Catalyst 3550 Series Switches High CPU Utilization

Document ID: 71990

Introduction

This document describes causes of high CPU utilization on the Cisco Catalyst 3550 Series Switches. This document also lists common network or configuration scenarios that can cause high CPU utilization on the Catalyst 3550 Series Switches.

Cisco Catalyst switches use the **show processes cpu** command in order to identify the causes of high CPU utilization. The **show processes cpu** command shows CPU utilization averaged over the past five seconds, one minute, and five minutes. CPU utilization numbers do not provide a true linear indication of the utilization with respect to the offered load. These are some of the major reasons:

- In a real world network, the CPU has to handle various system maintenance functions, such as network management.
- The CPU has to process periodic and event−triggered routing updates.
- There are other internal system overhead operations, such as polling for resource availability, that are not proportional to traffic load.

Prerequisites

Requirements

There are no specific requirements for this document.

Components Used

The information in this document is based on Catalyst 3550 Series 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 CPU utilization of 20% to 50% is normal on a Catalyst 3550 Switch, even under minimal load. CPU utilization does not reflect the total number of packets being switched or the total load on the switch. The CPU is responsible for the process of IP traffic (broadcast, telnet, SNMP) on the management VLAN, for the process of control packets Spanning−Tree Protocol (STP), Cisco Discovery Protocol (CDP), DDSN Transfer Protocol (DTP), Port Aggregation Protocol (PAgP), Link Aggregation Control Protocol (LACP), Unidirectional Link Detection (UDLD), address learning, routing protocols, port status and LED operations. If the CPU utilization is extremely high (around 90% − 99%), this does not directly affect the switching of data. However, high CPU utilization might start to affect protocols such as STPs. The CPU on the Catalyst 3550 is used for management purposes only. The CPU is not used to forward packets. This is handled by ASICs. The increase in the CPU utilization does not affect traffic forwarding.

The first step to troubleshoot the high CPU utilization is to check the Cisco IOS version release notes of your Catalyst 3550 Switch for the possible known IOS bug. This way you can eliminate the IOS bug from your troubleshooting steps. Refer to Cisco Catalyst 3550 Series Switches Release Notes for a description of new features, system requirements, limitations, restrictions, caveats, and troubleshooting information for a particular software release for Catalyst 3550 Switches.

Troubleshoot Common High CPU Utilization Problems

This section provides information about some of the common high CPU utilization problems on the Catalyst 3550 Switch.

High CPU Utilization After Enabling GRE Tunnels

Generic Routing Encapsulation (GRE) tunnels are not supported on the Cisco Catalyst 3550 Switch. Even though the CLI commands are there to configure the GRE, it is not officially supported. Refer to the Unsupported VPN Configuration Commands section of Unsupported CLI Commands for Catalyst 3550 for this information. The reason for this is that the Cisco Catalyst 3550 Switch uses hardware−based Cisco Express Forwarding (CEF) switching. There is no method to CEF−switch GRE packets. GRE packets must be encapsulated by the software. The hardware does not have the capability to encapsulate the packets. Consequently, this traffic is processed or software switched. The process or software switched traffic can quickly cause the CPU to spike.

High CPU Utilization When Pinging Own Interface

An extended ping from one interface to another interface on the same switch can cause high CPU utilization. This can occur when a large number of ping packets are sent and received. This is an expected behavior. The workaround is to not perform a ping from one interface to another on the same switch. Refer to Cisco bug ID CSCea19301 (registered customers only) for more information.

High CPU Utilization Due to VUR_MGR bg Process

The VUR_MGR bg process is the Vegas Unicast Routing Manager process which is a platform specific module that interfaces with IOS. This process implements the hardware independent fuctionality required in the platform for unicast routing. Each time an Address Resolution Protocol (ARP) is resolved for a destination, the corresponding entry needs to be programmed in hardware.

The VUR_MGR bg process is responsible for unicast routing and is high if the switch is learning routing information. It is also high if you see frequent routing changes. Issue the **clear ip route** command in exec mode to clear the condition. However, this does not prevent the condition to recur.

```
Cat-3550#show processes cpu
CPU utilization for five seconds: 99%/0%; one minute: 99%; five minutes: 99%
 PID  Runtime(ms)  Invoked  uSecs    5Sec   1Min   5Min TTY Process
   1           4   1103385      0   0.00%  0.00%  0.00%   0 Load Meter
   2       52592   5709333      9   0.00%  0.00%  0.00%   0 Spanning Tree
   3     1897604    550508   3447   0.00%  0.01%  0.00%   0 Check heaps
   4           0         1      0   0.00%  0.00%  0.00%   0 Chunk Manager
   5       12268     59126    207   0.00%  0.00%  0.00%   0 Pool Manager
   6           0         2      0   0.00%  0.00%  0.00%   0 Timers
   7          12         2   6000   0.00%  0.00%  0.00%   0 Entity MIB API
   8        1244   1101072      1   0.00%  0.00%  0.00%   0 HC Counter Timer
   9        3340     99041     33   0.00%  0.00%  0.00%   0 ARP Input
  10           0         2      0   0.00%  0.00%  0.00%   0 Net Input
  11           0         1      0   0.00%  0.00%  0.00%   0 Critical Bkgnd
  12       13172    591932     22   0.00%  0.00%  0.00%   0 Net Background
  13           0        61      0   0.00%  0.00%  0.00%   0 Logger
  14        7644   5465294      1   0.00%  0.00%  0.00%   0 TTY Background
  15        9608   5465305      1   0.00%  0.00%  0.00%   0 Per-Second Jobs
  16           0         2      0   0.00%  0.00%  0.00%   0 Vegas Storm Cont
  17          88 126356607      0   0.00%  0.00%  0.00%   0 Vegas LED Proces
  18           0         1      0   0.00%  0.00%  0.00%   0 SCQ_PROCESS
  19          20       238     84   0.00%  0.00%  0.00%   0 RAM Access
  20       84020  71566887      1   0.00%  0.00%  0.00%   0 SW Frame Ager
  21           4   1103386      0   0.00%  0.00%  0.00%   0 Compute load avg
  22      368684     91998   4007   0.00%  0.00%  0.00%   0 Per-minute Jobs
  23        6288  10814019      0   0.00%  0.00%  0.00%   0 L2TM Process
  24      132772  15688292      8   0.00%  0.00%  0.00%   0 Vegas Statistics
  25       10056  10814019      0   0.00%  0.00%  0.00%   0 L3TM
  26          40  70646868      0   0.00%  0.00%  0.00%   0 HMATM Learn proc
  27        4808   5465299      0   0.00%  0.00%  0.00%   0 HMATM Age proces
  28           8        42    190   0.00%  0.00%  0.00%   0 VL2MM
  29       12896   5465325      2   0.00%  0.00%  0.00%   0 L3MD
  30      306288  13114927     23   0.00%  0.00%  0.00%   0 L3MD_STAT
  31           0         1      0   0.00%  0.00%  0.00%   0 Vegas Bridging
  32         168  68417996      0   0.00%  0.00%  0.00%   0 VegasPM
  33   506543096   3512719 144203  99.67% 99.79% 99.95%   0 VUR_MGR bg proce

!--- Output suppressed.
```

Refer to Cisco bug ID CSCdx31480 (registered customers only) for more information. This bug is fixed in Cisco IOS Software Release 12.1(11)EA1 and later.

High CPU Utilization Due to IP Input Process

The Cisco IOS software process called IP input takes care of process−switching IP packets. If the IP input process uses unusually high CPU resources, the switch is process−switching a lot of IP traffic. Refer to Troubleshooting High CPU Utilization in IP Input Process for information on how to troubleshoot high CPU utilization due to the IP input process.

High CPU Utilization Due to GigaStack GBIC Modules

When you insert a GigaStack GBIC in a GBIC module slot, the CPU utilization increases by six percent. This increase occurs for each GigaStack GBIC added to the switch. The VegasPM process in the **show processes cpu** command output shows this CPU utilization. The VegasPM process manages the Gigastack GBIC operation on the switch.

```
Cat−3550#show processes cpu
CPU utilization for five seconds: 73%/0%; one minute: 67%; five minutes: 66%
 PID Runtime(ms)   Invoked      uSecs   5Sec   1Min   5Min TTY Process
   1           0         1          0  0.00%  0.00%  0.00%   0 Chunk Manager
   2        9008   2476157          3  0.00%  0.00%  0.00%   0 Load Meter
   3         700       722        969  0.00%  0.00%  0.00%   0 SpanTree Helper
   4     3607596   1260465       2862  0.08%  0.02%  0.00%   0 Check heaps
   5    10197340  22011722        463  0.00%  0.01%  0.06%   0 Pool Manager
   6           0         2          0  0.00%  0.00%  0.00%   0 Timers
   7          32        20       1600  0.00%  0.00%  0.00%   0 Entity MIB API
   8       55160   2475564         22  0.00%  0.00%  0.00%   0 HC Counter Timer
   9     8469276  82844936        102  0.40%  0.10%  0.08%   0 ARP Input
  10          24      1592         15  0.00%  0.00%  0.00%   0 Net Input
  11           0         1          0  0.00%  0.00%  0.00%   0 Critical Bkgnd
  12      338884   6258667         54  0.00%  0.00%  0.00%   0 Net Background
  13          64      1035         61  0.00%  0.00%  0.00%   0 Logger
  14      122036  12262545          9  0.00%  0.00%  0.00%   0 TTY Background
  15      623312  12262563         50  0.00%  0.00%  0.00%   0 Per−Second Jobs
  16       20084   2476157          8  0.00%  0.00%  0.00%   0 Compute load avg
  17     1048664    206875       5069  0.00%  0.00%  0.00%   0 Per−minute Jobs
  18           0         2          0  0.00%  0.00%  0.00%   0 Vegas Storm Cont
  19     8455544 279108955         30  0.08%  0.02%  0.00%   0 Vegas LED Proces
  20           0         1          0  0.00%  0.00%  0.00%   0 SCQ_PROCESS
  21        1800      2468        729  0.00%  0.00%  0.00%   0 RAM Access
  22     8175752 229227255         35  0.00%  0.04%  0.04%   0 SW Frame Ager
  23         456      1104        413  0.00%  0.00%  0.00%   0 VLAN Info Update
  24      205420  23909802          8  0.00%  0.00%  0.00%   0 L2TM Process
  25    44726932  24487858       1826  0.24%  0.37%  0.38%   0 Vegas Statistics
  26      617288  23909801         25  0.00%  0.01%  0.00%   0 L3TM
  27     4122052 141972837         29  0.00%  0.00%  0.00%   0 HMATM Learn proc
  28      197324  12262549         16  0.00%  0.00%  0.00%   0 HMATM Age proces
  29      495616   2291212        216  0.00%  0.01%  0.00%   0 VL2MM
  30     2175000  18368663        118  0.00%  0.00%  0.00%   0 L3MD
  31   142702548  78397696       1820  1.30%  1.13%  1.11%   0 L3MD_STAT
  32           0         1          0  0.00%  0.00%  0.00%   0 Vegas Bridging
  33  3607116392 397152327       9082 61.46% 63.81% 63.05%   0 VegasPM
  34      121392   6295553         19  0.00%  0.00%  0.00%   0 VUR_MGR bg proce

!−−− Output suppressed.
```

As a workaround, if the network design permits, use other types of GBICs such as fiber. These do not cause additional CPU utilization. Refer to Cisco bug ID CSCdx90515 (registered customers only) for more information.

High CPU Utilization Due to TTY Background Process

The TTY Background process is a generic process used by all terminal lines (console, aux, async, and so on). Normally there should not be any impact on the performance of the switch, because this process has a lower priority compared to the other processes that need to be scheduled by the Cisco IOS software.

If this process takes high CPU utilization, check whether logging synchronous is configured under line con 0. Refer to Cisco bug ID CSCdy01705 (registered customers only) for more information.

High CPU Utilization Due to SNAP Encapsulation of IPv4 Packets

On the Catalyst 3550 Switch, Layer 3 forwarding of IPv4 in the Subnetwork Access Protocol (SNAP) can only be done in the software. SNAP−encapsulated IPv4 packets that are directed to the router MAC address or the Hot Standby Router Protocol (HSRP) group MAC address (if this is the active router in the HSRP group) are forwarded to the switch CPU. This action can potentially cause high CPU utilization levels.

Packets received from media types that require SNAP encapsulation of IPv4 packets require the switch to forward SNAP−encapsulated packets. In general, Layer 2 forwarding of IPv4 in SNAP encapsulation takes place in the hardware, unless a VLAN map or port Access Control List (ACL) contains an IP ACL. However, this cannot take place on the Cisco Catalyst 3550 Switch.

This is a hardware limitation, and there is no workaround. Refer to Cisco bug ID CSCed59864 (registered customers only) for more information.

High CPU Utilization Due to IP Redirects

ICMP redirect messages are used by routers and switches to notify the hosts on the data link that a better route is available for a particular destination. By default, Cisco routers and switches send ICMP redirects.

You can expect the sourcing device to act on the ICMP redirect that the Catalyst 3550 sends, and to change the next hop for the destination. However, not all devices respond to an ICMP redirect. If the device does not respond, the Catalyst 3550 must send redirects for every packet that the switch receives from the sending device. These redirects can consume a great deal of CPU resources. The high CPU utilization is caused due to the high amount of ICMP redirect traffic that hits the CPU. The workaround is to disable ICMP redirects with the **no ip redirects** interface mode command.

This scenario can also occur when you have configured secondary IP addresses. When you enable the secondary IP addresses, the IP redirect is automatically disabled. Make sure you do not manually enable the IP redirects when you have configured secondary IP addresses.

High CPU Utilization Due to Broadcast Storm

The CPU utilization of a switch can go up if a large number of broadcast packets is received.

```
3550−1>show int gi0/1
GigabitEthernet0/1 is up, line protocol is up (connected)
  Hardware is Gigabit Ethernet, address is 0014.698a.bb31 (bia
0014.698a.bb31)
  MTU 1500 bytes, BW 1000000 Kbit, DLY 10 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)

!−−− Part of the output elided.
```

```
  5 minute input rate 3660000 bits/sec, 382 packets/sec
  5 minute output rate 7413000 bits/sec, 2119 packets/sec
     25774872 packets input, 1350686943 bytes, 0 no buffer
     Received 25774872 broadcasts (0 multicast)
     0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored

!--- Rest of the output elided.
```

If these broadcast storms are frequent, then you might have to look into the design of the network. If the broadcast storms are occasional, you can configure the Storm Control feature in order to equip the device against the storm. Refer to Configuring Port−Based Traffic Control for more information.

NetPro Discussion Forums − Featured Conversations

Networking Professionals Connection is a forum for networking professionals to share questions, suggestions, and information about networking solutions, products, and technologies. The featured links are some of the most recent conversations available in this technology.

NetPro Discussion Forums − Featured Conversations for LAN
Network Infrastructure: LAN Routing and Switching
Network Infrastructure: Getting Started with LANs

Related Information

- **The show processes Command**
- **Troubleshooting High CPU Utilization due to Processes**
- **High CPU Utilization on Catalyst 2900XL/3500XL Switches**
- **Catalyst 3750 Series Switches High CPU Utilization Troubleshooting**
- **Catalyst 6500/6000 Switch High CPU Utilization**
- **Troubleshooting High CPU Utilization on Cisco Routers**
- **LAN Product Support Pages**
- **LAN Switching Support Page**
- **Technical Support & Documentation − Cisco Systems**

Updated: Jul 30, 2007 Document ID: 71990