



# **"In-flight" Encryption in Service Provider Networks**

## Abstract

Driven by the exponential growth in traffic and the mass migration of enterprise and personal data to the cloud, the security of electronic communications is a priority issue faced by the service providers tasked with delivering the services driving today's connected global economy. To address this problem, service providers must find a way to efficiently encrypt traffic 'in-flight', as it transits global networks, without compromising service quality. Securing end-to-end transport networks using OTN encryption brings a compelling new solution to the table. It offers a low latency, service & protocol agnostic implementation that makes efficient use of network bandwidth. Innovation in OTN processing silicon, uniquely led by Microsemi's DIGI-G4, delivers against these requirements, while at the same time supporting flexible and scalable encryption service delivery models spanning OTN switched or point-to-point WDM optical transport networks.

## Contents

"In-Flight" Encryption in Service Provider Networks.....	3
Introduction.....	3
"In-Flight" Network Encryption – Market Drivers.....	3
Encryption – The Basics.....	4
Service Provider Requirements.....	4
Complexity and Cost of End-to-End Encryption.....	4
Network Latency / Performance.....	5
Network Efficiency.....	5
Multi-Service Support.....	5
Deployment Flexibility and Scalability.....	5
"In-Flight" Network Encryption.....	6
IPsec (L3) Encryption.....	6
MACsec (L2) Encryption.....	7
OTN Encryption.....	7
Optimized Network Efficiency & Latency.....	8
Multi-Service Capability.....	8
Maximum Network Deployment Flexibility and Scalability.....	8
Bulk/Wavelength OTN Encryption.....	9
Sub-Wavelength OTN Encryption.....	9
Recap.....	10
OTN Encryption: Operator Uses Cases.....	11
Encrypted High-Bandwidth Datacenter Interconnect (DCI) Services.....	11
Encrypted Private Line & Cloud Connect Transport-as-a-Service (TaaS).....	12
Encrypted MPLS Transport.....	13
Operator Use Case Recommendations.....	13
DIGI-G4: Securing Cloud and Communication Service Provider Transport Networks.....	14
Conclusion.....	14

## "In-Flight" Encryption in Service Provider Networks

### Introduction

*"What last year's revelations showed us was irrefutable evidence that unencrypted communications on the Internet are no longer safe. Any communications should be encrypted by default."*

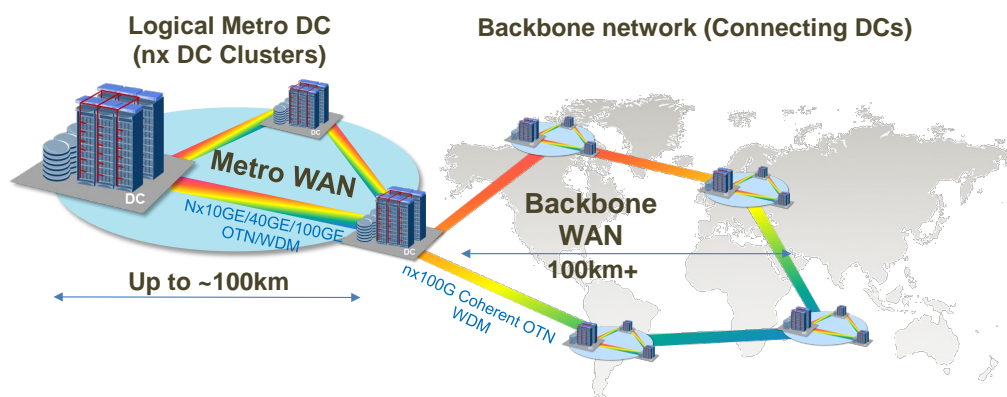
Edward Snowden – *The Guardian Interview, Moscow July 2014*<sup>1</sup>

Driven by a number of recent high-profile events, the security of electronic communications is a priority issue faced by the service providers tasked with delivering the services driving today's connected global economy. The exponential growth in traffic and the distribution of enterprise and personal data to cloud services are the primary flash points for this issue. To address this problem, service providers must find a way to efficiently encrypt traffic "in-flight", as it transits global networks, without compromising service quality. This paper discusses network encryption, the options, benefits and drawbacks, and makes a case for OTN encryption as a means for service providers to accomplish these goals.

### "In-Flight" Network Encryption – Market Drivers

Global cloud IP traffic is forecast to quadruple over the next 5 years, reaching 6.5 Zettabytes by 2018, while global Mobile data traffic will grow nearly 10-fold to almost 300 Exabytes by 2019.<sup>2,3</sup> The adoption of these new services by Enterprises and Consumers alike is fundamentally changing the make-up and traffic patterns of the underlying of the networks that are tasked with delivering content and communications to the end users. In response, service providers are transitioning their networks to a distributed model for interconnecting end-users to content. As shown in [Figure 1](#), single-site, hyper-scale datacenters (DC) in densely populated urban areas are making way for clusters of smaller datacenters interconnected by high-speed optical fiber WAN links. This is driven primarily by real-estate and power requirement constraints. These clusters form logical "Metro Datacenters" that are in turn interconnected globally by high-speed optical backbone WAN links, predominantly based on 100Gbit/s technology.

**Figure 1 • Next-Generation Datacenter Interconnect Network**



These new network topologies result in a large increase in the number of Metro and Backbone WAN optical fiber links, thereby compounding the risk that data is compromised while "in flight"—be it from datacenter to datacenter (also known as Datacenter Interconnect traffic or DCI), from datacenter/central office to end-user, or from end-user to the Internet.<sup>4</sup> A simple example of this exposure of data: even

<sup>1</sup> The Guardian, "Edward Snowden urges professionals to encrypt client communications", July 17, 2014.

<sup>2</sup> "Cisco Global Cloud Index: Forecast and Methodology, 2013-2018.", Cisco, 2014.

<sup>3</sup> "Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2014-2019.", Cisco, 2015.

<sup>4</sup> The Guardian "GCHQ taps fibre-optic cables for secret access to world's communications", June 21, 2013

for the most basic of user interaction with the "cloud", a Google Search, data travels on average 1,500 miles to be processed.<sup>5</sup>

The costs associated with not securing data "in-flight" on these links can be significant. The annual cost to the global economy of cybercrime has been estimated to be more than \$400 Billion USD, and it has been estimated that the U.S. cloud computing industry may lose \$22 to \$35 billion USD over the next three years as a result of the recent high-profile reports.<sup>6, 7</sup> Conservatively, it has been estimated that more than 800 Million individuals had their personal information compromised by cybercrime activities in 2013.<sup>8</sup> Customers have responded accordingly—data security is of paramount importance when evaluating service provider offerings. In a recent survey, more than 70% of service providers stated that security is the number one concern for their customers in making decisions to leverage cloud services.<sup>9</sup>

The challenge faced by service providers is in selecting a holistic end-to-end security strategy that fits well with current and future technology choices for the underlying network tasked with delivering content to end-users.

## Encryption – The Basics

Encryption generally involves two processes:

1. Encryption of data
2. Authentication for anti-forgery protection

Encryption of data is done via a cryptographic algorithm, or block cipher, transforming the data into cryptotext that can only be decrypted by an end-user with a valid 'key'. The most common block ciphers adhere to the Advanced Encryption Standard (AES) specified by the U.S. National Institute of Standards and Technology (NIST) in 2001.<sup>10</sup> AES is based on ciphers with block sizes of 128-bits with three different key size lengths: 128-bit, 192-bits and 256-bits.<sup>11</sup> For example, AES-256, the most common and most secure of the AES standards, refers to the AES block cipher with a 256-bit key length.

Authentication adds an additional layer of security to prevent forgery attacks that attempt to discover the encryption key. At the most basic level, authentication tags are sent along with the payload from the source to the destination. The destination validates the in-band authentication tag prior to decrypting the payload. If the tag is incorrect, the destination ignores the traffic.

## Service Provider Requirements

For communication and Internet content service providers, a number of factors must be weighed when evaluating network encryption solutions including:

- Complexity and cost
- Network latency
- Network throughput / utilization
- Multi-service support
- Deployment flexibility and scalability

## Complexity and Cost of End-to-End Encryption

Given the immense CAPEX and OPEX costs associated with deploying and managing networks, introducing any incremental complexity at any layer of the network must be carefully evaluated. There are several

<sup>5</sup> Samantha Murphy Kelly, "How a Google Search Travels Around the World", Mashable, June 13, 2012.

<sup>6</sup> Intel Security "Net Losses: Estimating the Global Cost of Cybercrime", June 2014.

<sup>7</sup> The Information Technology Innovation Foundation, "How Much Will PRISM Cost the U.S. Cloud Computing Industry?", August 2013.

<sup>8</sup> Intel Security "Net Losses: Estimating the Global Cost of Cybercrime", June 2014

<sup>9</sup> Infonetics Cloud Service Strategies: Global Service Provider Survey, January 2015

<sup>10</sup> Federal Information Processing Standards Publication 197 (FIPS 197). Advanced Encryption Standard (AES). Online: <http://csrc.nist.gov/publications/fips/fips197/fips-197.pdf>.

<sup>11</sup> NIST Special Publication 800-38D, November, 2007. Recommendation for Block Cipher Modes of Operation: Galois/Counter Mode (GCM) and GMAC. Online: <http://csrc.nist.gov/publications/nistpubs/800-38D/SP-800-38D.pdf>.



options for "in-flight" encryption corresponding to different layers in the OSI stack. Typically, the higher up the network stack traffic is encrypted the more complex and costly it will be to implement and manage end-to-end. Operators must therefore evaluate the layer at which they must implement security versus the cost and complexity of encrypting higher up the network stack.

### Network Latency / Performance

Latency is an important characteristic for a wide-variety of services—both for communications and cloud service providers. In addition to the well-known latency limits on voice communications, the user-experience of e-Commerce and Cloud-based Services is known to be related to network latency.

For example:

- Amazon estimated that every 100 msec delay in page loads results in -1% in sales revenue<sup>12</sup>
- Google Search algorithms factor web site speed factors into Google Search result rankings<sup>13</sup>
- Amazon AWS Elastic Cloud Compute (EC2) service aims for latencies of 2 milliseconds or less in order to ensure acceptable user experience.<sup>14</sup>

All network encryption techniques inherently increase network latency, however, the magnitude of this increase is a function of encrypted payload frame size, as well as the parameters of the encryption implementation used, including hardware architecture, cipher block size and key length, cipher mode and authentication mode.

### Network Efficiency

With the well documented divergence of network traffic growth and service revenues, maximizing the utilization of underlying network infrastructure is key for all service providers, as any incremental network capacity comes with significant capital and operational costs. Some "in-flight" encryption techniques force service providers to trade-off network efficiency for security. As such, incremental bandwidth expansion resulting from encryption must be weighed carefully against wasted network capacity costs.

### Multi-Service Support

Although IP/packet and Ethernet-based services represent the fastest growing segment of traffic traversing operator networks, operators also need to contend with the large install base of other protocols such as SONET/SDH, Infiniband and Fiber-Channel for datacenter connectivity. These client types are no less immune to security risk than IP or Ethernet traffic; therefore, selecting a network encryption solution that addresses all client and service types is an important consideration.

### Deployment Flexibility and Scalability

Service providers are also challenged to support a diverse range of service rates and formats. Service delivery rates today range from 1 Gbit/s to 100 Gbit/s and beyond while service offerings span Broadband to Private Wavelength Services. Additionally, emerging service offerings in datacenter and cloud services demand the ability to dynamically turn bandwidth up and down on demand. Supporting secure transport of a wide-range of granularities and service types drives a need for scalability in encrypted payload size and the ability to flexibly adapt encrypted payload sizes dynamically as customer demands change.

Architecturally, service providers must also contend with the wide-range of network topologies and technology choices end-to-end in their networks. In metro and backbone optical transport networks, this includes point-to-point WDM links as well as Layer 1 Switching environments, both of which are predominantly based on Layer 1 OTN technologies today. Layer 2 and Layer 3 Packet-based access and aggregation networks are equally diverse, with point-to-point Ethernet as well as MPLS-based networks being common. Effectively and efficiently encrypting traffic end-to-end therefore demands an encryption solution that can flexibly support all network architectures that may exist at that network layer.

In the subsequent sections, we will discuss the options available to service providers to implement "in-flight" encryption and discuss each in the context of key operator requirements.

<sup>12</sup> Greg Linden, "Make Data Useful.," [http://sites.google.com/site/glinden/Home/StanfordDataMining.20 06-11-28.ppt](http://sites.google.com/site/glinden/Home/StanfordDataMining.20%2006-11-28.ppt)

<sup>13</sup> Webmaster Central Blog, "Using site speed in web search ranking", Google, April 2010

<sup>14</sup> Amazon AWS Elastic Cloud Compute (EC2) 2014-10-02 Documentation:  
<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-regions-availability-zones.html>

## "In-Flight" Network Encryption

Encryption of traffic "in-flight" is defined as encrypting data as it traverses from source to destination in a network. This is not to be confused with encryption of data "at-rest", which refers to protecting access to data stored in hard drives on PCs or large storage arrays in datacenters and central offices.

The three most common methods for implementing "in-flight" network encryption are:

1. IPsec or Layer 3 (L3) encryption
2. MACsec or Layer 2 (L2) encryption
3. OTN or Layer 1 (L1) encryption

The above options are not mutually exclusive. Given the diverse set of requirements across all services, geographies and customers, operators may implement one or all in tandem.

### IPsec (L3) Encryption

IPsec, a suite of protocols standardized by the IETF, enables the encryption of individual packets that make up traffic flows in the IP domain. Authentication headers are added on a per packet basis and are used to validate access to the encrypted data. IPsec supports a number of cryptographic algorithms and authentication modes, with the most common being AES-CBC (Cipher-Block-Chaining) and AES-GCM (Galois Counter Mode).

As a large number of services and applications originate in the IP domain, IPsec offers a true standards-based end-to-end encryption solution that is agnostic to the underlying physical network equipment in place—routers, optical transport equipment, or otherwise.

However, IPsec has several potential limitations. Firstly, IPsec by definition does not support non-IP traffic flows, including datacenter storage protocols such as Fiber-Channel and Infiniband.

There is also a performance vs. cost and power trade off related to incremental processing (CPU or Network Processors) and associated memory resources required to process and encrypt packets with IPsec. This can place an upper limit on flow sizes and total bandwidth that can realistically be implemented by network equipment line cards supporting IPsec encryption. It is not uncommon to be restricted to maximum IPsec-enabled flows of 10 Gbit/s or maximum total IPsec bandwidth per card of 40 Gbit/s on a Router line card supporting a total bandwidth in excess of 100 Gbit/s. Where service needs dictate the requirement to encrypt larger bandwidth traffic flows, for example at 100G wavelength granularity, which is common from DCI applications, service providers may look to lower layer encryption solutions such as MACsec or OTN encryption that are more conducive to these traffic types.

Lastly, the additional overhead necessary to secure each packet results in packet size expansion which in turn results in increases in network latency and wasted bandwidth on optical fiber networks. It has been shown that for smaller frame sizes such as 64-byte packet-sizes, typical of voice or video, IPsec can result in up to 40% wasted bandwidth and incremental latencies on the order of 125 msec.<sup>15</sup> For larger packet sizes, network throughput increases however latency is further degraded, with latencies approaching 350 msec for large packet sizes (1420-bytes).<sup>16</sup> Such latencies violate the tolerances of most services, and service providers who are already struggling to adapt their networks to support the exponential growth in traffic simply cannot afford to absorb this wasted bandwidth.

<sup>15</sup> Marko Bobinac, "Layer 2 Network Encryption – Where safety is not an optical illusion", SafeNet, 2013.

<sup>16</sup> Ibid.

## MACsec (L2) Encryption

Moving down one layer in the OSI model, MACsec, or Media Access Control Security, is an IEEE standard that encrypts data contained within Ethernet frames on point-to-point Ethernet links.<sup>17</sup> MACsec secures data by adding up to 32 authentication and security bytes to the existing Ethernet frame. Both the header and tail are checked by the receiving end of the point-to-point link to ensure that data was not compromised while traversing the link. The MACsec standard supports both 128-bit and 256-bit AES block cipher algorithms.

When compared to IPsec, MACsec offers the following:

- Improved network efficiency and latency as bandwidth expansion is capped at 32-bytes
- Lower incremental cost and power, as the MACsec standard is an extension of standard Ethernet frame format supported by existing Ethernet MAC/PHY technology
- Less complex to manage than at the IP layer

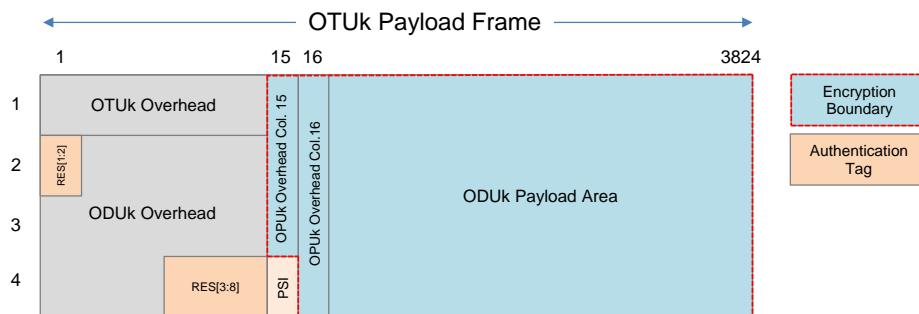
However, MACsec is not without its challenges. Similar to IPsec, MACsec excludes native support for all non-Ethernet client types, forcing service providers to rely on alternative options for securing these clients. MACsec is also constrained to the hop-by-hop architecture that defines the MAC layer. As such, traffic secured by MACsec cannot be transparently managed by downstream nodes. In order to maintain the security of this traffic, MACsec must be implemented at all nodes downstream from the point of encryption in the network, potentially adding cost, power and network planning complexity. Furthermore, MACsec is constrained to Ethernet-based Layer 2 networks. Therefore, for applications originating at higher layers, MACsec must rely on a higher-layer encryption solution, such as IPsec, to round out an end-to-end secure connection.

## OTN Encryption

Encrypting at the Layer 1 using OTN is emerging as a compelling third option available to service providers as it offers an efficient low-latency secure transport solution with a high degree of flexibility in service-type, rate, and network architecture.

As shown in [Figure 2](#), OTN encryption encrypts data contained within the existing OTN payload frame, known as an OPUk. Existing reserved bytes within the overhead of the OTN frame carry the authentication tag. The algorithm and authentication modes used vary and are implementation specific; however, AES-256 is a common block cipher with support for both GCM and CTR modes common.

**Figure 2 • Encrypted OTN Payload Frame**



OTN encryption is similar to MACsec in terms of cost, power and complexity to implement. The incremental cost and power of encrypting traffic is specific to the OTN framer hardware implementation. However, typical incremental processing requirements to encrypt and decrypt OPUk payload and insert/read authentication fields can be considered minimal and lower than the requirements to do so at IP layer, as limited buffering and complex filtering hardware is required. Additionally, complexity to manage secure transport at Layer 1 can be considered akin to Layer 2—less complex than at the IP layer.


<sup>17</sup> IEEE, "802.1AE Media Access Control (MAC) Security", Online: <http://www.ieee802.org/1/pages/802.1ae.html>

The following sections present several compelling advantages offered by OTN encryption.

### Optimized Network Efficiency & Latency

As neither the underlying payload nor the existing OTN frame are padded or extended in any way to facilitate the encryption and authentication process, securing the network with OTN encryption does not come at the expense of wasted precious fiber bandwidth. As shown in Figure 3, OTN encryption offers 100% throughput regardless of the underlying client type or frame-size of packet-based traffic.

**Figure 3 • Network Throughput vs. Encrypted Frame Size (Bytes)**



OTN encryption also offers a low latency encryption solution. Latency varies depending on hardware implementation and encryption mode, however sub-180nsec latencies are achievable for all OPUk frame sizes using an AES-256 block cipher. Incremental latencies of this magnitude leave plenty of margin in the available end-to-end budget for most service provider services.

### Multi-Service Capability

OTN as a Layer 1 transport protocol is a multi-protocol convergence layer. It is capable of transporting virtually all client types and protocols, from constant bit-rate services delivered by OTN, to IP/Packet, Ethernet and Storage-based client revenue streams. Unlike MACsec and IPsec, OTN encryption inherently provides service providers with a single encryption solution addressing all client types and protocols end-to-end within the transport network.

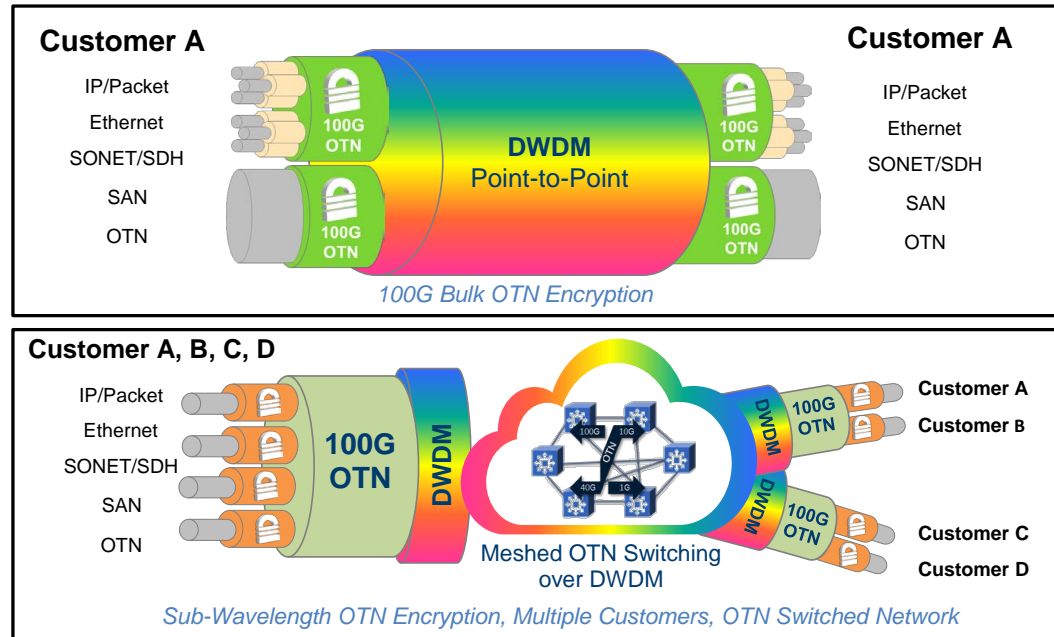
### Maximum Network Deployment Flexibility and Scalability

OTN encryption provides scalability and flexibility in deployment options end-to-end within Layer 1 transport networks. OTN offers the scalability to encrypt payloads with rates from 1.25 Gbit/s up to 100 Gbit/s as well as the flexibility to support multiplexing of lower-rate encrypted OTN signals into higher-rate un-encrypted OTN signals.

This enables operators to support two transport encryption service models:

1. Bulk or Wavelength OTN Encryption
2. Sub-Wavelength OTN Encryption

**Figure 4 • OTN Encryption Service Models: Bulk/Wavelength vs. Sub-Wavelength OTN Encryption**



### Bulk/Wavelength OTN Encryption

Bulk OTN encryption can be defined as encrypting at a wavelength level, such as 10G, 40G or 100G wavelengths, but it need not be limited to service with a 1:1 ratio of client signal to wavelength rate. A bulk OTN encryption service could consist of an encrypted 100 Gbit/s OTU4 signal carrying a 100G client over a single 100G wavelength, or an encrypted OTU4 signal carrying within it a mixture of unencrypted lower-rate clients. In both cases, the entire payload is encrypted at the 100G wavelength granularity.

Bulk encryption would typically be deployed in point-to-point WDM network configurations and single end-customer deployment scenarios. Example use-cases for bulk encryption would be a leased wavelength encrypted transport service or an encrypted Datacenter Interconnect (DCI) service.

### Sub-Wavelength OTN Encryption

Sub-wavelength OTN encryption can be defined as encrypting lower-rate clients prior to multiplexing into higher-rate unencrypted wavelengths.

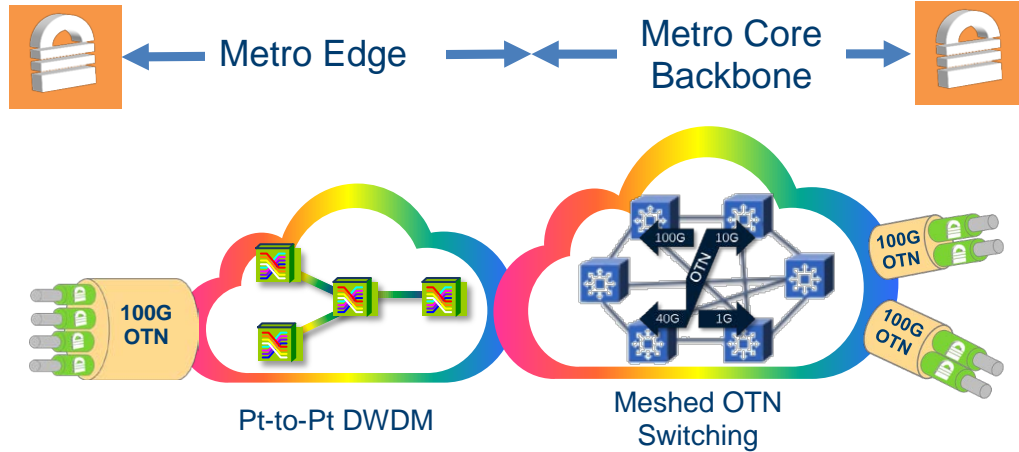
This model is emerging in importance with the growing adoption of Layer 1 OTN switching and converged Packet Optical Transport (P-OTP) network equipment by operators worldwide, as a means to aggregate and switch sub-wavelength traffic to ensure most efficient use of costly 100G optical infrastructure.<sup>18</sup> Sub-wavelength OTN encryption aligns with this model as sub-wavelength traffic flows, possibly from different customers, can be encapsulated into encrypted OTN containers and traverse the network independently, addressing the requirement for end-to-end security of individual traffic sources at no cost to network efficiency or performance.

Sub-wavelength OTN encryption is also not limited to deployment in Layer 1 OTN switching network configurations. This flexibility is especially important given that service provider networks today are a mix of point-to-point WDM and Layer 1 switched architectures. As shown in Figure 5, sub-wavelength encryption could be deployed in 100G Muxponders at the edge of the network to aggregate and encrypt lower-rate traffic onto higher-rate unencrypted 100G wavelengths. As this traffic transitions to the Metro Core and Backbone networks, predominantly based on mesh-based OTN switching networks, these individual sub-wavelength encrypted clients can continue to traverse the networks independently, without the need to decrypt higher-rate wavelengths and therefore compromise the security of

<sup>18</sup> Ibid.

sub-wavelength traffic. As the network evolves over time, pushing Layer 1 switching closer to the Metro Edge, the encryption model does not need to change.

**Figure 5 • Sub-Wavelength OTN Encryption: Deployment Flexibility and Future-Proof**



*Sub-Wavelength OTN Encryption: Future-Proof and Deployment Flexibility*

## Recap

All "in-flight" network encryption options have their merits and drawbacks. The trade-offs of implementing one or all of the above must be weighed against the incremental revenue and profit that can be driven through enhancements to existing service offerings and new revenues from emerging markets.

**Figure 6 • 'In-Flight' Encryption Solutions Scorecard**

Requirement	IPsec (L3)	MACsec (L2)	OTN (L1) Encryption
Complexity & Cost	High	Low	Low ✓
Latency	High	Low	Low ✓
Multi-Service	No	No	Yes ✓
Network Efficiency	Low	Medium	100% ✓
Scalable Encrypted Payload Size	Restricted	Restricted Standard MAC sizes	Flexible 1.25G – 100G ODUflex ✓
Flexible Deployment	-	Low Hop-to-Hop Only	High Wavelength Sub-Wavelength Point-to-Point WDM Layer 1 Switched ✓
End-to-End	IP Only	Layer 2 Only	Layer 1 OTN only

OTN encryption has clear advantages in that it supports flexible and scalable service delivery models without trading without trading off network latency, efficiency and the ability to support a diverse range of services and client types. The following section discusses the potential for incremental services revenue growth for network operators through several use cases for OTN encryption.



## OTN Encryption: Operator Uses Cases


### Encrypted High-Bandwidth Datacenter Interconnect (DCI) Services

Traffic flows between datacenters can be characterized as high bandwidth, packet, Ethernet, or Storage-based protocols, with reach requirements ranging from 100 to 1000s of kilometers depending on the geographical distribution of the datacenters. Examples of security-sensitive services represented by these traffic flows could be transport of trading data for financial institutions, data replication for enterprise disaster recovery / business continuity, or workload balancing for mission-critical applications hosting offered by cloud computing services such as Google Cloud Platform, Amazon AWS EC2, or Microsoft Azure.

Network efficiency is of particular importance for these services given the massive bandwidths that these services demand – in some cases on the order of 10 Terabits/s of interconnect between locations. Operators therefore require an encryption solution that does not trade-off network efficiency for security.

OTN transport over point-to-point DWDM networks leveraging 100G coherent optical transport technologies has emerged as a popular choice for operators to address the bandwidth and reach demands of these services. Encrypting at wavelength granularity, bulk OTN encryption provides a solution that is agnostic to the underlying mix of traffic, whether it be a mix of lower-rate client types or a client that is rate matched to the wavelength. Additionally, should the client type or rate change, bulk OTN encryption provides the flexibility to scale the encrypted payload size and the multi-service support to address changes on-demand.

**Figure 7 • Bulk OTN Encryption for DCI**



For example, as shown in Figure 7, a service provider needs to support a 100 Gbit/s link from DC-West to DC-East. Today, traffic flows drive the need to transport 10x 10GbE clients over a 100G wavelength, but the Service Provider expects this largely to transition over time to 100GbE. Bulk OTN encryption supports both models, encrypting an ODU4 that contains unencrypted 10x 10GbE clients or a single 100GbE client. Additionally, if at any point the customer requests support for transport of storage clients such as Fiber-Channel or Infiniband for data replication services across the same 100G WDM link, the bulk OTN encryption solution can adapt to encapsulate the storage clients into the right-sized ODUK container prior to encrypting the higher-rate 100G ODU4 wavelength.

## Encrypted Private Line & Cloud Connect Transport-as-a-Service (TaaS)

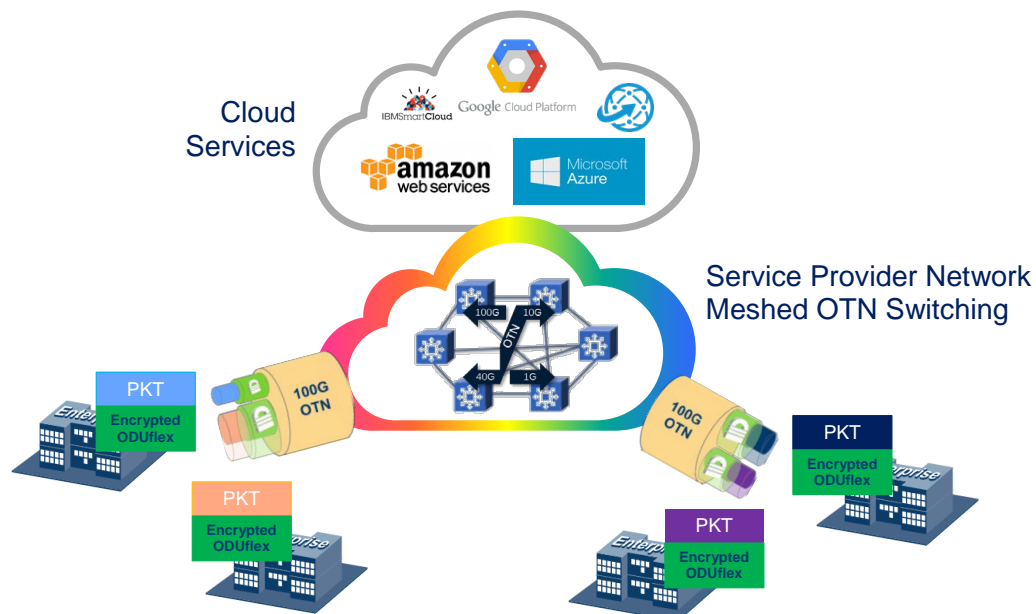
Although often overshadowed by growth in IP traffic from services such as Mobile and Broadband services, Private Line services still represent a significant source of revenue for operators. In the US alone, Private Line services represented a \$30B USD market in 2014.<sup>19</sup> In the same year, AT&T reported that wholesale wireline services represented approximately 10% of group revenues.<sup>20</sup>

The types of offerings making up these services are diverse and evolving—spanning traditional Ethernet Private Line, Wavelength and Leased Dark Fiber to emerging cloud interconnect services for Enterprises—each having a distinct set of SLAs and requirements. Generally, customers look to these services when they require constant, and in some cases high, bandwidth between sites, with their traffic kept separate from other traffic sources. Security is growing in importance, especially as more Enterprises are adopting cloud-based services for customer-facing applications and mission-critical workloads. As of 2014, 62% of enterprises are using cloud computing, with 70% of Verizon cloud customers using it for external-facing production applications.<sup>21</sup>

With OTN-based optical transport as the choice for most operators for delivering these services, sub-wavelength OTN encryption provides a value-add option for enhancing and differentiating these services in the market without limiting service types, scalability or performance. It offers a client and protocol agnostic solution with scalable, constant rate encrypted payload sizes up to 100 Gbit/s, as well as the flexibility to transport these services independently end-to-end in point-to-point or Layer 1 switched network.

Complementing sub-wavelength OTN encryption with the bandwidth on-demand characteristics offered by a Layer 1 OTN switched network additionally provides operators with a powerful secure offering to further drive adoption of high-value cloud-based services. Cloud-based workloads are inherently bursty and therefore do not fit well in an environment where the bandwidth between end-users and applications is fixed. Customers will also only be willing to move mission-critical workloads into the cloud with guaranteed security. The solution is a flexible encrypted cloud connect transport-as-a-service that can scale on demand as customer workloads change.

**Figure 8 • Encrypted Enterprise Cloud Transport-as-a-Service (TaaS)**



<sup>19</sup> "Private Line and Wavelength Services, 2014-2019", Insight Research Corp, October 2014.

<sup>20</sup> AT&T Q4'2014 Quarterly Results

<sup>21</sup> "State of the Market: Enterprise Cloud 2014", Verizon



With sub-wavelength OTN encryption in a Layer 1 switched network, private packet-based customer traffic can be GFP encapsulated and transported in encrypted flexibly-sized OTN transport containers, known as ODUflex, to and from datacenters housing Enterprise cloud-based applications. These encrypted ODUflex-based private circuits can scale the delivered bandwidth on-demand using automatic hitless adjustment techniques supported by the G.709 OTN standard (G.hao).

### Encrypted MPLS Transport

Up until now, in order to support encryption of MPLS traffic flows end-to-end in an MPLS network, operators have been forced into one of two corners:

1. Leverage IPsec to encrypt the underlying packets, resulting in trade-offs in latency and throughput
2. Leverage MACsec-based solutions, requiring a proprietary MACsec implementation to circumvent hop-by-hop constraints.

Today, more and more operators are looking to OTN-based optical transport equipment as the underpinning of an MPLS network to enable efficient transport of MPLS flows end-to-end. This has been further advanced by the introduction of Packet Optical Transport Platforms (P-OTPs / P-OTNs) and Hybrid Packet / OTN Line Card architectures that allow for native packet switching, encapsulation of packet-based flows into ODUflex containers and efficient grooming of many of these lower-rate flows onto 100G wavelengths.<sup>22</sup>

Sub-wavelength OTN encryption extends this capability, enabling encryption and transport of MPLS flows in a manner that is agnostic to the underlying network architecture at no cost to network performance or utilization. In this case, individual MPLS packet flows are mapped into encrypted ODUflex containers which in turn are multiplexed onto higher-rate wavelengths. These encrypted flows can then be routed independently throughout the transport network over both switched and point-to-point OTN-based transport networks.

### Operator Use Case Recommendations

OTN encryption provides service providers with a viable option to enhance and differentiate their transport service offerings without trading off quality, flexibility, performance, cost and complexity. To maximize this opportunity, service providers must select an encryption-enabled OTN processing silicon solution that provides all the tools to address a wide breadth of service types and application models, both at wavelength and sub-wavelength granularity. Not all OTN processors are created equal in this sense. The next section discusses technology advancements by Microsemi to enable operators to deliver a world-class secure transport solution.


<sup>22</sup> "Making 100G OTN Economical: OTN Switching & Packet-Optical Transport", Microsemi, May 2014.

## DIGI-G4: Securing Cloud and Communication Service Provider Transport Networks

Building on the OTN processing silicon innovation and leadership in its DIGI-120G device, Microsemi's DIGI-G4 400G OTN processor addresses cloud and communication service provider needs for a flexible, scalable SDN-ready, encrypted transport infrastructure. It is the industry's densest single-chip 4x 100G OTN processor offering 50 percent less power per port than the previous generation. The DIGI-G4 delivers the capacity, security and flexibility required for 400G line cards in packet optical transport platforms (P-OTP), ROADM/WDM and hyperscale data center interconnect platforms.

**Figure 9 • DIGI-G4: Low-Latency, Multi-Service, Flexible and Scalable Deployment**

Requirement	OTN (L1) Encryption		Microsemi DIGI-G4 OTN Encryption
Complexity & Cost	Low		✓
Latency	Low		✓ <180ns
Multi-Service	Yes		✓
Network Efficiency	100%		✓
Scalable Encrypted Payload Size	Flexible	1.25G – 100G ODUflex	✓
Flexible Deployment	High	Wavelength Sub-Wavelength Point-to-Point WDM Layer 1 Switched	✓



The DIGI-G4 integrates rate-agile, protocol-agnostic, NIST FIPS 197 certified OTN encryption functionality with less than 180ns encryption latency, allowing service providers to encrypt services without sacrificing network performance and efficiency.

The DIGI-G4 supports flexible encrypted service delivery and network deployment models, allowing service providers to deploy encryption-enabled transport platforms end-to-end in the transport network, spanning point-to-point WDM to Layer 1 Switched OTN networks and bulk to sub-wavelength encryption service delivery models. This flexibility means it protects investments as network architectures and deployment models evolve.

## Conclusion

The case to encrypt data in Service Provider networks is clear. The proliferation of mobile and cloud services has exponentially increased the amount of sensitive enterprise and personal data that transits service provider networks. Several "in-flight" network encryption solutions provide a solution to this problem, but with different solutions operating at correspondingly different layers in the network, they share very different characteristics in terms of performance, cost, and implementation flexibility. To this end, the cost, complexity and impact to network performance and service availability must be carefully considered when deciding how to encrypt confidential data.

Securing end-to-end transport networks using OTN encryption brings a compelling new solution to the table. It offers a low latency, service & protocol agnostic implementation that makes efficient use of network bandwidth—a critical requirement for service providers fighting to stay ahead of the deluge of data. Additionally, innovation in OTN processing silicon, uniquely led by Microsemi's DIGI-G4, enables service providers to support multiple flexible and scalable network models including OTN switched or point-to-point WDM. This not only allows service providers to enhance existing offerings but also pursue new revenue and profit generating services to differentiate themselves in the market, without the need to deploy corresponding greenfield networks.



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

Within the USA: +1 (800) 713-4113  
Outside the USA: +1 (949) 380-6100  
Sales: +1 (949) 380-6136  
Fax: +1 (949) 215-4996  
E-mail: [sales.support@microsemi.com](mailto:sales.support@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 Corporation (Nasdaq: MSCC) offers a comprehensive portfolio of semiconductor and system solutions for communications, defense & security, aerospace 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 Communications 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](http://www.microsemi.com).

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.