AmLight ExP
Americas Lightpaths Express & Protect

CI Lunch and Learn / Mar 18th, 2022

The 2021-2025 NSF IRNC AmLight-ExP project: Moving towards an Autonomic Network Architecture (ANA)

Jeronimo Bezerra <jab@amlight.net>

Outline

- What is AmLight?
- Network Connectivity
- Introduction to Autonomic Network Architecture (ANA)
- New AmLight SDN Architecture: Plane by Plane
- ANA at AmLight: a use case
- Conclusion

What is AmLight?

- A distributed academic exchange point built to enable collaboration among Latin America, Africa, and the U.S.
- Supported by NSF, OAC, and 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, built upon layers of trust and openness by sharing:
  - Infrastructure resources
  - 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…

Network Connectivity

- 600Gbps of upstream capacity between the U.S. and Latin America, and 100Gbps to Africa
  - +200Gbps from Sao Paulo to Boca Raton (2023)
- 2+ Tbps of international connectivity
  - AmLight Express (green), AmLight Protected (white), plus waves provided by RedClara and SANReN/TENET (pink)
- NAPs: Florida(3), Brazil(2), Chile, Puerto Rico, Panama, and South Africa
  - Soon: Atlanta, Georgia
- Infrastructure provided by collaborators with a centralized service provisioning solution

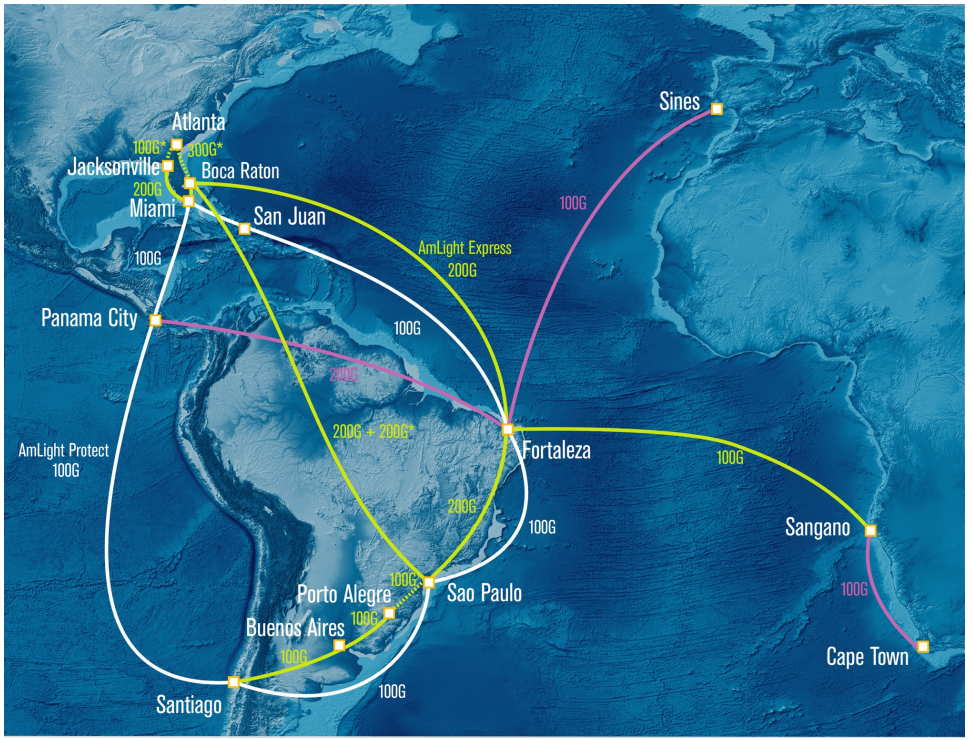

* and dashed lines represent on going efforts

AmLight ExP
Americas Lightpaths Express & Protect

AmLight: Collaboration at its finest!

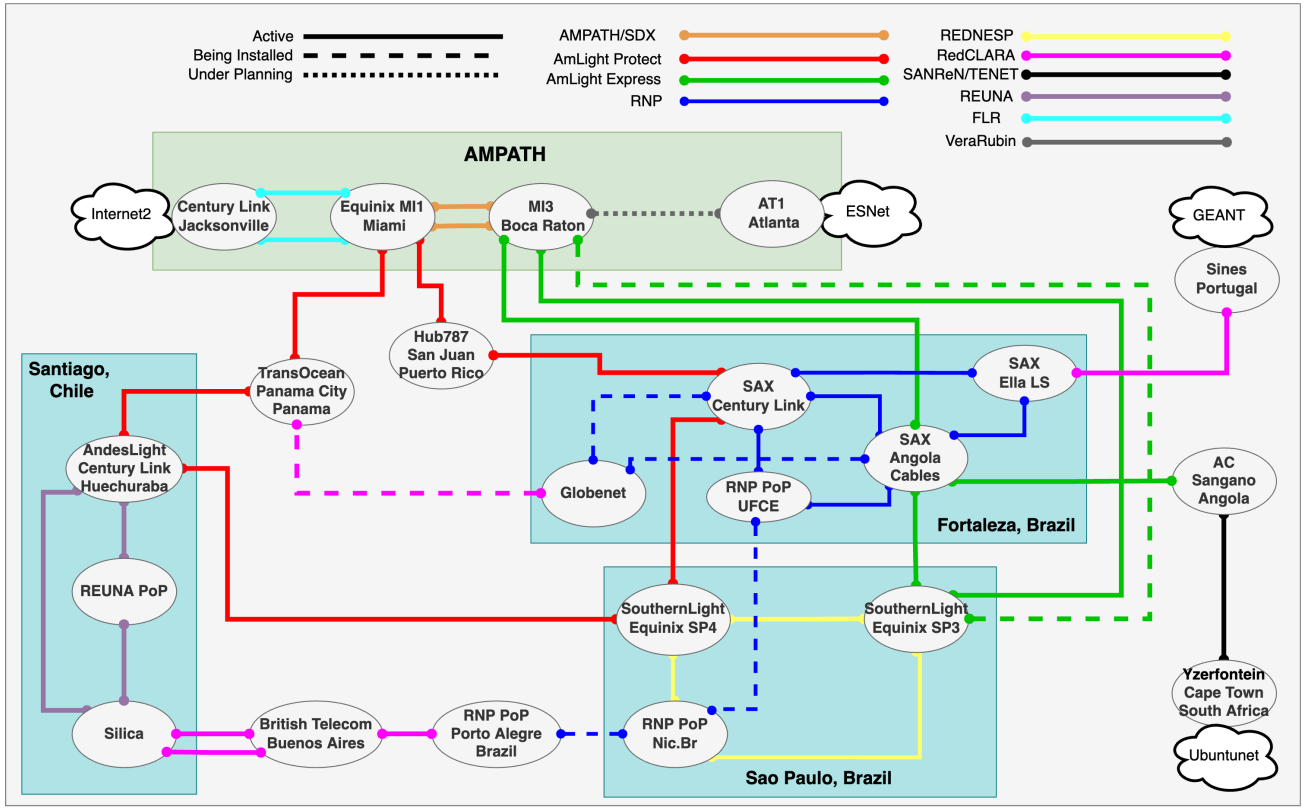

AmLight ExP
Americas Lightpaths Express & Protect

Autonomic Network Architecture (ANA)…

Autonomic Network Architecture (ANA)

- Autonomic Networking Architecture (ANA) creates the definitions and design goals for a self-managed network.
- Standardized via Internet Engineering Task Force (IETF) RFC 7575
- Autonomic systems were first described in a manifesto by IBM in 2001. The fundamental concept involves eliminating external systems from a system's control loops and closing of control loops within the autonomic system itself.
- *Self-management* is comprised of several "self-x" properties

Autonomic Network Architecture (ANA) [2]

- **Self-configuration**: Functions do not require configuration, by either an administrator or a management system. They configure themselves, based on self-knowledge.
- **Self-healing**: Autonomic functions adapt on their own to changes in the environment and heal problems automatically.
- **Self-optimizing**: Autonomic functions automatically determine ways to optimize their behavior against a set of well-defined goals.
- **Self-protection**: Autonomic functions automatically secure themselves against potential attacks.

Autonomic Network Architecture (ANA) [3]

	Automatic	Automation	Closed-Loop Orchestration	Autonomic
Description	User runs a script to change a service or configuration	User runs a "playbook" to change multiple services and configuration of multiple nodes at the same time	Application changes multiple services and configuration of multiple nodes. Nodes export state and counters. Application reacts to the new state and performs or not new changes in a forever loop.	Application discoveries assets, policies, and intents. Configure devices from scratch based on policies and intents. Minimal to no-user interaction. Resolution of conflicts defined by administrators
User-Input	Scripts, inputs, topology, destination	Scripts, inputs, inventory	Scripts, inputs, inventories, policies/conditions/triggers	Policies and intents

Autonomic Network Architecture (ANA) [3]

	Automatic	Automation	Closed-Loop Orchestration	Autonomic
Description	User runs a script to change a service or configuration	User runs a "playbook" to change multiple services and configuration of multiple nodes at the same time	Application changes multiple services and configuration of multiple nodes. Nodes export state and counters. Application reacts to the new state and performs or not new changes in a forever loop.	Application discoveries assets, policies, and intents. Configure devices from scratch based on policies and intents. Minimal to no-user interaction. Resolution of conflicts defined by administrators
User-Input	Scripts, inputs, topology, destination	Scripts, inputs, inventory	Scripts, inputs, inventories, policies/conditions/triggers	Policies and intents

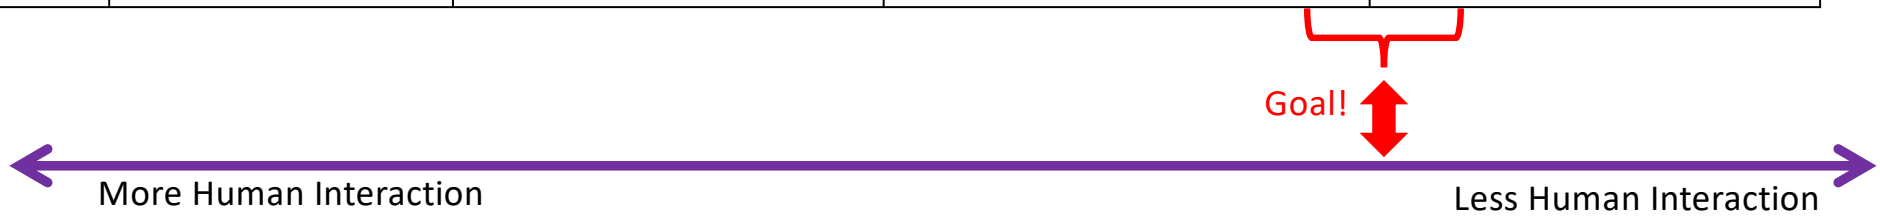

AmLight SDN Architecture…

AmLight SDN Architecture – 2015 - 2020

- From 2014-2020, we followed the "basic" SDN layers as in [1]
  - Application, Management, and Control Planes were very coupled under the same controller/orchestrator:
    - OESS, ONOS, Kytos, and Ryu
    - Each SDN Plane was operated as modules of the SDN controllers' software stack
  - Data Plane was a blend of vendors (Dell, Corsa, Brocade) with mixed support
  - Southbound Interface was OpenFlow 1.0 and OpenFlow 1.3

- From the ANA perspective, only *self-healing* was in place
  - Fiber cuts and device outages were handled by finding backup paths.

- Source of Truth was the SDN controllers' databases.
  - Extra service data had to be documented in an external repository, for instance, contacts in case of issues with the service.

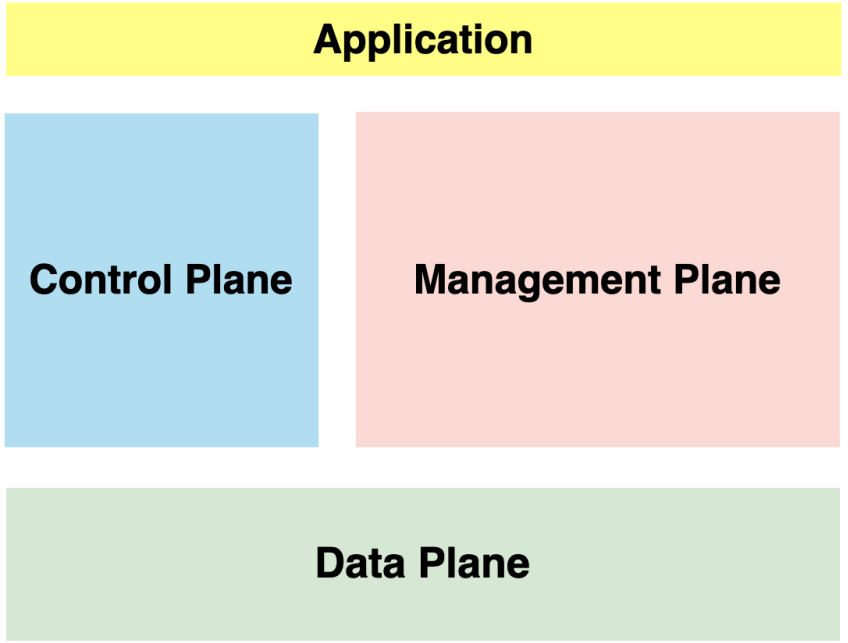

[1] IETF RFC 7626 & RFC 8597

AmLight ExP
Americas Lightpaths Express & Protect

AmLight SDN Architecture – 2021-2025

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

- Documentation Plane (Source of Truth) decoupled from the Control Plane:
  - NetBox for Source of Truth and Kytos-ng for Control Plane
- Management Plane (utilization counters) decoupled from the SDN Controller:
  - Instead of SNMP and OpenFlow counters, AmLight will use In-band Network Telemetry for counters
- With the Documentation Plane and network telemetry, a new plane was created to add intelligence to our network:
  - Intelligence Plane will focus on learning the network state and create a closed-loop control for self-optimization
  - Intelligence Plane will use In-band Network Telemetry to audit how the Control Plane implements the user requests
- Application Plane will be decoupled from the SDN Controller
  - To support self-configuration, the Application Plane will interface only the Documentation Plane for services, inventory, and policies. Control Plane will be notified of user requests to act

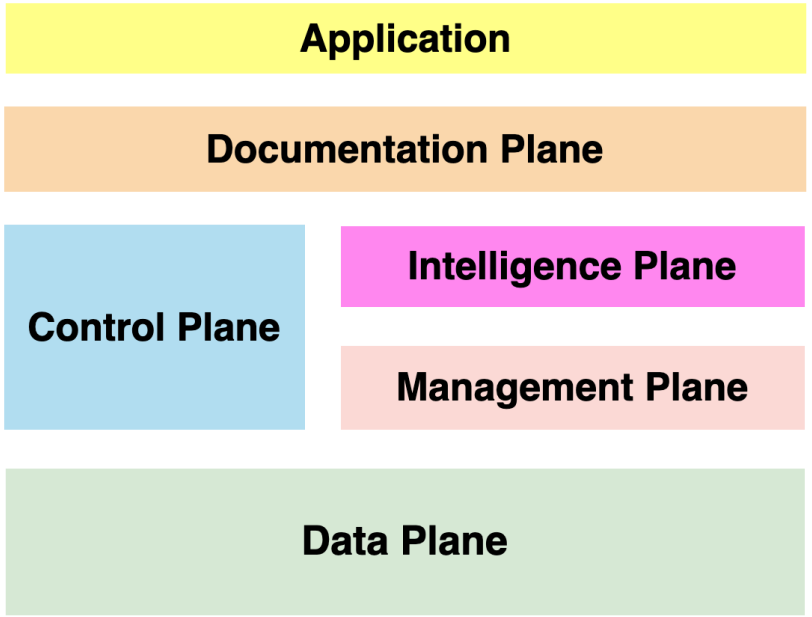

AmLight ExP
Americas Lightpaths Express & Protect

AmLight SDN Architecture – Big Picture

- The Autonomic Network Architecture specification enables partial deployment and support:
  - AmLight won't try to support all ANA functions and requirements
  - Focus will be on a few Autonomic Functions (AF)
- The first Autonomic Function planned is supporting L2VPNs fully managed by this architecture:
  - Administrators will get involved just in case of conflicts and when new policies are needed
- Users will have access to all details related to their services:
  - Topology, counters, telemetry, policies, number of events, and maintenance windows

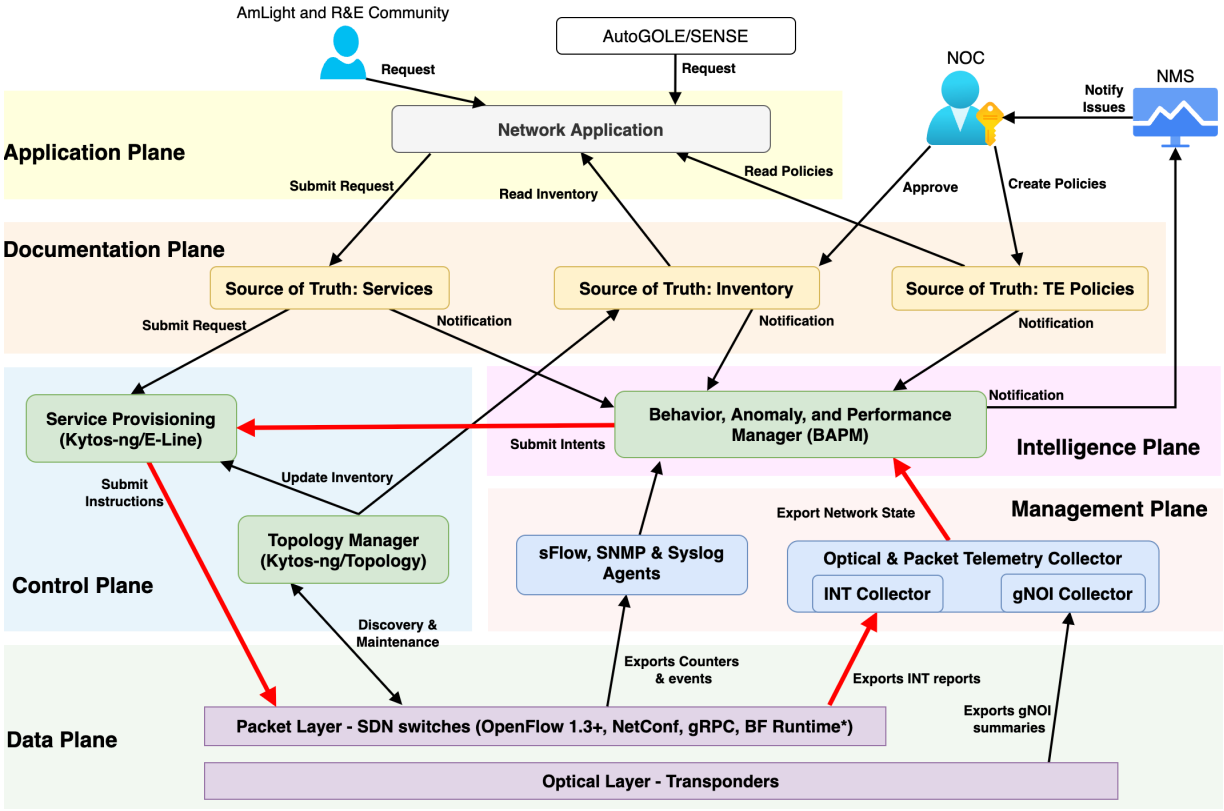

AmLight ExP
Americas Lightpaths Express & Protect

Application Plane

- Application Plane interacts with the Documentation Plane to:
  - Gather the traffic engineering and per-user policies
  - Understand the topology and resources available
  - Submit service requests to be stored
- Project Goals for all network applications:
  - Enhance network and topology visibility
  - Add per-user path constraints/restrictions
  - Support per-packet network telemetry
  - Support QoS policies + bandwidth reservation
- Network Applications:
  - L2VPN + AutoGOLE/NSI + SENSE
  - L3VPN or VRF
  - Cloud Services/DTN
  - IPv4/v6 Transit
  - DDoS Mitigation*

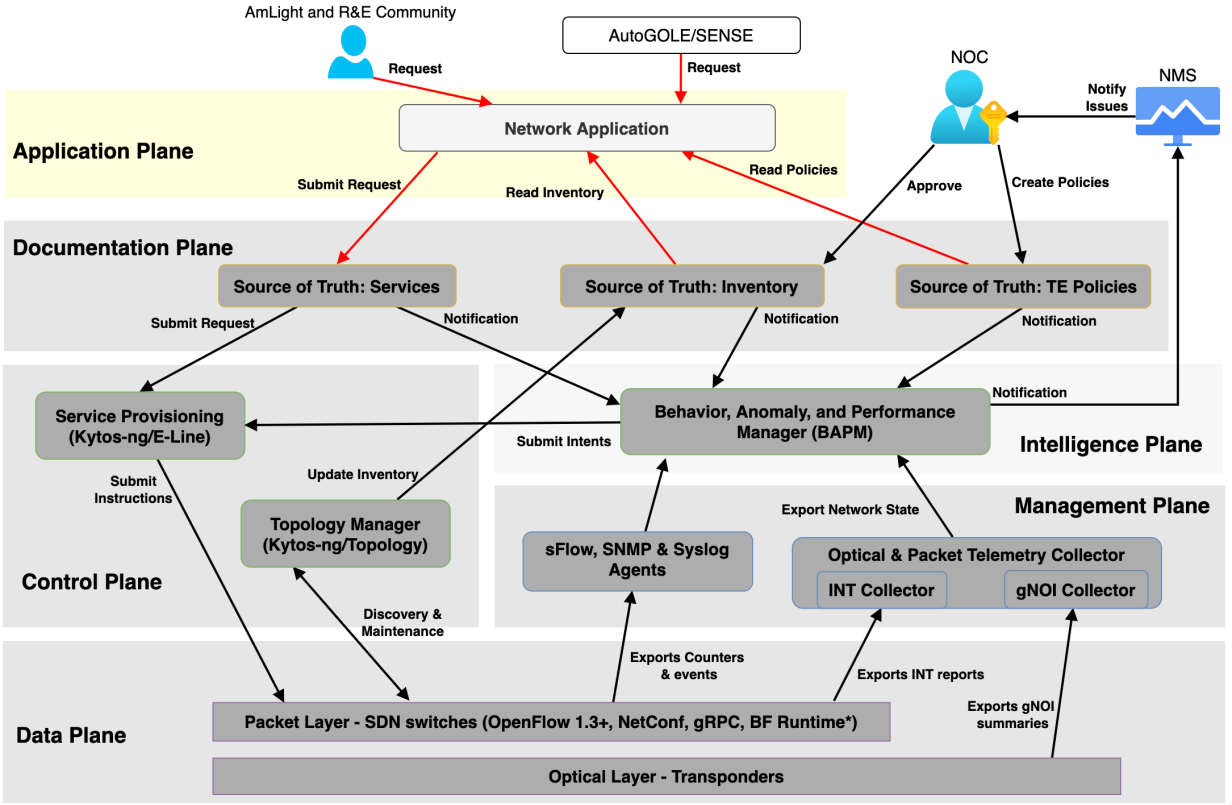

AmLight ExP
Americas Lightpaths Express & Protect

Documentation Plane

- New SDN Plane for AmLight
- To support ANA + self-configuration, three new source of truth will be developed
  - Services
  - Inventory
  - Traffic Engineering policies
- Data models to be hosted via NetBox plugins
- Netbox will trigger consumers to retrieve updates
- NOC's responsibilities:
  - Create TE policies
  - Create inventory items
  - Approve discovered inventory items

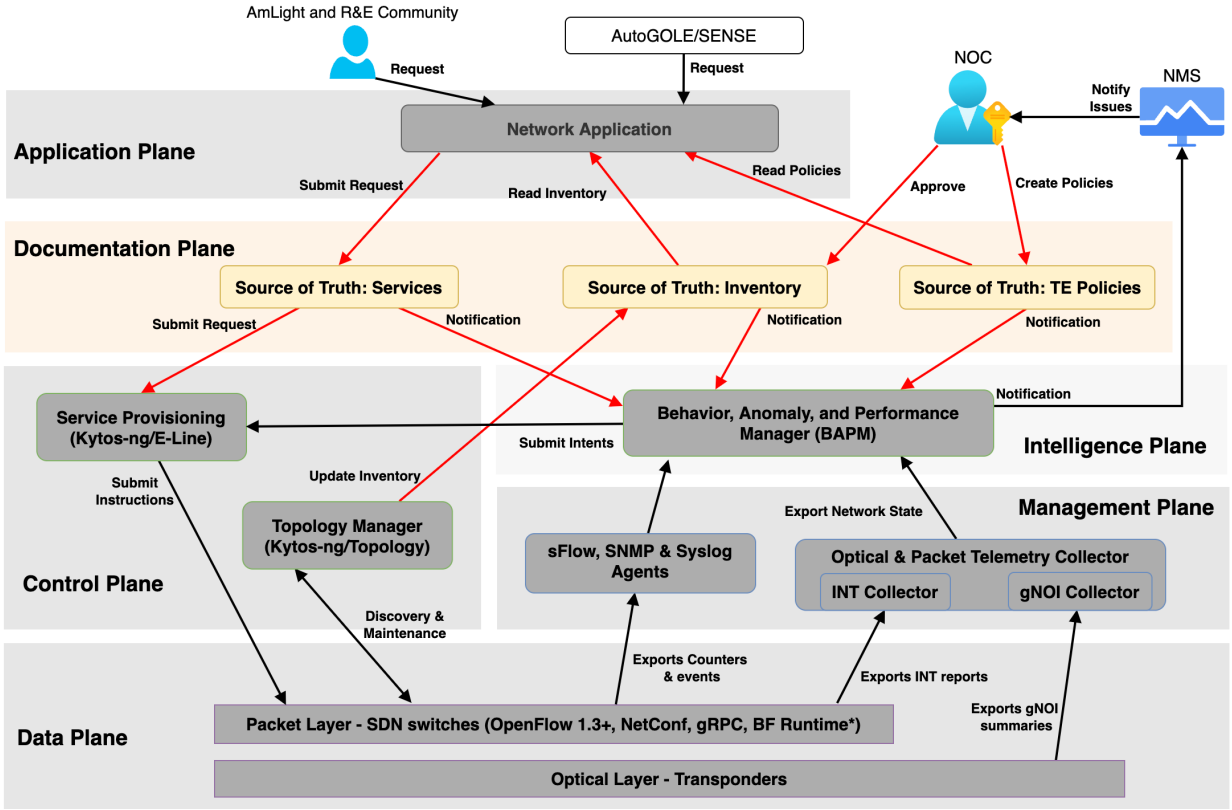

AmLight ExP
Americas Lightpaths Express & Protect

Control Plane

- Triggered by the Documentation Plane for ANA/self-configuration
- Major transformation compared to the previous AmLight-ExP project
  - From multiple open-source solutions (ONOS, OESS, FloodLight, and Ryu) to a new solution (Kytos-ng)
- Brand new SDN controller built from scratch
- 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

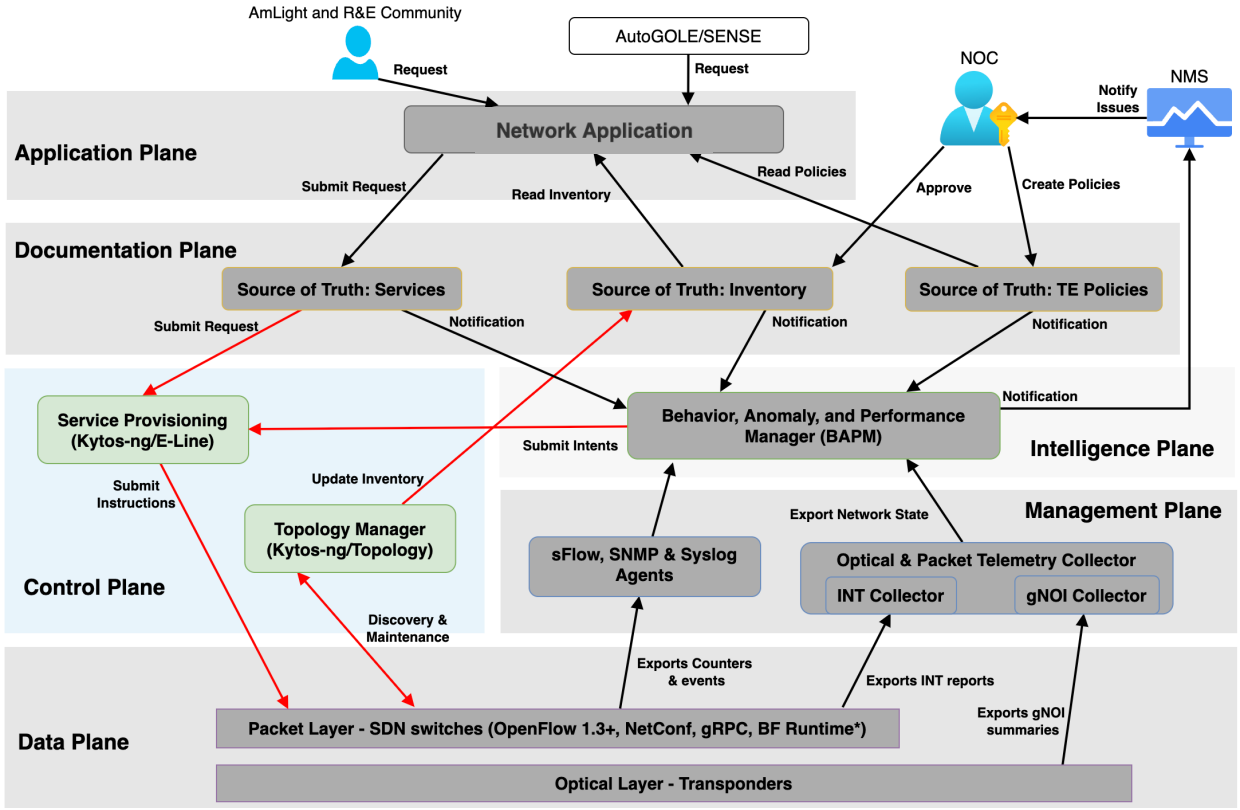

AmLight ExP
Americas Lightpaths Express & Protect

Control Plane – New Contoller: Kytos-ng

- Development focused on the AmLight operation requirements:
  - Simple REST API following OpenAPI 2.0 specs
  - Web UI with no more than 3 clicks for any action
  - Pathfinder with support for multiple metrics and restrictions:
    - # of hops, minimum delay, max bandwidth, ownership, reliability, priority, average bandwidth utilization
  - Supports OpenFlow 1.3+ and gRPC (to support BFRuntime or P4Runtime in the future)
  - Fully compatible with NoviFlow OpenFlow experimenter actions
  - Supports Ethernet Private Line (EPL) and Ethernet Virtual Private Line (EVPL)
  - Integration with Zabbix NMS for monitoring of resources
  - Simple and fast development of new applications
  - Roadmap for 2022:
    - Integration with In-band Network Telemetry to add per-packet telemetry
    - Support per-flow Bi-directional Forwarding Detection (BFD)
    - Support for VLAN range
- First Kytos-ng release was launched on Feb 15th.

*Kytos-ng is maintained by **FIU** and **rednesp** since May 2021*
https://github.com/kytos-ng

Data Plane

- *Gets instructions from the Control Plane and exports telemetry to the Management Plane*
- Second major transformation compared to the previous AmLight-ExP project
- Packet Layer:
  - Legacy switches replaced by NoviFlow switches (Tofino ASIC)
  - 32x100G ports per switch
  - Virtual Machine embedded
  - OpenFlow 1.3+1.4 & P4Runtime for southbound
  - INT for per-packet telemetry
  - Legacy routers replaced by Juniper MX204 routers
- Optical Layer:
  - Ciena Waveserver Ai transponders
  - Point-to-point circuits with API license
  - Streaming telemetry-capable devices
  - REST API being used for monitoring
- $$$ saved w/ colocation, power, and maintenance savings

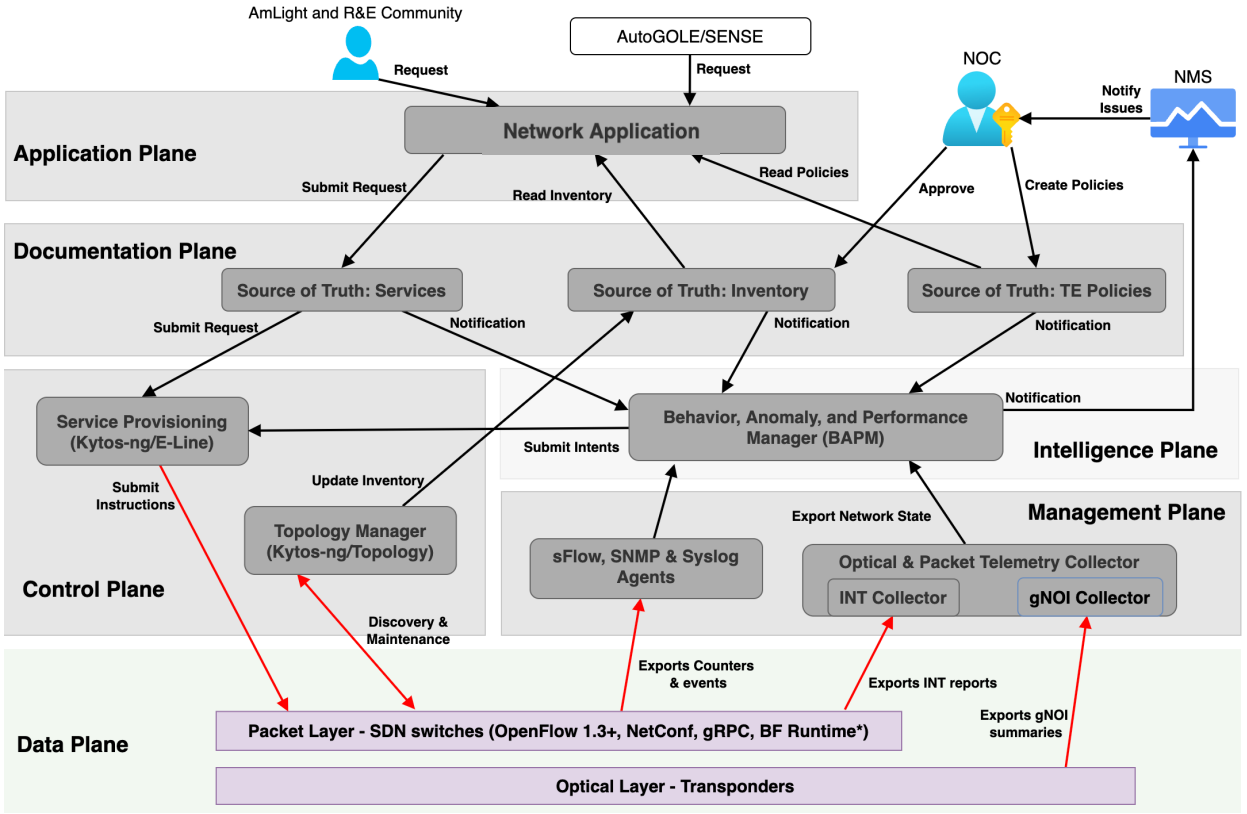

AmLight ExP
Americas Lightpaths Express & Protect

Management Plane

- *Gets telemetry data from the Data Plane and exports consolidated telemetry reports to the Intelligence Plane*
- Innovation of the AmLight-INT project: OPTC
- The Optic and Packet Telemetry Collector (OPTC) will collect, and process telemetry data exported by SDN switches and optical transponders
  - Juniper JTI added as we go.
  - *INT Collector solution presented on Feb 25th:*
    - *https://www.youtube.com/watch?v=CRnKKuP9I3Y*
- Telemetry data from Syslog, sFlow and SNMP will be used as well.
- OPTC will collect per packet delay, jitter, queue occupancy, top talkers, path taken, fiber attenuation, post-FEC, event logs, and counters and, based on thresholds defined by NOC, export reports to the Intelligence Plane

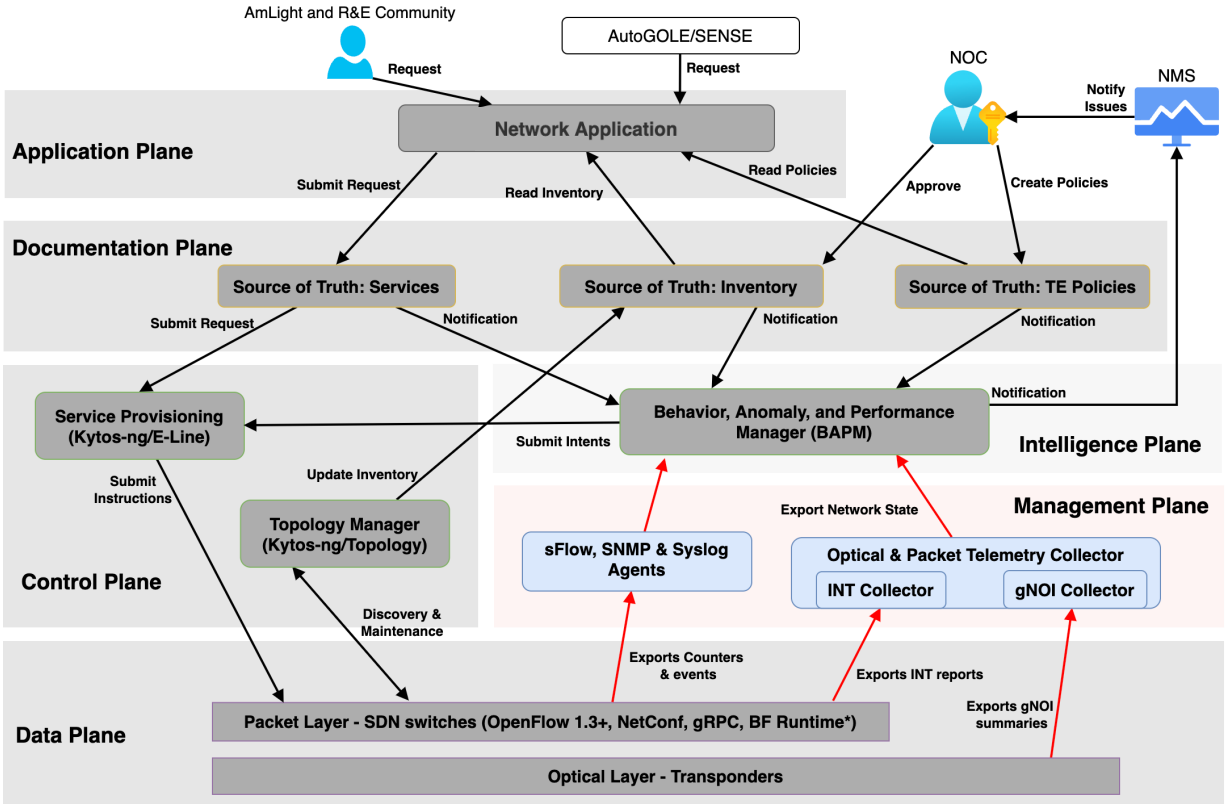

AmLight ExP
Americas Lightpaths Express & Protect

Intelligence Plane

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 anomalies,*

Change of mindset compared to the previous AmLight-ExP project

- Creates the closed-loop control for self-optimization
- Goal is to be prepared for 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

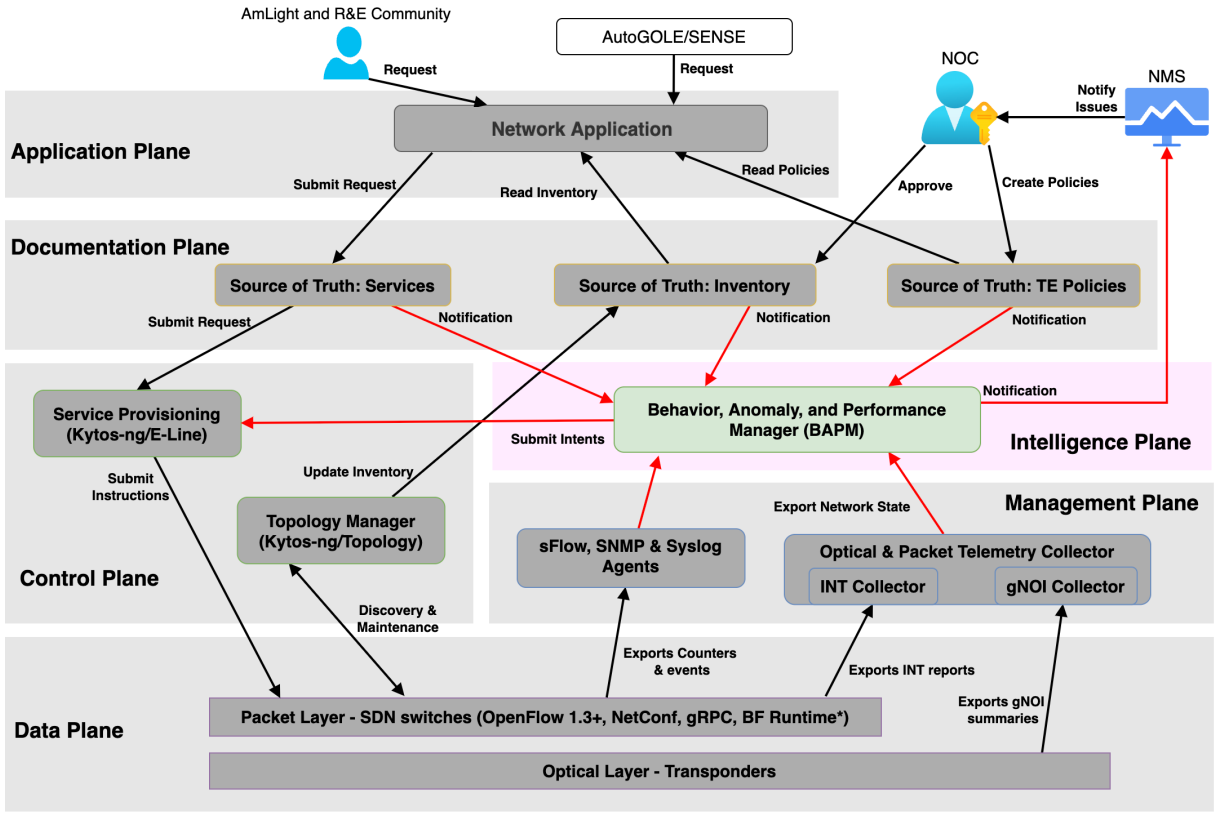

AmLight ExP
Americas Lightpaths Express & Protect

ANA at AmLight: a use case

AmLight ExP
Americas Lightpaths Express & Protect

ANA @ AmLight: a use case

- *The first Autonomic Function (AF): L2VPN*
- Topology Manager discovers the substrate and submits it for approval by NOC
- NOC approves the topology elements
- NOC creates Traffic Engineering policies:
  - For instance, average utilization over 2 minutes should be < 65% of the link capacity
  - L2VPNs with higher priority shouldn't be moved
  - Only move L2VPNs up to 3 times every 10 minutes
  - Don't move L2VPNs that have opted out for no protection

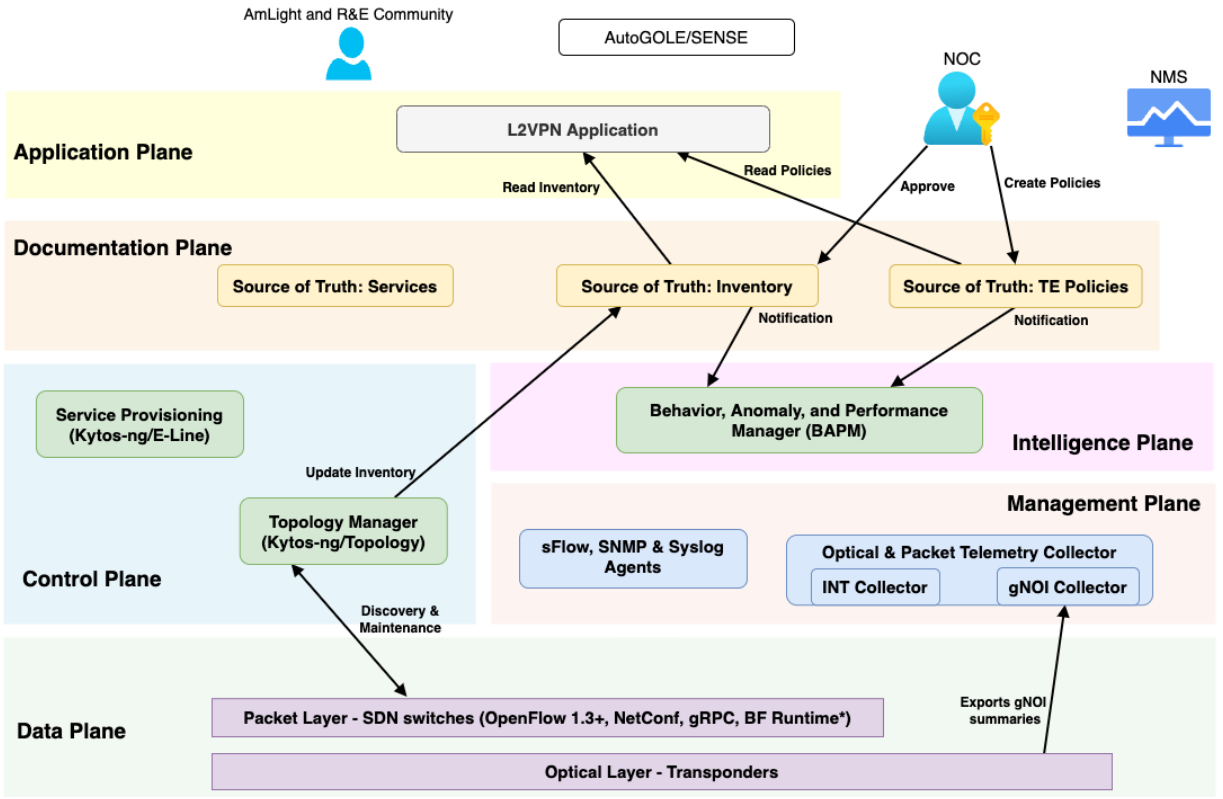

AmLight ExP
Americas Lightpaths Express & Protect

ANA @ AmLight: a use case [2]

The first Autonomic Function: L2VPN

AmLight user requests a L2VPN via Application Plane

1. Application Plane:
   - submits request to Documentation Plane that stores all metadata provided and triggers an event
2. Control Plane (Kytos E-Line):
   - gets triggered by the event,
   - retrieves the provided metadata,
   - computes a path based on the metrics provided
   - pushes flows to Data Plane
3. Data Plane (SDN switches):
   - Creates flow entries
   - Forwards user packets as requested

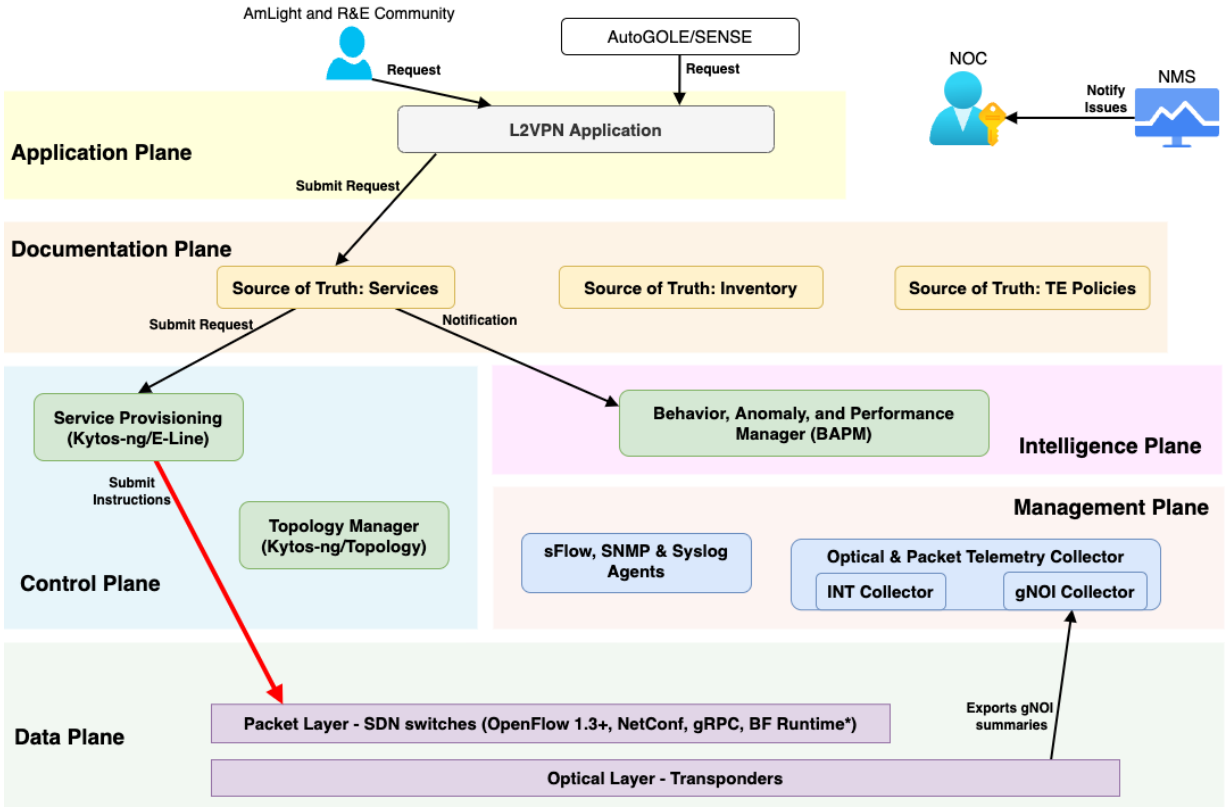

AmLight ExP
Americas Lightpaths Express & Protect

ANA @ AmLight: a use case [3]

The first Autonomic Function: L2VPN

AmLight user requests a L2VPN via Application Plane

4. Data Plane (SDN switches):
   - Forwards user packets as requested
   - Exports per-packet telemetry data to Management Plane

5. Management Plane (INT Collector):
   - Parses telemetry data looking for thresholds
     - Bandwidth utilization, queue occupancy, hop delay, path taken, …
   - If a threshold is reached, submit a network state report to Intelligence Plane

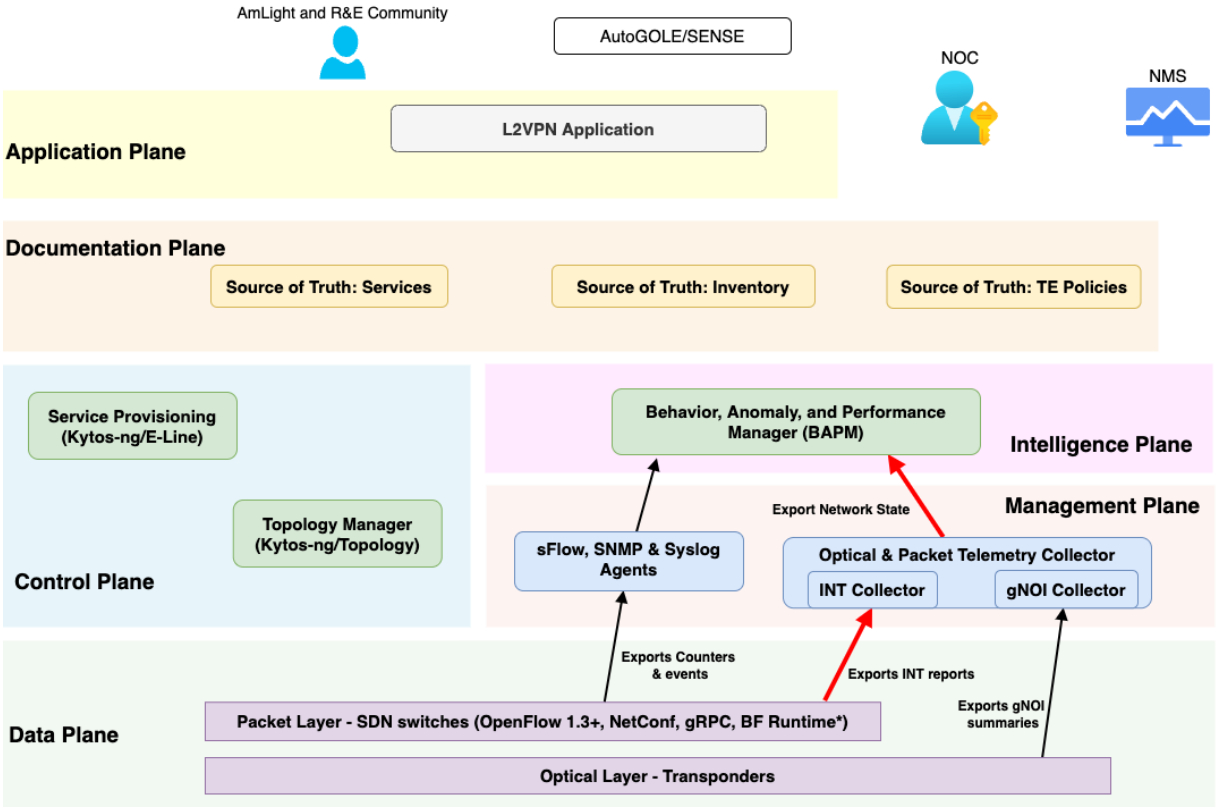

AmLight ExP
Americas Lightpaths Express & Protect

ANA @ AmLight: a use case [4]

The first Autonomic Function: L2VPN

AmLight user requests a L2VPN via Application Plane

6. Intelligence Plane (BAPM)
   - BAPM analyzes the network state report and policies
   - If a nonconformity is identified, BAPM uses the Services source of truth to identify if/which L2VPNs sharing the same resources (interface, queue, switch) should be relocated.
   - BAPM sends an intent to Kytos-ng E-Line asking to move L2VPNs [A, B, C] out of resource [Z].

7. Control Plane (Kytos-ng E-Line)
   - retrieves the provided metadata,
   - computes a path based on the metrics provided
   - Pushes flows to Data Plane

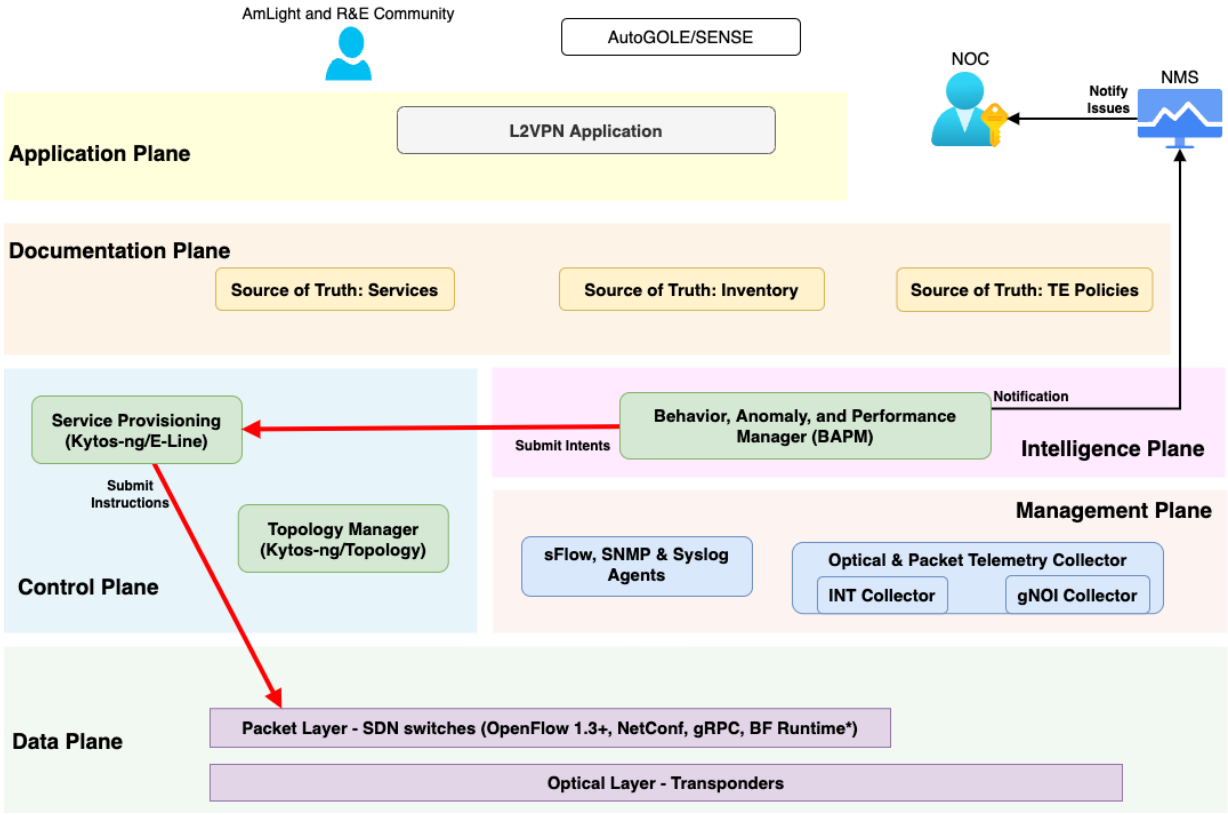

AmLight ExP
Americas Lightpaths Express & Protect

ANA @ AmLight: a use case [5]

The first Autonomic Function: L2VPN

AmLight user requests a L2VPN via Application Plane

8. Forever loop as long as policies aren't changed or L2VPN decommissioned.

Roadmap: Self-Optimizing the network:

- Year 2: < 5 seconds
- Year 3: < 2 seconds
- Year 4: < 1 second
- Year 5: < 500 ms

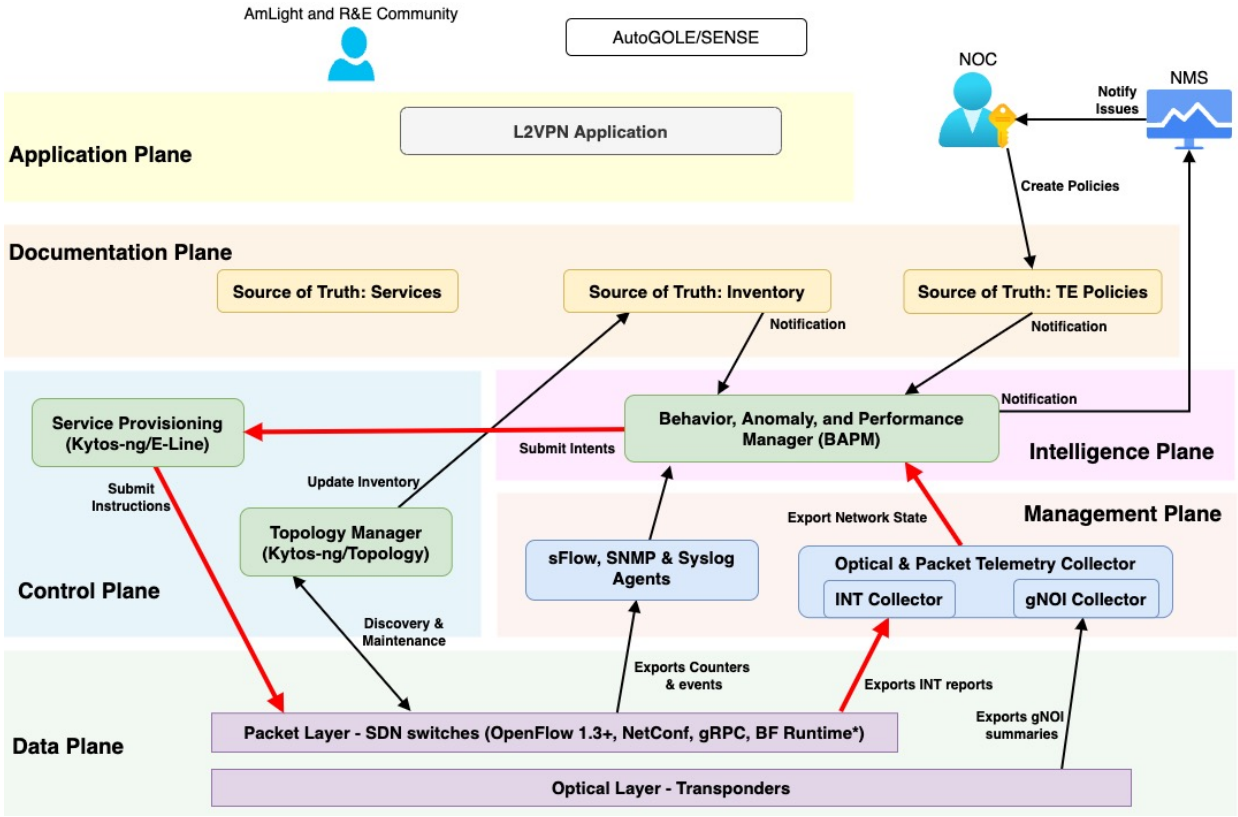

AmLight ExP
Americas Lightpaths Express & Protect

How will ANA support our goals?

- AmLight has many links and multiple paths between its sites:
  - From Chile to Jacksonville, there are more than 25 possible paths to take
  - With the new architecture, we expect to properly load balance network services across links, while respecting user constraints and requirements

- AmLight has an SLA-driven packet-loss-intolerant and sub-minute-response-time-expected science application:
  - With per-packet telemetry and sub-second network profiling capacities, AmLight will be prepared to react to network conditions under 1 second
  - With optical telemetry, AmLight will anticipate issues with the substrate and steer traffic out of the substrate before adverse events happen

- AmLight engineering team prefers to focus on engineering and new services than manual activities:
  - With the closed loop control, some time-consuming operational activities will be performed without human intervention
  - In the end, we will have more hands developing and automating routines

AmLight ExP
Americas Lightpaths Express & Protect

Thank You! Questions?

AmLight SDN Eng. Team <sdn@amlight.net>

The new 2021-2025 NSF IRNC AmLight-ExP project: Moving towards the Autonomic Network Architecture (ANA)