F5 Networks
V9 Performance
Report

A Broadband-Testing Report

First published January 2005 (V1.0)

Published by Broadband-Testing
La Calade, 11700 Moux, Aude, France

Tel : +33 (0)4 68 43 99 70
Fax : +33 (0)4 68 43 99 71
E-mail : info@broadband-testing.co.uk
Internet : http://www.broadband-testing.co.uk

TABLE OF CONTENTS

EXECUTIVE SUMMARY ... 1

THE AIMS OF THIS REPORT ... 2

Overview ... 3

THE TESTING ... 3

Test Focus ... 3
Transparency ... 3
Depth of Detail ... 3
Real-world ... 4
Total System Performance ... 4
The Test Validation Elements ... 5
The Tests Described ... 6
The L4 Tests ... 8
The L7 Tests ... 8
Compression Tests ... 8
SSL Tests ... 8
Total System Performance Tests ... 9
The L4 Test Results ... 9
The L7 Test Results ... 12
The Compression Test Results ... 14
The SSL Test Results ... 16
The Total System Performance Test Results ... 18
The DDoS Test Summarised ... 19

SUMMARY OF TEST RESULTS ... 20

TEST SECTION 2: THE OTHER TEST RESULTS ... 21

TABLE OF FIGURES

Figure 1 – SSL Transactions Per Second Example Test ... 2
Figure 2 – The Test bed Topology ... 6
Figure 3 – L4 Connections per Second ... 10
Figure 4 – L4 Throughput ... 10
Figure 5 – L4 Response Time ... 11
Figure 6 – L7 Connections Per Second ... 12
Figure 7 – L7 Throughput ... 13
Figure 8 – L7 Response Time ... 13
Figure 9 – Compression CPS Test ... 14
Figure 10 – Compression Throughput Test ... 15
Figure 11 – Compression Response Time ... 15
Figure 12 – SSL TPS – 1 Session ID Use – RC4-MD5 Cipher ... 16
Figure 13 – SSL TPS – 10 Session ID Use – RC4-MD5 Cipher ... 17
Figure 14 – SSL TPS – 1 Session ID Use – DES-CBC3-SHA Cipher ... 17
Figure 15 – SSL TPS – 10 Session ID Use – DES-CBC3-SHA Cipher ... 18
Figure 16 – Total System Performance Test ... 18
Figure 17 – DDoS Attack Test ... 19
Figure 18 – L4+7 CpS Test ... 21
Figure 19 – L4+7 Throughput Test ... 21
Figure 20 – L4+7 Response Time ... 22
Figure 21 – L7 Requests Per Second – 10 Transactions Per Connection ... 22
Figure 22 – L7 Throughput – 10 Transactions Per Connection ... 23
Figure 23 – L7 Response Time – 10 Transactions Per Connection ... 23
Figure 24 – L7 Requests Per Second – Unlimited Transactions Per Connection ... 24
Figure 25 – L7 Throughput – Unlimited Transactions Per Connection ... 24
Figure 26 – L7 Response Time – Unlimited Transactions Per Connection ... 25
Figure 27 – L4 Throughput ... 25
Figure 28 – L7 Throughput ... 26
Figure 29 – L4+7 Throughput ... 26
Figure 30 – SSL Throughput (DES-CBC3-SHA) ... 27
Figure 31 – SSL Throughput (RC4-MD5) ... 27
Figure 32 – Compression Throughput ... 28
Figure 33 – Compression/SSL Throughput (DES-CBC3-SHA) ... 28
Figure 34 – Compression/SSL Throughput (RC4-MD5) ... 29
Figure 35 – Latency Test – L4 Response Time ... 30
Figure 36 – Latency Test – L7 Response Time ... 30
Figure 37 – Optimal SSL Performance – Client Side Encryption Only ... 31

Broadband-Testing

Broadband-Testing is Europe's foremost independent network testing facility and consultancy organisation for broadband and network infrastructure products.

Based in the south of France, Broadband-Testing offers extensive labs, demo and conference facilities. From this base, Broadband-Testing provides a range of specialist IT, networking and development services to vendors and end-user organisations throughout Europe, SEAP and the United States.

Broadband-Testing is an associate of the following:

- *NSS Network Testing Laboratories (specialising in security product testing)*
- *Broadband Vantage (broadband consultancy group)*
- *Limbo Creatives (bespoke software development)*

Broadband-Testing Laboratories are available to vendors and end-users for fully independent testing of networking, communications and security hardware and software.

Broadband-Testing Laboratories operates an **Approval** scheme which enables products to be short-listed for purchase by end-users, based on their successful approval.

Output from the labs, including detailed research reports, articles and white papers on the latest network-related technologies, are made available free of charge on our web site at http://www.broadband-testing.co.uk

The conference centre in Moux in the south of France is the ideal location for sales training, general seminars and product launches, and Broadband-Testing can also provide technical writing services for sales, marketing and technical documentation, as well as documentation and test-house facilities for product development.

Broadband-Testing Consultancy Services offers a range of network consultancy services including network design, strategy planning, Internet connectivity and product development assistance.

EXECUTIVE SUMMARY

There are a number of different vendors in the traffic management market, each attempting to stake a claim to performance leadership.

This report was produced to put these claims to the test. In creating this report, over 150 different tests were run, making this the most comprehensive test ever performed of Layer 4-7 traffic management products. These tests have revealed that F5's BIG-IP® product line, both hardware and software, has a significant performance lead on the competition, a lead that looks to be very difficult for competitors to match anytime soon.

The test featured products from Cisco Systems®, Nortel Networks®, Radware®, NetScaler® and Redline Networks®, in addition to the F5 BIG-IP device, all with the latest available software revisions and comparable configurations at the time of testing. Broadband-Testing audited and validated the comprehensive test plan and a broad series of tests carried out by F5 over several months. The result is that the F5 product came out in front, often by a wide margin. This held true, despite – in some cases – the other vendors own "non-marketing" published performance figures being met or exceeded during the testing.

The 150 tests were run in order to ensure that a very wide range of user scenarios were covered and that the results were as meaningful as possible. Each test in itself was extremely detailed, and each test suite ran for an excess of seven hours. The focus throughout the testing was on simulating real-world conditions, using a multi-device test bed consisting of traffic generators from Ixia® and Spirent Communications™.

This report is focused on five key test areas, which have been driven by both existing and prospective customers of these types of products. These key test areas are: L4, L7, compression, SSL, and a total system performance "combo" test, where we examined the performance patterns when all the features were set to "on". This "combo" test is one that most vendors tend to shy away from when putting their products forward for evaluation. But Broadband-Testing wasn't shying away from anything in these tests.

One notable breakthrough was in the SSL transaction performance achieved (see Figure 1 overleaf).

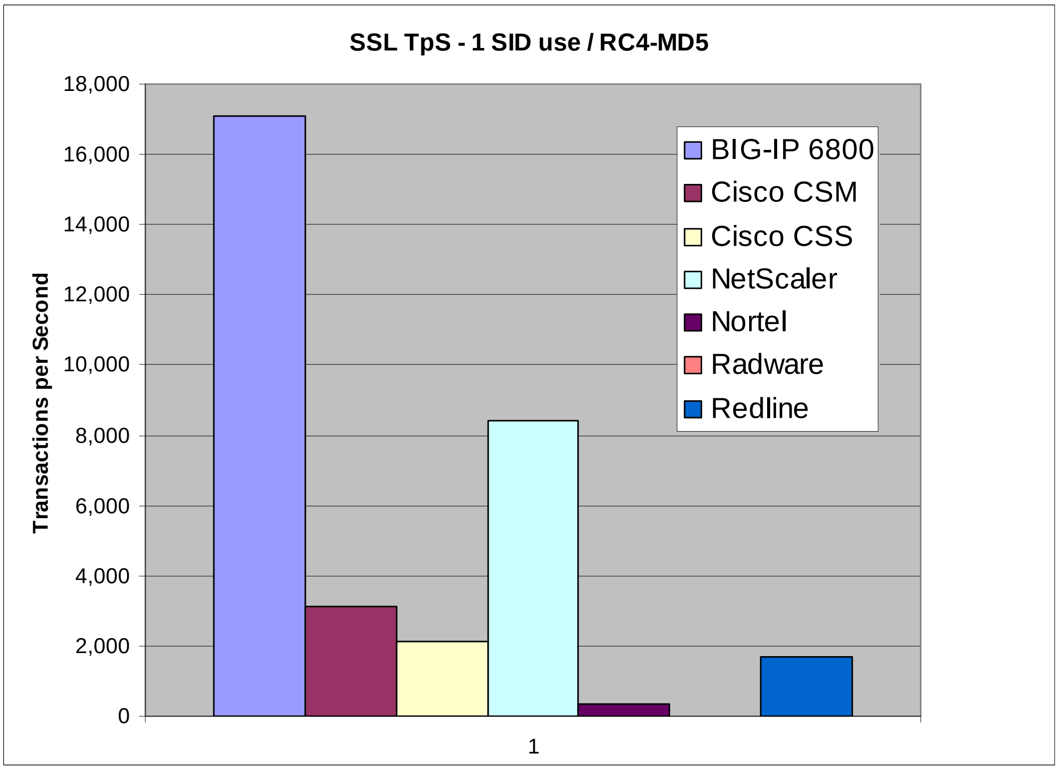

Figure 1 – SSL Transactions Per Second Example Test

Here we were achieving the figures claimed by all the vendors, with the exception of Redline whose published marketing numbers don't match, yet the BIG-IP 6800 doubled up on previous performance bests. This pattern – including where these test results even exceeded the other vendors' published performance claims – continued throughout the testing of the BIG-IP device and the validation of the 3rd party product tests.

What it points to is that F5 has gained a significant performance advantage on the competition with its new product range, based on the version 9 (v9) software and the new hardware platforms. Given that this is not a simple upgrade for F5, but a complete, ground-up reinvention of the BIG-IP products, it is not so surprising to see such a dramatic improvement in performance. A significant amount of development time has clearly gone into making this achievement possible.

The testing follows Broadband-Testing's earlier report on the intelligence and adaptability of the v9 software, most notably the new Traffic Management Operating System (TM/OS) that is present throughout the new BIG-IP product range. Here we found many functions unique to the BIG-IP device, as well as an exceptionally well-designed architecture which not only provides real depth of application traffic management now, but also allows for the company to easily expand its feature set and performance well into the future. The other report is available from F5's website at http://www.f5.com/solutions/v9_Functionality.pdf and, in conjunction with this one, provides a complete assessment of the BIG-IP products.

THE AIMS OF THIS REPORT

Overview

Within the scope of this report we've looked to test the performance of F5 Networks' BIG-IP product, featuring the v9 and TM/OS software release, by creating a wide-ranging series of tests that reflect many different real-world situations.

All are based around real application simulations and real application traffic. So the results really do mean something to a network manager or anyone who has to deal with the problems of application traffic management on a daily basis. Broadband-Testing has also verified some internal testing that F5 has carried out exhaustively, on the performance capabilities of some of their competitor's products, as part of a project that ran over several months.

This report follows our analysis of the intelligence and adaptability that V9 brought to the BIG-IP product range, which is available for free download from the F5 website, as outlined above.

Anyone wishing to follow up on any aspects of the report with the author, is welcome to contact me by email at sbroadhead@broadband-testing.co.uk

THE TESTING

Test Focus

There are four key areas the testing focuses on. These are:

- Transparency
- Depth of detail
- Real-world
- Total System Performance

Transparency

The testing we are carrying out here is completely open. All the test specifications are freely available at: http://www.f5.com/solutions/v9_performanceconfiguration.pdf, as are the test configurations for each device tested. The test tools we used from Ixia and Spirent can be purchased on the open market. This means that you are equally able to reproduce these tests in exact detail.

Depth of Detail

The test plan was devised over a long period of assessment and re-assessment. It covers over 150 tests per device, and these are all extensive tests running for over seven hours. So, total test time equates to many days of non-stop testing. This is in stark contrast to some tests you might read which, in some cases, literally run for a minute or less. These were true tests of the devices, both from a performance and stress perspective.
We also used two distinct sets of test tools – one from Ixia and one from Spirent Communications – in order to ensure we covered the breadth of testing we were looking to achieve, with the correct tools for the job. So the

Spirent WebAvalanche/WebReflector devices focus on application-oriented testing, while the Ixia devices focus on pure performance tests.

Real-world

The test schedule we put together was based around feedback from F5 customers, who told us what kind of tests they would like to see the devices subjected to. So the aim is for each and every result to be meaningful.

The test equipment we used generates real clients, real servers, and real traffic, so we only use the word “simulation” loosely; this was a lab-based test after all, but it is as close as we can come to invading the offices of an enterprise for several months and commandeering their network, which would be nice, but not really feasible.

The Ixia / Spirent test tools are both using a full TCP stack – these are not simply packet-generators, but are instead real clients whose performance varies based on latency and packet loss. Normal packet-level load testers can achieve a desired rate of traffic no matter if packets are dropped, or out of order, or received after substantial delays. Using clients with a full TCP stack ensures the resulting performance numbers are helpful when applied in the real-world, not just during benchmark testing.

All of these tests used standard Internet Explorer 5.x client headers, and full server response headers. This means that in a given test, such as 128byte page size, the actual size of the response will probably be closer to 400bytes, because default server headers are approximately 250bytes, just as you would find in the real world -- there are no tricks here.

In all tests, the throughput is measured in Ixia client applications bytes, not raw L2 throughput. It's common for vendors to show the throughput as measured at the network layer, which adds approximately 52 bytes to every packet, leading to an overall 8 - 15% inflation in the throughput numbers. In these tests, we are only measuring client applications bytes (L7 data), and multiplying this number by 8 to generate the commonly used bits per second measurement used for all networking equipment. What this means to an end-user is that a test which shows 2 gigabits per second, is actually showing 2 gigabits per second of L7 data, which is actually closer 2.3 gigabit per second in raw network bandwidth. Or 3.5 gigabits per second, is actually line-rate 4 gigabit per second.

As stated, these tests are designed around customer feedback – F5 asked their field engineers and customers what the average usage of a given feature is in a typical deployment, and used that feedback when constructing our Total System Performance Test for deciding what mix of various features should be used, and what the mix should be. In fact, customer feedback is how F5 decided which tests to run – these tests are not only tailored to answer common questions, the tests themselves were chosen based on customer demand.

Total System Performance

So what happens during a test when every feature is set to “on”? This is the test that most vendors like to avoid because it exercises all the key components on the device all at once. It is the single test that is most likely to show software bugs, or performance limitations that aren’t normally apparent in benchmark testing – it shows the performance one can expect in the real world where many features are used concurrently. This is exactly why we're running it, to see how the device under test copes in a full-on, real-world situation.

The typical vendor practice has been to publish performance numbers for specific functions that are tested in isolation while using unrealistic response sizes of say 1 byte. For example, the numbers you may see published for SSL performance are created while the device is doing nothing else but SSL encryption and decryption with an unlimited number of SSL ID (SID) reuses (to artificially inflate performance numbers). In my 8 years of experience with this market I have yet to see a device purchased and utilized solely for that singular function and I have never seen a useful application returning a response of 1 byte or secured with SSL utilizing unlimited SID reuse. In the real-world users tend to have a mix of functions and features, different traffic types, and a variety of request / response sizes. So the question becomes not only what is real-world individual SSL, L4, or L7 performance, but also what is the total performance of the device while running these functions concurrently with varying traffic types.

It has been my experience that users have all too often been mislead by vendors' inflated performance claims, and the intent of this particular test, as well as this whole report, is to shed light on this practice, get at the truth, and cut through the vendor marketing speak.

The Test Validation Elements

Broadband-Testing undertook testing of the F5 BIG-IP and validated testing by F5 of a number of competing devices from Cisco, Nortel, Redline Networks, NetScaler and Radware.

All these devices were tested running up-to-date, globally available software as available at the time of testing, with fully compatible configurations. One exception was with Redline Networks, where we tried unsuccessfully on several occasions to get the latest version (4.0) up and running, despite aid from its technical support team. We further validated the F5 code release by confirming its successful use (over five million connections) over an extended period of time, with an independent customer.

The line up was as follows:

- F5 BIG-IP 6800 Application Switch
- Cisco Content Services Switch (CSS) 11506
- Cisco CSM (WS-X6066-SLB-APC) and SSL Module (WS-SVC-SSL-1)
- NetScaler Application Switch 9950
- Nortel Application Switch 2424-SSL
- Radware Application Switch III (AS3)
- Redline Networks Web Application Processor E|X 3650

We used the following traffic generation devices to run the tests:

- Ixia 1600T (two fully-populated chassis), IxOS version 3.80, and IxLoad 1.10 SP2
- Spirent WebAvalanche 2500 and WebReflector 2500, version 6.2

These enabled us to generate multi-gigabit per second real web traffic flows with almost unlimited numbers of client and server simulations available. However, in order to keep the tests easily managed, repeatable and controllable, we standardised on 192 virtual client IP's and 32 servers. This gave us the ability to really push the devices under test, while ensuring that bottlenecks did not occur in the test bed itself, such as the core switch, but always at the device under test.

The actual test bed design was also kept as simple as possible, in order to ensure accuracy of results and easy repetition, with a client-side network and a server-side network and the device under test between the two. This configuration was repeated for each test and each device under test, though for the Redline device, which only has a single interface, we had to create an equivalent configuration which, in practice, worked in exactly the same way as for the other devices, the difference being that all devices, client and server side, had to be on the same network.

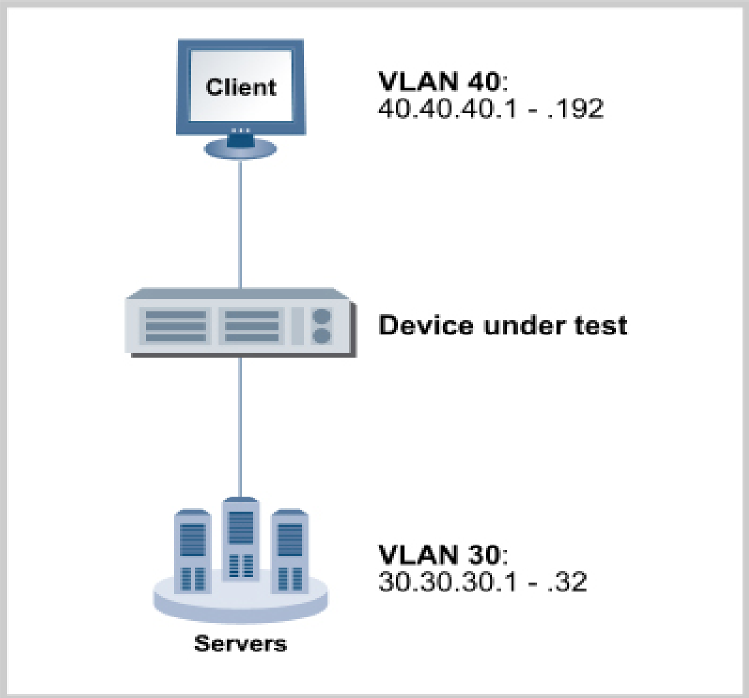

Figure 2 – The Test bed Topology

The Tests Described

In total, over 150 tests were run against each device. For reporting purposes we did pick out certain key tests that are the most applicable to most users in the real-world. These were L4, L7, compression, SSL, and the total system performance tests. We will also describe some Denial of Service (DoS) attack tests we ran. We are publishing the results of every other test performed, including the complete test configurations for each device under test, but that is not part of this particular report. You can get this other information by contacting F5 Networks directly.

Before standardising on the tests run, F5 researched any testing each vendor might have done themselves directly or via a 3rd party and tested many configurations before deciding on what exactly gave each device the best performance. Individual optimisations were also carried out for each device in order to ensure that they ran to their maximum potential, in many cases using the vendors own technical support team to do the optimisations. All this information is being published in a supplementary report, so that these tests can be accurately reproduced by anyone with access to the test bed devices. This again is in the spirit that these tests should be transparent, fair, and reproducible to satisfy even the greatest of cynics.

Once the tests were standardised, they were run on each device more than four times, in order to guarantee that the results were reproducible and accurate. Any time there was more than a 1% deviation in expected numbers, based on previous tests or vendor claims, the tests were rerun and, where necessary, settings were changed as part of the optimisation process, or vendor support was engaged to fix any problems.

Marketing Figures Versus Real-World Figures

During the testing, the performance results obtained for many devices in many tests equalled or even exceeded the vendors own performance claims, which is a trend for this market where too much misleading information has been published. As mentioned earlier, we did not succumb to “marketing figure” tests where special configurations are used – such as a reduced TCP stack (for example, using “perf-mode” or limited state TCP) – in order to achieve inflated results. It was our goal to define tests in a very clear way that use parameters that would be very meaningful to a customer for use as a sizing-guide in a purchasing decision. It's not uncommon to see a vendor claim "connections per second" by measuring "null connections", or by only measuring the new connections established (i.e. SYN's per second, not full connections where there is an actual transfer of data) in order to boost metrics.

In short, the test configurations were designed to help users with their capacity planning requirements, not win benchmarking world records, so the results recorded are meaningful in the real-world.

For each of the test areas we focused on a number of metrics such as connections per second, requests per second, transactions per second, throughput, maximum concurrent connections per second and latency. Obviously, there are some specifics for certain tests, such as compression and SSL, that are not globally applicable, but in the main report we tried to keep the test types and metrics recorded as consistent as possible.

The latency metric is critical, and is the most overlooked element in performance testing both from a vendor and a customer perspective. Vendors will typically make maximum performance claims without mentioning the cost of that performance in terms of latency. If they mention that they can do 8,000 new SSL sessions per second but leave out the fact that the latency cost per session in achieving this number is 2 seconds or higher, is this acceptable? For many users the latency performance hit is not worth the benefit of a high number of new sessions per second. Hence the need to measure such elements as Time To Last Byte (TTLB).

As stated earlier, in all tests, the throughput is measured in Ixia client applications bytes, not raw L2 throughput. It's common for vendors to show the throughput as measured at the network layer, which adds approximately 52 bytes to every packet, leading to an overall 8 - 15% inflation in the throughput numbers. In these tests, we are only measuring client applications bytes (L7 data), and multiplying this number by 8 to generate the commonly used bits per second measurement used for all networking equipment. What this means to an end-user is that a test which shows 2 gigabits per second, is actually showing 2 gigabits per second of L7 data, which is actually closer 2.3 gigabit per second in raw network bandwidth. Or 3.5 gigabits per second, is actually line-rate 4 gigabit per second.

Each test suite ran for over seven hours, and each device under test recorded several weeks of actual test time. This enabled us to evaluate the products in a real-world, day-in day-out usage pattern – a far more reliable assessment than simply running a one-off, five-minute test against a device and claiming that to be a worthwhile exercise. The history of network product testing is littered with examples of this kind. There was no manual intervention during the tests. Just as in the real world, there was nobody “on hand” to toggle parameters at specific moments that might optimise a device for one type of traffic, just for a given period of time. So it was a golden rule that the exact same configuration must be usable throughout all tests.

The L4 Tests

While the marketers like to talk about the “sexy” L7 stuff that’s going on, the reality is that there are still a huge number of vital applications being run at L4. Hence, we included a series of tests at this layer to give a good baseline performance metric for real-world L4 applications.

The L7 Tests

Layer 7 is where the real intelligence of these devices sits, and intelligent rule making inevitably puts the L7 device under more stress. So it is an excellent test of product design, reliability and scalability. With more and more use being made of L7 applications now, this test is more important than ever and will be even more representative of day-in, day-out performance capabilities in the future. What makes this test L7 instead of just L4 is our use of a common rule which inspects the file type from the HTTP URI.

Compression Tests

In environments where bandwidth is limited – notably any Internet-related application – compression can be a very useful feature for accelerating performance. However, if the device which is doing the acceleration now becomes a bottleneck, compression can’t provide any value – that’s why these tests are designed to show the maximum compression throughput that a given device can achieve.

SSL Tests

The amount of HTTPS (SSL) traffic now on the Internet is many times that of a couple of years ago. In fact, many users are adopting the position that encryption should be utilized everywhere, even inside the corporate network. As a result, SSL traffic is set to dramatically increase as more users and enterprises require a secure connection for credit card transactions, personal information, and similar applications. As more and more of these services appear on the Internet, more and more SSL transactions will occur. As a result, real-world SSL transaction performance sizing, whether terminating at the device under test, or then re-encrypting to send the request onto a target server, is increasingly vital, hence the tests included here.

Total System Performance Tests

As we outlined earlier, it is a true test of a device to enable all features – after all what are they there for if not to be used – and test multiple, concurrent services running against the device. Any design weaknesses, or under-performing components, such as CPU or memory, will be brutally exposed by this kind of test.

The L4 Test Results

At L4 we chose to run a series of tests measuring connections per second, throughput and response time. This gave us effective information on a device's maximum throughput, bandwidth handling capabilities (and where bottlenecks in the system are limiting these) as well as the overall user experience. These tests were repeated across a range of file requests sizes, from 128bytes to 512KB. Keep in mind that while this information is useful in sizing, the test performed was done as a single function, where a traffic direction and load balancing decision was made based on an IP:Port combination, transferring real data across a variety of request / response sizes. Very few customers run just L4 operations alone, but we start here to build up a complete picture of total system performance.

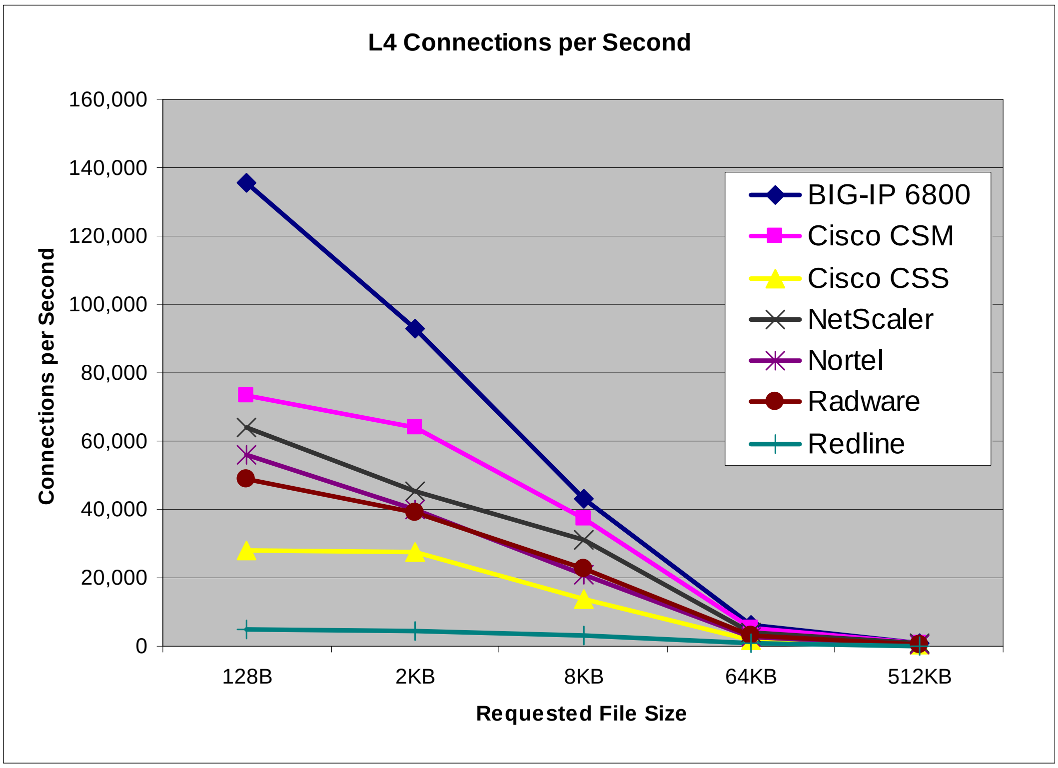

Figure 3 – L4 Connections per Second

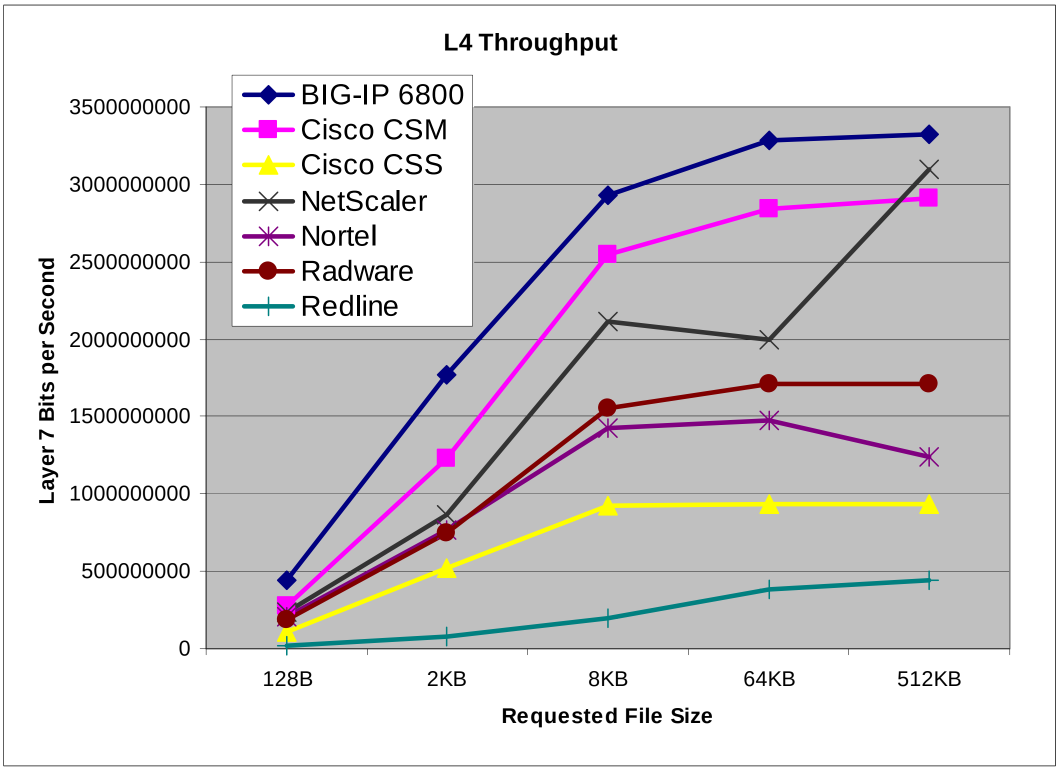

Figure 4 – L4 Throughput

For both connections per second and throughput, we see a consistent set of results. The BIG-IP 6800 proved fastest in both tests, peaking at 135,721 connections per second (128byte file size) while the next best was the Cisco CSM, at 73,439.

Throughput results were closer matched at the top end – where the Cisco CSM came close at nearly 70% of the BIG-IP device's capacity, but the BIG-IP device still topped the charts at all file request sizes, peaking at a total throughput of 3.52Gbps (512KB).

The implication here is that the BIG-IP product has less internal bottlenecks in its architecture, being able to perform well at all file size request points, and maintain that level of performance regardless of the metric being measured. This assessment is further supported by the response time test shown next. At smaller file request sizes, the BIG-IP system response time is effectively zero (as far as the user is concerned), while even at the 512KB file size it never rises about 239ms, still effectively instantaneous as far as a user is concerned.

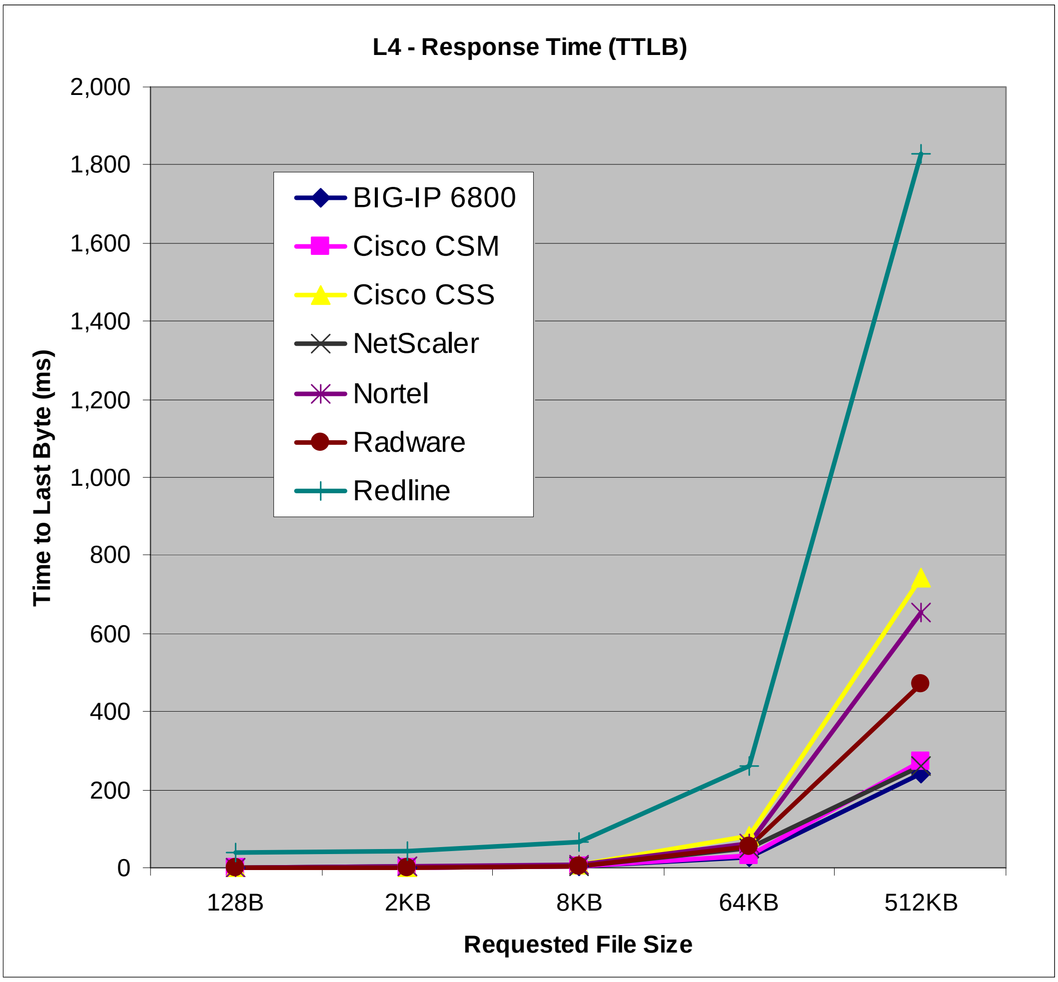

Figure 5 – L4 Response Time

The L7 Test Results

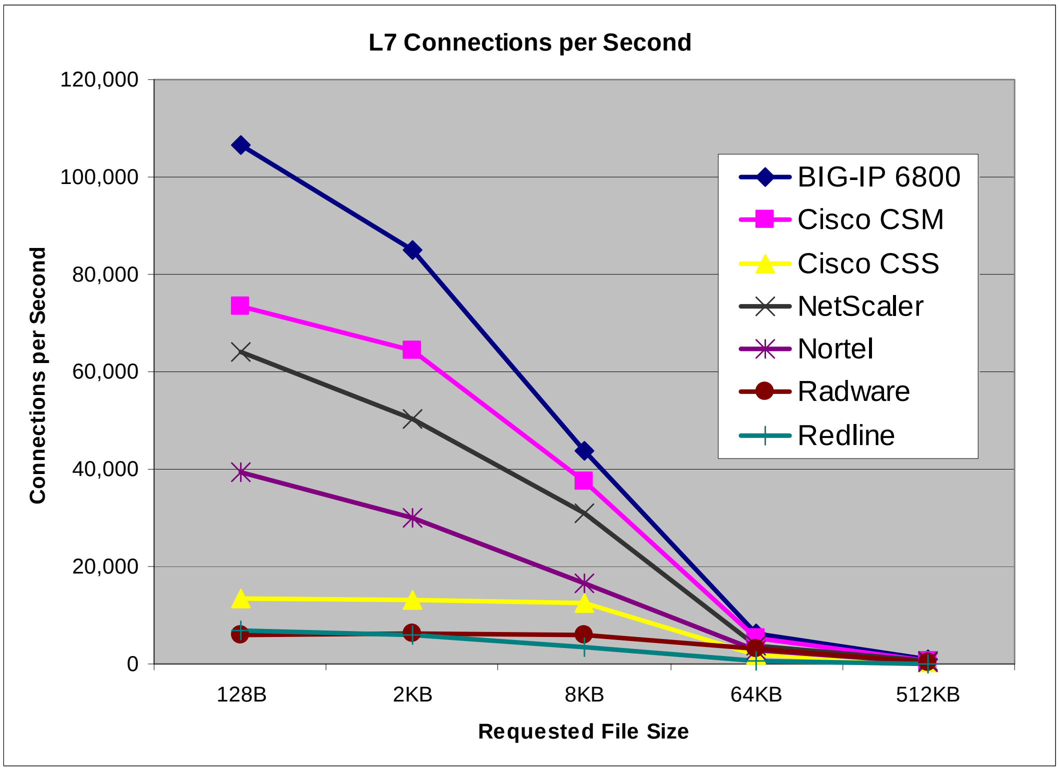

Figure 6 – L7 Connections Per Second

For L7, we chose the same set of metrics as for the L4 tests. Again, the BIG-IP 6800 emerged comfortably ahead on all tests, with the Cisco CSM and NetScaler 9950 performing as the best of the rest at 69% and 60% of the BIG-IP device's connection rate respectively, and Redline, Radware, and Nortel performing at a fraction of that.

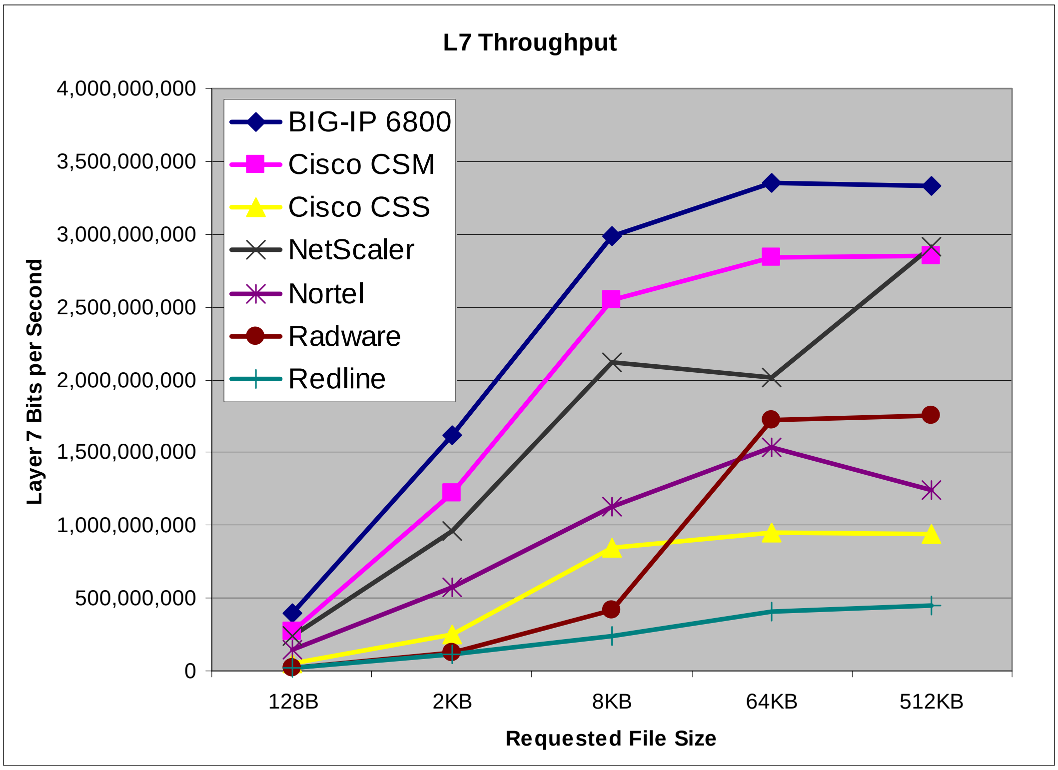

Figure 7 – L7 Throughput

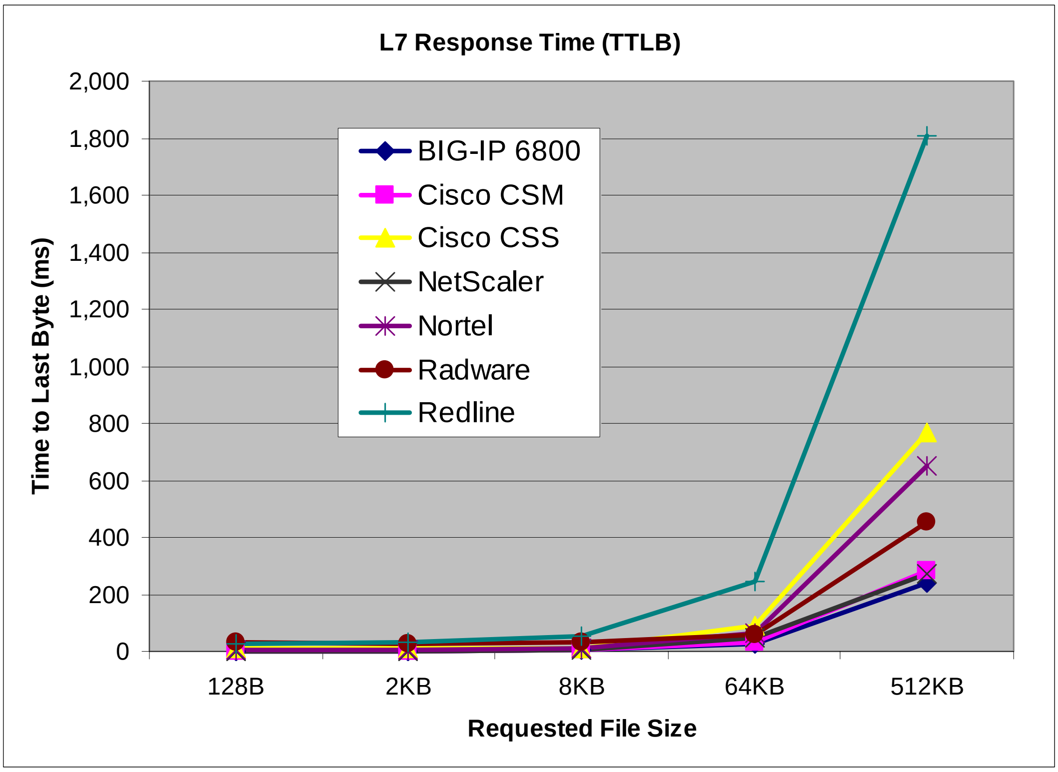

Figure 8 – L7 Response Time

The BIG-IP device's peak connection per second rate of 106,497 shows that the product's performance holds up well moving from L4 to L7, while throughput and response time results were actually slight improvement on the device's L4 performance. This augurs well for using the BIG-IP product in a complex L7 environment with rules to contend with; meaning performance and intelligence can go hand in hand.

The Compression Test Results

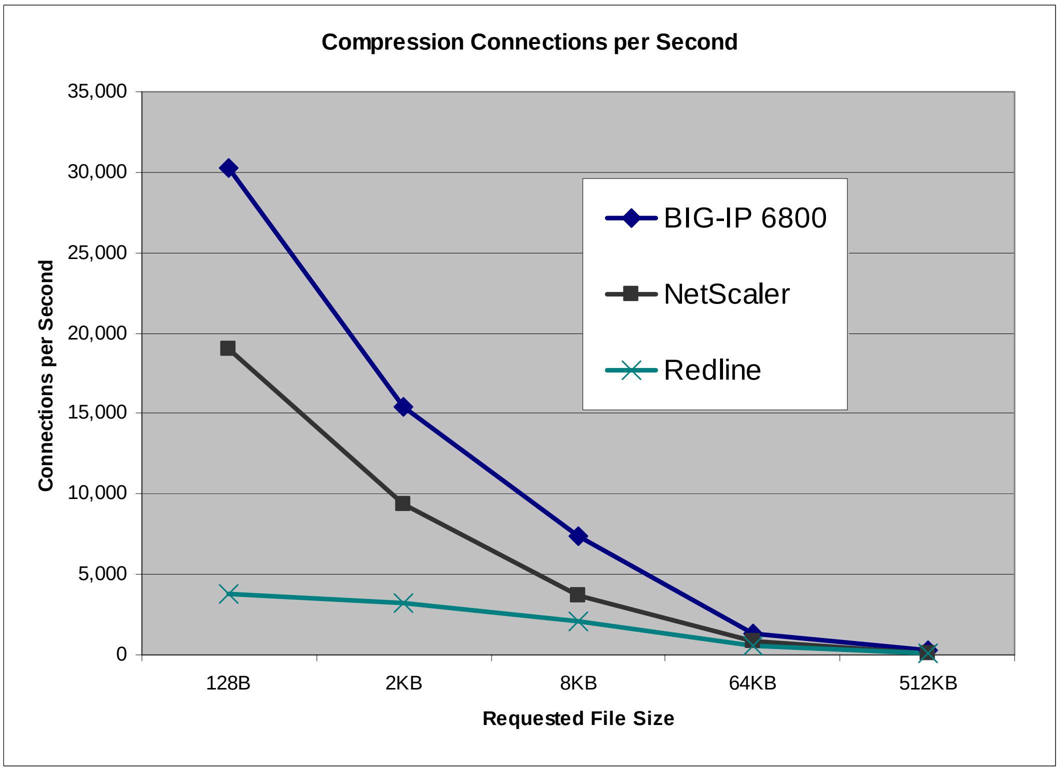

Figure 9 – Compression CPS Test

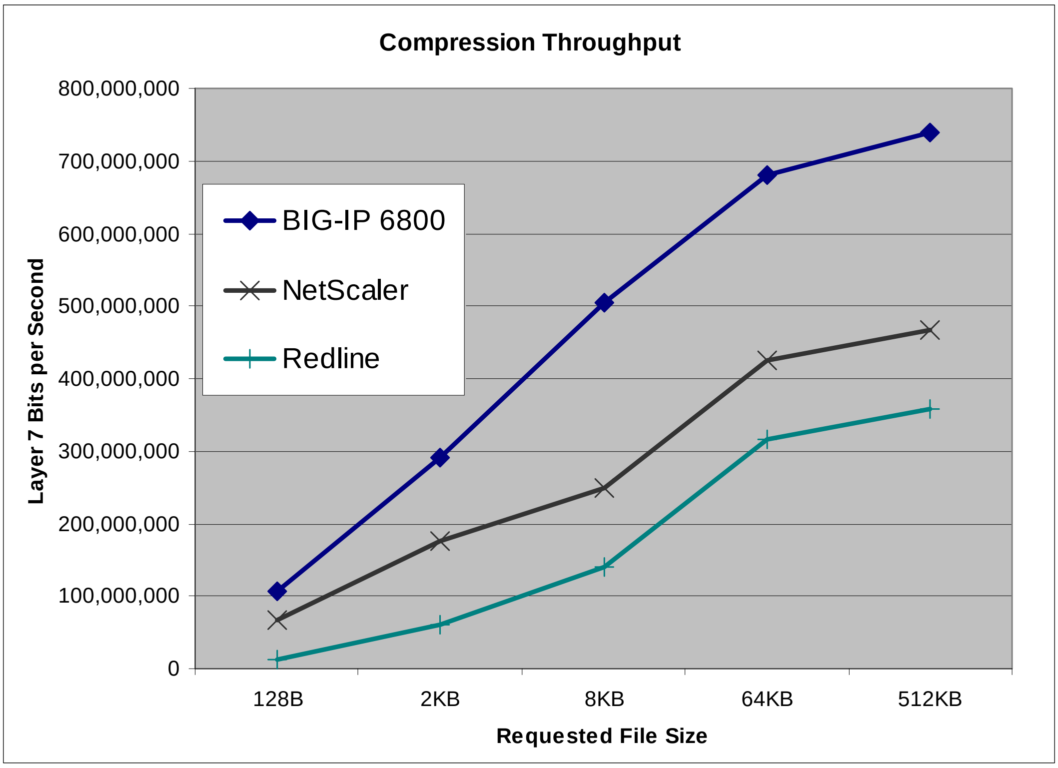

Figure 10 – Compression Throughput Test

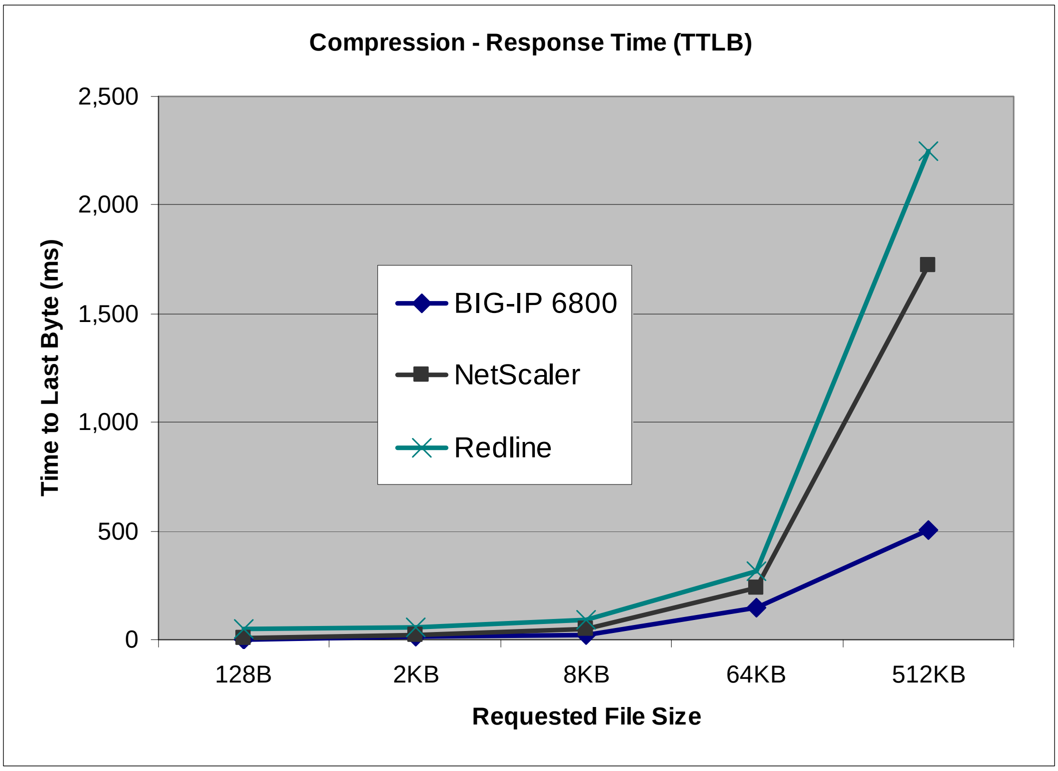

Figure 11 – Compression Response Time

Only three of the devices under test actually support compression – those from F5, NetScaler and Redline. In both the CpS and throughput tests, the

BIG-IP device emerges as a very clear leader. Peaking at 30,274 CpS (128bytes), the next closest is the NetScaler at 18,977.

In the throughput test, the BIG-IP product is even more dominant, topping out at well over four times the performance of its nearest rival.

The margin is similarly impressive in the response times test, where the BIG-IP device was between two and ten times better that its competition.

It is important to note that if you are compressing larger page sizes, (those greater than 16KB, which is a very common scenario for many web applications), you and your users may notice a performance hit depending upon the vendor solution in terms of throughput and response times. This can eliminate the benefit of compression in the first place where the implementation is about improving performance and reducing bandwidth. What's dramatic here is understanding how this individual performance test is affected with other functions enabled as well like SSL and L7 decisions.

The SSL Test Results

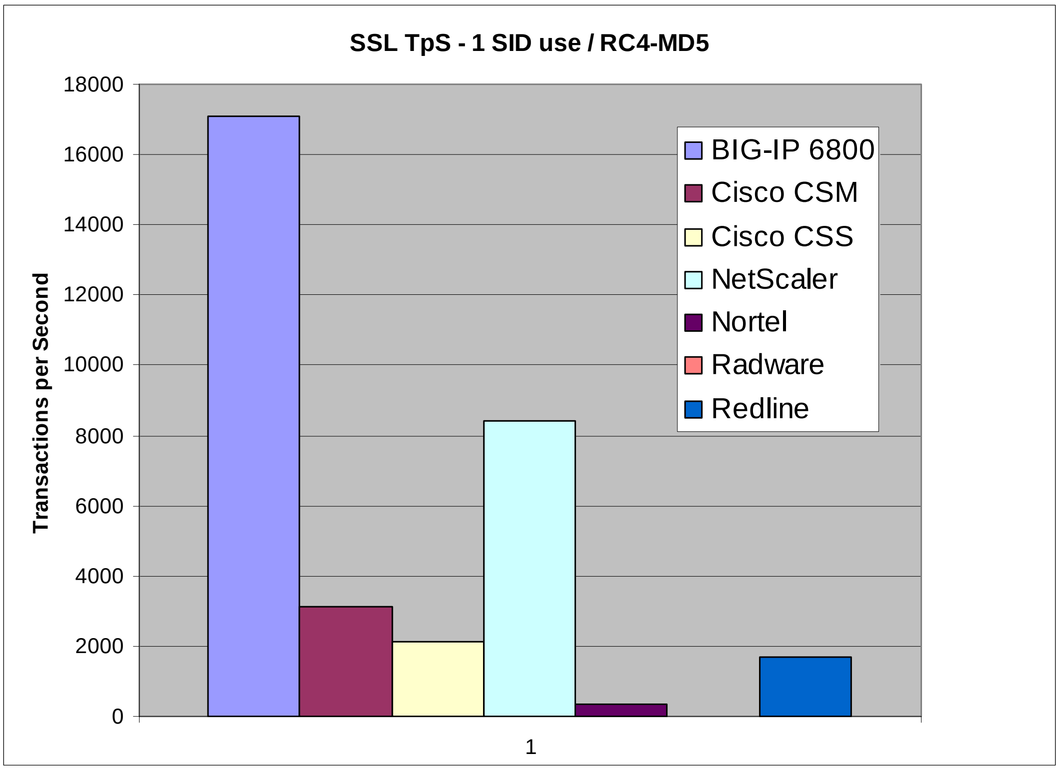

Figure 12 – SSL TPS – 1 Session ID Use – RC4-MD5 Cipher

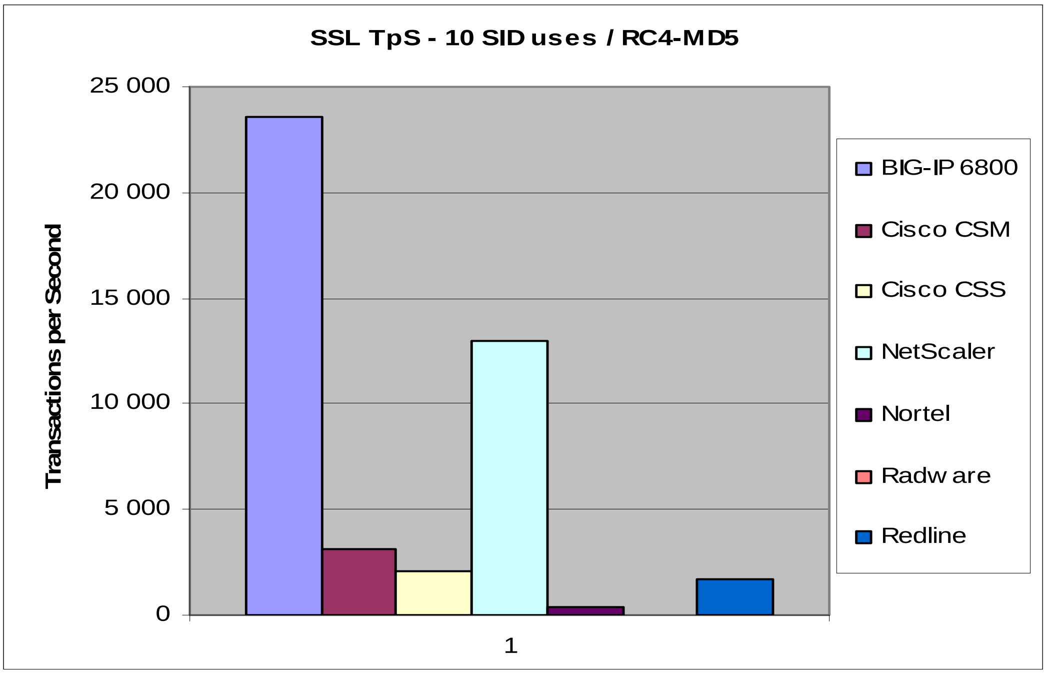

Figure 13 – SSL TPS – 10 Session ID Use – RC4-MD5 Cipher

The SSL tests – both for the 1 Session ID use and 10 Session ID use configurations – gave us record figures for the BIG-IP device, more than double anything that Broadband-Testing has seen, or indeed what has been published to date. With the RC4-MD5 cipher, the BIG-IP device peaked at 17,093 TpS with a single SID and a colossal 23,533 TpS with 10 SID users. Figures for the DES-CBC3-SHA cipher tests were only marginally behind.

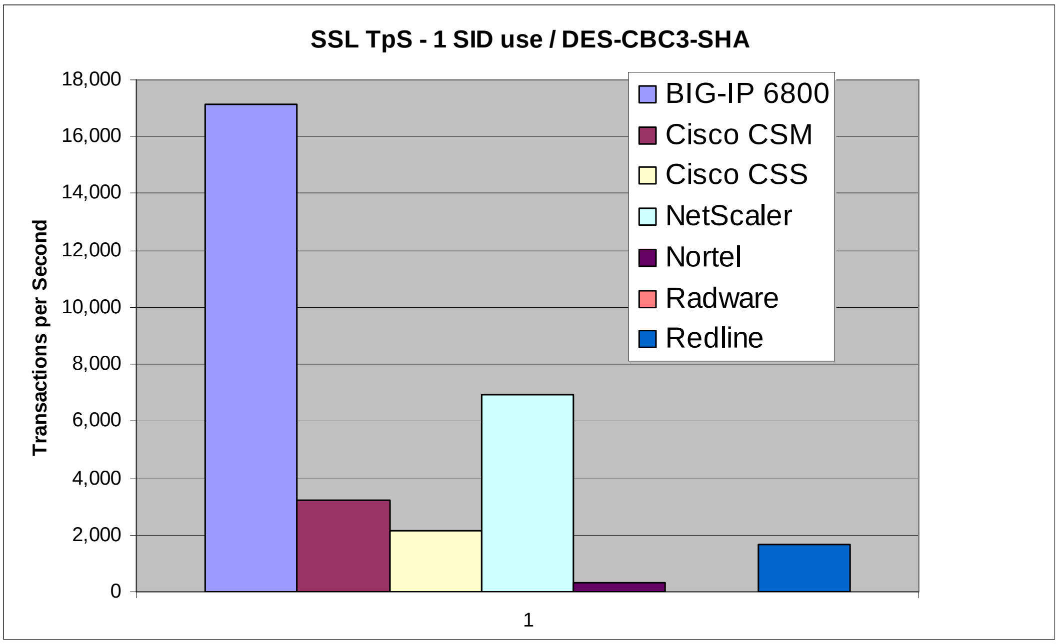

Figure 14 – SSL TPS – 1 Session ID Use – DES-CBC3-SHA Cipher

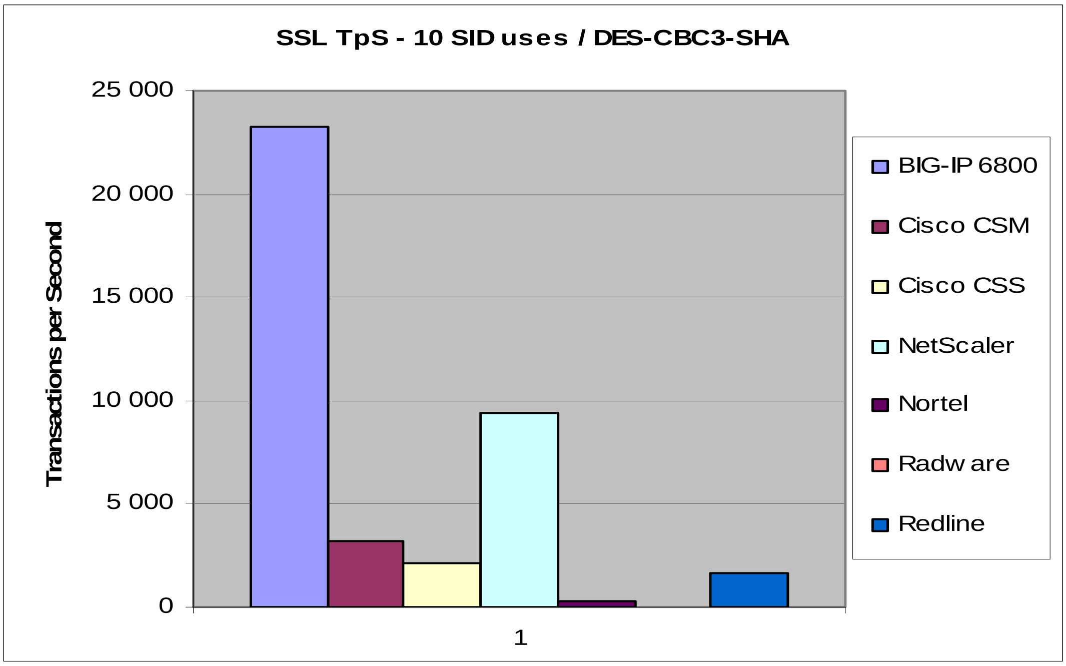

Figure 15 – SSL TPS – 10 Session ID Use – DES-CBC3-SHA Cipher

The Total System Performance Test Results

Because only the NetScaler and Redline products had the full feature set of the F5 BIG-IP device, just these representatives were valid for the "combo" test, with all features set to "on". This test consists of 33% compressed mixed traffic, 33% uncompressed mixed traffic, and 33% L4 traffic. "Mixed" traffic is 50% L7 SSL, and 50% L7 (no SSL).

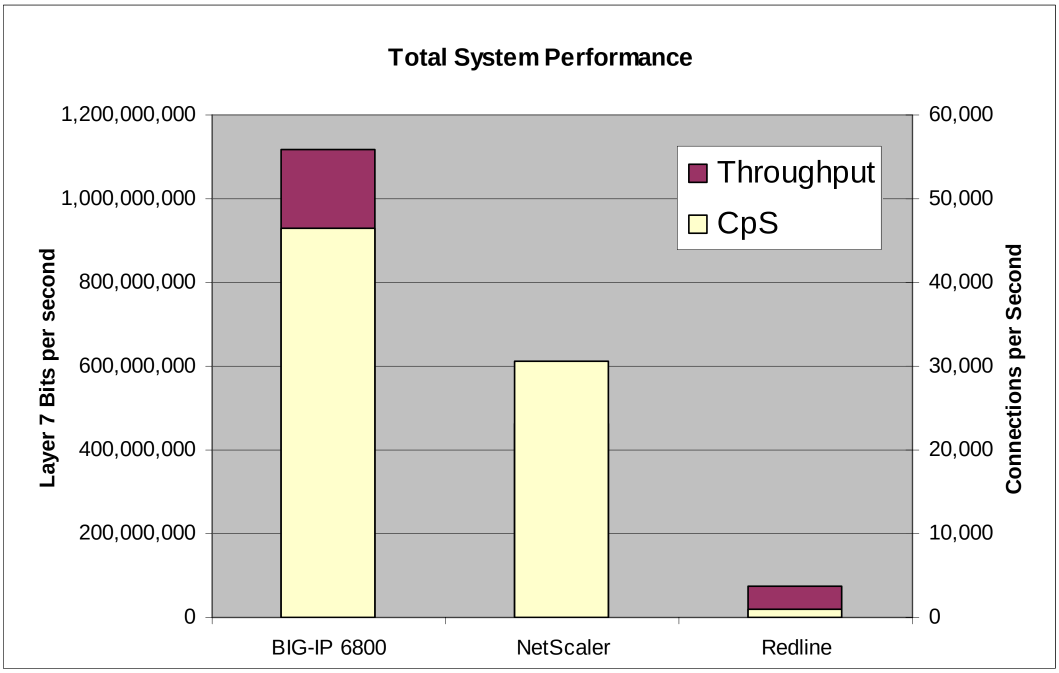

Figure 16 – Total System Performance Test

Performance was measured in both L7 bits per second recorded and connections per second. In both cases the F5 came out clearly ahead of the rest with NetScaler achieving approximately 40% of the BIG-IP device's total throughput and approximately 60% of BIG-IP's connection rate.

Given that, in the real world, a network manager simply wants to set the device to "on" for all features and forget about any other re-configuration, this is a realistic assessment of day-to-day performance expectations.

The DDoS Test Summarised

In addition to the pure, performance-oriented tests, a Distributed Denial of Service (DDoS) test was also run using Spirent's WebAvalanche to generate the simulated DDoS attack, a SYN flood.

We measured how the device maintained performance – an L7 http session with a 128byte page request size – under attack, with DDoS prevention functions enabled on all devices. Again, the BIG-IP device emerges well ahead of the competition

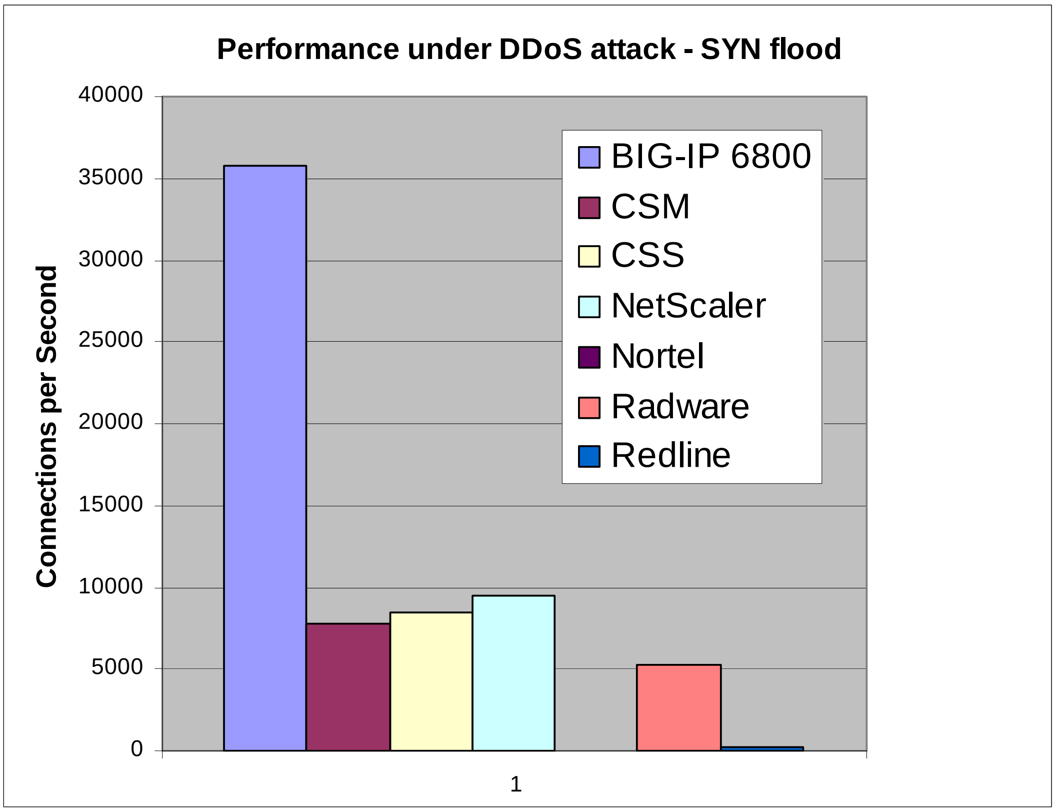

Figure 17 – DDoS Attack Test

SUMMARY OF TEST RESULTS

It is easy here to simply say, “Let the results speak for themselves” – which show the F5 BIG-IP 6800 is currently well ahead of the competition. While network administrators need a product that scales and runs all these functions without adding excessive latency, it is also important that the solution is smart. After all, application intelligence is what you’re expecting from this class of products. Performance is only part of the picture – but an important consideration for most companies that expect growth.

However, it is important to re-iterate just why the BIG-IP device’s performance is, in most cases, so significantly ahead of its rivals here. Development of the all-new BIG-IP platform with the v9 code release was not an overnight achievement, but easily the largest development project undertaken by F5.

As indeed was the testing carried out on the BIG-IP product and on the other L4-7 products in this test, as validated by Broadband-Testing. Testing of this scale has never been accomplished before. This is truly a first-of-its-kind type of event. F5 has invested an enormous amount of time and money to be the first vendor to run performance tests in as close to real world as possible, with scenarios covering a huge range of applications, all at peak utilisation of each device under test. The performance results obtained for many of the devices under test, in many of the tests achieved or even exceeded the vendors’ own performance claims, which very concretely validates the testing from our own perspective at Broadband-Testing.

We have therefore no hesitation in heartily recommending the F5 BIG-IP 6800 on the basis of its outstanding performance results in this test, that have really set new standards for the genre and which merit the gold award in this case – a first for F5 Networks.

Broadband Testing Gold Award

TEST SECTION 2: THE OTHER TEST RESULTS

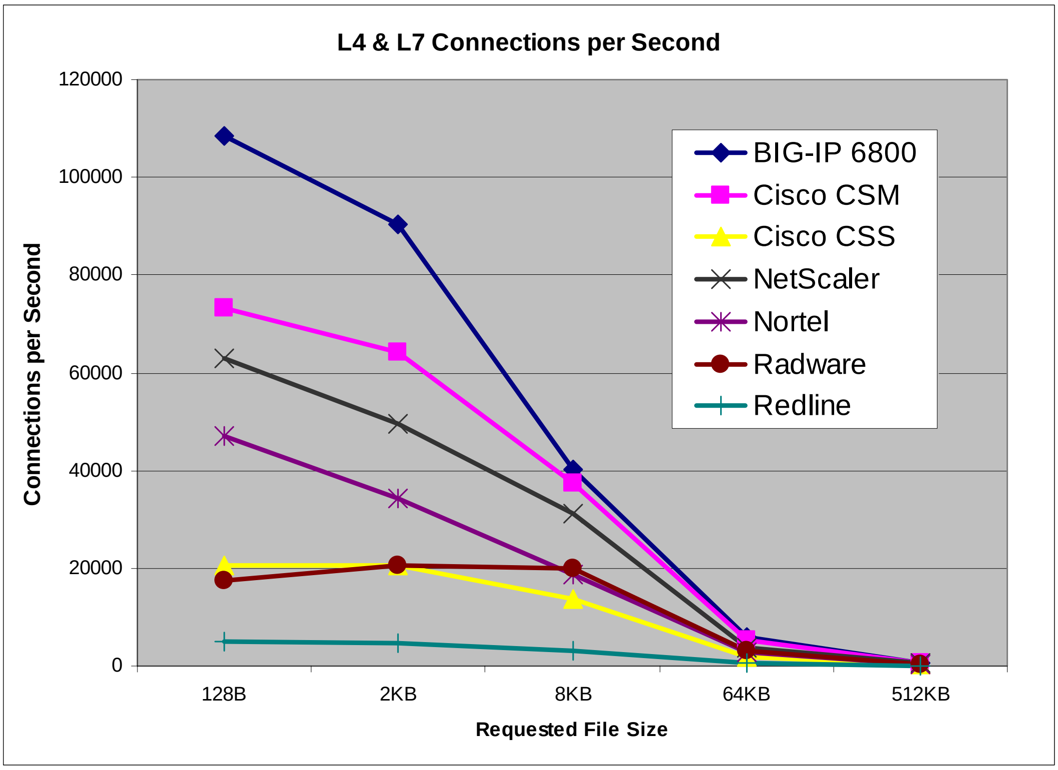

Figure 18 – L4+7 CpS Test

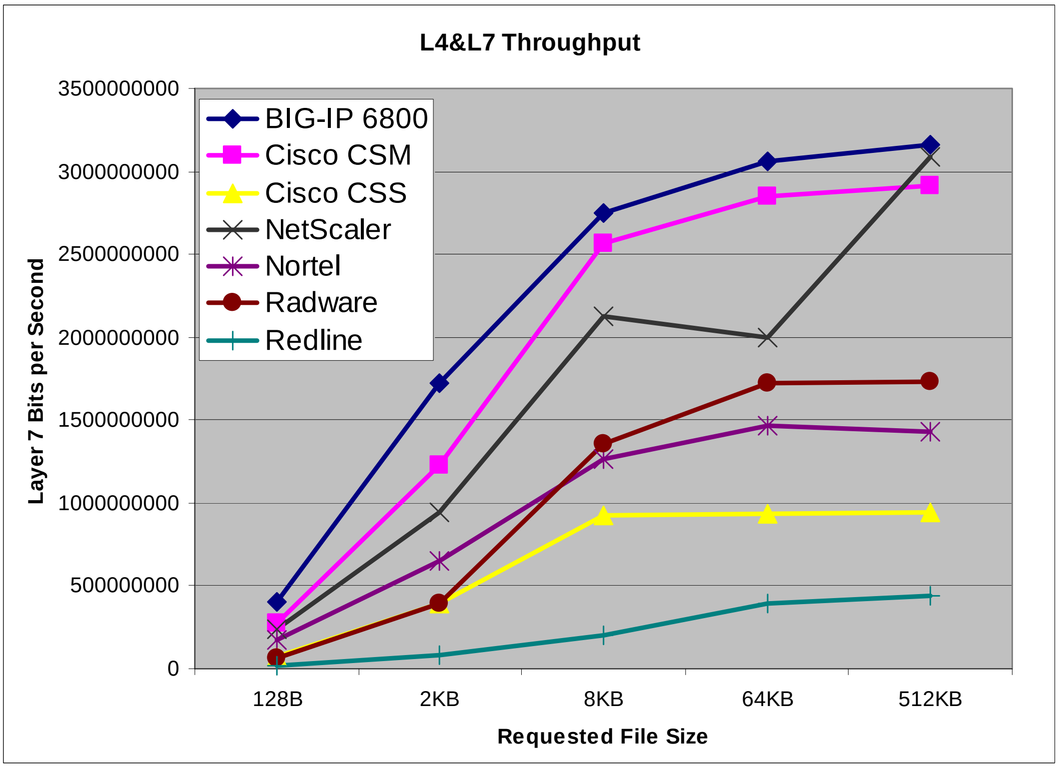

Figure 19 – L4+7 Throughput Test

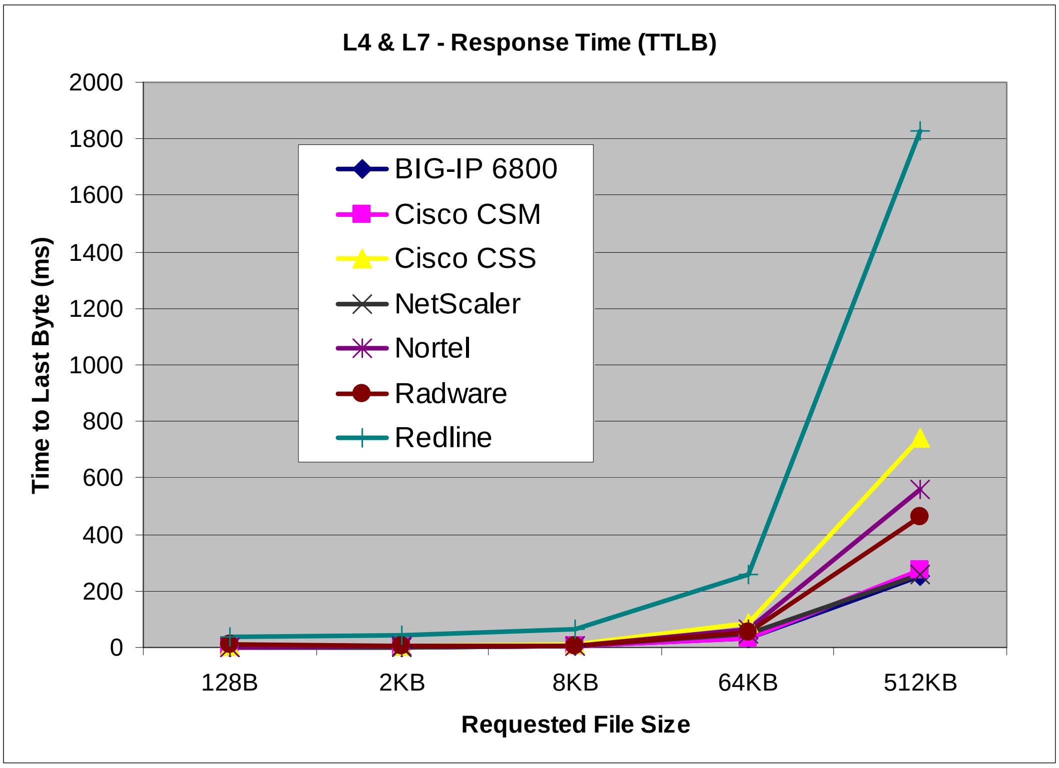

Figure 20 – L4+7 Response Time

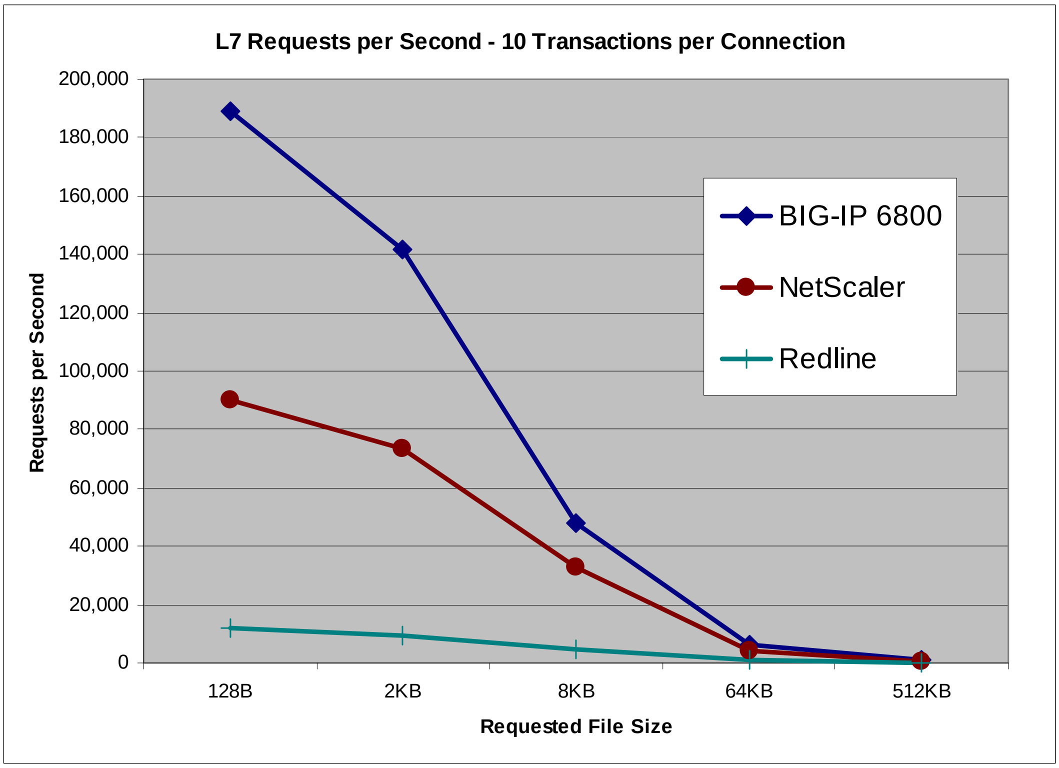

Figure 21 – L7 Requests Per Second – 10 Transactions Per Connection

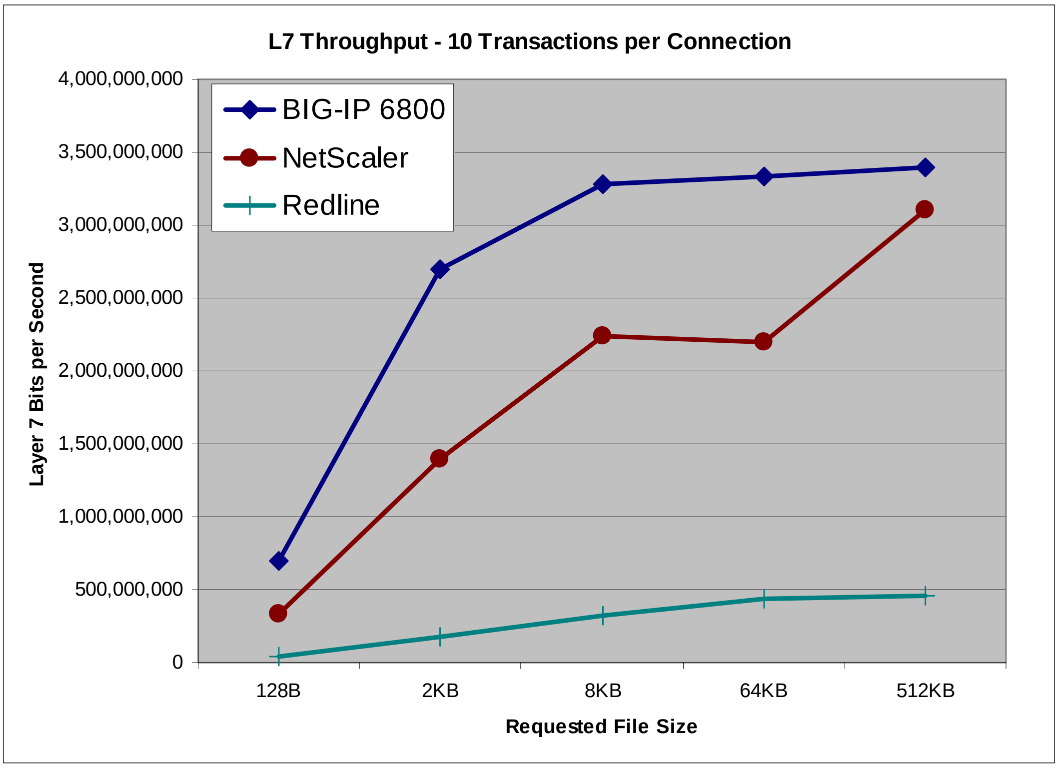

Figure 22 – L7 Throughput – 10 Transactions Per Connection

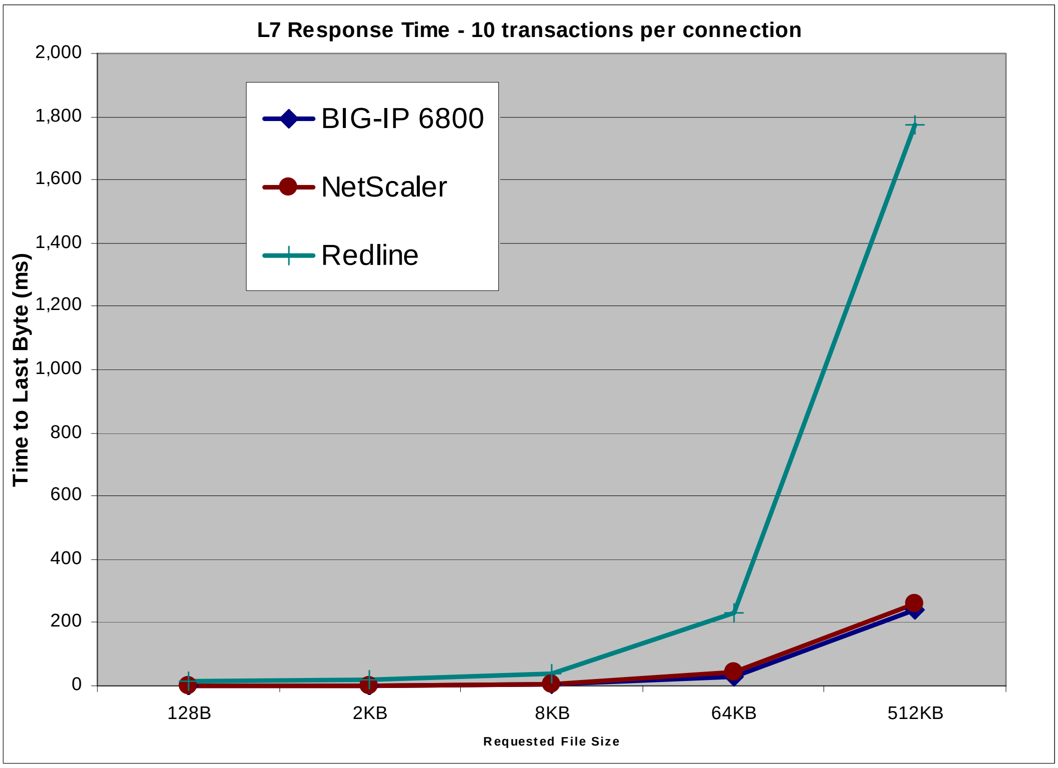

Figure 23 – L7 Response Time – 10 Transactions Per Connection

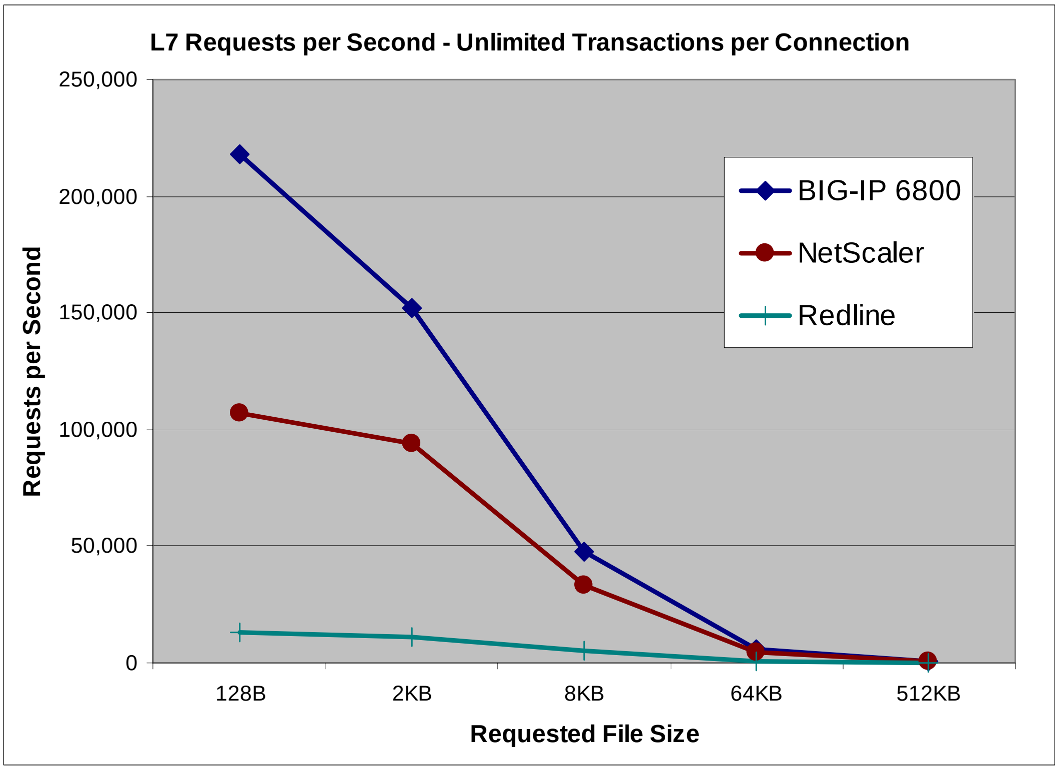

Figure 24 – L7 Requests Per Second – Unlimited Transactions Per Connection

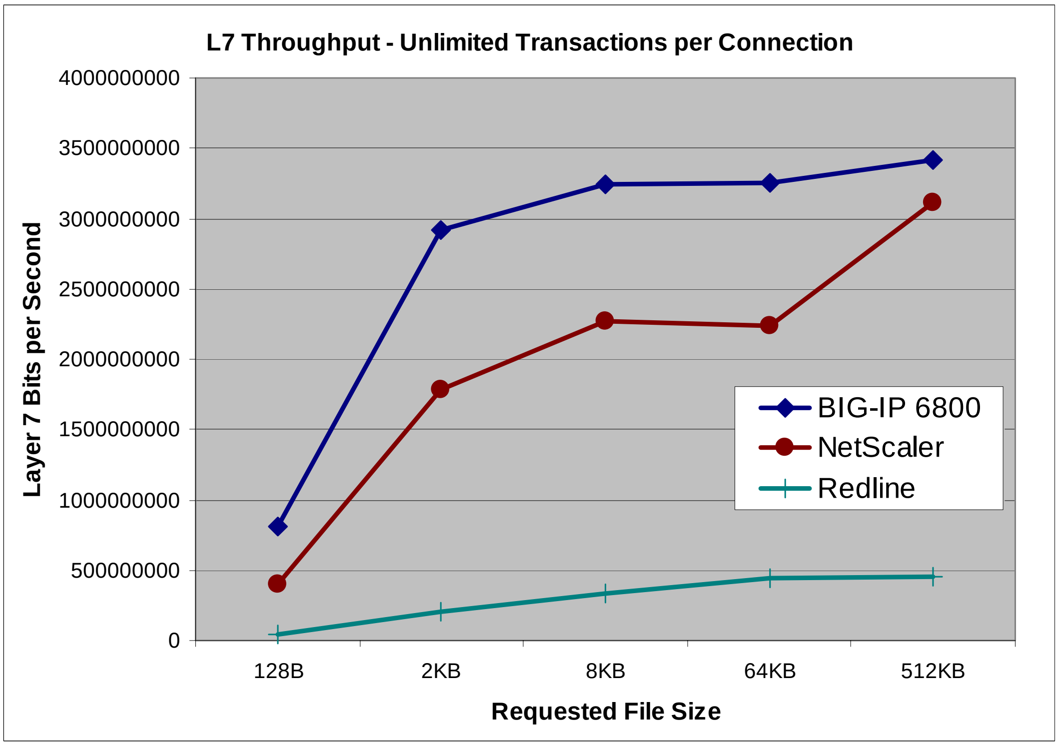

Figure 25 – L7 Throughput – Unlimited Transactions Per Connection

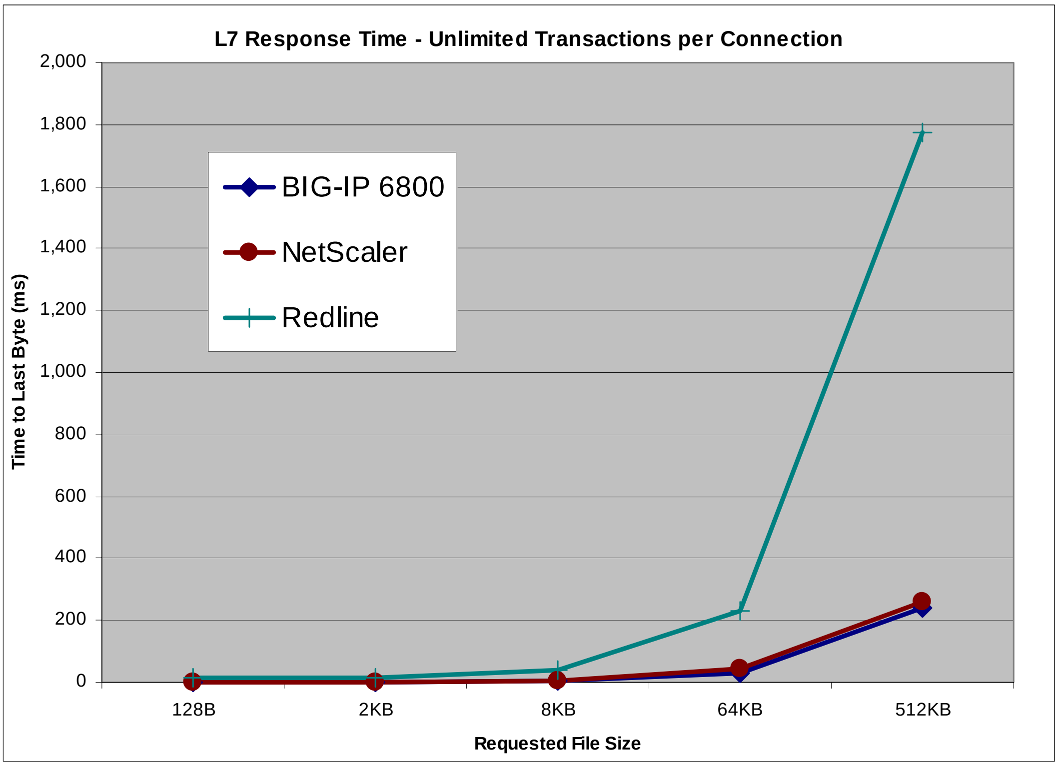

Figure 26 – L7 Response Time – Unlimited Transactions Per Connection

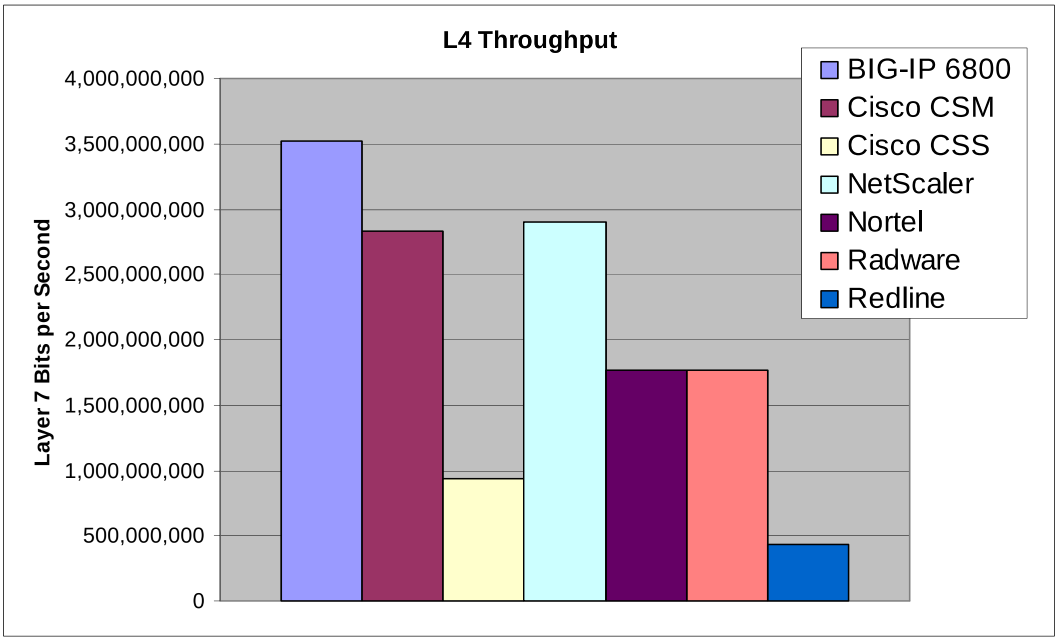

Figure 27 – L4 Throughput

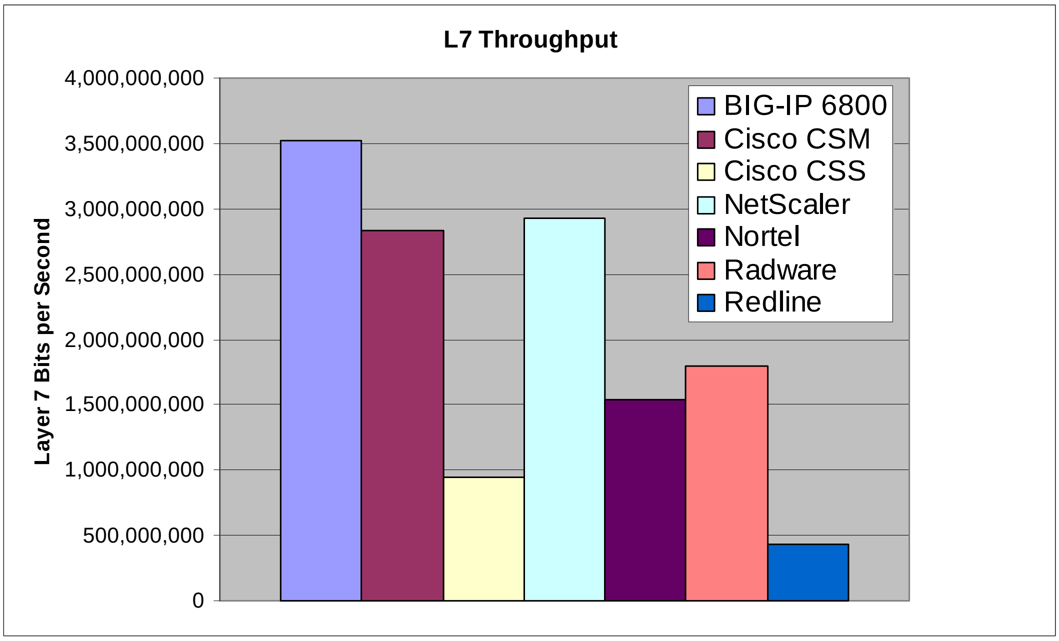

Figure 28 – L7 Throughput

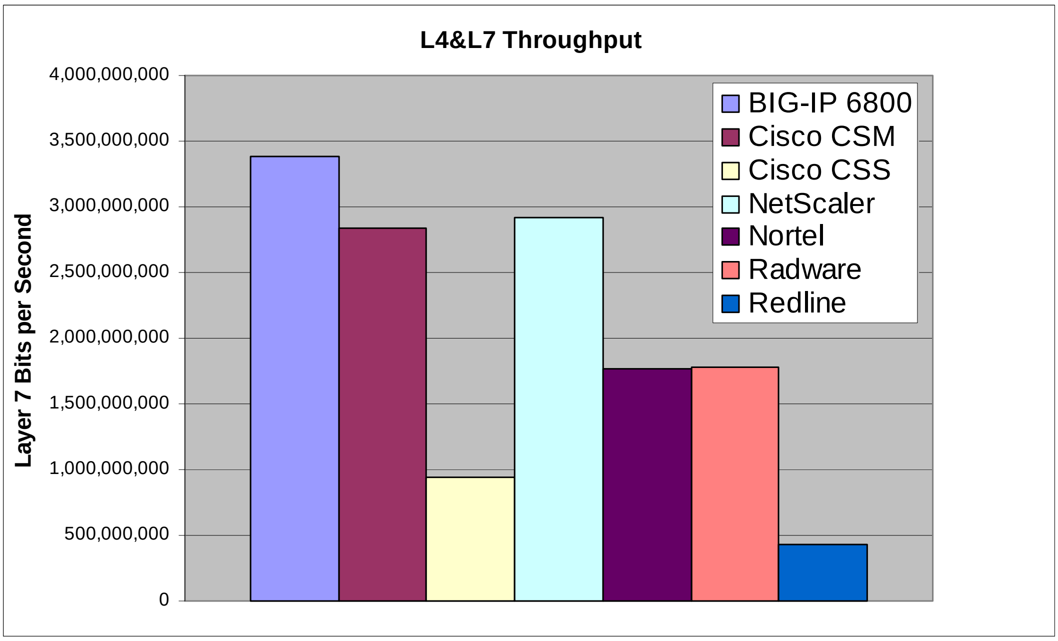

Figure 29 – L4+7 Throughput

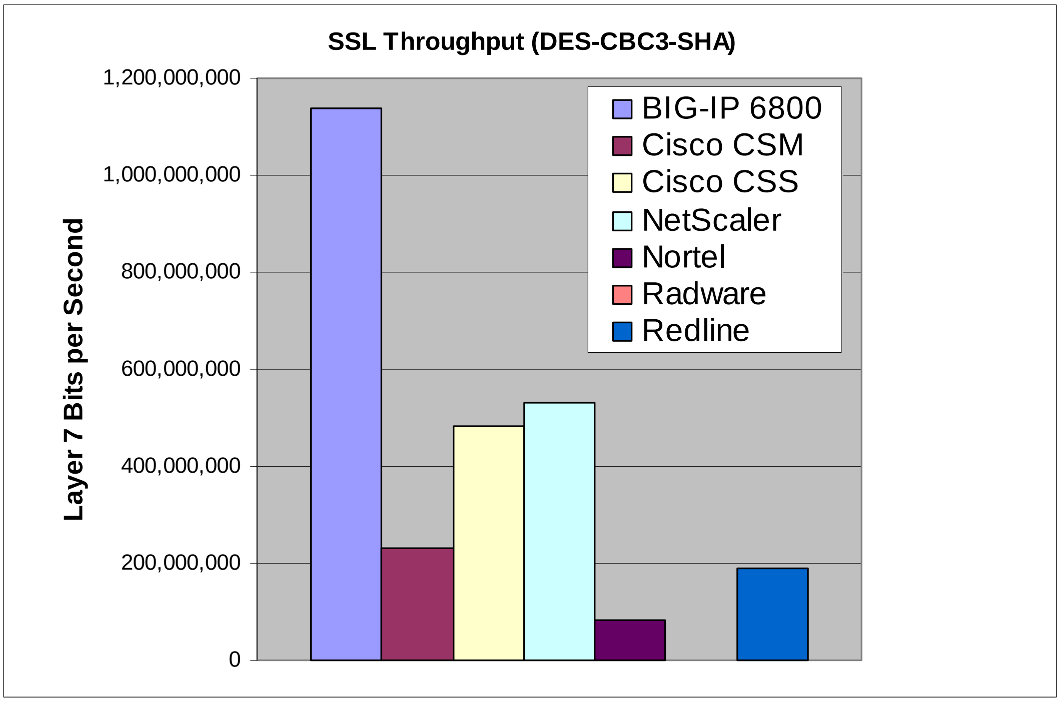

Figure 30 – SSL Throughput (DES-CBC3-SHA)

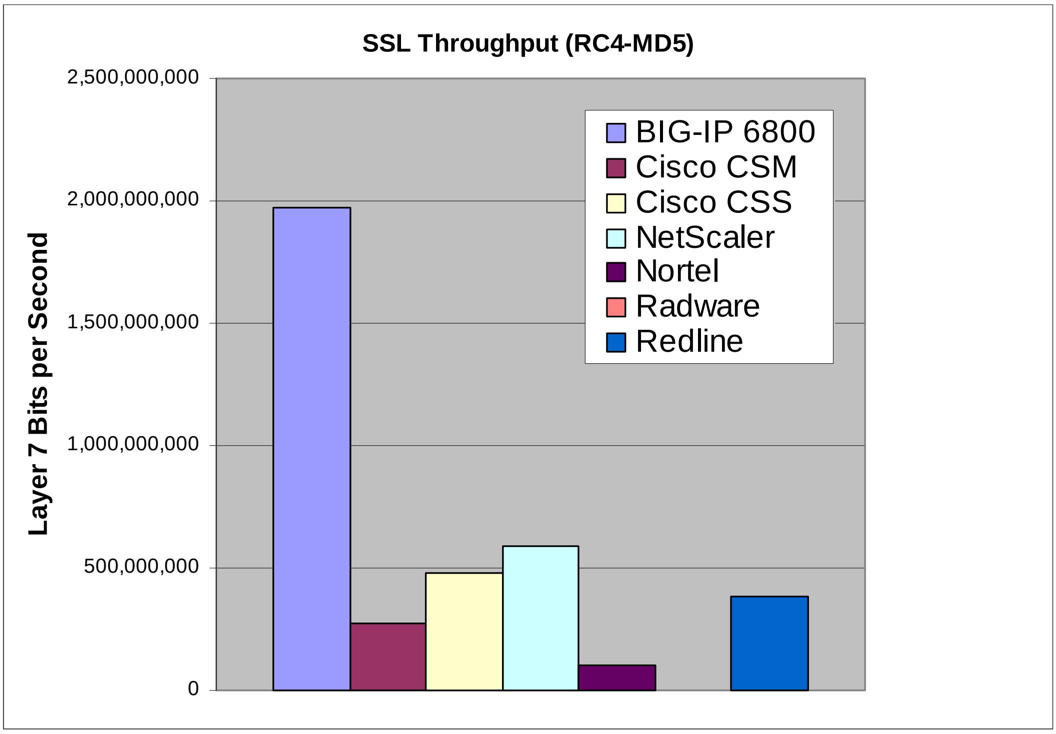

Figure 31 – SSL Throughput (RC4-MD5)

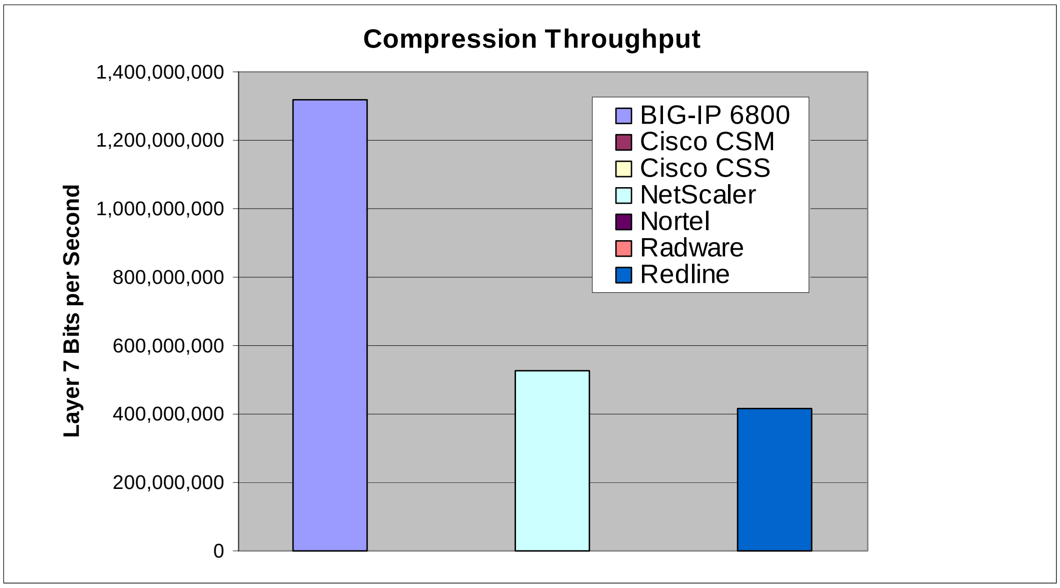

Figure 32 – Compression Throughput

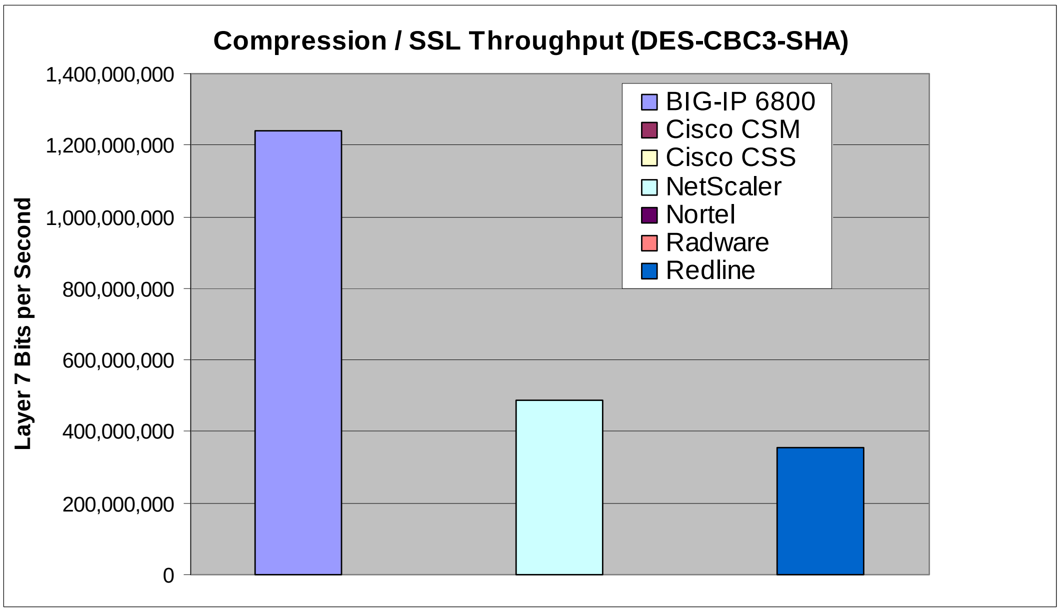

Figure 33 – Compression/SSL Throughput (DES-CBC3-SHA)

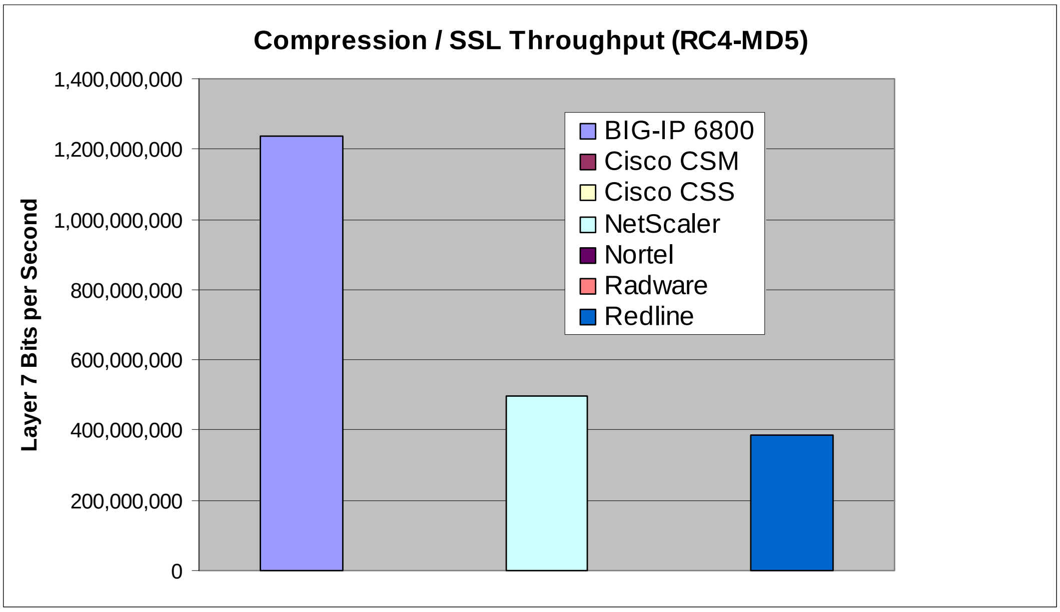

Figure 34 – Compression/SSL Throughput (RC4-MD5)

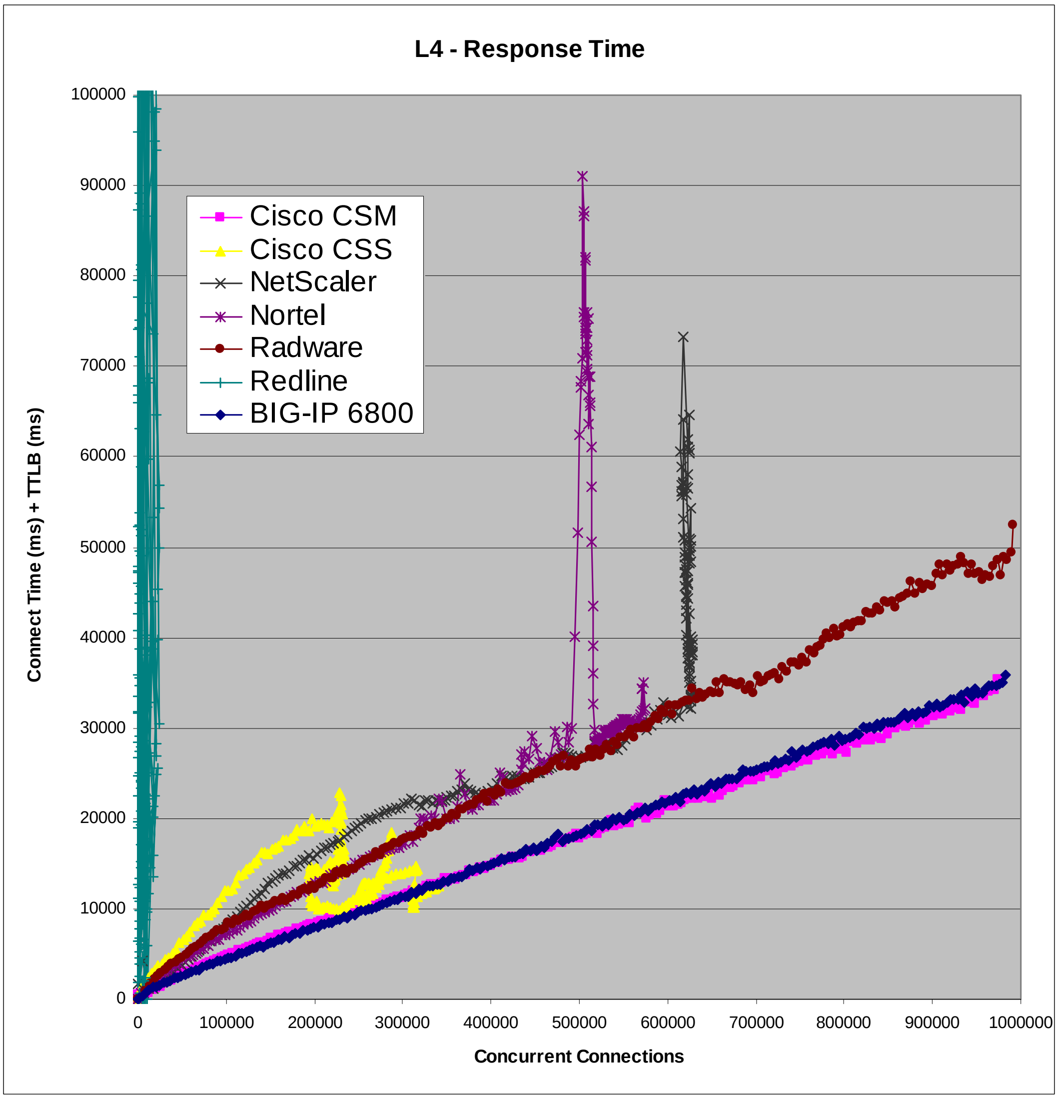

Figure 35 – Latency Test – L4 Response Time

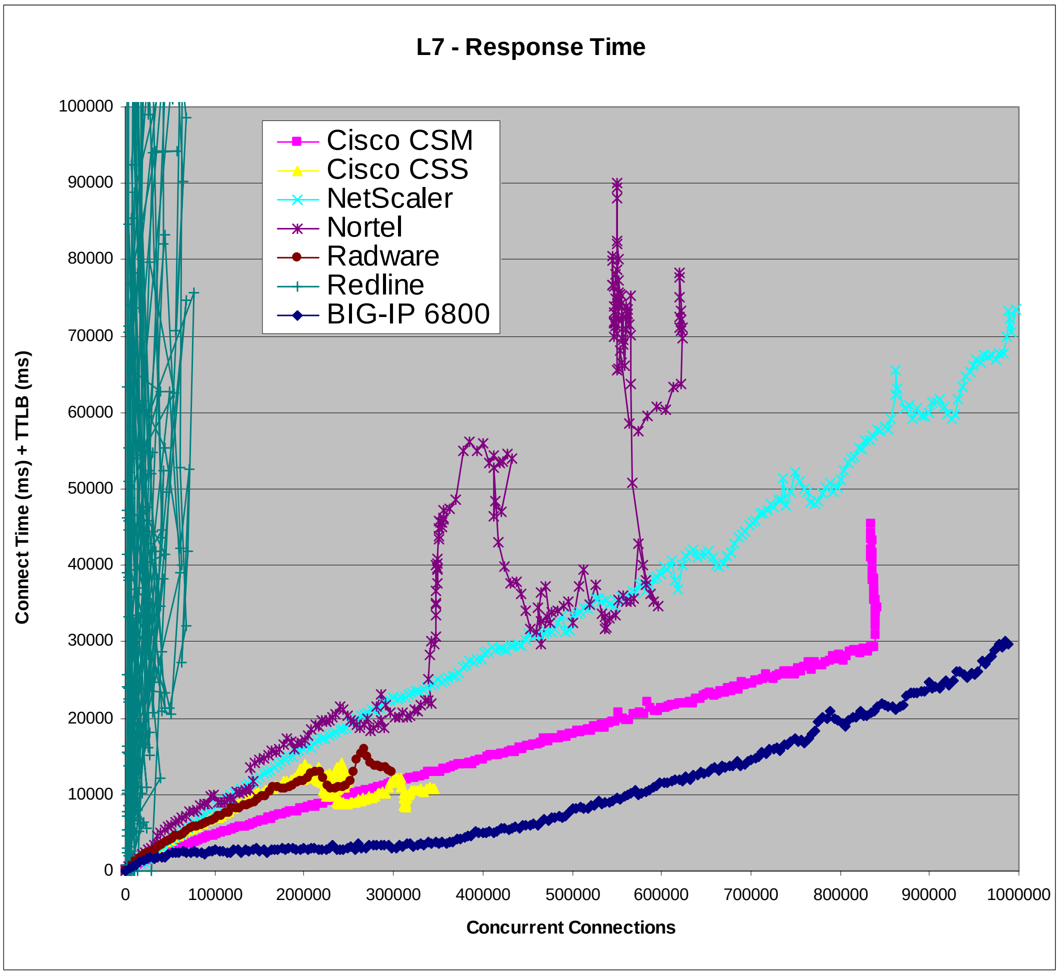

Figure 36 – Latency Test – L7 Response Time

Note: *In this test, both Radware and Cisco CSS can only maintain approximately 300,000 unique session table entries before other session table entries are expired, this is why their lines stop.*

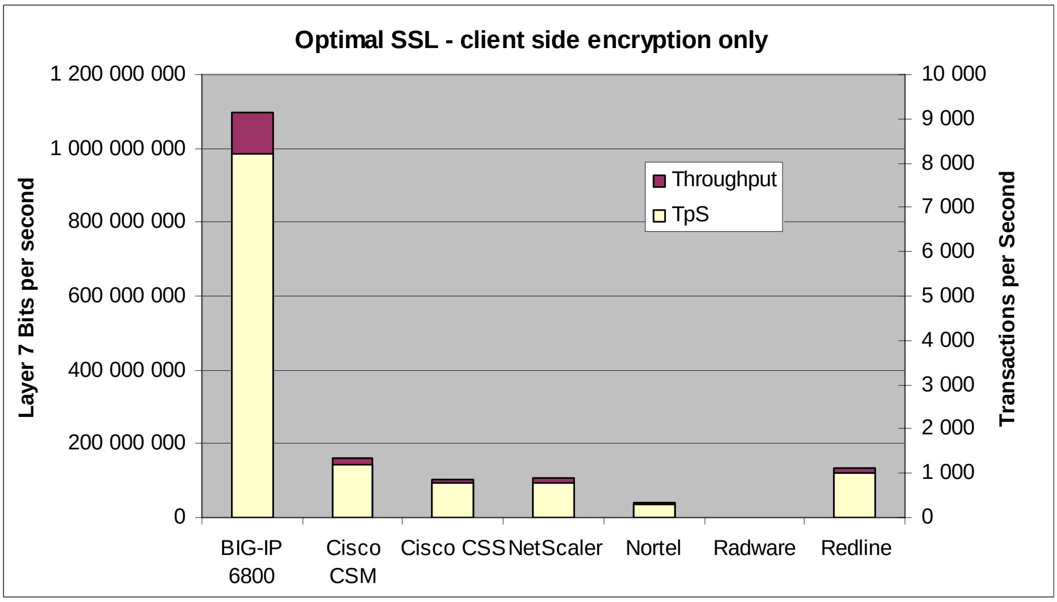

Figure 37 – Optimal SSL Performance – Client Side Encryption Only