

Cisco Nexus 9500 R-Series Line Cards and Fabric Modules

Contents

Introduction	3
Overview of the Cisco Nexus 9500 R-Series Portfolio	3
Cisco Nexus 9500 R-Series ASIC Architecture	5
Large External Buffers	5
Cell-Based Fabric Forwarding.....	5
Forwarding Scalability	5
Telemetry and Visibility	6
Network Capabilities	6
Fabric Module	7
Cisco Nexus 9500 R-Series Line Cards	8
Cisco Nexus N9K-X9636C-R Line Card	8
Cisco Nexus N9K-X9636Q-R Line Card	9
Packet Forwarding on the Cisco Nexus 9500 R-Series Portfolio	10
Packet Forwarding Pipeline	10
Ingress Receiver Packet Processor	11
Packet Header Parsing	11
Tunnel Service Termination	12
Layer 2 and Layer 3 Forwarding Lookup.....	12
Ingress Traffic Classification and Destination Resolution.....	13
Ingress Traffic Manager	13
Ingress VOQ Buffer	13
Scheduling and Flow Control	13
Ingress Transmit Packet Processor	13
Egress Receive Packet Processor.....	14
Egress Traffic Manager.....	14
Egress Transmit Packet Processor.....	14
Multicast Packet Forwarding.....	14
Multistage Replication	15
Conclusion	15
For More Information	15

Introduction

In 2016, with the introduction of 25, 50, and 100 Gigabit Ethernet, the data center switching industry shifted to new standards with greater capacity and new functions. The new 25, 50, and 100 Gigabit Ethernet supplements the previous 10 and 40 Gigabit Ethernet standards at a similar price and power efficiency, providing an approximately 250 percent increase in capacity.

The Cisco Nexus[®] 9500 R-Series line cards and switching modules enable the customer to build higher-performance data center networks with advanced features. Built with Broadcom Jericho application-specific integrated circuits (ASICs), the new R-Series products offer numerous network innovations to address the challenges involved in supporting data center evolution, converged and hyperconverged infrastructure, and virtualized and containerized applications.

This document discusses the hardware architecture of the Cisco Nexus 9500 R-Series portfolio and operation in Cisco NX-OS mode, based on Cisco NX-OS software. The R-Series architecture is an extension to the overall hardware architecture of the Cisco Nexus 9500 platform switches that is published on Cisco.com and that can be downloaded [here](#).

Overview of the Cisco Nexus 9500 R-Series Portfolio

The Cisco Nexus 9500 platform switches offer industry-leading 1-, 10-, 25-, 40-, 50-, and 100 Gigabit Ethernet port density; performance; power efficiency; and cost effectiveness. With the new R-Series line cards and fabric modules, Cisco now also offers increased forwarding scalability and a large off-chip buffer space.

The R-Series portfolio includes the products listed in Table 1.

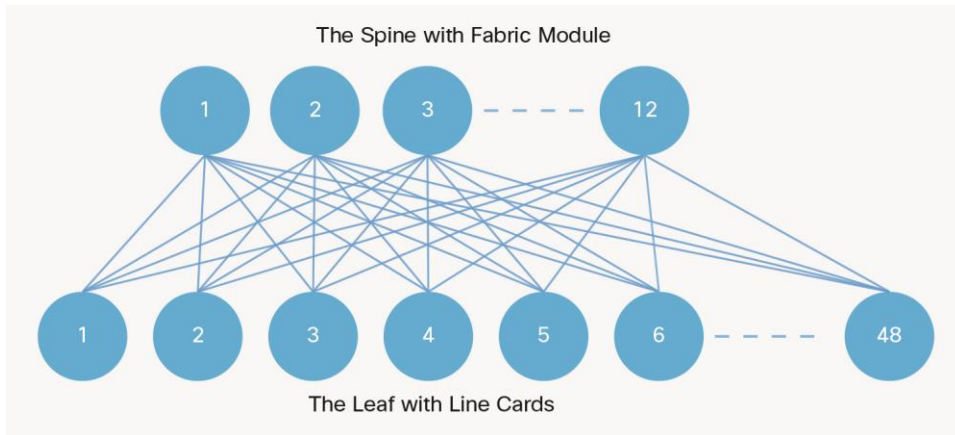
Table 1. Cisco Nexus 9500 R-Series Portfolio

Product	Description
Cisco Nexus N9K-X9636C-R	36 x 100-Gbps Quad Small Form-Factor Pluggable 28 (QSFP28) with 24 GB of packet buffer space
Cisco Nexus N9K-X9636Q-R	36 x 40-Gbps Enhanced QSFP (QSFP+) with 12 GB of packet buffer space
Cisco Nexus N9K-C9508-FM-R	Fabric module for the Cisco Nexus 9508 Switch

The R-Series is fully supported in the existing Cisco Nexus 9508 Switch chassis. This portfolio uses the common components of the Cisco Nexus 9500 chassis: supervisors, system controllers, chassis fan trays, power supply modules, etc. The R-Series portfolio requires the Cisco[®] N9K-SUP-B as the switch supervisor to support the enhanced route scale enabled by the R-Series portfolio. This requirement helps provide investment protection for customers who want to use their current investment in Cisco Nexus 9500 platform switches while using this platform for multiple different applications and requirements.

The Cisco Nexus 9500 platform switches internally implement Clos architecture (often referred to as a fat-tree topology) to connect R-Series fabric modules to the line cards. As shown in Figure 1, the ASICs on the fabric modules form the spine layer, and the ASICs on the line cards form the leaf layer. For example, using the 8-slot Cisco Nexus 9508 Switch, this internal Clos topology can grow up to 12 spine switches and 48 leaf switches (two ASICs per fabric module on the 8-slot chassis and six ASICs per line card).


Figure 1. Internal Folded Clos Topology of Cisco Nexus 9508 Switches



The Clos topology keeps the internal switch architecture simple and consistent with the overall data center network architecture. The R-Series line cards and fabric modules use cell-based forwarding.

Cisco Nexus 9500 platform fabric modules and line cards are interconnected physically through direct attachment with connecting pins. Line cards are horizontal and fabric modules are vertical, so line cards and fabric modules have an orthogonal orientation in the chassis, with each fabric module connected to all line cards, and all line cards connected to all fabric modules. This direct attachment alleviates the need for a switch chassis midplane. Figure 2 depicts the orthogonal interconnection of line cards and fabric modules and the midplane-free chassis of a Cisco Nexus 9500 platform switch.

Figure 2. Cisco Nexus 9500 Platform Line Card and Fabric Module Interconnection



This midplane-free chassis design offers multiple advantages. It allows the most direct and efficient cooling airflow. It increases chassis reliability by eliminating one moving component: the chassis midplane. It also makes upgrading to new line cards or fabric modules easier by eliminating the need to upgrade the midplane for a higher interconnection speed as well as the potential need to replace the chassis. The Cisco Nexus 9500 platform switch chassis thus can support multiple generations of line cards and fabric modules, providing a longer life span for the switch chassis and investment protection for the switch chassis and the common components.

Cisco Nexus 9500 R-Series ASIC Architecture

The Cisco Nexus 9500 portfolio provides a wide range of line cards to meet data center requirements for bandwidth, port density, and high scale, as well as numerous requirements needed by virtual machines and containers. To complement the Cisco Nexus 9500 portfolio, the R-Series includes a large packet buffer to meet the most demanding application requirements for packet buffer depth. The R-Series also supports demanding scalability requirements; cell-based forwarding over fabric; advanced features such as Virtual Extensible LAN (VXLAN) routing, segment routing, and multicast capabilities; and traditional network features such as virtual port channel (vPC).

Large External Buffers

With the new generation of applications in the data center, traffic patterns are unpredictable. In rare cases, these new traffic patterns can require more buffer space to keep sensitive traffic from dropping. The R-Series line cards, in addition to a 16-MB on-chip buffer, have up to 24 GB of graphics double-data-rate 5 (GDDR5) dynamic RAM (DRAM)-based off-chip memory that is used to absorb traffic during congestion. Each forwarding ASIC provides 4 GB of off-chip memory. The off-chip memory is used for the ingress virtual output queue (VOQ) to provide additional buffer depth. Table 2 shows the available buffer capacity for each R-Series line card.

Table 2. Cisco Nexus 9500 R-Series Buffer Capacity

Product	N9K-X9636C-R	N9K-X9636Q-R
Available buffer per ASIC	16-MB on-chip and 4-GB off-chip buffer	16-MB on-chip and 4-GB off-chip buffer
Available buffer per line card	96-MB on-chip and 24-GB off-chip buffer	48-MB on-chip and 12-GB off-chip buffer

The R-Series uses ingress VOQ architecture. VOQ architecture emulates egress queues in the ingress buffer with virtual queues. Each egress port has eight queues for unicast traffic and eight queues for multicast traffic. Traffic can be classified into traffic classes based on the class-of-service (CoS) or differentiated services code point (DSCP) value in the packets and then queued in the corresponding virtual queue for that traffic class.

The R-Series uses a distributed credit mechanism to transfer traffic over the fabric. Before a packet is scheduled to leave the VOQ, the ingress buffer scheduler requests a credit for the specific port and priority in the egress buffer. Credit is requested from an ingress credit scheduler for the destination port and priority. If buffer space is available, the egress scheduler grants access and sends the credit grant to the ingress buffer scheduler. If no buffer space is available in the egress buffer, the egress schedule will not grant a credit, and traffic will be buffered in the VOQ until the next credit is available.

Cell-Based Fabric Forwarding

The R-Series uses cell-based fabric to perform traffic forwarding between the ingress and egress line cards. To provide efficiency for the fabric modules, frames are segmented into variable length cells on the ingress line cards and transferred over the fabric to the egress line card, where they are reassembled into frames.

Efficient use of fabric capacity is achieved with variable-length cells of 64 to 256 bytes. A small frame can use only one cell, and multiple small frames can be packed into one 256-byte or smaller cell.

Forwarding Scalability

The changes in application development and the associated growth in the use of Linux containers and microservers affect many aspects of data center design, including scale requirements. As the servers are disaggregated into their component services—for example, each process or thread now becomes an endpoint—the number of addressable endpoints increases by an order of magnitude or more.

The R-Series line card ASIC responds to these requirements with greater scale.

The R-Series can store 750,000 host entries in the longest exact match (LEM) table, which is shared with the MAC address table. IPv4/IPv6 prefixes are stored in the longest prefix match (LPM) table, which has a capacity of 250,000 entries. The capacity of the routing tables allows storage of a large number of routes: a requirement to support the high-density of virtual machines and containers in today's data centers.

To address the desire for enhanced policy enforcement, the ASIC in the R-Series line cards provides a large access control list (ACL) ternary content-addressable memory (TCAM) table, supporting up to 24,000 entries. Access Lists (ACLs), quality-of-service (QoS) policy, policy-based routing (PBR), and control-plane policing (CoPP) share these TCAM resources.

Table 3 shows the forwarding table scale and ACL TCAM table scale of the forwarding ASIC in Nexus 9500 R-Series line cards.

Table 3. Cisco Nexus 9500 R-Series Line Card Scalability

Entry Type	Cisco Nexus 9500 R-Series ASIC
Layer 2 MAC addresses	750,000*
IPv4 host routes	750,000*
IPv4 LPM entries	250,000*
IPv6 LPM/32 entries	64,000*
TCAM entries	24,000

* Shared entries

Note: The Cisco Nexus 9500 R-Series ASIC is designed and built with the hardware scale described here. For supported software scale in a given software release, refer to the corresponding software release notes and verified scalability guide.

Telemetry and Visibility

The shift to efficient cloud-based provisioning requires organizations to alter the diagnostic and operational characteristics of the data center network. The inherent knowledge of the operations team, with teams knowing the locations of servers and the operations the servers are performing, started to alter with the early implementation of virtualization and will change entirely as the use of Linux containers increases. IT teams need to understand far more about the state of the infrastructure and the state of the applications running on the infrastructure, requiring more telemetry information than has traditionally been available. The R-Series supports several new sources of analytics information such as buffer monitoring and expanded counters to complement the diagnostic functions that exist in the current generation of switches such as 32 local Cisco Switched Port Analyzer (SPAN) and Encapsulated Remote SPAN (ERSPAN) sessions.

Network Capabilities

The changing requirements of the network are not limited to increased capacity and better operational characteristics. More sophisticated tunneling and forwarding capabilities, which had been limited to expensive backbone routers, are needed in today's high-speed data center networks.

VXLAN is an established industry-standard overlay technology for building scalable and reliable data center fabric networks that support application workload mobility and network agility. To move business-critical applications to the VXLAN overlay fabric, organizations must be able to help ensure high performance for both Layer 2 bridging and Layer 3 routing in the VXLAN overlay network. The R-Series forwarding ASICs have a built-in VXLAN routing capability to provide uncompromised forwarding bandwidth and latency performance for VXLAN routed traffic.

With the R-Series, switch features such as segment routing and full-featured VXLAN overlays are possible at greater scale:

- Single-pass VXLAN tunnel gateway
- VXLAN routing
- Segment routing
- Push 3 Multiprotocol Label Switching (MPLS) labels for segment routing
- vPC
- Buffer carving by traffic class

Note: The Cisco Nexus 9500 R-Series ASIC is designed and built with the features described here. However, the software support for some of the features may not be immediately available. For supported features in a given software release, refer to the corresponding software release notes.

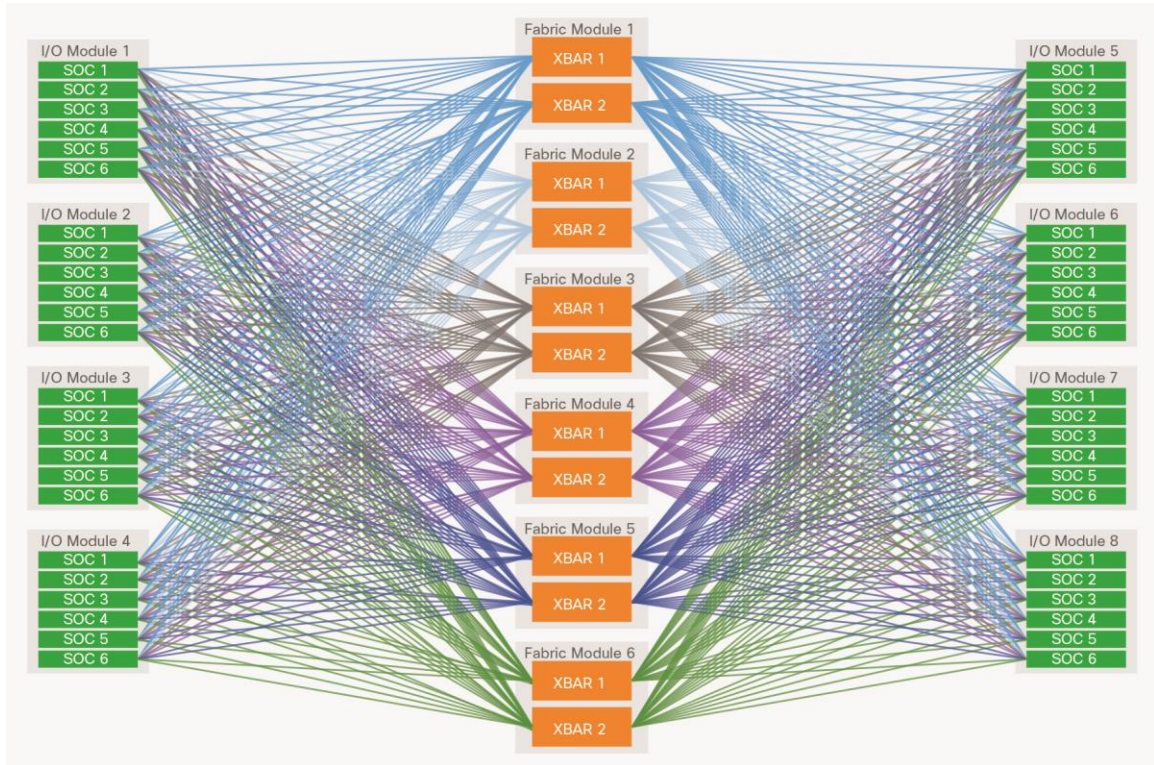
Fabric Module

The Cisco Nexus 9508 Switch uses up to six Cisco Nexus 9508 R-Series fabric modules to provide nonblocking internal bandwidth between line cards (Figure 3). The new fabric modules are built with the Broadcom (FE 3600) fabric engine ASIC.

The fabric module of the R-Series switches performs the following functions in the modular chassis architecture:

- The fabric modules provide high-speed nonblocking cell-based forwarding connectivity for the line cards. All links are active and used in active data paths. Cisco Nexus 9508 Switches can hold six Cisco Nexus 9508 R-Series fabric modules, with each additional fabric module adding 900 Gbps per line-card slot. With all fabric modules in place, the total bandwidth per line-card slot is 5.4 terabits per second (Tbps). With all the available bandwidth, the Cisco Nexus 9508 can provide N+1 redundancy. If a fabric module partially or completely fails, the front-panel port bandwidth is not affected.
- The fabric module performs multicast lookup and packet replication to send copies of multicast packets to the receiving egress ASICs on the line cards.

Figure 3. Cisco Nexus 9508 Switch Internal Connectivity with R-Series Line Cards and Fabric Modules



Cisco Nexus 9500 R-Series Line Cards

The R-Series 100-Gbps line card consists of six ASICs, and the 40-Gbps card consists of three ASICs. Both line cards provide a high density of 36 ports.

Each line card also has a built-in dual-core CPU to offload or accelerate control-plane tasks such as programming hardware table resources and collecting and sending line-card counters and statistics. This feature significantly improves system control-plane performance.

Cisco Nexus N9K-X9636C-R Line Card


The N9K-X9636C-R line card provides 36 100 Gigabit Ethernet QSFP28 ports. Figure 4 shows the front face of the line card.

Figure 4. Front Face of the Cisco Nexus N9K-X9636C-R Line Card



The N9K-X9636C-R line card is built with six Broadcom ASICs as its forwarding engines. Each ASIC provides six 100 Gigabit Ethernet ports. The forwarding capacity of the forwarding ASIC is 720 million packets per second (mpps). The ASIC provides full line rate on 24 of the 36 ports for frame sizes between 64 and 92 bytes, or full line rate on all 36 ports for frame sizes greater than 92 bytes with the folded Clos architecture of six fabric modules. From each line-card ASIC, thirty-six 25-Gbps lanes—for a total of 900 Gbps—interconnect each line card to a fabric module. With a total of six fabric modules, the maximum bandwidth available to a line card is 5.4 Tbps. This bandwidth can be used for data-plane traffic, multicast ingress replication, and SPAN traffic. Figure 5 illustrates the line-card internal architecture.

Figure 5. Cisco Nexus N9K-X9636C-R Line Card Architecture




This line card uses a physical-layer-free (PHY-less) design. This design reduces data transport latency on the port, decreases port power consumption, and improves reliability because it uses fewer active components.

In addition to 100 Gigabit Ethernet, the front-panel ports can operate at 40 Gigabit Ethernet speeds. When a port is populated with a QSFP28 transceiver, it can operate as a single 100 Gigabit Ethernet port. When a port has a QSFP+ transceiver plugged in, it can run as a single 40 Gigabit Ethernet port.

Cisco Nexus N9K-X9636Q-R Line Card


The N9K-X9636Q-R line card provides thirty-six 40 Gigabit Ethernet QSFP+ ports. Figure 6 shows the front face of the line card.

Figure 6. Front Face of the Cisco N9K-X9636Q-R Line Card



The N9K-X9636Q-R line card is built with three Broadcom ASICs as its forwarding engines. Each ASIC has twelve 40 Gigabit Ethernet ports and provides full line-rate performance for all packet sizes with the folded Clos architecture of six fabric modules. From each ASIC, eighteen 25-Gbps lanes—for a total of 450 Gbps—interconnect each line card to a fabric module. With a total of six fabric modules, the total bandwidth available to a line card and fabric module is 2.7 Tbps. This bandwidth can be used for data-plane or monitoring traffic. Figure 7 illustrates the line-card internal architecture.

Figure 7. Cisco Nexus N9K-X9636C-R Line Card Architecture



This line card features a PHY-less design. This design reduces data transport latency on the port, decreases port power consumption, and improves reliability because it uses fewer active components.

Front-panel ports can operate at 40 Gigabit Ethernet speeds. When a port has a QSFP+ transceiver plugged in, it can run as a single 40 Gigabit Ethernet port.

Packet Forwarding on the Cisco Nexus 9500 R-Series Portfolio

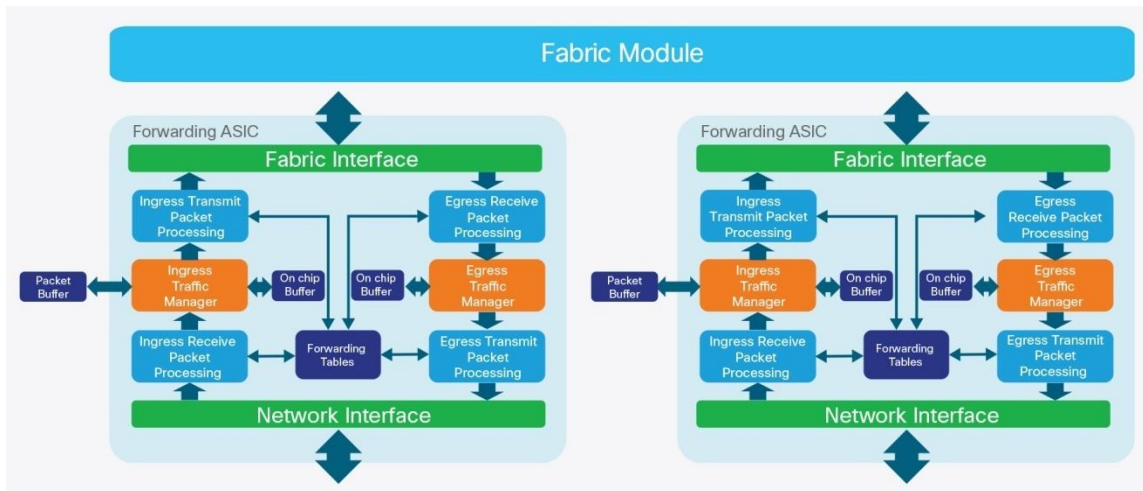
This section describes the packet forwarding process on the R-Series line cards and fabric modules.

Packet Forwarding Pipeline

The data-plane forwarding architecture of the Cisco Nexus 9500 R-Series switches includes the ingress pipeline on the ingress line card, fabric module forwarding, and the egress pipeline on the egress line card. The ingress and egress pipelines can be run on the same line card, or even on the same ASIC if the ingress and egress ports are on the same ASIC.

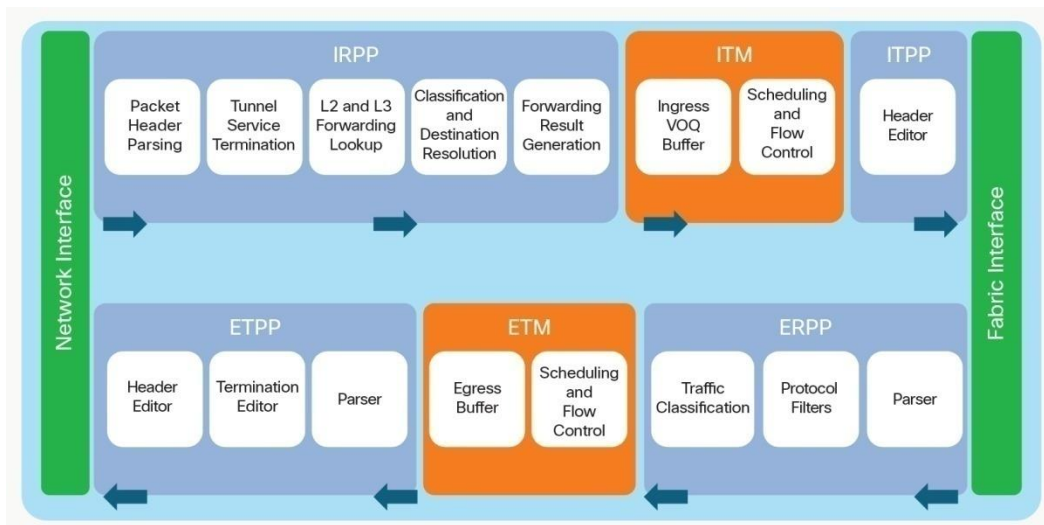
As shown in Figure 8, the forwarding pipeline consists of the network interface, ingress receive packet processor, ingress traffic manager, ingress transmit packet processor, fabric interface, egress receive packet processor, egress traffic manager, and egress transmit packet processor.

Figure 8. Forwarding Pipeline with Cisco Nexus 9500 R-Series Line Cards and Fabric Modules



Each pipeline element shown in Figure 8 can consist of multiple forwarding processing steps. Figure 9 illustrates the detailed steps in the ASIC pipelines when the switch is operating in NX-OS mode.

Figure 9. Detailed Forwarding Steps in the ASIC Pipelines



Ingress Receiver Packet Processor

The ingress receiver packet processor (IRPP) receives the packet from the network interface, parses the packet headers, and performs a series of lookups to decide whether to accept the packet and how to forward it to its intended destination. It also generates instructions to the data path for buffering and queuing of the packet.

Packet Header Parsing

When a packet arrives through a front-panel port, it goes through the ingress pipeline, and the first step is packet header parsing. The flexible packet parser parses the first 128 bytes of the packet to extract and save information such as the Layer 2 header, EtherType, Layer 3 header, and TCP/IP protocols.

These parsed fields are used in a series of forwarding table and ACL lookups to determine the:

- Destination output interfaces (based on Ethernet learning, IP host route entries, longest-prefix match, etc.)
- Compliance of switching and routing protocols (spanning tree, VXLAN, Open Shortest Path First [OSPF], Fabric Shortest Path First [FSPF], Intermediate System to Intermediate System [IS-IS], redirection, IP packet check, etc.)
- Policies (network access rights, permit order, security, etc.)
- Control-plane redirection and copying (Bridge Protocol Data Unit [BPDU], Address Resolution Protocol [ARP], Internet Group Management Protocol [IGMP], gleaning, etc.)
- System CoS classification (input queue, output queue, IEEE 802.1p tagging, etc.)
- Service rates and policers
- SPAN (ingress, egress, drop, etc.)
- Statistics (flow and interface packet and byte counters, etc.)
- Network flow-based load balancing (multipathing, port channel, etc.)
- Flow samplers (M of N bytes, M of N packets, etc.)
- Packet header rewrites (next-hop address, overlay encapsulation, time to live [TTL], etc.)
- Flow table to implement Cisco NetFlow

Tunnel Service Termination

After going through the parser, the packet goes through a series of lookups in which the goal is to forward traffic to a desirable destination. Sometimes, packets are encapsulated in additional headers to form an MPLS, generic routing encapsulation (GRE), or VXLAN header. To perform a lookup for traffic inside a tunnel, additional headers are terminated, and the lookup can be performed on the original header.

Layer 2 and Layer 3 Forwarding Lookup

As the packet goes through the ingress pipeline, it is subject to Layer 2 switching and Layer 3 routing lookups. First, the forwarding lookup examines the destination MAC (DMAC) address of the packet to determine whether the packet needs to be Layer 2 switched or Layer 3 routed. If the DMAC address is the switch's own router MAC address, the packet is passed to the Layer 3 routing lookup logic. If the DMAC address doesn't belong to the switch, a Layer 2 switching lookup based on the DMAC address and VLAN ID is performed. If a match is found in the MAC address table, the packet is sent to the egress port. If there is no match for DMAC address and VLAN combination, the packet is forwarded to all ports in the same VLAN.

In the Layer 3 lookup logic on the line card, the destination IP (DIP) address is used for searching in the Layer 3 host table. This table stores forwarding entries for directly attached hosts or learned /32 host routes. If the DIP address matches an entry in the host table, the entry indicates the destination port, next-hop MAC address, and egress VLAN. If there is no match to the DIP address in the host table, a Layer 3 lookup is performed in the LPM routing table.

When Layer 2 switching and Layer 3 host routing is performed, if the egress port is local to the ASIC, packets will be forwarded locally without going to fabric modules.

Ingress Traffic Classification and Destination Resolution

In addition to forwarding lookups, the packet undergoes ingress ACL processing. The ACL TCAM is checked for ingress ACL matches. Each ASIC has an ingress ACL TCAM table of 48,000 entries per ASIC to support system internal ACLs and user-defined ingress ACLs. These ACLs include port ACLs (PACLs), routed ACLs (RACLs), VLAN ACLs (VACLs), QoS ACLs, CoPP ACLs. This process allows maximum utilization of the ACL TCAM within R-Series line cards.

R-Series line cards support ingress traffic classification. On an ingress interface, traffic can be classified based on the address fields, IEEE 802.1q CoS, and IP precedence or DSCP in the packet header. The classified traffic can be assigned to one of the eight QoS groups. The QoS groups provide internal identification of the traffic classes that is used for subsequent QoS processes as packets go through the system, such as VOQ queueing.

After the Layer 2 and Layer 3 lookups, ACL processing, and traffic classification, which can influence the destination of the packet, resolution of the destination port is performed. The result of this operation identifies the destination port and queue values that will be used in the VOQ block to queue and schedule traffic. In this stage, reverse path forwarding (RPF) protection is performed for unicast and multicast traffic, and equal-cost multipath (ECMP) and port-channel load balancing.

The final step in the ingress forwarding pipeline is to collect all the forwarding metadata generated earlier in the pipeline and pass it to the downstream blocks through the data path.

A 72-byte internal header is stored along with the incoming packet in the packet buffer. The internal header provides information about the traffic class, drop precedence, source port, destination port, and whether the packet needs to be forwarded or whether mirroring and snooping actions need to be applied.

Ingress Traffic Manager

The ingress traffic manager (ITM) is the next block in the ingress pipeline. It performs steps related to queueing traffic in to VOQ, scheduling traffic for transmission over the fabric, and managing credits.

Ingress VOQ Buffer

The ingress VOQ buffer block manages both the on-chip buffer and the off-chip packet buffer. Both buffers use VOQ architecture, and traffic is queued based on the information from the IRPP. A total of 96,000 VOQs are available for unicast and multicast traffic.

Scheduling and Flow Control

Before a packet is transmitted from the ingress pipeline, the packet needs to be scheduled for transfer over the fabric. The ingress scheduler sends a credit request to the egress scheduler located in the egress traffic manager block. When the ingress traffic manager receives the credit, it starts sending traffic to the ingress transmit packet processor. If the egress buffer is full, traffic will be buffered in the dedicated queue represented by the egress port and traffic class.

Ingress Transmit Packet Processor

The ingress transmit packet processor (ITPP) performs the last step in the ingress pipeline. At this stage, a packet's internal header is updated, and packets are segmented into cells. Cells are of variable length and can be between 64 and 256 bytes. After the internal headers are transmitted in the cells, the cells are transferred over the fabric.

Egress Receive Packet Processor

The egress receive packet processor (ERPP) assembles the packets that are transferred over the fabric and processes them. To start, the first 128 bytes of the packet is parsed, checking for changes on the headers. In the next step in the ERPP, filtering is implemented, with traffic filtered based on VLAN, spanning tree, maximum transmission unit (MTU), and private VLAN. Before traffic goes to the egress traffic manager, it is classified. In this block, egress ACL classification is performed.

Egress Traffic Manager

The egress traffic manager (ETM) performs egress buffer accounting, packet queuing, and scheduling. All ports dynamically share the egress buffer resource. The output data path controller also performs packet shaping. The Cisco Nexus 9500 platform switches use a simple egress queuing architecture. An R-Series line card supports eight user-defined classes identified by QoS-group IDs. Each user-defined class has a unicast queue and a multicast queue for each egress port.

The ETM performs egress credit scheduling. In response to a credit request, the ETM block checks buffer availability on the destination port for the traffic class. If buffer space is available, credit will be granted to the ingress side for the packet from the VOQ. If buffer space is not available for the packet to be received, the egress buffer will not grant the credit.

The R-Series uses a distributed credit architecture in which the packet from the ingress VOQ requires communication with the egress scheduler to check and grant access to the fabric.

Egress Transmit Packet Processor

The egress transmit packet processor (ETPP) gets the input packet from the egress traffic manager and is responsible for all packet rewrites. It extracts the internal header and various packet header fields from the packet, performs a series of lookups, and generates the rewrite instructions. If packets belong to a tunnel, the ETPP can remove unwanted headers. In the header editor block, forwarding headers are updated with the TTL information or traffic class. If the packet needs to be encapsulated, the header editor will add a new header and encapsulated packet in the new tunnel header. Before the packet goes on the wire, the header editor builds or updates the Ethernet header, updates the VLAN if needed, and filters Spanning Tree Protocol if the new VLAN requires it. After all the changes have been made, the packet is transferred onto the wire.

Multicast Packet Forwarding

For multicast packets, the system needs to determine whether packets must undergo Layer 2 multicast forwarding or Layer 3 multicast forwarding. A packet will go through Layer 3 multicast forwarding if it meets the following criteria:

- It is an IP packet with a multicast address.
- It is not a link-local packet.
- Multicast routing is enabled on the bridge domain.

For Layer 2 IP multicast packets, forwarding lookups can be performed using either the Ethernet or IP headers. By default, R-Series line cards use the IP address for the Layer 2 multicast forwarding lookup, but the fan-out is limited to the same bridge domain.

For broadcast packets, the ASIC floods the packet in the bridge domain. The ASIC maintains separate per-bridge domain fan-out lists for broadcast and for unknown unicast traffic.

IP multicast forwarding in Application Spine Engine 2 (ASE2) and Leaf-and-Spine Engine (LSE) ASICs relies on the forwarding information base (FIB) table. For both Layer 2 and Layer 3 IP multicast, the switch looks up the IP address in the FIB table. The source address is used for RPF check, and the destination address is used to determine the outgoing interface list. If the FIB destination address search doesn't find a match, the packet is classified as unknown multicast. If IGMP and Multicast Listener Discovery (MLD) snooping is enabled, the packet is forwarded to all router ports on which the incoming bridge domain is present. If snooping is disabled, the packet is flooded in the bridge domain.

Multistage Replication

Multicast packets go through the same ingress and egress processing pipelines as unicast packets. However, one difference in the packet lookup and forwarding process is that the R-Series line cards and fabric modules perform three-stage distributed multicast lookup and replication. The multicast routing table is stored on all the line cards and fabric modules. The ingress ASIC on the line card performs the first lookup to resolve local receivers. If there are any local receivers, a copy is sent to the local ports. The ingress ASIC then sends a copy of the incoming packet to the fabric module. On receiving the cell, the fabric module performs the second lookup to find the egress line cards. The fabric module replicates the cell to each egress line card on which there are receivers. The egress line card performs the third lookup to resolve its local receivers and replicates the packet on those ports. This multistage multicast lookup and replication process provides the most efficient way to replicate and forward multicast traffic.

Conclusion

The Cisco Nexus 9500 R-Series portfolio continues to deliver high-density and high-performance 100 and 40 Gigabit Ethernet ports with buffer capacity that meets demanding application requirements. The R-Series portfolio provides additional capabilities to meet the new requirements in the data center, such as increased forwarding capacity, a high-performance VXLAN overlay solution, and segment routing, plus a set of traditional networking technologies. The R-Series complements the existing Cisco Nexus 9500 portfolio, which is a leading modular switch platform for building cloud-scale data center networks and supporting converged and hyperconverged infrastructure.

For More Information

Read more about the [Cisco Nexus 9500 architecture](#), or contact your local account representative.

Read more about the Cisco Nexus 9500-R Series in the [data sheet](#).



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)