



# Internet2 & NA-REX Update GRP4 - 10 October 2023

---

**Matt Zekauskas**  
Senior Researcher  
Network Services

# What is Internet2?

Not-for-profit computer networking consortium led by members from the research and education communities, industry, and government.

500+ members including:

- 251 institutions of higher education

- 9 partners and 76 members from industry

- 100 research and education networks or connector organizations

- 67 affiliate members

In general, Internet2:

- Provides the Internet2 Network - including an optical fiber footprint and cloud connectivity

- Operates international exchange points on the east coast of the United States; helps fund international links and connectivity

- Provides trust and identity services

- Organizes community events

- Engages programmatically in key technology development activities and security topics

# Approach Driven by Community Input

Coordinate with worldwide networks such that they function as systemic, best-in-class resources - ideally coupled with edge computing and storage systems

Provide programmatic support and infrastructure for:

- Grant funded programs. Specifically, experimental deployment of new capabilities using testbeds; which broaden impact and participation through data network-centric activities.
- Data-intensive sciences like LHC; other science and industry R&D program (high bandwidth)
- Cloud-centric applications supporting high-availability applications
- Integrated security, specifically routing integrity (MANRS, DDoS, RPKI-ROV)

Provide software, automation, and APIs which allow for integration of network resources into a global, end-to-end fabric that flexibly allocates, balances and conserves the available network resources

Support regional caches/data lakes and access to network overlays with intelligent control & data planes (e.g. FABRIC, Open Science Data Federation (OSDF))

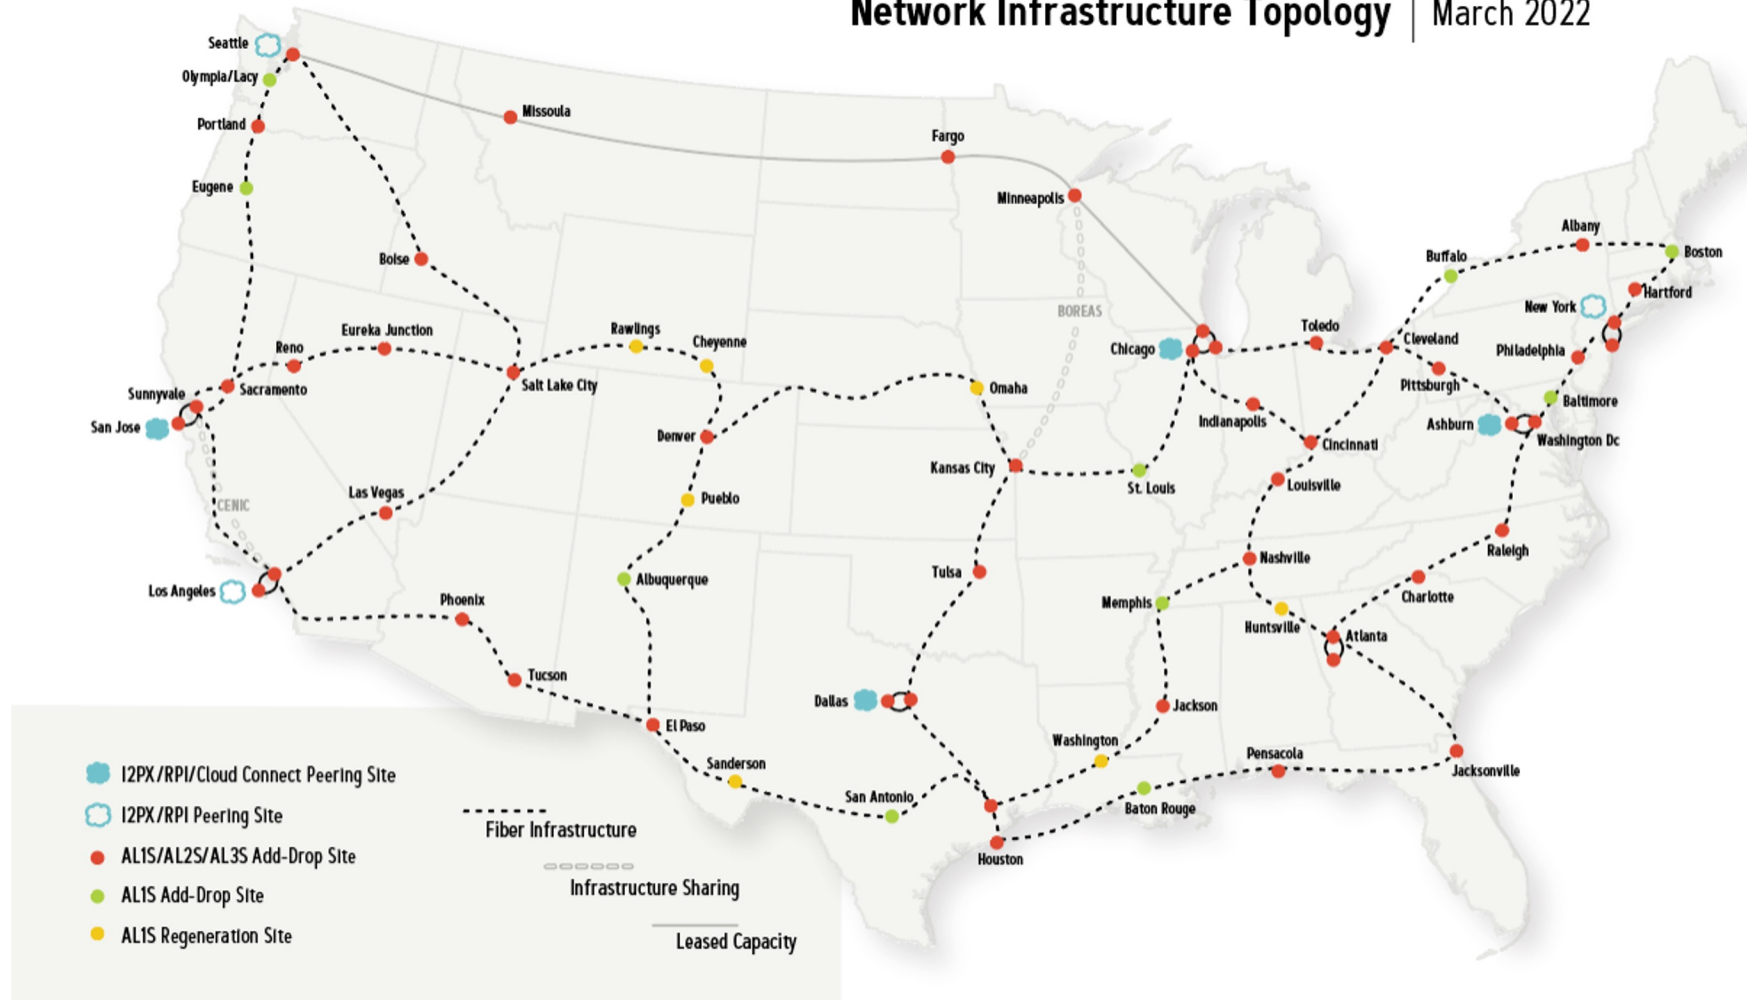
Support experiments with fully programmable components (P4, PINS; SRv6; 5G) and operations platforms (NRP; global SENSE Testbed)

# Some examples of things we are working on...

Engagement and Innovation	Collaborate Domestically (NA-REX) and Internationally (GREN, GNA-G)
Sustainability	Vendor Engagement, Lower Operating Costs Leverage Low Power Chassis, ZR+ Optics
Data Movement, Scalability	Build 400G+ for Infrastructure & Exchange Points High Scale at Edge / Cloud
Security	Promote Route Integrity (RPKI, ROAs)
Data Lakes, Caching, GPUs	Open Science Data Federation (OSDF); Testbeds
Programmability / Topology Visibility	Insight Console, APIs, Extensions supporting NSI, SENSE Enable core with SR-TE, BGP Classful Transport
Performance Assurance	Test Resources, Access to Platform
Programmatic Reporting (IRNC)	Flow Tools, Reporting

# Network and Service Updates

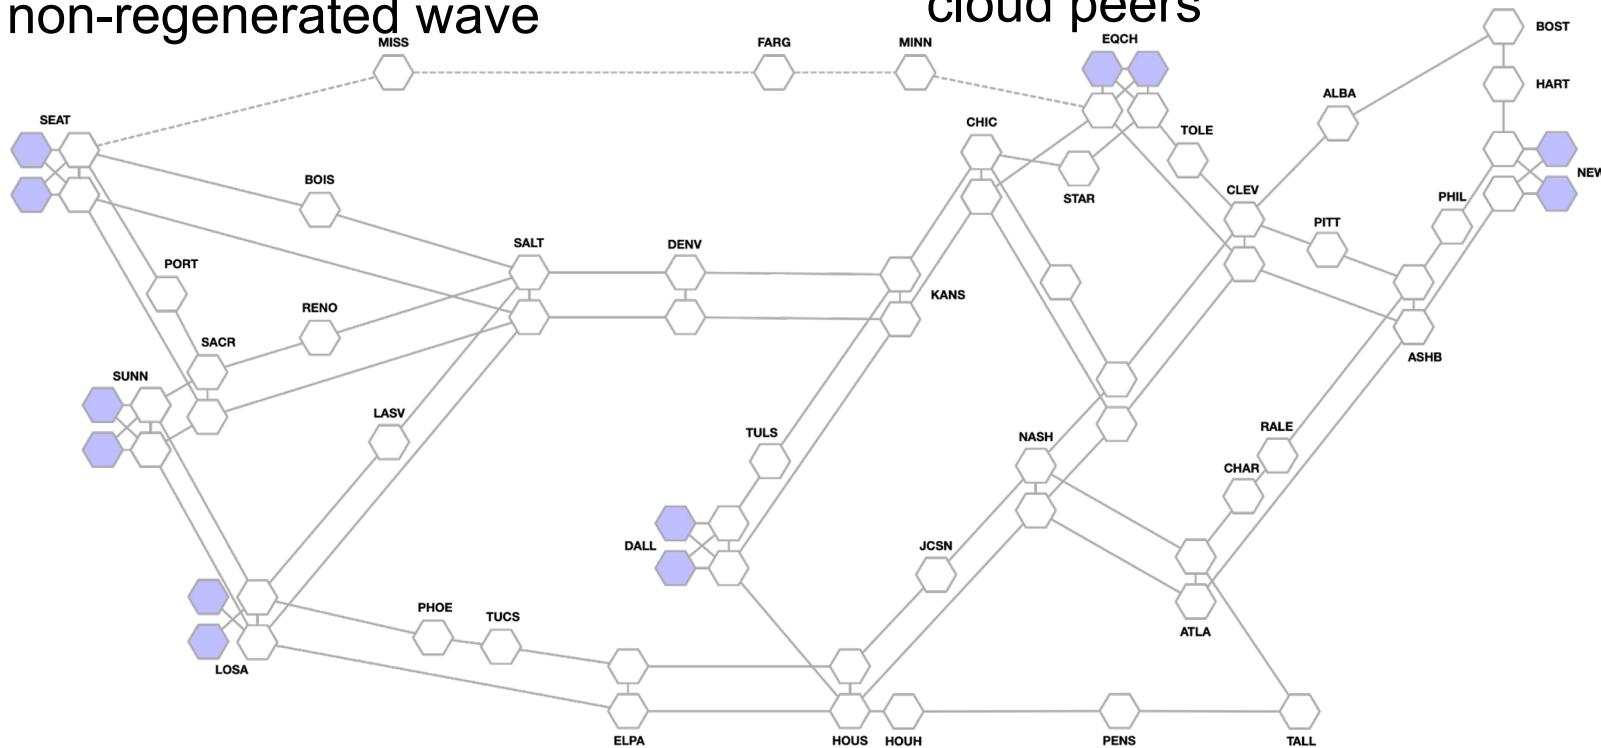
Network Infrastructure Topology | March 2022



# Network and Service Updates

- Fifth generation of the Internet2 backbone
- Ninety-four 400 Gbps Backbone links
- 27,600 Tbps of deployed capacity
- 1.6 Tbps available coast to coast
- Each link is on non-regenerated wave

- Disaggregated Switching/Routing Platform  
Cisco 8201/8202 - 77 Routers, 47 Sites
- Cisco Network Services Orchestrator (NSO)
- Redundant/resilient routers, dual-connected cloud peers



# Network and Service Updates

- 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



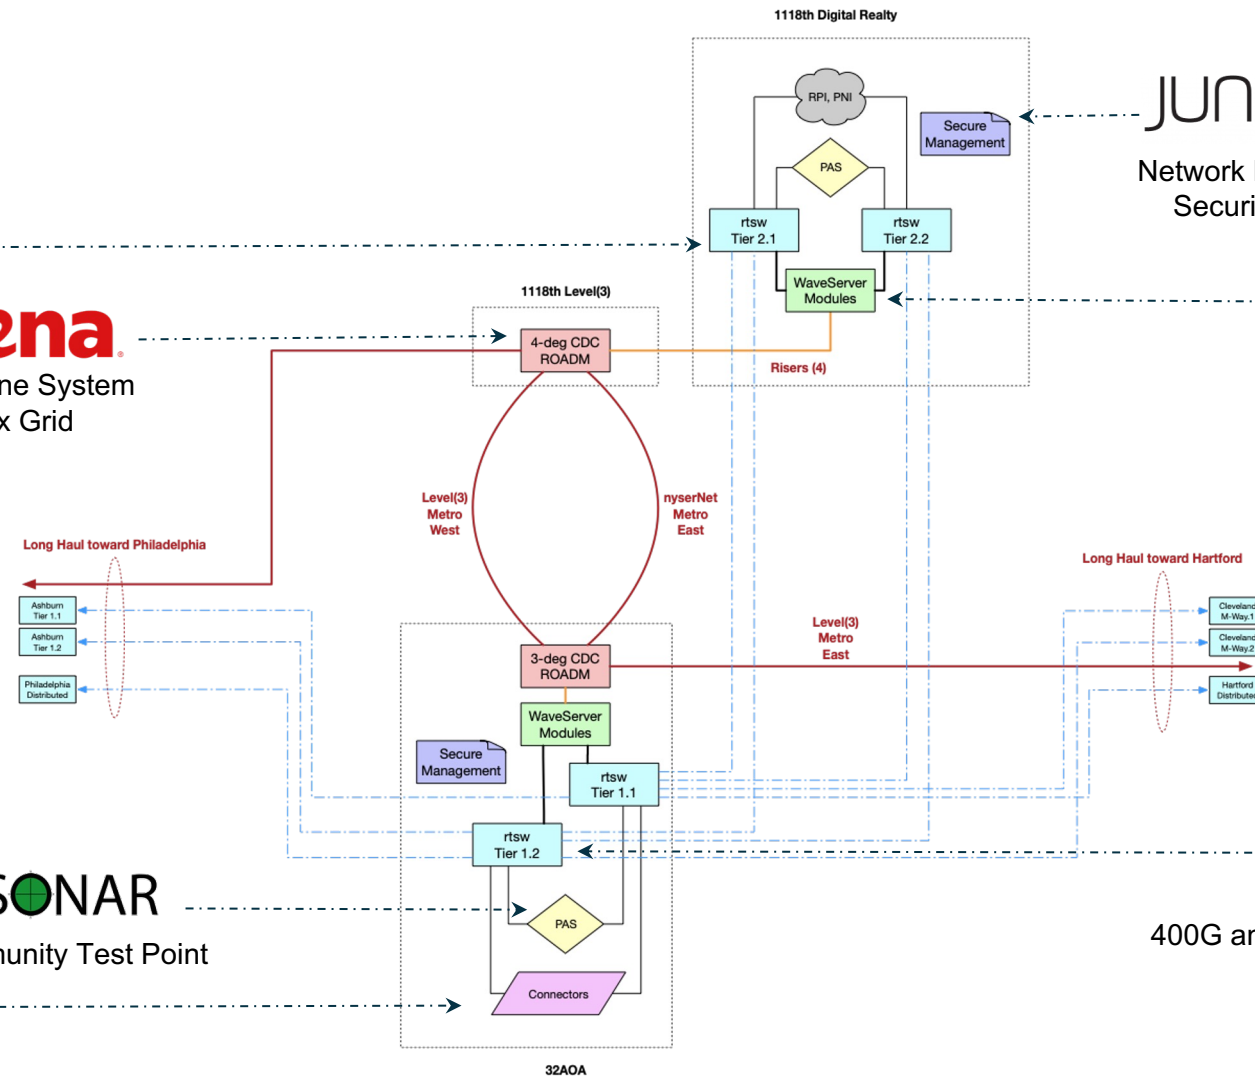
# Network and Service Updates

**CISCO**  
Terabit-Scale Cloud Edge

**ciena**  
Open Line System  
Flex Grid

**ciena**  
400G - 800G Transponders  
Long Haul and Metro

**JUNIPER**  
NETWORKS  
Network Management  
Security Overlay



**CISCO**  
400G and 100G Community Access  
All I2 Services

**ARISTA**  
400G Global Exchange Points

**perfSONAR**  
2x100 Community Test Point

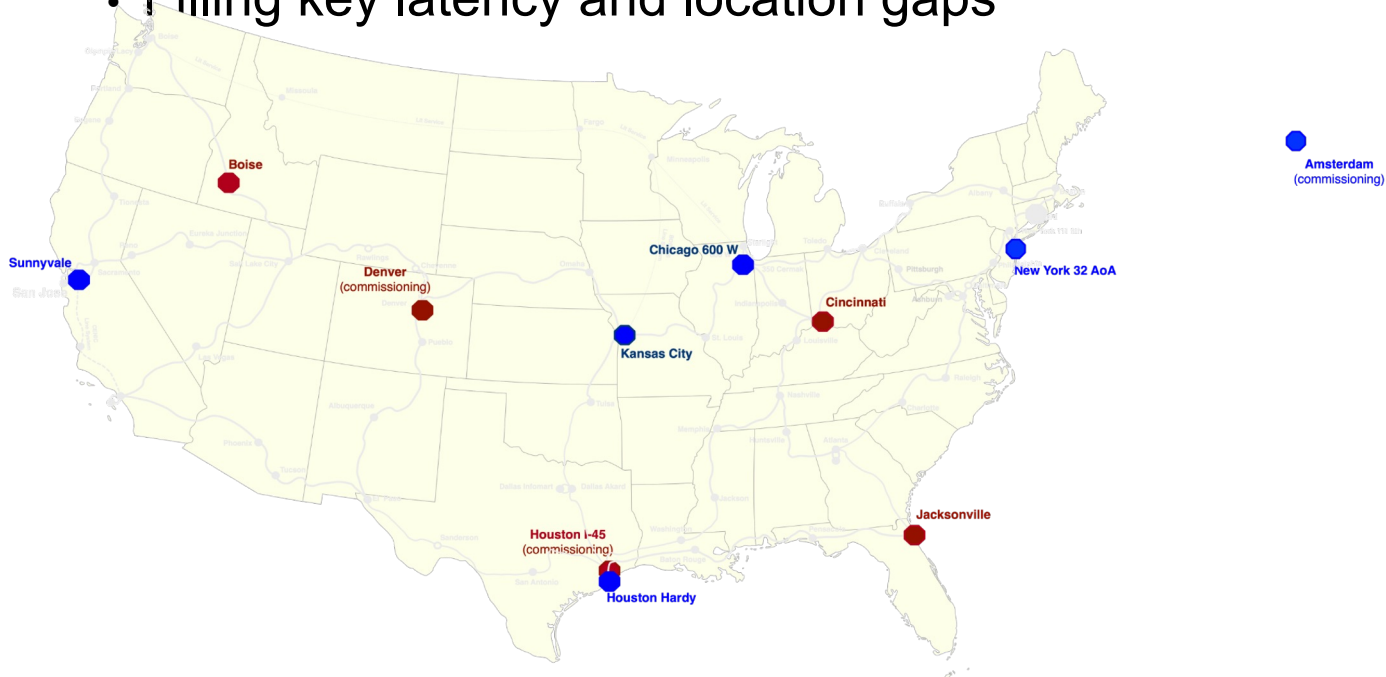
**GRP4**  
10 October 2023



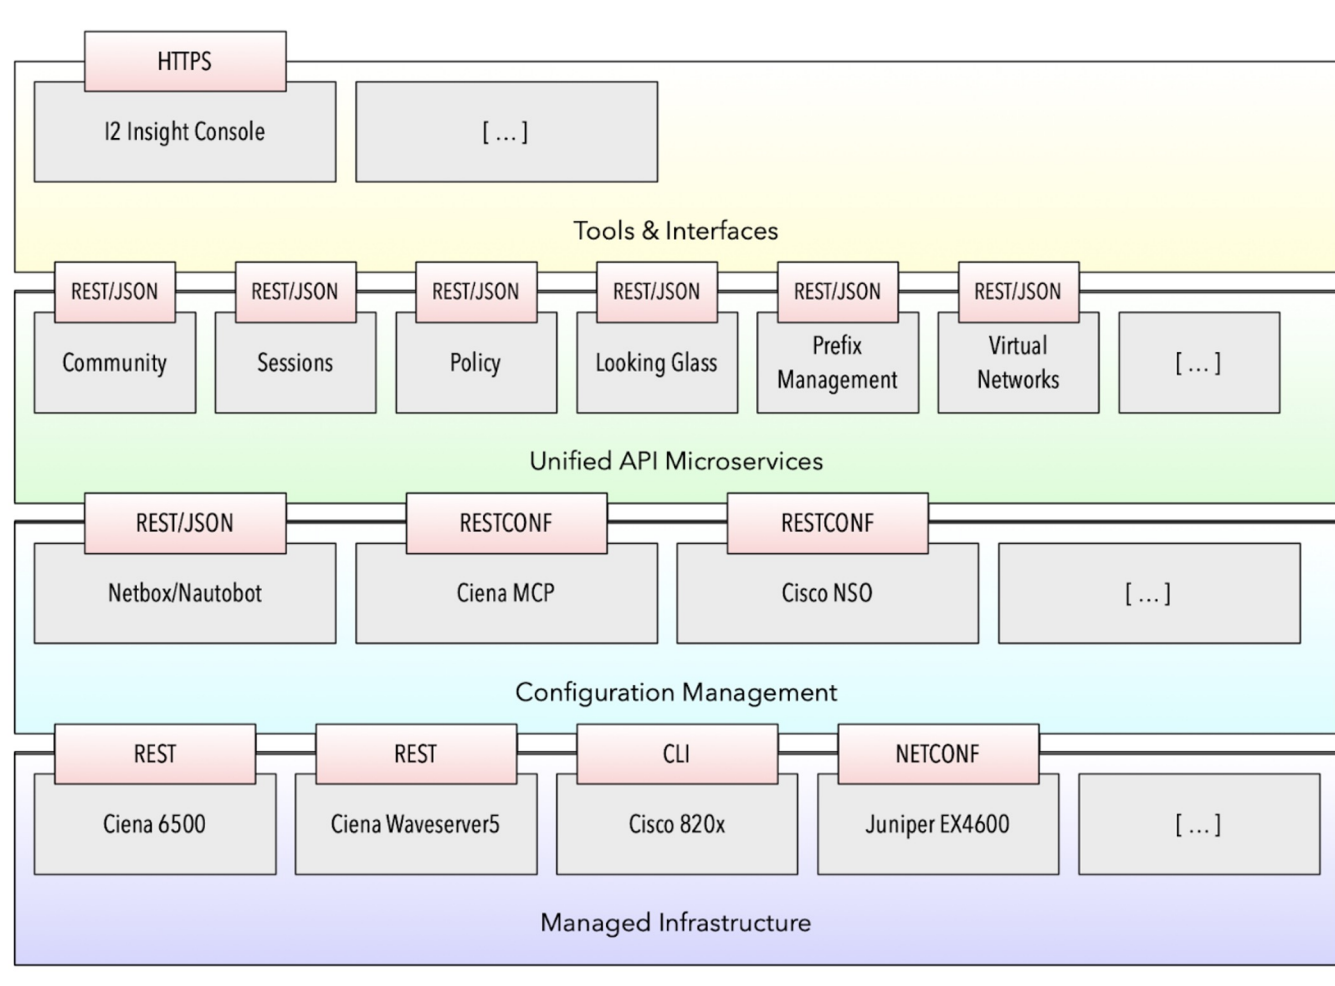
# Network and Service Updates

## 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



# Backbone and Service Updates



# Backbone and Service Updates

The screenshot displays the Insight Console interface for a virtual network. The main area shows four virtual routers arranged in a 2x2 grid, connected via Internet2 links. Each router is associated with an IPv4 address, an ASN, and a subinterface. The provisioning status for all routers is 'Pending'. The network space sidebar on the right provides details for the 'OCI-FastConnect' network, including its title, owner, last modified date, and creation date. It also features buttons to 'Add Virtual Router' and 'Add Virtual Switch', and a section for 'Collaborators' which currently shows 'No collaborators'.

Insight Console

Search Docs, Organizations, & Virtual Networks

Usability Testing Provide Feedback

Community Interfaces **Virtual Networks** Looking Glass

Matt Zekauskas | Sign out

Virtual Router VNRROUTER-10089

BGP Status IPv4 Up Connection **Live** Details

Internet2 New York, NY

IPv4 10.221.10.9/29

ASN

ASN 55038

Subinterface

Provisioning Status Pending

BGP Status IPv4 Up Connection **Live** Details

Internet2 Cleveland, OH

IPv4 10.221.10.6/29

ASN

ASN 55038

Subinterface

Provisioning Status Pending

BGP Status IPv4 Up Connection **Live** Details

Oracle US East (Ashburn) Internet2 Ashburn, VA

IPv4 10.221.10.25/29

ASN

ASN 31898

ASN 55038

Subinterface VLAN 101 Bundle-Ether6 agg3.ashb

Provisioning Status Pending

BGP Status IPv4 Up Connection **Live** Details

Internet2 Ashburn, VA Oracle US East (Ashburn)

IPv4 10.221.10.22/29

ASN

ASN 55038

Subinterface VLAN 102 Bundle-Ether6 agg4.ashb

Provisioning Status Pending

Virtual Network Space

Title OCI-FastConnect

Owner

Last Modified 2023-09-26T21:59:52.015996+00:00 by OESS

Created 2023-09-26T21:59:52.015987+00:00 by OESS

Virtual Devices

Add Virtual Router Add Virtual Switch

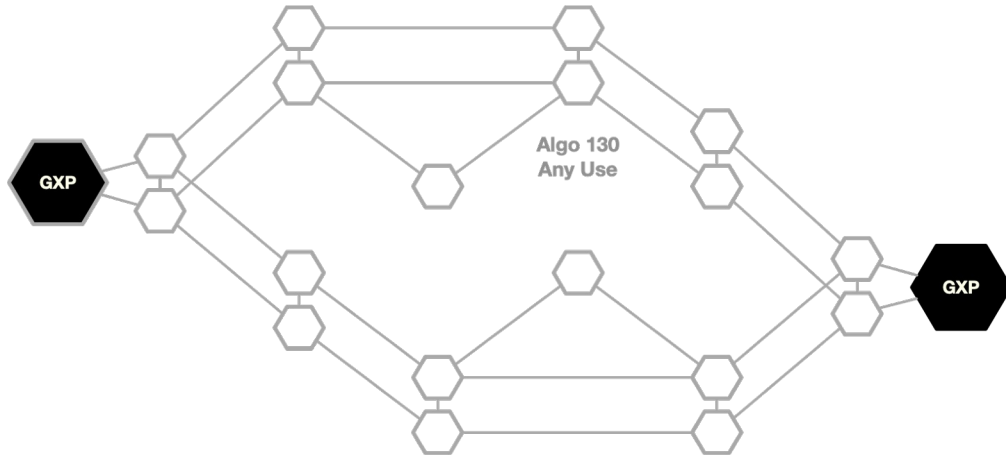
Collaborators

No collaborators

# Backbone and Service Updates

The screenshot displays a 'Layer 3 Connection (BGP Peering)' configuration page, which is currently 'Live'. At the top, a diagram shows two vertical server racks representing 'Internet2' and 'Oracle', connected by a yellow horizontal line. Below this, the configuration is split into three main sections: 'Internet2', 'BGP Status IPv4 Up', and 'Oracle'. The 'Internet2' section includes fields for IPv6 Address, IPv4 Address (10.221.10.30/29), ASN (55038), and Interface (Oracle US East (Ashburn) Cloudconnect, Bundle-Ether6.101 agg3.ashb). The 'BGP Status IPv4 Up' section shows Description, MTU (9000), Max Bandwidth, Display Position (3), BFD (Enabled), and BGP Key (\*\*\*\*\*). The 'Oracle' section includes IPv6 Address, IPv4 Address (10.221.10.25/29), and ASN (31898). At the bottom, there are tabs for 'Statistics' and 'Provisioner', and a 'Subinterface\*' label.

# Sidebar: Experimental Technologies Under Study




## Flexible Algorithm (Flex-Algo)

- Custom Prefix-SID's → "Algo"
- Nodes can participate in any algorithm or multiple algorithms
- Algorithms propagated through IGP
- Algorithms deployed by provisioning system such as NSO
  - allows for **SRTE without controller**
- Example of disaggregated architecture with all nodes in same Algo

## International and DIS Example: constraining traffic, differentiating paths

- Constrain / move any-use traffic to certain paths (Algo 130)
- Provide dedicated links for specific uses (Algo 131)
- Provide N and S differentiation for backup (Algo 128+129)
- Traffic "color" (**BGP-CT** or **BGP-CAR**) or label can be assigned to particular Algo at ingress
- Theoretically, traffic can be dynamically drained off of certain segments for DIS burst use



# International-Related Areas of Effort

**GXPs to 400G** - The upgrade of Internet2-operated global exchange points (GXPs) to support 400 gigabit technologies and community-driven automation.

**400G R&E transoceanic link** - Implementation of the first 400 gigabit links to connect Europe and the UK to the continental United States.

**NA-REX** - A series of related, partner driven activities intended to knit together exchange point operations, technologies, and features in the domestic United States.

- Common GXP features

- Experimental Links

- Internet2 Core Integration

- Software, Automation, and APIs

**Programmability and data caching** - Targeted implementation of P4 infrastructure and expansion of Open Science Data Federation (OSDF) caches on both domestic and international footprints

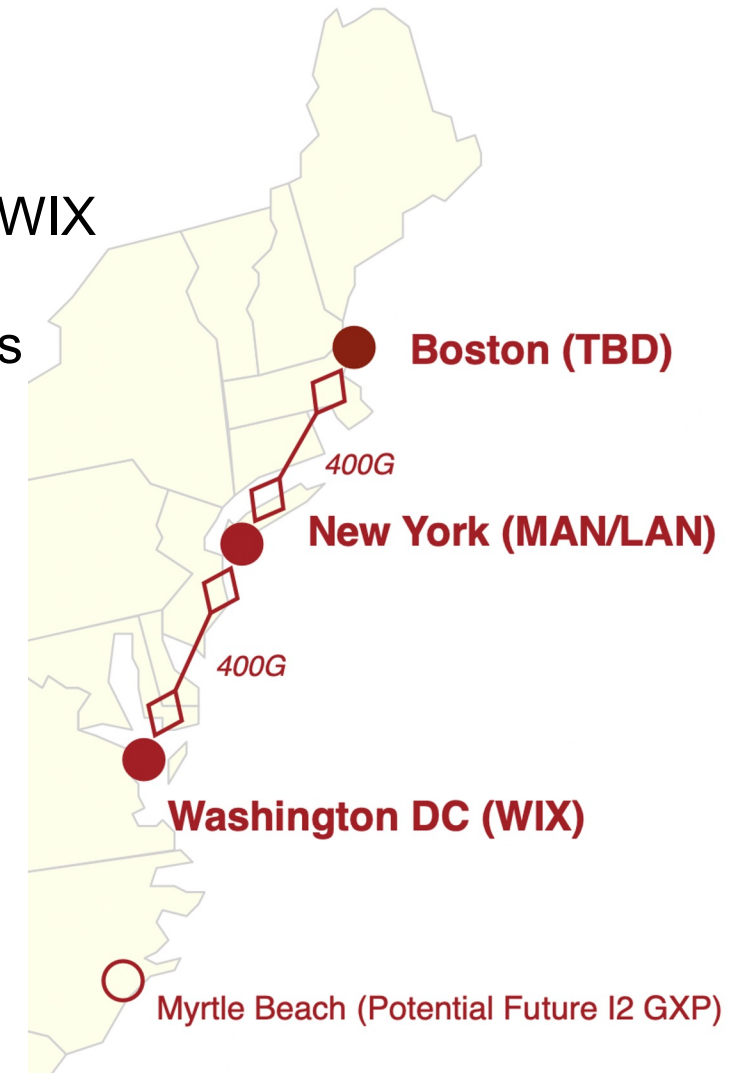
**These are best considered as a systemic approach toward support for data-intensive science!**



# Expanding 400G to Global Exchange Points

## Automating and Expanding Exchange Point Functionality; Match Community Development

- Hardware improvements to support 400G, including MAN/LAN and WIX
  - Arista DCS-7280PR3K-24 400G Switch
- 400 Gbps dedicated links between all three exchange point switches
  - ZR+ Optics
- Protocols
  - EVPN MPLS for L2 services (VXLAN Legacy)
  - SR MPLS for inter-node connectivity
- NSO Integration
- ISS Console Integration
  - Dashboard and Health Monitoring
  - Service provisioning






# First 400G Transoceanic Links

## 400G Transatlantic Capacity Additions/Upgrades on Amiti  cable

- 1 x 400G for Internet2/CANARIE
- 2 x 400G for ESnet
- Expected by end October
- Add Boston as open exchange point
- Exploring options on other cable systems, including Pacific Ocean



# NA-REX Vision

- Create Domestic Consortium of Exchange Point and Link Operators
- Improve Coordination of Activities - GXP's may operate differently but we can all work together!
- Enable / optimize funding vehicles and grant opportunities for all parties
- Provide for efficient use of resources for transcontinental traffic, including:
  - Leverage common cores for routing production traffic to minimize operating costs
  - Provide dedicated links for experimental and meeting specific use (SC)
- Support research testbeds (e.g. FABRIC, BRIDGES)
- Encourage consistent set of operating principles and software features, such as:
  - NSI/AutoGOLE/SENSE
  - P4 and related instrumentation
  - Performance Assurance Services (PAS) test infrastructure (e.g. perfSONAR)
  - Measurement, monitoring, and reporting applications (e.g. NetSage, iGROK, Stardust)

# NA-REX Vision

**Automation and Software Expertise**  
**(Stardust, SENSE, etc)**  
**Transport Infrastructure**  
**Research Programs**  
**Technology Research**

**Midwestern Exchange Points**  
**Automation and Software Expertise (SDN, SDX, SDI...)**  
**Transport Infrastructure**  
**Research Programs**  
**Technology Research**  
**Strong Track Record of IRNC Awards**

**Pacific Coast Exchange Points**  
**Pacific Connectivity**  
**Transport Infrastructure**  
**Research Programs**  
**Automation and Software Expertise**  
**Technology Research**  
**Strong Track Record of IRNC Awards**

**East Coast Exchange Points**  
**Atlantic Connectivity**  
**Transport Infrastructure**

**East Coast Exchange Points**  
**Atlantic Connectivity**  
**Automation and Software Expertise**  
**Transport Infrastructure**

**Europe and Asia-Pacific Connectivity**  
**Research Programs**  
**Automation and Software Expertise**  
**Strong Track Record of IRNC Awards**

**Gulf Coast Exchange Points**  
**South America, Africa, and Caribbean Connectivity**  
**Research Programs**  
**Automation and Software Expertise**



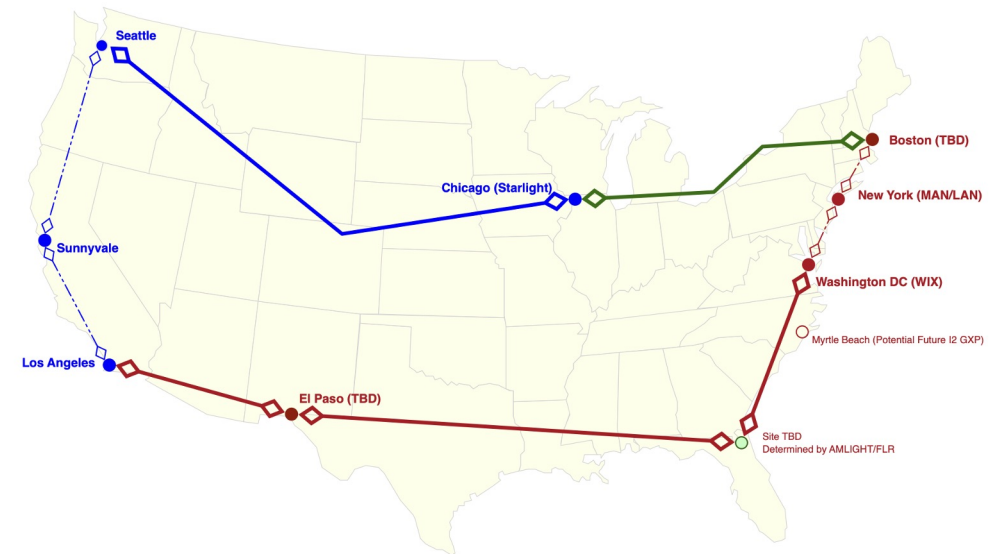
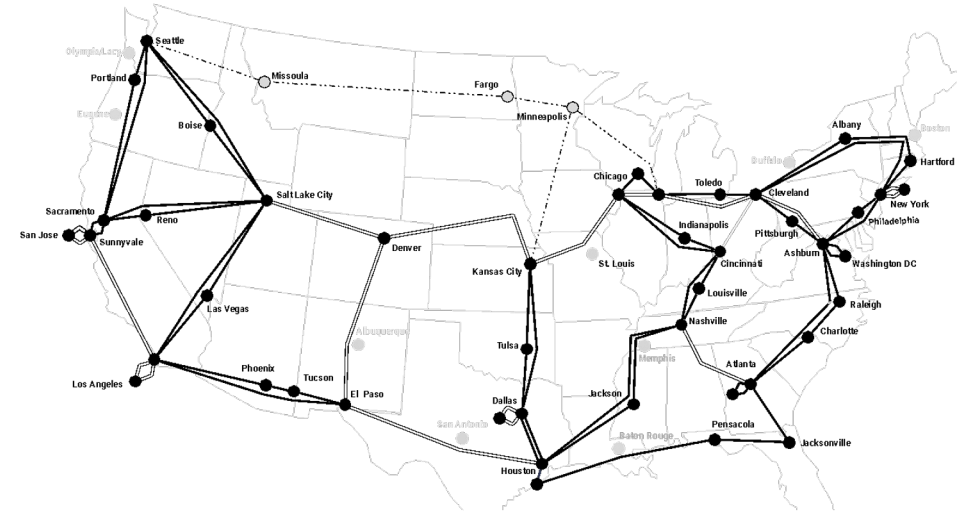
# Initial NA-REX Transport Architecture Concepts

## Leverage Internet2 NGL Core Network for Routine GXP Transit Traffic

- Routine traffic; commodity purposes
- Persistent production quality transport
- Path selection coming in future Cisco IOS-XR releases (along the lines of FLEX-ALGO)
- Potential support for multiple traffic classes to be developed (along the lines of BGP-CT, traffic-marking)

## Leverage Dedicated Waves for Programmatic Experimentation

- Example: Supercomputing, Data Mover Challenge
- Part of NA-REX Consortium Efforts, some potentially funded grants
- Persistent but movable
- 400 Gbps Native



# WHAT'S NEXT

Complete Deployment of New 400G Links and GXP's (Q4 of 2023)

Deliver Insight Console, APIs, and NSI-functionality to new Software Stack (Q4 of 2023; NSI Q1-Q2 2024)

Experiment with Expanded Core Network Functionality to Support Data-Intensive Use Cases (2023+2024)

Expand NA-REX Domestic US Consortium for Global Exchange Point Coordination  
MOUs, Roadmaps, Initial Work (2023)  
Expand Footprint, Add Features, Iterate (2024+)

Begin to merge improvements in core network technologies and software stack developments into general use (2024+)