

Network Simplification with Juniper Networks Virtual Chassis Technology

Table of Contents

Executive Summary	3
Introduction.....	3
Data Center Network Challenges.....	3
Design and Operations.....	3
Performance	4
Resource Consumption.....	4
Enterprise Network Challenges.....	4
Bandwidth Utilization	5
Manageability.....	5
Juniper Networks Virtual Chassis Technology	5
Simplifying the Data Center	5
Reduced Complexity	6
Simplifying the Enterprise Network	6
Cost-Efficient Use of Resources.....	8
Improved Performance	8
Product Portfolio.....	8
Conclusion.....	9
About Juniper Networks.....	10

List of Figures

Figure 1: Legacy data center architecture.....	3
Figure 2: Traditional enterprise network.....	4
Figure 3: Data center EX Series and QFX Series Virtual Chassis technology	6
Figure 4: Small to medium-sized enterprises with Virtual Chassis technology.....	7
Figure 5: Large enterprise network with Virtual Chassis technology.....	7

Executive Summary

The proliferation of mobile devices—combined with trends such as cloud, analytics, and social media—is fundamentally changing user behavior and network usage patterns, driving demand for bandwidth and a growing need for resiliency and security. These changes are making enterprise and data center networks increasingly complex, adding to IT's economic burden by creating environments that are inherently difficult to manage and operate.

Network complexity is by far the single biggest roadblock to scalability for data center and enterprise networks. Juniper Networks® Virtual Chassis technology—available on Juniper Networks EX Series Ethernet Switches and QFX Series switches—offers an innovative, unique solution for deploying and growing data center and enterprise networks, helping organizations address this growing complexity by making network architectures simpler, more reliable, and more cost-effective to manage and maintain. This paper describes how Virtual Chassis technology on EX Series and QFX Series switches reduces network complexity by simplifying network operations and architecture, increasing scalability and performance while reducing operational costs.

Introduction

Data Center Network Challenges

Most data center networks today employ legacy three-tier architectures. These networks are plagued by complexity and inflexibility in the various network layers, leading to poor performance and excessive resource consumption (Figure 1).




Figure 1: Legacy data center architecture

Design and Operations

Data center networks typically use a variety of design architectures to provide access layer connectivity. Two of the methods are top-of-rack and end-of-row deployments. While these deployment methods have their advantages, they also suffer from excessive resource consumption, inefficient bandwidth utilization, and costly cable management. In addition, server virtualization, which has been adopted in today's data center, demands networks that are nimble and adaptable enough to react quickly to changes and maintain efficient delivery of mission-critical services. Virtual servers and the workloads running on them need to be migrated either within or across data centers based on demand and load. The network has to be able to support this migration and adapt to the changes quickly with no risk of service disruption.

Such changes require that devices appear to be seamlessly connected on the same network, regardless of their physical proximity to each other. Legacy network architectures lack the flexibility to react to these types of changes, and this adversely affects the operational efficiency of the entire data center.

In addition to the way boxes are physically connected and oriented, maintaining the network becomes a challenge when existing technologies prevent network engineers from doing their jobs efficiently. Routine tasks such as monitoring devices, troubleshooting, providing configuration management, and maintaining software upgrades become increasingly difficult as the number of independent devices in the network increases.

Such operational challenges are further compounded if these devices are running different versions of software or have different configurations, since software must be carefully managed across devices to ensure consistent functionality and limit exposure to defects or other vulnerabilities. Special training or expertise might also be needed to support these configurations. As a result, the effort and resources required to adequately operate, maintain, and troubleshoot the unique requirements of each network device can be enormously time-consuming and expensive.

Performance

The complexities of today's data center architectures lead to increased latency, delays in network convergence, and limited bandwidth availability.

- **Latency caused by the network architecture:** Approximately 75% of all traffic in today's data center is server-to-server, which means it travels laterally, or east to west, across the infrastructure. However, due to the multilayered architecture employed by most data center networks, this traffic must first travel north and south from the access layer up to the aggregation and core layers and then back down again before it reaches its final destination—a costly, inefficient use of network assets that adds latency and complexity to each transaction.
- **Suboptimal use of access and uplink ports:** In today's data center, approximately 50% of access layer switch ports are used for inter-switch connections to higher-layer devices in the hierarchical tree, limiting the bandwidth available for supporting customer connections.
- **Layer 2 control plane scaling:** Spanning Tree Protocol (STP) is typically employed to prevent network loops from occurring in the data center. However, STP can take up to 50 seconds to converge in a network following a failure—even the Rapid Spanning Tree Protocol (RSTP) can require tens of seconds to converge in some topologies. Plus, both STP and RSTP render half the ports in the core and aggregation layers unusable, leading to inefficient bandwidth utilization.

Virtualized servers compound these problems, since they, too, require high performance and low latency.

Resource Consumption

Typically, the quest for higher bandwidth in modern data centers involves operators adding more network devices that end up consuming additional rack space, power, and cooling. However, this extra resource consumption does not necessarily translate to the required scale in bandwidth. This is an inefficient approach since the data center can rapidly run out of power due to excessive resource consumption and the bandwidth challenge now becomes multifold while the operating costs rapidly escalate.

Enterprise Network Challenges

Many enterprise access networks are also built with a three-tier architecture (Figure 2) and are constrained by the same complexities that plague legacy data centers. As the enterprise network grows, the sheer number of devices that need to be managed grows exponentially, increasing the burden on network administrators.

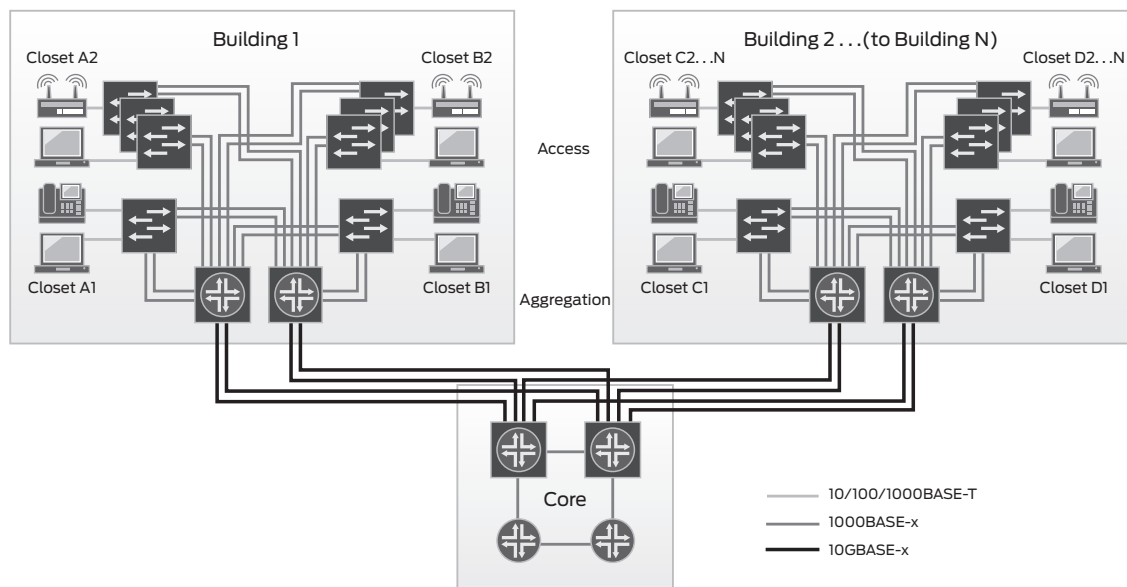


Figure 2: Traditional enterprise network

Bandwidth Utilization

Today's applications and business processes are bandwidth hungry and eventually lead to network performance degradation. Unfortunately, the traditional enterprise network architecture does not utilize bandwidth efficiently. For instance, while access layer switches connect to the aggregation layer using redundant links, Spanning Tree Protocol effectively blocks one of those links, reducing available bandwidth by 50%. Compounding the problem, as more ports are used to connect to aggregation layer switches, the number of ports available to connect desktops, IP phones, and access points is reduced as well.

Manageability

As enterprise networks grow, managing them becomes increasingly difficult. In traditional enterprise networks, even a small campus of 500 users needs 10 access switches. Everyday tasks such as configuration management, image upgrades, and monitoring increase operational expenses. It also becomes difficult for networks to scale as business locations grow in size.

Juniper Networks Virtual Chassis Technology

Juniper Networks Virtual Chassis technology, available on most Juniper Networks EX Series Ethernet Switches and QFX Series Ethernet switches, addresses many data center and campus challenges by allowing multiple interconnected switches to behave, operate, and be managed as a single logical, high-bandwidth device. Virtual Chassis technology simplifies the network by reducing the number of managed devices, helping networks scale without the operational overhead associated with maintaining a system of independent switches.

Virtual Chassis technology allows multiple EX Series and QFX Series switches to be interconnected over a high-speed backplane using dedicated Virtual Chassis ports, or through optional 10GbE or 40GbE fiber ports configured as Virtual Chassis ports (see Table 1). Overall system maintenance and management is greatly simplified, since up to 10 interconnected switches can be managed as a single entity, through a single management interface. In other words, Virtual Chassis technology can reduce the number of managed devices by a factor of up to 10, dramatically lowering operational expenses. Juniper Networks Junos® Space Network Director, a consolidated network management platform, can be used to configure and manage the Virtual Chassis configuration as a single, logical entity.

Virtual Chassis technology also delivers the following benefits:

- **Greater availability:** Enabling Virtual Chassis technology increases network availability to 99.999%. Connectivity remains uninterrupted since the Virtual Chassis master selection process happens automatically, with no manual intervention required. Additionally, the combination of Virtual Chassis technology with NSSU (nonstop software upgrade) and VRRP (Virtual Router Redundancy Protocol) virtually eliminates network downtime and enables nonstop business operations.
- **Better performance, scale, and flexibility without trade-offs:** The ability to span access and aggregation network tiers and interconnect up to 10 switches in a Virtual Chassis configuration enables flexible scaling as business requirements change. Pay-as-you-grow scalability on fixed configuration switches—for instance, from 24 to 480 10/100/1000BASE-T ports with Juniper Networks EX4200 Ethernet Switches, and from 32 to 480 10GbE small form-factor pluggable transceiver (SFP) ports on the Juniper Networks EX4500 Ethernet Switches—allows flexible growth as requirements change. Location flexibility can be derived by extending Virtual Chassis configurations across vertical or horizontal distances up to 80 km with redundant fiber links.
- **Consistent control plane feature implementation:** All EX Series and QFX Series switches run the same modular Juniper Networks Junos operating system, making network learning, administration, and maintenance much simpler.
- **Lower latency:** Since a Virtual Chassis configuration operates as a single logical device, latency is greatly reduced.
- **Reduced OpEx:** Lower rack space, power, and cooling resource consumption reduce operating costs, while the elimination of Spanning Tree Protocol extracts better link utilization and application performance from the network.

Simplifying the Data Center

In the data center, Virtual Chassis technology simplifies the network by collapsing tiers and flattening the network from three to two layers (see Figure 3). This is accomplished by interconnecting Virtual Chassis switch members via high-speed backplane connections (Virtual Chassis ports), conserving valuable access ports, and effectively merging what would normally be many LANs into one. As a result, the layers of switching required for network access are reduced. This flexibility extends a single Layer 2 access network beyond a single rack, reducing the effort required for network changes such as live server migrations. Moreover, a Virtual Chassis configuration can extend Layer 2 access between sites up to 80 km apart. This simplistic and innovative approach to networking, along with high-performance packet forwarding capabilities, greatly minimizes the effort required to deploy new services in today's virtualized data centers.




Figure 3: Data center EX Series and QFX Series Virtual Chassis technology

Reduced Complexity

Juniper's Virtual Chassis technology reduces complexity and delivers the benefits of both top-of-rack and end-of-row access switch deployments without businesses having to choose between the two. Virtual Chassis technology does the following:

- Simplifies cable management, since customer connections are intra-rack (copper or fiber)
- Makes configuration changes easier
- Reduces maintenance efforts, because there is only one IP address, one image, and one configuration file to manage
- Provides flexibility with uplinks in the core

Simplifying the Enterprise Network

Virtual Chassis technology also enables businesses to simplify their enterprise networks. Small to medium-sized campuses (up to 5,000 access ports) can collapse their aggregation and core layers using Virtual Chassis technology (see Figure 4). Fewer 10GbE uplinks are required when using this solution. Up to 10 Juniper Networks EX3300, EX4200, EX4300, EX4500, and EX4550 Ethernet Switches, in any combination, can be interconnected using Virtual Chassis technology—further simplifying the network by reducing the number of managed devices. Alternatively, up to four EX2200 switches with Virtual Chassis technology can be interconnected in low-density wiring closets.

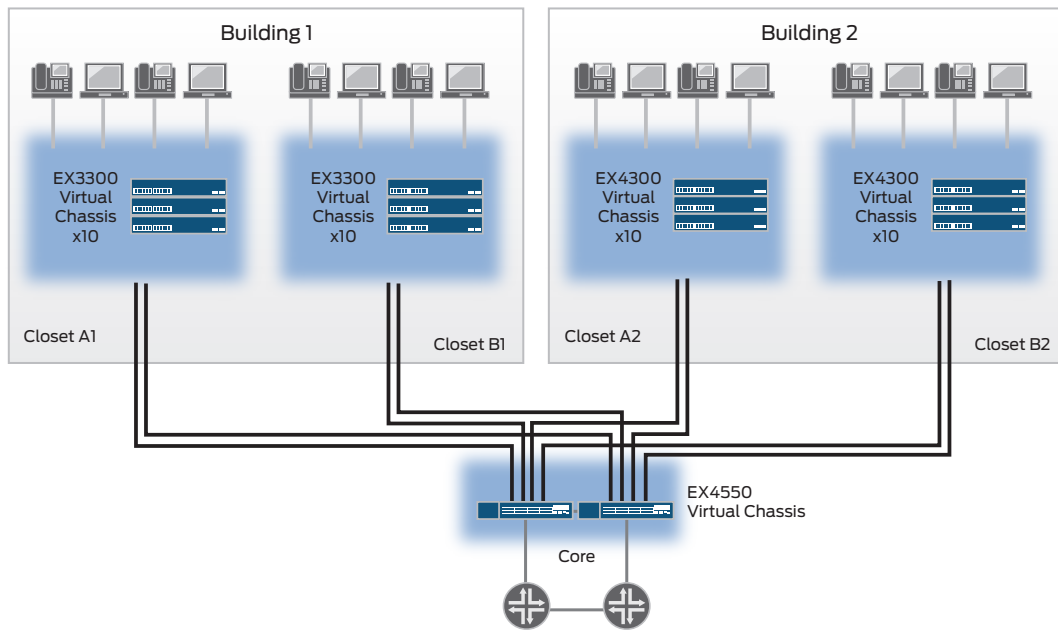


Figure 4: Small to medium-sized enterprises with Virtual Chassis technology

In large enterprise networks, the EX9200 or EX8200 lines of switches with Virtual Chassis technology can be deployed in the core and aggregation layers, respectively, as shown in Figure 5. Using Virtual Chassis technology here reduces the number of devices that need to be managed. Collapsing several individual devices into a single logical device eliminates the need for STP, eliminating delays resulting from network convergence while enhancing performance and improving resource utilization.

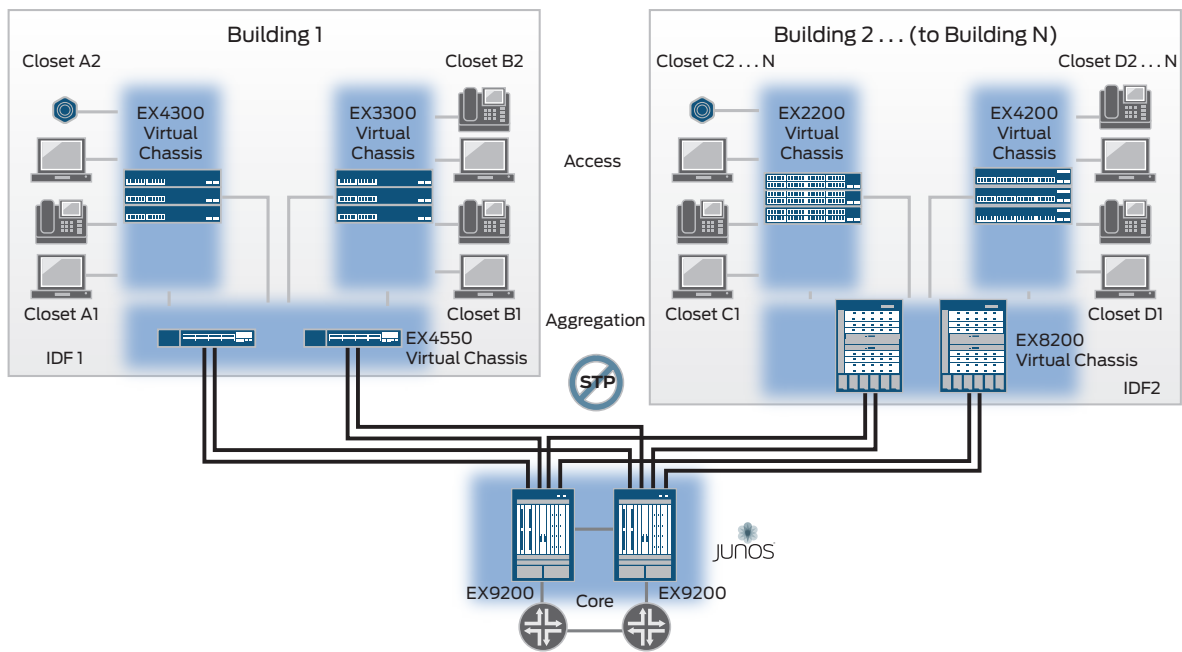


Figure 5: Large enterprise network with Virtual Chassis technology

Cost-Efficient Use of Resources

Virtual Chassis technology makes more efficient use of both access and uplink ports, optimizing rack space, power, and cooling resources. Also, by eliminating unnecessary switching layers, a Virtual Chassis configuration dramatically reduces the equipment and the resources required to plan, deploy, implement, and operate today's corporate network.

Improved Performance

Virtual Chassis technology reduces latency by flattening the network. Inter-switch traffic is routed over a dedicated Virtual Chassis backplane at line rates for all packet sizes, rather than flooding traffic over access ports, preserving valuable bandwidth. These backplane inter-switch connections reduce the number of devices, conserving valuable data center resources. Also, with Virtual Chassis technology, node and link failover times are measured in subseconds, without the need for an external Layer 2 control plane protocol like STP, creating a loop-free topology.

Product Portfolio

Virtual Chassis technology is available on EX Series and QFX Series switches designed for every location in the network, from access to core (see Table 1).

Table 1: Virtual Chassis Availability

EX Series Switch	Details	Maximum Number of Virtual Chassis Members	Maximum Virtual Chassis Backplane Capacity	Virtual Chassis Attributes
EX2200-C	12-port 10/100/1000BASE-T switch	4	NA	Front-panel fiber or copper GbE uplink ports
EX2200	24-port or 48-port 10/100/1000BASE-T switch with 4 SFP uplinks	4	NA	Front-panel fiber or copper GbE uplink ports
EX3300	24-port or 48-port 10/100/1000BASE-T switch	10	40-80 Gbps	Two of the four uplinks are automatically configured for Virtual Chassis configurations. They are implemented via the LCD without entering a single CLI command.
EX4200	24-port or 48-port 10/100/1000BASE-T switch with 2-port or 4-port uplinks	10	128 Gbps	A dedicated Virtual Chassis port is on each device. Optional 1GbE or 10GbE uplink ports and connections extend across 50 Km. A single logical switch consisting of devices resides in different wiring closets, buildings, or cities.
EX4300	24-port or 48-port 10/100/1000BASE-T switch with 4-port GbE/10GbE fiber uplink module 32-port 1000BASEX switch with 4x10GbE, 2x40GbE	10	320 Gbps	Standard 10GbE/40GbE ports can be used for Virtual Chassis configuration. This allows an entire building or multiple buildings to be managed as one device in small enterprise access networks.
EX4500	40 10GbE fiber ports and two expansion ports (48 ports total)	10	128 Gbps	Any combination of up to 10 EX4550, EX4500, and EX4200 switches can be interconnected in the same Virtual Chassis configuration
EX4550	32-48 10GbE SFP/SFP+ or 10GbE RJ-45, with 2 expansion slots (8-port uplink or Virtual Chassis module)	10	256 Gbps	Any combination of up to 10 EX4550, EX4500, and EX4200 switches can be interconnected in the same Virtual Chassis configuration
EX8200	8-slot or 16-slot modular chassis; range of line-card options	2	NA	EX8208 (8-slot) and EX8216 (16-slot) chassis can be interconnected using standard line-rate 10GbE interfaces. Connection between any two chassis in a Virtual Chassis configuration can either be a single line-rate 10GbE link or a link aggregation group (LAG) with up to 12 10GbE line-rate links. Deployment is done using Juniper Networks XRE200 External Routing Engine, which externalizes the control plane functionality and provides true control and data plane separation.

EX Series Switch	Details	Maximum Number of Virtual Chassis Members	Maximum Virtual Chassis Backplane Capacity	Virtual Chassis Attributes
EX9200	4-slot, 8-slot or 14-slot modular chassis; range of line-card options	2	NA	No XRE is required. Supports Virtual Chassis and MC-LAG to provide deployment flexibility required in core and aggregation layers.
QFX3500	63 x 10GbE ports—36 dual-mode 1GbE/10GbE ports; 12 dual-mode 10GbE or 2 Gbps, 4 Gbps or 8 Gbps Fibre Channel ports, 15 using QFSP+ to-SFP+ direct attach copper (DAC) or QSFP+ to SFP+ fiber splitter cables and optics	10	N/A	A mix of EX4300, QFX3500, QFX3600, or QFX5100 can be deployed in a single Virtual Chassis configuration. Interconnect with standard front-panel 10GbE or 40GbE ports.
QFX3600	64 ports (using QSFP+ to SFP+ direct attach copper [DAC] or QSFP+ to SFP+ fiber splitter cables and optics); 16 40GbE QSFP+ ports	10	N/A	Interconnect with standard front-panel 10GbE or 40GbE ports
QFX5100-24Q	24 40GbE QSFP+ ports; two expansion slots for hot-swappable 4 x 40GbE expansion modules 104 ports; 26 QSFP+ to SFP+ direct attach copper (DAC) or QSFP+ to SFP+ fiber splitter cables and optics on base system and hot-swappable 4 x 40GbE expansion module	10	N/A	A mix of EX4300, QFX3500, QFX3600, or QFX5100 can be deployed in a single Virtual Chassis configuration. Interconnect with standard front-panel 10GbE or 40GbE ports.
QFX5100-48S	48 1GbE/10GbE SFP/SFP+ ports; 6 40GbE QSFP+ ports 72 ports (48 10GbE SFP/SFP+ ports + 24 10GbE ports using QSFP+ to SFP+ direct attach copper [DAC] or QSFP+ to SFP+ fiber splitter cables and optics)	10	N/A	A mix of EX4300, QFX3500, QFX3600, or QFX5100 can be deployed in a single Virtual Chassis configuration. Interconnect with standard front-panel 10GbE or 40GbE ports.
QFX5100-96S	96 1GbE/10GbE SFP/SFP+ ports; 8 40GbE QSFP+ ports 104 ports (96 10GbE SFP/SFP+ ports + 8 10GbE ports using 2 QSFP+ to SFP+ direct attach copper [DAC] or QSFP+ to SFP+ fiber splitter cables and optics)	10	N/A	A mix of EX4300, QFX3500, QFX3600, or QFX5100 can be deployed in a single Virtual Chassis configuration. Interconnect with standard front-panel 10GbE or 40GbE ports.

Conclusion

With the proliferation of mobile devices and the corresponding rise in rich media consumption, demands for resiliency and security from today's networks are on the rise. To adapt to these challenges, modern-day enterprise and data center networks are becoming more complex and inherently difficult to manage and operate. The EX Series and QFX Series with Virtual Chassis technology deliver a highly scalable solution that reduces network complexity, increases flexibility, improves performance, and reduces resource consumption. This technology—together with market-leading port densities and the consistent, reliable, stable Juniper Networks Junos operating system—increases operational efficiencies by improving resource and network asset utilization. The result is lower operational, maintenance, and troubleshooting costs, which translates into a more cost-effective solution for both data center and enterprise networks.

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

Copyright 2015 Juniper Networks, Inc. All rights reserved. Juniper Networks, the Juniper Networks logo, Junos and QFabric are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners. Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.