Universal Aggregation For Service Convergence: Residential, Mobility & Business

A Technical Paper prepared for SCTE by

Michael Ting Wang
Network Architect III
Shaw Communications Inc.
2728 Hopewell Place NE, Calgary, AB,
Canada T1Y 7J7
+1 (403) 303-4054
Michael.Wang@sjrb.ca

Table of Contents

List of Figures

List of Tables

1. Introduction

Network operators must modernize their networks to support the growing demand for faster internet connections, mobile backhaul and fronthaul, and business services. Universal aggregation is a converged networking approach that uses a shared coherent Dense Wavelength Division Multiplexing (DWDM) core with Internet Protocol (IP) routing and switching platforms.

As is the case for other Multiple-System Operators (MSOs), Shaw must remain competitive while striving to meet the intense challenges of delivering a better user experience during the COVID-19 pandemic. The pandemic has put increasing pressure on many operators and universal aggregation could prove to be a valuable tool for their current and post COVID-19 challenges. It provides the right speeds, footprint, aggregation, and automation capabilities with the ability to integrate Internet Protocol/Multi-Protocol Label Switching (IP/MPLS) with transport infrastructure.

This paper outlines Shaw's work to provide universal aggregation under its modular, coherent DWDM platform while providing greater capacity and significant network cost reduction. The legacy approach to deploying multiple interconnected and complex platforms for each type of service—such as residential, mobility, and business—often results in increased operational cost and difficulty troubleshooting.

For operators, universal aggregation creates a smarter, simpler, and more agile network. Shaw's universal aggregation network supports 100 Gigabits Ethernet (100GbE), Optical Transport Unit 4 (OTU4), 10 Gigabits Ethernet (10GbE), and Optical Transport Unit 2/2e (OTU2/2e) over 100Gbps & 200Gbps wavelengths. It uses a Zero Touch Provisioning (ZTP) system that simplifies operations and allows Shaw engineers to automate most of the deployment tasks. Currently, its maximum bandwidth capacity is 1 Terabit Per Second (Tbits/s) in a single rack unit at 224 watts. The paper will present a comprehensive overview of the implementation process and important considerations for the industry to replicate Shaw's success.

2. Drivers for Convergence

Operators' economic environment has changed significantly over the last two years.

Firstly, there has been a dramatic increase in traffic due to the explosion of residential data/video/voice, mobile xHaul, and business services. It has significantly reduced the spare capacity in the deployed networks and making it necessary to overbuild and introduce new transport technologies to satisfy projected traffic demand. Secondly, most operators are experiencing a need to reduce their operational costs, particularly in building and managing complex multilayer, multivendor networks. To simultaneously satisfy these two trends, operators need to integrate residential, mobility, and business traffic under a single transport network while also providing a simple and automated way to manage that network.

In network infrastructure, where in the past it may have been necessary to use separate equipment for optical transport (layer 1), Ethernet switching (layer 2), and IP/MPLS routing (layer 3), there is an evolution toward a model where all three tiers are integrated. The trends show a convergence of DWDM with Ethernet/client Optical Transport network (OTN)-Synchronous Optical Networking (SONET) as the initial step, and the integration of DWDM/Ethernet-OTN-SONET with IP/MPLS in the future. Separately, there is also a trend toward integrating three types of network services onto common infrastructure: residential, mobility, and business. Furthermore, network-layer convergence will integrate photonic with Ethernet/client OTN/client SONET and IP/MPLS.

Tables 1 and 2 demonstrate the concepts of traffic-type convergence and network-layer convergence, respectively.

Table 1 – Traffic Convergence

	Traffic Categories
Traffic Convergence	Residential
	Mobility
	Business

Table 2 – Network-Layer Convergence

	Network Layers
Network-Layer Convergence	Tier 1 - Optical Photonic
	Tier 2 - Carrier Ethernet
	Tier 3 - IP/MPLS (Future)

2.1. Convergence of Traffic Types

Since March 2020, cable networks have seen a 30.8% growth in downstream traffic and a 54.8% growth in upstream traffic (**ref. [9]**), and the enormous growth in the three service types is exacerbating network silos for operators. The convergence drivers for different traffic types will be discussed in detail below.

2.1.1. Residential Traffic

Over the years, the MSO landscape has proven itself to be dynamic. Residential triple-play once dominated revenue opportunities. Although cable operators are turning their attention to new revenue-generating opportunities in wireless and business services, residential is still a crucial revenue source. Now, operators are focused on aggressively reducing the operational costs of delivering residential services while continuing to meet customer demands for faster speeds.

Residential broadband services are generally carried over a separate DWDM/MPLS network from other types of services, managed by distinct network monitoring systems and, most importantly, live in their own priority structure. They have their own service tier, service-level agreement (SLA), latency requirements, and so on. The residential offerings act as an independent silo with the potential for significant operational inefficiencies.

If operators continue to operate their residential broadband products as separate silos they are at risk of failing to achieve the operating margins necessary to remain competitive due to a duplication of efforts and lack of focus. This could allow telcos that *can* offer multiple services on a converged platform to take their market share.

2.1.2. Mobility Traffic

After significant growth and success in video, data, and land line voice services, wireless is the next frontier for cable. Quarter over quarter, the US cable mobile virtual network operator (MVNO) business continues to see mobile subscriber growth (**ref. [1]**). As of 2021, less than four years after the launch of the first MVNO by Comcast, three US cable MVNOs combined have amassed millions of customers. Comcast and

Charter MVNOs utilize Verizon as the mobile network operator (MNO). Recently, Cox Communication also demonstrated an interest in starting an MVNO. The momentum is there, and executive leadership at cable companies has consistently shown strong support for and interest in growing the wireless business.

Comcast and Charter have been signaling for some time that they intend to build Citizens Broadband Radio Service (CBRS) based mobile networks in their existing cable footprints in an effort to reduce the compensation given to Verizon and other MVNO partners for use of their networks **(ref. [2])**. Their MVNO operations were intended as a way to build a subscriber base and brand in advance of owning their own wireless networks, despite incurring consistent Earnings Before Interest, Taxes, Depreciation, and Amortization (EBITDA) losses for an extended period of time. Cox—which had entered the wireless space a decade ago, only to exit after disappointing results—has signaled its intention to re-enter the wireless market through the purchase of a significant number of CBRS licenses across its cable footprint. The largest cable operators already have a dense network of millions of Wi-Fi hotspots that can easily be turned into 5G small cells to create a mobile network.

In contrast to the US, virtually all of Canada's largest cable and telco operators offer residential wireline, business wireline, and mobile services on their own infrastructure (**ref. [1]**). Rogers, Canada's largest cable and mobile operator, has been offering mobile services since 1985, with Vidéotron launching its wireless services in 2010, and Shaw acquiring Wind Mobile in 2015. The Canadian market faces strong competition from Canada's large incumbent telco operators, which have invested heavily in fiber to the home, connecting more than 60% of their broadband homes directly to fiber, and leveraging a robust Radio Access Network (RAN) sharing agreement to minimize their infrastructure costs.

Today's cable operators are tomorrow's mobile operators **(ref. [3])**, and behind every efficient wireless network, there must be a converged wireline network. The industry has reached the consensus that optical DWDM transport's capability of transmitting a large number of information streams simultaneously over a single optical fibre makes it the best possible solution for the most demanding xHaul needs. Operating separate residential and mobile optical cores is prohibitively costly for operators, while converging optical networks into a single, common core will significantly cut costs and improve network agility. A converged network will manage peak traffic load more effectively, improving both speed and performance.

2.1.3. Business Traffic

Operators of all sizes have found growing opportunities to take part in the lucrative and ever-evolving business services space. While the battle for the wireless dollar has gotten most of the ink in the trade press recently, the traditional telcos and cable MSOs are waging a war on another, less publicized front: the small and medium-sized business market. The telcos are moving aggressively into the residential video services market, where cable operators currently enjoy a high penetration rate. To compensate for the loss of residential video customers, operators have chosen to pursue perennially underserved small- and medium-sized businesses. And they are winning over these customers with Carrier Ethernet.

Operators typically deliver business services over disparate architectures. A high-security application, for example, runs over its own DWDM wavelength. Schools, financial institutions, and telephony that require fast switching (50-millisecond) protected services would be typically served via a SONET network. And low-priced Ethernet services run over shared bandwidth tunnels using Layer 2 or 3 architectures. In addition, business customers may favour symmetrical services, while residential customers are well served by DOCSIS, which is currently asymmetrical, typically offering 5:1–10:1 ratios in the downstream/upstream bandwidth. With business customers driving cable operators to better utilize existing fibre plant and invest in new FTTP deployments, the trend is to offer multiple broadband pathways over existing fibre using DWDM.

2.2. Convergence of Network Layers

The layer-1 photonic network, the layer-2 carrier Ethernet network, and the layer-3 IP/MPLS network are designed, built, and expanded independently, and are often maintained by teams working in complete silos. This results in each layer being over-provisioned to cope with uncertain network demands (**ref. [4]**) and the networks being overprotected due to overlapping and redundant resiliency schemes on each layer.

This also results in additional Capital Expenditure (CAPEX) and Operational Expenditure (OPEX). In terms of CAPEX, transponders, switch ports, and router ports—the most expensive parts of the core of the network—become layered and the siloed infrastructure relies on large volumes of line cards for traffic hand-off between networking layers. Additionally, OPEX is increased because multiple teams are required to handle the provisioning and maintenance of a single service. Inefficiencies can also be seen in other areas, including a loss of time because of the need to deploy multiple processes and teams, and high complexity due to multiple independent network management systems (NMS) associated with each network layer, which results in high costs and poor network resource utilization and leads to poor monetization.

Multiple trends are driving operators to integrate DWDM optical transport with Ethernet/client OTN-SONET technology, primarily in metro and more recently in long haul networks. These trends include the demand for more bandwidth and greater agility with key applications, including cloud connect services, fixed broadband aggregation, mobile xHaul, and SONET migration.

2.3. Six Mandatory Attributes of The Unified Transport Platform

Operators have been asking for simpler, more cost-efficient converged network architectures that will enable them to concentrate on innovating revenue-generation services (**ref. [5]**). Based on Shaw's modeling, the converged agile optical framework must have following six attributes:

1) High bandwidth densification per rack space

For mobility operators, network densification means adding more cell sites to increase the amount of available capacity, which demands high bandwidth densification per rack space of the unified transport platform that contains Mobility xHaul, with a minimum threshold of 800Gbits/s per 1 Rack Unit (RU).

2) Building-block-approach scalability

Scalability is the ability of a system to expand without major modifications to its architecture. The concept implies the ability of a network system to accommodate a sudden increase in traffic volume gracefully and rapidly. After meticulous modeling by Shaw, it was found that the unified transport platform must be able to handle the unpredictable growth pattern of residential, mobility, and business traffic with building-block-approach scalability.

3) Full-set client interface support

The unified transport platform is required to support residential, mobility, and business services. Such a platform must be able to provide all required standard client interfaces, including Ethernet, OTN, and SONET.

4) Full-featured General Communication Channel (GCC0)

The unified transport platform has to integrate business services. Because of the power and rack space restrictions of business customer premises, installing an out-of-band management switch is usually not a viable option, or at least a very expensive one. Shaw's assessment concludes that in-band full-featured

GCC0 is a vital attribute of the converged agile optical framework. For a detailed explanation of GCC0, please go to Section 3.4.

5) Zero Touch Provisioning (ZTP)

Because the unified transport platform converges all service types, deployments are almost always multi-sited. Traditional turn-up based on manual initial configuration requires highly trained field personnel who are often difficult to find, and therefore multi-site deployment without ZTP tends to be slow, costly, and highly error prone. However, Shaw's analysis indicates that ZTP is a mandatory attribute for the unified transport platform.

6) Flex-grid

The concept of Flex-grid is related to the first attribute—high level bandwidth densification per rack space. To achieve high bandwidth densification, operators have to look beyond 100Gbps using a fixed grid. Only Flex-grid DWDM systems support 400G+ with high spectral efficiency. While fixed grid DWDM systems can still support "fat channels", it is at the expense of significantly lower spectral efficiency. Spectral efficiency is an important measure of how effectively or efficiently a fiber network transmits information.

3. Classic Optical Network

Classic Optical Networks are based on Legacy DWDM technology, which are hostile environments for service convergence. It is exceedingly difficult to implement service convergence in a classic optical network, if not entirely impossible. Legacy DWDM networks have the following defining features:

- It uses a chassis-based platform with low level bandwidth densification per rack space.
- Its scaling is inefficient.
- It is limited to SONET and 10GbE/40GbE client interfaces.
- It does not support line-side GCC0.
- It does not support ZTP.
- It uses a Fixed Grid (50 – 100GHz).

All the above characteristics present severe restrictions for network convergence. These characteristics are discussed in detail in the following sections.

3.1. Chassis-Based Platform and Low-Level Bandwidth Densification

All classic DWDM equipment manufacturers have chosen a chassis-based platform. In such a platform, each DWDM component is devised as a card to be installed in an empty slot of the chassis. For example, Arrayed Waveguide Gratings (AWG); Wavelength Selective Switches (WSS); Erbium-Doped Fiber Amplifier (EDFA); and most transponders/muxponders are typically constructed as individual cards. Figure 1 shows a typical 14 RU monolithic chassis.

A transponder/muxponder card is commonly installed vertically but still roughly takes about 1 RU. A typical legacy transponder/muxponder only has the bandwidth density of up to 100Gbps per 1RU. 100Gbps is significantly less than the minimum requirement of 800Gbit/s per 1 RU for a converged platform.

The transponders/muxponders work in pairs in a point-to-point system. At a specific source, the destinations for residential service, mobility service, or business service are usually different. Therefore, in a classic system this would require three pairs of transponder/muxponder cards at the source and destination, resulting in 3RU per site.

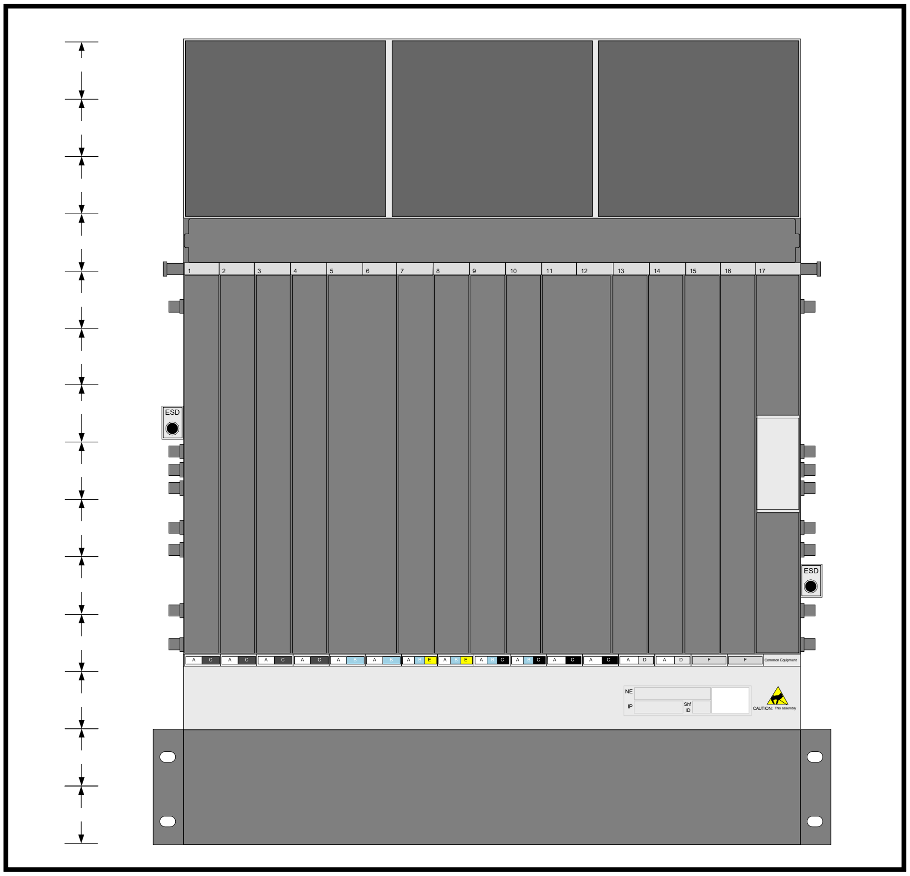

Figure 1 – Chassis Based Platform

3.2. Inefficient Scaling

The legacy monolithic-chassis DWDM systems are able to handle future traffic volumes on day one, but also require a large up-front investment for capacity that may not be needed for several years. This is particularly true for chassis equipped with specialized hardware such as terabit scale switching fabrics, as the fabric is part of the initial installation, even though transponders may be added over time. As for operators, space and power are also always at a premium, and classic optical networks require large space and power budgets at the beginning of deployment. The issue of inefficient scaling is particularly acute for remote mobility sites and business customer premises.

In order to achieve service convergence for mobility and business, MSOs need to examine modular DWDM systems with smaller footprints. For example, 1RU pizza-box blades that offer the advantage of efficient scaling for all operators.

3.3. Client Interface Limitations

Classic DWDM systems usually only support SONET client interfaces and low-rate Ethernet client interfaces. For a legacy transponder/muxponder, only three client formats are typically supported, including SONET OC192, 1GbE, and 10GbE, with only a few legacy DWDM vendors offering 40GbE and 100GbE client interface support. While these three client interfaces are often sufficient for residential traffic, they fall short of the requirements for mobility and business traffic, which both require 100GbE and OTU4. In addition, many business customers also need OTU2/2e.

Because of the limited client formats, MSO operators are forced to move mobility and business traffic away from legacy optical networks and onto IP/MPLS networks. The traffic migration to IP/MPLS is considered suboptimal due to DWDM offering the most efficient use of optical fibre throughput capacity, as well as lower price points. As far as the client ports are concerned, classic optical network focuses on residential customers, while overlooking mobility and business customers. The client interface format limitation is a serious roadblock to service convergence.

Table 3 and Table 4 list the key parameters for OTN Frames and SONET OC192, for reference.

Table 3 – OTN Frames

OTUk	Bit Rate (Gbps)	Payload Rate (Gbps)	Payload Types
OTU4	111.809973	104.355975	100GbE
OTU2e	11.095730	10.356012	10GbE LAN, 10GFC (TTT)
OTU2	10.709255	9.995277	10GbE WAN, 10GbE LAN (GFP-F), STM-64/STS-192

Table 4 – SONET OC192

Acronym	Bit Rate (Gbps)	DS0	DS1	DS3
OC-192	9.9953	129,024	5,376	192

3.4. Client-Side GCC0 Only

General Communication Channel 0 (GCC0) is used for overhead communication between network nodes. GCC0 has operations, administration, and management (OAM) functions such as performance monitoring, fault detection, signaling and maintenance commands in support of protection switching, fault sectionalization, service-level reporting, and control plane communications. GCC0 has two bytes within OTN overhead. GCC0 is terminated at every 3R (re-shaping, re-timing, re-amplification) point and used to carry management information and GMPLS signaling protocol.

Legacy DWDM transponders/muxponders only support client-side GCC0. Optical Transport Network (OTN) standards define the GCC0; it is an in-band channel used to carry management information between transponder/muxponder pairs. GCC0 is a critical method to manage an optical device for a business customer site. Because of power limitations and rack space restrictions, it is very difficult to install a layer-2 management switch for out-of-band management in a business customer site. Figure 2 represents a schematic diagram on client-side GCC0.

Although client-side GCC0 is a useful feature, it is considered as only a partial implementation of GCC0. In client-side GCC0, the network interface device (NID) at the far-end site encapsulates management information into the GCC0 bytes of each OTU2/OTU2e frame. The NID sends out the OTU2/OTU2e frames from its network-side port, which has a grey (black and white) wideband pluggable, most likely at 1310nm. When the near-end transponder/muxponder receives the frames on its client-side port from the far-end NID, it decapsulates the GCC0 bytes from the OTU2/OTU2e frames and recovers management information from the GCC0 bytes.

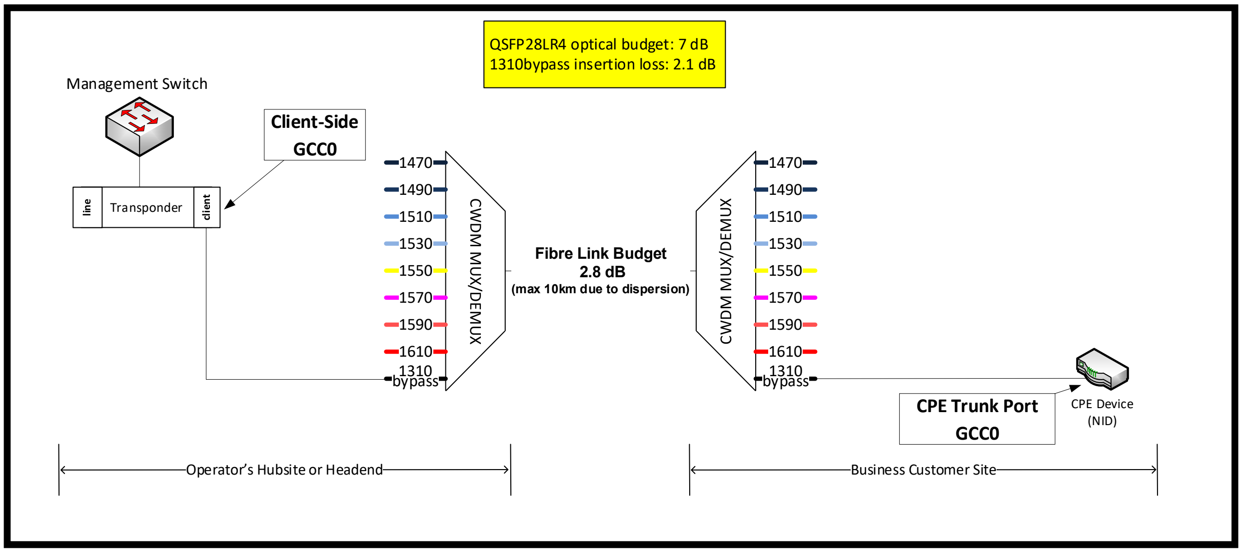

Figure 2 – Client-Side GCC0

While a management switch is necessary for the near-end site, it can be skipped for the far-end site when client-side GCC0 is supported on the near-end transponder/muxponder. It should be noted that the limiting factor for client-side GCC0 is the distance between the near-end and far-end. If there are only dark fibres between the near-end and far-end, the maximum distance for client-side GCC0 is approximately 10km, depending on the fibre rating. If CWDM is inserted between the near-end and far-end, the maximum distance is reduced to 5km. Though client-side GCC0's maximum distance is adequate for residential service and mobility service, it is insufficient for most cases of business service.

In summary, due to the aforementioned factors, legacy DWDM devices' lack of line-side GCC0 support considerably hampers the potential for service convergence.

3.5. Manual Turn-up Initial Configuration

As the name suggests, the goal of Zero Touch Provisioning (ZTP) is to install a networking appliance without the need for local configuration by a trained individual, making it possible for a new or replacement device to be sent to a site, physically installed, and powered up by a locally present employee without technical skills. The ZTP feature carries out the configuration and connection to the management system.

Legacy photonic-layer devices do not support ZTP, and therefore extensive manual configuration is required to turn up a new device in a classic optical network. Manual configuration is laborious, prone to errors, costly, and time-consuming. In this scenario, an individual with basic configuration skills and a laptop has to go onsite and configure the device for basic operation before its configuration can be completed remotely using the central management system **(ref. [6])**. Alternatively, the device can first be sent to a central location where it is staged before being sent to its final location for installation. This is also costly as it requires shipping the device twice, which could potentially mean passing through customs twice. And there is the risk of accidentally shipping an appliance with an IP-address destined for a different site.

For residential service, ZTP is helpful but it is not always needed. However, for mobility service at a remote site and business service at a customer premise, ZTP is critically important. With ZTP, a device can be shipped directly from any warehouse to the remote site or customer premise and installed as soon as it arrives and be up and running within minutes of installation. This dramatically reduces the lead time, time spent on an installation, and number of configuration errors, which are significant benefits for mobility and business services.

Because legacy DWDM's chassis does not support ZTP, operators must look into next generation optical devices for service convergence. ZTP is becoming more widely supported as next generation optical equipment vendors realize their equipment can now be installed by local techs who may not be trained on provisioning and configuration.

3.6. Fixed-Grid Photonics

Older generation optical networks are based on 100 Ghz or 50 Ghz spaced photonic systems. These gridded networks can offer forty-eight or ninety-six fixed grid optical channels within the total 4800 Ghz C-Band Spectrum. This fixed grid spacing is based on the International Telecommunication Union (ITU) standard and has been the norm for most photonic systems for over 20 years. These systems use passive, ITU grid filters to provide wavelength ingress/egress. Second and third generation optical transponders running at 35 GBaud typically require 37.5 Ghz of optical spectrum, which fits perfectly into these ITU gridded filters. This version of the DWDM network has served the industry well for years.

As next generation transponders are moving towards 400GbE, they bring with them the need for larger per channel spectrum. This larger channel spectrum requirement exceeds that of what is available on these ITU gridded filters. The Fixed-Grid paradigm cannot support them, making this method obsolete.

4. Universal Aggregation

Universal Aggregation is a converged networking approach that enables the aggregation of traffic from SONET, Ethernet, and OTN services using shared fibre and the same optical DWDM platform. Under the Universal Aggregation paradigm, the DWDM systems are considered as both core technology and access technology. Behind the SONET, IP/Ethernet, and OTN services, the traffic sources include residential, wireless, and commercial.

In a traditional network, it is no wonder that performance and reliability are increasingly valued as operators must configure, shape, and optimize each service type individually on separate platforms, with each platform having its own vendor-provided technical support. Sometimes, this means dealing with multiple platforms with the same vendor. However, in the majority of cases, operators must deal with several different platforms from separate vendors. Managing multiple interconnected hardware deployments across diverse vendor environments significantly increases operational cost and makes troubleshooting a very complex task. Traditional separated network architectures are simply unable to scale quickly enough, negatively impacting time-to-market for new service revenues.

In traditional network architectures, each of the three network layers has its own network silo. There is growing evidence against siloed networks, contributing to its inherent weakness as part of the network architecture. Disadvantages of network silos include:

- An inability to better utilize or share the silo's resources.
- The extended period of time needed to deploy, manage, and upgrade in the siloed environment.
- Rising operating costs, given each silo's unique processes.
- The redundant building of each silo's own protection systems to meet availability requirements.

Recent developments in coherent optical technologies provide greater opportunities for operators to move toward comprehensive network convergence. In network converged systems, residential, mobility and business services coexist over the same fibres, forwarding tables, data planes, servers, etc. However, getting to this point will take time and careful planning. To help support the shift towards full network convergence, operators need a robust platform framework.

Universal Aggregation breaks up the aforementioned silos by converging all into a single, unified platform. This means aggregating mobile, residential, and enterprise traffic vertically, while at the same time aggregating SONET, Ethernet, and OTN services horizontally. Universal Aggregation enables a simple, compact, scalable, and efficiently converged infrastructure with coherent optic transport, saving operators from having to build different networks for different services. By supporting all services on a unified platform, Universal Aggregation reduces operational cost considerably and expands competitiveness substantially.

Table 5 – Universal Aggregation

	Traffic-Type Convergence	**Network-Layer Convergence**
Universal Aggregation	Residential	Tier 1 - Optical Photonic
	Mobility	Tier 2 – Carrier Ethernet
	Business	Tier 3 - IP/MPLS (Future)

Shaw has successfully implemented Universal Aggregation on a blade-centric architecture. It is a converged framework spanning access, transport, packet, and optical, and combining fast, simple deployment with a wide range of options for right-sizing capacity. Residential, mobility and commercial services are aggregated into one unified platform supporting SONET, Ethernet, and OTN. Shaw's implementation of Universal Aggregation has the following technical characteristics:

- High Bandwidth Density per Rack Space
- Efficient Scalability
- Full-Set Client Support - SONET, Ethernet, and OTN
- Line-side GCC0
- Zero Touch Provisioning
- Flexible Grid

Shaw's successful experience indicates that the most efficient unified technical platform of Universal Aggregation is a modular, coherent DWDM system with Ethernet/OTN/SONET client interfaces. In the section below, we will discuss the blade-centric platform in detail.

4.1. Overview of Blade-Centric Platform

As mentioned previously in Section 3, over-engineered, restrictive, chassis-based systems are not adaptable in the converged environment. A typical chassis is shown on the left side of Figure 3, and a typical modular blade is shown on the right side. The typical DWDM chassis is of 14RU, while the typical modular pizza-box is only of 1RU.

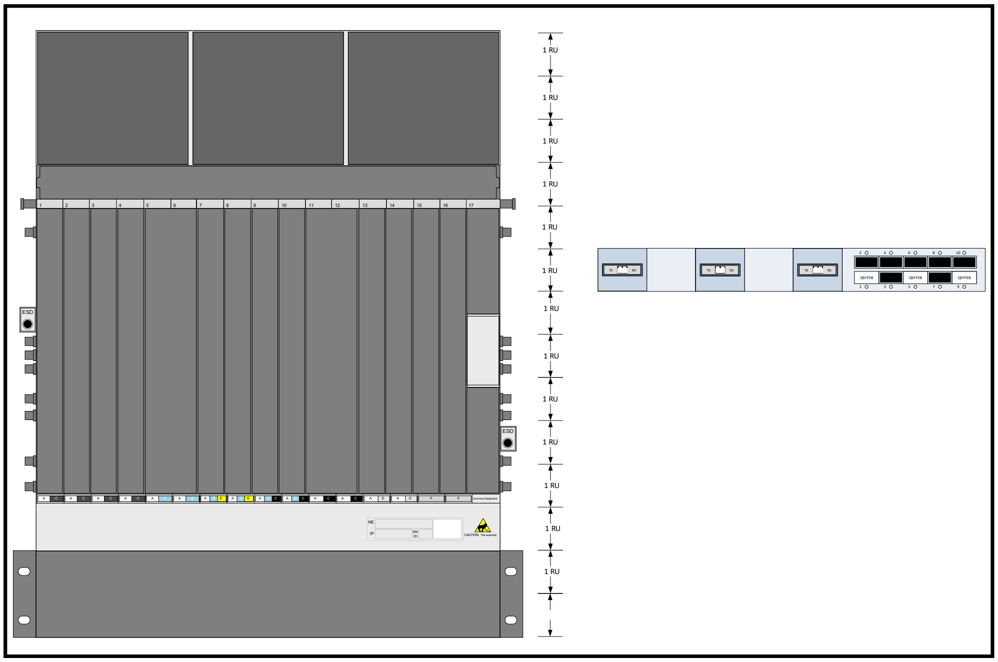

Figure 3 – Chassis vs Blade

The modular product line divides DWDM functions into three areas: ROADM blades, Splitter/Coupler blades, and Transponder blades, as shown in Figure 4. The ROADM-on-a-blade provides wavelength selective switching and amplification. While Figure 5 only shows two ROADM blades, our platform currently supports a maximum of eight ROADM blades (one main blade and up to seven tributary blades) that can be interconnected as a single Network Element (NE). A Splitter/Coupler blade has the main function of channel add/drop and replaces legacy static filters in CDC (Colorless, Directionless, Contentionless) configurations. Splitter/coupler blades are interconnected with the ROADM blades as one NE, and the Transponder blade is a transceiver with very high port density on both network and client sides, which in the future could be part of a single NE as well.

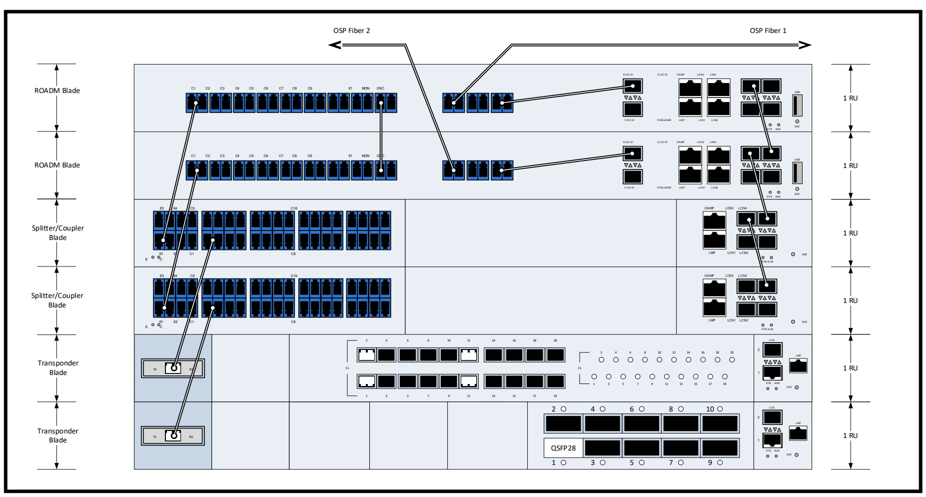

Figure 4 – ROADM Blades, Splitter/Coupler Blades, and Transponder Blades

This isolation approach allows vendors to rapidly implement and deploy new features in each secluded functional area when and as much as needed. Because of the modularity, vendors tend to be much more responsive to feature requests from operators.

4.2. High Network Densification

The blade-based platform is a simpler, more agile approach to creating an optical network for Universal Aggregation of residential, mobility, and business traffic. The blade paradigm is the only feasible platform for universal aggregation at this point in time due to the exceptionally high bandwidth density per 1RU. Figure 5 shows a typical implementation of Universal Aggregation on a modular platform. The transponder blade at the bottom of Figure 5 aggregates residential, mobility, and business into one 1RU pizza-box at the source. To achieve the same goal, the typical chassis-based platform would have to use three transponder cards, which means approximately three times of power consumption and rack space.

The transponder blade's line port bandwidth capacity can either be 100Gbits/s or 200Gbits/s, depending on the modulation on the line port, which ultimately depends on the fibre span loss and OSNR (Optical Signal to Noise Ratio) of the system. If the distance between the source and destination is shorter than 230km, it is likely that 200Gbits/s would suffice. If the distance is longer than 230km, however, capacity would have to be 100Gbits/s.

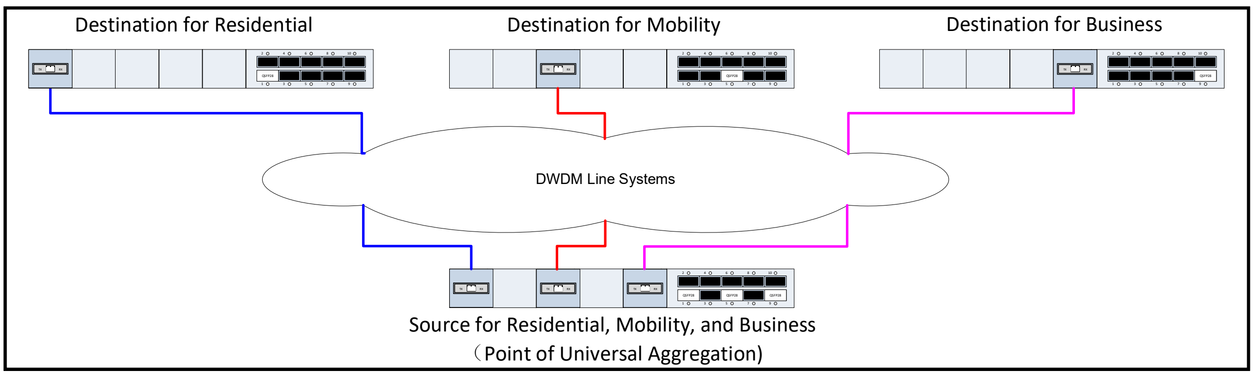

Figure 5 – Universal Aggregation on A Modular Platform

Figure 6 shows a typical transponder module with QSFP28 client interfaces. It has five line ports, and each line port has a maximum bandwidth capacity of 200Gbits/s. The total maximum capacity for the 1RU module is 5x200Gbits/s, which is 1Tbits/s.

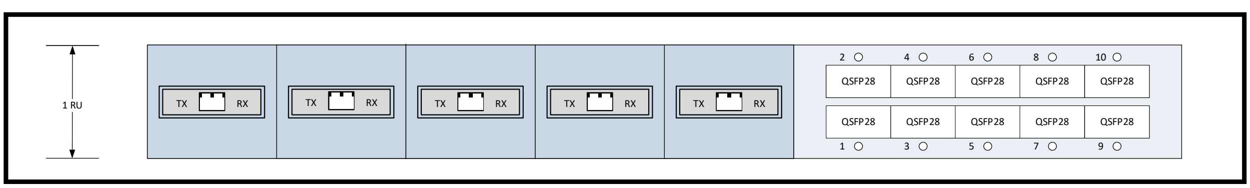

Figure 6 – Maximum Bandwidth Capacity of a Transponder Blade

While 1Tbits/s bandwidth capacity per pizza-box is considered extraordinary for a 1RU rack space, the power consumption is about 224 watts, which is on par with shelf-based transponder cards. In the blade-centric paradigm, the transponder blades consume a larger percentage of total power as compared to ROADM blades and splitter/coupler blades; this fact enables vendors to focus on independent improvements in power for transponder blades without the constraints of traditional converged shelves.

Most operators should be able to achieve an overall 30% power savings per 100Gbps bandwidth compared to legacy systems.

4.3. Efficient Scalability with Building-block-like Infrastructure

Traditional chassis-based platforms represent a significant initial capital outlay and the chassis backplane is inherently inflexible. On the contrary, a blade-based modular system requires a smaller, granular initial investment while allowing for a pay-as-you-grow approach. With single rack unit sized blades, the system is also space efficient and flexible.

The paradigm offers the advantage of efficient scaling for all operators. Scaling efficiently is one of the key network requirements for operators in the era of Universal Aggregation. The building-block approach to hardware allows for a low initial spend for year-one deployments with the ability to grow incrementally as traffic increases and more capacity is required. Many converged, monolithic-chassis DWDM systems, by contrast, are able to handle future traffic volumes on day one, but also require a large up-front payment for that capacity even when the capacity may not be needed for several years. This is particularly true for chassis equipped with specialized hardware such as terabit scale switching fabrics, as the fabric is part of the initial installation, even though transponders may be added over time.

4.4. Full Set Client Support - SONET, Ethernet, and OTN

Each service type demands a unique set of client interfaces from DWDM transponders. For residential landline telephony, SONET OC192 is needed. For residential Internet, both 10GbE and 100GbE have to be supported to communicated with CMTS. For mobility xHaul, 100GbE and OTU4 are required. Business customers request the widest range of client interfaces: SONET OC192, 10GbE, 100GbE, OTU4, OTU2/2e.

Table 6 summarizes the client interface formats required by each service type.

Table 6 – Service Type/Client Interfaces

Service Type	Client Interface Formats
Residential	SONET OC192,10GbE, 100GbE
Mobility	100GbE, OTU4.
Business	SONET OC192, 10GbE, 100GbE, OTU4, OTU2/2e

Chassis-based legacy platform only supports a limited subset of the client interfaces listed in the above table, ruling out monolithic chassis as a platform for Universal Aggregation. However, the blade-centric platform supports all client interfaces in Table 6. As listed in Table 7, Module A supports SONET OC192, 10GbE, and OTU2/2e, while Model B supports 100GbE and OTU4. Two models of modular blades are also shown in Figure 7, below.

Table 7 – Model/Client Interfaces

Model Type	Client Interface Formats
Model A	SONET OC192,10GbE, OTU2/2e
Model B	100GbE, OTU4.

Comparing Tables 6 and 7 confirms that together, Models A and B will provide full set client interfaces for residential, mobility, and business services, and that modular blades are ideal for Universal Aggregation of these three types of traffic.

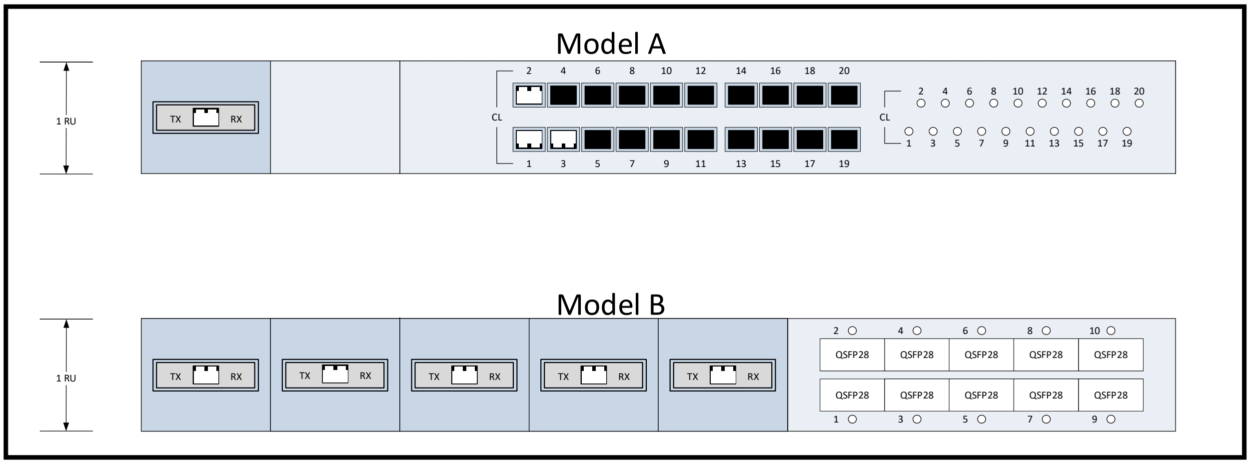

Figure 7 – Model A and Model B

4.5. Line-Side GCC0

As discussed in Section 3.3, legacy DWDM systems only support client-side GCC0. Blade-based next-generation DWDM systems support both client-side GCC0 and line-side GCC0.

For line-side GCC0, the transponder/muxponder in the far-end site encapsulates management information into the GCC0 bytes of each OTU4 frame. The OTU4 frame then goes through the electrical-to-optical conversion and is sent out of the line port of the far-end transponder/muxponder on a specific DWDM wavelength. Upon receiving each OTN frame by the near-end transponder/muxponder, the GCC0 bytes are extracted and the management information is retrieved. In using line-side GCC0, there is no need for a management switch in the far-end site, although a management switch is still necessary for the near-end site. The maximum distance for line-side GCC0 is about 80km depending on the fibre rating, even with CWDM inserted. This distance is adequate for most business services.

Figure 8 below is the schematic diagram of line-side GCC0.

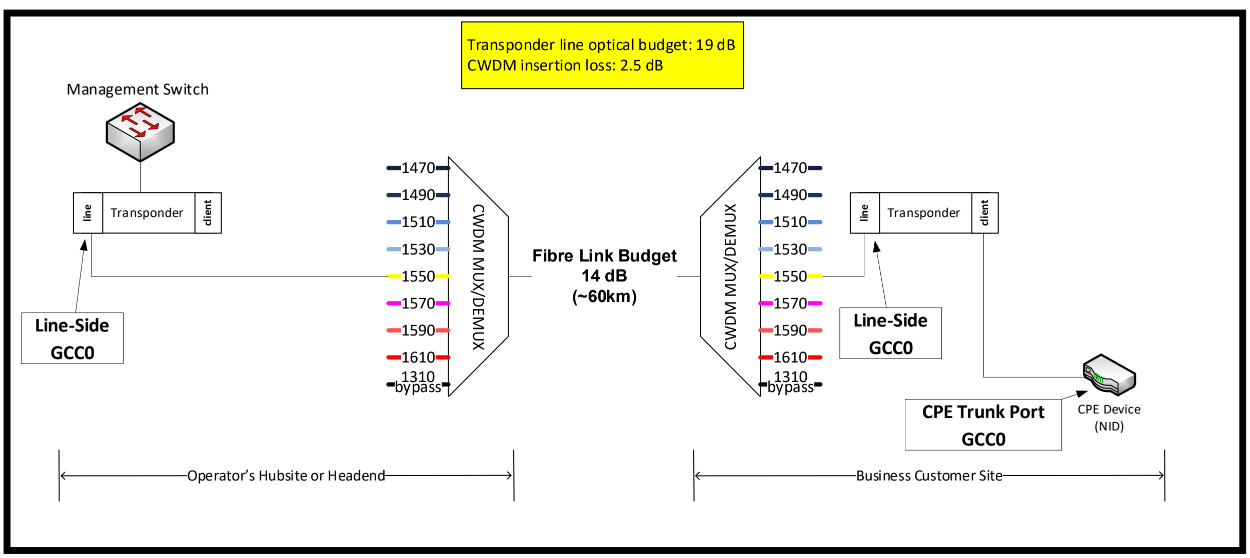

Figure 8 – Line-Side GCC0

4.6. Zero Touch Provisioning (ZTP)

The blade-centric next-generation platform also supports ZTP. Shaw has set up several pre-staging centres in various regions. Before shipping a blade to the site, Shaw engineers will pre-provision the blade and save the configurations to the USB key attached to the blade. The blade will then be shipped with the USB key to local technicians. When the local technicians receive the blade and USB key, all that is required is a plug-and-play. Figure 9 shows a typical set-up of a pre-staging lab.

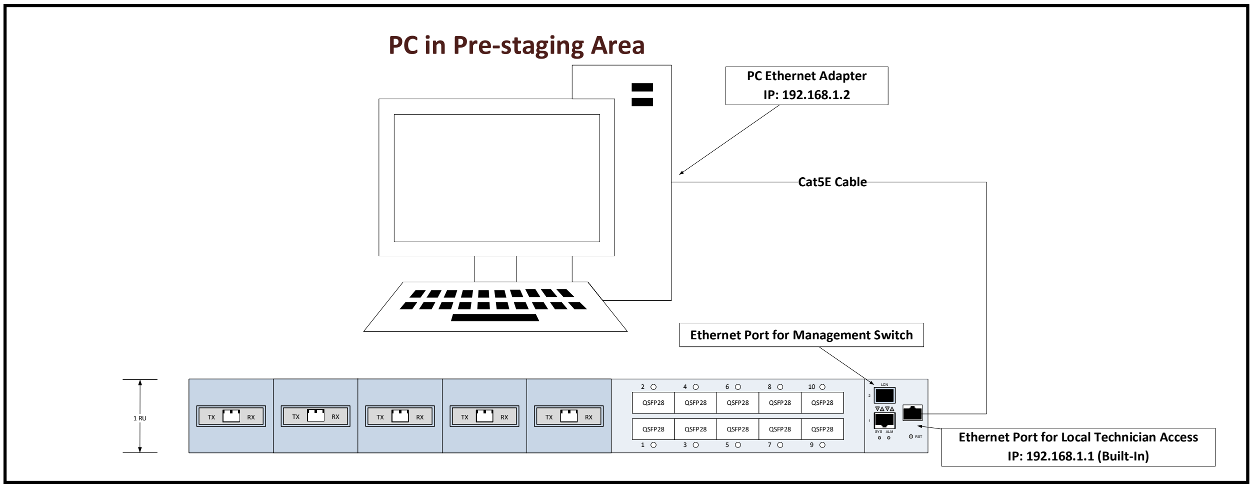

Figure 9 – Zero Touch Provisioning

4.7. Flexible Grid

Flexible Grid, also written as Flex-Grid, is in contrast to Fixed-Grid which was discussed previously in section 3.6. Flex-Grid combines two concepts together: finer wavelength granularity and the ability to join adjacent wavelength slots together to form arbitrary sized channels. From hardware perspective, Liquid Crystal on Silicon (LCoS) based Wavelength Selective Switch (WSS) makes a ROADM Flex-Grid.

Universal Aggregation requires 400Gbps transponder/muxponder because residential, mobility, and business each normally requires 100G client bandwidth. The blade-centric next-generation modular platform provides a 400G transponder blade with Flex-Grid. With the 400G blade, operators will be able to transport more than twice the information within the same spectrum on Fixed-Grid. Operators can actually further increase the spectral efficiency by 25% with the use of Flex-Grid by allowing the subcarriers in DWDM superchannels to be squeezed more closely together. While Flex-Grid drastically improves spectral efficiency, the modular 400G transponder pizza-box does not require Flex-Grid as absolute pre-requisite, it will also work on existing legacy Fixed-Grid for backward compatibility.

The superchannels enable operational scaling by allowing operators to turn up optical capacity in larger increments with the same effort. Going forward, next-generation optical transport networks will need to make the most use of the flexibility of advanced coherent modulation technologies. More advanced flexible coherent modulations will support a wide range of modulations tailored to specific applications of residential, mobility, and business. Flexible DWDM grids will enable more efficient and flexible use of optical spectrum to maximize capacity.

5. Conclusion

This paper has outlined the benefits of Universal Aggregation on a modular platform, which is a promising paradigm for Cable operators.

Under the paradigm of Universal Aggregation, it is unnecessary for operators to build separate networks for different traffic types of residential, mobility and business. By aggregating all three types of service on a unified platform, Universal Aggregation reduces operational cost considerably and expands competitiveness substantially.

In traditional systems, each network layer has their own network silo. In the Universal Aggregation model, these silos are all broken down. The consistent unified platform across optical, Ethernet, and IP/MPLS will smash down all network silo walls.

Shaw's Universal Aggregation utilizes a blade-centric platform with the following features:

- It has astoundingly high bandwidth density per 1RU with a 90% increase in available system bandwidth comparing with legacy systems.
- It uses a building block approach to scalability, which only requires a small initial spend while allowing for continuous growth through pay-as-you-go.
- It supports 100GbE, OTU4, 10GbE, and OTU2/2e over 100Gbps and 200Gbps wavelengths.
- Its line-side GCC0 offers in-band management for remote mobility sites and business customer sites.
- It implements a Zero Touch Provisioning (ZTP) system that simplifies operations and allows engineers to automate most of the deployment tasks.
- It supports Flex-Grid which increase the spectral efficiency by 25% for 400G blades.

Abbreviations

3R	Re-shaping, Re-timing, Re-amplification
10GbE	10 Gigabits Ethernet
100GbE	100 Gigabits Ethernet
CAPEX	Capital Expenditure
CBRS	Citizens Broadband Radio Service
CPE	Customer Premise Equipment
DSP	Digital Signal Processor
DWDM	Dense Wavelength-Division Multiplexing
EBITDA	Earnings Before Interest, Taxes, Depreciation, and Amortization
EDFA	Erbium-Doped Fiber Amplifier
GCC	General Communication Channel
GMPLS	Generalized Multi-Protocol Label Switching
IP	Internet Protocol
IP/MPLS	Internet Protocol/ Multi-Protocol Label Switching
ITU	International Telecommunication Union
LCoS	Liquid Crystal on Silicon
MSO	Multiple-System Operators
MVNO	Mobile Virtual Network Operator
NE	Network Element

NID	Network Interface Device
NMS	Network Management System
OAM	Operations, Administration and Management
OPEX	Operational Expenditure
OSNR	Optical Signal to Noise Ratio
OTN	Optical Transport Network
OTU2/2e	Optical Transport Unit 2/2e
OTU4	Optical Transport Unit 4
RAN	Radio Access Network
ROADM	Re-configurable Optical
RU	Rack Unit
SONET	Synchronous Optical Networking
Tbits/s	Terabit Per Second
UA	Universal Aggregation
WSS	Wavelength Selective Switch
ZTP	Zero Touch Provisioning

Bibliography & References

[1] *Jennifer Andréoli-Fang, John T. Chapman*, *"Cable and Mobile Convergence - A Vision from the Cable Communities Around the World", SCTE•ISBE, 2020.*

[2] *Jeff Heynen, "Cable's Fixed-Mobile Convergence Future", www.delloro.com, 2020*

[3] *Jeff Baumgartner, "Cable's evolutionary path leads to mobile, convergence", www.lightreading.com, 2020.*

[4] *Fernando X. Villarruel, "Framework for Convergence of Services on The MSO Network - Using the Principles of Network Slicing", SCTE•ISBE, 2020.*

[5] *Kevin Bourg, Sergey Ten, Peter Wigley "An Overview Of Optical Architectures Necessary To Achieve 5G's Key Performance Indicators, SCTE•ISBE, 2020.*

[6] *Timothy Maenpaa, "Addressing Unrelenting Growth In Backbone Fiber Systems Using Next Generation Photonics And Automation", SCTE•ISBE, 2020.*

[7] *Chad Andrews, Steve Canepa, Bob Fox, Marisa Viveros, "The end of communications services as we know them - How 5G and edge computing will help define who wins in the booming digital economy", IBM Institute for Business value, 2021*

[8] *Koby Reshef, "Paving the Path to 400G Migration", www.pipelinepub.com, 2021.*

[9] *National Cable & Telecommunications Association, "The Asymmetric Nature of Internet Traffic", www.ncta.com, 2021.*

Acknowledgements

I would like to express my thanks and gratitude to Damian Poltz (SVP, Wireline Technology & Strategy) and Felipe Arroyo (Manager, Optical Networks), who provided the opportunity to undertake this project, as well as their guidance and input throughout. I would also like to express my gratitude to Lili Ti (Advisor, Communications) for her help with editing this paper.