Juniper Apstra System

Athanasios Douitsis
January 2022

Juniper Networks | Engineering Simplicity

Agenda

- Intro to data center networking
  - Datacenter Fabric and Switching
  - VXLAN / EVPN / Clos topologies
- The Apstra Intent Based approach
  - The importance of the intent
  - Closed loop network management
- Demo

Problem: Creation of a Datacenter Network

- Need of a switch / router with many hundreds of ports
- Limit to the number of ports of a physical device (even the biggest ones)
  - Physical limitations
  - Geographical limitations (can't cover the entire campus with a single device)
  - Single point of failure limitations
  - Bandwidth limitations
- Usage of many switches connected to each other
  - **Spanning tree** makes active-active balancing hard
  - Changes of topology cause disruptions
  - End host change of port causes short disruptions

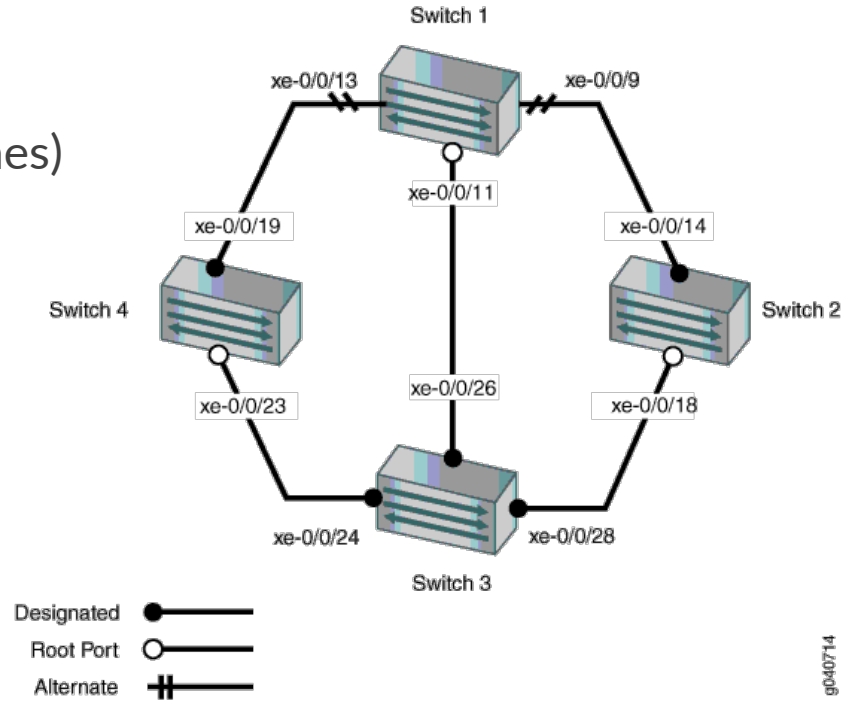

- Need for a conceptual fabric emulating a switch, with arbitrarily large bandwidth between any endpoint pair
  - Overlay a fake ethernet fabric over an IP underlay

Underlay: IP-based Clos network (3-stage, 5-stage)

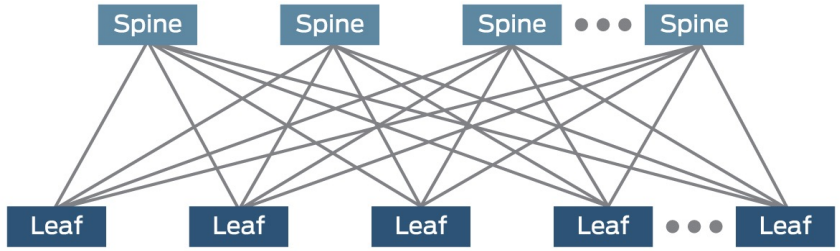

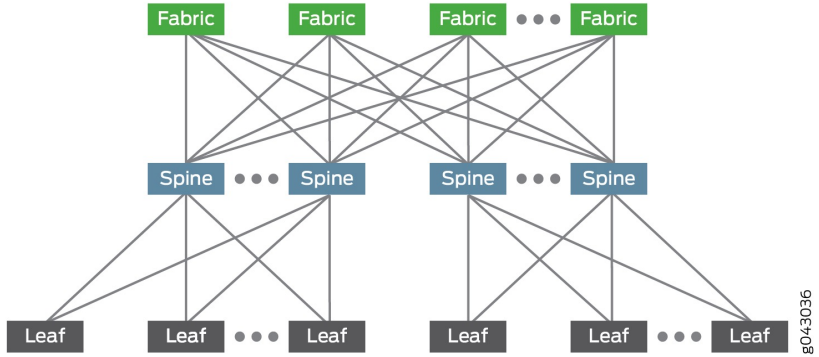

Overlay: VXLAN for the data plane

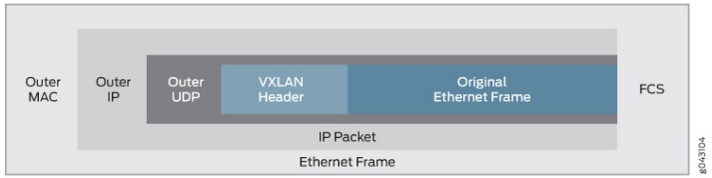

- Allows encapsulation of layer-2 (ethernet) packets in UDP
- So it becomes possible to use an already existing IP network as a switch
- Instead of real Ethernet frames between links, now there are UDP packets in the IP links
- Routing and traffic engineering is now possible, load balancing, quick recovery, etc.

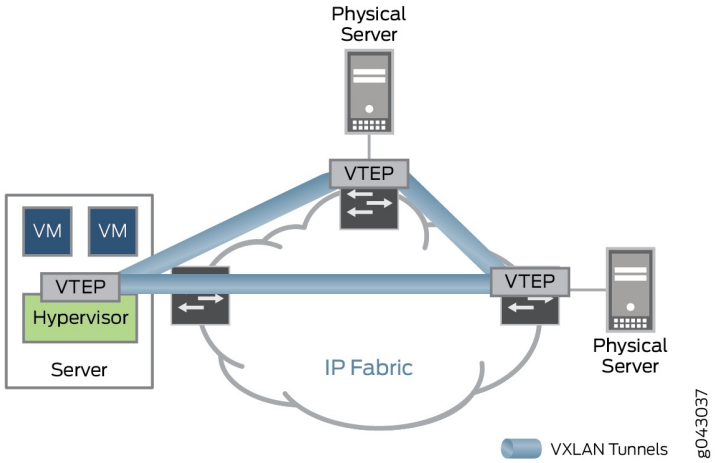

Overlay: BGP EVPN for the fabric control plane

- Switching control plane: 802.1d mac learning, broadcast
- Additionally: ARP / NDP, IPv4 / IPv6 routing

- Use BGP as the protocol to make the control plane communicate
- L2VPN/EVPN BGP address family to emulate the control plane functions (and beyond) of a real switch

- EVPN route types:
  - Type 1: Ethernet segement Identifier
  - Type 2: Mac route or Mac/IP route (includes ARP)
  - Type 3: BUM traffic delivery
  - Type 5: Pure IP routes (to cross between virtual segments)

Management of a fabric

- A traditional switch is relatively easy to configure and maintain, no serious monitoring necessary (example, a small home)
- A switch topology with VLANs is somewhat tricky to configure, easy to maintain and hard to monitor (example: the NTUA campus)
- An IP based fabric using EVPN / VXLAN is hard to configure, maintain and monitor

- Conclusion: A DC fabric cannot be approached using the same management principles
  - *Need for **automation***

“Traditional” Automation

- Configuration
  - Ansible
  - Chef
  - Puppet
  - SaltStack
  - Perl / Python / Ruby and libraries
- Monitoring and data collection
  - MRTG, Kibana, Gafana,
  - Prometheus
  - InfluxDB

(apologies for any fine tools and libraries not mentioned here)

Problems:

- Too much focus on how instead of why and what
- A human is needed as a CPU to parse the provided information – The real state of the network is stored in (one or more) humans
- Difference between ***syntax*** and ***semantics***
- The elephant in the room: Need for an operations-centric approach with a single source of truth

Operations day in the life

Chapter 1

New Launch!
Set up of VLANS, VMs, etc. in minutes

Chapter 3

Postponed Launch
All network rollback in less than a minute

Chapter 2

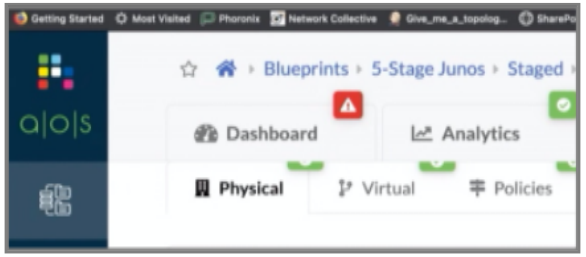

Corrupted Config
Quick visibility to root cause in intuitive dashboard

Chapter 4

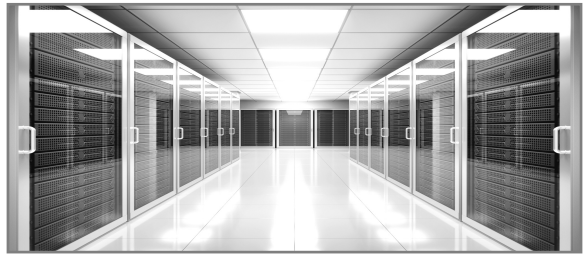

Scale to Meet New Demand
Scale new resources with pre-validated templates

Challenges of Day 2+ Operations

Network teams have too much to do….

- Monitoring tool proliferation and # of devices/components mean **'needle in a haystack'** challenge to pinpoint issues
- **Cross-functional finger-pointing**—networking teams on the defense and must prove innocence
- Length of time to **roll back** a change when issues arise
- **Change review** is onerous—delays new services and important fixes
- Too many **CLI** touchpoints for just one change
- **Lack of visibility** to grey failures to get ahead of device issues and prevent user impact
- **Security patches** and NOS updates can take long to plan and require (or trigger) outages
- **Lack perspective** of the whole network to understand what's going on
- **Multi-vendor** creates challenges in setup, visibility, and trouble identification
- **Networking skills scarcity** can make hiring challenging

Juniper Apstra difference

Operate the Network as One System

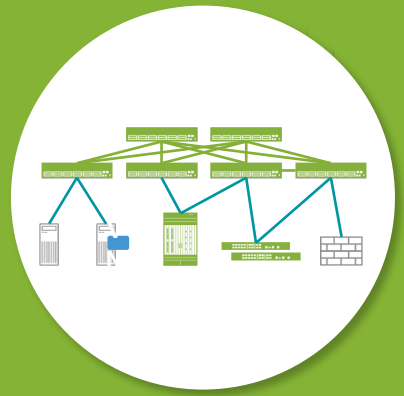

A Unified, Intent-based Approach

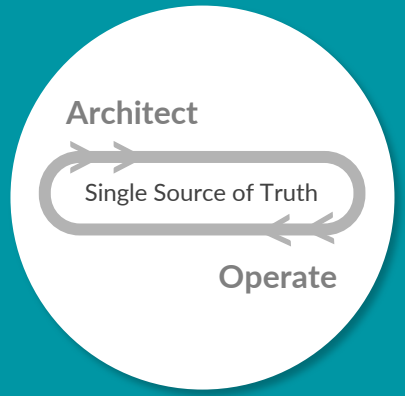

Open and Multi-vendor

One Unified Solution, Consistent Experience

Describe the **WHAT**

Software delivers the **HOW**

Know **WHEN and WHY**

Architect

Operator

Intent

continuous automation and closed-loop validation

Analytics

Day 0 **Design**

Day 1 **Build Deploy**

Day 2+ **Operate Assure**

Automate every day

Automate and Assure Your Data Center

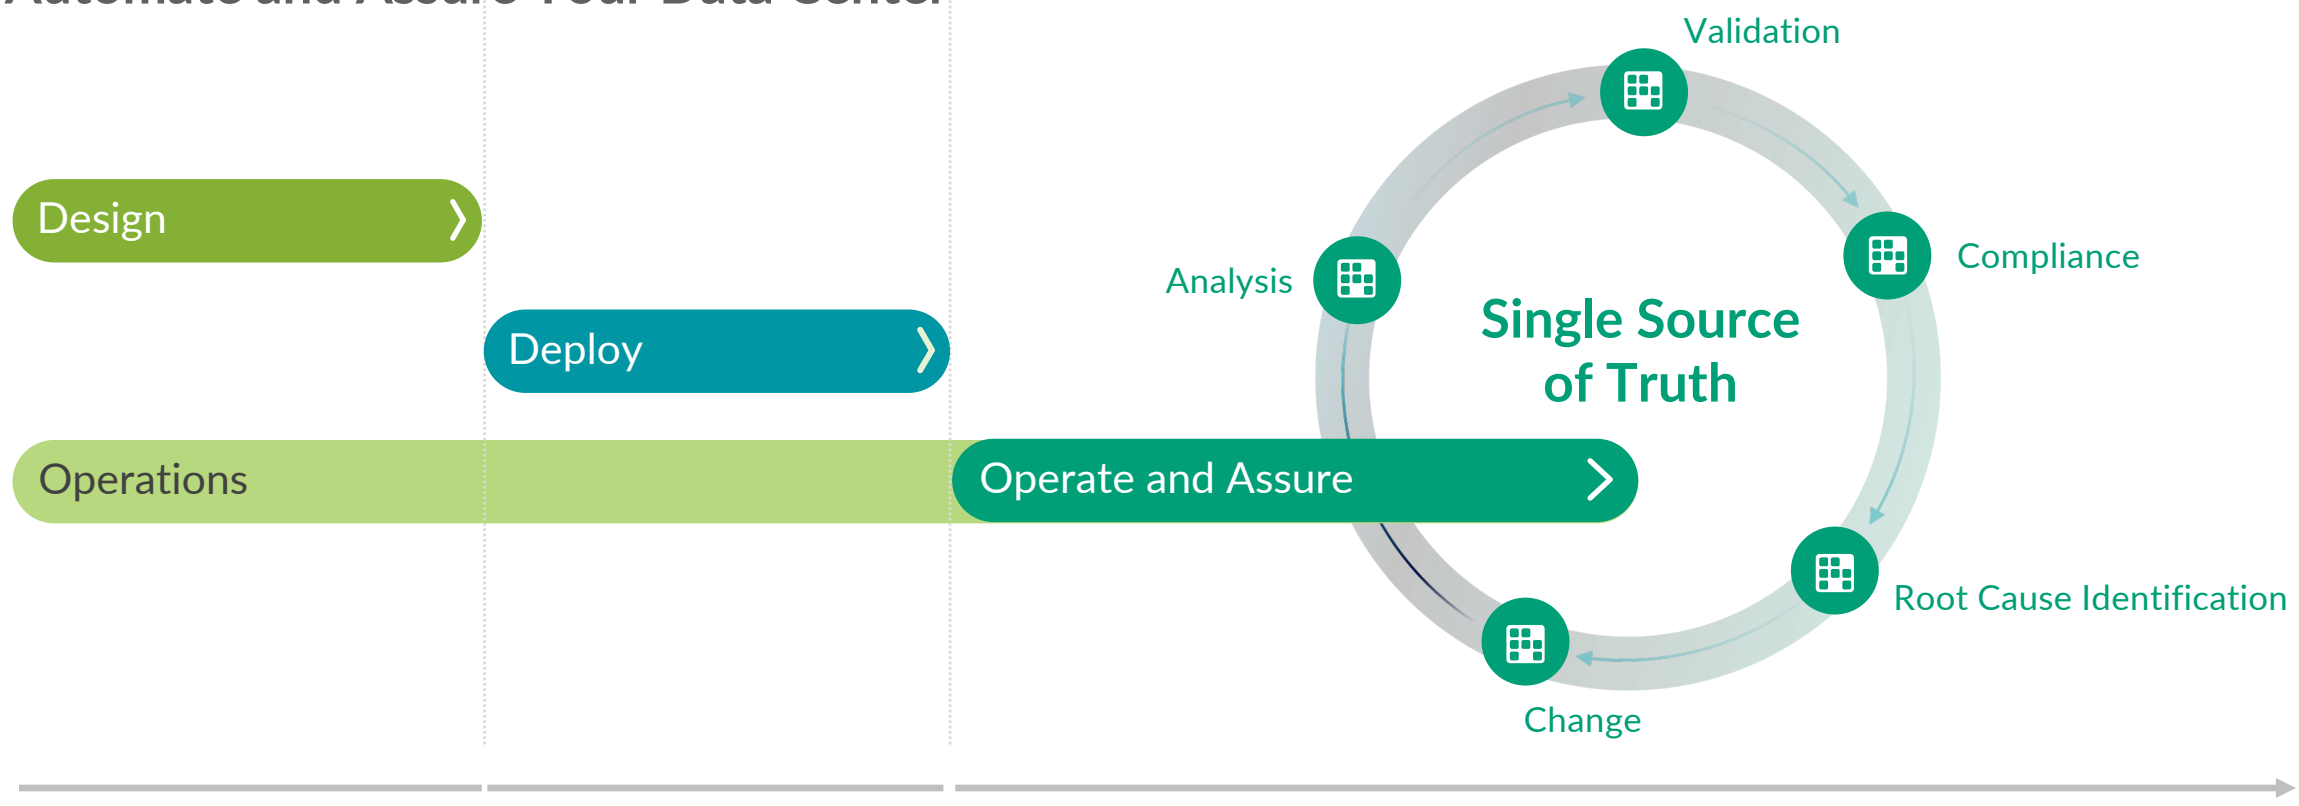

Day 0

Reference design
Pre-validation
Install

Day 1

Zero-touch provisioning
Test
Validate

Day 2+

Visibility / Analytics / Insights
Troubleshoot/ Remediate
Optimize

Change management
Maintain / Update
Compliance / Audit

Apstra: Intent Based Networking Solution details

- Standardised **Reference Design** Solutions
  - Works across Tier-1 vendors such as Junos (+Junos Evolved), SONiC, NXOS, EOS, Cumulus
- OPEX saving advantage
  - Abstract scalable **Blueprint** for DC networks design (CLOS)
    - Template can be replicated across large DCs
  - Dynamic configuration generation following the **Intent**
    - Graceful handling of day-2 operations
  - Closed loop device management
    - Device ***expected*** state monitoring by **telemetry** components
    - Verification of **Intent**, detection of **deviations**
- Fast problem resolution
  - Intent-Based Analytics
  - Root Cause Identification

Apstra Key Technologies

Intent-Based Networking

Benefit: Simplify effort of architects and operators to design, deploy and operate

Outcome: Transformed focus on the business results with insights for continuous improvement

Single Source of Truth

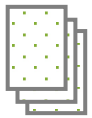

Benefit: Speed operations actions with repeatable, vendor-agnostic blueprints and knowledge graphs

Outcome: Faster migration/change with more time on value (not the arcane semantics of management)

Closed-Loop Validation

Benefit: Assure with continuous verification, proactive insights and root cause analysis

Outcome: Reduce problems, outages and mean time to repair while raising operational efficiency

Time Voyager Rollback

Benefit: Avoid change issues with visibility, fast rollback and system-documented change control

Outcome: Reduce business impact of errors and assure compliance, auditing and knowledge retention

Maintenance/Upgrade Mode

Benefit: Separate HW/SW upgrade cycles to reduce maintenance windows and planned downtime

Outcome: Increased commitments to SLAs and user satisfaction and lowered risks of outdated software

Flexible Integrations

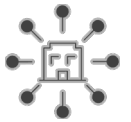

Benefit: Support existing and future cross-organizational workflows and new vendors

Outcome: Quick compliance to changing business operations and lower cost of technology adoption

Demonstration

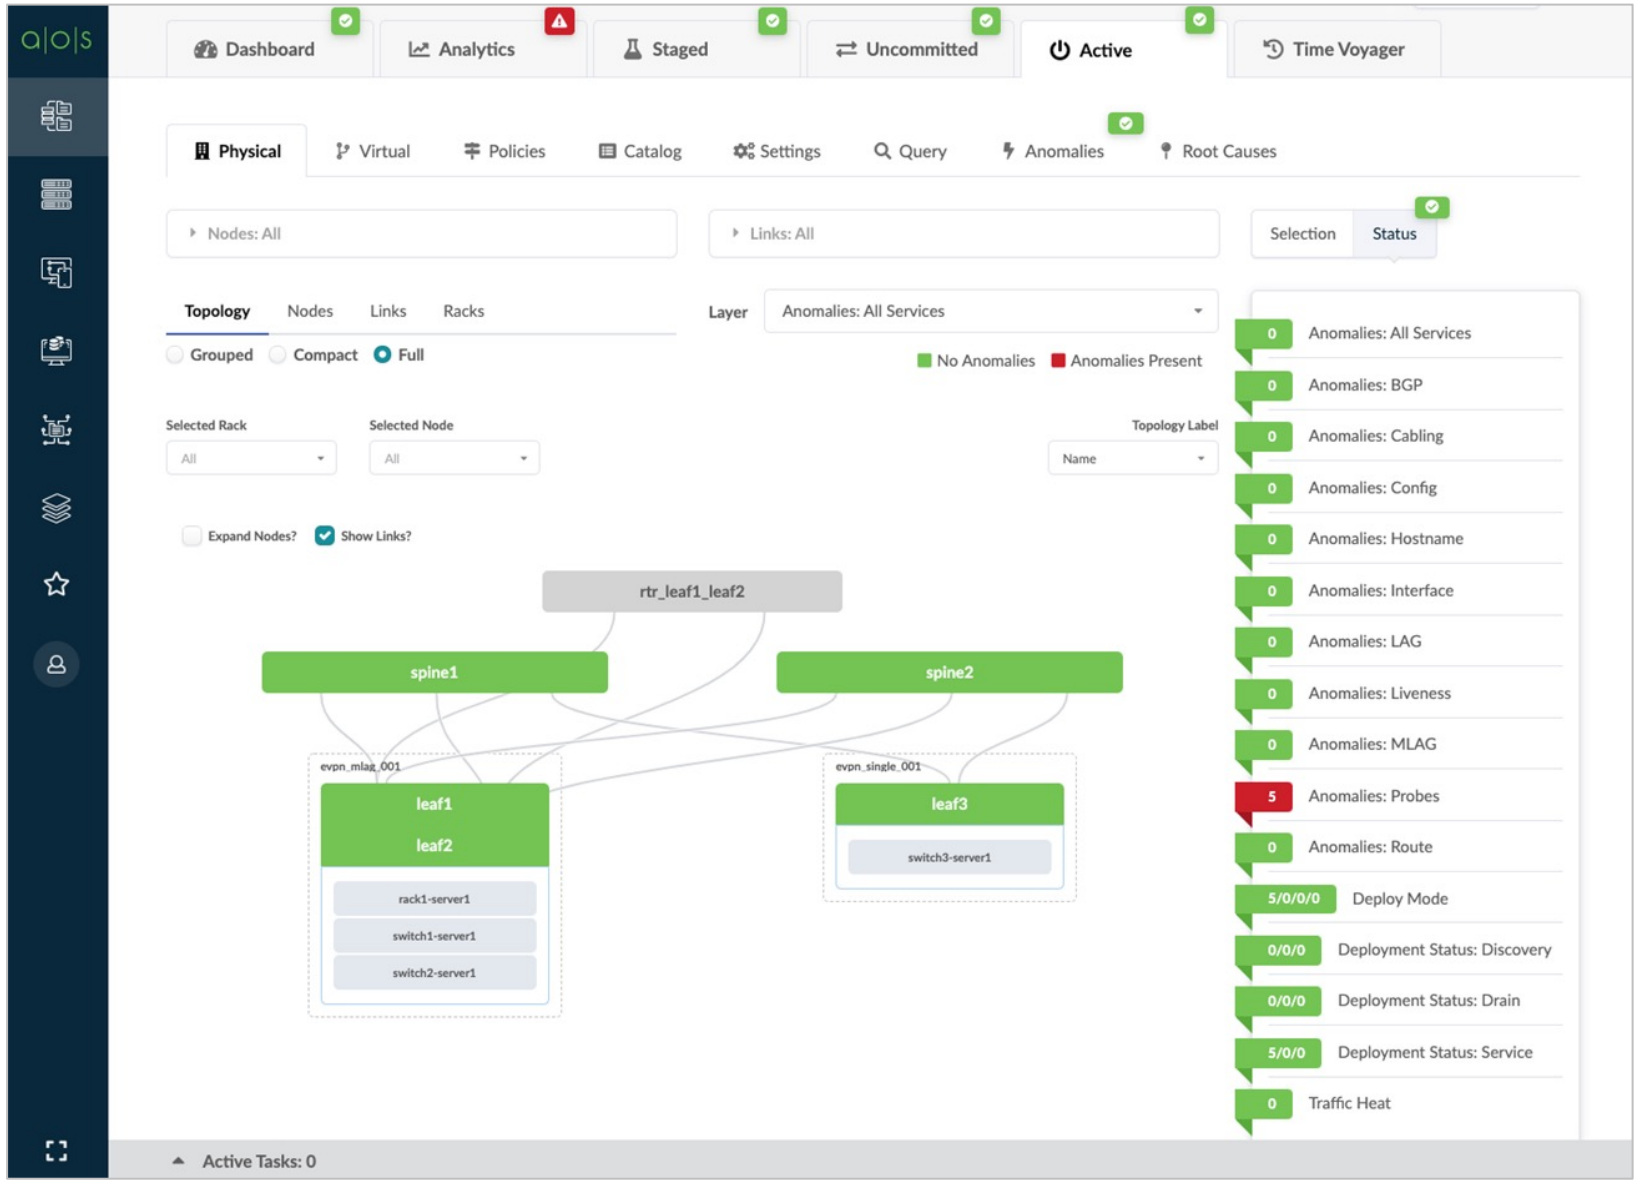

Learn It. Try It. (for free)

Apstra Academy

https://apstra.com/products/apstra-academy

YouTube

https://juniper.net/apstra-playlist

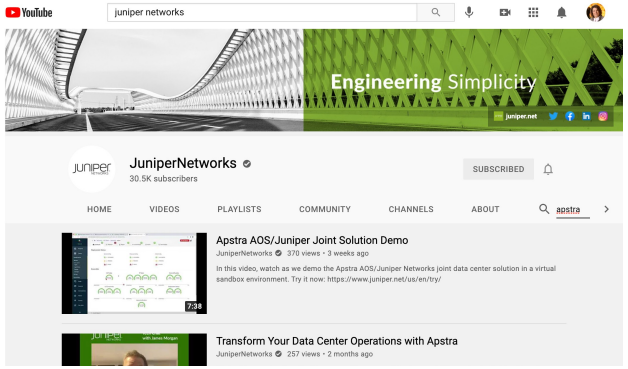

Juniper vLabs

https://www.juniper.net/us/en/forms/apstra-free-trial/

- Cloud-based lab environment
- Virtualized, pre-built network topologies
- Available for free!

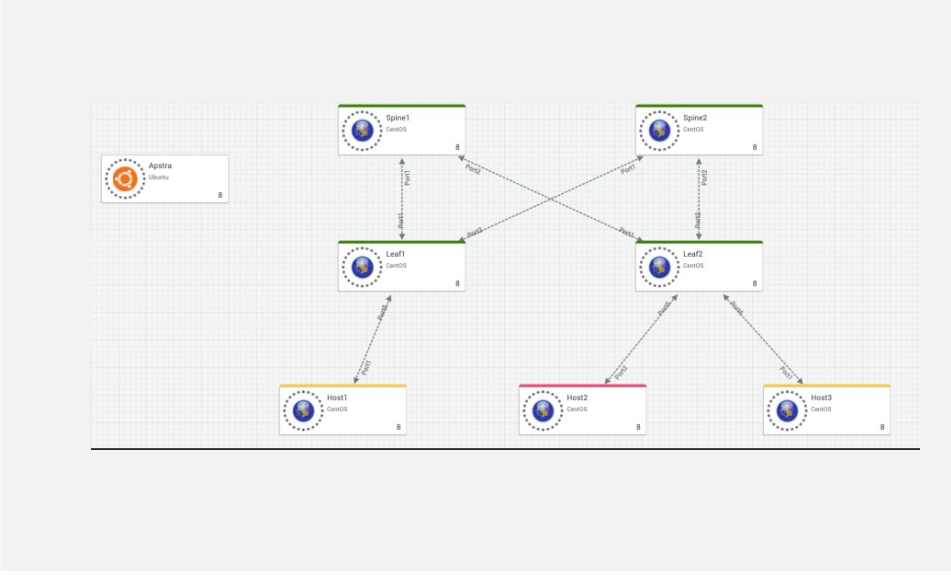


Thank you
JUNIPER NETWORKS
Engineering Simplicity
adouitsis@juniper.net