White Paper

The Need for OTN in Data Center Interconnect (DCI) Transport

July, 2016

Power Matters.™

Abstract

Explosive bandwidth growth from on-demand cloud services is forcing network operators to rethink the way in which data centers are interconnected globally with optical transport networks. These new data center interconnect, or DCI, networks must meet these massive capacity needs and scale efficiently, in power and cost, without compromising performance and manageability of the network infrastructure. OTN-based optical transport technology is purpose-built to scale DCI network capacity in the 100G Cloud Era, while delivering the best blend of features, deterministic performance and transport-grade Operations, Administration and Management (OA&M). This whitepaper presents a case for OTN-based DCI and discusses how Microsemi's innovation and leadership in OTN processing silicon raises the bar for OTN-based DCI transport networks.

Contents

The Rise of Data Center Interconnect (DCI) Transport

Network operators globally are struggling to cope with two evolutionary problems: 1) the exponential growth in network traffic, and 2) the insatiable demand of consumers and enterprises for on-demand, anytime-anywhere, cloud-based services. Netflix data is a proxy for this issue—monthly streaming hours grew by 1000x from 2008 to 2015[1]. These services and new cloud-enabled business paradigms are fundamentally changing the way in which people and enterprises access and store information in the cloud, and as a result are forcing dramatic changes in the underlying networks that support these services.

Traditional North-South traffic, also known as "Machine-to-User" traffic, is large and growing. However, East-West traffic flows, or "Machine-to-Machine" traffic, are several orders of magnitude larger—driven by the need to cache content closer to end users, to distribute compute and storage resources in support of cloud-based workloads and to meet the requirements for redundancy and global reach expected from cloud services. It is estimated that by 2019, machine-to-machine traffic will represent approximately 85% of total data center traffic, or 8,500 exabytes per year[2]. To support these massive machine-to-machine flows, often distributed across many geographically dispersed data centers, high-bandwidth data center interconnect network connectivity is required. Facebook reported that inter-data center traffic represents 18% of total network traffic, while it is estimated that global inter-data center traffic flows will grow at a rate of 31% annually through 2019 to almost 1,000 exabytes and almost 10% of global data center traffic.[3,4] To service this massive bandwidth need, a new breed of optical network is emerging.

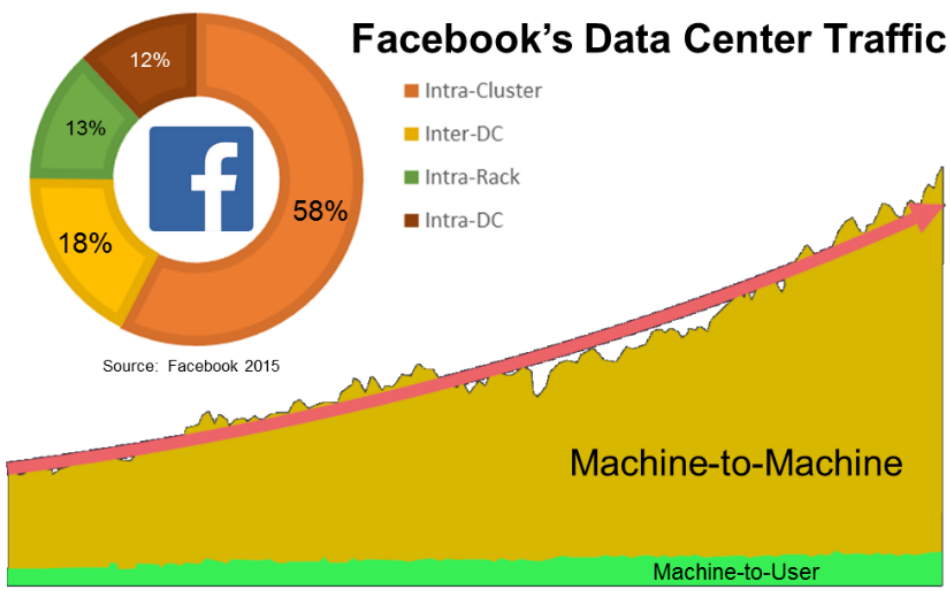

Requirements For DCI Optical Transport Networks

Interconnecting today's data centers offers new challenges that drive unique requirements compared to traditional optical transport networks.

DCI Challenge	DCI Network Requirement
Ultra High-Bandwidth Reach: 10-1000s of Kilometers	□ Dense 100G DWDM optical links □ Coherent Transmission, optimized for Metro & Long-Haul
Diverse Client Types & Rates on Servers, Storage, Switches	□ Flexible Rates, Efficient Aggregation from 10G to 100G □ Multi-Service: Ethernet, Fiber-Channel, etc…
Real-Time Access Mission Critical Apps	□ Low Latency transport networking protocol □ Robust Fault Isolation □ Aligned with traditional transport OA&M
Data Security Cyber-crime, Regulatory requirements	□ Encryption of data 'in-flight' between data centers
SDN Management & Automation Efficient, Dynamic, Programmable	□ SDN-aware Transport: Open APIs, Topology Discovery
Diverse geographic placement of DCs	□ Flexible Deployment: Leased Line and Dark Fiber
High real-estate & cooling costs	□ Low power per bit transport □ High-bandwidth per rack-unit

Multiple transport technology choices are available to address the unique needs of DCI networks. OTN-based DCI fulfills the key requirements above, while leveraging the extensive expertise and heritage in optical transport platform development, and network operations and management that exist today at OEMs and network operators.

The Value of OTN-Based DCI

Interconnecting data centers with OTN-based transport meets the key requirements for DCI while offering several compelling advantages when compared to alternative solutions.

The Need for Efficient, Multi-Service OTN Aggregation

OTN is the de-facto networking protocol for next-generation communications service provider 100G Metro and Core optical transport networks worldwide. The industry uses the term OTN loosely and there is a misperception among some market verticals that OTN only implies OTN switching. 100G OTN networking can be divided into three important functions:

1. Transponding, or wrapping of 100Gbps client signals with OTN over point-to-point 100G Wavelength Division Multiplexing (WDM) links
2. Muxponding, or multiplexing of multiple lower-rate client signals into higher-rate 100G links
3. Switching, grooming or cross-connection of lower-rate signals from node to node in an optical network

OTN muxponding and transponding have been the underpinning for point-to-point ROADM/WDM networks for several generations now. These Metro and Core networks are now transitioning to mesh-based OTN switching architectures, driven by the need to more efficiently pack 100G fiber infrastructure to handle the growing bandwidth problem across all service types.

Transponding of 100GE clients onto 100G wavelengths using OTN is a natural fit for the needs of DCI applications, as OTN is the de-facto standard framing format in use in all Digital Signal Processors (DSPs) required for 100G+ Coherent Optical Transmission. OTN provides a lightweight wrapping for the underlying 100GE client signal, in addition to a rich set of OA&M features to rapidly detect and isolate faults, and manage the performance of optical links.

However, the client switch or router ports facing DCI networks are not limited to 100 Gbps rates, nor is the client networking protocol limited to Ethernet. Different classes of data centers drive different requirements. Large hyperscale data centers operated by the big cloud service and internet content providers (Google, Amazon, and Microsoft, to name a few) have migrated most of the optical transport facing ports originating from top-of-rack data center switches from SFP+ to QSFP-capable ports, compatible with either 4x10GE / 1x40GE-capable QSFP+ or 1x100GE-capable QSFP28 pluggable optical modules. Given that QSFP28 modules are still in the early stages of availability and deployment, 10GE and 40GE clients will continue to play a big role feeding DCI optical transport platforms.

While there's a great deal of focus on the needs of hyperscale data centers, there is also a sizeable and growing market of smaller data center sites that still require high-bandwidth optical connectivity. These data centers are more numerous and are distributed throughout metro regions, servicing high-value enterprise verticals, including financial, healthcare and government, and the co-location, internet exchange or carrier neutral provider markets. These sites have a wide range of requirements depending on the function and type of service offered. In these environments, 10G SFP+ ports still represent a large percentage of the ports from top-of-rack switches to DCI transport. Additionally, other client types, such as Fibre Channel, are commonly required by the financial and healthcare industries, where business continuity applications such as synchronous data replication are critical.

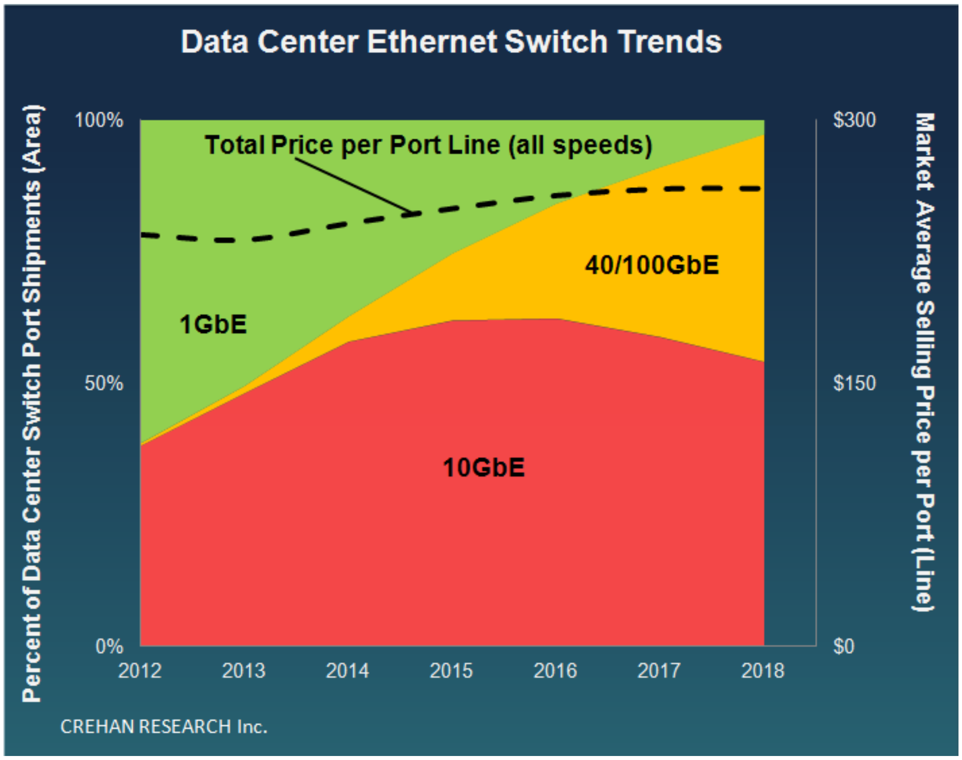

In both cases, OTN muxponding and switching are the only solutions that are purpose-built to enable multiplexing of all client protocols and rates onto 100G wavelengths. Additionally, OTN multiplexing delivers maximum network efficiency and lowest latency as it avoids throughput reduction due to the additional overhead needed in alternative Layer 2-based implementations.

Security in OTN Transport Networks

The rapid growth of cybercrime and the well-documented costs to the enterprise world associated with addressing data breaches has brought the security of data to the forefront. The privacy of data is most exposed as it transits between data centers, where extensive physical security measures at these sites are no longer effective. It is no surprise that "in-flight" network layer encryption is quickly becoming a key requirement, and in some cases mandatory, for DCI optical transport.

Encrypting data at the OTN layer is the natural solution to this problem as this secures traffic closest to the exit of the data center. OTN encryption delivers strong AES-256-bit encryption without impacting link latency or available optical fiber bandwidth regardless of client type or rate. Alternative encryption approaches are available, however these solutions are limited to Ethernet or IP traffic and they reduce the available optical fiber bandwidth by greater than 20% when compared to OTN encryption[5]. Leveraging OTN encryption on DCI optical transport equipment also eliminates the unnecessary cost and power burden of adding encryption functionality to higher volume DC top-of-rack switch Ethernet ports.

OTN transport networks are also inherently secure by design when compared to alternative transport technologies, adding an additional layer of security to mission critical Dense Wavelength Division Multiplexing (DWDM) networks. Links in an OTN transport network are hard-partitioned into dedicated circuits. This segregation of network traffic makes it difficult to intercept data transferred between nodes—this ensures that no one client can compromise other clients on a link.

Best-In-Class Performance, Reliability and OA&M Capabilities

Low-cost and power-per-bit transport cannot come at the expense of the performance, reliability and manageability of DCI optical links. Supporting distributed, scalable, cloud-based mission-critical applications with global reach means that data center network operators must have the ability to:

- Deliver deterministic performance
- Provision services simply and transparently
- Reliably manage real-time status of, and troubleshoot and isolate faults in, DCI optical links

OTN delivers consistent full line-rate, end-to-end performance, regardless of client or service type. Each client is assigned a specific and configurable bandwidth depending on the rate or protocol, ensuring guaranteed jitter, throughput, latency and availability through the network. This eliminates contention or packet drops that may exist in alternative solutions. Provisioning services in an OTN-based DCI network is simple and transparent, each client is network-timing and rate-independent over the same 100G wavelength, eliminating the need for unnecessary operational complexities such as managing forwarding tables, scheduling, and packet buffers.

Designed explicitly to ensure carrier-grade QoS and manageability, OTN-based DCI provides superior OA&M capabilities. OTN offers full performance monitoring and MAC statistics on a per client basis as well as full line-side monitoring features (section, path and tandem connection monitoring), enabling end-to-end status monitoring, fault isolation, and network troubleshooting down to the client level.

Additionally, OTN-based networks are resistant to Layer 2 or Layer 3-based denial-of-service attacks that are becoming more prevalent. Decoupling the DCI transport network from vulnerabilities at these layers ensures maximum reliability and performance for mission-critical applications.

Flexible Scaling of DCI Network Capacity

The availability of existing network infrastructure is a key consideration when selecting the right technology for DCI applications. Greenfield networks leveraging privately-owned dark fiber offer maximum flexibility in the choice of DCI transport technology. However, dark fiber is not universally available in all geographies. The trend to distribute data centers, and therefore content and services, closer to the user further exacerbates this challenge. Leasing either a Wavelength or Ethernet transport service, from a traditional Communications Service Provider, may be the only cost-effective option available to deliver connectivity between sites for a particular geography. As such, a hybrid approach, leveraging both privately-built dark-fiber networks and leased transport services, is a common choice today for most cloud service and internet content providers.

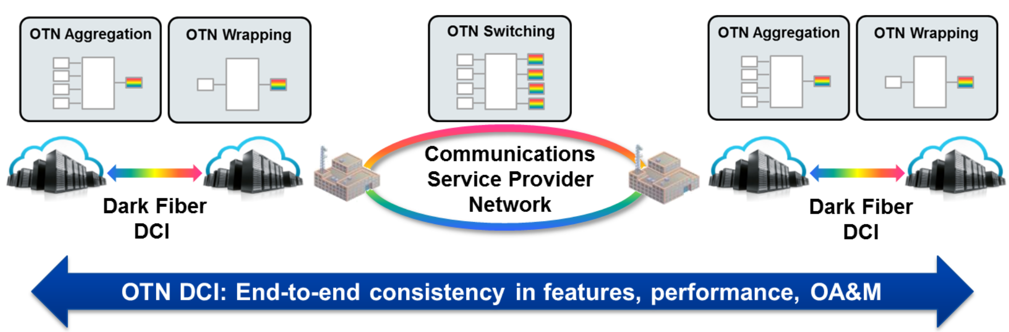

Using this approach, OTN-based DCI provides data center operators with:

- The flexibility to scale site-to-site connectivity based on the most cost-effective service available for a particular geography, and
- Consistent features, performance and OA&M capabilities in all DCI network segments, lowering network operations complexity and therefore overall OPEX spending

DCI Bandwidth-On-Demand

OTN-based DCI also offers the novel approach to dynamically and cost-effectively scale DCI network capacity. Bandwidth-intensive tasks, such as caching or database mirroring, are generally not time sensitive and therefore can be deferred to off-peak, overnight hours when network capacity and utility costs are lowest. Utilizing G.HAO dynamic and hitless resizing of packet-based services capabilities unique to OTN, leased line transport services can scale available bandwidth based on these cyclical cost patterns. This is a win-win for the entire supply chain—data center operators can leverage the extensive footprint of Metro and Core OTN switch-based Communication Service Provider networks worldwide to obtain additional capacity when they need it at the lowest cost per bit, Communications Service Providers are presented with a new means to monetize those extensive investments, and OEMs achieve economies of scale in using a common technology across DCI and traditional telecommunications networks, further supporting healthy industry innovation.

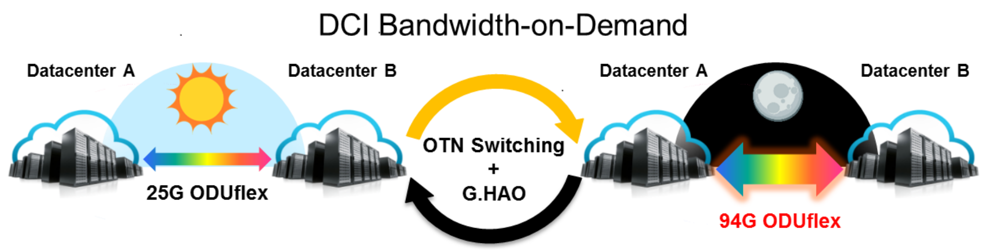

Microsemi's DIGI-G4: Best-of-Breed DCI Transport

Building upon Microsemi's OTN processing silicon innovations and leadership from DIGI-120G, the DIGI-G4 400G OTN Processor raises the bar for OTN-based DCI — meeting all key requirements and bringing new purpose-built, industry-first innovations to market that make high-capacity, scalable, flexible, SDN-ready, encrypted DCI transport a reality.

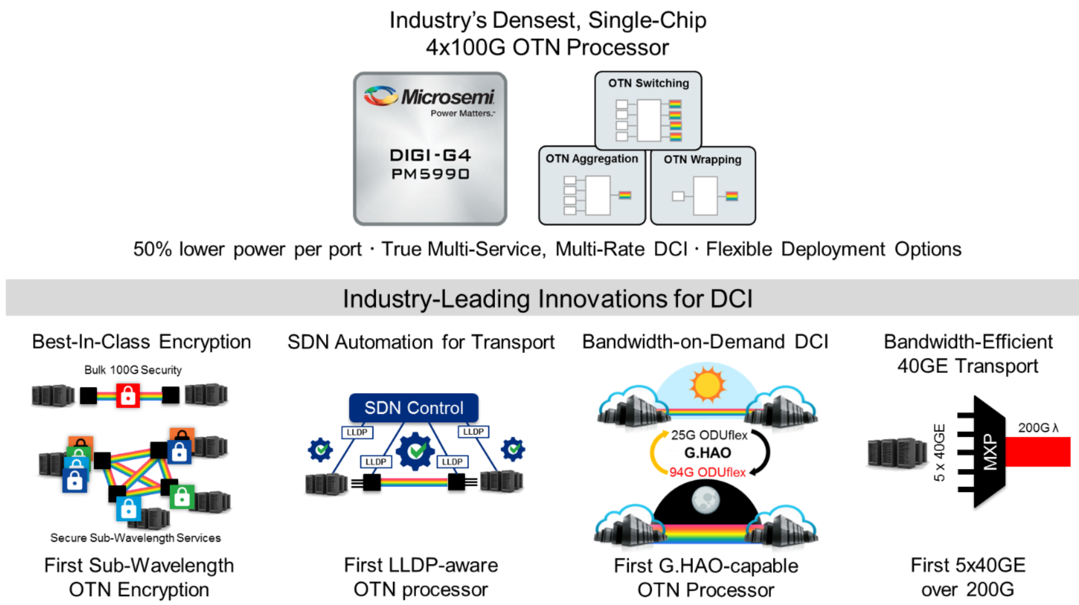

The DIGI-G4 is the industry's densest single-chip 4x100G OTN processor delivering true multi-service, multi-rate OTN transponding, muxponding, and switching capabilities at 50 percent less power per port than the previous generation. This makes low power per bit and high-capacity per rack-unit transport possible without trading off the required flexibility to support all necessary traffic types, rates and network deployment options.

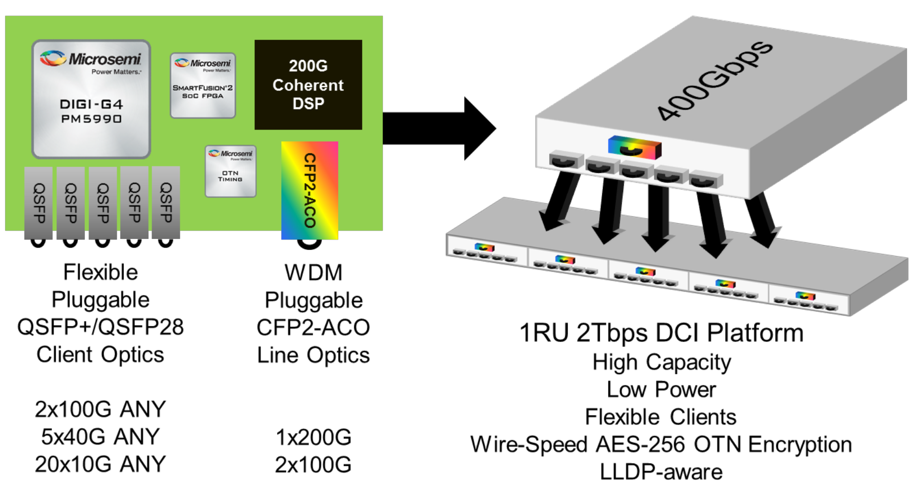

The DIGI-G4 addresses the industry need to secure DCI optical networks with the integrated OTN encryption capabilities. Outperforming protocol-specific alternatives, the DIGI-G4's OTN encryption engine delivers ultra-low latency (sub-180ns), strong standards-based (industry's first FIPS 197 AES-256 certified solution) 'in-flight' 100G optical encryption, regardless of client or traffic type without compromising network performance and service quality. The DIGI-G4 also brings to market the industry's first sub-wavelength optical encryption solution, allowing traffic to be secured end-to-end at "service-layer" granularity. These secure services are compatible with, and can be routed efficiently through, 100G Metro OTN Switched networks without the need to decrypt and disaggregate "bulk" encrypted 100G links at every node in the network, thereby exposing sensitive data in transit. This provides an option when dark fiber is scarce to leverage "pay-as-you-grow" encrypted leased line DCI services that deliver the required bandwidth with deterministic performance and guaranteed end-to-end security. The combination of DIGI-G4 with Microsemi's SmartFusion2 SoC FPGA completes the solution for encryption-enabled DCI transport platforms. SmartFusion2's integrated secure host central processing unit (CPU) supports Public Key Infrastructure (PKI)-based authentication architectures, provides the secure key storage required to set up and manage end-to-end encrypted optical links, delivers secure boot functionality for existing non-secure host CPUs and protects against the threat of side-channel attacks with differential power analysis (DPA) countermeasures.

The DIGI-G4 is the industry's first OTN solution to extend SDN awareness into optical transport. A key vision of SDN is to enable programmatic control physical network resources. Supporting monitoring of the Link-Layer Discovery Protocol (LLDP) on client ports, the DIGI-G4 allows SDN control infrastructure to maintain a real-time view of the connectivity and topology of the underlying DCI transport network resources. This eliminates the need to manually maintain static connectivity tables between DC Switch/Router and optical transport equipment, a costly and often error-prone exercise.

The DIGI-G4 is the first industry solution to deliver bandwidth-efficient transport of 5x40GE clients over 200G of optical capacity. The standard approach to this problem wastes 20% of the available optical bandwidth, supporting a maximum of 4x40G of client ports. The DIGI-G4's innovative solution to 40G aggregation, ensures no wasted bandwidth/optical spectrum and in doing so allows for an additional port 40GE client traffic to be supported on each line card deployed. This is of significant value where optical fiber capacity/spectrum is scarce and highest bandwidth is a priority.

Lastly, Microsemi's DIGI family of OTN processors is the industry's first and only solution in production today capable of supporting G.HAO in OTN transport networks, enabling the industry to leverage investments in OTN switch-based Metro and Core networks to offer bandwidth-on-demand DCI transport services.

Conclusion

Today’s new breed of cloud services demand dedicated, high-bandwidth, secure optical data center-to-data center connectivity. To address this need, DCI optical networks must scale efficiently, in bandwidth, power and cost, without compromising performance or manageability. OTN as a networking protocol is purpose-built to meet these requirements: delivering low latency, bandwidth-efficient, rate and protocol-agnostic 100G optical transport with best-in-class "in-flight" encryption capabilities, and new and compelling use-cases for the industry to collectively scale DCI bandwidth on-demand in a cost-effective way.

Microsemi’s DIGI-G4 OTN processor raises the bar for OTN-based DCI networks – delivering industry-first innovations purpose-built to ensure that high-capacity, scalable, flexible, SDN-aware, encrypted DCI transport platforms become a reality.

References

1. *Completing the Netflix Cloud Migration*, Netflix, February 2016.
   https://media.netflix.com/en/company-blog/completing-the-netflix-cloud-migration
2. *Cisco Global Cloud Index: Forecast and Methodology*, 2014–2019", Cisco, April 2016.
3. Ibid.
4. A. Roy et al. "Inside the Social Network's (Datacenter) Network", SIGCOMM 2015, August 2015.
5. *'In-flight' Encryption in Service Provider Networks*, Microsemi Corporation, 2015.

Power Matters.™

Microsemi Corporate Headquarters
One Enterprise, Aliso Viejo,
CA 92656 USA

Within the USA: +1 (800) 713-4113
Outside the USA: +1 (949) 380-6100
Fax: +1 (949) 215-4996
Email: sales.support@microsemi.com
www.microsemi.com

© 2016 Microsemi Corporation. All rights reserved. Microsemi and the Microsemi logo are trademarks of Microsemi Corporation. All other trademarks and service marks are the property of their respective owners.

Microsemi makes no warranty, representation, or guarantee regarding the information contained herein or the suitability of its products and services for any particular purpose, nor does Microsemi assume any liability whatsoever arising out of the application or use of any product or circuit. The products sold hereunder and any other products sold by Microsemi have been subject to limited testing and should not be used in conjunction with mission-critical equipment or applications. Any performance specifications are believed to be reliable but are not verified, and Buyer must conduct and complete all performance and other testing of the products, alone and together with, or installed in, any end-products. Buyer shall not rely on any data and performance specifications or parameters provided by Microsemi. It is the Buyer's responsibility to independently determine suitability of any products and to test and verify the same. The information provided by Microsemi hereunder is provided "as is, where is" and with all faults, and the entire risk associated with such information is entirely with the Buyer. Microsemi does not grant, explicitly or implicitly, to any party any patent rights, licenses, or any other IP rights, whether with regard to such information itself or anything described by such information. Information provided in this document is proprietary to Microsemi, and Microsemi reserves the right to make any changes to the information in this document or to any products and services at any time without notice.

Microsemi Corporation (Nasdaq: MSCC) offers a comprehensive portfolio of semiconductor and system solutions for aerospace & defense, communications, data center and industrial markets. Products include high-performance and radiation-hardened analog mixed-signal integrated circuits, FPGAs, SoCs and ASICs; power management products; timing and synchronization devices and precise time solutions, setting the world's standard for time; voice processing devices; RF solutions; discrete components; enterprise storage and communication solutions; security technologies and scalable anti-tamper products; Ethernet solutions; Power-over-Ethernet ICs and midspans; as well as custom design capabilities and services. Microsemi is headquartered in Aliso Viejo, California, and has approximately 4,800 employees globally. Learn more at **www.microsemi.com**.