

Juniper Networks Virtual Chassis Fabric Technology

A Fast, Flat Network for Small and Medium-Sized Data Centers

Table of Contents

Executive Summary	3
Introduction.....	3
The Evolution of Juniper Networks' Fabric Technologies	4
Virtual Chassis Fabric for Medium and Large Data Centers	5
Key Features and Benefits.....	6
A Deep Dive Look at the Virtual Chassis Fabric.....	8
Virtual Chassis Fabric Use Cases	9
Conclusion—Purpose-Built for Modern Applications in Today's Data Center	9
About Juniper Networks.....	10

List of Figures

Figure 1: Virtual Chassis configuration	4
Figure 2: Virtual Chassis Fabric configuration.....	5
Figure 3: Autodiscovery with Virtual Chassis Fabric technology	6
Figure 4: Virtual Chassis Fabric integrated control plane	7

Executive Summary

Today's enterprise data centers need a network architecture designed to support virtual machines (VMs) running highly distributed, modular applications. Traditional three-tier network designs introduce too much latency for modern applications and require cumbersome device-by-device configuration and management. To address today's data center requirements, Juniper Networks has pioneered fabric technologies that deliver the high performance that modern applications need and the business agility and reduced costs that come with management simplicity.

Juniper's fabric technologies have evolved from Virtual Chassis, for interconnecting up to 10 switches to create a single logical device, to its latest offering, Virtual Chassis Fabric, designed for small and medium-sized data centers, is optimized for mixed 1GbE/10GbE/40GbE environments and can connect up to 20 switches. Like its predecessors, Virtual Chassis Fabric provides a high-performance, flat network topology and the management simplicity of a single logical device.

For small data centers, Virtual Chassis Fabric technology can support the entire switching infrastructure. In larger data centers, IT can link multiple Virtual Chassis Fabric pods to support shared resource pools. Virtual Chassis Fabric also delivers the high-speed, low-latency connectivity that modern applications need at a scale and price designed for mid-sized application deployments. With its single point of management, Virtual Chassis Fabric reduces management costs and simplifies change, enabling IT to readily adapt to new application and business needs.

Through the use of common building blocks such as the Juniper Networks QFX5100 line of switches, Juniper fabric solutions offer a cost-effective approach to data center networking with complete investment protection. For example, enterprises can start small and build up to a full-scale Virtual Chassis Fabric configuration. With this flexible architecture, you can even build more than one fabric and connect them.

Introduction

Thanks to virtualization and distributed application architectures, organizations of all sizes can bring up new applications and services quickly and easily. Unfortunately, many data center networks don't let you fully capitalize on the business agility that virtualization and modern application architectures provide. Traditional network architectures are too slow and too cumbersome to configure. For true agility, enterprises need a flat, high-performance, low-latency network that can be managed like a single, logical switch.

Today's data centers are built with high-performance blade and rack servers, which typically run multiple VMs, which in turn run increasingly modular, web- and cloud-based applications. These modern applications are driving increased traffic levels and different traffic patterns, which place specific demands on the data center network.

Critical business applications such as enterprise resource planning (ERP) and customer relationship management (CRM) are divided into multiple modules. As a result, relatively simple operations such as entering an order can trigger 100 or more individual transactions, each imposing its own latency. Some applications, including e-commerce applications, even behave dynamically, spinning up new instances or migrating workloads in response to traffic loads.

Their distributed nature means that modern applications are spread across racks of servers, each served by multiple switches. The applications generate a tremendous amount of server-to-server, or east-west traffic as the various modules communicate with one another. Multitier network architectures aren't well matched to modern applications. They force this east-west traffic to first travel north and south, up and down the network tree, before arriving at its ultimate destination, adding significant latency that can cause application performance to degrade under load.

Latency adds up quickly in a multitier architecture. A few milliseconds at each hop become seconds of latency per transaction, which seriously impacts application performance—and user experience. Research has shown that even two seconds of delay prompts a majority of customers on e-commerce websites to abandon their shopping carts¹.

Multitier network architectures are also too complex and expensive to operate. Each time a new application is deployed, network administrators must configure the network device by device, driving up operations overhead and slowing down application rollouts. With multitier network architectures, administrators have too many switches to manage, too many management tools to learn, and too many manual processes to execute.

To support today's application environment, enterprises need a high-performance, low-latency data center network that costs less to operate, automates management functions, and simplifies change. Over the past few years, fabric technologies have become popular solutions for implementing efficient data center networks, reducing traditional three-tier networks to flat systems featuring predictable performance and deterministic low latency along with plug-and-play deployment and the ease of managing a single logical switch.

To date, fabric technologies have been primarily targeted at large (i.e., global enterprise-sized) data centers. At the low end, clustering technologies have been available to deliver operational and management simplicity. Thanks to technology advances and engineering innovations, Juniper Networks has developed a new fabric-based data center networking solution purpose-built for small and medium-sized data centers running today's modern applications.

Juniper's Virtual Chassis Fabric technology delivers a high-speed, flat data center networking solution with the deterministic low latency and predictable performance that applications need, and the low operational overhead of a single logical device. In small data centers, a single Virtual Chassis Fabric configuration can provide the entire switching infrastructure. In larger data centers, IT can link multiple Virtual Chassis Fabric pods to support shared resource pools; Virtual Chassis Fabric can also support special projects that require dedicated pools. With Virtual Chassis Fabric technology, customers benefit from improved application performance and the agility that comes with management simplicity.

The Evolution of Juniper Networks' Fabric Technologies

Juniper Networks has pioneered the development of network fabric technologies for both data center and campus networks. In 2008, Juniper introduced its Virtual Chassis technology on its EX Series Ethernet Switches for top-of-rack deployments. Virtual Chassis technology allows up to 10 switches, interconnected via a high-speed backplane, to be managed and operated as a single, logical device with a single IP address. Virtual Chassis technology supports up to 480 ports and can extend Layer 2 access between sites up to 40 kilometers apart.




Figure 1: Virtual Chassis configuration

This architecture conserves valuable access ports and flattens the network from three to two tiers, reducing latency and improving network convergence. Virtual Chassis technology also greatly simplifies management by effectively reducing the number of managed devices by up to a factor of 10, minimizing the effort required to deploy new services and lowering operational expenses. Virtual Chassis technology further reduces operational expenses by automating many network-related tasks.

Simple, flexible, and easy to manage, Virtual Chassis technology is available on most Juniper Networks switches, including the new EX4300 and QFX5100 lines.

Virtual Chassis Fabric for Small and Medium-Sized Data Centers

Both Virtual Chassis and Virtual Chassis Fabric technologies deliver a high-performance, flat network topology with the management simplicity of a single logical device. Thanks to more powerful yet less expensive chipsets and continued software innovation, Juniper has developed a new kind of fabric technology that brings these same benefits to both small enterprise data centers as well as medium-sized data centers that want a dedicated resource pool for mid-sized application deployments.

Juniper's Virtual Chassis Fabric technology enables up to 20 interconnected switches to operate as a two-tier, low-latency, high-performance data center fabric. Deployed in a spine-and-leaf configuration, a Virtual Chassis Fabric deployment features two to four 10/40GbE QFX5100 switches in the spine and up to 16 leaf nodes, which can include any mix of 1/10/40GbE ports on EX4300, QFX3500, QFX3600, and QFX5100 switches. For large multi-dimensional scale configurations, Virtual Chassis Fabric deployments with four spine devices and up to 16 leaf devices has also been qualified. The versatile QFX5100 is the universal building block for all Juniper data center fabric architectures, giving enterprises the flexibility to choose the architecture that best meets their immediate needs and migrate easily from Virtual Chassis to Virtual Chassis Fabric deployments as the business grows and/or they need to deploy new applications. This flexibility is a key aspect of Juniper's fabric evolution.




Figure 2: Virtual Chassis Fabric configuration

Today's data centers are primarily using 1GbE servers but are moving up to 10GbE links as they deploy new applications. Consequently, Virtual Chassis Fabric is optimized for mixed 1GbE/10GbE (server side) and 40GbE (uplinks) data center environments, accommodating both speeds in the same architecture. Virtual Chassis Fabric supports up to 2,688 10GbE ports and provides any-rack-to-any-rack deterministic throughput with less than 2 microseconds of latency.

Like all Juniper fabric offerings, Virtual Chassis Fabric significantly simplifies network operations through a single point of management, effectively reducing the number of managed devices by up to a factor of 20. In addition, Virtual Chassis Fabric technology offers plug-and-play functions such as autodiscovery of spines and leafs and the automatic election of a master Routing Engine (RE).

Key Features and Benefits

Virtual Chassis Fabric technology delivers the following features and benefits:

- **A single entity to manage:** Virtual Chassis Fabric technology gives IT a single logical point of management for all interconnected switches, presenting one management IP address that is accessible from any member of the fabric. One spine switch acts as the master device, handling all communications and control for the fabric. In the event the master goes down, an automated election process based on graceful Routing Engine switchover (GRES) seamlessly transfers control to a new master with no interruption in service.
- **Zero-touch, plug-and-play provisioning:** Juniper designed Virtual Chassis Fabric configurations to be much easier to deploy and manage than a traditional three-tier network. For example, by factory default, supported devices automatically join the fabric. In addition, the master RE can push out configuration information to other switches, automatically provisioning spine-and-leaf nodes. Configuration and image synchronization are also supported.
- **Smart network management:** Customers can manage a Virtual Chassis Fabric configuration using standard CLI commands as well as the Junos® Space Network Director network management application. Integrated with the QFX5100 line, Insight Technology for Analytics helps IT make better network design decisions and identify network hotspots by providing dynamic buffer utilization monitoring and reporting. With an interval of 10 milliseconds, Insight provides microburst and latency details, capturing and reporting microburst events that exceed defined thresholds. The data can be viewed via CLI, system log, or streamed to external servers for more analysis—for example, to Network Director.

Junos Space Network Director is a comprehensive, automated network management solution that lets you visualize, analyze, and control the entire enterprise network through a single pane of glass. For example, Network Director helps synchronize physical and virtual environments, ensuring that network policies follow workloads as they move from server to server or from VM to VM. It also automates routine management tasks such as network provisioning and troubleshooting, dramatically improving operational efficiency and reliability. Network Director comes with a set of RESTful APIs that provide a single-point interface to orchestration tools such as OpenStack and CloudStack for end-to-end configuration and management of network services.

- **L2 and L3 capabilities on the same device:** Virtual Chassis Fabric provides rich L2 and L3 functionality, which gives IT the flexibility to use the connectivity method that works best for the applications being supported. Unlike competitors that charge for individual features, Juniper offers license bundles that are cost-effective and simplify the purchasing process. For example, the base license for Virtual Chassis Fabric includes L2 functionality along with L3 routing for IPv4 and IPv6, while the advanced feature license gives you MPLS, BGP, and IS-IS on all Virtual Chassis Fabric ports.

In addition, Virtual Chassis Fabric includes:

- L2 and L3 multicasting for up to 60,000 routes using a unified forwarding table and Bidirectional Multicast Distribution Trees, providing predictable latency and replication points and automatic load rebalancing if there's a topology change
- Fibre Channel over Ethernet (FCoE) capabilities, so organizations can easily consolidate storage networks onto the data network
- Support for standards-based network virtualization protocols such as Virtual Extensible LAN (VXLAN), which enables pods to be connected over a WAN, Open vSwitch Database (OVSDB) protocol, Network Virtualization using Generic Routing Encapsulation (NVGRE), and integration with Juniper Networks Contrail and VMware NSX Layer 2 Gateway Services functionality to programmatically enable connectivity between VMware virtual networks and physical network environments.
- **Resiliency and high availability:** Virtual Chassis Fabric technology has robust redundancy features that ensure reliability and uptime as well as nondisruptive maintenance operations. These features include redundant Routing Engines operating in an active and hot-backup mode that ensure forwarding continues during RE failover; data plane redundancy with active/active uplink forwarding; and uplink redundancy as well as server multihoming. On the individual switches, hot-swappable components such as power supplies, fans, and expansion modules provide additional levels of redundancy and availability at the hardware level.
- **Nonstop software upgrade (NSSU):** A mechanism for upgrading Junos OS on EX Series and QFX Series Ethernet switches through a single command with minimal traffic disruption, NSSU is available on any Virtual Chassis Fabric with redundant Routing Engines. NSSU leverages underlying high availability features such as GRES, NSR, and NSB to upgrade the Junos OS version running on a switch or in a Virtual Chassis Fabric without disrupting the control plane, and only minor, sub-second interruptions to the data plane. Member switches are upgraded individually, permitting traffic to continue flowing through devices that are not being upgraded. (For more information about NSSU in a Virtual Chassis Fabric, visit http://www.juniper.net/documentation/en_US/junos14.1/topics/concept/nssu-vcf.html.)
- **Investment protection:** Juniper designed its new fabric architecture with flexibility and investment protection in mind. For example, customers who have deployed the EX4300 Ethernet Switch, or the QFX3500, QFX3600, or QFX5100 switches in a Virtual Chassis configuration can easily migrate to Virtual Chassis Fabric. Because the QFX5100 was designed to be a universal building block for all of Juniper's data center fabric architectures, customers can use this switch in Virtual Chassis and Virtual Chassis Fabric deployments, providing investment protection and an easy migration path as needs change.

A Deep Dive Look at the Virtual Chassis Fabric

In addition to the switch hardware, a Virtual Chassis Fabric configuration includes robust implementations of the management, control, and forwarding planes that deliver management simplicity, high availability, and exceptional performance.

- **Management plane:** The management plane for the Virtual Fabric Chassis enables up to 20 switches to be managed as a single device, greatly simplifying all aspects of management, from initial configuration and deployment to ongoing operations and upgrades. Key capabilities include:
 - Auto discovery: All switches attached to the Virtual Chassis Fabric configuration automatically find each other by exchanging "hello" messages using Link Layer Discovery Protocol (LLDP).




Figure 3: Autodiscovery with Virtual Chassis Fabric technology

- Plug-and-play auto-configuration: With this capability, IT only needs to preconfigure spine switches; all directly connected leaf switches in factory default mode will automatically join the Virtual Chassis Fabric as “line cards.” Also known as zero-touch provisioning, this capability provides policy-driven provisioning and network bring-up that simplifies and speeds network deployment and reduces downtime due to human error.
- Auto upgrade: The system will detect the models of each switch node and the master RE will automatically push the right software image to them. Likewise, any configuration changes or updates are also handled automatically by the master RE, which also ensures configuration and image synchronization.
- **Control plane:** Virtual Chassis Fabric technology features a fully redundant, integrated control plane and up to four integrated, redundant REs that provide centralized control for network ports. The control plane provides automatic fabric topology discovery and loop-free forwarding based on shortest path calculation for unicast traffic, and bidirectional multicast distribution trees for multicast and broadcast traffic. To protect control traffic, the system maintains a separate queue with its own buffer threshold and priority scheduling.

Virtual Chassis Fabric technology employs quad RE redundancy. In this model, if the master RE fails, the backup takes over the master role and one of the remaining spine switches is promoted to backup. A key benefit of this model is that forwarding remains intact during RE failovers. GRES, NSR, and nonstop bridging (NSB) are all supported, along with NSSU, enabling you to upgrade Junos OS releases using a single command with minimal traffic disruption.




Figure 4: Virtual Chassis Fabric integrated control plane

- **Data Plane:** The Virtual Chassis Fabric data plane provides low latency and predictable performance for both L2 and L3 by supporting local switching on all ports and having all fabric links operating in active/active mode, with traffic load balanced on all links. Virtual Chassis Fabric provides 550 nsec in-rack latency and 1.8 usec inter-rack (port to port) latency, with no more than three hops rack to rack.

In addition to active/active uplink forwarding, data plane redundancy includes uplink redundancy and 16-way server multihoming. The data plane also supports smart trunks, whereby fabric links are automatically aggregated into trunks, including next-hop trunks, from local to direct neighbors, remote destination trunks, and from local to a remote destination.

Virtual Chassis Fabric Use Cases

Virtual Chassis Fabric was designed to address the requirements of financial services, federal government, healthcare, and other organizations that are deploying modern applications and need a high-performance network that also helps keep costs under control through simplified management. By providing any-to-any connectivity between switch ports, Virtual Chassis Fabric maximizes performance for physical servers, virtual servers, and storage in the data center. Virtual Chassis Fabric also supports VM mobility, ensuring that the network can keep up with the dynamic nature of today's applications.

In addition to high-performance application environments, Virtual Chassis Fabric is ideally suited for FCoE transit deployment and VM mobility beyond a pod.

Conclusion—Purpose-Built for Modern Applications in Today's Data Center

Juniper Networks Virtual Chassis Fabric technology addresses the needs of today's small and medium-sized data centers for a high-performance, easy-to-manage network architecture. Optimized for mixed 1GbE/10GbE/40GbE environments, Virtual Chassis Fabric lets you interconnect up to 20 switches in a high-speed, low-latency fabric, where all switches are managed as a single device.

Virtual Chassis Fabric is a key offering in Juniper's growing fabric portfolio. Through the use of common building blocks such as the QFX5100 line of switches, Juniper fabric solutions offer an evolutionary, cost-effective approach to data center networking with investment protection. Enterprises can start small and build up to a full-scale Virtual Chassis Fabric configuration, even build more than one fabric and interconnect them.

Virtual Chassis Fabric delivers the performance that modern applications need at a scale and price designed for mid-size application deployments in small to medium-sized data centers. With its single point of management, Virtual Chassis Fabric reduces management costs and simplifies change, so your organization can readily adapt to application and business needs—today, tomorrow, and beyond.

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 2016 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.

JUNIPER
NETWORKS