Define FC-IP Routing with IBM SAN42B-R

This document can be found on the web, www.ibm.com/support/techdocs
Search for document number WP##### under the category of "White Papers.

Version Date: July 08, 2016

IBM Switzerland

Björn Steiner / Michael Mettler
Adv. IT Specialist SSIS Switzerland
bstr@ch.ibm.com / mmet@ch.ibm.com

Table of Contents

Define FC-IP Routing with IBM SAN42B-R 1
Version Date: July 08, 2016 1
IBM Switzerland 1
Björn Steiner / Michael Mettler 1
Introduction 3
1. Overview 4
   Requirements 4
   Used components 4
2. FC-IP 4
   Architecture – High Level Overview 5
   Technical solution 6
3. Caveats 8
   Virtual fabric 8
   Isolation and path failover 8
   Dependency of ports and ASICs 8
   Interopability with Cisco 8
4. Implementation 9
   Preparation Core SAN 9
   Preparation FCIP-Routers 9
   Limitations and solutions 9
      VEx-Port 9
      Traffic isolation 9
      Solution 9
   Overview – FCIP Switches 10
      DC1 10
      DC2 10
   Overview – Fabric 1 11
   Overview – Fabric 2 12
   Step-by-Step configuration 13
   Verification 18

Table of Figures

Figure 1: Architecture – High Level Overview 5
Figure 2: Differentiation between ports 6
Figure 3: Routing concept 7
Figure 4: Overview – Fabric 1 11
Figure 5: Overview – Fabric 2 12

Introduction

This paper describes the implementation of an edge-backbone-edge architecture based on SAN42B-R products. It is intended to provide all required implementation steps on SAN42B-R product as well as the architectural decisions made to achieve the customers and products requirements.

The audiences for this document are Storage Specialists and IT Architects.

1. Overview

This chapter provides a brief overview of the given requirements and the used components.

Requirements

The following requirements were given:

- A four site design based on a SVC enhanced stretched cluster within 1 location (two sites) and a Disaster Recovery relationship to the other location (two sites) with another SVC enhanced stretched cluster.
- There is no native Fibre Channel available between DR locations (about 20km)
- connection between DR locations may also be used by different SANs (Fabrics)
- Separation of SAN connectivity between DR locations as well as with other SANs in each location
- The SAN environments are based on Brocade products (IBM or different OEMs)

Used components

The following components were used:

- SVC (v7.4.05); 8-node Cluster enhanced stretched cluster (ESC); DH8 with 4*2port 16GB card
- SAN384B-2 as Core directors
- SAN42B-R as FC-IP Router
- V7000 as extended Quorum for the ESC

2. FC-IP

Due there is no native Fibre Channel available between two datacenters, a technology is needed to convert Fibre Channel protocol to the available IP protocol used in the LAN environment.
FC-IP is the technology to achieve this and supports the used components.

Architecture – High Level Overview

For this concept only the connection between the 2 datacenter is relevant. The connections of the local SAN and requirements for SVC enhanced stretched cluster are not covered in this document. The picture below (Figure 1) shows a high level overview of the applied solution - an **edge-backbone-edge** architecture.

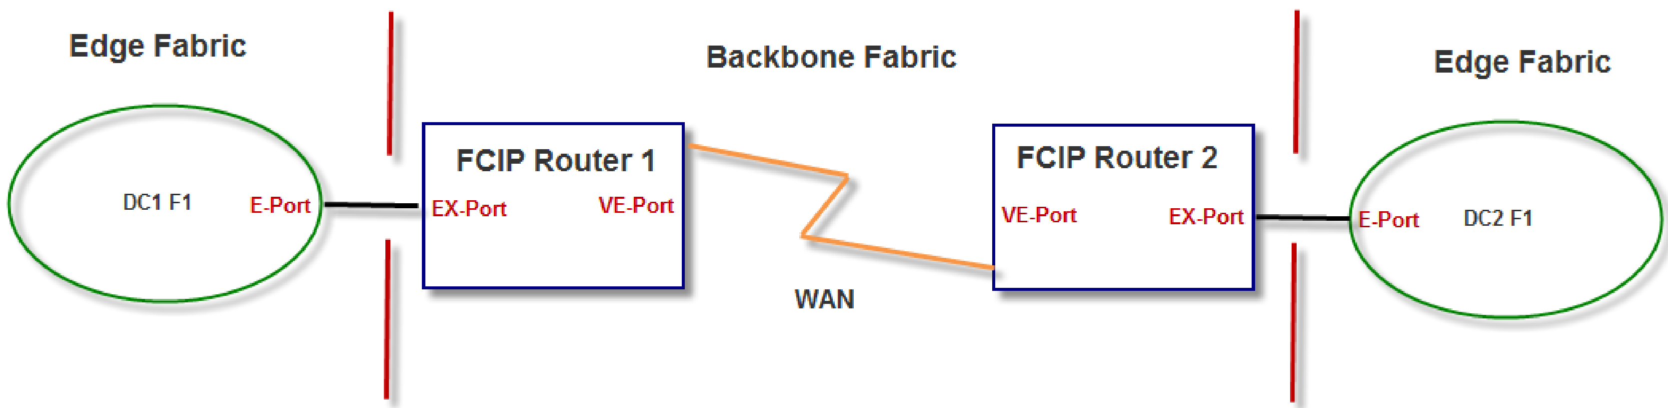

Figure 1: Architecture – High Level Overview

Technical solution

The solution is based on a SAN42B-R product.

FC-IP Router in general

FC-IP Routers are used to convert FC-protocol to IP-protocol.
Connections to the SAN are provided by FC-ports, connections to LAN or WAN by Ethernet-ports.
This product is integrated as a normal FC switch to the logical SAN infrastructure. Interoperability to dedicated SAN or LAN infrastructure should be checked in advance.

SAN42B-R

At the time of this documentation this is the smallest FC-IP router based on Brocade that provide 10GB Ethernet and 16Gb FC-ports.

The product provides following ports:

- 24 Fibre Channel Ports for 16Gb
- 16 Ethernet Ports for 1 Gb or 10Gb
- 2 Ethernet Ports for 40Gb

Nomenclature FC-IP Router

E /Ex -port	Connection Port between two FC-Switches
VE /VEx -port	emulates an E port on each end of an IP tunnel
circuit	Connection between two IP- addresses
tunnel	Connection between FC-Port (VE/Vex) and Ethernet port (circuit)
route	Logical definition for usage of a tunnel
ISL	Inter Switch Link (connection between to SAN switches)

Differentiation between E and Ex Port resp. VE and Vex port

An InterswitchLink is a connection between two physical ports of two SAN directors, defined as E-Ports.
EX_Port is a firewall when used in conjunction with FCR at a designated FC E_Port.
VE_Port emulates an E_Port on either end of a FCIP tunnel.
VEX_Port functions as an E_Port, handles mapping between FC and TCP/IP plus acts as a firewall. Fabrics will not merge.

E port types	No FCE	FCR
Native FC	E_Port	EX_Port
Extended over tunnel	VE_Port	VEX_Port

Figure 2: Differentiation between ports

NB: There are many possibilities to define a Routed Fabric. It has to be considered at which point routing function should be implemented.

Configuration Concept

In the concept of this document the routing is in fact not an IP-routing but it is an extension of a Fabric with IP. In our example, routing is implemented between the EX and VE-Port, as shown in Figure 3.

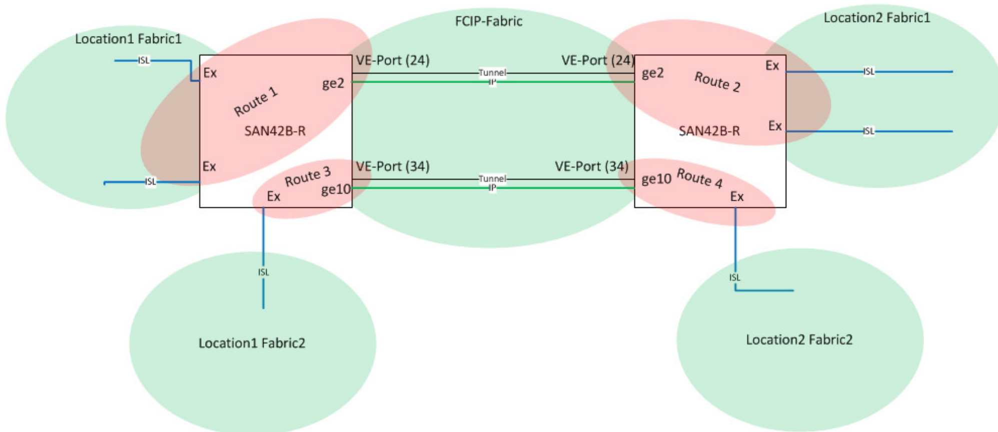

Figure 3: Routing concept

Configuration FC-IP Router

The FC-IP Routers are configured as a normal switch within the existing fabric. A fabric merge between the two DCs is avoided by EX port configuration. The connection to the SAN is in a manner that only one path (IP) is available between the DCs. To have the redundancy over the two HA-location in one fabric, the router is placed within site1 on both DCs and in the other fabric it is placed within the site2. With this configuration there will be no disruption between the two locations if one of the sites fails.

Configuration SAN

To have redundancy for the connection between router and director 2 connections will beadded to the director. These connections will be placed on each blade to have the ability for online maintenance on a blade regarding the fabric connections.

Only the Public Fabrics are connected to FC-IP Routers. The private fabrics are used only within a dedicated site.

3. Caveats

This chapter covers some limitations and issues faced in this implementation.

Virtual fabric

If an usage of virtual fabrics is required, the limitations and restrictions of this feature have to be considered before implementation due they can affect the design.

- Use of VEx ports
- Use of TI Zones

Isolation and path failover

To isolate the traffic between different tunnels, TI-Zones (traffic isolation zones) has to be implemented. This is an additional zoning that is implemented on the device where the IP connection is done.
During a fail over of a tunnel within a dedicated TI zone, the traffic replication is covered by the next available path within the same TI zone.If there is no second path available within a specific TI zone, the next available dedicated or non-dedicated path will be used instead.

A dedicated path belongs to a specific TI zone, a non-dedicated path does not belong to a specific TI zone. To avoid a fail over from a path within a dedicated TI zone to another dedicated path within another TI zone (or to a non-dedicated path), path fail over has to be disabled. As a result, no paths are available within a specific TI zone, the traffic replication will be covered by the FCIP router within the other fabric.

Since this setup is only available if virtual fabrics are disabled, use of this function must be excluded.

Dependency of ports and ASICs

In the SAN42B-R there are 2 ASICs (dp0 and dp1) that are responsible for workload on the Ethernet ports.
Therefore check which ports must be redundant to each other.
During a firmware update, the ASICs will be shut down (each after the other). Since a tunnel may consist of more than one IP-channel, redundancy can be achieved that way and failover happens within the tunnel. Or redundancy on the environment is available over a second fabric.

Interopability with Cisco

Behaviour of an Ethernet port on a Brocade FCIP-router is defined by standards and do not fully operate with proprietary functions of the network infrastructure.
Cisco Nexus are used in this environment as network switch for the IP connection between the locations. On the Cisco ports the function for UDLD is activated.
Brocade recommends disabling UDLD when connecting SAN42B-R switches to Cisco Nexus Switches.
If UDLD is used and cannot be disabled on these ports, the workaround is to disable the VE (VEx) port used for this tunnel and then disable the tunnel. With these steps Cisco recognizes a defined state down and not an unexpected one and the port on the Cisco stays up.

4. Implementation

This chapter provides informations for the design and implementation as well as a complete example how this design can be achieved.

Preparation Core SAN

The Core SAN do not have any special prerequisites before starting the routing. With this configuration there is also no additional feature necessary since the FCIP-router will join the fabric as normal switch.

Preparation FCIP-Routers

The FCIP-router must be configured as normal switch.
In our case it's necessary that the virtual Fabric functionality is disabled to use the failover functionality.

Limitations and solutions

The following limitations had to be considered:

VEx-Port

VEx ports are not supported on this kind of Router. Only the Director Blade supports VEx ports.

Traffic isolation

Traffic isolation is possible by creating virtual Fabrics or TI-Zones.
Virtual Fabrics could not be used since VEx ports are not supported on the used products.
TI-Zones with failover disabled are not allowed in virtual Fabrics (base switches).

Solution

First concept was based on virtual Fabrics to isolate all workload logically. Since this was no longer possible because of the product is not supporting Vex Ports and due the requirement to prohibit tunnel failover, virtual Fabrics were disabled on the FCIP-Routers.
Technically this is possible since the FCIP-Router builds its own Fabric and tunnels isolate the traffic.
Isolation of traffic was a requirement to monitor the performance related to the environment in case of impacts and billing.

Overview – FCIP Switches

In the following example, there are 2 FCIP Switches per data center. Each FCIP router is used for a dedicated fabric. DC1 FCIP 1 and DC2 FCIP 1 are connected to SAN directors within fabric 1, DC1 FCIP 2 and DC2 FCIP 2 are connected to SAN directors within fabric 2.

DC1

Switch name	**IP interface**
DC1 FCIP1	10.10.1.11
DC1 FCIP1	10.10.1.12

Switch name	**IP interface**
DC1 FCIP2	10.10.1.13
DC1 FCIP2	10.10.1.14

DC2

Switch name	**IP interface**
DC2 FCIP1	10.20.1.11
DC2 FCIP1	10.20.1.12

Switch name	**IP interface**
DC2 FCIP2	10.20.1.13
DC2 FCIP2	10.20.1.14

Overview – Fabric 1

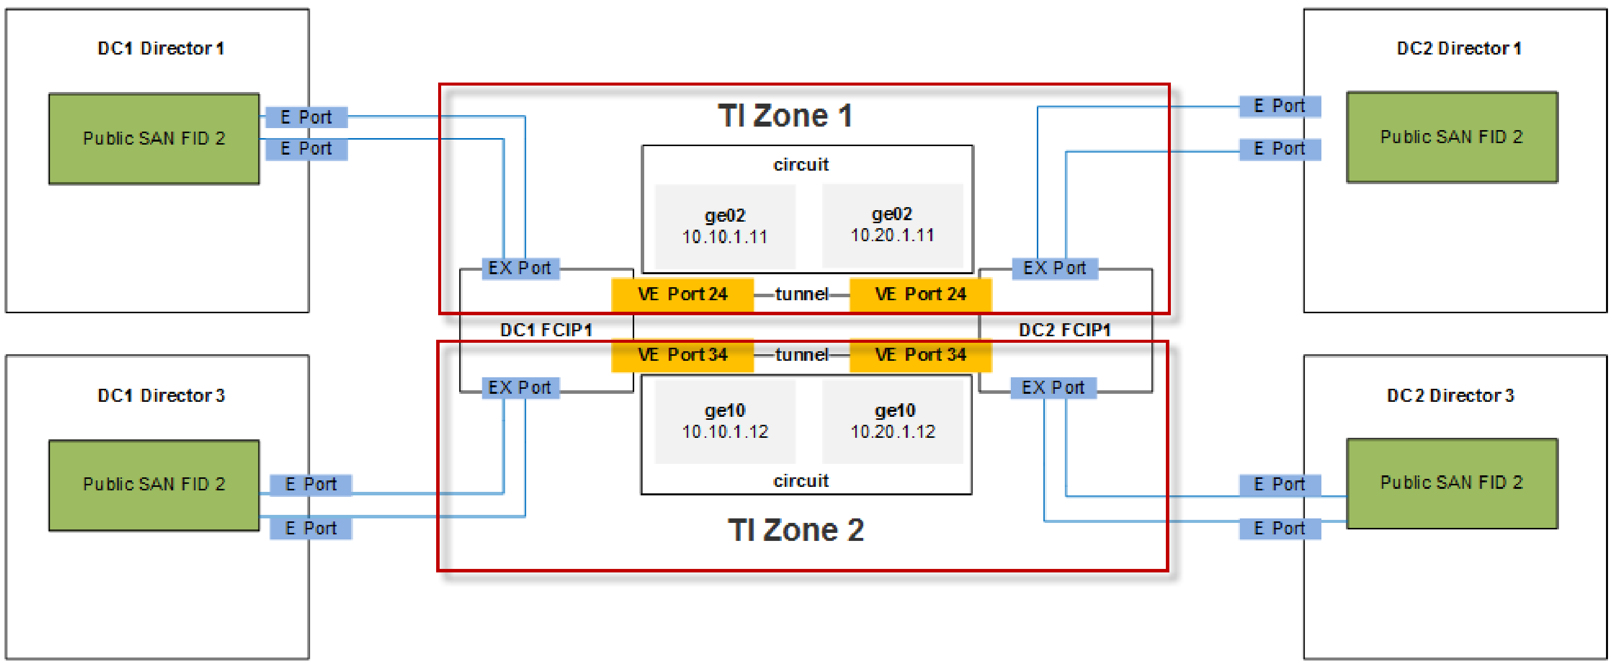

Figure 4: Overview – Fabric 1

Overview – Fabric 2

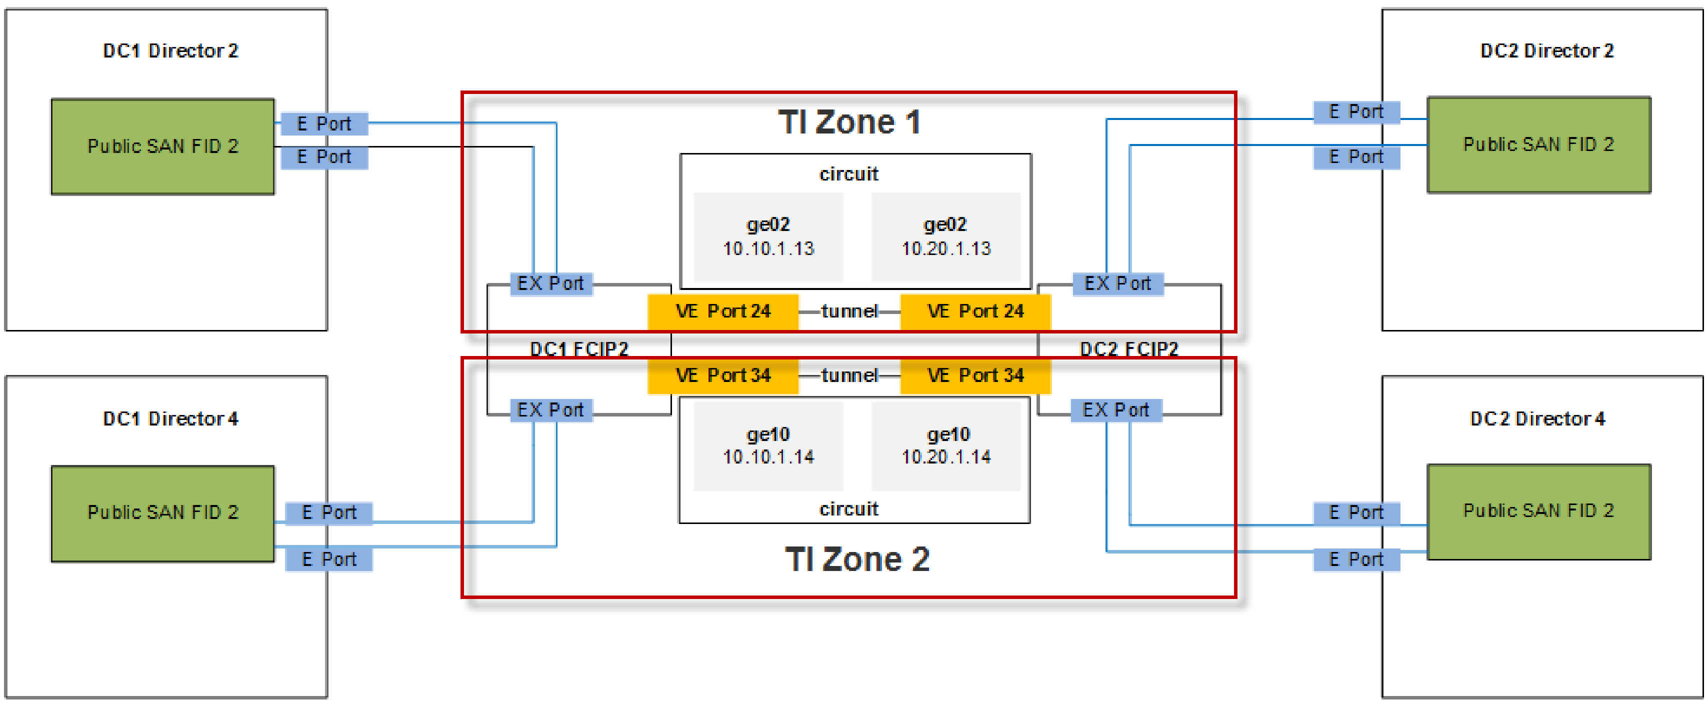

Figure 5: Overview – Fabric 2

Step-by-Step configuration

This Step-by-Step shows only the relevant configuration steps. Basic-knowledge about zoning and configuration of SANs must exist.

- Setup the FCIP routers as normal switches
- Per default these switches are configured with virtual fabrics on. This function must be disabled. Before start remove all virtual Fabrics if already configured
  - fosconfig –show (check state of virtual Fabrics)
  - fosconfig --disable vf (disables the virtual Fabrics)
    - ➔ this command requires a reboot of the switch
- Apply the required licenses
  - Integrated Routing must be available
  - WAN rate upgrade may be required
  - Additional Ports usage may be required
- Enable routing (if not already done)
  - fosconfig –show (check state of routing)
  - fosconfig –enable fcr
- Create the IP interfaces
  - **FCIP1 location1**
  - portcfg ipif ge2.dp0 create 10.10.1.11/27 vlan 1206 mtu 9000
  - portcfg ipif ge10.dp1 create 10.10.1.12/27 vlan 1206 mtu 9000
  - **FCIP2 location1**
  - portcfg ipif ge2.dp0 create 10.10.1.13/27 vlan 1206 mtu 9000
  - portcfg ipif ge10.dp1 create 10.10.1.14/27 vlan 1206 mtu 9000
  - **FCIP1 location2**
  - portcfg ipif ge2.dp0 create 10.20.1.11/27 vlan 323 mtu 9000
  - portcfg ipif ge10.dp1 create 10.20.1.12/27 vlan 323 mtu 9000
  - **FCIP2 location2**
  - portcfg ipif ge2.dp0 create 10.20.1.13/27 vlan 323 mtu 9000
  - portcfg ipif ge10.dp1 create 10.20.1.14/27 vlan 323 mtu 9000
- Create the IP routes
  - **FCIP1 location1**

- portcfg iproute ge2.dp0 create 10.20.1.0/27 10.10.1.11
- portcfg iproute ge10.dp1 create 10.20.1.0/27 10.10.1.12

- **FCIP2 location1**
- portcfg iproute ge2.dp0 create 10.20.1.0/27 10.10.1.13
- portcfg iproute ge10.dp1 create 10.20.1.0/27 10.10.1.14
- **FCIP1 location2**
- portcfg iproute ge2.dp0 create 10.10.1.0/27 10.20.1.11
- portcfg iproute ge10.dp1 create 10.10.1.0/27 10.20.1.12
- **FCIP2 location2**
- portcfg iproute ge2.dp0 create 10.10.1.0/27 10.20.1.13
- portcfg iproute ge10.dp1 create 10.10.1.0/27 10.20.1.14

- Test IP-connections
  - **FCIP1 location1**
  - portcmd –ping ge2.dp0 –s 10.10.1.11 –d 10.20.1.11
  - portcmd –ping ge2.dp0 –s 10.10.1.12 –d 10.20.1.12
  - **FCIP2 location1**
  - portcmd –ping ge2.dp0 –s 10.10.1.13 –d 10.20.1.13
  - portcmd –ping ge2.dp0 –s 10.10.1.14 –d 10.20.1.14
  - **FCIP1 location2**
  - portcmd –ping ge2.dp0 –s 10.20.1.11 –d 10.10.1.11
  - portcmd –ping ge2.dp0 –s 10.20.1.12 –d 10.10.1.12
  - **FCIP2 location2**
  - portcmd –ping ge2.dp0 –s 10.20.1.13 –d 10.10.1.13
  - portcmd –ping ge2.dp0 –s 10.20.1.14 –d 10.10.1.14
- Create FCIP-Tunnel

  Fabric 1:

  - **FCIP1 location1**
  - portcfg fciptunnel 24 create --local-ip 10.10.1.11 --remote-ip 10.20.1.11 -b 1500000 -B 1500000 -c deflate --ipsec none

- portcfg fciptunnel 34 create --local-ip 10.10.1.12 --remote-ip 10.20.1.12 -b 1500000 -B 1500000 -c deflate --ipsec none
- **FCIP1 location2**
- portcfg fciptunnel 24 create --local-ip 10.20.1.11 --remote-ip 10.10.1.11 -b 1500000 -B 1500000 -c deflate --ipsec none
- portcfg fciptunnel 34 create --local-ip 10.20.1.12 --remote-ip 10.10.1.12 -b 1500000 -B 1500000 -c deflate --ipsec none
    - -b 1500000 –B 1500000 limits the Bandwidth to 1,5MB
    - The tunnels between VE-Port 24 in both locations are created in both direction
    - The tunnels between VE-Port 34 in both locations are created in both direction

Fabric 2:

- **FCIP2 - location1**
- portcfg fciptunnel 24 create --local-ip 10.10.1.13 --remote-ip 10.20.1.13 -b 1500000 -B 1500000 -c deflate --ipsec none
- portcfg fciptunnel 34 create --local-ip 10.10.1.14 --remote-ip 10.20.1.14 -b 1500000 -B 1500000 -c deflate --ipsec none
- **FCIP2 - location2**
- portcfg fciptunnel 24 create --local-ip 10.20.1.13 --remote-ip 10.10.1.13 -b 1500000 -B 1500000 -c deflate --ipsec none
- portcfg fciptunnel 34 create --local-ip 10.20.1.14 --remote-ip 10.10.1.14 -b 1500000 -B 1500000 -c deflate --ipsec none

- The Backbone is now ready and merges to one single Fabric
- Configure all ports as Ex-ports that are connected to the fabrics in location1 and 2
  - portcfgexport 1 –a 1 –f 100
  - portcfgexport 2 –a 1 –f 110
    - -f is the fabricID the port should connect to
- The E-Ports to the existing fabrics in location 1 and 2 can be done
- The Fabrics should “merge” with the Ex-Port they are connected to
  - On all switches /directors
  - fcrfabricshow
    - This command shows on all switches/directors its routed fabrics
- With LSAN-Zoning the routed zones can be implemented:
  - In this case this are the ports between the SVC for Mirroring and the Ports of the v7000 to act as Quorum
    - LSAN Zones are normal zones which start with “LSAN_” before the zone name , all rules for normal zones are also valid for LSAN-Zones

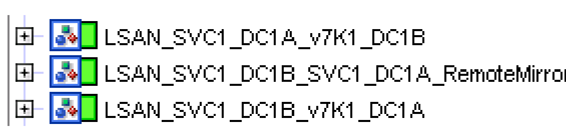

  - The LSAN-Zones have to be defined in both Fabrics in location1 and location2
- After verification of all the connections the TI-Zones must be implemented for traffic isolation and to prohibit failover:

- The TI Zones start with “TI_” before the zone name and contain the port (domain,port) of the members in these zones.
- zone --create -t ti -o dn "TI-SAP " -p "1,0; 1,1; 1,2; 1,3; 2,0; 2,1; 2,2; 2,3; 1,24; 2,34"
  - -t defines that it act as a TI-Zone
  - -o dn denies failover
  - Port 1,x are of FCIP location1 (where 1,24 is the tunnel)
  - Port2,x are of FCIP location2 (where 2,24 is the tunnel)

```
TI Zone Name:    TI_SAP
Port List:       1,0; 1,1; 1,2; 1,3; 2,0; 2,1; 2,2; 2,3; 1,24; 2,24

Configured Status: Activated / Failover-Disabled
Enabled Status: Activated / Failover-Disabled
```

  - Since a zone cannot exist without a config but a TI-Zone is not a normal Zone that can be enabled and disabled as the other zones there must be a zone configured to add to the config.
  - In this case all the LSAN-Zones are added to the Backbone fabrics as well. These zones were added to the cfg and activated. With this activation automatically the TI-Zone will be activated.

Verification

In the following references, DC1A corresponds to DC1 and DC1B corresponds to DC2.

```
DC1A_FCIP1:bstr> fcrfabricshow
FCrouter WWN: 10:00:50:eb:1a:99:be:6f, Dom ID:   1,
Info: 10.16.193.142, "DC1A_FCIP1"
   EX_Port        FID    Neighbor Switch Info (enet IP, WWN, name)
   ------------------------------------------------------------------------
      0           2  xxx.xxx.xxx.xx1  10:00:50:eb:1a:[redacted]  "DC1A_Director1_Public"
      1           2  xxx.xxx.xxx.xx2  10:00:50:eb:1a:[redacted]  "DC1A_Director1_Public"
      2           2  xxx.xxx.xxx.xx3  10:00:50:eb:1a:[redacted]  "DC1A_Director3_Public"
      3           2  xxx.xxx.xxx.xx4  10:00:50:eb:1a:[redacted]  "DC1A_Director3_Public"

FCrouter WWN: 10:00:50:eb:1a:b5:c2:b1, Dom ID:   2,
Info: 10.17.236.203, "DC1B_FCIP1"
   EX_Port        FID    Neighbor Switch Info (enet IP, WWN, name)
   ------------------------------------------------------------------------
      0          20  xxx.xxx.xxx.xx5  10:00:50:eb:1a:[redacted]  "DC1B_Director1_Public"
      1          20  xxx.xxx.xxx.xx6  10:00:50:eb:1a:[redacted]  "DC1B_Director1_Public"
      2          20  xxx.xxx.xxx.xx7  10:00:50:eb:1a:[redacted]  "DC1B_Director3_Public"
      3          20  xxx.xxx.xxx.xx8  10:00:50:eb:1a:[redacted]  "DC1B_Director3_Public"
```

Verifying final device routing.

The following output of the lsanzoneshow –s command confirms that the LSAN configuration is valid in the backbone fabric. Devices in the edge fabric DC1 exisit with the FID 2 and devices from the edge fabric DC2 with the FID 20 are imported.

```
DC1A_FCIP1:bstr> lsanzoneshow -d
Fabric ID: 2 Zone Name: LSAN_SVC1_DC1A_V7K1_DC1B
        50:05:07:68:0c:11:[redacted]  EXIST in FID 2
        50:05:07:68:0c:11:[redacted]  EXIST in FID 2
        50:05:07:68:0c:21:[redacted]  EXIST in FID 2
        50:05:07:68:0c:11:[redacted]  EXIST in FID 2
        50:05:07:68:0c:11:[redacted]  EXIST in FID 2
        50:05:07:68:0c:21:[redacted]  EXIST in FID 2
        50:05:07:68:0c:21:[redacted]  EXIST in FID 2
        50:05:07:68:0c:21:[redacted]  EXIST in FID 2
        50:05:07:68:0c:11:[redacted]  EXIST in FID 2
        50:05:07:68:0c:21:[redacted]  EXIST in FID 2
        50:05:07:68:0c:21:[redacted]  EXIST in FID 2
        50:05:07:68:0c:11:[redacted]  EXIST in FID 2
        50:05:07:68:0c:21:[redacted]  EXIST in FID 2
        50:05:07:68:0c:11:[redacted]  EXIST in FID 2
        50:05:07:68:0c:11:[redacted]  EXIST in FID 2
        50:05:07:68:0c:21:[redacted]  EXIST in FID 2
        50:05:07:68:0b:31:[redacted]  Imported from FID 20
        50:05:07:68:0b:33:[redacted]  Imported from FID 20
        50:05:07:68:0b:31:[redacted]  Imported from FID 20
        50:05:07:68:0b:33:[redacted]  Imported from FID 20
Fabric ID: 2 Zone Name: LSAN_SVC1_DC1B_V7K1_DC1A
        50:05:07:68:0c:11:[redacted]  Imported from FID 20
        50:05:07:68:0c:11:[redacted]  Imported from FID 20
        50:05:07:68:0c:11:[redacted]  Imported from FID 20
        50:05:07:68:0c:21:[redacted]  Imported from FID 20
        50:05:07:68:0c:21:[redacted]  Imported from FID 20
        50:05:07:68:0c:11:[redacted]  Imported from FID 20
        50:05:07:68:0c:21:[redacted]  Imported from FID 20
        50:05:07:68:0c:21:[redacted]  Imported from FID 20
        50:05:07:68:0c:11:[redacted]  Imported from FID 20
        50:05:07:68:0c:21:[redacted]  Imported from FID 20
        50:05:07:68:0c:11:[redacted]  Imported from FID 20
        50:05:07:68:0c:21:[redacted]  Imported from FID 20
        50:05:07:68:0c:11:[redacted]  Imported from FID 20
        50:05:07:68:0c:21:[redacted]  Imported from FID 20
        50:05:07:68:0c:11:[redacted]  Imported from FID 20
        50:05:07:68:0c:21:[redacted]  Imported from FID 20
        50:05:07:68:0b:31:[redacted]  EXIST in FID 2
        50:05:07:68:0b:33:[redacted]  EXIST in FID 2
        50:05:07:68:0b:31:[redacted]  EXIST in FID 2
        50:05:07:68:0b:33:[redacted]  EXIST in FID 2
```

References

- IBM System Storage b-type Multiprotocol Routing
- Fabric OS Extension Administrator's Guide - Fabric OS v7.4.0
- Product Guide IBM System Storage SAN42B-R Extension Switch