



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

GSEC Practical Exam Submission

Name: Laura Keadle

Version: 2.0 (revised August 13, 2001)

Comparison Shopping for Scalable Firewall Products

Introduction

Have you ever really read the User License on your network firewall software? If you're like most of us, you just wiz past it on your way to the "accept" button on the bottom. You may not realize that the user license for your firewall probably says something along the lines of:

This product is not warranted to actually do anything or to be depended upon in any way and therefore it is illegal to publish any test results related to it's performance unless we (the manufacturer) do the testing.

No Network Designer worth their salt would dream of purchasing a router or switch without demanding benchmark test results on throughput and subscription rates. After all, routers and switches represent choke points on the network where over-subscription can reduce a gigabit backbone to 10mbps crawl due to failed connections or latency caused by re-transmittal of lost packets. A really poorly chosen exterior router or switch can even cause a cascade failure of the entire network by propagating, or failing to contain, broadcast storms. At best, an over-subscribed exterior device will choke Internet connections and waste money paid for expensive bandwidth.

Router and switch performance is commonly measured in accordance with standardized testing procedures (such as RFC 2544) for meeting established packet loss criteria and throughput under a variety of load conditions for both steady state and bursty traffic.

It's important to understand the full capabilities of these critical network devices when they are idling as compared to when they are stressed. Otherwise, it's impossible to plan for future growth, let alone understand the expected behavior of the device when the next network revolution changes the demands on our Internet connections. This key issue, of planning for growth as well as the unexpected, is called scalability.

First and foremost, a firewall is essentially a networking device. It represents no less critical a potential point of failure than a router or switch. So why don't manufacturers evaluate firewalls for throughput under load? Or more precisely, why don't they evaluate these products in terms of benchmarking standards that allow the consumer to compare them against each other, and more importantly, understand the products in terms of likely performance in the networks for which they're being selected?

Evaluating Firewall Performance

Vendor documentation of firewall throughput is scarce at best, and when it is offered it's often expressed as a "rating" which is entirely subjective within the vendor's own product line or as a number of concurrent UDP sessions or VPN connections.

The most quotable vendor white paper I've come across, though certainly not the only vendor to inadequately substantiate their test methodology is from a brief available from Check Point at www.checkpoint.com/products/firewall-1/pbrief.html. CheckPoint states:

"Measuring Throughput

Throughput, the ability of a firewall to inspect and then forward network traffic, is one of the key performance indicators for a firewall. The process of inspecting network traffic and matching it to a "rule", which either permits or denies the traffic, should impose as little delay as possible on the speed of the traffic. If the firewall process cannot keep up with the speed of a given network link, bandwidth is wasted.

About The Results

FireWall-1 throughput performance results are not meant to be used as a comparison between vendor platforms. No attempt was made to match similar systems in price/performance categories. The systems tested may represent any model in a vendor's product line, from a high-end system to a low-end system. The numbers listed here are simply representative of certain models within a particular vendor's product line. (Emphasis is mine.)

Occasionally, a more forthcoming vendor will provide data comparing their own product to another by "packets per second". What troubles me about these claims is that there is not standard methodology or supporting documentation as to what kind of packet is being passed. While these kinds of measurements are not entirely without value, it's clear that there is no fair way to differentiate one vendor's claims from another.

Let's look at some of the throughput issues when selecting a network device:

Packet Size - A large packet is generally going to represent less overhead per bits of data transmitted through a networking device. This is true of routers and switches as these devices must examine the frame, in the case of a switch, and the IP header, in the case of a router, in order to select the next interface or hop for the packet's journey.

A firewall device looks further inside the packet to the TCP/UDP and transport header (or higher in the case of application level firewalls). This means that the issue of overhead and packet size is even more critical. Each and every packet must be examined by the firewall, and so when a vendor makes a claim of 100Mbps throughput, were their tests based on 1518 byte FTP file transfer packets, or were they 64 byte ICMP packets?

Another issue with packet size is that the memory buffers or queues of firewall hardware products, especially appliances or router based products, tend to be optimized for certain packet sizes. What this means is that a reputable vendor may have engineered the hardware platform to optimize a throughput for medium size packet flows in anticipation of matching the actually range of packets on an average Ethernet network.

A product which optimizes the packet flow for the size packets you need to pass, may not look as well statistically when compared to other vendors in product reviews, even though they may meet your needs better than a product which passes very large packets more quickly but performs substantially slower when processing smaller sized packets.

Bi-directional vs. Aggregate Flow – Do the firewall throughput statistics offered by the vendor represent full-duplex connectivity or half-duplex ? Most of us work with networks that are designed with Cat. 5 or better cabling and are capable of, or upgradeable to, full-duplex operation. If your internet connection is full-duplex then you probably need to know the firewall product's performance under full-duplex.

When a firewall vendor makes a claim to meet a certain throughput speed it's easy to assume that the claim is meant as full-duplex. That's often not the case. The fact is that throughput claims for 100Mbps, may mean that each of two interfaces is roaring along at 50 Mbs half-duplex and the vendor is counting both directions.

In such cases, the vendor is providing you with a number that is actually the aggregate throughput. Not that aggregate throughput isn't a useful thing to know when comparing products. After all, the vast majority of firewall products (routers, too) don't double their throughput just by adding an interface. If you have a firewall product with three 100 Mbs interfaces and each interface is capable of 100 Mbs full-duplex throughput, but the firewall is only capable of 200 Mbs aggregate throughput, adding another interface is not going to increase the throughput available for your exterior connection.

The plain truth is that the interface cards in the firewall are rarely the limiting factor in aggregate throughput. This is something you need to know before you buy the upgrade to the gigabit quad card. For example, if you have a firewall product with two gigabit interfaces, just because the interface cards are capable of gigabit speeds doesn't mean that the firewall has the processing power to pass packets fast enough to match them. The CPU, ASICS (Application Specific Integrated Circuits), motherboard bus speed and memory are more likely the limiting factors and will all be affected by the load on the device.

Number of Sessions - While a router does keep track of sessions, and often is used for some basic packet filtering, the number of concurrent sessions is vastly more critical to comparing throughput for statefull inspection firewalls. The critical issue here is the number of concurrent connections that the product can keep in it's state tables - plus it's ability to buffer incoming packets to avoid packet loss so that the right users stay connected and the bad guys get stopped.

This is one of the test conditions which is most badly in need of standardization. Performance issues must be examined in terms of load and the reliability of the connections under loaded conditions. Routers and switches are expected to meet performance criteria of 0.001% packet loss (considered zero loss) at their published throughput. Gigabit throughput speeds on a firewall are of little benefit if the network Internet connection is bogged down by excessive latency caused by dropped sessions or lost packets which must be retransmitted.

Number of VPN's and Encryption - One could easily be swayed by vendor claims that a particular firewall has exceptional processing power (and is therefore very fast) because it can maintain a high number of VPN's using 3DES and SHA-1, IPSec or other encryption schemes. If only it were that easy to assess processing power in this manner. In most cases of the highest performance with VPN's and encryption come from built-in or add-on encryption accelerator cards.

Not that these kinds of speed issues aren't important also in some cases, but when encryption is off loaded from the CPU to an accelerator card, any latency involving the encryption process tends to be limited to the encrypted traffic. It doesn't usually have much effect on the throughput of the normal traffic.

When comparing vendors, first ask yourself what you need your firewall to do for you. There have been several good papers and an entire SANS class about firewall selection that covers feature sets very well. My point in this case is not to confuse the issues of needed features with basic performance when comparing vendors.

There are two key performance issues relating to VPN's and encryption. The first issue is that if you need to handle VPN's and encryption at the firewall, by all means compare the various products for throughput under all kinds of load conditions, including number of sessions, number of VPN's, and various types of encryption.

The second performance issue concerning VPN's and encryption goes back to packet size. A program using a large packet size, for example a file transfer, may generate a 1518 byte Ethernet packet. The VPN device or firewall taking this packet through the encryption process will create a packet size which exceeds 1518 and must be fragmented.

Adding an IPSec Authentication Header, for example, may only add 4 bytes to the 1518 byte packet, but the result is still an over size packet. The first router that touches that packet on it's trek through the internet will fragment it, resulting in at least two packets of lesser size, one of which has a TCP/UDP port specified and the remainder which don't.

The performance issue you need to evaluate is whether or not the firewall products you are considering can handle the fragmentation – both going out and coming in from other networks. I was astounded to find that at least one major vendor's router based product can't handle fragmentation.

If the product you are most interested in can't handle fragmentation, then you'll want to make sure that the programs at your user end that send large packets will tolerate the latency inherent in the windowing negotiation. In addition, evaluate whether the retransmission traffic will affect your network or Internet connection. How much bandwidth are you potentially wasting due to retransmission of fragmented packets refused by the firewall? Evaluate whether the performance loss for your users in latency and the loss of exterior bandwidth for retransmission of large packets makes the performance of other products better suited for your needs.

Rule Sets - It's seemingly easy to agree that the firewall rule set, it's length and complexity will affect a product's throughput and packet loss performance. The tough question is how to decide on a standard rule set for testing purposes, so that when we look at firewall product performance criteria we're not trying to compare apples and oranges.

Vendors that are providing any performance data at all specify, usually in very tiny print, that the firewall was tested with one rule. This rule typically examines then passes UDP packets. While this seems to have the advantage of providing a somewhat level playing field between vendors, it doesn't provide the kind of information a network engineer needs to realistically evaluate product performance.

After all, the entire purpose of the firewall is centered around it's ability to filter undesirable traffic without denying bandwidth to legitimate traffic. Evaluating a firewall's performance with a single UDP rule is like evaluating the capacity of a dump truck by hauling around a bushel of leaves.

The closer test conditions can be made to match realistic network situations the more likely it is that network engineers can select the correct products and be certain about their performance. How to establish agreement about what might constitute a typical or basic rule-set for the purposes of performance testing is a subject that deserves a separate paper, as it's certainly beyond the scope of what can be included here.

Bench Marking

What's really apparent in any attempt to do comparison shopping for firewall products is the need to establish benchmarking standards and test procedures. The lack of established standards is as disadvantageous to reputable established vendors as it is to vendor's of new products and the innovation of new products.

In our earlier discussion of packet size, full or half duplex and aggregate throughput performance, we've touched on some of the issues that may actually place the more ethical vendor's at a disadvantage. Let's look at another example where this may happen.

In looking at some of the new firewall appliance products, it's clear that some of these products can process long and complex rule-sets, multiple VPN's and tens of thousands of sessions with only 10 to 25% throughput performance degradation and still meet the .001% packet loss criteria required of routers and switches. These innovative new products do this with the use of ASICS which off load the rule-set, session and VPN processing on to secondary processing chips built on to the interface cards.

In comparing one of these products to a software product (however reputable) residing on an Intel box, a consumer may very well get the impression that both of these products process packets at the same 100 Mbs throughput rate. When neither device has any load, similar bus speeds, and interface card speeds may make them perform very similarly. But the Intel based software firewall placed under a moderate load identical to the newcomer appliance may have a performance degradation of 60 to 70%.

The reason for this is that all of the Intel/software firewall's processing must be carried by the finite speed of a single CPU. As the load on the CPU increases, competing processes increase the interrupts to the processor slowing the packet throughput even further, potentially overflowing the L1 and L2 cache and causing packet loss, and subsequent dropped connections. This is essentially a denial of service condition.

What this means to the Security Engineer with the T-3 line (45Mbs) and a moderate load on the firewall, is that the one of these firewalls will provide an effective bandwidth of 75 to 80 Mbs and the other may be choking his 45 Mbs connection down to 30 to 40 Mbs. If you plan to grow in your network at all, that's going to really hurt, as well as waste money.

Disclaimer: The case of an appliance being superior to a software/Intel firewall may not hold true when comparing any specific two software or appliance devices under specific conditions. This is just a theoretical example used to illustrate how a single test result (throughput) can be very misleading without standard test conditions that are valid for comparing **both** very different hardware architectures.

In our earlier discussion of packet size, full or half duplex and aggregate throughput performance, we've touched on some of the issues which may actually place the more ethical vendor's at a disadvantage and is the dark side of benchmarking.

In my research for this paper, I contacted several major vendor's marketing representatives and was privileged to speak with some of the higher level management. Their caution was that an unethical vendor can also use poorly defined benchmarking and vague testing procedural standards to design products which excel at passing the testing, but are inferior in real network settings.

An example given to me that illustrated this well is where a standard, such as RFC 2544, specifies that the device be tested at certain packet sizes, 64, 512, 1024 and 1518 bytes. An unscrupulous vendor could design the hardware to maximize throughput for those precise packet sizes with narrowly defined memory buffers (queues), and look very good statistically in the test. While in an actual network where the data stream demands flexibility, too many of the packets would fall outside the narrow margins of the hardware engineering and packet loss in a real network would be much higher than the test would indicate.

Of course, the cure for this objection is to make sure that the test criteria include measurements for performance under loaded conditions, with a standard rule-set, varying numbers of sessions, types of protocols transmitting, and VPN's employing a variety of different standard encryption schemes.

Conclusions

The general consensus from the vendor responses I've received, and feedback from the security and network engineers I've polled here at the Dahlgren Navy base, is that is that the task of defining testing standards for firewall performance is far more complex than simply adapting RFC 2544 and other benchmarking standards for network devices. However, it's a good place to start.

In evaluating firewall throughput, Security Engineers ought to begin in the same place that a Network Designer would begin, by understanding the traffic flow through the network at the point where you want to put the device.

A relatively low speed internet connection, for example through an ISDN line, may only require a throughput of 128 kbs. This connection speed is easily handled by most of the low end firewall products on the market. In fact, in working with small business customers with throughput demands less than a T-1 line (1.5 Mbs) it's reasonable to assume that firewall feature sets, ease of configuration and vendor support are more critical concerns than throughput.

For networks with multiple T-1 lines, T-3 lines or greater, or firewalls between departments within the LAN, it's increasingly critical that the firewall product which connects the protected network to the rest of the world should not restrict the bandwidth your company is paying for and needs, or it's ability to add foreseeable future connections or bandwidth.

In the absence of established evaluation standards, keep in mind some of the issues discussed here by digging for whatever information vendor's do have available which can affect throughput under load and packet loss performance. Take a look at the benchmarking RFC's for routers and switches and consider setting up your own (non-publishable) testing procedures that would more closely and fairly match your own network conditions and architecture. For the cost of purchasing the wrong product, or the cost of mis-diagnosing a slow internet connection, the time spent by the engineers in the testbed may be well worth it.

Until we have a accepted test standards and a "Consumer Reports" for networking devices to test them objectively, we'll need to be smart and assertive consumers.

References

Brader, S., and McQuaid, J.: "Benchmarking Methodology for Network Interconnect Devices" RFC 2544; March 1999; URL: <ftp://ftp.isi.edu/in-notes/rfc2544.txt>

Dawkins, S., et. al.; "End-to-end Performance Implications of Links with Errors", RFC 3155; Aug. 2001; URL: <ftp://ftp.isi.edu/in-notes/rfc3155.txt>

ICSA Labs; "ICSA 3rd Annual Firewall Buyer's Guide", copyright 1998; URL: http://www.icsalabs.com/html/communities/firewalls/buyers_guide/index.shtml

ICSA Labs; "Firewall Product Certification Criteria Version 3.0a", Dec. 29, 2000; URL: http://www.icsalabs.com/html/communities/firewalls/certification/criteria/criteria_3.0a.shtml

Network World Fusion, "Firewall Buyer's Guide"; July 1999, URL: <http://www.nwfusion.com/bq/firewalls/firewalls.jsp>

Newman, D. ; “Benchmarking Terminology for Firewall Performance”, RFC 2647; Aug. 1999;
URL: <ftp://ftp.isi.edu/in-notes/rfc2647.txt>

Paxson, V. et.al; “Framework for IP Performance Metrics”, RFC 2330; May 1998;
URL: <ftp://ftp.isi.edu/in-notes/rfc2330.txt>

Shultz, K.; “Firewall Appliances – Look Ma! No Moving Parts”; July 5, 1999, Internet Week Magazine, URL: <http://www.internetwk.com/reviews01>

Smith, Gary; “A brief Taxonomy of Firewalls – Great Walls of Fire”, SANS Reading Room, May, 18, 2001; URL: http://www.sans.org/infosecFAQ/firewall/firewall_list.htm

Snyder, J.; “Hot Firewalls Finding New Niches”; Network World, July 7, 1999;
URL: <http://www.nwfusion.com/reviews/0719rev.html>

Sun, N. Freed; “Behavior of and Requirements for Internet Firewalls”, RFC 2979; Oct. 2000;
URL: <ftp://ftp.isi.edu/in-notes/rfc2979.txt>

The Tolly Group; “Firewalls RFP”; Network World, July 1999;
URL: <http://www.nwfusion.com/reviews/0719rfp.html>

Zwicky, E. D., Cooper, S., and Chapman, D. B.; “Building Internet Firewalls, 2nd Edition”; copyright June 2000, O'Reilly & Associates, Inc.; ISBN 1-56592-871-7

Credits:

Research for this paper was made possible by the assistance of the many knowledgeable sales engineers and excellent white papers, on-line brochures, and product guides available from the following companies:

www.3com.com
www.cisco.com
www.federated.com
www.gnatbox.com
www.lucent.com
www.netopia.com
www.sonicssys.com
www.watchguard.com
www.nfr.com
www.toplayer.com
www.netscreen.com
www.checkpoint.com
www.truesecure.com

and others.