

Catalyst 3750 Series Switches High CPU Utilization Troubleshooting

Document ID: 68461

Introduction

Prerequisites

- Requirements
- Components Used
- Conventions

Background Information

Troubleshoot Common High CPU Utilization Problems

- High CPU Due to a Storm of IGMP Leave Messages
- High CPU Due to a GRE Tunnel
- High CPU During a Configuration Change
- High CPU Due to Excessive ARP Requests
- High CPU Due to IP SNMP Process
- High CPU Due to SDM Template
- High CPU Due to Policy Based Routing
- High CPU Due to Excessive ICMP Redirects

Related Information

Introduction

This document describes causes of high CPU utilization on the Cisco Catalyst 3750 Series Switches. Similar to Cisco routers, switches use the **show processes cpu** command in order to show CPU utilization to identify the causes of high CPU utilization. However, due to the differences in architecture and forwarding mechanisms between Cisco routers and switches, the typical output of the **show processes cpu** command differs significantly. This document also lists some common symptoms that cause high CPU utilization on the Catalyst 3750 Series Switch.

Prerequisites

Requirements

There are no specific requirements for this document.

Components Used

The information in this document is based on Catalyst 3750 Switches.

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, make sure that you understand the potential impact of any command.

Conventions

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

Background Information

Before you look at the CPU packet-handling architecture and troubleshoot high CPU utilization, you must understand the different ways in which hardware-based forwarding switches and Cisco IOS® Software-based routers use the CPU. The common misconception is that high CPU utilization indicates the depletion of resources on a device and the threat of a crash. A capacity issue is one of the symptoms of high CPU utilization on Cisco IOS routers. However, a capacity issue is almost never a symptom of high CPU utilization with hardware-based forwarding switches.

The first step to troubleshoot the high CPU utilization is to check the Cisco IOS version release notes of your Catalyst 3750 Switch for the possible known IOS bug. This way you can eliminate the IOS bug from your troubleshooting steps. Refer to Cisco Catalyst 3750 Series Switches Release Notes for the list of release notes for Catalyst 3750 Switches.

Troubleshoot Common High CPU Utilization Problems

This section covers some of the common high CPU utilization problems on the Catalyst 3750 Switch.

High CPU Due to a Storm of IGMP Leave Messages

One of the common reasons for high CPU utilization is that the Catalyst 3750 CPU is busy with the processing storm of Internet Group Management Protocol (IGMP) leave messages. If a stack of Catalyst 3750 Switches that run Cisco IOS Software Release 12.1(14)EA1a are connected to another switch, such as a Cat6500 that runs CatOS, which generates MAC-based IGMP queries with IP options, the 3750 experiences high CPU utilization in the IGMP SN (snopping) process. This is a result of the MAC-based query packets looping within the stack. You can also see a high CPU with the **HRPC hl2mm request** process. If you have EtherChannel configured on the Catalyst 3750 stack with Cisco IOS Software Release 12.1(14)EA1a, a storm of IGMP leave messages might be created.

The Catalyst 3750 receives many IGMP queries. This makes the IGMP query counter start to increment by hundreds per second. This leads to high CPU in the Catalyst 3750 Switch. Refer to Cisco bug ID CSCeg55298 (registered customers only) . The bug was identified in Cisco IOS Software Release 12.1(14)EA1a and is fixed on Cisco IOS Software Releases 12.2(25)SEA and later. The permanent solution is to upgrade to the latest Cisco IOS version. The temporary workaround is to disable IGMP snooping on the Catalyst 3750 stack, or disable MAC-based query on the switch connected to the 3750 stack.

This is a sample output from the **show ip traffic** command which shows IP packets with bad options and alerts that rapidly increment:

```
Switch#show ip traffic
Rcvd: 48195018 total, 25628739 local destination
 0 format errors, 0 checksum errors, 10231692 bad hop count
 0 unknown protocol, 9310320 not a gateway
 0 security failures, 10231 bad options, 2640539 with options
Opts: 2640493 end, 206 nop, 0 basic security, 2640523 loose source route
 0 timestamp, 0 extended security, 16 record route
 0 stream ID, 0 strict source route, 10231 alert, 0 cipso, 0 ump
 0 other
Frgs: 16 reassembled, 0 timeouts, 0 couldn't reassemble
32 fragmented, 0 couldn't fragment
Bcast: 308 received, 0 sent
Mcast: 4221007 received, 4048770 sent
Sent: 25342014 generated, 20710669 forwarded
Drop: 617267 encapsulation failed, 0 unresolved, 0 no adjacency
 0 no route, 0 unicast RPF, 0 forced drop
 0 options denied, 0 source IP address zero
```

!--- Output suppressed.

The **show processes cpu** command displays information about the active processes in the switch and their corresponding CPU utilization statistics. This is a sample output of the **show processes cpu** command when the CPU utilization is normal:

```
switch#show processes cpu
```

```
CPU utilization for five seconds: 8%/4%; one minute: 6%; five minutes: 5%
```

PID	Runtime(ms)	Invoked	uSecs	5Sec	1Min	5Min	TTY	Process
1	384	32789	11	0.00%	0.00%	0.00%	0	Load Meter
2	2752	1179	2334	0.73%	1.06%	0.29%	0	Exec
3	318592	5273	60419	0.00%	0.15%	0.17%	0	Check heaps
4	4	1	4000	0.00%	0.00%	0.00%	0	Pool Manager
5	6472	6568	985	0.00%	0.00%	0.00%	0	ARP Input
6	10892	9461	1151	0.00%	0.00%	0.00%	0	IGMPSN

!--- CPU utilization at normal condition.

7	67388	53244	1265	0.16%	0.04%	0.02%	0	CDP Protocol
8	145520	166455	874	0.40%	0.29%	0.29%	0	IP Background
9	3356	1568	2140	0.08%	0.00%	0.00%	0	BOOTP Server
10	32	5469	5	0.00%	0.00%	0.00%	0	Net Background
11	42256	163623	258	0.16%	0.02%	0.00%	0	Per-Second Job
12	189936	163623	1160	0.00%	0.04%	0.05%	0	Net Periodic
13	3248	6351	511	0.00%	0.00%	0.00%	0	Net Input
14	168	32790	5	0.00%	0.00%	0.00%	0	Compute load a
15	152408	2731	55806	0.98%	0.12%	0.07%	0	Per-minute Job
16	0	1	0	0.00%	0.00%	0.00%	0	HRPC hI2mm req

!--- Output suppressed.

This is a sample output of the **show processes cpu** command when the CPU utilization is high due to the IGMP snooping process:

```
switch#show processes cpu
```

```
CPU utilization for five seconds: 8%/4%; one minute: 6%; five minutes: 5%
```

PID	Runtime(ms)	Invoked	uSecs	5Sec	1Min	5Min	TTY	Process
1	384	32789	11	0.00%	0.00%	0.00%	0	Load Meter
2	2752	1179	2334	0.73%	1.06%	0.29%	0	Exec
3	318592	5273	60419	0.00%	0.15%	0.17%	0	Check heaps
4	4	1	4000	0.00%	0.00%	0.00%	0	Pool Manager
5	6472	6568	985	0.00%	0.00%	0.00%	0	ARP Input
6	10892	9461	1151	100	100	100	0	IGMPSN

!--- Due to high CPU utilization.

7	67388	53244	1265	0.16%	0.04%	0.02%	0	CDP Protocol
8	145520	166455	874	0.40%	0.29%	0.29%	0	IP Background
9	3356	1568	2140	0.08%	0.00%	0.00%	0	BOOTP Server
10	32	5469	5	0.00%	0.00%	0.00%	0	Net Background
11	42256	163623	258	0.16%	0.02%	0.00%	0	Per-Second Job
12	189936	163623	1160	0.00%	0.04%	0.05%	0	Net Periodic
13	3248	6351	511	0.00%	0.00%	0.00%	0	Net Input
14	168	32790	5	0.00%	0.00%	0.00%	0	Compute load a
15	152408	2731	55806	0.98%	0.12%	0.07%	0	Per-minute Job

!--- Output suppressed.

High CPU Due to a GRE Tunnel

The General Routing Encapsulation (GRE) tunnel is not supported by the Cisco Catalyst 3750 Series Switches. Even though this feature can be configured with CLI, the packets can be neither switched by hardware, nor by software, which increases the CPU utilization.

Note: Only Distance Vector Multicast Routing Protocol (DVMRP) tunnel interfaces are supported for multicast routing in the Catalyst 3750. Even for this, packets cannot be switched with hardware. The packets routed through this tunnel must be switched through software. The larger number of packets forwarded through this tunnel increases CPU utilization.

There is no workaround for this problem. This is a hardware limitation in the Catalyst 3750 Series Switches.

High CPU During a Configuration Change

If Catalyst 3750 Switches are connected in a stack, and if there are any configuration changes made to a switch, the **hulc running config** process wakes up and generates a new copy of the running configuration. Then, it sends to all the switches in the stack. The new running configuration is CPU-intensive. Therefore, the CPU usage is high when building a new running configuration process and when forwarding the new configurations to other switches. However, this high CPU usage should exist only for the same amount of time it takes to perform the building configuration step of the **show running-configuration** command.

There is no need of a workaround for this problem. The CPU usage is normally high in these situations.

This is a sample output of the **show processes cpu** command when the CPU utilization is high due to the **hulc running** process:

```
switch#show processes cpu

CPU utilization for five seconds: 63%/0%; one minute: 27%; five minutes: 23%

PID Runtime(ms)   Invoked  uSecs      5Sec   1Min   5Min   TTY Process
  1      384         32789    11         0.00%  0.00%  0.00%  0  Load Meter
  2      2752         1179    2334         0.73%  1.06%  0.29%  0  Exec
  3  318592         5273   60419         0.00%  0.15%  0.17%  0  Check heaps
  4         4           1     4000         0.00%  0.00%  0.00%  0  Pool Manager
  5     6472         6568    985         0.00%  0.00%  0.00%  0  ARP Input
  6    10892         9461   1151         0.00%  0.00%  0.00%  0  IGMPSN
  7    67388        53244   1265         0.16%  0.04%  0.02%  0  CDP Protocol
  8   145520       166455    874         0.40%  0.29%  0.29%  0  IP Background
  9     3356         1568   2140         0.08%  0.00%  0.00%  0  BOOTP Server
 10         32         5469     5         0.00%  0.00%  0.00%  0  Net Background
 11    42256       163623    258         0.16%  0.02%  0.00%  0  Per-Second Job
 12   189936       163623   1160         0.00%  0.04%  0.05%  0  Net Periodic
 13     3248         6351    511         0.00%  0.00%  0.00%  0  Net Input
 14      168        32790     5         0.00%  0.00%  0.00%  0  Compute load
 15   152408        2731   55806         0.98%  0.12%  0.07%  0  Per-minute Job
 16         0           1         0         0.00%  0.00%  0.00%  0  HRPC hI2mm req
 17   85964         426   201793        55.72% 12.05%  5.36%  0  hulc running
```

!--- Output suppressed.

High CPU Due to Excessive ARP Requests

High CPU utilization of the Address Resolution Protocol (ARP) input process occurs if the router has to originate an excessive number of ARP requests. ARP requests for the same IP address are rate-limited to one request every two seconds. Therefore, an excessive number of ARP requests has to originate for different IP addresses. This can occur if an IP route has been configured and points to a broadcast interface. An obvious example is a default route, such as:

```
ip route 0.0.0.0 0.0.0.0 FastEthernet0/0
```

In this case, the router generates an ARP request for each IP address that is not reachable through more specific routes, which means that the router generates an ARP request for almost every address on the Internet. Refer to [Specifying a Next Hop IP Address for Static Routes](#) for more information on how to configure the next hop IP address for static routing.

Alternatively, an excessive amount of ARP requests can be caused by a malicious traffic stream which scans through locally attached subnets. An indication of such a stream is the presence of a very high number of incomplete ARP entries in the ARP table. Because incoming IP packets that trigger ARP requests have to be processed, troubleshooting this problem is essentially the same as troubleshooting high CPU utilization in the IP Input process.

High CPU Due to IP SNMP Process

In the latest Cisco IOS versions for Catalyst 3750, Simple Network Management Protocol (SNMP) requests are handled by the SNMP ENGINE. It is normal for the CPU to go high due to this SNMP ENGINE process. The SNMP process has a low priority and should not affect any functionality on the switch.

Refer to [IP Simple Network Management Protocol \(SNMP\) Causes High CPU Utilization](#) for more information on high CPU utilization caused by the SNMP ENGINE process.

High CPU Due to SDM Template

The Switch Database Management (SDM) on the Catalyst 3750 Series Switches manages the Layer 2 and Layer 3 switching information that is maintained in the Ternary Content Addressable Memory (TCAM). The SDM templates are used to configure system resources in the switch in order to optimize support for specific features, which depends on how the switch is used in the network. The SDM templates can be selected in order to provide maximum system usage for some functions, or to use the default template in order to balance resources. The templates prioritize system resources in order to optimize support for these types of features:

- **Routing** The routing template maximizes system resources for unicast routing, typically required for a router or aggregator in the center of a network.
- **VLANs** The VLAN template disables routing and supports the maximum number of unicast MAC addresses. This is typically selected for a Layer 2 switch.
- **Access** The access template maximizes system resources for access control lists (ACLs) to accommodate a large number of ACLs.
- **Default** The default template gives balance to all functions.

There are two versions of each template: a desktop template and an aggregator template.

Note: The default template for desktop switches is the default desktop template. The default template for the Catalyst 3750-12S is the default aggregator template.

Select an appropriate SDM template that provides maximum system usage for the feature used. An

inappropriate SDM template can overload the CPU and severely degrade the switch performance.

Issue the **show platform tcam utilization** command to see how much TCAM has now been utilized and how much is available.

```
Switch#show platform tcam utilization
```

CAM Utilization for ASIC# 0	Max Masks/Values	Used Masks/values
Unicast mac addresses:	784/6272	12/26
IPv4 IGMP groups + multicast routes:	144/1152	6/26
IPv4 unicast directly-connected routes:	784/6272	12/26
IPv4 unicast indirectly-connected routes:	272/2176	8/44
IPv4 policy based routing aces:	0/0	0/0
IPv4 qos aces:	528/528	18/18
IPv4 security aces:	1024/1024	27/27

Note: Allocation of TCAM entries per feature uses a complex algorithm. The above information is meant to provide an abstract view of the current TCAM utilization

If the TCAM utilization is close to maximum for any of the parameters, check if any of the other template features can optimize for that parameter.

```
show sdm prefer access | default | dual-ipv4-and-ipv6 | routing | vlan
```

```
Switch# show sdm prefer routing
```

"desktop routing" template:

The selected template optimizes the resources in the switch to support this level of features for 8 routed interfaces and 1024 VLANs.

number of unicast mac addresses:	3K
number of igmp groups + multicast routes:	1K
number of unicast routes:	11K
number of directly connected hosts:	3K
number of indirect routes:	8K
number of policy based routing aces:	512
number of qos aces:	512
number of security aces:	1K

In order to specify the SDM template to use on the switch, issue the **sdm prefer** global configuration command.

Note: The switch reload is required to use the new SDM template.

High CPU Due to Policy Based Routing

Policy Based Routing (PBR) implementation in Cisco Catalyst 3750 switches has some limitations. If these restrictions are not followed, it can cause high CPU utilization.

- You can enable PBR on a routed port or an SVI.
- The switch does not support route-map deny statements for PBR.
- Multicast traffic is not policy-routed. PBR applies only to unicast traffic.
- Do not match ACLs that permit packets destined for a local address. PBR forwards these packets, which can cause ping or Telnet failure or route protocol flapping.

- Do not match ACLs with deny ACEs. Packets that match a deny ACE are sent to the CPU, which can cause high CPU utilization.
- In order to use PBR, you must first enable the routing template with the **sdm prefer routing** global configuration command. PBR is not supported with the VLAN or default template.

For a complete list, refer to the PBR Configuration Guidelines.

High CPU Due to Excessive ICMP Redirects

You can get ICMP dropped redirects when one VLAN (or any Layer 3 port) receives a packet where the source IP is on one subnet, the destination IP is on another subnet, and the next hop is on the same VLAN or layer 3 segment.

Here is an example:

You can see this message in **show log**:

```
51w2d: ICMP-Q:Dropped redirect disabled on L3 IF: Local Port Fwding L3If:Vlan7 L2If:Gigabi
TPFFD:EDC10041_02C602C6_00B0056A-000000B4_EBF6001B_0D8A3746
```

This occurs where the packet is received on VLAN 7 with source IP 64.253.128.3 and tries to reach 208.118.132.9, the destination IP. You can see that the next hop configured in the switch (64.253.128.41, in this case) is also on the same VLAN 7.

Related Information

- [Understanding EtherChannel Inconsistency Detection](#)
- [Multicast Does Not Work in the Same VLAN in Catalyst Switches](#)
- [CPU Utilization on Catalyst 4500/4000, 2948G, 2980G, and 4912G Switches That Run CatOS Software](#)
- [Catalyst 6500/6000 Switch High CPU Utilization](#)
- [LAN Product Support Pages](#)
- [LAN Switching Support Page](#)
- [Technical Support & Documentation – 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.](#)

Updated: Jun 26, 2009

Document ID: 68461
