

DLSw+ TCP Performance

Overview

Cisco Systems performed a series of tests to determine the percentage of the CPU utilized on various Cisco router platforms as a function of data frames transported between two Data-Link Switching Plus (DLSw+) TCP peers. This data can help customers make a ballpark comparison between the processing capabilities of different router model types and a determination of how many data frames per second (fps) a particular router platform can support.

The routers tested were configured for very little other than DLSw+ (for instance, loopback interfaces were not used because that would have required the use of a routing protocol). Obviously in any real network environment there will be additional overhead for other configured features and protocols that must be factored in.

In addition, the results represent the load on the router in a very static controlled lab environment. Different frame sizes can be packaged by TCP differently, and other factors such as explorer load can also affect router CPU utilization (even for the same frame rate, router performance could be better or worse).

Sizing for determining the appropriate Cisco router platform is almost completely dependent on the amount of traffic and packet rate. From a historical perspective with Cisco DLSw+, this is predicated on how many packets per second of Logical Link Control, type 2 (LLC2) traffic a router has to forward. Increasing the number of SNA physical units (PUs) or DLSw+ peers and having no LLC2 traffic has very little impact on router CPU utilization. However, when sending frames across DLSw+ peers, router CPU utilization rises proportionally to the number of frames per second being sent inbound and outbound.

The most important information required to make an accurate router sizing determination is the peak rate of transactions per second (tps) and the transaction profile (number of bytes inbound and outbound). If the transaction rate is combined with the SNA response unit (RU) size information, then the rate of frames per second and the number of frames per transaction can be calculated. For example, if through analysis of the transaction profile it is determined that the transaction size is 2000 bytes outbound and 1000 bytes inbound with an RU size of 1024 bytes, then basically you have two frames out and one frame in for a total of three frames.

Lastly, although CPU utilization on most of the router platforms tested was driven to 70 percent router CPU utilization (and higher), note that it is recommended to design DLSw+ routers to utilize no more 50 percent of the CPU. This is especially true if there are multiple active peers configured on the DLSw+ routers for availability and redundancy. If redundancy is important, determine an acceptable fault-domain, double it, and peer that number of sites to a pair of redundant peers (half on each) that back each other up. Use peer load balancing


between the redundant peers if load balancing between DLSw+ peers is also desired. For example, if 250 is determined to be the acceptable fault-domain, then put 250 peers on each redundant Cisco DLSw+ router so that effectively they back each other up. If one fails, you then have 500 sites on the remaining DLSw+ peer router.

For more information about designing DLSw+ networks, consult the *DLSw+ Design and Implementation Guide* at www.cisco.com/warp/public/cc/pd/ibsw/ibdlsw/prodlit/toc_rg.htm.

The Test Environment


Figure 1 illustrates the basic Token Ring test environment that was used. Router 1 was the unit under test and Router 2 was either a Route Switch Processor 4 (RSP4) in the case of the Network Processing Engine (NPE)-150 and all smaller routers, or an RSP8 with the remaining routers. Wavetek Wandel Goltermann Domino was the test tool used to generate SNA traffic. We discovered during our testing that SNA sessions would not remain stable on the Domino when the data rates exceeded 700 fps on a Token Ring LAN. As a workaround to this problem, one Domino was attached to two Token Ring interfaces (400 fps per each interface) when it was required to drive data rates in excess of 700 fps.

Figure E-1 Token Ring: Branch Routers



The Ethernet topology (shown in Figure 2) was the same except both the LAN and WAN were replaced with Ethernet. There was no problem with running all the Ethernet traffic over one interface; however, at high data rates, we discovered that sessions could not be sustained using 10 MB Ethernet. For the NPE-300 and RSP8, 10 MB Ethernet was replaced with 100 MB Ethernet and these problems were eliminated. The CPU utilization was measured on both 10 MB and 100 MB Ethernet for frame rates up to 2000 fps to ensure that results were consistent between them.

Figure E-2 Ethernet: Branch Routers





Because the Cisco 1600 and 1700 Series routers do not have multiple LAN interfaces, a V.35 serial cable was used for the WAN connection. Data was also taken with a serial WAN on the Cisco 2600 and 3600 Series routers to tie the Cisco 1600 and 1700 data to the rest. For the small routers, Ethernet hubs and Token Ring media access units (MAUs) were used on the LAN segments. However, because of the problem with Token Ring and to simplify the physical setup of the test bed, the MAUs were replaced with Catalyst® switches. A Catalyst 5500 was used for the Token Ring LAN network and a Catalyst 3900 was used for the Token Ring WAN network. A Catalyst 2900XL was used for the Ethernet WAN, and two Catalyst 2900XLs trunked together were used for the LAN interface. This was only to simplify the physical setup. All Dominos were attached to one Catalyst 2900 in one lab, and all routers were attached the second Catalyst 2900 in the other lab. Then only one wire needed to be run between the two labs.

For the Route Switch Module (RSM), Multilayer Switch Feature Card (MSFC), and Route Switch Feature Card (RSFC), the Catalyst 2900 was trunked not to another Catalyst 2900, but to its respective switch. Because of the problem with Token Ring, separate VLANs were used with each Domino on the Catalyst 3900. However, in the Ethernet environment one VLAN was used for all the Dominos on the host side, and a separate VLAN was used for all the Dominos on the terminal side.

Nothing was configured on the routers except what was essential to take measurements. Therefore, the configurations for all the routers were the same. Each router was configured with local and remote peer statements. For Ethernet, a bridge group was configured on the LAN interface. For Token Ring, a separate ring group was configured for each LAN interface. Additionally, the load interval on all interfaces was set to 30 seconds, and a clock rate of 800,000 bits per second (bps) was used on the serial interfaces. No tuning was done on the routers; all parameters (LLC2, TCP, DLSw+, and so on) were set to the default.

Version 2.4 of the Domino Analyzer program and Version 1.2 of the SNAGEN application, which runs under the analyzer, were used.

The following LLC2 parameters were used on the Domino:

- T1 timer: 2s
- Ti timer: 30s
- Retries: 8 times
- Test frame: 10s
- Activation delta: 50ms

The host side Domino sent 128-byte data frames, and a 128-byte definite response frame was sent by the terminal side Domino. The host side Domino was always on the RSP4 and RSP8 side, and the terminal side Domino was connected to the unit undergoing testing. A different host address was assigned to each host Domino. The Dominos exchanged fixed (PU 2.0) exchange identification (XID) packets. PUs and logical units (LUs) were combined to generate data in the following way. For all data rates less than 800 fps, one LU per PU was used. For data rates that were multiples of 800 fps, four LUs per PU were used. For intermediate data rates, four LUs per PU were used up to the multiple of 800, and one LU per PU was used for the remaining. For example, the following frame rates consist of the following PUs per LUs. Remember that for each LU there is one command frame and one response frame sent per second.

- 24 fps: 12 PUs, one LU per PU
- 50 fps: 25 PUs, one LU per PU
- 100 fps: 50 PUs, one LU per PU
- 800 fps: 100 PUs, four LUs per PU
- 900 fps: 100 PUs, four LUs per PU; 50 PUs, one LU per PU
- 1600 fps: 200 PUs, four LUs per PU
- 1700 fps: 300 PUs, four LUs per PU; 50 PUs, one LU per PU

Initially the CPU utilization data was taken from the router manually by executing the `sh proc cpu` command. For the branch routers the CPU load stabilized quickly. Checks were done to ensure that the five-minute average converged to the one-minute average. After doing this we used the one-minute average after seven minutes, unless there was a large discrepancy between the five-second and one-minute averages. In those cases we waited for the one-minute and five-minute averages to converge.

For the larger routers, however, there was more variation in the CPU load. For these, a script recorded all the statistics, along with the interface and TCP statistics approximately every 20 seconds (we then performed manual averaging of the data). The interface statistics for the WAN and LAN were used to confirm that DLSw+ router was processing the correct amount of data.

Results

Figure 3 compares the number of Token Ring data frames per second processed with the corresponding router CPU utilization for the Cisco 3640, 4700, Catalyst 5000 RSM, RSP2, and NPE-150 hardware platforms. Figure 4 compares the number of Token Ring data frames per second processed with the corresponding router CPU utilization for the Cisco 4700, NPE-150, NPE-200, NPE-300, RSP4, and RSP8 hardware platforms.

Figure E-3 DLSw+ TCP Performance (Token Ring): Platforms Set A


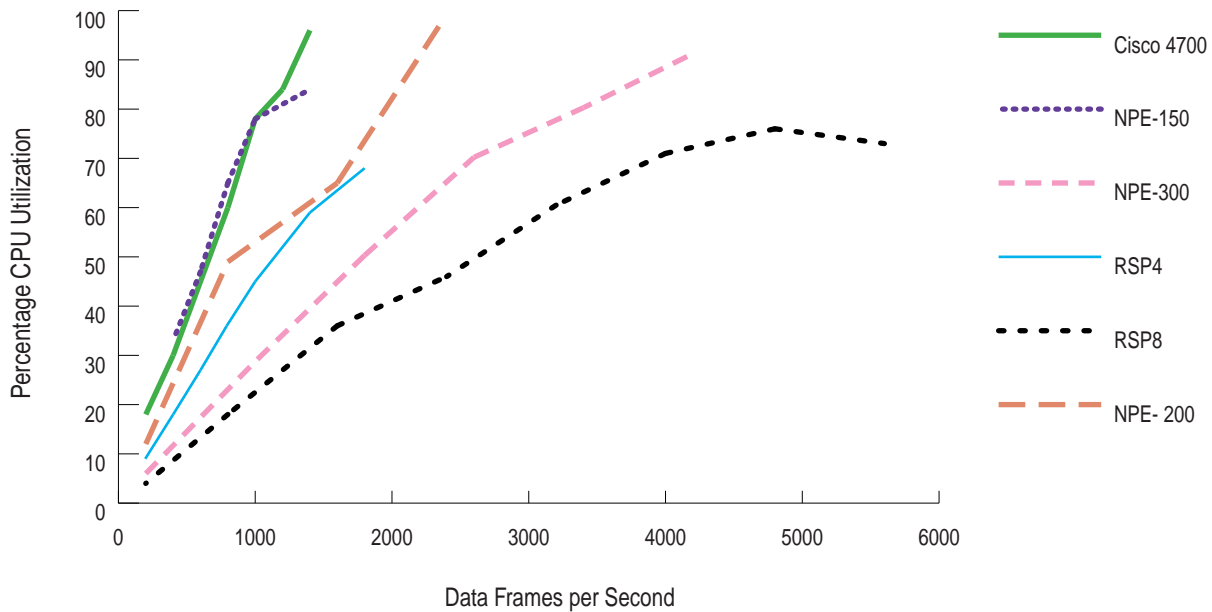


Figure E-4 DLSw+ TCP Performance (Token Ring): Platforms Set B



Note: When the RSP4 was tested, there were not enough Dominos available to fully utilize the router. Therefore, a CPU utilization result of less than 100 percent does not indicate that the router is incapable of handling more traffic.

There was a decrease in the CPU load when using switches instead of hubs, because frames arrived faster via the switch and TCP was able to encapsulate more data frames in one TCP packet. However, this phenomenon was not observed for data rates less than 1000 fps, so there was no inconsistency between the data taken over hubs and that taken over switches.

Figure 5 compares the number of Ethernet data frames per second processed with the corresponding router CPU utilization for various low-end Cisco hardware platforms (Cisco 1600 with serial WAN, 2600, 2600 with serial WAN, 3640, 3640 with serial WAN, RSP2, RSM, and Cisco 1720 with serial WAN). Figure 6 compares the number of Ethernet data frames per second processed with the corresponding router CPU utilization for a second set of low-end Cisco hardware platforms (RSM, RSP2, Cisco 4700, RSFC, MSFC, NPE-200, RSP4, NSFC2, NPE-300, and RSP8).

Figure E-5 DLSw+ TCP Performance (Ethernet): Platforms Set A


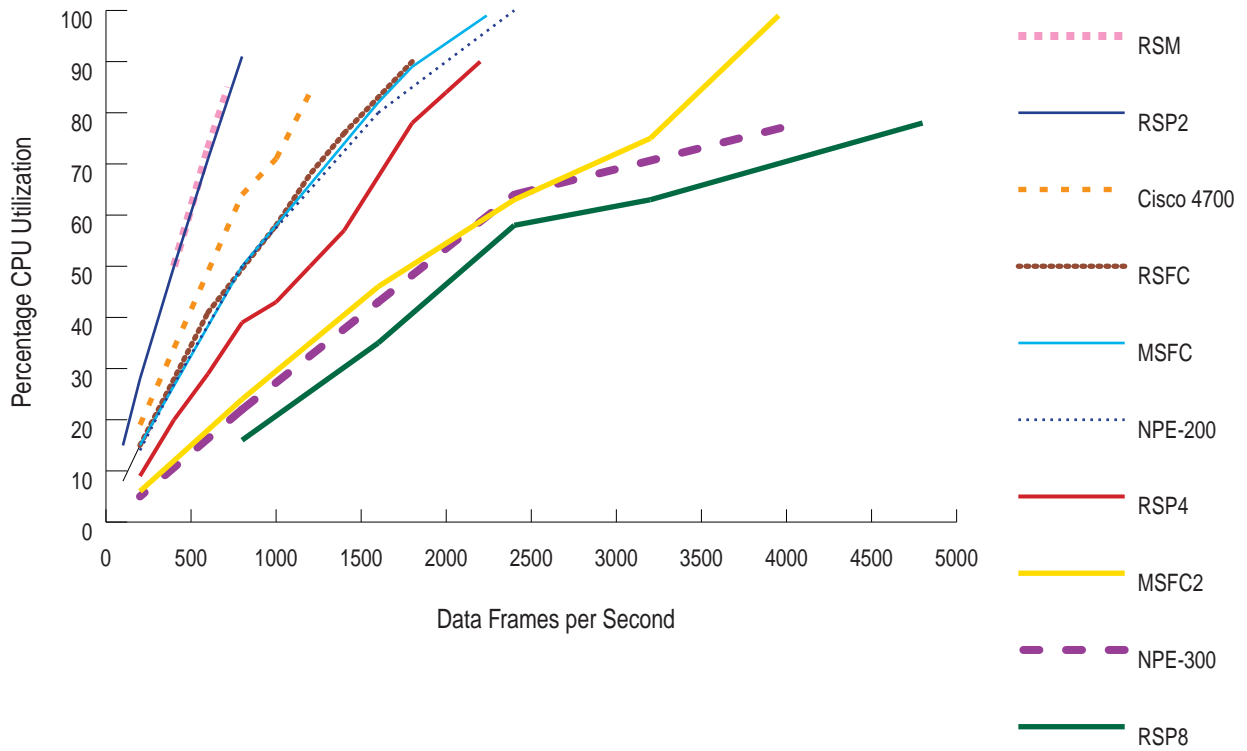


Figure E-6 DLSw+ TCP Performance (Ethernet): Platforms Set B





Version Information

For the Cisco 1720 modular access router, Cisco IOS Release 12.1(3.3) was used; for all other routers Release 12.1(3.1) was used. Because there was no official image for the MSFC, a special image built off the Release 12.1(3.1) code was used.

