Internet2 Advanced Networking

NGN-SIG Meeting – April 2024

Christian Todorov, Network Services, Internet2

Agenda

- Next Generation Networking
- Network Performance Assurance
- Optical Infrastructure Sharing
- Router Sharing

Internet2 Networking Overview

Next Generation Infrastructure

The Internet2 Network

- Internet2 Network
  - 400G native networking – one of the first networks globally to deploy 400G at scale
  - 18,700 miles of dark fiber. Replaced ~ 80% I2-IRU nationwide LEAF fiber with new low-loss SMF-28 Ultra single mode fiber (through 2042)
  - 32Tbps of optical capacity per link
  - Assured supply chain for all networking vendors; equipment certifications:
    - 6500: JITIC, NIAP/CC, CMVP
    - Waveserver 5: NIAP/CC, CMVP
    - Cisco 8200: NIAP/CC, CMVP;
    - WS5 & 8200: JITIC - in process
  - 24x7x365 Production-level engineering and NOC support

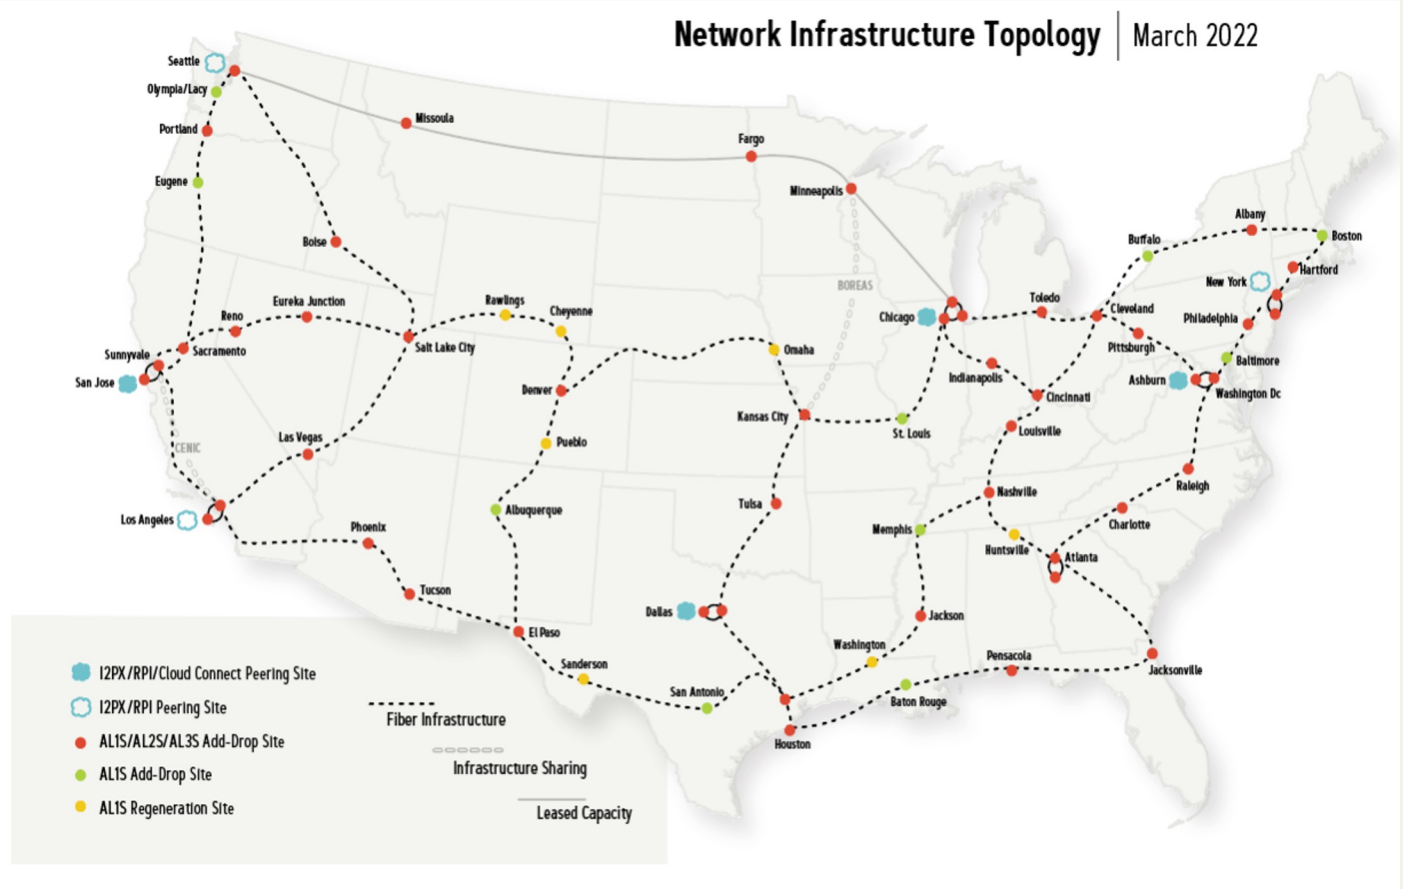

JITIC: Joint Interoperability Test Command
NIAP/CC: National Information Assurance Partnership/Common Criteria
CMVP: Cryptographic Module Validation Program – Supported capability; no current demand for encryption

MASSIVE SCALE WITH SMALLER CARBON FOOTPRINT

	2010	NGI Online 2022	Impact
Total Annual Traffic	104 PB	2,785 PB	Increase of 26x
Backbone Link Capacity	250 Gbps	127,700 Gbps	Increase of 500x
Total Device Capacity	453 Gbps	810,000 Gbps	Increase of 1800x
Footprint Power Utilization	300,000 Watts	100,000 Watts	Decrease by 2/3x

RPI, I2CC, DDoS, I2IC Network Numbers Entering 2024

2024 Stats

- I2 Cloud Connect
  - **539** Connections
- I2 RPI
  - **42** 10G Interfaces
  - **2** 100G Interfaces
- I2 DDOS
  - **14** Subscribers
  - **95** Tenants
- I2 Insight Console (Since 11/23)
  - **217** Layer 2 Connections Created or Updated
  - **90** Layer 3 Connections Created or Updated
  - **48,000** Looking Glass Queries Executed
- Routing Integrity
  - LRSA executed by **~300** R&E institutions

New Internet2 Performance Assurance Service

47 Node 2x100G dedicated perfSONAR nodes - Deployment complete

100G interface for internal PAS mesh and other 100G interface for ad hoc and external testing

- Performance tuning:
  - Initial configuration in place
  - Final tuning is pending

A Plethora of Public perfSONAR Points in PoPs

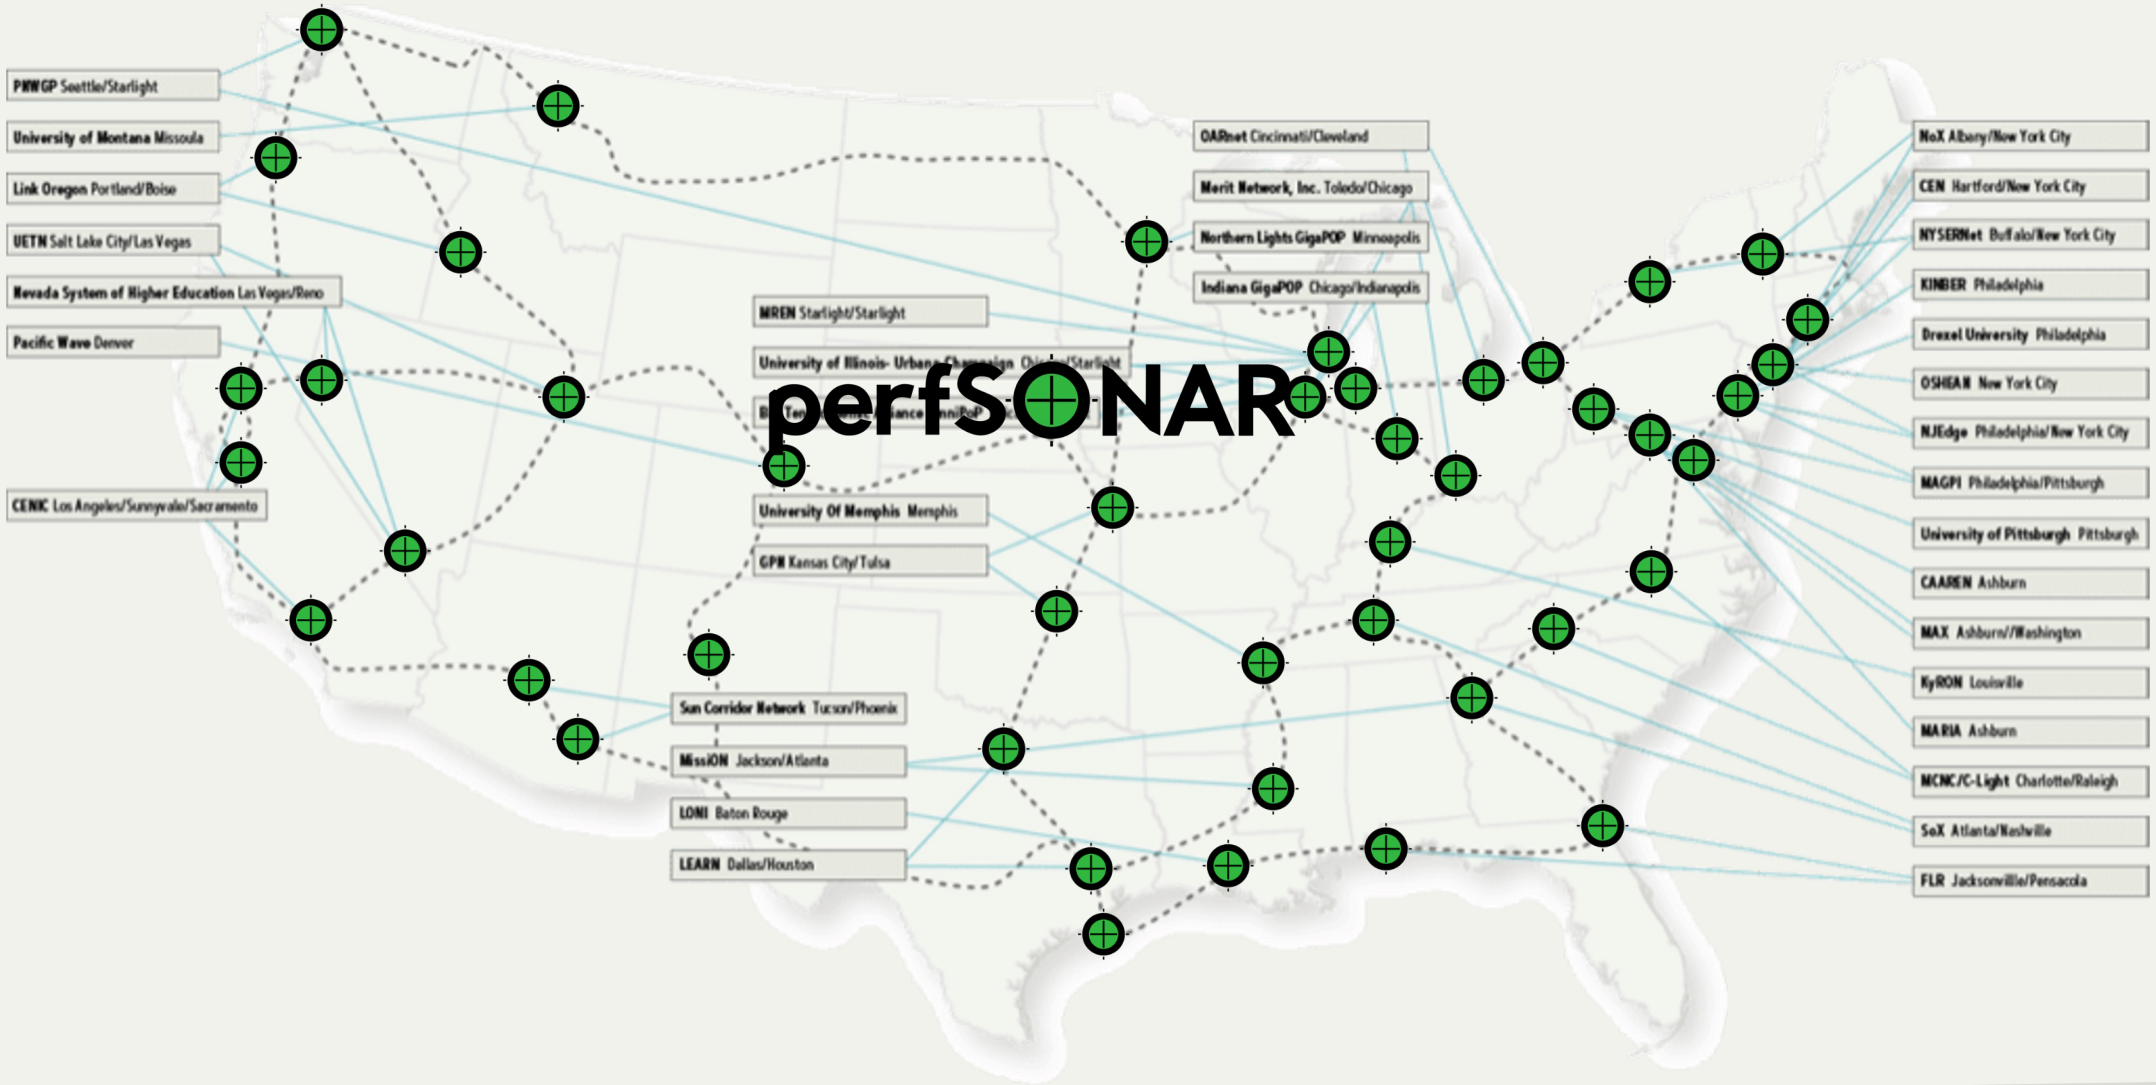

Network Attached Distributed Computing

Distributed Computing

- Participate in Open Science Data Federation (OSDF)
- Provides Distributed high-throughput computing (dHTC) in support of open science
- Internet2 operates 12 cache / compute nodes
- Filling key latency and location gaps

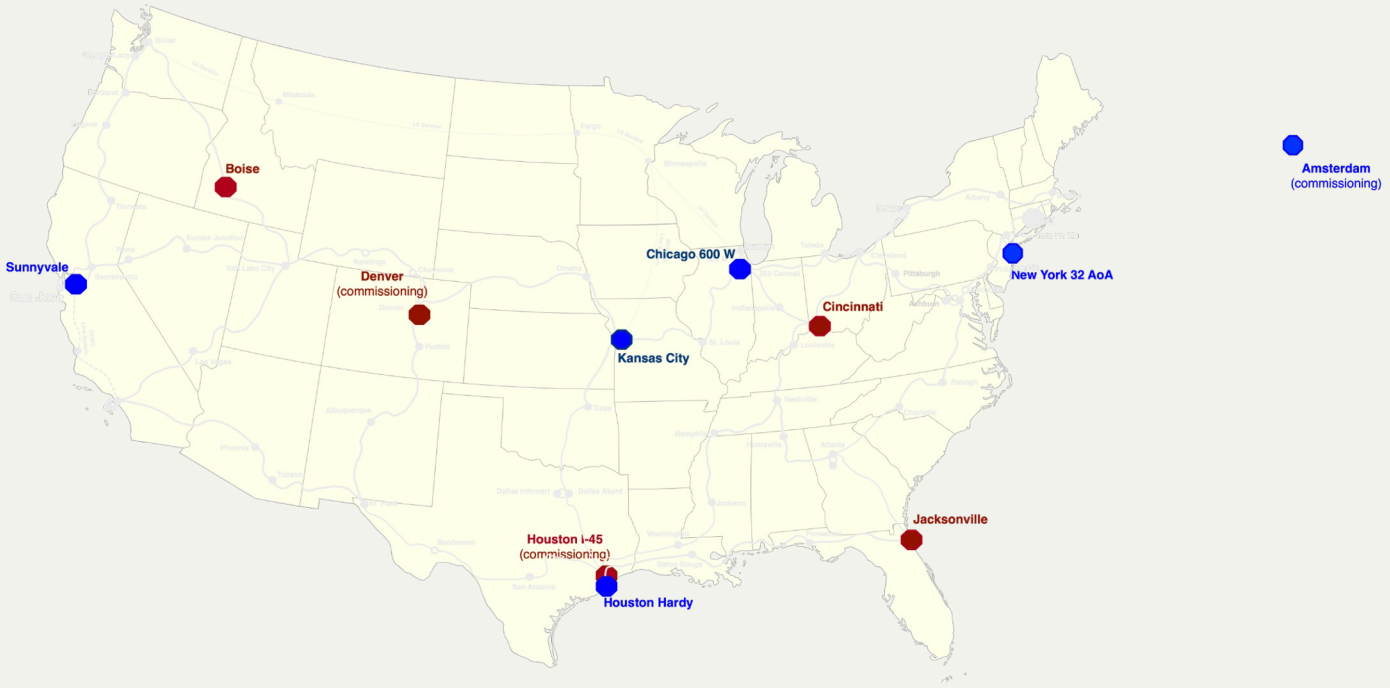

400G Trans-Atlantic

400G Transatlantic Capacity Additions/Upgrades on Amitié cable

- 1 x 400G for Internet2/CANARIE
- 2 x 400G for ESnet
- **Active Now!**
- Added Boston as open exchange point
- Early effort to acquire spectrum services from commercial providers – lessons learned

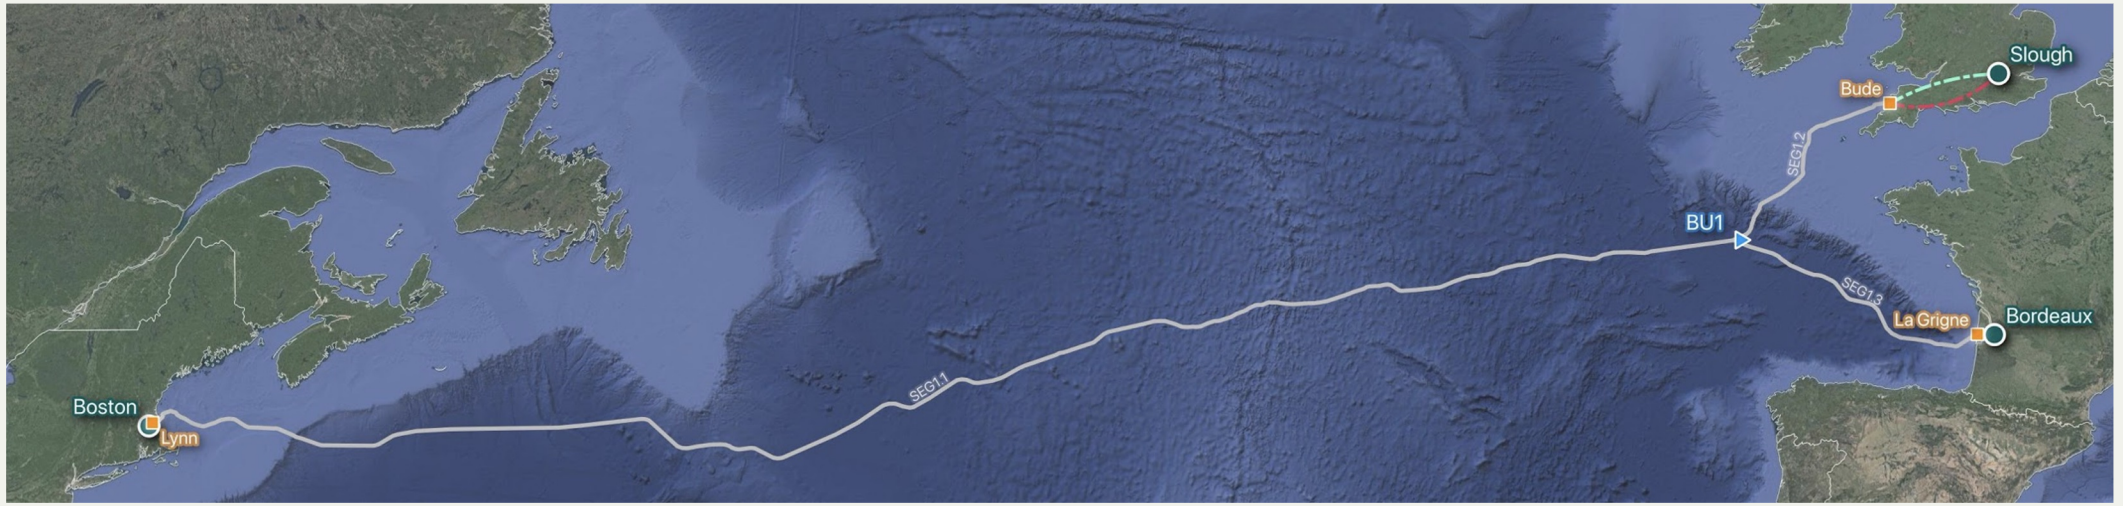

Highlight: Developing Areas of Community Engagement & Contribution

Internet2 NETWORK SERVICES

International Collaboration: Architecture Development

- Confluence of Advanced North Atlantic (ANA) Engineering and Global Network Advancement Group (GNA-G) Advanced Engineering Working Group
  - ANA has 400G capacity coming online 2024,2025 and needs improved link management, analytics, and interoperability with DOE networks

- A system-engineering approach to support:
  - Global traffic management (example: primary and secondary paths)
  - Varied philosophies and business models of participants
  - Support Multiple Use Cases

- Likely will focus on ANA initially, but push larger best practices out to community through the GNA-G

- Additional GNA-G activity on “digital twin” testbeds and demo activities (led by Marijke Kaat and the team at the University of Amsterdam)

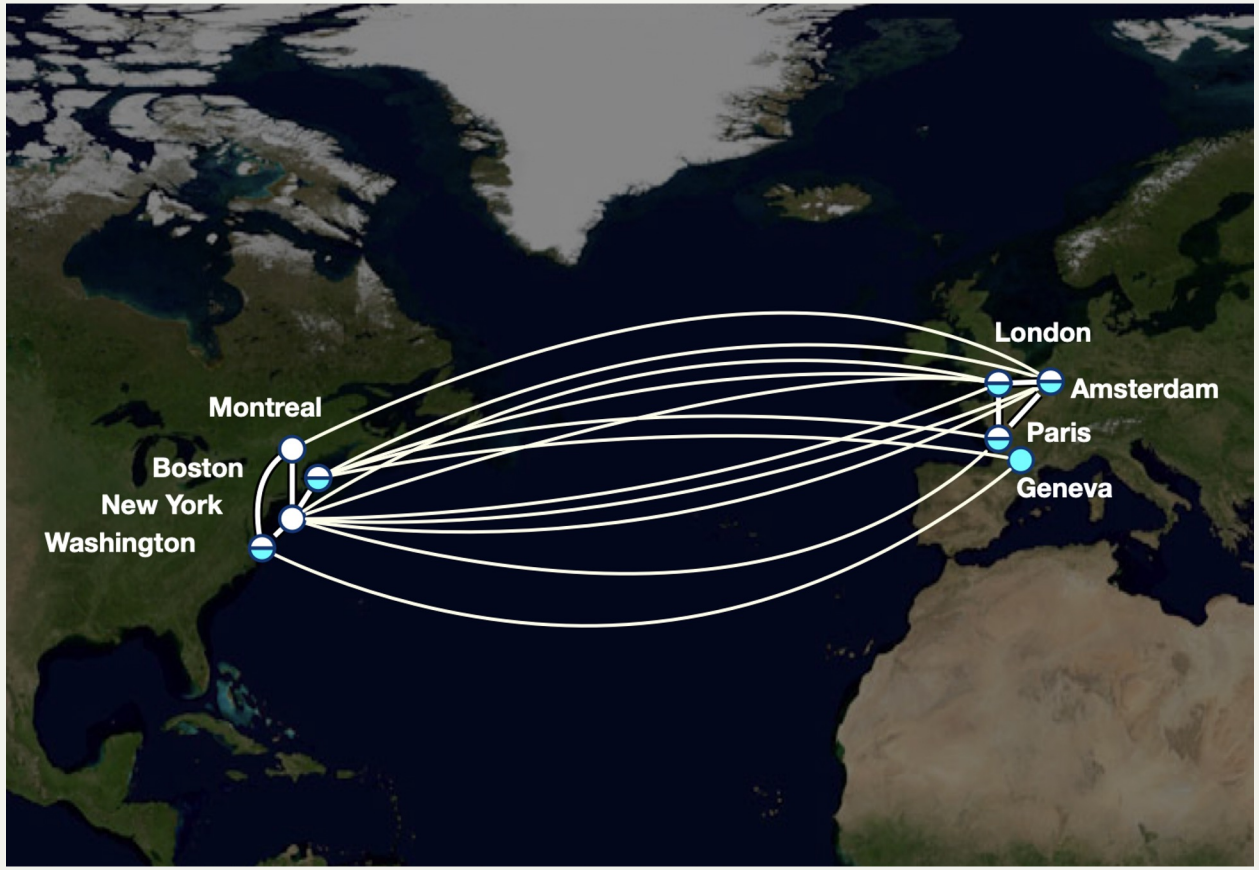

Highlight: Developing Areas of Community Engagement & Contribution

Internet2 NETWORK SERVICES

International Collaboration: NA-REX

- North American consortium of Exchange Point Operators, Connectivity Providers, Science Networks
- Improve coordination, establish common operating principles
- Leverage common toolsets, provisioning mechanisms
- Support dedicated links for experimental traffic, network research
- Initial links online:
  - Chicago and Seattle (dedicated)
  - Chicago and Los Angeles (shared)
- Links were successfully leveraged by Supercomputing NREs (Network Research Experiments)
- Currently planning for 2024:
  - Deployment of dashboard based on NetSage
  - Bring additional 400G links online
  - Restore API functionality to I2 Exchange Points (Q2)

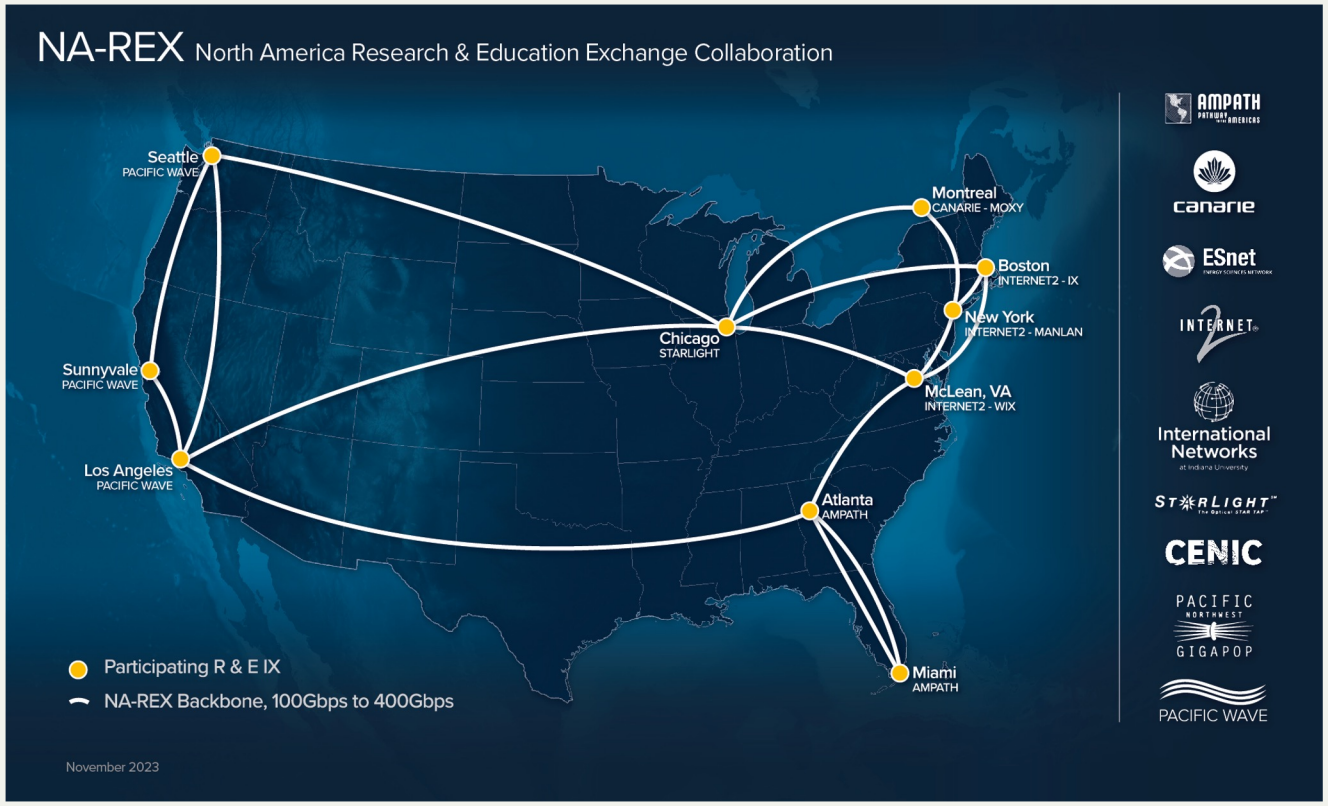

Internet2 Optical Networking

Spectrum Sharing

Internet 2 Optical Network

- 26,000 kilometers of optical fiber
- Ultra SMF28 based
- Designed to be flexible with multiple add/drop sites
- Open line system, not just locked into a single vendor for transponders

Spectrum Sharing in a Nutshell

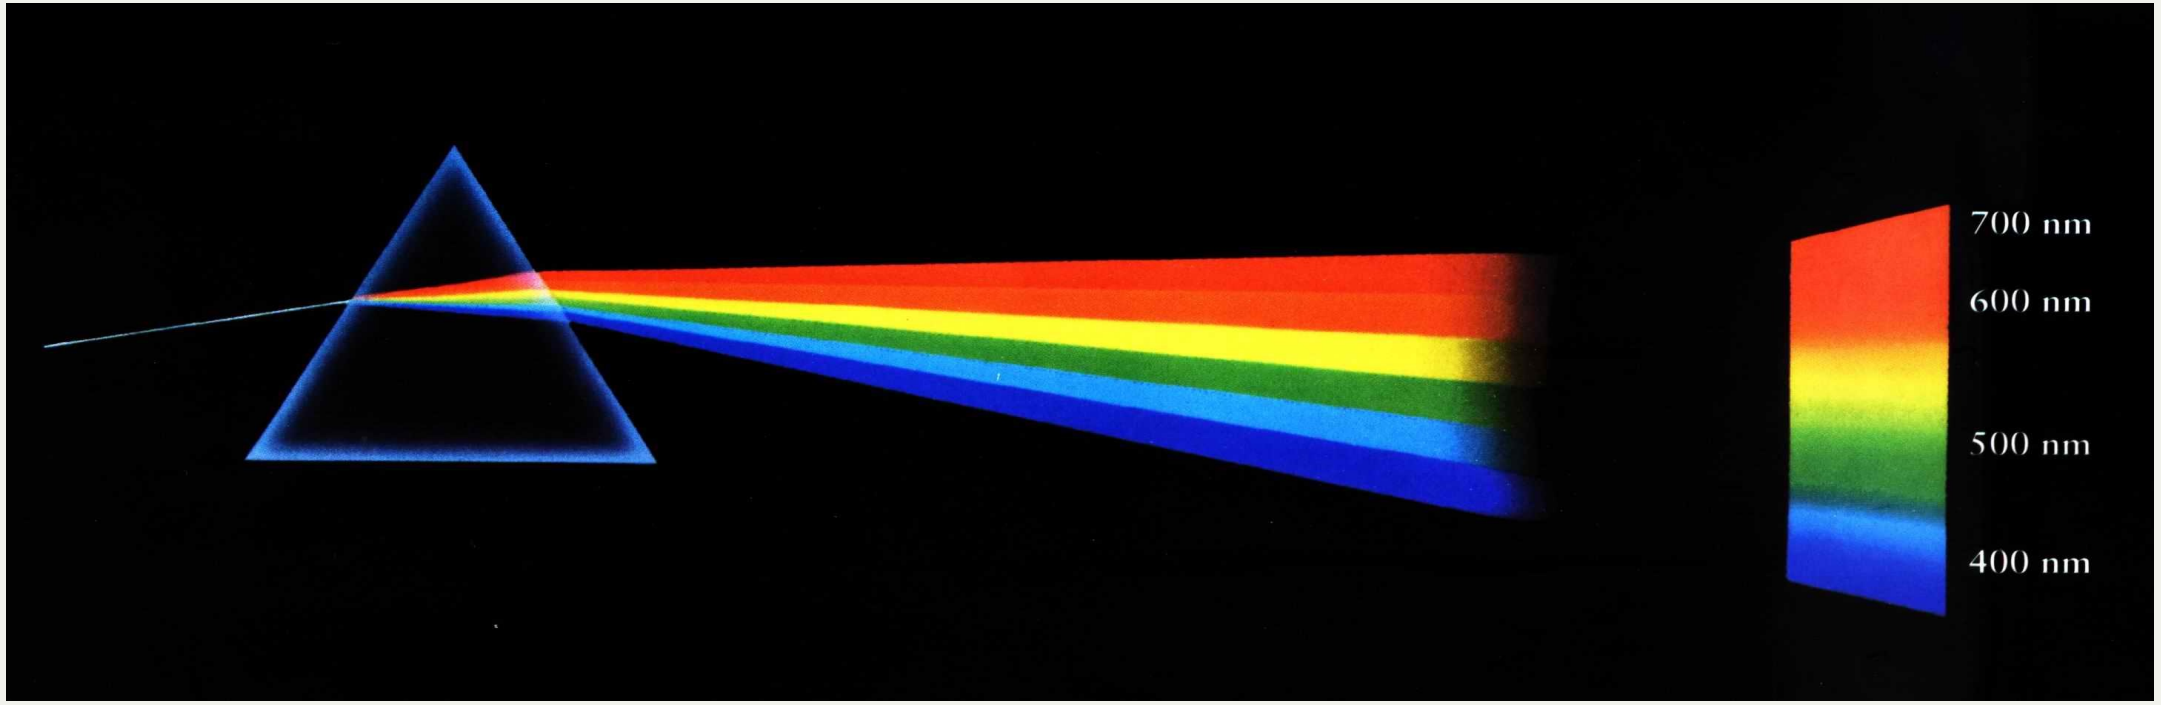

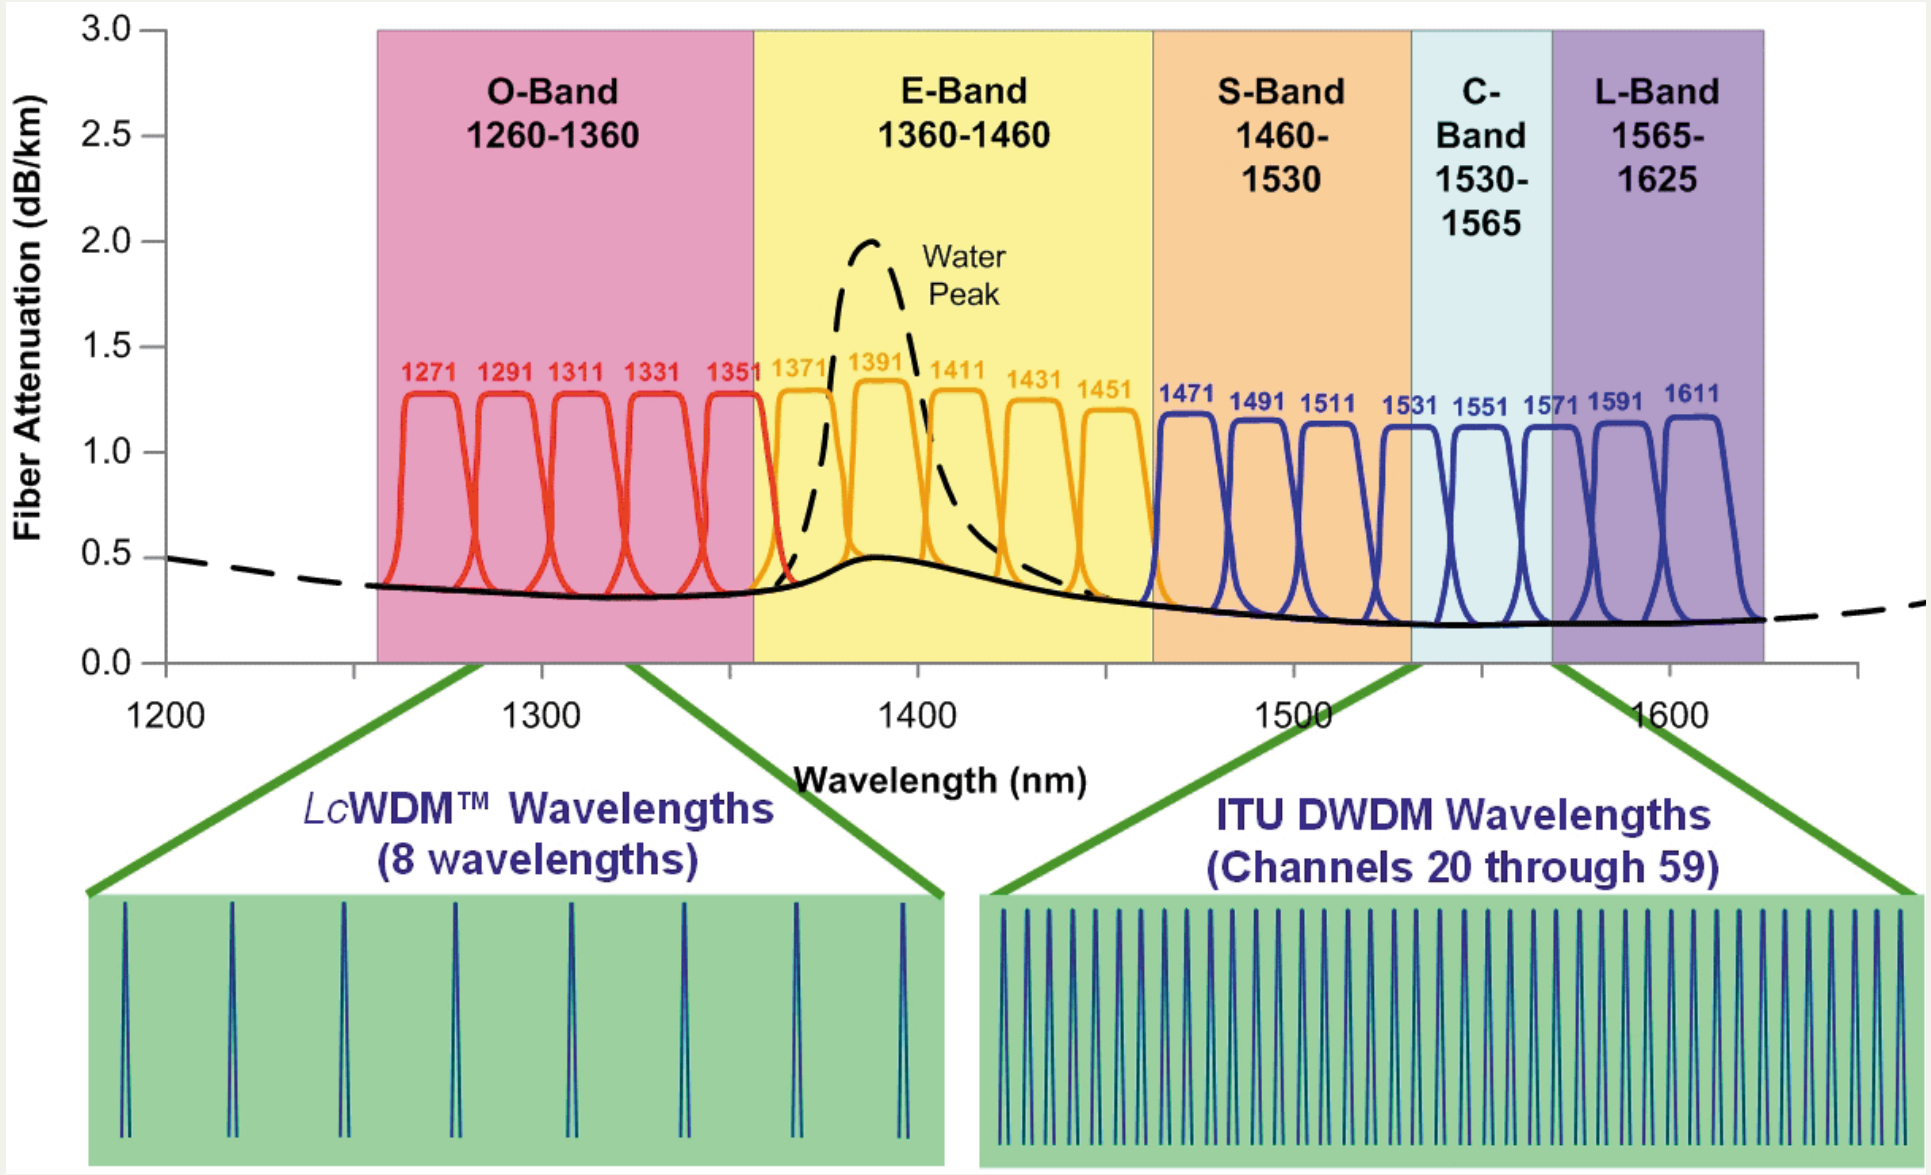
Fiber Attenuation (dB/km)
3.0
2.5
2.0
1.5
1.0
0.5
0.0
O-Band
1260-1360
E-Band
1360-1460
S-Band
1460-
1530
C-
Band
1530-
1565
L-Band
1565-
1625
Water
Peak
1271 1291 1311 1331 1351 1371 1391 1411 1431 1451
1471 1491 1511 1531 1551 1571 1591 1611
1200
1300
1400
1500
1600
Wavelength (nm)
LcWDM™ Wavelengths
(8 wavelengths)
ITU DWDM Wavelengths
(Channels 20 through 59)

Spectral Isolation: “Waves in Fiber”

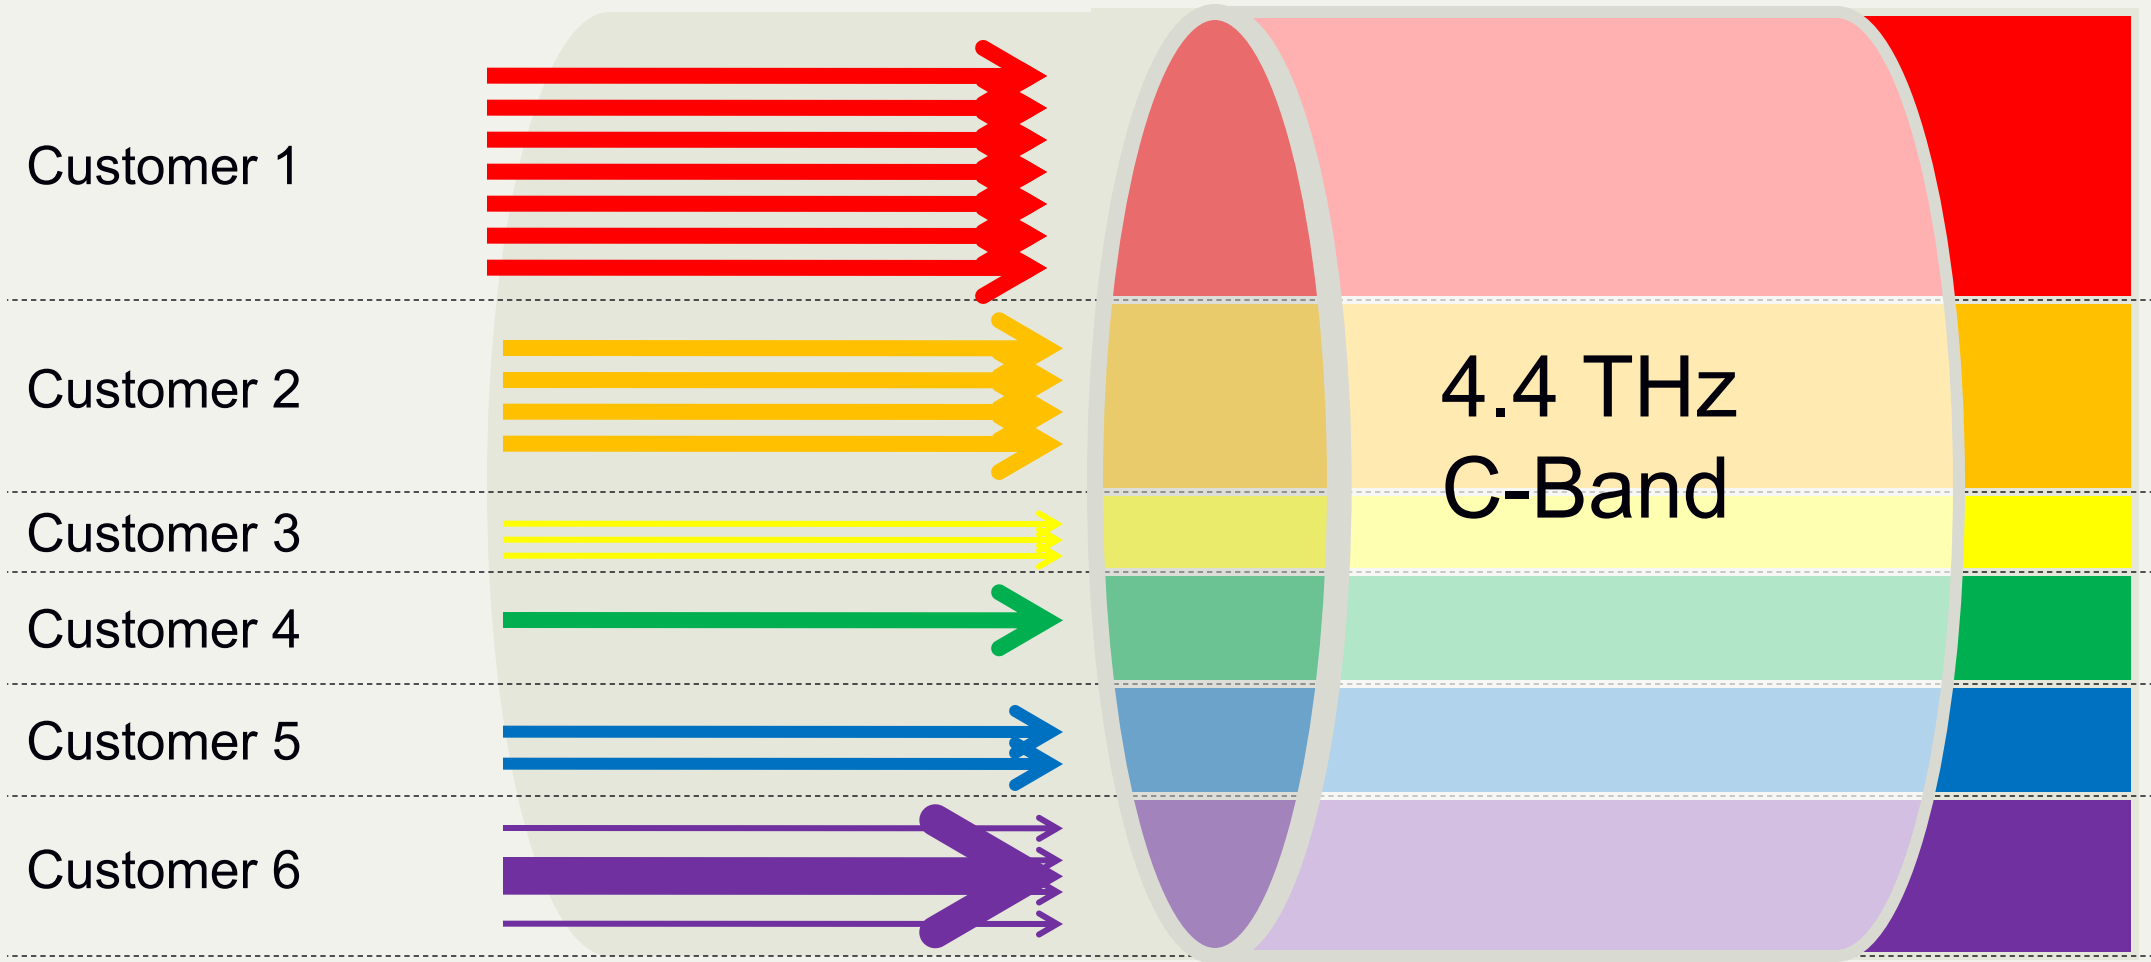

Optical Networking: Spectrum Sharing Concept

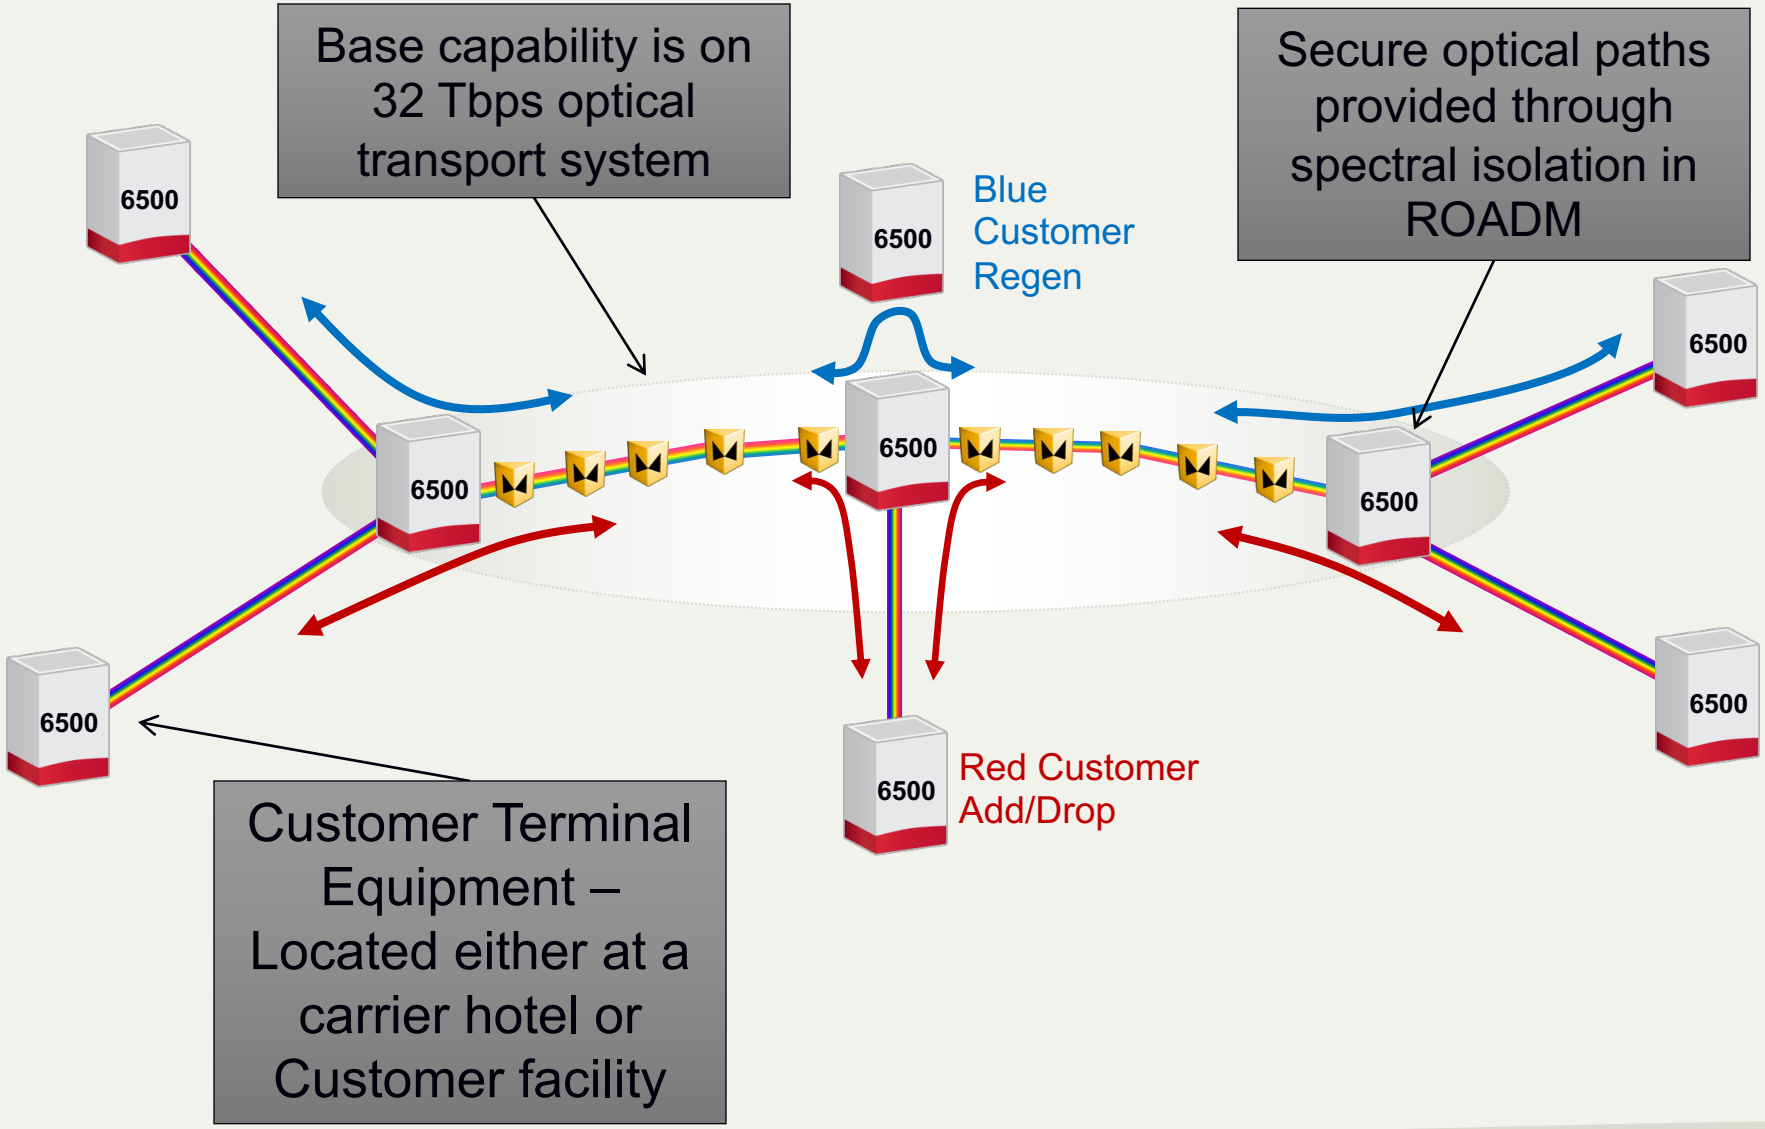

Managed wave services and foreign wave spectrum management

Internet2, post-NGI deployment, can take in flexible (flex-grid) foreign spectrum on-network via any of our ROADM sites.
This has already been put to good use.

There are some challenges, however.

- Channel allocation and control.
- Input power level management.
- Hard to automate both via internal line system controls or external resources (yet!).

Given that, is there a better way?

Managed wave services and foreign wave spectrum management

Solution: Internet2 managed Ciena Waveserver 5 platforms.
Example: Connector network needed to upgrade several of their 10g long-haul links to 100G.

- They had already started the acquisition process to obtain several Waveserver 5 chassis with WaveLogic5 modules to provides this.
- They also had contracted with Internet2 to provide spectrum for these links.
- Originally, they were to manage their hardware plugged into our line system.
- Given some of the distances involved, **optimization of the power levels is crucial**.

While we can hand-tune this, this could change over time. Any spectral management for spectrum defragmentation would have to be managed.

The overall solution for this was very simple, in addition to the fiber cross-connects, we also ran copper cross-connects – management port.
This provided management access to the connector's Waveservers, and Internet2 assumed management of the member devices.
This allowed for better end-to-end visibility across the entire line system and endpoints for management and required minimal additional effort.

This isn't to say we're not in the foreign spectrum business!

Other managed wave services.

In 2022, Internet2 was asked to provide 400G redundant wave services for a long-haul DCI project.
Emphasis was placed by the end-user on reliability, redundancy and diversity of the service provided.

Challenge: One end was not currently co-located in a facility that housed an I2 point-of-presence.

What we did have, however, was available dark fiber between facilities, but had limited space for equipment.

This meant that a full ROADM site was out of the question.

Thankfully, it turns out that the Waveservers have an in-band remote management mode. We decided to sub-tend the remote node off the local I2 POP Waveserver.

The waves were then regenerated on the POP Waveserver via redundant modules and sent across the line system to I2 Waveservers in the destination POPs where the end-user was already co-located with us.

This has proven to be highly reliable and is an option that is available to the community if they need it.

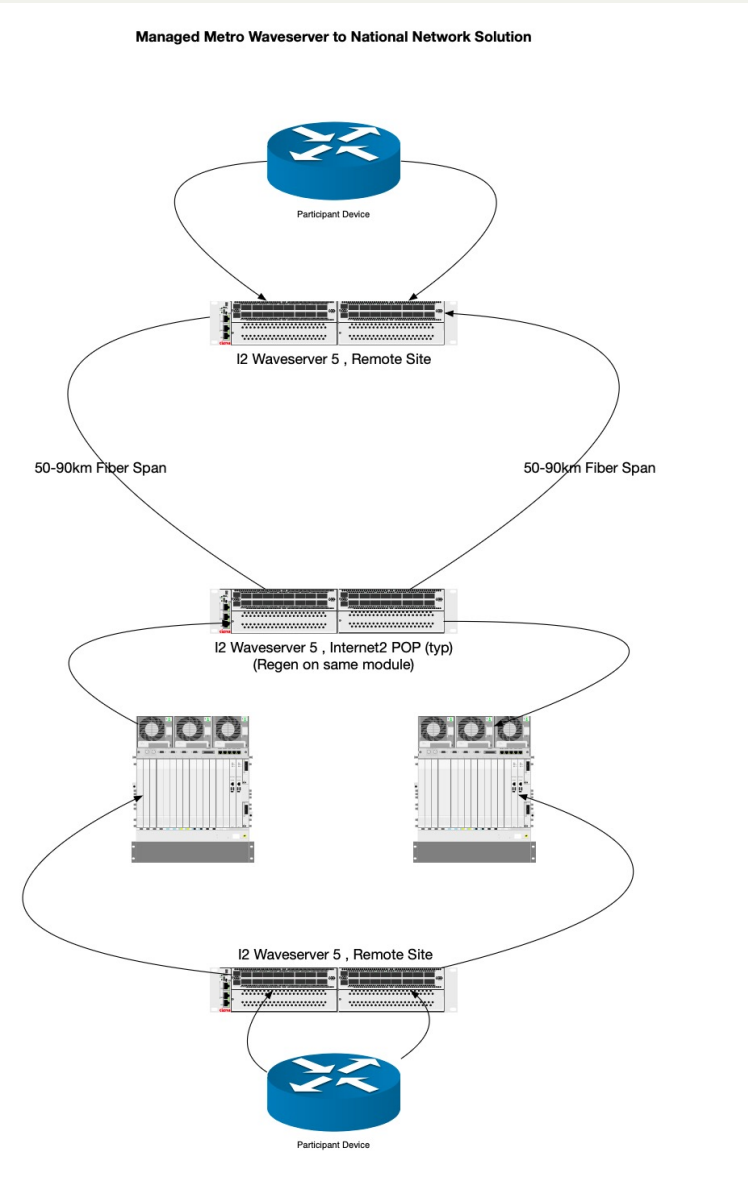

Managed wave services

What other strategies could we be doing utilizing?

Well, given the advent of coherent pluggables in the market, what if instead of managed remote Waveservers, we simply interconnected router-to-router with coherent optics in locations where additional connectivity diversity is desired?

This is very useful, again, in cases where it does not make economic sense for one network or the other to co-locate equipment in the same POP when dark fiber is available between two nearby POPs occupied by the interconnecting networks, especially if it’s under 80-100km!

Coherent Pluggables:

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.

Implementation as of 3/1/2024:

New Boston pop (BOST2) is now online, and connected purely with Cisco 400G-ZR+ “bright” pluggables to adjoining POPs (Hartford and Albany). No Waveservers were deployment when the pop was rolled out. Ciena WaveLogic Nano still to be tested, this is dependent on host platform support being fully in place.

Coherent Pluggables, continued:

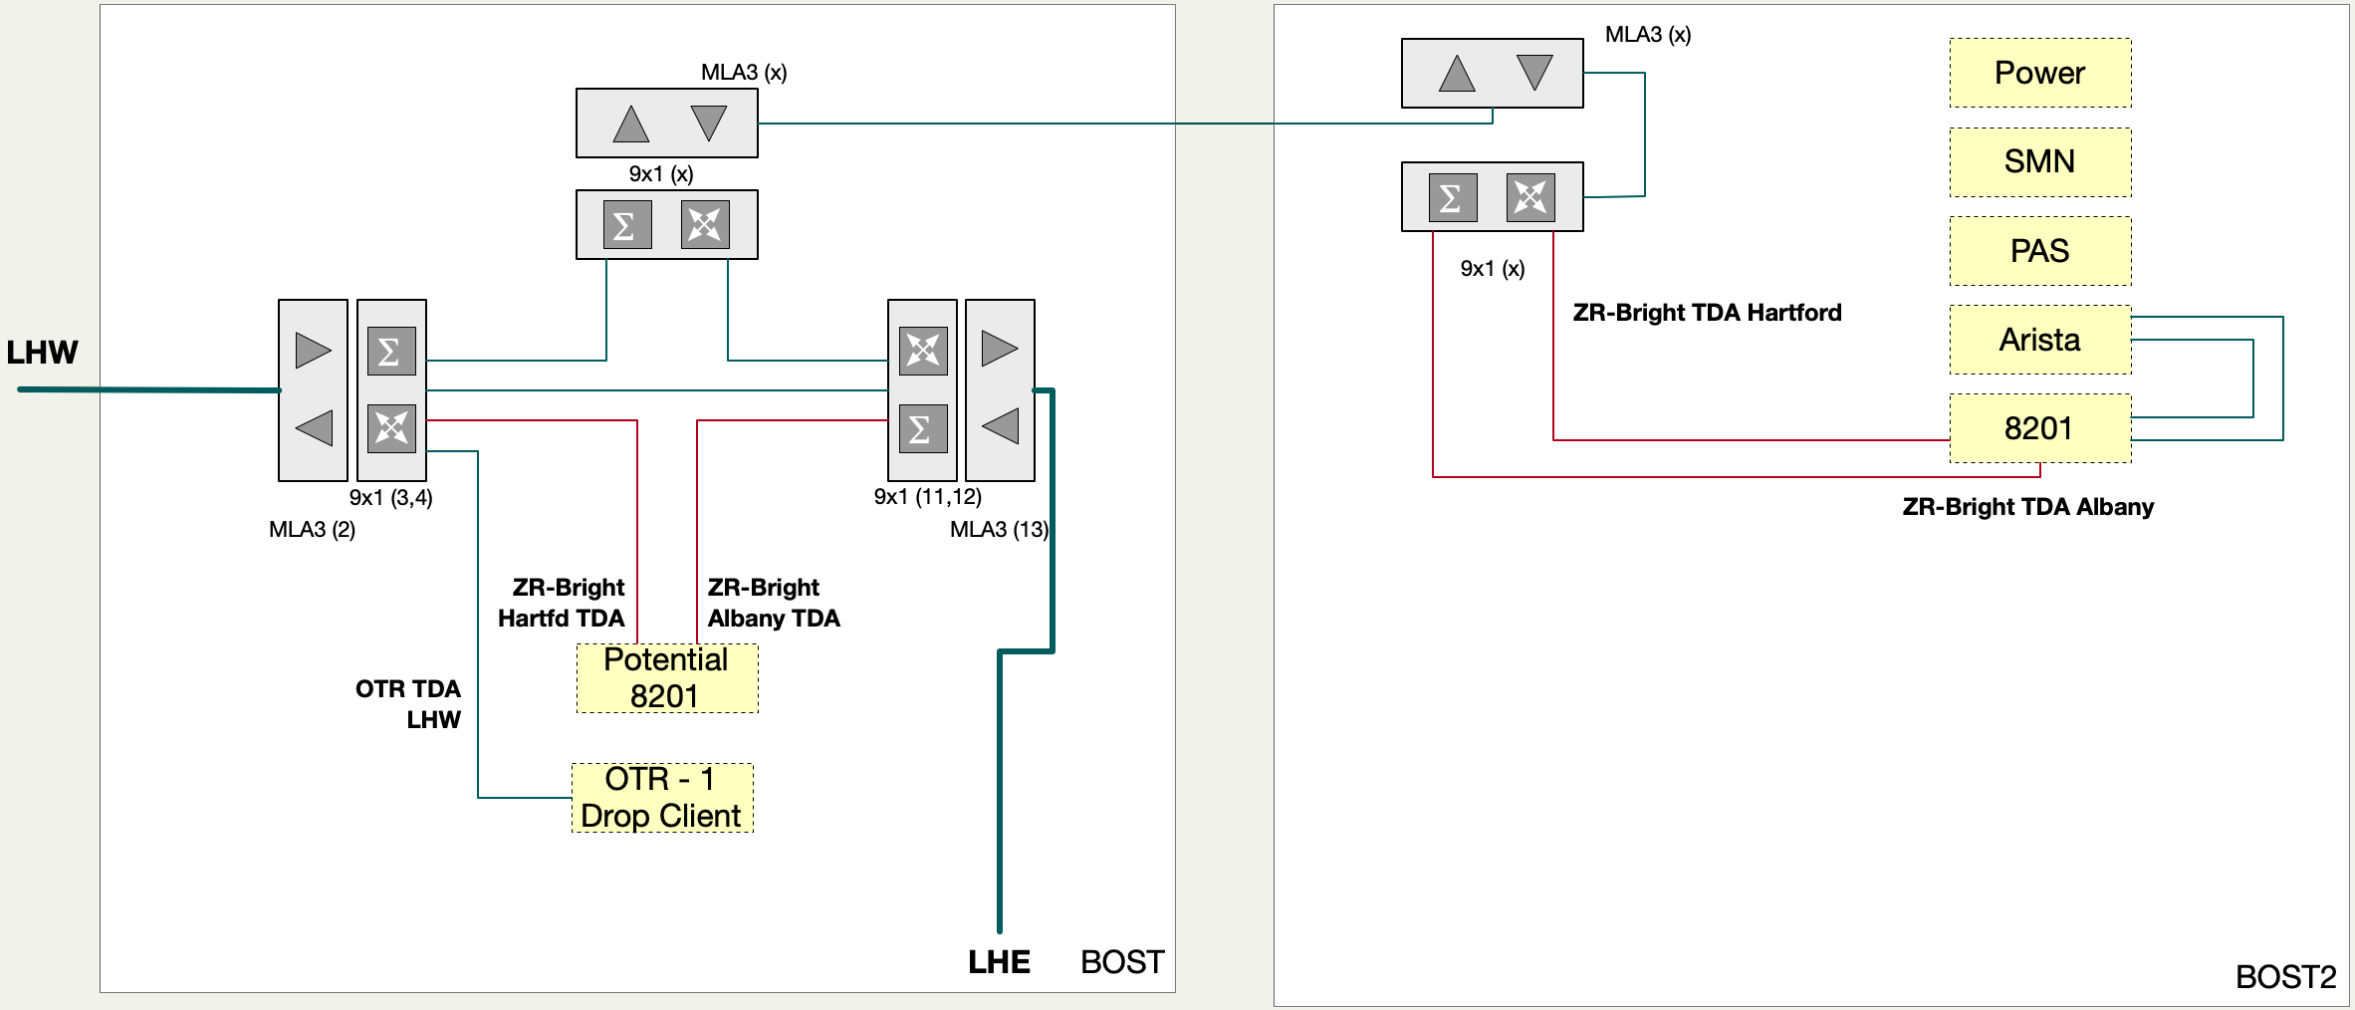

Potential Use Case

Long-haul research exchange support using large spectrum channels

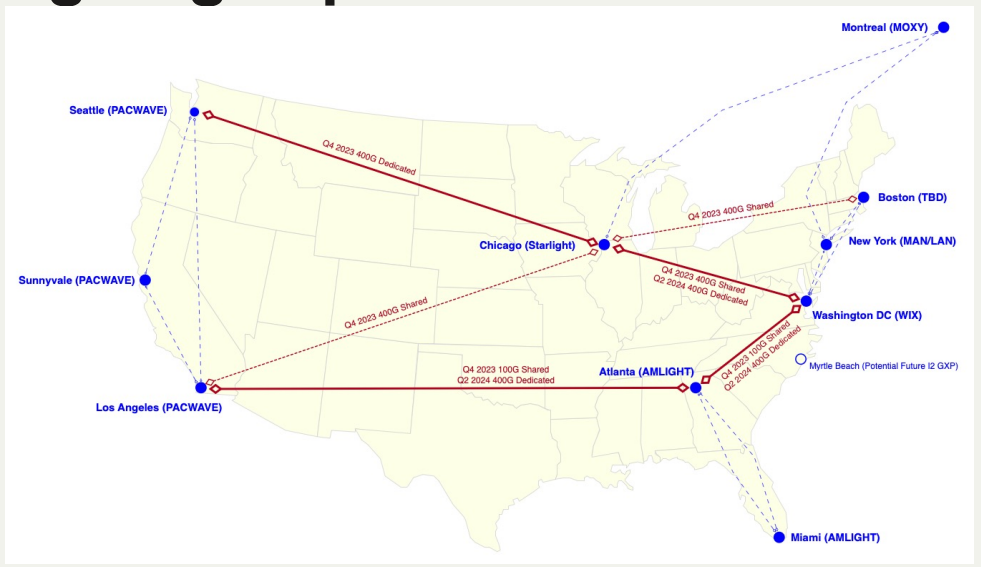

- As part of SC ‘23, we rolled out a new very long-haul link on the network, connecting Seattle (PACWAVE) to Chicago (Starlight).
- This link is now permanently in place.
- Composed of 2x200 waves on a Ciena WaveLogic 5 modules, over a distance of 4900km
- This is mostly due to distances involved, the SNR margins were just shy of good enough to run as a single 400G wave.

- New connection turned up as of 3/5!
- Starlight to Los Angeles 4561km, unregenerated, single 400g wavelength, with 2dB of margin (this is very good given the distances

Why do we run these links as unregenerated?

Latency, and to keep space/power/cooling open in intermediary POPs.

LARGE PLATFORM FLEX

Current Platform Model

NGI Small

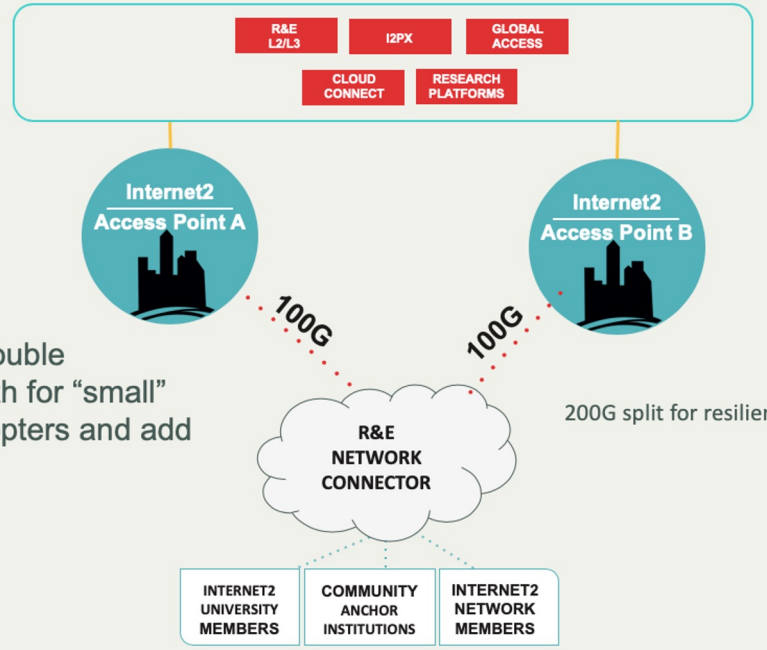

NGI could double the bandwidth for "small" Platform adopters and add resiliency

200G split for resiliency

NGI Large
(Research Intensive)

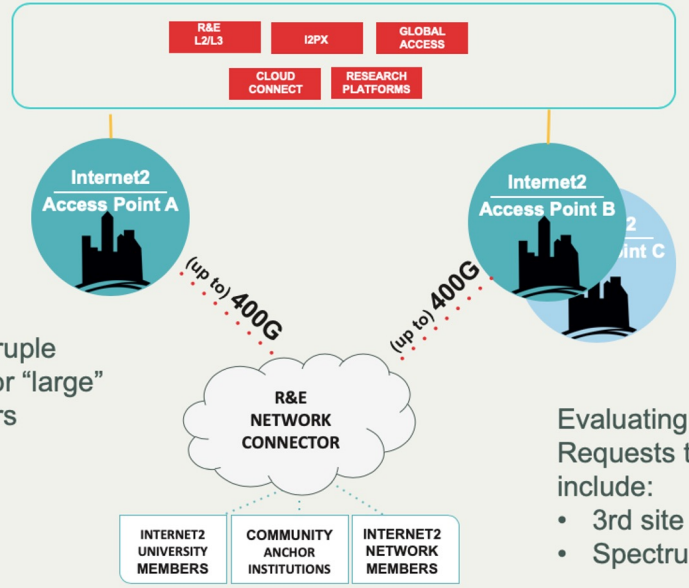

NGI could quadruple the bandwidth for "large" platform adopters

Evaluating Requests to include:
- 3rd site
- Spectrum

Large Platform Flex

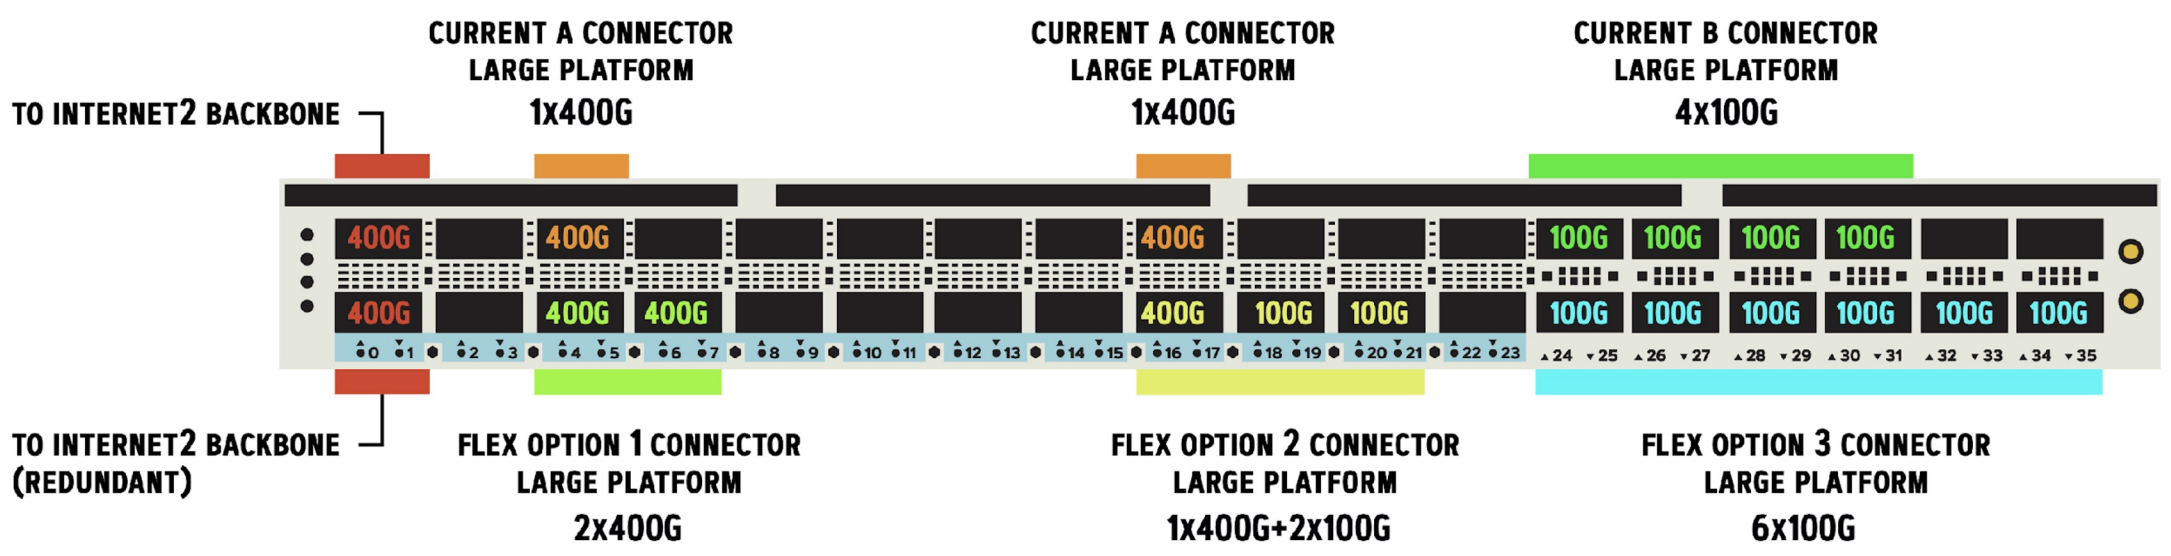

Flex Ports are managed via NSO and are individually configurable!

Thank you!

Questions:

Christian Todorov ctodorov@internet2.edu