

# CSNA Tuning Guide

Document ID: 17534

---

## **Introduction**

### **Before You Begin**

Conventions

Prerequisites

Components Used

### **Problem Identification**

Determine If Router Is Fast Switching or Process Switching

Identify Problem When CIP Is Process Switching Packets

Identify Problems When CIP Is Fast Switching Packets

### **LLC2 Tuning Parameters**

### **Conclusion**

### **Related Information**

---

## **Introduction**

With the introduction of Cisco Systems Network Architecture (CSNA) on the Channel Interface Processor (CIP), Cisco 7000/7500 routers are required to provide Systems Network Architecture (SNA) support for more connections, sessions, and data traffic than ever before.

Therefore, the router may not be capable of providing the desired SNA support with the default tuning parameters. Because the CSNA feature can switch packets from the CIP to the Routing Processor (RP) at a high speed, it's possible that packets will be dropped using the default tuning parameters. The router can adjust and recover from packet drops if the frequency is low; however, session loss may occur when the frequency of drops is high.

The purpose of this document is to:

- Help CSNA users identify problem areas with CSNA performance or session loss, or both.
- Provide guidelines for tuning router parameters and identifying problems.

## **Before You Begin**

### **Conventions**

For more information on document conventions, see the Cisco Technical Tips Conventions.

### **Prerequisites**

There are no specific prerequisites for this document.

### **Components Used**

This document is not restricted to specific software and hardware versions.

# Problem Identification

If you are experiencing problems with inadequate performance or session loss, or both, in the CSNA environment, you can examine trouble areas within the router to determine if tuning can help.

## Determine If Router Is Fast Switching or Process Switching

Problems within the CSNA router can be isolated depending on the type of packet switching with which the CIP is primarily involved. The CIP supports two types of switching:

- Fast switching.
- Process switching.

Fast switching allows the CIP to switch packets directly to another interface processor without queuing the packet to the RP. Process switching is slower because the CIP must copy the packet into a buffer and queue the buffer for processing by the RP. It is necessary to isolate problems differently depending on whether the CIP is fast switching or process switching. Often, it can be determined whether the CIP is fast switching or process switching packets based on the router features configured and used.

Several router features are always process switched:

- Advanced Peer-to-Peer Networking (APPN)
- Data-Link Switching Plus (DLSw+)/TCP
- Downstream Physical Unit (DSPU)/RSRB
- Remote Source-Route Bridging (RSRB)/TCP
- RSRB/TCP with local-ack
- Synchronous Data Link Control Logical Link Control (SDLLC)

Other router features are always fast switched:

- DLSw+/direct
- DLSw+/Fast-Sequenced Transport (FST)
- RSRB/direct
- RSRB/FST
- Source-Route Bridging (SRB)
- Source-Route Translational Bridging (SR/TLB) (as of IOS 11.2)

The definitive method for determining the CIP packet switching type is to issue the following command:

```
show interface Channel x/2 stats
```

Example.

```
dspu-7k#show interface channel 4/2 stats
```

```
Channel4/2
Switching path  Pkts In  Chars In  Pkts Out  Chars Out
Processor      35126   1270810   35110     1108845
Route cache         0         0         0         0
Autonomous/SSE    0         0         0         0
Total          35126   1270810   35110     1108845
```

In the **show interface channel** command output:

- The **Processor** counts indicate the number of process-switched packets on the interface.

- The **Route cache** counts indicate the number of fast-switched packets on the interface.
- The **Autonomous/SSE** counts indicate the number of autonomously-switched packets on the interface.

**Note:** CSNA does not support autonomous packet switching.

## Identify Problem When CIP Is Process Switching Packets

Problems with performance or session loss, or both, are most apparent when process switching large numbers of sessions or high data transfer rates, or both.

There are several problem areas in the process switch path that are most common:

1. Buffer misses/failures
2. Input queue overflow on an interface
3. Output queue overflow on an interface

Each of these areas is discussed in a separate document.

## Identify Problems When CIP Is Fast Switching Packets

Problems with performance or session loss, or both, when the CIP is fast switching should be rare, but may occur under certain conditions. When packets are dropped during fast switching, the **show interface** command output shows increasing numbers of output queue drops that cannot be accounted for by process switching drops.

You can tune the TX queue limit for the interface using the **tx-queue-limit** interface configuration command. The **show controller cxbus** command output shows the current TX Queue Limit for each interface.

Example.

```
dspu-7k#show controller cxbus
```

```
Switch Processor 5, hardware version 11.1, microcode version 10.12
  Microcode loaded from system
  512 Kbytes of main memory, 128 Kbytes cache memory
  16 256 byte buffers, 4 1024 byte buffers, 51 1520 byte buffers, 38 2068 byte buffers, 73
  Restarts: 0 line down, 0 hung output, 0 controller error
CIP 4, hardware version 4.1, microcode version 180.40
  Microcode loaded from flash nightly/cip180-40.OffCTA_950922
  Controller Sync: 5 timeouts, 5 resyncs 0 failures, 0 max phase count
  EPROM version 165.0, VPLD version 4.28
  CPU utilization 39%, sram 848/512K, dram 59542704/64M
Interface 32 - Channel 4/0
  10 buffer RX queue threshold, 18 buffer TX queue limit, buffer size 4496
  ift 0007, rql 5, tq 0000 0000, tql 18
  Transmitter delay is 0 microseconds
Interface 33 - Channel 4/1
  10 buffer RX queue threshold, 18 buffer TX queue limit, buffer size 4496
  ift 0007, rql 6, tq 0000 0000, tql 18
  Transmitter delay is 0 microseconds
Interface 34 - Channel 4/2
  10 buffer RX queue threshold, 18 buffer TX queue limit, buffer size 4496
  ift 0007, rql 5, tq 0000 06D8, tql 18
  Transmitter delay on all interfaces is 0 flags
```

**Note:** All packets (explorers) must be process switched at least once before fast switching takes over.

**Note:** If the nature of CSNA problems is with establishing connections, consider isolating problems with process switching rather than fast switching

## LLC2 Tuning Parameters

The following Logical Link Control, type 2 (LLC2) session characteristics can be tuned using the **llc2** configuration command:

Command	Description
<b>N2</b> (default: eight retries)	<p>The number of times the local LLC2 should retry various operations, including sending unacknowledged frames or retrying poll of busy workstation.</p> <p>When N2 retries are attempted without success, the LLC2 connection is terminated.</p>
<b>ack-delay-time</b> (default: 100 ms)	<p>The maximum amount of time the local LLC2 allows incoming I-frames to remain unacknowledged. Normally, the local LLC2 station continues to receive I-frames.</p> <p>Once the <b>ack-max</b> number of I-frames has been received, the local LLC2 station acknowledges all frames at once with a single acknowledgement.</p> <p>However, there may be situations where the remote LLC2 station will not send any additional I-frames until an acknowledgment is received from the local LLC2 station.</p> <p>In this scenario, the <b>ack-max</b> number of I-frames will never be received and acknowledgment for received I-frames is never sent. However, the LLC2 station avoids this deadlock condition by starting a timer upon receipt of each I-frame.</p> <p>If the <b>ack-delay</b> timer expires for an I-frame that remains unacknowledged, the LLC2 station sends an acknowledgment of all received I-frames even though the <b>ack-max</b> number of frames were not received.</p>
<b>ack-max</b> (default: three packets)	<p>The maximum number of I-frames received before an acknowledgment must be sent.</p>

	<p>The <b>ack-max</b> is essentially the window of I-frames received that may remain unacknowledged.</p> <p>Increasing the <b>ack-max</b> may increase throughput for the LLC2 sessions because more frames can be received and processed before stopping to send the acknowledgement. In addition, more I-frames received can be acknowledged with a single acknowledgment.</p> <p>However, increasing the <b>ack-max</b> for the receiver LLC2 station, without increasing the <b>local-window</b> of the sender LLC2 station, may result in lower performance because of idle time where the <b>ack-delay</b> timer must expire before acknowledgment is sent to re-start transmission of I-frames from remote LLC2 station.</p> <p>The <b>ack-max</b> is the receiver LLC2 station equivalent of the <b>local-window</b> on the sender LLC2 station.</p>
<b>dynwind</b>	The congestion control with dynamic window.
<b>idle-time</b> (default: 60000 ms for CIP)	<p>The frequency of polls during periods of idle traffic.</p> <p>Decreasing the <b>idle-time</b> may help jump-start problematic LLC2 sessions for false idles. Increasing the <b>idle-time</b> decreases session traffic during true session idle.</p>
<b>local-window</b> (default: seven packets)	<p>The maximum number of I-frames to send before waiting for an acknowledgment. The <b>local-window</b> is essentially the window of I-frames sent that may remain unacknowledged.</p> <p>Increasing the <b>local-window</b> may increase throughput for the LLC2 sessions because more frames can be sent to the receiver before stopping to wait for the acknowledgment. In addition, the receiver can acknowledge more I-frames with a single acknowledgment.</p> <p>However, decreasing the <b>local-window</b> for the sending LLC2 station, without decreasing the <b>ack-max</b> of the receiver</p>

	<p>LLC2 station, may result in lower performance because of idle time waiting for the receiver's <b>ack-delay</b> timer to expire before acknowledgment is received to re-start transmission of I-frames from sender LLC2 station.</p> <p>The <b>local-window</b> is the sender LLC2 station equivalent of the <b>ack-max</b> on the receiver LLC2 station.</p>
<p><b>t1-time</b> (default: 1000 ms)</p>	<p>The length of time LLC2 waits for an acknowledgment to transmitted I-frames.</p> <p>Once the local LLC2 station has transmitted up to its <b>local-window</b> number of I-frames, an acknowledgment of these I-frames is expected from the remote LLC2 station. If the T1 timer expires before this acknowledgment is received, the local LLC2 station retransmits these I-frames.</p> <p>Increasing the <b>t1-time</b> helps the LLC2 stations avoid unnecessary re-transmissions of I-frames where latency times are such that acknowledgment may have been delayed (but not dropped) in the network.</p>
<p><b>tbusy-time</b> (default: 9600 ms)</p>	<p>The length of time the local LLC2 waits after the remote LLC2 station has indicated "busy" state before attempting to poll state again.</p> <p>If the TBUSY timer expires before the remote LLC2 station has indicated a clear state, the local LLC2 station will poll the remote again for current status. Decreasing the <b>tbusy-time</b> may help jump-start LLC2 sessions out of busy state sooner.</p> <p>Increasing the <b>tbusy-time</b> will decrease session traffic but remote LLC2 stations may take longer to recover from a busy state.</p>
<p><b>tpf-time</b> (default: 1000 ms)</p>	<p>The length of time the local LLC2 waits for a final response to a poll before re-sending the poll. If TPF timer expires before the remote LLC2 station has responded to a poll, the local LLC2 station re-transmits its poll.</p>

	<p>Increasing the <b>tpf-time</b> helps LLC2 stations avoid unnecessary re-transmissions of polls where latency times are such that the responses may have been delayed (but not dropped) in the network.</p>
<p><b>trej-time</b> (default: 3200 ms)</p>	<p>The length of time local LLC2 waits for a resend of a rejected frame before sending the reject command.</p> <p>If TREJ timer expires before the remote LLC2 station has re-sent a rejected I-frame, the local LLC2 station re-transmits its reject.</p> <p>Increasing the <b>trej-time</b> helps LLC2 stations avoid unnecessary re-transmissions of rejects where latency times are such that the re-transmitted I-frames may have been delayed (but not dropped) in the network.</p>
<p><b>txq-max</b></p>	<p>The queue for holding <b>llc2</b> information frames.</p> <p>The <b>txq-max</b> is the maximum size of the transmit queue for an LLC2 station. If the number of frames to be transmitted exceeds the current number of frames on the LLC2 TXQ, the additional frames are dropped.</p> <p>Increase the <b>txq-max</b> when the local LLC2 station must consistently transmit more frames than the LLC2 TXQ allows. For example, if the <b>local-window</b> is 128 and the <b>txq-max</b> is 100, 128 I-frames may likely be transmitted by the local LLC2, but 28 frames are dropped because of queue overflow. The <b>txq-max</b> should be increased to at least 128.</p>
<p><b>xid-neg-val-time</b> (default: 0 ms)</p>	<p>The frequency of Exchange of</p>
<p><b>xid-retry-time</b> (default: 60000 ms)</p>	<p>Identification (XID).</p> <p>How long router waits for reply to XID.</p>

**Note:** Each adapter under the CSNA internal LAN is represented by a separate LLC2 that may be tuned independently from the other adapters.

Example.

```
!
interface Channel4/2
```

```
no ip address
no keepalive
max-llc2-sessions 1000
lan TokenRing 0
source-bridge 666 1 99
adapter 0 4000.b0ca.7000
llc2 t1-time 3000
adapter 1 4000.b0ca.7001
llc2 t1-time 5000
adapter 2 4000.b0ca.7002
llc2 t1-time 10000
```

**Note:** When the RSRB/TCP local-ack feature is configured, the additional LLC2 representing LAK may also require tuning at the channel virtual interface level.

Example.

```
!
interface Channel4/2
no ip address
no keepalive
llc2 t1-time 5000

!--- Tuning for RSRB/TCP local-ack LLC2.

max-llc2-sessions 1000
lan TokenRing 0
source-bridge 666 1 99
adapter 0 4000.b0ca.7000
llc2 t1-time 3000

!--- Tuning for CSNA adapter LLC2.
```

## Conclusion

The first step of router tuning is to identify the problem area and determine what to tune. Using snapshots of the router's **show** command output helps to highlight the problem areas.

A single snapshot may identify a particular problem area from the past that overshadows the current problem area.

Example.

In the following **show buffer** output, it may seem that the primary problem area is **Big** buffers:

```
dspu-7k#show buffer

Buffer elements:
  500 in free list (500 max allowed)
  234 hits, 0 misses, 0 created

Public buffer pools:
Small buffers, 104 bytes (total 25, permanent 10):
  5 in free list (5 min, 10 max allowed)
  1359 hits, 28 misses, 17 trims, 32 created
  28 failures (0 no memory)
Middle buffers, 600 bytes (total 90, permanent 90):
  90 in free list (10 min, 200 max allowed)
  171 hits, 0 misses, 0 trims, 0 created
  0 failures (0 no memory)
```

```
Big buffers, 1524 bytes (total 341, permanent 90):
  120 in free list (5 min, 300 max allowed)
  113128 hits, 1323 misses, 959 trims, 1210 created
  1318 failures (0 no memory)
VeryBig buffers, 4520 bytes (total 10, permanent 10):
  10 in free list (0 min, 300 max allowed)
  20 hits, 0 misses, 0 trims, 0 created
  0 failures (0 no memory)
Large buffers, 5024 bytes (total 10, permanent 10):
  10 in free list (0 min, 30 max allowed)
  0 hits, 0 misses, 0 trims, 0 created
  0 failures (0 no memory)
Huge buffers, 18024 bytes (total 2, permanent 0):
  0 in free list (0 min, 13 max allowed)
  2 hits, 2 misses, 0 trims, 2 created
  0 failures (0 no memory)
```

The following output shows that the **Big** buffers are stable and the **Small** buffer pool is the current problem area:

```
dspu-7k#show buffer
```

```
Buffer elements:
  500 in free list (500 max allowed)
  234 hits, 0 misses, 0 created

Public buffer pools:
Small buffers, 104 bytes (total 42, permanent 10):
  7 in free list (5 min, 10 max allowed)
  1576 hits, 35 misses, 17 trims, 39 created
  34 failures (0 no memory)
Middle buffers, 600 bytes (total 90, permanent 90):
  90 in free list (10 min, 200 max allowed)
  171 hits, 0 misses, 0 trims, 0 created
  0 failures (0 no memory)
Big buffers, 1524 bytes (total 341, permanent 90):
  120 in free list (5 min, 300 max allowed)
  113128 hits, 1323 misses, 959 trims, 1210 created
  1318 failures (0 no memory)
VeryBig buffers, 4520 bytes (total 10, permanent 10):
  10 in free list (0 min, 300 max allowed)
  20 hits, 0 misses, 0 trims, 0 created
  0 failures (0 no memory)
Large buffers, 5024 bytes (total 10, permanent 10):
  10 in free list (0 min, 30 max allowed)
  0 hits, 0 misses, 0 trims, 0 created
  0 failures (0 no memory)
Huge buffers, 18024 bytes (total 2, permanent 0):
  0 in free list (0 min, 13 max allowed)
  2 hits, 2 misses, 0 trims, 2 created
  0 failures (0 no memory)
```

---

## Related Information

- [Technical Support – Cisco Systems](#)

---

[Contacts & Feedback](#) | [Help](#) | [Site Map](#)

© 2008 – 2009 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

---

