Transparent Layer 2 Protocol Tunneling

This chapter introduces you to Transparent Layer 2 Protocol Tunneling to help initiate control packets from a local customer edge (CE) device to a remote CE device.

Transparent Layer 2 Protocol Tunneling

Table 1: Feature History Table

Feature Name	Release Information	Feature Description
Transparent Layer 2 Protocol Tunneling	Release 24.4.1	Introduced in this release on: Fixed Systems (8200, 8700); Modular Systems (8800 [LC ASIC: P100]) (select variants only*) * The Layer 2 protocol tunneling functionality is now extended to: • 8212-48FH-M • 8711-32FH-M • 88-LC1-52Y8H-EM • 88-LC1-12TH24FH-E • 88-LC1-36EH • 8712-MOD-M

Feature Name	Release Information	Feature Description
Transparent Layer 2 Protocol Tunneling	Release 7.3.2	This feature allows Layer 2 protocol data units (PDUs) to be kept intact and delivered across the service-provider network to the other side of the customer network. Such delivery is transparent because the VLAN and Layer 2 protocol configurations are maintained throughout. With this feature, service providers can send traffic from multiple customers across a core network without impacting the traffic of other customers. This feature is enabled by default.

This feature allows Layer 2 protocol data units (PDUs) to be tunneled across the core network without being interpreted and processed by intermediary network devices. Any packet on the L2 network is forwarded without any change. This feature is enabled by default.

You must configure the supported protocols only main and bundle interface. If you want a specific protocol packets to be punted over bundle members or subinterfaces, that protocol has to be enabled on the main interface as well.

When you want to use CFM and PVRST protocols, you must enable these protocols on a subinterface.

You can use this feature with the following protocols:

- Link Layer Discovery Protocol (LLDP)

- Cisco Discovery Protocol (CDP)

- Multiple Spanning Tree Protocol (MSTP)

- Per-VLAN Rapid Spanning Tree (PVRST)

- Connectivity Fault Management (CFM)

- Link Aggregation Control Protocol (LACP)

- Operation, Administration, Management (OAM)

- Synchronized Ethernet (SyncE)

- MAC security

- Priority Flow Control (PAUSE)

All packets on PW VPLS or VPWS are always tunnelled and no packet is sent to the CPU for processing (punted).

The following table depicts the behavior of the router when you enable a specific protocol on an interface.

L2 Protocol	Untagged Packet	Tagged Packet
Cisco Protocols	If Cisco protocols are enabled on the physical port, the traffic is sent to the CPU for processing. If Cisco protocols are disabled, the traffic is tunneled.	Traffic is always tunneled.
LLDP	If this protocol is enabled on the physical port, the traffic is sent to the CPU for processing. If this protocol is disabled, the traffic is tunneled.	Traffic is always tunneled.
PVRST/PVRST+	If this protocol is enabled on the main port, the traffic is sent to the CPU for processing. If this protocol is disabled, the traffic is tunneled.	If this protocol is enabled on the subinterface, the traffic is sent to CPU for processing. If it is disabled, the traffic is tunneled.
MSTP	If this protocol is enabled on the physical port, the traffic is sent to the CPU for processing. If this protocol is disabled, the traffic is tunneled.	Traffic is always tunneled.
CFM	If this protocol is enabled on the physical port, the traffic is sent to the CPU for processing. If this protocol is disabled, the traffic is tunneled.	If this protocol is enabled on the Xconnect, the traffic is sent to CPU for processing. If it is disabled, the traffic is tunneled.
LACP/SyncE/LOAM	If this protocol is enabled on the physical port, the traffic is sent to the CPU for processing. If this protocol is disabled, the traffic is tunneled.	Traffic is always tunneled.
PFC	If this protocol is enabled on the physical port, the traffic is sent to the CPU for processing. If this protocol is disabled, the traffic is tunneled.	Traffic is always tunneled.

Configuration

You cannot disable transparent tunneling, this feature is enabled by default.

To display the protocols that are enabled per interface, use the **show ofa objects ethport base location 0/1/CPU0** command.

```
Router# show ofa objects ethport base location 0/1/CPU0
ethport element 0 (hdl:0x308f38e360):
  base
  |-- dpd_slf -   pending(cr/up/dl):0/0/0, sibling:0x3093b811c8, child:2, num_parents:3,
parent-trans_id:1523, visits:0
   color_mask:0, last_bwalk_id:0 num_bwalks_started:0
  |-- keylen - 4
  |-- trans_id - 489153
  |-- create_trans_id - 1523
  |-- obj_handle - 0x308f38e360
  |-- flag - 10
  |-- reason - 0
  |-- table_operation - 6
  |-- total_obj_size - 632
  |-- idempotent - 0
  |-- inflight - 0
  |-- table_prop - jid=169 mtime=(GMT)2021.Jan.09 13:05:46.670570
  |--  (cont'd)  - replayed=0times
  `-+ npu_results
    |-- npu0 - 0:Success
    |-- npu1 - 0:Success
    |-- npu2 - 0:Success
    `-- npu3 - 0:Success
  ofa_npu_mask_t npu_mask =>
…
  ofa_bool_t remote_chain_in_use => TRUE
  ofa_bool_t local_chain_in_use => TRUE
  uint8_t copc_profile => 0
  ofa_bool_t lldp_enable => FALSE
  ofa_bool_t slow_proto_enable => FALSE
  ofa_bool_t cdp_enable => (not set)
  ofa_bool_t pvrst_enable => FALSE
  ofa_bool_t mstp_enable => FALSE
  ofa_bool_t macsec_enable => FALSE
  ofa_bool_t mka_enable => FALSE
  ofa_bool_t pfc_enable => FALSE
  ofa_bool_t cfm_enable => FALSE
  dpa_npu_mask_t npu_bmap => (not set)


Router# show ofa objects l2if base location 0/1/CPU0
l2if element 0 (hdl:0x3094ba70a8):
  base
  |-- dpd_slf -   pending(cr/up/dl):0/0/0, sibling:0x308f8087c8, child:1, num_parents:1,
visits:0
   color_mask:0, last_bwalk_id:0 num_bwalks_started:0
  |-- flag - 10
     |-- flag.is_id_alloced - 0x1
  |-- keylen - 4
  |-- trans_id - 18311
  |-- create_trans_id - 18299
  |-- obj_handle - 0x3094ba70a8
  |-- obj_rc - 0x0
  |-- reason - 0
  |-- table_operation - 6
  |-- total_obj_size - 776
  |-- idempotent - 1
  |-- inflight - 0
  |-- table_prop - jid=137 mtime=(GMT)2021.Jun.21 14:53:56.644917
  |--  (cont'd)  - replayed=0times
```

```
  `-- obj_rc - 0:Success
 ofa_npu_mask_t npu_mask => 0 (not set)
 uint32_t member_count => 1
@dpa_intf_t intf => 0x0f00000a
…
ofa_l2vpn_fwd_state_type fwd_state => (not set)
  ofa_bool_t cfm_enable => FALSE
  ofa_bool_t pvrst_enable => TRUE
  dpa_npu_mask_t npu_bmap => 1
```

To verify whether the L2 packet is flooded or forwarded by NPU, look at the interface counters. In case of flood, like multicast MAC, you will notice an increment in the output counters of the interface. When the traffic is forwarded with unicast MAC, you will notice an increment in the output counters only on the egress interface.

The following output displays the interface counters:

```
Router# show interface hundredGigE 0/0/2/0 accounting

HundredGigE0/0/2/0
  Protocol            Pkts In          Chars In      Pkts Out       Chars Out
  CDP                       0                0        163608        21923472
```