Cox Next Generation 400G IP+OLS Architecture for Maximum Network Optimization and Cost Benefits

A Technical Paper prepared for SCTE by

Saurabh Patil
Principal Architect
Cox Communications
6305 Peachtree Dunwoody Rd, Atlanta, GA, 30328
404-269-8364
Saurabh.patil@cox.com

Jason Bishop
Senior Manager Design Engineering
Cox Communications
6305 Peachtree Dunwoody Rd, Atlanta, GA, 30328
404-269-3994
Jason.Bishop@cox.com

Table of Contents

List of Figures

1. Introduction

Cox Communications has extensive optical fiber infrastructure comprising over 25,632 route miles of National Optical backbone and in Metro Markets. Cox Communications relies on Backbone to connect all their Metro markets, urban centers, and peering centers to keep everything running smoothly and efficiently. It's an essential tool that ensures seamless connectivity and enables them to serve their customers better. Fiber is an integral part of Cox Communication's assets, and we in Engineering work constantly to optimize and strategically utilize our fibers for the lowest Total Cost of Ownership (TCO) and better customer experience to transmit Terabits of data.

As we approach the Shannon limit to exhaust fiber capacity, what other options do we have as a service provider to keep up with ever-increasing bandwidth demands and capacity pressure from innovation and customer consumption? There is another 15-20% scope to improve spectral efficiency before fiber exhaustion, but that is just a stop-gap for 1-2 years for capacity growth.

This problem needs not only engineering innovation from optical vendors and Internet Protocol (IP) routers but also from service providers to rethink how we architect our network and use surgical techniques to use appropriate technology in certain parts of the networks where it's needed.

In Cox, we use embedded highly expensive coherent transponders extensively in Backbone, metro, and some cases in access networks as a solution for adding capacity, even if it's short distances. The saying "When you only have a hammer, every problem can start to look like a nail" is undeniably true. Having limited tools or resources can lead to a narrow-minded way of thinking. It's crucial to acknowledge that there are multiple approaches to a problem, and sometimes, it's necessary to step out of our comfort zone and explore unfamiliar tools and routines to find the optimal solution. Here, we use high-power transponder solutions for all our problems with heavy reliance on over-engineering. The consequence is paying a very high significant cost and increasing power consumption. With the introduction and maturity of new coherent pluggable optics, we at Cox Communications are learning to use pluggable optics in Backbone, Metro, and Access to leverage lower pricing and new capabilities. Utilizing these new coherent pluggable optics on high-capacity short spans can offer exceptional spectral efficiency at a reasonable cost. This approach is a strategic method to enhance performance while staying within budget.

In this paper, we will discuss how we are re-architecting Cox's Backbone and metros to overcome Shannon limits and use coherent pluggable optics in Cox's Next generation Open Line System (OLS) that supports both pluggable (400G ZR, ZR+, ZR++) and traditional transponder architectures for maximum network optimization and cost benefits.

Also, we share our experience of doing live production field testing of 400G ZR++ coherent pluggable optics on very challenging spans and its operational impacts and results.

2. Cox Evolution of Optical Infrastructure

The Cox Communications network backbone was established circa 2001, consolidating leased OC-48, OC-12, and OC-3 circuits. Previously, each Cox market operated independently with its transit peering connections. Over the years, Cox experienced tremendous growth in markets and data, which made financial sense to build our own national Backbone to control our capacity planning without using leased circuits for better customer experience. The goal was to connect all metro markets to peering transit locations, reduce leased circuit costs, and provide resiliency.

2.1. Cox Optical Footprint

In 2007, a fiber optic network was established using Dense Wavelength Division Multiplexing (DWDM) technology, gradually expanding to cover all metropolitan markets and peering locations. Initially, the DWDM network utilized 80 channels with a fixed spacing of 50Ghz in the C Band, with each channel capable of carrying 10G/OC-192/1G. To carry non-return-to-zero (NRZ) 10G signals, dispersion compensation modules (DCM) were employed on the line systems. It wasn't until 2012 that Cox made the move to 100G Ethernet technology on the same fixed 80-channel DWDM system.

The national optical backbone links all national peering centers (NPC), national data centers (NDC), regional caching centers (RCC), and Cox regional data centers (RDC) using an optical mesh topology. Meanwhile, a metro core optical network connects all Cox hub locations as a hub and spoke topology, which aggregates at the RDC.

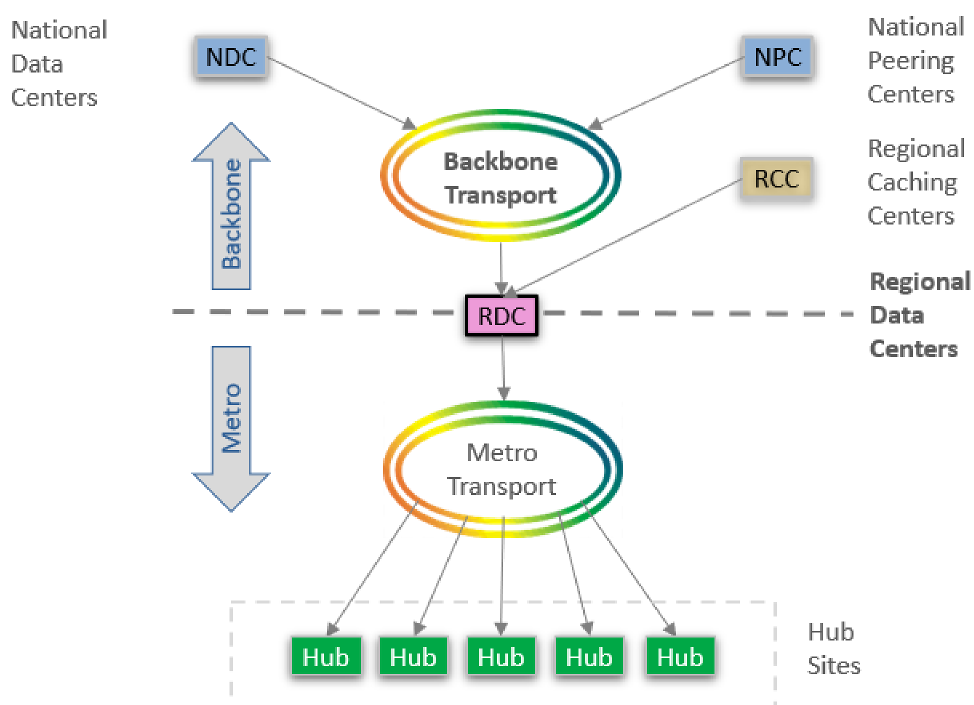

Figure 2 – Cox Backbone/Metro as of 2023

2.2. Problem Statement

The current Backbone and Metro core network architecture will not support future capacity requirements and service demands. To simplify the deployment and its ongoing support, re-architecting the entire optical transport system is needed, especially the line system, to accommodate future technological evolution, e.g., 400G/600G/1Tbps/1.2Tbps/1.6Tbps single wavelengths and pluggable optics such as 400GZR/ZR+/ZR++ for internet protocol over dense wavelength division multiplexing (IPoDWDM) architecture.

In summary, the current architecture in Cox's metro cores and national optical Backbone cannot handle advanced data rates. Many current line systems have limited technologies, such as fixed channel Reconfigurable Optical Add Drop Multiplex (ROADM) filters with a capacity of 100Ghz/50Ghz and 40, 80, or 96 channels. Some systems also use Fixed Optical Add Drop multiplex (FOADM) for optical transport, while others still require Dispersion Compensation Modules (DCM) for 10G wavelengths.

Approximately 70% of the optical Backbone and metro infrastructure require replacement to support 400G to 1.2Tbps waves and 400G ZR+ pluggable optics, while the remaining 30% needs updates in software and hardware to facilitate pluggables.

Let's dive into what the current transport network looks like.

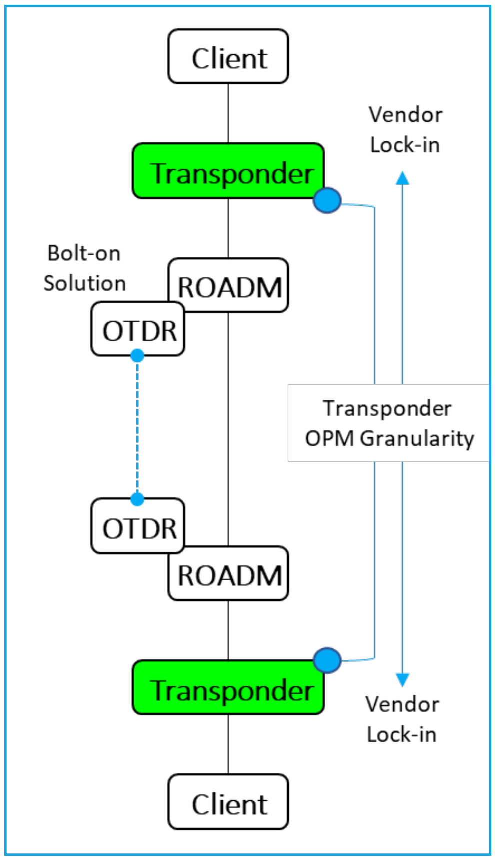

Figure 1 – Current Transport Network

Above is a typical example of a transportation network wherein a single vendor offers a comprehensive system encompassing the optical line system, with in-line amplifiers (ILA), ROADM, and transponders that convert client traffic to DWDM wavelengths on the line system. The advantage of a single vendor is that it has tighter power management between the line system and transponders, along with an element management system (EMS) for end-to-end visibility. This architecture allows a single vendor complete control over technology and innovation in service provider networks, leading to slow functional improvement and limited cost compression opportunities.

Introducing new technology on the transponder where most of the innovation is happening to maximize the spectral efficiency fully will require a whole line system upgrade or change to an entirely new vendor, which is cost-prohibitive and time-consuming.

3. Open Line System (OLS)

As the name suggests, an Open Line System (OLS) is an optical transport Dense Wavelength Division Multiplexing (DWDM) system without technology and vendor lock-in. The line system is meticulously designed to provide reliable and efficient support for high-capacity signal transponders. It also has pluggable optics that enable seamless access and facilitate the smooth operation of metro and long-haul backbones.

There are multiple interpretations of OLS, but at Cox, we use the OLS Architecture which means that a single vendor provides the optical and physical line system with open APIs, while different vendors supply the transponders and coherent pluggable optics with their open APIs. This type of architecture is sometimes referred to as partial disaggregation.

Without the technology lock-in, our line system can quickly adapt to support a wide range of technologies. It can scale up to accommodate future speeds of 400/800/1.6T coherent signals.

Furthermore, by avoiding vendor lock-in, the line system can support various transponder technologies from different vendors, even those that may be more advanced than the current vendor's technology. This flexibility allows service providers to choose the best vendors for their needs and mix and match them to take advantage of each vendor's strengths.

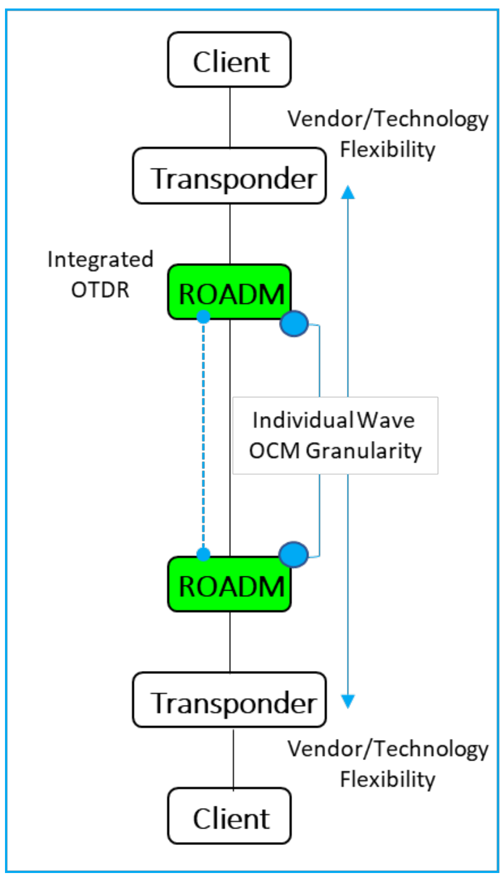

Figure 2 – OLS Transport Network

3.1. OLS Requirements

The requirements for implementing OLS in various networks may differ, but at Cox, we ensure the inclusion of necessary components in our network. These components are mandatory to ensure the smooth and efficient functioning of the OLS system.

Reconfigurable Optical Add Drop Multiplexer ROADM-based architecture must have flex grid capability. ROADMs are capable of wavelength routing, power balancing between channels, and spectral bandwidth allocation. To achieve this, they require the ability to measure fine granular optical channels using precise Optical Channel Monitors (OCMs). These OCMs can measure external alien wavelengths and power, mainly when using pluggable optics in OLS.

Flex Grid or Gridless is central technology that allows complete flexibility in packing any signal width (i.e., 50 GHz, 75 GHz, 112.5 GHz, or 150 GHz) as closely as possible. The current grid structure has been modified to allow for greater flexibility in frequency allocation. The grid can be adjusted in increments of either 6.25 GHz or 12.5 GHz, providing more excellent customization options for users.

Spectral efficiency is a critical factor in determining the extent to which available spectrum is utilized. This measure is indicative of the effectiveness with which fiber is being utilized, which is of utmost importance. Higher spectral efficiency allows service providers to assess how well transponders can be used to cram more bandwidth into a minimal fiber spectrum. This, in turn, expands the use of single-line systems, reducing overall network costs while simultaneously increasing capacity. Therefore, it is crucial to maintain and improve spectral efficiency as it contributes significantly to the network's overall performance.

$$\text{Spectral Efficiency (SE)} = \frac{\text{Information rate (bit/sec)}}{\text{Bandwidth (Hz)}}$$

A 10G signal sent over 50 GHz fixed spectrum has a SE of 0.2 b/s/Hz.

A 100G signal sent over 50 GHz has a much better SE of 2 b/s/Hz.

A 400G signal sent over 75 GHz has a better SE of 5.33 b/s/Hz.

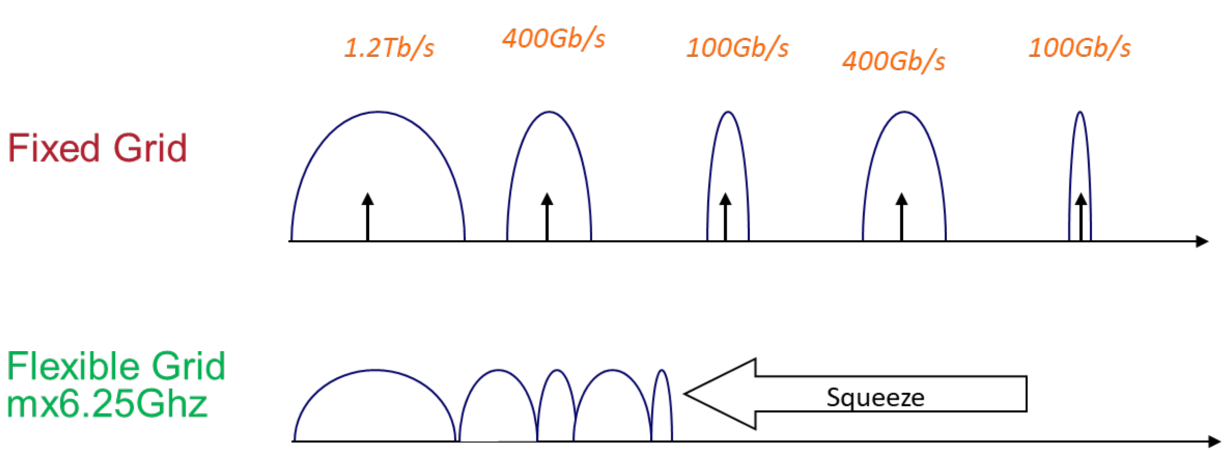

Figure 3 – Flex Grid or Gridless Example

Our simulations show a 30% improvement in capacity as channels only use as much bandwidth as required and nothing additional, as seen in the figure above. The Flex Grid feature is needed in the OLS architecture, which makes the optical transport modulation signal agnostic and a secure investment in the ROADM-based architecture.

An Optical Time Domain Reflectometer (OTDR) is a highly efficient and advanced device that is utilized for the purpose of detecting any fiber measurements and cuts in real-time. It provides the operator with real-time loss measurements and any possible deviations, which are essential for ensuring accurate design parameters. The OTDR is an indispensable tool for professionals working in fiber optic communication, as it enables them to quickly and easily identify any issues that may arise and take necessary corrective actions expeditiously. This device has proven to be highly reliable and effective, making it a popular choice among professionals in the industry. In the new optical line systems, OTDRs are integrated into the ROADMs, which give higher integration.

C+L band support starts with C-Band and then expands to L-Band to double the fiber capacity for future-proofing networks and being scalable.

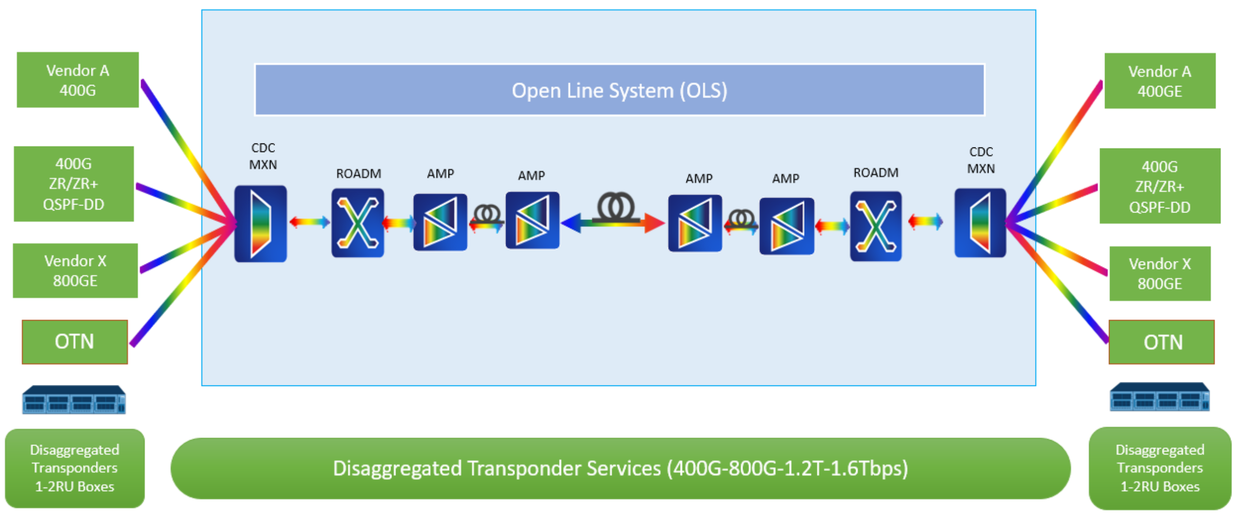

Figure 4 – Open Line System with Disaggregated Transponders (OLSDT)

The above figure shows the OLS with disaggregated transponders from different vendors, including OTN and pluggable optics. The line system is disaggregated from transponders and technology, giving Cox flexibility to scale and use each vendor's technology strengths.

3.2. OLS Benefits

As seen in the above section, a summary of the benefits are listed below:

- Maximize metro and core fiber pair bandwidth with C+L.
- Client and line disaggregated (different vendors)
- Separate transponder and ROADM development
- Highest opportunity for cost compression.
- Remove traditional client and line integration.
- Remove traditional IP and Optical boundaries (400G ZR/ZR+)
- Create maximum flexibility with client/service aggregation.
- Reduce client/service aggregation costs.

4. 400G ZR++ Pluggables

The history of optical transmission has seen many technical milestones, such as 40G and 100G coherent optical transmission. However, the latest innovation of the 400G ZR coherent optical plug is significant and is expected to revolutionize the optical industry. Service providers must adjust their architecture to take advantage of the plug's small form factor and reduced cost-per-bit savings. The 400G ZR is a highly advanced and reliable coherent optical plug that has been ingeniously modified to fit the compact QSFP-DD form factor. The innovative Optical Interoperability Forum (OIF) has designed this cutting-edge standard for data transmission as one of the first standards to establish an interoperable 400G interface. This cost-effective and practical solution is rapidly gaining popularity in the industry due to its unparalleled performance and efficiency in transmitting data.

OIF 400G ZR is limited to 80 to 120 km pt-pt links and best suits the Data Center Interconnect DCI environment. The typical launch is -10dbm, and this is an issue launching into ROADM-based architecture without external amplification.

To take a step further, 400G OpenZR+ MSA expanded 400ZR reaches far beyond 120km by utilizing higher-gain forward error-correction (FEC) coding and increased compensation for chromatic and polarization mode dispersion. This enabled the use of 400G pluggables not only in hyperscaler environments but also in metro/regional service provider network environments. Service providers have thus focused on these "plus" versions of 400G modules for their network needs.

The 400ZR++ is a pluggable solution that makes use of coherent optics. It is capable of extending beyond 120km, with the only difference being a higher Tx power range of -1 dBm to +5 dBm. Cox plans to utilize this solution as it is perfect for both Brownfield and Greenfield ROADM architecture. The power output is sufficient to counteract high insertion loss of Add/Drop Colorless -Directionless-Contentionless (CDC) architecture and transmit almost 1000km on the same OLS line system as vendor-specific embedded transponders.

Figure 5 – 400G ZR ++ plug QSFP-DD form factor.

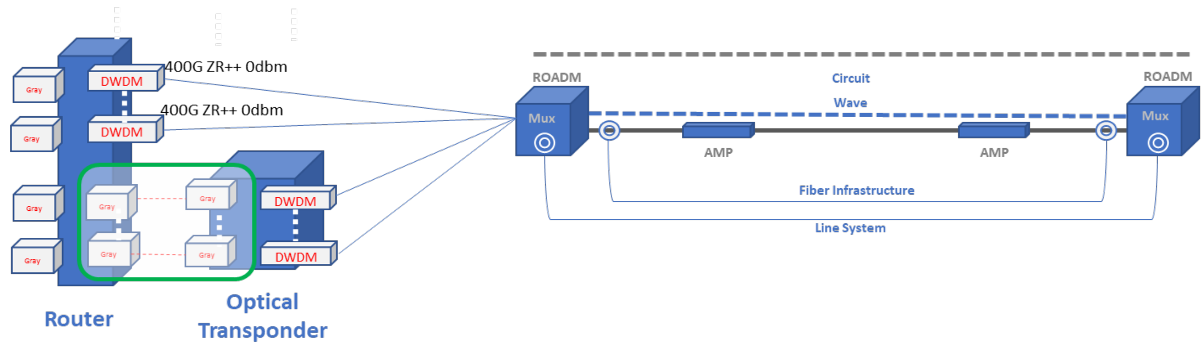

Figure 6 – 400G ZR ++ in Prod

In the Cox production, the implementation displayed in the figure above demonstrates where the 400G ZR++ pluggables will be utilized along with the Open Line System (OLS). The 400G ZR++ will be directly plugged into the line card of the routers, eliminating the need for an optical transponder and grey 400GE Ethernet client back-to-back plugs.

5. Live Production testing and Details.

In Cox, we put the 400G ZR ++ high power optics with -1dbm to +5dbm to test in a live Backbone production environment on very challenging routes over two different vendor line systems: one is OLS with Cox standard, and the other is a fixed older system with a Multicast Switch (MCS) add-drop architecture.

1. LA-LV route 522 km

On the Los Angeles to Las Vegas route, we have a combination of Large Effective Area Fiber(LEAF) fiber type and Single Mode Fiber (SMF) fiber type with distances between amplifiers of 110km each with the combination of Erbium Doped Fiber Amplifier (EDFA) and RAMAN amplifier chains which use stimulated Raman scattering.

The line system used was vendor A with OLS line system and twin Wavelength Selective Switches (WSS) which is commonly named as MXN add drop architecture.

Line System vendor B with MCS add/drop architecture with fixed spectrum.

This route was over 4 ROADM hop architecture.

2. LA-PHX route 1130 km

On the Los Angeles to Phoenix route, a Combination of TWC, LEAF, and SMF-28 fiber with Optical Ground Wire (OPGW) aerial fibers with 8 ROADM hops.

The same line system characteristics are utilized as in the previous LA-LV link.

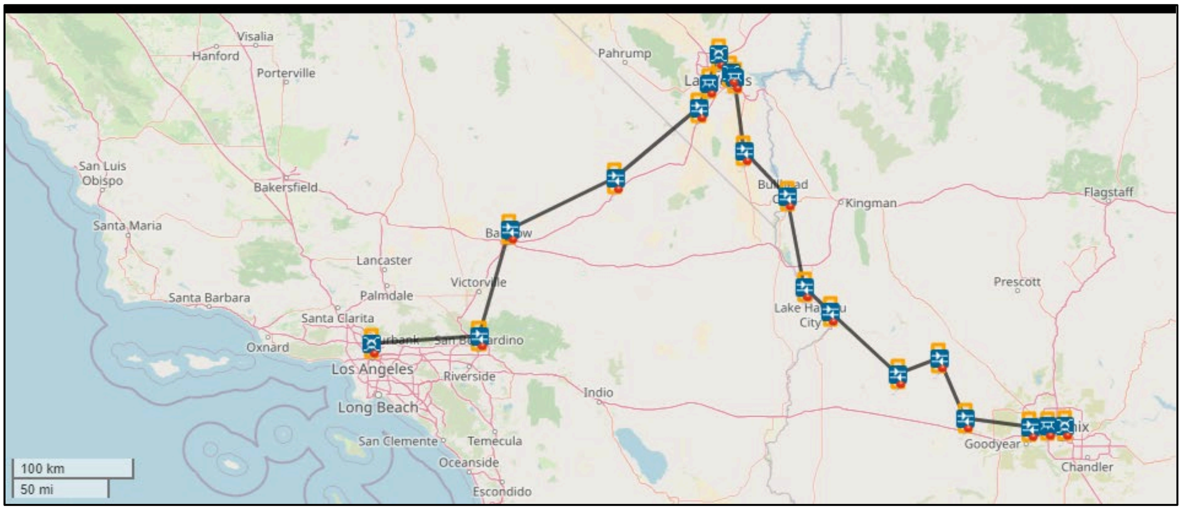

Figure 7 – Route LA-PHX Production test

Each pluggable vendor was allowed to bring their own host (i.e., the router-like devices where the 400G ZR++ are inserted), which they know will operate without issues and can be configured.

We used a 100G test set in Las Vegas and Phoenix to measure the ethernet frames for data integrity with a loopback at the LA side.

The testing methodology used a designated port on the Add/drop structure on each line system and tested frequencies in the different ranges – starting in the lower frequencies around 193Thz, gradually moving to the center C band frequencies, and then to higher frequencies around 196Thz – while keeping power constant at around -1 dBm to 0 dBm for optimal performance.

Also, we kept the spectral width – including guard bands – constant at 75Ghz, although some vendors needed an extra width of 87.5Ghz with their proprietary FEC and higher baud rate to get the highest performance from the plug.

5.1. Results of the test

Vendor 400G ZR++ 0dbm	Freq Low 193-194Thz	Freq Center 194-195Thz	Freq High 195-196Thz	Ckt Width (GHz)	Notes
Q	Pass	Pass	Fail	75	300G Test pass on higher freq
Y	Pass	Pass	Fail	87.5	300G Test pass on higher freq
G	Pass	Pass	Fail	87.5	300G Test pass on higher freq

Figure 8 – Results LA-LV

Vendor 400G ZR++ 0dbm	Freq Low 193-194Thz	Freq Center 194-195Thz	Frequency High 195-196Thz	Ckt Width (GHz)	Notes
Q	Fail	Fail	Fail	75	300G Test pass
Y	Fail	Fail	Fail	87.5	300G Test pass
G	Pass	Pass	Fail	87.5	300G pass on failed High freq

Figure 9 – Results LA-PHX

In the test over a live production environment, we observed that the pluggable vendors could close the span consistently. We were very impressed with the performance of every vendor we tested.

Before we dive into the results, let's review a couple of important metrics that Cox uses to determine its field deployable standard. The Pre-FEC BER threshold is the highest error rate at which the FEC algorithm can still provide nearly error-free communication after decoding. A Pre-FEC BER of 1.8 E-3 or higher is considered good and approved for field deployment. The Optical Signal to Noise Ratio (OSNR) measures the ratio of signal power to noise power in an optical channel. A higher OSNR is better, and generally, OSNR values above 24 dB are acceptable and approved for field deployment.

The Los Angeles to Las Vegas span was short in distance, and all vendors passed the lower and center frequencies with good Pre-Fec BER and OSNR margins. The results could have been more consistent in the higher frequencies, with some vendors passing the test error-free, but the OSNR values were not up to Cox standard to be field deployable.

The challenging span between Los Angeles to Phoenix is 1130km, where only one vendor could close the link with acceptable Q value but failed at higher frequencies, i.e., 195Thz and above. Then, the data rate was shifted to 300G, and all plugs could close the link with acceptable Pre-fec BER.

The biggest issue we discovered over our both-line system is that the 4000G ZR++ failed over higher frequencies, i.e., 195Thz range and above. The Pre FEC and Q values were not production-ready to be deployed and were very unstable, with inconsistent data integrity.

This issue can occur due to greater nonlinear penalties and ripple effects at higher frequencies, which causes Optical Signal to Noise Ratio (OSNR) to drop.

6. Architecture Changes

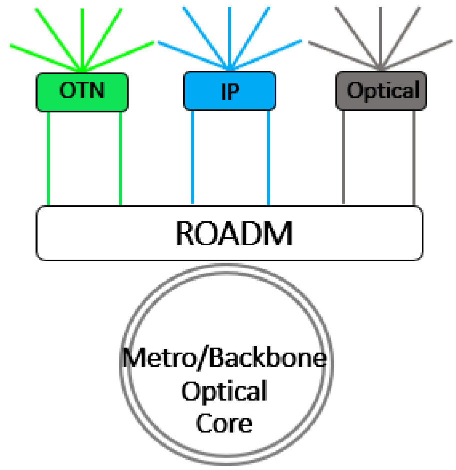

Figure 8 – Cox OLS Network Architecture

Cox's architecture has undergone a major change with the introduction of the Open Line System (OLS) at the foundational optical transport layer Zero (L0). The OLS must be scalable and future-proof to support the next generation of high-capacity, higher baud rate embedded transponders and pluggable optics. OLS includes key architecture components, such as ROADMs with Flex grid capabilities that can accommodate any data rates and modulation format. Additionally, C+L technology is used to future-proof the network from the ever-increasing bandwidth demands, providing 2x the bandwidth in the same fiber strands.

We will keep the OTN layer, IP layer with routers, and Optical transponders layer separate to eliminate any single point of failure while maintaining service termination and delivery separate.

We can optimize each layer separately and scale it according to its lifecycle and capacity demands. This gives Cox the advantage of being flexible and highly competitive if the demands or scope changes.

A ROADM optimized core optical layer is the best architecture for Cox multi-layer and multi-service network, which warrants flexibility and not vendor lock-in or a particular architecture locked in compared to Router HOP-HOP architecture.

400G ZR is a technology option and not an architecture change, and it's another transponder technology that will be used wherever the distances are less than 1000km in Backbone depending on the ROADM hops. Embedded transponders, which are highly designed with long haul links for distances above 1000 km and can overcome multiple ROADM hops filtering penalties, will be used for distances above 1000km and challenging multiple ROADM links.

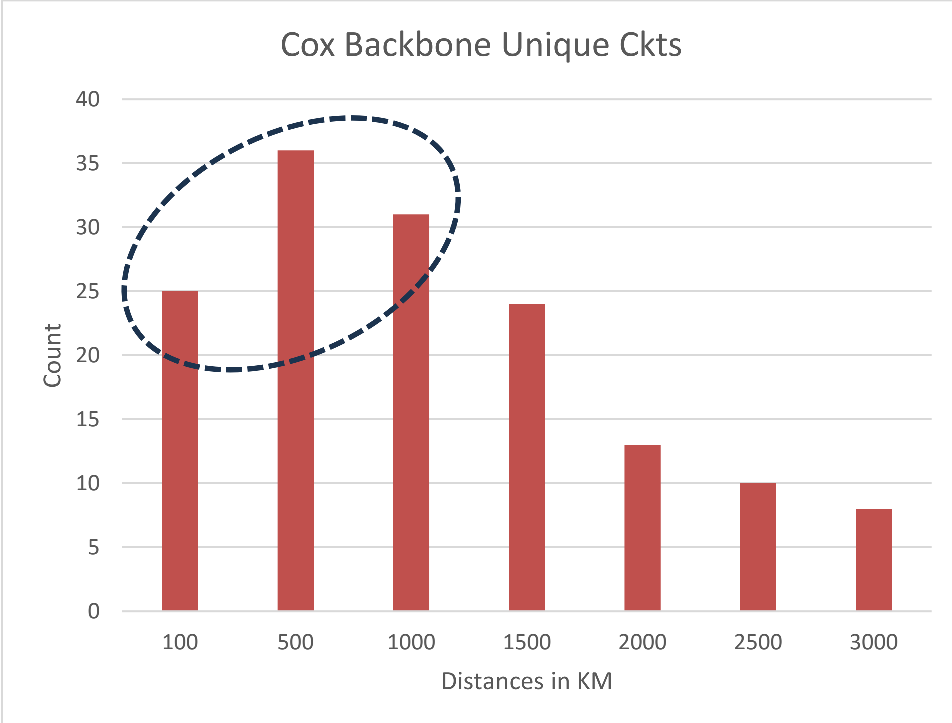

Figure 10 – Target implementation of 400G ZR ++ in Production

As we have tested, we can close the links within close to 1000km, but it all depends on the filtering penalties and the number of ROADM hops it takes.

The above Graph shows that close to 55% of the Cox Backbone circuits are candidates to close with 400G ZR++. It can significantly change how we design our circuits going forward because of the price compression, space, and power reduction we get using pluggable optics.

6.1. IP Layer Architecture Change

With the move to 400G optics, Cox is reevaluating its backbone router network topology. Today, the routed Backbone mixes point-to-point links and express circuits. All the links comprise various quantities of bundled 100G interfaces, with the highest around 30x100G (3Tbps). Moving to 400G interfaces will reduce the number of physical links by 75%. Some of the express circuits may be removed to take advantage of the space, power, and cost savings of using 400G ZR+. Another benefit of eliminating express circuits will be reducing the complexity of Shared Risk Link Group (SRLG) configurations on the routers. When you have multiple A-Z router bundles traversing the same physical span, your SRLG configurations need to be 100% accurate so you can avoid issues during failure events.

7. Operational Challenges

We have observed that in live production testing on different frequency spectrums, the performance is different because of nonlinear penalties at higher frequencies and ripple effects at different frequencies. This makes engineering alien 400G ZR++ waves through the OLS system challenging and will limit the distances it operates.

Line system vendors will need to develop robust engineering planning tools to provide engineering feasibility of these optical trunks with provided PSD (power spectral density).

End-to-end connectivity visibility is lost between the transponder, 400G ZR++ pluggable optics, and ADD/Drop card, making it harder for the operational team to troubleshoot. Also, there is no LLDP support since the grey ethernet optics with back-to-back are eliminated.

Currently, there exists no algorithm to manage power distribution between the line system and pluggable optics. This poses a challenge in adjusting power levels automatically in the event of a loss due to fiber issues and patch panel challenges. To minimize any fiber impairments, it is advisable to position the pluggable optics in the router as close as possible to the OLS transport.

Standards need to be accepted and implemented by both Optical and IP router vendors and work together to expose Open Config Common Management Interface Specification (CMIS) standards consistently.

Service provisioning in OLS and IP config is a challenge where the end-to-end to point and click experience through EMS is lost.

Service and Fault management is broken in optical vendor EMS, and it needs to be stitched together through an orchestrator/controller or home-grown tool to make it operational and workable.

Host router application selection (appsel) data for 400G ZR++, i.e., vendor-specific configuration for maximum performance with proprietary FEC and baud rate, must be supported and verified by the router vendor. Hence, it is part of the software configuration.

8. Conclusion

The optimal strategy for Cox is a ROADM architecture with a Flex grid and C+L system with optimized ROADM bypass. Using 400G ZR++ in the router directly will enable IPoDWDM, saving space and power and reducing cost per bit.

Cox will continue to deploy OLS ROADM infrastructure and use 400G ZR++ where applicable, i.e., for distances below 1000 km.

400G ZR QSFP-DD form factor pluggables are mature and ready for production deployment. Cox is confident that this will be deployed in large numbers in the coming years and integrated into the Cox Backbone and Metro networks.

This architecture with OLS gives Cox communications maximum flexibility by using embedded optics transponders for long-haul challenging links over 1000 km with multiple ROADMs and 400GZR ++ for short-haul links, which make the bulk of Cox Backbone and Metro links for maximum optimized gains and lowest Total Cost of Ownership.

Abbreviations

appsel	application selection
bps	bits per second
CMIS	common management interface specification
DCM	dispersion compensation module
DWDM	dense wavelength division multiplexing
DCI	data center interconnect
EMS	element management system
EDFA	erbium doped fiber amplifier
FEC	forward error correction
Gbps	gigabits per second
Ghz	gigahertz
HD	high definition
IP	internet protocol
IPoDWDM	internet protocol over dense wavelength division multiplexing
ILA	inline amplifier
L0	layer 0
LEAF	large effective area fiber
MCS	multicast switch
NRZ	non return to zero
NPC	national peering center
NDC	national data center
OLS	open line system
OSNR	optical signal to noise ratio
OTN	optical transport network
OPGW	optical ground wire
RDC	regional data center
ROADM	reconfigurable optical add-drop module
RCC	regional caching center
SCTE	society of cable telecommunications engineers
SMF	single mode fiber
THz	terahertz
Tbps	terabit per second
TCO	total cost of ownership