

Troubleshooting High CPU Utilization Caused by the HyBridge Input Process on Routers With ATM Interfaces

Document ID: 10448

Contents

Introduction

Prerequisites

- Requirements

- Conventions

Standard Bridging Architecture

Typical Symptoms

Troubleshooting

Workarounds

Related Information

Introduction

This document explains how to troubleshoot high CPU utilization in a router due to the HyBridge Input process. ATM interfaces can support a large number of permanent virtual circuits (PVCs) configured to use request for comments (RFC) 1483 bridged-format protocol data units (PDUs) with standard Cisco IOS[®] bridging and integrated routing and bridging (IRB). This approach relies heavily on broadcasts for connectivity to remote users. As the number of remote users and PVCs increases, the number of broadcasts among these users also increases. Under certain circumstances, these broadcasts produce high CPU utilization on the router.

Prerequisites

Requirements

There are no specific requirements for this document.

Conventions

Refer to Cisco Technical Tips Conventions for more information on document conventions.

Standard Bridging Architecture

The TRFC 1483 specifies that a transparent bridge (which includes a Cisco router configured for bridging) must be able to flood, forward, and filter bridged frames. Flooding is the process by which a frame is copied to all possible appropriate destinations. An ATM bridge floods a frame when it explicitly copies the frame to each virtual circuit (VC), or when it uses a point-to-multipoint VC.

With standard Cisco IOS bridging, frames such as Address Resolution Protocols (ARPs), broadcasts, multicasts, and spanning-tree packets must go through this flooding process. Cisco IOS bridging logic handles every such packet:

1. Runs through the list of interfaces and subinterfaces configured in the bridge group.

2. Runs through the list of VCs configured on the member interfaces in the bridge group.
3. Replicates the frame to each VC.

The Cisco IOS software routines that handle replication need to run in a loop to duplicate the packet on each PVC. If the router supports a large number of bridged-format PVCs, the replication routines run for an extended period, which drive up the CPU. A capture of the **show process cpu** command displays a large "5sec" value for HyBridge input, which is responsible for forwarding packets that use the process switching method of packet forwarding. Cisco IOS needs to process-switch such packets as spanning tree bridge protocol data units (BPDUs), broadcasts, and multicasts that cannot be multicast fast-switched. Process switching can consume large amounts of CPU time since only a limited number of packets are processed per invocation.

When a single interface supports many VCs, traversal of the VC list can overwhelm the CPU. Cisco Bug ID CSCdr11146 resolves this problem. When the bridging logic runs in a loop to replicate the broadcasts, it relinquishes the CPU intermittently. Relinquishment of the CPU is also called suspension of the CPU.

Note: Configuration of many subinterfaces in the same bridge group can also overwhelm the CPU.

Typical Symptoms

If your bridged PVCs result in high CPU utilization on the router, the first thing to look for is a high number of broadcasts on your interface:

```
ATM_Router# show interface atm1/0
ATM1/0 is up, line protocol is up
  Hardware is ENHANCED ATM PA
  MTU 4470 bytes, sub MTU 4470, BW 44209 Kbit,    DLY 190 usec,
    reliability 0/255, txload    1/255, rxload 1/255
  Encapsulation ATM, loopback not set
  Keepalive not supported
  Encapsulation(s): AAL5
  4096 maximum active VCs, 0 current VCCs
  VC idle disconnect time: 300 seconds
  77103 carrier transitions
  Last input 01:06:21, output 01:06:21, output    hang never
  Last clearing of "show interface" counters    never
  Input queue: 0/75/0/702097 (size/max/drops/flushes);    Total output drops: 12201965
  Queueing strategy: Per VC Queueing
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
    59193134 packets input,    3597838975 bytes, 1427069 no buffer
    Received 463236 broadcasts,    0 runts, 0 giants, 0 throttles
    46047 input errors, 46047    CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
    91435145 packets output,    2693542747 bytes, 0 underruns
    0 output errors, 0 collisions,    4 interface resets
    0 output buffer failures,    0 output buffers swapped out
```

As a side effect, you can see a high number of drops on the interface. Under this situation, the problem can be anywhere from slow response on the router to the complete inaccessibility of the router. If you bring the interface down or disconnect the cable from the ATM interface, it should bring the router back.

If the broadcast traffic is bursty, which results only in CPU spikes for short periods of time, the problem can be alleviated if you change the input hold queue on the interface to accommodate the bursts. The default hold queue size is 75 packets and can be changed with the **hold-queue <queue length> in/out** command. Typically, the size of the hold queue must not be increased above 150 because this causes more process-level load on the CPU.

Troubleshooting

If you encounter problems with high CPU utilization caused by HyBridge input, capture this output when you contact the Cisco Technical Assistance Center (TAC). To capture this output, use these commands:

- **show process cpu** – If you notice high CPU utilization, use the **show process CPU** command to isolate which process is at fault. See Troubleshooting High CPU Utilization on Cisco Routers.
- **show stacks {process ID}** – You can also use this command to see what processes are operative and look for potential problems. Paste the output of this command in the Output Interpreter Tool (registered customers only) . Once the processes have been decoded, you can search for possible bugs with the Software Bug Toolkit.

Note: You need to register for a CCO account and be logged on to use both of these tools.

- **show bridge verbose** – Use this show command to determine how many subinterfaces are put in the same bridge group, as well as to see if the interface is overwhelmed.

```
router#show process cpu
```

```
CPU utilization for five seconds: 100%/26%; one minute: 94%; five minutes: 56%
```

PID	Runtime(ms)	Invoked	uSecs	5Sec	1Min	5Min	TTY	Process
1	44	38169	1	0.00%	0.00%	0.00%	0	Load Meter
2	288	733	392	0.00%	0.00%	0.00%	0	PPP auth
3	44948	19510	2303	0.00%	0.05%	0.03%	0	Check heaps
4	4	1	4000	0.00%	0.00%	0.00%	0	Chunk Manager
5	2500	6229	401	0.00%	0.00%	0.00%	0	Pool Manager
[output omitted]								
86	4	1	4000	0.00%	0.00%	0.00%	0	CCSWOFR
87	3390588	1347552	2516	72.72%	69.79%	41.31%	0	HyBridge Input
88	172	210559	0	0.00%	0.00%	0.00%	0	Tbridge Monitor
89	1139592	189881	6001	0.39%	0.42%	0.43%	0	SpanningTree

```
router#show stacks 87
```

```
Process 87: HyBridge Input Process
Stack segment 0x61D15C5C - 0x61D18B3C
FP: 0x61D18A18, RA: 0x60332608
FP: 0x61D18A58, RA: 0x608C5400
FP: 0x61D18B00, RA: 0x6031A6D4
FP: 0x61D18B18, RA: 0x6031A6C0
```

```
router#show bridge verbose
```

```
Total of 300 station blocks, 299 free
Codes: P - permanent, S - self
```

BG	Hash	Address	Action	Interface	VC Age	RX count	TX count
1	8C/0	0000.0cd5.f07c	forward	ATM4/0/0.1	9 0	1857	0

Flood ports (BG 1)	RX count	TX count
ATM4/0/0.1	0	0

In addition, shut down the Bridge Group Virtual Interface (BVI) and monitor CPU utilization with several captures of output from the **show process cpu** command.

Workarounds

Cisco recommends that you implement these workarounds as a solution to high CPU utilization caused by standard bridging:

- Implement the Cisco IOS x Digital Subscriber Line Bridge Support feature, which configures the router for intelligent bridge flooding through subscriber policies. Selectively block ARPs, broadcasts,

multicasts and spanning-tree BPDUs.

- Break up the VCs on a few multipoint interfaces, each with a different IP network.
- Configure the aging timer of IP ARP and bridging table entries to the same value. Otherwise, you can see unnecessary flooding of traffic in your links. The default ARP timeout is four hours. The default bridge aging-time is 10 minutes. For a remote user that has been idle for 10 minutes, the router purges the user's bridge table entry only and retains the ARP table entry. When the router needs to send traffic downstream to the remote user, it checks the ARP table and finds a valid entry to point to the MAC address. When the router checks the bridge table for this MAC address and fails to find it, the router floods the traffic out every VC in the bridge group. Use these commands to set the ARP and bridge table aging times.

```
router(config)#bridge 1 aging-time ?  
<10-1000000> Seconds
```

```
router(config)#interface bvi1
```

```
router(config-if)#arp timeout ?  
<0-2147483> Seconds
```

- Replace standard bridging and IRB with routed bridge encapsulation (RBE) or bridged-style PVCs at the head-end ATM interface. RBE increases forwarding performance as it supports Cisco Express Forwarding (CEF) and runs IP packets only through a routing decision and not through a bridging decision. On the 12.1(1)T train, the packets can be software switched. If so, you can see this error message:

```
%FIB-4-PUNTINTF: CEF punting packets switched to ATM1/0.100 to next slower pa  
%FIB-4-PUNTINTF: CEF punting packets switched to ATM1/0.101 to next slower pa
```

The problem is documented in CSCdr37618, and the fix is to upgrade to 12.2 mainline. Refer to Routed Bridged Encapsulation Baseline Architecture and Configuring Bridged-Style PVCs on ATM Interfaces in the GSR and 7500 Series for more information.

Related Information

- [Troubleshooting High CPU Utilization on Cisco Routers](#)
- [Technical Support & Documentation – Cisco Systems](#)
- [Tools and Utilities – Cisco Systems](#)

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

© 2014 – 2015 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

Updated: Jun 05, 2005

Document ID: 10448
