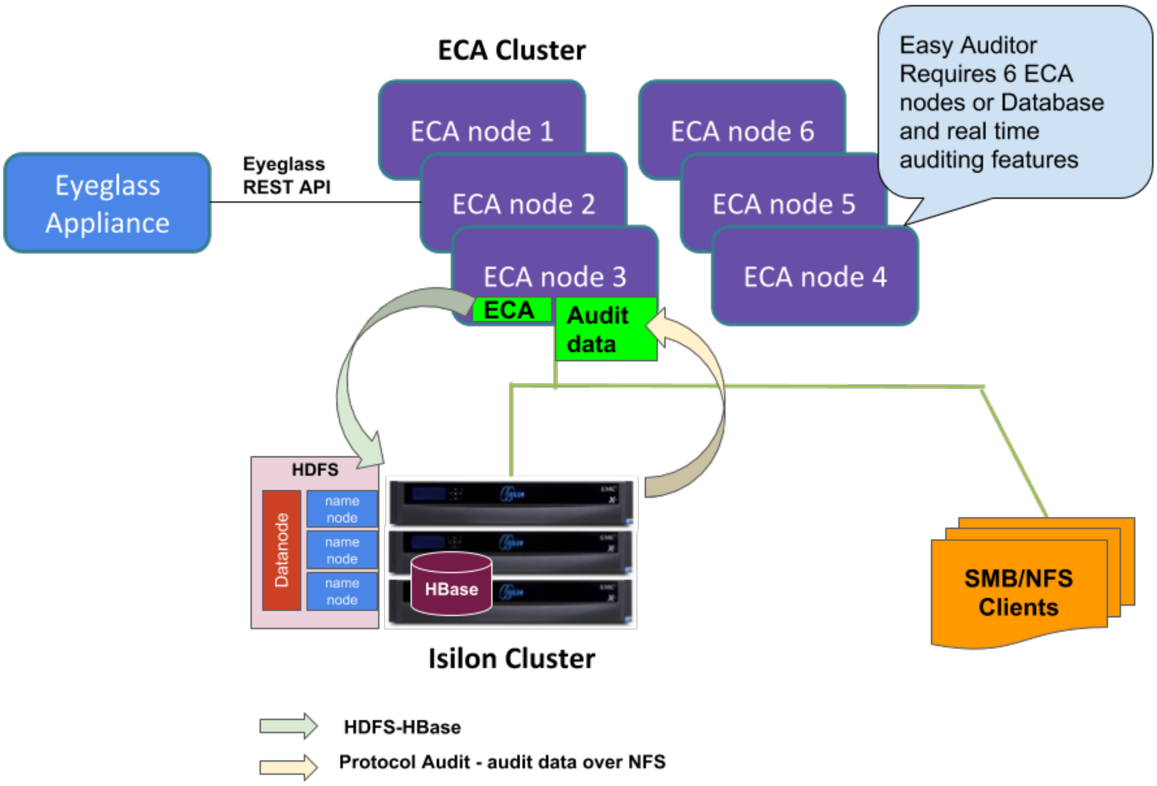

ECA Cluster Sizing and Performance Considerations

ECA Cluster Size by Application (Ransomware Defender, Easy Auditor, Performance Auditor)

ECA clusters are 3-12 nodes or greater depending the applications running on the cluster and the number of events per second or by the number of cluster nodes that generate audit events. The minimum ECA node configurations that are supported for all deployments are documented below. **NOTE: New applications or releases with features that require more resources will require ECA cluster to expand to handle multiple clusters or new application services.**

Application Configuration	Number of ECA	ESX hosts to split	ECA Node VM	Network Latency NFS	Easy Auditor Databa	Host Hardware

	VM nodes Required	VM Workload and Ensure High Availablity	Size	mount For Ransomware Defender & Easy Auditor	se Network Latency between ECA and Power Scale storing the DB	Configuration Requirements
Ransomware Defender only [3,6,8, 9]	3 ECA node cluster (1 to 2 managed clusters OR < 6000 audit events per second) 6 ECA node cluster (> 2 managed clusters OR > 6000 EVTS)	2[6]	4 x vCPU, 16G Ram, 30G OS partition + 80G disk	< 10 ms RTT	NA	2 socket CPU 2000 GHZ or greater, Disk IO latency average read and write < 20 ms
Easy Auditor	6 ECA	2 [6]	4 x	< 10 ms	< 5 ms	2 socket

Only [2, 3, 5,7,8, 9]	node cluster		vCPU, 16G Ram, 30G OS partition + 80G disk	RTT	RTT	CPU 2000 GHZ or greater, Disk IO latency average read and write < 20 ms
Ransomware Defender And Easy Auditor Unified deployment **(< 18K events per second)** [3,5,7,8, 9]	6 ECA node cluster	2 [6]	4 x vCPU, 16G Ram, 30G OS partition + 80G disk	< 10 ms RTT	< 5 ms RTT	2 socket CPU 2000 GHZ or greater, Disk IO latency average read and write < 20 ms
Very High IO rate clusters (> 18 K events per second) Ransomware Defender And Easy Auditor Unified deployment [3,5,7,8,9]	9 ECA node cluster	3[6]	4 x vCPU, 16G Ram, 30G OS partition + 80G disk	< 10 ms RTT	< 5 ms RTT	2 socket CPU 2000 GHZ or greater, Disk IO latency average read and write < 10 ms
Large node count clusters > 20 nodes. The more nodes to monitor audit data in real-time requires	**20 -30 nodes = 9 VM's** **> 30 nodes**	3 [6]	4 x vCPU, 16G Ram, 30G OS partition	< 10 ms RTT	< 5 ms RTT	2 socket CPU 2000 GHZ or greater, Disk IO latency average

more VM's to maintain throughput for a supported configuration. [3,5,7,8,9]	**= 12 VM's**		+ 80G disk			read and write < 10 ms
Unified Ransomware Defender, Easy Auditor and Performance Auditor Deployments [3,4,5,7,8,9]	6-9 ECA VM's depending on event rates	3 [6]	6 x vCPU, **20G Ram**, 30G OS partition + 80G disk	< 10 ms RTT	< 5 ms RTT	2 socket CPU 2000 GHZ or greater, Disk IO latency average read and write < 10 ms

[1] VMware OVA, Microsoft Hyper-v VHDX are available appliance platforms

[2] Contact support for reduced footprint configuration with 3 VM's only for low event rate environments.

NOTE: OVA default sets resource limit of 18000 MHZ for the OVA shared by all ECA VM nodes in the cluster. This limit can be increased if audit event load requires more CPU processing. Consult support before making any changes in vmware.

[3] NOTE: The ECA cluster is a real-time distributed processing and analysis platform which requires the ECA VM's to be on the same layer 2 subnet and low latency between the VM's. The VM's communicate between each other for many functions that requires low latency VM to VM communications. It is not supported to split the VM's in a single ECA cluster between data centers. The only supported distributed mode requires the Mini-ECA deployment architecture covered in this guide.

[4] NOTE: Unified Ransomware Defender, Easy Auditor and Performance Auditor requires additional resources in addition to the event rate sizing requirements. Each new application requires resources to maintain real time performance.   Additional RAM is required on ECA nodes 2-N.   The additional resources are 4 G RAM per ECA node and 2 additional vCPU per ECA node. If the event rate is very high resources may need to be increased beyond these initial settings. The Eyeglass VM also has increased RAM requirements for Performance Auditor. Consult the EyeGlass Scalability table for RAM upgrade requirements.

[5] NOTE: Audit Data retention > 1 year will increase the database size. As the database grows the number of VM's required to maintain the larger database will also need to increase. Any data retention > 1 year will require additional 3 ECA VM's added to maintain support for larger databases. Data retention set to greater than 365 days will require additional resources and a minimum of 3 additional ECA VM's to expand the cluster size.

[6] NOTE: Supported HA requires multiple physical Hosts to split VM's across hosts based on the size of the ECA cluster. ECA clusters with 3 VM's can tolerate N-1 VM failures, ECA clusters with 6 VM's can tolerate N-2 failures and ECA clusters > than 6 can tolerate N-3 VM failures.

[7] NOTE: All customers running Onefs 8.2 or later must disable directory open and directory close to reduce audit rate and reduce VM

footprint for a supported configuration. See instructions here to maintain support for your deployment.

8 NOTE: Storage vmotion or SDRS and DRS should be disabled since the ECA vm's are a real-time processing system.

9 NOTE: Archiving old gz files that collect on Onefs nodes is required to maintain performance of audit data ingestion. The cluster maintains old audit data in gz files once the active audit log reaches 1 GB. These files will collect for ever and NFS ingestion performance is impacted once the total GZ file count exceeds 5000 and will continue to degrade above this level.   It is recommended to follow the procedures here or with Onefs 9.x use the auto archive of audit data feature.

IP Connection and Pool Requirements for Analytics database

Requires HDFS on the Cluster (Easy Auditor)

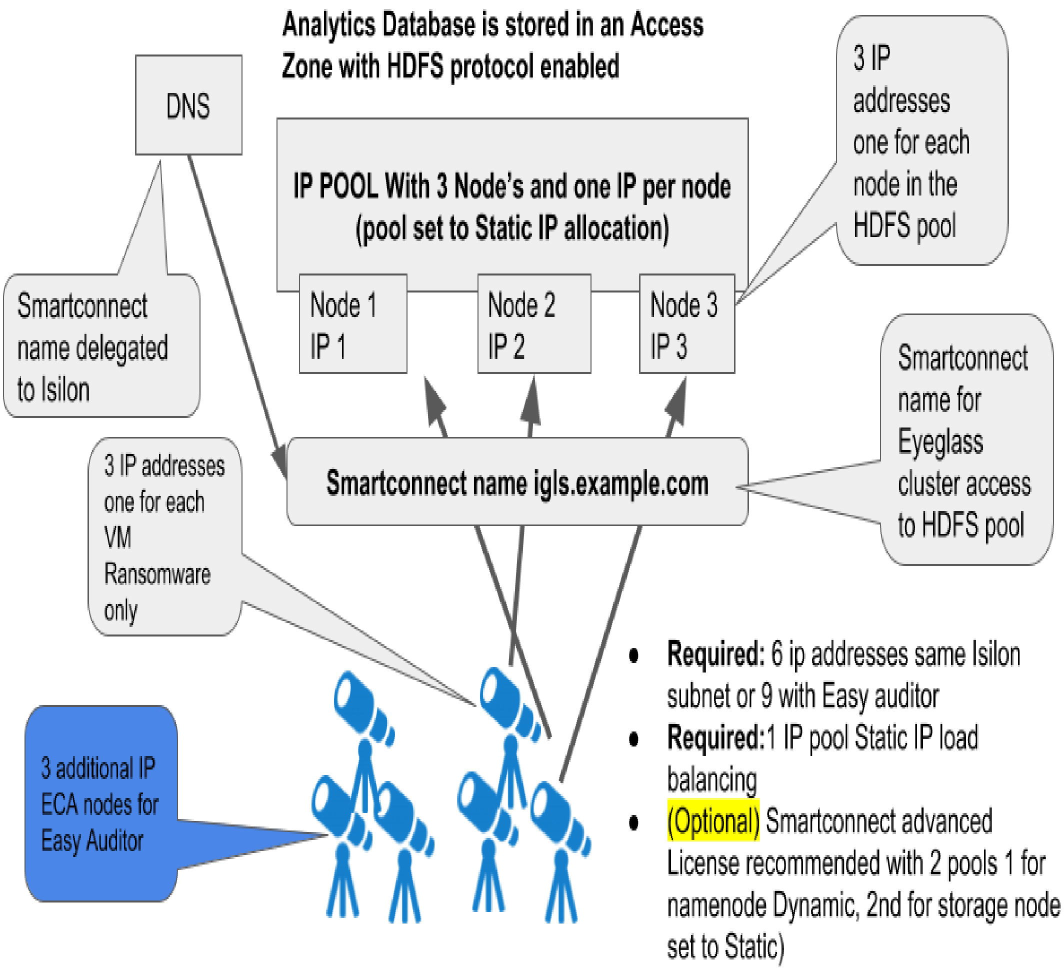
Analytics Database is stored in an Access Zone with HDFS protocol enabled
DNS
IP POOL With 3 Node's and one IP per node (pool set to Static IP allocation)
3 IP addresses one for each node in the HDFS pool
Smartconnect name delegated to Isilon
Node 1 IP 1
Node 2 IP 2
Node 3 IP 3
Smartconnect name for Eyeglass cluster access to HDFS pool
3 IP addresses one for each VM Ransomware only
Smartconnect name igls.example.com
3 additional IP ECA nodes for Easy Auditor

- **Required:** 6 ip addresses same Isilon subnet or 9 with Easy auditor
- **Required:**1 IP pool Static IP load balancing
- (Optional) Smartconnect advanced License recommended with 2 pools 1 for namenode Dynamic, 2nd for storage node set to Static)

ECA Cluster Network Bandwidth Requirements to PowerScale

(Ransomware Defender, Easy Auditor, Performance Auditor)

Each ECA node process audit events and writes data to the analytics database using HDFS on the same network interface. Therefore the combined TX and RX constitutes the
peak bandwidth requirement per node. The table below is is a minimum bandwidth requirements per ECA VM example calculation.
HDFS Bandwidth estimates and guidelines for Analytics database network bandwidth access to PowerScale.

Product Configuration	**Audit Event rate Per Second**	**Peak Bandwidth requirement**	
		Events per second per ECA cluster (input NFS Reading events from PowerScale to ECA cluster)	**Audit data Writes Mbps per ECA cluster (output HDFS writing events)**
Ransomware Defender only	2000 evts	Input to ECA → 50 Mbps	Out of ECA <-- < 150 Mbps
Unified Ransomware and Easy Auditor - **Steady state storing events**	> 4000 evts	Input to ECA → 125 Mbps	Out of ECA ← 500 Mbps - 1.2 Gbps
Easy Auditor Analysis Reports (long running reports)	NA	Input to ECA (HDFS from PowerScale) ← 800 Mbps - 1.5 Gbps while report runs	

VMware ESX Host Compute Sizing for ECA nodes

(Ransomware Defender, Easy Auditor, Performance Auditor)

Audit data is a real-time intensive processing task. Auditing workload increases with file IO, and the number of users is a good metric to estimate file IO workload per user. The table below is based on an assumption of 1.25 events per second per user with a peak of 1.5 events per second and can be used as a guideline to help determine how many events per second your environment will produce. This will help you to determine the sizing of the VM and placement on ESX hardware.

VMware or Hyper-V Host Requirements

1. **NOTE: Vmware environments with DRS and SDRS should excempt the ECA and vApp from dynamic relocation as a best practise. As a real-time application with time skew requirements between VM's for processing and database operations, DRS movement of running VM's . For maintenance purposes it is ok migrate vm's as needed.**

Number of active concurrent Users per cluster 1	**ECA VM per Physical Host Recommendation**	**Estimated Events Guideline**
1 to 1000	1 Host	=5000 * 1.25 = 6,250 events per second
5000 - 10000	2 Host	=10,000 * 1.25 = 12,500 events per second
> 10000	3 Host	= Number of users * 1.25 events/second

1 Active tcp connection with file IO to the cluster

ECA Cluster Deployment Topologies with PowerScale Clusters

Considerations & Requirements to Select a Deployment Option for the ECA Cluster:

Centralized ECA deployment is easier to manage and monitor with distributed clusters. This requires a central ECA cluster to mont audit folder using NFS at remote locations. Follow the requirement guideles below to determine which deployment model should be used.

- **When to use Mini ECA solution**: latency between the site where the ECA database for Easy Auditor will reside or the site where Ransomware Defender Analytics cluster will reside and the site with a cluster that will be managed by Easy Auditor or Ransomware Defender is > 8 ms ping time should use Mini ECA with the intall guide here.
    - if latency is < 10 ms ping times then use NFS mount to the remote cluster for Easy Auditor/Ransomware Defender monitoring

2.

Security - Firewall Port Requirements Ransomware Defender , Easy Auditor and Performance Auditor

Firewall Rules and Direction Table

NOTE: These rules apply to traffic into the VM and out of the VM. All ports must be open between VM's, private VLAN's and firewall between VM's is not supported.

To harden the ECA we recommend firewalling the ports between Eyeglass VM and the ECA vm. No external access to ECA is required other than SSH access for management access.   This is the #1 method to secure the ECA. We also recommend securing the trouble shooting GUI's (HBASE, Spark and Kafka) to a management subnet.

Eyeglass GUI VM (Applies to Ransomware Defender & Easy Auditor)		
Consult the EyeGlass firewall ports Required ports that must be in place for Addon products Ransomware and Easy Auditor are listed in the table for specific features.		
Ransomware Defender Only		
Port	**Direction**	**Function**
Operating System Open Suse 15.1		It is customer responsibility to patch the operating system and allow Internet repository access for automatic patching. The OS is not covered by the support agreement.
2181 (TCP)	Eyeglass → ECA	zookeeper
9092 (TCP)	Eyeglass → ECA	Kafka
5514 (TCP) as of	ECA → Eyeglass	syslog

2.5.6 build 84		
443 (TCP)	ECA → Eyeglass	TLS messaging
443 (HTTPS)	2.5.7 Eyeglass→ Internet	Downloading file extension list whitelist this url https://storage.googleapis.com/rwdefender.superna.net/supernaRansomwareFilters.json
NFS (UDP & TCP)	ECA → Power Scale	NFS export mounting audit data folder on managed clusters
NTP (UDP) 123	ECA → NTP server	time sync
Additional Ports for Easy Auditor		
8020 or 585 (TCP)	ECA → Power Scale	HDFS
16000, 16020	Eyeglass → ECA	hbase
6066 (TCP)	Eyeglass →	Spark job engine

	ECA	
9092 (TCP)	Eyeglass --> ECA	Kafka broker
4040 (TCP)	Admin browser → ECA	Running jobs monitor
8081 (TCP)	Admin browser → ECA	Spark Workers UI
8080 (TCP)	Admin browser → ECA	Spark Master UI
16010 (TCP)	Admin browser → ECA	HBase Master UI
16030 (TCP)	Admin browser → ECA	HBase Regionserver UI
18080 (TCP)	Eyeglass VM→ ECA	Spark History Report for monitoring running searches in Easy Auditor
9000 (TCP)	Admin browser → ECA	Kafka UI

Eyeglass VM Pre-requisites - Mandatory Step

Eyeglass License Requirements

1. Eyeglass must be deployed with or upgraded to the correct compatible release for the ECA release that is being installed.
2. Eyeglass Licences for Easy Auditor or Ransomware Defender must be added to Eyeglass VM.
   a. Login to Eyeglass
   b. Open Licence manager Icon
   c. Follow how to download license key instructions using the email license token provided with your purchase.
   d. Upload the license key zip file from Step #3
   e. Web page will refresh
   f. Open License manager
   g. Select Licensed devices tab
   h. Set the license status for each product to **user licensed** for clusters that should be monitored by Ransomware Defender or Easy Auditor (depending on what license keys you purchased).
   i. Set the license status for each product for each cluster that should not be licensed to **Unlicensed**. This is required to

ensure licences are applied to the correct cluster and blocked from being applied to the incorrect cluster.

j. 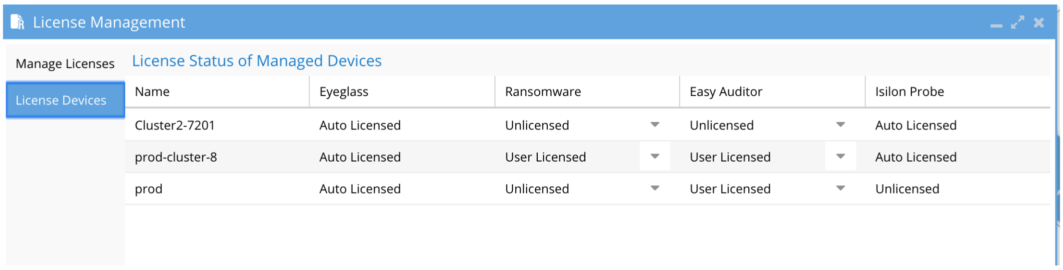

License Management

Manage Licenses
License Devices

License Status of Managed Devices

Name	Eyeglass	Ransomware	Easy Auditor	Isilon Probe
Cluster2-7201	Auto Licensed	Unlicensed	Unlicensed	Auto Licensed
prod-cluster-8	Auto Licensed	User Licensed	User Licensed	Auto Licensed
prod	Auto Licensed	Unlicensed	User Licensed	Unlicensed

Deployment Overview

The Eyeglass appliance is required to be installed and configured. The ECA Cluster runs in a separate group of VM's from Eyeglass.

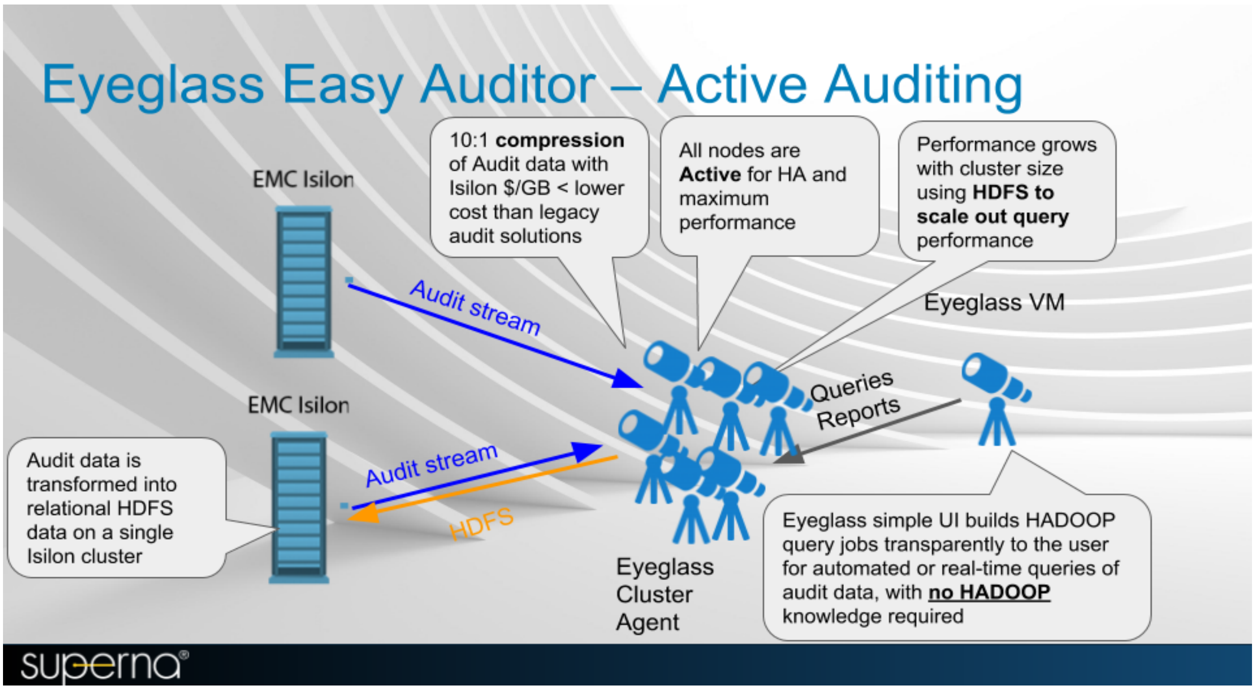

Eyeglass Easy Auditor – Active Auditing

10:1 **compression** of Audit data with Isilon $/GB < lower cost than legacy audit solutions

All nodes are **Active** for HA and maximum performance

Performance grows with cluster size using **HDFS to scale out query** performance

Audit data is transformed into relational HDFS data on a single Isilon cluster

Eyeglass simple UI builds HADOOP query jobs transparently to the user for automated or real-time queries of audit data, with **no HADOOP** knowledge required

Eyeglass will be responsible for taking action against the cluster and notifying administrators.

- PowerScale cluster stores analytics database (this can be the same cluster that is monitored for audit events)

- Eyeglass appliance with Ransomware Defender agent licenses or Easy Auditor Agent Licenses Or Performance Auditor Licenses
- PowerScale cluster with HDFS license to store the Analytics database for Easy Auditor only (Ransomware Defender only deployments no longer need HDFS pool as of release 2.5.5 or later)

Summary Overview of steps to install and configure Easy Auditor or Unified Ransomware Defender ,Easy Auditor and Performance Auditor

1. Configure Access Zone for Analytics database using an Access Zone with HDFS enabled
2. Configure SmartConnect on the Access Zone for the HDFS Database
3. Configure Smartconnect name for NFS mount access to each managed cluster
4. Create Eyeglass api token for ECA to authenticate to Eyeglass
5. Install ECA cluster
6. Configure ECA cluster master config
7. Push config to all nodes from master with ECA CLI
8. Start up the cluster software
9. Verify cluster is up and database is created

10. Verify Eyeglass Service heartbeat and ECA cluster nodes have registered with Eyeglass in the Managed Services Icon
11. Test Features

Isilon/PowerScale Protocol Audit Configuration (Required Step) (Ransomware Defender , Easy Auditor, Performance Auditor)

Overview

This section configures PowerScale file auditing required to monitor user behaviors.   The Audit protocol can be enabled on each Access Zone independently that requires monitoring.

1. Enable Protocol Access Auditing OneFS GUI
    a. Click Cluster Management > Auditing
    b. In the Settings area, select Enable Configuration Change Auditing and Enable Protocol Access Auditing checkbox.
    c. In the Audited Zones area, click Add Zones.
    d. In the Select Access Zones dialog box, select one or more access zones, and click Add Zones (do not add Eyeglass access zone).

Disable High Rate Audit Events Onefs 8.2 and later (Mandatory Step)

Directory Open and Directory close events generate unnecessary load on the cluster to log these event types, these event types are not used by Ransomware Defender or Easy Auditor (Default settings do not store these events in the Database). These events also cause performance issues on the cluster and high cluster overhead for these 2 events. It is required to disable these events.

Procedure to Disable High Rate Events

1.

- Login to the Onefs cluster over ssh as the root user.
    - Replace yellow highlight with access zone names that are enabled for auditing. This change takes effect immediately and will reduce audit overhead and increase auditing performance.
- **isi audit settings modify --zone=<zone_name> --remove-audit-success=open_directory,close_directory**

Preparation of Analytics Database or Index (Easy Auditor) (Required Step)

Prepare the PowerScale Cluster for HDFS

Prerequisites

1. **Easy Auditor only**
    2. **Must add minimum 3 PowerScale nodes added to new IP pool and assign the pool to the access zone created created for the audit database**
    3. **Must configure smartconnect zone name with FQDN**

4. **Must complete DNS delegation to the FQDN assigned to the new pool for HDFS access**
5. **Must Enable HDFS protocol on the new access zone (protocols tab in OneFS gui) Easy Auditor only**
6. **Must have HDFS license applied to the cluster**
7. **Must configure Snapshot schedule on the access zone path below every day at midnight with 30 day retention**
8. **Optional - Create SyncIQ policy to replicate the db to a DR site.**

1. Activate a license for HDFS. When a license is activated, the HDFS service is enabled by default.
2. Create "eyeglass" Access Zone with path "`/ifs/data/igls/analyticsdb`" for the HDFS connections from hadoop eyeglass compute clients (ECA) and under Available Authentication Providers, select only the Local System authentication provider.
3. Select create create zone base directory

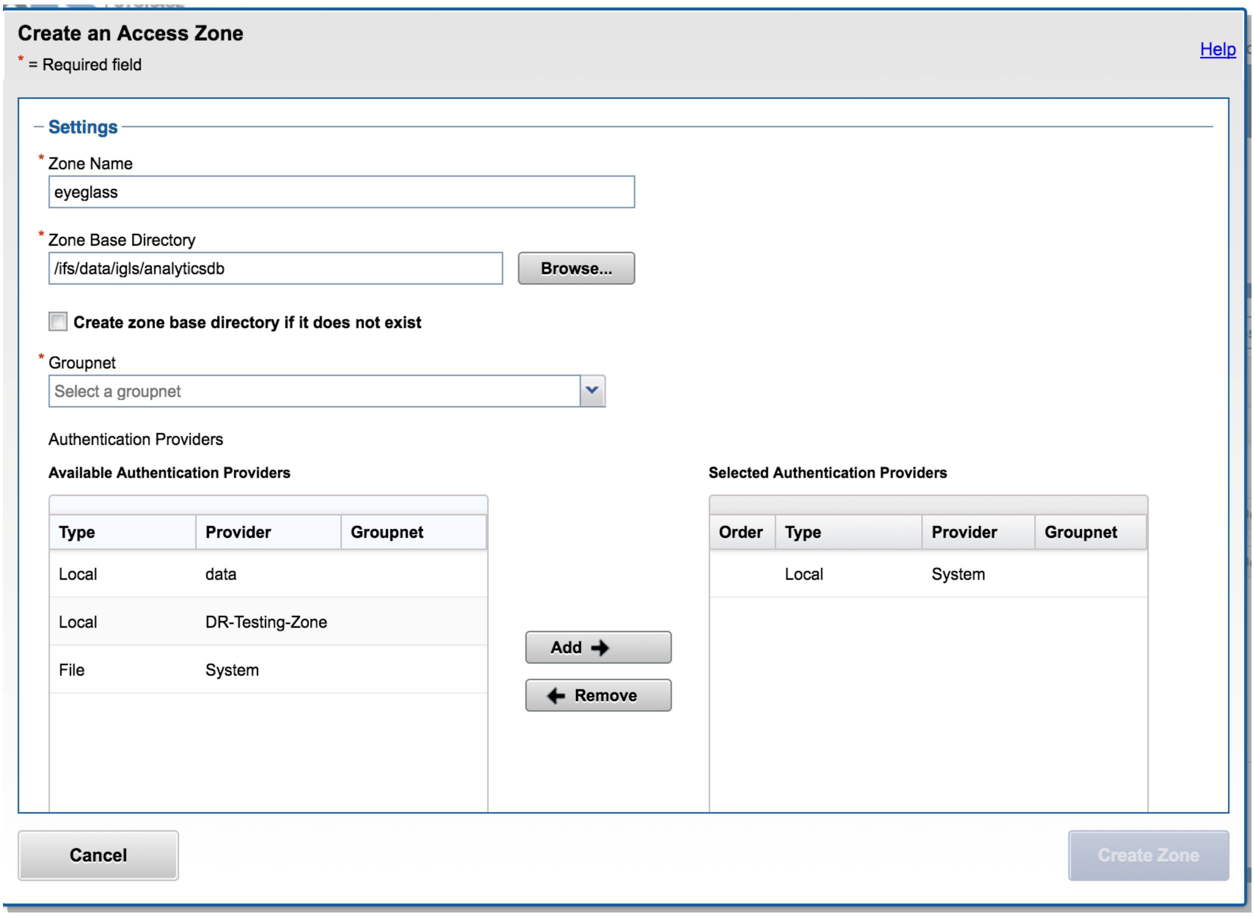

Create an Access Zone

* = Required field

Settings

* Zone Name: eyeglass

* Zone Base Directory: /ifs/data/igls/analyticsdb

□ Create zone base directory if it does not exist

* Groupnet: Select a groupnet

Authentication Providers

Available Authentication Providers

Type	Provider	Groupnet
Local	data	
Local	DR-Testing-Zone	
File	System	

Selected Authentication Providers

Order	Type	Provider	Groupnet
	Local	System	

NOTE: Ensure that Local System provider is at the top of the list. Additional AD providers are optional and not required.

NOTE: In OneFS 8.0.1 the Local System provider must be added using the command line. After adding, the GUI can be used to move the Local System provider to the top of the list.

```
isi zone zones modify eyeglass --add-auth-providers=local:system
```

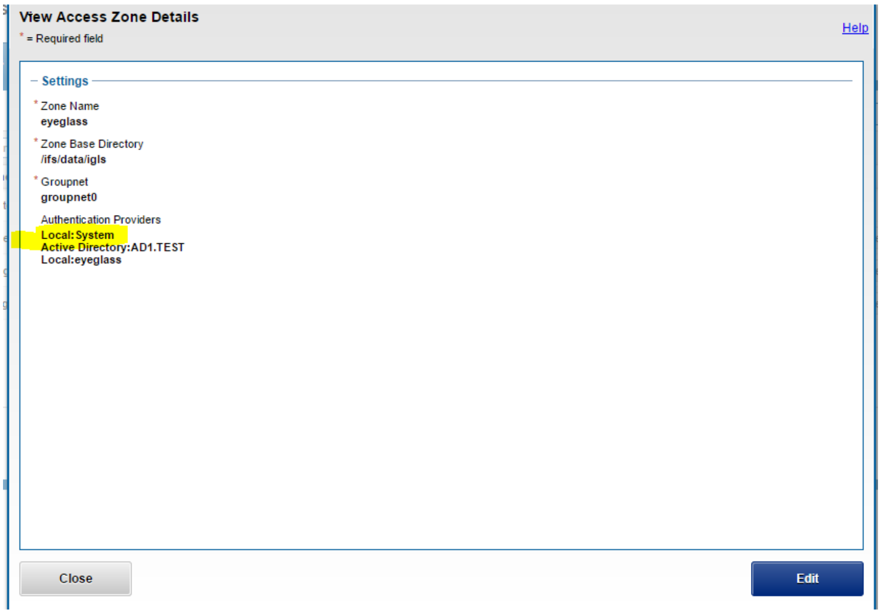

View Access Zone Details

* = Required field

Settings

* Zone Name: eyeglass

* Zone Base Directory: /ifs/data/igls

* Groupnet: groupnet0

Authentication Providers:
Local:System
Active Directory:AD1.TEST
Local:eyeglass

1. Set the HDFS root directory in eyeglass access zone that supports HDFS connections.

```
(Onefs 8.x)
    isi hdfs settings modify --root-
directory=path_to_hdfs_root_dir --
zone=access_zone_name_for_hdfs
    Example:
    isi hdfs settings modify --root-
directory=/ifs/data/igls/analyticsdb/  --
zone=eyeglass
```

1. Create **One** IP pool for HDFS access with at least 3 nodes in the pool to ensure high availability access to each ECA node, the Pool will be configured with round robin load balancing. This will be used for datanode and storage node access by the ECA cluster for the Analytics database.

   Command:

```
(Onefs 8.0)
      isi network pools create
      groupnet0.subnet0.hdfspool  --
      ranges=172.22.1.22-172.22.1.22 --ifaces
      1-4:10gige-1  --access-zone eyeglass --
      sc-dns-zone hdfs-mycluster.ad1.test --
      alloc-method static
```

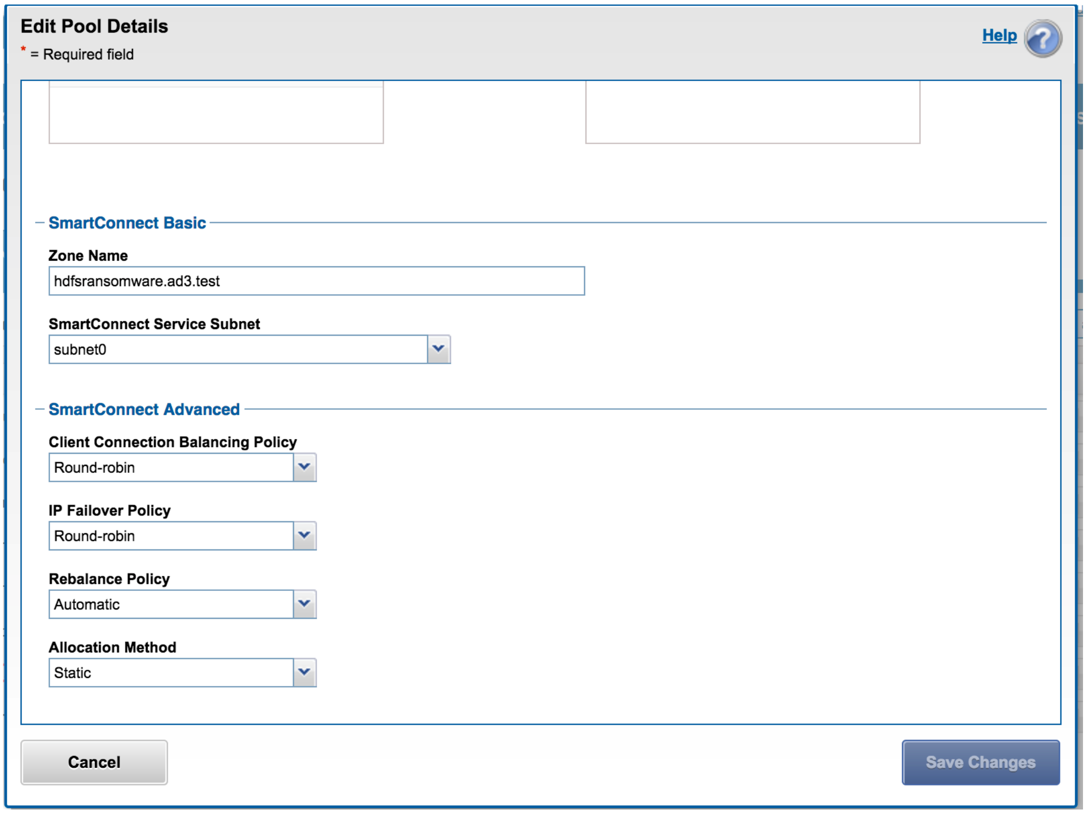
Edit Pool Details

* = Required field

SmartConnect Basic

Zone Name: hdfsransomware.ad3.test

SmartConnect Service Subnet: subnet0

SmartConnect Advanced

Client Connection Balancing Policy: Round-robin

IP Failover Policy: Round-robin

Rebalance Policy: Automatic

Allocation Method: Static

A virtual HDFS rack is a pool of nodes on the PowerScale cluster associated with a pool of Hadoop compute clients. To configure virtual HDFS racks on the PowerScale Cluster:

NOTE: The ***ip_address_range_for_client = the ip range used the by the ECA cluster VM's.***

Command:

```
(Onefs 8.0)
   isi hdfs racks create /hdfs_rack_name --
   zone=access_zone_name_for_hdfs --client-
   ip-ranges=ip_address_range_for_client --
   ip-pools=subnet:pool

   Example:
   isi hdfs racks create /hdfs-iglsrack0 --
   client-ip-ranges=172.22.1.18-172.22.1.20 -
   -ip-pools=subnet0:hdfspool --zone=eyeglass
   isi hdfs racks list --zone=eyeglass
   Name            Client IP Ranges          IP
   Pools
```

```
------------------------------------------
-------------------
/hdfs-rack0 172.22.1.18-172.22.1.20
subnet0:hdfspool
------------------------------------------
-------------------
Total: 1
```

1. Create local Hadoop user in the System access zone.

   NOTE: User ID must be *eyeglasshdfs*.

   Command:

```
(Onefs 8.0)
    isi auth users create --name=eyeglasshdfs -
-provider=local --enabled=yes --zone=system
    Example:
    isi auth users create --name=eyeglasshdfs -
-provider=local --enabled=yes --password-
expires=no --zone=system
```

1.

2. Login via SSH to the PowerScale cluster as the **root user** to change the ownership, permissions and block inherited permissions from parent folders on the HDFS path that will be used by Eyeglass ECA clusters.

   1. **chown -R eyeglasshdfs:'Isilon Users' /ifs/data/igls/analyticsdb/**
   2. **chmod -R 755 /ifs/data/igls/analyticsdb/**
   3. **chmod -c +dacl_protected /ifs/data/igls/analyticsdb**
   4. **NOTE: if using a cluster in compliance mode do not run the commands above and run the command below.**

5. `chmod 777 /ifs/data/igls/analyticsdb/`

1. **Analytics Cluster setup Complete.**

Deployment, Installation and Configuration ECA Cluster (Required Step)

OVA Install Prerequistes

Configuration Item	**Notes**
see scaling section	The OVA file will deploy 3 vm's. to build a 6 node cluster, deploy the OVA twice and move the VM's into the first Cluster object in vcenter. See instructions below to correctly move VM's into a single vapp in vcenter.
vSphere 6.x or higher	
1x ip address on the same subnet for each node	
Gateway	
Network Mask	
DNS IP	
NTP server IP	
IP Address of Eyeglass	
API token from Eyeglass	
Unique cluster name (lower case no special characters)	

Installation Procedure of the ECA Vmware OVA

1. The deployment is based on three node ECA appliances.
2. Download the Superna Eyeglass™ OVF from https://www.supernaeyeglass.com/downloads
3. Unzip into a directory on a machine with vSphere client installed
4. Install the OVA using steps below with **HTML vCenter web interface**.
    a. WARNING: Access vCenter with an FQDN DNS name NOT an IP address, a bug in vCenter will generate an error during OVA validation.

5. NOTE: IF DEPLOYING A 6 OR 9 NODE CLUSTER FOR EASY AUDITOR THE 3 VM vAPP OVA STEPS BELOW WILL BE DOWN TWICE FOR 6 NODE AND THREE TIMES FOR 9 NODE CLUSTER.   THIS WILL CREATE 2 OR 3 vAPP'S AND THE VM'S FROM EACH vAPP'S CAN BE MOVED INTO A SINGLE COMMON vAPP OBJECT IN VCENTER AND REMOVE THE EMPTY vAPP OBJECTS IN VCENTER
6. **NOTE: The ECA name on the 2nd or 3rd vAPP deployment does not need to match the first vAPP ECA name. Once completed the ECA name used for the first ECA cluster will be synced to all VM's defined on node 1 ECA cluster master configuration file.**
7. **MANDATORY STEP: POWER ON THE VAPP AFTER DEPLOYMENT SO THAT IP ADDRESSES GET ASSIGNED. DO NOT REMOVE VM'S FROM VAPP BEFORE POWER ON.**
    a. Make sure the first boot steps completes by reviewing running this command and repeat on each ECA vm to ensure it has completed all first boot steps.
        i. sudo systemctl status superna-on-boot (enter admin password and verify the first boot process completes)
        ii. Verify the process has completed and exited the processing.
        iii. cat /var/log/superna-on-boot.log, it must show done before the boot process is completed. Do not proceed until this steps finishes.

    b. **Procedures**
        i. For the 2nd/3rd ECA OVA deployment power on the vapp
        ii. Ping each VM ip in the cluster until it responds to ping. (this allows first boot scripts to run)

iii. Once the VM's ping you can move the VM's from the vApp to the 1st ECA vapp deployed, repeat for each VM and once done the empty vApp can be deleted. Drag and drop the VM from the vapp to the ECA 1 vapp.

iv. Repeat for each ECA OVA deployed AFTER the first ECA OVA.

8. 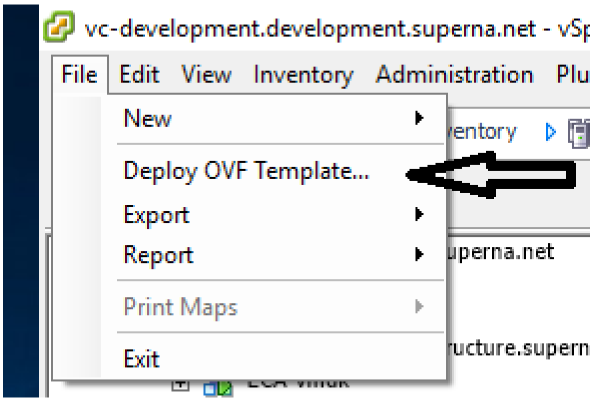

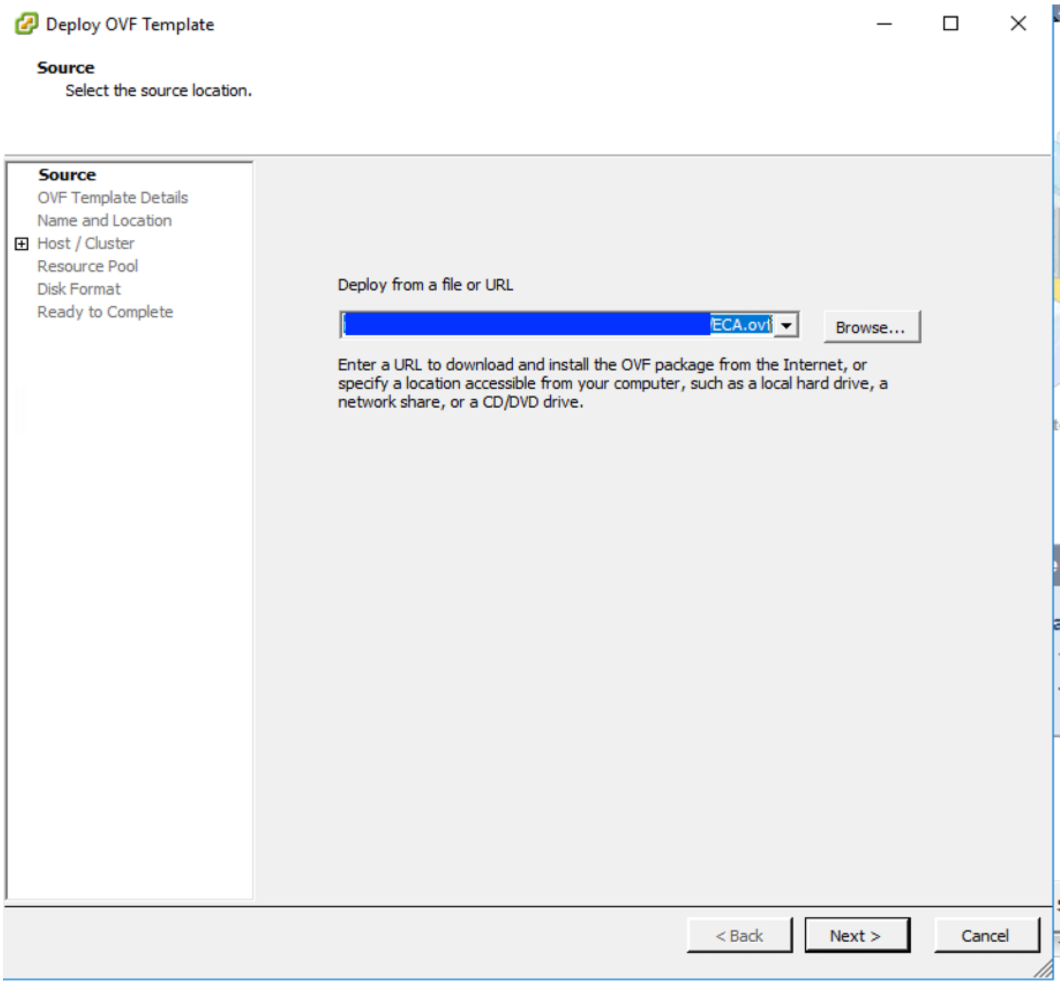

Source

Select the source location.

Deploy from a file or URL

Enter a URL to download and install the OVF package from the Internet, or specify a location accessible from your computer, such as a local hard drive, a network share, or a CD/DVD drive.

9. **vCenter HTML Example**

10.

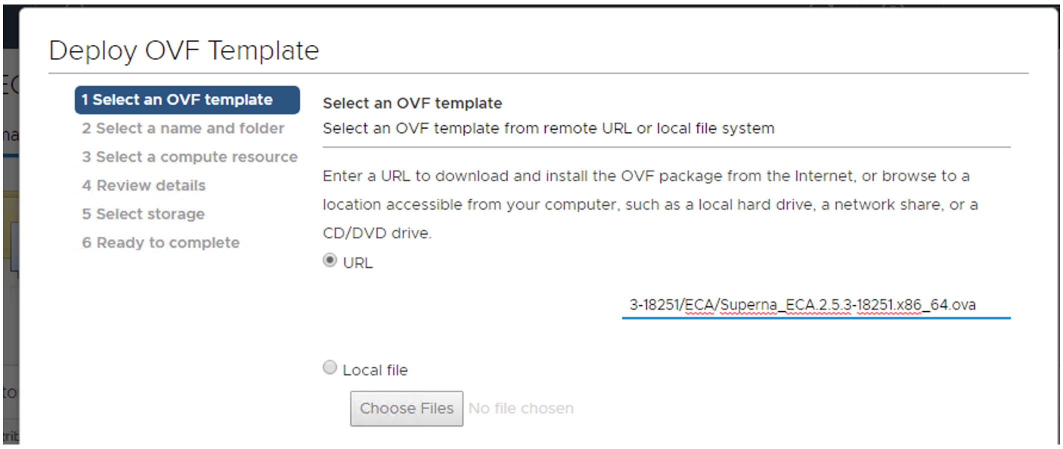

Deploy OVF Template

Select an OVF template

Select an OVF template from remote URL or local file system

Enter a URL to download and install the OVF package from the Internet, or browse to a location accessible from your computer, such as a local hard drive, a network share, or a CD/DVD drive.

● URL

3-18251/ECA/Superna_ECA.2.5.3-18251.x86_64.ova

○ Local file

11. Deploy from a file or URL where the OVA was saved

12. Using vCenter client set required VM settings, for datastore, networking. **NOTE: Leave setting as Fixed IP address**

13. Complete the networking sections as follows:

    a. ECA Cluster name **(NOTE: must be lowercase < 8 characters and no special characters, with only letters)**

    b. IMPORTANT: ECA Cluster name cannot include _ as this will cause some services to fail

    c. All VM are on the same subnet

    d. Enter network mask (will be applied to all VM's)

    e. Gateway IP

    f. DNS server (must be able to resolve the igls.<your domain name here>) (Use nameserver IP address)

    g. NOTE: Agent node 1 is the master node where all ECA CLI commands are executed for cluster management

14. vCenter Windows client example

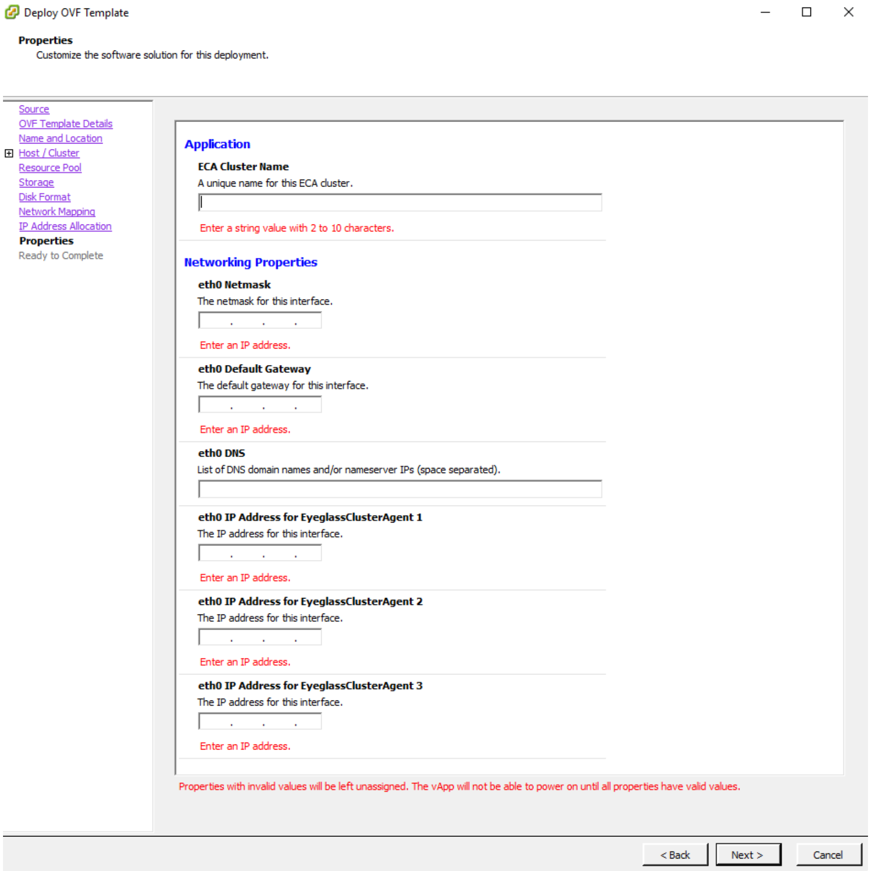

Deploy OVF Template

Properties
Customize the software solution for this deployment.

Application

ECA Cluster Name
A unique name for this ECA cluster.

Enter a string value with 2 to 10 characters.

Networking Properties

eth0 Netmask
The netmask for this interface.

Enter an IP address.

eth0 Default Gateway
The default gateway for this interface.

Enter an IP address.

eth0 DNS
List of DNS domain names and/or nameserver IPs (space separated).

eth0 IP Address for EyeglassClusterAgent 1
The IP address for this interface.

Enter an IP address.

eth0 IP Address for EyeglassClusterAgent 2
The IP address for this interface.

Enter an IP address.

eth0 IP Address for EyeglassClusterAgent 3
The IP address for this interface.

Enter an IP address.

Properties with invalid values will be left unassigned. The vApp will not be able to power on until all properties have valid values.

a.

b. vCenter HTML Client Example

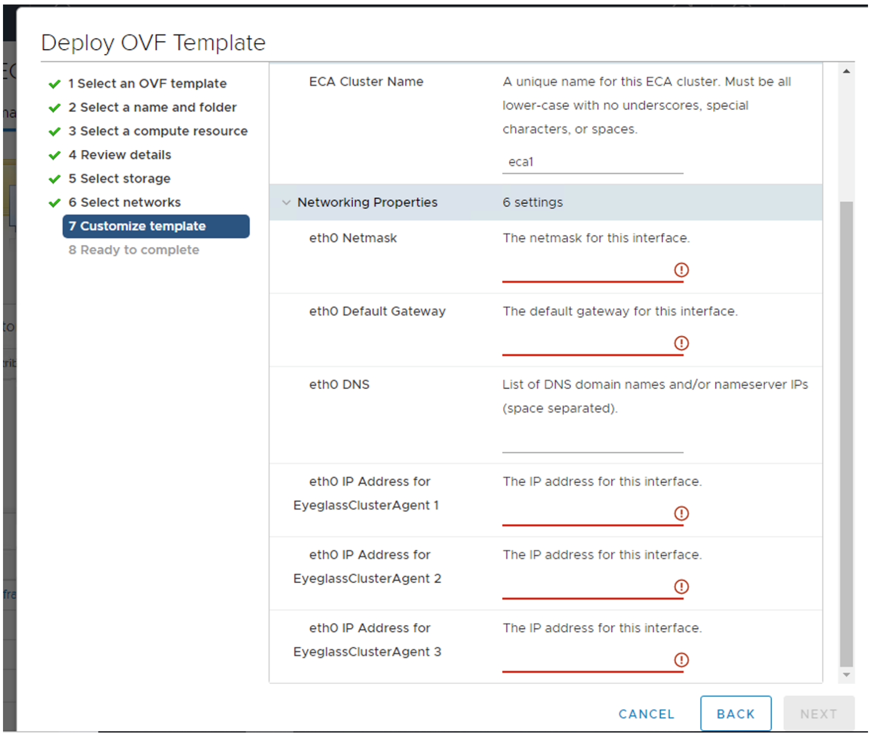

Deploy OVF Template

ECA Cluster Name	A unique name for this ECA cluster. Must be all lower-case with no underscores, special characters, or spaces. eca1
Networking Properties	6 settings
eth0 Netmask	The netmask for this interface.
eth0 Default Gateway	The default gateway for this interface.
eth0 DNS	List of DNS domain names and/or nameserver IPs (space separated).
eth0 IP Address for EyeglassClusterAgent 1	The IP address for this interface.
eth0 IP Address for EyeglassClusterAgent 2	The IP address for this interface.
eth0 IP Address for EyeglassClusterAgent 3	The IP address for this interface.

c.

15. Example OVA vAPP after deployment

a. 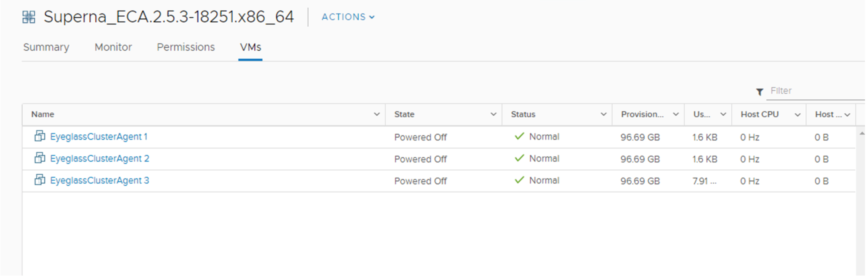

Name	State	Status	Provision...	Us...	Host CPU	Host ...
EyeglassClusterAgent 1	Powered Off	✓ Normal	96.69 GB	1.6 KB	0 Hz	0 B
EyeglassClusterAgent 2	Powered Off	✓ Normal	96.69 GB	1.6 KB	0 Hz	0 B
EyeglassClusterAgent 3	Powered Off	✓ Normal	96.69 GB	7.91 ...	0 Hz	0 B

16. OPTIONAL If you are deploying a 6 or 9 node ECA cluster repeat the deployment again following the instractions above and set the ip addresses on the new VM's to expand the overall cluster ip range to 6 or 9 VM's. The ECA name can be any value since this will be synced from node 1 of the first OVA cluster that was deployed.
    a. **After deployment of the 2nd or 3rd ECA, open the vAPP and rename the vm's as follows:**
        i. **6 or 9 Node ECA:**
            1. EyeglassClusteredAgent 1 to EyeglassClusteredAgent 4
            2. EyeglassClusteredAgent 2 to EyeglassClusteredAgent 5
            3. EyeglassClusteredAgent 3 to EyeglassClusteredAgent 6
            4. **ONLY** If a 9 node ECA cluster continue to rename the 3rd OVA VM's inside the vAPP
            5. EyeglassClusteredAgent 1 to EyeglassClusteredAgent 7
            6. EyeglassClusteredAgent 2 to EyeglassClusteredAgent 8

7. EyeglassClusteredAgent 3 to EyeglassClusteredAgent 9

ii. Now drag and drop the vm inside each of the vAPP's into the vAPP created for the first 3 VM's deployed.

iii. Once completed you can delete the empty vAPP deployed for VM's 4-9.

iv. Once done the initial vAPP will look like this (9 node ECA shown).

v. 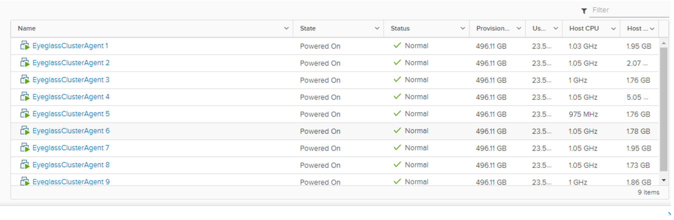

Filter

Name	State	Status	Provision...	Us...	Host CPU	Host ...
EyeglassClusterAgent 1	Powered On	✓ Normal	496.11 GB	23.5...	1.03 GHz	1.95 GB
EyeglassClusterAgent 2	Powered On	✓ Normal	496.11 GB	23.5...	1.05 GHz	2.07 ...
EyeglassClusterAgent 3	Powered On	✓ Normal	496.11 GB	23.5...	1 GHz	1.76 GB
EyeglassClusterAgent 4	Powered On	✓ Normal	496.11 GB	23.5...	1.05 GHz	5.05 ...
EyeglassClusterAgent 5	Powered On	✓ Normal	496.11 GB	23.5...	975 MHz	1.76 GB
EyeglassClusterAgent 6	Powered On	✓ Normal	496.11 GB	23.5...	1.05 GHz	1.78 GB
EyeglassClusterAgent 7	Powered On	✓ Normal	496.11 GB	23.5...	1.05 GHz	1.95 GB
EyeglassClusterAgent 8	Powered On	✓ Normal	496.11 GB	23.5...	1.05 GHz	1.73 GB
EyeglassClusterAgent 9	Powered On	✓ Normal	496.11 GB	23.5...	1 GHz	1.86 GB

9 items

vi. **Done**

17. **After Deployment is complete Power on the vAPP**

a. Ping each ip address to make sure each node has finished booting

b. Login via SSH to **the Master Node** (Node 1) using the “ecaadmin” account default password 3y3gl4ss and run the following command:

c. **ecactl components configure-nodes** (this command sets up keyless ssh for the ecaadmin user to manage the cluster)

d. **On Eyeglass Appliance:** generate a unique API Token from Superna Eyeglass REST API Window. Once a token has been generated for the ECA Cluster, it can be used in that ECA’s startup command for authentication.

e. Login to Eyeglass goto main menu Eyeglass REST API menu item. Create a new API token. This will be used in the startup file for the ECA cluster to authenticate to the Eyeglass VM and register ECA services.

f. 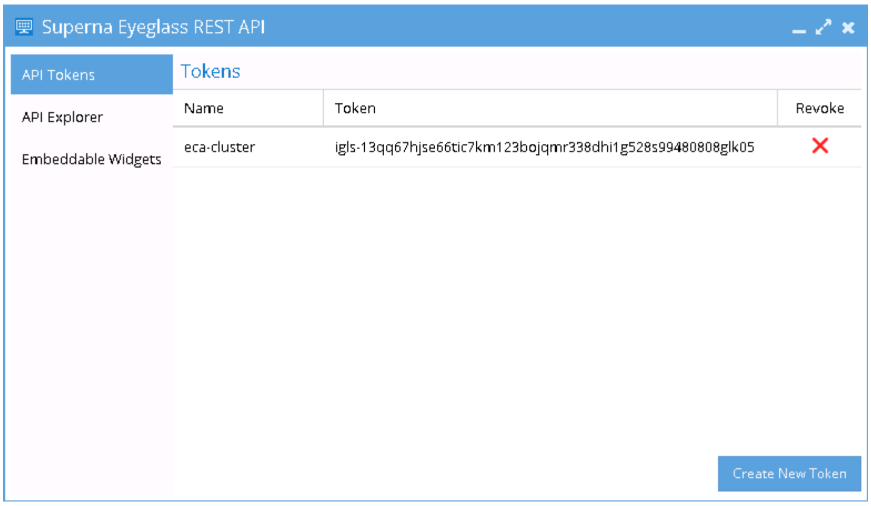

Name	Token	Revoke
eca-cluster	igls-13qq67hjse66tic7km123bojqmr338dhi1g528s99480808glk05	✕

18. On ECA Cluster Master node ip 1

a. Login to that VM using assh as the ecaadmin user default password 3y3gl4ss. From this point on, commands will only be executed on the master node.

b. On the master node, edit the file nano /opt/superna/eca/eca-env-common.conf , and change these five settings to reflect your environment. Replace the variables accordingly.

c. Set the IP address or FQDN of the Eyeglass appliance and the API Token (created above), uncomment the parameter lines before save file. I.e:

   i. export EYEGLASS_LOCATION=*ip_addr_of_eyeglass_appliance*

   ii. export EYEGLASS_API_TOKEN=*Eyeglass_API_token*

d. Verify the IP addresses for the nodes in your cluster. **It is important that NODE_1 be the master**, (i.e. the IP address of the node you're currently logged into.)

i. NOTE: add additional ECA_LOCATION_NODE_X=x.x.x.x for additional node in the ECA cluster depending on ECA cluster size. All nodes in the cluster must be listed in the file. Copy a line and paste to add additional ECA nodes and make sure to change the node number example to add the 4th ECA VM it would like like this export ECA_LOCATION_NODE_4=

ii. export ECA_LOCATION_NODE_1=*ip_addr_of_node_1 (set by first boot from the OVF)*

iii. export ECA_LOCATION_NODE_2=*ip_addr_of_node_2 (set by first boot from the OVF)*

iv. export ECA_LOCATION_NODE_3=*ip_addr_of_node_3 (set by first boot from the OVF)*

e. Set the HDFS path to the SmartConnect name setup in the Analytics database configuration steps. Replace the FQDN *hdfs_sc_zone_name* with <your smartconnect FQDN >.

f. NOTE: Do not change any other value. Whatever is entered here is created as a subdirectory of the HDFS root directory that was set earlier.

g. export ISILON_HDFS_ROOT='hdfs://*hdfs_sc_zone_name*:8020/*eca1*'

19. **Done: Continue on to the Cluster Auditing Configuration Section**

Time Configuration PowerScale, Eyeglass, ECA cluster (Required Step) (Ransomware Defender , Easy Auditor, Performance Auditor)

Overview: To get accurate auditing features for Ransomware or Easy Auditor time sync between all components is critical step. NTP should be used on all VM's and use the same NTP source.

1. Verify PowerScale clusters being monitored are using an NTP server. Many Internet time sources exist or internal Enterprise server IP address
2. Enable NTP on all PowerScale clusters
3. Eyeglass VM configure the same NTP servers used by PowerScale by following.
4. On each ECA VM repeat the YAST steps above to configure NTP on each VM.

If NTP and ESX host time sync conflict it may be necessary to disable ESX host time sync to the ECA nodes to allow ECA nodes to get NTP time versus esx host. This ensure that DB queries and each node has consistent time in sync across Eyeglass VM's and ECA nodes.

How to disable VMware vSphere ESXi host time sync (Mandatory Step)

For ECA:

1. Initiate ecactl cluster down
2. Power down ECA vApp
3. Change VMware time sync configuration as below:
4. Click on Virtual Machine
5. Right click on ECA node1
6. Click Edit Settings..

7. Click on Option
8. Click VMware Tools
9. Uncheck ‘Synchronize guest time with host’
10. Click OK
11. Power up vApp
12. Initiate ecactl cluster up
13.

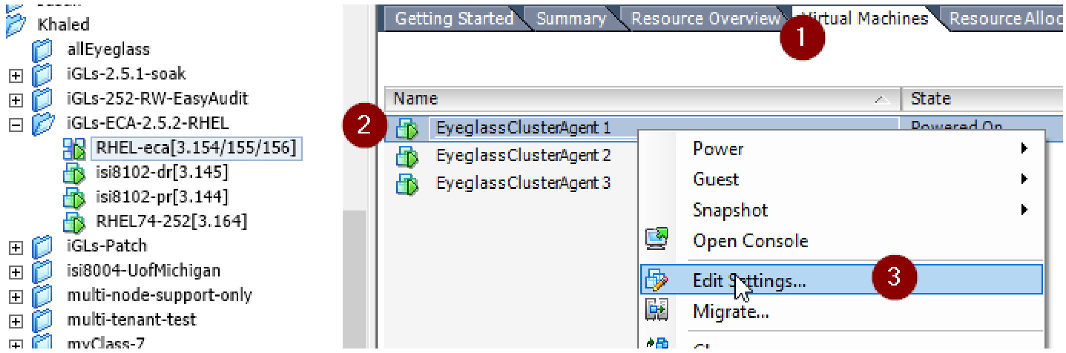

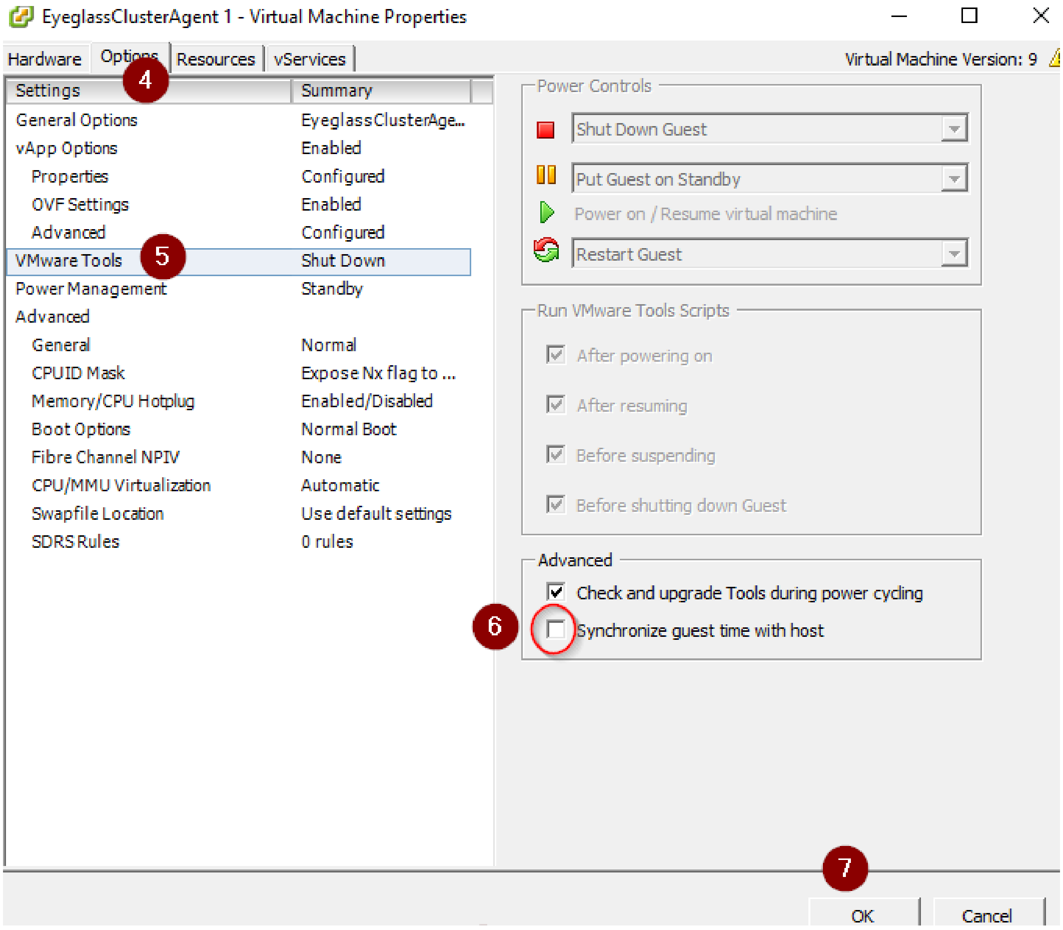
EyeglassClusterAgent 1 - Virtual Machine Properties
Hardware
Options
Resources
vServices
Virtual Machine Version: 9
4
5
6
7
OK
Cancel

Settings	Summary
General Options	EyeglassClusterAge...
vApp Options	Enabled
Properties	Configured
OVF Settings	Enabled
Advanced	Configured
VMware Tools	Shut Down
Power Management	Standby
Advanced	
General	Normal
CPUID Mask	Expose Nx flag to ...
Memory/CPU Hotplug	Enabled/Disabled
Boot Options	Normal Boot
Fibre Channel NPIV	None
CPU/MMU Virtualization	Automatic
Swapfile Location	Use default settings
SDRS Rules	0 rules

Power Controls

Shut Down Guest

Put Guest on Standby

Power on / Resume virtual machine

Restart Guest

Run VMware Tools Scripts

☑ After powering on

☑ After resuming

☑ Before suspending

☑ Before shutting down Guest

Advanced

☑ Check and upgrade Tools during power cycling

□ Synchronize guest time with host

-

NOTE: Apply this change on ALL ECA nodes. Perform same steps for Eyeglass Appliance if needed
New changes may take up to 15 mins in some cases, you may need to restart ntpd after cluster up

How to Configure Audit Data Ingestion on the Isilon/Powerscale (Ransomware Defender, Easy Auditor, Performance Auditor) (Required Step)

Prerequisites for Isilon/Powerscale Audit Data NFS Export

1. A Smartconnect name configured in the system zone for the NFS export created on /ifs/.ifsvar/audit/logs (steps below)
2. Isilon/Powerscale IP pool set to dynamic mode for NFS mount used by the ECA cluster nodes for a highly available NFS mount.

    1.

Create a read-only NFS Export on the Isilon/PowerScale cluster(s) that will be managed

1. Note all managed clusters by this ECA cluster will require an export created for audit event processing

- In the example below replace <ECA_IP_1> with the IP address of ECA nodes 1, 2, 3, if a larger ECA cluster you MUST add all IP addresses in the ECA cluster to the NFS export root and client lists.

    1. **isi nfs exports create /ifs/.ifsvar/audit/logs --root-clients="*<ECA_IP_1>,<ECA_IP_2>,<ECA_IP_3>*" --**

clients="*<ECA_IP_1>,<ECA_IP_2>,<ECA_IP_3>*" --read-only=true -f --description "Easy Auditor Audit Log Export"

1.

Configure an NFS mount point on each ECA node to Read Audit data from Isilon/PowerScale (Required)

1. Audit events are ingested over NFS mounts on ECA nodes 1 - X (where X is the size of your ECA cluster). Follow the steps below to add the export to each of the VM's.
2. **What you will need to complete this step on nodes 2 - x (where x is the last node IP in the cluster):**
   a. Cluster GUID and cluster name for each cluster to be indexed
   b. Cluster name as shown on top right corner after login to OneFS GUI NOTE: this name is case sensitive and the mount MUST match the case of the cluster name.
      i. See example of where to get this information.

ii. 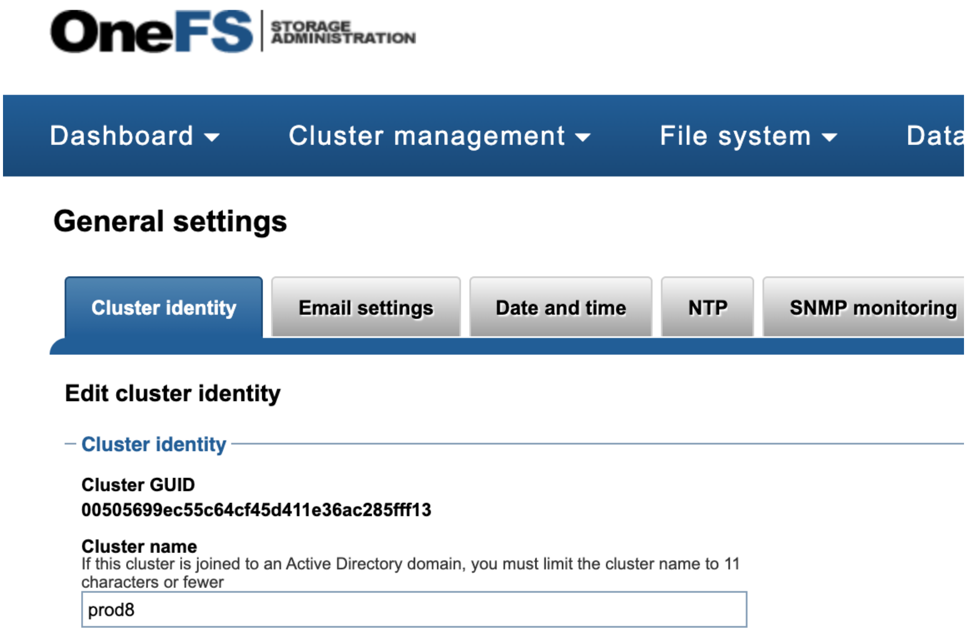

Edit cluster identity

Cluster identity

Cluster GUID
00505699ec55c64cf45d411e36ac285fff13

Cluster name
If this cluster is joined to an Active Directory domain, you must limit the cluster name to 11 characters or fewer

prod8

3. **Login to ECA node 1**

   a. **ssh ecaadmin@x.x.x.x (where x.x.x.x is node 1 IP of the ECA cluster)**

4. **Create local mount directory and sync to all nodes**

   a. **ecactl cluster exec "sudo mkdir -p /opt/superna/mnt/audit/GUID/clusternamehere/"**

      i. replace GUID and clusternamehere with correct values from above (note cluster name is case sensitive and must match the cluster name case as shown in Onefs)

      ii. enter the admin password when prompted on each ECA node

   b. Verify the folder exists on all ECA nodes

      i. ecactl cluster exec " ll `ls` /opt/superna/mnt/audit"

5. **NFS Mount Setup with Centralized Mount file for all Nodes with Auto-Mount**

a. **NOTE: This option will mount on cluster up using a centralized file to control the mount. This simplifies changing mounts on nodes and provides cluster up mount diagnostics.**

b. **Configuration Steps for Auto mount**

   i. **nano /opt/superna/eca/eca-env-common.conf**

   ii. **add a variable to ensure the cluster up stops if the NFS mount fails copy this line and paste into the file export STOP_ON_AUTOMOUNT_FAIL=true**

   iii. ssh to ECA node 1 as ecaadmin user to enable auto mount and make sure it starts on OS reboot. **NOTE: for each node you will be prompted for the ecaadmin password.**

   iv. **ecactl cluster exec "sudo systemctl unmask autofs"**

   v. **ecactl cluster exec "sudo systemctl start autofs"**

   vi. Check and ensure the service is running

      1. **ecactl cluster exec "sudo systemctl status autofs"**

c. Add new entry to auto.nfs file on ECA node 1

   i. **NOTE: the FQDN should be a smartconnect name for a pool in the System Access Zone IP Pool. <NAME> is the cluster name collected from the section above. GUID is the cluster GUID from the General settings screen of Onefs**

   ii. **NOTE: Add 1 line for each Isilon/Powerscale cluster that will be monitored from this ECA cluster.**

iii. **Fix the command below with correct values highlighted in yellow and paste into the SSH session. Repeat for each Isilon/Powerscale cluster mount**

1. echo -e "\n/opt/superna/mnt/audit/<GUID>/<NAME> --fstype=nfs,nfsvers=3,ro,soft <FQDN>:/ifs/.ifsvar/audit/logs" >> /opt/superna/eca/data/audit-nfs/auto.nfs

iv. Verify the auto.nfs file contents

1. **cat /opt/superna/eca/data/audit-nfs/auto.nfs**

d. Push the configuration to all ECA nodes

i. **ecactl cluster push-config**

e. **Start Auto mount and verify the mount**

i. **ecactl cluster exec "sudo systemctl restart autofs" (note you will be asked to enter the ecaadmin password for each ECA node)**

ii. **Check the mount by typing command below**

1. mount

f. The cluster up command will read the mount file and mount on each ECA node during cluster up.

Start up the ECA Cluster (Required)

1. At this point you can start up the cluster
2. SSH to ECA node 1 as ecaadmin and run the command below.
3. **ecactl cluster up (Note can take 5-8 minutes to complete)**

4. **See trouble shooting section below for commands to verify startup and post startup status.**

Verify ECA Remote Monitoring Connection from the Eyeglass Appliance (Required Step) (Ransomware Defender , Easy Auditor, Performance Auditor)

1. Login to Eyeglass as admin user
2. Check the status of the ECA Cluster, click ‘Manage Service’ Icon and click on + to expand the container or services for each eca node review image below.
3. Verify the ip addresses of the ECA nodes are listed.
4. Verify all cluster nodes show and all docker containers show green health.
5. **NOTE: Hbase status can take 5 minutes to transition from warning to Green.**
6.

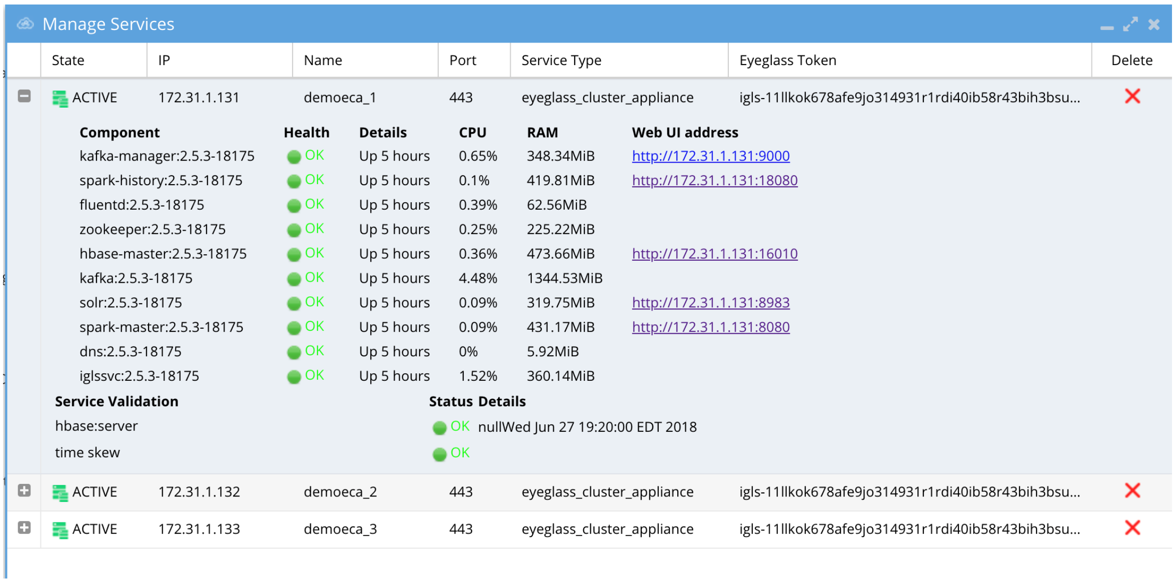

Manage Services

	State	IP	Name	Port	Service Type	Eyeglass Token	Delete
⊟	ACTIVE	172.31.1.131	demoeca_1	443	eyeglass_cluster_appliance	igls-11llkok678afe9jo314931r1rdi40ib58r43bih3bsu...	✕

Component	Health	Details	CPU	RAM	Web UI address
kafka-manager:2.5.3-18175	● OK	Up 5 hours	0.65%	348.34MiB	http://172.31.1.131:9000
spark-history:2.5.3-18175	● OK	Up 5 hours	0.1%	419.81MiB	http://172.31.1.131:18080
fluentd:2.5.3-18175	● OK	Up 5 hours	0.39%	62.56MiB	
zookeeper:2.5.3-18175	● OK	Up 5 hours	0.25%	225.22MiB	
hbase-master:2.5.3-18175	● OK	Up 5 hours	0.36%	473.66MiB	http://172.31.1.131:16010
kafka:2.5.3-18175	● OK	Up 5 hours	4.48%	1344.53MiB	
solr:2.5.3-18175	● OK	Up 5 hours	0.09%	319.75MiB	http://172.31.1.131:8983
spark-master:2.5.3-18175	● OK	Up 5 hours	0.09%	431.17MiB	http://172.31.1.131:8080
dns:2.5.3-18175	● OK	Up 5 hours	0%	5.92MiB	
iglssvc:2.5.3-18175	● OK	Up 5 hours	1.52%	360.14MiB	

Service Validation	Status Details
hbase:server	● OK nullWed Jun 27 19:20:00 EDT 2018
time skew	● OK

	State	IP	Name	Port	Service Type	Eyeglass Token	Delete
⊞	ACTIVE	172.31.1.132	demoeca_2	443	eyeglass_cluster_appliance	igls-11llkok678afe9jo314931r1rdi40ib58r43bih3bsu...	✕
⊞	ACTIVE	172.31.1.133	demoeca_3	443	eyeglass_cluster_appliance	igls-11llkok678afe9jo314931r1rdi40ib58r43bih3bsu...	✕

How to Backup and Protect the Audit Database with SnapshotIQ (Easy Auditor) (Required Step)

Use the PowerScale native SnapshotIQ feature to backup the audit data. Procedure documented here.

Installation Complete

At this point in the document, installation should be complete. Review the CLI guides for customizing different products.

Ransomware and Easy Auditor IGLS CLI command Reference

Additional configuration can be completed from the CLI to customize the default behavior.

- See Eyeglass CLI commands for Ransomware Defender and Easy Auditor
- See ECA CLI Commands

How to Upgrade the ECA cluster Software For Easy Auditor , Ransomware Defender and Performance Auditor

NOTE: Contact support first before upgrading the cluster for compatibility with the Eyeglass version. Both Eyeglass and ECA MUST be running the same version.

NOTE: Upgrade assistance is scheduled and is a service not covered under 7/24 support. Please review the EULA terms and conditions.

NOTE: Always take a VM level snapshot before any upgrade steps to provide a roll back to the previous release

Steps to upgrade:

1. Disable Ransomware Defender, Easy Auditor, Performance Auditor functionality before beginning upgrade - **required first step**
    a. Log in to ECA Node 1 using "ecaadmin" credentials.
    b. Issue the following command: ecactl cluster down
    c. Please wait for the procedure to complete on all involved ECA nodes.
    d. Done!
2. Upgrade Eyeglass VM first and downlod the latest release from here https://www.supernaeyeglass.com/downloads
    a. **NOTE: Eyeglass and ECA cluster software must be upgraded to the same version.**
    b. Follow the guide here
    c. Double check licenses are assigned to the correct clusters based on information here.
    d. Double check Ransomware Defender, Easy Auditor and Performance Auditor settings are correct and match the settings you had before the upgrade.

3. Then Downlod latest GA Release for the ECA upgrade following instructions here https://www.supernaeyeglass.com/downloads
4. Log in to ECA Node 1 using "ecaadmin" credentials.
5. Note: ECA is in down state - ecactl cluster down already done in step 1
6. Verify by executing command: ecactl cluster stats
7. No containers should be running
8. If containers running stop by executing the command and waiting for it to complete on all nodes: ecactl cluster down
9. Once the above steps are complete
    a. Use WinSCP to transfer run file on node 1 (Master Node) in /home/ecaadmin directory
    b. ssh to ECA node 1 as ecaadmin user example ssh ecaadmin@x.x.x.x
    c. cd /home/ecaadmin
    d. chmod +x ecaxxxxxxx.run file (xxxx is name of file)
    e. ./ecaxxxxxx.run file (xxxx is name of file)
    f. Enter password for ecaadmin when prompted
    g. wait for installation to complete
    h. Capture the upgrade text to provide to support if needed
10. Software upgrade is completed
11. Now bring up the ECA cluster
    a. **ecactl cluster up**

b. wait until all services are started on all nodes. If there are any errors copy the upgrade log referenced when the upgrade starts and use Winscp to copy the log file to your pc or copy and paste to a support case.

12. Once completed, login to Eyeglass open the managed services icon, verify all ECA nodes show green and online. If any services show warning or inactive wait at least 5 minutes, if the condition persists, open a support case.

a. 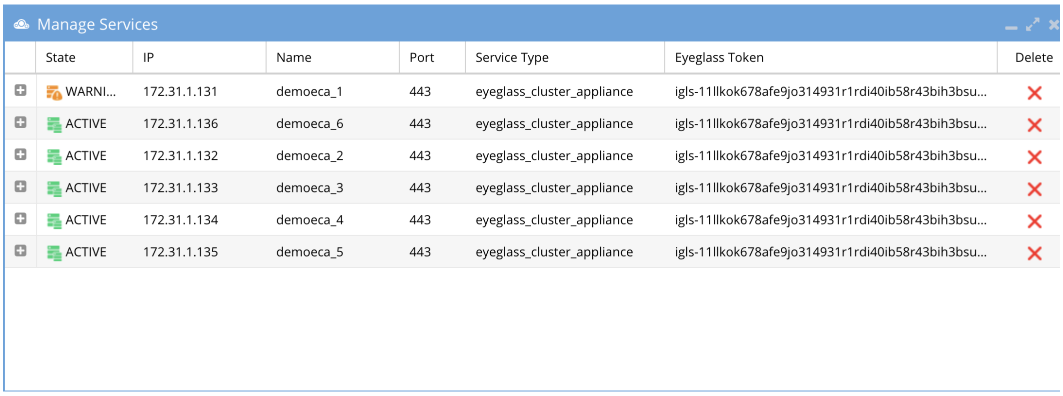

Manage Services

	State	IP	Name	Port	Service Type	Eyeglass Token	Delete
⊞	WARNI...	172.31.1.131	demoeca_1	443	eyeglass_cluster_appliance	igls-11llkok678afe9jo314931r1rdi40ib58r43bih3bsu...	✕
⊞	ACTIVE	172.31.1.136	demoeca_6	443	eyeglass_cluster_appliance	igls-11llkok678afe9jo314931r1rdi40ib58r43bih3bsu...	✕
⊞	ACTIVE	172.31.1.132	demoeca_2	443	eyeglass_cluster_appliance	igls-11llkok678afe9jo314931r1rdi40ib58r43bih3bsu...	✕
⊞	ACTIVE	172.31.1.133	demoeca_3	443	eyeglass_cluster_appliance	igls-11llkok678afe9jo314931r1rdi40ib58r43bih3bsu...	✕
⊞	ACTIVE	172.31.1.134	demoeca_4	443	eyeglass_cluster_appliance	igls-11llkok678afe9jo314931r1rdi40ib58r43bih3bsu...	✕
⊞	ACTIVE	172.31.1.135	demoeca_5	443	eyeglass_cluster_appliance	igls-11llkok678afe9jo314931r1rdi40ib58r43bih3bsu...	✕

13. If the above steps pass and all ECA nodes show green online

a. Then use the Security guard test in Ransomware Defender or run the RoboAudit feature run in Easy auditor, to validate audit data ingestion is functioning.

14. Consult the admin guide of each product to start a manual test of these features.

How to Migrate ECA cluster settings to a new ECA cluster deployment - To upgrade Open Suse OS.

To upgrade an ECA cluster OS, it is easier to migrate the settings to a new ECA cluster deployed with the new OS. Follow these steps to deploy a new ECA cluster and migrate configuration to the new ECA cluster.

Prerequisites

1. The ECA cluster has a logical name between the nodes, when deploying the new OVA the deployment promtps for the ECA cluster name and this should use the same name as the previous ECA cluster.
    a. How to get the ECA cluster name
    b. Login to eca node 1 via ssh ecaadmin@x.x.x.x and then run the command below:
    c. cat /opt/superna/eca/eca-env-common.conf | grep ECA_CLUSTER_ID
    d. Use the value returned after the = sign when deploying the new ECA cluster.
    e. Using winscp utility copy the following files from ECA node 1 of the **existing** ECA cluster, login using ecaadmin user
        i. /opt/superna/eca/eca-env-common.conf
        ii. /opt/superna/eca/docker-compose.overrides.yml
        iii. /opt/superna/eca/conf/common/overrides/ThreatLevels.json
        iv. /opt/superna/eca/data/audit-nfs/auto.nfs

f. **NOTE: This procedure assumes the IP addresses will stay the same so that the cluster NFS export does nont need to be changed and no impact to any firewall rules.**

2. Deploy a new OVA ECA cluster using the latest OS OVA by downloading from instructions here.
3. Follow deployment instructions in the install guide to deploy the new OVA and use the ECA cluster name captured from the prerequisites when prompted during the installation process of the OVA. The install guide for deploying the OVA is here.
   a. **NOTE: Use the same ip addresses as the current ECA cluster**
4. Shutdown the old ECA cluster
   a. Login to node 1 as ecaadmin
   b. **ecactl cluster down**
   c. wait for the shutdown to finish
   d. Using vcenter UI power off the vapp
5. Startup new ECA cluster
   a. Using vcenter UI power on the vapp
   b. ping each ip address in the cluster until each VM responds. **NOTE: Do not continue if you cannot ping each VM in the cluster**.
   c. Using winscp and login ecaadmin copy the files from the steps above into the new ECA OVA cluster
   d. On node 1 replace these files with the backup copies
      i. /opt/superna/eca/eca-env-common.conf

ii. /opt/superna/eca/docker-compose.overrides.yml

iii. /opt/superna/eca/conf/common/overrides/ThreatLevels.json

iv. /opt/superna/eca/data/audit-nfs/auto.nfs

e. On Nodes 1 to node X (where x is the last node in the cluster)

i. On each node complete these steps

ii. ssh ecaadmin@x.x.x.x (ip of each eca node)

1. sudo -s (enter ecaadmin password when prompted)
2. mkdir -p /opt/superna/mnt/audit/cluster GUID/cluster name/
    a. example only /opt/superna/mnt/audit/0050569960fcd70161594d21dd22a3c10cbe/prod-cluster-8
3. Repeat for each cluster managed by this ECA cluster. View the contents of this file to get the cluster GUID and name **/opt/superna/eca/data/audit-nfs/auto.nfs**

iii. Restart the Autofs process to read the auto.nfs file and mount all the clusters

1. **ecactl cluster exec "sudo systemctl restart autofs"**
2. Test the mount worked on each node

3. **ecactl cluster exec "mount"**
4. Verify the mount is present on all nodes in the output from the mount command.

f. Startup the new ECA cluster

i. login to eca node 1 as ecaadmin

ii. **ecactl cluster up**

iii. review statup messages for errors.

g. Done.

Monitor Health and Performance of an ECA Cluster (Optional)

The sections below cover how to check the vitals of an ECA cluster. Always check with support before taking any actions. ECA clusters are designed to consume CPU for most operations and is expected to see high CPU on all nodes most of the time.

How to Verifying ECA Cluster Status

1. **On the master node run these commands**:

- run the following command: **ecactl db shell**
- Once in the shell execute command: **status**
- Output should show 1 active master , 2 backup master server

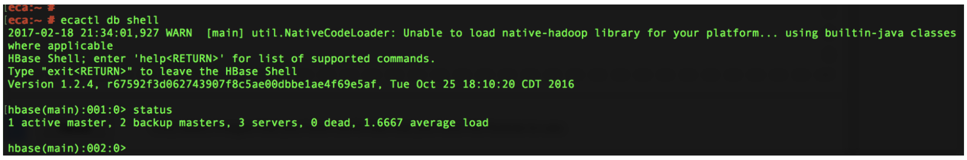

```
eca:~ #
eca:~ # ecactl db shell
2017-02-18 21:34:01,927 WARN  [main] util.NativeCodeLoader: Unable to load native-hadoop library for your platform... using builtin-java classes
where applicable
HBase Shell; enter 'help<RETURN>' for list of supported commands.
Type "exit<RETURN>" to leave the HBase Shell
Version 1.2.4, r67592f3d062743907f8c5ae00dbbe1ae4f69e5af, Tue Oct 25 18:10:20 CDT 2016

hbase(main):001:0> status
1 active master, 2 backup masters, 3 servers, 0 dead, 1.6667 average load

hbase(main):002:0>
```

How to Verify ECA containers are running

1. Command: “ecactl containers ps”

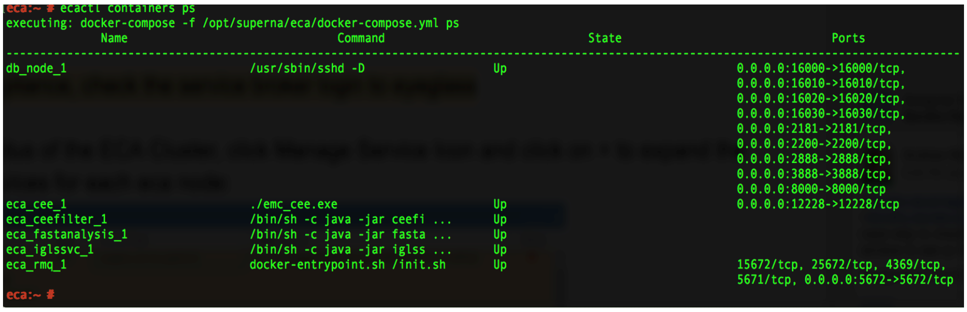

```
eca:~ # ecactl containers ps
executing: docker-compose -f /opt/superna/eca/docker-compose.yml ps
              Name                                Command                State                                Ports
-----------------------------------------------------------------------------------------------------------------------------------
db_node_1                            /usr/sbin/sshd -D                    Up                                   0.0.0.0:16000->16000/tcp,
                                                                                                               0.0.0.0:16010->16010/tcp,
                                                                                                               0.0.0.0:16020->16020/tcp,
                                                                                                               0.0.0.0:16030->16030/tcp,
                                                                                                               0.0.0.0:2181->2181/tcp,
                                                                                                               0.0.0.0:2200->2200/tcp,
                                                                                                               0.0.0.0:2888->2888/tcp,
                                                                                                               0.0.0.0:3888->3888/tcp,
                                                                                                               0.0.0.0:8000->8000/tcp
eca_cee_1                            ./emc_cee.exe                        Up                                   0.0.0.0:12228->12228/tcp
eca_ceefilter_1                      /bin/sh -c java -jar ceefi ...       Up
eca_fastanalysis_1                   /bin/sh -c java -jar fasta ...       Up
eca_iglssvc_1                        /bin/sh -c java -jar iglss ...       Up
eca_rmq_1                            docker-entrypoint.sh /init.sh        Up                                   15672/tcp, 25672/tcp, 4369/tcp,
                                                                                                               5671/tcp, 0.0.0.0:5672->5672/tcp

eca:~ #
```

Check cluster status and that all analytics tables exist (Ransomware Defender, Easy Auditor, Performance Auditor) (Optional Step)

- Command: ‘ecactl cluster status’
- This command verifies all containers are running on all nodes and verifies each node can mount the tables in the Analytics database.
- **If any error conditions open a support case to resolve or retry with steps below:**
- ecactl cluster down
- ecactl cluster up
- Send ECA cluster startup text to support

How to Check ECA node Container CPU and Memory Usage (Optional)

1. Login to the eca node as ecaadmin
2. Type cli command to see real time view of container resources utilization
3. **ecactl stats**

How to Enable Real-time Monitor ECA cluster performance (If directed by support)

Use this procedure to enable container monitor to determine if cpu GHZ are set correctly for query and writing to PowerScale performance.

1. To enable cadvisor, add the following line to eca-env-common.conf:
2. export LAUNCH_MONITORING=true
3. This will launch cadvisor on all cluster nodes.
4. If you want to launch it on a single node, login to that node and execute:
5. ecactl containers up -d cadvisor
6. Once the cadvisor service is running, login to http://<IP OF ECA NODE>:9080 to see the web UI.
7. Done.

ECA Cluster Modification Procedures (optional)

How to expand Easy Auditor cluster size

NOTE: Support will determine if your installation requires expansion. Always contact support.

Follow this steps to add 3 or 6 more VM's for analytics performance increase for higher event rate or long running queries on a large database. Deploy the ECA OVA again, copy the new VM's into the vAPP, remove the vAPP created during the deployment. **NOTE: The**

ECA name will not matter during the new OVA deployment since it will be synced from the existing ECA cluster during cluster up procedures.

1. Login to the master ECA node
2. **ecactl cluster down**
3. deploy one or two more eca clusters. No special config needs to be added on the newly deployed ECA OVA.
4. nano /opt/superna/eca/eca-env-common.conf to add more nodes:
5. ECA_LOCATION_NODE_4: <IP>
6. ECA_LOCATION_NODE_5: >IP>
7. add anywhere from nodes 4 to 9 depending on the number of VM's added to the cluster.
8. **ecactl components configure-nodes**
9. **ecactl cluster up**
10. This will expand HBASE and Spark containers for faster read and analytics performance
11. Login to eyeglass and open managed services
    a. 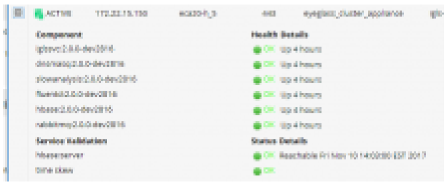

ACTIVE 172.22.15.150 eca20-fi_3 443 eyeglass_cluster_appliance igls-	
Component	**Health Details**
iglssvc:2.0.0-dev2816	● OK Up 4 hours
[illegible]	● OK Up 4 hours
[illegible]	● OK Up 4 hours
fluentd:2.0.0-dev2816	● OK Up 4 hours
hbase:2.0.0-dev2816	● OK Up 4 hours
rabbitmq:2.0.0-dev2816	● OK Up 4 hours
Service Validation	**Status Details**
hbaseserver	● OK Reachable Fri Nov 10 14:00:00 EST 2017
time skew	● OK

12. Now HBASE needs to balance the load across the cluster for improved read performance.
    a. Now login to the Region Master vm typically node 1
    b. http://x.x.x.x:16010/ verify that each region server (6 total) are visible in the UI
    c. Verify each has assigned regions
    d. Verify requests are visible to each region server
    e. Verify tables section shows no regions offline, no regions in other column,

f. Example screenshots of 6 region servers with regions and normal looking table view

g. 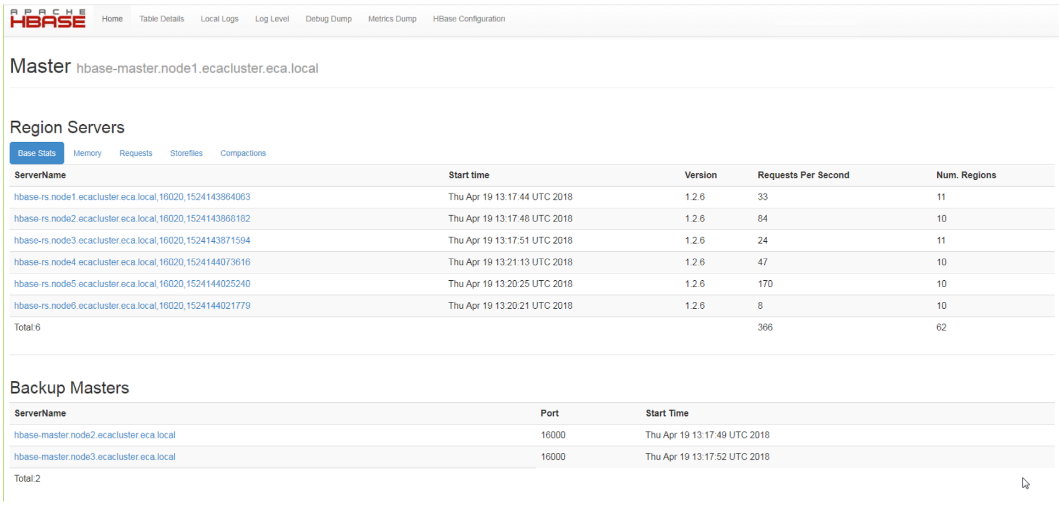

Master hbase-master.node1.ecacluster.eca.local

Region Servers

ServerName	Start time	Version	Requests Per Second	Num. Regions
hbase-rs.node1.ecacluster.eca.local,16020,1524143864063	Thu Apr 19 13:17:44 UTC 2018	1.2.6	33	11
hbase-rs.node2.ecacluster.eca.local,16020,1524143868182	Thu Apr 19 13:17:48 UTC 2018	1.2.6	84	10
hbase-rs.node3.ecacluster.eca.local,16020,1524143871594	Thu Apr 19 13:17:51 UTC 2018	1.2.6	24	11
hbase-rs.node4.ecacluster.eca.local,16020,1524144073616	Thu Apr 19 13:21:13 UTC 2018	1.2.6	47	10
hbase-rs.node5.ecacluster.eca.local,16020,1524144025240	Thu Apr 19 13:20:25 UTC 2018	1.2.6	170	10
hbase-rs.node6.ecacluster.eca.local,16020,1524144021779	Thu Apr 19 13:20:21 UTC 2018	1.2.6	8	10
Total:6			366	62

Backup Masters

ServerName	Port	Start Time
hbase-master.node2.ecacluster.eca.local	16000	Thu Apr 19 13:17:49 UTC 2018
hbase-master.node3.ecacluster.eca.local	16000	Thu Apr 19 13:17:52 UTC 2018
Total:2		

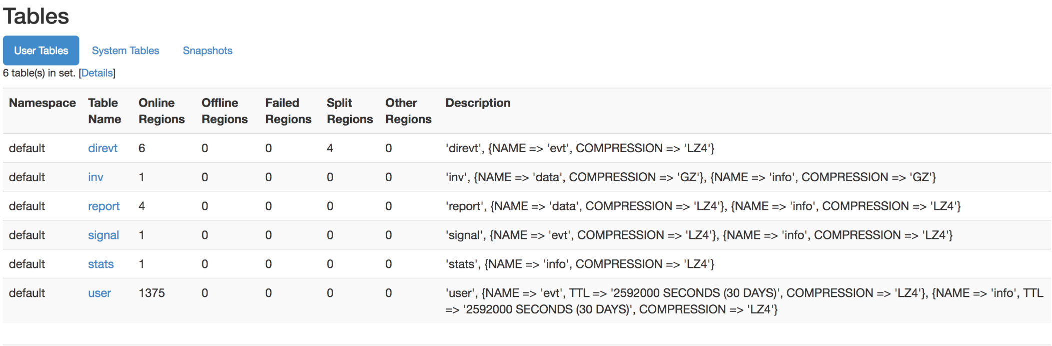

Tables

6 table(s) in set. [Details]

Namespace	Table Name	Online Regions	Offline Regions	Failed Regions	Split Regions	Other Regions	Description
default	direvt	6	0	0	4	0	'direvt', {NAME => 'evt', COMPRESSION => 'LZ4'}
default	inv	1	0	0	0	0	'inv', {NAME => 'data', COMPRESSION => 'GZ'}, {NAME => 'info', COMPRESSION => 'GZ'}
default	report	4	0	0	0	0	'report', {NAME => 'data', COMPRESSION => 'LZ4'}, {NAME => 'info', COMPRESSION => 'LZ4'}
default	signal	1	0	0	0	0	'signal', {NAME => 'evt', COMPRESSION => 'LZ4'}, {NAME => 'info', COMPRESSION => 'LZ4'}
default	stats	1	0	0	0	0	'stats', {NAME => 'info', COMPRESSION => 'LZ4'}
default	user	1375	0	0	0	0	'user', {NAME => 'evt', TTL => '2592000 SECONDS (30 DAYS)', COMPRESSION => 'LZ4'}, {NAME => 'info', TTL => '2592000 SECONDS (30 DAYS)', COMPRESSION => 'LZ4'}

h.

13. done.

Advanced Configurations (Optional)

How to Configure a Ransomware Defender Only Configuration (Skip if running multiple products)

Make this change before starting up the cluster to ensure docker containers that are not required are auto disabled on startup.

1. Login to node 1 over ssh as ecaadmin
2. nano /opt/superna/eca/eca-env-common.conf
3. add a variable

    a. export RSW_ONLY_CFG=true

4. :wq (to save the file)
5. Continue with startup steps below

6.1. Eyeglass Clustered Agent ECA Hyper-V Installation Guide

Home Top

- Read first
- Create ECA Hyper-V Virtual Machine
- Configure ECA data disk
- Configuration of ECA cluster

Read first

1. ECA appliance uses 2 disks. 1 for OS and 1 for data
2. OS disk requires **20 GB** [default disk]
3. Data disk requires **80 GB** [read below on how to create]
4. ECA install and upgrade GUIDE HERE
5. ECA admin guide HERE

Create ECA Hyper-V Virtual Machine

1. Download ECA Hyper-V vhdx file from https://support.superna.net portal
2. Deploy a new `Virtual Machine`

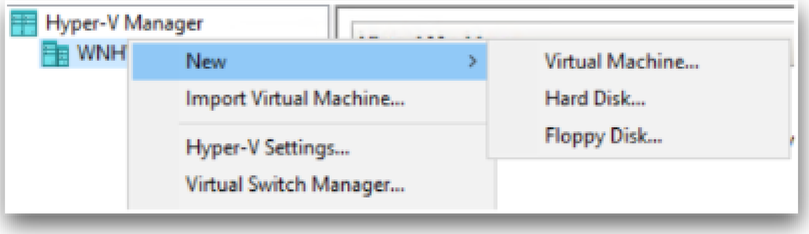

3. Enter `Name` for the VM

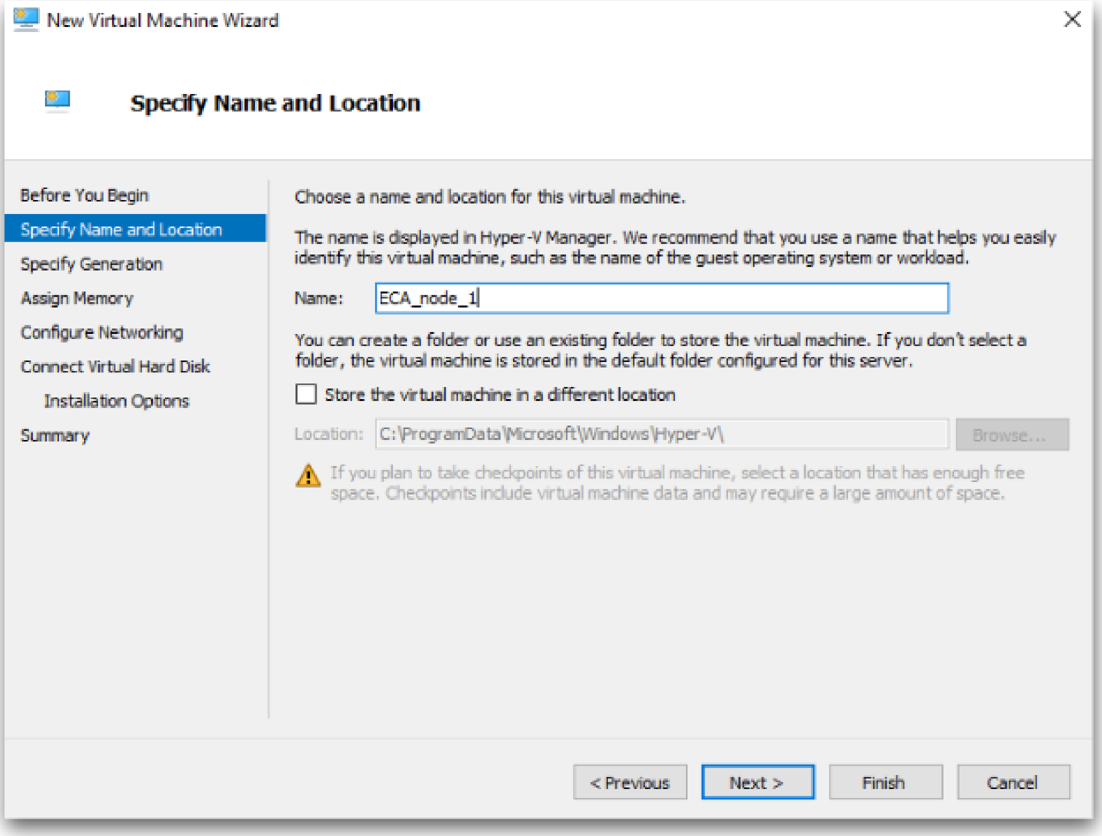
New Virtual Machine Wizard
Specify Name and Location
Before You Begin
Specify Name and Location
Specify Generation
Assign Memory
Configure Networking
Connect Virtual Hard Disk
Installation Options
Summary

Choose a name and location for this virtual machine.

The name is displayed in Hyper-V Manager. We recommend that you use a name that helps you easily identify this virtual machine, such as the name of the guest operating system or workload.

Name: ECA_node_1

You can create a folder or use an existing folder to store the virtual machine. If you don't select a folder, the virtual machine is stored in the default folder configured for this server.

□ Store the virtual machine in a different location

Location: C:\ProgramData\Microsoft\Windows\Hyper-V\

If you plan to take checkpoints of this virtual machine, select a location that has enough free space. Checkpoints include virtual machine data and may require a large amount of space.

4. Check `Generation 1`

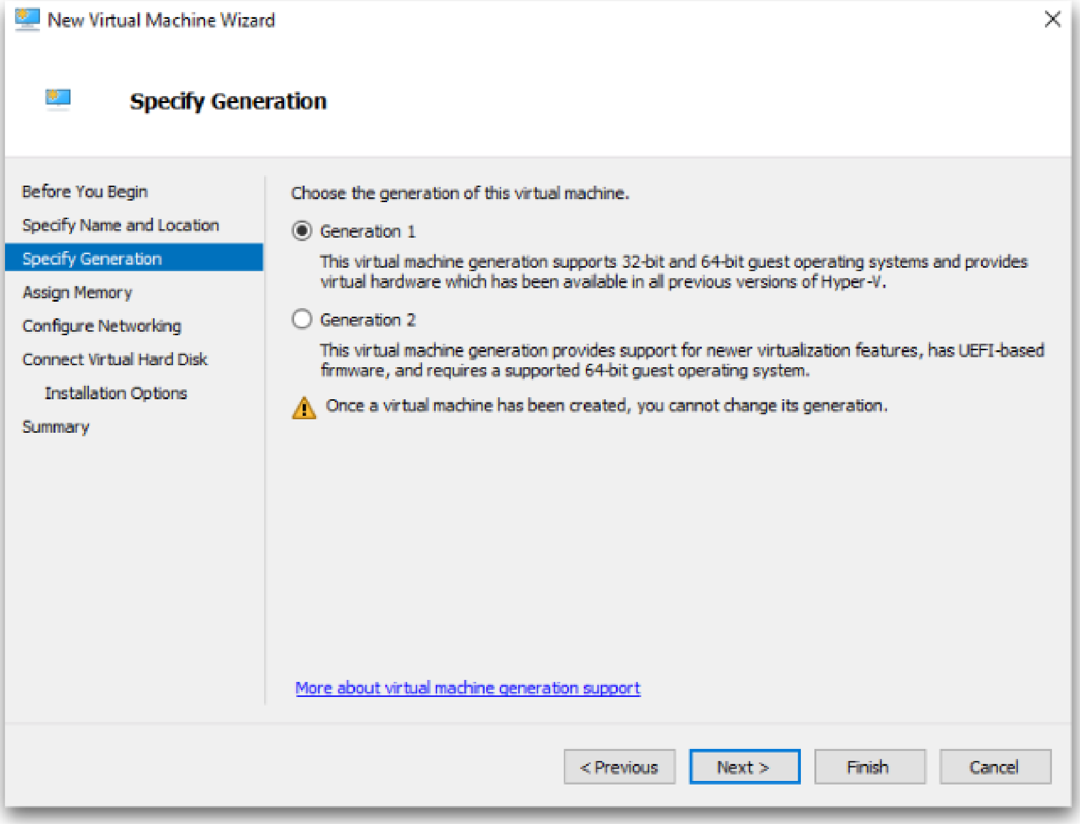

Choose the generation of this virtual machine.

● Generation 1

This virtual machine generation supports 32-bit and 64-bit guest operating systems and provides virtual hardware which has been available in all previous versions of Hyper-V.

○ Generation 2

This virtual machine generation provides support for newer virtualization features, has UEFI-based firmware, and requires a supported 64-bit guest operating system.

Once a virtual machine has been created, you cannot change its generation.

More about virtual machine generation support

5. Startup memory **16384 MB [16 GB]**

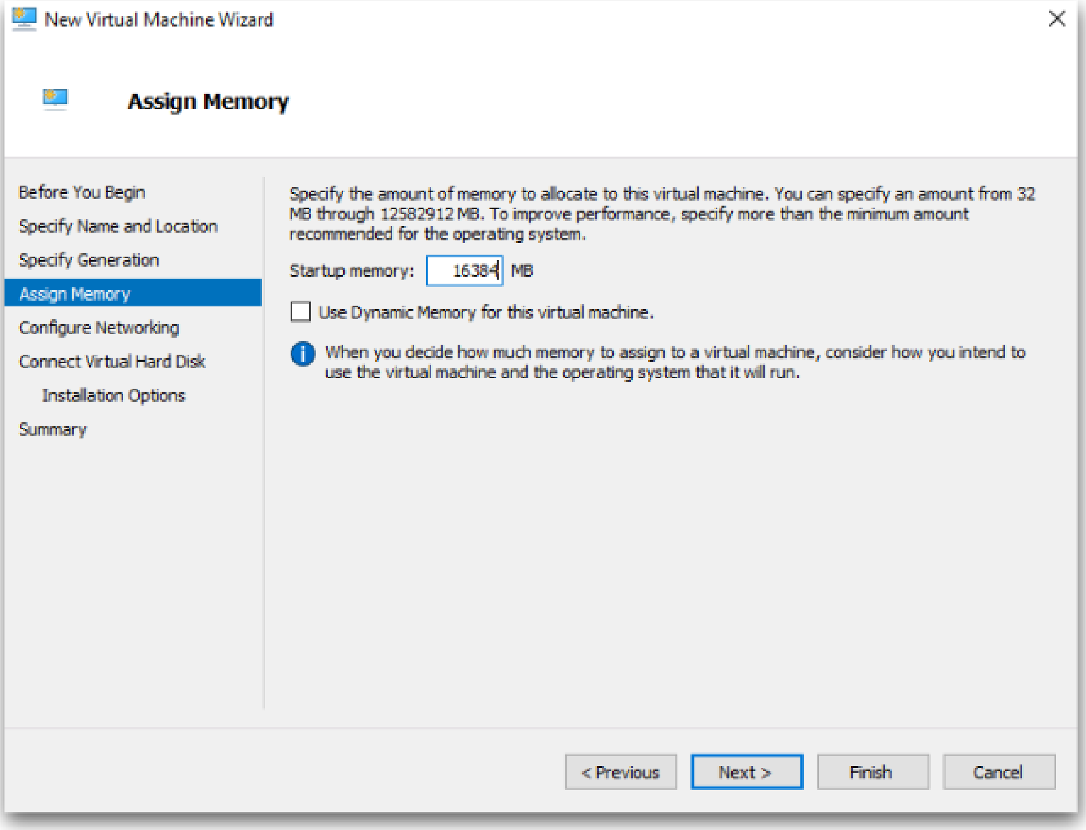

New Virtual Machine Wizard

Assign Memory

Specify the amount of memory to allocate to this virtual machine. You can specify an amount from 32 MB through 12582912 MB. To improve performance, specify more than the minimum amount recommended for the operating system.

Startup memory: 16384 MB

□ Use Dynamic Memory for this virtual machine.

When you decide how much memory to assign to a virtual machine, consider how you intend to use the virtual machine and the operating system that it will run.

6. Select `Network Adapter`

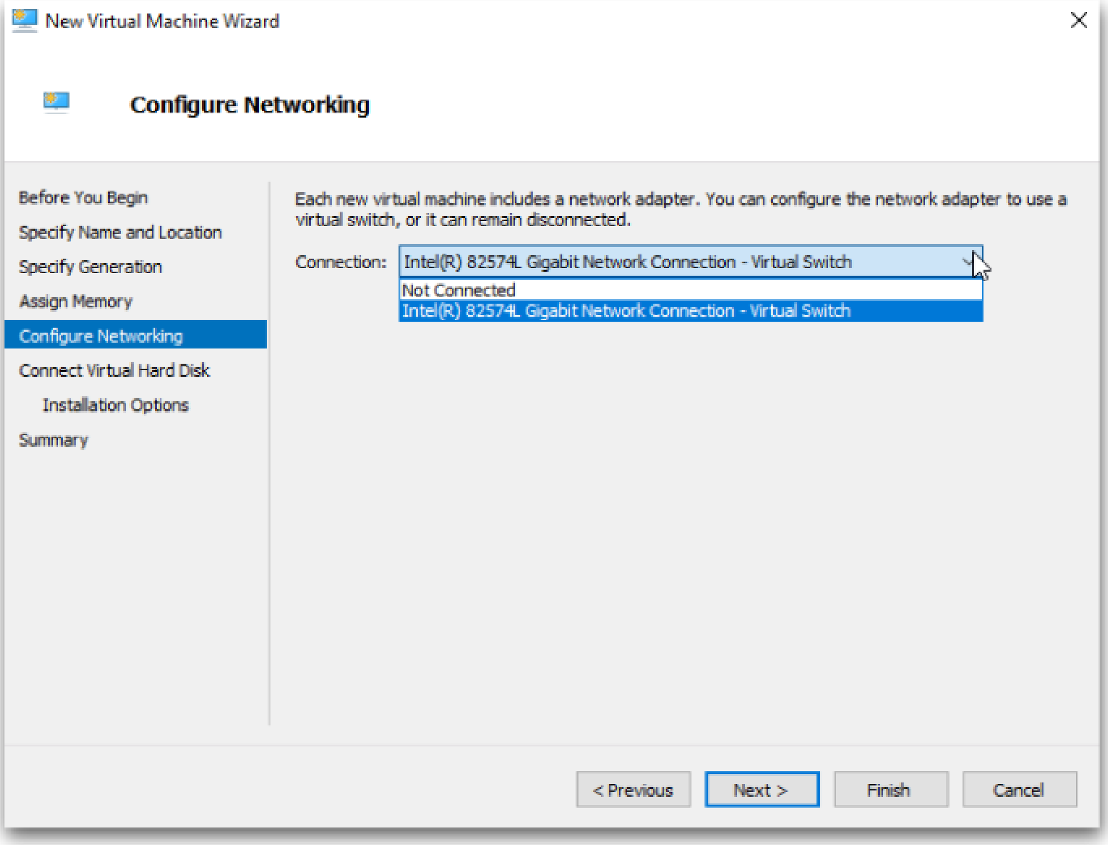

New Virtual Machine Wizard

Configure Networking

Each new virtual machine includes a network adapter. You can configure the network adapter to use a virtual switch, or it can remain disconnected.

Connection: Intel(R) 82574L Gigabit Network Connection - Virtual Switch

7. Use an existing `Virtual Hard disk` → Browse to newly downloaded ECA `vhdx` file

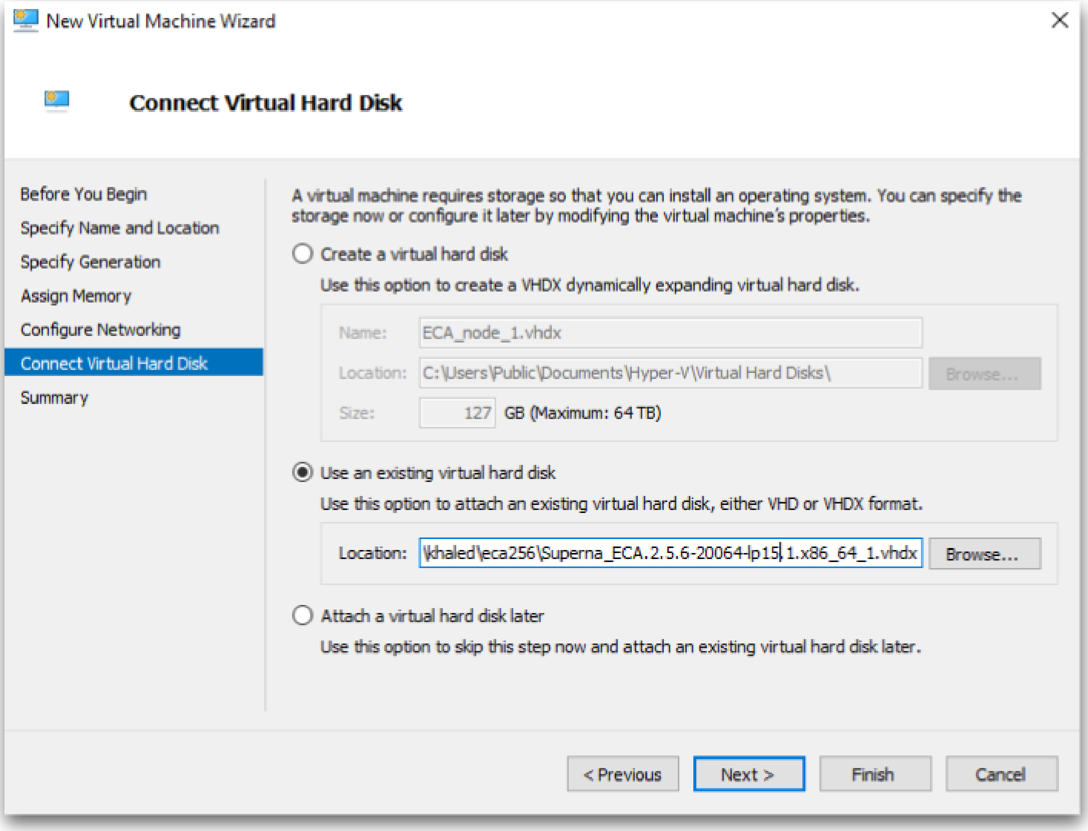

New Virtual Machine Wizard

Connect Virtual Hard Disk

A virtual machine requires storage so that you can install an operating system. You can specify the storage now or configure it later by modifying the virtual machine's properties.

○ Create a virtual hard disk

Use this option to create a VHDX dynamically expanding virtual hard disk.

Name: ECA_node_1.vhdx

Location: C:\Users\Public\Documents\Hyper-V\Virtual Hard Disks\

Size: 127 GB (Maximum: 64 TB)

● Use an existing virtual hard disk

Use this option to attach an existing virtual hard disk, either VHD or VHDX format.

Location: \khaled\eca256\Superna_ECA.2.5.6-20064-lp151.x86_64_1.vhdx

○ Attach a virtual hard disk later

Use this option to skip this step now and attach an existing virtual hard disk later.

8. Complete the Wizard

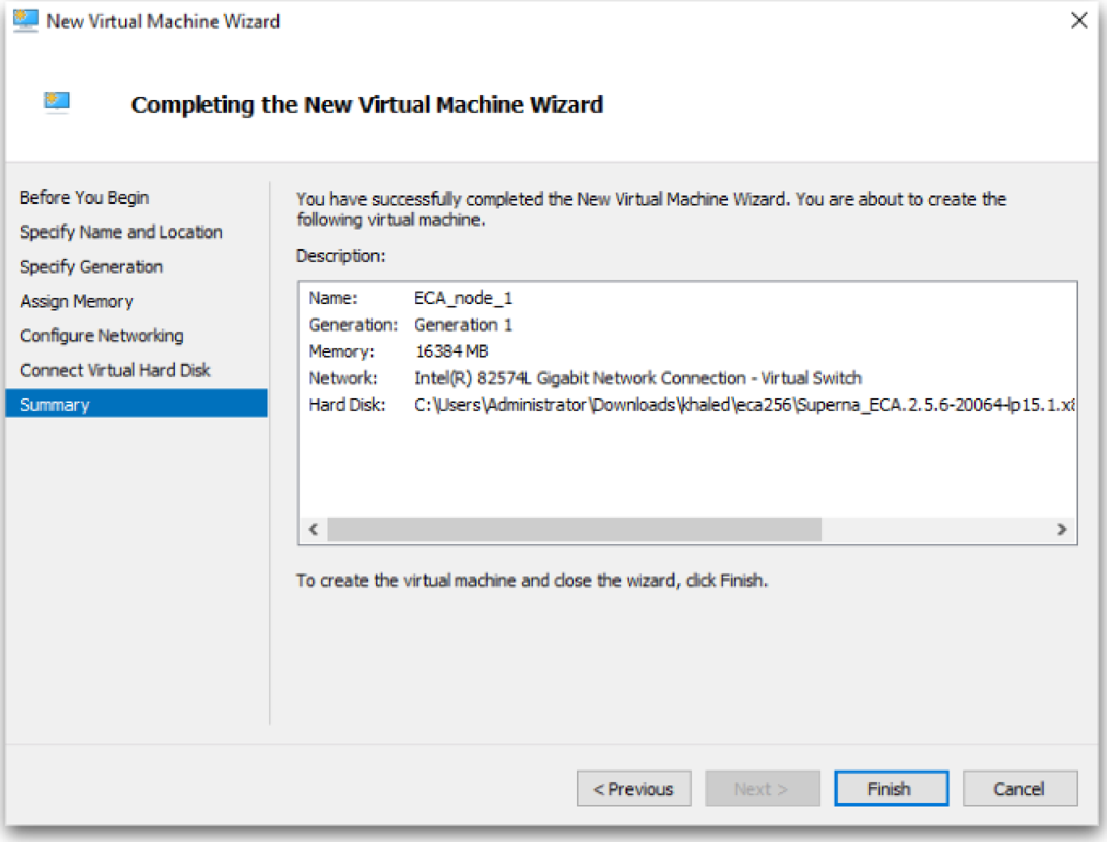

Completing the New Virtual Machine Wizard

You have successfully completed the New Virtual Machine Wizard. You are about to create the following virtual machine.

Description:

Name: ECA_node_1
Generation: Generation 1
Memory: 16384 MB
Network: Intel(R) 82574L Gigabit Network Connection - Virtual Switch
Hard Disk: C:\Users\Administrator\Downloads\khaled\eca256\Superna_ECA.2.5.6-20064-lp15.1.x

To create the virtual machine and close the wizard, click Finish.

Configure ECA data disk

1. After deploying, go to the new VM → Right Click → Settings

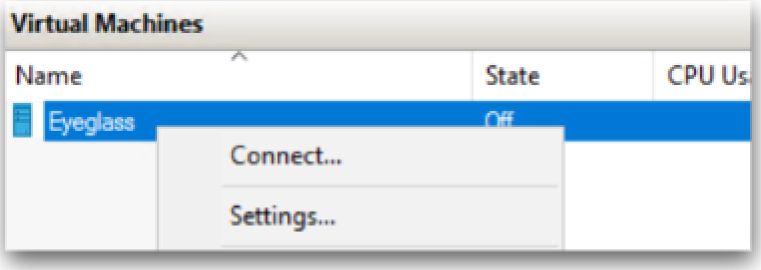

2. From `IDE Controller 0` → Add a `Hard Drive`

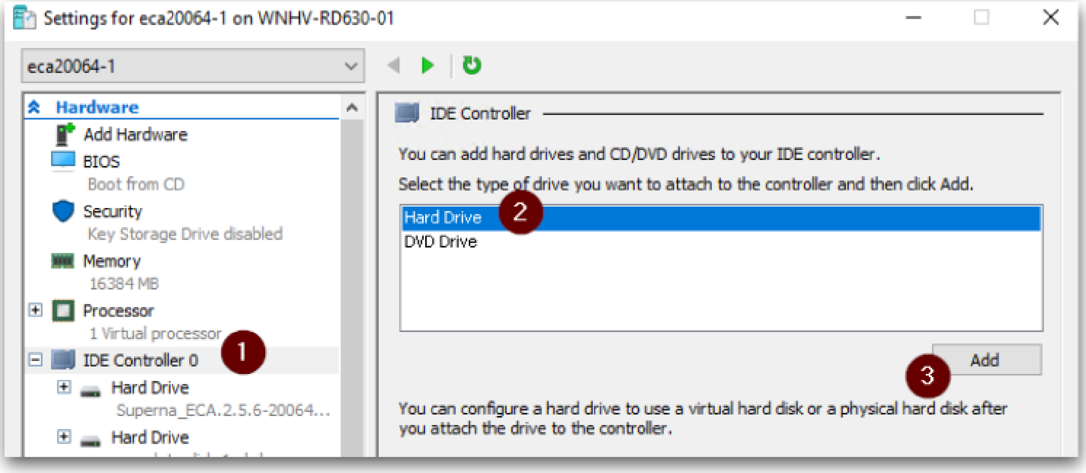

Settings for eca20064-1 on WNHV-RD630-01

eca20064-1

IDE Controller

You can add hard drives and CD/DVD drives to your IDE controller.

Select the type of drive you want to attach to the controller and then click Add.

Hard Drive
DVD Drive

You can configure a hard drive to use a virtual hard disk or a physical hard disk after you attach the drive to the controller.

3. Create New

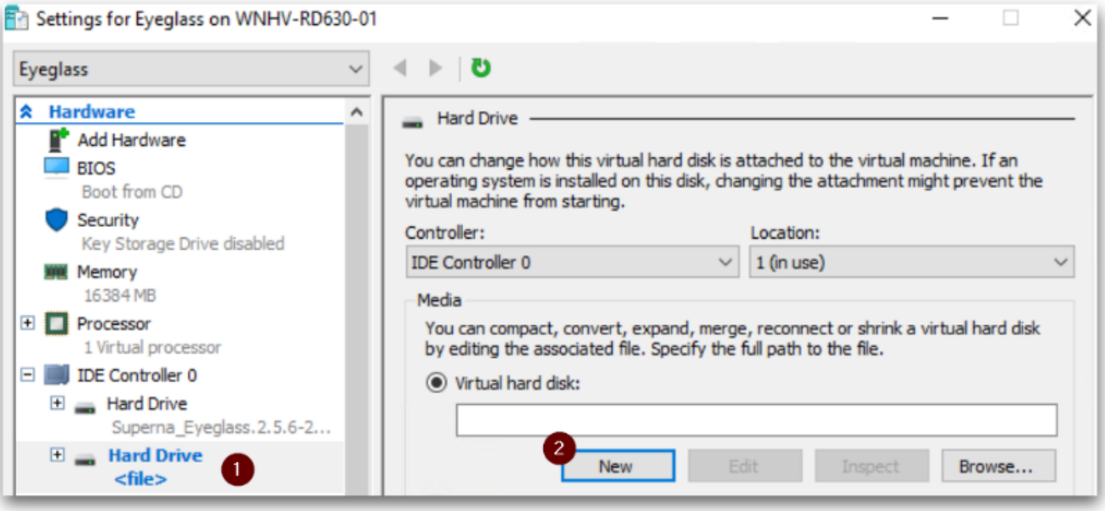

Settings for Eyeglass on WNHV-RD630-01

Eyeglass

Hard Drive

You can change how this virtual hard disk is attached to the virtual machine. If an operating system is installed on this disk, changing the attachment might prevent the virtual machine from starting.

Controller: IDE Controller 0

Location: 1 (in use)

Media

You can compact, convert, expand, merge, reconnect or shrink a virtual hard disk by editing the associated file. Specify the full path to the file.

● Virtual hard disk:

4. Choose Disk Format → `VHDX`

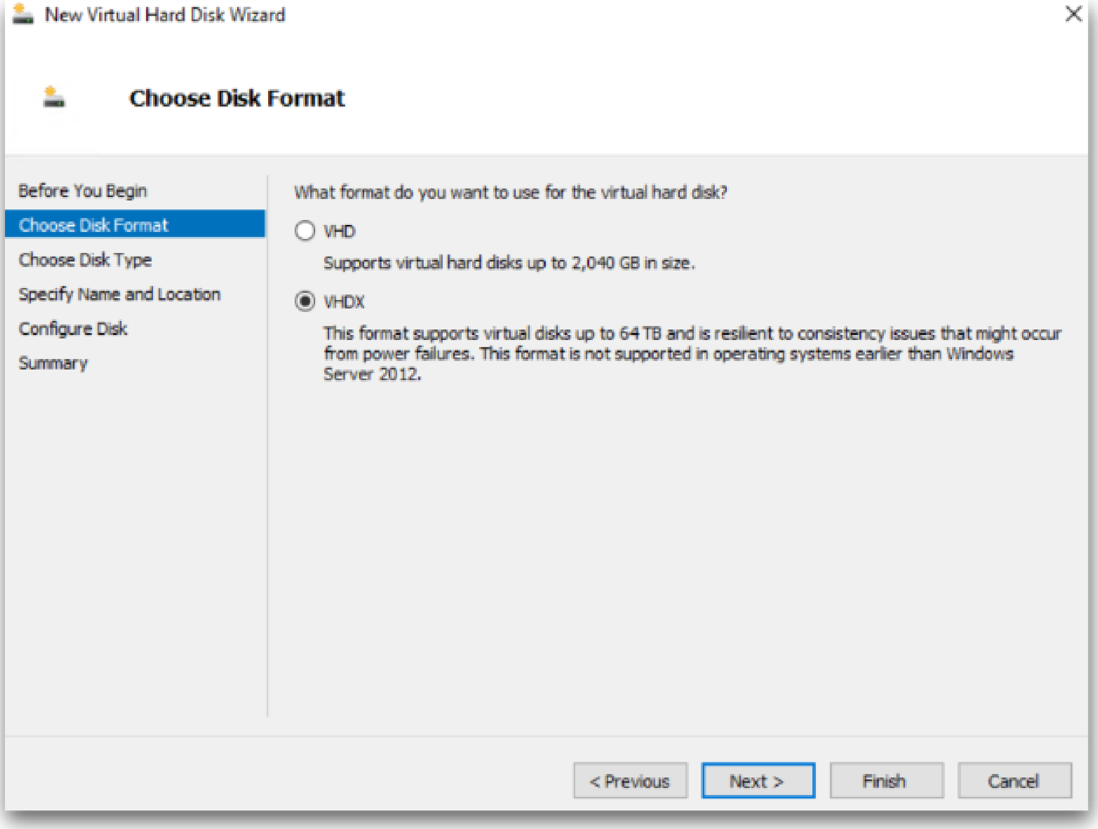

New Virtual Hard Disk Wizard

Choose Disk Format

What format do you want to use for the virtual hard disk?

○ VHD

Supports virtual hard disks up to 2,040 GB in size.

● VHDX

This format supports virtual disks up to 64 TB and is resilient to consistency issues that might occur from power failures. This format is not supported in operating systems earlier than Windows Server 2012.

5. Choose Disk Type → `Fixed size`

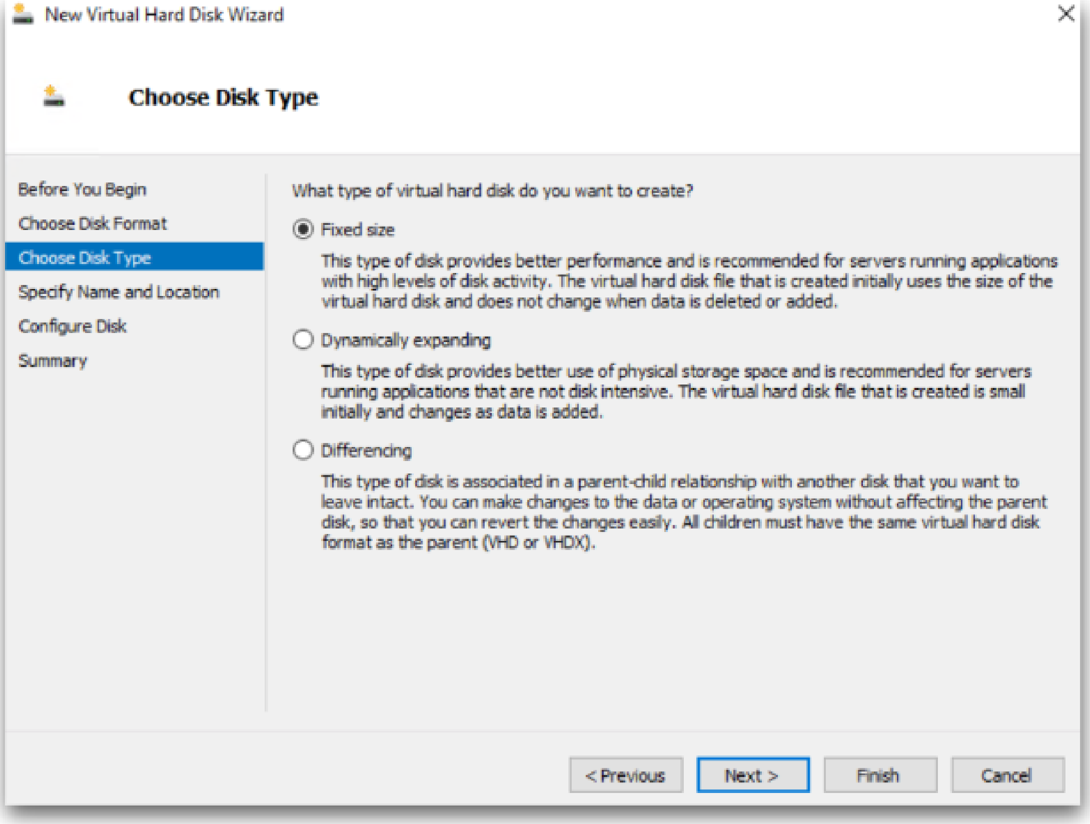

New Virtual Hard Disk Wizard

Choose Disk Type

What type of virtual hard disk do you want to create?

● Fixed size

This type of disk provides better performance and is recommended for servers running applications with high levels of disk activity. The virtual hard disk file that is created initially uses the size of the virtual hard disk and does not change when data is deleted or added.

○ Dynamically expanding

This type of disk provides better use of physical storage space and is recommended for servers running applications that are not disk intensive. The virtual hard disk file that is created is small initially and changes as data is added.

○ Differencing

This type of disk is associated in a parent-child relationship with another disk that you want to leave intact. You can make changes to the data or operating system without affecting the parent disk, so that you can revert the changes easily. All children must have the same virtual hard disk format as the parent (VHD or VHDX).

6. Name the data disk

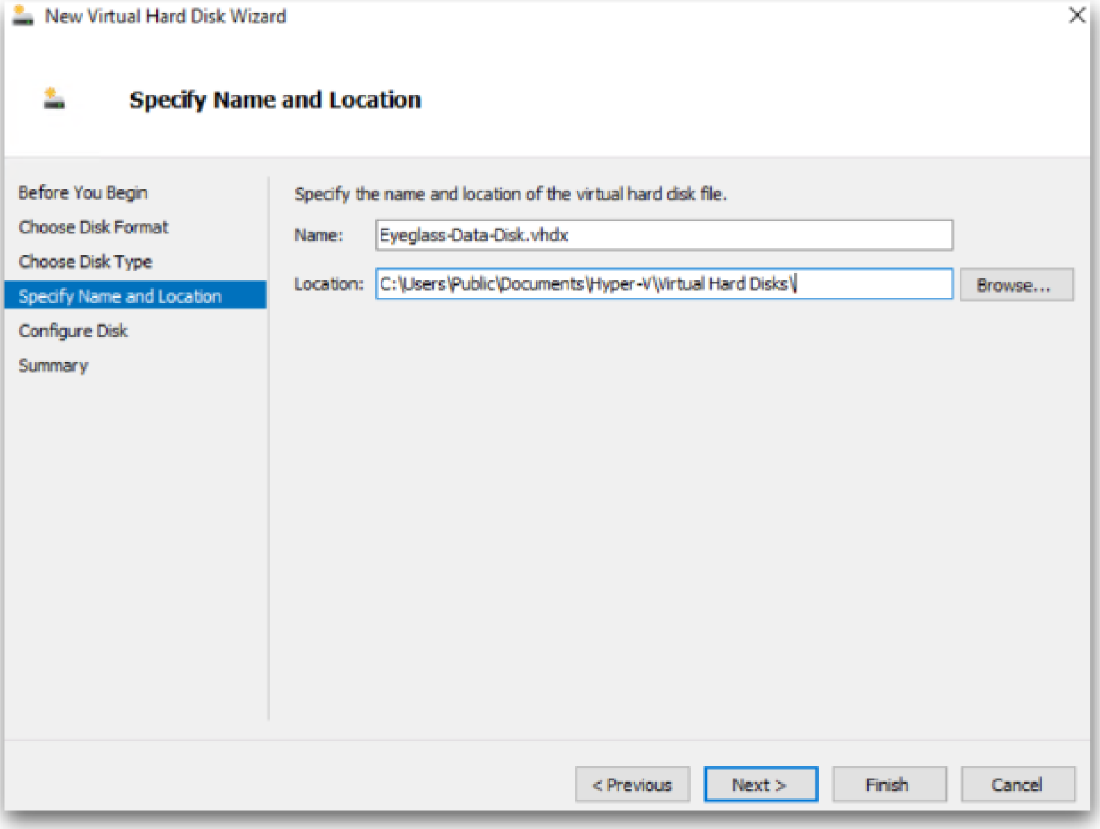

Specify Name and Location

Specify the name and location of the virtual hard disk file.

Name: Eyeglass-Data-Disk.vhdx

Location: C:\Users\Public\Documents\Hyper-V\Virtual Hard Disks\

7. Create a new blank virtual hard disk : **80 GB**

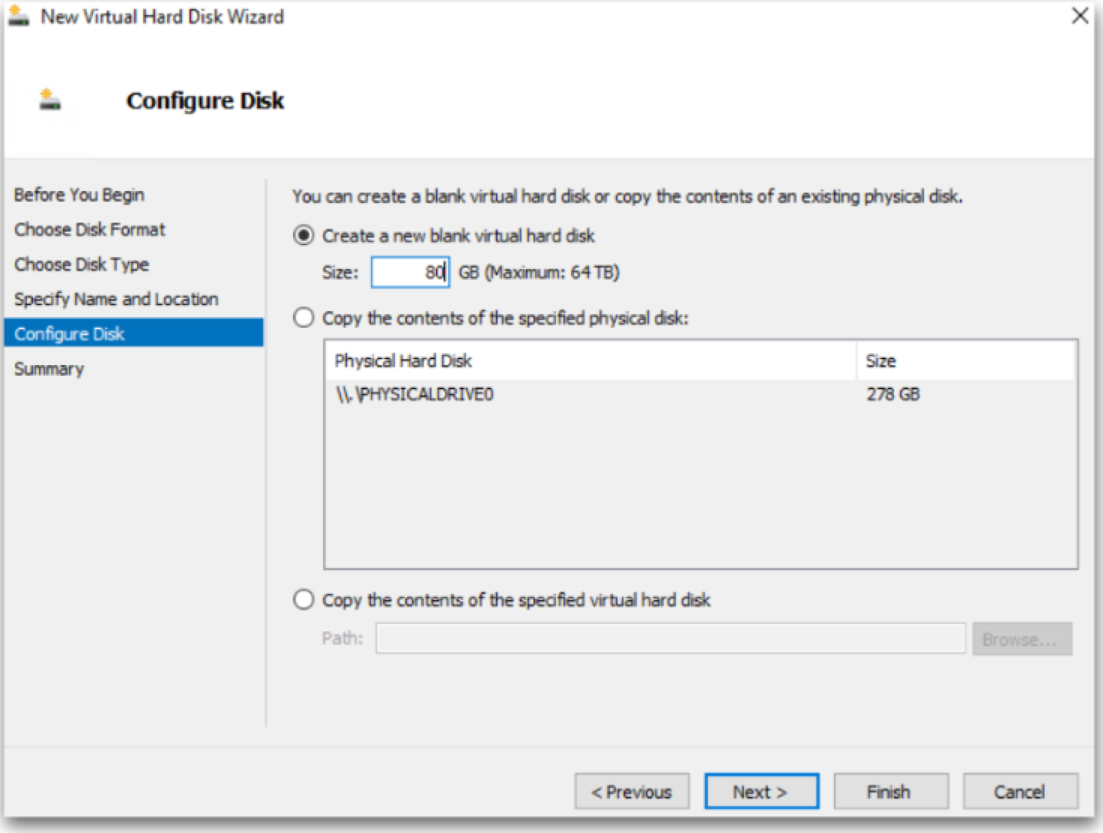

Configure Disk

You can create a blank virtual hard disk or copy the contents of an existing physical disk.

● Create a new blank virtual hard disk

Size: 80 GB (Maximum: 64 TB)

○ Copy the contents of the specified physical disk:

Physical Hard Disk	Size
\\.\PHYSICALDRIVE0	278 GB

○ Copy the contents of the specified virtual hard disk

Path:

8. Complete the data disk Wizard

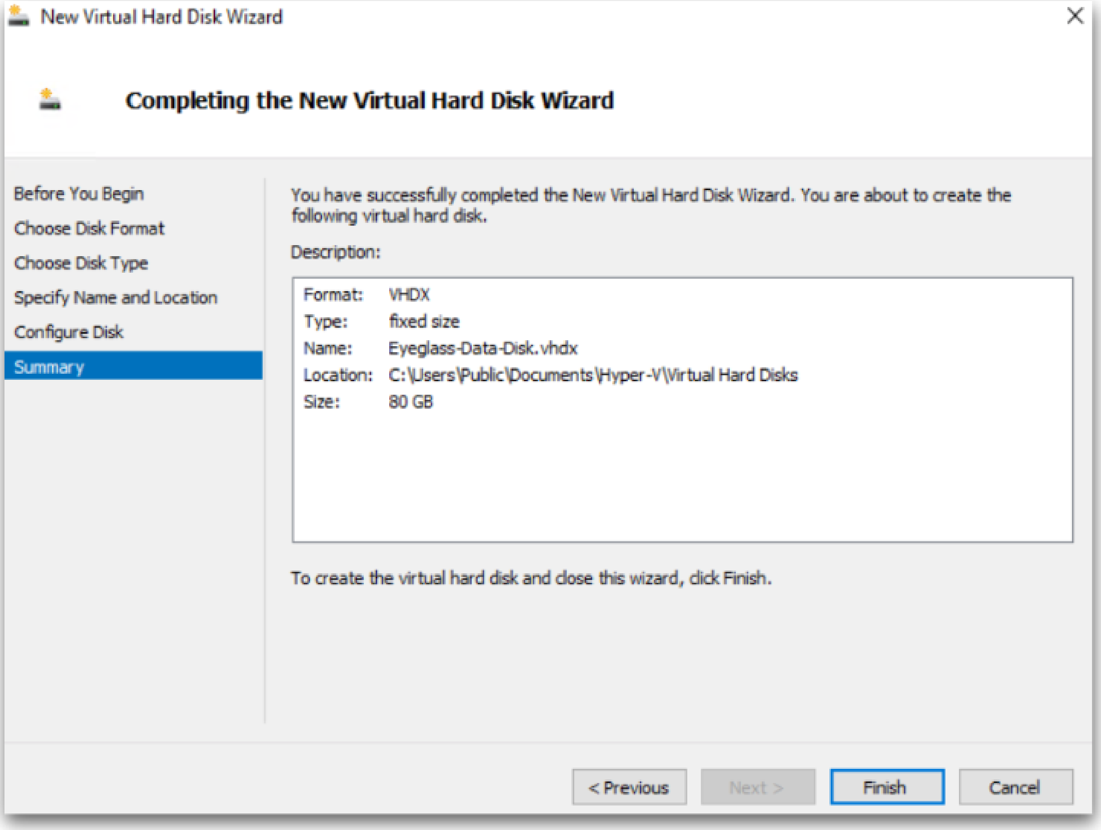

You have successfully completed the New Virtual Hard Disk Wizard. You are about to create the following virtual hard disk.

Description:

Format: VHDX
Type: fixed size
Name: Eyeglass-Data-Disk.vhdx
Location: C:\Users\Public\Documents\Hyper-V\Virtual Hard Disks
Size: 80 GB

To create the virtual hard disk and close this wizard, click Finish.

Configuration of ECA cluster

SSH username: ecadmin
SSH password: 3y3gl4ss

1. Power up the VM and wait 5-10 minutes to allow Superna on-boot script to run. tail the superna-on-boot.log and wait for it to finish. Then follow the on-screen instruction

```
tail -2 /var/log/superna-on-boot.log
```

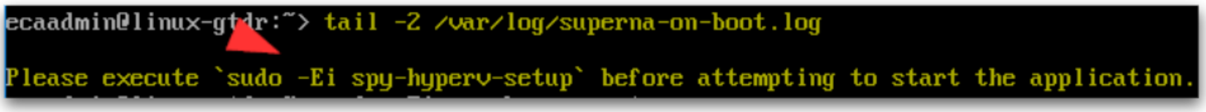

```
ecaadmin@linux-gtdr:~> tail -2 /var/log/superna-on-boot.log

Please execute `sudo -Ei spy-hyperv-setup` before attempting to start the application.
```

2. Run the command to setup your ECA Hyper-V node 1

```
sudo spy-hyperv-setup
```

When prompted, enter `admin` user password [default password is: 3y3gl4ss] and enter IP, Netmask, Gateway, Hostname, DNS, NTP info

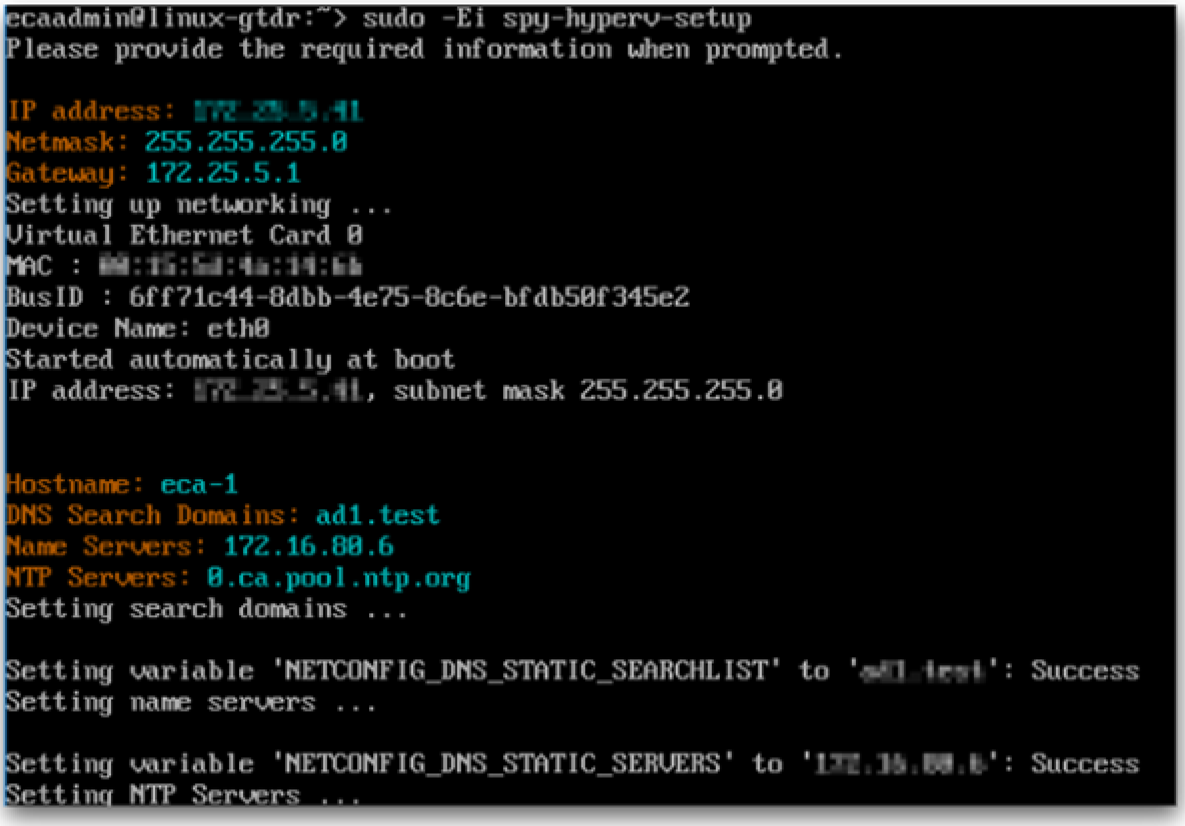

```
ecaadmin@linux-gtdr:~> sudo -Ei spy-hyperv-setup
Please provide the required information when prompted.

IP address: [redacted]
Netmask: 255.255.255.0
Gateway: 172.25.5.1
Setting up networking ...
Virtual Ethernet Card 0
MAC : [redacted]
BusID : 6ff71c44-8dbb-4e75-8c6e-bfdb50f345e2
Device Name: eth0
Started automatically at boot
IP address: [redacted], subnet mask 255.255.255.0


Hostname: eca-1
DNS Search Domains: ad1.test
Name Servers: 172.16.80.6
NTP Servers: 0.ca.pool.ntp.org
Setting search domains ...

Setting variable 'NETCONFIG_DNS_STATIC_SEARCHLIST' to '[redacted]': Success
Setting name servers ...

Setting variable 'NETCONFIG_DNS_STATIC_SERVERS' to '[redacted]': Success
Setting NTP Servers ...
```

3. [**STOP: go to STEP#4 and STEP#5**] Node 1 is the ECA Master node. **Do not press `y` until node 2-N prompts are marked as `n`.** Go to the next step.

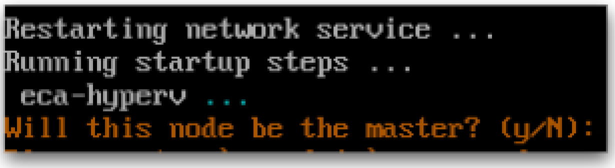

```
Restarting network service ...
Running startup steps ...
 eca-hyperv ...
Will this node be the master? (y/N):
```

4. Repeat STEP#1 and STEP#2 on node 2-N nodes where 'N' is the number of nodes you want to deploy for ECA

5. [**IMPORTANT**] On Node 2-N, when prompted for `master` node, Enter `n`

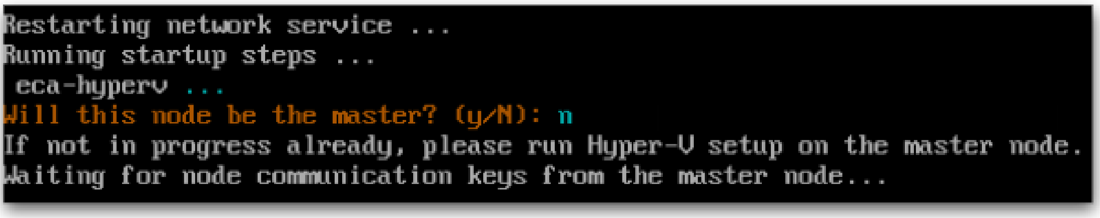

```
Restarting network service ...
Running startup steps ...
 eca-hyperv ...
Will this node be the master? (y/N): n
If not in progress already, please run Hyper-V setup on the master node.
Waiting for node communication keys from the master node...
```

6. **[IMPORTANT]** Go to ECA Master node 1 and press `y` to complete the master node setup
   - Enter ECA cluster name. Keep it simple.No Uppercase, no underscore, no special characters allowed.
   - Enter child nodes IP [space separated]

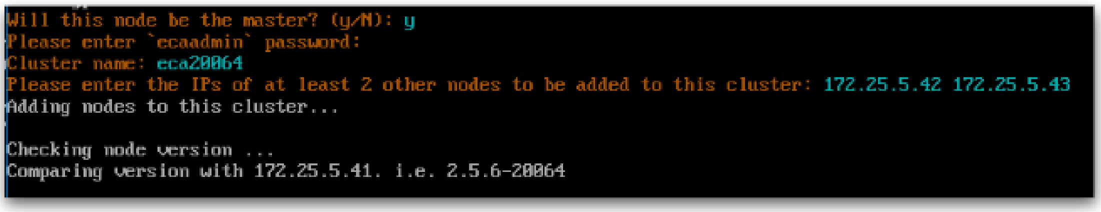

```
Will this node be the master? (y/N): y
Please enter `ecaadmin` password:
Cluster name: eca20064
Please enter the IPs of at least 2 other nodes to be added to this cluster: 172.25.5.42 172.25.5.43
Adding nodes to this cluster...

Checking node version ...
Comparing version with 172.25.5.41. i.e. 2.5.6-20064
```

7. When all done - you will see setup complete message [check all nodes]

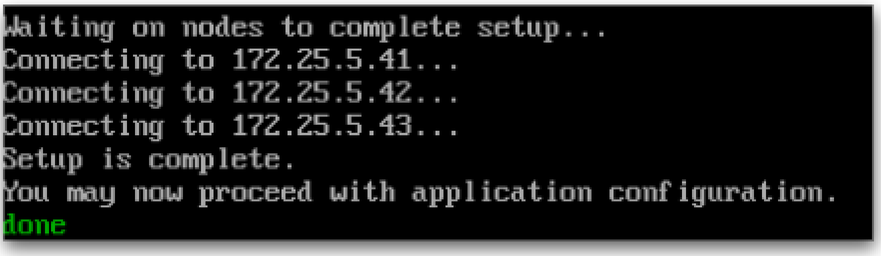

```
Waiting on nodes to complete setup...
Connecting to 172.25.5.41...
Connecting to 172.25.5.42...
Connecting to 172.25.5.43...
Setup is complete.
You may now proceed with application configuration.
done
```

6.2. Mini-ECA Installation and Configuration to Enable Distributed Cluster Mode

Home Top

Overview

New in release 2.5.5 is mini-eca deployment mode that enables centralized security processing of cluster audit data. This reduces the cost and simplifies deployments with a large central analytics and database ECA cluster and distributed single VM or VM pairs at remote locations that collect audit data locally and forward to the central cluster for processing.

Requirements

1. The latency between the main ECA cluster and the remote mini ECA's must be below a ping time of 80 ms. Anything over this may not be supported.
2. The FSTAB method of mounting the cluster audit folder is required (see instructions below)

Firewall for Mini ECA

Port	Direction	Comments
22 ssh	ECA main cluster <-> mini eca admin pc --> mini ECA	
2181, 2888 TCP	ECA main cluster <-> mini eca	
9092 , 9090 TCP	ECA main cluster <-> mini eca	
5514 (TCP) as of 2.5.6 build 84	mini ECA --> Eyeglass	
443 (TCP)	mini ECA --> Eyeglass admin pc --> mini eca	
NFS (UDP & TCP) version 3	mini ECA --> cluster	
NTP (UDP) 123	mini ECA --> NTP server	
DNS UDP 53	mini ECA --> DNS server	
TCP port 5000 for node 1 ECA (**during upgrades only**)	all ECA and mini ECA --> node 1 main ECA cluster IP	

Network Impact Calculation

1. To Calculate the bandwidth requriement requires knowing the audit event rate for the cluster.
2. To get the audit event rate run this command to get disk operations average per PowerScale node
    a. isi statistics query current --nodes=all --stats=node.disk.xfers.rate.sum
    b. This command returns the average per node at the bottom of the results. Use this value in the calcuation below.
    c. Take the average per node and multiple by the number of nodes.
        i. example 2200 (average as reported with command above) * 7 (nodes) = 15,400. Divide this number by 1.83415365 (ratio of audit events to disk transfers).   15,400 / 1.83415365 = 8396 events per second
        ii.
    d. Now use the calculation below to compute the network bandwith required to on average to forward events to the central site for processing.
3. 5 Mbps of network traffic @ 1000 events/sec example 8396 events/sec ÷ 1000 * 5 Mbps = 40 Mbps

Deployment Diagram

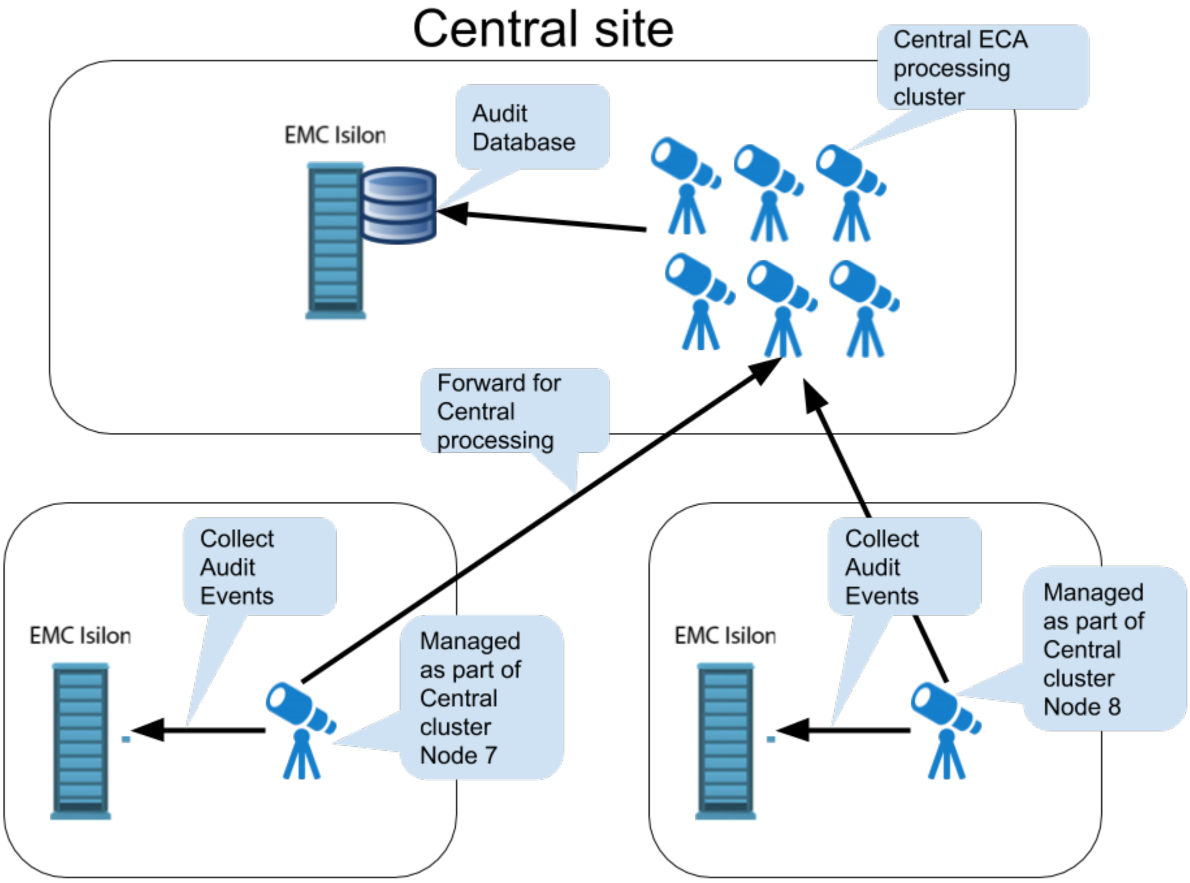
Central site
Central ECA processing cluster
Audit Database
EMC Isilon
Forward for Central processing
Collect Audit Events
EMC Isilon
Managed as part of Central cluster Node 7
Collect Audit Events
EMC Isilon
Managed as part of Central cluster Node 8

How to Deploy Mini-ECA VM's

1. Deploy the OVA as normal following ECA OVA deployment instructions
2. Delete ECA node 2 and 3 for a single Mini-ECA deployment (NOTE: mini-ECA does support HA configurations and can operate with ECA nodes 1 and 2 , this would require only node 3 is deleted from the vApp)
3. Done

How to Configure Mini-ECA nodes to Join Central ECA cluster

1. Login to ECA Central node 1 as ecaadmin
2. vim /opt/superna/eca/eca-env-common.conf

3. Add additional ECA nodes , add a line for each mini-ECA at each remote site and increase the node ID for each new line
    a. export ECA_LOCATION_NODE_7=x.x.x.x
4. run 'ecactl components configure-nodes' to add mini for passwordless ssh
5. vi /opt/superna/eca/data/common/neOverrides.json to add clusters for specific nodes copy the text below and edit based on your configuration and then save the file with :wq
    a. replace cluster name below in yellow with the mini-ECA cluster name, replace the nodes mapping to align the cluster name at a remote site to the ECA node ID configured in the eca-env-common.conf file in the step above.
    b. **NOTE: This mapping must be done correctly to ensure events are processed for the correct cluster and tagged correctly.**
    c. **NOTE: the example file below lists 2 clusters, you will need an entry for each mini-ECA that is deployed.**
    d. **NOTE: the nodes section is identifying ECA node identifier configured in the eca-env-common.conf**
    e. **NOTE: only a single node should be listed if only 1 mini-ECA is used, the example below shows a 2 node mini-ECA and a single node mini-ECA**

```
[{          "name": "SC-8100A",          "nodes":
["2", "3"] }, {          "name": "SC-8100B",
"nodes": ["7"] }]
```

6. Now configure services on the mini-ECA nodes by using the overrides file to specify mini-ECA nodes using the template
7. cp /opt/superna/eca/templates/docker-compose.mini_7_8_9.yml /opt/superna/eca/docker-compose.overrides.yml
    f. **NOTE: This file configures any mini-ECA with the correct services automatically for mini-ECA nodes 7-9 if they exist, no additional configuration is required.**

How to configure NFS mount on Mini-ECA

Each mini ECA will need to mount the cluster it has been assigned.

1. The steps to create the export are the same as this section here.
2. The steps to add to /ets/fstab are the same as this section here.
3. Done.

How to verify the configuration

1. Startup the cluster will start up on all nodes
2. ecactl cluster up
3. verify any start up issues on all nodes
4. Generate test events on each cluster
5. Use wiretap feature to view these events

7. Eyeglass Clustered Agent vApp Install Guide (Ransomware Defender for ECS)

Home Top

Overview

This guide covers installation of the ECA VM's to protect Dell ECS storage. A unified Isilon/Powerscale and Dell ECS installation is also possible. This guide covers unified and standalone deployment. The ECS to ECA forwarding supports Active passive processing with HA and backup ECA VM processing.

This guide should be followed after ECA deployed for Isilon or Powerscale or if deploying ECA nodes for ECS only.

ECA Cluster Sizing and Performance Considerations

1. A single 3 VM ECA cluster can protect up to 4 ECS clusters

Firewall Requirements

NOTE: All other ECA Ransomware Defender ports are required in addition to the ports below.

Port	Direction	Comments
rsyslog TCP 514	ECS --> ECA VM's (all 6 VM's)	Each ECS node requires these ports open
HTTPS TCP 443 TCP	Eyeglass --> ECS management IP	needed for API access from Eyeglass to managed ECS nodes

ECS Setup Steps

This procedure must be replicated on each ECS node

1. Create an rsyslog file at /etc/rsyslog.d/push-dataheadsvc-access-log.conf following the example below with the following edits:
    a. two instances of Target="ip_address" with the IP address of your ECA nodes 2 and 3 respectively (yellow highlight). The first line (node 2) will be the active listener. The second line will take over if the first ECA VM is down.
    b. Tag="vdc1" will need to be updated if you change the name of your VDC

```
#$DebugFile /home/admin/rsyslog.debug
#$DebugLevel 2
module(load="imfile" PollingInterval="1") #needs to be done just once
ruleset(name="ecsaccesslogs") {
```

```
action(type="omfwd" Target="172.25.1.6" Port="514" Protocol="tcp")
action(type="omfwd" Target="172.25.1.7" Port="514" Protocol="tcp"
action.execOnlyWhenPreviousIsSuspended="on")
stop
}

input(type="imfile" ruleset="ecsaccesslogs"
File="/var/log/vipr/emcvipr-object/dataheadsvc-access.log"
Tag="vdc1"
Severity="info"
Facility="local7"
StateFile="ecstosyslog")
```

1. Restart the rsyslog process for the changes to take effect
    a. sudo systemctl restart rsyslog
2. **NOTE Mandatory**: Repeat the above rsyslog configuration on each ECS node in the ECS cluster.

ECA Setup Steps

1. ECS forwards syslog messages to the ECA VM's running syslog
2. Make sure the tuboaudit syslog server is enabled in eca-env-common.conf before cluster up:
3. Login to eca node 1
4. nano /opt/superna/eca/eca-env-common.conf
    a. Add variable below line to the file

    b. export TURBOAUDIT_ECS_SERVER_ENABLED=true
    c. control+x and yes to save

5. If the cluster is running shutdown and restart
    a. ecactl cluster down
    b. ecactl cluster up
6. done.
7. Complete Licensing and remaining configuration following the guide here.

8. Eyeglass Hyper-V Installation Guide

Home Top

- Tested on
- Important Read
- Create Eyeglass Hyper-V Virtual Machine
- Configure Eyeglass data disk
- Configuration of Eyeglass

Tested on

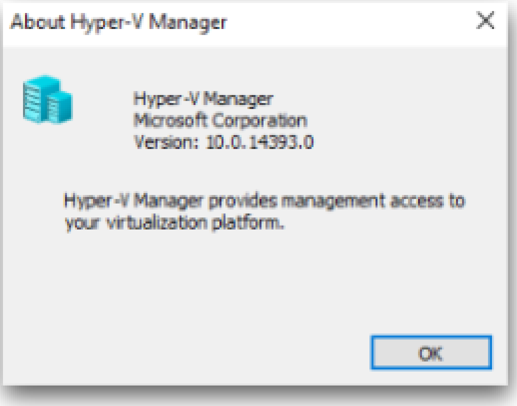

About Hyper-V Manager

Hyper-V Manager
Microsoft Corporation
Version: 10.0.14393.0

Hyper-V Manager provides management access to your virtualization platform.

Important Read

1. Eyeglass appliance uses 2 disks. 1 for OS and 1 for data
2. OS disk requires 20 GB [default disk]
3. Data disk requires 80 GB [read below on how to create]

Create Eyeglass Hyper-V Virtual Machine

1. Download vhdx from https://support.superna.net portal
2. Deploy a new `Virtual Machine`

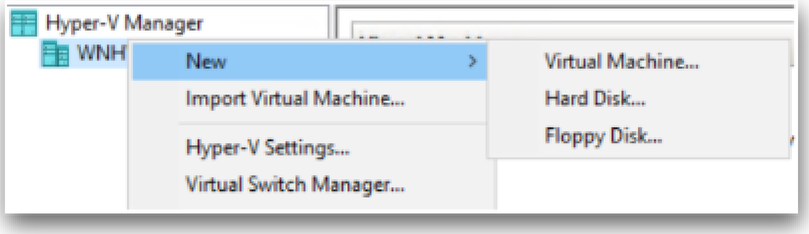

3. Enter `Name` for the VM

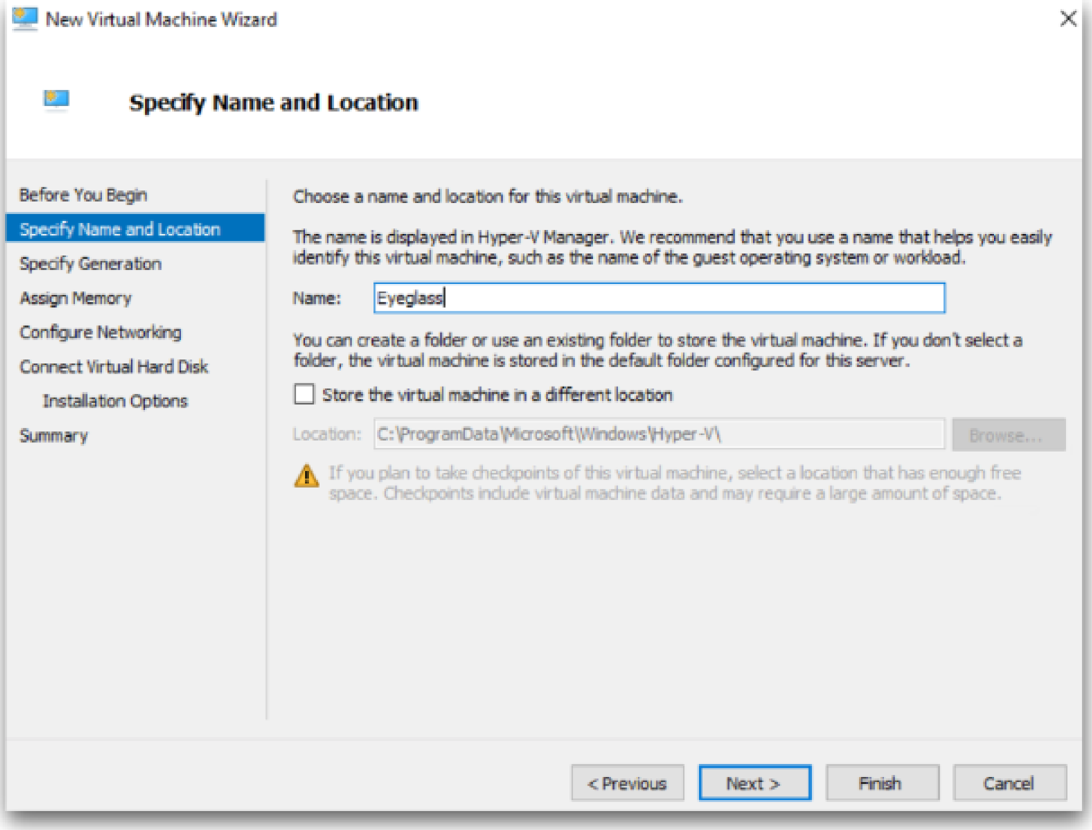

Specify Name and Location

Choose a name and location for this virtual machine.

The name is displayed in Hyper-V Manager. We recommend that you use a name that helps you easily identify this virtual machine, such as the name of the guest operating system or workload.

Name: Eyeglass

You can create a folder or use an existing folder to store the virtual machine. If you don't select a folder, the virtual machine is stored in the default folder configured for this server.

□ Store the virtual machine in a different location

Location: C:\ProgramData\Microsoft\Windows\Hyper-V\

If you plan to take checkpoints of this virtual machine, select a location that has enough free space. Checkpoints include virtual machine data and may require a large amount of space.

4. Check `Generation 1`

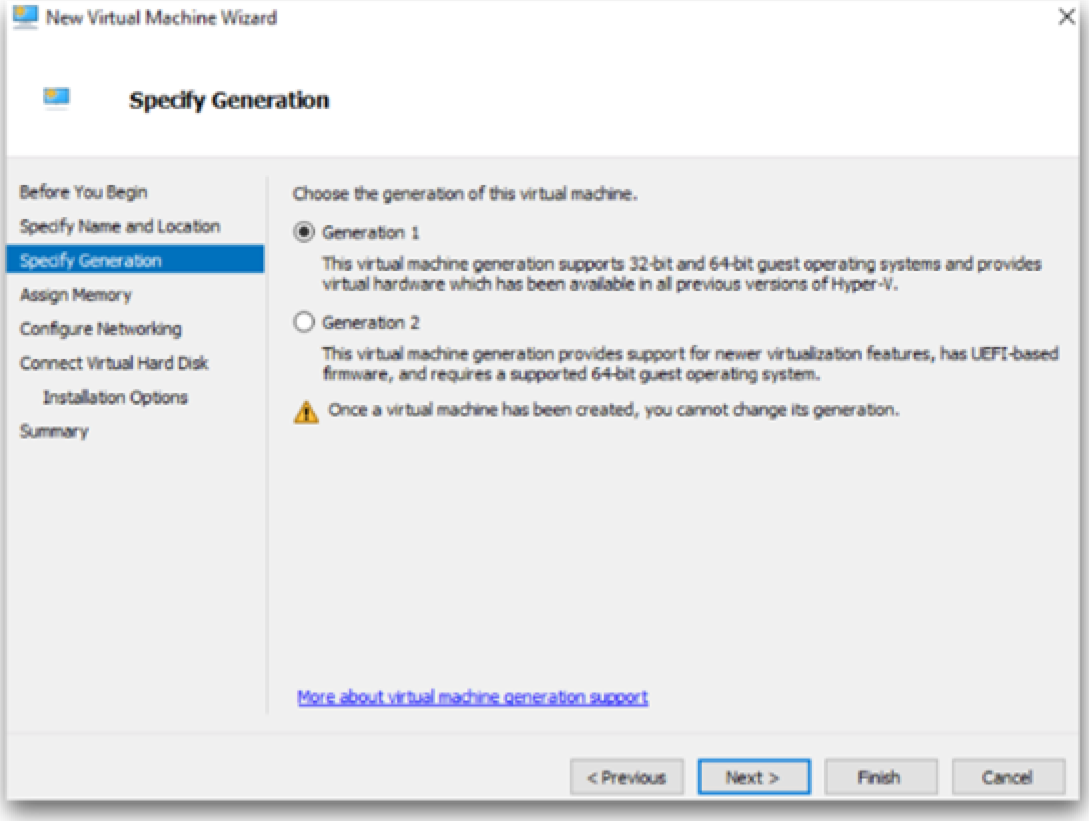

Specify Generation

Choose the generation of this virtual machine.

● Generation 1

This virtual machine generation supports 32-bit and 64-bit guest operating systems and provides virtual hardware which has been available in all previous versions of Hyper-V.

○ Generation 2

This virtual machine generation provides support for newer virtualization features, has UEFI-based firmware, and requires a supported 64-bit guest operating system.

Once a virtual machine has been created, you cannot change its generation.

More about virtual machine generation support

5. Startup memory 16384 MB [16 GB]

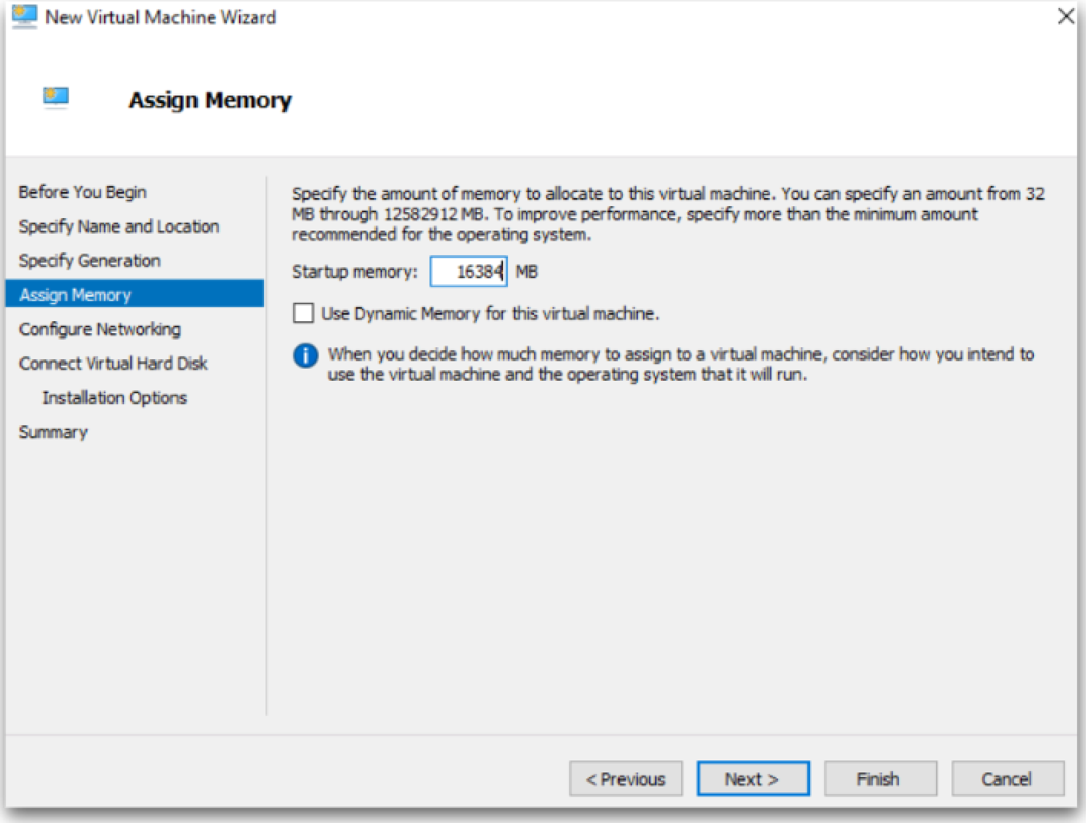

New Virtual Machine Wizard

Assign Memory

Specify the amount of memory to allocate to this virtual machine. You can specify an amount from 32 MB through 12582912 MB. To improve performance, specify more than the minimum amount recommended for the operating system.

Startup memory: 16384 MB

□ Use Dynamic Memory for this virtual machine.

When you decide how much memory to assign to a virtual machine, consider how you intend to use the virtual machine and the operating system that it will run.

6. Select `Network Adapter`

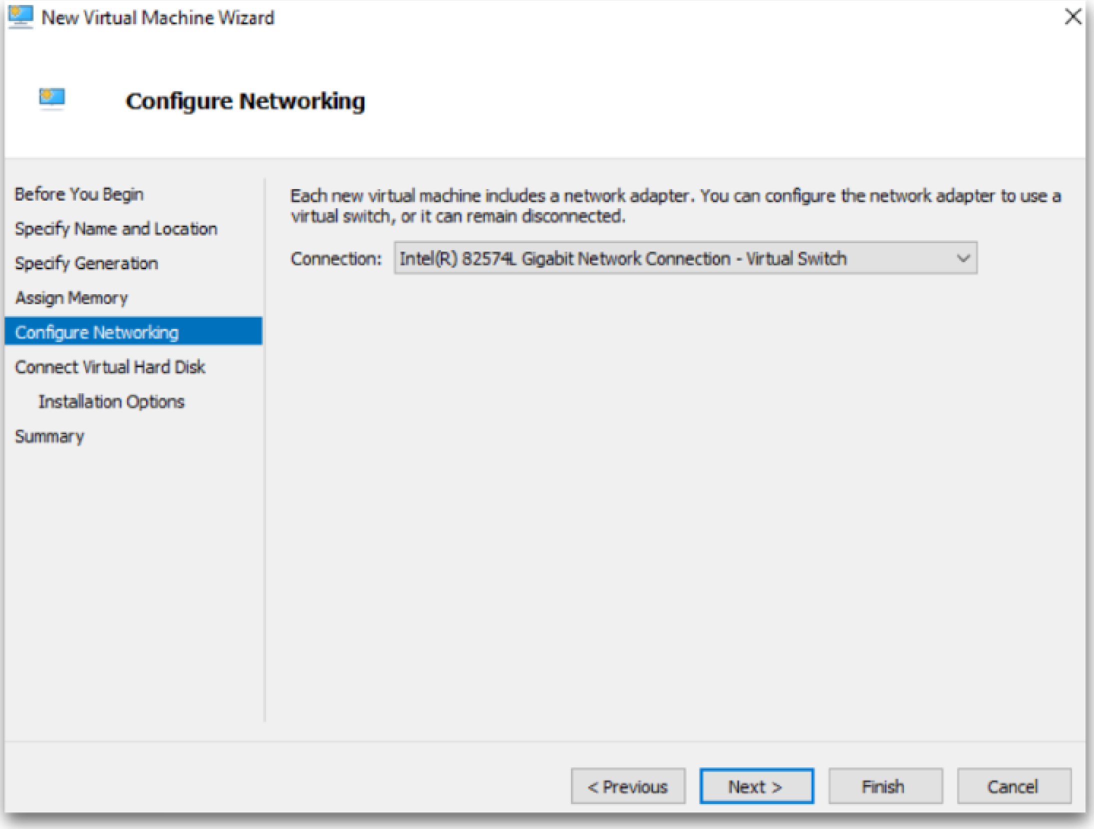

New Virtual Machine Wizard

Configure Networking

Each new virtual machine includes a network adapter. You can configure the network adapter to use a virtual switch, or it can remain disconnected.

Connection: Intel(R) 82574L Gigabit Network Connection - Virtual Switch

7. Use an existing `Virtual Hard disk` → Browse to newly downloaded Eyeglass `vhdx` file

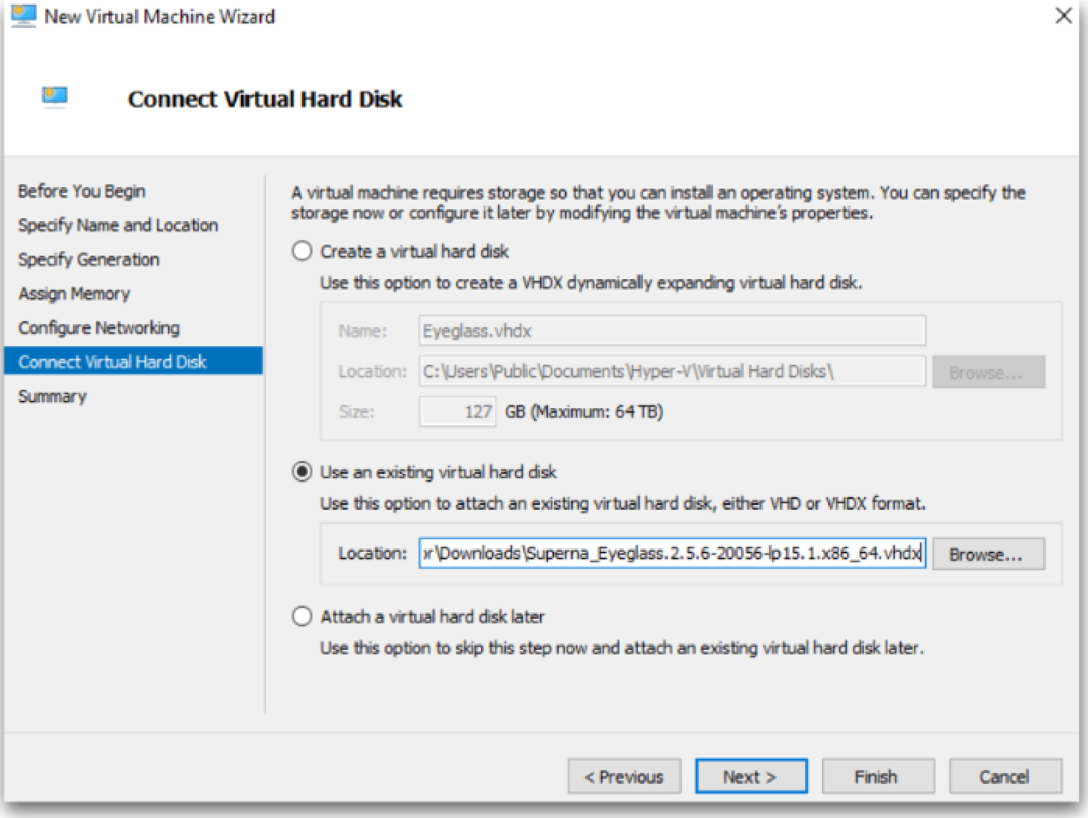

Connect Virtual Hard Disk

A virtual machine requires storage so that you can install an operating system. You can specify the storage now or configure it later by modifying the virtual machine's properties.

○ Create a virtual hard disk

Use this option to create a VHDX dynamically expanding virtual hard disk.

Name: Eyeglass.vhdx

Location: C:\Users\Public\Documents\Hyper-V\Virtual Hard Disks\

Size: 127 GB (Maximum: 64 TB)

● Use an existing virtual hard disk

Use this option to attach an existing virtual hard disk, either VHD or VHDX format.

Location: r\Downloads\Superna_Eyeglass.2.5.6-20056-lp15.1.x86_64.vhdx

○ Attach a virtual hard disk later

Use this option to skip this step now and attach an existing virtual hard disk later.

8. Complete the Wizard

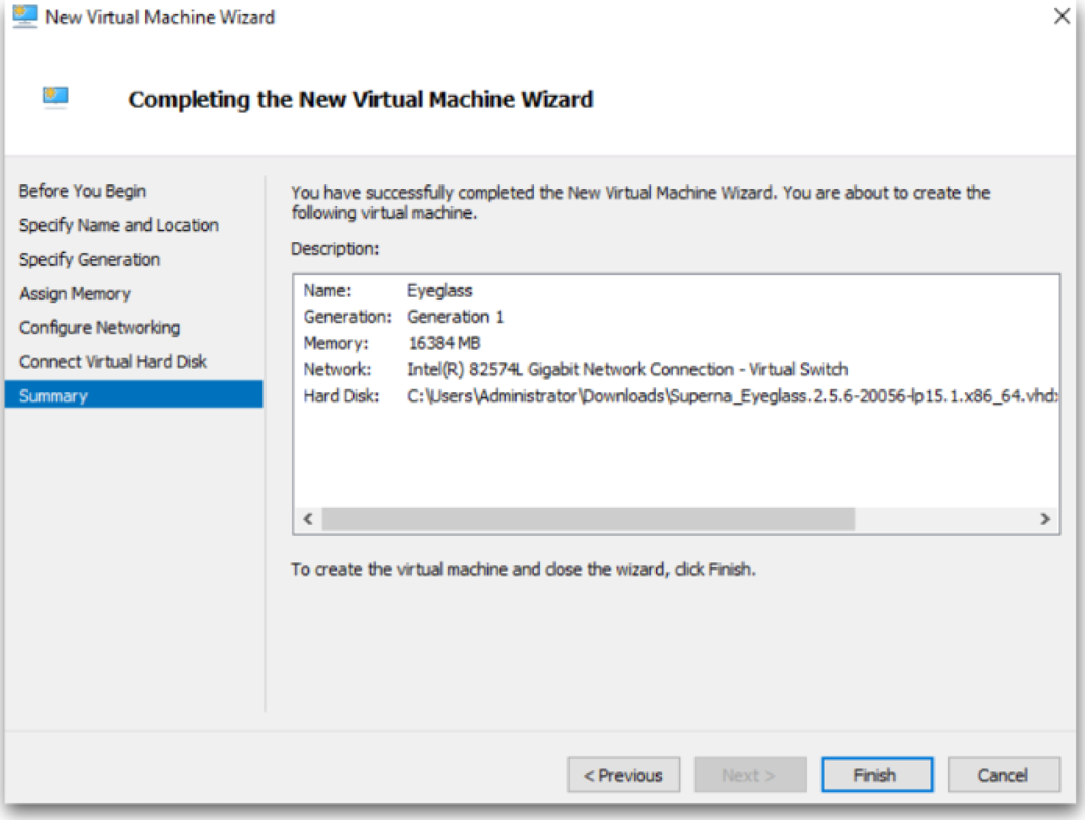

Completing the New Virtual Machine Wizard

You have successfully completed the New Virtual Machine Wizard. You are about to create the following virtual machine.

Description:

Name: Eyeglass
Generation: Generation 1
Memory: 16384 MB
Network: Intel(R) 82574L Gigabit Network Connection - Virtual Switch
Hard Disk: C:\Users\Administrator\Downloads\Superna_Eyeglass.2.5.6-20056-lp15.1.x86_64.vhd:

To create the virtual machine and close the wizard, click Finish.

Configure Eyeglass data disk

1. After deploying, go to the new VM → Right Click → Settings

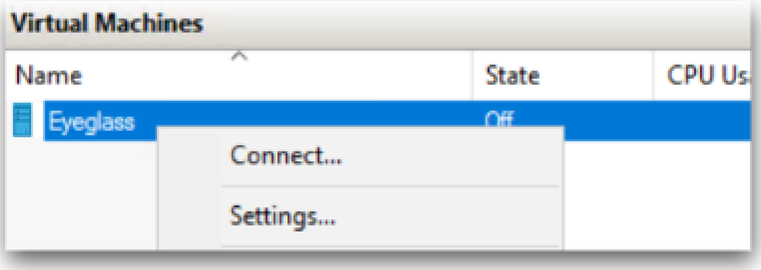

2. From `IDE Controller 0` → Add a `Hard Drive`

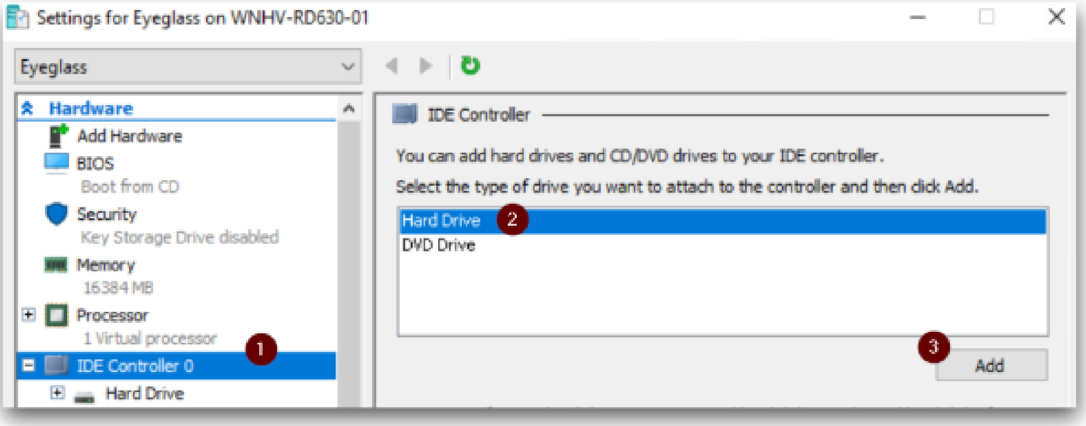

You can add hard drives and CD/DVD drives to your IDE controller.

Select the type of drive you want to attach to the controller and then click Add.

3. Create New

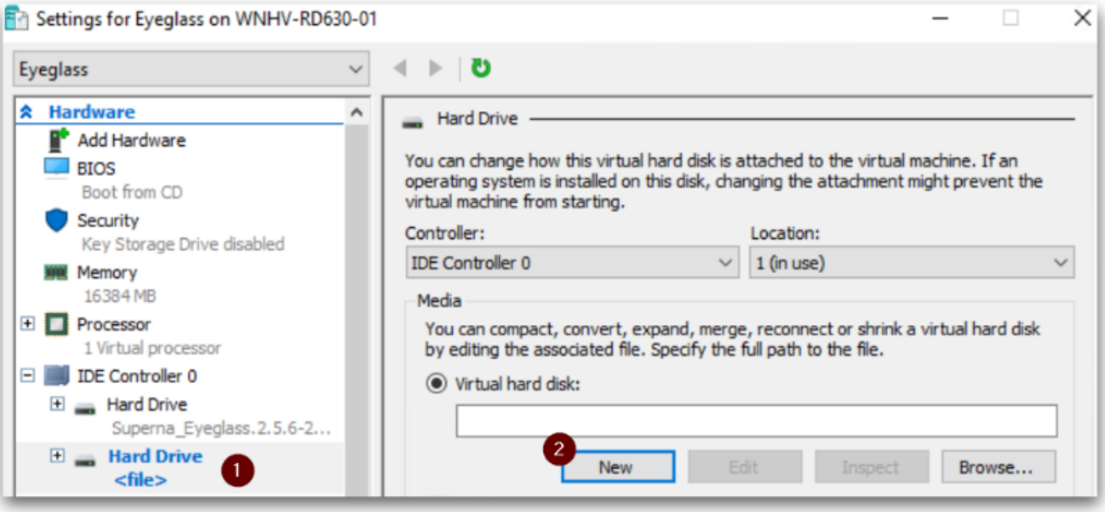

You can change how this virtual hard disk is attached to the virtual machine. If an operating system is installed on this disk, changing the attachment might prevent the virtual machine from starting.

Controller: IDE Controller 0

Location: 1 (in use)

You can compact, convert, expand, merge, reconnect or shrink a virtual hard disk by editing the associated file. Specify the full path to the file.

● Virtual hard disk:

4. Choose Disk Format → `VHDX`

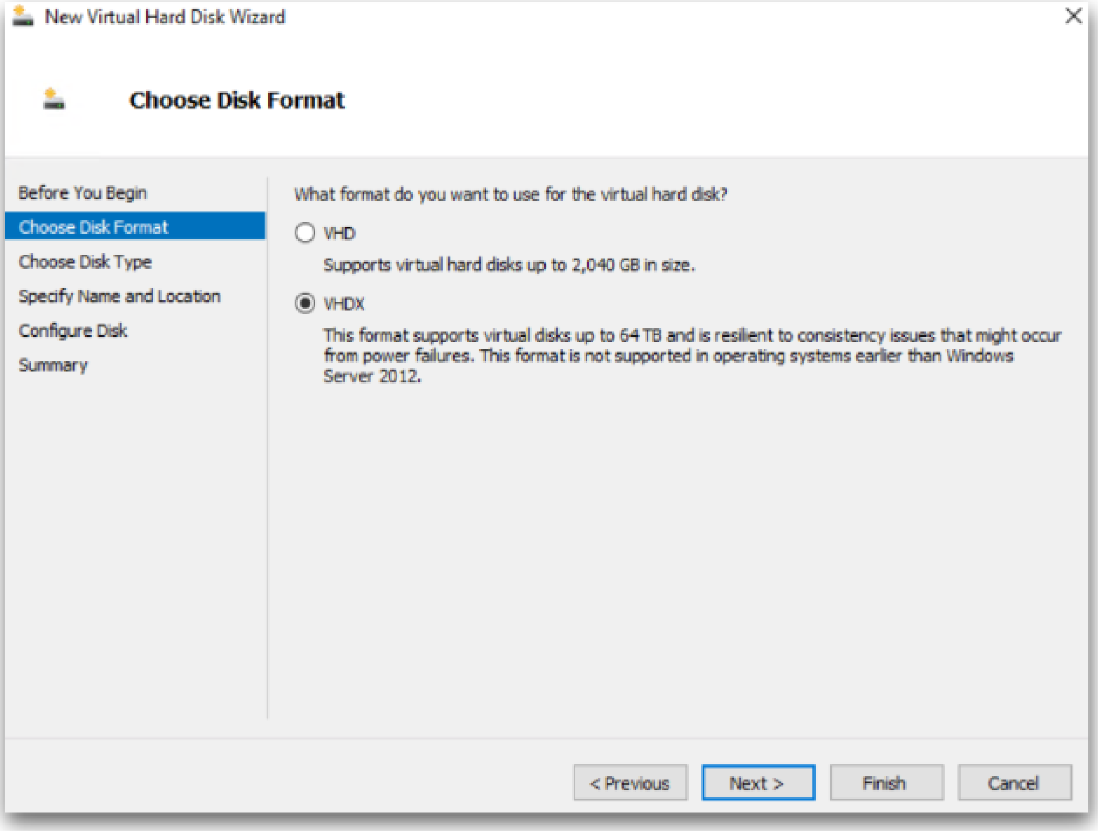

New Virtual Hard Disk Wizard

Choose Disk Format

What format do you want to use for the virtual hard disk?

○ VHD

Supports virtual hard disks up to 2,040 GB in size.

● VHDX

This format supports virtual disks up to 64 TB and is resilient to consistency issues that might occur from power failures. This format is not supported in operating systems earlier than Windows Server 2012.

5. Choose Disk Type → `Fixed size`

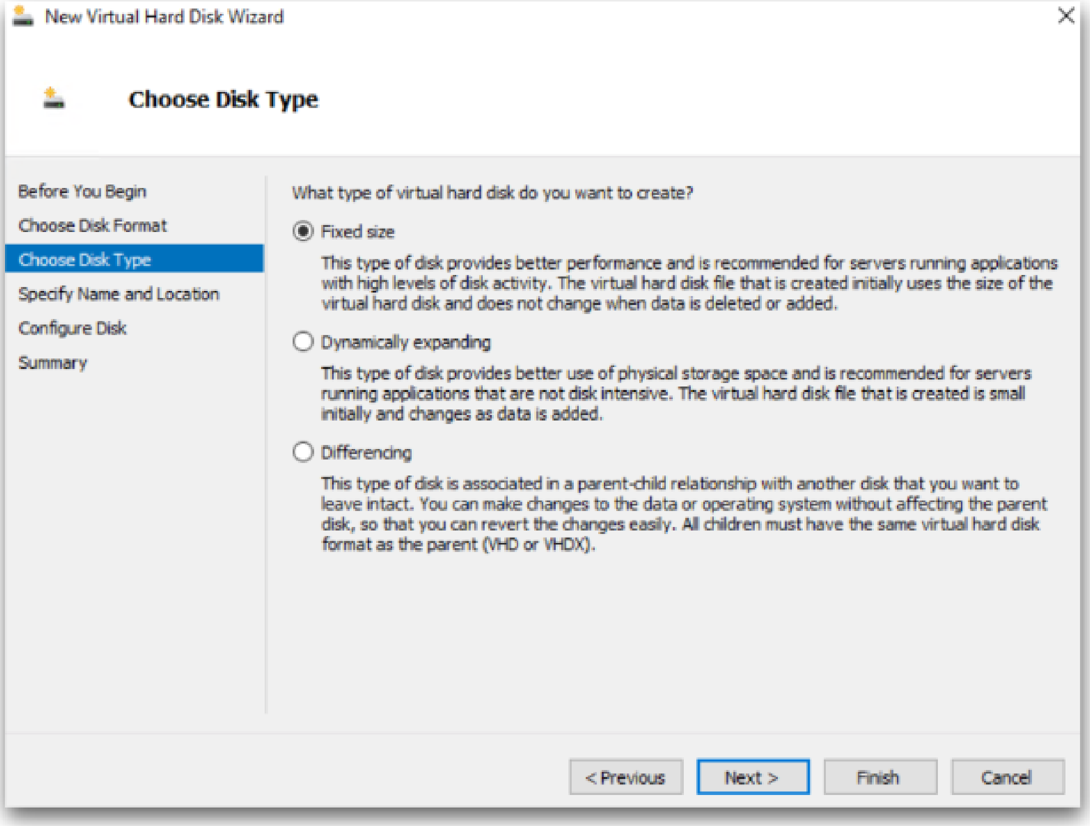

New Virtual Hard Disk Wizard

Choose Disk Type

What type of virtual hard disk do you want to create?

● Fixed size

This type of disk provides better performance and is recommended for servers running applications with high levels of disk activity. The virtual hard disk file that is created initially uses the size of the virtual hard disk and does not change when data is deleted or added.

○ Dynamically expanding

This type of disk provides better use of physical storage space and is recommended for servers running applications that are not disk intensive. The virtual hard disk file that is created is small initially and changes as data is added.

○ Differencing

This type of disk is associated in a parent-child relationship with another disk that you want to leave intact. You can make changes to the data or operating system without affecting the parent disk, so that you can revert the changes easily. All children must have the same virtual hard disk format as the parent (VHD or VHDX).

6. Name the data disk

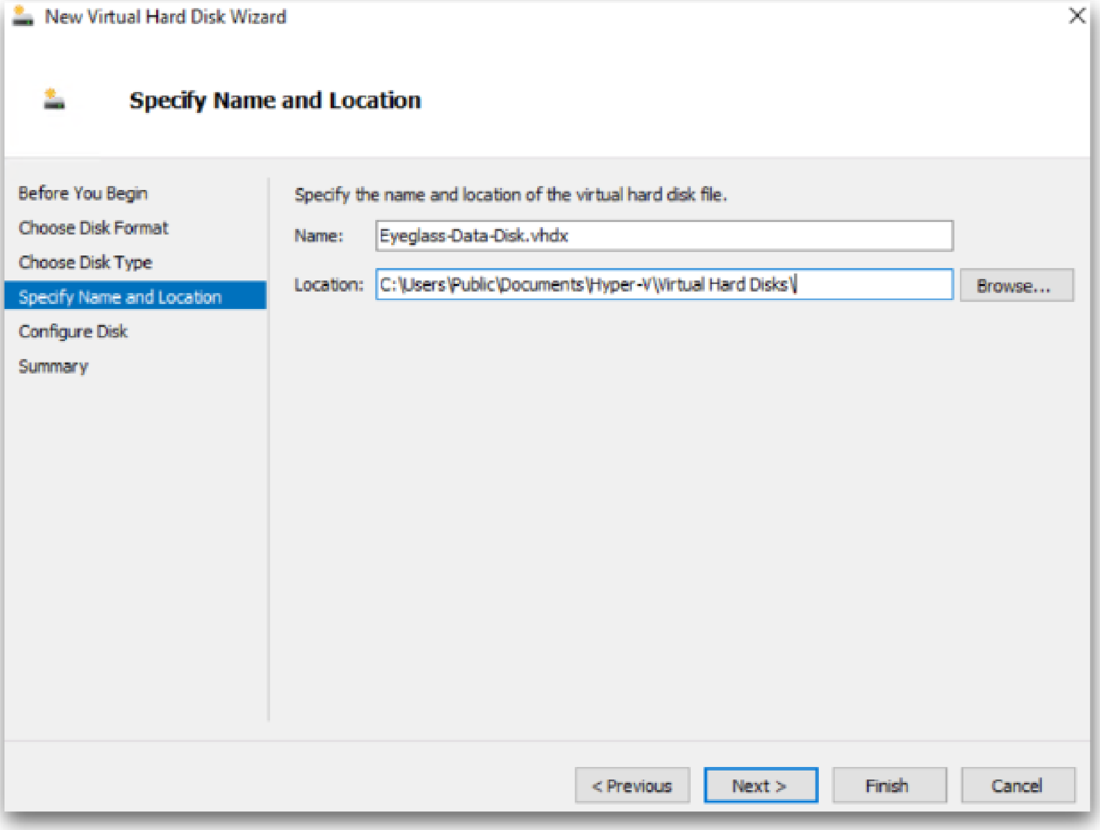

New Virtual Hard Disk Wizard

Specify Name and Location

Specify the name and location of the virtual hard disk file.

Name: Eyeglass-Data-Disk.vhdx

Location: C:\Users\Public\Documents\Hyper-V\Virtual Hard Disks\

7. Create a new blank virtual hard disk : **80 GB**

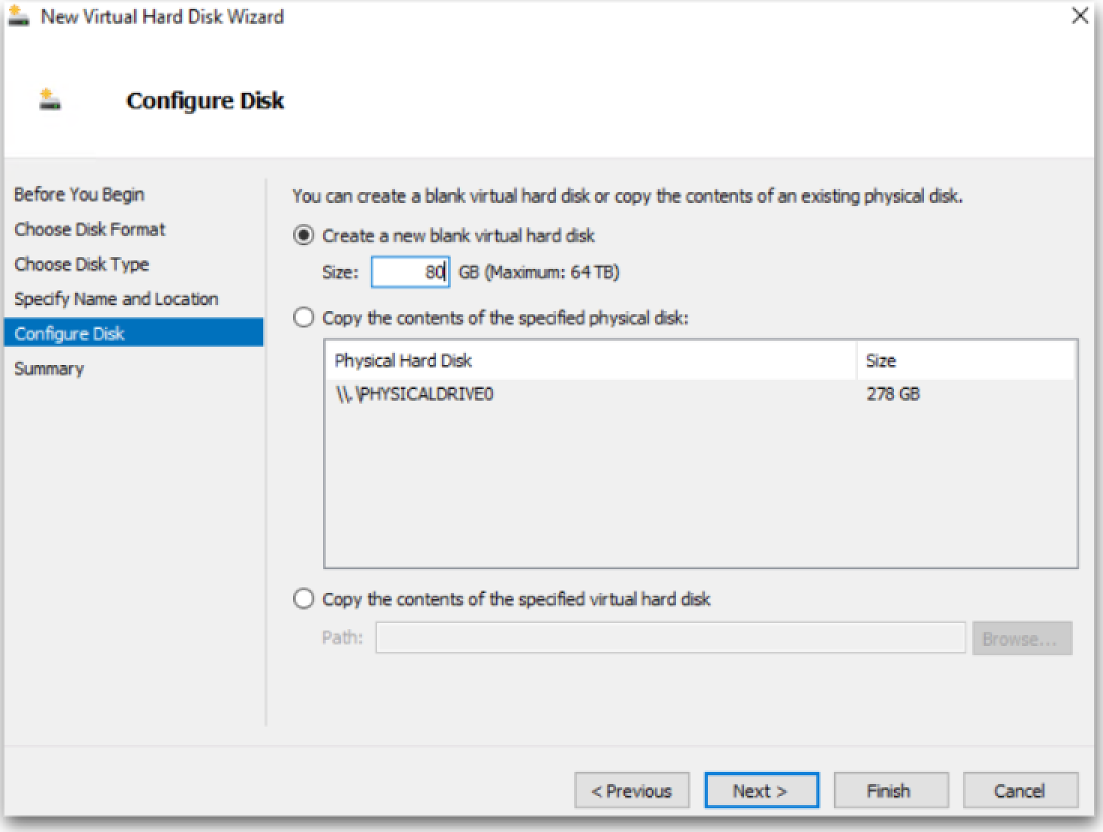

New Virtual Hard Disk Wizard

Configure Disk

You can create a blank virtual hard disk or copy the contents of an existing physical disk.

● Create a new blank virtual hard disk

Size: 80 GB (Maximum: 64 TB)

○ Copy the contents of the specified physical disk:

Physical Hard Disk	Size
\\.\PHYSICALDRIVE0	278 GB

○ Copy the contents of the specified virtual hard disk

Path:

8. Complete the data disk Wizard

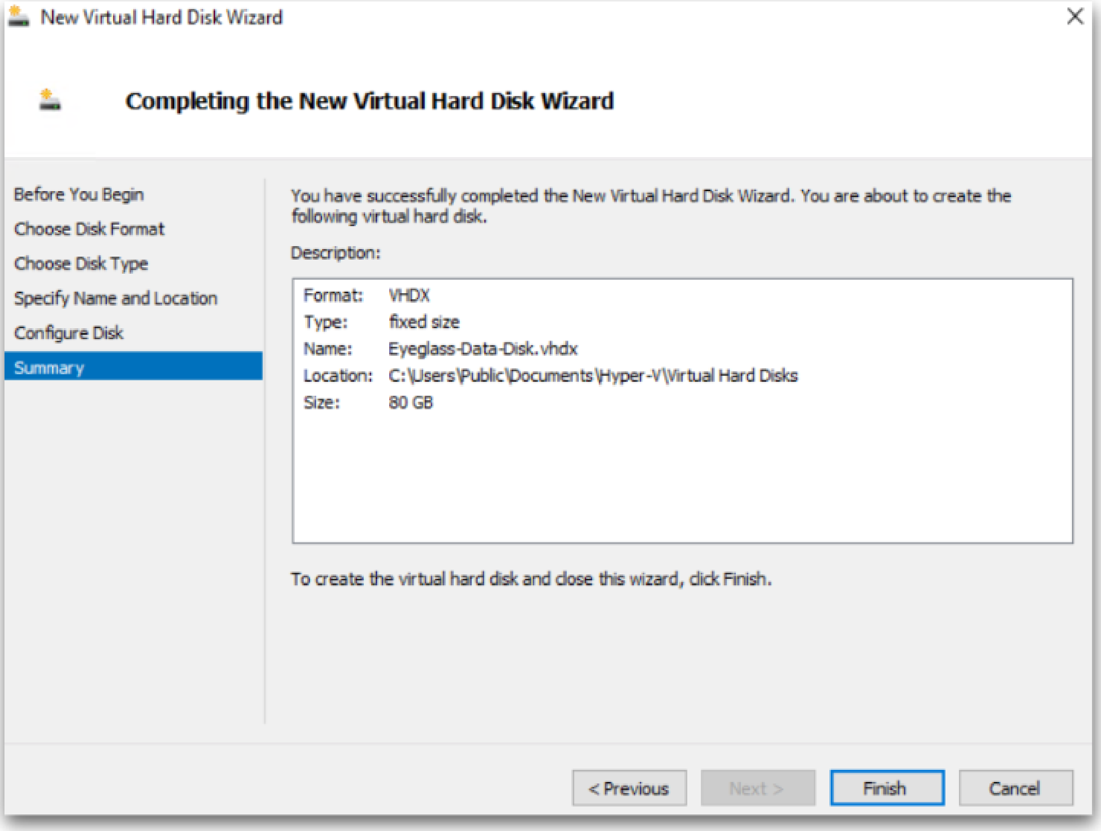

You have successfully completed the New Virtual Hard Disk Wizard. You are about to create the following virtual hard disk.

Description:

Format: VHDX
Type: fixed size
Name: Eyeglass-Data-Disk.vhdx
Location: C:\Users\Public\Documents\Hyper-V\Virtual Hard Disks
Size: 80 GB

To create the virtual hard disk and close this wizard, click Finish.

Configuration of Eyeglass

1. Power up the VM and wait 5-10 minutes to allow Superna on-boot script to run. tail the superna-on-boot.log and wait for it to finish. Then follow the on-screen instruction

```
tail -2 /var/log/superna-on-boot.log
```

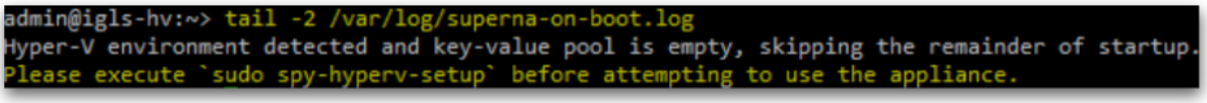

```
admin@igls-hv:~> tail -2 /var/log/superna-on-boot.log
Hyper-V environment detected and key-value pool is empty, skipping the remainder of startup.
Please execute `sudo spy-hyperv-setup` before attempting to use the appliance.
```

2. Run the command to setup your Eyeglass Hyper-V appliance

```
sudo spy-hyperv-setup
```

When prompted, enter `admin` user password [default password is: 3y3gl4ss] and enter IP, Netmask, Gateway, Hostname, DNS, NTP info

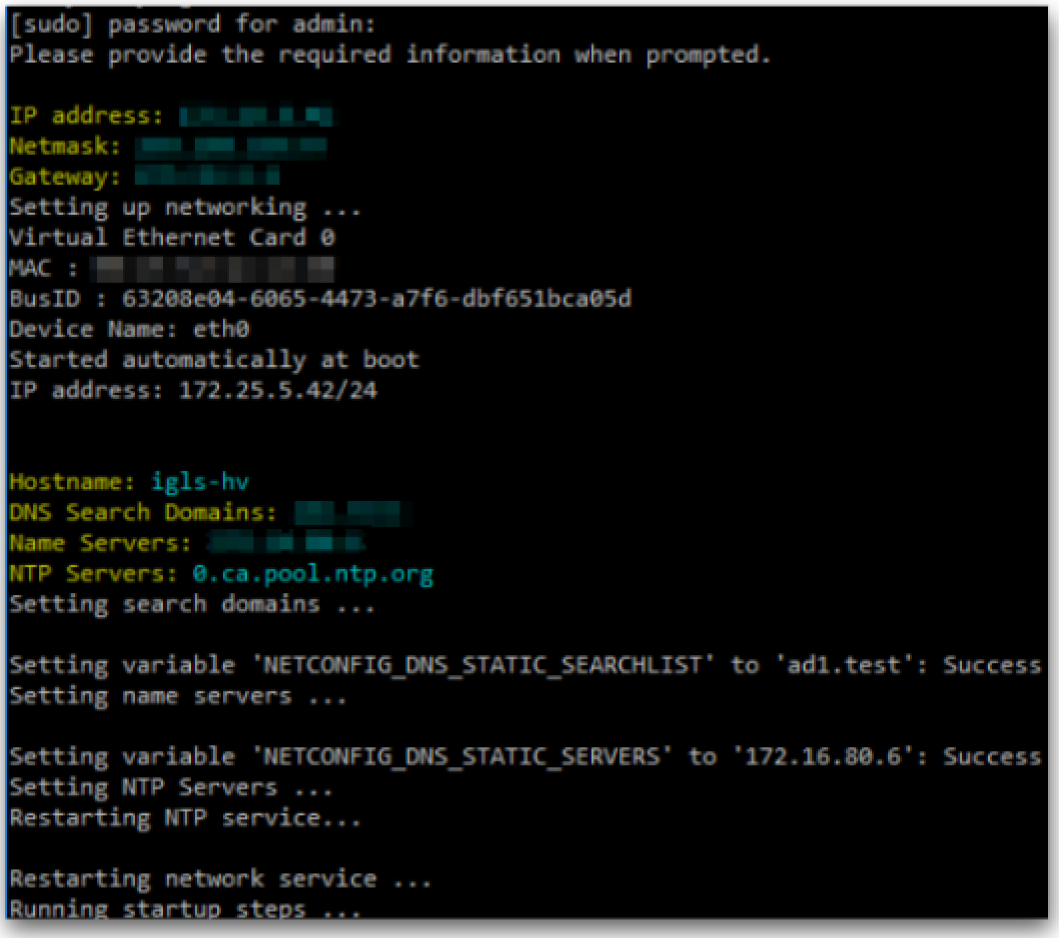

```
[sudo] password for admin:
Please provide the required information when prompted.

IP address: [redacted]
Netmask: [redacted]
Gateway: [redacted]
Setting up networking ...
Virtual Ethernet Card 0
MAC : [redacted]
BusID : 63208e04-6065-4473-a7f6-dbf651bca05d
Device Name: eth0
Started automatically at boot
IP address: 172.25.5.42/24


Hostname: igls-hv
DNS Search Domains: [redacted]
Name Servers: [redacted]
NTP Servers: 0.ca.pool.ntp.org
Setting search domains ...

Setting variable 'NETCONFIG_DNS_STATIC_SEARCHLIST' to 'ad1.test': Success
Setting name servers ...

Setting variable 'NETCONFIG_DNS_STATIC_SERVERS' to '172.16.80.6': Success
Setting NTP Servers ...
Restarting NTP service...

Restarting network service ...
Running startup steps ...
```

---

3. Now check the service status and ensure they are in **active (running)** state

```
systemctl status sca scadb lighttpd tomcat sera
```

4. Open Google Chrome browser and go to the following link

```
https://<Eyeglass_IP>
```

9. Eyeglass PowerScale Edition Quick Start Guide and Upgrade Guide for RHEL and Centos 7.6, 7.7 and 8.x

Home Top

Read Me First

This installer option requires a purchased RPM install license key, a trial key option is available. The OS key is required to use this installer option in production. This is not part of enterprise keys and excluded from maintenance contracts.

This license key allows customers to build their own appliance.
 Email sales@superna.net for a assistance with ordering.

Eyeglass Quick Start

Use this document to get your new Eyeglass installation up and running fast with all the best options.

For planning DR and understanding design choices with Eyeglass use the Eyeglass Start Here First Guide

System Requirements

Operating System:

- CentOS Version 7.6, 7.7, 8.x
- Red Hat Enterprise Linux Version 7.6, 7.7, 8.x
- NOTE: The OS itself is not covered under the support contract.

Appliance requirements:

- 4 vCPU
- 16G RAM
- RAM please see scalability table.

*Note requires shutting down the VM and editing RAM and restart)

- 130 GB disk
- Chrome Browser (preferred), Browser must support Websockets, Internet Explorer not supported. ,
- Eyeglass Port Requirements: Eyeglass-Ports-Requirements

Support limitations

1. The Operating system maintenance of patches and updates is customer responsibility

2. Installation of the Linux Eyeglass installer combines application software and tested and supported dependencies including the following:
    a. Sudoer configuration
    b. Lighttpd

1. **Support statement:** This yum package manager will not be allowed to upgrade these components unless forced overwrite option is used. The application versions installed by the Eyeglass dependency rpm is version controlled and is the supported version. Customers that want to update these packages own the risk of breaking application functionality. Support of versions other than the provided version is excluded from support contract coverage. Customers will be asked to downgrade the version if application functionality is impacted by the customer forced override of a supported package version.
2. See Appendix A for a controlled list of files set as owned by Eyeglass RPM dependencies and how to force update the affected applications.

Supported OneFS releases

Please refer to the Release Notes for the Eyeglass PowerScale Edition version that you are installing.

Feature Release Compatibility

Please refer to the Release Notes for the Eyeglass PowerScale Edition version that you are installing.

Eyeglass Scalability Limits

Please refer to the Eyeglass Admin Guide Scalability limits.

New Eyeglass Installation

If you are doing a new Eyeglass installation, continue following steps in this document.

Download Eyeglass

Request download Eyeglass RHEL/CENTOS from Superna website Latest Appliance Code Download

Deploy the virtual machine with Centos or RHEL

Eyeglass is delivered as RPM installer and has dependencies that must be installed

1. Subnet and network required so that appliance will have IP connectivity to the PowerScale clusters that it is managing and the users that are using it
2. IP address for the appliance
3. Gateway
4. DNS server
5. RPM package dependencies required before installation:
    a. RHEL 7.x
        i. nodejs, syslog-ng, lighttpd, protobuf-python, tomcat, net-tools, nfs-utils, python-setuptools, zip , unzip, net-tools-deprecated, python34-pip, rsync, sysstat, fail2ban, perl and epel-release
        ii. The rhel-7-server-optional-rpms repository will have to be enabled for Red Hat Enterprise Linux 7 installations. Additionally the epel-release package will have to be manually added. Be sure to follow the detailed instructions below for procedure.
    b. RHEL 8.x
        i. nodejs syslog-ng lighttpd python3-protobuf fontconfig net-tools nfs-utils python3-setuptools zip

unzip python3-pip rsync sysstat fail2ban perl lighttpd

ii. The epel-release-latest-8.noarch.rpm repository will have to be enabled for Red Hat Enterprise Linux 8 installations. Additionally the epel-release package will have to be manually added. Be sure to follow the detailed instructions below for procedure.

6. **NOTE: If you are using a hostname or FQDN for the target cluster in your SyncIQ policies, the DNS information entered here must be able to resolve it back to a discovered cluster ip address (should resolve to a synciq smartconnect zone ip pool ip address) in order for Eyeglass to perform configuration replication. Eyeglass will not create the associated configuration replication job. If the hostname or FQDN cannot be resolved.**

Steps to Install

Appliance Deployment steps :

Step 1 : Install Centos or RHEL VM as per above requirements

Step 2 : Download RHEL/Centos RPM installer file: Download the correct run file to your VM using the 7.x or 8.x install file. Follow steps to download the installer https://www.supernaeyeglass.com/downloads   and NOTE: You will need the password to download the RHEL licensed installer.

Step 3 : Login as root.

Step 4 : SCP install file to the VM.

Step 5 : chmod 755 install-file-name

Step 6: Edit the file /etc/sysconfig/selinux and set up the "SELINUX=permissive" parameter, after reboot the vm.

Step 7: Install dependencies

RHEL 7.x

1. **If you have RHEL**, you must enable the “*Red Hat Enterprise Linux 7 Server - Optional (RPMs)*” repository by executing the following command:

```
subscription-manager repos --enable rhel-7-server-optional-rpms
```

2. Enable the *Extra Packages for Enterprise Linux* repository:

```
curl -Ok https://dl.fedoraproject.org/pub/epel/epel-release-latest-7.noarch.rpm
yum install ./epel-release-latest-7.noarch.rpm
```

1. All supported versions should now complete dependency install:

```
yum install nodejs syslog-ng lighttpd protobuf-python tomcat net-tools nfs-utils python-setuptools zip unzip net-tools-deprecated python34-pip rsync sysstat fail2ban perl -y
```

Note:	The syslog-ng has a conflict with rsyslog, you need to remove the rsyslog package to be able install syslog-ng package by executing the following command: **yum remove rsyslog**

RHEL 8.x

```
curl -Ok https://dl.fedoraproject.org/pub/epel/epel-release-latest-8.noarch.rpm
yum install ./epel-release-latest-8.noarch.rpm
yum install nodejs syslog-ng lighttpd python3-protobuf fontconfig net-tools nfs-utils python3-setuptools zip unzip python3-pip rsync sysstat fail2ban perl lighttpd -y
```

Note:	The syslog-ng has a conflict with rsyslog, you need to remove the rsyslog package to be able install syslog-ng package by executing the following command: **yum remove rsyslog**

Step 8: ./install-file-name (use the 7.x file or the 8.x install file)

Step 9: verify output to verify installation completes without error, dependency checks will fail and indicate which packages are not installed

Step 10: Check services are running correctly

- systemctl status sca
- systemctl status -l scadb
- systemctl status lighttpd

Step 11: Send install log file in same directory as install file to http://support.superna.net if installation fails, by opening a support case and attaching file to the case

Step 12: Eyeglass appliance sudoers file needs to be updated with this information:

vi sudo configuration file
 sca ALL=(ALL) NOPASSWD: /opt/bin/yum, /opt/superna/bin/kill_packagekit.sh, /opt/superna/sbin/*, /usr/bin/systemctl restart syslog

Eyeglass Initial Configuration

Your Eyeglass initial configuration steps are:

1. Login to the Eyeglass UI
2. Install License
3. Create Eyeglass service account first for each PowerScale cluster with minimum permissions (if not done configure Clusters in Eyeglass using root user)
4. Add Clusters ()

Login to the Eyeglass UI

To login to the Eyeglass web UI, enter the following URL into your browser (Chrome preferred) replacing <Eyeglass IP address> with the real IP address assigned to the appliance:

https://<Eyeglass IP address>

You have 2 options for login authentication:
Local - Select Auth Type “Local” and use the admin user and password configured on the appliance
Default user/password: **admin / 3y3gl4ss**

Install License

Retrieve your Eyeglass License keys (instructions provided here).

NOTE: You will require a CENTOS or RHEL OS license key in addition to other Eyeglass product keys

Upload the license zip file provided to you by Superna:

IMPORTANT: Do not unzip the license file. Upload the zip file.

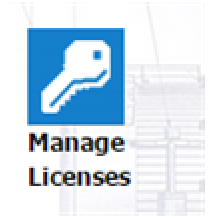

IMPORTANT: You will be asked to accept the Eyeglass EULA and Phone Home after selecting the Upload button. License will not be loaded unless EULA is accepted.

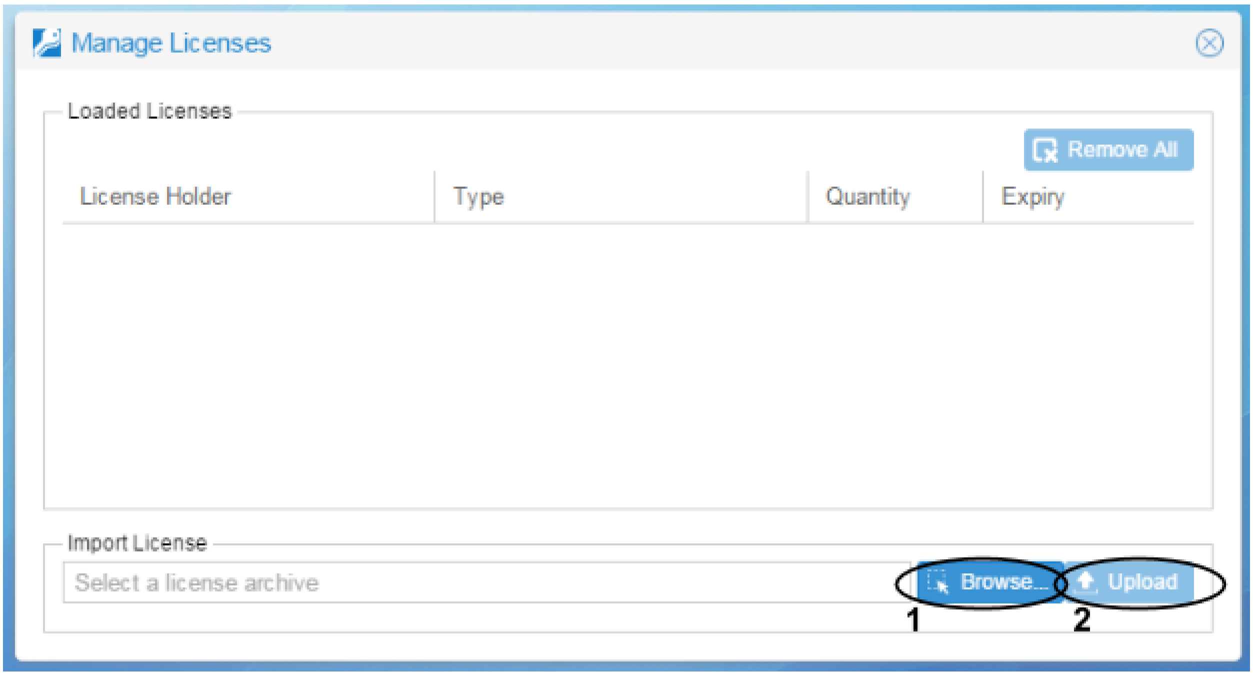

Add PowerScale Clusters

NOTE No Auto Refresh Inventory View

This window does not auto refresh after adding a cluster. You must click the refresh button bottom right to verify when a cluster has finished discovery. This process can take 5-10 minutes typically.

NOTE Cluster DNS Setup and Add Cluster to Inventory:

If discovery takes a very long time to complete (> 10 minutes), then it's important to check cluster configuration data can resolve external URL. Cloud pools uses a URL to a storage bucket and if this URL can not complete DNS lookup to IP address API calls that discovery cloud pools will take too long to complete and will timeout the cluster discovery.

Important After Discovery of a Cluster’s SyncIQ policies all eyeglass configuration jobs are disabled automatically

Configuration Replication Jobs for zones, shares, exports and nfs alias protected by SyncIQ Policy automatically created and are in USERDISABLED state after successful provisioning in Eyeglass. Enabling these Jobs will be part of the installation steps.

Important Clusters on source target must be in the support feature matrix

PowerScale cluster replication pairs must be running supported OneFS version as documented in the System Requirements / Feature Release Compatibility matrix.

Before you add a cluster to Eyeglass verify SyncIQ FQDN Name resolution

This step is important to allow eyeglass to automatically build configuration replication jobs correctly. Eyeglass will resolve the FQDN of the SyncIQ policy and then compare the returned ip address to all PowerScale clusters added to the eyeglass appliance. If no match is found, Config Sync jobs will fail be be added to the jobs window, until name resolution works correctly. A system alarm is also raised that indicates no matching clusters found for the SyncIQ policies on Cluster named X.

1. Login to eyeglass
2. open eyeglass shell from eyeglass main menu (bottom left)
3. login as admin with default password 3y3gl4ss
4. Get list of SyncIQ policies from the source cluster you are adding and record the FQDN target host value used in the policy
5. validate the FQDN will resolve correctly on eyeglass
6. nslookup FQDN
    1. If an ip address does not get returned you MUST fix this using YAST utility to add DNS to eyeglass (see admin guide for instructions)
    2. OR you must sudo to root with sudo -s (enter admin password)
    3. vi /etc/hosts and add an entry for the FQDN value that does not resolve correctly

4. **NOTE: DNS is the preferred solution to resolve entries, hosts file can be used as a work around on the appliance for each smartconnect zone that does not resolve to an ip address**

7. Repeat nslookup step for each FQDN used on each cluster you want to add to eyeglass for DR management

From the Eyeglass UI add the PowerScale Clusters between which Eyeglass will be replicating the share and export configuration data.

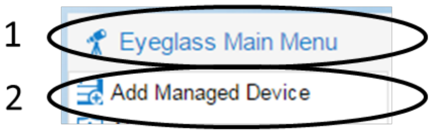

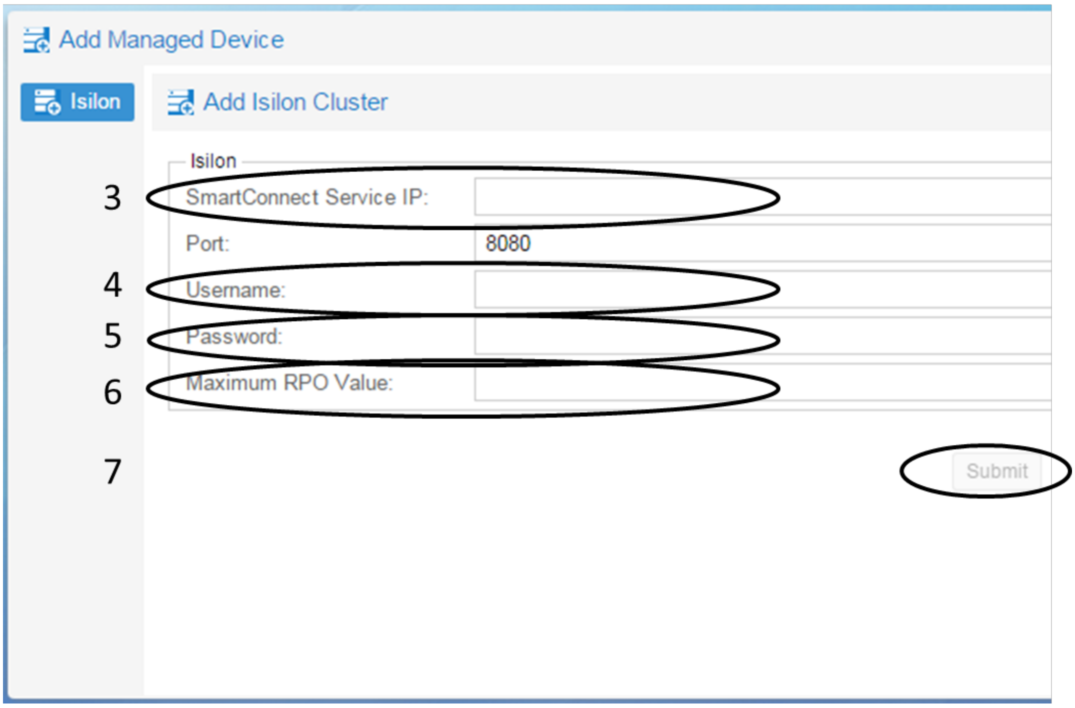

Note:

- SmartConnect Service must be IP address format.
- Maximum RPO Value is the Recovery Point Objective for the cluster in **minutes**. If you are using the RPO feature, this target is used during RPO analysis. More information about Eyeglass RPO analysis can be found here.
- To create an Eyeglass service account with minimum privileges follow the instructions provided here.

Once the PowerScale is added, Eyeglass will automatically run an inventory task to discover the PowerScale components. When completed, the discovered inventory can be seen in the Inventory View.

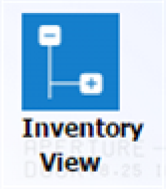

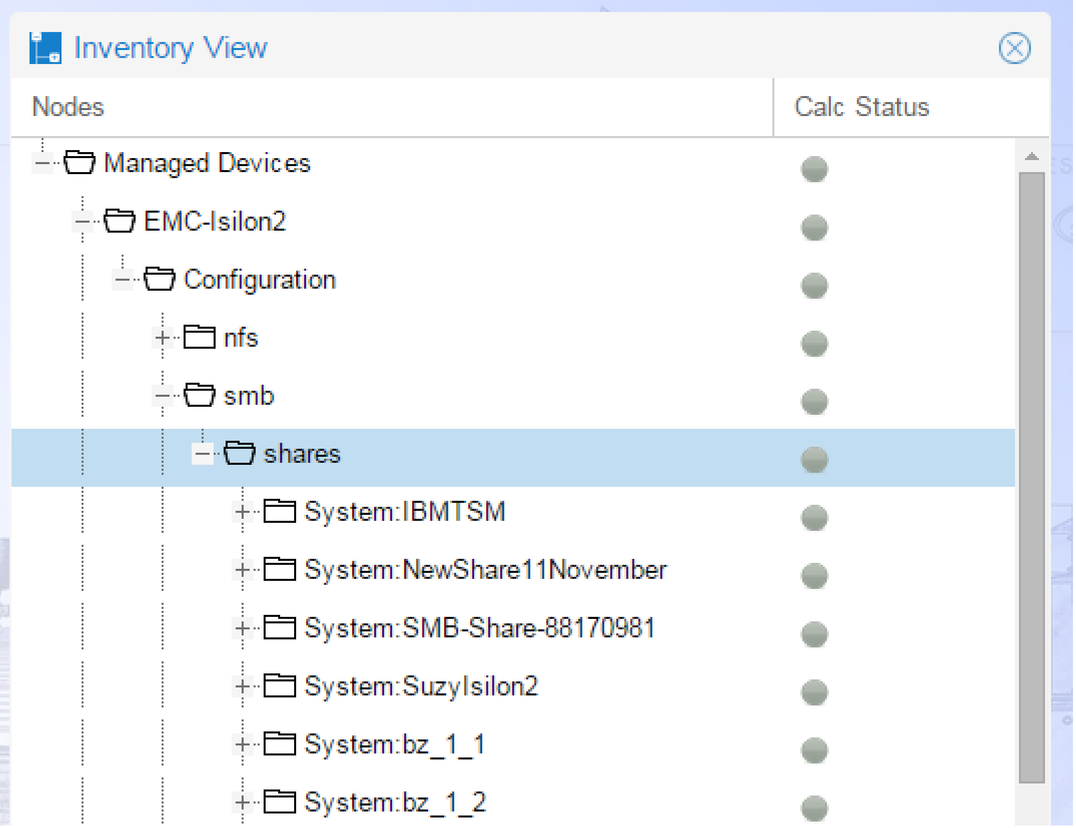

Enable Eyeglass Jobs

Once you have configured your PowerScale cluster pair and the Inventory task has completed, 3 Eyeglass Jobs are automatically created per SyncIQ Policy to replicate between the SyncIQ Policy defined source and target. In addition to the Configuration Replication Jobs, Failover Readiness Jobs are created between replicating clusters that monitor the configuration and readiness of Access Zones.

Note: These jobs are disabled by default (see admin guide on how to change default to enable via the CLI). Once enabled they will raise alarms if all configuration for Access Zones is not created or prerequisites completed.

Prerequisite for Enabling Configuration Replication

1. If you have an Active - Active Replication Topology (for data), confirm that you do not have an unsupported share or nfs alias environment described in the diagram below:

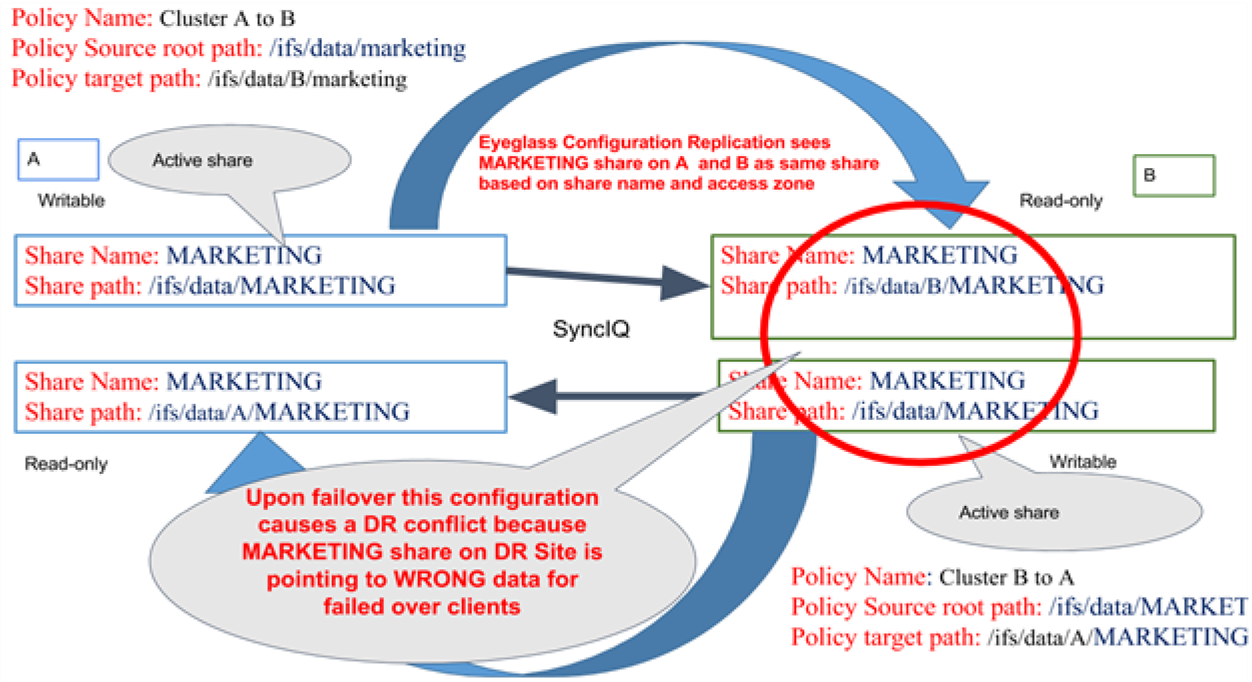

Policy Name: Cluster A to B
Policy Source root path: /ifs/data/marketing
Policy target path: /ifs/data/B/marketing
A
Writable
Active share
Eyeglass Configuration Replication sees MARKETING share on A and B as same share based on share name and access zone
B
Read-only
Share Name: MARKETING
Share path: /ifs/data/MARKETING
Share Name: MARKETING
Share path: /ifs/data/B/MARKETING
SyncIQ
Share Name: MARKETING
Share path: /ifs/data/A/MARKETING
Share Name: MARKETING
Share path: /ifs/data/MARKETING
Read-only
Writable
Active share
Upon failover this configuration causes a DR conflict because MARKETING share on DR Site is pointing to WRONG data for failed over clients
Policy Name: Cluster B to A
Policy Source root path: /ifs/data/MARKET
Policy target path: /ifs/data/A/MARKETING

1. Review Eyeglass Admin Guide Jobs description to understand what the Configuration Replication Jobs will do.
2. Review Eyeglass Admin Guide for Configuration Replication Pre-requisites
3. Review how Eyeglass determines uniqueness for configuration items and what properties are replicated.

Enable Jobs for Configuration Replication

Next step is to enable your Share, Export, NFS Alias (AUTO) Jobs for Configuration Replication. This can be done on a Job by Job basis by following these steps:

Select the Configuration Replication Job to be enabled.

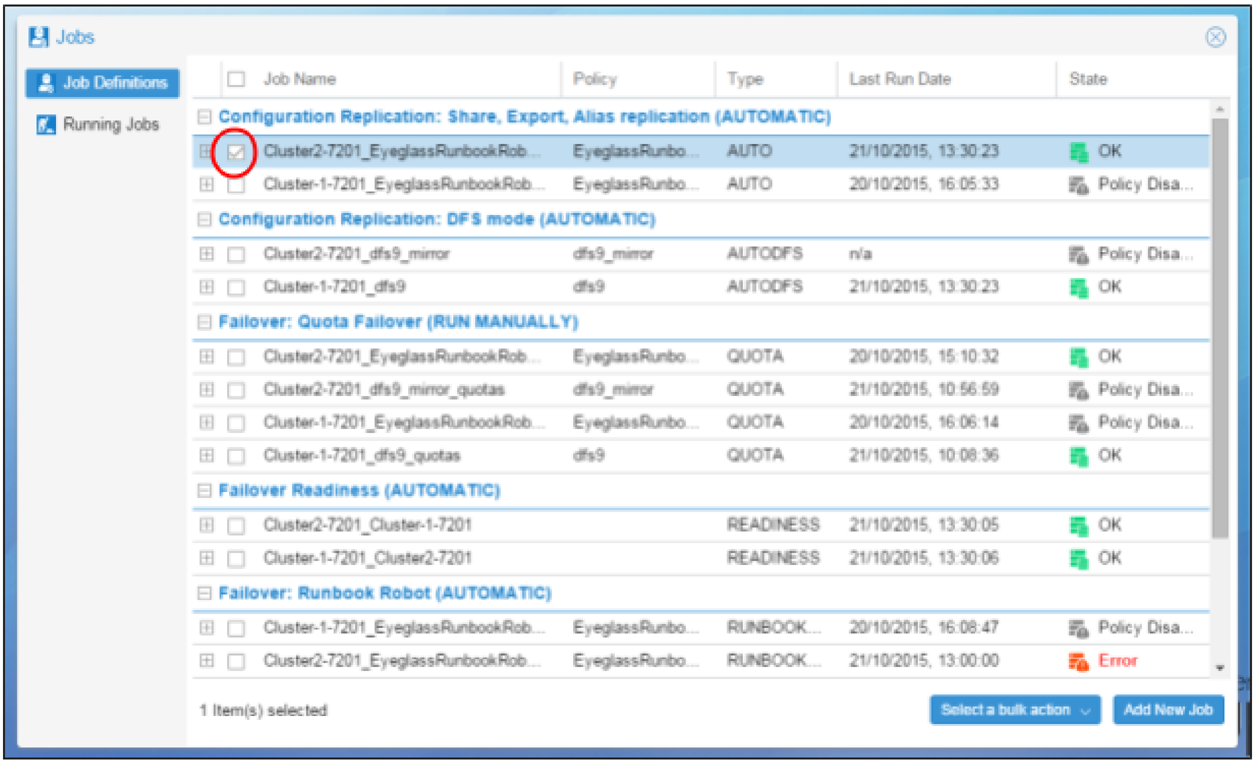

□	Job Name	Policy	Type	Last Run Date	State
	Configuration Replication: Share, Export, Alias replication (AUTOMATIC)				
☑	Cluster2-7201_EyeglassRunbookRob...	EyeglassRunbo...	AUTO	21/10/2015, 13:30:23	OK
□	Cluster-1-7201_EyeglassRunbookRob...	EyeglassRunbo...	AUTO	20/10/2015, 16:05:33	Policy Disa...
	Configuration Replication: DFS mode (AUTOMATIC)				
□	Cluster2-7201_dfs9_mirror	dfs9_mirror	AUTODFS	n/a	Policy Disa...
□	Cluster-1-7201_dfs9	dfs9	AUTODFS	21/10/2015, 13:30:23	OK
	Failover: Quota Failover (RUN MANUALLY)				
□	Cluster2-7201_EyeglassRunbookRob...	EyeglassRunbo...	QUOTA	20/10/2015, 15:10:32	OK
□	Cluster2-7201_dfs9_mirror_quotas	dfs9_mirror	QUOTA	21/10/2015, 10:56:59	Policy Disa...
□	Cluster-1-7201_EyeglassRunbookRob...	EyeglassRunbo...	QUOTA	20/10/2015, 16:06:14	Policy Disa...
□	Cluster-1-7201_dfs9_quotas	dfs9	QUOTA	21/10/2015, 10:08:36	OK
	Failover Readiness (AUTOMATIC)				
□	Cluster2-7201_Cluster-1-7201		READINESS	21/10/2015, 13:30:05	OK
□	Cluster-1-7201_Cluster2-7201		READINESS	21/10/2015, 13:30:06	OK
	Failover: Runbook Robot (AUTOMATIC)				
□	Cluster-1-7201_EyeglassRunbookRob...	EyeglassRunbo...	RUNBOOK...	20/10/2015, 16:08:47	Policy Disa...
□	Cluster2-7201_EyeglassRunbookRob...	EyeglassRunbo...	RUNBOOK...	21/10/2015, 13:00:00	Error

1 Item(s) selected

Select a bulk action and then select the Enable/Disable option.

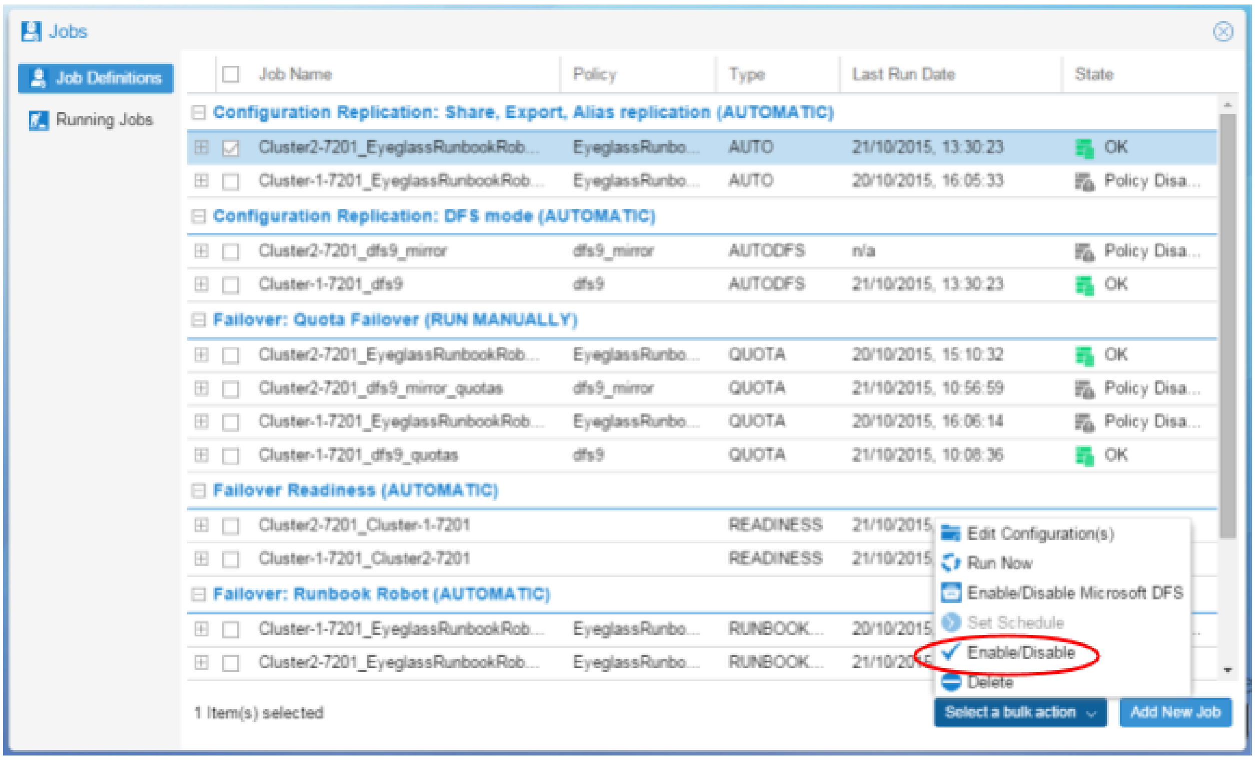

	Job Name	Policy	Type	Last Run Date	State
⊟	**Configuration Replication: Share, Export, Alias replication (AUTOMATIC)**				
⊞ ☑	Cluster2-7201_EyeglassRunbookRob...	EyeglassRunbo...	AUTO	21/10/2015, 13:30:23	OK
⊞ □	Cluster-1-7201_EyeglassRunbookRob...	EyeglassRunbo...	AUTO	20/10/2015, 16:05:33	Policy Disa...
⊟	**Configuration Replication: DFS mode (AUTOMATIC)**				
⊞ □	Cluster2-7201_dfs9_mirror	dfs9_mirror	AUTODFS	n/a	Policy Disa...
⊞ □	Cluster-1-7201_dfs9	dfs9	AUTODFS	21/10/2015, 13:30:23	OK
⊟	**Failover: Quota Failover (RUN MANUALLY)**				
⊞ □	Cluster2-7201_EyeglassRunbookRob...	EyeglassRunbo...	QUOTA	20/10/2015, 15:10:32	OK
⊞ □	Cluster2-7201_dfs9_mirror_quotas	dfs9_mirror	QUOTA	21/10/2015, 10:56:59	Policy Disa...
⊞ □	Cluster-1-7201_EyeglassRunbookRob...	EyeglassRunbo...	QUOTA	20/10/2015, 16:06:14	Policy Disa...
⊞ □	Cluster-1-7201_dfs9_quotas	dfs9	QUOTA	21/10/2015, 10:08:36	OK
⊟	**Failover Readiness (AUTOMATIC)**				
⊞ □	Cluster2-7201_Cluster-1-7201		READINESS	21/10/2015,	
⊞ □	Cluster-1-7201_Cluster2-7201		READINESS	21/10/2015,	
⊟	**Failover: Runbook Robot (AUTOMATIC)**				
⊞ □	Cluster-1-7201_EyeglassRunbookRob...	EyeglassRunbo...	RUNBOOK...	20/10/2015,	
⊞ □	Cluster2-7201_EyeglassRunbookRob...	EyeglassRunbo...	RUNBOOK...	21/10/20	

1 Item(s) selected

On the next Configuration Replication cycle, the enabled Job will be run.

Initial state for Jobs

You can change the default behavior so that these Jobs are enabled by default using the cli commands for here.

Setup Eyeglass for Email Notification

1. Configure SMTP
2. Configure Email Recipients

Configure SMTP

1. Enter the information for your email server in the **Notification Center** / **Configure SMTP** tab.

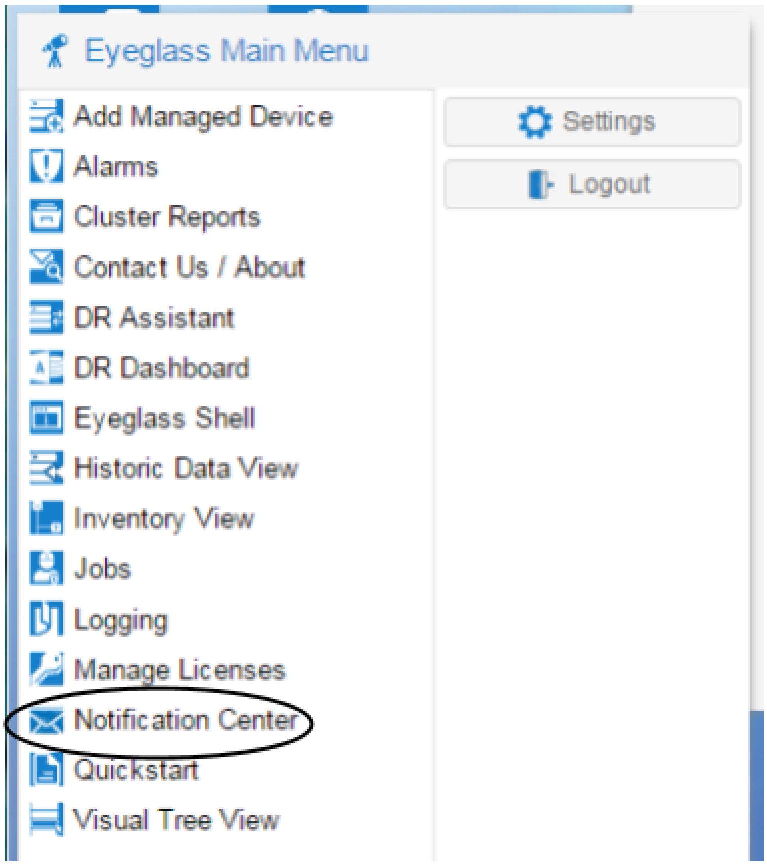

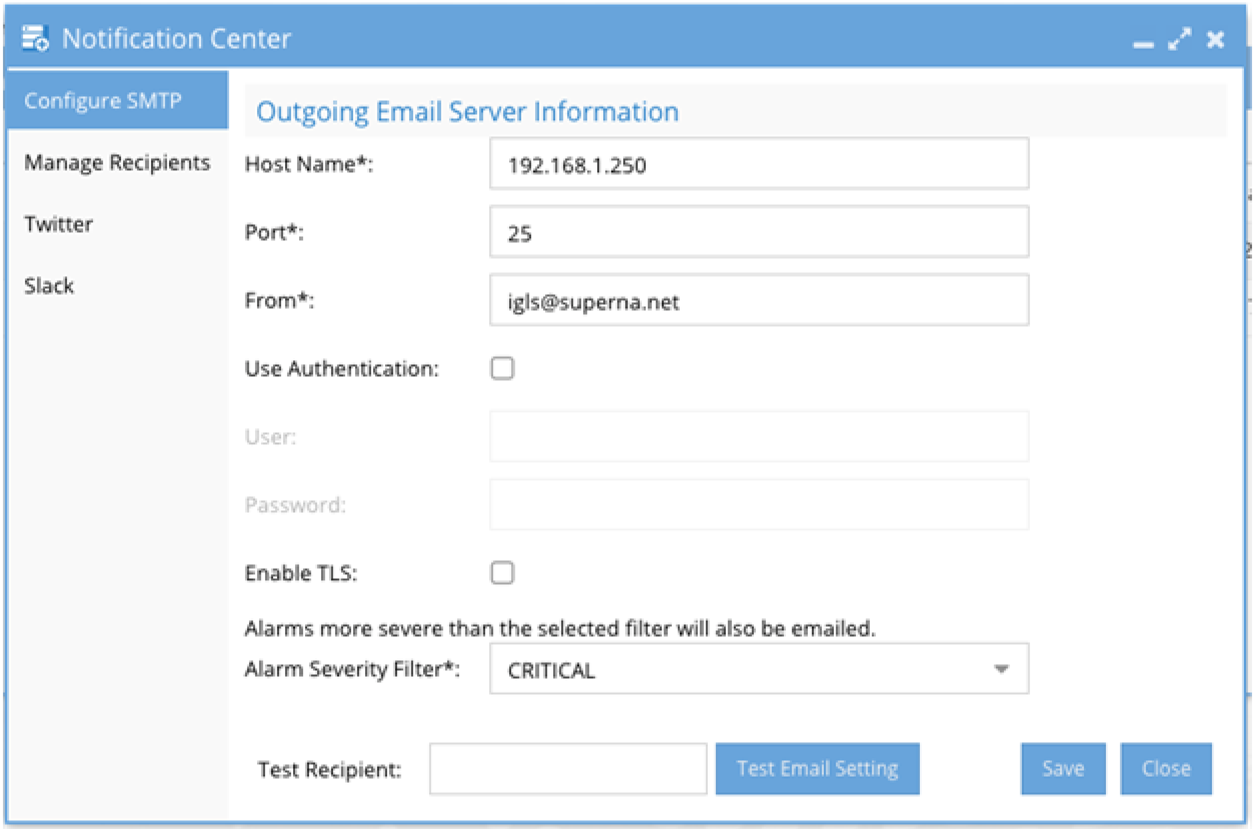

Outgoing Email Server Information

Host Name*: 192.168.1.250

Port*: 25

From*: igls@superna.net

Use Authentication: □

User:

Password:

Enable TLS: □

Alarms more severe than the selected filter will also be emailed.

Alarm Severity Filter*: CRITICAL

Test Recipient:

- Host name: Enter the host name for your email server
- Port: Enter the port which should be used for sending email

- From: Enter the email address of the sender of the email. Typically this is required to be a valid email address recognized by the email server.
- Use Authentication: Select if email server requires an authenticated login
    - User: User or email address for authentication
    - Password: Password for authentication
- Enable TLS: Select the Enable TLS check box if your email server expects TLS communication.
- Alarm Severity Filter: Select level of alarms for which you would like to receive email.

2. Use the **Test Email Setting** button to check that the email server information added is correct. If an error occurs, you will get error codes from the SMTP connection. The "no error" response indicates successful connection. If error is returned the debug response should be sent to support.superna.net.

3. **Save** your changes.

Configure Email Recipients

1. Enter the information for your email server in the **Notification Center** / **Manage Recipients** tab.
    - Email Recipient: Enter the email address that emails will be sent to.
2. Select the **Add** button.

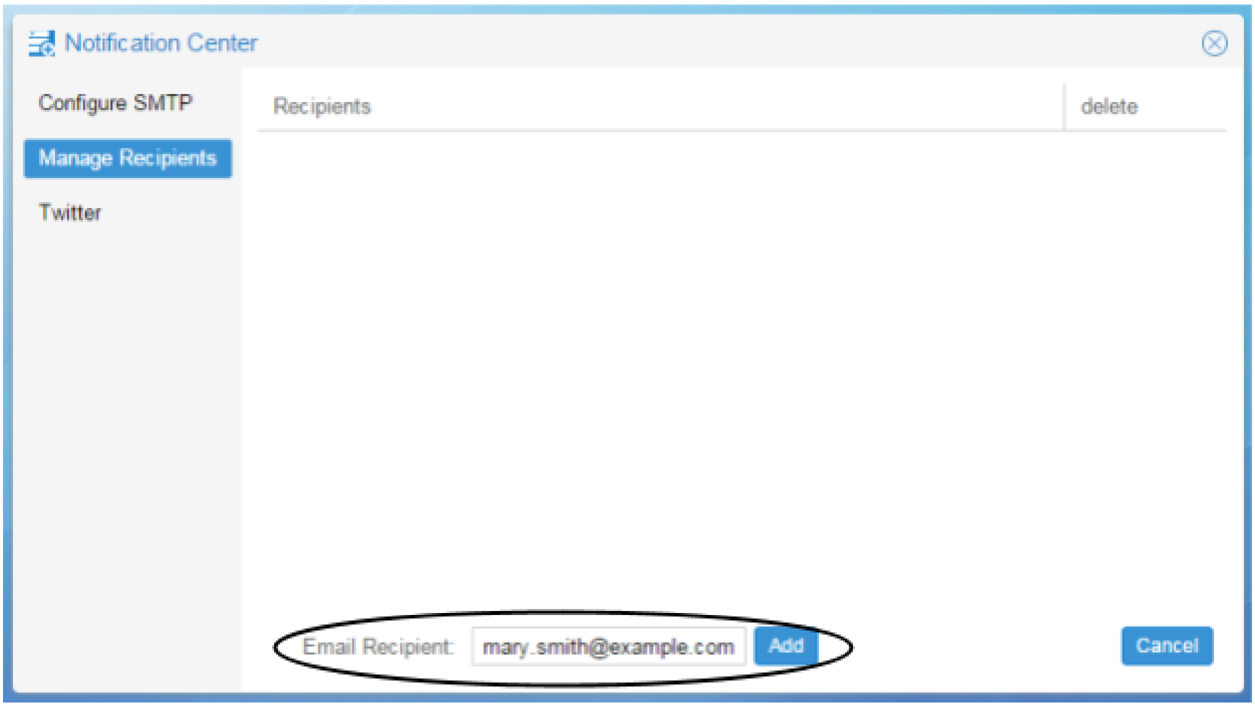

Email Recipient: mary.smith@example.com

Appendix A

Use this procedure to force update Eyeglass controlled packages.

List of files controlled by Eyeglass

1. /etc/lighttpd/conf.d/proxy.conf
2. /etc/lighttpd/lighttpd.conf
3. /etc/lighttpd/modules.conf
4. /etc/motd
5. /etc/sudoers.d/admin

These above mentioned files need to be replaced after you upgrade your RHEL packages. We strongly recommend copying the above mentioned files to a location prior to performing RHEL package upgrade so that they can be copied back safe after upgrade.

Eyeglass Upgrade Instructions for RHEL and Centos

Use these steps to upgrade an existing RHEL or Centos Eyeglass Appliance:

IMPORTANT: Before upgrading to latest Eyeglass code, please take a VMware snapshot of the Eyeglass appliance. This is the only way to roll back if upgrade fails.

Pre-requisites:

1) To download the RHEL/Centos upgrade file open a support ticket to request the download link and password.

2) You may need internet connectivity for this upgrade. If you do not have internet access from the Eyeglass appliance you need to download the extra packages. You are responsible for installing the RHEL packages since RHEL is a licensed product. The installer will detect the missing packages and prompt you on how to install.

We recommend that you run the installer to find the missing packages and proceed accordingly.

Upgrade Instructions

1. Download the file to your local machine.

2. ssh to the Eyeglass appliance and assume root user

3. Copy the file to the Eyeglass appliance (ie in the /tmp directory).

4. Make the file executable (replacing name with specific run file name)

   ```
   chmod +x eyeglass_RHEL_file_name.run
   ```

5. Run the file (replacing name with specific run file name)

```
./eyeglass_RHEL_file_name.run
```

6. Upgrade completed.

Post Upgrade Instructions

1. After upgrade, using Eyeglass CLI run the following command

```
rpm -qa | grep -i 'eyeglass' | grep '<PRE_UPGRADE_MAJOR_EYEGLASS_VERSION>' | xargs rpm -e
```

example after upgrade from 2.5.4 to 2.5.5 the command would be:

```
rpm -qa | grep -i 'eyeglass' | grep '2.5.4' | xargs rpm -e
```

2. Login to the Eyeglass Web UI.

3. Check version number by going to About/Contact

Installed Package List

eyeglass_authservice-2.5.7-21105.x86_64.rpm
eyeglass_db-2.5.7-21105.x86_64.rpm
eyeglass_deps-2.5.7-21105.rhel7.x86_64.rpm
eyeglass_licencing-2.5.7-21105.x86_64.rpm
eyeglass_logparser-2.5.7-21105.x86_64.rpm
eyeglass_pygls-2.5.7-21105.rhel7.x86_64.rpm
eyeglass_rest-2.5.7-21105.x86_64.rpm
eyeglass_sca-2.5.7-21105.x86_64.rpm

eyeglass_sera-2.5.7-21105.x86_64.rpm
eyeglass_servicebroker-2.5.7-21105.x86_64.rpm
eyeglass_ui-2.5.7-21105.x86_64.rpm

Top level OS Dependencies

```
"base":
"lighttpd",
"nfs-utils",
"nodejs",
"net-tools",
"sysstat",
"rsync",
"fail2ban"

"RHEL 7": [
"python34-pip",
"protobuf-python",
"python-setuptools",
"rsyslog"
],
"RHEL 8": [
"python3-protobuf",
"python3-setuptools",
"(syslog-ng or rsyslog)"
]
```

10. Eyeglass Probe PowerScale installation guide

Home Top

10.1. How to Install Video

Home Top

How to Install Video

10.2. System Overview

Home Top

System Overview

Eyeglass Probe for CA UIM leverages Eyeglass DR Orchestration platform software to provide a single integration point for CA UIM and native PowerScale REST API. CA UIM with Eyeglass Probe offers alarm and DR status readiness for PowerScale multi cluster monitoring of hardware, software and DR readiness.

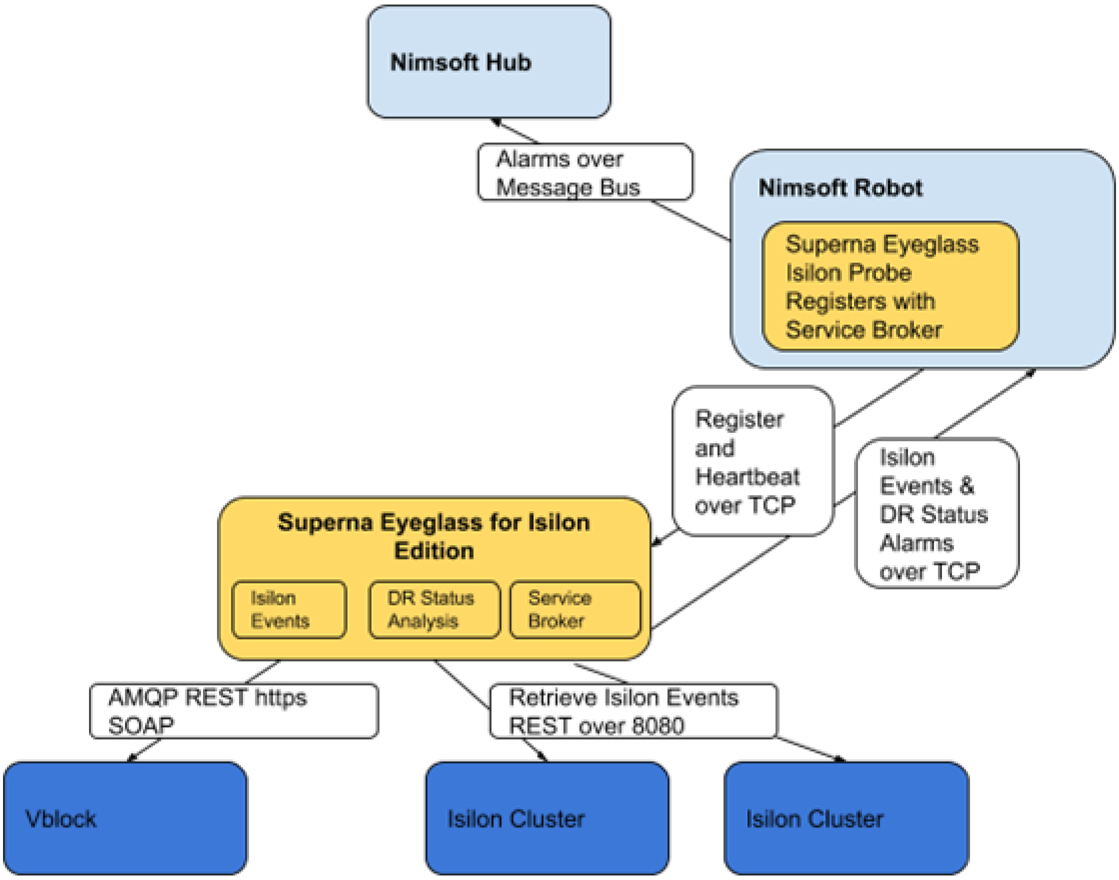
Nimsoft Hub
Alarms over Message Bus
Nimsoft Robot
Superna Eyeglass Isilon Probe Registers with Service Broker
Register and Heartbeat over TCP
Isilon Events & DR Status Alarms over TCP
Superna Eyeglass for Isilon Edition
Isilon Events
DR Status Analysis
Service Broker
AMQP REST https SOAP
Retrieve Isilon Events REST over 8080
Vblock
Isilon Cluster
Isilon Cluster

Eyeglass Probe: This probe provides inventory and alarm collection from PowerScale clusters and Vblock’s.

Eyeglass Probe Deployment and Management: Eyeglass Probe is a native UIM probe that is deployed and managed using the UIM Infrastructure Manager. The UIM Infrastructure Manager connects to an active Hub and allows you to control, configure and manage all robots and probes connected to that hub. It also provides a probe Archive. A probe can be deployed to any UIM hub once it has been imported into the Archive.

- Eyeglass Probe must be installed on a UIM Robot which has IP connectivity to the Eyeglass appliance

For additional documentation on CA UIM, please refer to support.ca.com.

© Superna LLC

10.3. System Requirements Eyeglass and Probe

Home Top

System Requirements Eyeglass and Probe

UIM

Eyeglass Probe developed using UIM SDK for Java.

Eyeglass Probe was validated using:

Table 1 UIM Validation List

UIM Component	Version
UIM	8.3, 8.4
Eyeglass appliance	1.5.4
UIM Probe	1.558
PowerScale	7.2.0.x or 7.2.1.x or 8.0.0.x
Vblock Vision IO	2.6
Vcenter	5.5 >

Eyeglass Probe

The Eyeglass Probe installation has the following System Requirements:

- Microsoft Windows Server 2008 R2 64 or greater
- Linux OS supported by UIM
- Hardware or Virtual Machine with 1 G RAM and Dual Core CPU 32 or 64 bit, 10G for the probe

The Eyeglass Probe requires that the following ports on the server be available:

Table 2 Eyeglass PowerScale appliance Requirements

Item	Minimum	Direct Documentation Link
CPU	4 x vCPU	Link
OS	ESXi > 5.0	Link
Disk	80GB	Link
NIC	1 x 10 Mbps	Link
Firewall Open Ports PowerScale cluster to Eyeglass appliance	See here	See here

Table 3 Eyeglass Probe Port Requirements for Vblock

Port	Direction	Protocol
8443	Outbound	https
5672	Outbound	AMQP
443	Outbound	https

Table 4 Eyeglass Probe Port Requirements

Port	Direction	Protocol
37356	outbound	TCP
23458 (default)	inbound	TCP

Table 4 Eyeglass Probe Robot Requirements

Item	Operating Range	Minimum
CPU	NA	1 vCPU
Memory	512MB	512MB
Thread Usage	2-10	2
Handle Usage	2-10	2
Loaded Classes	<15	NA
Heap Memory usage	< 512MB	NA

File Sizes	UIM Robot Manages log file size and rotation, no other files created	NA

10.4. Installation Overview

Home Top

Installation Overview

The installation steps for the Eyeglass Probe are:

1. Install the Eyeglass appliance
2. Add Probe license keys to the appliance
3. Add PowerScale clusters to the appliance
4. Configure Eyeglass Appliance Probe Authentication Token
5. Import Eyeglass Probe into UIM Infrastructure Manager Archive.
6. Deploy Eyeglass probe.
7. Configure the probe to connect to the eyeglass appliance with the api token
8. Send test event from PowerScale

10.5. Install Eyeglass Appliance

Home Top

Install Eyeglass Appliance

1. The appliance can be installed using the instructions here Eyeglass PowerScale Edition Quick Start Guide for Eyeglass

1. Installs appliance with ip address
2. Adds license keys
3. Adds clusters

10.6. Add PowerScale probe license keys to the appliance

Home Top

Add PowerScale probe license keys to the appliance

Overview: Install license keys, two types exist PowerScale monitoring only or DR enhanced keys available with eyeglass DR automation edition for failover and failback automation.

The DR keys will provide DR state information for replicating cluster pairs and failover status, along with RPO status and compliance.

The monitoring only keys will provide event monitoring on PowerScale clusters

Retrieve license keys

Install keys as per install guide here.

10.7. Add PowerScale clusters to the appliance

Home Top

Add PowerScale clusters to the appliance

Overview: This section is covering both DR monitoring and event monitoring license key functionality.

1. DR license key cluster addition instructions follow this guide
2. For event monitoring only login to eyeglass appliance ip address https://x.x.x.x with admin userID and default password 3y3gl4ss
3. Goto Add Managed Device
3. Enter cluster management ip address typically the smartconnect subnet ip address for HA ip address management

1. User name with minimum privileges below
2. Once done Open Inventory tree icon on the desktop and use Refresh button (Bottom right). To verify complete cluster inventory is completed.
3. If the text shows Adding.... This means it's in progress still or discovery error has occurred.   If Adding persists for more than 10 minutes, then open a support case http://support.superna.net.
4. If it looks like the screenshot above then you are done.

10.8. Configure Eyeglass Appliance Probe Authentication Token

Home Top

Configure Eyeglass Appliance Probe Authentication Token

Overview

The Eyeglass Probe needs to be configured with an API token created on the Eyeglass appliance monitoring the PowerScale clusters used for authentication . To create this token:

1. Login to Eyeglass appliance with admin user with https://x.x.x.x of the appliance
2. Open eyeglass main menu to select **Eyeglass REST API**. Follow screenshots below to complete token creation steps needed to configure the Probe to authenticate to eyeglass appliance

3.

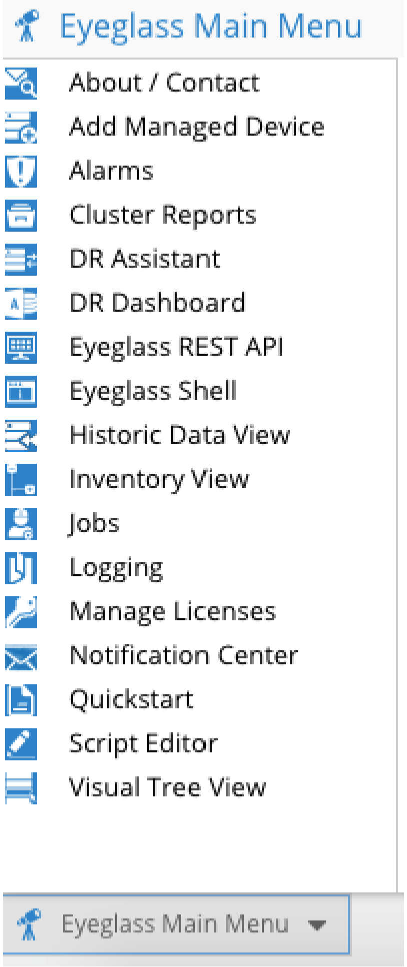

4.

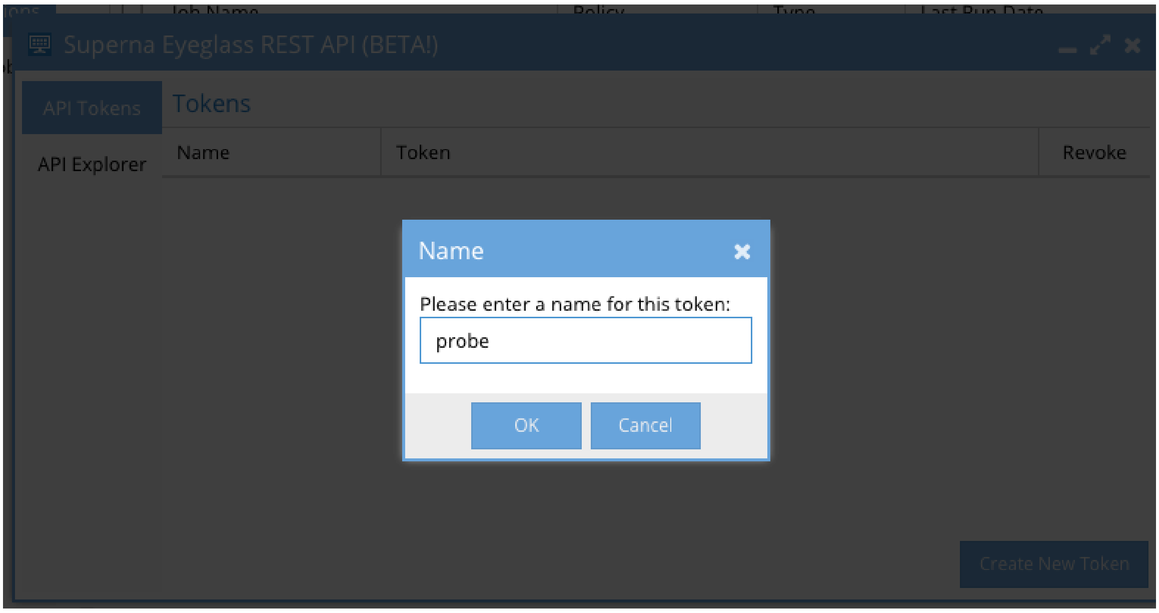

Please enter a name for this token: probe

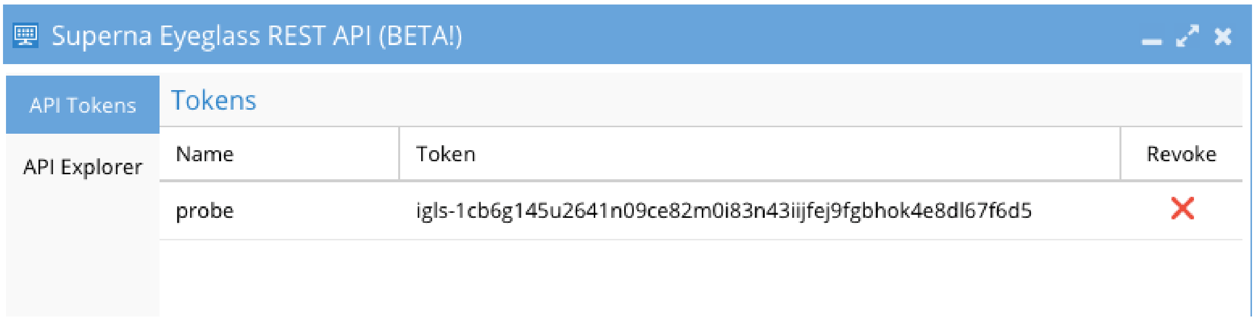

Tokens

Name	Token	Revoke
probe	igls-1cb6g145u2641n09ce82m0i83n43iijfej9fgbhok4e8dl67f6d5	✕

5. Done

10.9. Configure Eyeglass Firewall for Probe Service Broker Connection

Home Top

Configure Eyeglass Firewall for Probe Service Broker Connection

Overview:

The Probe running on UIM Robot over the network requires a firewall port opened from the probe to Eyeglass (one direction port opened). Note if using converged deployment model with co-resident Robot + probe then the firewall must be modified to open default HUB to Robot port of 48000.

1. Login to Eyeglass appliance via ssh
2. Type sudo -s
3. Enter admin password
4. Type yast
5. Arrow down to security and users and arrow to firewall enter
6. Arrow to Allowed Services
7. Tab to Advanced option

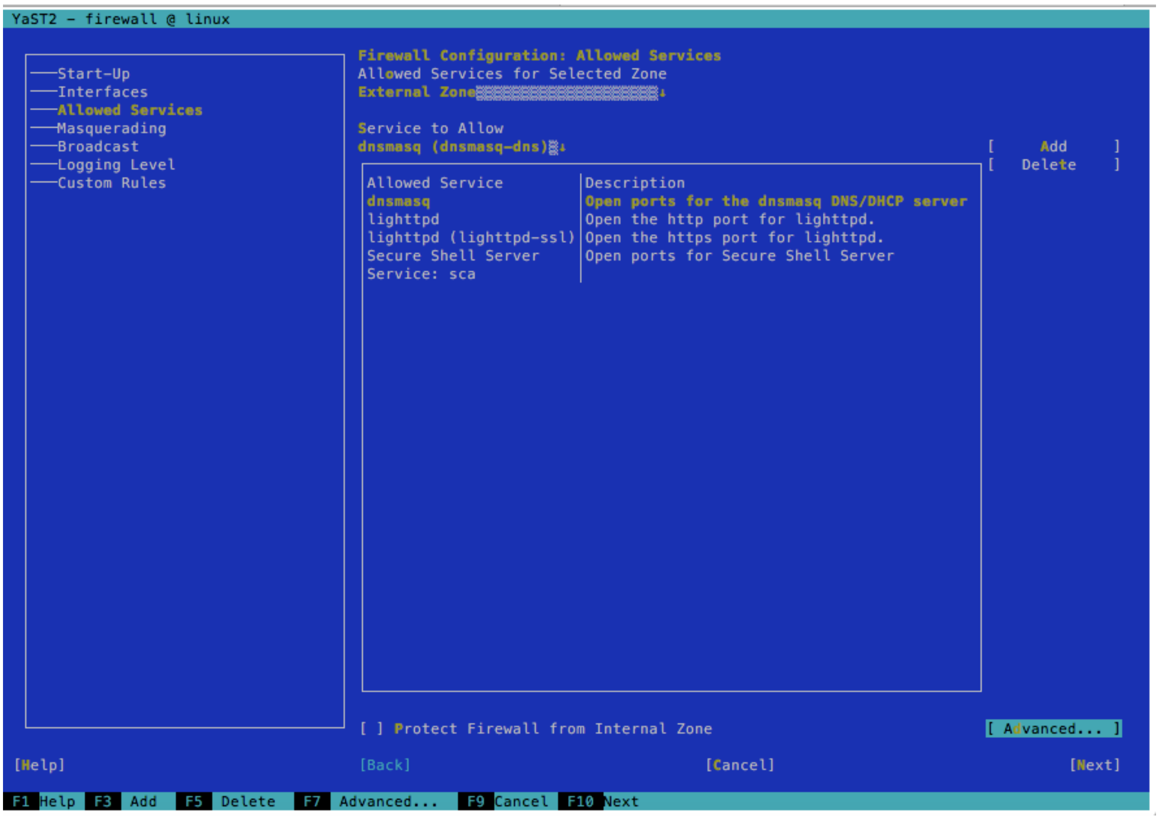

```
YaST2 - firewall @ linux

──Start-Up                    Firewall Configuration: Allowed Services
──Interfaces                  Allowed Services for Selected Zone
──Allowed Services            External Zone
──Masquerading
──Broadcast                   Service to Allow
──Logging Level               dnsmasq (dnsmasq-dns)                          [    Add    ]
──Custom Rules                                                               [  Delete   ]
                              Allowed Service          Description
                              dnsmasq                  Open ports for the dnsmasq DNS/DHCP server
                              lighttpd                 Open the http port for lighttpd.
                              lighttpd (lighttpd-ssl)  Open the https port for lighttpd.
                              Secure Shell Server      Open ports for Secure Shell Server
                              Service: sca

                              [ ] Protect Firewall from Internal Zone        [ Advanced... ]

[Help]                        [Back]               [Cancel]                  [Next]
F1 Help  F3 Add  F5 Delete  F7 Advanced...  F9 Cancel  F10 Next
```

1.

8. Enter ports as shown below with the following notes:

1. Enter 37356 for TCP when Service Broker port for UIM probe access is required

2. Also enter port 48000 if you have installed UIM robot onto the eyeglass appliance for converged deployment model where the UIM robot and probe will run on the eyeglass appliance. Follow Linux Robot install instructions to install a Robot on Eyeglass appliance.

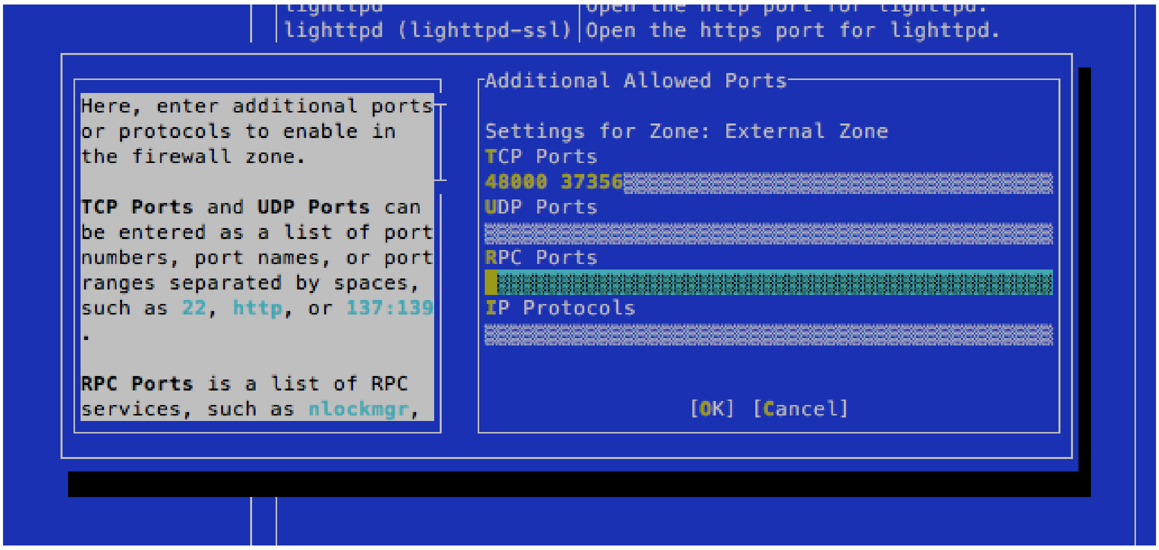

```
             lighttpd                 Open the http port for lighttpd.
             lighttpd (lighttpd-ssl)  Open the https port for lighttpd.

Here, enter additional ports   Additional Allowed Ports
or protocols to enable in
the firewall zone.             Settings for Zone: External Zone
                               TCP Ports
TCP Ports and UDP Ports can    48000 37356
be entered as a list of port   UDP Ports
numbers, port names, or port
ranges separated by spaces,    RPC Ports
such as 22, http, or 137:139
.                              IP Protocols

RPC Ports is a list of RPC
services, such as nlockmgr,              [OK]  [Cancel]
```

3.

4. The above image shows both ports setup on the appliance
5. Click ok, then exit yast settings to save.
6. Done

10.10. Import Eyeglass Probe to UIM Infrastructure Manager Archive

Home Top

Import Eyeglass Probe to UIM Infrastructure Manager Archive

To import Eyeglass Probe to the UIM Infrastructure Manager Archive:

1. Login to the UIM Infrastructure Manager.
2. Expand the Archive folder and select the hub which is running the Robot where the probe will be deployed.

Hint: For the Eyeglass Probe, select the hub which has a Robot that has IP connectivity to the Eyeglass appliance.

3. The list of existing probes archived on this hub is displayed. Click on the right mouse button in this main window pane.
4. Select the **Import** option in the menu that opens.

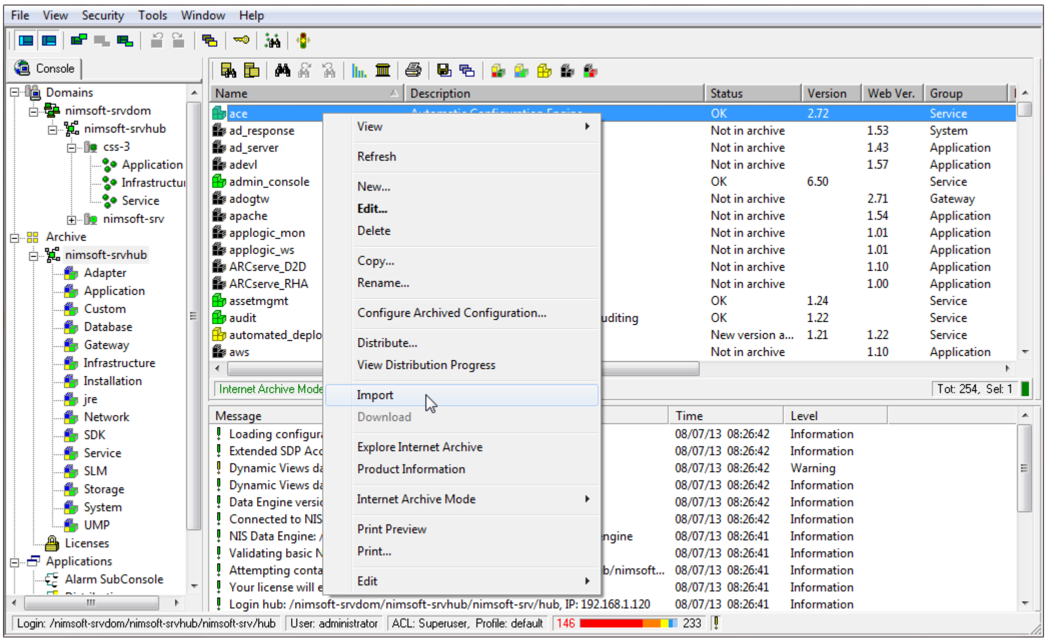

Name	Description	Status	Version	Web Ver.	Group
ace	Automatic Configuration Engine	OK	2.72		Service
ad_response		Not in archive		1.53	System
ad_server		Not in archive		1.43	Application
adevl		Not in archive		1.57	Application
admin_console		OK	6.50		Service
adogtw		Not in archive		2.71	Gateway
apache		Not in archive		1.54	Application
applogic_mon		Not in archive		1.01	Application
applogic_ws		Not in archive		1.01	Application
ARCserve_D2D		Not in archive		1.10	Application
ARCserve_RHA		Not in archive		1.00	Application
assetmgmt		OK	1.24		Service
audit	uditing	OK	1.22		Service
automated_deplo		New version a...	1.21	1.22	Service
aws		Not in archive		1.10	Application

Message	Time	Level
Loading configura	08/07/13 08:26:42	Information
Extended SDP Acc	08/07/13 08:26:42	Information
Dynamic Views da	08/07/13 08:26:42	Warning
Dynamic Views da	08/07/13 08:26:42	Information
Data Engine versio	08/07/13 08:26:42	Information
Connected to NIS	08/07/13 08:26:42	Information
NIS Data Engine: / ngine	08/07/13 08:26:41	Information
Validating basic N	08/07/13 08:26:41	Information
Attempting conta b/nimsoft...	08/07/13 08:26:41	Information
Your license will e	08/07/13 08:26:41	Information
Login hub: /nimsoft-srvdom/nimsoft-srvhub/nimsoft-srv/hub, IP: 192.168.1.120	08/07/13 08:26:41	Information

Login: /nimsoft-srvdom/nimsoft-srvhub/nimsoft-srv/hub | User: administrator | ACL: Superuser, Profile: default | 146 | 233

Figure 2 UIM Infrastructure Manager Archive Import Package

5. Select the probe package that you would like to add to the archive and then select the **Open** button.

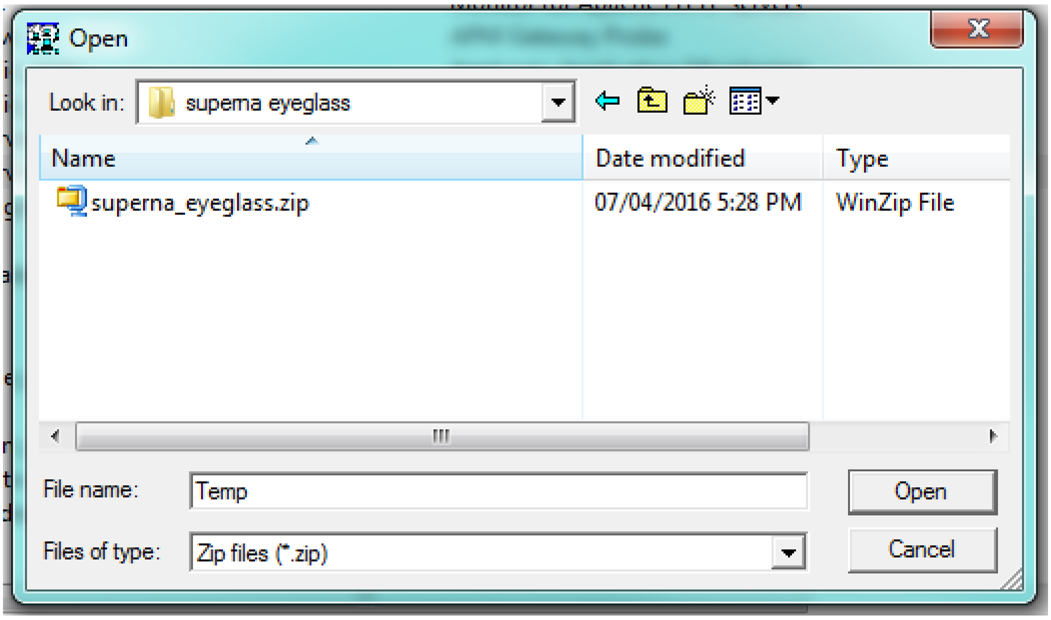

Look in: supema eyeglass

Name	Date modified	Type
superna_eyeglass.zip	07/04/2016 5:28 PM	WinZip File

File name: Temp

Files of type: Zip files (*.zip)

Figure 3 UIM Infrastructure Manager Archive Package Selection

1. The package(s) are added to your archive and appear in the main window pane for the hub.

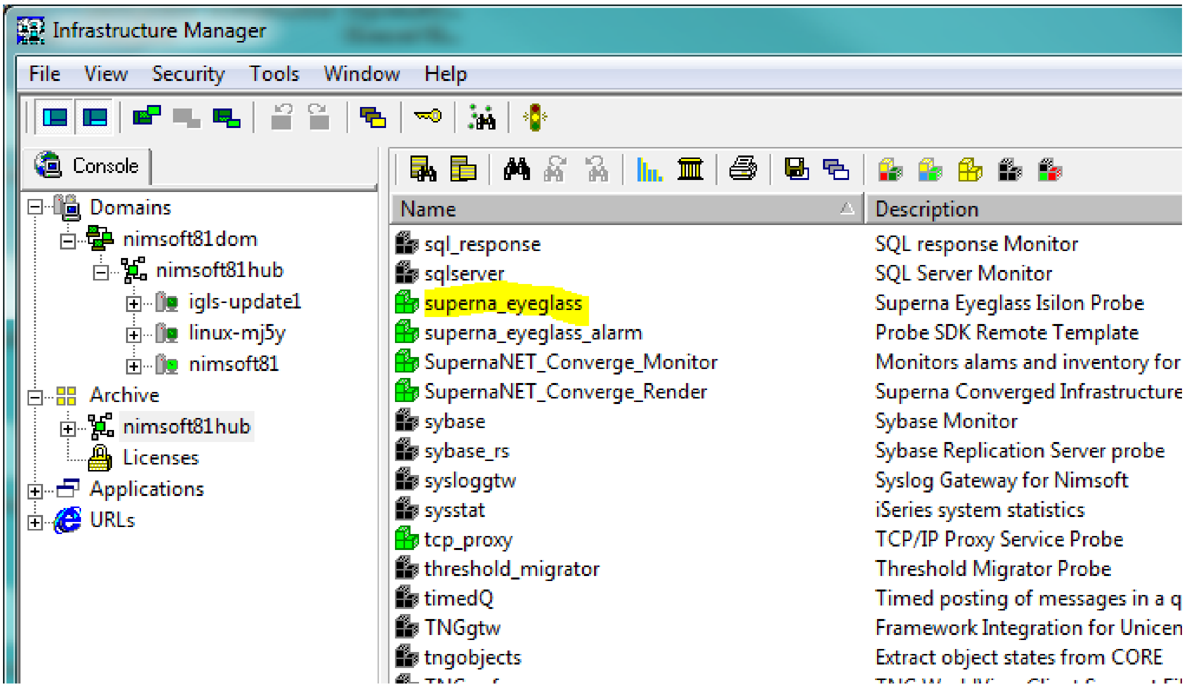

Name	Description
sql_response	SQL response Monitor
sqlserver	SQL Server Monitor
superna_eyeglass	Superna Eyeglass Isilon Probe
superna_eyeglass_alarm	Probe SDK Remote Template
SupernaNET_Converge_Monitor	Monitors alams and inventory for
SupernaNET_Converge_Render	Superna Converged Infrastructure
sybase	Sybase Monitor
sybase_rs	Sybase Replication Server probe
sysloggtw	Syslog Gateway for Nimsoft
sysstat	iSeries system statistics
tcp_proxy	TCP/IP Proxy Service Probe
threshold_migrator	Threshold Migrator Probe
timedQ	Timed posting of messages in a q
TNGgtw	Framework Integration for Unicen
tngobjects	Extract object states from CORE

Figure 4 UIM Infrastructure Manager Archive Package List

10.11. Deploy Eyeglass Probe

Home Top

Deploy Eyeglass Probe

Prerequisites

There are no prerequisite for installing Eyeglass Probe.

Deploy Eyeglass Probe

To deploy the Eyeglass Probe on a UIM Robot:

1. Login to the UIM Infrastructure Manager.
2. Select the Archive where the probe was imported.
3. Expand the Domains and hub folders until you can see the Robot where the probe will be deployed.

Hint: For the Eyeglass Probe, the Robot where it is deployed must have IP connectivity to the Eyeglass appliance.

4. Drag the Eyeglass Probe from the package archive and drop it onto the Robot.

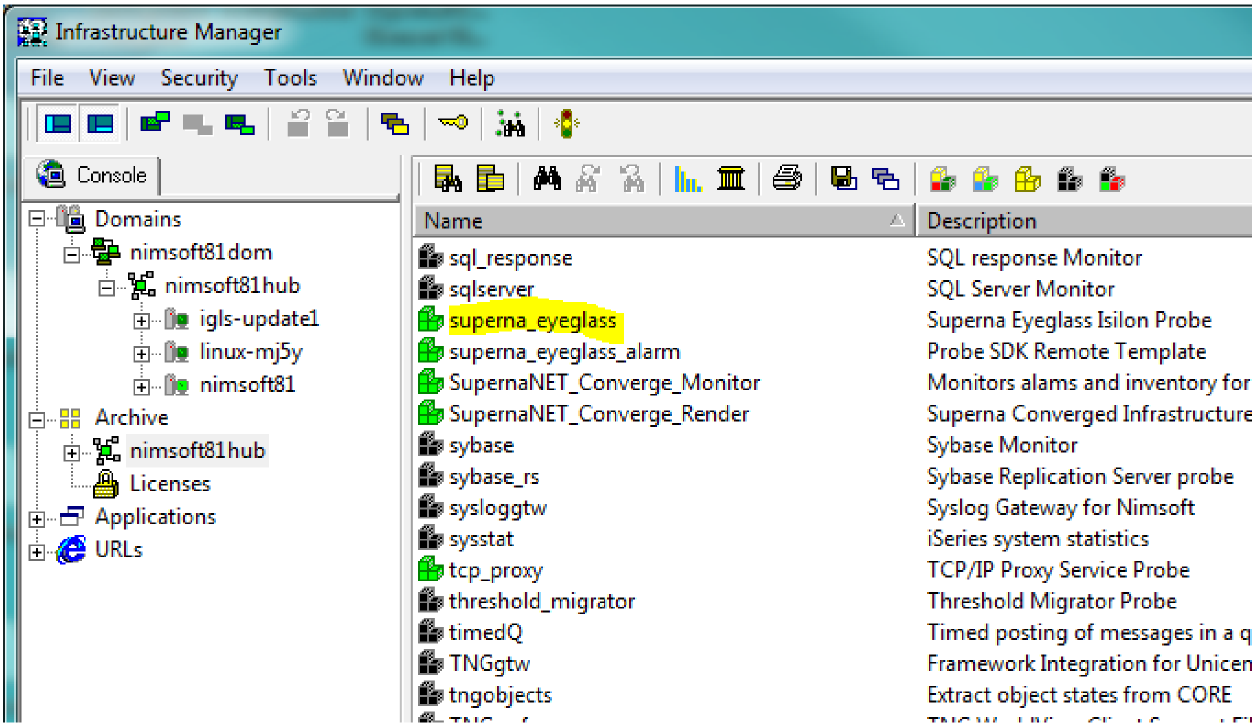

Name	Description
sql_response	SQL response Monitor
sqlserver	SQL Server Monitor
superna_eyeglass	Superna Eyeglass Isilon Probe
superna_eyeglass_alarm	Probe SDK Remote Template
SupernaNET_Converge_Monitor	Monitors alams and inventory for
SupernaNET_Converge_Render	Superna Converged Infrastructure
sybase	Sybase Monitor
sybase_rs	Sybase Replication Server probe
sysloggtw	Syslog Gateway for Nimsoft
sysstat	iSeries system statistics
tcp_proxy	TCP/IP Proxy Service Probe
threshold_migrator	Threshold Migrator Probe
timedQ	Timed posting of messages in a q
TNGgtw	Framework Integration for Unicen
tngobjects	Extract object states from CORE

Figure 5 UIM Infrastructure Manager Archive Package List

5. The Eyeglass Probe is automatically installed but not started until configured.

1. 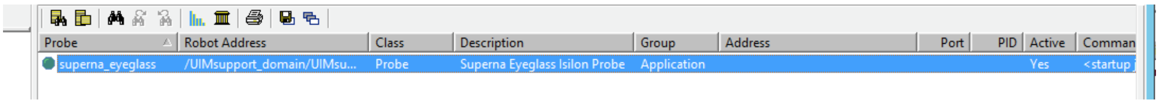

Probe	Robot Address	Class	Description	Group	Address	Port	PID	Active	Comman
superna_eyeglass	/UIMsupport_domain/UIMsu...	Probe	Superna Eyeglass Isilon Probe	Application				Yes	<startup

6. Configure probe connection to eyeglass appliance from the admin console interface by selecting the **Configure...** option

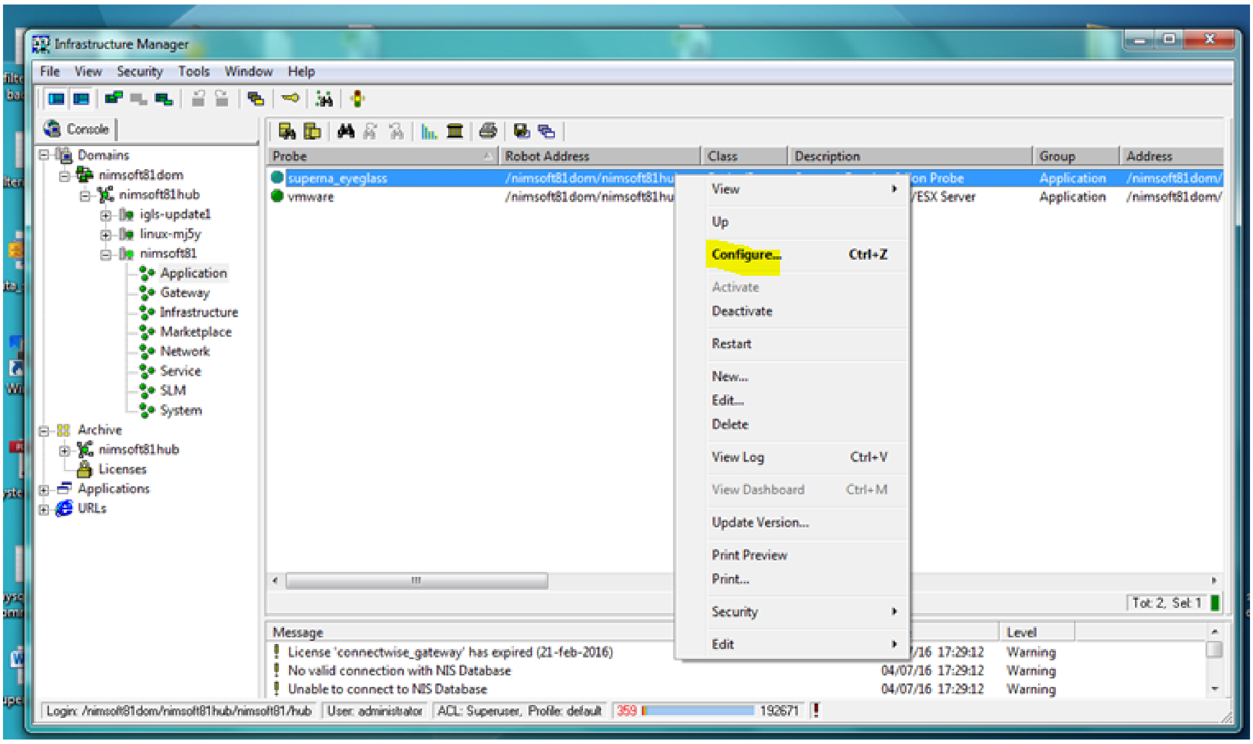

Probe	Robot Address	Class	Description	Group	Address
superna_eyeglass	/nimsoft81dom/nimsoft81hu		on Probe	Application	/nimsoft81dom/
vmware	/nimsoft81dom/nimsoft81hu		/ESX Server	Application	/nimsoft81dom/

Message		Level
License 'connectwise_gateway' has expired (21-feb-2016)	/16 17:29:12	Warning
No valid connection with NIS Database	04/07/16 17:29:12	Warning
Unable to connect to NIS Database	04/07/16 17:29:12	Warning

7. Add the information required for the superna eyeglass probe to connect to the superna eyeglass appliance by selecting the **service_broker** folder in the Raw Configure window:

ipaddress: Enter the IP address of the superna eyeglass appliance for the probe to connect

port: Enter 37356 (This is the port on which the superna eyeglass is expecting the probe to connect on and should not be changed)

Example:

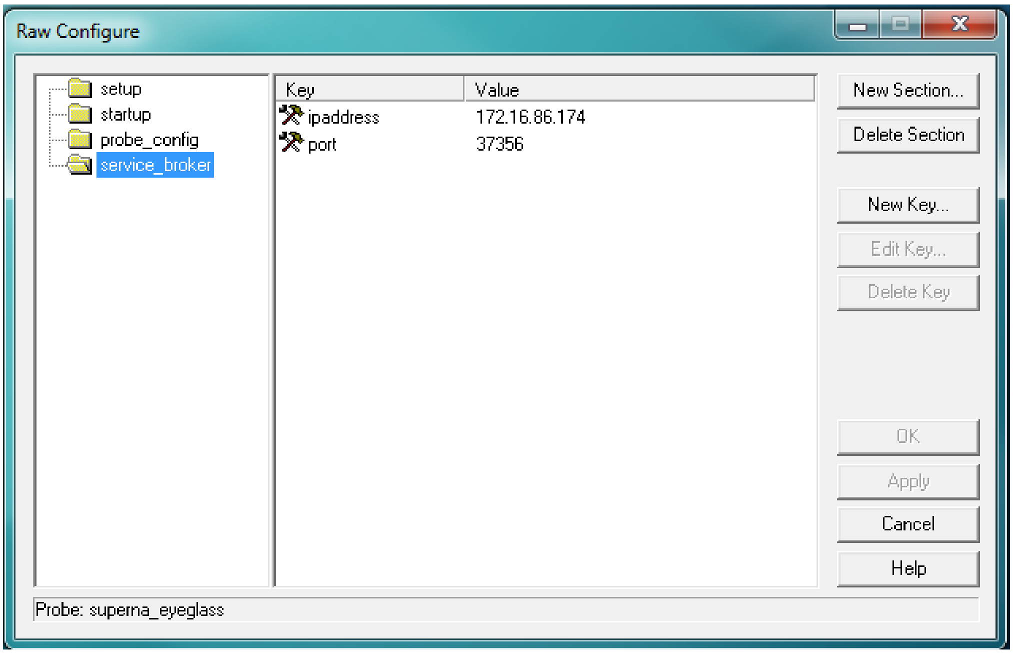

Key	Value
ipaddress	172.16.86.174
port	37356

Probe: superna_eyeglass

8. Add the information required for the superna eyeglass appliance to be able to connect to the superna eyeglass probe in order to pass the alarm and event information by selecting the **probe_config** folder in the Raw Configure window:

ipaddress: Enter the IP address of the robot on which the superna eyeglass probe is deployed (default provided but can be overridden)

port: Enter the port on which the probe will be listening for notifications from the superna eyeglass appliance (default provided but can be overridden - port must be outside the port range allocated to Nimsoft)

token: Enter the authentication token generated in previous step **Configure Eyeglass Appliance Probe Authentication Token**

Example:

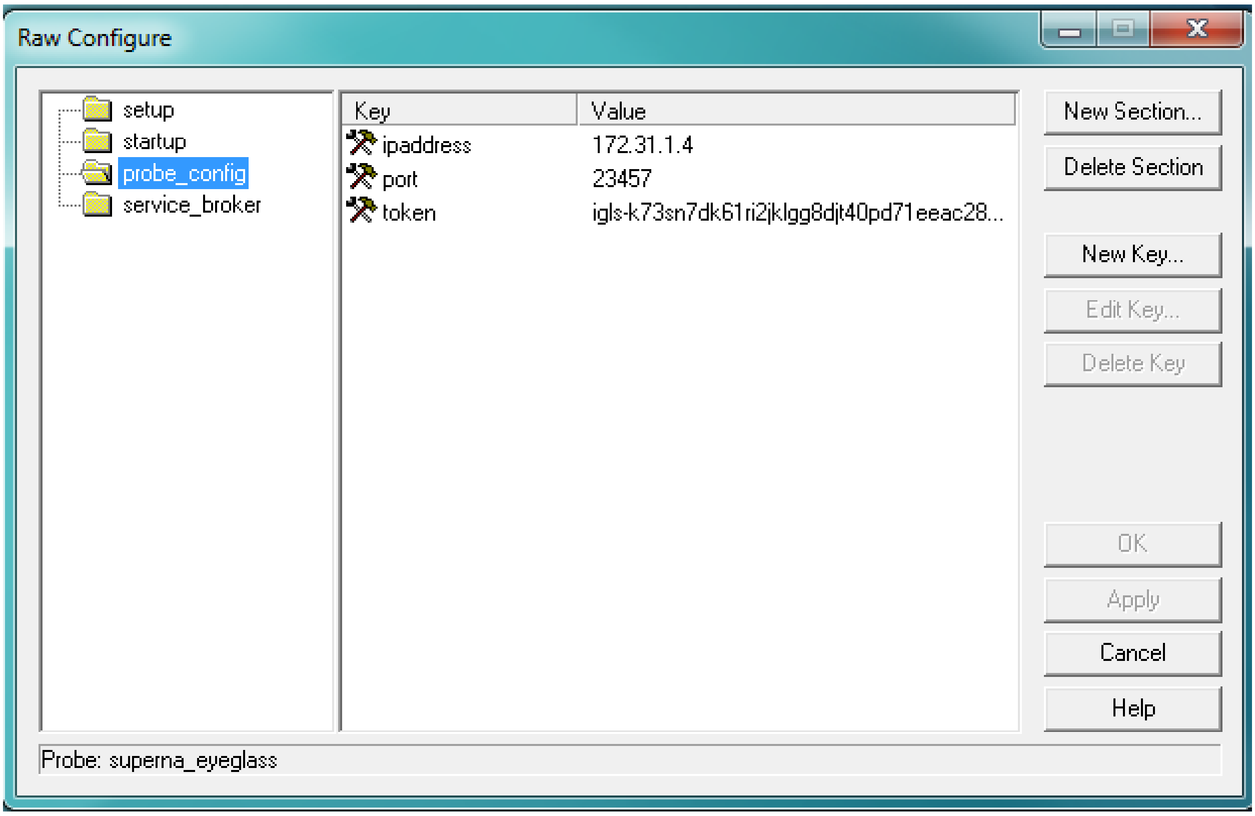

Key	Value
ipaddress	172.31.1.4
port	23457
token	igls-k73sn7dk61ri2jklgg8djt40pd71eeac28...

Probe: superna_eyeglass

9. After applying the configuration and click ok, the auto probe start should start the probe and attempt a connection to eyeglass appliance.

10. To verify the probe deployment, select the Robot where the probe was deployed. You should see the Eyeglass Probe in the list that is displayed with “Yes” in the Active column.

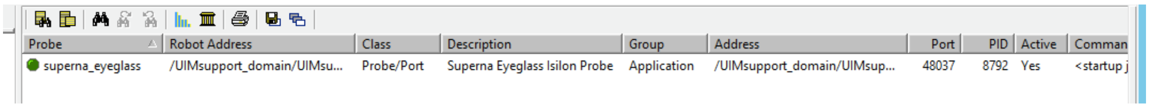

Probe	Robot Address	Class	Description	Group	Address	Port	PID	Active	Comman
superna_eyeglass	/UIMsupport_domain/UIMsu...	Probe/Port	Superna Eyeglass Isilon Probe	Application	/UIMsupport_domain/UIMsup...	48037	8792	Yes	<startup j

Figure 6 Eyeglass Probe in the UIM Infrastructure Manager Robot Probe List

7. If the probe does not start and show green with a process ID and port listed in UIM. Then proceed to troubleshooting probe startup section in Troubleshooting section

1. Its also recommended the eyeglass appliance service broker lists the probe connection. This can be verified by logging into eyeglass appliance as was done for the add cluster steps.
1. Then open Service Broker Icon on the desktop
2. The login registration from entered during probe configuration should be listed with Active state which means a TCP connection is setup with the probe.
3. If the Probe loses contact with the eyeglass appliance due to network issue the status will change to inactive.
4. 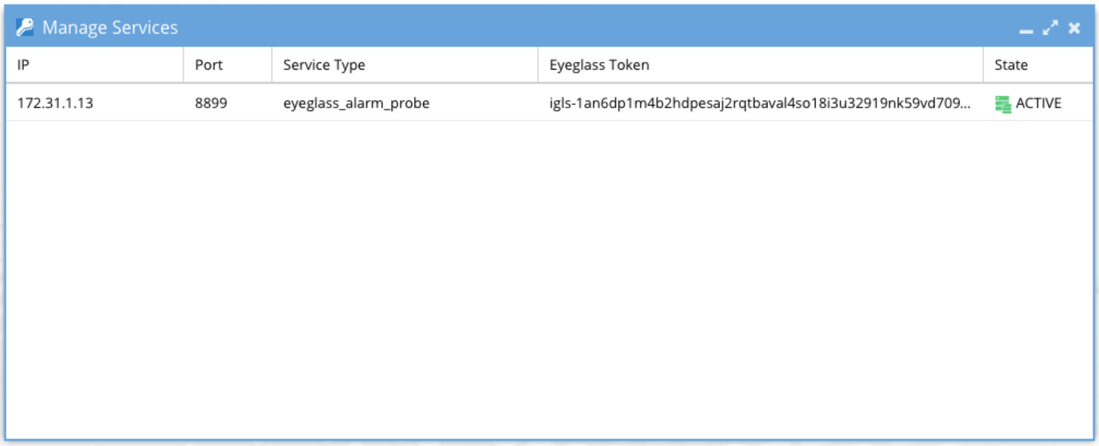

IP	Port	Service Type	Eyeglass Token	State
172.31.1.13	8899	eyeglass_alarm_probe	igls-1an6dp1m4b2hdpesaj2rqtbaval4so18i3u32919nk59vd709...	ACTIVE

5. 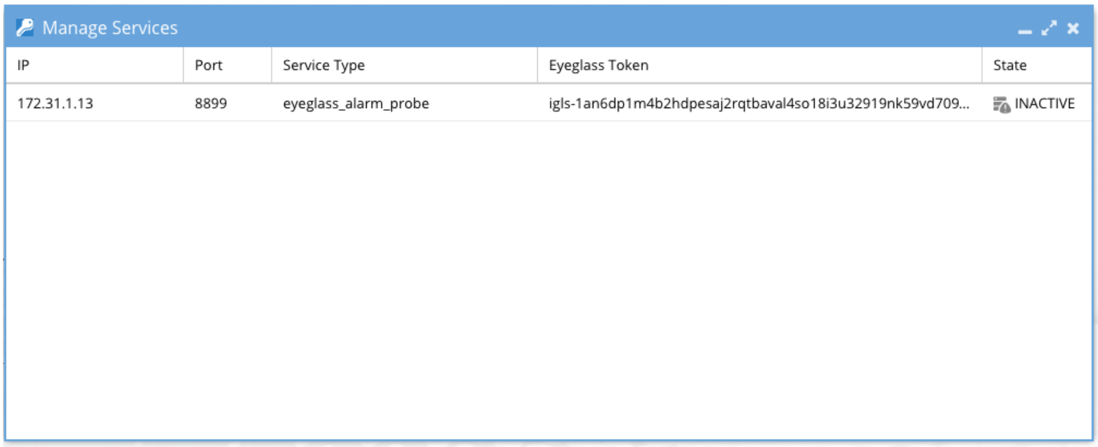

IP	Port	Service Type	Eyeglass Token	State
172.31.1.13	8899	eyeglass_alarm_probe	igls-1an6dp1m4b2hdpesaj2rqtbaval4so18i3u32919nk59vd709...	INACTIVE

8. Once this configuration is completed, the superna eyeglas probe will be ready to receive alarms from eyeglass appliance devices monitored once configured in eyeglass.

10.12. Configure Eyeglass Probe UIM Group of PowerScale clusters for monitoring

Home Top

Configure Eyeglass Probe UIM Group of PowerScale clusters for monitoring

To configure the Eyeglass Probe clusters managed by Eyeglass displayed in a group:

1. Log in as an administrator to the UIM UMP.
2. The UIM UMP home page appears.
3. Create (Dynamic folder of discovered Clusters by finding device names of Eyeglass appliance ip address)
4. This will place clusters into a group for all managed PowerScale clusters

10.13. Eyeglass Probe Administration

Home Top

Eyeglass Probe Administration

Overview

This chapter presents the following topics: Use the image below for all actions

1. Start the superna_Eyeglass probe (activate)

1. Stop the superna_Eyeglass probe (deactivate)

2. Restart the superna_Eyeglass

3. Delete the superna_Eyeglass probe

4. View debug log for startup errors

5. 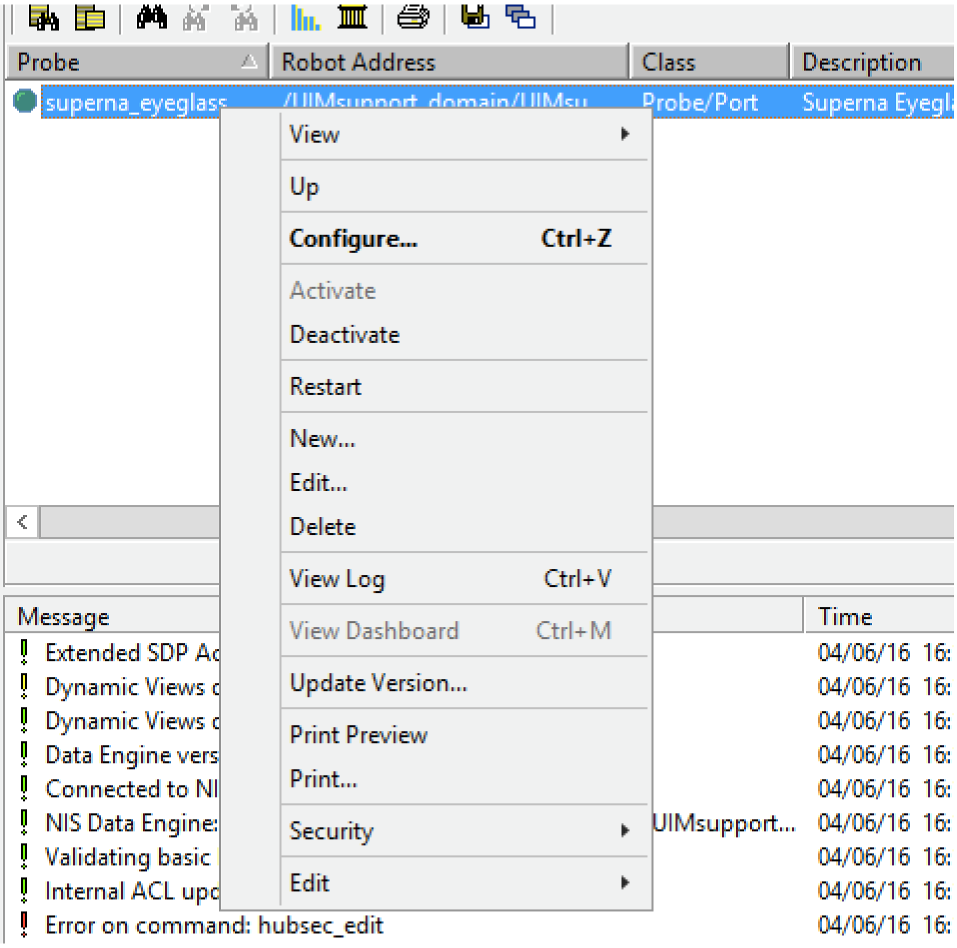

Probe	Robot Address	Class	Description
superna_eyeglass	/UIMsupport_domain/UIMsu	Probe/Port	Superna Eyegl

```
Message                                   Time
Extended SDP Ac                           04/06/16 16:
Dynamic Views c                           04/06/16 16:
Dynamic Views c                           04/06/16 16:
Data Engine vers                          04/06/16 16:
Connected to NI                           04/06/16 16:
NIS Data Engine:          UIMsupport...   04/06/16 16:
Validating basic                          04/06/16 16:
Internal ACL upc                          04/06/16 16:
Error on command: hubsec_edit             04/06/16 16:
```

10.14. Troubleshooting section

Home Top

Troubleshooting section

If the probe fails to start follow these steps

1. Check Service Broker registration listing in eyeglass appliance (see steps above for checking status) should show active

1. 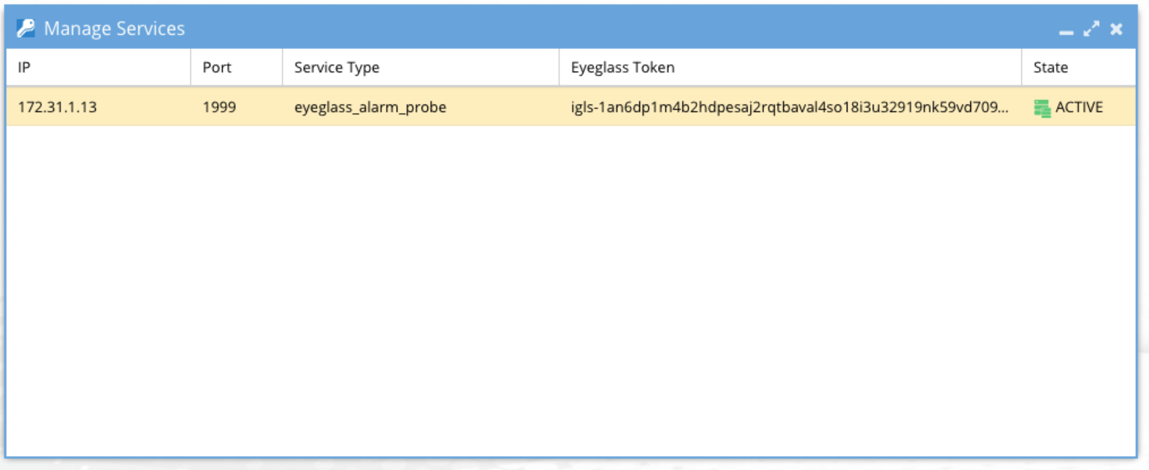

IP	Port	Service Type	Eyeglass Token	State
172.31.1.13	1999	eyeglass_alarm_probe	igls-1an6dp1m4b2hdpesaj2rqtbaval4so18i3u32919nk59vd709...	ACTIVE

2. View UIM Probe log

1. Select probe in UIM

2. Right click for menu and select view log and raise support case or resolve the error from the message

3. 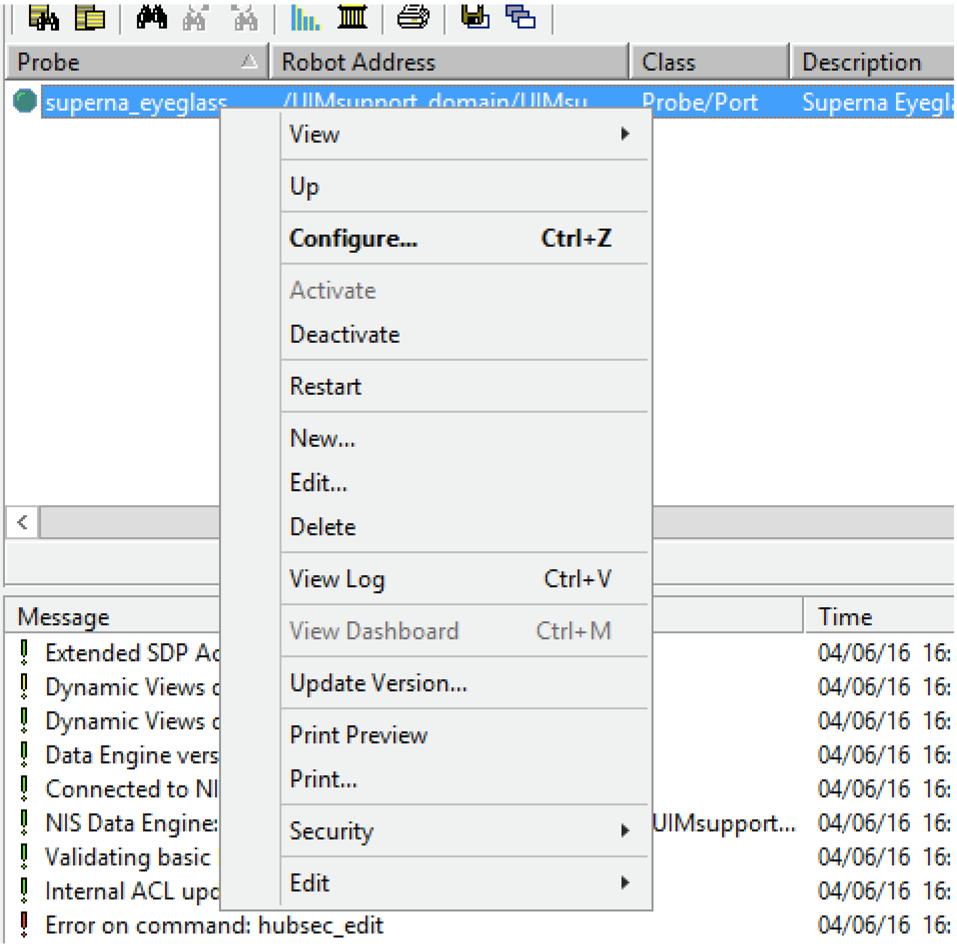

Probe	Robot Address	Class	Description
superna_eyeglass	/UIMsupport_domain/UIMsu	Probe/Port	Superna Eyegl

```
Message                                                   Time
Extended SDP Ac                                           04/06/16 16:
Dynamic Views c                                           04/06/16 16:
Dynamic Views c                                           04/06/16 16:
Data Engine vers                                          04/06/16 16:
Connected to NI                                           04/06/16 16:
NIS Data Engine:                          UIMsupport...   04/06/16 16:
Validating basic                                          04/06/16 16:
Internal ACL upc                                          04/06/16 16:
Error on command: hubsec_edit                             04/06/16 16:
```

3. Restart the broker service on eyeglass

1. Login via ssh as admin to eyeglass appliance

2. Sudo -s

3. Enter admin password

4. systemctl restart iglsservicebroker

10.15. Test Alarm/Event flow to with PowerScale Devices UIM

Home Top

Test Alarm/Event flow to with PowerScale Devices UIM

1. Login to OneFS UI goto Dashboard, Events, Notification Settings
2. Click Send test event
3. Open UIM console alarm view and verify the test event appears **NOTE: Can take up to 1 minute to appear**.

10.16. Eyeglass Probe Technical Support Playbook

Home Top

Eyeglass Probe Technical Support Playbook

Eyeglass Probe Capabilities and Components Monitored

Eyeglass Probe for CA UIM leverages PowerScale REST API and SSH to monitor PowerScale clusters and provide DR status for SyncIQ replication:

Eyeglass Probe: This probe provides inventory, state and alarm collection from individual clusters and replicating pairs of clusters with SyncIQ replication.

This subject of this document is the Eyeglass Probe.

The primary capabilities of the Eyeglass Probe are to:

- Collect cluster Events and convert to UIM alarms
- Tag clusters with name and node for each alarm
- Provide SyncIQ monitoring and DR status from Eyeglass DR dashboard and raise alarms on DR readiness changes

The Eyeglass Probe monitors the following components:

Component	**Version**
PowerScale clusters	7.2.0.x, 7.2.1.x, 8.0.0.x
SyncIQ replication	7.2.0.x, 7.2.1.x, 8.0.0.x

Overview of Architecture

Eyeglass Probe is solution that monitors PowerScale clusters and Vblocks . It collects the inventory and state information from these clusters and passes it to the Eyeglass PowerScale Dashboard. It also collects alarm information from these devices and passes it to Nimsoft UIM bus.

The data flow between Eyeglass Probe.

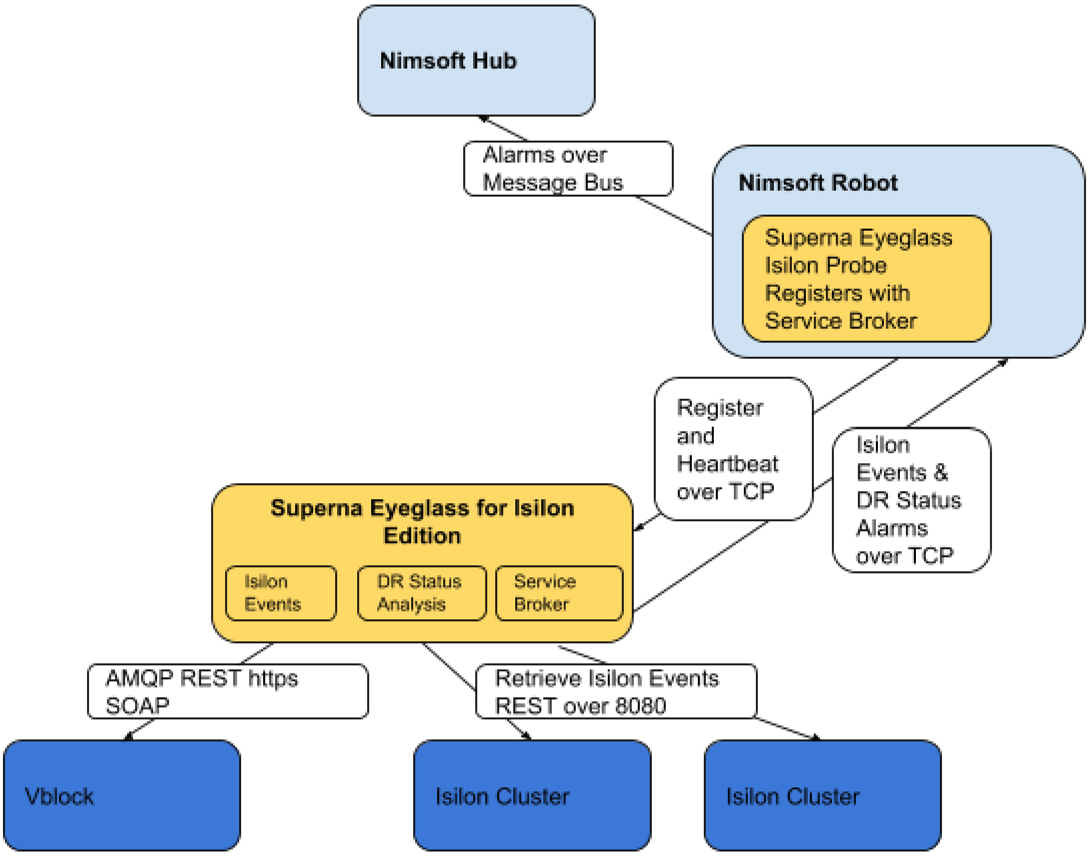
Nimsoft Hub
Alarms over Message Bus
Nimsoft Robot
Superna Eyeglass Isilon Probe Registers with Service Broker
Register and Heartbeat over TCP
Isilon Events & DR Status Alarms over TCP
Superna Eyeglass for Isilon Edition
Isilon Events
DR Status Analysis
Service Broker
AMQP REST https SOAP
Retrieve Isilon Events REST over 8080
Vblock
Isilon Cluster
Isilon Cluster

Description of Main Probe Tasks:

Task	Description
Inventory and Calculated (Calc) Status	• Protocol: REST • Scheduled retrieval of information every 5 minutes
Events	• Rest API
DR Status	Eyeglass REST API
Events Vblock	AMQP
Vblock Inventory	REST and SOAP

Configuration and Log Files

Eyeglass Probe

Configuration Items

Configuration files are located in the *C:\Program Files (x86)\Nimsoft\probes\marketplace\superna_eyeglass* directory. Following files have primary function in configuration of running probe:

File	Description
Superna_eyeglass.cfg .jar and .cfx	File containing parameters used by probe on initial startup.
superna_eyeglass.txt	Log file with errors on startup or exit

Log Files

Log files are located in th*e C:\Program Files (x86)\Nimsoft\probes\marketplace\superna_eyeglass*directory. Following log files contain the operational logging information for the probe:

File	Description
superna_eyeglass.txt	File containing logs related to: • Starting probe • Stopping probe • Connection to eyeglass appliance service broker process **Example log entry for Starting probe:** May 27 13:26:40:976 [main, superna_eyeglass] =========== START: Calling startupAndRestartInitialization(0) ======== May 27 13:26:40:976 [main, superna_eyeglass] =========== START: Reading configuration files and setting up monitors ======== May 27 13:26:40:976 [main, superna_eyeglass] START: GenResourceFactory.readAllQosDefs May 27 13:26:40:976 [main, superna_eyeglass] DONE: GenResourceFactory.readAllQosDefs entries: 0 May 27 13:26:40:976 [main, superna_eyeglass] =========== START: Scheduling DataCollectors for Resources ======== May 27 13:26:40:992 [main, superna_eyeglass] SCHEDULER: using

	scheduler_thread_pool_size = 3 May 27 13:26:40:992 [main, superna_eyeglass] =========== DONE: Scheduling DataCollectors for Resources ======== May 27 13:26:40:992 [main, superna_eyeglass] =========== DONE: Reading configuration files and setting monitors ======== May 27 13:26:40:992 [main, superna_eyeglass] =========== DONE: Finished startupAndRestartInitialization ======== May 27 13:26:44:918 [pool-1-thread-1, superna_eyeglass] Starting the eyeglass probe in 5 seconds... May 27 13:26:44:918 [pool-1-thread-1, superna_eyeglass] -------- Right before the registration call ----------- May 27 13:26:44:965 [pool-1-thread-1, superna_eyeglass] **Created service to listen on port 23457** **May 27 13:26:44:965 [pool-1-thread-1, superna_eyeglass] Service listening**

Debugging Facilities

None available.

Troubleshooting

Inventory Not Collected

Setup:

- Eyeglass Probe deployed and configured to Eyeglass appliance
- Eyeglass PowerScale licenses installed
- Eyeglass PowerScale connection to PowerScale Clusters

Problem:
In the Eyeglass PowerScale inventory tree   the cluster shows Adding…..

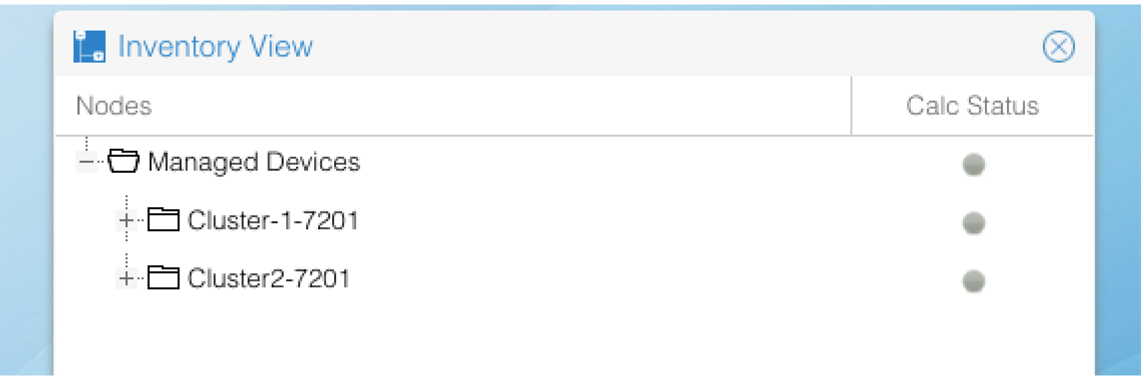

Troubleshooting:

1. Refresh the Eyeglass PowerScale Inventory tree.
2. Determine elapsed time since was added.   inventory may take 5 - 10 minutes depending on the state of the system when the was added:
    - Was inventory collection task in progress when was added? The in-progress inventory task has to complete before inventory for a new is added.
    - How many inventories are being completed simultaneously.
3. Look in main.log and determine whether there was an inventory failure. You will see this almost immediately after NDEUpdate task starts with an entry in main.log that looks something like this:

```
2013-11-04 15:39:14,898 NDEUpdateTask:collectData [211] Failed to get Inventory Data.
2013-11-04 15:39:14,898 RefreshNEData:doDataRefresh [183] Could not sync inventory for new nodes - Failed to get Shelf Data - Inventory failed
```

Troubleshooting Steps for connect to Eyeglass appliance:

1. Connect using https://x.x.x.x
2. Ensure ports 443, 2011 and 2012 are not blocked between browser and the appliance

Unable to Load Licenses Troubleshooting

Problem:
Try to upload the license zip file to the Eyeglass license manager .

Troubleshooting:

1. Browser
License upload only works with Chrome browser.

Alarms not appearing in UIM

Setup:

- Eyeglass Probe deployed and configured to Eyeglass appliance
- Eyeglass PowerScale licenses installed
- Eyeglass PowerScale connection to PowerScale Clusters

Problem:
Alarms do not appear to be processed in UIM from PowerScale clusters.

Troubleshooting

1. Login to eyeglass appliance and check status of the probe connection to the service broker process in eyeglass
2. On desktop open the Services
3. 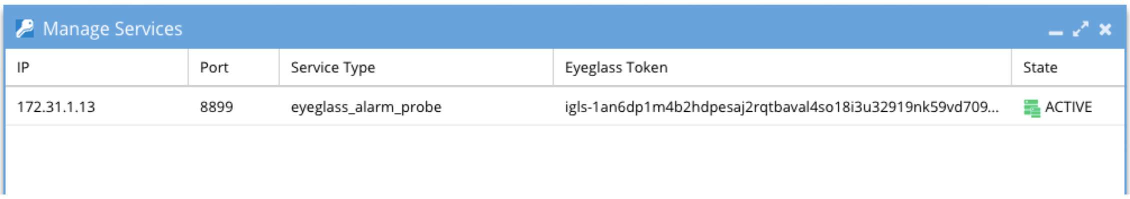

IP	Port	Service Type	Eyeglass Token	State
172.31.1.13	8899	eyeglass_alarm_probe	igls-1an6dp1m4b2hdpesaj2rqtbaval4so18i3u32919nk59vd709...	🟩 ACTIVE

4. If the State shows inactive it means the IP connection has failed between the probe machine and the eyeglass appliance
5. Check probe machine and test ping from this machine to the appliance
6. If this is successful then:
7. Login to eyeglass appliance via ssh using admin user account
8. Check tcp connections from the probe host machine using
9. Netstat -p | grep x.x.x.x (where x.x.x.x is ip address of the machine running the probe)
    a. If no tcp connections are seen from this source ip address it indicates connection issues on this port
10. Now try this command from the machine running the probe
    a. telnet y.y.y.y 37356 (y.y.y.y is the ip address of the eyeglass appliance and requires telnet on the probe machine)
    b. If connection is successful to the service broker port it indicates the connection is successful and telnet UI will appear to hang.
    c. If this fails this message will be seen
    d. Trying y.y.y.y...
    e. telnet: connect to address y.y.y.y: Connection refused
    f. telnet: Unable to connect to remote host
    g. If the command fails use another PC other than the probe host and try the same command. This will help identify if a firewall is blocking access to this port between the probe host and the eyeglass appliance

Can not Add PowerScale clusters to Eyeglass Appliance

Setup:

- Eyeglass Probe deployed and configured to Eyeglass appliance
- Eyeglass PowerScale licenses installed

Problem:
Can not add clusters to eyeglass for monitoring

Troubleshooting

1. Verify that ports https over 8080 and ssh 22 are open between the eyeglass appliance the PowerScale clusters and try again
2. Eyeglass will return password error if the password is incorrect. Failure to add will not be user or password related
3. Verify the required service broker ports are open between the Eyeglass appliance the robot running the probe. Verify the firewall on the eyeglass appliance has opened the ports for the Probe default port connect and register.
4. Verify ip reachability with ping and no firewall blocks ports
5. Open a case with support if known of the above resolve the adding… cluster.

11. Eyeglass Installation Procedure - Proxmox kvm qcow

Home Top

- Support Statement
- Superna Eyeglass OVx Conversion Procedure
- ECA OVx Conversion Procedure

Support Statement

This guide is as is and not supported by the support contract. This is provided as a guide example only to be used as a reference for the Proxmox hypervisor that uses qcow and KVM. These steps only apply to this hypervisor but is an example of how a OVF/OVA can be converted.

Superna Eyeglass OVx Conversion Procedure

Create a new VM in Proxmox, taking note of VM ID, you can give it a name for your ease.

General:

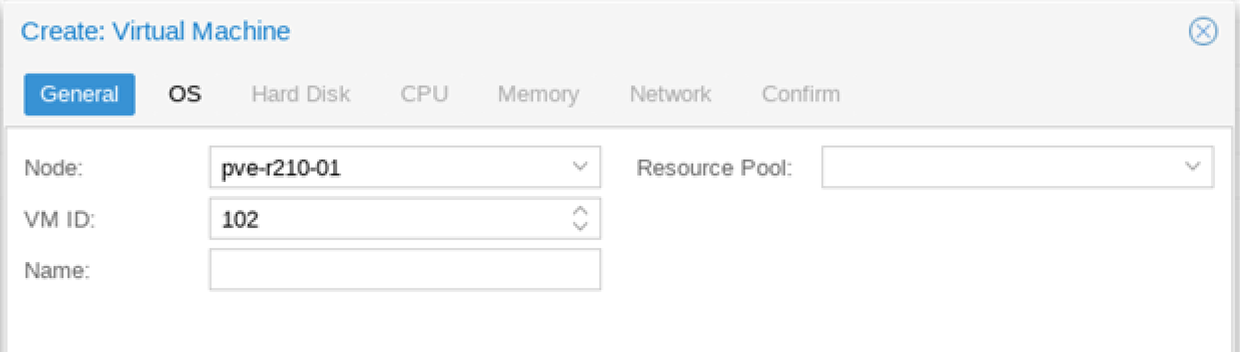

Node: pve-r210-01

VM ID: 102

Name:

Resource Pool:

OS: Don’t Use any any media, OS : Linux 4.x/3.x Kernel

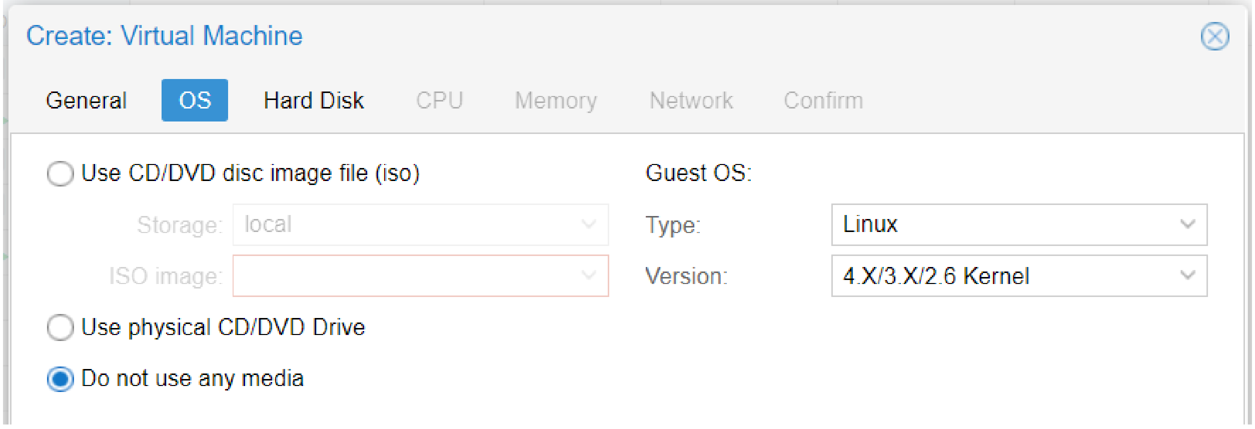

○ Use CD/DVD disc image file (iso)

Storage: local

ISO image:

○ Use physical CD/DVD Drive

● Do not use any media

Guest OS:

Type: Linux

Version: 4.X/3.X/2.6 Kernel

Hard Disk: Change ‘Bus/Device’ to SATA, Everything else can be left as defaults, this HD will be removed in a later step.

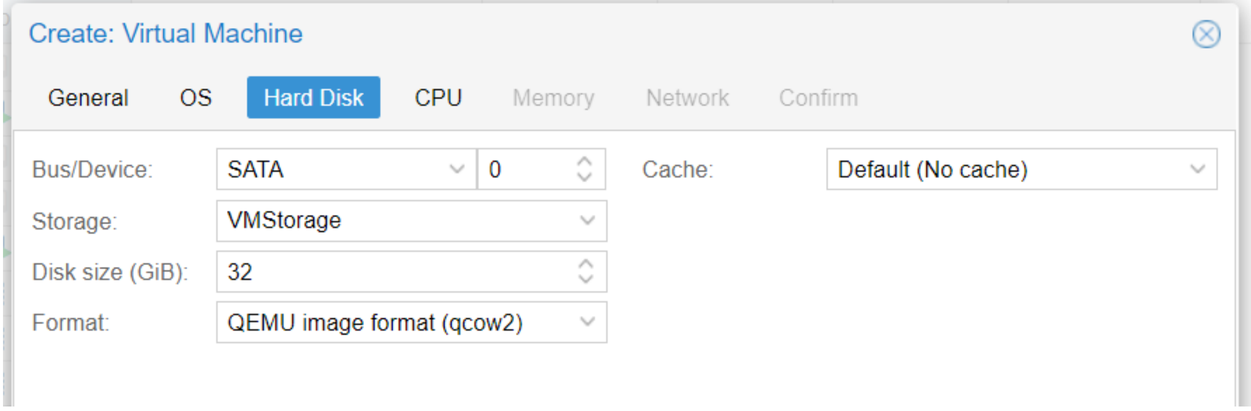

Bus/Device: SATA 0

Storage: VMStorage

Disk size (GiB): 32

Format: QEMU image format (qcow2)

Cache: Default (No cache)

CPU: 2 Sockets, 2 Cores. Default (kvm64) type was tested

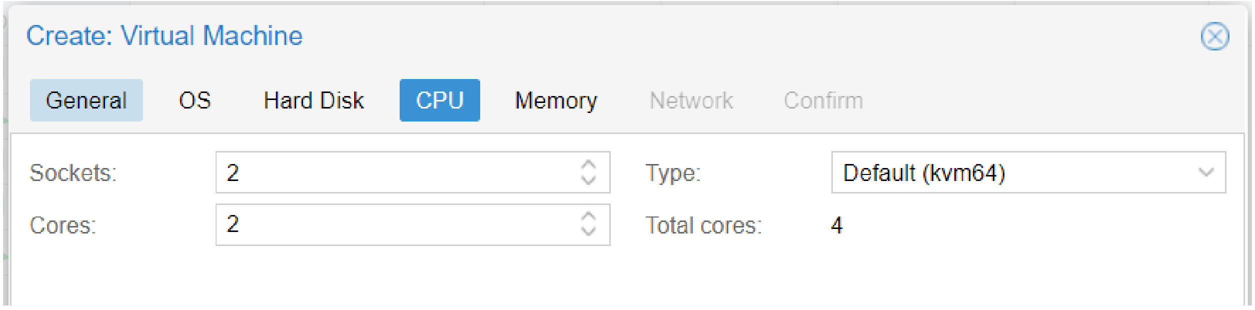

Sockets: 2

Cores: 2

Type: Default (kvm64)

Total cores: 4

Memory: 16384MB Ram (16G)

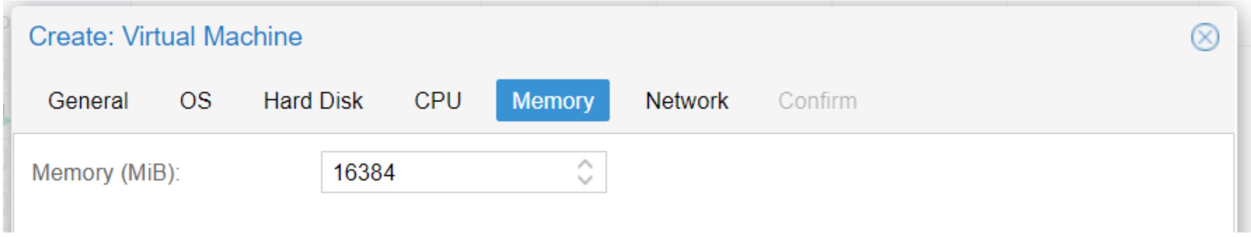

Memory (MiB): 16384

Network: Add an Ethernet adapter (VirtIO Paravirtualized was tested), and a **VLAN tag** if required.

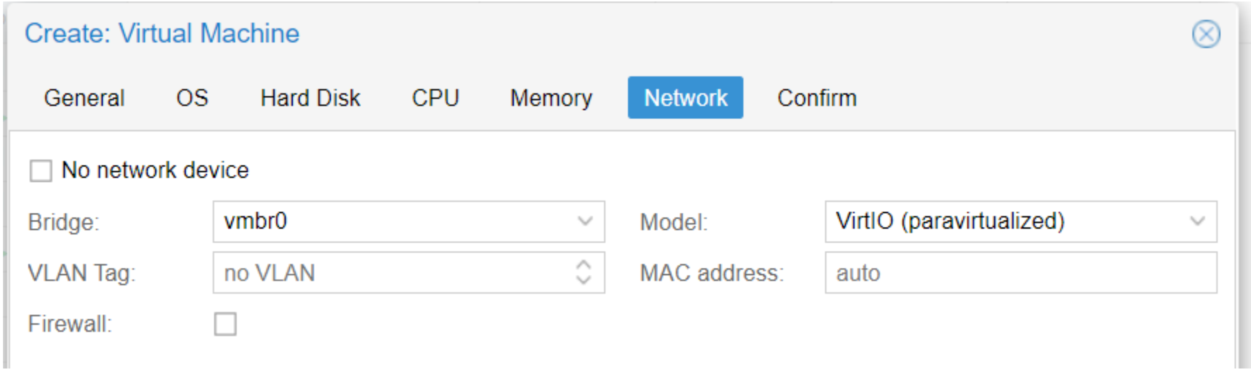

□ No network device

Bridge: vmbr0

Model: VirtIO (paravirtualized)

VLAN Tag: no VLAN

MAC address: auto

Firewall: □

Click on the newly created VM, click Hardware, and then select Hard Disk (scsi0), then click **Detach** and Yes when asked to confirm.

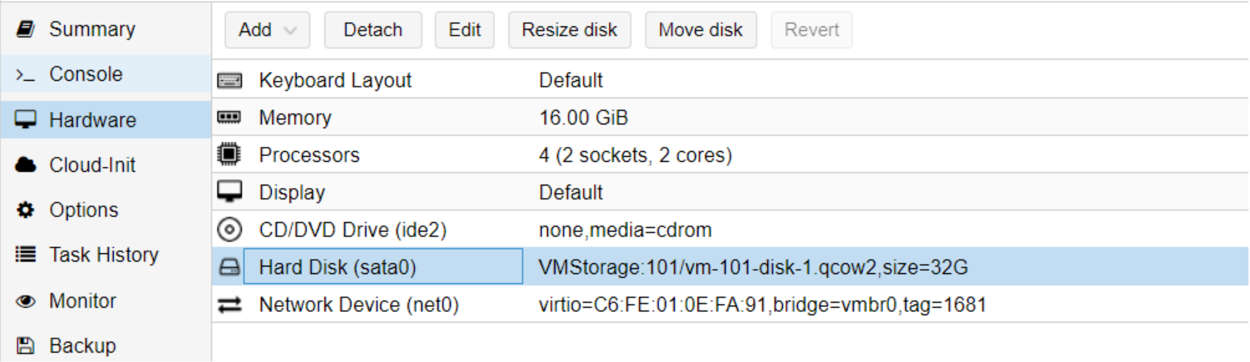

Keyboard Layout	Default
Memory	16.00 GiB
Processors	4 (2 sockets, 2 cores)
Display	Default
CD/DVD Drive (ide2)	none,media=cdrom
Hard Disk (sata0)	VMStorage:101/vm-101-disk-1.qcow2,size=32G
Network Device (net0)	virtio=C6:FE:01:0E:FA:91,bridge=vmbr0,tag=1681

The Hard Disk will drop into an Unused Disk state. Select this Unused Disk and select **Remove**

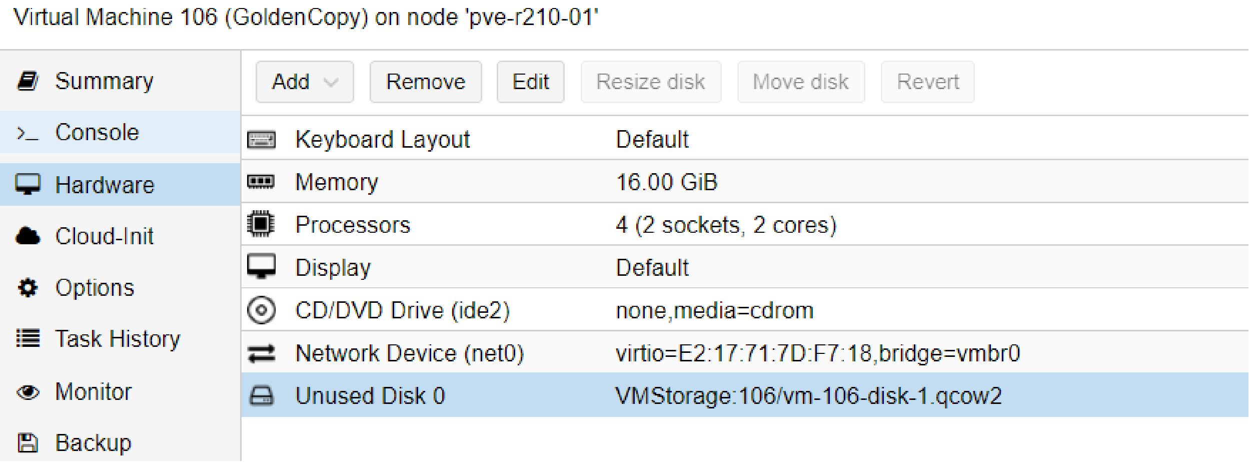
Virtual Machine 106 (GoldenCopy) on node 'pve-r210-01'

Keyboard Layout	Default
Memory	16.00 GiB
Processors	4 (2 sockets, 2 cores)
Display	Default
CD/DVD Drive (ide2)	none,media=cdrom
Network Device (net0)	virtio=E2:17:71:7D:F7:18,bridge=vmbr0
Unused Disk 0	VMStorage:106/vm-106-disk-1.qcow2

Option 1: OVF/VMDK only

Download the VMDK from repo. IF OVA, extract the vmdk from the OVA using an archive program. (7zip, winzip winrar all work).

Transfer the VMDK to Proxmox VE images folder (*/var/lib/vz/images*)

E.g.

```
scp Superna_Eyeglass.2.5.7-20285-lp15.1.x86_64-disk1-pve.vmdk root@proxmoxip:/var/lib/vz/images
```

SSH over to the Proxmox HV to perform the following steps.

Convert the image to qcow2:

```
qemu-img convert -f vmdk Superna_Eyeglass.2.5.7-20285-lp15.1.x86_64-disk1.vmdk
 -O qcow2 Superna_Eyeglass.2.5.7-20285-lp15.1.x86_64.qcow2
```

You can now **copy (or move)** this qcow to the VMID subfolder from above.

At this point, Proxmox should populate the WebUI with the new drive so it can be simply attached. If it doesn't, manually add the drive by editing the VM config file, stored by default at */etc/pve/qemu-server* and add the following line to the end of the config file:

```
unused0: local:<id>/<filename>
```

Where **<ID>** is the VMID noted above and **<filename>** is the name of qcow2 disk to be attached to the VM. E.g.

```
unused0: local:102/Superna_Eyeglass.2.5.3-18251.x86_64-disk1-pve.qcow2
```

Return to the Proxmox WebUI where you will find the eyeglass boot disk, designated as ‘unused’.

Double click on the unused disk to Attach it to the VM:

Bus/Device: SATA 0

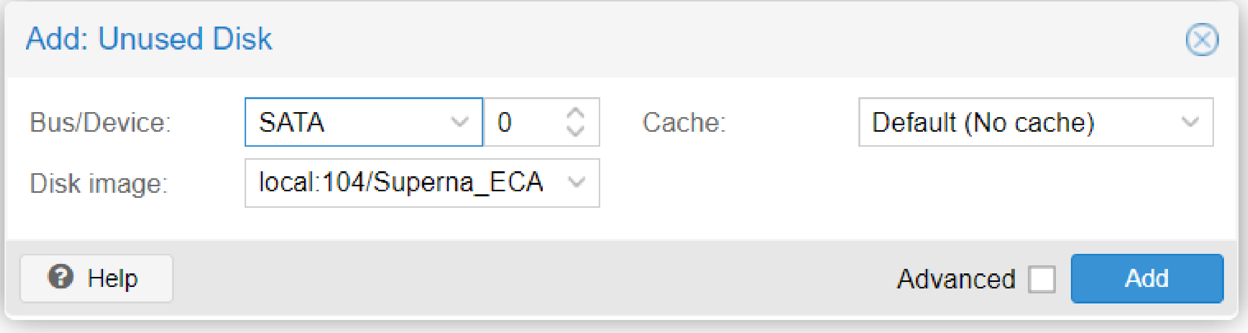

Add: Unused Disk

Bus/Device:	SATA 0	Cache:	Default (No cache)
Disk image:	local:104/Superna_ECA		

Click Add.

The configuration should look similar to this:

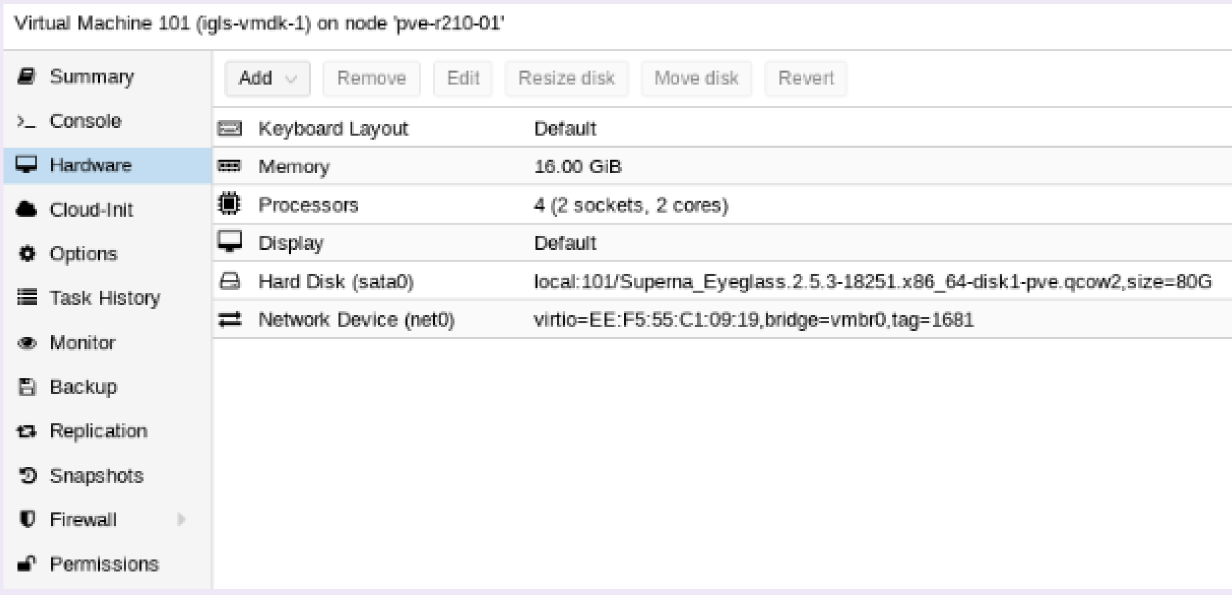

Virtual Machine 101 (igls-vmdk-1) on node 'pve-r210-01'

Keyboard Layout	Default
Memory	16.00 GiB
Processors	4 (2 sockets, 2 cores)
Display	Default
Hard Disk (sata0)	local:101/Superna_Eyeglass.2.5.3-18251.x86_64-disk1-pve.qcow2,size=80G
Network Device (net0)	virtio=EE:F5:55:C1:09:19,bridge=vmbr0,tag=1681

Start the VM, Connect to Console; you can login with user **ecaadmin** and default password **3y3gl4ss** and proceed with normal post install configuration.

ECA OVx Conversion Procedure

Create a new VM in Proxmox, taking note of VM ID, you can give it a name for your ease.

General:

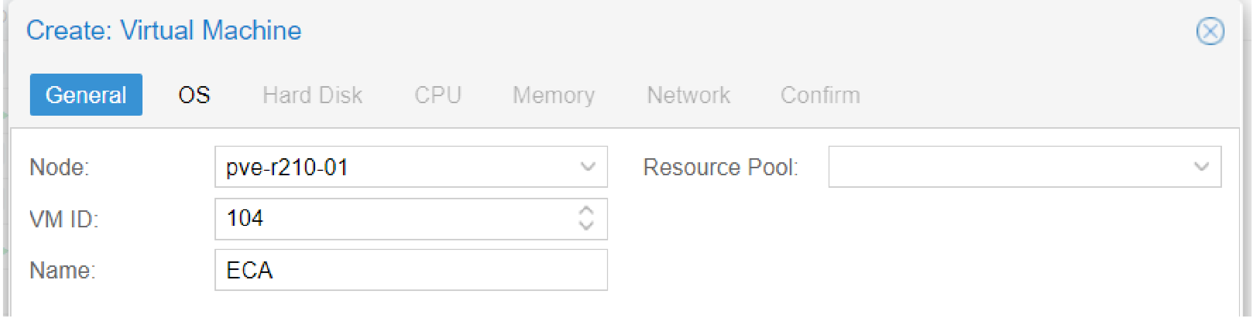

Create: Virtual Machine

Node: pve-r210-01

Resource Pool:

VM ID: 104

Name: ECA

OS: Don’t Use any any media, OS : Linux 4.x/3.x Kernel

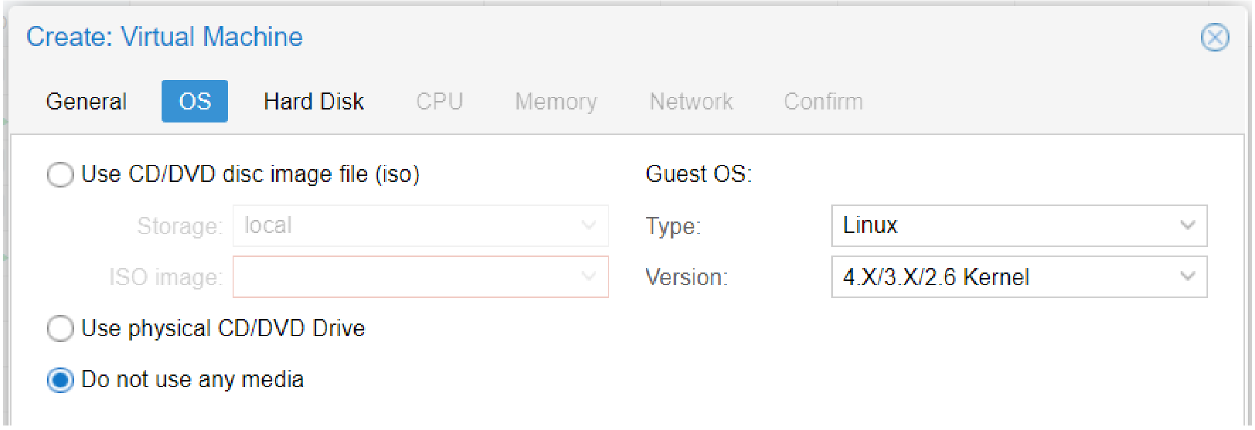

Create: Virtual Machine

○ Use CD/DVD disc image file (iso)

Storage: local

ISO image:

○ Use physical CD/DVD Drive

● Do not use any media

Guest OS:

Type: Linux

Version: 4.X/3.X/2.6 Kernel

Hard Disk: Change 'Bus/Device' to SATA, Everything else can be left as defaults, this HD will be removed in a later step.

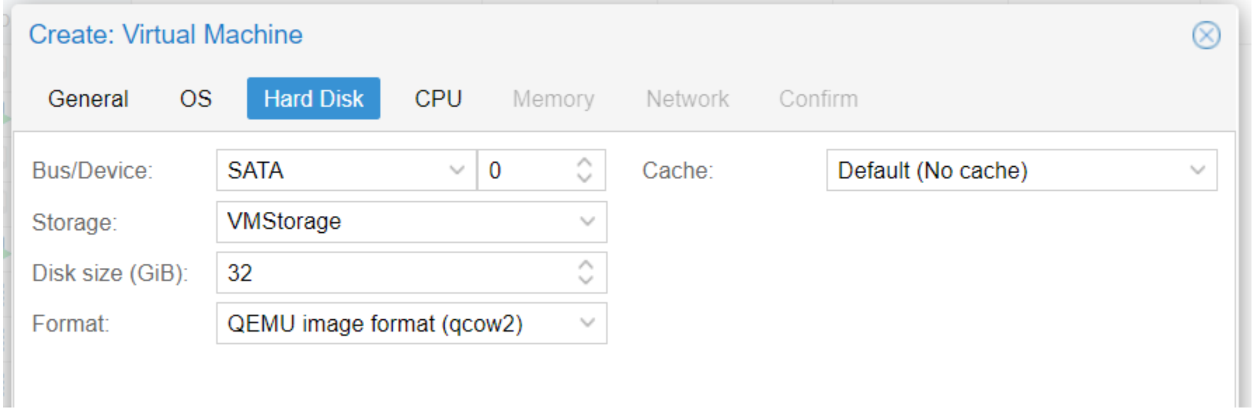

Create: Virtual Machine

Bus/Device: SATA 0

Storage: VMStorage

Disk size (GiB): 32

Format: QEMU image format (qcow2)

Cache: Default (No cache)

CPU: 2 Sockets, 2 Cores. Default (kvm64) type was tested

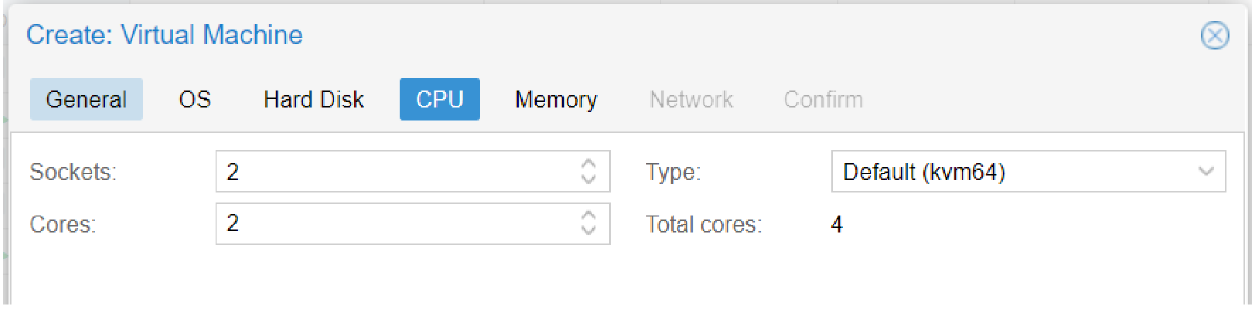

Create: Virtual Machine

Sockets: 2

Cores: 2

Type: Default (kvm64)

Total cores: 4

Memory: 16384MB Ram (16G)

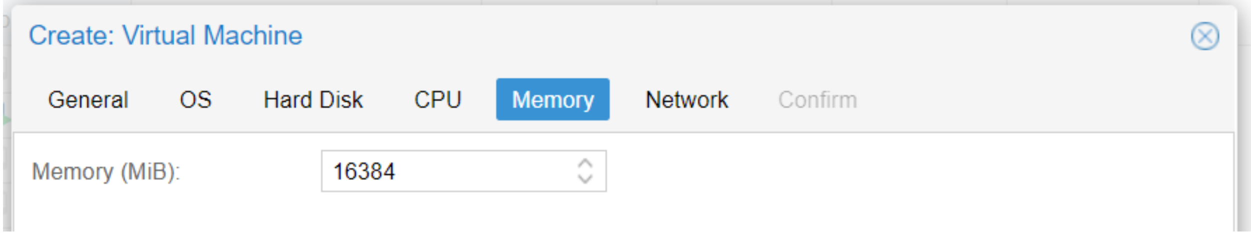

Memory (MiB): 16384

Network: Add an Ethernet adapter (VirtIO Paravirtualized was tested), and a **VLAN tag** if required.

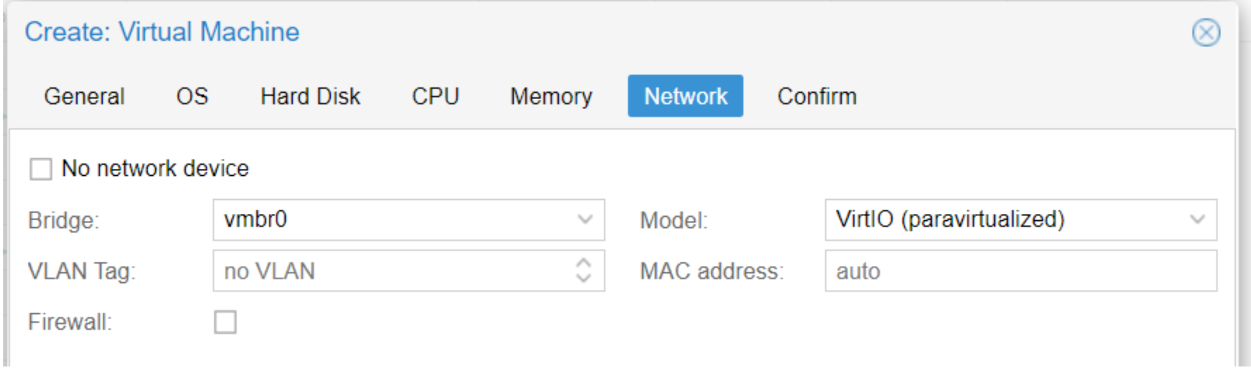

□ No network device

Bridge:	vmbr0	Model:	VirtIO (paravirtualized)
VLAN Tag:	no VLAN	MAC address:	auto
Firewall:	□		

Click on the newly created VM, click Hardware, and then select Hard Disk (scsi0), then click **Detach** and Yes when asked to confirm.

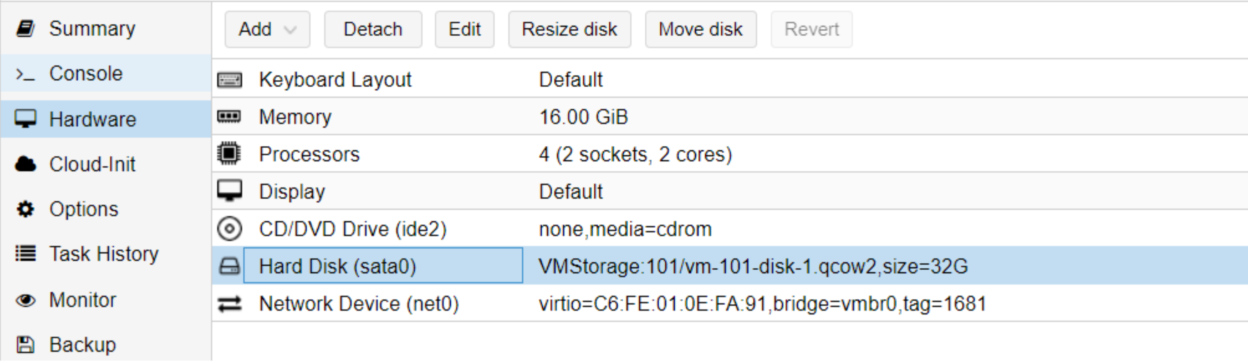

Keyboard Layout	Default
Memory	16.00 GiB
Processors	4 (2 sockets, 2 cores)
Display	Default
CD/DVD Drive (ide2)	none,media=cdrom
Hard Disk (sata0)	VMStorage:101/vm-101-disk-1.qcow2,size=32G
Network Device (net0)	virtio=C6:FE:01:0E:FA:91,bridge=vmbr0,tag=1681

The Hard Disk will drop into an Unused Disk state. Select this Unused Disk and select **Remove**

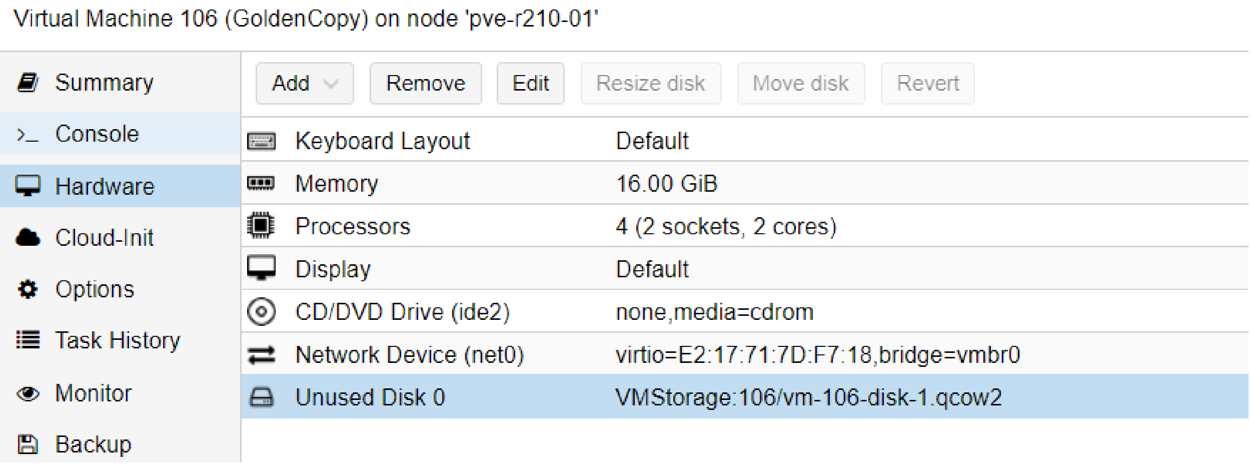

Keyboard Layout	Default
Memory	16.00 GiB
Processors	4 (2 sockets, 2 cores)
Display	Default
CD/DVD Drive (ide2)	none,media=cdrom
Network Device (net0)	virtio=E2:17:71:7D:F7:18,bridge=vmbr0
Unused Disk 0	VMStorage:106/vm-106-disk-1.qcow2

Option 1: OVF/VMDK only

Download the VMDK from repo. IF OVA, extract the vmdk from the OVA using an archive program. (7zip, winzip winrar all work).

Transfer the VMDK to Proxmox VE images folder (*/var/lib/vz/images*)

E.g.

```
scp Superna_ECA.2.5.7-20285-lp15.1.x86_64-disk1-pve.vmdk root@proxmoxip:/var/lib/vz/images
```

SSH over to the Proxmox HV to perform the following steps.

Convert the image to qcow2:

```
qemu-img convert -f vmdk Superna_ECA.2.5.7-20285-lp15.1.x86_64-disk1.vmdk
 -O qcow2 Superna_ECA.2.5.7-20285-lp15.1.x86_64.qcow2
```

You can now **copy (or move)** this qcow to the VMID subfolder from above.

At this point, Proxmox should populate the WebUI with the new drive so it can be simply attached. If it doesn't, manually add the drive by editing the VM config file, stored by default at */etc/pve/qemu-server* and add the following line to the end of the config file:

```
unused0: local:<id>/<filename>
```

Where **<ID>** is the VMID noted above and **<filename>** is the name of qcow2 disk to be attached to the VM. E.g.

```
unused0: local:102/Superna_Eyeglass.2.5.3-18251.x86_64-disk1-pve.qcow2
```

Return to the Proxmox WebUI where you will find the eyeglass boot disk, designated as 'unused'.

Double click on the unused disk to Attach it to the VM:

Bus/Device: SATA 0

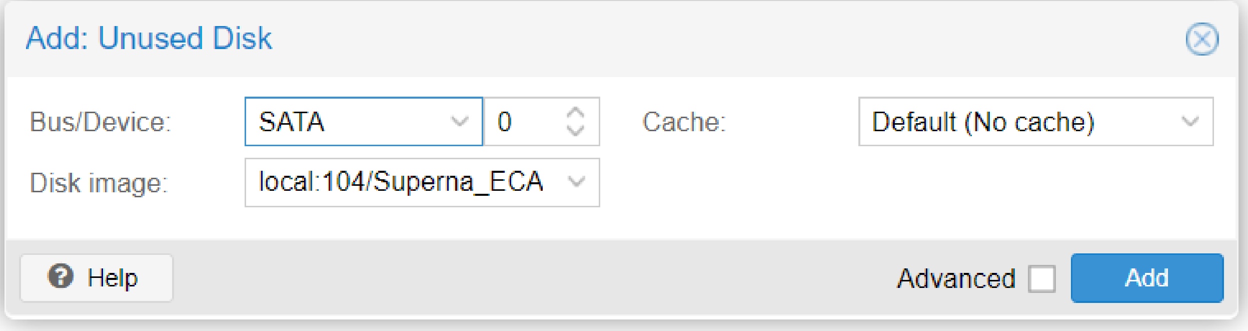

Add: Unused Disk

Bus/Device: SATA 0

Cache: Default (No cache)

Disk image: local:104/Superna_ECA

Click Add.

Select the boot order for adjusted device 1 to sata0. In addition to this we require a completely empty second disk. From UI Add > Hard Disk; make sure to choose SATA 1 and disk size minimum 80 GB.

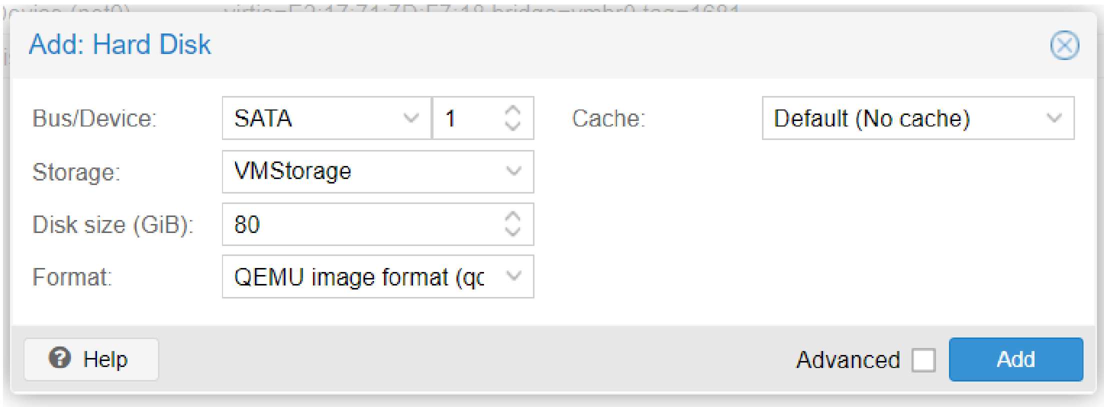

Add: Hard Disk

Bus/Device: SATA 1

Cache: Default (No cache)

Storage: VMStorage

Disk size (GiB): 80

Format: QEMU image format (qc

The configuration should look similar to this:

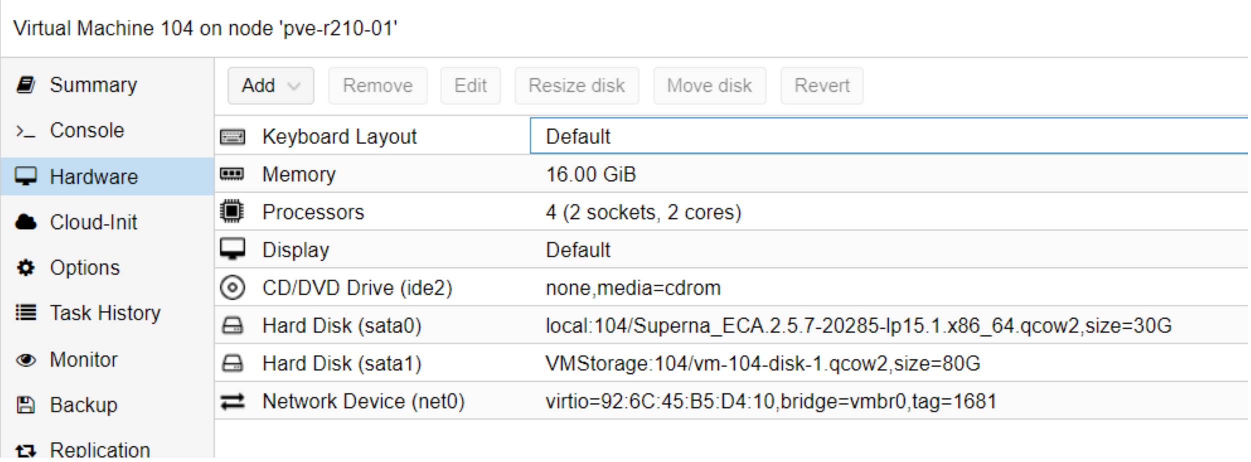

Virtual Machine 104 on node 'pve-r210-01'

Keyboard Layout	Default
Memory	16.00 GiB
Processors	4 (2 sockets, 2 cores)
Display	Default
CD/DVD Drive (ide2)	none,media=cdrom
Hard Disk (sata0)	local:104/Superna_ECA.2.5.7-20285-lp15.1.x86_64.qcow2,size=30G
Hard Disk (sata1)	VMStorage:104/vm-104-disk-1.qcow2,size=80G
Network Device (net0)	virtio=92:6C:45:B5:D4:10,bridge=vmbr0,tag=1681

Start the VM, Connect to Console; you can login with user **ecaadmin** and default password **3y3gl4ss** and proceed with post install configuration as follows.

Now, configure the networking setting by specifying an IP Address, Subnet Mask, Default Gateway, Cluster Name and comma-separated DNS Nameservers manually with

```
/opt/superna/bin/ovf set-value [-h] KEY=VALUE [KEY=VALUE ...]
```

For e.g.

```
/opt.superna.bin/ovf set-value net.eth0.ipv4.ip=192.168.2.227 net.eth0.ipv4.gateway=192.168.1.1
net.eth0.ipv4.netmask=255.255.255.0 vm.clustername=ecatest net.nameservers=172.16.80.7
```

You can run `/opt/superna/bin/ovf definitions` to show what properties are supported by this deployment. However you need a separate call for **mode** and **hostname** (which should be same name as vm.clustername)

```
/opt/superna/bin/ovf set-value --force net.hostname=’ecatest' mode='eca'
```

Once the OVF values are set, check the on boot log on

```
/var/log/superna_on_boot.log
```

If the tail message appearing like “Waiting for node communication keys from master node …”, you have to manually create */tmp/install-done*

Then run the following commands

```
touch /opt/superna/.firstboot
echo '3y3gl4ss' > /tmp/ecaadmin_pw
sudo /opt/superna/bin/on_boot.sh
```

Let’s wait for the firstboot script to complete. You can check the current progress on the screen or on

```
/var/log/superna_on_boot.log
```

During this **on-boot** process, the script will create a data filesystem on the second disk. Also, sets the networking properties and distributes the configurations and SSH/SSL keys.

Note: If the tail message appearing like “Waiting for node communication keys from master node …”, you have to manually create */tmp/install-done*

Now, you can verify the installed version with

```
ecactl version
```

You’ll see something like this:

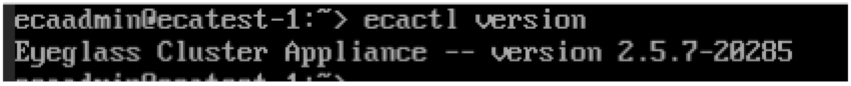

```
ecaadmin@ecatest-1:~> ecactl version
Eyeglass Cluster Appliance -- version 2.5.7-20285
```

You might need to configure /opt/superna/eca/eca-env-common.conf as determined by the environment like the Isilon IP address, etc.

12. Ransomware Defender Enterprise Airgap Agent VM Install Guide

Introduction to this Guide

Use this document to get your new Ransomware Defender single VM into a secure AirGap vault ESX host.

System Requirements

1. vSphere 6.0 or higher appliance appliance requires
2. 4 vCPU

3. 16 GB RAM
4. 30G OS partition plus 80 GB disk Total disk size in VMware 110G

IMPORTANT INFORMATION REGARDING ADDING CLUSTERS TO AirGap Agent VM READ-ME FIRST

-

  Supported OneFS releases

1. Please refer to the Release Notes for the Ransomware Defender AirGap Vault Agent.

Prerequisites

1. ESX host installed
2. ensure esx host time is set correctly and time zone is correct
3. Ransomware Defender AirGap Vault Agent VM should have time sync to VM checked

Ransomware Defender AirGap Vault Agent Firewall Port Requirements

1. The vault agent requires
   a. port 8080 https from vault agent --> to the vault cluster management pool
   b. port 23 ssh from vault agent --> to the vault cluster management pool

Download Eyeglass (Mandatory)

1. Download ECA zip from Superna web site following instructions here Latest ECA Download

2.

Deploy the Ransomware Defender AirGap Vault Agent Appliance (Mandatory)

Ransomware Defender AirGap Vault Agent is delivered in an OVF format for easy deployment in your esx environment. Deploy the OVF and then follow the wizard to setup networking for this Linux appliance. You will need to know:

1. subnet and network required so that appliance will have IP connectivity to the PowerScale clusters that it's managing, and the users that are using it
2. IP address for the appliance
3. Gateway

Steps to Deploy the OVF with vSphere Client (Mandatory)

OVF Deployment steps :
Step 1 : Download an ECA OVF zip file from Latest Appliance Download.
Step 2 : Unzip the contents of the zip file from Step 1 onto a computer with vSphere web url.

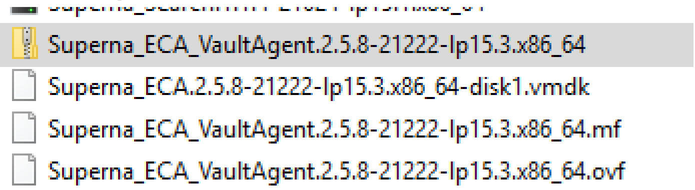

Superna_ECA_VaultAgent.2.5.8-21222-lp15.3.x86_64
Superna_ECA.2.5.8-21222-lp15.3.x86_64-disk1.vmdk
Superna_ECA_VaultAgent.2.5.8-21222-lp15.3.x86_64.mf
Superna_ECA_VaultAgent.2.5.8-21222-lp15.3.x86_64.ovf

Step 3 : Login to the esx host with appropriate login credentials.
Step 4 : Single click on VMware vSphere client on the Desktop. Login with appropriate login credentials.

Step 5 : Once logged in to VMware client, you can see different Menus on the top left of the application. Next, go to the File menu and select Deploy OVF Template.

Step 6 : Browse to the location of OVF files you've downloaded and unzipped in step 1 and 2. Select OK and then Next.

Next, You will see the OVF template details. Verify the details and proceed by selecting Next. Notice download size to be under allocated disk size limit.

Step 7 : Choose a unique name for the virtual machine and select Inventory location for the deployed template. Once done, select Next.

Step 8 : Select the host/cluster where you want to run the deployed template and then Next.

Step 9 : Select the Resource pool within which you wish to deploy the template.

Step 10 : Select a destination storage for virtual machine files, select Next

Step 11 : Select Disk Format for the datastore you selected in previous step.

Step 12 : Enter the networking properties for the Eyeglass appliance VM in the OVF properties display. Replace with correct settings for your environment.

Step 13 : When done, verify your settings and deploy the OVF

After deployment:

Step 1 : Power On the virtual machine.

1. The Ransomware Defender AirGap Vault Agent appliance is deployed with following default admin user password:
2. ssh to eyeglass vm as ecaadmin
    a. **sudo systemctl status superna-on-boot** (enter admin password and verify the first boot process completes)
    b. default login and password: **ecaadmin/3y3gl4ss**
    c. set the OVF mode
        i. ovf set-value -f mode=vault-agent
    d. **sudo -s**
    e. umount /opt/superna/mnt/zk-ramdisk

   f. remove the tmpfs /opt/superna/mnt/zk-ramdisk tmpfs nodev,nosuid,noexec,nodiratime,size=512M 0 0 line from /etc/fstab
      i. nano /etc/fstab
      ii. remove the line above
      iii. with ctrl key +k on the line
      iv. ctrl key + x to save and exit
      v. done
3. Can also be used to login to the SSH session
4. **Mandatory: It is highly recommended to reset the default password after the appliance is deployed.**
   a. **Type passwd**
   b. **enter new password and confirm password**
5. Setup Time zone (Mandatory)
   a. Follow Animated GIF below to set using YAST
   b. ssh as admin user,
   c. sudo -s
   d. Enter admin password
   e. type yast
   f. select menu system --> date and time
   g. set the time zone
   h. Done
6. Cluster startup
   a. exit if you are still root user
   b. whoami (make sure you are ecaadmin)
   c. ecactl cluster up
   d. done

Upgrade Vault Agent Procedures

This section covers how to upgrade the vault agent software. Two methods are available depending on vault operation mode. Method #1 in-band maintenance window option requires the vault agent to have in-band maintenance mode enabled, Method #2 requires console access or mouse keyboard access to esx host inside the vault, this requires physical access to the vault.

Method #1 - in-band maintenance mode upgrade procedure

1. Requirements:
    a. Enable Vault agent in band maintenance mode see guide.
    b. 2 Factor authentication configured on the eyeglass VM and the vault agent vm. See guide here.
2. Upgrade Procedures
    a. Download the vault agent upgrade file from support https://support.superna.net. Recored the md5 checksum from the download menu.
    b. Compute the md5 checksum after download and compare to the md5 checksum posted on the download site.
    c. Using winscp copy the run file to the production Powerscale connected to the vault network
    d. Request the vault maintenance mode window with this command for a 45 minute maintenance window
        i. igls airgap vaultaccessrequest --interval=45
        ii. The vault agent checks in every 2 hour for maintenance requests on the hour example 8 am, 10 am etc..
        iii. At the next time interval check prepare for the upgrade following steps below.

e. ssh to the production cluster with the user that copied the upgrade file to the cluster. This example assumes the admin cluster user with a home directory of /ifs/home/admin.

    i. 1 minute after the hour, test access to the with scp /ifs/home/admin/<upgrade file name> eyeglass@x.x.x.x:<upgrade file name> (x.x.x.x is IP of the vault cluster interface). If the vault door is open a password prompt will be presented, if a timeout message appears the vault has not opened for maintenance yet, retry the scp copy.

    ii. Once the file is copied to the vault login to the vault over ssh

        1. ssh eyeglass@x.x.x.x

        2. scp /ifs/home/eyeglass/<upgrade file name> ecaadmin@y.y.y.y:<upgrade file name> (y.y.y.y is the vault agent VM) enter ecaadmin password to complete the copy.

    iii. SSH to the vault agent VM

        1. ssh ecaadmin@y.y.y.y

    iv. Shutdown vault agent

        1. ecactl cluster down

    v. upgrade vault agent

        1. chmod 777 /home/ecaadmin/<upgrade file name>

        2. ./home/ecaadmin/<upgrade file name>

3. wait for upgrade to complete

vi. start vault cluster software

1. ecactl cluster up

vii. Check remaining time for the timed maintenance before the vault network is auto closed

1. ecactl airgap checkopen

viii. Verify upgrade

1. type docker ps (verify all containers are running and none are restarting)

ix. Check configuration

1. ecactl isilons list
2. ecactl airgap list

x. Close the vault to exit the ssh session

1. ecactl vault close
2. The ssh session will timeout once the vault closes

xi. Upgrade Complete

Method #2 - Physical Vault VM Access

1. **Requirements**
   - a. Bastion host with access to vault management switch inside the vault
   - b. OR physical access with laptop to management switch inside the vault

c. Secured vault laptop dedicated for vault access. This laptop should be physically secured when not in use with limited personnel access to the laptop. Change control to gain access to the secured laptop.

d. Dedicated USB stick for up

2. Upgrade Procedures

    a. Download the vault agent upgrade file from support https://support.superna.net. Recored the md5 checksum from the download menu.

    b. Compute the md5 checksum after download and compare to the md5 checksum posted on the download site.

    c. Copy the upgrade file to a USB stick.

    d. Copy the upgrade file from the USB stick to the to the management laptop.

    e. Winscp the upgrade file to the vault agent using the ecaadmin user

    f. SSH to the vault agent VM

        i. ssh ecaadmin@y.y.y.y (or use ssh utility)

    g. Shutdown vault agent

        i. ecactl cluster down

    h. upgrade vault agent

        i. chmod 777 /home/ecaadmin/<upgrade file name>

        ii. ./home/ecaadmin/<upgrade file name>

        iii. wait for upgrade to complete

    i. Start vault cluster software

    i. ecactl cluster up

j. Verify upgrade

    i. type docker ps (verify all containers are running and none are restarting)

k. Check configuration

    i. ecactl isilons list

    ii. ecactl airgap list

l. Test Vault Connectivity to production cluster

    i. ecactl airgap check --prod <protected cluster name>

    ii. Verify successful communications

m. Upgrade Complete