# DATA CENTER INTERCONNECT MARKET TRENDS AND REQUIREMENTS

## THE IMPACT OF MARKET TRENDS ON DATA CENTER INTERCONNECT FOR LARGE ENTERPRISES

STRATEGIC WHITE PAPER

The enterprise data center interconnect (DCI) market has traditionally focused on near real-time disaster recovery and business continuity between sites and has typically been driven by two concerns — bandwidth and latency. The preferred solution has been to deploy WDM optical equipment to meet these bandwidth and latency requirements. While this approach continues to be a valid option, other solutions are becoming available as the data center is being completely redefined. The rapid gains in data center efficiency, the emergence of mega data centers and the impact of cloud computing, as well as the virtualization of data center architecture and Software-Defined Networking (SDN) mean that large enterprises need to consider new approaches to meet current and future DCI needs.

# TABLE OF CONTENTS

# INTRODUCTION

The data center interconnect (DCI) market has traditionally focused on near real-time disaster recovery and business continuity between sites and has typically been driven by two concerns — bandwidth and latency.

Large enterprises such as banks, healthcare companies, retail and manufacturing companies need to securely backup and replicate mission-critical data and applications between multiple locations. The technical requirements for DCI are driven by the application and its tolerance for latency while the choice of network used to connect data centers depends on whether synchronous or asynchronous connectivity is needed, and whether the backup sites are located near to or far from the primary site.

A typical large enterprise might be a global financial company with trading operations in city centers and back-up data centers located tens of kilometers away in the same metropolitan area. For such large enterprises with the most stringent performance, regulatory and security requirements, optical WDM is a widely used DCI solution. This approach to interconnect different data centers continues to be a valid option where the security and integrity of data remains central to the company's business and operations.

Other solutions are becoming available as DCI is being completely redefined. This comes as a result of rapid gains in data center efficiency, the emergence of mega data centers and the impact of cloud computing, software defined networking (SDN) and virtualization on data center architecture. Emerging technologies in the cloud era mean that large enterprises must consider new approaches to meet future DCI needs.

DCI is expanding beyond the business continuity and disaster recovery role it has traditionally played. It is also becoming a way to connect data centers in a cloud of virtualized compute and storage using SDN to automatically allocate resources as required. The key question for large enterprises used to be: "How can I maximize the amount of bandwidth between sites at the lowest latency and cost?" Given the rapid and dramatic market changes, the key question now is: "How can I meet my current DCI needs with a solution that encompasses virtualization, SDN and the cloud?"

DEFINING LARGE ENTERPRISES

- Fortune Global 500 and Fortune 250 companies
- Large numbers of employees and many sites covering large geographical areas
- Scale, performance and reliability requirements similar to service providers
- Interested in adopting the web-scale approach used successfully by Internet companies

# MARKET TRENDS

## Enterprise market trends

As shown in Figure 1[1], several general market trends are affecting enterprises.

**Figure 1. General market trends that will affect large enterprises**

**37%**
of the global workforce of 1.3 billion will be mobile by 2015

**25%**
of enterprises will have an enterprise App store by 2017

**45%**
of employees think personal hardware is more useful than their company's

**50%**
growth in enterprise data by 2017 drives growth in data center and cloud

1 "Predicts 2014: Consumerization Impacts Enterprises' Social and Collaboration Strategies", Gartner, November 2013

These trends include:

- **Mobility:** Increasing numbers of mobile devices such as smart phones and tablets – including personal devices – need to connect securely to enterprise resources and applications, many of them running in data centers.
- **Consumerization:** Employees think personal smart phones and tablets are more useful than their company-provided PC or laptop, and many managers believe consumer applications, such as social media improve customer and employee engagement. The combination of smart devices, mobility and social media apps creates additional data that must be stored, processed and acted upon.
- **Enterprise apps:** Enterprise app stores are becoming the standard IT approach, but integrating these new apps with existing business-critical applications in the enterprise environment presents a major challenge.
- **Data explosion:** Big data and mobile are leading to further data center growth and the trend toward mega data centers and cloud computing. In addition, the data explosion is driving the need for higher security, analytics and disaster recovery, as well as business continuity.

As a consequence of these market trends, many enterprises are going mobile, adopting the "bring your own device" and "bring your own application" approaches to IT, as well as struggling with the growth of big data. Some enterprises are also adopting the "bring your own cloud" approach to IT, in which departments, workgroups or employees use public or third-party cloud services that are often quicker, easier and cheaper than using internal IT resources to fulfil specific needs. However, the lack of visibility into how these shadow IT services are being used and the effect on regulatory compliance have potential implications for enterprise data and data center security. Combined, these market trends are impacting future needs for enterprise data centers and DCI.

## Data center and DCI market trends

Driven by cloud computing, big data and the compelling economics of virtualization, the enterprise data center market is undergoing a period of rapid change, growth and consolidation. Some of the factors driving the data center market include:

- **Size and efficiency:** Large enterprises are consolidating branch office and remote site servers into larger data center facilities. This consolidation of data centers is driving the need to reduce energy consumption, carbon footprint and costs. With the cost of electricity as the largest single factor in ongoing data center operational costs, power utilization efficiency (PuE) has emerged as a leading benchmark. The energy efficiency gain in using more efficient server technology is a major incentive to upgrade data center facilities.
- **Future expansion and scale:** Big Data applications such as Hadoop data mining and data analytics can quickly scale up to multiple racks of high-density servers and SAN arrays. The need to manipulate ever larger business-critical data sets drives the decision to deploy new data center server and SAN equipment.
- **Server virtualization:** The rapid scale up/scale down possibilities of virtualized servers and data centers has created a strong incentive to consolidate virtualized servers into fewer data centers. Ensuring continuous availability and data mobility between data centers across the cloud will become increasingly important.
- **Cloud services:** Leading data center operators are rapidly adopting virtualization technologies to offer cloud services, such as computing-as-a-service, storage-as-a-service and infrastructure-as-a-service. Cloud services are attractive to many large enterprises but SLAs and guaranteed, secure access to these virtualized services remain challenges.

Server virtualization and cloud services have been adopted successfully over the last few years by Internet companies, such as Amazon and Google. This so-called web-scale approach uses a virtualized, cloud-based architecture built on commodity compute, storage and network components. Combined with software-defined networking, it enables capacity to be increased and decreased as needed by rapidly and automatically assigning and de-assigning resources. The web-scale approach provides on-demand access to compute and storage, as well as enabling more cost-effective and scalable outsourcing of IT and apps. It also provides very high resiliency that allow parts of the infrastructure to fail without disrupting services.

Many small-to-medium enterprises find it more cost effective to outsource their data centers by outsourcing or co-locating their servers in a shared data center facility managed by a telecom operator or service provider. Although many large enterprises outsource non-critical IT and network functions to increase efficiency and reduce costs, many prefer to own, lease and or operate their own data centers and DCI themselves. This is particularly the case where proprietary data is seen as business-critical and its security and integrity provide a competitive advantage and remain central to the company's business and operations. These large enterprises need an approach that can scale to meet future needs, respond to competitive pressures, and adapt to data center virtualization, SDN and the cloud.

A report by Gartner[2] predicts that by 2017 over 50% of large enterprises are expected to adopt the web-scale approach to IT that Internet companies have adopted.

Gartner[2] predicts that by 2017 more than 50 percent of large enterprises running their own data centers are expected to adopt the web-scale approach to IT because of the flexibility, agility, efficiency and cost savings offered by this approach. This will have a dramatic impact on large enterprises' future data center architecture, as well as future DCI requirements.

## DCI REQUIREMENTS

As noted, enterprise data centers are consolidating, growing and becoming more efficient, redundant and dynamic, driving the need for DCI. There is also a growing IT dependence on business continuity and disaster recovery as IT becomes an increasingly integral part of the business—linking customers, partners and suppliers. Therefore, DCI is becoming a key component not only of business continuity and disaster recovery but also of workload distribution. These factors are driving DCI needs for more bandwidth, capacity and intelligence to distribute the volume of traffic and transactional data between data centers.

According to IDC[3], the key business drivers for enterprise DCI vary by enterprise segment. However, the most common drivers include more efficient IT resource utilization (banking and financial services), improved end user access to applications (government/education) and simplified data center management (industrial/distribution). IDC found that the most common applications for enterprise DCI continue to be near real-time disaster recovery and business continuity. Future growth areas for enterprise DCI will be dynamic workload scheduling and asynchronous data replication, particularly for cloud-based applications.

According to IDC[3], key growth areas for DCI will be dynamic workload scheduling and asynchronous data replication, particularly for cloud-based applications.

2 "Strategic Technology Trend: Web-Scale Singularity Means Goodbye to Conventional IT Wisdom", Gartner, February 2014
3 "Highlights of IDC's 2013 Datacenter Interconnect Survey of Enterprise Datacenter Respondents", IDC, February 2014

Different networking technologies are available to address different DCI networking challenges and support applications with different performance and latency requirements. The following factors need to be considered for DCI:

- The type of applications, and whether they are synchronous or asynchronous, which will have some influence on the choice of technology and its implementation.
- Overall bandwidth requirements and distance between data centers, which could mean a trade-off between speed and cost.
- Latency requirements for business- and mission-critical applications, which usually dictates the fastest and most efficient technology for applications that demand the lowest latency.
- The types of client interfaces (both computing and storage), which may limit the choice of solution when other factors are taken into consideration.
- Operations such as infrastructure management and service provisioning, which may impact efficiency, costs and the ability to ensure compliance.

While most segments of the market have the same general DCI requirements — such as, high bandwidth, scalability, availability, reliability and adaptability — different segments have specific requirements or priorities based on their business requirements, as shown in Table 1.

**Table 1. DCI-specific requirements and priorities by enterprise segment**

| UTILITIES | OIL & GAS | TRANSPORT | GOVERNMENT | FINANCIAL | SCIENTIFIC/HPC | HEALTHCARE |
|---|---|---|---|---|---|---|
| • Secure data<br>• BC and DR<br>• Low latency<br>• SCADA | • Secure data<br>• BC and DR<br>• Low latency<br>• SCADA | • Low latency<br>• Video sharing<br>• Real-time data<br>• BC & DR | • Consolidation of data centers<br>• Cloud DC<br>• Real-time data<br>• Secure data<br>• BC & DR | • Low latency<br>• Real-time data<br>• High frequency trading<br>• Secure data<br>• BC & DR | • Time of day allocation<br>• Collaboration<br>• Real-time data<br>• BC & DR<br>• Low latency | • Collaboration<br>• Video sharing<br>• Low latency<br>• Real-time data<br>• Secure data<br>• BC & DR |

## Synchronous DCI

Synchronous DCI is used in applications where there is a need to securely mirror or replicate mission-critical data across data centers in real-time, which puts additional performance and availability requirements on the DCI network infrastructure. Examples of synchronous DCI applications are:

- Business continuity, where business-critical data needs to be constantly backed up and mirrored with primary data in as near real-time as possible to maintain data integrity and minimize the possibility of business-critical data getting out of sync.
- Disaster recovery, which involves switching operations to use a back-up data center or recovering backed-up data as quickly as possible following a major disruptive incident.
- High-volume financial trading, where high-value market trades need to be made in as near real-time as possible and with high reliability and integrity.
- Continuous data availability and mobility solutions, where virtualization of storage separates data from physical device limitations and enables business- and mission-critical applications to remain up and running under all circumstances, even during failures.

Big data synchronization for high performance computer cluster and grid computing environments, where huge data files stored on ultra-fast flash storage arrays and accessed by advanced high-performance clustered file systems, need to be constantly synchronized.

Point-to-point WDM is used to provide scalable synchronous DCI because it offers the best cost-performance ratio and lowest latency required for near real-time applications. Typically, large enterprises acquire, lease or build dark fiber connections between sites with optimal paths, and deploy and manage their own WDM equipment. Most enterprise synchronous DCI applications require 10 Gbps connections although a few niche applications — such as high-performance computing (HPC) clusters — are driving a need for 100 Gbps connections.

Most deployments are between sites where distances are limited, such as within a metro area, to reduce latency and cost. For example, financial companies requiring extremely low latency for financial trading typically deploy synchronous DCI over point-to-point WDM links where the data centers are usually within 30 km of the major financial market. At greater distances, additional techniques may need to be used to overcome timing and synchronization issues.

Ethernet transport, in which Ethernet is mapped over optical wavelengths or sub wavelengths, can also be used for some synchronous DCI applications depending on latency tolerance and distance between data centers. Carrier Ethernet (CE) services, such as Ethernet Line (E-Line), Ethernet Tree (E-Tree) and Ethernet LAN (E-LAN) services, defined by the Metro Ethernet Forum (MEF) CE 2.0 standard, can all be used to interconnect multiple data centers for synchronous DCI applications, subject to latency and distance between sites.

In some instances where multiple data centers as well as remote sites need to be connected, synchronous applications can be encapsulated in IP or Ethernet across an IP/MPLS network. One of several encapsulation techniques, such as Virtual Private LAN Service (VPLS) or IP Virtual Private Network (VPN) can be used. This approach allows multi-point connectivity between multiple data centers and remote sites, is more cost effective and easier to configure and can also support longer distances. But packet encapsulation increases latency by adding delay, and may cause delay variation and potentially out of order delivery, particularly over long distances. Although IP/MPLS QoS and traffic engineering mechanisms can be used, there may be a performance trade-off for synchronous DCI applications depending on distance, making IP/MPLS more suitable for asynchronous DCI applications.

## Asynchronous DCI

Asynchronous DCI is used in applications that do not need to be performed in real-time or where latency is less of a concern. Examples of asynchronous DCI applications are server-based replication, network attached storage (NAS) backup and remote tape backup of non-critical data, such as user files and email. Optical WDM and Carrier Ethernet transport services such as E-Line, E-Tree and E-LAN are suitable for many of these traditional asynchronous applications that have high-bandwidth, predictable performance, low latency and multi-protocol requirements. Asynchronous DCI is also being driven by the cloud phenomenon and the desire to extend the benefits of server virtualization to multiple data centers.

However, there is currently a difference of opinion as to whether asynchronous DCI, when used in cloud and server virtualization environments, should use Layer 2 Ethernet or Layer 3 IP. Moving virtual machines between data centers, domains and/or public/private clouds requires tight control and management of IP addresses within the data centers. One view is that this requires transparent IP connections between data centers. A counterview suggests that the network connection between data centers requires transparent Ethernet connections. In reality, large enterprises should ensure that any solution they implement for asynchronous DCI applications supports both IP and Ethernet connections.

MPLS is one solution as it supports both IP and Ethernet connections, provides speeds up to 100 Gbps and implements robust traffic engineering, fast route and link redundancy. MPLS supports Layer 2 Ethernet VPNs, such as Virtual Private LAN Service (VPLS) that can connect several data centers together in a cloud-based solution, effectively acting as a single LAN that spans the wide area. MPLS also supports Layer 3 IP VPNs that can be used to connect cloud-based data centers over the wide area network (WAN). IP VPNs can also be used to connect remote locations that require access to applications running in the data center and provide backup of remote data or integration with the large enterprise private WAN.

## Virtualization within the data center and software-defined networking

Server virtualization within data centers has made compute and storage more dynamically consumable and has paved the move to the cloud. This has also had a profound impact on network infrastructure within the data center. While application turn-up on virtual compute platforms takes only minutes, data center network configuration to support those platforms can take much longer to implement.

Ethernet-based VLANs are widely used within data centers but the limitations of VLAN technology mean that they have limited use for server virtualization. Also, the data center network needs to respond quickly and automatically if the goal of using SDN is to be realized to deliver automatic and on-demand provisioning at scale. This has led to the adoption of virtual extensible LAN (VXLAN) and network virtualization overlay (NVO) technologies for data center networking.

VXLAN is the de-facto overlay data plane standard for data center networking. It encapsulates Ethernet in IP, can be routed by IP and can be terminated on computer infrastructure or network equipment. The underlay network may be any IP network that uses existing routing, resiliency and load balancing mechanisms. Overlays can be viewed as a tunnel between two end points within the data center. They provide a number of benefits, including VPNs for multi-tenancy, network virtualization for location independence of resources within the data center, improved resource allocation and protection from topology or technology changes.

Interest in virtualization and SDN for next-generation data centers has been driven by the early success of web-scale Internet companies, such as Google's SDN data center network, which has attracted attention for its high network link utilization rate of 95 percent. With most implementations operating at much lower utilization levels today (typically 30 to 40 percent), virtualization and SDN are seen as a more cost-effective way of accommodating future growth in the near term by improving link efficiency rather that adding more connections or increasing interface speeds.

### Extending data center virtualization and SDN across the WAN

SDN promises to enable the full value of data center virtualization because it offers greater network agility and efficiency through multilayer resource discovery and control, policy-driven provisioning and dynamic path selection. SDN simplifies and automates service creation resulting in rapid service innovation and delivery. However, to achieve this means making the network inside the data center operate seamlessly and automatically with the wide area network (WAN) used to connect data centers, remote sites and end users.

One approach is to integrate the data center network with existing WAN networking technologies, such as IP/MPLS and emerging technologies, such as Ethernet VPN (EVPN). EVPN provides an overlay control plane and is based on well-established MPLS and BGP routing protocols and operational experience used in service provider networks. It also includes some improvements over existing techniques, such as Virtual Private LAN Service (VPLS) used to deliver Ethernet LAN services over MPLS. When combined with VXLAN as the overlay data plane, EVPN enables data center virtualization to be extended across the WAN between multiple data centers.

Extending virtualization and SDN over the WAN will achieve the goal of activating data center interconnect automatically and dynamically, depending on the application requirements. This approach will deliver unrestricted cloud-based data centers and greatly simplify operations, reduce costs and increase agility.

## OPTIMAL TECHNOLOGIES FOR DCI

Table 2 summarizes the optimal networking technology to support the requirements of different DCI applications. The applications are organized into tiers with Tier 1 applications requiring a solution with the highest performance and lowest latency. Tiers are then mapped to networking technologies that can meet the underlying requirements of the application.

**Table 2. Optimal networking technology to support different enterprise DCI applications**

| ENTERPRISE DCI APP | APPLICATION PERFORMANCE REQUIREMENTS | TYPICAL APPLICATIONS | NETWORKING TECHNOLOGY |
|---|---|---|---|
| Tier 1 | • < 5 milliseconds latency<br>• < 50 milliseconds restoration<br>• < 10 Gbps bandwidth (CWDM)<br>• ≤ 200 Gbps (DWDM)<br>• > 50 ports (Fibre Channel and FICON, InfiniBand transport) | • Synchronous applications<br>• Metro or regional <150 km<br>• Business continuity and disaster recovery<br>• Federated storage – data collaboration and migration (metro/regional)<br>• Content delivery and video caching synchronization | Optical WDM:<br>• Short reach CWDM < 50 km < 10 Gbps<br>• Short reach DWDM < 150 km ≤ 200 Gbps |
| Tier 2 | • < 50 milliseconds latency<br>• < 50 milliseconds restoration<br>• > 20 and < 100 Gbps bandwidth<br>• > 20 ports (Fibre Channel and Ethernet) | • Asynchronous and synchronous applications (depending on distance)<br>• Metro, regional, national >150 km<br>• Business continuity and disaster recovery<br>• Tape vaulting<br>• Bandwidth aggregation (data center consolidation)<br>• Federated storage – data collaboration and migration (regional/national) | Optical WDM (Medium- or long-reach DWDM) or Layer 2 Carrier Ethernet transport |

| ENTERPRISE DCI APP | APPLICATION PERFORMANCE REQUIREMENTS | TYPICAL APPLICATIONS | NETWORKING TECHNOLOGY |
|---|---|---|---|
| Tier 3 | • < 50 milliseconds latency<br>• < 50 milliseconds restoration<br>• < 20 Gbps bandwidth<br>• < 20 ports (Ethernet) | • Asynchronous and synchronous applications (depending on distance)<br>• Metro, regional or national<br>• Server-based replication<br>• NAS and remote tape backup<br>• Remote office access and backup | Layer 2 Carrier Ethernet transport |
| Tier 4 | • < 100 milliseconds latency<br>• < 100 milliseconds restoration<br>• < 10 Gbps bandwidth<br>• < 10 ports (Ethernet) | • Asynchronous applications<br>• Metro, regional, national or global<br>• NAS backup<br>• Remote location access and backup<br>• Integration with private WANs<br>• Extending DC virtualization and SDN over the WAN | • Layer 2 Ethernet, Layer 3 IP VPNs or Business Internet Access<br>• DC Gateway to extend DC virtualization and SDN over the WAN |

## Optical WDM for Tier 1 and Tier 2 DCI applications

Optical WDM is the technology of choice for Tier 1 and Tier 2 DCI to address the very high-bandwidth, predictable performance, very low latency and multi-protocol requirements of synchronous and many asynchronous applications. Different optical WDM technologies can be used to support different distances, depending on factors, such as the power budget of the optics, the wavelength used for transmission, type of fiber, and so forth. For latency-sensitive synchronous applications such as storage replication, the location of the data centers is a critical factor. This is because the speed of light inside an optical fiber determines the distance for such applications, limiting the deployment of synchronous DCI applications to distances of typically less than 150 km within metro or regional areas.

Data centers within a metro or regional area can be connected by building private optical WDM networks using short-reach Course Wave Division Multiplexing (CWDM) technology. For distances of more than 150 km, using medium- or long-reach Dense Wave Division Multiplexing (DWDM) technology, or using a leased or managed wavelength service from a service provider, may be the best approach.

## Private WDM networks and managed WDM services

Large enterprises have several options when implementing WDM for DCI, such as:

- Purchase a managed WDM service from a service provider — In this scenario, the enterprise purchases a complete solution from a service provider that installs and operates the network and WDM equipment on behalf of the enterprise, and provides a managed wavelength of sub-wavelength service.
- Building and running their own private WDM network — In this scenario, the enterprise purchases private, dedicated fiber from a dark fiber provider and purchases, installs and operates its own WDM equipment to create a private optical WDM network.
- Specify and purchase a solution from a network integrator or service provider — In this scenario, the enterprise may specify and purchase the WDM equipment to meet specific requirements, such as security. The enterprise rents a dark fiber service and outsources the operation of the network but maintains control of data encryption and security.

In general, where DCI requirements can be met by a small number of wavelengths, and where performance, security and regulatory requirements may be lower or less stringent, a managed WDM service can provide the lowest cost solution. However, a typical managed WDM service only provides a single wavelength or sub-wavelength between locations, which is less resilient and thus has reliability and availability implications. With this model, overall DCI costs can increase dramatically when large enterprises need to add wavelengths to add resilience or scale their DCI capacity.

Building a private WDM network incurs higher startup costs, but it also brings several benefits. For example, a private WDM network remains the preferred solution for large enterprises with the most stringent performance, regulatory or security requirements. Private dedicated fiber also delivers very high bandwidth and capacity for future DCI needs. Turning up additional wavelengths dynamically when needed incurs only marginal cost after the initial build has been completed. This scenario also provides the highest level of control and the highest level of reliability, particularly when backup paths are implemented. Business cases show that the ROI for a private fiber build can be as low as 12 to 18 months for a scalable, secure optical WDM solution for DCI.

For many large enterprises, the hybrid approach, in which they specify and purchase the solution but rent wavelength or sub-wavelength dark fiber services and outsource the running of the network, is the most appropriate solution. This approach offers more flexibility and control of capital and operational costs while maintaining security of business- and mission-critical data for regulatory and compliance purposes.

## Ethernet and IP for Tier 3 and Tier 4 DCI applications

For Tier 3 applications within metro, regional or national areas, Layer 2 Carrier Ethernet transport provides high-bandwidth, predictable performance and low latency for many asynchronous applications and some synchronous applications, depending on latency requirements and distance. It can be implemented over optical wavelengths or sub-wavelengths as a managed or unmanaged point-to-point E-Line, multipoint E-Tree or multipoint E-LAN service.

**Figure 2. DCI for Tier 1 through Tier 4 applications**

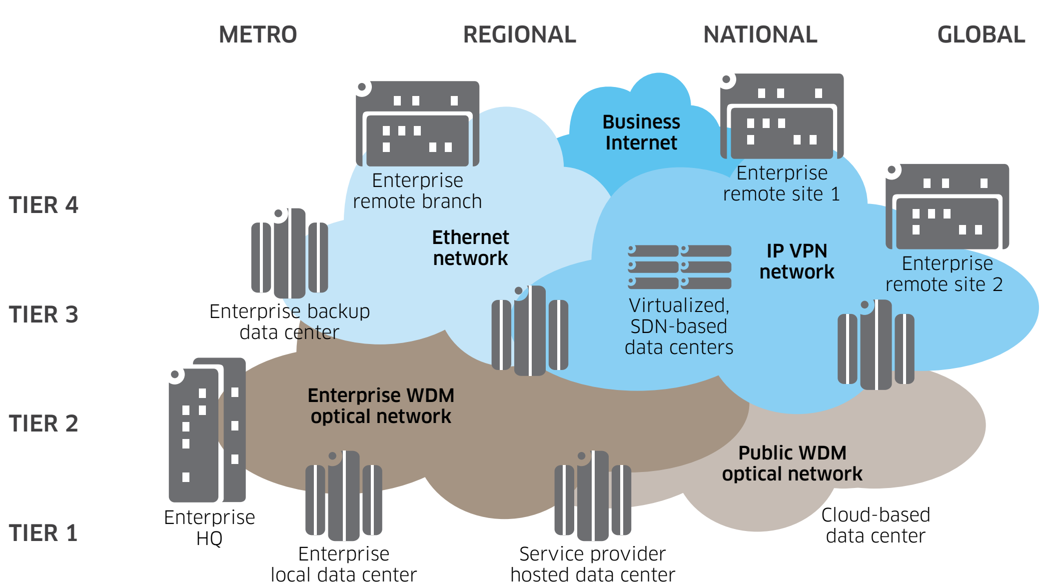

For Tier 3 and Tier 4 applications that have unbound latency, lower bandwidth and regional, national or international connectivity requirements covering multiple locations, a Layer 2 Ethernet or Layer 3 IP VPN solution based on IP/MPLS is the best option. Ethernet VPNs such as Virtual Private LAN Service (VPLS) can connect several data centers together in a cloud-based solution, effectively acting as a single LAN that spans the wide area. IP VPNs and SDN can also provide cloud-based solutions that enable virtualized, cloud-based data centers to be connected over the WAN. IP VPNs also offer easier connectivity for multiple remote locations that require access to applications running in the data center, and can be used to provide backup of remote data or integration with the large enterprise private WAN.

Ultimately, many enterprises will continue to rely on private WDM networks or leased wavelength capacity to support business- and mission-critical Tier 1 and Tier 2 DCI applications. Some enterprises will supplement this approach with Ethernet Transport or Ethernet and IP VPN services to support Tier 3 and Tier 4 applications and to connect to remote enterprise or public data centers. Many enterprises that are implementing virtualization and SDN within their data centers will also look to extend this capability over the WAN. This approach will help them achieve the goal of unrestricted cloud-based data centers that simplify operations, reduce costs and increase agility.

# SUMMARY

A number of market factors are redefining the DCI market. The rapid gains in data center efficiency, the emergence of mega data centers, the impact of cloud computing, as well as the virtualization of data center architecture and software-defined networking mean large enterprises need to consider new approaches to meet their current and future DCI needs. The web-scale approach to data centers that Internet companies have adopted successfully uses a virtualized, cloud-based architecture built on commodity compute, storage and network components. Combined with software-defined networking, it enables data center capacity to be increased and decreased as needed by rapidly and automatically assigning and de-assigning resources.

While optical WDM continues to be the preferred option for synchronous DCI applications, such as near real-time disaster recovery and business continuity, large enterprises need to consider the impact of emerging market trends on their future asynchronous DCI requirements. IP/MPLS is a suitable option for cloud-based data center environments as it provides multi-site connectivity, support for high-speed connections, advanced QoS, traffic engineering and high reliability and availability features.

Data center virtualization technologies, such as NVO and VXLAN, combined with emerging technologies such as EVPN, are based on established network and routing technologies, such as MPLS and BGP. They provide a solution that enables the connection of virtualized, cloud-based data centers over the WAN. When integrated with SDN, EVPN offers a simple, efficient, flexible and agile way of automatically connecting data centers and remote sites while also meeting large enterprises' future DCI application needs.

Alcatel-Lucent provides a wide choice of DCI solutions that help large enterprises de-risk the provisioning of data center interconnect, including:

- Optical WDM solutions that provide high performance, low latency and security for business- and mission-critical synchronous applications.
- IP/MPLS solutions that offer high performance, multi-site data center interconnect with advanced traffic engineering for asynchronous applications.
- SDN solutions that deliver virtualization and unrestricted data center networking within the data center while extending SDN across the WAN to greatly simplify operations, reduce costs and increase agility.

Alcatel-Lucent DCI solutions are used by many large enterprises in the financial, healthcare, consumer and industrial sectors for business- and mission-critical DCI applications. They are also widely deployed in the government, oil and gas, transportation and utility sectors, as well as by many leading service providers and network operators.[4]

For more information about Alcatel-Lucent products and solutions for DCI, please see:

**www.alcatel-lucent.com/products#data-center**

**www.alcatel-lucent.com/solutions/data-center-connect**

4 These solutions are discussed in the Alcatel-Lucent white paper, *Data Center Interconnect Solutions for Large Enterprises.*

# ACRONYMS

| | |
|---|---|
| CE | Carrier Ethernet |
| DCI | Data Center Interconnect |
| DWDM | Dense Wave Division Multiplexing |
| EVPN | Ethernet virtual private network |
| HPC | High-performance computing |
| LAN | Local Area Network |
| NAS | Network Attached Storage |
| NVO | Network Virtual Overlay |
| MEF | Metro Ethernet Forum |
| MPLS | Multi-Protocol Label Switching |
| QoS | Quality of Service |
| SAN | Storage Area Network |
| SDN | Software Defined Networking |
| VPLS | Virtual Private LAN Service |
| VPN | Virtual Private Network |
| VXLAN | Virtual Extensible LAN |
| WDM | Wave Division Multiplexing |