SDN-Controlled 400GbE end-to-end service using a CFP8 client over a deployed, commercial flexible ROADM system

L. E. Nelson[1], G. Zhang[1], N. Padi[1], C. Skolnick[1], K. Benson[1], T. Kaylor[1], S. Iwamatsu[2], R. Inderst[2], F. Marques[2], D. Fonseca[2], M. Du[2], T. Downs[2], T. Scherer[2], C. Cole[3], Y. Zhou[3], P. Brooks[4], and A. Schubert[4]
[1]*AT&T Labs, 200 Laurel Avenue South, Middletown, NJ, lenelson@att.com*
[2]*Coriant R&D GmbH, Munich, Germany; Coriant, Amadora, Portugal; Coriant USA, Naperville, IL*
[3]*Finisar Corporation, Sunnyvale, CA*
[4]*Viavi Solutions Deutschland GmbH, Eningen, Germany*

Abstract: We demonstrated a 400Gb/s Ethernet end-to-end circuit, inclusive of 400GbE client card with CPF8 interface and dual-carrier 16QAM line-side, on a production 100G core network segment between New York City and Washington DC. During the field trial, we demonstrated the feasibility of SDN-enabled creation, deletion, and re-routing of the 400G service.
OCIS codes: (060.4510) Optical communications; (060.1660) Coherent communications; (060.2360) Fiber optics links.

1. Introduction

While commercial deployments of 100G technology operating at 2bit/s/Hz are underway, standards already are focused on 400Gb/s Ethernet as the next-generation transport interface rate. In 2013, the IEEE launched its first activities aimed at exploring development of the 400GbE standard, and the IEEE 802.3bs 400GbE Task Force currently expects to ratify the standard by year-end 2017 [1]. Amongst the multiple options for pluggable client optics [2] (e.g. single or multiple fibers, single-mode or multi-mode fiber, single or multiple wavelengths, and modulation format), the 400GbE Task Force has adopted baseline proposals for 100m, 500m, 2km, and 10km link distances, based on either 16 or 8 electrical lanes. The industry has coalesced around CFP8 as the first generation form-factor for the 400G client pluggable, and the first samples were CFP8-SR16, based on 16 parallel lanes of 25Gbaud for short-reach. As the IEEE 802.3bs Working Group is at draft D2.2, and soon to be at D3.0, the first CFP8-FR8 and –LR8 samples, based on PAM-4, are now undergoing testing and verification of interoperability.

Optical transport at 400Gb/s has been demonstrated experimentally utilizing several options to optimize the trade-off between spectral efficiency and reach [3]. Reports of transmission of real-time 400G (and higher rates) over field fiber [4-6] date back to 2013; however when clients were transported by the lineside, they were 100GbE clients. Here we demonstrate, to our knowledge, the first end-to-end transmission of a true standards-based 400GbE service using a CFP8 client over a 400G channel in a commercial field system. The trial was conducted over a deployed route in a software-defined-networking (SDN)-enabled long-haul network, from New York City to Washington DC, with existing 100G channels carrying live customer traffic. 400GbE client cards and test-sets at both end points utilized CFP8-SR16 interfaces, and the received 400GbE client signals were error-free over 24 hour soak tests. In addition, we demonstrated the feasibility of SDN-enabled creation and deletion of the 400GbE channel on the direct path from NYC to WashDC, and Layer-1 re-routing of the 400GbE channel to an alternate path through Philadelphia.

2. 400GbE Transponder

The 400GbE transponder consists of a line card and a client card, which is based on the protocol set forth in Draft 2.2 of the emerging IEEE 802.3bs 400GbE standard, with 16 physical lanes at the electrical layer, as is the 400G test-set. The client card (Fig 1) contains a CFP8 (SR16, FR8, or LR8) cage as physical port for the 400GbE signal. The CFP8 connects through an electrical 400GAUI-16 interface to a Xlinx Ultrascale+™ FPGA, which implements the 400GbE FEC, PCS processing, and mapping of the 400GbE signal into four ODU4 data streams, which are connected to the line card through the backplane. Reed Solomon (544,514,10) with 10b symbol is the FEC in the 802.3bs 400GbE architecture, with pre-FEC BER threshold of ~2.0e-4. The framer FPGA also implements compensation of skew among the four ODU4 signals, arising from different propagation delays of the two 200Gb/s wavelengths across the optical system and fiber plant (i.e. due to dispersion in the fiber). A controller unit manages the framer FPGA and the CFP8 through its MDIO interface. The line card (CloudWave I02L200G-1) has two flexible optical interfaces and supports a throughput of up to 400Gb/s, where each line interface is configurable to 100Gb/s QPSK, 150Gb/s 8-QAM, or 200Gb/s 16-QAM formats. The line card maps the four ODU4 signals into two wavelengths, each with 200Gb/s 16-QAM modulation.

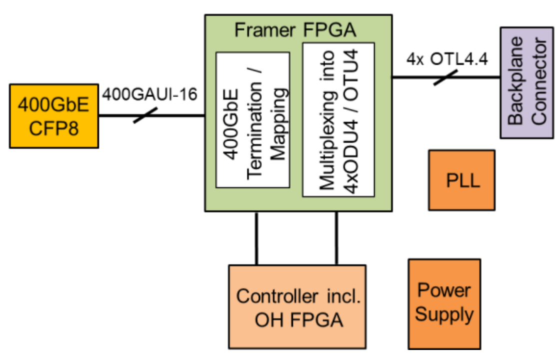

Fig 1: Block diagram of 400GbE client card

The CFP8-SR16 leverages the 100GBASE-SR4 technology, including existing uncooled 850nm VCSELs. The sixteen 26.6Gb/s electrical lanes of the 400GAUI-16 interface directly modulate the 16 optical lanes, without requiring multiplexing, translation, or de-skewing inside the module. The 16 optical lanes operate over 16 parallel multi-mode fibers, with maximum reach of 100m over OM4 fiber. The FEC implemented in the host card FPGA protects both the 400GAUI-16 electrical interface and the 400GBASE-SR16 optical link.

3. SDN-enabled 400G Service

An additional aspect of the 400GbE field trial was demonstrating, via an SDN Controller (SDN-C), the wavelength provisioning of three use cases: creation and deletion of a 400G channel, and re-routing an established 400G channel onto an alternate optical path while maintaining the existing client interfaces. SDN control is facilited in the deployed network by tunable, steerable ROADMs and the vendor-provided ROADM Network Controller (RNC).

For this field trial, we leveraged a centralized SDN-C that manages a typical core network [7]. Figure 2 shows how the SDN-C interacts with the RNC to ultimately establish the 400GbE channel. From the service provisioning portal, an end-user submits a request (create, delete, or re-route) with the appropriate parameters to set up the 400GbE service. SDN-C validates the request and verifies resource availability, and then initiates a wavelength operation request to RNC via the plug-in or standard RESTful interface. RNC then interacts with the ROADMs by performing a series of provisioning-oriented tasks. Service creation entails creating the end-points, setting the port mode, identifying the optimal path, tuning the two 200G coordinated wavelengths comprising the 400G channel, and provisioning the cross-connects. These operations are all done automatically through control and data plane interactions. As the wavelengths are being established, RNC provides status to SDN-C, which then makes decisions and informs the end-user. In the trial, the entire process from end-user's request initiation to successful provisioning of live traffic takes only a few minutes. Service deletion unprovisions the wavelengths and returns the end-point resources to the available pool for the next creation request. Service re-route simulates the case where the optical path transporting the 400GbE channel is disrupted or impaired, and on-demand re-routing of the wavelength path is warranted. The re-route scenario required specifying constraints to re-route around the simulated failure and utilized the same end-point resources, thus minimizing service interruptions.

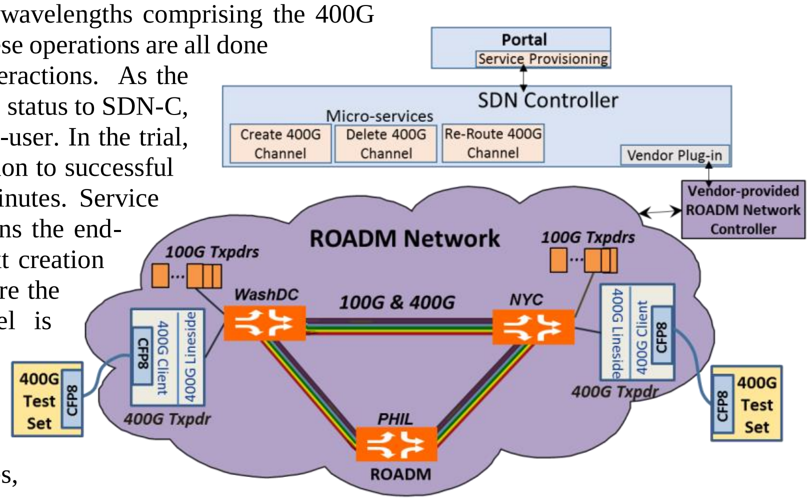

Fig. 2: Schematic of SDN controller architecture interacting with RNC.

4. Field Trial Route and Lab Verification of Line-Side Performance

The field trial was conducted over a recently commissioned direct route between NYC and WashDC (453km) consisting of six spans of TrueWave RS (TWRS) fiber with average length of 75.5km and average dispersion of 4.7ps/nm/km. The optical transport system (Coriant hiT 7300) supports up to ninety-six 100G channels at 50GHz spacing, utilizing erbium-doped fiber amplifiers and no optical dispersion compensation. Tunable, non-directional 50-GHz ROADMs based on 1x20 wavelength-selective switches are located in NYC and WashDC and were used to add/drop the two 200G 16-QAM wavelengths comprising the 400G channel to/from the existing traffic-bearing 100G polarization-multiplexed QPSK channels. The alternate route between NYC and WashDC is over 9 spans (540km) of standard single-mode fiber (SSMF), and the two 200G 16-QAM wavelengths were expressed through Philadelphia.

The 400G line-side performance was verified in the lab over an identical optical transport system and 6x80km TWRS and SSMF links, for two cases: with the 200G dual-carriers at channels 9 (1531.90nm) and 10, and at channels 91 (1564.68m) and 92, with the intention to cover the entire spectrum. At least two 100G channels were adjacent at shorter and longer wavelengths. Figure 3 shows the max pre-FEC BER of the poorer 200G carrier, as the 200G carriers' provisioned powers were varied relative to that of the 100G

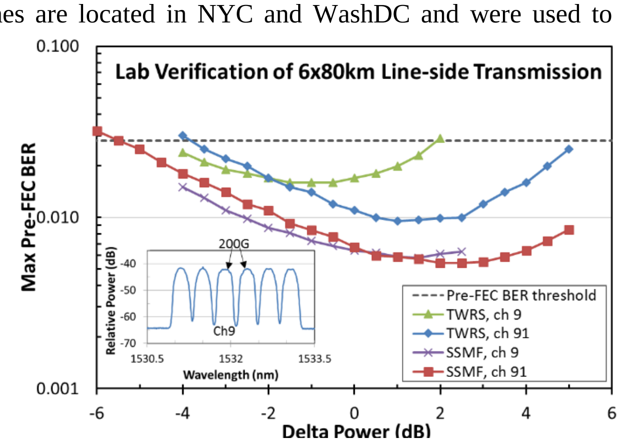

Fig. 3: 200G carrier line-side performance vs power relative to 100G channels. Inset shows spectrum at ch 9.

channels (i.e. delta power). For 0dB delta power, the OSNRs were measured to be approximately 23.4 and 22.8dB for 6x80km TWRS, and 25.0 and 24.6dB for 6x80km SSMF, for ch 91 and 9, respectively. For TWRS and 200G carriers at ch 9 and 10, the 100G (200G) channels incurred <0.3dB (<0.6dB) degradation in Q due to the adjacent 200G (100G) channels. The results indicated sufficient 400G performance and at least 2dB OSNR margin for both routes in the field, with the 200G carriers provisioned at the same power as 100G.

5. Field Trial Results

During the field trial, the first request to SDN-C was to create the 400GbE channel with no constraints on the route or wavelengths of the dual-carriers comprising the 400G. RNC chose the direct route between NYC and WashDC and channels 24 and 32 for the 200G carriers (see inset in Fig. 4), and creation of the 400G channel took 3 minutes. The 400GbE client signals over the end-to-end link were monitored on Viavi ONT-600 testsets with 400G CFP8 modules (see Fig. 2), and were error-free on all layers (PHYS, PCS, and MAC/IP) for 12 hours both in WashDC and NYC. SDN-C then instructed RNC to include the node in PHIL, and the 400GbE channel was re-routed on the alternate path (ch 24 and 32 were again chosen). Re-routing took approximately 2 min. The 400GbE client signal was error-free on all layers for 24 hours both in WashDC and NYC. Table 1 summarizes the 400GbE client performance and 200G carriers' line-side performance over the direct and alternate routes. No post-FEC line-side errors were recorded during either soak test.

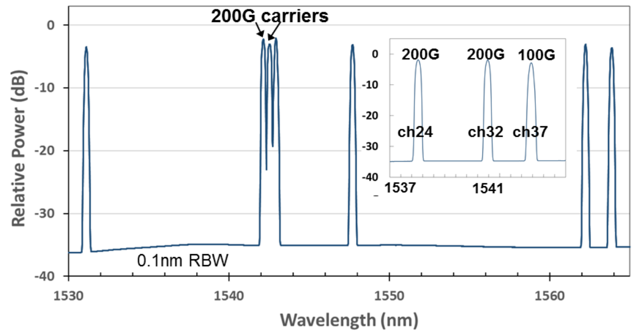

Fig. 4: Spectrum of channels on direct route between NYC and WashDC, with 200G carriers at ch 35 and 36. Inset shows case when RNC chose ch 24 and 32 for the 200G carriers.

Next, channel restrictions were applied to force the RNC to provision channels 35 and 36 as the 200G carriers, in order to assess the impact on the existing 100G channel 37. SDN-C created the 400GbE channel on the direct route (Fig. 4 shows the launched spectrum in NYC), and the 400GbE client was error-free on all layers (PHYS, PCS, and MAC/IP) during the 24 hour soak. As shown in Table 1, the maximum lineside BERs were 8.8e-3 for ch 35 in WashDC and 1.2e-2 for ch 36 in NYC. The 100G ch 37 incurred <0.2dBQ impairment due to the 400G dual-carrier signal, indicating nearly negligible performance impact. When the re-routing request to include PHIL was submitted to SDN-C, the 400GbE client signal was down for only 124sec in NYC and 111sec in WashDC; then the 400GbE clients soaked error-free on all layers for 24 hours. The maximum lineside pre-FEC BERs were 3.6e-3 for ch 35 in WashDC and 5.3e-3 for ch 36 in NYC. No post-FEC line-side errors occurred during either soak test.

Route and direction	Fiber type	Route length (km)	Soak test (hours)	Line-side			Client-side	
				Channel number of 200G carriers	Maximum pre-FEC BER's of 200G carriers	Uncorrected Bit Errors (after FEC)	PCS Corrected Code Word Errors	PCS Uncorrected Code Word Errors (after FEC)
NYC --> WashDC	TWRS	453	12	24 & 32	9.5e-3 (ch24) & 7.7e-3 (ch32)	0	0	0
WashDC --> NYC	TWRS	453	12	24 & 32	1.2e-2 & 8.5e-3	0	53	0
NYC --> PHIL --> WashDC	SSMF	540	24	24 & 32	4.9e-3 & 3.5e-3	0	0	0
WashDC --> PHIL--> NYC	SSMF	540	24	24 & 32	5.0e-3 & 2.8e-3	0	328	0
NYC --> WashDC	TWRS	453	24	35 & 36	8.8e-3 (ch 35) & 1.1e-2 (ch 36)	0	0	0
WashDC --> NYC	TWRS	453	24	35 & 36	8.5e-3 & 1.2e-2	0	153	0
NYC --> PHIL --> WashDC	SSMF	540	24	35 & 36	3.6e-3 & 5.2e-3	0	0	0
WashDC --> PHIL --> NYC	SSMF	540	24	35 & 36	2.9e-3 & 5.3e-3	0	271	0

Table 1: Summary of 400GbE performance during the field trial.

6. Conclusion

We have demonstrated the first, to our knowledge, end-to-end transmission of a 400GbE service using a true standards-based client in a commercial field system. The 400GbE client card utilizes a CPF8 module and is paired with a 400G dual-carrier 16QAM line-side for transmission over a production 453km network segment carrying live 100G traffic between New York City and Washington DC. In addition, we demonstrated the feasibility of SDN-enabled 400G channel creation, deletion, and re-routing from a direct to alternate route.

7. References

[1] IEEE P802.3bs 400 Gb/s Ethernet Task Force, http://www.ieee802.org/3/bs/public/index.html
[2] C. Cole, "Beyond 100G client optics," IEEE Communications Magazine, Vol. 50, No. 2, pp. S58-S66, Feb. 2012.
[3] See X. Zhou and L. Nelson, J. Lightwave Technol, Vol. 30, No. 24, pp. 3779-3792, 2012 and references therein.
[4] Y. Huang, et al., "Real-Time 400G Superchannel Transmission using 100-GbE based 37.5-GHz Spaced Subcarriers," OFC 2013, NW4E.1.
[5] Y. R. Zhou, et al., "Real-time gridless 800G superchannel transport field trial over 410km using coherent DP-16 QAM," OFC 2014, Tu2B.3.
[6] A. Pagano, et al., "400 Gb/s real-time trial using rate-adaptive transponders for next generation flexible-grid networks," OFC 2014, Tu2B.4.
[7] M. Birk, et al., "Evolving to an SDN-enabled ISP Backbone: Key Technologies and Applications", IEEE Comm Mag, 54(10): 129-135, 2016.