



LPTS Overview

Local Packet Transport Services (LPTS) maintains tables describing all packet flows destined for the secure domain router (SDR), making sure that packets are delivered to their intended destinations.

LPTS uses two components to accomplish this task: the port arbitrator and flow managers. The port arbitrator and flow managers are processes that maintain the tables that describe packet flows for a logical router, known as the Internal Forwarding Information Base (IFIB). The IFIB is used to route received packets to the correct Route Processor for processing.

LPTS interfaces internally with all applications that receive packets from outside the router. LPTS functions without any need for customer configuration. However, the policer values can be customized if required. The LPTS show commands are provided that allow customers to monitor the activity and performance of LPTS flow managers and the port arbitrator.

- [LPTS Policers, on page 1](#)
- [Per Port Rate Limiting of Multicast and Broadcast Punt Packets, on page 5](#)
- [LPTS Domain Based Policers, on page 7](#)

LPTS Policers

In Cisco IOS XR, the control packets, which are destined to the Route Processor (RP), are policed using a set of ingress policers in the incoming ports. These policers are programmed statically during bootup by LPTS components. The policers are applied based on the flow type of the incoming control traffic. The flow type is determined by looking at the packet headers. The policer rates for these static ingress policers are defined in a configuration file, which are programmed on the route processor during bootup. You can change the policer values based on the flow types of these set of ingress policers. You are able to configure the rate per policer per node.



Note

- You can get the default policer values and the current rates of the flow types from the output of the following show command:

```
show lpts pifib hardware police
```

Configuration Example

Configure the LPTS policer for the OSPF and BGP flow types with the following values globally for all nodes:

- ospf unicast default rate 3000
- bgp default rate 4000

```
Router#configure
Router(config)#lpts pifib hardware police
Router(config-pifib-policer-global)#flow ospf unicast default rate 3000
Router(config-pifib-policer-global)#flow bgp default rate 4000
Router (config-pifib-policer-global)#commit
```

Running Configuration

```
lpts pifib hardware police
flow ospf unicast default rate 3000
flow bgp default rate 4000
!
```

Verification

```
Router#show run lpts pifib hardware police
lpts pifib hardware police
flow ospf unicast default rate 3000
flow bgp default rate 4000
```

Configuration Example

Configure the LPTS policer for the OSPF and BGP flow types with the following values on an individual node - 0/RP0/CPU0:

- ospf unicast default rate 3000
- flow bgp default rate 4000

```
Router#configure
Router(config)#lpts pifib hardware police location 0/RP0/CPU0
Router(config-pifib-policer-per-node)#flow ospf unicast default rate 3000
Router(config-pifib-policer-per-node)#flow bgp default rate 4000
Router(config-pifib-policer-per-node)#commit
```

Running Configuration

```
lpts pifib hardware police location 0/RP0/CPU0
flow ospf unicast default rate 3000
flow bgp default rate 4000
```

Verification

The **show lpts pifib hardware police location 0/RP0/CPU0** command displays pre-Internal Forwarding Information Base (IFIB) information for the designated node.

```
Router#show lpts pifib hardware police location 0/RP0/CPU0
```

```
-----
Node 0/RP0/CPU0:
-----
Burst = 100ms for all flow types
-----
```

FlowType	Policer	Type	Cur. Rate	Burst	npu
OSPF-uc-default	32106	np	3000	1000	0
BGP-default	32118	np	4000	1250	0

Verification

The **show controllers npu stats traps-all instance all location 0/RP0/CPU0** command displays packets that are locally processed and packets that are dropped by the CPU.

Router# **show controllers npu stats traps-all instance all location 0/RP0/CPU0**

Trap Type	NPU	Trap	TrapStats	Policer	Packet	Packet
	ID	ID	ID		Accepted	Dropped
RxTrapMimSaMove (CFM_DOWM_MEP_DMM)	0	6	0x6	32037	0	0
RxTrapMimSaUnknown (RCY_CFM_DOWN_MEP_DMM)	0	7	0x7	32037	0	0
RxTrapAuthSaLookupFail (IPMC default)	0	8	0x8	32033	0	0
RxTrapSaMulticast	0	11	0xb	32018	0	0
RxTrapArpMyIp	0	13	0xd	32001	0	0
RxTrapArp	0	14	0xe	32001	11	0
RxTrapDhcpv4Server	0	18	0x12	32022	0	0
RxTrapDhcpv4Client	0	19	0x13	32022	0	0
RxTrapDhcpv6Server	0	20	0x14	32022	0	0
RxTrapDhcpv6Client	0	21	0x15	32022	0	0
RxTrapL2Cache_LACP	0	23	0x17	32003	0	0
RxTrapL2Cache_LLDP1	0	24	0x18	32004	0	0
RxTrapL2Cache_LLDP2	0	25	0x19	32004	1205548	0
RxTrapL2Cache_LLDP3	0	26	0x1a	32004	0	0
RxTrapL2Cache_ELMI	0	27	0x1b	32005	0	0
RxTrapL2Cache_BPDU	0	28	0x1c	32027	0	0
RxTrapL2Cache_BUNDLE_BPDU	0	29	0x1d	32027	0	0
RxTrapL2Cache_CDP	0	30	0x1e	32002	0	0
RxTrapHeaderSizeErr	0	32	0x20	32018	0	0
RxTrapIpCompMcInvalidIp	0	35	0x23	32018	0	0
RxTrapMyMacAndIpDisabled	0	36	0x24	32018	0	0
RxTrapMyMacAndMplsDisable	0	37	0x25	32018	0	0
RxTrapArpReply	0	38	0x26	32001	2693	0
RxTrapFibDrop	0	41	0x29	32018	0	0
RxTrapMTU	0	42	0x2a	32020	0	0
RxTrapMiscDrop	0	43	0x2b	32018	0	0
RxTrapL2AclDeny	0	44	0x2c	32034	0	0

Rx_UNKNOWN_PACKET	0	46	0x2e	32018	0	0
RxTrapL3AclDeny	0	47	0x2f	32034	0	0
RxTrapOamY1731MplsTp (OAM_SWOFF_DN_CCM)	0	57	0x39	32029	0	0
RxTrapOamY1731Pwe (OAM_SWOFF_DN_CCM)	0	58	0x3a	32030	0	0
RxTrapOamLevel	0	64	0x40	32023	0	0
RxTrapRedirectToCpuOamPacket	0	65	0x41	32025	0	0
RxTrapOamPassive	0	66	0x42	32024	0	0
RxTrap1588	0	67	0x43	32038	0	0
RxTrapExternalLookupError	0	72	0x48	32018	0	0
RxTrapArplookupFail	0	73	0x49	32001	0	0
RxTrapUcLooseRpfFail	0	84	0x54	32035	0	0
RxTrapMplsControlWordTrap	0	88	0x58	32015	0	0
RxTrapMplsControlWordDrop	0	89	0x59	32015	0	0
RxTrapMplsUnknownLabel	0	90	0x5a	32018	0	0
RxTrapIpv4VersionError	0	98	0x62	32018	0	0
RxTrapIpv4ChecksumError	0	99	0x63	32018	0	0
RxTrapIpv4HeaderLengthError	0	100	0x64	32018	0	0
RxTrapIpv4TotalLengthError	0	101	0x65	32018	0	0
RxTrapIpv4Ttl0	0	102	0x66	32008	0	0
RxTrapIpv4Ttl1	0	104	0x68	32008	0	0
RxTrapIpv4DipZero	0	106	0x6a	32018	0	0
RxTrapIpv4SipIsMc	0	107	0x6b	32018	0	0
RxTrapIpv6VersionError	0	109	0x6d	32018	0	0
RxTrapIpv6HopCount0	0	110	0x6e	32011	0	0
RxTrapIpv6LoopbackAddress	0	113	0x71	32018	0	0
RxTrapIpv6MulticastSource	0	114	0x72	32018	0	0
RxTrapIpv6NextHeaderNull	0	115	0x73	32010	0	0
RxTrapIpv6Ipv4CompatibleDestination	0	121	0x79	32018	0	0
RxTrapMplsTtl1	0	125	0x7d	32012	316278	2249
RxTrapUcStrictRpfFail	0	137	0x89	32035	0	0
RxTrapMcExplicitRpfFail	0	138	0x8a	32033	0	0
RxTrapOamp (OAM_BDL_DN_NON_CCM)	0	141	0x8d	32031	0	0

RxTrapOamEthUpAccelerated (OAM_BDL_UP_NON_CCM)	0	145	0x91	32032	0	0
RxTrapReceive	0	150	0x96	32017	125266112	0
RxTrapUserDefine_FIB_IPV4_NULL0	0	151	0x97	32018	0	0
RxTrapUserDefine_FIB_IPV6_NULL0	0	152	0x98	32018	0	0
RxTrapUserDefine_FIB_IPV4_GLEAN	0	153	0x99	32016	0	0
RxTrapUserDefine_FIB_IPV6_GLEAN	0	154	0x9a	32016	0	0
RxTrapUserDefine_IPV4_OPTIONS	0	155	0x9b	32006	0	0
RxTrapUserDefine_IPV4_RSVP_OPTIONS	0	156	0x9c	32007	0	0
RxTrapUserDefine	0	157	0x9d	32026	0	0
RxTrapUserDefine_BFD	0	163	0xa3	32028	0	0
RxTrapMC	0	181	0xb5	32033	0	0
RxNetflowSnoopTrap0	0	182	0xb6	32018	0	0
RxNetflowSnoopTrap1	0	183	0xb7	32018	0	0
RxTrapMimSaMove (CFM_DOWM_MEP_DMM)	1	6	0x6	32037	0	0
RxTrapMimSaUnknown (RCY_CFM_DOWN_MEP_DMM)	1	7	0x7	32037	0	0
RxTrapAuthSaLookupFail (IPMC default)	1	8	0x8	32033	0	0
RxTrapSaMulticast	1	11	0xb	32018	0	0
RxTrapArpMyIp	1	13	0xd	32001	0	0

Associated Commands

- lpts pifib hardware police
- flow ospf
- flow bgp
- show lpts pifib hardware police

Per Port Rate Limiting of Multicast and Broadcast Punt Packets

This feature enables rate limiting of multicast and broadcast punted traffic at the interface level. Currently, a rate limit is supported per NPU level. This feature supports rate limiting at the interface level so as to protect a port from receiving the multicast and broadcast storm of punted traffic. Rate limiting for all the L3 protocol punt packets and L2 protocol packets (only ERPS, and DOT1x) is supported on physical and bundle main interfaces.

Configuring a Rate Limit to the Multicast and Broadcast Punted Traffic

You can configure the multicast and broadcast rate limit in three levels:

- Interface level
- Global level
- Domain level

Along with rate limiting the multicast and broadcast punted traffic, you can configure rate limit to these protocol punted traffic:

- ARP
- CDP
- LACP

The protocol specific configurations are explained in the below section.

Limitation

When broadcast and multicast rate limit is configured along with ARP rate limit, the ARP packets increment broadcast and multicast counters.

Interface Level

This example shows how to configure the rate limit of 1000 pps for the multicast and broadcast punted traffic at the TenGig interface:



Note

A interface level rate limit configuration has the highest priority over a global and domain level configurations.

1. Router# configure
Enters the configuration mode.
2. Router(config)# lpts punt police 0/RP0/CPU0
Enters punt configuration mode.
3. Router(config-lpts-punt-policer)# interface TenGigE0/0/0/3
Enters per interface level policer configuration.
4. Router(config-lpts-punt-policer-global-if)# mcast rate 1000
Configures a rate limit of 1000 pps for multicast punted traffic.
5. Router(config-lpts-punt-policer-global-if)# bcast rate 1000
Configures a rate limit of 1000 pps for broadcast punted traffic.
6. Router(config-lpts-punt-policer-global-if)# commit
Commit the configuration.

Global Level

This example shows how to configure the rate limit of:

- 1000 pps for the multicast and broadcast punted traffic

1. Router# configure
Enters the configuration mode.
2. Router(config)# lpts punt police 0/RP/CPU0
Enters punt configuration mode.
3. Router(config-punt-policer-global)# mcast rate 1000
Configures multicast rate limit of 1000 pps.
4. Router(config-punt-policer-global)# bcast rate 1000
Configures broadcast rate limit of 1000 pps.
5. Router(config-punt-policer-global)# commit
Commit the configuration.

LPTS Domain Based Policers

You can configure a particular port, a group of ports, or a line card of a router with LPTS policers of a single domain. Configuration of port-based policers that belong to a particular domain enables better categorisation and control of different types of ingress traffic. For example, since iBGP traffic has a higher rate of traffic flow, the ports that handle iBGP traffic can be configured with higher policer rates compared to the ports that handle eBGP traffic.

Restrictions

- The policer rates that are configured for ports or line cards are carried forwards as policer rates of the domain after configuring the ports or line cards as part of a domain. For example, if port hundredGigE 0/0/0/1 and port hundredGigE 0/0/0/2 have policer rate of 3000 for ospf unicast known flow and if the ports are configured as part of domain CORE, then the policer rate of domain CORE for ospf unicast known flow is 3000 unless it is configured otherwise.
- You can configure only one domain per router.
- A Domain name can be any word but can have up to a maximum of 32 characters.

Configuration Example

To configure LPTS domain based policers, use the following steps:

1. Enter the LPTS hardware configuration mode and create a domain.
2. Configure the interfaces for the domain.
3. Enter the LPTS hardware configuration mode for the domain CORE, and then configure the ingress policer rates for the domain CORE at the global level.

4. Enter the LPTS hardware configuration mode for the domain CORE, and then configure the ingress policer rates for the domain CORE at the line card level.

Configuration

```

/* Enter the LPTS hardware ingress policer configuration mode and create a domain named
CORE. */
Router# config
Router(config)# lpts pifib hardware domain CORE

/* Configure the interfaces for the domain CORE. */
Router(config-lpts-domains-CORE)# interface hundredGigE 0/0/0/1
Router(config-lpts-domains-CORE)# interface hundredGigE 0/0/0/2
Router(config-lpts-domains-CORE)# commit
Router(config-lpts-domains-CORE)# exit

/* Enter the LPTS hardware configuration mode for the domain CORE, and then configure the
ingress policer rates for the domain CORE at the global level. */
Router(config)# lpts pifib hardware police domain CORE
Router(config-lpts-policer-global-CORE)# flow ospf unicast known rate 6000
Router(config-lpts-policer-global-CORE)# flow ospf unicast default rate 7000
Router(config-lpts-policer-global-CORE)# commit
Router(config-lpts-policer-global-CORE)# exit
Router(config-lpts-policer-global)# exit

/* Enter the LPTS hardware configuration mode for the domain CORE, and then configure the
ingress policer rates for the domain CORE at the line card level. */
Router(config)# lpts pifib hardware police location 0/0/CPU0 domain CORE
Router(config-lpts-policer-global-CORE)# flow ospf unicast known rate 7000
Router(config-lpts-policer-global-CORE)# flow ospf unicast default rate 8000
Router(config-lpts-policer-global-CORE)# commit

```

Running Configuration

```

lpts pifib hardware domain CORE
  interface HundredGigE0/0/0/1
  interface HundredGigE0/0/0/2
!
lpts pifib hardware police
  domain CORE
  flow ospf unicast known rate 6000
  flow ospf unicast default rate 7000
!

lpts pifib hardware police location 0/0/CPU0 domain CORE
  flow ospf unicast known rate 7000
  flow ospf unicast default rate 8000
!

```

Verification

Use the following command to verify information about the LPTS domains configured:

```

Router# show lpts pifib domains
Thu Nov 21 15:49:31.334 IST

Domains Information: 1 Configured
-----
Domain: [1] CORE
-----
interface [-----] HundredGigE0/0/0/1

```



```
interface [-----] HundredGigE0/0/0/2
          0 local of total 2 interfaces
```

