Table of Contents

1. Eyeglass Start Here First

Home Top

Eyeglass Start Here First

Eyeglass DR Planning and Implementation

- Overview of key functions to monitor for DR Readiness

Install

See: "Eyeglass Isilon Edition Quick Start Guide for Eyeglass Installation"

Upgrades

See: "Upgrade Guide"

DR 101 with Eyeglass

The Big Picture

As the diagram below shows, many moving parts are required to failover an application. This will be the reference diagram used to document best practices for each component that has specific needs for a DR design.
This reference solution will focus on Isilon, Eyeglass and file based DR planning and implementation.

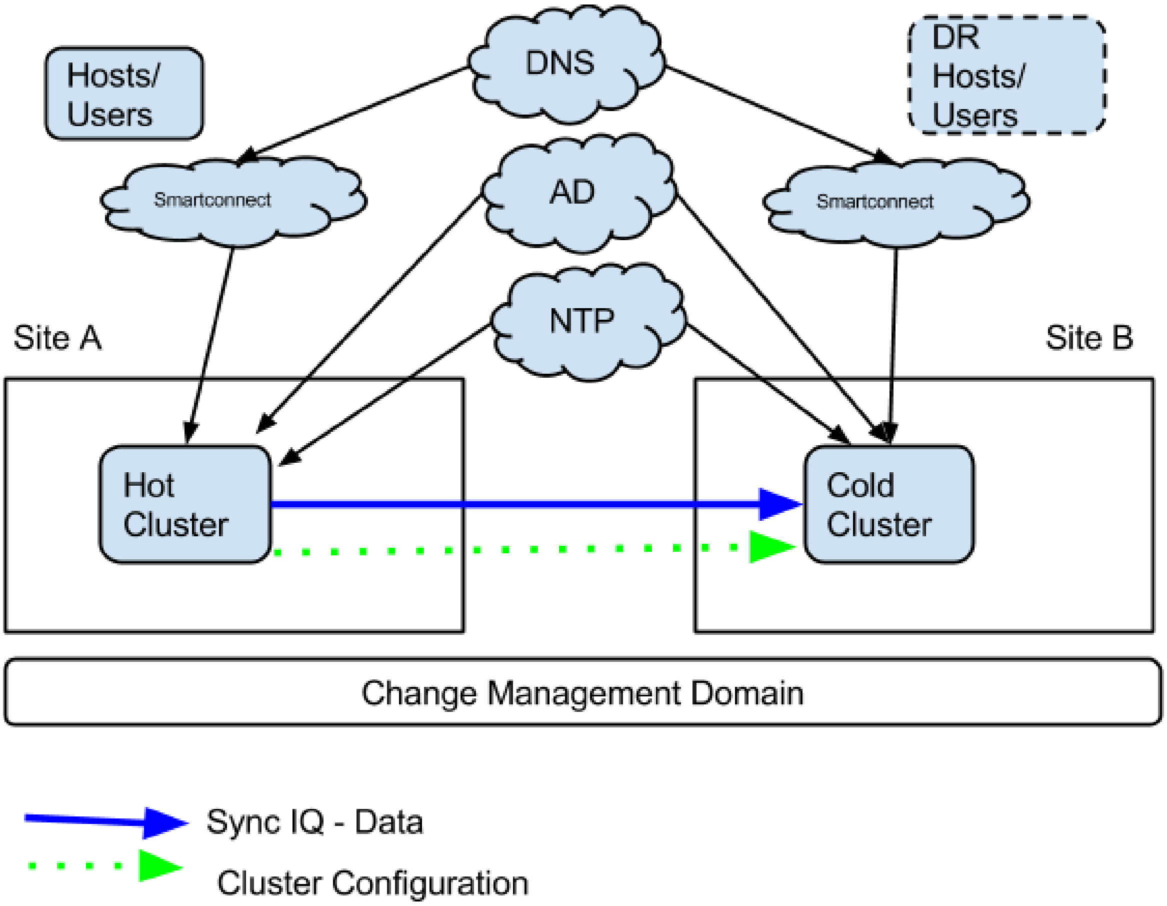

DR with Eyeglass means key decisions for each area identified below under the A, B, C's. Each area requires a decision and plan to implement the various design choices. Please note each Section A, B and C needs to have a plan based on the choices below or your DR failover will not be successful.

This section is designed to guide you through the options (with pros and cons of each) to make a selection based on your environment. You can not miss a section, as this will leave your DR design with gaps that will mean failover will not succeed unless each area has been implemented and tested.

DR assessment services are available for customers to get Superna technical teams to review and recommend a design. A list of services are located here.

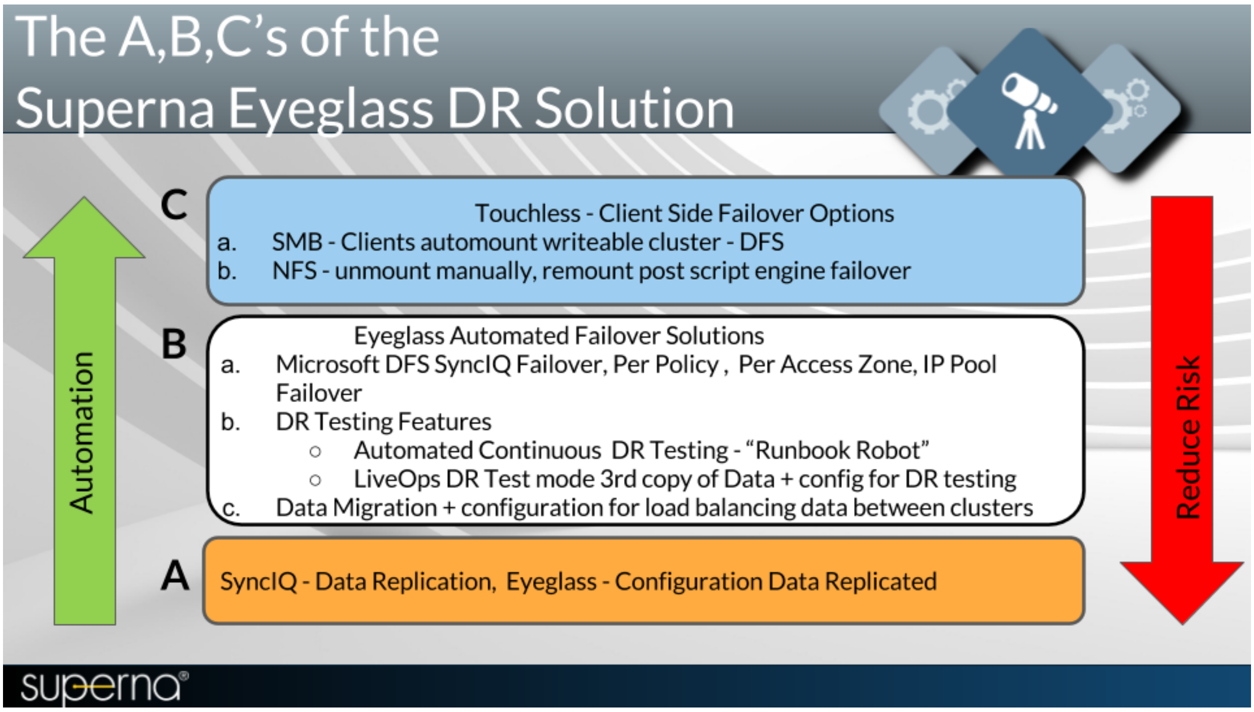

The A,B,C's of the
Superna Eyeglass DR Solution

C

Touchless - Client Side Failover Options

a. SMB - Clients automount writeable cluster - DFS
b. NFS - unmount manually, remount post script engine failover

B

Eyeglass Automated Failover Solutions

a. Microsoft DFS SyncIQ Failover, Per Policy , Per Access Zone, IP Pool Failover
b. DR Testing Features
   - Automated Continuous DR Testing - "Runbook Robot"
   - LiveOps DR Test mode 3rd copy of Data + config for DR testing
c. Data Migration + configuration for load balancing data between clusters

A

SyncIQ - Data Replication, Eyeglass - Configuration Data Replicated

A - Data Replication:

A. SyncIQ - Data Replication, Eyeglass - Configuration Data Replicated

If source configuration data (shares, exports, quotas) is not available you will not be successful at failing over. SyncIQ replicates data, Eyeglass syncs configuration data automatically with no user actions

required.
Review the slides below to ensure the basics of DR are understood before proceeding to Section B and C.

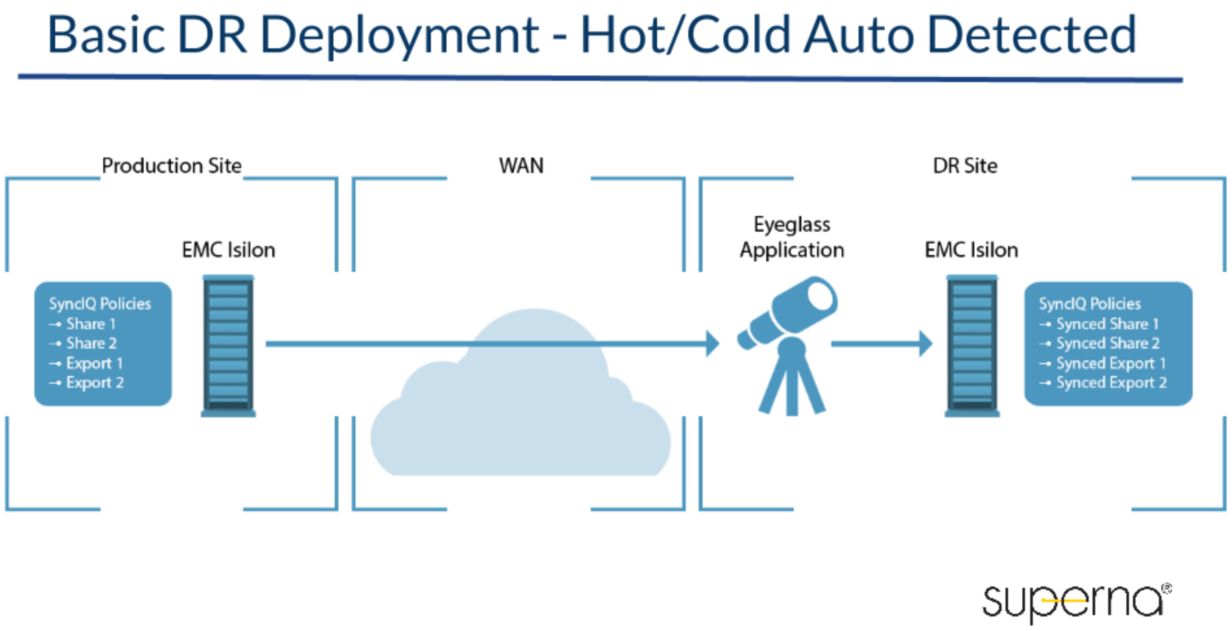

Basic DR Deployment - Hot/Cold Auto Detected

SyncIQ Policies
→ Share 1
→ Share 2
→ Export 1
→ Export 2

SyncIQ Policies
→ Synced Share 1
→ Synced Share 2
→ Synced Export 1
→ Synced Export 2

Data Replicates it one direction and only 1 cluster is writeable.

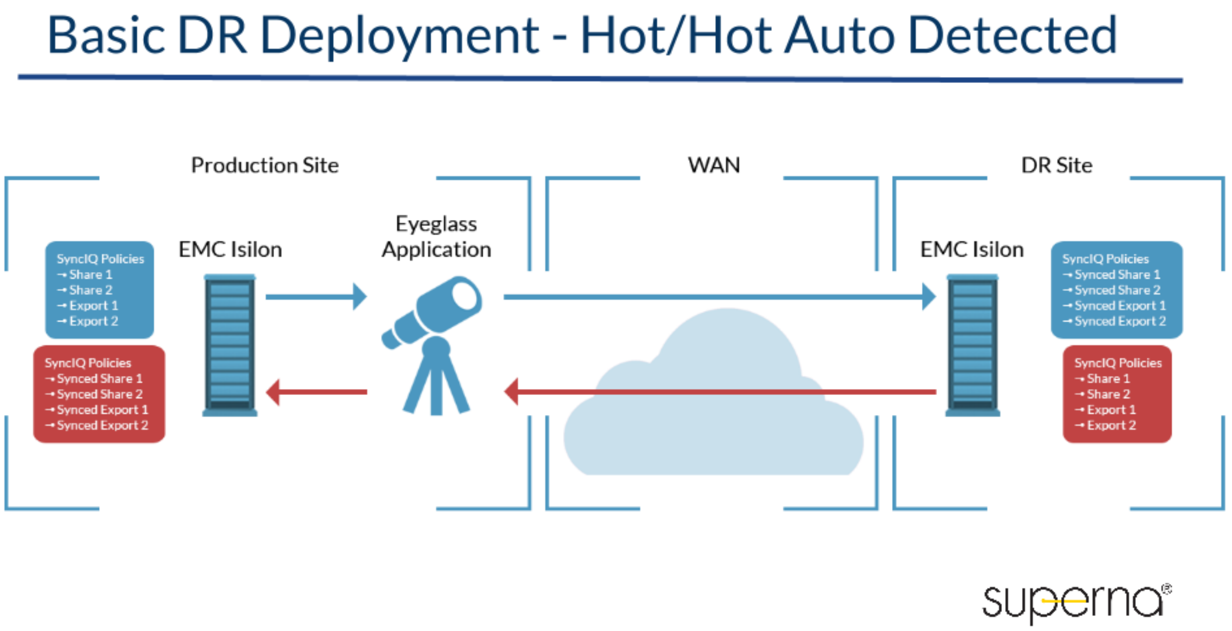

Basic DR Deployment - Hot/Hot Auto Detected

SyncIQ Policies
→ Share 1
→ Share 2
→ Export 1
→ Export 2

SyncIQ Policies
→ Synced Share 1
→ Synced Share 2
→ Synced Export 1
→ Synced Export 2

SyncIQ Policies
→ Synced Share 1
→ Synced Share 2
→ Synced Export 1
→ Synced Export 2

SyncIQ Policies
→ Share 1
→ Share 2
→ Export 1
→ Export 2

Data is replicated in both directions and both clusters are writeable.

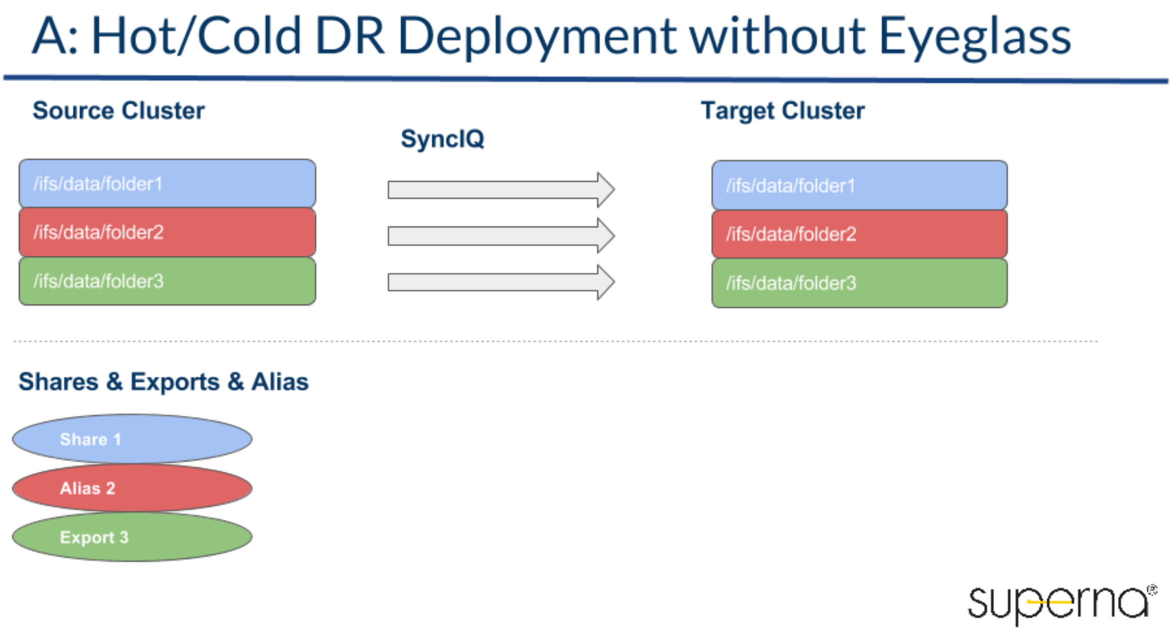

A: Hot/Cold DR Deployment without Eyeglass

Before Eyeglass is deployed configuration data is not protected or analyzed if both sides match.

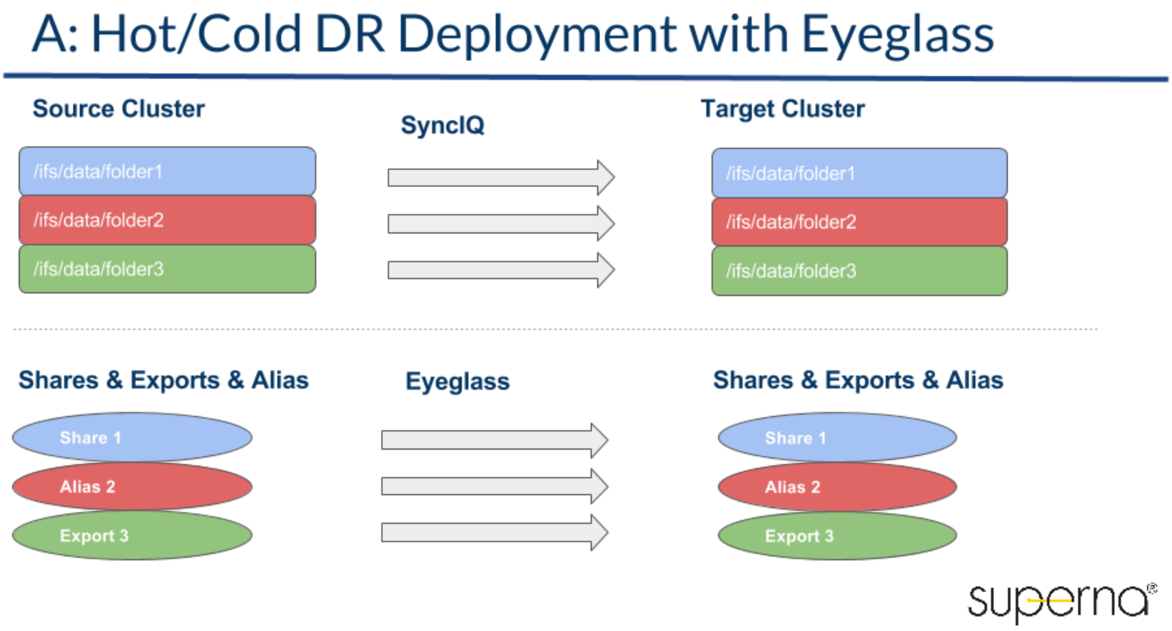

A: Hot/Cold DR Deployment with Eyeglass

After Eyeglass deployment all protected data has configuration data protected automatically, even if new configuration is added or modified or deleted.   Configuration is audited at the attribute level of shares, exports, quotas, aliases.

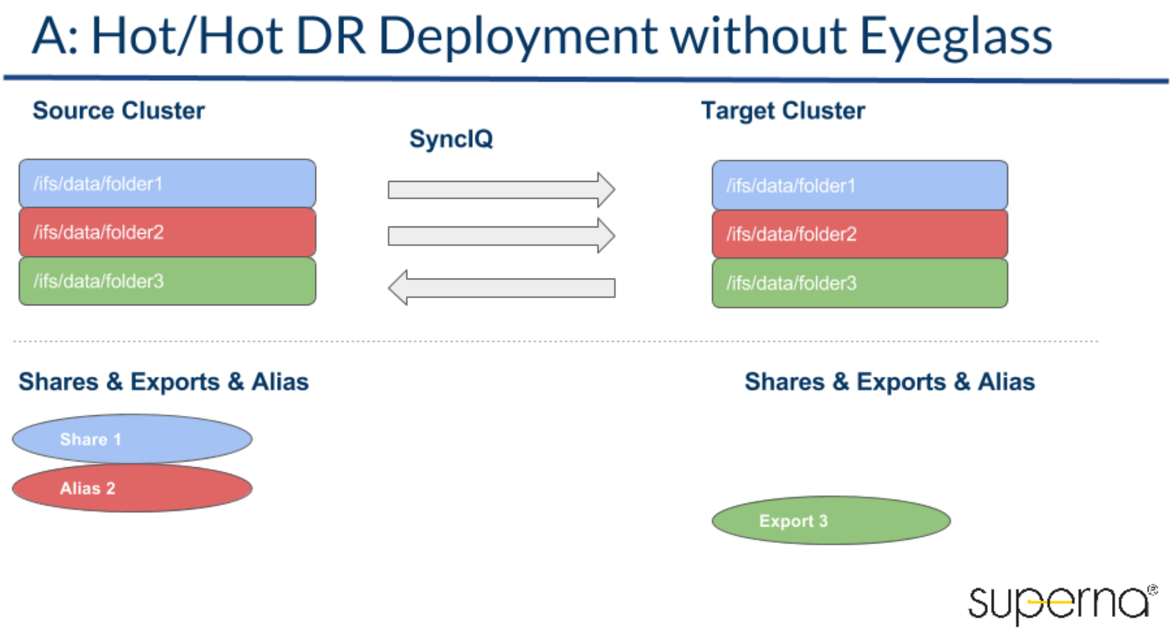

A: Hot/Hot DR Deployment without Eyeglass

Source Cluster

SyncIQ

Target Cluster

/ifs/data/folder1

/ifs/data/folder2

/ifs/data/folder3

/ifs/data/folder1

/ifs/data/folder2

/ifs/data/folder3

Shares & Exports & Alias

Share 1

Alias 2

Shares & Exports & Alias

Export 3

superna

Before Eyeglass deployment each cluster has configuration data that does not exist on the DR cluster (opposite cluster).

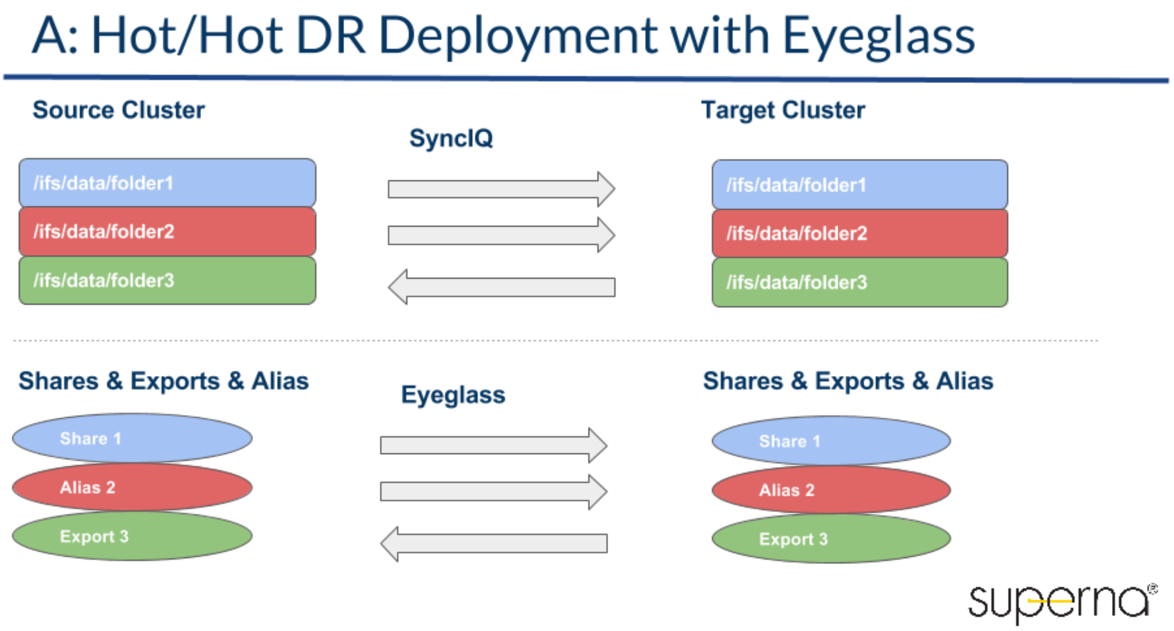

A: Hot/Hot DR Deployment with Eyeglass

Source Cluster

SyncIQ

Target Cluster

/ifs/data/folder1

/ifs/data/folder2

/ifs/data/folder3

/ifs/data/folder1

/ifs/data/folder2

/ifs/data/folder3

Shares & Exports & Alias

Eyeglass

Shares & Exports & Alias

Share 1

Alias 2

Export 3

Share 1

Alias 2

Export 3

superna

After Eyeglass deployment configuration data is synced in both directions to protect data that is replicated in both directions. Eyeglass correctly detects the direction of the replication of data and matches this direction when configuration data is synced.

Access Zone Failover Hot Hot Limitations

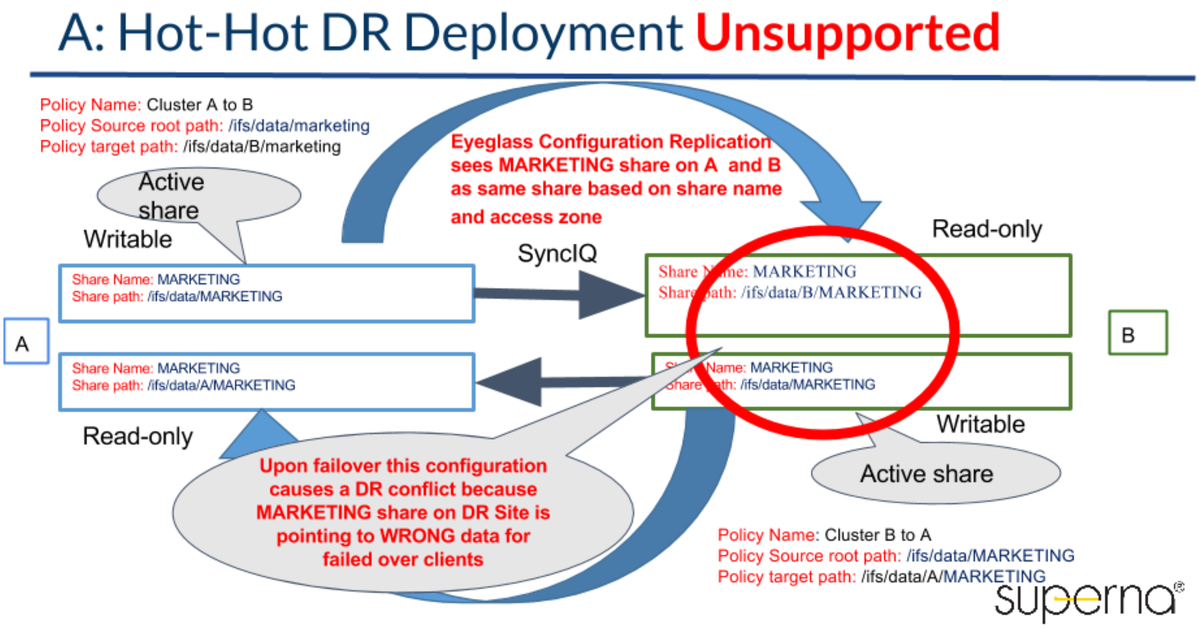

A: Hot-Hot DR Deployment **Unsupported**

Policy Name: Cluster A to B
Policy Source root path: /ifs/data/marketing
Policy target path: /ifs/data/B/marketing

Active share

Writable

Eyeglass Configuration Replication sees MARKETING share on A and B as same share based on share name and access zone

SyncIQ

Read-only

A

Share Name: MARKETING
Share path: /ifs/data/MARKETING

Share Name: MARKETING
Share path: /ifs/data/A/MARKETING

Share Name: MARKETING
Share path: /ifs/data/B/MARKETING

Share Name: MARKETING
Share path: /ifs/data/MARKETING

B

Read-only

Writable

Upon failover this configuration causes a DR conflict because MARKETING share on DR Site is pointing to WRONG data for failed over clients

Active share

Policy Name: Cluster B to A
Policy Source root path: /ifs/data/MARKETING
Policy target path: /ifs/data/A/MARKETING

superna

If you have an Hot - Hot Replication Topology (for data), confirm that you do not have an unsupported share or NFS Alias environment as shown in the above diagram.

IP Pool Failover Access Zone Hot Hot Support

This new feature allows an access zone to have child ip pools, failover independently. This allows hot hot access zone configurations (release 2.0 or later).

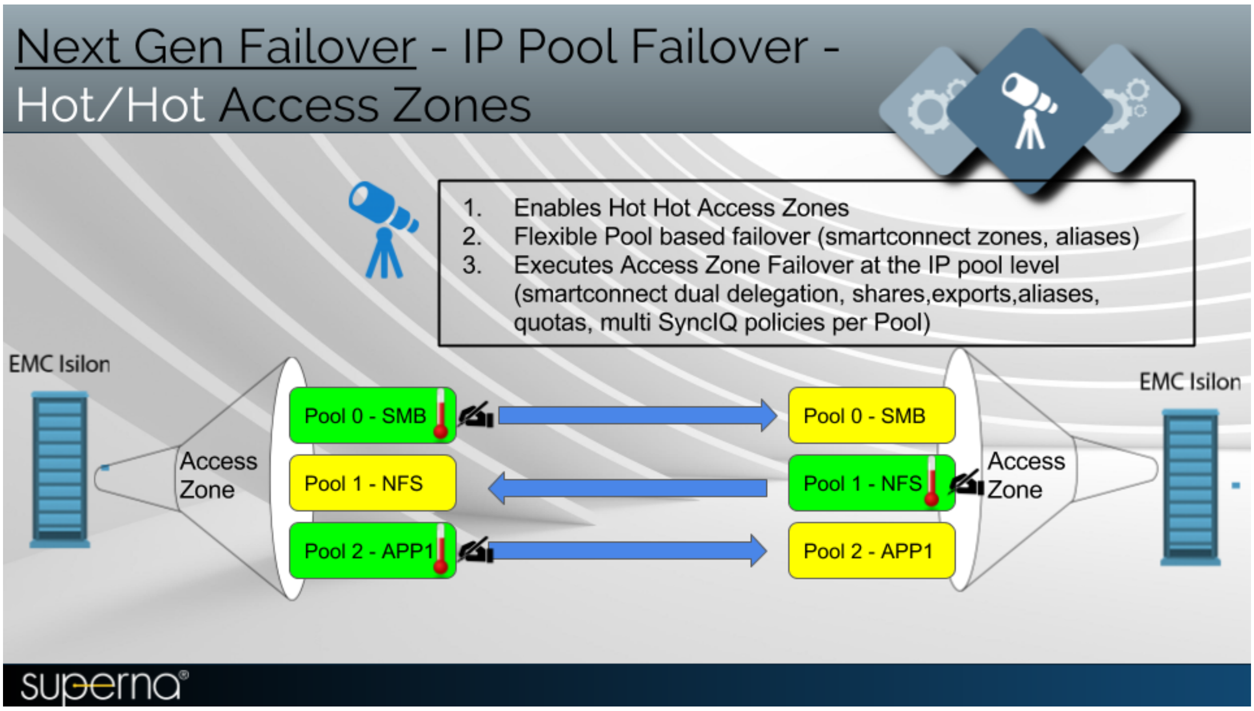

Next Gen Failover - IP Pool Failover - Hot/Hot Access Zones

1. Enables Hot Hot Access Zones
2. Flexible Pool based failover (smartconnect zones, aliases)
3. Executes Access Zone Failover at the IP pool level (smartconnect dual delegation, shares,exports,aliases, quotas, multi SyncIQ policies per Pool)

B - Failover Solutions :

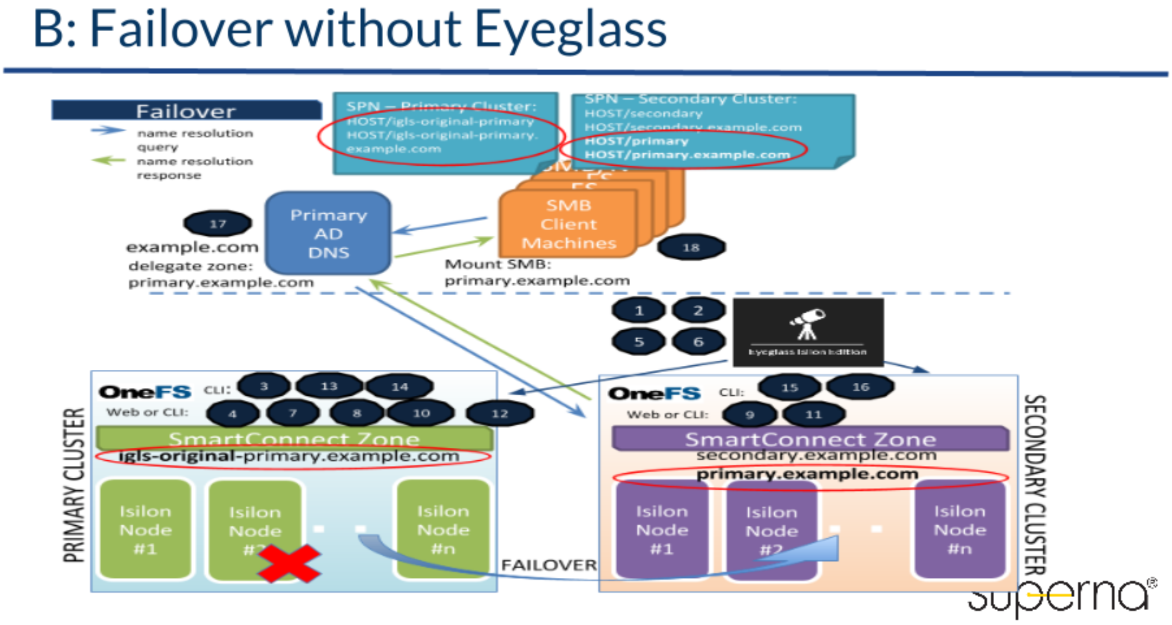

B: Failover without Eyeglass

Failover process is many manual steps, that are order dependant and error prone. The above diagram shows the complexity and number of steps that are need to be performed. This can easily reach > 100 CLI commands on different devices. This introduces risk to document test and execute manual steps.

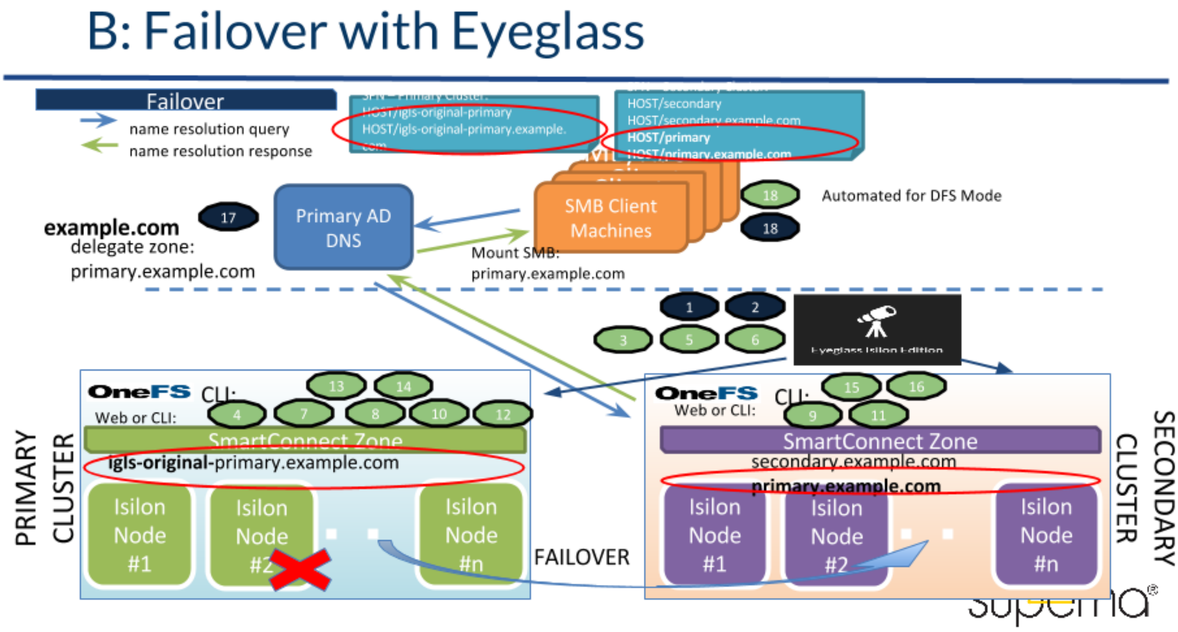

B: Failover with Eyeglass

Failover

name resolution query

name resolution response

HOST/igls-original-primary
HOST/igls-original-primary.example.com

HOST/secondary
HOST/secondary.example.com
HOST/primary
HOST/primary.example.com

example.com
delegate zone:
primary.example.com

Primary AD DNS

SMB Client Machines

Mount SMB:
primary.example.com

Automated for DFS Mode

Eyeglass Isilon Edition

PRIMARY CLUSTER

OneFS CLI:
Web or CLI:

SmartConnect Zone
igls-original-primary.example.com

Isilon Node #1

Isilon Node #2

Isilon Node #n

FAILOVER

OneFS CLI:
Web or CLI:

SmartConnect Zone
secondary.example.com
primary.example.com

Isilon Node #1

Isilon Node #2

Isilon Node #n

SECONDARY CLUSTER

superna

With Eyeglass Automated Failover Solutions, each of these steps is automated, logged and monitored continuously for readiness. The simple easy to view DR Dashboard shows status and alerts to anything that would block failover from being successful (see diagram below).

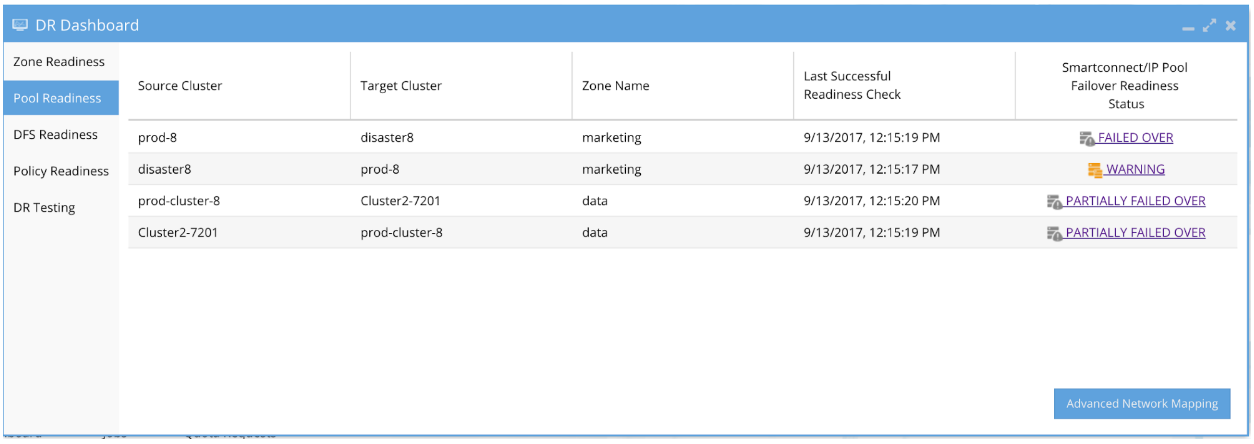

DR Dashboard

Source Cluster	Target Cluster	Zone Name	Last Successful Readiness Check	Smartconnect/IP Pool Failover Readiness Status
prod-8	disaster8	marketing	9/13/2017, 12:15:19 PM	FAILED OVER
disaster8	prod-8	marketing	9/13/2017, 12:15:17 PM	WARNING
prod-cluster-8	Cluster2-7201	data	9/13/2017, 12:15:20 PM	PARTIALLY FAILED OVER
Cluster2-7201	prod-cluster-8	data	9/13/2017, 12:15:19 PM	PARTIALLY FAILED OVER

Failover Option Selection - When to use each type of failover

1. DFS mode (DFS mounted shares)
2. Access zone (multi protocol SMB and NFS and can include DFS data)
    1. IP Pool failover - Enables hot hot data within an access zone, more granular multi protocol failover

3. Per SyncIQ policy (NFS or special cases where pre or post failover scripting can be used)

Note: Fully automated multi site cluster failover option is available for Access Zone or DFS mode

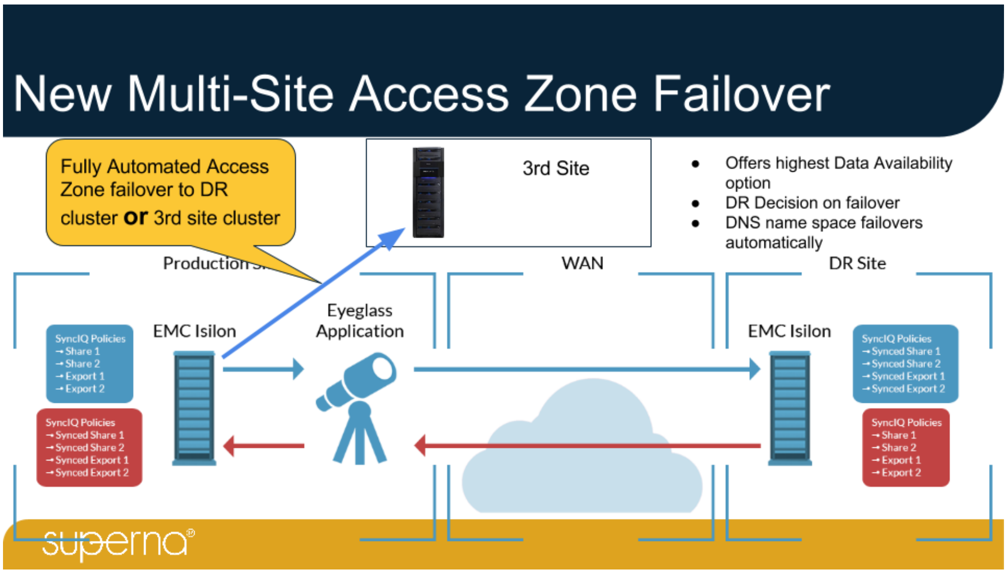

New Multi-Site Access Zone Failover

Fully Automated Access Zone failover to DR cluster **or** 3rd site cluster

- Offers highest Data Availability option
- DR Decision on failover
- DNS name space failovers automatically

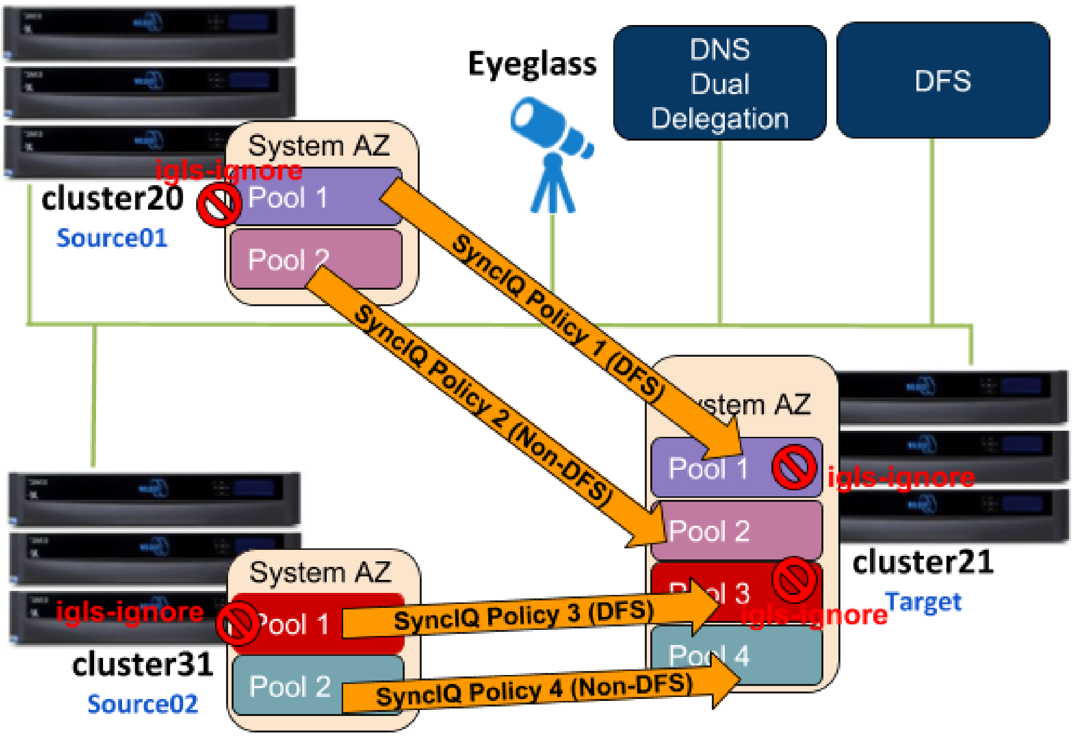

Microsoft DFS SyncIQ Failover :

- **Multi Site Support?**
  - Yes fully automated A to B or A to C failover is supported. See "Multi Site Failover Guide"

- **When to use it?**
  - For SMB data protection when zero touch client failover is needed. Other SMB failover options exist but NONE achieve a touchless seamless failover.
  - All steps are fully automated including the client side.
- **Why use it?**
  - **Superna highly recommends DFS integrated failover for all customers.** Not using DFS? Need to protect data easily? The effort to switch mounts to DFS is worth it.
  - Does not require DNS updates or suffer from DNS/IP mount caching issues on clients and servers.
  - Does not require SPN management during failover.
  - Does not require unmount or re-authentication post failover.
  - Allows SmartConnect Zone names to be different on source and destination clusters.
  - Supports **Active Active** cluster replication with SyncIQ .
- **What you need to know?**
  - It 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.
  - Easy to configure with How to Video "How to setup Microsoft DFS failover Mode with Superna Eyeglass"
  - See also: "Microsoft DFS Mode Failover Guide"
- **Still don’t want to use DFS?**
  - Then you will need to review the other failover options for SMB shares to compare the pros and cons.
- **Estimated knowledge and effort to configure**
  - Active Directory DFS basic knowledge
  - Share create on Isilon
  - SyncIQ setup on Isilon

        - **Overall knowledge and effort to implement: Low**

Per Access Zone Failover OR IP Pool failover

- **Multi Site Support?**
    - Yes fully automated A to B or A to C failover is supported. See "Eyeglass Multi Site Access Zone or DFS Failover Guide"
    - **NOTE: IP Pool failover not supported**
- **When to use it?**
    - For SMB or NFS data protection when it's OK that the **Entire** Access data is acceptable as the failover unit. (It means all data, shares, exports, quotas within the Access Zone failover together).
    - Automates DNS failover with dual delegation of NS records to both clusters. How it works is here: "Geographic Highly Available Storage solution with Eyeglass Access Zone Failover and Dual Delegation"
    - How to Video on configuring Access Zone is here: "Eyeglass Access Zone How To Setup and Configure Overview"
    - Can be setup with Active Active cluster configuration but this requires a unique non overlapping Access Zone on each cluster. Multiple Access Zones can be configured for independent failover.
    - **IP Pool Failover :**
        - **Allows active data within an access zone for hot hot configurations**
        - **More granularity for failover within an access zone**
        - **failover readiness is managed at the pool level versus the Access zone level**
- **Why use it?**
    - It automates more steps including automated DNS, Service Principal Name failover, SmartConnect Zone aliasing and supports SMB, exports, export aliases, quotas and SyncIQ

data failover and pre for failback.

    - Supports pre and post failover scripting to stop applications as needed, allows NFS hosts to have mounts updated using Eyeglass to remotely execute script to unmount and remount (For details see the following video "Eyeglass Scripting How To remount exports automation")
    - Allows different levels of services to applications teams or business units by using IP pool failover
- **What you need to know?**
    - It **does** manage SPN during failover for SMB direct mounts to ensure authentication functions post failover.
    - DFS enabled policies **CAN exist** in the Access Zone **AND** be failed over with the Access Zone.
        - DFS enabled policies, that also protect NFS exports, will require separate SmartConnect Zone name and IP pool for NFS export data access, since this name will failover but DFS smart connect name and IP pool will not. This will allow you to take advantage of Access Zone automation or manual steps / post failover scripting to update NFS client mounts.
        - Also requires an IP pool for DFS mount paths separate from direct mount SMB share SmartConnect Zone names
        - **NOTE: IP Pool failover requires each pool to have SyncIQ policies (one or more) to be assigned to the pool. This means all the file system data protected by these policies plus all smartconnect names and aliases assigned to the pool failover at the same time. Requires knowledge of how the data is mounted.**
    - Access Zone Failover does support **Active Active** failover designs with SyncIQ and Eyeglass.
        - Requires unique non overlapping Access Zones on each cluster

        - SmartConnect Zone Delegation is updated in DNS to both clusters at the same time using dual delegation. Can only be done for one cluster or the other but not both for a given SmartConnect Zone.   Dual Delegation is explained here.
        - Eyeglass does **NOT** support any Smartconnect Zones that have writeable data on the same Access Zone across two **DIFFERENT** clusters.
        - Active Active Clusters **IS** supported when each cluster has separate Access Zones that are replicated between the clusters. **Example cluster A - Access Zone Corp → Cluster B And Cluster B - Access Zone "Marketing" → Cluster A**
        - Script Engine can assist with pre and post failover scripting to remount NFS mounts see guide and video here: "Eyeglass Scripting How To remount exports automation"
        - **WARNING**: DNS/IP caching affects clients ability to remount, as does caching. This can impact remounts and name resolution, unless mount caches are cleared post failover on clients, and DNS servers are updated with new NS record.
    - **IP Pool Failover supports hot hot configuration within an access zone as long as the pools and policies can be mapped.**
        - **Access Zone failover can still be used to failover all IP pools.**
        - **More than one IP pool can be failed over at the same time.**
- **Estimated knowledge and effort to configure**
    - Host side validation of shares and exports post failover (Windows and Linux mount commands).
    - Active Directory ADSI Edit access to Isilon cluster machine accounts for post failover validation (Attributes not visible in

Users and computers snapin). Details here: "How to - Delegation of Cluster Machine Accounts with Active Directory".

- Share create on Isilon.
- Export create on Isilon.
- SyncIQ setup on Isilon.
- Networking with Subnet Service IP, SmartConnect Zones, routing between clients and Isilon clusters at both sites.
- DNS delegation knowledge, name resolution, debugging and testing.
- DNS server knowledge to verify name resolution switches correctly.
- Admin access to DNS.
- **Overall knowledge and effort to implement: High**

Per SyncIQ Policy Failover

- **When to use it?**
  - For SMB or NFS data protection when more control over what portion of the file system is failed over. This failover does require manual steps as per the run book break down located here: : "Eyeglass SyncIQ Policy Failover Guide".
- **Why use it?**
  - It allows for more control over what is failed over and supports quota failover as a single failover workflow.
  - It supports SMB shares and exports underneath the same policy path. It can support Active Active replication, but SmartConnect Zone planning is required to ensure if a policy is failed over, the SmartConnect Zone is also failed over **manually** using aliases (see EMC documentation for details on ISI commands)
  - It can be used with pre and post failover scripting for application specific shutdown and startup. Examples are applications that write data to both clusters for HA can

    benefit from this failover mode.

  -

- **What you need to know?**
  - It does not manage SPN during failover for SMB. SPN will need to be manually managed on the cluster Active Directory machine accounts for source and target clusters. (see EMC documentation for details on ISI commands).
    - Impacts Kerberos authentication if SPN are not managed. Read about what this means here: "How to - Delegation of Cluster Machine Accounts with Active Directory"
  - It requires DNS to be updated for the Smartconnect Zone Delegation to point at the new clusters subnet service IP (manually or with script engine post failover script).
    - **WARNING**: DNS/IP caching affects clients ability to remount as does caching. This can impact remounts and name resolution unless mount caches are cleared post failover on clients and DNS servers are updated with new NS record.
- **Estimated knowledge and effort to configure**
  - Active Directory ADSIEdit and SPN validation.
  - Share create on Isilon.
  - SyncIQ setup on Isilon.
  - DNS delegation knowledge.
  - SPN management ISI commands on Isilon.
  - Smartconnect alias ISI commands on Isilon.
  - **Overall knowledge and effort to implement: Medium to High**

Data Migration With Access Zone Migration

- **When to use it?**
  - For migration data + shares, exports, quotas between Access Zones on the same cluster or different clusters. It

simplies data migration and updates the configuration data paths to match the destination. See: "Access Zone Migration Admin Guide".

- **Why use it?**
    - To move application data and configuration between Access Zones or cluster to balance application across clusters or isolate in an Access Zone.
    - To get better granularity of failover by creating Access Zones for business units and migrate data into the new Access Zones.
- **What you need to know?**
    - During the SyncIQ copy step, where Eyeglass creates a SyncIQ policy, to copy the data changes to the source path will need to be re-synced to catch up changes. The old configuration data is left behind to be manually cleaned up.
        - **WARNING**: Users will still need to be remounting new SmartConnect name in the new Access Zone or cluster. This information is not migrated.
- **Estimated knowledge and effort to configure**
    - Isilon networking.
    - SMB , NFS mounting on clients post migration.
    - **Overall knowledge and effort to implement: Medium**

Runbook Robot Continuous DR Testing

- **When to use it?**
    - Testing DR capabilities on a continuous basis ensures the highest level of DR readiness. Provides DR training without production impact. Fails over and back automatically with Eyeglass acting as a cluster witness and writing data to the cluster to verify successful failover and failback. Tests all aspects of DR readiness. See configuration guide here: "RunBookRobot Admin Guide".
- **Why use it?**

    - Offers highest level of DR Readiness.
    - Enables DR operational training and testing without downtime.
- **What you need to know?**
    - Requires an Access Zone dedicated to DR testing.
    - Users NFS export to write data but can have shares created for testing as well.
    - Test data can be copied into the Access Zone for simulated DR testing.
- **Estimated knowledge and effort to configure**
    - Isilon networking.
    - DR concepts.
    - SMB , NFS mounting on clients post failover tests.
    - Access Zone failover guide knowledge.
    - **Overall knowledge and effort to implement: Medium**

LiveOPS DR Testing Mode

- **When to use it?**
    - Enables test and development testing on live copy of production data in an isolated Access Zone. Clones shares and exports from production into a test Access Zone. Provides writeable copy of production data using DR as source to make 3rd copy on the DR cluster.
    - Customers can complete DR procedure testing on live data without impact to production data.
    - Security is mirror of production.
    - Allows a subset of the DR data to be copied into DR testing access zones.
    - Multiple DR testing Access Zones supported for business units.
    - Near real-time copy lags by 5 minutes from production DR copy.

    - Changes to the files system are backed out after disabling DR test mode.
    - Dedupe license on Isilon can reduce on disk overlap to 0 bytes over a period of time, allowing a permanent 3rd copy for testing at any time.
    - See configuration guide here: "Live Ops - Continuous Operations".
- **Why use it?**
    - Offers highest risk free DR testing and devops on live copy of production for application upgrade testing or development, or DR procedure testing and compliance.
    - Enables DR operational training and testing without downtime.
- **What you need to know?**
    - Requires extra disk space or Dedupe license.
    - Requires new DNS mount namespace for DR testing.
- **Estimated knowledge and effort to configure**
    - Isilon networking.
    - DR concepts.
    - SMB , NFS mounting on clients post failover tests.
    - Access Zone failover guide knowledge.
    - **Overall knowledge and effort to implement: Medium**

C - Touchless - Client Side Failover Option

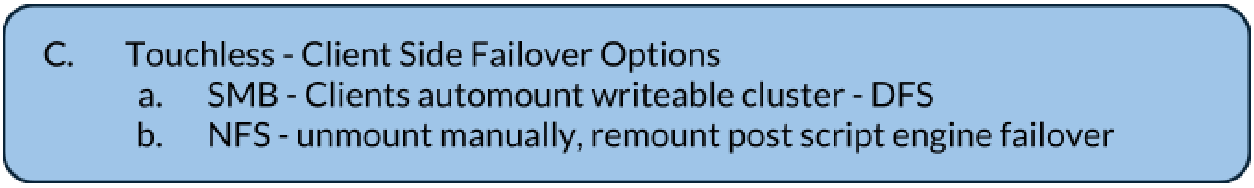

C. Touchless - Client Side Failover Options
  a. SMB - Clients automount writeable cluster - DFS
  b. NFS - unmount manually, remount post script engine failover

What is client side failover?

This the goal of failover is to have zero steps on clients post storage layer failover, and minimize data loss window under all conditions. This section outlines the choices and which option reduces the steps and data loss window the most with the least complexity and dependencies.

What are the client side options?

- **Eyeglass Integrated Microsoft DFS mode**
    - Provides the most automated solution for client side auto failover detection, and avoids all manual steps on the client side.
    - **Protocols: SMB Shares only.**
    - No DNS change required.
    - No reauthentication.
    - No unmount or remount required.
    - No SPN management required.
    - No SmartConnect Zone changes.
    - **Available with: DFS mode in Eyeglass.**
    - **Least Steps Choice.**
    - **Smallest Data Loss Window.**
- **Eyeglass automated configuration sync and Access Zone failover with automatic DNS Failover**
    - This option uses Eyeglass to pre-sync all changes to shares and exports to the target cluster. During failover name resolution switches when SmartConnect Zones failover using dual DNS delegation for SmartConnect Zone names.
    - **Protocols: SMB and NFS exports**
    - Requires remount of SMB and NFS exports to resolve SmartConnect name to the target cluster, and remount and reauthenticate.
    - Requires SPN management during failover to ensure Kerberos tickets are issued against the mount request, based on the target cluster AD machine account. Requires

SPN correctly registered against the target cluster machine account. (**Access Zone automates this step**)

- Requires IP pool to IP pool failover planning and mapping with Eyeglass Access Zone failover. (Access Zone setup guide explains how to use igls-xxx hints to map IP pools for failover)
- **Available with**: **Per SyncIQ failover Or Access Zone Failover**
- **Higher Data Loss Window exposure due to remount requirement**

- **Networking Solution:**
    - **Layer 2 failover and SmartConnect Zone Aliasing -** some customers have Layer 2 network path between data centers. This allows a different solution on failover that moves the Subnet Service IP from source cluster to target cluster during failover by editing the networking on the target cluster.
        - **Pros:**
            - Preserve DNS entries on failover so that no CNAME updates are needed. (**this issue is solved with Dual Delegation**)
            - Eyeglass Access Zone failover can be used to assist with SmartConnect Zone failover automation including DNS
            - Only suitable for NFS, due to SMB issues with failing over SPN's for Kerberos authentication means it's a poor choice to protect SMB data. This requires manual steps as SmartConnect names and related SPN's needed to be failover over to the other cluster computer object in AD. This moves the problem from DNS to active directory.
        - **Cons**:

- **No**t supported by Eyeglass workflow to move Subnet Service IP from Cluster Source to Target, which means a manual step is required. Dual Delegation solves the DNS issue simply with one time NS record added per SmartConnect Zone.
- **Cannot** pre-create SmartConnect Access Zones on target since it will create SPN conflicts in AD.
- Does **not** solve SPN management at failover issues, and still requires Eyeglass Access Zone failover feature to be implemented.
- **WARNING**: DNS/IP caching affects clients ability to remount. This can impact remounts and name resolution, unless mount caches are cleared post failover on clients, and DNS servers are updated with new NS record.
- NFS clients using the SAME FQDN, will **NOT** auto mount the exports after Layer 2 failover. Stale Mounts requiring a remount.
- NFS differences between clusters also causes stale mount and requires a remount to complete the failover.
- If Layer 2 to the DNS server IP address is present, ARP cache updates on clients and switch ports will **IMPACT** name resolution to the same IP address, and may require ARP cache flush on Ethernet switches.
- **Overall Assessment**: **Eliminates DNS IP delegation update BUT this is not necessary, Dual Delegation solves this without manual effort. This solution:**
  - **Did NOT avoid client side unmount and remount,**
  - **Did NOT simplify steps, and introduces**

manual steps to Add subnet IP, and then remove from source cluster during failover.

- **Recommendation: Not Recommended.**

- **NFS Host side unmount and remount automation planning**
    - The unmount command can be issued with -f or force option to avoid stale mounts that occur during either a controlled failover or uncontrolled scenario. Hosts with open files need to unmount and remount the new cluster post failover.
    - Eyeglass a offers post failover script engine. This will allow host side automation to run and unmount and remount Isilon storage by policy, with variables that can be used to select hosts that use the policy data being failed over. This is covered in more detail in the Script engine section of Eyeglass Isilon Edition Administration Guide

1. DR Readiness

What it means to be ready and how to configure and Monitor Readiness Status

Eyeglass not only syncs config data, tracks configuration changes, and executes automated failover, it offers a dashboard that monitors DR status and alarms on any changes detected that would impact the ability to failover.
Daily monitoring of clusters is no longer required with Eyeglass acting as a cluster Witness and monitoring key cluster events and running audits against the clusters to verify everything is ready for failover.

Overview of key functions to monitor for DR Readiness

- **SyncIQ Jobs monitoring:**
    - Eyeglass monitors all SyncIQ jobs and alarms if any are in failed state.
- **Configuration data sync and Audits:**
    - Eyeglass syncs shares, exports and stores quotas for

failover. If sync jobs fail for this metadata, an alarm is raised against the DR state for the data being shared.

- **RPO Summary reports:**
  - Eyeglass tracks SyncIQ averages over time and flags any SyncIQ jobs operating outside their normal 30 day operating range along with trending and reporting on current data loss window RPO per cluster.
- **Cluster Witness and DR Testing:**
  - Eyeglass can setup a test Access Zone, create test data on the cluster to test failover and failback operations, and test the data path to the cluster while orchestrating failover and failback on a daily bases without impacting production access zones.
- **LiveOps Dashboard:**
  - New in 1.6.3 and later the ability to sync Snapshot schedules and dedupe settings allows Eyeglass to sync cluster configuration needed to manage recovery points on either cluster.
  - Post failover, snapshots are available as recovery points and data is pre-deduped without any waiting
  - The new dashboard makes it easy to see if a planned failover can be scheduled with all snapshots and dedupe synced in advance
- **Cluster Storage monitor:**
  - Migrate data config and access zone aware same cluster inter access zone or between clusters
  - Monitor cluster usage centrally
  - Monitor hardware health centrally
  - Search for shares and exports across cluster and see disk utilization
  - Self Serve Quota portal (automated workflow to manage quotas)
  - Cluster Storage monitor information can be found here:

"Eyeglass Cluster Storage Monitor"

- **CA UIM Probe:**
    - Product information can be found here: Eyeglass Isilon Probe for CA UIM

2. Advanced topics

- **Eyeglass REST API**
    - API Guide
- **Post failover scripting**
    - Script Engine Overview
- **Advanced Configuration**
    - Eyeglass CLI Commands
- **Solutions Guides**
    - Warm Standby appliances
    - Active Active appliances
    -