

Troubleshoot TCP Slowness Issues due to MSS Adjustment in Catalyst 9K Switches

Contents

[Introduction](#)

[Information about TCP MSS Adjustment](#)

[Behaviour](#)

[Topology](#)

[Scenario](#)

[Initial Configuration & Behaviour](#)

[Behaviour After TCP MSS Adjustment](#)

[TCP MSS Adjustment Causing Slowness during Enormous Amount of TCP Traffic](#)

[Important Points](#)

Introduction

This document describes how a Catalyst 9K Switch does the TCP MSS adjustment, and how TCP slowness is linked to this feature.

Information about TCP MSS Adjustment

The Transmission Control Protocol (TCP) Maximum Segment Size (MSS) Adjustment feature enables the configuration of the maximum segment size for transient packets that traverse a router, specifically TCP segments with the SYN bit set. The `ip tcp adjust-mss` command is used in interface configuration mode to specify the MSS value on the intermediate router of the SYN packets in order to avoid truncation.

When a host (usually a PC) initiates a TCP session with a server, it negotiates the IP segment size by using the MSS option field in the TCP SYN packet. The MTU configuration on the host determines the value of the MSS field. The default MTU value for a NIC of the PC is 1500 bytes with a TCP MSS value of 1460 (1500 bytes - 20 bytes IP header - 20 bytes TCP header).


The PPP over Ethernet (PPPoE) standard supports an MTU of only 1492 bytes.

The disparity between the host and PPPoE MTU size can cause the router in between the host and the server to drop 1500-byte packets and terminate TCP sessions over the PPPoE network.

Even if the path MTU (which detects the correct MTU across the path) is enabled on the host, sessions can be dropped because system administrators sometimes disable the Internet Control Message Protocol (ICMP) error messages that must be relayed from the host for path MTU to work.

The `ip tcp adjust-mss` command helps prevent TCP sessions from being dropped by adjusting the MSS value of the TCP SYN packets. The `ip tcp adjust-mss` command is effective only for TCP connections passing through the router. In most cases, the optimum value for the `max-segment-size` argument of the `ip tcp adjust-mss` command is 1452 bytes.

This value plus the 20-byte IP header, the 20-byte TCP header, and the 8-byte PPPoE header add up to a 1500-byte packet that matches the MTU size for the Ethernet link.



Note: TCP MSS adjustment-based traffic is software switched in Catalyst 9K Switches. This document explains scenarios which assume that the TCP MSS adjustment-based traffic is software-switched. Refer to the Configuration Guide in order to confirm if a specific HW/SW software switches the TCP MSS adjustment-based traffic.

Behaviour

As mentioned earlier, TCP MSS adjustment-based traffic is always software-switched. This means that if you try and perform TCP adjustment, then the Switch sends the TCP traffic to the CPU for the MSS modification.

For example, if you modify the TCP MSS value on an interface, then all TCP traffic being received on that interface gets punted to the CPU. The CPU then changes the MSS value and sends the traffic out to the required interface where that TCP packet was headed.

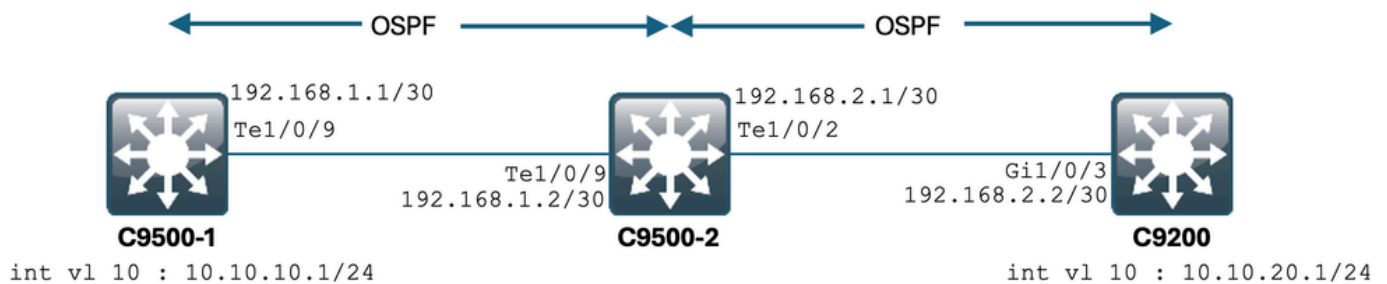
Due to this reason, if there is an enormous amount of TCP traffic with MSS adjustment, then this overloads the CPU Queue.

When a CPU Queue is overloaded, the Control Plane Policer (COPP) polices that traffic and drops packets to maintain the Queue Policer rate. This causes the TCP packets to get dropped.

Hence issues such as File Transfer Slowness, SSH session creations, and Citrix Application Slowness (if using TCP) are seen.

A real-life example of how this happens is shown here.

Topology



Scenario

You are going to SSH into the C9200 from the C9500-1.
SSH using C9500-1's VLAN 10 (10.10.10.1) as the Source.
The destination of the SSH is C9200's VLAN 20 (10.10.20.1/24).
SSH is TCP-based, hence any slowness in TCP also affects this SSH session creation.

There is a transit L3 Switch (C9500-2) between C9500-1 and C9200.
There are two transit /30 L3 links, one between C9500-1 & C9500-2, and one between C9500-2 & C9200.
OSPF is used for reachability across all three Switches, and all/30 subnets and SVIs are advertised into OSPF.
All the IPs shown earlier are reachable between themselves.

In C9500-2 Te1/0/9, the modification of the TCP MSS value is done.
When the SSH from the C9500-1 is initiated, a TCP 3-Way Handshake occurs.
The SYN packet hits the C9500-2 Te1/0/9 (Ingress), where the TCP MSS adjustment is performed.

Initial Configuration & Behaviour

An EPC capture on C9500-2 Te1/0/9 (both directions) was taken and SSH from C9500-1 to C9200 was started.

Here is the EPC configuration:

```
C9500-2#show monitor capture mycap
Status Information for Capture mycap
Target Type:
Interface: TenGigabitEthernet1/0/9, Direction: BOTH
Status : Inactive
Filter Details:
Capture all packets
```

Buffer Details:
Buffer Type: LINEAR (default)
Buffer Size (in MB): 80
File Details:
File not associated
Limit Details:
Number of Packets to capture: 0 (no limit)
Packet Capture duration: 0 (no limit)
Packet Size to capture: 0 (no limit)
Maximum number of packets to capture per second: 1000
Packet sampling rate: 0 (no sampling)
C9500-2#

Starting the EPC:

```
C9500-2#monitor capture mycap start
Started capture point : mycap
C9500-2#
```

Starting the SSH from C9500-1 to C9200:

```
C9500-1#ssh -l admin 10.10.20.1
Password:
```

Stopping the EPC:

```
C9500-2#monitor capture mycap stop
Capture statistics collected at software:
Capture duration - 6 seconds
Packets received - 47
Packets dropped - 0
Packets oversized - 0
Bytes dropped in asic - 0
Capture buffer will exists till exported or cleared
Stopped capture point : mycap
C9500-2#
```

And here are the EPC captured packets:

```
C9500-2#show monitor capture mycap buffer brief
Starting the packet display ..... Press Ctrl + Shift + 6 to exit
1 0.000000 10.10.10.1 -> 10.10.20.1 TCP 60 44274 -> 22 [SYN] Seq=0 Win=4128 Len=0 MSS=536
2 0.001307 10.10.20.1 -> 10.10.10.1 TCP 60 22 -> 44274 [SYN, ACK] Seq=0 Ack=1 Win=4128 Len=0 MSS=536
3 0.001564 10.10.10.1 -> 10.10.20.1 TCP 60 44274 -> 22 [ACK] Seq=1 Ack=1 Win=4128 Len=0
4 0.003099 10.10.20.1 -> 10.10.10.1 SSH 73 Server: Protocol (SSH-2.0-Cisco-1.25)
5 0.003341 10.10.10.1 -> 10.10.20.1 SSH 73 Client: Protocol (SSH-2.0-Cisco-1.25)
6 0.003419 10.10.10.1 -> 10.10.20.1 TCP 118 [TCP segment of a reassembled PDU]
```

```

7 0.003465 10.10.10.1 -> 10.10.20.1 TCP 118 44274 -> 22 [ACK] Seq=84 Ack=20 Win=4109 Len=64 [TCP segment]
8 0.003482 10.10.10.1 -> 10.10.20.1 TCP 118 44274 -> 22 [ACK] Seq=148 Ack=20 Win=4109 Len=64 [TCP segment]
9 0.003496 10.10.10.1 -> 10.10.20.1 TCP 118 44274 -> 22 [ACK] Seq=212 Ack=20 Win=4109 Len=64 [TCP segment]
10 0.003510 10.10.10.1 -> 10.10.20.1 TCP 118 44274 -> 22 [ACK] Seq=276 Ack=20 Win=4109 Len=64 [TCP segment]
11 0.003525 10.10.10.1 -> 10.10.20.1 TCP 118 44274 -> 22 [ACK] Seq=340 Ack=20 Win=4109 Len=64 [TCP segment]
12 0.004719 10.10.20.1 -> 10.10.10.1 TCP 60 22 -> 44274 [ACK] Seq=20 Ack=84 Win=4045 Len=0
~ Output Cut ~

```

You can see the TCP handshake happening in packet number 1,2,3.

Packet No.1 is the SYN packet.

You can see that it comes with an MSS value of 536.

The SYN, ACK packet (Packet No.2) is also seen coming from the C9200 with an MSS value of 536.

Here, the MSS value remains intact and is not changed by the Switch.

Behaviour After TCP MSS Adjustment

Here is the TCP MSS adjustment configuration on C9500-2 Te1/0/9:

```

C9500-2#sh run int te1/0/9
Building configuration...
Current configuration : 119 bytes
!
interface TenGigabitEthernet1/0/9
no switchport
ip address 192.168.1.2 255.255.255.252
ip tcp adjust-mss 512 -----> Here we are changing the MSS value to 512.

```

Now, take an EPC capture on C9500-2 Te1/0/9 (both directions), and start SSH from C9500-1 to C9200.

Here is the EPC configuration:

```

C9500-2#show monitor capture mycap
Status Information for Capture mycap
Target Type:
Interface: TenGigabitEthernet1/0/9, Direction: BOTH
Status : Inactive
Filter Details:
Capture all packets
Buffer Details:
Buffer Type: LINEAR (default)
Buffer Size (in MB): 80
File Details:
File not associated
Limit Details:
Number of Packets to capture: 0 (no limit)
Packet Capture duration: 0 (no limit)
Packet Size to capture: 0 (no limit)
Maximum number of packets to capture per second: 1000
Packet sampling rate: 0 (no sampling)
C9500-2#

```

Start the capture, SSH from C9500-1 to C9200, and stop the capture.

Here are the CPU-captured packets:

```
C9500-2#show monitor capture mycap buffer brief
Starting the packet display ..... Press Ctrl + Shift + 6 to exit
1 0.000000 b8:a3:77:ec:ba:f7 -> 01:00:0c:cc:cc:cc CDP 398 Device ID: C9500-1.cisco.com Port ID: TenGiga
2 0.636138 10.10.10.1 -> 10.10.20.1 TCP 60 53865 -> 22 [SYN] Seq=0 Win=4128 Len=0 MSS=536
3 0.637980 10.10.20.1 -> 10.10.10.1 TCP 60 22 -> 53865 [SYN, ACK] Seq=0 Ack=1 Win=4128 Len=0 MSS=512
4 0.638214 10.10.10.1 -> 10.10.20.1 TCP 60 53865 -> 22 [ACK] Seq=1 Ack=1 Win=4128 Len=0
5 0.639997 10.10.20.1 -> 10.10.10.1 SSH 73 Server: Protocol (SSH-2.0-Cisco-1.25)
6 0.640208 10.10.10.1 -> 10.10.20.1 SSH 73 Client: Protocol (SSH-2.0-Cisco-1.25)
7 0.640286 10.10.10.1 -> 10.10.20.1 TCP 118 [TCP segment of a reassembled PDU]
8 0.640341 10.10.10.1 -> 10.10.20.1 TCP 118 53865 -> 22 [ACK] Seq=84 Ack=20 Win=4109 Len=64 [TCP segmen
9 0.640360 10.10.10.1 -> 10.10.20.1 TCP 118 53865 -> 22 [ACK] Seq=148 Ack=20 Win=4109 Len=64 [TCP segme
10 0.640375 10.10.10.1 -> 10.10.20.1 TCP 118 53865 -> 22 [ACK] Seq=212 Ack=20 Win=4109 Len=64 [TCP segm
11 0.640390 10.10.10.1 -> 10.10.20.1 TCP 118 53865 -> 22 [ACK] Seq=276 Ack=20 Win=4109 Len=64 [TCP segm
12 0.640410 10.10.10.1 -> 10.10.20.1 TCP 118 53865 -> 22 [ACK] Seq=340 Ack=20 Win=4109 Len=64 [TCP segm
~ Output Cut ~
```

You can see the TCP handshake happening in packets number 2,3,4.

Packet No.2 is the SYN packet.

You can see that it comes with an MSS value of 536.

But the SYN, ACK packet (Packet No.3) is seen coming from the C9200 with an MSS value of 512.

This is because when the SYN packet reaches the C9500-2 Te1/0/9, it is sent to the C9500-2's CPU for TCP MSS modification from 536 to 512.

The C9500-2's CPU changes the MSS to 512 and sends the SYN packet out of Te1/0/2 towards C9200.

Then all following TCP transactions use the same modified MSS value.

Now let's deep dive as to how the SYN packet traverses through the Switch and the MSS change happens.

Once this SYN packet reaches the interface of the C9500-2, it is sent to the CPU for MSS modification.

It first goes through the FED (where you can capture it), and then goes to the CPU (where you can capture it as well).

Let's first take a FED Punt capture on C9500-2.

Here is the FED punt capture configuration:

```
C9500-2#debug platform software fed switch 1 punt packet-capture buffer limit 16384
Punt PCAP buffer configure: one-time with buffer size 16384...done
```

Starting the FED punt capture:

```
C9500-2#debug platform software fed switch 1 punt packet-capture start
Punt packet capturing started.
```

Starting the SSH from C9500-1 to C9200:

```
C9500-1#ssh -l admin 10.10.20.1
Password:
```

Stopping the FED punt capture:

```
C9500-2#debug platform software fed switch 1 punt packet-capture stop
Punt packet capturing stopped. Captured 3 packet(s)
```

And here are the FED punt captured packets:

```
C9500-2#show platform software fed switch active punt packet-capture brief
Punt packet capturing: disabled. Buffer wrapping: disabled
Total captured so far: 3 packets. Capture capacity : 16384 packets
```

```
----- Punt Packet Number: 1, Timestamp: 2024/07/31 01:29:46.373 -----
interface : physical: TenGigabitEthernet1/0/9[if-id: 0x00000040], pa1: TenGigabitEthernet1/0/9 [if-id: 0x00000040]
metadata : cause: 55 [For-us control], sub-cause: 0, q-no: 4, linktype: MCP_LINK_TYPE_IP [1]
ether hdr : dest mac: 0100.5e00.0005, src mac: b8a3.77ec.baf7
ether hdr : ethertype: 0x0800 (IPv4)
ipv4 hdr : dest ip: 224.0.0.5, src ip: 192.168.1.1
ipv4 hdr : packet len: 100, ttl: 1, protocol: 89
```

```
----- Punt Packet Number: 2, Timestamp: 2024/07/31 01:29:47.432 -----
interface : physical: TenGigabitEthernet1/0/9[if-id: 0x00000040], pa1: TenGigabitEthernet1/0/9 [if-id: 0x00000040]
metadata : cause: 11 [For-us data], sub-cause: 1, q-no: 14, linktype: MCP_LINK_TYPE_IP [1]
ether hdr : dest mac: 00a3.d144.4bf7, src mac: b8a3.77ec.baf7
ether hdr : ethertype: 0x0800 (IPv4)
ipv4 hdr : dest ip: 10.10.20.1, src ip: 10.10.10.1
ipv4 hdr : packet len: 44, ttl: 254, protocol: 6 (TCP)
tcp hdr : dest port: 22, src port: 35916
```

```
----- Punt Packet Number: 3, Timestamp: 2024/07/31 01:29:48.143 -----
interface : physical: TenGigabitEthernet1/0/1[if-id: 0x00000009], pa1: TenGigabitEthernet1/0/1 [if-id: 0x00000009]
metadata : cause: 96 [Layer2 control protocols], sub-cause: 0, q-no: 1, linktype: MCP_LINK_TYPE_LAYER2
ether hdr : dest mac: 0100.0ccc.cccc, src mac: 78bc.1a27.c203
ether hdr : length: 443
```

You can see that Packet No. 2 is the TCP SYN packet from 10.10.10.1 to 10.10.20.1, coming in from Te1/0/9.

The 'q-no' is important to note here. You can see that it chooses Queue No. 14 to go from the FED to the CPU.

Here you can see all the 32 queues present for traffic to move from the FED towards the CPU:

```
C9500-2#show platform hardware fed switch active qos queue stats internal cpu policer
```

```
CPU Queue Statistics
```

```
=====
(default) (set) Queue Queue
```

QId PlcIdx Queue Name Enabled Rate Rate Drop(Bytes) Drop(Frames)

```
-----  
0 11 DOT1X Auth Yes 1000 1000 0 0  
1 1 L2 Control Yes 2000 2000 0 0  
2 14 Forus traffic Yes 4000 4000 0 0  
3 0 ICMP GEN Yes 600 600 0 0  
4 2 Routing Control Yes 5400 5400 0 0  
5 14 Forus Address resolution Yes 4000 4000 0 0  
6 0 ICMP Redirect Yes 600 600 0 0  
7 16 Inter FED Traffic Yes 2000 2000 0 0  
8 4 L2 LVX Cont Pack Yes 1000 1000 0 0  
9 19 EWLC Control Yes 13000 13000 0 0  
10 16 EWLC Data Yes 2000 2000 0 0  
11 13 L2 LVX Data Pack Yes 1000 1000 0 0  
12 0 BROADCAST Yes 600 600 0 0  
13 10 Openflow Yes 200 200 0 0  
14 13 Sw forwarding Yes 1000 1000 0 0  
15 8 Topology Control Yes 13000 13000 0 0  
16 12 Proto Snooping Yes 2000 2000 0 0  
17 6 DHCP Snooping Yes 400 400 0 0  
18 13 Transit Traffic Yes 1000 1000 0 0  
19 10 RPF Failed Yes 200 200 0 0  
20 15 MCAST END STATION Yes 2000 2000 0 0  
21 13 LOGGING Yes 1000 1000 0 0  
22 7 Punt Webauth Yes 1000 1000 0 0  
23 18 High Rate App Yes 13000 13000 0 0  
24 10 Exception Yes 200 200 0 0  
25 3 System Critical Yes 1000 1000 0 0  
26 10 NFL SAMPLED DATA Yes 200 200 0 0  
27 2 Low Latency Yes 5400 5400 0 0  
28 10 EGR Exception Yes 200 200 0 0  
29 5 Stackwise Virtual OOB Yes 8000 8000 0 0  
30 9 MCAST Data Yes 400 400 0 0  
31 3 Gold Pkt Yes 1000 1000 0 0
```

As you can see overhead, Queue No. 14 is the 'Sw forwarding' queue.
In this case, this queue is used by the TCP traffic in order to get punted to the CPU.

Now, let's take a CPU (Control-Plane) capture on C9500-2.

Here is the CPU capture configuration:

```
C9500-2#sh mon cap test  
Status Information for Capture test  
Target Type:  
Interface: Control Plane, Direction: BOTH  
Status : Inactive  
Filter Details:  
Capture all packets  
Buffer Details:  
Buffer Type: LINEAR (default)  
Buffer Size (in MB): 80  
File Details:  
File not associated  
Limit Details:  
Number of Packets to capture: 0 (no limit)  
Packet Capture duration: 0 (no limit)  
Packet Size to capture: 0 (no limit)
```


Packet sampling rate: 0 (no sampling)
C9500-2#

You start the capture, SSH from C9500-1 to C9200, and stop the capture.

Here are the CPU-captured packets:

```
C9500-2#show monitor capture test buffer brief
Starting the packet display ..... Press Ctrl + Shift + 6 to exit
 1 0.000000 00:a3:d1:44:4b:81 -> 01:80:c2:00:00:00 STP 60 RST. Root = 32768/1/00:a3:d1:44:4b:80 Cost = 0
 2 0.000010 00:a3:d1:44:4b:a3 -> 01:80:c2:00:00:00 STP 60 RST. Root = 32768/1/00:a3:d1:44:4b:80 Cost = 0
 3 0.000013 00:a3:d1:44:4b:a4 -> 01:80:c2:00:00:00 STP 60 RST. Root = 32768/1/00:a3:d1:44:4b:80 Cost = 0
 4 0.000016 00:a3:d1:44:4b:a6 -> 01:80:c2:00:00:00 STP 60 RST. Root = 32768/1/00:a3:d1:44:4b:80 Cost = 0
 5 0.000019 00:a3:d1:44:4b:a7 -> 01:80:c2:00:00:00 STP 60 RST. Root = 32768/1/00:a3:d1:44:4b:80 Cost = 0
 6 0.000022 00:a3:d1:44:4b:a8 -> 01:80:c2:00:00:00 STP 60 RST. Root = 32768/1/00:a3:d1:44:4b:80 Cost = 0
 7 0.055470 c0:8b:2a:04:f0:6c -> 01:80:c2:00:00:0e LLDP 117 TTL = 120 SysName = bg118-cx-amx-b02-2.cisco
 9 0.220331 28:63:29:20:31:39 -> 00:01:22:53:74:20 0x3836 30 Ethernet II
10 0.327316 192.168.1.1 -> 224.0.0.5 OSPF 114 Hello Packet
11 0.442986 c0:8b:2a:04:f0:68 -> 01:80:c2:00:00:0e LLDP 117 TTL = 120 SysName = bg118-cx-amx-b02-2.cisco
12 1.714121 10.10.10.1 -> 10.10.20.1 TCP 60 23098 -> 22 [SYN] Seq=0 Win=4128 Len=0 MSS=536
13 1.714375 10.10.10.1 -> 10.10.20.1 TCP 60 [TCP Out-Of-Order] 23098 -> 22 [SYN] Seq=0 Win=4128 Len=0 MSS=512
14 2.000302 00:a3:d1:44:4b:81 -> 01:80:c2:00:00:00 STP 60 RST. Root = 32768/1/00:a3:d1:44:4b:80 Cost = 0
15 2.000310 00:a3:d1:44:4b:a3 -> 01:80:c2:00:00:00 STP 60 RST. Root = 32768/1/00:a3:d1:44:4b:80 Cost = 0
~ Output Cut ~
```

Packet No. 12 is the TCP SYN packet coming into the CPU (punt), with the default MSS value of 536.
Packet No. 13 is the TCP SYN packet sent out by the CPU (inject), after modifying the MSS value to 512.

You can also take a quick CPU debug in order to see this change happening.

Here is the CPU debug configuration:

```
C9500-2#debug ip tcp adjust-mss
TCP Adjust Mss debugging is on
```

Starting the SSH from C9500-1 to C9200:

```
C9500-1#ssh -l admin 10.10.20.1
Password:
```

Stopping the CPU Debug:

```
C9500-2#undebug all
All possible debugging has been turned off
```

Looking at the logs for the debugs:

```
C9500-2#show logging
Syslog logging: enabled (0 messages dropped, 2 messages rate-limited, 0 flushes, 0 overruns, xml disabled)
No Active Message Discriminator.
No Inactive Message Discriminator.
Console logging: disabled
Monitor logging: level debugging, 0 messages logged, xml disabled,
filtering disabled
Buffer logging: level debugging, 230 messages logged, xml disabled,
filtering disabled
Exception Logging: size (4096 bytes)
Count and timestamp logging messages: disabled
File logging: disabled
Persistent logging: disabled
No active filter modules.
Trap logging: level informational, 210 message lines logged
Logging Source-Interface: VRF Name:
TLS Profiles:
Log Buffer (102400 bytes):
*Jul 31 01:46:32.052: TCPADJMSS: process_enqueue_feature
*Jul 31 01:46:32.893: TCPADJMSS: process_enqueue_feature
*Jul 31 01:46:36.136: TCPADJMSS: process_enqueue_feature
*Jul 31 01:46:41.318: TCPADJMSS: process_enqueue_feature
*Jul 31 01:46:42.412: TCPADJMSS: process_enqueue_feature
*Jul 31 01:46:43.254: TCPADJMSS: process_enqueue_feature
*Jul 31 01:46:43.638: TCPADJMSS: process_enqueue_feature
*Jul 31 01:46:45.783: TCPADJMSS: Input (process)
*Jul 31 01:46:45.783: TCPADJMSS: orig_mss = 536 adj_mss = 512 src_ip = 10.10.10.1 dest_ip = 10.10.20.1
*Jul 31 01:46:45.783: TCPADJMSS: paktype = 0x7F83C7BCBF78
*Jul 31 01:46:50.456: TCPADJMSS: process_enqueue_feature
*Jul 31 01:46:51.985: TCPADJMSS: process_enqueue_feature
C9500-2#
```

You can see overhead that the Original MSS value of 536 is being adjusted to 512.

Finally, you can take an EPC capture on C9200 Gi1/0/3 in order to confirm that the TCP SYN packet is indeed coming with an MSS of 512.

Here is the EPC configuration:

```
C9200#sh mon cap mycap
Status Information for Capture mycap
Target Type:
Interface: GigabitEthernet1/0/3, Direction: BOTH
Status : Inactive
Filter Details:
Capture all packets
Buffer Details:
Buffer Type: LINEAR (default)
Buffer Size (in MB): 80
Limit Details:
Number of Packets to capture: 0 (no limit)
Packet Capture duration: 0 (no limit)
```

Packet Size to capture: 0 (no limit)
Packet sampling rate: 0 (no sampling)
C9200#

You start the capture, SSH from C9500-1 to C9200, and stop the capture.

Here are the CPU-captured packets:

```
C9200#sh mon cap mycap buff br
```


```
-----  
# size timestamp source destination dscp protocol  
-----  
0 118 0.000000 192.168.2.1 -> 224.0.0.5 48 CS6 OSPF  
1 64 0.721023 10.10.10.1 -> 10.10.20.1 48 CS6 TCP  
2 64 0.722015 10.10.10.1 -> 10.10.20.1 48 CS6 TCP  
3 77 0.728026 10.10.10.1 -> 10.10.20.1 48 CS6 TCP  
4 122 0.728026 10.10.10.1 -> 10.10.20.1 48 CS6 TCP  
5 122 0.728026 10.10.10.1 -> 10.10.20.1 48 CS6 TCP  
6 122 0.728026 10.10.10.1 -> 10.10.20.1 48 CS6 TCP  
7 122 0.728026 10.10.10.1 -> 10.10.20.1 48 CS6 TCP  
8 122 0.728026 10.10.10.1 -> 10.10.20.1 48 CS6 TCP  
9 122 0.728026 10.10.10.1 -> 10.10.20.1 48 CS6 TCP  
10 122 0.730025 10.10.10.1 -> 10.10.20.1 48 CS6 TCP  
~ Output Cut ~
```

In C9200, you are not able to see the packet details like in Wireshark, only the brief and hex details are available.

Hence you can export the earlier packets to a **pcap** file in the flash.

```
C9200#mon cap mycap export flash:Gi1-0-3-Both.pcapng  
Exported Successfully
```

Then you can copy this file via TFTP to your local PC, and open the file in Wireshark.
Here is the Wireshark capture.



You can see that the TCP MSS value of the SYN packet is 512.

TCP MSS Adjustment Causing Slowness during Enormous Amount of TCP Traffic

Now, let's assume a network has multiple devices using TCP traffic.

For example, they can be transferring files, or accessing a TCP-based application (such as a Citrix Server).

You have simulated it by connecting an IXIA (traffic generator) to C9500-2 Te1/0/37, sending TCP SYN packets at a high rate.

This IXIA device acts as a network segment, where multiple users are using TCP-based applications.

You have configured **ip tcp adjust-mss** CLI on Te1/0/37.


This causes all TCP traffic being received on Te1/0/37 to be punted to C9500-2's CPU.

This in turn chokes the C9500-2's COPP Policer's 'Sw forwarding' queue, as mentioned earlier in the document.

As a consequence, the SSH session establishment from C9500-1 to C9200 is affected.

Either the SSH session does not form, and get timed out, or is established after a delay.

Here is how the topology looks:



Let's see this in action.

Here is the configuration of C9500-2 Te1/0/37:

```
C9500-2#sh run int te1/0/37
Building configuration...
Current configuration : 135 bytes
interface TenGigabitEthernet1/0/37
no switchport
ip address 10.10.40.1 255.255.255.0
ip tcp adjust-mss 500
load-interval 30
end
```

Now you start sending enormous traffic from IXIA into the Te1/0/37 interface. Let's take a look at the incoming traffic rate:

```
C9500-2#sh int te1/0/37 | in rate
Queueing strategy: fifo
30 second input rate 6425812000 bits/sec, 12550415 packets/sec → We can see the enormous Input rate.
30 second output rate 0 bits/sec, 0 packets/sec
```

Let's try to SSH from C9500-1 to C9200 now:

```
C9500-1#ssh -l admin 10.10.20.1
% Connection timed out; remote host not responding
C9500-1#
```

You can clearly see that the C9500-1 was not able to SSH into the C9200.

This is because the TCP SYN packet being sent by the C9500-1 was getting dropped by the 'Sw forwarding' queue, which is being bombarded with traffic from Te1/0/37.

Let's take a look at the queue:

```
C9500-2#sh platform hardware fed switch active qos queue stats internal cpu policer
CPU Queue Statistics
```

```
=====
(default) (set) Queue Queue
QId PlcIdx Queue Name Enabled Rate Rate Drop(Bytes) Drop(Frames)
-----
0 11 DOT1X Auth Yes 1000 1000 0 0
1 1 L2 Control Yes 2000 2000 0 0
2 14 Forus traffic Yes 4000 4000 0 0
3 0 ICMP GEN Yes 600 600 0 0
4 2 Routing Control Yes 5400 5400 0 0
5 14 Forus Address resolution Yes 4000 4000 0 0
6 0 ICMP Redirect Yes 600 600 0 0
7 16 Inter FED Traffic Yes 2000 2000 0 0
8 4 L2 LVX Cont Pack Yes 1000 1000 0 0
9 19 EWLC Control Yes 13000 13000 0 0
10 16 EWLC Data Yes 2000 2000 0 0
11 13 L2 LVX Data Pack Yes 1000 1000 0 0
12 0 BROADCAST Yes 600 600 0 0
13 10 Openflow Yes 200 200 0 0
14 13 Sw forwarding Yes 1000 1000 39683368064 620052629 → We can see the huge number of dropped packets in t
15 8 Topology Control Yes 13000 13000 0 0
16 12 Proto Snooping Yes 2000 2000 0 0
17 6 DHCP Snooping Yes 400 400 0 0
18 13 Transit Traffic Yes 1000 1000 0 0
19 10 RPF Failed Yes 200 200 0 0
20 15 MCAST END STATION Yes 2000 2000 0 0
21 13 LOGGING Yes 1000 1000 0 0
22 7 Punt Webauth Yes 1000 1000 0 0
23 18 High Rate App Yes 13000 13000 0 0
24 10 Exception Yes 200 200 0 0
25 3 System Critical Yes 1000 1000 0 0
26 10 NFL SAMPLED DATA Yes 200 200 0 0
27 2 Low Latency Yes 5400 5400 0 0
28 10 EGR Exception Yes 200 200 0 0
29 5 Stackwise Virtual OOB Yes 8000 8000 0 0
30 9 MCAST Data Yes 400 400 0 0
31 3 Gold Pkt Yes 1000 1000 0 0
```

Let's collect the output multiple times in order to ensure the dropped count is increasing during the issue:

```
C9500-2#sh platform hardware fed switch active qos queue stats internal cpu policer | in Sw forwarding
14 13 Sw forwarding Yes 1000 1000 47046906560 735107915
14 13 21 Sw forwarding Yes
13 system-cpp-police-sw-forward : Sw forwarding/ LOGGING/ L2 LVX Data Pack/ Transit Traffic/
21 system-cpp-police-ios-feature : ICMP GEN/ BROADCAST/ ICMP Redirect/ L2 LVX Cont Pack/ Proto Snooping
C9500-2#
!
C9500-2#sh platform hardware fed switch active qos queue stats internal cpu policer | in Sw forwarding
14 13 Sw forwarding Yes 1000 1000 47335535936 739617752
```

```

14 13 21 Sw forwarding Yes
13 system-cpp-police-sw-forward : Sw forwarding/ LOGGING/ L2 LVX Data Pack/ Transit Traffic/
21 system-cpp-police-ios-feature : ICMP GEN/ BROADCAST/ ICMP Redirect/ L2 LVX Cont Pack/ Proto Snooping
C9500-2#
!
C9500-2#sh platform hardware fed switch active qos queue stats internal cpu policer | in Sw forwarding
14 13 Sw forwarding Yes 1000 1000 47666441088 744788145
14 13 21 Sw forwarding Yes
13 system-cpp-police-sw-forward : Sw forwarding/ LOGGING/ L2 LVX Data Pack/ Transit Traffic/
21 system-cpp-police-ios-feature : ICMP GEN/ BROADCAST/ ICMP Redirect/ L2 LVX Cont Pack/ Proto Snooping
C9500-2#

```

As you can see, the dropped count is increasing, and the SSH traffic (TCP SYN packet) is getting dropped here.

Now if you are unaware via which interface/SVI you are getting this inflow of traffic, you have a specific command to help out.

```

C9500-2#show platform software fed switch active punt rates interfaces
Punt Rate on Interfaces Statistics
Packets per second averaged over 10 seconds, 1 min and 5 mins
=====
| | Recv | Recv | Recv | Drop | Drop | Drop
Interface Name | IF_ID | 10s | 1min | 5min | 10s | 1min | 5min
=====
TenGigabitEthernet1/0/37 0x00000042 1000 1000 1000 0 0 0
-----
C9500-2#

```

The **show platform software fed switch active punt rates interfaces** command gives us the list of interfaces which are responsible for receiving huge amount of traffic that is being punted to the CPU. You can clearly see Te1/0/37 here, which is the interface through which you are getting the TCP traffic.

Now, if you want to see the amount of traffic hitting all the COPP Policer queues (that is being received on the earlier interface), you can use:

show platform software fed switch active punt rates interfaces <IF_ID from the above output>

Let's have a look:

```

C9500-2#show platform software fed switch active punt rates interfaces 0x42
Punt Rate on Single Interfaces Statistics
Interface : TenGigabitEthernet1/0/37 [if_id: 0x42]

Received Dropped
-----
Total : 2048742 Total : 0
10 sec average : 1000 10 sec average : 0
1 min average : 1000 1 min average : 0
5 min average : 1000 5 min average : 0

Per CPUQ punt stats on the interface (rate averaged over 10s interval)
=====
Q | Queue | Recv | Recv | Drop | Drop |

```

```

no | Name | Total | Rate | Total | Rate |
=====
0 CPU_Q_DOT1X_AUTH 0 0 0 0
1 CPU_Q_L2_CONTROL 7392 0 0 0
2 CPU_Q_FORUS_TRAFFIC 0 0 0 0
3 CPU_Q_ICMP_GEN 0 0 0 0
4 CPU_Q_ROUTING_CONTROL 0 0 0 0
5 CPU_Q_FORUS_ADDR_RESOLUTION 0 0 0 0
6 CPU_Q_ICMP_REDIRECT 0 0 0 0
7 CPU_Q_INTER_FED_TRAFFIC 0 0 0 0
8 CPU_Q_L2LVX_CONTROL_PKT 0 0 0 0
9 CPU_Q_EWLC_CONTROL 0 0 0 0
10 CPU_Q_EWLC_DATA 0 0 0 0
11 CPU_Q_L2LVX_DATA_PKT 0 0 0 0
12 CPU_Q_BROADCAST 0 0 0 0
13 CPU_Q_CONTROLLER_PUNT 0 0 0 0
14 CPU_Q_SW_FORWARDING 2006390 1000 0 0 -----> We can see high amount of traffic hitting the Sw forward
15 CPU_Q_TOPOLOGY_CONTROL 0 0 0 0
16 CPU_Q_PROTO_SNOOPING 0 0 0 0
17 CPU_Q_DHCP_SNOOPING 0 0 0 0
18 CPU_Q_TRANSIT_TRAFFIC 0 0 0 0
19 CPU_Q_RPF_FAILED 0 0 0 0
20 CPU_Q_MCAST_END_STATION_SERVICE 0 0 0 0
21 CPU_Q_LOGGING 34960 0 0 0
22 CPU_Q_PUNT_WEBAUTH 0 0 0 0
23 CPU_Q_HIGH_RATE_APP 0 0 0 0
24 CPU_Q_EXCEPTION 0 0 0 0
25 CPU_Q_SYSTEM_CRITICAL 0 0 0 0
26 CPU_Q_NFL_SAMPLED_DATA 0 0 0 0
27 CPU_Q_LOW_LATENCY 0 0 0 0
28 CPU_Q_EGR_EXCEPTION 0 0 0 0
29 CPU_Q_FSS 0 0 0 0
30 CPU_Q_MCAST_DATA 0 0 0 0
31 CPU_Q_GOLD_PKT 0 0 0 0
-----

```

Collecting the output multiple times in very short intervals:

```

C9500-2#show platform software fed switch active punt rates interfaces 0x42 | in SW_FORWARDING
14 CPU_Q_SW_FORWARDING 2126315 1000 0 0
C9500-2#
C9500-2#show platform software fed switch active punt rates interfaces 0x42 | in SW_FORWARDING
14 CPU_Q_SW_FORWARDING 2128390 1000 0 0
C9500-2#
C9500-2#show platform software fed switch active punt rates interfaces 0x42 | in SW_FORWARDING
14 CPU_Q_SW_FORWARDING 2132295 1000 0 0
C9500-2#

```

This clearly shows that the Sw forwarding queue is choked.

Once you remove the `ip tcp adjust-mss` command from the Te1/0/37, or if you stop this TCP traffic, the SSH access from C9500-1 to C9200 immediately gets re-established.

Let's have a look at the SSH session after shutting down C9500-2 Te1/0/37:


```
C9500-1#ssh -l admin 10.10.20.1  
Password:
```

You can see that the SSH access is again restored.

Hence, you can correlate the TCP Slowness here (SSH access blocked) due to the high amount of TCP traffic in the network, with TCP MSS adjustment.

Important Points

1. Whenever you have TCP slowness in your network, such as File transfer slowness, accessibility to TCP-related applications and so on, and you have TCP MSS adjustment configured on a Catalyst Switch, ensure to check the COPP Policer drops in order to check if there is a high amount of TCP traffic in the network or not.
2. If you have configured TCP MSS adjustment on a Catalyst Switch, ensure that the TCP traffic in your network does not oversubscribe the COPP Policer rate, otherwise, TCP-related issues (slowness, packet drops) are seen in your network.