AmLight ExP

Americas Lightpaths Express & Protect

WTR/PoP-Ba, Sep 10th, 2024

AmLight: Enabling research and collaboration between Latin America and the world: Past, Present, and Future

Jeronimo "Jab" Bezerra - FIU/AmLight

Outline

- Introduction
- What is AmLight?
- Network Connectivity
- The Challenge
- AmLight SDN
- Future

Why am I here?

- The PoP-BA/RNP operates regionally but it has a global influence:
  - Consulting, training, research, collaboration, and exporting talents (myself included)
  - UFBA/PoP-BA and FIU/CIARA are similar in many aspects
  - Also, AmLight's staff has former PoP-BA network engineers.

AmLight ExP
Americas Lightpaths Express & Protect

PoP BA
PONTO DE PRESENÇA
DA RNP BAHIA

Why am I here? [2]

- The WTR has been a great venue for AmLight to share its new developments

Experiência de conectividade no exterior
Humberto Galiza - AmLight
2017

Why am I here? [2]

- The WTR has been a great venue for AmLight to share its new developments

What is AmLight?

- A distributed academic exchange point built to enable collaboration among Latin America, Caribbean, Africa, and the U.S.

- Supported by the U.S. Nation Science Foundation via the IRNC program under award # OAC-2029283 for the 2021-2025.

- Partnerships with R&E networks in the U.S., Latin America, Caribbean and Africa.
  - Sharing infrastructure and human resources

AmLight ExP
Americas Lightpaths Express & Protect
(NSF Award # OAC-2029283)

2021-2025 AmLight-ExP Goals

- Vision: Continue enabling collaboration among researchers and network operators in Latin America, Africa, and the U.S. by providing reliable, sustainable, scalable, and high-performance network connectivity and services.

- Focus:
  - Supporting Service Level Agreement (SLA)-driven science applications
  - Improving network visibility and management
  - Enable integration between AmLight and network-aware science drivers
  - Add new network and cloud services
  - Minimize the human role in network operation

Network Connectivity

- 2.1+ Tbps of international connectivity
  - 600Gbps of upstream capacity between the U.S. and Latin America, and 100Gbps to Africa
    - Expanding to **1.1Tbps** in 2024/2025!
  - 400G activated between Sao Paulo and Buenos Aires

- By 2025, AmLight will reach 4.9 Tbps of total capacity

- NAPs:
  - Florida (Miami, Boca Raton, Jacksonville), Brazil (Sao Paulo, Fortaleza), Chile (Santiago), Puerto Rico (San Juan), Panama (Panama City), South Africa (Cape Town),
  - New: Georgia (Atlanta), Argentina (Buenos Aires)

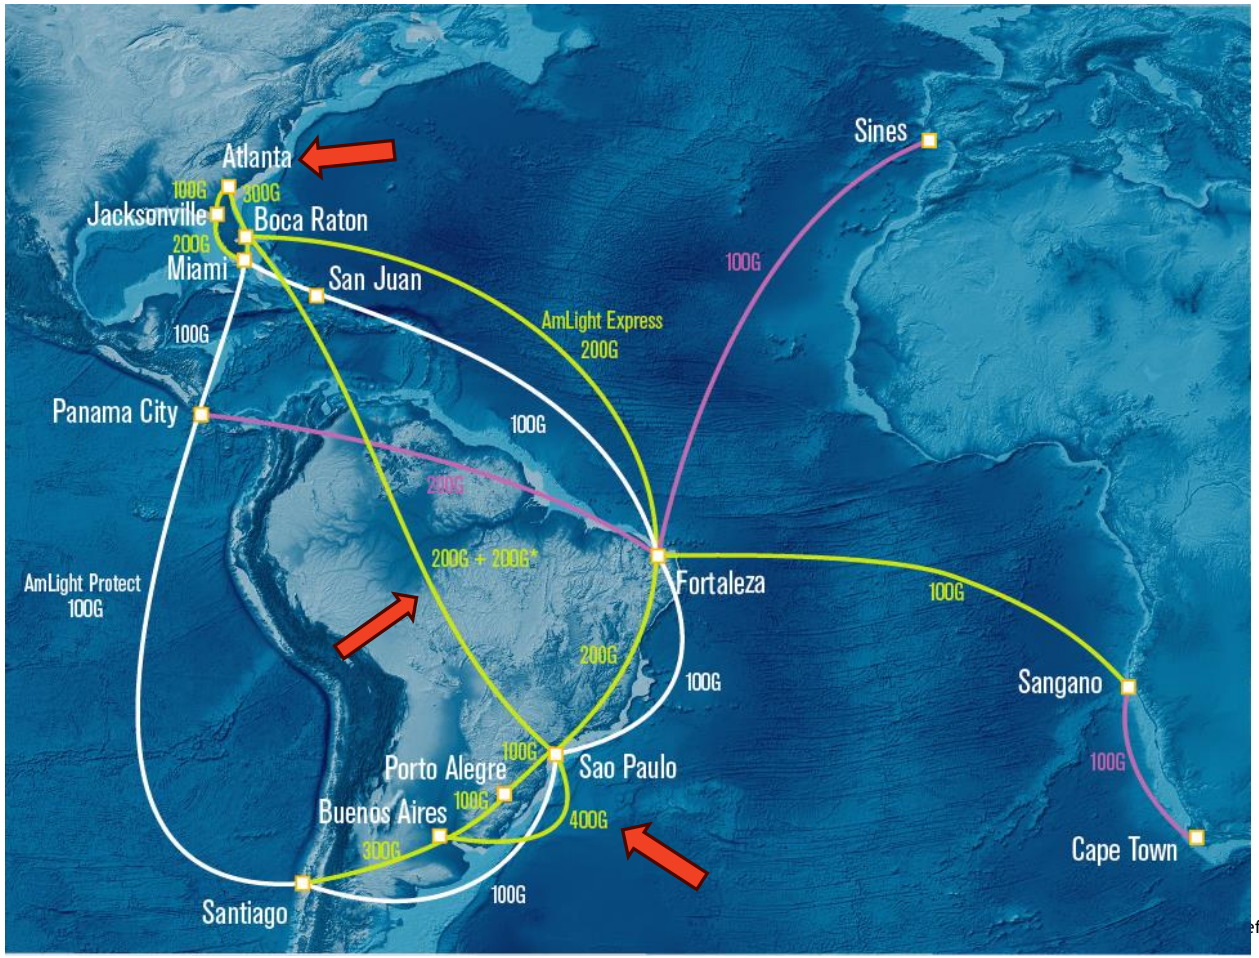

efforts

AmLight ExP
Americas Lightpaths Express & Protect

Network Connectivity [2]

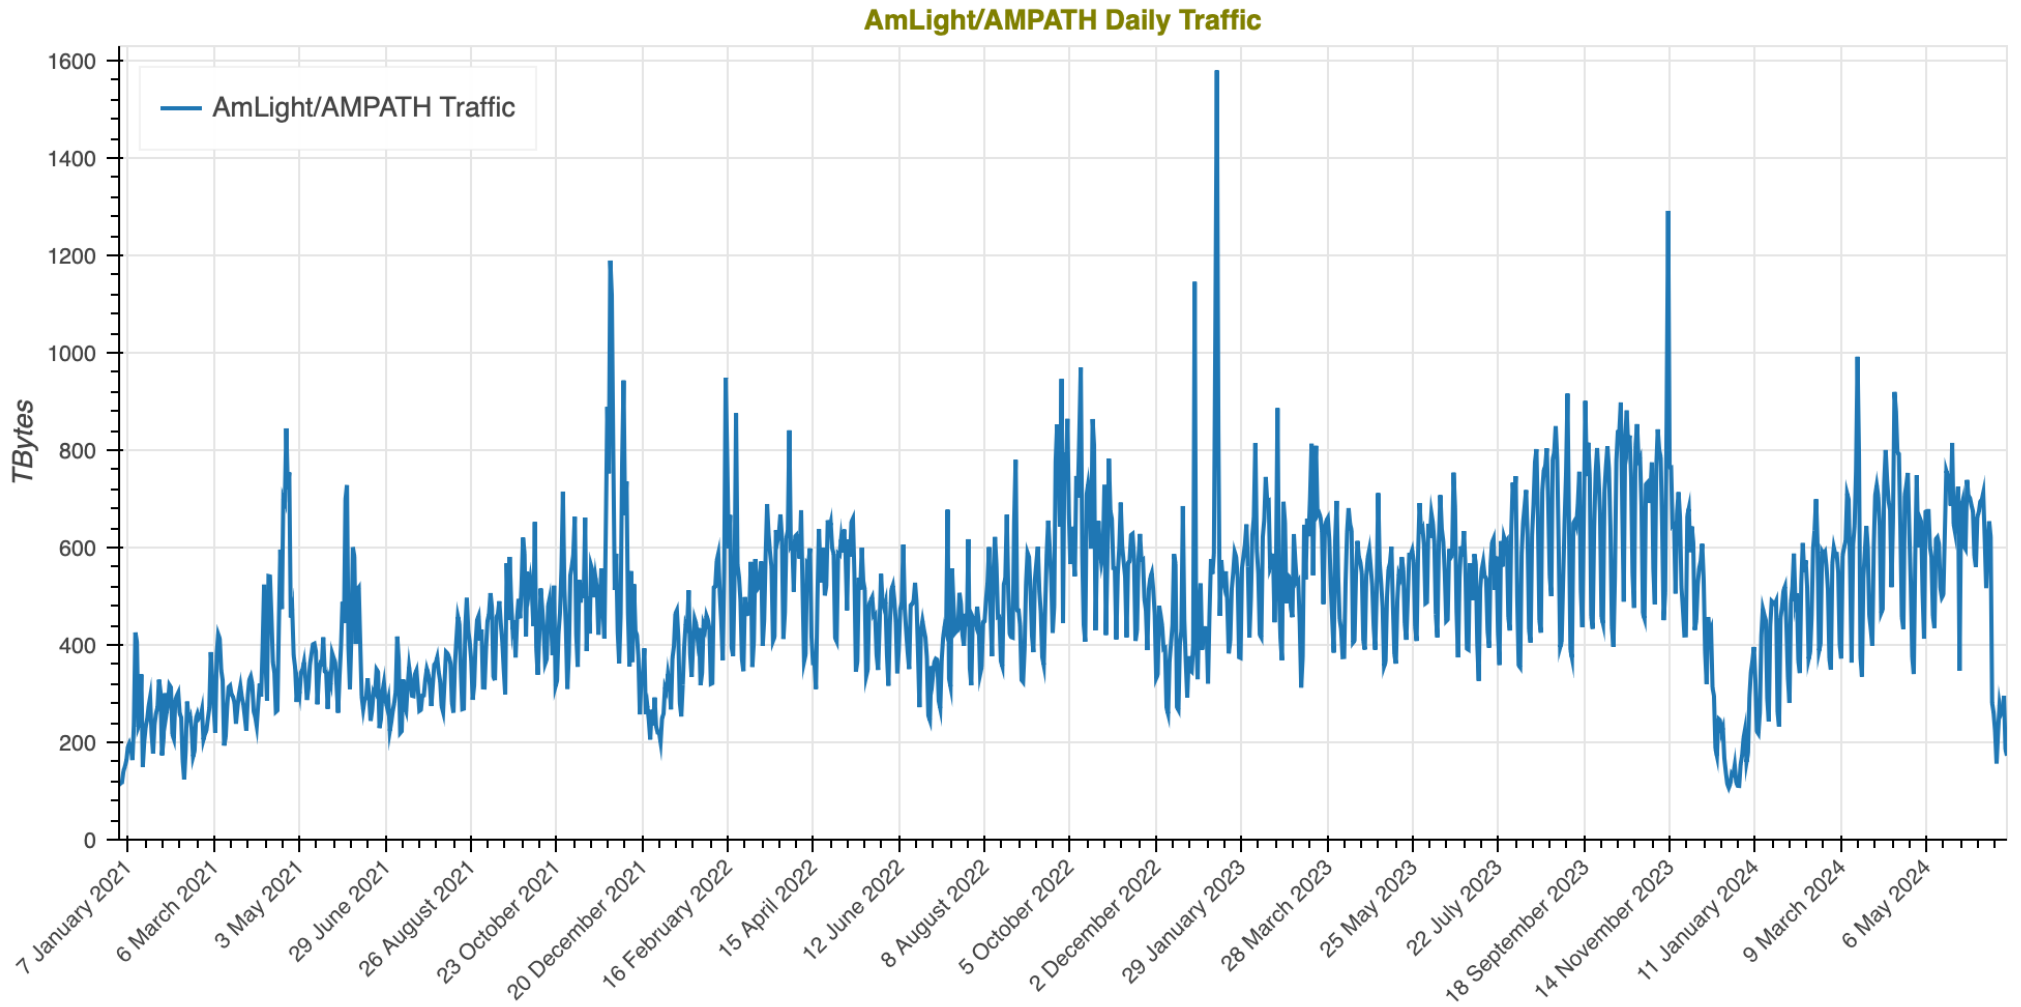

AmLight: Collaboration at its finest!

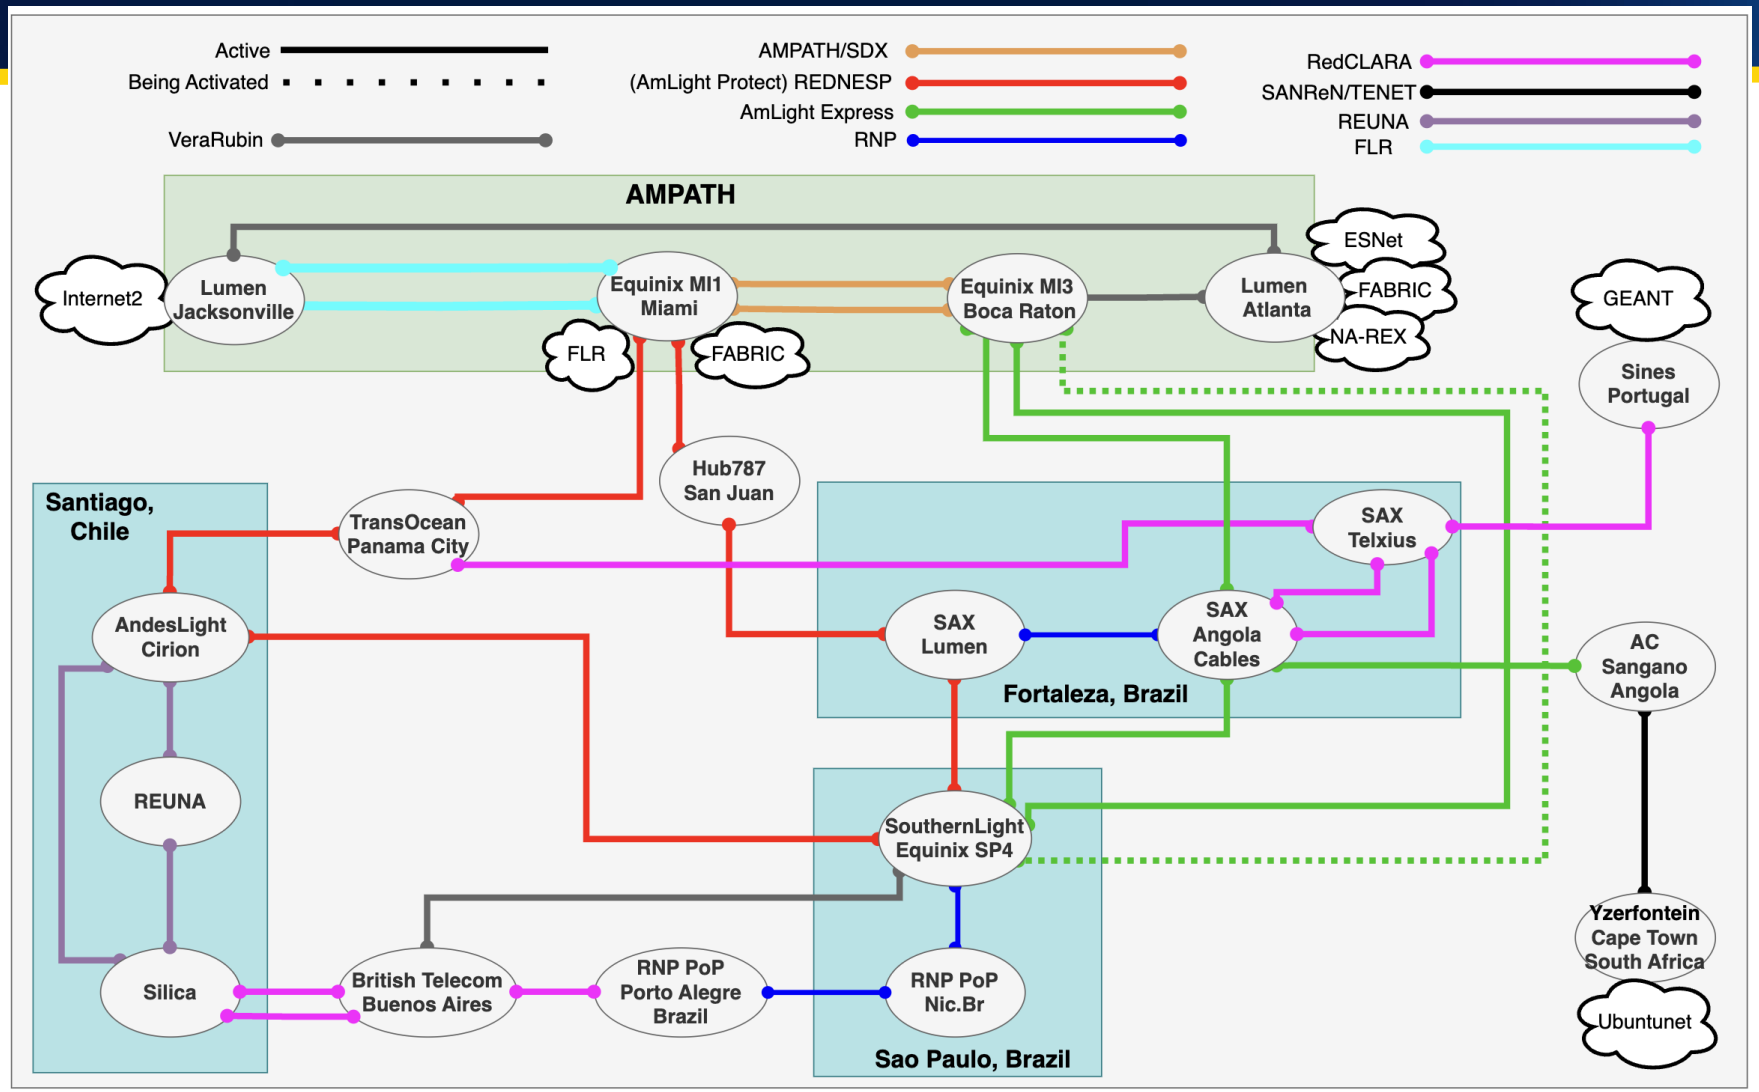

AmLight ExP
Americas Lightpaths Express & Protect

Network Connectivity in the U.S.

- 1.2 Tbps (3x400 Gbps) to be activated to Internet2 via NA-REX collaboration
- 300 Gbps activated with the Energy Science Network (Esnet)
- 200 Gbps activated for FABRIC testbed

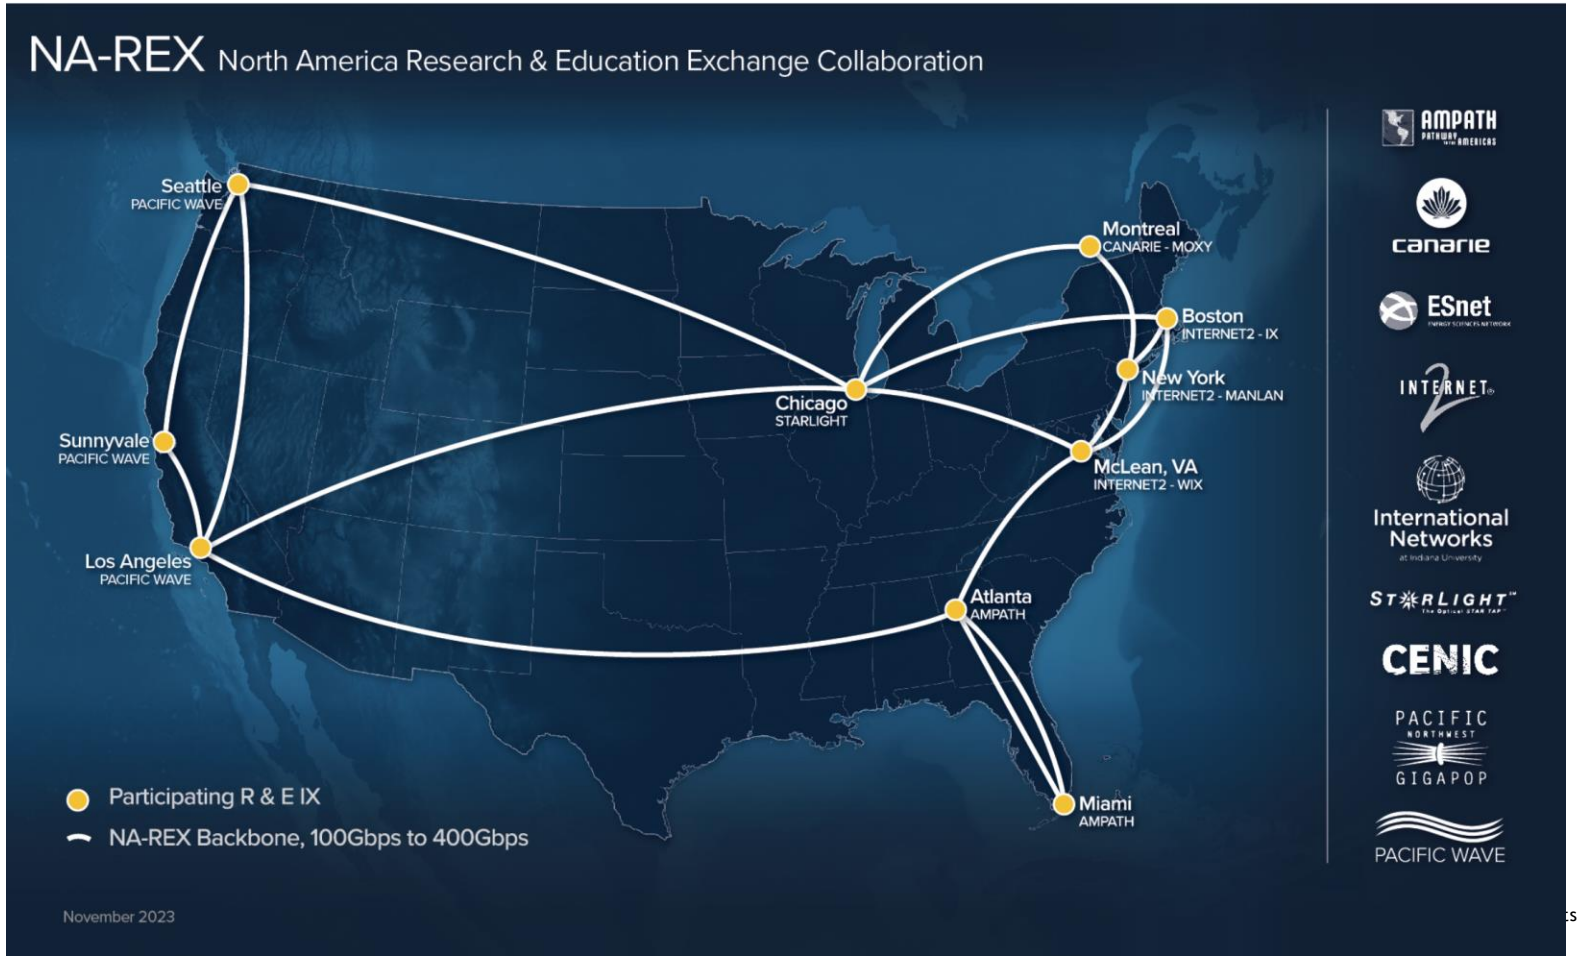

AmLight SDN Architecture…

Key motivations to move to SDN back in 2014

Improving operations efficiency

- *How to improve layer 2 multi-domain provisioning?*

Introducing network programmability

- *How to support network testbeds managing the network infrastructure in a secure way?*

	Average time to provision a new circuit		Avg. number of e-mails exchanged	
Domains involved in the path	***before SDN***	***after SDN***	***before SDN***	***after SDN***
RNP, ANSP, RedCLARA, AmLight, Internet2, ESnet	5 days	< 5 minutes	10	0
Other domains using OSCARS or NSI support	12 days	< 5 minutes	65	0

Reference: Benefits brought by the use of OpenFlow/SDN on the AmLight intercontinental research and education network, 2015

The Challenge

Vera Ru

nobservatory.org

Vera Rubin Obs's Data Transfers in a Nutshell

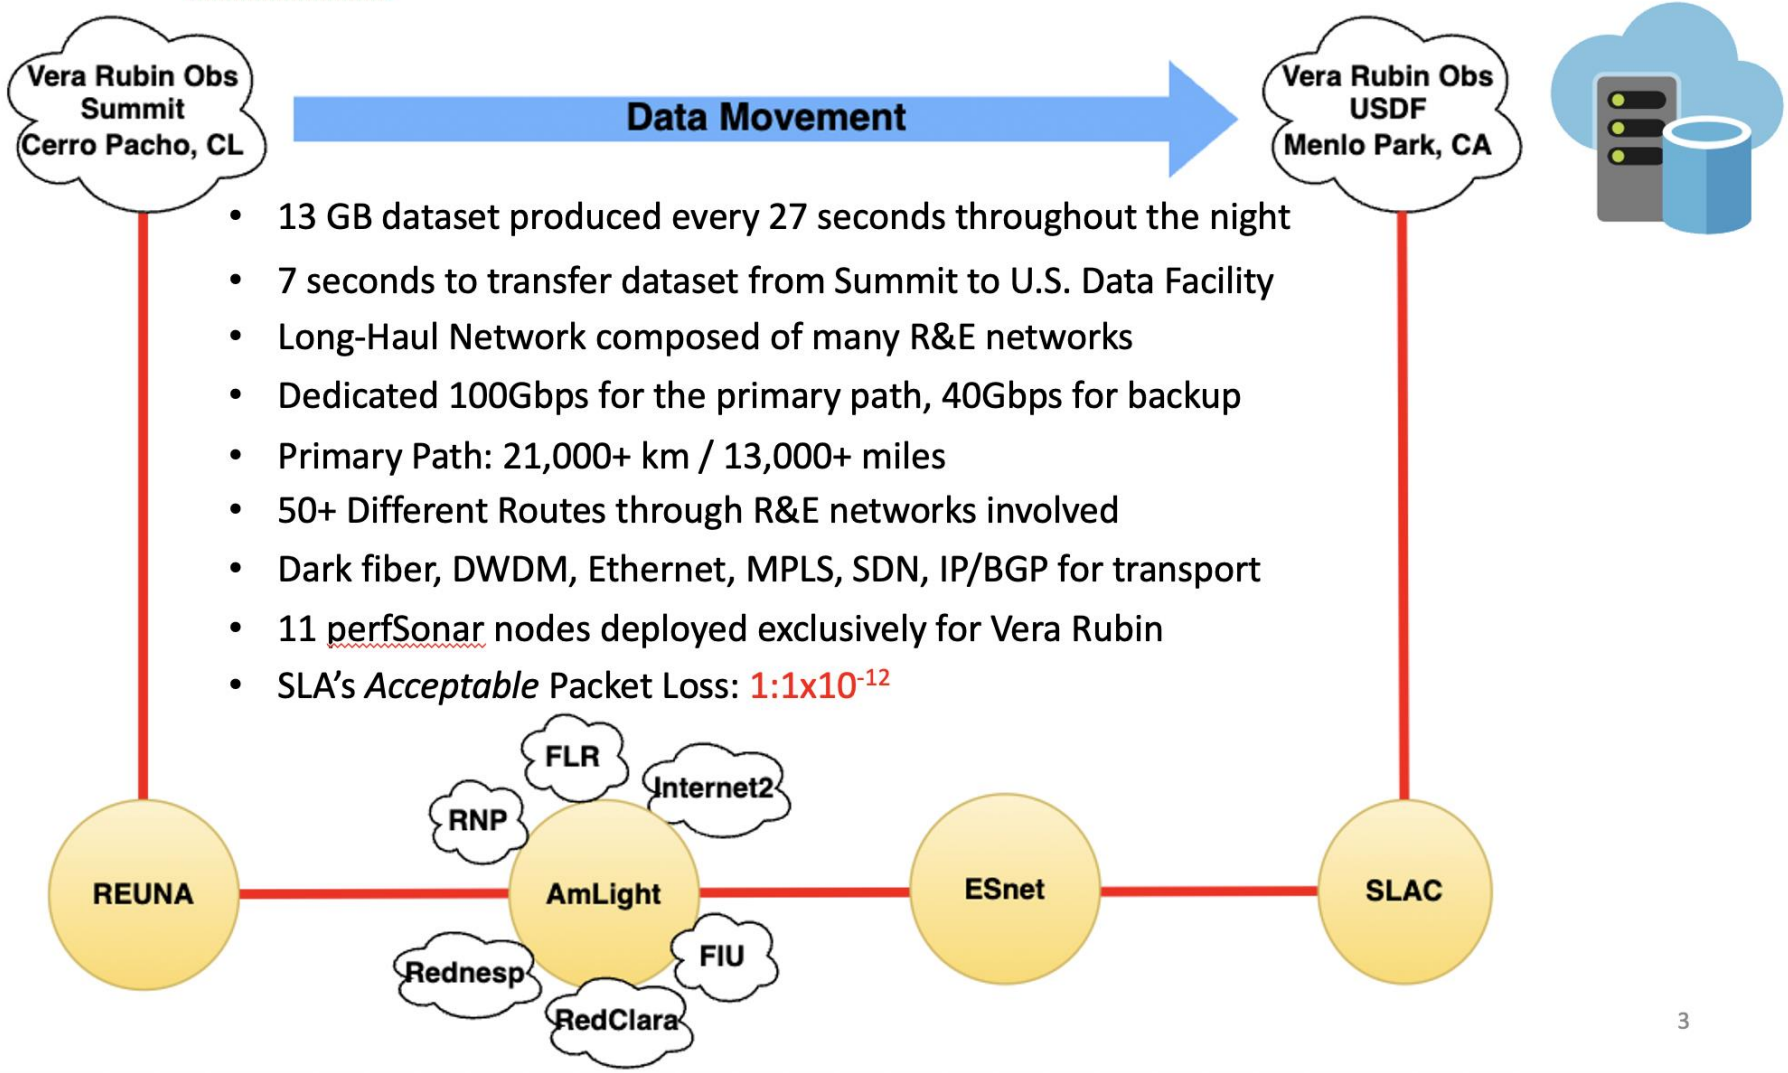

- 13 GB dataset produced every 27 seconds throughout the night
- 7 seconds to transfer dataset from Summit to U.S. Data Facility
- Long-Haul Network composed of many R&E networks
- Dedicated 100Gbps for the primary path, 40Gbps for backup
- Primary Path: 21,000+ km / 13,000+ miles
- 50+ Different Routes through R&E networks involved
- Dark fiber, DWDM, Ethernet, MPLS, SDN, IP/BGP for transport
- 11 perfSonar nodes deployed exclusively for Vera Rubin
- SLA's *Acceptable* Packet Loss: $1{:}1x10^{-12}$

AmLight SDN Architecture – 2021-2025

For 2021-2025, we will add specialized components per SDN Plane:

- New Control Plane:
  - Kytos-ng SDN Platform developed to address AmLight's and science drivers' requirements
- New Data Plane:
  - Programmable switches replacing legacy devices with customizable P4 pipeline
- New Management Plane:
  - Granular real-time network visibility with In-band Network Telemetry (INT)
- New Intelligence Plane:
  - Traffic engineering and optimizations based on granular telemetry reports in real-time

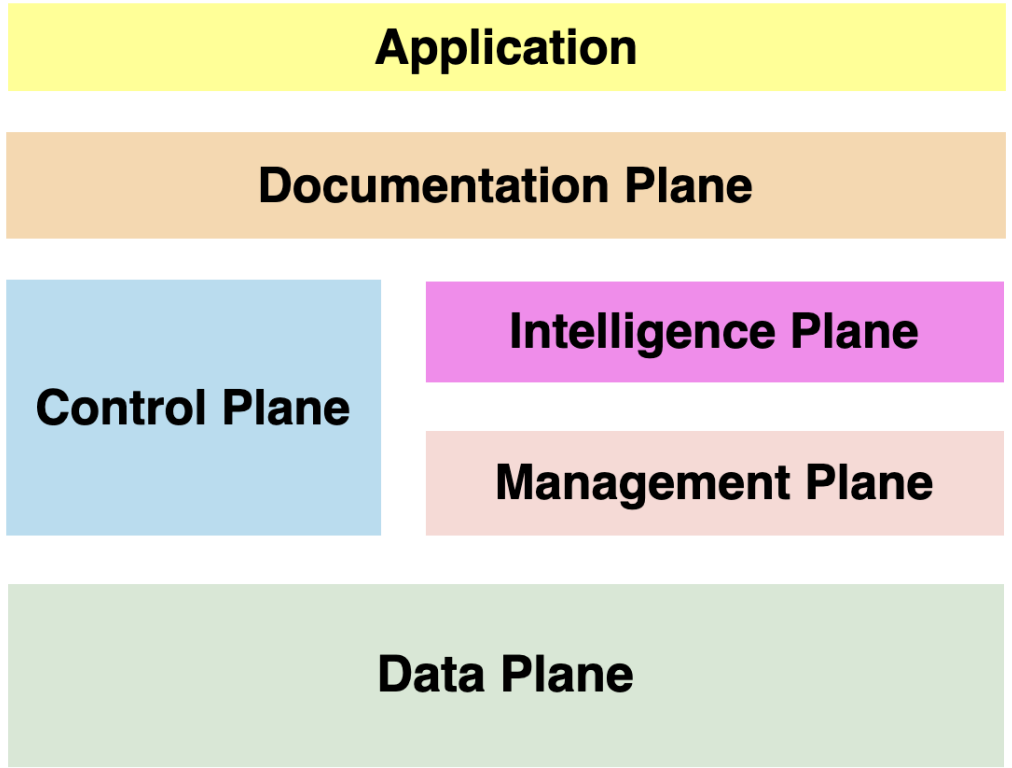

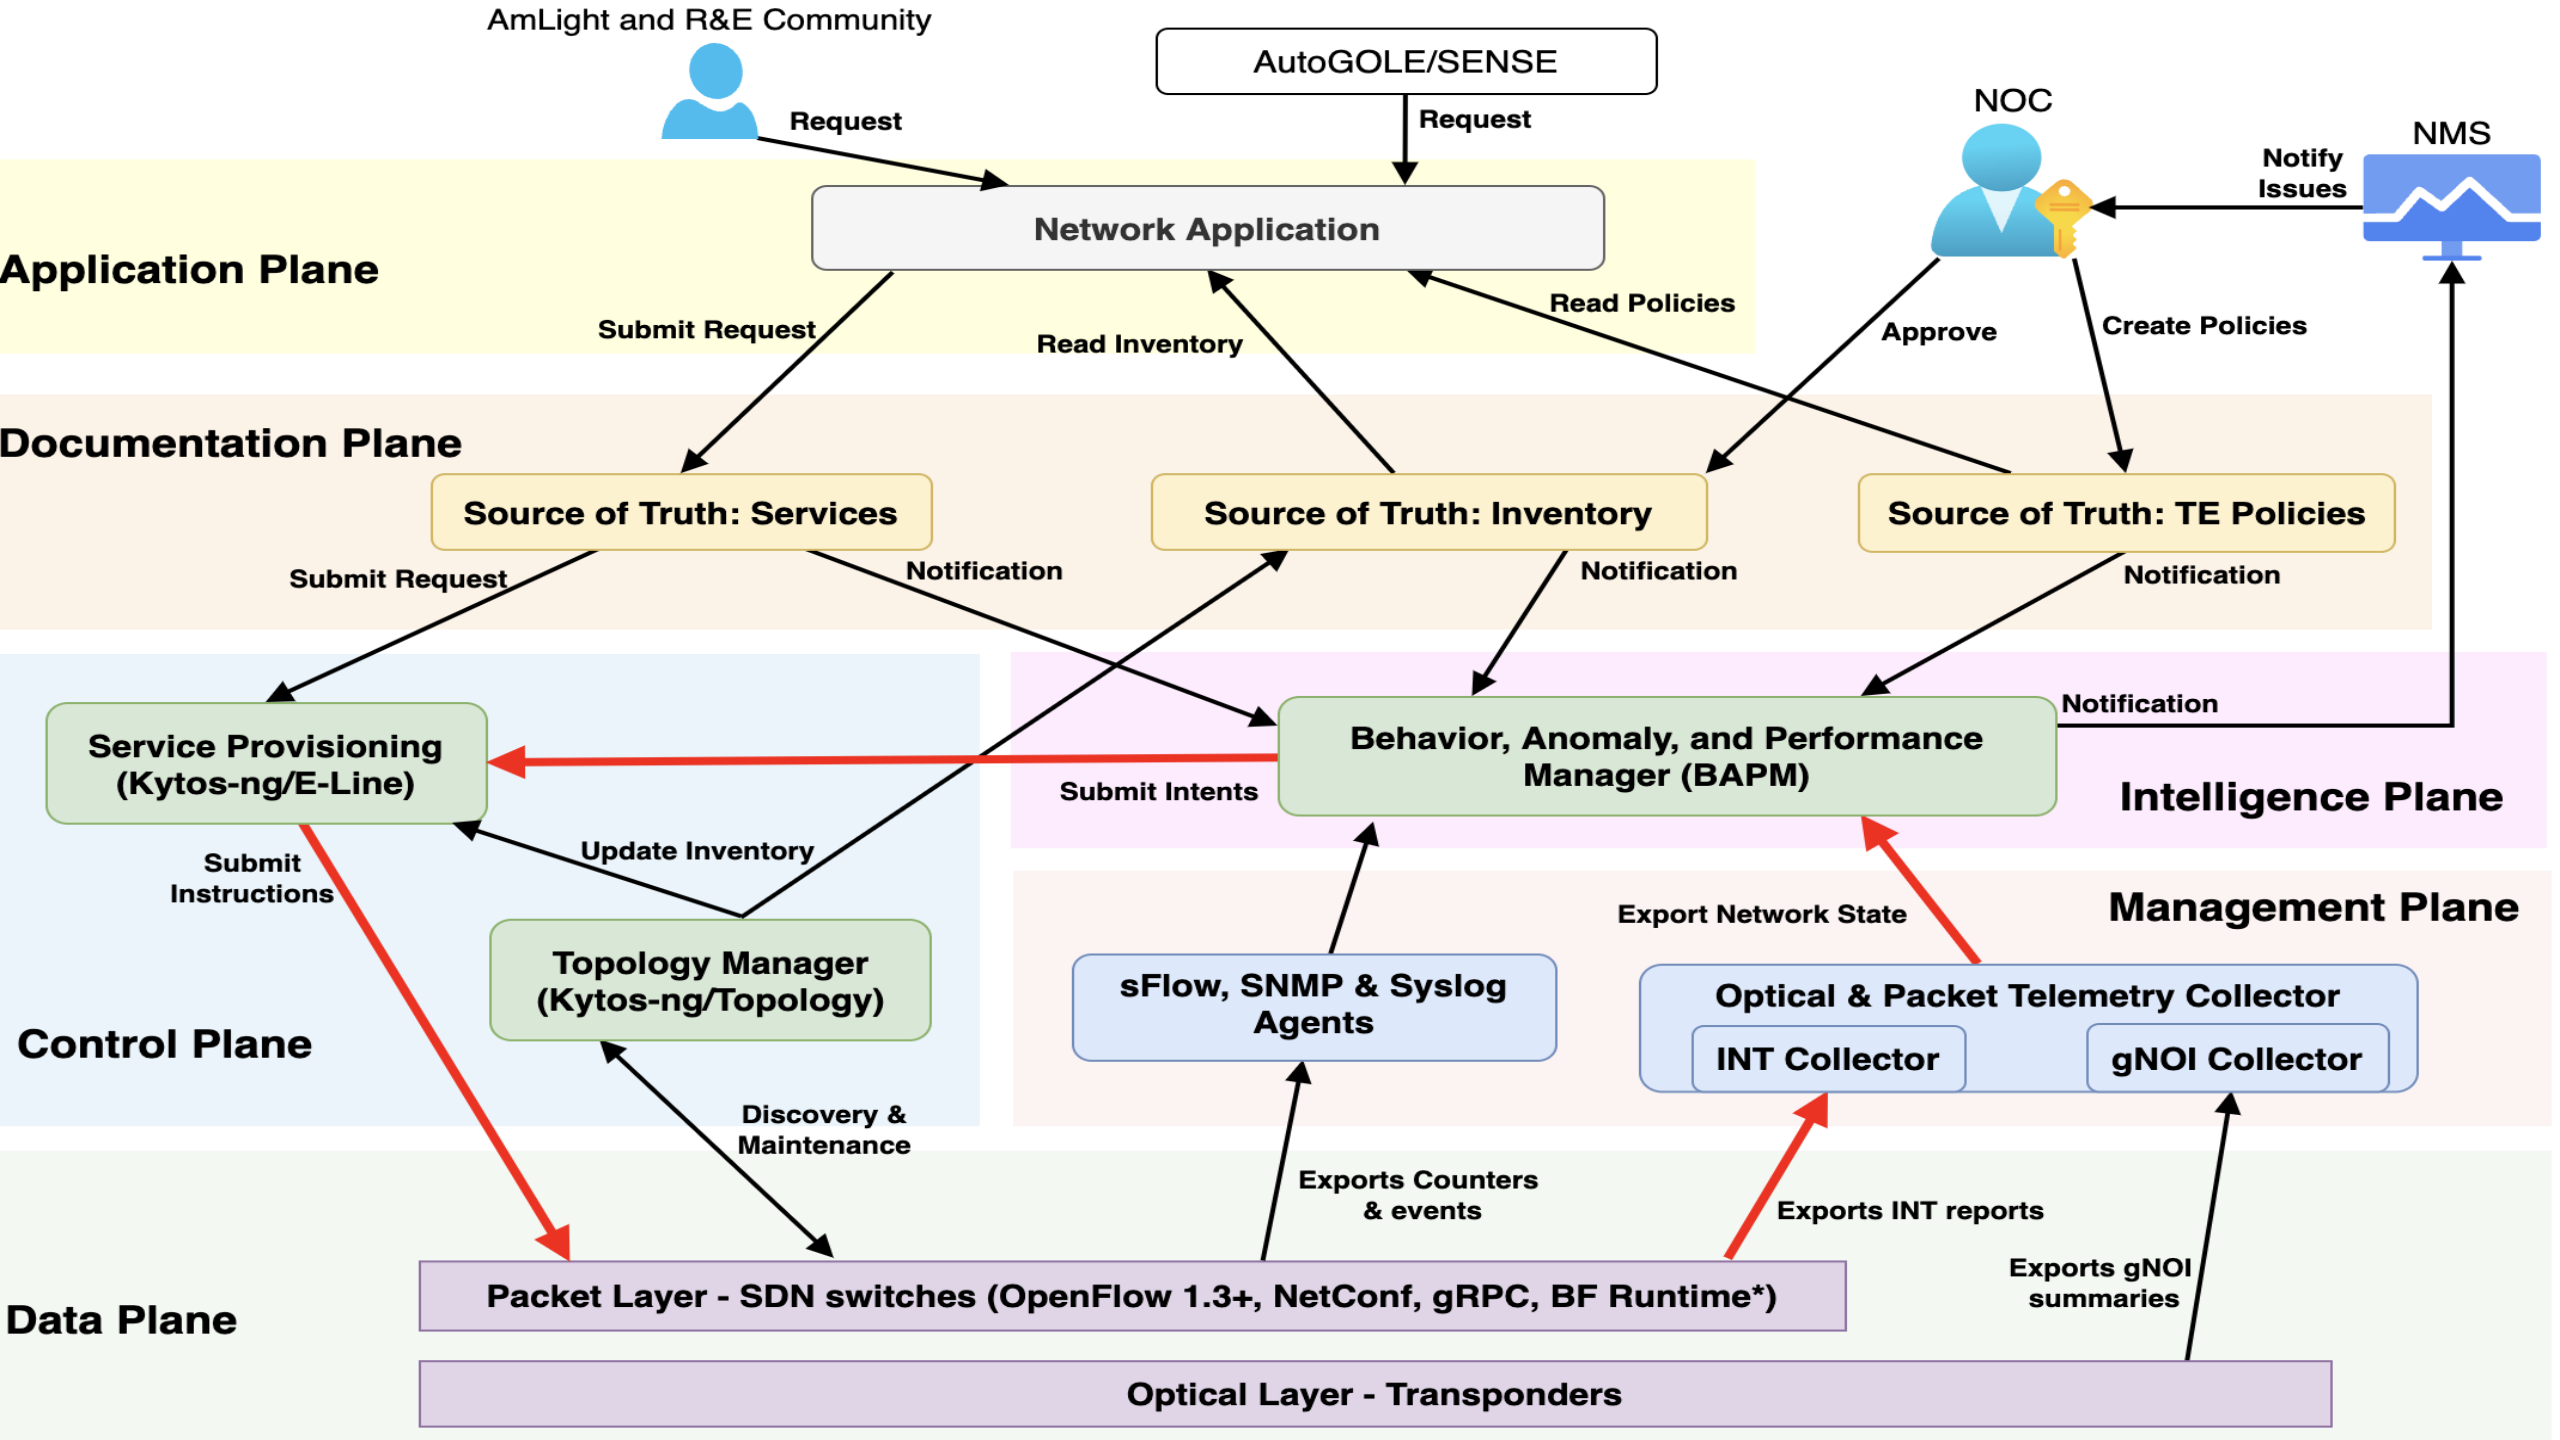
AmLight and R&E Community
AutoGOLE/SENSE
Request
Request
NOC
NMS
Notify Issues
Application Plane
Network Application
Submit Request
Read Inventory
Read Policies
Approve
Create Policies
Documentation Plane
Source of Truth: Services
Source of Truth: Inventory
Source of Truth: TE Policies
Submit Request
Notification
Notification
Notification
Service Provisioning (Kytos-ng/E-Line)
Behavior, Anomaly, and Performance Manager (BAPM)
Notification
Submit Intents
Intelligence Plane
Submit Instructions
Update Inventory
Management Plane
Export Network State
Topology Manager (Kytos-ng/Topology)
sFlow, SNMP & Syslog Agents
Optical & Packet Telemetry Collector
INT Collector
gNOI Collector
Control Plane
Discovery & Maintenance
Exports Counters & events
Exports INT reports
Exports gNOI summaries
Data Plane
Packet Layer - SDN switches (OpenFlow 1.3+, NetConf, gRPC, BF Runtime*)
Optical Layer - Transponders

Control Plane: Kytos-ng SDN Platform

- Major transformation compared to the previous AmLight-ExP project
- Brand new SDN controller built from scratch
  - https://github.com/kytos-ng
- Built to be fully compatible with the new AmLight Data Plane
- Addresses the requirements of the AmLight community and science drivers:
  - Telemetry and pathfinding options

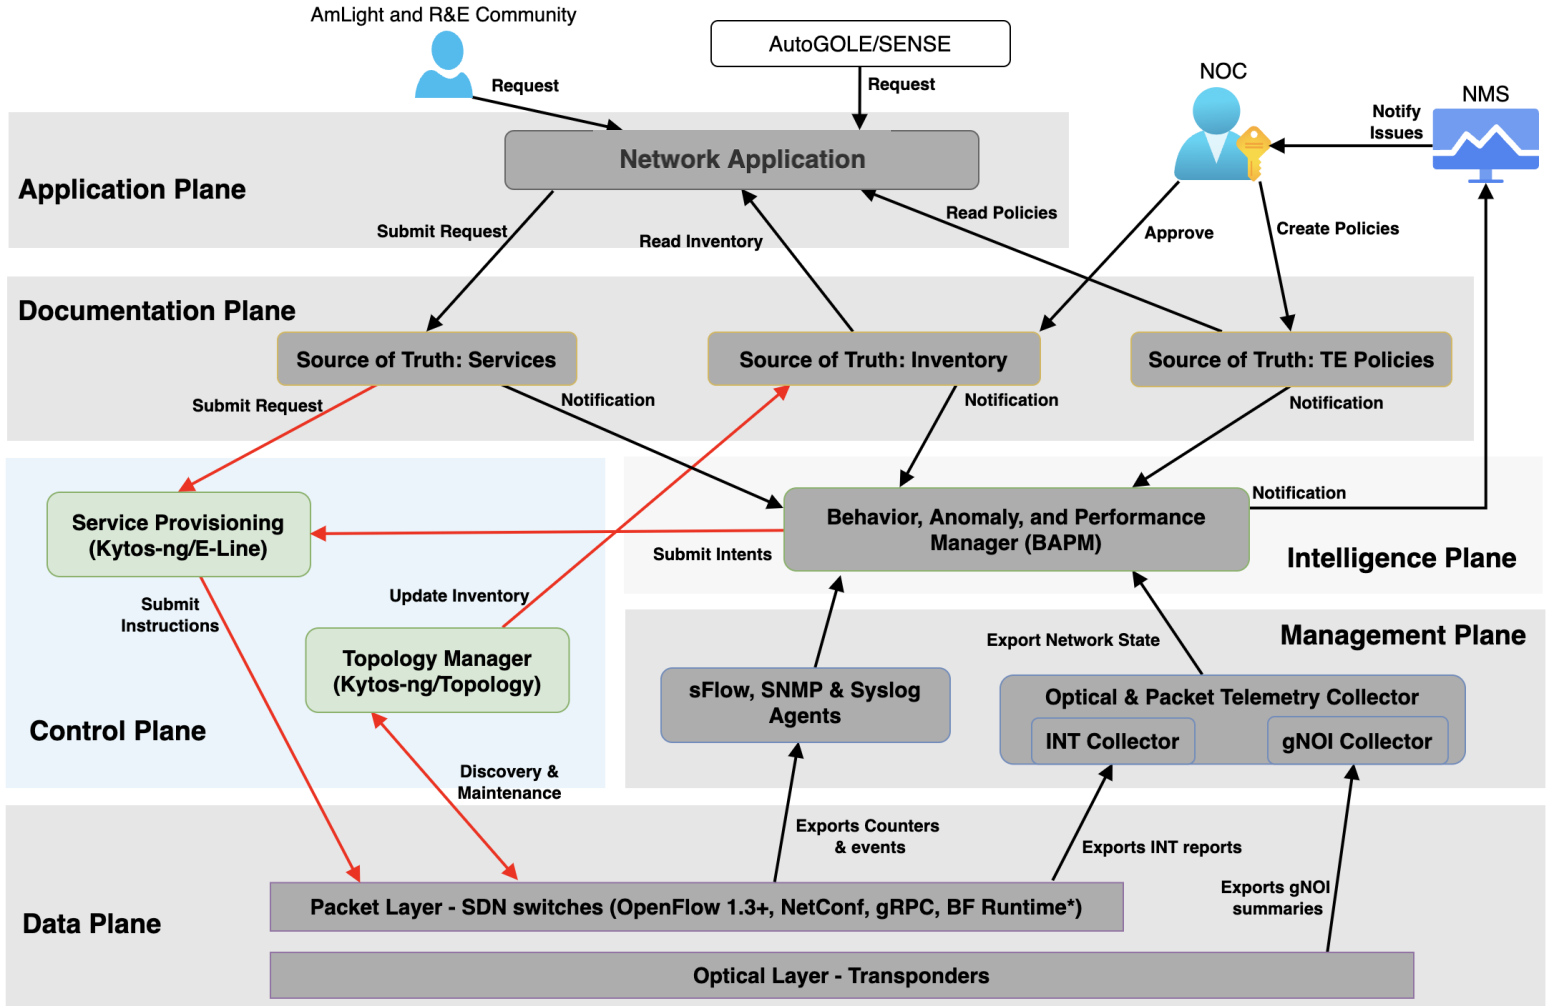

AmLight ExP
Americas Lightpaths Express & Protect

Data Plane – Forwarding and Telemetry

- Second major transformation compared to the previous AmLight-ExP project
- Programmable P4 switches
- Support for 100Gbps and 400Gbps
- Support for OpenFlow 1.3+1.4 & BFRuntime
- Supports In-band Network Telemetry (INT) for per-packet telemetry

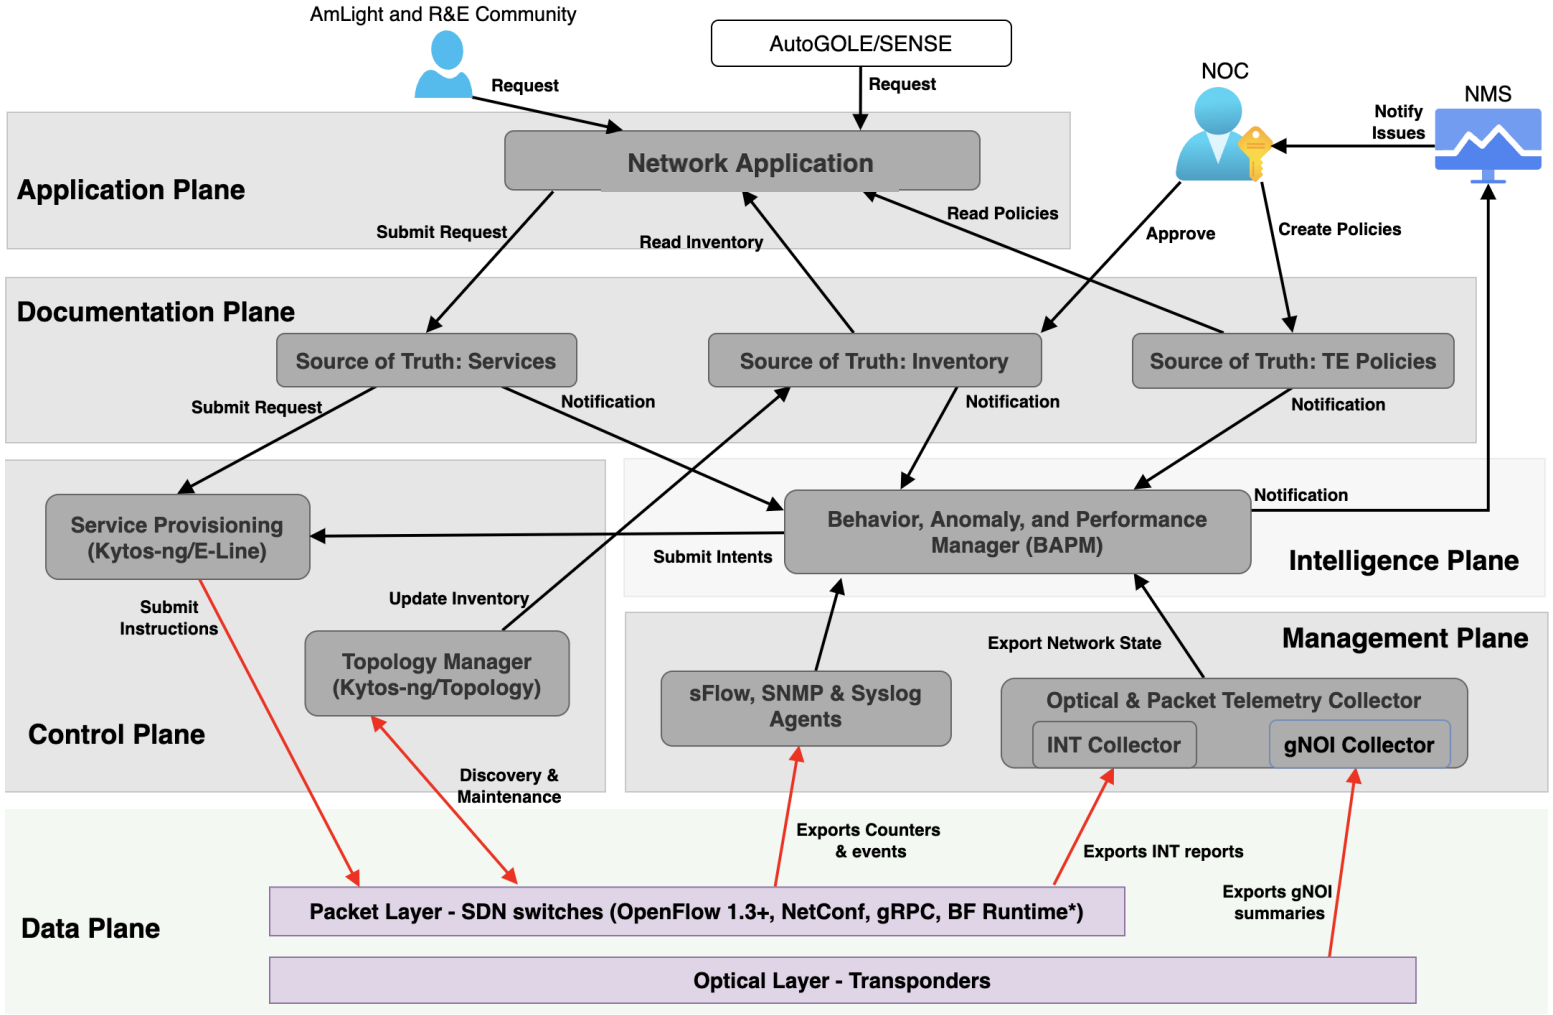

AmLight ExP
Americas Lightpaths Express & Protect

Management Plane

- In-band Network Telemetry (INT) exports telemetry reports per-second per-packet
  - Instantaneous bandwidth utilization
  - Instantaneous queue/buffer occupancy
  - Instantaneous hop and flow delays
  - Per-packet path trace
  - Microburst-detection
- Juniper Telemetry Interface (JTI) exports telemetry reports from Juniper MX204 routers:
  - Every 2 seconds for interface counters
  - Every 1 second for device's sensors

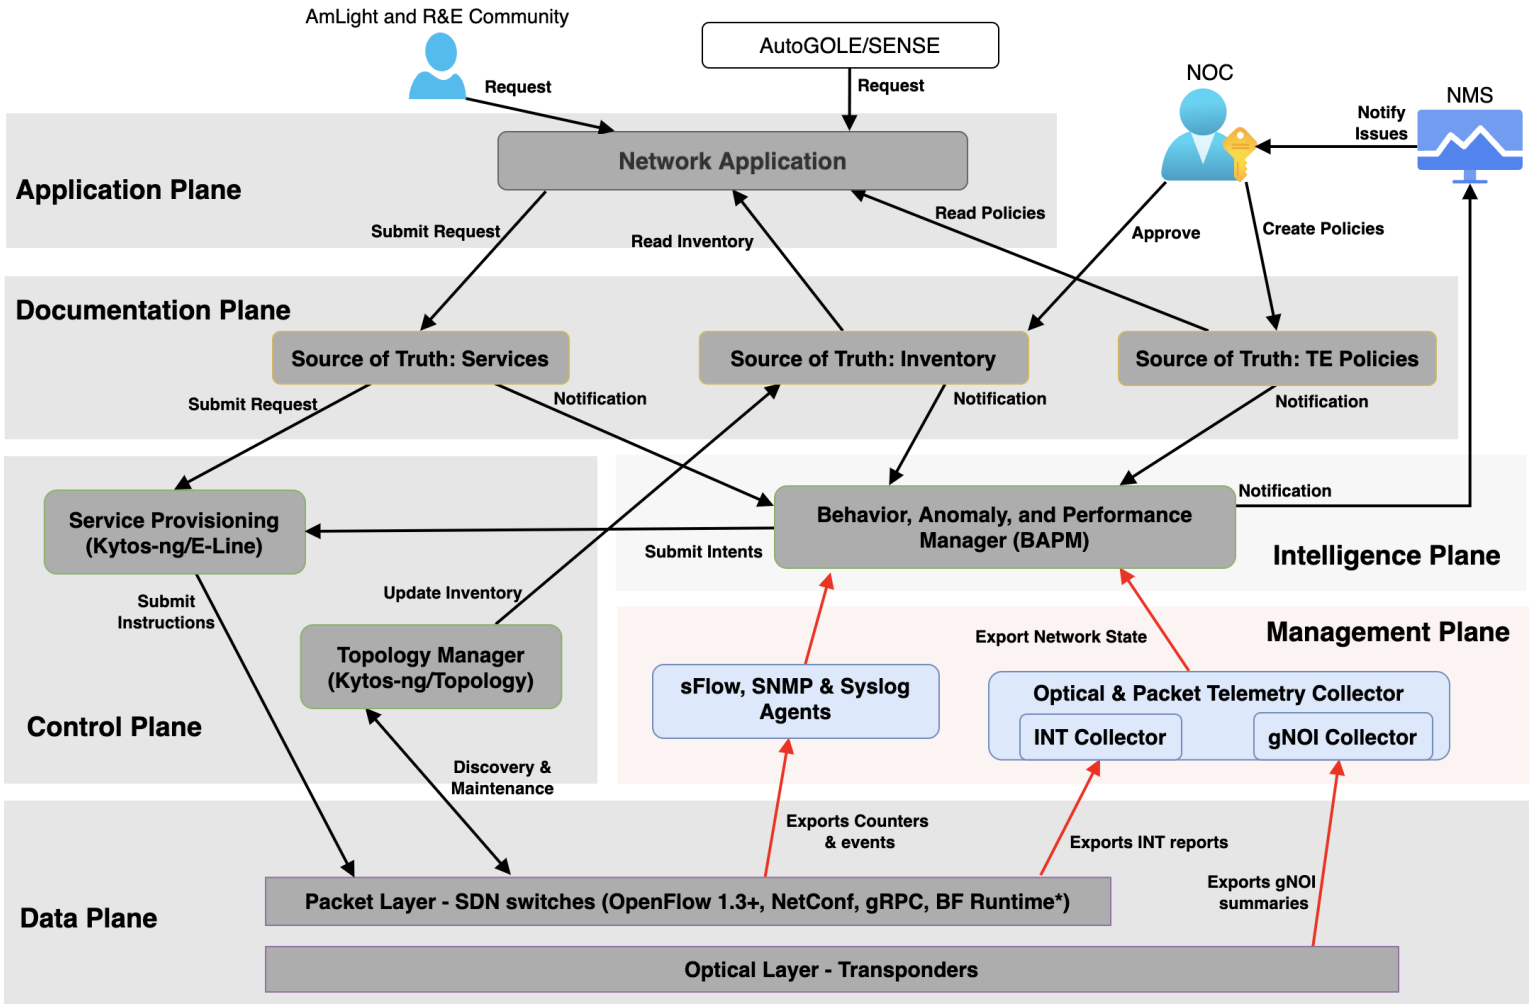

Management Plane: Network Utilization

- 5 data transfers/bursts of 40-50Gbps for 5 seconds.
- Top: INT metadata exported in real time, per packet
- Bottom: SNMP get running as fast as supported by the switch: 14 seconds.

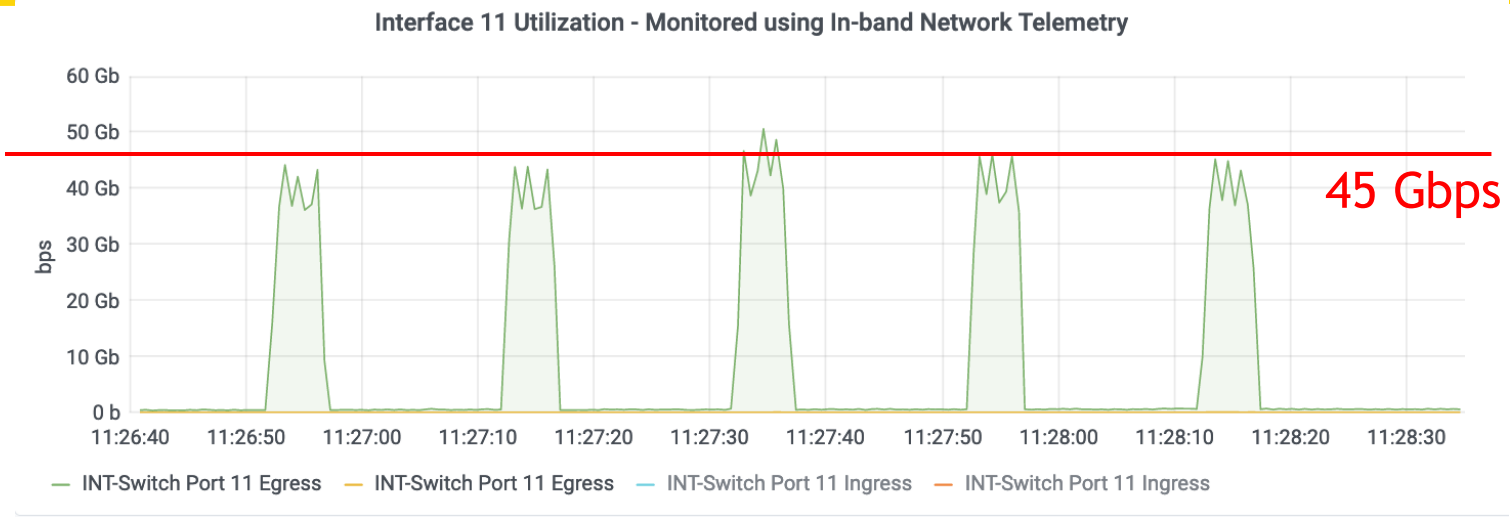

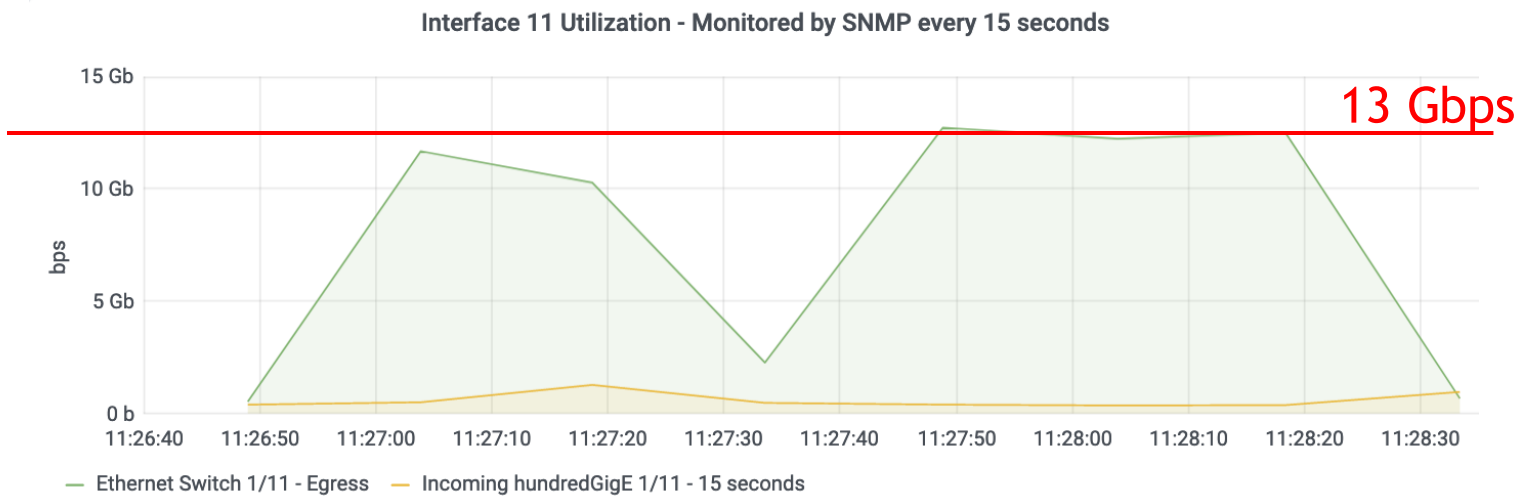

AmLight ExP
Americas Lightpaths Express & Protect

Intelligence Plane: Network Optimization

1. *Gets inventory, policies, and services from the Documentation Plane*
2. *Gets telemetry reports from the Management Plane*
3. *Profiles AmLight's traffic every 100-500ms*
   - *Discovers performance issues and traffic anomalies*
4. *Makes suggestions to the Control Plane*
   - *Steer traffic, Load balance services, rate-limit, etc.*

Creates a closed-loop for network optimization

- Goal is sub-second reaction and debugging
- Behavior, Anomaly, and Performance Manager (BAPM) is the component responsible for the intelligence

Example of policies:

- 80+% BW utilization >= 2s
- 50+% [Queue Occupancy] >= 2s
- Number of path changes >= 5 in 2h

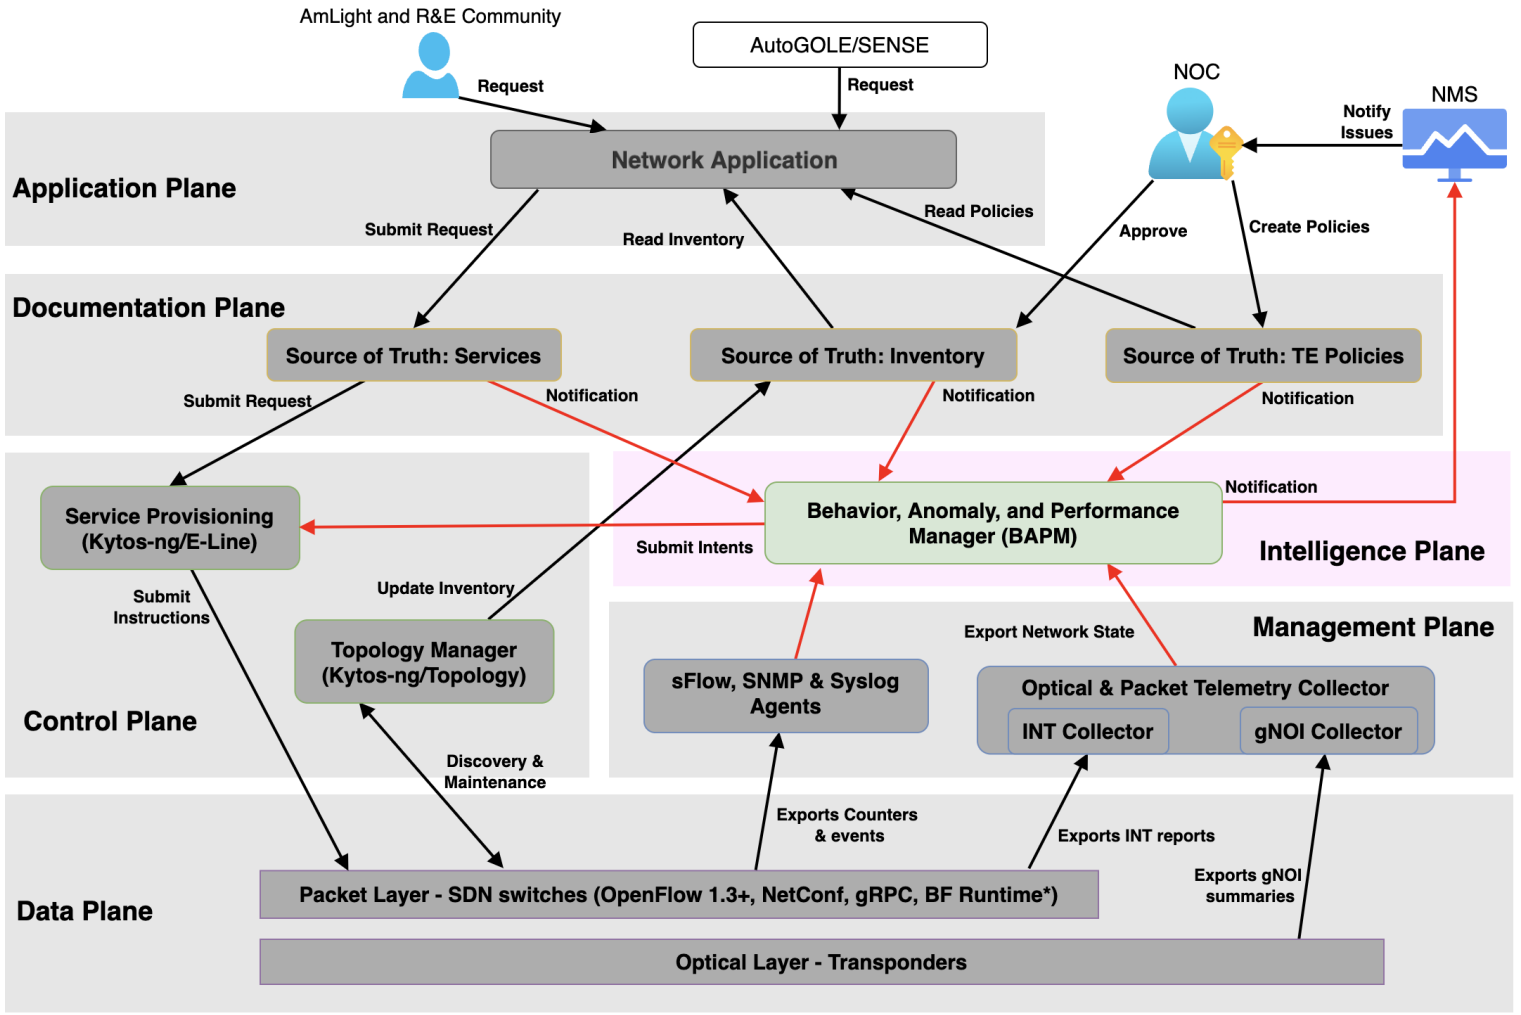

AmLight ExP
Americas Lightpaths Express & Protect

What comes next?

Future

- BERToD 1.0: An automated framework for frame loss detection (abstract accepted at the 2024 Internet2 Technology Exchange)
- Microburst detection (paper presented at the 2023 NOMS)
- DDoS identification (paper accepted at INDIS 2024) and mitigation on premises
- Supporting for in-network computing: How to enable AI/ML research at OXPs (reconfiguring DPUs and GPUs on demand)
- BERToD 2.0: Leveraging IA/ML to isolate faults
- AmLight 3.0 (2026-2030): Turning AmLight into a Lossless Network: Building a stateful network using FPGAs for packet duplication and deduplication

BERToD - Bit Error Rate Test on Demand

- An Automated frame loss detection framework
  - Leverages network state from BAPM (for instance, link utilization) to decide the max bandwidth to use
  - Leverages topology from Kytos-ng
  - Leverages policies and documentation from Netbox
  - Leverages physical loops in the programmable switches
  - Uses the EXFO NetBlazer packet generator to test
  - Uses Grafana to display the results
- Tests every link and path possible to isolate eventual frame loss
- Highly accurate and granular results

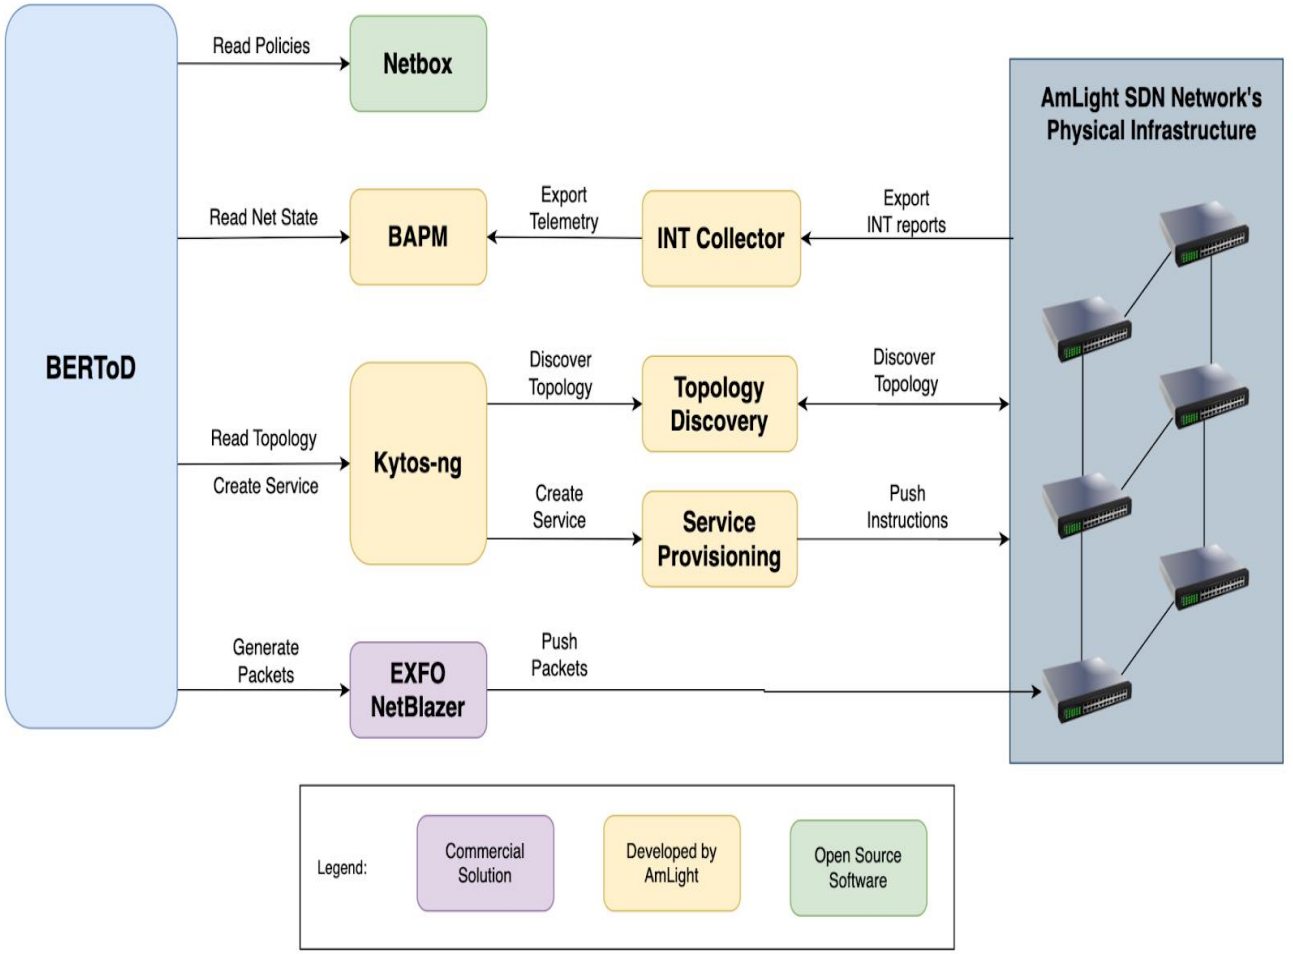

BERToD - Bit Error Rate Testing on Demand

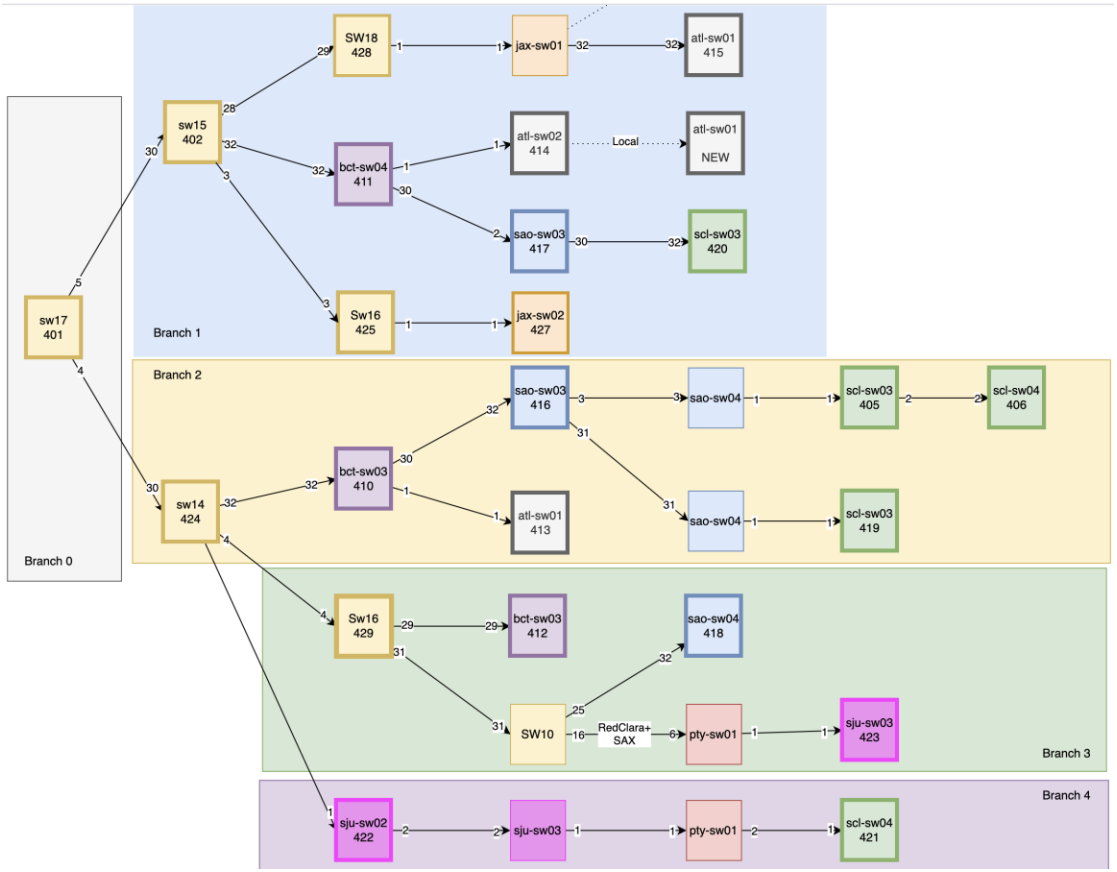

% of Frame Loss Grouped by Day

	08/28			08/30			09/01			09/03			09/05			09/07			09/09		
VLAN-401	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
VLAN-402	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
VLAN-405	0.0000	0.0000	0%	0.0168	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
VLAN-406	0.162%	0.0004	0.0003	0.0630	0.0000	0.0000	0.0000	0.0000	0%	0.0000	0%	0.0000	0.0000	0.0000							
VLAN-410	0%	0.0009	0.0041	0.0103	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
VLAN-411	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0.0000	0%	0%	0%	0%	0%	0%	0%	0%	0%
VLAN-412	0.0023	0.0073	0.0267	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
VLAN-413	0.0078	0%	0.0584	0.0211	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
VLAN-414	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
VLAN-415	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
VLAN-416	0%	0%	0.0442	0.0027	0.0002	0.0013	0.0159	0%	0%	0.0365	0.0427	0.0205	0%	0.0241	0.0637						
VLAN-417	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0.0000	0%	0%	0%	0%	0%	0%	0%	0%	0%
VLAN-418	0.0000	0.0735	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
VLAN-419	0.0262	0.0007	0.0020	0.0878	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
VLAN-420	0%	0%	0%	0%	0%	0.0141	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
VLAN-421	0.0000	0.0577	0.0021	0.0398	0%	0%	0%	0%	0%	0%	0.0000	0.0000	0.0000	0.0000							
VLAN-422	0.0095	0.0078	0.0103	0.0685	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
VLAN-424	0.0416	0.0148	0.0014	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
VLAN-425	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
VLAN-427	0.0042	0.0006	0.0002	0.0000	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
VLAN-428	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
VLAN-429	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%

< 0.00000100% | 0.00000100%+ | 0.000100%+ | 0.00100%+ | 0.0100%+ | 0.100%+

AmLight ExP
Americas Lightpaths Express & Protect

Thanks! Questions?