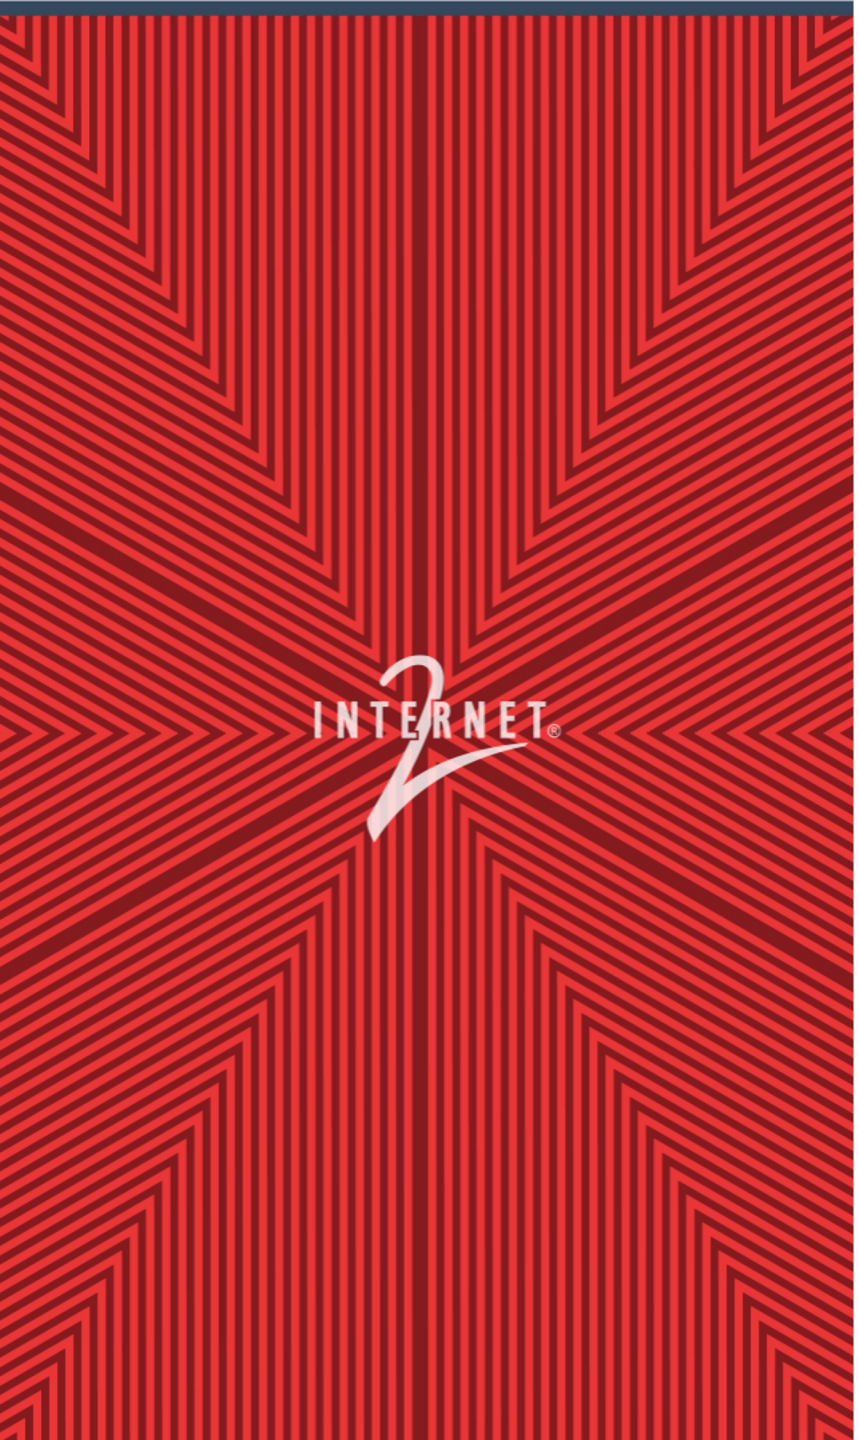

Internet2 Update

11th SIG-NGN: Large science support & network futures
Prague, Czechia
20 April 2023

Chris Wilkinson
Sr Director for Planning, Architecture, and Engineering
Network Services

Common Themes from Research Community, NRENs

Coordinate across worldwide networks to ensure 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 Targets Supporting Science

- “Traditional” Bulk Data, Scalability → 400G+ Infrastructure, Exchange Points
- Programmability / Topology Visibility → APIs / Portal
- Cloud-Access → High Scale at Edge
- Performance Assurance → Testing Resources, Access to Platform
- Programmatic Reporting (IRNC) → Potential Netsage Pipeline
- Data Lakes → Open Science Data Federation (OSDF) Caches On-Net

Network Topology / Service Delivery Points

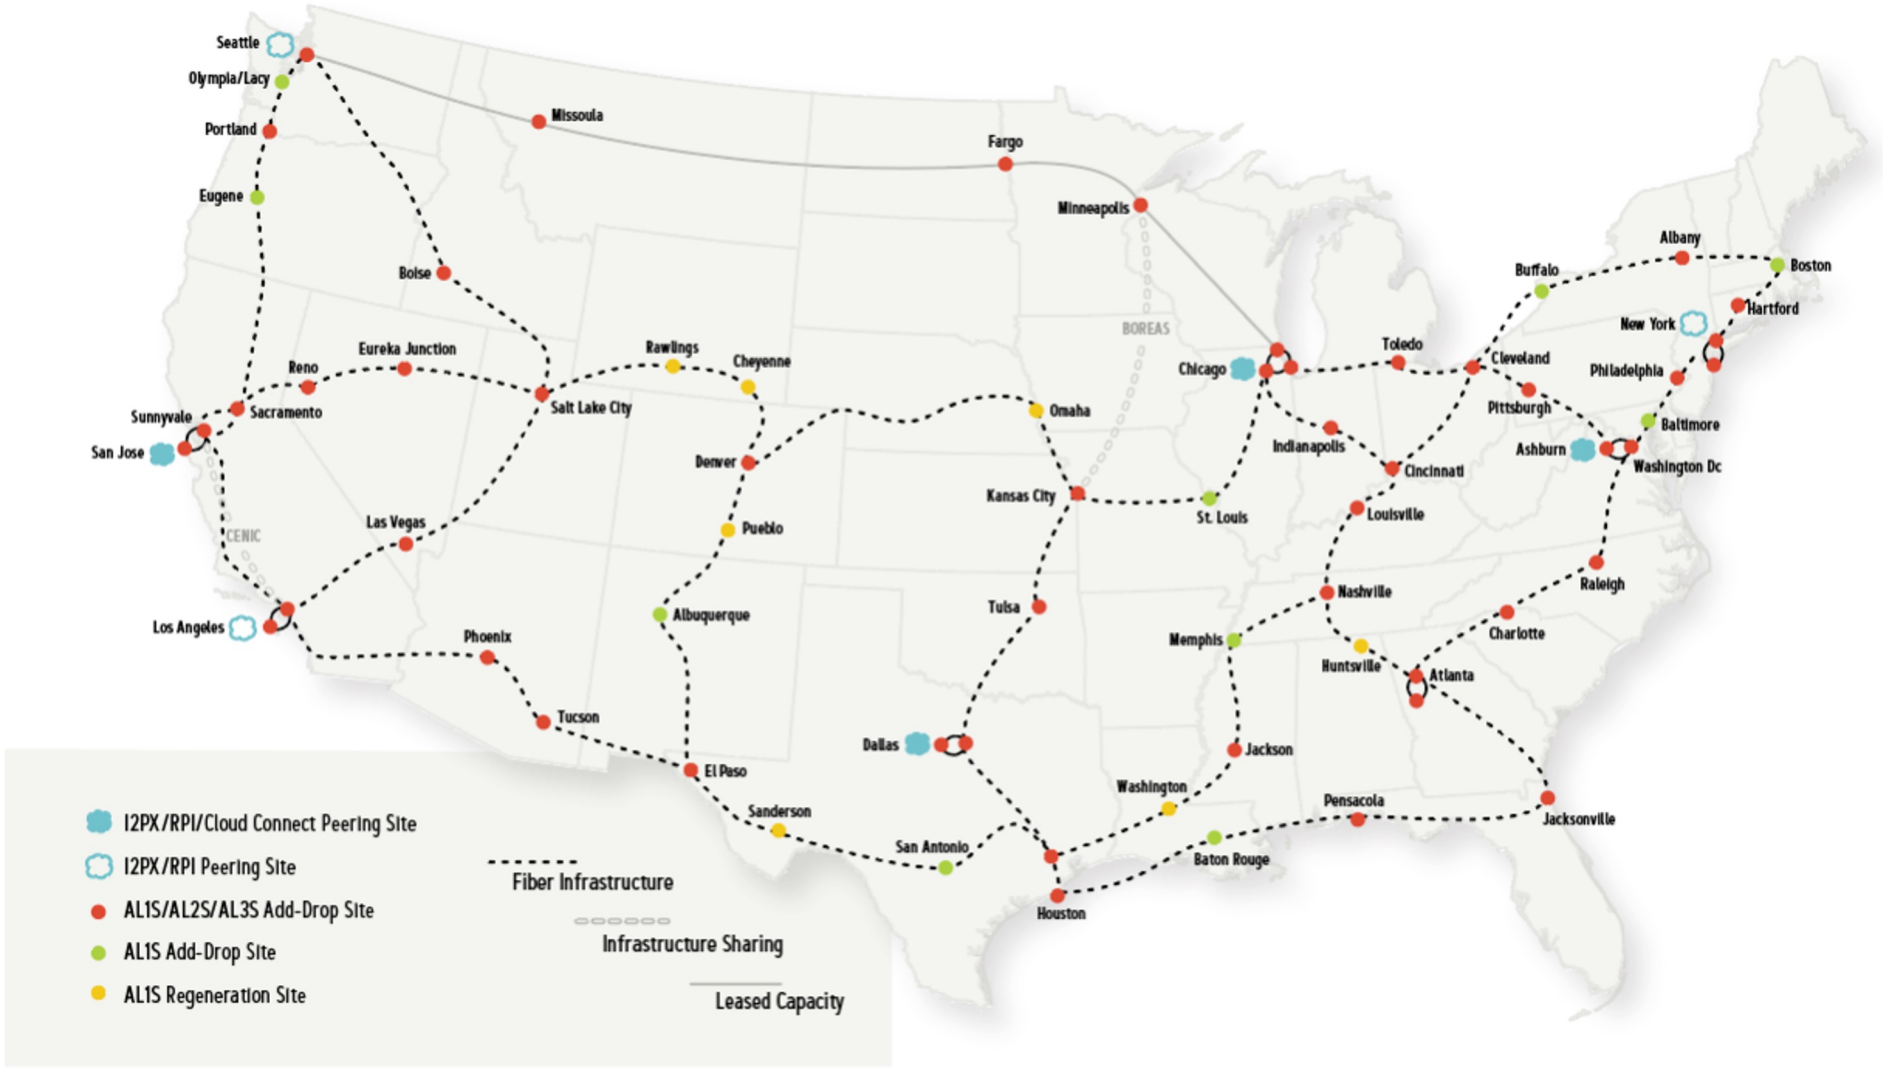

Vendors

perfSONAR
2x100 Cloud Edge Test Point

CISCO
Terabit-Scale Cloud Edge

ciena
Open Line System
Flex Grid

JUNIPER NETWORKS
Network Management
Security Overlay

perfSONAR
2x100 Community Test Point

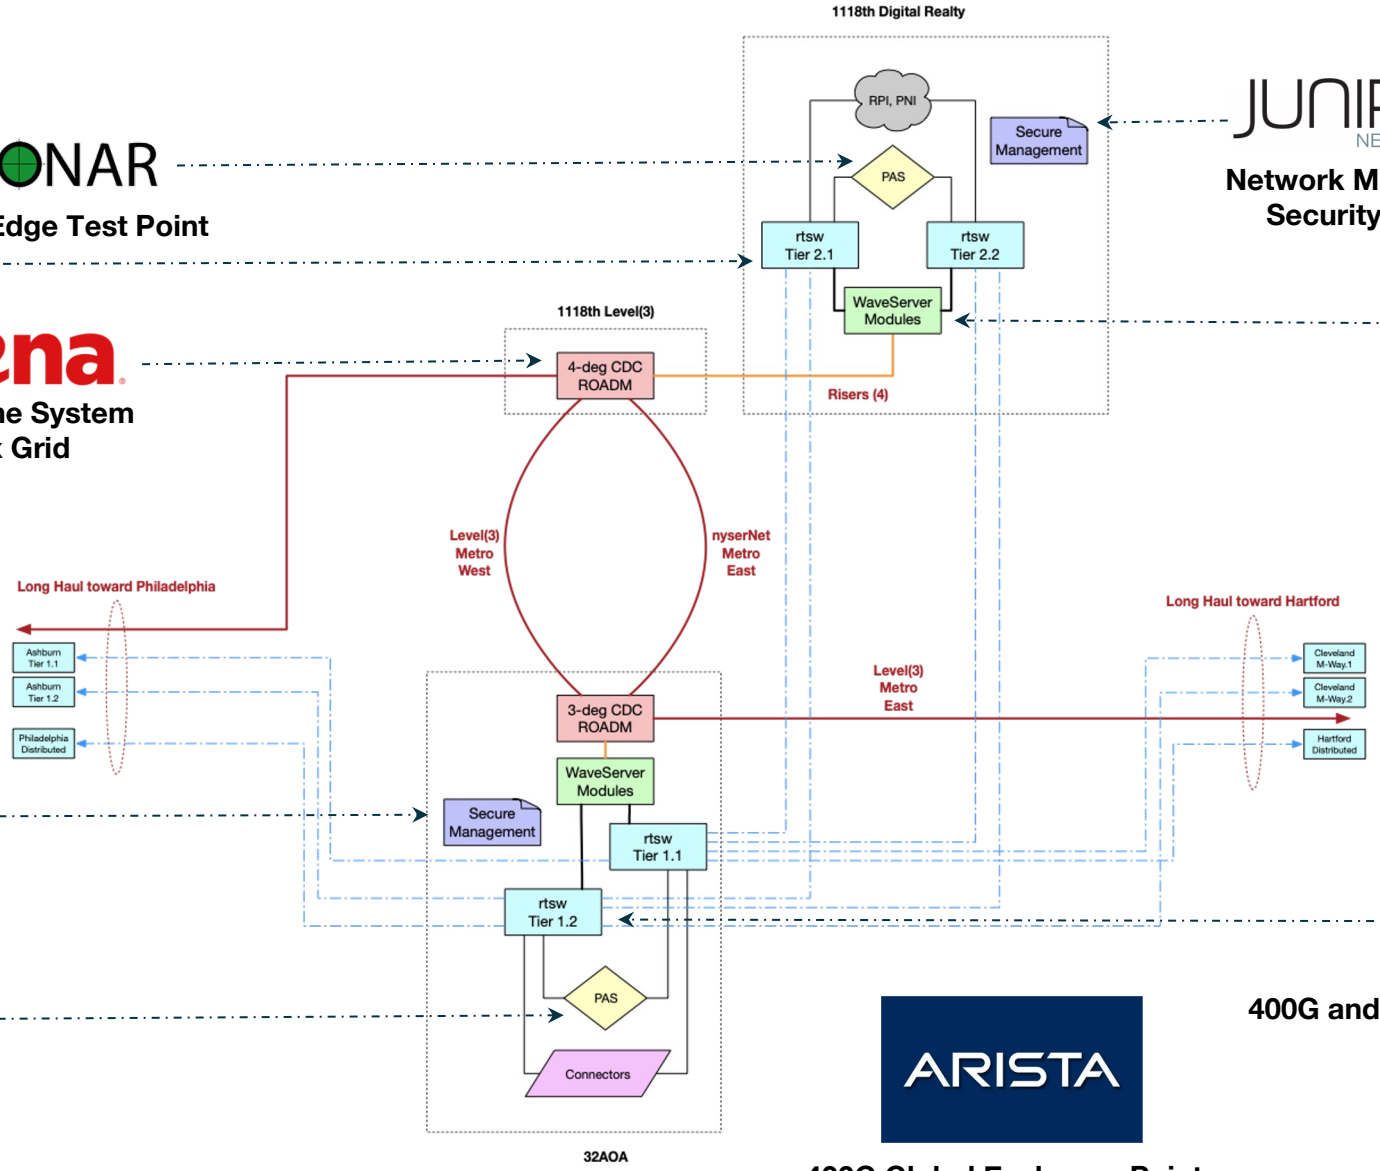

JUNIPER NETWORKS
Network Management
Security Overlay

CISCO
400G and 100G Community Access
All I2 Services

ciena
400G - 800G Transponders
Long Haul and Metro

ARISTA
400G Global Exchange Points
(Coming Soon)

Example based on New York Metro

Backbone Architecture

- Ninety-four 400 Gbps Backbone links
- 27,600 Tbps of deployed capacity
- 1.6 Tbps available contiguously coast to coast
- Depending on topology, all cities have 800G – 1.4 TB of exit capacity
- Each link is on non-regenerated wave

Northern Tier remains at 100G due to non-proliferation of 400G infrastructure on that corridor.

Backbone and Cloud Architecture

NGI Disaggregted Core Network is Operational and Performing Well

- SR-MPLS
- Shortest Path
- No Controller
- Path-Equivalency for Cloud
- Complexity Managed!

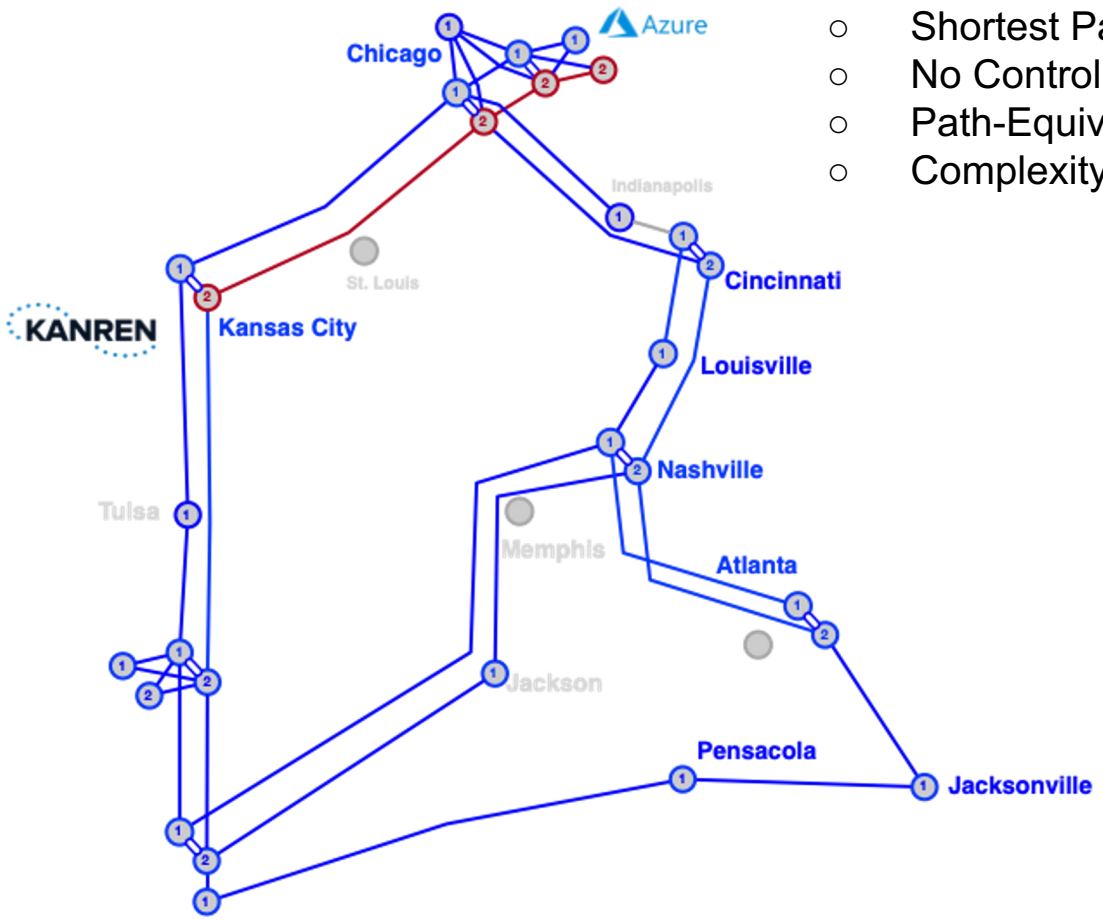

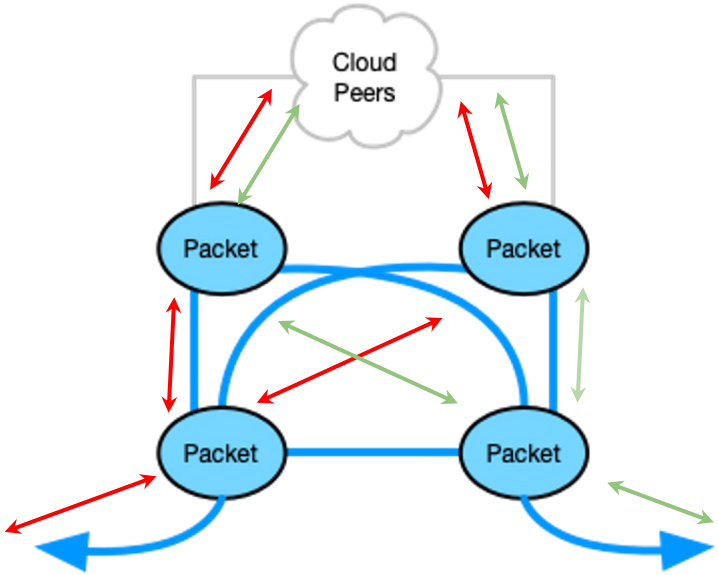

Backbone and Cloud Architecture

CDN / Peer	Gbps
Microsoft	1400
Amazon	800
Google	780
Akamai	700
Apple	700
Facebook	660
Comcast	630
Oracle	620
Fastly	560
Stackpath	400

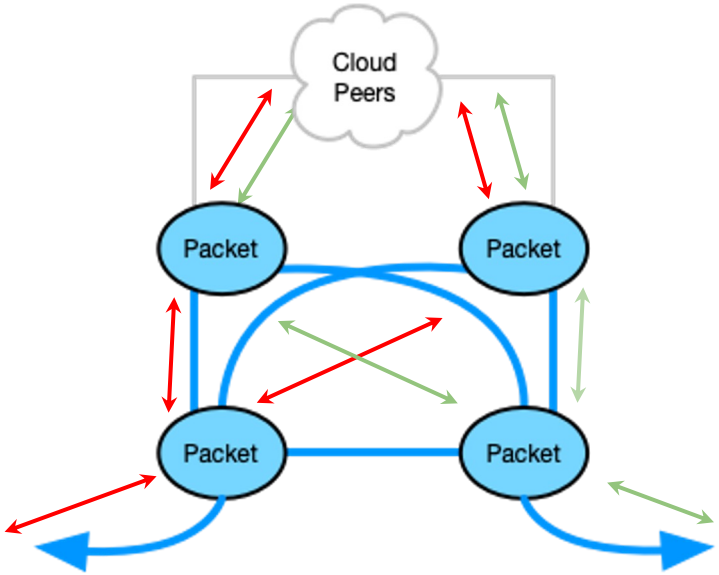

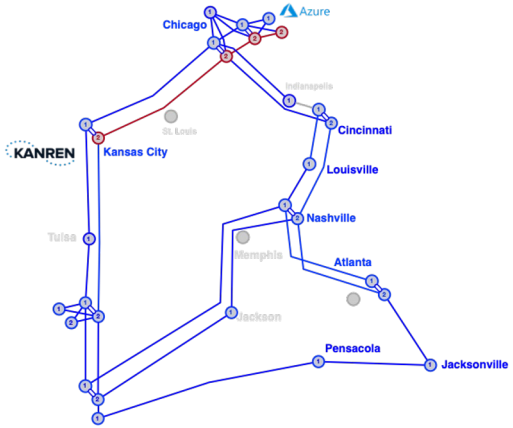

Open Science Data Federation (OSDF) Caches

National Research Platform

Open Cyberinfrastructure accessed by all 3,900 accredited degree granting higher education institutions in the US.

- Open Science
- Open Data
- Open Source
- Open Infrastructure

Currently 11 Nodes On-Net

- Target serving institutions which do not have on-site caches
- Storage, GPUs

APIs / Software Stack

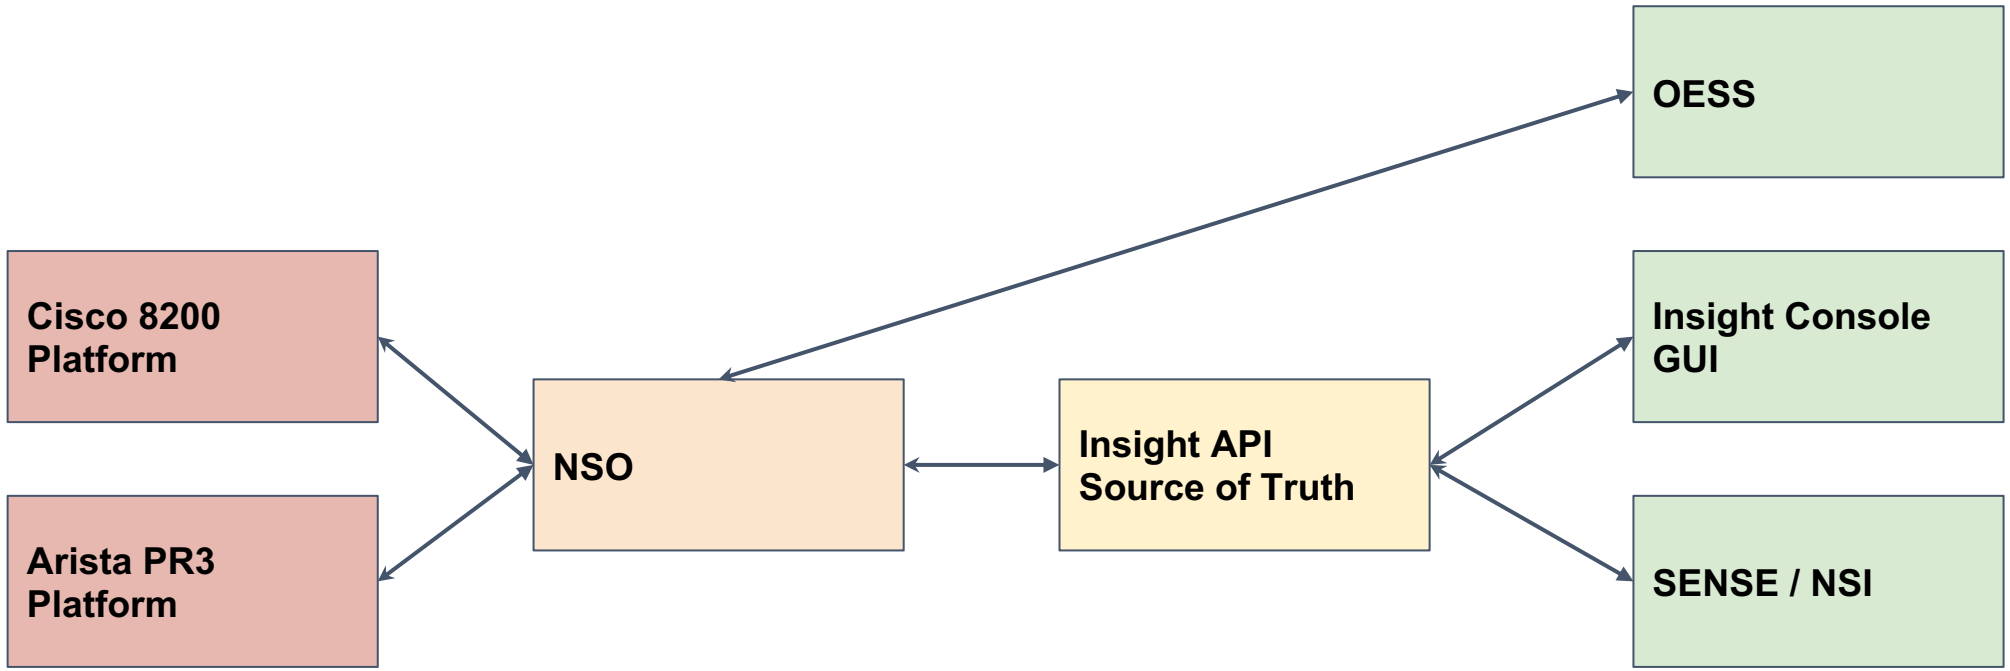

APIs / Insight Software Stack

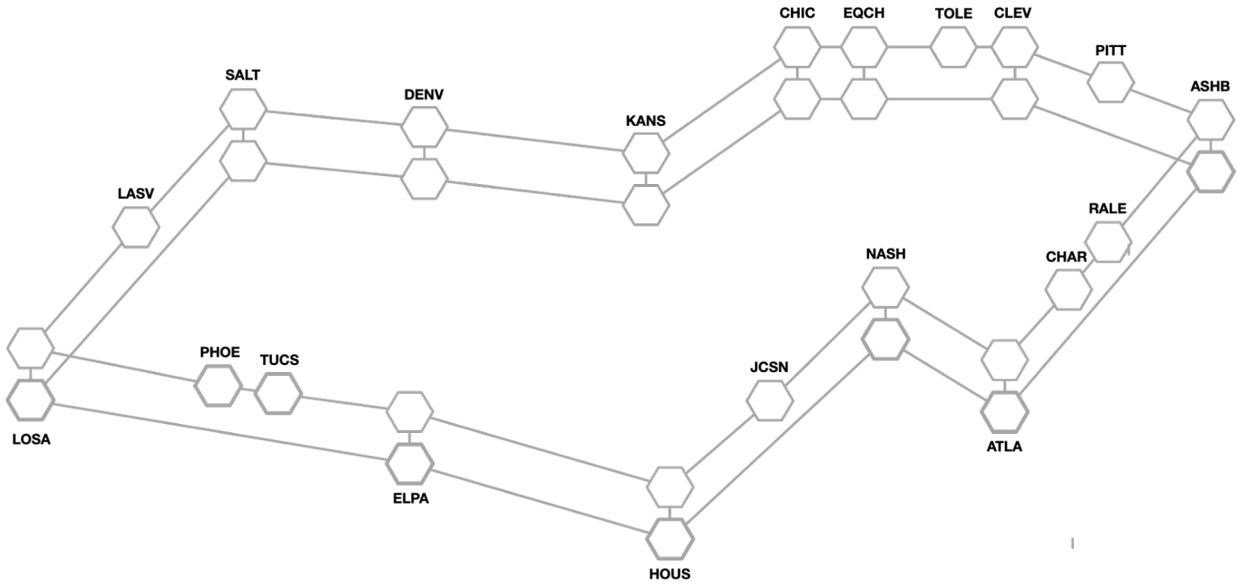

- Portal and APIs
  - Provision Layer2 and Layer3 Services
  - Allow for Path Selection, Optimization
  - Convey Complex Topology Information
  - Convey Key Statistics and Metrics
- Direct Access to Platform through Router Proxy
  - Debugging

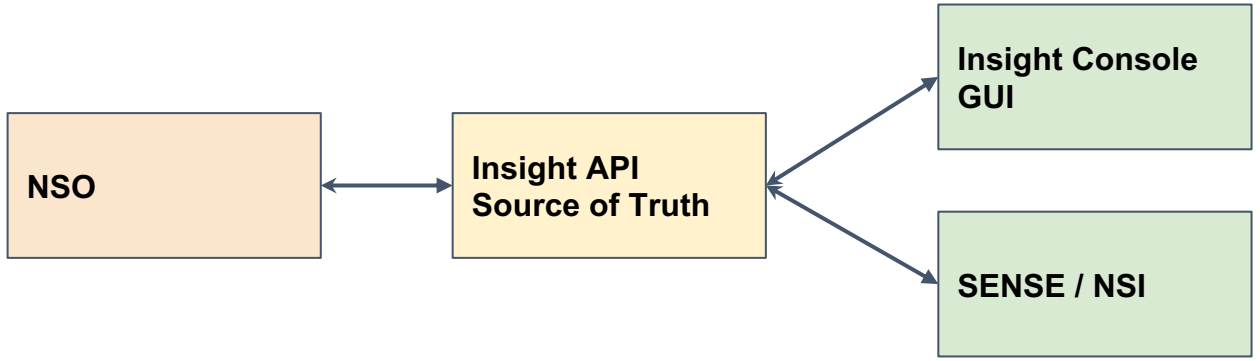

First 400G Transoceanic Link

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

- Exploring second 400G link into New York, Washington or other east-coast city TBD

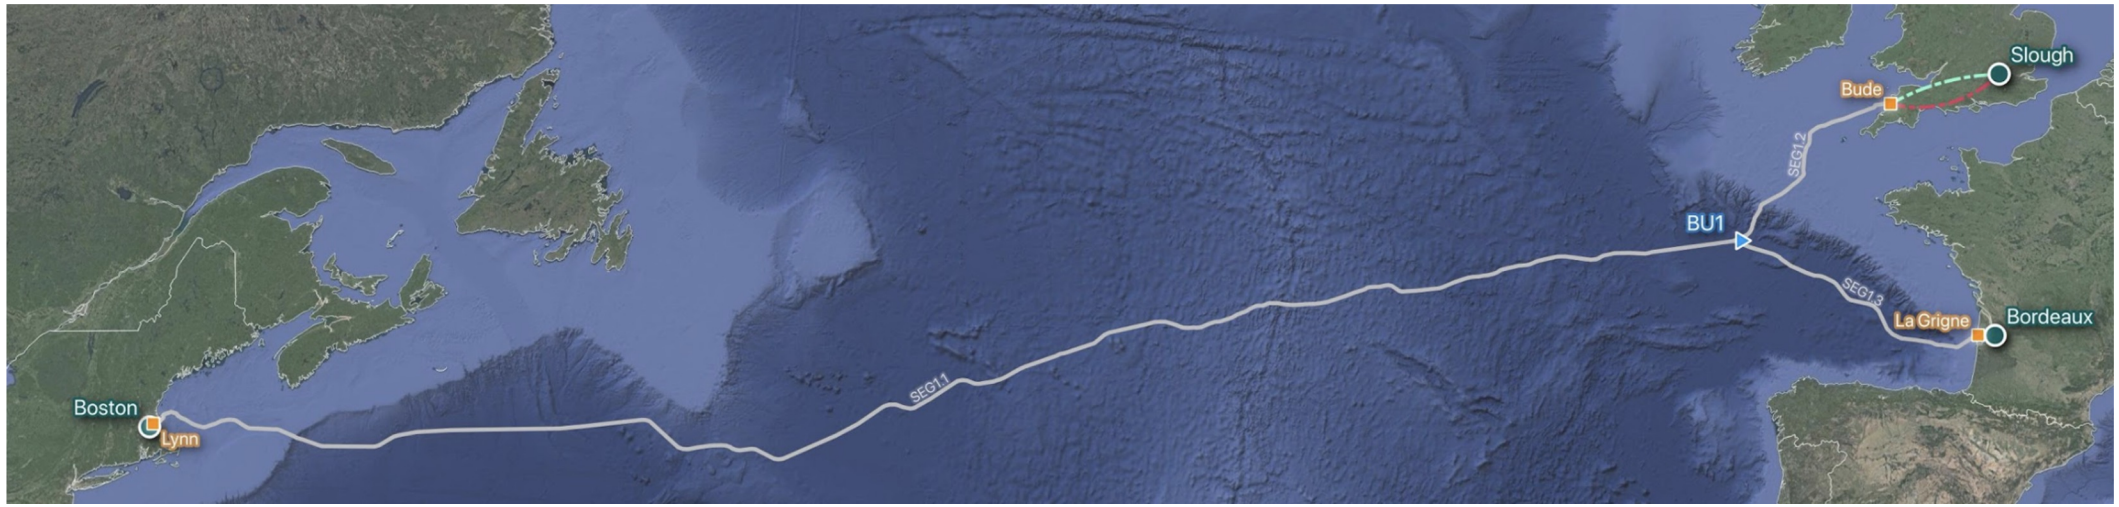

GXPs to 400G - Built with Automation

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

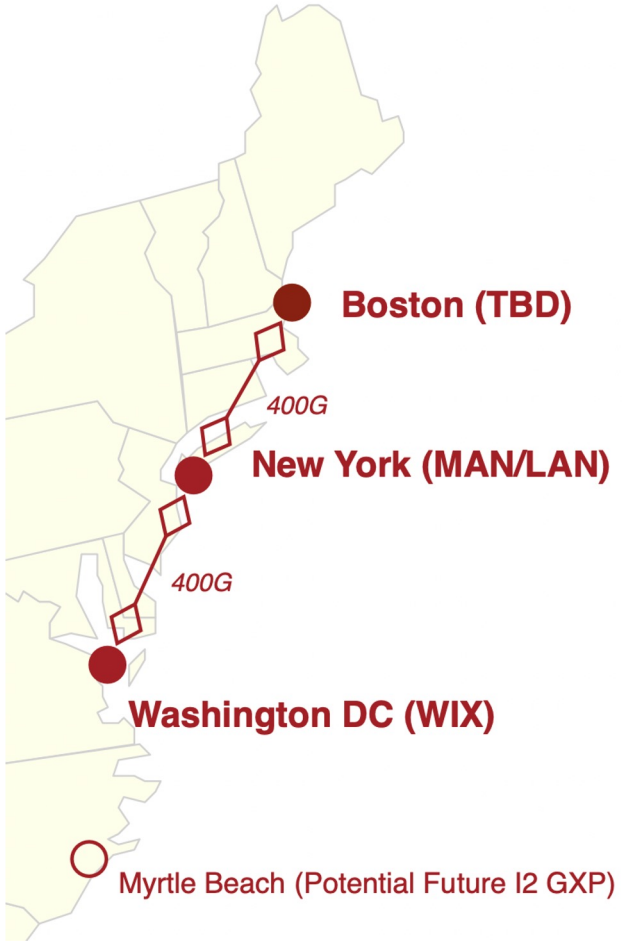

GXPs to 400G - Integrated into Core

Automating and Expanding Exchange Point Functionality; Match Community Development

- Maintain separate service layer for International Exchanges
- Integrate into core network SR-MPLS fabric
- Supports multi-vendor approach

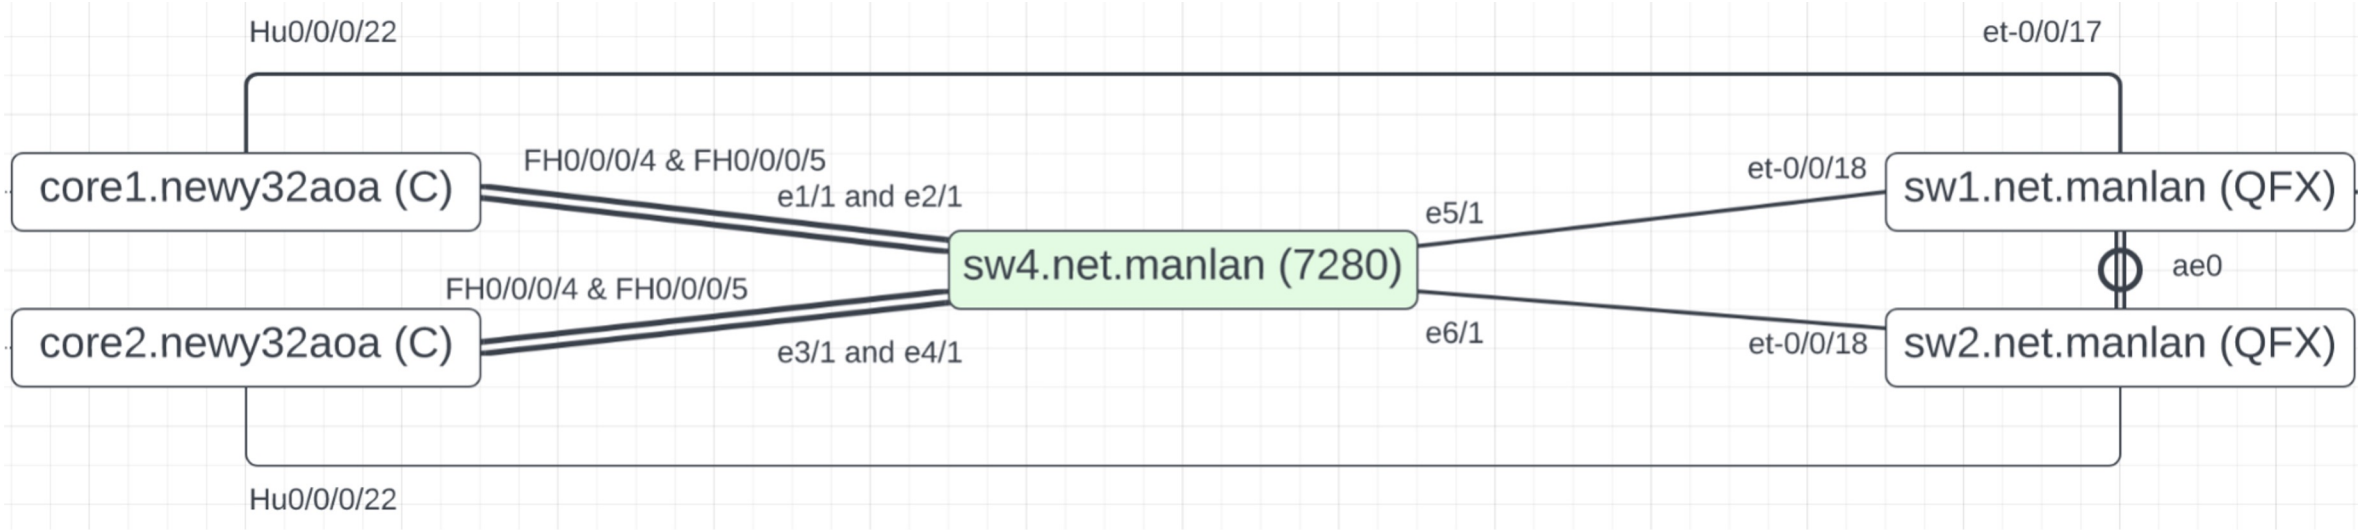

AP-REX 2.0 Vision

- Create Domestic Consortium of Exchange Point and Link Operators
- Improve Coordination of Activities - GXPs 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 Vision

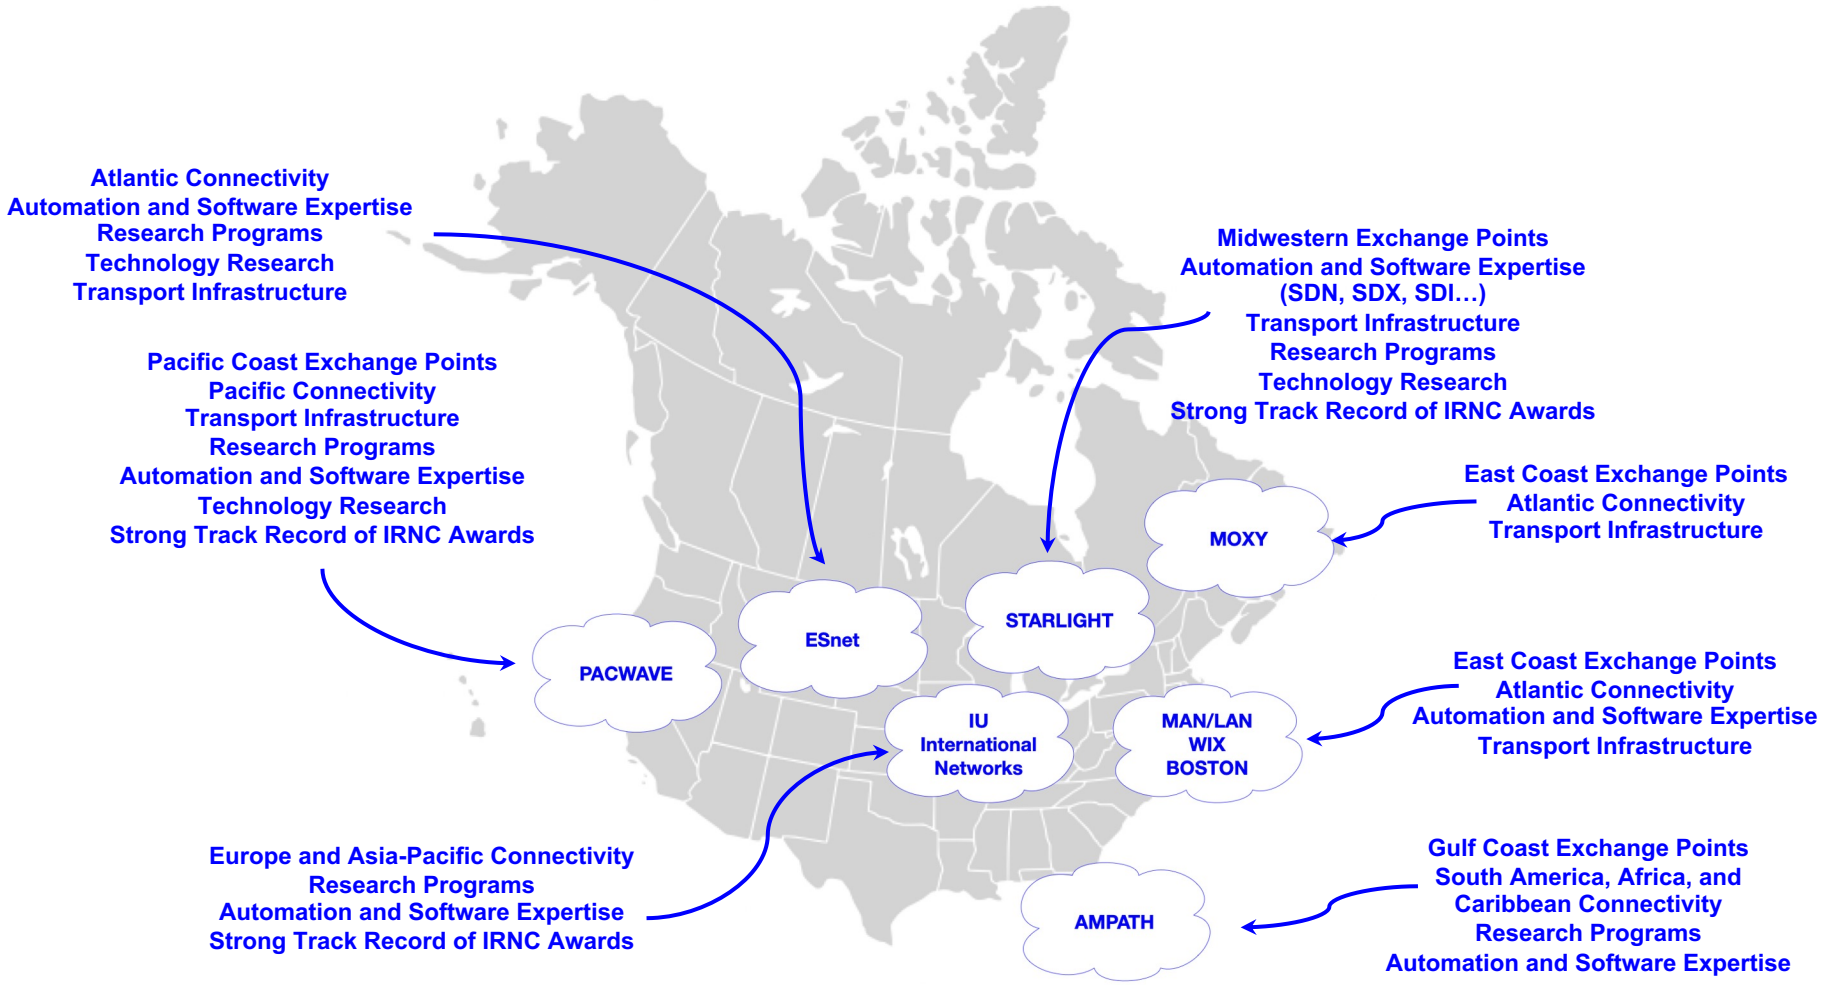

AP-REX 2.0 Vision

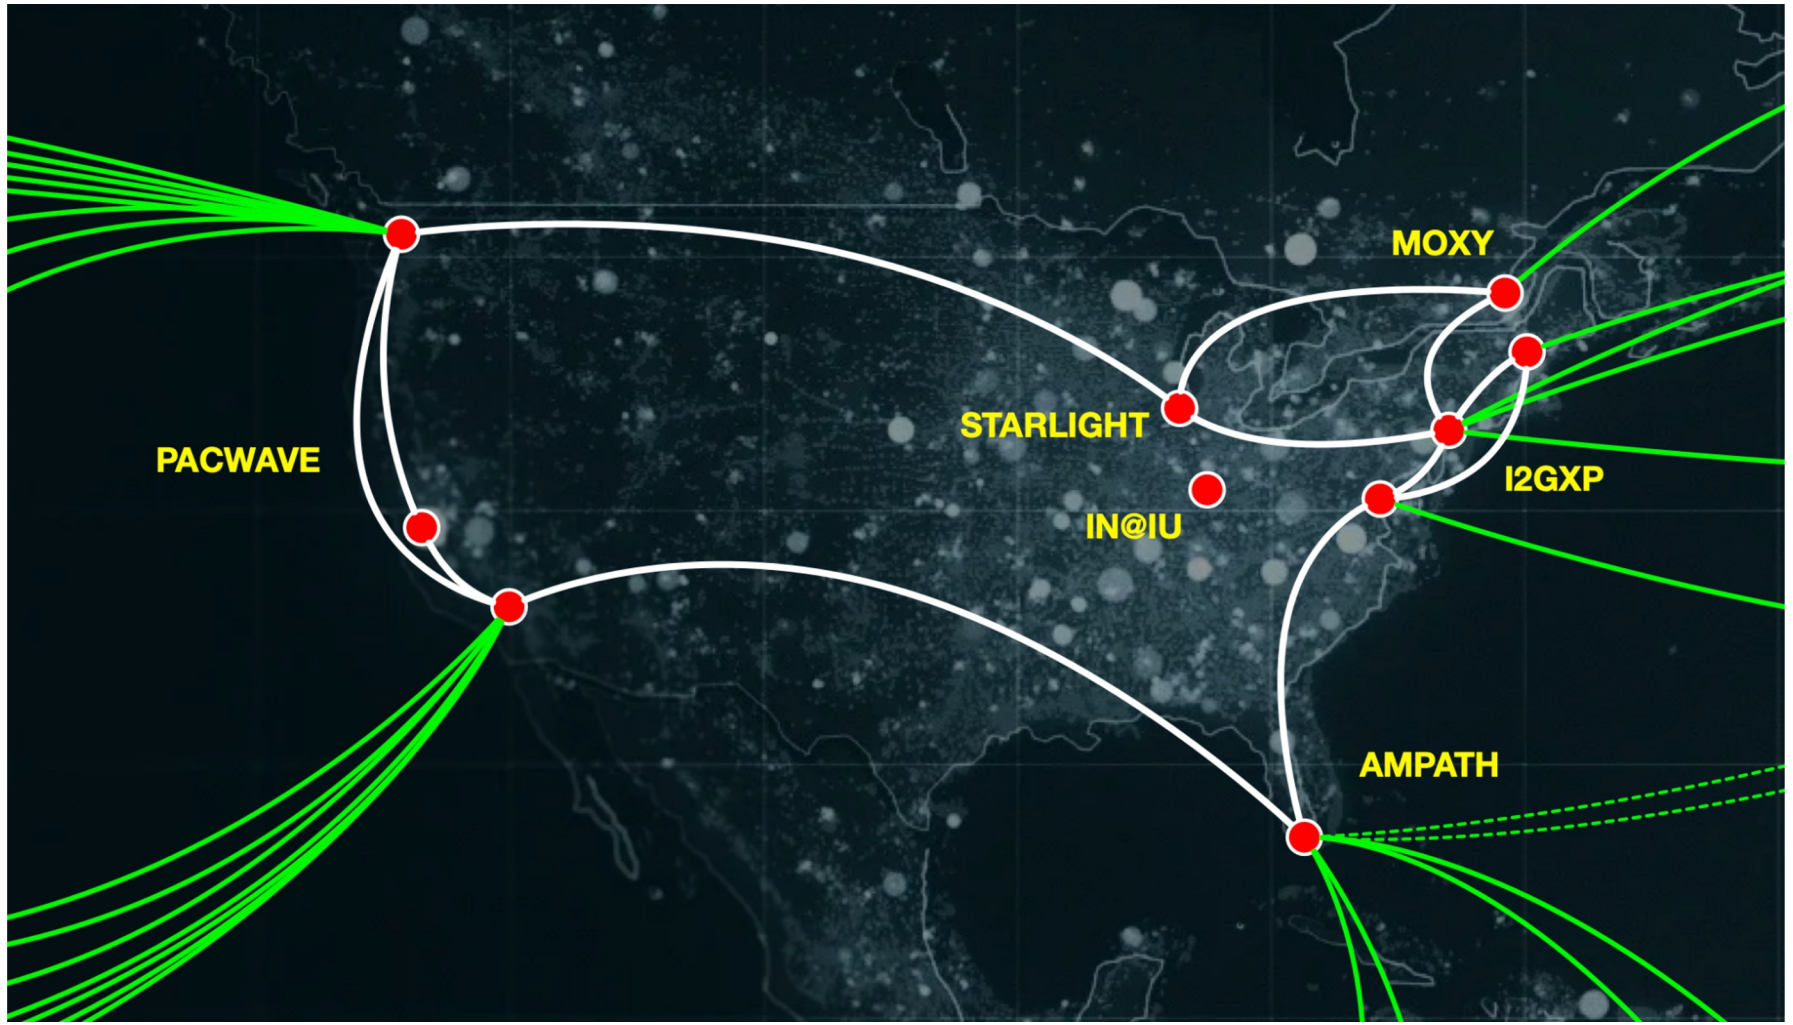

Initial AP-REX 2.0 Transport Architecture Concepts

Leverage Internet2 NGI Core Network for Routine GXP Transit Traffic

- Routine traffic; commodity purposes
- Provide Transit of US
- 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)

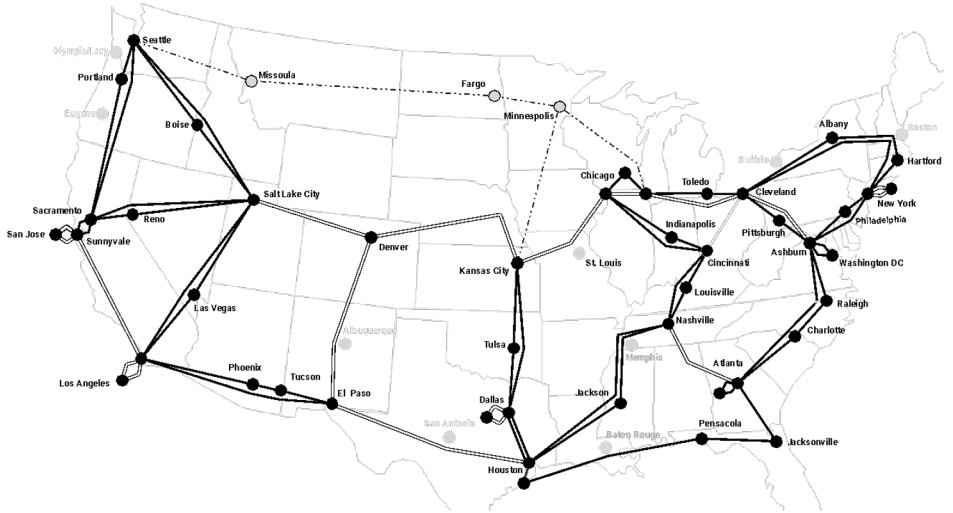

Leverage Dedicated Waves for Programmatic Experimentation

- 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!*

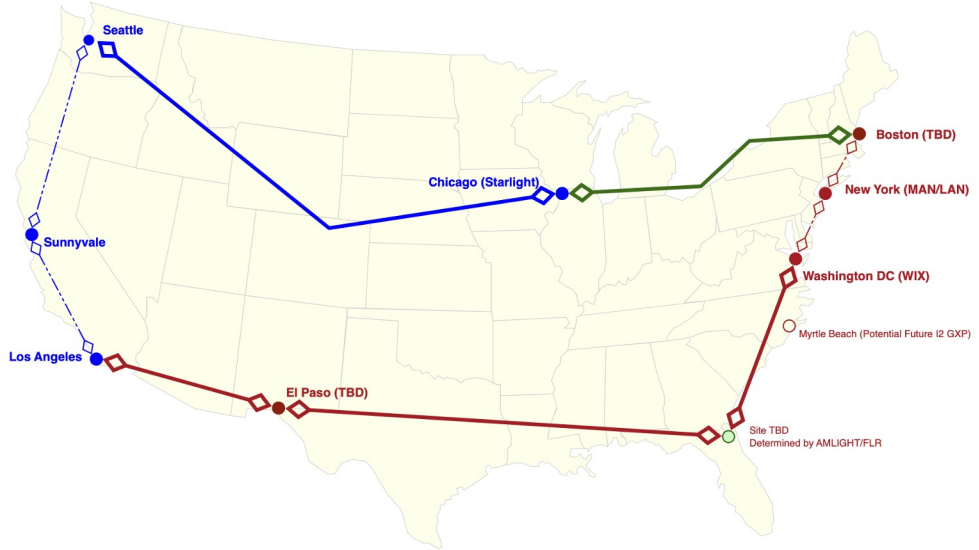

Core Network Experimentation

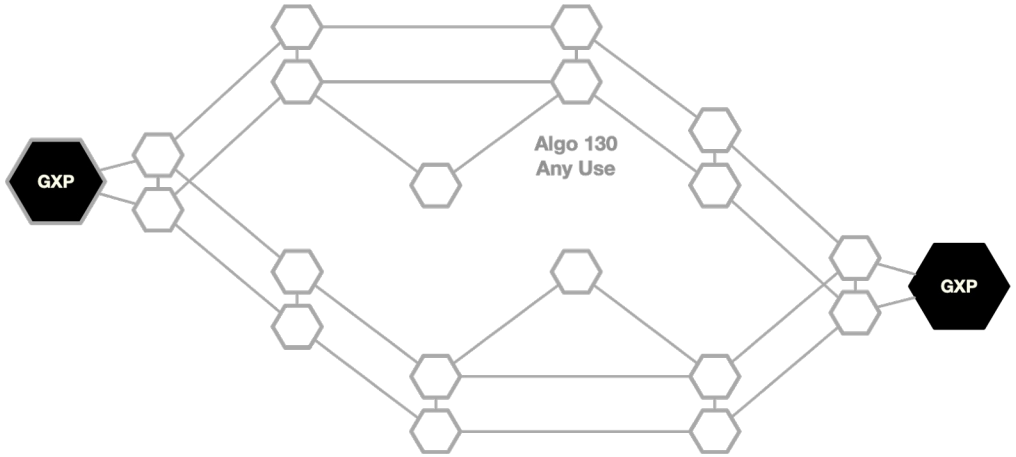

Flexible Algorithm (Flex-Algo)

- 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" or "label"** can be assigned to particular Algo at ingress
- Theoretically, traffic can be dynamically drained off of certain segments for DIS burst use

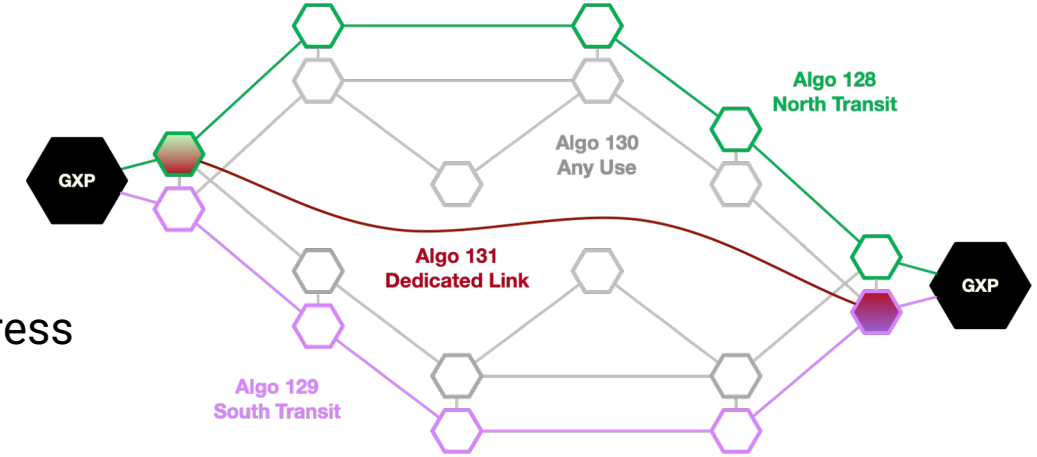

WHAT'S NEXT

Complete Deployment of New 400G Links and GXPs (Q2 - Q4 of 2023)

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

Experiment with Netsage (or similar) pipelines for reporting (late 2023, early 2024)

Experiment with Expanded Core Network Functionality to Support Data-Intensive Use Cases (2022-2023)
- Laboratory Experimentation with FLEX-ALGO and BGP-CT/BGP-CAR
- Potentially team up with group RNP and RENATER

Expand AP-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+)