

Comparitive Performance Report

Application Delivery Controllers - January 2014

aiScaler vs Citrix Netscaler vs F5 BIG-IP AAM, tested on AWS

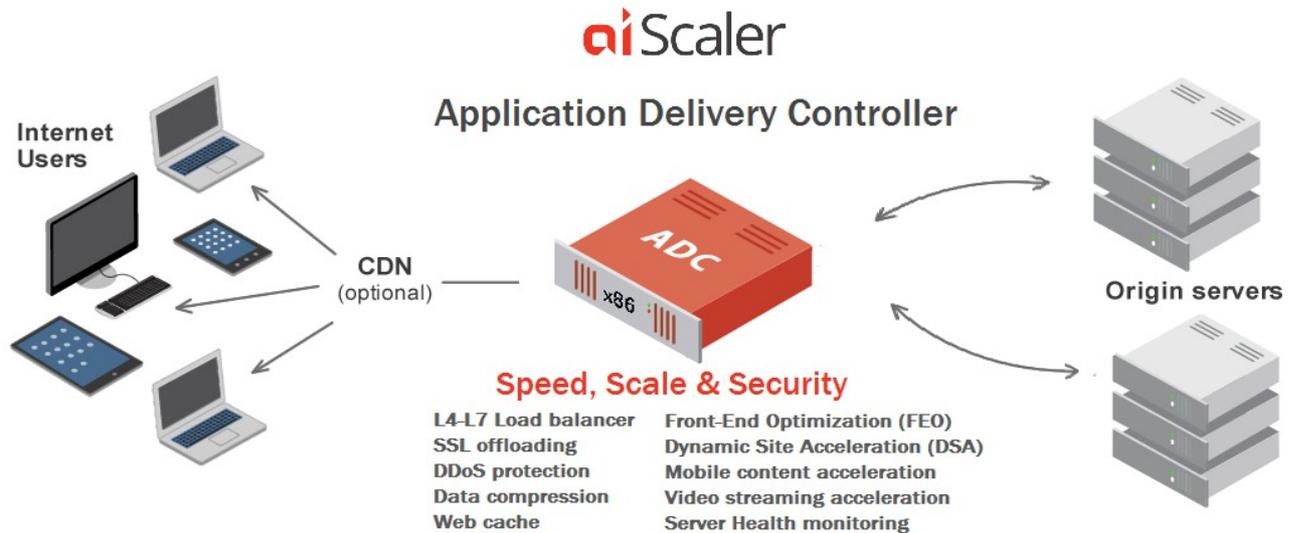
Content

1. Introduction.....	2
2. Features and pricing.....	3
3. Ease of installation.....	6
4. Performance testing.....	7
4.1 Apache Benchmark – Measuring RPS.....	7
4.2 Pingdom – Page loading times.....	8
4.3 Siege – Stress-test 1.....	9
4.4 Blitz.io – Stress-test 2.....	10
5. Conclusion.....	14



1. Introduction

Application Delivery Controllers are used by websites to scale and accelerate applications delivery. ADC's provide front end intelligence that simplifies responding to high high volume or complex customer requests. The ADC is placed between the web servers and CDN or public internet and generally replace the load balancer.



ADC's have historically been hardware-based providing only load balancing and simple application acceleration. Modern ADC's offer many more features such as Traffic management, web caching, Denial of Service protection, SSL offload, compression, Dynamic Site Acceleration (DSA), Front-End Optimization (FEO), mobile content acceleration and health monitoring.

ADC's are now available as software and can be installed on any x86 compatible environment, removing the need for proprietary hardware. These solutions can be downloaded and deployed in private data centers or rented by the hour from public cloud providers such as Amazon Web Services.

Using ADC's in distributed global locations, allows users to create private Content Delivery Networks, with control over edge locations, scaling, cache control, IP restrictions and DoS protection. This allows companies to provide globally available high speed sites without modifying their application.

Prior to the availability of virtual ADC's deployable on public clouds, you required expensive capital deployment or long term service contracts. These contracts provided by organizations like Akamai & Limelight have minimum use fees and termination penalties. Generic CDN services lack the ability to provide precise optimization, focused only on the specifics of your environment. The result is sub-par performance without direct insight or control of your customer experience.

Using public cloud services such as the AWS Marketplace, users can rent application acceleration software for as little as 19 cents per hour! Instances can be easily deployed worldwide, physically near your users. There are no minimum contracts and you control the configuration of your site. While CDN's decrease the latency of the internet and distribute static content, they do not increase the speed or capacity of the application. An ADC based network in conjunction with a CDN increases the speed, scale and stability of your site while reducing support and hosting costs.

We tested virtual ADC's from the top 3 popular vendors on the AWS Marketplace: aiScaler, Citrix and F5. The following report is the result of our testing.

2. Features and pricing

The first thing we notice: software is more expensive than hardware. On the AWS Marketplace, a customer pays an hourly fee to rent the software, and an hourly fee to Amazon for the computing power. From all tested products, the software fees are higher than the EC2 hardware fees.

We tested the Netscaler VPX Platinum edition, since it's the only solution in the VPX line offering clustering and “AppCache” which is Citrix's name for their web cache technology. The Marketplace states that their Standard and Enterprise edition also offer caching, but this seems to be a mistake, and is correctly specified on the Citrix website.

F5 licensing is built in a way that each component of your configuration has to be licensed. BIG-IP has 11 such components. To have a comparable configuration with caching, SSL termination and DDoS mitigation, you need at least the Local Traffic Manager (LTM) and Application Acceleration Manager (AAM).

	aiScaler		Citrix Netscaler VPX Platinum		F5 Big IP
Bandwidth	no software limit		10mbit license	200mbit license	200mbit license
Instance type	medium	xlarge	large	large	xlarge
Total hourly fee on AWS	\$0.41	\$0.79	\$1.404	\$2.314	N/a
Price of a dedicated license with dynamic caching	\$9,995		\$8,000	\$15,000	\$8,995 (AAM) - no DDoS mitigation
Surcharge SSL termination	\$2,000		included	included	\$7,995 (LTM) includes DDoS mitigation
Annual Support contract	\$7,995 per client		\$4,928 per license	\$9,240 per license	\$6,400-\$11,800 per license
Hourly support rate	\$175		N/a		\$375-\$475
Support included	2 hours		none		none
Minimum TCO based on: -acquisition of 2 production licenses in a clustered setup with DDoS mitigation and SSL support. -annual support contracts	\$31,985		\$25,856	\$48,480	\$46,780-\$57,580

AWS prices taken on Dec 2013 from the AWS data-center in N. Virginia. aiScaler and Citrix prices were taken from their respective websites, December 2013. Annual support contract are based on a clustered setup. F5 price taken from third party resellers, December 2013. Testing was done with a free trial with a 10mbit limit, with the same features. The testing license is not for sale for production, unlike the Netscaler 10mbit license. SKU's involved: EW2Z0000081 EW2Z0000084 F5-BIG-LTM-AWS-200M , F5-BIG-AM-AWS-200M, F5-SVC-BIGVESTDL13, F5-CST-S, F5-CST-P

Pricing in a high-bandwidth scenario with multiple domains

aiScaler licensing has no software limit on bandwidth or functionality. There are effectively only two licenses: with or without SSL termination -neither of them limited by bandwidth. This makes aiScaler considerably cheaper in high-bandwidth environments. Netscaler and F5 ask respectively \$30,000 and \$31990 for their 1GB licenses, in addition to more expensive support contracts. aiScaler license fees do not increase, even in high-bandwidth environments. aiScaler is also the only manufacturer that allows sharing of licenses for multiple domains. F5 and Citrix require clients to buy separate licenses for each domain, effectively doubling the costs for two domains. Additional domains for aiScaler cost \$1,995

	aiScaler Enterprise	Netscaler VPX Platinum 1Gbit	F5 BIG-IP 1Gbit, AAM +LTM
Minimum TCO based on: -Aquirement of licenses for a high-availability setup, with DDoS mitigation and SSL support -2 domains -a high-bandwidth requirement (1 Gbit) -annual support contracts	\$35,975	\$193,920	\$153,560-\$175,160

Feature comparison table

	aiScaler	Netscaler VPX Platinum	F5 BIG-IP AAM + LTM
Load balancing	✓	✓	✓
Clustering	✓	✓	✓
Compression	✓	✓	✓
SSL	✓	✓	✓
Dynamic cache	✓	✓	✓
Cookie driven caching	✓	✓	✓
Content driven caching	✓	✓	✓
Cache Expiration Using a Request	✓	✓	✓
Cache Expiration Using a File	✓	✗	✗
Cache Expiration Using a Form	✓	✓	✓
POST caching	✓	✓	✓
Unifying cache for diff. websites ⁴	✓	✗	✗
SPDY	✗	✓	✓

Response pre-fetching	✓	✓	✓
Mobile Devices Detection	✓	✗ ¹	✓
External origin	✓	✓	✗ ³
URL rewrites	✓	✓	✓
Geo-IP Location	✓	✓	✓
Origin based on request	✓	✗	✗
Web-Based Deployment	✓	✓	✓
Health monitoring	✓	✓	✓
Real-Time Stats	✓	✓	✓
CLI	✓	✓	✓
SNMP	✓	✓	✓
Nagios	✓	✗ ²	✗ ²
Application Firewall	✓	✓	✓
HTTP DoS Mitigation	✓	✓	✓
Malware Scanning	✗	✗	✗
Email alerts	✓	✓	✓
TCP Multiplexing	✓	✓	✓
SSL TPS Limit	unlimited	500	depends on license
HTTP throughput	unlimited	depends on license	depends on license
Customer plugins	✓	✗	✓ iRules
CSS/JS auto minify	✗	✗	✓
Free support	2 hours	none	none
AWS hourly billing	✓	✓	✗
AWS VPC required	No	Yes	Yes

¹ No out-of-the box solution, user has to set up own patterns.

² You have to use (community-made) plugins.

³ Origin must be within a VPC.

⁴ Included in the AWS aiScaler license. For dedicated licenses there is an extra fee of \$1,995 per domain for aiScaler. F5 and Netscaler require separate licenses for each domain, effectively doubling the costs for two domains.

3. Ease of installation

aiScaler

First, you need to launch a regular EC2 instance through the marketplace. You can use the web-based deployment tool (www.aiscaler.com/deploy) to edit and push the configuration file to the EC2 instance. Your site is then accessible through aiScaler; you can test it using the public DNS address of the EC2 instance. After completing the initial configuration, you can modify the caching patterns according to your application. The web based deployment tool offers caching pattern templates for the most popular CMS solutions.

Support offered: First 2 hours free, after that 175 usd/h

Time to install: 30 minutes

Skills required: basic AWS knowledge

Netscaler

Configuring Netscaler is more complicated, even though it does have a web GUI. In addition, you need a VPC, 2 Elastic IPs and 2 private IPs to run Netscaler, which are not required by aiScaler.

First, you have to set up a VPC with at least 2 subnets. Then start an instance within your new VPC. You cannot use a public DNS address to connect to a VPC instance; you must associate an Elastic IP or use a VPN. After logging in to the Netscaler management console, you should allow dynamic cache functionality, configure a service with your origin server and create a virtual server, which will include the service. Then add a subnet IP (SNIP) or mapped IP (MIP) to be able to enable the health checks on the service. Return to the AWS Management Console and add two private addresses to the instance - mapped IP and virtual server IP. Finally, associate a second Elastic IP for the private virtual server IP to make the service available from internet. Now you should be able to access your site. Do not forget to configure the cache patterns afterwards. One of the main drawbacks is the slow Java based web GUI.

Support offered: based on agreement or \$1,020/year (premier support for Netscaler 200 Mbps)

Time to install: 2 hours

Skills required: AWS VPC, TCP/IP

Note: Netscaler's caching mechanisms is easy to circumvent, when getting attacked by a DDoS attack. Netscaler completely bypasses the cache, when a request header does not contain an "Accept-Encoding field" to specify gzip compression. When a request is cachable, but not yet cached, Netscaler can open hundreds of connections to the origin by default, while aiScaler will open a single one, wait for the response and serve a cached version. In other words, you can attack the Netscaler, by waiting for the cache to expire and then the whole attack power will reach the origin, making it inaccessible and impossible to refresh the cache. This allows attackers to directly reach the origin server to DoS it, effectively making the protection useless.

F5 BIG-IP

F5 BIG-IP configuration is similar to Netscaler. It also has a web GUI and it also requires a VPC, 2 Elastic IPs and 2 Private IPs. You have to be careful how you design your infrastructure, because BIG-IP may be unable to connect to your origin, we did not have a similar issue with aiScaler or Netscaler.

First, you have to set up a VPC with at least 2 subnets. Then start an instance within your new VPC. You cannot use a public DNS address to connect to a VPC instance; you must associate an Elastic IP or use a VPN. After connecting to the instance, you will be asked to do some basic setup like service selection and licensing. Then you have to add VLANs and Self IPs. After this you have to create your application. One of the main drawbacks is the complicated application setup and especially caching configuration.

Netscaler and F5 are similar in the basic configuration (of adding a website). They both require a VPC, network interfaces and IPs setup. However F5 is much more complicated when you have to configure some advanced settings like caching patterns, origins, headers and logging.

Support offered: An annual support contract may be purchased separately from F5 Technical Support Services.

Time to install: 2 hours+

Skills required: AWS VPC, TCP/IP

4. Performance testing

Testing was performed with four different tools:

- Apache Benchmark to measure the maximum number of RPS.
- Pingdom to measure page loading times
- Siege, to measure stability and response time under a load of certain amount of users.
- Blitz.io, which is similar to Siege, but tests are performed from multiple servers around the world, for a more real-world simulation

For the Siege and Apache Benchmark we used Ubuntu 12.04, as load generator, running on a m1.large instance. For the origin server we used an Ubuntu based LAMP configuration. To be able to compare performance of the ADC's we removed the backend bottleneck by configuring both solutions to cache everything. In other words, after warming up the cache, tests didn't even reach the origin (backend) server. This allows us to use the smallest instance available (t1.micro) without any impact on the results.

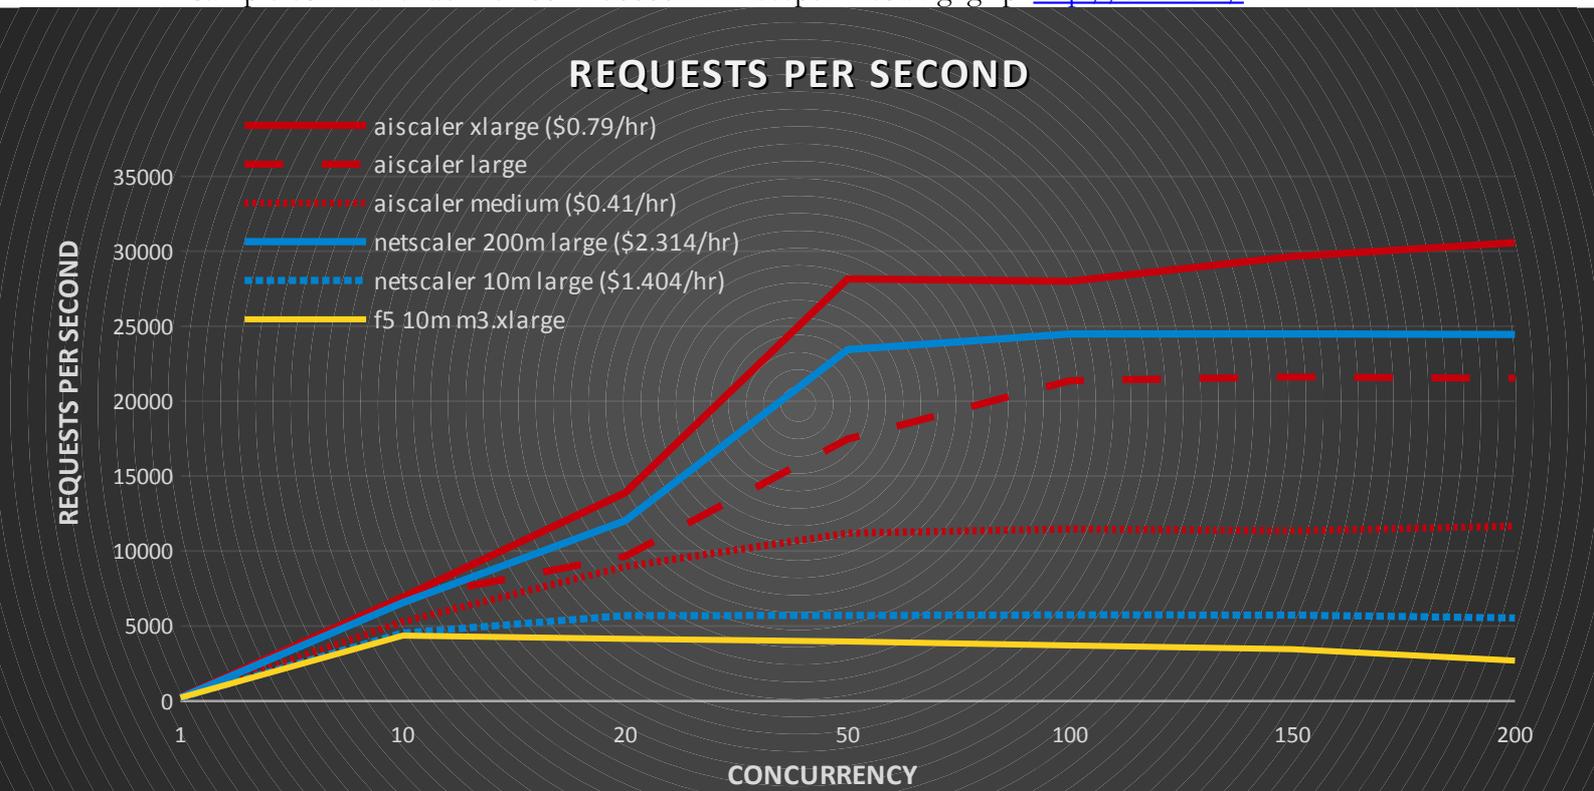
After the cache was warmed up, we kept monitoring the origin server to make sure traffic didn't reach Apache at all. All tests were done in threefold, to make sure the obtained data is correct. All instances were launched in the same AWS availability zone: us-east-1a

4.1 Apache Benchmark – Measuring RPS

This is a command line utility, which shows you how many requests per second your application is capable of serving.

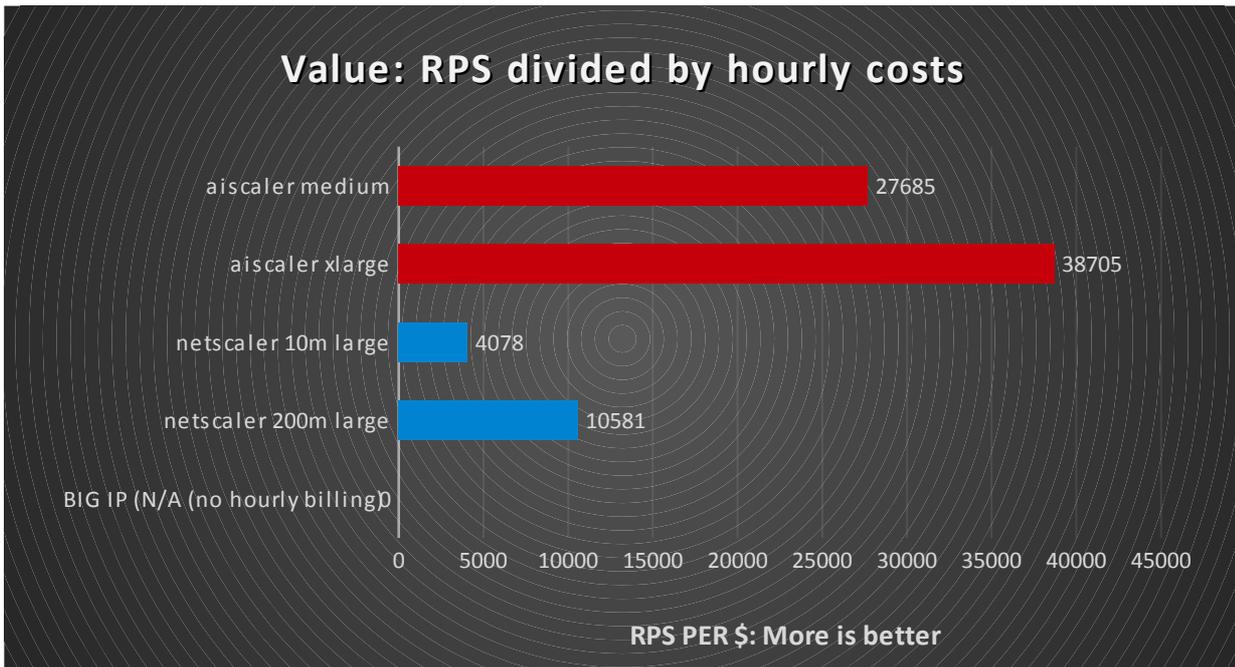
What is measured: number of maximum requests per second.

Sample command: `ab -kc 200 -n 50000 -H 'Accept-Encoding: gzip' http://site.com/`



While running on the same hardware (m1.large), Nginx 200mbps serves around 10% more requests than aiScaler. However you can run aiScaler on a much bigger instance, for the same price tag. The aiScaler xlarge (\$0.79/hr) is still cheaper than the Nginx 200mbit large (\$2.314/hr), and yet it serves more RPS than the Nginx instance.

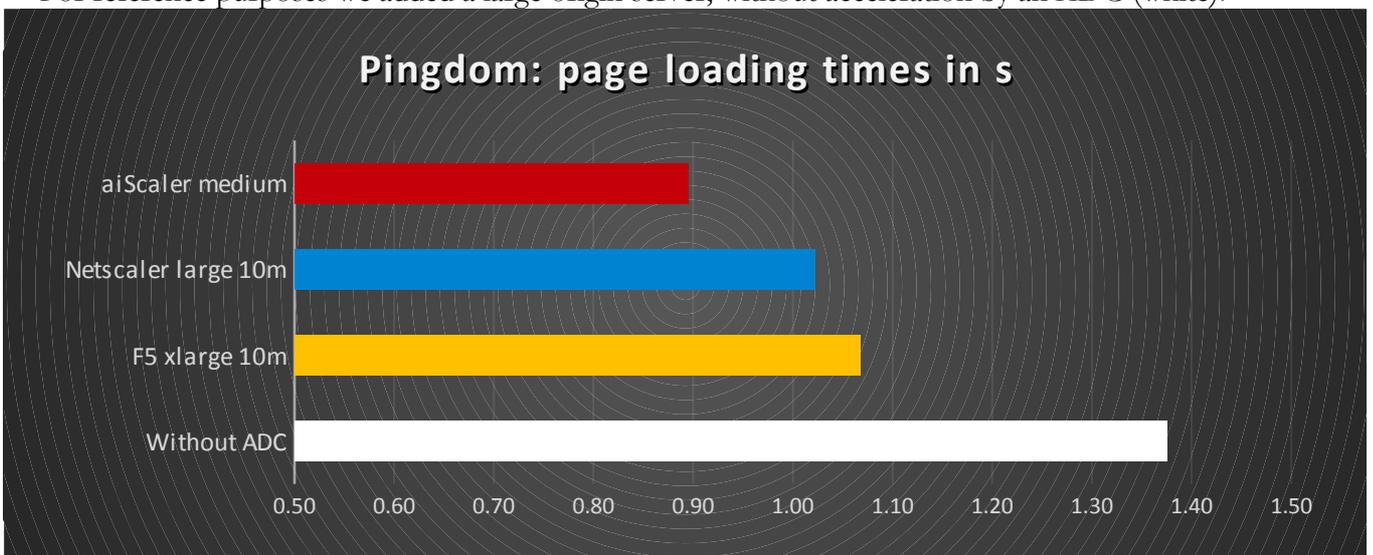
For a good value comparison, we divided the RPS by the hourly costs:



4.2 Pingdom – Full page speed test

Pingdom provides a third-party testing tool that can be accessed through: <http://tools.pingdom.com>

Netscaler and F5 offers the on-the-fly SPDY implementation, so we expected them to be ahead in this test. However, aiScaler implements a similar mechanism as SPDY, which actually ends up being faster. For reference purposes we added a large origin server, without acceleration by an ADC (white).



The page loading times in the table above are an average of a Wordpress template Index, Wordpress post, Magento template Index and a Magento product page. For the raw data go to:

http://aiscaler.com/wp-content/uploads/2014/01/Raw_data_page_loading_times.xlsx

The size of the instance or type of license does not matter much, when measuring page loading time of one page at a time. This test purely measures the speed of the internal architecture of each software.

4.3 Siege – Stress-test 1

Siege is an open-source command line utility, which allows one to hit a web server with a configurable number of concurrent simulated users. In addition, instead of testing against a single URL, Siege allows you to test against multiple. This allows for a more real-world simulation of how a user would use your system.

What is measured: server stability and response times under a load of certain amount of users. Sample command: `siege -v -i -c 300 -t 1m -f links.txt`

	aiScaler medium	aiScaler large	Netscaler 10mbit	Netscaler 200mbit	F5 10mbit m3.xlarge
Transactions:	35080 hits	35080 hits	34212 hits	35385 hits	35358 hits
Availability:	100.00%	100.00%	100.00%	100.00%	100.00%
Elapsed time:	59.65 secs	59.65 secs	59.64 secs	59.51 secs	59.16 secs
Data transferred:	84.14 MB	84.14 MB	82.09 MB	84.67 MB	4.89 MB
Response time:	0.00 secs	0.00 secs	0.01 secs	0.00 secs	0.01 secs
Transaction rate:	588.10 trans/sec	588.10 trans/sec	573.64 trans/sec	594.61 trans/sec	597.67 trans/sec
Throughput:	1.41 MB/sec	1.41 MB/sec	1.38 MB/sec	1.42 MB/sec	0.08 MB/sec
Concurrency:	30348	30348	33664	17564	36192
Successful transactions:	35080	35080	34212	35385	35358
Failed transactions:	0	0	0	0	0
Longest transaction:	0.06	0.06	0.22	0.11	1.00
Shortest transaction:	0.00	0.00	0.00	0.00	0.00

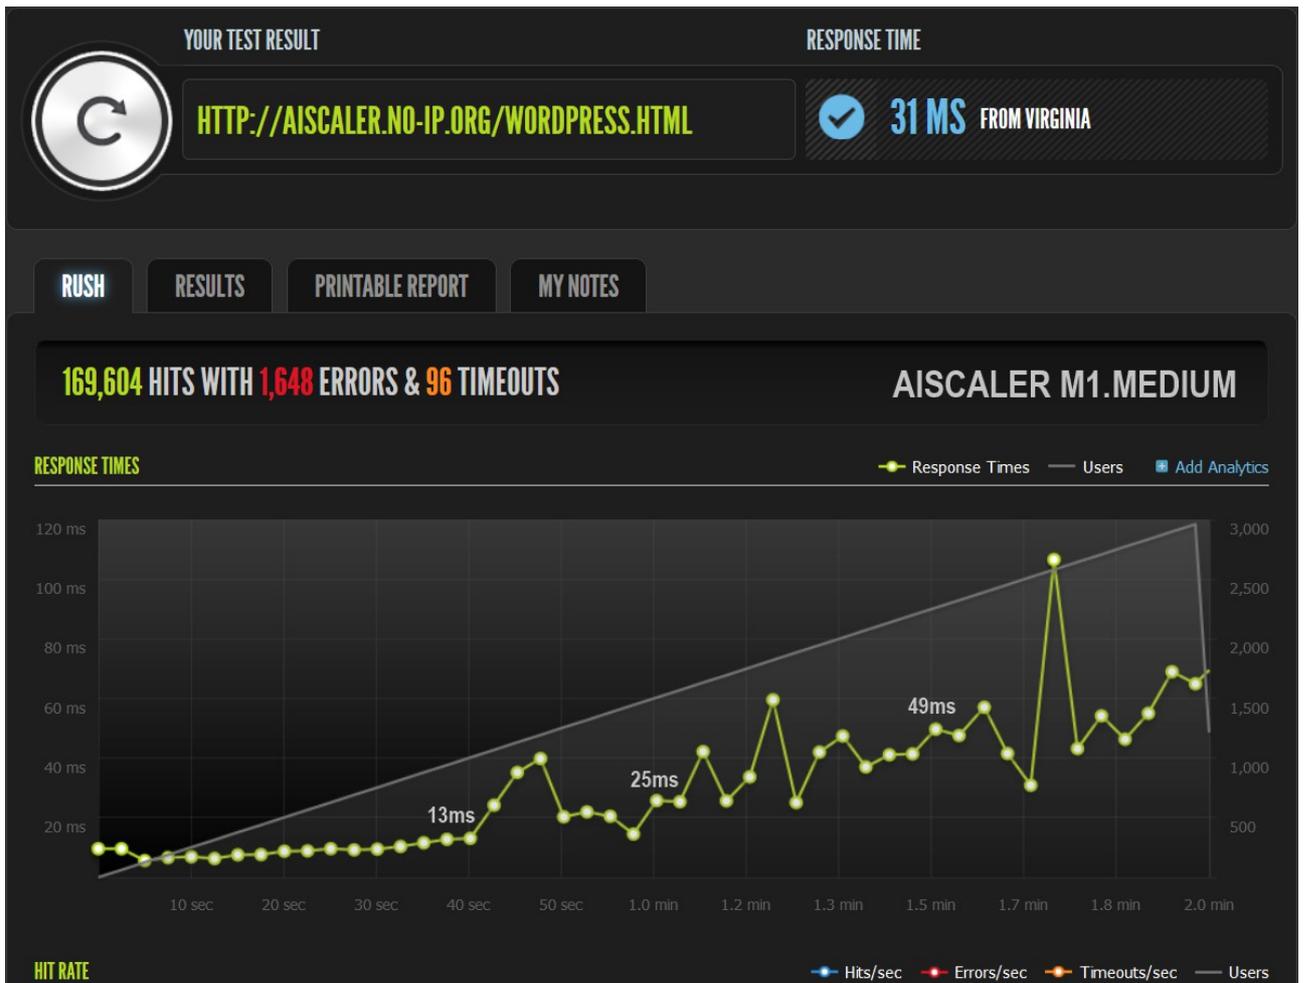
All ADC's were stable with a 100% availability and zero failed transactions. When it comes to transaction time, aiScaler beats Netscaler and F5, even when running on a smaller instance. These Siege results are largely irrelevant, as it clearly doesn't push any of the ADC's to the point of breaking down.

4.4 Blitz.io – Stress-test 2

This is a commercial service, similar to Siege, which offers a web interface and automatically generates graphs. Tests are performed from multiple servers around the world, for a more real-world simulation. It measures server stability and response times under a load of an increasing number of users. We kept on increasing the number of users, until the ADC would break down, as you can see from the number of errors.

Sample command: `-H "Accept-Encoding: gzip" -p 1-3000:120 http://site.com/`

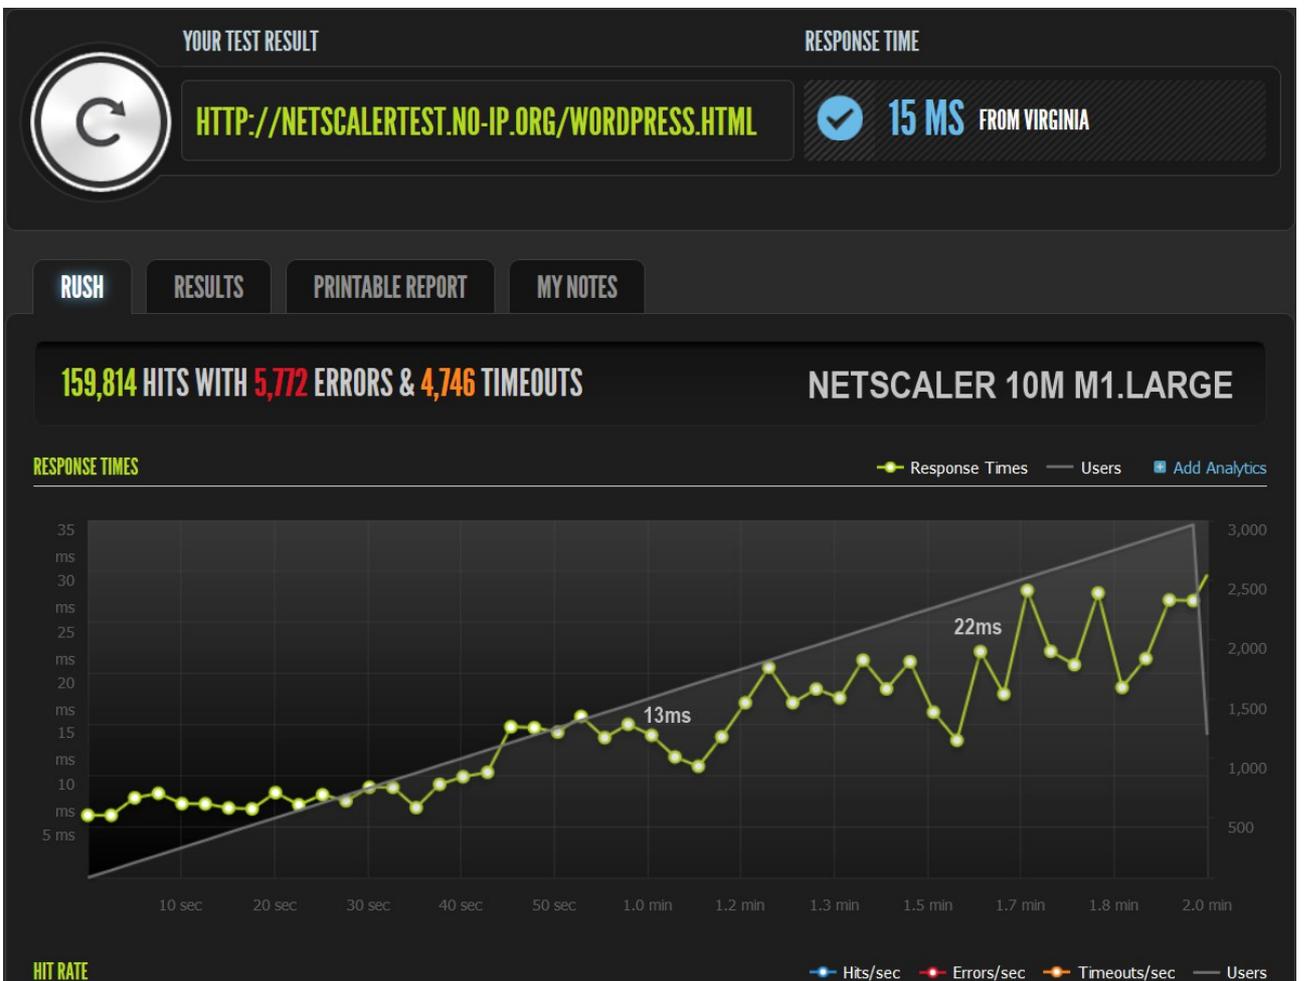
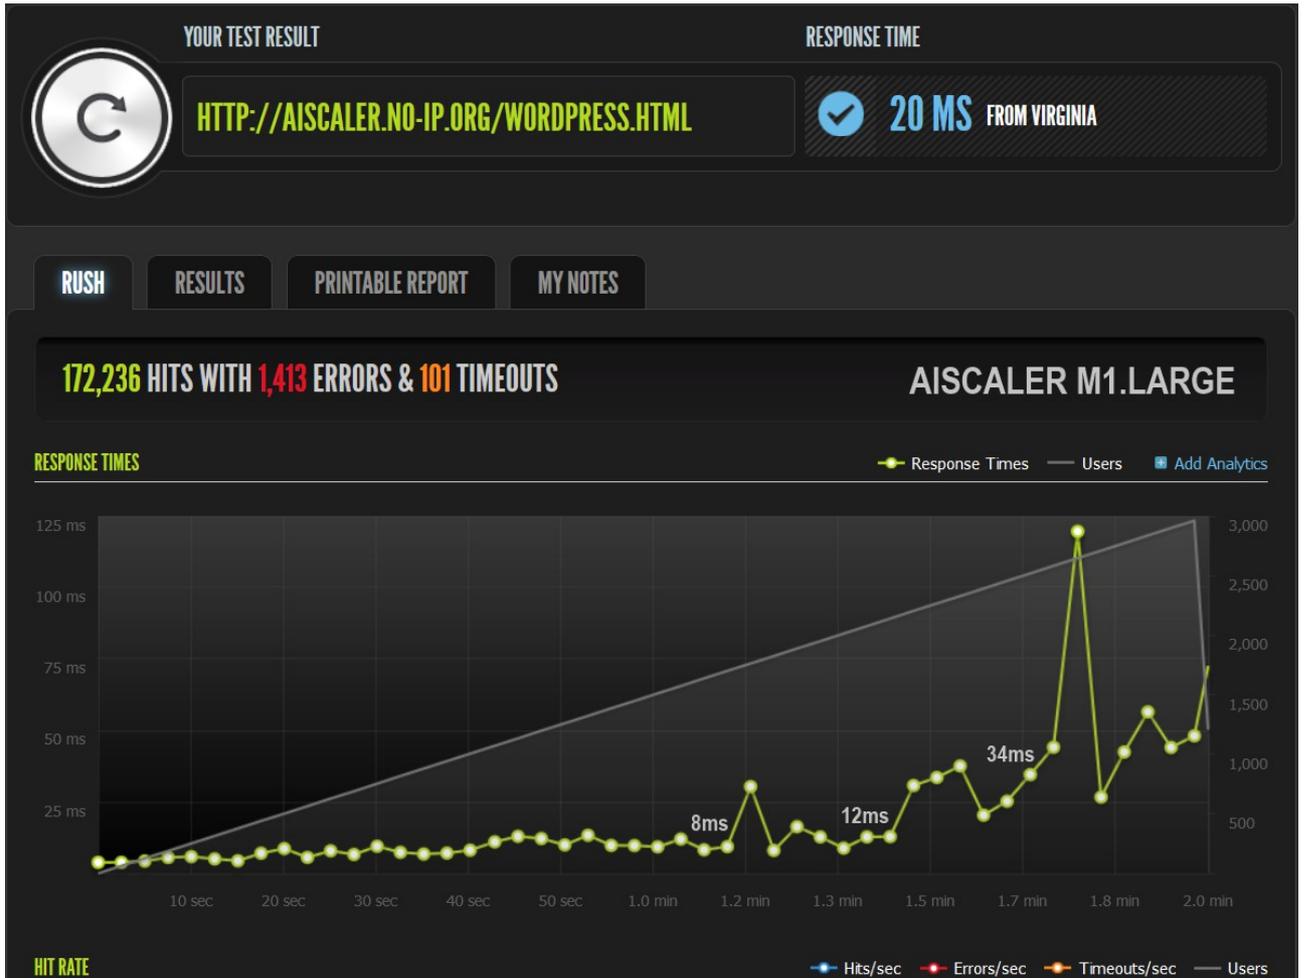
A typical screenshot of a blitz.io tests looks like this:

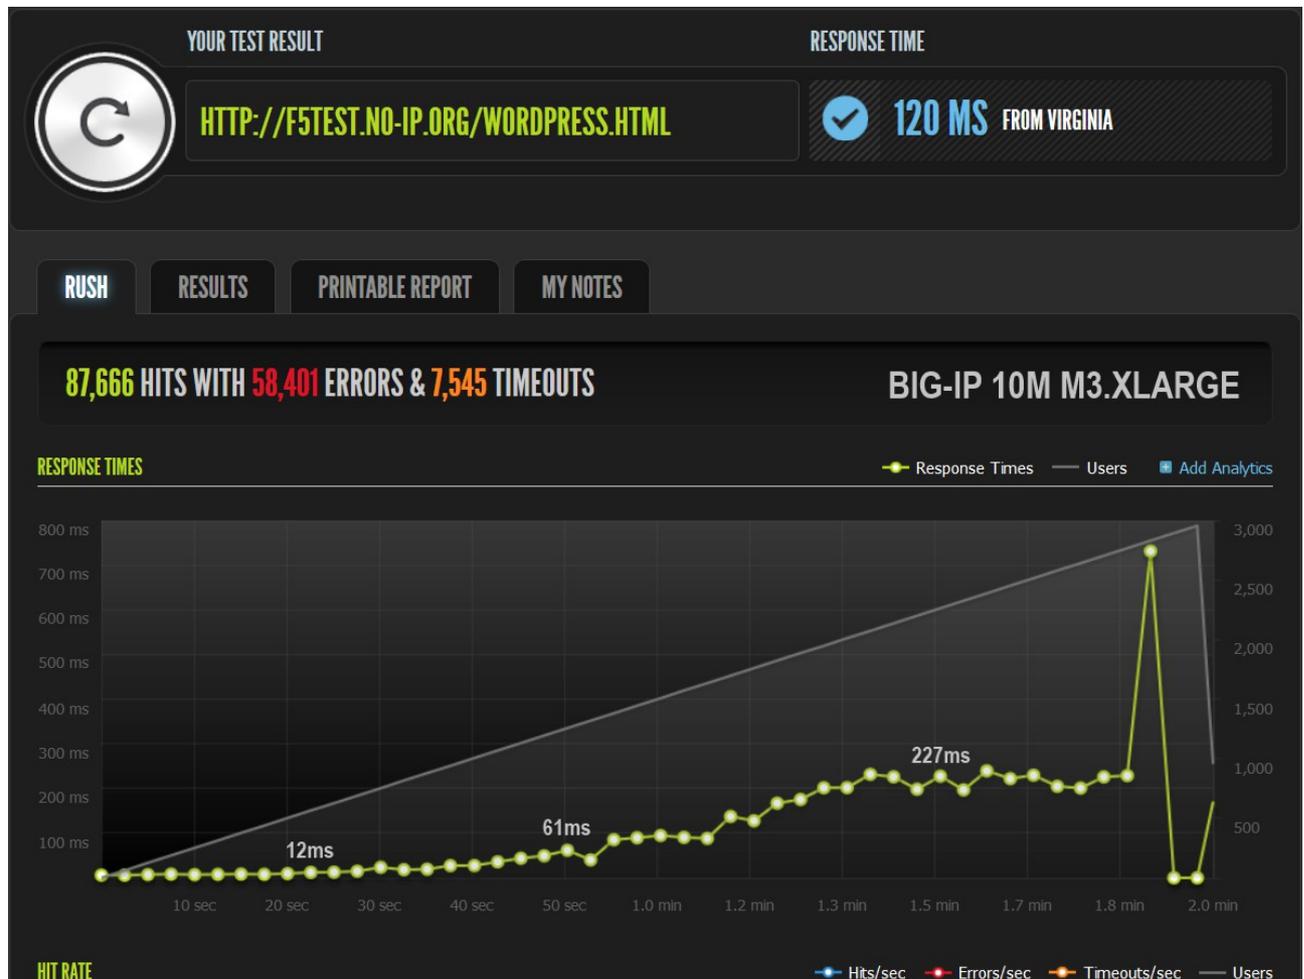
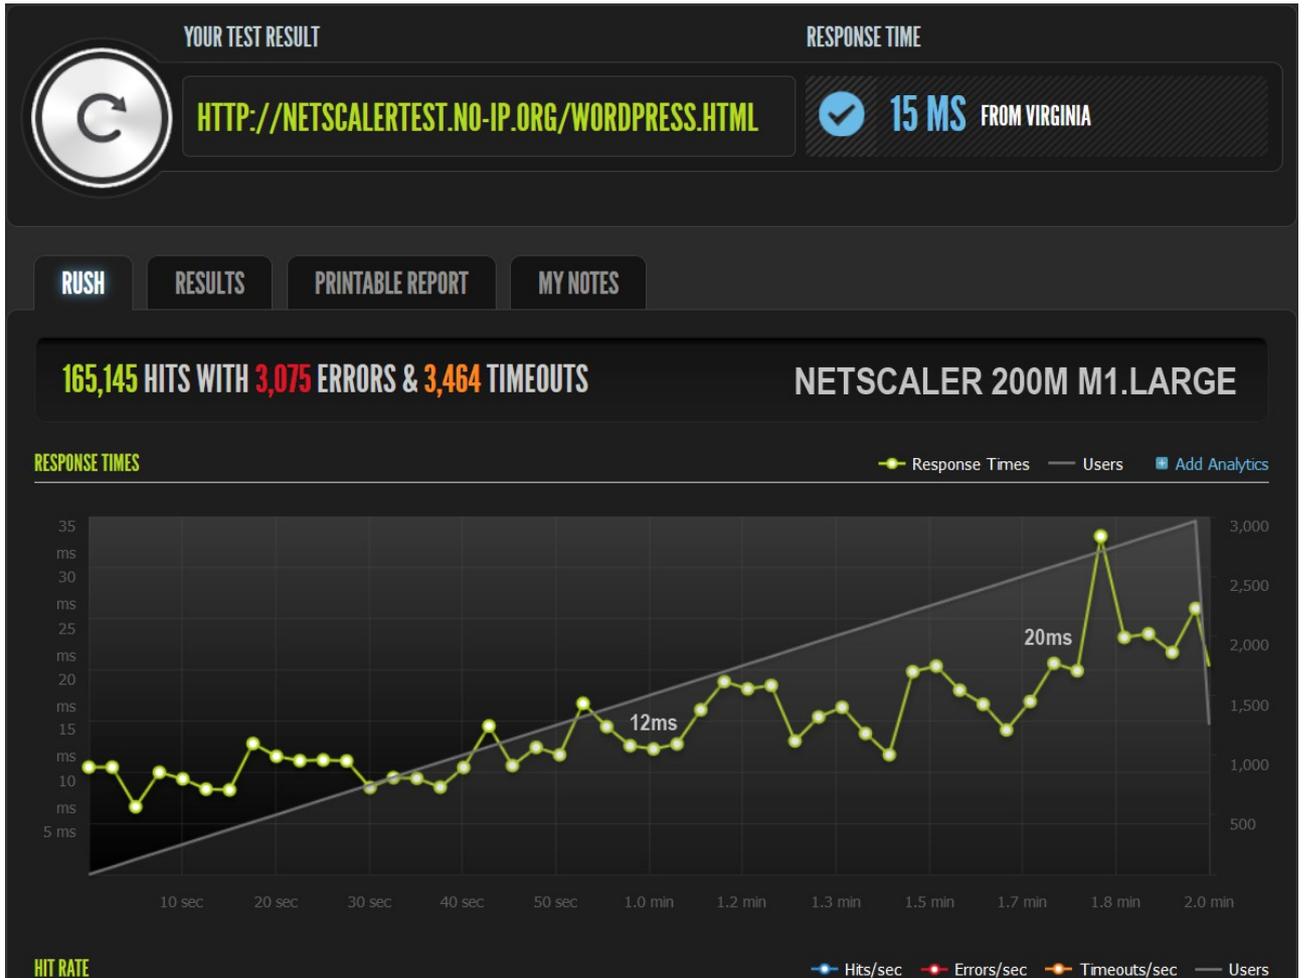


Blitz.io summary of results:

	Hits in 2 minutes	Errors & time-outs	Failure %
aiScaler m1.medium	169604	1744	1.0%
aiScaler m1.large	172236	1514	0.9%
Netscaler 10Mbit m1.large	159814	10518	6.6%
Netscaler 200Mbit m1.large	165145	6539	4.0%
F5 Big IP 10Mbit m3.xlarge	87666	65946	75.2%
Varnish	142964	15487	10.8%

Performing such a test is similar to a DoS attack, or simulating a situation with a peak of internet traffic. The most relevant measurements are the number of hits that the ADC was capable of accepting and the failure% under stress. For the fun of it, we also added test-results of Varnish. Even though Varnish is an open-source product - and is completely different in terms of features, support and ease of installation- it is possible to compare HTTP caching performance of a simple Wordpress site. On the following pages you can see the screen shots from all tests.







5. Conclusion

All products from aiScaler, F5, Citrix were easy to install and configure for an advanced network engineer. You can get started with aiScaler a lot faster, as it does not require setting up a mandatory Virtual Private Cloud network interfaces and several extra IP's addresses. It is also the only manufacturer to include 2 hours of free installation support.

After the initial configuration Citrix and Netscaler are relatively easy to configure, while F5 is more complicated when it comes to caching patterns, origins, headers and logging. F5 does have the community-driven plugins (iRules), which is offset by the lack of out-of-the box features.

Feature-wise the ADC's are comparable. Depending on your application, you might favor one product over the other. For example, for mobile content management and device detection, Netscaler seems a bad choice. F5 on the other hand, is not compatible with origins that are hosted outside of the AWS cloud. A huge advantage of aiScaler is that it can unify cache for several domains on the same instance. aiScaler lacks SPDY support and does not minify CSS and Javascript. Despite this aiScaler won 3 out of 4 benchmarks.

The real difference comes with the price tag and performance. F5 does not offer hourly pricing, which defeats one of the main advantages of modern cloud-computing. It requires a minimum investment of USD \$16,990, which still doesn't include computing costs. Netscaler starts at \$1.40/hr and aiScaler at \$0.41/hr.

Page loading times were lowest under aiScaler, followed by Netscaler, then F5. Only in ApacheBench, Netscaler managed to narrowly outperform aiScaler, and only when comparing performance on the same hardware. However, for the same total *costs* (hardware plus software), aiScaler instances easily outperform Netscaler instances. F5 -again- is the worst performer, and not just because their license was limited by bandwidth¹

Especially during DDoS attacks, aiScaler will be the best performer, as it has the best results in the stress-tests. F5 lacks performance and Netscaler, while better than F5, still can't nearly keep up with aiScaler and has exploitable defense mechanisms.

¹ It was hard to believe the F5 performance was consistency that bad, so we posted the results on the F5 support forum. We were told that these results are accurate for this license. Admittedly, the license was limited at 10mbit, but so was the Netscaler 10mbit license, which performed a lot better. Assuming that these results can be extrapolated to other F5 licenses, you need a more expensive license to achieve the same results as with Netscaler license, that is limited by the same bandwidth. Even then, a more expensive F5 license will not improve page loading times or the complicated configuration process.