Network devices: Is more expensive better?

LUCIANO AHUMADA,
MARTA BARRIA,
AND MARCELO MARABOLI

Network devices are fundamental items in computer network design. Switches are devices that filter and send packages between, from, and towards different segments of a computer network. These switches work in the second (and some in the third) layer of the Open Systems Interconnection (OSI) reference model. The price of commercial switches vary from ten to thousands of dollars, depending on the administration features, capacity, provided functions (Layer-2 or Layer-3 switching), technical support, mean time before failure (MTBF), among others. Hence, the election of the best device must be in terms of the ratio price/performance since there is generally a budget limit to build a computer network. Therefore, an engineer's goal is to build the best network staying within budget constraints. Our goal is to determine if typical low-priced Layer-2 switches experience the same packet loss rate, throughput, and jitter as the more expensive Layer-2 switches. Assuming that both types of devices develop the same performance and that the advanced administration capabilities are not a crucial point in the network design, engineers could potentially save thousands of dollars and obtain the same result.

What to compare

Empirical results presented in technical reports are a typical way to compare the performance of any device, but those results must follow strict guidelines to avoid external factors that may alter the results. Since our goal is to compare the packet loss rate (PLR), throughput, and jitter of switches from different vendors and with different prices, then a standard methodology to evaluate these parameters and make the final comparison is needed. The Benchmarking Methodology Working Group (BMWG), part of the Internet Engineering Task Force (IETF), published the "Benchmarking Methodology for Network Interconnect Devices," also known as RFC2544, which defines a specific set of tests that may be used to describe the performance characteristics of a network interconnecting device. The results of these tests will provide the user with comparable data to evaluate these devices. RFC2544 also discusses and defines specific formats for reporting the results of the tests and conditions that should be included for specific cases and also gives additional information about testing practices.

Following the RFC2544 recommendations, the first step is to choose the network testing configuration and the required software to generate the data. The elected network measurement setup is shown in Fig. 1.

We used the Iperf tool to send data according to the methodology exposed in the RFC2544 from the Iperf client to the device under test (DUT) and from the DUT to the Iperf server. Iperf is recognized as a powerful tool capable of determining statistics of different network configurations, being able to measure maximum sustained Transmission Control Protocol (TCP) bandwidth, allowing the tuning of various parameters and User Datagram Protocol (UDP) traffic measurements. This behavior represents a stress environment for a Layer-2 switch that emulates a high load user. Both the Iperf client and server were installed on separated desktop computers (Dell with a Celeron 2 Ghz processor), using an nonintegrated PCI Ethernet network interface card (Dlink DFE 550-TX) running Fedora Core 1 Linux O.S. We completed more than 5,000 measurements changing the DUT and offered load from 1 Mb/s to 100 Mb/s according to the recommendations of the RFC2544 using separate transmitting and receiving devices. Connections are made from the sending port of the Iperf client to the receiving port of the DUT

0278-6648/05/$20.00 © 2005 IEEE

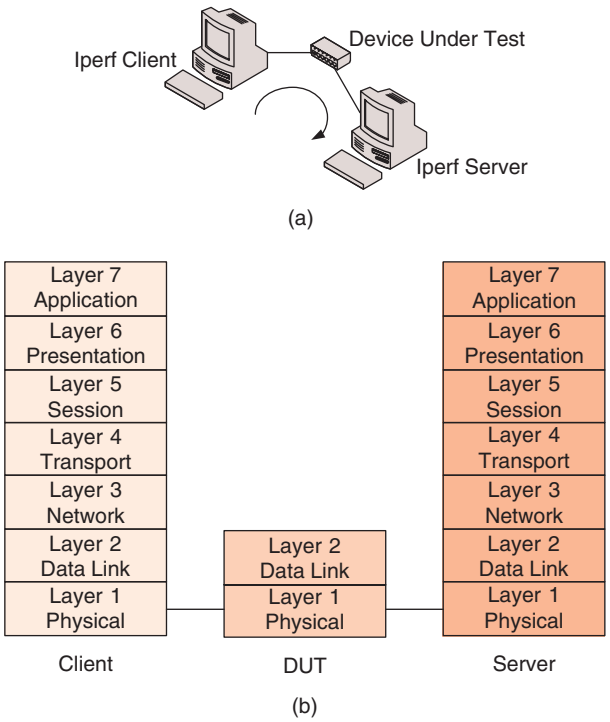

Fig. 1 Network configuration: (a) network measurement setup and (b) equivalent OSI model

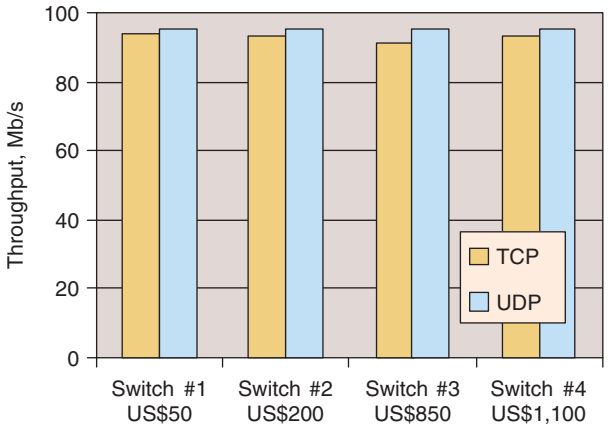

Fig. 2 TCP and UDP throughput for different devices

and from another port of the DUT to the Iperf server. Once the client stops the transmission and the Iperf server finishes its work, the Iperf server processes the throughput, PLR, and jitter statistics. Packet size and transmission rates used in this article followed the recommendations of RFC2544 for full duplex Ethernet 100BaseTX.

Table 1. Devices under test.

Item	Cost (US$)	Number of ports	Transmission rate (Mb/s)	SNMP	Web administration capabilities
Switch # 1	50	8	10/100	No	No
Switch # 2	200	24	10/100	No	Yes
Switch # 3	850	24	10/100	Yes	No
Switch # 4	1,100	24+2	10/100 + 2 10/100/1,000	Yes	Yes

Empirical results

We tested four different devices, with prices fluctuating from US$50 to US$1,100. The least expensive device has only eight ports without any kind of administration capabilities. The most expensive switch has 24 ports operating at 10/100 Mb/s, plus two 10/100/1,000 Mb/s ports. The details of these Layer-2 switches are shown in Table 1.

According to the previously explained procedure, we measured the maximum TCP/UDP throughput, jitter, and UDP packet loss rate. These results are shown in Figs. 2–4.

As can be seen in Fig. 2, all devices experimented, in practical terms, the same throughput (varying from 91 to 94 Mb/s). It is important to notice that the devices did not reach the maximum theoretical throughput (100 Mb/s) as a consequence of the ratio data/header (Layer-2, TCP/IP). Figure 3 shows the average value of the jitter obtained by means of the methodology described earlier. Hence, there is no significant basis to argue that the price is a relevant factor to obtain lower jitter values. In fact, low-priced switches show lower values of that parameter. In accordance with Fig. 4, PLR reports an asymptotic tendency in all cases (4%), including the analysis of different DUTs and packet sizes. According to our results, there is no empirical evidence to hold some level of correlation between price and throughput, jitter, or PLR. It was practically impossible to differentiate the effect of changing the DUT on these parameters. If PLR, throughput, or jitter are the decisive factors at the time of acquiring interconnection devices, paying US$50 or US$1,100 would generate almost identical results if benchmarking techniques are used to make the comparison.

Conclusions

There is no empirical basis to hold that a Layer-2 switch performance is related to its price if the design factors are the PLR, throughput, or jitter. Therefore, the difference in cost between Layer-2 switches must be related to administration features, number/type of ports, supported Layer-2 protocols, latency, and head-of-line blocking. Consequently, the best Layer-2 switch is not necessarily the most

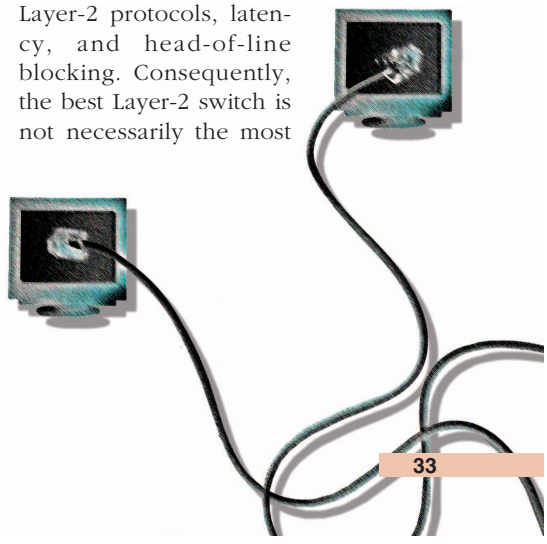

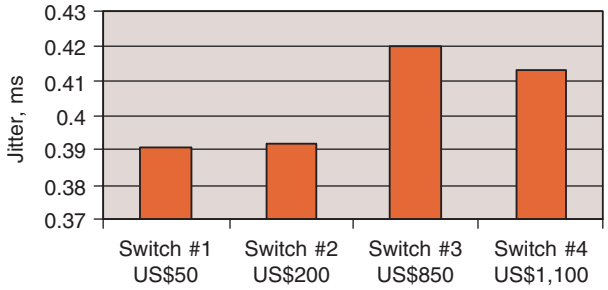

Fig. 3 Jitter statistics

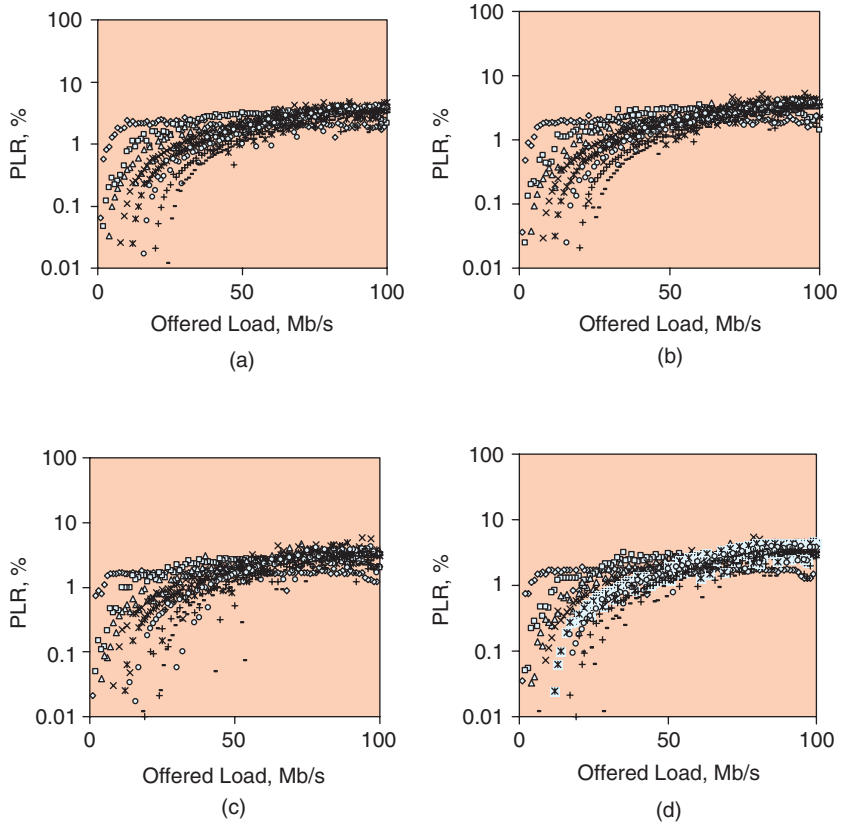

Fig. 4 UDP packet loss rate as a function of offered load for different devices and packet sizes according to RFC2544: (a) Switch # 1, (b) Switch # 2, (c) Switch # 3, (d) Switch # 4

◇ **packet size: 64 B** □ **packet size: 128 B**
△ **packet size: 256 B** × **packet size: 512 B**
∗ **packet size: 768 B** ○ **packet size: 1,024 B**
+ **packet size: 1280 B** – **packet size: 1,518 B**

expensive one, and depending on the administration features needed for each network, engineers could save hundreds or thousands of dollars by using a less expensive model and obtain the same result. For that reason, the most expensive is not necessarily better.

Read more about it

• Benchmarking Methodology Working Group [Online]. Available: http://www.ietf.org/html.charters/bmwg-charter.html

• S. Bradner and J. McQuair, "RFC2544: Benchmarking methodology for network interconnect devices," Internet RFC/STD/FYI/BCP Archives.

• A. Tirumala, F. Qin, J. Dugan, J. Ferguson, and K. Gibbs, "Iperf version 1.7.0 website," NLANR/DAST [Online]. Available: http://dast.nlanr.net/Projects/Iperf/

• P. Primet, R. Harakaly, and F. Bonnassieux, "Experiments of network throughput measurement and forecasting using the network weather service cluster" in *Proc. Computing and the Grid 2nd IEEE/ACM Int. Symp. CCGRID2002*, 2002, pp. 383–388.

• I. Elhanany and V. Tabatabaee, "Benchmarking next-generation switch fabrics," *Computer*, vol. 36, no. 10, pp. 109–110, Oct. 2003.

• A. Tanenbaum, *Computer Networks*. Englewood Cliffs, NJ: Prentice-Hall, 1996.

• R. Perlman, *Interconnections, Second Edition: Bridges, Routers, Switches, and Internetworking Protocols*. Reading, MA: Addison-Wesley, 2000.

• D. Comer, *Computer Networks and Internet*. Englewood Cliffs, NJ: Prentice-Hall, 1997.

About the authors

Luciano Ahumada received his Doctor in Electronics degree from the Universidad Técnica Federico Santa María, Chile, 2005. He is with the Escuela de Ingeniería Informática, Universidad Diego Portales, Chile. His research interests are related to wireless channel modeling and computer networks.

Marta Barría received her Ph.D. in informatics from the Pontificia Universidad Católica de Río de Janeiro, Brazil, 2001; her M.Sc. in systems engineering and informatics from the Universidad Federal de Río de Janeiro, Brazil, in 1992; and the civil engineer degree from the Universidad Técnica Federico Santa María, Chile, in 1982. She is with the Departamento de Computación, Universidad de Valparaíso, Chile. Her research interests are related to computers networks, quality of service, and performance analysis of computer networks.

Marcelo Maraboli is a M.Sc. student in electronic engineering at the Universidad Técnica Federico Santa María, Chile. He received the CISSP certification in Washington D.C, 2002. Since 1996 he has been in charge of the Network and Telecommunications Area, Universidad Técnica Federico Santa María, Chile. His research interests are related to the implementation of innovative Internet services, computer networks, and systems administration.