

Dell EMC VPLEX: Networking

Implementation and planning best practices

Abstract

This document is one of a set of technical notes containing current Dell EMC™ VPLEX™ implementation and best practices information. VPLEX and metro node networking best practices are discussed in this document.

January 2021

Revisions

Date	Description
January 2021	Initial release

Acknowledgments

Author: Michael Cram

Michael.cram2@dell.com

The information in this publication is provided "as is." Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

This document may contain certain words that are not consistent with Dell's current language guidelines. Dell plans to update the document over subsequent future releases to revise these words accordingly.

This document may contain language from third party content that is not under Dell's control and is not consistent with Dell's current guidelines for Dell's own content. When such third party content is updated by the relevant third parties, this document will be revised accordingly.

Copyright © January 2021 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners. [1/19/2021] [Best Practices] [H13552]

Table of contents

Revisions.....	2
Acknowledgments.....	2
Table of contents	3
Executive summary.....	4
1 Internal networks	5
2 Management network and internal management VPN	6
3 VPLEX Metro WAN cluster connectivity.....	7
3.1 Fibre Channel WAN cluster connectivity	7
3.1.1 Latency/roundtrip conditions.....	7
3.1.2 Buffer-to-buffer credits.....	8
3.1.3 Dual fabrics for inter-cluster communication	8
3.1.4 Zoning inter-cluster connectivity	9
3.2 WAN cluster connectivity for Metro over IP (10Gb/E)	11
3.2.1 Dual networks for inter-cluster communication	12
3.2.2 Performance tuning VPLEX Metro IP WAN	12
4 Rule sets.....	15
5 Metro node	16
5.1 Networking: Intra-cluster and Inter-cluster communications	16
5.2 Intra-cluster.....	16
5.3 Inter-cluster.....	17
5.3.1 IP WAN connectivity	17
A Technical support and resources	18

Executive summary

There are four general areas of connectivity for VPLEX:

1. Host to VPLEX connectivity.
2. VPLEX to array connectivity
3. WAN connectivity for either Fibre Channel or Ethernet protocols.
4. LAN connectivity.

This document will address the LAN and WAN connectivity requirements. See the Dell EMC™ VPLEX™ SAN connectivity document for host and array connectivity.

Metro Node LAN and WAN are also covered in this document. Metro node is a feature add-on for both Dell EMC PowerStore™ and Dell EMC Unity™ XT arrays that is based on VPLEX active/active architecture. It is built on the Dell R640 hardware and running the 7.0 code version.

These technical notes are for Dell EMC field personnel and partners and customers who will be configuring, installing, and supporting VPLEX. An understanding of these technical notes requires an understanding of the following:

- SAN technology and network design
- Fiber Channel block storage concepts
- VPLEX concepts and components

Note: This network best practices document does not cover host cluster network considerations as VPLEX does not provide solutions directly for host stretch cluster networks. Host stretch clusters may require layer 2 network connectivity so please refer to host requirements and network vendor products that provide the ability to separate cluster nodes across datacenters. Even though this document does not provide for host cluster network requirements they are still an integral part of the overall solution.

1 Internal networks

Within the each VPLEX cluster are two sets of private networks. One set is a redundant internal management IP network between cluster components and the cluster management server. The second set is a redundant internal network which provides director to director communications solely within the VPLEX cluster (intra-cluster). This internal Fibre Channel network (VS2) or InfiniBand network (VS6) is commonly referred to as LOCAL COM. A single-engine cluster simply connects corresponding Local COM ports together with FC/IB cables direct connected. For dual and quad engine VPLEX clusters the LOCAL COM includes two internal use FC or IB switches. **At no time should any of these internal networks be connected to customer networks.**

2 Management network and internal management VPN

Each VS2 VPLEX cluster has a single management port located on the management server.

The VS6 has two management ports, one on MMCS-A and the other on MMCS-B. Both of these ports must be configured and connected to the customer network. MMCS-A is the management port that will be accessed for all management and monitoring purposes.

MMCS-B must be configured so that processes such as ndu or configuration changes like setting cluster ID during installation can be persisted to MMCS-B or those processes will fail.

This port provides the connectivity to facilitate the management and configuration of the cluster. For VPLEX Metro, the management port is also utilized for a VPN which extends the internal management networks between all of the cluster components in both clusters and, if installed, the VPLEX Witness.

Refer the Security Guide for port information and the protocols that need to be allowed on the firewall (both in the outbound and inbound filters).

The VPLEX Security Configuration Guide is located here:

<https://solveonline.emc.com/solve/home/42>

3 VPLEX Metro WAN cluster connectivity

VPLEX Metro connectivity is defined as the communication between clusters in a VPLEX Metro. The two key components of VPLEX Metro communication are FC (FCIP, DWDM) or 10 GigE and VPN between management servers. Please refer to the VPLEX Internal Management VPN section for details. FC WAN is the Fibre Channel connectivity and 10 GigE is the Ethernet connectivity options between directors of each cluster. Please choose the section that applies to your configuration.

3.1 Fibre Channel WAN cluster connectivity

- Latency may be up to 10ms rtt, depending on the use case. Many systems require 5ms rtt or less. Please consult the latest VPLEX Simple Support Matrix found on: <https://solveonline.emc.com/solve/home/42>
- Each director's FC WAN ports must be able to see at least one FC WAN port on every other remote director (required).
- Independent FC WAN links are strongly recommended for redundancy
- Each director has two FC WAN ports that should be configured on separate fabrics (or FC/IP external hardware) to maximize redundancy and fault tolerance so that a single external VPLEX failure does not cause a full WAN communications failure to a single director
- Configure FC WAN links between clusters like ISLs between FC switches. This does not require a merged fabric between locations.
- Logically isolate VPLEX Metro traffic from other traffic using zoning, VSANs or LSAN zones. Additionally, physical isolation from the SAN can be achieved by connecting to separate switches used to connect the data centers without requiring connection to the local SAN.
- The director's local com port is used for communications between directors within the cluster.

Note: Cisco IVR is not supported on the VPLEX WAN COM fabric.

FC WAN connectivity utilizes Fibre Channel with standard synchronous distance limitations. Considerations for Fibre Channel include latency/round trip conditions and buffer-to-buffer credits as well as the BB_credits applied to distance. An excellent source for additional information is the Dell-EMC Symmetrix Remote Data Facility (SRDF) Connectivity Guide or Dell-EMC Networked Storage Topology Guide, available through E-Lab™ Interoperability Navigator at:

<http://elabnavigator.EMC.com>

Note: VPLEX short-write was introduced in VPLEX 5.2 SP1. The feature is not compatible with FCIP Fast Write (Brocade) and FCIP Write Acceleration (Cisco). This feature must be disabled on the FCIP switches to ensure that VPLEX works properly. Please refer to the latest VPLEX Simple Support Matrix – Distance Extension section for details.

3.1.1 Latency/roundtrip conditions

Latency is generally referred to in milliseconds (ms) as the combined roundtrip time (RTT) between local and remote clusters. A FC-Frame by itself takes approximately 1 ms to traverse a one-way distance of 100 km from primary-transmitter to secondary-receiver.

For example, if two locations are 100 km apart, since standard Fibre Channel protocol requires two round trips for a write I/O, then 4 ms of latency (2 x RTT) will be added to the write operation. VPLEX short-write, introduced in VPLEX 5.2 SP1, will reduce number of required round trips for a write I/O to the equivalent of 1.5 round trips. However, as more network components are attached to the configuration for pure Fibre Channel environments, latency will naturally increase. This latency can be caused by network components such as host initiators, switches, fiber optics, and distance extension devices, as well as factors such as cable purity. The VPLEX application layer will contribute additional delays on top of the network.

The supported network round trip latency is up to 10ms for VPLEX Metro. However, the actual supported round trip latency depends not only on VPLEX but on the server OS, clustering software, and host level applications. In addition, host virtualization, if applicable, may have requirements as well. Please see the VPLEX Simple Support Matrix, as well as release notes and documentation for all host components (OS, application, virtualization, and clustering).

Using performance monitoring tools, the roundtrip time can be determined for troubleshooting any WAN latency issues within the network.

3.1.2 Buffer-to-buffer credits

Fibre Channel uses the BB_Credits (buffer-to-buffer credits) mechanism for hardware-based flow control. This means that a port has the ability to pace the frame flow into its processing buffers. This mechanism eliminates the need of switching hardware to discard frames due to high congestion. On VPLEX, during fabric login, the VS2 8 Gbps FC ports advertise a BB credit value of up to 41. The switch will respond to the login with its maximum support BB credit value. The BB credit value returned by the switch is used by the FC port, up to a maximum of 41. Dell-EMC testing has shown this mechanism to be extremely effective in its speed and robustness. The BB credits between the two clusters must be configured for long distance.

Refer to the Dell-EMC Networked Storage Topology Guide for proper calculations and settings for your WAN connectivity.

3.1.3 Dual fabrics for inter-cluster communication

A VPLEX Metro should be set up with redundant and completely independent Fibre Channel connectivity between clusters. This provides maximum performance, fault isolation, fault tolerance, and availability.

Redundant fabrics are of critical importance due to the fact that when the directors in one cluster have inconsistent connectivity with the directors in the remote cluster, the two clusters will be logically split until the connectivity issues are resolved. This is by design. The firmware requires full connectivity among all directors for protocols such as cache coherence and inter-cluster communication. Every local director must be able to communicate with every remote director and vice versa. Without full connectivity, the director will continue to run but will bring the inter-cluster link down. The net result is that all volumes at the non preferred (or losing) site will become unavailable as per the pre-defined per volume cluster detach rules. Recovery is simple, but manual. It requires that connectivity be restored between all directors prior to the resumption of I/O operations. This scenario is much less likely if two independent and redundant connections are utilized between clusters.

The following is an example fabric configuration:

Site-1 has switches, switch-1A and switch-1B.

Site-2 has switches, switch-2A and switch-2B.

At Site-1:

- a) Connect "A2-FC00" and "B2-FC00" to switch-1A.
- b) Connect "A2-FC01" and "B2-FC01" to switch-1B

At Site-2:

- a) Connect "A2-FC00" and "B2-FC00" to switch-2A
- b) Connect "A2-FC01" and "B2-FC01" to switch-2B

Place ISLs between switch-1A and switch-2A and between switch-1B and switch-2B. The best practice is to have your ISL traffic travel through independent links (and/or carrier) between your sites.

3.1.4 Zoning inter-cluster connectivity

The Fibre Channel WAN COM ports are configured into two different port groups. Ports on the WAN COM module labeled FC00 will be in port group 0 and the ports labeled FC01 will be in port group 1. The inter-cluster connectivity will be spread across two separate fabrics for HA and resiliency. Port group 0 will be connected to one fabric and port group 1 will be connected to the other fabric.

A Fibre Channel WAN COM port in one cluster will be zoned to all of the FC WAN ports on the same port group at the remote site. This is roughly equivalent to one initiator zoned to multiple targets. This is repeated for each and every Fibre Channel WAN COM ports at both clusters.

This zoning provides additional fault tolerance and error isolation in the event of configuration error or a rogue fabric device (when compared to a single large zone).




Figure 1 Fibre Channel WAN COM connections on VS2 VPlex hardware

Though this requires more setup than a single zone, it is worth the effort and should not be considered out of the norm for a SAN administrator.

3.1.4.1 Zoning using director FC WAN COM modules A2 and B2

Assuming two fabrics and dual-engine systems for Cluster A and Cluster B, each fabric would be zoned as follows:

Key:

Director-1-1-A = cluster 1 engine director A1 : D11A

Director-2-1-B = cluster 2 engine 1 director B: D21B

WAN COM ports are FC00 or FC01

All combined (example): D11A:FC00

Zone configuration:

Zone 1: D11A:FC00-> D21A:FC00, D21B:FC00, D22A:FC00, D22B:FC00

Zone 2: D11A:FC01-> D21A:FC01, D21B:FC01, D22A:FC01, D22B:FC01

Zone 3: D12B:FC00 ->D21A:FC00, D21B:FC00, D22A:FC00, D22B:FC00

Zone 4: D12B:FC01 ->D21A:FC01, D21B:FC01, D22A:FC01, D22B:FC01

Zone 5: D21A:FC00 ->D11A:FC00, D11B:FC00, D12A:FC00, D12B:FC00

Zone 6: D21A:FC01 ->D11A:FC01, D11B:FC01, D12A:FC01, D12B:FC01

Zone 7: D21B:FC00 ->D11A:FC00, D11B:FC00, D12A:FC00, D12B:FC00

Zone 8: D21B:FC01 ->D11A:FC01, D11B:FC01, D12A:FC01, D12B:FC01

There would be 16 zones for the quad engine configuration. Please extrapolate from this example.

Note: If VPLEX is deployed with IP inter-cluster network (fcip), the inter-cluster network must not be able to route to the following reserved VPLEX subnets: 128.221.252.0/24, 128.221.253.0/24, and 128.221.254.0/24.

3.1.4.2 Checking cluster connectivity

To check for FC WAN connectivity, log in to the VPLEX CLI and run the following command:

```
connectivity director --director <director_name>
```

An example illustrating connectivity to director A WAN-COM module A2:

```
VPlexcli:/> connectivity director --director director-1-1-A
```

Directors discovered by director-1-1-A, UUID 0x000000003b201e0b:

Directors discovered by east_35, UUID 0x000000003b201e0b:

Director UUID	Protocol	Address	Ports
0x000000003b301fac	COMSCSI	0x50001442901fac21	A2-FC01
	COMSCSI	0x50001442901fac20	A2-FC00
0x000000003b201fac	COMSCSI	0x50001442801fac21	A2-FC01
	COMSCSI	0x50001442801fac20	A2-FC00
0x000000003b301f80	COMSCSI	0x50001442901f8031	A3-FC01
	COMSCSI	0x50001442901f8030	A3-FC00
0x000000003b201f80	COMSCSI	0x50001442801f8030	A3-FC00
	COMSCSI	0x50001442801f8031	A3-FC01
0x000000003b301e07	COMSCSI	0x50001442901e0721	A2-FC01
	COMSCSI	0x50001442901e0720	A2-FC00
0x000000003b201e07	COMSCSI	0x50001442801e0721	A2-FC01
	COMSCSI	0x50001442801e0720	A2-FC00
0x000000003b301e0b	COMSCSI	0x50001442901e0b31	A3-FC01
	COMSCSI	0x50001442901e0b30	A3-FC00

Check to make sure that the director has connectivity to the remote directors using the ports “A2-FC00” and “A2-FC01” shown in the “Ports” column. Connections to I/O Module 2 (A2 or B2) are WAN-COM connections (shown in **Bold**). The display will also show LOCAL-COM connections to ports in I/O Module 3 (A3 or B3).

Repeat this process for all the remaining directors in your system and check to make sure that they can reach the remote directors using both the FC WAN ports.

Validate using command:

```
vplexcli:>connectivity validate-wan-com
```

3.2 WAN cluster connectivity for Metro over IP (10Gb/E)

VPLEX Metro IP connectivity is defined as the communication between clusters in a VPLEX Metro over Ethernet. The two key components of VPLEX Metro communication are Ethernet (IP) WAN and VPN. (Please see the VPLEX Internal VPN section for details.) Ethernet WAN is the network connectivity between directors of each cluster and the VPN is connectivity between management servers for management purposes.




Figure 2 IP WAN COM connections on VS2 VPLEX hardware

- Each director's IP WAN ports must be able to see at least one WAN port on every other remote director (required)
- Independent WAN links are strongly recommended for redundancy
- Each director has two WAN ports that should be configured on separate hardware to maximize redundancy and fault tolerance
- Configure WAN links between clusters on network components that offer the same Quality of Service (QoS)
- VPLEX uses best available path load balancing and is capable of adapting to changing WAN latency
- Logically isolate VPLEX Director WAN traffic from other WAN traffic
- The supported network round trip latency is less than or equal to 10ms

3.2.1 Dual networks for inter-cluster communication

A VPLEX Metro should be set up with redundant and completely independent networks between clusters located over geographically different paths. This provides maximum performance, fault isolation, fault tolerance, and availability.

Redundant networks are of critical importance due to the fact that when the directors in one cluster have inconsistent connectivity with the directors in the remote cluster, the two clusters will be logically split until the connectivity issues are resolved. This is by design. The firmware requires full connectivity among all directors for protocols such as cache coherence and inter-cluster communication. Without full connectivity, the director will continue to run but will bring the inter-cluster link down. The net result is that all volumes at the losing site will become unavailable as per the pre-defined per volume or per consistency group cluster detach rules or through the control based on VPLEX Witness.

Maximum Transmission Unit (MTU) size is also a configurable attribute. Performance measurements for GeoSynchrony 6.0.1 and newer, do not show significant improvements with jumbo frames. It is recommended that the default MTU of 1500 be kept.

3.2.2 Performance tuning VPLEX Metro IP WAN

3.2.2.1 Setting the “max-transfer-size” for rebuilds:

The “max-transfer-size” parameter is used to determine the size of the region on a virtual volume that is locked while data is read for a distributed (or local) raid-1 rebuild as well as Data Mobility jobs when migrating devices and/or extents. This is set on a per device basis. Metro node uses the rebuild function when you first

create a distributed raid-1 or local raid-1 device to synchronize data on both mirror legs. During normal production operations a rebuild would be fairly *uncommon*. This setting can dynamically be adjusted at any time – even during a migration or rebuild. The 128K size is very conservative in terms of FE host impact and the appropriate size is a function of when you start to see impact from moving it higher and your overall time objectives during a rebuild or migration. In effect, max transfer is a tunable throttle that lets the user control the speed *and* host impact of distributed/local raid-1 rebuilds and migrations.

Under normal host read/write activities the max transfer size setting will not affect your replication speed across the wire. Metro node will take all the bandwidth that is allocated to it provided it has sufficient write traffic. A suggested next step would be to start some host IO against DR1 and test various levels of throughput per your normal IOPS testing procedures.

```
VPlexcli:/> rebuild show-transfer-size --devices *
device name                               transfer-size
-----
device_6006016018641e000e933e974fbae111_1 128K
device_6006016018641e000f933e974fbae111_1 128K
device_6006016018641e0010933e974fbae111_1 128K
device_6006016018641e0011933e974fbae111_1 128K
device_6006016018641e0012933e974fbae111_1 128K
```

```
VPlexcli:/> rebuild set-transfer-size --devices * --limit 1M
```

```
VPlexcli:/> rebuild show-transfer-size --devices *
device name                               transfer-size
-----
device_6006016018641e000e933e974fbae111_1 1M
device_6006016018641e000f933e974fbae111_1 1M
device_6006016018641e0010933e974fbae111_1 1M
device_6006016018641e0011933e974fbae111_1 1M
device_6006016018641e0012933e974fbae111_1 1M
```

```
VPlexcli:/> rebuild status
```

device	rebuild	rebuilder	rebuilt /total	percent finished	throughput	ETA
device_641e000e93_1	full	s2_0908_spa	5.59G/10.1G	55.47%	31.8M/s	2.4min
device_641e000f93_1	full	s2_0908_spa	5.59G/10.1G	55.47%	31.8M/s	2.4min
device_641e001093_1	full	s1_0a69_spb	3.87G/10.1G	38.40%	42.4M/s	2.48min
device_641e001193_1	full	s2_0908_spb	620M/10.1G	6.01%	16.8M/s	9.58min
device_641e001293_1	full	s1_0a69_spa	7.99G/10.1G	79.32%	34.5M/s	61sec

3.2.2.2 Set the Socket-Buf-Size:

The formula for setting the “socket-buf-size” is:

Bandwidth*RTT Latency/2 for each of two connections

Example:

100MB * .1sec/2=5MB socket buffer setting each (default is 1MB)

Syntax:

```
VPlexcli: /clusters/cluster-1/cluster-connectivity/option-sets/optionset-com-0>  
set socket-buf-size 5M  
VPlexcli: /clusters/cluster-1/cluster-connectivity/option-sets/optionset-com-1>  
set socket-buf-size 5M
```

```
VPlexcli: /clusters/Cluster-2/cluster-connectivity/option-sets/optionset-com-0>  
set socket-buf-size 5M  
VPlexcli: /clusters/Cluster-2/cluster-connectivity/option-sets/optionset-com-1>  
set socket-buf-size 5M
```

Make sure that you set the Maximum Transmission Unit (MTU) and the Socket Buffers according to your anticipated workloads. The following figures offer some suggested values to start your base lining process:

4 Rule sets

It is recommended that Consistency Groups be used for setting and maintaining detach rules. Cluster witness only works with Consistency Groups and does not arbitrate detach rules on volumes that are not in Consistency Groups.

As a minimum, set the detach timer to 5 seconds. Setting the detach delay lower than 5 seconds can result in unnecessary or numerous storage detaches during periods of network instability. Multiple detaches in a short period of time can also result in many unnecessary data rebuilds and subsequently in reduced performance.

Configure detach rules based on the cluster/site that you expect to continue I/O during any network outage.

Avoid conflicting detach situations. Each distributed device (or consistency group) must have a rule set assigned to it. When a cluster's distributed device detaches during a link outage or other communications issue with the other members of a distributed device, the detached device can resume I/O. Therefore, it is important to understand the nature of the outage and which cluster is set to automatically detach. It is a recommendation that the rule set configuration for each distributed device or group of devices be documented as well as plans for how to handle various outage types.

It is important to note that rule sets are applied on a distributed device basis or to a number of devices within a consistency group. It is within normal parameters for different distributed devices to resume I/O on different clusters during an outage. However, if a host application uses more than one distributed device, most likely all of the distributed devices for that application should have the same rule set to resume I/O on the same cluster. Hosts with multiple devices are best managed using Consistency Groups.

5 Metro node

5.1 Networking: Intra-cluster and Inter-cluster communications

Intra-cluster communications are the cluster communications within the cluster at one site whereas inter-cluster communications are the communication layer across clusters connected together in a Metro configuration.

Note: The network best practices section does not cover host cluster network considerations as metro node does not provide solutions directly for host stretch cluster networks. Host stretch clusters may require layer 2 network connectivity so please refer to host requirements and network vendor products that provide the ability to separate cluster nodes across datacenters. Even though this document does not provide for host cluster network requirements they are still an integral part of the overall solution.

5.2 Intra-cluster

Within each metro node cluster are two sets of private networks. One set is a redundant internal management IP network between cluster nodes color coded in green and purple. These Ethernet cables are direct connected to the corresponding ports between nodes. The second set is a redundant internal Ethernet based Twin-AX network which provides director to director communications solely within the metro node cluster (intra-cluster). This internal redundant Twin-AX network is commonly referred to as LOCAL COM. A Local COM network simply connects corresponding Local COM ports together with Twin-AX cables direct connected. **At no time should any of these internal networks be connected to customer networks.**

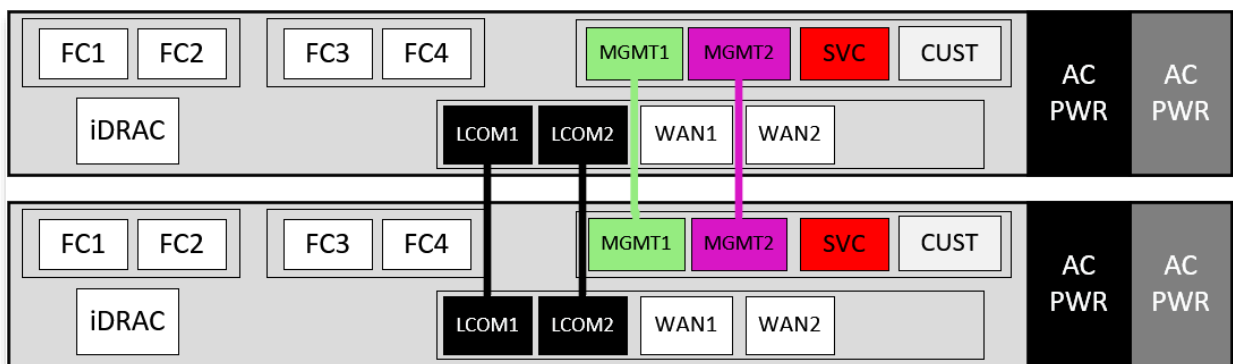


Figure 3 LCOM , MGMT1 & MGMT2 ports

Each node has an iDRAC port as well as service ports and customer network ports. The iDRAC port will be disabled from the factory. The SVC (service) ports will be used for connecting to the service laptop on IP address 128.221.252.2 using the service login. Follow the installation and configuration guide when performing initial install as it will direct you to which port to connect to during the procedure. After completing the installation, day to day operation may be performed from either node using the CUST (customer) ports.

The CUST ports will be connected to the customer network using customer supplied addresses. This port is used for day-to-day operations as well as file transfer and VPN secure connections between all nodes across a Metro.

5.3 Inter-cluster

Inter-cluster is the WAN COM communications layer between metro node clusters in a Metro configuration. There are two WAN ports on each of the four nodes and will be configured into two port groups. All WAN1 ports will be in one port group and all WAN2 ports will be in a different port group. These ports are 10GbE optical and are required to be connected via switched networks. Direct connect is not supported as each port on one node must communicate with both ports in the same port group on the remote cluster.

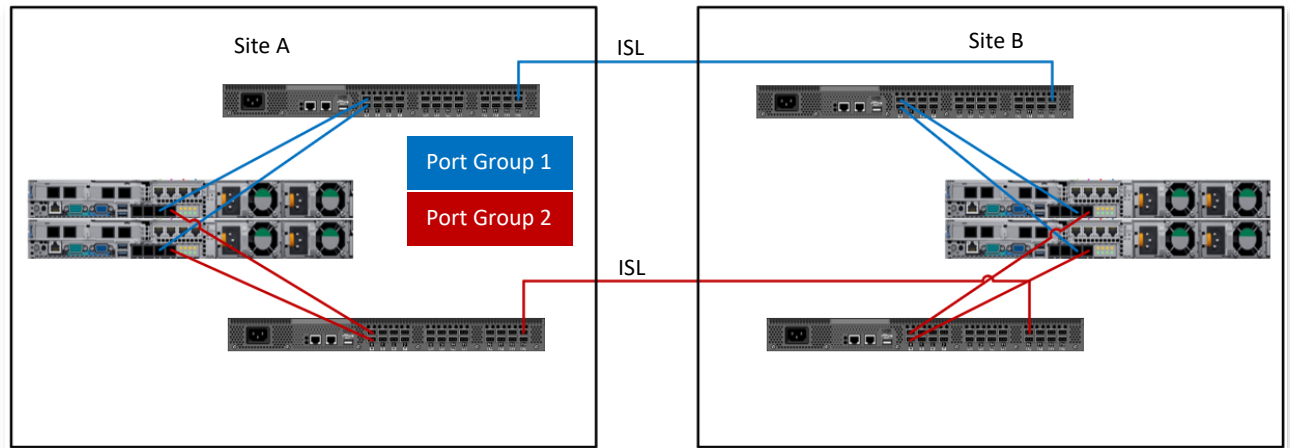


Figure 4 WAN network connectivity

Metro node will load balance between the two port groups using the InvACO algorithm which is essentially a best available path option. The best practice to take full advantage of this load balancing algorithm is to have “air gap” separation between the port groups. Any network configuration that has a shared component may have an impact on both port groups when there is an outage on that component or simply congestion that may be experienced by both port groups. Likewise, any component that is physically located next to its corresponding component on the other port group may experience outages due to environmental causes on both components such as the proverbial backhoe taking out both links at the same time.

5.3.1 IP WAN connectivity

- Each director’s IP WAN ports must be able to see at least one WAN port on every other remote director (required)
- The director’s local com port is used for communications between directors within the cluster
- Independent WAN links are strongly recommended for redundancy
- Each director has two WAN ports that should be configured on separate hardware to maximize redundancy and fault tolerance
- Configure WAN links between clusters on network components that offer the same Quality of Service (QoS)
- Metro node uses best available path load balancing and is capable of adapting to changing WAN latency
- Logically isolate metro node director WAN traffic from other WAN traffic
- The supported network round trip latency is less than or equal to 10 ms
- Each link must be sized to handle 100% of the WAN traffic to prevent an impact to the applications

Please see VPLEX Performance Tuning section for additional information.

A Technical support and resources

[Dell.com/support](https://dell.com/support) is focused on meeting customer needs with proven services and support.

[Storage and data protection technical documents and videos](#) provide expertise to ensure customer success with Dell EMC storage and data protection products.