Juniper Networks

Data Center Glossary

Published: 2015-05-01

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

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 Data Center Glossary

The information in this document is current as of the date on the title page.

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at http://www.juniper.net/support/eula.html. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

About This Glossary 1
Data Center Glossary 1

About This Glossary

This glossary provides definitions of data center networking related terms. For definitions of other networking terms, see the Juniper Networks Glossary.

Data Center Glossary

A

active flow monitoring Flow monitoring carried out on the same router that forwards the packets being monitored. In contrast, a passive monitoring router does not forward the packets being monitored—it receives mirrored packets from a router that performs the packet forwarding. *See also* flow monitoring.

ADVPN Auto Discovery VPN. Protocol that enables dynamic establishment of spoke-to-spoke VPN tunnels. When passing traffic from one spoke to another spoke, the hub can suggest that the spokes establish a shortcut between each other. Shortcuts can be established and torn down dynamically between spokes, resulting in better network resource utilization and less reliance on a centrally located hub.

Auto Discovery VPN ADVPN. Protocol that enables dynamic establishment of spoke-to-spoke VPN tunnels. When passing traffic from one spoke to another spoke, the hub can suggest that the spokes establish a shortcut between each other. Shortcuts can be established and torn down dynamically between spokes, resulting in better network resource utilization and less reliance on a centrally located hub.

autoinstallation Automatic installation and configuration of software on a device over the network from a preexisting configuration file stored on a configuration server—typically a Trivial File Transfer Protocol (TFTP) server. Autoinstallation occurs on a device that is powered on without a valid configuration (boot) file or that is configured specifically for autoinstallation. Autoinstallation is especially useful when multiple devices must be configured and deployed on a network.

automation A broad term that encompasses many levels of automating network functions. Automation can refer to managing virtual resources and/or physical resources. Network automation capabilities can include the configuration and provisioning of network devices, spinning up and spinning down network services and applications as needed, managing network devices and services, and enforcing service-level agreements. Automation reduces the operational overhead of network configuration, provisioning, and management.

B

BFD Bidirectional Forwarding Detection. Protocol that uses control packets and shorter detection time limits (than the default failure detection mechanisms for Layer 3 protocols) to more rapidly detect failures in a network.

BGP Border Gateway Protocol. BGP is used as a control plane protocol in IP fabrics to support the DC overlay network. BGP provides L3 connectivity between every host in the network that participates in the overlay network.

Bidirectional Forwarding Detection BFD. Protocol that uses control packets and shorter detection time limits to more rapidly detect failures in a network.

Blade server A thin server in a rack, generally dedicated to a single application.

Border Gateway Protocol BGP. BGP is used as a control plane protocol in IP fabrics to support the DC overlay network. BGP provides L3 connectivity between every host in the network that participates in the overlay network.

C

chassis cluster Physically connected and configured devices that provide redundancy and ensure service continuity in the event of partial or complete device failure. Chassis clusters provide a resilient system architecture, synchronizing session and kernel states across control and data planes to prevent a single point of failure from disabling the network.

Clos network fabric Multistage switching network in which switch elements in the middle stages are connected to all switch elements in the ingress and egress stages. Clos networks are well-known for their nonblocking properties—a connection can be made from any available input port to any available output port, regardless of the traffic load in the rest of the system.

cloud Internet based environment of virtualized computing resources, including servers, software, and applications that can be accessed by individuals or businesses with Internet connectivity. Cloud types include public, private, and hybrid.

cloud computing Cloud computing provides on-demand access to a shared pool of configurable network computing resources such as servers, storage, services, applications, and network devices. Cloud computing offers fast resource provisioning and scaling, that enables cloud computing customers to spin up services quickly, with minimal in-house management and resources. Cloud computing is the basis for Infrastructure as a Service (IaaS) and Software as a Service (SaaS). *See also* IaaS and SaaS.

control plane Virtual network path used to set up, maintain, and terminate data plane connections. *See also* data plane.

converged network adapter (CNA) Physical adapter that combines the functions of a Fibre Channel *host bus adapter (HBA)* to process FCoE frames and a *lossless Ethernet network interface card (NIC)* to process non-FCoE Ethernet frames. CNAs have one or more Ethernet ports. CNAs encapsulate Fibre Channel frames in Ethernet for FCoE transport and de-encapsulate Fibre Channel frames from FCoE to native Fibre Channel. *See also* host bus adapter.

D

data center A physical or virtual infrastructure that houses computer, server, data storage, management, and networking systems and components in a secure environment. Data centers often have redundant power, cooling, connectivity, storage, and network systems to prevent loss of service and data.

data center bridging DCB. Set of IEEE specifications that enhances the Ethernet standard to enable it to support converged Ethernet (LAN) and Fibre Channel (SAN) traffic on one Ethernet network. DCB features include priority-based flow control (PFC), enhanced transmission selection (ETS), Data Center Bridging Capability Exchange protocol (DCBX), quantized congestion notification (QCN), and full-duplex 10-Gigabit Ethernet ports.

Data Center Bridging Capability Exchange protocol DCBX. Discovery and exchange protocol for conveying configuration and capabilities among neighbors to ensure consistent configuration across the network. It is an extension of the Link Layer Data Protocol (LLDP, described in IEEE 802.1ab, *Station and Media Access Control Connectivity Discovery*).

Data Center Interconnect The connection between two or more data centers that are physically (geographically) separated.

data plane Virtual network path used to distribute data between nodes. *Also known as* transport plane. *See also* control plane.

DCB Data center bridging. Set of IEEE specifications that enhances the Ethernet standard to allow it to support converged Ethernet (LAN) and Fibre Channel (SAN) traffic on one Ethernet network. DCB features include priority-based flow control (PFC), enhanced transmission selection (ETS), Data Center Bridging Capability Exchange protocol (DCBX), quantized congestion notification (QCN), and full-duplex 10-Gigabit Ethernet ports.

DCBX Data Center Bridging Capability Exchange protocol. Discovery and exchange protocol for conveying configuration and capabilities among neighbors to ensure consistent configuration across the network. It is an extension of the Link Layer Data Protocol (LLDP, described in IEEE 802.1ab, *Station and Media Access Control Connectivity Discovery*).

Director group Entity that controls the QFabric System. A Director Group is composed of two Director devices (DG0 and DG1) that act in a synchronized master/slave relationship for redundancy. The Director Group hosts the necessary virtual machine (VM) components to run and maintain the system.

E

enhanced transmission selection ETS. A hierarchical scheduling mechanism that provides better bandwidth utilization and finer granularity of bandwidth management within a link, as described in IEEE 802.1Qaz.

ESI Ethernet segment identifier. A 10-octet value ranging from 0x00 through 0xFFFFFFFFFFFFFFFFFFFF that represents the Ethernet segment (ES). An ESI must be set to a networkwide, unique, nonreserved value when a customer edge (CE) device is multihomed to two or more provider edge (PE) devices. For a single-homed CE device, the reserved ESI value 0 is used. The ESI value of “all FFs” is also reserved.

ESX, VMWare ESXi Enterprise-level software hypervisors from VMware that do not need an additional operating system to run on host server hardware.

Ethernet PAUSE As defined in IEEE 802.3X, a flow control mechanism that temporarily stops the transmission of Ethernet frames on a link for a specified period. A receiving element sends an Ethernet PAUSE frame when a sender transmits data faster than the receiver can accept it. Ethernet PAUSE affects the entire link, not just an individual flow. An Ethernet PAUSE frame temporarily stops all traffic transmission on the link.

Ethernet segment Ethernet links between a customer edge (CE) device and one or more provider edge (PE) devices. In a multihomed topology, the set of links between the CE device and PE devices is considered a single Ethernet segment. Each Ethernet segment is assigned an identifier.

Ethernet segment identifier ESI. A 10-octet value ranging from 0x00 through 0xFFFFFFFFFFFFFFFFFFFF that represents the Ethernet segment (ES). An ESI must be set to a networkwide, unique, nonreserved value when a customer edge (CE) device is multihomed to two or more provider edge (PE) devices. For a single-homed CE device, the reserved ESI value 0 is used. The ESI value of “all FFs” is also reserved.

Ethernet tag identifier Tag that identifies the broadcast domain in an Ethernet VPN (EVPN) instance. The broadcast domain is a VLAN and the Ethernet tag identifier is the VLAN ID.

Ethernet VPN EVPN. Type of VPN that enables you to connect a group of dispersed customer sites by using a Layer 2 virtual bridge. As with other types of VPNs, an EVPN comprises customer edge (CE) devices (routers or switches) connected to provider edge (PE) devices. The PE devices can include an MPLS edge switch that acts at the edge of the MPLS infrastructure.

ETS enhanced transmission selection. A hierarchical scheduling mechanism that provides better bandwidth utilization and finer granularity of bandwidth management within a link, as described in IEEE 802.1Qaz.

EVI EVPN instance. Routing and forwarding instance configured on provider edge (PE) devices to create the EVPN service.

EVPN Ethernet VPN. Type of VPN that enables you to connect a group of dispersed customer sites by using a Layer 2 virtual bridge. As with other types of VPNs, an EVPN comprises customer edge (CE) devices (routers or switches) connected to provider edge (PE) devices. The PE devices can include an MPLS edge switch that acts at the edge of the MPLS infrastructure.

EVPN instance EVI. Defined on provider edge (PE) devices to create the EVPN service.

F

fabric Interconnection of network nodes using one or more network switches that function as a single logical entity.

fabric schedulers Schedulers that identify a packet as high or low priority based on its forwarding class, and associate schedulers with the fabric priorities.

FC Fibre Channel. High-speed network technology used for storage area networks (SANs).

FC forwarder FCF. Fibre Channel switch that has all physical Fibre Channel ports and the necessary set of services as defined in the T11 Organization *Fibre Channel Switched Fabric* (FC-SW) standards.

FCF FC forwarder, FCoE forwarder. The two types of forwarders are:

- FC forwarder. Fibre Channel switch that has all physical Fibre Channel ports and the necessary set of services as defined in the T11 Organization *Fibre Channel Switched Fabric* (FC-SW) standards.
- FCoE forwarder. Device that has the necessary set of services defined in the T11 Organization *Fibre Channel Switched Fabric* (FC-SW) standards and that has the Fibre Channel over Ethernet (FCoE) capabilities to act as an FCoE-based Fibre Channel switch, as defined by the Fibre Channel Backbone – 5 (FC-BB-5) Rev. 2.00 specification.

FCoE Fibre Channel over Ethernet. Standard for transporting FC frames over Ethernet networks. FCoE encapsulates Fibre Channel frames in Ethernet so that the same high-speed Ethernet physical infrastructure can transport both data and storage traffic while preserving the lossless CoS that FC requires. FCoE servers connect to a switch that supports both FCoE and native FC protocols. This enables FCoE servers on the Ethernet network to access FC storage devices in the SAN fabric on one converged network.

FCoE forwarder FCF. Device that has the necessary set of services defined in the T11 Organization Fibre Channel Switched Fabric (FC-SW) standards and that has the FCoE capabilities to act as an FCoE-based Fibre Channel switch, as defined by the Fibre Channel Backbone – 5 (FC-BB-5) Rev. 2.00 specification.

FCoE Initialization Protocol FIP. Layer 2 protocol that establishes and maintains Fibre Channel (FC) virtual links between pairs of FCoE devices such as server FCoE nodes (ENodes) and FC switches.

FCoE Initialization Protocol snooping FIP snooping. Security feature enabled for FCoE VLANs on an Ethernet switch that connects FCoE nodes to Fibre Channel switches or FCFs. The two types of FIP snooping inspect data in FIP frames and use that data to create firewall filters that are installed on the ports in the FCoE VLAN. The filters permit only traffic from sources that perform a successful fabric login to the Fibre Channel switch. All other traffic on the VLAN is denied. FIP snooping can also provide additional visibility into FCoE Layer 2 operation.

FCoE transit switch Switch with a minimum set of features designed to support FCoE Layer 2 forwarding and FCoE security. The switch can also have optional additional features. Minimum feature support is:

- Priority-based flow control (PFC)
- Enhanced transmission selection (ETS)
- Data Center Bridging Capability Exchange (DCBX) protocol, including the FCoE application TLV
- FIP snooping (minimum support is FIP automated filter programming at the ENode edge)

A transit switch has a Fibre Channel stack even though it is not a Fibre Channel switch or an FC forwarder.

FCoE VLAN Fibre Channel over Ethernet VLAN. VLAN dedicated to carrying FCoE traffic. FCoE traffic must travel in a VLAN. Only FCoE interfaces should be members of an FCoE VLAN. Ethernet traffic that is not FCoE traffic must travel in a different VLAN.

Fibre Channel FC. High-speed network technology used for storage area networks (SANs).

Fibre Channel fabric Network of Fibre Channel devices that provides communication among devices, device name lookup, security, and redundancy.

Fibre Channel over Ethernet FCoE. Standard for transporting FC frames over Ethernet networks. FCoE encapsulates Fibre Channel frames in Ethernet so that the same high-speed Ethernet physical infrastructure can transport both data and storage traffic while preserving the lossless CoS that FC requires. FCoE servers connect to a switch that supports both FCoE and native FC protocols. This enables FCoE servers on the Ethernet network to access FC storage devices in the SAN fabric on one converged network.

Fibre Channel over Ethernet VLAN FCoE VLAN. VLAN dedicated to carrying FCoE traffic. FCoE traffic must travel in a VLAN. All members of an FCoE VLAN must be FCoE interfaces. Ethernet traffic that is not FCoE traffic must travel in a different VLAN.

FIP FCoE Initialization Protocol. Layer 2 protocol that establishes and maintains Fibre Channel (FC) virtual links between pairs of FCoE devices such as server FCoE nodes (ENodes) and FC switches.

FIP snooping FCoE Initialization Protocol snooping. Security feature enabled for FCoE VLANs on an Ethernet switch that connects FCoE nodes to Fibre Channel switches or FCFs. The two types of FIP snooping inspect data in FIP frames and use that data to create firewall filters that are installed on the ports in the FCoE VLAN. The filters permit traffic only from sources that perform a successful fabric login to the Fibre Channel switch. All other traffic on the VLAN is denied. FIP snooping can also provide additional visibility into FCoE Layer 2 operation.

Firefly Perimeter (Now known as vSRX) Services Gateway application that delivers a complete virtual firewall solution, including advanced security, robust networking, and automated virtual machine life cycle management capabilities for service providers and enterprises. *See also* vSRX.

flow monitoring Application that monitors the flow of traffic and enables lawful interception of packets transiting between two routers. Traffic flows can be passively monitored by an offline router or actively monitored by a router participating in the network. *See also* active flow monitoring and passive flow monitoring.

G

graceful restart Process that enables a router whose control plane is undergoing a restart to continue forwarding traffic while recovering its state from neighboring routers. Without graceful restart, a control plane restart disrupts services provided by the router. Implementation varies by protocol. *Also known as* nonstop forwarding.

graceful Routing Engine switchover GRES. In a router that contains a master and a backup Routing Engine, enables the backup Routing Engine to assume mastership automatically, with no disruption of packet forwarding. *Also known as* Stateful Switchover (SSO).

graceful switchover Junos OS feature that enables a primary device, such as a Routing Engine, to start functioning as (or switch over to) the backup device without interrupting packet forwarding.

GRES graceful Routing Engine switchover. In a router that contains a master and a backup Routing Engine, enables the backup Routing Engine to assume mastership automatically, with no disruption to packet forwarding. *Also known as* Stateful Switchover (SSO).

H

HA high availability. Configuring devices to ensure service continuity in the event of a network outage or device failure. Used to provide fault detection and correction procedures to maximize the availability of critical services and applications. High availability provides both hardware-specific and software-specific methods to ensure minimal downtime and ultimately improve the performance of your network. *See also* chassis cluster.

high availability HA. Configuring devices to ensure service continuity in the event of a network outage or device failure. Used to provide fault detection and correction procedures to maximize the availability of critical services and applications. High availability provides both hardware-specific and software-specific methods to ensure minimal downtime and ultimately improve the performance of your network.

hypervisor In cloud computing, platform virtualization software that runs on a host computer, allowing multiple instances of operating systems, called guests, to run concurrently on the host within their own VMs and share virtualized hardware resources. A virtualized software layer that manages the relationships between VMs that run on its host and compete for its resources. A hypervisor controls and manages resource allocation. A hypervisor is said to run on bare metal, that is, directly on the hardware whose resources it shares. The term *hypervisor* was created by IBM to refer to software that is conceptually one level higher than an operating system's supervisor. *Also known as* Virtual Machine Manager (VMM).

I

IaaS Infrastructure as a Service. Physical or virtual cloud servers and other resources such as switches, routers, firewalls, storage devices, load balancers, and other network equipment leased to customers as needed. IaaS providers offer customers the ability to scale services up or down easily without having to make the capital investment in equipment and expertise.

in-service software upgrade ISSU. General term for one of several different ways that Juniper Networks platforms upgrade software versions with minimal disruption to network traffic. Unified ISSU is used for routing platforms, which operate at Layer 2 and Layer 3. Nonstop software upgrade (NSSU) is used for switching platforms that operate at Layer 2 and Virtual Chassis configurations. Topology-independent in-service software upgrade (TISSU) is used for virtual environments, where devices are not linked by a hardware-based topology. *See also* NSSU, TISSU, and unified ISSU.

Infrastructure as a Service IaaS. Physical or virtual cloud servers and other resources such as switches, routers, firewalls, storage devices, load balancers, and other network equipment leased to customers as needed. IaaS providers offer customers the ability to scale services up or down easily without having to make the capital investment in equipment and expertise.

integrated routing and bridging IRB. Provides simultaneous support for Layer 2 (L2) bridging and Layer 3 (L3) routing within the same bridge domain. Packets arriving on an interface of the bridge domain are L2 switched or L3 routed based on the destination MAC address. Packets addressed to the router's MAC address are routed to other L3 interfaces.

Interconnect device QFabric system component that acts as the primary fabric for data plane traffic traversing the QFabric system between Node devices. Examples of Interconnect devices include the QFX3008-I Interconnect device in a QFX3000-G QFabric system, the QFX5100-24Q configured as an Interconnect device, and the QFX3600-I Interconnect device in a QFX3000-M QFabric system.

IP fabric A Layer 3 network fabric (all network switches use only Layer 3 connectivity), often using BGP as the Layer 3 protocol.

IP VPN A Layer 3 VPN service implemented using BGP/MPLS IP VPNs (RFC 4364).

IRB integrated routing and bridging. Provides simultaneous support for Layer 2 (L2) bridging and Layer 3 (L3) routing within the same bridge domain. Packets arriving on an interface of the bridge domain are L2 switched or L3 routed based on the destination MAC address. Packets addressed to the router's MAC address are routed to other L3 interfaces.

ISSU in-service software upgrade. General term for one of several different ways that Juniper Networks platforms upgrade software versions with minimal disruption to network traffic. Unified ISSU is used for routing platforms, which operate at Layer 2 and Layer 3. Nonstop software upgrade (NSSU) is used for switching platforms that operate at Layer 2 and Virtual Chassis configurations. Topology-independent in-service software upgrade (TISSU) is used for virtual environments, where devices are not linked by a hardware-based topology. *See also* NSSU, TISSU, and unified ISSU.

J

Junos Fusion A method of significantly expanding the number of available network interfaces on a device—an *aggregation device*—by allowing the aggregation device to add interfaces through interconnections with *satellite devices*. The entire system—the interconnected aggregation device and satellite devices—is called a Junos Fusion. A Junos Fusion simplifies network administration because it appears to the larger network as a single, port-dense device that is managed using one IP address.

Junos Space Carrier-class network management system for provisioning, monitoring, and diagnosing Juniper Networks routing, switching, security, and data center platforms.

L

latency Time taken by a packet to reach its destination after it has left the source.

Leaf device In a spine-and-leaf network architecture, leaf devices act as the network access layer, and spine devices act as core and aggregation devices. Leaf devices connect to servers on the access side, and uplink to the spine devices in a mesh (each leaf device connects to every spine device). Spine and leaf architecture provides predictable behavior and latency characteristics because every device is the same number of hops from every other device.

line rate Total number of physically transferred bits per second, including useful data and protocol overhead, over a communication link. For example, if the line rate of a link is 10 Gbps, the link transmits 10 gigabits of data every second over its physical interface.

M

MAC-VRF MAC address virtual routing and forwarding table. This is the Layer 2 forwarding table on a provider edge (PE) device for an Ethernet VPN instance (EVI). *See also* EVI.

MetaFabric The Juniper architecture for next-generation data centers that simplifies and accelerates the deployment and delivery of applications within and across multiple data center locations.

Multiple VLAN Registration Protocol MVRP. A Layer 2 implementation of the Multiple Registration Protocol (MRP) that is defined in the IEEE 802.1ak standard. MRP and MVRP were designed to perform the same functions as Generic Attribute Registration Protocol (GARP) and GARP VLAN Registration Protocol (GVRP) while overcoming some GARP and GVRP limitations, in particular, limitations involving bandwidth usage and convergence time in large networks with large numbers of VLANs.

multitenancy In cloud computing, many virtual machines (VMs) can run concurrently on a virtualized host whose hypervisor manages how resources are shared among the VMs. The VMs are referred to as guest, or tenant, VMs. To users of these guest VMs, it is as if they are running a single, physical machine to which resources are dedicated. Each guest VM runs its own operating system image and user space applications.

MVRP Multiple VLAN Registration Protocol. A Layer 2 implementation of the Multiple Registration Protocol (MRP) that is defined in the IEEE 802.1ak standard. MRP and MVRP were designed by IEEE to perform the same functions as Generic Attribute Registration Protocol (GARP) and GARP VLAN Registration Protocol (GVRP) while overcoming some GARP and GVRP limitations, in particular, limitations involving bandwidth usage and convergence time in large networks with large numbers of VLANs.

N

NaaS Network as a Service. Virtualized network infrastructure leased by a cloud provider to customers as needed.

NEBS Network Equipment Building System. Set of guidelines originated by Bell Laboratories in the 1970s to assist equipment manufacturers in designing products that were compatible with the telecom environment.

Network as a Service NaaS. Virtualized network infrastructure leased by a cloud provider to customers as needed.

Network Director

A Junos Space application that provides a comprehensive automated network management solution for the enterprise data center and campus. It enables network and cloud administrators to visualize, analyze, and control their entire enterprise network—data center and campus, physical and virtual infrastructure, virtual overlays networks, and wired and wireless—through a single pane of glass.

Network Functions Virtualization

NFV. Standard IT virtualization technology that consolidates many network equipment types onto standard-architecture high-volume servers, switches, and storage. NFV involves designing, deploying, and managing network functions in software that can be moved to, or instantiated in, various locations in the network as required, without the need to install purpose-built hardware. Although NFV complements software-defined networking (SDN), NFV can be deployed without SDN and vice versa. *See also* software-defined networking.

Network Node Group

Each QFabric System has one network Node group and up to eight physical Nodes can be configured to be part of the network Node Group. The devices in the network Node Group connect to an external network. The network Node group relies on two external Routing Engines running on the Director group. These redundant network Node group Routing Engines run the routing protocols required to support the connections from the network Node group to external networks.

Network Virtualization using Generic Routing Encapsulation

NVGRE. A network virtualization technology that uses generic routing encapsulation (GRE) to encapsulate network protocols and tunnel data across Layer 3 networks to provide a network overlay (VXLAN and STT are examples of two other network overlay technologies). NVGRE offers the ability to create large numbers of VLANs to address the problems caused by the limited number of VLANs (4,096) that the IEEE 802.1Q specification permits. The larger number of VLANs means that larger networks can scale up to a larger number of separate virtual networks.

network-attached storage (NAS)

Dedicated file server and storage for multiple clients, with the ability to share files among multiple users.

NFV

Network Functions Virtualization. Standard IT virtualization technology that consolidates many network equipment types onto standard-architecture high-volume servers, switches, and storage. NFV involves designing, deploying, and managing network functions in software that can be moved to, or instantiated in, various locations in the network as required, without the need to install purpose-built hardware. Although NFV complements software-defined networking (SDN), NFV can be deployed without SDN and vice versa. *See also* software-defined networking.

Node device

Routing and switching device that connects to endpoints (such as servers or storage devices) or external network peers, and is connected to the QFabric system through an Interconnect device. You can deploy Node devices similarly to the way a top-of-rack switch is implemented. Examples of Node devices include the QFX3500 Node device, QFX3600 Node device, and QFX5100 Node device. *See also* Interconnect device.

Nonstop bridging (NSB)

Feature that enables a transparent switchover mechanism for Layer 2 protocol sessions.

nonstop software upgrade NSSU. Software upgrade for switching platforms with redundant Routing Engines and for most Virtual Chassis or Virtual Chassis Fabric from one Junos OS release to another with no disruption on the control plane and with minimal disruption to network traffic. A switching architecture requires a different approach than the one for a routing architecture to preserve control plane information. *See also* ISSU, TISSU, and unified ISSU.

NSR nonstop active routing. High availability feature that enables a routing platform with redundant Routing Engines to preserve routing information on the backup Routing Engine and switch over from the primary Routing Engine to the backup Routing Engine without alerting peer nodes that a change has occurred. NSR uses the graceful Routing Engine switchover (GRES) infrastructure to preserve interface, kernel, and routing information. *Also known as* nonstop routing (NSR).

NSSU nonstop software upgrade. Software upgrade for switching platforms with redundant Routing Engines and for most Virtual Chassis or Virtual Chassis Fabric from one Junos OS release to another with no disruption on the control plane and with minimal disruption to network traffic. A switching architecture requires a different approach than the one for a routing architecture to preserve control plane information. *See also* ISSU, TISSU, and unified ISSU.

NVGRE Network Virtualization using Generic Routing Encapsulation. A network virtualization technology that uses generic routing encapsulation (GRE) to encapsulate network protocols and tunnel data across Layer 3 networks to provide a network overlay (VXLAN and STT are examples of two other network overlay technologies). NVGRE offers the ability to create large numbers of VLANs to address the problems caused by the limited number of VLANs (4,096) that the IEEE 802.1Q specification permits. The larger number of VLANs means that larger networks can scale up to a larger number of separate virtual networks.

O

Open Virtualization Archive OVA. Compressed archive file (in TAR format) of an Open Virtualization Format (OVF) package. An OVF package contains the components (operating system, middleware, and software applications) needed to install a virtual appliance or a virtual machine. *See also* Open Virtualization Format.

Open Virtualization Format OVF. Platform-independent packaging and distribution method for software to be run on virtual machines (VMs). The OVF supports industry-standard content verification and integrity checking and provides a basic scheme for managing software licensing. As described by the standard, the OVF defines an *open, secure, portable, efficient, and extensible format for the packaging and distribution of software to be run in virtual machines*. An OVF package consists of several files placed in one directory. *See also* Open Virtualization Archive.

Open vSwitch Database Management Protocol OVSDB. OpenFlow management protocol for Open vSwitch implementations.

orchestration Orchestration manages cloud-based and on-premises networks in a coordinated, automated fashion, to align network resources with providing customer services and meeting business requirements. Orchestration uses automation to provide services by defining policies and service levels, then applying the policies and service levels using applications that have automated workflows. An example of orchestration is a web application that a customer uses to request new services that require network resources, in which the application automatically configures and implements the customer request.

Out-of-band management Use of a dedicated channel for managing network devices. Most Juniper Networks devices have a management port with an RJ-45 connector that you can use to connect the device to a management device for out-of-band management.

OVA Open Virtualization Archive. Compressed archive file (in TAR format) of an Open Virtualization Format (OVF) package. An OVF package contains the components (operating system, middleware, and software applications) needed to install a virtual appliance or a virtual machine. *See also* Open Virtualization Format.

Overlay network A logical, separate network that runs on top of an existing physical infrastructure. In a data center, an overlay network is typically used to create a virtual network by encapsulating traffic between virtual switches and tunneling the traffic over the physical network.

OVF Open Virtualization Format. Platform-independent virtual machines (VMs) packaging and distribution method. The OVF supports industry-standard content verification and integrity checking and provides a basic scheme for managing software licensing. As described by the standard, the OVF defines an *open, secure, portable, efficient, and extensible format for the packaging and distribution of software to be run in virtual machines*. An OVF package consists of several files placed in one directory. The Open Virtualization Archive (OVA) is an alternative method that uses a TAR file containing the OVF directory.

OVSDB Open vSwitch Database Management Protocol. OpenFlow management protocol for Open vSwitch implementations.

P

PaaS Platform as a Service. An entire computing platform (often including operating system, programming environment, web server, and database applications), leased by a cloud provider to customers as needed. Customers can use the computing platform to develop applications and run them in the cloud environment to provide services to their clients.

passive flow monitoring Flow monitoring carried out on a routing platform such as a monitoring station that is not participating in the network. A passive monitoring router does not forward the packets being monitored—it receives mirrored packets from a router that performs the packet forwarding. In contrast, in active flow monitoring, the monitoring is carried out on the same router that forwards the packets being monitored. *See also* flow monitoring.

peering Process in which two networks connect and exchange traffic.

PFC priority-based flow control. Link-level flow control mechanism defined by IEEE 802.1Qbb that enables independent flow control for each class of service to ensure that no frame loss from congestion occurs in data center bridging networks. PFC is an enhancement of the Ethernet PAUSE mechanism, but PFC controls classes of flows, whereas Ethernet PAUSE indiscriminately pauses all of the traffic on a link. *Also known as* priority flow control.

Platform as a Service PaaS. An entire computing platform (often including operating system, programming environment, web server, and database applications), leased by a cloud provider to customers as needed. Customers can use the computing platform to develop applications and run them in the cloud environment to provide services to their clients.

PMSI Provider multicast service interface. A logical interface in a provider edge (PE) device that is used to deliver multicast packets from a customer edge (CE) devices to remote PE devices in the same VPN, destined to CE devices.

port mirroring Method by which a copy of an IPv4 or IPv6 packet is sent from the routing or switching platform to an external host address or a packet analyzer for analysis. *Also known as* traffic mirroring, switch port analyzer (SPAN), and lawful intercept.

priority-based flow control PFC. Link-level flow control mechanism defined by IEEE 802.1Qbb that enables independent flow control for each class of service to ensure that no frame loss from congestion occurs in data center bridging networks. PFC is an enhancement of the Ethernet PAUSE mechanism, but PFC controls classes of flows, whereas Ethernet PAUSE indiscriminately pauses all of the traffic on a link. *Also known as* priority flow control. *See also* Ethernet PAUSE.

private cloud A type of cloud implemented in a proprietary network or data center that uses cloud computing technologies to create a virtualized infrastructure operated solely for a single organization, whether it is managed internally or externally. *See also* public cloud.

Provider multicast service interface PMSI. A logical interface in a provider edge (PE) device that is used to deliver multicast packets from a customer edge (CE) devices to remote PE devices in the same VPN, destined to CE devices.

public cloud A cloud type in which a hosting service provider makes resources such as applications, storage, and CPU usage available to the public. Public clouds must be based on a standard cloud computing model. *See also* private cloud.

Q

QCN Quantized Congestion Notification. Congestion management mechanism defined in IEEE 802.1Qau that sends a congestion notification message through the network to the ultimate source of the congestion. Instead of pausing transmission from the connected peer (as priority-based flow control does), QCN tries to stop congestion at its source—the network edge where the end host originates the congestion-causing flow. The idea is that instead of pushing a flow control message through the network one device at a time (like priority-based flow control), QCN tries to find the cause of congestion and stop the flow at the source.

QFabric System Highly scalable, distributed, Layer 2 and Layer 3 networking architecture that provides a high-performance, low-latency, and unified interconnect solution for next-generation data centers. A QFabric system collapses the traditional multi-tier data center model, enables the consolidation of data center endpoints (such as servers, storage devices, memory, appliances, and routers), and provides better scaling and network virtualization capabilities than traditional data centers.

Essentially, a QFabric system can be viewed as a single, nonblocking, low-latency switch that supports thousands of 10-Gigabit Ethernet ports or 2-Gbps, 4-Gbps or 8-Gbps Fibre Channel ports to interconnect servers, storage, and the Internet across a high-speed, high-performance fabric.

Quantized Congestion Notification QCN. Congestion management mechanism defined in IEEE 802.1Qau that sends a congestion notification message through the network to the ultimate source of the congestion. Instead of pausing transmission from the connected peer (as priority-based flow control does), QCN tries to stop congestion at its source—the network edge where the end host originates the congestion-causing flow. The idea is that instead of pushing a flow control message through the network one device at a time (like priority-based flow control), QCN tries to find the cause of congestion and stop the flow at the source.

R

Redundant Server Node Group RSNG. An RSNG consists of two physical Nodes. The Routing Engines on the Nodes operate in an active/backup fashion. You can configure multiple pairs of RSNGs within a QFabric system. These primarily connect to dual-NIC servers.

RSNG Redundant Server Node Group. An RSNG consists of two physical Nodes. The Routing Engines on the Nodes operate in an active/backup fashion (think of a virtual chassis with two member switches). You can configure multiple pairs of RSNGs within a QFabric system. These primarily connect to dual-NIC servers.

S

SaaS Software as a Service. Application software and databases, leased by a cloud provider to customers as needed. The cloud provider provides and runs the data center on which the application software runs.

SAN storage area network. Network whose primary purpose is the transfer of data between computer systems and storage devices. This term is most commonly used in the context of any network that supports block storage, usually iSCSI, Fibre Channel, and FCoE networks.

SDN software-defined networking. Approach to computer networking that uses methods of network abstraction, such as virtualization, to simplify and scale network components and uses software to define and manage network components. SDN separates the data plane, which forwards traffic, from the control plane, which manages traffic flow, and enables users to program network layers. SDN is often used with Network Functions Virtualization (NFV) to allow agile placement of networking services when and where they are needed. By enabling this level of programmability, SDN enables users to optimize their network resources, increase network agility, provide service innovation, accelerate service time-to-market, extract business intelligence, and ultimately enable dynamic, service-driven virtual networks. *See also* NFV.

Security Director

Junos Space application that enables administrators quickly manage all phases of the security policy life cycle for stateful firewall, unified threat management (UTM), intrusion prevention system (IPS), application firewall (AppFW), VPN, and Network Address Translation (NAT) through a centralized web-based interface.

Server Node Group

SNG. Default Node group in a QFabric system. Whenever a Node becomes part of a QFabric system, it comes up as an SNG. These primarily connect to servers that do not need any cross-Node redundancy. The most common examples are servers that have only one NIC.

SNG

Server Node Group. This is the default group and consists of one Node. Whenever a Node becomes part of a QFabric system, it comes up as an SNG. These primarily connect to servers that do not need any cross-Node redundancy. The most common examples are servers that have only one NIC.

Software as a Service

SaaS. Application software and databases, leased by a cloud provider to customers as needed. The cloud provider provides and runs the data center on which the application software runs.

software-defined networking

SDN. Approach to computer networking that uses methods of network abstraction, such as virtualization, to simplify and scale network components and uses software to define and manage network components. SDN separates the data plane, which forwards traffic, from the control plane, which manages traffic flow, and enables users to program network layers. SDN is often used with Network Functions Virtualization (NFV) to allow agile placement of networking services when and where they are needed. By enabling this level of programmability, SDN enables users to optimize their network resources, increase network agility, provide service innovation, accelerate service time-to-market, extract business intelligence, and ultimately enable dynamic, service-driven virtual networks. *See also* NFV.

Spine device

In a spine-and-leaf network architecture, spine devices act as core and aggregation devices, and leaf devices act as the network access layer. Leaf devices connect to servers on the access side, and uplink to the spine devices in a mesh (each leaf device connects to every spine device). Spine and leaf architecture provides predictable behavior and latency characteristics because every device is the same number of hops from every other device.

Spine-and-Leaf Network Architecture

Spine-and-leaf network architecture replaces a three (or more) tier network architecture with a two-tier architecture consisting of a spine tier and a leaf tier. Spine devices act as core and aggregation devices, and leaf devices act as the network access layer. Leaf devices connect to servers on the access side, and uplink to the spine devices in a mesh (each leaf device connects to every spine device). Spine and leaf architecture provides predictable behavior and latency characteristics because every device is the same number of hops from every other device.

storage area network

SAN. Network whose primary purpose is the transfer of data between computer systems and storage devices. This term is most commonly used in the context of any network that supports block storage, usually iSCSI, Fibre Channel, and FCoE networks.

T

TISSU

topology-independent in-service software upgrade. Software upgrade for virtual machine and top-of-rack environments from one software image to another with no disruption to traffic transiting the device. In *topology-independent* virtual environments, devices are not linked by a hardware-based topology and such environments require a different approach for software upgrade than the one for hardware-based environments, which include routers and switches. *See also* ISSU, NSSU, and unified ISSU.

top-of-rack switch

Switch installed on the top of a rack to which all the devices present in the rack are connected to. The top-of-rack switches in turn are connected to aggregation switches. This reduces cabling by avoiding direct connections from the devices in a rack to aggregation switches and makes it easier to identify the point of network failure to a specific rack.

topology-independent in-service software upgrade

TISSU. Software upgrade for virtual machine and top-of-rack environments from one software image to another with no disruption to traffic transiting the device. In *topology-independent* virtual environments, devices are not linked by a hardware-based topology and such environments require a different approach for software upgrade than the one for hardware-based environments, which include routers and switches. *See also* ISSU, NSSU, and unified ISSU.

U

Underlay network

The physical network (compute, storage, and network infrastructure) below an overlay network. Overlay networks abstract physical networks into logical networks, but the physical network that underlays the logical network still transports the data. *See also* overlay network.

unified in-service software upgrade

unified ISSU. Software upgrade for routing platforms from one Junos OS release to another with no disruption of the control plane and with minimal disruption of traffic. Unified ISSU is supported only on platforms with dual Routing Engines. A routing architecture requires a unified approach to preserve routing tables and control plane information. *See also* ISSU, NSSU, and TISSU.

unified ISSU

unified in-service software upgrade. Software upgrade for routing platforms from one Junos OS release to another with no disruption of the control plane and with minimal disruption of traffic. Unified ISSU is supported only on platforms with dual Routing Engines. A routing architecture requires a unified approach to preserve routing tables and control plane information. *See also* ISSU, NSSU, and TISSU.

V

vCenter

The VMware vCenter server, formerly known as VMware VirtualCenter, that centrally manages VMware vSphere environments, allowing administrators control over the virtual environment. The vCenter provides centralized control and visibility at every level of the virtual infrastructure. It manages clusters of ESX/ESXi hosts, including their VMs, hypervisors, and other parts of the virtualized environment.

VCF

Virtual Chassis Fabric. Evolution of the Virtual Chassis feature, which enables you to interconnect multiple devices into a single, logical device inside of a fabric architecture.

Virtual Chassis

Interconnected devices functioning as one logical device. Similar to a Virtual Switching System or a stack.

Virtual Chassis Fabric

VCF. Evolution of the Virtual Chassis feature, which enables you to interconnect multiple devices into a single, logical device inside of a fabric architecture.

Virtual Extensible LAN

VXLAN. A network virtualization protocol defined in RFC 7348 for running an overlay network on a Layer 3 infrastructure in order to connect multiple Layer 2 networks across Layer 3 connections. VXLANs address the scalability issues in large cloud computing deployments caused by the limited number of traditional VLANs. Instead of the 12-bit identifier VLANs use (4,096 unique values), VXLANs use a larger identification field that theoretically allows the creation of more than 16 million unique VXLANs. VXLANs encapsulate Layer 2 frames in Layer 4 UDP packets for transport across Layer 3 connections.

virtual machine

VM. A simulation of a physical machine such as a workstation or a server that runs on a host that supports virtualization. Many VMs can run on the same host, sharing its resources. A VM has its own operating system that can be different from that of other VMs running on the same host.

Virtual Router Redundancy Protocol

VRRP. Protocol that enables you to configure virtual default routers on Fast Ethernet and Gigabit Ethernet interfaces.

virtualization

Technology that abstracts the physical characteristics of a machine, creating a logical version of it, including creating logical versions of entities such as operating systems and various network resources.

VM

virtual machine. A simulation of a physical machine such as a workstation or a server that runs on a host that supports virtualization. Many VMs can run on the same host, sharing its resources. A VM has its own operating system that can be different from that of other VMs running on the same host.

vMotion

VMware technology that allows for transition of active, or live, virtual machines from one physical server to another, undetectable to the user. It enables VMware to migrate a *live* VM (that is, a VM that is still running with no downtime) from one ESXi host to another host on a different physical server. vMotion allows for system maintenance on hosts and offers improved performance if greater capacity is available on another host.

VMware NSX

A network virtualization platform that reproduces the entire network model in software, enabling virtual networks that can be programmatically provisioned and managed independently of the underlying hardware.

VMware vSphere

VMware cloud operating system that can manage large pools of virtualized computing infrastructure, including software and hardware.

VMware vSphere client

Application or software that administers VMware vSphere.

VNI

VXLAN network identifier. In the VXLAN protocol, the 24-bit numeric ID that identifies a VXLAN segment.

vNIC Virtualized network interface card that connects a VM to a vSwitch. A VM can have multiple vNICs. A vNIC presents the same media access control (MAC) interface that a physical interface provides.

VRRP Virtual Router Redundancy Protocol. Protocol that enables you to configure virtual default routers on Fast Ethernet and Gigabit Ethernet interfaces.

vSRX (Formerly known as Firefly Perimeter) Services Gateway application that delivers a complete virtual firewall solution, including advanced security, robust networking, and automated virtual machine life cycle management capabilities for service providers and enterprises. *See also* Firefly Perimeter.

vSwitch A virtualized switch that resides on a physical server and directs traffic among VMs and their virtualized applications. Network activity between colocated VMs transits it.

VTEP VXLAN tunnel endpoint. In the VXLAN protocol, the entity that performs the encapsulation and decapsulation of VXLAN packets.

VXLAN Virtual Extensible LAN. A network virtualization protocol defined in RFC 7348 for running an overlay network on a Layer 3 infrastructure in order to connect multiple Layer 2 networks across Layer 3 connections. VXLANs address the scalability issues in large cloud computing deployments caused by the limited number of traditional VLANs. Instead of the 12-bit identifier VLANs use (4,096 unique values), VXLANs use a larger identification field that theoretically allows the creation of more than 16 million unique VXLANs. VXLANs encapsulate Layer 2 frames in Layer 4 UDP packets for transport across Layer 3 connections.

VXLAN network identifier VNI. In the VXLAN protocol, the 24-bit numeric ID that identifies a VXLAN segment.

VXLAN tunnel endpoint VTEP. In the VXLAN protocol, the entity that performs the encapsulation and decapsulation of VXLAN packets.

Z

zero-touch provisioning Method of automatically provisioning new Juniper Networks switches in your network. When you physically connect a switch to the network and boot it with a factory default configuration, it attempts to upgrade the Junos OS software automatically and autoinstall a configuration file from the network.