Table of Contents

Understanding How Round Robin Algorithms Work on a Cisco CSS 11800

Introduction

The round robin algorithm may not appear to work correctly in a Cisco CSS 11800 series content services switch. An example of such a problem could occur when you issue the **show summary** command. The results may not show an even distribution across servers.

Prerequisites

Requirements

There are no specific requirements for this document.

Components Used

The information in this document is based on the following hardware versions:

- Cisco CSS 11800 series content services 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, see the Cisco Technical Tips Conventions.

Understanding How Round Robin Algorithms Work on a Cisco CSS 11800

The CSS 11800 combines centralized processing and memory resources for policy−based flow setup with distributed network processors that support wire−speed flow forwarding.

The CSS 11800 Content Policy Engine consists of the following:

- Four, one million instructions per second (MIPS) reduced instruction set computing (RISC) processors, which also are known as four switch fabric processors (SFPs)

- 512 MB of memory
- 20 Gbps of throughput

By distributing the incoming requests across the four SFPs, the CSS 11800 can handle high volumes of flows without sacrificing performance.

A CSS 11800 that is populated with two switch fabric modules (SFMs) gives the system a total of four SFPs, two of which are on each SFM. Incoming connections are distributed across the four SFPs for initial flow setup for greater overall performance. The exact distribution across SFPs is established by performing an X−OR on the protocol source and destination ports of the incoming request and binding it to one of the SFPs based on the last two bits of the result.

This distribution results in a statistical balance across the four SFPs, each of which runs its own instance of the load−balancing algorithm. In most cases, a perfect round robin load−balancing occurs because most applications use sequential ports, which results in perfect distribution across servers.

It may be possible to hit a different SFP, which could result in the same server being hit in succession because a separate instance of the load−balancing algorithm is being used per SFP. While this behavior can be noticed if a small number of nonsequential streams are used, it should result in better distribution as more streams are added. In real−world scenarios in which there are many flows, all four processors are used. Distribution across the servers should be fairly even.

Related Information

- **Cisco CSS 11000 Series Product Support Page**
- **Technical Support − Cisco Systems**