Table of Contents

1. Failover Release Notes

Home Top

- (All Releases) Latest Release
- (Release 1.6.3 >) Snapshot schedule expiration offset has OneFS API bug that adds extra time to creation of the snapshot schedule.
- (All Releases) SyncIQ file filters not supported
- (Release < 1.8.0) Technical Advisory #10
- (All Releases) OneFS Failover and Failback without waiting for quota scan job to complete
- (All Releases) SPNs not updated during failover for OneFS8 non-default groupnet AD provider (T3848)
- (Release 1.8.3) Failure to run Resync Prep step during DFS Failover Deletes Shares on Target Cluster (T4145)
- (Releases < 1.9.0) DR Assistant returns "Error Retrieving Zones Undefined" if many access zones exist
- (Releases => 1.9.0) DFS Failover Enhancement to handle partial or complete Share Rename failures
- (All Releases) Access Zone Failover Networking Roll back on failures
- (All Releases) PowerScale OneFS 8.0.0.5 API load sharing with Smartconnect issue

- (ALL RELEASES) Time to complete steps for Allow Writes and Preparation to Failback Unknown

All notes should be followed prior to any failover attempt

(All Releases) Latest Release

1. The latest version of Eyeglass has been installed, we add enhancements to each release based on customer failovers to prevent or document anything that will block or impact failover. Not upgrading to latest release affects your entitlement to support for a planned failover event. Excludes real DR failovers.
2. **If you are planning failover and want Eyeglass DR readiness assessment, we require 7 days advanced notice and support logs submitted.**
3. **If your Eyeglass installation is N-2 releases where N is the currently published GA release, you may choose to stay on a release that is N-2 without affecting support only if the following steps are completed:**
    a. Open a case with support and upload support logs. **Then follow instructions below**.
        i. State planned failover will use a release that is within N-2 and update the case.
        ii. State the **Failover Features** of N release has been reviewed here and update the case to state confirmation.

iii. Support will request confirmation in the case that N, N-1 and N-2 "failover section" of the release notes (example here) has been reviewed and that you are confirming and accepting the risk in your environment.

iv. **NOTE: Failure to complete #1-3, will affect support entitlement of using N-2 release for a planned failover.**

4. **Target DR Releases if already running one of these releases**

    a. **Check with support for target release or consult software matrix above.**

    b. **NOTE: If not running these releases upgrade to the latest GA release is required.**

    c. **NOTE: if using one of these releases, all release notes apply and assumed read and accepted prior to any planned failover.**

(Release 1.6.3 >) Snapshot schedule expiration offset has OneFS API bug that adds extra time to creation of the snapshot schedule.

1. This results in an expiration on the DR cluster, that can be greater than entered on the source cluster. example expire in 20 days will be 22 days on the target cluster. Different units of off set all result in a value greater than entered. After failover the DR (target cluster) value will be synced back to the source (Prod cluster). Thereby losing the original expiry off set and extending the expire time by a new offset from the API error. This has been raised with EMC as SR to resolve.

2. **Work around**: Before failover ensure a cluster report has been generated (cluster reports icon), or an existing emailed cluster report exists. Post Failover re-enter the original values on the DR snapshot schedules using the cluster report values from the source cluster as a reference.

    a. Another option is disable Snapshot Sync jobs in the jobs window if the above workaround does not meet your needs to preserve expiry of snapshot settings.

(All Releases) SyncIQ file filters not supported

1. File pattern filters are **NOT** synced on failover, these pattern filters can result in unprotected data during failover and failback. Failover and failback work flows require customer testing for their own use case. All file filter scenario's are untested without support for custom workflows related to file filters failover failback issue not present under normal failover and failback workflow

(Release < 1.8.0) Technical Advisory #10

1. For the case where PowerScale clusters have been added to Eyeglass using FQDN, uncontrolled failover for case where source cluster is not reachable does not start and gives the error ""Error performing zone failover: Cannot find associated source network element for zone". This issue will be addressed in a 1.8.1 patch. Eyeglass installations using FQDN to add clusters must upgrade to this patch once available. **Workaround**: please refer to Technical Advisory #10

(All Releases) OneFS Failover and Failback without waiting for quota scan job to complete

1. In OneFS 8 quota scan job is started as soon as a quota is created (cannot be disabled on OneFS 8). Resync Prep on failover or failback will fail when Quota scan job is active on a path on the target cluster. Do not add/edit quotas before or during failover. If you have Quotas with snapshot overhead enabled, deleting a snapshot may trigger a quota scan. Also, after Eyeglass failover quotas are created by Eyeglass and quota scan will start. If failback is attempted right away (typically testing only scenario) without waiting for quota scan to complete the resync prep step is blocked from running due to domain lock from the quota

scan. **Workaround: Wait for quota scan job to complete before attempting failover or failback. Use cluster running jobs UI to verify if quota scan is running or not before attempting to failover.**

(All Releases) SPNs not updated during failover for OneFS8 non-default groupnet AD provider (T3848)

1. For the case where OneFS 8 is configured with multiple groupnets and different AD provider between groupnets, the SPN update during failover does not succeed for non-default groupnet AD providers. SPN's are not deleted for source cluster and are not created for the target cluster. The failover log indicates success. This is due to a OneFS8 defect with multiple AD providers.
   NOTE: SPN delete / create for the AD provider defined in groupnet0 is successful. Workaround: Manually delete and create the SPN for the Smartconnect Zones that were moved from AD ADSI Edit interface.

(Release 1.8.3) Failure to run Resync Prep step during DFS Failover Deletes Shares on Target Cluster (T4145)

1. If during a DFS failover the Resync Prep does not run due to error prior to Resync Prep step or in the Resync Prep step itself, post failover Configuration Replication finds that the Eyeglass Job is still active on Failover source cluster and the replication of the renamed igls-dfs-<share> results in deletion of the <share> on the target cluster.
   a. **Workaround**: Prior to failover disable the Configuration Replication task. This does not affect the Configuration Replication step executed during failover.
   b. To disable the Eyeglass Configuration Replication task, execute the below command from the Eyeglass appliance command line:

i. igls admin schedules set --id Replication --enabled false

ii. Post successful failover, re-enable Eyeglass Configuration Replication task.

c. To disable the Eyeglass Configuration Replication task, execute the below command from the Eyeglass appliance command line:

i. igls admin schedules set --id Replication --enabled true

d. **Fixed in > 1.9.0 - any step fails the Jobs in eyeglass are left at user disabled state and will not run until manually enabled again. Ensuring SyncIQ policy issues can be recovered to correct state first and then user enable the policies in Eyeglass.**

(Releases < 1.9.0) DR Assistant returns "Error Retrieving Zones Undefined" if many access zones exist

1. This error can occur when attempting a failover when many access zones and many policies per access zone are configured. A database query times out return all data needed to validate the failover. This is addressed with optimized DB query in 1.9 release. The impact is inability to start a failover.
2. **Work Around: Increase timeout on browser to return all needed data from the database to start a failover.**
    a. For temporary fix, please follow the steps below and let us know the update via this case:
    b. SSH to the eyeglass appliance as admin user
    c. type password (default: 3y3gl4ss)
    d. sudo su - (default password: 3y3gl4ss)
    e. vi /srv/www/htdocs/eyeglass/js/eyeglass_globals.js
    f. please change the ajax_get_timeout_seconds value to 600.

g. Please refer the screenshot for details:

h. :wq! // save the changes //

i. login to the eyeglass webpage and open the DR assistant and check whether error still present or resolves. You may need to clear browser cache to ensure new java script is loaded to the browser that includes the new timeout.

j. Done.

(Releases => 1.9.0) DFS Failover Enhancement to handle partial or complete Share Rename failures

1. DFS mode uses parallel threads to rename shares for all policies involved in the failover.
2. If share renaming is failed for **all** the shares from a cluster, then failover status is **error**. Failover is stopped and Users are not redirected to target cluster. **Make writeable and Resync prep does not run and data is active on source cluster still.**
3. If share renaming is failed only for **some** shares from the source cluster, then failover status is **warning AND failover will continue to run make writeable and resync prep.**
4. **Summary:** In this scenario it is best to attempt the failover of some shares fail rename. If all fail abort failover and stop.

(All Releases) Access Zone Failover Networking Roll back on failures

1. This feature has been available for some time and should be understood how it works.

2. During make writeable step Eyeglass will send API to target cluster to start the make write able step.
3. At this point in the failover smartconnect networking and SPN failover has been completed and dual delegation will mean new mount requests will be handled by the target cluster and SPN authentication will be handled by the target cluster.
4. If the make writeable step Succeeds on at least ONE policy of N (of all policies involved in the Access zone), **the failover logic will continue. This means you are partially failed over for some of the data in the access zone. It also means all networking and SPN's are failed over. Next step is to resolve failed make write step on policies to get the file system writeable. This often requires EMC SR to resolve root cause of failover on the target cluster.**
5. If **NONE** of the policies pass the make writeable step **AUTOMATIC** rollback of Smartconnect networking and SPN's are reverted to the source cluster.
    a. The failover log shows if networking rollback is initiated. If you find this in the failover log, Your failover is aborted and all data remains writeable on the source cluster.
    b. Example Log entry **2017-06-08 22:49:00::260 INFO Starting Step: "Networking Updates Rollback"**
    c. To validate source cluster data access to the following:
        i. nslookup to the smartconnect name(s) involved in the failover (use failover log for full list). IP returned should be from the source cluster
        ii. Test share and NFS mount access to the source cluster and verify you can mount and write
        iii. This will validate SPN authentication for shares as well.
        iv. Determine root cause , which may require EMC SR to resolve before rescheduling the failover

6. **Summary: This logic determines the best option automatically. If some data succeeds to failover its best to resolve only the failed policies than aborting the entire failover. If no data succeeds at the make writeable step it is best to revert and abort the failover. Eyeglass handles this decision automatically.**

(All Releases) PowerScale OneFS 8.0.0.5 API load sharing with Smartconnect issue

1. Any issue with this oneFS release was found where API calls from eyeglass when clusters are added using FQDN smartconnect name, shows DFS mode share rename step uses parallel API calls to load shared across nodes results in HTTP 409 AEC error from the cluster when a share rename share fails.
2. The share is renamed correctly but the cluster does not remove the old share leaving the igls-dfs-sharename and sharename on the target cluster.
3. The HTTP 409 error is sent incorrectly by the cluster and Eyeglass treats this as a failed step, even though the rename was successful.
4. **Summary: Work around is to delete the cluster from eyeglass inventory window, re-add the cluster with subnet service IP to avoid this cluster bug. No known resolution for this issue at this time on OneFS. Impact of not switching to SSIP, is failed DFS failover when using FQDN cluster add with 8.0.0.5.**

(All Releases) Missing SPN Validations for Zone Readiness and Pool Readiness cause SPN create/delete to fail during failover

Zone Readiness and Pool Readiness SPN validations do not check for the conditions below.

IMPACT: These conditions will cause SPN delete/create to fail during a failover:

1) SPN has been created in AD with lower case host (example: host/SPN_name) instead of uppercase HOST (example: HOST/SPN_name)

2) SPN has been created in AD where SPN_name has different case than associated SmartConnect Zone name (example: for SmartConnectZone prod.example.com SPN is configured as HOST/Prod.Example.com)

Workaround: Modify SPN in AD that have above issues so that all SPNs

1) use upper case HOST in the SPN definition (HOST/SPN_name)

2) SPN name matches case of Smartconnect Zone name

(All Releases) User Quota creation fails on failover for multiple disjointed AD Domain environment

In an PowerScale environment that is configured to use multiple AD Domains and those Domains are not joined, user quota creation for the quotas related to the non-default AD Domain will fail with the error:

Requested persona was not of user or group type

Workaround: None available with Eyeglass.

(ALL RELEASES) Time to complete steps for Allow Writes and Preparation to Failback Unknown

Time to complete failover steps to make data writeable and prepare to failback (resync prep) can take a long time for some environments related to large number of files/directories and other factors and time is not predictable or deterministic.

(CURRENT RELEASE) Failover Known Issues

Failover related Known issues for the current release can be found here.

2. Failover Planning Guide and Checklist

Home Top

Failover Planning Guide and checklist

Introduction to this Guide

Overview

The Eyeglass PowerScale edition greatly simplifies DR with DFS. The solution allows DFS to maintain targets (UNC paths) that point at both source and destination clusters.

The failover and failback operations are initiated from Eyeglass and move configuration data to the writeable copy of the UNC target. Grouping of shares by SyncIQ policy allows Eyeglass to automatically protect shares added to the PowerScale. Quotas are also detected and protected automatically.

The following checklist will assist you in plan and test your configuration for Failover in the event DR (Disaster Recovery) is needed.

Chapter 2 - Checklist to plan for Failover

Steps <u>Before</u> failover Day	Task	Description	Completed
0	**Document DR Runbook plan**	• Organize steps, contacts, order of steps, contacts per step required on execution of failover day	
0A	**Submit support logs for failover readiness audit (7 days before planned event) (see image for case option to request assessment)** 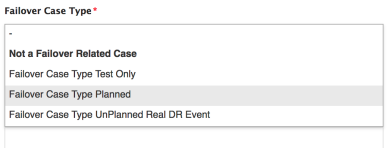	Failover Release Notes	
0B	Take failover training labs to practice execution	• https://www.supernaeyeglass.com/booking	
1A	**Review DR Design Best Practices** **Review Failover Release notes** **Warning: Mandatory Step for all customers DR Assistance requires acceptance before continuing**	• Eyeglass and PowerScale Failover Best Practices • Failover Release Notes	
1B	**Upgrade Eyeglass to latest version (Eyeglass releases includes failover rules engine updates that add rules found from other customer failovers that continuously improve or avoid known failover issues)** Failover Release Notes	• Eyeglass PowerScale Edition Upgrade Guide	
1C	**Test DR procedures**	• Setup Runbook robot feature for continuous DR testing • Test failover with Superna Eyeglass	

		• Test it again, again and again • Failback • Review results, logs to ensure steps that Superna Eyeglass executes are understood • Consulting documentation on failover mode you planned to implement • Execute test plan before failover day to validate procedures	
1D	**Benchmark Failover (access Zone)**	• Copy data into a test policy or the runbook robot access zone **(note Robot can only use 1 policy for testing, to complete multi policy testing a test access zone would need to be created and configured for access zone failover)** • Execute test failover and use failover log to find the make writable step time delta to the start of the log. This is the point at which failover is completed, and failback steps now begin to execute but clients are able to write data to target at this point. • Repeat above with 2 policies and a known quantity of data so that both policies sync data and failover. Record the make writable time difference log step to the beginning of the failover log time stamp • Repeat one more time with 3 policies same amount of data in each directory • Now average the 3 test run times to the make writeable step and use this value that is unique to your environment (clusters, WAN, nodes in replication, etc..) to use to calculate estimated	

		failover times if you have more than 3 policies. • **Note the test access zone should have all configuration completed (hints, spn, shares and exports and quotas) to ensure that the time estimates are as close to production configuration when estimating failover times.** • **Note: If change rate is expected to be zero before planned failover then skip step to create changed data before failover.** • **Note: The reason to create as many shares under each policy as in production is to get the time for the rename step to complete for each share, this step is parallel operation but should be benchmarked on your clusters** • **Note: failover logs include steps post failover to prepare for failback and complete audit of the clusters. The failover job time DOES NOT REPRESENT THE TIME IT TAKES TO FAILOVER. YOU MUST CALCULATE THE MAKE WRITABLE STEP IN THE LOGS**	
1E	**Benchmark Failover (DFS Mode)**	• Use the Access Zone with DFS mode policy or create a test DFS mode policy • Copy test data into path • Create one more shares into the path of test policy (if you have more than one share under a policy in production than create as many shares as you have in production policy configuration) • Create more than one policy as per above step example 3 to get	

		a good time average • Create changed data if you plan to failover with un-synced data (optional step) • Run DFS mode failover on 1 policy, then 2 then 3. Record the make writeable step time difference to the start of the failover log. • Calculate the average time per policy (based on your production configuration) • Use this number to estimate the time to complete your production failover times • **Note: The reason to create as many shares under each policy as in production is to get the time for the rename step to complete for each share, this step is parallel operation but should be benchmarked on your clusters** • **Note: failover logs include steps post failover to prepare for failback and complete audit of the clusters. The failover job time DOES NOT REPRESENT THE TIME IT TAKES TO FAILOVER. YOU MUST CALCULATE THE MAKE WRITABLE STEP IN THE LOGS**	
2	**Contact list for failover day**	• AD administrator • DNS administrator • Cluster storage Administrator • workstation, server administrators • Application team for dependant applications • Change Management case entered for outage window	

3	**Reduce failover and failback** time - Run manual domain mark jobs on all syncIQ policy paths (this will speed up failover because domain mark can take a long time to complete and elongates the failover time)	All policies run this procedure on all policies. Domain mark	
4	Count shares,exports, NFS alias, quotas on source and target with OneFS UI	Validates approximate config count is synced correctly (also verify Superna Eyeglass DR Dashboard) (there should be no quotas synced on target - only shares, exports and NFS alias)	
5	Verify dual delegation in DNS before failover	This verifies that DNS is pre-configured for failover for all Smartconnect Zones that will be failed over (Access Zone failover fails over all Smartconnect Zones on all IP pools in the Access Zone)	
6	DFS failover preparation	1. using dfsutil verify clients that will be failing over show two active paths to storage and that correct path is active 2. Verify all DFS mounts have both referrals configured dfsutil tool downloaded by OS type check path resolution 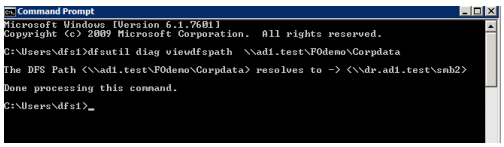 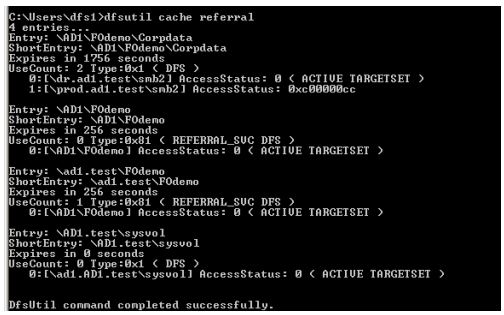 	

7	Communicate to application teams and business units that use the cluster the failover outage impact	1. Scheduled maintenance window with application and business units 2. Ensure to explain that data loss will occur if data is written passed the maintenance window start time	
8	Set all policies schedule to every 15 minutes or less 1 day prior to the failover to ensure data is staying in sync. This also ensures the failover speed will be optimized	This step is critical step to changne to avoid long running policies or long running jobs that will extend your failover and maintenance window. Specifically ensure run on change is never left enabled since policies that are running cannot be controled for failover.	
Steps on the failover Day	**Task**	**Description**	**Completed**
0	**SMB and NFS IO paused or stopped before failover start to avoid data loss**	For SMB protocol the 2.0 or later feature can be used to block IO to shares with DR assistant. This inserts a deny read permission dynamically before failover starts and removes after failover completes. 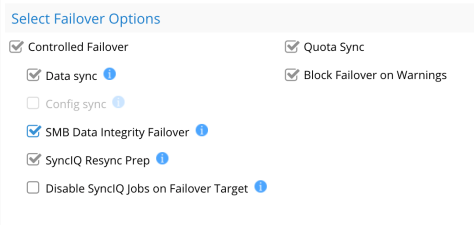 NFS requires the protocol to be disable to guarantee no IO. Exports should be unmounted before disabling the the protocol on the cluster.	
1	**Force run syncIQ policies 1 hour before planned failover**	Run each syncIQ policy before so that the failover policy run will less data to sync	

2	**Execute failover**	How to Execute A Failover with DR Assistant	
3	**Monitor failover**	How to Monitor the Eyeglass Assisted Failover	
4	If Required Data recovery guide	Failover Recovery Procedures	
5	Ensure Active Directory admin is available	ADSIedit recovery steps are required and needs Active Directory Administrator access to cluster machine accounts	
After Failover	Test Data Access	Use post failover steps guided steps How to Validate and troubleshoot A Successful Failover WHEN Data is NOT Accessible on the Target Cluster	

Planning Check List Excel Download

1. Superna Eyeglass Failover Planning Checklist

3. Failover Process and Customer Role

Home Top

Overview

This document describes:

- Roles and Responsibilities during Failover with Superna Eyeglass
- Day of Failover Support Process
- Superna Eyeglass Failover Planning Process - requires 7 days notice if you plan to use support services to validate your environment. If you chose to open a case with less notice, all validations covered by this process may not be completed in time for your failover. Customer must accept the risks of some validations or remediations not being completed in time.

Roles and Responsibilities during Failover with Superna Eyeglass

Superna Eyeglass Support Role in Failover

Superna Eyeglass Superna Eyeglass EULA and Support Services Agreement includes support for the Superna Eyeglass product. As related to failover this includes:

- Failover planning process including readiness health check and remediation prior to a planned event
- Failover log analysis **during** and post failover health check and failback assessment
- Root cause of issues during a failover
- Next steps required to complete a failover under all conditions, all recovery steps are documented.

Superna Eyeglass product support entitlement does not include

1. **Assisted failover of customer data.**
    a. Professional services from 3rd party Eyeglass Certified partners offer hands on failover services. **(**Eyeglass Professional Services**)**
2. **Decisions on data protection before, during and post failover.**
    a. These decisions reside with customers. Superna Eyeglass Support can not be responsible for legal reasons.
3. **Assistance with Data migration using the Data and Configuration tools in Eyeglass.**
    a. Support can answer questions on use cases, limitations, recommendations
    b. Support will not join a zoom or webex to assist with data migration and is consider a Data protection decision by customers to execute all steps.
    c. We suggest testing all options before attempting any data migrations
4. **Support for any external hardware and software vendors is excluded from support**
    a. Support agreements must be in place for all external hardware and software vendors. Functional recovery steps will be provided that requires customer subject matter experts to execute on the 3rd party vendor products.
    b. Customers must have access to support for (AD, DNS, Networking, Hosts, PowerScale), Superna Eyeglass support

can provide root cause of external component issues but is not primary support replacement for these components.

c. Superna Eyeglass Support is **legally not authorized** to take **control** of any **devices** or make business **decisions** on behalf of customers during a failover event **Or** provide specific technical advice that affects a 3rd party vendor where that vendor should be consulted.
d. Superna support cannot provide assistance that would violate customer support agreements with 3rd party vendors.
e. Superna Eyeglass Support can not take control or join a Webex or phone call to assist with hands on failover procedures or troubleshooting for 3rd party hardware and software **for legal reason**

Customer Expectations and Role in Failover

- Customers must provide or have access to all skills required to complete a failover and debug any issues in their IT environment which includes (AD, DNS resolution and updates if necessary, AD domain edit permissions to computer objects with ADSI Edit, PowerScale knowledge on SyncIQ operations, share/export management, Networking, firewalls, Windows logon process, Linux mount requirements, application specific knowledge that uses NAS shares).
- Customers must be logged in to the support.superna.net portal for purpose of uploading failover logs and communicating with Superna

Support team on any questions or issues that may arise during the failover.

How to Open a Failover support Case and Set the Correct Type

When a case is opened 4 choices are available:

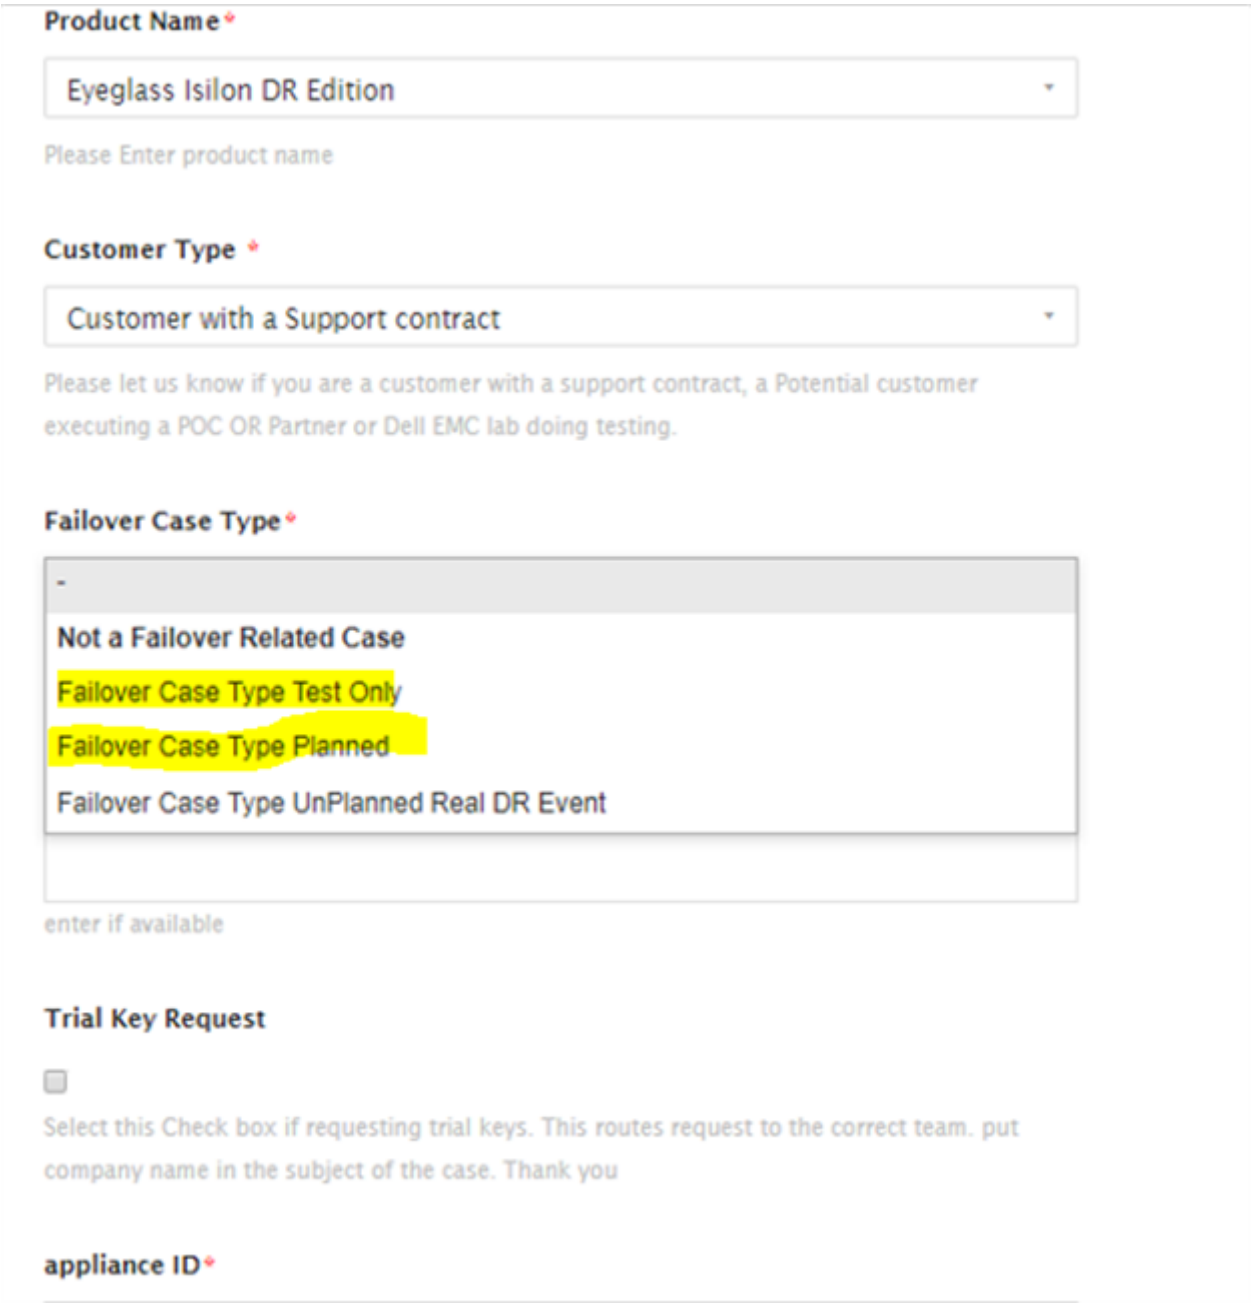

1. Not a failover case

    a. will not be treated as failover case if selected

2. Failover Case Type Test Only

    a. This is assumed to be none production data with no business impact.

    b. This will lower priority case if any higher priority cases existing depending on active case workload at the time the case is opened.

    c. Support will **automatically assume the failover health check support process** for this case type with suggested 7 days notice to allow the full process to be completed.

    d. You may opt out of this process. Please indicate this to support when opening the case to ensure communication and understanding on support.

3. Failover Case Type Planned

    a. This is assumed to be production data failover for business continuity testing.

    b. Support assumes this is planned and scheduled event and not a last minute decision to failover and you have a full planning phase for the event.

    c. Support will **automatically assume the failover support process will be used to health check your environment** and remidate any issues prior to the planned event. This process

works best if at least 7 days notice is provided. This advanced notice is based on our experience with many failovers and allows time to address many of the known issues that could impact your failover. **This process is designed to significantly reduce your risks for the planned event.**

d. **You may opt out of this process. Please indicate this to support when opening the case to ensure communication and understanding on support and risks you will assume as a result of opting out of this process that is included in your support contract.**

e. **If you are not running the latest target GA code then, please read all release notes for the Eyeglass release you are running, all these risks and known issues are assumed to be accepted and understood by you (release notes)**

f. **All process steps outlined on this page will be assumed to be understood and used for the planned event. You should schedule a meeting if you have questions on this process. We will happily review the details on a call to ensure both support and your organization are aligned on expectations and responsibilities before you execute a failover.**

4. Failover Case Type Unplanned Real DR Event

a. This assumes it affects production data.

b. Support will prioritize this case type above all others

c. The DR assistant controlled check box should **UNCHECKED** only if the source cluster shows as unreachable in the **Continous Operations Dashboard Icon**

d. Providing the failover log is a **mandatory** step to get support for an uncontrolled or Real DR event failover.

e. This was not a planned event.

f. **This case type should not be used for testing.** (for testing of DR events use this supported procedure (any other procedure is not supported and cases opened as DR event that are in fact a test will be switched to case type test).

g. The process on this page should be referenced so that support **exclusions** are understood and clear as it relates to the support contract and domain knowledge related to 3rd parties example AD, DNS, PowerScale, Windows, Linux must use experts in your organization or support contracts as required to follow steps and directions provided by Superna support throughout the process. Superna support is not a substitue for the knowledge and skils releated to 3rd party software and hardware.

h. **As stated on this page assistance or decisions related to your business data is not available with product support and decisions and execution of DR failover resides with customers IT staff.**

Day of Failover Support Process

Please find below steps for day of failover. **NOTE: Support will follow this process below exactly as written, this is fastest process to complete a failover.**

How to Receive the fastest Failover support during a failover

The following is the fastest process to get timely support. Steps below are based on 10,000 plus failovers executed globally and any deviation will negatively affect response time.

1. Update the case when you are about to start your failover so that the support engineer can expect the failover log soon.
2. Remain logged into our support portal with the case open on a browser tab and copy and paste the failover log from the running failover tab of the DR Assistant. This provides support realtime notification. **Do NOT use email to provide status.**
    a. If you want **feedback** on anything listed in the failover log, copy and paste the whole log to the case **EVEN** if the failover is not finished. We can provide feedback during the failover if you have questions.
    b. **DURING FAILOVER: Release > 2.5 will post a message to the failover log indicating when each policy is failed over and data access testing can be begin. Monitor the log and start testing as soon as the message is seen per policy OR post partial log to the case for review and confirmation.**
    Example:
    Failover steps for policy: <policy Name> completed in X minute(s). Data access manual testing should be completed, using this guide as a reference How to

Validate and troubleshoot A Successful Failover WHEN Data is NOT Accessible on the Target Cluster.

3. When asked by support to upload the completed failover log do this **immediately** (How to download failover log). This file is **mandatory** to receive support on **ANY** question you may have.
4. Create a full backup (how to here) and upload it **immediately after the failover. Our support system has automated failover debugging and can anaylze > 300 issues in minutes. We will not join a webex since our support engineer cannot review the debug analysis while on a webex. It is not technically possible to provide a faster response or analyze an issue without this support log analysis, delaying the upload to support WILL delay your support response to ANY question.**
5. **If the support engineer determines a rapid response is required, we have the ability to open a chat window to provide a faster response or directions and allows us to continue to review failover analysis provided by our support tools. This chat window feature is only supported if you are logged into our support website.**

Objective:

- Verify access to data on the target cluster post failover is supports priority in accessing failover logs

- Once data access validated and you have provided confirmation of data access to the case.. Support will then move on to assess failback steps if any errors in failover steps.

1. **BEFORE FAILOVER:** Onefs known bug with quota scan. Follow these Failover Release Notes before failover
2. **DURING FAILOVER**: Release > 2.5 will post a message to the failover log indicating when each policy is failed over and data access testing can be begin. Monitor the log and start testing as soon as the message is seen per policy. Failover steps for policy: <policy Name> completed in X minute(s).
3. **AFTER FAILOVER:** Attach failover log from eyeglass as a reply to this email.
4. **AFTER FAILOVER:** Test data access by following these instructions How to Validate and troubleshoot A Successful Failover WHEN Data is NOT Accessible on the Target Cluster
5. **AFTER FAILOVER:** Upload full backup from eyeglass appliance.
6. **AFTER FAILOVER:** Reply to this email with results from data access testing.

If you are planning to failback on same day, could you please upload a new set of eyeglass full backup (Create Full backup) logs by following this procedure to access the failback readiness: This is mandatory.

How to Raise an Eyeglass Support Request

Superna Eyeglass Failover Planning Process

The Failover Planning process is a series of steps designed to eliminate known risks during a failover event. It includes planning steps to be completed by the customer and a Superna Eyeglass failover readiness

health check and remediation prior to a planned event completed by Superna Eyeglass Support.

Step 1 in the process is to open a support case notifying the Superna Support team of the planned failover. The Superna Eyeglass Support team will post the Failover planning checklist to the case describing planning process and next steps.

Provide the date, time and time zone of a planned failover.

If this information is provided, we ensure a **dedicated** support engineer is scheduled for your failover to review support logs and information as it is posted to the case. If this information is **not** provided we cannot plan to support the failover with a dedicated support resource and normal support agreement response times apply.

Superna Eyeglass Support requires 7 days notice to allow this process to be followed by opening a case. Customers that do not follow the process are accepting all risks from release notes and known issues the planning process is designed to eliminate. Support level is reduced for customers that do not follow documented procedures to eliminate known risks.

Failover Planning Checklist

For a planned failover please review the following to maintain support for your installation. We provide a Failover Planning Guide and Checklist. We provide a complete planning process and customer role.

1. **STEP 1 We require written acknowledgement of the failover mandatory steps below to be posted to the case. This process ensures you have read and understood all steps for failover to ensure a successful failover and ownership of the failover process resides with your company.**
2. **STEP 2: Submit support logs for review of your environment (NOTE: Until step 1 is complete. support is unable to review logs.) NOTE: The readiness check only includes items listed at the end of this message.**
3. **STEP 3: Support will determine next steps based on review and post to the case.**
4. **NOTE1: Customers must have access to support for (AD, DNS, Networking, Hosts, PowerScale), Eyeglass support can provide root cause of external component issues but is not primary support replacement for these components.**
5. **NOTE2: support can not take control or join a Webex to assist with hands on failover procedures for 3rd party hardware and software for legal reasons.**
6. **NOTE3: Support agreements must be in place for all external hardware and software vendors. Functional recovery steps will be provided that requires customer subject matter experts to execute on the 3rd party vendor products.**
7. **NOTE4: Decisions on data protection before, during and post failover reside with customers. Support can not be responsible**
8. **NOTE5: For a permanent link to policies and support expecta tions for planned failovers refer to this Failover Process and Customer Role.**

Technical details for failover planning are provided below. We would like to schedule a 30 minute webex to review these failover planning mandatory steps as well as the day of failover support process and answer any questions you might have. Please provide us with your availability so that we can schedule.

PART 1 - Per Failover Acknowledgement Required

Mandatory steps in this section need to be acknowledged for **every** failover.

0) What is covered by readiness check

As per your request to review readiness for failover, we will review the following for failover readiness validation:
- Access Zone Failover hints
- Access Zone Readiness status / DFS Readiness status / Policy Readiness Status
- SPN errors
- Eyeglass misconfigurations (example Eyeglass version)
- SyncIQ Domain Mark
- Planned Failover Type
- Client Redirection Guide Followed (DNS Dual delegation/DFS Dual Delegation)
- Eyeglass appliance health check
- Cluster API response health check

Anything other than the above is not checked. Anything else not listed needs to be verified by yourselves in the context of your overall failover plan.

Acknowledged: Yes/No

1-1) Mandatory - Are you using Superna Eyeglass Easy Auditor? yes or no

1-2) Mandatory - Are you using Superna Eyeglass Ransomware Defender? yes or no

1-3) Mandatory - How many Access zones or policies will be failed over?

1-3A) Mandatory - Do you mount shares with DNS CNAME's?
If yes you must manually failover DNS CNAME SPN's or best practice create IP pool alias = to the CNAME to ensure SPN's are failed over by Eyeglass

1-4) Mandatory - Are you failing over and failback on the same day? yes or no

1-5) Mandatory - How long a maintenance window have you scheduled?

Note: If you are failing over less than or equal to 25 SyncIQ policies we recommend 4 hours of maintenance window but if greater than 25 SyncIQ policies then we recommend 6 hours of maintenance window.

1-6) Mandatory - Please create new RPO reports from "Reports On Demand" window and attach to the case.

This will help us to determine average time for SyncIQ policies to complete. Please use GIF animated procedure to create RPO reports: How to create RPO report using Reports on Demand

1-7) Mandatory Step: If also planning OneFS 7 to 8 upgrade review and execute this procedure after upgrade to OneFS 8 for ANY cluster PowerScale Upgrade Procedure with Eyeglass

Acknowledged: Yes/No or NA

1-8) Mandatory Step: Please upload a full Eyeglass Backup 7 days prior to failover for our review of your failover readiness.

Please upload a new set of eyeglass full backup logs (Create Full backup) by following this procedure:
How to Raise an Eyeglass Support Request

Acknowledged: Yes/No

PART 2 - One Time Acknowledgement Required

Mandatory steps in this section only need to be acknowledged once and therefore do not need to be re-acknowledged for every failover. Once Acknowledged support will assume these steps below are done for all failovers and all risks associated with these validations are accepted.

2-1) Mandatory Step: Use the latest Eyeglass release for planned failover or you are automatically accepting risks of using older release. N-2 release accepted under conditions covered in the Failover Release Notes: Failover Release Notes

Acknowledged: Yes/No

What about unplanned DR event?

Unplanned (real dr with source cluster destroyed) no it's not required to be on latest release. However, we always expect the software is upgraded to supported releases published here Software Releases. We provide *notifications to @eyeglassPowerScale or email distribution*

2-2) Mandatory Step: Review the Failover release notes and assess all risks to your environment have been assessed. If you have questions, now is the time to ask.

Acknowledged: Yes/No
>> We provide notice when a release has support removed and the requirement is to upgrade. We improve failover from all cluster failovers from all customers each release so without upgrading to latest release you don't benefit from other customers failovers.
>> We send all Eyeglass release notifications to all registered emails in the support portal.

2-3) Mandatory Step: Complete the failover planning checklist.

This document that has links to all key documents that you need to review and steps to execute before running your DR test:
Failover Planning Guide and Checklist
Acknowledged: Yes/No

2-4) Mandatory Step: Prepare a contact list in place for failover day.

Acknowledged: Yes/No

Customers must provide or have access to all skills required to complete a failover and debug any issues in their IT environment which includes (AD, DNS resolution and updates if necessary, AD domain edit permissions to computer objects with ADSI Edit, PowerScale knowledge on SyncIQ operations, share/export management, Networking, firewalls, Windows logon process, Linux mount requirements, application specific knowledge that uses NAS shares)

Superna Eyeglass Support agreement is for the Superna Eyeglass product and is not a replacement for skill or support agreement for all external hardware and software vendors.

2-5) Mandatory Step: Plan to stop IO to the source cluster before failover. This is a failover best practice.

Acknowledged: Yes/No

If yes, for SMB protocol the 2.0 or later feature can be used to block IO to shares with DR assistant. This inserts a deny read permission dynamically before failover starts and removes after failover completes.

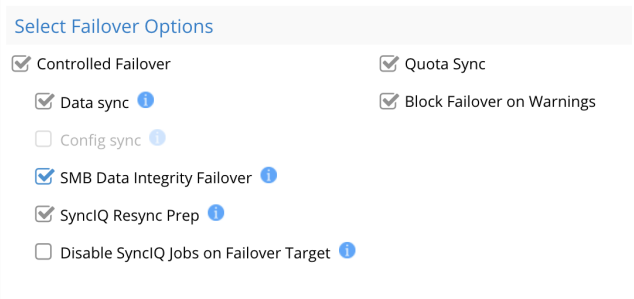

NFS requires the protocol to be disabled to guarantee no IO. Exports should be unmounted before disabling the protocol on the cluster.

If no, you will incur data loss.

2-6) Mandatory Step: Run domain mark jobs on all SyncIQ policies prior to failover.

Acknowledged: Yes/No

Consult explanation Eyeglass and PowerScale Failover Best Practices.

2-7) Mandatory Step: Review Operational steps and procedures for the day of failover

Acknowledged: Yes/No

1. Operational Steps for Eyeglass Assisted Access Zone Failover
2.
    a. How to Monitor the Eyeglass Assisted Failover
    b. How to Validate and troubleshoot A Successful Failover WHEN Data is NOT Accessible on the Target Cluster
    c. Post Access Zone Failover Steps
    d. Post Access Zone Failover Checklist
    e. Troubleshooting Failover
3. Operational Steps for Eyeglass Microsoft DFS Mode Failover
4.
    a. How to Monitor the Eyeglass Assisted Failover
    b. How to Validate and troubleshoot A Successful Failover WHEN Data is NOT Accessible on the Target Cluster
    c. Post Eyeglass Microsoft DFS Mode Failover Checklist

2-8) Mandatory Step: Review failover Recovery document

1. Failover Recovery Procedures

Acknowledged: Yes/No

2-9) Mandatory - Follow supported procedure for an uncontrolled failover where you are simulating source cluster unavailable

If YES, **the ONLY supported procedure** using Superna Eyeglass is documented here: Eyeglass Simulated Disaster Event Test Procedure.

ANY change to the documented procedure is NOT supported.

Recovery from a real uncontrolled failover is customer responsibility and is NOT covered by Superna Support. Refer to the Failover Design Guide - How to Execute A Failover with DR Assistant for more information on execution and responsibility for recovery on a real uncontrolled failover.

2-10) Mandatory - For multi-phase failover upload Eyeglass full backup between failovers for assessment of environment to reduce risk for subsequent failovers. Support level is reduced without backup being provided.

Acknowledged: Yes/No

4. Eyeglass and PowerScale Failover Best Practices

Home Top

- IMPORTANT READ this --- All Planned Failover Attempts MUST read this support statement
- IMPORTANT READ this --- Do not attempt failover without completing this step. Best Practise for Fast Failback and Pre Failover Steps
- IMPORTANT READ this --- Failover timeouts with Eyeglass - Cluster Operations that can take longer than planned
- Best Practice General:
- SyncIQ Performance Tuning Best Practices
- Data Loss Considerations
- Best practices for DFS mode Failover Design:
- Best practices for Access Zone and per SyncIQ mode Failover Design
- Best Practise DNS Configuration for Access Zone Failover
- Best Practice for Protecting Data for HA and Failover with Eyeglass
- Best Practice for PowerScale Networking
- Best Practice for Kerberos Service Principal Names (SPN's)
- Best practice to verify the following on all DNS

- Best practice DNS delegation of NS records
- Best practice - Do DR Testing with RunBook Robot for Access Zones

This section is a collection of best practices. Details on configuration is in the admin guide. This is section is aimed at quick short descriptions of best practices in one easy to read place, that covers Eyeglass and SyncIQ.

IMPORTANT READ this --- All Planned Failover Attempts MUST read this support statement

1. Planned failovers must use the latest software available. Each release has fixes, improvements and new error conditions blocked or warned that can prevent issues or robuts failover.
2. Always plan to upgrade appliance software as step before any planned failover.
3. This is a supportability **requirement** for customers and expected as basic step in keeping DR software updated as key component for planning and readiness.
4. This requirement **supersedes** any change management or IT policies that require upgrades to be planned, this must be factored into any planned failover. Trial keys are available for lab systems as are PowerScale Simulators for testing upgrades in advance of a planned failover event.
5. EULA requires customers to maintain current updated software as condition of the license when raising failover support cases.

IMPORTANT READ this --- Do not attempt failover without completing this step. Best Practise for Fast Failback and Pre Failover Steps

1. Run domain mark manually on all SyncIQ paths following instructions in online PowerScale documentation. Read this to

understand why its important to run it now (see section Prepare policy for accelerated failback performance)

Create a SyncIQ domain

1. You can create a SyncIQ domain to increase the speed at which failback is performed for a replication policy.
2. Failing back a replication policy requires that a SyncIQ domain be created for the source directory. OneFS automatically creates a SyncIQ domain during the failback process. However, if you intend on failing back a replication policy, it is recommended that you create a SyncIQ domain for the source directory of the replication policy while the directory is empty.

Create a protection domain Procedures

You can create replication or snapshot revert domains to facilitate snapshot revert and failover operations. You cannot create a SmartLock domain. OneFS automatically creates a SmartLock domain when you create a SmartLock directory.

1. Click **Cluster Management** > **Job Operations** > **Job Types**
2. In the Job Types area, in the DomainMark row, from the Actions column, select **Start Job**.
3. In the **Domain Root Path** field, type the path of the directory you want to create a protection domain for.
4. From the **Type of domain** list, specify the type of domain you want to create.
5. Ensure that the **Delete this domain** check box is cleared.
6. Click **Start Job**.
7. Confirm completed step
    1. Run this on source cluster isi_classic domain list
    2. Output should show SyncIQ domain on each syncIQ policy that has been created if you have successfully run domain mark on all policies
    3. If any paths are missing repeat step 4

IMPORTANT READ this --- Failover timeouts with Eyeglass - Cluster Operations that can take longer than planned

The following section is very important to review, If you have never failed over a policy than you have never run domain mark which eyeglass and PowerScale require to run domain mark job on the source cluster before failover. The following conditions WILL increase the time to run cluster operations and if you have policies that match this criteria then increase the timeout for Eyeglass failover jobs.
Policies criteria for increased timeout:

1. **Many TB of data protected by Single SyncIQ policy (many is not precise but if you think it's a lot of data for your environment then this applies to you)**
2. **Many small files (same as above if you know it has a lot then it likely does and this applies to you)**
3. **You have daily schedules for SyncIQ AND you have high change rate in GB's per day and policies take over 1 hour to run normally each day**

If you have policies as per above AND you have run domain mark in advance of a failover as recommended above as a MUST DO. Domain mark can take hours so read and please do this before failover.
When Eyeglass starts and cluster task (example start resync prep, run policy, even make writeable for policies that match the criteria above). Then the per task time should be increased. From the default of 180 minutes to a number greater than 180 minutes based on looking RPO graph or report of the policy you are planning to failover. Do this before attempting a failover or failback of a policy that matches the above criteria

How to change the timeout
igls adv failovertimeout set --minutes 360

Best Practice General:

1. Eyeglass - We recommend DFS mode for SMB share protection and DR

2. Eyeglass - We recommend Access Zone Failover when NFS and SMB data needs to failover together
3. Eyeglass - We recommend syncIQ policy mode failover for customers with small numbers of NFS exports and hosts for automation
4. Eyeglass We recommend Access zone when multi protocol SMB/NFS is required within a single Access zone OR when only NFS DR protection is required
5. Eyeglass NFS only failover - Use simpler per policy Failover with Eyeglass and unmount remount new DR Smartconnect zone name. It's faster and requires less planning and configuration than Access Zone Failover
6. Eyeglass Multi-protocol failover allows both protocols to failover together using Access Zone failover
7. PowerScale - For a syncIQ best practise for System level recovery you can refer to EMC document (PowerScale - Backup and recovery guide) https://www.emc.com/collateral/TechnicalDocument/docu56055_onefs-backup-recovery-guide-7-2.pdf
8. Eyeglass - Create smartconnect mapping alias hints on all ip subnet pools, hint the syncIQ smartconnect zone with ignore to ensure it's not failed over
9. Eyeglass - Delegate machine account credentials to cluster machine accounts in Active Directory
10. Eyeglass - Enable phone home support for faster support response times
11. Eyeglass - Configure Run Book Robot Access Zone and policies to ensure failover and failback is functioning daily
12. PowerScale - Always use FQDN on Smartconnect zone names
13. PowerScale - Create a SyncIQ Failback Domain to ensure fail back operations take less time
    1. Create a SyncIQ domain You can create a SyncIQ domain to increase the speed at which failback is performed for a replication policy. Because you can fail back only synchronization policies, it is not necessary to create SyncIQ domains for copy policies.
    2. Failing back a replication policy requires that a SyncIQ domain be created for the source directory. OneFS automatically creates a SyncIQ domain during the failback process. However, if you intend on failing back a replication policy, it is recommended that you create a SyncIQ domain for the source directory of the replication policy while the directory is empty. Creating a domain for a directory that contains less data takes less time.
    3. Procedure 1. Click Cluster Management > Job Operations > Job Types. 2. In the Job Types area, in the DomainMark row, from the Actions column,

select Start Job. 3. In the Domain Root Path field, type the path of a source directory of a replication policy. 4. From the Type of domain list, select SyncIQ. 5. Ensure that the Delete domain check box is cleared. 6. Click Start Job.

14. PowerScale - Create an IP and smartconnect pool that is only used for SyncIQ and create policies with run policy only on nodes subnet IP Pool/Smartconnect zone.
    1. Select option to Connect to nodes in the target smartconnect zone when creating policies
15. PowerScale - Don't mount data using the SyncIQ smartconnect zone, use other IP pools and smartconnect zones for users to mount data

This section covers key topics to review before planning DR with Eyeglass

SyncIQ Performance Tuning Best Practices

Consult the document below to turn SyncIQ job worker threads per node for high latency WAN and faster SyncIQ node operations (Syncing, make writeable, resync prep steps).
OneFS 7 and 8 are both covered in the document below.
https://www.emc.com/collateral/hardware/white-papers/h8224-replication-PowerScale-synciq-wp.pdf

Data Loss Considerations

When SyncIQ is set to a schedule or on changes mode it’s important to understand the impact to data loss on failover operations.

1. When a SyncIQ job is running and Eyeglass failover job is started the default behaviour will attempt to start a final data sync by running the SyncIQ policies in the job.
    1. If there is an existing SyncIQ Job running, Eyeglass failover will wait a maximum of 1 hour for the running SyncIQ Policy job to complete.
    2. For Urgent Failover requirements skip config sync and data sync option in the DR assistant UI by unselecting.
    3. If SyncIQ Job has not completed with an hour, an error is returned and the failover is aborted.
    4. Data Loss impact - Since SyncIQ is snapshot based, changes that have occurred since the start of the existing running job will be lost. Depending on the start time of the currently running job, this could represent a large amount of data

1. Mitigate Data Loss - Login to PowerScale to verify whether a SyncIQ Job is running for the policies being failed over.

1. Steps: Wait for the running job to complete and then start the failover. You may also consider disconnecting client access at this point to ensure that there is not a large amount of data that requires replication during SyncIQ Job run by the failover.
2. Set the SyncIQ Job schedule to manual before starting a failover. Eyeglass will run the SyncIQ policy as part of the failover procedure.

Best practices for DFS mode Failover Design:

1. DO Use Domain based DFS roots
2. DO Use DFS referral ordered list to select production UNC path as default first in the list to speed up referral processing and mount times
3. DO Use UNC path targets that point to SmartConnect zones
4. DO Name SmartConnect zones differently on source and target clusters so that debugging with dfsutil.exe is easier and smartconnect can load the cluster nodes during normal operations and after with failover
5. DO Group one or more SyncIQ policies by name and enable DFS mode in Eyeglass to failover related SyncIQ policies with DFS. (No hard rule requires this but it's easier to manage groups of related DFS failover if the names have similar prefix)
6. DO Create dedicated IP pools on source and target clusters for DFS protected data
    1. Within an Access Zone, create igls-ignore hints to ensure smartconnect zones are not failed over with Access Zone failover

Best practices for Access Zone and per SyncIQ mode Failover Design

Sub access Zone means a syncIQ policy within an access zone is used for failover of the data protected by the policy. This is supported but has limitations in amount of automation possible with this option.

1. Don't attempt Failover of a single SyncIQ policy within an Access zone unless you are prepared for manual steps below.
    1. There is no method to map a SyncIQ policy to a SmartConnect zone used by clients to mount the data. Incorrect configuration, or failing over a SmartConnect zone using an alias could impact other clients using the SmartConnect zone. Eyeglass can not failover SmartConnect zones without risk of causing inaccessible data on the production cluster unless ALL Smartconnect Zones are failed over to the target cluster.

    2. The storage admin is responsible to failover the SmartConnect zone manually in this scenario
    3. The SPN delete of the access zone and creation on the target cluster is also a manual step the storage admin must execute using ISI commands.
2. Do configure Access zone failover and design DR to failover all policies and SmartConnect zones in the access zone
3. Do all SyncIQ policies to be at the same level as the Access Zone base path or lower in the file system
4. Do create shares or exports underneath the path of SyncIQ policies to ensure they are automatically protected as well.
5. Do setup subnet:pool mappings for Access Zone failover using hints to map pools
6. Do setup Runbook Robot Advanced with Access zone configuration and verify it succeeds before attempting an Access zone failover
7. Do Use DFS mode for SMB within an Access Zone Failover Multi Protocol design
8. Don't Failover with Eyeglass per SyncIQ level failover unless you understand the limitations below.
    1. To allow partial, single SyncIQ policy(s) within an Access Zone the following constraints apply:
    2. Any smartconnect zones used are assumed to be manually failed over with aliases and DNS updates to point DNS at target cluster smartconnect ip address
    3. AD SPN creation on target and deletion on source cluster is manual, since Eyeglass does not know which smartconnect zones and SPN's are required on the source cluster after a policy is failed over leaving data accessible on the source cluster

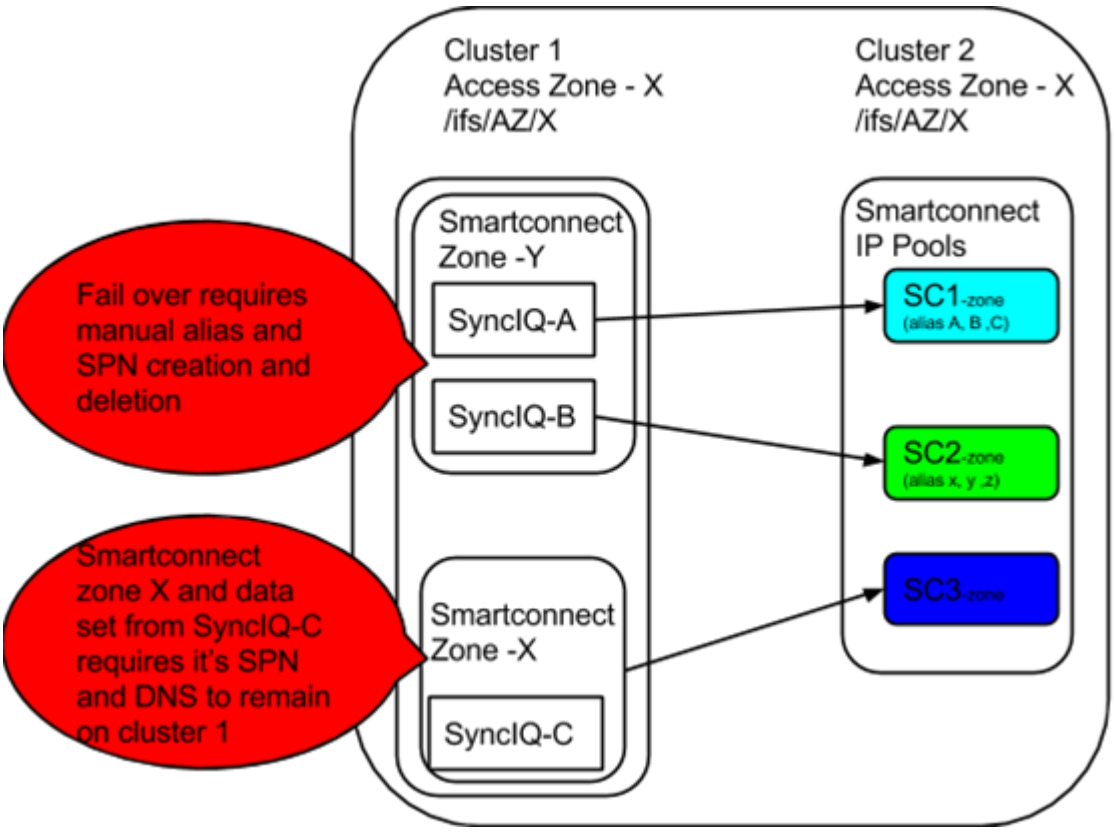

Best Practise DNS Configuration for Access Zone Failover

1. DNS that delegates NS records to Smartconnect Zones are the last step in the failover process to point the the failover Smartconnect Service IP on the target cluster (typically at the DR site).
2. This NS record is setup to point at the SSIP of the production cluster for the Smartconnect Zones within the Access Zone that will be failed over.
3. SmartConnect Zone aliases will also have NS records to delegate the alias entries as well to the SmartConnect Zone SSIP that has the alias assigned.
4. Delegation should use an A record for each SSIP but the Delegation for the NS should use a CNAME that points to the A record. This is best practice and simplifies the update on failover of the CNAME to point at the DR cluster SSIP A record

Best Practice for Protecting Data for HA and Failover with Eyeglass

1. DO - Organize Data into Protocol failover policies example policies for SMB and policies for NFS to take advantage of DFS mode
2. DO - Organize Data / SyncIQ Policies / Shares / Exports / Aliases / Quotas by Zone for failover

3. DO - Shares/Exports/Alias should be grouped into Zones based on which data sets that need to be failed over together.
4. DO - Map each subnet/pool clients use to access data to a target cluster subnet\pool using Eyeglass hint aliases
5. DON'T - Put SyncIQ policies at a level above the Access Zone root directory
6. DON'T - Use excludes and includes in your SyncIQ Policy. Excluded directory will be read-only after failover. For DFS mode, share on source cluster related to excluded path is not preserved

Best Practice for PowerScale Networking

1. It's best to use fewer ip pools to simplify DNS, Alias creation on failover and reduce updates to DNS required for failover.
    1. Example:
        1. SmartConnect Zone for Data
        2. SmartConnect Zone for SyncIQ
        3. SmartConnect Zone for management (Eyeglass and other applications)
        4. SmartConnect Zone for Backup

Best Practice for Kerberos Service Principal Names (SPN's)

1. Use Eyeglass DFS mode to limit kerberos authentication issues for cluster machine accounts
2. If NTLM fallback is disabled OR Microsoft patches or new OS's disable NTLM fallback, you don't want your DR strategy depending on authentication fallback to a legacy protocol. It's best to ensure SPN's are accurate for Kerberos authentication and use Access Zone failover as the unit of failover.

Best practice to verify the following on all DNS

1.

    a. To prevent giving out stale DNS entries, the DNS time-to-live (TTL) on the NS delegations should be set to zero, or as close to zero as possible, so that the DNS information is as fresh as possible.

    b. Certain clients perform DNS caching and might not connect to the node with the lowest load if they make multiple connections within the lifetime of the cached address.

c. Do not create reverse DNS entries, also known as pointer (PTR) records, for PowerScale SmartConnect service IP addresses or SmartConnect zone names. SmartConnect does not provide reverse lookups. **OR see #4 below as alternative.**

d. DO If A Records are used for PowerScale node IP's and SSIP's. Make sure forward and reverse lookups match example nslookup ip x returns host name Y and nslookup of y returns IP X. This is required to ensure TLS connections function correctly, since TLS will validate ip to name and name to ip address to protect against man in the middle attacks to TLS connections.

Best practice DNS delegation of NS records

This section describes best practices for DNS delegation for PowerScale clusters.

1. Delegate to address (A) records, not to IP addresses. The SmartConnect service IP on an PowerScale cluster must be created in DNS as an address (A) record, also called a host entry. An A record maps a URL such as www.superna.net to its corresponding IP address. Delegating to an A record means that if you ever need to failover the entire cluster, you can do so by changing just one DNS A record. All other nameserver delegations can be left alone. In many enterprises, it is easier to have an A record updated than to update a name server record, because of the perceived complexity of the process.
2. Use one name server record for each SmartConnect zone name or alias. We recommend creating one delegation for each SmartConnect zone name or for each SmartConnect zone alias on a cluster. This method permits failover of only a portion of the cluster's workflow—one SmartConnect zone—without affecting any other zones. This method is useful for scenarios such as testing disaster recovery failover and moving workflows between data centers.
3. Follow consistent mount paths
    1. Mount entries for any NFS connections must have a consistent mount point, in the format of sczonename.domain.com:/ifs/path. This way, when you fail over, you don't have to manually edit your fstab or automount entries.
4. Use Access Zones to compartmentalize your data based on importance. If your environment is OneFS 7.1.1 or later and you use access zones, you must define an access zone root path to help segment data into the appropriate access zone

and enable the data to be compartmentalized. This is similar to what Celerra or VNX administrators might do if they have a VDM that has its own root file system. So, in addition to the default System access zone, you must add another layer. For example: /ifs/clustername/accesszonename/

5. Recommend to your client system administrators that they turn off client DNS caching, where possible. To handle client requests properly, SmartConnect requires that clients use the latest DNS entries. If clients cache SmartConnect DNS information, they might connect to incorrect SmartConnect zone names. In this situation, SmartConnect might not appear to be functioning properly.
6. Do NOT: We do not recommend creating a single delegation for each cluster and then creating the SmartConnect zones as sub records of that delegation

SmartConnect service IPs Each cluster needs only one SmartConnect service IP (SSIP), as long as there are no firewalls between the infrastructure DNS servers, and the SSIP that block TCP and UDP port 53. It doesn't matter how many domains or subnets the cluster is joined to or participates in. SmartConnect is essentially a very selective DNS server that answers only for the SmartConnect zone names and SmartConnect zone aliases that are configured on it. A DNS server doesn't have to respond with an IP address from the subnet that the DNS server is in: it responds only with the correct IP address based on the name being looked up. Which subnet the DNS server resides in is irrelevant.

This above means that failover to the target cluster can update the A record to point to the SSIP of the target cluster using the hints mapping described below for Eyeglass to create aliases in the correct smartconnect subnet on the target.

Best practice - Do DR Testing with RunBook Robot for Access Zones

Note: Runbook Robot is Access Zone Failover and allows testing of Access Zone failover on non-production access zones

1. It is best practice to setup an environment with non-production data and shares / exports / quotas representative of the production environment and run Failover and Failback testing to understand the failover operation in your environment with Eyeglass DR Assistant.

2. It is best practice to set up SyncIQ Robot for regular automated Failover and Failback for non-production data and shares / exports / quotas in your environment.