Buffer Tuning for all Cisco Routers

Contents

Introduction

Prerequisites

Requirements

Components Used

Conventions

General Overview

Low-End Platforms (Cisco 1600, 2500, and 4000 Series Routers)

<u>High-End Platforms (Route Processors, Switch Processors, Silicon Switch Processors, and Route/Switch Processors)</u>

Particle-based Platforms

Buffer Tuning

Reserving Fabric Buffer

Buffer Leaks

Related Information

Introduction

This document describes an overview of buffer tuning based on current platforms, and gives general information about the **show buffers** command.

Prerequisites

Requirements

There are no specific requirements for this document.

Components Used

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

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

Conventions

For more information on document conventions, see the <u>Cisco Technical Tips Conventions</u>.

General Overview

Buffer tuning allows you to modify the way a router allocates buffers from its available memory, and helps prevent packet drops during a temporary burst of traffic.

To determine whether your router needs to have its buffer tuned, use the **show interfaces** and **show buffers** commands.

If you have the output of the **show interfaces** and **show buffers** commands, or the output of the **show technical-support** (from enable mode) from your Cisco device, you can use <u>Cisco CLI Analyzer</u> to display potential issues and fixes. To use <u>Cisco CLI Analyzer</u>, you must be a <u>registered</u> customer, be logged in, and have JavaScript enabled.

Below is an example of the **show interfaces** command output:

```
<#root>
```

```
Output queue 0/40, 1041 drops; input queue 0/75, 765 drops 35252345 packets input, 547082589 bytes,
```

940 no buffer

- The input and output drops are due to the input and output queues being overrun by a burst of traffic. This is not related to a buffer problem, but rather to a process switching performance limitation.
- "No buffer" represents the number of packets dropped because there is no free buffer to copy the packet.

Using the **show buffers** command, look at the buffer size corresponding to the maximum transmission unit (MTU) of the interface:

```
Middle buffers, 600 bytes (total 150, permanent 25):
147 in free list (10 min, 150 max allowed)
61351931 hits, 137912 misses, 51605 trims, 51730 created
91652 failures (0 no memory)
```

The table below explains the output:

Keyword	Description
total	The total number of buffers in the pool, including both used and unused buffers.
permanent	The permanent number of allocated buffers in the pool. These buffers are always in the pool, and cannot be trimmed away.
in free list	The number of buffers currently available in the pool that are free for use.
min	The minimum number of buffers that the router should attempt to keep "in free list." If the number of buffers "in free list" falls below the "min" value, the router should attempt to create more buffers for that pool.
max allowed	The maximum number of buffers allowed "in free list." If the number of buffers "in free list" is greater than the "max allowed" value, the router should attempt to trim buffers from the pool.
hits	The number of buffers successfully allocated from the free list.
misses	The number of times a buffer has been requested, but no buffers are available in the free list, or when there are fewer than "min" buffers in the free list.
trims	The number of buffers that have been trimmed from the pool when the number of buffers "in free list" exceeds the number of "max allowed" buffers.

created	The number of buffers that have been created in the pool when the number of buffers "in free list" was less than "min."
no memory	The number of times the router tried to create new buffers, but could not due to insufficient free memory in the router.
Tailures	The number of failures to grant a buffer to a requester under interrupt time (remember that the router can create new buffers at process switching level, so "failure" does not occur unless there is "no memory"). The number of "failures" represents the number of packets that have been dropped due to buffer shortage.

How Buffers Are Handled by the Router

The number of buffers "in free list" is the number of available buffers. When a buffer request comes in, a buffer from the "in free list" is allocated.

The Cisco IOS® buffers are used for two major reasons:

- To handle traffic terminating at the router.
- When packets are being process switched.

If there are no buffers available, and fast switching is enabled, there is a buffer failure and the packet is dropped. When the buffer pool manager process detects a buffer failure, it "creates" a new buffer to avoid future failures.

The router does not create a new buffer if the number "in free list" equals the "max allowed" value. If there is not enough memory in the router to create a new buffer, this is recorded as "no memory". If the number "in free list" is greater than the "max allowed" number, the router "trims" some excess buffers.

The number of "failures" and "no memory" are the only areas of concern. Failures may occur, but these should stabilize after a while. The router creates or trims buffers as necessary to stabilize the number of failures. If the number of failures continues to increase, buffer tuning might be necessary.

If there is not enough memory to create new buffers, look for a <u>buffer leak</u>, or a more <u>general memory problem</u>. Buffers are not created in the fast-switching path, so if the router tries to fast-switch a packet and there is no buffer available, the packet is dropped, and a failure is reported. A new buffer is created the next time the buffer pool manager is run.

Low-End Platforms (Cisco 1600, 2500, and 4000 Series Routers)

Fast-switched packets and process-switched packets share the same buffers. These buffers are located in the shared memory. The shared memory is located in dynamic RAM (DRAM) in Cisco 1600 and 2500 Series Routers, or in shared RAM (SRAM) for the Cisco 4000, 4500, and 4700 Series Routers.

The first lines of the **show memory** command tell you how much shared memory you have, how much is currently used, and its lowest point. When a packet cannot be fast-switched, a pointer to the packet is inserted in the process switching input queue, but the packet itself is not copied.

Here is the output of the show buffers command on a low-end platform (Cis	sco 4500)
--	-----------

<#root>

router#

show buffers

```
Buffer elements:
     471 in free list (500 max allowed)
     870696495 hits, 0 misses, 0 created
Public buffer pools:
Small buffers, 104 bytes (total 50, permanent 50):
     49 in free list (20 min, 150 max allowed)
     27301678 hits, 23 misses, 20 trims, 20 created
     0 failures (0 no memory)
Middle buffers, 600 bytes (total 150, permanent 25):
     147 in free list (10 min, 150 max allowed)
     61351931 hits, 137912 misses, 51605 trims, 51730 created
     91652 failures (0 no memory)
Big buffers, 1524 bytes (total 67, permanent 50):
     67 in free list (5 min, 150 max allowed)
     46293638 hits, 455 misses, 878 trims, 895 created
     0 failures (0 no memory)
VeryBig buffers, 4520 bytes (total 96, permanent 10):
     79 in free list (0 min, 100 max allowed)
     11818351 hits, 246 misses, 98 trims, 184 created
     243 failures (0 no memory)
Large buffers, 5024 bytes (total 10, permanent 0):
     10 in free list (0 min, 10 max allowed)
     4504003 hits, 873040 misses, 759543 trims, 759553 created
     873040 failures (0 no memory)
Huge buffers, 18024 bytes (total 0, permanent 0):
     0 in free list (0 min, 4 max allowed)
     O hits, O misses, O trims, O created
     0 failures (0 no memory)
Interface buffer pools:
TokenRingO buffers, 4516 bytes (total 48, permanent 48):
     0 in free list (0 min, 48 max allowed)
     3099742 hits, 9180771 fallbacks
     16 max cache size, 1 in cache
TokenRing1 buffers, 4516 bytes (total 48, permanent 48):
     0 in free list (0 min, 48 max allowed)
     335172 hits, 403668 fallbacks
     16 max cache size, 16 in cache
Serial1 buffers, 1524 bytes (total 96, permanent 96):
     63 in free list (0 min, 96 max allowed)
     33 hits, 0 fallbacks
     0 max cache size, 0 in cache
Serial2 buffers, 1524 bytes (total 96, permanent 96):
     63 in free list (0 min, 96 max allowed)
     701370936 hits, 268 fallbacks
     0 max cache size, 0 in cache
Serial3 buffers, 1524 bytes (total 96, permanent 96):
     63 in free list (0 min, 96 max allowed)
     33 hits, 0 fallbacks
     0 max cache size, 0 in cache
SerialO buffers, 4546 bytes (total 96, permanent 96):
     28 in free list (0 min, 96 max allowed)
     346854 hits, 5377043 fallbacks
     32 max cache size, 27 in cache
```

The interface buffer pools are used by the interfaces for input/output (I/O). When there are no more buffers in the interface buffer free list, the router goes to the public buffer pools as a fallback. There is no performance hit for a fallback.

Caching is a software manipulation that speeds up the availability of buffers for interrupt level driver code by bypassing some overhead.

Note: Normally, interface buffers should not be tuned.

High-End Platforms (Route Processors, Switch Processors, Silicon Switch Processors, and Route/Switch Processors)

Here is the output of the **show buffers** command on a high-end platform:

```
<#root>
Router#
show buffers
Buffer elements:
 498 in free list (500 max allowed)
 326504974 hits, 0 misses, 0 created
Public buffer pools:
Small buffers, 104 bytes (total 150, permanent 150):
 140 in free list (30 min, 250 max allowed)
  564556247 hits, 148477066 misses, 16239797 trims, 16239797 created
 29356200 failures (0 no memory)
Middle buffers, 600 bytes (total 120, permanent 120):
 116 in free list (20 min, 200 max allowed)
 319750574 hits, 85689239 misses, 9671164 trims, 9671164 created
 26050704 failures (0 no memory)
Big buffers, 1524 bytes (total 100, permanent 100):
 98 in free list (10 min, 300 max allowed)
 20130595 hits, 14796572 misses, 251916 trims, 251916 created
 11813639 failures (0 no memory)
VeryBig buffers, 4520 bytes (total 15, permanent 15):
  14 in free list (5 min, 300 max allowed)
 22966334 hits, 3477687 misses, 13113 trims, 13113 created
 2840089 failures (0 no memory)
Large buffers, 5024 bytes (total 12, permanent 12):
 12 in free list (0 min, 30 max allowed)
 849034 hits, 1979463 misses, 1028 trims, 1028 created
 1979456 failures (0 no memory)
Huge buffers, 18024 bytes (total 6, permanent 5):
 4 in free list (2 min, 13 max allowed)
 338440 hits, 1693496 misses, 1582 trims, 1583 created
 1640218 failures (0 no memory)
```

Header Pools

The public buffer pools are located in DRAM, and are called system buffers. The shared memory on the Route/Switch Processor (RSP) is called system packet memory (MEMD), and allows 2 MB of memory. On the Route Processor (RP) and the Switch Processor (SP) (or Silicon Switch Processor - SSP), system buffers are located on the RP, and the MEMD is located on the SP (or SSP).

When a packet comes in, it is saved in MEMD by the receiving interface processor (except in the case of the

Versatile Interface Processor - VIP). If it cannot be fast-switched, the whole packet is copied into a system buffer in DRAM. Therefore, the buffers you see in the **show buffers** command are the system buffers located in DRAM.

The <u>show controllers cbus</u> command shows you the interface buffers in MEMD. Again, it is not advised to tune the interface buffers. When a packet cannot be fast-switched, and is copied to a system buffer, the packet is dropped, and a failure is counted if there is no system buffer available.

Particle-based Platforms

The Cisco 3600 and 7200 Series Routers use particles. The interface buffers are atomic buffers, called particles, into which the packets are split. When a packet cannot be fast-switched, the router has to reassemble it in one system buffer, because the process switching code cannot handle particles.

Below is the output of the **show buffers** command on a Cisco 3600:

```
<#root>
Router#
show buffers
Buffer elements:
     499 in free list (500 max allowed)
     136440 hits, 0 misses, 0 created
Public buffer pools:
Small buffers, 104 bytes (total 50, permanent 50):
     49 in free list (20 min, 150 max allowed)
     4069435 hits, 141 misses, 73 trims, 73 created
     52 failures (0 no memory)
Middle buffers, 600 bytes (total 25, permanent 25):
     25 in free list (10 min, 150 max allowed)
     628629 hits, 21 misses, 21 trims, 21 created
     3 failures (0 no memory)
Big buffers, 1524 bytes (total 50, permanent 50):
     50 in free list (5 min, 150 max allowed)
     9145 hits, 0 misses, 0 trims, 0 created
     0 failures (0 no memory)
VeryBig buffers, 4520 bytes (total 10, permanent 10):
     10 in free list (0 min, 100 max allowed)
     O hits, O misses, O trims, O created
     0 failures (0 no memory)
Large buffers, 5024 bytes (total 0, permanent 0):
     0 in free list (0 min, 10 max allowed)
     0 hits, 0 misses, 0 trims, 0 created
     O failures (O no memory)
Huge buffers, 18024 bytes (total 0, permanent 0):
     0 in free list (0 min, 4 max allowed)
     O hits, O misses, O trims, O created
     O failures (O no memory)
Interface buffer pools:
CD2430 I/O buffers, 1524 bytes (total 0, permanent 0):
     0 in free list (0 min, 0 max allowed)
     0 hits, 0 fallbacks
Header pools:
```

```
Header buffers, 0 bytes (total 265, permanent 256):
     9 in free list (10 min, 512 max allowed)
     253 hits, 3 misses, 0 trims, 9 created
     0 failures (0 no memory)
     256 max cache size, 256 in cache
Particle Clones:
     1024 clones, 0 hits, 0 misses
Public particle pools:
F/S buffers, 256 bytes (total 384, permanent 384):
     128 in free list (128 min, 1024 max allowed)
     256 hits, 0 misses, 0 trims, 0 created
     O failures (O no memory)
     256 max cache size, 256 in cache
Normal buffers, 1548 bytes (total 512, permanent 512):
     356 in free list (128 min, 1024 max allowed)
     188 hits, 0 misses, 0 trims, 0 created
     0 failures (0 no memory)
     128 max cache size, 128 in cache
Private particle pools:
Ethernet0/0 buffers, 1536 bytes (total 96, permanent 96):
     0 in free list (0 min, 96 max allowed)
     96 hits, 0 fallbacks
     96 max cache size, 64 in cache
SerialO/O buffers, 1548 bytes (total 14, permanent 14):
     0 in free list (0 min, 14 max allowed)
     14 hits, 0 fallbacks
     14 max cache size, 14 in cache
BRIO/O buffers, 1548 bytes (total 14, permanent 14):
     0 in free list (0 min, 14 max allowed)
     14 hits, 0 fallbacks
     14 max cache size, 14 in cache
BRIO/0:1 buffers, 1548 bytes (total 14, permanent 14):
     0 in free list (0 min, 14 max allowed)
     14 hits, 0 fallbacks
     14 max cache size, 14 in cache
BRIO/0:2 buffers, 1548 bytes (total 14, permanent 14):
     0 in free list (0 min, 14 max allowed)
     14 hits, 0 fallbacks
     14 max cache size, 14 in cache
TokenRingO/O buffers, 1548 bytes (total 64, permanent 64):
     0 in free list (0 min, 64 max allowed)
     64 hits, 0 fallbacks
     64 max cache size, 64 in cache
     4 buffer threshold, 0 threshold transitions
```

Private particle pools are used by the interfaces, and should not be tuned. When no buffer is available in the free list, the router falls back to the public particle pools.

Header buffers are used to record a list of all the particles belonging to a packet.

Note: System buffers are used for process switching. On the Cisco 3600, all these buffers are in the I/O memory which is located in DRAM. You can specify the amount of I/O memory using the **memory-size iomem** command. On the Cisco 7200, the interface particle buffer pools for the high bandwidth port adapters (PAs) are located in SRAM.

Buffer Tuning

Below is an example from the **show buffers** command:

```
Middle buffers, 600 bytes (total 150, permanent 25):
147 in free list (10 min, 150 max allowed)
61351931 hits, 137912 misses, 51605 trims, 51730 created
91652 failures (0 no memory)
```

In this example, the middle buffers have a lot of failures. This is not a serious problem, as it represents only 0.1 % of the hits. These numbers should easily be improved with some buffer tuning.

Buffer tuning is done only when packets cannot be CEF switched.

Depending on the architecture of the router, the buffers you tune generally belong either to I/O memory (low-end), or main memory (high-end). Before tuning the buffers, first check whether you have enough free I/O or main memory using the first lines of the **show memory** command.

Here are some general values that you can use:

- **permanent:** take the number of total buffers in a pool and add about 20%.
- min-free: set min-free to about 20-30% of the permanent number of allocated buffers in the pool.
- max-free: set max-free to something greater than the sum of permanents and minimums.

In the buffer tuning example above, we could enter these commands in the global configuration mode:

```
buffers middle permanent 180
buffers middle min-free 50
buffers middle max-free 230
```

Normally, Cisco IOS® Software creates buffers dynamically so that these settings are fine. However, in case of a traffic burst, the router may not have enough time to create the new buffers, and the number of failures may continue to increase. Use the **buffers** command to change the default buffer pool settings. Ensure that changes in the buffer values are made with caution since improper buffer settings can affect the system performance. If you would like to clear the buffer counters, the router will have to be reloaded.

There are two kinds of traffic bursts:

- **Slow burst**: In this case, the router has sufficient time to create new buffers. Increase the number of min-free buffers. By using free buffers, you can reach the min-free value, and then create new buffers.
- **Fast burst**: With fast traffic bursts, the router does not have enough time to create new buffers, so you should use the free buffers. To do this, modify the number of permanent buffers.

Conclusion: If the create counter increases after the initial tuning, increase the min-free (slow burst). If the failure counter increases, but not the create counter (fast burst), increase the permanent value.

Reserving Fabric Buffer

You can enter the **fabric buffer-reserve** command to improve the system throughput and reserve the ASIC buffers.

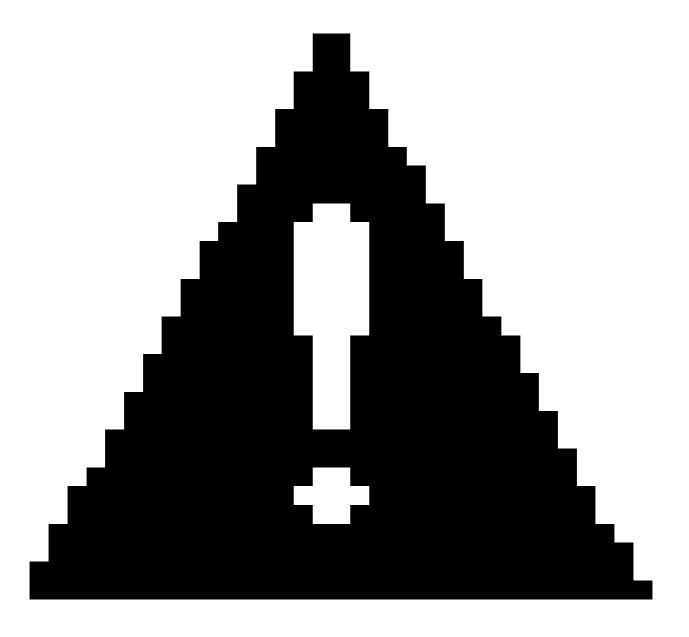
This command is supported on these modules:

- WS-X6704-10GE
- WS-X6748-SFP
- WS-X6748-GE-TX
- WS-X6724-SFP

This command is not supported on the Cisco 7600 Series routers that are configured with a Supervisor Engine 32.

<#root>

fabric buffer-reserve [high | low | medium | value]



Caution: Use this command only under the direction of Cisco TAC.

These are common circumstances where this command is useful:

- Line protocol goes down for multiple interfaces
- Overruns are seen on multiple interfaces
- Ports frequently leave and join EtherChannel
- TestMacNotification test repeatedly fails for line cards with DFC

Buffer Leaks

Below is an example of the output of the **show buffers** command:

43670437 hits, 5134 misses, 0 trims, 1506 created 756 failures (0 no memory)

This output indicates a buffer leak in the big buffers pool. There is a total of 1556 big buffers in the router and only 52 are in the free list. Something is using all the buffers, and not freeing them. For more information on buffer leaks, see <u>Troubleshooting Buffer Leaks</u>.

Related Information

- Troubleshooting Memory Problems
- Troubleshooting Buffer Leaks
- Basic System Management Commands
- Modifying the System Buffer Size
- Technical Support Cisco Systems