WHITE PAPER

New Economics of Datacenter Interconnect

Sponsored by: Cisco Systems

Nav Chander
February 2016

EXECUTIVE SUMMARY

IDC completed a total cost of ownership (TCO) analysis that compares an existing datacenter interconnect (DCI) solution with a new hyperscale DCI solution from Cisco that scales to accommodate the rapidly growing wide area network (WAN) requirements of hyperscale datacenters. This solution addresses the next-generation communications service provider (CSP) and hyperscale WAN requirements in datacenters for modular routing software, high-density 100G ports for optical interconnect, new software extensibility, and analytical tools and automation that enable CSP and Web-scale companies to customize their services. IDC findings confirm that the Cisco DCI hyperscale solution yields:

- 55% improved TCO through next-generation datacenter and WAN infrastructure
- 54% increased operational efficiencies through automation
- 70% link utilization through telemetry, data analytics, and application-engineered routing
- 49% transport capex reduction from DC and WAN convergence (IP and optical)
- 24% reduction in power consumption

Tier 1 carriers and hyperscale cloud service providers have an insatiable need for highly scalable and automated optical datacenter interconnect solutions. This is required to cope with the expected tripling in global IP data traffic from 2015 to 2019. By 2019, more than 65% of all IP traffic will remain in metro networks due to increased DCI, cloud services, and content delivery applications within the metro area. The increasing number and type of data applications across datacenters will also require significant automation to allow DCI operations to scale.

OVERVIEW OF HYPERSCALE CLOUD DATACENTER INTERCONNECT INFRASTRUCTURE

IDC believes that datacenter interconnect infrastructure and associated DCI services represent one of the most significant growth areas for network equipment vendors, CSPs, and Internet content providers (ICPs) alike. A combination of industry trends — digitization, cloud computing, content distribution, virtualization, and the explosion in video traffic — is making datacenter interconnect a critical piece of the network strategy for CSPs, ICPs, and vendors for the following reasons:

- Companies, such as Google, Microsoft, and Facebook, can acquire their own optical networking infrastructure instead of leasing datacenter interconnect services from CSPs.
- There is growing bandwidth demand for links between datacenters and interconnection points.

- A 31% projected CAGR for inter-DC traffic from 2015 to 2020 will stress existing DCI infrastructure.
- A 2014 IDC survey of 200 enterprise CIOs indicates that over 70% of these companies have virtualized more than 50% of their servers and plan to increase the use of virtualized IT servers and services, such as cloud compute and storage, over a two-year period.

Connectivity between datacenters is becoming a critical enterprise asset as organizations seek to scale infrastructure and push data and application processing closer to the edge. This is leading to a wealth of investment around the datacenter interconnect, particularly in new optical platforms that strive to recreate local data performance over wide area infrastructure.

Datacenter interconnect equipment helps connect one datacenter to another geographically separate datacenter in a metro or wide area network. Datacenter interconnect networking equipment (routers, Ethernet switches, and optical switches) has a direct effect on how content is delivered. An example is streaming video. The best way to ensure video will stream in high definition is to provide it from the datacenter nearest to the viewer's location. However, if the specific video content doesn't exist in that datacenter, it has to stream from another location through a datacenter interconnect.

The datacenter interconnect may link an organization's own datacenter to a third-party datacenter that is being used for disaster recovery or other services. In most cases, content providers lease the fiber-optic connectivity between datacenters. However, in recent years, companies such as Microsoft, Facebook, and Google began investing in their own long-haul fiber-optic networks. By controlling the fiber itself, a company can respond more rapidly to changes in traffic than when it is leasing capacity.

The increasing use of direct connections to cloud exchanges by enterprises is indirectly impacting CSPs, as more cloud-based technologies become part of enterprises' solutions. IDC expects demand for 100G DCI capacity from cloud-scale providers to grow 60-80% per year for the next five years based on the evolving use of cloud, big data, and software as a service (SaaS).

The overview discusses some important considerations for the requirements of a next-generation DCI infrastructure.

- Labor cost of operating a DCI infrastructure
- Power consumption
- Network utilization
- Capex requirements for 100G infrastructure
- Service automation and troubleshooting

Operations

DCI applications require a high density of high-capacity physical ports, and the inherent desire to minimize downtime risk is extremely costly. There is an expectation that the specialized skills and experience required to operate these complex DCI networks are expensive and scarce. Both communications service providers and cloud service providers must continuously upgrade training and technical resources to keep up with the evolving network virtualization and optical and routing technology advances of vendor DCI solutions. This can slow down the implementation of network changes to increase scaling and to upgrade switches and optical networking equipment and routers. The lack of automation and tools to improve upgrade cycles and capacity increases impacts opex. Figure 1 depicts the importance of automation for provisioning new devices and software in large-scale DCI networks.

FIGURE 1

Automating Provisioning

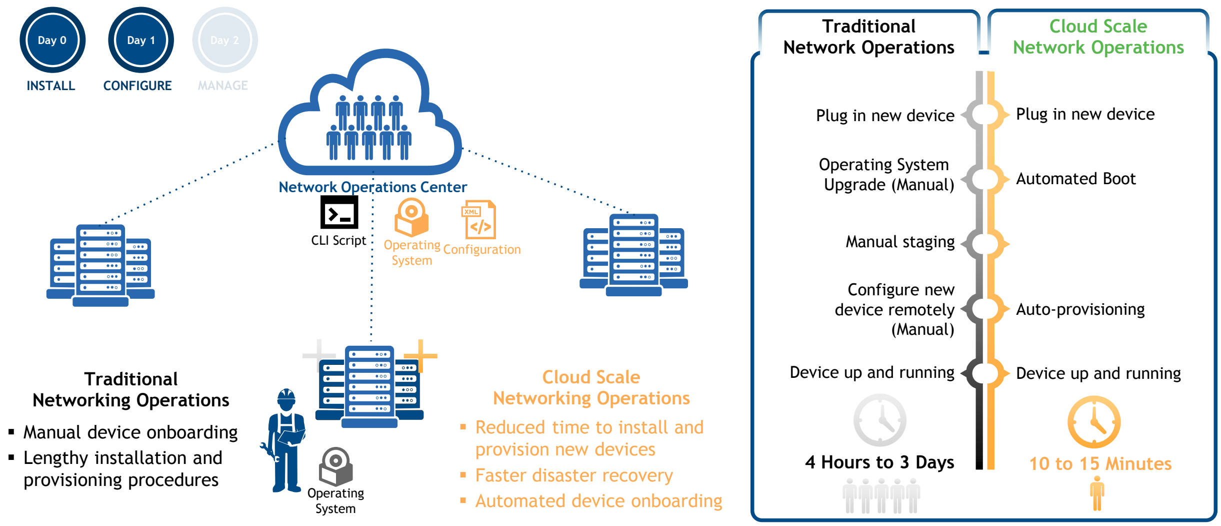

Source: Cisco, 2016

Network Scaling

Rapid, dynamic growth of cloud-scale networking infrastructure is driving increased use of automation software and Web-based tools to help provision and scale networks for expanding DCI use cases. In a cloud-scale environment where there may be 4 million lines of configuration files, there can be over 30,000 configuration changes a month, with 20,000 command-line interface (CLI) commands issued and deleted every 5 minutes. CLI scripts are complex to develop, and changes in a vendor's operating system (OS) can impact them. Employing platforms from multiple vendors compounds this complexity, leading to longer cycle times for provisioning capacity changes. Such delays are not acceptable in an on-demand environment. The traditional telecom networking environment is quite different. It has more predictable and static network configurations, and bandwidth changes have 30- to 90-day provisioning cycles for most enterprise and wholesale networking services. To remain competitive, CSPs will have to transform their networks to become more agile and be able to support a new generation of flexible cloud connectivity and on-demand services.

For hyperscale and very large CSP environments, new data models, such as YANG, can help reduce the complexity of CLI scripts. Combined with automated provisioning and auto-boot capabilities, they can reduce the provisioning of a new device, a service, or even an upgrade to 10–15 minutes compared with 3–5 hours.

Network Utilization

Cloud-scale networks must operate with utilization significantly higher than the typical 40–45% that many CSPs experience today because upgrading capacity to keep up with 60–70% DCI bandwidth growth cannot be met simply by spending equivalent capex. However, more intelligent visibility and control are necessary to increase utilization to 70% and higher. An integrated suite of features, such as segment routing, telemetry, and SD-WAN policies, can automatically reroute or balance traffic flows across the network in near real time to boost network utilization to 75% or higher. Figure 2 depicts how

Cisco integrates the needed capabilities into its DCI solution. Telemetry-enabled network devices that leverage YANG data models can help communications service providers and cloud service providers dynamically match the requirements of application traffic flows to network capacity and topology using software that automatically optimizes the network paths they use.

FIGURE 2

Fine-Grained Visibility and Control Optimize Utilization

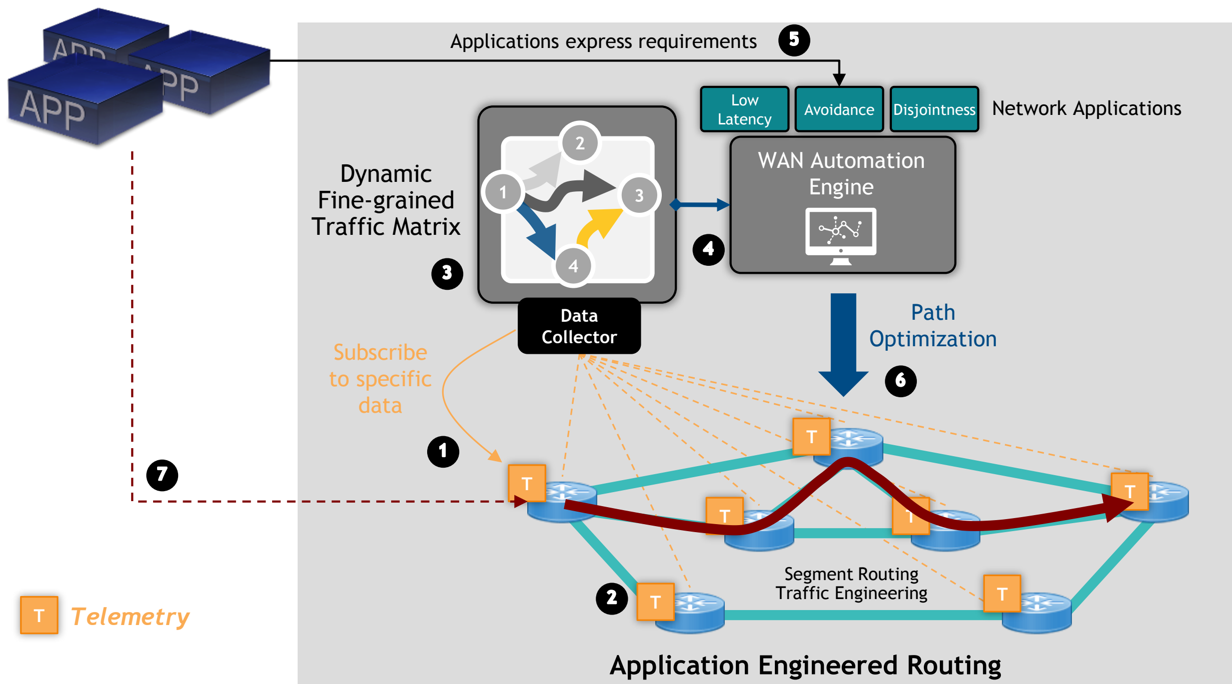

Source: Cisco, 2016

Service Agility

One of the biggest challenges for communications service providers and cloud service providers has been the ability to rapidly develop, integrate, test, and deploy new services and new capabilities within existing services. Many cloud service providers have resorted to building their own operational and network management tools to overcome the limitations of their networking vendors' products. These vendors provide monolithic software releases that require communications service providers and cloud service providers to conduct extensive manual qualification testing, lasting 6–9 months, simply because interdependencies in the software mean that new features can impact functions almost anywhere in the existing software. Any additional maintenance software fixes have to be retested, delaying the whole service life-cycle release, which can take up to 12 months.

Tier 1 carriers and hyperscale cloud service providers need solutions that enable them to implement only the features they select and cut software release testing by at least 50%. Such solutions need to offer granular packaging of modular software components, programmability of the platform, and the use of a Linux host OS. This capability enables CSPs to innovate and experiment with new service releases by providing them the extensibility to develop their own features and incorporate other third-party software innovations at a rapid pace, independent from the equipment vendors' major release cycles.

Power Consumption

High-capacity DCI applications require a significant density of 100G ports, and lowering the cost of space and power requirements has to be an important consideration for next-generation deployments. For routing and optical ports, newer technology can significantly reduce power consumption. Newer 100G coherent optical port designs can consume just 0.8W/Gbps compared with current solutions using 1.7W/Gbps for standalone DWDM.

DCI ECONOMICS — BUSINESS CASE ANALYSIS

An economic comparison of current (2015–2016) DCI capacity requirements for 100G DCI and DCI capacity requirements for 2016–2019 examines a typical cloud-scale provider use case with capacity increasing at a 180% CAGR. It compares Present Mode of Operation (PMO) and Future Mode of Operation (FMO) within a modern hyperscale datacenter infrastructure.

Table 1 shows the networking equipment in use for both alternatives. Figure 3 and Figure 4 show the PMO and FMO comparison of equipment focused on the inter-DC requirements.

The PMO solution is made up of a 40G top-of-rack (ToR) architecture with tier 1, tier 2, and tier 3 datacenter switches and optical DCI equipment. IDC focused on the tier 3 switch and optical networking costs for the DCI, as shown in Figure 3, supporting 500 40G routed ports. The total optical DCI bandwidth capacity is 1,725 100G ports.

TABLE 1

DCI Equipment

Equipment Category	PMO	FMO
Spine	40G Switch	NCS 5508
Edge Gateway	Core router (40G and 100G interfaces)	
Optical DCI	Standard LH and Metro DWDM products	NCS 1002

Source: IDC, 2016

The FMO solution is an alternative scalable 100G ToR architecture and hyperscale DCI solution consisting of the NCS 5508 high-density 100GE routing platform. It also employs the compact NCS 1002 optical 2 RU system that supports up to 2Tbps of client and 2Tbps of trunk traffic with advanced 100G coherent optics hardware and software. The FMO solution supports 25G server configurations and scales to support 512 routed 100G ports and 865 100G DWDM ports. Figure 4 depicts the FMO alternative.

FIGURE 3

PMO DCI Hyperscale Configuration with 40G Fabric

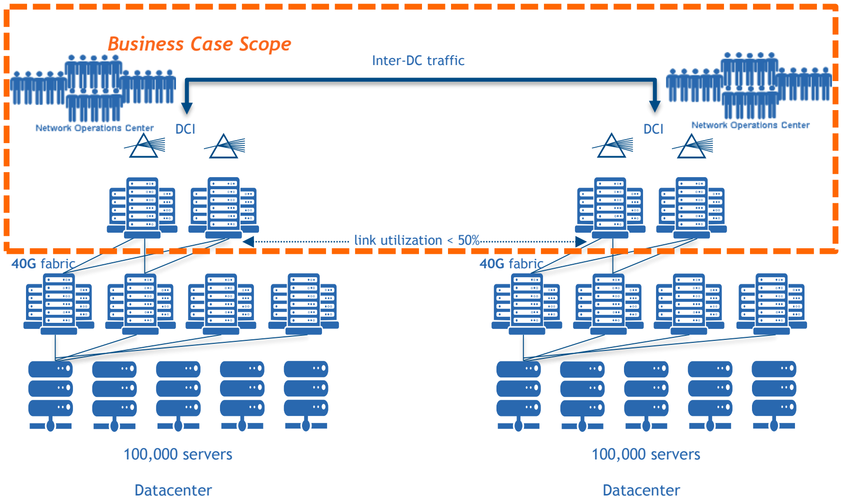

Present Mode of Operation

- Limited automation
- Manual onboarding of new devices
- Multiple network outages (due to human configuration errors)

Source: IDC, 2016

FIGURE 4

FMO DCI Hyperscale Configuration with 100G Fabric

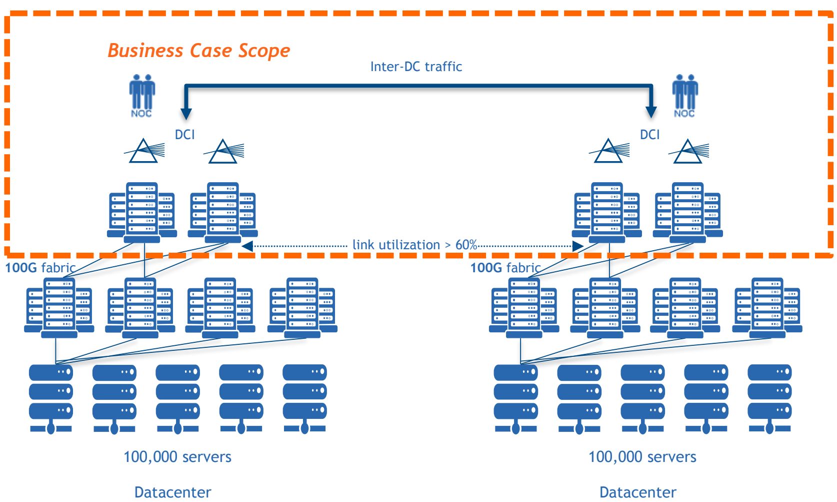

Future Mode of Operation

- High-scale automation
- Automated onboarding of new devices
- Scarce network outages (automation benefits)

Source: IDC, 2016

Operational

To compare the operational efficiency of the PMO and the FMO, IDC examined key attributes such as costs for labor to provision devices, downtime based on the outages, and IT administration staff ratios for maintaining network infrastructure. Table 2 provides those assumptions for each alternative.

The operational cost of the FMO was 65% lower than the operational cost of the PMO, saving $2.3 million per year. The FMO alternative leverages the software automation capabilities and WAN automation policies shown previously (refer back to Figure 1 and Figure 2), significantly reducing IT staff costs.

TABLE 2

Operational Attributes

IT admin to device ratio	
Present Mode of Operation	30
Future Mode of Operation	200
Time to install and provision devices (in minutes)	
Present Mode of Operation	360
Future Mode of Operation	15
Gains from cost reduction of application failures resulting from increased quality (GQL)*	
Average minutes to recover difference	28.3
Revenue per minute	$7,908.00
Gain from reducing IT operations headcount waste*	
Average FTE salary - NOC operations	$120,000.00
Savings of 7.2 hours per week out of 45 total hours per week	16%

**How to Calculate the ROI of Continuous Delivery*, Zend report, 2014

Source: IDC, 2016

Transport Cost

The cost of DCI WAN connections is also a significant part of the TCO analysis. The FMO solution leverages integrated 100G coherent optics and high-density 100G routing interfaces on the NCS 5508 compared with the PMO solution, which uses a 40G switching architecture and separate optical switches. Comparing the PMO and the FMO alternatives reveals that the total bandwidth capacity of optical connections increases from 172,000Gbps to 409,000Gbps and the WAN switching capacity increases from 262,400Gbps to 614,400Gbps. Even with these increases, the overall capex per gigabits per second decreases by 49%, as shown in Figure 5. IDC assumed a link utilization of 40% for the PMO alternative and 60% for the FMO alternative. Greater TCO savings may be achieved with link utilizations that exceed 60% and approach 70–75% for the FMO, which is feasible with the Cisco equipment used in the IDC analysis.

FIGURE 5

PMO Versus FMO Transport Cost Comparison

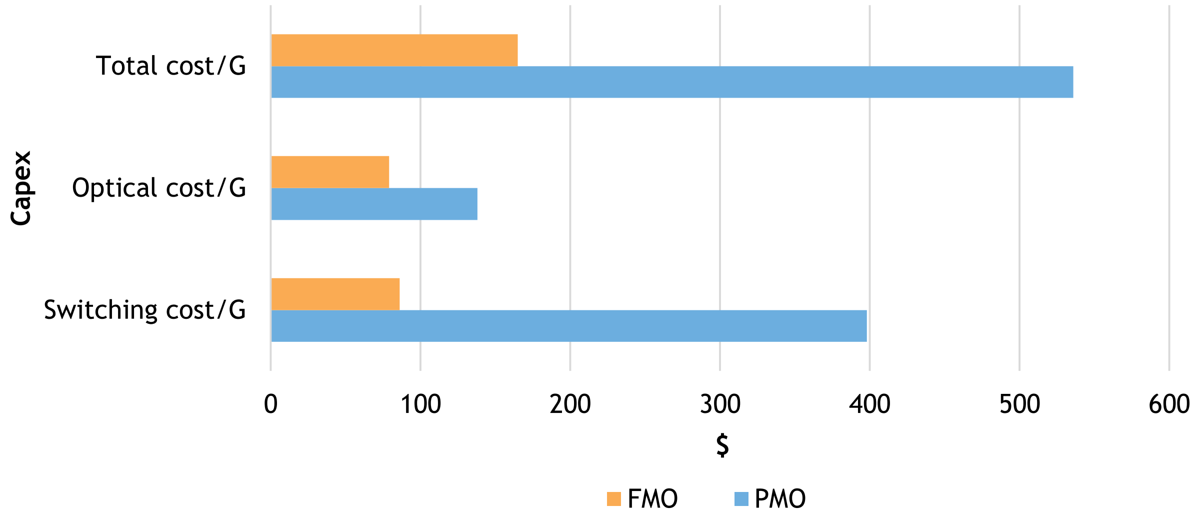

Source: IDC, 2016

Power Consumption

Lower power consumption per gigabits per second is also a key area of focus for IT operations, particularly for hyperscale DCI applications where many high-density optical ports are in close proximity to each other and power consumption can be as much as 1.6–1.8W/Gbps for high-capacity ports. For the total switching (tier 2, tier 3, and DCR) and optical capacity of the PMO alternative, an average of 7W/Gbps is consumed.

For the FMO alternative, which has approximately five times the available bandwidth capacity of the PMO, the optical port power consumption with new coherent optics and compact form factors is reduced to 0.8–1.25W/Gbps. The total switching (tier 2 and tier 3) and optical capacity for the FMO alternative consumes an average of 2.5W/Gbps.

Three-Year TCO Analysis

Figure 6 and Figure 7 show the three-year comparison of capex and opex for the PMO and FMO options, with a 55% improvement in TCO/G using the FMO. The analysis considers additional factors, such as datacenter outage costs and significant reduction in IT admin costs from 11% of the TCO to 1% in the FMO alternative. Capex includes the cost of procuring the optical, switching, and routing hardware and software and assumes a five-year amortization for the networking equipment, averaged over the three-year TCO period. Opex assumptions are quite conservative and assume a flat IT labor rate for the PMO and FMO alternatives and also flat power rates, both of which are likely to increase over time. IDC estimates that these two factors can combine to further enhance TCO savings by 1–2%.

FIGURE 6

Three-Year PMO TCO

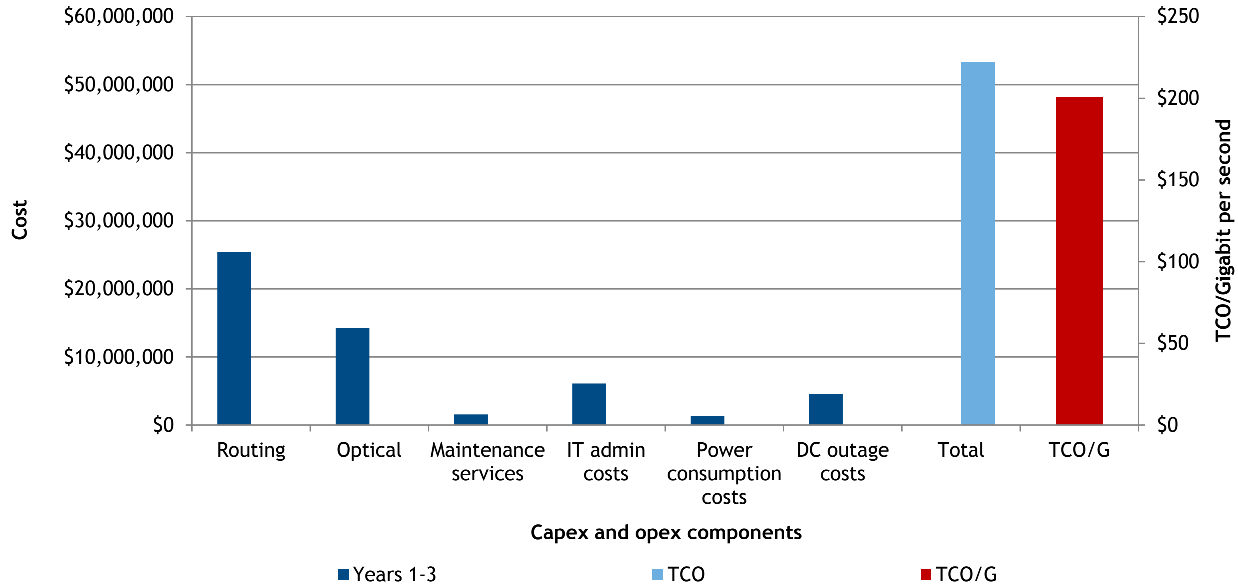

Source: IDC, 2016

FIGURE 7

Three-Year FMO TCO

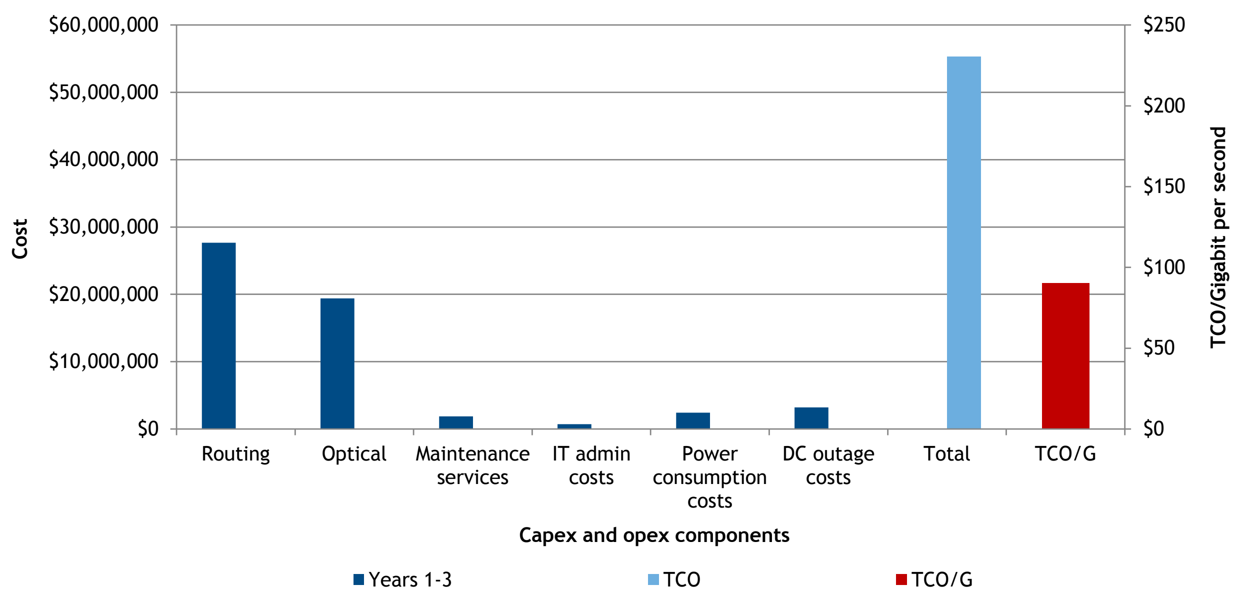

Source: IDC, 2016

CONCLUSION

IDC concludes that the new hyperscale DCI solution from Cisco clearly demonstrates a 55% TCO improvement compared with existing DCI solutions and that this new solution will scale to accommodate the rapidly growing WAN requirements of the hyperscale datacenter that will be required to support the exponential growth of cloud-scale networking. The integration of high-speed optical networking, routing, and switching technologies into a DCI platform addresses the next-generation CSP and hyperscale WAN requirements expected during the next three to five years in datacenters. The analysis also clearly shows the economic and competitive advantages of an open, programmable platform with modular routing software, as it offers hyperscale service providers the flexibility to customize their environments. CSP and Web-scale companies gain the agility and economic efficiency of high-density 100G ports for optical interconnect, new software extensibility, analytical tools, and real-time automation tools to adapt their DCI services more dynamically to changing market demand.

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

Global Headquarters

5 Speen Street
Framingham, MA 01701
USA
508.872.8200
Twitter: @IDC
idc-community.com
www.idc.com