



# Data Center Interconnect Essentials

## Connecting Data Centers



Expert Series

# Data Center Interconnect Essentials

## Connecting Data Centers

### Table of Contents

Introduction—Understanding Data Center Interconnect (DCI) .....	3
Rise of the data center and DCI .....	3
Data Center Interconnection Challenges .....	4
Leveraging Technological Advances to Overcome DCI Challenges .....	5
DCI Deployment Scenarios .....	6
DCI Applications .....	8
Simplifying the Deployment and Management of the DCI Network .....	9
Enhancing DCI with Open APIs .....	10
Creating Custom Software Applications for DCI .....	11
Orchestrating DCI Networks .....	12
Improving DCI Reliability via a Control Plane .....	13
Securing the Cloud Between Data Centers .....	13
Selecting a DCI solution to deploy .....	14
Summary .....	15

## Introduction—Understanding Data Center Interconnect

To meet insatiable demand for video, data, and cloud-based content and services, Data Center Interconnect (DCI) technologies have evolved to enable ultra-high capacity, massive scalability, power efficiency, and management simplicity.

As the name suggests, DCI is the technology that connects two or more data centers together over short, medium, or long distances using high-speed packet-optical connectivity. It connects routers and servers to each other or to multi-tenant data centers, so end-users and data can connect to storage and compute resources, cloud applications, and cloud services. It's a broad concept that comprises connectivity solutions across different layers of the network. DCI can be built from the optical layer, with packet-optical technology, or the packet layer, with switches and routers. It can be part of a managed service provided by an operator, or it can be a solution enterprises build for themselves. Networks sit at the heart of DCI, providing the connectivity to meet evolving requirements for operational simplicity, scalability, and automated performance. Leveraging advanced DCI technologies, organizations can speed the delivery of new services, reduce operating costs, and improve flexibility and efficiency.

DCI connects two or more data centers together, so end-users and data can connect to the right cloud applications and services.

By interconnecting data centers, users can consume content on demand, and gain access to information and entertainment anywhere, at any time. The recent growth of DCI has been fueled by content and cloud service providers, such as Google, Amazon, Netflix, and Facebook. These companies have evolved how data centers are built and interconnected to provide content delivery to billions of users. Additionally, DCI helps hospitals meet rigorous business and clinical needs, seamlessly share data among providers, safely back up and store millions of records, and prepare for growth. It provides the networking infrastructure banks need to support a full range of digital products, from digital money transfers to real-time payments, and open banking via APIs. In the education arena, campus IT networks rely on DCI to access and deliver digital

content between data centers and between campus facilities and the data center.

Many enterprises are moving IT resources to multi-tenant data centers and public clouds. In 2018, companies connected to roughly five public and private clouds on average, and 66 percent of cloud users planned to grow their cloud spending by 20 percent or more<sup>1</sup>. As this trend accelerates, connectivity from enterprise data centers and between cloud data centers is expected to grow in lockstep. As the migration to cloud services continues, even standalone enterprise data centers do not exist in isolation. Enterprise data centers must talk to each other—sharing data and content, providing backups, and enabling greater resiliency. With the right combination of solutions, enterprises can ensure the smooth transport of critical assets across any distance and among any number of data centers. With speedy and reliable connections in place, geographically separate data centers can easily share resources and balance workloads.

## Rise of the data center and DCI

A data center is a facility used to store, manage, and process data. It typically occupies anywhere from an entire floor of a building (such as a basement or dedicated floor) to a specially built facility with stringent access controls or even underground tornado-proof fortifications. The number of data centers is growing to match the widespread use of cloud services and virtualization—and these trends are only expected to grow. By 2021, interconnection bandwidth will grow to over 8,200 Tb/s, which outpaces Internet traffic by almost two times the growth rate<sup>1</sup>. These interconnections are shaping and scaling the global digital economy, enabling the rapid growth of multi-cloud services, giving rise to new ecosystems such as the Internet of Things (IoT), and propelling digital business to outgrow the overall global economy by three times.

Multiple trends are driving the need for more data centers. Due to rising content demands, some organizations are moving data centers closer to users, while others choose to locate data centers where power and/or real estate costs are lower. The growth of 'hyperscale' data centers that leverage shared resources and cloud computing, along with stringent security protections, are poised for explosive growth, replacing traditional enterprise Web server farms with cloud services that deliver far greater automation. Hyperscale data center growth is poised to expand, from nearly \$87 billion in 2016 to an estimated \$360 billion in 2023,<sup>2,3</sup> and hyperscale data

<sup>1</sup> RightScale 2018 State of the Cloud Report summary, <https://www.rightscale.com/p/state-of-the-cloud>

<sup>2</sup> Equinix Global Interconnection Index, Vol. 2, 2018, p8 <https://www.equinix.com/resources/whitepapers/global-interconnection-index/>

<sup>3</sup> Hyperscale Data Centers: Market Strategies and Forecasts, Worldwide, 2017-2023, January 2017, Wintergreen Research, Inc. <https://www.researchandmarkets.com/reports/4033103/hyperscale-data-centers-market-strategies-and#rela0>

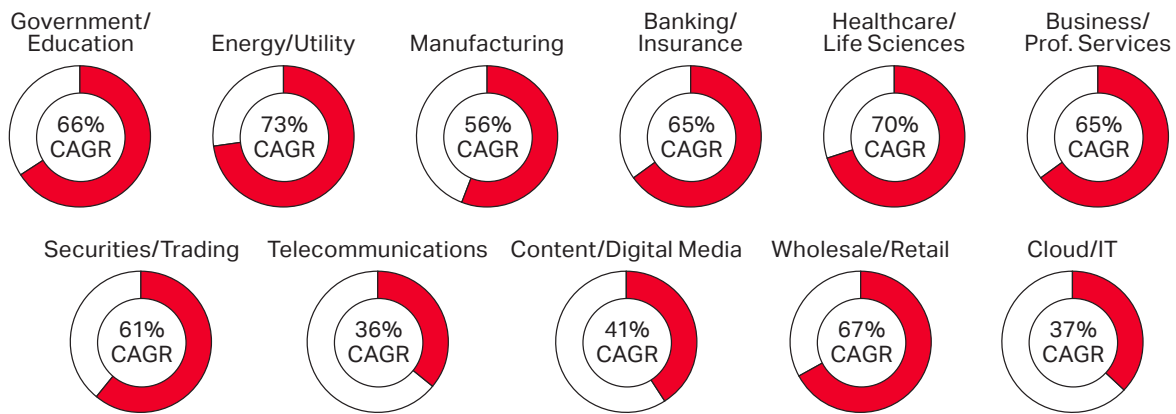


Figure 1. Interconnect bandwidth growth forecast through 2021 by industry<sup>4</sup>

centers are poised to light thousands of fiber pairs over the coming years, driving huge bandwidth requirements between locations that are often geographically diverse.

Cloud and content providers are continuing to build out metro data centers to bring content closer to the end-consumer or enterprise. Additionally, resiliency and digital records costs also drive enterprises and IT service providers to deploy data centers for backup and mirroring applications as part of ongoing Business Continuity/Disaster Recovery (BC/DR) plans. As enterprises move IT applications to the cloud, they require a mix of public cloud, private cloud, and multiple Software-as-a-Service (SaaS) vendors. The data centers for these applications may be located within proximity of each other for synchronous applications or separated by hundreds or thousands of kilometers across oceans for asynchronous applications. Multinational organizations have built data centers around the globe to provide quick and secure access to data for R&D, sales, support, billing, and back-office functions.

As shown in Figure 1, a variety of industries are expected to experience double-digit growth requirements for interconnect bandwidth over the coming years. Whether it's connectivity to the data center or interconnect between campuses, offices, hospitals, or industrial locations, companies require high-speed interconnect for mission-critical applications and projects.

Additionally, cloud services are driving data center and interconnect growth as they gain broad acceptance. Cloud-based services and applications enable convenient, secure, and on-demand access to a shared pool of configurable computing resources. Data centers underpin nearly all cloud services, which is why a cloud service is only as good as the network that interconnects its data centers. Service providers

are projected to grow their cloud interconnectivity capacity by 70 percent annually (CAGR) through 2020 as they work to provide more and better cloud-native digital services.

The demand for instant access to content has ignited a data center evolution to embrace shifting traffic patterns (anywhere, any time, on any device), meet requests for data, and provide the necessary bandwidth to support a wide array of applications and services. Without interconnected data centers, cloud-based applications and content delivery would not function properly.

### Data center interconnection challenges

As the number of data centers has grown and the DCI market has evolved, network operators have had to deploy new capacity rapidly to keep up with the growing demand. However, rapidly scaling capacity has its own set of challenges. The technical and physical challenges in interconnecting data centers include:

- **Capacity and scalability limitations**—Often, the size of application data sets entering or leaving a data center can be quite large, from hundreds of gigabytes to terabytes. This amount of data requires networking equipment capable of providing reliable, high-capacity data center connections with advanced programmability to scale to higher rates as needed.
- **Distance limitations**—As distances extend out of the metro, into the long-haul, and across the ocean, network operators are challenged to provide global connectivity. Costs can dramatically increase as distances between data centers increase, evoking the need for signal regeneration, and the amount of capacity that can be provided per wavelength decreases as a function of distance. These factors increase the overall cost/bit for service delivery. Some services may

<sup>4</sup> Equinix Global Interconnection Index, Vol. 1, 2017, p11 <https://www.equinix.com/resources/whitepapers/global-interconnection-index/>

require a connection with minimum latency to maintain synchronization between the server sending information and the storage device used to save it. Although choosing the shortest physical route can minimize fiber-induced latency, software- and equipment-induced latency must be kept to a minimum through proper design practices.

- **High cost**—With the anticipated traffic growth between data centers pegged at a compound growth rate of 30 percent per year, network-related costs due to rapid scalability must remain under control. Data center space, power, and cooling have on-going operational costs that increase as demand grows. Plus, DCI solution forecasting, planning, and deployment can be cumbersome and complex, especially for smaller operators or enterprises with limited engineering resources and expertise.
- **Slow, manual provisioning of end-to-end services**—To provision connectivity services between network elements located in data centers, operations personnel must use different management systems, unique to each vendor's domain. Building connections between domains/vendors becomes manually intensive, requiring significant time and effort, adding the risk of configuration errors. Additional management systems may exist for each layer in the network, further exacerbating the problem. Connecting two data centers should be a rapid and reliable process. Managing the connection between data centers should not require ongoing manual interventions.
- **Complex, cumbersome platform integration**—Integrating DCI platforms into existing back-office and management software can be difficult and time-consuming. Operators want tools to ease integration efforts and provide access to more data from the network through open APIs and streaming telemetry.
- **Lack of network health information after deployment**—Network operators lack the information and tools to perform network analytics to optimize capacity deployed in the network and diagnose problems before they occur. This increases operations costs and slows the scalability of the DCI network.
- **Security exposure**—Information stored in data centers, including financial transactions, personal records, and corporate data, is mission-critical and confidential, creating a requirement to ensure data center network connections are trusted, reliable, and secure. Legislation increasingly is requiring more stringent identification and notification for security breaches. Depending on the type of data or records being transported, regulations also may require encryption

of the data. Companies can face heavy regulatory fines if sensitive data is compromised, which can affect a company's overall profitability.

Without a simple-to-scale DCI solution that readily adapts, network operators risk decreased competitive differentiation and sub-optimal end-user experience due to lower-capacity, lower-speed connections. Revenue growth may be impeded, due to longer time to market, if DCI networks are unable to rapidly scale and offer faster service velocity. Furthermore, revenue opportunities may be missed if the operator is unable to sufficiently scale capacity across different distances—from metro to regional, long-haul, and subsea. Interconnect capacity must scale to keep up with double-digit industry growth, and it must do so for any distance application.

### Leveraging technological advances to overcome DCI challenges

The industry has responded to these challenges with new, highly scalable technologies and products to meet the ever-growing bandwidth needs of the most demanding and Over-the-Top (OTT) cloud applications and services. The following technologies are being deployed by network operators to resolve their DCI challenges:

- **Overcoming capacity limitations with coherent optics**—Coherent optics have paved the way for the successful transmission of data over almost any distance at rates from 100 Gb/s. Coherent modem technology has unleashed a dramatic increase in transport capacity, which is a key DCI requirement. Cutting-edge coherent optical chipsets increase capacity and/or reach with programmable modulation and advanced soft-decision Forward Error Correction (FEC) techniques. However, not all coherent implementations deliver the same level of performance. It is important to understand tradeoffs between capacity, distance, power, and space to select the best coherent technology for a given application.
- **Increasing fiber capacity with C- and L-bands**—For subsea DCI networks and some fiber-constrained terrestrial DCI networks, lighting both the C-band and L-band can unlock additional capacity across limited fiber resources, doubling overall capacity for DCI services. The need for L-band is driven by the cost of the fiber plant, availability of additional fiber for growth, and sustained bandwidth growth. Coherent modem technology is approaching the maximum fiber efficiency possible within the C-band. Expanding into the L-band opens additional spectrum to maximize return on the fiber plant investment.

- **Overcoming distance limitations using coherent optical processing**—Fiber impairments such as chromatic or polarization mode dispersion have long been roadblocks to high-bandwidth connectivity across great distances. But these limitations are increasingly addressed through breakthroughs in Digital Signal Processing (DSP) using coherent optical processing technology. This has enabled networking equipment providers to introduce packet-optical networking platforms capable of automatically and intelligently compensating for fiber-optic transmission impairments. These advances allow large data flows to be carried over thousands of kilometers using different fiber types without compromising speed or performance.
- **Overcoming latency with high-performance and ultra-high-speed optoelectronics**—Sophisticated hardware design, optimized software engines, and high-performance optoelectronics have significantly reduced network equipment-related latency. Minimizing latency is crucial to the successful implementation of data center applications such as database mirroring.
- **Overcoming high costs through application-optimized platforms**—Today's leading optical networking platforms are meticulously designed and purpose-built for DCI applications. Simple planning, ordering, and installation processes make it possible to quickly interconnect data centers. Full programmability allows data center operators to design and build applications for their specific operational needs. The ability to provide high-speed performance in a small footprint connects data centers at the lowest cost per bit. Small footprints and low power consumption positively impact operating costs, while the modularity of networking components enables scaling to multiple terabits of transport capacity without massive increases in Capital Expenditure (CAPEX) and Operating Expenses (OPEX).
- **Automating manual operations**—Data center networks are constantly changing, leading to traffic trends that are difficult to predict given the current level of spontaneous access to the resource pool by a wide range of users and applications. Many manual tasks can now be automated with APIs and associated applications. Management and orchestration tools or custom applications can be created to execute bandwidth increase requests, set up new connections between two end-points, modify existing connections, and perform daily tasks required in data center-to-data center operations without the need for human intervention.
- **Simplifying integration with increased platform openness**—Platforms must support open APIs to ease back-office integration efforts; however, that is no longer enough. Operators want access to more data from the network through streaming telemetry. Requirements are also emerging for new software architectures that allow installation of either vendor-specific or third-party software components to optimize solutions for any DCI application.
- **Simplifying management complexity in multi-layer/multi-vendor networks**—SDN introduced the concepts of abstraction, centralized control, and software programmability. By abstracting hardware specifics using a defined information model and interfaces, centralized control systems can perform end-to-end management across multi-layer, multi-domain, or even multi-vendor networks. Centralized software control brings high levels of programmability, enabling automation of inefficient operational tasks required for end-to-end service provisioning across multi-vendor domains. Leveraging open APIs, service orchestration tools integrate with other Business and Operational Support Systems (B/OSS) so business objectives can be realized in a holistic, streamlined, automated manner.
- **Extracting real-time network health information**—Advanced coherent processors provide operators with more access to information and data through streaming telemetry, and new intelligent tools can be used to perform analytics on the data to help network operators optimize bandwidth and capacity within the DCI network, or notify them where problems may develop in the network before they occur.
- **Reducing security exposure with in-flight wire-speed encryption**—Increasing data breaches and security regulations have put a strong emphasis on requirements to ensure data is safe whether at rest or in flight, including when it moves from one data center to another. Modern DCI platforms incorporate sophisticated, standardized encryption algorithms, making in-flight data encryption a reality. This provides greater security protection from the moment data leaves a data center until it enters another data center over an interconnecting network.

## DCI deployment scenarios

The proliferation of services and applications that require reliable 'instant' access to user content has been a primary driver behind the growing demand for data centers in metropolitan areas, especially when pushing content closer

to users. Enterprise migration to cloud services (private, hybrid, or public) has fueled the deployment of data centers that serve a wide variety of applications, including mirroring or backup services. Other typical business services, such as email or Web pages, may not require instant access, so they can be hosted from data centers in less costly, more remote regions.

Some typical DCI deployment scenarios include:

- **Inside large data centers**—Although DCI is used to connect two geographically diverse data centers, the same technology can be deployed within a single location, such as a large data center building or a multi-tenant data center, to interconnect floors or co-location areas. High-speed coherent optics can reduce the number of fiber connections required between these areas to lower fiber connection fees and ongoing costs. Connecting server clusters within a single data center can help improve an enterprise's ability to manage big data analytics, for example.
- **Campus interconnect/carrier-neutral provider/data center farms**—DCI can connect data centers inside a campus or a wholesale data center premises. In this scenario, data centers from different companies are co-located on the same campus. Research and education can use DCI to interconnect facilities within a large university campus or research campus. Interconnecting data centers across campus environments ensures flexible, fast connectivity while maximizing investment in deployed fiber and existing photonic line infrastructure. DCI speeds service activation, and platforms optimized for DCI can reduce footprint, power, and cooling across data centers campus-wide. Very often, high-capacity interconnect services are needed between the various privately- owned data centers (Figure 2).

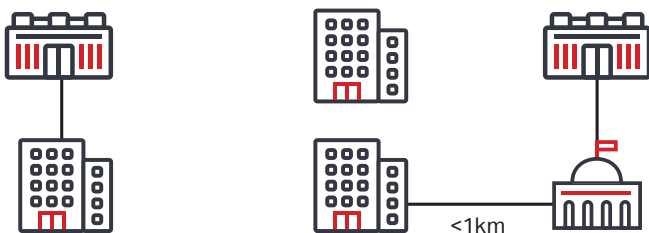


Figure 2. Campus DCI/data center farms

- **Metro DCI**—As cloud and content providers attempt to move content and services closer to the enterprise and end-user, data centers are popping up all over major metropolitan areas. Connecting data centers in metro areas can bring challenges, such as restricted space, limited access to fiber, stringent power consumption and heat dissipation requirements, a mix of old and new fiber plants, and limited operational

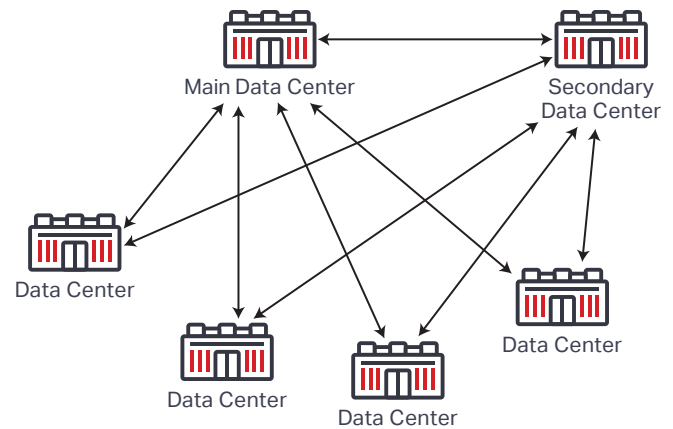


Figure 3. Hub-and-spoke DCI architecture

windows. Metro interconnect distances typically vary from a few kilometers to 100 kilometers, and typically are deployed across simple, point-to-point links.

If multiple sites are connected to a larger central data center, then a hub and spoke architecture may be used. The hub location may be mirrored with a secondary data center for operational continuity purposes (Figure 3). This provides data replication from various corporate sites, branch offices, or smaller data centers to primary/secondary data centers for disaster recovery and continuity applications.

Recently, data centers have evolved to make use of spine-and-leaf architectures to arrange servers and switches to optimize data flows, so they are more predictable, consistent, and fault tolerant. In some cases, these architectures extend across links between nearby data centers, and packet technology is used across the interconnect paths. Layer 2 and Layer 3 packet switching can segregate data based on its community of interest and the location of the data within the cloud. Combining packet-based techniques with optical network technology helps make efficient use of links (intra- or inter-data center) to maximize performance and minimize expense.

Connecting data centers across metropolitan regions can deliver strong competitive advantages, along with cost and process efficiencies. Deploying a new interconnect architecture to link bandwidth exchanges to cloud exchanges within a metropolitan area brings enterprise customers and cloud providers together. Data center operators can improve enterprise end-user experiences through metro DCI by delivering higher-capacity and higher-speed connections.

- **Long-haul DCI**—Large, mega data centers often reside in remote locations with access to power sources and lower-cost real estate. These mega data centers provide massive amounts of storage and compute resources and are poised

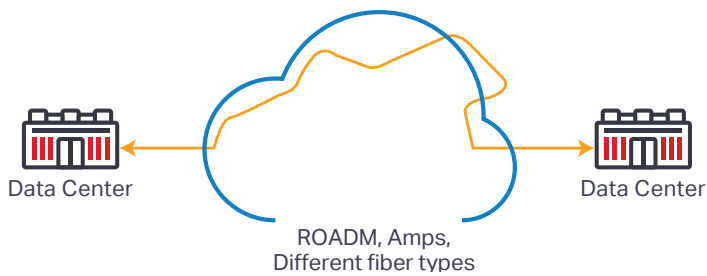


Figure 4. Point-to-point DCI over ROADM-based line system

to light thousands of fiber pairs over the coming years, driving huge bandwidth requirements between locations that are often geographically diverse. They may be interconnected via point-to-point links (Figure 4). ROADM may be deployed as part of the photonic line infrastructure, giving more flexibility when lighting connections between locations.

As the long-haul DCI network becomes more robust, it may be migrated to a mesh-based architecture to provide protection and restoration for important, high-value services (Figure 5). This topology consists of connecting each data center site to at least two other sites to increase the network's overall survivability.

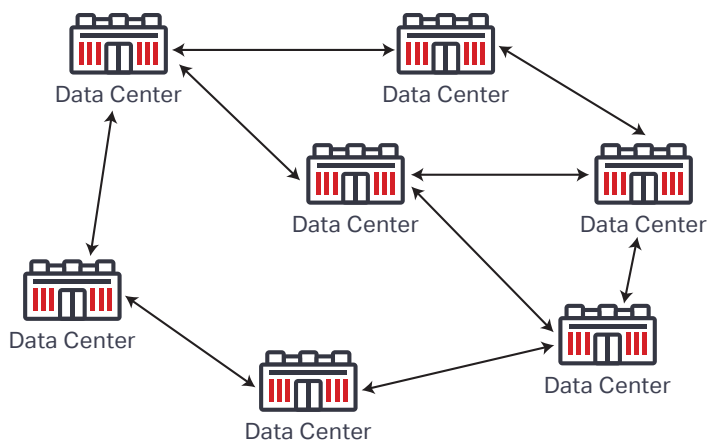


Figure 5. Mesh-based DCI network

- **Subsea DCI**—Data center operators have invested in continental-scale and international networks. Globalization, a booming digital economy, access to lower-cost energy, and real estate costs are fueling the creation of data centers and content hubs around the globe in strategic locations. As business and connectivity requirements become more global, there is a need to interconnect these data centers across distances that stretch from the terrestrial networks across the ocean (Figure 6).

This type of ultra-long-haul connectivity requires high-performance DCI technology to maximize the capacity provided across the subsea cable and enable longer-distance subsea connectivity across the Atlantic or Pacific.

These DCI deployment scenarios cover a broad spectrum of applications across a variety of distances, from a few meters between floors to thousands of kilometers spanning oceans. The requirements for each deployment scenario can vary, depending on parameters such as distance, capacity, scalability, and network flexibility.

### DCI applications

Today's cloud and enterprise applications rely on high-capacity DCI. As companies move more processes and services to the cloud, the demand for capacity is increasing. Some common DCI applications include:

- **Business Continuity/Data Recovery (BC/DR)**—These applications strive to minimize the impact of data center downtime in the aftermath of an emergency by developing a backup plan to transfer data, access control, and provide other activities to a single or multiple secondary data centers. Interconnecting data centers for recovery operations is crucial to minimizing the impact of downtime on an organization. Once a failure or disruption is detected, recovery time varies, ranging from a few seconds—to automatically



Figure 6. Subsea DCI networking

transfer control and reroute traffic from primary to secondary data centers—to several minutes, when a human response or interaction may also be required. It is critical to ensure primary and secondary data centers are equipped with sufficient bandwidth to handle traffic for backups and restoration services and accommodate client-server, storage, and other recovery requirements. The bandwidth required is dictated by each organization's critical applications and the amount of data covered in the BC/DR plan.

- **Off-site/off-hours data backup**—This is the recurring operational procedure required to back up enterprise data after hours every day to a remote data center. Typically, a large amount of data is automatically duplicated at another data center overnight during a backup window that spans from 30 minutes to several hours. DCI bandwidth dictates the length of this data backup window.
- **Data mirroring**—Used by financial institutions and brokerage firms, this application instantaneously duplicates the same set of processes and transactions from a primary or master data volume, such as a financial transaction mainframe, to a secondary volume or mirroring site.
- **Enterprise cloud applications**—Other cloud applications, beyond data backup and disaster recovery, are becoming more widely adopted by enterprises and businesses. These applications vary across different verticals, from business to education, healthcare, and finance-specific applications. Basic enterprise applications include apps for email management, document management and storage, payroll, HR services, employee benefits, procurement, IT services, and more. Cloud-based banking applications can be used to verify transactions, protect against fraud, and automate processes. All require fast, reliable connectivity to the data center and, in many cases, between data centers to interconnect cloud applications, computing, and storage solutions to form a complete enterprise cloud solution.
- **Live Virtual Machine (VM) migrations**—As enterprises expand globally, a DCI network can be used for real-time migration of VMs. Setting up and provisioning required connections and capacities to migrate VMs may be based on time of day and time zones in a 'follow the sun' manner for workload balancing. VMs may also be migrated leveraging 'less-used' capacity sites or to move processing power closer to applications or users, as required.
- **Edge computing and big data analytics**—IoT devices and things like autonomous vehicles will drive more data to the data center and require more storage and computing for analysis. Analytics for large experimental data sets

and big data require more resources and larger bandwidth connectivity to and between data centers. As data and applications are pushed closer to the user, scaled-down data centers will pop up at the edge of the network, driving additional interconnect bandwidth requirements. Localized edge compute resources will be used to provide higher performance and reduce latency across a range of applications.

- **Content distribution**—Data centers are where content lives, so streaming video providers, online music providers, social media companies, ad servers, and other content providers rely on connected data centers to deliver content to their audiences over a variety of devices including mobile phones, smart TVs, tablets, and laptops.

### Simplifying deployment and management of the DCI network

To simplify deployment, enhance scalability, and enable content delivery and cloud applications, operators are turning to new, purpose-built DCI systems. These systems enable rapid and massive scalability to grow capacity for cloud applications and content distribution. In addition to scale, they reduce the cost per bit of DCI connectivity with advanced coherent interfaces and reduced footprint and power utilization.

Software management and orchestration tools can be deployed to manage the network as a dynamic pool of resources for various cloud and content distribution applications. Cloud applications, such as off-hours data backup, VM migration, and content delivery may request a change in the bandwidth allocation of one or multiple links, set up a new link between two data center end-points, or tear down a link when it is no longer needed. With the right tools and intelligence, the network can readily adapt and become more proactive about making real-time changes in the event of deterioration in Quality of Service (QoS). For example, if there is an increase in bit-error rates across some spans, the network can automatically take steps to prevent performance degradation or outage by re-routing services across different paths within the DCI network.

The dynamic coupling between the network that connects data centers and the cloud delivers more efficient, scalable, and flexible data center-to-data center connectivity and enables the deployment of new and differentiated services such as Bandwidth on Demand (BoD), with which users and their applications can request bandwidth and networking resources in a highly automated and customized environment.

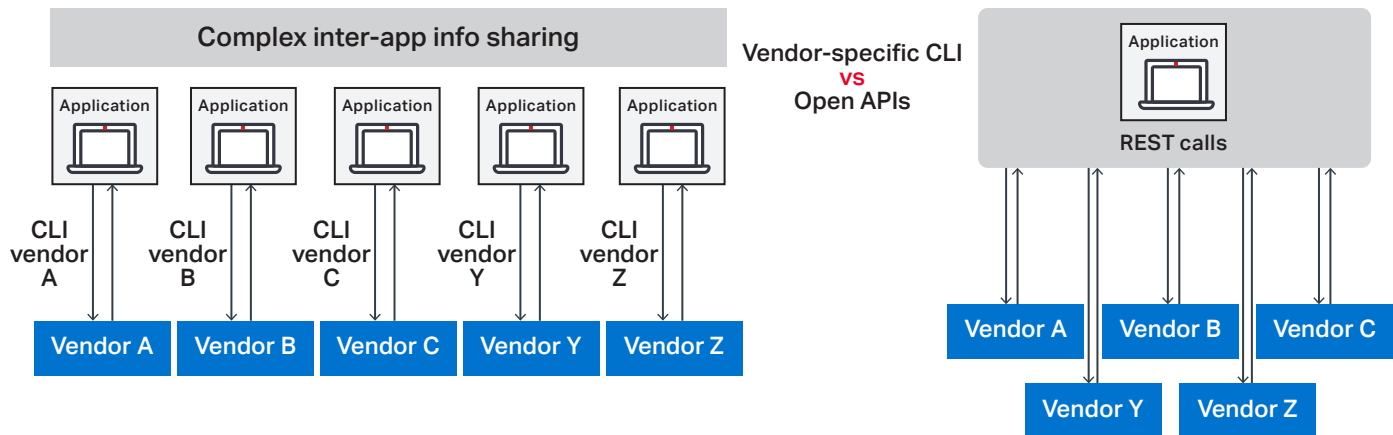


Figure 7. APIs enabling multi-platform/multi-vendor application development

Service providers may also offer Infrastructure-as-a-Service (IaaS), Platform-as-a-Service (PaaS), and Network-as-a-Service (NaaS) subscriptions that deliver the highest levels of resiliency and performance.

New capabilities are continually being brought to market to evolve DCI platforms for greater scalability, easier installation and deployment, and simplified management and orchestration to improve service velocity and enable new types of interconnect services.

### Enhancing DCI with open APIs

The shift to open industry standards has improved the design and building of next-generation network equipment. This shift reduces and ultimately eliminates dependency on various function-specific software tools built and maintained by equipment suppliers. This approach has also helped accelerate the evolution to open cloud networking and virtualization.

Open, standards-based networks and applications use a modular approach for designing, building, and managing networks that interconnect data centers to enhance performance, increase virtualization, and accelerate automation. Open Application Programming Interfaces (APIs) are an open set of published routines, protocols, and tools for building software applications that specify how software components and services should interact. These open APIs can be used to build customized software tools and applications for a network operator's specific needs, or they can be used to integrate DCI hardware and software solutions into existing back-office systems.

Development APIs enable both Development (Dev) and Operations (Ops) teams to create customized application

and automation tools that can help enable and optimize DCI. Benefits of using open APIs include:

- Faster, easier multi-platform/multi-vendor application development**—In the present mode of operation, data center operators rely heavily on sets of protocols and commands provided by networking vendors. Each Network Element (NE) type comes with its own list of Command Line Interfaces (CLIs), which can be used to perform all NE configuration and management tasks. Often, data center operators must deal with NEs from numerous vendors, or even multiple NEs from the same vendor, which feature completely different CLIs and increase operational complexity. These limitations are eliminated when using open APIs, as NE management and configuration are no longer bonded to each NE and vendor-specific CLI. DevOps teams may now create software applications that can interact with different NE types operating on different network layers built by different vendors. Information can be shared across multiple software applications to create intelligent tools and capabilities for advanced network operating tasks (Figure 7).
- Easy integration with IT tools**—Open APIs simplify integration of hardware and software components into existing IT tools, such as programs and scripts. Open APIs are the glue between DCI products and the network operator's B/OSS.
- Efficient IT resource utilization**—Quite often, the use of an Element Management System (EMS) or Network Management System (NMS) requires the creation of numerous and long-lasting remote access sessions that put a significant strain on IT resources to handle session accessibility and management. Development APIs like RESTCONF and NETCONF were built to be very light on IT resources. Commands can be sent to the network element, and a reply with the requested data is

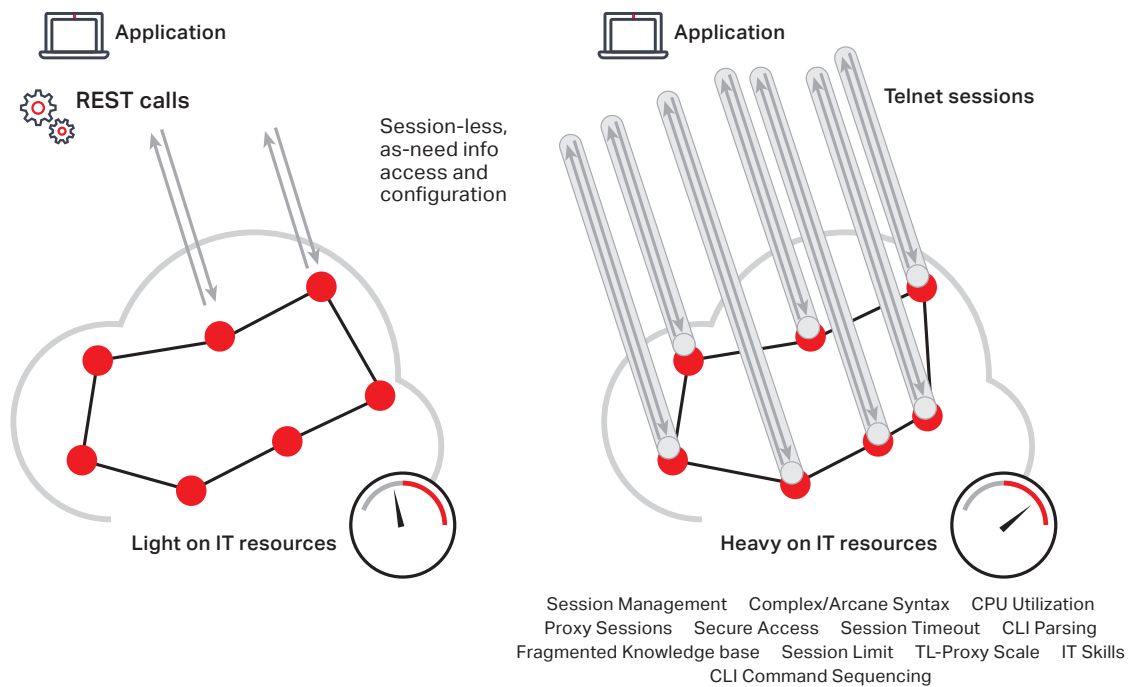


Figure 8. APIs enable efficient IT resource utilization

returned over a short-duration session that doesn't require heavy IT resources (Figure 8).

This is like the way a webpage loads when a user types the URL of a website in an Internet browser. The page loads after a short period of time, and no other information is sent until the user clicks on a button or a link. Also, the use of APIs translates into CAPEX savings by eliminating the use of servers and other compute devices between software applications and the network elements being managed.

Not only must APIs be open and published rather than closed and proprietary, they must be robust and offer access to more information and data from each network element. The success of building a software application using APIs depends on the variety and types of data that can be accessed and interpreted to perform various tasks required by the end-user.

Think about a car dashboard that can only display speed, where the only sensor in the vehicle is the speedometer. Other information about fuel, oil pressure, coolant temperature, and so on are unavailable to the driver. Advanced features that detect lane deviation or potential impact, which could save the driver's life in an emergency, would also be missing.

APIs are underutilized unless key data can be accessed and interpreted by tools for operational benefits. In other words, DCI platforms must possess features and capabilities to provide a

wide set of data points (sensors) for intelligent interpretation and analysis. These are referred to as rich APIs. Using the car analogy, rich APIs are represented by the car's ability to provide more complete information, such as engine temperature, fuel level, current fuel consumption, tire pressure, and other measurements, readily available via the car's on-board computer and displayed in an easy-to-understand manner. A next-generation DCI platform should be able to provide access to a wide spectrum of data points to interface into smarter, customized software applications.

### Creating custom software applications for DCI

Demand for greater flexibility and programmability is increasing to help accommodate traffic demands that arise from cloud services, quick, high-quality video and content distribution, and IT virtualization. Open APIs enable advanced programmability for faster multi-platform/multi-vendor application development, easier integration with IT tools, and more efficient utilization of IT resources. Content providers and data center operators are starting to leverage this new model, developing software applications and tools based on specific needs.

Cloud-based application development portals unlock the potential of open APIs to simplify integration activities and speed the development of new tools. Developers and IT teams can test applications without investing in IT infrastructure or stressing IT budgets and resources. Much like a smartphone and its related application development environment,

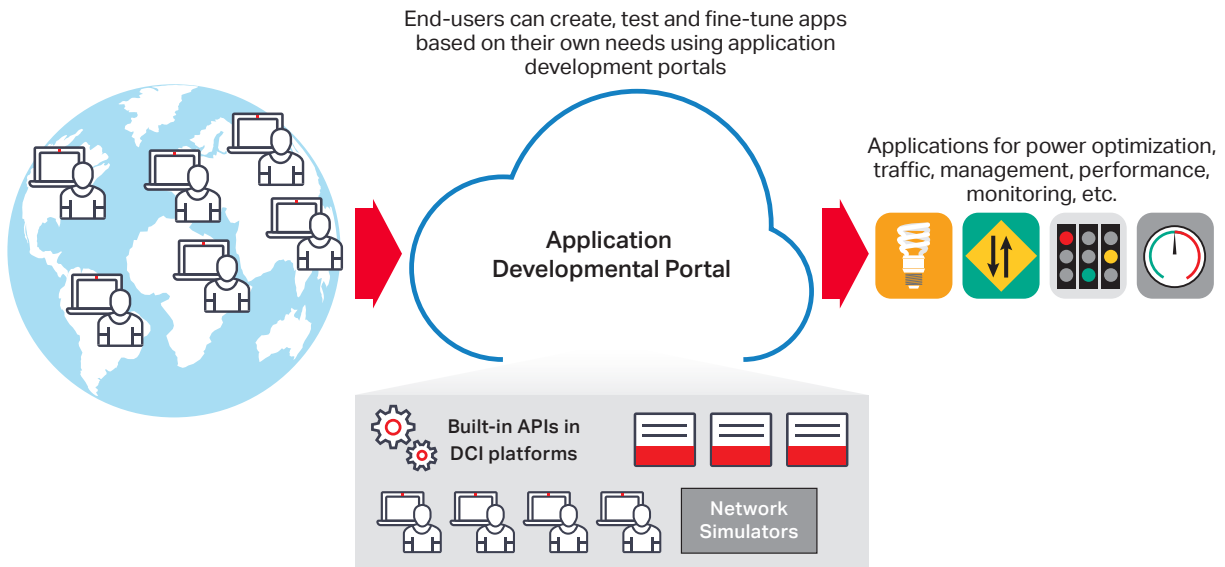


Figure 9. Application development portal for DCI

application development portals can be used to develop and test applications that leverage open APIs, such as apps for enhanced network visualization, fault and capacity event notification, data center cluster management, or even detailed performance monitoring.

These application development portals allow developers to innovate, experiment, and test new service models on virtual network resources, resulting in dramatically lower entry and exit expenses. Network and data center operators can leverage external development resources to scale and accelerate the development of software applications, emulate specific network configurations, model what-if scenarios, and overlay different information sources to create intelligent software tools that can help streamline operations (Figure 9).

Ciena's Emulation Cloud™ helps network operators exploit Ciena's open APIs by providing a Web portal to quickly and easily develop, test, and integrate new applications, without investing in IT and network infrastructure.

### Orchestrating DCI networks

Often, the infrastructure used to connect data centers is composed of multiple technological layers and specialized domains that span access, metro, and long-haul backbone networks. Connecting every element along a path between data centers is traditionally a manual process that entails updating multiple vendor and domain-specific element managers and SDN controllers, then integrating these elements with back-end Operational Support Systems (OSSs). This deployment model is complex and error-prone and preserves the management 'silos' of legacy platforms, making the migration to SDN and NFV more complex and less efficient.

Multi-Domain Service Orchestration (MDSO) provides an open software layer that eliminates management silos and enables data center operators to automate end-to-end provisioning and orchestration of all network resources, even across network equipment from different vendors. Leveraging open APIs

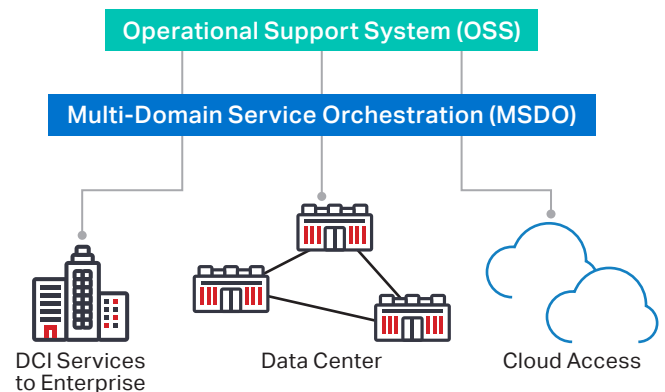


Figure 10. DCI network orchestration

and model-driven templates, an efficient MDSO integrates with third-party SDN controllers, elements, and network management systems, along with other platforms, to manage/orchestrate services that include physical and virtual resources and use multiple heterogeneous technologies and vendor domains (Figure 10).

MDSO is used by organizations to get their connectivity services up and running quickly to support enhanced and innovative applications for customers and API-based services for partners. MDSO helps enterprises allocate bandwidth dynamically and eliminates complexity in service delivery processes. Powerful automation capabilities minimize the potential for error associated with manual setup and configuration. It can also interwork with machine learning platforms to improve network reliability by predicting points of hardware failure before they occur.

### Improving DCI reliability via a control plane

Downtime in content distribution networks translates to losses of millions of dollars in advertising revenue and can seriously damage a provider's reputation and brand image. Similarly, a network failure can have a disastrous impact on cloud

applications such as data mirroring/backup or BC/DR. Therefore, network operators may choose to leverage an intelligent control plane to operate as the brain of the network, reacting to network changes in real time, without human intervention. Network changes may include anything from multiple simultaneous failures to an increase in latency across any one of the network's critical spans. The control plane improves network availability and protects against failures such as fiber cuts or hardware failures that would otherwise impact connectivity between data centers. It makes the network interconnecting data centers completely autonomous and self-healing, even in the event of multiple fiber cuts or hardware failures.

When coupled with virtualization capabilities, a DCI network may be partitioned to allow specific links, wavelengths, sub-wavelengths, or even nodes to be dedicated for a specified use, with preset thresholds for latency, bandwidth, and resiliency. Some virtual networks are configured for high capacity and/or high resiliency for data center-to-data center connectivity. Other virtual networks are created with less stringent requirements for user access or other applications (Figure 11).

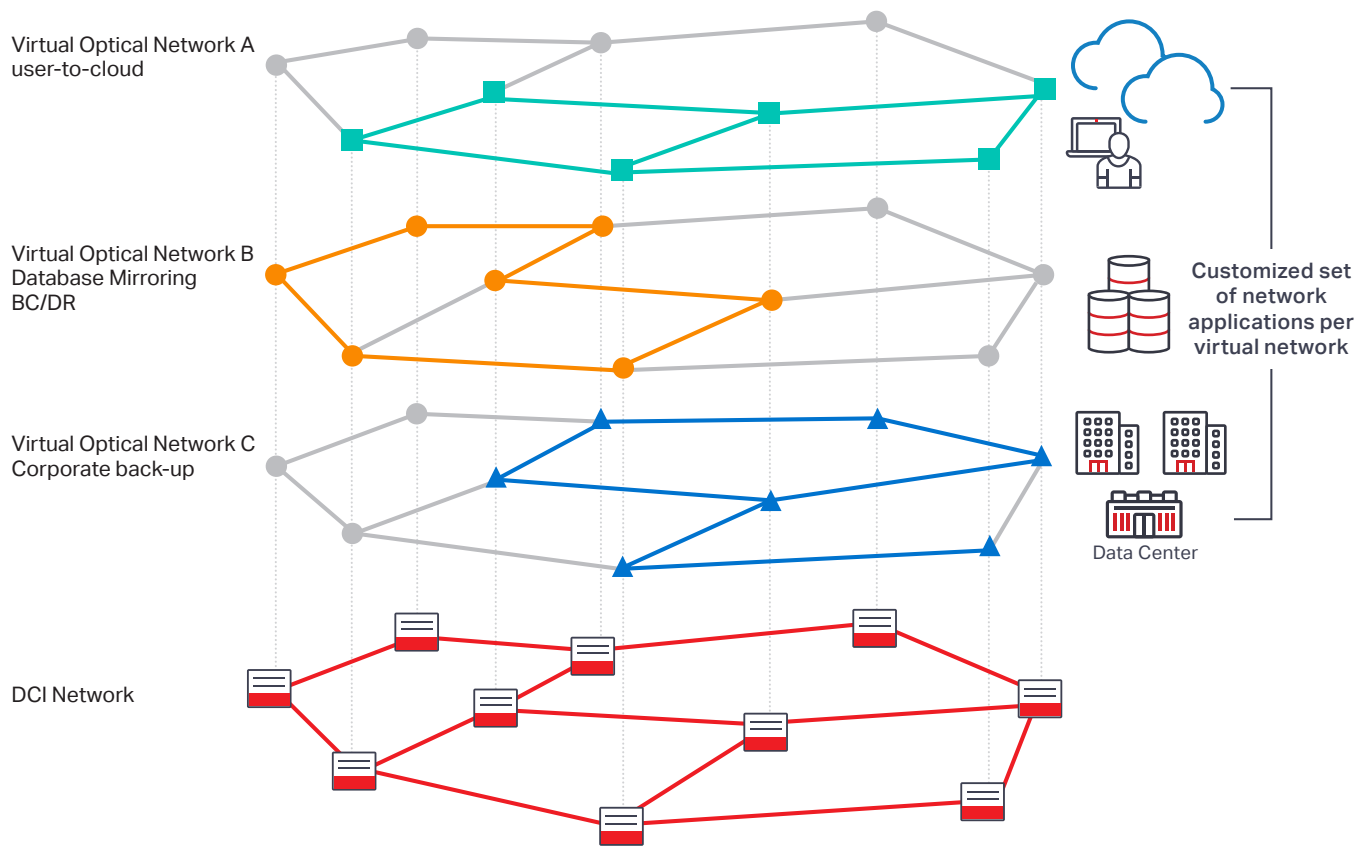


Figure 11. Virtual DCI network partitions support different application requirements

## Securing the cloud between data centers

Cloud usage streamlines business operations and makes it easier to access critical information and back up data. However, as applications increasingly send and receive data between different locations, organizations must ensure that data is secure. The security of the communications link between applications in different locations is often overlooked, because the fiber networks that interconnect data centers have been viewed as impenetrable. But data often travels across third-party and/or service provider networks, so it is not under a single organization's full control. A simple, scalable way to secure data traveling across the network must be adopted to ensure confidentiality. Today's DCI platforms can provide encryption at Layer 1 for bulk encryption of all in-flight data passing across the encrypted link. This can be a simple first step for organizations implementing a security policy between data centers or in the cloud. Because it encrypts all data on the link, it even secures communications between legacy apps that may not fit into other security frameworks. Furthermore, it can be used as an additional layer of defense in case an upper-level security compromise occurs. Organizations have the flexibility and choice to use Layer 1 encryption as a catch-all at the bottom layer or as part of a more comprehensive security plan.

## Selecting a DCI solution to deploy

Choosing the best solution involves multiple factors that impact business transactions and internal operations. Depending on traffic levels and capacity needs, scalability may be a top requirement, or pay-as-you grow licensing may be more appropriate. If space is a priority, then compact, dense platforms that maximize capacity in a small footprint will apply.

In addition to the technology, it is important to select a vendor that has proven solutions that de-risk the DCI deployment. Market leadership in both DCI and small form-factor devices leads to depth of experience that is especially useful when deploying a new solution set. For organizations with limited engineering resources, a vendor with robust services offerings can help to accelerate the DCI deployment.

Depending on the specific application, a DCI solution should address some or all of the following:

- **Peace of mind for traffic growth**—High capacity and seamless scalability to multiple terabits per second can help data center operators meet traffic growth requirements for the foreseeable future without massive injections of CAPEX or network disruptions due to ongoing capacity increases. Highly programmable coherent interfaces allow network

operators to closely match capacity to system margin and support applications, from metro to long-haul to subsea, with a single platform. Smaller operators should select platforms that offer scalability for lower traffic volumes and provide pay-as-you-grow options to rapidly react to changing demand or bandwidth growth.

- **Simple operations and quick turn-up**—Purpose-built DCI platforms are specifically designed for operational simplicity, and a good DCI platform should enable hassle-free execution of DCI projects. Easy planning, fast turn-up and configuration, rapid service provisioning, and intuitive management can all help lower operational costs and speed execution velocity. Whether the end goal is to underpin content delivery or offer a DCI service to enterprise customers, a platform with simple operations and quick turn-up will improve business performance and time to revenue.
- **Reduce recurring/operating costs**—Compact, DCI-optimized designs with low power consumption and small footprints allow data center operators to reduce electricity, cooling, and real estate costs. Simple product architecture also helps lower management, sparing, licensing, and training costs.
- **Mitigate security risks**—DCI platforms can provide advanced encryption algorithms to protect in-flight data as it travels from one data center to another. Organizations can mitigate security risks that arise from interconnecting applications, data, networking, and security controls across regions and across the globe.
- **Openness and programmability for easy integration with back-office tools for task automation**—Northbound Interfaces (NBIs) and REST APIs allow data center operators to automate labor-intensive manual tasks to significantly reduce maintenance windows and free staff from repetitive, error-prone manual tasks. Application development portals can be leveraged for education on open APIs, and data center operators can develop and test unique and customized software tools in the cloud, without investing in physical hardware. Service orchestration with MDSO brings the power of software-defined programmability to enable automated lifecycle operations and resource orchestration for DCI networks.
- **Increase flexibility**—A good DCI platform should be designed to be flexible by supporting various connections and interfaces, protocol rates, modulation schemes, and deployment scenarios, including those run over new or existing photonic lines. Capabilities such as packet aggregation and switching also enable data center operators

to meet growing Web-scale needs and improve port utilization to reduce the number of expensive router interfaces required.

- **Improve adaptability**—Networks are changing dynamically to meet growing demands. A comprehensive DCI solution must collect key performance indicators and real-time telemetry from the network and analyze them to provide suggestions for proactive actions that will keep the network up and running. DCI hardware and software solutions that provide a view of the resources and health of the network can help optimally provision wavelengths and services to ensure the DCI build-out is most effective. Such a solution requires underlying infrastructure instrumented to collect real-time metrics, leveraging analytics and automation to process those metrics and drive informed, efficient decisions.
- **Field proven**—Proven solutions de-risk DCI, and these solutions leverage field-proven technologies to carry mission-critical traffic. Vendors with market share leadership and leadership in small form-factor devices have deployment experience, and vendors with a wide range of customers across a variety of industry segments have the knowledge to assist with any DCI deployment—big or small.

**Ciena is the undisputed market leader and partner of choice for DCI.**

- Ciena has more experience in DCI than any other vendor in the market.
- Ciena de-risks DCI with proven, innovative solutions for DCI in any application space.
- Ciena is a leader in small form-factor devices, purpose-built for DCI deployments in nearly 120 networks globally.
- Ciena's highly programmable infrastructure and intelligent software solutions combine to make the DCI network adapt to change more readily, transforming it into a dynamic pool of resources for on-demand cloud and DCI applications.

## Summary

The world's Internet population is estimated at more than 3.7 billion. We now live in a global culture that demands content that's consumable on demand, anywhere, at any time, with the highest levels of quality. Losing access to a data center can be disastrous to revenue, resulting in hundreds of millions of dollars lost in online transactions or thousands of flights grounded. Interconnecting these data centers is key to enabling applications to keep business running smoothly and delivering content to keep end-users happy.

Advancements in DCI hardware and software have set new benchmarks for operational simplicity, scalability, and virtualization to accelerate the pace of deployment and time to market. Network operators can reduce operational costs and improve IT flexibility and efficiency with highly programmable platforms that leverage open APIs. In addition, openness and programmability speed and simplify integration with existing back-office tools and applications. DCI networks are leveraging these new technologies to evolve into networks that can easily adapt as well as deliver scale along with automation and ease of deployment to facilitate DCI modernization without risk.

