



Implementing LPTS

- [LPTS Overview, on page 1](#)
- [LPTS Policers, on page 1](#)
- [LPTS and NPU Traps, on page 3](#)
- [Defining Dynamic LPTS Flow Type , on page 5](#)

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

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 effective 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 flowtypes 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-lpts-policer-global)#flow ospf unicast default rate 3000
Router(config-lpts-policer-global)#flow bgp default rate 4000
Router(config-lpts-policer-global)#commit
```

Running Configuration

```
Router#show running-config lpts
lpts pifib hardware police
  flow ospf unicast default rate 3000
  flow bgp default rate 4000
!
```

Verification

```
Router#show lpts pifib hardware police
```

```
-----
Node 0/RP0/CPU0:
-----
```

FlowType	Policer	Type	Cur. Rate	Burst	Accepted	Dropped	npu
Fragment	2	np	542	1000	0	0	0
OSPF-mc-known	3	np	1627	1000	0	0	0
OSPF-mc-default	4	np	1084	1000	0	0	0
OSPF-uc-known	5	np	542	1000	0	0	0
OSPF-uc-default	6	np	2878	1000	0	0	0
BFD-default	10	np	8136	1000	0	0	0
BFD-MP-known	11	np	8136	1000	0	0	0
BGP-known	16	np	17000	1000	0	0	0
BGP-cfg-peer	17	np	1627	1000	0	0	0
BGP-default	18	np	3880	1000	0	0	0

Configuration Example

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

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

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

Running Configuration

```
Router#show running-config lpts
lpts pifib hardware police location 0/0/CPU0
  flow ospf unicast default rate 3000
  flow bgp default rate 4000
!
```

Verification

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

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

```
-----
Node 0/0/CPU0:
-----
```

FlowType	Policer	Type	Cur. Rate	Burst	Accepted	Dropped	npu
Fragment	2	np	542	1000	0	0	0
Fragment	2	np	542	1000	0	0	1
OSPF-mc-known	3	np	1627	1000	0	0	0
OSPF-mc-known	3	np	1627	1000	0	0	1
OSPF-mc-default	4	np	1084	1000	0	0	0
OSPF-mc-default	4	np	1084	1000	0	0	1
OSPF-uc-known	5	np	542	1000	0	0	0
OSPF-uc-known	5	np	542	1000	0	0	1
OSPF-uc-default	6	np	2878	1000	0	0	0
OSPF-uc-default	6	np	2878	1000	0	0	1
BFD-default	10	np	8136	1000	0	0	0
BFD-default	10	np	8136	1000	0	0	1
BFD-MP-known	11	np	8136	1000	0	0	0
BFD-MP-known	11	np	8136	1000	0	0	1
BGP-known	16	np	17000	1000	0	0	0
BGP-known	16	np	17000	1000	0	0	1
BGP-cfg-peer	17	np	1627	1000	0	0	0
BGP-cfg-peer	17	np	1627	1000	0	0	1
BGP-default	18	np	3880	1000	0	0	0
BGP-default	18	np	3880	1000	0	0	1

Associated Commands

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

LPTS and NPU Traps

Network Processing Unit (NPU) traps are raised by the routers for inspection. NPU traps are raised in response to the type of packets received by the router and can indicate either exception packets, error packets, or non-LPTS control packets.

- Examples of exception packets include glean adjacency traffic or packets with IPv4 options.
- Examples of error packets include IPv4 packet with bad checksum or IPv6 packets with a hop count of zero.

- Examples of non-LPTS control packets include those packets that do not get processed through LPTS (for example, LACP, LLDP and other L2 control packets).

Each of the NPU traps are policed at a rate that is pre-programmed by the router's system design. Packets are policed per NPU and excess traffic is dropped by the NPU with respect to the system design. Some NPU trap packets that are allowed by NPU policers are sent to the CPU if they need additional processing. Others that exceed the NPU policer rate are dropped by the NPU.

Verification

Use the command `show controllers npu stats traps-all instance NPU-Number|all location RP|LC` command to check the NPU trap statistics for all the NPUs or per NPU of a router.

For fixed systems, the NPU trap statistics is available for the location `0/RP0/CPU0` and is provided through the command `show controllers npu stats traps-all instance all location 0/RP0/CPU0`. For distributed systems, NPU trap statistics is available for the line card locations and is provided through the command `show controllers npu stats traps-all instance all location 0/1/CPU0`. You can use the command `clear controller npu stats traps-all instance NPU-Number|all location RP|LC`

In the following example:

- **(D)** indicates the trap packets that are dropped in the NPU.
- **(D*)** indicates the trap packets that are dropped in NPU but are available for analysis.
- The **Accepted** count in the output indicates the ones that are available for analysis.

`RP/0/RP0/CPU0:router#show controllers npu stats traps-all instance all location 0/RP0/CPU0`

Trap Type	NPU ID	Trap ID	TrapStats ID	Policer Rate	Policer	Packet Accepted	Packet Dropped
ETHERNET_ACL_DROP (D)	0	0	0x0	1	0	0	0
ETHERNET_ACL_FORCE_PUNT (D*)	0	1	0x0	1	0	0	0
ETHERNET_VLAN_MEMBERSHIP (D*)	0	2	0x0	1	0	0	0
ETHERNET_ACCEPTABLE_FORMAT	0	3	0x0	258	100	0	0
UNKNOWN_VLAN_OR_BUNDLE_MEMBER (D*)	0	4	0x0	259	100	0	0
NOT_MY_MAC (D*)	0	5	0x0	260	100	0	0
ETHERNET_NO_SIP_MAPPING (D*)	0	6	0x0	1	0	0	0
ETHERNET_NO_VNI_MAPPING (D*)	0	7	0x0	1	0	0	0
ETHERNET_NO_VSID_MAPPING (D*)	0	8	0x0	1	0	0	0
ARP	0	9	0x0	264	542	0	0
ETHERNET_SA_ERROR (D*)	0	11	0x0	266	100	0	0
ETHERNET_DA_ERROR (D*)	0	12	0x0	1	0	0	0
ETHERNET_SA_MULTICAST (D*)	0	13	0x0	268	100	0	0
DHCPV4_SERVER	0	14	0x0	269	542	0	0
DHCPV4_CLIENT	0	15	0x0	270	200	0	0
ETHERNET_INGRESS_STP_BLOCK (D*)	0	18	0x0	1	0	0	0
PTP_OVER_ETHERNET	0	19	0x0	274	4000	0	0
.
.
.
.
OAMP_BFD_INCORRECT_TTL (D*)	0	157	0x0	412	100	0	0
OAMP_BFD_INVALID_PROTOCOL (D*)	0	158	0x0	413	100	0	0
OAMP_BFD_INVALID_UDP_PORT (D*)	0	159	0x0	414	100	0	0

OAMP_BFD_INCORRECT_VERSION (D*)	0	160	0x0	415	100	0	0
OAMP_BFD_INCORRECT_ADDRESS (D*)	0	161	0x0	416	100	0	0
OAMP_BFD_MISMATCH_DISCR	0	162	0x0	417	500000	0	0
OAMP_BFD_STATE_FLAG_CHANGE	0	163	0x0	418	500000	0	0
OAMP_BFD_SESSION_RECEIVED (D)	0	164	0x0	419	100	0	0
OAMP_PFC_LOOKUP_FAILED (D*)	0	165	0x0	420	100	0	0
OAMP_PFC_DROP_INVALID_RX (D*)	0	166	0x0	1	0	0	0
APP_SGACL_DROP (D*)	0	168	0x0	1	0	0	0

Defining Dynamic LPTS Flow Type

The Dynamic LPTS flow type feature enables you to configure LPTS flow types and also enables you to define the maximum LPTS entries for each flow type in the TCAM. The dynamic LPTS flow type configuration is on per line card basis, hence you can have multiple profiles configured across line cards.

When the router boots, the default LPTS flow types are programmed in the TCAM. For each flow type the maximum flow entries are predefined. Later, at runtime, you have an option to choose the flow type based on network requirements and also configure the maximum flow entry value. The maximum flow entry value of zero denotes that a flow type is not configured.



Note You can get the default maximum flow values for both configurable flow and non-configurable flow from the output of the following show command:

```
show lpts pifib dynamic-flows statistics location <location specification>
```

The list of configurable and non-configurable flow types are listed in below tables. You can also use **show lpts pifib dynamic-flows statistics location** command to view the list of configurable and non-configurable flow types:



Note The sum of maximum LPTS entries configured for all flow types must not exceed 16000 entries per line card.

Configuration Example

In this example you will configure the BGP-known and ISIS-known LPTS flow type in the TCAM and define the maximum flow entries as 1800 and 500 for node location 0/1/CPU0. As the new maximum values are more than the default values, we have to create space in the TCAM by disabling other flow types so that the sum of maximum entries for all flow types per line card does not exceed 8000 entries. Hence RSVP-known flow type is set to zero in our example:

The maximum dynamic scale for any flow type should be configured such that all LPTS entries for that flow type are in hardware. One way to achieve that is to increase the dynamic scale. This may help avoid session flaps for NSR-enabled protocols like BGP and OSPF in case of triggers like RP fail overs.

```
Router#configure
Router(config)#lpts pifib hardware dynamic-flows location 0/1/CPU0
Router(config-pifib-flows-per-node)#flow bgp known max 1800
Router(config-pifib-flows-per-node)#flow rsvp known max 0
Router(config-pifib-flows-per-node)#commit
```

Running Configuration

```
Router#show running-config lpts pifib hardware dynamic-flows location 0/1/CPU0
lpts pifib hardware dynamic-flows location 0/1/CPU0
  flow bgp known max 1800
  flow rsvp known max 0
!
```

Verification

This show command displays dynamic flow statistics. You can see that the flow types BGP-known and ISIS-known are configured in the TCAM with newly configured maximum flow entry value. You can also see that the RSVP-known flow type is disabled:

```
Router#show lpts pifib dynamic-flows statistics location 0/1/CPU0
```

```
Dynamic-flows Statistics:
```

```
-----
(C - Configurable, T - TRUE, F - FALSE, * - Configured)
Def_Max - Default Max Limit
Conf_Max - Configured Max Limit
HWCnt - Hardware Entries Count
ActLimit - Actual Max Limit
SWCnt - Software Entries Count
P, (+) - Pending Software Entries
```

FLOW-TYPE	C	Def_Max	Conf_Max	HWCnt/ActLimit	SWCnt	P
-----	-	-----	-----	-----/-----	-----	-
Fragment	F	2	--	2/2	2	
OSPF-mc-known	T	600	--	2/600	2	
OSPF-mc-default	F	4	--	4/4	4	
OSPF-uc-known	T	300	--	1/300	1	
OSPF-uc-default	F	0	--	0/0	1	+
BFD-default	F	2	--	2/2	2	
BFD-MP-known	T	40	--	1/40	0	
BGP-known	T*	2400	1800	6/900	6	
BGP-cfg-peer	T	900	--	0/900	0	
BGP-default	F	4	--	4/4	4	
PIM-mcast-default	F	40	--	0/40	0	
PIM-mcast-known	T	300	--	0/300	0	
PIM-ucast	F	40	--	2/40	2	
IGMP	T	1200	--	0/1200	0	
ICMP-local	F	4	--	4/4	4	
ICMP-control	F	5	--	5/5	5	
LDP-TCP-known	T	300	--	0/300	0	
LDP-TCP-cfg-peer	T	300	--	0/300	0	
LDP-TCP-default	F	40	--	0/40	0	
LDP-UDP	T	300	--	0/300	0	
All-routers	T	300	--	0/300	0	
RSVP-default	F	4	--	1/4	1	
RSVP-known	T*	300	0	0/0	1	+
SNMP	T	300	--	8/300	8	
SSH-known	T	40	--	0/40	0	
SSH-default	T	1	--	1/1	2	+
HTTP-known	T	40	--	0/40	0	
SHTTP-known	T	40	--	0/40	0	
TELNET-known	T	40	--	0/40	0	
TELNET-default	T	1	--	1/1	1	
UDP-known	T	0	--	0/0	0	
UDP-default	F	2	--	2/2	2	
TCP-known	T	40	--	0/40	0	
TCP-default	F	2	--	2/2	2	

```
Raw-default          F          2    --    2/2          2
GRE                  F          4    --    0/4          0
VRRP                 T        150    --    0/150        0
DNS                  T         40    --    0/40         0
NTP-known            T         40    --    0/40         0
DHCPv4               T         40    --    0/40         0
DHCPv6               T         40    --    0/40         0
TPA                  T        1000    --    0/1000       0
PM-TWAMP             T          10    --    0/10         0
-----
Active TCAM Usage : 13421/16000 [Platform MAX: 16000]
HWCnt/SWCnt       : 65/88
-----
```

In the above show command output, the last column **P** specifies the pending software flow entries for the flow type.

