Configure and Troubleshoot Catalyst 8000 CPU Resource Allocation

Contents

Introduction

This document describes Catalyst 8000 CPU core allocation including how to configure resource templates distribution and verify its utilization.

Components Used

This document applies to Catalyst 8000 platforms utilizing an x86-based software data plane (vQFP).

- All the commands were run on a C8500L.
- This document applies for C8500L, C8300, C8200 and C8000v.

Note: The number of cores and their IDs vary based on the model and the core distribution configuration.

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.

Configure

The Catalyst 8000 Series employs resource templates to partition physical and logical (hyperthreaded) cores. This partitioning prevents resource contention between background management tasks and high-priority packet forwarding or containerized services.

1. Resource Templates

Based on the deployment, you can choose from several templates:

- **Service Plane (SP) Heavy:** Allocates additional cores to services like AppQoE and Unified Threat Defense (UTD/Snort). This is the default mode for Cisco SD-WAN "controller" mode operation on platforms like the C8500L.
- **Control Plane (CP) Heavy:** Prioritizes routing protocol processing. Recommended for Route Reflector roles or high-scale VPN head-ends (for example FlexVPN).
- **Data Plane (DP) Heavy:** The default template for standard routing. It maximizes the cores dedicated to packet forwarding to achieve peak throughput.

2. Configuring the Template

To apply a resource template, enter global configuration mode.

<#root>

```
Router(config)#
```

platform resource ?

```
control-plane-extra-heavy Use Control Plane Extra Heavy template
control-plane-heavy Use Control Plane Heavy template
data-plane-heavy Use Data Plane Heavy template
data-plane-normal Use Data Plane Normal template
service-plane-heavy Use Service Plane Heavy template
service-plane-medium Use Service Plane Medium template

Router(config)#
```

platform resource service-plane-heavy

Note: Changing the platform resource template requires a write-memory and a reload to take effect.

Verify and Interpret CPU Utilization

When monitoring the CPU on a Catalyst 8000, the output of show process cpu platfrom sorted can show utilization near 100% on many cores. This is often by design.

Understanding ucode_pkt_PPE0 and "Hot-Spinning"

The process ucode_pkt_PPE0 represents the microcode running on the Packet Processing Engines (PPE).

- **Polling Architecture:** Unlike the Control Plane, which "sleeps" when idle, the Data Plane cores use a "polling" (or "hot-spin") mechanism. They constantly poll the hardware interfaces for new packets to process to ensure the lowest possible latency.
- **High Utilization is Normal:** Because of this polling, it is normal for Data Plane cores to show ~100% utilization even when traffic throughput is low.
- **Aggregate Percentage:** In the process list, ucode_pkt_PPE0 can show values past 100% (for example 1400%). This is an aggregate total of all cores assigned to the Data Plane.

 **Caution**: Example run on 8500L, on other platforms the core distribution can look a little different.

<#root>

```
Router#
```

show process cpu platform sorted

```
CPU utilization for five seconds: 71%, one minute: 71%, five minutes: 71%
Core 0: CPU utilization for five seconds:  2%, one minute:  1%, five minutes:  1%
```

<-- Control Plane (Idle/Normal)

```
Core 1: CPU utilization for five seconds:  1%, one minute:  1%, five minutes:  1%
```

<-- Control Plane (Idle/Normal)

```
Core 2: CPU utilization for five seconds: 99%, one minute: 98%, five minutes: 99%
```

<-- Data Plane (Hot-Spinning)

```
Core 3: CPU utilization for five seconds: 99%, one minute: 99%, five minutes: 99%
```

<-- Data Plane (Hot-Spinning)

```
Core 4: CPU utilization for five seconds: 99%, one minute: 99%, five minutes: 99%
```

<-- Data Plane (Hot-Spinning)

Core 5: CPU utilization for five seconds: 99%, one minute: 98%, five minutes: 99%

<-- Data Plane (Hot-Spinning)

Core 6: CPU utilization for five seconds: 99%, one minute: 99%, five minutes: 99%

<-- Data Plane (Hot-Spinning)

Core 7: CPU utilization for five seconds: 99%, one minute: 99%, five minutes: 99%

<-- Data Plane (Hot-Spinning)

Core 8: CPU utilization for five seconds: 100%, one minute: 99%, five minutes: 100%

<-- Data Plane (Hot-Spinning)

Core 9: CPU utilization for five seconds: 100%, one minute: 99%, five minutes: 100%

<-- Data Plane (Hot-Spinning)

Core 10: CPU utilization for five seconds: 21%, one minute: 22%, five minutes: 21%

<-- Service Plane (Active Workload)

Core 11: CPU utilization for five seconds: 7%, one minute: 4%, five minutes: 4%

<-- Service Plane (Active Workload)

Core 12: CPU utilization for five seconds: 1%, one minute: 1%, five minutes: 1%

<-- Control Plane (Idle/Normal)

Core 13: CPU utilization for five seconds: 2%, one minute: 1%, five minutes: 1%

<-- Control Plane (Idle/Normal)

Core 14: CPU utilization for five seconds: 99%, one minute: 99%, five minutes: 99%

<-- Data Plane (Hot-Spinning)

Core 15: CPU utilization for five seconds: 99%, one minute: 99%, five minutes: 99%

<-- Data Plane (Hot-Spinning)

Core 16: CPU utilization for five seconds: 99%, one minute: 98%, five minutes: 99%

<-- Data Plane (Hot-Spinning)

Core 17: CPU utilization for five seconds: 99%, one minute: 99%, five minutes: 99%

```
<-- Data Plane (Hot-Spinning)

Core 18: CPU utilization for five seconds: 99%, one minute: 99%, five minutes: 99%

<-- Data Plane (Hot-Spinning)

Core 19: CPU utilization for five seconds: 99%, one minute: 99%, five minutes: 99%

<-- Data Plane (Hot-Spinning)


   Pid    PPid    5Sec    1Min    5Min  Status        Size  Name
-------------------------------------------------------------------------------
 14571   14564   1442%   1437%   1440%  R           883704  ucode_pkt_PPE0
```

- **Cores 2-9 and 14-19:** Showing ~99-100% utilization indicates these cores are dedicated to the Data Plane and are actively polling for packets.
- **ucode_pkt_PPE0 at 1442%:** This confirms that 14 cores are currently allocated to the Data Plane/PPE and are in operational "hot-spin" mode.
- **Cores 0, 1, 12, 13:** Showing low utilization (1-2%) indicates the Control Plane is healthy and not under stress.

For specific core distribution of the Catalyst 8000 platform series it can be reviewed the links:

8200/8300 Core Distribution.

8000v Core Distribution

Check CPU Allocation

To verify how the cores are currently partitioned use this verification command:

<#root>

```
Router#

show platform software cpu allocation



CPU alloc information:
Control plane cpu alloc: 0-1,12-13
Data plane cpu alloc: 2-11,14-19
Service plane cpu alloc: 0
Slow control plane cpu alloc:
Template used: default-data_plane_heavy
```

Troubleshoot

Measuring True Data Plane Load

Because the CPU cores dedicated to the Data Plane show 100% utilization, you must use this command to see the actual processing load on the Quantum Flow Processor (QFP) :

<#root>

```
Router#
```

show platform hardware qfp active datapath utilization

```
CPP 0: 5 secs 1 min 5 min 60 min
Input: Total (pps) 62 71 75 73
(bps) 399280 514352 572520 559440
Output: Total (pps) 61 71 75 73
(bps) 391904 514648 573408 560424
Processing: Load (pct) 7 8 8 8

Crypto/IO
Crypto: Load (pct) 0 0 0 0
RX: Load (pct) 0 0 0 0
TX: Load (pct) 10 9 9 9
Idle (pct) 90 90 90 90
```

What to look for:

- **Processing: Load (pct):** This is the most critical metric. In the example previous, the load is only **7-8%**. This means that despite the CPU cores showing 100% (hot-spinning), the router actually has over 90% of its data plane capacity remaining.
- **Crypto: Load (pct):** Shows the utilization of the hardware encryption engines. If this is high, the device is heavily loaded with VPN/IPsec traffic.
- **Input/Output (pps/bps):** Use these to correlate traffic spikes with the Processing Load.

Identifying Congestion

- **QFP Drops:** If the "Processing: Load (pct)" is consistently high (>80%), check for drops using show platfrom hardware qfp active statistics drop.
- **Control Plane Health:** Cores 0, 1, 12, and 13 do not hot-spin. If these cores show high utilization, it indicates high Cisco IOS features or routing protocol load (for example BGP convergence, SNMP polling, voice signaling, etc).
- **Service Plane Monitoring:** Cores 10 and 11 (in the example) show actual workload for services like Snort. If these reach 100%, the Service Plane is saturated, even if the Data Plane (QFP) load is low.