

CHAPTER 5

Network Telemetry

In order to operate and ensure availability of a network, it is critical to have visibility and awareness into what is occurring on the network at any one time. Network telemetry offers extensive and useful detection capabilities which can be coupled with dedicated analysis systems to collect, trend and correlate observed activity.

Baseline network telemetry is both inexpensive and relatively simple to implement. This section highlights the baseline forms of telemetry recommended for network infrastructure devices, including:

- Time Synchronization
- Local Device Traffic Statistics
- System Status Information
- CDP Best Common Practices
- Syslog
- SNMP
- ACL Logging
- Accounting
- Archive Configuration Change Logger
- Packet Capture

More information on network telemetry and the critical role it plays in security can be found in the whitepaper *How to Build a Cisco Security Operations Center*. This paper provides an overview of the principles behind security operations, along with guidance on how to build a security operations center. The whitepaper is available at the following URL:

http://www.cisco.com/en/US/solutions/collateral/ns341/ns524/ns546/ns310/net_implementation_white _paper0900aecd80598c16.html

CSF Methodology Assessment

The results of applying the CSF methodology for baseline network telemetry are presented in the table below and highlight the technologies and features identified and integrated in Network Security Baseline.

Visibility and Awareness

Identify	Monitor	Correlate
• CDP	• NTP	
• SNMP	• Local device statistics	
• Syslog	• System Status Information	
	- Memory/CPU/processes	
	 CPU and memory threshold notification 	
	CDP Best Common Practices	
	• Logging	
	– Syslog	
	– SNMP	
	– Accounting	
	 Configuration change notification and logging 	
	• Packet capture	
	– SPAN/RSPAN	
	- Copy/capture VACLs	

 Table 5-1
 CSF Methodology Assessment – Visibility and Awareness

Control and Containment

No control and containment features for CSF methodology.

Time Synchronization

When implementing network telemetry, it is important that dates and times are both accurate and synchronized across all network infrastructure devices. Without time synchronization, it is very difficult to correlate different sources of telemetry.

Enabling Network Time Protocol (NTP) is the most common method of time synchronization.

General best common practices for NTP include:

- A common, single time zone is recommended across an entire network infrastructure in order to enable the consistency & synchronization of time across all network devices.
- The time source should be from an authenticated, limited set of authorized NTP servers.

Detailed information on NTP and NTP deployment architectures is available in the *Network Time Protocol: Best Practices White Paper* at the following URL:

http://www.cisco.com/warp/public/126/ntpm.pdf

Timestamps and NTP Configuration

In Cisco IOS, the steps to enable timestamps and NTP include:

- **Step 1** Enable timestamp information for debug messages.
- **Step 2** Enable timestamp information for log messages.
- **Step 3** Define the network-wide time zone.
- **Step 4** Enable summertime adjustments.
- **Step 5** Restrict which devices can communicate with this device as an NTP server.
- **Step 6** Restrict which devices can communicate with this device as an NTP peer.
- **Step 7** Define the source IP address to be used for NTP packets.
- **Step 8** Enable NTP authentication.
- **Step 9** Define the NTP servers.
- **Step 10** Define the NTP peers.
- **Step 11** Enable NTP to update the device hardware clock

The Cisco IOS commands to achieve the above steps are provided below.

Timestamp information for debug messages can be enabled with the following global configuration command:

service timestamps debug datetime localtime show-timezone msec

Timestamp information for log messages can be enabled with the following global configuration command:

service timestamps log datetime localtime show-timezone msec

The network-wide time zone, shown in this example as PST, can be enabled with the following global configuration command:

```
clock timezone EST -5
```

Summertime adjustments, shown this example for PDT, can be enabled with the following global configuration command:

clock summer-time EDT recurring

A list of allowed NTP servers and peers can be enforced with an ACL:

```
access-list 10 remark ACL for NTP Servers and peers
access-list 10 permit <NTPserver1>
access-list 10 permit <NTPserver2>
access-list 10 permit <NTPpeer1>
access-list 10 deny any log
!
ntp access-group peer 10
```

NTP clients can be restricted with an ACL:

access-list 15 remark ACL for NTP clients access-list 15 permit <Client1> access-list 15 permit <Client2> access-list 15 deny any log

```
!
ntp access-group serve-only 15
```

The source IP address to be used for NTP packets can be defined with the following global configuration command:

ntp source <Loopback or OOB interface>

The first step in enabling NTP authentication is to define an MD5 key to be used for NTP transactions: ntp authentication-key <key#> md5 <strong8charkey>

The keys to be accepted for NTP authentication are subsequently defined with the following command: ntp trusted-key <key#>

NTP authentication is enforced with the following global configuration command:

ntp authenticate

NTP is used to update the device hardware clock with the following global configuration command:

```
ntp update-calendar
```

The NTP servers are defined with the following global configuration command:

```
ntp server <NTPserver1>
ntp server <NTPserver2>
```

Any NTP peers are defined with the following global configuration command:

```
ntp peer <NTPpeer1>
ntp peer <NTPpeer2>
```

For more information on configuring NTP, refer to the following URL:

http://www.cisco.com/en/US/docs/ios/netmgmt/configuration/guide/nm_basic_sys_manage_ps6350_T SD_Products_Configuration_Guide_Chapter.html#wp1001170

Local Device Traffic Statistics

Local device statistics are the most basic and ubiquitous form of telemetry available. They provide baseline information such as per-interface throughput and bandwidth statistics, enabled features and global per-protocol traffic statistics.

In Cisco IOS, this information is accessed from the command line interface (CLI). The format of a command output, as well as the command itself and its options, vary by platform. It is important to review and understand these differences. The most commonly used commands can be aliased to enable greater operational ease of use.

Per-Interface Statistics

In Cisco IOS, per-interface statistics are available which include throughput (pps) and bandwidth (bps) information. Per-interface statistics can be accessed with the **show interface** command:

```
Router#show interface gigabitEthernet 4/48
GigabitEthernet4/48 is up, line protocol is up (connected)
Hardware is C6k 1000Mb 802.3, address is 0013.5f21.6c80 (bia 0013.5f21.6c80)
Description: cr17-3845-1 fe0
```

```
Internet address is 10.139.5.8/31
 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
  input flow-control is off, output flow-control is off
  Clock mode is auto
 ARP type: ARPA, ARP Timeout 04:00:00
 Last input 00:00:03, output 00:00:00, output hang never
 Last clearing of "show interface" counters never
 Input queue: 0/75/15005/235 (size/max/drops/flushes); Total output drops: 1
 Queueing strategy: fifo
Output queue :0/40 (size/max)
 5 minute input rate 4751000 bits/sec, 3006 packets/sec
5 minute output rate 4499000 bits/sec, 2755 packets/sec
L2 Switched: ucast: 19841909032 pkt, 3347755205145 bytes - mcast: 96885779 pkt,
5131184435 bytes
L3 in Switched: ucast: 27282638229 pkt, 5095662463006 bytes - mcast: 94 pkt, 5191 bytes
mcast
L3 out Switched: ucast: 43107617667 pkt, 7275264441541 bytes
  47118207406 packets input, 9306459456266 bytes, 0 no buffer
  Received 83653389 broadcasts, 0 runts, 0 giants, 0 throttles
   0 input errors, 0 CRC, 0 frame, 649 overrun, 0 ignored
   0 input packets with dribble condition detected
   43210876182 packets output, 8089398934796 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 output buffer failures, 0 output buffers swapped out
```

For more information on the **show interface** command, refer to the following URL:

http://www.cisco.com/en/US/docs/ios/12_3/interface/command/reference/int_s3g.html

Cisco IOS routers are set by default to use a 5 minute decaying average for interface statistics. Setting the decaying average to one minute provides more granular statistics. The length of time for which data is used to compute load statistics can be changed by using the load-interval interface configuration command.

```
Router(config)# interface <interface-type number>
Router(config)# load-interval 60
```

For more information on the **load-interval interface** command, refer to the following URL:

http://www.cisco.com/en/US/docs/ios/interface/command/reference/ir_11.html#wp1014158

The Cisco IOS **pipe** command and its parsing options may also be used to target specific information in the interface output. For example, to quickly view the one minute input and output rates on an interface:

Router#show interface <interface-type numer> | include 1 minute

¹ minute output rate 119223000 bits/sec, 23936 packets/sec



High input or output rates over a period of a minute or so can be very helpful in detecting anomalous behavior.

Clearing the interface counters is often necessary to see what is occurring in a particular instance. However, ensure useful information is not being discarded prior to doing so. To clear interface counters:

Router#clear counters <interface-type number>

¹ minute input rate 54307000 bits/sec, 17637 packets/sec

Per-Interface IP Feature Information

In Cisco IOS, per-interface feature information provides information about the IP features configured on an interface. In particular, this command is useful to identify the number or name of the ACL being enforced, in order to check the ACL counter hits. Per-interface feature information can be accessed with the **show ip interface** command:

```
Router#show ip interface <interface-type number>
T
Router#show ip interface FastEthernet 2/0
FastEthernet2/0 is up, line protocol is up
  Internet address is 198.133.219.6/24
  Broadcast address is 255.255.255.255
  Address determined by setup command
 MTU is 1500 bytes
 Helper address is not set
  Directed broadcast forwarding is disabled
  Outgoing access list is not set
  Inbound access list is 110
  Proxy ARP is disabled
  Local Proxy ARP is disabled
  Security level is default
  Split horizon is enabled
  ICMP redirects are never sent
  ICMP unreachables are always sent
  ICMP mask replies are never sent
  IP fast switching is enabled
  IP fast switching on the same interface is disabled
  IP Flow switching is disabled
  IP CEF switching is enabled
  IP Feature Fast switching turbo vector
  IP Feature CEF switching turbo vector
  IP multicast fast switching is enabled
  IP multicast distributed fast switching is disabled
  IP route-cache flags are Fast, CEF
  Router Discovery is disabled
  IP output packet accounting is disabled
  IP access violation accounting is disabled
  TCP/IP header compression is disabled
  RTP/IP header compression is disabled
  Policy routing is disabled
  Network address translation is disabled
  BGP Policy Mapping is disabled
  WCCP Redirect outbound is disabled
  WCCP Redirect inbound is disabled
  WCCP Redirect exclude is disabled
Router#
```

The **show ip interface** command also provides per-interface uRPF dropped packet statistics. The Cisco IOS **pipe** command and its parsing options can be used to quickly access this information, as shown below.

```
Router#show ip interface <interface-type number> | include 1 verification !
Router#show ip interface FastEthernet 2/0| include veri
IP verify source reachable-via ANY
794407 verification drops
1874428129 suppressed verification drops
```

For more information on the show ip interface command, refer to the following URL:

http://www.cisco.com/en/US/docs/ios/12_3/interface/command/reference/int_s3g.html#wp1205362

Global IP Traffic Statistics

In Cisco IOS, global IP statistics provide a lot of useful information, including per-protocol counts for ICMP, TCP, UDP, and multicast traffic. Global IP traffic statistics can be accessed with the **show ip traffic** command:

```
Router#show ip traffic
IP statistics:
  Rcvd: 4744853 total, 4650886 local destination
        0 format errors, 0 checksum errors, 0 bad hop count
        0 unknown protocol, 0 not a gateway
        0 security failures, 0 bad options, 0 with options
  Opts: 0 end, 0 nop, 0 basic security, 0 loose source route
        0 timestamp, 0 extended security, 0 record route
        0 stream ID, 0 strict source route, 0 alert, 0 cipso, 0 ump
        0 other
  Frags: 0 reassembled, 0 timeouts, 0 couldn't reassemble
        0 fragmented, 0 fragments, 0 couldn't fragment
  Bcast: 1432237 received, 9 sent
  Mcast: 3156376 received, 3147383 sent
  Sent: 3213086 generated, 284 forwarded
  Drop: 42692 encapsulation failed, 0 unresolved, 0 no adjacency
        0 no route, 0 unicast RPF, 0 forced drop
        0 options denied
  Drop: 0 packets with source IP address zero
 Drop: 0 packets with internal loop back IP address
ARP statistics:
 Rcvd: 1419832 requests, 4643 replies, 300822 reverse, 0 other
 Sent: 1057 requests, 4897 replies (0 proxy), 0 reverse
```

This command is very useful for general troubleshooting, as well as for detecting anomalies.

The **show ip traffic** command also provides global uRPF dropped packet statistics. The Cisco IOS pipe command and its parsing options may be used to quickly access this information, as shown below.

```
Router#show ip traffic | include RPF
0 no route, 124780722 unicast RPF, 0 forced drop
```

For more information on the show ip traffic command, refer to the following URL:

http://www.cisco.com/en/US/docs/ios/12_3/ipaddr/command/reference/ip1_s2g.html#wp1081111

System Status Information

Memory, CPU and Processes

A basic indication of a potential issue on a network infrastructure device is high CPU.

In Cisco IOS, information about CPU utilization over a 5-second, 1-minute, and 5-minute window is available with the command:

Router#show processes cpu

The Cisco IOS **pipe** command and its parsing options may be used to exclude information which is not consuming any CPU.

```
Router#show processes cpu | exclude 0.00%_0.00%_0.00%
CPU utilization for five seconds: 38%/26%; one minute: 40%; five minutes: 43%
```

PID	Runtime(ms)	Invoked	uSecs	5Sec	1Min	5Min	TTY	Process
5	192962596	13452649	14343	0.00%	0.52%	0.44%	0	Check heaps
15	4227662201	L540855414	274	0.65%	0.50%	0.49%	0	ARP Input
26	2629012683	3680473726	71	0.24%	0.29%	0.36%	0	Net Background
50	9564564	11374799	840	0.08%	0.07%	0.08%	0	Compute load avg
51	15291660	947844	16133	0.00%	0.03%	0.00%	0	Per-minute Jobs
58	15336356	92241638	166	0.08%	0.02%	0.00%	0	esw_vlan_stat_pr
67	10760516	506893631	21	0.00%	0.01%	0.00%	0	Spanning Tree
68	31804659682	2556402094	1244	7.02%	7.04%	7.75%	0	IP Input
69	25488912	65260648	390	0.00%	0.03%	0.00%	0	CDP Protocol
73	16425564	11367610	1444	0.08%	0.02%	0.00%	0	QOS Stats Export
81	12460616	1020497	12210	0.00%	0.02%	0.00%	0	Adj Manager
82	442430400	87286325	5068	0.65%	0.73%	0.74%	0	CEF process
83	68812944	11509863	5978	0.00%	0.09%	0.11%	0	IPC LC Message H
95	54354632	98373054	552	0.16%	0.12%	0.13%	0	DHCPD Receive
96	61891604	58317134	1061	1.47%	0.00%	4.43%	0	Feature Manager

High CPU utilization values for the IP Input process is a good indicator that traffic ingressing or egressing the device is contributing meaningfully to CPU load. The amount of process-driven traffic versus interrupt-driven traffic is also important.

Understanding the network devices deployed in your network and their normal status is key to establishing a baseline, from which anomalies may be detected.

For more information on the **show proc cpu** command, refer to the following URL:

http://www.cisco.com/en/US/docs/ios/12_3/configfun/command/reference/cfr_1g09.html#wp1042641

Memory Threshold Notifications by Syslog

Cisco IOS offers the ability to send a notification upon memory thresholds being exceeded. A syslog message is sent when memory utilization falls below a configurable low watermark and when free memory once again reaches the configured threshold. Low watermarks can be defined for both processor and input/output (I/O) memory with the following commands:

Router(config) #memory free low-watermark processor <kilobytes threshold> Router(config) #memory free low-watermark io <kilobytes threshold>

Once the thresholds are configured, the router will issue a notification every time the available free memory falls below the specified threshold, and every time the available free memory rises to 5 percent above the specified threshold.

Example output when available free processor memory less than the specified threshold:

000029: *Aug 12 22:31:19.559: %SYS-4-FREEMEMLOW: Free Memory has dropped below 2000k Pool: Processor Free: 66814056 freemem_lwm: 204800000

Example output when available free processor memory recovered to more than the specified threshold 000032: *Aug 12 22:33:29.411: %SYS-5-FREEMEMRECOVER: Free Memory has recovered 2000k Pool: Processor Free: 66813960 freemem_lwm: 0

For more information on the Memory Threshold Notifications feature, refer to the following URL:

http://www.cisco.com/en/US/docs/ios/12_2s/feature/guide/fs_memnt.html

Reserving Memory for Critical Notifications

Cisco IOS offers the ability to preserve critical system logging when a device is overloaded and system resources are low. This feature reserves a region of memory on the device which is only available for critical system logging. Critical system logging memory reservation is enabled with the following command:

```
Router(config) # memory reserve critical <kilobytes>
```

```
<u>Note</u>
```

The amount of memory reserved for critical notifications cannot exceed 25 percent of total available memory.

For more information on the **memory reserve critical** command, refer to the following URL:

http://www.cisco.com/en/US/docs/ios/12_2s/feature/guide/fs_memnt.html#wp1057054

CPU Threshold SNMP Trap Notification

Cisco IOS offers the ability to send a notification upon CPU thresholds being exceeded. An SNMP trap can be sent when CPU utilization exceeds a configurable high-water mark, and when CPU utilization falls below a configurable low-water mark, within a configurable window.

Sudden increases in CPU load on routers and switches often indicate that some event is taking place, however high CPU is not always an indicator of malicious activity. Therefore, analysis and correlation of other event information is highly recommended.

CPU threshold notification can be defined for:

- Total CPU utilization
- CPU process utilization
- CPU interrupt utilization

CPU Thresholding Notification can be configured with the **process cpu threshold** command:

Router(config)# process cpu threshold type {total | process | interrupt} rising <percentage> interval <seconds> [falling <percentage> interval <seconds>] For example, to send an SNMP trap upon total CPU utilization exceeded 80% for more than 5 seconds and being below 20% for more than 5 seconds:

Router(config)# snmp-server enable traps cpu threshold Router(config)# snmp-server host 172.26.150.206 traps public cpu Router(config)# process cpu threshold type total rising 80 interval 5 falling 20 interval 5

It is also a good practice to set the process entry limit and the size of the history table for CPU utilization statistics with the **process cpu statistics** command:

Router(config)# process cpu statistics limit entry-percentage number [size seconds]

The following example shows how to set an entry limit at 40 percent and a size of 300 seconds:

Router(config) # process cpu statistics limit entry-percentage 40 size 300

For more information on the process cpu threshold type command, refer to the following URL:

http://www.cisco.com/en/US/docs/ios/fundamentals/command/reference/cf_m1.html#wp1011517

For more information on the **process cpu statistics limit entry-percentage** command, refer to the following URL:

http://www.cisco.com/en/US/docs/ios/fundamentals/command/reference/cf_m1.html#wp1011684

MAC Address Table Status

Information on which MAC addresses are currently connected to which port can be useful for traceback upon anomalous behavior being detected.

Cisco IOS offers the ability to view the status of the MAC address table on Catalyst Switches using the following command:

```
Router#show mac-address-table
```

The output includes information on currently stored MAC addresses, if they are static or dynamically learnt, their age, and associated VLAN and interface.

This command can be used to trace a specific MAC address, as shown below for the MAC address 0100.5e00.0128 on a Catalyst 6500 with a Supervisor Engine 720:

```
Router# show mac-address-table address 0100.5e00.0128
Legend: * - primary entry
      age - seconds since last seen
      n/a - not available
 vlan mac address type learn
                                 age
                                                ports
_____+
                                                    _____
Supervisor:
   44 0100.5e00.0128 static Yes

    Fa6/44, Router

   1 0100.5e00.0128 static Yes
                                         Router
Module 9:
  44 0100.5e00.0128 static Yes
                                         Fa6/44,Router
                                     _
   1 0100.5e00.0128 static Yes
                                     _
                                         Router
```

The number of MAC addresses currently stored in the MAC address table and the amount of space remaining can be viewed with the mac-address-table count command. An sample output for a particular slot on a Catalyst 6500 is shown below.

```
Router# show mac-address-table count slot 1
MAC Entries on slot 1 :
Dynamic Address Count: 4
Static Address (User-defined) Count: 25
Total MAC Addresses In Use: 29
Total MAC Addresses Available: 131072
```

MAC address table entries will age out according to the configured aging time. Dynamically learnt MAC addresses can be cleared with the following command:

clear mac-address-table dynamic

For more information on the show mac-address-table command, refer to the following URL:

http://www.cisco.com/en/US/docs/ios/12_3/switch/command/reference/swi_s3.html#wp1074753

Open Ports and Sockets

The ports and sockets open on a device should be reviewed to ensure that unused or unneccesary ports and sockets are disabled. In Cisco IOS, open ports and sockets can be viewed with the show control-plane host open-ports command:

Router# show control-plane host open-ports							
Active in	nternet connections	(servers and established)					
Prot	Local Address	Foreign Address		Service	State		
tcp	*:22	*:0	SSE	I-Server	LISTEN		
tcp	*:23	*:0		Telnet	LISTEN		
tcp	*:63771	172.26.150.206:49	IOS host	service	ESTABLIS		
udp	*:49	172.26.150.206:0	TACACS	service	LISTEN		
udp	*:67	*:0	DHCPD	Receive	LISTEN		

Router#

The show control-plane host open-ports command was introduced in 12.3(4)T. For earlier versions of Cisco IOS follow the steps below:

To check open UDP ports use the **show ip sockets** command:

Router#show ip sockets

A sample output is shown below:

Router#show ip sockets

Prot	o Remote	Port	Local	Port	In	Out	Stat	TTY	OutputIF
17	0.0.0.0	0	198.133.219.6	67	0	0	2211	0	
88	listen		any	100	0	0	0	0	
17	172.26.150.206	49	172.26.159.165	49	0	0	21	0	
17	listen		any	161	0	0	1	0	
17	listen		any	162	0	0	11	0	
17	listen		any	56443	0	0	1	0	
17	172.26.150.206	514	172.26.159.165	55759	0	0	210	0	
Rout	er#								

In earlier Cisco IOS versions, open TCP ports can be viewed in two steps with the following commands:

Router#show tcp brief all Router#show tcp tcb

Use the show tcp brief all command to see the IP source and destination IP addresses and state of TCP sessions. This command also provides the transmission control block (TCB), an internal identifier used by the router/switch to identify the connection. The TCB values are then used to identify the ports associated with the connections.

Router#show tcp brief all

Local Address TCB 661BB46C 172.26.159.165.49128 6612A398 198.133.219.6.179 20465FC8 172.26.159.165.22 50711308 198.133.219.6.16422 661B9248 172.26.159.165.19110 6612ACC4 *.179 661294C0 *.179 Router# Use the **show tcp tcb** command to identify the source and destinations sockets for a given TCP session:

Foreign Address (state) 172.26.150.206.49 ESTAB 198.133.219.10.11003 172.26.159.164.15774 198.133.219.5.179 172.26.150.206.49 198.133.219.5.* 198.133.219.10.*

ESTAB ESTAB ESTAB CLOSEWAIT LISTEN LISTEN

Router#show tcp tcb 20465FC8

Connection state is ESTAB, I/O status: 1, unread input bytes: 0 Connection is ECN Disabled, Mininum incoming TTL 0, Outgoing TTL 255 Local host: 172.26.159.165, Local port: 22 Foreign host: 172.26.159.164, Foreign port: 15774

```
Connection tableid (VRF): 0 ...
```

CDP Best Common Practices

Cisco Discovery Protocol (CDP) is a device discovery protocol that runs over Layer 2 and allows Cisco devices to exchange information with each another, including adjacency information. This can be useful for traceback in the case of a network incident.

CDP information is transferred in clear text and communication is unauthenticated. Consequently, CDP is vulnerable to abuse, including sniffing to learn sensitive information about a network infrastructure device, such as IP address, software version, router model, and network topology.

CDP is, however, very useful to OPSEC (operational security) personnel for traceback during an incident, enabling hop-by-hop investigation.

The best common practice for CDP is currently:

- Enable CDP on point-to-point infrastructure inks
- Disable CDP on edge devices or interfaces where is it not required and where such as service may represent a risk, including:
 - LAN access edge
 - Internet transit edge
 - Extranet edge
 - Any public-facing interfaces

It should be noted that CDP is required for some Cisco applications, products and features, such as Cisco IP Telephony for network management.

In the rare case CDP is not used for troubleshooting or security operations, the service can be globally disabled using the **no cdp run** command:

Router(config) #no cdp run

The CDP service may be disabled on a per- interface basis using the **no cdp enable** interface command:

```
Router(config)# interface <interface-type number>
Router(config-if)# no cdp enable
```

For more information on the **cdp run and cdp enable** command, refer to the following URLs:

http://www.cisco.com/en/US/docs/ios/12_3/configfun/command/reference/cfr_1g01.html#wp1032125 http://www.cisco.com/en/US/docs/ios/12_3/configfun/command/reference/cfr_1g01.html#wp1031940

CDP Neighbor Information

To view the detailed information CDP discovers about neighboring devices, use the following command: Router#show cdp neighbors

A sample output is shown below:

```
cr18-7301-1#sh cdp neighbors
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
S - Switch, H - Host, I - IGMP, r - Repeater
```

Device ID	Local Intrfce	Holdtme	Capability	Platform	Port ID
cr17-2821-1	Gig 0/2	154	RSI	2821	Gig 0/0
cr18-6500-2.cisc	co.com				
	Gig 0/1	150	RSI	WS-C6506	Gig 2/1
cr18-6500-1.cisc	co.com				
	Gig 0/0	123	RSI	WS-C6506-	EGig 2/1
cr18 = 7301 = 1 #					

For more information on the show cdp neighbors command, refer to the following URL:

http://www.cisco.com/en/US/docs/ios/12_3/configfun/command/reference/cfr_1g07.html#wp1032872

Syslog

Syslog is a UDP-based logging facility enabling detailed messages to be sent from a device to a syslog server. Syslog packets are reactively sent based on the occurrence of specific events on a device and provide invaluable operational information, including system status, traffic statistics and device access information.

Syslog data is critical to recording day-to-day event and debugging information, as well as notifying operational staff of critical system alerts. Cross-network data aggregation to a central syslog server enables detailed and behavioral analysis of the data which is key to incident handling and attack forensics, as well as general network visibility and routine troubleshooting.

Due to is invaluable role in cross-network data aggregation, syslog should be enabled not just on routers, firewalls and switches, but also on hosts and applications, such as DNS, TACACS+ and RADIUS.

Syslog Best Common Practices

One of the challenges with syslog is that the amount of syslog data can create significant load on both the device sending the data and the syslog server. Cisco IOS, by default, sets syslog logging to level 6 (informational), which can generate a lot of output and can possibly impact the device CPU. Consequently, it is critical to ensure that the following considerations are taken into account:

- Log syslog messages to a central server
- Ensure the syslog server has adequate storage and processing capacity
- Syslog rate-limiting is recommended, where available
- Selectively enabled more detailed logging on critical systems or systems exposed to external users
- Use facility numbers to enable the syslog output to be more easily organized on the syslog server, by default, Cisco routers export syslog as facility 'local7'
- Do not log to the console on Cisco IOS devices
- Define the source IP address to be used for syslog messages. This will typically be a loopback or an OOB interface to ensure consistency.
- Ensure timestamps are enabled per the guidelines in the section Timestamps and NTP
- Ensure syslog messages are stored in a searchable database
- The syslog server should be properly hardened and cryptographic techniques should be considered to protect the logs

• Syslog communication is not authenticated and data is not encrypted. Consequently, syslog packets are susceptible to sniffing. If the confidentiality, integrity and reliability of syslog messages is a concern, an IPSec tunnel/connection from devices to the syslog servers may be considered.

Syslog to a Central Server

In Cisco IOS, syslog messages can be sent to a central server by defining the source IP address to be used, the syslog server(s) to which to send the messages, and the logging trap level.

Router(config)#logging source-interface <interface>
Router(config)#logging host <syslogserver>
Router(config)#logging trap <level>

Numeric Severity Level	Equivalent Word	Description	
0	emergencies	System unusable	
1	alerts	Immediate action needed	
2	critical	Critical conditions	
3	errors	Error conditions	
4	warnings	Warning conditions	
5	notifications	Normal but significant condition	
6	informational	Informational messages only	
7	debugging	g Debugging messages	

Table 5-2 Error Message Severity Levels, Equivalent Text, and Descriptions

Cisco Security Monitoring, Analysis, and Response System (MARS) can be used to collect, analyze and correlate event information generated from a diverse se of network devices and host applications, from Cisco and other vendors. MARS security monitoring can be used in combination of other forms of telemetry, helping reduce false positives and improving threat identification and mitigation responses.

For more information on the logging source-interface command, refer to the following URL:

http://www.cisco.com/en/US/docs/ios/netmgmt/command/reference/nm_07.html#wp1015076

For more information on the logging host command, refer to the following URL:

http://www.cisco.com/en/US/docs/ios/netmgmt/command/reference/nm_07.html#wp1020426

For more information on the logging trap command, refer to the following URL:

http://www.cisco.com/en/US/docs/ios/netmgmt/command/reference/nm_07.html#wp1015177

Syslog Named Facilities

The use of multiple named syslog facilities can be used to logically separate and store syslog messages. Each facility name can be stored and processed in different directory location on the syslog server. For instance, a logging facility directory structure may be implemented based on the role of a device, such as core router, internal edge device, external edge device, etc. Each individual facility may then be processed, reviewed and acted upon according to operational needs.

In Cisco IOS, syslog named facilities are configured using the following command:

Router(config)#logging facility <facility>

For more information on the logging facility command, refer to the following URL:

http://www.cisco.com/en/US/docs/ios/netmgmt/command/reference/nm_07.html#wp1012726

Syslog Rate-Limiting

Syslog rate-limiting, if available, is recommended to ensure that syslog messages do not impact the CPU of either the sending device or the syslog server. As the name indicates, syslog rate-limiting limits the rate of messages logged per second.

In Cisco IOS versions supporting syslog rate-limiting, syslog is limited to 10 messages per second by default. This default behavior can be changed with the logging rate-limit command. The logging rate-limit command allows you set a different rate limit for all severity levels, or for severity levels above an specified value.

In the following example, messages of any severity level are limited to 5 per second:

Router(config) #logging rate-limit 5

In this example, messages at level 3 and above are limited to one message per second, while syslog messages at levels 0-2 (emergencies, alerts and critical) are not rate-limited.

Router(config) #logging rate-limit 1 except 2

For more information on the **logging rate-limit** command and syslog facilities, refer to the following URL:

http://www.cisco.com/en/US/docs/ios/netmgmt/command/reference/nm_07.html#wp1014866

Common Syslog Servers

Some syslog servers in common use include:

• Simple Event Correlator (SEC)

A powerful open source tool which enables parsing and correlation of syslog output. It requires some scripting skills to use but is quite useful.

Cisco CS-MARS

Takes syslog input from routers, switches, firewalls, IDS, VPN concentrators and combines it with other forms of telemetry in order to provide anomaly-detection and event correlation. It can be configured to accept syslog output from many different types of network devices, servers, etc.

• Sawmill

A commercial Web-based tool used to analyze many different kinds of syslog output and render HTML reports and graphs.

Kiwi Syslog Daemon: http://www.kiwisyslog.com/kiwi-syslog-daemon-overview/

Open source databases such as MySQL and PostGres are often used to store syslog output for post-processing and searching.

SNMP

Simple Network Management Protocol (SNMP) is a popular network management protocol that provides a range of information useful for network telemetry.

A general overview of SNMP, including information on the different versions and their associated security mechanisms, as well as general guidelines, can be found in the section SNMP Access.

Common SNMP Servers

Some tools and management platforms in common use which leverage SNMP telemetry include:

- Open source tools including NET-SNMP, MRTG, Cricket, RRDTool, Nagios
- Cisco MARS and Arbor PeakFlow DoS which combine SNMP polling of interface speeds with NetFlow telemetry in order to perform traffic profiling and anomaly-detection
- NMS platforms such as HP OpenView and Nagios Often support alerting based upon low-water-marks as well as high-water-marks

ACL Logging

ACL logging can be used to log basic information related to:

- · Successful access attempts from authorised communicators
- Failed access attempts from non-authorised communicators

The use of an ACL with the 'log' keyword provides information on the IP addresses and port numbers associated with a packet. The use of an extended ACL with the 'log-input' keyword extends this information to also include the input interface and source MAC or VC. Consequently, the 'log-input' keyword should be used whenever extended ACLs are available.

It is recommended that ACL logging is only enabled on an ACL that will not normally match a very large number of packets. This ensures the log remains manageable and does not overflow the restricted log file size. ACL log messages are rate-limited and so any anomalies in this expected behavior will not severely impact the device availability.

- Router Access Control List (ACL) logging is very common; however, it can generate a lot of information in a short span of time, and have a negative impact on performance. ACL logging can be rate-limited, but there are other forms of telemetry such as NetFlow which can provide more information about network traffic
- The Cisco ASA firewall appliance and the Firewall Services Module (FWSM) for the 6500/7600 also generate ACL logs; the performance impact of ACL logging on these devices is much less than on routers

Accounting

Accounting is a critical element of network telemetry and is covered in Infrastructure Device Management Access Logging, page 2-25.

Configuration Change Notification and Logging

Cisco IOS offers a Configuration Change Notification and Logging (Config Log Archive) feature which tracks commands executed in configuration mode through the CLI or HTTP. This feature is covered in Infrastructure Device Management Access Logging, page 2-25.

Packet Capture

Packet capture is generally undertaken after a macro-level indication of an anomaly, for instance via SNMP or syslog, in order to enable more detailed analysis.

General guidelines include:

- Packet capture should take place at key points in the topology such as distribution gateways, IDC switch meshes, desktop access switch meshes, and in some cases, the core.
- It is important to be as specific as possible when capturing packets; at high rates of speed, the amount of information can be overwhelming.
- It is extremely important to ensure that traffic is captured bidirectionally, or, if this is not possible, to identify the unidirectional nature of the capture and take this into account when analyzing the captured traffic.
- Conversely, it is important to avoid capturing duplicate traffic, especially in complex topologies.

SPAN/RSPAN

Switchport Analysis and Remote Switchport Analysis (SPAN/RSPAN) are Cisco IOS features that enable packets to be passed to traffic analysis systems.

SPAN sessions allow you to monitor traffic on one or more ports, or one or more VLANs, and send the monitored traffic to one or more destination ports.

SPAN/RSPAN does not a have measurable performance impact on the network device and does not affect the switching of traffic on source ports or VLANs. SPAN sends a copy of the packets received or transmitted by the source ports and VLANs to the destination port. The destination port must be dedicated for use by SPAN.

For more information on SPAN on Catalyst 6500 Series Switches, see the configuration guide at the following URL:

http://www.cisco.com/en/US/docs/switches/lan/catalyst6500/ios/12.2SXF/native/configuration/guide/s pan.html

Copy/Capture VLAN ACLs

The Cisco IOS VLAN ACL (VACL) feature may be leveraged to enable packets to be passed to traffic analysis systems. The VACL action clause "forward capture" is used to allow packets to continue to flow but to copy them to ports configured as "capture ports"; a port configured to capture VACL-filtered traffic.

An example of how to define and apply a copy/capture VACL access map to forward and capture all IP traffic matching net_10 is shown below.

```
Router(config) # vlan access-map capture1 10
```

```
Router(config-access-map)# match ip address net_10
Router(config-access-map)# action forward capture
Router(config-access-map)# exit
Router(config)# vlan filter capture1 vlan-list 2, 4-6
```

An interface can be configured as a capture port with the following command:

```
Router(config)# interface interface
Router(config-if)# switchport capture
```

The capture action sets the capture bit for the forwarded packets so that ports with the capture function enabled can receive the packets. Only forwarded packets can be captured.

A capture port supports only egress traffic. No traffic can enter the switch through a capture port. VACLs do not a have measurable performance impact on the network device.

For more information on VACLs on Catalyst 6500 Series Switches, refer to the following URL:

http://www.cisco.com/en/US/docs/switches/lan/catalyst6500/ios/12.2SXF/native/configuration/guide/v acl.html

The Cisco NAM-2 captures packets via SPAN/RSPAN or copy/capture VACLs on the 6500/7600. It can perform basic on-board analysis but captures are typically saved and downloaded for use by a dedicated traffic analysis device, such as Ethereal or Network General Sniffer.

General Network Telemetry Indicators

- CPU
 - Spikes in CPU load on network infrastructure devices are often an indication that an event is taking place.
 - High CPU is not always an indicator of malicious activity. Trending is important.
 - Correlating CPU utilization with other information such as network traffic statistics, routing-table changes, etc. is very useful.
 - A baseline of CPU utilization over time is a good idea from a network management standpoint, and also allows operational staff to determine if further investigation is warranted.
- Traffic Rates
 - Excessive bandwidth (bps) and/or throughput (pps) can be an indicator of undesirable traffic.
 - DoS attacks and DoS-like worms can cause high amounts of traffic.
 - High amounts of traffic are not always indicative of problems, since it is dependent on the situation at any particular time. This is another reason why a telemetry baseline is important.
 - Most Network Management Systems (NMS) can alert on high bps and/or high pps. It can be just as important to be notified of a drop in link speed (or number of routes, etc.) as it is to learn of an increase.
- Link Flaps
 - Link-flaps can indicate that something is amiss.
 - Link flaps are often a sign of miss-configuration, backhole and other similar incidents. However, link flaps can also result from malicious activity, such as a DoS attack.
 - Routers and switches can be configured to notify monitoring systems when a link flap occurs.
 - Link flaps are not always indicative of problems, since it is dependent on the situation at any particular time. This is another reason why a telemetry baseline is important.