



Configuring TCP

First Published: October 23, 2006

Last Updated: November 19, 2008

The Transmission Control Protocol (TCP) is a protocol that specifies the format of data and acknowledgments used in data transfer. TCP is a connection-oriented protocol because participants must establish a connection before data can be transferred. By performing flow control and error correction, TCP guarantees reliable, in-sequence delivery of packets. It is considered a reliable protocol because if an IP packet is dropped or received out of order, TCP will request the correct packet until it receives it.

Finding Feature Information

Your software release may not support all the features documented in this module. For the latest feature information and caveats, see the release notes for your platform and software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the [“Feature Information for TCP” section on page 19](#).

Use Cisco Feature Navigator to find information about platform support and Cisco IOS and Catalyst OS software image support. To access Cisco Feature Navigator, go to <http://www.cisco.com/go/cfn>. An account on Cisco.com is not required.

Contents

- [Prerequisites for TCP, page 2](#)
- [Information About TCP, page 2](#)
- [How to Configure TCP, page 7](#)
- [Configuration Examples for TCP, page 13](#)
- [Additional References, page 17](#)
- [Feature Information for TCP, page 19](#)
- [Glossary, page 24](#)



Americas Headquarters:
Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

© 2008 Cisco Systems, Inc. All rights reserved.

Prerequisites for TCP

TCP Time Stamp, TCP Selective Acknowledgment, and TCP Header Compression

Because TCP time stamps are always sent and echoed in both directions and the time-stamp value in the header is always changing, TCP header compression will not compress the outgoing packet. To allow TCP header compression over a serial link, the TCP time-stamp option is disabled. If you want to use TCP header compression over a serial line, TCP time stamp and TCP selective acknowledgment must be disabled. Both features are disabled by default. Use the **no ip tcp selective-ack** command to disable TCP selective acknowledgment once it is enabled.

Information About TCP

To configure TCP, you should understand the following concepts:

- [TCP Services, page 2](#)
- [TCP Connection Establishment, page 3](#)
- [TCP Connection Attempt Time, page 3](#)
- [TCP Selective Acknowledgment, page 3](#)
- [TCP Time Stamp, page 4](#)
- [TCP Maximum Read Size, page 4](#)
- [TCP Path MTU Discovery, page 4](#)
- [TCP Window Scaling, page 5](#)
- [TCP Sliding Window, page 5](#)
- [TCP Outgoing Queue Size, page 6](#)
- [TCP Congestion Avoidance, page 6](#)
- [TCP Explicit Congestion Notification, page 6](#)
- [TCP MSS Adjustment, page 6](#)
- [TCP Applications Flags Enhancement, page 7](#)
- [TCP Show Extension, page 7](#)

TCP Services

TCP provides reliable transmission of data in an IP environment. TCP corresponds to the transport layer (Layer 4) of the OSI reference model. Among the services TCP provides are stream data transfer, reliability, efficient flow control, full-duplex operation, and multiplexing.

With stream data transfer, TCP delivers an unstructured stream of bytes identified by sequence numbers. This service benefits applications because they do not have to chop data into blocks before handing it off to TCP. Instead, TCP groups bytes into segments and passes them to IP for delivery.

TCP offers reliability by providing connection-oriented, end-to-end reliable packet delivery through an internetwork. It does this by sequencing bytes with a forwarding acknowledgment number that indicates to the destination the next byte the source expects to receive. Bytes not acknowledged within a specified

time period are retransmitted. The reliability mechanism of TCP allows devices to deal with lost, delayed, duplicate, or misread packets. A time-out mechanism allows devices to detect lost packets and request retransmission.

TCP offers efficient flow control, which means that, when sending acknowledgments back to the source, the receiving TCP process indicates the highest sequence number it can receive without overflowing its internal buffers.

TCP offers full-duplex operation and TCP processes can both send and receive at the same time.

TCP multiplexing allows numerous simultaneous upper-layer conversations to be multiplexed over a single connection.

TCP Connection Establishment

To use reliable transport services, TCP hosts must establish a connection-oriented session with one another. Connection establishment is performed by using a “three-way handshake” mechanism.

A three-way handshake synchronizes both ends of a connection by allowing both sides to agree upon initial sequence numbers. This mechanism also guarantees that both sides are ready to transmit data and know that the other side is ready to transmit as well. The three-way handshake is necessary so that packets are not transmitted or retransmitted during session establishment or after session termination.

Each host randomly chooses a sequence number used to track bytes within the stream it is sending. Then, the three-way handshake proceeds in the following manner:

- The first host (Host A) initiates a connection by sending a packet with the initial sequence number (X) and synchronize/start (SYN) bit set to indicate a connection request.
- The second host (Host B) receives the SYN, records the sequence number X, and replies by acknowledging the SYN (with an ACK = X + 1). Host B includes its own initial sequence number (SEQ = Y). An ACK = 20 means the host has received bytes 0 through 19 and expects byte 20 next. This technique is called forward acknowledgment.
- Host A acknowledges all bytes Host B sent with a forward acknowledgment indicating the next byte Host A expects to receive (ACK = Y + 1). Data transfer then can begin.

TCP Connection Attempt Time

You can set the amount of time the Cisco IOS software will wait to attempt to establish a TCP connection. Because the connection attempt time is a host parameter, it does not pertain to traffic going through the device, just to traffic originated at the device. To set the TCP connection attempt time, use the **ip tcp synwait-time** command in global configuration mode. The default is 30 seconds.

TCP Selective Acknowledgment

The TCP Selective Acknowledgment feature improves performance in the event that multiple packets are lost from one TCP window of data.

Prior to this feature, with the limited information available from cumulative acknowledgments, a TCP sender could learn about only one lost packet per round-trip time. An aggressive sender could choose to resend packets early, but such resent segments might have already been successfully received.

The TCP selective acknowledgment mechanism helps improve performance. The receiving TCP host returns selective acknowledgment packets to the sender, informing the sender of data that have been received. In other words, the receiver can acknowledge packets received out of order. The sender can then resend only the missing data segments (instead of everything since the first missing packet).

Prior to selective acknowledgment, if TCP lost packets 4 and 7 out of an 8-packet window, TCP would receive acknowledgment of only packets 1, 2, and 3. Packets 4 through 8 would need to be resent. With selective acknowledgment, TCP receives acknowledgment of packets 1, 2, 3, 5, 6, and 8. Only packets 4 and 7 must be resent.

TCP selective acknowledgment is used only when multiple packets are dropped within one TCP window. There is no performance impact when the feature is enabled but not used. Use the **ip tcp selective-ack** command in global configuration mode to enable TCP selective acknowledgment.

Refer to RFC 2018 for more detailed information about TCP selective acknowledgment.

TCP Time Stamp

The TCP time-stamp option provides better TCP round-trip time measurements. Because the time stamps are always sent and echoed in both directions and the time-stamp value in the header is always changing, TCP header compression will not compress the outgoing packet. To allow TCP header compression over a serial link, the TCP time-stamp option is disabled. Use the **ip tcp timestamp** command to enable the TCP time-stamp option.

Refer to RFC 1323 for more detailed information on TCP time stamp. Refer to the [“Configuring TCP Header Compression”](#) chapter of the *Cisco IOS Quality of Service Solutions Configuration Guide* for more information about TCP header compression.

TCP Maximum Read Size

The maximum number of characters that TCP reads from the input queue for Telnet and rlogin at one time is a very large number (the largest possible 32-bit positive number) by default. To change the TCP maximum read size value, use the **ip tcp chunk-size** command in global configuration mode.

We do not recommend that you change this value.

TCP Path MTU Discovery

Path MTU Discovery is a method for maximizing the use of available bandwidth in the network between the endpoints of a TCP connection, which is described in RFC 1191. IP Path MTU Discovery allows a host to dynamically discover and cope with differences in the maximum allowable maximum transmission unit (MTU) size of the various links along the path. Sometimes a router is unable to forward a datagram because it requires fragmentation (the packet is larger than the MTU you set for the interface with the interface configuration command), but the “don’t fragment” (DF) bit is set. The intermediate gateway sends a “Fragmentation needed and DF bit set” ICMP message to the sending host, alerting it to the problem. Upon receiving this ICMP message, the host reduces its assumed path MTU and consequently sends a smaller packet that will fit the smallest packet size of all the links along the path.

By default, TCP Path MTU Discovery is disabled. Existing connections are not affected when this feature is enabled or disabled.

Customers using TCP connections to move bulk data between systems on distinct subnets would benefit most by enabling this feature. Customers using remote source-route bridging (RSRB) with TCP encapsulation, serial tunnel (STUN), X.25 Remote Switching (also known as XOT or X.25 over TCP), and some protocol translation configurations might also benefit from enabling this feature.

Use the **ip tcp path-mtu-discovery** global configuration command to enable Path MTU Discovery for connections initiated by the router when it is acting as a host.

For more information about Path MTU Discovery, refer to the “[Configuring IP Services](#)” chapter of the *Cisco IOS IP Application Services Configuration Guide*.

TCP Window Scaling

The TCP Window Scaling feature adds support for the Window Scaling option in RFC 1323. A larger window size is recommended to improve TCP performance in network paths with large bandwidth-delay product characteristics that are called Long Fat Networks (LFNs). The TCP Window Scaling enhancement provides that support.

The window scaling extension in Cisco IOS software expands the definition of the TCP window to 32 bits and then uses a scale factor to carry this 32-bit value in the 16-bit window field of the TCP header. The window size can increase to a scale factor of 14. Typical applications use a scale factor of 3 when deployed in LFNs.

The TCP Window Scaling feature complies with RFC 1323, *TCP Extensions for High Performance*. The maximum window size has been increased to 1,073,741,823 bytes. The larger scalable window size will allow TCP to perform better over LFNs. Use the **ip tcp window-size** command in global configuration mode to configure the TCP window size.

TCP Sliding Window

A TCP sliding window provides more efficient use of network bandwidth because it enables hosts to send multiple bytes or packets before waiting for an acknowledgment.

In TCP, the receiver specifies the current window size in every packet. Because TCP provides a byte-stream connection, window sizes are expressed in bytes. A window is the number of data bytes that the sender is allowed to send before waiting for an acknowledgment. Initial window sizes are indicated at connection setup, but might vary throughout the data transfer to provide flow control. A window size of zero means “Send no data.” The default TCP window size is 4128 bytes. We recommend you keep the default value unless you know your router is sending large packets (greater than 536 bytes). Use the **ip tcp window-size** command to change the default window size.

In a TCP sliding-window operation, for example, the sender might have a sequence of bytes to send (numbered 1 to 10) to a receiver who has a window size of five. The sender then places a window around the first five bytes and transmits them together. The sender then waits for an acknowledgment.

The receiver responds with an ACK = 6, indicating that it has received bytes 1 to 5 and is expecting byte 6 next. In the same packet, the receiver indicates that its window size is 5. The sender then moves the sliding window five bytes to the right and transmit bytes 6 to 10. The receiver responds with an ACK = 11, indicating that it is expecting sequenced byte 11 next. In this packet, the receiver might indicate that its window size is 0 (because, for example, its internal buffers are full). At this point, the sender cannot send any more bytes until the receiver sends another packet with a window size greater than 0.

TCP Outgoing Queue Size

The default TCP outgoing queue size per connection is 5 segments if the connection has a TTY associated with it (such as a Telnet connection). If no TTY connection is associated with a connection, the default queue size is 20 segments. Use the **ip tcp queuemax** command to change the 5-segment default value.

TCP Congestion Avoidance

The TCP Congestion Avoidance feature enables the monitoring of acknowledgment packets to the TCP sender when multiple packets are lost in a single window of data. Previously the sender would exit Fast-Recovery mode, wait for three or more duplicate acknowledgment packets before retransmitting the next unacknowledged packet, or wait for the retransmission timer to slow start. This could lead to performance issues.

Implementation of RFC 2581 and RFC 3782 addresses the modifications to the Fast-Recovery algorithm that incorporates a response to partial acknowledgments received during Fast Recovery, improving performance in situations where multiple packets are lost in a single window of data.

This feature is an enhancement to the existing Fast Recovery algorithm. There are no commands used to enable or disable this feature.

To monitor the acknowledgment packets, the output of the **debug ip tcp transactions** command has been enhanced to show the following conditions:

- TCP entering Fast Recovery mode.
- Duplicate acknowledgments being received during Fast Recovery mode.
- Partial acknowledgments being received.

TCP Explicit Congestion Notification

The TCP Explicit Congestion Notification (ECN) feature provides a method for an intermediate router to notify the end hosts of impending network congestion. It also provides enhanced support for TCP sessions associated with applications that are sensitive to delay or packet loss including Telnet, web browsing, and transfer of audio and video data. The benefit of this feature is the reduction of delay and packet loss in data transmissions. Use the **ip tcp ecn** command in global configuration mode to enable TCP ECN.

TCP MSS Adjustment

The TCP MSS Adjustment feature enables the configuration of the maximum segment size (MSS) for transient packets that traverse a router, specifically TCP segments with the SYN bit set. Use the **ip tcp adjust-mss** command in interface configuration mode to specify the MSS value on the intermediate router of the SYN packets 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 value of the MSS field is determined by the maximum transmission unit (MTU) configuration on the host. The default MSS value for a PC is 1500 bytes.

The PPP over Ethernet (PPPoE) standard supports a 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 may be dropped because system administrators sometimes disable the ICMP error messages that must be relayed from the host in order 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.

See the [“Configuring the MSS Value and MTU for Transient TCP SYN Packets”](#) section on page 9 for configuration instructions.

TCP Applications Flags Enhancement

The TCP Applications Flags Enhancement feature enables the user to display additional flags with reference to TCP applications. There are two types of flags: status and option. The status flags indicate the status of TCP connections such as retransmission timeouts, application closed, and synchronized (SYNC) handshakes for listen. The additional flags indicate the state of set options such as whether or not a virtual private network (VPN) routing and forwarding (VRF) instance is set, whether or not a user is idle, and whether or not a keepalive timer is running. Use the **show tcp** command to display TCP application flags.

TCP Show Extension

The TCP Show Extension feature introduces the capability to display addresses in IP format instead of hostname format and to display the virtual private network (VPN) routing and forwarding (VRF) table associated with the connection. To display the status for all endpoints with the addresses in IP format, use the **show tcp brief numeric** command.

How to Configure TCP

This section contains the following procedures:

- [Configuring TCP Performance Parameters, page 7](#)
- [Configuring the MSS Value and MTU for Transient TCP SYN Packets, page 9](#)
- [Verifying TCP Performance Parameters, page 10](#)

Configuring TCP Performance Parameters

Perform the following task to configure TCP performance parameters.

Prerequisites

- Both sides of the link must be configured to support window scaling or the default of 65,535 bytes will apply as the maximum window size.
- To support ECN, the remote peer must be ECN-enabled because the ECN capability is negotiated during a three-way handshake with the remote peer.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **ip tcp synwait-time** *seconds*
4. **ip tcp path-mtu-discovery** [**age-timer** {*minutes* | **infinite**}]
5. **ip tcp selective-ack**
6. **ip tcp timestamp**
7. **ip tcp chunk-size** *characters*
8. **ip tcp window-size** *bytes*
9. **ip tcp ecn**
10. **ip tcp queuemax** *packets*

DETAILED STEPS

	Command or Action	Purpose
Step 1	enable Example: Router> enable	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal Example: Router# configure terminal	Enters global configuration mode.
Step 3	ip tcp synwait-time <i>seconds</i> Example: Router(config)# ip tcp synwait-time 60	(Optional) Sets the amount of time the Cisco IOS software will wait to attempt to establish a TCP connection. The default is 30 seconds.
Step 4	ip tcp path-mtu-discovery [age-timer { <i>minutes</i> infinite }] Example: Router(config)# ip tcp path-mtu-discovery age-timer 11	(Optional) Enables Path MTU Discovery. <ul style="list-style-type: none"> • age-timer—Time interval, in minutes, TCP reestimates the path MTU with a larger maximum segment size (MSS). The default is 10 minutes. The maximum is 30 minutes. • infinite—Disables the age timer.
Step 5	ip tcp selective-ack Example: Router(config)# ip tcp selective-ack	(Optional) Enables TCP selective acknowledgment.

	Command or Action	Purpose
Step 6	<code>ip tcp timestamp</code> Example: Router(config)# ip tcp timestamp	(Optional) Enables the TCP time stamp.
Step 7	<code>ip tcp chunk-size characters</code> Example: Router(config)# ip tcp chunk-size 64000	(Optional) Sets the TCP maximum read size for Telnet or rlogin. Note We do not recommend that you change this value.
Step 8	<code>ip tcp window-size bytes</code> Example: Router(config)# ip tcp window-size 75000	(Optional) Sets the TCP window size. The <i>bytes</i> argument can be set to an integer from 0 to 1073741823. To enable window scaling to support LFNs, the TCP window size must be more than 65535. The default window size is 4128 if window scaling is not configured.
Step 9	<code>ip tcp ecn</code> Example: Router(config)# ip tcp ecn	(Optional) Enables ECN for TCP.
Step 10	<code>ip tcp queuemax packets</code> Example: Router(config)# ip tcp queuemax 10	(Optional) Sets the TCP outgoing queue size.

Configuring the MSS Value and MTU for Transient TCP SYN Packets

Perform this task to configure the maximum segment size (MSS) for transient packets that traverse a router, specifically TCP segments with the SYN bit set, and to configure the MTU size of IP packets.

If you are configuring the `ip mtu` command on the same interface as the `ip tcp adjust-mss` command, we recommend that you use the following commands and values:

- `ip tcp adjust-mss 1452`
- `ip mtu 1492`

SUMMARY STEPS

1. `enable`
2. `configure terminal`
3. `interface type number`
4. `ip tcp adjust-mss max-segment-size`
5. `ip mtu bytes`
6. `end`

DETAILED STEPS

	Command or Action	Purpose
Step 1	enable Example: Router> enable	Enables privileged EXEC mode. <ul style="list-style-type: none"> Enter your password if prompted.
Step 2	configure terminal Example: Router# configure terminal	Enters global configuration mode.
Step 3	interface <i>type number</i> Example: Router(config)# interface ethernet0/1	Configures an interface type and enters interface configuration mode.
Step 4	ip tcp adjust-mss <i>max-segment-size</i> Example: Router(config-if)# ip tcp adjust-mss 1452	Adjusts the MSS value of TCP SYN packets going through a router. The <i>max-segment-size</i> argument is the maximum segment size, in bytes. The range is from 500 to 1460.
Step 5	ip mtu <i>bytes</i> Example: Router(config-if)# ip mtu 1492	Sets the MTU size of IP packets, in bytes, sent on an interface.
Step 6	end Example: Router(config-if)# end	Exits to global configuration mode.

Verifying TCP Performance Parameters

This task shows you how to verify configured TCP performance parameters.

SUMMARY STEPS

- show tcp** [*line-number*] [**tcb** *address*]
- show tcp brief** [**all** | **numeric**]
- debug ip tcp transactions**

DETAILED STEPS

Step 1 **show tcp** [*line-number*] [**tcb** *address*]

Displays the status of TCP connections. The arguments and keyword are as follows:

- line-number*—(Optional) Absolute line number of the Telnet connection status.
- tcb**—(Optional) Transmission control block (TCB) of the ECN-enabled connection.

- *address*—(Optional) TCB hexadecimal address. The valid range is from 0x0 to 0xFFFFFFFF.

The following is sample output from the **show tcp tcb** command that displays detailed information by hexadecimal address about an ECN-enabled connection:

```
Router# show tcp tcb 0x62CD2BB8

Connection state is LISTEN, I/O status: 1, unread input bytes: 0
Connection is ECN enabled
Local host: 10.10.10.1, Local port: 179
Foreign host: 10.10.10.2, Foreign port: 12000

Enqueued packets for retransmit: 0, input: 0 mis-ordered: 0 (0 bytes)

Event Timers (current time is 0x4F31940):
Timer           Starts    Wakeups           Next
Retrans         0         0                 0x0
TimeWait        0         0                 0x0
AckHold         0         0                 0x0
SendWnd         0         0                 0x0
KeepAlive       0         0                 0x0
GiveUp          0         0                 0x0
PmtuAger        0         0                 0x0
DeadWait        0         0                 0x0

irs:           0 snduna:         0 sndnxt:         0   sndwnd:         0
irs:           0 rcvnxt:         0 rcvwnd:         4128 delrcvwnd:    0

SRTT: 0 ms, RTTO: 2000 ms, RTV: 2000 ms, KRTT: 0 ms
minRTT: 60000 ms, maxRTT: 0 ms, ACK hold: 200 ms
Flags: passive open, higher precedence, retransmission timeout

TCB is waiting for TCP Process (67)

Datagrams (max data segment is 516 bytes):
Rcvd: 6 (out of order: 0), with data: 0, total data bytes: 0
Sent: 0 (retransmit: 0, fastretransmit: 0), with data: 0, total data
bytes: 0
```

Cisco IOS Software Modularity

The following is sample output from the **show tcp tcb** command from a Software Modularity image:

```
Router# show tcp tcb 0x1059C10

Connection state is ESTAB, I/O status: 0, unread input bytes: 0
Local host: 10.4.2.32, Local port: 23
Foreign host: 10.4.2.39, Foreign port: 11000
VRF table id is: 0

Current send queue size: 0 (max 65536)
Current receive queue size: 0 (max 32768) mis-ordered: 0 bytes

Event Timers (current time is 0xB9ACB9):
Timer           Starts    Wakeups           Next (msec)
Retrans         6         0                 0
SendWnd         0         0                 0
TimeWait        0         0                 0
AckHold         8         4                 0
KeepAlive       11        0                 7199992
PmtuAger        0         0                 0
GiveUp          0         0                 0
Throttle        0         0                 0

irs:   1633857851 rcvnxt: 1633857890 rcvad: 1633890620 rcvwnd: 32730
```

```

iss:      4231531315  snduna: 4231531392  sndnxt: 4231531392  sndwnd:  4052
sndmax: 4231531392  sndcwnd:  10220

SRTT: 84 ms,  RTTO: 650 ms,  RTV: 69 ms,  KRRT: 0 ms
minRRT: 0 ms,  maxRRT: 200 ms,  ACK hold: 200 ms

Keepalive time: 7200 sec, SYN wait time: 75 sec
Giveup time: 0 ms, Retransmission retries: 0, Retransmit forever: FALSE

State flags: none

Feature flags: Nagle

Request flags: none
Window scales: rcv 0, snd 0, request rcv 0, request snd 0
Timestamp option: recent 0, recent age 0, last ACK sent      0

Datagrams (in bytes): MSS 1460, peer MSS 1460, min MSS 1460, max MSS 1460
Rcvd: 14 (out of order: 0), with data: 10, total data bytes: 38
Sent: 10 (retransmit: 0, fastretransmit: 0), with data: 5, total data bytes: 76

Header prediction hit rate: 72 %

Socket states: SS_ISCONNECTED, SS_PRIV

Read buffer flags: SB_WAIT, SB_SEL, SB_DEL_WAKEUP
Read notifications: 4

Write buffer flags: SB_DEL_WAKEUP
Write notifications: 0
Socket status: 0

```

Step 2 **show tcp brief [all | numeric]**

(Optional) Displays addresses in IP format.

Use the **show tcp brief** command to display a concise description of TCP connection endpoints. The keywords are as follows. Use the optional **all** keyword to display the status for all endpoints with the addresses in a Domain Name System (DNS) hostname format. Without this keyword, endpoints in the LISTEN state are not shown. Use the optional **numeric** keyword to display the status for all endpoints with the addresses in IP format.

The following is sample output from the **show tcp brief** command while a user is connected to the system by using Telnet:

```

Router# show tcp brief

TCB          Local Address          Foreign Address        (state)
609789AC     Router.cisco.com.23   cider.cisco.com.3733  ESTAB

```

The following example shows the IP activity by using the **numeric** keyword to display the addresses in IP format.

```

Router# show tcp brief numeric

TCB          Local Address          Foreign Address        (state)
6523A4FC     10.1.25.3.11000      10.1.25.3.23         ESTAB
65239A84     10.1.25.3.23         10.1.25.3.11000     ESTAB
653FCBCC     *.1723 *.* LISTEN

```

Step 3 debug ip tcp transactions

Use the **debug ip tcp transactions** command to display information about significant TCP transactions such as state changes, retransmissions, and duplicate packets. This command is particularly useful for debugging a performance problem on a TCP/IP network that you have isolated above the data-link layer.

The following is sample output from the **debug ip tcp transactions** command:

```
Router# debug ip tcp transactions

TCP: sending SYN, seq 168108, ack 88655553
TCP0: Connection to 10.9.0.13:22530, advertising MSS 966
TCP0: state was LISTEN -> SYNRCVD [23 -> 10.9.0.13(22530)]
TCP0: state was SYNSENT -> SYNRCVD [23 -> 10.9.0.13(22530)]
TCP0: Connection to 10.9.0.13:22530, received MSS 956
TCP0: restart retransmission in 5996
TCP0: state was SYNRCVD -> ESTAB [23 -> 10.9.0.13(22530)]
TCP2: restart retransmission in 10689
TCP2: restart retransmission in 10641
TCP2: restart retransmission in 10633
TCP2: restart retransmission in 13384 -> 10.0.0.13(16151)]
TCP0: restart retransmission in 5996 [23 -> 10.0.0.13(16151)]
```

The following line from the **debug ip tcp transactions** command output shows that TCP has entered Fast Recovery mode:

```
fast re-transmit - sndcwnd - 512, snd_last - 33884268765
```

The following lines from the **debug ip tcp transactions** command output show that a duplicate acknowledgment is received when in Fast Recovery mode (first line) and a partial acknowledgment has been received (second line):

```
TCP0:ignoring second congestion in same window sndcwn - 512, snd_1st - 33884268765
TCP0:partial ACK received sndcwnd:338842495
```

Configuration Examples for TCP

This section provides the following configuration examples:

- [Verifying the Configuration of TCP ECN: Example, page 13](#)
- [Configuring the TCP MSS Adjustment: Examples, page 15](#)
- [Configuring the TCP Application Flags Enhancement: Example, page 17](#)
- [Displaying Addresses in IP Format: Example, page 17](#)

Verifying the Configuration of TCP ECN: Example

The following example shows how to verify that ECN is configured:

```
Router# show running-config

Building configuration...
.
.
.
ip tcp ecn ! ECN is configured.
.
.
```

The following example shows how to verify that TCP is ECN enabled on a specific connection (local host):

```
Router# show tcp tcb 123456A

!Local host
!
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Connection is ECN Enabled
Local host: 10.1.25.31, Local port: 11002
Foreign host: 10.1.25.34, Foreign port: 23
```

The following example shows how to display concise information about one address:

```
Router# show tcp brief

!
TCB           Local address           Foreign Address           (state)
609789C       Router.cisco.com.23      cider.cisco.com.3733     ESTAB
```

The following example show how to enable IP TCP ECN debugging:

```
Router# debug ip tcp ecn
!
TCP ECN debugging is on
!
Router# telnet 10.1.25.31

Trying 10.1.25.31 ...
!
01:43:19: 10.1.25.35:11000 <---> 10.1.25.31:23   out ECN-setup SYN
01:43:21: 10.1.25.35:11000 <---> 10.1.25.31:23   congestion window changes
01:43:21: cwnd from 1460 to 1460, ssthresh from 65535 to 2920
01:43:21: 10.1.25.35:11000 <---> 10.1.25.31:23   in non-ECN-setup SYN-ACK
```

Before a TCP connection can use ECN, a host sends an ECN-setup SYN (synchronization) packet to a remote end that contains an Echo Congestion Experience (ECE) and Congestion window reduced (CWR) bit set in the header. Setting the ECE and CWR bits indicates to the remote end that the sending TCP is ECN capable, rather than an indication of congestion. The remote end sends an ECN-setup SYN-ACK (acknowledgment) packet to the sending host.

In the example above, the “out ECN-setup SYN” text means that a SYN packet with the ECE and CWR bit set was sent to the remote end. The “in non-ECN-setup SYN-ACK” text means that the remote end did not favorably acknowledge the ECN request and, therefore, the session is not ECN capable.

The following debug output shows that ECN capabilities are enabled at both ends. In response to the ECN-setup SYN, the other end favorably replied with an ECN-setup SYN-ACK message. This connection is now ECN capable for the rest of the session.

```
Router# telnet 10.10.10.10

Trying 10.10.10.10 ... Open
Password required, but none set
!
1d20h: 10.1.25.34:11003 <---> 10.1.25.35:23   out ECN-setup SYN
1d20h: 10.1.25.34:11003 <---> 10.1.25.35:23   in ECN-setup SYN-ACK
```

The following example shows how to verify that the hosts are connected:


```
Router# show debugging
!
TCP:
  TCP Packet debugging is on
  TCP ECN debugging is on
```

```

!
Router# telnet 10.1.25.234
!
Trying 10.1.25.234 ...
!
00:02:48: 10.1.25.31:11001 <---> 10.1.25.234:23 out ECN-setup SYN
00:02:48: tcp0: O CLOSED 10.1.25.234:11001 10.1.25.31:23 seq 1922220018
          OPTS 4 ECE CWR SYN WIN 4128
00:02:50: 10.1.25.31:11001 <---> 10.1.25.234:23 congestion window changes
00:02:50: cwnd from 1460 to 1460, ssthresh from 65535 to 2920
00:02:50: tcp0: R SYNSENT 10.1.25.234:11001 10.1.25.31:23 seq 1922220018
          OPTS 4 ECE CWR SYN WIN 4128
00:02:54: 10.1.25.31:11001 <---> 10.1.25.234:23 congestion window changes
00:02:54: cwnd from 1460 to 1460, ssthresh from 2920 to 2920
00:02:54: tcp0: R SYNSENT 10.1.25.234:11001 10.1.25.31:23 seq 1922220018
          OPTS 4 ECE CWR SYN WIN 4128
00:03:02: 10.1.25.31:11001 <---> 10.1.25.234:23 congestion window changes
00:03:02: cwnd from 1460 to 1460, ssthresh from 2920 to 2920
00:03:02: tcp0: R SYNSENT 10.1.25.234:11001 10.1.25.31:23 seq 1922220018
          OPTS 4 ECE CWR SYN WIN 4128
00:03:18: 10.1.25.31:11001 <---> 10.1.25.234:23 SYN with ECN disabled
00:03:18: 10.1.25.31:11001 <---> 10.1.25.234:23 congestion window changes
00:03:18: cwnd from 1460 to 1460, ssthresh from 2920 to 2920
00:03:18: tcp0: O SYNSENT 10.1.25.234:11001 10.1.25.31:23 seq 1922220018
          OPTS 4 SYN WIN 4128
00:03:20: 10.1.25.31:11001 <---> 10.1.25.234:23 congestion window changes
00:03:20: cwnd from 1460 to 1460, ssthresh from 2920 to 2920
00:03:20: tcp0: R SYNSENT 10.1.25.234:11001 10.1.25.31:23 seq 1922220018
          OPTS 4 SYN WIN 4128
00:03:24: 10.1.25.31:11001 <---> 10.1.25.234:23 congestion window changes
00:03:24: cwnd from 1460 to 1460, ssthresh from 2920 to 2920
00:03:24: tcp0: R SYNSENT 10.1.25.234:11001 10.1.25.31:23 seq 1922220018
          OPTS 4 SYN WIN 4128
00:03:32: 10.1.25.31:11001 <---> 10.1.25.234:23 congestion window changes
00:03:32: cwnd from 1460 to 1460, ssthresh from 2920 to 2920
00:03:32: tcp0: R SYNSENT 10.1.25.234:11001 10.1.25.31:23 seq 1922220018
          OPTS 4 SYN WIN 4128
!Connection timed out; remote host not responding
    
```

Configuring the TCP MSS Adjustment: Examples

Figure 1 Example Topology for TCP MSS Adjustment



170888

The following example shows how to configure and verify the interface adjustment value. Configure the interface adjustment value on router B:

```
Router_B(config)# interface ethernet2/0
Router_B(config-if)# ip tcp adjust-mss 500
```

Telnet from router A to router C, with B having the MSS adjustment configured.

```
Router_A# telnet 192.168.1.1
Trying 192.168.1.1... Open
```

Observe the debug output from router C:

```
Router_C# debug ip tcp transactions

Sep 5 18:42:46.247: TCP0: state was LISTEN -> SYNRCVD [23 -> 10.0.1.1(38437)]
Sep 5 18:42:46.247: TCP: tcb 32290C0 connection to 10.0.1.1:38437, peer MSS 500, MSS is
500
Sep 5 18:42:46.247: TCP: sending SYN, seq 580539401, ack 6015751
Sep 5 18:42:46.247: TCP0: Connection to 10.0.1.1:38437, advertising MSS 500
Sep 5 18:42:46.251: TCP0: state was SYNRCVD -> ESTAB [23 -> 10.0.1.1(38437)]
```

The MSS gets adjusted to 500 on Router_B as configured.

The following example shows the configuration of a PPPoE client with the MSS value set to 1452:

```
vpdn enable
no vpdn logging
!
vpdn-group 1
request-dialin
protocol pppoe
!
interface Ethernet0
 ip address 192.168.100.1.255.255.255.0
 ip tcp adjust-mss 1452
 ip nat inside
!
interface ATM0
 no ip address
 no atm ilmi-keepalive
 pvc 8/35
 pppoe client dial-pool-number 1
!
dsl equipment-type CPE
dsl operating-mode GSHDSL symmetric annex B
dsl linerate AUTO
!
interface Dialer1
 ip address negotiated
 ip mtu 1492
 ip nat outside
 encapsulation ppp
 dialer pool 1
 dialer-group 1
 ppp authentication pap callin
 ppp pap sent-username sohodyn password 7 141B1309000528
!
ip nat inside source list 101 Dialer1 overload
ip route 0.0.0.0.0.0.0.0 Dialer1
access-list permit ip 192.168.100.0.0.0.0.255 any
```

Configuring the TCP Application Flags Enhancement: Example

The following output shows the flags (status and option) displayed using the **show tcp** command.

```
Router# show tcp
.
.
.
Status Flags: passive open, active open, retransmission timeout
App closed

Option Flags: vrf id set
IP Precedence value: 6
.
.
.
SRTT: 273 ms, RTTO: 490 ms, RTV: 217 ms, KRTT: 0 ms
minRTT: 0 ms, maxRTT: 300 ms, ACK hold: 200 ms
```

Displaying Addresses in IP Format: Example

The following example shows the IP activity by using the **numeric** keyword to display the addresses in IP format.

```
Router# show tcp brief numeric

TCB           Local Address      Foreign Address    (state)
6523A4FC      10.1.25.3.11000    10.1.25.3.23      ESTAB
65239A84      10.1.25.3.23      10.1.25.3.11000   ESTAB
653FCBBC      *.1723 *.* LISTEN
```

Additional References

The following sections provide references related to TCP.

Related Documents

Related Topic	Document Title
IP addressing and services configuration tasks	<i>Cisco IOS IP Addressing and Services Configuration Guide</i>
IP application services commands: complete command syntax, command mode, command history, defaults, usage guidelines, and examples	<i>Cisco IOS IP Application Services Command Reference.</i>
Path MTU Discovery	<i>Configuring IP Services</i>
TCP security features	<i>TCP Out-of-Order Packet Support for Cisco IOS Firewall and Cisco IOS IPS</i> <i>Configuring TCP Intercept (Preventing Denial-of-Service Attacks)</i>
TCP Header Compression, Class-based TCP Header Compression	<i>Configuring Class-Based RTP and TCP Header Compression</i> <i>Configuring TCP Header Compression</i>
Troubleshooting TCP	“ <i>Troubleshooting TCP/IP</i> ” part of the <i>Internetwork Troubleshooting Handbook</i>

Standards

Standard	Title
No new or modified standards are supported, and support for existing standards has not been modified.	—

MIBs

MIB	MIBs Link
CISCO-TCP-MIB	To locate and download MIBs for selected platforms, Cisco IOS releases, and feature sets, use Cisco MIB Locator found at the following URL: http://www.cisco.com/go/mibs

RFCs

RFC	Title
RFC 793	<i>Transmission Control Protocol</i>
RFC 1191	<i>Path MTU discovery</i>
RFC 1323	<i>TCP Extensions for High Performance</i>
RFC 2018	<i>TCP Selective Acknowledgment Options</i>
RFC 2581	<i>TCP Congestion Control</i>
RFC 3168	<i>The Addition of Explicit Congestion Notification (ECN) to IP</i>

RFC	Title
RFC 3782	<i>The NewReno Modification to TCP's Fast Recovery Algorithm</i>
RFC 4022	<i>Management Information Base for the Transmission Control Protocol (TCP)</i>

Technical Assistance

Description	Link
<p>The Cisco Support website provides extensive online resources, including documentation and tools for troubleshooting and resolving technical issues with Cisco products and technologies.</p> <p>To receive security and technical information about your products, you can subscribe to various services, such as the Product Alert Tool (accessed from Field Notices), the Cisco Technical Services Newsletter, and Really Simple Syndication (RSS) Feeds.</p> <p>Access to most tools on the Cisco Support website requires a Cisco.com user ID and password.</p>	<p>http://www.cisco.com/techsupport</p>

Feature Information for TCP

Table 1 lists the features in this module and provides links to specific configuration information. Only features that were introduced or modified in Cisco IOS Release 12.2(1) or a later release appear in the table.

For information on a feature in this technology that is not documented here, see the “[Cisco IOS IP Application Services Features Roadmap](#)” or the “[FHRP Features Roadmap](#).”

Not all commands may be available in your Cisco IOS software release. For release information about a specific command, see the command reference documentation.

Use Cisco Feature Navigator to find information about platform support and software image support. Cisco Feature Navigator enables you to determine which Cisco IOS and Catalyst OS software images support a specific software release, feature set, or platform. To access Cisco Feature Navigator, go to <http://www.cisco.com/go/cfn>. An account on Cisco.com is not required.



Note

Table 1 lists only the Cisco IOS software release that introduced support for a given feature in a given Cisco IOS software release train. Unless noted otherwise, subsequent releases of that Cisco IOS software release train also support that feature.

Table 1 **Feature Information for TCP**

Feature Name	Releases	Feature Information
TCP Application Flags Enhancement	12.4(2)T, 12.2(31)SB2 Cisco IOS	<p>The TCP Applications Flags Enhancement feature enables the user to display additional flags with reference to TCP applications. There are two types of flags: status and option. The status flags indicate the status of TCP connections; for example, retransmission timeouts, application closed, and synchronized (SYNC) handshakes for listen. The additional flags indicate the state of set options; for example, whether or not a virtual private network (VPN) routing and forwarding (VRF) identification is set, whether or not a user is idle, and whether or not a keepalive timer is running.</p> <p>The following sections contain information about this feature:</p> <ul style="list-style-type: none"> • TCP Applications Flags Enhancement, page 7 • Verifying TCP Performance Parameters, page 10 • Configuring the TCP Application Flags Enhancement: Example, page 17 <p>The following command was modified by this feature: show tcp.</p>

Table 1 Feature Information for TCP (continued)

Feature Name	Releases	Feature Information
TCP Congestion Avoidance	12.3(7)T	<p>The TCP Congestion Avoidance feature enables the monitoring of acknowledgment packets to the TCP sender when multiple packets are lost in a single window of data. Previously the sender would exit Fast-Recovery mode, wait for three or more duplicate acknowledgment packets before retransmitting the next unacknowledged packet, or wait for the retransmission timer to slow start. This could lead to performance issues.</p> <p>Implementation of RFC 2581 and RFC 3782 addresses the modifications to the Fast-Recovery algorithm that incorporates a response to partial acknowledgments received during Fast Recovery, improving performance in situations where multiple packets are lost in a single window of data.</p> <p>This feature is an enhancement to the existing Fast Recovery algorithm. There are no commands used to enable or disable this feature.</p> <p>To monitor the acknowledgment packets, the output of the debug ip tcp transactions command has been enhanced to show the following conditions:</p> <ul style="list-style-type: none"> • TCP entering Fast Recovery mode. • Duplicate acknowledgments being received during Fast Recovery mode. • Partial acknowledgments being received. <p>The following sections provide information about this feature:</p> <ul style="list-style-type: none"> • TCP Congestion Avoidance, page 6 • Verifying TCP Performance Parameters, page 10 <p>The following command was modified by this feature: debug ip tcp transactions.</p>

Table 1 Feature Information for TCP (continued)

Feature Name	Releases	Feature Information
TCP Explicit Congestion Notification	12.3(7)T 12.2(31)SB2	<p>The TCP Explicit Congestion Notification (ECN) feature provides a method for an intermediate router to notify the end hosts of impending network congestion. It also provides enhanced support for TCP sessions associated with applications that are sensitive to delay or packet loss including Telnet, web browsing, and transfer of audio and video data. The benefit of this feature is the reduction of delay and packet loss in data transmissions.</p> <p>The following sections provide information about this feature:</p> <ul style="list-style-type: none"> • TCP Explicit Congestion Notification, page 6 • Configuring TCP Performance Parameters, page 7 • Verifying TCP Performance Parameters, page 10 • Verifying the Configuration of TCP ECN: Example, page 13 <p>The following commands were introduced or modified by this feature: debug ip tcp ecn, ip tcp ecn, show debugging, show tcp.</p>
TCP MSS Adjust	12.2(4)T 12.2(8)T 12.2(28)SB 12.2(33)SRA 12.2(18)ZU2 12.2(33)SXH	<p>The TCP MSS Adjust feature enables the configuration of the maximum segment size (MSS) for transient packets that traverse a router, specifically TCP segments in the SYN bit set.</p> <p>In 12.2(4)T, this feature was introduced.</p> <p>In 12.2(8)T, the command that was introduced by this feature was changed from ip adjust-mss to ip tcp adjust-mss.</p> <p>In 12.2(28)SB and 12.2(33)SRA, this feature was enhanced to be configurable on subinterfaces.</p> <p>The following sections provide information about this feature:</p> <ul style="list-style-type: none"> • TCP MSS Adjustment, page 6 • Configuring the MSS Value and MTU for Transient TCP SYN Packets, page 9 • Configuring the TCP MSS Adjustment: Examples, page 15 <p>The following command was introduced by this feature: ip tcp adjust-mss.</p>

Table 1 Feature Information for TCP (continued)

Feature Name	Releases	Feature Information
TCP Show Extension	12.4(2)T 12.2(31)SB2	<p>The TCP Show Extension feature introduces the capability to display addresses in IP format instead of hostname format and to display the virtual private network (VPN) routing and forwarding (VRF) table associated with the connection.</p> <p>The following sections contain information about this feature:</p> <ul style="list-style-type: none"> • TCP Show Extension, page 7 • Verifying TCP Performance Parameters, page 10 • Displaying Addresses in IP Format: Example, page 17 <p>The following command was modified by this feature: show tcp brief.</p>
TCP Window Scaling	12.2(8)T 12.2(31)SB2	<p>The TCP Window Scaling feature adds support for the Window Scaling option in RFC 1323. A larger window size is recommended to improve TCP performance in network paths with large bandwidth, long-delay characteristics that are called Long Fat Networks (LFNs). This TCP Window Scaling enhancement provides that support.</p> <p>The following sections provide information about this feature:</p> <ul style="list-style-type: none"> • TCP Window Scaling, page 5 • Configuring TCP Performance Parameters, page 7 • Verifying TCP Performance Parameters, page 10 <p>The following commands were introduced or modified by this feature: ip tcp window-size.</p>

Glossary

LFN—Long Fat Networks. Large bandwidth, long-delay networks where the throughput is high and the transmission distance is long. Networks with satellite connections are one example of an LFN. Satellite links always have high propagation delays and typically have high bandwidth.

TCP—Transmission Control Protocol. Connection-oriented transport layer protocol that provides reliable full-duplex data transmission. TCP is part of the TCP/IP protocol stack.

CCDE, CCENT, CCSI, Cisco Eos, Cisco HealthPresence, Cisco IronPort, the Cisco logo, Cisco Lumin, Cisco Nexus, Cisco Nurse Connect, Cisco Pulse, Cisco StackPower, Cisco StadiumVision, Cisco TelePresence, Cisco Unified Computing System, Cisco WebEx, DCE, Flip Channels, Flip for Good, Flip Mino, Flipshare (Design), Flip Ultra, Flip Video, Flip Video (Design), Instant Broadband, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn, Cisco Capital, Cisco Capital (Design), Cisco:Financed (Stylized), Cisco Store, and Flip Gift Card are service marks; and Access Registrar, Aironet, AllTouch, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, Continuum, EtherFast, EtherSwitch, Event Center, Explorer, Fast Step, Follow Me Browsing, FormShare, GainMaker, GigaDrive, HomeLink, iLYNX, Internet Quotient, IOS, iPhone, iQuick Study, IronPort, the IronPort logo, Laser Link, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerKEY, PowerPanels, PowerTV, PowerTV (Design), PowerVu, Prisma, ProConnect, ROSA, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0908R)

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

© 2008 Cisco Systems, Inc. All rights reserved.