

Resilient IT Infrastructure (Infra)

By H Raghavan

Overview

- Introduction
- What is a resilient infra ?
- Purpose of resilient IT infra
- The journey, how it started and reached the destination
- The schematic of the resilient architecture
- Key features of the architecture
- Other IT services that need resiliency
- Challenges faced while executing the project
- Emerging challenges

Introduction

- Dictionary meaning of resilience: The positive capacity of an organizational system to adapt and return to equilibrium after a crisis, failure or any kind of disruption. Including: an outage, natural disasters, man-made disasters, terrorism or similar
- How can we check this on ourselves ?
- How resilient are you ?

- **What is a resilient infra ?**

- The term “resilient” refers to the ability of such infrastructure systems (including their interconnected ecosystems and social systems) to absorb disturbance and still retain their basic function and structural capacity.
- Ex.: Earthquake resilient structures, like dams and residential buildings
 - Twister resilient structures
 - A hybrid car which can run on dual fuel
 - An UPS

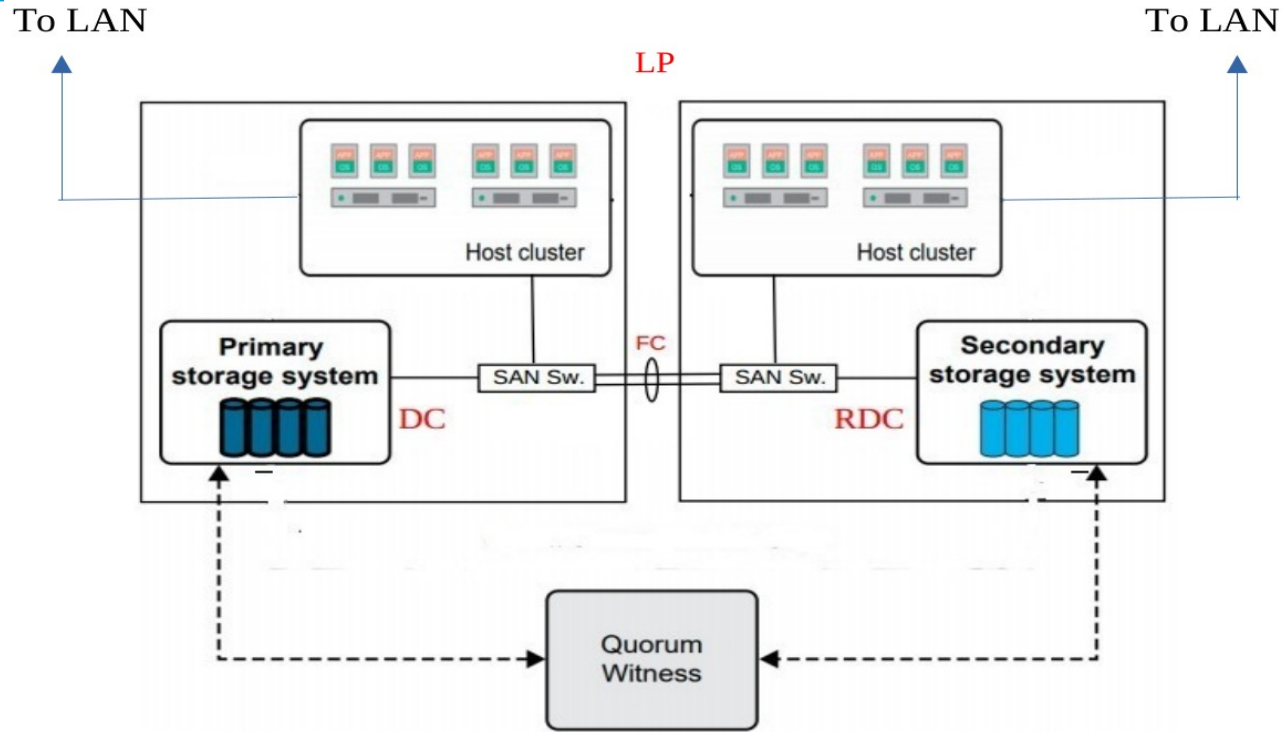
- Purpose of resilient IT infra

- To maintain acceptable service levels of business operations, critical processes or IT eco systems.
- To ensure business continuity even during a disaster
- To maintain IT services during breakdown and upgrade activities

• The journey, how it started and reached the destination

- Need of the hour, tech refresh was due in the year 2020. Bottlenecks were seen and breakdowns were increasing in the running system.
- Platform was already virtualized using VmWare and CC team was not ready to lose the edge
- Discussed with leading OEM's on the available technologies to meet the requirements. OEM's Nutanix, Hitachi, Dell, HPE and Lenovo (very time consuming process)
- Zeroed on Nutanix and VmWare
- Dropped Nutanix due to high cost and to avoid a new virtualization platform
- The final resilient solution was orchestrated on HPE storage platform which is compatible with VmWare along with data backup and restore solution

The schematic of the resilient architecture



DC = Data Center
RDC = Redundant Data Center
LP = Logical Partitioning
FC = Fiber Channel for Remote Copy

- **Key features of the architecture**


- Compression and De-duplication at storage level
- Real time data sync between DC and RDC
- Completely redundancy at storage, virtual machine and at network level
- To increase reliability, storage systems are with SSD drives
- Pro-active support from OEM
- Dedicated FC SAN network
- VmWare features are heavily utilized to achieve resilience
- The single solution we implemented cater to SAN and NAS requirements
- Scalable on demand

Concepts and definitions

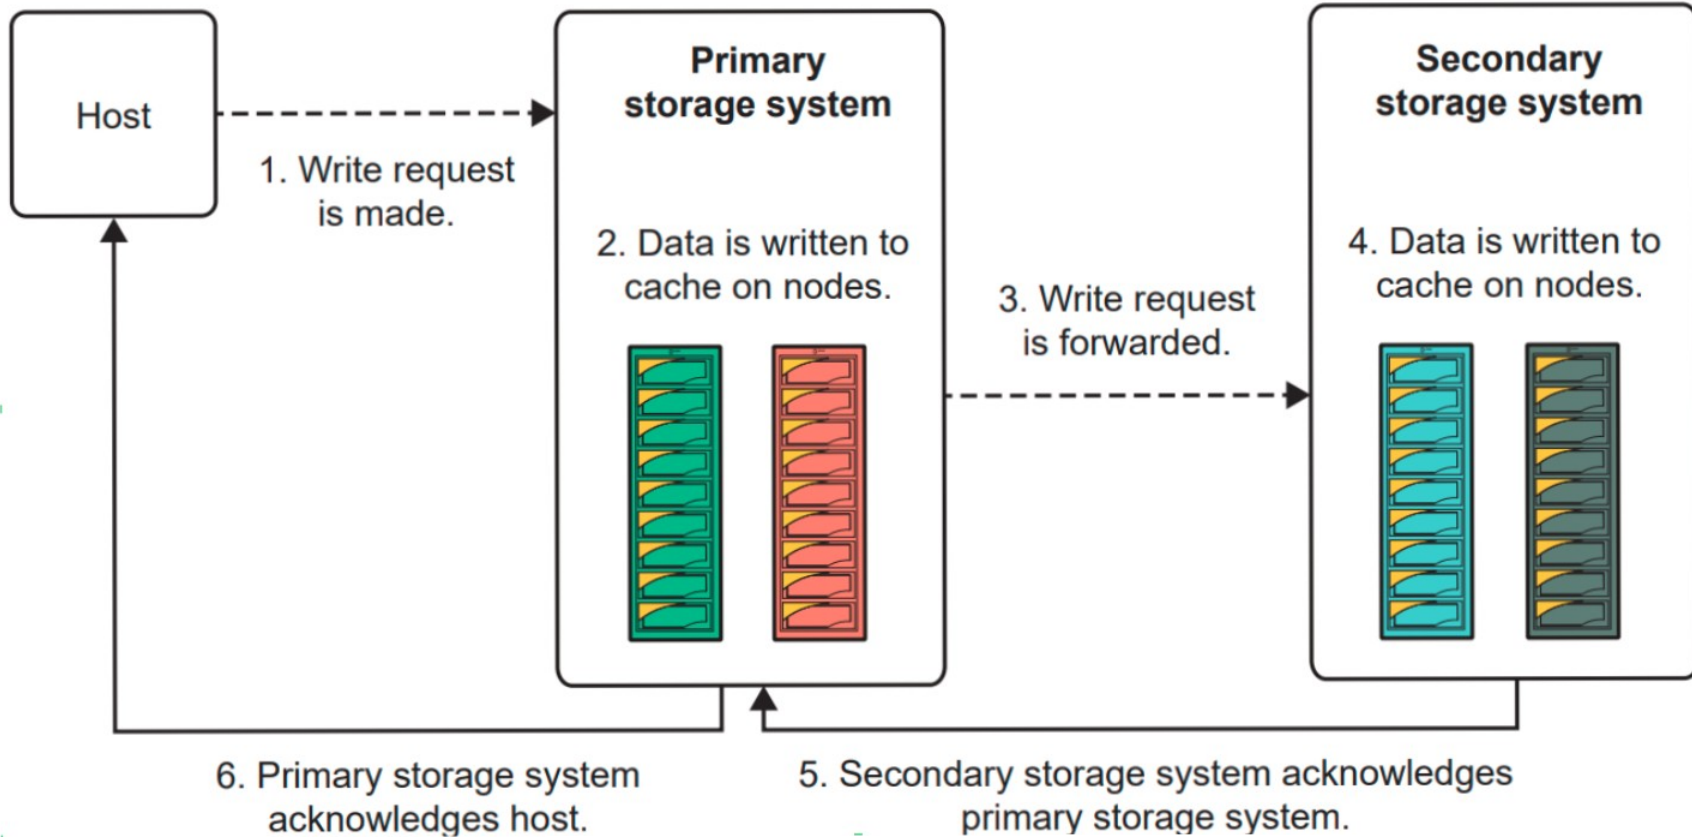
- De-duplication
- Real time data synchronization *
- Recovery Time Objective (RTO) *
- Recovery Point Objective (RPO) *
- Why multi-tenant architecture

- Tier-1 system characteristics

- Hardware redundancy, advanced replication capabilities, massive scalability, host connectivity, secure administrative segregation of users, hosts and application data, support mixed workloads

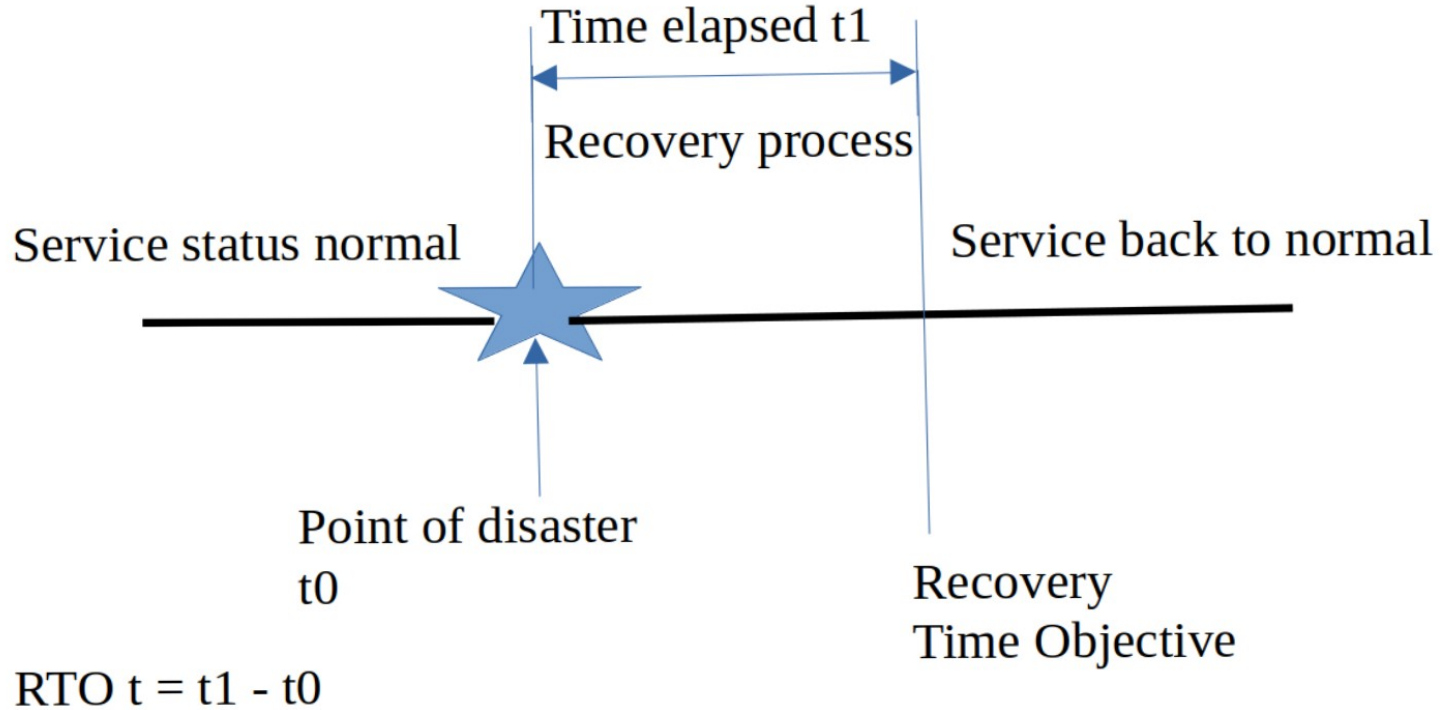


Realtime data synchronization



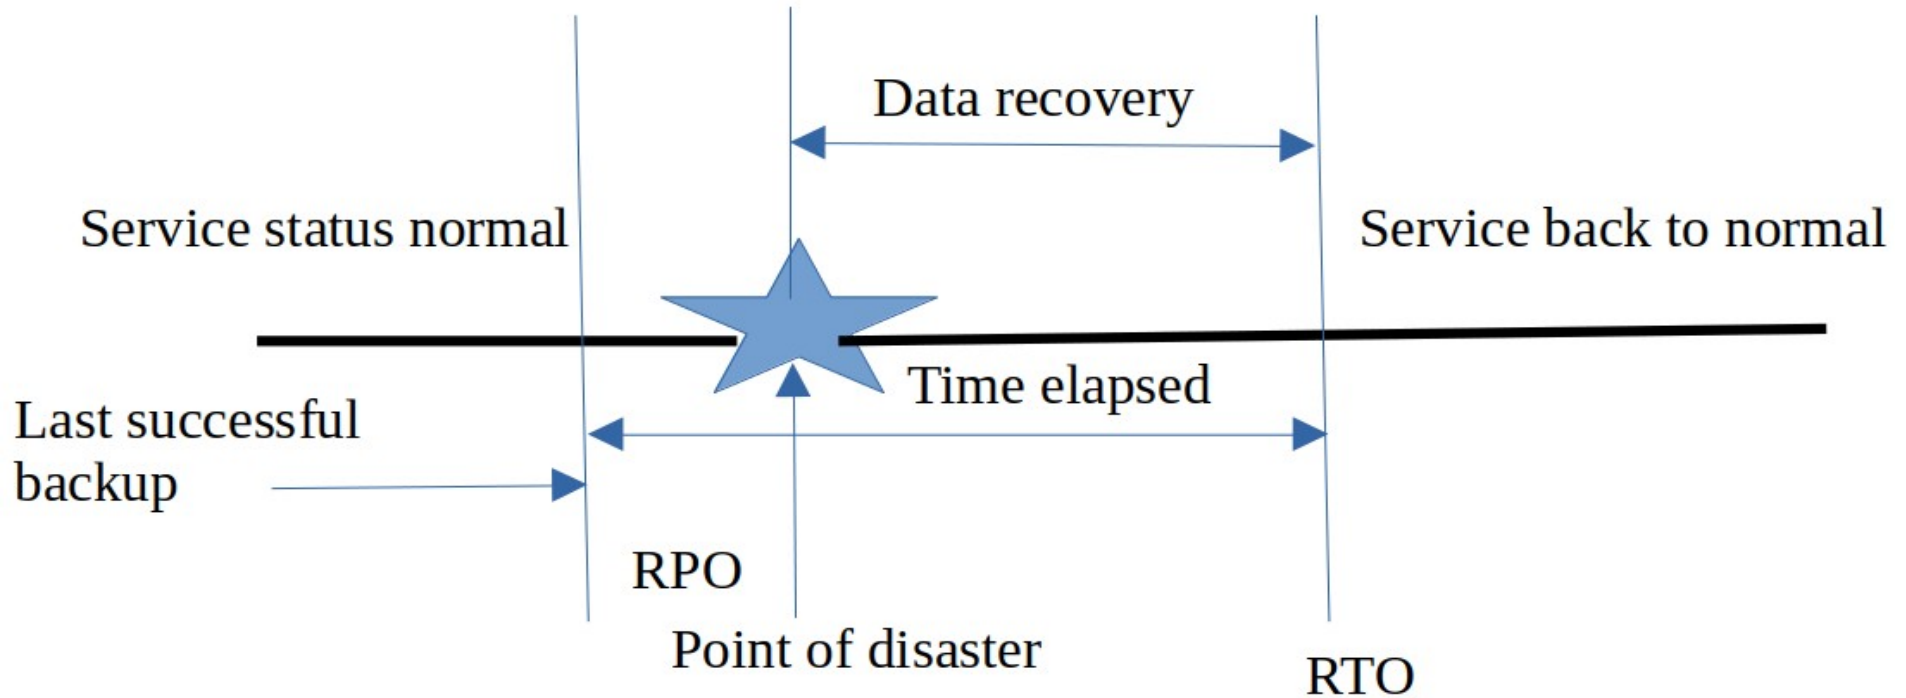
Recovery Time Objective (RTO)

- RTO



Recovery Point Objective (RPO)

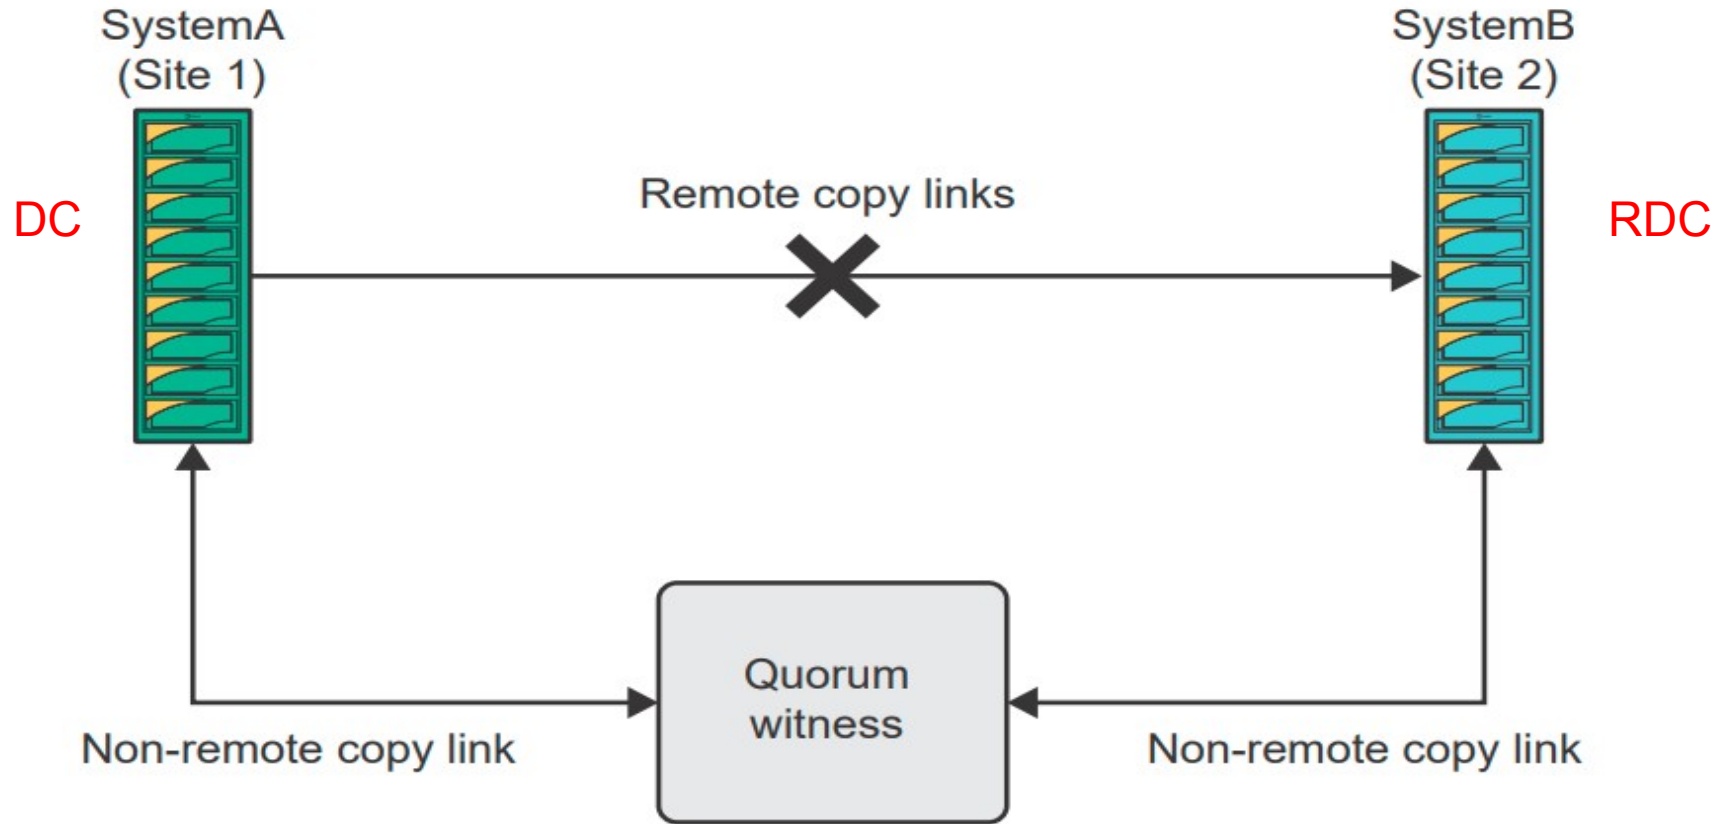
- RPO



Conditions to analyze failure scenarios

- The following conditions are considered for analysis
- During the remote copy link and communication failures scenarios, there are two sites with a quorum witness at a third site.
- There is an array at each site. DC (Site 1 or System A) is the source array and RDC (site 2 or System B) is the target in normal operations.
- The virtual machines are in stretched cluster configuration

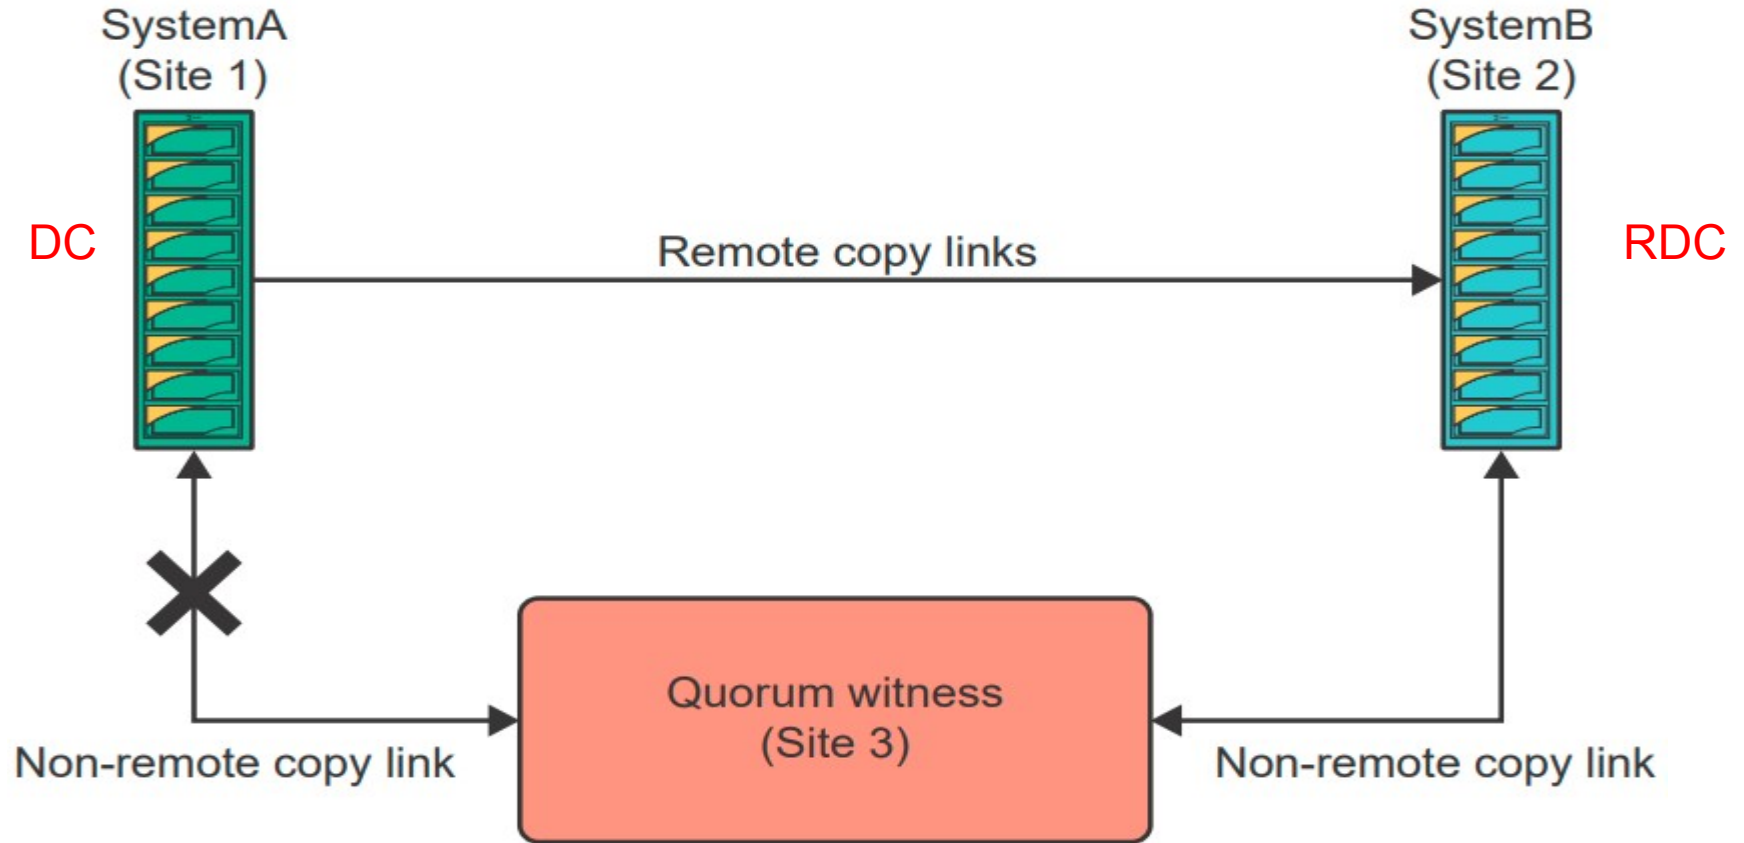
Failure scenario - 1



Failure scenario - 1

- Communication failure between array to array (DC & RDC)
 - Storage system failover does not occur
 - VM hosts continue to access storage volumes in DC
 - Replication of data between DC and RDC will stop
- Continued

Failure scenario - 2



Failure scenario - 2

- Communication failure between storage and quorum witness

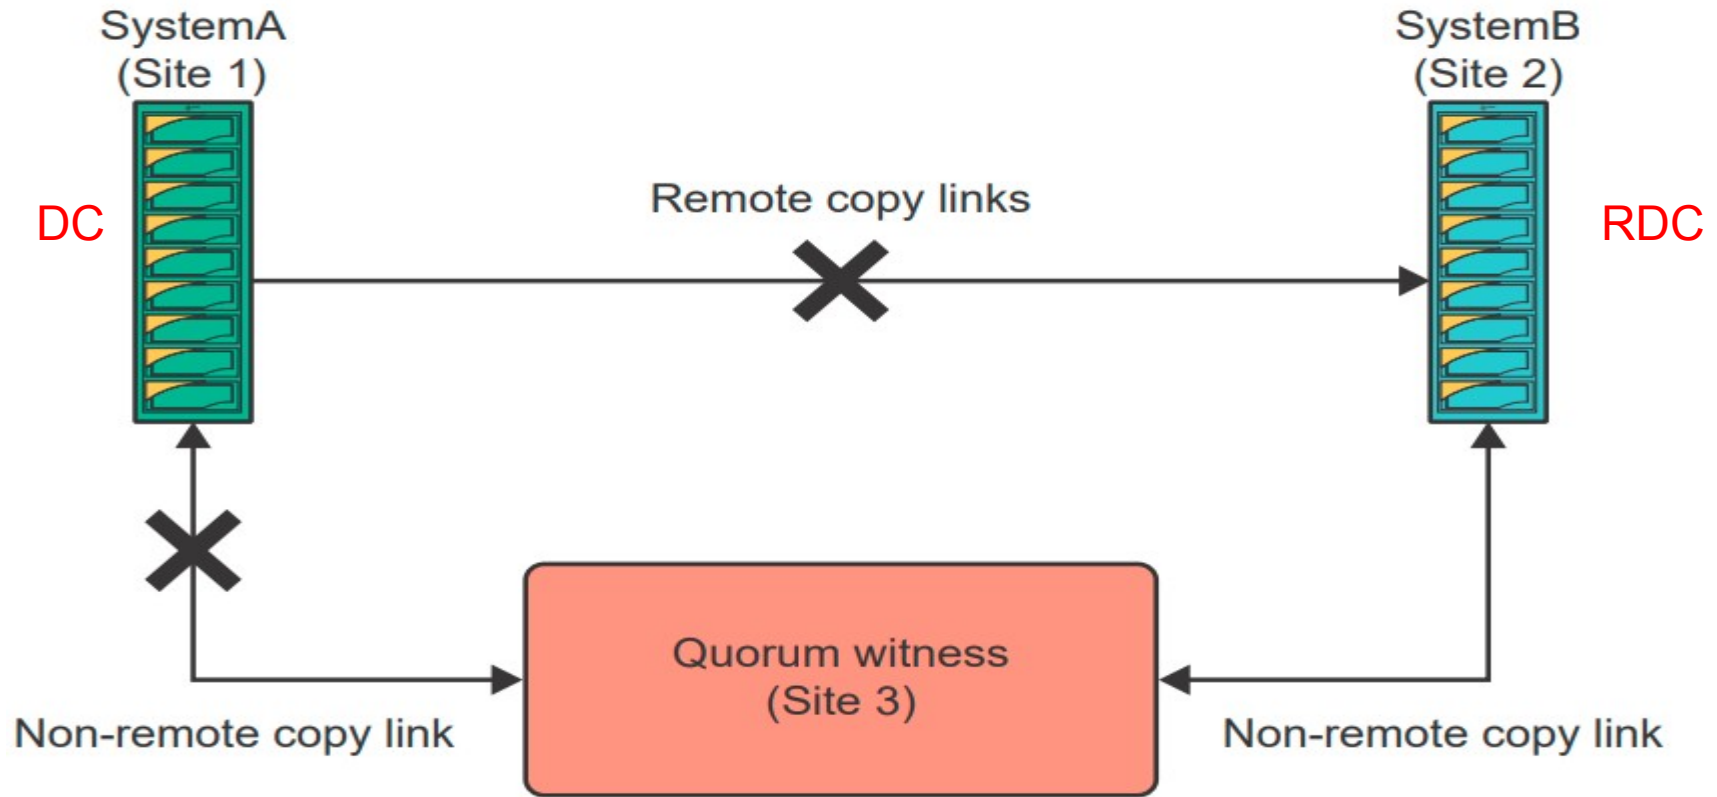
Storage system failover does not occur

VM hosts continue to access storage volumes in DC, Link between DC and RDC is available

Replication of data between DC and RDC will continue

- Continued

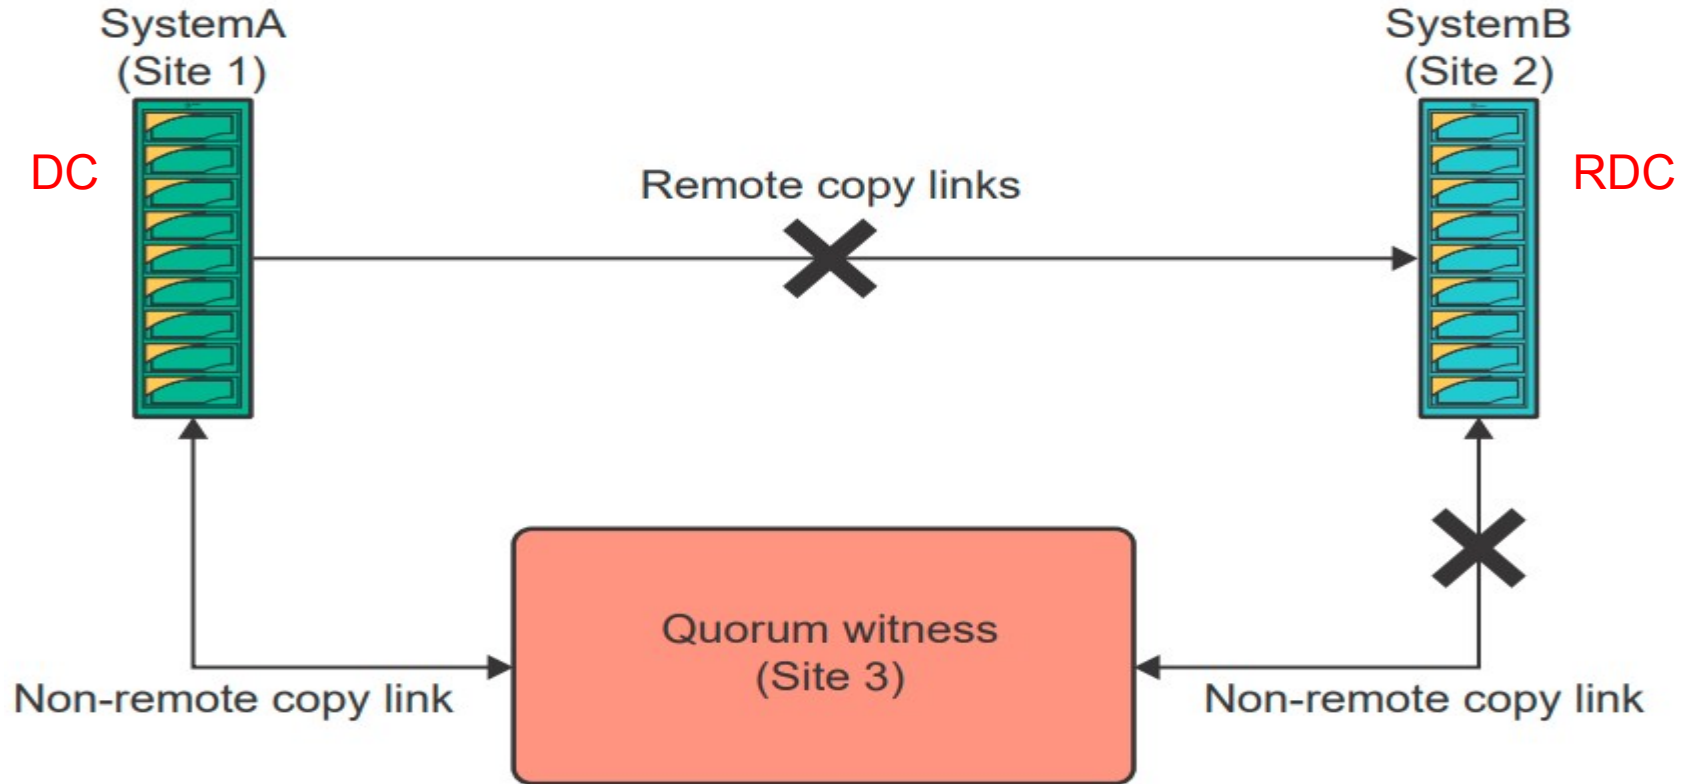
Failure scenario - 3



Failure scenario - 3

- Communication failure: DC to quorum witness and DC to RDC
 - Failover takes place due to dual network failure
 - VM host data requests goes to the NEW primary storage at RDC
 - Replication of data between DC and RDC will stop due to communication failure
- Continued

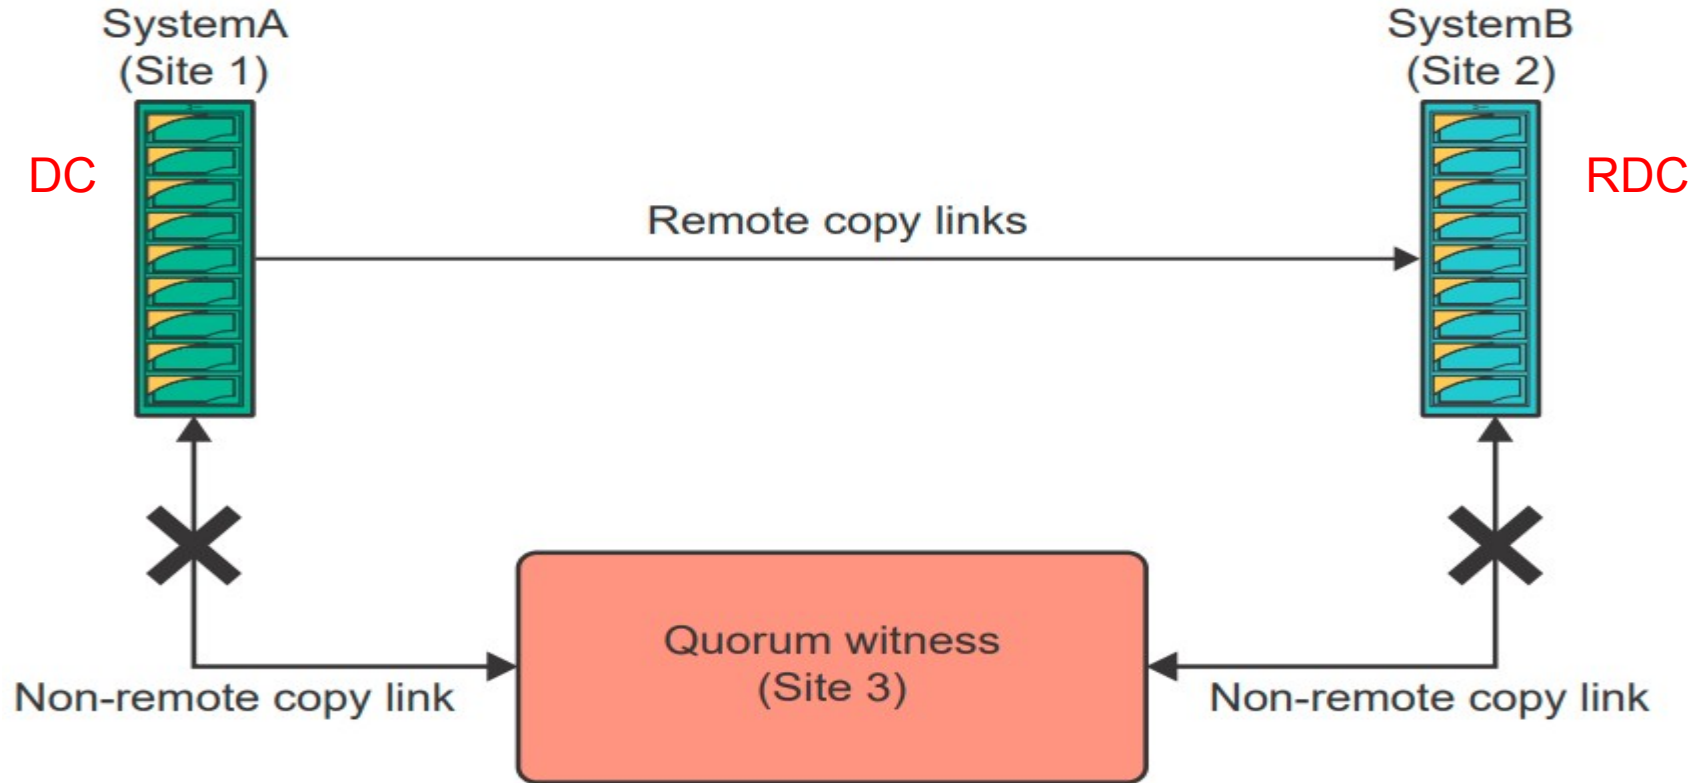
Failure scenario - 4



Failure scenario - 4

- Communication failure: RDC to quorum witness and DC to RDC
 - Storage system failover does not occur
 - VM hosts continue to access storage volumes in DC
 - Replication of data between DC and RDC will stop due to communication failure
- Continued

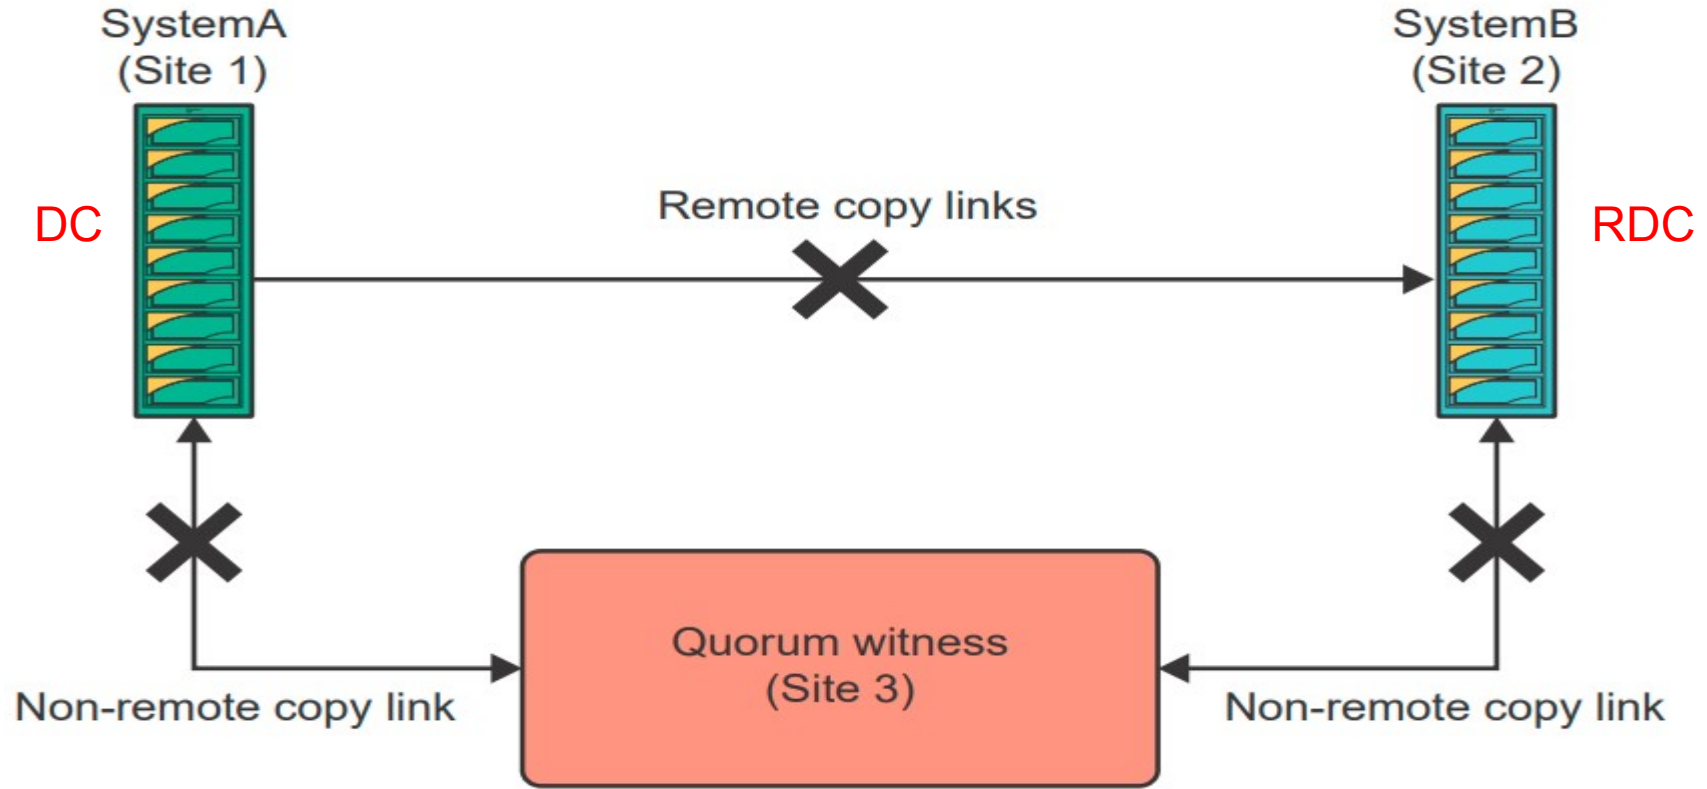
Failure scenario - 5



Failure scenario - 5

- Communication failure to quorum witness for both DC and RDC
 - Storage system failover does not occur
 - VM hosts continue to access storage volumes
 - Replication of data between DC and RDC will continue
- Continued

Failure scenario - 6



Failure scenario - 6

- Communication failure: DC and RDC to QW (Quorum Witness) and DC to RDC
 - Failover does not occur
 - Storage system at DC performs a FAILSAFE operation
 - Storage system at RDC will be the target
 - VM hosts data requests are stopped
 - Replication of data between DC and RDC will stop due to communication failure

- **Other IT services that need resiliency**

- We need to pay attention on the network part
 - Consider dual ISP links
 - Consider dual core switch
 - Have two DC's and have dual uplinks to all the access level switches
- Conditioned and emergency power
 - Run UPS's in parallel to improve the reliability
 - Modular UPS's help in scaling and helps to improve the reliability
 - Seek power from two different sources
 - Most important, consider DG power while designing your DC's

- **Challenges faced while executing the project**

- Had difficulties due to couple of staff was infected by Covid
- Severe shortcomings by OEM and non-compliance to RFP
- Arguments and meetings went on for three months
- Lack of sound technical manpower at the supplier end
- With all the challenges and headwind, installation took 12 months to accept and give the final signoff

- **Emerging challenges**

- AMC cost of VmWare has gone up by 15x.
- Very difficult to shift to new platform, hardware is still under warranty
- The current platform is tailor made taking advantage of HPE and VmWare product features
- Enormous time, planning and efforts are required to shift to new platform
- New platform needs approval and funding

- Thank you for your attention







Storage Systems 2 >



Status

Performance >



Status


Total Capacity > ⓘ



Total Capacity

Allocated	97,214 GiB
Free	45,778 GiB
Total	142,992 GiB

Allocated Capacity >



Allocated

Block	84,768 GiB
System	12,446 GiB
Total	97,214 GiB

Capacity Efficiency >



Efficiency

Savings	955,943 GiB
Used	66,441 GiB
Usable	1,022,384 GiB

15.4:1
Overall compaction


Device Type Capacity >



Status

SSD (142,992 GiB) 68%

Historical Capacity 2 >



Top Systems By Saturation

Performance insights is not enabled

Top App Vol Sets By Throughput

Performance insights is not enabled

Top App Vol Sets By Latency

Performance insights is not enabled

Virtual Volume Protection >

Unprotected	4 Volumes	200.00 GiB
Protected	44 Volumes	338,352.00 GiB
Local snapshots	7 Volumes	77,824.00 GiB

RMC Protection Jobs

[Configure RMC](#)

RMC Protected Volumes >



Unprotected (48 V...)	338,552.00 GiB
Protected (0 Vol...)	0.00 GiB