NETWORKWORLD Reprint

THE CONNECTED ENTERPRISE ≡ JULY 22, 2013

F5 data center firewall aces performance test

BIG-IP 10200v combines firewall, load balancer, DoS protection in one fast, scalable appliance

BY DAVID NEWMAN, NETWORK WORLD

Huge data center, check. Multiple 10G Ethernet pipes, check. Load balancer, check. Firewall? Really? Do network architects need to buy yet another box, and likely take a performance hit?

Not according to F5 Networks, which says its BIG-IP 10200v with Advanced Firewall Manager (AFM) can handle traffic at 80-Gbps rates while screening and protecting tens of millions of connections, and simultaneously load-balancing server traffic.

In this exclusive Clear Choice test, we put those claims to the test. The F5 firewall came up aces, maxing out network capacity while also offering sophisticated filtering and attack protection capabilities. In some cases, traffic rates were higher with the firewall in place than without, probably because the F5 device managed server loading more efficiently.

Although F5 is mainly known for its BIG-IP application delivery controllers (ADC), the company has steadily been adding to its security suite, especially for the data center. The BIG-IP 10200v, introduced early this year, is the second largest member of the family, with 16 10G Ethernet interfaces and two 40G Ethernet interfaces in a 2RU form factor. The only larger unit is the chassis-based Vipiron 4800.

NETRESULTS

PRODUCT	BIG-IP 10200v
COMPANY	F5 Networks
PRICE	With ADF bundle, $118,995; 10G transceivers, $12,800; price as tested, $131,795
PROS	High performance; no need for separate load balancer; solid DoS protection
CONS	Minor number of connection failures; vulnerability scanner sold separately

While most of the attention in next-generation firewalls has focused on client protection, F5 targets the BIG-IP 10200v mainly for data-center use, protecting servers. Adding stateful firewall capability at very high rates is one part of that strategy.

Another part is iRules, an existing feature that allows users to inspect, modify, and reroute traffic based on HTTP and HTTPS headers. iRules use a syntax similar to many scripting languages. For network managers without scripting skills, the F5 appliance includes some canned iRules for common tasks, such as redirecting HTTP requests to HTTPS, or preventing Windows Mobile users from being locked out when they've changed their passwords in Active Directory but not on their mobile devices.

Both firewall and iRules can be configured via command-line or Web interfaces. The Web interface will look familiar to anyone who's used F5 load balancers. We don't have a lot of experience with F5 gear, but found the Web UI generally easy to navigate.

Another server protection feature is built-in denial-of-service attack (DoS) protection. The device includes nearly 40 DoS filters, all enabled by default. These filters work at layers 2-4, and cover both IPv4 and IPv6. (The firewall also works with IPv6 traffic, but time constraints limited us to testing with IPv4 traffic.)

More DoS protection comes from the IP-Intelligence feature, which identifies and blocks IP addresses for various classes of threats. Using information from a worldwide sensor network, IP-Intelligence can block traffic from botnets, Windows exploits, phishing exploits, and other classes of threats. IP-Intelligence is not enabled by default, and we did not use it in performance testing.

These features are included as part of F5's Advanced Firewall Manager (AFM) package. F5 separately sells an Application Security Manager (ASM) package, which includes application inspection and intrusion detection, but we did not test this. So, the BIG-IP 10200v is best suited to end users who want to merge firewall and load balancer in one appliance. However, if you're looking for an all-in-one security device, you'll need to buy the additional ASM package.

HOW FAST?

We tested firewall performance in terms of speed and scalability (see "How We Did It" below). In some tests, the Spirent Avalanche traffic generator/analyzer offered fixed object sizes, which are useful in determining absolute maximum speeds. We also configured Avalanche to offer a mix of Web object sizes and content types just as network managers would find in production networks.

In one of the fixed-object tests, Spirent Avalanche exclusively offered 10-kbyte objects. Numerous studies have shown the average object size of all Web transactions is somewhere near that figure. If anything, that average is trending downward, driven by AJAX-heavy Web apps. And to determine the highest possible rates, we also conducted tests using 512-kbyte objects. At this size and up, the transaction overhead involved in HTTP is negligible.

Since more and more Web traffic uses encryption, we ran all the speed tests with plaintext traffic and again encrypted with Secure Sockets Layer/Transport Layer Security (SSL/TLS), using HTTPS. We then repeated the SSL/TLS tests with decryption enabled.

In tests involving static object sizes, the F5 firewall came close to maxing out our test bed's network capacity. With 10-kbyte plaintext Web objects, the F5 firewall moved traffic at 78.630Gbps, almost saturating the 80-Gbit/s pipes between clients and servers. With 512-kbyte plaintext Web objects, the rate was 80.519Gbps. (These forwarding rates are aggregates of traffic in both directions, so it's possible to exceed 80Gbps due to TCP acknowledgements sent back from clients to servers.)

The F5 firewall moved static objects over SSL at rates that met or exceeded the capacity of the Avalanche test tool, moving 10- and 512-kbyte objects at 17.288G and 20.919Gbps respectively. Both numbers are at least 1Gbps faster than those for the Avalanche tool running back to back with no firewall inline.

The most plausible explanation for the difference is that, like all BIG-IP appliances, the 10200v is a load balancer. By performing web server health checks and distributing requests accordingly, the F5 firewall is able to distribute workloads more efficiently than clients and servers can do on their own.

In the mixed-object tests, the BIG-IP 10200v moved plaintext traffic at 37.486Gbps. That's almost 99.5% the capacity of the Spirent Avalanche traffic generator when running the same test in a back-to-back configuration.

When running the same test with SSL traffic, the F5 firewall moved traffic at 12.874Gbps, about 99.8% the capacity of the Avalanche test tool running back to back. Thus, in both tests, the 10200v moved traffic almost as fast as it was offered.

TAKING A PEEK AT SSL

With all the recent news about government wiretaps and corporate espionage, it's easy to assume that decrypting SSL traffic is automatically a bad thing. That assumption would be false.

Organizations have several good reasons for wanting to decrypt SSL traffic. Some industries have regulations that require traffic inspection. Others may want to obfuscate certain strings in traffic (for example, credit card or Social Security numbers). Others may simply want to break down application percentages, or troubleshoot server or network problems. Whatever the reason, there are legitimate reasons for organizations to terminate SSL connections; decrypt the traffic and pass it along to external devices for further analysis; and then re-encrypt it and send it on its way.

The problem, as past test results have shown, is that SSL decryption can introduce a big performance hit. In past tests, we've seen rates nosedive from tens of gigabits well down into the megabit range when decryption is enabled. Given the computationally intense nature of decryption and encryption, those concerns about performance only increase as traffic rates rise.

In the case of the F5 firewall, there is a performance cost to SSL decryption, but it's nowhere near as steep as we've seen in past tests. For example, the 10-kbyte Web object test ran at a tad over 17Gbps with SSL traffic; with decryption, that rate fell to 11.188Gbps. So, there's certainly a performance hit with SSL decryption, but it's hardly the nosedive into megabit territory we've seen in previous tests.

HOW HIGH?

Another key measure of firewall performance is scalability, which in turn has two dimensions: capacity and rate. We tested the F5 firewall both in terms of maximum concurrent TCP connections and maximum connection setup rate.

Connection capacity is important because a single user request can involve many TCP connections. For example, a single request for the home page of many news sites can involve 100 or more TCP connections due to web design trends, ad servers, streaming media servers, and other factors.

Connection rate matters because web sites may be hit with huge bursts of traffic. One common example is flash mobs, where some event (e.g., availability of a new product or concert tickets) causes a huge spike in connection request rates. Another common use case is disaster recovery, where the loss of one set of servers causes traffic to be migrated to a new set of servers.

In the capacity tests, we configured Spirent Avalanche never to request one web object per connection and then do nothing for the rest of the test. Since Avalanche doesn't age out TCP connections by default, we were able to build up progressively larger connection counts, well into the tens of millions.

F5 claims the BIG-IP 10000v supports 36 million concurrent connections. We validated that claim, sustaining 36,000,291 unique TCP connections for a 60-second period.

In the rate tests, we used HTTP 1.0 to ensure each new Web request would force a new TCP connection. Here again, F5 exceeded its rated capacity of 850,000 connections per second. In our tests, the BIG-IP sustained an average of 869,183 new connections per second for a 60-second period.

We did find a couple of fit and finish issues in the F5 firewall, both minor. The firewall failed to process a minuscule percentage of TCP connections – on the order of dozens to hundreds of failures out of millions to tens of millions of transactions. We'd configured Spirent Avalanche to abort any transaction taking more than 1 second, which is an eternity at 10G Ethernet rates. For a tiny number of attempts, TCP handshakes never completed. (All tests ran without errors between a pair of Avalanche C100 appliances.)

In an even smaller number of cases, the F5 firewall transmitted an extra TCP reset (RST) packet during connection shutdown. This is odd considering we'd configured Spirent Avalanche to close connections with TCP finished (FIN) and not RST flags. F5's explanation is that connection state between the firewall's client and server sides wasn't synchronized for a tiny number of connections, and in these cases the firewall sent a gratuitous RST packet. (Older versions of Windows – Windows XP and earlier, and Windows Server 2003 and earlier – tear down TCP connections with a RST rather than a FIN packet. This saves a little memory on the client, but it's a terrible idea for intermediate devices like firewalls, since they will continue to try to track connection state). Again, though, we consider both issues minor annoyances.

Protecting a data center's servers when rates climb into the dozens of gigabits is a significant challenge. With its high-speed rates, its high scalability, and its server protection features, F5's BIG-IP 10200v with the Advanced Firewall Manager (AFM) package is up to that challenge.

F5 Firewall Aces Performance Testing

Turning on firewall features can sometimes result in a significant performance hit, but in our performance testing, the F5 firewall came up aces, maxing out network capacity while also offering sophisticated filtering and attack protection capabilities. In some cases, traffic rates were higher with the firewall in place than without, probably because the F5 device managed server loading more efficiently. (Numbers below represent Gbit/s.)

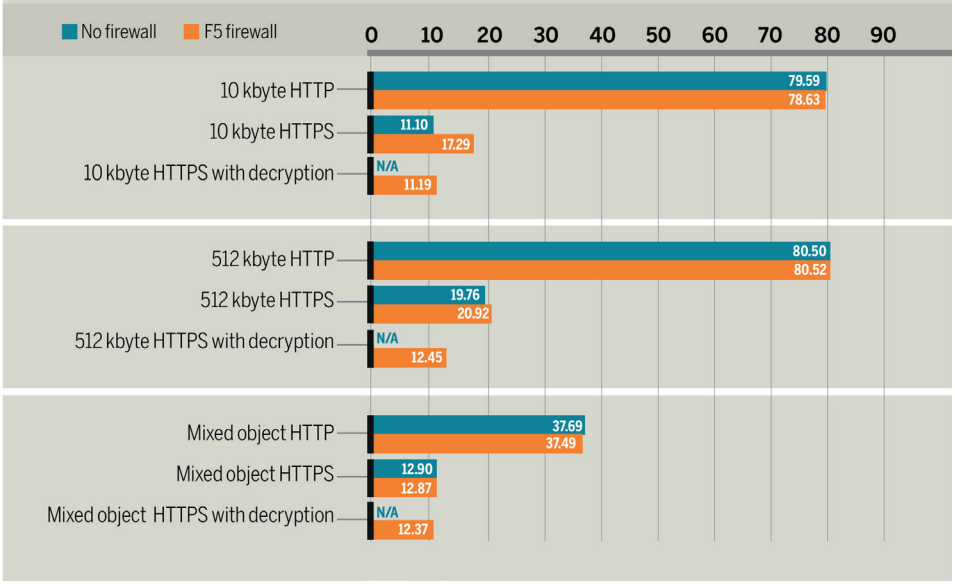

Thanks

Network World gratefully acknowledges the assistance of Spirent Communications, which supplied its Spirent Avalanche C100 traffic appliances. Spirent's Michelle Rhines and Jeff Brown also provided engineering support for this project.

Newman is a member of the Network World Lab Alliance and president of Network Test, an independent test lab and engineering services consultancy. He can be reached at dnewman@networktest.com.

How We Did It

We assessed performance using three sets of tests, covering forwarding rates with mixed HTTP content; rates with static HTTP content, and TCP connection behavior. Two pairs of Spirent Avalanche C100 traffic generator/analyzers, each equipped with eight 10G Ethernet interfaces, served as the primary test tool.

For the forwarding rate tests, we configured each of the F5 firewall's 16 10G Ethernet interfaces to act as a gateway for a different IP subnet. We also installed more than 500 access rules on each firewall. We configured Spirent Avalanche to emulate 2,048 clients and up to 80 servers, distributed across the 16 subnets.

In the mixed-content tests, we offered the same combination of HTTP object types and sizes as in previous Network World tests of next-generation firewall performance. Object types included text, images, and other binary content such as PDF files. Object sizes ranged from 1 kbyte to 1,536 kbytes, all requested over HTTP. We also reran the same tests using SSL with an RC4-MD5 cipher.

The static-content tests also used HTTP and SSL, but in this case involved separate tests with 10- and 512-kbyte text objects. For both mixed- and static-content tests, we averaged forwarding rates over a 60-second steady-state period with no failed requests.

To determine concurrent TCP connection count, we configured each new client emulated by Spirent Avalanche to request one object and then do nothing, building up progressively larger numbers of connections. The maximum concurrent connection count was determined to be the largest count at which the firewall serviced all requests with no failed requests.

To determine connection setup rate, we configured clients and servers emulated by Spirent Avalanche to use HTTP version 1.0, forcing the use of a new TCP connection for each HTTP request. Using a binary search, we determined the maximum rate at which the firewall could service requests for 60 seconds with no failed transactions.