



Implement LPTS

- [LPTS Overview, on page 1](#)
- [LPTS Policers , on page 1](#)
- [Per Port Rate Limiting of Multicast and Broadcast Punt Packets, on page 6](#)
- [LPTS Domain Based Policers, on page 13](#)
- [Defining Dynamic LPTS Flow Type, on page 14](#)

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 necessary. 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 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

```
RP/0/RP0:hostname#configure
RP/0/RP0:hostname(config)#lpts pifib hardware police
RP/0/RP0:hostname(config-pifib-policer-global)#flow ospf unicast default rate 3000
RP/0/RP0:hostname(config-pifib-policer-global)#flow bgp default rate 4000
RP/0/RP0:hostname (config-pifib-policer-global)#commit
```

Running Configuration

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

Verification

```
RP/0/RP0:hostname#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 - 0/0/CPU0:

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

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

Running Configuration

```
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.

```
RP/0/RP0:hostname#show lpts pifib hardware police location 0/0/CPU0
-----
Node 0/0/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 0/0/CPU0** command displays packets that are locally processed and packets that are dropped by the CPU.

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

Trap Type	NPU ID	Trap ID	TrapStats ID	Policer	Packet	Packet
					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 An interface level rate limit configuration has the highest priority over a global and domain level configurations.

1. `RP/0/RP0:hostname# configure`
Enters the configuration mode.
2. `RP/0/RP0:hostname(config)# lpts punt police`
Enters punt configuration mode.
3. `RP/0/RP0:hostname(config-lpts-punt-policer)# interface TenGigE0/0/0/8/0`
Enters per interface level policer configuration.

4. `RP/0/RP0:hostname(config-lpts-punt-policer-global-if)# mcast rate 1000`
Configures a rate limit of 1000 pps for multicast punted traffic.
5. `RP/0/RP0:hostname(config-lpts-punt-policer-global-if)# bcast rate 1000`
Configures a rate limit of 1000 pps for broadcast punted traffic.
6. `RP/0/RP0:hostname(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. `RP/0/RP0:hostname# configure`
Enters the configuration mode.
2. `RP/0/RP0:hostname(config)# lpts punt police`
Enters punt configuration mode.
3. `RP/0/RP0:hostname(config-punt-policer-global)# mcast rate 1000`
Configures multicast rate limit of 1000 pps.
4. `RP/0/RP0:hostname(config-punt-policer-global)# bcast rate 1000`
Configures broadcast rate limit of 1000 pps.
5. `RP/0/RP0:hostname(config-punt-policer-global)# commit`
Commit the configuration.

Domain Level

This example shows how to configure the LPTS domain and apply a rate limit of:

- 1000 pps for the multicast and broadcast punted traffic

1. `RP/0/RP0:hostname# configure`
Enters the configuration mode.
2. `RP/0/RP0:hostname(config)# lpts punt police domain ACCESS`
Enters LPTS punt domain configuration mode.
3. `RP/0/RP0:hostname(config-lpts-punt-policer-global-ACCESS)# mcast 5000`
Configures multicast rate limit of 5000 pps.
4. `RP/0/RP0:hostname(config-lpts-punt-policer-global-ACCESS)# bcast 5000`
Configures broadcast rate limit of 5000 pps.
5. `RP/0/RP0:hostname(config-lpts-punt-policer-global-ACCESS)# exit`
Exits the domain ACCESS mode.

6. `RP/0/RP0:hostname(config-lpts-punt-policer)# exit`
Exits the LPTS punt configuration mode.
7. `RP/0/RP0:hostname(config)# lpts pifib hardware domain ACCESS`
Enters LPTS hardware domain configuration mode.
8. `RP/0/RP0:hostname(config-pifib-domain-ACCESS)# interface TenGigE0/0/0/8/1`
Applies the domain ACCESS to the TenGigE0/0/0/8/1 interface node.
9. `RP/0/RP0:hostname(config-pifib-domain-ACCESS)# exit`
Exits LPTS domain mode.
10. `RP/0/RP0:hostname(config)# lpts punt police location 0/0/CPU0`
Enters LPTS punt police configuration mode.
11. `RP/0/RP0:hostname(config-lpts-punt-policer)# protocol arp rate 500`
Configures the rate limit of 500 pps for the ARP protocol packets.
12. `RP/0/RP0:hostname(config-lpts-punt-policer)# protocol cdp rate 500`
Configures the rate limit of 500 pps for the CDP protocol packets.
13. `RP/0/RP0:hostname(config-lpts-punt-policer)# exit`
Exits the LPTS punt policer configuration mode.
14. `RP/0/RP0:hostname(config)# lpts punt police location 0/4/CPU0`
Configures LPTS punt police at the node location 0/4/CPU0.
15. `RP/0/RP0:hostname(config)# commit`
Commits the configuration



Note After committing the configuration, verify if an error message is captured in the syslog with respect to the multicast and broadcast rate limit.

Protocol Punted Traffic

You can configure a rate limit to these protocol punted traffic - ARP, CDP, and LACP.

This example shows how to configure the following rate limit for protocol punted traffic at the global level:

- 500 pps for ARP and CDP protocols

1. `RP/0/RP0:hostname(config-punt-policer-global)# protocol arp rate 500`
Configures rate limit of 500 pps for protocol ARP packets.
2. `RP/0/RP0:hostname(config-punt-policer-global)# protocol cdp rate 500`
Configures rate limit of 500 pps for protocol CDP packets.
3. `RP/0/RP0:hostname(config-punt-policer-global)# commit`

Commit the configuration.

This example shows how to configure the following rate limit for protocol punted traffic at the domain level:

- 500 pps for ARP and CDP protocols

1. RP/0/RP0:hostname(config)# lpts pifib hardware domain ACCESS
Enters LPTS hardware domain configuration mode.
2. RP/0/RP0:hostname(config-pifib-domain-ACCESS)# interface TenGigE0/0/0/8/1
Applies the domain ACCESS to the TenGigE0/0/0/8/1 interface node.
3. RP/0/RP0:hostname(config-pifib-domain-ACCESS)# exit
Exits LPTS domain mode.
4. RP/0/RP0:hostname(config)# lpts punt police location 0/0/CPU0
Enters LPTS punt police configuration mode.
5. RP/0/RP0:hostname(config-lpts-punt-policer)# protocol arp rate 500
Configures the rate limit of 500 pps for the ARP protocol packets.
6. RP/0/RP0:hostname(config-lpts-punt-policer)# protocol cdp rate 500
Configures the rate limit of 500 pps for the CDP protocol packets.
7. RP/0/RP0:hostname(config-lpts-punt-policer)# exit
Exits the LPTS punt policer configuration mode.
8. RP/0/RP0:hostname(config)# lpts punt police location 0/4/CPU0
Configures LPTS punt police at the node location 0/4/CPU0.
9. RP/0/RP0:hostname(config)# commit
Commits the configuration

Running Config

```
lpts punt police
interface TenGigE0/0/0/8/0
  mcast rate 1000
  bcast rate 1000
!
mcast rate 1000
bcast rate 1000
protocol arp rate 700
protocol cdp rate 700
domain ACCESS
  mcast rate 5000
  bcast rate 5000
!
!
lpts pifib hardware domain ACCESS
interface TenGigE0/0/0/8/1
!
lpts punt police location 0/0/CPU0
protocol arp rate 500
protocol cdp rate 500
```

```
!
lpts punt police location 0/4/CPU0
!
```

Verification

In the below show command output, you should look for highlighted fields that confirms the rate limit configuration at domain, and interface level:

```
RP/0/RP0:hostname# show lpts punt statistics location 0/0/CPU0
Fri Nov 15 06:23:20.410 UTC
```

```
Lpts Punt Policer Statistics:
```

```
-----
Punt_Reason - Ingress Packets type to be Punt policed
Scope        - Configured scope - Global/Domain/IFH
State        - Current config state
Rate         - Policer rate in PPS
Accepted     - No of Packets Accepted
Dropped     - No of Packets Dropped
Domain      - Domain name
-----
```

```
-----
Interface Name      : any
Punt Reason       : ARP
Domain           : ACCESS
Scope              : Default
State              : Active
Configured Rate  : 1000
Operational Rate   : 986
Accepted           : 0
Dropped           : 0
Last Update (if any):
Punt Type          : ARP
Interface Handle    : 0x00000000
Is Virtual         : 0
Is Enabled         : 1
Packet Rate        : 1000
Domain             : 1
CreateTime        : Fri Nov 15 2019 06:22:42.237.188
Platform:
  PolicerID       : 32398
  NPU: TCAM-entry  StatsID
    0:            172 0x80001d54
    1:            297 0x80001dd0
    2:            172 0x80001d54
    3:            172 0x80001d54
    4:            172 0x80001d54
    5:            172 0x80001d54
-----
```

```
-----
Interface Name      : any
Punt Reason       : CDP
Domain           : ACCESS
Scope              : Default
State              : Active
Configured Rate  : 1000
Operational Rate   : 986
Accepted           : 0
Dropped           : 0
Last Update (if any):
Punt Type          : CDP
Interface Handle    : 0x00000000
-----
```

```

Is Virtual          : 0
Is Enabled         : 1
Packet Rate       : 1000
Domain            : 1
CreateTime        : Fri Nov 15 2019 06:22:42.258.192
Platform:
  PolicerID       : 32404
  NPU: TCAM-entry  StatsID
    0:            173 0x80001d55
    1:            298 0x80001dd1
    2:            173 0x80001d55
    3:            173 0x80001d55
    4:            173 0x80001d55
    5:            173 0x80001d55
-----
Interface Name     : any
Punt Reason      : ARP
Domain            : default
Scope             : Local
State             : Active
Configured Rate  : 500
Operational Rate  : 515
Accepted          : 980
Dropped          : 0
Last Update (if any):
Punt Type         : ARP
Interface Handle   : 0x00000000
Is Virtual        : 0
Is Enabled        : 1
Packet Rate       : 500
Domain            : 0
CreateTime        : Tue Nov 12 2019 06:31:25.136.800
Platform:
  PolicerID       : 32306
  NPU: TCAM-entry  StatsID
    0:            41 0x80001cd2
    1:            41 0x80001cd2
    2:            41 0x80001cd2
    3:            41 0x80001cd2
    4:            41 0x80001cd2
    5:            41 0x80001cd2
-----
Interface Name     : any
Punt Reason      : CDP
Domain            : default
Scope             : Local
State             : Active
Configured Rate   : 500
Operational Rate  : 515
Accepted          : 4292
Dropped          : 0
Last Update (if any):
Punt Type       : CDP
Interface Handle   : 0x00000000
Is Virtual        : 0
Is Enabled        : 1
Packet Rate     : 500
Domain            : 0
CreateTime        : Tue Nov 12 2019 06:31:25.513.897
Platform:
  PolicerID       : 32312
  NPU: TCAM-entry  StatsID
    0:            42 0x80001cd3
    1:            42 0x80001cd3

```

```

2:          42 0x80001cd3
3:          42 0x80001cd3
4:          42 0x80001cd3
5:          42 0x80001cd3

```

```

-----
Interface Name      : TenGigE0
Punt Reason        : MCAST
Domain             : default
Scope             : Global
State              : Active
Configured Rate    : 1000
Operational Rate   : 986
Accepted           : 0
Dropped           : 0
Last Update (if any):
Punt Type          : MCAST
Interface Handle   : 0x0800001c
Is Virtual         : 1
Is Enabled         : 1
Packet Rate        : 1000
Domain             : 0
CreateTime         : Tue Nov 12 2019 06:32:43.210.014
Platform:
  PolicerID        : 32396
  NPU: TCAM-entry  StatsID
    0:             170 0x80001d52
    1:             172 0x80001d53
    2:             170 0x80001d52
    3:             170 0x80001d52
    4:             170 0x80001d52
    5:             170 0x80001d52

```

```

-----
Interface Name      : TenGigE0
Punt Reason        : BCAST
Domain             : default
Scope             : Global
State              : Active
Configured Rate    : 1000
Operational Rate   : 986
Accepted           : 0
Dropped           : 0
Last Update (if any):
Punt Type          : BCAST
Interface Handle   : 0x0800001c
Is Virtual         : 1
Is Enabled         : 1
Packet Rate        : 1000
Domain             : 0
CreateTime         : Tue Nov 12 2019 06:32:43.227.279
Platform:
  PolicerID        : 32397
  NPU: TCAM-entry  StatsID
    0:             171 0x80001d53
    1:             173 0x80001d54
    2:             171 0x80001d53
    3:             171 0x80001d53
    4:             171 0x80001d53
    5:             171 0x80001d53

```

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. */
RP/0/RP0:hostname# config
RP/0/RP0:hostname(config)# lpts pifib hardware domain CORE

/* Configure the interfaces for the domain CORE. */
RP/0/RP0:hostname(config-lpts-domains-CORE)# interface hundredGigE 0/0/0/1
RP/0/RP0:hostname(config-lpts-domains-CORE)# interface hundredGigE 0/0/0/2
RP/0/RP0:hostname(config-lpts-domains-CORE)# commit
RP/0/RP0:hostname(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. */
RP/0/RP0:hostname(config)# lpts pifib hardware police domain CORE
RP/0/RP0:hostname(config-lpts-policer-global-CORE)# flow ospf unicast known rate 6000
RP/0/RP0:hostname(config-lpts-policer-global-CORE)# flow ospf unicast default rate 7000
RP/0/RP0:hostname(config-lpts-policer-global-CORE)# commit
RP/0/RP0:hostname(config-lpts-policer-global-CORE)# exit
RP/0/RP0:hostname(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. */
RP/0/RP0:hostname(config)# lpts pifib hardware police location 0/0/CPU0 domain CORE
RP/0/RP0:hostname(config-lpts-policer-global-CORE)# flow ospf unicast known rate 7000
RP/0/RP0:hostname(config-lpts-policer-global-CORE)# flow ospf unicast default rate 8000
RP/0/RP0:hostname(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:

```

RP/0/RP0:hostname# 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

```

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 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 that are configured for all flow types must not exceed 8000 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:

```
RP/0/RP0:hostname#configure
RP/0/RP0:hostname(config)#lpts pifib hardware dynamic-flows location 0/1/CPU0
RP/0/RP0:hostname(config-pifib-flows-per-node)#flow bgp known max 1800
RP/0/RP0:hostname(config-pifib-flows-per-node)#flow ISIS known max 500
RP/0/RP0:hostname(config-pifib-flows-per-node)#flow RSVP known max 0
RP/0/RP0:hostname(config-pifib-flows-per-node)#commit
```

Running Configuration

```
RP/0/RP0:hostname#
flow bgp known max 1800
flow isis known 500
flow RSVP known 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:

```
RP/0/RP0:hostname#

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	2	--	2/2	2	
ISIS-known	T	300	500	500/300	0	
ISIS-default	F	1	--	1/1	1	

BGP-known	T	900	1800	1800/900	0
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
ICMP-default	F	9	--	9/9	9
ICMP-app-default	F	2	--	2/2	2
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	0/300	0
SNMP	T	300	--	0/300	0
SSH-known	T	150	--	0/150	0
SSH-default	F	40	--	0/40	0
TELNET-known	T	150	--	0/150	0
TELNET-default	F	4	--	0/4	0
UDP-default	F	2	--	2/2	2
TCP-default	F	2	--	2/2	2
Raw-default	F	2	--	2/2	2
GRE	F	4	--	0/4	0
VRRP	T	150	--	150/150	0
DNS	T	40	--	0/40	0
NTP-default	F	4	--	0/4	0
NTP-known	T	150	--	0/150	0
TPA	T	5	--	0/5	0

 Local Limit : **7960**/8000 /*The sum of maximum flow entries configured for all flow types
 per line card is less than 8000*/

HWCnt/SWCnt : 45/51

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