



VENDOR AGNOSTIC DATACENTER SDN SOLUTION

HUNOG – October 2023

Istvan Laky – Senior Systems Engineer

ilaky@juniper.net

What we here is missing...



Simplicity

Reduce costs and complexity with cloud-like operations you control



Assurance

Provide reliability and security to maintain expected outcomes



Automation

Automate the routine and mundane to reduce risk and move faster



Experience

Improve experience with intelligent analysis, response and proactivity

Need for a handful tool which ...



- Provides a single software platform for datacenter networks that is used for Day 0 (design), Day 1 (build/deploy), and Day 2 (operate) phases of the lifecycle of the network.
- Integrates with multiple hardware platforms (Juniper, Cisco, Arista, and SONIC).
- The software continually monitors, automates and validates datacenter network switching infrastructure in real-time and streamlines and verifies change, significantly reducing OPEX cost and management complexities.

.. on Intent basis

What do we mean by Intent

Our Servers:

- How many NICs?
- Speed of NICs?
- Function of NICs?
- NICs connected to what TORs?

Our TORs/Leafs/Spines?

- How many TORs?
- NIC/TOR redundancy protocols?
- Bandwidth/oversubscription to spine?
- Form factor of TORs?

External Connections:

- How many external connections?
- Speed of external connections?
- What connections go where?
- Route summarization, filtering, injection?

Network Isolation:

Server facing VLAN/VXLANs?

- VXLAN control protocol? EVPN or static?
- Multitenancy - VRFs?
- Security policy?
- Default settings per device (MOTD, syslog, AAA)

Mission of an intent base management system

Scale, performance and function

- **Intent-based workflows** to continuously validate best practices are set up
- Apstra deployments scale to 1000's of networking devices.
- Multi-vendor (SONiC, etc.); Integrations: **VMware vCenter and NSX-T**

Ease of use and time savings

- Intent Time Voyager to **rollback** the entire network, single or multi-vendor, to any previously known state
- Full network lifecycle automation and assurance that is **easy and intuitive**

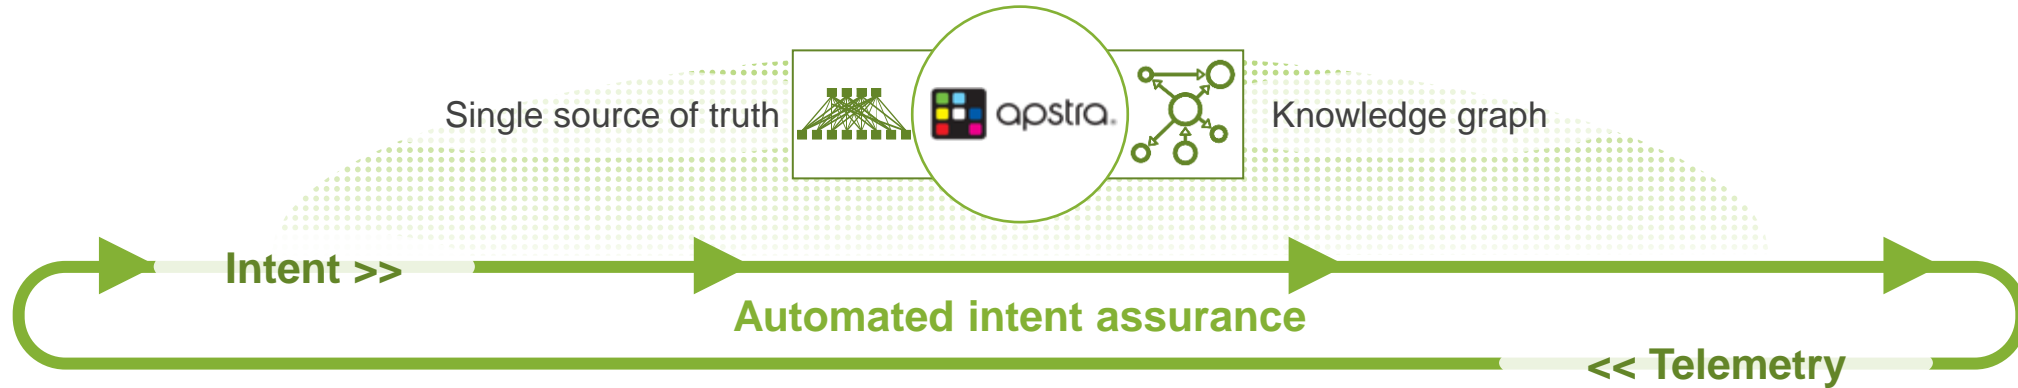
TCO Savings

- Automated provisioning with easy-to-use GUI and powerful intent-based assurance
- Users are quickly self-sufficient – easy to install and easy to use **without post-sales support.**

Self-driving network building blocks

- **Unique root-cause identification** from intent knowledge graph (single source of truth) in closed loop automation
- Avoid outages and ensure that network intent is achieved through predictive analytics
- Integrate to **self-service and DevOps** for provisioning with Ansible, REST APIs and CLI

Automated and Intent-assured Operations



Design

- Vendor-agnostic
- Multiple fabrics
- Multiple pods, racks, servers
- DC & DCI EVPN-VXLAN/IP

Build

- Vendor-specific blueprints
- Render cabling plan
- App/service modelling
- Architecture validations

Deploy

- Easy network ZTP/configs
- Time voyager rollback
- Upgrade/maintenance mode
- Fabric multi-tenancy

Assure

- Continuous validation
- Root-cause identification
- Service-level monitoring
- Customizable dashboards

Integrations:





Step - 1
DESIGN

Designing the Logical switches

Logical Device

- The port count / type required



Interface Map


- Which Physical ports used



Device Profile

- The physical switch itself

- 4 * 100G Spine
- 2 x 100G Peer MC-LAG
- 2 x 25G External Router
- 46 * 10G L2 / L3 server

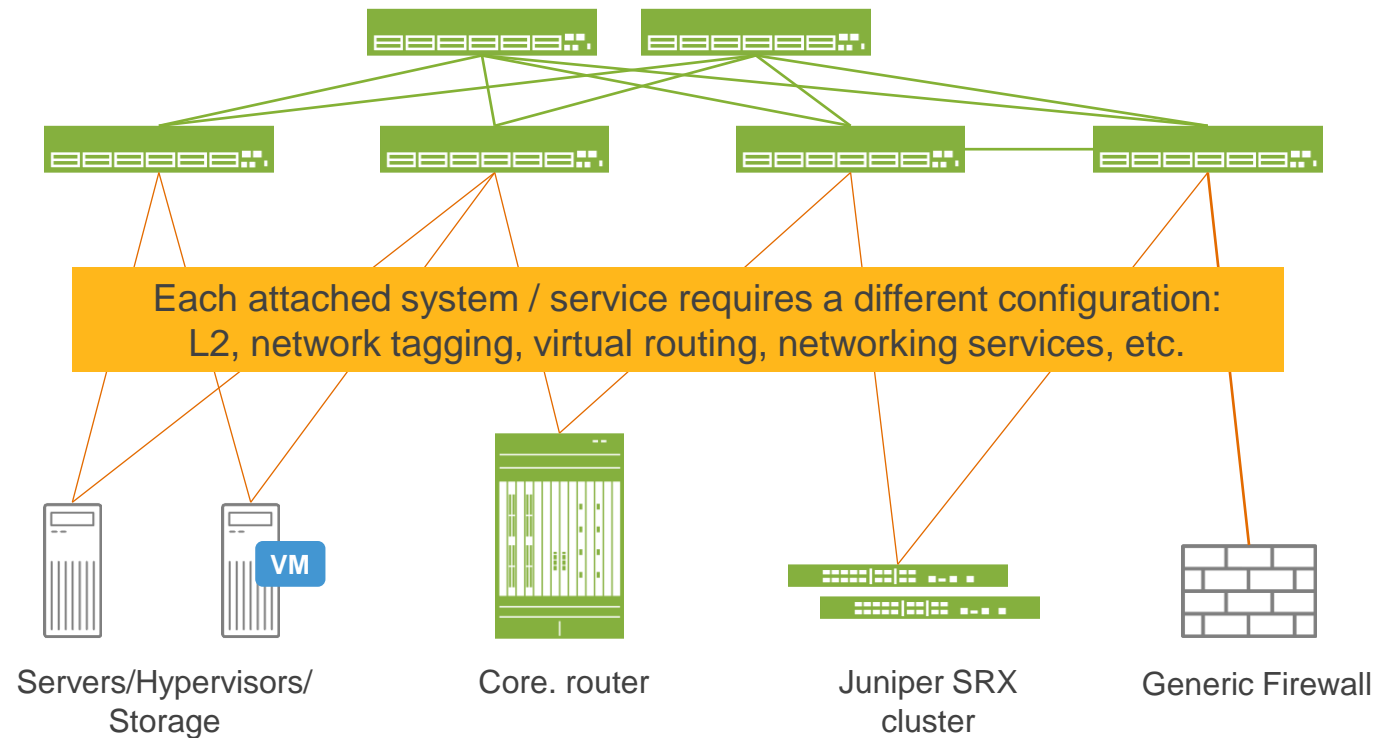


Connectivity Templates

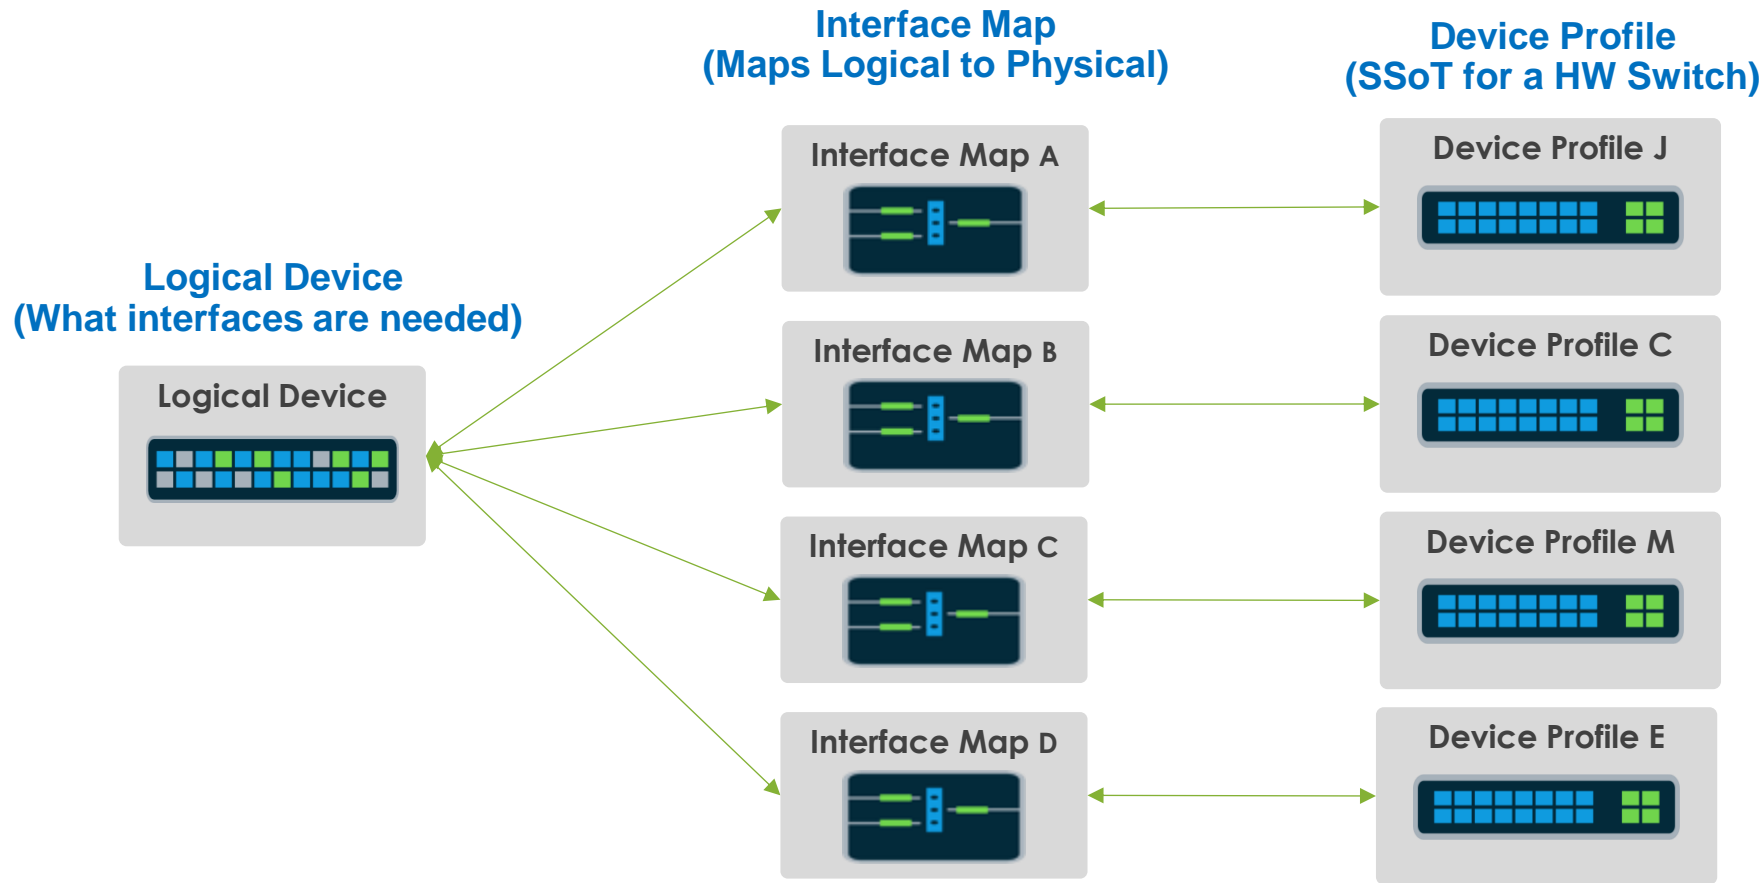
Simplified, validated, and repeatable connectivity design for attached systems

Standardize deployment in a customizable way for predictable operations for:

- Bulk, accurate adds across entire EVPN-VXLAN fabric in minutes
- Validation that everything in the network is functioning properly



The design / build process



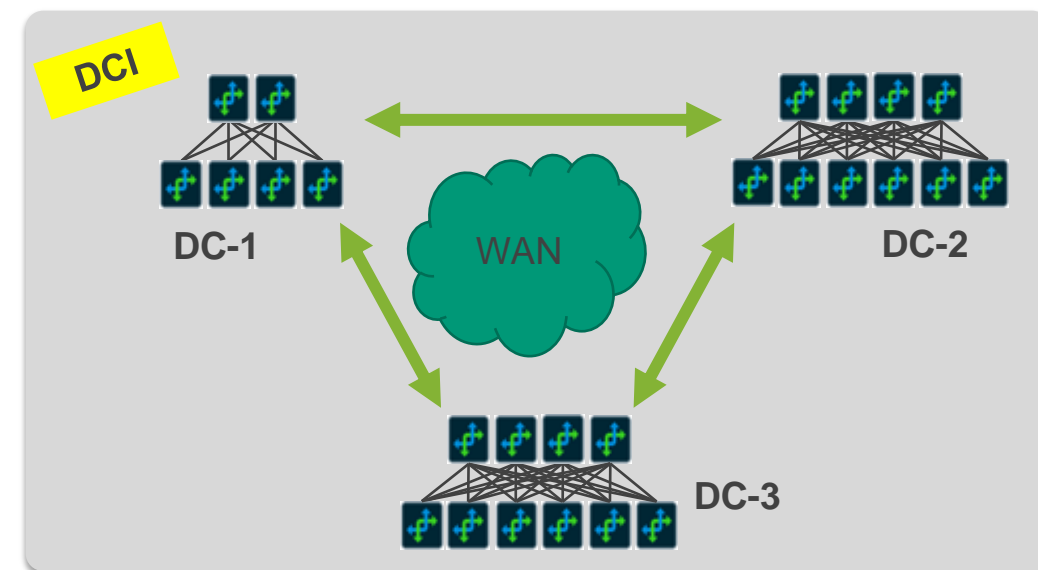

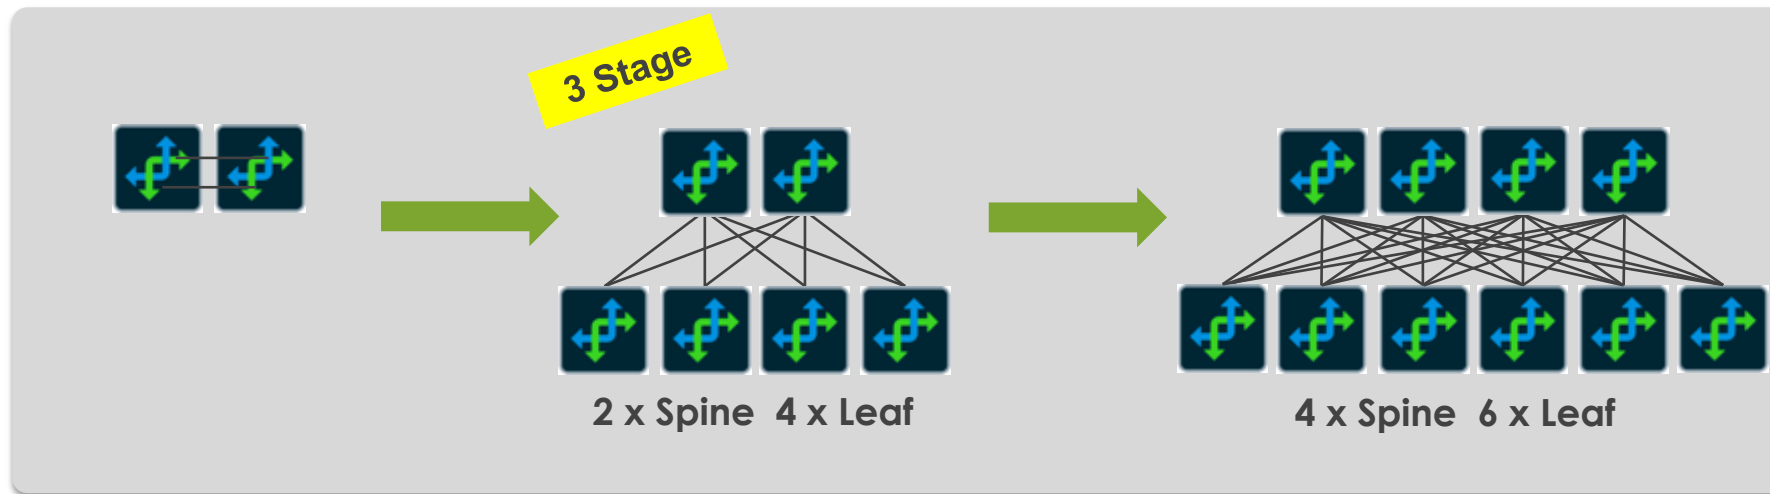
JUNIPER NETWORKS Junos EVO

ARISTA EOS

CISCO NXOS


DELL SONiC

IP Fabric - Any Topology | Scale | Vendor



Reference design – Edge Routing

- Prefer Bottom injected vs Top injected approach:
 - North-South traffic handled at the leaf layer.
 - Use Border-Leaf, and/or Border-POD in large 5Stage.
 - Benefits: traffic. fully symmetrical/equidistant and load-balanced over all links available.
- EVPN Routing:
 - Performed at the leaf layer: Edge Routed Bridging.
 - Allows better scaling – routing responsibilities distributed.
 - ARP Suppression at the leaf.



Reference design – eBGP as the routing protocol

- eBGP for the underlay & overlay - rfc7938
 - Well-known, well-documented and most inter-operable IP protocol.
 - Scale extremely well. iBGP doesn't scale well (full mesh required) – RR fixes this
 - Very predictable, endless customization.
 - Better at handling routing loops (AS-PATH) as opposed to iBGP
 - Path Management: AS-PATH is similar to a traceroute
 - eBGP satisfies the external router / DCI use case
 - Extensible: MP-BGP can carry more than IP prefix (EVPN AF) --> One protocol for underlay and overlay simplifies troubleshooting – quicker root cause

[\[Docs\]](#) [\[txt|pdf\]](#) [\[draft-ietf-rtgw...\]](#) [\[Tracker\]](#) [\[Diff1\]](#) [\[Diff2\]](#) [\[Errata\]](#)

INTERNET ENGINEERING TASK FORCE (IETF)
Request for Comments: 7938
Category: Informational
ISSN: 2070-1721

Errata Exist
P. Lapukhov
Facebook
A. Premji
Arista Networks
J. Mitchell, Ed.
August 2016

Use of BGP for Routing in Large-Scale Data Centers

Abstract

Some network operators build and operate data centers that support over one hundred thousand servers. In this document, such data centers are referred to as "large-scale" to differentiate them from smaller infrastructures. Environments of this scale have a unique set of network requirements with an emphasis on operational simplicity and network stability. This document summarizes operational experience in designing and operating large-scale data centers using BGP as the only routing protocol. The intent is to report on a proven and stable routing design that could be leveraged by others in the industry.

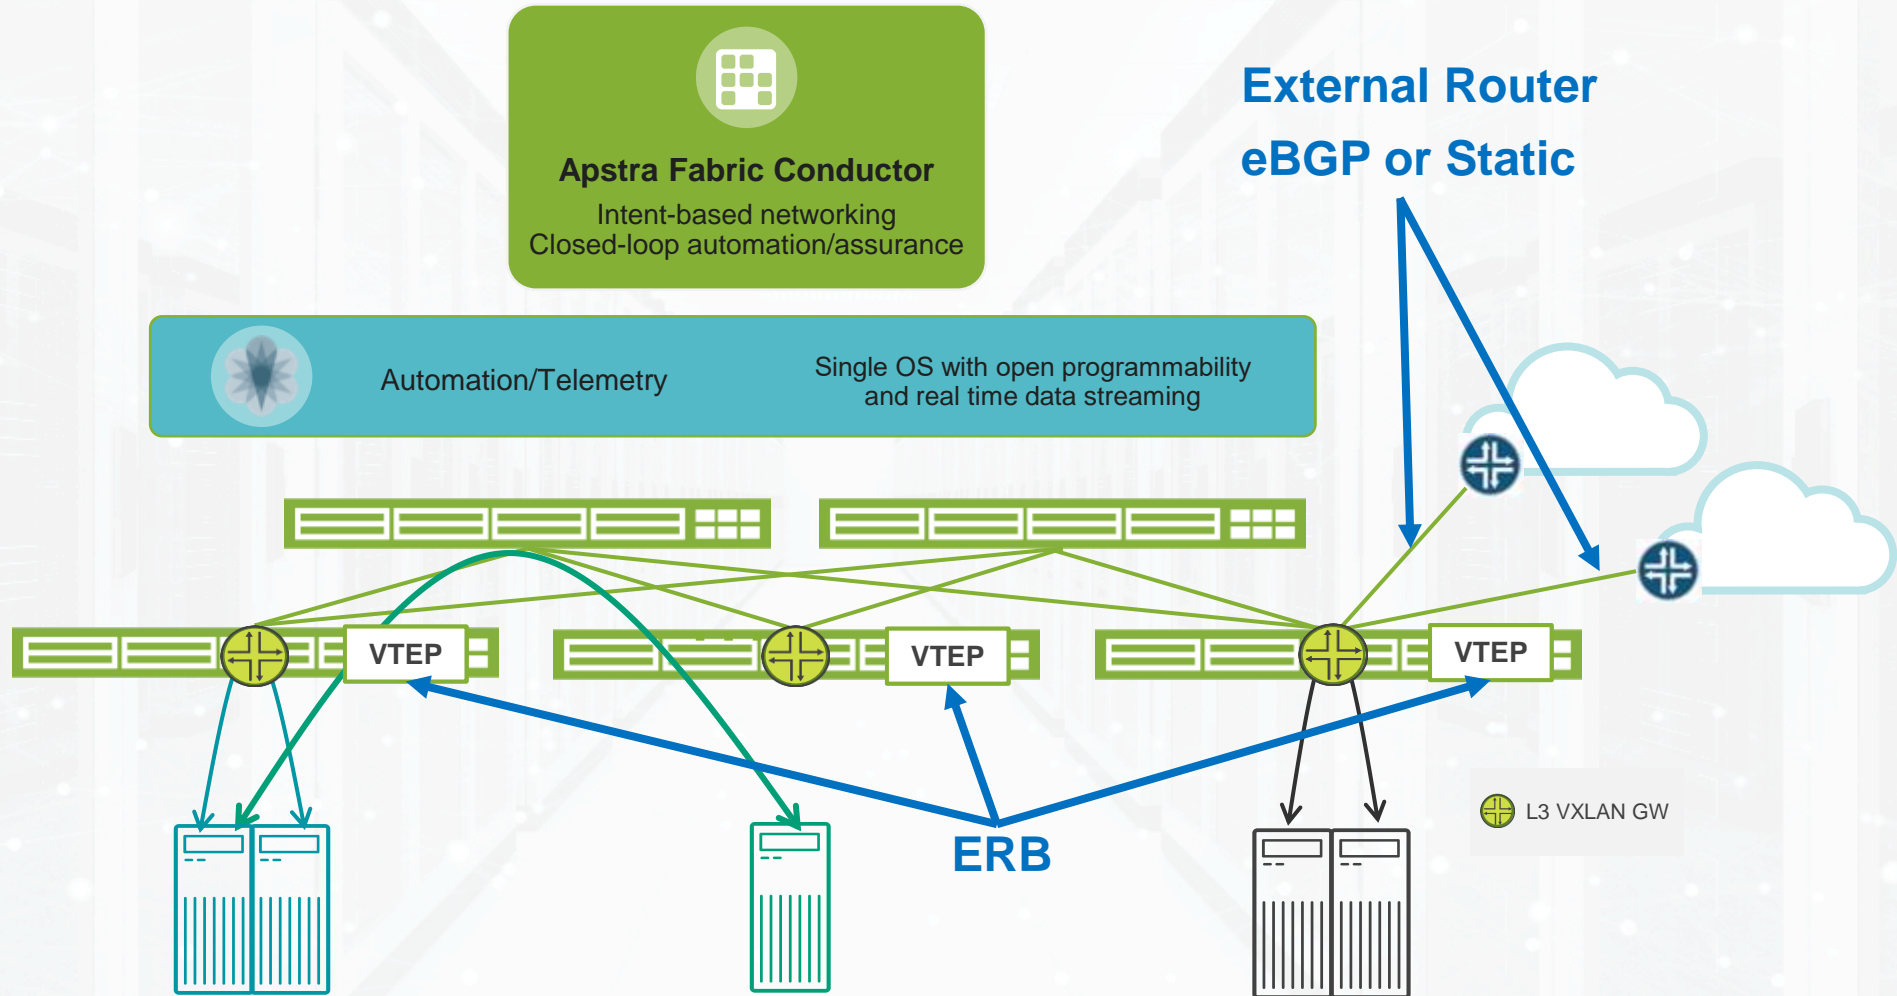
Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see [Section 2 of RFC 7841](#).

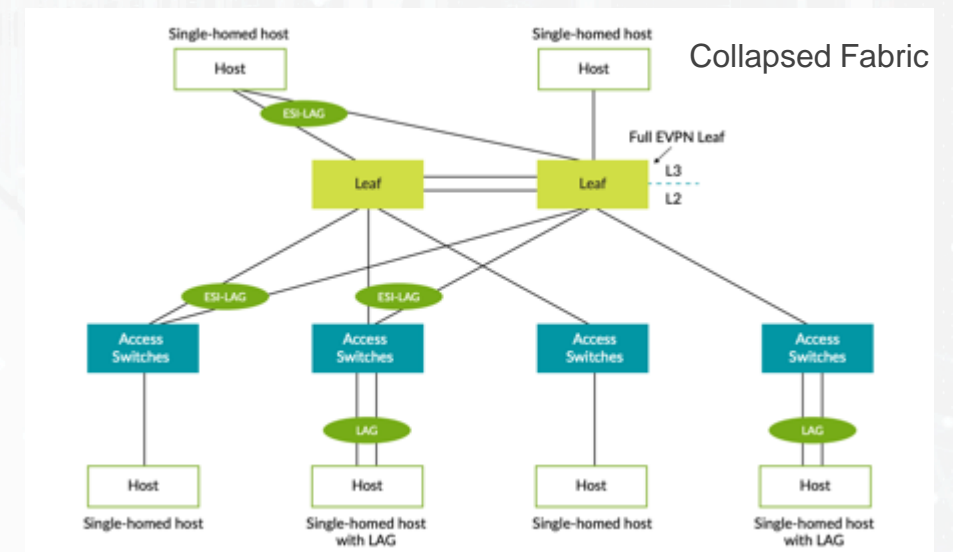
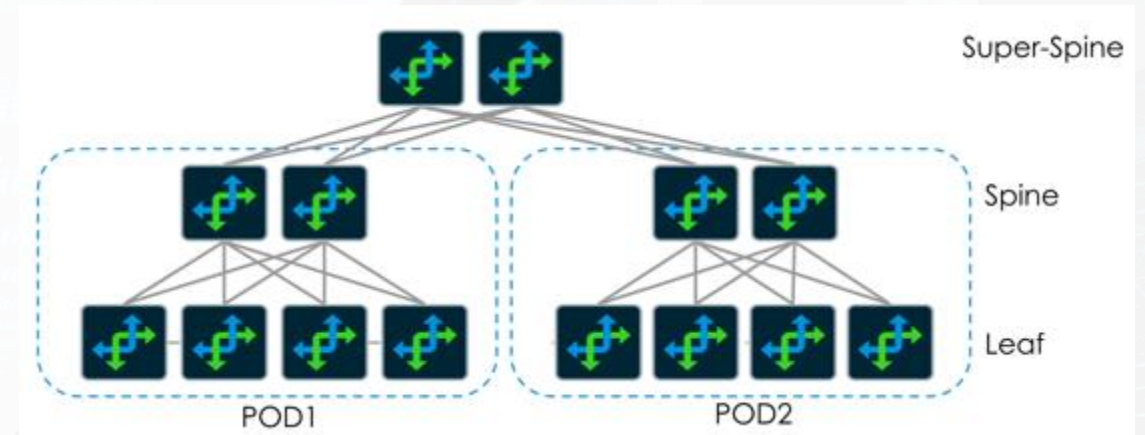
Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7938>.

Reference Design: Distributed routing model - scale



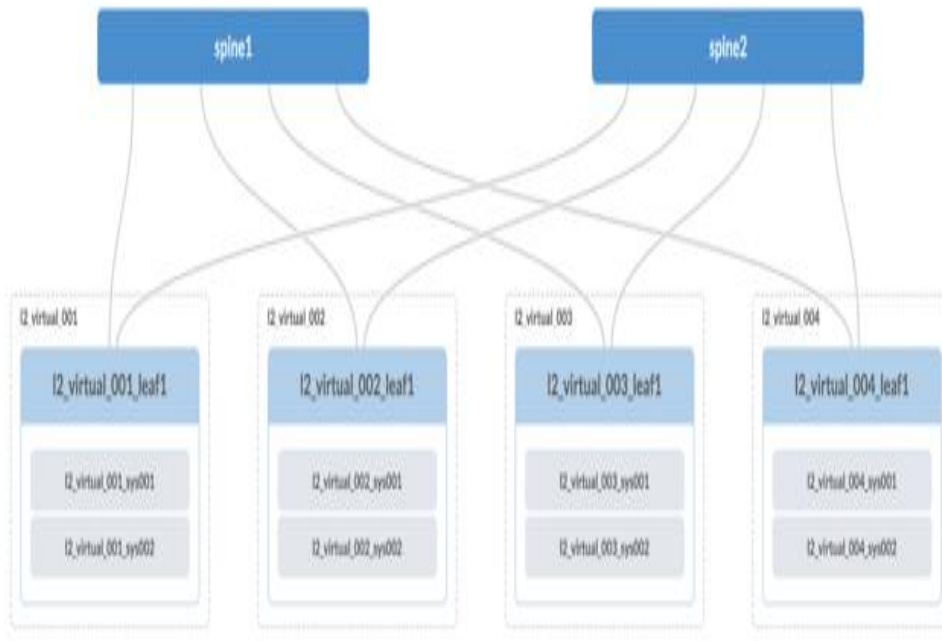
Reference design – Clos Topologies

- Edge-routed (ERB) EVPN-VXLAN (eBGP) overlays and DCI
- Collapsed Fabric
- Bridged overlay (BO) supported (ERB w/o IRB)
- 5-Stage introduces the notion of POD.
- Each POD being a 3-Stage Clos with uplinks from Spines to Super-Spine.
- You can have one or multiple Super-Spine planes.




User intent in a Graph Database

Expand Nodes? Show Links?




=


- Click to toggle
- logical_device
 - interface
 - link
 - system
 - virtual_network
 - domain
 - vn_instance
 - routing_policy
 - redundancy_group
 - I3_edge_ip_connectivity_policy



At scale



=





Step - 2










BUILD and DEPLOY

Time Voyager - Entire DC Rollback/Rollforward

Revisions


Query: All

1-3 of 3 Page Size: 25

Description	Created At	User	Actions
Added VXLAN Red100 to all 4 racks	2020-02-21, 12:52:34 current	admin	  
Added Rack 4	2020-02-21, 12:46:29	admin	  
Initial Config - 3 racks	2020-02-21, 12:39:56	admin	  

VMware NSX-T integration

Adds to pre-existing V-Sphere and NSX-T integration



How it works:

1. Apstra polls the NSX-T controller
2. A server admin creates a new portgroup
3. Apstra configures resources in the underlay
4. Apstra knows and exposes:
 - Server roles, name, IP address
 - Server uplink config
 - NSX-T N-VDS
 - Overlay to underlay maps

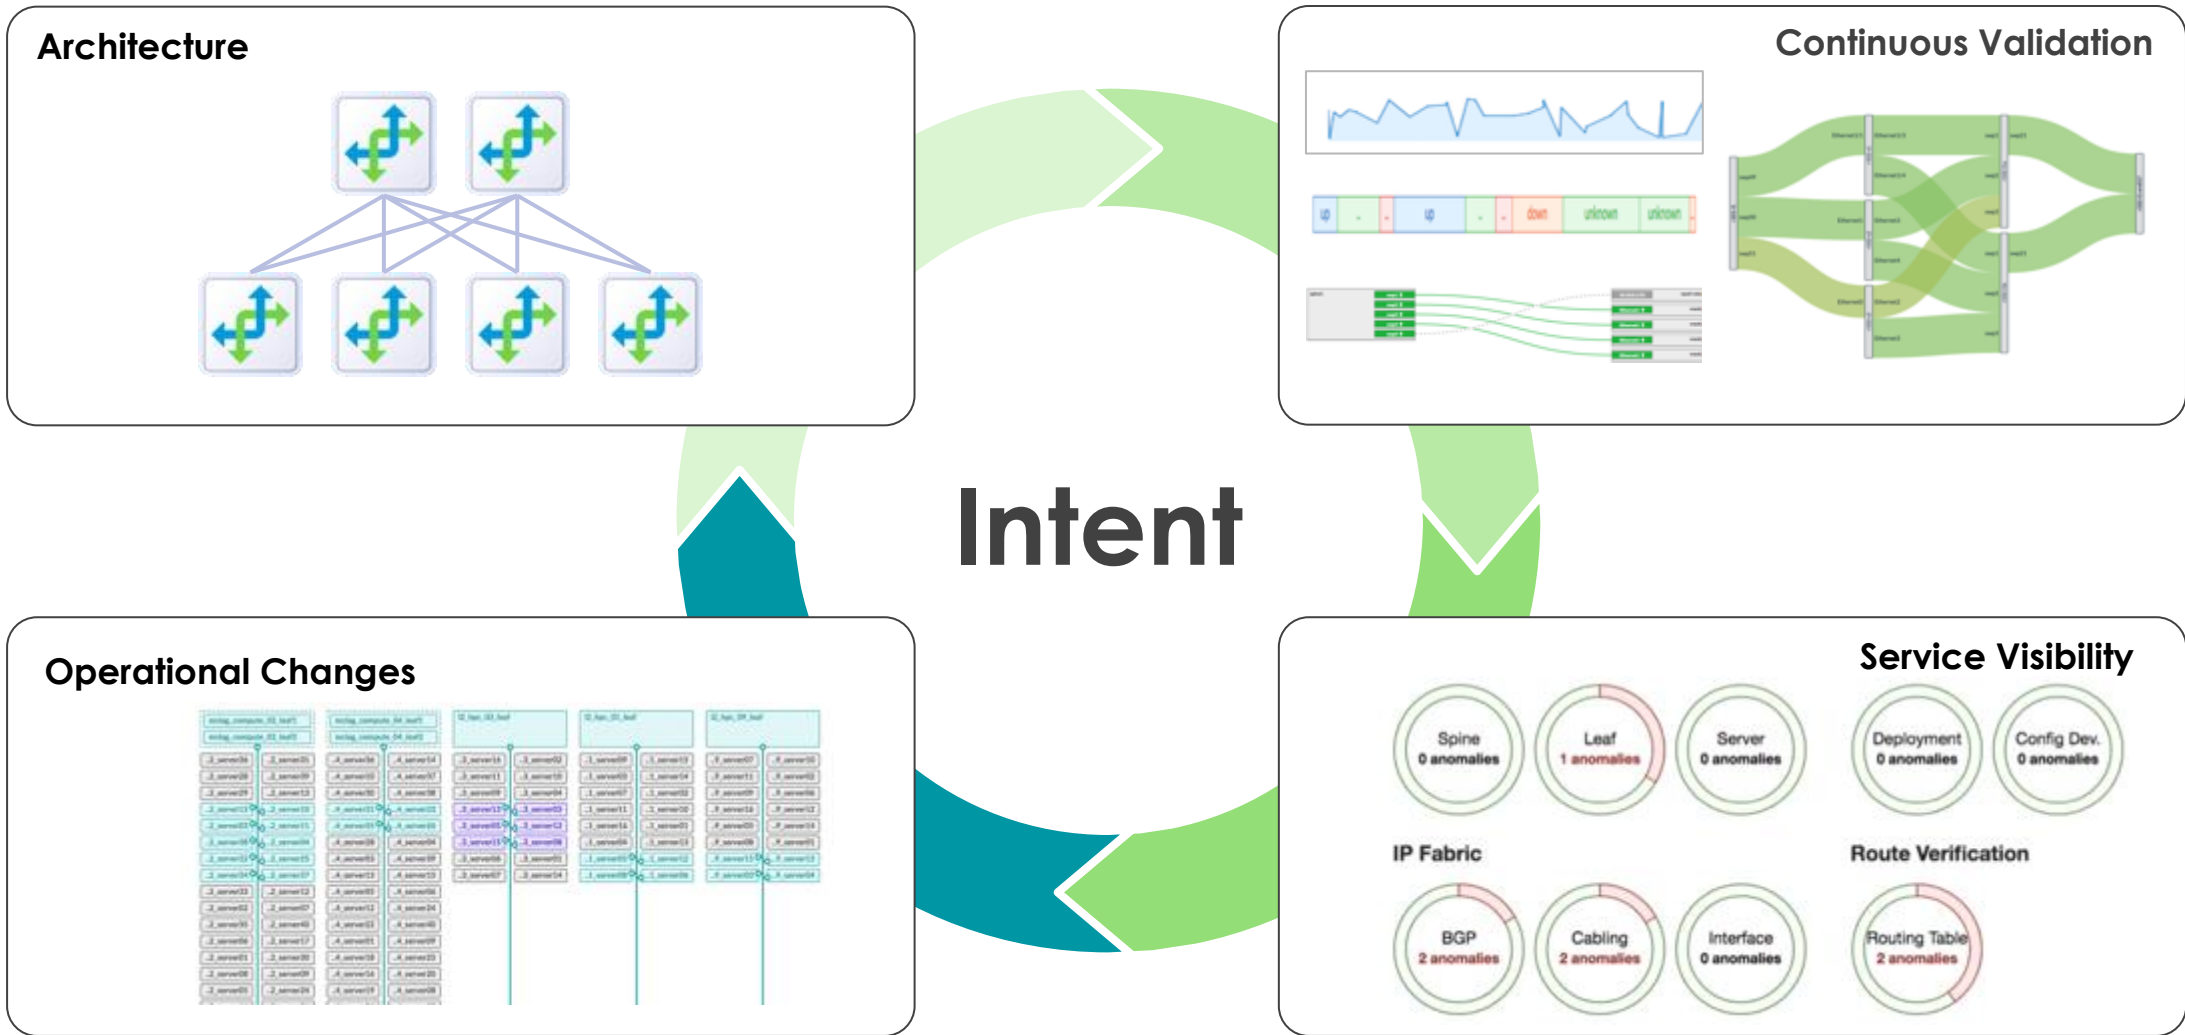
Results:

- Automate fabric setup for workloads—changes from up to weeks to minutes
- Faster mean time to innocence for the network during app issues
- Locate VMs within the fabric—avoid stranded VMs
- Optimized physical network for the NSX-T defined virtual overlay
- Integrated remediation workflows

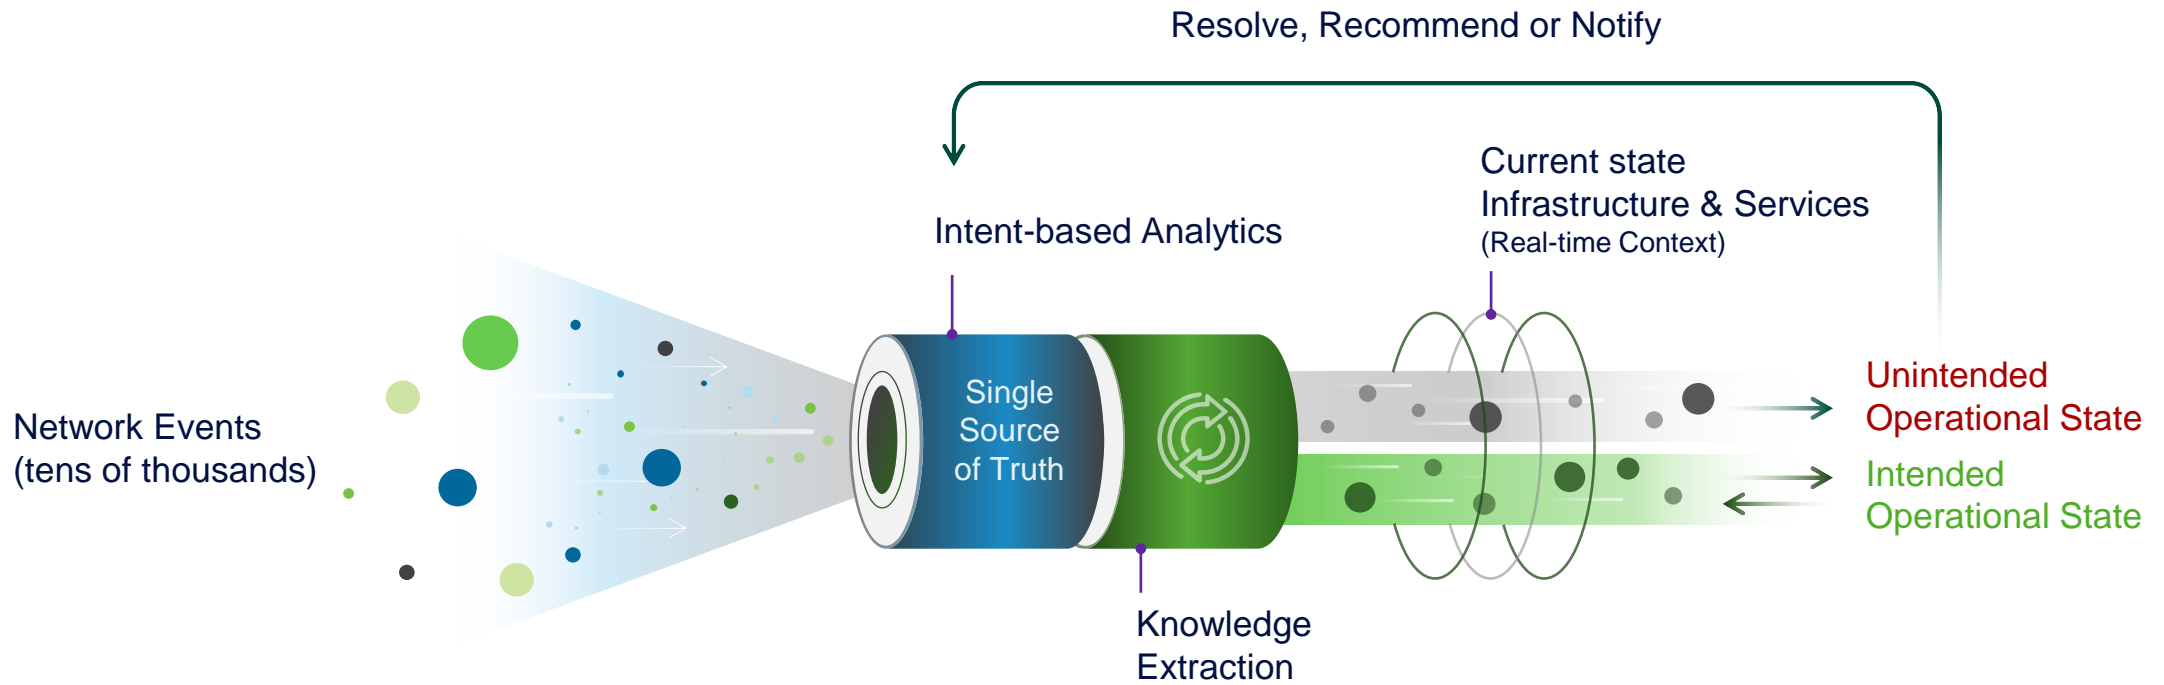


Step - 3
OPERATE

Closed Loop - Verify Health & Performance



Day 2+ Operations



Knowledge retention



Closed-loop validation



Eliminates outages

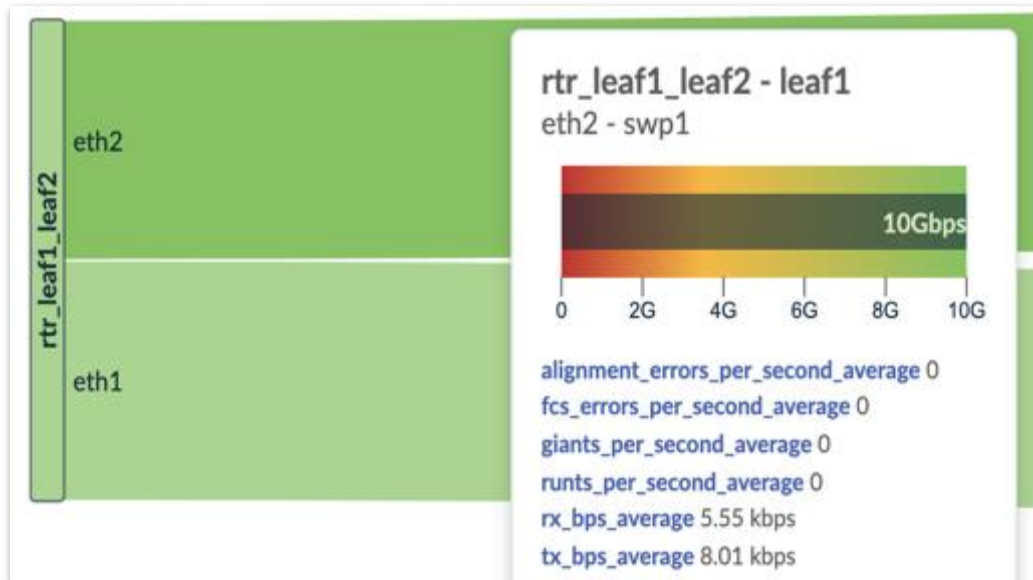


Lowers mean time to innocence

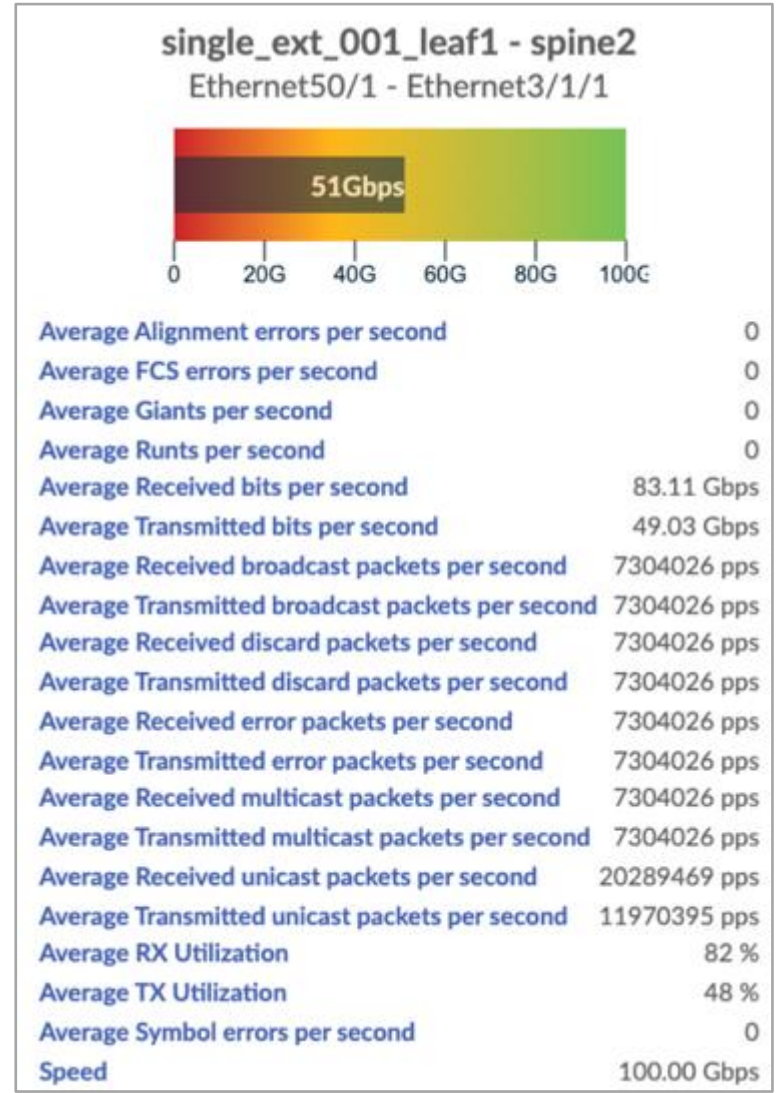
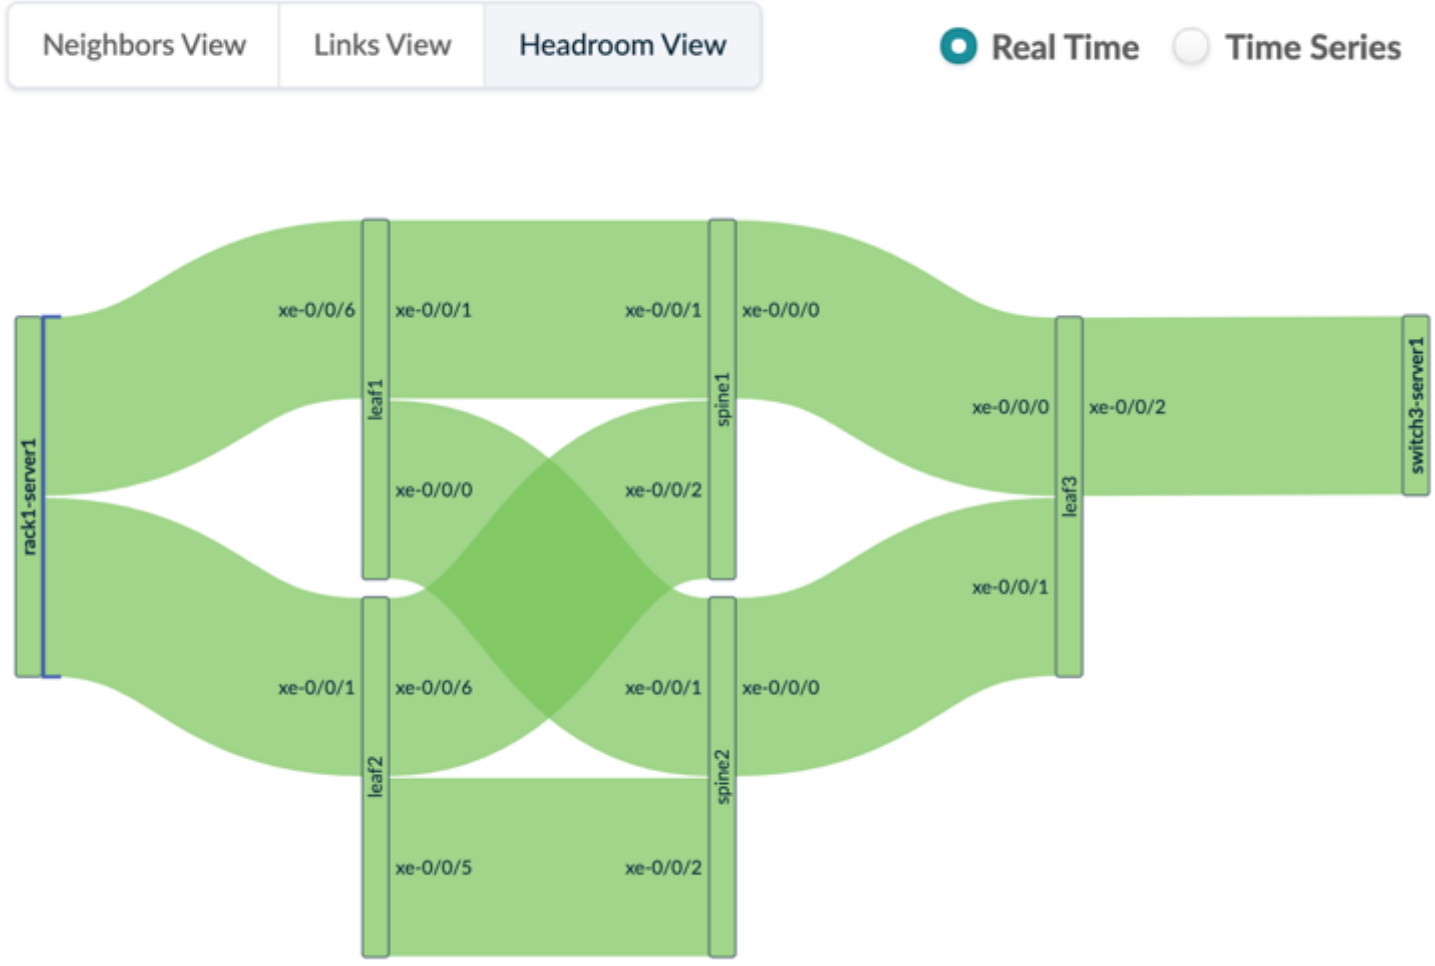


Change Management

Powerful Analytics



End to End Network Visibility



Topology View Cabeling

Blueprints

Staged Active

Physical Telemetry

Color: ok (green), violating intent (red), unintended (grey)

Link: present (solid), absent (dashed)

Interface: up (green), down (grey)

Show All Neighbors

Device	Interface	Link Status	Neighbor Device	Neighbor Interface	Neighbor Status
P18C26S01	Ethernet1/1	present	P18C27L01	Ethernet29/1	up
	Ethernet2/1	present	P18C27L02	Ethernet29/1	up
	Ethernet3/1	absent	P18C28L01	Ethernet29/1	down
	Ethernet4/1	present	P18C28L02	Ethernet29/1	up
	Ethernet5/1	present	P18C29L01	Ethernet29/1	up
	Ethernet6/1	present	P18C29L02	Ethernet29/1	up
	Ethernet7/1	present	P18C30L01	Ethernet29/1	up
	Ethernet8/1	present	P18C30L02	Ethernet29/1	up
	Ethernet9/1	present	P18C31L01	Ethernet29/1	up
	Ethernet10/1	present	P18C31L02	Ethernet29/1	up
	Ethernet11/1	present	P18C32L01	Ethernet29/1	up
	Ethernet12/1	present	P18C32L02	Ethernet29/1	up
	Ethernet13/1	present	P18C34L01	Ethernet29/1	up
	Ethernet14/1	present	P18C34L02	Ethernet29/1	up

Telemetry Device

- Probes
- All Services (14)
- Liveness
- Config
- Interface (1)
- Cabling (1)
- BGP (2)
- Route (10)
- Hostname

A photograph of a server room with rows of server racks. A vertical green bar is overlaid on the left side of the image. The text "Demo in 1 minute" is centered in white. The server racks are silver and filled with server units. The floor is light-colored and the ceiling has a grid pattern.

Demo in 1 minute

[+ Create Blueprint](#)



Not_Hotdog

 20 anomalies

3 days ago

Summary


Spines: 2, Leafs: 4, L2 Servers: 8

Deployment status: 

Recognition

Gartner

2020 Magic Quadrant **LEADER**
Data Center and Cloud Networking



Forrester

“The vendor has a **strong portfolio of hardware and software**, with feature depth and excellent automation capabilities. **As a result, the vendor can meet the technical needs of nearly all use cases in this market.**”

Juniper Networks recognizes that customers need more than technology and products. **Customers that want to automate the entire network** with a consistent OS, from the data center to the business edge, **should look to Juniper Networks.**”



apstra

Intent-based Networking Systems

Gartner Cool vendor in enterprise networking

IBNS can reduce network infra. delivery times by 50%-90%.



Best of VMworld winner

IBNS can reduce the number and duration of outages by at least 50%. - Gartner

Summary



Scale, Performance and Function

“Apstra's system gives us a single view into the entire media production network - regardless of our devices, switches, operating systems or vendors.”
- Bloomberg

Bloomberg



Ease of Use and Time Savings

“We went from POC to production in 1 weeks.”
- BeElastic

 **Beelastic**



TCO Savings

“In the previous solution there were 8 people managing the network, that number went down to 2.”
- Global 500 Manufacturer

YAHOO!
JAPAN



Self-driving Network

“Our success depends on rapidly introducing new services, quickly scaling to new demand and delivering the assured experiences. Juniper and Apstra joining forces as one company helps T-Systems”
- T-Systems

T-Systems



THANK YOU

JUNIPER NETWORKS | Engineering
Simplicity

Use Case EPFL



Vertical: University

Country: Switzerland

Incumbent: Cisco 3-Tier architecture

Requirements: Building a private Cloud for the complete organization to avoid move of workloads to public cloud. Being as much flexible as possible in regards to their different workloads (mainly HPC, big storage, big CPU)

POC: 8 weeks onsite

Solution: Switching: DELL
NOS: SONiC 3.3.

Size: 500 racks within 3 DC`s in a 5-stage CLOS

Use Case T-Systems



Vertical: Service Provider

Country: Germany – global roll out

Incumbent: Cisco 3-Tier architecture, Cisco ACI, Brocade, Juniper

Requirements: One central mgmt tool for all DC`s globally and for all vendors they have deployed.
Supporting at least Cisco, Juniper, SONiC
Supporting different workloads like shared and dedicated environments with different access profiles

POC: 12 weeks onsite

Solution: Switching: Cisco, Juniper, - SONiC in near future
NOS: NX-OS and JunOS

Size: approx. 2.000 devices within 17 DC`s in 3-stage and 5-stage CLOS

Use Case Raiffeisen IT



Vertical: Financial

Country: Austria

Incumbent: Cisco 3-Tier architecture

Requirements: Raiffeisen IT wanted to build a new DC to replace the old Cisco devices (EOL). They chose the central mgmt tool first and ran a RfP across all vendors to get the best solution and price.

POC: 4 weeks in CloudLabs

Solution: Switching: Juniper
NOS: JunOS

Size: 350 devices across 2 DC`s with DCI

Use Case BeElastic



Vertical: MSP

Country: Switzerland

Incumbent: Mellanox with Cumulus, mgmt with Ansible

Requirements: Moving from a classical L2 network to a new EVPN-VxLAN environment is very complex. BeElastic wanted a central mgmt tool for the new roll out. The self development effort with ANSIBLE was too complex and BeElastic chose for Apstra

POC: 2 weeks remote

Solution: Switching: NVIDIA / Mellanox
NOS: Cumulus

Size: 56 devices

Use Case Wortmann

WORTMANN AG

Vertical: Service Delivery MSP

Country: Germany

Incumbent: Arista

Requirements: Moving to EVPN-VxLAN Wortmann wanted to automate the operations as much as possible. Arista CloudVision was too much manual work and not out of the box. Installation with Apstra was done within 1 day

POC: 3 weeks remote

Solution: Switching: Arista
NOS: EOS

Size: 250 devices

Use Case advania

Vertical: MSP

Country: Iceland and Nordics

Incumbent: Cisco 3-Tier architecture

Requirements: Advania wanted to build new services based on EVPN-VxLAN. Their preferred vendor for Servers is DELL. Apstra showcased how to automate and operate the complete environment on SONiC.

POC: 2 weeks remote

Solution: Switching: DELL
NOS: SONiC

Size: 500 devices across 5 DC`s, 3-stage CLOS each with DCI

Use Case UMB



Vertical: MSP

Country: Switzerland

Incumbent: Cisco 3-Tier architecture

Requirements: UMB need to build a new private cloud to establish new end user services. For that UMB decided very early in the project, most of their daily tasks should be automated. After considering several tools UMB decided for Apstra due to the vendor agnostic approach and best solution

POC: 1 week onsite

Solution: Switching: Juniper
NOS: JunOS

Size: 120 devices in 2 DC`s with DCI

Yahoo Japan Corporation



Vertical: Media and Entertainment

Country: Japan

Incumbent: Cisco, Brocade, Juniper

Requirements: Needed to operate and manage a large scale of IP Clos fabric networks with a small number of staff. And needed the multivendor support for their data center environment.

URL : <https://www.juniper.net/content/dam/www/assets/case-studies/us/en/2021/yahoo-japan.pdf>

POC: 6 months

Solution: Switching: Arista, EdgeCore, Cisco
NOS: EOS, Cumulus, NXOS

Size: 600+ devices in 3 DC`s in 3-stage and 5 stage CLOSS