



Cisco Catalyst 4500 Series Supervisor Engine 6-E with Multicast Architecture

Why Should I Care About Multicast on the Cisco Catalyst 4500 Series Supervisor 6-E?

The new Cisco® Catalyst® 4500 Series Supervisor Engine 6-E features CenterFlex technology, which supports IPv4 and IPv6 multicast in hardware. The high-performance Supervisor Engine 6-E offers 320-Gbps nonblocking switching fabric and enables a rich set of security and quality-of-service (QoS) features. Like earlier supervisor engine models, the entire forwarding intelligence resides in the supervisor engine, leaving the line card design simple and easy to upgrade with new features.

The Supervisor Engine 6-E is designed for the Cisco Catalyst 4500 E-Series Switch, a modular wiring closet and aggregation Ethernet switch capable of forwarding Layer 2 and 3 multicast traffic at wire speed. The Catalyst 4500-E Series is ideal for Layer 2 and 3 multicast applications in a large enterprise wiring closet, medium-sized distribution layer, or small core network. By supporting IPv6 multicast forwarding in hardware, the Supervisor Engine 6-E protects your investment in the Catalyst 4500-E Series and extends the product lifecycle.

What Problems Need to Be Solved?

Many organizations are attempting to minimize employee travel in order to cut costs and help reduce carbon emissions. One approach is to conduct business over the Internet using applications such as video conferencing and distance e-learning. And multicast provides the most network-efficient solution for these applications.

As the network expands to support more application use, the network elements within it also need to scale to support such expansion. Because IPv4 multicast has limited multicast addresses, IPv6 multicast, with its virtually unlimited supply of addresses, becomes the best option.

Cisco Catalyst 4500 Series Supervisor Engine 6-E Multicast Architecture

The next-generation Supervisor Engine 6-E supports IPv4 and IPv6 multicast in hardware. Its CenterFlex technology is enabled by Cisco developed application-specific integrated circuits (ASICs) specific to the Supervisor Engine 6-E that deliver industry-leading centralized performance and configuration flexibility. Primary components of the CenterFlex ASICs include the Integrated Packet Processor (IPP), which performs switching decision for all ports; the Very-Fast Forwarding Engine (VFE) that performs packet forwarding; TCAM4, which allows 1 billion searches per second; and XgStub, which increases the slot bandwidth from 4 Gbps to 24 Gbps. These new CenterFlex ASICs supercharge the throughput to 250 million packets per second (mpps) for IPv4 and 125 mpps for IPv6 packets.

The new Supervisor Engine 6-E processes the IPv4 and IPv6 multicast packet similarly to how it processes the unicast packet except replication is done at the Replica Management Module (RMM). The RMM manages VLAN/port sets and tunneling adjacencies for replication of packets that have to be sent to multiple destinations. The RMM consists of two data structures: Replica Request Queue (RRQ) and Replica Expansion Table (RET). The RRQ stores the packets to be replicated and uses the RET to determine the set of ports/tunneling adjacencies on which the packet should be replicated and forwarded.

The best way to understand how the Supervisor Engine 6-E supports multicast is to examine the Day in the life of a multicast packet.

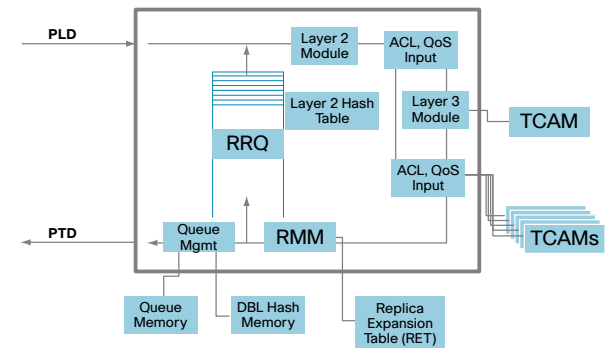
Day in the Life of a Multicast Packet in the Supervisor Engine 6-E

When a multicast packet arrives at the IPP, IPP parses the basic IP header, and any relevant extension headers for an IPv6 multicast packet, to ultimately parse the Layer 4 (TCP/UDP) data. It stores the packet in the shared packet

memory with 16 MB of shared bandwidth. The parsed flow label and a pointer to the start and end of the packet buffer are sent to the Very-Fast Forwarding Engine (VFE) in the Packet Lookup Descriptor (PLD).

Figure 1 depicts how the packet is processed after the PLD comes into the VFE.

Figure 1. VFE Forwarding for Multicast



As Figure 1 shows, when the PLD arrives at the VFE:

1. The packet goes through Layer 2 lookup. Spanning tree state is checked. Packet MAC source and MAC destination together with receive VLAN ID are looked up in the Layer 2 Hash Table. Layer 2 lookup also determines whether the packet is destined for router functionality.
2. Input classification is used to classify the packet via rules loaded into the Input Classification TCAM. This stage supports features such as input ACL, QoS, redirect, Policy-Based Routing and ACL-based copy. The results indicate whether the packet is denied or permitted, whether the packet should be forwarded to a specific port or through a specific adjacency entry, which input Policier to use, source or destination user group or both and so on.



Cisco Catalyst 4500 Series Supervisor Engine 6-E with Multicast Architecture

3. The packet is looked up in the Layer 3 lookup TCAM for IP bridging or routing lookup and multicast RPF check. Layer 3 lookup TCAM supports two parallel lookups per packet, one for multicast RPF and the other for destination address. Multicast RPF lookup data specifies the RPF interface(s) for the packet, actions for RPF failure, and source user group. Forwarding lookup data specifies an adjacency or Replica Expansion Table (RET) index for IP routing or bridging lookups or destination port or adjacency for Layer 2 MAC destination lookups.
4. The packet goes through output classification via rules loaded into the Output Classification TCAM. This stage supports features such as output ACL, QoS, redirect, and copy. The results indicate whether the packet is denied or permitted, whether the packet should be forwarded to a specific port, which output Policer to use, whether ToS and/or CoS of the packet is rewriteable at the end of VFE processing, what DBL actions to use, etc.
5. Output mapping determines final output QoS mappings for the packet after output policing and checks whether the packet is to be sniffed. Because the packet is to be replicated for multicast, a replica request descriptor is queued to the Replica Request Queue waiting for replication.
6. The RMM consults with the RET, the Replication Expansion Table, where the multicast forwarding information for high-speed access was stored. Upon a hit, the RET index and the RPF VLAN is sent to the VFE. The RET has the VLAN (as determined by Layer 3 lookup) and ports within each VLAN (as determined by IGMP/MLD Snooping) to which the packet needs to be copied.
7. After the necessary information is gathered, it will replicate the corresponding number of PLDs for the number of receivers. Each PLD will go through

the same pipeline stages as the original PLD except the input lookup stages.

8. VFE also performs the Dynamic Buffer Limiting (DBL) algorithm. DBL keeps per-flow buffer count as part of queue buffer management. DBL may drop the packet according to the congestion in the output queue.
9. A transmit descriptor is queued to the Queue Memory. The transmit queues are dynamically assigned to each line card and subsequently to each port based on chassis configuration. Each queue can have a maximum of 8184 packets independent of packet sizes.
10. When it is time to transmit the packet, the VFE sends the Packet Transmission Descriptor (PTD) to the IPP, which has the packet buffer pointer and all the information IPP needs to rewrite and forward the packet.
11. IPP reads the packet data from Packet Memory; in this case it needs to replicate the packet data corresponding to the number of PLDs that were replicated, perform necessary header rewrites, and transmit the packets out the corresponding ports. After all the packets are transmitted, IPP decrements the reference count of the packet buffer and frees the buffers.

The Supervisor Engine 6-E only replicates multicast packet headers while the body of the packet is stored in the Packet Memory, and every packet header goes through the same pipeline and each stage is processed in hardware. Thus the Supervisor Engine 6-E yields high yet consistent performance and low yet deterministic latency.

What Are the Benefits of Catalyst 4500 Series Supervisor Engine 6-E?

Through CenterFlex technology, the Cisco Catalyst 4500 Series Supervisor Engine 6-E not only supports scalable and high performing IPv4 multicast, it also supports IPv6 multicast in hardware with similar scalability and high performance. This means customers can integrate the Catalyst 4500 E-Series with the Supervisor Engine 6-E into their existing network for IPv4 multicast and take advantage of the easy migration path to IPv6 multicast in the future.