Benefits of OTN on Router Line Interfaces for Service Providers and OEMs

Introduction

The continued growth of big data, combined with enterprises migrating their workloads to the cloud, has resulted in IP traffic growing more than fivefold in the past 5 years[1]. These packet services are managed by IP routers, which rely on optical transport networks to create the mesh that connects the routers together across a wide area network, as shown in Figure 1. While the traffic being managed by routers is generally packet based, optical transport networks rely on the Optical Transport Network (OTN) protocol, which was optimized for the needs of the service providers that manage the networks. The requirement to manage both the packet-based router network and the OTN-based optical transport network can create multiple challenges and inefficiencies. These networks are often managed by separate groups and treated as separate networks. However, the industry is working hard, via technologies such as Software Defined Networking (SDN), to bring the physical layer networks together and truly manage them end to end as one network. As a first step, there are significant benefits to be had by extending the leading optical transport protocol, OTN, to the router interface, as it has more robust end-to-end fault detection, fault isolation and performance monitoring capabilities than Ethernet. These capabilities both increase operational efficiencies and reduce capital expenditures. This paper will focus on the connectivity of IP routers to the optical transport network and discuss trade-offs between different solutions and key features of the router line interface.

OTN Overview

OTN was developed within the ITU-T as a standard for Dense Wavelength Division Multiplexed (DWDM) signals, and is used by service providers worldwide to manage their DWDM networks. OTN has become the dominant protocol for optical transport networks for a number of reasons, including the following:

- Unlike SONET/SDH, OTN was designed to be an efficient transport layer for packet services such as Ethernet. At the same time, OTN is able to support the multiplexing of many different protocols including SONET/SDH, video, and storage protocols such as Fiber Channel. This allows service providers to support both legacy and new segments of their network on a single converged network.
- OTN has overhead for operations, administration and maintenance (OAM) that are optimized for optical transport.
- OTN enables end-to-end performance monitoring, simplifies fault isolation and provides protection switching for very fast restoration.
- OTN provides a standard-based mechanism for forward error correction (FEC), allowing for additional reach of optical connections.
- OTN can scale to very high data rates. Standards already exist for 100Gb/s, with work ongoing that will scale this to multiple hundreds of Gb/s and then to Terabits per second.

OTN’s performance monitoring and fault isolation capabilities are very powerful. As opposed to just knowing that there is a fault somewhere in the path, the end-to-end capabilities of OTN allow the Network Management System (NMS) to know at which link in the path it has occurred. OTN also allows this awareness to span different carrier networks, so that there are no ‘blind spots’. If the end-to-end path crosses carrier boundaries, OTN’s Tandem Connection Monitoring (TCM) capabilities allow each carrier to know the location of the fault, regardless of which carrier’s network the fault occurred on.

Additional background on OTN can be found in “A Tutorial on ITU-T G.709 Optical Transport Networks (OTN)”[2].

[1] Cisco VNI, May 2015

[2] A Tutorial on ITU-T G.709 Optical Transport Networks (OTN), PMC white paper PMC-2081250 by Steve Gorshe, June 2011

Router to Optical Transport Network Connectivity

Routers are typically peered with adjacent routers via the Ethernet protocol to form the layer 3 network upon which the internet is based. The optical transport network, which is a layer 1 network, provides the physical connectivity to facilitate this mesh of routers, as shown in Figure 1. The physical connection between the router and the optical transport network is typically done with a standard “grey” optical module. In most cases, Ethernet is used to make this connection. The optical transport network will take this Ethernet client and map it onto an OTN framing structure so it can be multiplexed and switched in the optical transport network. However, it is also possible to support OTN framing directly on the router, as shown in Figure 1, so that the router now presents itself as an OTN client to the optical transport network. Moving OTN onto the router simplifies the management and monitoring of the entire physical layer network from router port to router port by having a single, common protocol.

Figure 1: Connection between the Router and the Optical Transport Network

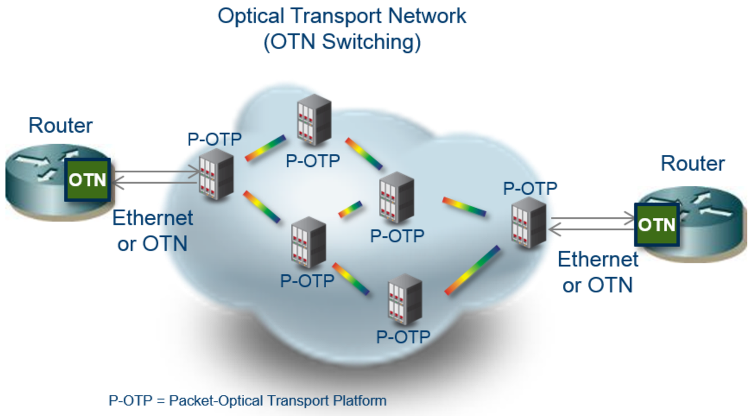

End to End Monitoring and Fault Isolation

If the packet and optical transport networks are managed and treated separately, it can be difficult to detect and isolate a fault, particularly if it is related to the connection between a router and the optical transport domain. However, if OTN is extended from the optical transport domain to the router, this task becomes more straightforward. An example of this is shown in Figure 2. In this case, there was a cut in the fiber that connected the egress traffic on the router to the optical transport network. If the routers were connected to the optical transport network with Ethernet, the following would occur:

- Router B would know there was an issue and send a remote fault to the Network Management System (NMS), but would not know anything about the location of the issue.
- Router A, whose egress connection is the problem, would not immediately have any information on the issue. It will periodically send out messages to check that it is still connected to the optical transport network, but won’t immediately learn of the fiber cut.
- The optical transport box that has Router A as a client will generate a client signal fail (CSF), but cannot communicate the information back to Router A.
- The inability for different nodes on the network to share information about the issue will delay the time for Router A to switch to its protect path.

However, if the router port supports the OTN protocol, the sharing of information will result in a faster resolution of the issues.

- The optical transport box that has Router A as an OTN client will generate a backward defect indicator (BDI) at the OTU layer, and communicate this to Router A
- After receiving the BDI, Router A will know there is an issue and switch to its protect path

Figure 2: OTN on the Router Enables Full End-to-end Monitoring and Fault Isolation

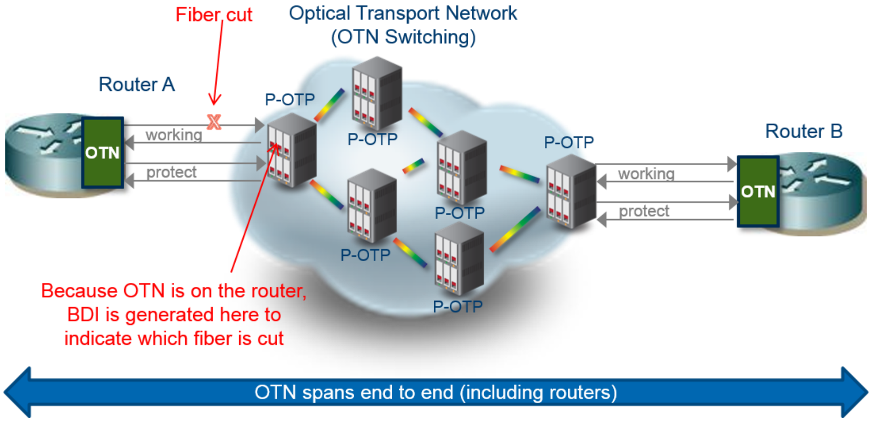

Performance Monitoring

The addition of the OTN protocol on the router also benefits service providers in terms of monitoring performance related to Service Level Agreements (SLAs). This is due to the end-to-end monitoring capabilities of OTN that leverage the standard-based FEC used on each link. With Ethernet connectivity, depending on which Ethernet rate and physical medium dependent sub-layer (PMD) is used, there may or may not be a FEC associated with the link. If no FEC is used for the Ethernet link, no bit error information will be available, as only packet and symbol errors will be counted. If a FEC is used, the error rate is only known for the one link and not communicated to the optical transport network. OTN, on the other hand, has the ability to monitor the full end-to-end error rate. In the example shown in Figure 3, Router A has issues with the link to the optical transport network and is injecting errors into the network. Because the router supports the OTN protocol on the port, the service provider will be able to monitor the end-to-end bit error rate (BER) and compare it with the SLA, and will see the issue. The service provider will also know the end-to-end error rate across the optical transport network, and knows that it is not degrading the error rate, and can quickly determine that the issue is in the connection to the router. It is now clear where a technician needs to be dispatched, and the problem can be quickly resolved.

Figure 3: OTN on the Router Eases Performance Monitoring Related to the SLA

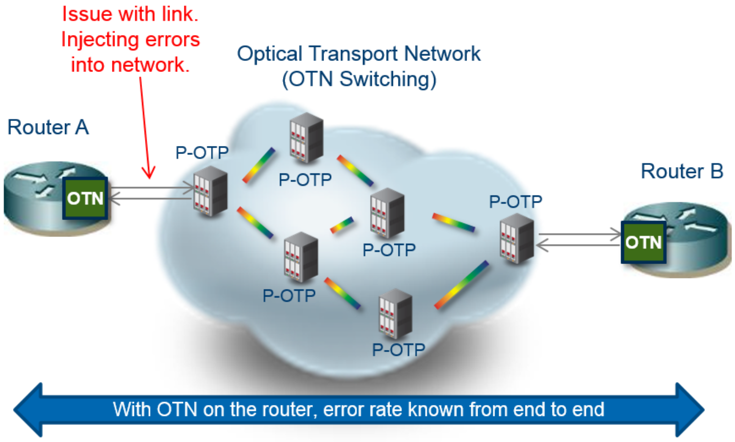

An additional monitoring feature of OTN is that it enables single-ended performance monitoring. This means that the error, defect and performance information from one end is echoed back to the other end so that both ends have the complete picture. This can be important in certain problem scenarios such as that shown in Figure 4. In this case, both routers support OTN, but the connection to the control plane on router B has gone down. If Ethernet were used for the router-optical transport interface, it would have been impossible to monitor the performance of the physical layer connecting to Router B, but because OTN echoes this information to Router A, polling Router A lets the NMS know that there are no issues with the dataplane and that traffic is flowing normally.

Figure 4: OTN Enables Single-ended Error Polling

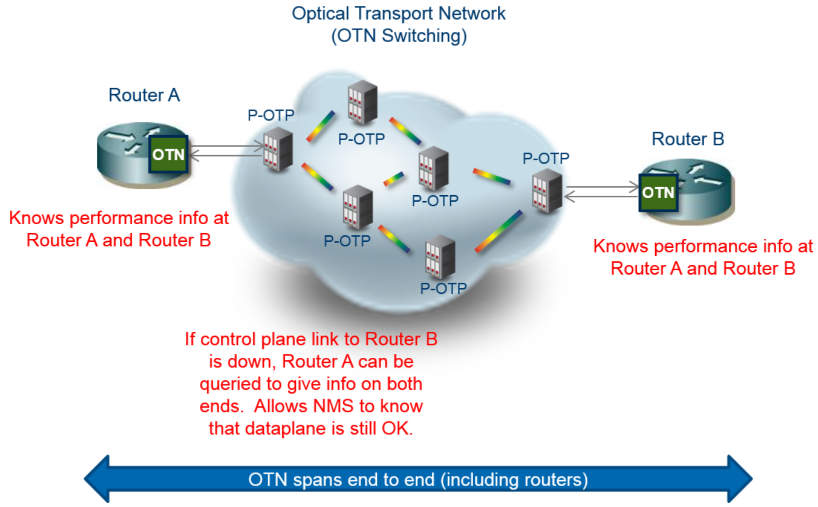
Optical Transport Network
(OTN Switching)
Router A
OTN
P-OTP
P-OTP
P-OTP
P-OTP
P-OTP
P-OTP
Router B
OTN
Knows performance info at Router A and Router B
Knows performance info at Router A and Router B
If control plane link to Router B is down, Router A can be queried to give info on both ends. Allows NMS to know that dataplane is still OK.
OTN spans end to end (including routers)

Capex Savings

In addition to operational efficiencies in management and monitoring, OTN on the router also enables capital expenditure (capex) savings, in part due to the integration of a standard-based FEC with OTN. When routers are peered together, they may be tens of kilometers apart. In some cases, an optical module may not support the length of fiber needed to connect the routers, and a repeater is required. However, the use of the OTN FEC allows for the extension of the optical link, eliminating the need for optical repeaters in some cases and saving the cost of that equipment.

Figure 5: Capex Savings Through the Use of OTN FEC

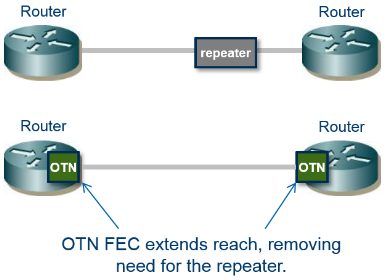

OTN Multiplexing

The use of OTN discussed so far has been for an OTN wrapper application where a single Ethernet client is transparently mapped into the OTN frame structure. A recent advance in optical network architecture has been the introduction of OTN switches, which are network elements capable of switching individual circuits at the sub-wavelength level. As OTN switches are deployed, a router supporting OTN could multiplex several separate clients into a single higher speed interface without sacrificing the ability for the optical transport network to switch each client independently. This feature on the router can reduce the number of ports used to connect the router to the optical transport network, reducing the total expenditure on optical interfaces and enabling increased line-card density while also simplifying the layer 2 feature requirements on the optical transport equipment as explained below.

Figure 6 shows different ways to connect many 10GE clients to the optical transport infrastructure. In the first case, as shown in Figure 6a, many 10GE optical connections are used. This allows for simple processing in the optical transport domain, as these can now be mapped into an OTN container, multiplexed together, and switched as required to send the clients across the network. However, this solution requires many 10GE connections with the cost of all of the optics, limits faceplate density, and requires that all flows consume 10G of optics whether or not they require a full 10G of bandwidth.

An alternative is shown in Figure 6b. In this case, the network processor on the router aggregates the flows into a single 100GE client. This reduces the number of optical connections required, but instead pushes the need for layer 2 functionality into the optical transport network so that the optical transport platform can reconstruct individual packet flows before mapping each into OTN circuits. There are also additional efficiency gains because the flows that are aggregated do not need to be 10GE. These flows can also be at lower Ethernet rates, which allows more than 10 flows to be supported.

A third solution is shown in Figure 6c. In this case, OTN multiplexing is used on the router. Like in b), fewer optical connections are required, saving the cost of the optics and allowing for higher density line cards. This also supports an efficient aggregation of less than 10G flows, as ODU0's can be used to support 1GE flows, and ODUflex's can be used to support other rates. Unlike in b), no layer 2 processing is required in the optical transport equipment, simplifying the system requirements. The optical transport equipment can simply de-multiplex the OTU4 into the multiple packet flows (each within its own ODUk container) and then switch these as required.

Figure 6: Benefits of OTN Multiplexing on the Router

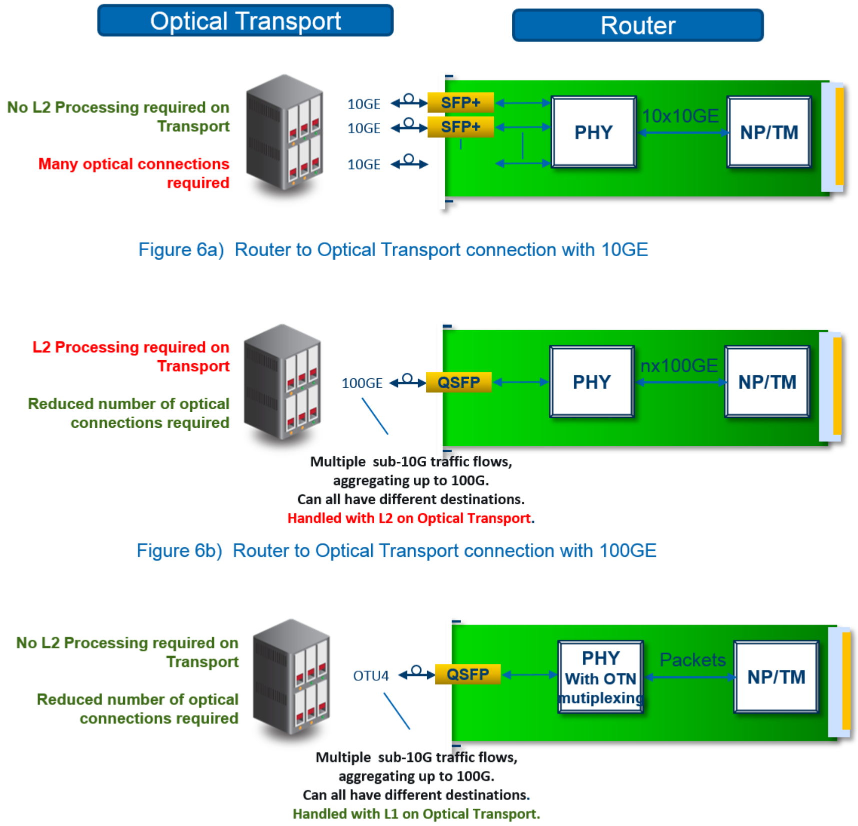
Optical Transport
Router
No L2 Processing required on Transport
Many optical connections required
10GE
10GE
10GE
SFP+
SFP+
PHY
10x10GE
NP/TM

Figure 6a) Router to Optical Transport connection with 10GE

L2 Processing required on Transport
Reduced number of optical connections required
100GE
QSFP
PHY
nx100GE
NP/TM
Multiple sub-10G traffic flows,
aggregating up to 100G.
Can all have different destinations.
Handled with L2 on Optical Transport.

Figure 6b) Router to Optical Transport connection with 100GE

No L2 Processing required on Transport
Reduced number of optical connections required
OTU4
QSFP
PHY With OTN mutiplexing
Packets
NP/TM
Multiple sub-10G traffic flows,
aggregating up to 100G.
Can all have different destinations.
Handled with L1 on Optical Transport.

Figure 6c) Router to Optical Transport connection with OTU4 and OTN Multiplexing

There are also Ethernet based solutions being proposed to accomplish similar functionality as that shown in Figure 6c. One option is to use a multi-link gearbox (MLG) to multiplex the 10GE clients into 100GE. The challenge with this solution is that MLG is not a common feature of optical transport interfaces, whereas multiplexed OTN interfaces are common. MLG also limits clients to 10GE and 40GE. In contrast, OTN can support clients in multiples of 1G. The use of MLG can also create issues if Precision Time Protocol (PTP) is to be supported.

Security and OTN Encryption

After many highly publicized security breaches, such as the documents leaked by Edward Snowden, there has been an increased focus on data security. Two of the most common ways to encrypt "in-flight" traffic on a router today are IPSec (layer 3) and MACSec (layer 2). With the addition of OTN support on the router, OTN encryption (layer 1) is enabled as a third option.

As shown in Figure 7, 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.

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 the IP layer, as limited buffering and 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.

Figure 7: Encrypted OTN Payload Frame

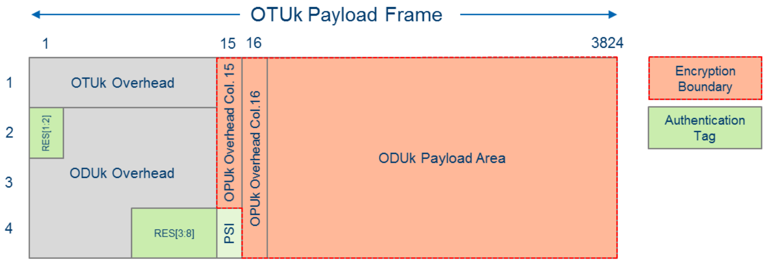

One benefit of OTN encryption compared to MACSec is OTN's support for end-to-end encryption. MACsec is constrained to the hop-by-hop architecture that defines the MAC layer. As such, traffic secured by MACsec cannot be transparently managed by downstream nodes. In order to maintain the security of this 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. With OTN encryption, it is sufficient for only the end points to encrypt and decrypt the traffic.

An additional benefit of OTN encryption is that is does not impact the network throughput. Because 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 8, OTN encryption offers 100% throughput regardless of the underlying client type or frame-size of packet-based traffic.

Figure 8: Network Throughput vs Encrypted Frame Size

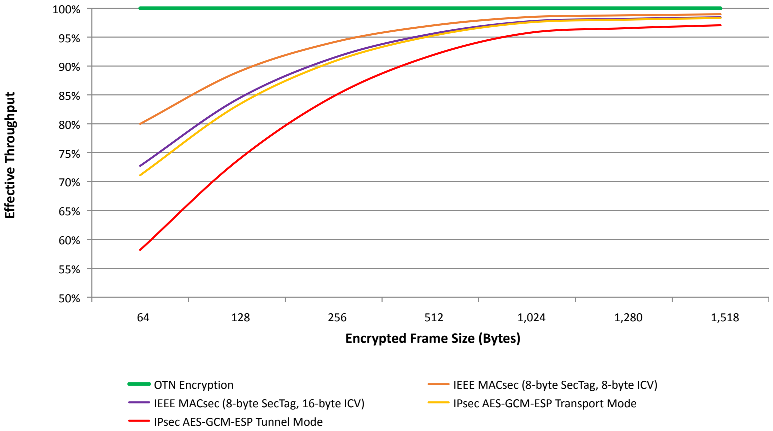

OTN encryption also offers a low latency encryption solution, making it a particularly valuable option for enterprise private-line and datacenter interconnect applications. Latency varies depending on the hardware implementation and encryption mode. However, sub-180nsec latencies are achievable for all 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.

Further reading on OTN encryption can be found in "'In-flight' Encryption in Service Provider Networks."[3]

Network Timing and PTP

In an Ethernet or IP network that relies on packet switching, timing synchronization is very important to ensure the quality of service, particularly for mobile users. In order to synchronize timing across the network, many IP routers are designed to support IEEE 1588v2 (also known as Precision Timing Protocol, PTP). There are two primary options to support PTP in conjunction with OTN on a router.

The first method is to transparently transport the PTP timing within a client that is carried over OTN, such as PTP-over-Ethernet-over-OTN. The benefit of this method is that it does not require any PTP functionality within the OTN switching network. However, to avoid degrading the quality of the distributed timing, this method requires the OTN switching network transport latency to have symmetric delays. Achieving this symmetric delay and maintaining it after a protection switching event is very challenging. As a result, this method is best suited for point to point applications where the delays are constant and can be compensated for.

A second method to distribute PTP timing over an OTN switched network is using the OTN overhead. The distributed timing accuracy can be very high because this method does not suffer from common impairments such as packet delay variation. Because of the quality of this timing, the carrier may be able to sell it as a service. PTP-over-OTN overhead is currently being standardized by ITU-T Study Group 15.

[3] 'In-flight' Encryption in Service Provider Networks, PMC white paper PMC-2150716

PMC's META Product Family: Enabling OTN on a Router

To enable OTN on a router, a device is needed that supports both Ethernet and OTN while also supporting additional functionality such as OTN encryption and PTP. In addition to these features, multi-rate support is also important. Supporting 10GE, 40GE and 100GE within the same device allows an OEM to design a single router line card to support multiple rates. This increases flexibility, reduces development cost and reduces the number of variants for both the OEM and Service Provider to support in inventory.

As shown in Figure 9, the router PHY is the interface to the optics on the linecard. Therefore, a key feature of the PHY is the ability to interface to a wide variety of optics at 10G and 25G lane rates, which are the most common today.

These requirements are all met by PMC's META family of Router PHY devices. The META product family has options supporting from 20Gb/s to 240Gb/s of bandwidth in a single device. The 240G device is able to simultaneously support 10GE/OTU2e, 40GE/OTU3 and 100GE/OTU4. The entire family of META devices also supports PTP.

An example of a router line card using PMC's META-240G device is shown in Figure 9. This example shows a line card supporting up to 800Gb/s of bandwidth. Clients can be configured as Ethernet or OTN on a port-by-port basis. As a QSFP cage can be compatible with 4x10G, 1x40G or 1x100G optics, the interfaces on the META-240G can also be programmed to support any of these rates. In the 100G case, the META-240G supports both 10x10G and 4x25G as it has integrated the gearbox functionality. With different optics, the line card below can also support up to 960Gb/s of traffic, with each META-240G supporting 24x10GE/OTU2e.

The META 240G supports OTN encryption with 20 encryption engines, enabling a low latency security option that does not impact the total throughput. It also supports both PTP-over-Ethernet-over-OTN and PTP-over-OTN Overhead options.

Figure 9: 800G Router Line Card based on META-240G

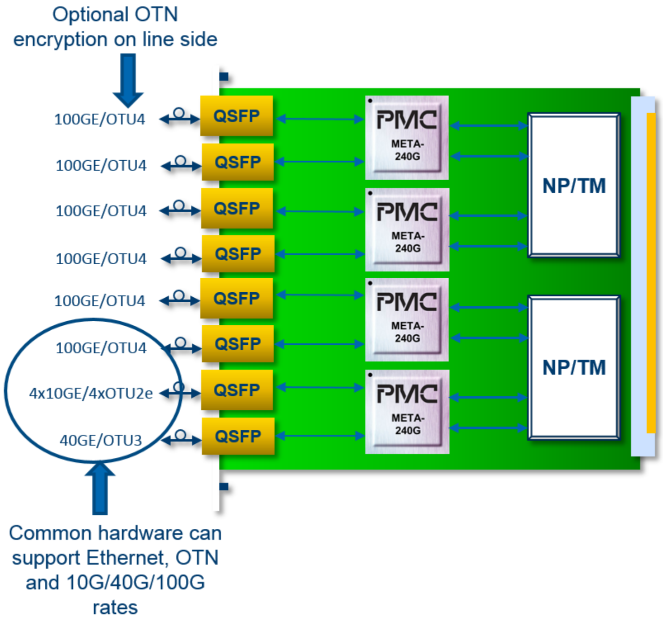

PMC also has a product line to support OTN multiplexing, called the DIGI product line. The latest device in the DIGI product line, the DIGI-G4, supports multiplexing of 10GE/OTU2e, 40GE/OTU3 and 100GE/OTU4 clients, and can multiplex the lower rate clients into an OTU4, as shown in Figure 10. An Interlaken interface can also be used to support GE rates to be multiplexed into an OTU4. The device supports OTN encryption and PTP as well. It uses the same Software Design Kit as the META product family. A common hardware design can support both the META and DIGI devices, allowing a single design to cover all requirements.

Figure 10: Router Line Card support OTN Multiplexing

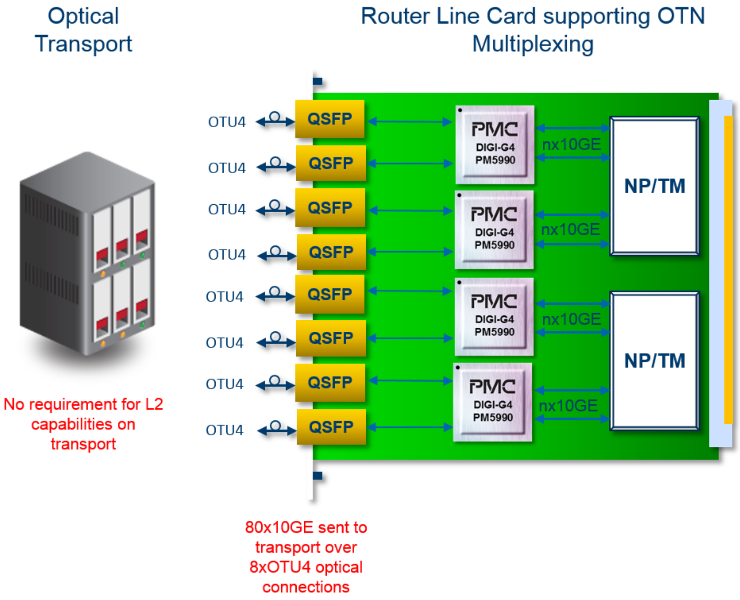
Optical Transport
Router Line Card supporting OTN Multiplexing
OTU4
QSFP
PMC DIGI-G4 PM5990
nx10GE
NP/TM
No requirement for L2 capabilities on transport
80x10GE sent to transport over 8xOTU4 optical connections

Conclusions

Routers rely on optical transport networks to create the mesh to connect to other routers in the network. OTN is the dominant protocol in optical transport networks, and there is significant value in extending this protocol to the router. OTN has very powerful end-to-end performance monitoring and fault isolation capabilities, increasing operational efficiencies and capex savings for the enterprises and service providers deploying the routers. By supporting the OTN protocol on the router ports that interface to the optical transport network, service providers can monitor end-to-end error rates and can more quickly recover from faults, improving the service level to their customers. They can also more quickly and accurately determine the location of faults, reducing operational expenditures tied to debugging network issues. Service providers can also achieve capital expenditure savings by reducing the number of repeaters used for reach extension when peering routers.

When OEMs design router line cards, a PHY device is generally required. PHY devices that support OTN can enable these benefits to their customers, the service providers, in a similar power and density profile to Ethernet only devices. PMC's META and DIGI families of OTN devices are ideally suited for these applications.