CPU Utilization on Catalyst 4500/4000, 2948G, 2980G, and 4912G Switches That Run CatOS Software

Document ID: 12032

Contents

Introduction

This document provides information about the output of the **show processes cpu** command when you issue the command on the Cisco Catalyst 4500/4000, 2948G, 2980G, and 4912G switches that run Catalyst OS (CatOS) system software. This document describes how to identify the causes of high CPU utilization on these switches. The document also lists some common network or configuration scenarios that cause high CPU utilization on the Catalyst 4500 series.

Note: If you run Cisco IOS Software–based Catalyst 4500/4000 series switches, refer to High CPU Utilization on Cisco IOS Software–Based Catalyst 4500/4000 Switches.

Note: In this document, the words switch and switches refer to the Catalyst 4500/4000, 2948G, 2980G, and 4912G switches.

Like Cisco routers, switches use the **show processes cpu** command in order to show CPU utilization for the switch Supervisor Engine processor. 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. The meaning of the output differs, too.

This document clarifies these differences. The document describes use of the CPU on the switches and how to interpret the **show processes cpu** command output.

Prerequisites

Requirements

There are no specific requirements for this document.

Components Used

The information in this document is based on the software and hardware versions for:

- Catalyst 4500/4000 switches that run CatOS
- Catalyst 2948G switch
- Catalyst 2980G and 2980G−A switches
- Catalyst 4912G switch

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

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

Understand CPU Utilization on Catalyst 4500/4000, 2948G, 2980G, and 4912G Switches

Cisco software−based routers use software in order to process and route packets. CPU utilization on a Cisco router tends to increase as the router performs more packet processing and routing. Therefore, the **show processes cpu** command can provide a fairly accurate indication of the traffic processing load on the router.

Catalyst 4500/4000 that run CatOS, 2948G, 2980G, and 4912G switches do not use the CPU in the same way. These switches make forwarding decisions in hardware, not in software. Therefore, when the switches make the forwarding or switching decision for most frames that pass through the switch, the process does not involve the Supervisor Engine CPU.

Instead, the Supervisor Engine CPU performs other important functions. The functions that it performs include:

- Assists in MAC address learning and aging

  Note: MAC address learning is also called path setup.
- Runs protocols and processes that provide network control

  Examples include Spanning Tree Protocol (STP), Cisco Discovery Protocol (CDP), VLAN Trunk Protocol (VTP), Dynamic Trunking Protocol (DTP), and Port Aggregation Protocol (PAgP).
- Handles network management traffic that is destined to the sc0 or me1 interfaces of the switch

  Examples include Telnet, HTTP, or Simple Network Management Protocol (SNMP) traffic.

The **show processes cpu** command provides information about the Supervisor Engine CPU; the switch hardware that makes the forwarding decisions does not provide this information. Therefore, the output of the command does not directly correlate to the switching performance or traffic load of the switches.

Typical show processes cpu Command Utilization

You can locate potential issues and fixes if you:

- Issue the **show–tech support** command or the **show processes cpu** command from your Cisco device.
- Use the Output Interpreter ⬀ (registered customers only) tool.

In some cases, even a switch that passes little or no traffic reports CPU utilization that is higher than is typical with other CatOS–based switches. Output of the **show processes cpu** command shows this high CPU utilization.

Note: Examples of other CatOS–based switches are the Catalyst 5500/5000 and 6500/6000 series switches.

On a Catalyst 4003, 4006, 2948G, 2980G, or 4912G switch, typical CPU utilization is 1›0 percent. On a Catalyst 4006 switch on which you have installed one or more WS–X4148–RJ45V modules, the typical utilization is higher. The typical utilization is usually 20 0 percent. The utilization is higher because these modules perform additional port monitoring in order to detect connected IP phones. The modules need to detect the connected phones so that inline power can be applied, if necessary.

As a rule, these percentages do not increase in proportion to the amount of traffic that passes through the switch. Therefore, whether the switch is completely idle or passes large amounts of traffic, the average CPU utilization percentages do not change significantly.

Typically, the highest percent utilization processes are the Switching Overhead and the Admin Overhead processes. This example shows the output of the **show processes cpu** command on a Catalyst 4006 switch with a Supervisor Engine II that runs CatOS:

Note: Some output has been suppressed for clarity.

```
Console> (enable) show processes cpu

CPU utilization for five seconds:  43.72%
                    one minute:  43.96%
                  five minutes:  34.17%

PID Runtime(ms) Invoked     uSecs    5Sec     1Min     5Min    TTY Process
--- ----------- ----------  -------- -------  -------  ------- --- ---------------
1   143219346   0           0        74.28%   56.04%   65.83% -2  Kernel and Idle
3   5237943     1313358     330000    2.84%    2.00%    2.00% -2  SynConfig
13  4378417     92798429    2000      1.97%    1.00%    1.00% -2  gsgScpAggregati
19  2692969     8548403     14000     1.23%    1.00%    1.00% -2  SptBpduRx
84  6702117     92798314    9000      2.77%    2.00%    2.00% 0   Console
97  9382372     16190292    12499     4.26%    4.22%    4.31% 0   Packet forwardi
98  23438905    7904296     9352     16.64%   19.57%   17.50% 0   Switching overh
99  2271479     1443242     57968     1.19%    1.04%    0.98% 0   Admin overhead

Console> (enable)
```

Switching Overhead is actually a process that consists of several subprocesses. The subprocesses handle these tasks:

- Address learning for new MAC addresses

  Note: MAC address learning is also called path setup.
- Normal host entry aging, as well as fast aging, due to the reception of STP Topology Change Notification (TCN) bridge protocol data units (BPDUs)
- Packet processing for control traffic, such as STP BPDUs, CDP, VTP, DTP, and PAgP
- Packet processing for management traffic, such as Telnet, SNMP, and HTTP, as well as broadcast and multicast packets in the sc0 or me1 subnets

Admin Overhead is a process for switch hardware management. Admin Overhead handles these tasks:

- Switch fabric application–specific integrated circuit (ASIC) and other hardware management
- Line card ASIC management
- Port monitoring

Causes of High CPU Utilization

As the Typical show processes cpu Command Utilization section of this document mentions, the typical CPU utilization on Catalyst 4500/4000 series switches is higher than on other CatOS–based switches. These other switches include the Catalyst 5500/5000 and 6500/6000.

However, in some cases, the Supervisor Engine CPU utilization can exceed this expected range. CPU utilization can exceed the typical ranges on the switch for these reasons:

- **Address learning** The first frame in any flow from a source MAC address to a destination MAC address is redirected to the Supervisor Engine CPU. With this redirection, address learning can occur. Once the CPU sets up the path in hardware, subsequent frames that use the same source and destination MAC addresses are switched in hardware. The CPU has no involvement. Therefore, if the CPU must learn a large number of MAC addresses in a short period of time, the CPU utilization can rise. Utilization rises during setup of the paths. The switch needs to learn a large number of MAC addresses in a short period at, for example, the start of the business day or right after lunch. At these times, many users power up their systems or log in to the network.
- **STP TCNs in the network** TCN BPDUs cause the switch to perform fast aging on MAC addresses that the switch has learned. As a typical result, many frames are sent to the CPU for address learning and path setup. Therefore, you must find the root cause of the TCNs and prevent the occurrence. These are some possible causes:

    - A port in the network that flaps
    - Hosts who power up and down on ports that do not have STP PortFast enabled
- **The receipt of excessive broadcast traffic on the management interfaces (sc0 or me1)** Broadcasts in the management subnets/VLAN must be raised high enough up the protocol stack on the switch to determine if the Supervisor Engine is the intended recipient of the traffic. Examples of traffic that can increase CPU utilization on the switch include:

    - Internetwork Packet Exchange (IPX) Routing Information Protocol/Service Advertising Protocol (RIP/SAP)
    - AppleTalk control traffic
    - Broadcast Network Basic Input/Output System (NetBIOS) frames
    - Legacy IP applications that use broadcast
- **Excessive management traffic** Certain management traffic can cause high CPU utilization on the switch. Particularly frequent SNMP polling is an example.
- **Software switched traffic** When you use the Layer 3 module, remember that all traffic that reaches the router on the native VLAN is routed in software. This situation has an adverse effect on the performance of the switch. The microcode on the WS–X4232–L3 does not process 802.1Q packets that come in on the native VLAN without tags. Instead, the packets go to the CPU, and the CPU processes the packets. This process results in high CPU utilization if the CPU receives packets without tags at a high rate on the native VLAN subinterfaces. Therefore, create a dummy VLAN (which does not contain any user traffic) as the native VLAN.

    Note: Create a dummy VLAN as the native VLAN on the trunk links between the router and the switch. The CPU routes in software all the traffic that sends on the native VLAN, which has an adverse effect on the performance of the switch. Create an additional VLAN that you do not use anywhere else in the network and make this VLAN the native VLAN for the trunk links between the

router and the switch.

Ping Latency

Another misconception is that ping response latency is the result of high CPU utilization on the switch Supervisor Engine. The response latency occurs when you ping the switch sc0 interface. Response latency is more than 10 ms.

Internet Control Message Protocol (ICMP) request and reply processing is a low−priority task on the Supervisor Engine. Many more−important tasks have precedence over ping response generation. Therefore, ping response times of 7™0 ms are typical, even on a completely idle switch. On a particularly busy switch, response times can be even longer.

However, pings through the switch are typically forwarded in hardware. In these cases, the switch sees the ICMP echo request and reply as simply data frames. The response latency consists of:

- The round−trip forwarding delay through the switch

  This is usually a very short delay, in the order of microseconds.
- The latency of the IP stacks in the process and response to the ping requests and replies
- Any other delay in the network that the ICMP packets must traverse

  An example of such delay is multiple router hops.
- Unnecessary IP redirects due to extensive use of static routing

Recommendations

The Supervisor Engine CPU utilization does not reflect the hardware forwarding performance of the switch. Still, you must baseline and monitor the Supervisor Engine CPU utilization.

1. Baseline the Supervisor Engine CPU utilization for the switch in a steady state network with normal traffic patterns and load.

   Note which processes generate the highest CPU utilization.
2. When you troubleshoot CPU utilization, consider these questions:

   - Which processes generate the highest utilization? Are these processes different from your baseline?
   - Is the CPU consistently elevated, over the baseline? Or are there spikes of high utilization, then a return to the baseline levels?
   - Are there TCNs in the network? Or are the redundant links properly configured with spanning tree parameters to avoid loops?

     Note: Flapping ports or host ports with STP PortFast disabled cause TCNs.
   - Is there excessive broadcast or multicast traffic in the management subnets/VLAN?
   - Is there excessive management traffic, such as SNMP polling, on the switch?
3. If possible, isolate the management VLAN from the VLANs with user data traffic, particularly heavy broadcast traffic.

   Examples of this type of traffic include IPX RIP/SAP, AppleTalk, and other broadcast traffic. Such traffic can impact the Supervisor Engine CPU utilization and, in extreme cases, can interfere with the normal operation of the switch.
4. Consider a switch upgrade.

For Catalyst 4500/4000 series Supervisor Engines and switches that run CatOS, consider a switch upgrade to release 5.5(7) or later.

These releases integrate several CPU−related optimizations, particularly in the area of the Switching Overhead subprocesses.

In CatOS release 6.4.4 and later, there is an extension of the management request timeout period. The timeout period extension can prevent many transient control packets timeouts that a busy CPU can cause.

Note: Releases 6.1(1) and later support the Catalyst 2980G−A.

Related Information

- **High CPU Utilization on Cisco IOS Software−Based Catalyst 4500/4000 Switches**
- **Catalyst 6500/6000 Switch High CPU Utilization**
- **Catalyst 3750 Series Switches High CPU Utilization Troubleshooting**
- **LAN Product Support**
- **LAN Switching Technology Support**
- **Technical Support & Documentation − Cisco Systems**

Updated: Aug 30, 2005   Document ID: 12032