Unlocking Resilience: How a Wall Street Bank Created a Secure and Automated 4-Site Replication Solution

Rudy Dussault, Vice President, Enterprise Z Storage Manager
John Wolfgang, Flash Storage Architect

Agenda

01 Our 4 Site Enterprise Z Environment

02 Our Desired Capabilities

03 CSM 4 Site Replication Enhanced Session

04 Automate Data Center Operations & Reduce Manual Tasks

05 Strengthen Security and Access Controls

OUR 4 SITE ENTERPRISE Z ENVIRONMENT

An Overview of our Mainframe Environment

SHARE
EDUCATE • NETWORK • INFLUENCE

Environment A

PROD
647 TB
12 DS8950F
4 IBM Z

Non-Prod
647 TB
8 DS8950F
4 IBM Z

Environment B

PROD
484 TB
8 DS8950F
4 IBM Z

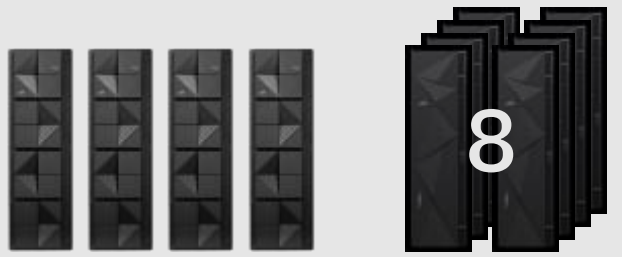

Non-Prod
161 TB
4 DS8950F
4 IBM Z

An Overview of our Mainframe Environment

Site 1

Environment A

PROD:
3 DS8950F
1 IBM Z

Non-Prod:
2 DS8950F
1 IBM Z

Environment B

PROD:
2 DS8950F
1 IBM Z

Non-Prod:
1 DS8950F
1 IBM Z

Site 2

Environment A

PROD:
3 DS8950F
1 IBM Z

Non-Prod:
2 DS8950F
1 IBM Z

Environment B

PROD:
2 DS8950F
1 IBM Z

Non-Prod:
1 DS8950F
1 IBM Z

Site 3

Environment A

PROD:
3 DS8950F
1 IBM Z

Non-Prod:
2 DS8950F
1 IBM Z

Environment B

PROD:
2 DS8950F
1 IBM Z

Non-Prod:
1 DS8950F
1 IBM Z

Site 4

Environment A

PROD:
3 DS8950F
1 IBM Z

Non-Prod:
2 DS8950F
1 IBM Z

Environment B

PROD:
2 DS8950F
1 IBM Z

Non-Prod:
1 DS8950F
1 IBM Z

OUR DESIRED CAPABILITIES

Desired Storage Capabilities

- Production
  - Low latency flash performance, high reliability
- High-Availability
  - Co-located synchronous copy
  - FICON-addressed for seamless workload transition
- Disaster Recovery
  - Out-of-region asynchronous copy with low RPO
- Cyber Resiliency
  - Out-of-region copy for backups
  - Protect from cyber attack or logical corruption
  - Space efficient and no/low effect on DR RPO
- Site Swap Capability
  - Fully-redundant and functionally-identical Data Centers
  - Rapid and repeatable site swaps via automation

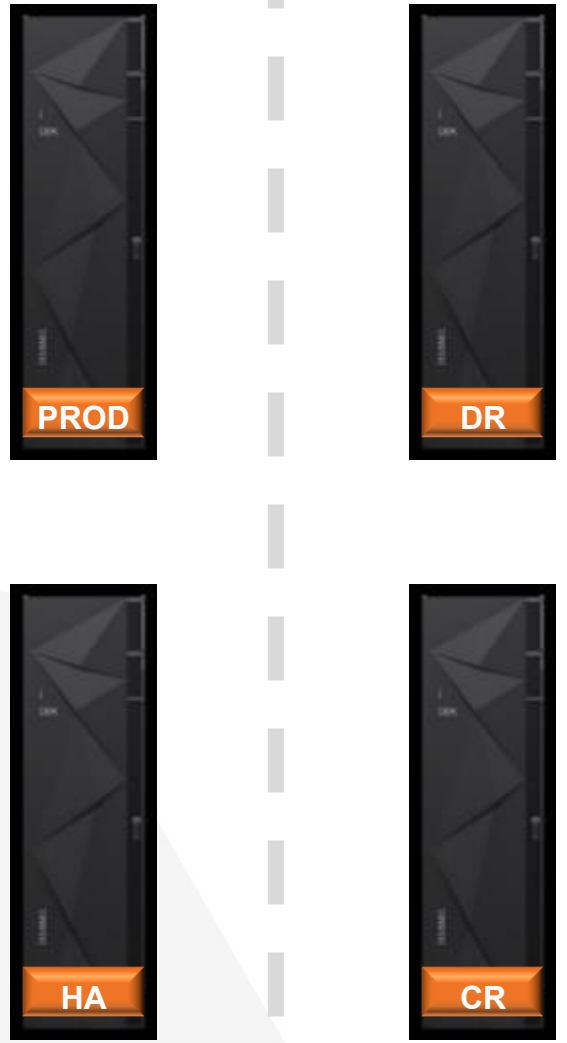

CSM 4 SITE REPLICATION ENHANCED SESSION

CSM 4 Site Replication Enhanced Session

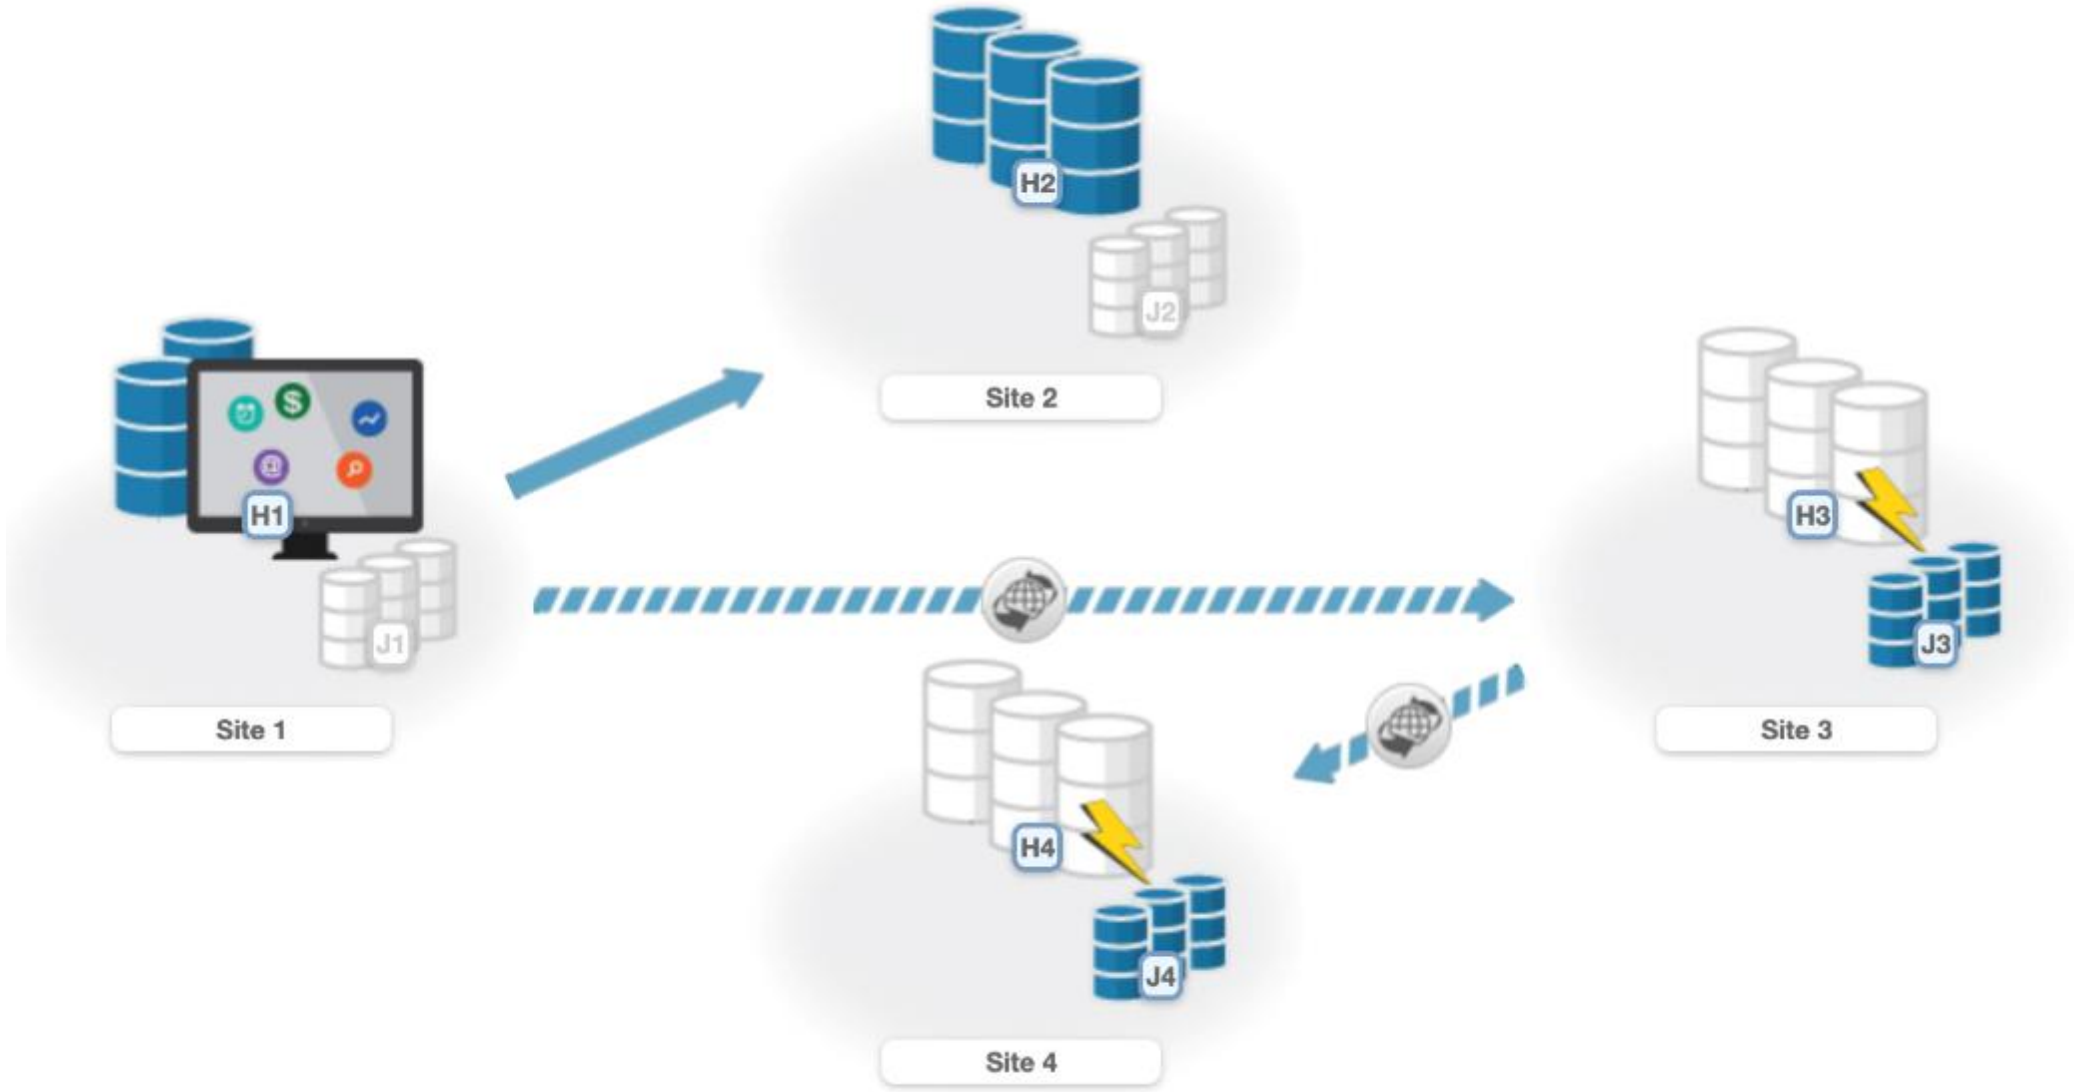

CSM 4 Site Replication Enhanced Session

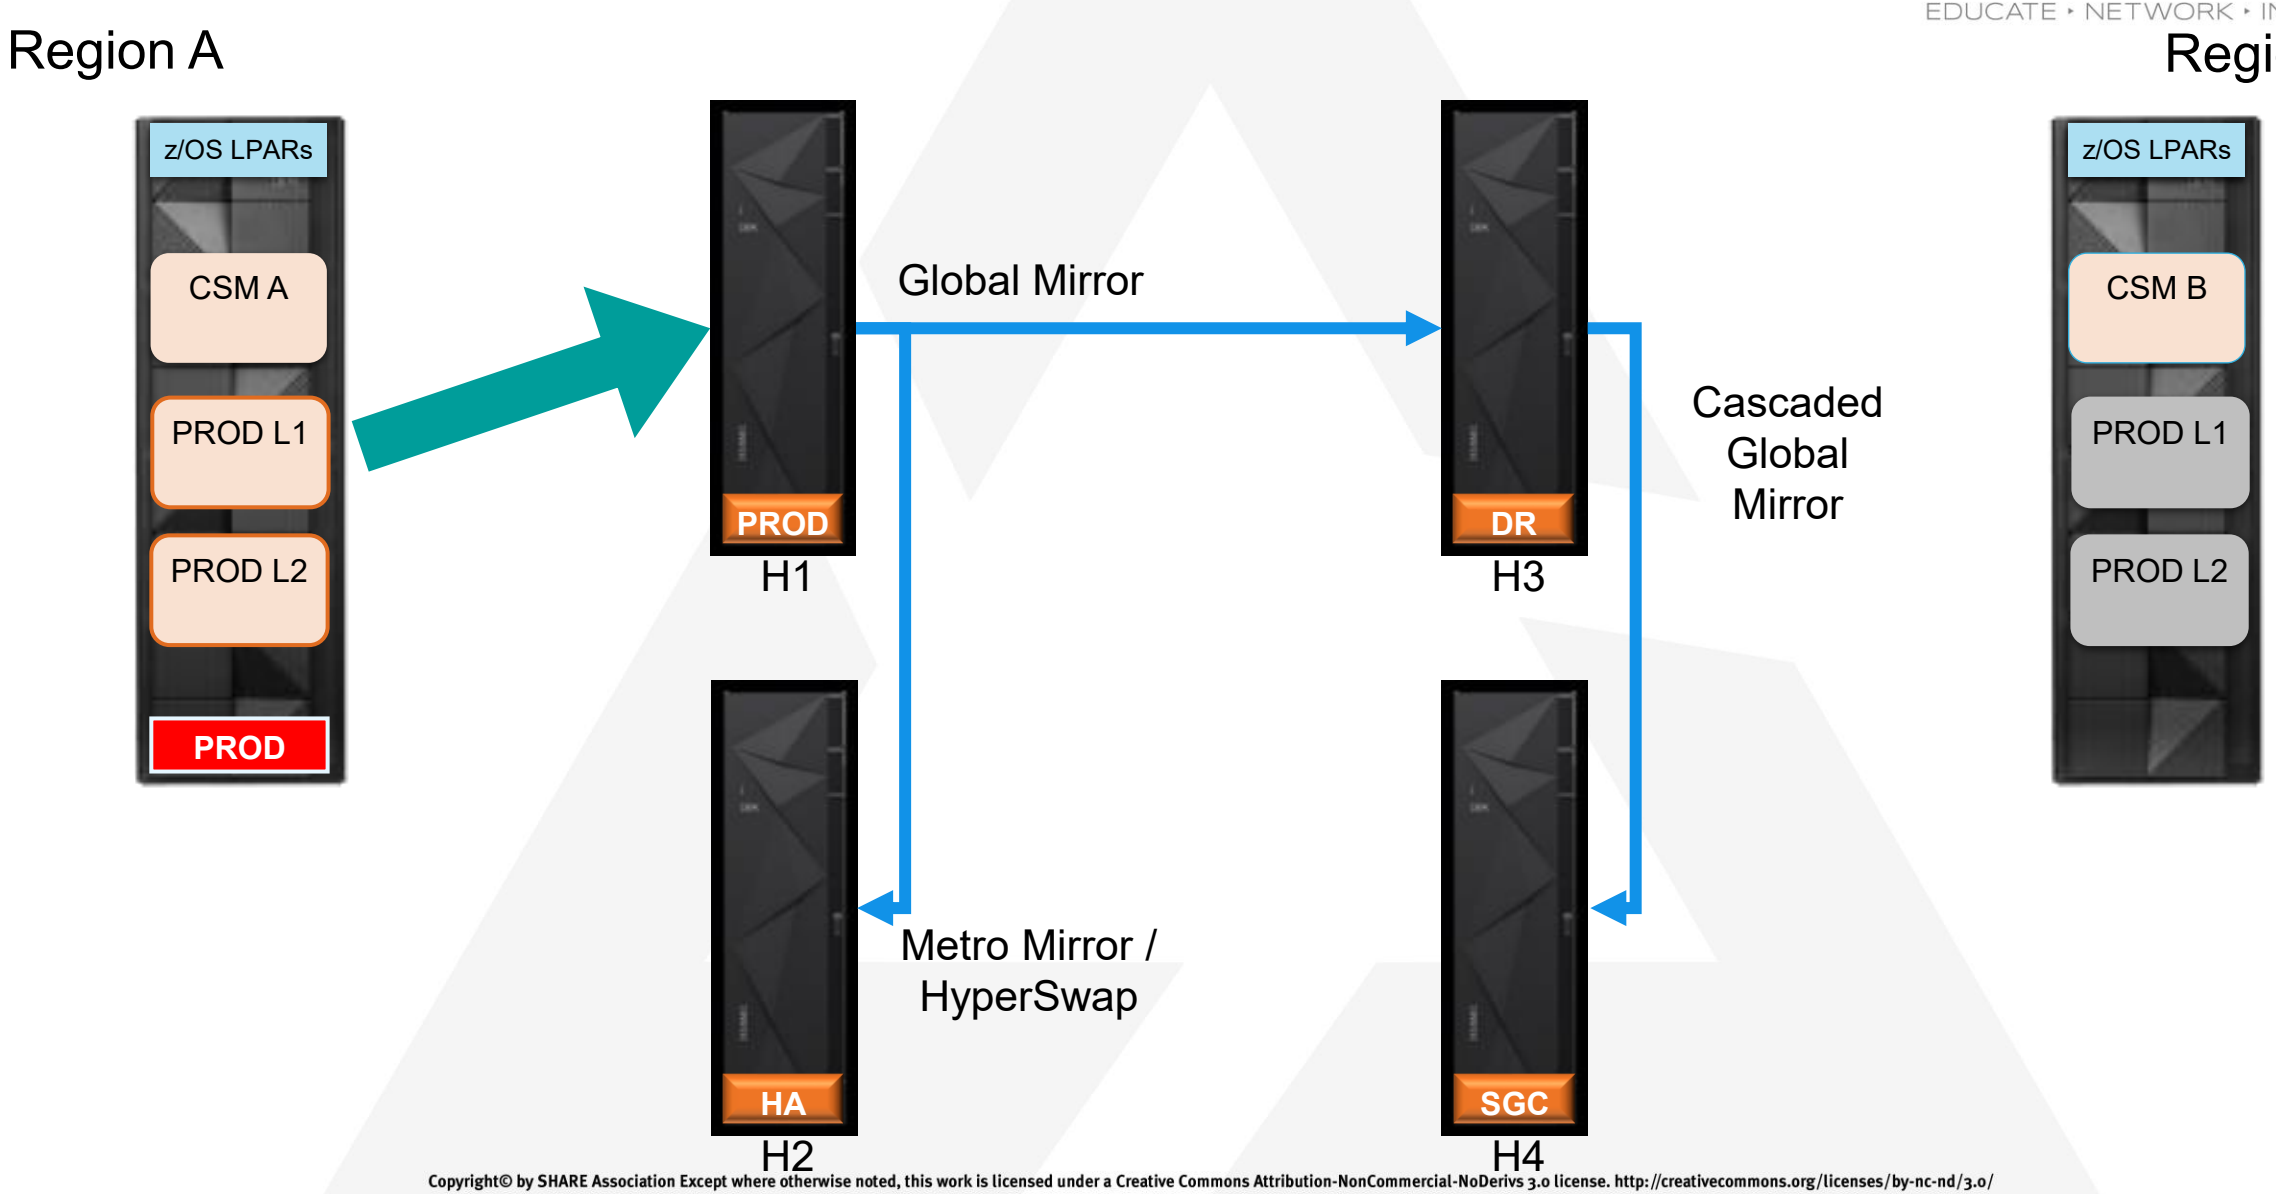

CSM 4 Site Replication Enhanced Site Swap

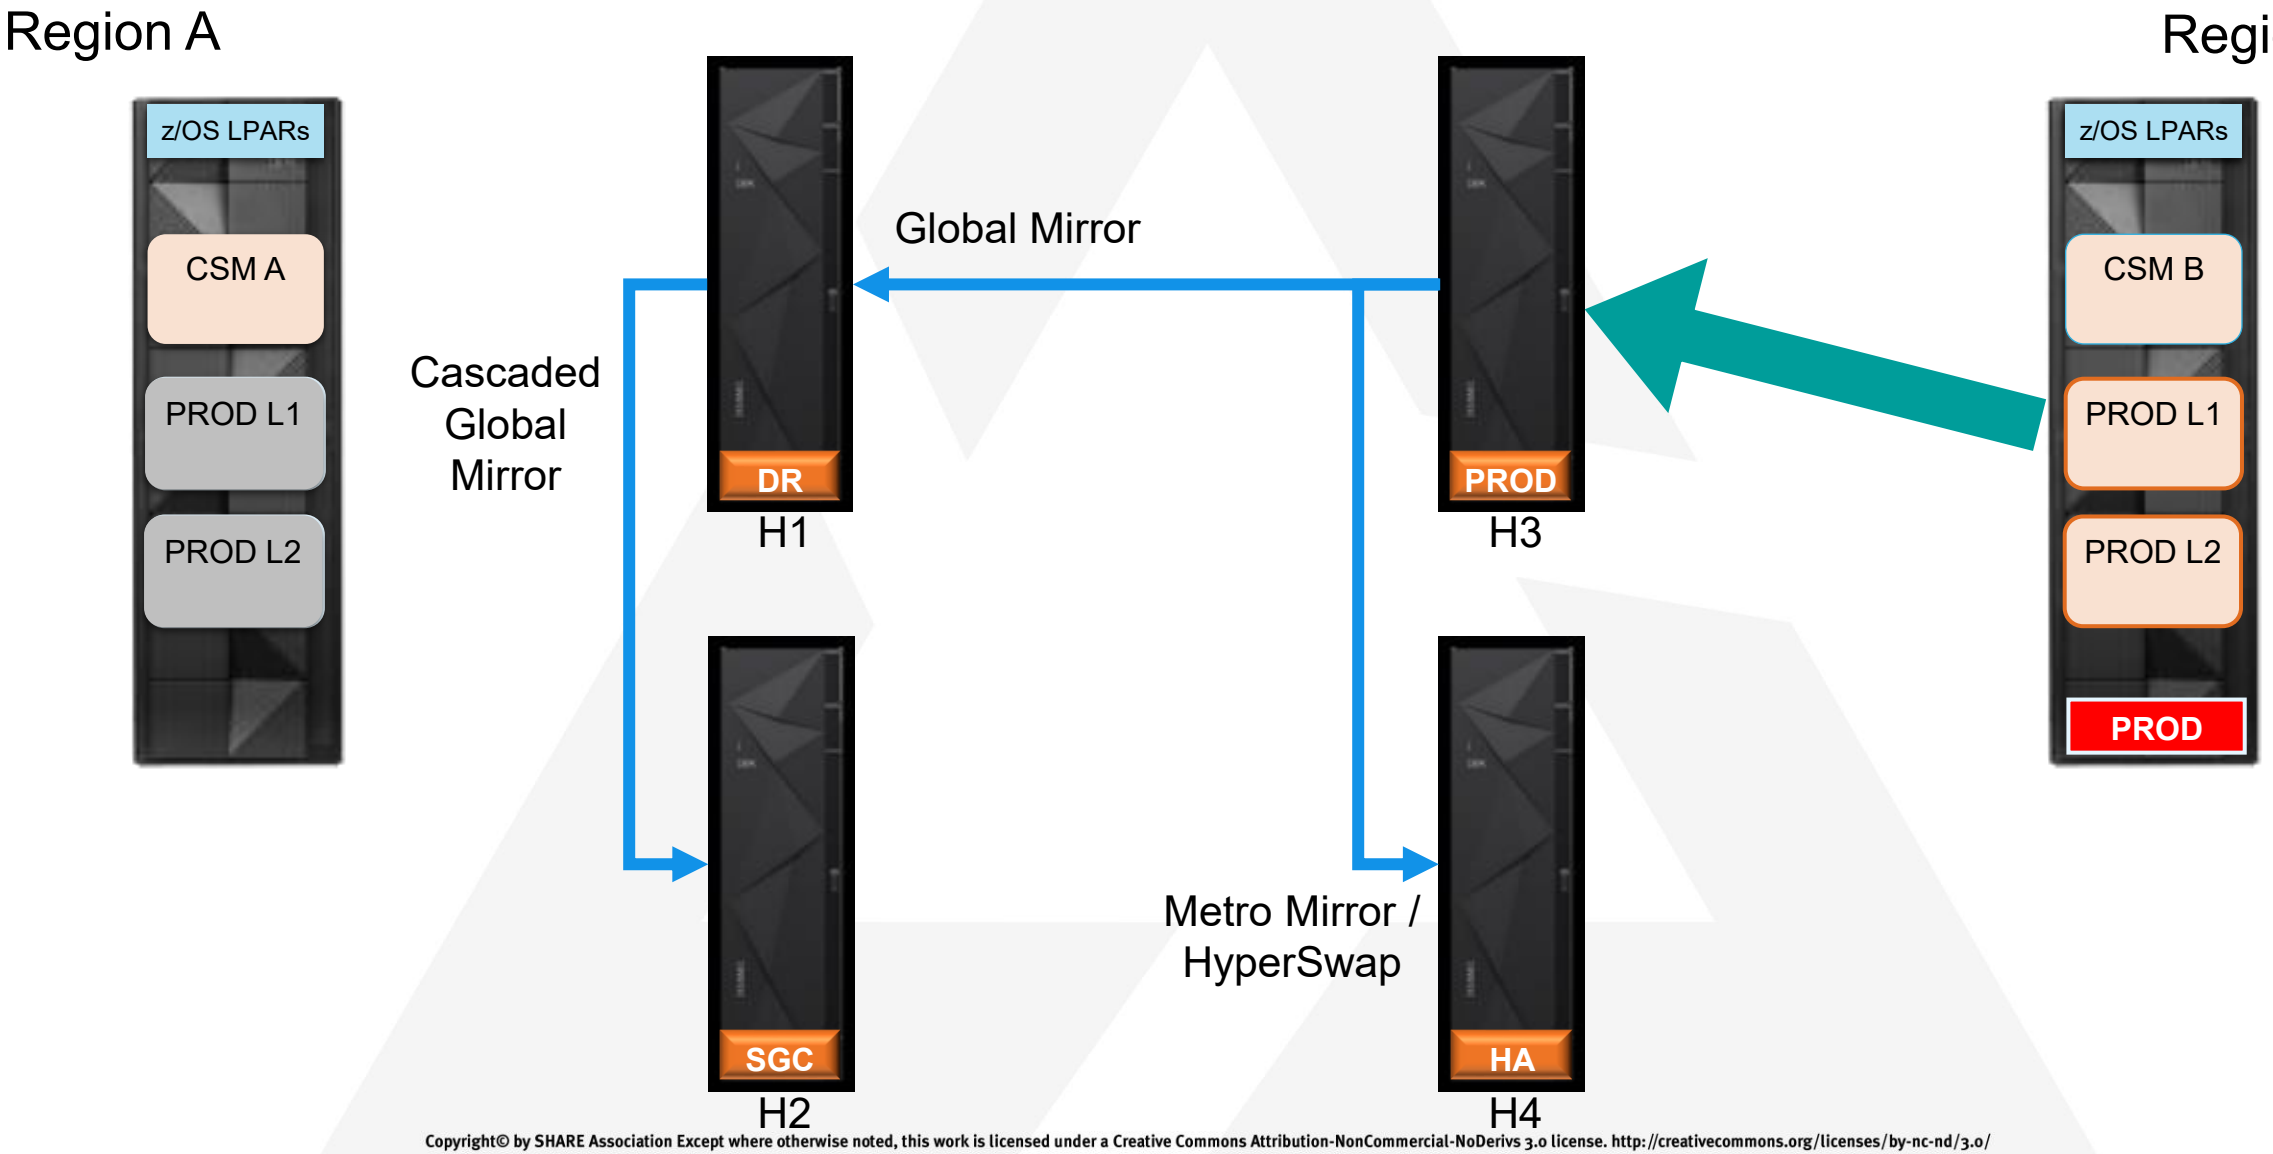

All DS8950Fs are Identical for Maximum Flexibility

- Site Swap capability requires each DS8950F to assume different roles
- Users typically configure PROD systems differently from SGC systems
- We want all systems to have capability to assume **ANY** role

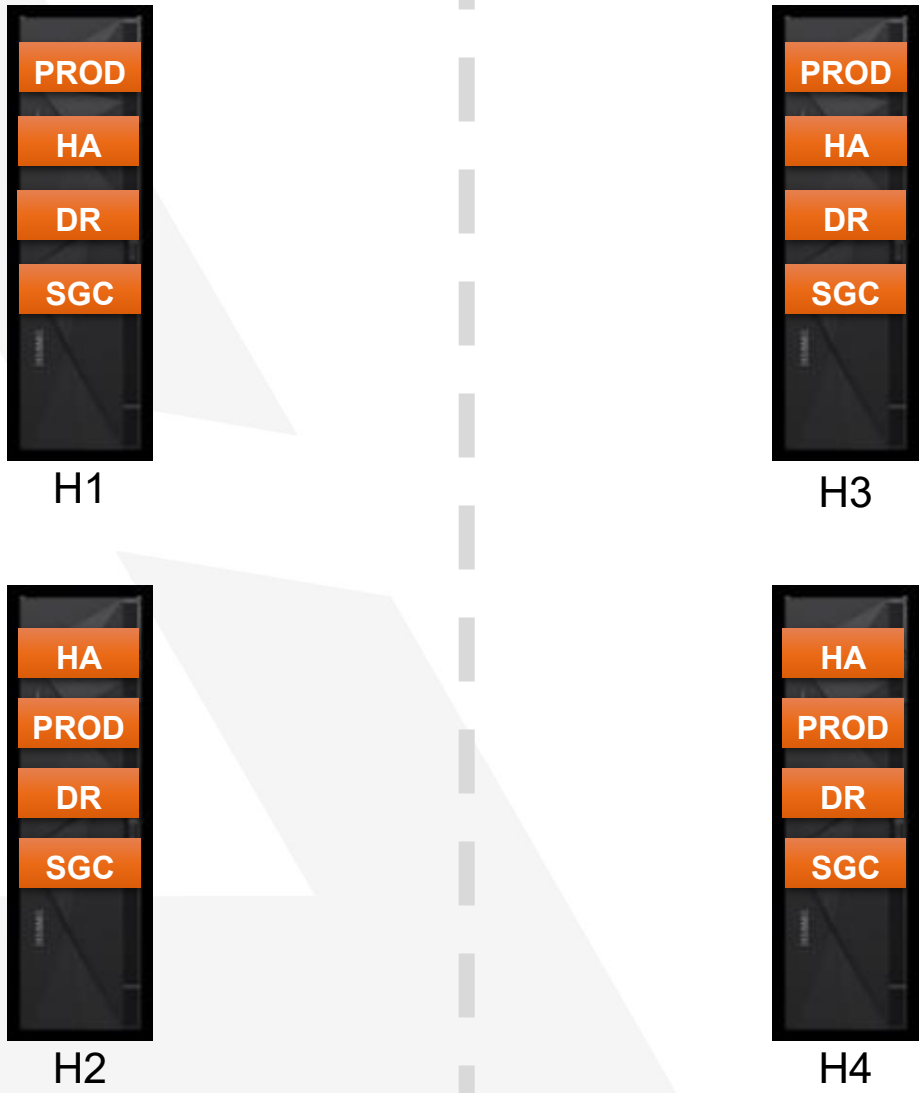

Subchannel Set Usage to Reduce Complexity

- Four copies of data that **must be accessible** in each region
- To manage this complexity, we use all four subchannel sets in each region

Usage	Location	Subchannel Set
Production	H1	0
High-Availability	H2	1
Practice	H1	2
Recovery	H2	3

H1

H2

AUTOMATE DATA CENTER OPERATIONS & REDUCE MANUAL TASKS

Automation on top of Automation

- CSM is an automation product with further automation capabilities which allow further customized automation inside of them
- We leverage these automation capabilities and then add even more external automation and monitoring

CSM Overview and Basic Functionality

- Simple, flexible, and powerful management tool for large copy services configurations
- Streamline and automate complex replication tasks such as:
  - Configurating and monitoring replication
  - Planned & Unplanned outage recovery
  - DR Testing
  - Safeguarded Copy
- Built-in customizable “Scheduled Task” automation capabilities
- Built-in security and data protection features
- Provides GUI, CLI, and API accessibility

Task	Storage Command Line Interface	Copy Services Manager
Set up Global Mirror	5 steps	1 command (START)
Global Mirror Recovery	4 steps	1 command (RECOVER)
Set up Metro Global Mirror	8-25 steps	1 command (START)
Make a practice copy	5 steps	1 command (FLASH)

Automation on top of Automation

- CSM is an automation product with further automation capabilities which allow further customized automation inside of them
- We leverage these automation capabilities and then add even more external automation and monitoring

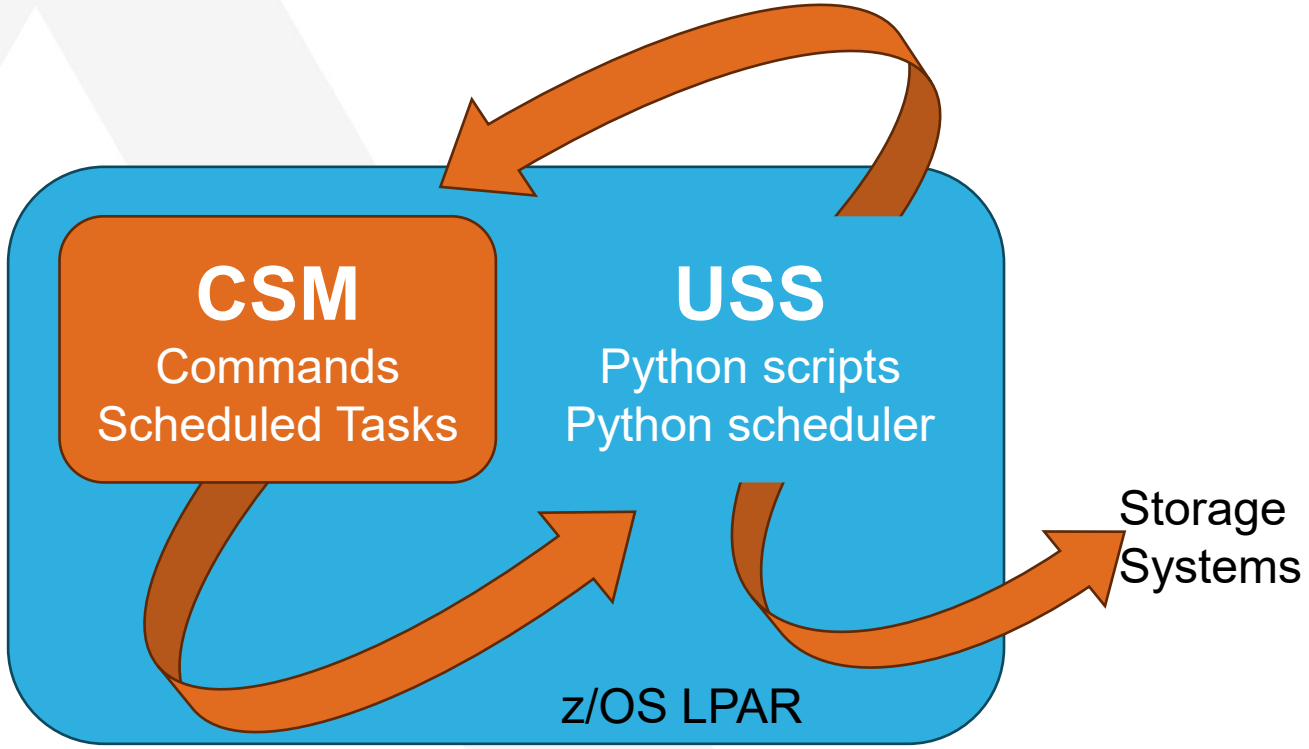

Cascaded Incremental Resynchronization to Handle Failure Scenarios

- Loss of DR DS8950F…
  - H1: PROD
  - H2: High-Availability
  - ~~H3: Disaster Recovery~~
  - H4: Safeguarded Copy

- IR from H1 to H4
  - H1: PROD
  - H2: High-Availability
  - H4: Disaster Recovery ***and*** Safeguarded Copy?

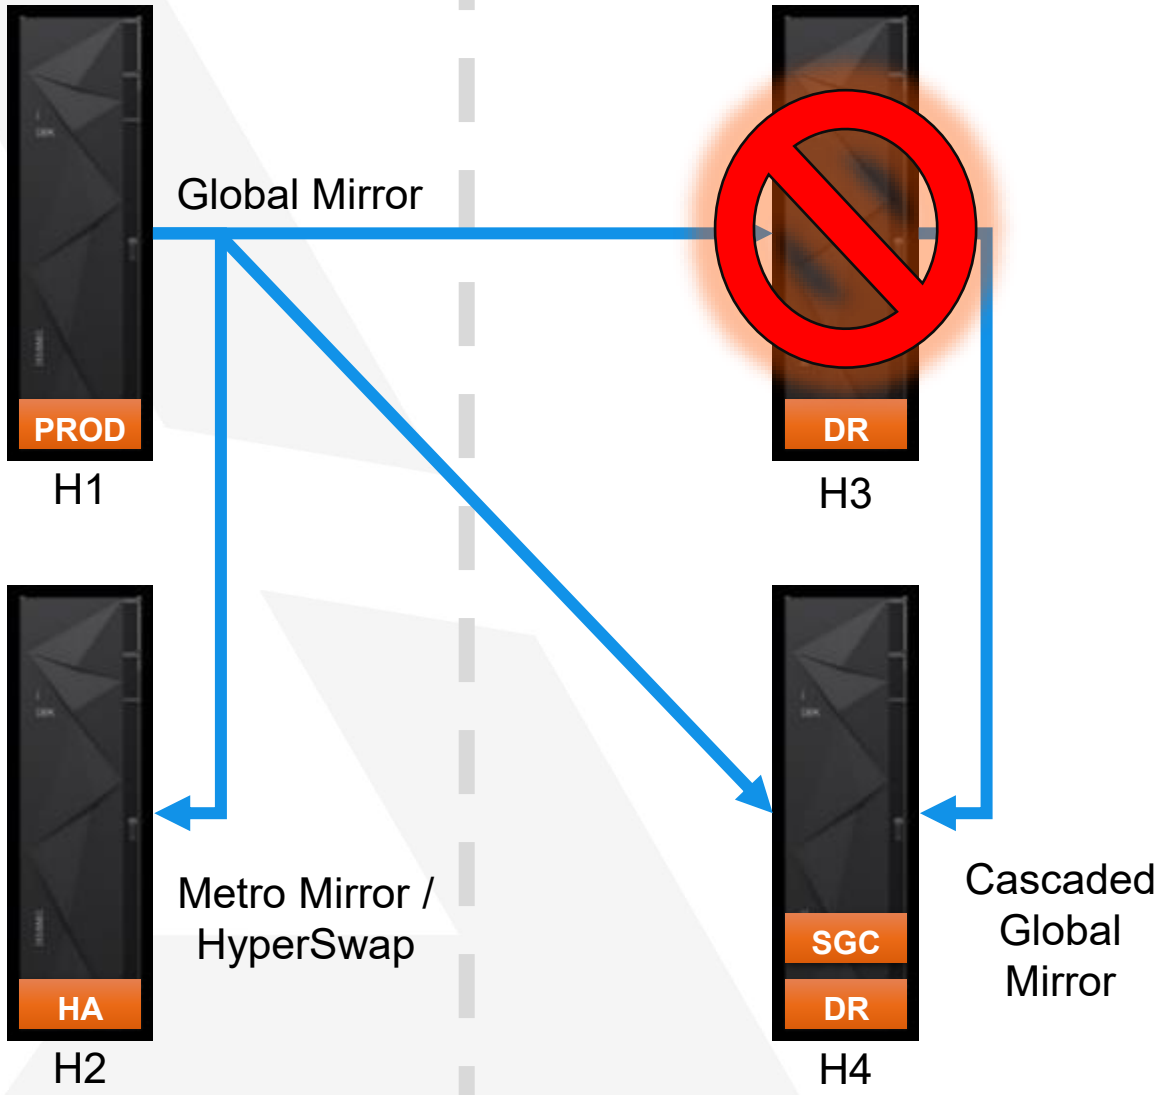

Multi-Target Incremental Resynchronization to Handle More Failure Scenarios

- Loss of PROD DS8950F…
  - ~~H1: PROD~~
  - H2: High-Availability
  - H3: Disaster Recovery
  - H4: Safeguarded Copy
- Automatically HyperSwaps to H2
- CSM automatically restarts replication from H2 to H3 (Incrementally)
  - H2: PROD
  - H3: Disaster Recovery
  - H4: Safeguarded Copy

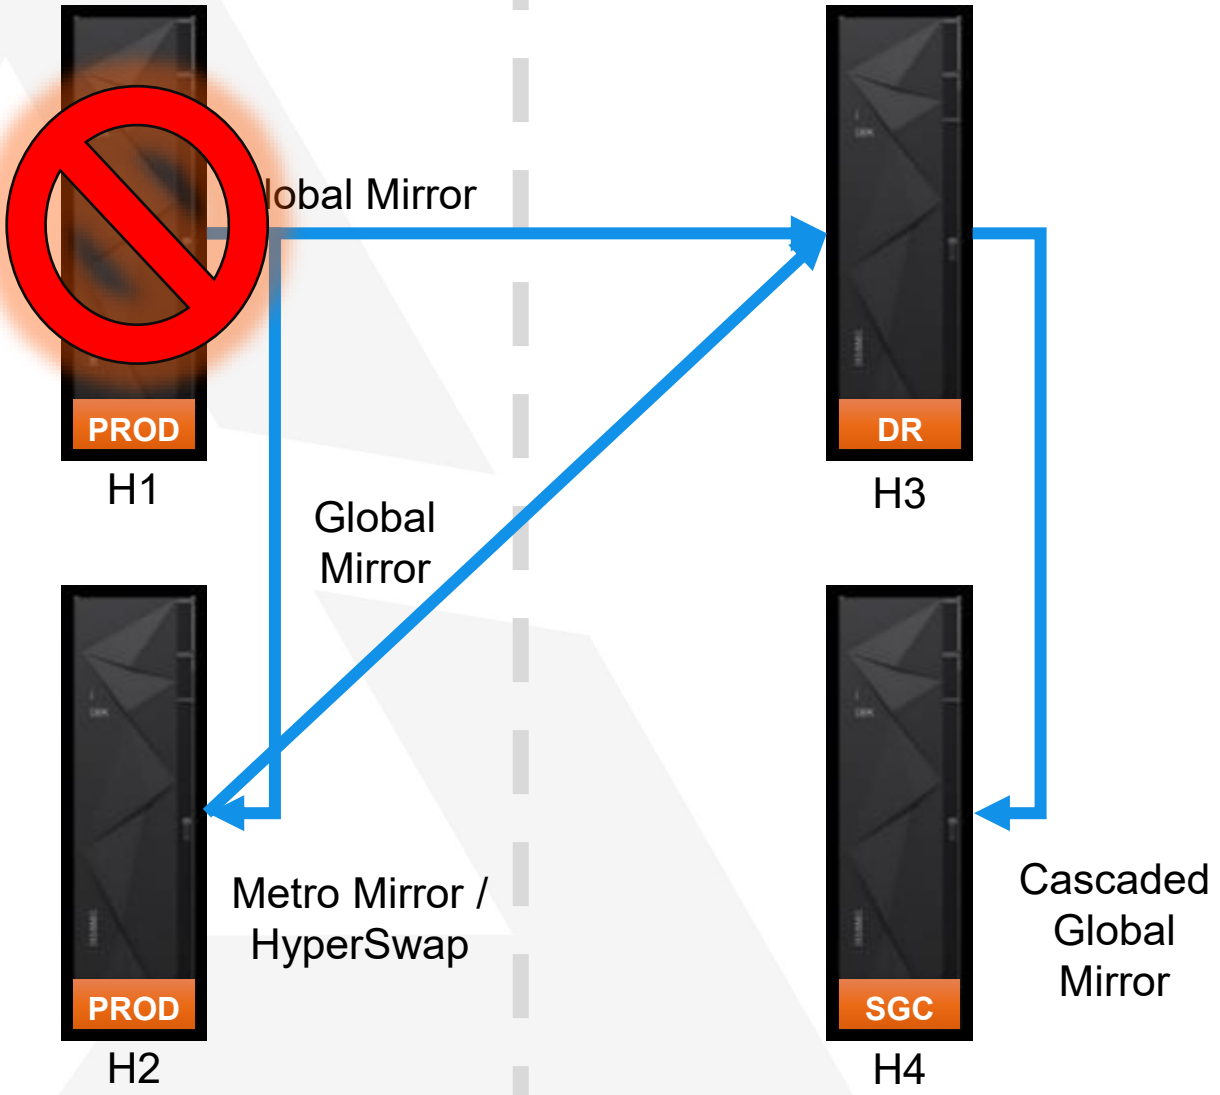

Create DR Test Volumes and SGC Backups
Session Associations

- "Flash 3rd Site" command
- "Backup 4th Site" command

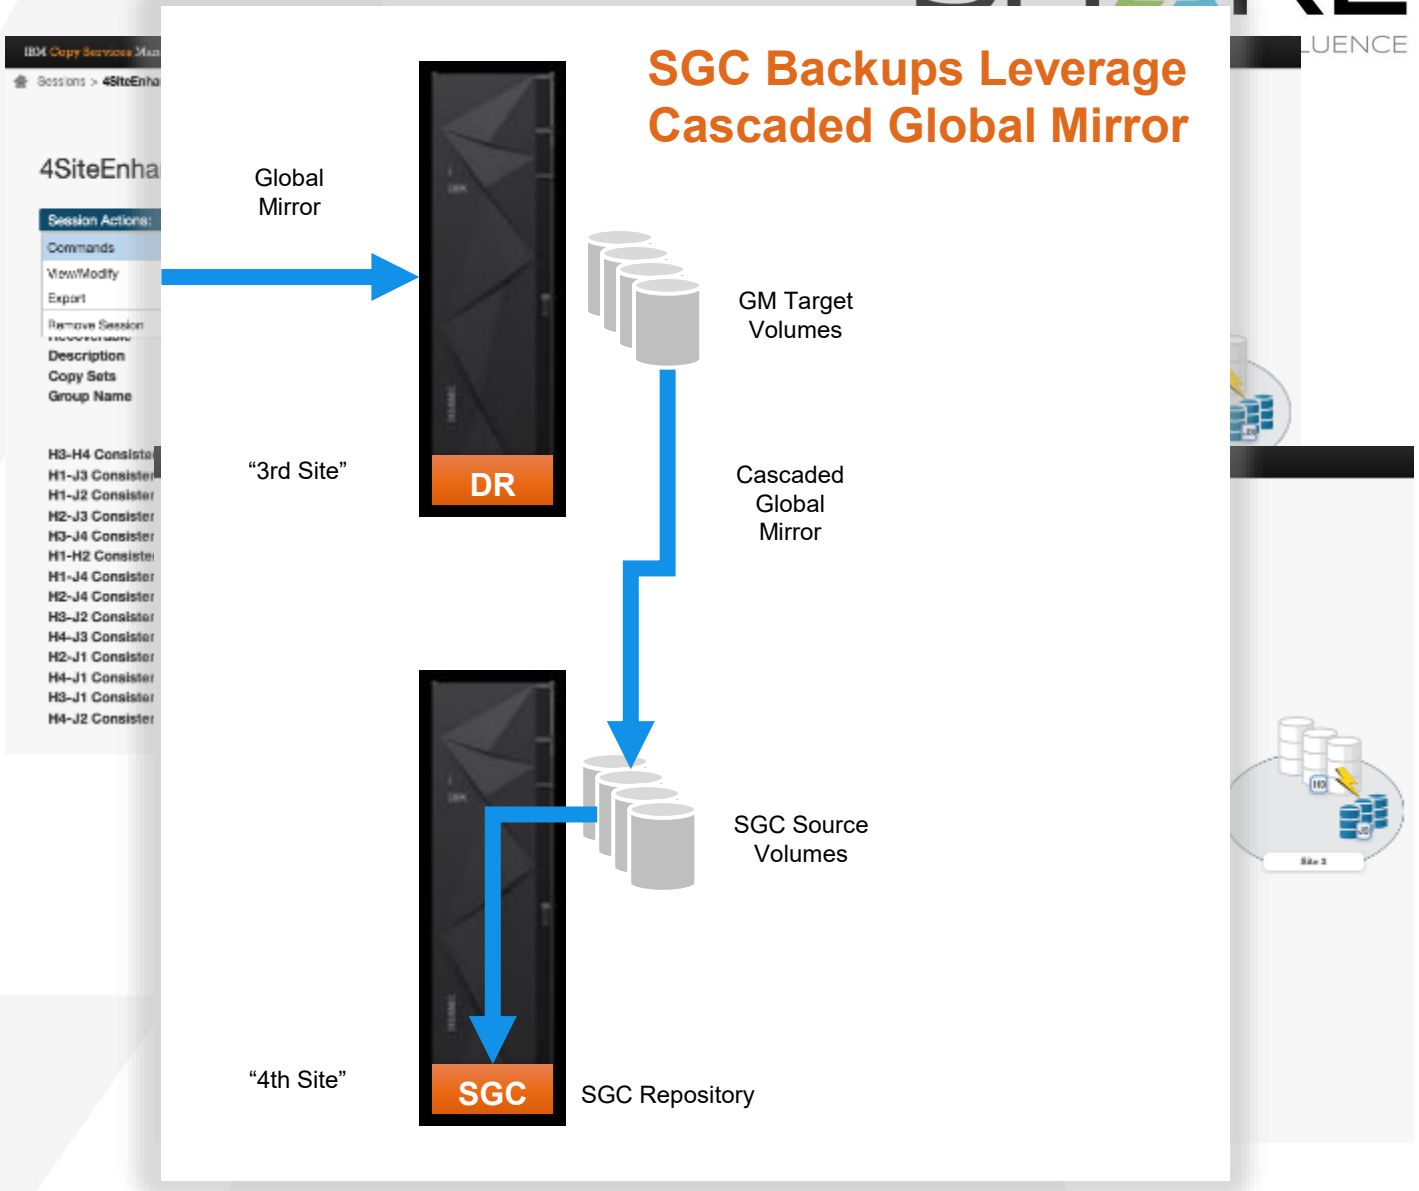

Automated SGC Backup Processing

Every 10 minutes & retained for 24 hours

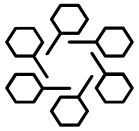

CSM Scheduled Task

- Backup 4th Site command
- Integrates all SGC & GM processing
- On Failure: Resume Primary GM

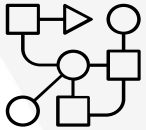

Integrated Status checks to minimize RPO

- Primary & Cascaded GM health
- Primary & Cascaded RPO
- Fail ASAP & restore GM ASAP

SYSPLEX Associations for Enhanced SGC Performance

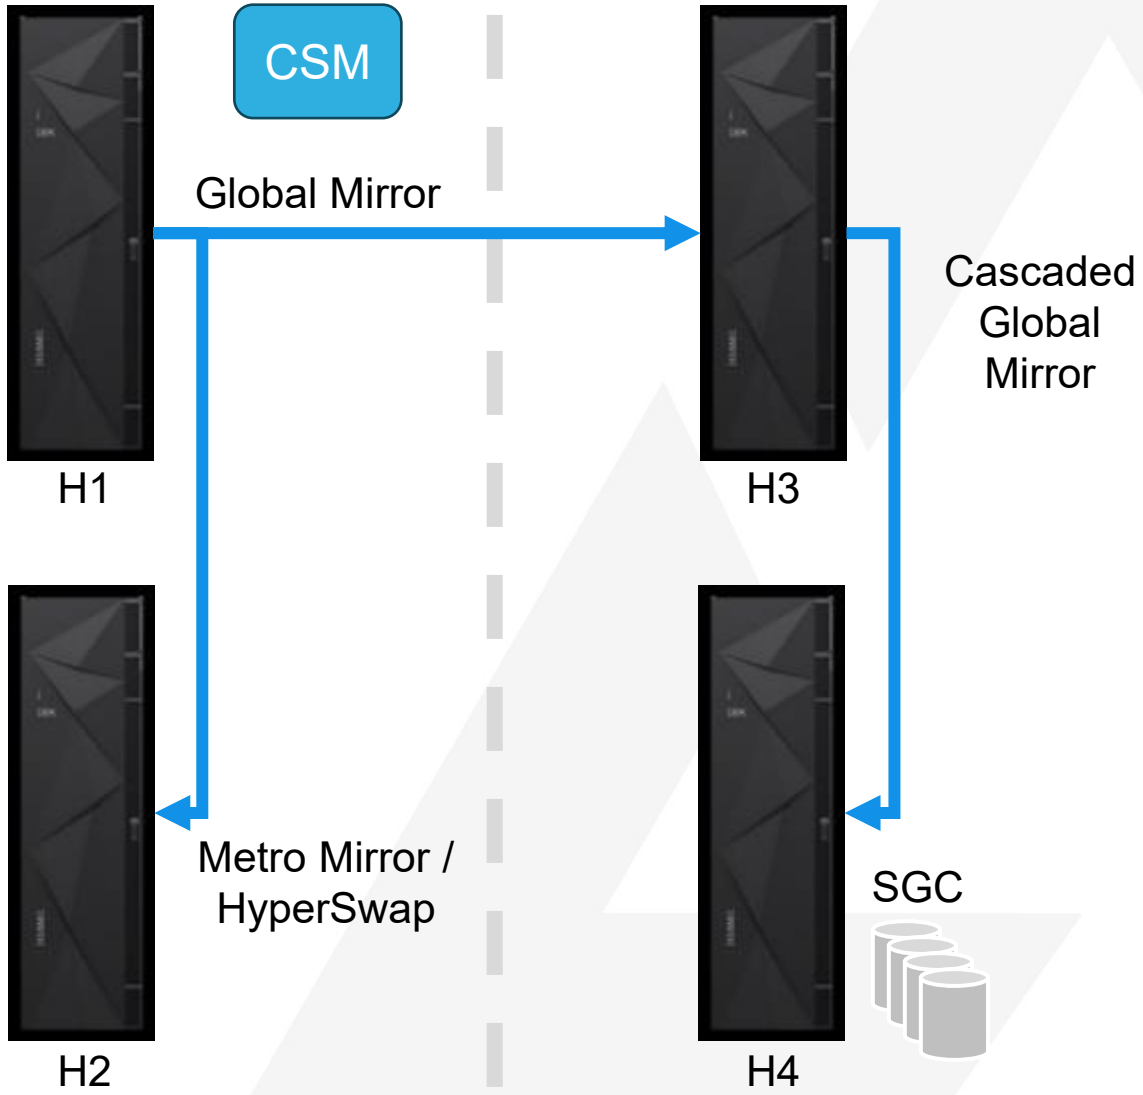

- Normally performance isn't critical when taking an SGC backup at a remote site that is cascaded off a GM relationship
- Since we create SGC backups every 10 minutes, we must create the backup AND restart the replication back up as quickly as possible to ensure suitable RPO
- When creating a backup using the IP connection to the HMC, longer delays can occur due to queries that need to be issued during backup processing
- So how can we create these SGC backups faster?

Offload SGC Processing to Alternate LPAR

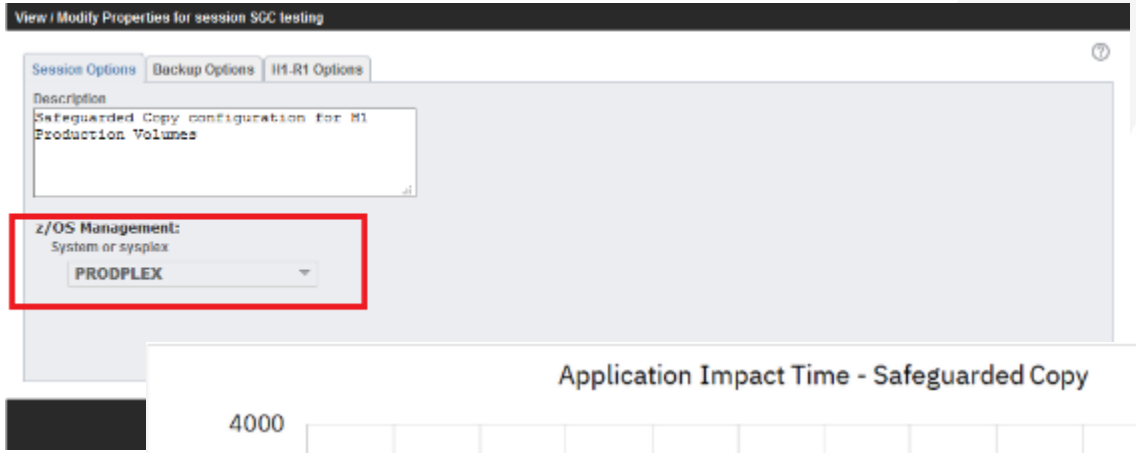

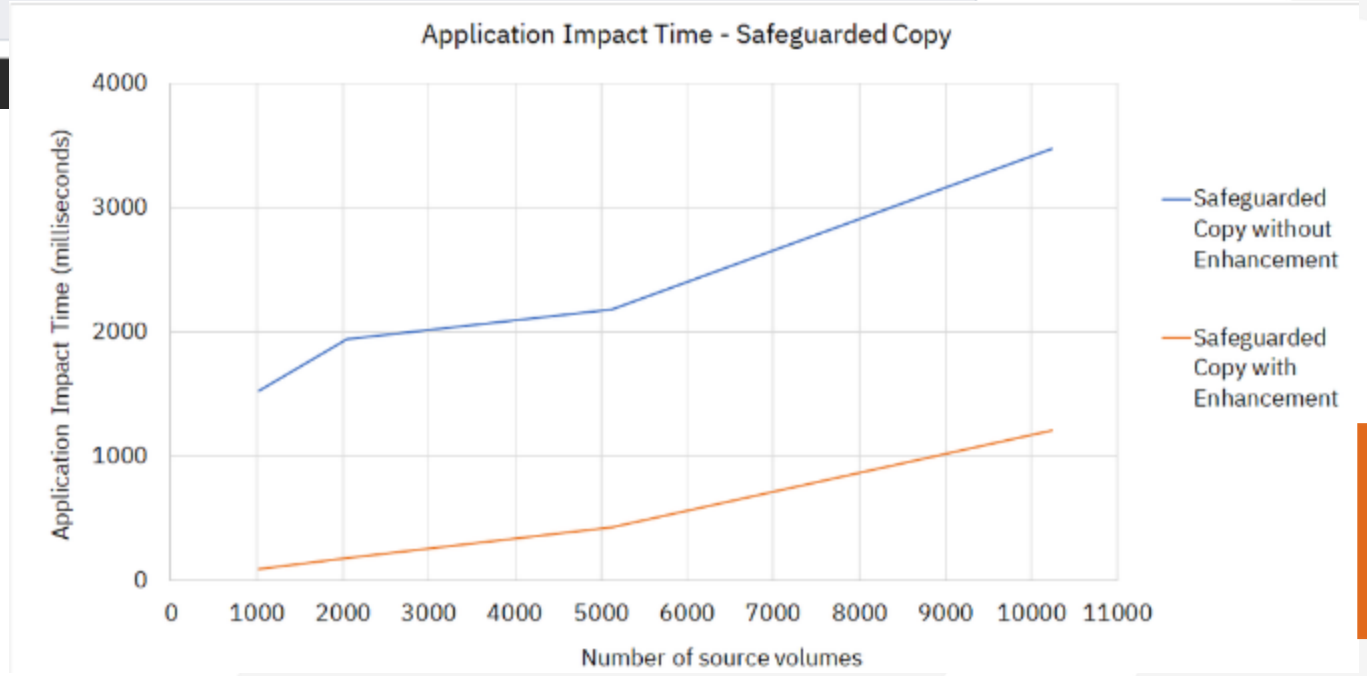

- When you associate a Sysplex to an SGC session, CSM will offload the command to z/OS to create the backup
- z/OS utilizes its FICON connection to the hardware and has proven to significantly decrease the time it takes to create an SGC backup

Reduces processing on primary CSM **and** improves SGC backup performance

Implementation through CSM IP to z/OS Connectivity

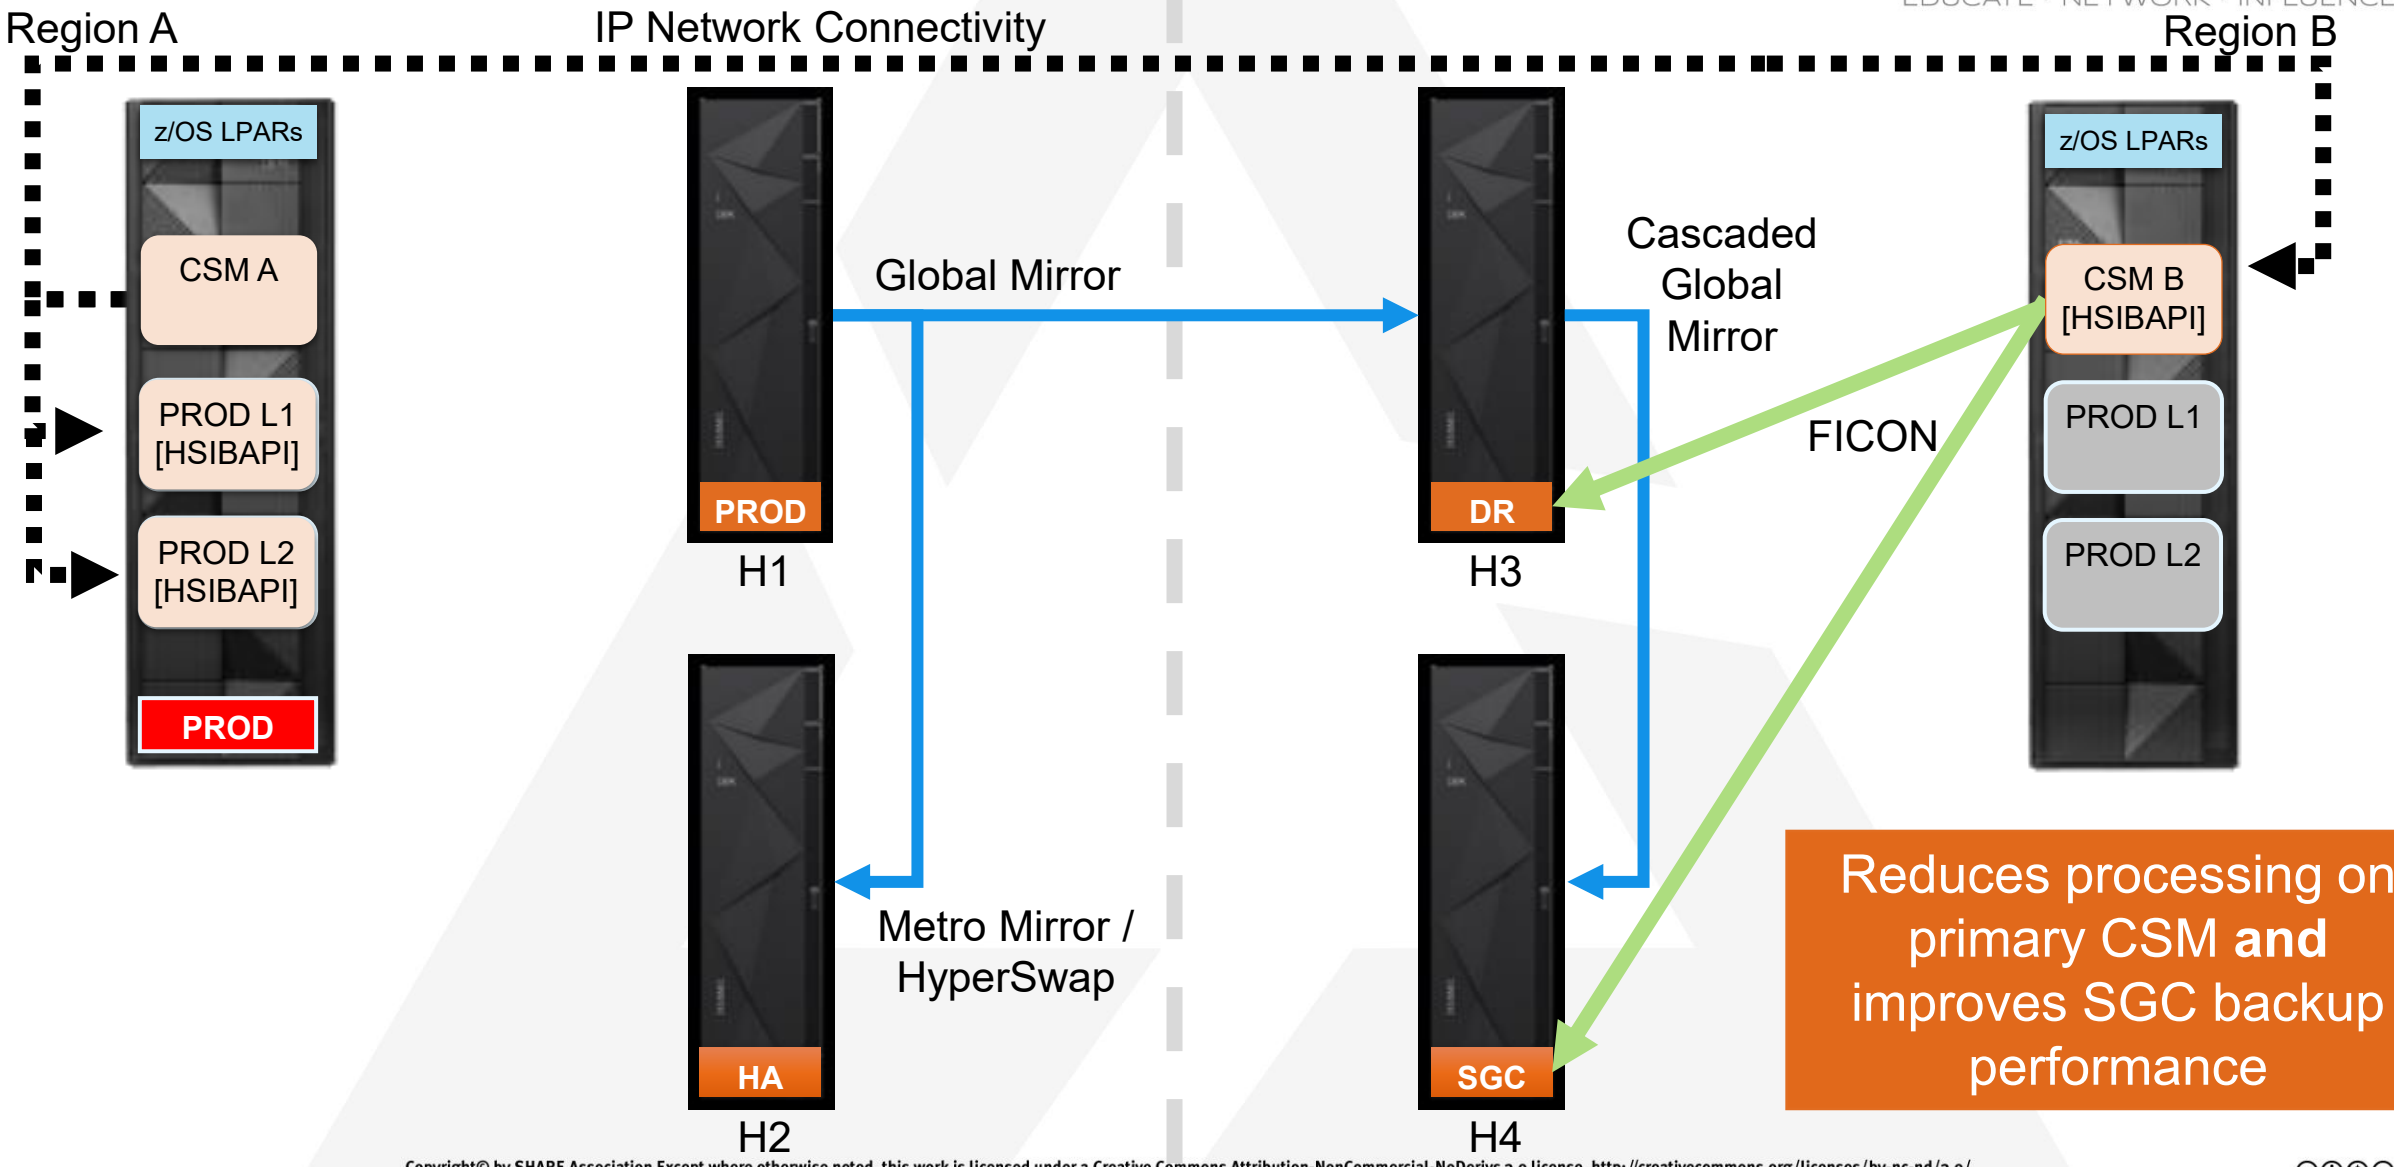

Reduces processing on primary CSM **and** improves SGC backup performance

Hold Weekend SGC Backups

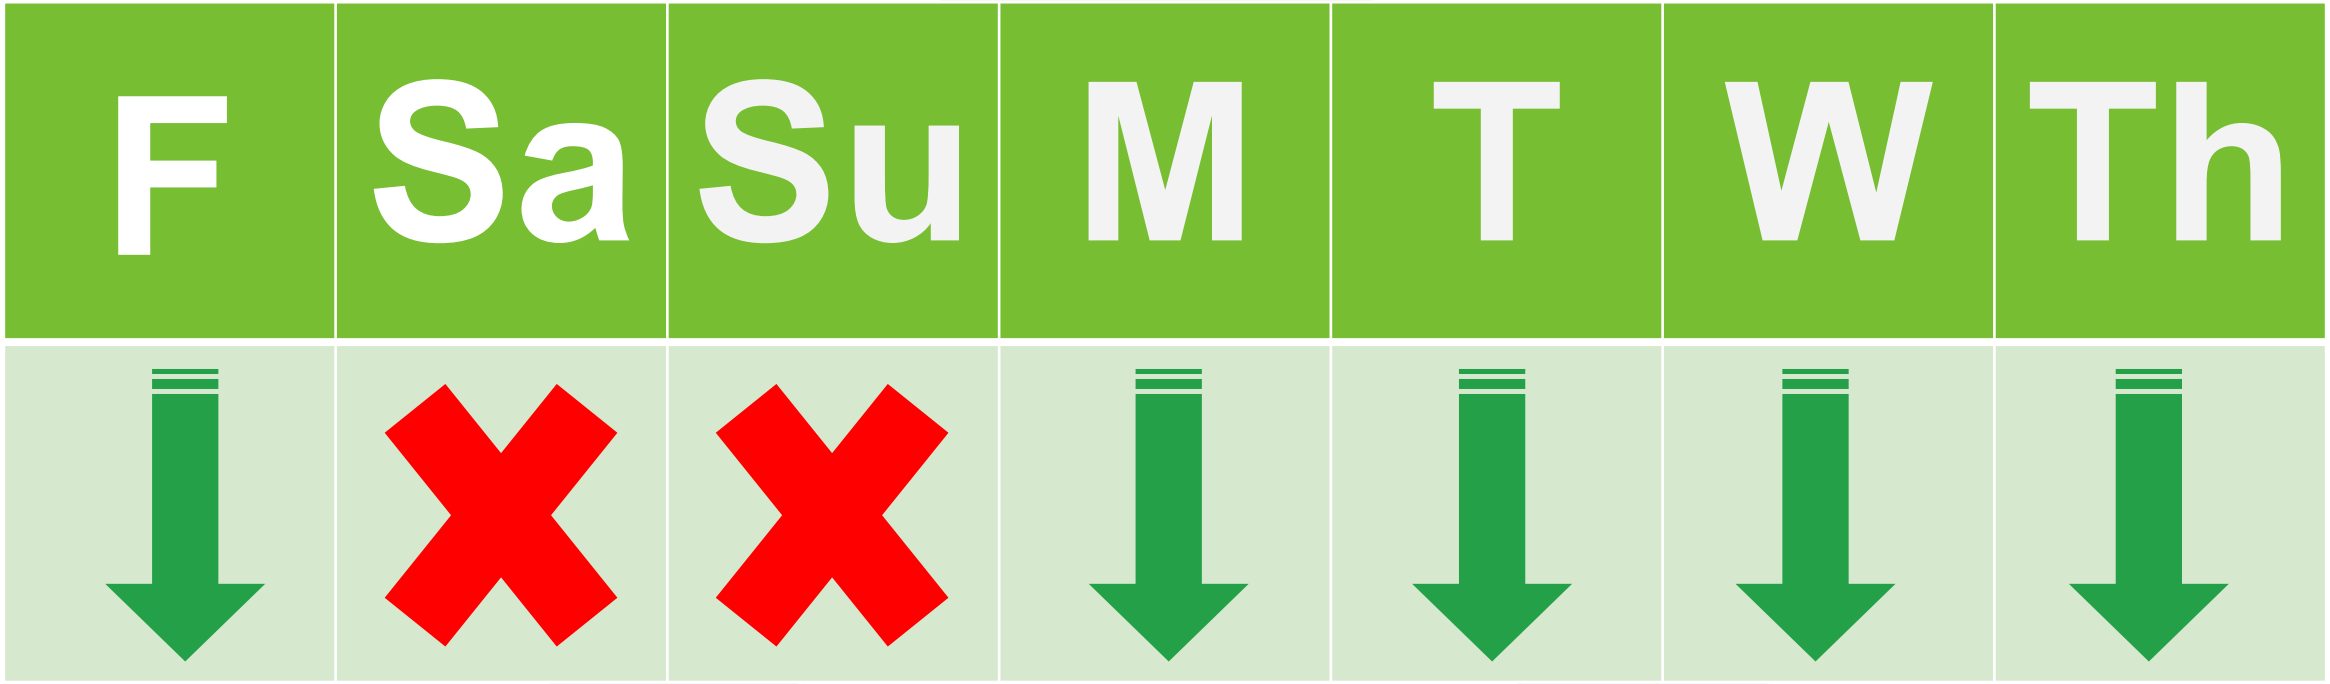

- We want to protect Business Data
- Very little Processing after Batch Completes Friday PM
- We Suspend the Backups and Expiration over the Weekends
- Restart SGC Backups Monday
- BUT: Normal expiration no longer works
- SO: Expiration must be automatically extended after Hold

Hold Weekend & Holiday SGC Backups

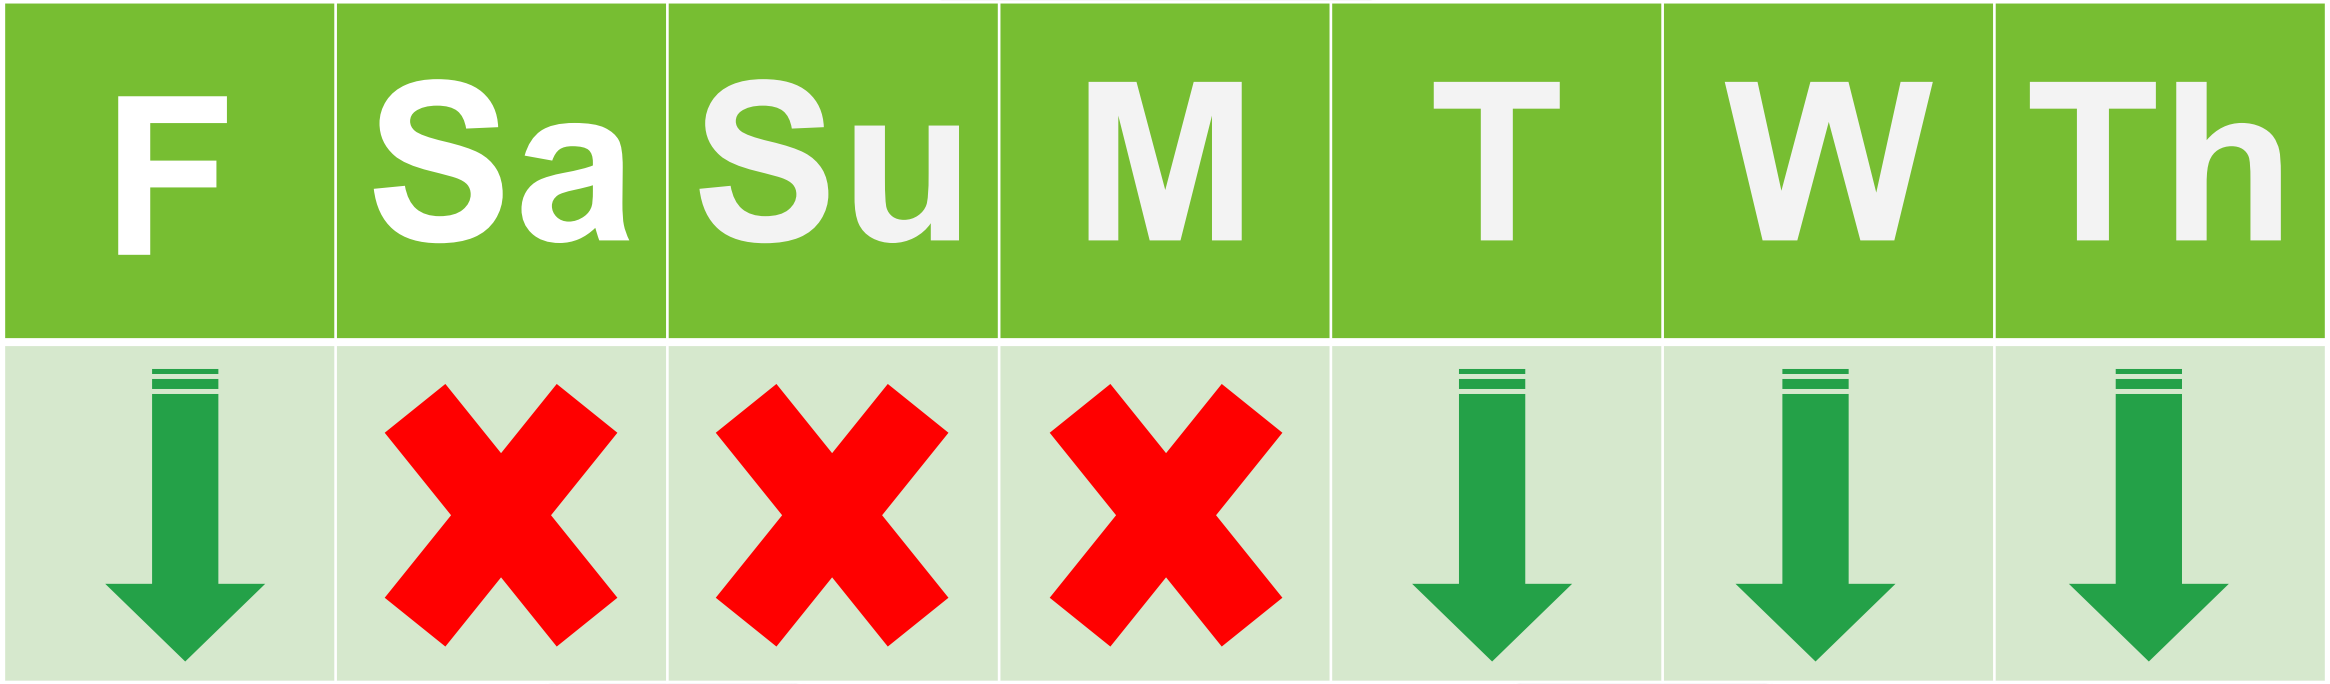

- We want to protect Business Data
- Very little Processing after Batch Completes Friday PM and during Holidays
- We Suspend the Backups and Expiration over the Long Weekends
- Restart SGC Backups Monday or Tuesday
- BUT: Normal expiration no longer works
- SO: Expiration must be automatically extended after Hold

Hold Weekend & Holiday SGC Backups

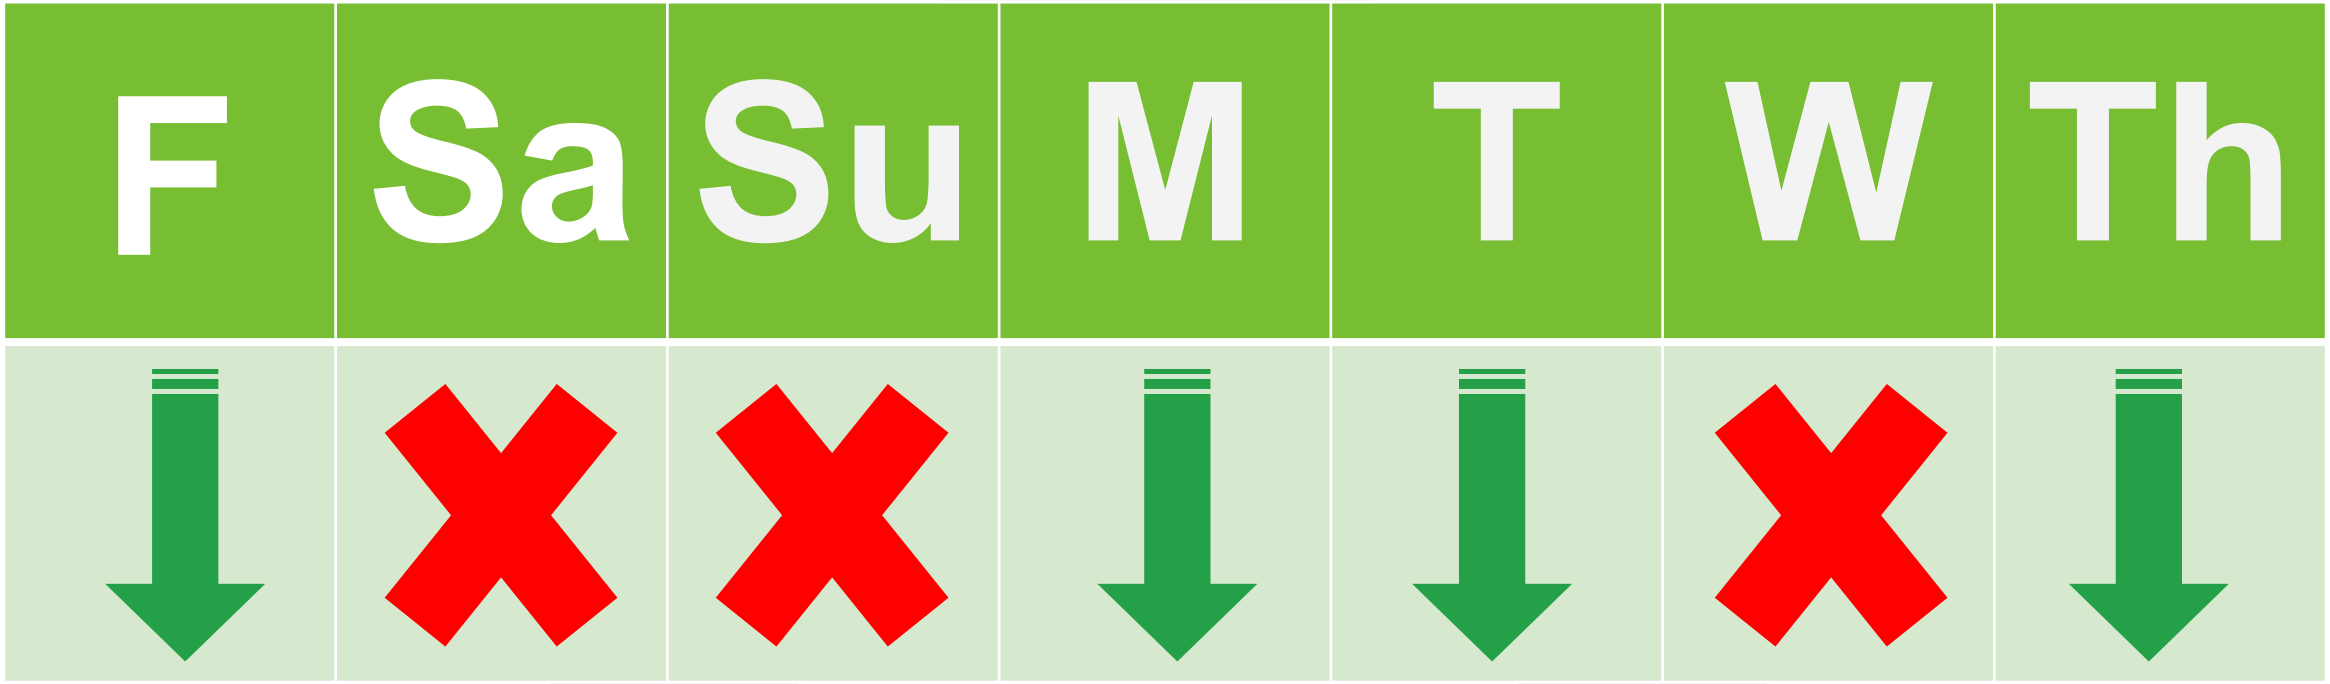

- We want to protect Business Data
- Very little Processing after Batch Completes Friday PM and during Holidays
- We Suspend the Backups and Expiration over the Weekends and mid-week Holidays
- Restart SGC Backups Monday and after the Holiday
- BUT: Normal expiration no longer works
- SO: Expiration must be automatically extended after Hold

Advanced Automation: CSM Scheduled Tasks

- Multiple commands
- Multiple sessions
- Status Checking
- External scripts
- Scheduled or ad-hoc

- Chain tasks together:
  - If Successful
  - On Failure

- Tasks created for:
  - SGC Backups
  - Site Swap
    - Check OOS
    - Prepare
    - GO
    - No-Go

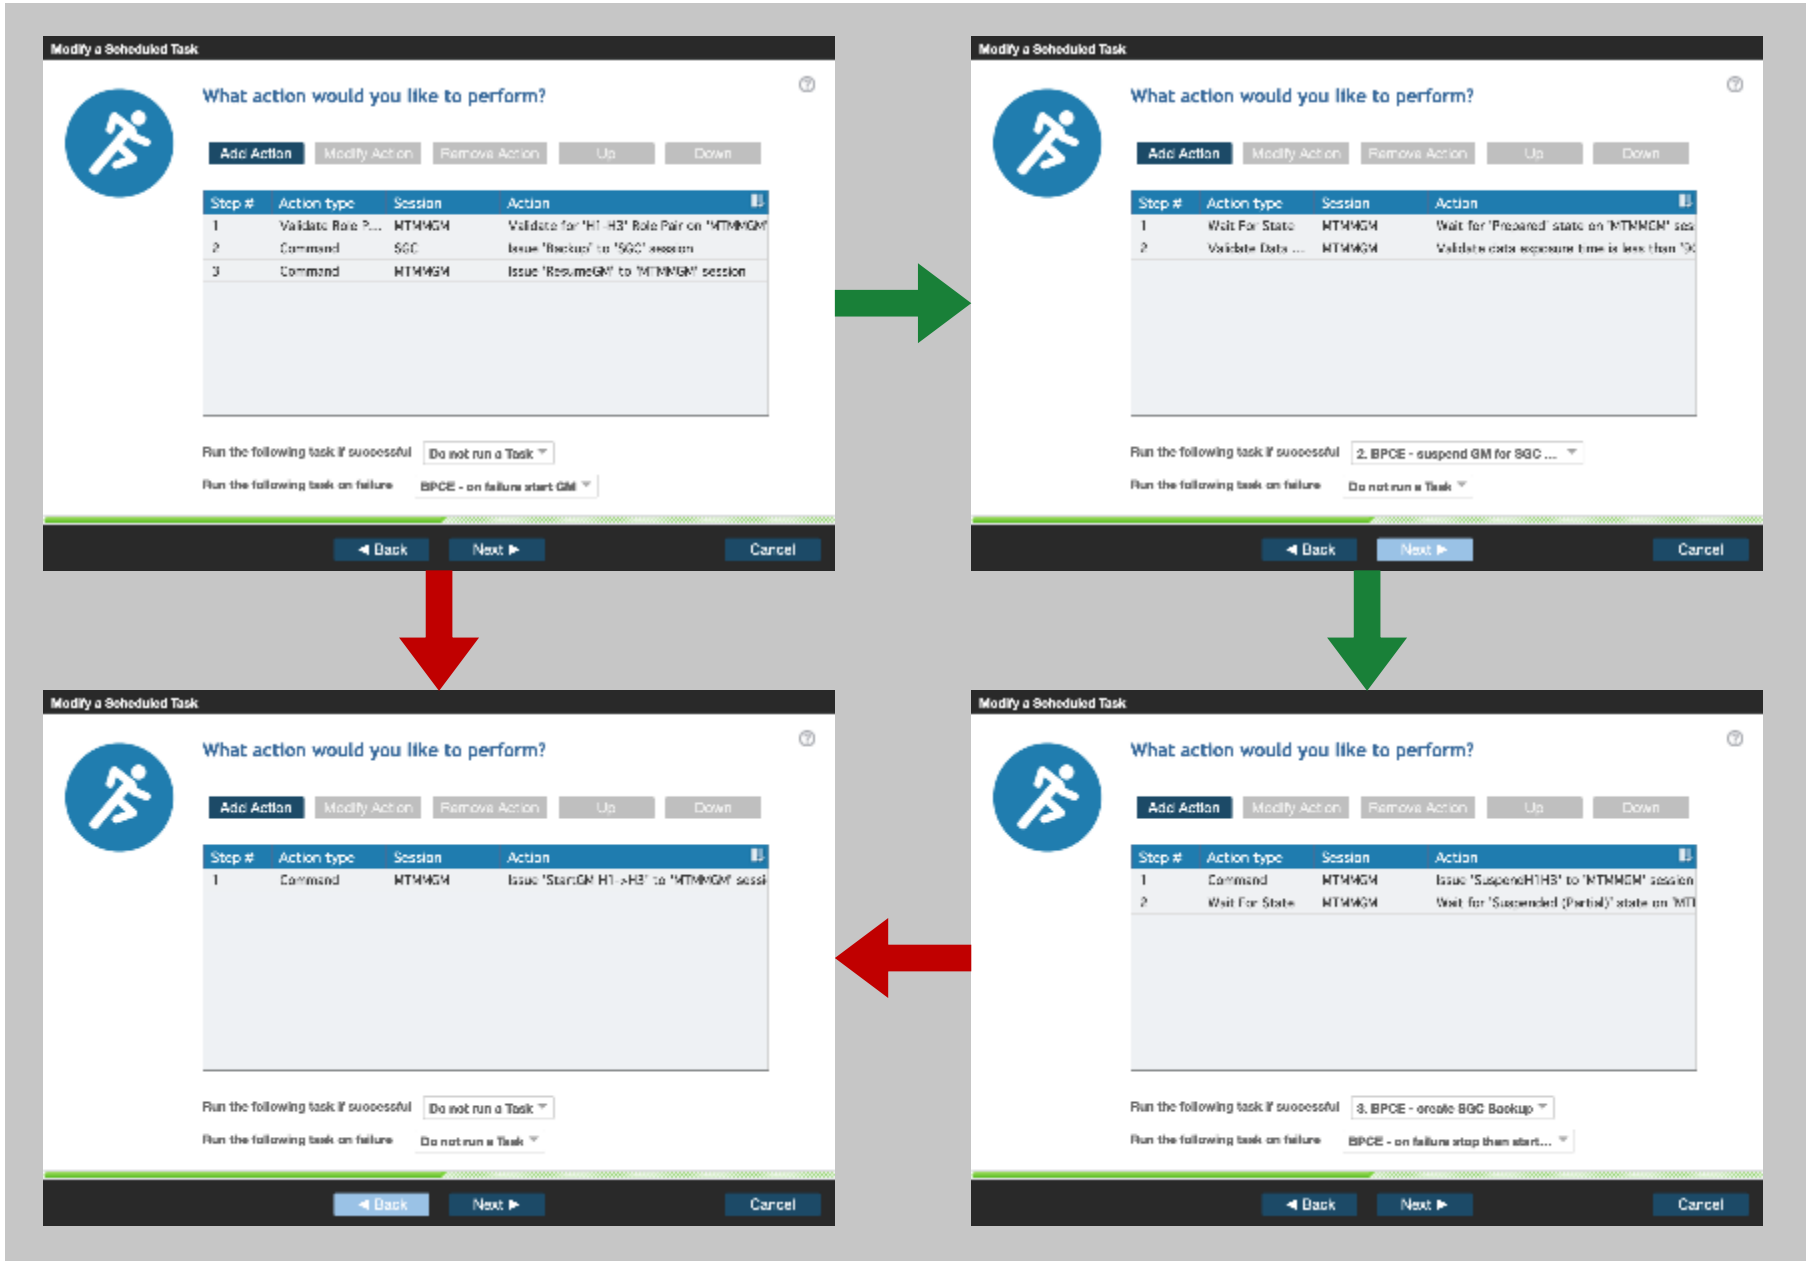

What action would you like to perform?

Add Action | Modify Action | Remove Action | Up | Down

Step #	Action type	Session	Action
1	Command		Issue 'Confirm Production at Site 1' to
2	Command		Issue 'Confirm Production at Site 1' to
3	Command		Issue 'FailoverH2' to sessi
4	Command		Issue 'Start H1->H2' to ses
5	Command		Issue 'StopH1H2' to sessi

Run the following task if successful: ALL-Start H1 to H3-GC

Run the following task on failure: Do not run a Task

What action would you like to perform?

Add Action | Modify Action | Remove Action | Up | Down

Step #	Action type	Session	Action
1	Command		Issue 'StartGC H1->H2 H1->H3' to
2	Command		Issue 'StartGM H1->H3' to
3	Run External Script		Run an external command on server localhc
4	Command		Issue 'StartGC H3->H4' to
5	Command		Issue 'StartGM H3->H4' to
6	Run External Script		Run an external command on server localhc
7	Run External Script		Run an external command on server localhc
8	Command		Issue 'StartGC H1->H2' to

Run the following task if successful: Do not run a Task

Run the following task on failure: Do not run a Task

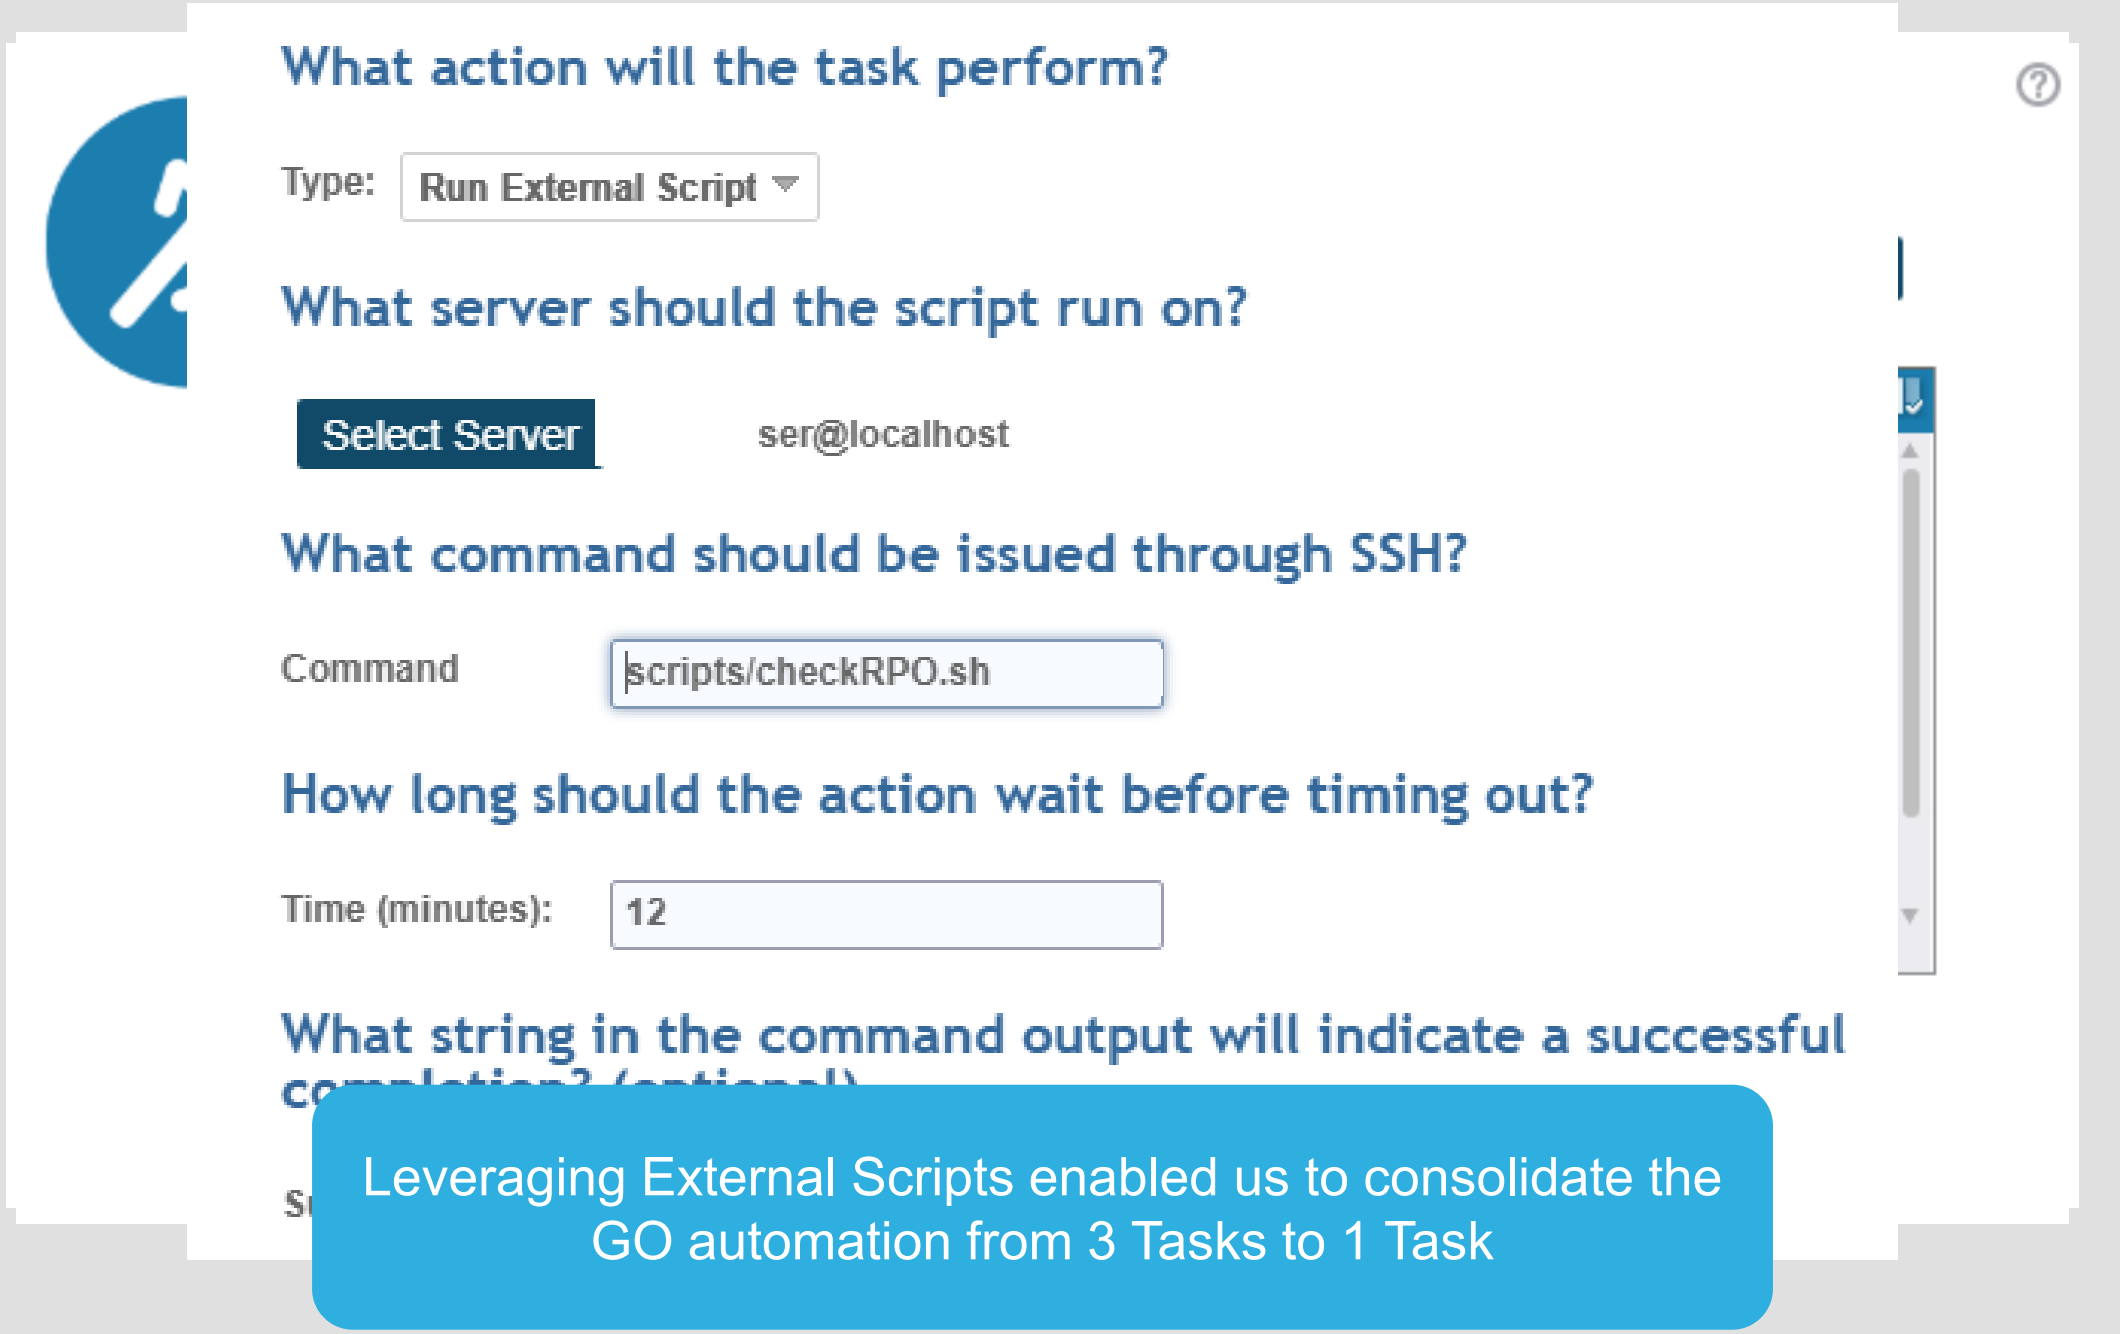
What action will the task perform?
Type:
Run External Script
What server should the script run on?
Select Server
ser@localhost
What command should be issued through SSH?
Command
scripts/checkRPO.sh
How long should the action wait before timing out?
Time (minutes):
12
What string in the command output will indicate a successful
Leveraging External Scripts enabled us to consolidate the GO automation from 3 Tasks to 1 Task

External Monitoring and Alerting

- All Storage system and CSM instances are configured to send alerts to z SYSLOG

- We have automation that actively monitors the SYSLOG
  - Notify if certain expected messages are not seen in certain time periods
  - Alert if certain unexpected messages are seen
  - Automatically trigger actions for certain conditions

External Script Infrastructure

- Monitoring system created (Python on USS) to periodically check various portions of the system and alert if necessary

- Example: Proactive SGC Multiplier monitoring and expansion
  - Automatically takes action to expand multipliers when necessary

- Example: Replication performance monitoring
  - Automatically transitions replication from synchronous to asynchronous if response times exceed a threshold

External Script Infrastructure

- Periodic Queries to Ensure Proper System Operation and Satisfy Audit Requirements
  - Hourly: Validation of successful SGC backups
  - Daily: Reports on SGC Backups and RPO Performance for previous 24 hours
  - Weekly: Parity checks of the environment
    - CKD Volumes, VOLSERs, ports, encryption status…

- Facilitate Command Center Control
  - Script to open and close support access
  - JCL that calls Python script to automatically generate support logs for all DS8950Fs in each environment

SYSPLEX Associations for HyperSwap

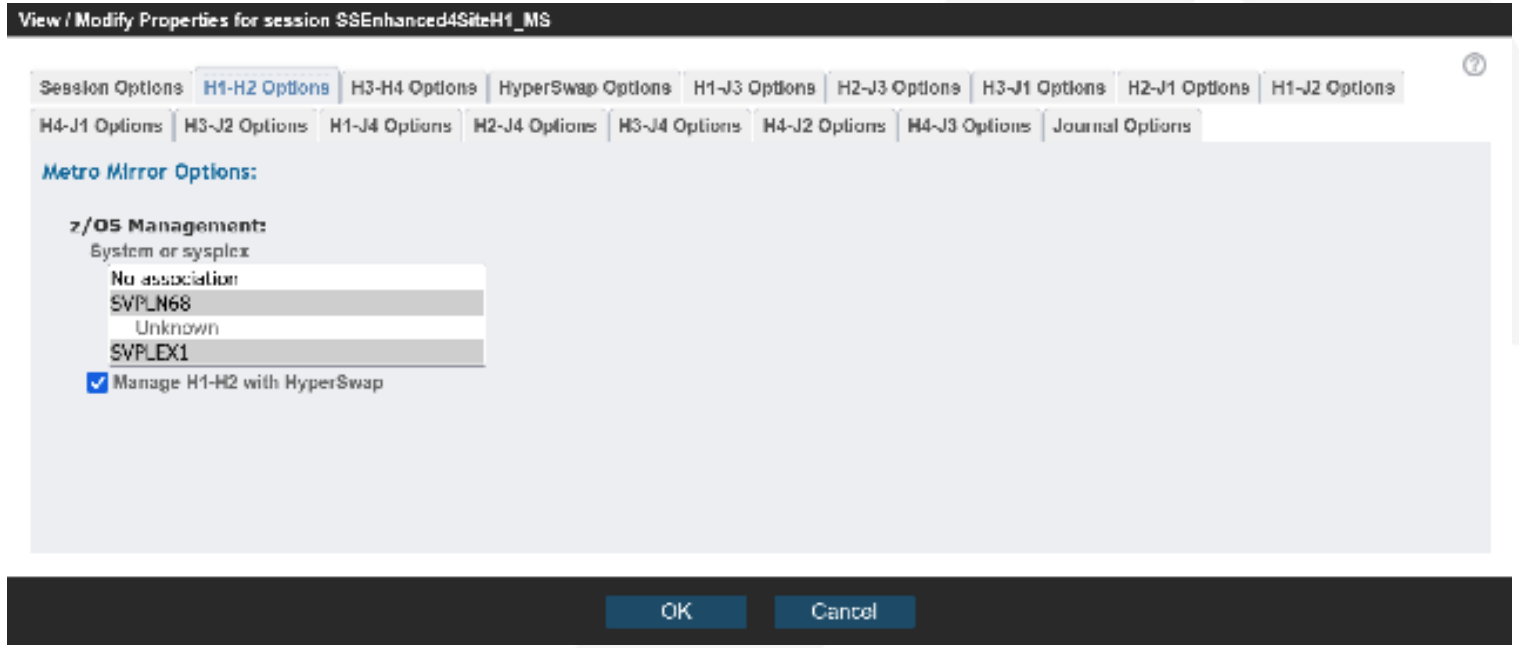

- New CSM feature manages multiple SYSPLEXes in a single session
- Now ONE CSM session coordinates and loads the configuration across all associated SYSPLEXes
- Hardware Reserves features on DS8000, and z/OS HyperSwap allow separate SYSPLEXes to HyperSwap if any of the other SYSPLEXes HyperSwap

STRENGTHEN SECURITY AND ACCESS CONTROLS

RACF Manages all Authentication

- RACF is installed on all C-systems

- All Authentication is managed by RACF and requires Multi-Factor Authentication
  - CSM
  - DS8950Fs
  - z/OS LPARs

Goal: No static known passwords anywhere!

PassTicket Usage for Security Without Known Passwords

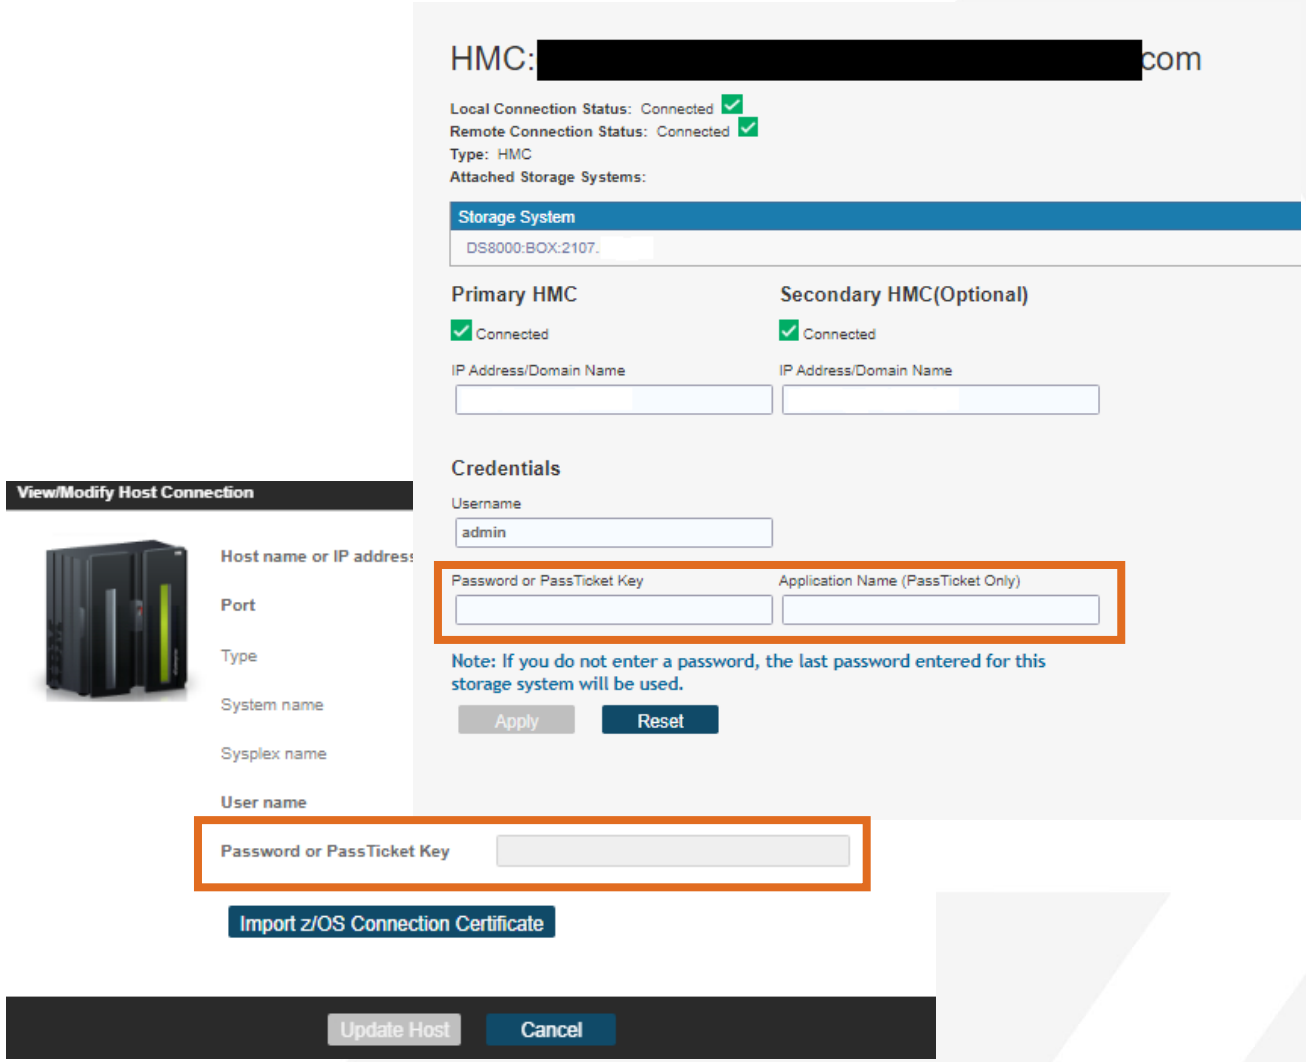

- Requirement to remove all static "known" passwords
- Automated apps cannot use MFA
- Solution is to use PassTickets
- Application generates a new password for each sign-in attempt, and it expires after 10 minutes.
- Implemented for connectivity between
  - CSM and DS8950Fs
  - CSM and z/OS LPARs
  - Automation scripts and DS8950Fs

More Security Without Known Passwords

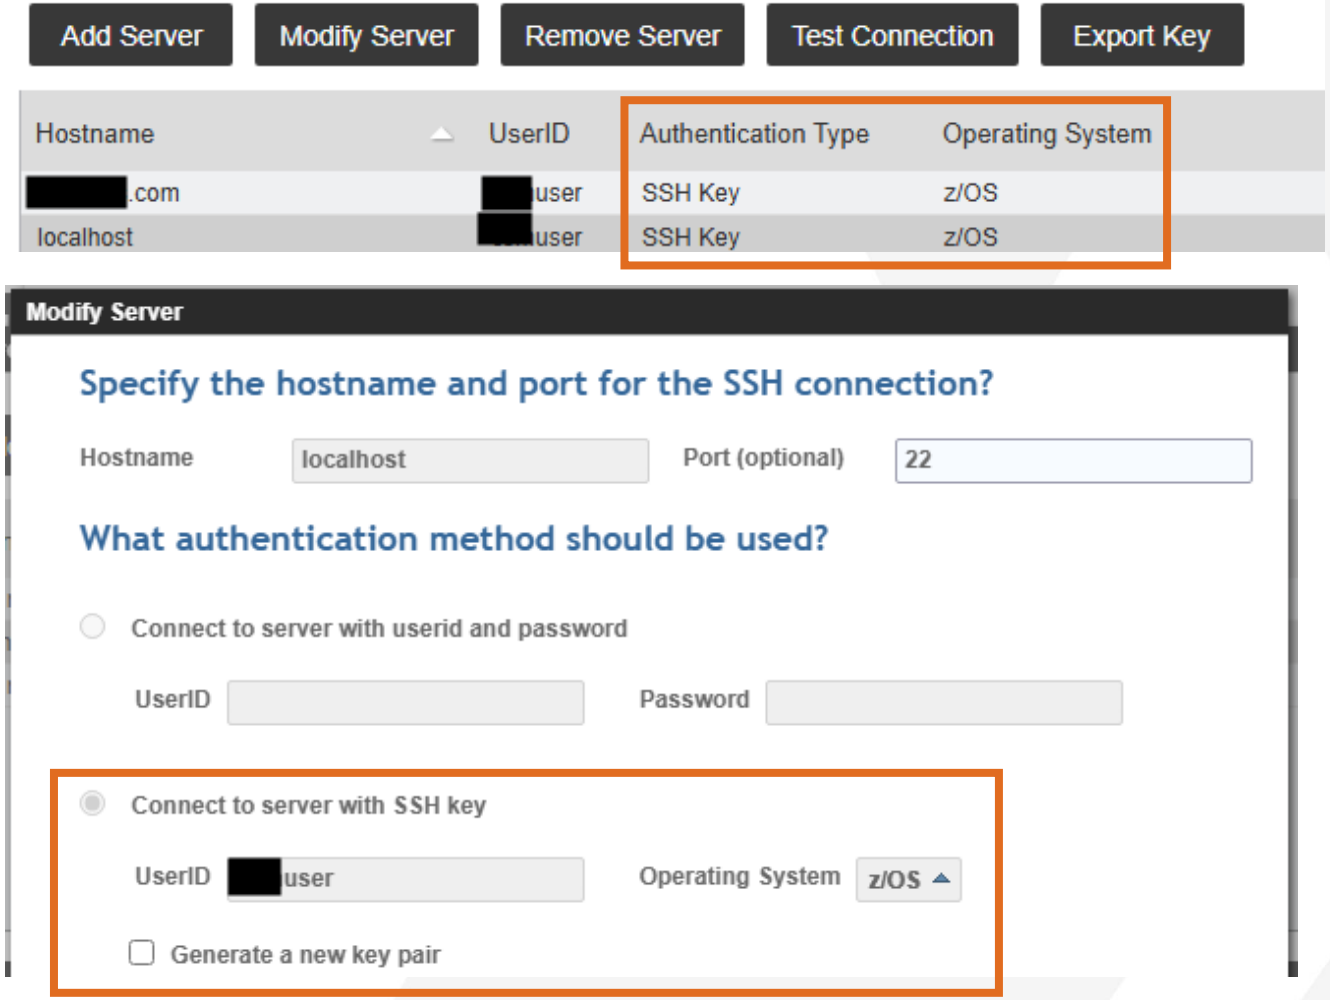

- Requirement to remove all static "known" passwords
- CSM utilizes SSH certificate authentication for CSM external script automation capability
- CSM leverages SAF authentication for CSMCLI connectivity from z/OS LPARs

SGC Dual Control for Enhanced Security

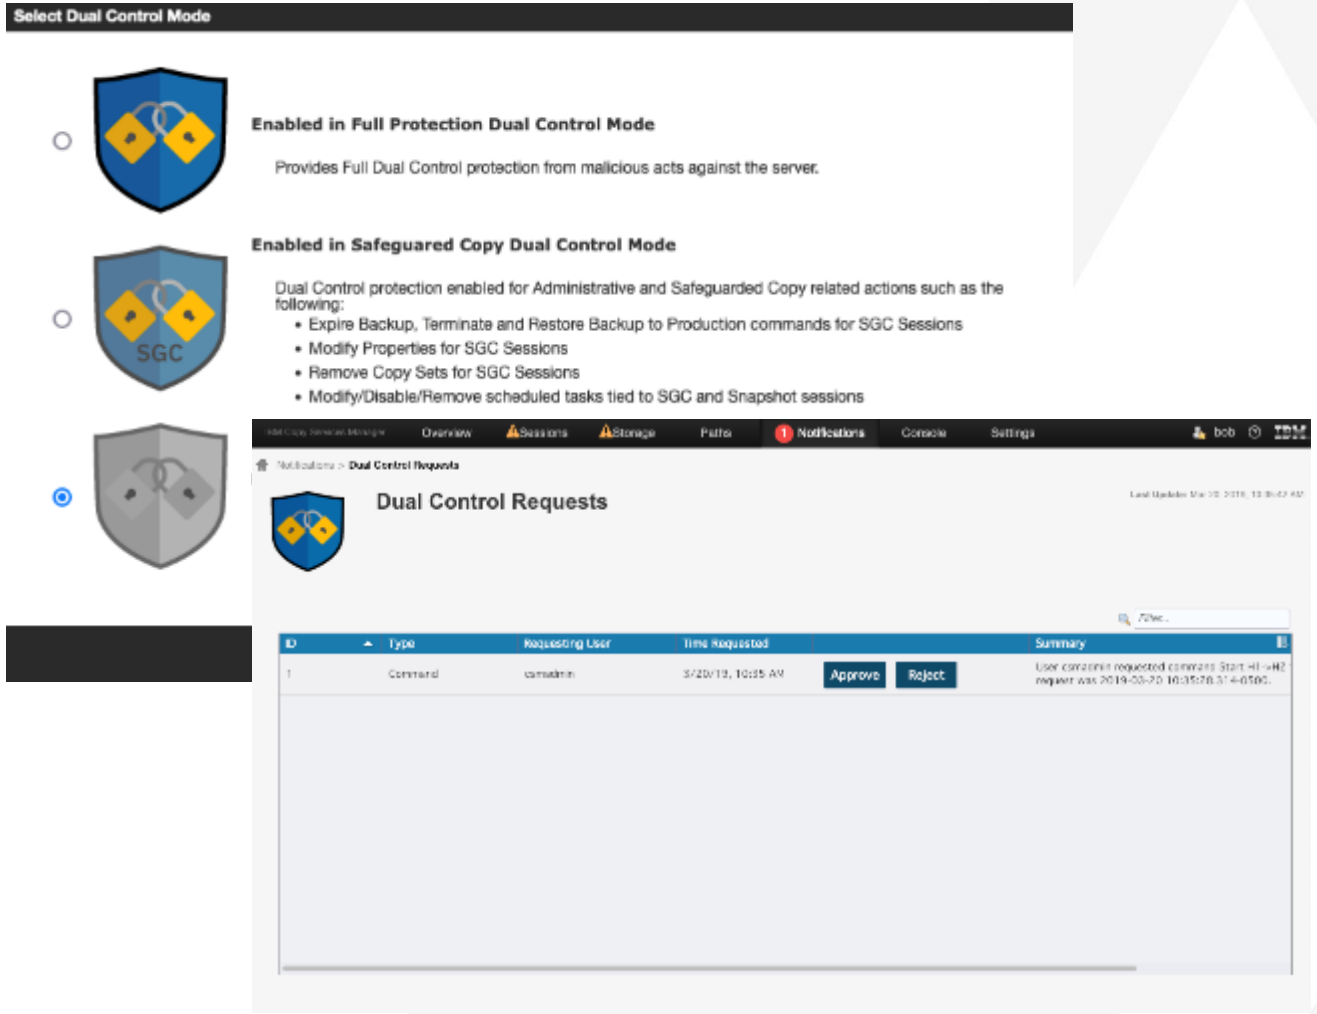

- CSM Dual Control protects customer environments by requiring two users to perform actions
- Full Dual Control mode however can lock down TOO much and make managing the replication too onerous
- Safeguarded Copy Dual Control mode is used in this environment to lock down SGC without the overhead of full Dual Control mode

SGC Immutability Mode

- Management of SGC Backups is split between storage system and CSM providing separation of duties

- CSM Time-Locked Expirations prevents SGC backups from being modified or removed until after the retention time
  - Our retention is 24 hours…
  - CSM will not expire the backups until 24 hours has elapsed
  - No manual expiration is allowed until 24 hours has elapsed
    - Even with dual confirmation

- CSM Immutability Mode sets a system-wide minimum setting for the retention value of SGC backups

- Once set, this value can **NEVER** be lowered - even for other SGC sessions
  - Also set to 24 hours thus guaranteeing a 24-hour retention

Protects against bad actors and innocent mistakes

Storage System “Airgap” Mode

- Storage Systems that contain the SGC Backups are put into “Airgap” mode

- All external inbound and outbound ports are closed by the storage system itself
  - GUI, API, Remote Support
  - The CLI port is the only open port

- An authenticated user can temporarily open the ports for service or enable/disable airgap mode during site swaps via the CLI

FUTURES / NEXT STEPS

Future Enhancements

- Enable the Command Center to Execute Site Swaps through Automation

- Automated Data Collection for Support Cases
  - Trigger storage system log collection
  - Open vendor support access
  - Create vendor case
  - Collect and FTP accessible software logs to the vendor
  - Trigger extended health check on our environment

- Automatically Start & Stop work sessions on weekends to ensure parity across all storage systems

- Implement the automated holiday SGC Holds

- Create an automated Selective Recovery Management process leveraging a Rocket Software product

Your feedback is important!

SHARE
EDUCATE • NETWORK • INFLUENCE

Submit a session evaluation for each session you attend:

www.share.org/evaluation

THANK YOU!