

Table of Contents

1. Failover Recovery Procedures.....	2
1.1. Overview.....	3
1.2. Failover Log Analysis.....	4
1.3. Replication Policy Failover Preparation (Data Sync and Config Sync).....	6
1.4. DFS Mode.....	9
1.5. Networking Operations.....	10
1.6. Replication Policy Failover - All policies.....	12
1.7. Replication Policy Failover - Recovery.....	17
1.8. Run Quota Jobs Now Failover Recovery Steps.....	21
1.9. Replication Policy Failover Finalize.....	23
1.10. Post Failover Script Execution.....	25
1.11. Post Failover.....	26
1.12. Check Client Access (Manual Step).....	28
2. Cluster Move Procedure to New Data Center.....	30
2.1. Cluster Move Procedure to New Site - Change Cluster IP Address.....	31
2.2. Cluster Move Procedure to New Site - Cluster IP Unchanged.....	34
3. PowerScale Upgrade Procedure with Eyeglass.....	37
3.1. ReadMe First - Prior to PowerScale Upgrade.....	38
3.2. Cluster OneFS Upgrade is between Major Release Version Eyeglass Procedures.....	39
3.3. Cluster OneFS Upgrade is between Minor Release Version Eyeglass Procedures.....	45
3.4. Functions Impacted DURING a Cluster OneFS Upgrade and Eyeglass, Ransomware Defender, Easy Auditor, Performance Auditor Procedures.....	47
4. Eyeglass Simulated Disaster Event Test Procedure.....	49
4.1. Introduction.....	50
4.2. Initial Environment Setup.....	52
4.3. Verify Environment Setup.....	55
4.4. Simulated Disaster Scenario - DFS Test SyncIQ Policy Failover.....	56
4.5. Simulated Disaster Scenario - EyeglassRunbookRobot Access Zone Failover.....	65
5. Appliance Operational Procedures.....	74
6. All Products Hardening Guide.....	80
7. InPlace Appliance Open Suse OS Upgrade.....	109

1. Failover Recovery Procedures

Home [Top](#)

- [Overview](#)
- [Failover Log Analysis](#)
- [Replication Policy Failover Preparation \(Data Sync and Config Sync\)](#)
- [DFS Mode](#)
- [Networking Operations](#)
- [Replication Policy Failover - All policies](#)
- [Replication Policy Failover - Recovery](#)
- [Run Quota Jobs Now Failover Recovery Steps](#)
- [Replication Policy Failover Finalize](#)
- [Post Failover Script Execution](#)
- [Post Failover](#)
- [Check Client Access \(Manual Step\)](#)

© Superna Inc

1.1. Overview

[Home](#) [Top](#)

Overview

This document covers failover recovery procedures, the steps to review failover logs and determine which step failed, along with the steps to recover. The tables in this document are labeled to identify the various steps in failover where an error would occur, along with an identification of the error, description of impact on failover, recovery steps and any special instructions or warnings. Failover Preparation and all the Policies, etc.

Each failover log downloaded from DR Assistant contains ALL sections in this document. For any step indicated as failed in the failover log, consult the section in this recovery guide for steps to remediate or recovery from a failed step.

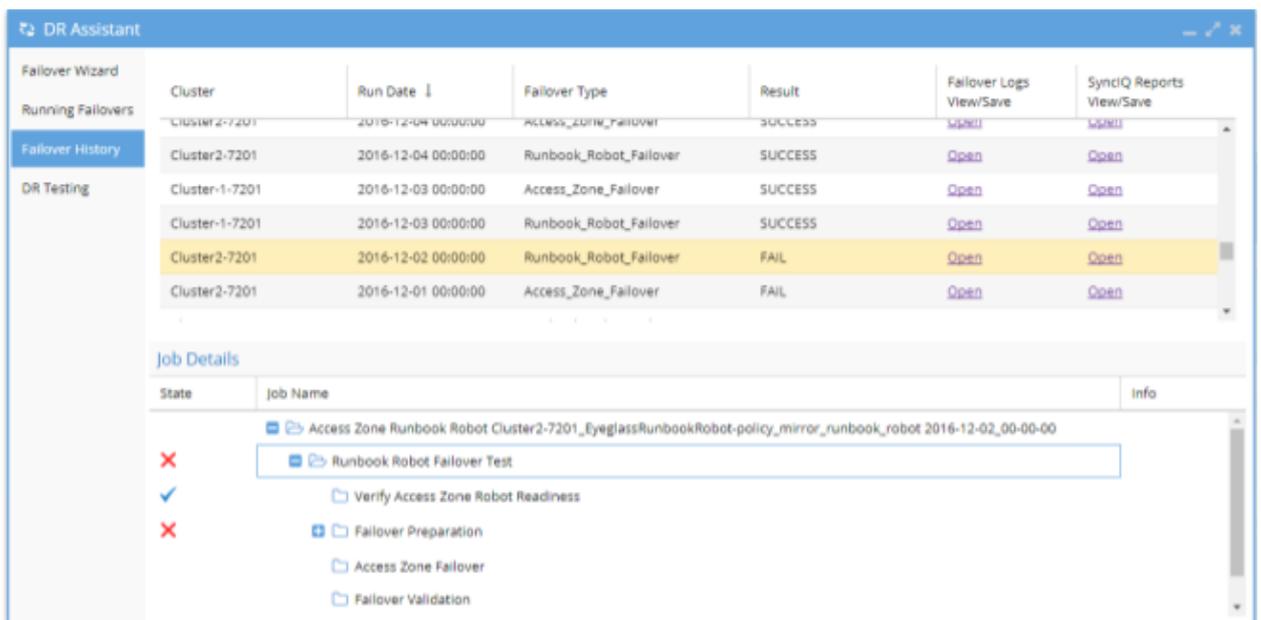
© Superna Inc

1.2. Failover Log Analysis

Home Top

Failover Log Analysis

The Failover Logs can be found by clicking on the DR Assistant on the Eyeglass desktop. Click on the Failover History tab, you'll see the various jobs that have been run, run date and results (Pass or Fail). By clicking on an individual job, the details of that job will appear in the window below (see screenshot below).



Steps to reading a failover log.

1. Identify the section in the log with an error message expand folders to find red X
2. Determine which step failed and which table applies for next step
3. See which scenario in each table applies to the scenario in the failover log

NOTE: Frequently SyncIQ policy issues are source of errors. For steps run policy, make writeable, resync prep (4 steps) and run mirror policy the related SyncIQ job report has details on the error condition.

4. As of 1.8 or later releases: The following SyncIQ job reports are now collected and presented in a separate log, making it easier to see failures on the cluster related to SyncIQ:

1. This error reporting if sent to Support will allow faster resolution and opening of a case with EMC if the root cause is SyncIQ policy failure that cannot be recovered or retried.

- Run Report;
- Resync Prep Report;
- Resync Prep Domain Mark Report;
- Resync Prep Restore Report;
- Resync Prep Finalize Report;

Report added for mirror policy:

- Run Report.

NOTE: If using advanced parallel mode expect more policy reports and errors will potentially exist and need review to see which policy failed. Also note that, sequential failover logs are in order but in parallel mode logging will not be sequential based on failover logic.

© Superna Inc

1.3. Replication Policy Failover Preparation (Data Sync and Config Sync)

Home Top

Replication Policy Failover Preparation (Data Sync and Config Sync)

Failure to Complete Step	Description of Impact to Failover	Recovery Steps	Special Instructions or Warnings
Wait for other failover jobs to complete	Eyeglass only runs one failover job at a time. Failover has not started.	<ol style="list-style-type: none"> 1. Wait until no Failover jobs are present in the running jobs window. 2. Restart failover 	This step will wait for up to two hours in the running state before timing out.
SOURCE get POLICY info	Eyeglass cannot communicate with source cluster. Failover fails.	<ol style="list-style-type: none"> 1. Validate connectivity between Eyeglass and the source cluster. 2. Restart failover 	<p>This step is not run during uncontrolled failover.</p> <p>Data Loss Impact - Failover cannot proceed, data loss scenario until resolved.</p>
Wait for existing policy jobs to complete.	Failover fails if this step times out. Default timeout failover timeout in the current release is 180 minutes, this can be increased with igls cli command see Admin Guide to increase timeout for each step in failover. Also see Best Practices for Failover with Eyeglass .	<ol style="list-style-type: none"> 1. Validate no other failover jobs are running on eyeglass. 2. Restart failover. 3. NOTE: SyncIQ policies that return error from the cluster will require EMC support case to be opened to resolve policies that will not start or run resync prep or execute make writeable operations 	<p>The timeout on this step read from the failover timeout value in minutes.</p> <p>Data Loss Impact - Failover cannot proceed, data loss scenario until resolved.</p>

<p>SOURCE remove schedule POLICY</p>	<p>Eyeglass cannot communicate with source cluster.</p> <p>Failover fails.</p>	<ol style="list-style-type: none"> 1. Validate connectivity between Eyeglass and the source cluster. 2. Restart failover. 	<p>This step is not run during uncontrolled failover.</p> <p>Data Loss Impact - Failover cannot proceed, data loss scenario until resolved.</p>
<p>Replication Policy Failover run all policies</p> <p>(SynclQ Data Sync option enabled on failover job and multiple policies failed over)</p>	<p>Final incremental sync of data on the filesystem has failed.</p> <p>Failover Aborted.</p> <p>Source and targets still in initial state meaning target cluster is read-only still</p>	<ol style="list-style-type: none"> 1. Use the Eyeglass Job Details tree to determine which policies ran successfully, and which failed. 2. Open OneFS on the source to determine the reason for policy failure(s). <ol style="list-style-type: none"> 1. If policy job is still running, wait for it to complete, or cancel the job in OneFS. 2. Manually run the policy, and see if it can be completed. 3. Open Case with EMC for SynclQ jobs failing to run. 4. Restart Failover. 	<p>Eyeglass waits for up to failover over timeout value in minutes for each policy to run. It will abort with a timeout if the incremental sync takes longer than an hour. Wait for this job to finish, then restart the failover job to try again.</p> <p>If the unsynced data on the source filesystem is unimportant and it's more important to get the target cluster file system writable, the Failover can be restarted with the "Data Sync" box unchecked. This will skip the final incremental sync, and will result in data loss for any data not previously replicated with SynclQ.</p> <p>NOTE: This step is not run during uncontrolled failover.</p>
<p>Run Configuration Replication now.</p> <p>(Config Sync enabled in failover job)</p>	<p>Final sync of configuration shares/exports/aliases has failed.</p> <p>Failover wil continue.</p> <p>Source and targets still in initial state meaning target cluster is read-only still</p>	<ol style="list-style-type: none"> 1. Open the Eyeglass Jobs window, and switch to the running jobs tab. 2. Select the most recent configuration 	<p>If the configuration on source and target are known to be identical, the Config Sync option can be unchecked when restarting failover to skip this step.</p> <p>If the unsynced configuration on the source is unimportant and it's more important to get the target cluster file system writable, the Failover can be</p>

		<p>replication job.</p> <ol style="list-style-type: none"> 3. Use the Job Details to determine the reason for failure. 4. Address the config replication failure OR make note of the unsynced configuration data 5. Restart failover. 	<p>restarted with the "Config Sync" box unchecked. This will skip this step.</p> <p>Note This step is not run during uncontrolled failover.</p>
--	--	--	--

© Superna Inc

1.4. DFS Mode

Home Top

DFS Mode

Failure to Complete Step	Description of Impact to Failover	Recovery Steps	Special Instructions or Warnings
DFS Share(s) rename failure on target or source	DFS clients will not switch clusters	<ol style="list-style-type: none">1. Manually remove igls-dfs from share(s) on the target cluster that did not get renamed (consult failover log), will complete the failover and clients will auto switch2. Manually add igls-dfs prefix to share(s) on the source cluster that did not get renamed (consult failover log), will block client access to source and switch to target <p>Following steps apply for Superna Eyeglass Release < 1.9 (2.0 and later will run this steps if a share rename fails)</p> <ol style="list-style-type: none">3. Manually run allow writes from OneFS for the policy(s) that were selected for failover4. Manually run related Quota Jobs from Eyeglass5. Manually run re-sync prep from OneFS for the policy(s) that were selected for failover6. Apply SyncIQ Policy schedule to the target cluster policies that were failed over.	

© Superna Inc

1.5. Networking Operations

Home Top

Networking Operations

Failure to Complete Step	Description of Impact to Failover	Recovery Steps	Special Instructions or Warnings
<p>Rename Source SC (SmartConnect) zone names & aliases</p>	<p>Failover failed during networking step.</p> <p>Auto Rollback of networking will be applied to place SmartConnect Zone names and aliases back on the source cluster.</p> <p>Source and target clusters result in same initial state as before failover and file system will be read/write on the source with SmartConnect Zones return to original configuration</p>	<ol style="list-style-type: none"> 1. Use the info link in the job details tree to determine the reason for failure. (At this point retrying access failover is likely to fail again. you are now switching to manual failover use this table as a guide on the order of the steps that must be run. See table here of steps. 2. You will need to know the steps for these procedures (consult EMC documentation on these steps) 3. Restart failover but only select SynclQ failover job type since Access zone failover will attempt the networking failover again and is likely to fail again. Select the policies that are required to failover. 4. Start failover with policies selected and review the table for manual step order. 	<p>This step is not run during uncontrolled failover.</p>
<p>Modify Source SPNs</p>	<p>SPN operation failure does not abort failover but are logged in the failover log</p> <p>Failover continues.</p>	<ol style="list-style-type: none"> 1. Post failover, open the failover log, and look for instructions listing which SPNs need to be fixed. 2. Manually create/delete SPNs on the source cluster. 3. Use ADSIedit tool and administrator access on the domain to edit source cluster machine account and remove SPN entries (short and long) for each SmartConnect zone that 	<p>SPN operations on the source are proxied through the target cluster ISI commands during failover operations, so source cluster availability does not affect Eyeglasses ability to fix SPN during failover to target</p>

		failed over)	
Rename target SC (SmartConnect) zone names & aliases	<p>Failover failed during networking step.</p> <p>Rollback will be applied.</p> <p>Source and target clusters result in same initial state as before failover with SmartConnect zones reverting to original configuration.</p>	<ol style="list-style-type: none"> 1. Use the info link in the job details tree to determine the reason for failure. (At this point retrying access failover is likely to fail again. you are now switching to manual failover use this table (in the Eyeglass Failover Design Guide) as a guide on the order of the steps that must be run: See table here of steps. 2. You will need to know the steps for these procedures (consult EMC documentation on these steps) 3. Restart failover but only select SyncIQ failover job type since Access zone failover will attempt the networking failover again and is likely to fail again. Select the policies that are required to failover. 4. Start failover with policies selected and review the table for manual step order. 	
Modify Target SPNs	<p>SPN operation failure does not abort failover.</p> <p>Failover continues.</p>	<ol style="list-style-type: none"> 4. Post failover, open the failover log, and look for instructions listing which SPNs need to be fixed. 5. Manually create/delete SPNs on the source cluster. 6. Use ADSIedit tool and administrator access on the domain to edit source cluster machine account and add SPN entries (short and long) for each SmartConnect zone that failed over) 	

1.6. Replication Policy Failover - All policies

Home Top

Replication Policy Failover - All policies

To assist use the failover log browser view to map each table to the correct section of the guide.

State	Job Name	Info
✓	<ul style="list-style-type: none"> [-] Replication Policy Failover - All policies [-] Replication Policy Failover EyeglassRunbookRobot-policy <ul style="list-style-type: none"> Cluster2-7201 allow writes at /ifs/data/robot [-] Replication Policy Failover - Recovery <ul style="list-style-type: none"> [-] Replication Policy Failover EyeglassRunbookRobot-policy <ul style="list-style-type: none"> Cluster-1-7201 resync prep EyeglassRunbookRobot-p... Cluster2-7201 run EyeglassRunbookRobot-policy_mir... Cluster-1-7201 set schedule for policy EyeglassRunbo... Run Quota Jobs Now [-] Replication Policy Failover finalize <ul style="list-style-type: none"> + Finalize quota for path: /ifs/data/robot Set configuration replication for policies to ENABLED 	<p>Info</p> <p>Info</p>

Failure to Complete Step	Description of Impact to Failover	Recovery Steps	Special Instructions or Warnings Or Data Loss Impact
Replication Policy Failover all policies	<p>One or more of the SyncIQ policies in the Failover job did not successfully complete its failover operation.</p> <p>This is the parent task containing all sub policies.</p>	<ol style="list-style-type: none"> 1. Use the table in the next section to determine why the SyncIQ policy did not successfully failover and address the issue. 2. Use the Job Details for the policy failover jobs to determine if at least one policy has successfully executed the step named CLUSTERNAME allow writes POLICY 	<p>This is a parent step containing sub-steps for each replication policy. Recovery from a failure in this step depends on how far the failover proceeded through the child steps.</p> <p>Some steps should be reviewed for data loss impact.</p>

		<p>PATH.</p> <ol style="list-style-type: none"> 1. If yes: the source cluster is failed over. To recover, create a new failover job of type Sync IQ, select all of the policies that failed or did not run, and start the SyncIQ failover job. 2. if no: the cluster is not failed over yet. After fixing the error for the SyncIQ policy, start a new Access Zone Failover job. 3. NOTE: SyncIQ policies that return error from the cluster will require EMC support case to be opened to resolve policies that will not start or run resync prep or execute make writeable operations 	
--	--	--	--

Replication Policy Failover <Policy Name>

Note: On Eyeglass, step names in the following table contain actual name of the cluster(s) and the policy name. These have been replaced with SOURCE, TARGET, and POLICY in the steps below.

To assist use the failover log browser view to map each table to the correct section of the guide.

State	Job Name	Info
✓	Policy failover	
✓	Start Policy failover	
✓	Policy failover -- zone: EyeglassRunbookRobot-AccessZone	
✓	Replication Policy Failover Preparation	
✓	Replication Policy Failover - All policies	
✓	Replication Policy Failover EyeglassRunbookRobot-policy	
✓	Cluster2-7201 allow writes at /ifs/data/robot	
✓	Replication Policy Failover - Recovery	

Failure to Complete Step	Description of Impact to Failover	Recovery Steps	Special Instructions or Warnings OR Data Loss Impact
TARGET allow writes POLICY PATH	<p>The target cannot be put into writes allowed state.</p> <p>NOTE: as of 1.6 all policies involved in a failover job, are issued the make writeable command before Re-sync prep is executed.</p> <p>NOTE: as of 2.0 an attempt to execute make writeable is made against all policies in the failover even in the case where 1 or more have an error on this step.</p> <p>Failover fails (only on releases <2.0).</p>	<ol style="list-style-type: none"> 1. Validate connectivity between Eyeglass and the target cluster. 2. From the OneFS UI, manually allow writes again for policy that failed on the target. Address any errors that may arise on OneFS from this step. 3. NOTE: SyncIQ policies that return error from the cluster will require EMC support case to be opened to resolve policies that will not start 	Data Access Impact - Users will only have access to read-only data for any policy where allow writes failed or was not run due to error.

		<p>or run resync prep or execute make writeable operations</p> <ol style="list-style-type: none"> 4. For any policies where allow writes was not attempted due to an error, manually allow writes for those policy that were not run. 5. If this is a controlled failover, manually run resync_prep on the policy on the the source cluster. 6. If this is a controlled failover, manually set the schedule on the mirror policy. 7. Consider this policy job as a success, and relaunch failover according to the logic in the previous table. 8. (once Resync prep is completed) Manually run related Quota Jobs from Eyeglass or contact support to 	
--	--	--	--

		<p>use post failover quota sync tool.</p> <p>9. NOTE: SynclQ policies that return error from the cluster will require EMC support case to be opened to resolve policies that will not start or run resync prep or execute make writeable operations</p>	
--	--	---	--

© Superna Inc

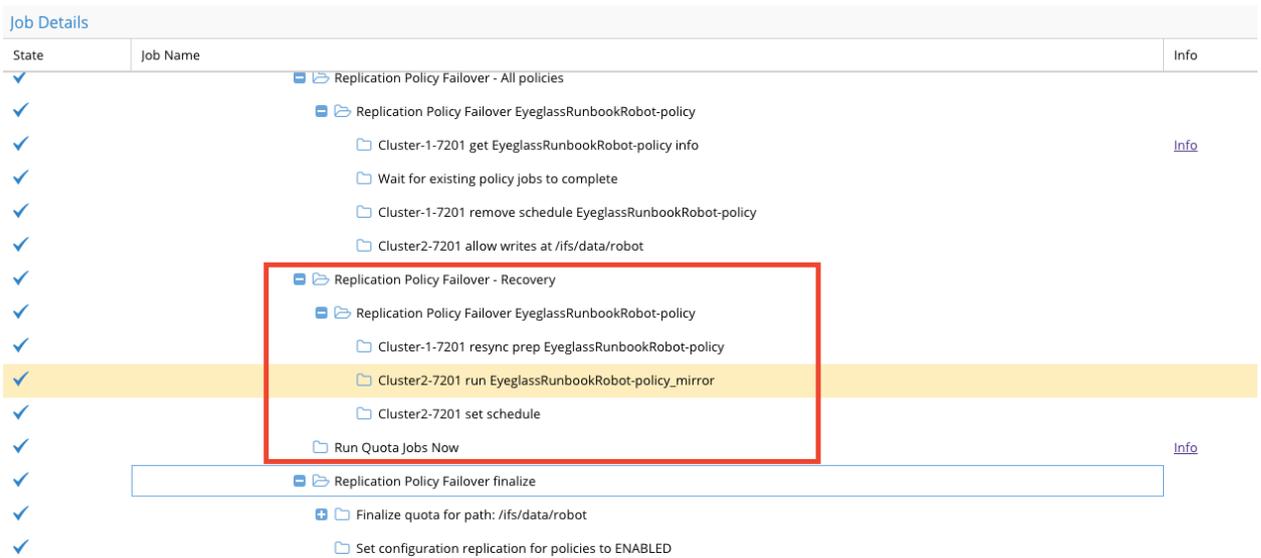
1.7. Replication Policy Failover - Recovery

Home Top

Replication Policy Failover - Recovery

Always review DR Assistant failover log to use this table below.

See screenshot below as reference to column heading name to see how to review the failover log



Failure to Complete Step	Description of Impact to Failover	Recovery Steps	Special Instructions or Warnings or Data loss Impact
Replication Policy Failover - Recovery	<p>One or more of the SyncIQ policies in the Failover job did not successfully complete its failover operation (multiple steps).</p> <p>This is the parent task containing all sub policy steps (see screenshot above to see the order of parent and child steps)</p> <p>IMPORTANT</p> <p>The rollback will not be applied in the event of a failure in this section of the failover. Previously networking and allow-writes steps have succeeded and the cluster is assumed as "failed over". Any policy with an error in this section on resync prep, apply schedule or run policy will</p>	<ol style="list-style-type: none"> 1. Use the table in the next section to determine why the SyncIQ policy did not successfully failover and address the issue. 2. NOTE: SyncIQ policies that return error from the cluster will require EMC support case to be opened to resolve policies that will not start or run resync prep and these steps will need to be completed 	<p>This is a parent step containing sub-steps for each replication policy. Recovery from a failure in this step depends on how far the failover proceeded through the child steps.</p> <p>Data Loss Impact - none (failover for any policy without this step running is completed BUT reprotecting the filesystem is blocked until the mirror policy is created and run successfully)</p>

	need to be completed manually and will likely require EMC technical support to fix root cause on the failure.	manually	
--	---	-----------------	--

Replication Policy Failover <Policy Name>

Note: On Eyeglass, step names in the following table contain actual name of the cluster(s) and the policy name. These have been replaced with SOURCE, TARGET, and POLICY in the steps below.

See screenshot below as reference to column heading name to see how to review the failover log

Failure to Complete Step	Description of Impact to Failover	Recovery Steps	Special Instructions or Warnings
SOURCE resync prep POLICY	<p>The mirror policy cannot be created.</p> <p>Policy is failed over. Target cluster is active.</p> <p>Overall fail back readiness status is failed.</p> <p>NOTE: as of 1.6 all policies involved in a failover job, are issued the make writeable command before Re-sync prep is executed.</p> <p>NOTE: as of 2.0 an attempt to execute resync</p>	<ol style="list-style-type: none"> 1. Login to OneFS on the source cluster. 2. Manually execute resync prep on the policy. Address any errors that may arise. 3. Manually set the schedule on the mirror 	<p>This step is not run during uncontrolled failover.</p> <p>Data Loss Impact - none (failover for any policy without this step running is completed BUT reprotecting the filesystem is blocked until the mirror policy is created and run successfully)</p>

	<p>prep is made against all policies in the failover where make writeable has previously succeeded. In the event of an error on a Resync Prep step failover logic continues and attempts Resync Prep for any policy where make writeable succeeded.</p>	<p>policy.</p> <ol style="list-style-type: none"> 4. Consider this policy job as a success. 5. (once Resync prep is completed) Manually run related Quota Jobs from Eyeglass or contact support to use post failover quota sync tool. 6. NOTE: SyncIQ policies that return error from the cluster will require EMC support case to be opened to resolve policies that will not start or run resync prep 	
<p>TARGET run POLICY_mirror</p>	<p>The mirror policy is unrunnable.</p> <p>Policy is failed over. Target cluster is active.</p> <p>Overall failover status is failure.</p>	<ol style="list-style-type: none"> 1. Login to OneFS on the target cluster. 2. Manually run the mirror policy. Address any errors. 3. Manually set the schedule on the mirror policy. 4. Consider this policy job as a success 5. Contact Superna support for 	<p>This step is not run during uncontrolled failover.</p> <p>Data Loss Impact- none (failover for any policy without this step running is completed BUT reprotecting the filesystem is blocked until the mirror policy is created and run successfully)</p>

		<p>quota run procedures.</p> <p>6. NOTE: SynclQ policies that return error from the cluster will require EMC support case to be opened to resolve policies that will not start or run resync prep</p>	
TARGET set schedule	<p>The schedule on the mirror policy cannot be set.</p> <p>Policy is failed over.</p> <p>Target cluster is active.</p> <p>Overall failover status is failure.</p>	<ol style="list-style-type: none"> 1. Login to OneFS on the target cluster. 2. Manually set the schedule on the mirror policy. 3. Consider this policy job as a success 4. Quotas will not failover contact Superna support. 	<p>Data Loss Impact - none (failover for any policy without this step running is completed BUT reprotecting the filesystem is blocked until the mirror policy is created and run successfully)</p>

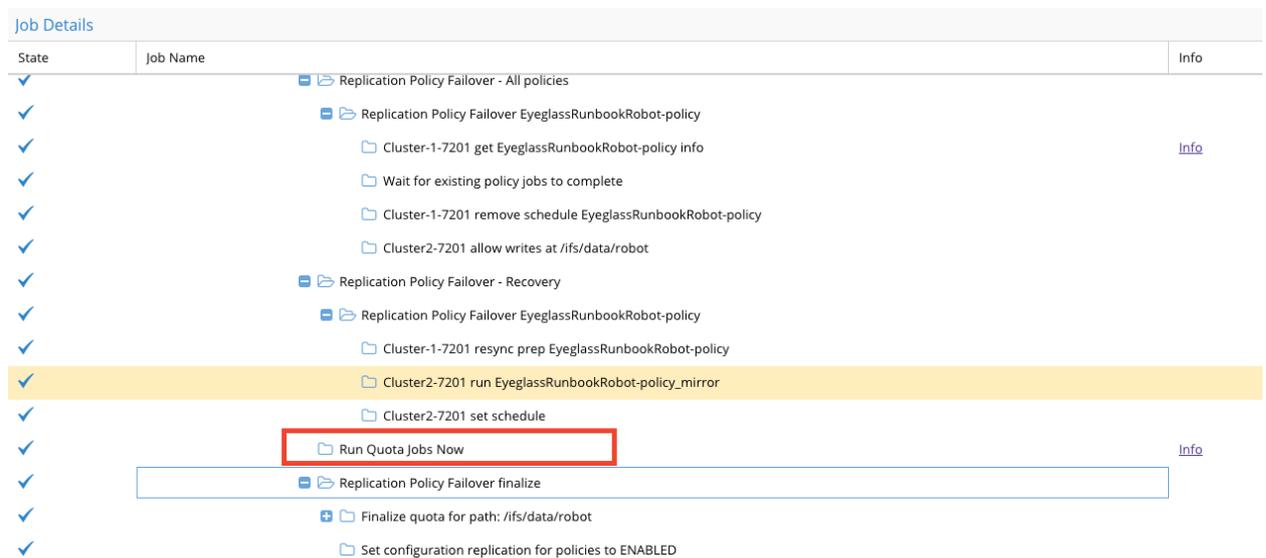
1.8. Run Quota Jobs Now Failover Recovery Steps

Home Top

Run Quota Jobs Now Failover Recovery Steps

NOTE: Quota jobs get applied automatically during failover. Contact Support to use post failover quota sync tool to recover quotas that did not failover.

See screenshot below as reference to column heading name to see how to review the failover log



Failure to Complete Step	Description of Impact to Failover	Recovery Steps	Special Instructions or Warnings
Quota partially failed over	<ol style="list-style-type: none"> 1. Disk limits not applied on target cluster. 2. compare quotas manually on source and target using oneFS UI consult errors in active alarms to assist with identifying quotas that did not 	<ol style="list-style-type: none"> 1. Recreate quotas manually and delete on source cluster 2. Open case to get list of failed quota creates or deletes - requires logs to be uploaded to support case 	<ol style="list-style-type: none"> 1. No data loss impact, quota step runs last in the failover logic

	failover		
--	----------	--	--

© Superna Inc

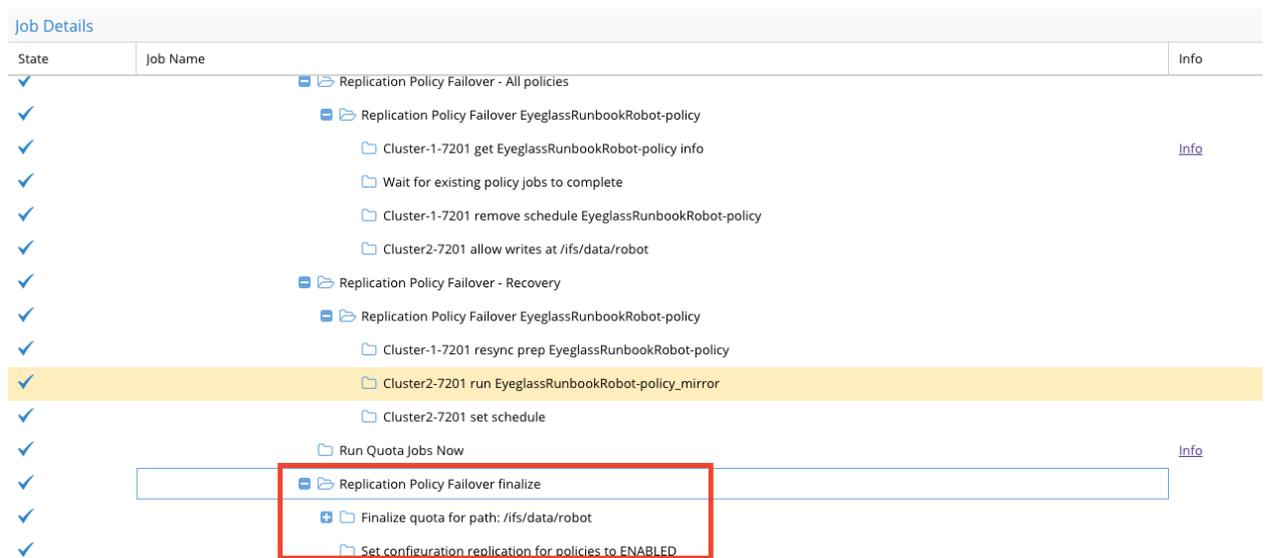
1.9. Replication Policy Failover Finalize

Home Top

Replication Policy Failover Finalize

Note: this step runs one child step per policy, listed as “Finalize quota for path <policy source path>”. The table below describes failures on those steps, and the following should be done for any failed steps, or steps that did not run.

See screenshot below as reference to column heading name to see how to review the failover log



Failure to Complete Step	Description of Impact to Failover	Recovery Steps	Special Instructions or Warnings or Data Loss Impact
Finalize quota for path: PATH Delete Quotas on Source	Could not delete quotas from source. Policy is failed over. Target cluster is active. Failover overall status is failure.	<ol style="list-style-type: none"> 1. On SOURCE OneFS, find all quotas on data protected by the SyncIQ policy. 2. Validate that those quotas are present on the TARGET cluster. 3. Delete these quotas from 	This step is not run during uncontrolled failover. Data Loss Impact - none (failover for any policy without this step running is completed BUT reprotecting can be impacted by the presence of quotas on the source cluster)

		the SOURCE.	
Set configuration replication for policies to ENABLED	<p>Could not enable Eyeglass Configuration Replication Jobs</p> <p>Policy is failed over. Target cluster is active.</p> <p>Failover overall status is failure.</p>	<ol style="list-style-type: none"> 1. Open the Eyeglass Jobs window. 2. Select configuration replication job and Enable. 3. Use log to determine the reason for failure. 	<p>This step enables the newly configured mirror policies post failover (if they did not exist). Eyeglass will detect the new policy and it will be enabled post failure to replicate configuration data back to the source cluster.</p> <p>Data Loss Impact - none (if this step fails, it block config from syncing back to source cluster and can be enabled manually from jobs window.</p>

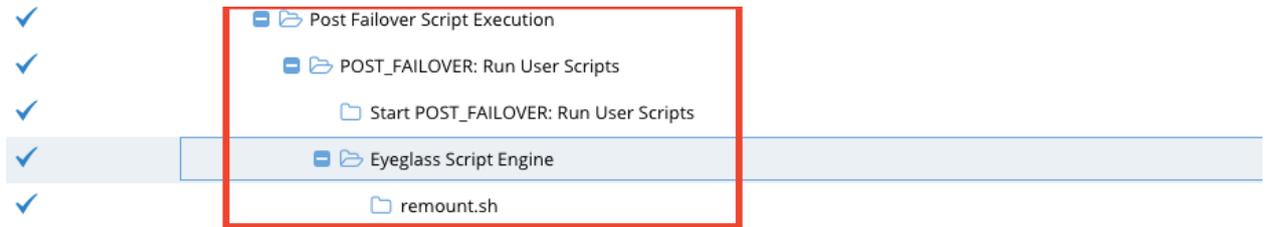
© Superna Inc

1.10. Post Failover Script Execution

Home Top

Post Failover Script Execution

See screenshot below as reference to column heading name to see how to review the failover log



Failure to Complete Step	Description of Impact to Failover	Recovery Steps	Special Instructions or Warnings or Data Loss Impact
Eyeglass Script Engine	<p>A user supplied post-failover script failed.</p> <p>Failover overall status is failure.</p>	<ol style="list-style-type: none"> 1. Use the script engine to fix errors in the failing scripts, and re-run those that failed. 2. Use test script function to validate output and error codes returned to failover jobs 	<p>This step relies on user-supplied implementations.</p> <p>Review the script output to verify if it executed correctly. Error codes set by the script should fail the failover job if set correctly. See Pre Post Failover Scripting Guide on proper script exit code values to indicate failure vs successful execution.</p> <p>Data Loss Impact - This step should only result in failure to remount or start applications post failover. Logs should be reviewed to ensure all steps completed and correct any script failures manually if they failed.</p>

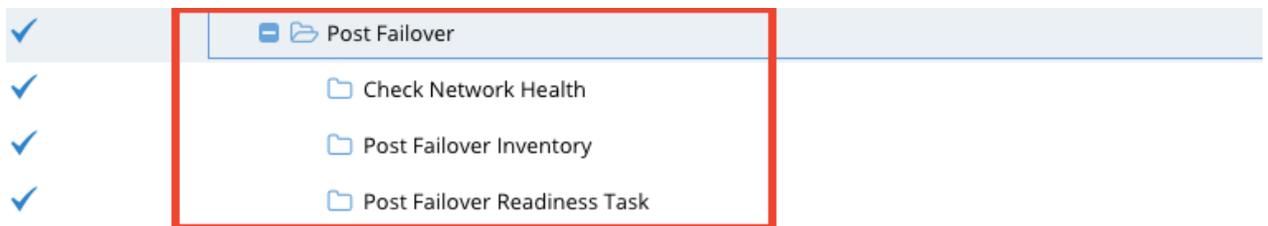
1.11. Post Failover

Home Top

Post Failover

This step completes after all critical steps are executed, if all steps passed to this point then no data loss condition can result from failures. This section is used for SmartConnect roll back to source cluster (no steps for DNS if dual delegation configured).

See screenshot below as reference to column heading name to see how to review the failover log



Failure to Complete Step	Description of Impact to Failover	Recovery Steps	Special Instructions or Warnings or Data Loss Scenario
Check Network Health	<p>The Networking Rollback job could not be initiated.</p> <p>Is an impact only if the failover did not reach the make writeable step of a policy.</p>	<ol style="list-style-type: none"> 1. Use the failover log to determine the networking operations that were performed during failover. 2. Use OneFS to manually revert the networking operations on SOURCE and TARGET as required. 	<p>The rollback logic is only executed if the failover job failed before the make writeable step on a policy.</p> <p>Data Loss Impact - In this scenario failover never completed and the source cluster data is still the production copy. Potential data loss scenario of the source data was deemed not usable. Contact EMC to get assistance correcting make writeable step also consult best practises which cover scenario's that block make writeable step. Best Practices for Failover with Eyeglass</p>
Post Failover Inventory	<p>The post failover Inventory job failed.</p> <p>Failover is successful, Eyeglass UI may be out of</p>	<ol style="list-style-type: none"> 1. Validate source and target cluster connectivity with Eyeglass. 	<p>This step is not run during uncontrolled failover.</p> <p>Data Loss Impact - None. This step updates the</p>

	date.	<ol style="list-style-type: none"> 2. Open the alarms window, and look for any alarms related to configuration replication. 3. Manually run configuration replication, or wait until the next automatic cycle. 	Eyeglass UI.
Post Failover Readiness Task	<p>The post failover Readiness job failed.</p> <p>Failover is successful, Eyeglass UI may be out of date.</p>	<ol style="list-style-type: none"> 1. Validate source and target cluster connectivity with Eyeglass. 2. Manually run the Access Zone Readiness job, or wait until the next automatic cycle. 	<p>This step is not run during uncontrolled failover.</p> <p>Data Loss Impact - None. This step updates the Eyeglass UI.</p>

1.12. Check Client Access (Manual Step)

Home Top

Check Client Access (Manual Step)

Failure to Complete Step	Description of Impact to Failover	Recovery Steps	Special Instructions or Warnings
DNS SmartConnect zone validation	Note: Dual delegation switches DNS automatically with networking failover available in Access Zone Failover	<p>nslookup SmartConnect zone name</p> <ol style="list-style-type: none"> 1. Confirm the ip address returned is the target cluster ip pool 2. Repeat for all SmartConnect zones that were failed over 3. If incorrect then update target cluster by using isi command to create the missing alias on the ip pool 4. If ip returned is still the source cluster, then rename source cluster SmartConnect zone name to ensure dual delegation will not use this cluster's DNS service to answer queries 	
Refresh session to pick up DNS change	SMB Client unable to access data on Failover Target cluster despite successful failover and DNS Updates and session refresh	<p>Check SPNs using ADSI Edit tool and confirm that</p> <ol style="list-style-type: none"> 1. Failover Source Cluster - SPN for SmartConnect Zone that Client is using does NOT exist. 2. Failover Target Cluster - SPN for SmartConnect Zone that Client is 	<p>You cannot create a missing SPN on the Active Cluster if it still exists for the Failed Over cluster. You need to remove from Failover Over cluster first and then add to active cluster.</p>

		<p>using DOES exist.</p> <p>3. If above condition is not met, using ADSI Edit to update SPN to be on correct cluster.</p>	
SMB Direct Mount Shares	Dual delegation updates DNS but requires clients to remount and query DNS to get a new cluster ip address	<pre>net use //sharename /delete</pre> <pre>net use //sharename</pre> <p>Or use map network drive in Explorer</p>	
NFS mounts	Dual delegation updates DNS but requires clients to remount and query dns to get a new cluster ip address	<pre>umount -fl /path of mount -a (reads fstab file and remounts, does force and lazy unmount to handle open files)</pre>	
DFS clients	auto switch	no action needed	

Copyright Superna LLC 2017

© Superna Inc

2. Cluster Move Procedure to New Data Center

Home [Top](#)

- [Cluster Move Procedure to New Site - Change Cluster IP Address](#)
- [Cluster Move Procedure to New Site - Cluster IP Unchanged](#)

© Superna Inc

2.1. Cluster Move Procedure to New Site - Change Cluster IP Address

Home Top

Cluster Move Procedure to New Site - Change Cluster IP Address

This procedure is used to move a cluster to a new site and failover, failback and a change of the Cluster IP address is required at the new site.

Steps need to follow while moving one cluster to a new site:

1. Shares or exports exist on active cluster (writeable data) and on the read only cluster (Cluster B)
2. Failover one policy or access zone at a time to the target cluster (Cluster B)
 1. **Expected Result:** Clients and applications are successfully accessing data from the new active cluster (Cluster B)
1. Completes without error
2. Shares and exports now exist on new active cluster (Cluster B) on the new read only cluster (cluster A)
3. Disable all Eyeglass jobs
4. Move source (Cluster A) to new site
5. Change IP address on Cluster A
6. Bring source (Cluster A) back online
7. Change source (Cluster A) IP address from Eyeglass inventory view using the below procedure:

[How to - Edit PowerScale IP Address in Eyeglass](#)

8. Eyeglass jobs should all be enabled at this step. Confirm this is the case.
 1. **Expected Result:** Eyeglass Configuration Replication Jobs are running without error

Note : If Eyeglass Configuration Replication Jobs are in error proceed to Step 11 to reset Eyeglass to pick up IP address changes for Jobs on the new active cluster (Cluster A)

9. Fallback one policy at a time or access zone to the source (Cluster A)
 1. **Expected Result:** Clients and applications are successfully accessing data from the new active cluster (Cluster A)
10. **Note :** If Eyeglass Replication Jobs are in error proceed to Step 11 to reset Eyeglass to pick up IP address changes for Jobs on the new active cluster (Cluster A) If the Eyeglass Configuration Replication Jobs and Fallback is working successfully on new active cluster (Cluster A) are working then the relocation procedure is complete.
11. If Eyeglass Replication Jobs are in error reset Eyeglass to pick up IP address changes for Jobs on the new active cluster (Cluster A) with the following procedure:
 1. SSH to Eyeglass appliance using admin: sudo -s enter (must use root) then use admin password (default password: 3y3gl4ss)
 2. Run command <igls appliance rediscover>
 3. Once reset completes, go to the chrome browser and refresh the browser and login with the credentials
 4. Login to jobs window and validate that all jobs are in the correct states, Eyeglass will restore all job status but should be double checked,
 1. Correct any if jobs in user disabled state or not in DFS state that should be DFS enabled.
 2. **IMPORTANT: You must enable DFS mode on an job that did not rediscover and shows in the DFS job section of jobs window before enabling the Job to prevent creation of active shares on target cluster. Enable DFS mode first, then enable.**
 5. Using Bulk Actions select a policy in jobs window and use Run now.
 1. Using Running jobs tab to watch and verify successful (check mark no red X) on the configuration sync
 6. If using Access Zone failover, verify that failover readiness job is enabled.
 1. Using bulk actions force run the readiness job wait until it completes (running jobs window)

7. Check DR Dashboard to verify it shows a good DR status

1.

© Superna Inc

2.2. Cluster Move Procedure to New Site - Cluster IP Unchanged

Home Top

Cluster Move Procedure to New Site - Cluster IP Unchanged

Steps need to follow while moving one cluster to a new location: using same network setting (same cluster IP address)

Before moving the cluster to a new location, configuration replication jobs in Eyeglass should be working properly and configuration synced to target cluster.

1. Disable configuration replication jobs in Eyeglass before turning off the cluster being moved.

Note: For Failover scenario, the failover away from the cluster that is going to be moved should be completed successfully

- config replication jobs that synced the configuration from source to target in disable mode

- Jobs on the active cluster are green.

Login to Eyeglass → click on *Jobs* icon in desktop → *Job Definitions* → check the box to select jobs in *Configuration Replication:Share, Export, Alias replication* and/or *Configuration Replication:Access Zone replication* (if applicable) and/or

Configuration Replication:DFS mode (if applicable) and/or

Configuration Replication:Snapshot Schedules (if applicable) and/or

Failover: Runbook Robot section (if applicable) → Select bulk a action → Enable/Disable to disable the jobs

2. Then Turn off the cluster to be moved.
3. After the cluster is moved to the new site and powered up, check the connectivity between the cluster and Eyeglass appliance. (Hint: you may ping the cluster IP from Eyeglass.)

4. Enable Configuration Replication:Share,Export, Alias replication jobs in Eyeglass.

Login to Eyeglass → click on *Jobs* icon in desktop → *Job Definitions* → check the box to select jobs in *Configuration Replication:Share, Export, Alias replication* section → Select bulk a action → Enable/Disable to enable the jobs.

Note: Enable any *Snapshot Schedules, Configuration Replication:Access Zone replication, Failover: Runbook Robot* jobs were enabled before cluster move.

5. Run the configuration for one replication cycle to verify configuration replication jobs working properly.

click on *Jobs* icon in desktop → *Job Definitions* → check the box to select jobs in *Configuration Replication:Share, Export, Alias replication* section → Select bulk a action → *Run now*

Steps need to follow while moving one cluster to a new location: using same network setting (different cluster IP address)

Before moving the cluster to a new location, configuration replication jobs in Eyeglass should be working properly and the configuration synced to the target cluster.

1. Disable configuration replication jobs in Eyeglass before turn the source cluster off. For Failover scenario. Failover should be completed successfully - config replication jobs that synced the the configuration in disable mode.

Login to Eyeglass → click on *Jobs* icon in desktop → *Job Definitions* → check the box to select jobs in *configuration replication:shares, exports, alias replication* section → Select bulk a action → Enable/Disable to disable the jobs

Notes:

DFS jobs should be disabled; if any, before turn off the cluster.

From Eyeglass 1.6.3, *Snapshot Schedules* and *Access Zone replication* should be disabled.

2. Turn off the Source cluster.
3. Reset the Eyeglass appliance.

Reset Eyeglass to re-add the cluster (different IP address will be added to Eyeglass for the moved cluster) .

SSH to Eyeglass appliance using admin: `sudo -s` enter (must use root) then use admin password (default password: 3y3gl4ss)

1. `cd /opt/superna/sbin`
2. `./reset.sh`
3. Once reset completes, login to chrome browser and refresh the browser and login with the credentials.
4. Now, add both of clusters using Management subnet SSIP.
4. After the cluster is moved to the new site and powered up, check the connectivity between the cluster and Eyeglass appliance. (Hint: you may ping the Eyeglass IP.)
5. Enable Configuration Replication:Share,Export, Alias replication jobs in Eyeglass.

Login to Eyeglass → click on *Jobs* icon in desktop → *Job Definitions* → check the box to select jobs in *Configuration Replication:Share, Export, Alias replication* section → Select bulk a action → Enable/Disable to enable the jobs.

Note: Enable any *Snapshot Schedules* and *Access Zone replication* jobs were enabled before cluster move.

6. Run the configuration for one replication cycle to verify configuration replication jobs are working properly.

click on *Jobs* icon in desktop → *Job Definitions* → check the box to select jobs in *Configuration Replication:Share, Export, Alias replication* section → Select bulk a action → *Run now*

Useful link:

[How to - Edit PowerScale IP Address in Eyeglass](#)

Copyright 2017 Superna LLC

© Superna Inc

3. PowerScale Upgrade Procedure with Eyeglass

[Home](#) [Top](#)

- [ReadMe First - Prior to PowerScale Upgrade](#)
- [Cluster OneFS Upgrade is between Major Release Version Eyeglass Procedures](#)
- [Cluster OneFS Upgrade is between Minor Release Version Eyeglass Procedures](#)
- [Functions Impacted DURING a Cluster OneFS Upgrade and Eyeglass, Ransomware Defender, Easy Auditor, Performance Auditor Procedures](#)

© Superna Inc

3.1. ReadMe First - Prior to PowerScale Upgrade

Home Top

Prior to PowerScale Upgrade Determine which process to follow

Check the Feature Release Compatibility table in the appropriate [Eyeglass Release Notes](#) to ensure cluster release compatibility listing is supported for the cluster configuration upgrade for the source and destination cluster.

There are procedures to complete before and after a OneFS upgrade, the different scenario's are upgrade OneFS to a Major new release or a Minor release. Follow the guide based on your scenario.

1. If your upgrade process will always have the source and destination cluster in a supported configuration AND upgrade does not change version first 2 digits (for example 7.2.1.1 upgrade to 7.2.1.2 or 8.0.x to 8.0.x, 9.x), follow instructions for *[Supported Eyeglass functionality when a Cluster OneFS Upgrade is between Minor Release Version](#)*.
2. If your upgrade process is for OneFS 7.2.x to OneFS 8.0.x or 8.2, 9.x follow instructions for *[Supported Eyeglass functionality when a Cluster OneFS Upgrade is between Major Release Version](#)*.
3. If your upgrade process will result in source and destination cluster being in an unsupported configuration at any point follow instructions for *[Eyeglass functionality that will be Impacted DURING a Cluster OneFS Upgrade](#)*.

© Superna Inc

3.2. Cluster OneFS Upgrade is between Major Release Version Eyeglass Procedures

Home Top

- [Description](#)
- [How To Video](#)
- [Read Me - Upgrades to 8.2 or later releases](#)
- [Upgrade Procedure](#)
 - [Upgrade Target Cluster First Step:](#)
 - [Upgrade Source Cluster Second Step:](#)
- [What you need to know in Mixed mode Onefs Releases on Source and Target Clusters](#)

Description

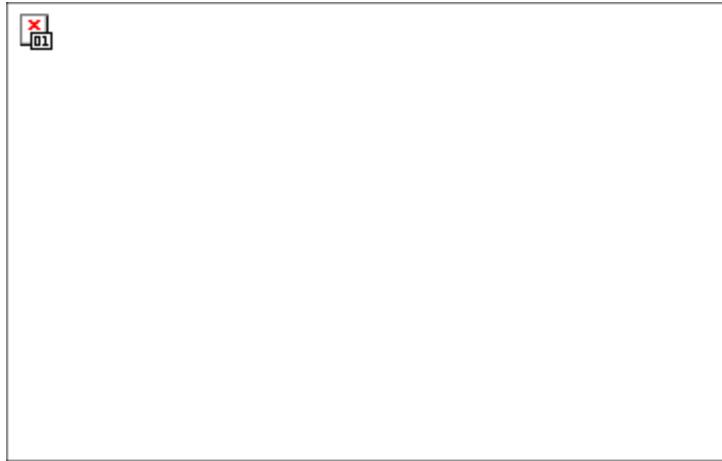
The following is the procedure for performing site cluster rolling upgrade from one release to a new major release of OneFS.

How To Video

Read Me - Upgrades to 8.2 or later releases

1. **NOTE:** As of OneFS 8.2 or later the cluster does not support API calls to the SSIP (subnet service IP). Please ensure you are using an IP address from a pool that has System Access Zone assigned to it (Example: IP from the Management Pool). The pool should be configured as a dynamic pool so that the IP address will failover to other nodes in the cluster.
2. You should change the Eyeglass IP before upgraded to 8.2 or later releases. Use the steps below.

- a. Login to Eyeglass WebUI
- b. Goto Inventory View
- c. Right click on cluster each cluster, a context menu should open
- d. Select Edit on the context menu



- e.
- f. In the Edit Managed Device window, select **ALL** text (do not edit the ip address replace all text) in the SmartConnect Service IP field and remove it and replace it with the appropriate IP address only for the given cluster.
- g. Enter the **eyeglass** service account password used to add the cluster
- h. Click Submit

 A screenshot of a web application window titled "Edit Managed Device". The window has a blue header and a light gray sidebar on the left with the word "Isilon" selected. The main content area is titled "Edit Isilon Cluster" and shows a form for "Isilon prod8". The form contains the following fields:

- SmartConnect Service IP: A text input field containing "prod8-172.31.1.106.igls.local" with a blue selection highlight over the entire text.
- Port: A text input field containing "8080".
- Username: A text input field containing "eyeglass".
- Password: An empty text input field.
- Maximum RPO Value: A text input field containing "5".

 At the bottom right of the form, there are two buttons: "Submit" and "Cancel".

- i.

Upgrade Procedure

To upgrade both site's clusters to a new Major OneFS release by performing **upgrade on 1 site at a time**. The plan is to **upgrade the Target Cluster first** and then **follow by upgrade the Source Cluster**.

Upgrade Target Cluster First Step:

1. Verify that prior to upgrade, the Eyeglass has been configured properly and Eyeglass Jobs statuses are OK. No error.
2. Prepare Target Cluster for Upgrade. Refer to EMC PowerScale documentation for detailed instructions, including:
3. Pre-upgrade checks
 - a. Supported upgrade paths from current OneFS release to the target release. Might need to upgrade to intermediate release before final upgrade to the target release.
 - b. Download and store the PowerScale installer to the Cluster (e.g. /ifs/data folder).
4. Disable Eyeglass Jobs temporary for upgrade process.
5. Perform OneFS upgrade on Target Cluster. Eyeglass will report that the cluster is unreachable.
6. Once upgrade process of Target Cluster has completed successfully, perform the following tasks to let Eyeglass manage the Target Cluster that now has new OneFS Release:
 - a. Update the Eyeglass Admin role privileges on this Target Cluster by adding the additional permissions list in the document below:
 - b. [PowerScale Cluster User Minimum Privileges for Eyeglass](#)
 - c. Update the sudo privileges for the Eyeglass service account on this Target Cluster and review the document above.
 - d. Restart Eyeglass sca service:
 - e. SSH as admin user.
 - f. Then sudo -s.
 - g. Enter admin password.
 - h. To stop: `systemctl stop sca.service`

- i. To start: `systemctl start sca.service`
 - j. Verify: `systemctl status sca.service`
 - k. Re-enable the Eyeglass Jobs.
7. Manually Run with Run Now bulk actions menu from Jobs Window or wait until the next Replication Job has started. Once Jobs have completed, verify Eyeglass status is OK. **No error.**
8. Verify all Jobs are green and DR Dashboard audit and readiness for zones status
 - a. **NOTE: At this point in the site cluster upgrade path, Eyeglass is now in mixed mode, which means OneFS API will use the lower cluster version api and takes priority when syncing configuration information from source to target cluster.**
9. Verify product functionality by following steps [here](#).

Upgrade Source Cluster Second Step:

1. Verify that prior to upgrade, the Eyeglass Jobs statuses are OK. No error.
2. Target Cluster for Upgrade. Refer to EMC PowerScale documentation for detailed instructions.
3. Perform Eyeglass Failover to Target Cluster. (see documentation for the mode of failover and follow the [Failover Planning Guide and checklist](#)).
1. Verify Failover has completed successfully. No errors in the failover log.
2. By default, all Eyeglass Jobs for Mirror Policies are in User Disabled state. If not disabled (e.g, have done failover and failback before), manually disable all Eyeglass Jobs temporarily during the cluster upgrade process.
3. Perform OneFS upgrade on Source Cluster. Eyeglass will report that the cluster is unreachable, this is expected.
4. Once upgrade process of Source Cluster has completed successfully, perform the following tasks to let Eyeglass manage the Source Cluster.

5. Update the Eyeglass Admin role privileges on this Source Cluster by adding additional permissions listed in the document below:
 - a. [PowerScale Cluster User Minimum Privileges for Eyeglass](#)
 - b. Update the sudo privileges for the Eyeglass service account on this Source Cluster:
 - c. **Additional sudo permissions should be reviewed in the document above:**
 - i. Restart Eyeglass sca service:
 1. SSH as admin user.
 2. Then sudo -s
 3. Enter admin password.
 - ii. To stop:
 1. `systemctl stop sca.service`
 - iii. To start:
 1. `systemctl start sca.service`
 - iv. Verify:
 1. `systemctl status sca.service`
 - v. Re-enable the Eyeglass Jobs from the Jobs icon
 - vi. Run or wait till the next Replication Jobs have started. Once Jobs have completed, verify Eyeglass status is OK. No error.
 - vii. Verify product functionality by following steps [here](#).

What you need to know in Mixed mode Onefs Releases on Source and Target Clusters

In this mixed mode (Any cluster version combination with 8.x 8.2 or 9.x added to the same eyeglass)

- Config is supported between clusters running lower version to higher version OneFS releases. To ensure you are running a supported release, consult the release notes for exact versions supported.
- Failover is supported.
- Failback is supported.

- Runbook Robot is supported.
- RPO reporting is supported.

NOTE: It is expect this mixed mode will be a temporary scenario before upgrading the source cluster (expectation is a period would be weeks to months at most). It is not intended to be a normal operating mode between replicating clusters

NOTE: if the only one pair of clusters managed by Eyeglass are mixed mode is still active for all clusters. Mixed mode API is disabled only once all clusters managed by Eyeglass are upgraded to the same Major release.

© Superna Inc

3.3. Cluster OneFS Upgrade is between Minor Release Version Eyeglass Procedures

Home Top

- [Follow all Steps below when upgrading between minor Onefs versions](#)

Follow all Steps below when upgrading between minor Onefs versions

1. Login to Eyeglass web page prior to upgrade and confirm that configuration replication is in a good state.
2. Begin the upgrade process. Eyeglass is in a running state during upgrade process.
3. If source or destination cluster is not reachable at any point during upgrade you may see failed Eyeglass Configuration Replication Jobs and their related alarms. Subsequent Eyeglass configuration replication Job execution with source and target clusters in reachable state should clear these alarms.
4. Eyeglass Alarms during Isilon/PowerScale Upgrade
 - a. The following provides a summary of the Eyeglass alarms that you might encounter during an PowerScale upgrade where source and/or target cluster are unreachable:
 - b. One or both clusters may not reachable when Eyeglass Replication Task is starting
 - i. Replication jobs failed.

- ii. DR Dashboard has Error status and readiness alarms will be sent
- iii. Multiple Alarms will be raised for each scheduled replication job and this is expected

5. Post Upgrade of Onefs

- a. Verify product functionality by following steps [here](#).
- b. It is recommended to leave Eyeglass running for 30 minutes after the upgrade procedure has been completed and then check the status of the Eyeglass Jobs to confirm that system has returned to pre-upgrade state. Any alarms that occurred can be seen in the Alarms History.

6. done

© Superna Inc

3.4. Functions Impacted DURING a Cluster OneFS Upgrade and Eyeglass, Ransomware Defender, Easy Auditor, Performance Auditor Procedures

Home Top

Eyeglass functionality that will be Impacted DURING a Cluster OneFS Upgrade

1. Configuration Sync will have errors during upgrade, they can be ignored until the Cluster upgrade is complete
2. Runbook Robot will not execute without errors
3. Ransomware Defender and Easy Auditor will not be able to collect events during upgrade.
4. Planned Failover execution will not be supported during a failover.
5. DR Readiness Dashboard will NOT be accurate
6. All Reports will not execute correctly

Post OneFS Upgrade Steps

1. Eyeglass DR Edition
 - a. Restart the SCA service
 - b. ssh to appliance as admin user
 - c. sudo -s
 - d. enter admin password
 - e. systemctl restart sca
 - f. Wait 30 minutes
 - g. Login to gui and verify no errors in jobs or DR dashboard
2. Ransomware Defender, Easy Auditor and Performance Auditor
 - a. login to node 1 of the ECA cluster

- i. ecactl cluster down
 - b. once all steps to shutdown complete
 - i. ecactl cluster up
 - c. Wait for all steps to complete
 - d. Login to Eyeglass, open Managed service icon and verify all nodes show green (you may need to wait 5 minutes)
 - i. Ransomware Defender run Security guard, verify it completes successfully
 - ii. Easy Auditor run Robot Audit, verify it completes successfully
 - iii. Performance Auditor - open the tool icon to verify performance data is shown
3. done

© Superna Inc

4. Eyeglass Simulated Disaster Event Test Procedure

[Home](#) [Top](#)

- [Introduction](#)
- [Initial Environment Setup](#)
- [Verify Environment Setup](#)
- [Simulated Disaster Scenario - DFS Test SyncIQ Policy Failover](#)
- [Simulated Disaster Scenario - EyeglassRunbookRobot Access Zone Failover](#)

© Superna Inc

4.1. Introduction

Home Top

Introduction

This document outlines Eyeglass simulated disaster test scenario for clients who want to perform disaster testing during a scheduled maintenance window:

1. Controlled failover with production SynclQ policies, and uncontrolled failover with a DFS mode Test SynclQ policy, **without** impacting or **exposing** production data to data loss or resync risks.
2. Controlled failover of production Access Zone, and uncontrolled failover with EyeglassRunbookRobot Access Zone, **without** impacting or **exposing** production data to data loss or resync risks.
3. Optional: this can be used to only perform simulated failover without any production data failed over within the same maintenance window.
4. Open a case with Superna support to review this document prior to attempting the procedures within this document

Support Statement on Use of this Procedure

1. **This procedure is the only supported process.** Any variant of the process that uses uncontrolled failover on production data is unsupported. If a support request is raised for a test that intentionally used uncontrolled failover on production data, the customer will have to take responsibility for recovery steps using documentation without support assistance.
2. **This procedure requires the “[Failover Planning Guide and checklist](#)” to be followed to maintain support as per the support contract for planned failovers.** The Failover Planning Guide and checklist will be requested for validation from Superna support when any case is opened regarding failover.

Important Note

Note: Production SynclQ policies should target and protect directories in Access Zones **other** than the EyeglassRunbookRobot Access Zone or the test DFS SynclQ policy.

4.2. Initial Environment Setup

[Home](#) [Top](#)

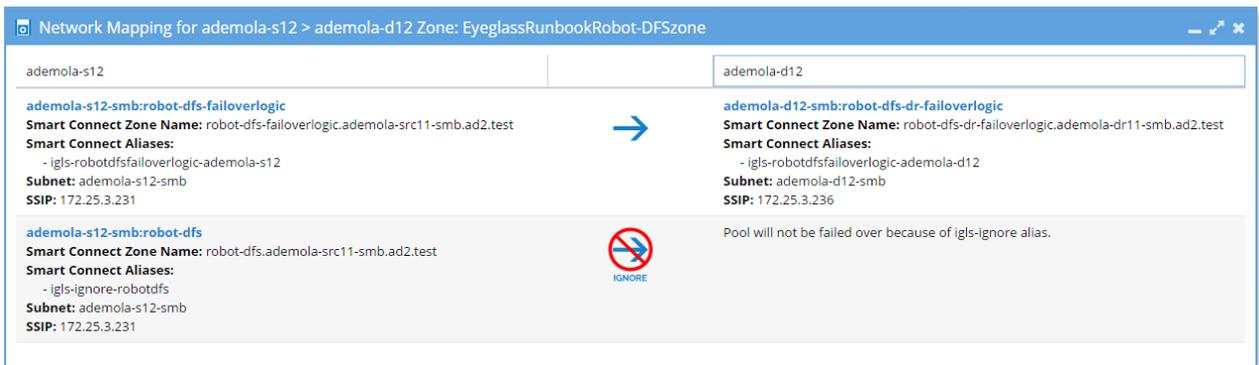
Initial Environment Setup

If you have already configured the Eyeglass RunbookRobot feature in your environment for Access Zone or DFS Continuous DR Testing, you may skip this initial environment setup section and proceed to [Verify Environment Setup](#) section.

NOTE: ONLY 1 RUNBOOK ROBOT SUPPORTED PER EYEGLOSS

1. **Both (DFS, Access Zone):** On both Cluster1 and Cluster2, create an Access Zone with name format "EyeglassRunbookRobot-XXXX", where XXXX is a string or number of your choice.
2. **Both (DFS, Access Zone):** On both Cluster1 and Cluster2, as best practice, create an IP pool dedicated for SynclQ data replication.
 1. Note that the *Replication IP pool should be in the System Access Zone*. While configuring your production or test SynclQ policies, for the "Restrict Source Nodes" option, make sure to select the second option that says "Run the policy only on nodes in specified subnet and pool", then select the IP pool that you have dedicated to SynclQ data replication.
 2. For the Replication IP pool, make sure to configure a SmartConnect Zone Alias with syntax `igls-ignore-xxxx`, where `xxxx` is a string that makes the alias (as a whole string) unique across your infrastructure.
3. **Both (DFS, Access Zone):** On both Cluster1 and Cluster2, create an IP pool for clients to access data in the EyeglassRunbookRobot-XXXX Access Zone.
 1. The IP pool for DFS client access will be configured with SmartConnect zone alias of the format `igls-ignore-xxxx`
 2. The Access Zone failover logic pool will be configured with SmartConnect zone alias of the format `igls-aaaa-bbbb`, where `aaaa` is the same string on Cluster1 and Cluster2, and `bbbb` make the alias (as a whole string) unique across your infrastructure. Eyeglass uses the first two sections (`igls-aaaa`) to map the pool from Cluster1 to Cluster2.

See below a sample pool mapping for the DFS and Access Zone as seen from Eyeglass Zone Readiness:



4. **Both (DFS, Access Zone)** Create Test SynclQ policy on Cluster1 with name format "EyeglassRunbookRobot-yyyy", where yyyy is a number or string of your choice.

Note that you can only have **one** SynclQ policy within each EyeglassRunbookRobot-XXXX Access Zone.

1. Make the test SynclQ policy target host the SmartConnect zone FQDN of the dedicated SynclQ replication IP Pool on Cluster2.
2. Restrict Source Nodes option should be selected when configuring the test SynclQ policy.
3. The test SynclQ policy source directory path should be below the EyeglassRunbookRobot access zone base directory.
4. Run the policy on OneFS UI once it has been created.
5. **Both (DFS, Access Zone)** On Cluster1 EyeglassRunbookRobot Access Zone, create test SMB shares (and quotas if required) at folder paths below the Test SynclQ policy source folder.

1. **DFS Mode:** In DFS snapin for AD configure the DFS folder targets using FQDN(s) for Cluster1 and Cluster2 IP pool for client access.

Example format for DFS folder targets are:

<\\fqdn-of-DFS-testdata-IPpool-on-cluster1\SMBsharename>, and <\\fqdn-of-DFS-testdata-IPpool-on-cluster2\SMBsharename>. Windows client will mount this DFS share using the path <\\AD domain name\dfsrootname\dfsfoldername>.

6. **DFS Mode:** On Eyeglass, change the configuration replication job associated with Test DFS policy to DFS mode. (See product documentation for details)

7. **Both (DFS, Access Zone)** In order to have a Test SynclQ mirror-policy in place (and also as best-practices), using Eyeglass, perform an initial **controlled failover** (Cluster1 to Cluster2) of the EyeglassRunbookRobot Access Zone or Test DFS SynclQ policy, followed by a **controlled fallback** (Cluster2 to Cluster1) of the EyeglassRunbookRobot Access Zone or Test DFS SynclQ policy.

See sample screenshot below showing that mirror policies are in place for all the Jobs.

Job Name	Policy	Type	Last Run Date	State
Configuration Replication: DFS mode (AUTOMATIC)				
ademola-s12_EyeglassRunbookRobot-DFS	EyeglassRunbookRobot-DFS	AUTODFS	4/7/2017, 12:05:23 A...	OK
ademola-s12_test-dfs-siq	test-dfs-siq	AUTODFS	4/6/2017, 10:21:29 A...	Policy Disabled
ademola-s12_dfs-az1-siq1	dfs-az1-siq1	AUTODFS	4/6/2017, 10:51:46 A...	Policy Disabled
ademola-d12_EyeglassRunbookRobot-DFS_mirror	EyeglassRunbookRobot-DFS_mirror	AUTODFS	4/6/2017, 3:56:56 PM	Policy Disabled
ademola-d12_test-dfs-siq_mirror	test-dfs-siq_mirror	AUTODFS	4/7/2017, 12:05:23 A...	OK
ademola-d12_dfs-az1-siq1_mirror	dfs-az1-siq1_mirror	AUTODFS	4/7/2017, 12:05:23 A...	OK

8. **Access Zone Mode:** Must have [DNS Dual Delegation](#) in place and [AD permissions for cluster to read and write SPNs](#).

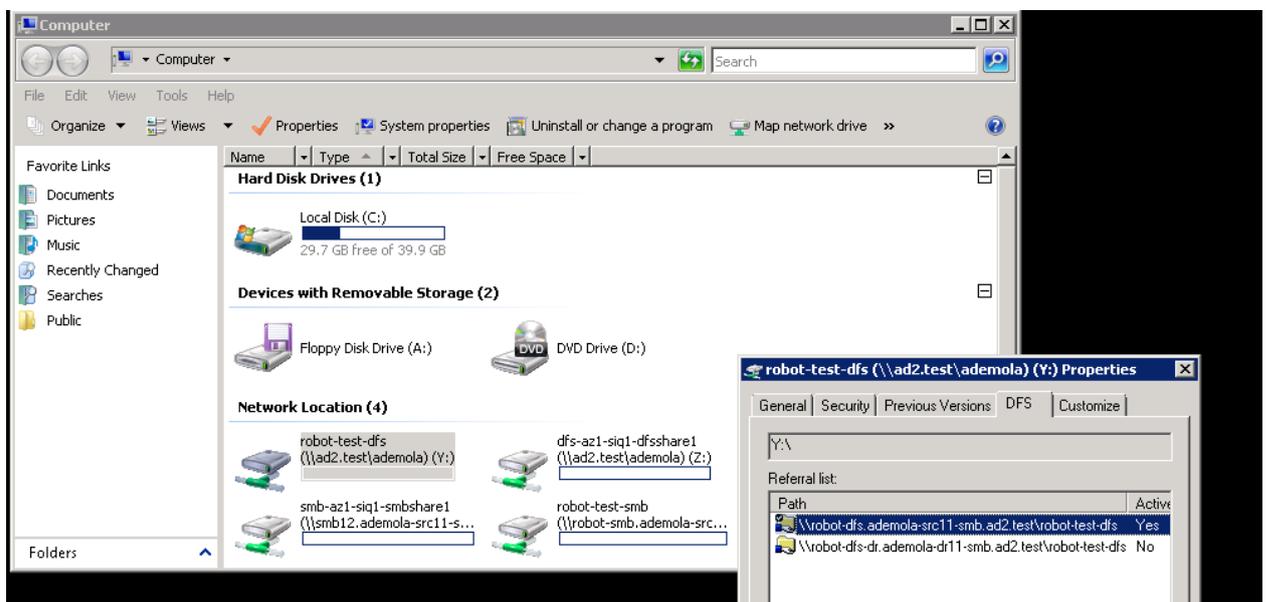
© Superna Inc

4.3. Verify Environment Setup

Home Top

Verify Environment Setup

1. **Verify production DFS policies** are setup correctly with dual folder targets.
2. **Verify test DFS policy:** Write data to the share added to the DFS namespace which is protected by the EyeglassRunbookRobot DFS SyncIQ Policy on a Windows Client before attempting the uncontrolled failover. Verify using DFS tab on properties of the DFS folder name in Explorer to confirm that Cluster 1 contains the active folder target.



3. **Verify test non-DFS policy:** Repeat write test of the non DFS policy in the Access Zone. This will verify DNS resolution to the correct cluster.
4. **Access Zone Mode:** Verify AD permissions by following the document: [How to Validate AD Cluster Delegation is Ready for Failover and failback of SPNs.](#)

© Superna Inc

4.4. Simulated Disaster Scenario - DFS Test SyncIQ Policy Failover

Home Top

Simulated Disaster Scenario - DFS Test SyncIQ Policy Failover

Use this procedure to simulate a DFS policy failover using uncontrolled mode to simulate a real DR event. This assumes a DFS mode policy has been enabled inside the Runbook Robot Access Zone.

Note: This test can be done with or without production data failover in the same maintenance window.

***Note:** Before implementing the following simulated disaster scenario DFS Test SyncIQ policy failover steps, please make sure you have followed instructions/steps in the [“Important Note”](#), [“Initial Environment Setup”](#), [“Verify Environment Setup”](#) and [“Support Statement”](#) sections.*

Pre Simulated Disaster - Cluster1 (prod cluster) is available - Controlled Failover

1. Review all steps in the [Failover Planning Guide and checklist](#) before beginning. This is required to maintain support for this procedure. See [support statement](#) above on planning guide requirement.
2. Perform Microsoft DFS *controlled failover* for Production SyncIQ policies (the uncontrolled test policy will NOT be failed over at this step) from Cluster1 to Cluster2 using Eyeglass.
3. On Eyeglass, enable the Production SyncIQ mirror-policy job in the Jobs window if it is in USERDISABLED state post failover.
4. Write data to production shares protected by Production SyncIQ Policy from DFS mount (confirm that Cluster2 share path is the active target) on Windows Client after controlled failover.
5. **Production data controlled failover completed.** See below sample DFS Readiness at this stage.

DR Dashboard							
Policy Readiness		Name	SyncIQ Policy	Source	Destination	Last Successful Readiness Check	DR Failover Status
Zone Readiness	<input type="checkbox"/>	ademola-s12_EyeglassRunbookRobot-DFS	EyeglassRunbookRobot-DFS	ademola-s12	ademola-d12	4/6/2017, 11:45:4...	
DFS Readiness	<input type="checkbox"/>	ademola-s12_test-dfs-siq	test-dfs-siq	ademola-s12	ademola-d12	4/6/2017, 11:45:4...	
DR Testing	<input type="checkbox"/>	ademola-s12_dfs-az1-siq1	dfs-az1-siq1	ademola-s12	ademola-d12	4/6/2017, 11:45:4...	
	<input type="checkbox"/>	ademola-d12_EyeglassRunbookRobot-DFS_mirror	EyeglassRunbookRobot-DFS_mirror	ademola-d12	ademola-s12	4/6/2017, 11:45:4...	
	<input type="checkbox"/>	ademola-d12_test-dfs-siq_mirror	test-dfs-siq_mirror	ademola-d12	ademola-s12	4/6/2017, 11:45:4...	
	<input type="checkbox"/>	ademola-d12_dfs-az1-siq1_mirror	dfs-az1-siq1_mirror	ademola-d12	ademola-s12	4/6/2017, 11:45:4...	

[Generate SyncIQ Job Charts](#)

Simulated Disaster - Cluster1 (prod cluster) becomes unavailable - Uncontrolled Failover

This procedure simulates a source cluster that has been destroyed or is unreachable on the network for a long period of time and requires a failover to the secondary site.

Note: This step will only operate against test policy and Access Zone created in [initial setup](#) section only, to maintain access to support for this procedure.

1. Simulate Cluster 1 failure:
 1. On Cluster1 OneFS UI, remove the node interfaces from dedicated IP pool used for *client access* to EyeglassRunbookRobot-DFSzone Access Zone (NOTE: perform only on this 1 IP pool). Consult EMC documentation. *The DFS folder target path from Cluster1 will now be failed when the node interfaces are removed from the pool.*
 2. Step 1a. above simulates DNS response failure to Cluster1 EyeglassRunbookRobot-DFSzone Access Zone, without actually impacting SSIP or normal DNS operations. At this point in the process, name resolution is down, and NetBIOS sessions are disconnected from the Cluster1 EyeglassRunbookRobot-DFSzone Access Zone.

```
ademola-igls4:/home/admin # nslookup robot-dfs.ademola-src11-smb.ad2.test
Server:          127.0.0.1
Address:         127.0.0.1#53

** server can't find robot-dfs.ademola-src11-smb.ad2.test: SERVFAIL
ademola-igls4:/home/admin #
```

Notice from the above screenshot, that name resolution to the SmartConnect zone name is not resolving as expected (SERVFAIL is returned). **At this point we have simulated a disaster** as Cluster1's EyeglassRunbookRobot-DFSzone Access Zone

SmartConnect zone name resolution is failing, and no shares can be access on Cluster1 on this Access Zone.

3. **Set the schedule for the EyeglassRunbookRobot DFS policy on the source cluster to manual.** As the policy won't be able to run anyway if the source cluster has been destroyed. **Do not proceed until this step is done.**
1. Perform DFS *uncontrolled failover* for EyeglassRunbookRobot Test DFS mode SyncIQ policy from Cluster1 to Cluster2 using Superna Eyeglass. (see documentation for more details)

Failover Wizard

The Failover Wizard will guide you through the process of failing over your data and configurations from a source to a target cluster.

Select your failover type, source cluster, and failover settings to begin.

Select Failover Type

SyncIQ Policy
 Access Zone
 Microsoft DFS

Select Source Cluster

Source Cluster:

Select Failover Options

Controlled failover

 Data sync

 Config sync

 SyncIQ Resync Prep

 Disable SyncIQ Jobs on Failover Target

Back Next

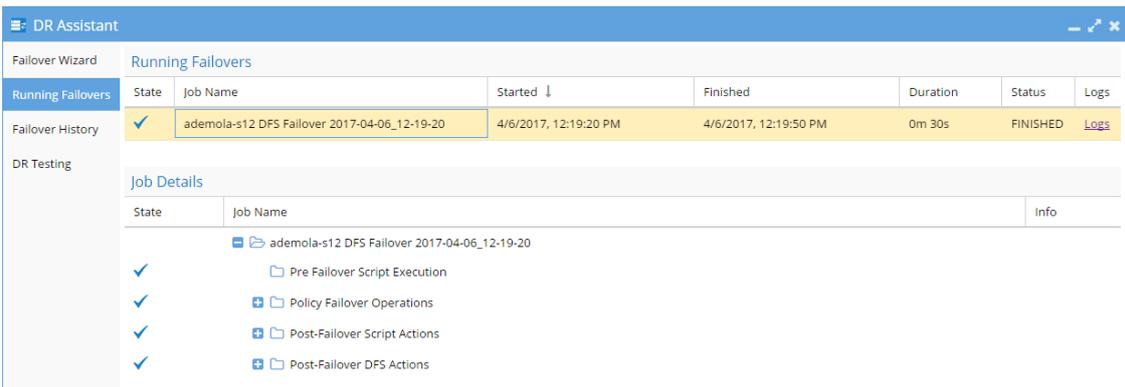
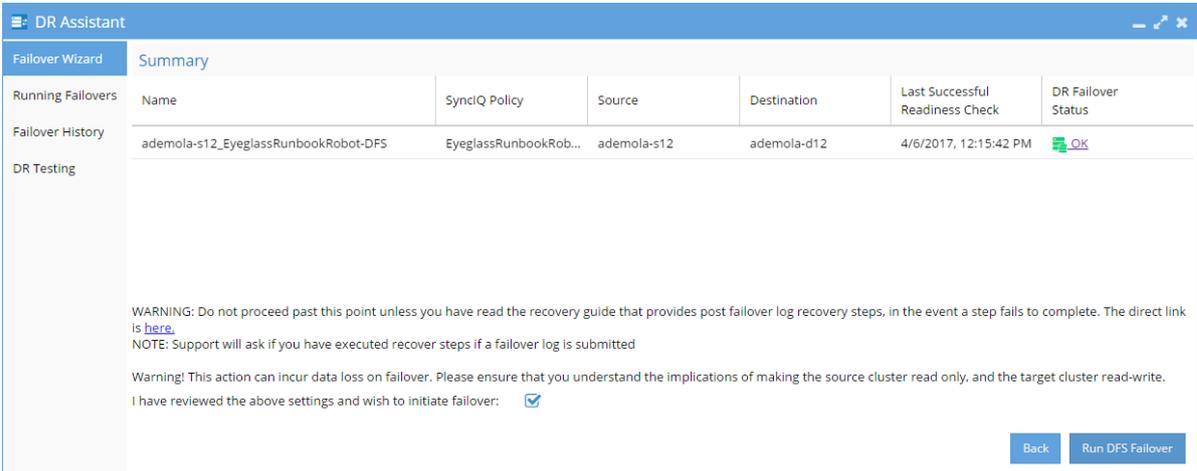
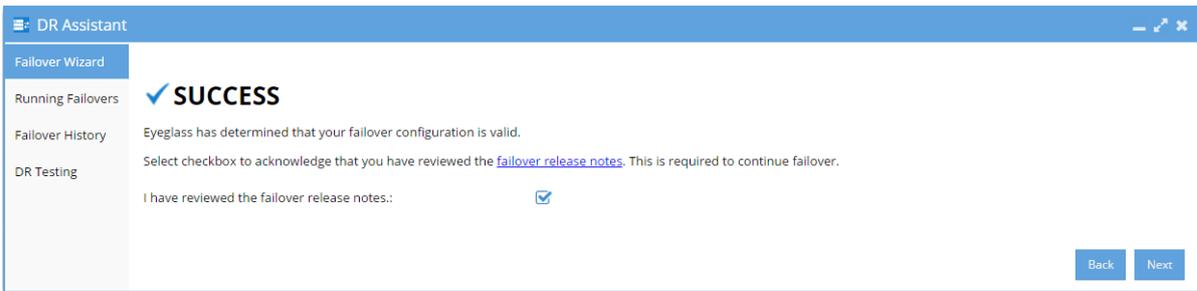
Best Practices

- Run domain mark manually in advance to ensure a fast failover. See [this url](#) for details.
- Failback operations are executed on the clusters (resync Prep), this runs 4 steps, two on the source and two on the target cluster that can take time to complete (see domain mark optimization). To increase performance and reduce the time taken to process the resync prep steps, it is recommended to increase the SyncIQ worker threads from the default to 10 or more.

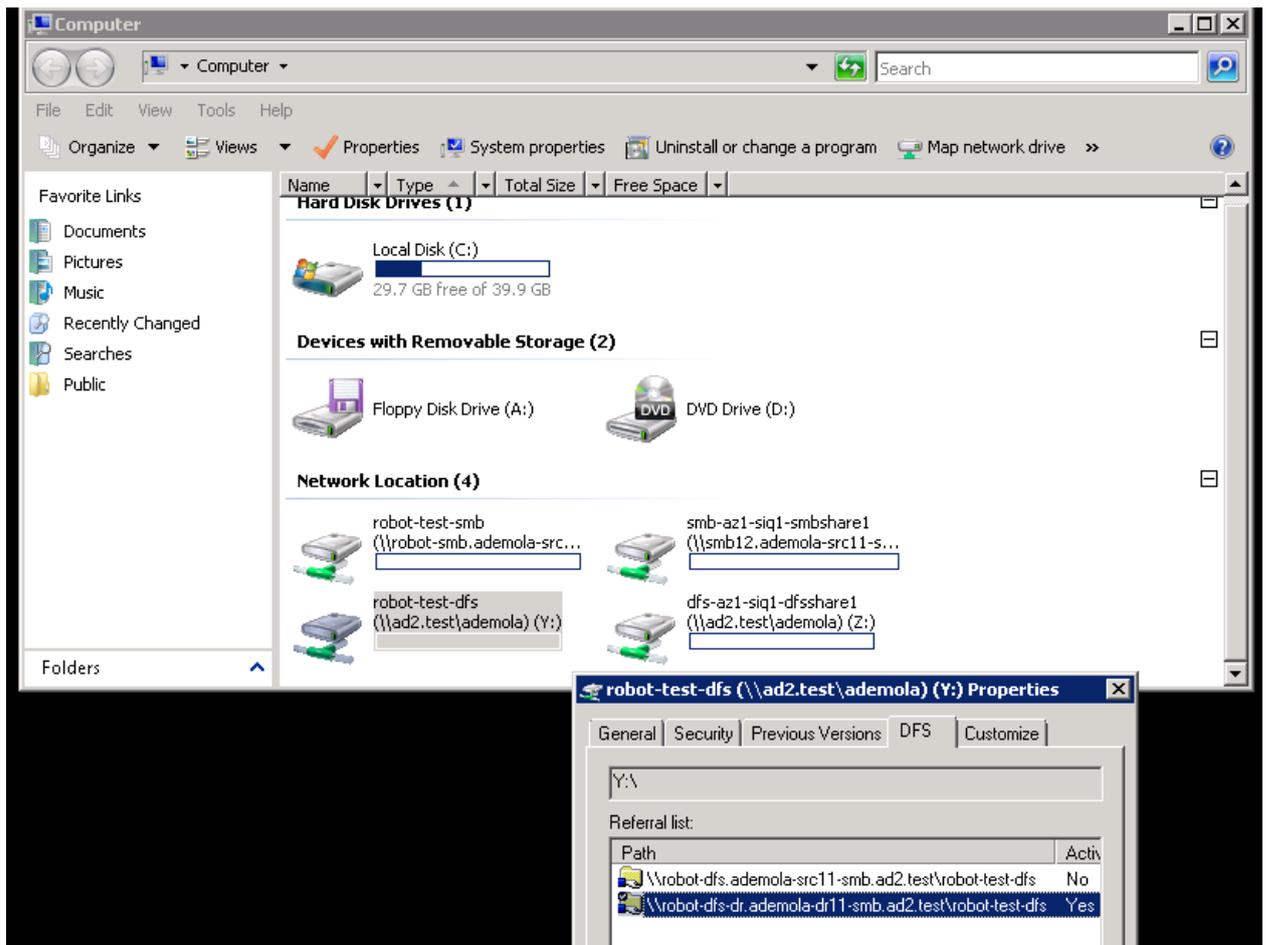
Back Next

	Name	SyncIQ Policy	Source	Destination	Last Successful Readiness Check	DR Failover Status
Running Failovers	<input checked="" type="checkbox"/> ademola-s12_EyeglassRunbookRobot-DFS	EyeglassRunbookRobot-DFS	ademola-s12	ademola-d12	4/6/2017, 12:15:42 PM	OK
Failover History	<input type="checkbox"/> ademola-s12_test-dfs-siq	test-dfs-siq	ademola-s12	ademola-d12	4/6/2017, 12:15:42 PM	FAILED_OVER
DR Testing	<input type="checkbox"/> ademola-s12_dfs-az1-siq1	dfs-az1-siq1	ademola-s12	ademola-d12	4/6/2017, 12:15:42 PM	FAILED_OVER

Back Next



2. Wait until the uncontrolled failover completes.
3. Write data to share protected by EyeglassRunbookRobot Test DFS SyncIQ Policy from the DFS mount (confirm that Cluster2 share path is the active target) on Windows Client after uncontrolled failover.



4. Uncontrolled DFS Failover is complete.

Post Simulated Disaster - Cluster1 (prod cluster) Recovery Steps for DFS

These steps are executed to restore the uncontrolled policies to a working state. The production data is currently failed over to Cluster 2 using controlled failover. Some customers may choose to stay on Cluster 2 as production for some period of time before planning a failback. The test policies can be recovered by following the steps below:

1. **Simulate Cluster 1 returning to Service :**
 1. On Cluster1 OneFS UI, rename shares within Test SyncIQ policy path to have igls-dfs-<sharename> format (this step should happen after "**uncontrolled failover**" step)

2. On Cluster1 OneFS UI, reconnect previously removed node interfaces back to IP pool used for DFS client access to test data on EyeglassRunbookRobot-DFSzone Access Zone.
2. On Cluster1 OneFS UI, run resync-prep on EyeglassRunbookRobot-DFS Test SynclQ policy (consult EMC Documentation).
3. Verify that resync-prep process was completed without error before proceeding to next steps.
 1. Check on OneFS SynclQ reports tab to make all steps pass successfully.
 4. Check the job state in Eyeglass
 1. From Eyeglass, verify both policies on Cluster 1 and Cluster 2 and re-enable the Eyeglass job for the EyeglassRunbookRobot-DFS Test Policy on Cluster 1 and the mirror policy on Cluster 2 in the Jobs icon.
 2. Allow Eyeglass Configuration Data Replication to run at least once.
 3. From Eyeglass Jobs-->"Running Jobs" window, verify that Eyeglass Configuration Data Replication in step 4b above has completed without errors.

Note: As stated in step 4b above, Eyeglass Configuration Data Replication task must complete before continuing with steps below.

4. Verify Eyeglass jobs show policy state correctly with **Cluster 1 policy** showing policy **Disabled** and **Cluster 2** showing **Enabled** and OK (green).
5. Wait for Config sync to correctly show the above state.
6. Do not continue until the above validations are done.
5. Perform Microsoft DFS-type **controlled failback** from Cluster 2 to Cluster 1 for EyeglassRunbookRobot DFS Test SynclQ mirror-policy using Superna Eyeglass DR Assistant.

DR Assistant

Failover Wizard

Running Failovers: The Failover Wizard will guide you through the process of failing over your data and configurations from a source to a target cluster.

Failover History: Select your failover type, source cluster, and failover settings to begin.

DR Testing

Select Failover Type

SyncIQ Policy
 Access Zone
 Microsoft DFS

Select Source Cluster

Source Cluster:

Select Failover Options

Controlled failover

- Data sync
- Config sync
- SyncIQ Resync Prep
- Disable SyncIQ Jobs on Failover Target

Back Next

DR Assistant

Failover Wizard

Running Failovers

Failover History

DR Testing

Best Practices

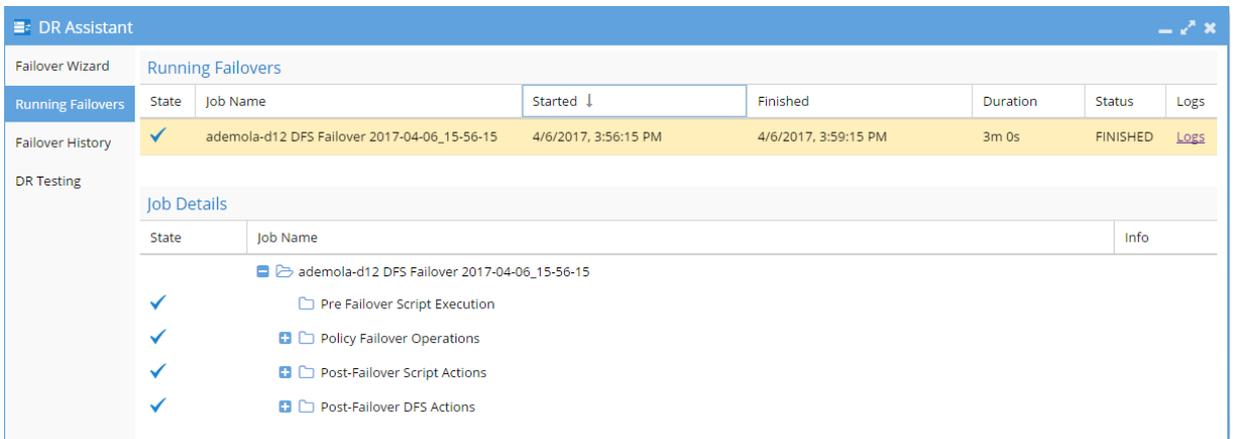
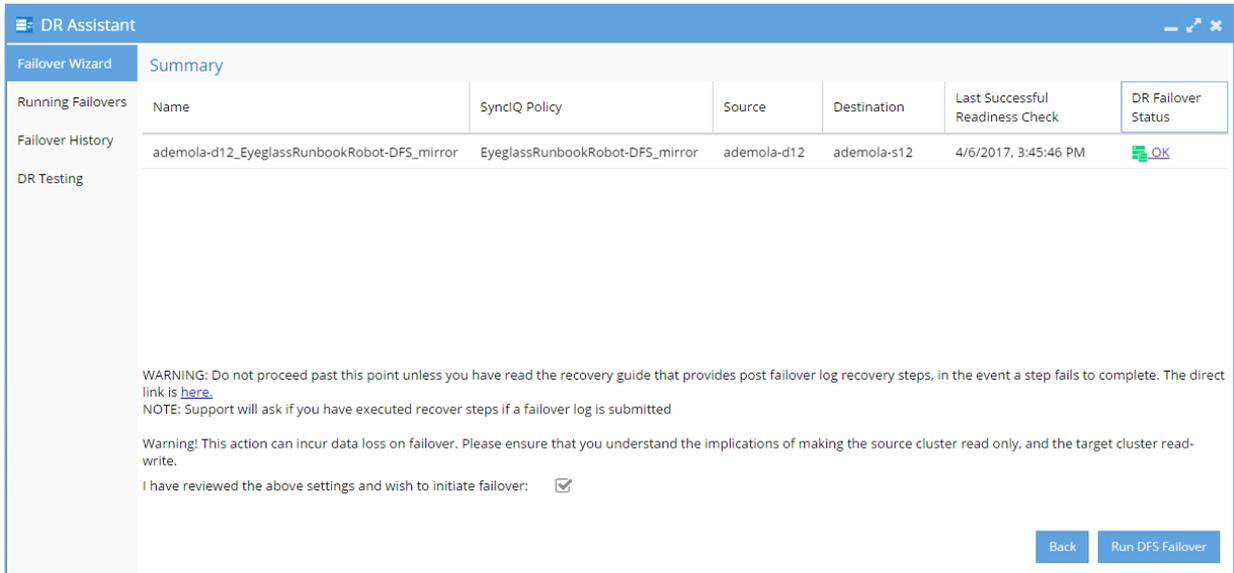
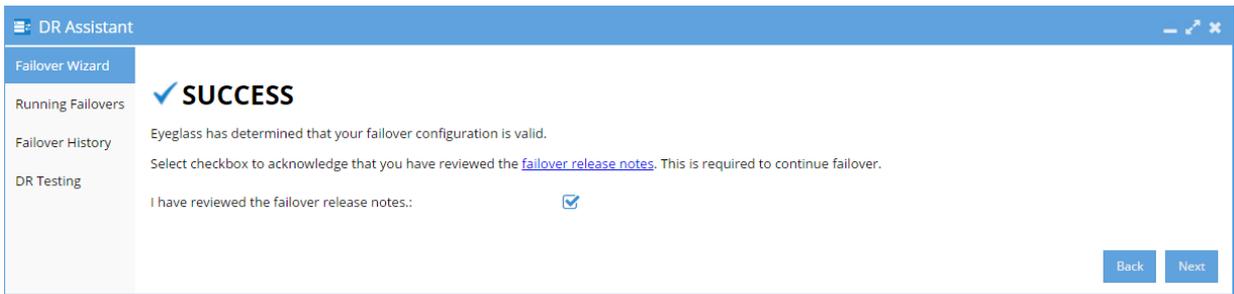
- Run domain mark manually in advance to ensure a fast failover. See [this url](#) for details.
- Failback operations are executed on the clusters (resync Prep), this runs 4 steps, two on the source and two on the target cluster that can take time to complete (see domain mark optimization). To increase performance and reduce the time taken to process the resync prep steps, it is recommended to increase the SyncIQ worker threads from the default to 10 or more.

Back Next

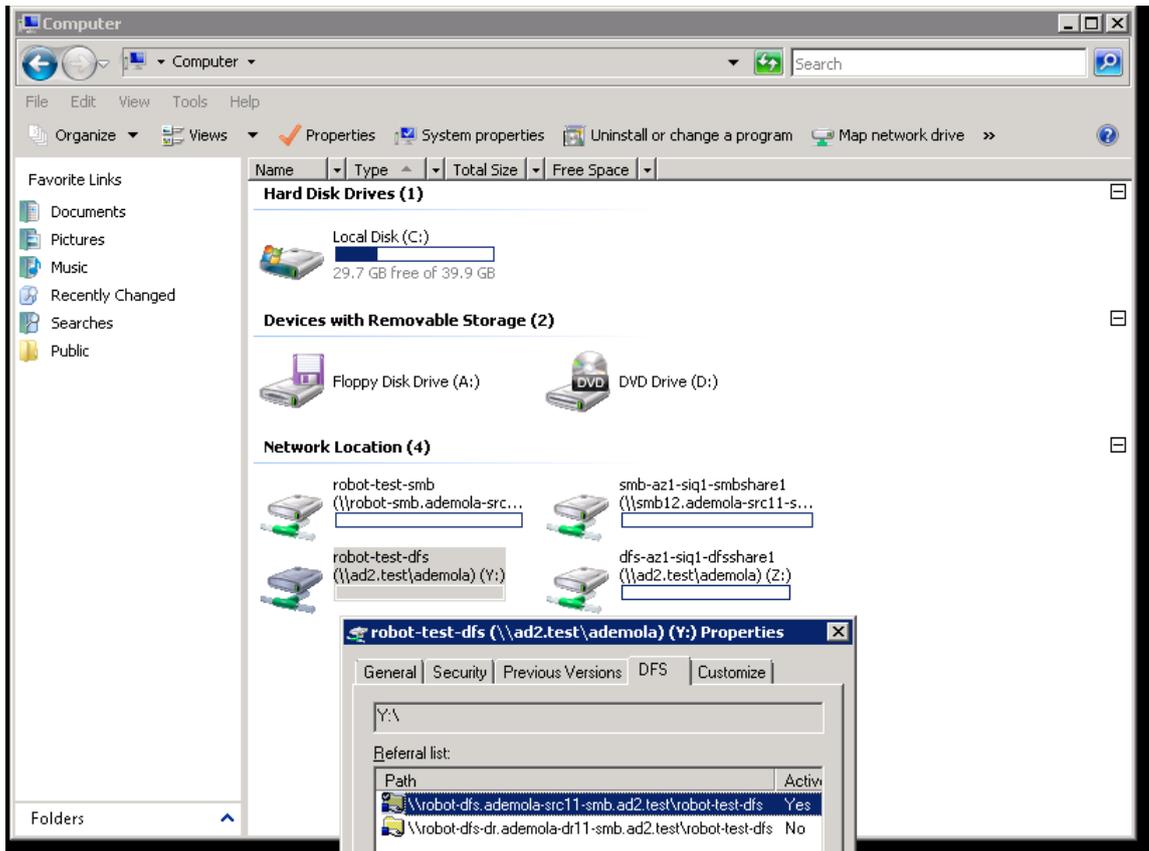
DR Assistant

Failover Wizard	<input type="checkbox"/> Name	SyncIQ Policy	Source	Destination	Last Successful Readiness Check	DR Failover Status
Running Failovers	<input checked="" type="checkbox"/> ademola-d12_EyeglassRunbookRobot-DFS_mirror	EyeglassRunbookRobot-DFS_mirror	ademola-d12	ademola-s12	4/6/2017, 3:45:46 PM	OK
Failover History	<input type="checkbox"/> ademola-d12_test-dfs-siq_mirror	test-dfs-siq_mirror	ademola-d12	ademola-s12	4/6/2017, 3:45:46 PM	OK
DR Testing	<input type="checkbox"/> ademola-d12_dfs-az1-siq1_mirror	dfs-az1-siq1_mirror	ademola-d12	ademola-s12	4/6/2017, 3:45:46 PM	OK

Back Next



6. Wait for Failover to complete,
7. Write data to the share protected by EyeglassRunbookRobot DFS Test SyncIQ Policy from a DFS mount (confirm that Cluster1 share path is the active target) on Windows Client after **controlled** failback.



8. Perform Microsoft-DFS-type *controlled failback* of all Production SyncIQ mirror-policies from Cluster 2 to Cluster 1 using Superna Eyeglass.
9. Write data to share protected by Production SyncIQ Policy from DFS mount (confirm that Cluster1 share path is the active target) on Windows Client after controlled failback.

© Superna Inc

4.5. Simulated Disaster Scenario - EyeglassRunbookRobot Access Zone Failover

Home Top

Simulated Disaster Scenario - EyeglassRunbookRobot Access Zone Failover

Use this procedure to simulate an Access Zone failover using uncontrolled mode, to simulate a DR event. This assumes dual delegation has been implemented and the Runbook Robot Access Zone is fully functional.

Note: This test can be done with or without production data failover in the same maintenance window.

Note: Before implementing the following simulated disaster scenario Access Zone failover steps, please make sure you have followed instructions/steps in the [“Important Note”](#), [“Initial Environment Setup”](#), [“Verify Environment Setup”](#) and [“Support Statement”](#) sections.

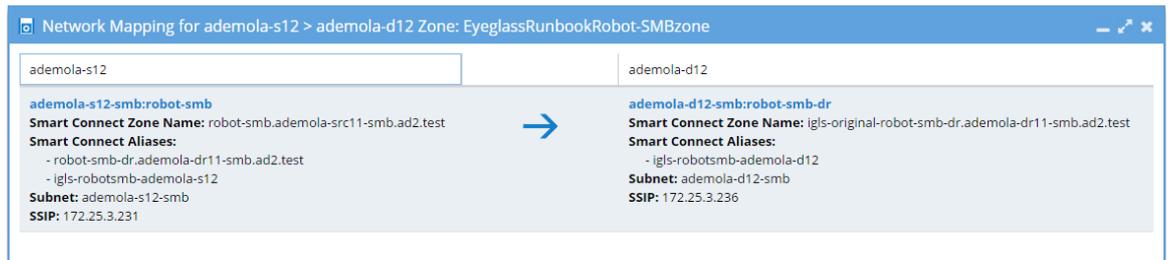
Pre Simulated Disaster - Cluster1 (prod) is available - Controlled Failover

1. Review all steps in the “Failover Planning Guide and checklist” before beginning. This is required to maintain support for this procedure. **See support statement above on planning guide requirement.**
2. Using Eyeglass, perform *controlled Failover* of your **Production Access Zone(s)** from Cluster1 to Cluster2.
3. On Eyeglass, enable each Production SyncIQ mirror-policy jobs for your Production Access Zone if they are in USERDISABLED state. Consult “Failover Planning Guide and checklist” to maintain support for this procedure.
4. Write data to share(s) protected by **Production** SyncIQ Policies from DFS mount (confirm that Cluster2 share path is the active target) on Windows Client after controlled failover of the Production Access Zone.
5. Do not proceed until above failover is validated as successful.
6. Do not fail over the EyeglassRunbookRobot-SMBzone Access Zone on Cluster1.

7. Procedure complete. Do not continue to next steps until successful **Controlled Failover** of your **Production Access Zone(s)** has completed successfully.

Simulated Disaster - Cluster 1 (prod) becomes unavailable: EyeglassRunbookRobot-SMBzone Access Zone

1. **Simulate Cluster 1 failure:** See below pool Cluster 1 to Cluster 2 IP pool mapping just pre-disaster:

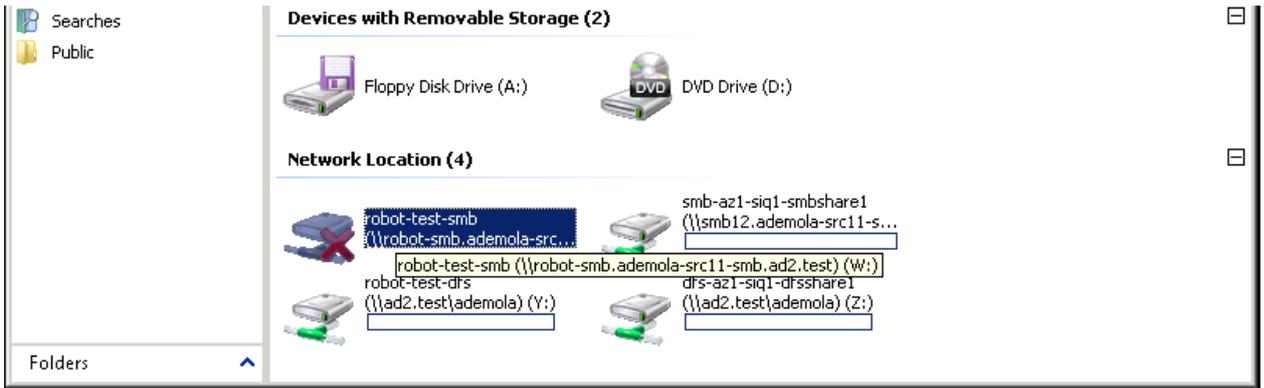


1. On Cluster1 OneFS UI, disconnect node interfaces from dedicated IP pool used for *client access* to test data on EyeglassRunbookRobot-SMBzone Access Zone (IP Pool assigned to the robot Access Zone). If SMB folder was properly set up, SMB folder target path from Cluster1 **will fail** when the node interfaces are removed from the pool. This is required to disconnect SMB session from clients to this pool and cause SMB mount failure.
2. Step 1a. above simulates DNS response failure to Cluster 1 as well without any IP's in the pool, without actually impacting SSIP or normal DNS operations. At this point in the process, name resolution is down, and NetBIOS sessions are disconnected from Cluster 1 EyeglassRunbookRobot-SMBzone Access Zone.

```
ademola-igls4:/home/admin # nslookup robot-smb.ademola-src11-smb.ad2.test
Server:      127.0.0.1
Address:     127.0.0.1#53

** server can't find robot-smb.ademola-src11-smb.ad2.test: SERVFAIL
```

Notice from the above screenshot, that name resolution to SmartConnect name is down as expected (SERVFAIL is returned). **At this point we have simulated a disaster** as Cluster1 EyeglassRunbookRobot-SMBzone Access Zone SmartConnect zone name resolution is failing, and no shares can be access on Cluster1 EyeglassRunbookRobot-SMBzone Access Zone.

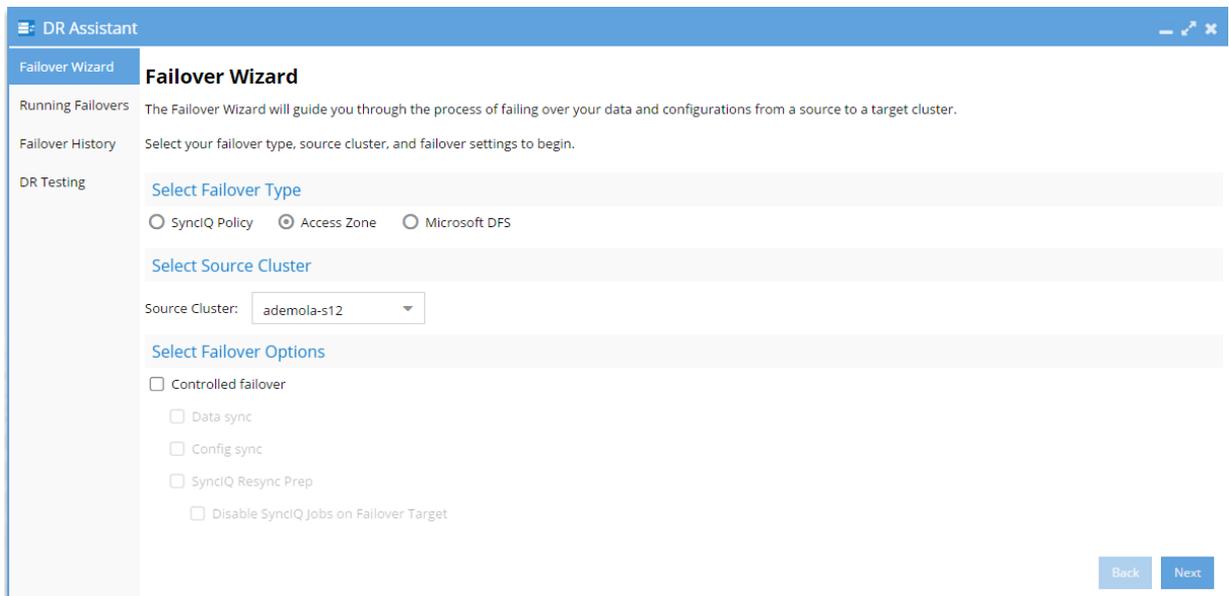


- Also, **set the schedule for the EyeglassRunbookRobot-SMB Test policy on Cluster 1 to manual**. As a policy won't be able to run anyway if the source cluster has been destroyed. **Do not proceed until this step is done.**

NOTE: in a real DR event, it is assumed the source cluster is unreachable on the network.

NOTE: Make note of the schedule, it will need to be reapplied at the end of this procedure.

- Perform Failover:** Using Eyeglass, perform *uncontrolled failover* for EyeglassRunbookRobot-SMBzone Access Zone from Cluster1 to Cluster2.



DR Assistant

Failover Wizard

Running Failovers

Failover History

DR Testing

Best Practices

- Run domain mark manually in advance to ensure a fast failover. See [this url](#) for details.
- Fallback operations are executed on the clusters (resync Prep), this runs 4 steps, two on the source and two on the target cluster that can take time to complete (see domain mark optimization). To increase performance and reduce the time taken to process the resync prep steps, it is recommended to increase the SyncIQ worker threads from the default to 10 or more.

Back Next

DR Assistant

Failover Wizard

Zone Selection

Access Zone Name ↑	Source Cluster	Target Cluster	Last Successful Readiness Check	DR Failover Status
<input checked="" type="checkbox"/> EyeglassRunbookRobot-SMBzone	ademola-s12	ademola-d12	4/6/2017, 4:49:13 PM	OK
<input type="checkbox"/> dfs-az1	ademola-s12	ademola-d12	4/6/2017, 4:49:13 PM	FAILED OVER
<input type="checkbox"/> nfs-az1	ademola-s12	ademola-d12	4/6/2017, 4:49:13 PM	FAILED OVER
<input type="checkbox"/> smb-az1	ademola-s12	ademola-d12	4/6/2017, 4:49:13 PM	FAILED OVER

Back Next

DR Assistant

Failover Wizard

Running Failovers

✓ SUCCESS

Failover History

Eyeglass has determined that your configuration is valid.

DR Testing

The following policies will be failed over: **ademola-s12_EyeglassRunbookRobot-SMB**
 The following policies will NOT be failed over: **No disabled policies found.**

Select checkbox to acknowledge that you have reviewed the [failover release notes](#). This is required to continue failover.

I have reviewed the failover release notes:

Back Next

DR Assistant

Failover Wizard

Summary

Running Failovers

Zone Selection

Access Zone Name ↑	Source Cluster	Target Cluster	Last Successful Readiness Check	DR Failover Status
<input checked="" type="checkbox"/> EyeglassRunbookRobot-SMBzone	ademola-s12	ademola-d12	4/6/2017, 4:49:13 PM	OK

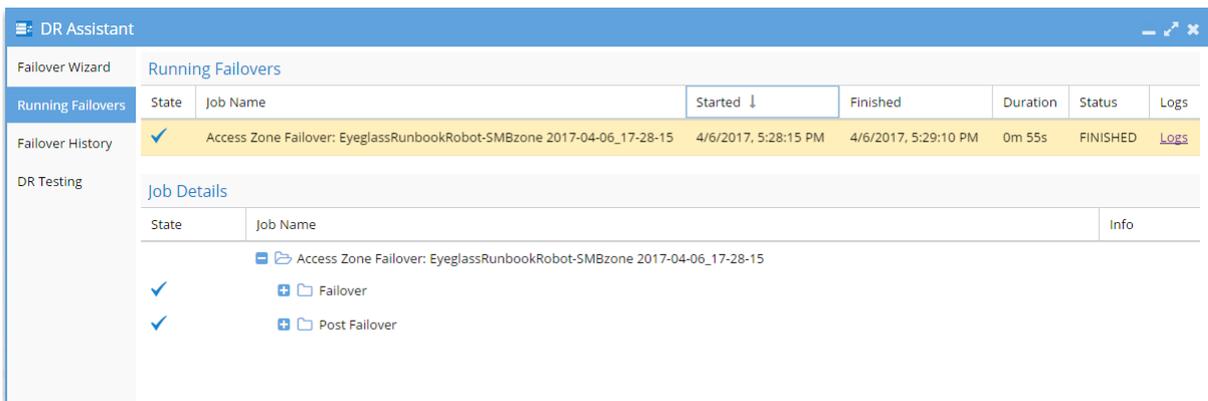
WARNING: Do not proceed past this point unless you have read the recovery guide that provides post failover log recovery steps, in the event a step fails to complete. The direct link is [here](#).

NOTE: Support will ask if you have executed recover steps if a failover log is submitted

Warning! This action can incur data loss on failover. Please ensure that you understand the implications of making the source cluster read only, and the target cluster read-write.

I have reviewed the above settings and wish to initiate failover:

Back Run Failover

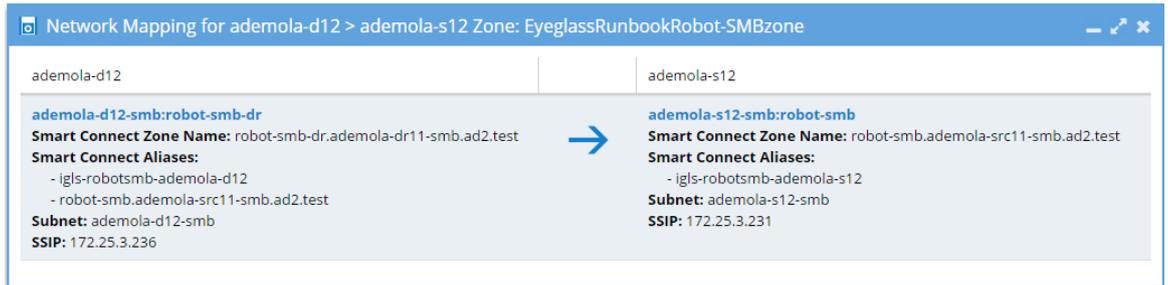


3. Wait until uncontrolled failover completes.
 1. Check SPN's are failed over in AD correctly using ADSI Edit.
 2. **Validation:** Test using nslookup to make sure DNS now resolves to Cluster 2.
 3. Correct or debug resolution of SmartConnect name before continuing.
 4. Test Client Access: **This step requires unmount and remount of the share to get new IP address.**
 1. Reboot the client machine that was used to validate the share pre-disaster to guarantee that the Netbios session to Cluster1 has not been preserved.
 2. Mount the share.
 3. Write data to share protected by EyeglassRunbookRobot-SMB Test SyncIQ Policy from SMB mount on Windows Client after the uncontrolled failover.
 5. Uncontrolled Access Failover complete.

Post Simulated Disaster - Cluster1 (prod) becomes available: EyeglassRunbookRobot-SMBzone Access Zone

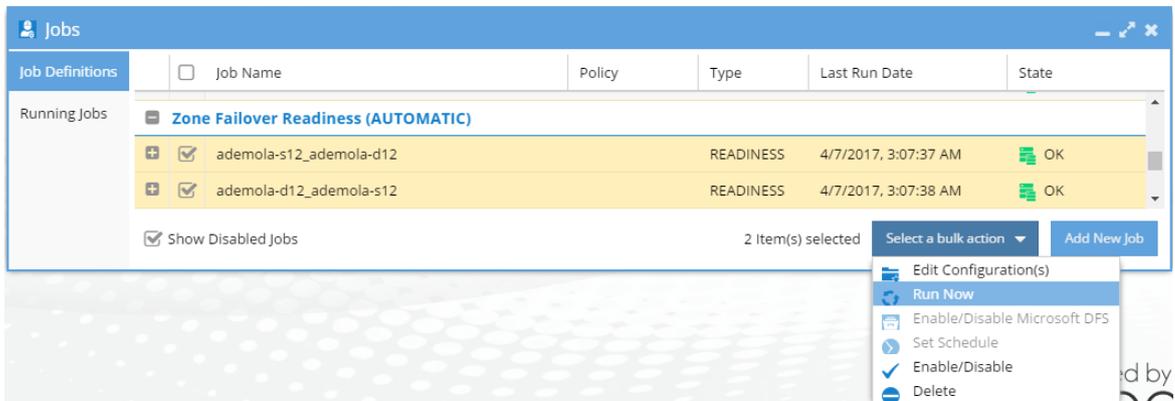
1. **Simulate Cluster 1 availability:** See below Cluster 2 to Cluster 1 EyeglassRunbookRobot-SMBzone Access Zone IP pool mapping just after Cluster1 is available. *Note that previously removed node interface have not been re-connected at this*

point:



1. One Cluster1 OneFS UI, edit source cluster EyeglassRunbookRobot-SMBzone Access Zone IP pool SmartConnect name and apply igls-original prefix to existing SmartConnect name. **This is required step before re-connecting the previously removed Cluster1 node interface.**
2. On Cluster1 OneFS UI, reconnect previously removed node interfaces back to IP pool used for client access to test data on EyeglassRunbookRobot-SMBzone Access Zone.
2. On Cluster1 OneFS UI, run resync-prep on EyeglassRunbookRobot-SMB Test SyncIQ policy. Consult EMC Documentation.
3. Verify that resync-prep process was completed without error before proceeding to next steps.
1. Resolve any errors before continuing. **Resync prep must have run successfully before you attempt to complete remaining steps. Check the cluster reports show no errors before continuing.**
4. Check Eyeglass job state
 1. From Eyeglass, verify EyeglassRunbookRobot-SMB Test Policy, and the mirror policy are in the correct state. **Mirror policy** should be **Enabled** and the **Cluster 1 policy** should be **Disabled** state.
 2. Allow Eyeglass Configuration Data Replication to run at least once.
 3. Note: Configuration Data Replication task must complete before continuing with steps below. Verify a config sync task has been run from running jobs window without errors.
 4. Verify Eyeglass jobs show policy state correctly with **Cluster 1 policy** showing policy **Disabled** and **Cluster 2 mirror policy** showing **Enabled** and OK (green).
 5. Wait for Config sync to correctly show the above state.
 6. Do not continue until the above validations are done.

- From the Eyeglass Jobs window, select Run Now from the Select a Bulk Action menu and run the Zone Failover Readiness jobs. This allows a new Access Zone Failover Readiness Audit to be computed.

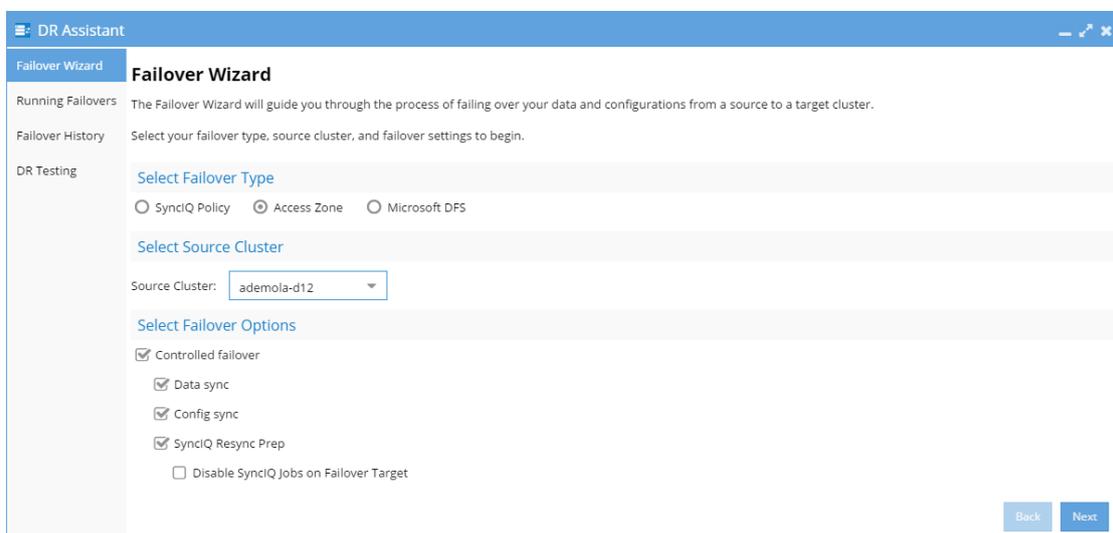


- On Eyeglass DR Dashboard, confirm that zone readiness looks good for the EyeglassRunbookRobot-SMBzone Access Zone.

The screenshot shows the 'DR Dashboard' window. It contains a table with columns: Policy Readiness, Zone Readiness, Source Cluster, Target Cluster, Zone Name, Last Successful Readiness Check, Network Mapping, and DR Failover Status. The table lists several readiness checks for different zones and clusters. The 'EyeglassRunbookRobot-SMBzone' is highlighted in red, indicating a 'FAILED OVER' status.

Policy Readiness	Zone Readiness	Source Cluster	Target Cluster	Zone Name	Last Successful Readiness Check	Network Mapping	DR Failover Status
DFS Readiness		ademola-s12	ademola-d12	EyeglassRunbookRobot-SMBzone	4/7/2017, 9:20:51 AM	View Map	FAILED OVER
DR Testing		ademola-d12	ademola-s12	EyeglassRunbookRobot-SMBzone	4/7/2017, 9:20:51 AM	View Map	OK
		ademola-s12	ademola-d12	System	4/7/2017, 9:20:51 AM	View Map	FAILED OVER
		ademola-d12	ademola-s12	System	4/7/2017, 9:20:51 AM	View Map	OK
		ademola-s12	ademola-d12	dfs-az1	4/7/2017, 9:20:51 AM	View Map	FAILED OVER
		ademola-d12	ademola-s12	dfs-az1	4/7/2017, 9:20:51 AM	View Map	OK
		ademola-s12	ademola-d12	nfs-az1	4/7/2017, 9:20:51 AM	View Map	FAILED OVER
		ademola-d12	ademola-s12	nfs-az1	4/7/2017, 9:20:51 AM	View Map	OK
		ademola-s12	ademola-d12	smb-az1	4/7/2017, 9:20:51 AM	View Map	FAILED OVER
		ademola-d12	ademola-s12	smb-az1	4/7/2017, 9:20:51 AM	View Map	OK

- Perform Controlled Failback:** Using Eyeglass, perform controlled failback of EyeglassRunbookRobot-SMBzone Access Zone from Cluster 2 to Cluster 1.



DR Assistant

Failover Wizard

Running Failovers

Failover History

DR Testing

Best Practices

- Run domain mark manually in advance to ensure a fast failover. See [this url](#) for details.
- Failback operations are executed on the clusters (resync Prep), this runs 4 steps, two on the source and two on the target cluster that can take time to complete (see domain mark optimization). To increase performance and reduce the time taken to process the resync prep steps, it is recommended to increase the SyncIQ worker threads from the default to 10 or more.

Back Next

DR Assistant

Failover Wizard

Zone Selection

Running Failovers

Failover History

DR Testing

Access Zone Name ↓	Source Cluster	Target Cluster	Last Successful Readiness Check	DR Failover Status
<input type="checkbox"/> smb-az1	ademola-d12	ademola-s12	4/7/2017, 9:20:51 AM	OK
<input type="checkbox"/> nfs-az1	ademola-d12	ademola-s12	4/7/2017, 9:20:51 AM	OK
<input type="checkbox"/> dfs-az1	ademola-d12	ademola-s12	4/7/2017, 9:20:51 AM	OK
<input type="checkbox"/> System	ademola-d12	ademola-s12	4/7/2017, 9:20:51 AM	OK
<input checked="" type="checkbox"/> EyeglassRunbookRobot-SMBzone	ademola-d12	ademola-s12	4/7/2017, 9:20:51 AM	OK

Back Next

DR Assistant

Failover Wizard

Running Failovers

Failover History

DR Testing

✓ SUCCESS

Eyeglass has determined that your configuration is valid.

The following policies will be failed over: **ademola-d12_EyeglassRunbookRobot-SMB_mirror**
 The following policies will NOT be failed over: **No disabled policies found.**

Select checkbox to acknowledge that you have reviewed the [failover release notes](#). This is required to continue failover.

I have reviewed the failover release notes.:

Back Next

DR Assistant

Failover Wizard

Summary

Running Failovers

Zone Selection

Failover History

DR Testing

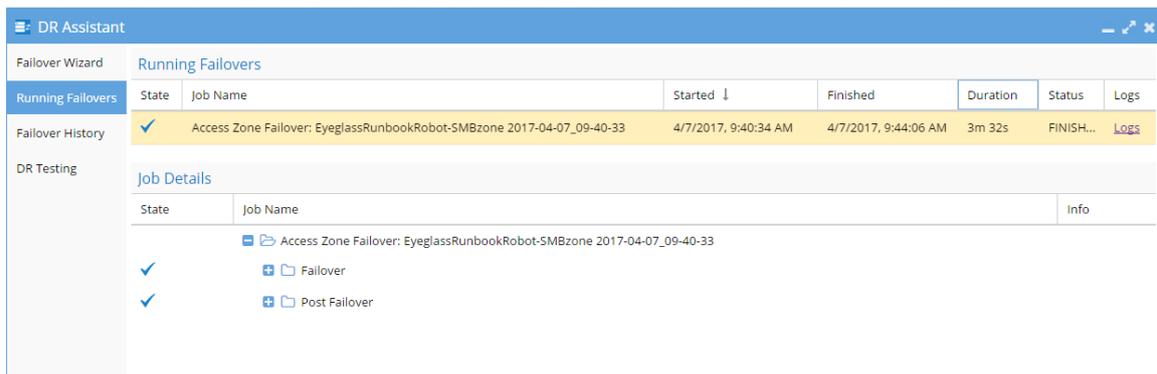
Access Zone Name ↓	Source Cluster	Target Cluster	Last Successful Readiness Check	DR Failover Status
<input checked="" type="checkbox"/> EyeglassRunbookRobot-SMBzone	ademola-d12	ademola-s12	4/7/2017, 9:20:51 AM	OK

WARNING: Do not proceed past this point unless you have read the recovery guide that provides post failover log recovery steps, in the event a step fails to complete. The direct link is [here](#).
 NOTE: Support will ask if you have executed recover steps if a failover log is submitted

Warning! This action can incur data loss on failover. Please ensure that you understand the implications of making the source cluster read only, and the target cluster read-write.

I have reviewed the above settings and wish to initiate failover:

Back Run Failover



8. Wait until controlled failover completes.
4. Check SPNs are failed over in AD correctly using ADSI Edit.
5. Validation: Test using nslookup to make sure DNS now resolves to Cluster 1.
6. Correct or debug resolution of SmartConnect name before continuing.
9. Test Client Access: This step requires unmount and remount of the share to get new IP address.
1. Reboot the client machine that was used to validate the share pre-disaster to guarantee that the Netbios session to Cluster2 has not been preserved.
2. Mount the share.
3. Write data to share protected by EyeglassRunbookRobot-SMB Test SyncIQ Policy from SMB mount on Windows Client after uncontrolled failover.
10. Controlled procedure complete.
11. If performing failback of Production data follow planning guide process to maintain support.

Copyright Superna LLC 2017

© Superna Inc

5. Appliance Operational Procedures

[Home](#) [Top](#)

- [Overview](#)
- [Eyeglass Common Procedures](#)
 - [Shutting down the appliance](#)
 - [Reboot the appliance](#)
 - [IP address change](#)
 - [vmotion to new ESX host](#)
 - [Password Change](#)
- [ECA cluster Common Operations \(Ransomware Defender, Easy Auditor, Performance Auditor, Search & Recover and Golden Copy\)](#)
 - [Health Monitor - Audit Data Ingestion and Audit Data Save to Database](#)
 - [Restart the cluster](#)
 - [Reboot a single node](#)
 - [ECA cluster ip address change](#)
 - [ECA password change](#)
 - [ECA cluster shutdown](#)
 - [Power Loss to VM or Recovery from a Reboot without graceful shutdown](#)

Overview

This guide covers common operational questions and procedures for Eyeglass and ECA clusters.

Eyeglass Common Procedures

1. Shutting down the appliance

- a. Login as admin using ssh
- b. sudo -s (enter admin password)
- c. systemctl stop sca
- d. shutdown

1. Reboot the appliance

- a. NOTE: This should not be used as a trouble shooting step. This is not required for any operational steps. If a reboot is required for an operating system patch follow these steps. For all other cases open a case with support and do not reboot the appliance unless directed by support
- b. login as admin
- c. sudo -s (enter admin password)
- d. systemctl stop sca
- e. reboot

2. IP address change

- a. login as admin
- b. sudo -s (enter admin password)
- c. yast
- d. use arrow keys to select networking and edit interface ip settings, including, DNS and default gateway

3. vmotion to new ESX host

- a. Eyeglass can be moved between hosts without shutdown
- b. PowerScale code upgrade [see guide](#)

4. Password Change

- a. Login as the admin user over ssh (also applies to other builtin user accounts for other products, rwdefend, auditor)

- b. type the command below:
 - i. passwd
- c. Enter new password, retype new password
- d. done

ECA cluster Common Operations (Ransomware Defender, Easy Auditor, Performance Auditor, Search & Recover and Golden Copy)

All of the above products share common operating procedures.

1. For additional scenarios on ECA based product admin guide [see the guide](#).

2. Health Monitor - Audit Data Ingestion and Audit Data Save to Database

- a. **NOTE: These graphs should be checked every time you have errors or issues with Ransomware Defender, Easy Auditor or Performance Auditor. Each product depends on audit data ingestion. The graph should never be a flat line for received or sent audit records.**
- b. **NOTE: The actual values are only intended to show general ingestion rates and are not intended for any other purpose. The graphs can be used to monitor if ingestion is balanced across each ECA VM's. The graphs and relative rates of ingestion are for sizing the number ECA nodes 6, 9 or more following documented sizing guide lines and to monitor ingestion rates across times of day.**
- c.
 - i. **The most common cause is network issues between the ECA VM' s and the cluster.**
- d. Login to Eyeglass GUI (Requires 2.5.8 or later)
- e. Click the Managed Services Icon

- f. Click the ECA monitor button. This will launch a new browser tab to display received audit data per ECA node rate of events per minute per node and sent audit data per ECA node to the Isilon / Powerscale HDFS database.
- g. **NOTE: The graphs should never be flat lines at the 0 events per second graph for any extended period of time. The graph shows events per minute per ECA node. 0 events per seconds is not a normal state and some positive number of events should be constantly seen.**



- h.
- i. Turbo Audit received is audit data ingestion
- j. evtarchive sent is audit data saving to HDFS (Easy Auditor Only)

3. Restart the cluster

- a. login to node 1 as ecaadmin over ssh
- b. ecactl cluster down
- c. then
- d. ecactl cluster up
- e. NOTE: can take 5-7 minutes to startup and shutdown all nodes in the cluster

4. Reboot a single node

- a. login to the node over ssh as ecaadmin
- b. sudo -s (enter admin password)
- c. reboot

- d. Now login to node 1 as ecaadmin
- e. `ecactl cluster up` (this will ensure all services are started on all nodes, even if the cluster is already running)

5. ECA cluster ip address change

- a. Review the eca admin [guide](#)

6. ECA password change

- a. Login as the ecaadmin user over ssh
- b. type the command below:
- c. `passwd`
- d. Enter new password, retype new password
- e. done

7. ECA cluster shutdown

- a. Use this procedure to stop the cluster software
- b. login to eca node 1 as ecaadmin
- c. type
 - i. `ecactl cluster down`

8. Power Loss to VM or Recovery from a Reboot without graceful shutdown

- a. login to node 1 as ecaadmin over ssh
- b. `ecactl cluster down`
- c. then
- d. `ecactl cluster up`

© Superna Inc

6. All Products Hardening Guide

[Home](#) [Top](#)

- [How to use this Guide](#)
- [Securing Eyeglass, ECA, Search & Recover and Golden Copy by Applying OS Patches](#)
 - [How to Subscribe Eyeglass OS Security updates](#)
 - [How to Use Superna Hosted Open Suse Repository for OS patching automation](#)
 - [Steps to Configure the Superna Open Suse Repository](#)
- [General Purpose OS Advanced Hardening \(All products\)](#)
- [How to add a Signed Certificate to the WebUI's \(Eyeglass, ECA Cluster, Golden Copy, Search & Recover\)](#)
- [How to Extend the HTTPS GUI Key Length to 4096](#)
- [Web Server HTTP Hardening Directives for Eyeglass and Search & Recover and Golden Copy](#)
 - [Eyeglass WebUI](#)
- [HTTPS Security Algorithm Hardening Eyeglass DR < 2.5.7](#)
- [HTTPS Security Algorithm Hardening Eyeglass DR 2.5.7 >](#)
- [HTTPs Web Server Hardening \(ECA, Search & Recover and Golden Copy\)](#)
- [How to turn off bash history \(Eyeglass, Golden Copy, ECA, Search & Recover\)](#)
- [Hardening Password Complexity \(All products\)](#)
- [Banning local user accounts after repeated failed login attempts](#)
 - [Configuration Steps for Eyeglass](#)
 - [Configuration Steps for Golden Copy & Search & Recover](#)
 - [How to edit the configuration defaults](#)
 - [Default configuration:](#)
- [Eyeglass WebUI Security API Auditing](#)
 - [How to monitor user UI actions and Authentication Login](#)
 - [How to view the GUI API Log](#)

- [ECA VM Hardened Virtual Secured Network \(Ransomware Defender, Easy Auditor, Performance Auditor\)](#)
- [2 Factor SSH Authentication for Eyeglass, Golden copy, Search & Recover or ECA VM's](#)
- [How to Mediate python vulnerability CVE-2022-35861](#)

How to use this Guide

This guide provides additional security hardening steps that are optionally applied to one or more products as indicated below. Not all sections apply to each product. OS Customizations are provided as is without support under the support contract. All OS customizations are not backed up and will not be migrated to a new appliance.

Securing Eyeglass, ECA, Search & Recover and Golden Copy by Applying OS Patches

1. Before scanning the appliance with security tools the following steps must be taken:
 - a. Upgrade to the latest OVA operating system using backup and restore to get web server configured with default hardening. [Upgrade guide](#). Follow the backup and restore steps. **NOTE: Os patching is not covered by the support contract and is customer responsibility.**
 - b. Patch the operating system
 - i. **NOTE: Requires internet access to the appliance to reach OS internet repositories**
 - ii. **NOTE: Always take a vmware level snapshot before starting.**
 - iii. **Mandatory steps - Backup sudoer file and set a root password**
 1. Login as admin or ecaadmin
 2. Change root user password (it is random password on all appliances), this is needed in case recovery steps are needed.

- a. `sudo -s`
- b. `passwd` (save the password securely)
3. `cp /etc/sudoers /etc/sudoers.bak`
4. Repeat the above steps on each vm (ECA, Golden Copy, Search & Recover)
- iv. Follow this guide to configure a proxy to reach the Internet repositories. [Proxy Guide Eyeglass](#)
- v. login to eyeglass as admin user or as ecaadmin (ECA, Golden Copy, Search & Recover)
- vi. `sudo -s` (enter admin password)
- vii. `zypper refresh` (updates repositories)
- viii. `zypper update` (applies patches)
- ix. Review any messages that indicate a reboot is required to have the update take effect.
- x. DR Eyeglass VM - Before reboot
 1. `systemctl stop sca`
 2. `reboot`
- xi. ECA, Golden Copy, Search & Recover
 1. `ecactl cluster down`
 2. `sudo -s`
 3. `reboot`
 4. After reboot
 5. Login again as ecaadmin
 6. `ecactl cluster up`
 7. verify normal application health and review cluster up log for errors
- c. Use Eyeglass service account and review all information to make sure permissions are up to date
 - i. Reference: [PowerScale Cluster User Minimum Privileges for Eyeglass](#)

How to Subscribe Eyeglass OS Security updates

1. Note: The appliance defaults to weekly automatic critical patches and security updates if Internet connection is allowed to the appliance. Security patches are customer responsibility to apply and manage and not covered under the support contract.
2. If customers would like email notification of OS updates follow this link and register. <http://lists.suse.com/mailman/listinfo/sle-security-updates>

How to Use Superna Hosted Open Suse Repository for OS patching automation

1. The repository is hosted on a static IP address the url and ip address should be authorized for outbound only https from Eyeglass --> url , IP below. **NOTE: The default OS will use default open Suse mirror repositories. There is NO need to use Superna repo. The only reason to use this repo is for controlled firewall and URL to the report versus the dynamic mirror used on public repositories. To repeat this statement. if you don't know why you need this repo. open a support case and ask.**
 - a. The url is <https://opensuserepo.supernaeyeglass.com>
 - b. The static ip address is 34.110.179.6
 - c. Requires outbound HTTPS on 443 to reach the Superna Repository
2. Follow Steps below to change the default repository for patching and switch to https.

Steps to Configure the Superna Open Suse Repository

1. **VMware level snapshot before you start. Do NOT start without a snapshot.**
2. MANDATORY: Download the 15.4 eyeglass or ECA upgrade file.
3. Login to eyeglass as admin
4. `sudo -s`

5. Disable repositories

- a. `zypper lr -E | tail -n +5 | awk '{print $1}' | xargs zypper mr -d`

6. Add Superna Repository

- a. `zypper ar -n "Superna Update repository of OpenSUSE Backports"
https://opensuserepo.supernaeyeglass.com/updates/leap/15.4/backports/ "superna-repo-backports-update"`
- b. `zypper ar -n "Superna Non-OSS Repository"
https://opensuserepo.supernaeyeglass.com/distribution/leap/15.4/repo/non-oss/ "superna-repo-non-oss"`
- c. `zypper ar -n "superna-repo-oss" https://opensuserepo.supernaeyeglass.com/distribution/leap/15.4/repo/oss superna-repo-oss`
- d. `zypper ar -n "Superna update repository with updates from SUSE Linux Enterprise 15"
"https://opensuserepo.supernaeyeglass.com/updates/leap/15.4/sle" superna-repo-sle-update`
- e. `zypper ar -n "Superna Main Update Repository" https://opensuserepo.supernaeyeglass.com/updates/leap/15.4/oss
superna-repo-update`
- f. `zypper ar -n "Superna Update Repository (Non-Oss)"
https://opensuserepo.supernaeyeglass.com/updates/leap/15.4/non-oss/ superna-repo-update-non-oss`

7. `zypper refresh` (updates and tests access to the repository)

8. `zypper dup` (upgrade OS from 15.3 to 15.4) **Long step that can take 20 minutes or more depending on Internet bandwidth**

9. Reboot

10. SCP the 15.4 upgrade file to the appliance

11. `chmod 777 <eyelgass upgrade file>`

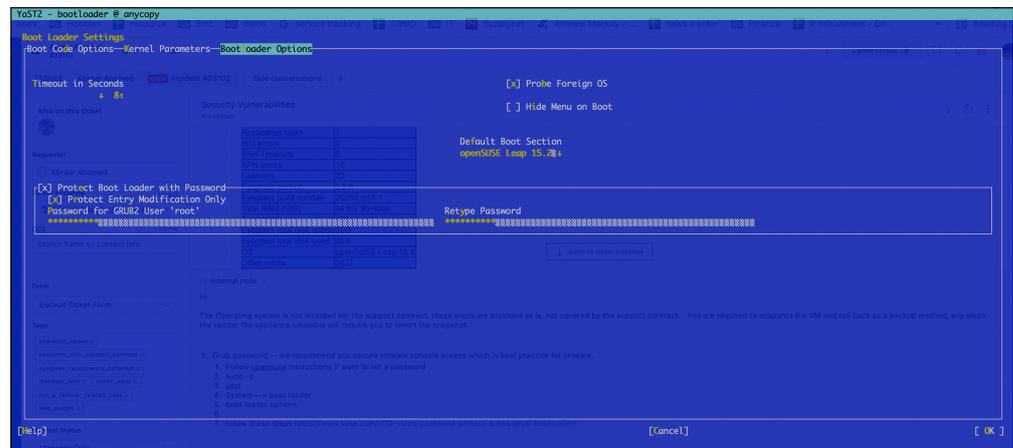
12. `./<eyeglass upgrade file>`

13. done

General Purpose OS Advanced Hardening (All products)

1. **Grub password** -- It is recommended to secure VM console access.
 - a. Follow opensuse instructions if want to set a password
 - b. `sudo -s`
 - c. `yast`
 - d. System --> boot loader --> boot loader options tab (use arrow keys to select)
 - e. enter a password

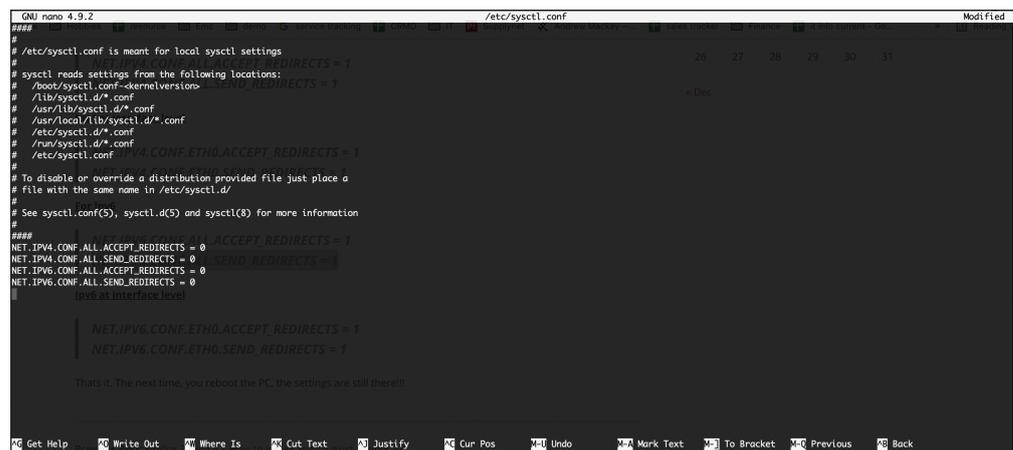
- f. NOTE: This will require password to make any changes to the boot loader it will not require a password to boot the OS.
- g. To enable boot password uncheck the option below "Protect Entry Modifications Only"



h.

2. Disable ICMP Redirects

- a. `sudo -s`
- b. `nano /etc/sysctl.conf`
- c. add these entries to the file and save the file, then control+x to save. Then reboot the OS with `reboot` command.
- d. `NET.IPV4.CONF.ALL.ACCEPT_REDIRECTS = 0`
- e. `NET.IPV4.CONF.ALL.SEND_REDIRECTS = 0`
- f. `NET.IPV6.CONF.ALL.ACCEPT_REDIRECTS = 0`
- g. `NET.IPV6.CONF.ALL.SEND_REDIRECTS = 0`



h.

3. User home directory Hardening

- a. `sudo -s`
- b. `cd /home`
- c. `chmod 750 admin ecaadmin screenshots`

How to add a Signed Certificate to the WebUI's (Eyeglass, ECA Cluster, Golden Copy, Search & Recover)

1. Eyeglass VM Steps
 - a. Follow the TLS cert steps in the admin guide [here](#).
2. For ECA, Golden Copy, Search & Recover follow these steps.
 - a. Access the WebUI from node 1 and create a DNS entry for node to create a FQDN to create a signed cert. The objective is to install the signed cert for nginx ECA Node-1
 - b. Create A record in DNS name for ECA Node-1 and verify with nslookup. Example eca1.domain.com
 - c. SSH to ECA Node-1 as ecaadmin
 - i. `cd /opt/superna/eca/conf/nginx`
 - ii. Verify that the nginx.key is there with `ls -la`
 - d. Create csr with that key file
 - i. Command: `openssl req -key nginx.key -new -out nginx.csr`
 - ii. SCP the nginx.csr file for signing
 - iii. Or type `cat nginx.csr` and copy and paste the text to submit for signing.
 - iv. When it is asked about the Common Name: provide the fqdn of ECA Node-1 (the name registered in DNS e.g. search.domain.com)
 - e. With that CSR certificate submit the request to Certificate Authority at your enterprise

- f. NOTE: These steps are CA specific, consult with your security team
- g. Once received the signed certificate encoded in PEM format
- h. scp (use WinSCP for Windows) and copy this file to ECA-1 under /opt/superna/eca/conf/nginx with name nginx.crt
- i. **NOTE: if not in PEM format, convert to PEM format or ask your Security team for pem format**
- j. Replace existing nginx.crt certificate with this new signed CA certificate.
- k. mv nginx.crt nginx.crt.bak (backup old file)
- l. cp /pathtonewfile/nginx.crt to /opt/superna/eca/conf/nginx/nginx.crt
- m. Restart nginx
- n. Bring down and up the ECA cluster to push the config to all the other ECA nodes
 - i. ecctl cluster down
 - ii. ecctl cluster up
- o. Verify the certificate when accessing the UI (e.g. https://FQDN)

How to Extend the HTTPS GUI Key Length to 4096

1. Reference: https://raymii.org/s/tutorials/Strong_SSL_Security_On_lighttpd.html (section Forward Secrecy & Diffie Hellman Ephemeral Parameters)
2. Login to Eyeglass
3. sudo -s
4. Ran: sudo openssl dhparam -out /etc/ssl/certs/dhparam.pem 4096

5. Took at least 10 -20 minutes to complete
6. Run: `sudo nano /etc/lighttpd/lighttpd.conf`
7. Find the definition for port 443 and add the following below `ssl.cipher-list` entry:
8. `ssl.dh-file = "/etc/ssl/certs/dhparam.pem"`
9. `ssl.ec-curve = "secp384r1"`
10. Restart the service: `sudo systemctl restart lighttpd.service`
11. Eyeglass GUI and ECA still posts health to Eyeglass without issue.

Web Server HTTP Hardening Directives for Eyeglass and Search & Recover and Golden Copy

1. This section has specific web server directives that address specific hardening http header responses and setting the Search & Recover TLS protocol requirements. **NOTE: Not required for Eyeglass 2.5.7 or later as these are set by default.**

2. Eyeglass WebUI

- a. For Eyeglass lighttpd https and http HEADER fix
- b. login as admin
- c. `sudo -s` (enter admin password)
- d. `nano /etc/lighttpd/lighttpd.conf`
- e. Add the following inside SERVER 443 block
- f. `control+w` and type `":443"` [enter key]
- g. Add the text below and replace the previous section for this between the () for this section `setenv.add-response-header`

```
setenv.add-response-header = (
    "Strict-Transport-Security" => "max-age=15768000",
```

```

"Content-Security-Policy" => "frame-ancestors 'self';",
"X-Content-Type-Options" => "nosniff",
"X-Frame-Options" => "DENY",
"X-XSS-Protection" => "1; mode=block"

```

)

1. See example

```

ssl.disable-client-renegotiation = "enable"

$SERVER["socket"] == ":443" {
    ssl.engine = "enable"
    ssl.pemfile = "/opt/superna/sca/.secure/ssl.pem"
    ssl.use-ssl2 = "disable"
    ssl.use-ssl3 = "disable"
    setenv.add-response-header = (
        "Content-Security-Policy" => "frame-ancestors 'self';",
        "X-XSS-Protection" => "1; mode=block",
        "Strict-Transport-Security" => "max-age=15768000",
        "X-Content-Type-Options" => "nosniff",
        "X-Frame-Options" => "DENY"
    )
    # redirect the root to /eyeglass
    $HTTP["host"] =~ "(.*)" {
        url.redirect = ("^/$" => "/eyeglass")
    }
}

```

a.

2. Now locate the section for http (only used to redirect to 443 port)

a. control+w and type **":80"** [enter key]

b. Add the text below and replace the previous section for this between the () **setenv.add-response-header**

c. setenv.add-response-header = (

```

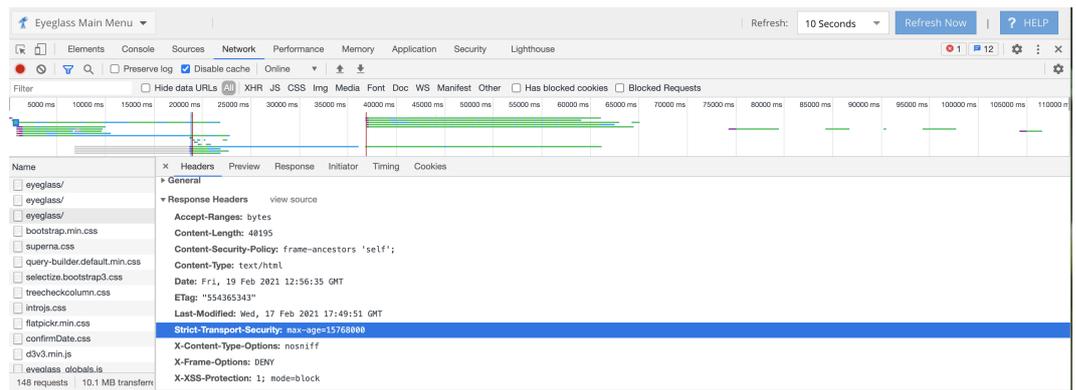
"Strict-Transport-Security" => "max-age=15768000",
"Content-Security-Policy" => "frame-ancestors 'self';",
"X-Content-Type-Options" => "nosniff",
"X-Frame-Options" => "DENY",
"X-XSS-Protection" => "1; mode=block"

```

)

```
## Redirect HTTP to HTTPS
$SERVER["socket"] == ":80" {
    setenv.add-response-header = (
        "Strict-Transport-Security" => "max-age=15768000",
        "Content-Security-Policy" => "frame-ancestors 'self';",
        "X-Content-Type-Options" => "nosniff",
        "X-Frame-Options" => "DENY",
        "X-XSS-Protection" => "1; mode=block"
    )
}
```

1.)
2. save the file with **control+x** answer yes to save and exit
3. Restart the web server
 - a. `systemctl restart lighttpd.service`
4. Verify with Google Chrome Developer tools (Press F12).
 - a. Login to eyeglass, select Network tab, select eyeglass web page on left side, click headers tab , expand response headers, verify "content-security-policy" and "strict transport security" Use the screen shot as per below.



HTTPS Security Algorithm Hardening Eyeglass DR < 2.5.7

1. Update Java HTTPS algorithms and certificate settings
2. `nano /opt/superna/java/jre/lib/security/java.security`
3. press control+w
4. type `jdk.tls.disabledAlgorithms` and the press enter
5. remove the # comment from this line see image below

```
# Example:
jdk.tls.disabledAlgorithms=MD5, SHA1, DSA, RSA keySize < 2048, SSLv2Hello, SSLv3, TLSv1, TLSv1.1
# Legacy algorithms for Secure Socket Layer/Transport Layer Security (SSL/TLS)
# processing in JSSE implementation
```

a.

6. Then press control+W
7. type `jdk.certpath.disabledAlgorithms` [enter]
8. repeat control+w [enter] 3 times until you locate the line with the # comment on the lines below
9. remove the # from **both** lines , refer to the image below.

```
jdk.certpath.disabledAlgorithms=MD2, MD5, SHA1 jdkCA & usage TLS, RSA keySize < 1024, DSA keySize < 1024, EC keySize < 224
```

a.

- 10.press control+x
11. answer yes to save the file and exit
12. For changes to take effect restart the sca
13. **systemctl restart sca**
14. done

HTTPS Security Algorithm Hardening Eyeglass

DR 2.5.7 >

1. ssh as admin to eyeglass
2. backup existing file
 - a. `mv /opt/superna/java/jre/lib/security/java.security /opt/superna/java/jre/lib/security/java.security.bak`
3. Copy enhanced security file
 - a. `cp /opt/superna/java/java.security.enhanced /opt/superna/java/jre/lib/security/java.security`
 - b. `sudo chown sca:users java.security`
4. `systemctl restart sca`

5. done

Overview: Port 80 is only used to redirect to port 443, it is not used for anything else. To block port 80 follow these steps

1. Login to Eyeglass appliance as admin user. Elevate to root using command:
sudo su -
2. Create Eyeglass port 80 firewall script:
 - a. nano /opt/superna/bin/firewall-rules.sh
3. Type 'i' to enter insert mode. Copy and paste the following to the file:
 - a. #!/bin/bash
 - b. iptables -I IN_public_deny -p tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable
4. Type ESC and :wq! to exit when file contents match above.
5. Change ownership and modify the file:
 - a. chown sca:users /opt/superna/bin/firewall-rules.sh ; if ["\$?" == 0]; then echo Success; fi
 - b. chmod u+x /opt/superna/bin/firewall-rules.sh ; if ["\$?" == 0]; then echo Success; fi
6. Create a service file to be run at boot after Network service is registered as RUNNING:
 - a. nano /etc/systemd/system/boot-firewall-rules.service
 - b. [Unit]
 - c. After=network.target
 - d. [Service]
 - e. ExecStart=/opt/superna/bin/firewall-rules.sh
 - f. [Install]
 - g. WantedBy=default.target
7. Type ESC and :wq! to exit when file contents match above.
8. Run the following commands for the changes to take effect (note: NO Reboot required. No impact):

9. `systemctl daemon-reload ; if ["$?" == 0]; then echo Success; fi`
10. `systemctl enable boot-firewall-rules.service`
11. `systemctl start boot-firewall-rules.service ; if ["$?" == 0]; then echo Success; fi`
12. `systemctl status boot-firewall-rules.service`
13. done

HTTPs Web Server Hardening (ECA, Search & Recover and Golden Copy)

1. Requirements:
 - a. < 2.5.7 ECA
 - b. 1.1.2 Search and Recover
 - c. 1.1.4 Golden Copy
 - d. NOTE: 2.5.7.1 and later have these changes already., Search 1.1.5 or later and Golden copy 1.1.6 or later.
2. Login to **ECA** or **Search & Recover** over ssh as ecaadmin
 - a. `nano /opt/superna/eca/conf/nginx/eca.conf.template`
3. Login to Golden Copy
 - a. `nano /opt/superna/eca/conf/nginx/eca.conf.simpletemplate`
4. Add the following inside server 443 block
 - a. `add_header Strict-Transport-Security "max-age=15768000" always;`
 - `add_header Content-Security-Policy "frame-ancestors 'self';" always;`
 - `add_header X-Content-Type-Options "nosniff" always;`
 - `add_header X-Frame-Options "DENY" always;`
 - `add_header X-XSS-Protection "1; mode=block" always;`
 - `ssl_protocols TLSv1.2 TLSv1.3 always;`

```
server {
    listen 443 ssl http2 default_server;
    listen [::]:443 ssl http2 default_server;

    server_name _;
    root /usr/share/nginx/html;

    add_header Strict-Transport-Security "max-age=15768000";
    add_header Content-Security-Policy "frame-ancestors 'self'";
    add_header X-Content-Type-Options "nosniff";
    add_header X-Frame-Options "DENY";
    add_header X-XSS-Protection "1; mode=block";

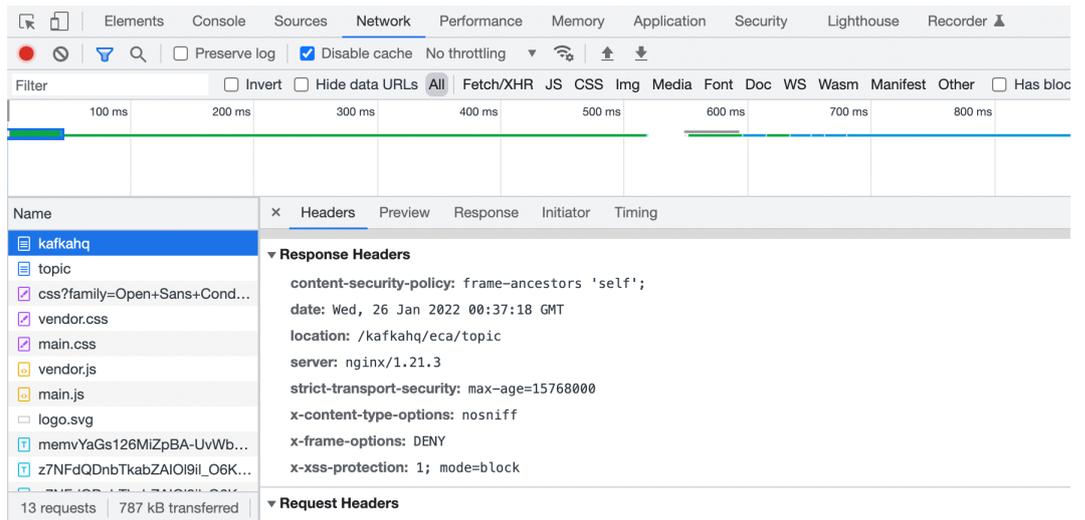
    ssl_certificate /etc/nginx/conf.d/nginx.crt;
    ssl_certificate_key /etc/nginx/conf.d/nginx.key;

    ssl_protocols TLSv1.2 TLSv1.3;

    location / {
        proxy_read_timeout 1800s;
        proxy_pass http://searchui:3000;
    }
}
```

b.

5. Push the config to all nodes
 - a. `ecactl cluster push-config`
6. Restart containers to read the new configuration
 - a. `ecactl cluster exec "ecactl cluster services restart --container nginx --all"`
7. Verify with Google Chrome developer tools
 - a. Login to the ECA vm's, open the managed services icon, Select node 1. Select kafkahq management url
 - b. login with default user id and password ecaadmin and 3y3gl4ss
 - c. Open developer tools in Chrome browser (press F12), Select the network tab, Select kafkahq on the left menu, change view to Headers. Verify the change using screen shot below.



d.

8. Done

How to turn off bash history (Eyeglass, Golden Copy, ECA, Search & Recover)

1. Bash history can contain access key commands. Disabling bash history disables command history.
2. Login as ecaadmin (Golden copy) or admin (eyeglass)
3. `history -c` (cleans current history)
4. `echo 'set +o history' >> ~/.bashrc`
5. `logout`
6. `done`

Hardening Password Complexity (All products)

Follow these steps to enable local password complexity of the builtin users admin, auditor and rwddefend. NOTE: These settings only apply to the local OS users, if using RBAC proxy login to PowerScale or AD use the password features of the PowerScale or AD to setup password complexity.

To set these password rules the - (minus number) means MUST have in the password. Use the definitions below to customize the example provided.

- Minimum password length should be x characters
 - value minlen
 - Password should have one UPPERCASE Character
 - value ucredit
 - Password should have one LOWERCASE Character
 - value lcredit
 - Password should have one Numeric Character
 - value dcredit
 - Password should have Special characters
 - value ocredit
 - Minimum Passwords to Remember or Password History
 - value pwhistory-remember
 - Accounts should be lockout after bad login attempts, see next section that blocks the source ip of the machine after failed local logins using fail 2 ban and firewall rules.
1. **Verify pam modules are installed (may not be required on all appliances depending on OS version, it may return no module found on 15.1 or later OS version which can be ignored and continue the steps)**
 2. login as admin
 3. sudo -s
 4. enter admin password
 5. zypper install pam-modules (this requires internet access to install additional pam modules)
 6. Answer yes to install new modules
 7. cd /etc/pam.d/
 8. cp common-password common-password.bak (backup old password file rules)
 9. **pam-config -a --cracklib --cracklib-minlen=6 --cracklib-lcredit=-1 --cracklib-ucredit=-1 --cracklib-dcredit=-1 --cracklib-ocredit=-1 --pwhistory --pwhistory-use_authtok --pwhistory-remember=3**
 - a. See definitions above for each value to customize

- b. This will generate a new common-password file
- c. When users try to change passwords they will require a password that matches these rules. **NOTE the root user can set a password for a user account that does not match these rules.**

Banning local user accounts after repeated failed login attempts

The appliance has several local users admin, auditor, and rwdefend used for builtin roles for different products. **NOTE: the root user password is randomized and sudo access to root should be used and leave the password randomized.**

To ban users that attempt brute force login attempts the following appliance enhancement allows control of lockouts and timed locked outs. This will setup firewall rules to block the ip of the user. The blocked login will cover ssh access and https to the WebUI. **NOTE: If proxy login is used to AD or PowerScale local users, using the RBAC features, these users will also be banned as well.**

1. Login as admin
2. sudo -s
3. enter admin password
4. zypper install fail2ban (requires Internet access to the appliance)
5. systemctl start fail2ban

Configuration Steps for Eyeglass

Highlevel:

- modified /etc/fail2ban/jail.conf [added 'eyeglass' section]

- enabled eyeglass filtering from /etc/fail2ban/jail.local
 - added 'eyeglass' custom filter file in /etc/fail2ban/filter.d/ directory
1. nano /etc/fail2ban/filter.d/eyeglass.conf (add the contents below to the file and save the file with :wq)

a. # Fail2ban filter for Superna Eyeglass

#

#

[INCLUDES]

before = common.conf

[Definition]

failregex = <HOST> \b\d{1,3}\.\d{1,3}\.\d{1,3}\.\d{1,3}\b - \[.* "POST
/RestClient/login/login HTTP/1.1" 500

datepattern = %%d/%%b/%%Y:%%H:%%M:%%S

ignoreregex =

2. nano /etc/fail2ban/jail.local

a. add The following to this file

i. [DEFAULT]

ignoreip = 127.0.0.1/8

bantime = 300

findtime = 300

maxretry = 3

[sshd]

enabled = true

[eyeglass]

enabled = true

- b. **Modify /etc/fail2ban/filter.d/sshd.conf file**
 - i. **sed -e /'spam_unix/s/^/#/g' -i /etc/fail2ban/filter.d/sshd.conf**
 - c. **Modify /etc/fail2ban/jail.conf to add eyeglass jail rule**
 - i. **sed -i "/HTTP servers/a[eyeglass]\n\nport = http,https\nlogpath = /var/log/lighttpd/access.log" /etc/fail2ban/jail.conf**
 - d. **restart the service**
 - i. **systemctl restart fail2ban**
 - ii. **check status**
 - iii. **systemctl status fail2ban**
 - e. **Optional - Find bantime and change default from 300 seconds to a value that meets your requirements**
 - f. **Optional - Find findtime and change default from 600 to a value that meets your requirements (A host is banned if it has generated "maxretry" during the last "findtime")**
 - g. **Optional - Find maxretry and change default from 3 to a value that meets your requirements**
3. Save the file after changes control+x answer yes to save
 4. done.

Configuration Steps for Golden Copy & Search & Recover

There is a fail2ban folder under /opt/superna/eca/conf/fail2ban in builds 1.1.4 > 20300. There is also a default jail.local, the following is its default content below. After modifying any conf under fail2ban folder the fail2ban container must be restarted to activate the new configuration.

How to edit the configuration defaults

1. ssh to the node as ecaadmin

2. nano /opt/superna/eca/conf/fail2ban/jail.local
3. edit the bold settings shown below to adjust ban time and retries for webui and ssh
4. control + x to save and exit
5. eectl container restart fail2ban (for changes to take effect)

Default configuration:

1. banned logins will be 300 seconds or 5 minutes
2. 5 retries of the password will be allowed before banning from the webui login
3. 5 retries of the ssh password will be allowed before banning.
4. nano /opt/superna/eca/conf/fail2ban/jail.local

[nginx-http-auth]

```
enabled = true
filter = nginx-http-auth
port = http,https
logpath = /var/log/nginx/error.log
maxretry = 5
bantime = 300
chain = DOCKER-USER
```

[sshd]

```
enabled = true
backend = systemd
filter = sshd
maxretry = 5
bantime = 300
```

[searchmw]

```
enabled = true
filter = searchmw
```

port = http,https
logpath = /var/log/searchmw/loginError.log
maxretry = 5
bantime = 300
chain = DOCKER-USER

Eyeglass WebUI Security API Auditing

This feature is available in 2.5.7 or later releases. This allows a user GUI audit log of which UI functions are accessed by logged in users. Combined with the web server access log the user name and ip address can be located for any UI actions taken by any user including proxy login users.

How to monitor user UI actions and Authentication Login

1. The audit log can be monitored from an ssh session on the eyeglass appliance.
2. Login as admin over ssh
3. Run this command to monitor user interface login in real-time
4. `tail -n 100 -f /opt/superna/sca/logs/apiaudit.log`
5. You can also search through this log with grep example
6. `grep "rsw" /opt/superna/sca/logs/apiaudit.log` **(this will locate all the api calls sent to the Ransomware Defender application UI icon. Each application has a name in the log that can be used to look for set or get or view functions.)**
7. Example log output

```

2020-12-18T11:21:02,856 API_AUDIT: [Login]/[Login][POST] - No session
2020-12-18T11:21:13,838 API_AUDIT: [Login]/[Login][GET] - Session: demo1@AD2.TEST / 4k2soc8krmq8126h0j24j6edo6
2020-12-18T11:21:13,841 API_AUDIT: [auth]/[Roles][GET] - Session: demo1@AD2.TEST / 4k2soc8krmq8126h0j24j6edo6
2020-12-18T11:21:13,892 API_AUDIT: [utils]/[Licenses][GET] - Session: demo1@AD2.TEST / 4k2soc8krmq8126h0j24j6edo6
2020-12-18T11:21:13,893 API_AUDIT: [eyeglass]/[AppGuideSetting][GET] - Session: demo1@AD2.TEST / 4k2soc8krmq8126h0j24j6edo6
2020-12-18T11:21:13,894 API_AUDIT: [utils]/[Licenses][GET] - Session: demo1@AD2.TEST / 4k2soc8krmq8126h0j24j6edo6
2020-12-18T11:21:14,066 API_AUDIT: [Login]/[Login][GET] - Session: demo1@AD2.TEST / 4k2soc8krmq8126h0j24j6edo6
2020-12-18T11:21:14,068 API_AUDIT: [Login]/[Login][GET] - Session: demo1@AD2.TEST / 4k2soc8krmq8126h0j24j6edo6
2020-12-18T11:21:14,080 API_AUDIT: [Login]/[Login][GET] - Session: demo1@AD2.TEST / 4k2soc8krmq8126h0j24j6edo6
2020-12-18T11:21:14,106 API_AUDIT: [dr]/[RetrieveDRParams][GET] - Session: demo1@AD2.TEST / 4k2soc8krmq8126h0j24j6edo6
2020-12-18T11:21:14,179 API_AUDIT: [rsw]/[RSWEventFiles][GET] - Session: demo1@AD2.TEST / 4k2soc8krmq8126h0j24j6edo6
2020-12-18T11:21:14,476 API_AUDIT: [mail]/[EmailHandler][GET] - Session: demo1@AD2.TEST / 4k2soc8krmq8126h0j24j6edo6
2020-12-18T11:21:24,740 API_AUDIT: [refresh]/[RefreshController][GET] - Session: demo1@AD2.TEST / 4k2soc8krmq8126h0j24j6edo6
2020-12-18T11:21:24,818 API_AUDIT: [services]/[Services][GET] - Session: demo1@AD2.TEST / 4k2soc8krmq8126h0j24j6edo6
2020-12-18T11:21:34,730 API_AUDIT: [refresh]/[RefreshController][GET] - Session: demo1@AD2.TEST / 4k2soc8krmq8126h0j24j6edo6
2020-12-18T11:21:44,776 API_AUDIT: [refresh]/[RefreshController][GET] - Session: demo1@AD2.TEST / 4k2soc8krmq8126h0j24j6edo6
2020-12-18T11:21:54,756 API_AUDIT: [refresh]/[RefreshController][GET] - Session: demo1@AD2.TEST / 4k2soc8krmq8126h0j24j6edo6
2020-12-18T11:22:00,042 API_AUDIT: [sera]/[Heartbeat_Post][POST] - No session
2020-12-18T11:22:00,046 API_AUDIT: [sera]/[Heartbeat_Post][POST] - No session
2020-12-18T11:22:00,054 API_AUDIT: [sera]/[Heartbeat_Post][POST] - No session
2020-12-18T11:22:00,058 API_AUDIT: [sera]/[Heartbeat_Post][POST] - No session
2020-12-18T11:22:00,068 API_AUDIT: [sera]/[Heartbeat_Post][POST] - No session
2020-12-18T11:22:00,107 API_AUDIT: [sera]/[Heartbeat_Post][POST] - No session
2020-12-18T11:22:00,170 API_AUDIT: [sera]/[Healthcheck][GET] - No session
2020-12-18T11:22:04,721 API_AUDIT: [refresh]/[RefreshController][GET] - Session: demo1@AD2.TEST / 4k2soc8krmq8126h0j24j6edo6
2020-12-18T11:22:04,799 API_AUDIT: [services]/[Services][GET] - Session: demo1@AD2.TEST / 4k2soc8krmq8126h0j24j6edo6
2020-12-18T11:22:14,603 API_AUDIT: [refresh]/[RefreshController][GET] - Session: demo1@AD2.TEST / 4k2soc8krmq8126h0j24j6edo6
2020-12-18T11:22:24,588 API_AUDIT: [refresh]/[RefreshController][GET] - Session: demo1@AD2.TEST / 4k2soc8krmq8126h0j24j6edo6
2020-12-18T11:22:34,748 API_AUDIT: [refresh]/[RefreshController][GET] - Session: demo1@AD2.TEST / 4k2soc8krmq8126h0j24j6edo6
2020-12-18T11:22:44,795 API_AUDIT: [refresh]/[RefreshController][GET] - Session: demo1@AD2.TEST / 4k2soc8krmq8126h0j24j6edo6

```

a.

8. How to view the web server access log to locate IP address information

- a. The web server access log can be used to locate the IP address of a logged user. This requires matching the time stamps in the api access log to the web server time stamps to find the source ip address of the user session.
- b. Using ssh to the eyeglass appliance as admin
- c. sudo -s (enter admin password to become root)
- d. cd /var/log/lighttpd
- e. ls (list the files and view the file with recent date stamp)
- f. Then use cat or grep or tail to view the log file
- g. Example log below

```

[anycopy@var/log/lighttpd] # tail -f access_log-20201202
10.100.239.2 igls2 - [18/Dec/2020:11:32:53 -0500] "GET /eyeglass/resources/icons/svg/blue/script_editor_blue_192.svg HTTP/1.1" 200 959 "https://igls2/eyeglass/resources/css/superna.css" "Mozilla/5.0 (Macintosh; Intel Mac OS X 11_0_0) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/87.0.4280.88 Safari/537.36"
10.100.239.2 igls2 - [18/Dec/2020:11:32:53 -0500] "GET /eyeglass/resources/icons/svg/blue/manage_service_blue_192.svg HTTP/1.1" 200 3007 "https://igls2/eyeglass/resources/css/superna.css" "Mozilla/5.0 (Macintosh; Intel Mac OS X 11.0.0) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/87.0.4280.88 Safari/537.36"
10.100.239.2 igls2 - [18/Dec/2020:11:32:53 -0500] "GET /eyeglass/resources/fonts/fontawesome-webfont.woff?v=4.7.0 HTTP/1.1" 200 77160 "https://igls2/eyeglass/resources/coveredit-all-2.css?dc=1608309173862" "Mozilla/5.0 (Macintosh; Intel Mac OS X 11.0.0) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/87.0.4280.88 Safari/537.36"
10.100.239.2 igls2 - [18/Dec/2020:11:32:53 -0500] "GET /eyeglass/resources/icons/svg/blue/jobs_blue_192.svg HTTP/1.1" 200 1395 "https://igls2/eyeglass/resources/css/superna.css" "Mozilla/5.0 (Macintosh; Intel Mac OS X 11.0.0) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/87.0.4280.88 Safari/537.36"
10.100.239.2 igls2 - [18/Dec/2020:11:32:53 -0500] "GET /eyeglass/resources/icons/svg/blue/user_roles_blue_192.svg HTTP/1.1" 200 1400 "https://igls2/eyeglass/resources/css/superna.css" "Mozilla/5.0 (Macintosh; Intel Mac OS X 11.0.0) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/87.0.4280.88 Safari/537.36"
10.100.239.2 igls2 - [18/Dec/2020:11:32:53 -0500] "GET /eyeglass/resources/icons/svg/blue/quickstart_blue_192.svg HTTP/1.1" 200 1108 "https://igls2/eyeglass/resources/css/superna.css" "Mozilla/5.0 (Macintosh; Intel Mac OS X 11.0.0) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/87.0.4280.88 Safari/537.36"
10.100.239.2 igls2 - [18/Dec/2020:11:32:53 -0500] "GET /eyeglass/resources/icons/svg/blue/easy_auditor_blue_192.svg HTTP/1.1" 200 1068 "https://igls2/eyeglass/resources/css/superna.css" "Mozilla/5.0 (Macintosh; Intel Mac OS X 11.0.0) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/87.0.4280.88 Safari/537.36"
10.100.239.2 igls2 - [18/Dec/2020:11:32:53 -0500] "GET /eyeglass/resources/icons/svg/blue/easy_auditor_blue_192.svg HTTP/1.1" 200 3619 "https://igls2/eyeglass/resources/css/superna.css" "Mozilla/5.0 (Macintosh; Intel Mac OS X 11.0.0) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/87.0.4280.88 Safari/537.36"
10.100.239.2 igls2 - [18/Dec/2020:11:32:54 -0500] "GET /RestClient/mail/EmailHandler_dc=1608309173862?page=1&start=0&limit=25 HTTP/1.1" 200 51 "https://igls2/eyeglass/" "Mozilla/5.0 (Macintosh; Intel Mac OS X 11.0.0) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/87.0.4280.88 Safari/537.36"
10.100.239.2 igls2 - [18/Dec/2020:11:32:55 -0500] "GET /eyeglass/resources/images/desktop-icons/window-information-48.png HTTP/1.1" 200 934 "https://igls2/eyeglass/" "Mozilla/5.0 (Macintosh; Intel Mac OS X 11.0.0) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/87.0.4280.88 Safari/537.36"

```

i.

How to view the GUI API Log

1. Login as admin
2. sudo -s (enter admin password)

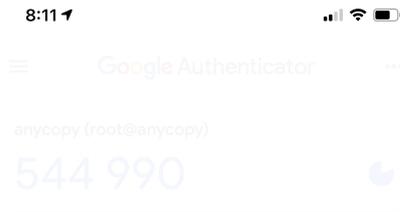
3. journalctl -u scagateway
 - a. This command will return all client browser api calls

ECA VM Hardened Virtual Secured Network (Ransomware Defender, Easy Auditor, Performance Auditor)

1. The Eyeglass and ECA installation and admin guides list firewall ports and directions required including management PC access to UI's. This feature will automatically secure the communications between Eyeglass and the ECA vm's. No open ports will be returned from ECA vm's with the exception of SSH and HTTPS. This creates a virtual secure network between Eyeglass and the ECA vm's with no external access to any ports. This feature is automatically enabled and configured.
2. Requirements:
 - a. 2.5.7 update 1
3. This release adds 2 new security features.
 - a. **Automatic Firewall for Superna VM's:** The ECA VM's need to be accessed by eyeglass over various ports. The installation of the ECA and cluster up process will apply IP tables firewall rules to only allow access to ECA ports from Eyeglass VM and between ECA VM's. This provides a secure network between Superna VM's without requiring customer infrastructure. This will be applied automatically.
 - b. **Authenticated Management UI's** - Various management UI's on ECA nodes will be accessed through a HTTPS proxy built into the ECA nodes that will require authentication. None of the ECA UI's will be directly accessible.

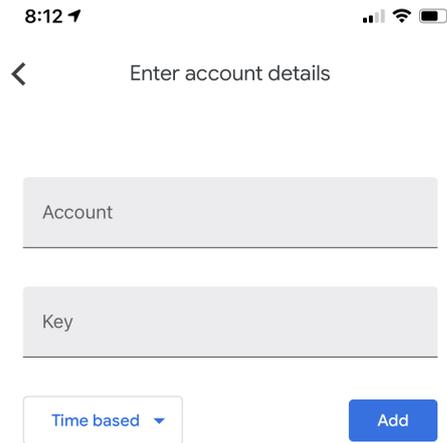
2 Factor SSH Authentication for Eyeglass, Golden copy, Search & Recover or ECA VM's

1. This procedure only secures SSH access to VMs.
2. Requirements:
 - a. Google Authenticator application on a mobile phone
 - i. [IOS](#)
 - ii. [Android](#)
3. **Installation**
 - a. ssh to the VM
 - b. sudo -s (enter password)
 - c. Install the pam module
 - d. zypper in google-authenticator-libpam
 - e. answer yes
 - f. To run the initialization app
 - i. google-authenticator
 - ii. Do you want authentication tokens to be time-based (y/n) y
 - iii. **NOTE: Very important step to complete**
 1. **You will be presented with a secret key used in the step below and multiple scratch codes. We strongly suggest saving these emergency scratch codes in a safe place, like a password manager. These codes are the only way to regain access if you lose your phone or lose access to your authenticator application, and each one can only be used once, so they really are in case of emergency.**
 - iv. Activate Google Authenticator application with output from the step above that shows the secret key.
 1. Click the plus to add a new profile and select add setup key



2.

3. Now enter the name of this VM for account name example Eyeglass and the secret key number output to the console from the step above.



4. _____

5. Your client is now configured

- v. Do you want me to update your "~/.google_authenticator" file (y/n)
y
- vi. **Do you want to disallow multiple uses of the same authentication token?** This restricts you to one login about every 30s, but it increases your chances to notice or even prevent man-in-the-middle attacks (y/n) (This part is a time-based login. We suggest answering 'yes' (y) here since this will prevent a replay attack, allowing you 30 seconds from the point of getting the code on your mobile to typing in your login prompt)
- vii. **By default, tokens are good for 30 seconds and in order to compensate for possible time-skew between the client and the server,**
we allow an extra token before and after the current time. If you

experience problems with poor time synchronization, you can increase

the window from its default size of 1:30min to about 4min.

Do you want to do so (y/n) (Answer yes for more secure, answer No to allow 8 valid codes in a 4:00-minute rolling window)

viii. If the computer that you are logging into isn't hardened against brute-force login attempts, you can enable rate-limiting for the authentication module. By default, this limits attackers to no more than 3 login attempts every 30s.

Do you want to enable rate-limiting (y/n) (yes is more secure)

g. Configuring OpenSSH

i. nano /etc/pam.d/sshd

ii. add this line to the file

1. auth required pam_google_authenticator.so

iii. control + x to save

iv. nano /etc/ssh/sshd_config

v. find this line and remove the comment at the front of the line

1. ChallengeResponseAuthentication yes

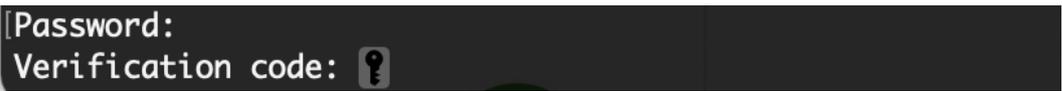
vi. control + x to save

h. Activate

i. systemctl restart sshd

i. Test login

j. ssh to the vm and enter the admin user password. You will now be prompted to enter a one time number called the verification code.

k.  Password:
Verification code: 

l. Use the Google Authenticator application and type in the code displayed in application. You have 30 seconds to enter the code to login successfully.

- m. **NOTE: The settings above determine what happens if you enter a bad or out of sync verification code. You may get rate limited to login if you have failed login attempts**
- n. done.

How to Mediate python vulnerability CVE-2022-35861

1. <https://nvd.nist.gov/vuln/detail/CVE-2022-35861>
2. Procedure locks the local user accounts so they cannot be used.
3. Login to eyeglass
4. `sudo -s`
5. These local users can be disabled to mitigate this low risk vulnerability
 - a. `screenshots:x:1001:100::/home/screenshots:/bin/bash`
 - b. `rwdefend:x:1002:100::/home/rwdefend:/bin/bash`
 - c. `auditor:x:1003:100::/home/auditor:/bin/bash`
6. `usermod -L rwdefend`
7. `usermod -L auditor`
8. `usermod -L screenshots`
9. Done.

7. InPlace Appliance Open Suse OS Upgrade

Home Top

- [Open Suse to SUSE Linux Enterprise Server \(SLES\) Migration](#)
 - [Summary](#)
 - [Prerequisites](#)
 - [How to Migrate Open Suse to SLES Licensed Subscription OS](#)
 - [How to Inplace OS Upgrade SLES licensed OS from 15.4 to 15.x](#)
- [In-Place ECA Appliance OS OpenSuse 15.4 to 15.5 Upgrade](#)
 - [Overview](#)
 - [Procedure](#)
- [In-Place Eyeglass Appliance OS OpenSuse 15.4 to 15.5 Upgrade](#)
 - [Overview](#)
 - [Procedure](#)

Open Suse to SUSE Linux Enterprise Server (SLES) Migration

Summary

This procedure allows customers to convert Open Suse to a licensed and supported SLES OS and allows inplace upgrade options using the YAST GUI for simplified OS patch management and major release upgrades. This subscription to SLES includes patches, upgrades and access to OS technical support if required.

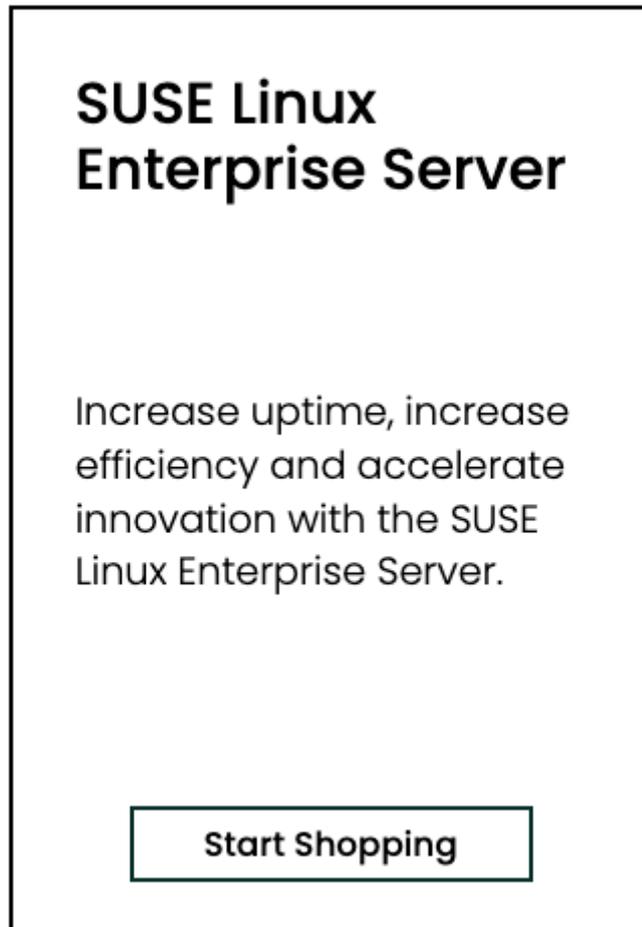
Prerequisites

1. Open Suse 15.4 or 15.5.
 - a. How to check.
 - i. Login as admin over ssh
 - ii. `cat /etc/os-release`
2. A SLES subscription purchased from Suse

3. Internet repository firewall configuration or proxy that allows https over port 443 to <https://updates.suse.com>

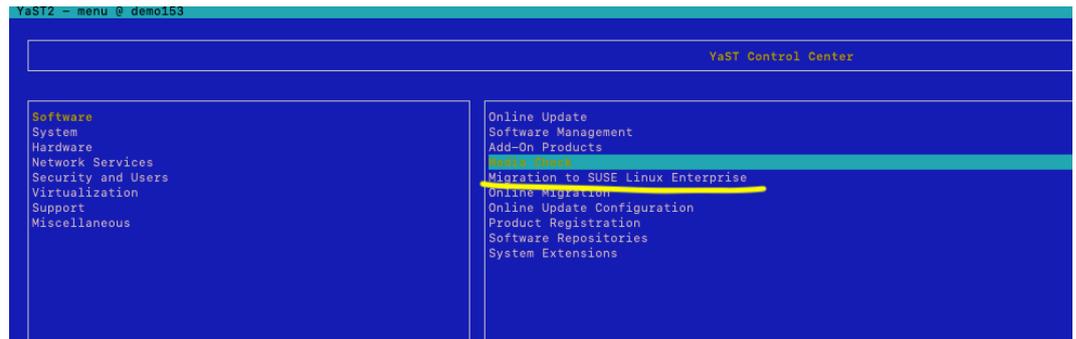
How to Migrate Open Suse to SLES Licensed Subscription OS

1. Purchase an OS subscription from <https://www.suse.com/shop/>



- a. Select
2. You will need to retain the email address used to purchase and the Activation Code
 - a. Login to the Eyeglass vm as admin user over ssh.
 - b. `sudo -s`
 - c. `zypper in yast2-migration-sle`
 - d. Once the installation is completed
 - e. `yast`

- f. select the menu on the right for **Migrate to SUSE Linux Enterprise**



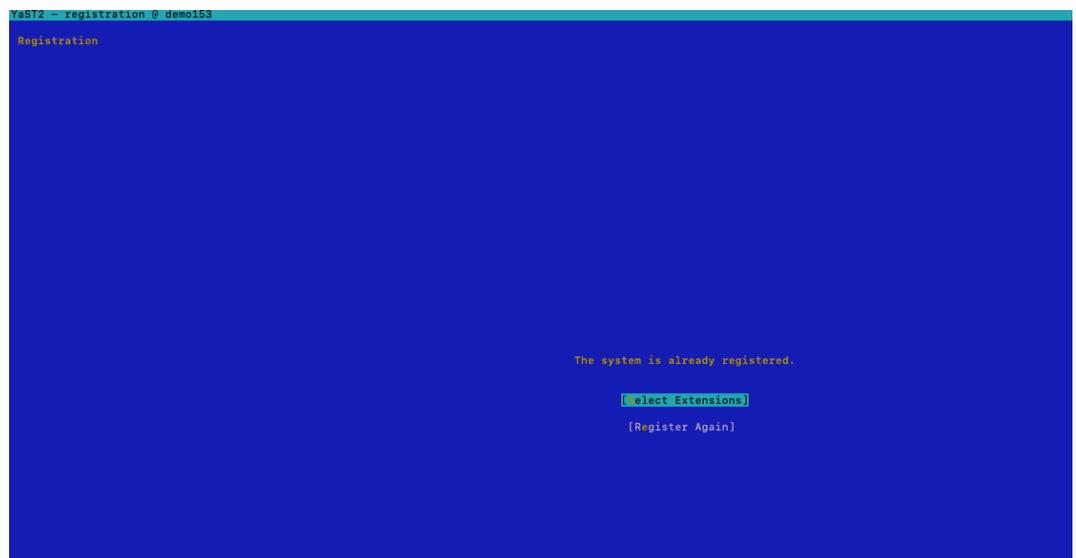
g.

- h. Follow the on screen steps to enter your email address and activation code.

3. Once this above steps are complete you should have a registered licensed SLES operating system, confirm this using this command

- a. `cat /etc/os-release`
- b. You will look for OS ID of **SLES**

4. Verify your VM is registered using the **Product Registration** menu



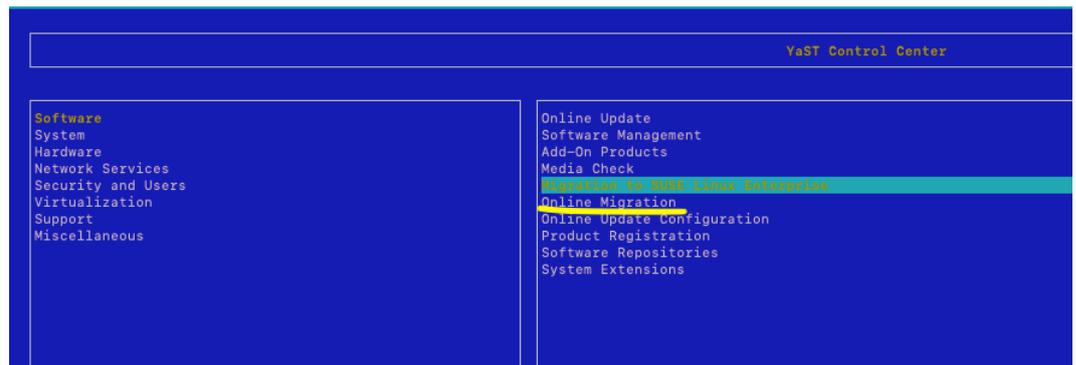
a.

5. **Done**

How to Inplace OS Upgrade SLES licensed OS from 15.4 to 15.x

- 1. NOTE: Only licensed SLES appliances offer the GUI upgrade inplace.
- 2. Login as admin over ssh
 - a. `sudo -s`
 - b. `yast`

- c. Select the menu option on the right **Online Migration**



d.

- e. Follow the online GUI steps to upgrade to 15.5,
- f. NOTE: During the upgrade 2 conflicts will be raised for dependencies, choose continue the upgrade.
- g. NOTE: The upgrade can take 30-60 minutes depending on download speeds. A reboot will be completed at the end of the upgrade.

3. Run Superna Eyeglass 15.5 installer run file after the OS upgrade to repair dependencies.

- a. Prepare for upgrade

- i. nano /etc/os-release

1. Change ID_LIKE from "suse" to "suse opensuse"

2. press control+x to save and exit

- b. Download the Eyeglass upgrade installer file (for 15.5) from support.superna.net portal and scp to Eyeglass appliance

- c. Change mode of the upgrade installer file

- i. `chmod 777 <name of upgrade file>`

- d. Execute the upgrade file

- i. `./<name of upgrade file`

- e. Once completed, login to Eyeglass Web UI to verify

In-Place ECA Appliance OS OpenSuse 15.4 to 15.5 Upgrade

Overview

In-place upgrades of the operating system should only be done if security patches are required for your deployment. Normally the OVA upgrade option is best to upgrade the OS and appliance at the same time. Only use this procedure if directed by support.

- i** NOTE: The operating system is the customer's responsibility, and assisted upgrades of the OS are not provided under the support contract. The supported method is using new OVA deployment and restoring the configuration for assisted upgrades to a new OS. **This procedure is provided "as is" with no support.**
- i** NOTE: Mandatory to take a VM-level snapshot before starting this procedure. The only recovery option on a failed OS upgrade is reverting a snapshot. Support is unable to recover a failed OS upgrade and will request the snapshot is reverted and schedule the OVA upgrade method.
- i** NOTE: This procedure is **only from 15.4 to 15.5**
- i** NOTE: **Recently tested with Search & Recover, see note in the procedure to avoid data loss of the index.**
- i** NOTE: This procedure requires an Internet connection

Procedure

1. Take VM snapshot of each ECA nodes VM

2. ssh to each ECA Appliance as `ecaadmin` user
3. `sudo -s`
4. `zypper refresh`
5. `zypper update`

a. It will prompt to continue - answer with “y”

...

```
The following package requires a system reboot:
kernel-default-5.14.21-150400.24.63.1
```

```
113 packages to upgrade, 3 new.
Overall download size: 577.8 MiB. Already cached: 0 B. After the
operation, additional 239.9 MiB will be used.
```

```
Note: System reboot required.
Continue? [y/n/v/...? shows all options] (y): y
```

...

b. It will also detect some conflicts - answer “yes” to continue:

...

```
Checking for file conflicts:
.....
.....
Detected 1 file conflict:
```

```
File /usr/share/bash-completion/completions/lsblk
from install of
    util-linux-systemd-2.37.2-150400.8.17.1.x86_64 (Update
repository with updates from SUSE Linux Enterprise 15)
conflicts with file from package
    util-linux-2.36.2-2.29.x86_64 (@System)
```

```
File conflicts happen when two packages attempt to install files
with the same name but different contents. If you continue,
conflicting files will be replaced losing the previous content.
Continue? [yes/no] (no): yes
```

...

c. Once completed, a reboot is required

d. Wait till all ECA nodes show that message (to indicate that the zypper update has completed). Before reboot, bring down the ECA cluster by running this command **only on ECA node 1**:

`exit` (to exit from sudo, and return to ecaadmin user)

Then execute cluster down command:

```
ecactl cluster down
```

e. After the `ecactl cluster down` process has completed, reboot each ECA nodes: (on ECA node 1, need to enter `sudo -s` mode, on other ECA nodes, there are still in sudo mode as per step #3. Execute:

```
reboot
```

6. After it is rebooted, login again to each ECA nodes as `ecaadmin`

```
sudo -s
```

7. New 4096 bit RSA signing key (Reference: [here](#))

a. The new 4096 bit RSA signing key was also introduced as part of openSUSE Leap 15.5 as well as 15.4 via a maintenance update, additionally a new 4096 bit RSA backports key was introduced.

b. Leap 15.4 users are expected to update their system prior to the upgrade of 15.5.

c. Users of 15.4 release currently need to import key manually with:

```
rpm --import /usr/lib/rpm/gnupg/keys/gpg-pubkey-29b700a4-62b07e22.asc
```

d. And also the new 4096 openSUSE Backports key needs to be imported:

```
rpm --import /usr/lib/rpm/gnupg/keys/gpg-pubkey-25db7ae0-645bae34.asc
```

8. Backup sudoers file

```
cp /etc/sudoers /etc/sudoers.bak
```

9. Update the repo

- a. Check if your Leap repos defined in `/etc/zypp/repos.d/` are using the `$releasever` variable already.

```
sudo grep baseurl /etc/zypp/repos.d/*.repo
```

- b. Example:

...

```
/etc/zypp/repos.d/repo-backports-debug-  
update.repo:baseurl=http://download.opensuse.org/update/leap/$rele  
asever/backports_debug/  
/etc/zypp/repos.d/repo-backports-  
update.repo:baseurl=http://download.opensuse.org/update/leap/$rele  
asever/backports/  
/etc/zypp/repos.d/repo-debug-non-  
oss.repo:baseurl=http://download.opensuse.org/debug/distribution/l  
eap/$releasever/repo/non-oss/  
...
```

...

- c. If they are still hard-coded with a particular Leap version number, then you need to modify them first

Example:

```
sed -i 's/15.4/${releasever}/g' /etc/zypp/repos.d/*.repo
```

10. Refresh with the new Repos of 15.5

```
zypper --releasever=15.5 refresh
```

11. Execute the full distribution upgrade to 15.5

```
zypper --releasever=15.5 dup
```

12. It will detect some dependency problems like the following. Choose the solution option to deinstallation the older package. Example:

...

3 Problems:

Problem: the to be installed samba-dsdb-modules-4.13.4+git.187.5ad4708741a-1.34.x86_64 requires 'libsamba-errors.so.1(SAMBA_ERRORS_1)(64bit)', but this requirement cannot be provided

Problem: the installed libsnp40-5.9.3-150300.15.8.1.x86_64 requires 'snmp-mibs >= 5.9.3', but this requirement cannot be provided

Problem: the installed samba-dsdb-modules-4.15.13+git.636.53d93c5b9d6-150400.3.23.1.x86_64 requires 'libldb.so.2(LDB_2.4.4)(64bit)', but this requirement cannot be provided

Problem: the to be installed samba-dsdb-modules-4.13.4+git.187.5ad4708741a-1.34.x86_64 requires 'libsamba-errors.so.1(SAMBA_ERRORS_1)(64bit)', but this requirement cannot be provided

deleted providers: samba-client-libs-4.15.13+git.636.53d93c5b9d6-150400.3.23.1.x86_64
not installable providers: libsamba-errors0-4.13.4+git.187.5ad4708741a-1.34.x86_64[openSUSE-Leap-15.3-1]

Solution 1: Following actions will be done:

keep obsolete samba-dsdb-modules-4.15.13+git.636.53d93c5b9d6-150400.3.23.1.x86_64

keep obsolete samba-client-libs-4.15.13+git.636.53d93c5b9d6-150400.3.23.1.x86_64

Solution 2: deinstallation of samba-client-libs-4.15.13+git.636.53d93c5b9d6-150400.3.23.1.x86_64

Solution 3: deinstallation of samba-dsdb-modules-4.15.13+git.636.53d93c5b9d6-150400.3.23.1.x86_64

Solution 4: break samba-dsdb-modules-4.13.4+git.187.5ad4708741a-1.34.x86_64 by ignoring some of its dependencies

Choose from above solutions by number or skip, retry or cancel [1/2/3/4/s/r/c/d/?] (c): 2

Problem: the installed libsnp40-5.9.3-150300.15.8.1.x86_64 requires 'snmp-mibs >= 5.9.3', but this requirement cannot be provided

not installable providers: snmp-mibs-5.9.3-150300.15.3.1.x86_64[repo-sle-update]

Solution 1: deinstallation of libsnp30-5.7.3-10.12.1.x86_64

```
Solution 2: deinstallation of libsnmp40-5.9.3-150300.15.8.1.x86_64
Solution 3: keep obsolete libsnmp30-5.7.3-10.12.1.x86_64
Solution 4: break libsnmp40-5.9.3-150300.15.8.1.x86_64 by ignoring some of its dependencies
```

```
Choose from above solutions by number or skip, retry or cancel [1/2/3/4/s/r/c/d/?] (c): 1
```

```
Problem: the installed samba-dsdb-modules-4.15.13+git.636.53d93c5b9d6-150400.3.23.1.x86_64 requires 'libldb.so.2(LDB_2.4.4)(64bit)', but this requirement cannot be provided
deleted providers: libldb2-2.4.4-150400.4.11.1.x86_64
Solution 1: Following actions will be done:
keep obsolete libldb2-2.4.4-150400.4.11.1.x86_64
keep obsolete samba-ldb-ldap-4.15.13+git.636.53d93c5b9d6-150400.3.23.1.x86_64
Solution 2: deinstallation of samba-ldb-ldap-4.15.13+git.636.53d93c5b9d6-150400.3.23.1.x86_64
Solution 3: deinstallation of yast2-adcommon-python-1.6-3.3.10.noarch
Solution 4: break samba-dsdb-modules-4.15.13+git.636.53d93c5b9d6-150400.3.23.1.x86_64 by ignoring some of its dependencies
```

```
Choose from above solutions by number or skip, retry or cancel [1/2/3/4/s/r/c/d/?] (c): 2
```

...

13. It will prompt us to continue. Answer with “y”

...

```
The following 2 packages require a system reboot:
kernel-default-5.14.21-150500.53.2 kernel-preempt-5.3.18-57.3
```

```
218 packages to upgrade, 38 to downgrade, 68 new, 1 to reinstall, 8 to remove.
```

```
Overall download size: 784.9 MiB. Already cached: 0 B. After the operation, additional 659.5 MiB will be used.
```

```
Note: System reboot required.
```

```
Continue? [y/n/v/...? shows all options] (y): y
```

...

14. It detects some conflicts, answer “yes” to continue

...

```
Detected 3 file conflicts:
```

```
File /etc/motd
  from install of
    openSUSE-release-15.5-lp155.286.1.x86_64 (Main Repository)
  conflicts with file from package
    eyeglass_cluster_appliance-2.5.12-23175.lp15.4.x86_64
  (@System)
```

```
File /etc/sudoers
  from install of
    sudo-1.9.12p1-150500.5.1.x86_64 (Main Repository)
  conflicts with file from package
    eyeglass_cluster_appliance-2.5.12-23175.lp15.4.x86_64
  (@System)
```

```
File /usr/share/bash-completion/completions/lsblk
  from install of
    util-linux-systemd-2.37.4-150500.7.2.x86_64 (Main Repository)
  conflicts with file from package
    util-linux-2.36.2-2.29.x86_64 (@System)
```

```
File conflicts happen when two packages attempt to install files
with the same name but different contents. If you continue,
conflicting files will be replaced losing the previous content.
Continue? [yes/no] (no): yes
```

...

15. Once completed, it will suggest a reboot. But **DO NOT** reboot first

...

```
Since the last system boot core libraries or services have been
updated.
Reboot is suggested to ensure that your system benefits from these
updates.
```

...

16. Before reboot, on each ECA nodes, copy `/etc/sudoers.bak` back to `/etc/sudoers`

```
cp /etc/sudoers.bak /etc/sudoers
```

17. Once done, reboot each ECA nodes:

```
reboot
```

18. After rebooted, re-login via ssh to each ECA nodes as `ecaadmin`

19. Verify OS version is **15.5** on each ECA nodes, from `cat /etc/os-release`

```
NAME="openSUSE Leap"
VERSION="15.5"
ID="opensuse-leap"
ID_LIKE="suse opensuse"
VERSION_ID="15.5"
PRETTY_NAME="openSUSE Leap 15.5"
ANSI_COLOR="0;32"
CPE_NAME="cpe:/o:opensuse:leap:15.5"
BUG_REPORT_URL="https://bugs.opensuse.org"
HOME_URL="https://www.opensuse.org/"
DOCUMENTATION_URL="https://en.opensuse.org/Portal:Leap"
LOGO="distributor-logo-Leap"
```

20. Download the latest ECA upgrade installer file (for 15.5) from **support.superna.net** portal and scp that file to **ECA node 1**

21. Change mode of the upgrade installer file

```
chmod +x <name of upgrade file>
```

22. Execute the upgrade file

```
./<name of upgrade file>
```

23. **NOTE: Only execute this step on ECA VM's. WARNING: Do not execute this step on Search & Recover, data loss will result.**

VM's. Once the upgrade process has completed, from **ECA node 1**, execute the following command (**Note**: the following command is a one-line command): :

```
ecactl cluster exec "sudo rm -rf /opt/superna/mnt/zk-ramdisk/*
/opt/data/kafka-logs/* \; sudo systemctl restart docker"
```

24. Once completed, then from **ECA node 1**, bring up the ECA cluster

```
ecactl cluster up
```

25. Wait till cluster up has fully completed. Once completed, login to Eyeglass WebUI and from Eyeglass WebUI Dashboard open the **Managed Services** icon and verify that all ECA nodes services are in an **Active** and healthy state.

In-Place Eyeglass Appliance OS OpenSuse 15.4 to 15.5 Upgrade

Overview

In-place upgrades of the operating system should only be done if security patches are required for your deployment. Normally the OVA upgrade option is best to upgrade the OS and appliance at the same time. Only use this procedure if directed by support.

- i** NOTE: The operating system is the customer's responsibility, and assisted upgrades of the OS are not provided under the support contract. The supported method is using new OVA deployment and restoring the configuration for assisted upgrades to a new OS. **This procedure is provided "as is" with no support.**
- i** NOTE: Mandatory to take a VM-level snapshot before starting this procedure. The only recovery option on a failed OS upgrade is reverting a snapshot. Support is unable to recover a failed OS upgrade and will request the snapshot is reverted and schedule the OVA upgrade method.
- i** NOTE: This procedure is **only from 15.4 to 15.5**
- i** NOTE: This procedure requires an Internet connection

Procedure

1. Take **VM snapshot** of **Eyeglass Appliance VM**
2. ssh to Eyeglass Appliance as **admin** user
3. `sudo -s`
4. `zypper refresh`
5. `zypper update`

- a. It will prompt to continue - answer with “y”

...

```
Note: System reboot required.  
Continue? [y/n/v/...? shows all options] (y): y
```

...

- b. It will also detect some conflicts - answer “yes” to continue:

...

```
Checking for file conflicts:  
.....  
.....[error]  
Detected 1 file conflict:
```

```
File /usr/share/bash-completion/completions/lsblk  
from install of  
    util-linux-systemd-2.37.2-150400.8.17.1.x86_64 (Update  
repository with updates from SUSE Linux Enterprise 15)  
conflicts with file from package  
    util-linux-2.36.2-2.29.x86_64 (@System)
```

```
File conflicts happen when two packages attempt to install  
files with the same name but different contents. If you  
continue, conflicting files will be replaced losing the  
previous content.
```

```
Continue? [yes/no] (no): yes
```

...

- c. Once completed, it will require a reboot
- d. Before reboot, stop sca service

```
systemctl stop sca
```
- e. After the sca service is stopped, reboot appliance:

```
reboot
```

6. After it is rebooted, re-login to Eyeglass SSH again as admin user

```
sudo -s
```

7. New 4096 bit RSA signing key (Reference: [here](#))

- a. The new 4096 bit RSA signing key was also introduced as part of openSUSE Leap 15.5 as well as 15.4 via a maintenance update, additionally a new 4096 bit RSA backports key was introduced.
- b. Leap 15.4 users are expected to update their system prior to the upgrade of 15.5.
- c. Users of 15.4 release currently need to import key manually with:

```
rpm --import /usr/lib/rpm/gnupg/keys/gpg-pubkey-29b700a4-62b07e22.asc
```

- d. And also the new 4096 openSUSE Backports key needs to be imported:

```
rpm --import /usr/lib/rpm/gnupg/keys/gpg-pubkey-25db7ae0-645bae34.asc
```

8. Backup sudoers file

```
cp /etc/sudoers /etc/sudoers.bak
```

9. Update the repo

- a. Check if your Leap repos defined in `/etc/zypp/repos.d/` are using the `$releasever` variable already.

```
sudo grep baseurl /etc/zypp/repos.d/*.repo
```

b. Example:

...

```
/etc/zypp/repos.d/repo-backports-debug-  
update.repo:baseurl=http://download.opensuse.org/update/leap/  
$releasever/backports_debug/  
/etc/zypp/repos.d/repo-backports-  
update.repo:baseurl=http://download.opensuse.org/update/leap/  
$releasever/backports/  
/etc/zypp/repos.d/repo-debug-non-  
oss.repo:baseurl=http://download.opensuse.org/debug/distribut  
ion/leap/$releasever/repo/non-oss/  
...
```

c. If they are still hard-coded with a particular Leap version number, then you need to modify them first

Example:

```
sed -i 's/15.4/${releasever}/g' /etc/zypp/repos.d/*.repo
```

10. Refresh with the new Repos of **15.5**

```
zypper --releasever=15.5 refresh
```

11. Execute the full distribution upgrade to **15.5**

```
zypper --releasever=15.5 dup
```

12. It will detect some dependency problems like the following. Choose the solution option to deinstallation the older package. Example

...

```
Problem: the installed libsnmp40-5.9.3-150300.15.8.1.x86_64  
requires 'snmp-mibs >= 5.9.3', but this requirement cannot be  
provided
```

```
not installable providers: snmp-mibs-5.9.3-  
150300.15.3.1.x86_64[repo-sle-update]
```

```
Solution 1: deinstallation of libsnmp30-5.7.3-10.12.1.x86_64
```

```
Solution 2: deinstallation of libsnmp40-5.9.3-150300.15.8.1.x86_64
Solution 3: keep obsolete libsnmp30-5.7.3-10.12.1.x86_64
Solution 4: break libsnmp40-5.9.3-150300.15.8.1.x86_64 by ignoring some of its dependencies

Choose from above solutions by number or cancel [1/2/3/4/c/d/?]
(c): 1
```

...

13. It will prompt us to continue. Answer with “y”

...

```
The following 2 packages require a system reboot:
  kernel-default-5.14.21-150500.53.2 kernel-preempt-5.3.18-57.3

217 packages to upgrade, 27 to downgrade, 45 new, 1 to reinstall,
4 to remove.
Overall download size: 748.9 MiB. Already cached: 0 B. After the
operation, additional 660.4 MiB will be used.

Note: System reboot required.
Continue? [y/n/v/...? shows all options] (y): y
```

...

14. It detects some conflicts, answer “yes” to continue

...

```
Checking for file conflicts:
.....
.....[error]
Detected 7 file conflicts:

File /etc/lighttpd/conf.d/magnet.conf
from install of
  lighttpd-mod_magnet-1.4.69-bp155.1.6.x86_64 (Main Repository)
conflicts with file from package
  eyeglass_deps-2.5.11-23140.lp15.4.x86_64 (@System)

File /etc/lighttpd/conf.d/proxy.conf
from install of
  lighttpd-1.4.69-bp155.1.6.x86_64 (Main Repository)
conflicts with file from package
```

```
eyeglass_deps-2.5.11-23140.lp15.4.x86_64 (@System)
File /etc/lighttpd/lighttpd.conf
  from install of
    lighttpd-1.4.69-bp155.1.6.x86_64 (Main Repository)
  conflicts with file from package
    eyeglass_deps-2.5.11-23140.lp15.4.x86_64 (@System)

File /etc/lighttpd/modules.conf
  from install of
    lighttpd-1.4.69-bp155.1.6.x86_64 (Main Repository)
  conflicts with file from package
    eyeglass_deps-2.5.11-23140.lp15.4.x86_64 (@System)

File /etc/motd
  from install of
    openSUSE-release-15.5-lp155.286.1.x86_64 (Main Repository)
  conflicts with file from package
    eyeglass_deps-2.5.11-23140.lp15.4.x86_64 (@System)

File /etc/sudoers
  from install of
    sudo-1.9.12p1-150500.5.1.x86_64 (Main Repository)
  conflicts with file from package
    eyeglass_deps-2.5.11-23140.lp15.4.x86_64 (@System)

File /usr/share/bash-completion/completions/lisblk
  from install of
    util-linux-systemd-2.37.4-150500.7.2.x86_64 (Main Repository)
  conflicts with file from package
    util-linux-2.36.2-2.29.x86_64 (@System)

File conflicts happen when two packages attempt to install files
with the same name but different contents. If you continue,
conflicting files will be replaced losing the previous content.
Continue? [yes/no] (no): yes
```

...

15. Once completed, it will suggest a reboot. But **DO NOT reboot first**

...

```
Since the last system boot core libraries or services have been
updated.
Reboot is suggested to ensure that your system benefits from these
updates.
```

...

16. Before reboot, copy `/etc/sudoers.bak` back to `/etc/sudoers`

```
cp /etc/sudoers.bak /etc/sudoers
```

17. Stop sca service

```
systemctl stop sca
```

18. Reboot Eyeglass appliance: `reboot`

19. After rebooted, re-login via ssh to Eyeglass as **admin** user

20. Verify OS version is **15.5** on Eyeglass appliance , from `cat /etc/os-release`

```
NAME="openSUSE Leap"
VERSION="15.5"
ID="opensuse-leap"
ID_LIKE="suse opensuse"
VERSION_ID="15.5"
PRETTY_NAME="openSUSE Leap 15.5"
ANSI_COLOR="0;32"
CPE_NAME="cpe:/o:opensuse:leap:15.5"
BUG_REPORT_URL="https://bugs.opensuse.org"
HOME_URL="https://www.opensuse.org/"
DOCUMENTATION_URL="https://en.opensuse.org/Portal:Leap"
LOGO="distributor-logo-Leap"
```

21. Elevate

```
sudo -s
```

22. Unlink symbolic links

```
unlink /opt/superna/document_generator/startOrientDb.sh
unlink /opt/superna/document_generator/databaseScriptSca.sh
unlink /opt/superna/document_generator/eyeglassAnalyzer
unlink /opt/superna/document_generator/dbmanager
unlink /opt/superna/document_generator/orientDBManager
```

23. Continue by providing “y” answer

...

```
The following 3 NEW packages are going to be installed:  
glibc-devel libxcrypt-devel linux-glibc-devel
```

```
3 new packages to install.  
Overall download size: 1.8 MiB. Already cached: 0 B. After the  
operation, additional 7.5 MiB will be used.  
Continue? [y/n/v/...? shows all options] (y): y
```

...

24. Download the Eyeglass upgrade installer file (for 15.5) from **support.superna.net** portal and scp to Eyeglass appliance

25. Change mode of the upgrade installer file

```
chmod 777 <name of upgrade file>
```

26. Execute the upgrade file

```
./<name of upgrade file>
```

27. Once completed, login to Eyeglass Web UI to verify

© Superna Inc