

# FIBRE CHANNEL SAN AUTOMATION AND ORCHESTRATION

*Rupin Mohan, Director R&D and CTO Storage Networking, HPE*

The IT world is moving toward Hybrid Cloud Infrastructure. Customers in large numbers are choosing applications that need to be deployed on premise but need some interactions and/or compute/data mobility with the cloud. Most tech savvy IT managers are aware of pros and cons of on-prem versus the cloud. Tier 0 and tier 1 applications with sensitivity to application up-time, failover disaster recovery, storage replication requirements continue to be deployed on premise in order to reduce business risk. Cost of downtime for most businesses continues to be very high. The cloud has certainly done an excellent job in providing speed and agility to customers. Cost advantages are debatable as customers are realizing over time that the cloud is not necessarily a cost-savings play. However, the speed and agility that the cloud brings to the table cannot be argued even with the best debaters in town.

Customers are looking for automation and orchestration for on-prem IT deployments including storage and SAN. One of the most important tasks SAN administrators perform in IT datacenters with Fibre Channel infrastructure is SAN zoning. The Fibre Channel Industry Association with T11 standards body has done some very innovative technical work to bring speed and agility to Fibre Channel. Let us explore more.

SAN plays a critical role in any data center by providing access and connectivity between storage arrays and servers via a dedicated network. Fibre Channel (FC) has been the dominant storage protocol that enjoys significant market share within SAN deployments. FC gained its popularity mainly because of its enterprise-class performance, availability and security.

Fibre Channel zoning is one key feature that adds to security and better management of the SAN by providing necessary segregation and allowing controlled communication among selected

devices within a large fabric. However, configuring zones still is a complex, tedious and error-prone operation in a majority of SAN installations. Therefore, there is a need for automating these operations as much as possible to avoid human errors and reduce potential SAN downtime.

Traditional zoning mechanisms have some limitations, and to overcome these, users have to create a large number of zones than necessary, which makes configuration of zones even more complex. To address these issues, the Fibre Channel Standards Technical Committee T11 developed a new type of zoning known as peer zoning. This has several advantages compared to traditional zoning.

## **Peer zoning**

As defined in Fibre Channel Generic Services (FC-GS) standards, peer zoning is a new way of enforcing zones in FC fabrics. By definition, peer zones are identified by a zone name and consist of principal members and peer members. As defined in FC standards, in a peer zone, there can be one or more principal members and similarly one or more peer members. Peer zones can coexist with traditional zones in the same active zone set.

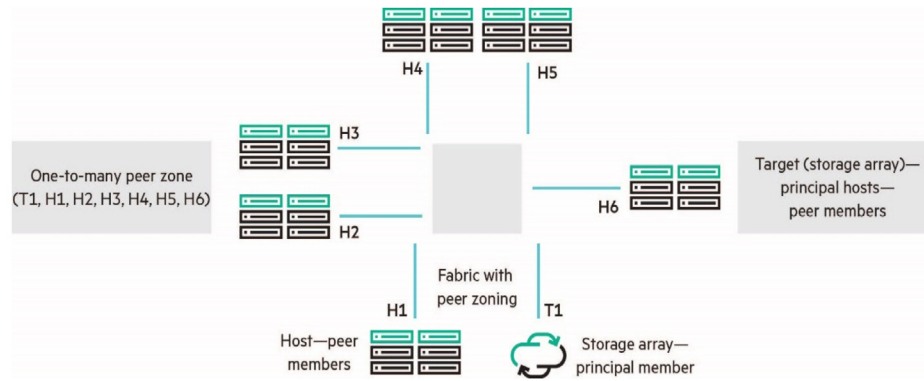
In a peer zone, by definition, communication between a principal member and a peer member is allowed, but communication is not allowed between any two peer members.

Peer zones can be configured by switch management tools or any device that has access to switch via in-band protocol. FC-GS standards define a set of three commands to read, add or update, and remove a peer zone from the active zone set. These commands are:

## **GAPZ—Get Active Peer Zone**

- This command is used to read peer zones from the existing switch database.

**Figure 1. Peer zoning example**



#### AAPZ—Add/Replace Active Peer Zone

- This command is used to configure a new zone or modify an existing zone.

#### RAPZ—Remove Active Peer Zone

- This command is used to remove a zone from an existing switch database.

There are other commands, but these are the three primary commands used for creation modification, and removal of peer zones.

#### Benefits of peer zoning

With traditional zoning, one-to-many zoning is not preferred, especially with one target and many hosts being in the same zone. The reason is that zoning permits communication between all members including hosts/initiators, which is unnecessary traffic that should be avoided. Often, fabric disruptions cause unnecessary change notification traffic that is irrelevant and is a hindrance for a majority of the members in that zone. To overcome this limitation, one-to-one (one target and one initiator) zones are recommended, but this quickly consumes valuable switch resources internally and also results in the creation of a large number of zones—resulting in inefficiency and wastage of time.

Peer zoning avoids all of these issues and at the same time, permits configuring one-to-many type zones as shown in Figure 1. In this example, there is one target port (T1) and multiple hosts (H1, H2, H3, H4, H5 and H6) that are part of the same zone. Here, communication is only allowed between a host and the target (storage array) but not between any two hosts. Similarly, any disturbance due to one of the host ports is only restricted to the target port and the rest of the host ports is shielded from receiving irrelevant traffic and change notifications.

In summary, peer zoning provides all of the benefits of initiator-based zoning (the most commonly used type) with less number of zones and optimal use of switch resources.

#### Target Driven Peer Zoning (TDPZ)

With simple peer zoning, zones still need to be configured manually either using the switch CLI or GUI, or some external management tool. The standard defines a very cool way of automating zoning driven by the storage array, target driven peer zoning (TDPZ).

With storage arrays implementing TDPZ, these zones can be automatically created by the array itself and thus avoid any zone pre-configuration requirement before hosts can be provisioned to the array. This significantly reduces the time needed to configure SANs and also makes them less error prone.

Without TDPZ, it is a requirement that an administrator must preconfigure zones on the switch before configuring hosts and VLUNs on the target side. With TDPZ, the user can now go directly to the element manager of the storage array (GUI/CLI) and start provisioning hosts to the array ports. All that the user needs to do is to ensure that peer zoning is enabled on the switches before starting provisioning of hosts.

Another major advantage of TDPZ, with minor extensions, is that users can read and display zones right from the storage array without going to the switch or any other management tool.

#### Summary

Customers can automate and orchestrate their SAN zoning in their data centers using peer zoning and TDPZ. This can result in a significant OPEX and time savings for them, depending on the size of their SAN due to the fact that number of zones created in their SAN is much lower number, plus the automation driven by TDPZ reduces the amount of time admins spend creating and managing SAN zoning during the lifecycle of the datacenter. Peer zoning and TDPZ, are clear examples of FCIA and T11 listening to customers pain points and complexities and their solving problems.