



Implementing LPTS

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.

For a complete description of the LPTS commands listed in this module, refer to the LPTS Commands module of *IP Addresses and Services Command Reference for Cisco NCS 6000 Series Routers*.

Feature History for Implementing LPTS

Release	Modification
Release 5.0.0	LPTS was introduced.
Release 5.2.1	Excessive ARP Punt Protection was supported.
Release 5.2.5	Excessive Punt Flow Trap Interface-based Flow feature was introduced.

- [Prerequisites for Implementing LPTS](#) , on page 1
- [Information About Implementing LPTS](#), on page 1
- [Configuring LPTS Policers](#), on page 4
- [Enabling the Excessive ARP Punt Protection](#), on page 6
- [Configuration Examples for Implementing LPTS Policers](#), on page 6
- [Additional References](#), on page 8

Prerequisites for Implementing LPTS

The following prerequisites are required to implement LPTS:

You must be in a user group associated with a task group that includes the proper task IDs. The command reference guides include the task IDs required for each command. If you suspect user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Information About Implementing LPTS

To implement LPTS features mentioned in this document you must understand the following concepts:

LPTS Overview

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). Pre-IFIB (PIFIB), which is an abbreviated copy of IFIB, is maintained by port arbitrator on route processor. The line card also downloads the PIFIB for fast lookup. While IFIB is only present on RP, PIFIB is present on both RP and LCs. The entries in PIFIB are used for a single lookup with an exact match. The IFIB, along with PIFIB are used to route received packets to the correct Route Processor or line card for processing.

LPTS interfaces internally with all applications that receive packets from outside the router. LPTS functions without any need for customer configuration. However, 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 line cards. 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 line card 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 (locally) and globally using the command-line interface (CLI); therefore, overwriting the static policer values.



Note

If two different ACLs with same ACEs are applied to an LPTS Policer, only the first ACL applied takes effect. When the first ACL is removed, the second ACL does not take effect on the LPTS Policer. If you want the second ACL to take effect on the LPTS Policer, reconfigure it on the LPTS Policer.

Excessive ARP Punt Protection

The Excessive ARP Punt Protection feature attempts to identify and mitigate control packet traffic from remote devices that send more than their allocated share of ARP control packet traffic. A remote device can be a device on a VLAN interface.

When remote devices send ARP control packet traffic to the router, the control packets are punted and policed by a local packet transport service (LPTS) queue to protect the router's CPU. If one device sends an excessive rate of control packet traffic, the policer queue fills up, causing many packets to be dropped. If the rate from one "bad actor" device greatly exceeds that of other devices, most of the other devices do not get any of their control packets through to the router. The Excessive ARP Punt Protection feature addresses this situation.



Note

Even when the Excessive ARP Punt Protection feature is not enabled, the "bad actors" can affect services for only other devices; they cannot bring down the router.

The Excessive ARP Punt Protection feature is supported on non-subscriber interfaces such as L2 and L3 VLAN sub-interfaces and bundle virtual interfaces (BVIS). If the source that floods the punt queue with

packets is a device with an interface handle, then all punts from that bad actor interface are penalty policed. The default penalty rate, which is non-configurable, is 10 packets per second (pps).

Functioning of Excessive ARP Punt Protection Feature

The Excessive ARP Punt Protection feature monitors ARP control packet traffic arriving from non-subscriber interfaces. These could be physical interfaces, sub-interfaces, or BVIs. It divides interfaces into two categories:

- "Parent" interfaces, which can have other interfaces under them.
- "Non-parent" interfaces, which have no interfaces under them.

A physical interface is always a parent interface because it has VLAN sub-interfaces. An L3 VLAN sub-interface can either be a parent or a non-parent interface.

When a flow is trapped, the Excessive ARP Punt Protection feature tries to identify the source of the flow. The first thing it determines is from which interface the flow came. If this interface is not a "parent" interface, then the feature assumes that it is the end-point source of the flow and penalty policing is applied only on the non-parent interface and not the parent interface. The software applies a penalty-policer in the case of a BVI interface also. If the trapped interface is a "parent" interface, then the entire interface is penalized, which would penalize all the interfaces under it.

For more information about enabling the Excessive ARP Punt Protection feature, see [Enabling the Excessive ARP Punt Protection, on page 6](#).



Note The Excessive ARP Punt Protection feature monitors all punt ARP traffic. You can exclude a particular interface on the router from the monitoring but a remote interface cannot be prevented from being flagged as bad if it is the source of excessive flows.

Bad actors are policed for ARP protocol. There is a static punt rate and a penalty rate for ARP protocol. For example, the sum total of all ARP punts from remote devices is policed at 1000 packets per second (pps) to the router's CPU. If one remote device sends an excessive rate of ARP traffic and is trapped, then ARP traffic from that bad actor is policed at 10 pps. The remaining (non-bad) remote devices continue to use the static 1000 pps queue for ARP.



Note The excessive rate required to cause an interface to get trapped has nothing to do with the static punt rate (that is, 1000 pps). The excessive rate is a rate that is significantly higher than the current average rate of other control packets being punted. The excessive rate is not a fixed rate, and is dependent on the current overall punt packet activity.

When an interface is trapped, it is placed in a "penalty box" for a period of time (a default of 15 minutes). At the end of the penalty timeout, it is removed from penalty policing (that is, packet dropping). If there is still an excessive rate of ARP control packet traffic coming from the remote device, then the remote interface is trapped again.

Interface-based Flow

For the Elephant Trap sampler, the MAC address is one of the key fields used to uniquely identify a flow. Certain cases of DoS attacks have dynamically changing source MAC addresses. An individual flow does not cross the threshold in such cases, and hence the Excessive Punt Flow Trap (EPFT) does not trap the flow.

With the interface-based flow feature, Elephant Trap does not consider MAC addresses as a key for uniquely identifying a flow. Hence, all packets received on a non-subscriber interface (irrespective of the source MAC address) are considered to be a part of a single flow. When excessive punts are received on the interface, EPFT does *ifhandle*-based trap, thereby penalty policing the punt traffic on that particular interface.

To enable interface-based flow, use the following command in global configuration mode:

lpts punt excessive-flow-trap interface-based-flow



Note

You cannot enable this command if EPFT is turned on for the subscriber-interfaces and non-subscriber-interfaces MAC, or vice versa. This is because interface-based flow feature is mutually exclusive with MAC-based EPFT on non-subscriber interface feature.

Restrictions

These restrictions apply to implementing Excessive ARP Punt Protection feature:

- This feature is non-deterministic. In some cases, the Excessive ARP Punt Protection feature can give a false positive, that is, it could trap an interface that is sending legitimate punt traffic.
- The Excessive ARP Punt Protection feature traps flows based on the relative rate of different flows; thus, the behavior depends on the ambient punt rates. A flow that is significantly higher than other flows could be trapped as a bad actor. Thus, the feature is less sensitive when there are many flows, and more sensitive when there are fewer flows present.
- Sometimes control packet traffic can occur in bursts. The Excessive ARP Punt Protection has safeguards against triggering on short bursts, but longer bursts could trigger a false positive trap.

Configuring LPTS Policers

This task allows you to configure the LPTS policers.

SUMMARY STEPS

1. **configure**
2. **lpts pifib hardware police** [location *node-id*]
3. **flow** *flow_type* {**default** | **known**} {**rate** *rate*}
4. **commit**
5. **show lpts pifib hardware policer** [location {**all** | *node_id*}]

DETAILED STEPS

	Command or Action	Purpose
Step 1	configure	

	Command or Action	Purpose
Step 2	<p>lpts pifib hardware police [location <i>node-id</i>]</p> <p>Example:</p> <pre>RP/0/RP0/CPU0:router(config)# lpts pifib hardware police location 0/2/CPU0 RP/0/RP0/CPU0:router(config-pifib-policer-per-node)# RP/0/RP0/CPU0:router(config)# lpts pifib hardware police RP/0/RP0/CPU0:router(config-pifib-policer-global)#</pre>	<p>Configures the ingress policers and enters pifib policer global configuration mode or pifib policer per node configuration mode.</p> <p>The example shows pifib policer per node configuration mode and global.</p>
Step 3	<p>flow <i>flow_type</i> {default known} {rate <i>rate</i>}</p> <p>Example:</p> <pre>RP/0/RP0/CPU0:router(config-pifib-policer-per-node)# flow ospf unicast default rate 20000</pre>	<p>Configures the policer for the LPTS flow type. The example shows how to configure the policer for the ospf flow type.</p> <ul style="list-style-type: none"> Use the <i>flow_type</i> argument to select the applicable flow type. For information about the flow types, see <i>IP Addresses and Services Command Reference for Cisco NCS 6000 Series Routers</i>. Use the rate keyword to specify the rate in packets per seconds (PPS). The range is from 0 to 4294967295. <p>Note LPTS policy for ntp-default flow type, supports a flow rate of 100 pps on Cisco ASR 9000 Series Router.</p> <p>Starting with Cisco IOS XR Release 6.1.3, LPTS policy for ntp-default flow type, supports a flow rate higher than 100 pps on Cisco ASR 9000 Series Router.</p> <p>Based on the number of NTP client scale requirement, you can increase the flow rate value to allow higher packets per second (PPS). For example,</p> <pre>lpts pifib hardware police location 0/0/CPU0 flow ntp default rate 1000 flow ntp known rate 1000</pre>
Step 4	commit	
Step 5	<p>show lpts pifib hardware policer [location {all <i>node_id</i>}]</p> <p>Example:</p> <pre>RP/0/RP0/CPU0:router# show lpts pifib hardware policer location 0/2/cpu0</pre>	<p>Displays the policer configuration value set.</p> <ul style="list-style-type: none"> (Optional) Use the location keyword to display pre-Internal Forwarding Information Base (IFIB) information for the designated node. The <i>node-id</i> argument is entered in the <i>rack/slot/module</i> notation. Use the all keyword to specify all locations.

Enabling the Excessive ARP Punt Protection

Perform this task to enable the Excessive ARP Punt Protection feature for non-subscriber interfaces. You can also set the penalty timeout and disable the protection on a particular interface in the router.

SUMMARY STEPS

1. **configure**
2. **lpts punt excessive-flow-trap non-subscriber-interfaces**
3. **lpts punt excessive-flow-trap penalty-timeout arp *time***
4. *(Optional)* **lpts punt excessive-flow-trap exclude *interface-type interface-id***
5. **commit**

DETAILED STEPS

	Command or Action	Purpose
Step 1	configure	
Step 2	lpts punt excessive-flow-trap non-subscriber-interfaces Example: RP/0/RP0/CPU0:router(config)# lpts punt excessive-flow-trap non-subscriber-interfaces	Enables the Excessive ARP Punt Protection feature on non-subscriber interfaces.
Step 3	lpts punt excessive-flow-trap penalty-timeout arp <i>time</i> Example: RP/0/RP0/CPU0:router(config)# lpts punt excessive-flow-trap penalty-timeout arp 10	Sets the penalty timeout value, which is a period of time that the interface trapped is placed in the penalty box, for a protocol. The penalty timeout value is in minutes and ranges from 1 to 1000. The default penalty timeout value is 15 minutes.
Step 4	<i>(Optional)</i> lpts punt excessive-flow-trap exclude <i>interface-type interface-id</i> Example: RP/0/RP0/CPU0:router(config)# lpts punt excessive-flow-trap exclude ethernet 0/0/0/1	Disables the monitoring of ARP punt protection feature on the specified interface.
Step 5	commit	

Configuration Examples for Implementing LPTS Policers

This section provides the following configuration example:

Configuring LPTS Policers: Example

The following example shows how to configure LPTS policers:

```

configure
lpts pifib hardware police
  flow ospf unicast default rate 200
  flow bgp configured rate 200
  flow bgp default rate 100
!
lpts pifib hardware police location 0/2/CPU0
  flow ospf unicast default rate 100
  flow bgp configured rate 300
!

```

The following is the show command and the sample output:

```
show lpts pifib hardware police location 0/2/CPU0
```

```
Node: 0/2/CPU0:
```

```

-----
flow_type          priority sw_police_id hw_policer_addr avgrate burst static_avgrate
avgrate_type
-----
-----
unconfigured-default low      0           580096           500      100      500           2
UDP-default         low      1           580608           500      100      500           2
TCP-default         low      2           581120           500      100      500           2
Mcast-default       low      3           581632           500      100      500           2
Raw-listen          low      4           582144           500      100      500           2
Raw-default         low      5           582656           500      100      500           2
Fragment            low      6           583168           1000     100      1000          2
OSPF-mc-known       high     7           583680           2000     1000     2000          2
ISIS-known          high     8           584192           2000     1000     2000          2
EIGRP               high     9           584704           1500     750      1500          2
RIP                 high    10           585216           1500     750      1500          2
OSPF-mc-default     low     11           585728           1500     1000     1500          2
ISIS-default        low     12           586240           1500     1000     1500          2
BGP-known           high    13           586752           2500     1200     2500          2
BGP-cfg-peer        mdeium  14           587264           100      1000     2000          0
BGP-default         low     15           587776           100      750      1500          1
PIM-mcast-default   mdeium  16           588288           23000    100      23000         2
PIM-ucast           low     17           588800           10000    100      10000         2
IGMP                mdeium  18           589312           3500     100      3500          2
ICMP-local          mdeium  19           589824           2500     100      2500          2
ICMP-app            low     20           590336           2500     100      2500          2
ICMP-default        low     21           590848           2500     100      2500          2
LDP-TCP-known       mdeium  22           591360           2500     1250     2500          2
LMP-TCP-known       mdeium  23           591872           2500     1250     2500          2
RSVP-UDP            mdeium  24           592384           7000     600      7000          2
RSVP-default        mdeium  25           592896           500      100      500           2
RSVP-known          mdeium  26           593408           7000     600      7000          2
IKE                 mdeium  27           593920           1000     100      1000          2
IPSEC-default       low     28           594432           1000     100      1000          2
IPSEC-known         mdeium  29           594944           3000     100      3000          2
MSDP-known          mdeium  30           595456           1000     100      1000          2
MSDP-cfg-peer       mdeium  31           595968           1000     100      1000          2
MSDP-default        low     32           596480           1000     100      1000          2
SNMP                low     33           596992           2000     100      2000          2
NTP-default         high    34           597504           500      100      500           2
SSH-known           mdeium  35           598016           1000     100      1000          2
SSH-default         low     36           598528           1000     100      1000          2
HTTP-known          mdeium  37           599040           1000     100      1000          2
HTTP-default        low     38           599552           1000     100      1000          2
SHTTP-known         mdeium  39           600064           1000     100      1000          2
SHTTP-default       low     40           600576           1000     100      1000          2
TELNET-known        mdeium  41           601088           1000     100      1000          2
TELNET-default      low     42           601600           1000     100      1000          2
CSS-known           mdeium  43           602112           1000     100      1000          2

```

CSS-default	low	44	602624	1000	100	1000	2
RSH-known	mdeium	45	603136	1000	100	1000	2
RSH-default	low	46	603648	1000	100	1000	2
UDP-known	mdeium	47	604160	25000	100	25000	2
TCP-known	mdeium	48	604672	25000	100	25000	2
TCP-listen	low	49	605184	25000	100	25000	2
TCP-cfg-peer	mdeium	50	605696	25000	100	25000	2
Mcast-known	mdeium	51	606208	25000	100	25000	2
LDP-TCP-cfg-peer	mdeium	52	606720	2000	1000	2000	2
LMP-TCP-cfg-peer	mdeium	53	607232	2000	1000	2000	2
LDP-TCP-default	low	54	607744	1500	750	1500	2
LMP-TCP-default	low	55	608256	1500	750	1500	2
UDP-listen	low	56	608768	4000	100	4000	2
UDP-cfg-peer	mdeium	57	609280	4000	100	4000	2
LDP-UDP	mdeium	58	609792	2000	1000	2000	2
LMP-UDP	mdeium	59	610304	2000	1000	2000	2
All-routers	high	60	610816	1000	500	1000	2
OSPF-uc-known	high	61	611328	2000	1000	2000	2
OSPF-uc-default	low	62	611840	100	100	1000	0
ip-sla	high	63	612352	10000	100	10000	2
ICMP-control	high	64	612864	2500	100	2500	2
L2TPv3	mdeium	65	613376	25000	100	25000	2
PCEP	mdeium	66	613888	100	200	100	2
GRE	high	67	614400	1000	1000	1000	2
VRRP	mdeium	68	614912	1000	1000	1000	2
HSRP	mdeium	69	615424	400	400	400	2
BFD-known	critical	70	615936	8500	300	8500	2
BFD-default	critical	71	616448	8500	100	8500	2
MPLS-oam	mdeium	72	616960	100	100	100	2
DNS	mdeium	73	617472	500	100	500	2
RADIUS	mdeium	74	617984	7000	600	7000	2
TACACS	mdeium	75	618496	500	100	500	2
PIM-mcast-known	mdeium	76	619008	23000	100	23000	2
BFD-MP-known	mdeium	77	619520	8400	1024	8400	2
BFD-MP-0	mdeium	78	620032	128	100	128	2
L2TPv2-default	mdeium	79	620544	700	100	700	2
NTP-known	high	80	621056	500	100	500	2
L2TPv2-known	mdeium	81	621568	2000	100	2000	2

Enabling Excessive ARP Punt Protection: Example

This example shows the Excessive ARP Punt Protection enabled for non-subscriber interfaces with the ARP penalty timeout set to two minutes and the protection disabled on one of the interfaces on the router.

```
!
configure
lpts punt excessive-flow-trap non-subscriber-interfaces
lpts punt excessive-flow-trap penalty-timeout arp 2
lpts punt excessive-flow-trap exclude interface TenGigE 0/0/0/0
end
!
```

Additional References

The following sections provide references related to implementing LPTS.

Related Documents

Related Topic	Document Title
Cisco IOS XR LPTS commands: complete command syntax, command modes, command history, defaults, usage guidelines, and examples	<i>Cisco LPTS Commands</i> module in the <i>IP Addresses and Services Command Reference for Cisco NCS 6000 Series Routers</i>

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: https://mibs.cloudapps.cisco.com/ITDIT/MIBS/servlet/index

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	http://www.cisco.com/techsupport

