

Table of Contents

1. Eyeglass Microsoft DFS Mode Admin Guide.....	3
1.1. Overview.....	4
1.2. Windows OS Compatibility.....	7
1.3. Requirements for Eyeglass Microsoft DFS Mode Failover.....	12
1.4. How DFS mode with Eyeglass Works.....	14
1.5. Considerations for Eyeglass Microsoft DFS Mode Failover.....	23
1.6. Preparing your System for the Eyeglass Microsoft DFS Mode Failover.....	25
1.7. Failover Planning and Checklist.....	41
1.8. Monitoring DR Readiness for Eyeglass Assisted Failover.....	42
1.9. Operational Steps for Eyeglass Microsoft DFS Mode Failover.....	45
1.10. Post Eyeglass Microsoft DFS Mode Failover Manual Steps for NFS Exports.....	47
1.11. Post Eyeglass Microsoft DFS Mode Failover Checklist.....	48
1.12. Advanced DFS mode Setup with Access Zones.....	54
2. Eyeglass Access Zone Failover Guide.....	56
2.1. Introduction to this Guide.....	58
2.2. Automated SMB Client Switch Testing Matrix.....	64
2.3. How to Setup and Configure Access Zone - Overview Video.....	66
2.4. Requirements for Eyeglass Assisted Access Zone Failover.....	67
2.5. Unsupported Data Replication Topology.....	78
2.6. Overlapping Access Zone Failover Supported Configurations.....	79
2.7. Recommendations for Eyeglass Assisted Access Zone Failover.....	83
2.7.1. Mixed DFS and None DFS Solution.....	89
2.8. Preparing your Clusters for Eyeglass Assisted Access Zone Failover.....	92
2.9. PowerScale Administration for Clusters Configured for Eyeglass Assisted Access Zone Failover..	114
2.10. Failover Planning and Checklist.....	116
2.11. Monitoring DR Readiness for Eyeglass Assisted Failover.....	117
2.12. Operational Steps for Eyeglass Assisted Access Zone Failover.....	122
2.13. Post Access Zone Failover Steps.....	123
2.14. Post Access Zone Failover Checklist.....	130
2.15. IP Pool Failover.....	138
2.16. Fan-In IP Pool Failover.....	153
2.17. Fan-Out IP Pool Failover.....	169
2.18. How to Configure Access Zone DNS Dual Delegation.....	173
2.19. Controlled Failover Option Results Summary.....	187
2.20. How to Configure Delegation of Cluster Machine Accounts with Active Directory Users and Computers Snapin.....	191
3. Eyeglass SyncIQ Policy Failover Guide.....	202
3.1. Introduction to this Guide.....	203

3.2. Requirements and for Eyeglass Assisted SyncIQ Policy Failover.....	204
3.3. Unsupported Data Replication Topology.....	207
3.4. Recommendations for Eyeglass Assisted SyncIQ Policy Failover.....	208
3.5. Preparing your System for the Eyeglass Assisted SyncIQ Policy Failover.....	211
3.6. Operational Steps for Eyeglass Assisted SyncIQ Policy Failover.....	214
3.7. Monitoring DR Readiness for Eyeglass Assisted Failover.....	225
3.8. Post SyncIQ Policy Failover Manual Steps.....	227
3.9. Controlled Failover Option Results Summary.....	231
4. Eyeglass Runbook Robot Guide.....	234
4.1. Introduction to this Guide.....	235
4.2. Runbook Robot (Automate DR Testing on a schedule).....	240
4.3. Basic DR Robot Configuration.....	242
4.3.1. Windows DFS Configuration Example For Basic Robot.....	255
4.4. Advanced DR Robot Configuration.....	259
4.5. Advanced Settings.....	270
5. Eyeglass and PowerScale Compliance Mode Admin Guide.....	281
6. Multi Site Failover Guide for Continuous Availability.....	287
6.1. Overview.....	288
6.2. Logical Diagram of Multi Site Failover.....	292
6.3. Access Zone Failover - SyncIQ Configuration for 3 site.....	300
6.4. 3 Site DFS Mode Failover.....	323

1. Eyeglass Microsoft DFS Mode Admin Guide

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

- [Overview](#)
- [Windows OS Compatibility](#)
- [Requirements for Eyeglass Microsoft DFS Mode Failover](#)
- [How DFS mode with Eyeglass Works](#)
- [Considerations for Eyeglass Microsoft DFS Mode Failover](#)
- [Preparing your System for the Eyeglass Microsoft DFS Mode Failover](#)
- [Failover Planning and Checklist](#)
- [Monitoring DR Readiness for Eyeglass Assisted Failover](#)
- [Operational Steps for Eyeglass Microsoft DFS Mode Failover](#)
- [Post Eyeglass Microsoft DFS Mode Failover Manual Steps for NFS Exports](#)
- [Post Eyeglass Microsoft DFS Mode Failover Checklist](#)
- [Advanced DFS mode Setup with Access Zones](#)

© Superna Inc

1.1. Overview

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

Overview

The Eyeglass DFS solution for PowerScale 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 fallback 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.

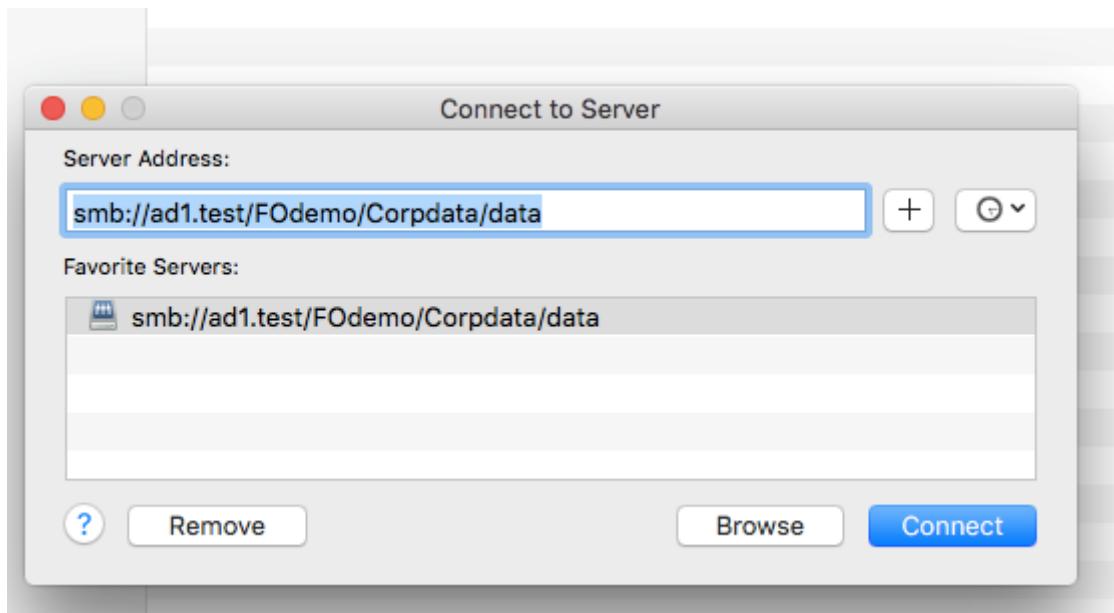
We recommend using Domain based DFS with Eyeglass as the most highly available solution for clients versus server based DFS roots.

Video How To - Overview

Key Values

- PowerScale Configuration synchronized to active cluster!
- Integrated with PowerScale SyncIQ and Eyeglass
Orchestrated Failover.

- Supports Quotas on PowerScale during failover and failback operations.
- Supports partial Access zone failover solutions. I.E per application level failover and failback within the Access zone (requires detected IP pool with Eyeglass ignore hints)
- NO DFS administration MMC tasks required for failover / failback!
- DFS referral list contains all targets (production and DR clusters).
- Automatically detected DFS targets protected.
- Kerberos ticket caching services failover without needing to clear Kerberos cache on clients.
- Avoids NTLM fallback authentication.
- NO CLIENT UNMOUNT needed!
- Clients should cache locally when selecting a target UNC path.
- Clients auto select writeable copy of the DFS mount.
- Apple OS X supports DFS mode!
- OS X Yosemite supports DFS with Active Directory if joined to the domain



© Superna Inc

1.2. Windows OS Compatibility

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

Windows DFS Client OS Compatibility and Release Notes

The following list of Windows OS clients have all been tested.

1. Windows 7, 8.x, 10
2. Windows Server 2008, 2012, 2012 R2, 2016
3. Samba version: Samba version 4.8.3 Configured for Microsoft DFS mounts and dual referral paths

Release Note

1. Onefs 8 has CA compatibility mode capability advertised called persistent file handles, This is required for Continuous Availability Mode support with PowerScale. Windows servers in cluster mode only advertise this capability when using CA mode.
1. An issue arise from Onefs 8 when NOT using CA mode this capability is still advertised to clients.
2. This can trigger the issue described in KB article below "**The computer can have more time to determine whether a shared folder is available if there is a failover of the shared folder.**"

3. NOTE: This issue only affected client machines with no active connection to PowerScale shares mounted over DFS. A delay was seen when mounting the DFS root that would complete within one minute. **NOTE: Actively mounted shares or mapped to DFS folders directly did not see this delay.**

2. Windows 10 and 2012 server have a registry setting described in the link below to correct this behavior. Windows 8 and 2012 server require a hotfix to be applied.

a. <https://support.microsoft.com/en-us/help/2820470/delayed-error-message-when-you-try-to-access-a-shared-folder-that-no-l>

DFS Feature Changes and Share names used on DFS synced Shares

To streamline how DFS mode with normal configuration sync, the feature has been enhanced as per the table below. This change allows new options for customers that require access to DR cluster data in read-only state, and preserves DFS mode functionality. It also reduces risk of issues during failover.

New feature to hide shares on the DR cluster or read only cluster are now available. After switching the tag, the next configuration sync cycle will apply the changes to the DR cluster share names.

Eyeglass Version	DFS Mode Sync Mode	DFS mode failover Behavior	Prefix on Shares	Post fix on shares for Security on DR Cluster
1.4.x and earlier	Delete shares of the same name on DR cluster	Create shares on DR cluster and delete on source cluster		
1.5 and beyond	Create shares on DR cluster with a renamed prefix added to the share name	Rename share on DR cluster and rename source cluster with a prefix added to the share name	<p>igls-dfs (this default can be changed and be changed to another string by editing the tag <dfsshareprefix>igls-dfs-</dfsshareprefix> in /opt/superna/sca/data/system.xml)</p> <p>WARNING: If DFS is enabled, changing this tag will NOT clean up shares with the previous name. Clean up is manual and new shares will be created with the new tag.</p>	
1.6 and beyond				<p>New feature to allow the source active cluster share to be visible BUT a \$ post fix can be applied on the Synced igl-dfs-sharename\$ to hide the share on the DR cluster. After failover the share names will flip and hide it on the Source cluster. This can be enabled by editing /opt/superna/sca/data/system.xml file. Change the tag to include \$ as shown below <dfssharesuffix>\$</dfssharesuffix></p> <p>NOTE: on an existing installation the old DFS renamed share will need to be manually deleted</p>

NOTE: When 1.5 DFS mode is enabled, all shares found on source will be created on the target cluster with the prefix

applied to the share name. If upgrading from 1.4.x DFS mode, no action is required and shares will be created using 1.5 logic.

DFS Fast Failover Mode - Superna Eyeglass 1.5.2 and beyond

Release	Speed Improvement	Notes
1.5.2 >	For DFS Failover (Microsoft DFS Mode or DFS enabled Job in an Access Zone Failover), the Renaming shares step occurs after Data sync and before the Policy Failover step (Allow Writes, Resync Prep). This ensures that the amount of time that DFS clients are directed to the failover source cluster is minimized once the failover has started and that the DFS clients are already directed to the target cluster when the filesystem becomes writeable.	NOTE: During failover, clients with open files will now receive a read-only error message if they attempt to save data once redirection has occurred but before the target is writeable. This is expected and gives the user feedback that writes will not be successful. Each application has different behaviour in how it returns a read-only file system error to the user.
1.6.0 >	Parallelized Rename - Now the rename process can use up to 10 threads at once to rename 10 shares in parallel across all policies in a failover job. This will provide a 10x speed improvement to redirect DFS clients faster under all failover conditions. With large share or policy count failovers getting accelerated by a factor of 10	

DFS Failover Enhancements

Release	Enhancement	Notes
1.9 >	<p>For DFS Failover (Microsoft DFS Mode or DFS enabled Job in an Access Zone Failover), following Share Rename Step Enhancement has been made:</p> <ul style="list-style-type: none"> • If share renaming is failed for all the shares for a cluster, then failover status is error and failover is stopped. This aborts the failover and leaves the data accessible on the source cluster. • If share renaming is failed only for 	Summary: This eliminates the possibility of data access outage from share rename step and ensures if some shares rename the failover will continue.

	some shares for a cluster, then failover status is warning and manual recovery on the shares that failed to rename is required.	
--	---	--

© Superna Inc

1.3. Requirements for Eyeglass Microsoft DFS Mode Failover

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

Requirements for Eyeglass Microsoft DFS Mode Failover

The requirements in this section must be met in order to initiate an Eyeglass Microsoft DFS Mode Failover. Failure to meet some of these conditions may block the Failover.

Cluster Version Requirements

Clusters participating in an Eyeglass Microsoft DFS Mode Failover must be running the supported PowerScale Cluster version for this feature. See the Feature Release Compatibility matrix in the Eyeglass Release notes specific to your Eyeglass version found [here](#).

SynIQ Policy Requirements - Blocks Failover

For a Microsoft DFS Mode failover with Eyeglass, **it is required that** the Eyeglass Configuration Replication Job for that **SynIQ Policy is in the Enabled state.**

Note: If the SynIQ Policy is disabled in OneFS or the corresponding Eyeglass Configuration Replication Job is disabled in Eyeglass the SynIQ Policy **failover will be blocked.**

Failover Target Cluster Requirements - Blocks Failover

For a SyncIQ Policy failover with Eyeglass, it is required that the PowerScale Cluster that is the target of the failover be IP reachable by Eyeglass with the required ports open.

Eyeglass Quota Job Requirements

For a SyncIQ Policy failover with Eyeglass, there are no Eyeglass Quota Job state requirements. Quotas will be failed over whether Eyeglass Quota Job is in Enabled or Disabled state.

Active Directory

- AD clients must have both paths of the folder target cached post failover.
- AD clients must be able to contact a Domain Controller.
- UNC paths to mount the DFS folder must use DFS UNC syntax \\domain name\dfsrootname\dfs folder name.
- SmartConnect zone names in the UNC targets must be delegated and resolvable by clients.
- SmartConnect zone name SPN's for folder target UNC paths must be correctly registered in AD.

© Superna Inc

1.4. How DFS mode with Eyeglass Works

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

- [Process Flow for DFS Failover with Eyeglass](#)
- [Normal](#)
- [Failover with Eyeglass DFS mode](#)
- [Failback with Eyeglass](#)
- [Eyeglass Microsoft DFS Mode Failover with NFS Export](#)
- [Linux use of DFS with Samba with DFS mode Failover](#)
- [How to Configure Samba with Linux with DFS mounts](#)

Process Flow for DFS Failover with Eyeglass

Normal

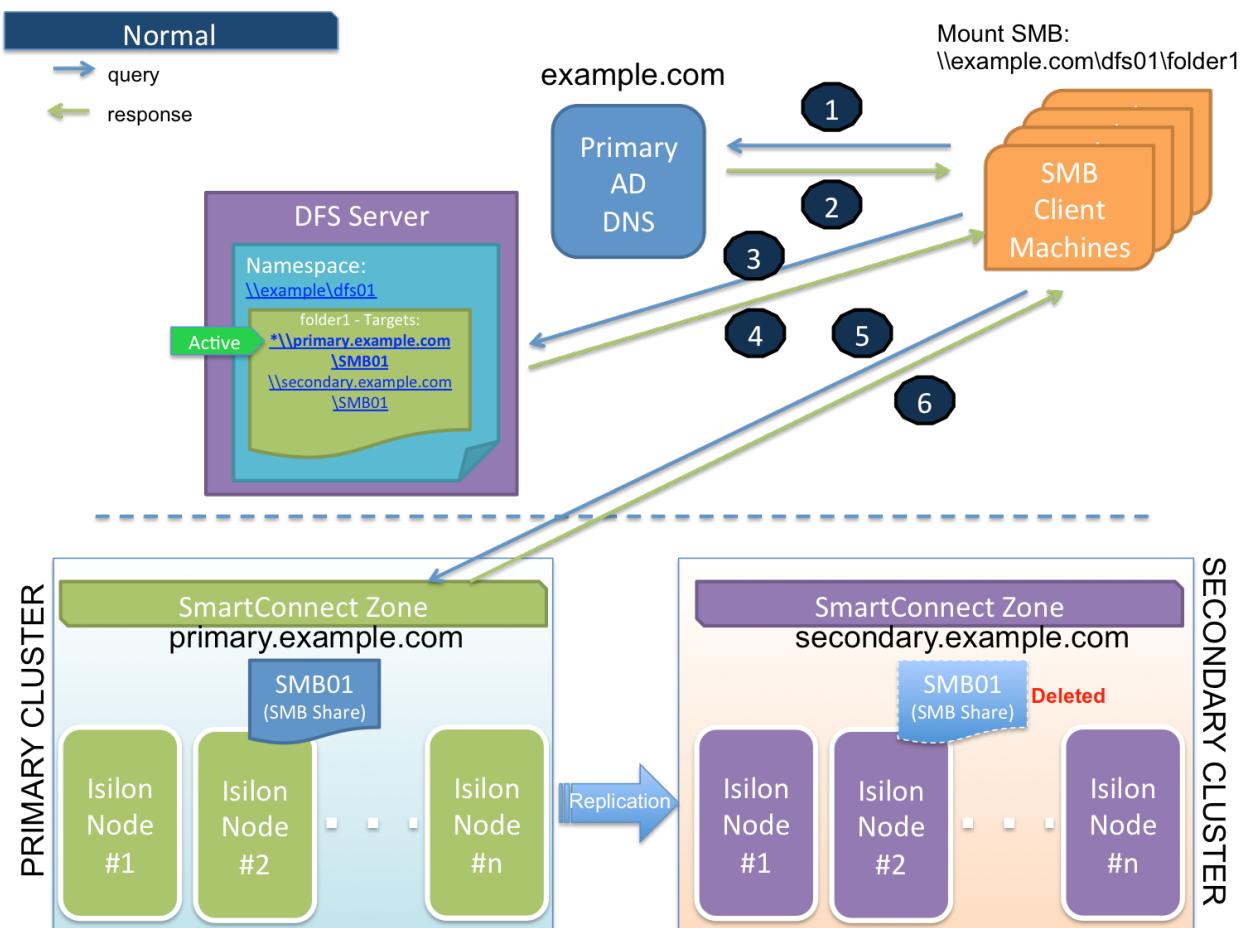
The following diagram displays the process flow in normal condition (before Failover)

Configuration:

DFS Target folder was configured to refer to both Primary Cluster and Secondary Cluster's SMB Shares, with the priority set to the

Primary Cluster. The Active Directory Sites and Subnets were also configured to let the clients resided on the same subnet and site as the PowerScale Primary Cluster IP addresses.

SMB Shares were created on Primary and Secondary Clusters. See [DFS Feature changes table \(DFS Feature Changes and Share names used on DFS synced Shares\)](#) for version behaviour.

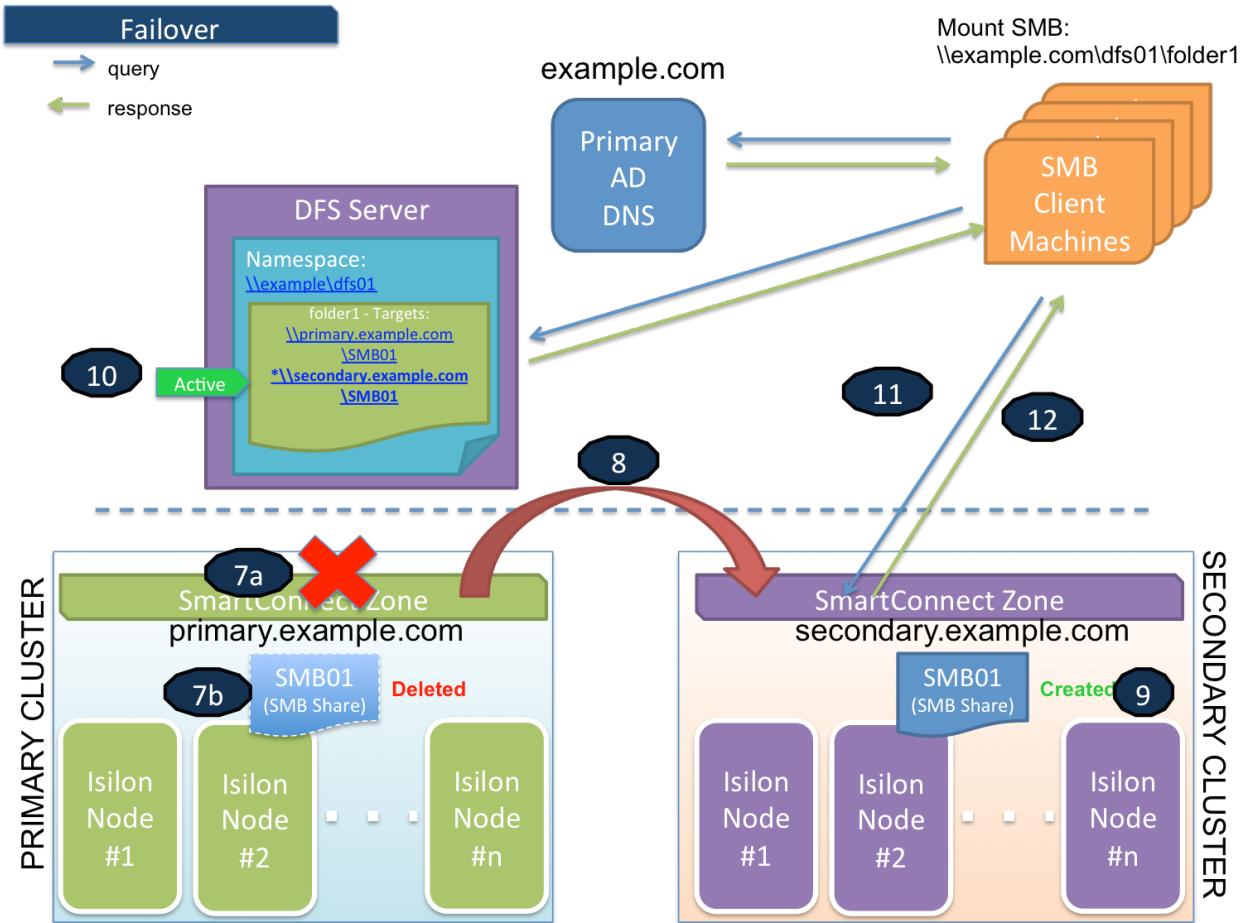


1. SMB Client is accessing a domain-based namespace (e.g. \\example.com\dfs01\folder1). This SMB client computer sends a query to the AD to discover a list of root targets for the namespace.

2. AD Controller returns a list of root targets defined for the requested namespace.
3. SMB client selects the root target from the referral list and sends a query to the root server for the requested link.
4. DFS root server constructs a list of folder targets in the referral. Order / priority of targets can be configured. Folder Target to the Primary Cluster is configured as the first path to refer to (higher priority) in the referral list. The SMB Client and Primary Cluster are also residing in the same Active Directory Site and Subnet. The SMB Share(s) on Secondary Cluster is not active (Deleted). DFS root server sends the referral information to the client. The active path is to the Primary Cluster.
5. SMB client tries to establish a connection to the selected target (the first priority / active target in the list).
6. PowerScale (Primary cluster) responds to this SMB connection.

Failover with Eyeglass DFS mode

The following diagram displays the process flow during Failover event.



Flow:

1. Primary Cluster failure or down detected,

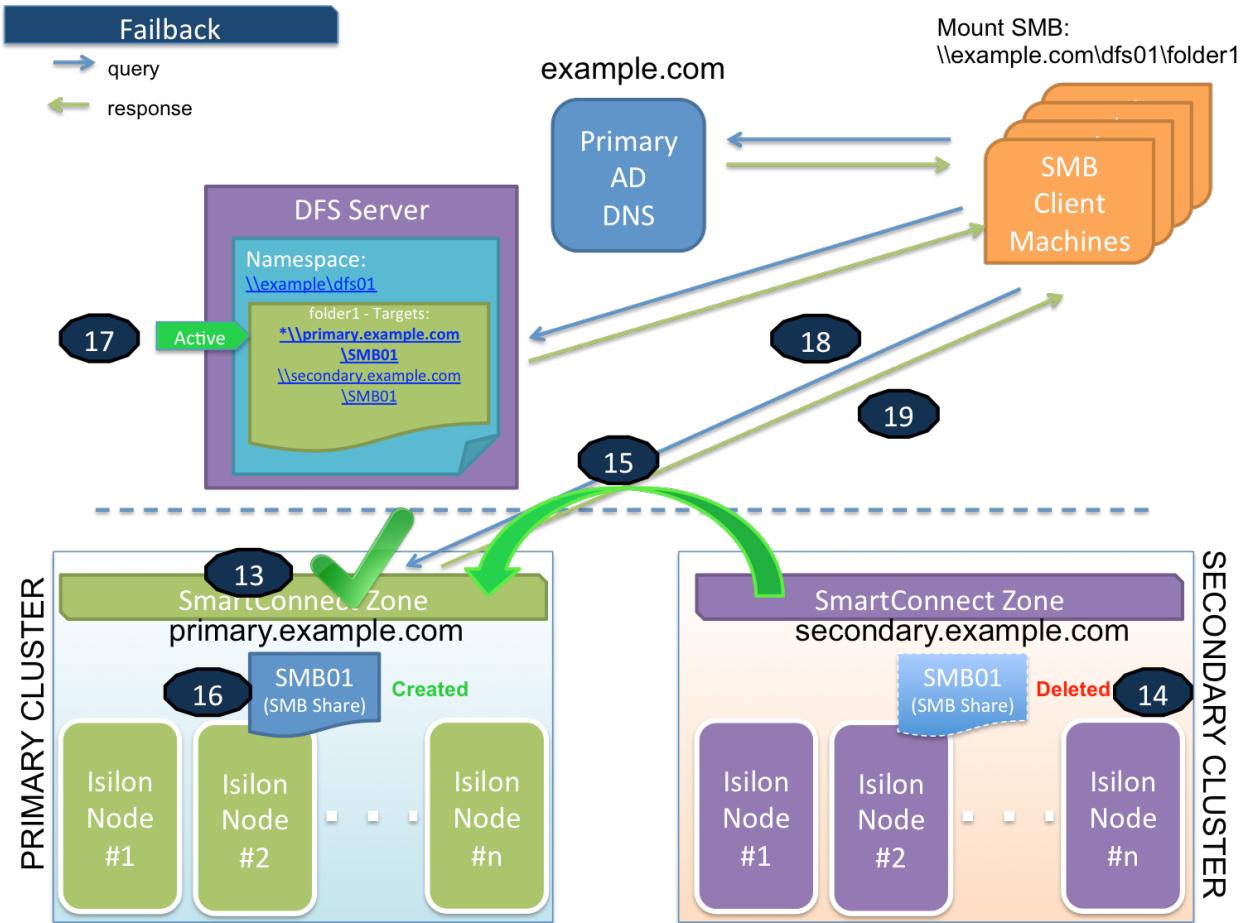
 1. Initiate Eyeglass Failover of DFS syncIQ policy protecting the DFS UNC targets (7a)

2. Eyeglass **deletes or renames** SMB share(s) on Primary Cluster (7b) **See DFS Feature changes table ([DFS Feature Changes and Share names used on DFS synced Shares](#)) for version behaviour.**
 1. Eyeglass **creates** quotas on Secondary cluster and **deletes** them from the Primary cluster
 2. Eyeglass performs SyncIQ Failover to Secondary Cluster

1. Eyeglass syncs data changes with SyncIQ to Secondary cluster move schedule over and runs re-sync prep (**8**)
3. Eyeglass **creates or renames** SMB share(s) on Secondary Cluster **See DFS Feature changes table ([DFS Feature Changes and Share names used on DFS synced Shares](#)) for version behaviour.**
4. DFS Target folder - path to the Secondary cluster will automatically be activated
5. SMB Client is connecting to the Secondary cluster
6. Secondary cluster is responding to SMB client

Fallback with Eyeglass

The following diagram displays the process flow during Fallback event.



1. Primary cluster is resumed and ready for failback. (**13, 14, 15, 16**)
1. Perform Eyeglass Failback of DFS SyncIQ policy protecting the DFS UNC targets (**15**)
2. Eyeglass syncs data changes with SyncIQ to Primary cluster (**15**)
3. Eyeglass performs SyncIQ Failover to Primary Cluster
4. Eyeglass **creates or renames** SMB share(s) on Primary Cluster (**16**) See **DFS Feature changes table (DFS Feature Changes and Share names used on DFS synced Shares)** for version behaviour.

5. Eyeglass **deletes or renames** SMB share(s) on Secondary Cluster **(14)** See **DFS Feature changes table ([DFS Feature Changes and Share names used on DFS synced Shares](#))** for version behaviour.
6. Eyeglass **creates** quotas on Primary cluster and **deletes** them from the Secondary cluster
7. DFS Target folder - path to the primary cluster will automatically be activated.
8. SMB Client is connecting to the primary cluster
9. Primary cluster is responding to SMB client

Eyeglass Microsoft DFS Mode Failover with NFS Export

Eyeglass Microsoft DFS Mode Failover can be used with SyncIQ policies that protect NFS exports by the same policy, but in this case manual steps or post-failover scripting must be used to update NFS client mounts.

Linux use of DFS with Samba with DFS mode Failover

Samba is a SMB port to Linux. This can mount SMB shares similar to how NFS mounts data.

This has been tested with DFS presented folder with dual referral paths. Testing showed that samba DFS folders did not auto switch

cluster smartconnect paths until the Interface removal method was used to remove the Interfaces from the source cluster IP pool. This procedure is fully described [here](#).

How to Configure Samba with Linux with DFS mounts

Install Linux DFS client

Environment

Centos Release: CentOS Linux release 7.6.1810 (Core)

Samba version: Samba version 4.8.3

Windows server version: 2016 Essentials

Microsoft Corporation DFS Management : Version: 6.0

Installation

- Install packages using yum. Package needed: samba samba-client cifs-utilskeyutils
 - yum install -y samba samba-client cifs-utils keyutils
- Enable firewall rules for SMB
 - firewall-cmd --permanent --zone=public --add-service=samba && firewall-cmd --reload && systemctl restart firewalld
- Modify /etc/request-key.conf file

- `sed -i "\$acreate cifs.spnego * * /usr/sbin/cifs.upcall %k\ncreate dns_resolver * * /usr/sbin/cifs.upcall %k" /etc/request-key.conf`
- Create Local directory for mount point and assign permission [tested with 777permission]
 - `mkdir -p /tmp/dfs && chmod 777 /tmp/dfs`
- Mount the DFS root [can use vers=3]
 - `mount -t cifs //DOMAIN/DFSRoot /tmp/dfs -o sec=ntlmv2, domain=DOMAIN, username=usr, vers=2.1`
- Browse to local mount directory and list your DFS Shares
 - `cd /tmp/dfs && ls -lha`
- Check your mount information
 - `df -h`

© Superna Inc

1.5. Considerations for Eyeglass Microsoft DFS Mode Failover

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

Considerations for Eyeglass Microsoft DFS Mode Failover

The following are highly recommended to ensure that all automated Eyeglass Microsoft DFS Mode failover steps can be completed.

SynIQ Policy Recommendations

- SynIQ Job in OneFS should have been completed without error and shows green.
- **Impact:**
- Failover will be blocked if SynIQ policies are in an error state on the cluster. Eyeglass will attempt to run the policy which will fail. Correct this on the OneFS cluster.
- Data loss due to unreplicated data.
- PowerScale does not support SynIQ Policy with excludes (or includes) for failover.
- **Impact:** Not a supported configuration for failback.
- PowerScale best practices recommend that SynIQ Policies utilize the Restrict Source Nodes option which requires an IP to be created with target SmartConnect zone..

- **Impact:** Subnet pool used for data replication is not controlled all nodes in the cluster can replicate data from all IP pools. This is hard to manage bandwidth and requires all nodes to have access to the WAN.

© Superna Inc

1.6. Preparing your System for the Eyeglass Microsoft DFS Mode Failover

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

Preparing your System for the Eyeglass Microsoft DFS Mode Failover

This section contains steps to configure an Eyeglass Microsoft DFS Mode Failover / Failback solution. Please refer to Microsoft DFS, PowerScale and Eyeglass documentation for their respective detailed configuration steps.

Overview

Step 1 - Configure PowerScale Smartconnect Zones.

Step 2 - Configure PowerScale SyncIQ Policies and SMB Shares (only source cluster needs the share created).

Step 3 - Configure DFS root (domain based).

Step 4 - Configure DFS folder.

Step 5 - Configure DFS Target (Primary cluster): \\SmartConnect Zone\share name.

Step 6 - Verify Client Access example:
\\domain.name\DFSrootname\

Step 7 - Eyeglass Setup and SyncIQ policy enabled for DFS Mode (enable mirror policy if it exists):

- run DFS Mode policy and verify its green OK.

Step 8 - Configure DFS Target (Secondary cluster): \\Smartconnect Zone\\share name.

Step 9 - Verify Client Access:

- mount DFS path: \\domain.name\\DFSrootname\\folder name.
- Write data successfully.

Preparation Steps

Step 1 - Configure PowerScale SmartConnect Zones:

DFS Folder Targets will be configured using UNC Path to Primary and Secondary SMB Share(s) by using their SmartConnect Zone names. If the PowerScale clusters do not already have SmartConnect Zone names configured they will need to be setup.

Step 2 - Configure PowerScale SyncIQ Policies and Shares:

Create the SyncIQ policies and shares required to protect and access the data if not already existing.

Design Considerations:

1. A single SyncIQ policy can have multiple DFS targets underneath it in the file system path.
 1. This means all DFS targets using this policy failover together.
 2. To achieve per application failover a single SyncIQ policy for each DFS target is required.

Step 3 - Configure DFS root

Recommend Domain based DFS for higher availability with replicated DFS root between domain controllers.

Step 4 - Create DFS folder

Step 5 - Configure Folder DFS Target (Primary cluster)

Configure the DFS Target for the Primary cluster (cluster which is currently active) by adding UNC path to share created on Primary cluster. This UNC path must use a SmartConnect Zone on the Primary cluster.

Example: \\production.example.com\SMBshare1

Step 6 - Verify Client Access

Check access from DFS enabled client and ensure that they have read / write access enabled to the Primary cluster as per the DFS Folder(s) and Target(s) configured.

IMPORTANT: You must use path including fully qualified active directory domain DNS name - for example:

\ad1.test\DFSRoot\ShareFolder

Tools help debug and plan DFS for DR before going into production. Test machines should install these tools which can be found here: ["Where to find DFSUTIL.EXE for Windows Server "](#)

Verify that the SMB client is accessing the DFS target folder's SMB Share from the primary cluster. DFS utility command to verify: (note DFS client utilities must be installed into the OS, they are not standard tools)

- dfsutil cache referral (ensure both targets are listed)
- dfsutil cache referral flush (clear the cache for debugging)
- dfsutil diag viewdfspath \\domainname\dfsroot\dfsfolder

Example:

C:\Windows\system32>dfsutil cache referral

4 entries...

Entry: \DFS-SVR-01\dfs02\folder2

ShortEntry: \DFS-SVR-01\dfs02\folder2

Expires in 270 seconds

UseCount: 1 Type:0x1 (DFS)

0:[\cluster11-z01.ad1.test\SMB-Share-002] AccessStatus: 0 (ACTIVE TARGETSET)

1:[\cluster12-z01.ad1.test\SMB-Share-002] AccessStatus: 0xc00000cc (TARGETSET)

Entry: \ad1.test\dfs02

ShortEntry: \ad1.test\dfs02

Expires in 270 seconds

UseCount: 0 Type:0x81 (REFERRAL_SVC DFS)

0:[\DFS-SVR-01\dfs02] AccessStatus: 0 (ACTIVE TARGETSET)

Entry: \AD1.test\sysvol

ShortEntry: \AD1.test\sysvol

Expires in 0 seconds

UseCount: 0 Type:0x1 (DFS)

0:[\ad1.AD1.test\sysvol] AccessStatus: 0 (ACTIVE TARGETSET)

Entry: \DFS-SVR-01\dfs02

ShortEntry: \DFS-SVR-01\dfs02

Expires in 270 seconds

UseCount: 0 Type:0x81 (REFERRAL_SVC DFS)

0:[\DFS-SVR-01\dfs02] AccessStatus: 0 (ACTIVE TARGETSET)

DfsUtil command completed successfully.

C:\Windows\system32>dfsutil diag viewdfspath

\ad1.test\dfs02\folder2

The DFS Path <\ad1.test\dfs02\folder2> resolves to ->

<\production.example.com\SMBshare1>

Done processing this command.

Step 7 - Eyeglass setup:

1. Install Eyeglass and add clusters as managed devices to Eyeglass.
2. Allow Eyeglass to complete initial discovery and creation of Eyeglass configuration replication Jobs.
3. Enable Eyeglass DFS mode on the Eyeglass Configuration Replication Jobs for the SyncIQ policies that will be using the Eyeglass DFS Failover solution.

NOTE: If a SyncIQ policy has already been failed over in OneFS and a mirror policy exists, enable Eyeglass DFS Mode for both the active and disabled policy.

Job Definitions	Job Name	Policy	Type	Last Run Date	State
Running Jobs	disaster8_marketing-shares-policy_m...	marketing-shar...	AUTO	10/2/2017, 9:15:32 AM	OK
	disaster8_EyeglassRunbookRobot-de...	EyeglassRunboo...	AUTO	10/1/2017, 9:05:09 PM	Policy Disa...
	disaster8_marketing-nfs-policy_mirror	marketing-nfs-p...	AUTO	10/2/2017, 9:15:32 AM	OK
Configuration Replication: DFS mode (AUTOMATIC)					
	prod-cluster-8_System-Zone-DFS	System-Zone-DFS	AUTODFS	10/2/2017, 9:15:32 AM	OK
	Cluster2-7201_data-zone-nfsdata_mir...	data-zone-nfsda...	AUTODFS	10/2/2017, 9:15:32 AM	OK
	Cluster2-7201_data-zone-shares_mirr...	data-zone-share...	AUTODFS	10/2/2017, 9:15:32 AM	OK
	Cluster2-7201_data-zone-dfs_mirror	data-zone-dfs_...	AUTODFS	10/2/2017, 9:15:32 AM	OK
	prod-8_marketing-nfs-policy	marketing-nfs-p...	AUTODFS	n/a	Policy Disa...
	prod-8_marketing-shares-policy	marketing-shar...	AUTODFS	n/a	Policy Disa...
	prod-8_marketing-dfs	marketing-dfs	AUTODFS	10/2/2017, 9:15:32 AM	OK
	disaster8_marketing-dfs_mirror	marketing-dfs_...	AUTODFS	10/2/2017, 9:50:31 AM	Policy Disa...
Disaster Recovery Testing (AUTOMATIC)					
	Cluster2-7201_Eyeglass-DR-Testing	Eyeglass-DR-Tes...	DRTESTING	10/2/2017, 9:15:33 AM	OK

Show Disabled Jobs 3 item(s) selected

Select a bulk action ▾ Add New Job

- Edit Configuration(s)
- Run Now
- Enable/Disable Microsoft DFS (selected)
- Set Schedule
- Enable/Disable
- Delete
- Enable/Disable Skip Config Replication

1. DFS Enabled SyncIQ policies will show in a new folder and be of Type AUTODFS.

Name:	Cluster2-7201_data-zone-dfs_mirror	data-zone-dfs_...	AUTO	10/2/2017, 9:50:32 AM	OK
Enabled/Disabled:	ENABLED				
Job Type:	AUTO				
Source:	Cluster2-7201				
Path:	/ifs/data/userdata/dfs1				
Target:	prod-cluster-8				
Path:	/ifs/data/userdata/dfs1				
Last Success:	10/2/2017, 9:50:32 AM				

2. Initially the state of the Eyeglass AUTODFS job will be pending. Once the next Eyeglass replication cycle has been executed and the new DFS job has been run the state will be updated. The state will be OK if no errors were encountered.

NOTE 1: (See DFS Feature changes table for version behaviour.) With Eyeglass DFS mode, the SMB shares only exist on the active cluster. This is expected and will not interfere with the failover. If Eyeglass detects that the share does exist on the standby cluster, Eyeglass DFS mode Jobs (AUTODFS) will delete it on the cluster that is currently the SyncIQ target (read-only) cluster.

NOTE 2: If the SyncIQ policy associated with an Eyeglass DFS mode (AUTODFS) job is renamed, Eyeglass considers this to be a new, previously unknown Job and it will be created as a regular Eyeglass Configuration Replication (AUTO) job. You must re-enable DFS for this job.

NOTE 3: See DFS Feature changes table for version behaviour.

Step 8 - Configure DFS Target (Secondary cluster):

Configure the DFS Target for the Secondary cluster (cluster which is NOT currently active) by adding UNC path to share created on

Secondary cluster. This UNC path must use a Smartconnect Zone on the Secondary cluster.

Example: \\dr.example.com\SMBshare1

!! IMPORTANT

You must use exactly the same share name as was used for Primary cluster DFS target.

Confirm that both DFS Targets for Primary and Secondary clusters are enabled.

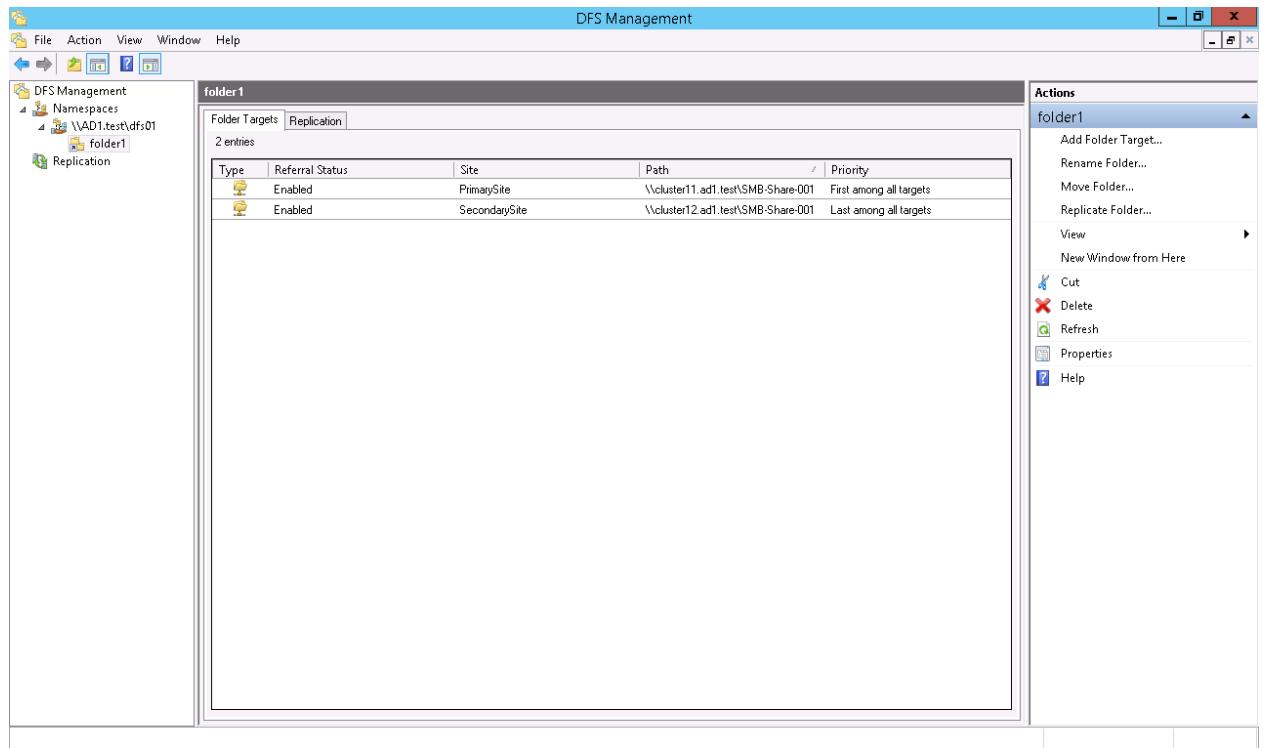
Design Considerations:

Configure Folder Targets with UNC Path to Primary and Secondary SMB Share by using their SmartConnect Zone names. **NOTE:**

With Eyeglass DFS mode, the shares only exist on 1 cluster at a time. You must create both target UNC's in DFS folder targets. The adding of the target cluster will not be able to detect the site information since the share does not exist.

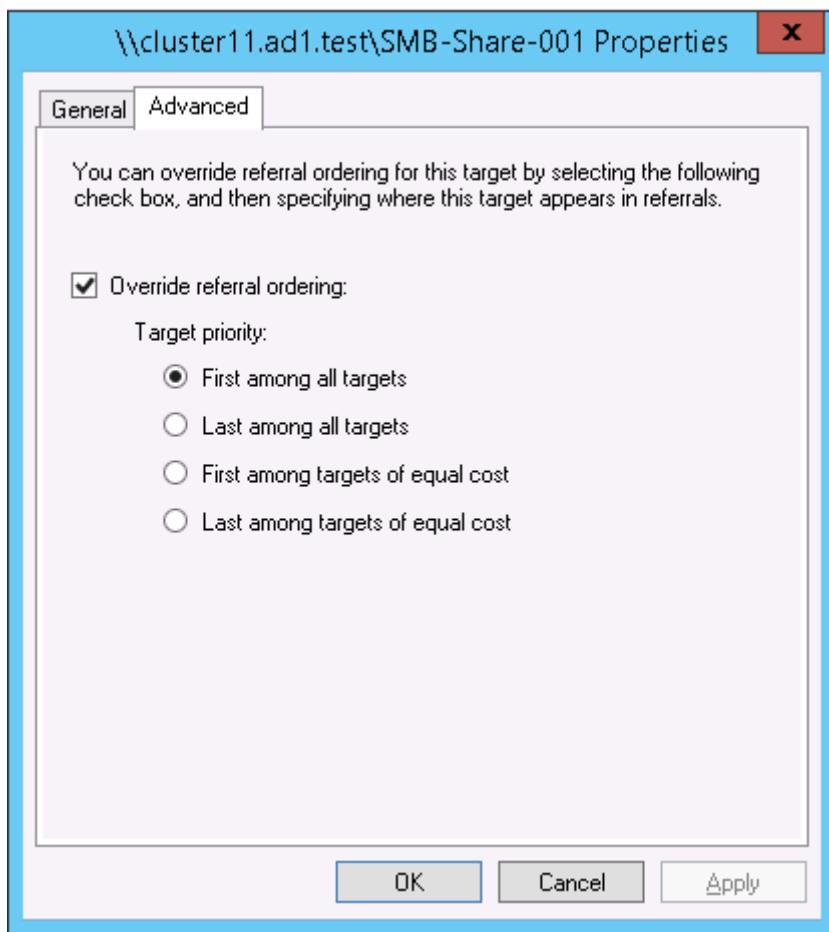
This is ok and will not interfere with the failover. (See DFS

Feature changes table for read-only share version behaviour.)

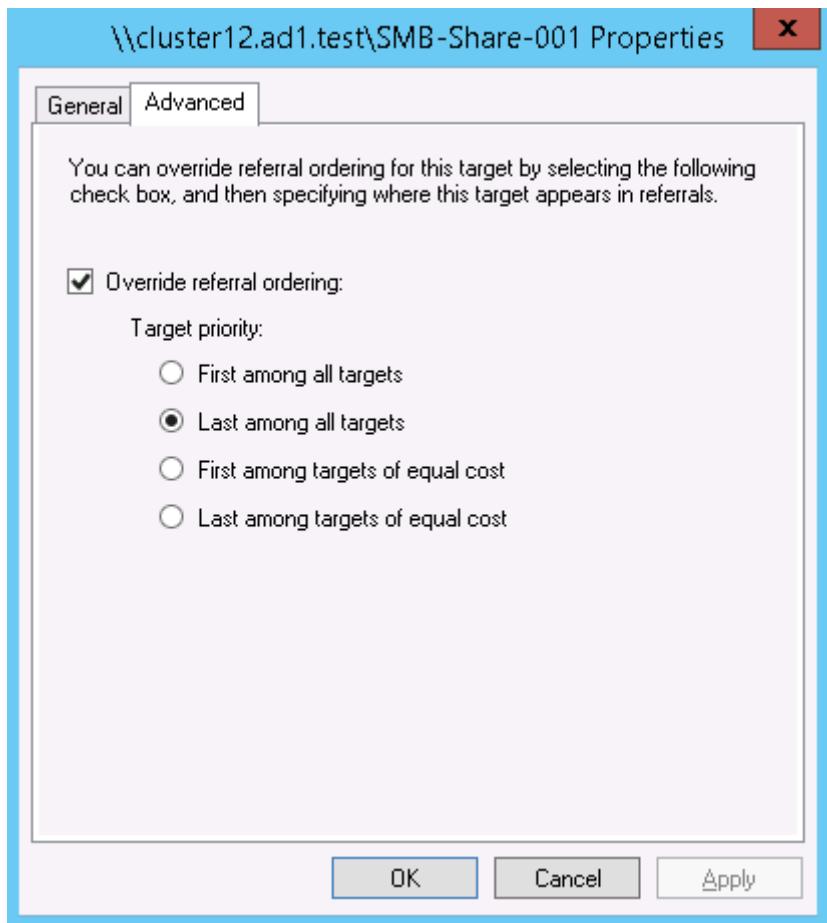


Set the Target Folder Priorities

Priority for Primary Cluster: First among all targets.



Priority for Secondary Cluster: Last among all targets.



Active Directory Sites and Subnets for PowerScale clusters and SMB Clients (Optional)

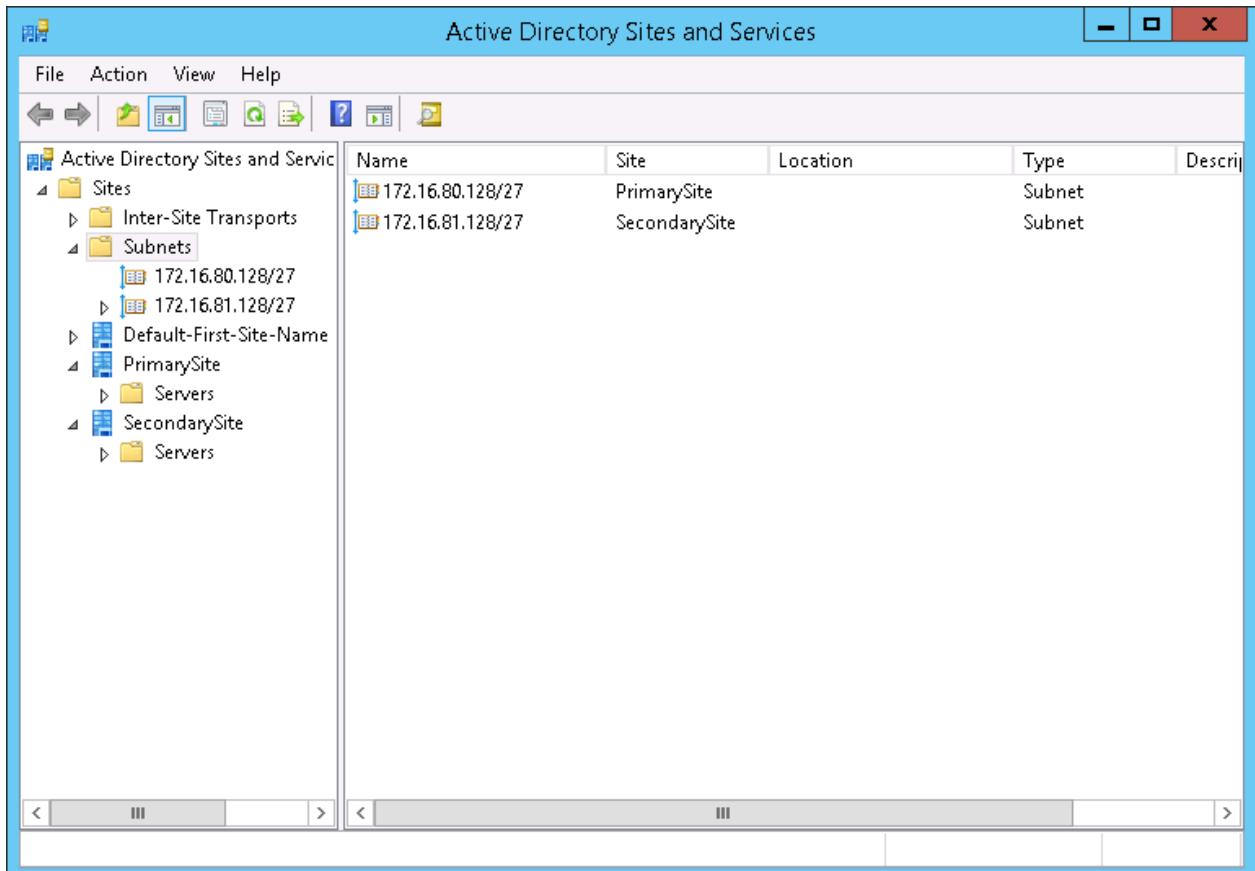
NOTE: It is not required for successful Eyeglass DFS mode failover.

Best practice:

This step is **optional** and helps optimize the client decision on which cluster to mount using IP subnets. This requires subnets are configured for all UNC targets and clients in Active Directory. For large networks this may not be practical. If not configured clients will pick a target based on the referral list without using sites.

Use two separate subnets and sites for the Primary cluster (PrimarySite) and Secondary cluster (SecondarySite). Example:

1. PrimarySite:Subnet 172.16.80.128/27
2. SecondarySite:Subnet 172.16.81.128/27



Place the SMB client's IP address in the Primary Site subnet's IP address range. With this setting, the Primary Cluster is referred to as the primary folder target path (higher priority).

Step 9 - Verify Client Access

Check access from DFS enabled client and ensure that they still have read / write access enabled to the Primary cluster ONLY as per the DFS Folder(s) and Target(s) configured.

IMPORTANT: You must use path including fully qualified active directory domain DNS name - for example:

\ad1.test\DFSRoot\ShareFolder

Tools help debug and plan DFS for DR before going into production. Test machines should install these tools.

[DFS util downloaded by OS type](#)

Now we have DFS folder target paths configured to have access to Primary and Secondary Cluster's SMB Share. But only the SMB Share on Primary Cluster is active (in normal condition / before failover). The SMB share on Secondary Cluster was created for DFS Folder Target creation as the second path. After it had been registered to DFS Folder setting, then it was deleted. With this initial state, you can verify that the SMB client is still accessing the DFS target folder's SMB Share from the primary cluster even though a DFS target for the Secondary cluster exists. DFS utility command to verify: (note DFS client utilities must be installed into the OS, they are not standard tools).

- dfsutil cache referral
- dfsutil diag viewdfspath

Example:

C:\Windows\system32>dfsutil cache referral

4 entries...

Entry: \DFS-SVR-01\dfs02\folder2

ShortEntry: \DFS-SVR-01\dfs02\folder2

Expires in 270 seconds

UseCount: 1 Type:0x1 (DFS)

0:[\cluster11-z01.ad1.test\SMB-Share-002] AccessStatus: 0 (ACTIVE TARGETSET

)

1:[\cluster12-z01.ad1.test\SMB-Share-002] AccessStatus:
0xc00000cc (TARGETSE

T)

Entry: \ad1.test\dfs02

ShortEntry: \ad1.test\dfs02

Expires in 270 seconds

UseCount: 0 Type:0x81 (REFERRAL_SVC DFS)

0:[\DFS-SVR-01\dfs02] AccessStatus: 0 (ACTIVE TARGETSET)

Entry: \AD1.test\sysvol

ShortEntry: \AD1.test\sysvol

Expires in 0 seconds

UseCount: 0 Type:0x1 (DFS)

0:[\ad1.AD1.test\sysvol] AccessStatus: 0 (ACTIVE TARGETSET)

Entry: \DFS-SVR-01\dfs02

ShortEntry: \DFS-SVR-01\dfs02

Expires in 270 seconds

UseCount: 0 Type:0x81 (REFERRAL_SVC DFS)

0:[\DFS-SVR-01\dfs02] AccessStatus: 0 (ACTIVE TARGETSET)

DfsUtil command completed successfully.

C:\Windows\system32>dfsutil diag viewdfspath

\ad1.test\dfs02\folder2

The DFS Path <\ad1.test\dfs02\folder2> resolves to ->

<\production.example.com\SMBshare1>

Done processing this command.

Configuration complete

1.7. Failover Planning and Checklist

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

Failover Planning and Checklist

Failover planning includes extended preparation beyond storage layer failover steps as related to the clients, application owners and any dependent systems such as DNS and Active Directory. A full Failover Plan is required taking this all into account. A [Failover Planning Guide and checklist](#) document is provided for input into your own Failover plan.

© Superna Inc

1.8. Monitoring DR Readiness for Eyeglass Assisted Failover

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

Monitoring DR Readiness for Eyeglass Assisted Failover

In addition to the Assisted Failover functionality, Eyeglass also provides the following features to monitor your Microsoft DFS Mode Failover Readiness:

- DR Readiness Validation

Eyeglass DR Readiness for DFS Enabled SyncIQ Policies

Eyeglass DFS Readiness DR Status provides a quick and easy way to assess the status on DFS mode enabled SyncIQ policies readiness for failover or fallback operations. DFS Readiness DR Status is “OK” when all of the conditions below are met:

- Your SyncIQ Policy is enabled.
- Your SyncIQ Policy Last Started and Last Success timestamp are identical.
- Your Eyeglass configuration replication Job is enabled.
- Your Eyeglass configuration replication Job Last Run and Last Success timestamp are identical.
- This validates that shares only exist on the writeable copy of the SyncIQ data. This is the normal operating mode for

Eyeglass DFS enabled policies. (See DFS Feature changes table for read-only share version behaviour.)

- Your Eyeglass configuration replication Job Audit Status is OK.

To check the DFS Readiness DR Status for your SyncIQ Policies in DFS mode:

1. Login to Eyeglass.
2. Open the **DR Dashboard**.
3. Select the **DFS Readiness** tab.

The DR Status is displayed per policy. It also indicates which pair of clusters are used in the DFS configuration and the associated SyncIQ policy.

DR Dashboard						
Zone Readiness	Name	SyncIQ Policy	Source	Destination	Last Successful Readiness Check	DR Failover Status
Pool Readiness	<input type="checkbox"/> prod-cluster-8_System-Zone-DFS	System-Zone-DFS	prod-cluster-8	Cluster2-7201	9/29/2017, 9:05:51 AM	INFO
DFS Readiness	<input type="checkbox"/> Cluster2-7201_System-Zone-DFS_mirror	System-Zone-DFS_mi...	Cluster2-7201	prod-cluster-8	9/29/2017, 9:05:54 AM	FAILED_OVER
Policy Readiness	<input type="checkbox"/> prod-8_marketing-dfs	marketing_dfs	prod-8	elasticsearch8	9/29/2017, 9:05:56 AM	FAILED_OVER

Expand a policy to see the details for the SyncIQ Policy and the Eyeglass Configuration Replication status:

- Last Run time of the SyncIQ policy and the status of the last run.

- Last Run time of the Eyeglass Configuration Replication job and the status of the last run and audit.

If new shares are created on the DFS mode policy, the next run of the Eyeglass Configuration Replication job in DFS mode will be aware of the new shares and ready to fail them over. It is important to check this job's status after creating more DFS shares under a policy.

The screenshot shows the DR Dashboard interface. The left sidebar has tabs for Zone Readiness, Pool Readiness, DFS Readiness (which is selected), and Policy Readiness. The main area displays 'DR Failover Operations' with two entries:

	Name	SyncIQ Policy	Source	Destination	Last Successful Readiness Check	DR Failover Status
Sync IQ Policy	System-Zone-DFS_mir...	System-Zone-DFS_mi...	Cluster2-7201	prod-cluster-8	10/2/2017, 10:45:58 ...	OK
Eyeglass Configuration Replication	Cluster2-7201_System-Zone-DFS_mir...					OK

Below the table, there is a 'Generate SyncIQ Job Charts' button.

© Superna Inc

1.9. Operational Steps for Eyeglass Microsoft DFS Mode Failover

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

Operational Steps for Eyeglass Microsoft DFS Mode Failover

Overview

The Eyeglass Microsoft DFS Mode Failover solution for PowerScale OneFS fully automates execution of these steps:

1. Performs SyncIQ Failover to Secondary Cluster and sync's outstanding data.
2. **Creates and Renames** ALL SMB Shares protected by the SyncIQ policy to the Secondary Cluster.
3. **Deletes or Renames** ALL SMB shares on the source cluster to ensure DFS clients can not mount the cluster **See DFS Feature changes table ([DFS Feature Changes and Share names used on DFS synced Shares](#)) for version behaviour.**
4. **Creates** ALL quotas on the target cluster protected by the SyncIQ policy.
5. **Deletes** ALL quotas on the source cluster to ensure failback and SyncIQ operations from Secondary cluster to primary.
6. **Copies** SyncIQ schedule to Secondary cluster.
7. Runs Sync Prep on Source cluster to become the target of SyncIQ policy replication.

Procedure for Running Eyeglass Microsoft DFS Mode Failover

IMPORTANT: Making any changes to the SyncIQ Policies or related Eyeglass Configuration Replication Jobs being failed over during the failover may result in unexpected results.

IMPORTANT: Eyeglass Assisted Failover has a 1 hour timeout on each failover step. Any step which is not completed within this timeout period will cause the failover to fail.

IMPORTANT: Deleting configuration data (shares, exports, quotas) or modifying Share name or NFS Alias name or NFS Export path on the target cluster before failing over without running Eyeglass Configuration Replication will incorrectly result in the object being deleted on the source cluster after failover. You must run Eyeglass configuration replication before the failover OR select the Config Sync checkbox on failover to prevent this from happening.

For detailed steps consult the failover guide table [here](#).

For detailed steps on execution and monitoring consult the [How to Monitor the Eyeglass Assisted Failover](#)

NOTE: On completion of the failover the DFS clients will no longer have a share on the primary cluster and will use the referral list to select a new mount which will be an active share on the target cluster.

1.10. Post Eyeglass Microsoft DFS Mode Failover Manual Steps for NFS Exports

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

Post Eyeglass Microsoft DFS Mode Failover Manual Steps for NFS Exports

For the case where the SyncIQ Policies that were failed over are also protecting NFS exports, manual steps, or post failover scripting is required to update NFS client mounts.

Please refer to the Eyeglass Admin Guide [Script Engine Overview](#).

© Superna Inc

1.11. Post Eyeglass Microsoft DFS Mode Failover Checklist

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

Post Eyeglass Microsoft DFS Mode Failover Checklist

After a DFS failover we recommend the following steps to confirm everything was successful.

IMPORTANT: If the failover was done with the **Controlled failover** option unchecked (a real DR event failover), this means the share rename step on source was not executed. The source cluster should **NOT** be allowed to be reachable on the network until the shares are renamed using onefs UI OR DFS referrals are edited to disable the folder target pointing to the source cluster. Consult Failover Recovery Guide.

1. Open DR Dashboard and select DFS Readiness tab.
2. Make sure the two jobs are shown 1) one for the Mirror policy which should be active and 2) for the original policy which should be disabled now.

DR Dashboard						
Zone Readiness	Name	SyncIQ Policy	Source	Destination	Last Successful Readiness Check	DR Failover Status
Pool Readiness	<input type="checkbox"/> prod-cluster-B_System-Zone-DFS	System-Zone-DFS	prod-cluster-B	Cluster2-7201	10/3/2017, 8:10:44 AM	FAILED_OVER
DFS Readiness	<input type="checkbox"/> Cluster2-7201_System-Zone-DFS_mirror	System-Zone-DFS_mirror	Cluster2-7201	prod-cluster-B	10/3/2017, 8:10:47 AM	OK
Policy Readiness						
DR Testing						

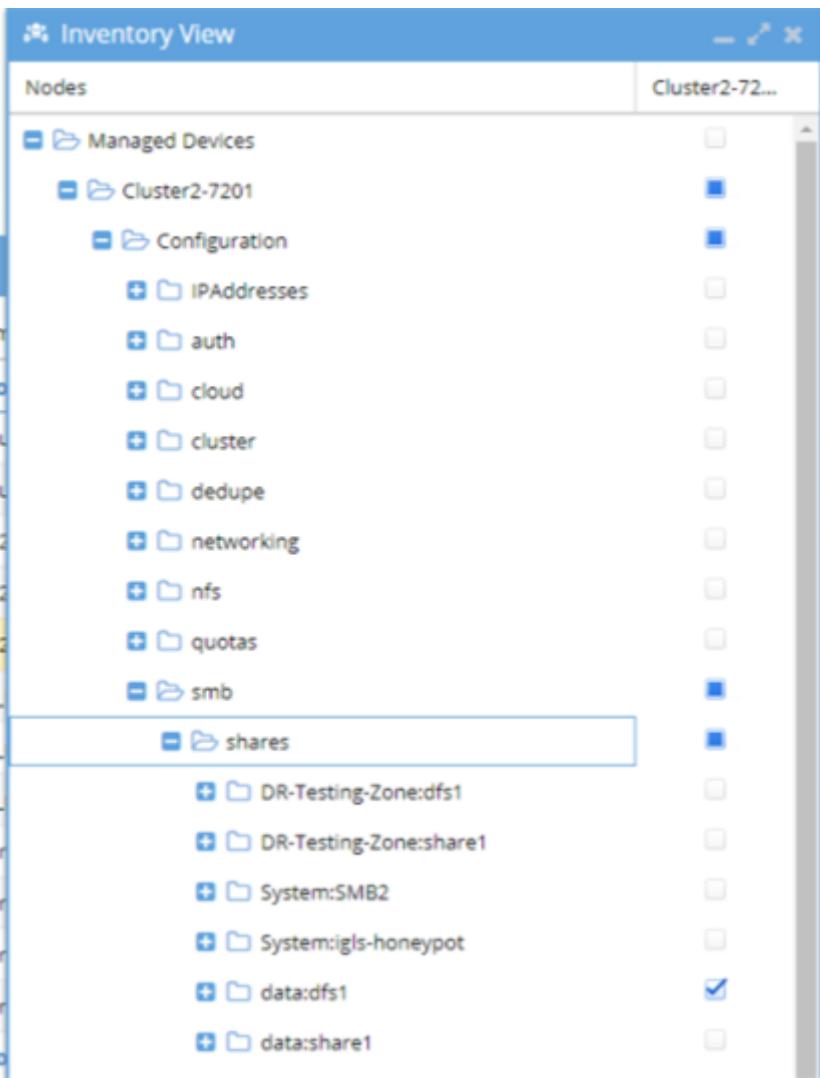
[Generate SyncIQ Job Charts](#)

3. Login to the source cluster (the one that you failed over from) and verify that all shares that are part of the failover are renamed with DFS prefix on the cluster. (See DFS Feature changes table for read-only share version behaviour).
1. You can get a list of shares that should have been failed over using this procedure.
2. Open the jobs window.
3. Select the DFS Mode SyncIQ policy job and select the check box then bulk actions and edit configuration option.

Jobs					
Job Definitions	Job Name	Policy	Type	Last Run Date	State
Running Jobs					
	prod-cluster-8_data-zone-dfs	data-zone-dfs	AUTOOFS	n/a	Policy Discrepancy
	prod-cluster-8_System-Zone-DFS	System-Zone-DFS	AUTOOFS	10/2/2017, 10:30:34 AM	Policy Discrepancy
<input checked="" type="checkbox"/>	Cluster2-7201_System-Zone-DFS_mir...	System-Zone-D...	AUTOOFS	10/3/2017, 8:10:20 AM	OK
	prod-8_marketing-nfs-policy	marketing-nfs-p...	AUTOOFS	n/a	Policy Discrepancy
Configuration Replication: DFS mode (AUTOMATIC)					
	Cluster2-7201_Eyeglass-DR-Testing	Eyeglass-DR-Tes...	DRTESTING	10/3/2017, 8:10:21 AM	OK
Disaster Recovery Testing (AUTOMATIC)					
	prod-cluster-8_data-zone-dfs_FILESYS...	data-zone-dfs	FILESYSTEM	10/2/2017, 5:40:32 AM	Policy Discrepancy
	prod-cluster-8_System-Zone-DFS_FILE...	System-Zone-DFS	FILESYSTEM	10/2/2017, 10:30:35 AM	Policy Discrepancy
	prod-cluster-8_data-zone-shares_FILE...	data-zone-shares	FILESYSTEM	9/26/2017, 6:35:27 AM	Policy Discrepancy
	prod-cluster-8_data-zone-nfsdata_FILE...	data-zone-nfsda...	FILESYSTEM	9/28/2017, 7:10:32 AM	Policy Discrepancy
	Cluster2-7201_System-Zone-DFS_mir...	System-Zone-D...	FILESYSTEM	10/3/2017, 8:10:21 AM	OK
	Cluster2-7201_data-zone-nfsdata_mir...	data-zone-nfsda...	FILESYSTEM	10/3/2017, 8:10:21 AM	OK
<input checked="" type="checkbox"/>	Show Disabled Jobs	1 item(s) selected			
<div style="border: 1px solid #ccc; padding: 5px; display: inline-block;"> Select a bulk action ▾ <ul style="list-style-type: none"> Edit Configuration(s) Run Now Enable/Disable Microsoft DFS Set Schedule <input checked="" type="checkbox"/> Enable/Disable Delete Enable/Disable Skip Config Replication </div>					

4. Expand the cluster configuration to look in the SMB folder.

All Shares that were part of the failover will have a blue check mark next to them.



5. For Eyeglass Release 1.5 and higher, check the Failover Log.

Look for the lines:

1. INFO Renamed <number> shares on source cluster.
2. INFO Renamed <number> shares on target cluster.
3. ERROR Policy: <policy name> Step: "Renaming shares on source and target clusters" Result: FAILURE: Failed to rename SMB shares for policy <policy name>. Failover not complete.

4. ERROR Cluster: <cluster name> Step: "Renaming shares on source and target clusters" Result:
FAILURE: Renaming shares failed.

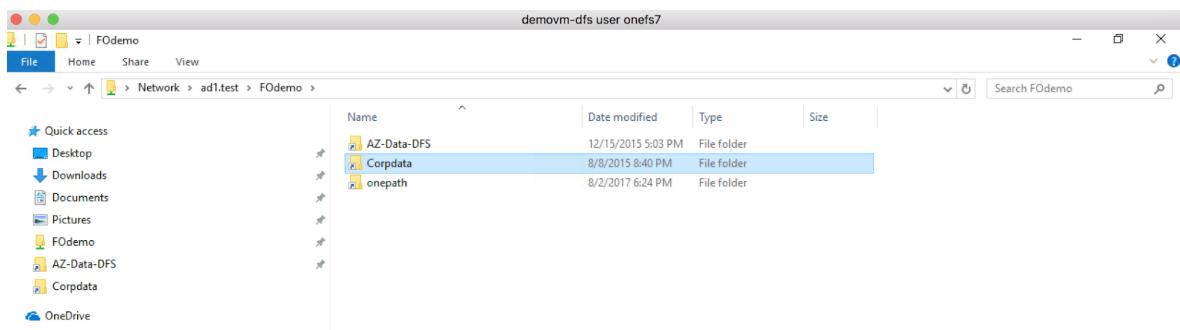
These numbers should be the same and should match the total number of shares protected by the SyncIQ Policy that was failed over, and there should be no ERROR related to share renaming. If this is not the case:

- Search on the failover **source cluster** for any shares protected by the SyncIQ Policy failed over that do not have "igls-dfs" prefix. Manually update the share name to add the "igls-dfs" prefix.
- Search on the failover **target cluster** for any shares protected by the SyncIQ Policy failed over that still have "igls-dfs" prefix. Manually update the share name to remove the "igls-dfs" prefix.

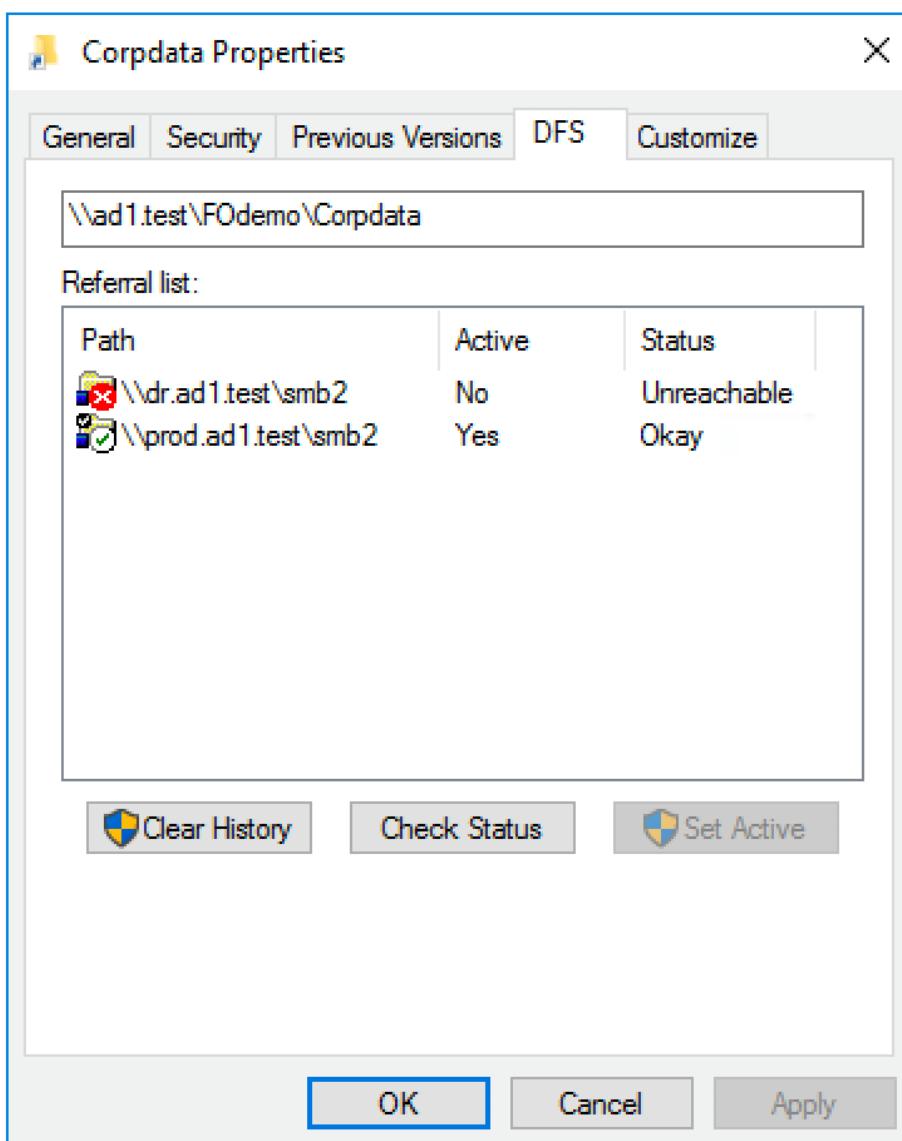
Procedure for Checking your SMB Clients Post DFS Failover:

Use this procedure to check that clients have picked up the change from one cluster to another.

1. From a Windows machine in the domain with access to the DFS mount.
2. Open with explorer \\<domain name>\<dfs root name>\



1. Right click a DFS folder involved in the failover select the DFS tab. Click check status and verify the active path is ok and active. Click the 2nd path and verify its unreachable.



1.12. Advanced DFS mode Setup with Access Zones

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

Advanced DFS mode Setup with Access Zones

One using DFS mode policies inside an Access zone that you intend to use Access Zone failover for other policies in the Access Zone. The following additional configuration must be done to ensure the SmartConnect Zones used for DFS folder UNC paths do not get failed over or require Eyeglass hints.

Configuration Steps

1. Create dedicated IP pool for DFS protected data and make the IP pool a member of the Access Zone.
2. Create shares for DFS UNC paths.
3. Create SyncIQ policy.
4. Discover the policy with Eyeglass and edit the job in the jobs definition window to Enable DFS mode.
1. Run the DFS mode job in Eyeglass to make sure it completes.
5. Create Eyeglass hints to ignore the DFS IP pool using a hint as a IP pool alias named "igls-ignore".
6. Create the hint on both IP pools: source and target cluster.

7. It should look like below after Zone Readiness job runs (it can be run with Run now option to update the DR dashboard).

Network Mapping for Cluster-1-7201 > Cluster2-7201 Zone: data

Cluster-1-7201

subnet0:dfsdata
Smart Connect Zone Name: dfsdata.ad1.test
Smart Connect Aliases:

- igls-ignore
- dfsdata-dr.ad1.test

Subnet: subnet0
SSIP: 172.31.1.200

subnet0:userdata
Smart Connect Zone Name: userdata.ad1.test
Smart Connect Aliases:

- somezone
- igls-user-prod

Subnet: subnet0
SSIP: 172.31.1.200

Cluster2-7201

Pool will not be failed over because of igls-ignore alias.

subnet0:userdata
Smart Connect Zone Name: igls-original-somezone
Smart Connect Aliases:

- igls-user-prod

Subnet: subnet0
SSIP: 172.31.1.201

IGNORE

Copyright Superna LLC 2017

© Superna Inc

2. Eyeglass Access Zone Failover Guide

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

- [Introduction to this Guide](#)
- [Automated SMB Client Switch Testing Matrix](#)
- [How to Setup and Configure Access Zone - Overview Video](#)
- [Requirements for Eyeglass Assisted Access Zone Failover](#)
- [Unsupported Data Replication Topology](#)
- [Overlapping Access Zone Failover Supported Configurations](#)
- [Recommendations for Eyeglass Assisted Access Zone Failover](#)
- [Preparing your Clusters for Eyeglass Assisted Access Zone Failover](#)
- [PowerScale Administration for Clusters Configured for Eyeglass Assisted Access Zone Failover](#)
- [Failover Planning and Checklist](#)
- [Monitoring DR Readiness for Eyeglass Assisted Failover](#)
- [Operational Steps for Eyeglass Assisted Access Zone Failover](#)
- [Post Access Zone Failover Steps](#)
- [Post Access Zone Failover Checklist](#)
- [IP Pool Failover](#)
- [Fan-In IP Pool Failover](#)
- [Fan-Out IP Pool Failover](#)
- [How to Configure Access Zone DNS Dual Delegation](#)

- Controlled Failover Option Results Summary
- How to Configure Delegation of Cluster Machine Accounts with Active Directory Users and Computers Snapin

© Superna Inc

2.1. Introduction to this Guide

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

Introduction to this Guide

The purpose of this document is to act as a guide to Access Zone Failover.

Overview

Access Zones keep all configuration data separated including authentication providers, this also provides segmentation for business units or application tiers. Access Zones allow data, configuration, IP pools and SmartConnect Zones to be associated to DNS delegations for data access.

With Access Zone failover, all SyncIQ policies, all Access Zones, all SMB Shares, all NFS Exports and all Quotas are failed over as a unit. Eyeglass can then alarm, detect and correct SPN entries automatically without the user being required to know in advance which SmartConnect zones match which SPN share mounts.

Shared filesystems with NFS and SMB that MUST failover together will benefit from Access Zone failover. If NFS only failover is required per SyncIQ policy, failover will meet your needs with less pre-configuration. Since NFS requires unmount and remount of data, it's just as easy to change the mount path name.

What's New with Access Zone Failover

Release 2.0

This release adds new Access Zone failover granularity option with IP pool failover. This new failover mode allows an Access Zone to support active data on 2 clusters within the same Access Zone.

The IP pool and all SmartConnect names or aliases on the pool to be selected for failover independently of other IP pools in the Access Zone. All of the same validations for DR Readiness now operate independently on IP pools.

DR Assistant has new SmartConnect/Pool failover option. One or more pools can be selected for failover.

Upgrading to this new failover mode from Access Zone failover configuration, requires mapping SyncIQ policies to the pools that they protect. Once all policies in the zone are mapped. The feature can be used.

DR Dashboard has a new tab showing IP pool failover readiness and configuration.

The failover logic and all previous requirements are exactly the same including dual delegation, SPN delegation, igls hints.

New requirement is that SyncIQ policies are mapped in DR dashboard Zone Readiness UI to one and ONLY one IP pool. A

pool can have more than one SyncIQ policy mapped to a single Pool. SyncIQ policies that are mapped should be protecting ONLY the SmartConnect names assigned to the pool. When the pool is failed over only the mapped policies will be selected for failover.

NOTE:

1. Zone Readiness will now validate each pool's readiness for failover independently.
2. Policies must be mapped to at least 1 pool to be failed over
3. The entire Access Zone can still be failed over in DR Assistant by selecting the zone and all pools within the zone will failover.
4. Multi site not supported in this release
5. Fan in topology for a shared DR cluster target with overlapping source cluster Access Zones is supported.

Release 1.8

Time skew validation added to check time differences between nodes and between Eyeglass and the clusters. This validation has an acceptable range that will not trigger a warning. This validation verifies that SyncIQ operations between clusters are not affected due to differences in the times between clusters. It runs during Zone Readiness and checks all time on each node in all clusters.

Zone Readiness for Cluster2-7201 > prod-cluster-8 Zone: System	
name	status
Zone Readiness Statuses	OK
OneFS SyncIQ Readiness	OK
Eyeglass Configuration Replication Readiness	OK
SPN Readiness	OK
SmartConnect/IP Pool Settings and Mappings Readin...	OK
Target Cluster Reachability	OK
Date-Time Validation	OK
Nodes Validation	OK
Eyeglass & Clusters Validation	OK
FQDN Alias Validation	OK
Additional Status Information	
Click on a row to view additional information.	

Release 1.6

As of release 1.6 and beyond multi site replication is now possible along 3 site fully automated failover.

This new option allows A to B and A to C site replication from the same source Access Zone. This provides a DR choice to failover to B site or C site depending on the DR event. In addition, this allows for failover and fallback operations from C back to A or B back to A site.

This feature will use triple site DNS delegation of SmartConnect zones and extends dual delegation to 3 NS records and allows the

DNS name space, to failover from site A to B or C and back if needed.

This extends SyncIQ to allow the highest, data and site location availability along with flexibility of “one button failover” to more than a single site.

Release 1.7

As of release 1.7 and beyond Access Zone Failover will restrict the number of parallel Job requests to the PowerScale cluster for the Run SyncIQ Policy data sync step based on cluster version:

- OneFS 7.2 - 5 parallel job requests (OneFS 7.x cluster have a limit of 5 concurrent policies). Eyeglass will monitor the progress for each Job and submit a new request as previously submitted requests are completed.
- OneFS 8 - parallel job requests limit based on Eyeglass appliance configuration (default 10)

Multi Site Access Zone Failover Guide

For detailed instructions on how to configure multi site automated 3 site failover see [Multi Site Failover Guide](#). This guide should be read and understood and implemented first with 2nd site, and then add 3rd site Access Zone failover setup.

Based on extensive testing for safe failovers, make writeable and resync prep are serialized steps. New parallel flag allows this step to run 10 threads wide but does not stop on failures. Use with caution see Failover Design Guide on configuration.

© Superna Inc

2.2. Automated SMB Client Switch Testing Matrix

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

- [Access Zone and IP Pool SMB client Testing Matrix](#)
- [Access Zone mixed DFS and none DFS client Testing Matrix](#)

Access Zone and IP Pool SMB client Testing Matrix

Operating System Version	SMB Protocol	OneFS Version
Windows 7 Enterprise	SMB2 DFS mounted with only 1 referral path	8.1.x.x
Windows 10 Enterprise	SMB 2, SMB 3, DFS mounted with only 1 referral path	8.0.0.x, 8.1.x.x
Window Server 2016 Essentials	SMB 2, SMB 3	8.0.0.x, 8.1.x.x
Windows Server 2012 R2	SMB 2, SMB 3	8.0.0.x, 8.1.x.x
Centos Release: CentOS Linux release 7.6.1810 (Core) Samba version: Samba version 4.8.3 Configured for Microsoft DFS mounts and dual referral paths configured	SMB2 , Basic support for SMB3 is included in Samba 4.0.0 and later	8.0.0.x, 8.1.x.x

NOTE: If not listed it has not been tested. We suggest testing specific combinations not listed in the table above. Support will not be able to test configurations.

Access Zone mixed DFS and none DFS client Testing Matrix

Operating System Version	SMB Protocol	OneFS Version
Windows 10 Pro or Enterprise	SMB3	8.1.x.x

© Superna Inc

2.3. How to Setup and Configure Access Zone - Overview Video

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

How to Setup and Configure Access Zone - Overview Video

The following video provides an overview tutorial on how to setup and configure Eyeglass Access Zone:

© Superna Inc

2.4. Requirements for Eyeglass Assisted Access Zone Failover

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

Requirements for Eyeglass Assisted Access Zone Failover

- Cluster Version Requirements - May Block Failover
- Access Zone Requirements - Blocks Failover
- SyncIQ Policy Requirements - Blocks Failover
- DFS Mode Requirements
- Shares / Exports / NFS Alias Requirements
- Eyeglass SmartConnect Requirements - Blocks Failover
- Eyeglass Failover Mapping Hints Requirements - Blocks Failover
- Failover Target Cluster Requirements - Block Failover
- Eyeglass Quota Job Requirements - Will not Block Failover

The requirements in this section must be met in order to initiate an Eyeglass Access Zone Assisted Failover. Failure to meet some of these conditions may **block** the Access Zone Failover.

Cluster Version Requirements - May Block Failover

Clusters participating in an Access Zone Failover must be running the supported PowerScale Cluster version for this feature. See the

Feature Release Compatibility matrix in the Eyeglass Release notes specific to your Eyeglass version found [here](#).

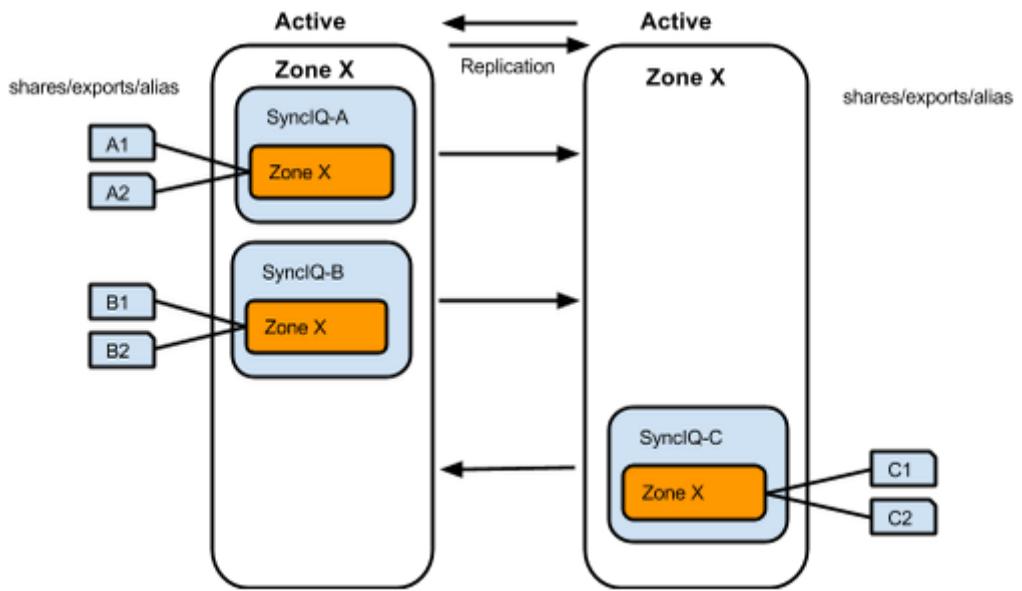
Access Zone Requirements - Blocks Failover

For an Access Zone failover with Eyeglass, the Access Zone must meet the following requirements:

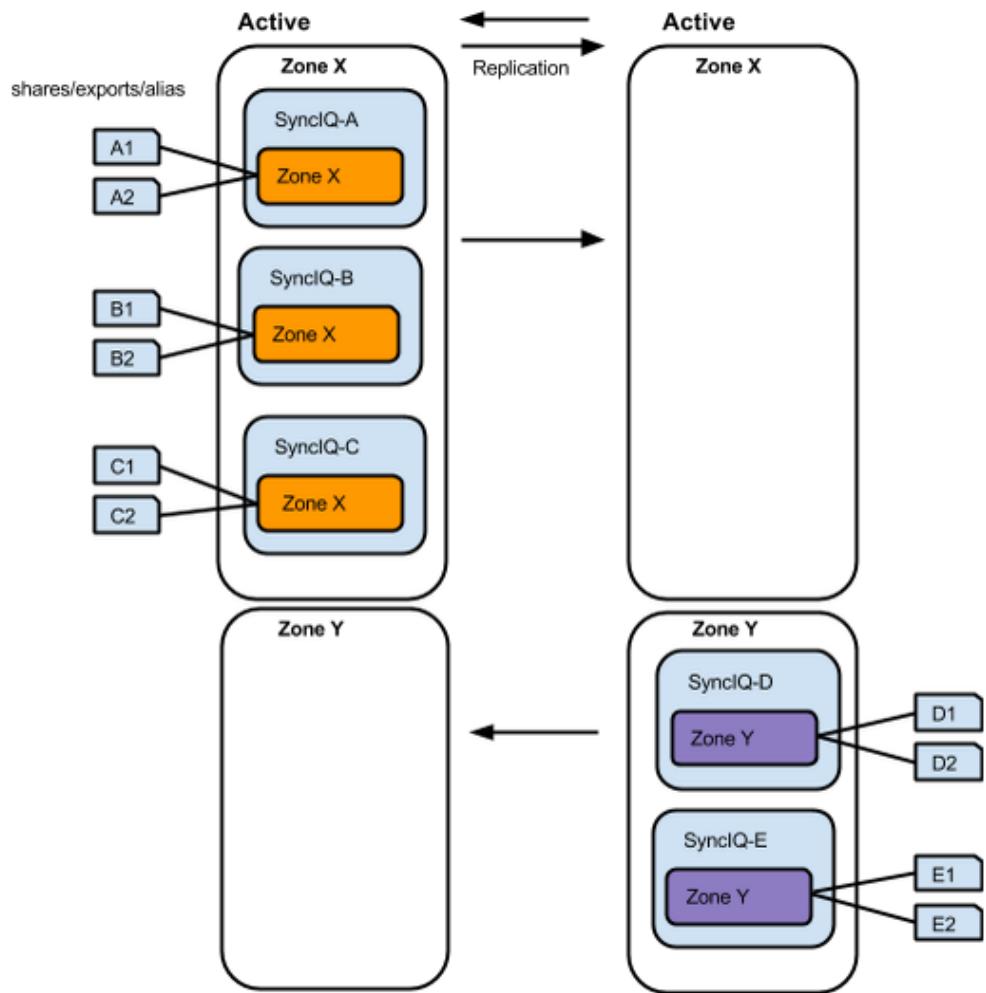
Release	System Access Zone Failover only Zone on the cluster	System Access Zone and Other non system Access Zones with failover
1.4 update 1 to 1.5.4	Supported	Not supported
1.6.0 >	Supported	Supported

- Access Zone must exist on Target Cluster with same name and authentication providers (**NOTE:** Access Zone Sync Jobs in Eyeglass can be used to create Access Zones automatically)
- Access Zone must be associated with at least **one** subnet IP pool
- Access Zone must be associated with one or more SyncIQ policies
- SyncIQ policies associated to Access Zone by path - the SyncIQ policy source path must be the same or below the Access Zone base path
- In Active-Active data replication topology, there is a dedicated Access Zone for each replication direction.
- Failover with Eyeglass of an Access Zone for Active-Active data replication topology that is shared by SyncIQ Policies on both clusters is **NOT SUPPORTED** as there is no “partial” fail back path to only fallback the subset of the SyncIQ policies that were originally failed over.

Example: Unsupported



Example: Supported



SyncIQ Policy Requirements - Blocks Failover

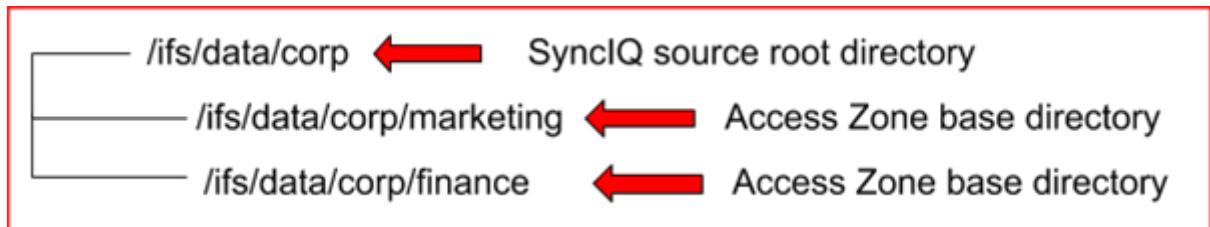
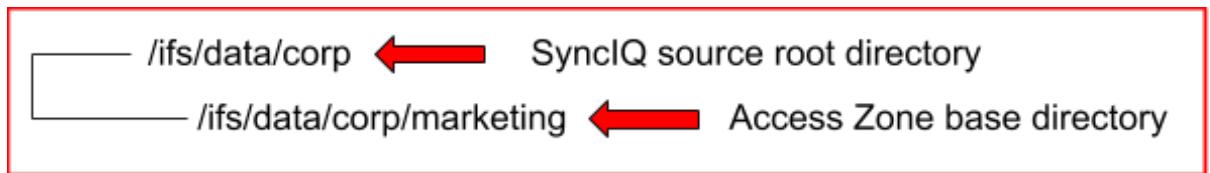
For an Access Zone failover with Eyeglass, it is required that the SyncIQ policy(s) identified as part of the Access Zone (based on Access Zone base path and SyncIQ Policy source path) meet the following requirements:

- All OneFS SyncIQ Policy(s) must have the same Target Cluster provisioned (**release 1.5.4 and earlier**)
- In Release 1.6 and later multi site replication allows an Access Zone policy to have a 3rd cluster target for multi site failover

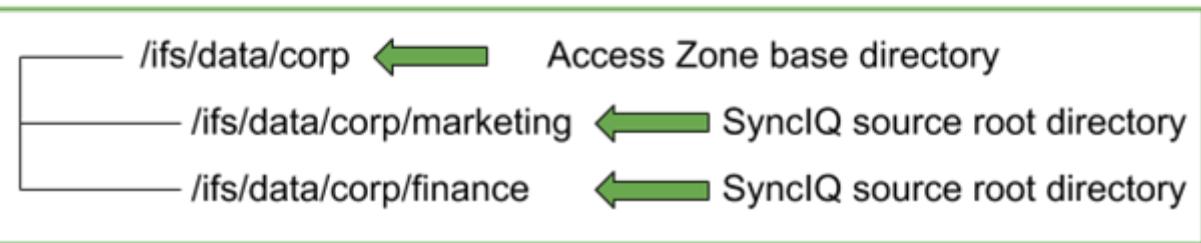
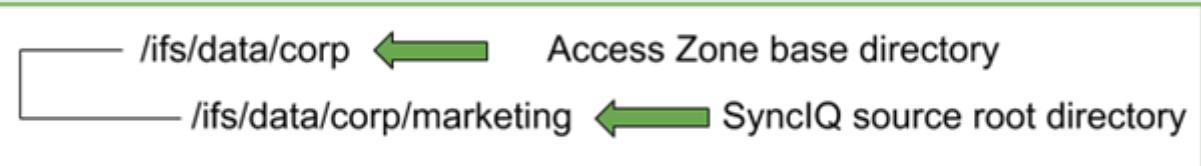
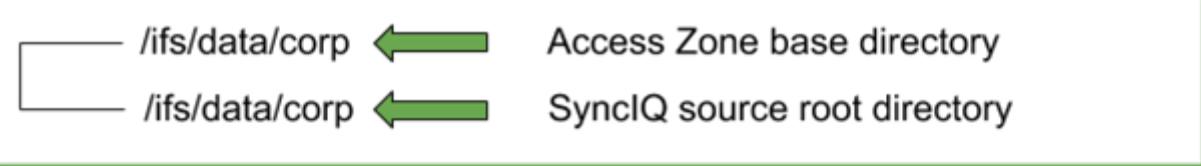
requirements. This can be a simple policy based failover, OR fully automated Access Zone Multi site failover policy if configured.

- OneFS SyncIQ Policy(s) Target Host SmartConnect Zone must be associated with a pool that is NOT going to be failed over
- SyncIQ Policy(s) source root directory must be at or below the Access Zone Base Directory

Example: Unsupported



Example: Supported



DFS Mode Requirements

DFS mode does **not** require SmartConnect zone names to failover.

If you have DFS mode SyncIQ policies that fall within the Access Zone root path, it is required to have a dedicated subnet:pool with SmartConnect Zones that are used for UNC paths for DFS folder targets. Please refer to the Eyeglass Microsoft DFS Mode Failover Guide [here](#) for more details on DFS mode setup and requirements.

Note:

DFS enabled policies that also protect NFS exports will require separate SmartConnect Zone for NFS export data access to take advantage of Access Zone automation or manual steps / post failover scripting to update NFS client mounts.

Shares / Exports / NFS Alias Requirements

Exports with multiple paths, where all paths are not associated with the same SyncIQ Policy, are not supported for Access Zone failover. This will result in a “A unique readiness data ...” error in Zone Readiness and the Overall Status cannot be displayed.

Eyeglass SmartConnect Requirements - Blocks Failover

- For an Access Zone failover with Eyeglass, the SmartConnect Zone FQDN must not exceed 50 characters.
- The Pool “SmartConnect service subnet” must be provisioned with the same Subnet that the Pool was created in.

Eyeglass Failover Mapping Hints Requirements - Blocks Failover

For an Access Zone failover with Eyeglass, a mapping between subnet pools on the Source and Target cluster is required to ensure that data is accessed from the correct SmartConnect Zone IP and node pool on the Target cluster after failover.

This mapping is used to failover SmartConnect zones from IP to IP pool and for fallback.

The Eyeglass mapping hints have the following requirements:

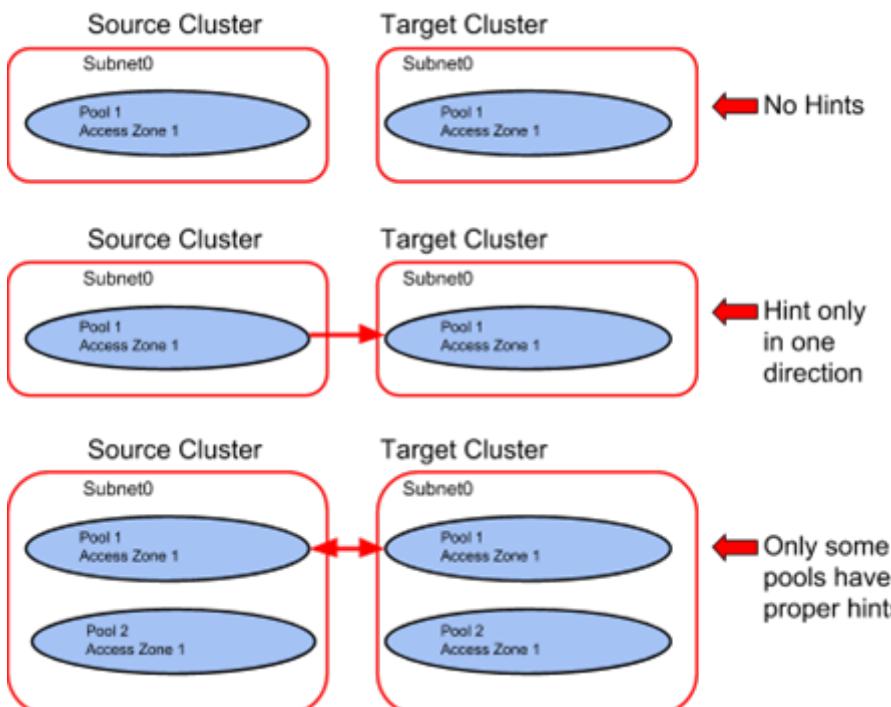
(NOTE: This is a one time setup process but needs to be repeated if new IP pools are created).

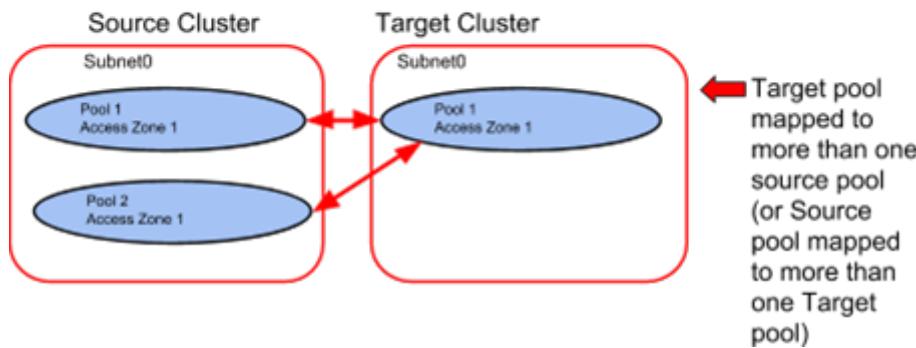
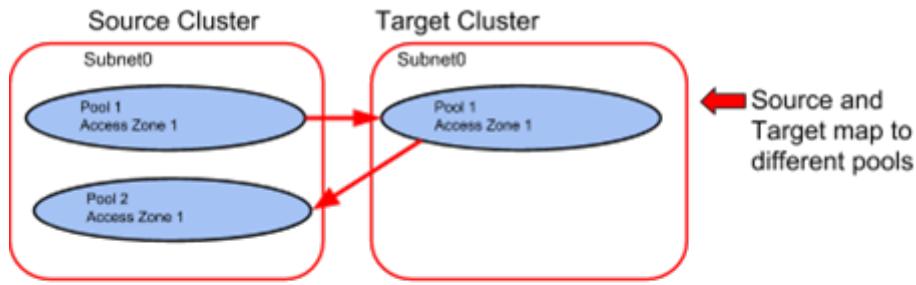
- Eyeglass Mapping Hints are simple SmartConnect aliases created with ISI or UI to map IP Pools and SmartConnect Name failover mapping between pairs of clusters. See mapping hints examples section.
- Every subnet IP pool associated with the Access Zone being failed over is required to have a mapping **PER IP pool** in the Access Zone. DR Dashboard will raise an error if mapping hints are not found or incorrectly created.
- Eyeglass mapping hints on both Source and Target cluster IP pools are created in **UNIQUE** pairs. See IGLS mapping hints section for syntax and examples.
- Incorrectly mapped pools will be alarmed in Access Zone Readiness in the DR Dashboard Access Zone Readiness Tab.
- DFS mode does not require SmartConnect zone names to failover. If you have DFS mode SyncIQ policies in the Access Zone, a dedicated subnet:pool with Eyeglass igls-ignore hint applied is required to retain SmartConnect zones on source and target clusters.
- The Subnet Pool which is used in the SyncIQ Policy Restrict Source Nodes option must NOT have an Eyeglass mapping hint and must have an igls-Ignore Hint applied (**NOTE: If misconfigured the SmartConnect zone used by SyncIQ would failover and would impact fallback operations and SyncIQ replication.)**
- The Subnet Pool which is associated with SyncIQ Policy Target Host property of a SyncIQ policy. This SmartConnect zone must **NOT** have an Eyeglass mapping hint and must have an Ignore Hint applied. This is the pool on the

target cluster used for SyncIQ replication. **NOTE: Failure to apply this hint can affect fallback operations and SyncIQ replication if the SmartConnect name is failed over by Eyeglass.** See mapping hints examples section.

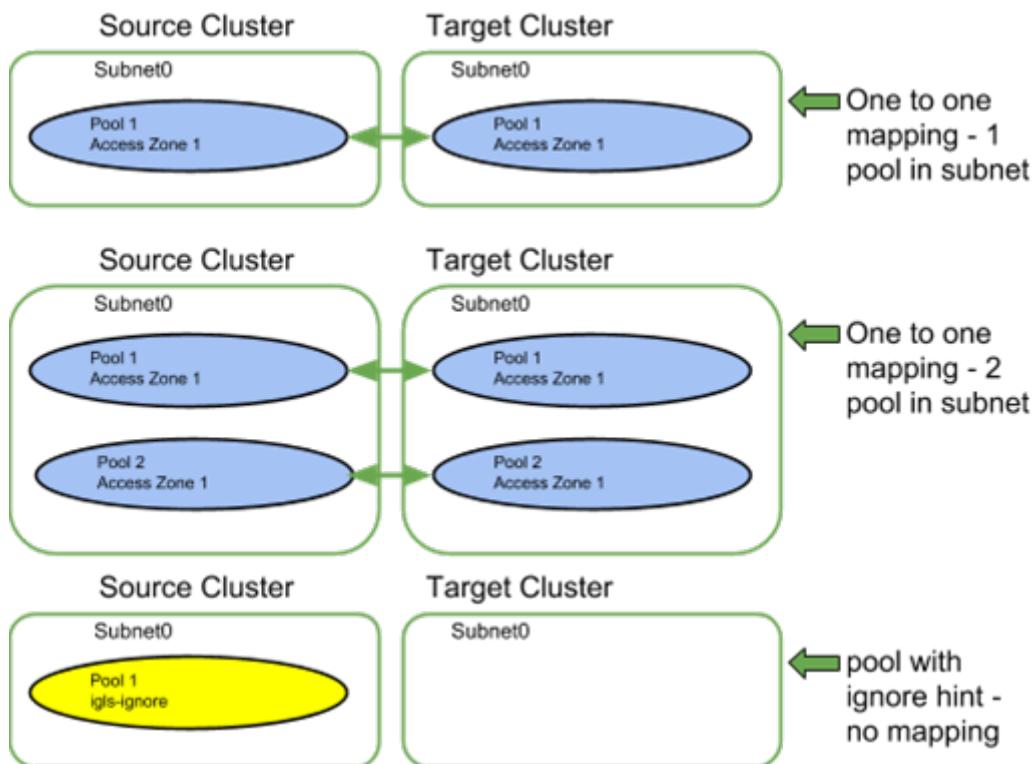
- Ignore hints are simply an alias with name of "igls-ignore" **NOTE: best practise to ensure unique hints by using a naming format that uses cluster name example igls-ignore-clustername.** This allows Eyeglass to match on igls-ignore while allowing the hint to be unique to avoid SPN collision in the Active Directory if the SmartConnect alias is added to AD with check or repair ISI command on the cluster. Since Hints are SmartConnect aliases they can be inserted to AD machine account but are not required for kerberos authentication since they are not used to mount shares. See [Active Directory Machine Account Service Principal Name \(SPN\) Delegation](#) in this document for detail.

Example: Unsupported





Example: Supported



Failover Target Cluster Requirements - Block Failover

For an Access Zone failover with Eyeglass, it is required that the PowerScale Cluster, that is the target of the failover, be IP reachable by Eyeglass with the required ports open.

Eyeglass Quota Job Requirements - Will not Block Failover

For an Access Zone failover with Eyeglass, there are no Eyeglass Quota Job state requirements. Quotas will be failed over whether Eyeglass Quota Job is in Enabled or Disabled state.

© Superna Inc

2.5. Unsupported Data Replication Topology

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

Unsupported Data Replication Topology

Replication topology with shares or NFS alias with the same name on both clusters, and protected by different SyncIQ policies, is not supported. Configuration Replication will overwrite the path on one cluster as the share / alias. It would attempt to have 2 SyncIQ policy on the same cluster with the same source path and failover will not succeed. **Note: This is an invalid DR configuration, this configuration means duplicate shares point to different data. This is not a good DR configuration and it will not be possible with or without Eyeglass to failover successfully.**

© Superna Inc

2.6. Overlapping Access Zone Failover Supported Configurations

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

- Requirements for Overlapping Access Zones
- Example 1
- Example 2
- How to Failover Overlapping Access Zones

Overlapping Access Zones has been added to 2.5.3 and will only be available in 2.5.5. **Release 2.5.6 will not support this configuration any longer and formal support for Overlapping Zone failover will be assessed in a future release.** NOTE: the following requirements that must be met to be supported. **If not listed below then it will not be supported.**

Requirements for Overlapping Access Zones

1. Overlapping Access Zones is defined as zones sharing the same base path example /ifs/data/zones has 2 or more access zones all configured with this path.

2. At least one syncIQ policy at or below the access zone path must exist.
3. If more than one SyncIQ policy exists under the access zone base path and shares or exports exist at or below each syncIQ path, then the policy will be auto assigned to the access zone where the share/export was created.
This cannot be changed.
4. NOT Supported:

 - nested Access Zones, example /ifs/data/zone1 and /ifs/data/zone1/childzone2
 - overlapping Access Zones with System Access Zone on /ifs

Example 1

2 different AZ with same base path

ZoneA--> /ifs/data/zonea

ZoneAA --> /ifs/data/zonea

3 different synciq policies that fall under the Access Zone path

pol1 --> /ifs/data/zonea/pol1

pol2 --> /ifs/data/zonea/pol2

pol3 --> /ifs/data/zonea/pol3

Example 2

2 different AZ with same base path

ZoneA--> /ifs/data/zonea

ZoneAA --> /ifs/data/zonea

1 synciq policy falls under the Access Zone path

pol1 --> /ifs/data/zonea/

How to Failover Overlapping Access Zones

Failover is 2 phase process. The first phase will failover each access zone that overlaps and this will complete the Networking failover (DNS, and SPN's) for each pool that is configured to failover. Phase 2 will failover the data.

1. Phase 1

- a. Enable parallel failover mode (check with support on how to enable this mode to allow concurrent failovers)
- b. Select Access Zone 1 and start a failover (NOTE: no SyncIQ policies will be failed over during this phase , this is expected)

- c. Restart DR Assistant and select each of the remaining access zones that overlap and start the failover one at a time.
 - d. Wait for the Access Zone failovers to complete
 - e. Submit the Eyeglass failover log to support to verify it was successful.
 - f. Repeat for each Overlapping access zone that was failed over.
- 2. Phase 2 - Do not start until Phase 1 is complete and successful**
- a. Open DR Assistant and select All SyncIQ policies that belong to the Overlapping Access Zones that were failed over in Phase 1.
 - b. Wait for the syncIQ failovers to complete
 - c. Submit the Eyeglass failover log to support to verify it was successful.
- 3. Done - Test Data Access to all data**

2.7. Recommendations for Eyeglass Assisted Access Zone Failover

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

Recommendations for Eyeglass Assisted Access Zone Failover

The conditions outlined in the following section are highly recommended to ensure that all automated Access Zone failover steps can be completed. If anyone of these conditions is not met it will result in a Warning.

- **Warnings Will Not block** Eyeglass Assisted Access Zone Failover, but potentially post failover will require additional manual steps to complete the failover.
- **Errors Will block** Eyeglass Assisted Access Zone Failover.

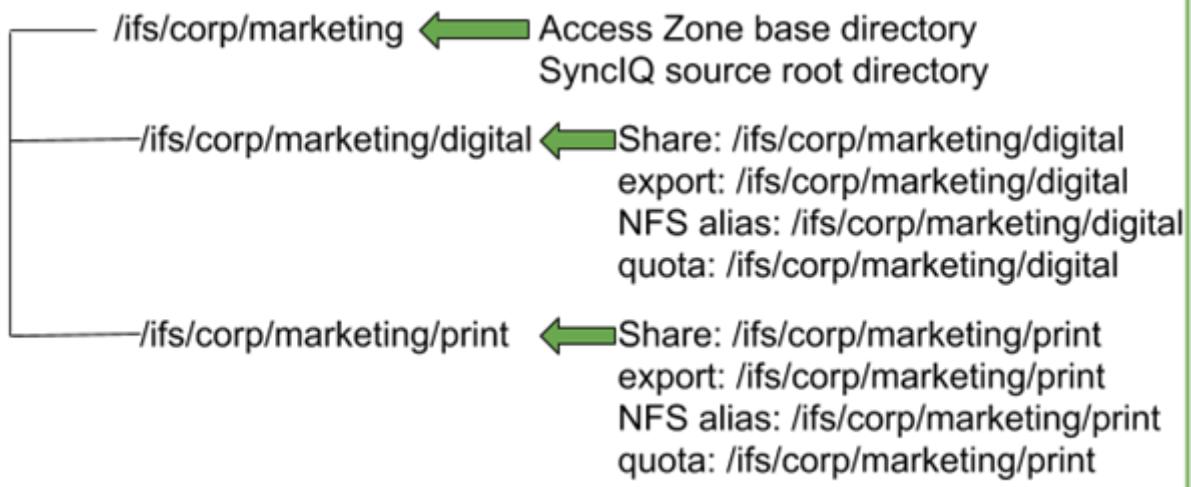
Shares / Exports / NFS Alias Recommendations

- All shares, exports and alias should be created in the Access Zone that is being failed over. It is not supported to have shares, exports and alias with a path that is outside (higher in the file system) than the Access Zone base path.
- **Impact - Data Access Outage:** The policy will not be selected for Failover based on path matching the Access Zone base path resulting in data that will NOT be failed over with the Access Zone.
- Access Zones Readiness in the DR Dashboard show which policies have been matched to the Access Zone and should

be verified to ensure all expected SyncIQ policies are present in the Access Zone Readiness.

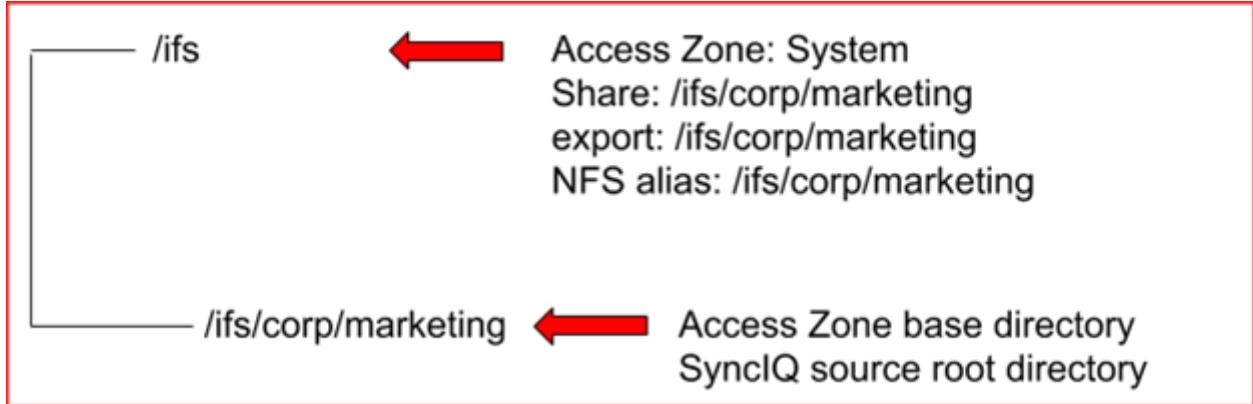
Example: Share / Export / NFS Alias Configuration

RECOMMENDED

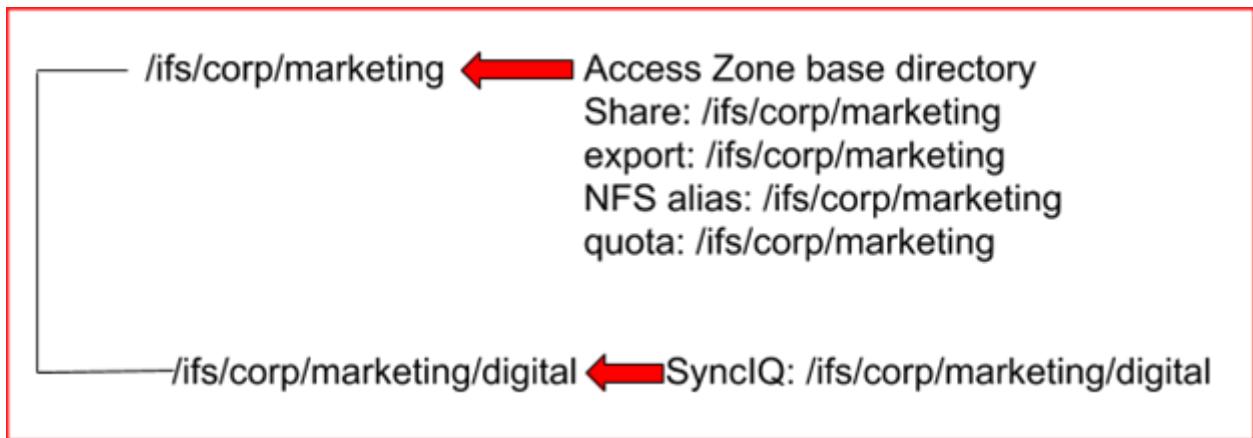


Example #1: Share / Export / NFS Alias Configuration NOT

RECOMMENDED



Example #2: Share / Export / NFS Alias Configuration NOT RECOMMENDED



- Eyeglass Configuration Replication Jobs for the SyncIQ Policies in the Access Zone being failed over should have been completed without error.
- **Impact - Data Access Outage:** Any missing or incorrect share / export / NFS alias information will prevent client access to data on the Target cluster. These configuration items will have to be corrected manually on the Target Cluster.

Service Principal Name Recommendations

For optimal Access Zone failover with Eyeglass where Access Zone contains SMB shares that are directly mounted using SmartConnect Zones, the following is recommended:

- Setup delegation for SPN add and delete to allow Eyeglass to automatically update SPNs based on SmartConnect changes made during failover ([How to Configure Delegation of Cluster Machine Accounts with Active Directory Users and Computers Snapin](#)).
- **Impact:** With no SPN updates, SMB share client authentication may not complete. SPNs will have to be updated manually for both the Source and Target cluster to enable Kerberos authentication again. NTLM fallback authentication should be verified with Active directory.

SynclIQ Policy Recommendations - Does not Block Failover

SynclIQ Policy last run should have been successful:

- OneFS SynclIQ Job(s) should have been run at least once.
- OneFS SynclIQ Job(s) for last run should have been successful.
- OneFS SynclIQ Job(s) should not be in a Paused or Cancelled state.
- **Impact:** depending on it's current status SynclIQ Policy MAY NOT be able to be run by Eyeglass assisted failover if the above recommendations have not been met . If it does not run, you will incur data loss during failover.
- Example 1: SynclIQ Policy has an error state. If it cannot be run from the OneFS, it will also not be able to run from Eyeglass.

- Example 2: SyncIQ Policy is paused. Eyeglass failover cannot RESUME a paused SyncIQ Policy - this must be resumed from OneFS.

You must investigate these errors and understand their impact to your failover solution.

PowerScale does not support SyncIQ Policy with excludes (or includes) for failover.

- **Impact:** Not a supported configuration for failback.

PowerScale best practices recommend that SyncIQ Policies utilize the Restrict Source Nodes option to control which nodes replicate between clusters.

- **Impact:** Subnet pool used for data replication is not controlled therefore, all nodes in the cluster can replicate data from all IP pools. This makes it hard to manage bandwidth and requires all nodes have access to the WAN.

Eyeglass failover will skip failover for any SyncIQ policies in the Access Zone which are in the disabled state.

- **Impact - Data Access Outage:** If SyncIQ Policies are disabled, the associated filesystem will be writeable on source and will **NOT** failover. Data on the source cluster will most likely will not be reachable by clients due to the fact that the networking and SmartConnect Zone required for data access will have been failed over, and the source SmartConnect zone is renamed to ensure clients can not mount it.

© Superna Inc

2.7.1. Mixed DFS and None DFS Solution

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

- [Overview](#)
- [Access Zone failover Configuration for Mixed Mode clients](#)
- [Eyeglass configuration](#)
- [What happens during failover](#)

Overview

A new solution with Access zone failover has been tested that addressed mixed DFS mount and none DFS mounts. Typically a single DFS referral path is used and DFS mode requires a 2nd referral path. This offers single policy failover granularity with DFS mode jobs in Eyeglass. This new solution allows single DFS referral paths to be used and retain the benefits of automatic client redirection during failover.

This solution simplifies scenario's with both DFS and non DFS mounted data and DFS is already configured with a single referral path.

Access Zone failover Configuration for Mixed Mode clients

1. The Access zone failover configuration requirements are the same for mixed mode. The IP pool used must be mapped for failover as described in this guide with aliases used as failover hints.
2. DFS configuration is unchanged using a smartconnect zone name in the zone and a single referral path
3. None DFS clients CAN use the same smarconnect name used by DFS clients or a different smartconnect name, both are supported
4. Eyeglass configuration
 - a. **Enable syncIQ jobs in the Eyeglass jobs icon for DFS mode for all data this is mounted over DFS**
 - b. **This will ensure SMB shares are renamed with igls-dfs on the target cluster**
 - c. **Any policy that is non DFS, can stay in the default configuration sync section of the Eyeglass jobs icon.**
5. What happens during failover
 - a. **The SMB shares are renamed which triggers DFS clients using a single referral path to check DNS for**

the smartconnect name which will now return a target cluster IP address using dual delegation smartconnet failover offered with Access zone failover.

- b. Non DFS client will need to remount**
- c. DFS clients will check DNS for new ip address and auto remount the DFS referral path and re-authenticate automatically.**

2.8. Preparing your Clusters for Eyeglass Assisted Access Zone Failover

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

Preparing your Clusters for Eyeglass Assisted Access Zone Failover

- Update PowerScale sudoer file for Eyeglass Service Account
- Active Directory Machine Account Service Principal Name (SPN) Delegation
- What is an SPN?
- What's the risk if I don't fix SPN's?
- Delegate SPN
- Configure Eyeglass Subnet IP Pool Mapping Hints
- Zone Aliases for Failover Overview
- When and how to NOT failover an IP pool SmartConnect name using hint: igls-ignore
- How to create Mapping Hints for IP pools between source and target clusters following best practise naming convention
- IGLS examples for ignore option on IP Pools
- OneFS 7 Example igls hints for ignore SyncIQ pool
- Example: Mapping with hints configured correctly will display the mapping
- OneFS 8 Example igls hints for ignore SyncIQ pool

- Hot-Cold Replication Topology Mapping Examples

The following steps described in this section are required to prepare your system for the Eyeglass Assisted Access Zone Failover:

- [How to - Delegation of Cluster Machine Accounts with Active Directory](#)

This is required to avoid Eyeglass requiring direct access Active Directory to synchronize the Service Principal Names (SPN) for production or DR clusters computer accounts.

- Service principal names are used by Kerberos authentication and machine accounts and a New SPN name pair is created each time a new SmartConnect Zone Alias is created.
- Mapping of SmartConnect Zones between source and target cluster.

This is required so that Eyeglass can create SmartConnect Zone names and aliases on your DR cluster automatically in the event of a DR failover.

- Update PowerScale sudoer file for Eyeglass Service Account.

This is required when the Eyeglass Service Account is being used to execute CLI commands that require root privileges (see the following section for detail).

Update PowerScale sudoer file for Eyeglass Service Account

Eyeglass Access Zone Failover requires some CLI commands that must run with root level access. Many customers also run the cluster in STIG or compliance mode for Smartlock WORM features. Root user account is not allowed to login and run commands. The **“SPN machine account maintenance before and after cluster failover”** command requires elevated permission to allow this user permissions across the cluster nodes. See “[Eyeglass Service account guide for minimum permissions](#)” for details on how to add sudo privileges to the Eyeglass cluster service account.

Active Directory Machine Account Service Principal Name (SPN) Delegation

What is an SPN?

It's used in Kerberos authentication from clients to network services for file serving. It's formed from SmartConnect Zones and has two forms: the NTLM netbios name and Kerberos name URL format.

Example:

When a client connects to \\data.example.com\sharename, the SPN for this authentication request to active directory uses the SPN name to authenticate. Kerberos is the default SPN request for

authentication and it uses the URL based request to the domain.

HOST\data

HOST\data.example.com

What's the risk if I don't fix SPN's?

Without SPN values set on the cluster machine accounts, Kerberos authentication will fail, but many Windows clients will fall back on NTLM authentication automatically (NTLM fallback can be disabled in the domain for higher level of security). NTLM is a legacy authentication protocol and considered less secure than Kerberos.

The Eyeglass solution aims at removing manual steps wherever possible. All prerequisites ensure manual steps are not required during a DR event. The Eyeglass solution also has a goal to remove dependencies between groups within IT, to reduce the potential for communication issues impacting DR failover.

SPN Delegation is a one time setup setup, that achieves simplified DR automation, and is required to use Eyeglass Access Zone failover feature. SPNs related to Source Cluster Zone and SmartConnect Zone AD providers will be deleted to avoid SPN collision in AD. During normal operating conditions, Eyeglass will

audit SPN's on source and destination clusters to insure they are correct and will remediate prior to any failover.

Delegate SPN

The steps outlined in "[How to - Delegation of Cluster Machine Accounts with Active Directory](#)" are required for each cluster machine account, for each AD provider that is added to each cluster. Example four different AD providers for different domains will require four delegations to be created, as per below, to each machine account name. Typically the cluster name is used when the cluster joins an active directory.

This one time setup avoids the requirement for failover operations to require ADS administrative permissions to successfully failover, and have SPN source and destination cluster values managed by Eyeglass. This reduce the risk of SPN authentication failures and ensures proper cluster self management of SPN fails required for proper Kerberos authentication for SmartConnect zones and aliases.

Note: Superna Eyeglass only manages SPN related to HOST. SPNs related to HDFS or NFS are not updated and will need to be manually repaired post failover.

Configure Eyeglass Subnet IP Pool Mapping Hints

This section covers why they are configured and how to configure mappings between IP pools for failover. IP pools serve data from SmartConnect names, the IP pool used to serve the name on failover is predetermined using ip pool mapping hints.

Zone Aliases for Failover Overview

Access Zone failover depends on dual DNS delegation to ensure no steps are required in DNS during a failover. The target cluster Subnet IP Pools require a SmartConnect Zone name set in the Onefs UI (must be in the UI and not an alias).

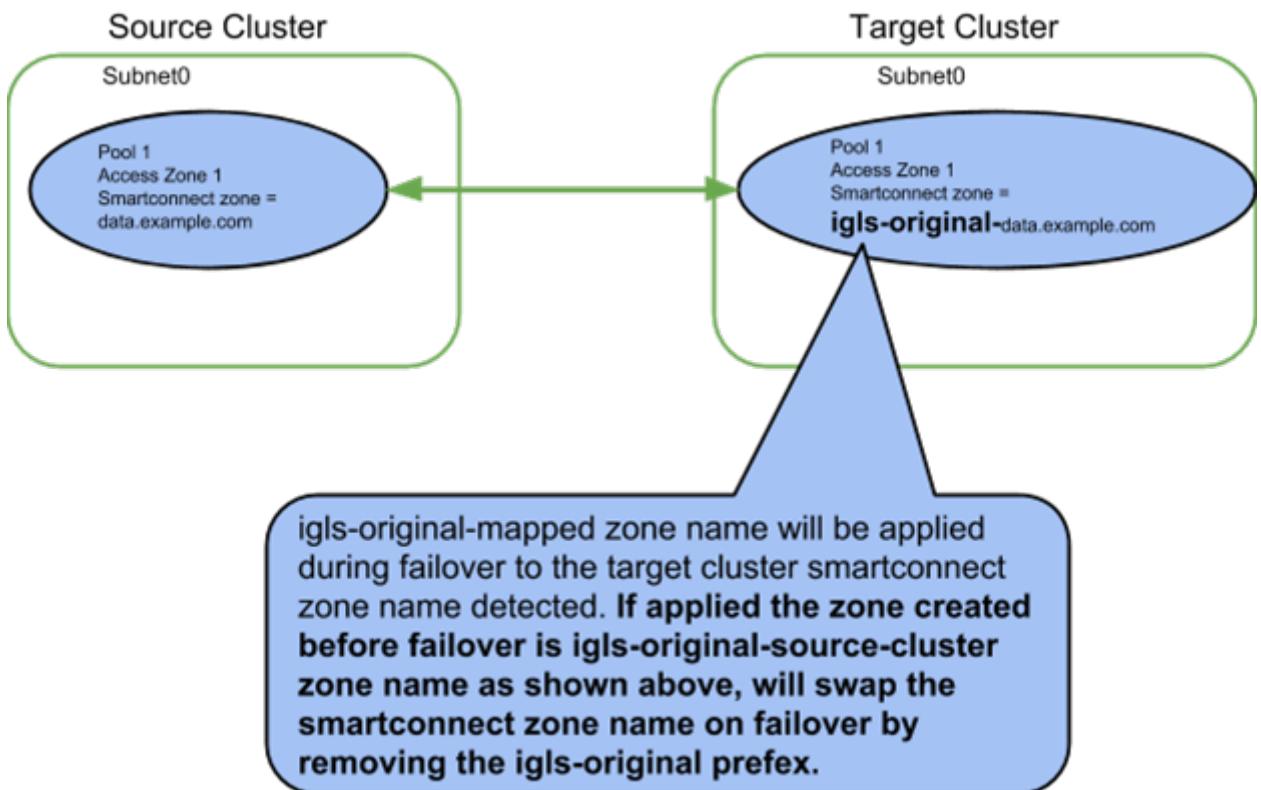
This SmartConnect name is not mounted or accessed on the DR cluster but it is required to configure the 2nd name server record in DNS to be setup.

The diagram below explains how we recommend SmartConnect Zone names to be entered to simplify the failover of a SmartConnect Zone name.

We recommend entering the source IP pool SmartConnect Zone name entered into the OneFS UI on the target mapped IP pool by applying a prefix of “**igls-original-<source cluster SmartConnect zone name>**

NOTE: The target ip pool MUST have a SmartConnect value.
The recommended value is as shown above or dual delegation responses will not function as expected. Blank or no SmartConnect name is NOT supported (no validation in DR Dashboard for checking if target cluster SmartConnect name is set correctly).

This simplifies failover visually in onefs UI and will **rename** on failover without an extra alias being created for failover purposes. It swaps the name from one side to the other during failover.



When and how to NOT failover an IP pool SmartConnect name using hint:
igls-ignore

The hint applied to an IP pool tells Eyeglass to not process this SmartConnect name and Aliases found on this IP pool.

Syntax of Igls-ignore:

This mapping hint can also be made unique using igls-ignore-xxxx where -xxx can be unique value to self document the purpose of the ignore. Examples below.

- Igls-ignore-repl (documents ignore on SmartConnect ip pool used for SyncIQ)
- Igls-ignore-dfsprod (documents prod hint for DFS pool used for DFS clients)
- Igls-ignore-mgmtclst1 (documents management FQDN pool used to manage the cluster)

When to apply ignore hints:

- For SyncIQ IP pools that are used for target host.
- For DFS IP pools so that no DNS updates are done for DFS target folders, also avoids SPN updates needed.
- For IP used for cluster management, and when Clusters are added to Eyeglass with this FQDN, it's required to apply an ignore hint so that Eyeglass will not lose access to the cluster during failover or fallback. This is also validated in the Zone Readiness screen and checked to make sure the FQDN used

for cluster add has an igls-ignore hint applied. This is a blocking condition for failover.

name	status
Zone Readiness Statuses	INFO
OneFS SyncIQ Readiness	INFO
Eyeglass Configuration Replication Readiness	OK
SPN Readiness	OK
SmartConnect/IP Pool Settings and Mappings Readin...	OK
Target Cluster Reachability	OK
Date-Time Validation	OK
FQDN Alias Validation	OK
172.31.1.104	OK
Zone Hot-Hot Validation	OK

Additional Status Information

If cluster was added to eyeglass with FQDN SmartConnect name for management, this SmartConnect zone must have an igls-ignore hint applied to avoid a failover impacting eyeglass access. Error means no igls-hint was found on the IP pool for smartconnect zone used for cluster management. ok means igls-ignore hint was found.

How to create Mapping Hints for IP pools between source and target clusters following best practise naming convention

Subnet:Pool Failover mapping between the Source and Target Cluster is done to ensure that data is accessed from the correct SmartConnect Zone IP and node pool on the Target cluster.

Mapping of IP address Pools should be completed after installation and will be audited by Eyeglass as part of the Failover Readiness validation. This is done using a SmartConnect Zone alias.

The hint alias on the IP Pool is of the form ***igls-xxx***, where “**xxx**” can be any string. *We recommend numbers to keep it simple.*

*Example; **igls-01-pool-name-prod** is self explanatory name and the DR mapping would be **igls-01-pool-name-dr**.*

NOTE: We also recommend this syntax to avoid SPN collision.

Eyeglass does not inject these hints but an admin could run ISI commands and inject them, no harm if they are present in AD computer account. For example “igls-01-prod” and “igls-01-DR” are matching hints since only “igls-xx” needs to match. Therefore the syntax form of “igls-xx-some-unique-string” with the trailing string “some-unique-string” allow them to be made unique and still match.

Eyeglass requires that the user decide which network pools are partnered during failover. Create the identical alias hint on the source network pools and their target network pools.

- A hint is a pre-fixed zone alias that instructs Eyeglass which source cluster network pool should failover to a specific target SmartConnect subnet pool. Eyeglass will detect when hints are missing and raises an alarm to correct it.
- An ignore hint is used to identify the SmartConnect Zone(s) used for SyncIQ replication. It is PowerScale best practice to have a dedicated SmartConnect Zone for this purpose, and

avoid using this zone name to mount data with clients.

During failover there is no need to failover the SmartConnect Zone used for SyncIQ. Eyeglass needs to know which zone name should be ignored during failover and readiness job assessment of Access Zones and SmartConnect Zones.

To add the mapping alias to the PowerScale cluster, ssh to the cluster and login as root, execute the following command: **(note subnet and pool names are case sensitive)**

- get list of pools “isi network list pools -v”
- isi network modify pool --name=<subnet:poolname> --add-zone-alias=<hint>

In the example below, we will execute the command on the source and target cluster to map pools to each other:

OneFS 7 Example igls hints for user data

- **Prod Cluster** isi network modify pool --name=subnet0:exampleProd --add-zone-aliases igls-01-prod
- **DR Cluster** isi network modify pool --name=subnet0:exampleDR --add-zone-aliases igls-01-dr

OneFS 8 Example igls hints for user data

Data Pool mapping example - Prod

Prod cluster data access IGLS hint applied to an Access Zone IP pool example. **NOTE: hint used to match is igls-marketing-marketingprod where marketingprod is used to identify the cluster the hint is applied. The marketingprod is not used to match the pools.**

The screenshot shows the 'Edit Pool Details' configuration page. It includes sections for SmartConnect Advanced settings (Client Connection Balancing Policy, IP Failover Policy, Rebalance Policy, Allocation Method), Advanced Settings (SmartConnect Auto Unsuspend Delay, SmartConnect DNS TTL options like 'No caching (recommended)' or 'Time to live'), and a SmartConnect Zone Aliases section. The 'igls-01-marketingprod' alias is highlighted with a red oval. At the bottom are 'Cancel' and 'Save Changes' buttons.

Data Pool mapping example - DR

DR cluster data access IGLS hint applied to an Access Zone IP pool example. **NOTE: hint used to match is igls-marketing-marketingdr where marketingdr is used to identify the cluster the hint is applied. The marketingdr is not used to match the pools**

Edit Pool Details

* = Required field

subnet0

SmartConnect Advanced

Client Connection Balancing Policy
Round-robin

IP Failover Policy
Round-robin

Rebalance Policy
Automatic

Allocation Method
Static

Advanced Settings

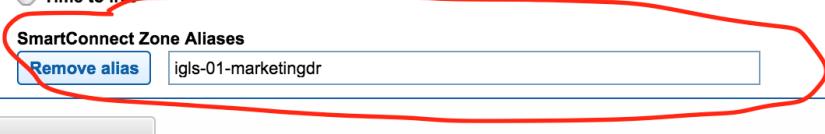
SmartConnect Auto Unsuspend Delay
0 Seconds

SmartConnect DNS TTL
 No caching (recommended)
 Time to live

SmartConnect Zone Aliases

Remove alias igls-01-marketingdr

Cancel **Save Changes**



IGLS examples for ignore option on IP Pools

In the example below, we will execute the command on the source and target cluster to pools that will be ignored for failover and are dedicated to SyncIQ replication or for DFS dedicated IP pools in the Access Zone :

OneFS 7 Example igls hints for ignore SyncIQ pool

- **Prod Cluster** isi network modify pool --name=subnet0:s iqProd --add-zone-aliases igls-ignore
- **DR Cluster** isi network modify pool --name=subnet0:s iqDR --add-zone-aliases igls-ignore

Example: Mapping with hints configured correctly will display the mapping

Network Mapping for Cluster2-7201 > prod-cluster-8 Zone: System		
Cluster2-7201	prod-cluster-8	
subnet0:pool0 Smart Connect Zone Name: dr.ad1.test Smart Connect Aliases: - igls-pool-dr Subnet: subnet0 SSIP: 172.31.1.201		subnet0:pool0 Smart Connect Zone Name: prod.ad1.test Smart Connect Aliases: - igls-pool-prod Subnet: subnet0 SSIP: 172.31.1.200
subnet0:SIQ-DR Smart Connect Zone Name: SIQ-DR.ad1.test Smart Connect Aliases: - igls-ignore-dr-siq Subnet: subnet0 SSIP: 172.31.1.201	 IGNORE	Pool will not be failed over because of igls-ignore alias.

OneFS 8 Example igls hints for ignore SyncIQ pool

SyncIQ Ignore hint replication Pool.

This Pool is used by SyncIQ for replication (restrict source and or target host for replication. **The igls-ignore- is used to ignore the pool the prod8 makes the hint unique and identifies the cluster the hint is applied using the cluster name prod.**

View Pool Details

* = Required field

SmartConnect Basic

Zone Name
repl.ad3.test

SmartConnect Service Subnet
subnet0

SmartConnect Advanced

Client Connection Balancing Policy
Round-robin

IP Failover Policy
Round-robin

Rebalance Policy
Automatic

Allocation Method
Static

Advanced Settings

SmartConnect Auto Unsuspend Delay
No delay

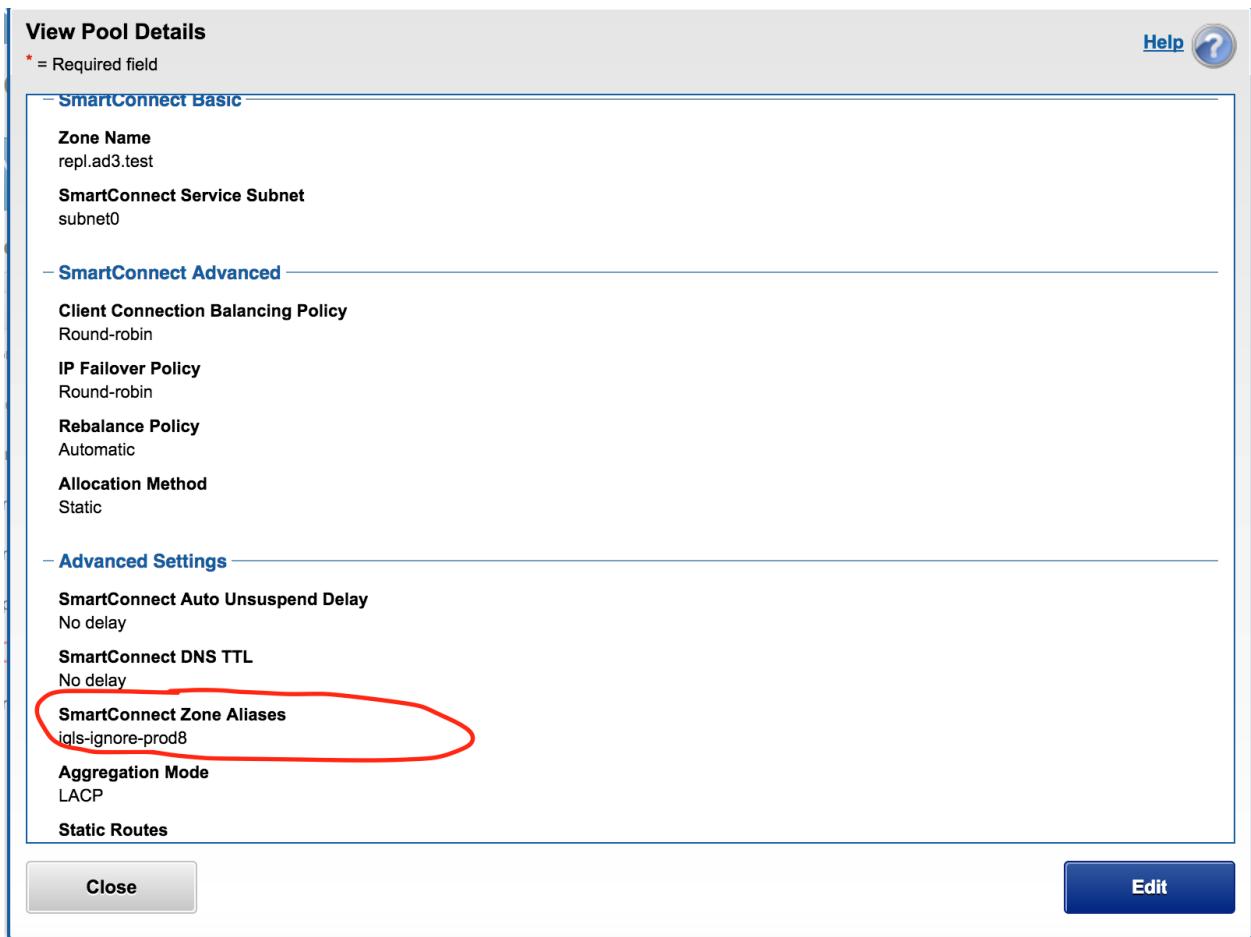
SmartConnect DNS TTL
No delay

SmartConnect Zone Aliases
igls-ignore-prod8

Aggregation Mode
LACP

Static Routes

Close **Edit**



DFS Ignore example - Prod

This is another ignore hint used for a **prod cluster** pool that protects DFS mounted data in the Access Zone. This IP pool and its SmartConnect names should not failover and uses an ignore hint igls-ignore-dfs01 where igls-ignore- is used to match and dfs01 is to make the hint unique.

Edit Pool Details

* = Required field

SmartConnect Advanced

Client Connection Balancing Policy
Round-robin

IP Failover Policy
Round-robin

Rebalance Policy
Automatic

Allocation Method
Static

Advanced Settings

SmartConnect Auto Unsuspend Delay
0 Seconds

SmartConnect DNS TTL
 No caching (recommended)
 Time to live

SmartConnect Zone Aliases

Remove alias igls-ignore-dfs01

+ Add a zone alias

Cancel **Save Changes**



DFS Ignore example - DR

This is another ignore hint used for a **DR cluster** pool that protects DFS mounted data in the Access Zone. This IP pool and its SmartConnect names should not failover and uses an ignore hint igls-ignore-dfs02 where igls-ignore- is used to match and dfs02 is to make the hint unique.

Edit Pool Details

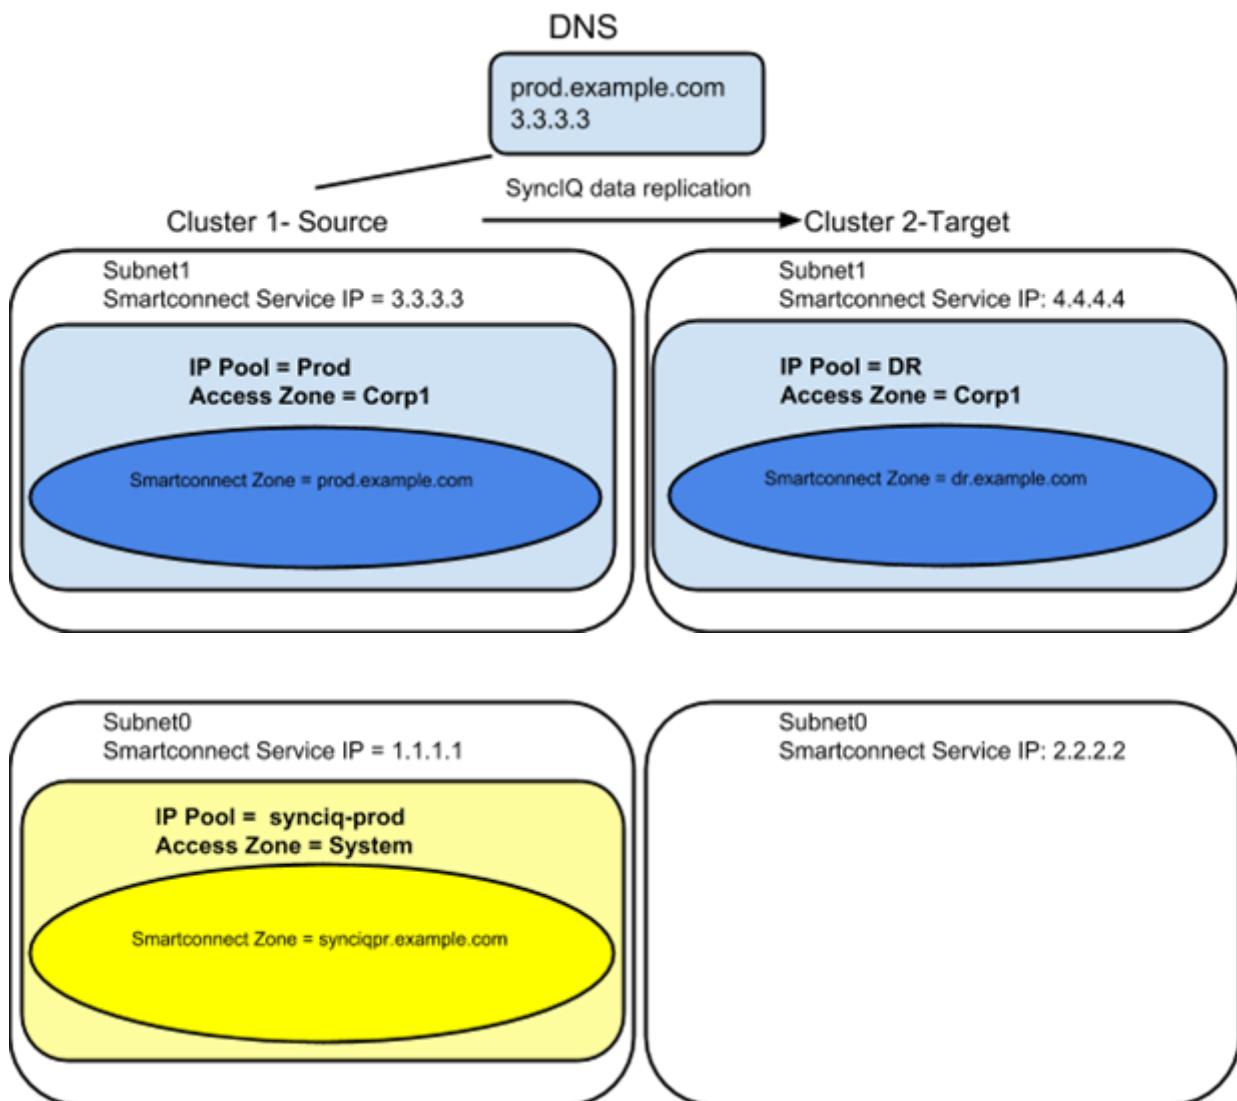
* = Required field

Automatic	<input type="button" value="▼"/>		
Allocation Method	<input type="button" value="Static"/>		
Advanced Settings			
SmartConnect Auto Unsuspend Delay			
0	Seconds		
SmartConnect DNS TTL			
<input checked="" type="radio"/> No caching (recommended) <input type="radio"/> Time to live			
SmartConnect Zone Aliases			
<input type="button" value="Remove alias"/>	igls-ignore-dfs\p2		
<input type="button" value="Add a zone alias"/>			
Aggregation Mode			
LACP	<input type="button" value="▼"/>		
Static Routes			
<input type="button" value="Add static route"/>			
Subnet	Netmask	Gateway	Actions
There are no static routes			

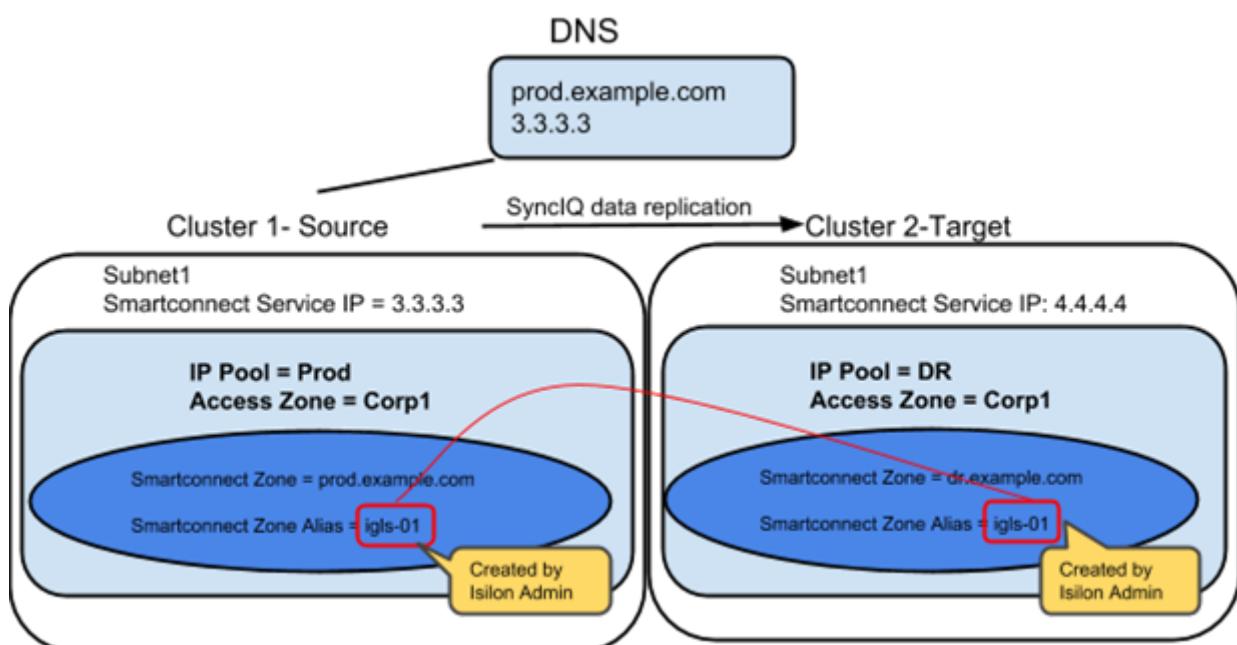
Hot-Cold Replication Topology Mapping Examples

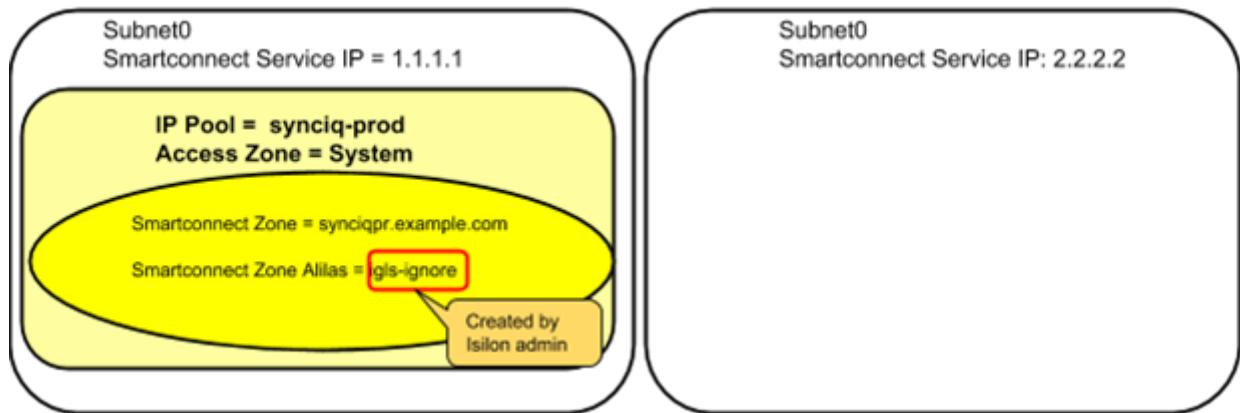
Example: Single Access Zone, Single IP Pool for Access, Single IP Pool for SyncIQ

Before Mapping:



After Mapping:

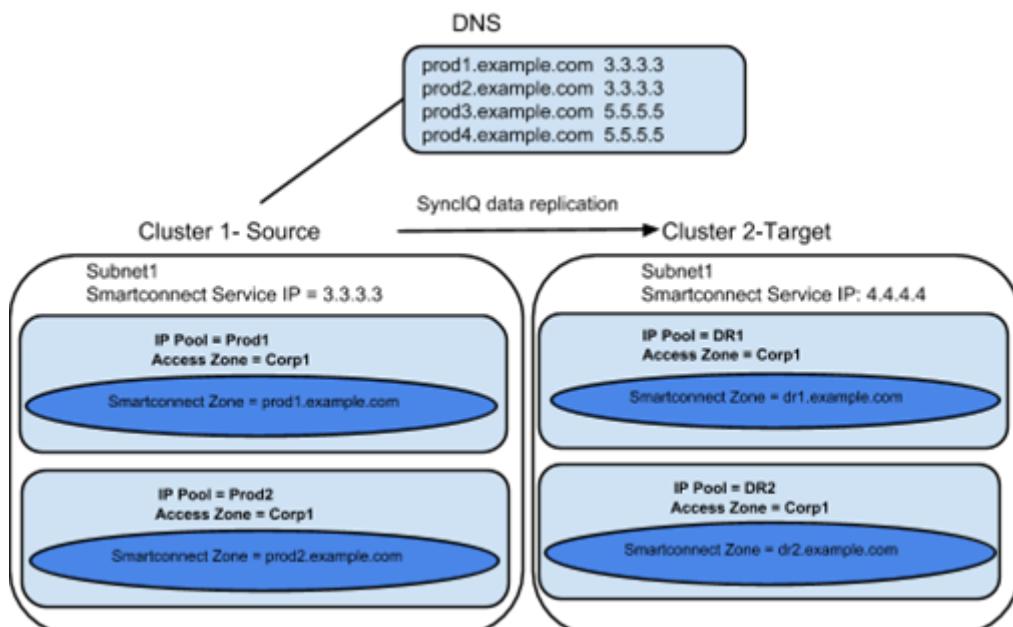


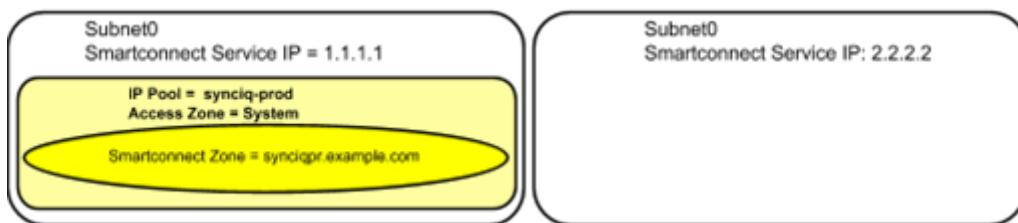
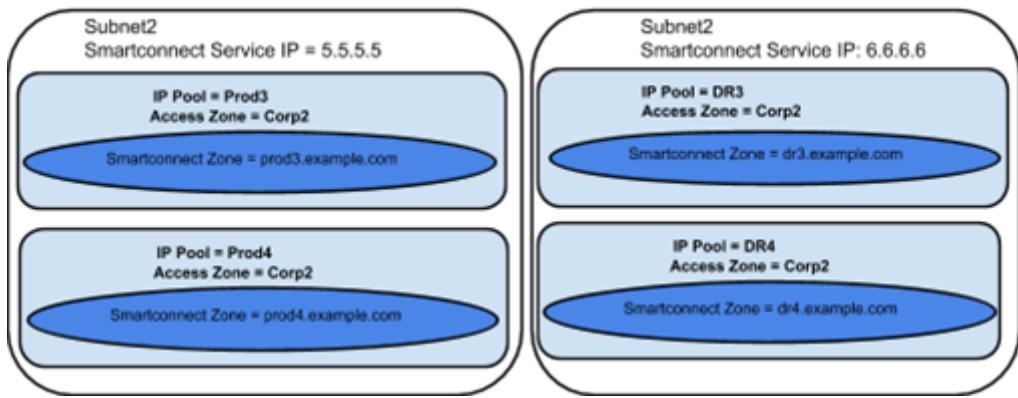


- Smartconnect Zone for SyncIQ aliased to ignore it during failover
igls-ignore
- Smartconnect Zone for User Access aliased for Eyeglass Failover Automation
igls-<unique string>
- Smartconnect Zone for DFS User Access aliased for Eyeglass Failover Automation
igls-<unique string>
- alias Smartconnect Zone alias pre-provisioned for Eyeglass Failover Automation

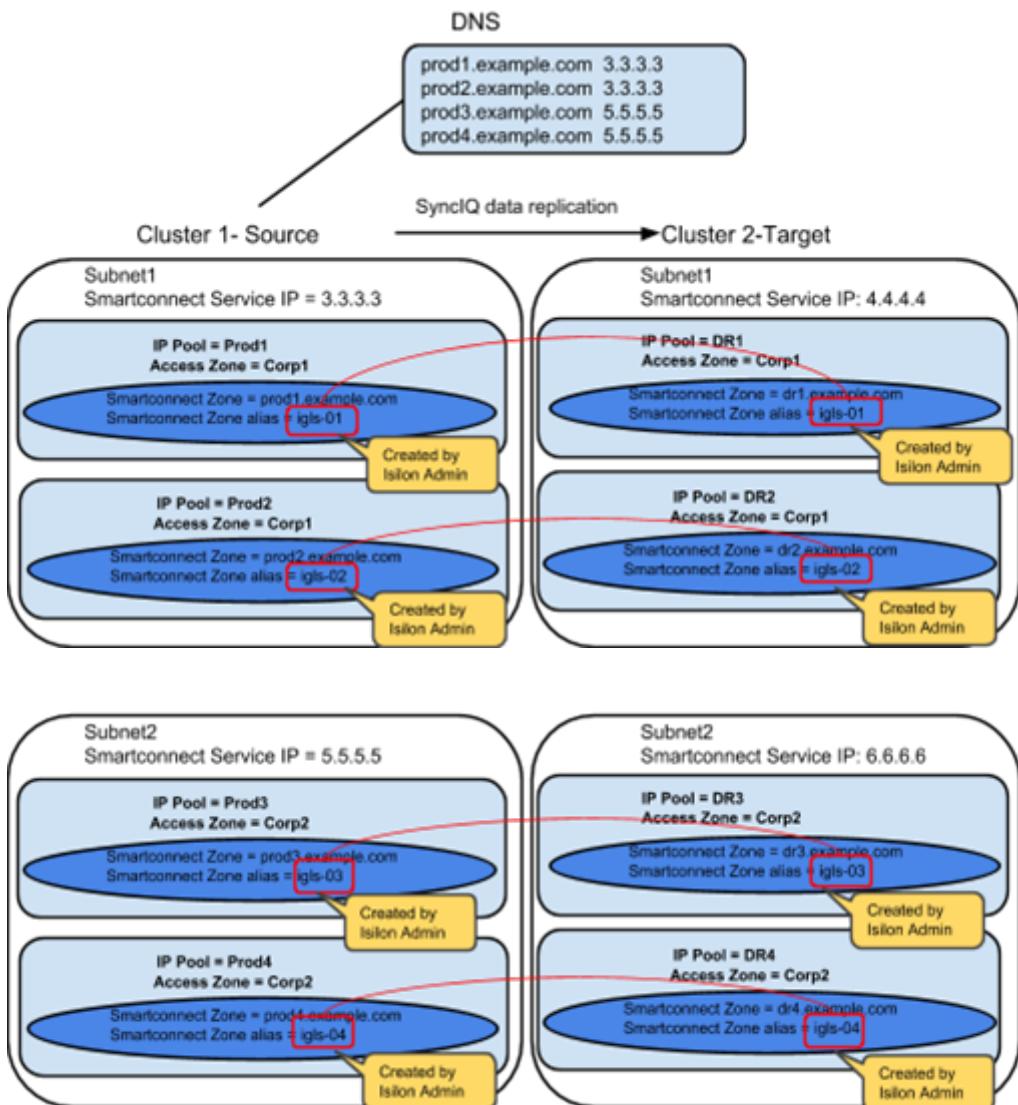
Example: Multiple Access Zone, Multiple IP Pool for Access, Single IP Pool for SyncIQ

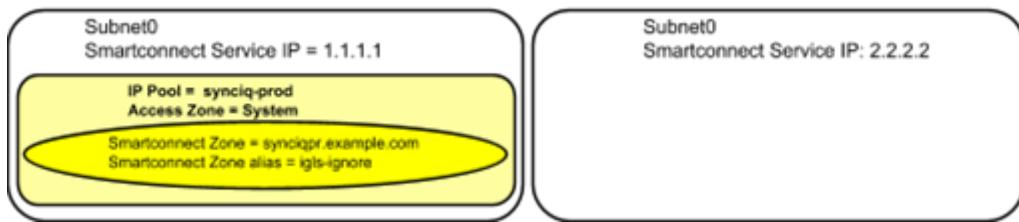
Before Mapping:





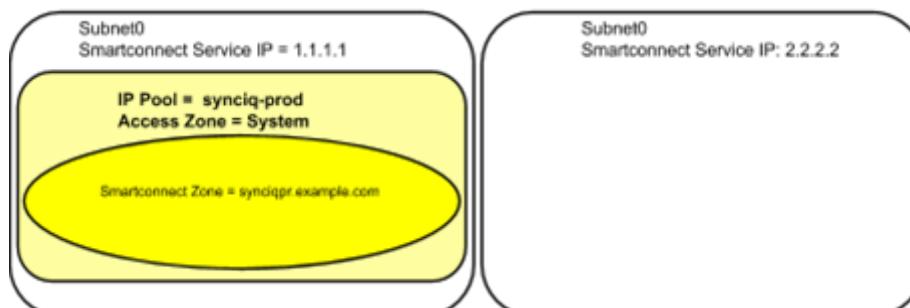
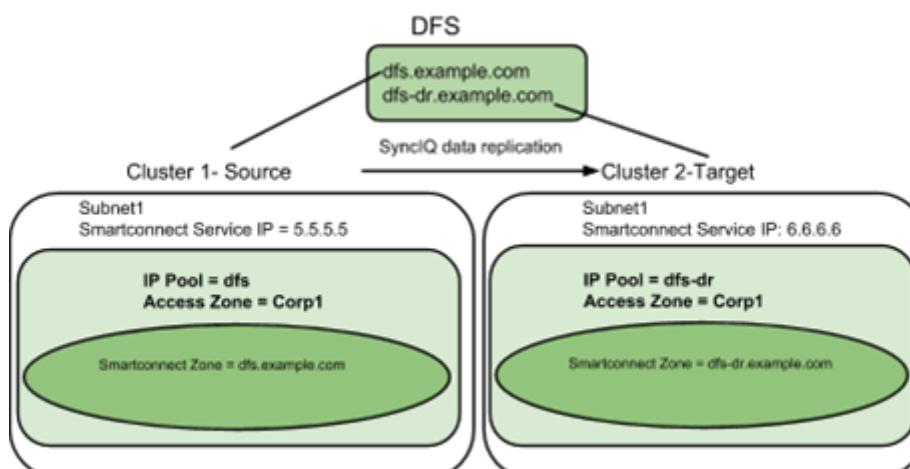
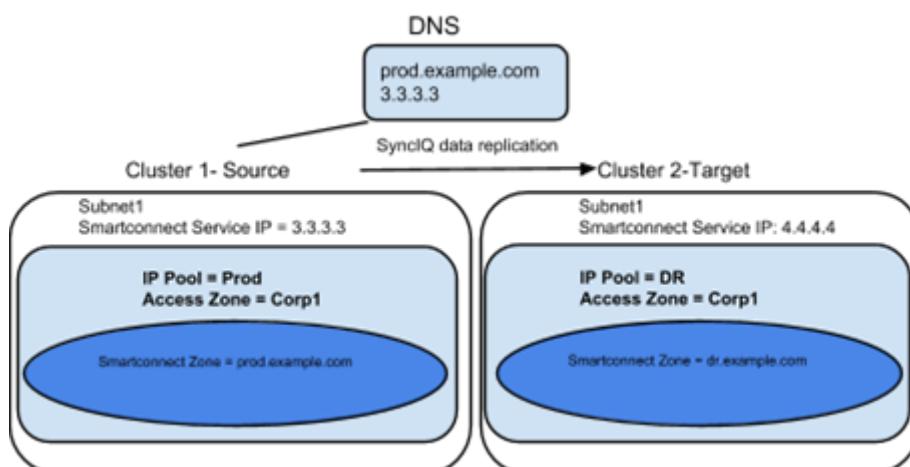
After Mapping:



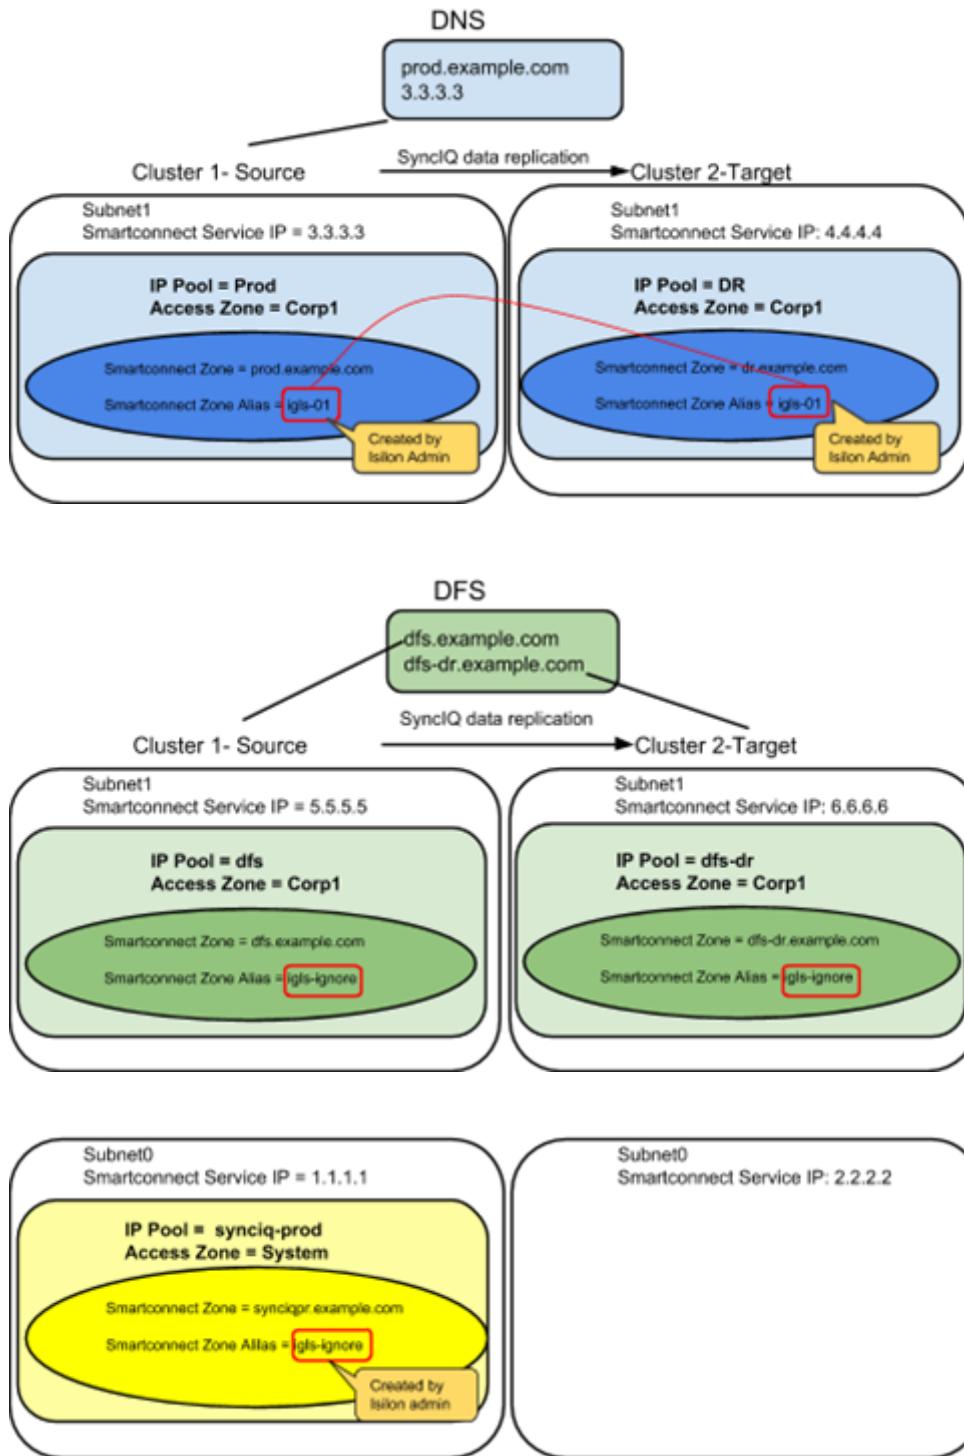


Example: Single Access Zone containing DFS, Single IP Pool for Access, Single IP Pool for SyncIQ

Before Mapping



After Mapping:



2.9. PowerScale Administration for Clusters Configured for Eyeglass Assisted Access Zone Failover

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

PowerScale Administration for Clusters Configured for
Eyeglass Assisted Access Zone Failover

SPN Management

With SPN Delegation configured as required in preparation for Eyeglass Assisted Access Zone Failover, Eyeglass will create SPNs related to SmartConnect Zones and Alias detected. **No manual SPN management is required.**

IMPORTANT:

Eyeglass does not create SPNs for any SmartConnect Zone or SmartConnect Zone Alias that are prefixed with igls. SPN check from Clusters configured with Eyeglass mapping hints will indicate that there are missing SPN's for these SmartConnect Zones and Aliases . This is expected as these SmartConnect Zones and Aliases are not used for Cluster access. **DO NOT EXECUTE SPN REPAiR** as it will fail if executed on both clusters because of conflict created by having identical mapping hints on clusters.

IMPORTANT:

Eyeglass does not remove “extra” SPN’s that do not correspond to detected SmartConnect Zones and Aliases. This must be done manually if required for these SPN to be removed.

Post Failover Automation

Many failover scenarios depend on extra steps performed on devices, software, and infrastructure external to the NAS cluster.

Using the Eyeglass script engine, these tasks can now be automated with output captured and appended to Eyeglass failover logs. For example:

- DNS updates post failover for SmartConnect zone CNAME editing.
- NFS host mount and remount automation.
- DNS server cache flushing.
- Application bring up and down logic to start applications post failover.
- Send alerts or emails.
- Run API commands on 3rd party equipment example load balancer, switch, router or firewall.

Please refer to the Eyeglass Admin Guide [Script Engine Overview](#) for more details.

2.10. Failover Planning and Checklist

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

Failover Planning and Checklist

Failover planning includes extended preparation beyond storage layer failover steps. A full Failover Plan will also take into account clients, application owners and any dependent systems such as DNS and Active Directory. The following link is a Failover Planning Checklist to help you develop your own Failover plan ([Failover Planning and Checklist](#)).

© Superna Inc

2.11. Monitoring DR Readiness for Eyeglass Assisted Failover

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

Monitoring DR Readiness for Eyeglass Assisted Failover

In addition to the Assisted Failover functionality, Eyeglass also provides the following features to monitor your Access Zone DR Readiness:

- Access Zone DR Readiness Validation.
- Runbook Robot.

Access Zone DR Readiness Validation

The DR Dashboard Zone Readiness tab provides a per Access Zone summary of all the key networking, Kerberos SPN, SmartConnect connect subnet\pool information along with SyncIQ status and Configuration replication validations performed to assess readiness for failover by Access Zone. The status for each are combined to provide an overall DR Status. The Zone Readiness is updated every 15 minutes by default ([See "igls cli commands" in the Eyeglass PowerScale Edition Administrative Guide to change this schedule](#)).

This information provides the best indicator of DR readiness for failover and allows administrators to check status on each

component of failover, identify status, errors and correct them, in order to get each Access Zone configured and ready for failover.

By default the Failover Readiness job which populates this information is disabled. Instructions to enable this Job can be found [here](#). Under Managing Eyeglass Jobs

If all of the Access Zone Requirements and Recommendations pass validation, the DR Dashboard status for the Access Zone is **green** indicating that the Access Zone is safe to failover.

If any of the Access Zone Requirements do **NOT** pass validation, the DR Dashboard status for the Access Zone is **red** indicating that the Access Zone is **NOT** ready to failover. In this state the DR Assistant will block you from starting the failover. Eyeglass will also issue a System Alarm for any of these conditions.

If any of the Access Zone Recommendations do **NOT** pass validation, the DR Dashboard status for the Access Zone is **orange (Warning)** indicating that the Access Zone can be failed over but there may be some additional manual steps required to complete the failover. In this state the DR Assistant will allow you to start the failover. Eyeglass will also issue a System Alarm for any of these conditions.

Additional information for Zone Readiness can be found in the Eyeglass Admin guide [here](#).

IMPORTANT:

If you make a change to your environment, the following Eyeglass tasks must run before the Zone Readiness will be updated:

- Configuration Replication.
- Failover Readiness.

IMPORTANT:

Readiness is **NOT** assessed for the Access Zone in the Failed Over state. This means the DR Dashboard Readiness provides a status, or Readiness, from the current active cluster to the DR target cluster **ONLY**. The reverse direction “Fail back” status is not assessed until failover to the target cluster.

DR Dashboard						
Zone Readiness	Source Cluster	Target Cluster	Zone Name	Last Successful Readiness Check	Network Mapping	DR Failover Status
Pool Readiness	ProdBx	DfBx	System	9/20/2017, 8:30:13 AM	View Map	ERROR
DFS Readiness	DRBx	ProdBx	System	9/20/2017, 8:30:13 AM	View Map	ERROR
Policy Readiness	prod-B	disasterB	EyeglassRunbookRobot	9/20/2017, 8:30:23 AM	View Map	FAILED OVER
DR Testing	disasterB	prod-B	EyeglassRunbookRobot	9/20/2017, 8:30:23 AM	View Map	WARNING
	prod-cluster-B	Cluster2-7201	System	9/20/2017, 8:30:29 AM	View Map	INFO
	Cluster2-7201	prod-cluster-B	System	9/20/2017, 8:30:28 AM	View Map	OK

Runbook Robot (Automate DR Testing on a schedule)

Overview

Many organizations schedule DR tests during maintenance windows and weekends, only to find out that the DR procedures did not work, or documentation needed to be updated. The Eyeglass Run Book Robot feature automates DR run book

procedures that would normally be scheduled in off peak hours, and avoids down time to validate DR procedures, providing Failover and Failback automation tests with reporting.

This level of automation provides a high level of confidence that your PowerScale storage is ready for failover with all of the key functions executed on a daily basis. In addition to automating failover and failback, Eyeglass operates as a cluster witness.

Eyeglass uses Access Zone mount paths to mount storage on both source and destination clusters the same way the cluster users and machines mount storage externally.

Run Book Robot Failover Coverage

The following validations are all performed on a daily basis, and the DR dashboard updated along with any failures sent as critical events. This is the best indicator that your cluster is ready for a failover.

- API access to both clusters is functioning - **Validated**.
- API access allows creation of export, share, quota - **Validated**.
- NFS mount of data external to the cluster functions - **Validated**.
- DNS resolution for SmartConnect is checked when Eyeglass configures itself to use SmartConnect service IP as its DNS resolver on the source, in order to verify SmartConnect zone functionality on mount of data requests - **Validated**.

- SyncIQ policy replication completes between source and destination cluster when data is written to the source - **Validated**.
- Configuration replication of test configuration from source to destination - **Validated**.
- SyncIQ failover to target cluster - **Validated**.
- Test data access on target cluster post failover - **Validated**.
- Verify data integrity of the test data on target cluster - **Validated**.
- Configuration Sync of quotas from source to target on failover - **Validated**.
- Delete Quotas on source cluster - **Validated**.
- SyncIQ Failback from target to source cluster - **Validated**.

Refer to the Eyeglass “[RunBookRobot Admin Guide](#)” for instructions on setting up and running the Runbook Robot.

© Superna Inc

2.12. Operational Steps for Eyeglass Assisted Access Zone Failover

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

Operational Steps for Eyeglass Assisted Access Zone Failover

Access Zone Workflow Steps Overview

For Ordered Steps for - Access Zone Failover consult [this table](#) in the Eyeglass Failover Design Guide.

Access Zone Execution Steps

For detailed steps on execution and monitoring consult the [Eyeglass Failover Design Guide](#).

© Superna Inc

2.13. Post Access Zone Failover Steps

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

Post Access Zone Failover Steps

- [Test Dual Delegation](#)
- [DNS SmartConnect name tests](#)
- [Check for SPN Errors](#)
- [Automated SMB Connection switch to target cluster after Failover](#)
- [Manual SMB connection switch to target cluster after Failover](#)
- [Refreshing NFS Mounts after Failover](#)

Test Dual Delegation

Update DNS is not required with dual delegation. If NSLOOKUP testing verifies the response is from target cluster, then no extra steps are required. If not check to insure DNS configuration is correctly using target cluster SSIP in the delegation records in DNS.

See the following reference for details on Dual Delegation:

[Geographic Highly Available Storage solution with Eyeglass Access Zone Failover and Dual Delegation](#)

DNS SmartConnect name tests

Verify that DNS Updates were completed correctly:

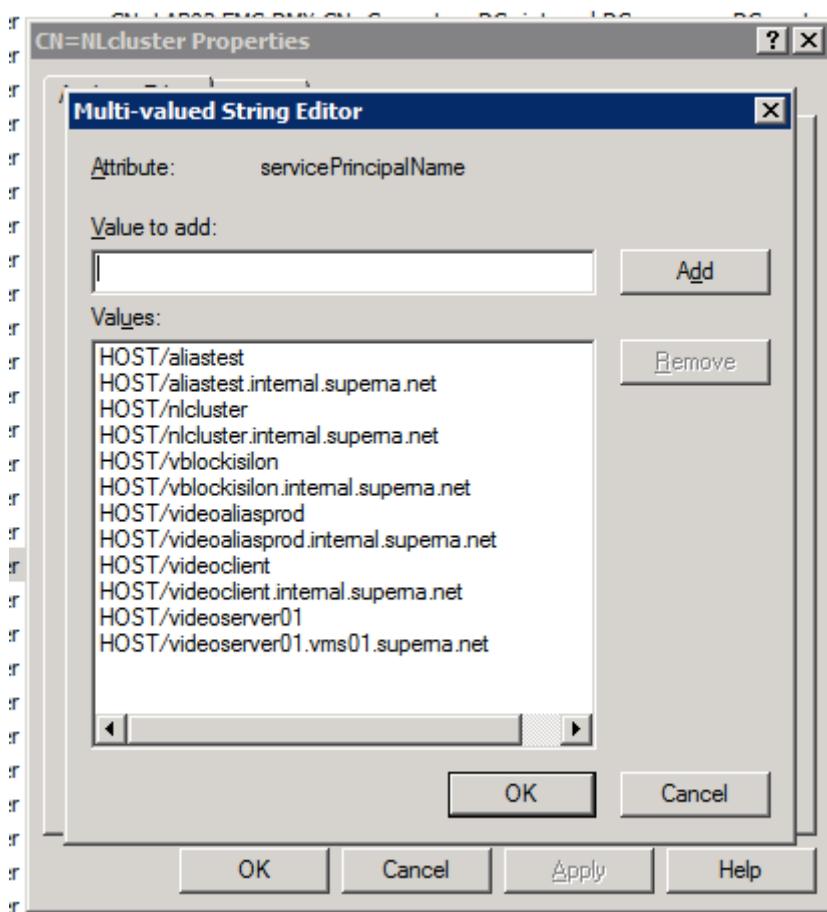
1. SSH admin@eyeglass ip address.
2. nslookup [enter].
3. server x.x.x.x [enter] (ip of subnet service ip on the TARGET cluster).
4. somesmartconnect.zone.name from the SOURCE cluster that was failed over [enter].
5. Expected response should return an IP address from the TARGET clusters ip pool that was mapped.
6. If expected output is from target cluster IP pool then failover of SmartConnect delegation to the TARGET cluster is correct.
7. Now repeat the server x.x.x.x command using production DNS server that has modified CNAME .
8. Now repeat above tests using production DNS client ip address example “server y.y.y.y” (where y is ip address of updated SOA primary DNS server where the delegation record was changed).
9. Verify the output returns ip address from TARGET cluster ip pool that was mapped.

Check for SPN Errors

Review the failover log and verify all SPN steps were successful. If any SPN steps show as failed manual recovery of the SPN will be

required using ADSIedit AD tool to perform delete and add of SmartConnect name or Alias names. The failover log contains all SPN smartconnect names that were included in the failover.

The Microsoft ADSIedit tool is the simplest method to make computer account SPN changes post failover. This tool requires Microsoft permissions to the computer account for the cluster being edited. Consult Microsoft documentation on ADSIedit usage.



Automated SMB Connection switch to target cluster after Failover

This procedure has been tested with specific OneFS versions and SMB protocol versions across different Windows OS's. Review the testing matrix of OS, Onefs and protocol combinations.

NOTE: You cannot use this procedure if you have disabled the SMB protocol on the cluster.

1. Complete the failover as normal with Eyeglass DR Assistant.
2. Complete the steps below for each Access Zone IP pool that was part of the failover AND has SMB connections that should switch to the target cluster.
 - a. View/edit the IP pool
 - b. Record the Interfaces that are members of the pool, this will be required to reconfigure the pool.

The screenshot shows the 'Pool Interface Members' configuration screen. At the top, there is a dropdown for 'Access Zone' set to 'data'. Below it is an 'IP Range' input field containing '172.31.1.112' and a 'Save Changes' button. The main area is divided into two sections: 'Available' and 'In Pool'. The 'Available' section contains a table with six rows, each with an LNN and an interface name: 1(ext-2), 1(ext-3), 1(ext-4), 1(ext-5), and 1(ext-6). The 'In Pool' section contains a table with one row: 1(ext-1). A red oval highlights the 'In Pool' section. Between the two sections are 'Add ➔' and '⬅ Remove' buttons. At the bottom are 'Cancel' and 'Save Changes' buttons.

- i. c. Select all interfaces and click Remove

- d. Save the pool with no interfaces as members.
- e. Once the above step is completed, any connected SMB clients will query DNS for the smartconnect name and will re-mount and re-authenticate to the target cluster.
- f. View/edit the same IP pool
- g. Re-add the interfaces that were recorded in the step above.
- h. Save the pool again.
- i. REPEAT the steps above for each IP pool included in the failover
- j. Test data access to the target cluster to verify SMB clients have switched and can write data. If any test access fails, use the [test data access debug guide](#).

Manual SMB connection switch to target cluster after Failover

This section describes steps to refresh an SMB connection post failover.

1. If the client was connected to the share during the failover.

2. Unmount the share (disconnect).
3. Remount the share (connect).
4. Test read/write against newly mounted shares.
5. If step 3 fails, the original connection information is likely cached on the client machine. The data in this case would continue to be available, however it would be Read-Only. Writes would fail. To remove the cached connection information reboot the pc or use net use command or Windows Explorer to unmount the share.
5. Test read/write against newly mounted share again.
6. If previous step fails, use the [Test Data Access procedures](#) to diagnose the issue.

Refreshing NFS Mounts after Failover

NFS mounts require an unmount and remount on the host.

1. To unmount an export with open files use umount -fl option (force and lazy flag)
2. Remount the export or if configured in fstab mount -a to remount any unmounted entries in the file.

3. **NOTE:** For automated NFS export remounts consider using the script engine feature to ssh to hosts and remount post failover. Guide is [here](#).

© Superna Inc

2.14. Post Access Zone Failover Checklist

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

Post Access Zone Failover Checklist

The following sections outline what can be checked post Access Zone failover to verify execution of all of the steps.

IMPORTANT:

If the failover was done with the **Controlled failover** option unchecked, then some steps on the failover SOURCE cluster will not have been executed.

SyncIQ Policy Updates

On the failover SOURCE cluster (the cluster you failed over FROM), for the SyncIQ Policies in the Access Zone that were failed over:

- SyncIQ Policies are Disabled in OneFS .
- Eyeglass configuration replication jobs related to these SyncIQ Policies are in Policy Disabled state.
- SyncIQ Policies in OneFS have their schedule set to manual.

On the failover TARGET cluster (the cluster you failed over TO), for the SyncIQ Policies in the Access Zone that were failed over:

- SyncIQ Policies are Enabled in OneFS.

- The corresponding Eyeglass Configuration Replication Jobs are also in Enabled state.
- SyncIQ Policies have same schedule that was originally set for the policy on the failover SOURCE cluster.

NOTE: If you have Eyeglass INITIALSTATE property set to disabled for AUTO jobs (check this using the command in the [Eyeglass Administration Guide](#)), the Eyeglass Configuration Replication job for the mirror SyncIQ Policy created during the first failover will be in User Disabled state. This job should be enabled following the instructions in the [Eyeglass Administration Guide](#).

Quota Updates

After the upgrade, there should be no quotas on the failover SOURCE cluster for the SyncIQ Policies in the Access Zone that were failed over. On the failover TARGET cluster you should find all quotas for the SyncIQ Policies in the Access Zone that were failed over.

SPN Updates

After the upgrade, there should be SPNs for all SmartConnect Zones and SmartConnect Zone Aliases related to the subnet pools associated with the Access Zone that was failed over.

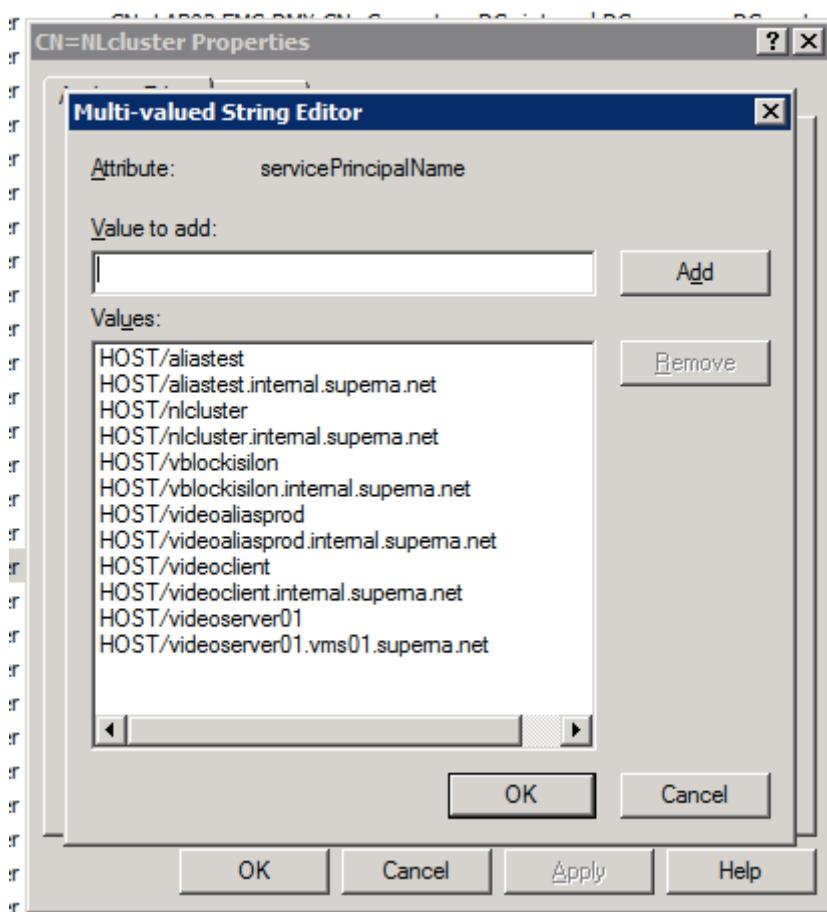
Note: SPNs are not created for SmartConnect Zones or SmartConnect Zone Aliases that are prefixed with “igls”.

Note: SPNs are not created for HDFS or NFS. These will need to be repaired manually.

IMPORTANT:

Due to an PowerScale issue, it may occur that executing the SPN repair step results in an error (both for execution by Eyeglass and from the PowerScale command line directly). In this case the SPNs will have to be repaired manually after which the SPN repair command will resume as expected.

Use the ADSIedit tool to verify the machine account SPN's.



SmartConnect Zone Updates

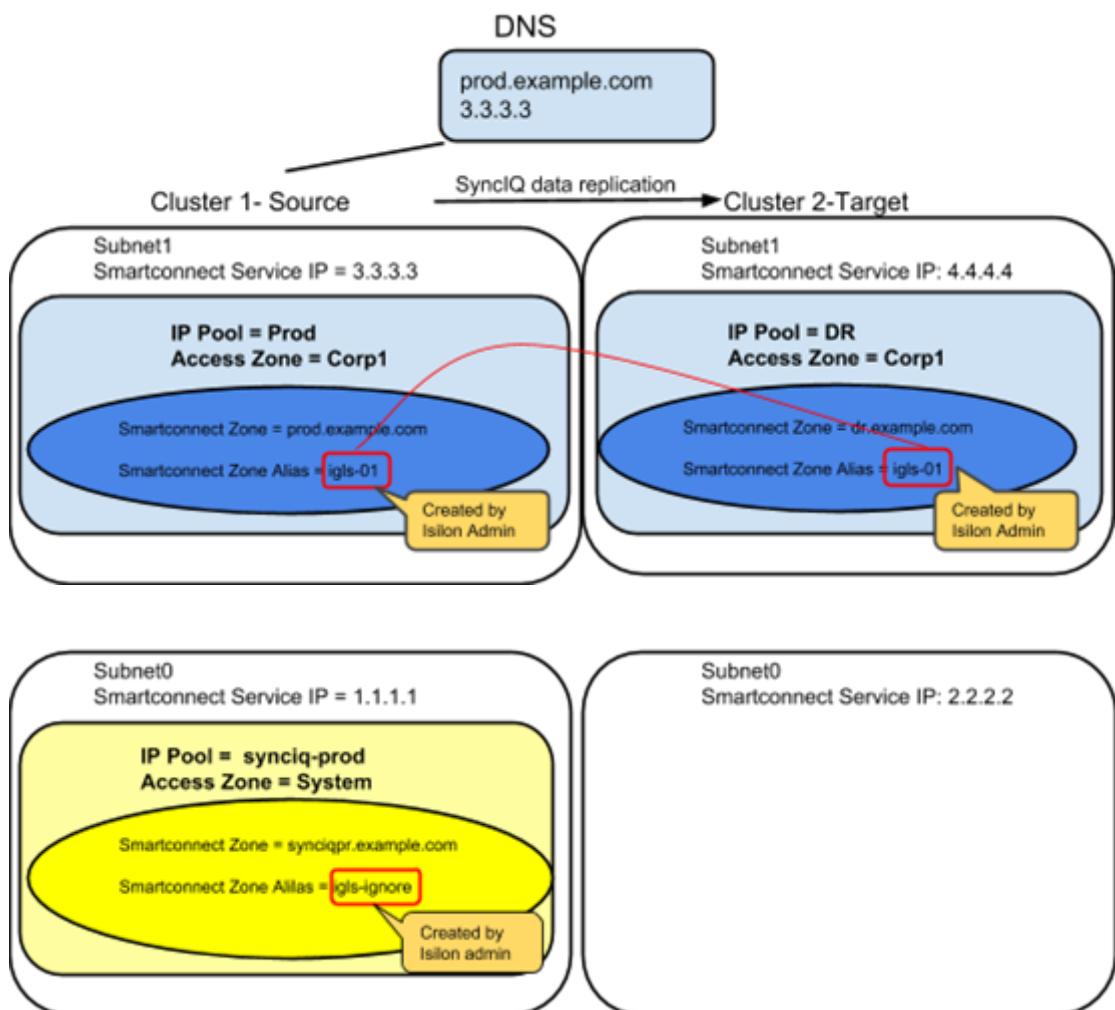
The following changes can be checked post failover for the SmartConnect Zones and aliases related to the subnet pools associated with the Access Zone that was failed over:

1. Eyeglass creates SmartConnect Zone alias on failover TARGET cluster with the same name as SmartConnect Zone on the failover SOURCE cluster partner IP Pool.
2. Eyeglass updates failover SOURCE cluster SmartConnect Zone name with the prefix “igls-original” .

3. Alias for failover TARGET SmartConnect Zone is removed from failover SOURCE cluster.
4. After Failover completed, DNS Admin or post failover scripting updates DNS entry for the SmartConnect Zone name to use the SmartConnect Service IP address from CLUSTER 2.

Example: SmartConnect Zone Update

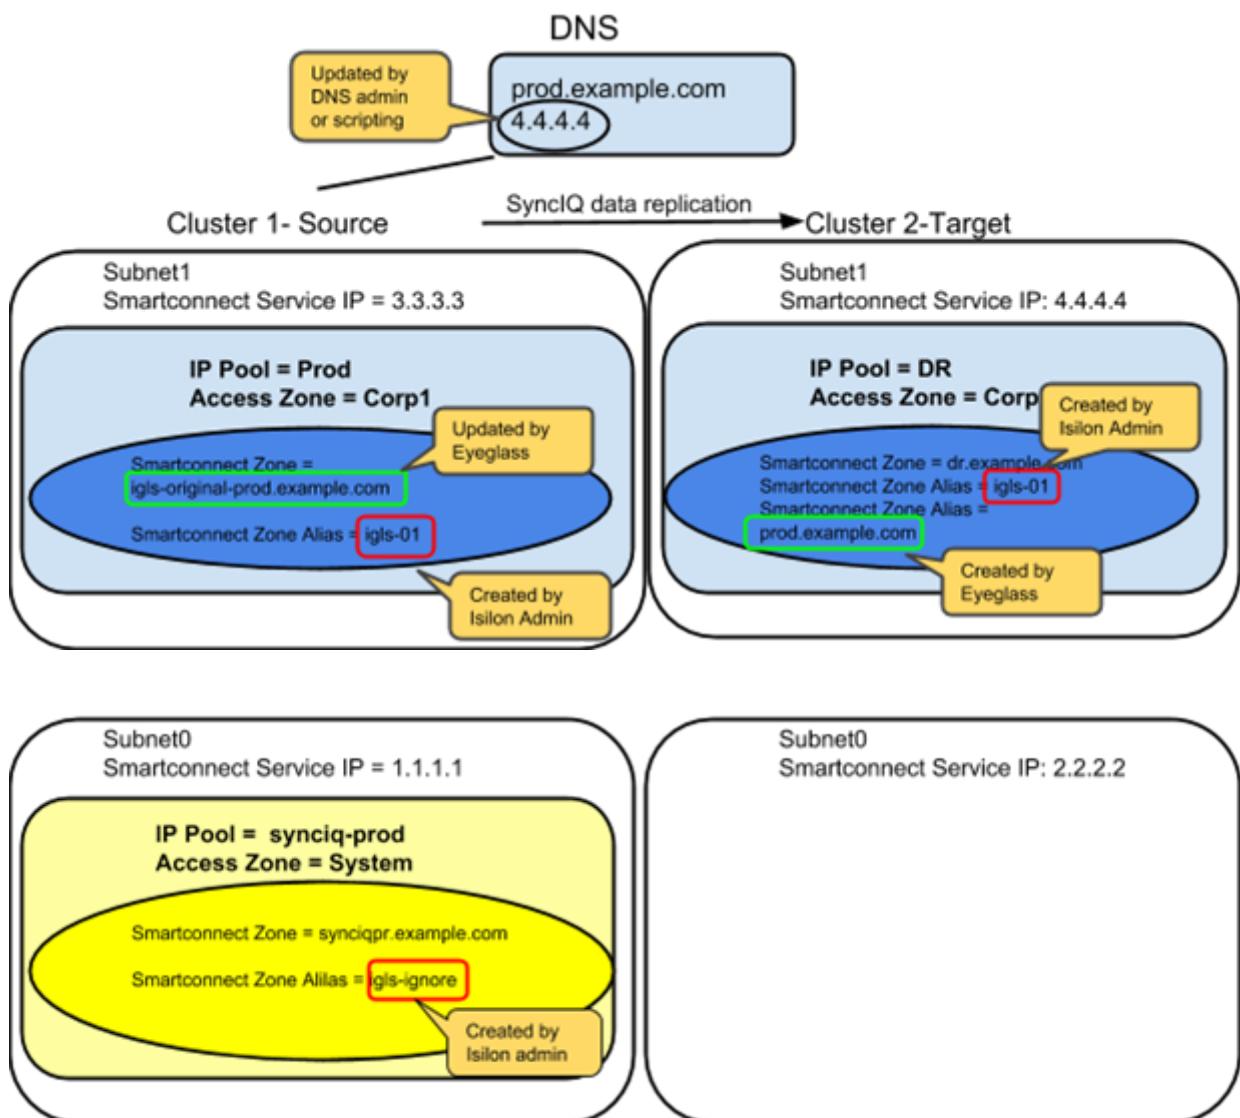
Initial Mapping Setup:



- Smartconnect Zone for SyncIQ aliased to ignore it during failover
igls-ignore
- Smartconnect Zone for User Access aliased for Eyeglass Failover Automation
igls-<unique string>
- alias Smartconnect Zone alias pre-provisioned for Eyeglass Failover Automation

FAILOVER from CLUSTER 1 to CLUSTER 2

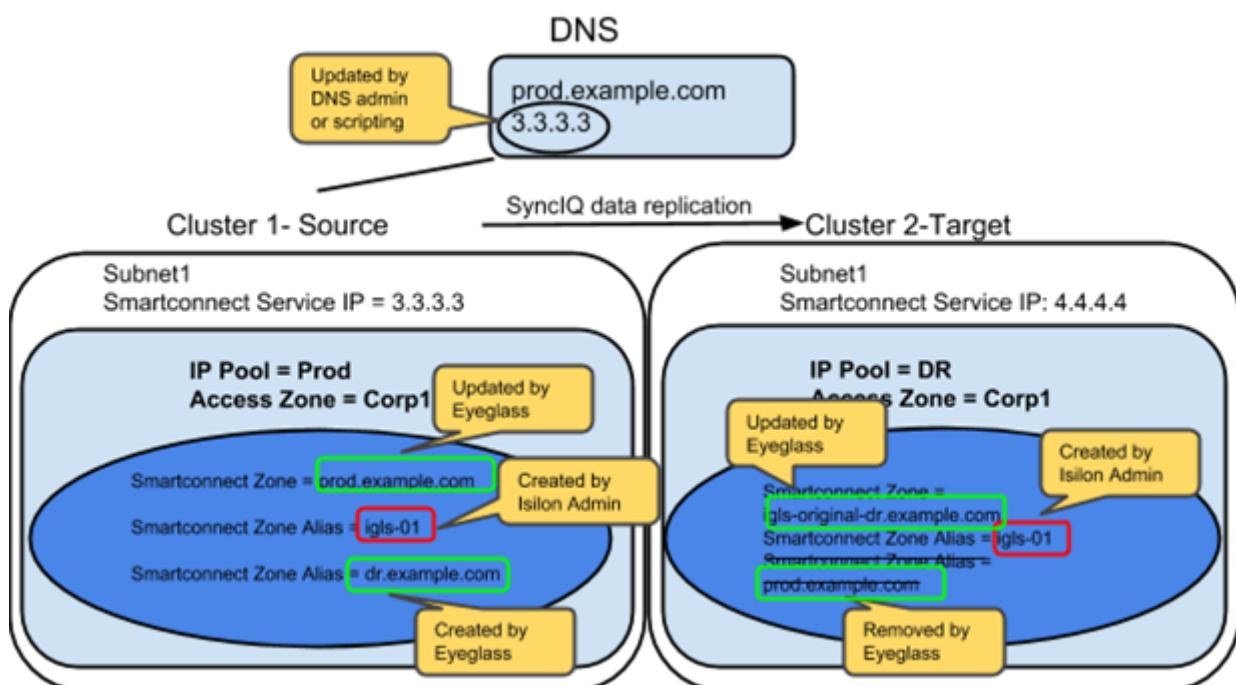
	Cluster 1	Cluster 2
subnet1:Prod	Eyeglass renames SmartConnect Zone prod.example.com to igls-original-prod.example.com	
subnet1:DR		Eyeglass creates SmartConnect Zone alias prod.example.com
subnet0:synciq-prod	no changes	

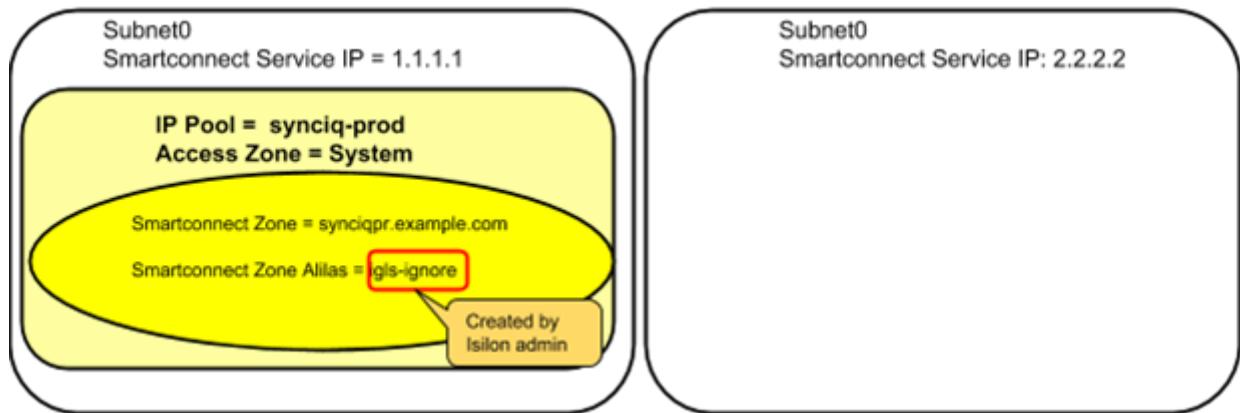


- Smartconnect Zone or Smartconnect Zone alias created or updated by Eyeglass during Failover
- DNS entry manually updated by DNS Admin after failover completed to change IP address for Smartconnect Zone to the Smartconnect Service IP on the new Source Cluster

FAILOVER AGAIN - CLUSTER 2 to CLUSTER 1

Subnet and Pool	Cluster 1	Cluster 2
subnet1:Prod	<p>Eyeglass renames SmartConnect Zone igls-original-prod.example.com to prod.example.com</p> <p>Eyeglass creates SmartConnect Zone alias dr.example.com</p>	
subnet1:DR		<p>Eyeglass renames SmartConnect Zone dr.example.com to igls-original-dr.example.com</p> <p>Eyeglass removes SmartConnect Zone alias prod.example.com created by previous failover</p>
subnet0:synciq-prod	no changes	





Smartconnect Zone or Smartconnect Zone alias created or updated by Eyeglass during Failover



DNS entry manually updated by DNS Admin after failover completed to change IP address for Smartconnect Zone to the Smartconnect Service IP on the new Source Cluster

© Superna Inc

2.15. IP Pool Failover

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

- [Overview](#)
- [Limitations](#)
- [How to Setup and Configure IP Pool Failover - Overview Video](#)
- [IP Pool Failover](#)
- [Overview](#)
- [Prerequisites](#)
- [Configuration Diagram](#)
- [Policy - Pool Mapping Diagram](#)
- [Configuration Steps:](#)
 - [Example of IP Pool Failover](#)
 - [Pool Failover Source \(Pool2\) ⇒ Target \(Pool2\)](#)
 - [Pool Failback Target \(Pool2\) ⇒ Source \(Pool2\)](#)

[IP Pool Failover](#)

This feature is available in Eyeglass version > 2.0.

Overview

Pool based failover, all SmartConnect names and aliases on a pool use standard Access Zone failover logic to failover.

Limitations

1. All IP pools must be within the same groupnet and use the same subnet service IP

How to Setup and Configure IP Pool Failover - Overview Video

The following video provides an overview tutorial on how to setup and configure IP pool failover feature:

IP Pool Failover

Overview

This feature provides greater flexibility to failover data allowing active/active data within an Access Zone on 2 clusters. This feature can provide Access Zone failover logic for DFS,SMB and NFS data. It has all the same requirements as Access Zone failover plus some additional steps to align SyncIQ policies to SmartConnect names (IP pools). The DR Dashboard has been updated to show failover status per IP pool. DR Assistant has been updated to allow 1 or

more pools to be selected for failover. DR Readiness validations are now completed per IP pool per Access Zone.

Support Requirements:

1. Each Access zone MUST have an equal number of pools on the source and target cluster. Odd number IP pools between cluster pairs is not supported. Example 2 pools on source cluster and 1 pool on the target cluster
2. **ALL Pools MUST be configured for pool failover or the Access Zone is in unsupported failover configuration.**
3. **FAN-IN of a mutiple IP pools on cluster A sharing a common IP pool on cluster B is NOT supported**
4. It is not supported to configure “partial” Pool Failover for an Access Zone. You must configure Pool Failover for all pools and SyncIQ Policies in an Access Zone if you are using this failover mode.
5. Once Pool Failover is configured for an Access Zone, that Access Zone will only appear in the Pool Readiness section of the DR Dashboard. It will not be displayed in the DR Dashboard Zone Readiness section. The policies in that Access Zone will also only be displayed in Pool Readiness view. They are not displayed in the Policy Readiness or DFS Readiness view.

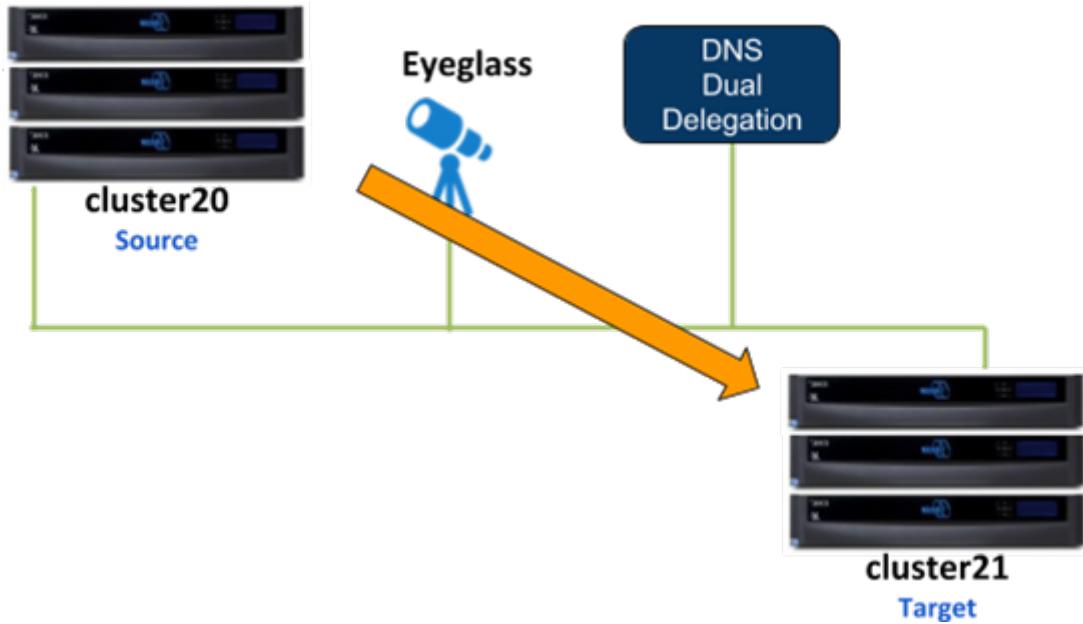
6. Once an Access Zone is configured for Pool Failover it must be failed over in SmartConnect/IP Pool failover mode from the DR Assistant. It will not be available for failover from the DR Assistant in Access Zone mode nor will the associated SyncIQ policies be able to be failed over in SyncIQ Policy or Microsoft DFS mode.
7. Basic Runbook Robot is not supported for Pool Failover System Access Zone as it is not available for pool to SyncIQ Policy mapping. Advanced DR Robot configuration for Access Zone failover must be used.
8. Pool Failover is not supported for multi-site A -> B and A -> C failover.

Prerequisites

- The IP Pool Failover is available in Eyeglass release > 2.0.
- SyncIQ policies are mapped to IP pools. This requirement means all SmartConnect names assigned to the pool will failover with the SyncIQ policies mapped to the pool. One or more policies can be mapped to the pool in Eyeglass. **Note: This may require changing your policy design to allow pool based failover.**

Configuration Diagram

The following diagram illustrates the basic configuration for the case of 1 Source cluster and 1 Target cluster.



As an example for this IP Pool failover setup, dual non-overlap IP Pools are configured on the source cluster and on the target cluster. All IP Pools are assigned to the System Access Zone.

SMB share

- **s-smb01** is configured on **Source** cluster

SMB shares are managed through Microsoft DFS server.

NFS export:

- **/ifs/data/s-nfs01** export is configured on **Source** cluster

SyncIQ Policies:

SyncIQ Policies are configured as follow:

SyncIQ Policy	Data	From	To
s01-t01-synciq01	SMB	Source	Target

s01-t01-synciq02	NFS	Source	Target
------------------	-----	--------	--------

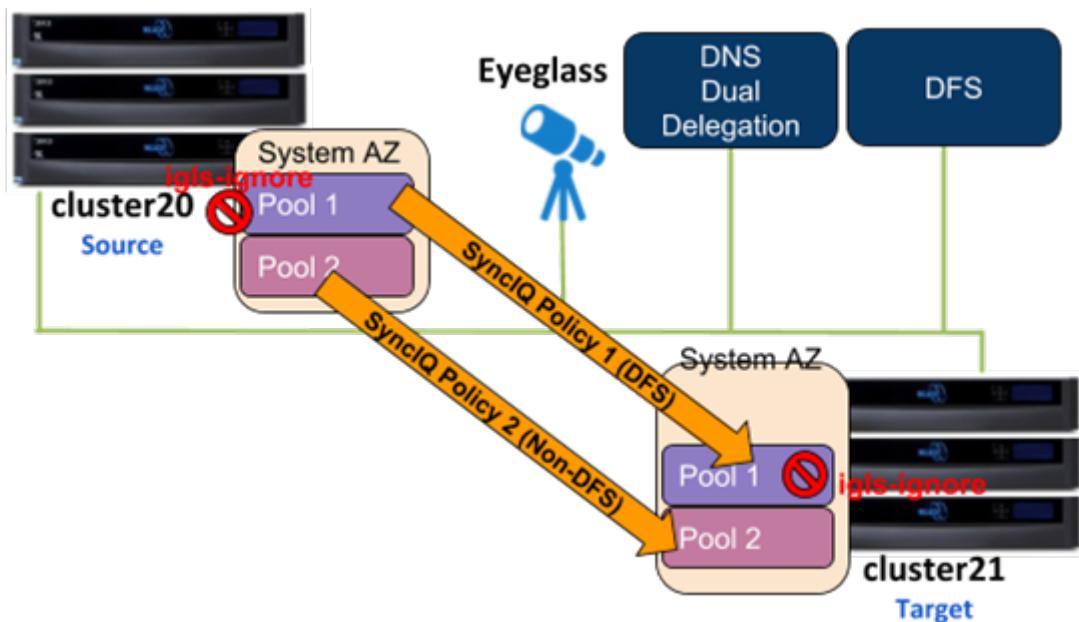
Policy - Pool Assignment

This table illustrates the SyncIQ Policies to Pools assignments:

From		To		Policy	Data
Cluster	Pool	Cluster	Pool		
Source	Pool 1	Target	Pool 1	s01-t01-synciq01	SMB (DFS)
Source	Pool 2	Target	Pool 2	s01-t01-synciq02	NFS

Policy - Pool Mapping Diagram

The following diagram shows the IP Pools and Policies mapping setup between Source and Target clusters

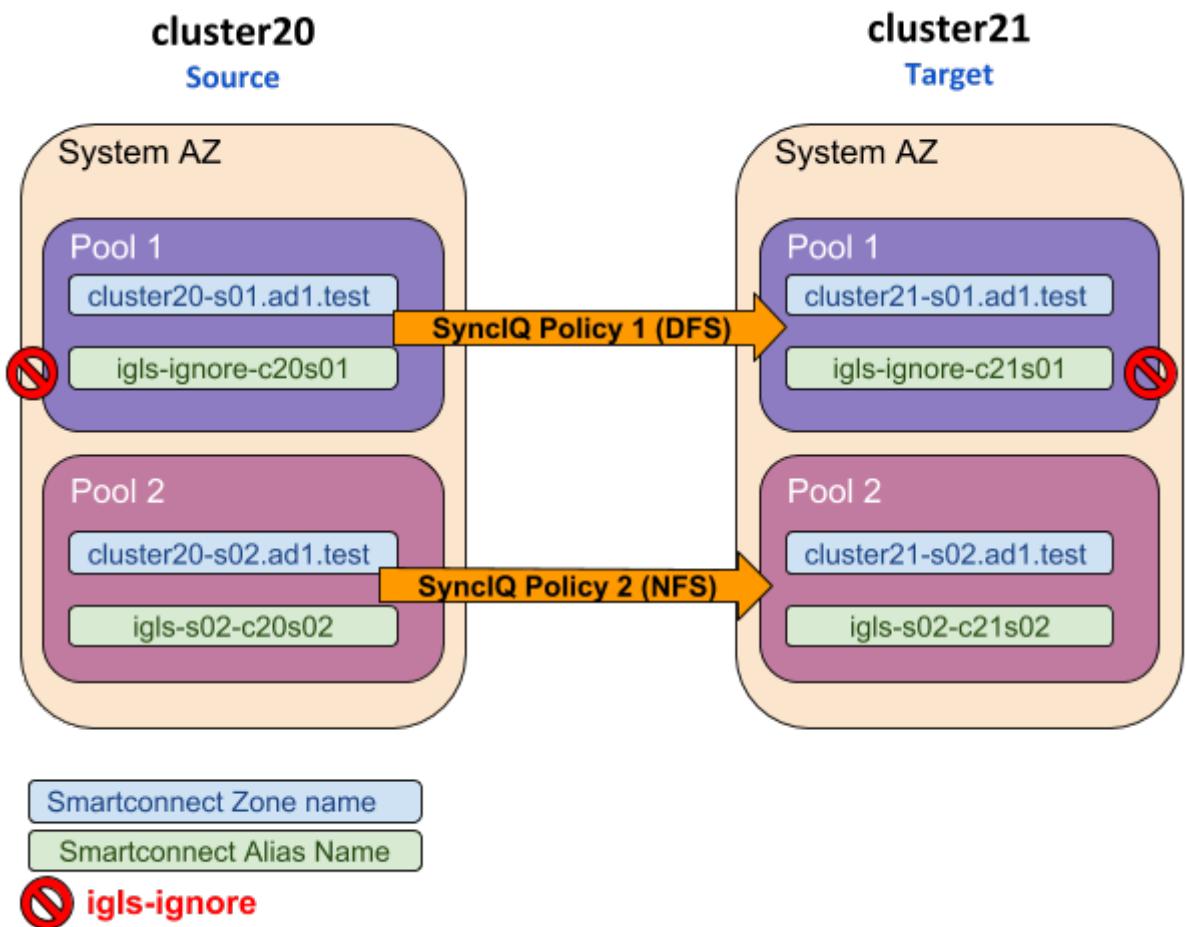


The *igls* mapping hints including *igls-ignore* also need to be configured for this dual IP Pool setup.

Based on that diagram, the following is the example of the *igls* and *igls-ignore* hint mappings:

From			To		
Cluster	Pool	SmartConnect Name / Alias (igls hints)	Cluster	Pool	SmartConnect Name / Alias (igls hints)
Source	Pool 1	cluster20-s01.ad1.test	Target	Pool 1	cluster21-s01.ad1.test
		igls-ignore-c20s01			igls-ignore-c21s01
Source	Pool 2	cluster20-s02.ad1.test	Target	Pool 2	cluster21-s02.ad1.test
		igls-s02-c20s02			igls-s02-c21s02

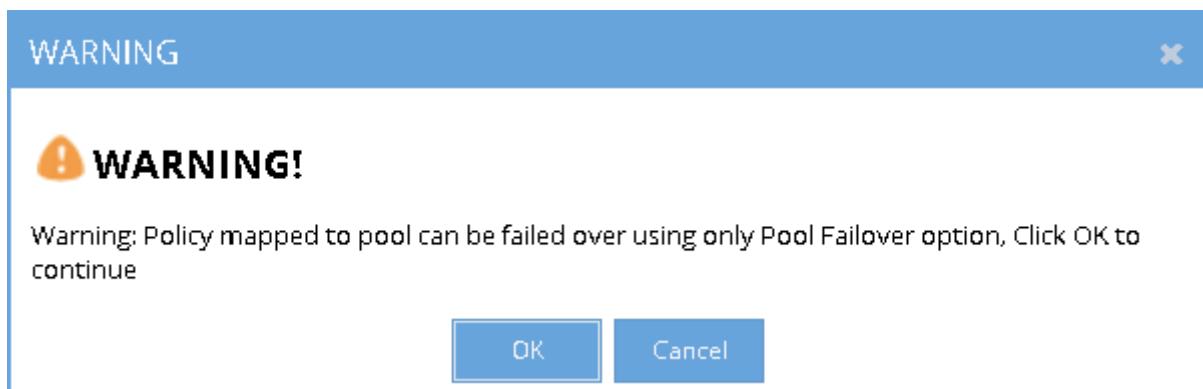
That configuration also can be seen in the following diagram:



Configuration Steps:

1. Configure the required `igls` and `igls-ignore` mapping hints for all the Pools in the Access Zone.

2. Verify the **DR Dashboard - Zone Readiness** that Network Mappings have been configured correctly.
3. Configure the Advanced Network Mapping to map the Policies to the Pools as specified in the Policies - Pool Mapping table. To assign the policy to the pool, from **DR Dashboard - Pool Readiness** click **Advanced Network Mapping** button. Eyeglass will prompt a warning for configuring this Advanced Network Mapping: "*Warning: Policy mapped to pool can be failed over using only Pool Failover option. Click OK to continue.*" Click OK to continue.

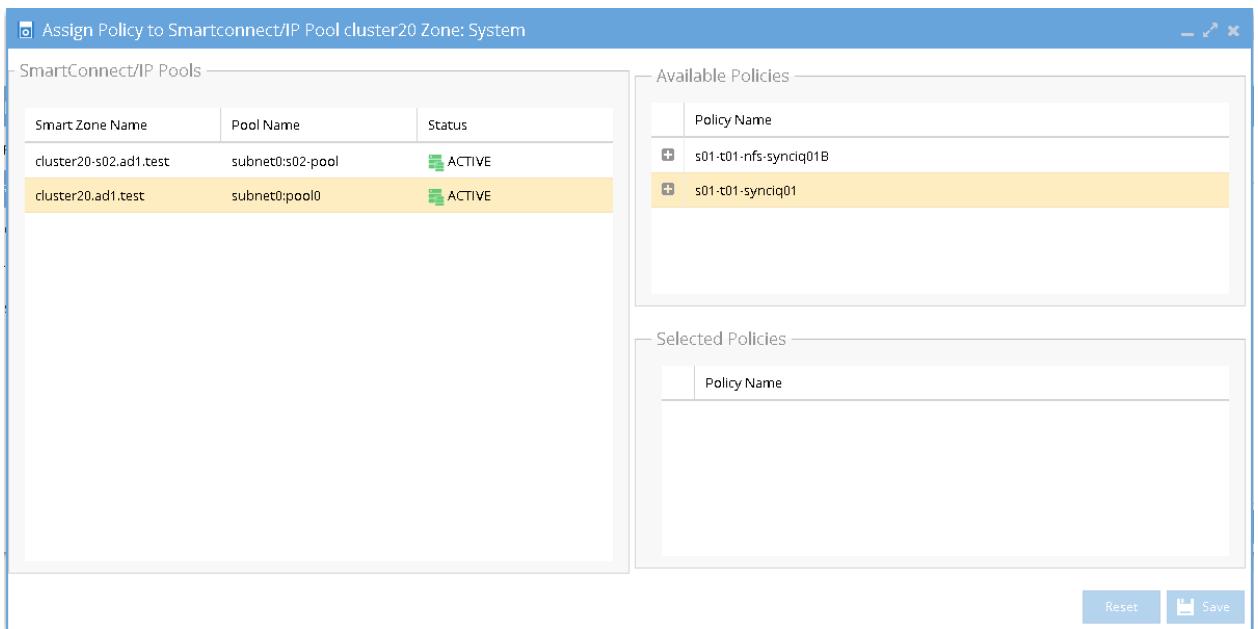


4. Select the Access Zone for specific Source-Target Cluster Pair to configure and then click Next button. In the **Assign Policy to SmartConnect/IP Pool** configuration window, select a SmartConnect/Pool from the list and then drag and drop the correct policy from the Available Policies section (under the Policy Name column) to the Selected Policies section. Click Save to save the modification for each policy as it is selected.

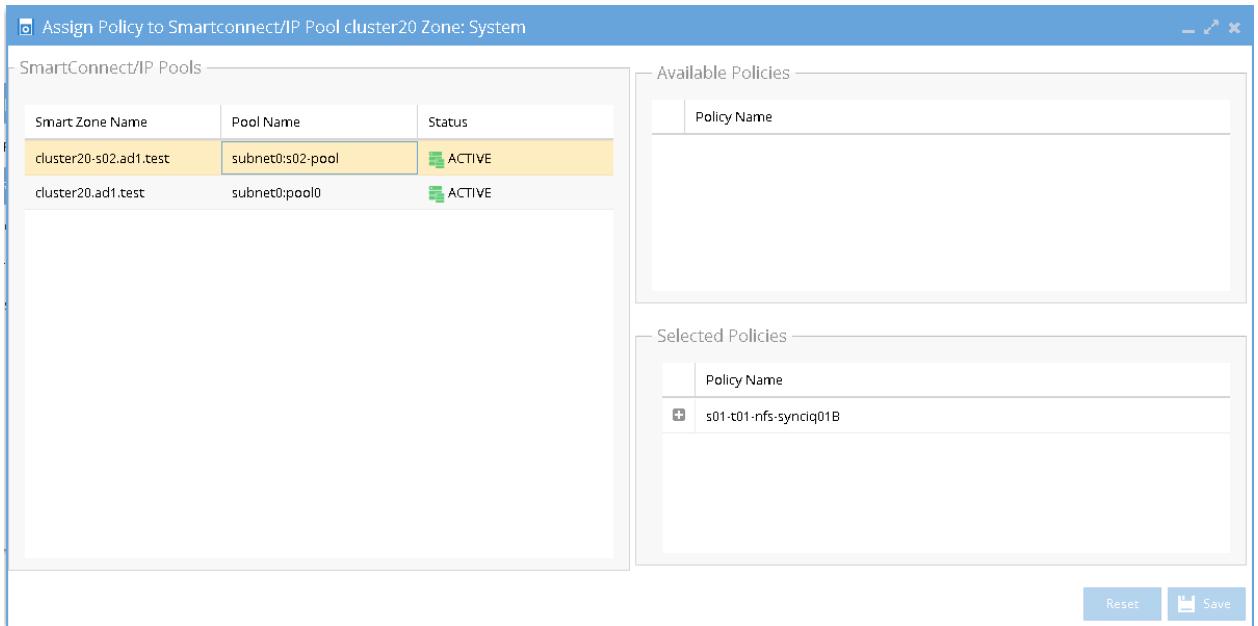
IMPORTANT: Please ensure the policies are mapped to the correct pools.

IMPORTANT: Policies where the Eyeglass Configuration Replication mode is set to DFS will only be displayed for selection when the Pool has an igls-ignore hint.

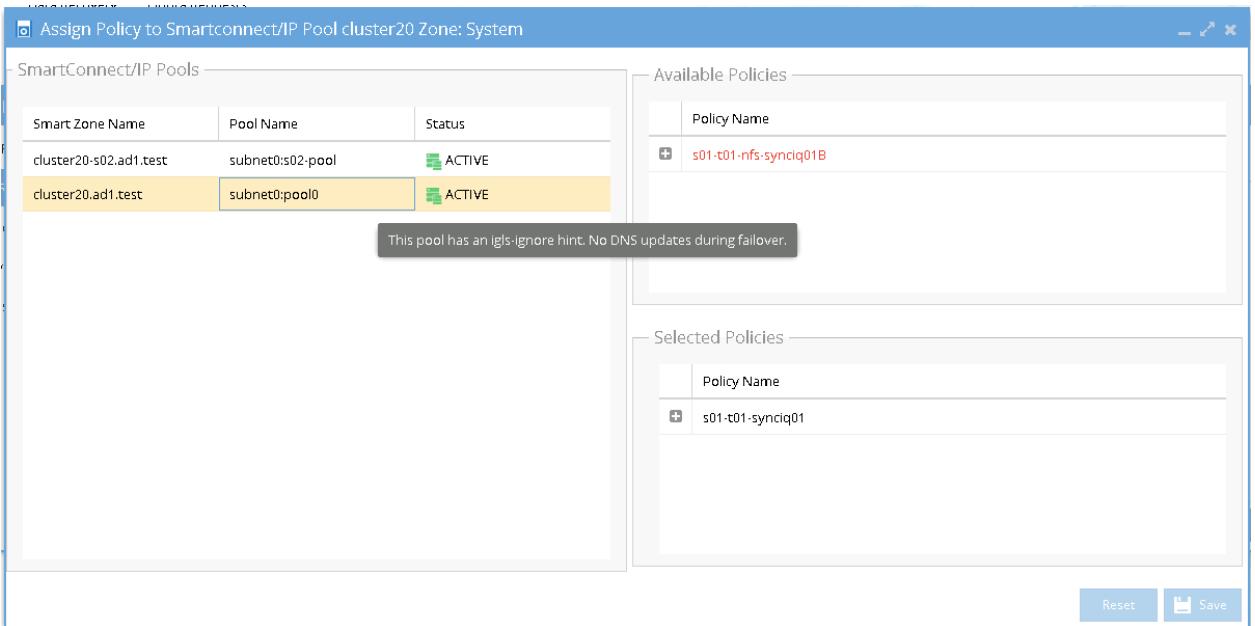
Example:



Once assigned we can see the following mappings:



The IP Pool that has the *igls-ignore* alias is listed in the SmartConnect/IP Pools section with description: "The pool has an *igls-ignore* hints. No DNS updates during failover". This description can be seen by hovering mouse over that IP pool area.



- Verify that in DR Dashboard - Pool Readiness the **DR Failover Status** is showing no error.

Example:

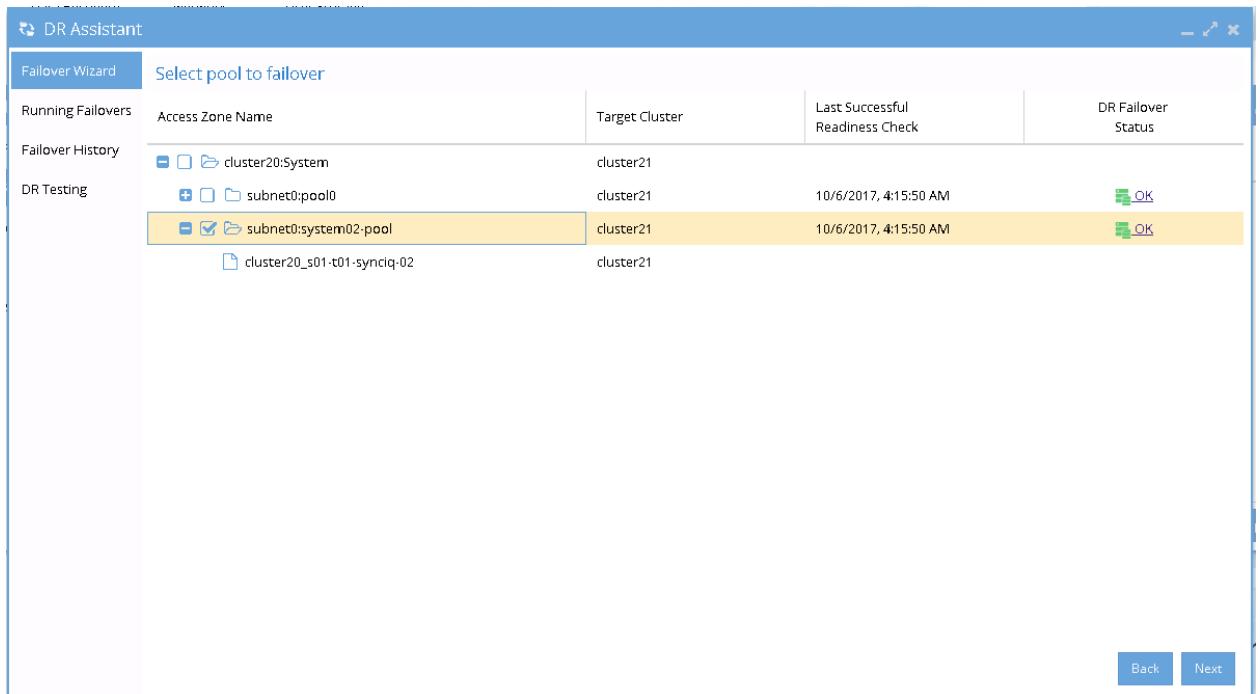
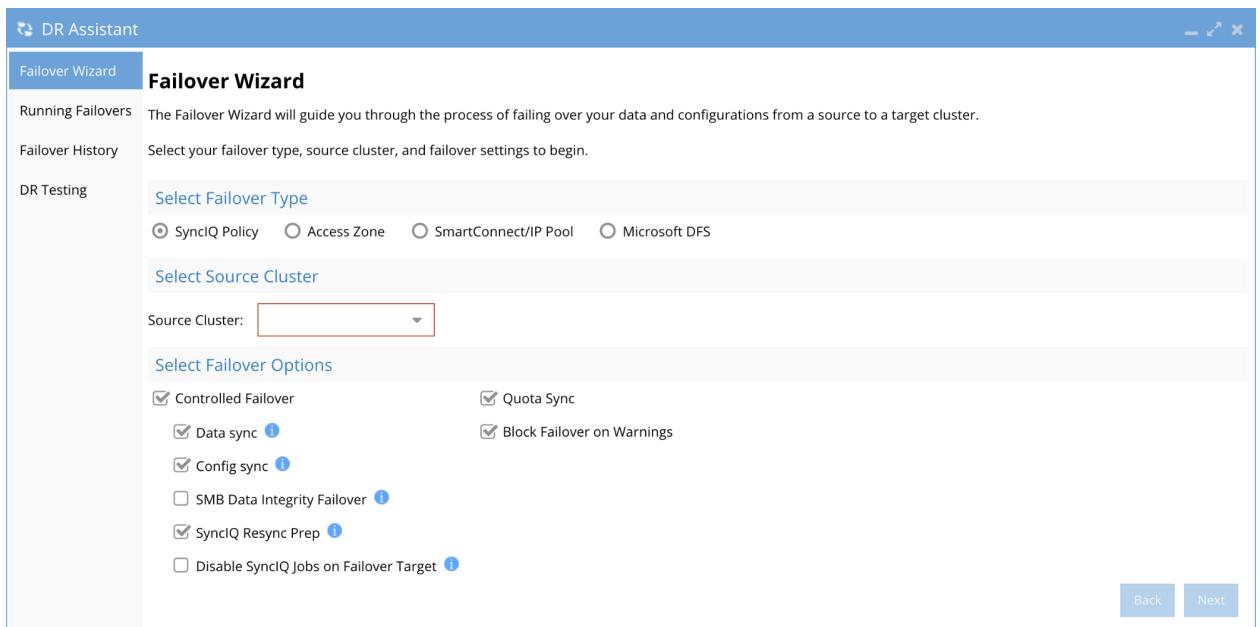
Zone Readiness	Access Zone Name	Pool Mapping	Target Cluster	Last Successful Readiness Check	Map Policy to Pool	DR Failover Status
Pool Readiness	cluster20:System + cluster20:subnet0:pool0		cluster21	10/6/2017, 4:00:50 AM	Map Now	OK
DFS Readiness	+ cluster20:subnet0:pool0	View Map	cluster21	10/6/2017, 4:00:50 AM	Map Now	OK
Policy Readiness	+ cluster20:subnet0:system02-pool	View Map	cluster21	10/6/2017, 4:00:50 AM	Map Now	OK
DR Testing						

Advanced Network Mapping

6. SmartConnect/IP Pool Failover

To perform IP Pool failover, in the DR Assistant Failover Wizard select failover type as **SmartConnect/IP Pool** Failover, and then select IP pool to be failed over .

Example:



7. If we want to failover all the pools within the Access Zone with this IP Pool failover, select all the pools from the list of Select pool to failover.

Example:

Failover Wizard		Select pool to failover		
Running Failovers	Access Zone Name	Target Cluster	Last Successful Readiness Check	DR Failover Status
Failover History	cluster20:System	cluster21		
DR Testing	<input checked="" type="checkbox"/> <input checked="" type="checkbox"/> <input type="checkbox"/> subnet0:pool0 <input checked="" type="checkbox"/> <input checked="" type="checkbox"/> <input type="checkbox"/> subnet0:system02-pool	cluster21	10/6/2017, 4:15:50 AM	
		cluster21	10/6/2017, 4:15:50 AM	

8. After Pool failover job has completed, we can see which pool has been failed over from the **Pool Readiness - DR Failover Status**.

Zone Readiness	Access Zone Name	Pool Mapping	Target Cluster	Last Successful Readiness Check	Map Policy to Pool	DR Failover Status
Pool Readiness	cluster20:System		cluster21			
DFS Readiness	+ cluster20:System	View Map	cluster21	10/6/2017, 4:23:57 AM	Map Now	
Policy Readiness	+ subnet0:pool0	View Map	cluster21	10/6/2017, 4:23:57 AM	Map Now	
DR Testing	+ subnet0:system02-pool	View Map	cluster21	10/6/2017, 4:23:57 AM	Map Now	
	cluster21:System		cluster20			

Example of IP Pool Failover

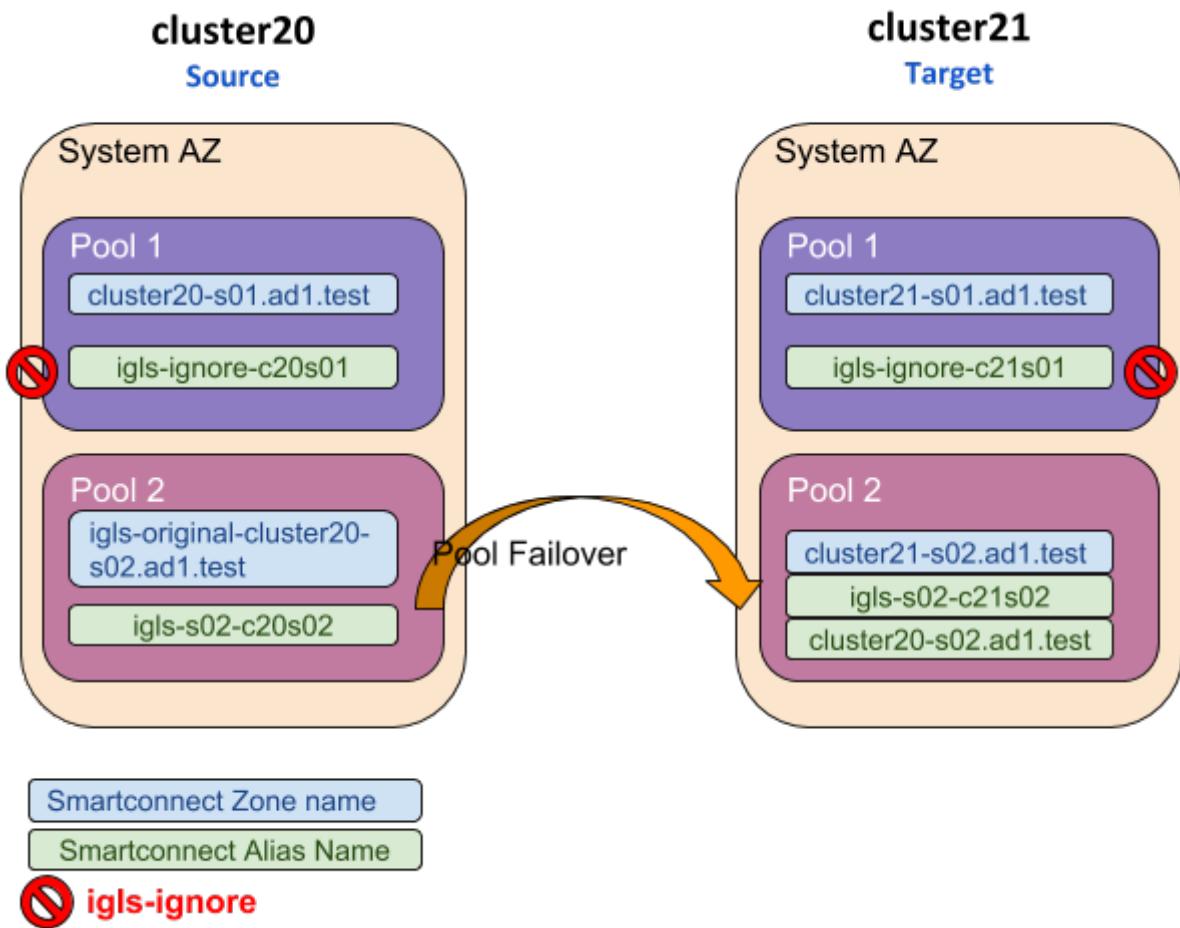
The following examples illustrate the changes of the SmartConnect zone names and alias names for IP Pool failover in the following sequences:

1. Pool Failover Source (Pool 2) to Target (Pool 2)
2. Pool Failback Target (Pool 2) to Source (Pool 2)

Pool Failover Source (Pool2) \Rightarrow Target (Pool2)

The following table and diagram illustrate the SmartConnect zone names and alias names after performing Failover from Source (Pool 2) to Target (Pool 2).

From			To		
Cluster	Pool	SmartConnect Name / Alias (igls hints)	Cluster	Pool	SmartConnect Name / Alias (igls hints)
Source	Pool 1	cluster20.ad1.test	Target	Pool 1	cluster21.ad1.test
		igls-ignore-c20s01			igls-ignore-c21s01
Source	Pool 2	igls-original-cluster20-s02.ad1.test	Target	Pool 2	cluster21-s02.ad1.test
		igls-s02-c20s02			igls-s02-c21s02 cluster20-s02.ad1.test

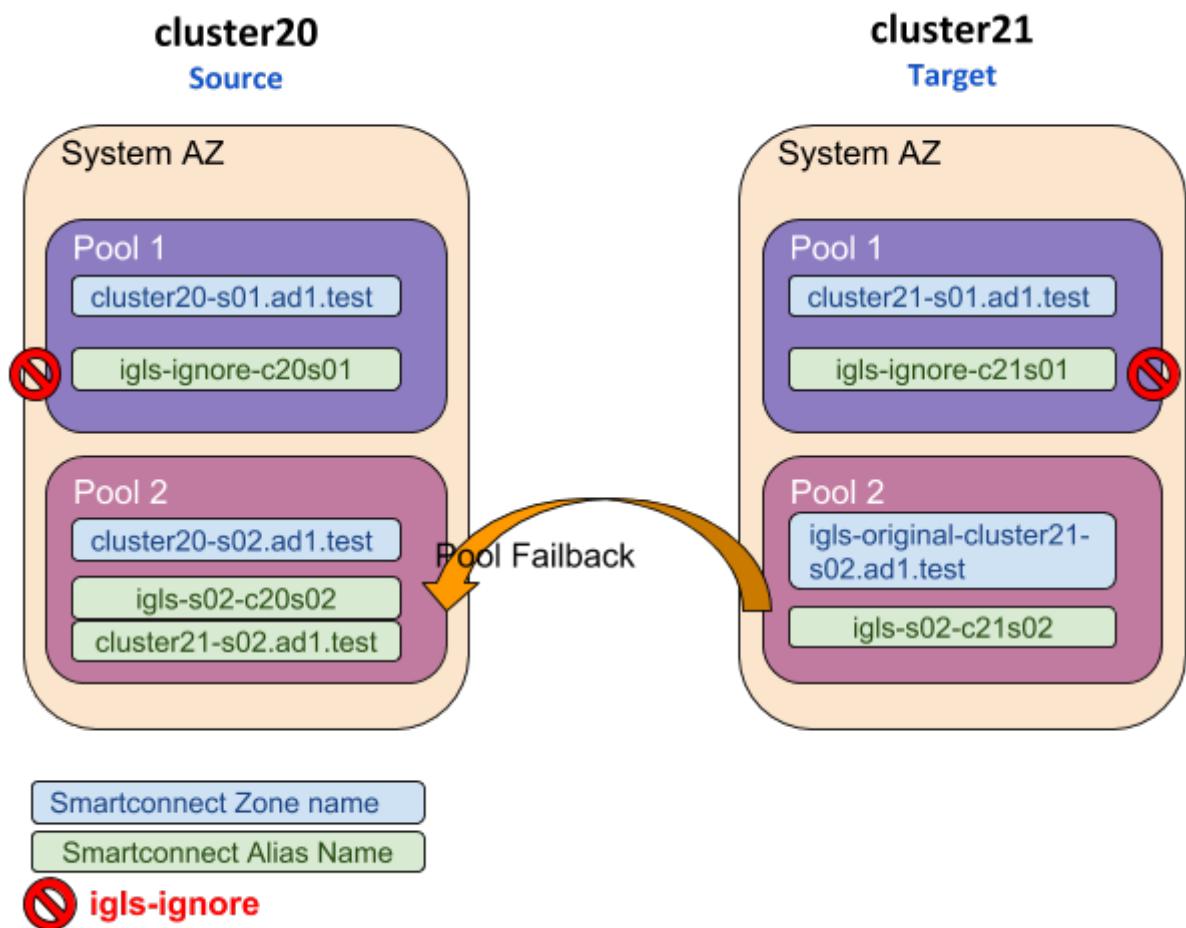


Pool Failback Target (Pool2) \Rightarrow Source (Pool2)

The following table and diagram illustrate the SmartConnect zone names and alias names after performing Failback from Target (Pool 2) to Source (Pool 2).

From			To		
Cluster	Pool	SmartConnect Name / Alias (igls hints)	Cluster	Pool	SmartConnect Name / Alias (igls hints)
Target	Pool 1	cluster21.ad1.test	Source	Pool 1	cluster20.ad1.test
		igls-ignore-c21s01			igls-ignore-c20s01
Target	Pool 2	igls-original-cluster21-s02.ad1.test	Source	Pool 2	cluster20-s02.ad1.test

		igls-s02-c21s02			igls-s02-c20s02
					cluster21-s02.ad1.test



© Superna Inc

2.16. Fan-In IP Pool Failover

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

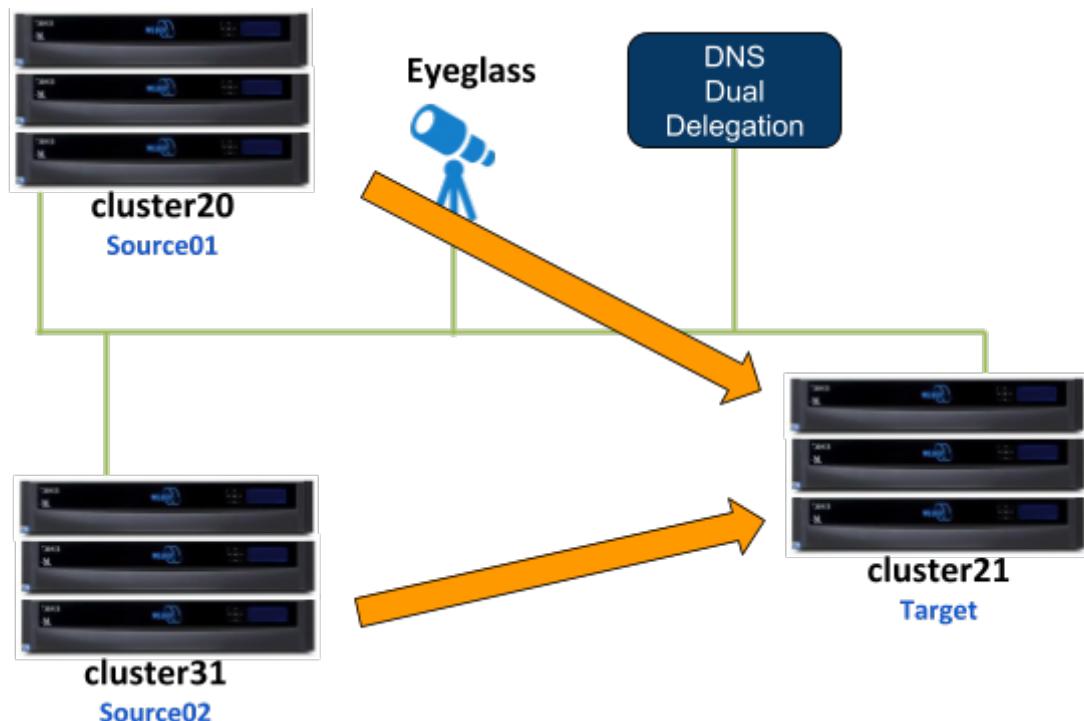
Fan-In IP Pool Failover

Prerequisites

- The target cluster must have a dedicated IP pool for each source cluster IP pool as a failover target. Example 3 source clusters with 1 IP each will require the target cluster to have 3 IP pools.

Fan-In configuration Diagram

The following diagram gives an example of Fan-In configuration topology



This configuration consists of 3 clusters, 2 of them are the source clusters and third one is the target cluster. For this example:

- **Cluster20** is the **Source01** Cluster
- **Cluster31** is the **Source02** Cluster
- **Cluster21** is the **Target** Cluster

For the IP Pool failover setup, dual non-overlap IP Pools are configured on the source clusters and four non-overlap IP Pools are configured on the target cluster. All IP Pools are assigned to the System Access Zone.

SMB shares

- **source01-smb01** is configured on **Source01** cluster
- **source02-smb01** is configured on **Source02** cluster

SMB shares are managed through Microsoft DFS server.

NFS exports:

- /ifs/data/source01-nfs01 export is configured on **Source01** cluster
- /ifs/data/source02-nfs01 export is configured on **Source02** cluster

SyncIQ Policies:

Configured the following SyncIQ Policies:

SyncIQ Policy	Data	From	To
s01-t01-synciq01	SMB	Source01	Target
s01-t01-synciq02	NFS	Source01	Target
s02-t01-synciq03	SMB	Source02	Target
s02-t01-synciq04	NFS	Source02	Target

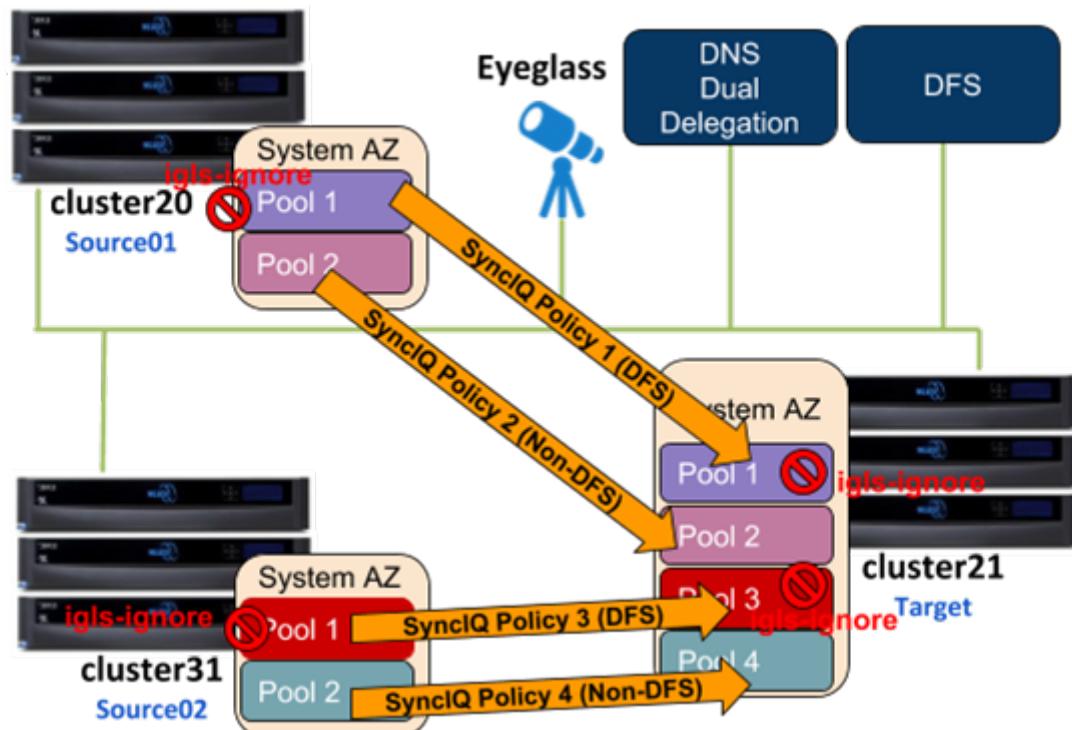
SyncIQ Policy to Pool Assignments

The following table illustrates the SyncIQ Policies to Pools assignments

From		To		Policy	Data
Cluster	Pool	Cluster	Pool		
Source01	Pool 1	Target	Pool 1	s01-t01-synciq01	SMB (DFS)
Source01	Pool 2	Target	Pool 2	s01-t01-synciq02	NFS
Source02	Pool 1	Target	Pool 3	s02-t01-synciq03	SMB (DFS)
Source02	Pool 2	Target	Pool 4	s02-t01-synciq04	NFS

SyncIQ Policies - Pools Mapping Diagram

The following diagram shows the SyncIQ Policies and Pools mapping setup

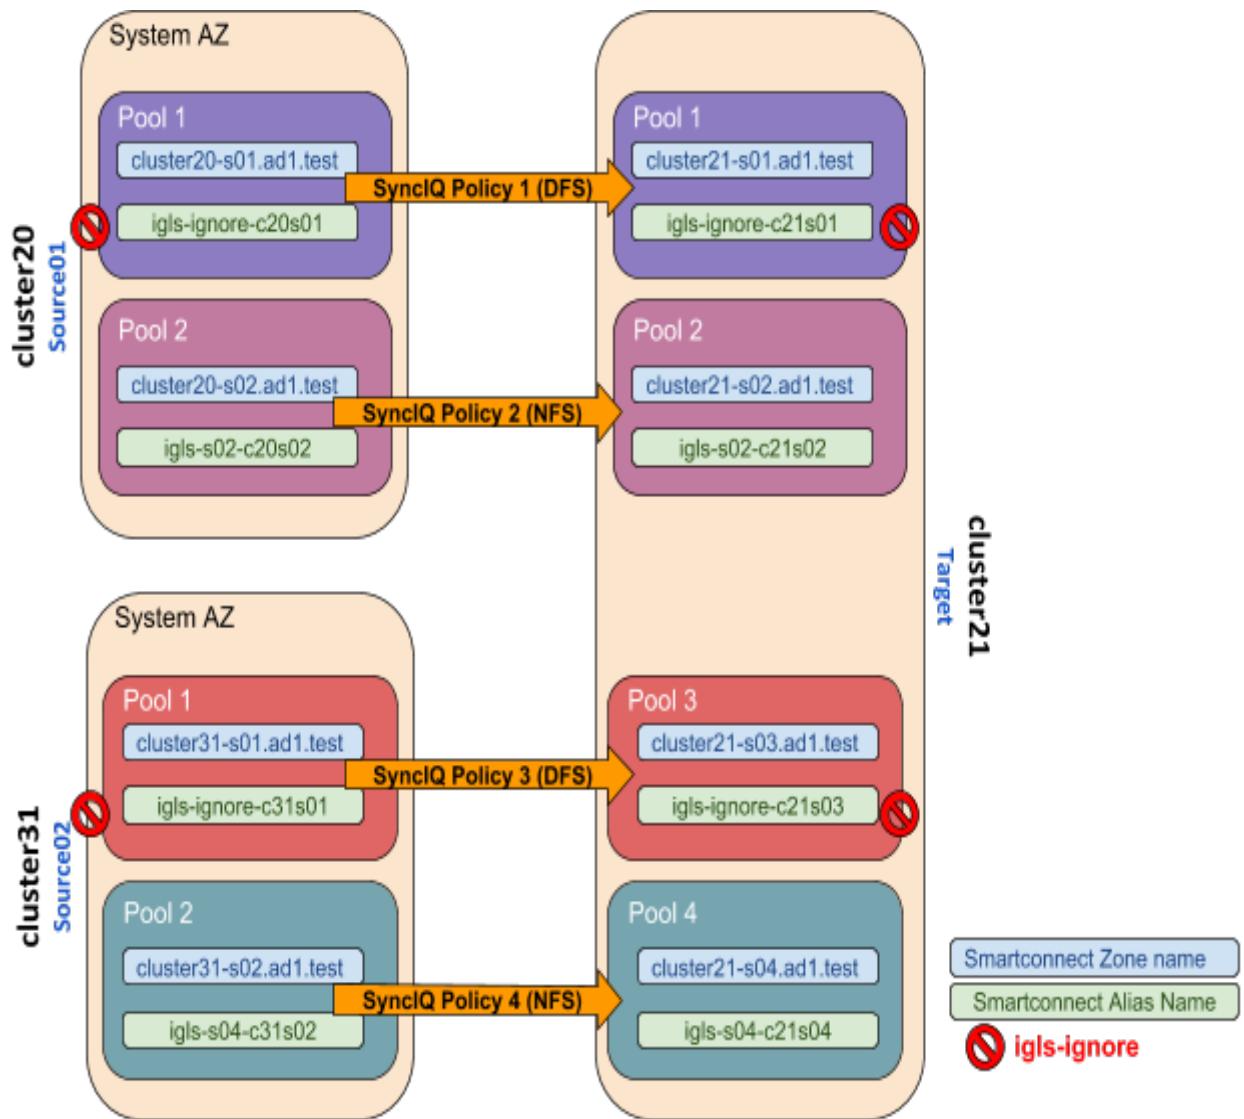


The *igls* mapping hints including *igls-ignore* also need to be configured for this dual IP Pool setup.

Based on the above diagram, the following is the example of the *igls* and *igls-ignore* mapping hints:

From			To		
Cluster	Pool	SmartConnect Name / Alias (igls hints)	Cluster	Pool	SmartConnect Name / Alias (igls hints)
Source01	Pool 1	cluster20-s01.ad1.test	Target	Pool 1	cluster21-s01.ad1.test
		igls-ignore-c20s01			igls-ignore-c21s01
Source01	Pool 2	cluster20-s02.ad1.test	Target	Pool 2	cluster21-s02.ad1.test
		igls-s02-c20s02			igls-s02-c21s02
Source02	Pool 1	cluster31-s01.ad1.test	Target	Pool 3	cluster21-s03.ad1.test
		igls-ignore-c31s01			igls-ignore-c21s03
Source02	Pool 2	cluster31-s02.ad1.test	Target	Pool 4	cluster21-s04.ad1.test
		igls-s04-c31s02			igls-s04-c21s04

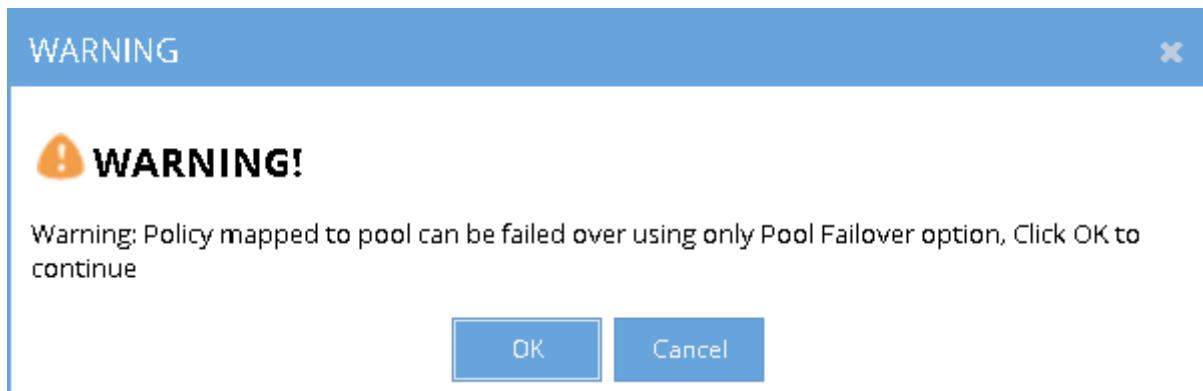
The mapping also can be seen in the following diagram:



Configuration Steps:

1. Configure the required `igls` and `igls-ignore` mapping hints for all the Pools in the Access Zone.
2. Verify from **DR Dashboard - Zone Readiness** that Network Mappings have been configured correctly.
3. Configure the **Advanced Network Mapping** to map the Policies to the Pools as specified in the Policies - Pool Mapping table. To assign the policy to the pool, from **DR Dashboard - Pool Readiness** click **Advanced Network Mapping** button. Eyeglass will prompt a warning for configuring this

Advanced Network Mapping: "Warning: Policy mapped to pool can be failed over using only Pool Failover option. Click OK to continue." Click OK to continue.



4. Select the Access Zone for specific Source-Target Cluster Pair to configure and then click Next button. In the **Assign Policy to SmartConnect/IP Pool** configuration window, select a SmartConnect/Pool from the list and then drag and drop the correct policy from the Available Policies section (under the Policy Name column) to each pool. Click Save to save the modification. Note: please ensure the policies are mapped to the correct pools.

Example:

A screenshot of the 'Assign Policy to Smartconnect/IP Pool' configuration window. The window title is 'Assign Policy to Smartconnect/IP Pool cluster31 Zone: System'.
SmartConnect/IP Pools:

Smart Zone Name	Pool Name	Status
cluster31-s02.ad1.test	subnet0:s02-pool	ACTIVE
cluster31.ad1.test	subnet0:pool0	ACTIVE

Available Policies:

Policy Name

Selected Policies:

Policy Name
s02-t02-nfs-synciq02B

Buttons:

- Reset
- Save

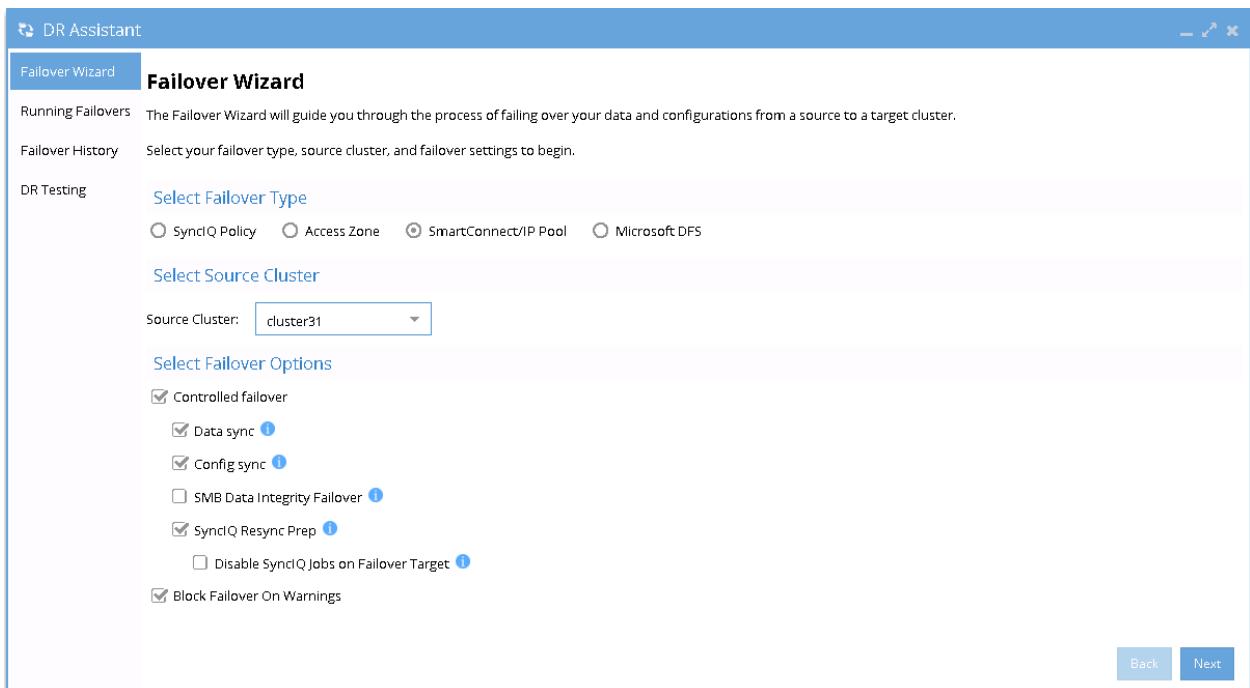
- Verify that the **SmartConnect/IP Pool Failover Readiness Status** is showing no error.

Example:

Zone Readiness	Access Zone Name	Pool Mapping	Target Cluster	Last Successful Readiness Check	Map Policy to Pool	DR Failover Status
Pool Readiness	cluster20:System		cluster21			
DFS Readiness	+ cluster21:subnet0:pool0	View Map	cluster21	10/6/2017, 4:30:37 AM	Map Now	OK
Policy Readiness	+ cluster21:subnet0:s02:pool	View Map	cluster21	10/6/2017, 4:30:37 AM	Map Now	OK
DR Testing	+ cluster21:System + cluster21:subnet0:s03:pool + cluster21:subnet0:s04:pool	View Map View Map	cluster31 cluster31 cluster20	10/6/2017, 4:30:32 AM 10/6/2017, 4:30:32 AM	Map Now Map Now	FAILED OVER FAILED OVER
	+ cluster21:cluster21:subnet0:pool0 + cluster21:cluster21:subnet0:s02:pool	View Map View Map	cluster20 cluster20	10/6/2017, 4:30:36 AM 10/6/2017, 4:30:36 AM	Map Now Map Now	FAILED OVER FAILED OVER
	+ cluster21:cluster21:subnet0:pool0 + cluster21:cluster21:subnet0:s02:pool	View Map View Map	cluster21 cluster21	10/6/2017, 4:30:30 AM 10/6/2017, 4:30:30 AM	Map Now Map Now	OK OK

6. SmartConnect/IP Pool Failover.

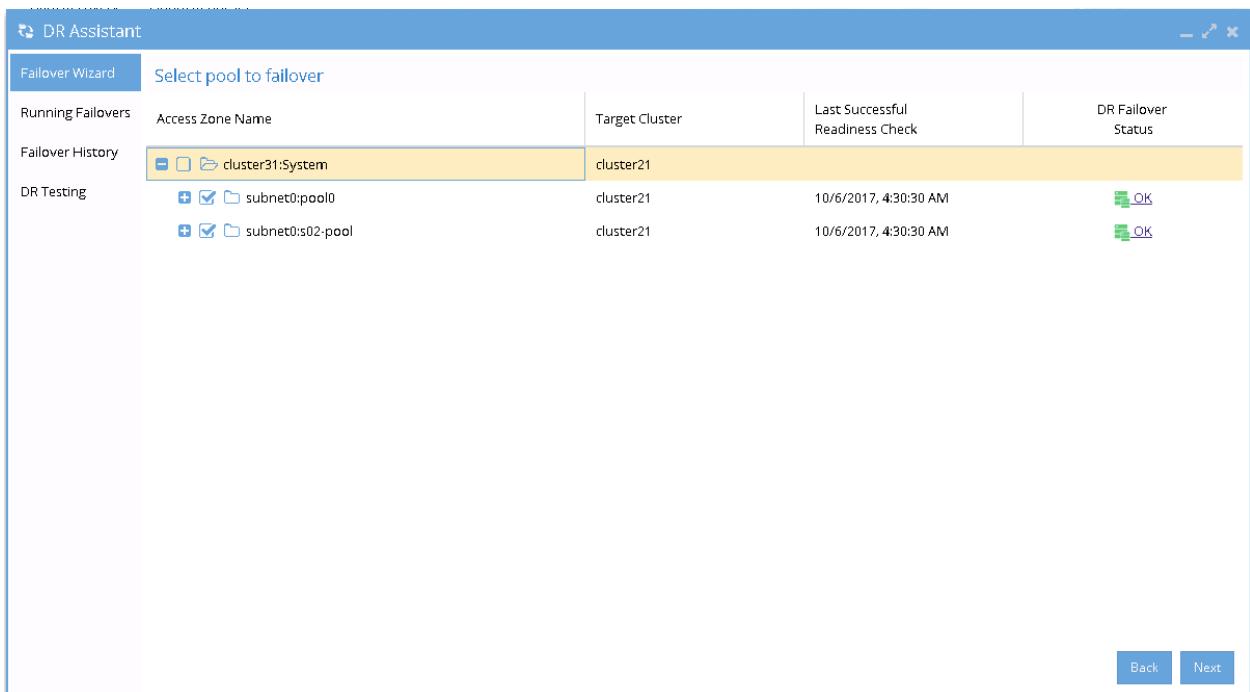
To perform IP Pool failover, in the DR Assistant Failover Wizard select failover type as **SmartConnect/IP Pool Failover**, and then select IP pool to be failed over.



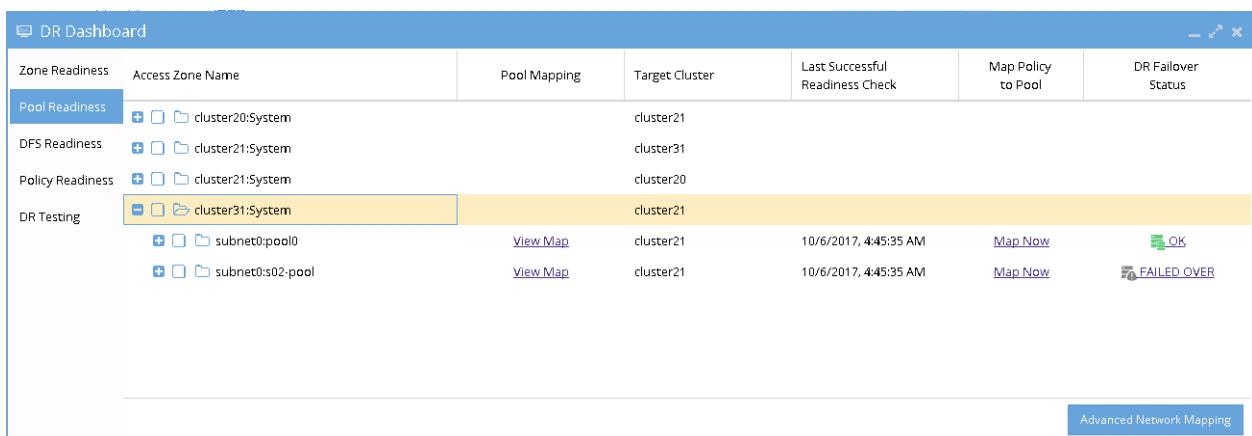
Select pool to failover				
Access Zone Name	Target Cluster	Last Successful Readiness Check	DR Failover Status	
cluster31:System	cluster21	10/6/2017, 4:30:30 AM		
subnet0:pool0	cluster21	10/6/2017, 4:30:30 AM		
subnet0:s02-pool	cluster21	10/6/2017, 4:30:30 AM		

Back Next

7. If we want to failover all the pools within the Access Zone with this IP Pool failover, select all the pools from the list of **Select pool to failover**.



8. After Pool failover job has completed, we can see which pool has been failed over from the **Pool Readiness - DR Failover Status**.



Example of Pool Failover

The following examples illustrate the changes of the SmartConnect zone names and alias names for IP Pool failover in the following sequences:

1. Pool Failover Source01 (Pool 2) to Target (Pool 2)

2. Pool Failover Source02 (Pool 2) to Target (Pool 4)

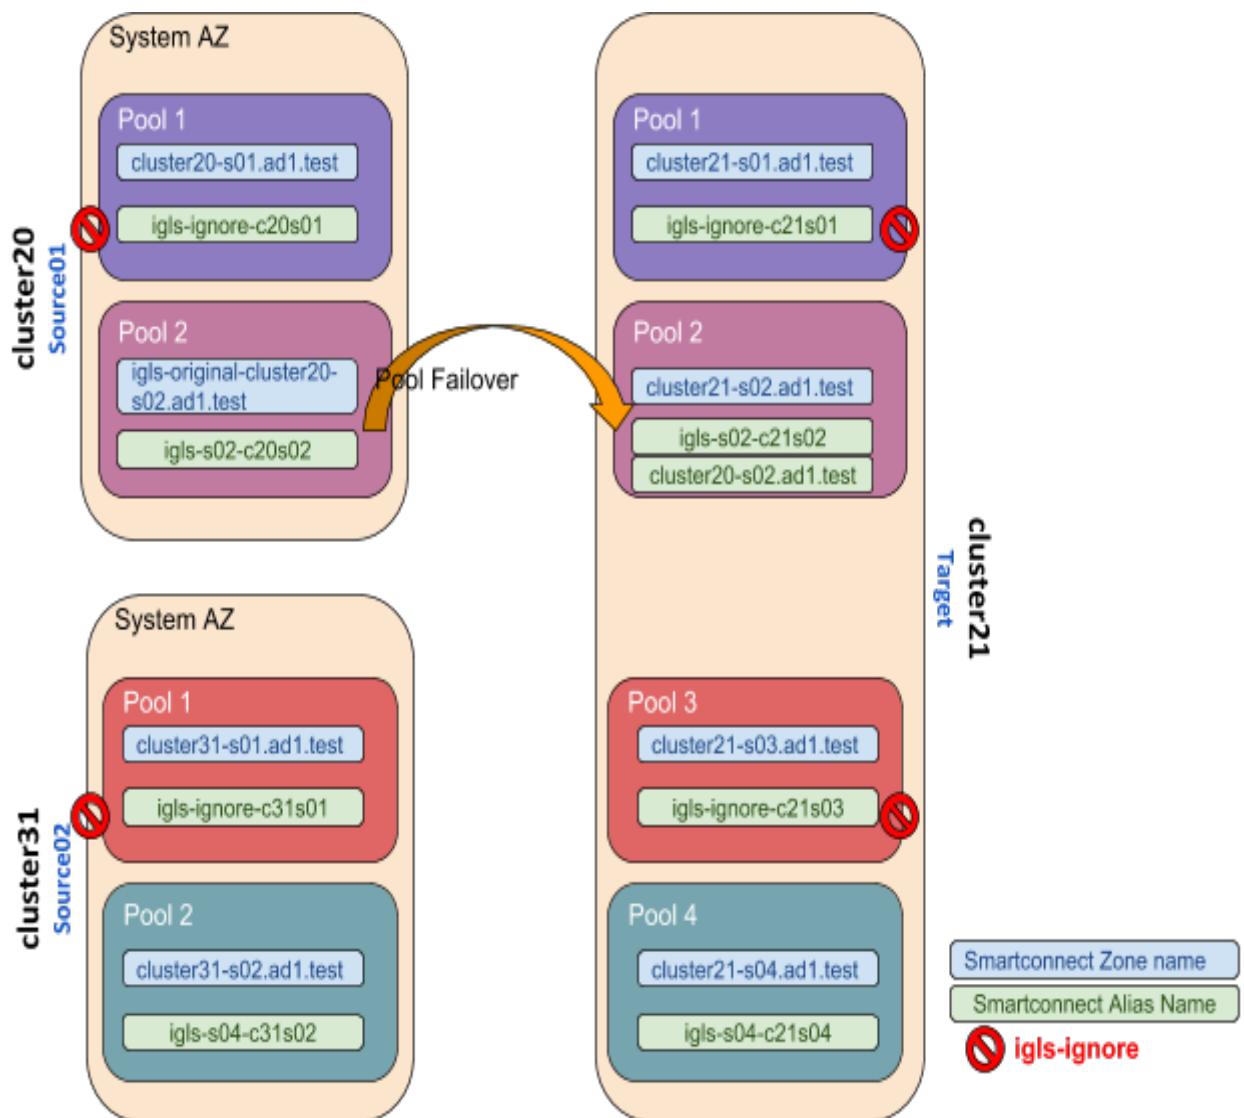
3. Pool Failback Target (Pool 2) to Source01 (Pool 2)

4. Pool Failback Target (Pool 4) to Source02 (Pool 2)

Pool Failover Source01 (Pool 2) ⇒ Target (Pool 2)

The following table and diagram illustrate the SmartConnect zone names and alias names after performing Failover from Source01 (Pool 2) to Target (Pool 2).

From			To		
Cluster	Pool	SmartConnect Name / Alias (igls hints)	Cluster	Pool	SmartConnect Name / Alias (igls hints)
Source01	Pool 1	cluster20-s01.ad1.test	Target	Pool 1	cluster21-s01.ad1.test
		igls-ignore-c20s01			igls-ignore-c21s01
Source01	Pool 2	igls-original-cluster20-s02.ad1.test	Target	Pool 2	cluster21-s02.ad1.test
		igls-s02-c20s02			igls-s02-c21s02 cluster20-s02.ad1.test
Source02	Pool 1	cluster31.ad1.test	Target	Pool 3	cluster21-s03.ad1.test
		igls-ignore-c31s01			igls-ignore-c21s03
Source02	Pool 2	cluster31-s02.ad1.test	Target	Pool 4	cluster21-s04.ad1.test
		igls-s04-c31s02			igls-s04-c21s04

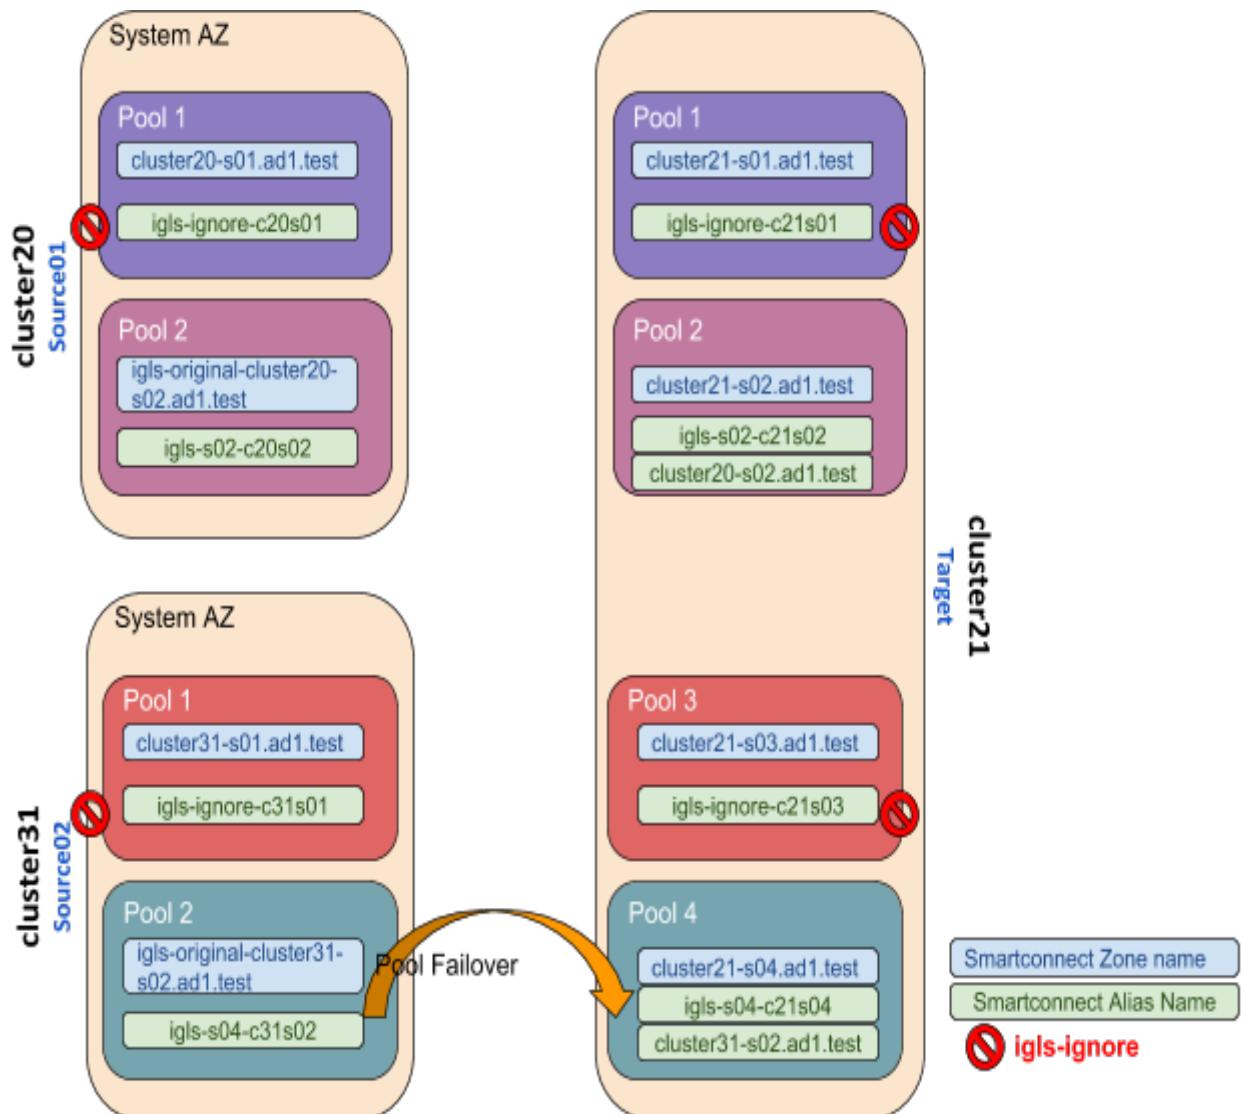


Pool Failover Source02 (Pool 2) ⇒ Target (Pool 4)

The following table and diagram illustrate the SmartConnect zone names and alias names after performing Failover from Source02 (Pool 2) to Target (Pool 4).

From			To		
Cluster	Pool	SmartConnect Name / Alias (igls hints)	Cluster	Pool	SmartConnect Name / Alias (igls hints)
Source01	Pool 1	cluster20-s01.ad1.test	Target	Pool 1	cluster21-s01.ad1.test
		igls-ignore-c20s01			igls-ignore-c21s01

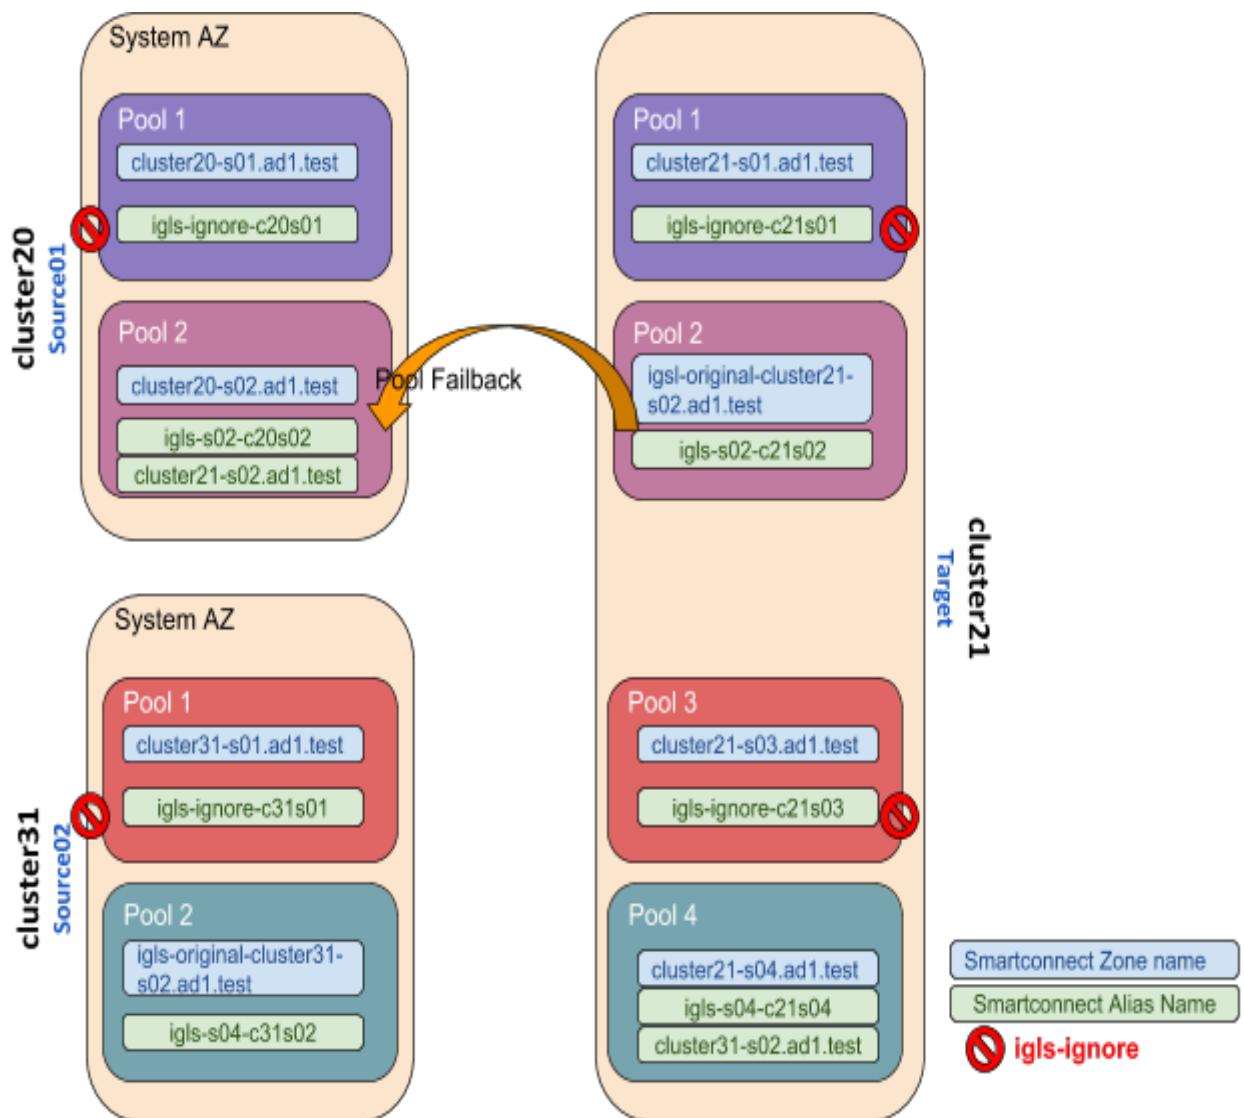
Source01	Pool 2	igls-original-cluster20-s02.ad1.test igls-s02-c20s02	Target	Pool 2	cluster21-s02.ad1.test igls-s02-c21s02 cluster20-s02.ad1.test
Source02	Pool 1	cluster31.ad1.test	Target	Pool 3	cluster21-s03.ad1.test
		igls-ignore-c31s01			igls-ignore-c21s03
Source02	Pool 2	igls-original-cluster31-s02.ad1.test	Target	Pool 4	cluster21-s04.ad1.test
		igls-s04-c31s02			igls-s04-c21s04 cluster31-s02.ad1.test



Pool Fallback Target (Pool 2) ⇒ Source01 (Pool 2)

The following table and diagram illustrate the SmartConnect zone names and alias names after performing Failback from Target (Pool 2) to Source01 (Pool 2).

From			To		
Cluster	Pool	SmartConnect Name / Alias (igls hints)	Cluster	Pool	SmartConnect Name / Alias (igls hints)
Target	Pool 1	cluster21-s01.ad1.test	Source01	Pool 1	cluster20-s01.ad1.test
		igls-ignore-c21s01			igls-ignore-c20s01
Target	Pool 2	igls-original-cluster21-s02.ad1.test	Source01	Pool 2	cluster20-s02.ad1.test
		Igls-s02-c21s02			igls-s02-c20s02 cluster21-s02.ad1.test
Target	Pool 3	cluster21-s03.ad1.test	Source02	Pool 1	cluster31.ad1.test
		igls-ignore-c21s03			igls-ignore-c31s01
Target	Pool 4	cluster21-s04.ad1.test	Source02	Pool 2	igls-original-cluster31-s02.ad1.test
		Igls-s04-c21s04 cluster31-s02.ad1.test			igls-s04-c31s02

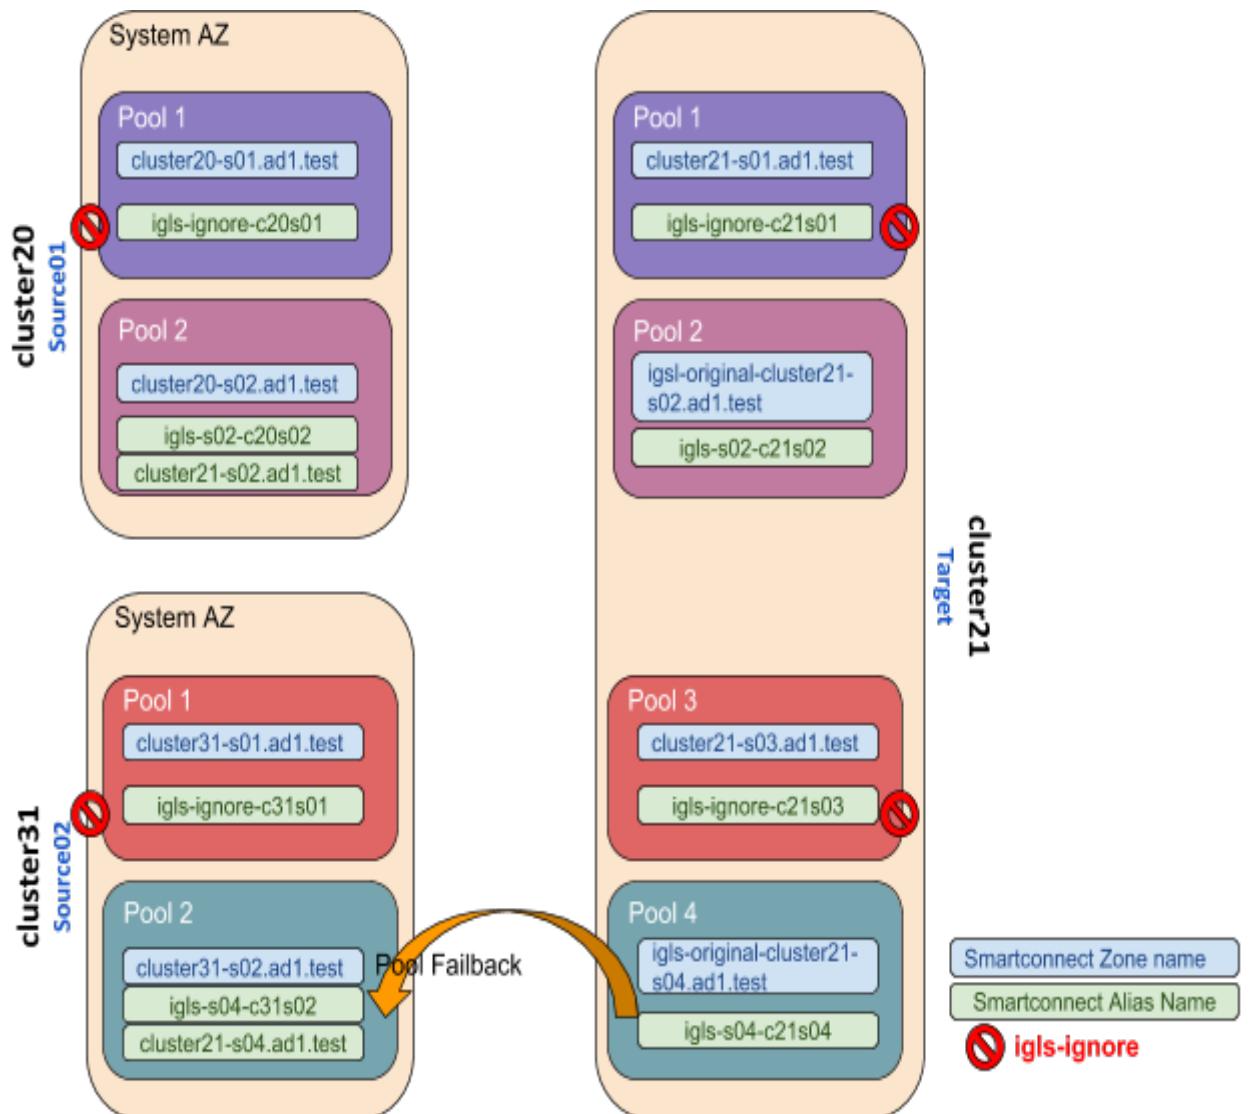


Pool Failback Target (Pool 4) \Rightarrow Source02 (Pool 2)

The following table and diagram illustrate the SmartConnect zone names and alias names after performing Failback from Target (Pool 4) to Source02 (Pool 2).

From			To		
Cluster	Pool	SmartConnect Name / Alias (igls hints)	Cluster	Pool	SmartConnect Name / Alias (igls hints)
Target	Pool 1	cluster21-s01.ad1.test	Source01	Pool 1	cluster20-s01.ad1.test
		igls-ignore-c21s01			igls-ignore-c20s01
Target	Pool 2	igls-original-cluster21-s02.ad1.test	Source01	Pool 2	cluster20-s02.ad1.test

		s02.ad1.test			
		igls-s02-c21s02			
Target	Pool 3	cluster21-s03.ad1.test	Source02	Pool 1	cluster31.ad1.test
		igls-ignore-c21s03			igls-ignore-c31s01
Target	Pool 4	igls-original-cluster21-s04.ad1.test	Source02	Pool 2	cluster31-s02.ad1.test
		igls-s04-c21s04			igls-s04-c31s02 cluster21-s04.ad1.test



© Superna Inc

2.17. Fan-Out IP Pool Failover

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

- Requirements
- Supported Configurations
- Example Diagrams
- Same Access Zone
- Different Access Zones

Fan out IP pool failover allows Cluster A to replicate data to Cluster B and cluster C using IP pool failover mode.

Requirements

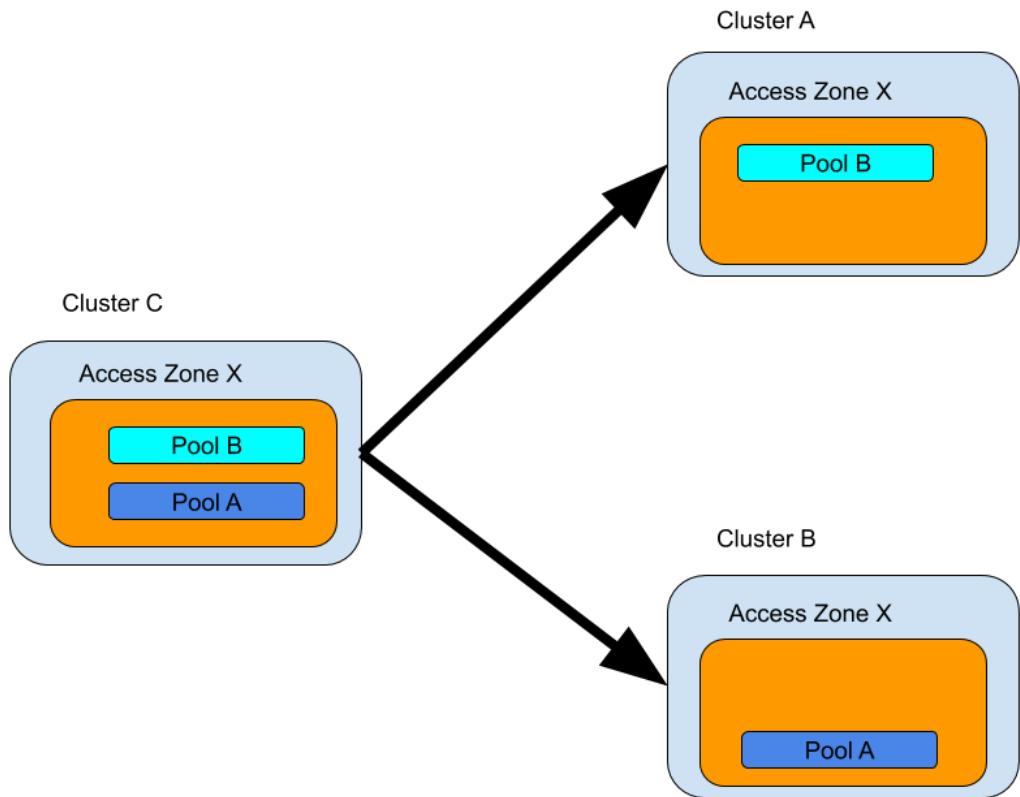
1. Dedicated pools on all clusters, a pool on the central cluster can only replicate to one other cluster, it cannot be shared to replicate to 2 clusters.
2. Cluster A pool A must be mapped to Cluster B Pool A (names of pools are examples only)
3. Cluster A pool B must be mapped to Cluster C Pool B

Supported Configurations

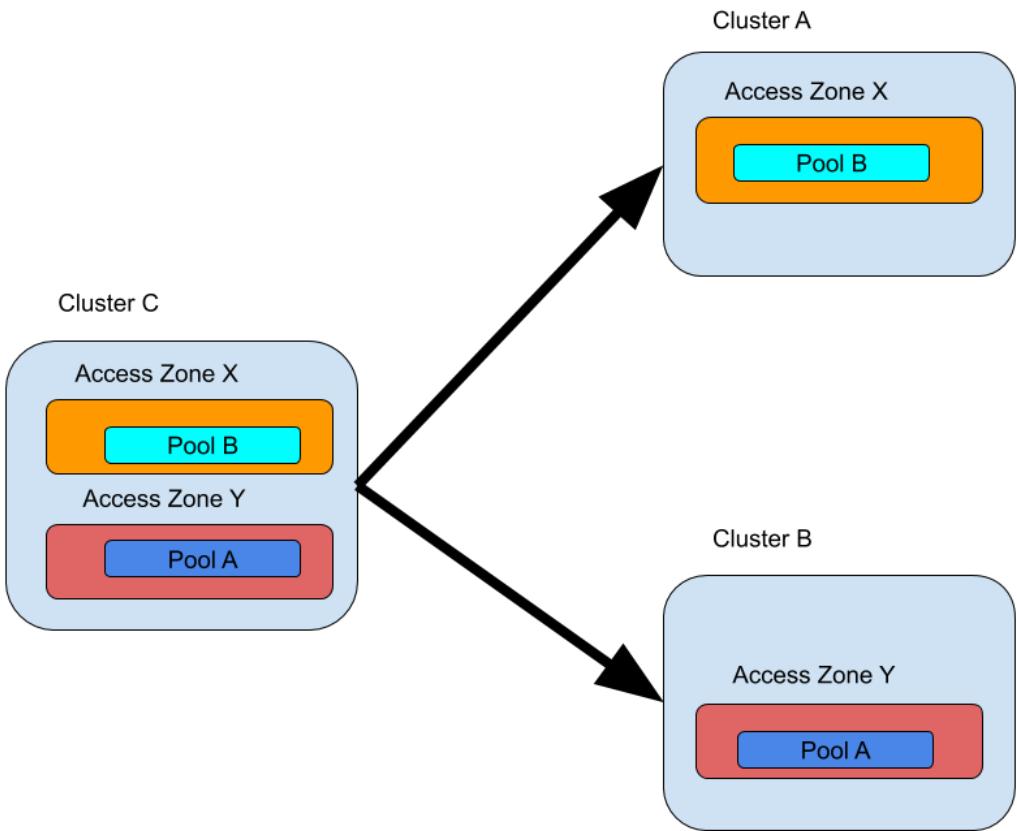
1. A single access zone on the central cluster can have pool A replicated to cluster A and pool B in the same access zone replicated to cluster C
2. More than one access zone on central cluster can replicate an access zone to Cluster A or cluster B but must use dedicated pools for each replication relationship

Example Diagrams

Same Access Zone



Different Access Zones



© Superna Inc

2.18. How to Configure Access Zone DNS Dual Delegation

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

- [Abstract:](#)
- [Unsupported Configurations](#)
- [Delegation](#)
- [FAQ Questions](#)
- [Does support include my DNS vendor?](#)
- [My DNS uses forwarding feature is this the same thing as Dual delegation?](#)
- [Should I delegate with IP address?](#)
- [Can two NS records be added to the delegation to use Primary and Secondary Subnet Service IP ahead of time to simplify failover and remove the DNS manual step post failover?](#)
- [Can dual delegation be used without Eyeglass?](#)
- [Will Dual Delegation support SMB and NFS failovers?](#)
- [How to setup the dual delegation?](#)
- [How does this work with DNS?](#)
- [DNS Return codes](#)
- [DNS Return Message](#)
- [DNS Response Code](#)
- [Function](#)

Abstract:

This technical note covers Dual Delegation, answers key questions, how to set it up and how it works, with DNS.

[Unsupported Configurations](#)

1. DNS Forwarding is untested and not recommended. Name server delegation is the recommended method to configure your DNS servers. This is often used with Infoblox.
2. Do not disable recursive queries on your DNS server. This is untested and not recommended.
3. Do not use CNAME's that point to Smartconnect names and do not create a loop that uses CNAME's within other resource records.
 - a. Dell Isilon documentation states CNAME's should not be used <https://www.delltechnologies.com/resources/us/asset/white-papers/products/storage/h16463-isilon-advanced-networking-fundamentals.pdf>
 - i. 8.10.2 SmartConnect zone aliases as opposed to CNAMEs
 - ii. SmartConnect zone aliases as opposed to CNAMEsA Canonical Name (CNAME) record is a DNS resource mapping one domain to another domain. CNAMEs are not recommended with OneFS, as it is not possible to discover which CNAME points to a given SmartConnect zone name. **During a disaster recovery scenario, CNAMEs complicate and extend the failover process, as many CNAMEs must be updated. Further, Active Directory**

Kerberos does not function with CNAMEs. Zone aliases are the recommended alternative. OneFS provides an option for creating SmartConnect zone aliases. As a best practice, a SmartConnect zone alias should be created in place of CNAMEs. To create a SmartConnect zone alias, use the following command: `isi networks modify pool --add-zone-aliases=` Once the SmartConnect zone alias is provisioned, a matching delegation record must be created in the site DNS, pointing to a SmartConnect Service IP (SSIP).

iii. The DNS RFC states

<https://tools.ietf.org/html/rfc1912>

2.4 CNAME records

A CNAME record is not allowed to coexist with any other data

1. RFC 1912 Common DNS Errors February 1996

Don't use CNAMEs in combination with RRs which point to other names

like MX, CNAME, PTR and NS. (PTR is an exception if you want to

implement classless in-addr delegation.) For example, this is strongly discouraged:

2. podunk.xx. IN MX mailhost

mailhost IN CNAME mary

mary IN A 1.2.3.4

Delegation

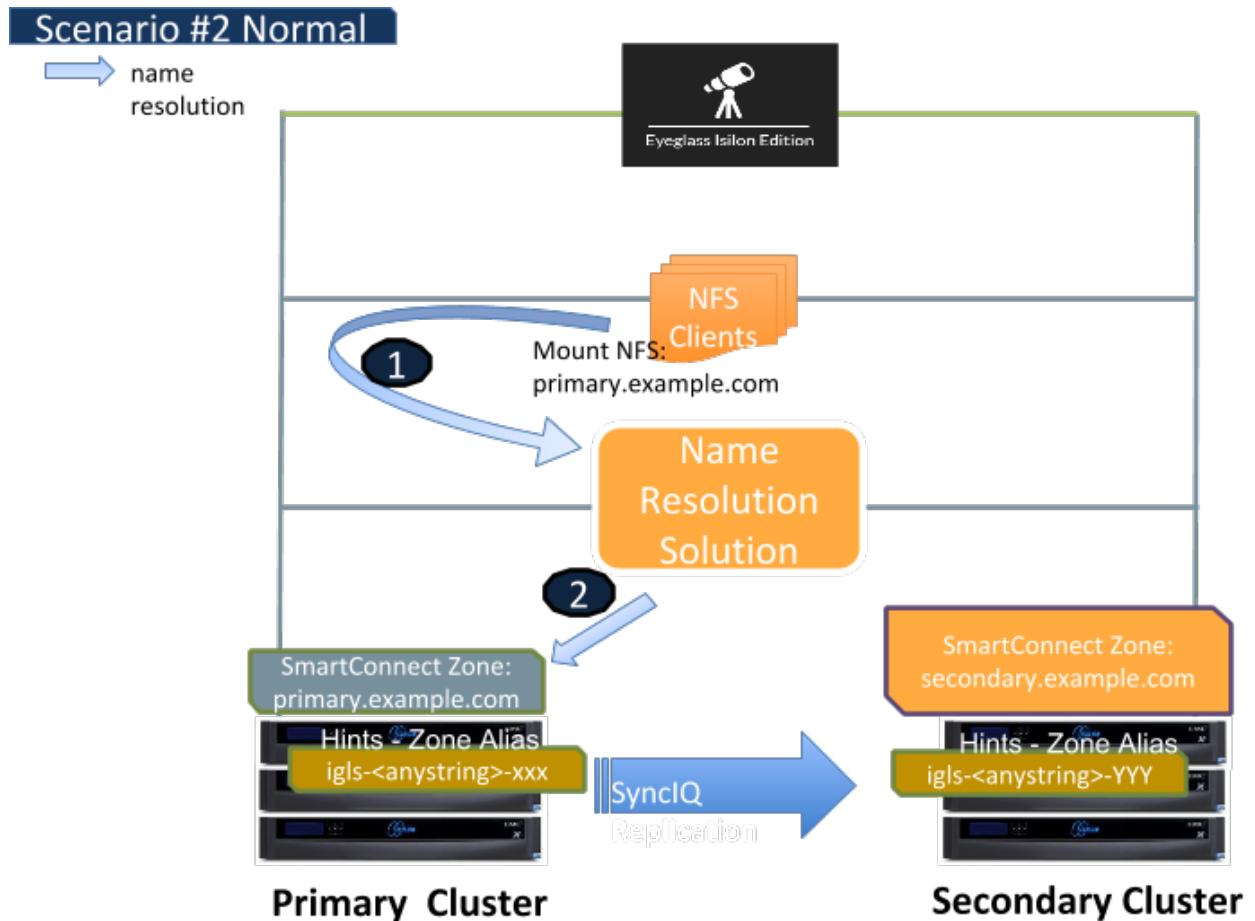
Let's review Access Zone failover with Eyeglass at a high level before explaining how "Dual Delegation" works. It involves SmartConnect zones applied to IP pools failing over. This means creating an alias on the target PowerScale that existed on the source cluster IP pool. Once this is completed, DNS delegation records are based on NS or name server records that point at the Subnet Service IP servicing the IP pool involved in the failover.

The IP Pools are updated to forward SmartConnect zone lookups to the newly active clusters subnet Service ip with the newly created IP alias on the IP pool.

Skip the reading and watch the video 5M on how to setup Dual Delegation with Microsoft DNS (other DNS vendors work as well, the concept is the same).

Continue reading below for details and how it works.

Eyeglass automates this process during failover BUT the source cluster SmartConnect Zone is renamed (Preserved for fallback operations), and leaves a simple breadcrumb zone that remains that is prefixed with `igls-original-whatever-the-zone-name-was`. This rename operation has a second benefit, in that the source cluster Service IP will no longer answer any queries that are sent to the source cluster Subnet Service IP. (see diagram below) DNS resolution **BEFORE** failover where typically on only a single NS record points at Primary cluster.



(see diagram below) DNS resolution **AFTER** manually updating the NS record to point at the Service IP of the secondary cluster.

FAQ Questions

Does support include my DNS vendor?

Support examples are provided as examples using Microsoft DNS, all other DNS vendor solutions should use the vendor documentation for creating Name server records. Support contract does not include support for DNS itself and how to guides for procedures should come from the DNS vendor on how to create name server records.

1. My DNS uses forwarding feature is this the same thing as Dual delegation?
2.
 - a. No DNS forwarding is not the same thing and is not based on DNS standards and is implemented differently by DNS vendors. You can use DNS forwarding but this guide will no longer be of any value and no DNS debug tools like nslookup or DIG will work to validate DNS forwarding configurations. Eyeglass Dual DNS validation will need to be disabled if you choose to use DNS forwarding and support will not be able to assist

to validate your configuration. All DNS vendors can use Name Server records. This is our recommendation for the above reasons.

3. Should I delegate with IP address?

4.

- a. The examples show IP address for simplicity but A records should be used in the name server delegation records to follow DNS best practices.

5. Can two NS records be added to the delegation to use Primary and Secondary Subnet Service IP ahead of time to simplify failover and remove the DNS manual step post failover?

1. Answer: Yes, this is possible and removes yet another step in the failover when using Access Zone Failover

6. Can dual delegation be used without Eyeglass?

1. Answer: no without Eyeglass, this configuration can not be used. This is because the Primary cluster SmartConnect Zone needs to be manually removed or renamed, Eyeglass turns this into an Atomic automated process during failover when the Secondary cluster file system is made writeable (the only time you need this functionality)

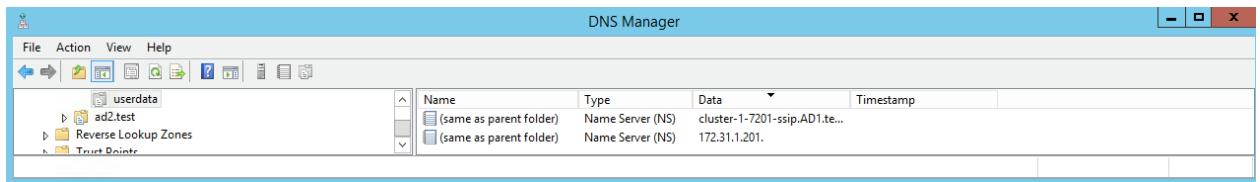
4. Will Dual Delegation support SMB and NFS failovers?

1. Answer: Yes, both protocols benefit from this feature with SMB handling this better since retrying failed mounts requires clicking on the drive letter or accessing the UNC again to look up the name to ip address. NFS clients will require unmount and

remount since they mount an ip address after name resolution has completed, even if used in the /etc/fstab file with an FQDN. Script Engine can now be focused entirely on host side automation without needing any DNS updates

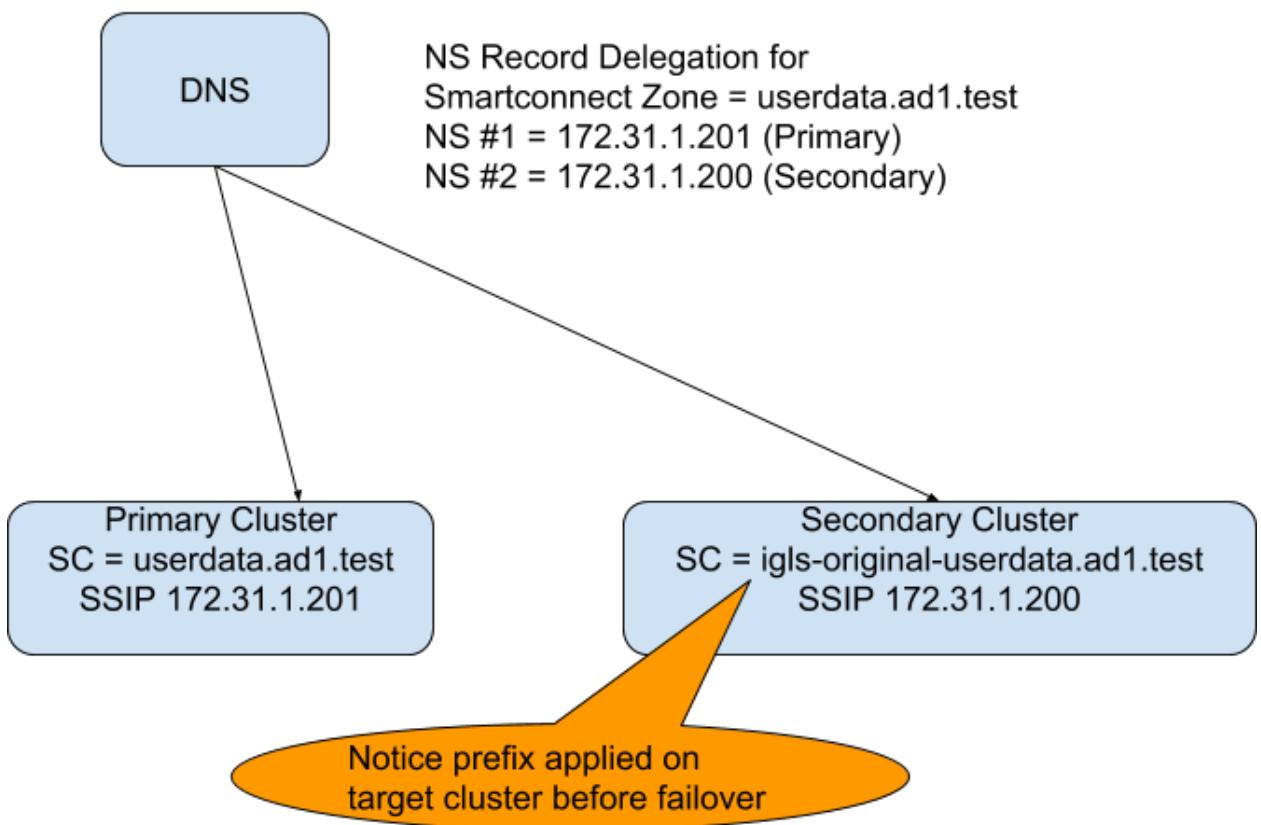
How to setup the dual delegation?

Simple, delegate the SmartConnect Zone with two NS records #1 to primary cluster and #2 the Secondary cluster SSIP that answers DNS queries for the IP pool.

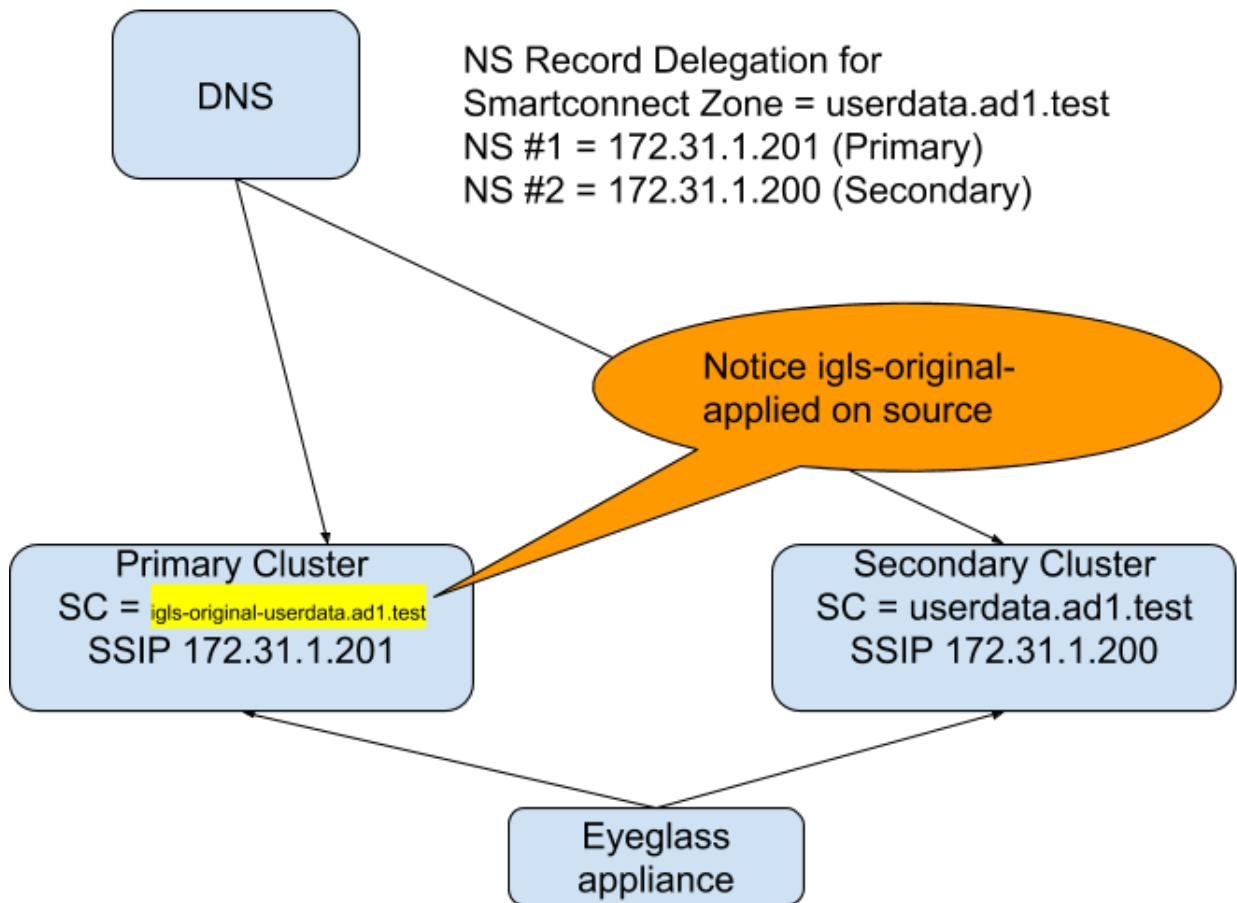


Let's review how this works. The following two diagrams show the DNS setup before and after failover with Eyeglass.

Before Access Zone Failover



After Access Zone Failover



How does this work with DNS?

Answer:

1. The DNS server can issue queries for SmartConnect Zone userdata.ad1.test to either Name server record.
2. If a query is sent to the Secondary Cluster before failover the PowerScale answers the query code 5 or Refused.

NOTE: apply igls-original-<production cluster smartconnect name>, dual delegation requires a SmartConnect name to exist on the target IP pool. We recommend the syntax above. Also

NOTE detection of the igls-original prefix on an Access Zone pool will update the DR dashboard with FAILEDOVER state.

7. This tells the DNS server to re-issue to the second name server record to satisfy the query.
8. Since the Primary cluster is configured for this SmartConnect Zone, it will answer the query from one of IP addresses in the IP pool as expected.
9. The DNS server returns the IP address provided to the client that issued the query.
10. done.

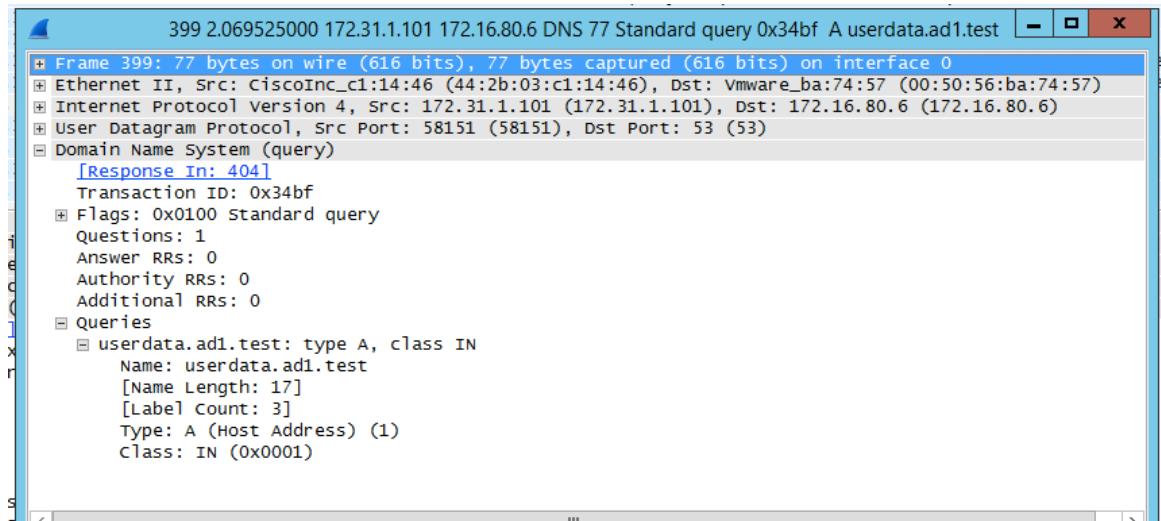
Failover Steps

1. Same as above except Eyeglass has disabled the Primary cluster SmartConnect Zone name with prefix igls-original-xxxx, and the Primary cluster will respond with DNS return code 5 Refused
2. DNS server re-issues the query to the Secondary cluster SSIP to get the query answered.
3. The above all assumes a TTL of 0 on NS records to avoid caching name to ip addresses which works with the exception of Linux mount commands.
4. If a real-DR event occurs versus a controlled failover where both clusters are reachable, the last step would be to ensure the Primary cluster is not ip reachable (if partial disaster), and prevent this cluster from coming up again post DR Event, so that it does not answer DNS queries again or

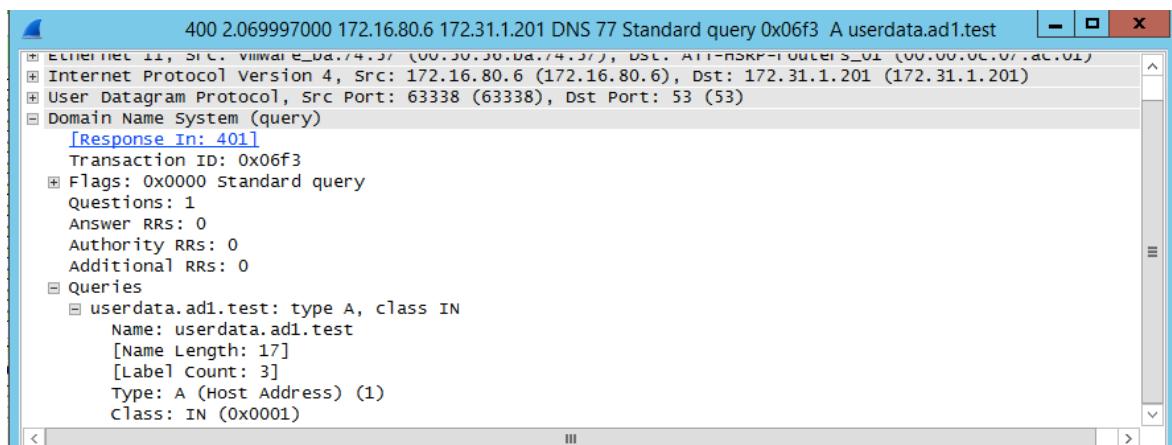
simply remove the entry from DNS.. This is standard practice but mentioned for completeness.

Let's review the wireshark traces below of a failover from a DNS view of a Linux Client.

1. Linux Client 172.31.1.101 issues query to userdata.ad1.test



2. DNS server source sends query to Primary cluster SSIP (no longer the active cluster)



3. Primary Cluster (ip 172.31.1.200) with igl-original renamed Smartconnect Zone answers the query from DNS server (172.16.80.6)

```

+ Frame 401: 77 bytes on wire (616 bits), 77 bytes captured (616 bits) on interface 0
  Ethernet II, Src: CiscoInc_c1:14:46 (44:2b:03:c1:14:46), Dst: VMware_BA:74:57 (00:50:56:ba:74:57)
  Internet Protocol Version 4, Src: 172.31.1.201 (172.31.1.201), Dst: 172.16.80.6 (172.16.80.6)
  User Datagram Protocol, Src Port: 53 (53), Dst Port: 63338 (63338)
  Domain Name System (response)
    [Request In: 400]
    [Time: 0.001966000 seconds]
    Transaction ID: 0x06f3
    Flags: 0x8405 Standard query response, Refused
    Questions: 1
    Answer RRs: 0
    Authority RRs: 0
    Additional RRs: 0
    Queries
      userdata.ad1.test: type A, class IN
        Name: userdata.ad1.test
        [Name Length: 17]
        [Label Count: 3]
        Type: A (Host Address) (1)
        Class: IN (0x0001)

```

5. DNS Server re-issues query to 2nd NS record in this case the Secondary cluster SSIP

```

+ Frame 402: 77 bytes on wire (616 bits), 77 bytes captured (616 bits) on interface 0
  Ethernet II, Src: VMware_BA:74:57 (00:50:56:ba:74:57), Dst: All-HSRP-routers_01 (00:00:0c:07:ac:01)
  Internet Protocol Version 4, Src: 172.16.80.6 (172.16.80.6), Dst: 172.31.1.200 (172.31.1.200)
  User Datagram Protocol, Src Port: 63338 (63338), Dst Port: 53 (53)
  Domain Name System (query)
    [Response In: 403]
    Transaction ID: 0x06f3
    Flags: 0x0000 Standard query
    Questions: 1
    Answer RRs: 0
    Authority RRs: 0
    Additional RRs: 0
    Queries
      userdata.ad1.test: type A, class IN
        Name: userdata.ad1.test
        [Name Length: 17]
        [Label Count: 3]
        Type: A (Host Address) (1)
        Class: IN (0x0001)

```

6. Secondary cluster responds with new IP address from the target failover over IP pool mapped by Eyeglass for failover

```

+ Frame 403: 93 bytes on wire (744 bits), 93 bytes captured (744 bits) on interface 0
  Ethernet II, Src: CiscoInc_c1:14:46 (44:2b:03:c1:14:46), Dst: VMware_BA:74:57 (00:50:56:ba:74:57)
  Internet Protocol Version 4, Src: 172.31.1.200 (172.31.1.200), Dst: 172.16.80.6 (172.16.80.6)
  User Datagram Protocol, Src Port: 53 (53), Dst Port: 63338 (63338)
  Domain Name System (response)
    [Request In: 402]
    [Time: 0.000334000 seconds]
    Transaction ID: 0x06f3
    Flags: 0x8400 Standard query response, No error
    Questions: 1
    Answer RRs: 1
    Authority RRs: 0
    Additional RRs: 0
    Queries
      userdata.ad1.test: type A, class IN
        Name: userdata.ad1.test
        [Name Length: 17]
        [Label Count: 3]
        Type: A (Host Address) (1)
        Class: IN (0x0001)
    Answers
      userdata.ad1.test: type A, class IN, addr 172.31.1.110
        Name: userdata.ad1.test
        Type: A (Host Address) (1)
        Class: IN (0x0001)
        Time to live: 0
        Data length: 4
        Address: 172.31.1.110 (172.31.1.110)

```

7. Client mount can now succeed using new ip address on the Secondary cluster IP pool to mount shares or exports
8. Done.

DNS Return codes

DNS Return Message	DNS Response Code	Function
NOERROR	RCODE:0	DNS Query completed successfully
FORMERR	RCODE:1	DNS Query Format Error
SERVFAIL	RCODE:2	Server failed to complete the DNS request
NXDOMAIN	RCODE:3	Domain name does not exist.
NOTIMP	RCODE:4	Function not implemented
REFUSED	RCODE:5	The server refused to answer for the query
YXDOMAIN	RCODE:6	Name that should not exist, does exist
XRRSET	RCODE:7	RRset that should not exist, does exist
NOTAUTH	RCODE:8	Server not authoritative for the zone
NOTZONE	RCODE:9	Name not in zone

Copyright 2017 Superna LLC

© Superna Inc

2.19. Controlled Failover Option Results Summary

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

Controlled Failover Option Results Summary

The following table indicates which failover steps are executed based on whether or not the **Controlled failover** option was selected when the Access Zone failover was initiated.

Steps	Description	Executed on	Access Zone	Controlled Failover selected	Controlled Failover NOT selected
1 - Ensure that there is no live access to data	Check for open files. If Open files found, decide whether to failover or wait to be closed.	Source	Manual	Not applicable - manual step	Not applicable - manual step
2 - Begin Failover	Initiate Failover from Eyeglass	Eyeglass	Manual	Not applicable - manual step	Not applicable - manual step
3 - Validation	Wait for other Eyeglass Failover jobs to complete	Eyeglass	Automated by Eyeglass	Step Executed	Step Executed
4 - Synchronize data	Run all OneFS SyncIQ policy jobs related to the Access Zone being failed over	Source	Automated by Eyeglass (all policies in the Access Zone)	Step Executed	Step NOT Executed
5 - Synchronize configuration (shares/export/alias)	Run Eyeglass configuration replication 1	Eyeglass	Automated by Eyeglass (based on matching Access Zone base path)	Step Executed	<p style="color: orange;">Step Executed based on last known data in Eyeglass</p> <p>*If you do not want this, uncheck the</p>

					"Config Sync" option to skip this step
6 - Synchronize quota(s)	Run Eyeglass Quota Jobs related to the SyncIQ Policy or Access Zone being failed over	Eyeglass	Automated by Eyeglass (based on matching Access Zone base path)	Step Executed	Step Executed based on last known data in Eyeglass
7 - Record schedule for SyncIQ policies being failed over	Get schedule associated with the SyncIQ policies being failed over on OneFS	Source	Automated by Eyeglass	Step Executed	Step NOT Executed
8 - Prevent SyncIQ policies being failed over from running	Set schedule on the SyncIQ policy(s) to manual on source cluster	Source	Automated by Eyeglass	Step Executed	Step NOT Executed
9 - Provide write access to data on target	Allow writes to SyncIQ policy(s) related to failover2	Target	Automated by Eyeglass (only for policies that match the Access Zone Base path)	Step Executed	Step Executed
10 - Disable SyncIQ on source and make active on target	Resync prep SyncIQ policy related to failover (Creates MirrorPolicy on target) from OneFS	Source	Automated by Eyeglass	Step Executed	Step NOT Executed
11 - Set proper SyncIQ schedule on target	Set schedule on MirrorPolicy(Target) using schedule from step 6 from OneFS for policy(s) related to the Failover	Target	Automated by Eyeglass	Step Executed	Step NOT Executed
12 - Remove quotas on directories that are target of SyncIQ (PowerScale best practice)	Delete all quotas on the source for all the policies	Source	Automated by Eyeglass	Step Executed	Step NOT Executed
13 - Change SmartConnect Zone on Source so not to resolve by Clients	Rename SmartConnect Zones and Aliases (Source)	Source	Automated by Eyeglass (Requires IP pool hints are configured See docs)	Step Executed	Step NOT Executed
14 - Avoid SPN Collision	Sync SPNs in all AD providers to current	Source	Automated by Eyeglass (AD)	Step Executed	Step Executed

	SmartConnect zone names and aliases (Source)		delegation must be completed as per install docs)		
15 - Move SmartConnect zone to Target	Add source SmartConnect zone(s) as Alias(es) on (Target)	Target	Automated by Eyeglass (Requires IP pool hints are configured See docs)	Step Executed	Step Executed
16 - Update SPN to allow for authentication against target	Sync SPNs in all AD providers to current SmartConnect zone names and aliases (Target)	Target	Automated by Eyeglass (Requires IP pool hints are configured See docs)	Step Executed	Step Executed
17 - Repoint DNS to the Target cluster IP address	Update DNS delegations for all SmartConnect Zones that are members of the Access Zone	DNS	Potentially Automated by Eyeglass (See Script Engine Docs)	Not applicable - manual or scripted step	Not applicable - manual or scripted step
18 - Refresh session to pick up DNS change	Remount the SMB share(s)	SMB Client Machines	Manual on clients (NOTE: DNS servers and clients cache DNS entries which will require touching the client or intermediate DNS servers to clear DNS caches (on Windows ipconfig /flushdns)	Not applicable - manual step	Not applicable - manual step

1. Initiates Eyeglass Configuration Replication task for all Eyeglass jobs
2. SyncIQ does NOT modify the ACL (Access control settings on the file system). It locks the file system. `ls -l` will be identically on both source and target

© Superna Inc

2.20. How to Configure Delegation of Cluster Machine Accounts with Active Directory Users and Computers Snapin

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

- [Abstract:](#)
- [Overview](#)
- [How to prepare you cluster for Eyeglass automated DR failover](#)
- [Key Design used by Eyeglass is proxy on failover SPN Management](#)
- [Summary of Permissions](#)
- [Automated Solution with Eyeglass Computer Object Level Method:](#)
- [Automated Solution with Eyeglass Organizational Unit \(OU\) Method:](#)
- [How to Use Active Directory Delegation of Control Wizard to Delegate Service Principal Name Permissions to the cluster](#)
- [Example of SPN in ADSedit Tool](#)
- [How to check cluster SPN permissions are set correctly](#)

Abstract:

In order to automate DR with SyncIQ and Eyeglass with Active Directory and SMB shares, its important to ensure Service Principal Names (SPN) are synchronized with the machine account used by the DR cluster. This technical note provides a methodology to restrict, the AD permissions needed for automated SPN management during failover, audit and remediation in Eyeglass

Overview

In order to automate DR with SyncIQ and Eyeglass with Active Directory and SMB shares, it's important to ensure Service Principal Names (SPN) are synchronized with the machine account used by the DR cluster.

Service Principal Names are used by Kerberos authentication and machine accounts and a new SPN name pair is created each time a new SmartConnect Zone Alias is created.

Note: Superna Eyeglass only manages SPN related to HOST. SPNs related to HDFS or NFS are not updated and will need to be manually repaired post failover.

How to prepare your cluster for Eyeglass automated DR failover

Eyeglass will create SmartConnect Zone names and aliases required on your DR cluster automatically in advance of a DR failover. This is done by mapping Subnet on cluster A to Subnet on Cluster B in the Eyeglass UI and once set all new SmartConnect Zone's created or Alias on the Production cluster will be synced to the DR cluster network pool and subnet. It's important to setup AD and your cluster in advance of failover to eliminate authentication issues due to missing SPN entries on the machine account.

Key Design used by Eyeglass is proxy on failover SPN Management

During failover SPN deletes must occur against the source cluster AD machine, but during a real DR event the source cluster is not reachable to issue SPN commands. Eyeglass solves this by issuing proxy SPN update commands to the DR cluster, but references the source cluster machine account name. This means that the Eyeglass can correct SPN entries on the source cluster even when it's not reachable.

Note: This proxy SPN management solution depends on the Delegation being done as stated below with an OU used for the cluster machine accounts and allowing each cluster to update the others SPN using ISI proxy commands.

Summary of Permissions

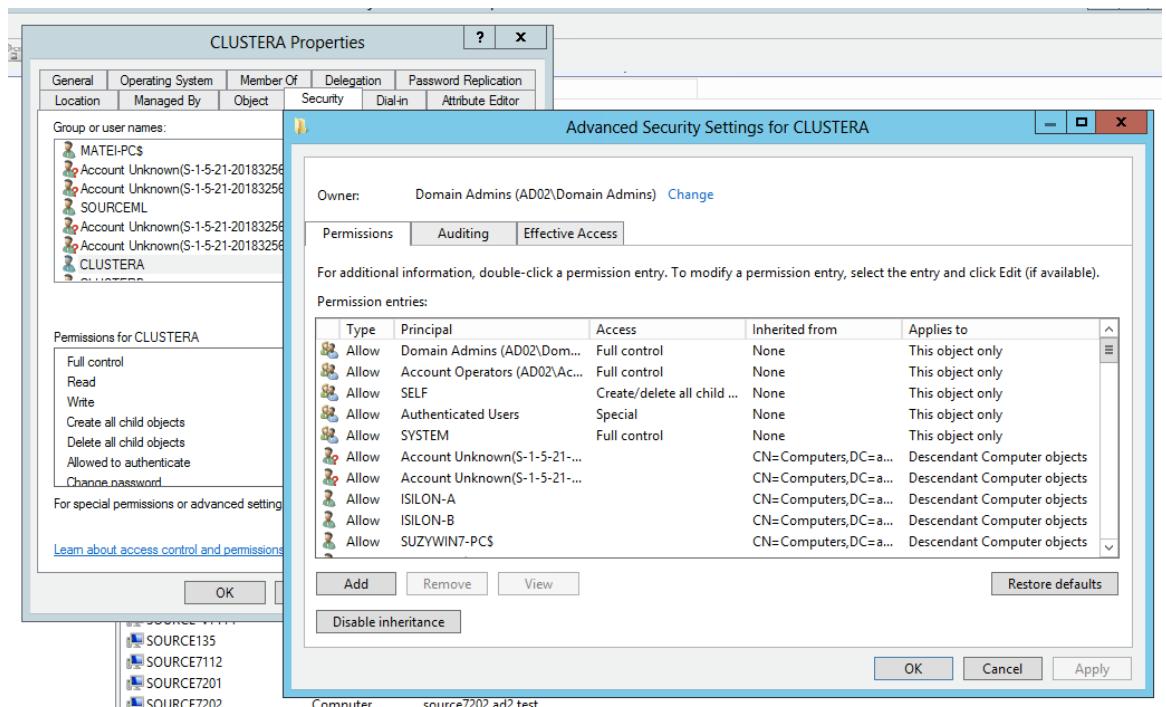
1. Each cluster must have permissions to read and write to the SPN property of its own computer object
2. Each cluster must have permissions to read and write to the SPN property of the opposite cluster computer object

Automated Solution with Eyeglass Computer Object

Level Method:

Use this method to restrict, at the object level, the AD permissions needed for automated SPN management during failover and audit and remediation features in Eyeglass. Recommended with a pair of clusters. Use the Organization Unit (OU) method described in the next section if more clusters objects are involved.

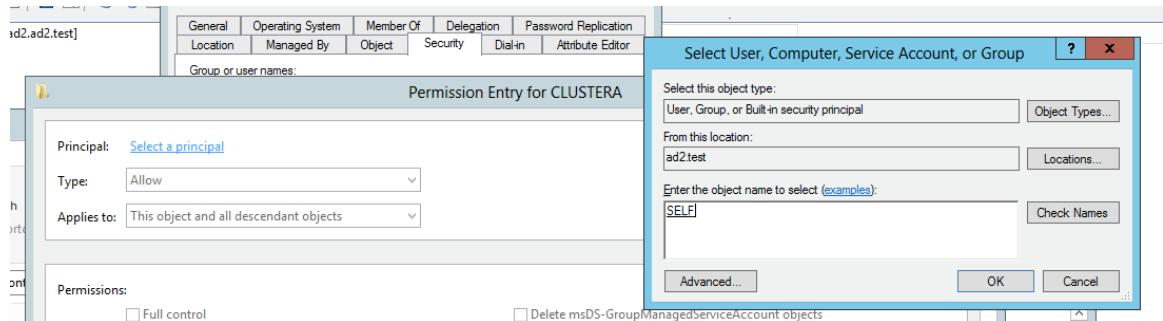
1. Select the first cluster in the pair in Users and Computers snappin as administrator user. Select properties and security tab and then the “Advanced” button of the dialog box. See below:



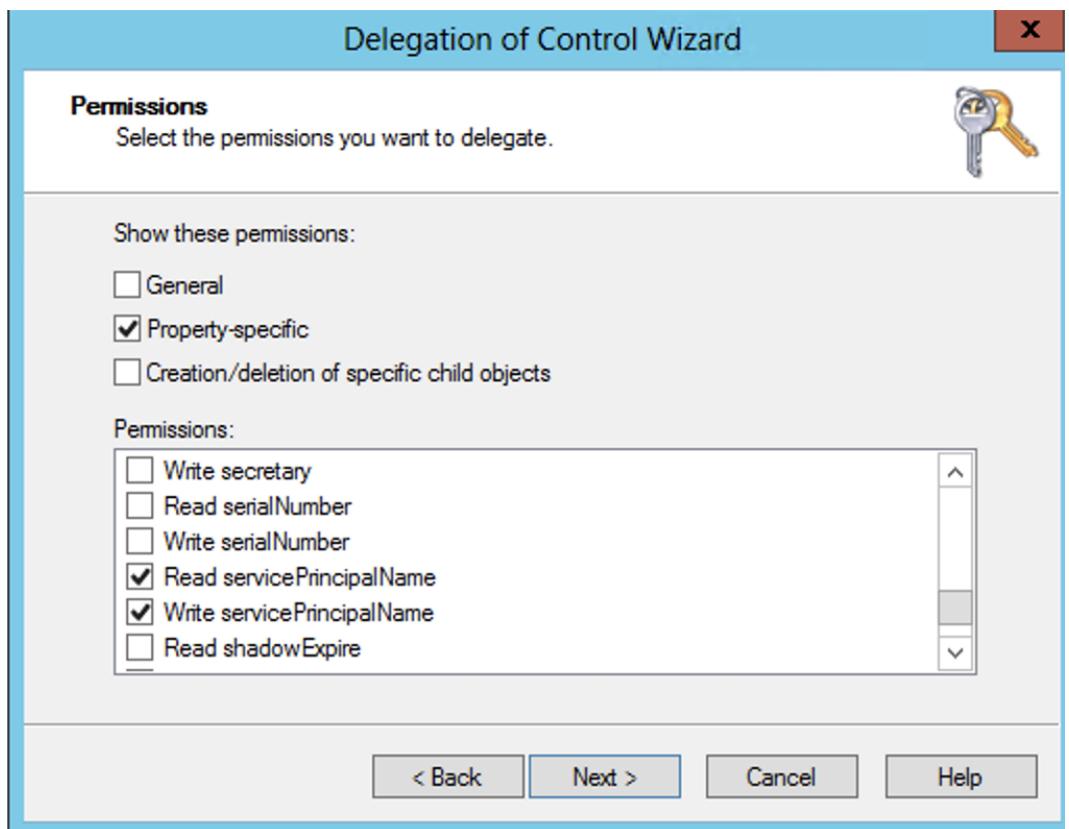
3. Click the add button on the Permissions window (see above)

4. For Principal click select

5. Then type “SELF” and check button and OK (see below)



6. Scroll down the list of permissions and select read and write Service Principal Name (see below)

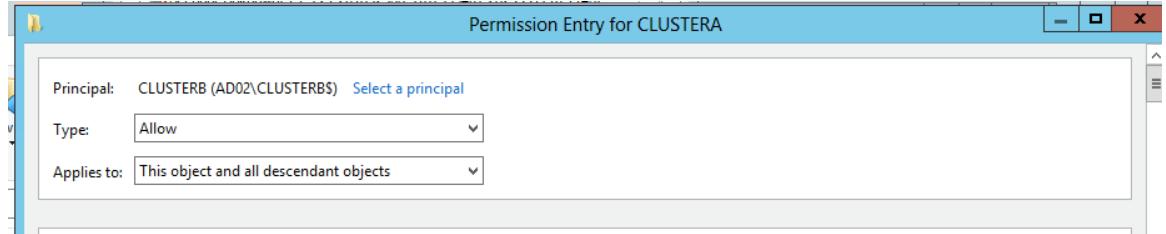


7.

8. Now we need to allow the other cluster machine account to access the SPN properties of the cluster machine account you have selected first . This is for failover of SPN proxy feature in Eyeglass that ensures SPN's can be managed even if a cluster is not reachable. This is called Cross permissions delegation.

9. Click the add button again but this time when selecting the principal to grant permissions, enter the “other” cluster in the replication pair in the dialog box. In the example below you can

see the dialog box title is “CLUSTERA” but the principal selected for the grant is “CLUSTERB”. Find the same read and write properties for service principal name as done above and apply the permissions to the CLUSTERB machine account.

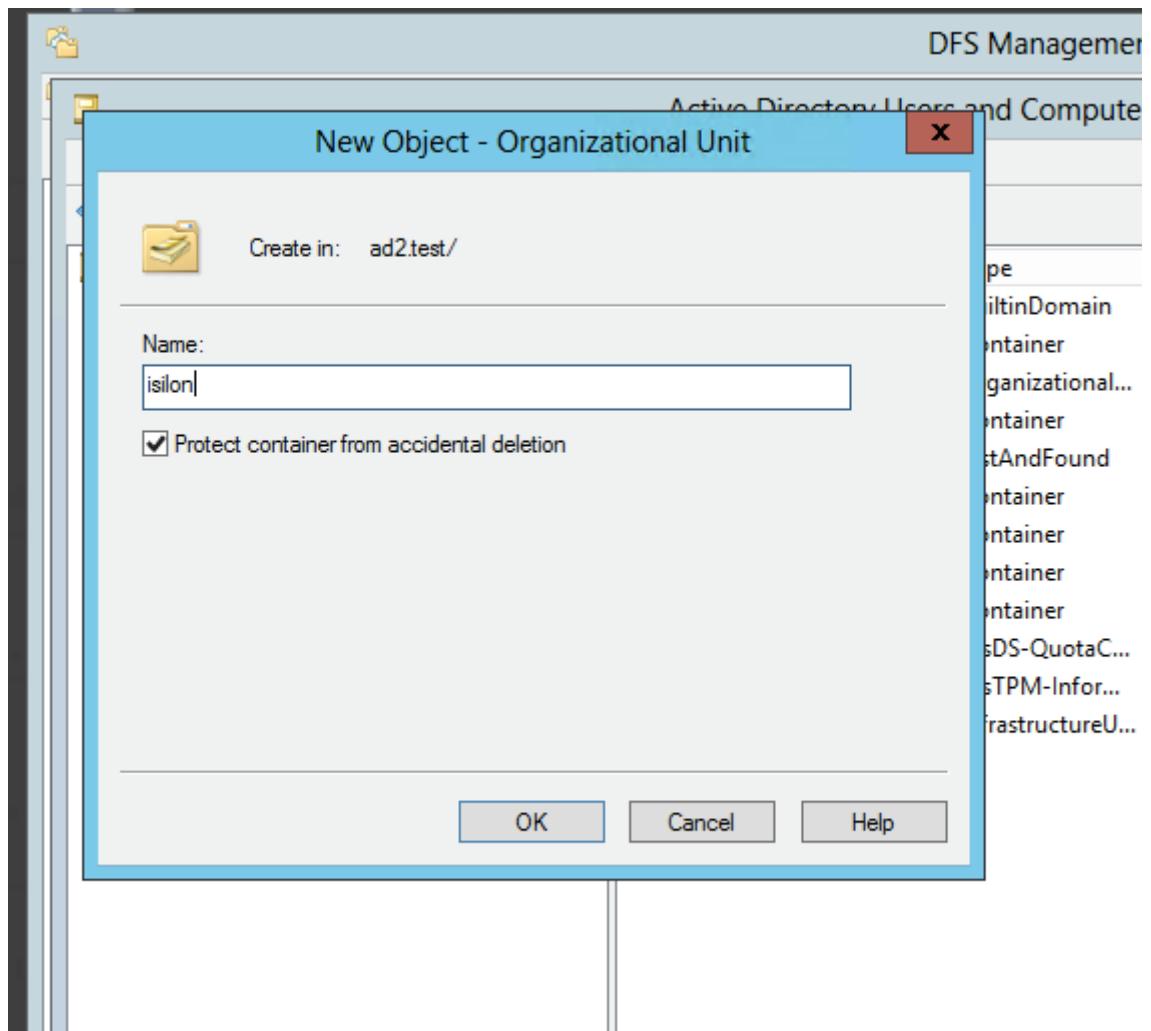


10. The above steps ONLY applied object level permissions to one cluster and both clusters must have these permissions for automated failover and fallback.
11. REPEAT above steps again by select the 2nd cluster of the pair and apply two sets of permissions
12. The Delegation permissions are now completed.

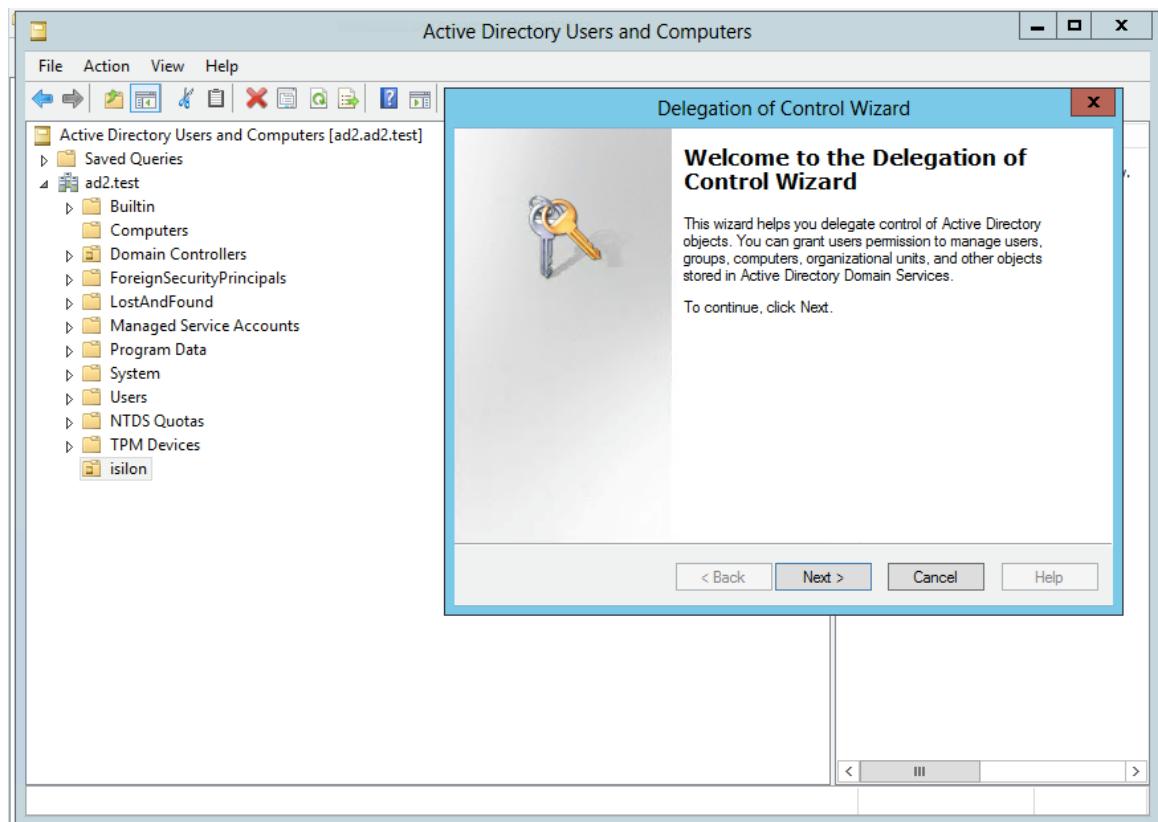
Automated Solution with Eyeglass Organizational Unit (OU) Method:

Use this method when more than one cluster pair is replicating. This method saves steps by doing the delegation once at the OU level versus at the object level. THis method is recommended with greater than 2 clusters to delegate. To avoid Eyeglass requiring the administrator AD account to synchronize the SPN for production or DR clusters, the following one time steps MUST be executed:

1. Using Active Directory Users and Computers Snap In admin tool, create an OU for the PowerScale cluster computer accounts.

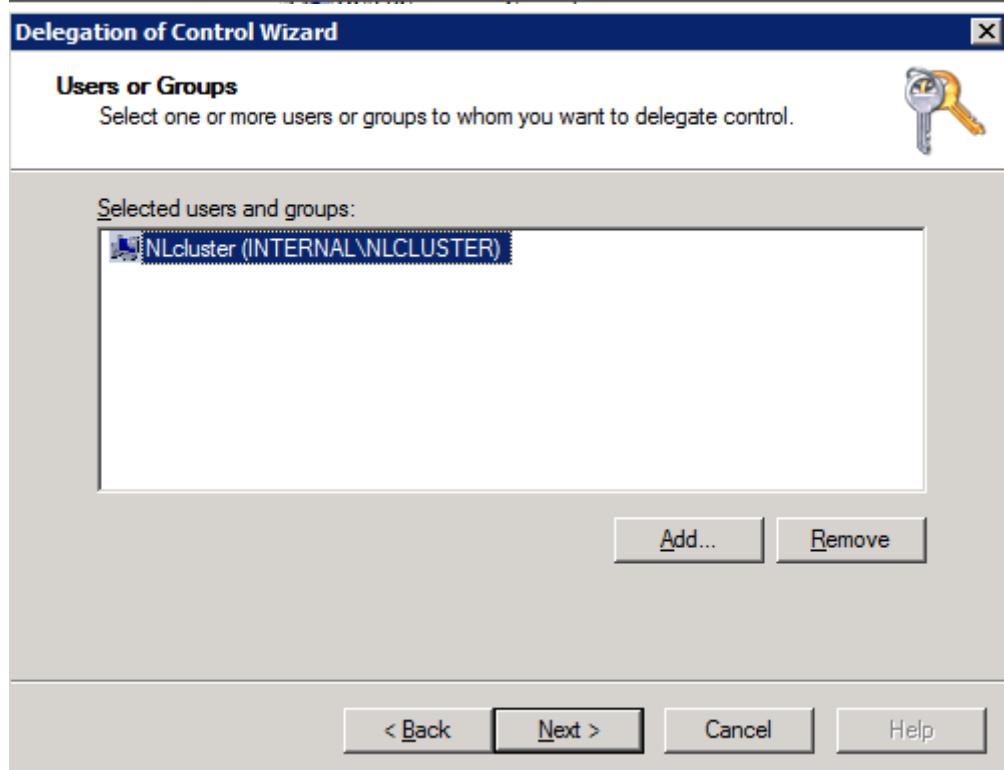


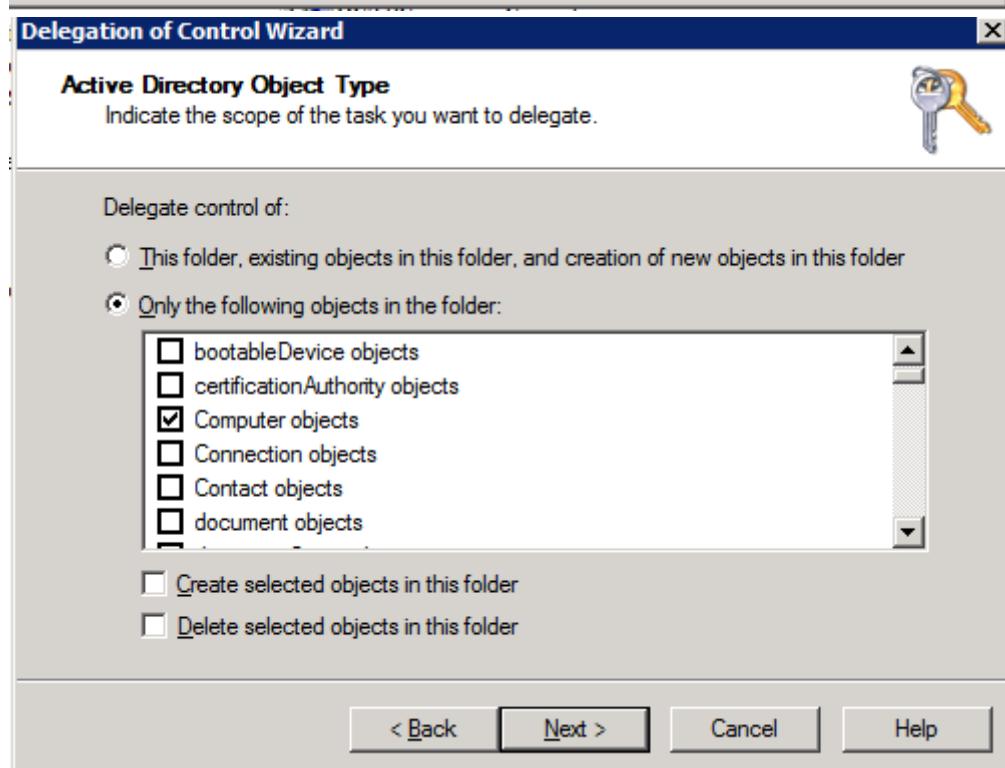
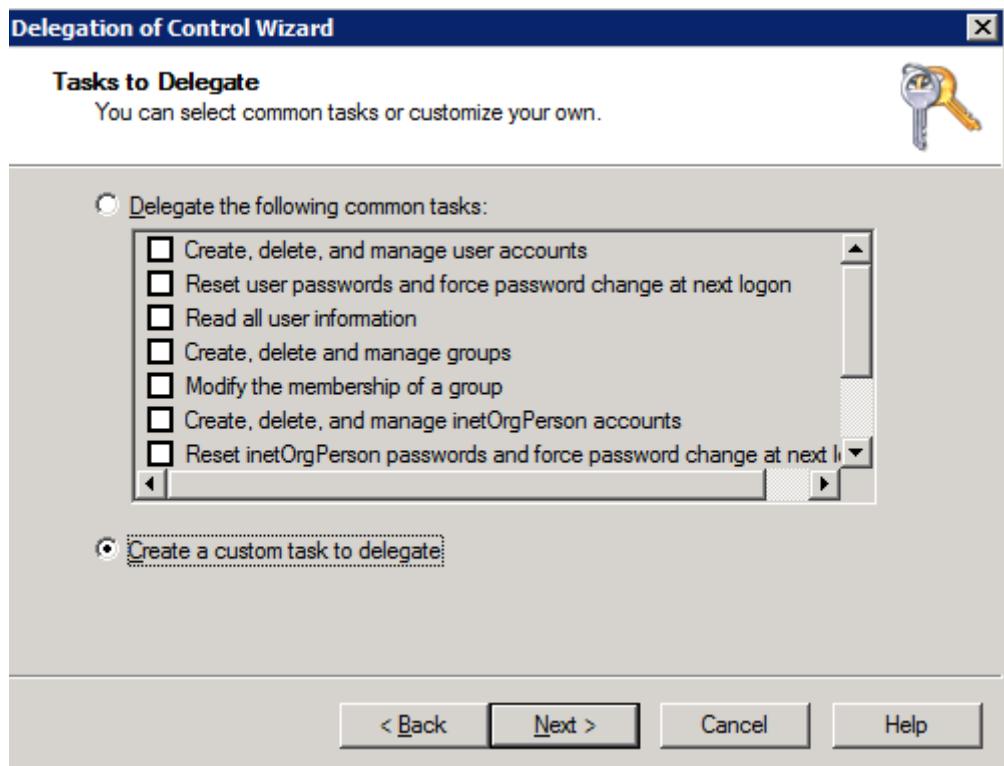
2. Move the cluster AD computer objects with drag and drop into the OU created above.
3. Right click the and select Delegate Control (note this applies to all computers accounts in this folder or OU).
4. Select the Delegation option.

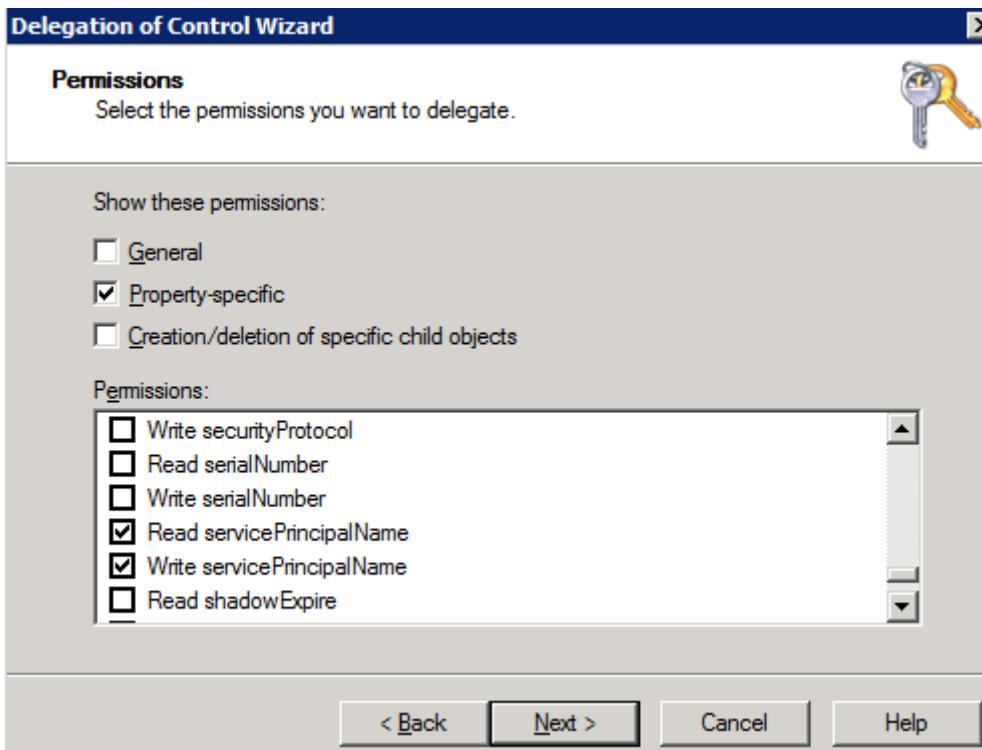


5. Follow screen shots below that assign the cluster permissions to read and write the service principal name.

How to Use Active Directory Delegation of Control Wizard to Delegate Service Principal Name Permissions to the cluster

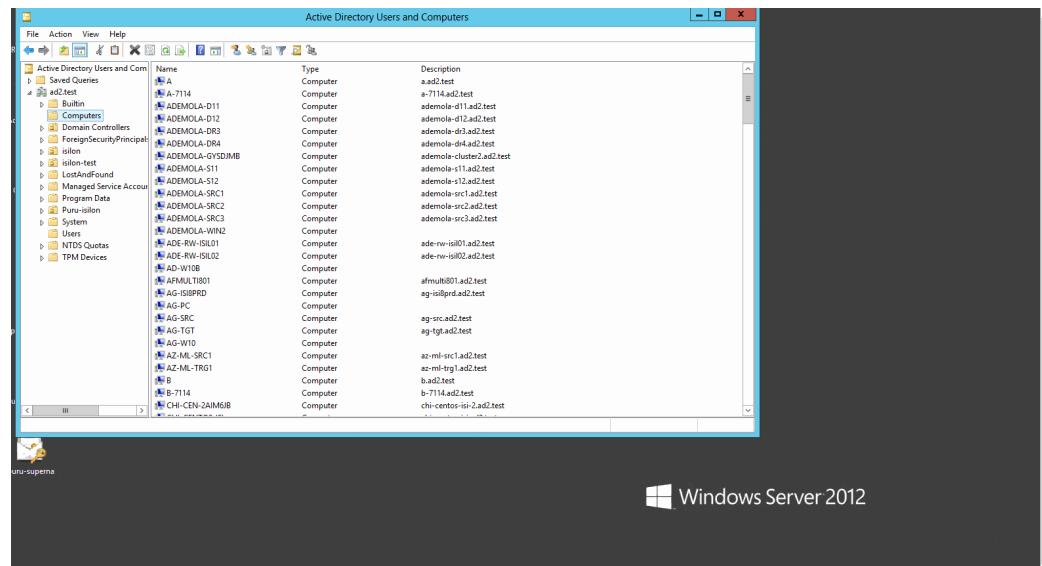






Example of SPN in ADSIedit Tool

How to check cluster SPN permissions are set correctly



6.

7.

1.

© Superna Inc

3. Eyeglass SyncIQ Policy Failover Guide

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

- [Introduction to this Guide](#)
- [Requirements and for Eyeglass Assisted SyncIQ Policy Failover](#)
- [Unsupported Data Replication Topology](#)
- [Recommendations for Eyeglass Assisted SyncIQ Policy Failover](#)
- [Preparing your System for the Eyeglass Assisted SyncIQ Policy Failover](#)
- [Operational Steps for Eyeglass Assisted SyncIQ Policy Failover](#)
- [Monitoring DR Readiness for Eyeglass Assisted Failover](#)
- [Post SyncIQ Policy Failover Manual Steps](#)
- [Controlled Failover Option Results Summary](#)

© Superna Inc

3.1. Introduction to this Guide

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

Introduction to this Guide

The purpose of this guide is to assist you with Eyeglass Assisted SyncIQ Policy Failover.

Overview

Eyeglass Assisted SyncIQ Policy Failover provides users with the tools and information to determine the DR readiness of their configuration. The Eyeglass DR Dashboard Policy Readiness tab provides a per SyncIQ Policy summary of all SyncIQ, OneFS and Eyeglass Configuration replication Job statuses. The status for each are combined to provide an overall DR Status.

This information provides the best indicator of DR readiness for failover and allows administrators to check status on each component of failover, identify status, errors and correct them to get each SyncIQ Policy configured and ready for failover.

The following document is a guide to help you with the requirements, considerations, system preparation, operational steps, and monitoring for Eyeglass assisted SyncIQ Policy failover.

© Superna Inc

3.2. Requirements and for Eyeglass Assisted SyncIQ Policy Failover

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

Requirements and for Eyeglass Assisted SyncIQ Policy Failover

Following conditions are required for Eyeglass Assisted SyncIQ Policy Failover.

Cluster Version Requirements

Clusters participating in SyncIQ Policy Failover must be running the supported PowerScale Cluster version for this feature. See the Feature Release Compatibility matrix in the “Eyeglass Release Notes” specific to your Eyeglass version found [here](#).

SyncIQ Policy Requirements

For a SyncIQ Policy failover with Eyeglass, it is required that the Eyeglass Configuration Replication Job for that SyncIQ Policy is in the Enabled state.

Note: If the SyncIQ Policy is disabled in OneFS or the corresponding Eyeglass Configuration Replication Job is disabled in Eyeglass the SyncIQ Policy failover will be blocked.

Failover Target Cluster Requirements

For a SyncIQ Policy failover with Eyeglass, it is required that the PowerScale Cluster that is the target of the failover be IP reachable by Eyeglass with the required ports open.

Eyeglass Quota Job Requirements

For a SyncIQ Policy failover with Eyeglass, there are no Eyeglass Quota Job state requirements. Quotas will be failed over whether Eyeglass Quota Job is in Enabled or Disabled state.

Service Principal Name Update Requirements

DFS mode and Access Zone failover manage SPN updates automatically. This is only required when SyncIQ policy failover includes direct mount shares.

Eyeglass will allow SyncIQ policy failover. However, Eyeglass requires the storage administrator to selectively fix SPN values on the source and destination cluster. This will ensure that data sets that remain on the source cluster still process SPN authentications via Kerberos using the source cluster machine account. This is the reason the user will be required to manually make the SPN deletes and additions on the target cluster, as the PowerScale cluster does not know which SPN and SmartConnect Zones are related to the SMB Shares.

SPN deletes from the source, followed by SPN add on the target. This must be done manually by reviewing the Smartconnect Zones required for the failover.

Consult the EMC documentation of ISI commands required to add or delete SPN's on computer accounts.

© Superna Inc

3.3. Unsupported Data Replication Topology

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

Unsupported Data Replication Topology

Replication topology with shares or NFS alias with the same name on both clusters and protected by different SyncIQ policies is not supported. Configuration Replication will overwrite the path on one cluster, as the share / alias can not be distinguished. The failover will not succeed, as it would attempt to have 2 SyncIQ policies on the same cluster with the same source path. **Note:**

This is an invalid DR configuration. This configuration means duplicate shares point to different data. Failover will be unsuccessful with or without Eyeglass.

© Superna Inc

3.4. Recommendations for Eyeglass Assisted SyncIQ Policy Failover

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

Recommendations for Eyeglass Assisted SyncIQ Policy Failover

These are important recommendations to ensure that all automated SyncIQ Policy failover steps can be completed. In some cases if the condition is not met it will result in an Error.

Notes:

- Errors **will not** block Eyeglass Assisted SyncIQ Failover from starting.
- If the recommendations are not followed it may result in an error during failover causing the failover to not complete
- Potential data loss will be incurred depending on the step that failed
- Post failover may require additional manual steps to complete the failover with application hosts
- Post failover scripting should be used to automate custom failover requirements per policy

Shares / Exports / NFS Alias Recommendations

- Eyeglass Configuration Replication Jobs for the SyncIQ Policies being failed over should have been completed (so not in Pending state) without error before failover is started

- **Impact:**
- Failover will not complete if configuration sync fails during failover
- Any missing or unsynced share / export / NFS alias information will block client access to data on the Target cluster. These configuration items will have to be corrected manually on the Target Cluster.

SynclIQ Policy Recommendations

- SynclIQ Job in OneFS should have been completed without error and shows green.
- **Impact:**
- Failover will not complete if errors occur during failover
- Data loss due to unreplicated data
- PowerScale does not support SynclIQ Policy with excludes (or includes) for failover.
- **Impact:** Not a supported configuration for failback
- PowerScale best practices recommend that SynclIQ Policies utilize the Restrict Source Nodes option which requires an IP to be created with target smartconnect zone.
- **Impact:** Subnet pool used for data replication is not controlled all nodes in the cluster can replicate data from all IP pools. This is hard to manage bandwidth and requires all nodes to have access to the WAN.

Smartconnect Recommendations

It is recommended to have a dedicated SmartConnect Zone or SmartConnect Zone Alias for data access per SyncIQ Policy instead of sharing SmartConnect Zone/Alias between policies. This will allow clients to continue to use same SmartConnect Zone for data access post failover and simplify the associated networking updates.

© Superna Inc

3.5. Preparing your System for the Eyeglass Assisted SyncIQ Policy Failover

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

Preparing your System for the Eyeglass Assisted SyncIQ Policy Failover

The following steps described in this section are required to prepare your system for the Eyeglass Assisted SyncIQ Policy Failover:

Update PowerScale sudoer file for Eyeglass Service Account with PowerScale Cluster version 7.1.1.0 or 7.1.1.1

Eyeglass SyncIQ Policy Failover requires some CLI commands that must run with root level access. Many customers also run the cluster in STIG or compliance mode for Smartlock WORM features, and the root user account is not allowed to login and run commands. Updating the PowerScale sudoer file allows the Eyeglass service account user to run the command without having root access.

Please refer to the “Update PowerScale sudoer file for Eyeglass service User (Root Level Commands Needed for Failover)” section in the Eyeglass PowerScale Edition Tech Notes [here](#) for details.

Post Failover Automation

Many failover scenarios depend on extra steps performed on devices, software, infrastructure external to the NAS cluster. Using the Eyeglass script engine, these tasks can now be automated with output captured and appended to Eyeglass failover logs. For example:

- DNS updates post failover for SmartConnect zone CNAME editing
- NFS host mount and remount automation
- DNS server cache flushing
- Application bring up and down logic to start applications post failover
- Send alerts or emails
- Run API commands on 3rd party equipment example load balancer, switch, router or firewall

Please refer to the Eyeglass Admin Guide [Script Engine Overview](#) for more details.

Failover Planning and Checklist

Failover planning includes extended preparation beyond storage layer failover steps as related to the clients, application owners and any dependent systems such as DNS and Active Directory. A full Failover Plan is required taking this all into account. A [Failover](#)

[Planning and Checklist](#) document is provided to help you develop your own Failover plan.

© Superna Inc

3.6. Operational Steps for Eyeglass Assisted SyncIQ Policy Failover

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

Operational Steps for Eyeglass Assisted SyncIQ Policy Failover

Eyeglass DR Assisted Failover SyncIQ policy Workflow Overview:

This workflow has dependencies that require manual steps to avoid data access loss. These conditions are addressed automatically with Access Zone Failover.

1. User selection of PowerScale Cluster which is being failed over “from” (Source).
2. User selection of one or more SyncIQ policies for failover.

Note: Quotas failover automatically when they match the SyncIQ policy paths

3. Eyeglass will update the target cluster SyncIQ policy to make it writeable and the final sync status is shown to the user when completed.

4. **SmartConnect Manual steps:** The user must ensure that SmartConnect Zone Aliases are created manually on the target cluster. (**Note:** Make sure the SmartConnect Zone was only in use by the SyncIQ policies selected for failover)

1. This is because an alias created on the target cluster is required for SmartConnect to answer DNS queries from clients expecting to mount the

production cluster SmartConnect Zone or alias name. In addition, SPN collisions in AD can result in authentication failures if the SPN step is skipped. (see Admin Guide section on [Active Directory Machine Account Service Principal Name \(SPN\) Delegation](#)).

2. This can be done from the OneFS CLI and is executed against the SmartConnect Zone on the source and target cluster.
5. **SMB only - SPN Step:** Since SmartConnect Zones and associated SPN entries **cannot be deleted on the source cluster** with Single policy failover workflow (because this may remove access to any data and mounts that are not failed over), no hints are required and it's up to the administrator to ensure SPN's that must be deleted at the source cluster are removed from the OneFS CLI and the target cluster SPN's are created using the OneFS CLI.
6. **Final Step:** DNS updates for SmartConnect Zones can done. The SmartConnect Service IP is used for delegation from DNS to PowerScale. This step is manual but can be scripted with post failover scripting (requires scriptable DNS server)
 1. This step is manually implemented in your company's DNS system that delegated the SmartConnect Zones to the PowerScale cluster.
 7. **Verify** clients can connect to existing SmartConnect Zone shares and exports using NSLOOKUP to verify DNS and then test mounting from a client machine.
 8. Failover Complete.

IMPORTANT: Making any changes to the SyncIQ Policies or related Eyeglass Configuration Replication Jobs being failed over during the failover may result in unexpected results.

IMPORTANT: Eyeglass Assisted Failover has a 1 hour timeout on each failover step. Any step which is not completed within this timeout period will cause the failover to fail.

IMPORTANT: Deleting configuration data (shares, exports, quotas) or modifying Share name or NFS Alias name or NFS Export path on the target cluster before failing over without running Eyeglass Configuration Replication will incorrectly result in the object being deleted on the source cluster after failover. You must run Eyeglass configuration replication before the failover OR select the Config Sync checkbox on failover to prevent this from happening.

IMPORTANT: More than one SyncIQ Policy can be failed over in a single Failover Job but an error in any one of the SyncIQ Policy failovers will cause failover of all SyncIQ Policies to be halted.

For detailed steps consult the failover guide table [here](#).

For detailed steps on execution and monitoring consult the [Failover Design Guide](#)

SyncIQ Policy Failover Procedure with DR Assistant

This is for application failover when host side automation scripts are required. This is recommended with NFS failover when no SPN or SmartConnect Zone automation is required.

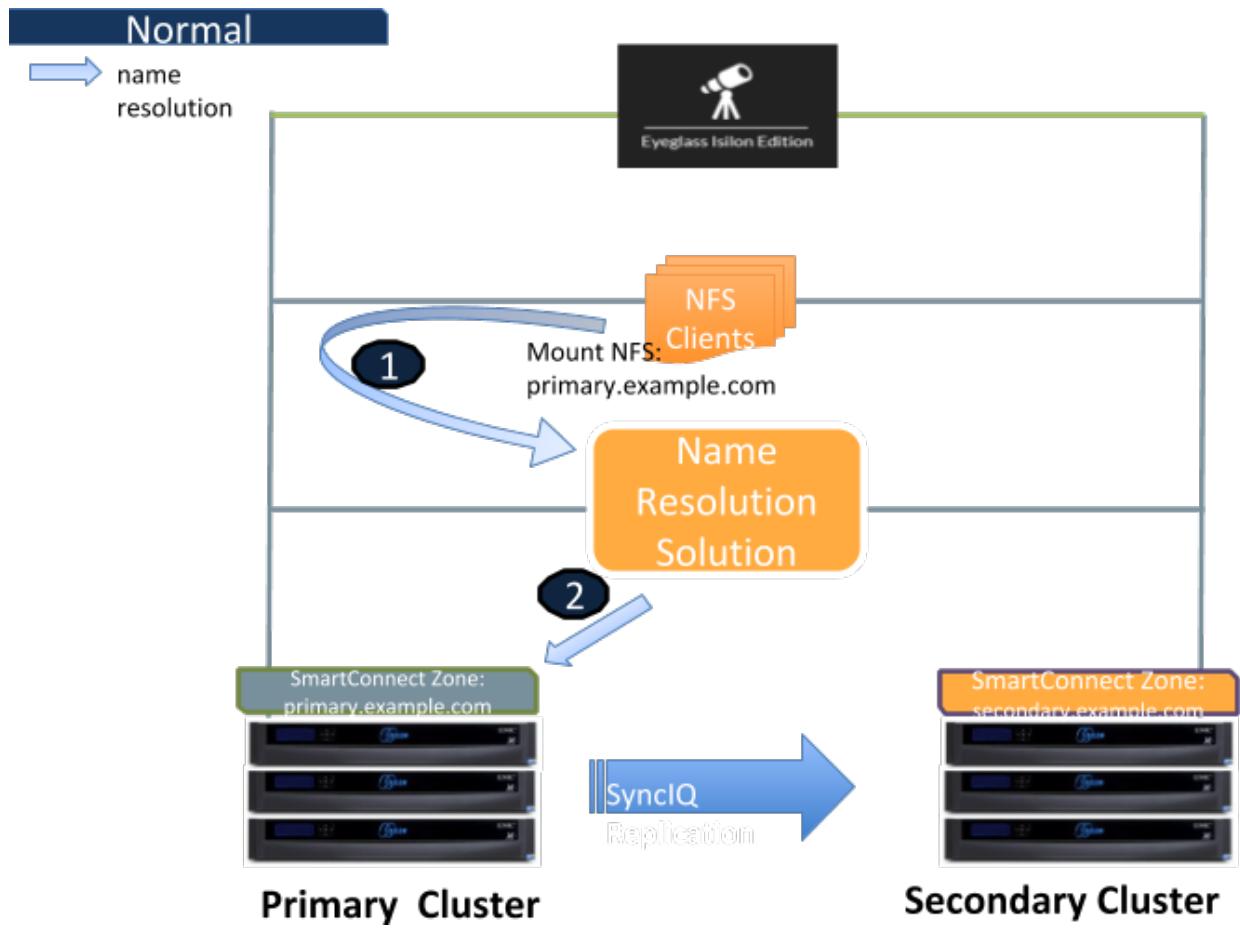
It offers flexibility to design a failover solution customized for an application use case, while using DFS mode or Access Zone failover for standardized failover workflows.

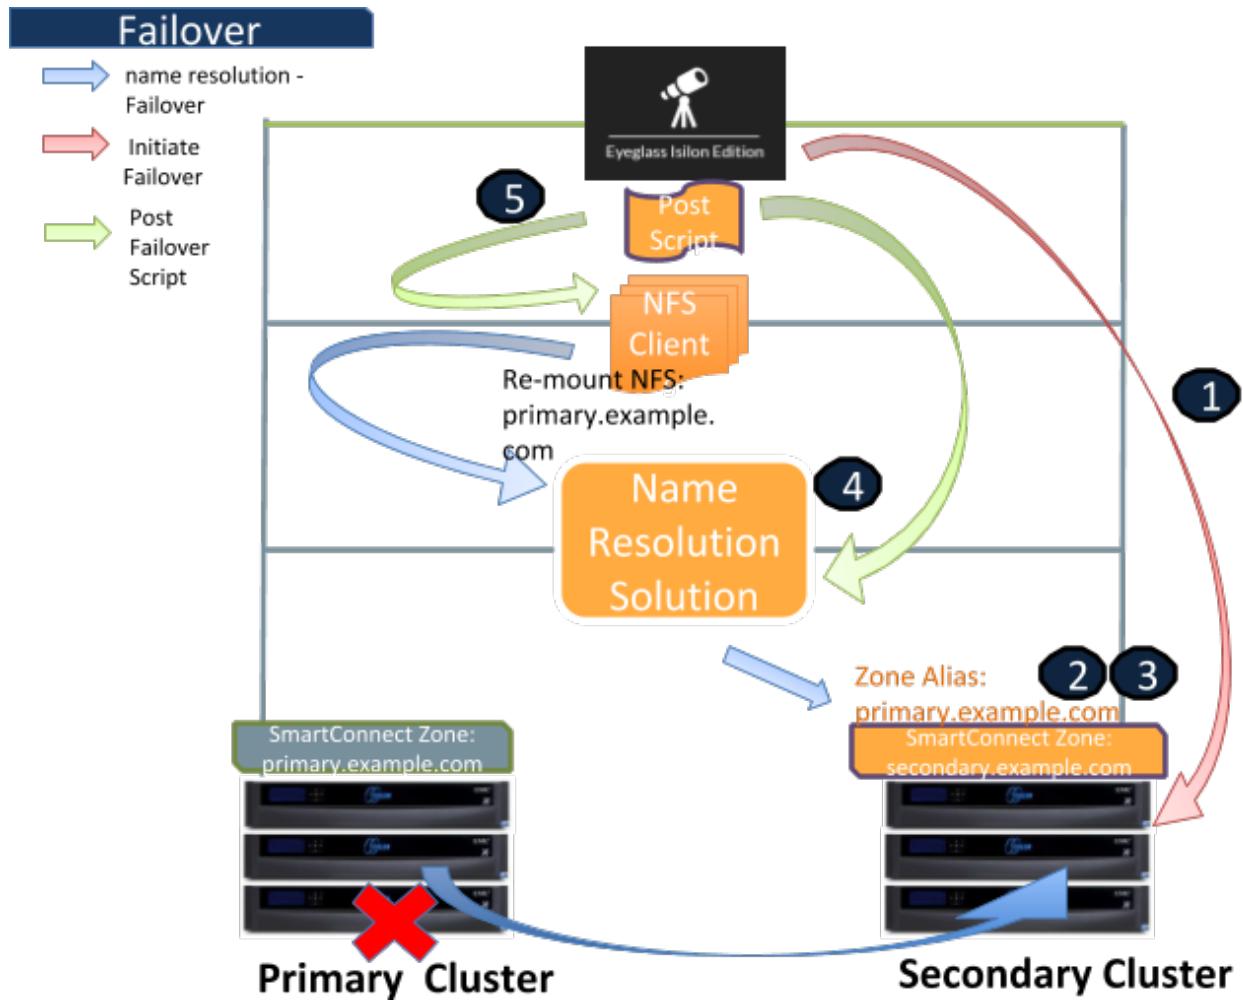
This failover mode can be used with SMB however, manual steps are required for SPN and SmartConnect Zones that are not required with DFS mode or Access Zone modes.

Recommended Use Cases:

1. NFS failover with post script host automation
2. Application failover with script engine to assist with application automation steps
3. Failover without DNS updates using SmartConnect zone names and unmount and remount automation

The diagrams below show the flow and steps of failover.





Failover Steps

Step 1: Initiate Failover SyncIQ Policy with Eyeglass

Step 2: Secondary Cluster - Create SmartConnect zone alias name - Manual (automated with Access Zone Failover)

Command:

```
"isi networks modify pool <subnet>:<zone-name> --add-zone-aliases <zone-alias-name>"
```

Example:

```
"isi network modify pool subnet0:zone01-pool --add-zone-aliases  
primary.ad1.test"
```

Step 3: Secondary Cluster - Ignore missing SPN - Manual

Command:

```
"isi auth ads spn check --domain <domain-name>"
```

When we run this command on the secondary cluster, we can see missing SPN. As this scenario is using local provider as the authentication provider, we can ignore this missing SPN message.

Step 4: Update Name Resolution – IP Address Automated by Eyeglass (with script engine)

We can use the Post Failover Eyeglass script to update the name resolution.

Step 5: Re-mount the NFS Mount Automated by Eyeglass (with script engine)

Unmount and remount the NFS mounting folder on the client to refresh the connection to the NFS export (accessed through the secondary cluster).

Command:

```
"umount -fl /mnt/z01-nfs01"
```

and then

For NFSv3:

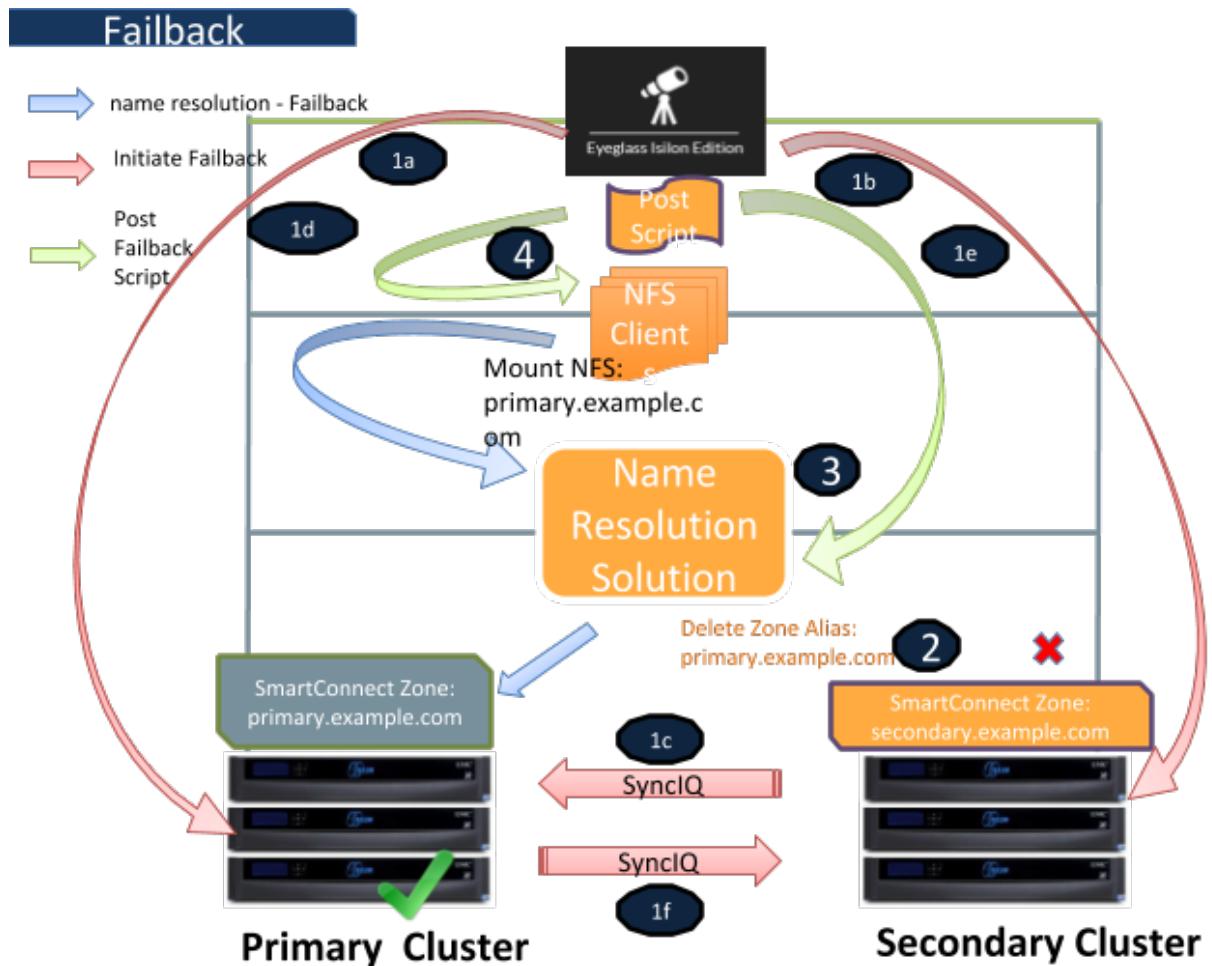
```
"mount -t nfs -o vers=3 primary.ad1.test:/ifs/data/zone01/z01-nfs01 /mnt/z01-nfs01"
```

For NFSv4:

```
"mount -t nfs -o vers=4 primary.ad1.test:/ifs/data/zone01/z01-nfs01 /mnt/z01-nfs01"
```

Might need to unmount the export with force / force and lazy options to solve the NFS stale issue

Also for NFSv4, we might need to kill the process that was accessing the export mount before the failover.



Fallback

Step 1: Initiate Fallback with SyncIQ policy with Eyeglass

1.a. Automated by Eyeglass

1.b. & 1c. Automated by Eyeglass

On Secondary - replicate data to the primary cluster with the mirror policies

1.d. Automated by Eyeglass

On the primary cluster, allow writes to the target directories of the mirror policies

1.e. & 1f. Automated by Eyeglass

On the secondary cluster, complete the failback process

Step 2: Secondary Cluster - Delete SmartConnect zone alias name (automated with Access Zone Failover)

Command:

```
"isi networks modify pool <subnet>:<zone-name> --remove-zone-aliases <zone-alias-name>"
```

Example:

```
"isi network modify pool subnet0:zone01-pool --remove-zone-aliases primary.ad1.test"
```

Step 3: Update Name Resolution - IP Address Automated by Eyeglass (with script engine)

We can use the Post Failover Eyeglass script to update the name resolution.

Step 4: Re-mount the NFS Mount Automated by Eyeglass (with script engine)

We can use the Post Failover Eyeglass script to update the name resolution.

Unmount and remount the NFS mounting folder on the client to refresh the connection back to the primary cluster's NFS export.

Command:

" umount /mnt/z01-nfs01"

and then:

For NFSv3:

"mount -t nfs -o vers=3 primary.ad1.test:/ifs/data/zone01/z01-nfs01 /mnt/z01-nfs01"

For NFSv4:

"mount -t nfs -o vers=4 primary.ad1.test:/ifs/data/zone01/z01-nfs01 /mnt/z01-nfs01"

© Superna Inc

3.7. Monitoring DR Readiness for Eyeglass Assisted Failover

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

Monitoring DR Readiness for Eyeglass Assisted Failover

In addition to the Assisted Failover functionality, Eyeglass also provides the following features to monitor your SyncIQ Policy DR Readiness:

- Policy DR Readiness Validation from the DR Dashboard
- Runbook Robot from the DR Dashboard

Policy DR Readiness Validation

The DR Dashboard Policy Readiness tab provides a per SyncIQ Policy summary of all SyncIQ OneFS and Eyeglass Configuration replication Job statuses. The status for each are combined to provide an overall DR Status. The Policy Readiness is updated each time the Eyeglass Configuration Replication task runs.

This information provides the best indicator of DR readiness for failover and allows administrators to check status on each component of failover, identify status, errors and correct them to get each SyncIQ Policy configured and ready for failover.

If all of the SyncIQ Policy Failover Recommendations pass validation, the DR Dashboard status for the Policy is green indicating that the Policy is safe to failover.

If any of the SyncIQ Policy Failover Recommendations do NOT pass validation, the DR Dashboard status for the SyncIQ Policy is red indicating that the SyncIQ Policy is NOT ready to failover. Eyeglass will also issue a System Alarm for any of these conditions.

Additional information can be found in the “Policy Readiness and DFS Readiness” section of the Eyeglass Admin Guide [here](#).

IMPORTANT:

If you make a change to your environment, the following Eyeglass tasks must run before the Policy Readiness will be updated:

- Configuration Replication - Job Definitions Icon

© Superna Inc

3.8. Post SyncIQ Policy Failover Manual Steps

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

Post SyncIQ Policy Failover Manual Steps

Update Networking for Client Access

Make the necessary networking changes required to redirect your clients to the Failover Target Cluster. This may involve:

- PowerScale SmartConnect Zone or Zone Alias updates
- DNS Updates

Update SPNs

For the case where SMB shares and Active Directory are being used for client access, any add or delete of SmartConnect Zone or Zone Alias will require an SPN update.

Refreshing SMB connection after Failover completed and DNS updated

This section describes steps to refresh an SMB connection post failover and DNS update.

1. If client was connected to the share during the failover, unmount the share (disconnect).
2. Remount the share (connect). Check Kerberos and SPN on the target cluster machine account.

3. Test read/write against newly mounted shares.
4. If step 3 fails, the original connection information is likely cached on the client machine. The data in this case would continue to be available, however it would be Read-Only. Writes would fail. To remove the cached connection information, open a Window cmd window and type the command:

```
ipconfig /flushdns
```

5. Test read/write against newly mounted share again.
6. If step 5 fails, it is likely a DNS cache issue and needs to be looked at upstream DNS servers.

Refreshing NFS connection after Failover completed and DNS updated

Please refer to document section steps to refresh an NFS connection post failover and DNS update.

Post SyncIQ Policy Failover Checklist

The following sections outline what can be checked post SyncIQ Policy failover to verify execution of all of the steps.

IMPORTANT: If the failover was done with the **Controlled failover** option unchecked, then some steps on the failover

SOURCE cluster will not have been executed. This is outlined in Appendix 1 in this document.

File System Updates

On the failover SOURCE cluster (the cluster you failed over FROM), the directories and sub-directories corresponding to the SyncIQ Policies that were failed over are read-only.

On the failover TARGET cluster (the cluster you failed over TO), the directories and sub-directories corresponding to the SyncIQ Policies that were failed over are writeable.

SyncIQ Policy Updates

On the failover SOURCE cluster (the cluster you failed over FROM), for the SyncIQ Policies that were failed over:

- SyncIQ Policies are Disabled in OneFS
- Eyeglass configuration replication jobs related to these SyncIQ Policies are in Policy Disabled state
- SyncIQ Policies in OneFS have their schedule set to manual

On the failover TARGET cluster (the cluster you failed over TO), for the SyncIQ Policies that were failed over:

- SyncIQ Policies are Enabled in OneFS

- The corresponding Eyeglass Configuration Replication Jobs are also in Enabled state
- SyncIQ Policies have same schedule that was originally set for the policy on the failover SOURCE cluster

NOTE: If you have Eyeglass “INITIALSTATE” property set to disabled for AUTO jobs (check this using the “igls adv initialstate show” command in the Eyeglass Admin Guide [here](#)), the Eyeglass Configuration Replication job for the mirror SyncIQ Policy created during the first failover will be in User Disabled state. This job should be enabled following the instructions in the Eyeglass Admin Guide [here](#).

Quota Updates

After the upgrade, there should be no quotas on the failover SOURCE cluster for the SyncIQ Policies that were failed over. On the failover TARGET cluster you should find all quotas for the SyncIQ Policies that were failed over.

© Superna Inc

3.9. Controlled Failover Option Results Summary

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

Controlled Failover Option Results Summary

The following table indicates which failover steps are executed based on whether or not the **Controlled failover** option was selected when the SyncIQ Policy failover was initiated.

Steps	Description	Executed on	SyncIQ Policy	Controlled Failover selected	Controlled Failover NOT selected
1 - Ensure that there is no live access to data	Check for open files. If Open files found, decide whether to failover or wait to be closed.	Source	Manual	Not applicable - manual step	Not applicable - manual step
2 - Begin Failover	Initiate Failover from Eyeglass	Eyeglass	Manual	Not applicable - manual step	Not applicable - manual step
3 - Validation	Wait for other Eyeglass Failover jobs to complete	Eyeglass	Automated by Eyeglass	Step Executed	Step Executed
4 - Synchronize data	Run all OneFS SyncIQ policy jobs related to the Access Zone being failed over	Source	Automated by Eyeglass	Step Executed	Step NOT Executed
5 - Synchronize configuration (shares/export/alias)	Run Eyeglass configuration replication 1	Eyeglass	Automated by Eyeglass (configuration exists on source and target)	Step Executed	Step Executed based on last known data in Eyeglass

					*If you do not want this, uncheck the “Config Sync” option to skip this step
6 - Synchronize quota(s)	Run Eyeglass Quota Jobs related to the SyncIQ Policy or Access Zone being failed over	Eyeglass	Automated by Eyeglass (deleted on source cluster and created on target cluster)	Step Executed	Step Executed based on last known data in Eyeglass
7 - Record schedule for SyncIQ policies being failed over	Get schedule associated with the SyncIQ policies being failed over on OneFS	Source	Automated by Eyeglass	Step Executed	Step NOT Executed
8 - Prevent SyncIQ policies being failed over from running	Set schedule on the SyncIQ policy(s) to manual on source cluster	Source	Automated by Eyeglass	Step Executed	Step NOT Executed
9 - Provide write access to data on target	Allow writes to SyncIQ policy(s) related to failover2	Target	Automated by Eyeglass	Step Executed	Step Executed
10 - Disable SyncIQ on source and make active on target	Resync prep SyncIQ policy related to failover (Creates MirrorPolicy on target) from OneFS	Source	Automated by Eyeglass	Step Executed	Step NOT Executed
11 - Set proper SyncIQ schedule on target	Set schedule on MirrorPolicy(Target) using schedule from step 6 from OneFS for policy(s) related to the Failover	Target	Automated by Eyeglass	Step Executed	Step NOT Executed
12 - Remove quotas on directories that are target of SyncIQ (PowerScale best practice)	Delete all quotas on the source for all the policies	Source	Automated by Eyeglass	Step Executed	Step NOT Executed
13 - Change SmartConnect Zone on Source so not to	Rename SmartConnect Zones and Aliases	Source	Manual	Not applicable - manual	Not applicable - manual

resolve by Clients	(Source)			step	step
14 - Avoid SPN Collision	Sync SPNs in all AD providers to current SmartConnect Zone names and aliases (Source)	Source	Manual (deletes smartconnect SPN from source cluster machine account)	Not applicable - manual step	Not applicable - manual step
15 - Move SmartConnect Zone to Target	Add source SmartConnect Zone(s) as Alias(es) on (Target)	Target	Manual	Not applicable - manual step	Not applicable - manual step
16 - Update SPN to allow for authentication against target	Sync SPNs in all AD providers to current SmartConnect Zone names and aliases (Target)	Target	Manual (adds new SmatrConnect alias SPN's to target cluster machine account)	Not applicable - manual step	Not applicable - manual step
17 - Repoint DNS to the Target cluster IP address	Update DNS delegations for all SmartConnect Zones that are members of the Access Zone	DNS	Manual	Not applicable - manual step	Not applicable - manual step
18 - Refresh session to pick up DNS change	Remount the SMB share(s)	SMB Client Machines	Manual on clients	Not applicable - manual step	Not applicable - manual step

1. Initiates Eyeglass Configuration Replication task for all

Eyeglass jobs

2. SyncIQ does NOT modify the ACL (Access control settings on

the file system). It locks the file system. `ls -l` will be

identically on both source and target

4. Eyeglass Runbook Robot Guide

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

- [Introduction to this Guide](#)
- [Runbook Robot \(Automate DR Testing on a schedule\)](#)
- [Basic DR Robot Configuration](#)
- [Advanced DR Robot Configuration](#)
- [Advanced Settings](#)

© Superna Inc

4.1. Introduction to this Guide

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

- [Overview](#)
- [Runbook DR Robot FAQ's Requirements](#)

The purpose of this guide is to assist you in configuring and testing DR readiness using the Eyeglass Runbook Robot.

Overview

To gain the most from Eyeglass Readiness features, the Eyeglass Runbook Robot allows customers to set up and have continued DR operating between pairs of clusters. This feature also allows testing of application failover logic by creating configuration data and copying data into the Robot Access zone.

This feature runs with a test Access Zone to eliminate impact on production Access Zones. The feature uses specially named Access Zone and SyncIQ policies so that there is no conflict with production Access Zone and policy data.

The following test validations are all done on a daily basis. The DR dashboard will be updated along with any failures sent as critical events. These daily Runbook Robot validation tests are your best indicator that your cluster is ready for a failover.

1. API access to both clusters is functioning - **Validated**
2. API access allows creation of export, share, quota - **Validated**
3. NFS mount of data external to the cluster functions - **Validated**
4. DNS resolution for SmartConnect is checked when Eyeglass configures itself to use SmartConnect service IP as its DNS resolver on the source to verify SmartConnect zone functionality on mount of data requests - **Validated**
5. SyncIQ policy replication completes between source and destination cluster when data is written to the source - **Validated**
6. Configuration replication of test configuration from source to destination - **Validated**
7. SyncIQ failover to target cluster - **Validated**
8. Test data access on target cluster post failover - **Validated**
9. Verify data integrity of the test data on target cluster - **Validated**
10. Configuration Sync of quotas from source to target on failover - **Validated**
11. Delete of Quotas on source cluster - **Validated**
12. SyncIQ Failback from target to source cluster - **Validated**

The setup of a Robot DR job is as simple process that has 3 modes of operation.

1. **Basic DR Robot** - This mode only requires a SyncIQ policy with a certain pre-defined name to exist. This allows Eyeglass to use this policy to automate writing data, failing over and failing back between the clusters the policy is configured to use.

DR Coverage: This basic mode does not test all the possible functions exhaustively and only validates basic DR readiness checks.

2. **Advanced DR Robot** - This mode requires the same SyncIQ policy to be created in an Access Zone with a pre-defined name AND to be the only policy in the Access Zone.

DR Coverage: Full API, DNS, SmartConnect Zone aliasing, data access, mounting of data, SyncIQ, data replication, configuration creation, and configuration replication.

3. **Advanced DR DFS mode Robot** - This mode only requires a single policy to exist and can be used with DFS mode. It's the same configuration as the Advanced Robot but allows DFS only customers to configure DFS clients to test write access to the Robot data for testing DFS switching and failover operations. It will still write data over an NFS export but also shares can be created to present a DFS folder for testing with the Robot data to make sure the DFS folder is accessible daily.

DR Coverage: mounting of data, SyncIQ, Data replication, configuration creation, configuration replication

Validates DFS switching and AD configuration in DFS if configured.

Configuration only changes by enabling DFS mode on the Robot Policy after creation. See the DFS guide on how to enable DFS mode on existing policies

If Zone Readiness is not green, the Robot job will not run. If an error occurs during the Robot DR job the job will be stopped at the failed step. Use the Failover Log to learn about failed steps.

Runbook DR Robot FAQ's Requirements

The following are some Frequently Asked Questions that users have asked in configuring the Runbook Robot:

1. Can more than one Runbook Robot be setup on a single Eyeglass appliance?
No, currently only a single pair of clusters are supported with Runbook Robot feature.
2. Does the Robot support Access Zones? **yes see Advanced setup.**
3. Can I create a Robot SyncIQ policy in the system zone? **Yes.**
4. Does Run Book Robot support SMB access when mounting the cluster ? **no, not at this time.**

5. Can I copy data into the Robot Policy folder to get more data replicated? **Yes**
this can be used to test failover with Eyeglass as well.

© Superna Inc

4.2. Runbook Robot (Automate DR Testing on a schedule)

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

Runbook Robot (Automate DR Testing on a schedule)

Overview

Many organizations schedule DR tests during maintenance windows and weekends, only to find out that the DR procedures did not work or documentation needed to be updated. Eyeglass Run Book Robot feature automates DR run book procedures that would normally be scheduled in off peak hours, and avoids down time to validate DR procedures, providing Failover and Fallback automation tests with reporting.

This level of automation provides high confidence that your PowerScale storage is ready for failover with all of the key functions executed on a daily basis. In addition to automating failover and fallback, Eyeglass operates as a cluster witness and mounts storage on both source and destination clusters (the same way the cluster users and machines mount storage externally using access zone mount paths).

Run Book Robot Failover Coverage with Basic Setup

1. API access to both clusters is functioning - **Validated**
2. API access allows creation of export, share, quota - **Validated**

3. NFS mount of data external to the cluster functions - **Validated**
4. SyncIQ policy replication completes between source and destination cluster when data is written to the source - **Validated**
5. Configuration replication of test configuration from source to destination - **Validated**
6. SyncIQ failover to target cluster - **Validated**
7. Test data access on target cluster post failover - **Validated**
8. Verify data integrity of the test data on target cluster - **Validated**
9. Configuration Sync of quotas from source to target on failover - **Validated**
10. Delete of Quotas on source cluster - **Validated**
11. SyncIQ Failback from target to source cluster - **Validated**

The above validations are all done daily and the DR dashboard updated along with any failures sent as critical event. When all the above validations pass on a daily basis, this is the best indicator that your cluster is ready for a failover.

Refer to the [RunBookRobot Admin Guide](#) for instructions on setting up and running the Runbook Robot.

© Superna Inc

4.3. Basic DR Robot Configuration

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

Basic DR Robot Configuration

The basic option requires a non-production, Eyeglass specific SyncIQ policy to exist. The Runbook Robot will exercise the failover and fallback of the data for this SyncIQ policy. This configuration does not fully exercise all automation required for failover but it's quick and easy to setup.

Prerequisites

1. SyncIQ policy in any Access Zone BUT only "***EyeglassRunbookRobot***" prefixed policies will execute the failover logic.
2. This type of Robot policy will not failover SmartConnect names, alias, SPN's.
3. This is designed to ensure no SmartConnect Zones used by non Robot policies will be failed over in the Access Zone:
2. The Basic Robot is intended to operate on non-production data. The steps below include directions to create the Runbook Robot specific SyncIQ policy on the cluster for the Basic Runbook Robot operation.
3. If you create shares, exports or quotas in the path of the Robot SyncIQ policy, they will be failed over as well using normal configuration sync jobs. It's a good way to test the

whole failover of the configuration. The configuration names of shares, exports can be any name.

Note: Eyeglass will modify the description, Root Client and Map Root User settings as required for Robot operation on an existing export with same path as the Robot SyncIQ Policy root path.

4. Quotas will be failed over during the Robot execution and deleted on source automatically if quotas have been created within the SyncIQ policy used for the Robot.
6. Eyeglass Hostname on the Network Card Setup and the Network Settings (setup using yast during initial Eyeglass appliance configuration) must be identical.

Configuration Steps for Basic DR Robot

The steps are the following:

1. Log into the PowerScale Source cluster via OneFS.
2. Click on **Data Protection**, then **SyncIQ**, then on the **Policies** tab.
3. Click on **+ Create a SyncIQ Policy**.



SyncIQ

Summary Policies Reports Local Targets Performance Rules Settings

SyncIQ Policies

[+ Create a SyncIQ Policy](#)

<input type="checkbox"/>	Policy Name	State	Last Known Good	Schedule	Source Directory	Target Host : Directory	Actions
<input type="checkbox"/>	Eyeglass-Runbo...	Enabled	No success	Manual	/ifs/data/testing/test	172.16.82.131 : /ifs/data/testing/test	View / Edit More ▾
<input type="checkbox"/>	valid_policy_1	Enabled	06/17/2015	Manual	/ifs/data/testing/1	172.16.82.131 : /ifs/data/testing/1	View / Edit More ▾
<input type="checkbox"/>	valid_policy_10	Enabled	06/17/2015	Manual	/ifs/data/testing/10	172.16.82.131 : /ifs/data/testing/10	View / Edit More ▾
<input type="checkbox"/>	valid_policy_11	Enabled	06/17/2015	Manual	/ifs/data/testing/11	172.16.82.131 : /ifs/data/testing/11	View / Edit More ▾
<input type="checkbox"/>	valid_policy_12	Enabled	06/17/2015	Manual	/ifs/data/testing/12	172.16.82.131 : /ifs/data/testing/12	View / Edit More ▾
<input type="checkbox"/>	valid_policy_13	Enabled	06/17/2015	Manual	/ifs/data/testing/13	172.16.82.131 : /ifs/data/testing/13	View / Edit More ▾
<input type="checkbox"/>	valid_policy_14	Enabled	06/17/2015	Manual	/ifs/data/testing/14	172.16.82.131 : /ifs/data/testing/14	View / Edit More ▾

4. In the **Settings** section of the **Create SyncIQ Policy** window, enter

EyeglassRunbookRobot-XYZ (where XYZ is a random string) in the **Policy Name** field.

IMPORTANT: This name format must be followed exactly in order for the Basic Runbook Robot job to execute.

Create SyncIQ Policy

* = Required field

— Settings —

* Policy Name

Description

Enable this policy

Action

Copy
 Synchronize

Run Job

Only manually
 On a schedule
 Whenever the source is modified

5. In **Source Cluster section**, select a Source Root Directory by clicking **Browse** next to the box and navigating to the desired folder.

IMPORTANT: This should be a directory containing NON production data as it will be failed over by default once a day.

Create SyncIQ Policy
* = Required field

– Source Cluster

* Source Root Directory
 Browse... (circled in red)

Included Directories
 Browse...

+ Add another directory path

Excluded Directories
 Browse...

+ Add another directory path

i The File Matching Criterion of "Path," "Modified," "Accessed," and "Created" are only valid for Policies with the Action "Copy" and will cause Policies with the Action "Synchronize" to fail.

File Matching Criteria

– IF Condition

–Select Filter Type-- **▼**

[Delete this block](#) | [Add an "And" condition](#)

+ Add an "Or" condition

Restrict Source Nodes

Run the policy on all nodes in this cluster

6. In **Target Cluster section** under **Target Host**, enter the target cluster SmartConnect zone name used for SyncIQ replication of the Target cluster , and under **Target Directory** the path on the Target cluster to where the data will be replicated .

Create SyncIQ Policy

* = Required field

– Target Cluster

* Target Host
172.16.82.131

* Target Directory
/ifs/data/testing/test2

Connect only to the nodes within the target cluster SmartConnect Zone

7. Do not enter or modify anything in the **Target Snapshots** and **Advanced Settings** sections.

9. Click on **Create Policy**. The policy will appear in the **SyncIQ Policies** chart.

SyncIQ Policies							+ Create a SyncIQ Policy
	Policy Name	State	Last Known Go	Schedule	Source Directory	Target Host : Directory	Actions
<input type="checkbox"/>	EyeglassRunbookRobot-test	Enabled	No success	Manual	/ifs/data/testing/test	172.16.82.131 : /ifs/data/testing/test	View / Edit More ▾
<input type="checkbox"/>	EyeglassRunbookRobot-test1	Enabled	No success	Manual	/ifs/data/testing/test1	172.16.82.131 : /ifs/data/testing/test1	View / Edit More ▾
<input type="checkbox"/>	EyeglassRunbookRobot-test2	Enabled	No success	Manual	/ifs/data/testing/test2	172.16.82.131 : /ifs/data/testing/test2	View / Edit More ▾
<input type="checkbox"/>	valid_policy_1	Enabled	06/17/2015	Manual	/ifs/data/testing/1	172.16.82.131 : /ifs/data/testing/1	View / Edit More ▾
<input type="checkbox"/>	valid_policy_10	Enabled	06/17/2015	Manual	/ifs/data/testing/10	172.16.82.131 : /ifs/data/testing/10	View / Edit More ▾
<input type="checkbox"/>	valid_policy_11	Enabled	06/17/2015	Manual	/ifs/data/testing/11	172.16.82.131 : /ifs/data/testing/11	View / Edit More ▾
<input type="checkbox"/>	valid_policy_12	Enabled	06/17/2015	Manual	/ifs/data/testing/12	172.16.82.131 : /ifs/data/testing/12	View / Edit More ▾

9. Click on **More** and select **Start Job** on newly created policy.

11. Ensure that the policy completes successfully.

12. Log into Eyeglass.

13. Open the **Jobs** window and select the **Job Definitions** section.

14. Under the Configuration Replication jobs, look for a Job Name that resembles *SourceClusterName_EyeglassRunbookRobot-XYZ* (where XYZ is the random string chosen when creating the policy). Make sure it's State is OK.

Notes:

1. This Job will be created after Eyeglass discovery has found the SyncIQ Policy created in the previous steps. If the Job is not present, wait for the next Eyeglass discovery cycle and then check again.
2. **NOTE: You can enable DFS mode on the basic robot policy. Eyeglass will still use NFS to write data to the cluster but DFS mount can be created to verify failover and fallback operations over DFS by manually check the DFS mount each day. Eyelgass will still execute normal NFS failover and fallback data integrity check.**

 - a. **How to Enable DFS mode Basic Robot**
 - b. **Find the policy the Configuration replication section of the Jobs icon.**
 - c. **Select the checkbox on the left-hand side for the Job mentioned above and click the Select a bulk action button at the bottom of the Jobs window. This time, click the Enable/Disable option.**
 - d. **Select the checkbox on the left-hand side for the Job mentioned above and click the Select a bulk action button at the bottom of the Jobs window and click the Enable/Disable Microsoft DFS option.**
3. Inventory runs with every Eyeglass Replication cycle (every 5 minutes by default) so you may need to wait for the next Eyeglass Replication cycle for the new Job to appear in the jobs window.

<input type="checkbox"/>	Job Name ↑	Policy	Type	Last Run Date	State
Configuration Replication: Share, Export, Alias replication (AUTOMATIC)					
<input type="checkbox"/>	Isilon-Sim-7202-Source_EyeglassRunbookRobot-test	EyeglassR...	AUTO	29/06/2015, 10:40:12	OK
<input type="checkbox"/>	Isilon-Sim-7202-Source_EyeglassRunbookRobot-test1	EyeglassR...	AUTO	29/06/2015, 10:40:12	OK
<input checked="" type="checkbox"/>	Isilon-Sim-7202-Source_EyeglassRunbookRobot-test2	EyeglassR...	AUTO	29/06/2015, 10:40:12	OK
<input type="checkbox"/>	Isilon-Sim-7202-Source_valid_policy_1	valid_polic...	AUTO	26/06/2015, 15:57:49	User Disa
<input type="checkbox"/>	Isilon-Sim-7202-Source_valid_policy_10	valid_polic...	AUTO	26/06/2015, 15:57:52	User Disa
<input type="checkbox"/>	Isilon-Sim-7202-Source_valid_policy_11	valid_polic...	AUTO	26/06/2015, 15:57:59	User Disa
<input type="checkbox"/>	Isilon-Sim-7202-Source_valid_policy_12	valid_polic...	AUTO	26/06/2015, 15:57:56	User Disa
<input type="checkbox"/>	Isilon-Sim-7202-Source_valid_policy_13	valid_polic...	AUTO	26/06/2015, 15:57:47	User Disa
<input type="checkbox"/>	Isilon-Sim-7202-Source_valid_policy_14	valid_polic...	AUTO	26/06/2015, 15:57:50	User Disa
<input type="checkbox"/>	Isilon-Sim-7202-Source_valid_policy_15	valid_polic...	AUTO	26/06/2015, 15:57:57	User Disa
<input type="checkbox"/>	Isilon-Sim-7202-Source_valid_policy_16	valid_polic...	AUTO	26/06/2015, 15:57:46	User Disa
<input type="checkbox"/>	Isilon-Sim-7202-Source_valid_policy_17	valid_polic...	AUTO	26/06/2015, 15:57:55	User Disa
<input type="checkbox"/>	Isilon-Sim-7202-Source_valid_policy_18	valid_polic...	AUTO	26/06/2015, 15:57:56	User Disa
<input type="checkbox"/>	Isilon-Sim-7202-Source_valid_policy_19	valid_polic...	AUTO	26/06/2015, 15:57:58	User Disa
<input type="checkbox"/>	Isilon-Sim-7202-Source_valid_policy_2	valid_polic...	AUTO	26/06/2015, 15:58:00	User Disa
<input type="checkbox"/>	Isilon-Sim-7202-Source_valid_policy_20	valid_polic...	AUTO	26/06/2015, 15:58:00	User Disa

14. If the SyncIQ Policy Name that was created has been entered properly, a new type of Job category will appear in Eyeglass: Failover: Runbook Robot (AUTOMATIC).

<input type="checkbox"/>	Job Name ↑	Policy	Type	Last ...	State
Configuration Replication: Share, Export, Alias replication (AUTOMATIC)					
Failover: Quota Failover (RUN MANUALLY)					
Failover: Runbook Robot (AUTOMATIC)					
<input type="checkbox"/>	Isilon-Sim-7202-Source_EyeglassRunbookRobot-test1_runbook_robot	EyeglassR...	RUNBOOK...	n/a	Pending
<input checked="" type="checkbox"/>	Isilon-Sim-7202-Source_EyeglassRunbookRobot-test2_runbook_robot	EyeglassR...	RUNBOOK...	n/a	Pending
<input type="checkbox"/>	Isilon-Sim-7202-Source_EyeglassRunbookRobot-test_runbook_robot	EyeglassR...	RUNBOOK...	n/a	Pending

15. By default, the Job will be queued to run at midnight every day.

1. If you want to test it after creating the policy, select the job in the jobs windows and use bulk action to run now.

16. When the Runbook Robot Job runs, in the **Running Jobs** section a job will appear with the Name *Runbook Robot*, followed by a time stamp.

The screenshot shows the 'Jobs' interface with the 'Running Jobs' tab selected. The main area displays a table of running jobs:

...	Job Name	Started...	Finished
X	Runbook Robot 1435608900010	29/06/2...	29/06/2...
X	Runbook Robot 1435608600000	29/06/2...	29/06/2...
✓	Configuration Replication 1435608600001	29/06/2...	29/06/2...

Below the table, a 'Job Details' section shows a hierarchical tree view of the 'Runbook Robot Failover Test' job:

- Runbook Robot Failover Test
 - + Failover Preparation
 - + Sync IQ Policy Failover
 - + Failover Validation
 - Runbook Robot Failover Cleanup
 - Isilon-Sim-7202-Source_EyeglassRunbookRobot-test_runbook_robot
 - Isilon-Sim-7202-Source_EyeglassRunbookRobot-test1_runbook_robot

17. The Runbook Robot job will run a three-stage Failover Test, which consists of a Failover Preparation, SyncIQ Policy Failover, and Failover Validation. Once this is done, it will run a Failover Cleanup.

1. The Failover Preparation component consists of the following steps:

2. The job will create a new Eyeglass export or update an existing Eyeglass export with the following parameters:
 - c. Description: Superna Eyeglass Runbook Robot export
 - Appliance ID
 - d. **Root Clients:** The Eyeglass IP address used
 - e. **Map Root User** settings
 - f. **Directory Paths:** The Source Root Directory chosen when creating the SyncIQ Policy
 - g. This export is then mounted on the Eyeglass appliance.
 - h. Data is then written to the export, which is dependent on the SyncIQ Policy Path and timestamp.
 3. The export is then unmounted
 4. **To view the exports created:** open OneFS with the IP address of the source cluster, click on the Protocols tab, then click on UNIX Sharing (NFS), followed by the NFS Exports tab.

The screenshot shows the OneFS Storage Administration interface. The top navigation bar includes links for Review recent events, Log out, and Help. Below the navigation is a cluster status bar indicating Cluster Name: Isilon-Sim-7202-Source (OneFS Version: 7.2.0.2). The main menu has tabs for DASHBOARD, CLUSTER MANAGEMENT, FILE SYSTEM, DATA PROTECTION, ACCESS, and PROTOCOLS. Under PROTOCOLS, sub-tabs include Windows Sharing (SMB), UNIX Sharing (NFS) (which is selected and highlighted in blue), FTP Settings, HTTP Settings, and ACLs. A sub-menu for UNIX Sharing (NFS) shows Current Access Zone: System. Below this are buttons for NFS Exports, NFS Aliases, Export Settings, and Global Settings. The NFS Exports section displays a table with the following data:

Export ID / Path	Description	Action
110 Path: /ifs/data/testing/test2	Superna Eyeglass Runbook Robot export. Appliance ID:	View details Delete
109 Path: /ifs/data/testing/test	Superna Eyeglass Runbook Robot export. Appliance ID:	View details Delete
106 Path: /ifs/data/testing/21/dir_5		View details Delete
105 Path: /ifs/data/testing/21/dir_4		View details Delete
104 Path: /ifs/data/testing/21/dir_3		View details Delete
103 Path: /ifs/data/testing/21/dir_2		View details Delete

- 5.
6. Once the preparation is successfully completed, Eyeglass performs a SyncIQ Policy Failover.
7. On completion of the failover, Runbook Robot will run a Failover Validation in order to make sure that the data was failed over without issues. The Failover Validation component consists of the following steps:
 - c. The export created is mounted on the Eyeglass appliance using the IP address of the failover target cluster.
 - d. The timestamp is read to confirm it is identical to what was written during the Failover Preparation step.
 - e. A new timestamp is written to the Failover.ts file.

- f. This new timestamp is then read to confirm it has been written correctly.
 - g. This export is then unmounted from the Eyeglass appliance.
- 8. Once the Failover Test has been completed, the Runbook Robot job will run a Failover Cleanup, which consists of unmounting the export from the Eyeglass appliance.
- 9. The default Robot jobs run every day at midnight and executes a failover Robot policy.
 - c. Creates export.
 - d. Mounts the cluster writes test data using the export created.
 - e. It failovers the policy.
 - f. Runs the policy.
 - g. Mounts the data on synced export on target.
 - h. Unmounts.
 - i. Moves the schedule on policy to the target.
 - j. Runs resync prep on source.
 - k. Goes to sleep until time to fallback.
- 10. Methods to verify it was successful:

- c. Check the DR Dashboard policies tab and verify it's green.
- d. Check the cluster SyncIQ policy status on both clusters and make sure the policy moves from one cluster to the other each day.
- e. Make sure the test file exists with a current date and time stamp on the active cluster (**Hint**: look in the policy path root file system for the test file).
- f. If a quota was applied to Robot policy path, make sure the quota moved to the target cluster AND was deleted on the source cluster (Eyeglass moves policies on failover).
- g. Review the Failover Log (DR Assistant / Failover History / Open log file). Failover log may also be downloaded from the Failover Log Viewer using the Download File link.

DFS Mode manual Validation of Basic Robot

© Superna Inc

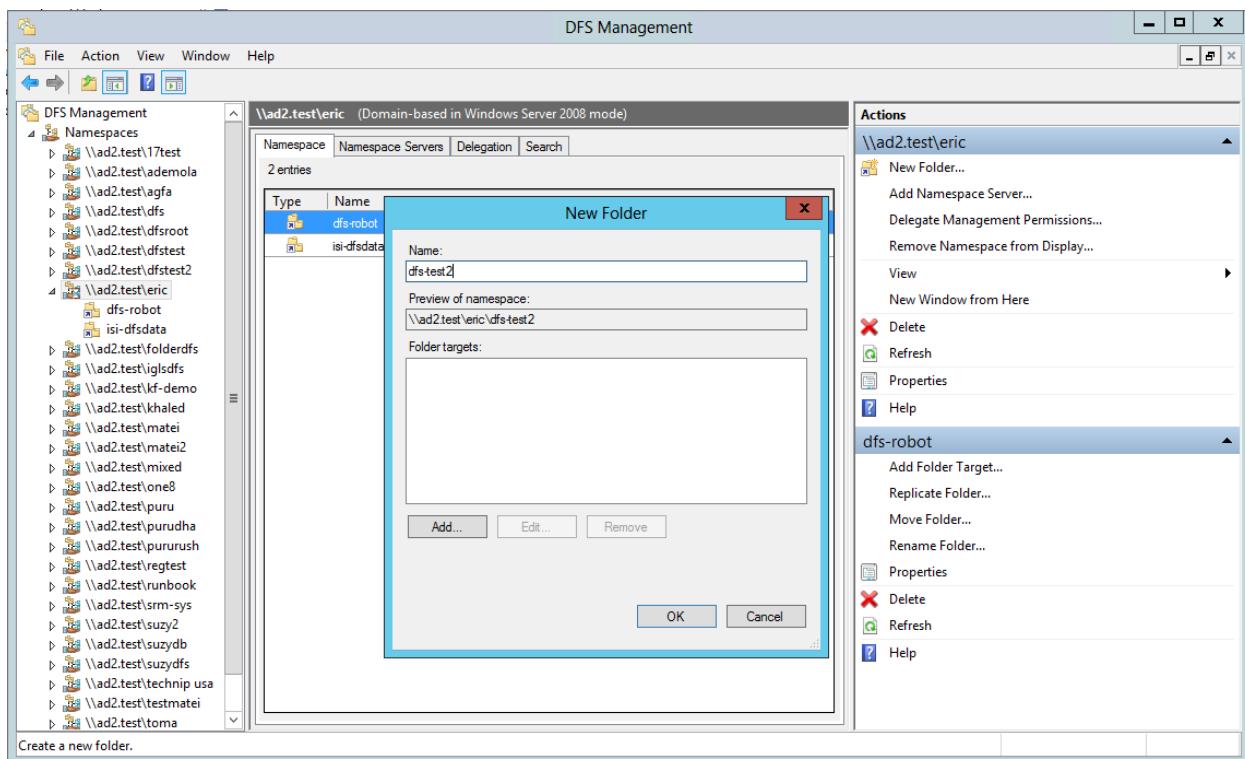
4.3.1. Windows DFS Configuration Example For Basic Robot

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

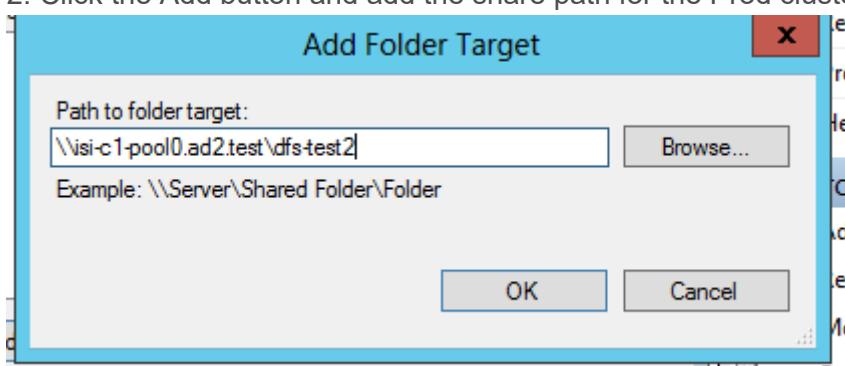
Windows DFS Configuration Example For Basic Robot

If you want see how the DFS Mode Runbook Robot changes the DFS Target from Production cluster to DR cluster follow the steps below:

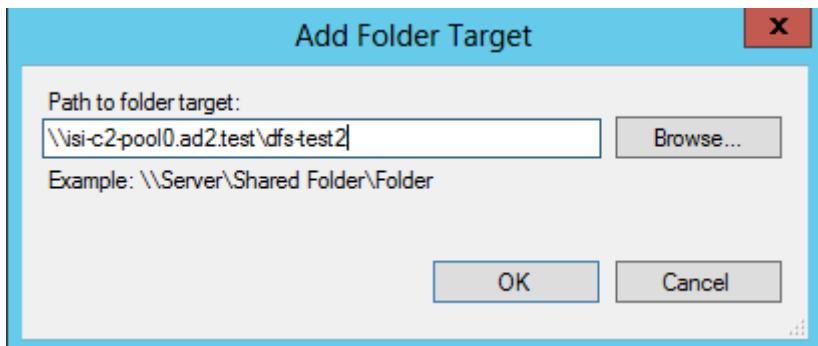
1. Open your DFS Management console and navigate to the Namespace you want to add your new folder to and select “New Folder” from the right hand menu and enter the name of the new DFS Share in the Name field.



2. Click the Add button and add the share path for the Prod cluster and click Ok.

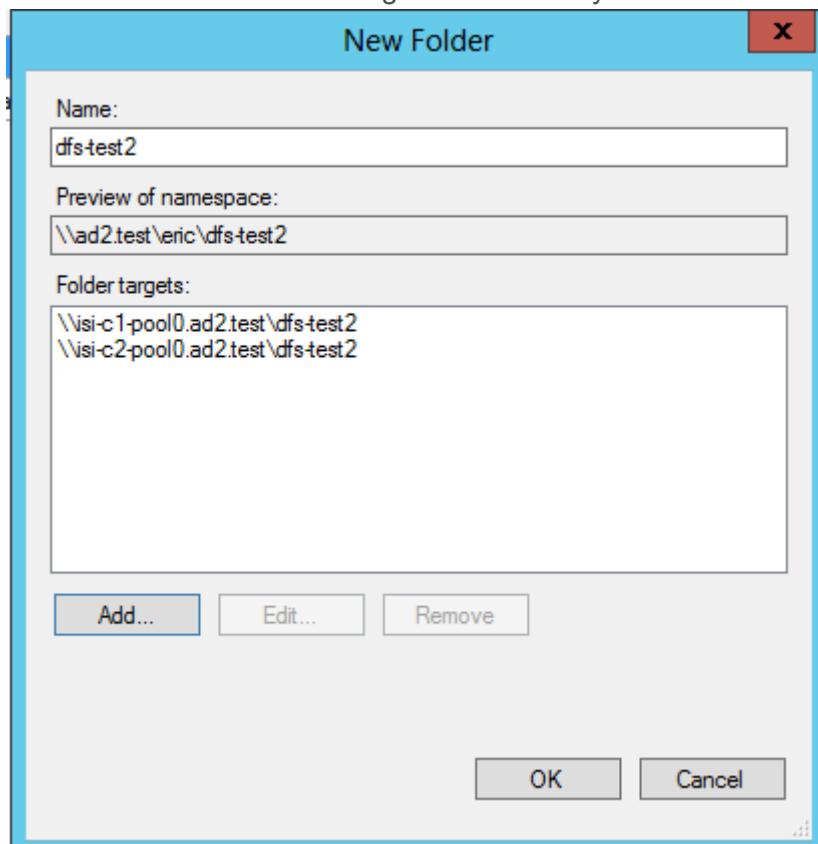


3. Click the Add button and add the share path for the DR cluster and click Ok.

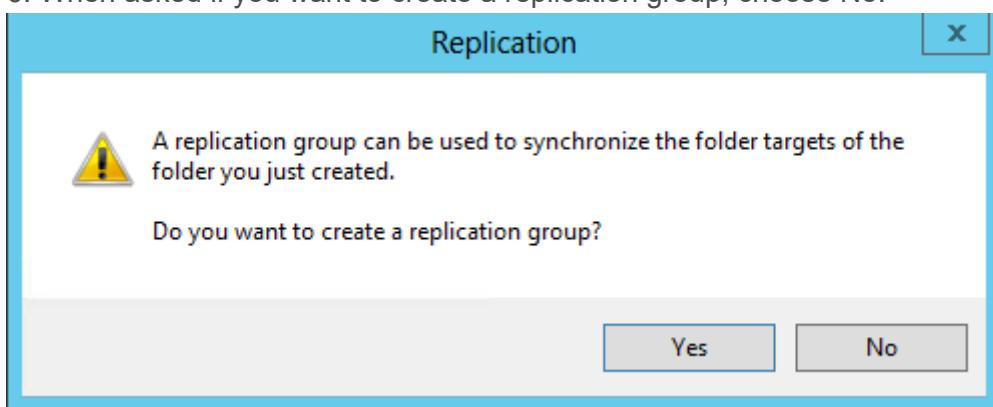


4. When adding the DR cluster, you may experience a long wait as the share on the DR cluster may not already exist, please wait until another window has come up. This will likely happen if you haven't ran the Configuration Replication for the DFS Mode Job you recently added. **If you are prompted to create the folder (if it does not exist), choose No.**

5. You will then have both targets added and you can click Ok.

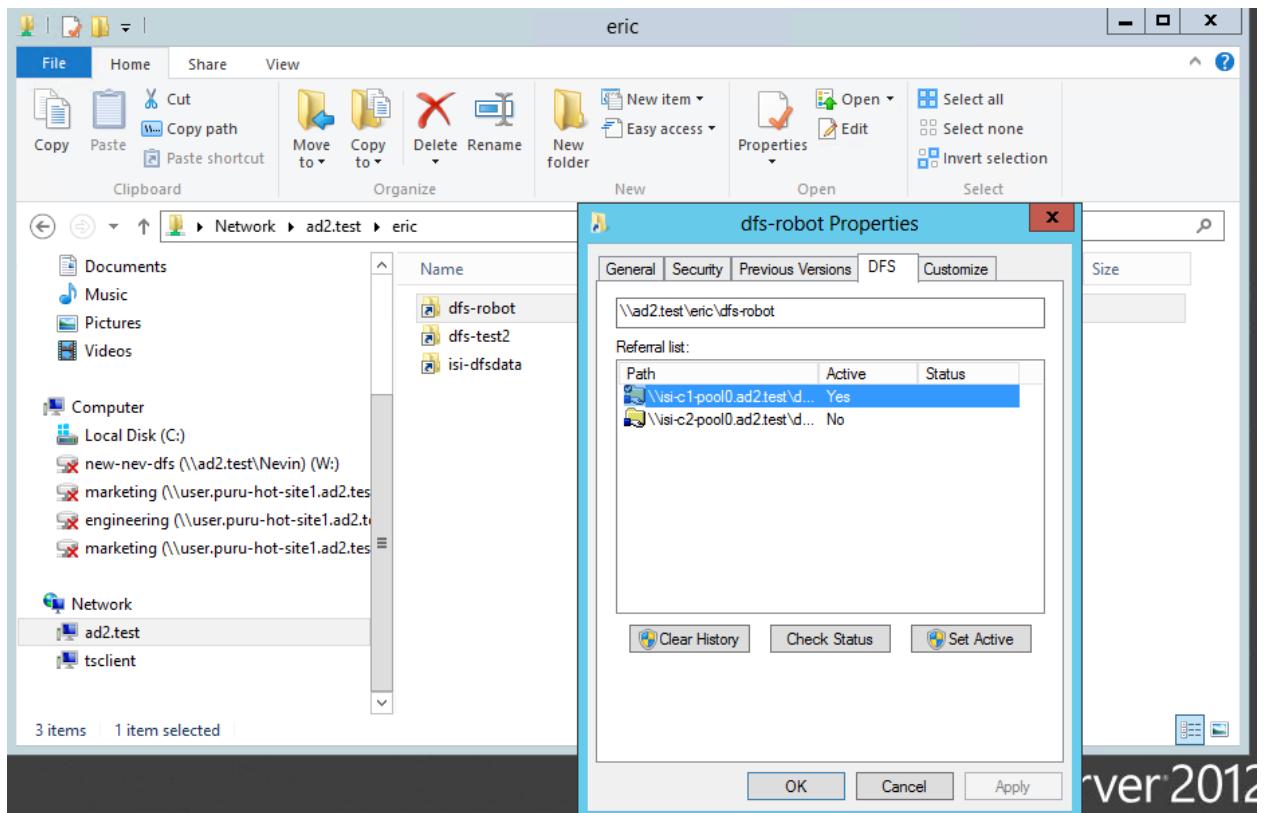


6. When asked if you want to create a replication group, choose No.

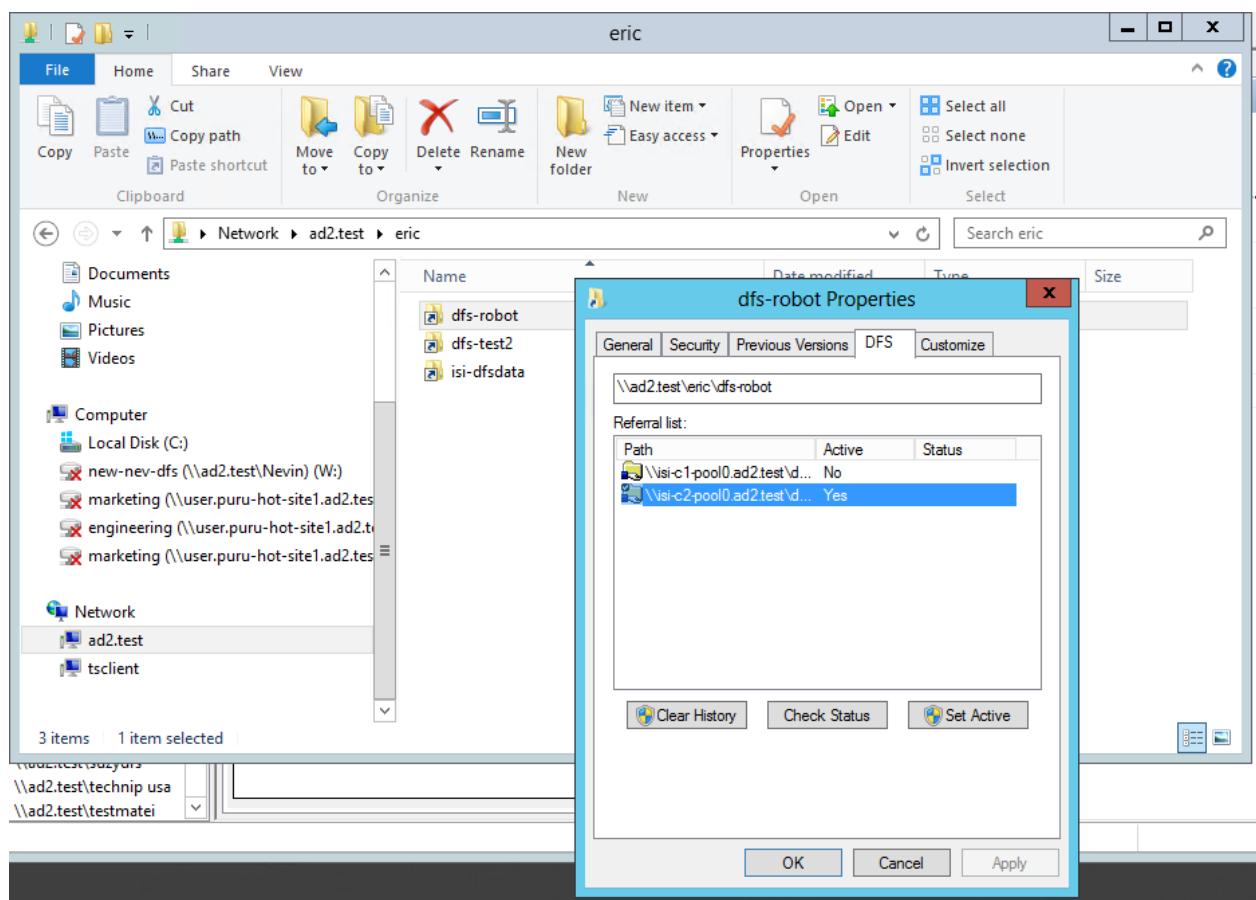


7. To view the active DFS Target through Windows Explorer:

- Open Windows Explorer and navigate to your DFS Namespace path where your newly created folder is in (this example: \\ad2.test\eric).
- Right-click on the folder you had just created and select Properties.
- Goto the DFS tab and you will see the two DFS targets you had created.
- The active target will have a Yes under the “Active” column. See image below:



- After a failover of a Sync IQ policy in DFS Mode, if you check the DFS Properties window as show above, you will see that the Target list will now show the DR cluster target is now showing as Active. See image below:



© Superna Inc

4.4. Advanced DR Robot Configuration

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

Advanced DR Robot Configuration

This option exercises all Eyeglass Failover automation and more closely follows the steps for an Access Zone failover and fallback operation. This takes more time to set up, but offers the highest level of confidence everything required for failover is in place for production Access Zones.

Prerequisites

1. Dedicated IP pool added as member to the Robot Access zone.
1. Create SmartConnect Zone name of your choosing on source and target cluster IP pools.
2. SyncIQ policy in an Access Zone with Runbook Robot prefixed name AND only one Runbook Robot prefixed policy can exist in this Access Zone. If any other SyncIQ policies are detected the Robot will disable itself and stop functioning. This is designed to ensure no production data ends up getting failed over in the Access Zone.
3. Share or export created in the path of the Robot SyncIQ policy, They will be failed over as well using normal configuration sync jobs. It's a good way to test the whole

failover of configuration. The configuration names of shares, exports can be any name.

Note: *Eyeglass will modify the description, Root Client and Map Root User settings as required for Robot operation on an existing export with the same path as the Robot SyncIQ Policy root path.*

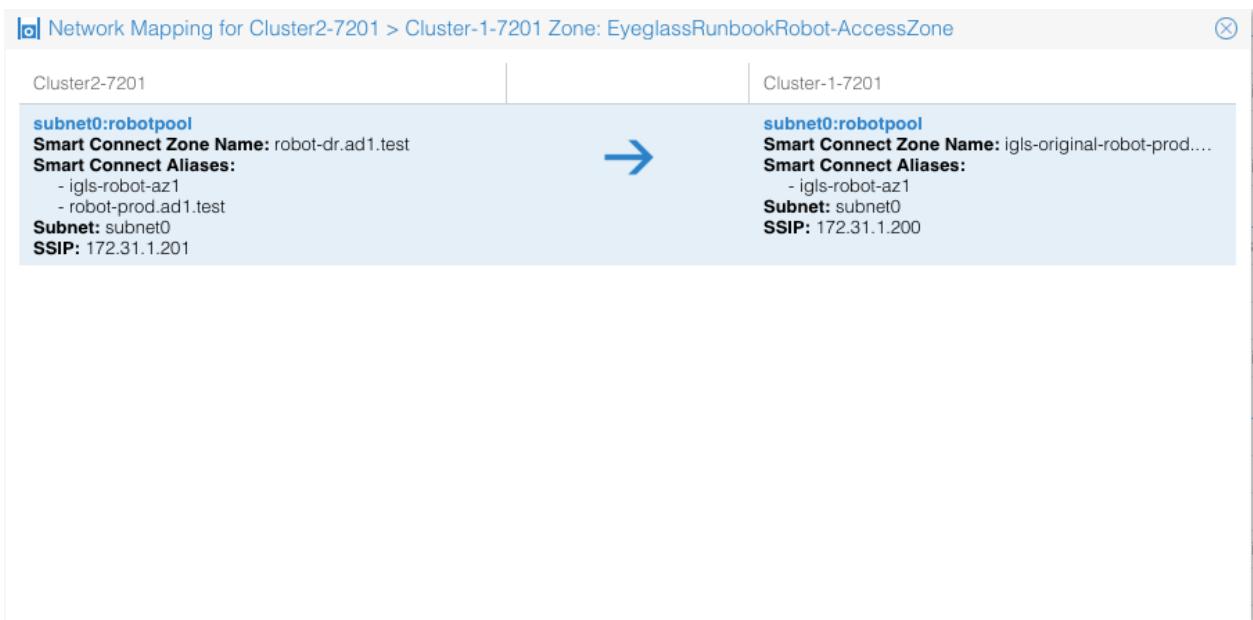
4. Quotas will be failed over during the Robot execution and deleted on the source automatically if quotas have been created within the SyncIQ policy used for the Robot.
5. Zone Readiness task must be enabled (default state is disabled). The Zone Readiness jobs can be enabled from the jobs windows. or the webshell CLI with (igls admin schedules set --id Readiness --enabled true)
6. Zone Readiness task must have run and Zone Readiness status for Robot Zone must be OK. Then Configuration Replication for the Robot Job must have been run successfully.
7. Eyeglass Hostname on the Network Card Setup and the Network Settings (setup using yast during initial Eyeglass appliance configuration) must be identical.

Failover Readiness (AUTOMATIC)			
<input checked="" type="checkbox"/>	Cluster-1-7201_Cluster2-7201	READINESS	9/16/2015, 4:11:10 PM 
<input type="checkbox"/>	Cluster2-7201_Cluster-1-7201	READINESS	9/16/2015, 4:11:10 PM 
Failover Replication Robot (AUTOMATIC)			

Preparation and planning instructions for Zone Readiness can be found in the [Eyeglass Access Zone Failover Guide](#).

Configuration Steps for the Advanced DR Robot

1. Create an Access Zone with name beginning with "**EyeglassRunbookRobot**" on both **source** and **target** clusters to be tested for DR. (Note: more than one pair of clusters can be tested with Runbook Robot).



2. Create a SyncIQ policy with the well known name "**EyeglassRunbookRobot-xxxx**" where *xxx* is a number or string of your choosing. Run the policy once it has been created.

- Use a path that is a **child of the access zone base path**. any path will do.
- See Basic setup above for detailed steps.

Job Name	Policy	Type	Last Run Date	State
Configuration Replication: Share, Export, Alias replication (AUTOMATIC)				
Cluster2-7201_EyeglassRunbookRobot-1	EyeglassRunboo...	AUTO	12/12/2015, 8:50:45 PM	OK
Name: Cluster2-7201_EyeglassRunbookRobot-1 Enabled/Disabled: ENABLED Job Type: AUTO Source: Cluster2-7201 Path: /ifs/data/robot/runbookrobotaccesszone Target: Cluster-1-7201 Path: /ifs/data/robot/runbookrobotaccesszone Last Success: 12/12/2015, 8:50:45 PM				
Cluster2-7201_EyeglassRunbook...	EyeglassRunb...	Cluster2-7201	Cluster-1-7201	OK
<u>Sync IQ Policy</u> OK Job Name: EyeglassRunbookRobot-1 Last Started: 12/12/2015, 7:03:10 PM Last Success: 12/12/2015, 7:03:10 PM Enabled: true <u>Eyeglass Configuration Replication</u> OK Job Name: Cluster2-7201_EyeglassRunbookRobot-1 Last Run: 12/12/2015, 8:50:45 PM Last Success: 12/12/2015, 8:50:45 PM Audit Status: AUDITSUCCEEDED Enabled: true				

3. Create a subnet pool (example Robotpool) and make it a member of the Access Zone created in Step #1, example "**EyeglassRunbookRobot**".

Note: the IP address space used should be reachable by Eyeglass to mount with NFS export. Create this pool on the source and target clusters.

1. **Example:** subnet0:Robotpool
 - Now create mapping alias for the Robot to move SmartConnect Zones from one pool to another. (see detailed section on hints for explanations in the Access Zone guide that describes syntax and how to make them unique to avoid spn collisions).

- **Example on Source cluster:** isi network modify pool --name=subnet0:Robotpool --add-zone-aliases=igls-01-prod
- **Example on target cluster:** isi network modify pool --name=subnet0:Robotpool --add-zone-aliases=igls-01-dr

NOTE: In this release, for the Access Zone Robot to show up in the Runbook Robot jobs, some configuration data needs to exist in the Access Zone. We recommend creating a share with no permissions anywhere in the Access Zone that is under the SyncIQ policy created within the Access Zone root path. Example “**Robotshare**”

5. After inventory runs (5 minutes default interval) you should see the Runbook Robot Jobs section showing the SyncIQ policy created in this section

Failover: Runbook Robot (AUTOMATIC)						
<input type="checkbox"/>	<input type="checkbox"/>	Cluster-1-7201_EyeglassRunbookRobot-001_runbook_ro...	...	RUNBOOK...	8/17/2015, 10:04:49 PM	Policy Disa...
<input type="checkbox"/>	<input type="checkbox"/>	Cluster-1-7201_EyeglassRunbookRobot-ZONEROBOT_r...	...	RUNBOOK...	8/17/2015, 9:57:14 PM	OK
<input type="checkbox"/>	<input type="checkbox"/>	Cluster2-7201_EyeglassRunbookRobot-001_mirror_runb...	...	RUNBOOK...	8/17/2015, 9:33:06 PM	Error
<input type="checkbox"/>	<input type="checkbox"/>	Cluster2-7201_EyeglassRunbookRobot-ZONEROBOT_mi...	...	RUNBOOK...	8/17/2015, 10:13:01 PM	Policy Disa...

6. To get the Zone Readiness updated (**default interval 6 hours**), run this job manually (See igls CLI in the [Eyeglass CLI](#) to change schedule of all jobs)
5. After the Zone Readiness job has completed, wait for the next Configuration Replication cycle to complete.

Job Name	...	Type	Last Run Date	Status
Cluster-1-7201_DFS-Demo-2-policy1_mirror	...	AUTO	8/18/2015, 7:15:59 AM	OK
Cluster-1-7201_EyeglassRunbookRobot-001	...	AUTO	8/17/2015, 10:01:08 PM	Policy Disa...
Cluster-1-7201_EyeglassRunbookRobot-ZONEROBOT	...	AUTO	8/18/2015, 7:15:59 AM	OK
Cluster2-7201_EyeglassRunbookRobot-001_mirror	...	AUTO	8/18/2015, 7:15:59 AM	OK
Cluster2-7201_DFS-Demo-2-policy1	...	AUTO	n/a	Policy Disa...
Cluster2-7201_EyeglassRunbookRobot-ZONEROBOT_mi...	...	AUTO	8/17/2015, 10:10:26 PM	Policy Disa...
Failover: Quota Failover (RUN MANUALLY)				
Failover Readiness (AUTOMATIC)				
Cluster-1-7201_Cluster2-7201	...	READINESS	8/18/2015, 2:00:03 AM	OK
Cluster2-7201_Cluster-1-7201	...	READINESS	8/18/2015, 2:00:03 AM	OK
Failover: Runbook Robot (AUTOMATIC)				
Cluster-1-7201_EyeglassRunbookRobot-001_runbook_ro...	...	RUNBOOK...	8/17/2015, 10:01:08 PM	OK
Cluster-1-7201_EyeglassRunbookRobot-ZONEROBOT_r...	...	RUNBOOK...	8/17/2015, 10:01:08 PM	OK
Cluster2-7201_EyeglassRunbookRobot-001_mirror_runb...	...	RUNBOOK...	8/17/2015, 10:10:26 PM	OK
Cluster2-7201_EyeglassRunbookRobot-ZONEROBOT_mi...	...	RUNBOOK...	8/17/2015, 10:10:26 PM	OK
Configuration Replication: Access Zone replication (AUTOMATIC)				

1 Item(s) selected

Action Buttons: Select a bulk action ▾ | Add New Job

8. Now open the DR Dashboard and select Zone Readiness to view the Robot Zone Readiness status

Source Cluster	Target Cluster	Zone Name	Last Success	Network Mapping	Overall Status ↑
Cluster2-7201	Cluster-1-7201	robot	Aug 18, 2015 1...	View Map	ERROR
Cluster-1-7201	Cluster2-7201	Robot	Aug 18, 2015 1...	View Map	ERROR
Cluster-1-7201	Cluster2-7201	System	Aug 18, 2015 1...	View Map	OK
Cluster2-7201	Cluster-1-7201	System	Aug 18, 2015 1...	View Map	OK

9. If you have errors related to the Robot zone, click on the link to view which areas are not correctly setup and correct the errors by reviewing documentation.

Zone Readiness for Cluster2-7201 > Cluster-1-7201 Zone: robot X

Name	Status
- Zone Readiness Statuses	ERROR
+ OneFS SyncIQ Readiness	OK
- Eyeglass Configuration Replication Readiness	ERROR
Cluster2-7201_EyeglassRunbookRob...	ERROR
- Smartconnect Zone Failover Mapping Readiness	ERROR
subnet0:robotpool	ERROR
- Zone Configuration Replication Readiness	ERROR
Cluster2-7201_EyeglassRunbookRob...	ERROR

Additional Status Information

Eyeglass Configuration Replication Readiness provides the status of the last Configuration Replication Job(s) for the policies in the Access Zone being failover. Status of OK indicates that the Job(s) completed successfully. Status of

10. To view the SmartConnect to pool mapping click the **View Map** link, correctly mapped with alias hints will look like the image below

Robot Access Zone with mapping example:

DR Dashboard

Policy Readiness	Source Cluster	Target Cluster	Zone Name	Last Success	Network Mapping	Overall Status	...
Zone Readiness	Cluster-1-7201	Cluster2-7201	System	Aug 18, 2015 1:...	View Map	OK	
DFS Readiness	Cluster-1-7201	Cluster2-7201	Robot	Aug 18, 2015 1:...	View Map	OK	

Network Mapping for Cluster-1-7201 > Cluster2-7201 Zone: Robot

Cluster-1-7201	Cluster2-7201
subnet0:robotpool Smart Connect Zone Name: robot-prod.ad1.test Smart Connect Aliases: - igls-mirror-Cluster2-7201.subnet0.robotpool Subnet: subnet0 SSIP: 172.31.1.200	subnet0:robotpool Smart Connect Zone Name: robot-dr.ad1.test Smart Connect Aliases: - igls-mirror-Cluster1-7201.subnet0.robotpool Subnet: subnet0 SSIP: 172.31.1.201

11. Confirm that your Robot zone is setup with all Zone Readiness Status indicators green (it needs to be all green)

Name	Status
Zone Readiness Statuses	OK
OneFS SyncIQ Readiness	OK
Eyeglass Configuration Replication Readiness	OK
Smartconnect Zone Failover Mapping Readiness	OK
Zone Configuration Replication Readiness	OK

Overall Status: OK

12. The default Robot jobs run every day at midnight and executes both Access Zone based failover Robot policies which moves all networking, SPN updates (delete, add) and aliases on subnet pools using the mapping hints each. It also executes any Runbook Robot policies in other Access Zones like system.

Robot jobs are configured without Continue on Error so initial validation check must not have any errors.

13. The default Robot jobs run every day at midnight and executes failover Robot policies which fails over the policy.

1. Creates export.

2. Mounts the cluster writes test data using export created.

3. It failovers the policy.
4. Runs the policy.
5. Deletes SPN's on source.
6. Renames SmartConnect alias on source with igls-original.
7. Creates SmartConnect alias on target based on mapping setup before hand.
8. Creates new SPN's against new cluster machine account.
9. Mounts the data on synced export on target.
10. Unmounts.
11. Moves the schedule on policy to the target.
12. Runs resync prep on source.
13. Goes to sleep until time to failback time.
14. Methods to verify it was successful
15. Check the DR Dashboard policies tab and verify it's green.
16. Check the cluster SyncIQ policy status on both clusters and make sure the policy moves from one cluster to the other each day.
17. Make sure the test file exists with a current date and time stamp on the active cluster (**Hint**; look in the policy path root file system for the test file).

18. If quota was applied to Robot policy path make sure the quota moved to the target cluster AND was deleted on the source cluster (Eyeglass moves policies on failover).
19. Review the Failover Log (DR Assistant / Failover History / Open log file). Failover log may also be downloaded from the Failover Log Viewer using the Download File link.
20. Check the DR Dashboard Zone Readiness Screen to make sure the Robot runs successfully and shows Failover or Green depending on which cluster the policies are active.
21. **NOTE:** *Initially after failover, the Zone Readiness will show an error for Eyeglass Configuration Replication and Zone Configuration Replication until the next Configuration Replication task and Zone Readiness task have completed.*
22. You can also check the SmartConnect Zone alias on each cluster and look for igls-original prefix on SmartConnect Zone name failed over.

The screenshot shows the configuration of a SmartConnect pool named "robotpool".

Basic settings

- IP range (low-high): 172.31.1.109 - 172.31.1.109
- Access zone: Robot

SmartConnect settings

- Zone name: igls-original-robot-dr.ad1.test
- Connection policy: Round Robin
- SmartConnect service subnet: subnet0
- IP allocation method: Static

Pool members

- Aggregation mode: Round-Robin Tx
- ext-1, Node 01 (172.31.1.109)

- 23.

24. The image above shows this clusters SmartConnect pool zone name was renamed on failover and indicates the dr.ad1.test zone has been moved to the production cluster based on the alias hints mapping. Cool!
25. The other way to check status is see which cluster has the enabled SyncIQ policy for the Robot zone

© Superna Inc

4.5. Advanced Settings

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

Advanced Settings

How to Change the Robot Scheduled interval

See Eyeglass Administration Guide [igls adv failover timeout](#) section

Manual Export Create for Runbook Robot

In some cases, creating the NFS export used by the Robot should be disabled on each run of the Robot job and allow manual export creation to be done.

Follow these steps **only when directed by support:**

1. Open Eyeglass shell from main menu.
2. Enter command “igls adv runbookrobot set --createExport=false”.
3. This will require manual export create on the Robot policy path with root client set to the IP address of the Eyeglass appliance.
4. Configuration sync will sync the export once created on the cluster with the enabled policy

Multiple Robot Feature Support (Hub and Spoke)

Testing multi-site replication topology:

A -> B

A -> C

set it up as follows:

cluster A

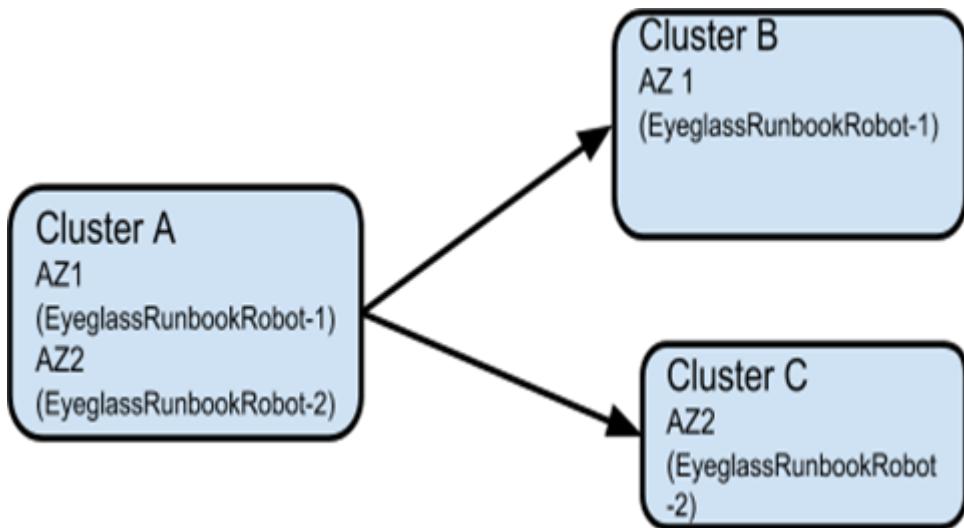
policy 1 = source path 1 to cluster B

policy 2 = source path 2 to cluster C

Topology	Access Zone path	Access Zone name	Policy source path	Policy target path
A → B	ifs/AZ1/	EyeglassRunbookRobot -1	Clstr A: ifs/AZ1/P1	Clstr B: ifs/AZ1/P1
A → C	ifs/AZ2/	EyeglassRunbookRobot -2	Clstr A: ifs/AZ2/P2	Clstr C; ifs/AZ2/P2

Hints pool mapping setup:

	Source	Target
A → B	igls-01-prod	igls-01-dr
A → C	igls-02-prod	igls-02-dr



Multiple Robot Feature Support (Chain)

Testing multi-site replication topology:

A -> B

B -> C

set up as follows:

cluster A

policy 1 = source path 1 to cluster B ,

cluster B

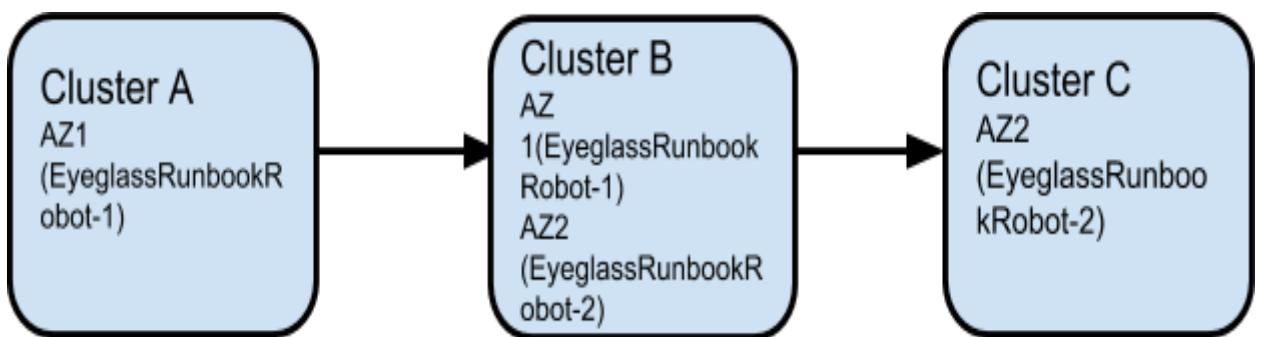
policy 2 = source path 2 to cluster C ,

Topology	Access Zone path	Access zone name	Policy source path	Policy target path
A → B	ifs/AZ1/	EyeglassRunbookRobot-1	Clsr A: ifs/AZ1/P1	Clsr B: ifs/AZ1/P1

B → C	ifs/AZ2/	EyeglassRunbookRobot -2	Clstr B: ifs/AZ2/P2	Clstr C: ifs/AZ2/P2
-------	----------	----------------------------	------------------------	------------------------

Hints pool mapping setup:

	Source	Target
A → B	igls-01-prod	igls-01-dr
B → C	igls-02-prod	igls-02-dr



Run Multi site Runbook job

1. Further details for prerequisite step before execute Runbook Robot DR
Automation can be found in [prerequisite](#) section in this document.
2. Create subnet pool (Robot-pool) for Runbook Access Zone, set mapping alias for the Robot pool to point SmartConnect Zones between source and target Robot-pool. More information for best practice creating Runbook Robot policy and other configuration (like, zone readiness and igs hints alias, etc can be found in above sections or in this [link](#) for Runbook Access Zone advance configuration.
3. Enable the Runbook job, Run the config replication policy job.
4. Run the Runbook Failover and check the running job to find the Failover steps

Multiple Robot Feature Support (Multiple Instances on Cluster Pairs)

Replicating in pairs:

Cluster A -> Cluster B

Cluster C -> Cluster D

1 eyeglass managing 4 clusters

Set it up as follows: Different Access Zone name and path on each pair

Summary of setup:

	Access Zone Path	SyncIQ Policy	Mapping Hint on Robot Subnet Pool
Cluster A	EyeglassRunbookRobot -1 /ifs/data/Robot	EyeglassRunbookRobot- Puru Source path: /ifs/data/Robot Target path: /ifs/data/Robot	igls-Robot-source8002
Cluster B	EyeglassRunbookRobot -1 /ifs/data/Robot		igls-Robot-target8002
Cluster C	EyeglassRunbookRobot	EyeglassRunbookRobot-	igls-Robot1-gbisi01

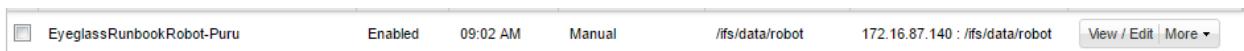
	-2 /ifs/data/Robot2	Robot2 Source path: /ifs/data/Robot2 Target path: /ifs/data/Robot2	
Cluster D	EyeglassRunbookRobot -2 /ifs/data/Robot2		igls-Robot1-gbisi02

Example:

cluster A (**Source-8002**)

policy 1 (**EyeglassRunbookRobot-Puru**) = source path:

/ifs/data/Robot to target path: **/ifs/data/Robot**



Cluster C (**SourceRobot7201**)

policy 1 (**EyeglassRunbookRobot-Robot2**) = source path:

/ifs/data/Robot2 to target path: **/ifs/data/Robot2**



Detailed configuration:

1. Create an Access Zone with name beginning with "EyeglassRunbookRobot" on all four clusters (i.e two source and two target cluster to be tested for DR.).

Note: Two source Access Zone name, basepath need to be unique and not overlapping.

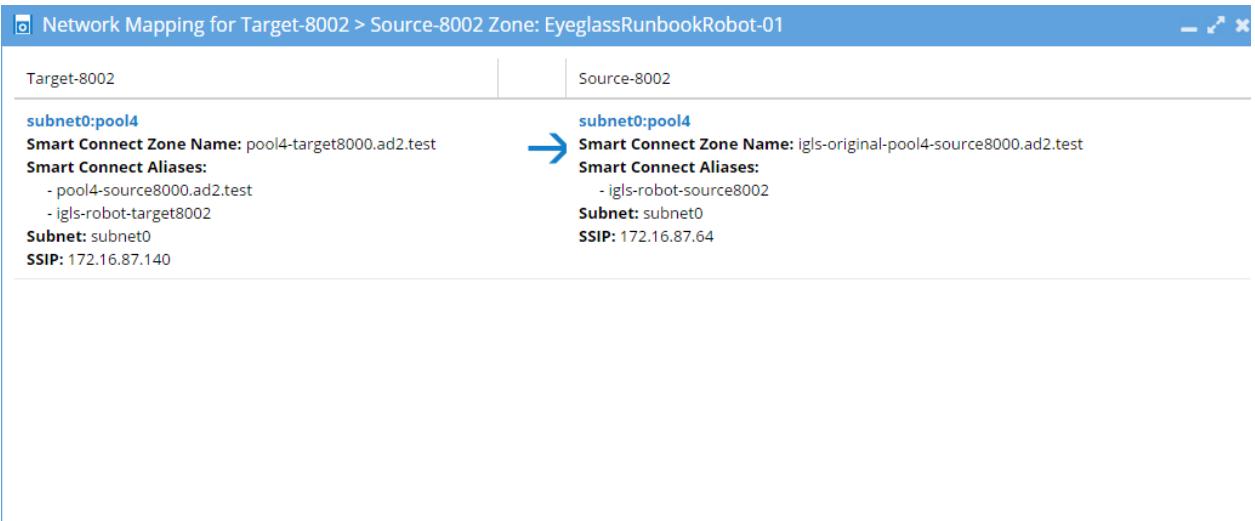
The screenshot shows the Eyeglass interface with two Access Zones listed. The first, 'EyeglassRunbookRobot-01', has a base path of '/ifs/data/robot' and is associated with 'groupnet0'. The second, 'EyeglassRunbookRobot-03', also has a base path of '/ifs/data/robot' and is associated with 'groupnet0'. A modal window is open for 'EyeglassRunbookRobot-03', displaying its details: Access Zone Name (EyeglassRunbookRobot-03), Zone Base Directory (/ifs/data/robot2), and Authentication Providers (Use all authentication providers). There are edit links for each of these fields. The modal also includes 'Hide details | Delete' and 'Close' buttons.

2. Create a SyncIQ policy on two source cluster See Basic setup above for detailed steps.
3. Create a subnet pool and make it a member of the Access Zone.

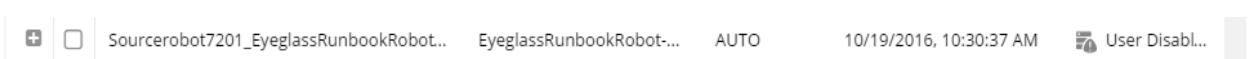
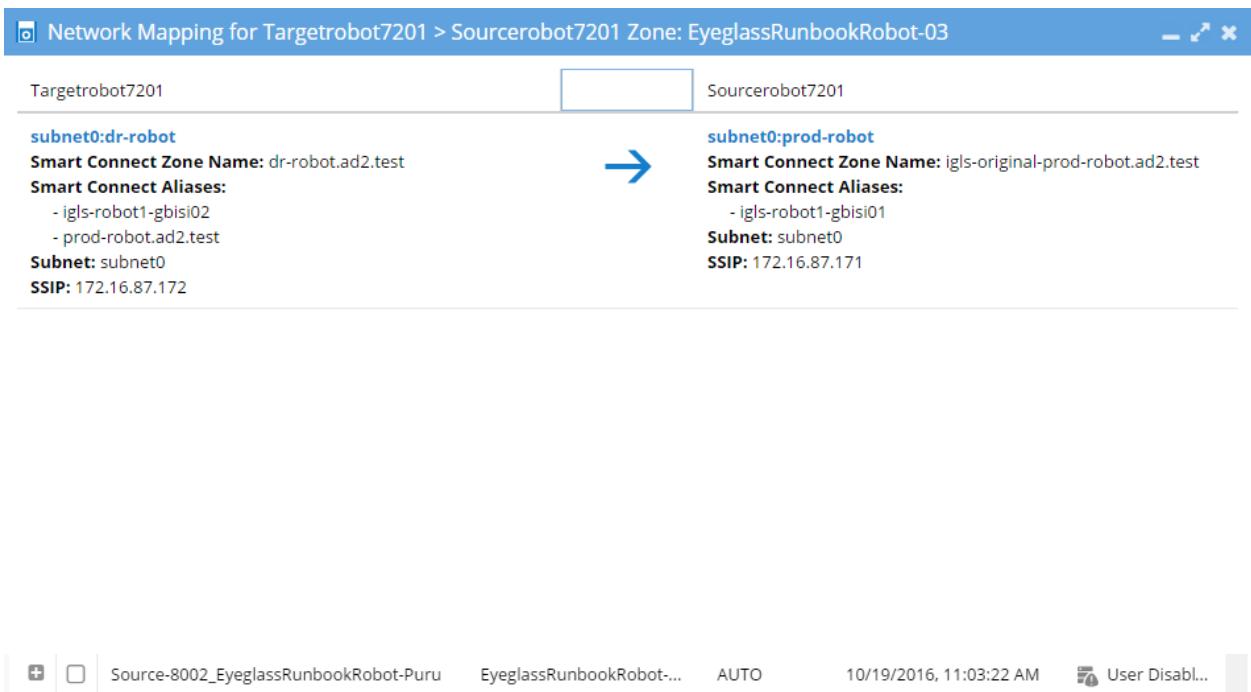
Note: the IP address used should be reachable by Eyeglass to mount with NFS export. Create this pool on all four clusters.

4. Dual DNS delegation needs to be done on both Source cluster Smartconnect Zone name and make sure it is resolving. (See [Geographic Highly Available Storage solution with Eyeglass Access Zone Failover and Dual Delegation](#))
5. Now create mapping alias for the Robot to move SmartConnect Zones from one pool to another. (See Eyeglass Access Zone Failover Guide "[Configure Eyeglass Subnet IP Pool Mapping Hints](#)" section.)
 1. Example on Source cluster A: isi network pools modify groupnet0.subnet0.pool4 --add-sc-dns-zone-aliases=igls-Robot-source8002

2. Example on Source cluster C: isi network modify pool --name=subnet0:prod-Robot After inventory runs (5 minutes default interval) you should see the Robot Jobs in "userdisabled" state therefore need to enable it and run it
3. --add-zone-aliases=igls-Robot1-gbisi01
4. Example on target cluster B: isi network pools modify groupnet0.subnet0.pool4 --add-sc-dns-zone-aliases=igls-Robot-target8002
5. Example on target cluster D: isi network modify pool --name=subnet0:dr-Robot --add-zone-aliases=igls-Robot1-gbisi02

A screenshot of a Network Mapping interface. The title bar says "Network Mapping for Target-8002 > Source-8002 Zone: EyeglassRunbookRobot-01". There are two main sections: "Target-8002" and "Source-8002". An arrow points from the Target section to the Source section.

Target-8002	Source-8002
subnet0:pool4 Smart Connect Zone Name: pool4-target8000.ad2.test Smart Connect Aliases: <ul style="list-style-type: none"> - pool4-source8000.ad2.test - igls-robot-target8002 Subnet: subnet0 SSIP: 172.16.87.140	subnet0:pool4 Smart Connect Zone Name: igls-original-pool4-source8000.ad2.test Smart Connect Aliases: <ul style="list-style-type: none"> - igls-robot-source8002 Subnet: subnet0 SSIP: 172.16.87.64



Jobs

Job Definitions	Job Name	Policy	Type	Last Run Date	State
Configuration Replication: Share, Export, Alias replication (AUTOMATIC)					
Running Jobs	+ Source-8002_Policy-accesszone	Policy-accesszone	AUTO	10/19/2016, 9:34:00 AM	OK
	+ Source-8002_PremPolicy	PremPolicy	AUTO	10/19/2016, 9:34:00 AM	OK
	+ Source-8002_Policy-policy	Policy-policy	AUTO	10/19/2016, 9:34:00 AM	OK
	+ Target-8002_Policy-policy_mirror	Policy-policy_mi...	AUTO	n/a	Policy Disa...
	+ Target-8002_Policy-accesszone_mirror	Policy-accesszo...	AUTO	n/a	Policy Disa...
	+ Sourcerobot7201_nyviewstore2	nyviewstore2	AUTO	n/a	Policy Disa...
	+ Sourcerobot7201_engg-saturn	engg-saturn	AUTO	10/19/2016, 9:34:00 AM	OK
	+ Sourcerobot7201_NYRADIO	NYRADIO	AUTO	n/a	Policy Disa...
	+ Sourcerobot7201_system	system	AUTO	10/19/2016, 9:34:00 AM	OK
	+ Sourcerobot7201_NYSUPPORT01	NYSUPPORT01	AUTO	10/19/2016, 9:34:00 AM	OK
	+ Sourcerobot7201_nycloudnas	nycloudnas	AUTO	10/19/2016, 9:34:00 AM	OK
	+ Sourcerobot7201_nybitsnas4	nybitsnas4	AUTO	10/19/2016, 9:34:00 AM	OK
	+ Sourcerobot7201_EyeglassRunbookR...	EyeglassRunboo...	AUTO	10/19/2016, 9:30:49 AM	Policy Disa...
	+ Sourcerobot7201_New-policy	New-policy	AUTO	10/19/2016, 9:34:00 AM	OK
	+ Targetrobot7201_NYRADIO_mirror	NYRADIO_mirror	AUTO	10/19/2016, 9:34:00 AM	OK
	+ Targetrobot7201_Systemdr	Systemdr	AUTO	10/19/2016, 9:34:00 AM	OK
	+ Targetrobot7201_nyviewstore2_mirror	nyviewstore2_m...	AUTO	10/19/2016, 9:34:00 AM	OK
	+ Targetrobot7201_EyeglassRunbookRo...	EyeglassRunboo...	AUTO	10/19/2016, 9:34:00 AM	OK
	+ Source-8002_EyeglassRunbookRobot-...	EyeglassRunboo...	AUTO	10/19/2016, 9:34:00 AM	Policy Disa...
	+ Target-8002_EyeglassRunbookRobot-...	EyeglassRunboo...	AUTO	10/19/2016, 9:00:56 AM	OK
Disaster Recovery Testing (AUTOMATIC)					
<input checked="" type="checkbox"/> Show Disabled Jobs		0 Item(s) selected		Select a bulk action ▾	Add New Job

6. To get the Zone Readiness updated See Advanced setup above for detailed steps

7. Now open the DR Dashboard and select Zone Readiness to view the Robot Zone Readiness status

The image displays two side-by-side screenshots of the DR Dashboard interface, showing the Zone Readiness status for two different cluster pairs.

Screenshot 1 (Top): This screenshot shows the Zone Readiness status for the Target-8002 > Source-8002 zone. The left panel lists various readiness checks, many of which are marked as "WARNING". The right panel shows the DR Dashboard with the "Zone Readiness" tab selected. It lists two entries: "Source-8002" and "Target-8002". The "Source-8002" entry shows "Last Successful Readiness Check: 10/17/2016, 1:11:3..." and "Network Mapping: FAILED OVER". The "Target-8002" entry shows "Last Successful Readiness Check: 10/17/2016, 1:11:3..." and "Network Mapping: WARNING". Below the dashboard, a "Network Mapping for Target-8002 > Source-8002 Zone: EyeglassRunbookRobot-01" window is open, showing network details for both the Target and Source clusters.

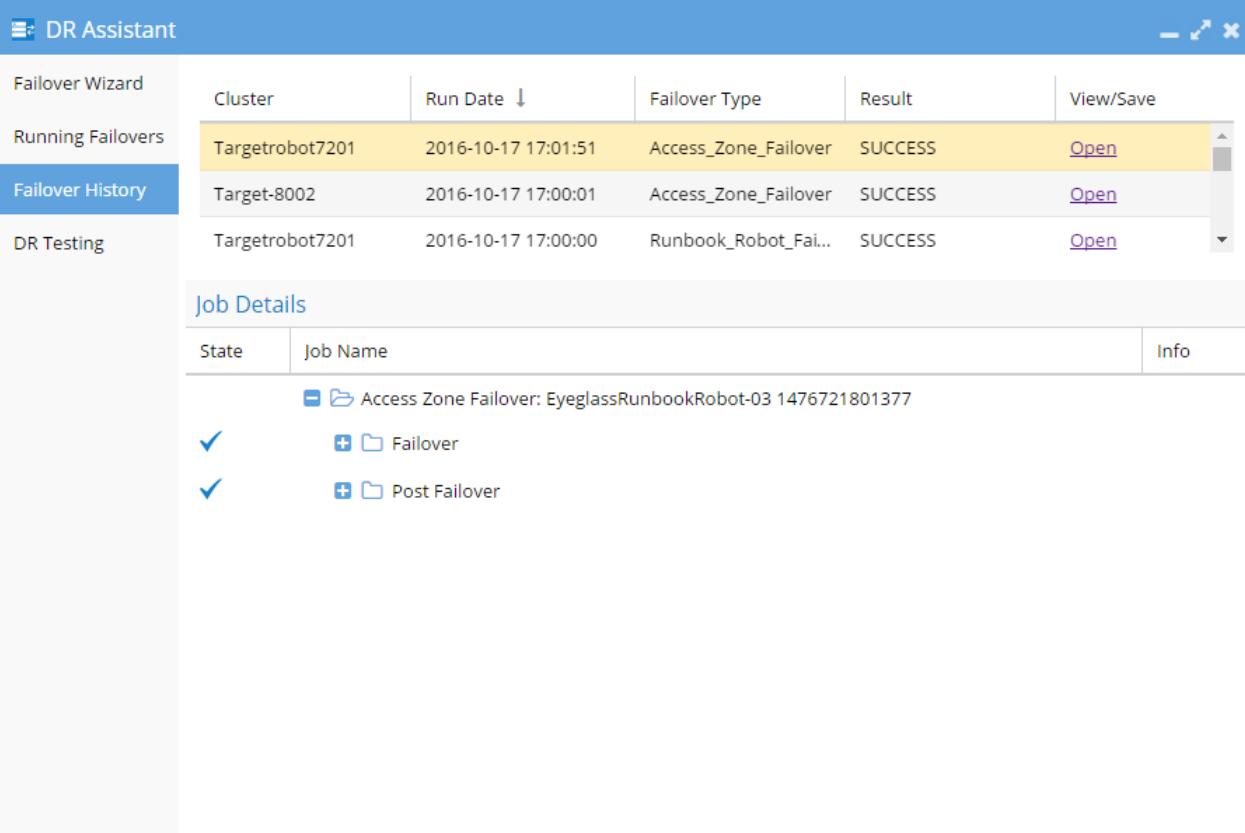
Screenshot 2 (Bottom): This screenshot shows the Zone Readiness status for the Sourcerobot7201 > Targetrobot7201 zone. The left panel lists readiness checks, with several marked as "WARNING". The right panel shows the DR Dashboard with the "Zone Readiness" tab selected. It lists two entries: "Sourcerobot7201" and "Targetrobot7201". The "Sourcerobot7201" entry shows "Last Successful Readiness Check: 10/17/2016, 1:11:3..." and "Network Mapping: ERROR". The "Targetrobot7201" entry shows "Last Successful Readiness Check: 10/17/2016, 1:11:3..." and "Network Mapping: ERROR". Below the dashboard, a "Network Mapping for Sourcerobot7201 > Targetrobot7201 Zone: EyeglassRunbookRobot-03" window is open, showing network details for both the Sourcerobot and Targetrobot clusters.

8. After verifying Zone readiness status, we have to run the Failover: Runbook Robot (AUTOMATIC) job to start failover

Failover: Runbook Robot (AUTOMATIC)						
[+]	<input type="checkbox"/>	Sourcerobot7201_EyeglassRunbookR...	EyeglassRunboo...	RUNBOOK...	10/19/2016, 9:33:19 AM	 Policy Disa...
[+]	<input type="checkbox"/>	Targetrobot7201_EyeglassRunbookRo...	EyeglassRunboo...	RUNBOOK...	10/19/2016, 9:06:13 AM	 OK
[+]	<input type="checkbox"/>	Source-8002_EyeglassRunbookRobot-...	EyeglassRunboo...	RUNBOOK...	10/19/2016, 9:36:06 AM	 Policy Disa...
[+]	<input type="checkbox"/>	Target-8002_EyeglassRunbookRobot-...	EyeglassRunboo...	RUNBOOK...	10/19/2016, 9:03:10 AM	 OK

9. After Failover you can check the SmartConnect Zone alias on each cluster and look for igls-original prefix on SmartConnect Zone name failed over. See Advanced setup above for detailed steps

10. Here is the Failover History



Failover Wizard					View/Save
Running Failovers	Cluster	Run Date	Failover Type	Result	
	Targetrobot7201	2016-10-17 17:01:51	Access_Zone_Failover	SUCCESS	Open
	Target-8002	2016-10-17 17:00:01	Access_Zone_Failover	SUCCESS	Open
	Targetrobot7201	2016-10-17 17:00:00	Runbook_Robot_Fai...	SUCCESS	Open

Job Details

State	Job Name	Info
✓	Access Zone Failover: EyeglassRunbookRobot-03 1476721801377	
✓	+ Failover	
✓	+ Post Failover	

11. Done

5. Eyeglass and PowerScale Compliance Mode Admin Guide

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

Overview

When Compliance Mode is enable on PowerScale cluster it prevents users from modifying or deleting files for compliance purposes.

Note: This guide applies to PowerScale Clusters with OneFS version 8.0.1.x and later.

Eyeglass appliance installation

Follow the installation guide until “Add PowerScale Cluster” section in the [install guide](#).

Use “**compmadmin**” user account to register the PowerScale Clusters on Eyeglass Appliance (Step 4).

Clusters with compliance mode active do not allow files owned by root to be modified. In regular installations “eyeglass” user account is used to register the PowerScale cluster and minimum privileges are required to set up, sudoers file cannot be edited to add eyeglass account so this is the reason the compadmin user in compliance mode is required to register the cluster.

1 Eyeglass Main Menu

2 Add Managed Device

3 SmartConnect Service IP:

4 Port: 8080

5 Username:

6 Password:

7 Maximum RPO Value:

Submit

Consideration replicating SmartLock directory:

You can create two types of SmartLock directories: enterprise and compliance.

Enterprise directories allow you to protect files without restricting the cluster with the regulation rule 17a-4. Complies directories are protected with the regulation rule 17a-4 and delete is not available.

Note: This guide is focused on compliance directories.

Limitations for source directory and target directory of a SyncIQ policy:

Source directory type	Target directory type	Allowed
Non-SmartLock	Non-SmartLock	Yes
Non-SmartLock	SmartLock enterprise	Yes
Non-SmartLock	SmartLock compliance	No
SmartLock enterprise	Non-SmartLock	Yes; however, retention dates and commit status of files will be lost.
SmartLock enterprise	SmartLock enterprise	Yes
SmartLock enterprise	SmartLock compliance	No
SmartLock compliance	Non-SmartLock	No
SmartLock compliance	SmartLock enterprise	No
SmartLock compliance	SmartLock compliance	Yes

Requirements for compliance directories before creating a syncIQ policy:

- The smartlock directory must be created on target prior to run the SyncIQ policy.
- If the compliance directory is not created on target the replication job will fail.
- Smartlock directory configuration settings are not replicated. If you change setting on source you must change them in target.
- The SyncIQ policy and SmartLock compliance directory must be configured at the same root directory level. A SmartLock compliance directory cannot be nested inside a SyncIQ policy.

Operations with compliance mode:

Set compliance clock

The compliance clock must be configured before creating SmartLock compliance directories. It is important to follow the EMC best practices when enabling a compliance mode. NTP configuration is recommended to set up on all nodes of source and target clusters to ensure all clock nodes are synchronized.

To set compliance mode follow these steps:

1. Start ssh session to PowerScale cluster and login with compadmin user account.
2. Set the compliance clock by running the following command

```
isi worm cdate set
```

3. Check if the compliance clock is running with the following command

```
isi worm cdate view
```

```
isi801-scr-1%  
isi801-scr-1% isi worm cdate view  
Compliance Clock: 2016-09-24T03:09:57  
isi801-scr-1%
```

Creation of SmartLock directory on Source

Once you create SmartLock directories you can commit files in that directory to WORM state.

1. Click File System > SmartLock > WORM.

2. Click Create Domain.

Define type (compliance) and directory path.

You can define the default retention period, minimum and/or maximum retention period. Also an autocommit time period.

Using CLI Command:

1. Open a secure shell (SSH) connection to any node in the cluster and log in.

2. Run the isi worm domains create command.

For Example:

The following command creates a compliance directory with a default retention period of 5 years, a minimum retention period of 4 years, a maximum retention period of 6 years, and an autocommit time period of 30 minutes.

```
isi801-scr-1%  
isi801-scr-1% isi worm domains create /ifs/picturezone/veggie --compliance --default-retention 5Y --min-retention 4Y --max-retention 6Y --autocommit-offset 30m  
You have 1 warnings:  
Once you specify a WORM domain as Compliance, it cannot be changed back.  
Are you sure? (yes/[no]): yes  
isi801-scr-1%  
isi801-scr-1%
```

View SmartLock directory settings:

Use this command to view SmartLock Directory settings:

1. Open a secure shell (SSH) connection to any node in the cluster and log in.
2. Run the isi worm domains list command.

```
isi801-scr-1%  
isi801-scr-1% isi worm domain list  
ID      Path          Type  
-----  
65536  /ifs/picturezone/fruit  compliance  
65538  /ifs/picturezone/veggie  compliance  
-----  
Total: 2  
isi801-scr-1%
```

To view details about the SmartLock directory perform isi worm domains view <domain-directory>

```
isi801-scr-1%  
isi801-scr-1% isi worm domain view /ifs/picturezone/veggie  
    ID: 65538  
    Path: /ifs/picturezone/veggie  
    Type: compliance  
    LIN: 4298702855  
Autocommit Offset: 30m  
    Override Date: -  
Privileged Delete: disabled  
Default Retention: 5Y  
    Min Retention: 4Y  
    Max Retention: 6Y  
isi801-scr-1%
```

View a file WORM status:

To check the WORM status of the file, perform the follow command:

1. Start ssh session to PowerScale cluster and login.
2. Check the WORM status of the file by running the following command

```
isi worm files view <file>
```

For example,

```
isi801-scr-1%  
isi801-scr-1% isi worm files view /ifs/picturezone/fruit/asia-region/file2.txt  
WORM Domains  
ID      Root Path  
-----  
65536  /ifs/picturezone/fruit  
  
WORM State: COMMITTED  
  Expires: 2016-10-05T01:22:26  
isi801-scr-1%
```

Failover and Failback operations with Eyeglass

According with your environment and requirements follow the Failover Configuration Guides on Superna documentation website.

Before failover or failback operations check the following documentation and [Failover Planning Guides](#):

Using the compadmin user account on Eyeglass failover and failback operations has the following known restrictions with system permissions and privileges.

6. Multi Site Failover Guide for Continuous Availability

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

- [Overview](#)
- [Logical Diagram of Multi Site Failover](#)
- [Access Zone Failover - SyncIQ Configuration for 3 site](#)
- [3 Site DFS Mode Failover](#)

© Superna Inc

6.1. Overview

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

- [Video How To - Overview Multi site Access Zone Failover](#)
- [Overview of Multi Site DR and Continuous Availability](#)
- [Overview of Multi Site Access Zone DR and Continuous Availability](#)
- [Pre-requisites](#)
- [Supported Access Zone Failover Operations](#)

Overview

The highest level of data protection comes from multi site replication, where a source clusters data is replicated to clusters at 2 different sites. Typically, 2 clusters are in metro location and the 3rd cluster is outside a power grid failure zone.

The goal of this failover design is to provide a choice of sites to failover and have fully automated failover between sites. In addition, fallback and failover again to the same or different site will be possible, but may require some manual steps to avoid SyncIQ policies that will block the failover.

In this configuration Eyeglass can be used with Access Zone failover solution to protect the data in an Access Zone and allow a

failover choice of target cluster 1 or target cluster 2 (at the 3rd site). This solution can operate at the Access Zone level and allows one or more Access Zones to be 3 site protected and other Access Zones only 2 site protected.

Video How To - Overview Multi site Access Zone Failover

Overview of Multi Site DR and Continuous Availability

This solution will support Access Zone fully automated and DFS mode multi site failover and fallback. The document covers Access Zone and DFS mode in separate sections.

Overview of Multi Site Access Zone DR and Continuous Availability

This solution allows one or more independent Access Zones to have 2 possible replication targets and sites to failover. This offers maximum data protection and full automation for DNS mount path failover for SMB and NFS.

Pre-requisites

1. DNS also requires triple Name Server Delegation records.

This is similar to Dual delegation where the SmartConnect Zones involved in the 3 site failover must now have 3 NS records pointing to all Subnet Service IP's for each cluster and subnet involved in the failover

2. Igls-hints provisioned on all 3 clusters for each Access Zone must be configured
 3. Recommended to apply on all target clusters the SmartConnect zone placeholder using igls-original-<SmartConnect Zone name on source cluster>
 4. Replicate the same data to site B and C to keep it simple using overlapping source path policies that replicate to cluster B and C with A being the source cluster
 5. Ensure IGLS-hints are globally unique example this avoids SPN collection issues on AD machine accounts but allows pools to be matched on the igls-xxxx-
1. A cluster - igls-poolx-Cluster-A
 2. B cluster - igls-poolx-Cluster-B
 3. C cluster - igls-poolx-Cluster-C
6. Use the same AD provider in all Access Zones
 7. Ensure all clusters are in the same AD Forest
8. **NOTE** When following AD Delegation you MUST extend the delegation to the 3rd cluster and make sure each cluster is given Write Service principal Name permissions to its own Computer Account AND the other 2 cluster machine accounts.

Supported Access Zone Failover Operations

- A to B or A to C

- If B is Active Cluster:
 - A to C Access Zone failover not supported
 - A to C per SyncIQ Policy failover supported (manual steps for networking)
- If C is Active Cluster:
 - A to B Access Zone failover not supported
 - A to B per SyncIQ Policy failover supported (manual steps for networking)

© Superna Inc

6.2. Logical Diagram of Multi Site Failover

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

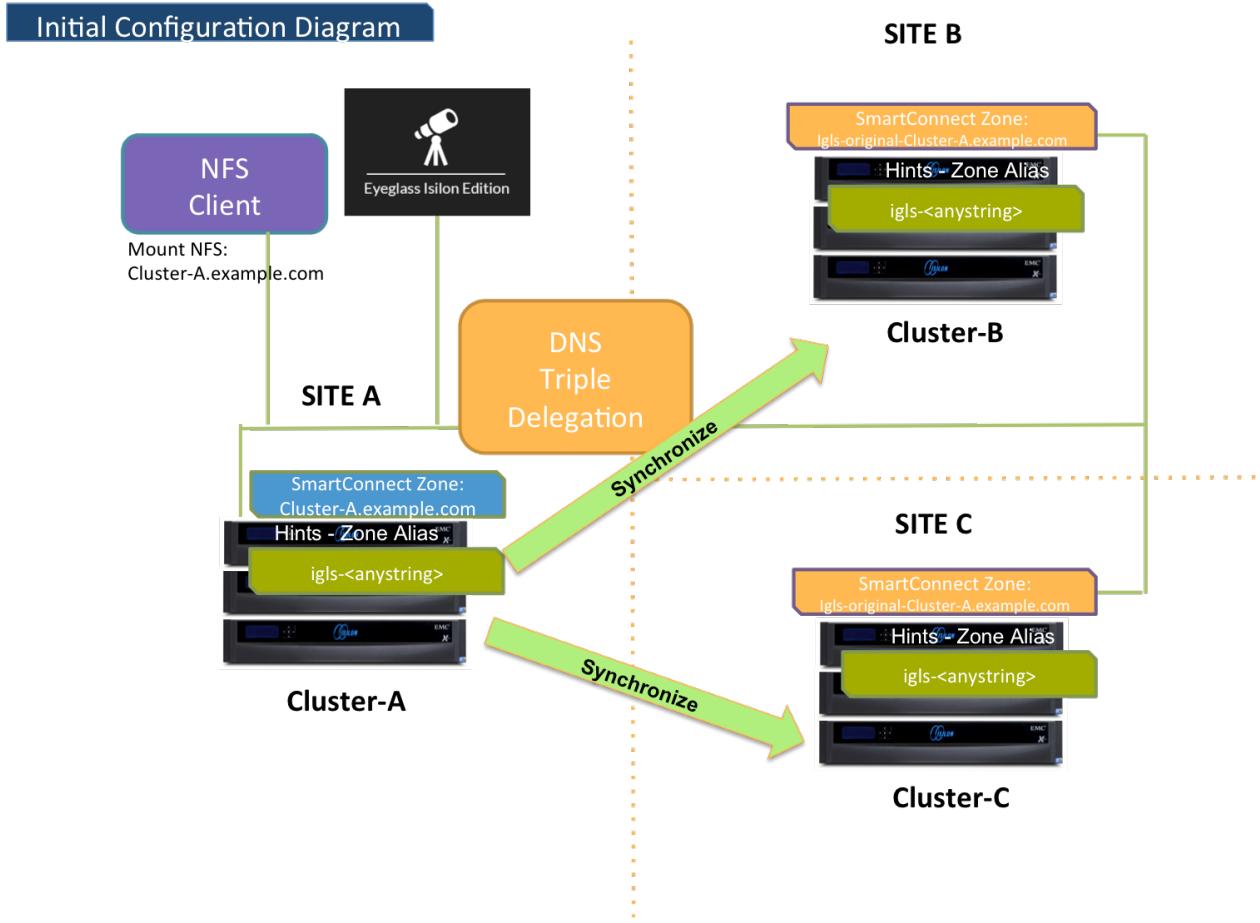
Logical Diagram of Multi Site Failover

- [Initial Configuration / Before Failover Diagram](#)
- [Failover A ⇒ B](#)
- [Fallback B ⇒ A](#)
- [Failover A ⇒ C](#)
- [Fallback C ⇒ A](#)
- [Eyeglass Access Zone Failover Steps](#)
- [Eyeglass Access Zone Fallback Steps](#)

Initial Configuration / Before Failover Diagram

The following diagram displays the initial configuration / before failover for 3-sites Failover.

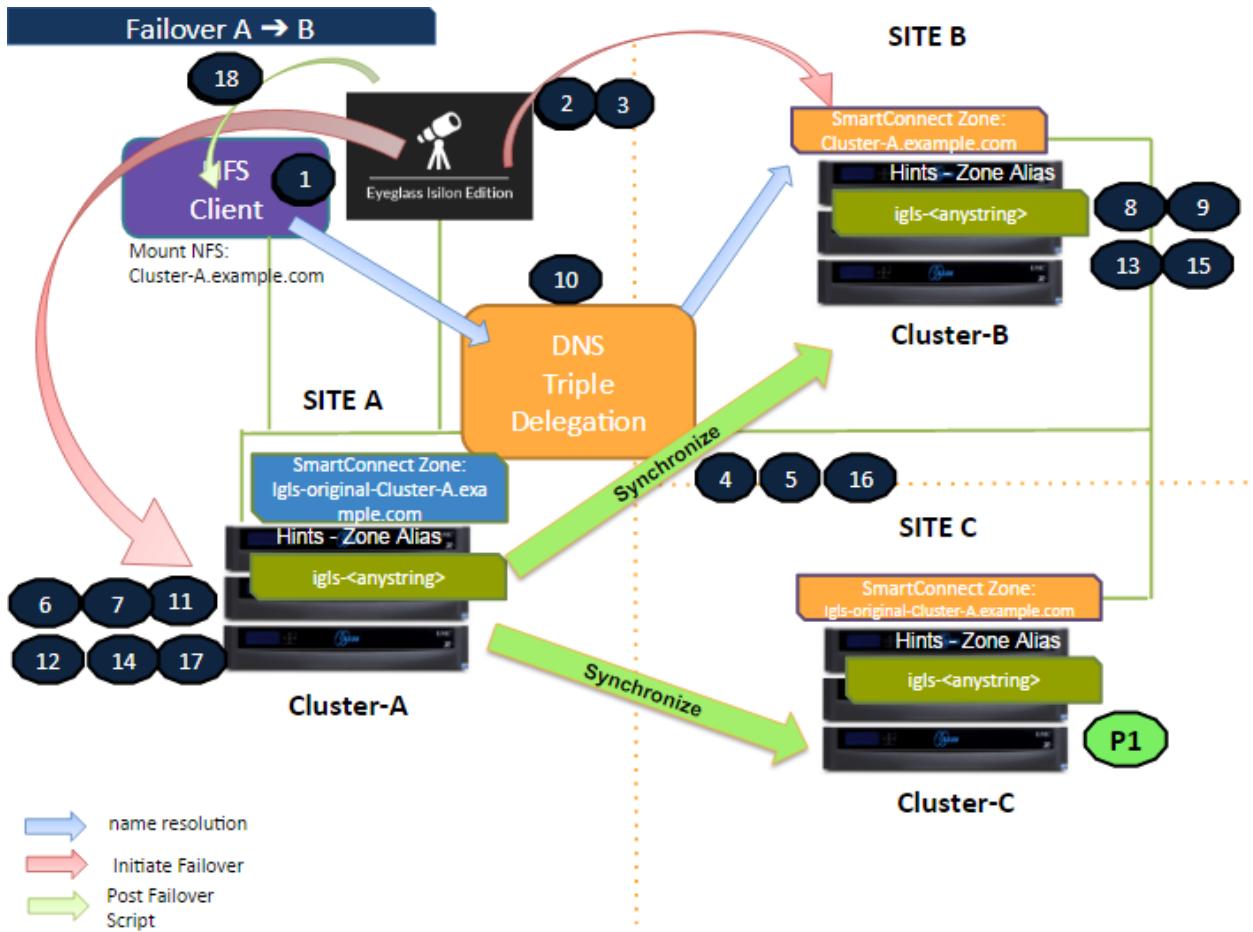
- Site A: Primary Site
- Site B: Secondary Site #1
- Site C: Secondary Site #2



Failover A \Rightarrow B

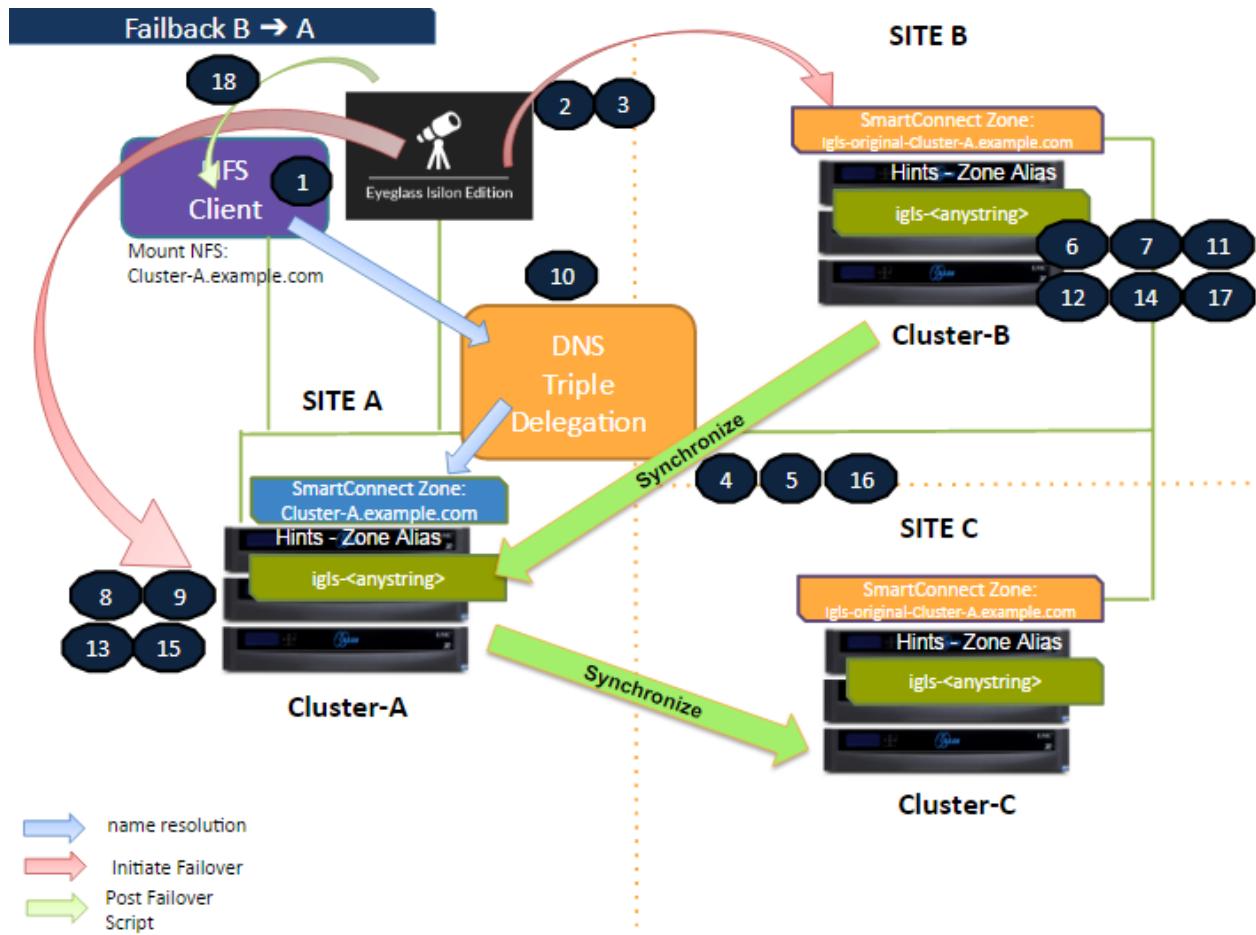
The following diagram illustrates the workflow for failover from A to B. Take note step P1 (Preparation Step - prior to initiate Eyeglass Access Zone Failover) - refer to the [procedure section](#) for details.

Refer to [this table](#) for the list of the numbered steps shown in this diagram.



Failback B ⇒ A

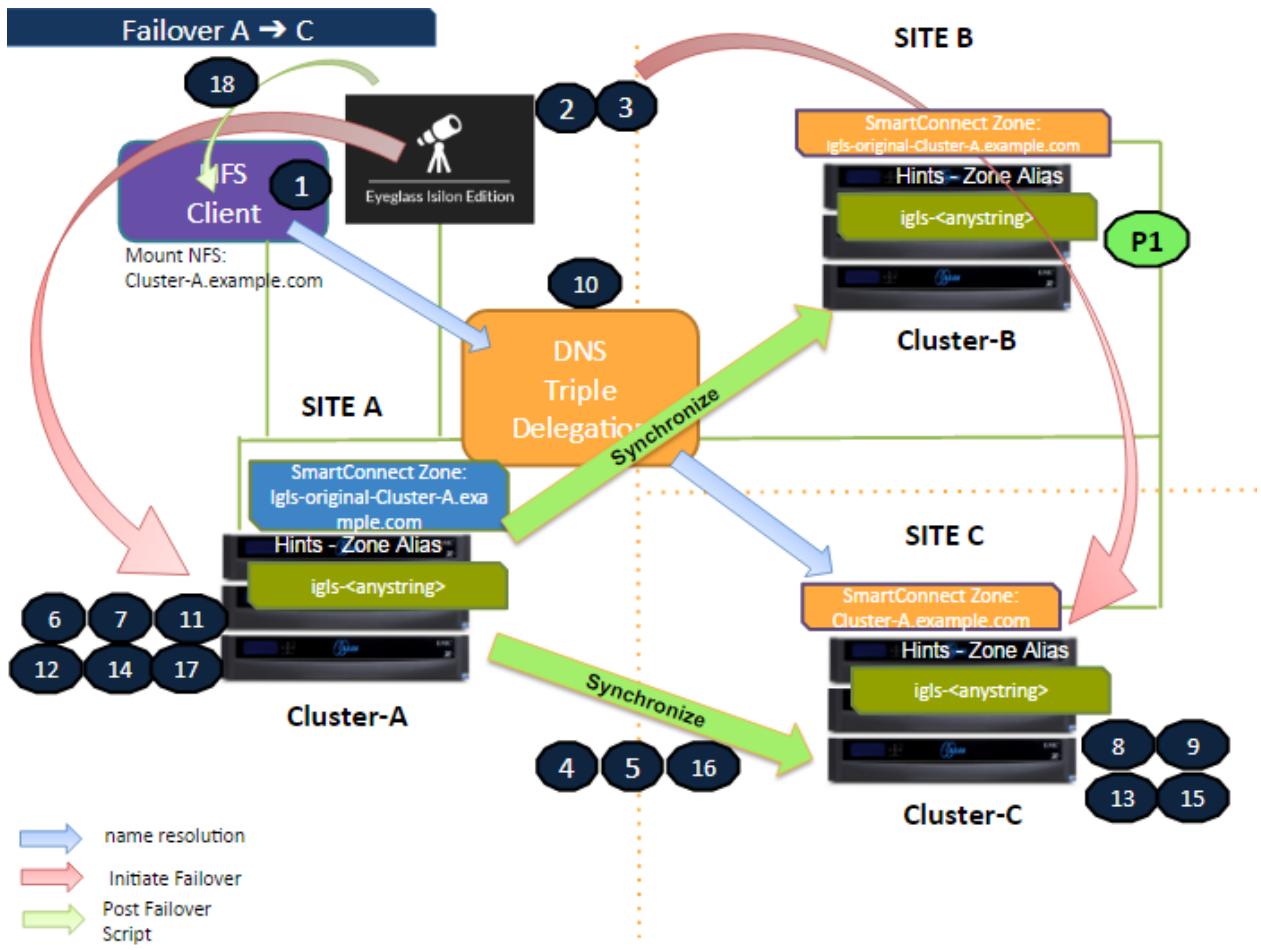
The following diagram illustrates the workflow for failback from B to A. Refer to [this table](#) for the list of the numbered steps shown in this diagram.



Failover A ⇒ C

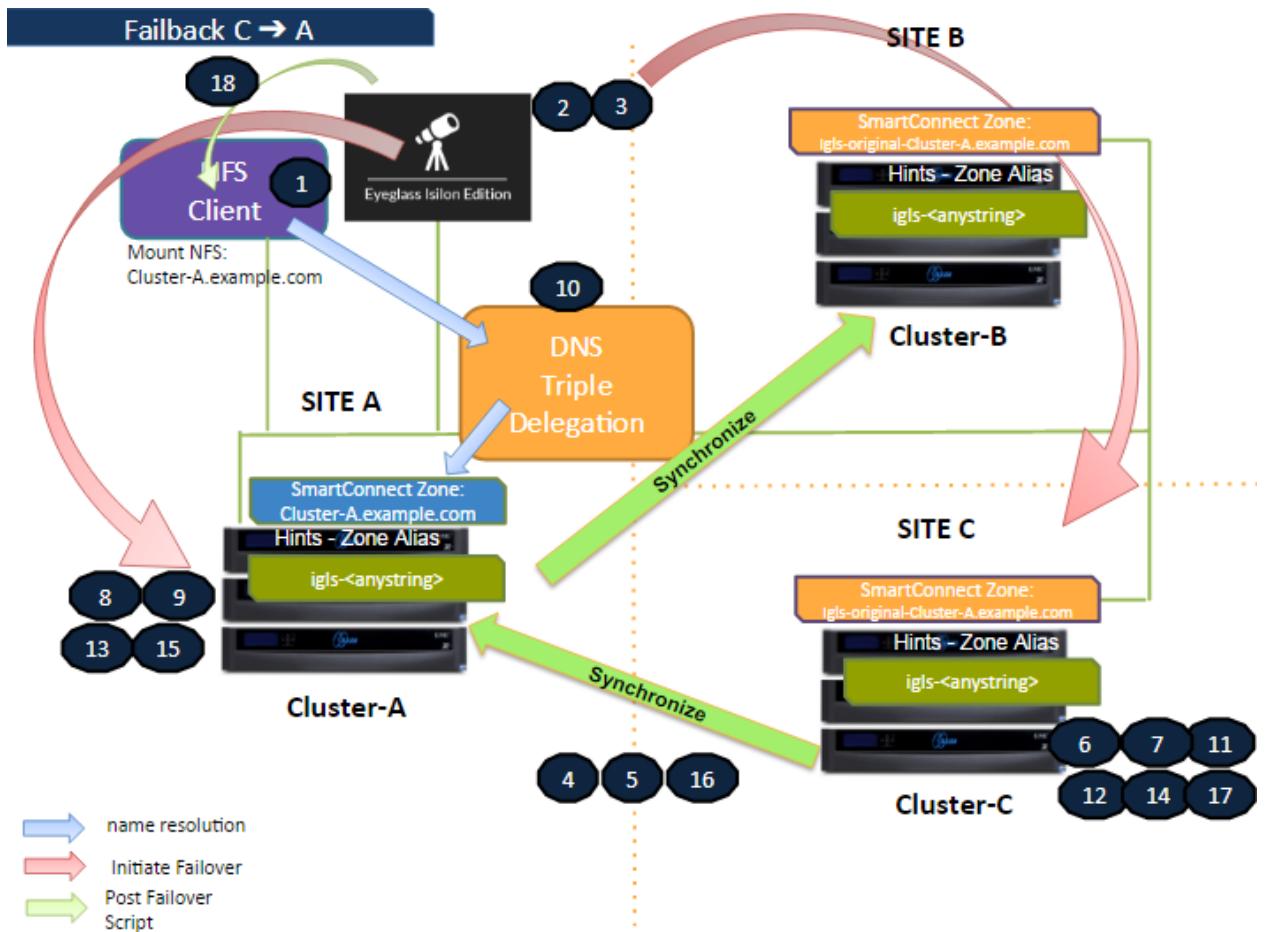
The following diagram illustrates the workflow for failover from A to C. Take note step P1 (Preparation Step - prior to initiate Eyeglass Access Zone Failover) - refer to the [procedure section](#) for details.

Refer to [this table](#) for the list of the numbered steps shown in this diagram.



Failback C ⇒ A

The following diagram illustrates the workflow for failback from C to A. Refer to [this table](#) for the list of the numbered steps shown in this diagram.



Eyeglass Access Zone Failover Steps

This table lists the Eyeglass Access Zone Failover steps with numbers as shown in the above Failover diagrams.

No	Step
P1	<p>Preparation.</p> <ul style="list-style-type: none"> • Failover A ⇒ B: Ensure there is no existing Mirror Policies between C to A. If there is existing Mirror Policies between C to A, delete first, before initiate Failover from A to B. • Failover A ⇒ C: Ensure there is no existing Mirror Policies between B to A. If there is existing Mirror Policies between B to A, delete first, before initiate Failover from A to C.

1	Ensure that there is no live access to data
2	Begin Failover
3	Validation
4	Synchronize data
5	Synchronize configuration (shares/export/alias)
6	Change SmartConnect Zone on Source so not to resolve by Clients
7	Avoid SPN Collision
8	Move SmartConnect Zone to Target
9	Update SPN to allow for authentication against target
10	Repoint DNS to the Target Cluster - DNS Triple Delegation
11	Record schedule for SyncIQ policies being failed over
12	Prevent SyncIQ policies being failed over from running
13	Provide write access to data on target
14	Disable SyncIQ on source and make active on target
15	Set proper SyncIQ schedule on target
16	Synchronize quota(s)
17	Remove quotas on directories that are target of SyncIQ (PowerScale best practice)
18	Refresh session to pick up DNS change (use post failover script)

Eyeglass Access Zone Failback Steps

This table lists the Eyeglass Access Zone Failback steps with numbers as shown in the above Failback diagrams.

No	Step
1	Ensure that there is no live access to data
2	Begin Failback
3	Validation
4	Synchronize data

5	Synchronize configuration (shares/export/alias)
6	Change SmartConnect Zone on Source (Secondary Cluster) so not to resolve by Clients
7	Avoid SPN Collision
8	Move SmartConnect zone to Target (Primary Cluster)
9	Update SPN to allow for authentication against target
10	Repoint DNS to the Target Cluster - DNS Triple Delegation
11	Record schedule for SyncIQ policies being failed back
12	Prevent SyncIQ policies being failed back from running
13	Provide write access to data on target
14	Disable SyncIQ on source (Secondary Cluster) and make active on target (Primary Cluster)
15	Set proper SyncIQ schedule on target (Primary Cluster)
16	Synchronize quota(s)
17	Remove quotas on directories that are target of SyncIQ (on Secondary Cluster) (PowerScale best practice)
18	Refresh session to pick up DNS change (use post failover script)

© Superna Inc

6.3. Access Zone Failover - SyncIQ Configuration for 3 site

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

Access Zone Failover - SyncIQ Configuration for 3 site

- SyncIQ Policy multi Site support Matrix
- Multi Site Failover and Failback Behaviours
- Zone Readiness
- Zone Readiness - Initial Configuration / before Failover / after Failback
- Before Failback B \Rightarrow A
- Before Failback C \Rightarrow A
- Summary of Network Mappings
- Access Zone Failover and Failback Procedures
- Failover from A to B Procedure:
- Failback from B to A Procedure:
- Failover from A to C Procedure:
- Failback from C to A Procedure

The source cluster SyncIQ policy path must fall inside the Access Zone for Target cluster 1, a Second SyncIQ policy can use the same source path (**recommended**). See the configuration guide below.

Not Recommended

Examples where failover leaves some data behind after failover and moves networking (SmartConnect Zones) to Target cluster.

After Failover **both** SmartConnect zones would failover **together** leaving the namespace and data stranded since the SyncIQ policies only failed over a portion of the data.

SyncIQ Policy multi Site support Matrix

Access Zone Path for these examples is /ifs/data/AZ1

Unsupported configuration

	Source Cluster	Target Cluster 1	Target Cluster 2
SyncIQ policy 1	/ifs/data/AZ1/data	x	
SyncIQ policy 2	/ifs/data/AZ1/marketing		x
SmartConnect Zone 1	data.example.com	x	x
SmartConnect Zone 2	marketing.example.com	x	x

A supported configuration requires that all data and all DNS name space fails over together to achieve fully automated Access Zone failover.

Recommended Policy Configuration

	Source Cluster	Target Cluster 1	Target Cluster 2
SyncIQ policy 1	/ifs/data/AZ1/data	x	
SyncIQ policy 2	/ifs/data/AZ1/marketing		x
SyncIQ policy 3	/ifs/data/AZ1/data		x
SyncIQ policy 4	/ifs/data/AZ1/marketing	x	
SmartConnect Zone 1	data.example.com	x	x
SmartConnect Zone 2	marketing.example.com	x	x

Multi Site Failover and Failback Behaviours

Operation	Direction	Supported	Require Manual Step Prior to Initiate Eyeglass Access Zone Failover
Failover	A ⇒ B	Yes	Yes - refer to this diagram and Access Zone Failover - SyncIQ Configuration for 3 site
Failback	B ⇒ A	Yes	No - refer to this diagram and Access Zone Failover - SyncIQ Configuration for 3 site
Failover	A ⇒ C	Yes	Yes - refer to this diagram and Access Zone Failover - SyncIQ Configuration for 3 site
Failback	C ⇒ A	Yes	No - refer to this diagram and Access Zone Failover - SyncIQ Configuration for 3 site

Zone Readiness

This section gives example of the Zone Readiness status and Network Mapping between Source-Target#1 and Source-Target#2 pairs.

For the purpose of this example we use the following names:

Site	Cluster Name
A (Source)	cluster20

B (Target#1)	cluster21
C (Target#2)	cluster31

There are two Access Zones on all those three clusters: zone01 and zone03

Zone Readiness - Initial Configuration / before Failover / after Failback

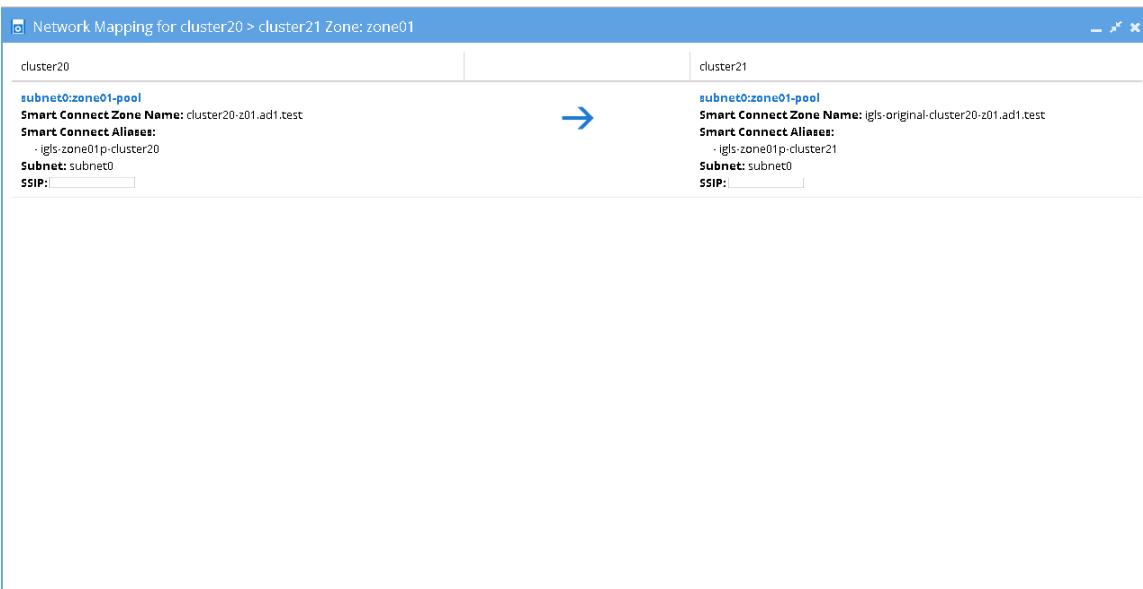
This is the Zone Readiness for Initial Configuration / before Failover / after failback state. As we can see from this figure, that both Source-Target Pairs (A - B and A - C) are listed in this DR Dashboard' zone readiness window.

This shows that a failover choice can be made to any target cluster in Green OK state (Warning status also allowed).

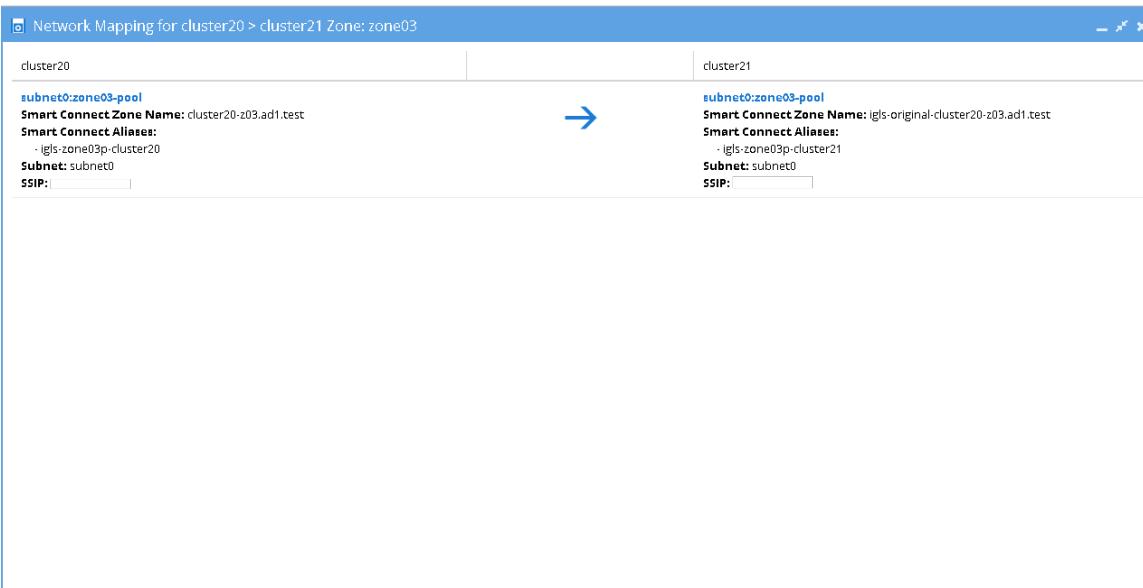
The screenshot shows the DR Dashboard interface with the 'Zone Readiness' tab selected. The table displays the following data:

	Source Cluster	Target Cluster	Zone Name	Last Successful Readiness Check	Network Mapping	Overall Status
DFS Readiness	cluster20	cluster21	zone03	6/15/2016, 11:48:17 P...	View Map	OK
DR Testing (Beta)	cluster20	cluster21	zone01	6/15/2016, 11:48:17 P...	View Map	OK
	cluster20	cluster31	zone01	6/15/2016, 11:48:17 P...	View Map	OK
	cluster20	cluster31	zone03	6/15/2016, 11:48:17 P...	View Map	OK

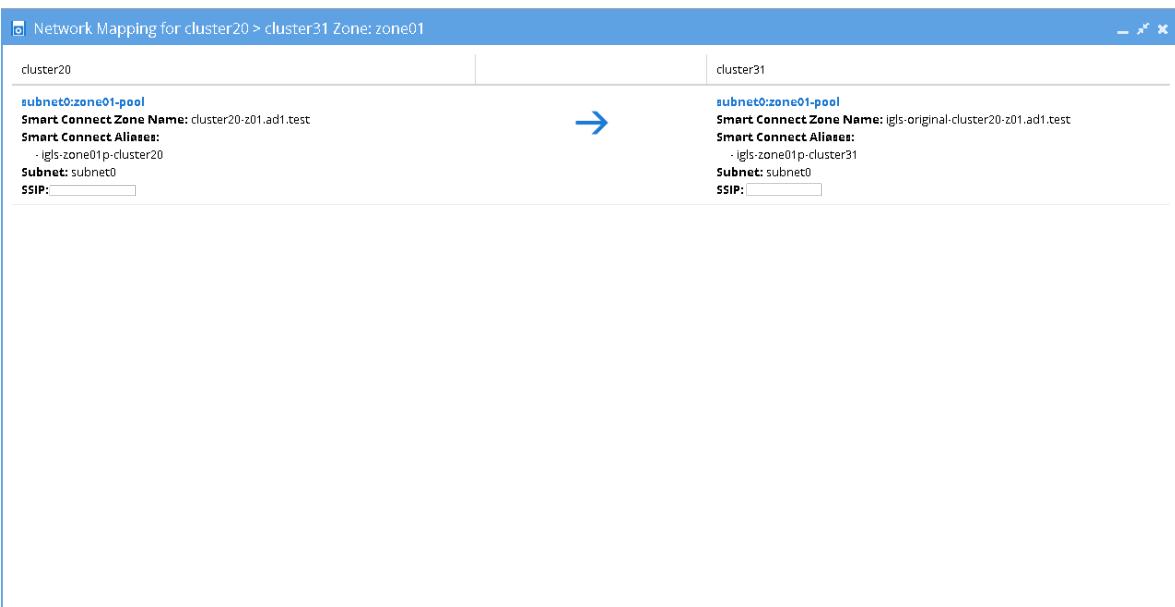
Network Mapping - Initial Configuration: A \Rightarrow B Zone01



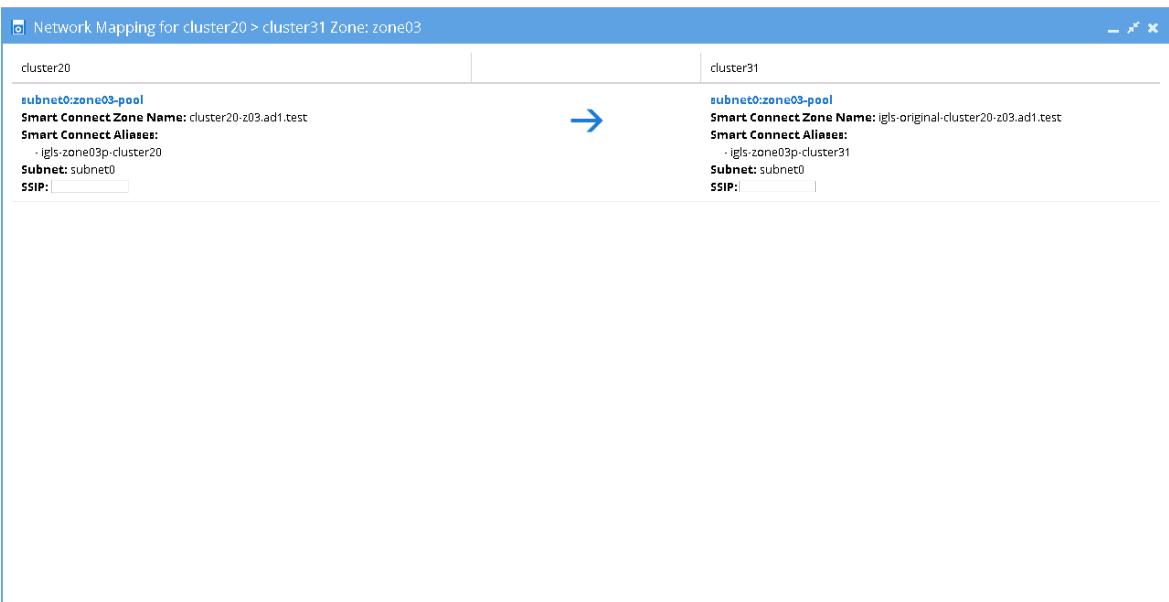
Network Mapping - Initial Configuration: A \Rightarrow B Zone03



Network Mapping - Initial Configuration: A \Rightarrow C Zone01



Network Mapping - Initial Configuration: A ⇒ C Zone03



This table shows the SmartConnect Zone Name and SmartConnect Alias Name mappings for this Initial Configuration / before failover / after fallback states:

Source - Target Pair	SyncIQ direction	Zone Name	SmartConnect Zone Name		SmartConnect Alias Mapping	
			Source Cluster	Target Cluster	Source Cluster	Target Cluster

cluster20 - cluster21	A \Rightarrow B	zone01	cluster20-z01.ad1.test	igls-original-cluster20-z01.ad1.test	igls-zone01p-cluster20	igls-zone01p-cluster21
		zone03	cluster20-z03.ad1.test	igls-original-cluster20-z03.ad1.test	igls-zone03p-cluster20	igls-zone03p-cluster21
cluster 20 - cluster31	A \Rightarrow C	zone01	cluster20-z01.ad1.test	igls-original-cluster20-z01.ad1.test	igls-zone01p-cluster20	igls-zone01p-cluster31
		zone03	cluster20-z03.ad1.test	igls-original-cluster20-z03.ad1.test	igls-zone03p-cluster20	igls-zone03p-cluster31

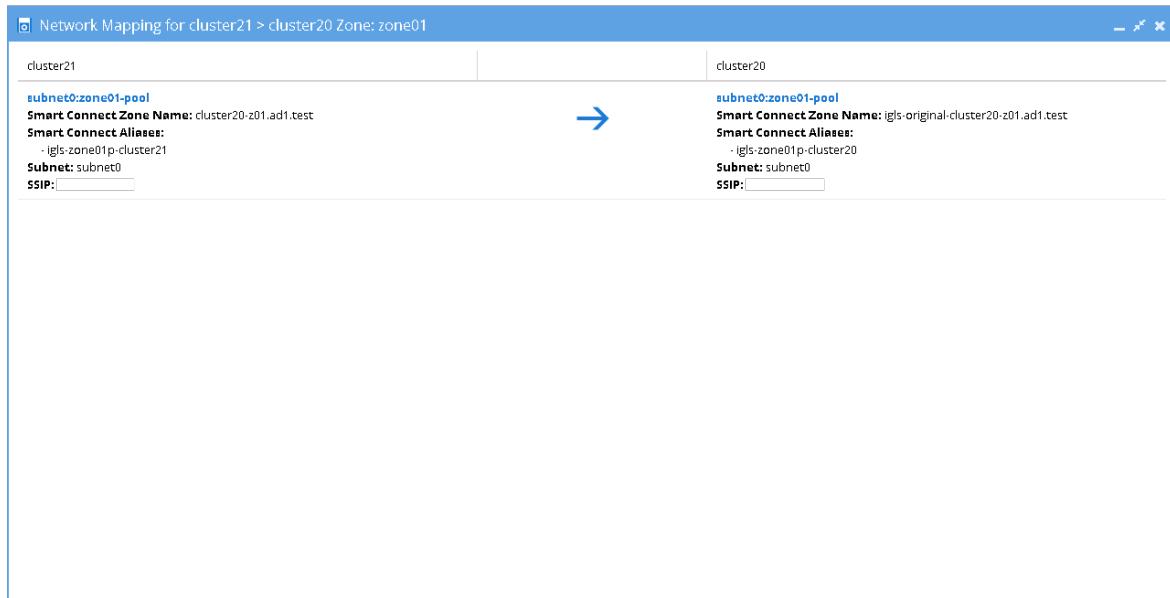
Before Failback B \Rightarrow A

Zone Readiness

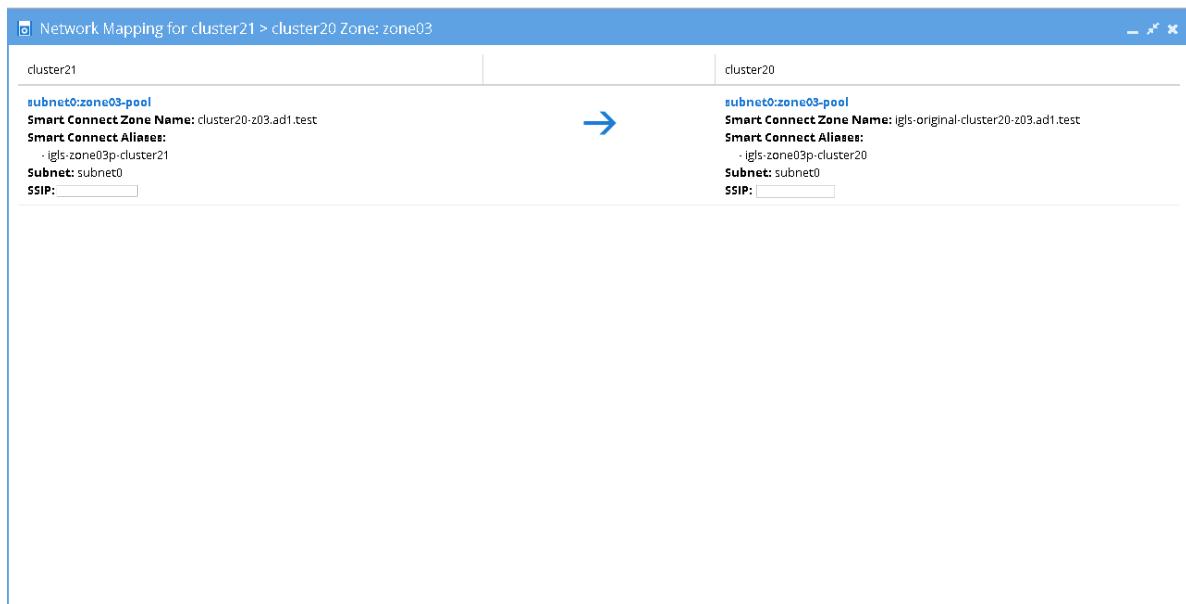
This Zone Readiness is for the state before Failback from B to A. As we can see from this figure, that only TargetB(cluster21)-SourceA(cluster20) pairs are listed as available in this DR Dashboard's zone readiness window. The other pairs (SourceA(cluster20)-TargetB(cluster21) and SourceA(cluster20)-TargetC(cluster31)) are stated as FAILED-OVER.

DR Dashboard						
Policy Readiness	Source Cluster	Target Cluster	Zone Name	Last Successful Readiness Check	Network Mapping	Overall Status
Zone Readiness	cluster20	cluster21	zone03	6/12/2016, 10:44:50 P...	View Map	FAILED-OVER
DFS Readiness	cluster20	cluster21	zone01	6/12/2016, 10:44:50 P...	View Map	FAILED-OVER
DR Testing (Beta)	cluster20	cluster31	zone03	6/12/2016, 10:44:48 P...	View Map	FAILED-OVER
	cluster20	cluster31	zone01	6/12/2016, 10:44:48 P...	View Map	FAILED-OVER
	cluster21	cluster20	zone03	6/12/2016, 10:44:48 P...	View Map	OK
	cluster21	cluster20	zone01	6/12/2016, 10:44:48 P...	View Map	OK

[Network Mapping - Before Failback B ⇒ A] B ⇒ A Zone01



[Network Mapping - Before Failback B ⇒ A] B ⇒ A Zone03



This table shows the SmartConnect Zone Name and SmartConnect Alias Name mappings for this Before Failback B ⇒ A state:

Source -	SynclIQ	Zone	SmartConnect Zone Name	SmartConnect Alias Mappings
----------	---------	------	------------------------	-----------------------------

Target Pair	direction	Name	Source Cluster	Target Cluster	Source Cluster	Target Cluster
cluster20 - cluster21	A \Rightarrow B	zone01	STATUS: FAILED OVER			
		zone03	STATUS: FAILED OVER			
cluster20 - cluster31	A \Rightarrow C	zone01	STATUS: FAILED OVER			
		zone03	STATUS: FAILED OVER			
cluster21 - cluster20	B \Rightarrow A	zone01	cluster20-z01.ad1.test	igls-original-cluster20-z01.ad1.test	igls-zone01p-cluster21	igls-zone01p-cluster20
		zone03	cluster20-z03.ad1.test	igls-original-cluster20-z03.ad1.test	igls-zone03p-cluster21	igls-zone03p-cluster20

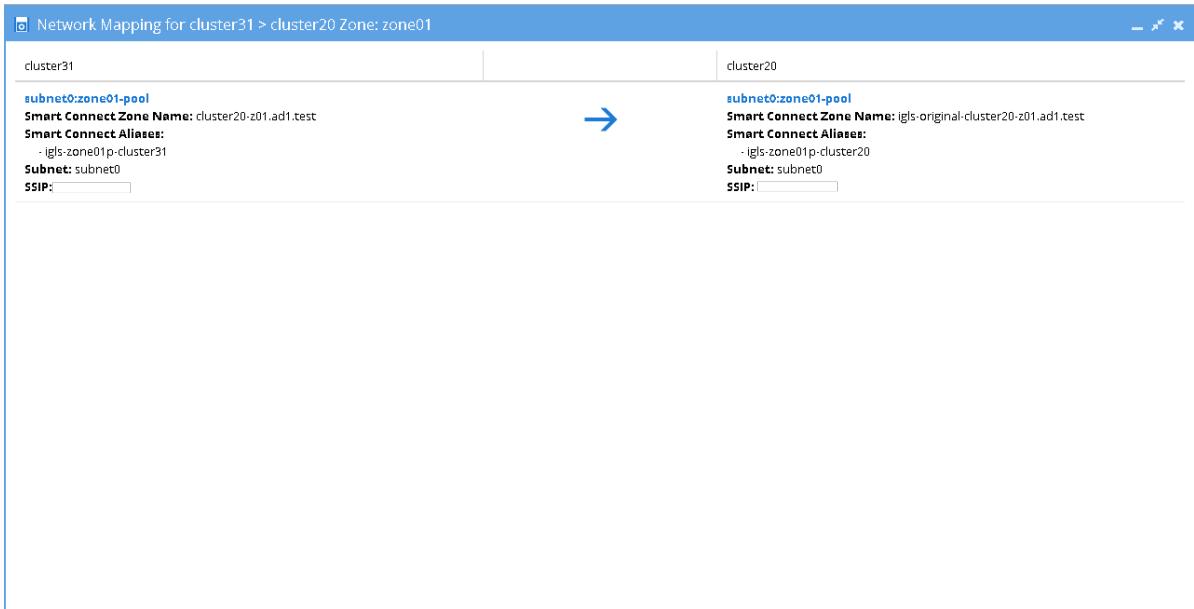
Before Failback C \Rightarrow A

Zone Readiness

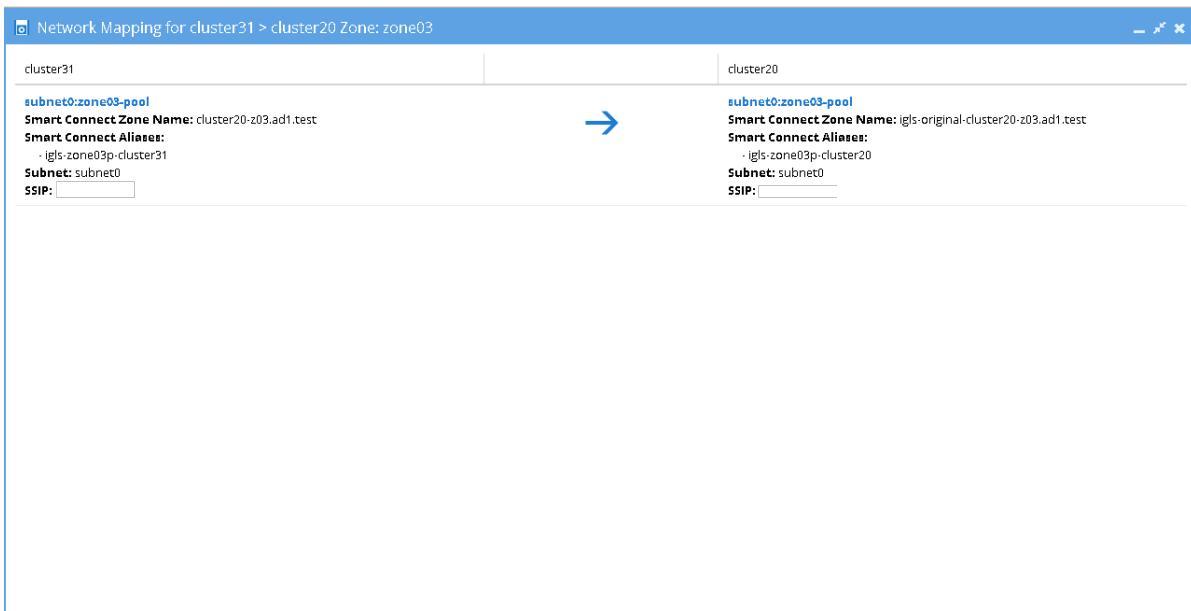
This Zone Readiness is for the state before Failback from C to A. As we can see from this figure, that only TargetC(cluster31)-SourceA(cluster20) pairs are listed as available in this DR Dashboard's zone readiness window. The other pairs (SourceA(cluster20)-TargetB(cluster21) and SourceA(cluster20)-TargetC(cluster31)) are stated as FAILED-OVER.

DR Dashboard						
Policy Readiness	Source Cluster	Target Cluster	Zone Name	Last Successful Readiness Check	Network Mapping	Overall Status
Zone Readiness	cluster20	cluster21	zone03	6/12/2016, 11:53:15 P...	View Map	FAILED OVER
DFS Readiness	cluster20	cluster21	zone01	6/12/2016, 11:53:15 P...	View Map	FAILED OVER
DR Testing (Beta)	cluster20	cluster31	zone01	6/12/2016, 11:53:16 P...	View Map	FAILED OVER
	cluster20	cluster31	zone03	6/12/2016, 11:53:16 P...	View Map	FAILED OVER
	cluster31	cluster20	zone01	6/12/2016, 11:53:15 P...	View Map	OK
	cluster31	cluster20	zone03	6/12/2016, 11:53:15 P...	View Map	OK

[Network Mapping - Before Failback C ⇒ A] C ⇒ A Zone01



[Network Mapping - Before Failback C ⇒ A] C ⇒ A Zone03



This table shows the SmartConnect Zone Name and SmartConnect Zone Alias Name mappings for Before Failback C ⇒ A state:

Source - Target Pair	SyncIQ direction	Zone Name	SmartConnect Zone Name		SmartConnect Alias Mappings	
			Source Cluster	Target Cluster	Source Cluster	Target Cluster
cluster20 - cluster21	A ⇒ B	zone01	STATUS: FAILED OVER			
		zone03	STATUS: FAILED OVER			
cluster20 - cluster31	A ⇒ C	zone01	STATUS: FAILED OVER			
		zone03	STATUS: FAILED OVER			
cluster31 - cluster20	C ⇒ A	zone01	cluster20-z01.ad1.test	igls-original-cluster20-z01.ad1.test	igls-zone01p-cluster31	igls-zone01p-cluster20
		zone03	cluster20-z03.ad1.test	igls-original-cluster20-z03.ad1.test	igls-zone03p-cluster31	igls-zone03p-cluster20

Summary of Network Mappings

Based on the above example, the following table summarizes the network mappings with zone names and zone alias names:

Initial Configuration / Before Failover / After Failback:

State	Access Zone	Name	Cluster20 (A)	Cluster21 (B)	Cluster31 (C)
Initial Config	zone01	Zone Name	cluster20-z01.ad1.test	igls-original-cluster20-z01.ad1.test	igls-original-cluster20-z01.ad1.test
		Zone Alias Hint	igls-zone01p-cluster20	igls-zone01p-cluster21	igls-zone01p-cluster31
	zone03	Zone Name	cluster20-z03.ad1.test	igls-original-cluster20-z03.ad1.test	igls-original-cluster20-z03.ad1.test
		Zone Alias Hint	igls-zone03p-cluster20	igls-zone03p-cluster21	igls-zone03p-cluster31

After Failover

This table shows the zone names and zone alias names after failover A \Rightarrow B / after failover A \Rightarrow C:

State	Access Zone	Name	Cluster20 (A)	Cluster21 (B)	Cluster31 (C)
After Failover A \Rightarrow B	zone01	Zone Name	igls-original-cluster20-z01.ad1.test	cluster20-z01.ad1.test	igls-original-cluster20-z01.ad1.test
		Zone Alias	igls-zone01p-cluster20	igls-zone01p-cluster21	igls-zone01p-cluster31
	zone03	Zone Name	igls-original-cluster20-z03.ad1.test	cluster20-z03.ad1.test	igls-original-cluster20-z03.ad1.test
		Zone Alias	igls-zone03p-cluster20	igls-zone03p-cluster21	igls-zone03p-cluster31
After Failover A \Rightarrow C	zone01	Zone Name	igls-original-cluster20-z01.ad1.test	igls-original-cluster20-z01.ad1.test	cluster20-z01.ad1.test
		Zone Alias	igls-zone01p-cluster20	igls-zone01p-cluster21	igls-zone01p-cluster31
	zone03	Zone Name	igls-original-cluster20-z03.ad1.test	igls-original-cluster20-z03.ad1.test	cluster20-z03.ad1.test
		Zone Alias	igls-zone03p-cluster20	igls-zone03p-cluster21	igls-zone03p-cluster31

Access Zone Failover and Failback Procedures

It is recommended to create SyncIQ Policies that will be used for multi site replications (e.g. to replicate from Site A to Site B and also from Site A to Site C) with names that reflect the Source-Target pairs.

The following table is an example for 2 SyncIQ Policies per Source-Target pairs:

SyncIQ Policy Name	SyncIQ Pairs
AB-synciq-01 AB-synciq-02	A and B
AC-synciq-01 AC-synciq-02	A and C

This name format will help us to identify which Access Zones that we want to failover.

Failover from A to B Procedure:

1. Prior to initiating Eyeglass Access Zone Failover from A to B, **we need to ensure that there is no existing SyncIQ Mirror Policies from C to A. The recovery resync prep step of this Failover A to B will create Mirror Policies from B to A with same Mirror Target Paths as the C to A (Mirror Target Paths are overlaps). This will make the Mirror Policies from B to A unrunnable and the Eyeglass Failover Job will fail.** If there are existing

ones, we need to delete them first. Refer to step P1 in the Failover workflow diagrams.

- a. NOTE: The above step MUST be completed before A to C failover, the order matters since the domain mark will be deleted on cluster A once the step above is completed.
2. Now run a domain mark job on each SyncIQ policy on cluster A. This is a required step since no domain mark exists and will be created during failover process from A to B. The best practise is to run domain mark before failover to ensure the resync prep step does not take a long time to complete. Domain mark can run longer if a the path has a large number of files. NOTE: During the time while domain mark job is running no failover from A to C should be executed until the domain mark job completes on ALL sync polices involved in the failover. Monitor progress from the PowerScale Cluster jobs UI.
3. Then we can perform Eyeglass Access Zone Failover as per normal. In DR Assistant Wizard, after we selected the source cluster (Cluster A (for this example: name cluster20)) the next wizard screen display the list of available Failover options based on Source-Target-Zone pairs.

Failover Wizard	Source Cluster	Target Cluster	Zone Name	Last Successful Readiness Check	Network Mapping	Overall Status
Running Failovers	<input type="checkbox"/> cluster20	cluster21	zone03	6/9/2016, 3:25:3...	View Map	OK
Failover History	<input type="checkbox"/> cluster20	cluster21	zone01	6/9/2016, 3:25:3...	View Map	OK
DR Testing (Beta)	<input type="checkbox"/> cluster20	cluster31	zone03	6/9/2016, 3:25:3...	View Map	OK
	<input type="checkbox"/> cluster20	cluster31	zone01	6/9/2016, 3:25:3...	View Map	OK

A to B

A to C

[Back](#) [Next](#)

3. We need to be careful to select the correct Target Cluster that we want to Failover (A to B or A to C). For this case we want to failover from A to B. Select a zone that we want to Failover from cluster20 (source) - cluster21 (target) pairs.
4. The next screen will give warning to highlight that this wizard will only perform Access Zone failover from A to B. The other policy on the same Access Zone (A to C) will **not** be failed over.

The screenshot shows the DR Assistant interface with the 'Failover Wizard' selected in the sidebar. A prominent yellow warning icon with the word 'WARNING!' is displayed. The text below it states: 'Eyglass has determined that your configuration is valid. Disabled Policies were found on the zone01 zone and will not be failed over.' It also lists specific policies: 'The following policies will be failed over: cluster20_AB-synciq-01' and 'The following policies will NOT be failed over: cluster20_AC-synciq-01'. A section titled 'Best Practices' provides several recommendations:

- Run domain mark manually in advance to ensure a fast failover. See [this url](#) for details.
- If scheduled SyncIQ policies start running during a failover, it can cause failover job to fail. If your SyncIQ schedule has a risk of starting during a planned failover, then the SyncIQ schedule should be set to manual before initiating failover.
- 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.

At the bottom right of the page are 'Back' and 'Next' buttons.

5. Proceed this Access Zone Failover as per normal. Refer to Eyeglass Access Zone Failover Guide for details.
6. Repeat the same procedure for failover other Access Zones from A to B.

Fallback from B to A Procedure:

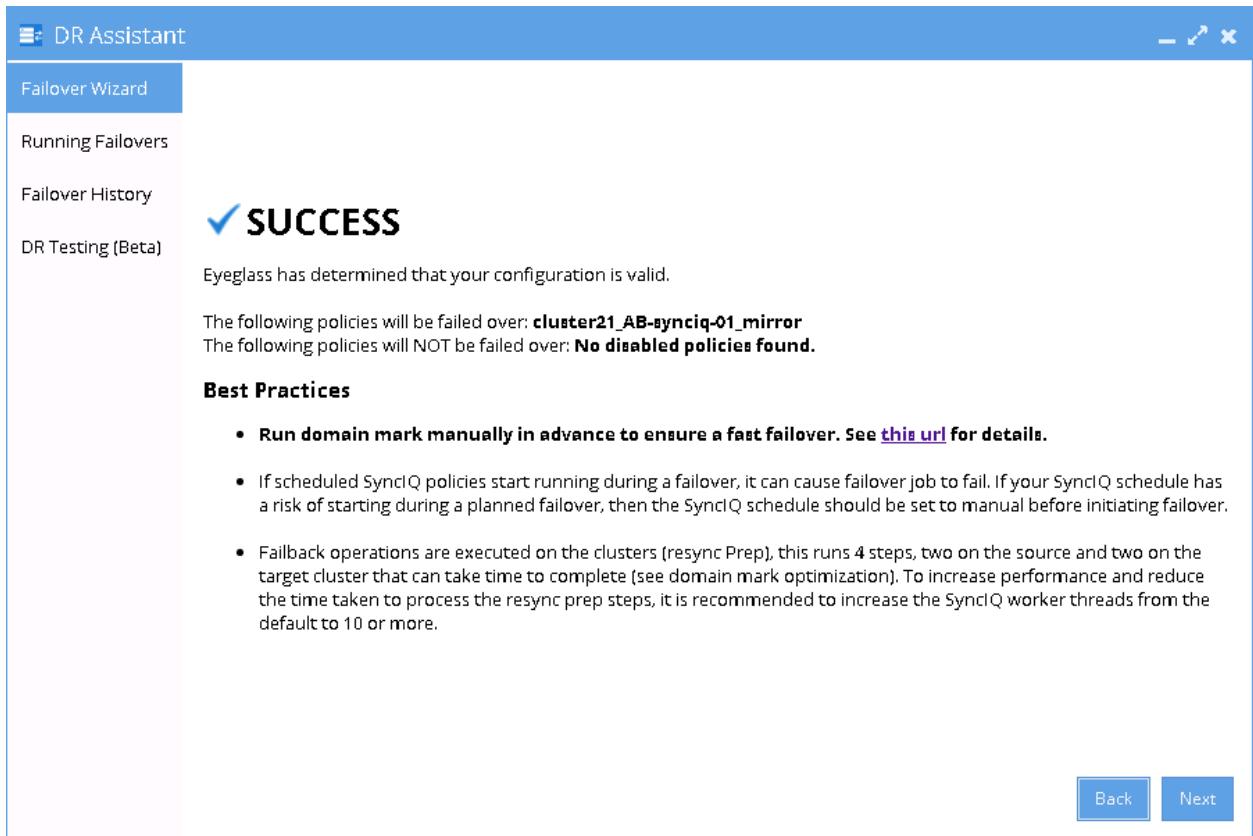
1. We can perform Eyeglass Access Zone Fallback as per normal.
2. In DR Assistant Wizard, after we selected the source cluster (Cluster B (for this example: name cluster21)) the next wizard screen will only display the Failover options From Cluster B (cluster21) to Cluster A (cluster20).

DR Assistant

Failover Wizard	Source Cluster	Target Cluster	Zone Name	Last Successful Readiness Check	Network Mapping	Overall Status
Running Failovers	<input type="checkbox"/> cluster21	cluster20	zone01	6/9/2016, 5:22:4...	View Map	OK
Failover History	<input type="checkbox"/> cluster21	cluster20	zone03	6/9/2016, 5:22:4...	View Map	OK
DR Testing (Beta)						

Back Next

3. Select the Access Zone to be failed back and the next screen will not highlight any warning about other SyncIQ policies that will not failed back.



4. Proceed this Access Zone Failback as per normal. Refer to
Eyeglass Access Zone Failover Guide for details.
3. Repeat the same procedure for failback other Access Zones
from B to A.

Failover from A to C Procedure:

1. Prior to initiate Eyeglass Access Zone Failover from A to C, **we need to ensure that there is no existing SyncIQ Mirror Policies from B to A. The recovery resync prep step of this Failover A to C will create Mirror Policies from C to A with same Mirror Target Paths as the B to A (Mirror Target Paths are overlaps). This will make the Mirror Policies from C to A unrunnable and the Eyeglass Failover Job will fail.** If there are existing ones, we need to delete them first. Refer to step P1 in the Failover workflow diagrams.

- a. NOTE: The above step MUST be completed before A to C failover, the order matters since the domain mark will be deleted on cluster A once the step above is completed.
2. Now run a domain mark job on each SyncIQ policy on cluster A. This is a required step since no domain mark exists and will be created during failover process from A to B. The best practise is to run domain mark before failover to ensure the resync prep step does not take a long time to complete. Domain mark can run longer if a the path has a large number of files. NOTE: During the time while domain mark job is running no failover from A to C should be executed until the domain mark job completes on ALL sync policies involved in the failover. Monitor progress from the PowerScale Cluster jobs UI.
3. Then we can perform Eyeglass Access Zone Failover as per normal. In DR Assistant Wizard, after we selected the source cluster (Cluster A (for this example: name cluster20)) the next wizard screen display the list of available Failover options based on Source-Target-Zone pairs.

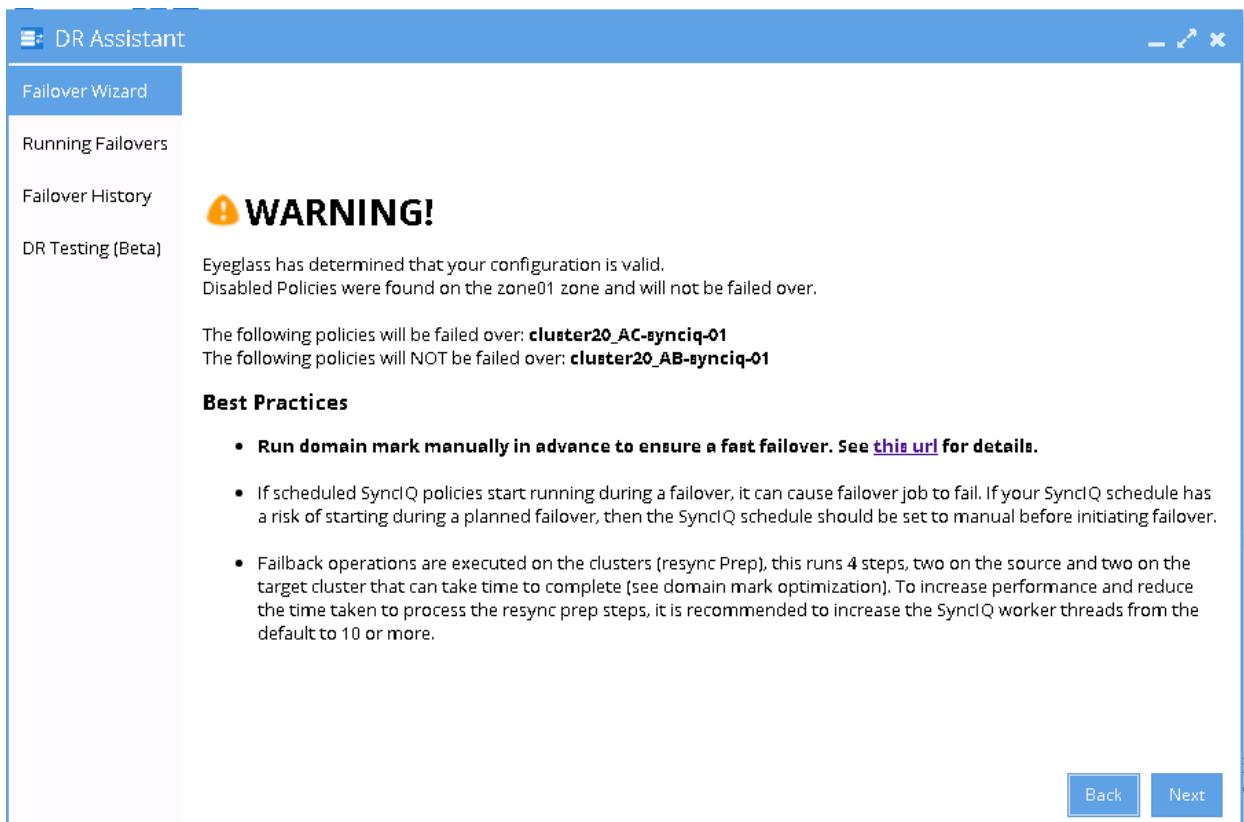
DR Assistant

Failover Wizard	Source Cluster	Target Cluster	Zone Name	Last Successful Readiness Check	Network Mapping	Overall Status
Running Failovers	<input type="checkbox"/> cluster20	cluster21	zone03	6/9/2016, 3:25:3...	View Map	OK
Failover History	<input type="checkbox"/> cluster20	cluster21	zone01	6/9/2016, 3:25:3...	View Map	OK
DR Testing (Beta)	<input type="checkbox"/> cluster20	cluster31	zone03	6/9/2016, 3:25:3...	View Map	OK
	<input type="checkbox"/> cluster20	cluster31	zone01	6/9/2016, 3:25:3...	View Map	OK

A to B A to C

Back Next

3. We need to be careful to select the correct Target Cluster that we want to Failover (A to B or A to C). For this case we want to failover from A to C. Select a zone that we want to Failover from cluster20 (source) - cluster31 (target) pairs.
4. The next screen will give warning to highlight that this wizard will only perform Access Zone failover from A to C. The other policy on the same Access Zone (A to B) will not be failed over.



5. Proceed this Access Zone Failover as per normal. Refer to Eyeglass Access Zone Failover Guide for details.
6. Repeat the same procedure for failover other Access Zones from A to C.

Failback from C to A Procedure

1. We can perform Eyeglass Access Zone Failback as per normal.
2. In DR Assistant Wizard, after we selected the source cluster (Cluster C (for this example: name cluster31)) the next wizard screen will only display the Failover options From Cluster C (cluster31) to Cluster A (cluster20).

DR Assistant

Failover Wizard	Source Cluster	Target Cluster	Zone Name	Last Successful Readiness Check	Network Mapping	Overall Status
Running Failovers	<input checked="" type="checkbox"/> cluster31	cluster20	zone03	6/9/2016, 6:31:5...	View Map	 OK
Failover History	<input type="checkbox"/> cluster31	cluster20	zone01	6/9/2016, 6:31:5...	View Map	 OK
DR Testing (Beta)						

Back Next

3. Select the Access Zone to be failed back and the next screen will not highlight any warning about other SyncIQ policies that will not fail back.

DR Assistant

Failover Wizard

Running Failovers

Failover History

DR Testing (Beta)

SUCCESS

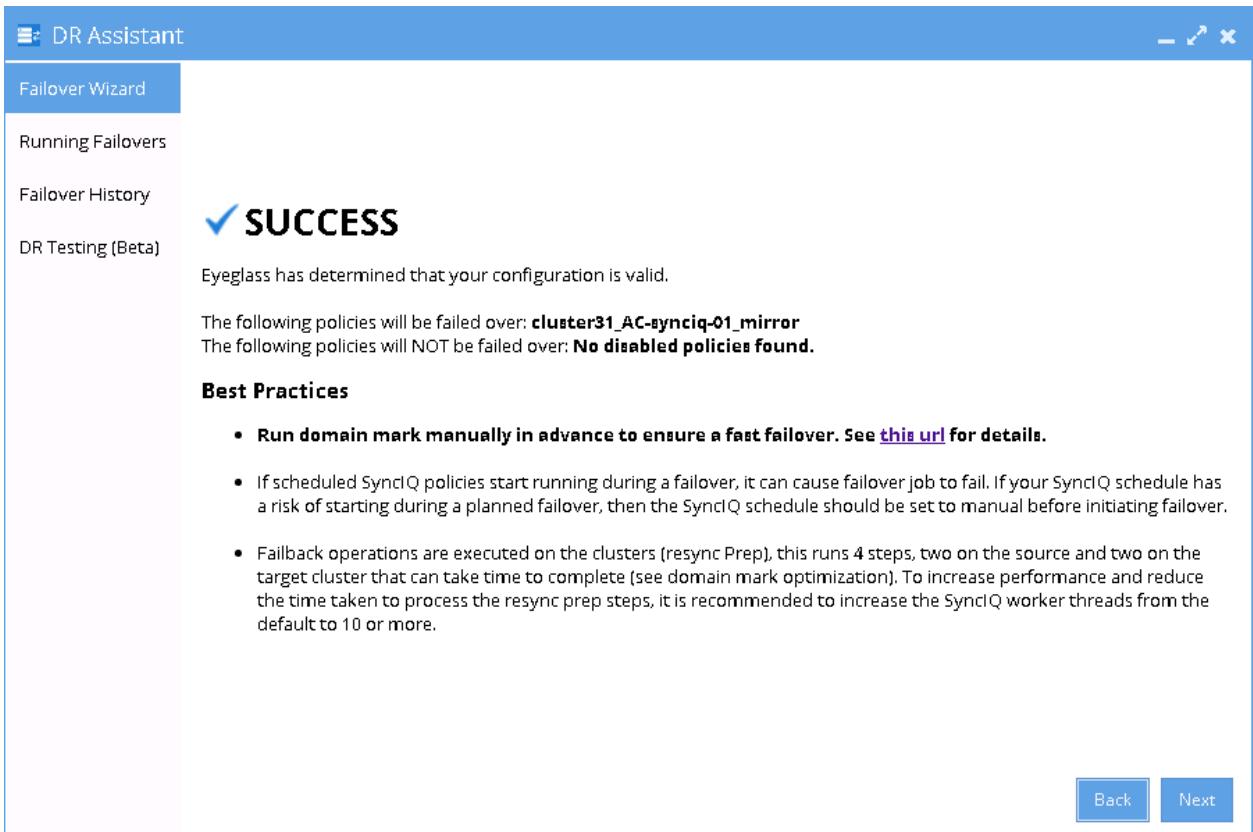
Eyeglass has determined that your configuration is valid.

The following policies will be failed over: **cluster31_AC-synciq-01_mirror**
The following policies will NOT be failed over: **No disabled policies found.**

Best Practices

- Run domain mark manually in advance to ensure a fast failover. See [this url](#) for details.
- If scheduled SyncIQ policies start running during a failover, it can cause failover job to fail. If your SyncIQ schedule has a risk of starting during a planned failover, then the SyncIQ schedule should be set to manual before initiating failover.
- 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



4. Proceed this Access Zone Failback as per normal. Refer to Eyeglass Access Zone Failover Guide for details.
5. Repeat the same procedure for failback other Access Zones from C to A.

© Superna Inc

6.4. 3 Site DFS Mode Failover

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

3 Site DFS Mode Failover

- [Overview](#)
- [Video How to - Overview Multi site DFS mode Failover](#)
- [Configuration](#)
- [DFS Mode Initial Configuration / Before Failover Diagram](#)
- [DFS Mode Failover A ⇒ B Diagram](#)
- [DFS Mode Failback B ⇒ A Diagram](#)
- [DFS Mode Failover A ⇒ C Diagram](#)
- [DFS Mode Failback C ⇒ A Diagram](#)
- [Eyeglass DFS Mode Failover Steps](#)
- [Eyeglass DFS Mode Failback Steps](#)
- [DFS Configuration](#)
- [DFS Readiness](#)
- [DFS Readiness - Initial Configuration / Before Failover](#)
- [DFS Readiness - Before Failback B ⇒ A](#)
- [DFS Readiness - Before Failback C ⇒ A](#)
- [Share Names and DFS Paths](#)
- [DFS Mode Failover and Failback Procedures](#)

- DFS Mode Failover from A to B Procedure:
- DFS Mode Failback from B to A Procedure:
- DFS Mode Failover from A to C Procedure:
- DFS Mode Failback from C to A Procedure:

This section will explain the configuration, failover and fallback workflows for 3 Sites DFS Mode Failover with Eyeglass for PowerScale. As explained in the previous sections of this document, there are 3 Sites for this setup: Site A (Source), Site B (Target #1) and Site C (Target #2).

Overview

This solution offers simply 2 site target with clients automatically redirected to the correct site.

- No DNS change
- No SPN changes
- Quotas follow shares as required to each site on failover and failback
- 3 DFS targets per folder

- Highest availability option for data with zero touch failover between sites

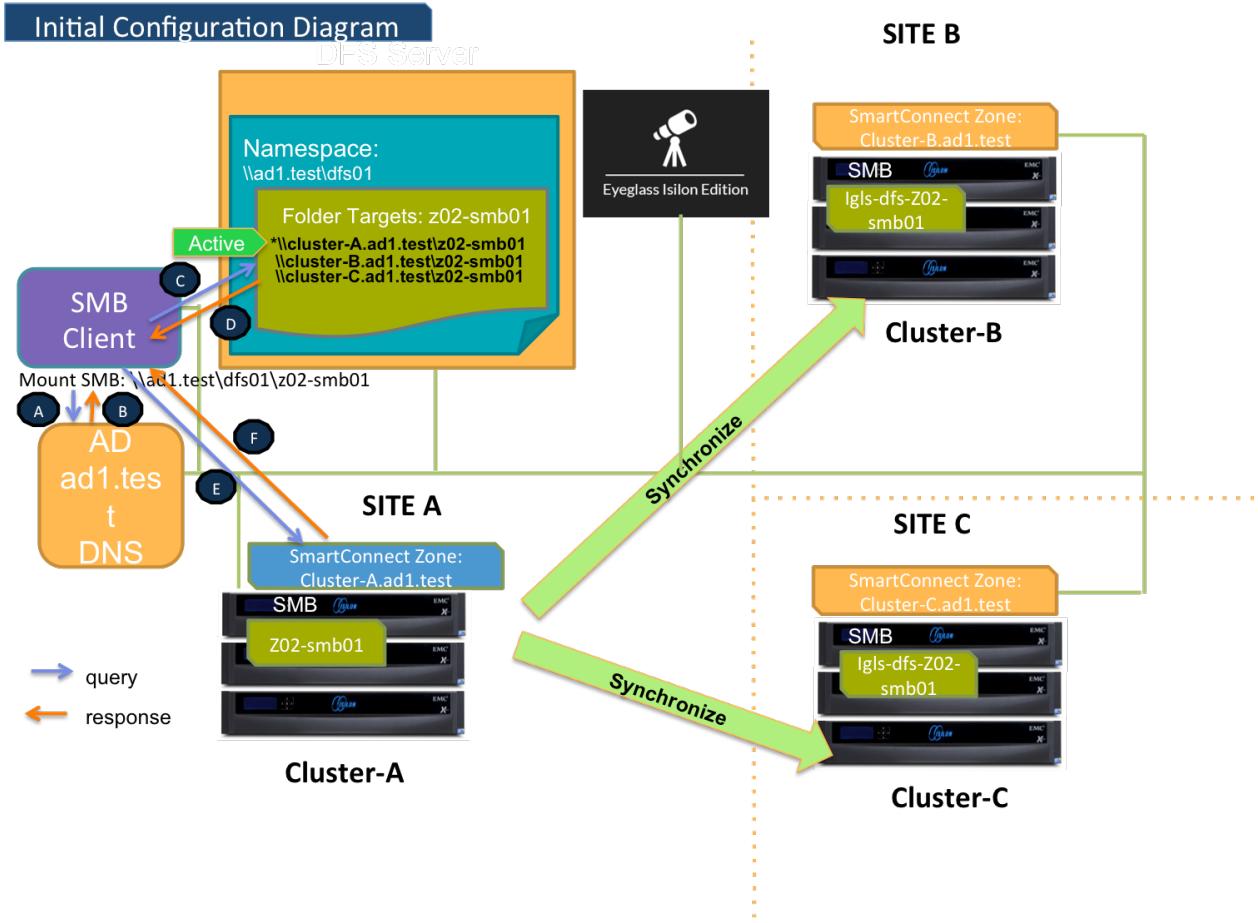
Video How to - Overview Multi site DFS mode Failover

Configuration

For this 3 Sites DFS Mode Failover, we need to configure the DFS Target Folder to have 3 referrals to 3 PowerScale Clusters. Data on the SMB folders referred as the DFS Target Folder is replicated from Site A to Site B, and also from Site A to Site C by using PowerScale SyncIQ replication.

DFS Mode Initial Configuration / Before Failover Diagram

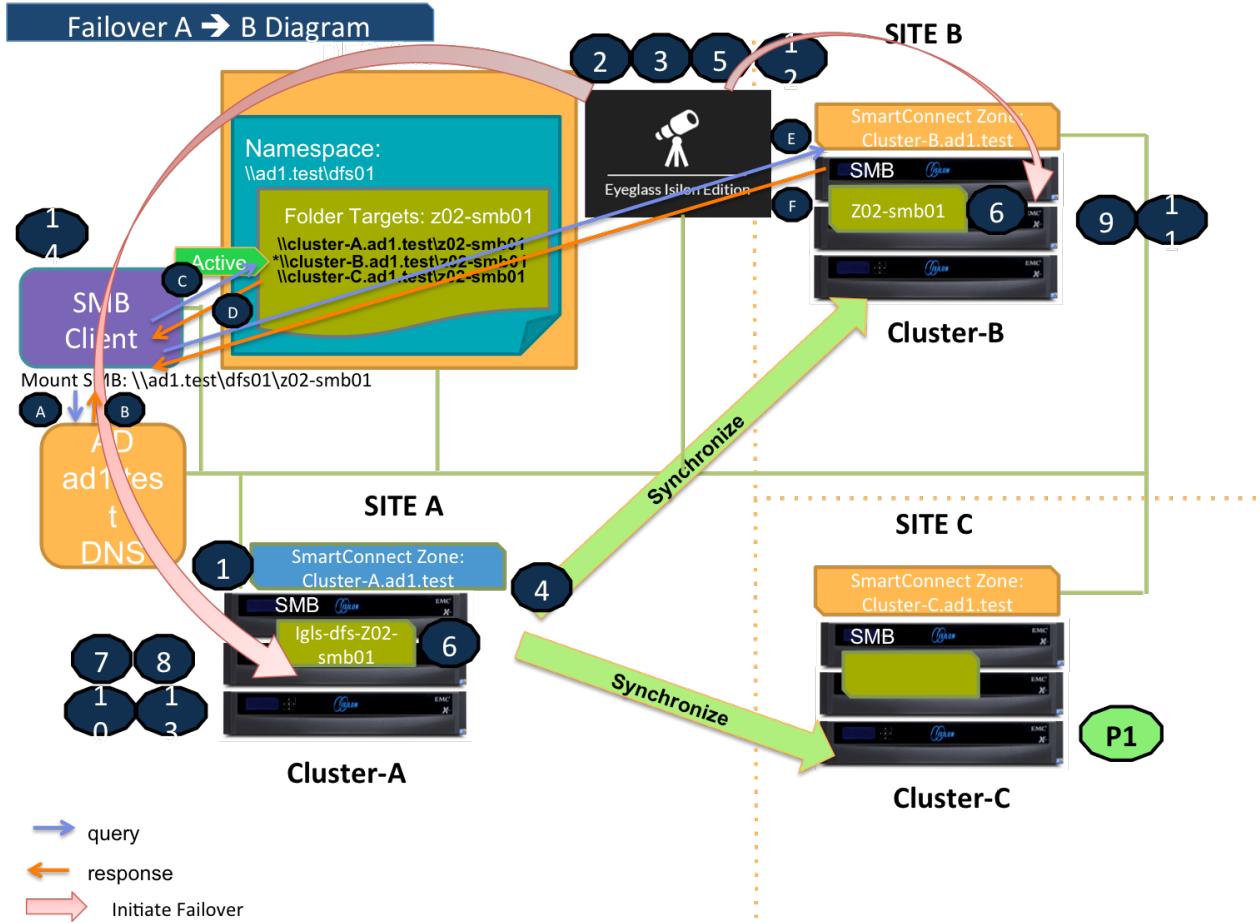
This diagram displays the initial configuration for this 3 Sites DFS Mode Failover.



DFS Mode Failover A ⇒ B Diagram

This diagram shows the Failover workflow from A to B. Take note step P1 (Preparation Step - prior to initiate Eyeglass DFS Mode Failover) - refer to the [procedure section](#) for details.

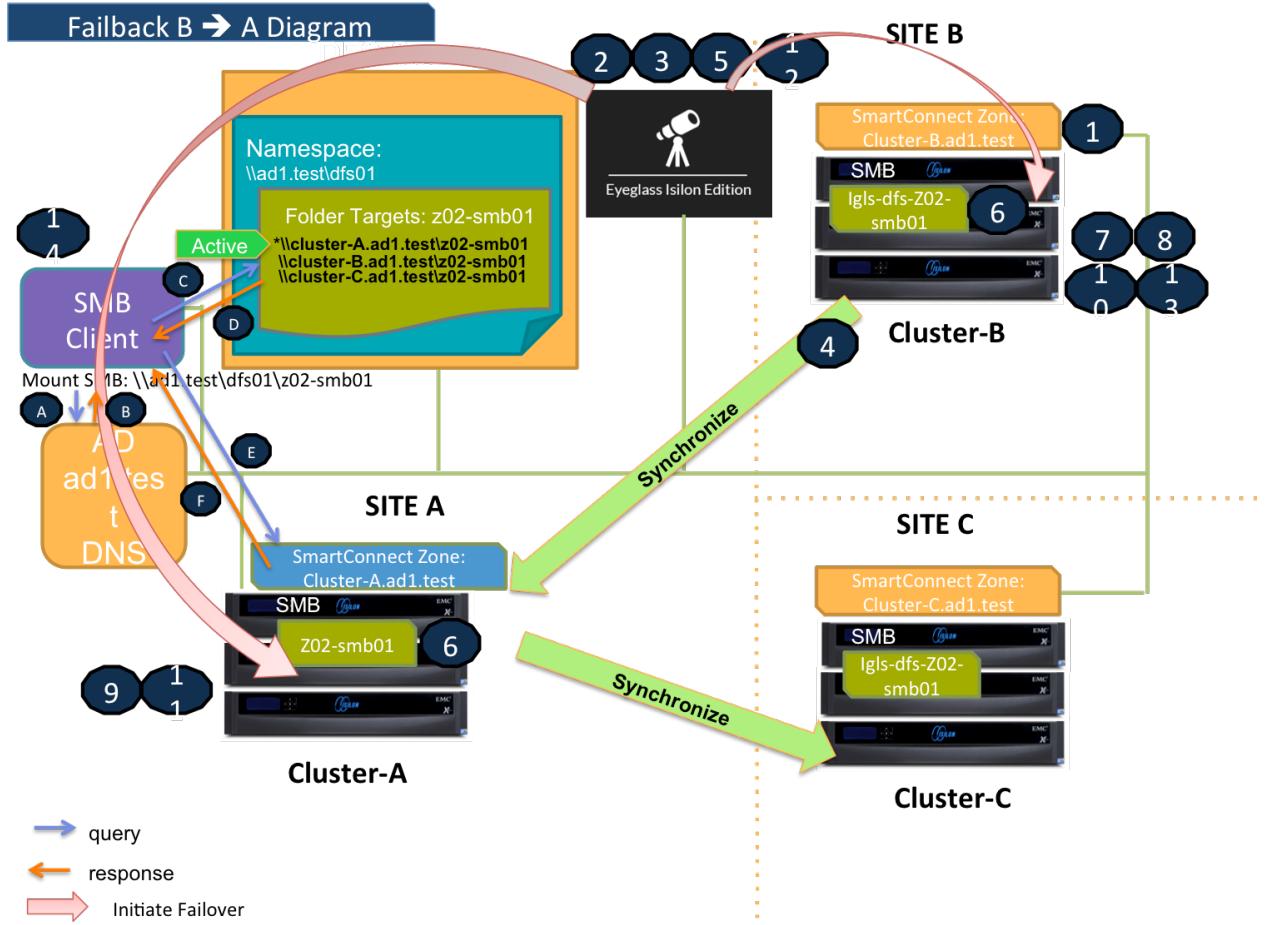
Refer to [this table](#) for the list of the numbered steps shown in this diagram.



DFS Mode Failback B ⇒ A Diagram

This diagram shows the Failback workflow from B to A.

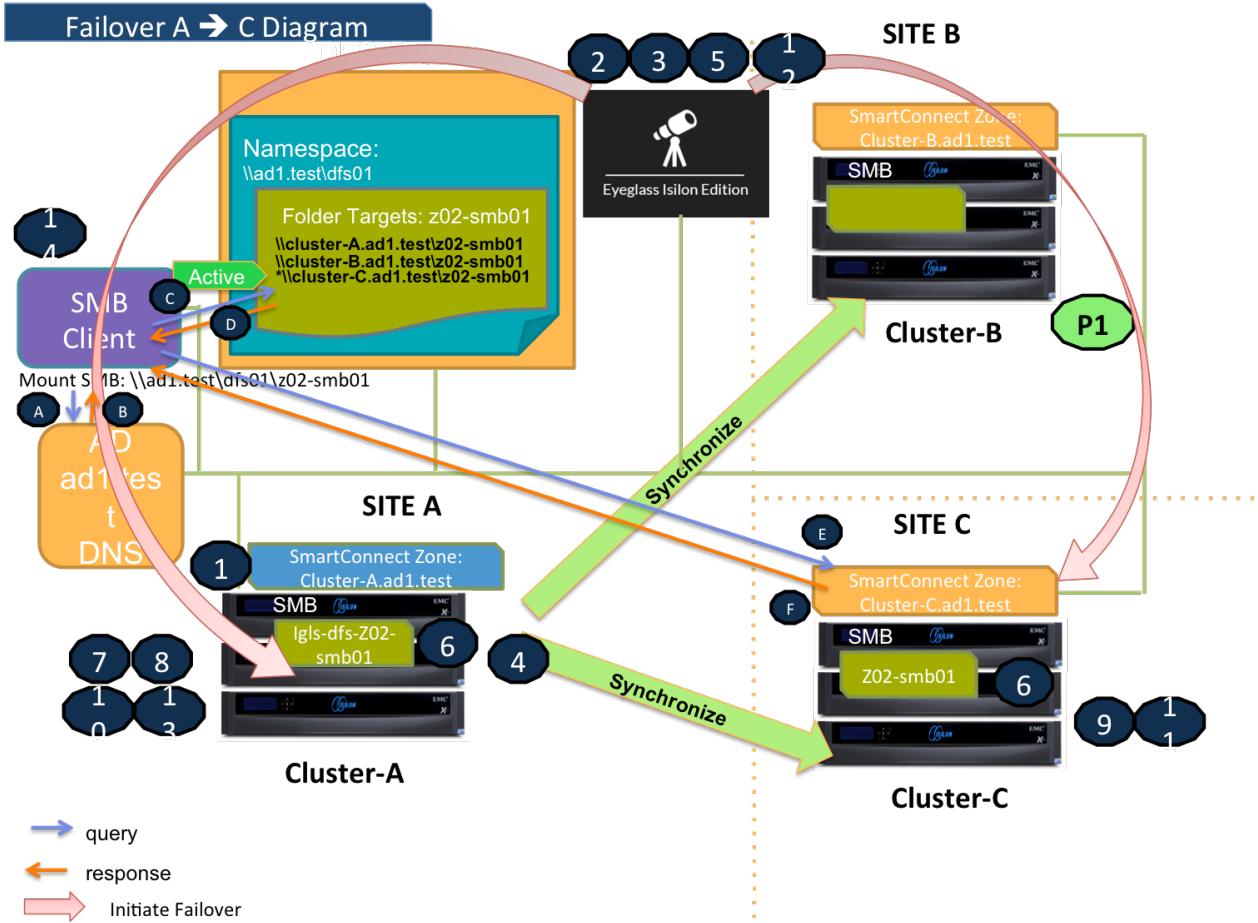
Refer to [this table](#) for the list of the numbered steps shown in this diagram.



DFS Mode Failover A ⇒ C Diagram

This diagram shows the Failover workflow from A to C. Take note step P1 (Preparation Step - prior to initiate Eyeglass DFS Mode Failover) - refer to the [procedure section](#) for details.

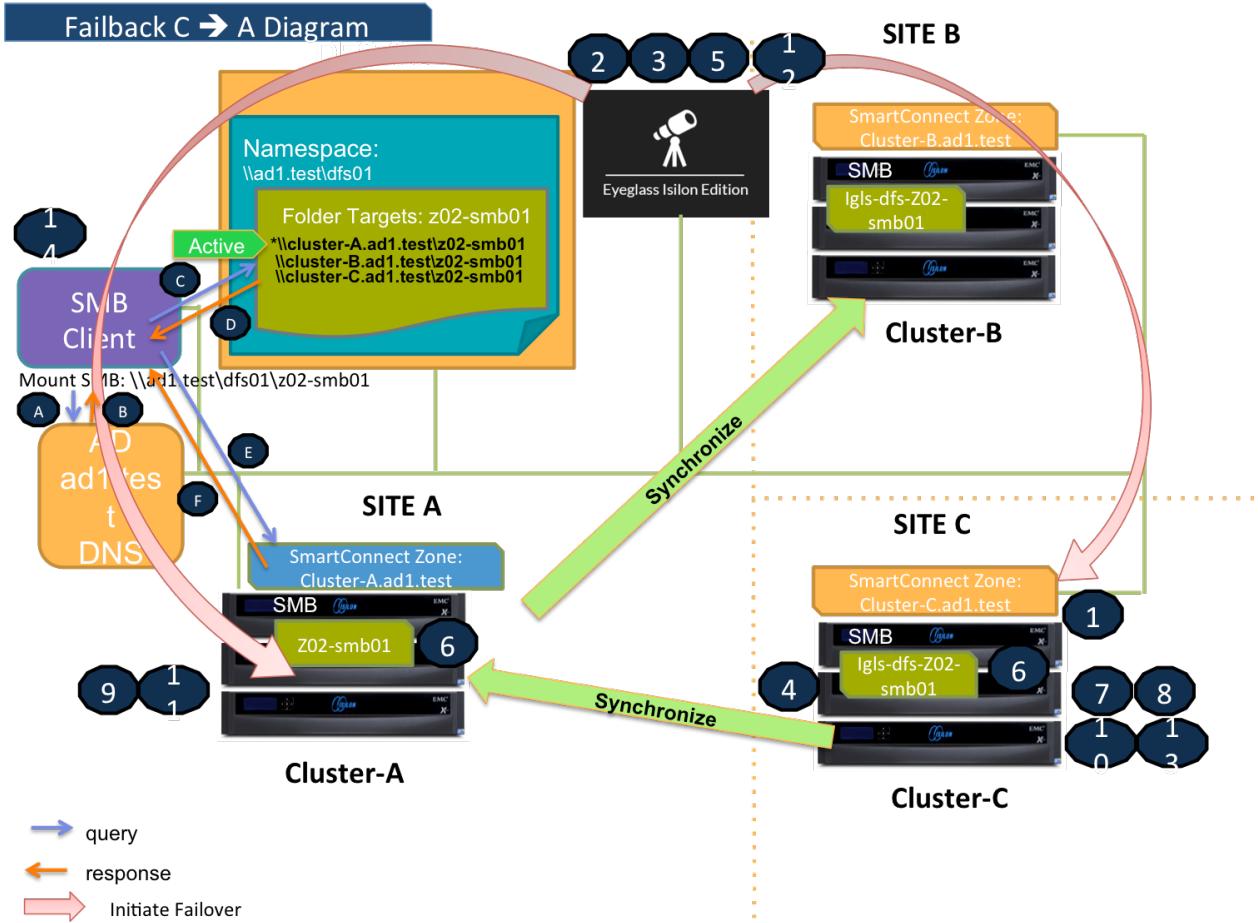
Refer to [this table](#) for the list of the numbered steps shown in this diagram.



DFS Mode Failback C ⇒ A Diagram

This diagram shows the Failback workflow from C to A.

Refer to [this table](#) for the list of the numbered steps shown in this diagram.



Eyeglass DFS Mode Failover Steps

This table lists the Eyeglass DFS Mode Failover steps with numbers as shown in the above Failover diagrams.

No	Steps
P1	<p>Preparation Step.</p> <p>Failover A ⇒ B: Ensure there is no existing Mirror Policies between C to A. If there is existing Mirror Policies between C to A, delete first, before initiate Failover from A to B.</p> <p>Failover A ⇒ C: Ensure there is no existing Mirror Policies between B to A. If there is existing Mirror Policies between B to A, delete first, before initiate Failover from A to C.</p>
1	Ensure that there is no live access to data

2	Begin Failover
3	Validation
4	Synchronize data
5	Synchronize configuration (shares/export/alias)
6	Renaming Shares
7	Record schedule for SyncIQ policies being failed over
8	Prevent SyncIQ policies being failed over from running
9	Provide write access to data on target
10	Disable SyncIQ on source and make active on target
11	Set proper SyncIQ schedule on target
12	Synchronize quota(s)
13	Remove quotas on directories that are target of SyncIQ (PowerScale best practice)
14	<p>Refresh SMB session to pick up DFS change:</p> <ol style="list-style-type: none"> 1. SMB Client is accessing a domain-based namespace (e.g. \\ad1.test\dfs01\z02-smb01) . This SMB client computer sends a query to the AD to discover a list of root targets for the namespace. 2. AD Controller returns a list of root targets defined for the requested namespace. 3. SMB client selects the root target from the referral list and sends a query to the root server for the requested link. 4. DFS root server constructs a list of folder targets in the referral. <p>1. Failover A ⇒ B:</p> <ol style="list-style-type: none"> 1. The SMB Share(s) on Cluster-A is not active (Renamed with igls-dfs- prefix), 2. The SMB Share(s) on Cluster-C is not active. (Deleted).

	<p>3. The active path is to the Cluster-B (Renamed to the actual name). DFS root server sends this referral information to the client.</p> <p>2. Failover A ⇒ C:</p> <ol style="list-style-type: none"> 1. The SMB Share(s) on Cluster-A is not active (Renamed with igls-dfs- prefix), 2. The SMB Share(s) on Cluster-B is not active. (Deleted). 3. The active path is to the Cluster-C (Renamed to the actual name). DFS root server sends this referral information to the client. 5. SMB client tries to establish a connection to the selected target (the active target in the list). 6. PowerScale with Active Target responds to this SMB connection.
--	---

Eyeglass DFS Mode Failback Steps

This table lists the Eyeglass DFS Mode Failback steps with numbers as shown in the above Failback diagrams.

No	Steps
1	Ensure that there is no live access to data
2	Begin Failback
3	Validation
4	Synchronize data
5	Synchronize configuration (shares/export/alias)
6	Renaming Shares
7	Record schedule for SyncIQ policies being failed back
8	Prevent SyncIQ policies being failed back from running

9	Provide write access to data on target
10	Disable SyncIQ on source and make active on target
11	Set proper SyncIQ schedule on target
12	Synchronize quota(s)
13	Remove quotas on directories that are target of SyncIQ (PowerScale best practice)
14	<p>Refresh SMB session to pick up DFS change:</p> <ol style="list-style-type: none"> 1. SMB Client is accessing a domain-based namespace (e.g. \\ad1.test\dfs01\z02-smb01) . This SMB client computer sends a query to the AD to discover a list of root targets for the namespace. 2. AD Controller returns a list of root targets defined for the requested namespace. 3. SMB client selects the root target from the referral list and sends a query to the root server for the requested link. 4. DFS root server constructs a list of folder targets in the referral. <p>1. Fallback B ⇒ A:</p> <ol style="list-style-type: none"> 1. The SMB Share(s) on Cluster-B is not active (Renamed with the igls-dfs- prefix), 2. The SMB Share(s) on Cluster-C is not active. (Renamed with the igls-dfs- prefix). 3. The active path is to the Cluster-A (Renamed to the actual name). DFS root server sends this referral information to the client. <p>2. Fallback C ⇒ A:</p> <ol style="list-style-type: none"> 1. The SMB Share(s) on Cluster-C is not active (Renamed with the igls-dfs- prefix), 2. The SMB Share(s) on Cluster-B is not active.

	<p>(Renamed with the igls-dfs- prefix).</p> <ol style="list-style-type: none"> 3. The active path is to the Cluster-A (Renamed to the actual name). DFS root server sends this referral information to the client. 5. SMB client tries to establish a connection to the selected target (the active target in the list). 6. PowerScale with Active Target responses to this SMB connection.
--	--

DFS Configuration

Configure the DFS Target Folder to have 3 referrals - Site A, Site B and Site C.

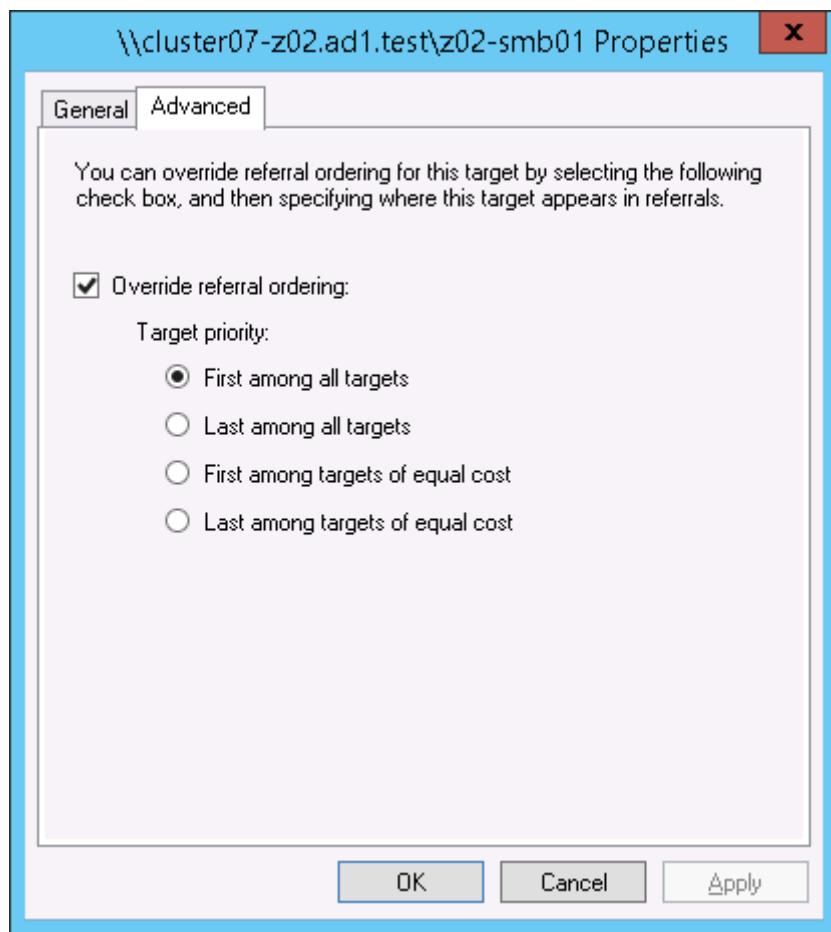
For an example we have configured the DFS Target Folder to have these three referrals:

1. Source (Site A): \\cluster07-z02.ad1.test\z02-smb01
2. Target#1 (Site B): \\cluster08-z02.ad1.test\z02-smb01
3. Target#2 (Site C): \\cluster06-z02.ad1.test\z02-smb01

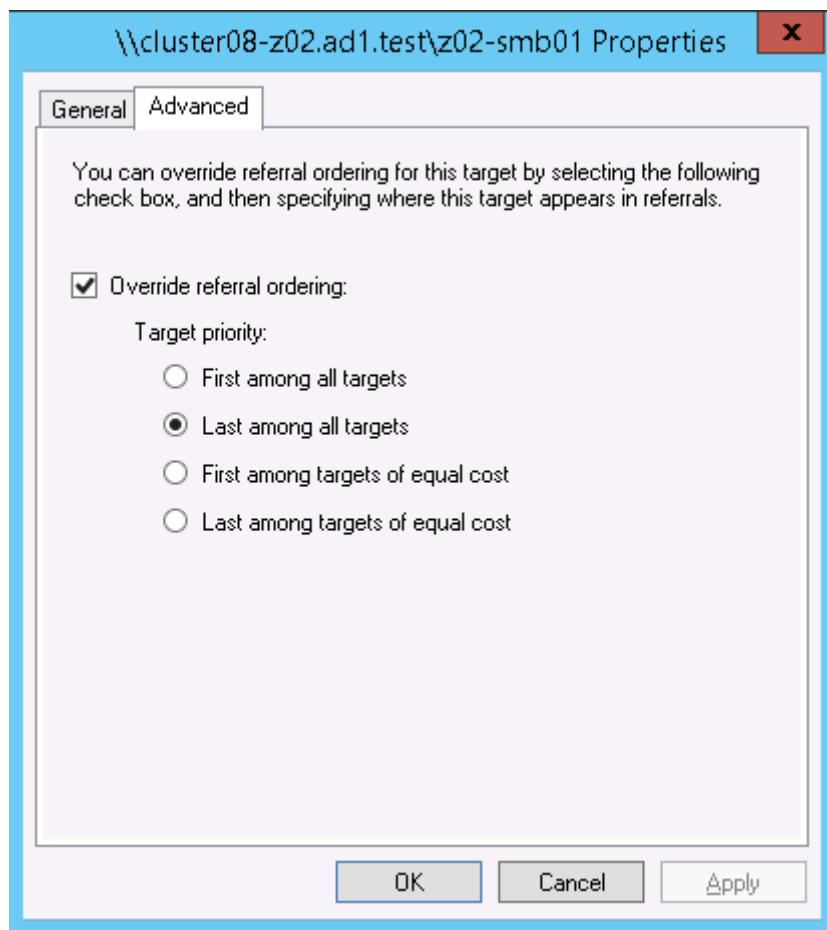
z02-smb01			
Folder Targets		Replication	
Type	Referral Status	Site	Path
Enabled	Enabled	Site-C	\\\cluster06-z02.ad1.test\z02-smb01
Enabled	Enabled	PrimarySite	\\\cluster07-z02.ad1.test\z02-smb01
Enabled	Enabled	SecondarySite	\\\cluster08-z02.ad1.test\z02-smb01

We have also configured the following target priority referral ordering:

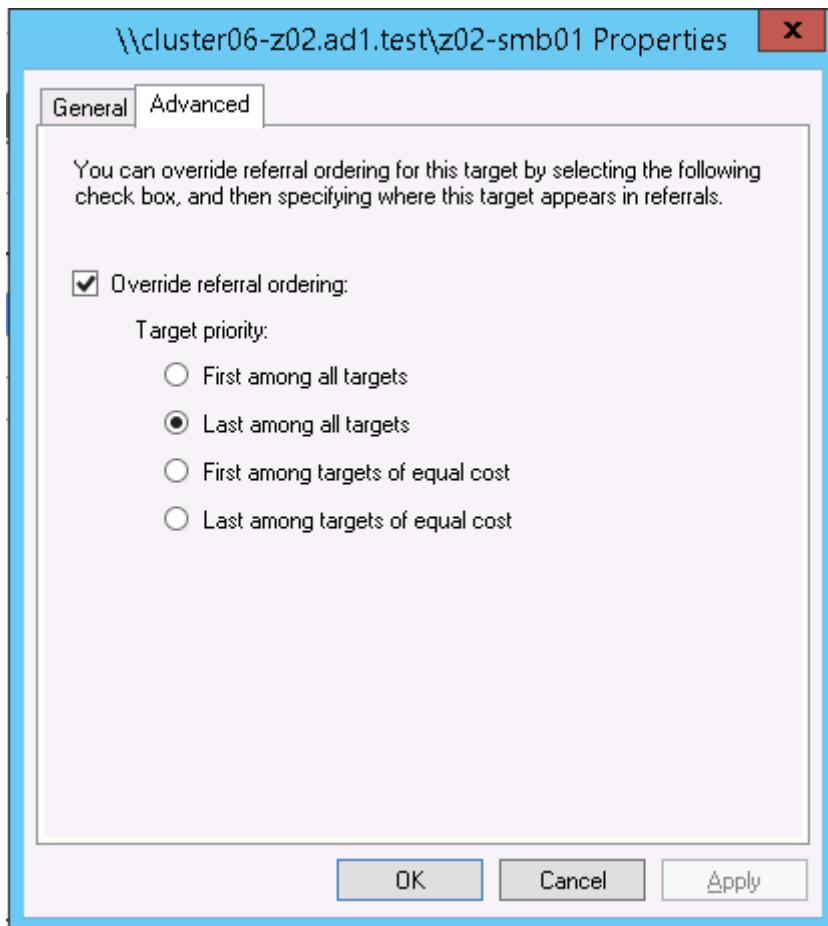
Source Cluster (Site A)



Target Cluster #1 (Site B)



Target Cluster #2 (Site C)



DFS Readiness

This section explains the different states of DFS Readiness for this 3 Sites DFS Mode Failover / Failback.

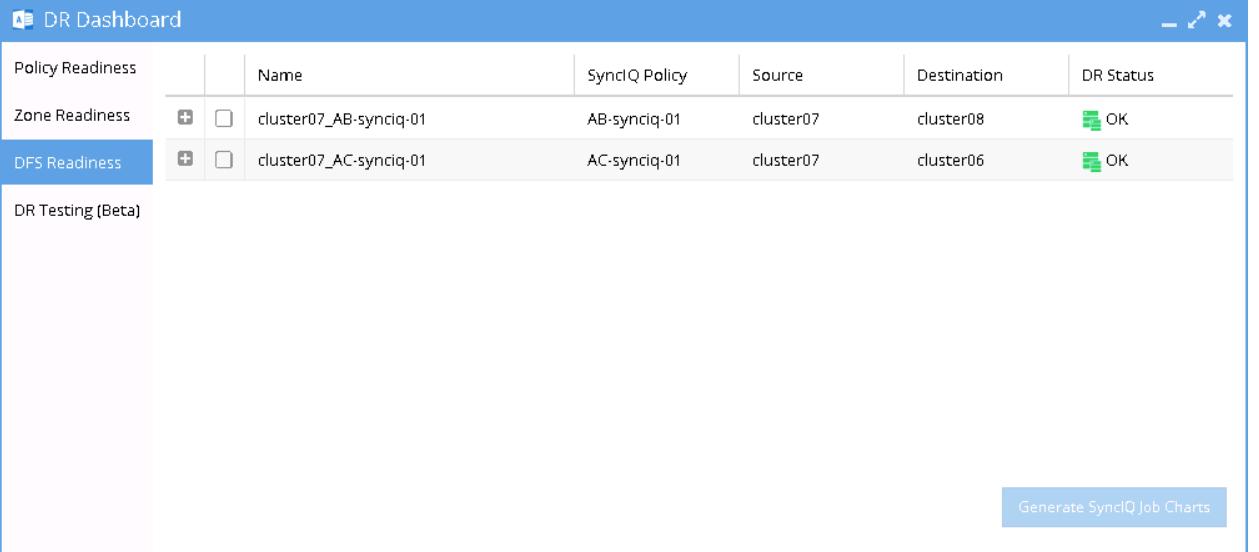
For the purpose of this example we use the following names:

Site	Cluster Name
A (Primary / Source)	cluster07
B (Secondary#1 / Target #1)	cluster08
C (Secondary#2 / Target #2)	cluster06

DFS Readiness - Initial Configuration / Before Failover

This is the DFS Readiness for Initial Configuration / before Failover state. As we can see from this figure, that both Source-Target Pairs (A - B and A - C) are listed in this DR Dashboard' DFS readiness window.

This shows that a DFS Mode failover choice can be made to any target cluster in Green OK state (Warning status also allowed).



Policy Readiness	Name	SyncIQ Policy	Source	Destination	DR Status
Zone Readiness	[+]	cluster07_AB-synciq-01	AB-synciq-01	cluster07	cluster08
DFS Readiness	[+]	cluster07_AC-synciq-01	AC-synciq-01	cluster07	cluster06

DR Testing (Beta)

Generate SyncIQ Job Charts

DFS Readiness - Before Failback B \Rightarrow A

This DFS Readiness is for the state before Failback from B to A.

DR Dashboard						
Policy Readiness		Name	SyndIQ Policy	Source	Destination	DR Status
Zone Readiness	[+]	<input type="checkbox"/> cluster07_AB-synciq-01	AB-synciq-01	cluster07	cluster08	DISABLED
DFS Readiness	[+]	<input type="checkbox"/> cluster07_AC-synciq-01	AC-synciq-01	cluster07	cluster06	OK
DR Testing (Beta)	[+]	<input type="checkbox"/> cluster08_AB-synciq-01_mirror	AB-synciq-01_mir...	cluster08	cluster07	OK

[Generate SyndIQ Job Charts](#)

Warning: As shown in this DR Readiness Dashboard that both AB Mirror Policy and AC Policy have DR Status OK. During Failback from B to A, we need to carefully select Cluster B as the source. Do not select the Cluster A as the source, as this will direct the process as Failover from A to C.

DFS Readiness - Before Failback C \Rightarrow A

This DFS Readiness is for the state before Failback from C to A.

DR Dashboard						
Policy Readiness		Name	SyndIQ Policy	Source	Destination	DR Status
Zone Readiness	[+]	<input type="checkbox"/> cluster07_AB-synciq-01	AB-synciq-01	cluster07	cluster08	OK
DFS Readiness	[+]	<input type="checkbox"/> cluster07_AC-synciq-01	AC-synciq-01	cluster07	cluster06	DISABLED
DR Testing (Beta)	[+]	<input type="checkbox"/> cluster06_AC-synciq-01_mirror	AC-synciq-01_mir...	cluster06	cluster07	OK

[Generate SyndIQ Job Charts](#)

Warning: As shown in this DR Readiness Dashboard that both AC Mirror Policy and AB Policy have DR Status OK. During Fallback from C to A, we need to carefully select Cluster C as the source. Do not select the Cluster A as the source, as this will direct the process as Failover from A to B.

Share Names and DFS Paths

Based on the above example, the following table describes the SMB Share Names and DFS Paths for various states:

Initial Configuration / Before Failover

	Cluster07 (A)	Cluster08 (B)	Cluster06 (C)
Share Name	z02-smb01	igls-dfs-z02-smb01	igls-dfs-z02-smb01
DFS Path Resolves to	\cluster07-z02.ad1.test\z02-smb01		
Access Status	0 (ACTIVE TARGETSET)	0xc00000cc (TARGETSET)	0xc00000cc

After Failover / After Fallback

		Cluster07 (A)	Cluster08 (B)	Cluster06 (C)
After Failover A => B	Share Name	igls-dfs-z02-smb01	z02-smb01	*1
	DFS Path Resolves to		\cluster08-z02.ad1.test\z02-smb01	
	Access Status	0xc00000cc (TARGETSET)	0 (ACTIVE TARGETSET)	0xc00000cc
After Fallback B => A	Share Name	z02-smb01	igls-dfs-z02-smb01	igls-dfs-z02-smb01*2
	DFS Path Resolves to	\cluster07-z02.ad1.test\z02-smb01		

	Access Status	0 (ACTIVE TARGETSET)	0xc00000cc (TARGETSET)	0xc00000cc
After Failover A => C	Share Name	igls-dfs-z02-smb01	*3	z02-smb01
	DFS Path Resolves to			\cluster06-z02.ad1.test\z02-smb01
	Access Status	0xc00000cc (TARGETSET)	0xc00000cc (TARGETSET)	0 (ACTIVE)
After Failback C => A	Share Name	z02-smb01	igls-dfs-z02-smb01*4	igls-dfs-z02-smb01
	DFS Path Resolves to	\cluster07-z02.ad1.test\z02-smb01		
	Access Status	0 (ACTIVE TARGETSET)	0xc00000cc (TARGETSET)	0xc00000cc

Remarks for Intermediate and Final States:

*1: States:

1. After Failover A \Rightarrow B process has just Completed: **igls-dfs-z02-smb01** (Intermediate State)
2. The 1st cycle of Configuration Replication (A \Rightarrow C) after failover A \Rightarrow B: **igls-dfs-igls-dfs-z02-smb01** (Intermediate State)
3. The 2nd cycle of Configuration Replication (A \Rightarrow C) after failover A \Rightarrow B: <empty> SMB shares deleted. (Final state)

*2: States:

1. After Failback B \Rightarrow A just Completed: <empty> SMB shares is not created (Intermediate State)

2. The 1st cycle of Configuration Replication (A \Rightarrow C) after failback B \Rightarrow A: **igls-dfs-z02-smb01** (Final State)

*3: States:

1. After Failover A \Rightarrow C just Completed: **igls-dfs-z02-smb01** (Intermediate State)
2. The 1st cycle of Configuration Replication (A \Rightarrow B) after failover A \Rightarrow C: **igls-dfs-igls-dfs-z02-smb01** (Intermediate State)
3. The 2nd cycle of Configuration Replication (A \Rightarrow B) after failover A \Rightarrow C:
<empty> SMB shares deleted. (Final state)

*4: States:

1. After Failback C \Rightarrow A just Completed: <empty> SMB shares is not created
(Intermediate State)
2. The 1st cycle of Configuration Replication (A \Rightarrow B) after failback C \Rightarrow A: **igls-dfs-z02-smb01** (Final State).

Based on that table we can see that after failover, it takes 2 cycles of Configuration Replication as waiting time for the SMB share name on the 3rd cluster to have its final state.

For the case of failback, it takes 1 cycle of Configuration Replication process as waiting time for the SMB share name on the 3rd cluster to have its final state.

DFS Mode Failover and Failback Procedures

It is recommended to create SyncIQ Policies that will be used for multi site replications (e.g. to replicate from Site A to Site B and also from Site A to Site C) with names that reflect the Source-Target pairs.

The following table is an example:

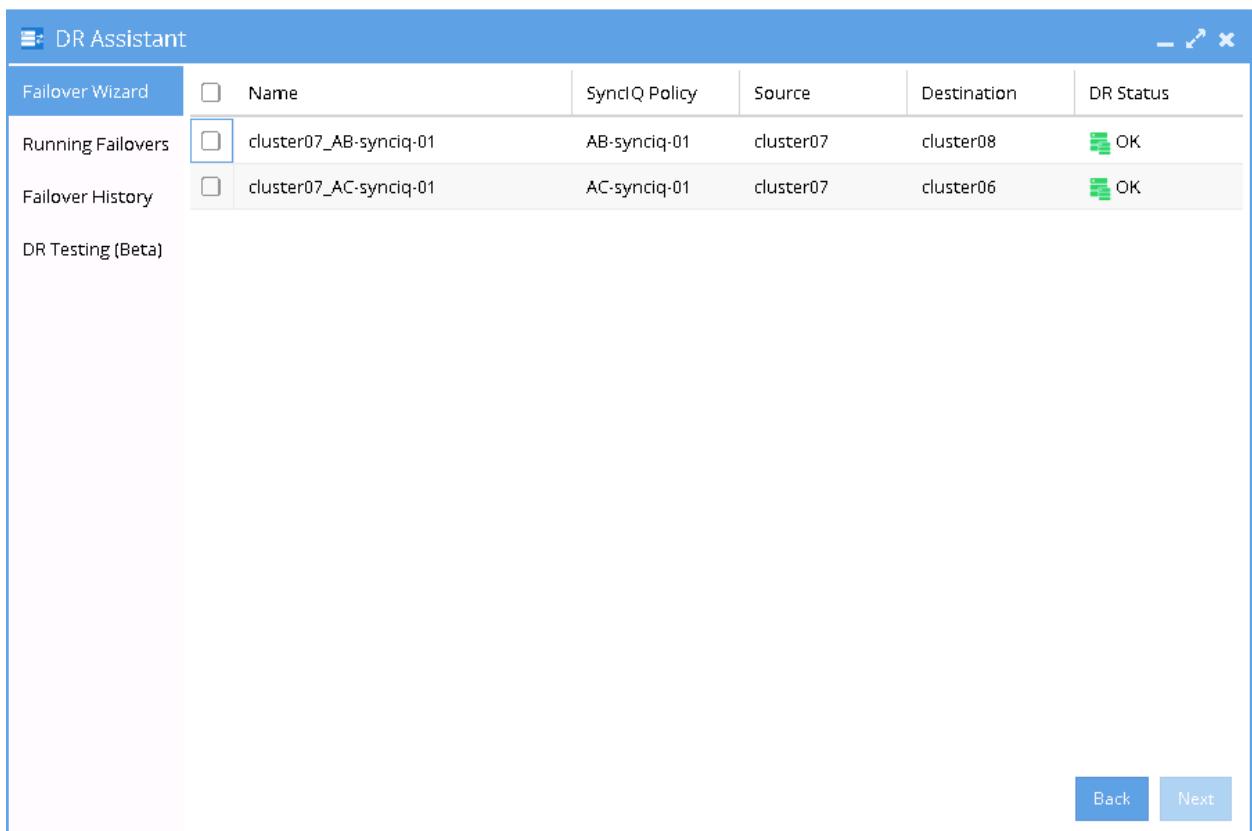
SyncIQ Policy Name	SyncIQ Pairs
AB-synciq-01	A and B
AC-synciq-01	A and C

This name format will help us to identify which SyncIQ Pairs that we want to failover.

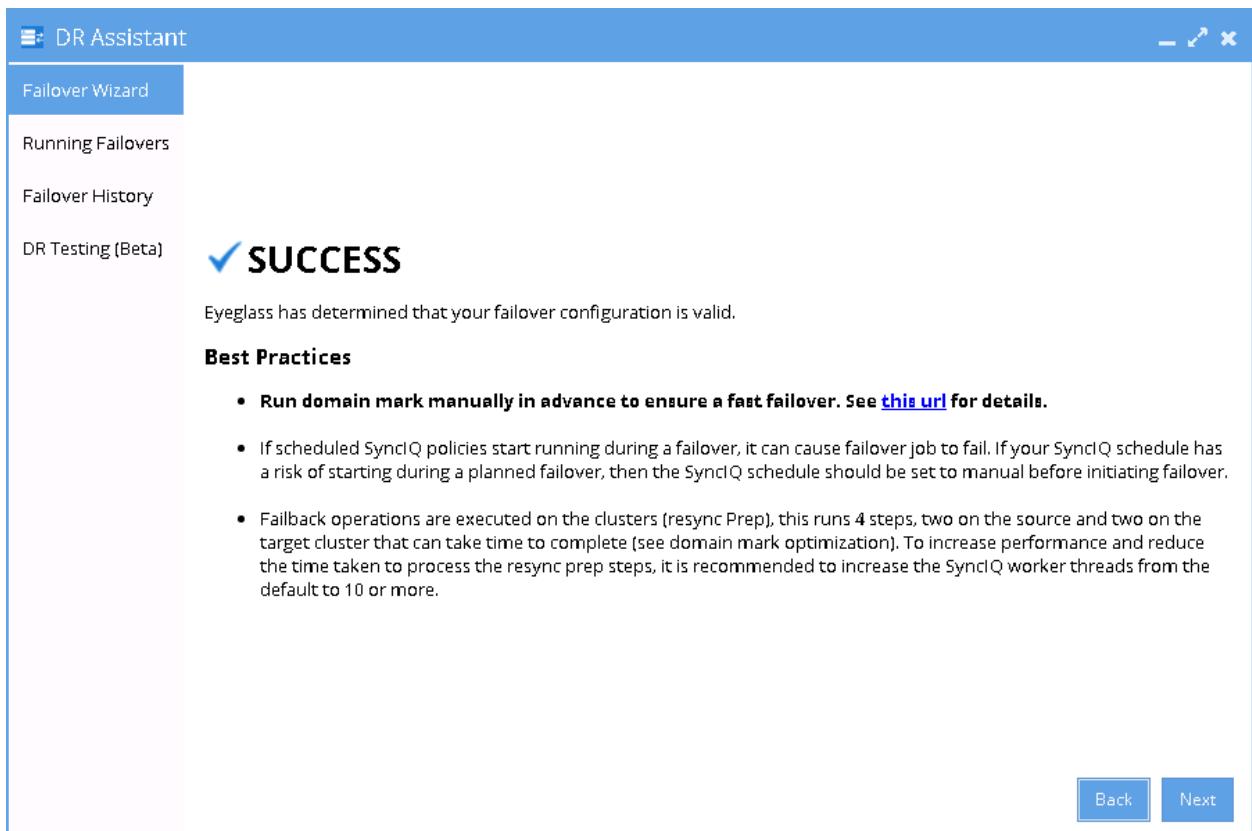
DFS Mode Failover from A to B Procedure:

1. Prior to initiate Eyeglass DFS Mode Failover from A to B, **we need to ensure that there is no existing SyncIQ Mirror Policies from C to A. The recovery resync prep step of this Failover A to B will create Mirror Policies from B to A with same Mirror Target Paths as the C to A (Mirror Target Paths are overlaps). This will make the Mirror Policies from B to A unrunnable and the Eyeglass Failover Job will fail.** If there are existing ones, we need to delete them first. Refer to step P1 in the Failover workflow diagrams.
 - a. **NOTE: The above step MUST be completed before A to C failover, the order matters since the domain mark will be deleted on cluster A once the step above is completed.**

2. Now run a domain mark job on each SyncIQ policy on cluster A. This is a required step since no domain mark exists and will be created during failover process from A to B. The best practise is to run domain mark before failover to ensure the resync prep step does not take a long time to complete. Domain mark can run longer if a the path has a large number of files. NOTE: During the time while domain mark job is running no failover from A to C should be executed until the domain mark job completes on ALL sync polices involved in the failover. Monitor progress from the PowerScale Cluster jobs UI.
3. Then we can perform Eyeglass DFS Mode Failover as per normal. In DR Assistant Wizard, after we select the source cluster (Cluster A (for this example: cluster07)) the next wizard screen display the list of available Failover options based on Source-Target pairs (A to B or A to C).



3. We need to be careful to select the correct Target Cluster that we want to Failover (A to B or A to C). For this case we want to failover from A to B. Select the AB Source-Target Pair.
2. The next screen will give validation whether the failover configuration is valid.



5. Proceed this DFS Mode Failover as per normal. Refer to
Eyeglass DFS Mode Failover Guide for details.

DFS Mode Failback from B to A Procedure:

1. We can perform Eyeglass DFS Mode Failback as per normal.
2. In DR Assistant Wizard, **ensure we select the correct source cluster B (name: cluster08). At this stage (After Failover A to B and before Failback from B to A), there are 2 available options to perform as also displayed in the DR Dashboard DFS Readiness. Do not select cluster A (name : cluster07) as the source, as this will lead to Failover from A to C instead.**
 - a. **NOTE: The above step MUST be completed before A to C failover, the order matters since the domain mark will**

be deleted on cluster A once the step above is completed.

3. Now run a domain mark job on each SyncIQ policy on cluster A. This is a required step since no domain mark exists and will be created during failover process from A to B. The best practise is to run domain mark before failover to ensure the resync prep step does not take a long time to complete. Domain mark can run longer if a the path has a large number of files. NOTE: During the time while domain mark job is running no failover from A to C should be executed until the domain mark job completes on ALL sync policies involved in the failover. Monitor progress from the PowerScale Cluster jobs UI.
4. After we select the correct source cluster (Cluster B (for this example: cluster08)) the next wizard screen will only display the Failback option From Cluster B (cluster08) to Cluster A (cluster07).

DR Assistant

Failover Wizard

Name	SynclQ Policy	Source	Destination	DR Status
cluster08_AB-synciq-01_mirror	AB-synciq-01_m...	cluster08	cluster07	OK

Running Failovers

Failover History

DR Testing (Beta)

Back Next

4. Select the AB mirror policy to fallback. The next screen will give validation whether the failover configuration is valid.

DR Assistant

Failover Wizard

Running Failovers

Failover History

DR Testing (Beta)

SUCCESS

Eyeglass has determined that your failover configuration is valid.

Best Practices

- Run domain mark manually in advance to ensure a fast failover. See [this url](#) for details.
- If scheduled SynclQ policies start running during a failover, it can cause failover job to fail. If your SynclQ schedule has a risk of starting during a planned failover, then the SynclQ schedule should be set to manual before initiating failover.
- 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 SynclQ worker threads from the default to 10 or more.

Back Next

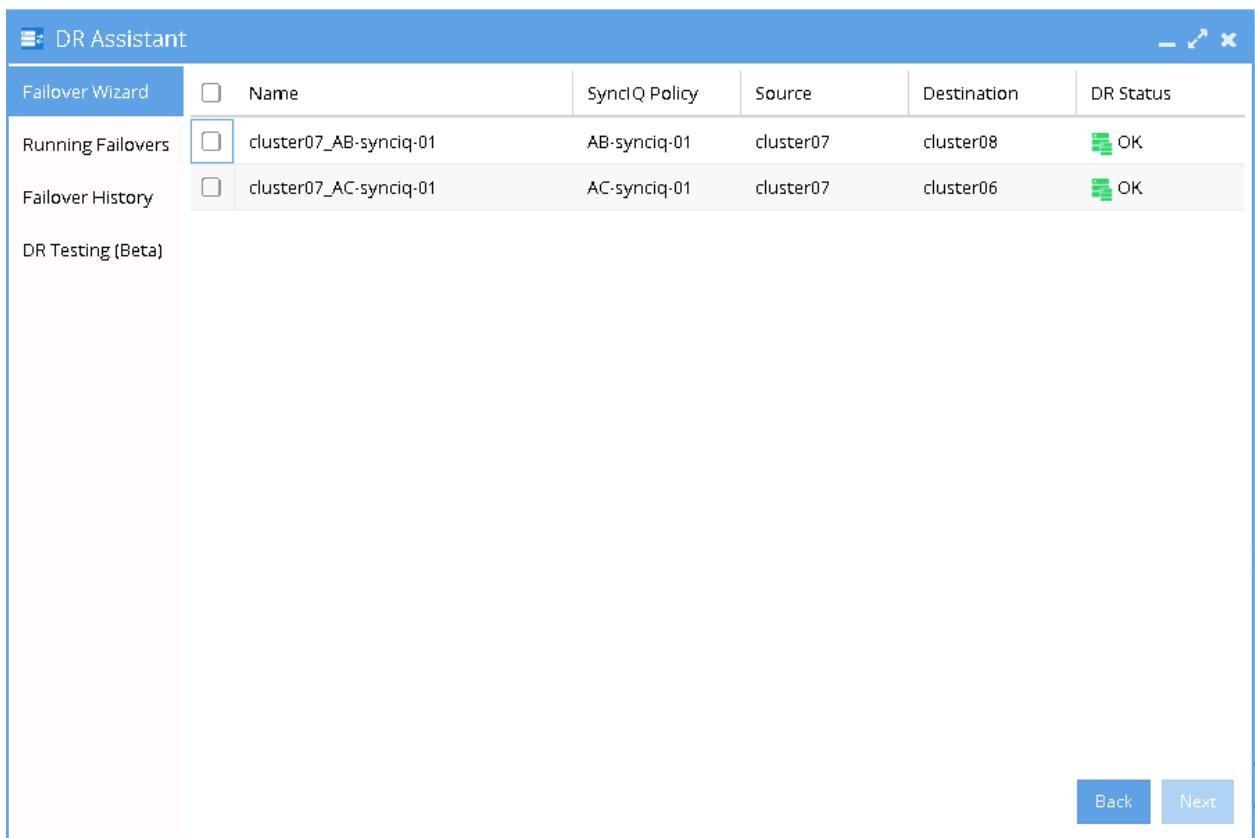
5. Proceed this DFS Mode Failback as per normal. Refer to Eyeglass DFS Mode Failover Guide for details.

DFS Mode Failover from A to C Procedure:

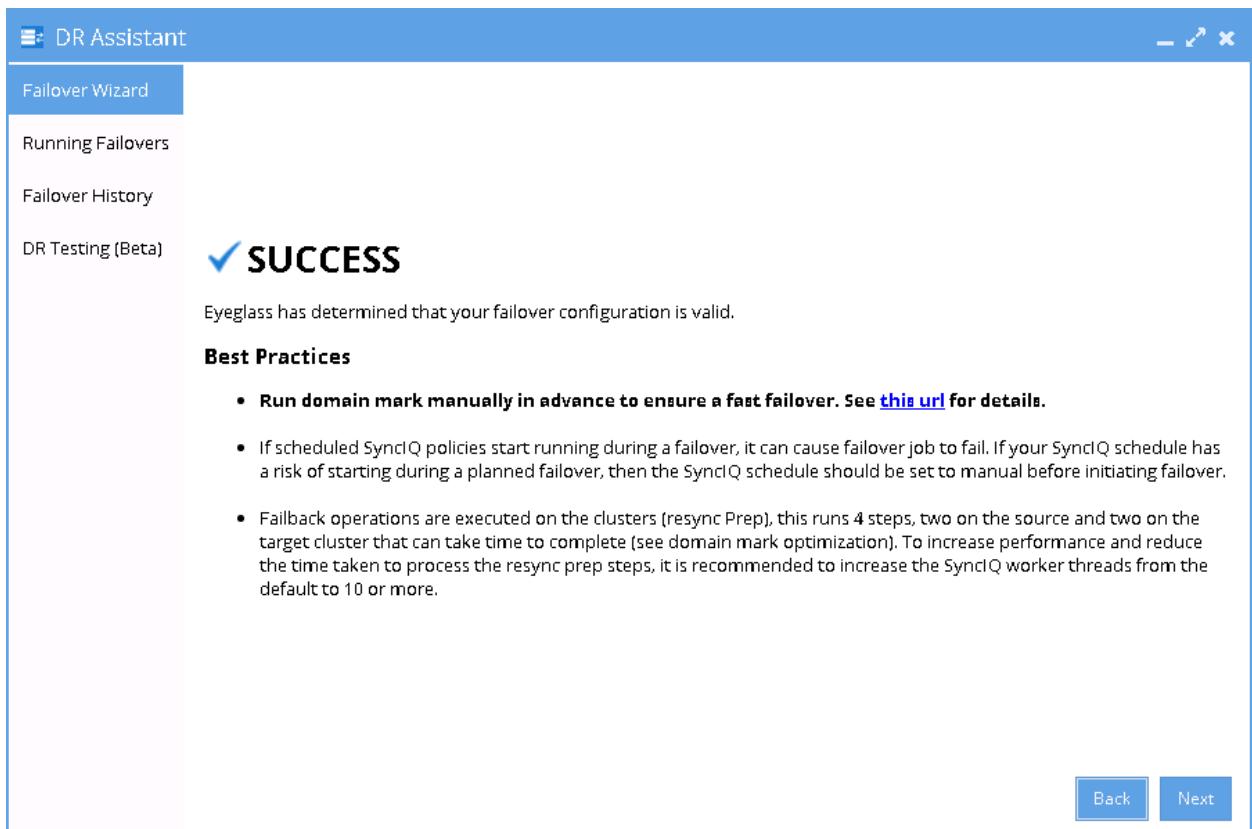
1. Prior to initiate Eyeglass DFS Mode Failover from A to C, **we need to ensure that there is no existing SyncIQ Mirror Policies from B to A. The recovery resync prep step of this Failover A to C will create Mirror Policies from C to A with same Mirror Target Paths as the B to A (Mirror Target Paths are overlaps). This will make the Mirror Policies from C to A unrunnable and the Eyeglass Failover Job will fail.** If there are existing ones, we need to delete them first. Refer to step P1 in the Failover workflow diagrams.
 - a. NOTE: The above step MUST be completed before A to C failover, the order matters since the domain mark will be deleted on cluster A once the step above is completed.
2. Now run a domain mark job on each SyncIQ policy on cluster A. This is a required step since no domain mark exists and will be created during failover process from A to B. The best practise is to run domain mark before failover to ensure the resync prep step does not take a long time to complete. Domain mark can run longer if a the path has a large number of files. NOTE: During the time while domain mark job is running no failover from A to C should be executed until the domain mark job completes on ALL sync polices involved in

the failover. Monitor progress from the PowerScale Cluster jobs UI.

3. Then we can perform Eyeglass DFS Mode Failover as per normal. In DR Assistant Wizard, after we select the source cluster (Cluster A (for this example: cluster07)) the next wizard screen display the list of available Failover options based on Source-Target pairs (A to B or A to C).



3. We need to be careful to select the correct Target Cluster that we want to Failover (A to B or A to C). For this case we want to failover from A to C. Select the AC Source-Target Pair.
4. The next screen will gives validation whether the failover configuration is valid.



5. Proceed this DFS Mode Failover as per normal. Refer to
Eyeglass DFS Mode Failover Guide for details.

DFS Mode Failback from C to A Procedure:

1. We can perform Eyeglass DFS Mode Failback as per normal.
2. In DR Assistant Wizard, **ensure we select the correct source cluster C (name: cluster06). At this stage (After Failover A to C and before Failback from C to A), there are 2 available options to perform as also displayed in the DR Dashboard DFS Readiness. Do not select cluster A (name : cluster07) as the source, as this will lead to Failover from A to B instead.**
 - a. NOTE: The above step MUST be completed before A to C failover, the order matters since the domain mark will

be deleted on cluster A once the step above is completed.

3. Now run a domain mark job on each SyncIQ policy on cluster A. This is a required step since no domain mark exists and will be created during failover process from A to B. The best practise is to run domain mark before failover to ensure the resync prep step does not take a long time to complete. Domain mark can run longer if a the path has a large number of files. NOTE: During the time while domain mark job is running no failover from A to C should be executed until the domain mark job completes on ALL sync policies involved in the failover. Monitor progress from the PowerScale Cluster jobs UI.
4. After we select the correct source cluster (Cluster C (for this example: cluster06)) the next wizard screen will only display the Failback option From Cluster C (cluster06) to Cluster A (cluster07).

DR Assistant

Fallover Wizard

	Name	SynIQ Policy	Source	Destination	DR Status
Running Failovers	cluster06_AC-synciq-01_mirror	AC-synciq-01_m...	cluster06	cluster07	OK

Fallover History

DR Testing (Beta)

Back Next

4. Select the AC mirror policy to fallback. The next screen will give validation whether the failover configuration is valid.

DR Assistant

Fallover Wizard

Running Failovers

Fallover History

DR Testing (Beta)

SUCCESS

Eyeglass has determined that your failover configuration is valid.

Best Practices

- Run domain mark manually in advance to ensure a fast failover. See [this url](#) for details.
- If scheduled SynIQ policies start running during a failover, it can cause failover job to fail. If your SynIQ schedule has a risk of starting during a planned failover, then the SynIQ schedule should be set to manual before initiating failover.
- 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 SynIQ worker threads from the default to 10 or more.

Back Next

5. Proceed this DFS Mode Failback as per normal. Refer to
Eyeglass DFS Mode Failover Guide for details.

© Superna Inc