Understand Express Forwarding

Contents

Introduction

Prerequisites

- Requirements
- Components Used

Overview

CEF Core Components

CEF Operation Modes

CEF Basic Packet Flow

Related Information

Introduction

This document describes Cisco Express Forwarding (CEF).

Prerequisites

Requirements

There are no specific requirements for this document.

Components Used

This document is not restricted to specific software and hardware versions.

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

Overview

Cisco Express Forwarding (CEF) switching is a proprietary form of scalable switching intended to tackle the problems associated with demand caching. CEF is an advanced Layer-3 IP-switching architecture built

into Cisco IOS that replaces the old route-cache model. It maintains its own highly optimized forwarding table plus separate adjacency tables, enabling line-rate packet forwarding with minimal CPU involvement. Every time a packet arrives on a router or switch, the device must decide where to forward it. Once the packet's validity is confirmed and it has passed any access-control lists or other drop conditions, the device consults its routing table to determine the best outbound interface. Cisco routers and switches can perform this lookup in three ways: process switching, fast switching, or Cisco Express Forwarding.

With CEF, the device proactively computes all routes and installs them into two optimized data structures, the Forwarding Information Base (FIB) and the adjacency tables, rather than waiting for each packet to trigger a lookup. On modern platforms, these tables are synchronized with the hardware forwarding plane, this yields lower per-packet CPU overhead, higher packet-per-second throughput, and consistent forwarding behavior, which is why CEF is enabled by default on modern Cisco platforms and remains fundamental to high-performance IP networks.

CEF Core Components

- Forwarding Information Base (FIB) table - CEF uses FIB to make IP destination prefix-based decisions, in other words, FIB is the who-to-reach database for CEF. FIB is a one-to-one mirror of the router's IP Routing Information Base (RIB), but reorganized for fastest possible prefix lookup in hardware. Whenever the IP routing table changes (due to routing-protocol updates, static-route changes, and so on), those changes are immediately reflected in the FIB so that forwarding always uses up-to-date topology information. Each FIB entry consists of a destination prefix and one or more next-hop pointers; when ECMP (equal-cost multipath) is configured, the FIB holds multiple pointers per prefix to support per-packet or per-destination load-balancing. By containing every known route (rather than just a demand-cache of recently used routes), the FIB removes the need for route-cache maintenance and prevents cache thrash during topology churn, delivering stable, predictable, forwarding performance.

- Adjacency table - Nodes in the network are said to be adjacent if they can reach each other with a single hop across a link layer. In addition to the FIB, CEF uses adjacency tables to append Layer 2 addressing information (MAC address, VLAN tag, interface, and so on) needed to forward the packet on that link. The adjacency table maintains Layer 2 next-hop addresses for all FIB entries, these entries are populated as adjacencies are discovered (for example, through mechanisms like Address Resolution Protocol (ARP)).

CEF Operation Modes

- Central CEF mode - When central CEF mode is enabled, the FIB and adjacency tables reside on the route processor, and the route processor performs the express forwarding. You can use CEF mode when line cards are not available for CEF switching, when you need to use features not compatible with distributed CEF switching, or when you are running on a non-distributed platform.

- Distributed CEF (dCEF) mode - When dCEF mode is enabled, line cards maintain identical copies of the FIB and adjacency tables. The line cards can perform the express forwarding by themselves, and this relieves the main processor which enhances system performance. This mode is intended to increase scalability and is available in distributed or modular platforms that can spread processing tasks across two or more line cards. dCEF uses an Inter-Process Communication (IPC) mechanism to ensure synchronization of FIBs and adjacency tables on the route processor and line cards.

CEF Basic Packet Flow

The next list of steps provides an explanation on how a packet is CEF switched:

1. The device's interface processor reads the packet from the network media and stores it in buffer memory, called the interface's receive (RX) ring.
2. The device's interface processor sends a receive interrupt to the main processor. The rest of the switching occurs during this interrupt.
3. The device performs a longest-match lookup for the destination in the FIB, using the destination IP address as the search key. If the FIB lookup fails, the packet is dropped. If the lookup succeeds, a FIB path is selected and the device use its pointer to the corresponding adjacency entry.
4. The device rewrites the Layer 2 header using the encapsulation string from the adjacency table and places the packet into the correct output queue for transmission on the outbound interface.
5. The successfully switched packet is then enqueued on the outbound interface's transmit (TX) ring.

Related Information

- Technical Support - Cisco Systems
- Understand and Troubleshoot CEF on Cisco IOS XE Routers
- Cisco Express Forwarding