Cisco Catalyst 6500 Series Supervisor Engine 2T: Management and Monitoring

White Paper

Contents

Introduction

The management capabilities of a switch contribute greatly to the ease of deployment and maintenance of a network. Cisco® Catalyst® 6500 Series Switches include several monitoring features for Supervisor Engine 720-based systems. The next-generation Cisco Catalyst 6500 Series Supervisor Engine 2T, recently released with 12.2(50)SY software, introduces a host of new features in this area. Supervisor Engine 2T brings in the capability to view traffic statistics on a more granular level, which extends to interface counters, quality of service (QoS) packet counts and statistics related to control-plane policing. These enhancements, along with Cisco IOS NetFlow developments, help network administrators monitor traffic flows through the switch more closely. In addition, the new supervisor engine's baseboard allows for remote console access and manageability to the switch via a Connectivity Management Processor (CMP). The Supervisor Engine 2T software also incorporates Embedded Event Manager (EEM) Version 3.0, an enhanced version of EEM. This paper discusses the key enhancements related to manageability, delivered by the Supervisor Engine 2T with Cisco IOS 12.2(50)SY.

Connectivity Management Processor (CMP)

A basic requirement in setting up and maintaining a switch is to have access to the device and to ensure reliable connectivity to it. The primary methods of accessing a switch in order to configure, verify or troubleshoot it are through a console or a Telnet/Secure Shell (SSH) connection. It is possible, however, for network issues to render both of these links inaccessible. For example, a high CPU condition or a denial-of-service (DoS) attack can effectively take down a Telnet connection. A hung Route Processor (RP) could lock up the console, leaving it impossible to recover or determine the state of the switch with no way to login to it. In such situations, network administrators may have to resort to resetting the switch to recover it at the cost of losing the ability to efficiently correct the problem. This, in turn, significantly contributes to network downtime.

A solution in such critical situations is to have a 'backdoor' access into the switch to help restore it. This is addressed on next-generation supervisor engines through a new processor called the Connectivity Management Processor (CMP), which exists in conjunction with the primary Route Processor (RP).

The CMP is an independent processor dedicated to switch management and has its own RAM, bootflash and front panel management Ethernet port. While the CMP and RP share the same console, a multiplexer enables switching between the two, providing access to the CMP even if the primary RP is hung. From within the CMP, a host of corrective actions can be taken to restore the switch. Examples of how the CMP can be used include system recovery (of the control plane), system resets and reboots and the copying of Cisco IOS image files in the event that the primary IOS image gets corrupted or deleted.

The Supervisor Engine 2T is the first supervisor in the Cisco Catalyst 6500 Series to provide this functionality. As part of the Multilayer Switch Feature Card 5 (MSFC5) complex that houses the CPU and serves as the control plane on the supervisor, CMP functionality has been added on the board (Figure 1). More details on the hardware can be found at the Cisco Catalyst 6500 Supervisor 2T Architecture White Paper.

Figure 1. Cisco Catalyst 6500 Supervisor Engine 2T

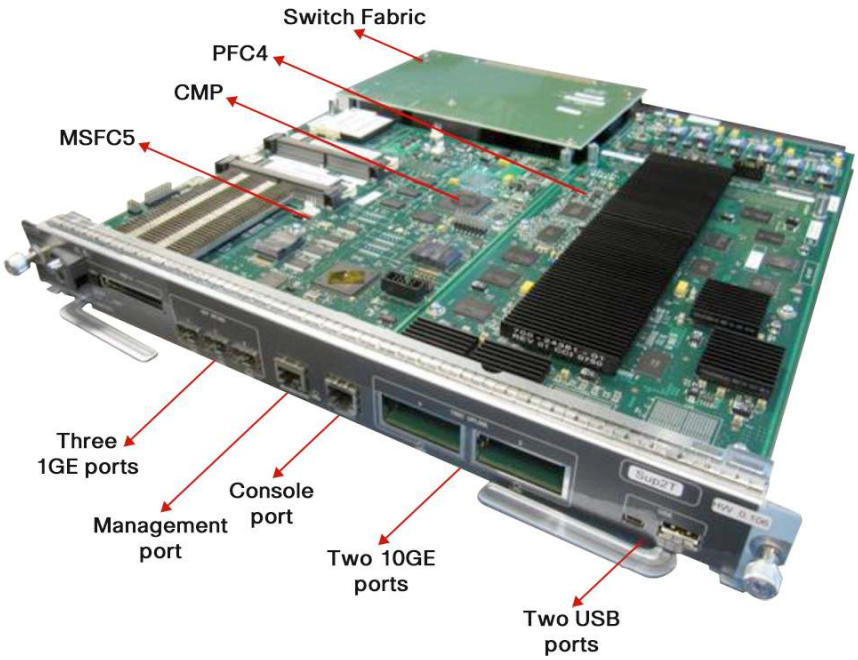

The CMP and the RP share the same console through a programmable multiplexer, which by default has the RP console active. The multiplexer intercepts specific escape sequences that instruct it to switch from one console to the other if required. The CMP runs its own image based on Linux, and the CMP and RP software run independent of each other. The block diagram of the MSFC5 in Figure 2 illustrates the structure of the RP and CMP.

Figure 2. Block Diagram of MSFC5

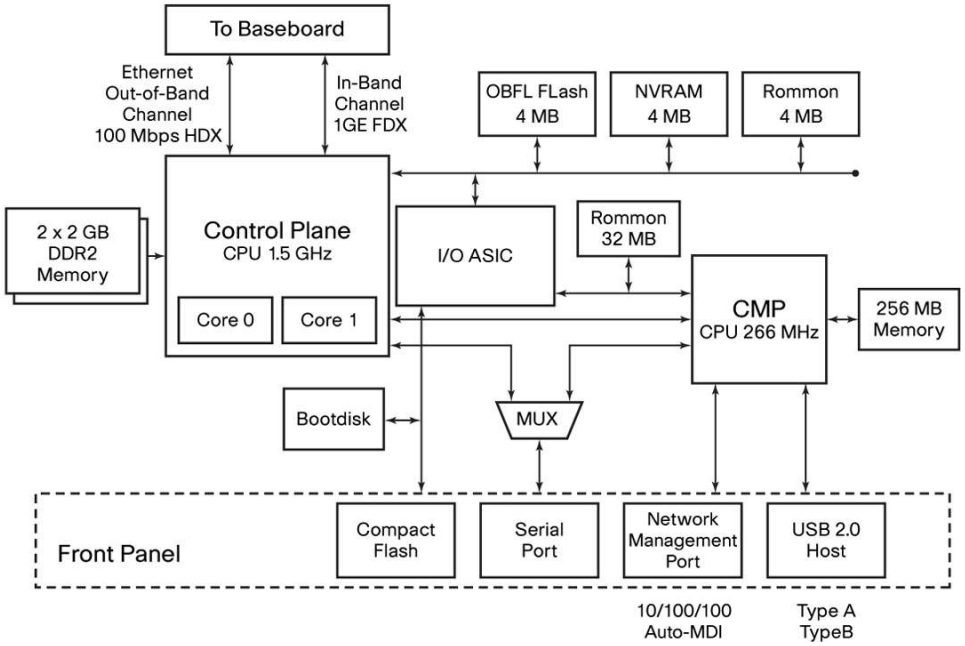

The CMP on the Supervisor Engine 2T supports a front panel 10/100/1000 management port. This port supports auto negotiation, detection, and correction of pair swaps (MDI crossover) and support for jumbo frames up to 9216 bytes in size. It also uses a bi-color LED on the front panel port to show activity and link status.

The following sections highlight some of the device management functions that the CMP brings to the table.

Remote Console

- Ability to access the console remotely (SSH/Telnet) without a terminal server for troubleshooting and recovery (even when the RP is not responsive).

This can enable the CMP to be used as a console of last resort if the primary console locks up. Access to the CMP can be via either the console or the management port. On the console, **Ctrl-C, Shift-M** three times switches to the CMP console, and **Ctrl-R, Shift-M** three times switches back to the RP console. The dedicated out-of-band Ethernet management port can also be configured (from within the CMP) with an IP address, to enable SSH or Telnet access to the CMP. Access to the CMP requires a username and password. Once connected to the CMP, it is possible to attach to the RP as well.

Note: SSH is enabled by default on the management interface. Telnet is disabled by default, and has to be enabled as shown below before a Telnet session can be started.

The following example shows how to log in to the CMP:

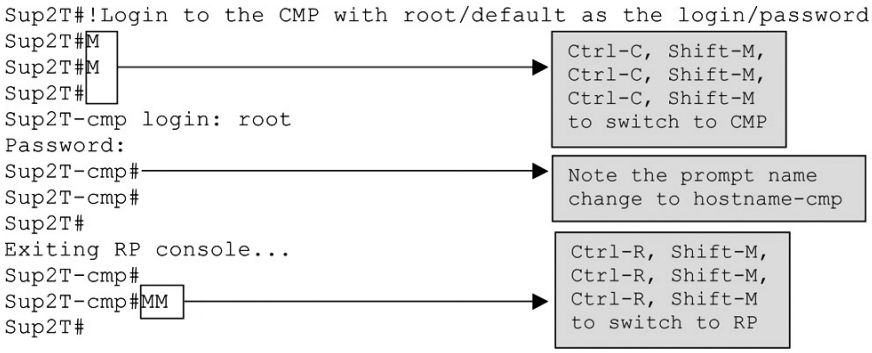

From within the CMP, the management interface can be configured with an IP address and set up for Telnet access, as in the following example:

```
Sup2T-cmp# configure terminal
Sup2T-cmp(config)# interface cmpmgmt
Sup2T-cmp(config-if)# ip address 192.168.0.2/24
Sup2T-cmp(config-if)# ip default gateway 192.168.0.1
Sup2T-cmp(config-if)# exit
Sup2T-cmp(config)# telnet server enable
Sup2T-cmp(config-if)# end
Sup2T-cmp# write memory

Building configuration...
[OK]
Sup2T#
Sup2T-cmp# show interface cmpmgmt
eth0      Link encap:Ethernet  HWaddr 00:1D:E5:E9:11:90
          inet addr:192.168.0.2  Bcast:192.168.0.255  Mask:255.255.255.0
          inet6 addr: fe80::21d:e5ff:fee9:1190/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:20629 errors:0 dropped:0 overruns:0 frame:0
          TX packets:19944 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
```

```
          RX bytes:7245840 (6.9 MiB)  TX bytes:1507837 (1.4 MiB)
          Base address:0x4000
Sup2T-cmp#
Sup2T-cmp# show running-config
ipaddr=192.168.0.2
netmask=255.255.255.0
gatewayip=192.168.0.1
hostname=Sup2T-cmp
telnet=enable
Sup2T-cmp#
```

You can Telnet directly to the CMP and attach to the RP from there if required. The attach option only works with Telnet or SSH, and not from the front panel console.

```
Sup2T-cmp#
Sup2T-cmp# attach
Attaching to RP console, use CTRL X to exit console
Sup2T>
Sup2T>en
Password:
Sup2T#
```

Here is a list of CMP commands available from enable mode and config mode:

```
Sup2T-cmp#
  attach        Attach RP console
  cd            Change current directory
  configure     Configuration from vty interface
  copy          Copy from one file to another
  delete        Delete a file
  detach        Detach RP console
  diagnostic    Diagnostic Command
  dir           List files on a filesystem
  disable       Turn off privileged mode command
  end           End current mode and change to enable mode
  exit          Exit current mode and down to previous mode
  help          Description of the interactive help system
  list          Print command list
  ping          Send echo messages
  pwd           Display current working directory
  quit          Exit current mode and down to previous mode
  reload        Reload System
  show          Show running system information
  traceroute    Trace route to destination
  upgrade       upgrade utilities
  write         Write running configuration to memory, network, or terminal
Sup2T-cmp#

Sup2T-cmp(config)#
  autoboot   Autoboot Enable/Disable
  end        End current mode and change to enable mode
  exit       Exit current mode and down to previous mode
  help       Description of the interactive help system
```

```
  hostname   Set system's network name
  interface  Select cmpmgmt interface to configure
  list       Print command list
  no         Negate a command or set its defaults
  password   Change password
  telnet     Telnet Enable/Disable
Sup2T-cmp(config)#
```

Remote Image Recovery

- Ability to quickly recover the IOS image remotely via a Trivial File Transfer Protocol (TFTP) server without the need for a terminal server or on-site personnel.

If IOS crashes and the IOS image on the Compact Flash is corrupted or deleted, the CMP can be used to boot an image from ROM monitor (ROMmon) mode. Once the **cmpmgmt** port is configured with an IP address and connectivity is established to a TFTP server, you can boot the image from the TFTP server using the following command in ROMmon:

boot tftp://<tftp server address>/<file path>/<file name>

```
rommon 4 > boot tftp:
 Syntax://xx.xx.xx.xx:/filename
Usage: boot tftp://xx.xx.xx.xx/<file name>
rommon 5 >
rommon 5 > boot tftp://9.20.1.19/Software/s2t54-adventerprisek9-mz.122-50.SY
Download Start
CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
CC
Download Completed! Booting the image.
Self extracting the image... [OK]
Self decompressing the image : ########################################
<snip>
```

Remote Reset of Switch

- Ability to reset the Supervisor Engine 2T remotely without any external power management devices.

The Supervisor Engine 2T can be reset from the CMP using the following commands:

reload route-processor soft: CMP sends a message to the RP to gracefully reset itself as it would with a normal reload.

reload route-processor hard: Toggles the RP reset line provided by an FPGA to cause the RP to reset.

The reset reason can be verified using the **show version** command on the RP, and with the **show system reset-reason** command on the CMP.

```
Sup2T# show system reset-reason
Jan  1 00:02:33 cmp [1362]: CMP_RST_REASON = warm reset
Jan  1 00:00:21 cmp [1368]: CMP_RST_REASON = software reset
Jan  1 00:00:44 cmp [1368]: CMP_RST_REASON = software reload
Jan  1 00:02:16 cmp [1355]: CMP_RST_REASON = cold reset
```

```
Jun 16 23:51:58 cmp [1355]: RP_RST_REASON = power on
Jun 16 23:50:58 cmp [1355]: CMP_RST_REASON = cold reset
Nov 30 00:02:12 cmp [1363]: RP_RST_REASON = s/w reset
Nov 30 00:05:04 cmp [1363]: RP_RST_REASON = reload
Sup2T#
```

Remote Console/RP Message Logging

- Ability to log messages from the RP to the CMP

In a situation where the primary console is unresponsive, CMP provides the ability to access the switch logs, in order to determine the possible causes. This would help in taking appropriate corrective action to restore the switch. The capability to view RP logs exists even when the RP is hung, since the messages previously sent by the RP are stored by the CMP in a flash partition reserved for syslogs.

```
Sup2T-cmp# show logging route-processor console
Jun 28 12:00:28 C9E :  *Jun 28 11:59:56.113: %C6KENV-4-MINORTEMPALARM: RP 7/0
device-1 temperature crossed threshold #1(=60C). It has exceeded normal operating
temperature range.
Jun 28 12:16:46 C9E :  *Jun 28 12:00:27.645: %C6KENV-4-MINORTEMPALARMRECOVER: RP
7/0 device-1 temperature crossed threshold #1(=60C). It has returned to normal
operating temperature range.
Jun 28 12:17:18 C9E :  *Jun 28 12:16:45.437: %C6KENV-4-MINORTEMPALARM: RP 7/0
device-1 temperature crossed threshold #1(=60C). It has exceeded normal operating
temperature range.
Sup2T-cmp#
```

While exporting logs to a syslog server is a feature available only from the RP, the CMP stores these files using a Linux syslog server. Linux "logrotate" is used for archiving the old files in a compressed format. It is then possible to transfer the archived files or the current file to a TFTP or Secure Copy (SCP) server using the copy command.

CMP implementation does not support features such as SNMP and Network Time Protocol (NTP). However, the CMP gets regular clock updates from Cisco IOS Software, which itself uses NTP.

Embedded Event Manager (EEM) 3.0

Cisco IOS Embedded Event Manager (EEM) is a unique subsystem within Cisco IOS Software. EEM is a powerful and flexible tool to automate tasks and customize the behavior of Cisco IOS Software and the operation of the device. Customers can use EEM to create and run programs or scripts directly on a router or switch. The scripts are referred to as EEM policies and can be programmed using a simple CLI-based interface or using a scripting language called Tool Command Language (Tcl). EEM allows customers to harness the significant intelligence within Cisco IOS Software to respond to real-time events, automate tasks, create customer commands and take local automated action based on conditions detected by the Cisco IOS Software itself.

EEM was first introduced on the Cisco Catalyst 6500 Series in Cisco IOS Software Release 12.2(18)SXF4. The software version 12.2(33)SXI on the Supervisor 720 supports EEM Version 2.4. The Supervisor Engine 2T with 12.2(50)SY software supports EEM Version 3.0 which includes a set of new event detectors, among other features.

The Cisco IOS Embedded Event Manager is a product-independent software feature consisting of a series of event detectors, an Embedded Event Manager Server and interfaces that allow action routines, called policies, to be invoked. There are also internal application programming interfaces that enable other Cisco IOS subsystems to take advantage of the EEM subsystem. Figure 3 illustrates the components of EEM.

Figure 3. Embedded Event Manager Architecture

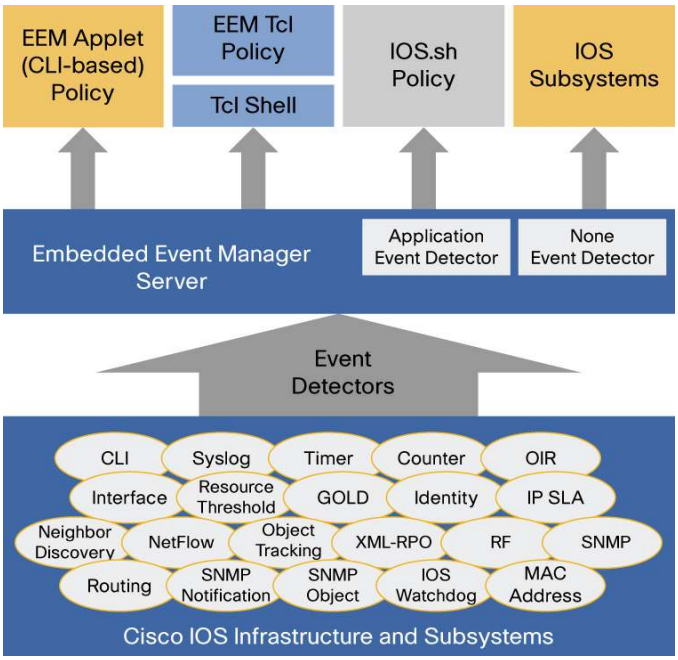

Notice that there are two types of EEM policies:

- Applet policies: Applet policies have an easy-to-use interface and are defined using the configuration CLI. They can be saved to the startup configuration file.
- Tcl policies: These policies are more flexible and have extensive capabilities. Tcl policies are defined using the Tcl programming language.

Once one or more policies are defined, the Event Detector software will watch for the conditions that match those defined by the policy. When a condition occurs, the event is passed to the Event Manager Server. The server in turn invokes any policy that has registered for that particular event. The actions defined within the policy are then carried out.

The current version of the EEM subsystem supported on the Supervisor Engine 2T is EEM Version 3.0. This version ushers in a significant number of enhancements over previous versions, including enhanced performance, increased feature integration, new capabilities and extended flexibility, enabling EEM to be used in new and exciting ways.

Event Detectors with EEM 3.0

A list of event detectors originally supported on the platform can be found here: http://www.cisco.com/en/US/prod/collateral/switches/ps5718/ps708/prod_white_paper0900aecd805457c3.html.

Included in EEM 3.0 are four new event detectors:

Routing Event Detector

This event detector monitors the events relative to the Routing Information Base (RIB). Events are raised for conditions such as when a particular route is added or removed, or when a route is modified.

For example, here is how a notification can be logged if a static route is added:

```
Sup2T(config)#event manager applet staticRouteAdd
Sup2T(config-applet)#event routing network 192.168.1.0/24 protocol Static type add
Sup2T(config-applet)#action 1.0 syslog msg “Static Route add”
Sup2T(config-applet)#end
```

```
Sup2T#
```

Trigger the policy by configuring a static route:

```
Sup2T(config)#ip route 192.168.1.0 255.255.255.0 10.0.0.3
```

Verify that the log message is generated:

```
%HA_EM-6-LOG: staticRouteAdd: Static Route add
```

Flexible NetFlow (FNF) Event Detector

This event detector can be used to detect events related to Flexible NetFlow. It provides a powerful set of triggers to detect and react to real-time network activity. It triggers policies based on the detection of flows that match particular criteria, such as when a new flow is seen with a particular destination IP address and port number. This event detector also detects conditions such as when the rate of new flow entries exceeds a defined threshold.

Figure 4 shows a Flexible NetFlow Event Detector that is configured to be triggered any time a TTL=0 packet is detected through NetFlow. Once this packet is detected, a syslog message is generated to inform the network operations team that such a flow record has been detected.

Figure 4. EEM and FNF Detect Malformed Packets

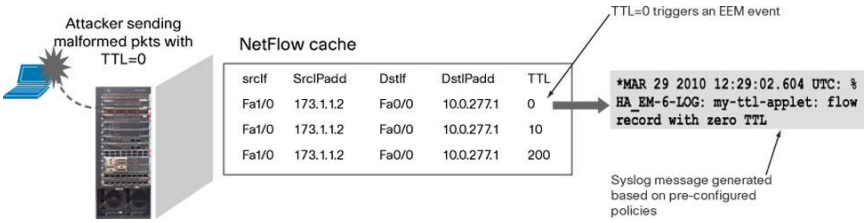

IP SLA Event Detector

This event detector provides event triggers based on Cisco IOS IP Service Level Agreements (SLA) operation results. It integrates IP SLA directly with the EEM subsystem and provides an event-driven mechanism to take immediate action when an IP SLA operation fails. Figure 5 shows a sample use case with the IP SLA Routing Event Detector.

Figure 5. Sample Use Case for EEM with the IP SLA Event Detector

Customer	Large enterprise customer
Problem	Need to provide reliable connectivity for servers behind fast-moving mobile access routers with dual link (Wi-Fi and digital cell). DHCP and static routing is used for both primary and backup links. Link failover to backup when repeated failure detected on primary and stay on backup until it fails. Hundreds of static Network Address Translation (NAT) config tied to primary interface needs automatic translation during failover.
Solution	Use Cisco IP SLAs for end-to-end link state monitoring and triggering EEM events. EEM policy is provided for static route and Network Address Translation (NAT) configuration changes.
Benefit	Increased network availability and reduced OpEx No alternative solution available today
Category	High availability
Additional Information	This use case can be applied generically to any wired network as well.

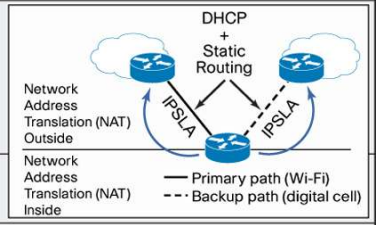

Enhanced CLI Event Detector

This event detector offers enhancements to make creation of custom CLI commands easier and more powerful. It provides new event triggers when special characters like “Tab”, “?”, and the “Enter” key are detected. It also provides a way to offer “Help” for the new commands and make them similar to commands developed by Cisco.

Once the CLI Event Detector is configured, there are commands available to verify that the EEM applet has been registered successfully. Here is an example to verify the registration of the applet, with the event CLI for the “?” character:

```
Sup2T#config terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Sup2T(config)#event manager applet CustomCliED_FF_applet_2_ques
Sup2T(config-applet)#event cli pattern “show cdp” questionmark
Sup2T(config-applet)#action 2 syslog msg “CustomCliED_FF_applet_2_ques”
Sup2T(config-applet)#end
Sup2T#

Sup2T#show event manager policy registered | include CustomCliED_FF_applet_2_ques

9  applet    user    cli        Off   Sat Nov 13 18:00:33 2010
CustomCliED_FF_applet_2_ques
 action 2 syslog msg “CustomCliED_FF_applet_2_ques”
```

Features of EEM 3.0

EEM 3.0 introduces several other features, noteworthy among which are the applet enhancements and high performance Tcl policies. Applets now include support for variables and logical functions, and if-then-else constructs. Here is a list of the major new functions added for customer usability:

- Class based scheduling
- High-performance Tcl policies
- Interactive CLI support with applets
- Variable logic for applets
- Digital signature support
- Support authenticating e-mail servers
- SMTP IPv6 support
- SNMP Tcl extensions (Get, Set and Notify for local and remote hosts)
- SNMP Proxy IPv6 support
- CLI Library XML-PI support

More information on these features can be found in the following document: http://www.cisco.com/en/US/docs/ios/netmgmt/configuration/guide/nm_eem_overview_ps6441_TSD_Products_Configuration_Guide_Chapter.html.

Interface Counters

Traditionally, interface counters viewed under the show interfaces command have reflected packet and byte counts for the following classes of traffic:

- L2 Bridged Unicast
- L2 Bridged Multicast

- L3 IN Unicast
- L3 IN Multicast
- L3 OUT Unicast
- L3 OUT Multicast

Here, L2 packets refer to those that arrive and leave on the same Layer 2 VLAN, and L3 packets imply that the input packet before forwarding and the output packet after forwarding are on different Layer 2 VLANs. These six counters reflect packet and byte counts, as seen in the following output:

```
Sup720#show interfaces GigabitEthernet 2/2
GigabitEthernet2/2 is up, line protocol is up (connected)
<snip>
5 minute input rate 1000 bits/sec, 1 packets/sec
5 minute output rate 0 bits/sec, 0 packets/sec
L2 Switched: ucast: 0 pkt, 0 bytes - mcast: 35 pkt, 2274 bytes
L3 in Switched: ucast: 0 pkt, 0 bytes - mcast: 0 pkt, 0 bytes mcast
L3 out Switched: ucast: 0 pkt, 0 bytes mcast: 0 pkt, 0 bytes
137 packets input, 14035 bytes, 0 no buffer
Received 137 broadcasts (0 IP multicasts)
0 runts, 0 giants, 0 throttles
0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
0 watchdog, 0 multicast, 0 pause input
0 input packets with dribble condition detected
9 packets output, 2902 bytes, 0 underruns
0 output errors, 0 collisions, 0 interface resets
0 babbles, 0 late collision, 0 deferred
0 lost carrier, 0 no carrier, 0 PAUSE output
Sup720#
```

This structure has been changed to now include **eight counters** on the Supervisor Engine 2T with per-protocol statistics, due to the implementation of Logical Interfaces (LIFs).

Bridge Domains and Logical Interfaces

Two new concepts are introduced on the Supervisor Engine 2T: Bridge Domains (BDs) and Logical Interfaces (LIFs). Bridge domains are used to represent Layer 2 broadcast domains, including VLANs. Logical interfaces are used to represent all types of Layer 3 interfaces, including (but not limited to) Switched Virtual Interfaces (SVIs), tunnel interfaces, routed ports, subinterfaces and Layer 3 EtherChannel interfaces.

A LIF on the Supervisor Engine 2T essentially maps to a port-VLAN combination, allowing for the capability to provide different services for different customers per interface. Prior to the introduction of Supervisor Engine 2T, VLANs were used internally by the system to represent not only Layer 2 VLANs but also Layer 3 interfaces, and the number of available VLANs was limited to 4K. With the introduction of LIFs on the Supervisor Engine 2T, the number of available routed interfaces has increased to 128K and the number of broadcast domains (equivalent of VLANs) has increased to 16K (4K at first software release). Figures 6 and 7 show the differences between Supervisor Engine 2T, and Supervisor Engines 32 and 720.

Figure 6. VLAN Usage for Supervisor Engines 32 and 720

Figure 7. VLAN Usage for Supervisor Engine 2T

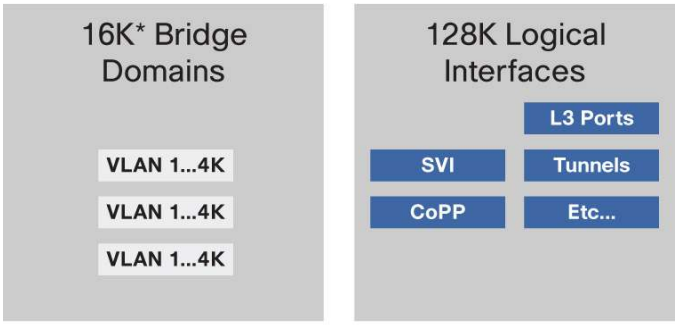

*4K bridge domains will be supported in the first release of software

LIF and Adjacency Statistics

Compared to its Supervisor Engine 720 counterpart where every Layer 3 interface was associated with a VLAN, the Supervisor Engine 2T associates each Layer 3 interface with a unique logical interface value or LIF. So while statistics were previously collected on a per-VLAN basis, statistics are now collected on a per-LIF basis.

There are 128K primary LIFs available on the Supervisor Engine 2T, and the forwarding engine has a component dedicated to capturing LIF statistics. The Policy Feature Card 4 (PFC4) on the Supervisor Engine 2T contains a Layer 2 forwarding engine and a Layer 3 forwarding engine, with additional components as shown in Figure 8.

Figure 8. Block Diagram of PFC4

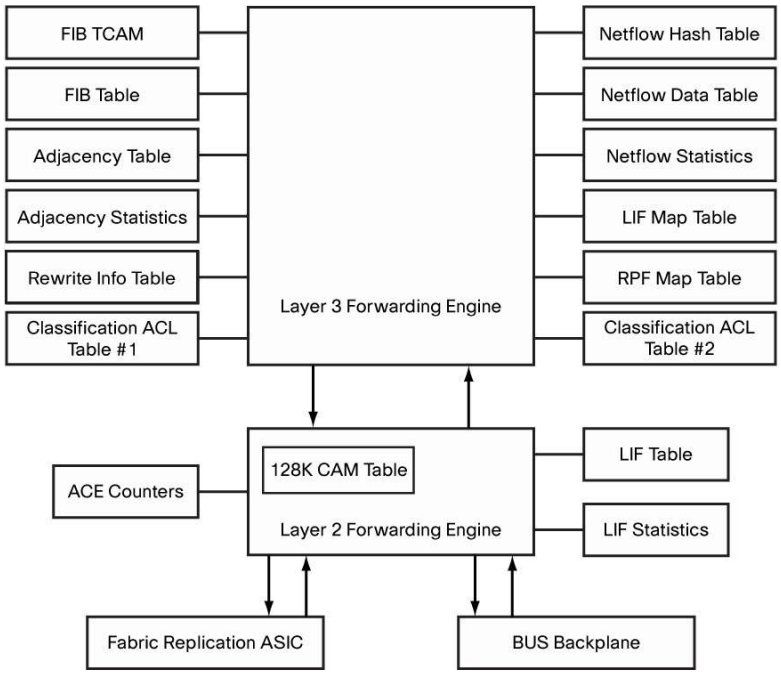

The LIF statistics module is attached to the Layer 2 forwarding engine and is used to collect packet counts for all **Layer 3 hardware switched non-recirculated** packets. This includes statistics for the following:

- SVIs
- Physical routed ports
- Layer 3 port-channels
- Subinterfaces

The support for counters under subinterfaces is a new feature on the Supervisor Engine 2T. It is important to note that LIF statistics are not available for tunnel interfaces due to the above definition regarding non-recirculated packets. Recirculated packets refer to packets that require multiple lookups to be forwarded through the switch (for example, GRE-encapsulated packets). Also, since a loopback interface does not have a LIF associated with it, these statistics are not available for loopback interfaces either.

For each of the primary LIFs, there are now eight pairs of counters (packet and byte) being maintained, for IPv4, IPv6, multicast and others. These counters are explained in detail in the section Per-Protocol Interface Counters. The counters are obtained from the LIF statistics module shown in Figure 8.

For the first 16K LIFs that map to the bridge domains (BDs), there are an additional two pairs of counters (packet and byte) obtained from the adjacency statistics module, attached to the Layer 3 forwarding engine, as shown in Figure 8. This includes counters for Layer 2 bridged packets (one pair of counters for unicast, and a second pair for multicast).

Thus, a Layer 3 interface would be mapped to eight pairs of counters (from the LIF stats) and a VLAN and SVI would be associated with 10 pairs of counters (eight pairs from LIF stats and two pairs from adjacency stats).

The logical flow in collecting LIF statistics is as follows: When a packet arrives on the ingress port, a LIF is associated with the (port, VLAN) combination. This is called the Ingress LIF. The packet then passes through Layer 2 and Layer 3 engine look-ups for the forwarding decisions. For Layer 2 packets, since they are switched within the same VLAN (or BD), the bridged counters will be incremented. For Layer 3 packets, the FIB (Forwarding Information Base) table returns an Adjacency pointer which contains the Egress LIF information. The Egress LIF consists of a VLAN and a destination index (which maps to the egress port). For packets that egress on an SVI, the Egress LIF corresponds to the BD-LIF of the VLAN. After processing by the forwarding engines, the LIF statistics module, which is attached to the Layer 2 forwarding engine, collects the statistics for both the Ingress LIF and Egress LIF. The adjacency statistics module maintains the counters for Layer 2 packets.

Per-Protocol Interface Counters

For each of the primary LIFs, there are now eight pairs of counters being maintained, which can be viewed through CLI outputs. Each pair of counters consists of a packet counter and a byte counter, and the following classes of traffic are now identified by default with the new structure:

- L3 IN IPv4 Unicast
- L3 OUT IPv4 Unicast
- L3 IN IPv4 + IPv6 Multicast
- L3 OUT IPv4 + IPv6 Multicast
- L3 IN IPv6 Unicast
- L3 OUT IPv6 Unicast
- L3 IN MPLS
- L3 OUT MPLS

Figure 9 shows a logical view of the LIF statistics counters. Each pair of counters is represented with 72 bits, comprising a 31-bit packet counter, a 40-bit byte counter and 1-bit parity for data integrity.

Figure 9. Logical View of LIF Statistics Counters

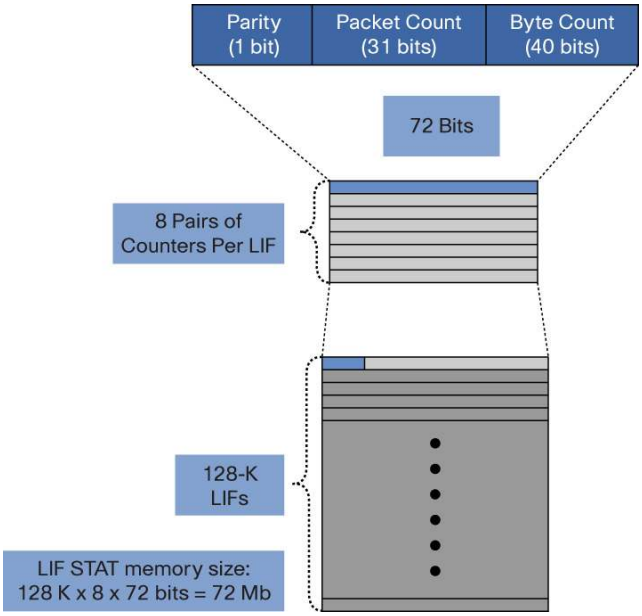

Correspondingly, the interface counters reflect the new fields as can be seen in the subsequent output. Packet and byte counts can be viewed for Layer 2 and Layer 3 switched traffic in both directions. Counters are also available for unicast IPv4 and IPv6, and for multicast traffic (IPv4 and IPv6 combined). Customers who are running a dual-stack IPv4 and IPv6 environment can now monitor these flows individually. In addition, support for Multiprotocol Label Switching (MPLS) counters has been added.

The following CLI outputs are grouped into three sets:

a) Layer 3 interface counters

b) VLAN counters

c) Interface accounting and interface stats

The counters for a Layer 2 switchport remain the same as in previous implementations.

Layer 3 Interface Counters

The following output highlights the new counters available:

```
Sup2T#show interfaces gigabitEthernet 7/1 detail
  GigabitEthernet7/1 is up, line protocol is up (connected)
  Hardware is C6k 1000Mb 802.3, address is 0024.c4dc.d600 (bia 0024.c4dc.d600)
  MTU 1500 bytes, BW 1000000 Kbit, DLY 10 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  Full-duplex, 1000Mb/s, media type is 1000BaseSX
  input flow-control is off, output flow-control is off
  Clock mode is auto
  ARP type: ARPA, ARP Timeout 04:00:00
```

```
  Last input 00:00:31, output never, output hang never
  Last clearing of “show interface” counters 3d19h
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
  L2 Switched: ucast: 0 pkt, 0 bytes - mcast: 0 pkt, 0 bytes
  L3 in Switched: ucast: 0 pkt, 0 bytes - mcast: 0 pkt, 0 bytes
  L3 out Switched: ucast: 0 pkt, 0 bytes - mcast: 0 pkt, 0 bytes
  Additional Counters:
  IPv4 in Switched: ucast: 0 pkt, 0 bytes
  IPv6 in Switched: ucast: 0 pkt, 0 bytes
  IPv4 + IPv6 in Switched: mcast: 0 pkt, 0 bytes
  MPLS in Switched: 0 pkt, 0 bytes
  IPv4 out Switched: ucast: 0 pkt, 0 bytes
  IPv6 out Switched: ucast: 0 pkt, 0 bytes
  IPv4 + IPv6 out Switched: mcast: 0 pkt, 0 bytes
  MPLS out Switched: 0 pkt, 0 bytes
     1091 packets input, 99251 bytes, 0 no buffer
     Received 1091 broadcasts (0 IP multicasts)
     0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
     0 watchdog, 0 multicast, 0 pause input
     0 input packets with dribble condition detected
     1091 packets output, 99311 bytes, 0 underruns
     0 output errors, 0 collisions, 0 interface resets
     0 babbles, 0 late collision, 0 deferred
     0 lost carrier, 0 no carrier, 0 PAUSE output
     0 output buffer failures, 0 output buffers swapped out
Sup2T#
```

VLAN Counters

VLAN counters consist of 10 pairs of counters: eight pairs from LIF statistics that map to the same counters as above, and two pairs from adjacency statistics. In the following outputs, the highlighted counters (Layer 3) are from LIF statistics, and the remaining counters (Layer 2) are from adjacency statistics:

```
Sup2T#show vlan counters
* L2 counters include multicast and broadcast packets

Vlan Id                                    : 1
L2 Unicast Packets                         : 65524
L2 Unicast Octets                          : 31441390
L3 Input Unicast Packets                   : 0
L3 Input Unicast Octets                    : 0
L3 Output Unicast Packets                  : 0
L3 Output Unicast Octets                   : 0
L3 Output Multicast Packets                : 0
L3 Output Multicast Octets                 : 0
L3 Input Multicast Packets                 : 0
```

```
L3 Input Multicast Octets                          : 0
L2 Multicast Packets                               : 0
L2 Multicast Octets                                : 0
<snip>
```

The **show vlan counters detail** output shows all the eight LIF counters:

```
Sup2T# show vlan id 10 counters detail
* L2 counters include multicast and broadcast packets

Vlan Id                                            : 10
L3 Input IPv4 Unicast Packets                      : 0
L3 Input IPv4 Unicast Octets                       : 0
L3 Output IPv4 Unicast Packets                     : 0
L3 Output IPv4 Unicast Octets                      : 0
L3 Input IPv4 + IPv6 Multicast Packets             : 0
L3 Input IPv4 + IPv6 Multicast Octets              : 0
L3 Output IPv4 + IPv6 Multicast Packets            : 0
L3 Output IPv4 + IPv6 Multicast Octets             : 0
L3 Input IPv6 Unicast Packets                      : 0
L3 Input IPv6 Unicast Octets                       : 0
L3 Output IPv6 Unicast Packets                     : 0
L3 Output IPv6 Unicast Octets                      : 0
L3 Input MPLS Packets                              : 0
L3 Input MPLS Octets                               : 0
L3 Output MPLS Packets                             : 0
L3 Output MPLS Octets                              : 0

L2 Known Bridging Unicast Packets                  : 8472535300
L2 Known Bridging Unicast Octets                   : 1084505509929
L2 Known Bridging Multicast Packets                : 0
L2 Known Bridging Multicast Octets                 : 0

Sup2T#
```

Interface Accounting and Interface Stats

Another command that makes use of LIF statistics is the **show interfaces accounting** command. The command can be used to get ingress Layer 3 switched protocol counters for IPv4, IPv6 and MPLS, as the following output shows:

```
Sup2T#show interfaces GigabitEthernet 2/2 accounting
GigabitEthernet2/2
                Protocol    Pkts In   Chars In   Pkts Out  Chars Out
                   Other      75102   24718090      20374    7742120
                      IP          0          0         18       2076
                    IPv6          0          0         18       2076
                 DEC MOP       3913     301301       1018     131322
                     ARP     140311    8418660         53       5936
                     CDP      19949    8148285      11331    4928985
Sup2T#
```

The totals shown above include both hardware switched and software switched packets.

The **show interfaces stats** command includes packet counts that tally with 'show interfaces accounting', with the ingress counters for Distributed cache being pulled from LIF statistics.

```
Sup2T#show interfaces gigabitEthernet 7/1 stats
GigabitEthernet7/1
          Switching path    Pkts In   Chars In   Pkts Out  Chars Out
               Processor       2005     162405       2005     154385
             Route cache          0          0          0          0
       Distributed cache          0          0          0          0
                   Total       2005     162405       2005     154385
Sup2T#
```

Here, the breakdown of the switching paths is as follows:

Processor: Packets that are software processed (non-interrupt path)

Route cache: Packets that take the software interrupt path (CEF switched packets)

Distributed cache: Packets forwarded in hardware

For the Distributed cache statistics, Ingress stats (Pkts In and Chars In) are obtained from LIF statistics and Egress stats (Pkts Out and Chars Out) are obtained from Adjacency statistics.

NetFlow

NetFlow is a process designed to collect information about traffic flows that pass through the switch. The Supervisor Engine 2T is equipped with a PFC4 that supports all the NetFlow features that existed on the PFC3 (Supervisor Engine 720) and adds these new capabilities:

- Sampled NetFlow: Allows users to have NetFlow records created based on a sample of traffic matching the flow. The Supervisor Engine 2T with PFC4 performs NetFlow sampling in hardware, whereas the PFC3-based supervisors performed NetFlow sampling in software.
- Increased support for NetFlow entries: Up to 512K NetFlow entries (no ingress or egress limitation) can now be stored in the PFC4. Up to 1 million NetFlow entries (512K for ingress and 512K for egress) can now be stored in the PFC4XL. Both of these represent a quadrupling of the capabilities of their PFC3-based predecessors.
- Improved NetFlow hash: The hash efficiency is improved to 99 percent, allowing a greater percentage of the NetFlow table to be utilized.
- Egress NetFlow: Provides support for collecting flow statistics for packets after they have had ingress processing applied to them and the final destination of the packets has been determined.
- Layer 2 NetFlow: NetFlow hardware is capable of creating and tracking bridged IP flows as opposed to traditional NetFlow where flows would be created and tracked only for IP traffic that gets routed.
- Flexible NetFlow: This is Cisco's next-generation NetFlow technology, featuring customized traffic identification and simultaneous tracking of multiple NetFlow applications. It supports the NetFlow Version 9 record format, including new fields for IPv6 and multicast information.
- TCP flags: TCP flags (SYN, FIN, RST, ACK, URGENT, PUSH) are now collected as part of a flow record.
- NetFlow with Embedded Event Manager: As described earlier, EEM Version 3.0 also supports a new event detector for Flexible NetFlow.

More details on the NetFlow features available on Supervisor Engine 2T can be found in the Cisco Catalyst 6500 Supervisor Engine 2T: NetFlow Enhancements paper on Cisco.com.

Control-Plane Policing and Hardware Rate Limiters

Control plane policing (CoPP) is a critical feature in ensuring that the CPU of the supervisor engine is protected and regulated from traffic spikes. The Supervisor Engine 2T introduces new hardware rate limiters (HWRL) and CoPP features to further aid in control-plane protection. Detailed information on CoPP and HWRL can be found at the Borderless Networks Security: Cisco Catalyst 6500 Series Control plane Protection Techniques for Maximum Uptime paper on Cisco.com.

From the monitoring point of view, the Supervisor Engine 2T adds packet-based counters to CoPP in addition to byte-based counters. The HWRL have been enhanced with several new rate limiters to make a total of 31 Layer 3 rate limiters and 26 Layer 2 rate limiters on the platform (as compared to eight Layer 3 rate limiters and four Layer 2 rate limiters on PFC3-based systems). Supervisor Engine 2T includes HWRL counters for forwarded, dropped, and leaked packets. Another feature that has been introduced is the ability to apply CoPP on exceptions such as IP-option, TTL-failure, etc. The rate-limiters available are shown in the output below.

```
Sup2T#show platform rate-limit
 State : ON - enabled but not sharing, ON/S - enabled and sharing
 Share : NS - not sharing, G - group, S - static sharing, D - dynamic sharing
       : P/sec - Packets/sec, B/sec - Bytes/second, BP - Burst period (microsec)
   Rate Limiter Type   State    P/sec  P/burst      B/sec    B/burst BP      Share
Leak
 -------------------- ----- -------- -------- ---------- ---------- ------- ------
           CEF RECEIVE   OFF        -        -          -          -       -     -
 CEF RECEIVE SECONDARY    ON    15000        -        -         - 1000000   . NS  OFF
             CEF GLEAN    ON     1000        -        -         - 1000000     NS  OFF
             IP ERRORS    ON     1000      100        -         -     100     NS  OFF
       UCAST IP OPTION    ON   700000        -        -         -    1000 G: 0, S   ON
         ICMP ACL-DROP    ON   700000        -        -         -    1000 G: 0, S   ON
         ICMP NO-ROUTE    ON      100        -        -         - 1000000     NS  OFF
         ICMP REDIRECT   OFF        -        -        -         -       -     -
           RPF FAILURE    ON      100        -        -         - 1000000     NS  OFF
          ACL VACL LOG    ON     2000        -        -         - 1000000     NS  OFF
        ACL BRIDGED IN   OFF        -        -          -          -       -     -
       ACL BRIDGED OUT   OFF        -        -          -          -       -     -
        ARP Inspection   OFF        -        -          -          -       -     -
      DHCP Snooping IN   OFF        -        -          -          -       -     -
           IP FEATURES   OFF        -        -          -          -       -     -
            MAC PBF IN   OFF        -        -          -          -       -     -
           CAPTURE PKT    ON   200000        -        -         - 1000000     NS  OFF
  IP ADMIS. ON L2 PORT   OFF        -        -          -          -       -     -
 MCAST IPV4 DIRECTLY C   OFF        -        -          -          -       -     -
   MCAST IPV4 FIB MISS   OFF        -        -          -          -       -     -
       MCAST IPV4 IGMP   OFF        -        -          -          -       -     -
    MCAST IPV4 OPTIONS   OFF        -        -          -          -       -     -
        MCAST IPV4 PIM   OFF        -        -          -          -       -     -
 MCAST IPV6 DIRECTLY C   OFF        -        -          -          -       -     -
        MCAST IPV6 MLD   OFF        -        -          -          -       -     -
 MCAST IPV6 CONTROL PK   OFF        -        -          -          -       -     -
           MTU FAILURE   OFF        -        -          -          -       -     -
           TTL FAILURE   OFF        -        -          -          -       -     -
 MCAST BRG FLD IP CNTR   OFF        -        -          -          -       -     -
```

```
      MCAST BRG FLD IP   OFF        -        -          -          -       -     -
             MCAST BRG   OFF        -        -          -          -       -     -
         MCAST BRG OMF   OFF        -        -          -          -       -     -
   UCAST UNKNOWN FLOOD   OFF        -        -          -          -       -     -
          LAYER_2 PDU    OFF        -        -          -          -       -     -
           LAYER_2 PT    OFF        -        -          -          -       -     -
       LAYER_2 PORTSEC   OFF        -        -          -          -       -     -
     LAYER_2 SPAN PCAP   OFF        -        -          -          -       -     -
       DIAG RESERVED 0    ON 33554431      -        -        -       1     NS  OFF
       DIAG RESERVED 1    ON 33554431      -        -        -       1     NS  OFF
       DIAG RESERVED 2    ON 33554431      -        -        -       1     NS  OFF
   DIAG RESERVED LIF 0    ON 33554431      -        -        -       1     NS  OFF
   MCAST REPL RESERVED    ON        0      -        -        -       0     NS  OFF
Sup2T#
Sup2T#show platform rate-limit hw-details
  : P/sec - Packets/sec, B/sec - Bytes/second, BP - Burst period
  HW_ID State    P/sec  P/burst      B/sec    B/burst BP      Index
  ----- ----- -------- -------- ---------- ---------- ------- ------
L3 Rate Limiters:
      1    ON     2000        -          -          - 1000000 0x0
      2   OFF        -        -          -          -       -      -
      3   OFF        -        -          -          -       -      -
      4   OFF        -        -          -          -       -      -
      5   OFF        -        -          -          -       -      -
      6   OFF        -        -          -          -       -      -
      7   OFF        -        -          -          -       -      -
      8   OFF        -        -          -          -       -      -
      9    ON     1000      100          -          -     100 0x0
     10    ON   700000        -          -          -    1000 0x0
     11    ON      100        -          -          - 1000000 0x0
     12   OFF        -        -          -          -       -      -
     13   OFF        -        -          -          -       -      -
     14   OFF        -        -          -          -       -      -
     15   OFF        -        -          -          -       -      -
     16   OFF        -        -          -          -       -      -
     17   OFF        -        -          -          -       -      -
     18   OFF        -        -          -          -       -      -
     19   OFF        -        -          -          -       -      -
     20   OFF        -        -          -          -       -      -
     21   OFF        -        -          -          -       -      -
     22   OFF        -        -          -          -       -      -
     23   OFF        -        -          -          -       -      -
     24   OFF        -        -          -          -       -      -
     25   OFF        -        -          -          -       -      -
     26   OFF        -        -          -          -       -      -
     27   OFF        -        -          -          -       -      -
     28   OFF        -        -          -          -       -      -
     29   OFF        -        -          -          -       -      -
     30   OFF        -        -          -          -       -      -
     31   OFF        -        -          -          -       -      -
```

```
L2 Input Rate Limiters:
     32    ON 33554431        -          -          -       1 0x0
     33    ON 33554431        -          -          -       1 0x0
     34    ON 33554431        -          -          -       1 0x0
     35    ON 33554431        -          -          -       1 0x0
     36    ON        0        -          -          -       0 0x0
     37   OFF        -        -          -          -       -      -
     38   OFF        -        -          -          -       -      -
     39   OFF        -        -          -          -       -      -
     40   OFF        -        -          -          -       -      -
     41   OFF        -        -          -          -       -      -
     42   OFF        -        -          -          -       -      -
     43   OFF        -        -          -          -       -      -
     44   OFF        -        -          -          -       -      -
     45   OFF        -        -          -          -       -      -
     46   OFF        -        -          -          -       -      -
     47    ON     1000        -          -          - 1000000 0x0
     48    ON      100        -          -          - 1000000 0x0
     49    ON    15000        -          -          - 1000000 0x0
     50    ON   200000        -          -          - 1000000 0x0
     51   OFF        -        -          -          -       -      -

L2 Output Rate Limiters:
     52   OFF        -        -          -          -       -      -
     53   OFF        -        -          -          -       -      -
     54   OFF        -        -          -          -       -      -
     55   OFF        -        -          -          -       -      -
     56   OFF        -        -          -          -       -      -
     57   OFF        -        -          -          -       -      -
Sup2T#
```

It is also possible to monitor control-plane traffic on a per-flow basis using Flexible NetFlow to develop realistic traffic rates, which can then be used in developing custom control-plane service policies. This is described in detail in the CoPP paper referenced above.

SNMP Support and MIBs

With Cisco IOS Software Release 12.2(50)SY, the Supervisor Engine 2T adds enhancements to several of the existing MIBs, and introduces a few new ones, as described in Table 1.

Table 1. New Platform-Specific MIBs Introduced with Cisco IOS Software Release 12.2(50)SY

MIB	Description
Cisco-Switch-NetFlow-MIB	Provides SNMP access to the Switch NetFlow component
Cisco-Switch-Stats-MIB	Provides SNMP access to the configuration and monitoring of traffic statistics on Cisco's switching devices
Cisco-TrustSec-Interface-MIB	Provides SNMP access to the Cisco TrustSec® component
Cisco-TrustSec-MIB	Provides SNMP access to the Cisco TrustSec component
Cisco-TrustSec-Policy-MIB	Provides SNMP access to the Cisco TrustSec Role-based Access Control (RBAC), Layer 3 transport and policy feature
Cisco-TrustSec-Server-MIB	Provides SNMP access to the Cisco TrustSec component
Cisco-TrustSec-SXP-MIB	Provides SNMP access to the Cisco TrustSec component

Summary

The Cisco Catalyst 6500 Series Supervisor Engine 2T with Cisco IOS Software Release 12.2(50)SY is a platform with an extensive set of monitoring features. The forwarding engine on Supervisor Engine 2T allows for increased capabilities and traffic analysis in areas such as QoS, NetFlow and SNMP. The platform's support for EEM 3.0 puts it on par with router platforms that run Cisco IOS Software, paving the way for a consistent and scalable method to customize device monitoring and automate tasks. Release 12.2(50)SY has also enhanced existing Cisco management innovations such as Cisco Smart Call Home, Generic Online Diagnostics GOLD, and Onboard Failure Logging (OBFL), which can be used to simplify monitoring, fault detection, and fault isolation. Overall, these features can be used to assist in comprehensive capacity planning, application assessment, network troubleshooting and security operations.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**

Printed in USA C11-663635-00 07/11