2023 INTERNET2 TECHNOLOGY exchange

Leveraging coherent optics and open line systems in production scenarios

JoAnne Bender:
Sr. Network Engineer, Internet2

Talk overview

- What is a coherent pluggable?
- Why are we opting to use coherent pluggables?
- Considerations in the Internet2 network for utilizing them
- Open Line System considerations
- Deployment strategies for coherent optics

What is a coherent pluggable?

Useful to review what a traditional pluggable optic for route/switch is:

- Traditional SFP/SFP+/QSFP-[28,DD] are amplitude modulated.
  - Typically RZ/NRZ/PAM4 modulation schemes.
  - For RN/NRZ, One bit per symbol (1 baud)
    - PAM4 doubles this as 2 bits per symbol, so higher efficiency, but still has the same issues as other amplitude schemes.
  - Fixed filters required in most cases, transmit lasers are tunable, but receivers are non-selective. Can’t really amplify without dispersion compensation

What is a coherent pluggable?

Coherent optical scheme:

- Instead of amplitude modulation, it uses a phase modulation approach.
- BPSK/QPSK/QAM (so much higher bit density per symbol)
- FEC built in.
- Automatically compensates for fiber dispersion
- Can be deployed with only amplifiers every 75-100km
- Can also use multiple light polarizations to transmit even more bits.

Previously these were very proprietary schemes via modem manufacturers, so interoperability between vendors was not possible. Something needed to be done, so the The OpenZR+ Multi-Source Agreement (MSA) was created.

What is a coherent pluggable?

Current standards:

	400ZR	Open ROADM	OpenZR+
Target Application	Edge DCI	Carrier ROADM Mesh	Regional/LH DCI and Carrier
Client Traffic	400GbE only	100-400GbE & OTN	100-400GbE Multirate
Target Reach @ 400G	120km	500km	Regional/Long-haul
Form Factor	QSFP-DD/OSFP	CFP2 or other	QSFP-DD/OSFP
SD-FEC	CFEC	oFEC	oFEC
Standards/MSA	OIF	Open ROADM MSA	OpenZR+ MSA

What is a coherent pluggable?

References:

Open ZR+ MSA Technical Specification:
https://openzrplus.org/site/assets/files/1075/openzrplus_1p0.pdf

Common Management Interface Specification (CMIS)
http://www.qsfp-dd.com/wp-content/uploads/2021/11/CMIS5p1.pdf

QSFP-DD/QSFP-DD800/QSFP112 Hardware Specification for
QSFP DOUBLE DENSITY 8X AND QSFP 4X PLUGGABLE TRANSCEIVERS
www.qsfp-dd.com/wp-content/uploads/2022/07/QSFP-DD-Hardware-Rev6.3-final.pdf

Implementation aspects

ZR+ power consumption:

- Not all systems can take a full load of ZR+ (especially the high powered “bright” versions!) unless designed for it.

- Often means that every other port, or densities as low as every four ports due to either power consumption or heat dissipation.

- Because of this, if you are planning a 400G-ZR+ or ZR+ high-powered deployment, you will need to consider planning your other pluggables carefully!

Implementation aspects

- Software support.
  - Plug parameters (eg: frequency, power, framing, symbol rates,etc) *must be set* in order to operate.
  - Not just setting shut/no shut, and possibly tuning frequency and power anymore!
  - This is due to all of the modes it can support:

This means that your device must have software support to support the CMIS commands to set this parameters.

This support is now very common on many 400G devices, but it is very essential now that you check with your platform vendor and optics vendors to find out minimum supported software versions for the features you may need.

Especially if you go outside the OpenZR spec and use vendor proprietary modes (such as Ciena PacketMAX) which require things like sending CMIS AppSel commands to the plug!

Table 1-3: OpenZR+ Line Encoding, Modulation and Symbol Rates

OpenZR+ Format	SFF-8024 Media ID	Payload Rate	Framing Format	Symbol Baud Rate (+/- 20ppm)	Modulation	FEC	Net Coding Gain (NCG) (dB)	Pre-FEC BER	Reference Standard
400ZR+	46h	400G	ZR400-OFEC-16QAM	60 138 546 798	16QAM	OFEC	11.6	2.0E-2	OpenZR+
300ZR+	47h	300G	ZR300-OFEC-8QAM	60 138 546 798	8QAM				
200ZR+	48h	200G	ZR200-OFEC-QPSK	60 138 546 798	QPSK				
100ZR+	49h	100G	ZR100-OFEC-QPSK	30 069 273 399	QPSK				

Developing a coherent coherent strategy for Internet2:

Internet2:

- Mostly a medium to longer haul network.
- Some metro networks and shorter haul distances within the network exist, however.
- Completely muxponded network for 400G today (primarily via Waveserver5)
- Existing muxponders additionally support 100G via capacity on the modems past the 400G requirements.

Constraints:

- Not desirable post-NGI to greatly expand the space/power/cooling footprint unless a major expansion is warranted.
- Wavelogic 6 is right around the corner, so we want to be conservative in how we want to allocate resources obtain more WL5 based modules.

Developing a coherent coherent strategy for Internet2:

Drivers/Inputs for coherent adoption:

- Growing need for additional 100G capacity.
  - This is largely provided by 100G ports on WL5 today.
  - WL5 line rates often could be higher as the line system will support it, but the configuration (Requiring both OTL4.4 and Ethernet) limits us to 1x 400G and 1x OTL4.4 and 1 100GE port today.
    - Therefore we have stranded capacity as we're not able to fully utilize the available SNR.
  - So in order to support additional 100G demands, we previously would be looking at installing additional WL5 sleds in POPs.
    - JoAnne, didn't you just say that you didn't want to acquire more WL5 sleds?
    - Well, if we can *move* WL5 sleds/chassis around the network a bit, we can get better equipment utilization without additional spend.
    - This also makes facilities and opts happy that we are lowering power consumption at smaller POPs.

Developing a coherent coherent strategy for Internet2:

Drivers/Inputs for coherent adoption:

- We also have considerations for additional bandwidth demands on 400G as well.
  - Some of this is due to additional redundancy we've considered adding to the network.
  - NA-REX connectivity is in progress.
  - Some of these NA-REX links are rather long, so NOT a good fit for a coherent pluggable.
  - Now I need more 400GE capacity on WL5 for longer-haul applications.
    - …And we're back to the not wanting to take up space with more WL5 units.
    - But, again, we can move them around!
  - What about new sites?

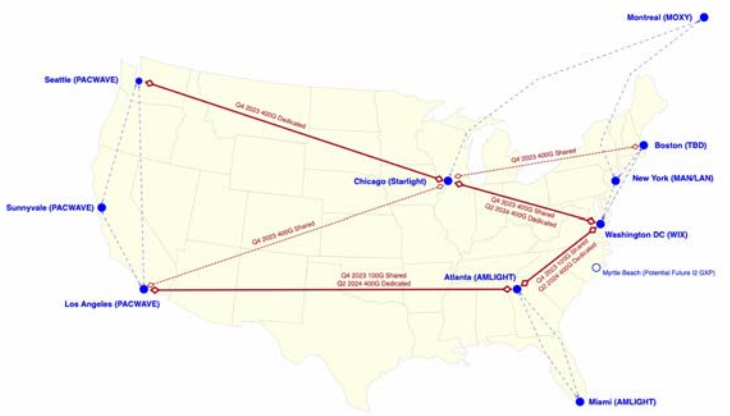

Developing a coherent coherent strategy for Internet2:

Hardware selection:

We have selected hardware from Ciena and Cisco for our coherent pluggable implementation.

Ciena:

400G Universal QSFP-DD
PN: 176-3370-900
Also has extended proprietary modes. (Packet Max) 60/70 gbaud

Cisco:

Cisco 400G QSFP-DD High-Power (Bright) Optical Module
PN: DP04QSDD-HE0
+1dBm launch power, for extended range and performance.

Implementation scenarios

Given these inputs, we've developed a few potential deployment scenarios:

Metro network capacity reclamation

Connector Interconnection

Medium-haul backbone capacity augmentation

New spur POP implementation

Implementation scenario #1: Metro network capacity reclamation

- Internet 2 has a few metro areas (Chicago, DC, NYC, LA, et al) where we are running 600G-800G locally for 2x 400 or 1x400 2x100G
- That pesky stranded capacity issue again.
- But these metros are also our larger optical sites and have larger amounts of drop capacity.
- So, if we pull 400G back out of the Waveservers for metro circuits, and use coherent direct attached, we could:
  - Convert to 800G line sides on the metro side, and run 8x100G across the metro.
    - This helps support the 100G wave expansions across long haul or metro.
  - Or reclaim the 400G capacity to support long-haul operations for the NA-REX mission.
  - Or even move the sleds to other markets where additional redundancy or capacity is needed.

Implementation scenario #2: Connector Integration

Connector wanting to add additional diverse ports as part of Large/Small platform flex:

- Is located in a POP within 100km of an I2 POP and has fiber access between POP locations.
- Does not want to build out a cabinet presence in the I2 POP, and possibly wants multiple ports from I2 from that POP as well.

Deployment solution:

- Pair of 400GE coherent pluggables, one in the I2 router, one in the connector fiber across the network
- 100GHz passives if more than one channel is desired.
- Could be extended to amplified scenarios (Raman, EDFA) if greater distances are desired.

Implementation scenario #3: Medium-haul backbone capacity augmentation

Additional 400G connectivity is required due to one or more of the following:

- Additional resiliency is required along a route to prevent outages.
- Additional bandwidth is needed at a particular location.
- A new “express path” from a POP to a major POP is required for latency reduction or dedicated path

Distance is within 550-600km.

Implementation is straightforward. Addition of coherent pluggable via the line system and build a new path.

Implementation Scenario #4: New spur POP

- Building out a new POP in Boston to land the new 400G interconnection with GÉANT.
- The problem is, that connection lands in 1 Summer, and Internet2 isn't in that location. But we are in Boston, and we can get fiber between the POPs!
- Additionally, there is not room in the existing Boston pop 6500 chassis to provide an additional drop degree toward that POP. But an existing local drop degree does exist today for a single 100G wave in that site (On OTR).
- What to do? More waveservers in this case may not be the answer. We're low on room. Now what?

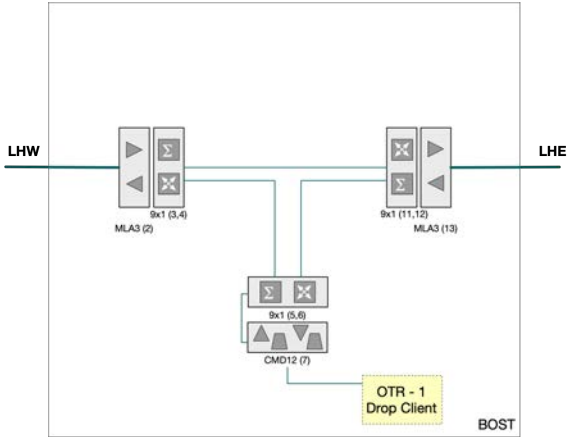

Current - BOST / BOST2 POP

Implementation Scenario #4: New spur POP

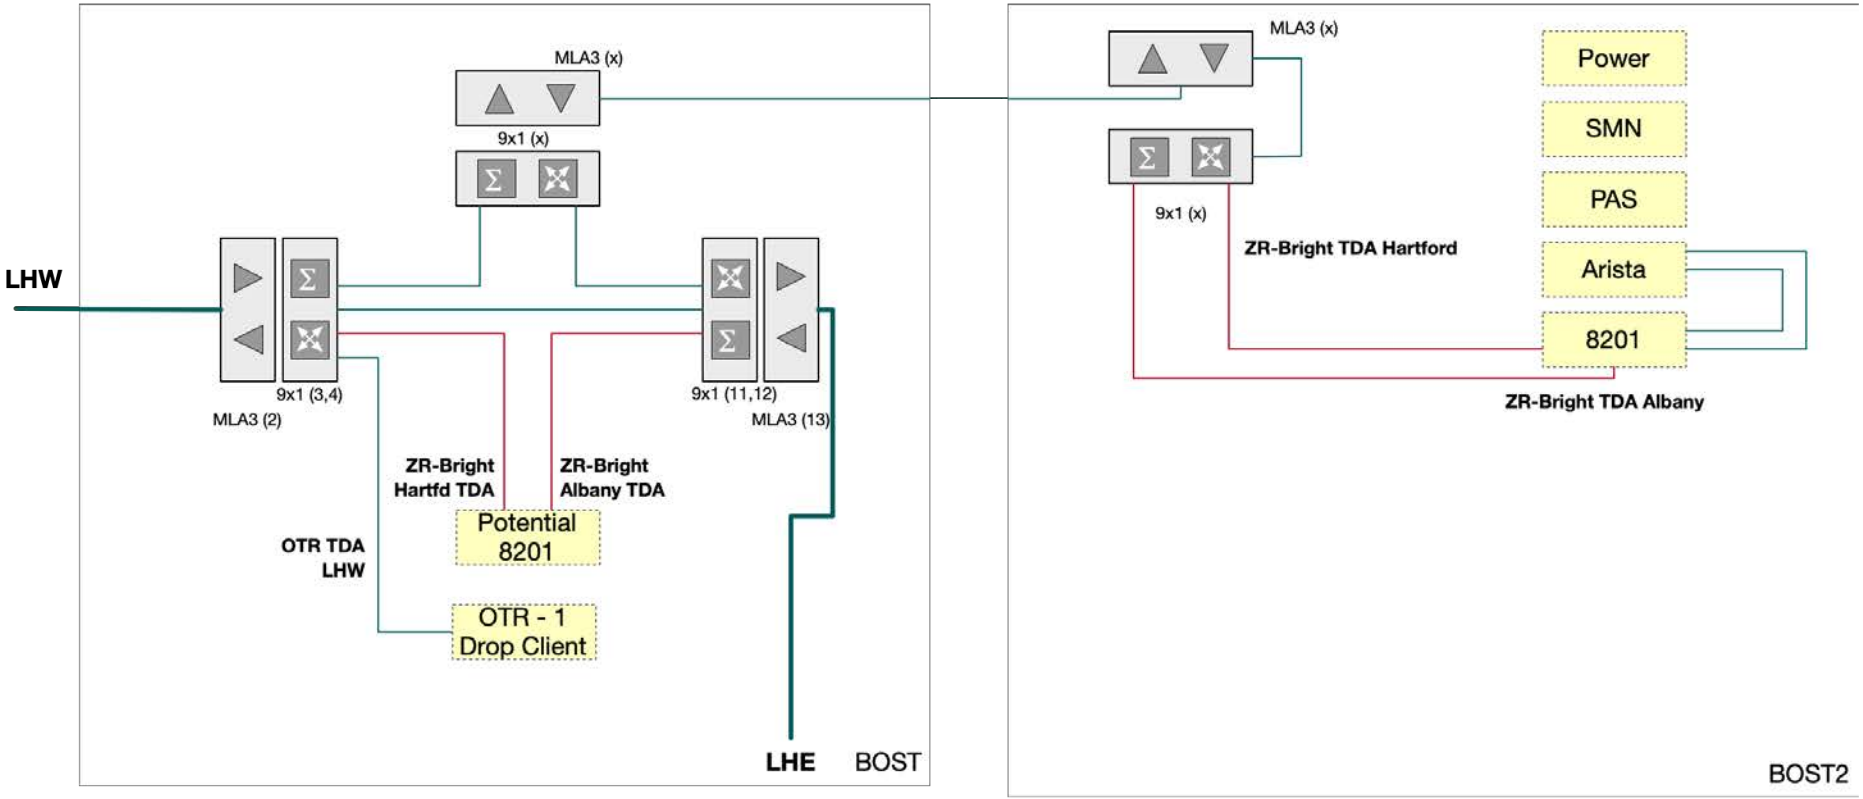

Optical SIG Announcement

Starting an optical SIG no later than 12/1

As interconnection becomes increasingly complex, we as a community should be discussing the layer 0 and layer 1 issues.

I am looking for a community member to chair this SIG.

If you're interested in participating, please see me afterwards, or send your contact info to jbender@internet2.edu

THANK YOU!

Questions?