TEST REPORT

Tolly.

#209127
June 2009

Commissioned by Citrix Systems, Inc.

Citrix NetScaler with nCore™ Outperforms F5 BIG-IP v10
NetScaler MPX 17000 ADC: Latency versus BIG-IP 8800

EXECUTIVE SUMMARY

Latency is an important metric that should be considered when evaluating web application delivery solutions. Excessive latency (i.e. delay) results in slower application performance which can jeopardize time-sensitive transactions and negatively impact a company's bottom line.

Citrix NetScaler exhibited very low latencies and was at least 3X faster than F5 BIG-IP in almost all test scenarios.

TEST HIGHLIGHTS

1. Delivers web traffic at least 3X faster than the F5 BIG-IP in most tests
2. Consistently demonstrates lower latency & CPU utilization and higher throughput
3. Under heavy loads and high connection counts the F5 BIG-IP adds considerable latency that could degrade user experience

Latency at Various Connection Loads - Basic L7
Maximum Supported 10GbE Interfaces
8,192-byte Response Object as Reported by Ixia
(lower numbers are better)

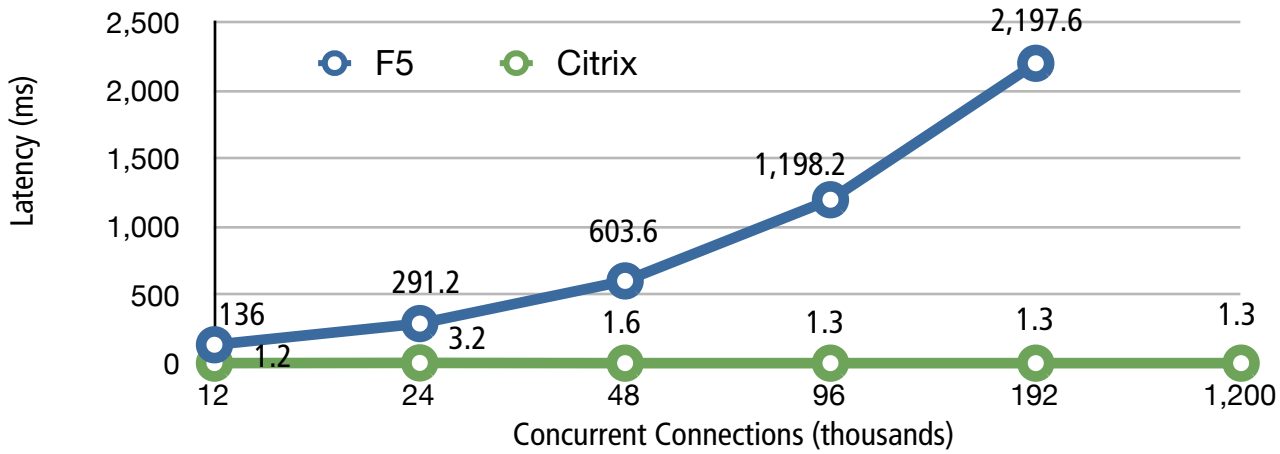

Note: The F5 device could not sustain a load of 1.2 million connections.

Source: Tolly, June 2009

Figure 1

Overview

Citrix Systems, Inc. commissioned Tolly to evaluate its NetScaler MPX 17000 application delivery controller running software release 9.1 with nCore technology. This was compared with the F5 Networks BIG-IP 8800 running its newly-released version 10 software. Testing measured system latency across a number of layer 7 load-balancing and content-switching scenarios. Each device uses multi-core processing to improve performance and scalability.

Both devices were benchmarked using the maximum number of 10GbE ports supported.

Across the suite of tests involving loads ranging from 12,000 connections to over 1 million concurrent connections, the Citrix NetScaler exhibited dramatically lower latency than F5 BIG-IP. Excessively high latency can have a significant negative impact on web-based applications and businesses. (See sidebar "Impact of Latency on the Bottom Line" for more information.)

Test Results

Basic & Advanced L7

In tests of basic L7 policy configurations and retrieving an 8,192-byte response object, NetScaler completed transactions, on average, in less than 3.2 milliseconds for all tests.

F5 BIG-IP, under the identical load and test scenarios, required at least 136 milliseconds. For high loads of 192,000 concurrent connections, BIG-IP introduced a delay of 2,197.6 milliseconds, on average, for each transaction -- over 2 seconds slower than NetScaler to deliver the same traffic. (See Figure 1.)

For the advanced L7 policy configuration, NetScaler again demonstrated faster performance with average transactions completing in less than 3 milliseconds for all tests. F5 BIG-IP, on the other hand, added at least 164 ms of latency under the exact same conditions. At 192,000 connections, F5 latency exceeded 2.5 seconds, on average, for each transaction. (See Figures 2 and 3.)

Citrix Systems, Inc.

NetScaler MPX 17000 with 9.1 nCore

Competitive Performance

Tested June 2009

Latency at Various Connection Loads - Advanced L7
Maximum Supported 10GbE Interfaces
8,192-byte Response Object as Reported by Ixia
(lower numbers are better)

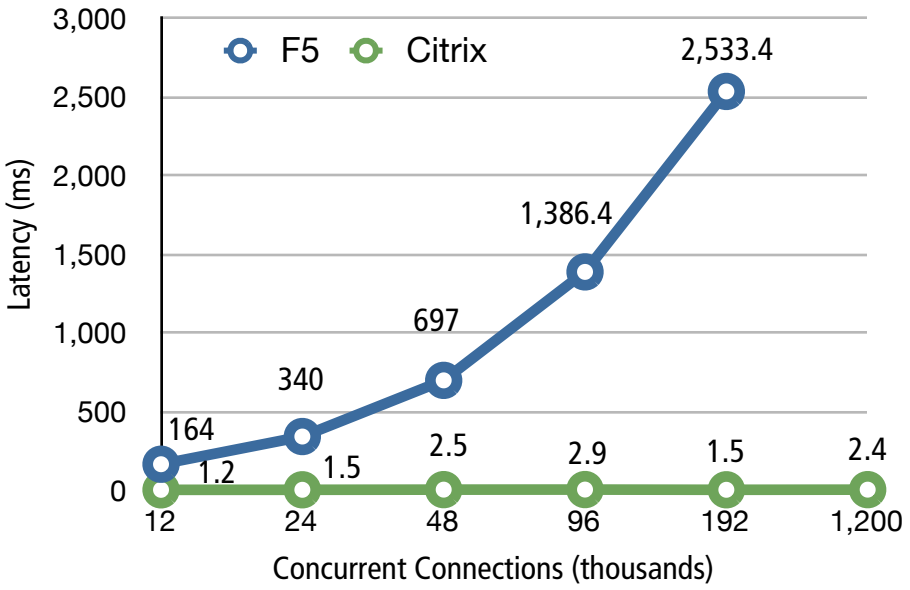

Note: The F5 device could not sustain a load of 1.2 million connections.

Source: Tolly, June 2009

Figure 2

Advanced L7 Test - Additional Measurements

Maximum Supported 10GbE Interfaces
8,192-byte Response Object

Connections (k)	Latency (ms)		Goodput (Gbps)		CPU (%)	
	Citrix	F5	Citrix	F5	Citrix	F5
12	1.2	164	8.4	5.0	30.6	90
24	1.5	340	8.3	4.8	28.6	90
48	2.5	697	8.3	4.6	29.2	90
96	2.9	1,386.4	8.3	4.5	29.7	90
192	1.5	2,533.4	8.3	4.3	28.6	90
1,200	2.4	N/A	8.0	N/A	30.8	N/A

Note: Latency and throughput as measured by Ixia. CPU utilization as reported by each device's console.

Source: Tolly, June 2009

Figure 3

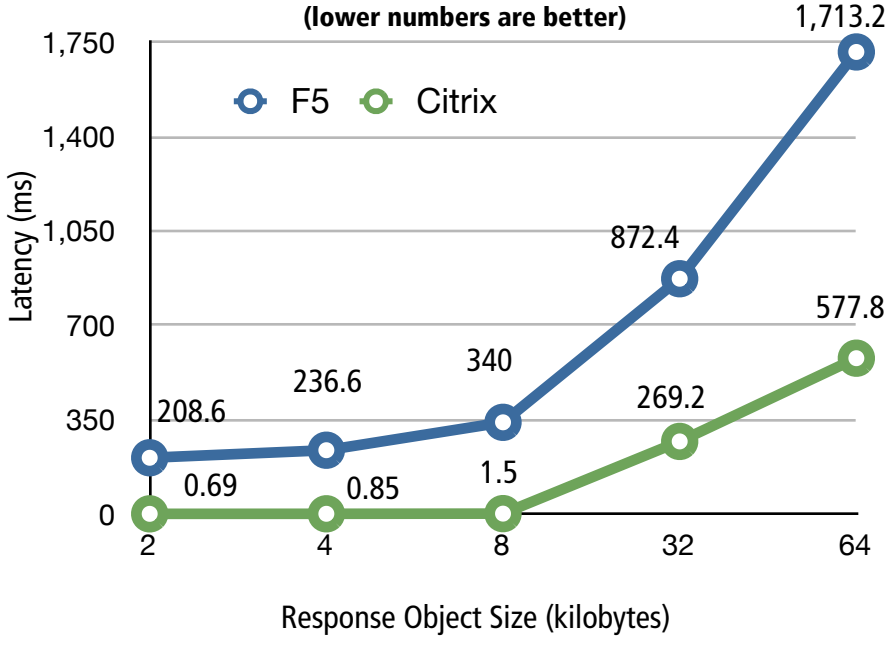

Source: Tolly, June 2009

Figure 4

It is noteworthy that at the maximum test load, 1.2 million connections, NetScaler's latency remained below 2.5 milliseconds. The F5 BIG-IP was unable to complete this test as the device could not sustain the load without excessive traffic loss.

Even at the lowest connection load tested, 12,000 connections, the BIG-IP's latency exceeded the NetScaler's highest latency at any load by more than 40X. (See Figures 1 and 2.)

Varying Response Sizes

To eliminate the possibility that NetScaler's faster performance was only evident at a single, fixed response size of 8,192 bytes, additional tests were run at varying response sizes. Load was set at a

relatively moderate level of 24,000 concurrent connections.

Again, NetScaler outperformed BIG-IP in every test demonstrating significantly lower latency. For example, with object sizes up to 8,192 bytes, the NetScaler exhibited latency of 1.5 ms or less while BIG-IP required between 200 and 340 ms to support the same traffic. For larger size objects, NetScaler's performance advantage (i.e. lower latency) only improved. (See Figure 4.)

Exploring CPU Utilization

Testers noted that the CPUs of the F5 device were fully taxed even at the lowest level of concurrent connections, registering 90% utilization at only 12,000 connections. (According to F5 documentation, 10% of each CPU's capacity is reserved for management functions.)

To characterize BIG-IP performance at lower CPU levels testers throttled back the load on the BIG-IP, measuring CPU utilization until a load was found that would drive the CPU to less than 70%. Testers then ran the Advanced L7 test again.

Even at lower CPU utilization, F5 BIG-IP still introduced significantly more latency than NetScaler.

At the highest level tested for BIG-IP, the latency was measured at 751 ms - still far longer than the 1.54 ms measured for NetScaler. (See Figure 5.)

Impact of Latency on the Bottom Line

Reducing latency is critical to successfully delivering web applications. High application latency results in long web page load times, poor user experience and adoption, customer dissatisfaction, poor public perception and negative impact to a company's bottom line.

Examples:

Amazon found every 100ms of latency cost them 1% in sales.

Google found an extra .5 seconds in search page generation time dropped traffic by 20%.

A broker could lose $4 million in revenues per millisecond if their electronic trading platform is 5 milliseconds behind the competition.

Source of examples: Todd Hoff - Highscalability.com Used with permission

Latency at Various Connection Loads - Advanced L7 Reduced Load for F5 (~70% CPU)
Maximum Supported 10GbE Interfaces
8,192-byte Response Object as Reported by Ixia
(lower numbers are better)

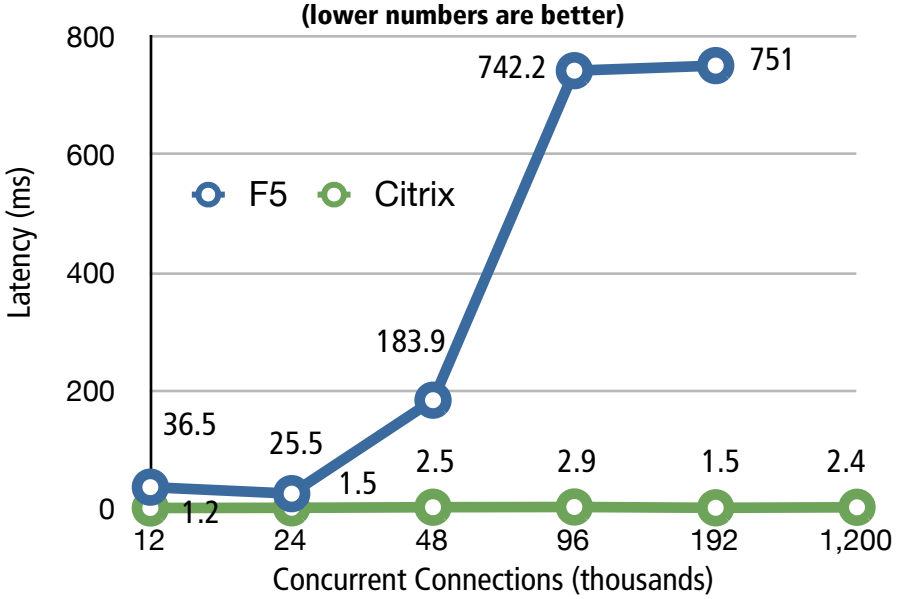

Note: The F5 device could not sustain a load of 1.2 million connections. A wait time was introduced to lower the CPU load on the F5 device. Citrix results are the same as test shown in Figure 2 and are shown here to make comparisons easier.

Source: Tolly, June 2009 Figure 5

Finally, to see whether NetScaler's superior performance was maintained at high CPU utilization, testers measured the NetScaler system latency when the CPUs hit maximum utilization. To do this, compression was enabled for both devices.

As the NetScaler implements a software-based compression algorithm these tests would push the CPUs to their maximum. (According to Citrix, one of the 8 CPUs is reserved for management functions, with the remaining 7 being used for processing application traffic.)

Despite running at max CPU load, NetScaler consistently exhibited lower latency than BIG-IP across all connection count levels. Additionally, as BIG-IP's CPU utilization increased with higher concurrent connection levels, the difference in response time between the two devices increased dramatically.

At 90% CPU utilization versus NetScaler at 100% CPU utilization, BIG-IP introduced almost 2.5X more latency than NetScaler: 1202.4 ms vs. 490.0 ms. (See Figure 6.)

Test Methodology

Basic & Advanced Layer 7 - 8KB Response Objects

In these tests, HTTP requests for 8KB objects were sent across the DUT. Tests were run with concurrent connection counts ranging from 12,000 to 1.2 million.

Latency with NetScaler at 100% CPU Utilization
Maximum Supported 10GbE Interfaces
32KB Response Object as Reported by Ixia
(lower numbers are better)

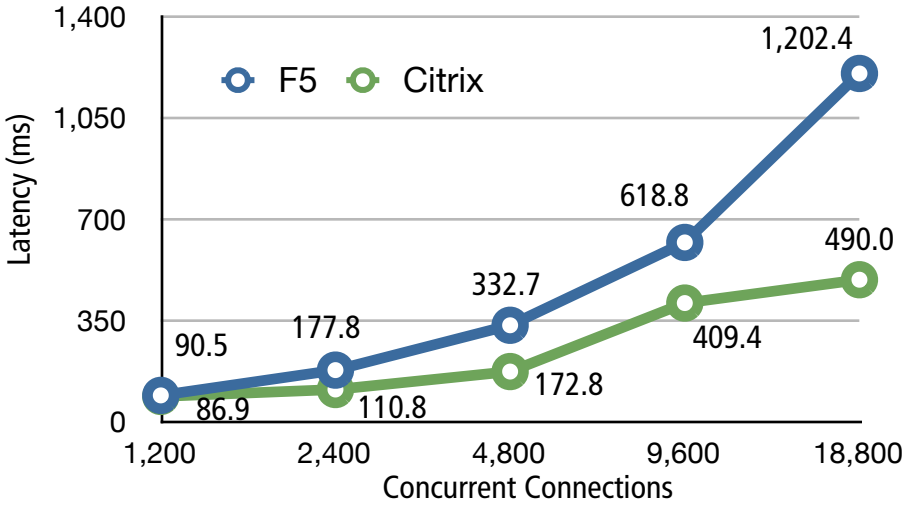

Note: The F5 device was not running at 100% CPU. It ranged from 35 to 90% CPU utilization. For F5, version 10.0.1 was used.

Source: Tolly, June 2009

Figure 6

Test Bed Topology

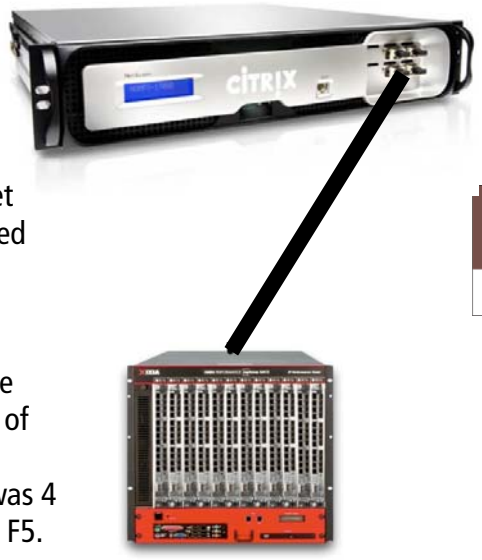

10 Gigabit Ethernet interfaces connected Ixia to the DUT.

Each DUT was configured with the maximum number of 10 GbE interfaces supported which was 4 for Citrix and 2 for F5.

PERFORMANCE VERIFIED By IXIA

Source: Tolly, June 2009

Figure 7

In the basic test, each DUT was configured to have 24 virtual IP addresses (VIP), round robin load balancing was used, and TCP connection multiplexing was enabled for each device.

In the advanced tests, each DUT was configured to send HTTP requests to different back-end servers based on the content of the URL. Additionally, URL rewriting was enabled for 50% of the traffic. Client IP insertion took place on all requests and source IP persistence was enabled. Engineers measured latency and also observed system CPU as reported by the DUT's console.

Testers noted that during these tests the BIG-IP reported average CPU utilization of 90% compared to NetScaler, which maintained utilization between 24% and 38%.

In order to see the impact on BIG-IP results when the system CPU was run at a lower utilization level, the load was throttled back. For the advanced L7 test with 12,000 connections, BIG-IP was run at approximately 70% CPU utilization.

Varying Object Sizes

For this test, the concurrent connection count was fixed at 24,000 while the response object size from the server was varied from 2Kbytes to 64Kbytes.

NetScaler Maximum CPU

As noted in earlier tests, F5 BIG-IP's CPU utilization was measured at 90% for many of the test scenarios while Citrix NetScaler, as noted earlier, operated in a range of 24%-38% CPU utilization for the same tests.

For this test, engineers wanted to determine if Citrix NetScaler can maintain its latency advantage even under full load while performing at maximum CPU utilization.

Engineers enabled dynamic GZIP compression for each device. As NetScaler implements compression in software, it is a CPU-intensive function for NetScaler. (Apparently BIG-IP uses a hardware-assist for this function.)

Testers ran a series of compression tests at varying connection counts ranging from 1,200 to 18,800 concurrent connections for both devices with a response size of 32KB. With compression enabled, Citrix NetScaler ran at 100% CPU utilization while F5 BIG-IP ranged from 35% - 90% utilization.

Test Bed

The test bed consisted of an Ixia OptIxia XM12 chassis with 4x 10G Application Performance Modules, emulating both the client and server sides of the sessions. (See Figure 7.)

A Citrix NetScaler MPX 17000 application delivery appliance was running version 9.1 with nCore technology. The NetScaler was outfitted with 4 10GbE connections. The F5 Networks BIG-IP 8800 was running software version 10.0. It was outfitted with the maximum supported 2 10GbE ports. (The F5 was also outfitted with Gigabit Ethernet ports but those were not used in the test.) The 10GbE ports represented the maximum supported for each DUT.

The NetScaler tested ran eight cores and had 32GB of memory. The F5 device is also a multi-processor device. The unit tested ran four cores and had 4GB of memory. Both appliances tested are fixed configurations. That is, customers cannot opt for different processor and/or memory configurations for these models.

Both DUTs were reset to factory defaults. Changes to the configurations were made only for the compression test. Each device was set to compress to GZIP level 5 in order that the load on each DUT would be equivalent. GZIP ranges from level 1 to level 9, lowest to highest compression levels.

About Tolly…

The Tolly Group companies have been delivering world-class IT services for 20 years. Tolly is a leading global provider of third-party validation services for vendors of IT products, components and services.
You can reach the company via E-mail at sales@tolly.com, or via telephone at +1 561.391.5610.

Visit Tolly on the Internet at:
http://www.tolly.com

Interaction with F5

In accordance with Tolly's Fair Testing Charter, Tolly personnel invited representatives from F5 to be involved in the testing. F5 declined that invitation.

For more information on the Tolly Fair Testing Charter, visit:
http://www.tolly.com/FTC.aspx

Terms of Usage

This document is provided, free-of-charge, to help you understand whether a given product, technology or service merits additional investigation for your particular needs. Any decision to purchase a product must be based on your own assessment of suitability based on your needs. The document should never be used as a substitute for advice from a qualified IT or business professional. This evaluation was focused on illustrating specific features and/or performance of the product(s) and was conducted under controlled, laboratory conditions. Certain tests may have been tailored to reflect performance under ideal conditions; performance may vary under real-world conditions. Users should run tests based on their own real-world scenarios to validate performance for their own networks.

Reasonable efforts were made to ensure the accuracy of the data contained herein but errors and/or oversights can occur. The test/audit documented herein may also rely on various test tools the accuracy of which is beyond our control. Furthermore, the document relies on certain representations by the sponsor that are beyond our control to verify. Among these is that the software/hardware tested is production or production track and is, or will be, available in equivalent or better form to commercial customers. Accordingly, this document is provided "as is", and Tolly Enterprises, LLC (Tolly) gives no warranty, representation or undertaking, whether express or implied, and accepts no legal responsibility, whether direct or indirect, for the accuracy, completeness, usefulness or suitability of any information contained herein. By reviewing this document, you agree that your use of any information contained herein is at your own risk, and you accept all risks and responsibility for losses, damages, costs and other consequences resulting directly or indirectly from any information or material available on it. Tolly is not responsible for, and you agree to hold Tolly and its related affiliates harmless from any loss, harm, injury or damage resulting from or arising out of your use of or reliance on any of the information provided herein.

Tolly makes no claim as to whether any product or company described herein is suitable for investment. You should obtain your own independent professional advice, whether legal, accounting or otherwise, before proceeding with any investment or project related to any information, products or companies described herein. When foreign translations exist, the English document is considered authoritative. To assure accuracy, only use documents downloaded directly from Tolly.com.

No part of any document may be reproduced, in whole or in part, without the specific written permission of Tolly. All trademarks used in the document are owned by their respective owners. You agree not to use any trademark in or as the whole or part of your own trademarks in connection with any activities, products or services which are not ours, or in a manner which may be confusing, misleading or deceptive or in a manner that disparages us or our information, projects or developments.

209127-djqsbE-jft-30JUN09-veri