



Internet2 Advanced Networking

LHCONE Meeting – April 18, 2023

Christian Todorov, Network Services, Internet2



Agenda

- Next Generation Networking & 400G
- Network Performance Assurance
- Network Automation and Orchestration
- Cloud Connectivity

New Faces

James Deaton

Vice President of Internet2's Network Services

- Lead both strategic direction and ongoing operations
- Manage all domestic and international service portfolios and relationships
 - Includes security, software, business development, operations, and network development
- Strong background in research engagement, software and automation
- Starting March 20

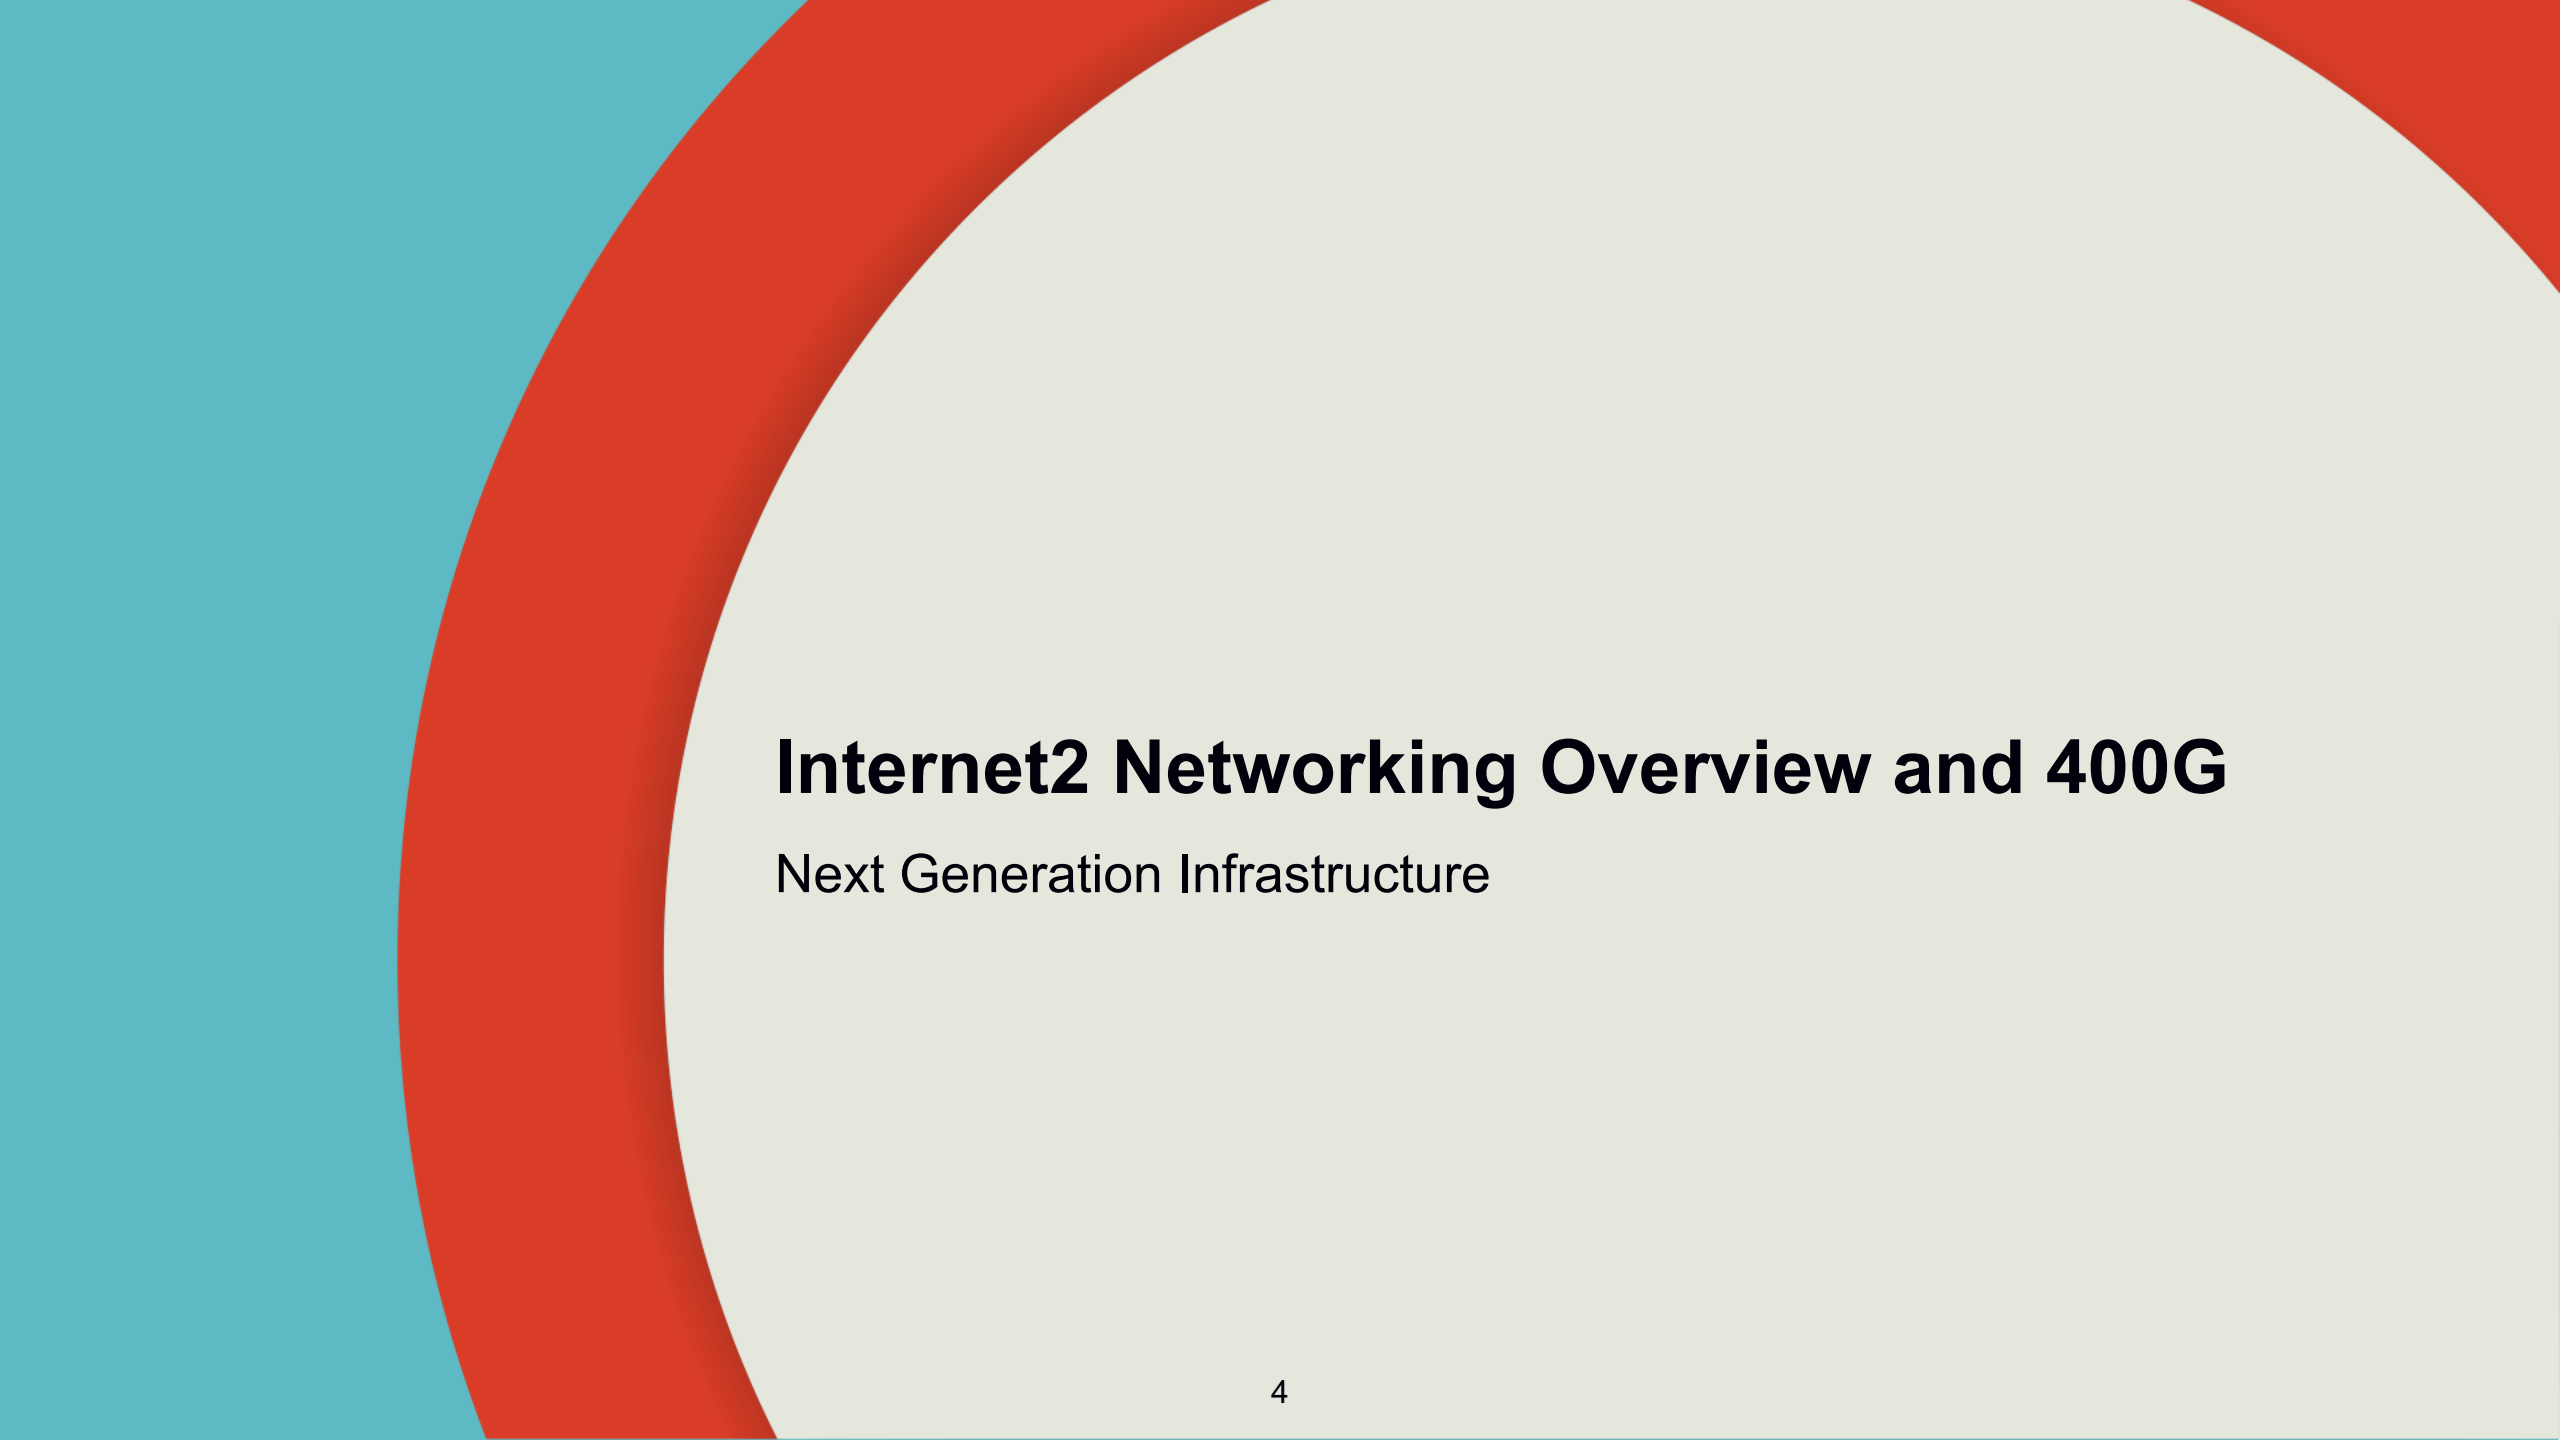


Thomas Araneta

International Engineering Program Manager

- Primary coordinator and point-of-contact for International Engagement
- Includes provisioning and service development
- Deep technical and community experience
- Started January 2nd





Internet2 Networking Overview and 400G

Next Generation Infrastructure

The NGI project has three phases:

- An upgrade of the underlying fiber, and the optical equipment to support 400G
- A complete replacement of the packet network and installation of a new Performance Assurance System (PAS)
- Implementation of Arroyo platform to support automation and orchestration of the new network and deliver new services to the community

NGI Principles

The Next Generation Infrastructure Program is a full set of activities to review and update the services, value and supporting technology of the Internet2 infrastructure portfolio (and relationships in the larger ecosystem)

- Includes the services and service models through which the community adopts Internet2 infrastructure services
- Includes new features, primarily driven by software, automation and systems virtualization to allow the infrastructure to be more readily integrated into the broader enterprise, campus, regional and cloud environment around us
- **Greenfield approach was selected to ensure optimal hygiene and best value of delivered system**



Data-Intensive Research

Enhanced Cloud Access

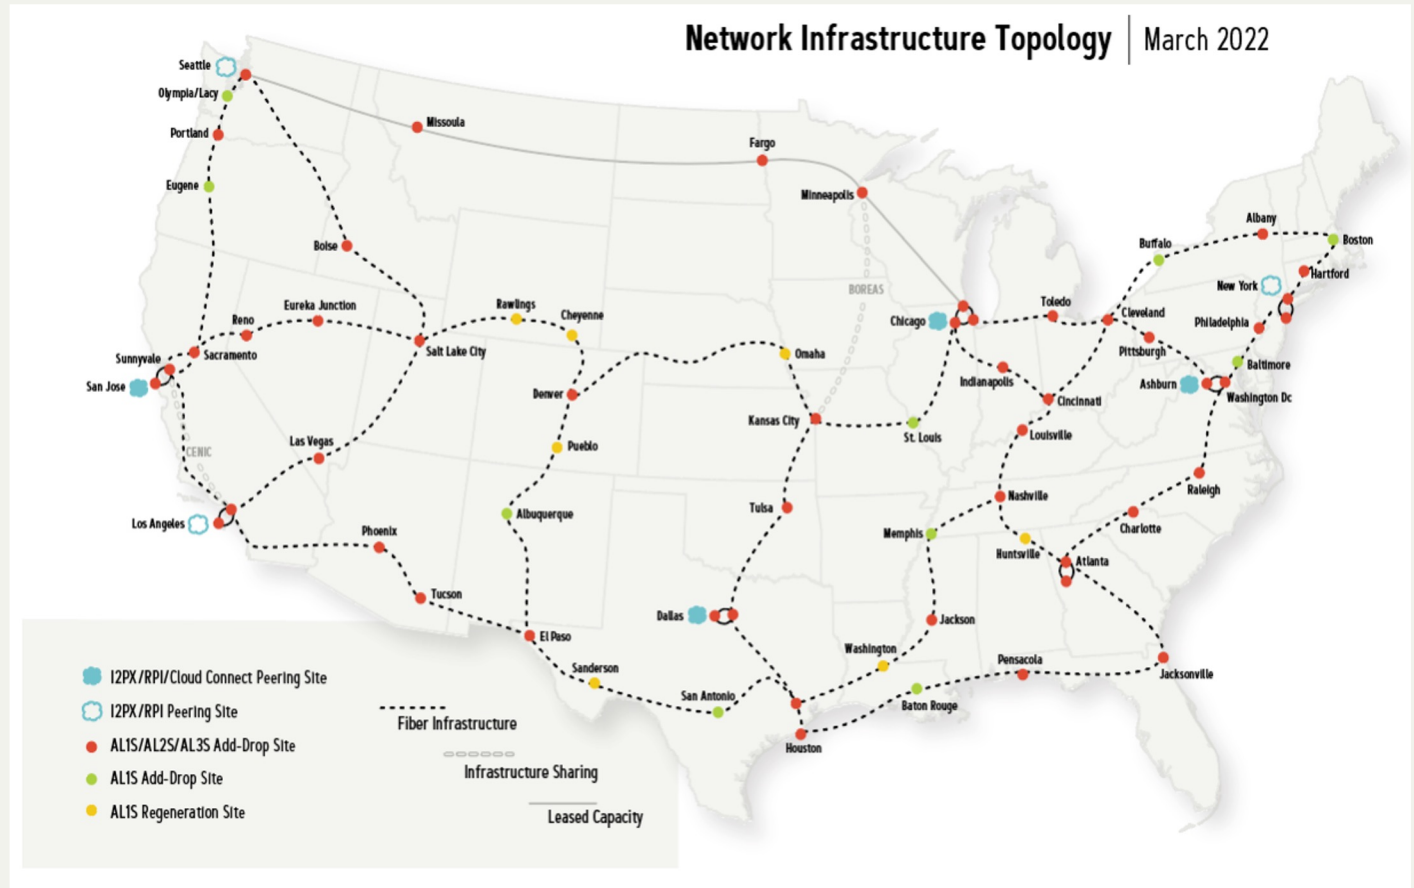
Software-Driven Infrastructure

Sustainable Economics

Infrastructure Sharing

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

GREENFIELD TECHNOLOGY BUILDOUT



Secure Management Network

Terminal Server
Firewall / Switch



NVIDIA

Performance Assurance Node

Dell R6515 Server
Mellanox ConnectX-5



NGI Packet Node

(2) Cisco 8201 – 24x400G & 12x100G



Experience. Outcomes.

NGI Optical Platform

Ciena WaveServer 5
32-slot C6500 Open Line System




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

400G Trans-Atlantic

400G Transatlantic Capacity Additions/Upgrades on Amitié cable

- 1 x 400G for Internet2/CANARIE
- 2 x 400G for ESnet
- **Mid 2023** (wet-plant in Q2 and terrestrial in Q3)
- Add Boston as open exchange point
- Early effort to acquire spectrum services from commercial providers – lessons learned
- Exploring second 400G link into New York, Washington or other east-coast city TBD



400G GXPs with Automation

Automating and Expanding Exchange Point Functionality; Match Community Development

- Hardware improvements to support 400G, including MAN/LAN and WIX
 - Arista DCS-7280PR3K-24
- 400 Gbps links between all three exchange point switches
- Protocols – bring up to NGI standards, modernize
 - EVPN MPLS for L2 services (VXLAN Legacy)
 - SR MPLS for inter-node connectivity
- NSO Integration
 - Improved consistency
 - Improved manageability
 - Build service models supporting services on GXPs
 - Integrate NSO models with Internet2's core network models
- ISS Console Integration
 - Dashboard and Health Monitoring
 - L2 and L3 service provisioning
 - End to end service monitoring within Console
 - Intend to expand support for pushing data to partners



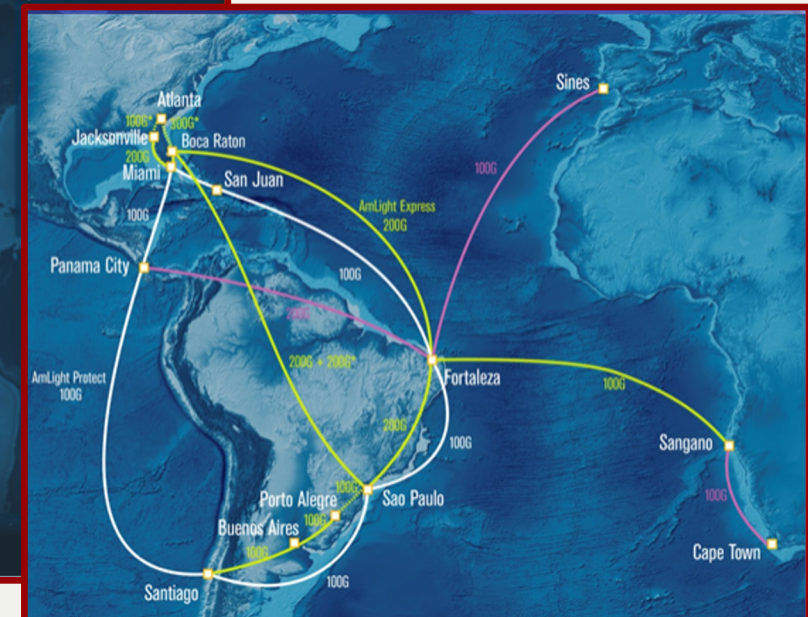
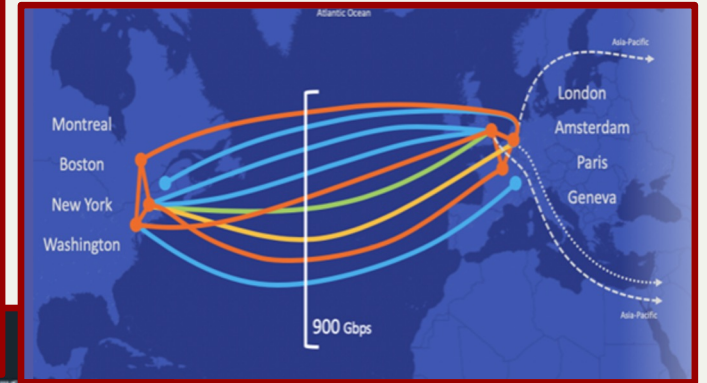
Facilitating and Partnering to Expand and Enhance Global Connectivity

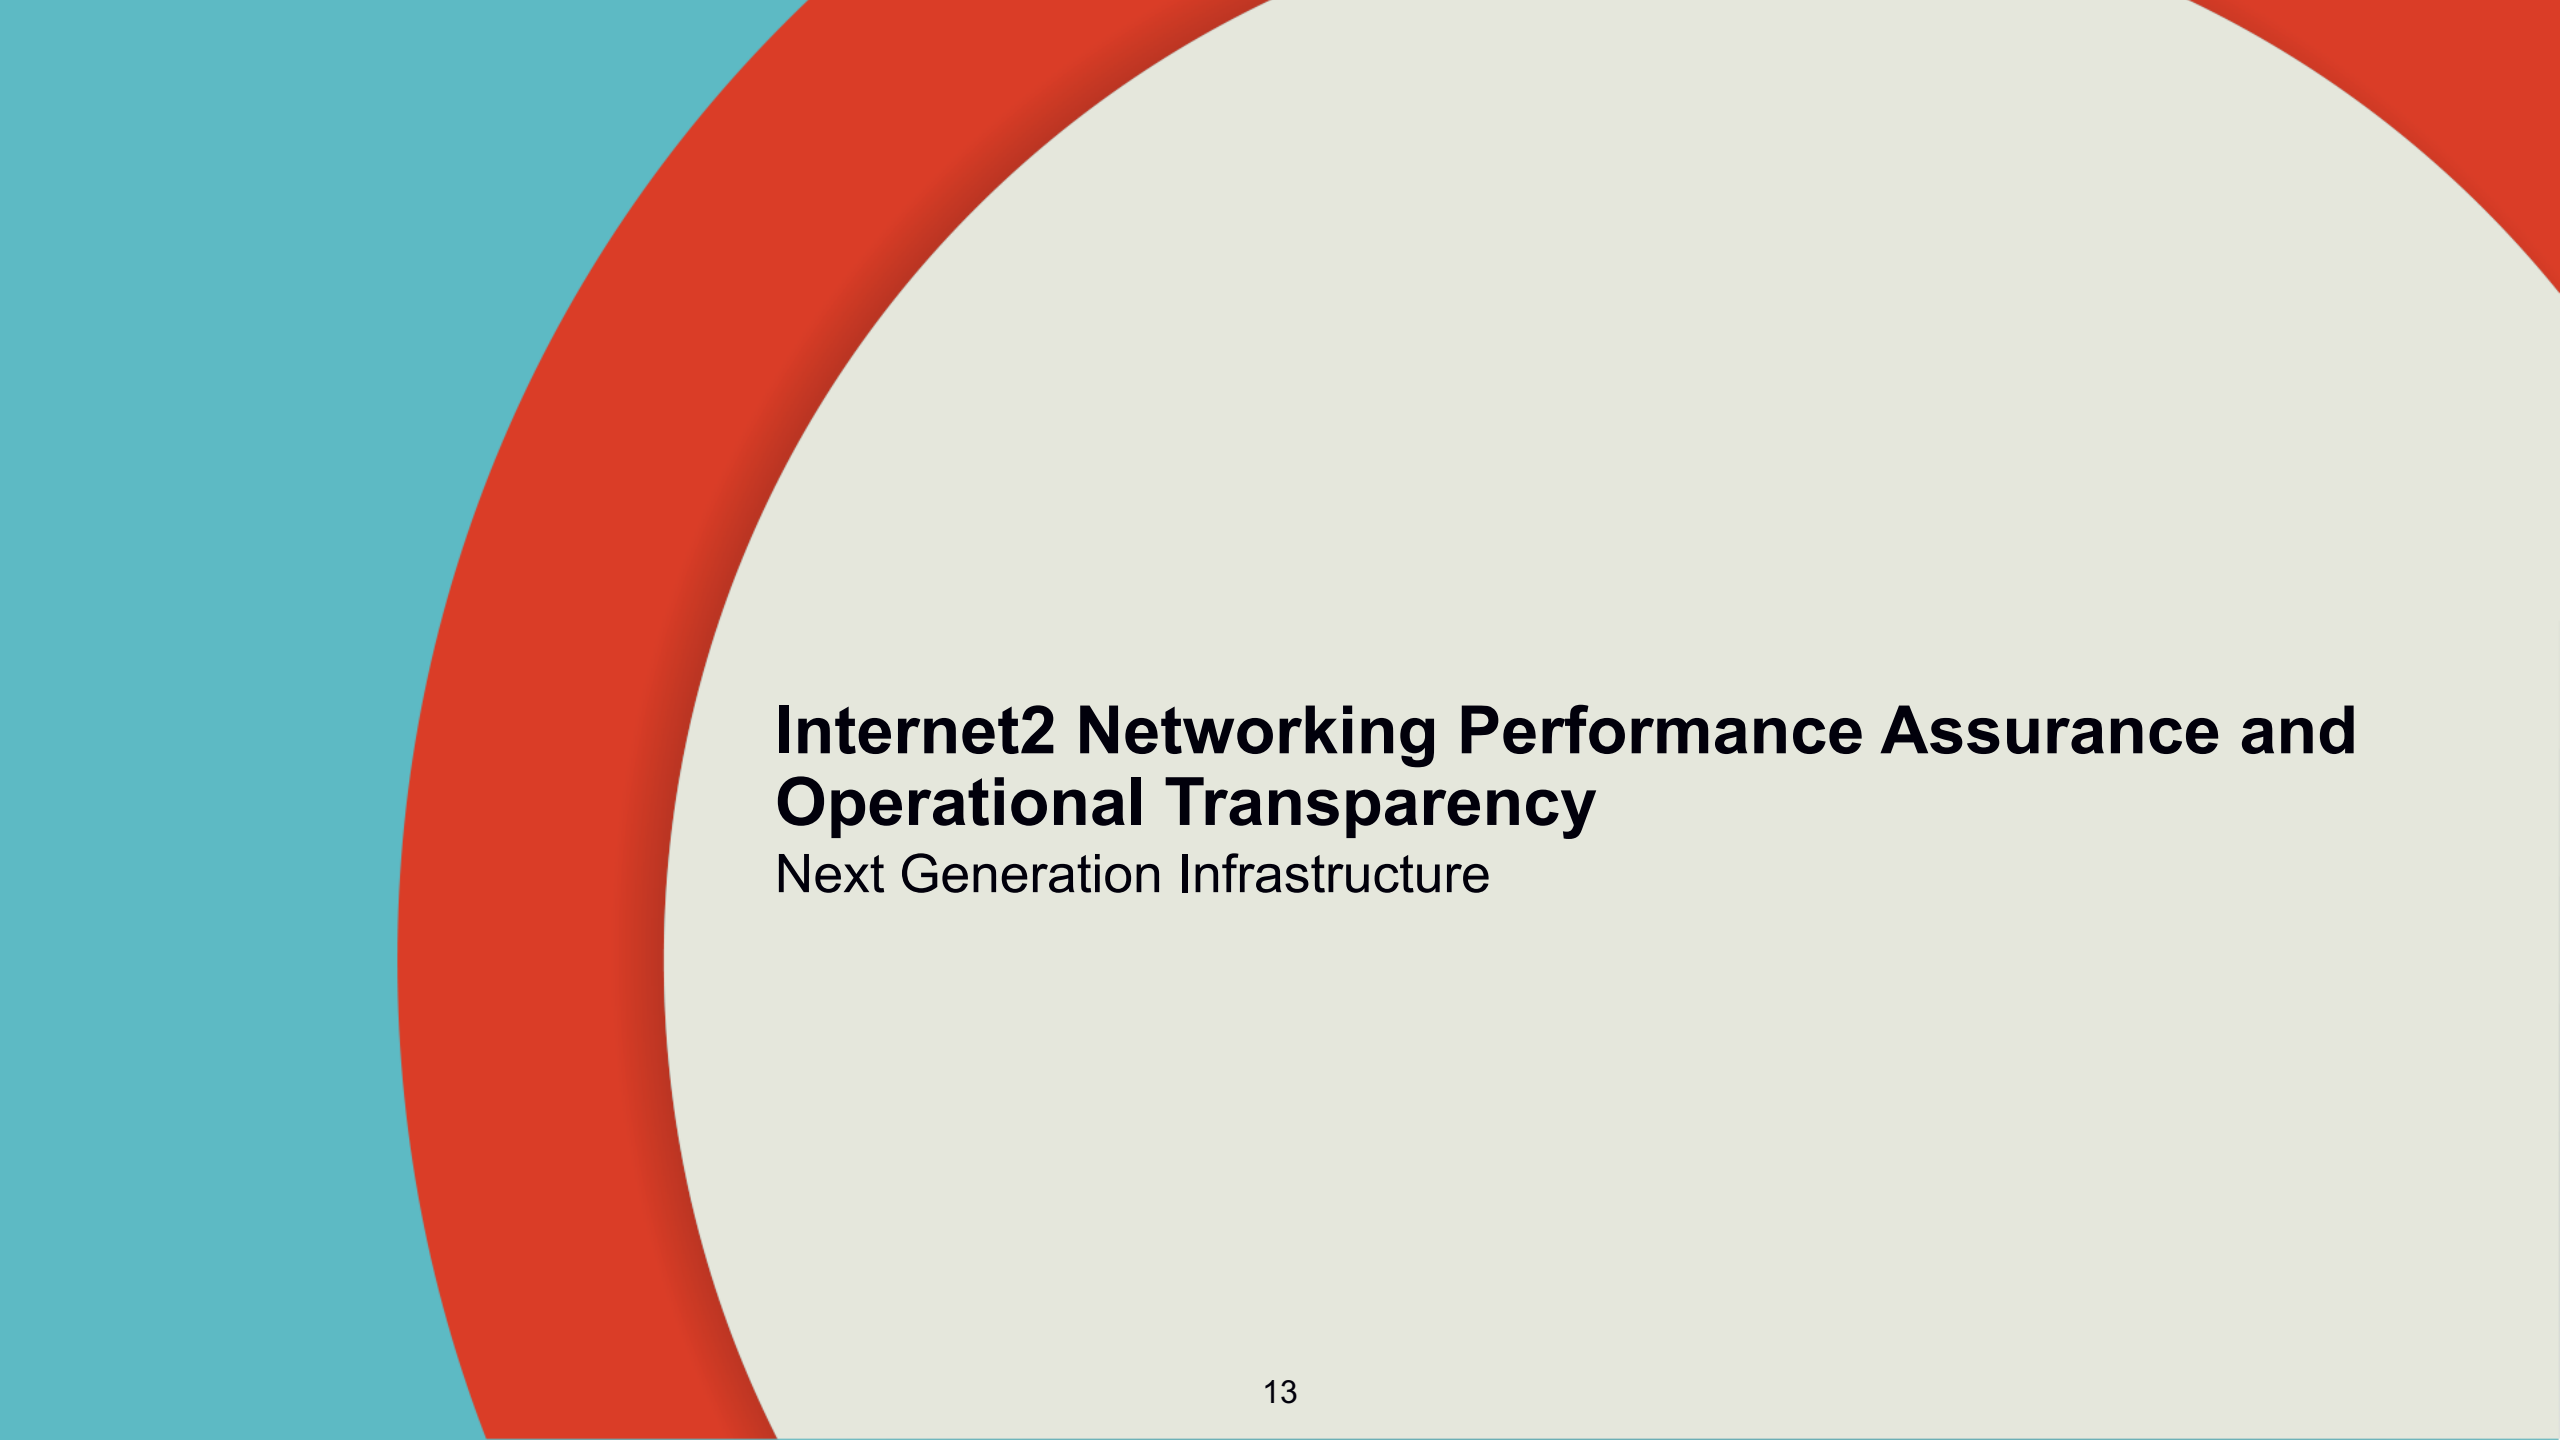
Provide support to community for adding 400G diversity / triversity in all regions:

- Logistical
- Contractual
- Financial
- Staff Time
- Relationships with vendors

In tandem with...

- APOnet
- ANA
- FIU
- et al





Internet2 Networking Performance Assurance and Operational Transparency

Next Generation Infrastructure

The NGI Project Second Phase:

- ✓ • An upgrade of the underlying fiber, and the optical equipment to support 400G
- A complete replacement of the packet network and installation of a new Performance Assurance System (PAS)
- Implementation of Arroyo platform to support automation and orchestration of the new network and deliver new services to the community

Operational Transparency

Mission

Provide security-preserving operational transparency to the network operator, user and application developer communities to aid in application development and troubleshooting, and for network visibility and validation.

Moving Network Operations into the Light

- PAS Nodes:
 - Published, persistent link-based network health and performance tests, e.g., jitter, loss, latency, throughput, etc.
 - Policy-managed, user-controlled ad hoc tests across the backbone up to 100G
- Looking Glass:
 - Policy-restricted, Role Based Access Control and abstracted access to, and restricted output from Internet2 routers for curated “show” commands, enabling enhanced network visibility to network operators, researchers and application developers
- Published Maintenance Reporting and Live Status

New Internet2 Performance Assurance Service

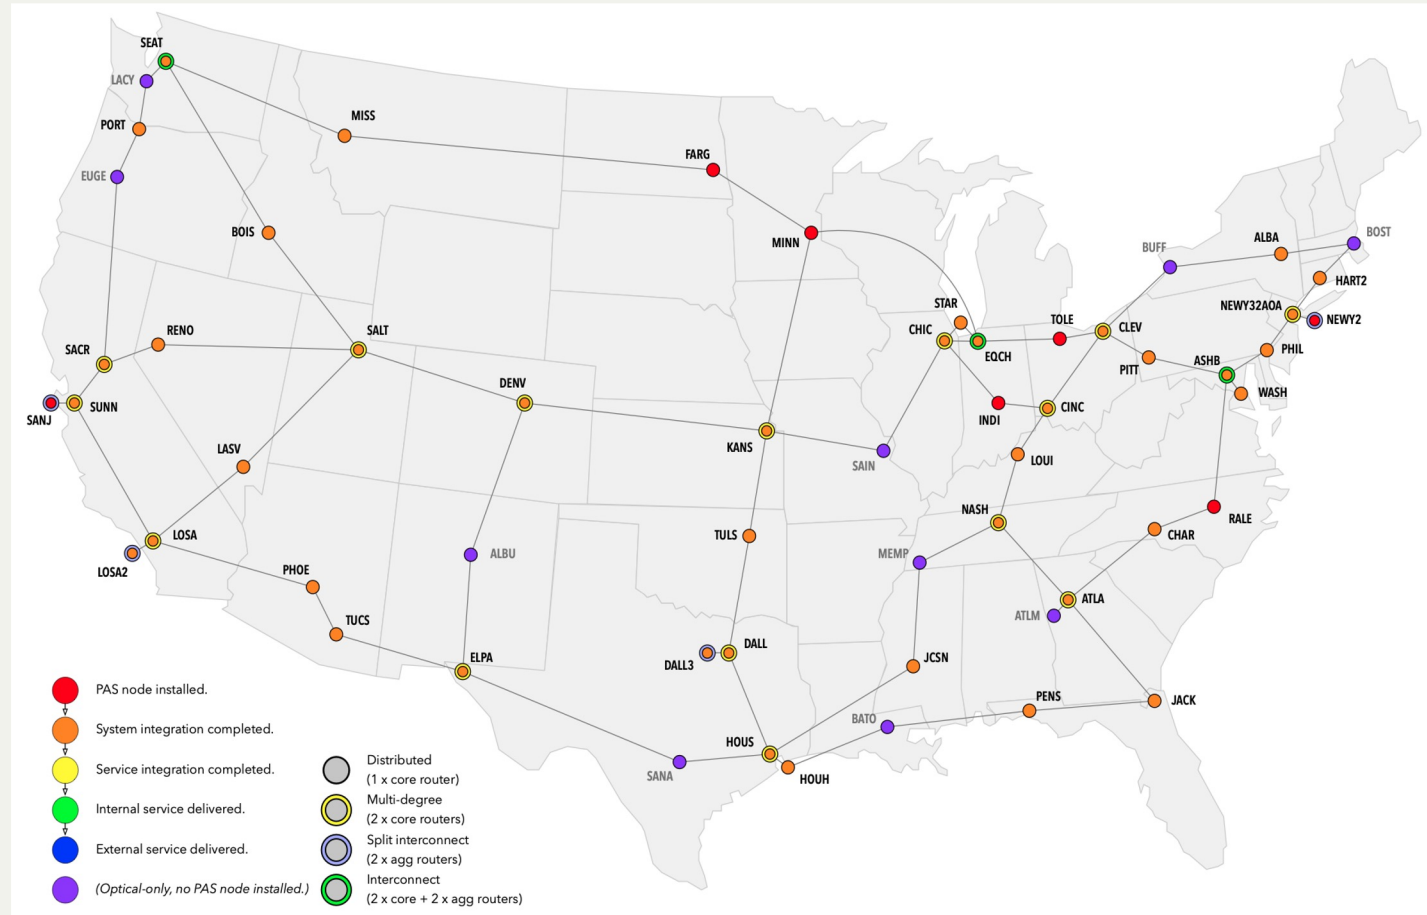
47 Node 2x100G dedicated perfSONAR nodes

- 43 nodes installed and reachable
- 3 nodes pending installation

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

Initial deployment projected to be complete by end of year

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





Internet2 Networking Software Architecture

Next Generation Infrastructure

The NGI project moves to its third phase:

- An upgrade of the underlying fiber, and the optical equipment to support 400G
- A complete replacement of the packet network and installation of a new Performance Assurance System (PAS)
- Implementation of Arroyo platform to support automation and orchestration of the new network and deliver new services to the community



- Early efforts will focus on:

- ✓ • Deployment of NSO for network service modeling, automation, orchestration, and management
- ✓ • Development of NS Insight Console
- Incorporating the ability to manage people and their roles across our ecosystem using IAM resources

Internet2 Insight Console - Delivering the next generation of network-based services for the R&E Community

- Leverages the automation and orchestration capabilities of NGI
- Delivers on requirements from enterprise and researchers for ecosystem-wide solutions
- Incorporates the ability to manage people and things across our ecosystem through IAM, so important in a software-based world
- Addresses need to improve routing integrity within the R&E community

Three Phases of Insight Console Development

- Implementation of existing services
 - Looking Glass
 - Routing Portal
 - Cloud Connect Portal/Virtual Network Portal
 - Introduction of Role Based Access (IAM)
- Community input to Insight Console roadmap during Fall 2022 for implementation by end of 2023
 - e.g. What tools and services will be needed by the Members to meet their needs?
- Community input to a sustainability model for the Insight Console by 2024
 - e.g. What tools and services will be needed by the campuses & others to meet their needs ?

AUTOMATION as APPLIED TO NGI

NSO service models heavily leveraged for service deployments

- Minimal, simple parameters needed to produce complex device configuration
- All changes are atomic across network

Layer 3 (BGP) service migrations successful with NSO

- Imported legacy network/DB data to produce NSO-generated service configurations
- Allowed migration of 1800 peerings in ~30 days (up to 150 / night)

Validation tooling under NSO determined success (quality assurance)

- BGP prefix acceptance/rejection monitoring before and after migrations
- Allows rapid validation & success of migrated services

Test Driven Development

- Rapid iteration of NSO service models with testing to deter regression
- Decrease time needed for new service deployments

Configuration Control / Monitoring

- Manual changes on devices are identified and flagged for reintegration into NSO models
- Ensures minimal drift

Time savings is not the entire picture


Consistency

Reduce engineer fatigue

Queue up changes in advance

Reliability (no fat fingering)

Tooling also provided validation input



I2 INSIGHT CONSOLE AND API DEVELOPMENT

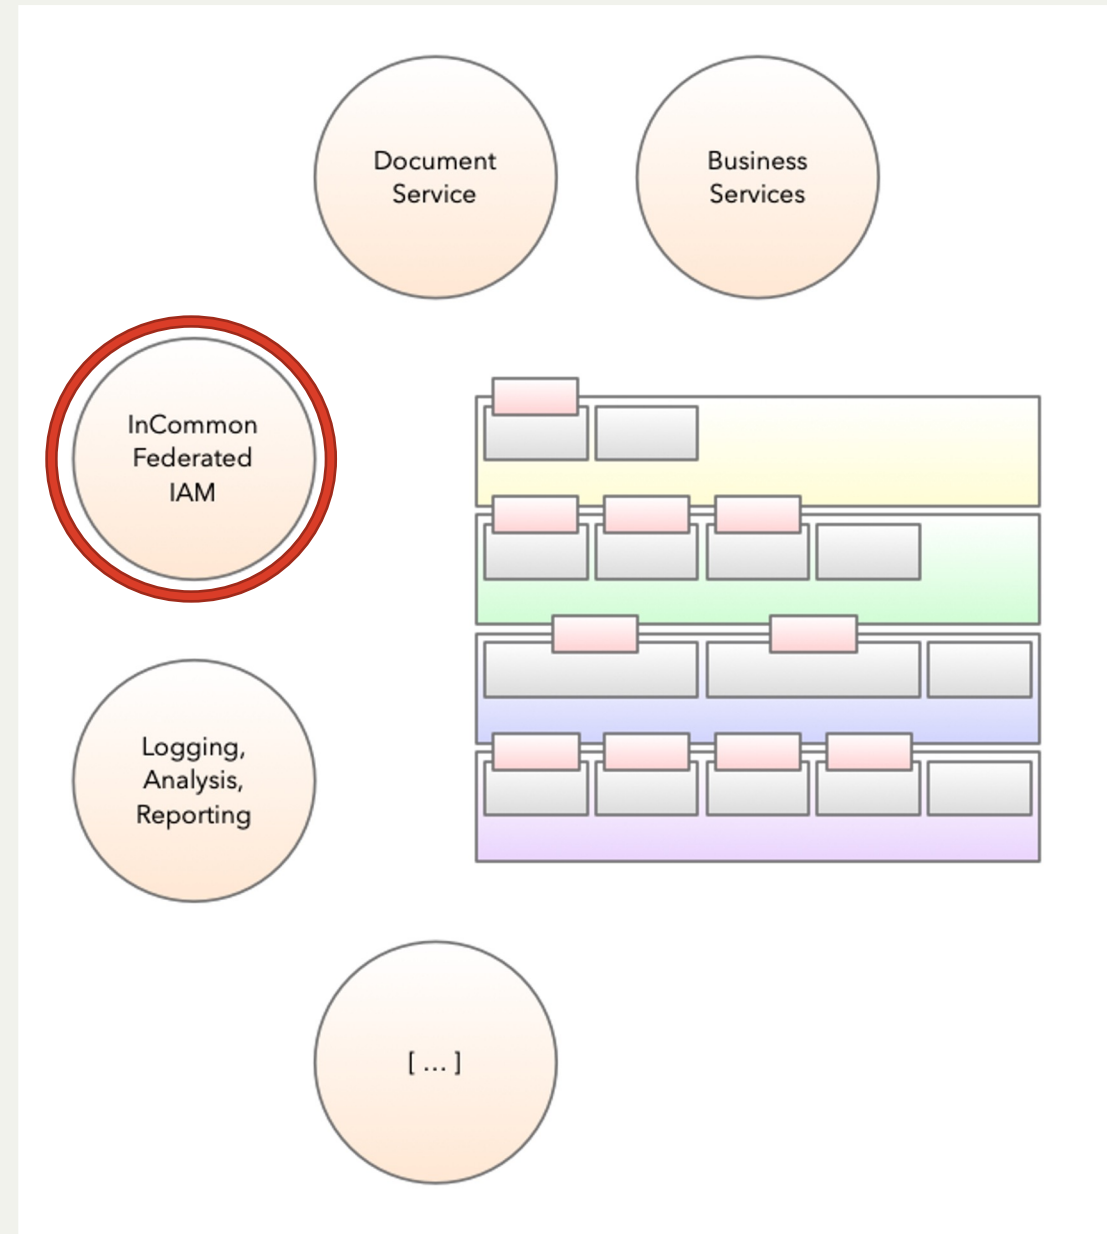
The screenshot displays the I2 Network Looking Glass web interface. On the left, there is a navigation menu with various network management tools. The main area is divided into a device selection pane and a terminal window. The terminal window shows the output of a 'show interfaces' command, displaying details for several loopback interfaces (Loopback0, Loopback1, Loopback8, Loopback11) on the device 'agg1.eqch.aa'. The output includes interface state, hardware information, and IP addresses.

“Internet2 Insight”

- Access to many (or all) of our software-controlled services.
- Single unified front-end which may access external applications and APIs as needed.
- All console functions use an underlying API layer, which may be exposed in the future through other tools or directly to members who want to integrate their own tools.
- Will eventually encompass international provisioning.

Arroyo External Integrations

- **InCommon Federated IAM:**
 - Fundamental to our ability to support role-based end to end access to services
 - Foundational to determining who is authorized to access the services that will be developed by the community
 - A vital component in supporting security needs such as Zero Trust Network environments



The InCommon integration:

- Technically, integration with the **Internet2 Identity Services** platform.
- Authentication:
 - Via **home identity provider**, using familiar SSO credentials, for any eduGAIN/InCommon-affiliated organization.
 - Alternatively, via an **Internet2 Guest Account**, if you don't currently have a home identity provider.
 - Note:
 - If you're interested in getting into the business of running an identity provider, Internet2 has identified several **InCommon Catalysts** to help you get started.
- Identity:
 - Authentication links you to your **Internet2 identity**.
 - msimpson @ Internet2 -> mike.simpson@at.internet2.edu
 - aadams @ UA -> alice.adams@at.internet2.edu
- Affiliation and attribution:
 - Identity is associated with **organizational affiliations and roles**:
 - "Alice Adams is a network engineer at the University of Arizona."
- Policy and authorization:
 - Affiliations and roles intersect with **policy configuration** to determine **authorizations** for specific functions in the API and tools:
 - "Network engineers can view BGP peering session information associated with their organization."

INTERNET² Identity Services

Select an Identity Provider

Please select a trusted identity provider from the list below by beginning to type in the edit box.

If you do not find your organization in this list, there are two options:

1. You can [login with your Google account](#).
2. You can [login with your Internet2 Guest account](#). If you are new, you may [create an Internet2 Guest account](#) (see [Guest Account Guide](#)).

Enter your organization's name

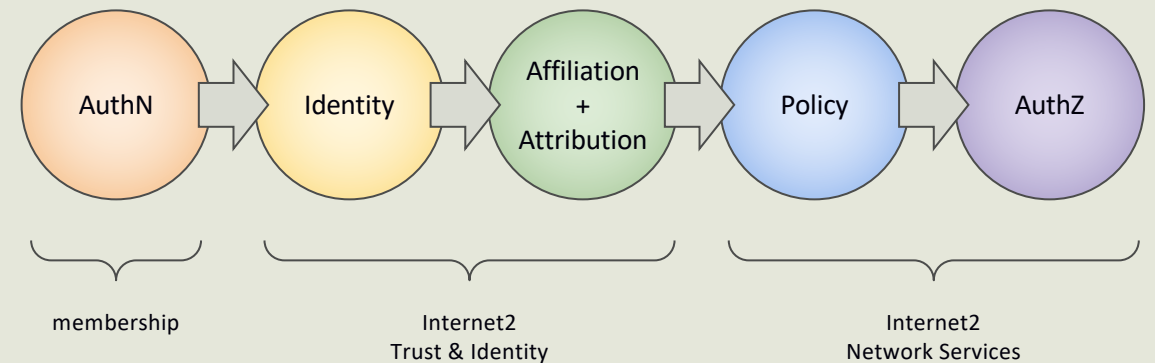
Enter the name of the organization you are affiliated with...

Continue

Remember selection for this web browser session.
 Remember selection permanently and bypass this step from now on.

[Help](#)

[I'm having trouble logging in](#)



The intent:

new relic

Search
+ Add data
All capabilities
All entities
APM & services
Query your data
Apps
Browser
Dashboards
Alerts & AI
Errors inbox
Hosts
Infrastructure
Logs
Mobile
Synthetic monitoring

Logs
lookingglass
1,242 logs found
Expand logs
timestamp

GlobalNOC

Filter navigator
Internet2 Lightning
Accounts
Network Connectors
37 Items - Sorted by Account Name - Filtered by All accounts - Network R...

Account Name ↑

- 3ROX (Three Rivers Optical Exchange)
- Big Ten Academic Alliance OmniPoP
- CAAREN (Capital Area Advanced Research and Education Network)
- CEN (Connecticut Education Network)
- CENIC (Corporation for Education Network Initiatives in California)
- Drexel University
- FLR (Florida LambdaRail, LLC)
- GPN (Great Plains Network)
- Indiana GigaPoP
- KINBER (Keystone Initiative for Network Based Education & Research)
- KyRON (Kentucky Regional Optical Network)
- LEARN (Lonestar Education and Research Network)
- Link Oregon
- Louisiana Board of Regents / LONI (Louisiana Optical Network)
- MAGPI (Mid-Atlantic GigaPoP)
- MARIA (Mid-Atlantic Research Infrastructure Alliance)
- MAX (Mid-Atlantic Crossroads)
- MCNC
- Merit Network, Inc.
- MISSION (Mississippi Optical Network)
- MREN (Metropolitan Research & Education Network)
- Nevada System of Higher Education - System Office
- NJEdge
- Northern Lights GigaPoP
- NoX (Northern Crossroads)
- NYSERNet, Inc.
- OARnet (Ohio Academic Resources Network)
- OSHEAN, Inc.

INTERNET2 Identity Services

Select an Identity Provider

Please select a trusted identity provider from the list below by beginning to type in the edit box.

If you do not find your organization in this list, there are two options:

- You can [login with your Google account](#).
- You can [login with your Internet2 Guest account](#). If you are new, you may [create an Internet2 Guest account](#) (see [Guest Account Guide](#)).

Enter your organization's name

Enter the name of the organization you are affiliated with...

Remember selection for this web browser session.
 Remember selection permanently and bypass this step from now on.

[Help](#)

[I'm having trouble logging in](#)

Requestor name	Req	Organization	Routing Table	ASN	Status	Created
Chris Wopat	BT	National Oceanic & Atmospheric Administr...	desired routing table L...	RITM0084257	Approved	2022-21:45
carl.lindahl	MC	University of Washington	69.91.192.0/18 The desired routing table...	RITM0084103	Approved	2022-23:04
Brian Mollo	No	University Corporation for Atmospheric R...	164.92.0.0/18 The desired routing table ...	RITM0084054	Approved	2022-13:06
Jeff Harrington	NY	Minnesota Statewide Education Group Network	2606:e540::/32 The desired routing table...	RITM0083954	Approved	2022-19:16
Jeff Harrington	NY	Indiana GigaPop	64.189.100.0/24 The desired routing tabl...	RITM0083674	Approved	2022-17:40
Pete Siemsen	BT	NYSERNet	192.30.126.0/23 The desired routing tabl...	RITM0083560	Approved	2022-19:11
John Harrington	NY	NYSERNet	192.30.126.0/23 The desired routing tabl...	RITM0083561	Approved	2022-19:11
David Farmer	NY	National Oceanic & Atmospheric Administr...	2601:a000::/20 The desired routing table...	RITM0082956	Approved	2022-01:26
Sean Carver	NY	Education Networks of	2610:128:4200::/40 The desired routing...	RITM0082835	Approved	2022-
Kyle Emmons	NY	Front Range GigaPop				

California State University

ASN 65000

PDP SANJ-CONN-MICROSOFT-2 | 12-SANJ-SANJ-100GE-192369

Entity CrossConnection-SiliconValley

California State University endpoint for Microsoft Azure

Router config

Endpoint description

172.20.252.113/30

IPv4 Address

65000

ASN

BGP Key

Delete Peering



AP-REX 2.0

Next Generation Research Supporting
Infrastructure

Atlantic-Pacific Research Exchange (AP-REX) 2.0 Objectives

- 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)

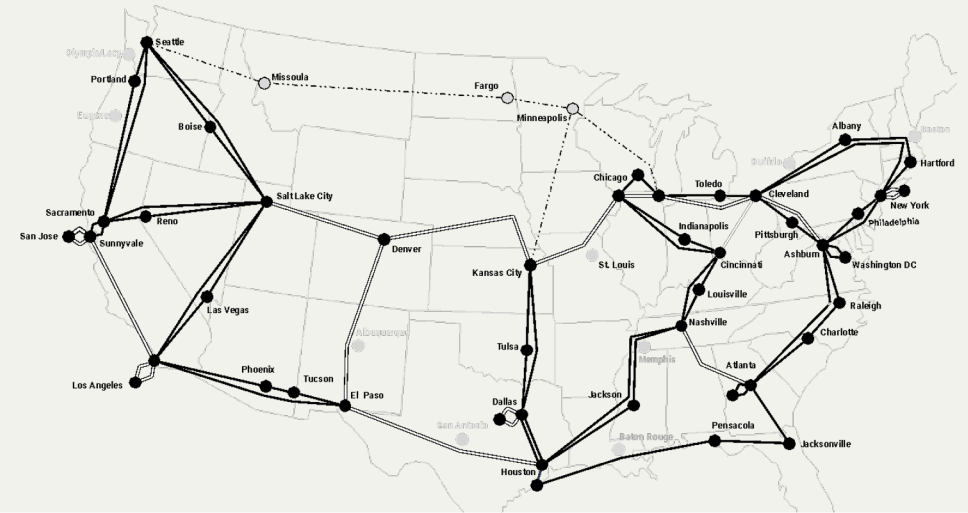
AP-REX 2.0 Concept Infrastructure



AP-REX 2.0 – Architectural & Transport Concepts

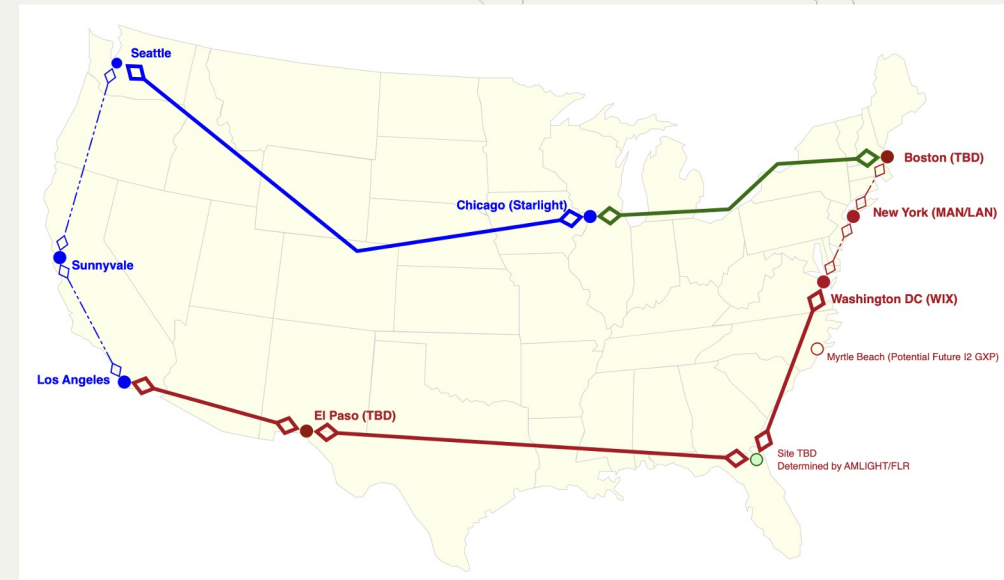
Leverage Internet2 NGI Core Network for Routine GXP Transit Traffic

- 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 AP-REX Consortium Efforts, some potentially funded grants
- Persistent but movable
- *Need to test to determine if there is actual benefit to dedicated waves!*






Internet2 Networking Cloud Connectivity

Next Generation Infrastructure

Getting to the Cloud



Internet2 Peer Exchange I2PX

Use of existing 3Tbps of peering capabilities to the major cloud providers for access to cloud SaaS services (e.g., Zoom or Office 365) – **shared** service

Internet2 Cloud Connect I2CC

Enables connectors to use the Internet2 infrastructure to obtain up to 5Gbps of “direct-connect” **private** Layer 2 and Layer 3 access to Microsoft, Amazon and Google cloud platforms **at no additional fee**. Extending your data center to the cloud. (Cloud provider fees apply)

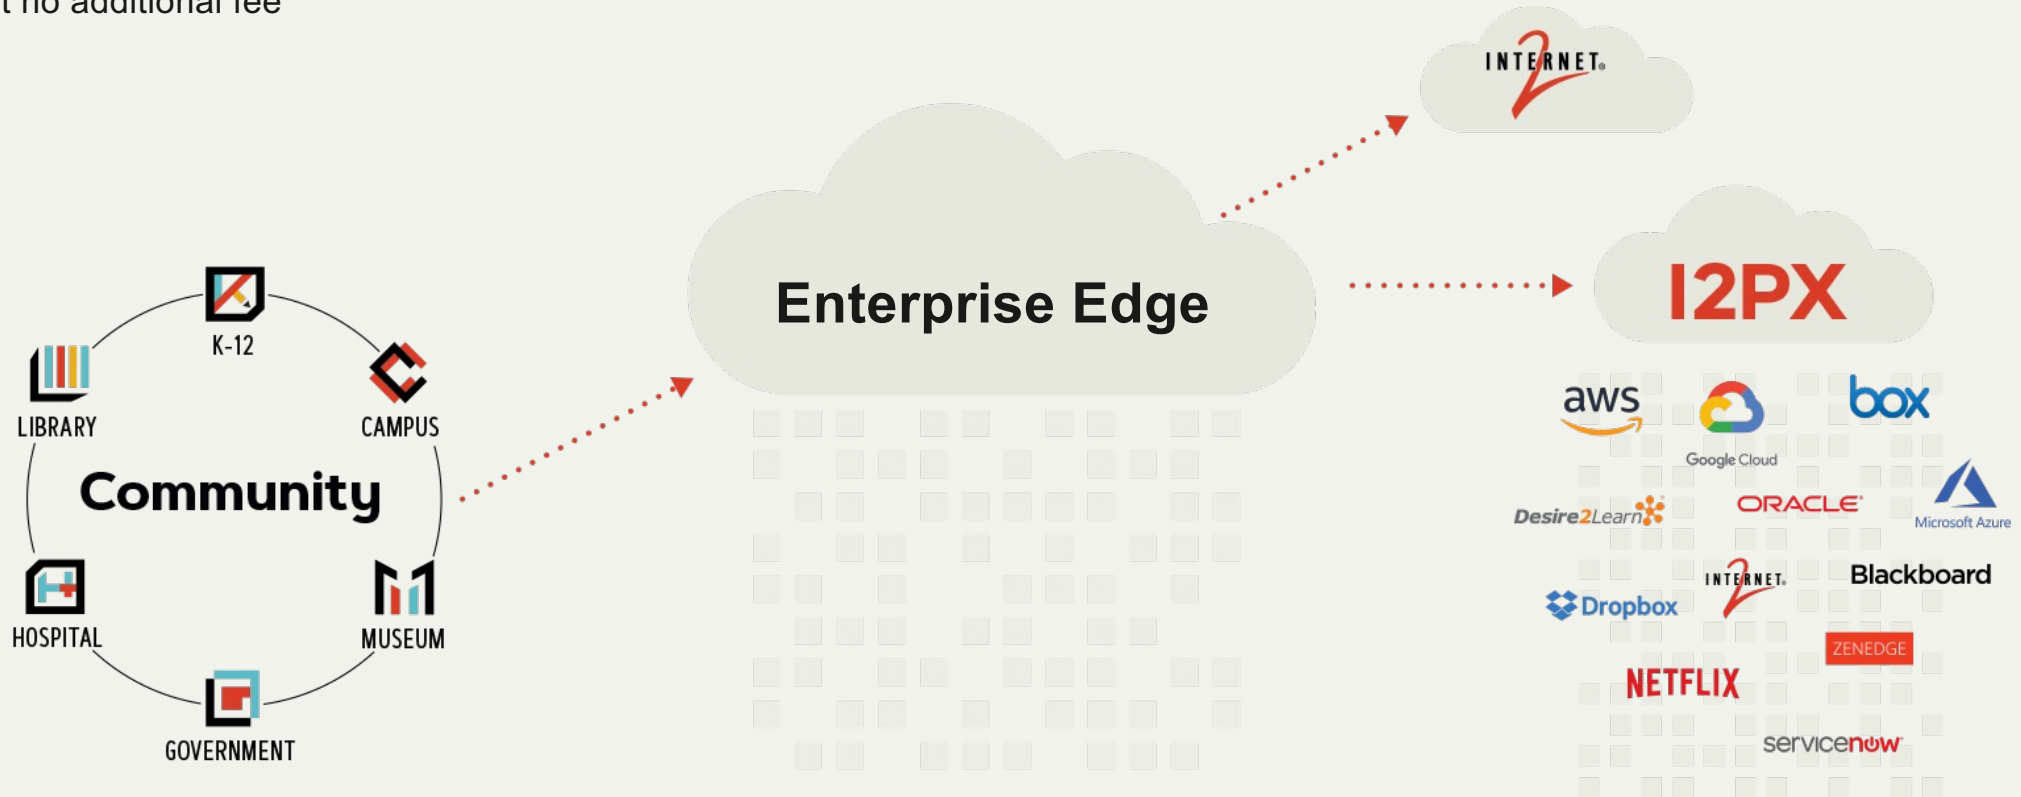
Internet2 Rapid Private Interconnect I2RPI

Provides **private dedicated** 10G interconnections at major peering points across the country at low annual rates. Leverages current Internet2 infrastructure to reach cloud providers, for dedicated access or improved resiliency. May be used to connect to any commercial provider located at the peering point

Designed to Deliver

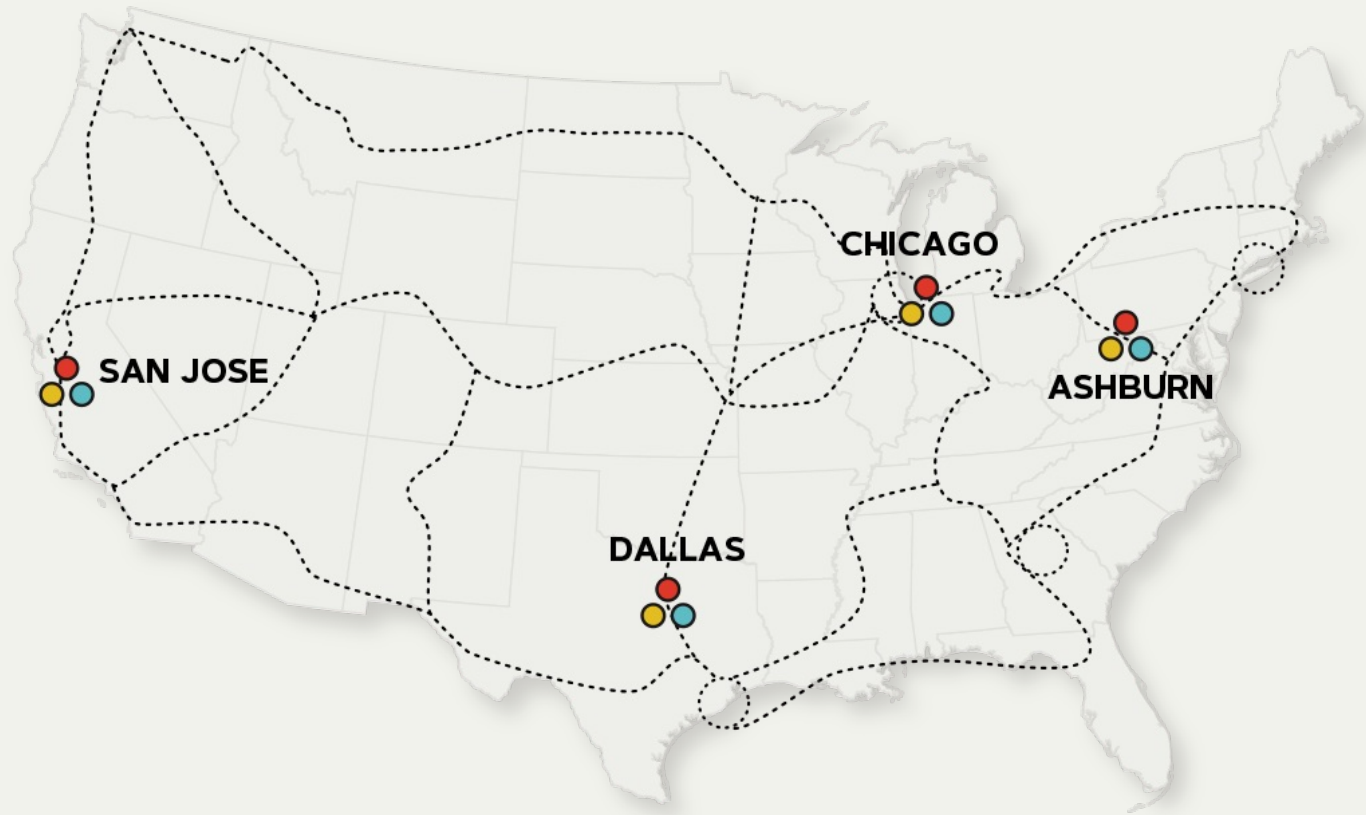
Internet2 Peer Exchange **I2PX**

- Allows a connector to have high performing on-net shared access to cloud service providers, avoiding the commodity internet
- Designed from the ground up to focus on hosting cloud providers most valued by the Research and Education community
- Available to at no additional fee



Nationwide Connectivity

Internet2 Cloud Connect I2CC



- **Microsoft ExpressRoute**
 - 2x1x100G in Ashburn + GovCloud
 - 2x1x10G in Chicago
 - 2x1x10G in Dallas
 - 2x1x100G in San Jose + GovCloud
- **Amazon Direct Connect**
 - 4x1x10G in Ashburn
 - 2x1x10G in Chicago
 - 2x1x10G in Dallas
 - 2x1x10G in San Jose
- **Google Cloud Interconnect**
 - 2x1x10G in Ashburn
 - 2x1x10G in Chicago
 - 2x1x10G in Dallas
 - 2x1x10G in San Jose

Nationwide Connectivity

Internet2 Rapid Private Interconnect **I2RPI**

RPI can be used to provide private dedicated direct connects to any provider with some examples being 10G connections to SIP service providers, esports exchanges, or Oracle Cloud.



Internet2 Networking - Tomorrow

The Future of Internet2 Next Generation
Infrastructure

Next Steps

Complete Deployment of New 400G Links and GXPs (Q2 and Q3 of 2023)

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

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

Expand AP-REX Domestic US Consortium for Global Exchange Point Coordination

- MOUs, Roadmaps, Initial Work (2023)

- Expand Footprint, Add Features, Iterate (2024+)

- Explore controller-less networking for advanced features

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



Thank you!

Questions:

Christian Todorov ctodorov@internet2.edu