Juniper Networks | Driven by Experience™

White Paper

ELIMINATING THE PAIN OF DATA CENTER MIGRATION

Relocate critical applications quickly and reliably with Juniper Apstra

TABLE OF CONTENTS

EXECUTIVE SUMMARY

Migrating a data center is often fraught with headaches. But it doesn't have to be that way. Juniper® Apstra lets network architects easily pre-stage and validate network designs and configurations without requiring any physical or virtual infrastructure resources. By applying this simple and reliable approach, data center network teams can significantly reduce migration time, cost, and risks.

This white paper explains how Apstra eases the pain of data center migration. It also explores specific use cases to demonstrate the common migration issues that Apstra addresses.

Introduction

Data center migrations are inevitable. New applications, evolving performance needs, and updated reliability requirements—among other changes—typically require a greater number of pods and sites, including those designed for disaster recovery. The number of data centers in use, whether through natural IT growth or as the result of an acquisition, may increase. Conversely, as data centers age and are decommissioned, or as the business divests itself of certain IT requirements or applications, the number of data centers may shrink. This natural ebb and flow of assets means IT managers should always be prepared for their next data center project while simultaneously seeking opportunities to consolidate.

When a migration is required, relocating critical applications and data quickly, reliably, and with minimal or no disruption is a top priority. It's a daunting task for network architects to add migration and/or conversion strategies onto their everyday responsibilities of designing, deploying, and operating the new data center infrastructure.

Juniper Apstra uses intent-based networking to ease the pain of data center migrations for both business leaders and network architects by reducing the time, risk, and cost associated with such actions.

Intent-Based Networking with Juniper Apstra

As a powerful automation and abstraction software solution for data center network infrastructures, Apstra saves hundreds of person-hours and dramatically reduces operational anomalies during a migration, particularly in the design phase and critical change-management window. Its intent-based networking approach helps network architects and operators automate designing, building, deploying, and validating data center networks. Apstra translates high-level business requirements (called "intent") into a fully operational data center network environment.

Intent-based networking represents a technological shift in the way the full network service life cycle (Days 0/1/2) is managed based on business requirements. It represents a consistent and streamlined operational model for architects and operators, with continuous validations at each step regardless of the switching platform or operating system. Users can choose to run a homogeneous network with a traditional networking vendor, or gradually introduce disaggregated open-networking and open-source options.

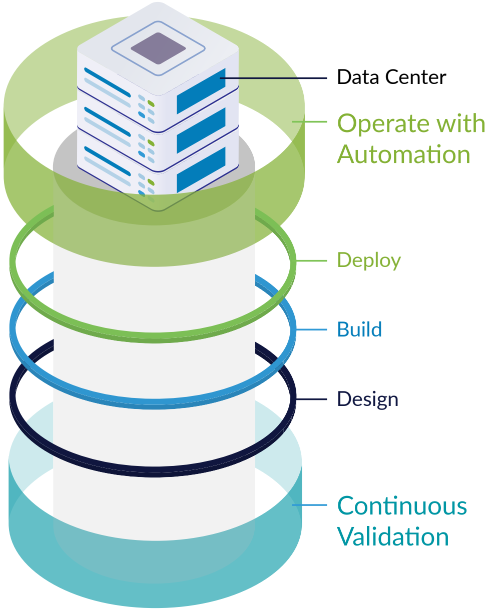

Figure 1: Intent-based networking automates and validates the full network life cycle.

Apstra helps network operators reduce CapEx and OpEx while decreasing the risks inherent in networking by providing a fully tested and validated design pattern. The definition of network service is key to the notion of intent. As a result, Apstra can effectively understand all known and unknown network conditions that directly impact business service levels.

The Apstra intent-based networking architecture includes the following benefits:

- Real-time preconditioned validation
- Real-time expectations validation
- Service validation
- Reliable change management through real-time queryable intent and operational context

Additionally, Apstra offers unparalleled best-in-class features to ensure the data center network is reliable and easy to operate.

Automated Troubleshooting with Intent-Based Analytics and Root-Cause Identification

Intent-based analytics (IBA) is an Apstra feature that automates complex troubleshooting, allowing you to gain more knowledge of your infrastructure and helping operators deal with operational status changes in their infrastructure. IBA extracts knowledge about the network state and performs validations on all resources, all the time, in real time. This "closed-loop" capability is always in sync with the network, even in the presence of change.

Root-cause identification (RCI) is a subcomponent of IBA that classifies conditions into actionable root causes, separating what is actionable/important (signal) from what is not (noise).

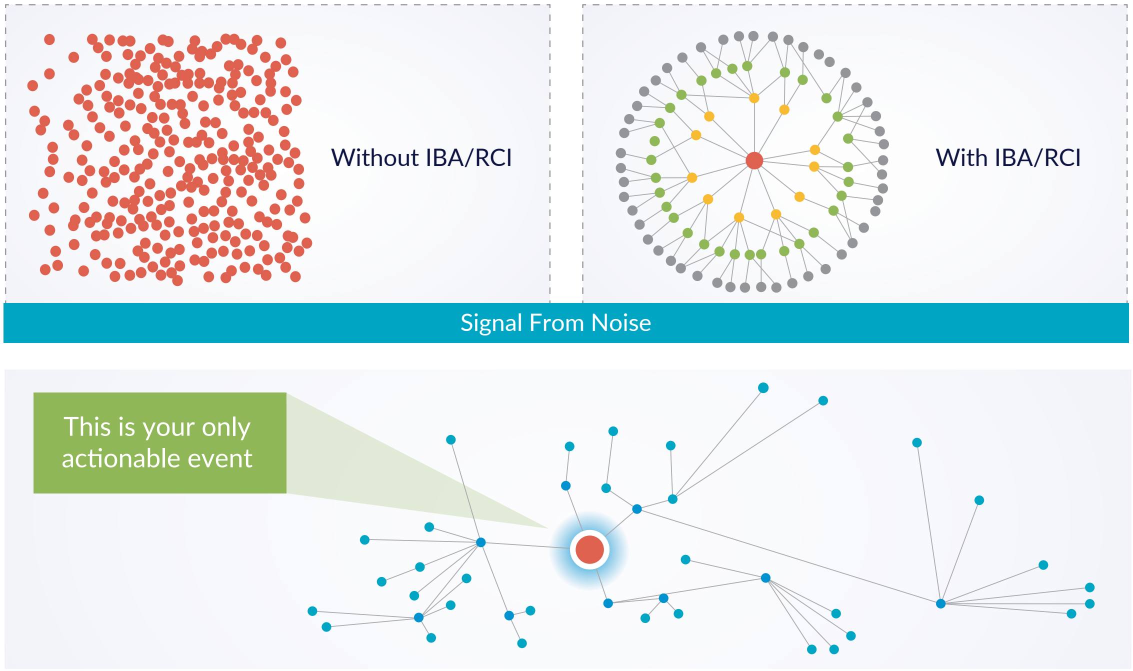

Figure 2: IBA separates actionable conditions from noise.

Group-Based Policy Design System

Apstra manages network security and workload isolation through its group-based policy (GBP) design system. This feature allows a network architect to create policies that are decoupled from enforcement mechanisms and to specify the intent in an implementation-independent way.

GBP simplifies and normalizes network device syntax across data center devices, implements firewall filters, and ensures that security configuration is easy and secure. For migration use cases, users can design/redesign and pre-stage security policies and perform policy validations in advance to reduce outages and application impact. This capability significantly reduces the time spent verifying, optimizing, and correcting security policies in maintenance windows.

GBP also performs policy validation checks and supports conflict identification and remediation (automated and manual) for overlapping or conflicting security policies. Apstra automatically optimizes the rendered policies, improving the size and device ternary content-addressable memory (TCAM) of firewall filter rules.

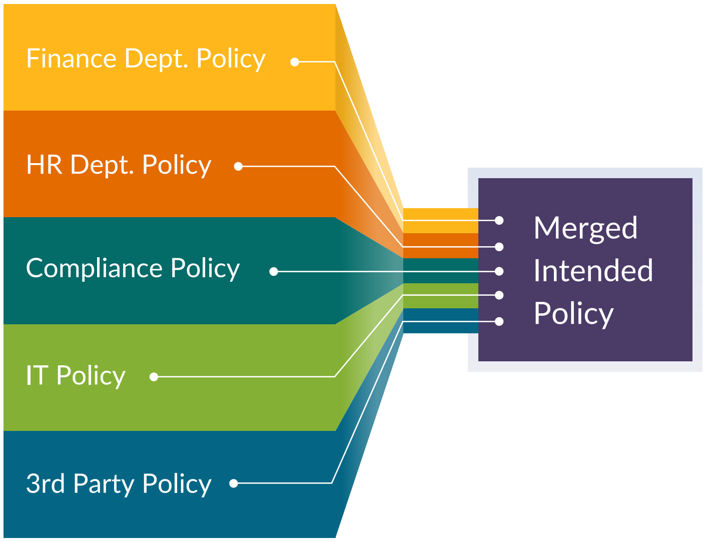

Figure 3: Apstra's group-based policy system optimizes policies.

Apstra EVPN-VXLAN Reference Design

Business leaders are routinely migrating from legacy three-tier architectures to IP Clos fabrics, also referred to as "spine-and-leaf" configurations. Why the need for this new architecture? Modern enterprise architects are responsible for delivering highly scalable, highly redundant fabrics to support the connection of large numbers of servers in an application hosting environment. The simplicity, scale, and deterministic design of the IP Clos fabric is particularly attractive for this application. Apstra provides validated IP Clos reference designs that employ standards-based protocols.

Apstra uses BGP-EVPN for a control plane and Virtual Extensible LAN (VXLAN) for Layer 2 application connectivity. These protocols were selected for their reliability, scaling attributes, and broad support in vendor hardware and software. Users can deploy three-stage and five-stage IP Clos fabrics with the option of implementing network virtualization overlays (NVOs) on top of the general IP network.

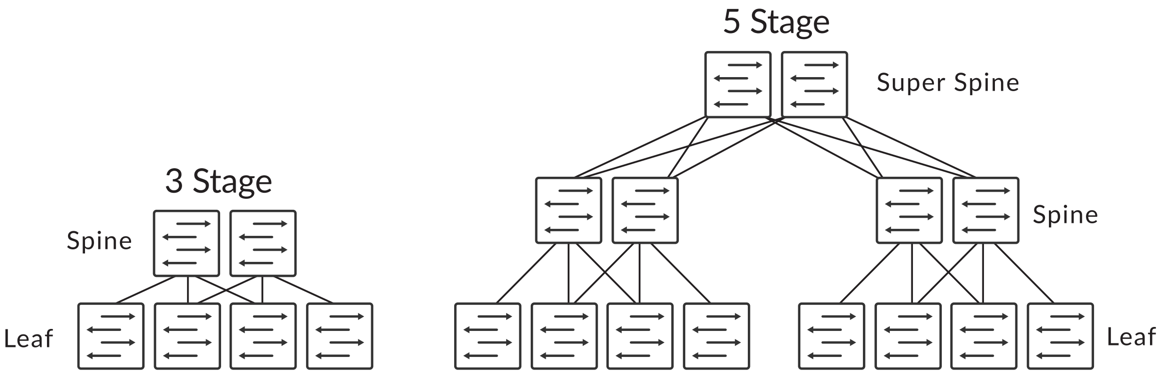

Figure 4: Apstra supports both three-stage and five-stage IP Clos fabrics.

The advantages of using a BGP-EVPN and VXLAN fabric include:

- **Interoperability**: A BGP-EVPN and VXLAN fabric leverages industry-standard protocols (RFC-7432).
- **Multitenancy**: A BGP-EVPN and VXLAN fabric supports L2 and L3 VPN.
- **Resiliency**: A BGP-EVPN and VXLAN fabric supports smaller fault zones, allowing architects to build networks with many small switches versus a few large chassis.
- **Efficient resource utilization**: A BGP-EVPN and VXLAN fabric uses equal-cost multipath (ECMP), Address Resolution Protocol (ARP) suppression, and Anycast Gateway to optimize traffic.
- **L2 mobility**: A BGP-EVPN and VXLAN fabric provides host mobility over VXLAN tunnels.

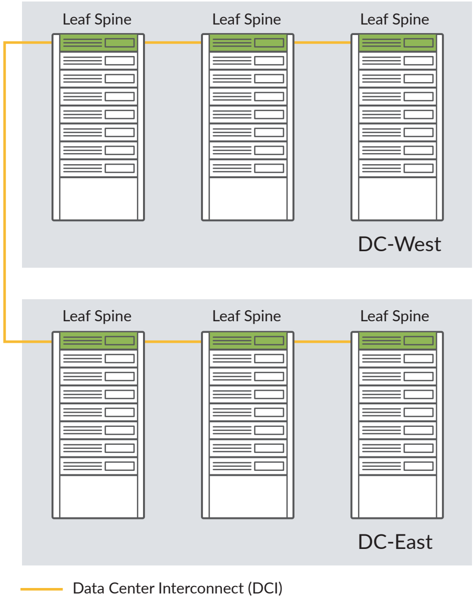

Figure 5: Apstra manages all IP fabric egress points when connecting multiple data centers.

Superior Migration Advantages

Prior to beginning a migration, architects can design the precise network and services they require using the Apstra graphical interface. Concurrently, Apstra validates the design and auto-generates precise configurations and service expectations based on the exact blueprint, network models, and software. This novel approach gives network teams responsible for the design and migration activity an unprecedented level of assurance to accurately stage physical and logical configurations ahead of maintenance windows.

Network operators can develop:

- Multitenancy isolation plans
- Subnet allocation schemas
- L2 application segments
- L3 routing domains
- External connectivity
- Security policies

Remarkably, Apstra can do this without the need for any physical or virtual network resources, saving thousands to millions of dollars in testing infrastructure investments.

Why is pre-staging so important? Because pre-staging:

- Removes deployment risks by accurately and reliably building and validating an entire design that works the first time
- Eliminates the CapEx and OpEx required for physical or virtual testing infrastructure
- Enables validation designs and redesigns without impact
- Reduces time spent in maintenance windows
- Allows for the verification and optimization of the security list
- Removes conflicting security policies
- Reverts full data center state in seconds, if required
- Ensures that auto-produce is always in sync with documentation
- Prepares real-time analytics that are specialized to migration-specific service-level agreements (SLAs)

The result is a rapid, low-risk migration to or from an interoperable data center using best practices and protocols.

Common Migration Issues and How Apstra Addresses Them

Use Case 1: Legacy Brownfield

Many customers are looking to migrate from an older hierarchical three-tier network (core, aggregation, access) that relies heavily on Layer 2 and Spanning Tree Protocol (STP) to modern spine-and-leaf (or Clos) architectures. Apstra can ensure the new fabric is designed and built—with the configuration validated—before the migration has even begun. This improves the speed of implementation and eliminates Day 0 implementation faults.

A well-planned, multi-phased strategy is vital to a successful migration. Each phase ensures devices can communicate internally and externally at L2 and L3 with minimal disruption.

A successful migration process follows these steps:

1. Build and prepare a new network.
2. Move all devices in a given L2 domain, one at a time, to the new fabric.
3. Migrate the middleware (L4-L7) services: firewall, load balancer, and so on.
4. Move the L3 default gateway and related security policy.

Step 1: Build and Prepare the New Network

Start with the following stages:

- Pre-stage your new data center network in Apstra to better understand your ultimate design and how it will integrate with your virtual infrastructure and external services
- Stand up a new three-stage or five-stage BGP-EVPN and VXLAN fabric managed by Apstra
- Establish a routed Layer 3 connection to the upstream external router for external connectivity
- Establish a Layer 2 connection to the legacy/brownfield implementation
- Create the security zones required to support virtual routing and forwarding (VRF) in a brownfield network.
- Create virtual networks to match L2 segments in the brownfield network
- Extend the virtual networks from a leaf device, or multichassis link aggregation group (MC-LAG) pair for redundancy, to the interface(s) connecting the brownfield to the greenfield networks

Requirement 1: External L3 Connection

The L3 connection will be the new exit point for the new eBGP/EVPN/VXLAN data center. This will be utilized by:

- Applications, whether new or existing, that are fully migrated for communicating to the outside world
- Applications not fully migrated between the old and new environment where the default gateway remains on the old environment

This routed connection could attach to the upstream router, or WAN device, for connectivity.

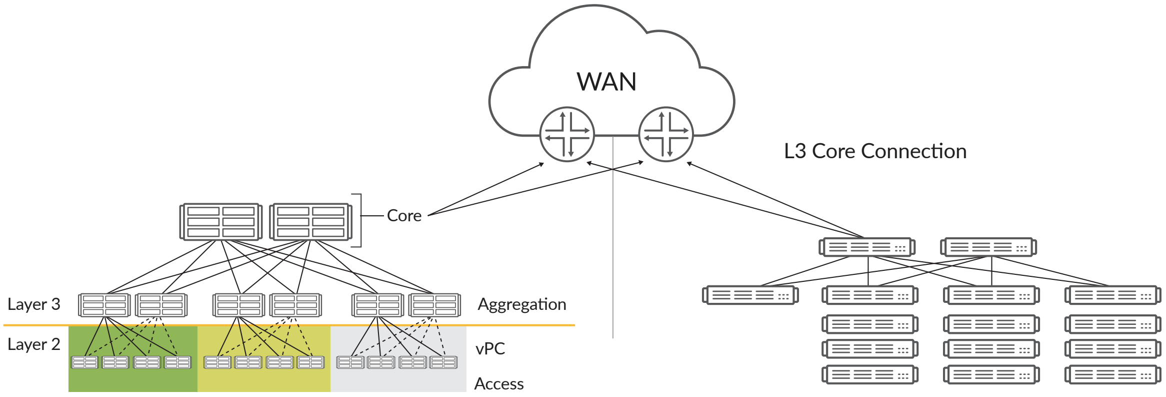

Figure 6: Example of an external L3 connection.

Requirement 2: L2 Connection Between the Old Network and the New BGP-EVPN VXLAN Fabric

For applications requiring a L2 connection for workload migration, the recommendation is to create a connection between a dedicated pair of MC-LAG devices on each fabric for redundancy. Each vendor has its own proprietary implementation: virtual path connection (VPC), MLAG, MC-LAG, CLAG, VLT, and others. The most typical connection method used when migrating from older three-tier networks is VLAN-to-VLAN between the fabrics. However, connecting both fabrics using BGP-EVPN and VXLAN is an option if the old environment has edge devices capable of supporting it. This option is covered in the Data Center Interconnect section.

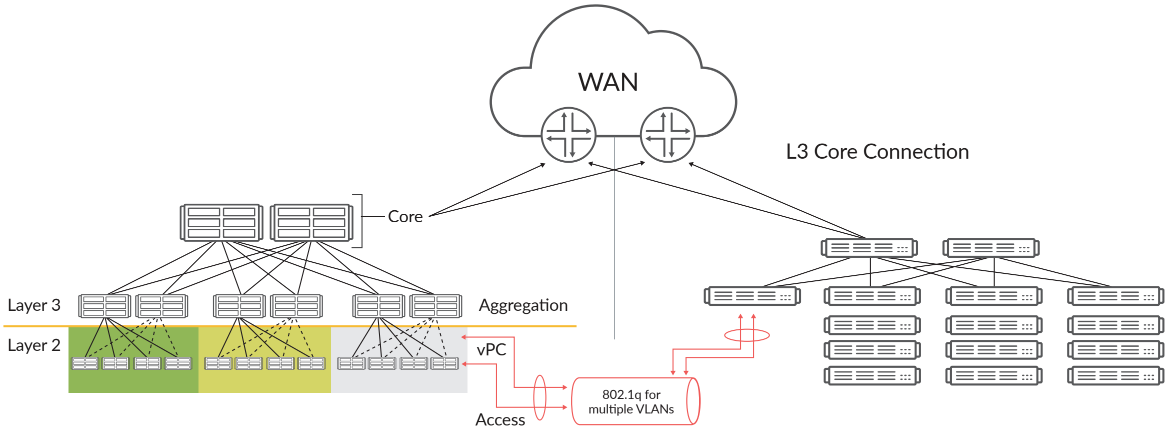

Figure 7: Example of an L2 connection between an old network and a new BGP-EVPN VXLAN fabric.

Requirement 3: L2 Loop Avoidance

EVPN does not provide any integration with STP and doesn't forward bridge protocol data units (BPDUs), so there is no STP blocking. In addition to using MC-LAG for a loop-free redundant connection, BPDU guard and root guard should be used to further protect against L2 loops. The old fabric should remain the root bridge. The L2 interface on the Apstra managed fabric will employ BPDU guard using configlets based on the vendor details chosen.

Step 2: Move All Devices in a Given L2 Domain to the New Fabric

The recommendation is to migrate all devices in an L2 domain before moving the gateway, as the gateway cannot simultaneously be active in both the legacy and the new Apstra-managed data center. The migration should be done one network or domain at a time versus an all-at-once approach.

During this phase of the migration, continue to use the original default gateway in the old fabric until all devices in the given L2 domain have migrated to the new fabric. Most, if not all, current IP addressing will remain the same during the migration, including the default gateway address for each tenant.

Step 3: Migrate the Middleware (L4-L7) Services

Move firewalls and other middleware devices after all the devices are relocated to the Apstra-managed fabric. The procedure for relocating these devices depends on their capabilities, whether they are L2 or L3 connected, and whether they are in an active/active or active/passive configuration.

Active/Standby

Migrating or relocating active/standby devices requires multiple steps. First, relocate the standby device to the Apstra managed fabric. This step doesn't disrupt or change the application traffic. The heartbeats (keepalive) messages between the active and standby devices will traverse the L2 connection above.

Once the QA team has tested and certified the applications are working, deactivate the device in the old network and make the device in the Apstra managed fabric the active one. This state change eliminates unnecessary cross-fabric traffic for all locally attached nodes in the new data center. Once QA validates the applications are performing as expected, remove and re-home the remaining device to the new data center.

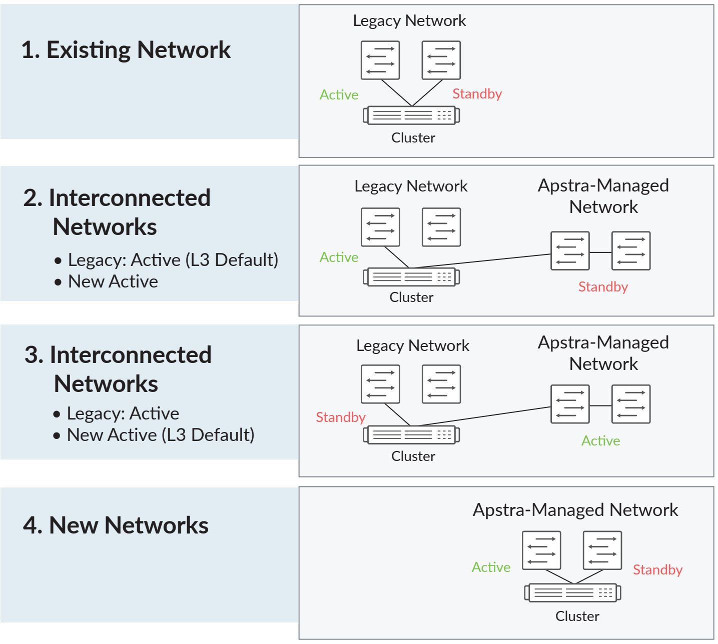

Figure 8: Example of an active/standby deployment.

Active/Active

Active/active deployments are connected to the new BGP-EVPN fabric using BGP as the dynamic routing protocol to provide deterministic next- hop selection and load balancing. The BGP peering, IP addressing, and route policy are pre-staged and validated before the maintenance windows in Apstra. For redundancy purposes, connect the L4-L7 devices to a pair of leaf switches in the EVPN fabric running an MC-LAG protocol.

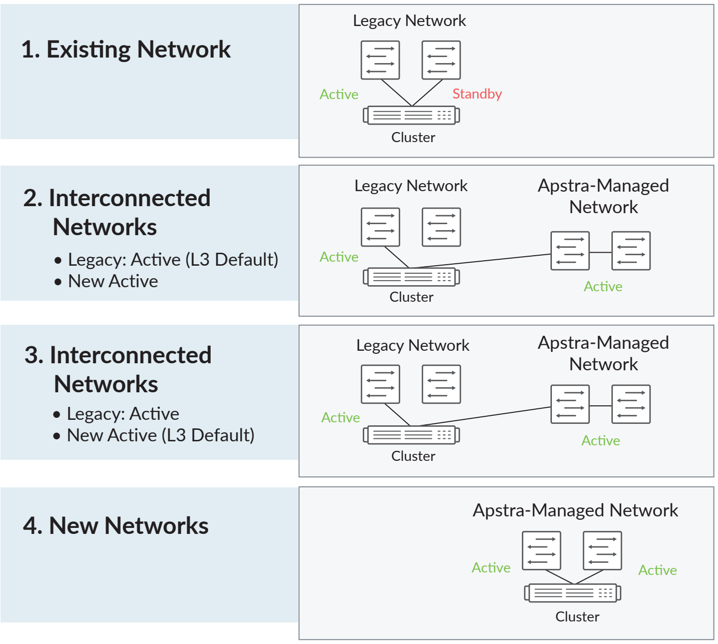

Figure 9: Example of an active/active deployment.

Step 4: Move the L3 Default Gateway

Once all (or an acceptable critical mass of) devices in the L2 domain have migrated to the new Apstra-managed fabric, it's time to move the default gateway to the new fabric. Apstra uses the Anycast Gateway feature to optimize the internal traffic, locating the first hop default gateway to the local leaf switches and removing inefficient cross-fabric trombone routing.

Enabling the L3 endpoint and Anycast Gateway feature in Apstra is a simple one-checkbox, one-field operation. During the maintenance windows, deactivate the legacy network's default gateway and make sure the subnet's default gateway only resides in Apstra.

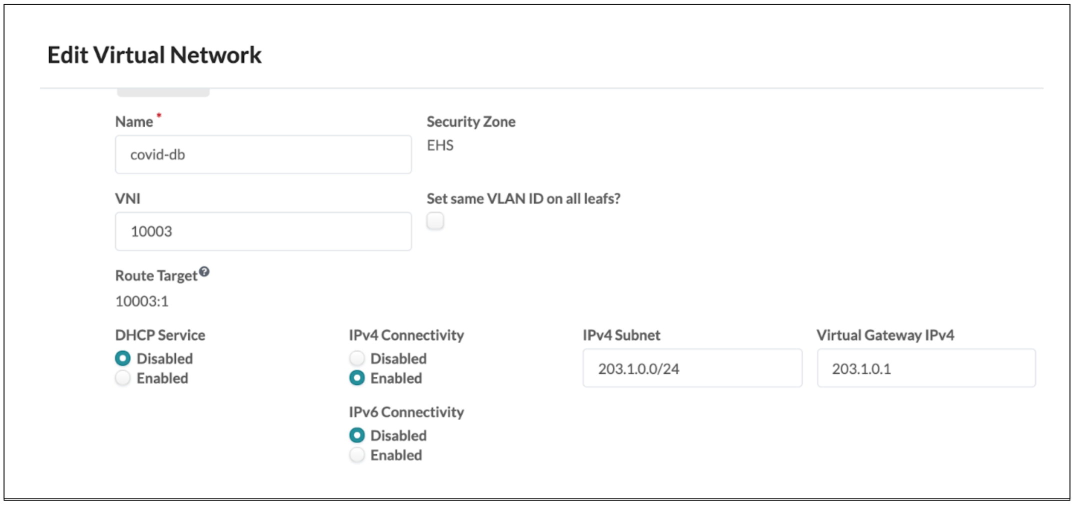

Edit Virtual Network

Name* covid-db

Security Zone EHS

VNI 10003

Set same VLAN ID on all leafs? □

Route Target 10003:1

DHCP Service ● Disabled ○ Enabled

IPv4 Connectivity ○ Disabled ● Enabled

IPv4 Subnet 203.1.0.0/24

Virtual Gateway IPv4 203.1.0.1

IPv6 Connectivity ● Disabled ○ Enabled

Figure 10: Apstra makes it easy to enable the L3 and Anycast Gateway feature.

Security Access Lists

During the maintenance windows, Apstra makes sure that the network does not enter an open (free-flowing traffic) state while the rules are being provisioned. To minimize the potential impact on policy deployments' existing traffic flows, Apstra performs an incremental firewall filter deployment process to deploy policy changes at each enforcement point.

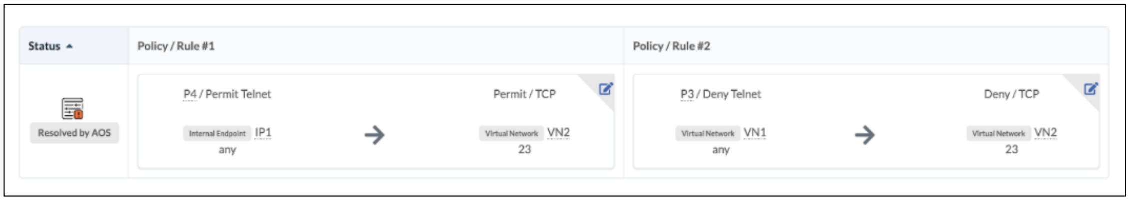

Status ▲	Policy / Rule #1	Policy / Rule #2
Resolved by AOS	P4 / Permit Telnet → Permit / TCP Internal Endpoint IP1 any → Virtual Network VN2 23	P3 / Deny Telnet → Deny / TCP Virtual Network VN1 any → Virtual Network VN2 23

Figure 11: Apstra deploys policy changes incrementally.

Use Case 2: Relocation or Hybrid Cloud

Applications frequently have to be moved between two geographically separated data centers or from a private data center to a cloud-hosted infrastructure—without disrupting business. In the past, such relocations were driven by business continuity, disaster recovery, or continuity-of-operations requirements.

With the rise of highly virtualized software-defined data centers (SDDC), cloud computing, and more recently edge computing, other likely scenarios have arisen, including:

- Collocation expansion: Share compute and storage resources with collocation data center facilities
- Resource pooling: Share and shift applications between data centers or public cloud resources to increase efficiency or improve end-user experience
- Rapid scalability: Expand capacity from a resource-limited location to another facility or data center
- Legacy migration: Gracefully move applications and data from older and inefficient equipment and architectures to more efficient, higher performing, and more cost-effective architectures

Apstra is uniquely positioned to give businesses the flexibility to extend their network services and security in a consistent and uncomplicated workflow to any number of locations, private or public, based on business needs.

Data Center Interconnect

Data Center Interconnect (DCI) technology is often used as a building block to connect geographically separated data centers at Layer 2 for disaster recovery, continuity of operations, and business continuity. Apstra allows users to deploy and manage a vendor-inclusive DCI solution that is simple, flexible, and intent based.

Using standards-based MP-BGP EVPN with VXLAN, the Apstra DCI feature is open and flexible, with three different deployment models:

- DCI using over-the-top
- DCI using gateways
- DCI using autonomous system border router

Public Cloud Connections

Users have an abundance of network connectivity options for connecting their on-premises and public cloud networks. These services range from dedicated interconnections to native VPNs and third-party VPN services. The most popular choice is to leverage internet connections coupled with VPNs for their low cost and high availability.

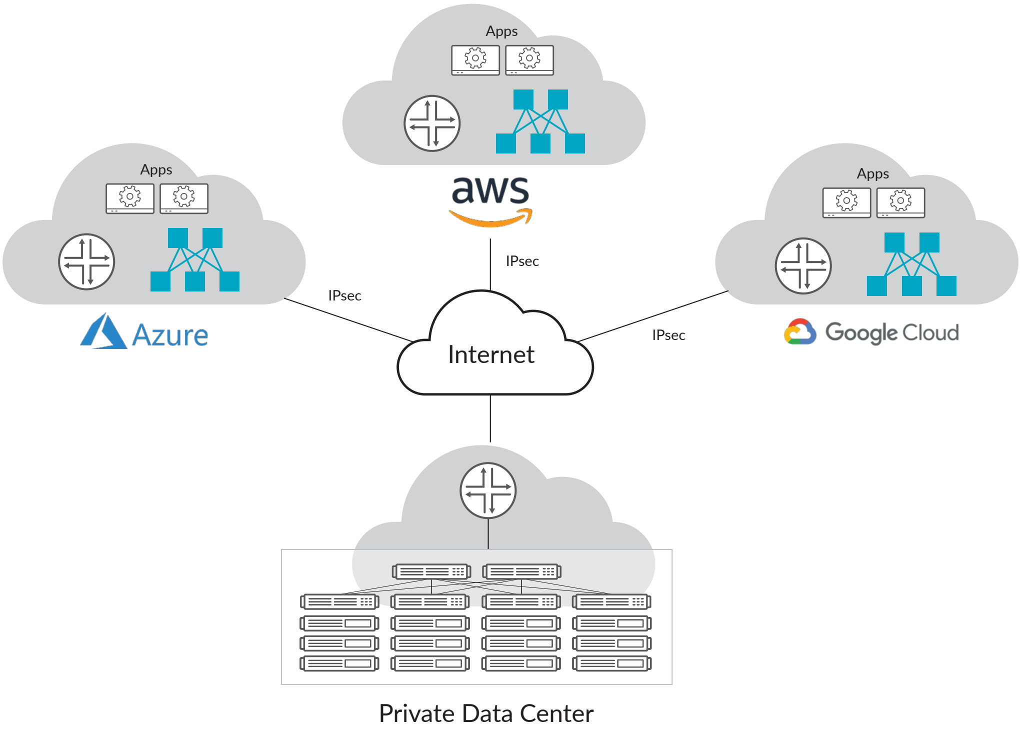

Figure 12: Users have choices for connecting on-premises and public cloud networks.

Apstra provides an elegant way to connect a private data center to a virtualized data center using the cloud provider of choice and a standards-based solution. Secure IPsec tunnels to remote networks can be established using native options or numerous third-party gateways. Apstra then automates the provisioning of BGP EVPN for reachability and VXLAN tunnels from existing leaf nodes to virtualized cloud-hosted devices for workload mobility.

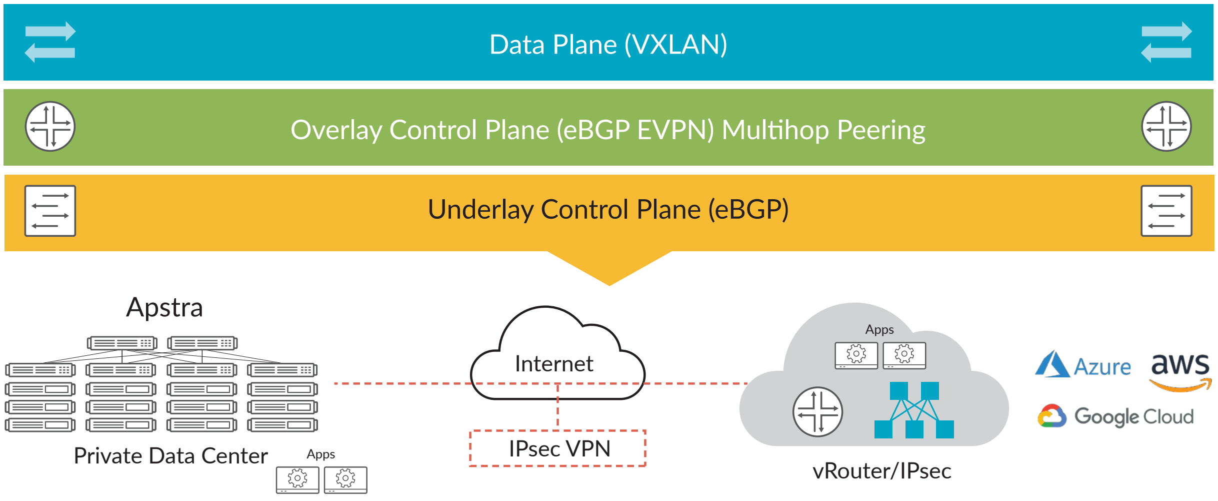

Figure 13: Apstra automates provisioning of BGP EVPN and VXLAN tunnels.

Use Case 3: Network Conversions

While not an application or data migration use case, network conversions are increasingly in demand. Apstra can minimize the disruptions caused by this use case by designing, building, deploying, and validating a precise replica of the existing topology down to individual tenants, hardware, and ports offline. There is no need for extra equipment or for the user to re-architect or reverse engineer an equivalent replacement configuration. All application addressing will remain the same.

The Apstra Reference Design supports legacy L2 connectivity models as well as homogenized containerized L3 options. Both models are fully supported by industry-standard protocols and services, ensuring that an Apstra system is not dependent on any particular vendor or proprietary feature. Businesses may be concerned about ongoing support, costs, or stability issues, or they may simply want to divest from a particular vendor's solution. These pre-existing network fabrics could be an existing BGP-EVPN fabric or a proprietary software-defined network solution.

As in other use cases, a well-planned, multi-phased strategy is vital to ensure a successful conversion. However, unlike other use cases, a no-downtime guarantee isn't always possible with conversions since conversions require rebooting or, in some cases, re-imaging the network switches in place. The existing application and network architecture will dictate the amount of network disruption. If the hosted applications are designed for failure and built with application-level resiliency, or if the IP Clos fabric has a redundancy design to support legacy L2 applications, the disruption can be minimal.

Automating device initialization and installing new OS images can be accomplished without human intervention using Apstra's multivendor zero-touch provisioning (ZTP) solution. ZTP saves time and reduces critical errors and maintenance window times.

One approach for migrating an existing three-stage network is to move half the spines and leafs to a parallel fabric under Apstra management, followed by the remaining switches. Using the assumption that the network is servicing resilient applications, this approach decreases the overall cross-sectional bandwidth of the initial network by half but avoids any total loss of service.

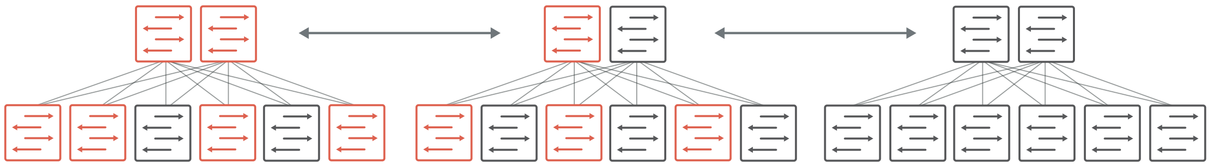

Figure 14: Migrating half of spines and leafs prevents total loss of service.

Conclusion

Apstra offers unprecedented advantages to organizations planning business-critical data and workload migrations that enable them to achieve business and technical goals faster than ever before.

Users can pre-stage and validate network designs and configurations to eliminate risk throughout the entire process. Without any dependency on physical or virtualized infrastructure, Apstra reduces planning times, as well as OpEx and CapEx. Apstra's ability to prepare multitenancy, subnet allocation, L2 application connectivity, L3 routing and service instantiation, external routing connectivity, and security policies in advance dramatically minimizes risk and disruption during critical change windows. The result is a simple, reliable approach that takes a tenth of the time of traditional migrations.

About Juniper Networks

At Juniper Networks, we are dedicated to dramatically simplifying network operations and driving superior experiences for end users. Our solutions deliver industry-leading insight, automation, security and AI to drive real business results. We believe that powering connections will bring us closer together while empowering us all to solve the world's greatest challenges of well-being, sustainability and equality.

Corporate and Sales Headquarters
Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, CA 94089 USA
Phone: 888.JUNIPER (888.586.4737)
or +1.408.745.2000
www.juniper.net

APAC and EMEA Headquarters
Juniper Networks International B.V.
Boeing Avenue 240
1119 PZ Schiphol-Rijk
Amsterdam, The Netherlands
Phone: +31.207.125.700

Driven by Experience™