

Data Center Interconnect for the Enterprise Hybrid Cloud

Introduction

The world is moving to the cloud. Everything from entertainment and consumer mobile applications to enterprise software and government services is already living in the cloud or will be headed there soon.

But in the rush to proclaim the inevitable dominance of the cloud, it's easy to forget that there are many clouds, including public and private clouds, and many paths to the future. For enterprises in particular, moving to the cloud can be a challenging multi-year journey. As they make this journey, most enterprises are planning hybrid cloud strategies, incorporating both public and private cloud infrastructures and using multiple cloud services.

In a hybrid multi-cloud environment, the data center interconnect (DCI) networks connecting the clouds become a critical component of IT infrastructure. Increasingly, enterprises are looking to high-capacity optical interconnection technologies to ensure that they can connect all the pieces of their hybrid cloud with security, scalability, performance, and control.

Enterprise Hybrid Cloud Momentum

A large majority of enterprises already use or plan to use a hybrid cloud infrastructure. For example, in one survey,¹ 67 percent of enterprise survey respondents are hybrid cloud users, employing both private clouds and public cloud services, meaning infrastructure as a service (IaaS) and platform as a service (PaaS) offerings such as those from Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP).

The same survey notes increasing interest in using multiple public clouds. The implication of this is clear: few enterprises are ready to move all their applications to a single public cloud.

The reasons for choosing a hybrid cloud strategy are numerous, including:

- Enterprises want to maintain tight control over mission-critical data
- Enterprises want to evaluate cloud services methodically, validating performance, security, and operational processes on some applications before moving to others
- With dozens of applications to consider and limited IT staff, migration to the cloud takes time
- Some legacy applications are hard to “cloudify” (i.e., to modify to run in the cloud)

But even though enterprises may hesitate to put everything into a public cloud environment, they are also reluctant to invest further in data centers that they build and run themselves. That leads to a growing market for colocation and managed or hosted cloud infrastructure services that provide many of the benefits of a private cloud with more flexibility and reduced investment risk. Figure 1 highlights the reasons for using varying infrastructures in a hybrid cloud strategy.

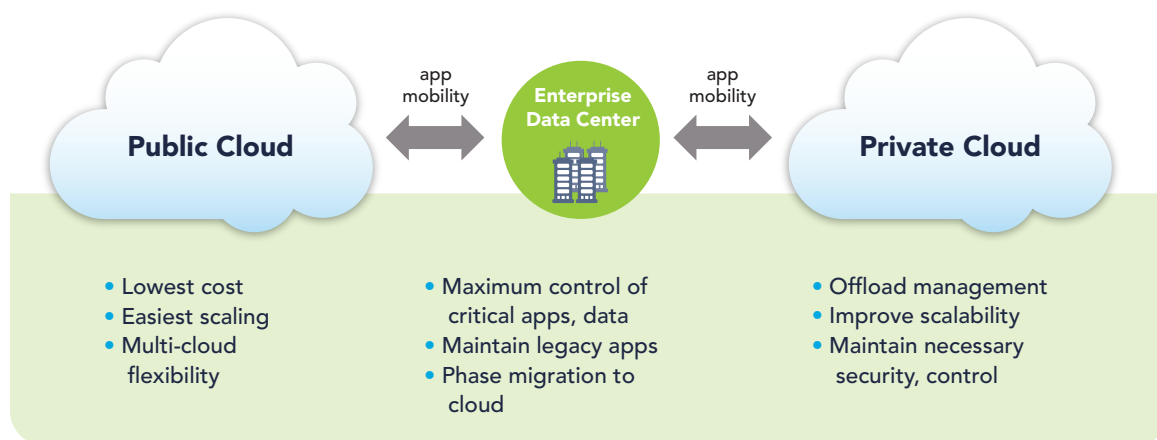


Figure 1: Advantages of different infrastructures in a hybrid cloud



Figure 2: Data center types that play a role in hybrid clouds

As a result, for many enterprises, the hybrid cloud runs in multiple different types of data centers, as depicted in Figure 2:

- Enterprise-owned and -managed data centers
- Enterprise-managed infrastructure in colocation facilities
- Hosted or managed private cloud service provider data centers, referred to as “private cloud” data centers
- Public cloud service provider data centers, often reached via direct connections at an internet exchange (IX) data center or a similar carrier-neutral facility

Connecting a hybrid cloud means connecting many or all these types of data centers. Let’s examine the requirements for such data center interconnection.

DCI Requirements for the Hybrid Cloud

Data center interconnect for a hybrid cloud can have very different requirements than traditional data center interconnect, or even than some early approaches to public cloud connectivity.

Historically, DCI requirements have varied widely depending on the application. The most demanding DCI scenarios typically involved high-capacity and low-latency synchronous data mirroring, in which storage area networks (SANs) were directly connected across fiber optic networks to enable all data to be replicated in two locations to guarantee zero data loss in the event of a failure in one location. Such DCI links must be no longer than around 50 to 100 kilometers (km) to limit the delay experienced by applications waiting for a remote disk write operation to be completed. And since fiber optic networks were historically complex and costly, only the most sophisticated enterprises in financial services and a few other industries could afford to build and operate them. A much larger proportion of historical DCI links were set up with lower capacity and less stringent latency requirements to support asynchronous backup and replication, and occasional data or application migration, with more manageable costs.

Looking at connectivity to public cloud services, many enterprises started off experimenting with connecting via virtual private network (VPN) over the internet. This was sufficient for small applications with limited need to exchange data outside of the cloud environment, such as software development, compute-intensive simulations, or batch processing jobs. As other applications requiring more reliable connectivity and performance moved to the cloud, enterprises have tended to move cloud connectivity to managed connectivity services such as Multiprotocol Label Switching (MPLS) VPNs. But as hybrid cloud DCI bandwidth requirements continue to grow, these services can become very expensive and may not provide the consistently high performance needed.

Looking forward, hybrid clouds will need far more from their networks than some of these past DCI applications. Several trends are driving new DCI network demands:

- **Distributed application architectures:** In a hybrid cloud environment, applications may be composed of multiple components, with a front-end web server in a public cloud, specialized application components running in a hosted private cloud environment, and a back-end database with highly sensitive data running in a tightly controlled enterprise data center. Communication between these components can require several times more bandwidth than the user-to-application traffic, and the application's performance can be highly sensitive to network latency or packet loss in between components.
- **Dynamic application mobility and scalability:** A hybrid cloud offers the potential to use multiple cloud infrastructures for the same application, whether to scale out compute capacity beyond the limits of a single data center, move application workloads based on cost or performance, or simply migrate applications to a different environment over time without incurring any application downtime. Application mobility can be particularly valuable in providing mobile end users the best experience as they travel by moving their applications and data, such as an email inbox, to a cloud data center near their current location.
- **Cloud-based data replication:** Hybrid clouds based in multiple physical locations offer new business continuity opportunities. For many hybrid cloud users, the first step in taking advantage of these opportunities is to use cloud-based data backup and replication to complement other data protection schemes.
- **Continuous application availability:** More sophisticated enterprise users are moving toward a completely new model for continuous global application availability that is not based on bulk data replication and recovery plans, but rather on distributed systems architectures (Figure 3). By running applications in multiple locations and continuously synchronizing critical data across locations, an enterprise can ensure that the applications are protected against the loss of any server or application instance, any copy of application data, or even the loss of an entire location. These architectural shifts dictate two key requirements for hybrid cloud DCI: predictable low latency and loss, and high base capacity with easy scalability.

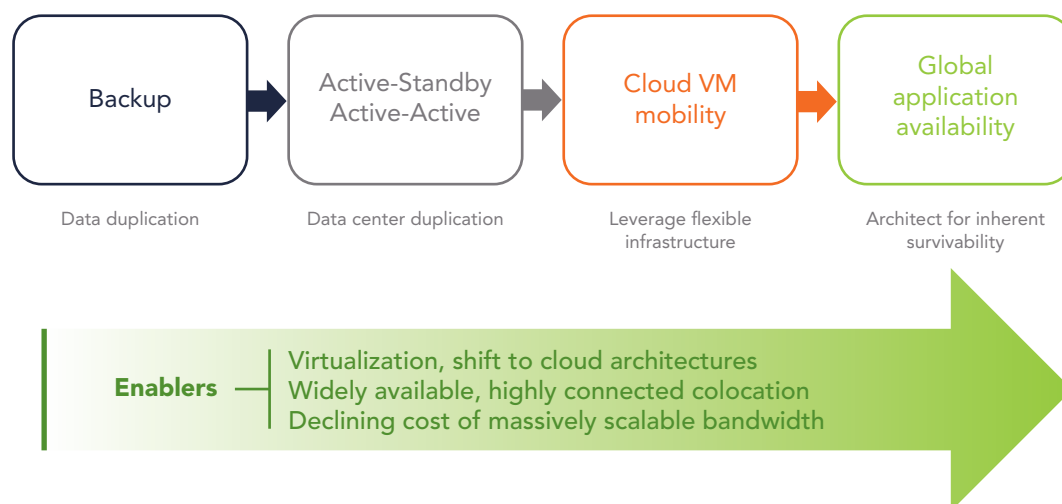


Figure 3: Evolving business continuity and application availability strategies

Another critical requirement for hybrid cloud DCI is security. Overall, security is one of the top concerns for enterprises moving applications to the cloud, and security of DCI links must be considered within that broader concern. To ensure DCI security, in-flight data encryption is an emerging requirement for many enterprises.

One final requirement for hybrid cloud DCI is operational efficiency, which translates into low operational expenses due to monthly recurring telecom service charges, equipment space and power requirements, and operations staff costs.

Below is a summary of DCI requirements for hybrid clouds:

1. Predictable low latency and packet loss
2. High base capacity with easy scalability
3. In-flight encryption of all data
4. Low recurring service costs
5. Space and power efficiency
6. Simple operations

Hybrid Cloud DCI Solutions

Looking at the DCI requirements outlined above, it's clear that internet-based VPNs won't meet capacity and performance needs. MPLS VPNs and other shared packet network services may be sufficient for some enterprise customers, but as capacity and performance demands grow, those services may fall short in terms of insufficient performance service level agreements (SLAs), high recurring service costs, lack of security, or all these factors.

In those cases, the best alternative may be dedicated fiber-based networks with wavelength-division multiplexing (WDM) systems, which can deliver the highest capacity, scalability, performance, and security while giving the enterprise more control as its needs change.

But what about cost and operational complexity? It's true that dedicated optical networks were traditionally too costly and complex for most customers, but those barriers are substantially lower today. In many metro areas, dark fiber is readily available and relatively inexpensive, and newer high-capacity WDM systems are increasingly compact, power-efficient, cost-effective, and much simpler to install and use than older optical network equipment.

Point-to-point and Point-to-multipoint Compact DCI Platforms


To achieve maximum DCI bandwidth with minimum cost, space, power, and complexity, Infinera introduced the first compact DCI platform in late 2014, and adoption has been growing rapidly ever since. Additional platforms have since been added to the portfolio, which now includes the Cloud Xpress (CX) Series, the Groove (GX) Series G30, and the most recent addition, the GX G40, to further reduce the cost of transport and significantly increase capacity while enabling network automation. These platforms support various mixes of client interfaces like 10 gigabits per second (Gb/s), 40 Gb/s, 100 Gb/s, and 400 Gigabit Ethernet, as well as multi-service support including OTN, Fibre Channel, SONET/SDH, and many others over flexible line interfaces from 100 Gb/s to multiple terabits per second (Tb/s). Figure 4 depicts Infinera's compact modular portfolio.



XT-3300



CX-1200



GX G30, G40, G60

Figure 4: Infinera compact modular DCI platforms

Infinera's compact modular DCI platforms offer true plug-and-play operations, with an all-in-one system design requiring no external components; a simple 1-2-3 configuration process; support for automation, including zero-touch provisioning; and the ability to scale capacity with the click of a mouse, adding 100 Gb/s of capacity at a time. Enterprises with no prior experience operating WDM-based optical networks have been successful deploying such systems without the need to go through extensive training.

Flexible Chassis-based WDM Platforms

In some cases, the simplicity of a compact DCI platform is not the best fit for an enterprise's needs, so a more flexible chassis-based WDM platform can be used. Benefits of such platforms include:

- Flexible topologies, including point-to-point, multipoint, and mixes of the two
- Flexible traffic routing and grooming with switching at the optical layer (reconfigurable optical add-drop multiplexer or ROADM), Optical Transport Network (OTN) layer, or Layer 2 (Ethernet switching)
- Cost-efficient support for lower aggregated capacities under 100 Gb/s

Even when such a flexible WDM system is needed, resulting in more configuration steps compared to compact DCI platforms, the overall level of system simplicity and manageability can still be a good fit for many enterprise DCI requirements, as numerous deployments have demonstrated.

Building Hybrid Cloud Networks with Optical DCI

Some hybrid clouds may only need the scalability, performance, and security benefits of dedicated optical DCI on selected links, but as hybrid cloud demands grow, many enterprises will benefit from a completely optical DCI network, built with either compact modular DCI appliances, flexible WDM platforms, or a mix of the two, as shown in Figure 5.

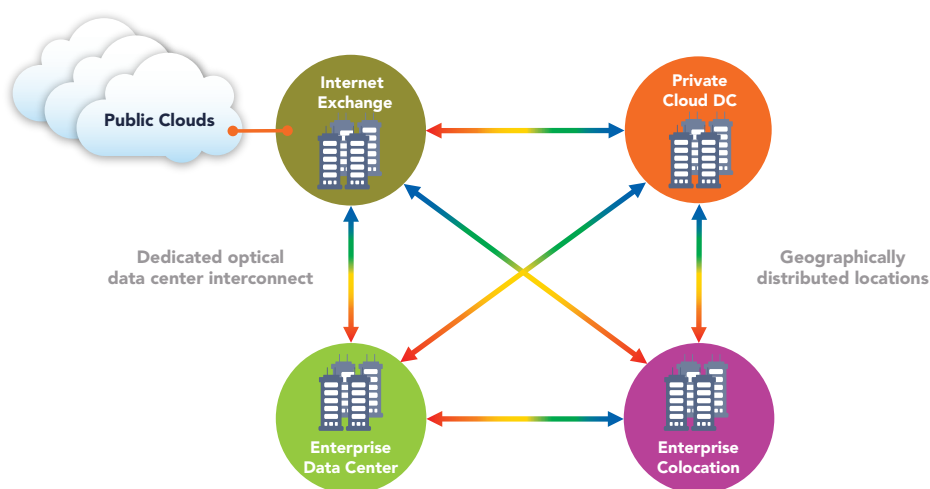


Figure 5: Hybrid cloud architecture using dedicated optical DCI

Summary

Many enterprises implementing hybrid cloud strategies will need to enhance their DCI networks to support their application performance, security, and availability needs. While in the past, traditional dedicated optical networks were too costly and complex for all but a few enterprises, increased dark fiber availability and the emergence of simple, efficient optical DCI platforms now make it feasible for a much wider range of enterprises to implement dedicated optical DCI with the security, scalable capacity, high performance, and control needed to optimize their hybrid clouds and satisfy their end users.

1 RightScale 2017 State of the Cloud Report (<https://www.rightscale.com/lp/state-of-the-cloud>)