Know Your Traffic When Evaluating Equipment

By Peter Sevcik and John Bartlett
April 2002

Enterprises today make extensive use of the Internet, either as a resource or as a way to distribute content to their customers, suppliers and partners. The Internet connection has become a substantial asset, bringing improved communications and efficiencies. But it brings challenges as well. The challenges of constant availability, performance, and security mean a host of new functions are needed to monitor and control the enterprise Internet connection.

For each new piece of equipment, the enterprise IT manager must ensure that the device will meet functional requirements, including performance. If it were as easy as checking a specification for bit speed on each box, performance would be easy to ensure. Newer devices, however, are working at many layers of the communications stack, and as a result, performance is very dependent on the higher-level responsiveness of the device, the number of users involved, and on the applications that are being run across the Internet connection.

This report describes the characteristics of typical enterprise applications, and how to map the application mix of the enterprise to the performance of appliances required at the Internet interface.

Network Device Performance Evaluation

There is much that has to be considered when evaluating a device. First is the location and primary function in network topology. Is the device in the data path or a sideline service? In each case there will be a required throughput in both packets and bytes required to perform its function.

Inside a Network Box

The design of most networking gear has two distinct data paths, the fast path and the slow path. The fast path is designed with high-speed hardware, perhaps ASIC devices or a network processor. This path handles all the most common cases of traffic flowing through the box, and is designed to run at the maximum bit-per-second and packet-per-second capacity of the system, typically at line speeds. The number of cases that the fast path can handle is determined by the complexity (read: cost) of the implementation.

The slow path is the data path that handles any case not implemented in the fast path hardware. The fast path recognizes all cases it knows how to handle. When it finds a packet type or flow (unique packet sequence) that is beyond its capabilities, the packet processing is shunted to the slow path, which is typically implemented by a general-purpose processor. This processor has software routines that determine how to handle the particular packet or flow.

Network systems are designed so that the large majority of cases are handled by the fast path hardware, ensuring that high system throughput is usually achieved. However, particular applications may exercise the slow path of your network appliance. If this is a heavily used application, overall network performance may suffer.

NetForecast Report 5060

The Testing Methodology

Whether testing a piece of networking gear by using an industry tester or relying on a test report for a reference, it is important to match the test to the environment in which the system will be used. A simple throughput test (the fast path) may be meaningless if the limiting factor in your environment is one of the slow path algorithms.

A typical test set up is to connect two controllable traffic sources (packet engines) into each side of a device under test. One packet engine simulates client traffic while the other simulates server traffic. The devices are typically interconnected by dedicated (switched) Ethernet ports (10 Mbps, 100 Mbps, or 1 Gbps).

A simple test may be meaningless. Tests that are designed for an appropriate application mix and user demographics are better. The result of proper testing will provide capacity results in:

- Data throughput: bits per second (bps)
- Packet processing rate: packets per second (pps)
- TCP connection rate: flows per second (fps)
- Application or user state change rate: users per second (ups)

Applications Matter

Each application has a characteristic network usage pattern. How the application was written and the environment in which it is being used determines this pattern. Each application can be profiled to determine how much data it moves across the network, how many TCP or UDP connections it opens and closes, the size and size distribution of its packets, and the burstiness of its network use. These profiles are important, because network equipment is stressed by the parameters of the application, user behavior, and location in the network.

Packet Mix: Size and Frequency

It is difficult to visualize all the parameters describing the complex packet profile of an application. NetForecast has developed a unique visualization methodology called the Packet Mix Chart (PMC). Figure 1 shows such a chart for the simplest, test where the client test engine and the server test engine both drive the device under test with 64 byte packets.

The horizontal axis is actually two separate axes that are a mirror image of one another. Packet length in bytes is measured to the left of center for client packets and to the right of center for server packets. The vertical axis is marked in 10 percent increments. Each unique major packet type that is found in the application is shown in the chart. A unique packet type is a cluster of packet sizes that appear at least 1 percent of the time. These charts are derived from actual packet traces taken when an experienced user was operating the application.

A lot of information is conveyed in a PMC. The width of each “box” represents the size of the packet cluster in bytes. The height of each “box” represents the number of packets which had that size as a percentage of all packets. The “boxes” are stacked upon each other in the charts. The sum of the heights of all the “boxes” (client and server) always equals 100 percent.

The total areas of the boxes on the client side vs. the server side directly indicate the relative traffic volume each side represents in bytes transmitted.

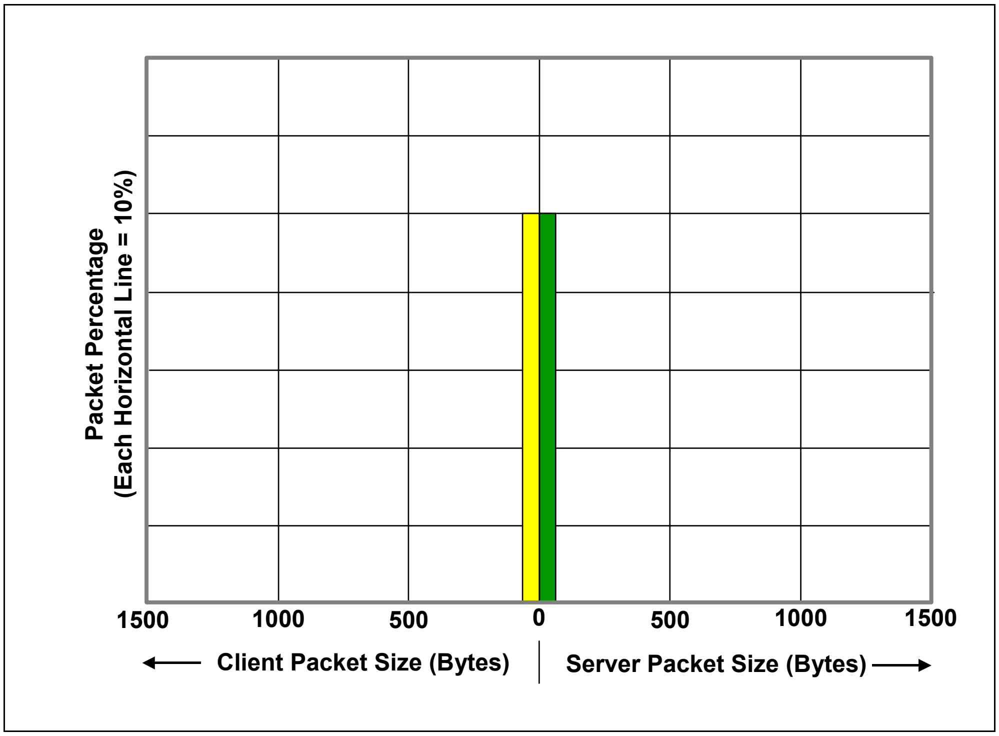

Figure 1 – Equal 64 Byte Packet Streams PMC

Figure 2 shows another typical test pattern where the packet sizes are doubled from the starting point of 64 Bytes until the MTU of Ethernet is filled (1514 Bytes).

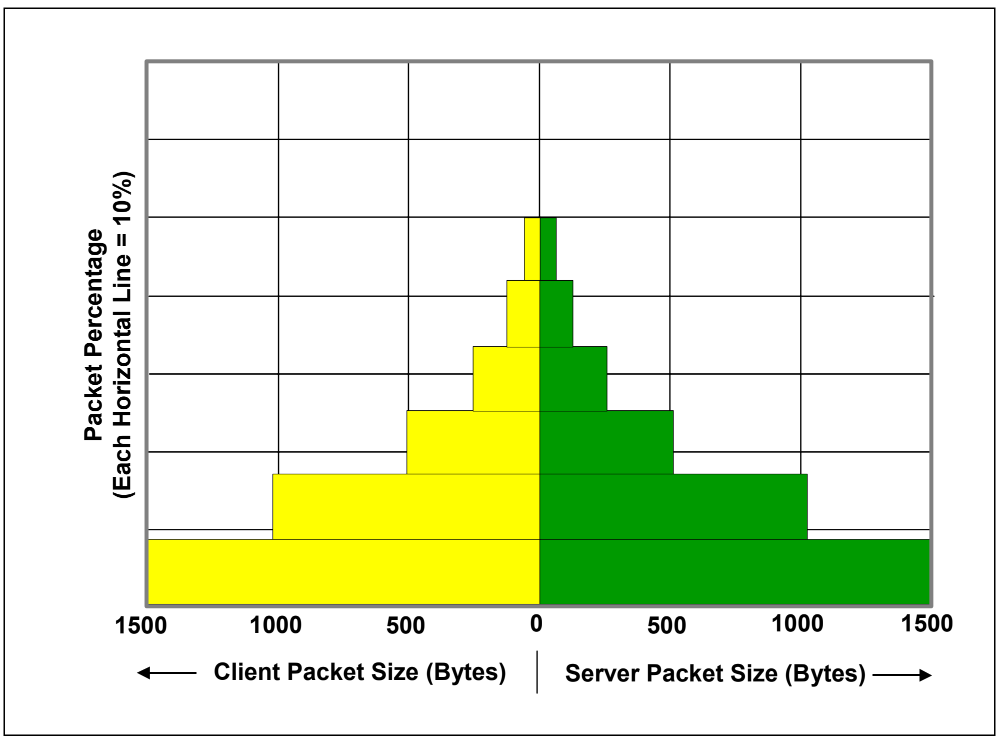

Figure 2 – Progressive Doubling PMC

Note the symmetry to the charts in Figures 1 and 2. Both client and sever send the same size packets equally often.

Figure 3 is a PMC for loading a typical Web page, the most common application on the Internet today. The Web page profile used here is a composite of the Keynote Business 40 as described in *Understanding Web Performance* [1]. The PMC in Figure 3 is substantially different than the PMCs shown in Figures 1 and 2. This illustrates how incorrect a simple test can be from reality of a true application.

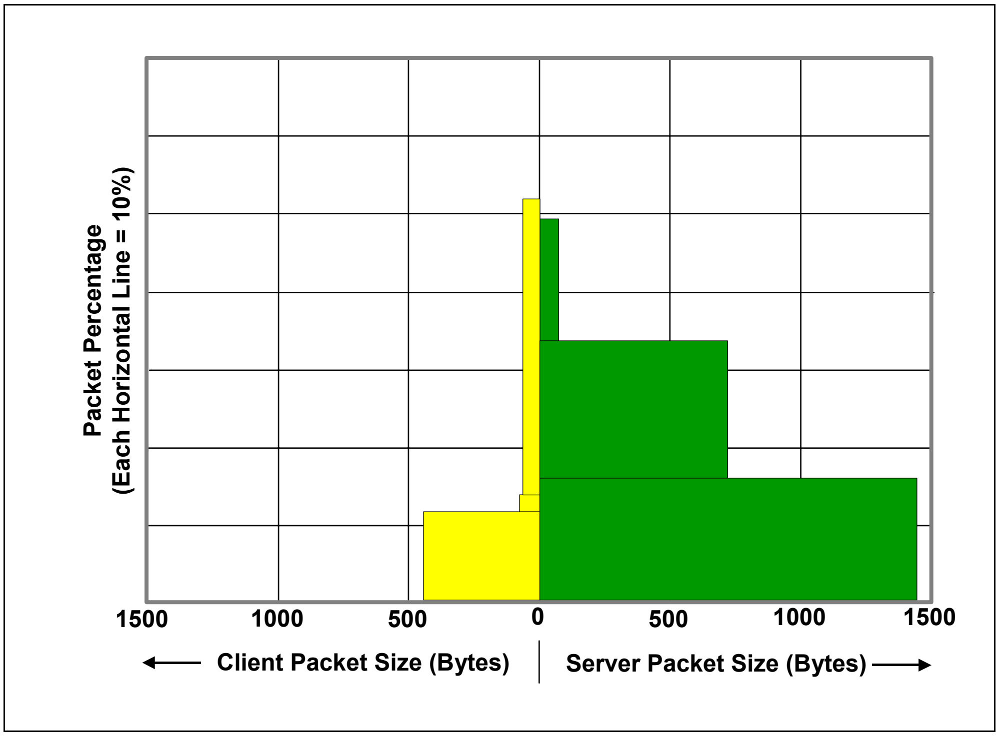

Figure 3 – Web Page PMC

Web traffic is one of the most complex forms of traffic used by the typical enterprise customer. Each click of the mouse creates many different TCP connections to numerous different computers on the network. Each request brings back a small portion of the page data, which is then assembled for the user's viewing. The next mouse click starts the process over again, often to an entirely new set of computers. The performance criteria for any network appliance that is interacting with Web traffic must not only consider data rate and packet rate, but also the number of TCP connections, and how fast they can be established.

If Web traffic is a significant portion of the enterprise Internet traffic, this packet distribution is important. For an enterprise with many users who are going to the Web for information, small packets will dominate outgoing traffic, and incoming traffic will have significantly larger packet sizes. For a content distribution enterprise, the opposite will be true.

Enterprise applications are another staple of the corporate network. Many of these applications are focused on the direct interaction of clients with a few servers. Thus the traffic for these applications has many fewer TCP connections, and often do not move much data either. However, the responsiveness of these applications is very important to their users, so the overall latency between client and server is critical. Figure 4 shows the Packet Mix Charts of four leading enterprise application platforms.

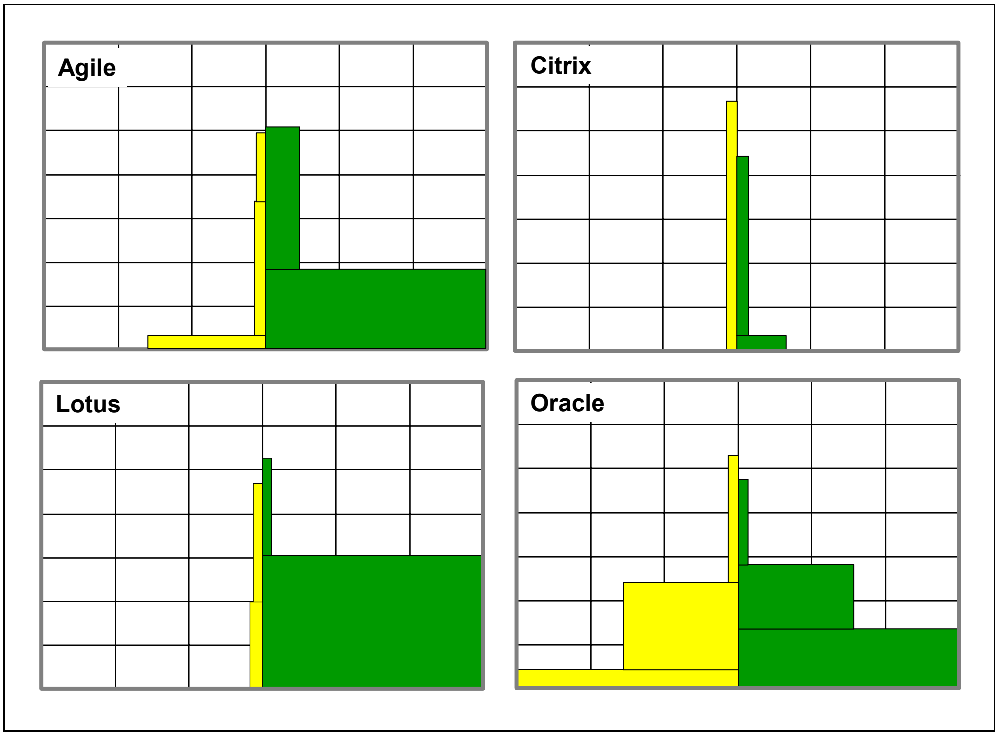

Figure 4 – Packet Mix Charts for Various Enterprise Applications

It is clear from all these charts that applications have unique profiles that are often very different from one another and vastly different from a simple single packet test.

Flow Rates

As network devices climb the protocol stack to perform more sophisticated functions than simply forwarding a packet, the work gets more complex. It is important to account for the workload associated with each function. For example, the packet rates and throughput rates are affected in peak loading by the distance between the client and server. Higher distance produces longer latency due to the speed of light limit of propagation delay and added equipment that must be traversed in longer paths.

That added latency directly affects the application response time for each task. A long task time stretches out the same traffic demand over more time. This leads to differing peak packet and bit rates per user as shown in Table 1.

Again note that each application has a very different traffic behavior in Table 1. Most notably, the peak-to-average ratios for packet rates and bit rates are very different for each application. This peak-to-average ratio is a good indication of how significantly bursts of traffic will impact traffic loads on the network.

Table 1 – Traffic Rates

	Flow Rate	Packet Rate		Bit Rate			Burstiness
Application	TCP Starts per User per min	Burst Packet Rate C-S (pps)	Burst Packet Rate S-C (pps)	Burst Bit Rate C-S (Kbps)	Burst Bit Rate S-C (Kbps)		Peak/Ave Ratio Over all Rates
National Network							
Web	24.00	15.1	26.3	15.9	131.9		2.6
Citrix	0.05	5.7	3.8	4.0	3.7		1.9
Lotus Replication	0.03	8.9	17.4	4.7	97.8		12.7
Agile	0.03	4.4	9.2	3.0	43.9		13.1
Oracle	0.13	4.5	5.3	5.7	15.8		5.3
Metro Network							
Web	24.00	47.8	83.6	50.6	418.9		8.3
Citrix	0.05	7.5	5.0	5.2	4.8		2.5
Lotus Replication	0.03	27.1	52.9	14.4	296.8		38.6
Agile	0.03	8.4	17.5	5.8	83.8		25.1
Oracle	0.13	6.4	7.5	8.1	22.3		7.5
Campus Network							
Web	24.00	86.3	150.8	91.3	755.2		15.0
Citrix	0.05	8.0	5.3	5.6	5.2		2.7
Lotus Replication	0.03	46.3	90.6	24.7	508.1		66.2
Agile	0.03	10.3	21.6	7.1	103.1		30.9
Oracle	0.13	7.0	8.2	8.9	24.4		8.2

Another intriguing aspect of flow start rates is that since remote users spread their peak demand over longer task times, they produce a lower peak demand. The consequence of this phenomenon is that you can support more individual users in the same bandwidth access line. Figure 5 shows the effect of distance on flow rates when filling an OC3 (155 Mbps) access line from a Web site.

The curves in Figure 5 are illustrative examples based upon several assumptions:

- The web servers are not a limiting factor
- All supported users are active in the busy hour
- The access line is peak utilized at 70 percent
- The access line is average utilized at 30 percent

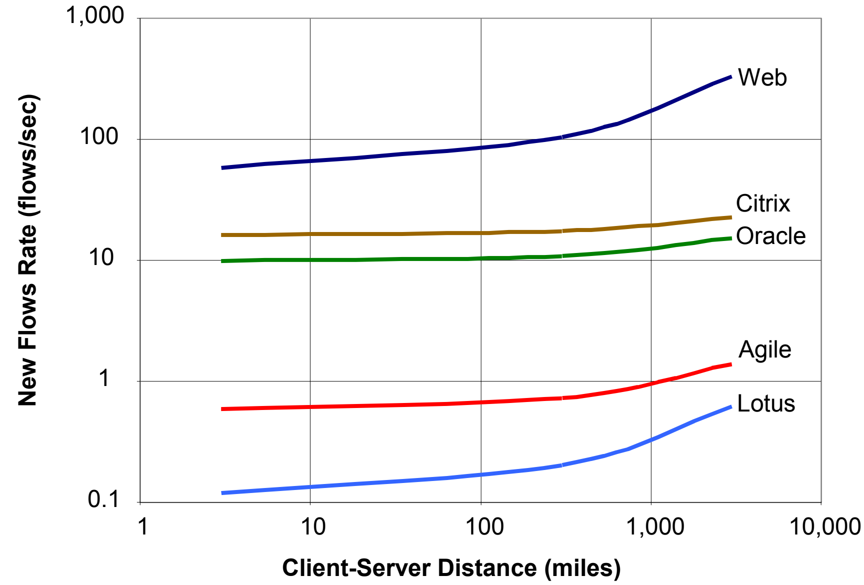

Figure 5 – Estimated Flow Rates for Active Users Filling an OC3 Circuit

Flow rates and packet rates are vastly different for each application. The Web is the most demanding application in generating new flows (TCP connections). This is but one way in which Web-based applications require careful planning.

Summary

Performance is not a single number specification. Network planners are advised to map application profiles and user characteristics into a comprehensive performance specification when selecting network equipment. If a site is supporting several applications, then the analysis must be performed upon a proper mix of profiles and users leading to a range of performance requirements. ¤

References

1 - "Understanding Web Performance" by Peter Sevcik and John Bartlett, NetForecast Report 5055, October 2001.

Peter Sevcik is President of NetForecast and a leading authority on Internet traffic and performance. He has contributed to the design of more than 100 networks. Peter led the project that divided the Arpanet into multiple networks in 1984, which was the beginning of today's Internet. He can be reached at peter@netforecast.com.

John Bartlett is a Principal at NetForecast, and has 24 years of experience designing at the chip, board, system and network levels, with a focus on performance. John led the team that built the first VLAN implementation, one of the first ATM switches, and he is a leading authority on the behavior of real-time traffic on the Internet. He can be reached at john@netforecast.com.