



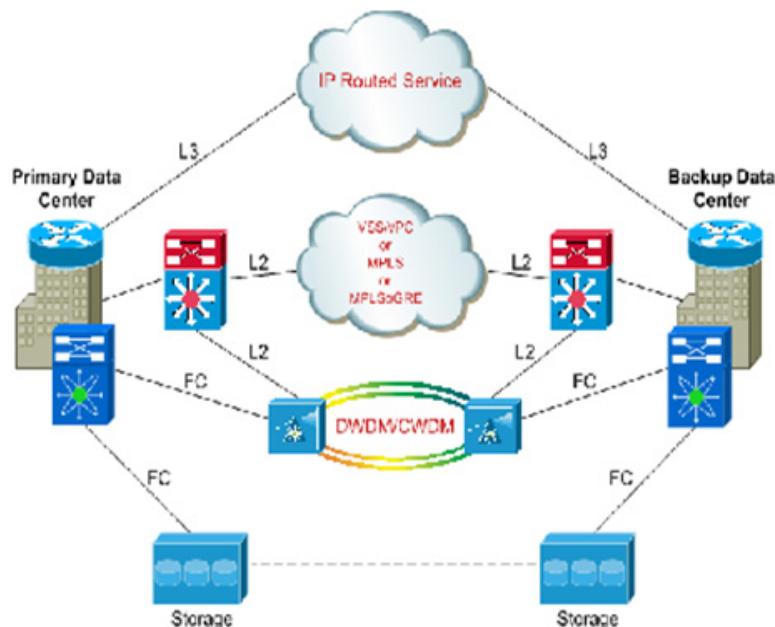
# CHAPTER 2

## Data Center Interconnect Solution Overview

The term DCI (Data Center Interconnect) is relevant in all scenarios where different levels of connectivity are required between two or more data center locations in order to provide flexibility for deploying applications and resiliency schemes.

Figure 2-1 summarizes the three general types of connectivity required for a DCI solution.

**Figure 2-1 DCI Connectivity Overview**



- **LAN Extension:** Provides a single Layer 2 domain across data centers. The data center applications are often legacy or use embedded IP addressing that drives Layer 2 expansion across data centers. Layer 2 Extension provides a transparent mechanism to distribute the physical resources required by some application frameworks such as the mobility of the active machine (virtual or physical).
- **Layer 3 Extension:** Provides routed connectivity between data centers used for segmentation/virtualization and file server backup applications. This may be Layer 3 VPN-based connectivity, and may require bandwidth and QoS considerations.

- **SAN Extension:** This presents different types of challenges and considerations because of the requirements in terms of distance and latency and the fact that Fibre Channel cannot natively be transported over an IP network.

In addition to the 3 functional component listed above, a holistic DCI solution usually leverages an additional building block. This is usually referred to as Path Optimization and deals with the fact that every time a specific VLAN (subnet) is stretched between two (or more) locations that are geographically remote, specific considerations need to be made regarding the routing path between client devices that need to access application servers located on that subnet. Same challenges and considerations also apply to server-to-server communication, especially for multi-tier application deployments. Path Optimization includes various technologies that allow optimizing the communication path in these different scenarios. Integration of network services (as FW, load-balancers) also represents an important design aspect of a DCI solution, given the challenges brought up by the usual requirement of maintaining stateful services access while moving workloads between data center sites.

LAN Extension represents a very important component of DCI and is the main focus of this document. The following sections of this chapter present the most common business requirements for LAN Extension, listing also its main technical requirements. The last section provides an overview of Cisco LAN Extension solution offering.

## LAN Extension Business Drivers

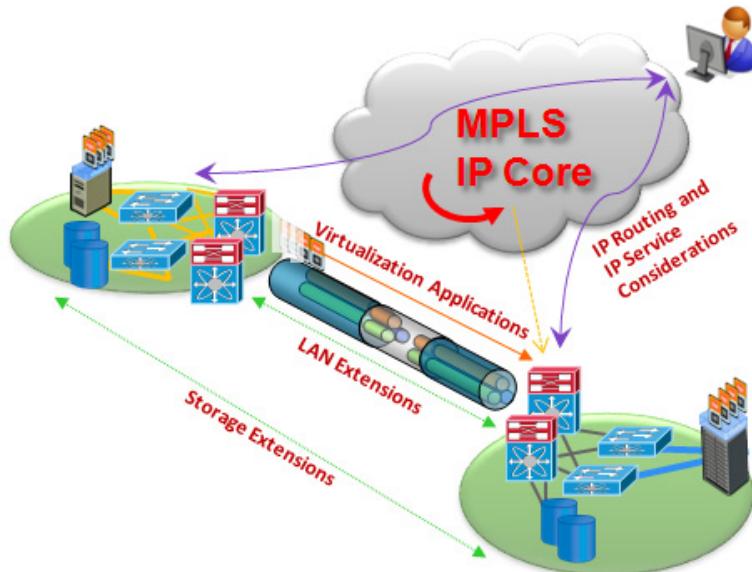
There are various business reasons driving the deployment of DCI solutions. Traditionally, Cisco recommends isolating and reducing Layer 2 networks to their smallest scope, usually limiting them to the access layer.

However, in some situations Layer 2 must be extended beyond the single data center, specifically when the framework or scenario developed for a campus has been extended beyond its original geographic area and over multiple data centers across long distances. Such scenarios are becoming more prevalent as high-speed service provider connectivity becomes more available and cost effective.

High-availability clusters, server migration, and application mobility are some important use cases that require Layer 2 extension.

### Workload Mobility (Active/Active Data Centers)

The deployment of LAN Extension technologies can also facilitate and maximize a company's server virtualization strategy, adding flexibility in terms of where compute resources (workload) reside physically and being able to shift them around geographically as needs dictate.

**Figure 2-2 Workload Mobility across DC Sites**

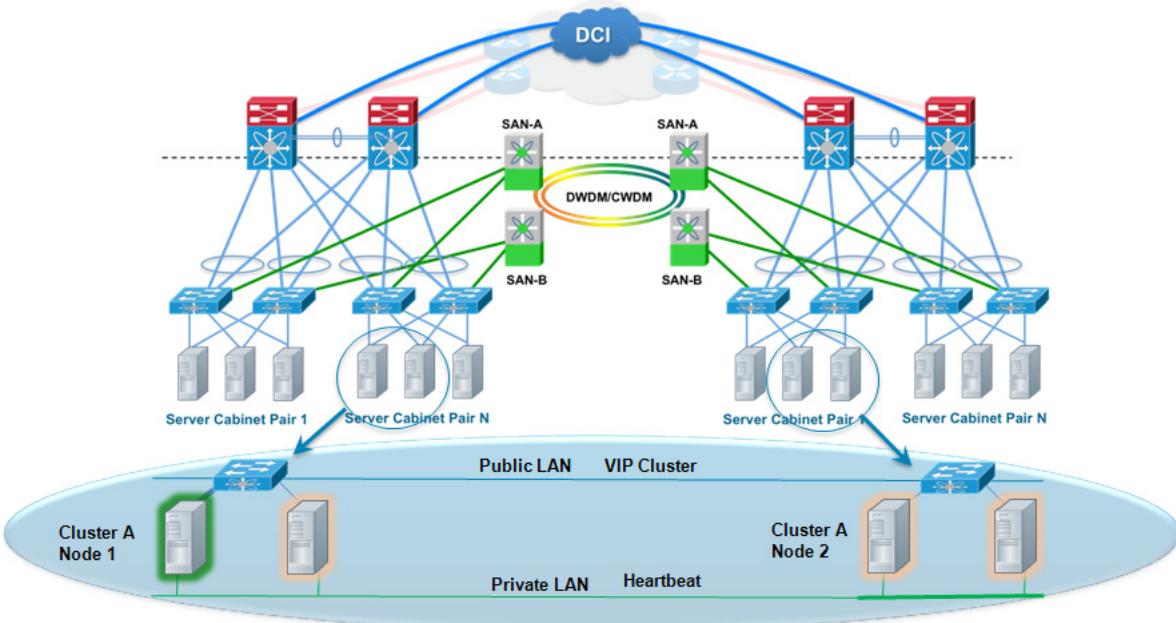
Some applications that offer virtualization of operating systems allow the move of virtual machines between physical servers separated by long distances. To synchronize the software modules of the virtual machines during a software move and to keep the active sessions up and running, the same extended VLANs between the physical servers must be maintained.

Migrating virtual machines between data centers provides compute power from data centers closer to the clients (“follow the sun”) or to load-balance across multiple sites. Enterprises with multiple sites can also conserve power and reduce cooling costs by dynamically consolidating virtual machines in fewer data centers.

#### **Business Continuance: High-Availability Clusters**

Despite the fact that application clustering is evolving and has started supporting deployments across L3 boundaries, there are still a long list of applications that require L2 adjacency between the cluster nodes. These applications include:

- Private inter-process communication (such as heartbeat and database replication) used to maintain and control the state of the active node.
- Public communication (from the client to the virtual IP of the cluster).

**Figure 2-3 Multi-Site HA Cluster**

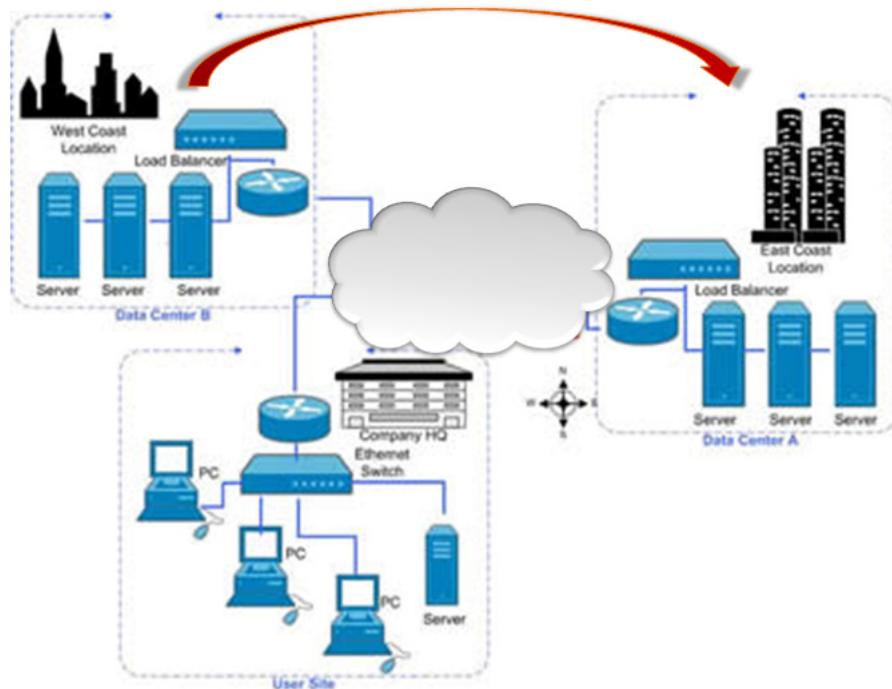
Despite the fact that application clustering is evolving and has started supporting deployments across L3 boundaries, there are still a long list of applications that require L2 adjacency between the cluster nodes. These applications include:

- Microsoft MSCS
- Veritas Cluster Server (Local)
- Solaris Sun Cluster Enterprise
- VMware Cluster (Local)
- Oracle Real Application Cluster (RAC)
- IBM HACMP
- EMS/Legato Automated Availability Manager
- NetApp Metro Cluster
- HP Metrocluster

#### **Data Center Migration and Consolidation**

Migration of servers (physical or virtual) between data centers is becoming more popular. This is often driven by different and sometimes opposite requirements, like the consolidation of a large number of DC sites into fewer ones, or the expansion of services from a few sites to multiple locations.

**Figure 2-4 Data Center Migration**



Providing a L2 path between sites is often desirable in these cases to ensure a smooth migration and minimize the experienced application down time. It is also important to keep in mind additional considerations:

- IP renumbering of servers to be moved is complex and costly. Avoiding IP address renumbering makes physical migration projects easier and reduces cost substantially.
- Some applications may be difficult to readdress at Layer 3 (mainframe applications, for example). In this case, it is easier to extend the Layer 2 VLAN outside the access layer, to keep the original configuration of the systems after the move.
- During phased migration, when only part of the server farm is moving at any given time, Layer 2 adjacency is often required across the whole server farm for business-continuity purposes.

## LAN Extension Considerations

As mentioned above, LAN extension solutions are commonly used to extend subnets beyond the traditional Layer 3 boundaries of a single data center. Stretching the network space across two or more data centers can accomplish many things. Doing so also presents a challenge, since providing these LAN extension capabilities may have an impact on the overall network design. Simply allowing Layer 2 connectivity between sites that were originally connected only at Layer 3 would have the consequence of creating new traffic patterns between the sites: STP BPDUs, unicast floods, broadcasts, ARP requests, and so on. This can create issues, some of them related to attacks (ARP or flood storms), others related to stability issues (size of STP domain) or scale (ARP caches or MAC address table sizes). How does an extended spanning-tree environment avoid loops and broadcast storms? How does a core router know where an active IP address or subnet exists at any given time?

## LAN Extension Technical Requirements

For deploying a solid LAN extension solution, it is important to keep into considerations two main following requirements:

- Spanning-Tree (STP) Isolation: the first basic requirement is to isolate the Spanning Tree domains between the data center sites belonging to the extended Layer 2 network. This is important to protect against any type of global disruptions that could be generated by a remote failure, and to mitigate the risk of propagating unwanted behavior such as topology change or root bridge movement from one data center to another. These packets could be flooded throughout the Layer 2 network, making all remote data centers and resources unstable, or even inaccessible.
- End-to-End loop prevention: In each data center site, the deployment of redundant physical devices providing LAN extension services is recommended to improve the overall resiliency of the LAN Extension solution. Therefore, a solution must eliminate any risk of creating an end-to-end Layer 2 loop; STP cannot be used for this purpose, given the previous requirement of isolating the STP domains between remote DC sites.

In addition to these, other requirements to be considered are:

- **WAN Load Balancing:** Typically, WAN links are expensive, so the uplinks need to be fully utilized, with traffic load-balanced across all available uplinks.
- **Core Transparency:** The LAN extension solution should ideally be transparent to the existing enterprise core, to minimize the operational impact.
- **Data Center Site Transparency:** The LAN extension solution should not affect the existing data center network deployment.
- **VLAN Scalability:** The solution must be able to scale to extend up to hundreds (sometimes few thousands) of VLANs.
- **Multisite Scalability:** The LAN extension solution should be able to scale to connect multiple data centers.
- **Hierarchical Quality of Service (HQoS):** HQoS is typically needed at the WAN edge to shape traffic when an enterprise subscribes to a substrate service provider service or a multipoint Ethernet virtual private line (EVPL) service.
- **Encryption:** The requirement for LAN extension cryptography is increasingly prevalent, to meet federal and regulatory requirements.

## Cisco LAN Extension Solutions

Cisco LAN Extension solutions can be divided in three categories:

### Ethernet Based Solutions

This category includes technologies like virtual Port-Channel (vPC) or Virtual Switching Systems (VSS), originally deployed for intra Data Center designs, but readapted for DCI deployments. The idea is to extend VLANs between remote sites by leveraging Multi Chassis Etherchannels (MCECs) established between devices deployed in different sites. As such, this solution mostly applies to point-to-point deployments, where the sites are connected via dedicated dark fiber links or protected DWDM optical circuits.

**Note**

More information around Ethernet Based LAN Extension solutions (VSS/vPC) can be found in the following paper:

[http://www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns224/ns949/ns304/ns975/data\\_center\\_interconnect\\_design\\_guide.pdf](http://www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns224/ns949/ns304/ns975/data_center_interconnect_design_guide.pdf)

It is worth noticing that an emerging technology, Cisco FabricPath, originally positioned to deploy large L2 domains inside a data Center network, could also be considered for LAN Extension purposes.

### MPLS Based Solutions

This category includes MPLS technologies providing L2 connectivity services over a L3 network service. Depending on the nature of the transport infrastructure between data center sites and the number of data center sites to be interconnected, different technologies can address the connectivity requirements. EoMPLS and VPLS are usually positioned for point-to-point and multipoint deployments respectively over native MPLS based infrastructure. This is often the case with large enterprise or SP deployments. When only a generic IP service is available to interconnect different DC sites which is usually the case for small and medium enterprises acquiring connectivity services from one or more SPs, the same EoMPLS/VPLS technologies can be deployed over a logical overlay connection built leveraging GRE tunnels also known as EoMPLSoGRE or VPLSoGRE deployments.

**Note**

More information on MPLS Based LAN Extension solutions can be found in the following paper:

[http://www.cisco.com/en/US/docs/solutions/Enterprise/Data\\_Center/DCI/DCI2\\_External.pdf](http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/DCI/DCI2_External.pdf)

### IP Based Solutions

Part of this category is an emerging Cisco technology, called Overlay Transport Virtualization (OTV). OTV is an IP based functionality that has been designed from the ground up to provide Layer 2 extension capabilities over any transport infrastructure: Layer 2 based, Layer 3 based, IP switched, label switched, and so on. The only requirement from the transport infrastructure is providing IP connectivity between remote data center sites. In addition, OTV provides an overlay that enables Layer 2 connectivity between separate Layer 2 domains while keeping these domains independent and preserving the fault-isolation, resiliency, and load-balancing benefits of an IP-based interconnection.

As of this writing, OTV is supported only on the Nexus 7000. However, there are plans to extend OTV support to other Cisco platforms in the future.

**Note**

More information on OTV can be found in the following paper:

[http://www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns224/ns949/ns304/ns975/OTV\\_intro\\_wp.pdf](http://www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns224/ns949/ns304/ns975/OTV_intro_wp.pdf)

**■ LAN Extension Considerations**