

CCNP BCMSN Notes - Multilayer Switching

Interfaces on multilayer switch are designated as switch ports (layer 2) with switchport or routed ports (layer 3) with no switchport. *Switched Virtual Interfaces (SVIs)* can be defined to provide a routed interface to a VLAN. Cisco Express Forwarding (CEF) Traditional multilayer switching ("route once, switch many", also known as NetFlow switching or route cache switching) was done through the combination of a route processor and a switching engine. CEF is the second generation of multilayer switching, and is enabled by default in hardware which supports it. CEF operation relies on two components working in tandem: the layer 3 engine (routing) and the layer 3 forwarding engine (switching). The layer 3 forwarding engine contains the Forwarding Information Base (FIB) and its Adjacency Table. Forwarding Information Base (FIB) The FIB is an optimized copy of the routing table, with more-specific routes listed first. Each entry in the FIB has layer 2 and 3 next-hop addressing information associated with it. FIB entries can be examined with `show ip cef`. Packets meeting certain conditions cannot be CEF-switched and will be punted to the layer 3 engine for traditional software routing:

- Expired TTL
- MTU exceeded
- ICMP redirect required
- Unsupported encapsulation type
- Compression and/or encryption is necessary
- An ACL log entry must be generated

Accelerated CEF (aCEF) can be implemented in some hardware to cache portions of the FIB on each line card. Distributed CEF (dCEF) stores the entire FIB on all capable line cards. **Adjacency Table** The adjacency table is the portion of the FIB which contains layer 2 next-hop information (MAC addresses which correspond to the layer 3 next-hop addresses). Similar to how the FIB is built from the routing table, the adjacency table is built from the ARP table. Adjacency information can be examined with `show adjacency`. Adjacency table entries with missing or expired layer 2 addresses are placed in the CEF glean state; packets must be punted to the L3 engine so an ARP request/reply can be generated. When a route is placed in the glean state, incoming packets will be dropped for up to two seconds as the switch awaits an ARP reply. Other adjacency states include:

- Null** - Represents a null interface (black hole)
- Drop** - Indicates packets cannot be forwarded to the destination and should be dropped
- Discard** - An ACL or other policy mandates that packets be dropped
- Punt** - Further processing is required by the layer 3 engine

Packet Rewrite The packet rewrite engine reconstructs the incoming packet with the appropriate next hop address information. Fields rewritten include:

- Layer 2 destination
- Layer 2 source
- IP TTL
- IP Checksum
- Layer 2 frame checksum

Fallback Bridging Non-IP protocols are not supported by CEF. Each SVI carrying nonroutable traffic can be assigned to a bridge group and bridged transparently, separate from normal L2 switching. A special type of STP known as VLAN-bridge is run on these bridge groups. Fallback bridging must be manually configured:

```
Switch(config)# bridge-group <group> protocol vlan-bridge
Switch(config)# interface vlan 10
Switch(config-if)# bridge-group <group>
Switch(config-if)# interface vlan 20
Switch(config-if)# bridge-group <group>
```

Verifying Multilayer Switching

- `show interface switchport` ("Disabled" verifies layer 3 operation)
- `show ip cef [detail]`
- `show bridge group`