



2010 Application Delivery Controller Performance Report

Contents

Letter of Introduction	3
Executive Summary	4
Introduction	5
Products tested	5
Test Results - Layer 4	6
Test Results – Layer 7	8
Analysis – L7 Performance	11
SSL Processing	11
Analysis – SSL Performance	14
HTTP Compression	14
Analysis – HTTP Compression	16
HTTP Caching and Compression	16
Analysis – HTTP Caching and Compression	17
Cookie Persistence	18
Analysis – Cookie Persistence	19
Power Efficiency	19
Analysis – Power Efficiency	20
Appendix A: Additional Testing Details	21
Load Generating Equipment	21
Test Specific Details	21
DUT to Switch Connections	22
Appendix B: Questions and Answers	23
Glossary	26

Letter of Introduction

The following performance report from F5 is a comparative and not a competitive report, which is an important distinction.

The guiding principles of a comparative report are that it must be accurate, transparent, and reproducible. It must provide as much factual information as possible so the customer can make an informed product decision. All test methodology and configuration information must also be freely available so anyone can validate the veracity of the data produced. A comparative report requires complete transparency in how results were derived so that, as in science, the results are reproducible.

On the other hand, the purpose and intent of a competitive report is for use as a sales tool, which does not assist in making an informed product decision. Reports of this nature tend to be the official-looking documents released by third-party test firms who have been paid by a specific vendor. Customers can check whether a report is comparative or competitive by finding out whether anyone has full access to the complete test methodology and the device configurations involved in the report. The methodology, if you have access to it, should also mimic real-world use scenarios and not artificial ones that inflate numbers. If none of these attributes are present, it is a competitive report and should be viewed with a great deal of skepticism.

Unfortunately discerning between the two approaches isn't always easy for customers. Third-party reports have the air of authenticity and impartiality and look like objective comparisons between vendors. Many customers do not know the vendor paying for the test often designs a methodology so their product "wins". The third-party test firm does not validate whether the tests are meaningful or artificial, they simply run the tests and validate the results. These reports have the appearance of being meaningful and relevant, but can be very misleading. Below are two simple and recent examples of this practice.

- One vendor recently released a report on Layer 7 performance versus F5. The numbers derived simply did not match our own performance benchmarking, which we go to great lengths to ensure is as accurate and factual as possible. As best we could discern, (the test methodology was not freely available) it seems the methodology used generated HTTP traffic, but was processed only at Layer 4; no Layer 7 processing was actually done. However, the claims were touted as Layer 7 sessions per second "implying" actions, like content switching, were being performed. The report omitted this important detail. The result of the test misleads customers because Layer 4 processing is easier than Layer 7 processing and yields higher results. Unfortunately, only after the customer buys the competitor's product would they become aware they would really get less than half the L7 performance doing real-world HTTP content switching, redirection, or persistence than the report would lead them to expect.
- Another recent report made some pretty outrageous SSL TPS claims between themselves and F5. The results made little sense to our performance experts. Like the previous example, the methodology was not freely available. After some intensive research, we were able to reproduce the "SSL" numbers reported by a third-party test firm paid to validate them. The problem was the test seemed to have been designed to artificially inflate SSL TPS. Rather than measure the resource-intensive work of setting up SSL connections, the test measured how many HTTP requests could be pipelined over a single SSL connection. The test results are factual for what was measured, but incorrectly reported as SSL TPS.

Examples like these are why such reports should never be viewed as a truly comparative evaluation but as a competitive sales tool. It is also the primary motivation for the following comparative report. F5 stands behind all results in the following pages. We conducted these tests with our own performance experts and we intentionally did not pay for or hire an "independent" third-party test organization to hide behind. However, we know we are not perfect and make mistakes at times. So, if in some way our tests are flawed, do not mimic real-world scenarios, are not reproducible, or if the product configurations are not optimized appropriately, let us know and we will correct our mistakes and update this report. Unlike others, we want our results to be open, honest and repeatable.

Sincerely,

Karl Triebes
SVP and CTO, F5 Networks



Executive Summary

The 2010 F5 Performance Report documents the performance of Application Delivery Controllers from the three top Application Delivery Controllers vendors: F5, Cisco® Systems, and Citrix® (based on market share). The top-end devices from each vendor, along with two mid-range devices, were tested for this report.

The market for Application Delivery Controllers (ADCs) is very competitive, with nearly every vendor claiming a performance advantage in one scenario or another. Unfortunately, the claims from each vendor rest on differing definitions of performance criteria. Each vendor has their own definitions of various terms (such as *Layer 7* and *connection*), preferred configuration settings for the different devices, and the presentation of results. These factors significantly reduce the value of typical vendor performance claims. With the inconsistent definitions between vendors, especially in the context of published data sheets, performance metrics cannot be fairly compared between vendors.

Many vendors publish reports from performance tests that they have hired third parties to produce. The configuration files for the devices tested and the testing equipment are rarely made available. It is difficult to determine the fairness of the tests or their applicability without this information.

In 2007, F5 introduced the industry's most transparent and robust [performance testing guide](#) into the public domain. With this publicly available guide, customers were given the framework for evaluating the performance of multiple ADC products with consistent definitions and evaluation criteria. Also in 2007, F5 published a Performance Report using this testing methodology. This 2010 Performance Report also follows those guidelines.

For this reason, F5 has published the configuration files for each device included in this report as well as the configuration files for the testing equipment. This allows anyone with the equivalent testing gear to reproduce these tests independently. The configuration files are available on F5's DevCentral web site:

http://devcentral.f5.com/downloads/perf/PerformanceReport2010_configs.zip.

The devices tested for this report span a broad range of performance capabilities and price. In some cases it is not appropriate to directly compare two devices because of these differences. It is left to the reader to evaluate the results in this report relative to each vendors published performance specifications and product pricing (see Questions 1 and 2 in "Appendix B: Questions and Answers" on page 23 for more information).

Still, there are several observations that can be made after reviewing these test results.

- The F5 VIPRION, a chassis based platform, scales near linearly as blades are added to the chassis. Each blade adds both processing capacity and network interfaces. This near linear scaling demonstrates the strength and value of F5's Clustered MultiProcessing (CMP) technology.
- Citrix's claim that their nCore firmware increases the performance of NetScaler® 4x - 7x over the Classic firmware is not demonstrated in any of the test results. In fact, the improvement is usually less than double.
- Devices are limited by their ability to process requests or by their internal bandwidth. It is generally easy to determine from these test results when a device is being limited by one or the other. When a device is bandwidth limited, it usually has available processing capacity that could be used for more complex traffic management.
- The F5 devices support much higher connection and request rates relative to their maximum throughput when compared to the other products tested.

In conjunction with this report, F5 is also publishing a Functionality Report that compares the features and capabilities of the products included in this report. The information in both reports is based on the current software versions as of March 15th, 2010.


Application Delivery Controller

2010 PERFORMANCE REPORT

Introduction

This document contains the results of testing conducted according to the previously published [performance testing methodologies guide](#). To ensure complete transparency, F5 has also [published the configurations](#) for all devices involved in the testing, including the load generation equipment. A glossary of terminology used in this document is available in “Appendix A: Additional Testing Details” on page 21.

The following diagram illustrates a high-level view of the environment in which the devices were tested.



Products tested

The following table lists the products tested for this report and the vendor’s published performance specifications for each device. The Cisco and Citrix products tested in this report were of new manufacture and purchased through regular retail channels.

In the following graphs and tables of results, the labels “VIPRION x 1” and “VIPRION x 2” indicate the number of blades (model PB200) installed in the VIPRION chassis.

Vendor Published Performance Specifications

Vendor	Device	Layer 4 Connections Per Second	Layer 4 Throughput (Gbps)	Layer 7 Requests Per Second	SSL TPS (RC4)	SSL Bulk Throughput (Gbps)	Compression (Gbps)
F5	VIPRION (PB200) x 2	1,400,000	36.0	3,200,000	100,000	18.0	24.0
F5	VIPRION (PB200) x 1	700,000	18.0	1,600,000	50,000	9.0	12.0
F5	BIG-IP 8900	400,000	12.0	1,200,000	58,000	9.6	8.0
F5	BIG-IP 3900	175,000	4.0	400,000	15,000	2.4	3.8
Cisco	ACE20	348,000	16.0	¹	20,000	3.3	N/A
Cisco	ACE4710	120,000	4.0	¹	7,500	1.0	2.0
Citrix	NetScaler MPX-17000 nCore	¹	18.0	1,500,000	80,000	6.5	6.0
Citrix	NetScaler MPX-17000 Classic	¹	15.0	340,000	48,000	6.0	6.0

Note on throughput measurements

F5 typically reports throughput that counts all the data sent across the interfaces. Stated another way, our standard method counts all the bits in every Ethernet frame. This is the industry standard method for measuring throughput and is referred to as the *Layer 2 throughput*.

¹ We were unable to find published numbers for these categories

Application Delivery Controller

2010 PERFORMANCE REPORT

This report is an exception to that practice. Here we are reporting *Layer 7 throughput*. This is due to a limitation of the load generating equipment used during these tests, which does not report Layer 2 throughput. For more information, see Questions 3 and 4 in "Appendix B: Questions and Answers" on page 23.

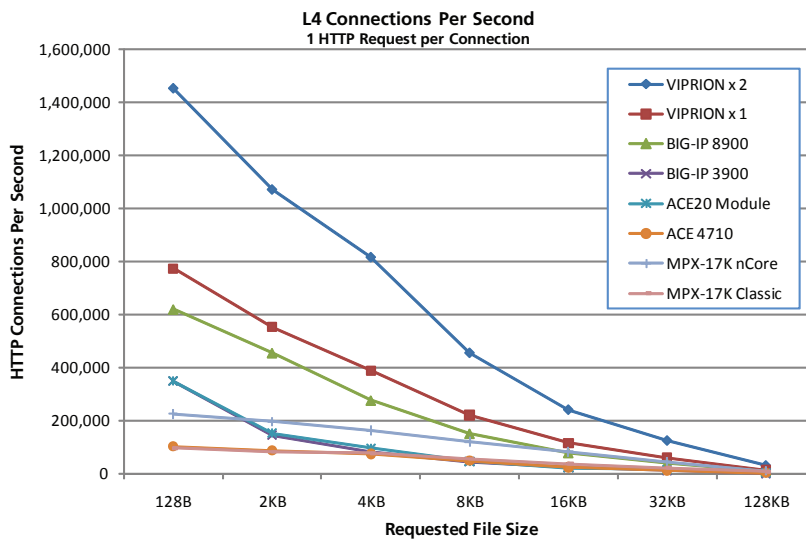
Test Results - Layer 4

Layer 4 (L4) performance is a measure of basic TCP/IP load balancing, a baseline configuration with the minimum set of features enabled. L4 performance is most relevant for applications that deal with a lot of bulk data, where little application awareness is required. Load balancing FTP servers and some types of streaming media are common scenarios for L4 load balancing.

All devices tested have configuration options for a L4-only (or TCP only) mode, shown for comparison in the following results. L4 tests often show the highest connections per second and/or throughput results that are possible for a given Application Delivery Controller. This makes L4 testing appropriate for use in baseline capacity planning; as it is unlikely that performance under more complex scenarios (i.e. with additional features enabled) will be higher than the baseline L4 results.

There are two Layer 4 tests shown in the following graphs. The first test has 1 HTTP request per TCP connection. Even though the ADC is not processing the HTTP traffic, the purpose of the HTTP request is to create a complete interaction between the client and the ADC. This includes setting up the TCP connection from the client to the ADC, the client requesting a file, the ADC returning the file to the client from the server and the TCP connection being closed.

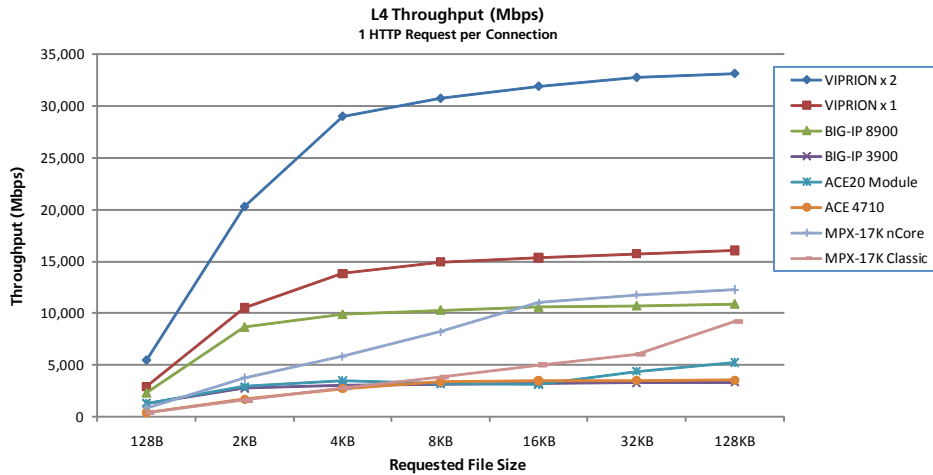
This tests the ability of the ADC to perform the computationally-intensive TCP connection process.




Device	128B	2KB	4KB	8KB	16KB	32KB	128KB
Viprion x 2	1,453,032	1,071,709	816,798	456,442	242,425	126,139	32,168
Viprion x 1	773,979	554,124	389,204	221,729	116,422	60,549	15,592
BIG-IP 8900	621,082	456,545	277,821	152,675	80,094	41,186	10,555
BIG-IP 3900	349,642	146,695	85,308	46,706	24,348	12,572	3,234
ACE20 Module	350,000	154,535	97,375	47,330	23,608	16,857	5,067
ACE 4710	103,657	88,308	76,183	50,104	26,181	13,504	3,445
MPX-17k nCore	226,170	197,214	164,022	121,773	83,457	45,358	11,913
MPX-17k Classic	97,927	83,439	78,433	56,381	37,960	23,280	8,938

Application Delivery Controller

2010 PERFORMANCE REPORT



Device	128B	2KB	4KB	8KB	16KB	32KB	128KB
Viprion x 2	5,453	20,297	29,020	30,747	31,919	32,776	33,154
Viprion x 1	2,904	10,502	13,825	14,938	15,363	15,739	16,057
BIG-IP 8900	2,329	8,650	9,873	10,286	10,572	10,697	10,871
BIG-IP 3900	1,310	2,779	3,035	3,145	3,212	3,264	3,326
ACE20 Module	1,312	2,926	3,460	3,189	3,116	4,379	5,217
ACE 4710	388	1,676	2,707	3,374	3,459	3,507	3,544
MPX-17k nCore	848	3,736	5,829	8,203	11,014	11,785	12,272
MPX-17k Classic	366	1,583	2,787	3,797	5,005	6,045	9,202



Device	128B	2KB	4KB	8KB	16KB	32KB	128KB
Viprion x 2	21,756	33,668	34,041	34,096	33,477	33,544	33,581
Viprion x 1	10,955	16,867	17,008	17,096	17,029	17,018	17,088
BIG-IP 8900	8,157	10,757	10,891	10,933	10,919	10,906	10,967
BIG-IP 3900	2,250	3,247	3,338	3,363	3,416	3,430	3,377
ACE20 Module	2,625	11,727	12,120	12,634	12,590	12,681	12,915
ACE 4710	2,959	3,405	3,537	3,536	3,544	3,561	3,575
MPX-17k nCore	4,004	12,848	12,591	13,697	13,654	13,607	13,721
MPX-17k Classic	1,533	4,364	4,824	5,186	5,169	5,353	5,346

Application Delivery Controller

2010 PERFORMANCE REPORT

Analysis – L4 Performance

The L4 performance tests can demonstrate the difference between a device's ability to setup TCP connections and its available bandwidth.

The ACE20 module shows clear signs of being limited by its ability to setup TCP connections. At small file sizes, its performance is matched or exceeded by the BIG-IP 3900. When the tests get to the larger file sizes where not as many setups are needed, the ACE pulls ahead of the 3900 in throughput.

This pattern is repeated with the NetScaler MPX-17000 running the nCore firmware when compared to the BIG-IP 8900 (or even the 3900). The BIG-IP 3900 sets up over one and a half times more TCP connections than the MPX-17000 nCore at the smallest file size and the BIG-IP 8900 can setup three times as many connections as the MPX-17000 nCore. In contrast, the throughput of the MPX-17000 nCore is 20% more than the BIG-IP 8900 and four times that of the BIG-IP 3900.

These two tests also demonstrate how the VIPRION scales linearly as blades are added to the chassis. At the larger file sizes, two blades actually performed slightly more than twice what one blade did. Also of note is that although both a single VIPRION blade and the MPX-17000 running nCore have a specified throughput of 18Gbps, the single VIPRION blade actually delivers 20% more than the MPX-17000.

The performance of the NetScaler device when running the Classic firmware is generally less than half what it is when running the nCore firmware. This performance difference repeats in all the tests and is greater in many of them.

The test with infinite requests is effective at finding each device's maximum throughput and we can clearly see that in these results.

Test Results – Layer 7


Layer 7 (L7) performance tests measure basic HTTP-aware load balancing; every HTTP request is inspected and then directed to the appropriate group of servers. This technology is commonly used to allow servers to host more specialized content. For example, one group of servers may store small images, another group might store large zip files, and a third could handle requests for dynamic web pages. This test performs a simple inspection of the HTTP URI to identify requests for images and direct them to a different group of servers.

L7 performance is relevant to most applications being deployed today – it is the performance metric most often referenced when comparing Application Delivery Controllers. The most important difference between a L4 test and a L7 test is that the Device Under Test (DUT) must inspect the application-layer data transferred between the clients and servers. Because every client request must be inspected, an increase in requests sent by the clients means additional stress placed on the DUT. Additionally, HTTP request multiplexing (connection reuse) is enabled during these tests to provide server offload and ensure that all HTTP requests in a given connection are inspected.

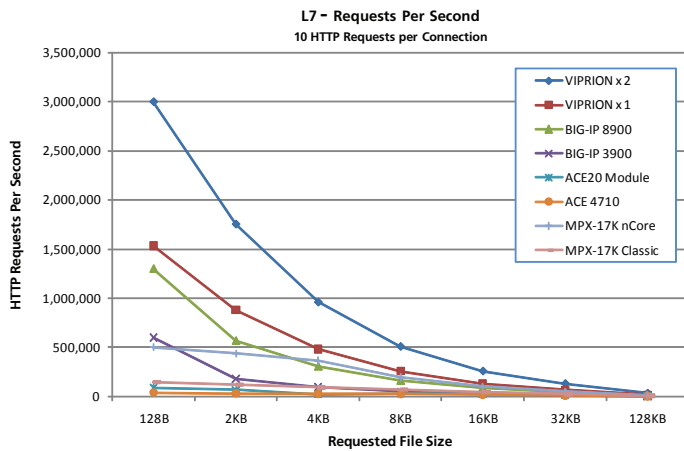
The following results include two very similar tests, with each test varying the number of HTTP requests per TCP connection. The slight difference in the tests demonstrates the effect of TCP/IP processing versus HTTP processing on the DUT. With one request per connection, there is an equal amount of TCP/IP processing and HTTP processing per request (a single TCP/IP connection is setup, a single request is sent, response received, and TCP/IP connection teardown). Adding additional HTTP requests per connection requires less TCP/IP processing, placing more stress on the L7 inspection capabilities of the DUT. Different applications have different needs with respect to requests per connection, but it is important to know that modern web browsers attempt to send as many requests per connection as possible.

Application Delivery Controller

2010 PERFORMANCE REPORT



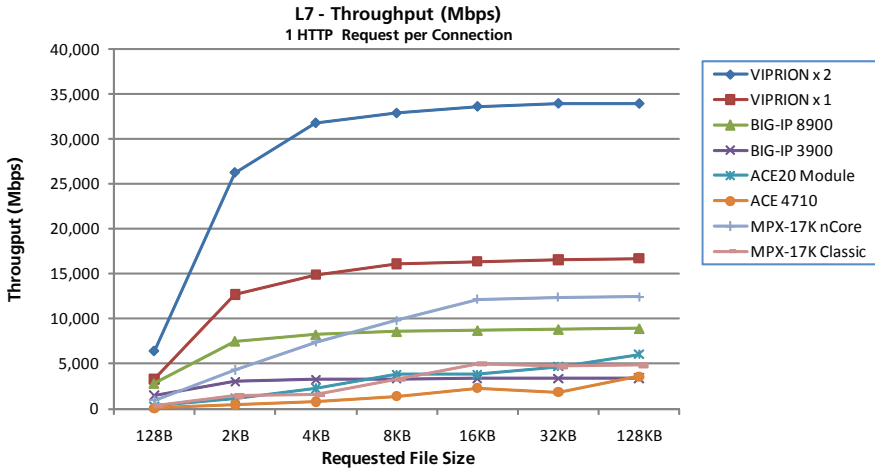
Device	128B	2KB	4KB	8KB	16KB	32KB	128KB
Viprion x 2	1,706,265	1,383,005	894,145	488,051	254,043	130,597	32,903
Viprion x 1	882,229	670,999	419,649	239,387	124,193	63,726	16,241
BIG-IP 8900	749,314	397,096	231,877	127,562	66,058	34,024	8,673
BIG-IP 3900	389,178	160,920	91,495	49,353	25,540	13,050	3,318
ACE20 Module	78,857	62,000	63,802	56,714	28,872	18,111	5,887
ACE 4710	26,444	23,735	22,700	20,833	17,592	7,160	3,490
MPX-17k nCore	253,954	228,203	209,065	146,161	92,445	47,802	12,089
MPX-17k Classic	89,444	75,971	45,209	48,509	37,545	18,185	4,701




Device	128B	2KB	4KB	8KB	16KB	32KB	128KB
Viprion x 2	3,000,043	1,753,923	960,049	508,591	258,206	130,724	32,738
Viprion x 1	1,535,425	879,202	482,520	256,616	130,387	66,516	16,833
BIG-IP 8900	1,299,486	568,903	308,154	163,600	83,340	42,469	10,749
BIG-IP 3900	600,001	178,383	96,234	50,636	25,799	13,183	3,337
ACE20 Module	86,222	66,329	21,168	29,338	27,775	17,306	6,124
ACE 4710	39,157	31,771	28,406	25,416	20,592	9,211	3,468
MPX-17k nCore	500,728	438,718	363,442	195,397	101,647	49,573	12,152
MPX-17k Classic	144,091	119,422	91,294	70,314	42,621	24,405	6,665

Application Delivery Controller

2010 PERFORMANCE REPORT



Device	128B	2KB	4KB	8KB	16KB	32KB	128KB
Viprion x 2	6,414	26,260	31,778	32,866	33,607	33,943	33,928
Viprion x 1	3,316	12,720	14,917	16,133	16,388	16,565	16,735
BIG-IP 8900	2,816	7,529	8,243	8,595	8,720	8,842	8,925
BIG-IP 3900	1,465	3,050	3,251	3,326	3,369	3,388	3,411
ACE20 Module	295	1,174	2,267	3,820	3,815	4,702	6,062
ACE 4710	98	448	806	1,403	2,324	1,865	3,589
MPX-17k nCore	954	4,326	7,431	9,849	12,204	12,421	12,449
MPX-17k Classic	335	1,439	1,609	3,269	4,969	4,722	4,837



Device	128B	2KB	4KB	8KB	16KB	32KB	128KB
Viprion x 2	11,284	33,138	34,107	34,292	33,938	33,969	33,732
Viprion x 1	5,772	16,672	17,155	17,313	17,226	17,277	17,336
BIG-IP 8900	4,887	10,784	10,949	11,022	11,016	11,054	11,068
BIG-IP 3900	2,255	3,381	3,421	3,412	3,407	3,423	3,434
ACE20 Module	323	1,256	753	1,976	3,664	4,492	6,307
ACE 4710	147	602	1,009	1,712	2,716	2,388	3,568
MPX-17k nCore	1,882	8,316	12,918	13,163	13,416	12,879	12,517
MPX-17k Classic	541	2,263	3,244	4,736	5,624	6,339	6,862

Analysis – L7 Performance

Some devices demonstrate an imbalance between their ability to process requests and their maximum throughput. The most extreme example is the ACE20 module, which is limited to 86,000 HTTP Requests Per Second (RPS), resulting in only 323Mbps of throughput with 128 byte files. Its maximum L7 throughput of 6.3Gbps is less than half of what it was able to achieve in the L4 throughput test.

All of the F5 devices deliver high rates of connection and request processing. Even the BIG-IP 3900 outperforms all of the other vendor's products when tested with small files.

The BIG-IP 3900 and the ACE 4710 both have a specified throughput of 4Gbps and they are priced similarly. Both devices demonstrate they are 4Gbps platforms, but the ACE 4710 can only achieve this at the largest file size. The BIG-IP 3900's maximum RPS is 14 to 15 times that of the ACE 4710. This enables the 3900 to deliver higher throughput than the 4710 at almost all file sizes tested.

The scalability of the VIPRION platform is demonstrated again in these L7 tests. And the single VIPRION blade processes three times as many requests than the MPX-17000 nCore and has 25% higher throughput.

Both the BIG-IP 8900 and a single VIPRION blade outperform the ACE20 module in both RPS and throughput at all file sizes.

SSL Processing

SSL is used around the world to secure communications between users and applications. SSL is a standard encryption protocol available in every major operating system, web browser, smart phone, and so on. SSL technology helps make online shopping secure, enables secure remote access (SSL VPN) and much more – SSL is ubiquitous in commercial and consumer networking security solutions. SSL provides security using a combination of public key cryptography (typically RSA®), and symmetric encryption (commonly RC4, 3DES, or AES). Both RSA and the various symmetric encryption algorithms are computationally-intensive, and require specialized hardware to achieve acceptable performance or large scale in the nearly all commercial uses of SSL (such as web based email, online stores, secure logins, online banking web sites, and SSL VPNs).

SSL Transactions Per Second (TPS) performance is a measure of encryption offload capability. For small response sizes, this primarily measures the RSA handshake operations that occur at the start of every new SSL session. This RSA operation is computationally-intensive; all major SSL offload vendors use specialized hardware to accelerate this task. For larger responses, the computational cost of the RSA operation is less relevant. Because the RSA operation only occurs once at the beginning of a session, the true comparison of performance is the throughput of encrypted traffic, also known as symmetric encryption or *bulk crypto*. Bulk crypto is a measure of the amount of data that can be encrypted in a given second. If a vendor has SSL offload hardware that can process bulk crypto, it is apparent in tests with large response sizes.


Tests were conducted across a range of file sizes to demonstrate the performance of both public key operations (small files) and bulk crypto (large files).

Tests were run with using the RC4-MD5 and AES256-SHA ciphers. RC4-MD5 is an older cipher and generally less secure than newer ciphers such as the AES based ones. RC4-MD5 is one of the most frequently used ciphers for SSL connections but companies are increasingly setting their default cipher to one based on AES.

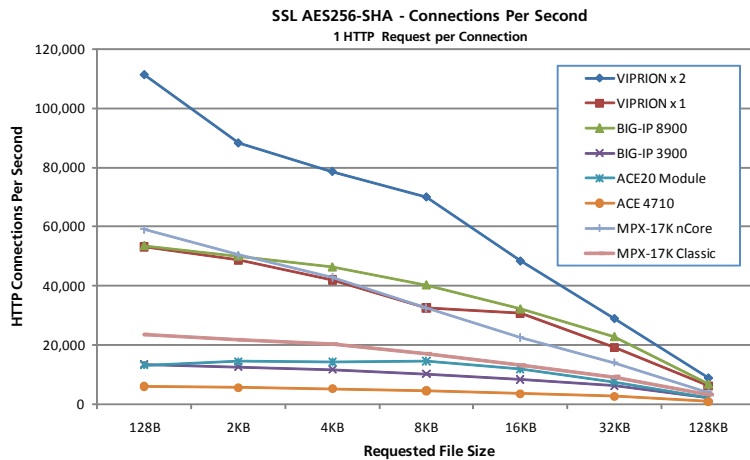
The test with one HTTP Request per connection measures SSL TPS while the infinite requests per connection test measures the devices bulk SSL throughput capability.

Application Delivery Controller

2010 PERFORMANCE REPORT




Device	128B	2KB	4KB	8KB	16KB	32KB	128KB
Viprion x 2	119,658	109,297	99,368	83,003	62,811	43,194	14,466
Viprion x 1	59,789	54,564	49,538	41,052	31,360	21,610	7,225
BIG-IP 8900	59,592	54,369	49,322	41,413	31,545	21,509	6,927
BIG-IP 3900	15,002	13,733	12,540	10,577	8,223	5,776	2,031
ACE20 Module	13,971	14,122	14,045	14,377	11,488	7,879	2,663
ACE 4710	6,026	5,583	5,027	4,173	3,219	2,283	832
MPX-17k nCore	76,161	72,600	66,145	52,464	35,949	20,750	5,542
MPX-17k Classic	24,354	22,835	20,773	17,275	13,339	9,391	3,285



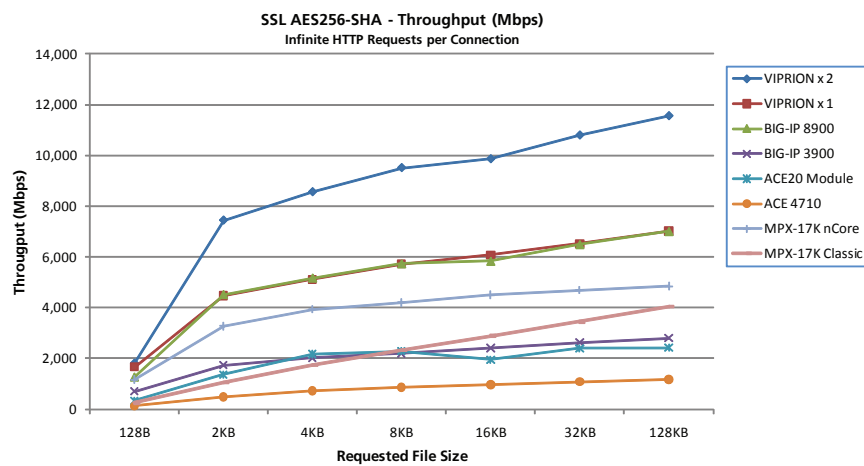
Device	128B	2KB	4KB	8KB	16KB	32KB	128KB
Viprion x 2	111,221	88,207	78,602	69,996	48,390	28,999	8,962
Viprion x 1	53,249	48,831	42,053	32,614	30,687	19,194	6,224
BIG-IP 8900	53,469	49,960	46,303	40,267	32,169	22,863	6,881
BIG-IP 3900	13,507	12,647	11,783	10,335	8,392	6,220	2,344
ACE20 Module	13,272	14,510	14,275	14,628	12,011	7,679	2,199
ACE 4710	6,067	5,691	5,330	4,582	3,652	2,706	998
MPX-17k nCore	59,133	50,466	42,934	32,561	22,583	14,081	4,141
MPX-17k Classic	23,679	21,746	20,332	17,194	13,177	9,267	3,247

Application Delivery Controller

2010 PERFORMANCE REPORT



Device	128B	2KB	4KB	8KB	16KB	32KB	128KB
Viprion x 2	3,504	10,678	12,070	12,589	13,564	14,913	15,945
Viprion x 1	1,769	5,343	5,985	6,298	6,777	7,450	7,955
BIG-IP 8900	1,336	4,922	6,216	6,382	6,779	7,250	7,595
BIG-IP 3900	756	1,756	1,933	1,982	2,111	2,261	2,355
ACE20 Module	334	1,353	2,560	2,293	2,062	2,300	2,855
ACE 4710	133	486	725	860	962	1,076	1,172
MPX-17k nCore	1,632	4,960	5,735	5,870	6,076	6,158	6,117
MPX-17k Classic	287	1,206	1,917	2,466	3,020	3,471	3,976



Device	128B	2KB	4KB	8KB	16KB	32KB	128KB
Viprion x 2	1,810	7,442	8,559	9,508	9,856	10,790	11,551
Viprion x 1	1,665	4,474	5,116	5,727	6,082	6,526	7,014
BIG-IP 8900	1,264	4,524	5,176	5,737	5,835	6,493	7,010
BIG-IP 3900	702	1,736	2,024	2,206	2,410	2,630	2,796
ACE20 Module	335	1,369	2,166	2,283	1,953	2,396	2,417
ACE 4710	133	486	725	860	962	1,076	1,172
MPX-17k nCore	1,167	3,271	3,931	4,198	4,502	4,693	4,848
MPX-17k Classic	249	1,054	1,744	2,316	2,886	3,442	4,040

Application Delivery Controller

2010 PERFORMANCE REPORT

Analysis – SSL Performance

As in other tests, the linear scalability of the VIPRION platform is clearly demonstrated, with the results for two blades being almost exactly twice that of the single blade results.

Only the F5 devices matched or exceeded their specified SSL TPS using the RC4-MD5 cipher. Excluding the VIPRION with two blades installed, the MPX-17000 nCore had the highest TPS at 76,000.

AES256-SHA is a more computationally-intensive encryption cipher than RC4-MD5 and that can be seen in the lower TPS rates of the AES tests compared to the RC4 tests. The degree of difference varies widely depending on the device.

The ACE 4710 actually performed the same on both tests, while the ACE20 performed 5% less on the AES test than on the RC4 test. All of the F5 devices showed a maximum TPS 10% lower on the AES tests than on the RC4 tests. The NetScaler MPX-17000 running nCore experienced the largest performance degradation with maximum TPS for AES being 22% less than for RC4.

In the RC4-MD5 throughput test, the MPX-17000 nCore fell behind the BIG-IP 8900 starting at the 4KB file size. This is despite the fact that its maximum TPS is 20,000 more. In the AES256-SHA throughput test, the MPX-17000 nCore never matches the BIG-IP 8900.

The maximum performance of the 2-blade VIPRION was not reached in these tests because it exceeded the capacity of our testing equipment by 1,000 to 2,000 transactions per second.

HTTP Compression

HTTP Compression performance is a measure of the standard compression algorithms supported in all modern web browsers. In situations where the bandwidth between clients and servers is limited, compression can provide significant performance benefits to end users. Compression can also help companies to achieve cost savings by reducing the bandwidth required to serve web-based applications.


The benefits of compression are widely understood, but compression is not universally used because it's very computationally intensive for servers. As a result, it is increasingly common to offload HTTP compression functionality to Application Delivery Controllers.

The most important metric when measuring compression is throughput. More data sent from the servers directly correlates with more compression work for the DUT.


In this test, 10 HTTP requests are sent per TCP connection to increase potential throughput. The ACE20 Module is not included because it does not support compression.

Application Delivery Controller

2010 PERFORMANCE REPORT



Device	128B	2KB	4KB	8KB	16KB	32KB	128KB
Viprion x 2	240,570	135,549	114,024	60,913	49,352	26,716	10,244
Viprion x 1	124,392	68,364	57,638	30,938	24,827	13,219	5,094
BIG-IP 8900	221,310	211,983	150,313	86,662	57,623	29,164	8,280
BIG-IP 3900	47,561	31,034	28,081	14,085	11,440	5,983	2,274
ACE 4710	31,618	25,239	23,966	19,979	15,548	8,919	2,274
MPX-17k nCore		108,767	58,936	38,959	27,018	13,272	4,032
MPX-17k Classic		53,175	47,384	34,660	25,784	16,621	6,018



Device	128B	2KB	4KB	8KB	16KB	32KB	128KB
Viprion x 2	903	2,558	4,049	4,103	6,510	6,940	10,546
Viprion x 1	467	1,296	2,045	2,087	3,276	3,435	5,234
BIG-IP 8900	831	4,017	5,344	5,837	7,610	7,576	8,520
BIG-IP 3900	178	588	996	948	1,509	1,553	2,340
ACE 4710	118	477	853	1,345	2,051	2,317	2,338
MPX-17k nCore		2,061	2,095	2,625	3,565	3,446	4,144
MPX-17k Classic		1,008	1,687	2,337	3,403	4,322	6,199

Application Delivery Controller

2010 PERFORMANCE REPORT

Analysis – HTTP Compression

What stands out in the compression test is the performance of the BIG-IP 8900. Except for the 2-blade VIPRION at 128B and 128KB file sizes, the 8900's throughput exceeds all of the other devices. The primary reason for this is that it contains dedicated compression hardware.

The results for the MPX-17000, (both Classic and nCore) do not include numbers for the 128B file size. This is due to a incompatibility between the testing equipment and the way the NetScaler devices write the Content-Length HTTP header. HTTP Headers have a format of "Header-Name: Value". When the NetScaler compresses files and the response is not chunked, it places multiple spaces between the header name and its value. The resulting header looks like this "Content-Length: 321". While the HTTP standard allows any amount of white space between the header name and its value, it recommends against it as almost all web servers and applications place a single space between the colon and the header value (see sections 2.2 and 4.2 of RFC 2616, the HTTP standard).

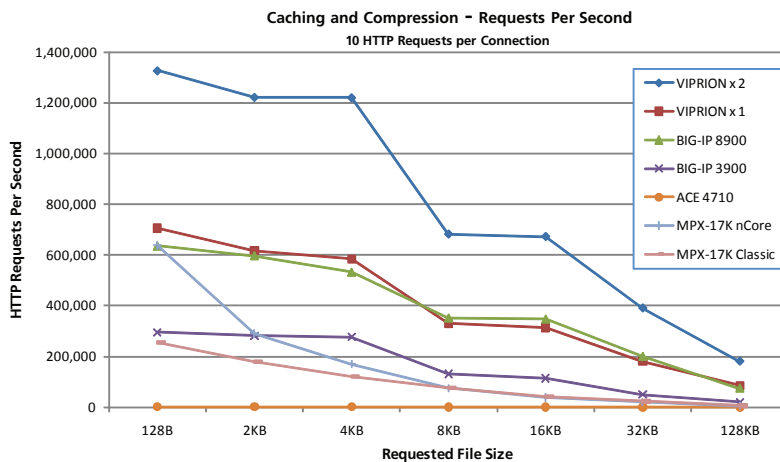
Because this extra spacing is so uncommon and not recommended, the Ixia equipment interpreted these as bad headers. This illustrates the potential interoperability problems applications may have with the behavior of the NetScaler devices.

Although not at the top, the MPX-17000 Classic had a stronger showing in the compression tests relative to the other devices than it did in all the other tests.

HTTP Caching and Compression

HTTP Caching performance is a measure of how quickly the DUT can serve HTTP responses from its own internal HTTP cache. HTTP Caching helps reduce server load by handling the majority of requests for static content, allowing the servers to focus resources on the business-specific aspect of the application. The use of HTTP Caching can also improve end user performance, because the servers have more free resources to process user requests.

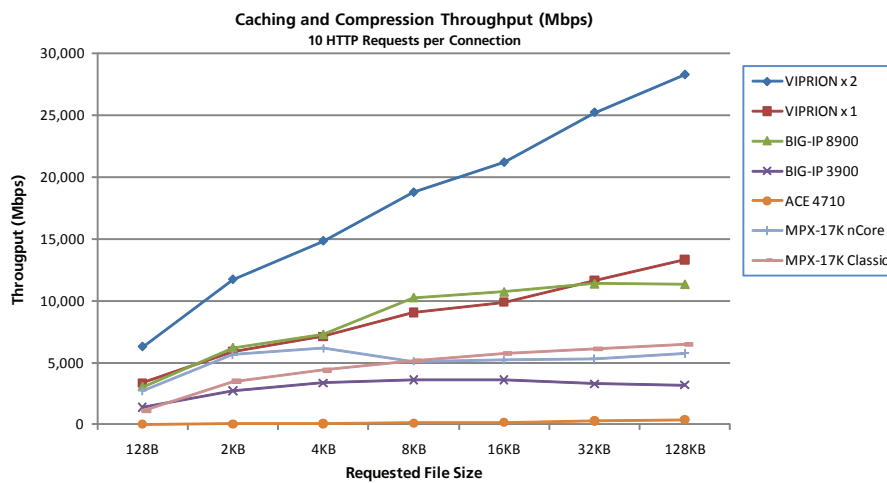
Another valuable but less obvious aspect of HTTP Caching is the fact that the devices providing the caching functionality incur lower load increases for serving cached objects compared to requesting content from the backend servers. By serving content from its own cache, a caching device avoids the round-trip-time required when requesting content from the servers, thus improving response time. HTTP Caching, when coupled with HTTP Compression can further lower resource requirements by reducing the number of times the same content has to be compressed.



Application Delivery Controller

2010 PERFORMANCE REPORT

Device	128B	2KB	4KB	8KB	16KB	32KB	128KB
Viprion x 2	1,326,332	1,221,523	1,220,099	682,430	672,955	390,872	180,991
Viprion x 1	706,029	616,926	584,265	329,339	314,193	180,576	85,392
BIG-IP 8900	635,427	597,431	532,465	351,395	347,236	199,936	74,310
BIG-IP 3900	295,829	282,281	276,852	131,287	114,380	50,790	20,346
ACE 4710	1,902	1,673	1,867	1,530	1,219	1,073	388
MPX-17k nCore	637,406	289,944	170,989	75,021	39,282	20,426	5,579
MPX-17k Classic	254,451	178,756	121,612	75,915	43,287	23,427	6,306



Device	128B	2KB	4KB	8KB	16KB	32KB	128KB
Viprion x 2	6,291	11,726	14,835	18,763	21,194	25,224	28,276
Viprion x 1	3,349	5,921	7,103	9,054	9,894	11,652	13,339
BIG-IP 8900	3,049	6,202	7,279	10,268	10,725	11,379	11,342
BIG-IP 3900	1,402	2,708	3,366	3,608	3,601	3,276	3,177
ACE 4710	8	32	66	104	160	278	396
MPX-17k nCore	2,727	5,647	6,167	5,093	5,205	5,318	5,749
MPX-17k Classic	1,088	3,481	4,386	5,155	5,735	6,098	6,497

Analysis – HTTP Caching and Compression

The ACE 4710 performed extremely poorly in this test, but this is consistent with Cisco’s documentation, which states: “Application acceleration performance on the ACE is 50 to 100 Mbps throughput. With typical page sizes and browser usage patterns, this equates to roughly 1,000 concurrent connections.”¹

Because this test does not use all of the ACE 4710’s acceleration techniques, it performed better than the documentation indicates it would, however it still only achieved 396 Mbps.

In comparison, the BIG-IP 3900’s performance was between 8 and 155 times greater than the performance of the 4710 in this test.


¹ Ace 4710 version A3(2.2) Device Manager GUI Configuration Guide, page 11-1

Cookie Persistence

In many applications it is important that requests from any one user are always sent to the same server. This is called persistence. Persistence is frequently required if the application maintains state for each user. ADCs use a number of techniques to implement persistence, with cookie persistence being one of the most commonly used methods.

With Cookie persistence, the ADC inserts a cookie in the HTTP header sent to a client. The client returns this cookie with all subsequent requests. The cookie sent to each client is unique and the ADC maintains a table that uniquely identifies which server should be used to handle requests for that client.


This test measures the ADCs performance when required to do the extra work of creating, inserting, tracking, and then parsing received cookies. The key metric is Requests Per Second.



Device	128B	2KB	4KB	8KB	16KB	32KB	128KB
Viprion x 2	2,266,901	1,628,105	955,242	502,768	256,241	130,070	32,761
Viprion x 1	1,160,740	770,273	465,940	246,480	126,923	66,010	16,770
BIG-IP 8900	968,203	413,226	211,689	123,797	69,353	34,753	7,889
BIG-IP 3900	500,812	168,849	93,021	49,633	25,589	13,066	3,332
ACE20 Module	80,013	55,594	58,545	25,839	29,212	18,270	6,078
ACE 4710	37,489	30,374	27,378	24,388	20,126	10,068	3,462
MPX-17k nCore	407,607	285,824	259,279	168,293	93,954	48,843	13,017
MPX-17k Classic	100,672	71,766	69,331	49,754	32,261	19,931	6,133

Application Delivery Controller

2010 PERFORMANCE REPORT



Device	128B	2KB	4KB	8KB	16KB	32KB	128KB
Viprion x 2	8,526	30,872	34,023	33,879	33,817	33,826	33,744
Viprion x 1	4,366	14,598	16,552	16,628	16,735	17,168	17,296
BIG-IP 8900	3,639	7,828	7,524	8,341	9,148	9,027	8,122
BIG-IP 3900	1,882	3,200	3,304	3,348	3,378	3,397	3,430
ACE20 Module	300	1,053	2,080	1,739	3,854	4,745	6,257
ACE 4710	140	575	972	1,642	2,655	2,620	3,564
MPX-17k nCore	1,531	5,417	9,219	11,339	12,398	12,690	13,426
MPX-17k Classic	378	1,360	2,463	3,353	4,257	5,177	6,314

Analysis – Cookie Persistence

The process of inserting cookies in the HTTP header demonstrates the limited L7 processing capacity of the ACE20 and ACE 4710. All the other devices outperform them at smaller file sizes.

As should be expected, all devices had lower performance in this test than in the L7 10 Requests per Connection test. The differences in performance varied not only from device to device but even from one file size to another.

Power Efficiency

The amount of electricity used by data center equipment is being scrutinized more than ever due to the increasing costs, both direct and indirect, of every watt consumed. A simple way to evaluate the energy efficiency of products is to calculate the work units produced per watt consumed.

Several ADC vendors have published energy efficiency numbers using the formula “HTTP RPS/watt.” F5 also believes this is an appropriate method to measure the energy efficiency of Application Delivery Controls.

The power draw of each device in this report was measured under two conditions. The first measurement was with the device configured and connected as it was for all tests but with no tests running. This is the minimum electricity the device will consume when powered on and connected to the network. The second measurement was taken while the device was under full load during the L7 10 Requests per Connection test.

Application Delivery Controller

2010 PERFORMANCE REPORT

Device	Power Draw (Watts)		Performance Max L7 RPS	Efficiency TPS/ Watt
	Idle	Full Load		
VIPRION (PB200) x 2	1100	1320	3,000,043	2,273
VIPRION (PB200) x 1	660	880	1,535,425	1,745
BIG-IP 8900	340	390	1,299,486	3,332
BIG-IP 3900	97	149	600,001	4,027
ACE20	754	807	86,222	107
ACE4710	107	110	39,157	356
MPX-17K nCore	440	470	500,728	1,065
MPX-17K Classic	416	424	144,091	340

Analysis – Power Efficiency

An ADC is more energy efficient the higher its RPS/watt rating. Using this metric, the F5 devices have 63%–3700% greater energy efficiency compared to the other vendor's products.

These tests represent two extremes. One is the worst case, where the device is on, but doing nothing (idle). The other extreme is when the device is processing the most HTTP requests it can.

While these results are very favorable to F5, the reality is that most applications use a range of file sizes, so even under full load, an ADC is unlikely to reach these RPS rates.

Application Delivery Controller

2010 PERFORMANCE REPORT

Appendix A: Additional Testing Details

Vendor	Product	Model	Software Version	Uplink to Switch
F5	BIG-IP	3900	10.1	4 x 1 Gbps
F5	BIG-IP	8900	10.1	2 x 10 Gbps
F5	VIPRION	PB200	10.1	4 x 10 Gbps
Cisco	Application Control Engine (ACE)	ACE20	A2(3.0)	2 x 10 Gbps
Cisco	Application Control Engine (ACE)	4710	A3(2.4)	4 x 1 Gbps
Citrix	NetScaler	MPX-17000	9.1 Build 102.8 nCore	2 x 10 Gbps
Citrix	NetScaler	MPX-17000	9.1 Build 102.8 Classic	2 x 10 Gbps

Load Generating Equipment

Load testing equipment from Ixia was used to generate and measure traffic to and from the ADCs. The hardware consisted of two XM12 Chassis, 20 ASM1000XMv12x-01 load modules and four ACCELERON-NP load modules. Software version 5.30.450.31GA was used.

Test Settings

The following test settings were changed from the defaults in the Ixia software:

- File Sizes: 128B, 2Kb, 4KB, 8KB, 16KB, 32KB, 128KB
- Number of Client IP addresses: 192
- Connections per user: 1
- HTTP Requests per connect: 1, 10, infinite (depending on type of test)
- Simulated Users: 512, 1024, 1536 or 3072 (depending on DUT)
- Servers: 72

Results reported are for tests with the SimUsers set as follows:

- 512 SimUsers: BIG-IP 3900, ACE 4710, NetScaler MPX-17000
- 1024 SimUsers: ACE20 Module
- 1536 SimUsers: BIG-IP 8900, VIPRION w/ (1) PB200 blade
- 3072 SimUsers: VIPRION w/ (2) PB200 blades

Test Specific Details

SSL Performance

- SSL Ciphers: RC4-MD5, AES256-SHA
- SSL Certificates: 1024 bit Key

Application Delivery Controller

2010 PERFORMANCE REPORT

Compression

- Compression method: gzip
- Files: all files were HTML, crafted to be moderately compressible, of 60-80% depending on file size (as is typically seen on the internet).

Caching and Compression

- The same files were used as in the compression tests.

DUT to Switch Connections

In all the tests for every device, all the interfaces connected from the DUT to the switch used link aggregation and 802.1Q VLAN tags. The link aggregation was manually configured with LACP disabled. Each DUT was connected to the switch with enough interfaces so the available bandwidth met or exceeded its specified performance capability.

Appendix B: Questions and Answers

1: Why don't the test results match the specifications published by the vendor?

There are several possible reasons why the results published in this report do not match the vendor's performance specifications. One of the most common reasons is vendors using different definitions for a performance measurement. L7 HTTP Requests Per Second (RPS) is another measurement vendors have defined differently. F5 defines L7 HTTP RPS as: Full TCP connection establishment (three way handshake), HTTP request & response (complete HTTP transaction) and TCP close (FIN, ACK, FIN, ACK), with the device inspecting each HTTP request and making a decision based on it. If only one HTTP Request is sent per TCP connection, F5 calls that Connections per second (CPS). If more than one request is sent per TCP connection, that is a measurement of RPS.

Some vendors define L7 HTTP RPS as the number of HTTP requests transmitted across connections of which the device is only doing L4 load balancing (no inspection or decisions based on HTTP content). In this scenario, the device is unaware of the HTTP requests or responses, as it is only processing data up to L4.

In order to ensure devices are inspecting the HTTP requests, F5 creates a rule (or policy) on the device being tested. This rule examines the URI (the part of the URL after the name of the server) of the HTTP request. If the URI ends in .png, then it directs the request to a different pool of servers than all other requests.

Another common reason a device may not reach its specified performance level is differences in test settings. There are many configurable parameters on both the device being tested and on the testing equipment that can affect performance. For many of these there is not one "correct" setting for a given type of test. Multiple settings may be equally valid.

The Ixia Simulated Users setting is a perfect example of this. One vendor might use a setting of 512 simulated users while a second vendor might set it at 5,000 and a third vendor might not set it at all and instead let the Ixia equipment adjust it dynamically during the test. All are valid options, however, one might be more appropriate than the others depending on what scenario the test is trying to simulate.

2: Can the device meet the vendor's published performance specifications?

Extensive tests were conducted to try and achieve some of the performance specifications of the ACE and NetScaler products. In some cases, they could be reached and in others we were not able to replicate them. The performance measurements we focused on were L4 Connections per Second (CPS), L4 throughput, L7 RPS and L7 throughput.

NetScaler MPX-17000 nCore: Citrix's published specifications for this device are 18Gbps of throughput and 1.5M HTTP Requests per Second. The only way to achieve the 1.5M HTTP RPS was to L4 load balance and send an unlimited number of requests across each TCP connection. In this scenario, the NetScaler was not doing any processing of the HTTP.

In order to achieve the specified L2 throughput, we had to: enable TCP Window Scaling and Selective Acknowledgement on the NetScaler and Ixia, increase the Ixia's TCP Buffer setting from 4KB to 64KB, use a L4 virtual server, use a file size of 512KB and send an unlimited number of HTTP requests across each TCP connection.

ACE20: We were unable to reach the ACE20's specified throughput of 16Gbps with any configuration settings we tried.
ACE20 and ACE 4710: The maximum L7 RPS results in the L7 1 Request per Connection test are the highest we were able to achieve with any combination of settings.

The decision to adjust the test settings for each device was made in order to represent each device at its best. If all the devices were in the same performance range, the decision would most likely have been to test all devices with the same SimUsers setting.

Application Delivery Controller

2010 PERFORMANCE REPORT

3: Why is throughput reported as L7 throughput instead of L2 throughput?

As noted earlier in this report, F5 typically reports throughput that counts all the data sent across the interfaces. Stated another way, by counting all the bits in every Ethernet frame. This is the industry standard method for measuring throughput and is referred to as Layer 2 throughput. This report is an exception to that practice. Here we are reporting Layer 7 throughput, due to a limitation of the load generating equipment used during these tests. That equipment does not report L2 throughput.

All vendors use L2 throughput in their published performance specifications. This difference in published specifications using L2 throughput and the test results in this report using L7 throughput should be kept in mind if comparing the numbers.

4: How much difference is there between L2 and L7 throughput?

There is not an absolute answer to this question, but only some general guidelines. The difference is affected by many factors such as the test type, files size, if VLAN tags are used, HTTP errors, and TCP resets and retransmits. The L2 throughput can be calculated if the assumption is made that there were no errors of any kind.

This example of the difference between L2 and L7 throughput is based on the L7 10 Requests Per Connection test with 128 byte file size. 802.1Q VLAN tags were used. Setting up the TCP connection, making 10 HTTP requests, getting the responses and closing the TCP connection, requires 25 Ethernet frames, totaling 6,190 bytes. Of those 6190 bytes, 4700 bytes are TCP payload, which are what Ixia counts for L7 throughput. In this example, L2 throughput is 31.7% more than L7 throughput ($4700 * 1.317 = 6190$).

5: How is compression throughput affected by the compressibility of files?

How compressible the files are has a dramatic influence on the throughput of an ADC when compressing traffic. In these tests, we used files that were 60%-70% compressible. To demonstrate the effect of compressibility on throughput, we ran additional tests using a 128KB file that is 90%+ compressible, and then using one that is less than 20% compressible. We only ran the tests on the MPX-17000 nCore and a single VIPRION blade because they are both rated for 18Gbps of throughput and use software compression. The table below shows the results along with the results from the main compression tests.

File Compressibility	Throughput (Gbps)	
	VIPRION x 1	MPX-17000 nCore
High	11.9	7.9
Moderate	5.2	4.1
Minimal	3.3	2.4

6: How did you decide what settings to use in the tests?

In our 2007 Performance Report, we tested with file sizes of 128B, 2KB, 8KB, 64KB and 512KB. We recently received feedback from some of our largest customers. They felt the 512KB file size was too large and that we should test more smaller file sizes. Based on that feedback, we chose to use file sizes of 128B, 2KB, 4KB, 8KB, 16KB, 32KB and 128KB.

Application Delivery Controller

2010 PERFORMANCE REPORT

7: Why was the Ixia “simulated users” setting changed for some devices?

Extensive preliminary tests were conducted to verify DUT configurations and identify settings that might be limiting performance. We know from experience that devices usually perform best when the Ixia was set to one number of simulated users (SimUsers) compared to other numbers.

During this preliminary phase, the L4 and L7 tests were run with SimUsers setting of 512, 1024, 1536, 2048, 3072 and 4094. If performance degraded from the 1024 to the 1536 settings, and again from 1536 to 2048 SimUsers, then tests were not run at the 3072 and 4096 SimUsers settings.

Because the Ixia equipment is trying to make connections and requests as quickly as it can, one simulated user generates a load similar to several hundred or several thousands of real users.

The results presented in this report are from tests with the SimUsers setting at which each ADC performed best. As an example, the BIG-IP 3900 and NetScaler MPX-17000 both performed best when the tests were set to 512 SimUsers and those are the test results included in this report. The BIG-IP 8900 performed best with 1,536 SimUsers so the results are from those tests.

One surprising result was that the ACE20 module performed almost exactly the same with SimUsers set to 512, 1024, 1536 and 2048.

The decision to adjust the test settings for each device was made in order to represent each one at its best, and was necessary to compensate for the varying performance ranges of the tested devices.

8: Why was the Test Objective setting changed for some devices?

Ixia tests require setting an objective, such as Requests Per Second (RPS) or throughput. Our typical practice is to set the objective higher than any of the devices can reach and test all of them with that setting. The objective is usually 1.5M RPS (called Transactions Per Second in the Ixia software). The results are usually the same as when testing a device and setting the objective to the maximum of what that device is capable.

In the case of the BIG-IP 8900 and the VIPRION, we had to set the objective higher than 1.5M RPS for some tests because they are capable of more than this.

During testing of the NetScaler, we observed large swings in performance on a second-to-second basis if the test objective was higher than the NetScaler was capable. The swings increased in size as the test objective increased.

Because of this behavior, we changed the Ixia Objective setting when testing the NetScaler to be only a little higher than what the NetScaler was able to achieve.

9: Why were only two interfaces on the Citrix NetScaler MPX-17000 connected to the switch?

The NetScaler MPX-17000 model tested has four 10Gigabit Ethernet interfaces. They are arranged two-over-two, numbered from left to right and top to bottom. The ports are identified as 1/1, 1/2, 1/3 and 1/4.

During testing, we found the performance of the NetScaler changed based on which interfaces were connected to the switch. The best performance in one test was seen when only one interface was connected (test was L7 RPS, infinite requests per connection and server connection reuse; 9% better than with ports 1/1 and 1/3 trunked). In all other tests, the best performance was achieved when the trunk from the MPX-17000 to the switch contained one interface from the top two and one interface from the bottom two (i.e. 1/1 and 1/3).

Application Delivery Controller

2010 PERFORMANCE REPORT

If the trunk consisted of two interfaces from the top, two from the bottom, or all four interfaces, performance was 7% - 9% lower than when the trunk had one port from the top and bottom. This behavior was consistent across all combinations of ports, multiple tests of each combination, multiple test types and various configuration settings on the NetScaler.

In order to present the best performance of the NetScaler, the results in this report are all from tests with ports 1/1 and 1/3 trunked together.

Glossary

Application Delivery Controller (ADC)

ADC's provide functionality such as load balancing, application acceleration, and server offload.

Denial of Service (DoS) or Distributed Denial of Service (DDoS)

Attacks that try to exhaust system resources by overloading network or authentication stacks on servers. Distributed attacks often originate from clients that have been compromised and are controlled by a central master node (a botnet), thereby amplifying the number of simultaneous connections/requests that can target a specific site. SYN Flood attacks are also common DoS attacks because they exhaust the network stack because each SYN packet opens a socket on the receiving server or device that consumes memory while it remains open.

HTTP Request Multiplexing (Connection reuse)

Refers to an ADC's capability to keep an HTTP connection open between the ADC and the server, even while there is no connection between the ADC and the client. Additionally, it specifies that these connections be used to handle as many HTTP requests from as many clients as possible, such that a small number of connections will be open between the ADC and servers, while many connections are opened and closed between the ADC and client. The net result is that servers are offloaded from having to process many TCP connection setup/teardown events. This is called OneConnect™ by F5.

Layer 7 (L7)

Refers to the top layer of the OSI protocol model. HTTP is an example of a layer 7 protocol.

Server

A computer running an application or service that is accessed by client computers. When used in the context of load balancing or application delivery, servers are the computers to which the ADC distributes the traffic.

Secure Sockets Layer (SSL)

An SSL transaction is defined as: TCP connection setup, SSL session setup, HTTP request, HTTP response, SSL shutdown, and TCP teardown.

Virtual Server

Virtual servers are a specific combination of IP address and TCP or UDP port number that are associated with a specific application. The Virtual Server is configured on an ADC that will load balance the traffic for the Virtual Server to the servers running the application.

F5 Networks, Inc. 401 Elliott Avenue West, Seattle, WA 98119 888-882-4447 www.f5.com

F5 Networks, Inc.
Corporate Headquarters
info@f5.com

F5 Networks
Asia-Pacific
info.asia@f5.com

F5 Networks Ltd.
Europe/Middle-East/Africa
emeainfo@f5.com

F5 Networks
Japan K.K.
f5j-info@f5.com

