

THE ENTERPRISE DATA CENTER

EVPN-VXLAN-based IP fabric

Challenge

Despite widespread adoption of cloud-enabled applications, most enterprise data centers continue to host legacy applications as well. This means that, for the foreseeable future, modern enterprise data centers must be able to handle both types of applications.

Solution

Juniper's data center fabric solution, based on a VXLAN overlay with an EVPN control plane, offers an efficient, scalable way to interconnect multiple data centers, campuses, and public clouds, letting operators deploy consistent architectures independent of scale and end point.

Benefits

- Control plane-based L2/L3 information exchange
- Internet-grade resiliency—no single point of failure
- Efficient host mobility
- Multitenancy
- Open, nonproprietary solution
- Scalability at all network layers
- Faster convergence
- Secure and flexible architecture

Traditionally, data centers have used Layer 2 technologies such as Spanning Tree Protocol (STP) and multichassis link aggregation groups (MC-LAG) to connect compute and storage resources. While these architectures work for small and medium-sized data centers where services are limited to a single network and cater to traditional requirements, they are simply too rigid to support the scalability needs of virtualized, multitenant next-generation data centers spread across geographically dispersed public, private, and hybrid cloud environments.

The Challenge

Deploying, securing, and connecting data centers is a complex task. As they evolve to include scale-out multitenant networks, these data centers need a new architecture that decouples the underlay (physical) network from a tenant overlay network.

Security also poses a unique challenge. Modern enterprise data centers want security—which is no longer just a perimeter problem—to be embedded within their network architectures. Not just inside the data center, but through segmentation, tenant separation, and policies extended across the entire organization.

The Juniper Networks EVPN-VXLAN IP Fabric Solution

While old-school data centers used legacy applications requiring L2 connectivity, the current best practice is to build scalable, highly available data centers based on a Layer 3 IP fabric. To bridge this gap, next-generation data centers must be able to support L2 connectivity services on top of the L3 IP fabric. Ethernet VPN (EVPN) is the key.

Solution Components

A data center's physical underlay network is designed to provide an L3 IP fabric. Also known as a Clos network, it is the fabric's responsibility to provide unicast IP connectivity from any physical device (server, storage, router, or switch) to any other physical device. An ideal underlay network provides low-latency, nonblocking, high-bandwidth connectivity from any point in the network to any other point. IP fabrics can vary in size and scale; a typical solution uses two layers—spine and leaf—to form what is known as a three-stage IP fabric, where every leaf is connected to every spine device. As the fabric grows, it may evolve into a five-stage IP fabric where "super spines" are added for inter-pod communication.

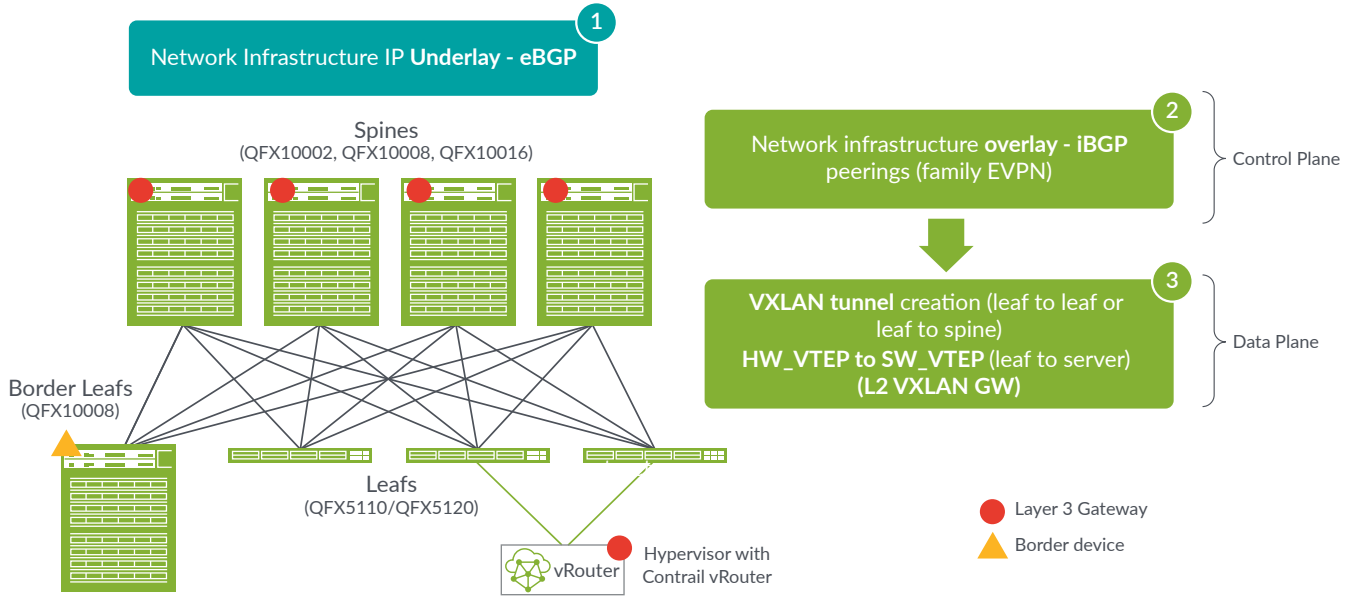


Figure 1: Main building blocks of an EVPN-VXLAN data center fabric

A network virtualization overlay—a virtual network transported over an IP underlay network—is a functional building block that enables multitenancy within a network. This allows you to share a single physical network across multiple tenants while isolating each tenant’s separate network traffic.

An EVPN-virtual extensible LAN (VXLAN)-based solution decouples the overlay network from the underlay using dependable and scalable Exterior Border Gateway Protocol (EBGP) as the routing protocol (see Figure 1). Each spine and leaf device is assigned its own autonomous system with a unique 32-bit autonomous system number, which is required to support EBGP. Other routing protocols, such as OSPF/ISIS, can also be used in the data center underlay network. Internal BGP (IBGP) is a routing protocol that exchanges reachability

information across an IP network. When combined with Multiprotocol BGP (MP-IBGP), IBGP allows EVPN to exchange reachability information with virtual tunnel endpoint (VTEP) devices. This capability is required to establish the inter-VTEP VXLAN tunnels used for overlay connectivity services.

The Juniper Networks EVPN-VXLAN Overlay Solution Choices

Juniper Networks EVPN-VXLAN-based IP fabric provides multiple overlay service options that let operators create virtualized L2 and L3 networks that satisfy both legacy and modern applications running on the IP underlay without introducing the complexity of MPLS (which was required in the past). Figure 2 shows the various EVPN-VXLAN overlay reference architectures that Juniper currently supports.

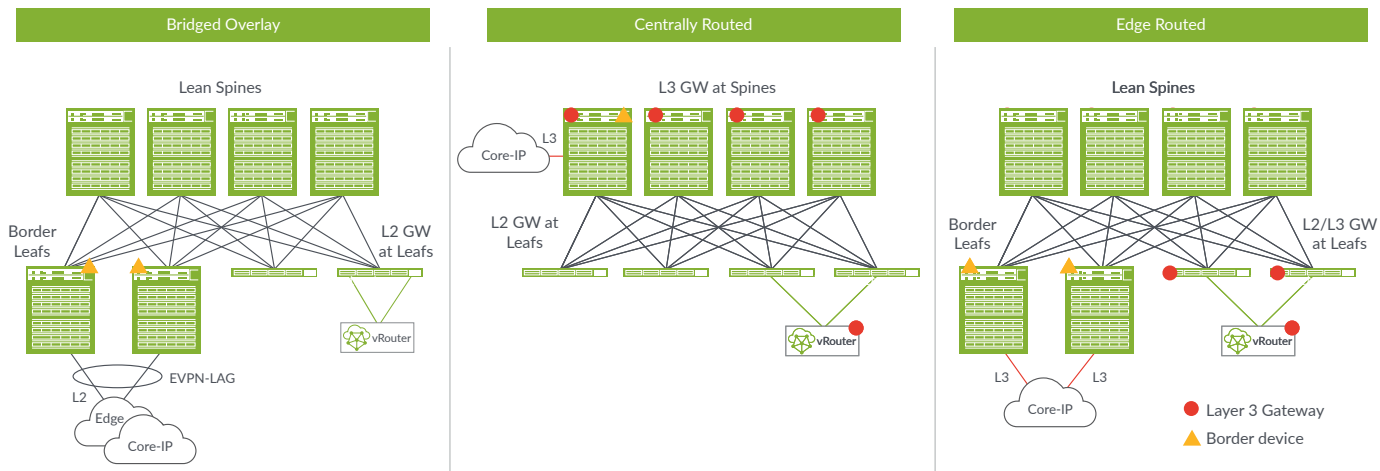


Figure 2: EVPN-VXLAN reference architectures in the data center

Bridged Overlays

A bridged overlay provides Ethernet bridging between leaf devices in an EVPN network, extending VLANs across VXLAN tunnels. With a bridged overlay approach, there is no need to migrate the IP gateway, which is managed by the external tenant.

Central Routing

With a centrally routed approach, routing occurs at a central gateway (the spine in this example). Traffic routed between hosts connected to the same leaf hairpins at this central gateway device. This is desirable when the majority of traffic is inter-rack or north-south. In this environment, tenant IP/VRF management is conducted centrally.

Edge Routing

Edge routing, which occurs at the edge access device where end systems are connected, ensures traffic between hosts is routed as close as possible to the end device at the leaf layer, or at the VRouter layer if a virtual machine (VM) is connected to the VRouter. Edge routing, which employs distributed tenant management, is best when traffic is primarily east-west and heavily segmented within the pod.

EVPN-VXLAN Enterprise Data Center Solution Features and Benefits

QFX Series Device Roles: Centrally Routed Architecture

Some key points to remember when using a centrally routed architecture:

- Centrally routed gateway devices should be feature-rich and support scale—for instance, the Juniper Networks QFX10000 line of switches.
- Centrally routed gateway devices can also serve as border gateways that advertise the prefix routes of the local tenancies to north-facing networks.
- Since centrally routed access devices don't require routing capabilities, Juniper Networks QFX5000 or QFX10000 switches can be used.

Table 1: QFX Series Device Roles: Centrally Routed Architecture

Device Roles	Edge Routed
Centrally Routed Access@ Server Leaf (L2 GW @Server Leaf)	QFX5100, QFX5110, QFX5200-32C, QFX10002-36Q/72Q, QFX5200-48Y, QFX5120-48Y
Centrally Routed_Gateway @Spine (L3 GW @Spine)	QFX10002-36Q/72Q, QFX10008/16
DCGW @ Leaf (Border Leaf)	QFX10002-36Q/72Q, QFX10008/16
DCGW@ Spine (Border Spine)	QFX10002-36Q/72Q, QFX10008/16
Null @ Spine (Lean Spine)	QFX5200, QFX5110, QFX10002-36Q/72Q, QFX5210-64C, QFX10002-60C

QFX Series Device Roles: Edge Routed Architecture

Some key points to remember when using an edge routed architecture:

- Since lean spines used in an edge routed architecture act as pure IP transit devices, they don't necessarily have to be feature rich.
- Edge routed access leaf devices are generally Trident-based platforms (such as Juniper Networks QFX5110 and QFX5120 switches) or Juniper Networks QFX10000 switches.

Table 2: QFX Series Device Roles: Edge Routed Architecture

Device Roles	Edge Routed
Edge Routed_GW @ Leaf (L2/L3 GW @Server Leaf)	QFX5110, QFX10002-36Q/72Q, QFX5120-48Y
DCGW @ Leaf (Border Leaf)	QFX10002-36Q/72Q, QFX10008/16
DCGW@ Spine (Border Spine)	QFX10002-36Q/72Q, QFX10008/16
Null @ Spine (Lean Spine)	QFX5000, QFX10000

Automating Data Center EVPN-VXLAN Fabrics

EVPN-VXLAN can be configured from the device command line interface (CLI), as well as with automation tools such as Ansible and Saltstack. It is also possible to orchestrate underlay and overlay setup through multicloud orchestrators like Juniper Networks Contrail Enterprise Multicloud and its unified management interface, Contrail Command. Figure 3 shows the options available for customers automating their EVPN-VXLAN fabric.

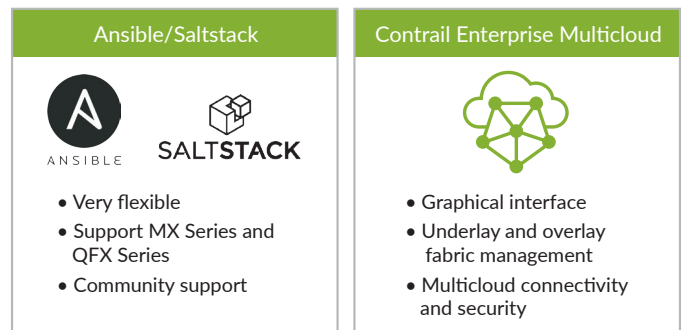


Figure 3: Automating EVPN-VXLAN fabric

The role of Juniper Networks Contrail® Enterprise Multicloud is to unite everything, allowing for automated operations, visibility, and system integration throughout the multicloud environment. For this to occur, all devices must be part of the fabric, enabling data center and private cloud environments to extend to the public cloud and create a unified architecture across disparate domains. This creates a homogeneous management environment where resources can be consumed the same way everywhere, regardless of location (see Figure 4).

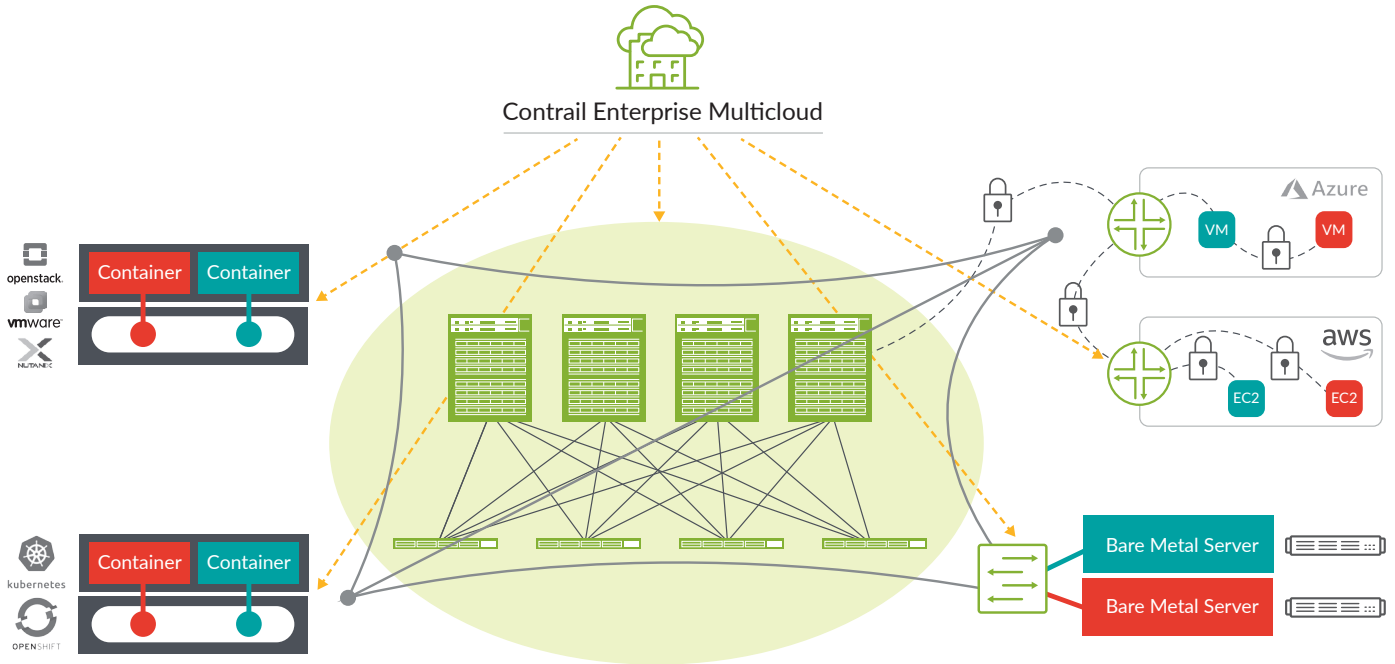


Figure 4: Building blocks from Juniper's data center blueprint

Telemetry Edge

Juniper AppFormix provides end-to-end visibility into multicloud environments (see Figure 5), eliminating any potential network issues and rendering operations simpler and more effective.

AppFormix helps you visualize and analyze both physical and virtual environments, using monitoring and intent-based analytics to transform raw data into information that can be used immediately.




Figure 5: AppFormix provides cross-layer visibility and analytics

Summary: The Enterprise Data Center Must Embrace EVPN-VXLAN

Juniper's enterprise data center solutions, based on a VXLAN overlay with an EVPN control plane, is an efficient and scalable way to build and interconnect multiple next-generation campuses, data centers, and public clouds. With a robust BGP/EVPN implementation on all QFX Series, Juniper is uniquely positioned to help EVPN technology realize its full potential by providing optimized, seamless, and standards-compliant L2 and L3 connectivity, both within and across today's evolving enterprise networks.

For more information about how to design and deploy overlay networks, please read the [Cloud Data Center Blueprint Architecture Components](#) in the Tech Library. To learn more about Juniper's EVPN-VXLAN solutions, read the [IP Fabric EVPN-VXLAN Reference Architecture](#).

About Juniper Networks

Juniper Networks brings simplicity to networking with products, solutions and services that connect the world. Through engineering innovation, we remove the constraints and complexities of networking in the cloud era to solve the toughest challenges our customers and partners face daily. At Juniper Networks, we believe that the network is a resource for sharing knowledge and human advancement that changes the world. We are committed to imagining groundbreaking ways to deliver automated, scalable and secure networks to move at the speed of business.

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

Fax: +1.408.745.2100

www.juniper.net

APAC and EMEA Headquarters

Juniper Networks International B.V.
Boeing Avenue 240
1119 PZ Schiphol-Rijk

Amsterdam, The Netherlands

Phone: +31.0.207.125.700

Fax: +31.0.207.125.701

JUNIPER NETWORKS | Engineering
Simplicity

