Simplify the Data Center with Junos Fusion

Juniper Networks Fabric Technology

Table of Contents

List of Figures

Executive Summary

For too long, the challenges associated with provisioning and managing the network have prevented organizations from realizing the full potential of data center virtualization. From a management and scaling perspective, traditional multitier network architectures are complex, static, and expensive to build and operate. The need to manually provision each network device has also impeded business agility.

Juniper Networks® Junos® Fusion addresses the challenges posed by traditional network architectures and provides customers with a bridge from legacy networks to software-defined cloud networks. This innovative architecture is based on three design principles: simplicity at scale, smart, and flexible.

A highly scalable fabric, Junos Fusion collapses multitier architectures into a single tier, reducing the number of devices in the data center network and cutting CapEx. Junos Fusion is a centrally managed fabric that features plug-and-play provisioning and auto-configuration capabilities, which greatly simplifies operations at scale and reduces OpEx while accelerating the deployment of new applications and services.

Junos Fusion's industry-leading scale, performance, and density are ideally suited to the needs of highly virtualized business-critical IT data centers and software-defined cloud data centers.

Introduction: Network Challenges in the Data Center

The data center is at the heart of business operations for enterprises and service providers alike. Over the past few years, server and storage virtualization have brought efficiencies to the data center in the form of rapid application deployment, as well as application resiliency and scale through virtual machine (VM) mobility.

However, virtualization of the network has lagged. Solutions such as Juniper Networks QFabric® System and Virtual Chassis technology have greatly simplified network provisioning and operations, but aspects of networking, such as VLAN configuration, still must be done manually.

For organizations with traditional multilayer network architectures, the challenges are particularly acute. IT must manually configure each switch and router across the entire network, a time-consuming and error-prone process that drives up operational expenses and impedes application deployment. While it takes the application team only minutes to bring up a new server in support of a business-critical application, it can take the network team hours or even days to make the changes required to support that application.

In this operational model, networks are static—which is completely at odds with the organization's need for responsiveness and agility. To unleash the full potential of virtualization and make data center operations truly agile, businesses need to fundamentally rethink how networks are designed. Traditional device-centric, three-tier network architectures are a poor match for today's virtualized and cloud-based data centers. They present IT with a number of network challenges, including:

- **Operational complexity:** IT organizations operate like a service provider internally. They are under pressure to meet business agility needs, including rapid application deployment, as well as performance requirements. Today, most of the network group's time is tied up with repeated manual bring-up and configuration activities that prevent them from performing more business-critical tasks, such as application performance optimization.

  Within the data center, it is estimated that 90 percent of the problems that arise are due to manual errors. A key culprit for these problems is traditional multitier networks, which require that every device be configured independently. This task is complex, time-consuming, and difficult to complete accurately for the entire network. Even the interaction between the application and network teams is essentially manual, requiring an IT "help desk" ticket.

  Likewise, software upgrades are a huge challenge for network operators because all devices must be upgraded to the same software version at once, which requires a long maintenance window. If there are any problems with the new software, it's a huge task to roll back to the previous version, resulting in application downtime.

  To address this provisioning, configuration, and maintenance complexity, data center operators need a network solution that is automated; supports plug-and-play provisioning and software upgrade versioning; and has a single point of management for the entire data center network.

- **Mix of legacy, cloud, and virtualized applications:** Today's applications are more distributed than ever and highly virtualized, utilizing a variety of hypervisors. They are also developed differently today than they were in the past, with a focus on rapid delivery and resiliency.

  While new applications are written to be cloud-ready, most IT organizations have legacy applications that are tied to the network topology but must be hosted alongside highly virtualized cloud applications.

  The new cloud-ready applications, which are highly virtualized with dozens of VMs per host, place demands on the network for high logical scale, including support for very high numbers of host media access control (MAC) addresses, routes, and VLAN instances. However, network devices historically have supported limited logical scale, and this has constrained the scale of data centers.

In the cloud era, workloads communicate using virtual tunnels with technologies such as Virtual Extensible LAN (VXLAN), MPLS, and generic routing encapsulation (GRE). The data sent across these tunnels requires policy enforcement and intelligent tunnel routing, switching, and termination.

IT needs a flexible network topology that supports both legacy and new applications with universal connectivity and uniform performance across all applications. Such a topology must accommodate the mix of bare-metal servers, appliances, and virtualized applications found in many data centers. It must also accommodate the VM mobility requirement that any VLAN or network segment be available anywhere in the network. Network devices must be smart enough to switch and route tunnels across even the largest, most host dense data centers.

- **Multiple layers of networking:** Historically, as data centers grew from small to large, the network became more complex, with multiple tiers of network devices arrayed in a tree-like structure. Having tiers of hardware to support access, application, core, and edge layers is expensive, requiring considerable capital investment in equipment along with cabling and power costs. Multilayer architectures are static and inefficient in how they forward traffic, which hurts business agility and application performance.

  Data center operators need an elastic, smart network architecture that can collapse legacy multitier designs into a simple one-tier design that can scale up and down as needed, and expand in a pay-as-you-grow manner. As the need for capacity grows, IT should be able to simply drop in new switches to scale the network and boost the number of endpoints it can support. In addition, this scalable network architecture should support high bandwidth, allowing all links to operate in active/active mode and deliver line-rate connectivity.

  Collapsing multiple network layers into a single layer by using a flexible set of building blocks reduces data center CapEx and OpEx, bringing consistency to operations.

To meet software-defined cloud data center business objectives, data center operators need a network that's flexible, scalable, simple to operate, and supports the full range of applications in use. The Junos Fusion architecture was developed specifically to address these network challenges.

Introducing Juniper Networks Junos Fusion

Junos Fusion is an innovative fabric architecture designed from the ground up to provide a bridge from legacy networks to software-defined cloud networks. Based on standard protocols and open application programming interfaces (APIs), Junos Fusion provides industry-leading scale, performance, and simplicity, virtualizing the network by treating it as a logical entity. By providing a single point of management for all devices in a data center switching network, this flexible architecture greatly simplifies network operations at scale.

Junos Fusion is built with the new Juniper Networks QFX10000 line of modular Ethernet switches and either EX4300 Ethernet switches and/or QFX5100 switches, providing a highly scalable, high-density network foundation for today's most demanding data center and cloud environments. Customers can start small—with as little as one rack—and expand the fabric to include up to 128 racks. Whether a data center has 300 or 6,000 servers, Junos Fusion can significantly cut operations overhead and enable faster application delivery.

The Junos Fusion architecture consists of two major components: "aggregation" devices and "satellite" devices. These components work together as a single switching system, flattening the network to a single tier without compromising resiliency. Data center operators can build individual Junos Fusion "pods" comprised of a pair of aggregation devices and a set of satellite devices; each pod is a discrete unit, managed as a single device. Pods can be small—for example, a pair of aggregation devices and a handful of satellites—or large with up to 128 satellite devices based on the needs of the data center operator.

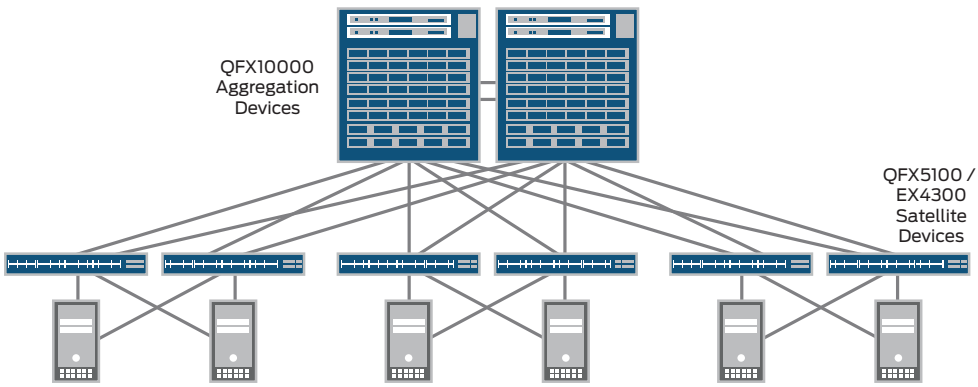

Figure 1: Junos Fusion deployments are comprised of aggregation and satellite devices.

The flexible Junos Fusion architecture is ideal for a variety of use cases, including:

- Highly virtualized business-critical IT data centers, allowing operators to quickly create high-performance, predictable pods
- Software-defined cloud data centers, supporting self-provisioning of services

Junos Fusion Key Features

The Junos Fusion architecture is based on three design principles that pave the way to next-generation software-defined cloud networking. These principles are simplicity at scale, smart, and flexible.

Simplicity at Scale

- Junos Fusion offers a single point of management for up to 64 racks with the capability to expand up to 128 racks.
- Junos Fusion enables plug-and-play provisioning for satellite devices without any user intervention. Devices are auto-discovered and download the correct software image and configuration; likewise, VLANs are auto-sensed and automatically provisioned on the appropriate switch ports.
- Software Upgrade Groups allow customers to perform staggered upgrades, providing tremendous flexibility and ensuring the highest application availability by making it easy for operators to qualify new software prior to a full-blown rollout.

Smart

- Junos Fusion provides a bridge to software-defined cloud networking with no protocol or vendor lock-in. It also supports native L2, L3, MPLS, and overlay tunnel routing capabilities with Open vSwitch Database (OVSDB) and Ethernet VPN (EVPN).
- Junos Fusion is programmable with REST APIs for management, orchestration, and event monitoring.
- Junos Fusion is driven by open, standard protocols, delivering a loosely coupled system with distributed forwarding at scale.
- The architecture provides industry-leading logical scale, performance, and density.
- Multiple control planes ensure high application resiliency.

Flexible

- The EX4300 and QFX5100 satellite devices used by the Junos Fusion architecture are fully functional switches optimized for data centers. They can be reused or redeployed in another architecture to protect investments.
- Junos Fusion supports seamless migration from 1GbE to 100GbE, with access choices ranging from 100BASE-T to 40GbE, and uplinks between satellite and aggregation devices ranging from 10GbE to 100GbE.

Inside Junos Fusion

The following sections detail the key components and technologies underpinning Junos Fusion, as illustrated in Figure 2.

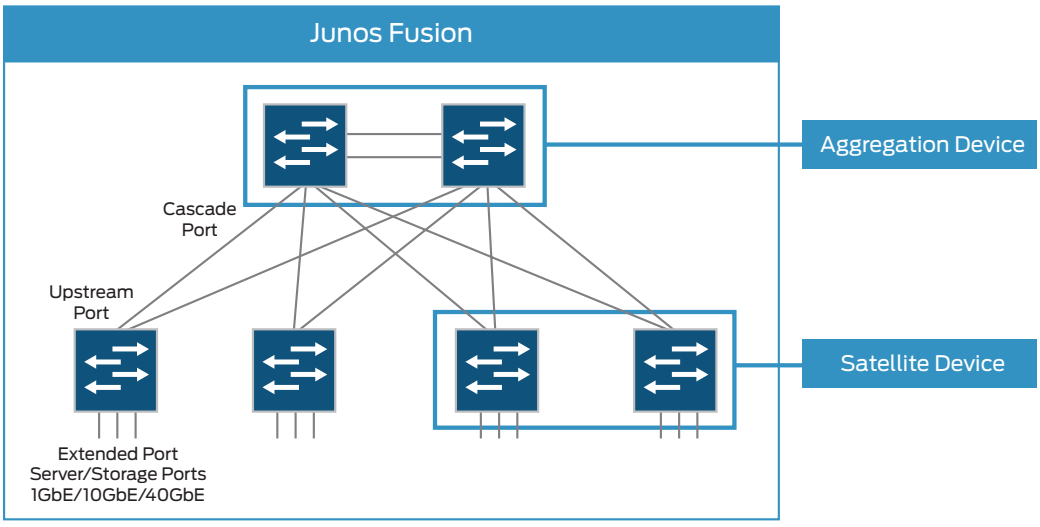

Figure 2. Junos Fusion architecture

Aggregation Devices

In the Junos Fusion architecture, QFX10000 switches function as aggregation devices, supporting speeds ranging from 10GbE to 100GbE. As the foundation of Junos Fusion deployments, the QFX10000 switches provide:

- Industry-leading line-rate 40GbE and 100GbE density, supporting up to 480 100GbE ports in a single chassis
- The highest logical Layer 2/Layer 3 scale, supporting up to 1 million MACs, 2 million host routes, and a 256,000-entry forwarding information base (FIB), also known as forwarding table
- Deep buffers, with 50 ms delay bandwidth
- No head-of-line blocking with a virtual output queue (VoQ)-based architecture
- High availability with topology-independent in-service software upgrade (TISSU)

The aggregation device runs Juniper Networks Junos operating system, extending rich Junos OS software functionality to the interfaces on the satellite devices. Aggregation devices manage all satellite devices centrally, eliminating the need to manage satellite switches individually, which simplifies operations and significantly reduces the overhead associated with configuring, monitoring, and upgrading these devices.

Satellite devices are connected to an aggregation device via a cascade port over which control and network traffic is sent and received. Network operators can set up one or more cascade port connections with one or more satellite devices. When multiple cascade ports are connected to a satellite device, traffic is automatically load-balanced across those ports. In addition, a cascade port can be included in a link aggregation group (LAG) or multichassis link aggregation group (MC-LAG).

For maximum resiliency, many data center operators will deploy two aggregation devices. These are connected via MC-LAG, which is automatically provisioned and configured. Operating as peers, the pair of aggregation switches provide dual control planes with active/active operation, ensuring maximum up time. The Junos Fusion architecture also allows IT to upgrade aggregation devices separately, one at a time. Customers benefit from flexibility and peace of mind, with the ability to roll back to the old software version if needed, without losing communication between the two aggregation switches.

Satellite Devices

Junos Fusion satellite devices—specifically the QFX5100 and EX4300 switches—are open, standards-based access platforms that support the IEEE 802.1BR protocol. Customers who have already deployed these switches can easily migrate to the Junos Fusion architecture with a simple Junos OS upgrade, protecting their investment.

These Juniper switch lines offer customers a flexible range of access speeds, including 1/10/40/100GbE today, with 25/50GbE options in the near future. Uplink speeds range from 10GbE to 100GbE.

Junos Fusion satellite switches run a Linux-based operating system. Satellite switches appear as line cards to the aggregation switch, with all of their ports visible, inheriting the full features and functionality of the aggregation device.

Each satellite device has at least one direct uplink connection to its associated aggregation device(s), and can be single- or dual-homed to aggregation devices. When a satellite switch has multiple uplink ports to the same aggregation device, traffic is automatically load-balanced among the uplink ports. In addition, the Junos Fusion fabric supports multihoming of hosts, such as servers and appliances, to a pair of satellite devices.

Juniper provides an API between satellite and aggregation devices that can be used to build applications that perform various functions on the satellite switches. For example, using this API, users can enable local switching for traffic whose source and destination are on the same satellite device. Any traffic whose destination is not local to the satellite switch is sent to the aggregation device for forwarding in a multipath configuration. Local switching is ideal for enabling two servers in the same rack to communicate directly, ensuring high-performance, low-latency forwarding.

In addition, the API supports application-specific uplink selection, or “port pinning,” which allows IT to configure the specific uplink to be used by an application or application flow. This capability is ideal for isolating short-lived “mice” flows, such as transactional data traffic, from long-lived “elephant” flows, such as storage traffic. This open API can be used by Juniper as well as third parties to develop applications that perform functions such as port mirroring, interface statistics collection, software image management, environmental monitoring, and centralized event monitoring.

Connecting Multiple Junos Fusion Systems

Data center operators have the flexibility to deploy small or large Junos Fusion pods and interconnect them using EVPN with VXLAN encapsulation. These standards-based network virtualization technologies are commonly used for transporting data within and across data centers. With EVPN and VXLAN support, customers can seamlessly connect Junos Fusion pods within and across data centers with optimal traffic forwarding. Customers also benefit from seamless connectivity in a multivendor environment.

Maximize Business Agility with Junos Fusion

An innovative fabric architecture, Junos Fusion paves the way to next-generation software-defined cloud networking. By flattening the network to a single tier and treating it as one logical block, Junos Fusion brings unparalleled simplicity, scalability, and flexibility to the data center. Customers benefit from a new level of agility within the data center, enabling them to rapidly deploy applications and services.

Junos Fusion also significantly lowers CapEx and OpEx, provides investment protection, and future-proofs the network with a pay-as-you-grow model. Junos Fusion greatly simplifies network operations at scale by giving operators a single point of management along with plug-and-play provisioning and autodiscovery. In addition, Junos Fusion's use of standard protocols and open APIs eliminates vendor lock-in.

With Junos Fusion, customers can finally unleash the full potential of data center virtualization.

About Juniper Networks

Juniper Networks is in the business of network innovation. From devices to data centers, from consumers to cloud providers, Juniper Networks delivers the software, silicon and systems that transform the experience and economics of networking. The company serves customers and partners worldwide. Additional information can be found at www.juniper.net.

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

2000597-001-EN Feb 2015