



# ***Cloud Computing Testing***

**April 2010**

## The New Data Center

With Local Area Network (LAN) and Internet bandwidth becoming less and less of a limiting factor, computer and network architectures are experiencing a paradigm shift. Data centers have been reassessed regarding both their architecture and the services they are able to provide. Enterprises with giant data centers have thus realized that with certain technology advances they can monetize on their investment in storage and computing resources in their data center by offering them up to customers as a new type of service: public cloud computing. Enterprises building their own private data centers, or private clouds, are also interested in the benefits of these advances.

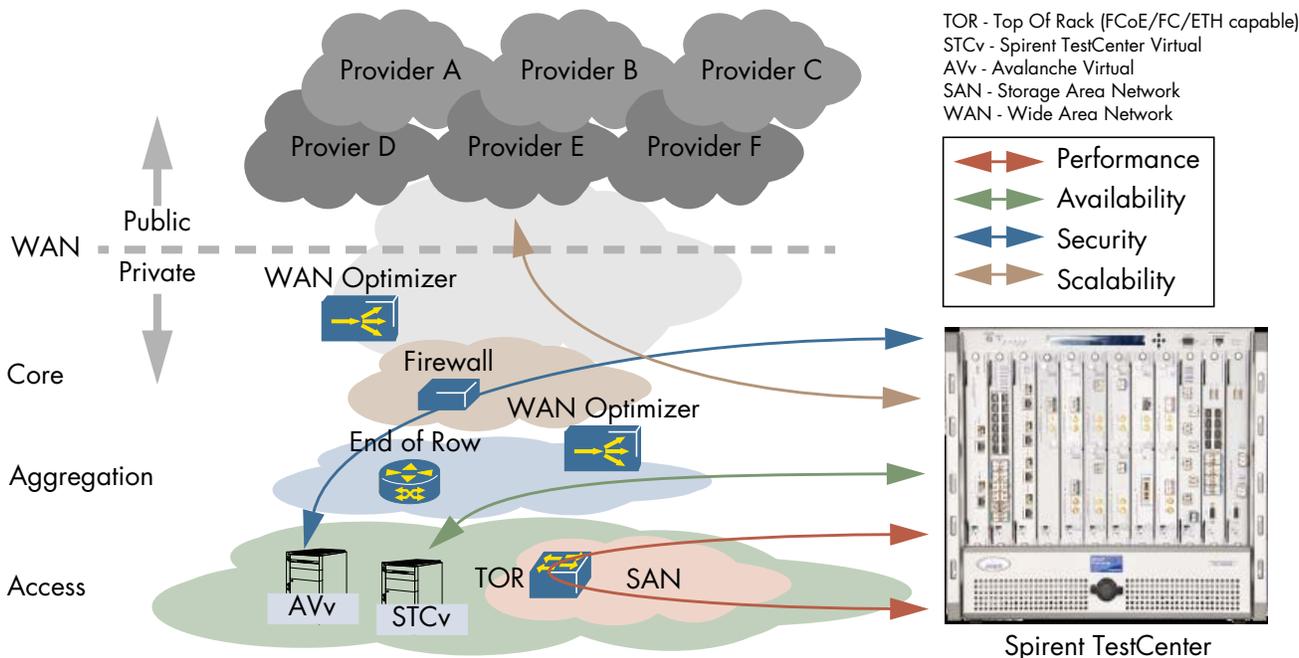
Such advances have transformed the data center to a converged and flexible environment. One of the building blocks of the transformation has been Fibre Channel over Ethernet (FCoE), which allowed the convergence of LAN links (typically Ethernet) with Storage Area Networks (SAN) links, most of which use Fibre Channel. SAN has abstracted storage from servers for years, however, a more recent trend has yielded the power and flexibility of abstracting the physical server resource from its users - virtualization.

This report details the methodology and results for a series of tests conducted by EANTC together with Spirent Communications. Each test focuses on a certain piece of the data center infrastructure. With one exception of a test conducted on public cloud enterprise services, all tests were performed within a

private mock data center network built at Spirent premises in Calabasas, CA. During the testing EANTC and Spirent engineers were addressing the following common questions:

- Data Center Benchmarking - How can I measure the performance of my new hybrid Data Center interconnection devices?
- Firewall Performance and Scalability - How does my firewall affect the performance of my data center when the number of user sessions increase?
- Virtual Security - How do I verify that my Virtual Machines (VMs) are secure and block connectivity to other VMs?
- Availability - How can I see my Service Level Agreement (SLA) is honored by my provider? Do advanced features such as VM migration affect availability to the customers to whom I provide?
- WAN Optimization - How do I prove that such systems really increase my application performance and provide return on investment (ROI)?
- Public Cloud Computing - Do my cloud computing services provide me the resources and reliability on par with my own data center?

This paper focuses on methodologies for testing potential issues arising as companies shift to a hybrid model of both public and private cloud computing.



**Figure 1: Generic Data Center Context**

## Private Cloud Data Center Test Scenario

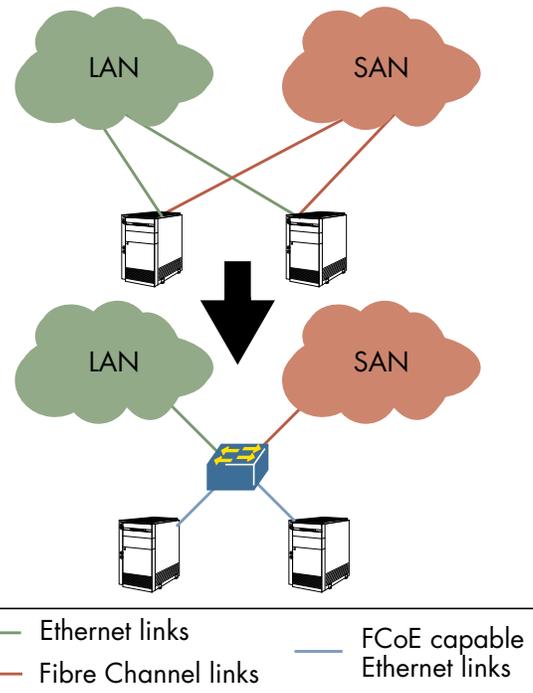
The tests were grouped into two sections - testing a private in-house data center, and testing a cloud service from a customer perspective. For the private data center, we assembled enough devices to build a small emulated data center and conduct tests on the equipment we wanted to hone-in on such as the Firewall, Data Center Bridging, Virtualization, and WAN optimization. In each specific test section below we describe the test setup, however Figure 1 shows how they fit together to build the data center puzzle. Since this is a public report with the purpose of discussing test methodology, vendor and model names tested will remain anonymous.

## Data Center Benchmarking: Queueput

Performance of packet switching devices continues to be a crucial and staple test for any enterprise or service provider installing a new switch or router. The IETF published RFCs 2544 and 2889 some 10 years ago to define a benchmarking methodology for IP and Ethernet (respectively) forwarding devices. These days, as Data Center Bridging (DCB) forwards data in a significantly differently way than LAN switches or routers, Spirent has jointly proposed an update to the benchmarking RFCs focusing on methodology for benchmarking data center bridges in an individual draft<sup>1</sup>.

When one speaks of a Storage Area Network (SAN), most likely one refers to either iSCSI or Fibre Channel. While these are seen by some as competing technologies, it is not debated that Fibre Channel has an extremely large install base - the IDC (International Data Corporation) tells us that \$9.2 billion was spent on Fibre Channel, FCoE, and Infiniband equipment in 2009. In the past years, as network protocols continue to reap the benefits of convergence, the T11 committee from the InterNational Committee for Information Technology Standards (INCITS) has defined a way to converge Fibre Channel and Ethernet networks - Fibre Channel over Ethernet (FCoE). The standard facilitates reduction of cables and ports in the data center. There was one major issue with putting Fibre Channel over Ethernet - Ethernet is a lossy protocol. Fibre Channel is not.

1. <http://tools.ietf.org/html/draft-player-dcb-benchmarking-01>



**Figure 2: Data Center Migration**

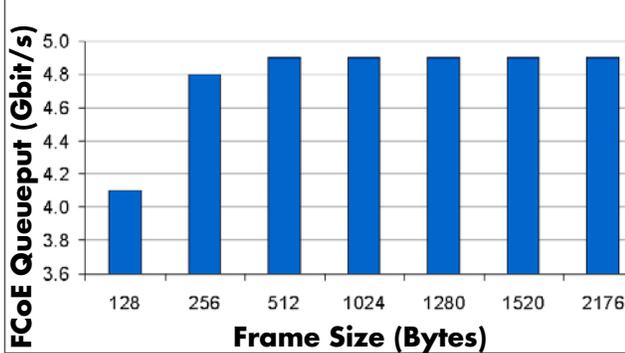
The solution came from the IEEE working group for Data Center Bridging (DCB) that defined control protocols such as DCBX (Data Center Bridging Capabilities Exchange Protocol) for automatic discovery of features, but mostly defined enhancements to enable lossless Ethernet. This is done by using Priority Flow Control (PFC, defined in 802.1Qbb) where a device uses a special frame - a PAUSE frame - to tell the upstream device that traffic rates are approaching a point where loss could occur. This allows the upstream device to stop transmitting traffic and thus avoids loss altogether.

Since RFCs 2544 and 2889 define throughput as the highest data rate achievable without frame loss, the same terminology could not apply to DCBs. The data center benchmarking draft proposed by Spirent now defines *queueput* as the highest data rate achievable on a queue before observing any PAUSE frames on it. Otherwise an algorithm similar to the original RFCs is used. Since this is done for each priority queue, it has been dubbed not as throughput, but *queueput*. Additional test methodologies are defined in the IETF draft, however, *queueput* measurements were performed here.

Since most disk arrays and SANs still have native Fibre Channel interfaces only, hardware vendors have started to produce switches capable of converging

Ethernet and Fibre Channel by means of FCoE. These are typically called “Top of Rack” switches (placed at the top of a data center rack) so they may now connect to servers, the LAN, and the SAN, leaving much more flexibility in the data center rack design. Such a device, was our DUT for this test.

While the queueput test can be used on any queue configured on the device, our test focused on the FCoE queue. Typically only one queue is configured for this use, and is considered the utmost highest priority (it is also usually priority code point = 3 in the VLAN header. The Spirent TestCenter connected two 10 Gigabit Ethernet ports to our DUT and transmitted two flows - a “low priority” Ethernet queue statically set to 5 Gbit/s, and the FCoE queue which was configured to increase its bandwidth steadily until PAUSE frames were received.

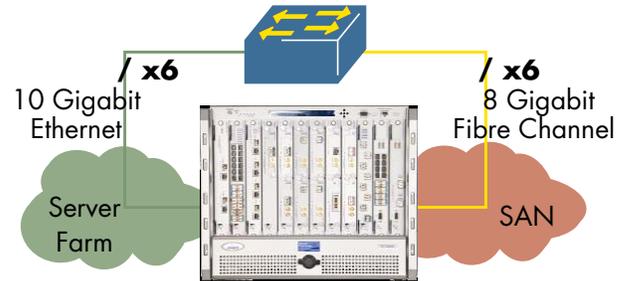


**Figure 3: Queueput Results**

The standard recommends testing 128, 265, 512, 1024, 1280, 1520, and 2176 Byte frames. Traffic of 256 Byte frames reached 4.8 Gbit/s without sending PAUSE frames, and only 4.1 Gbit/s for 128 Byte frames. The rest of the frame sizes achieved a queueput of 4.9 Gbit/s, seeing flow control kick in at 5.0 Gbit/s. In fact, early tests conducted before adding the “low priority” traffic showed this queueput result, and was the reason for statically sending 5 Gbit/s of “low priority”. This is not a major cause for concern as Fibre Channel interface typically send large frames. This tells the user of such a device that they should not expect to make use of the full 8 Gbit/s on the native Fibre Channel interfaces which would normally translate to 6.8 Gbit/s of good data (after encoding).

As a follow up “sanity check” test, we also tested sending storage traffic from FCoE interfaces to Fibre Channel interfaces. We connected 6 ports of each on the DUT to 6 ports of each on the Spirent TestCenter and transmitted 10% of line rate bidirectionally in a

paired port configuration (Ethernet port 1 to Fibre Channel port 1, Ethernet port 2 to Fibre Channel port 2, etc). At this low rate we did not observe any loss or PFC (PAUSE) frames.



**Figure 4: FCoE to Fibre Channel Topology**

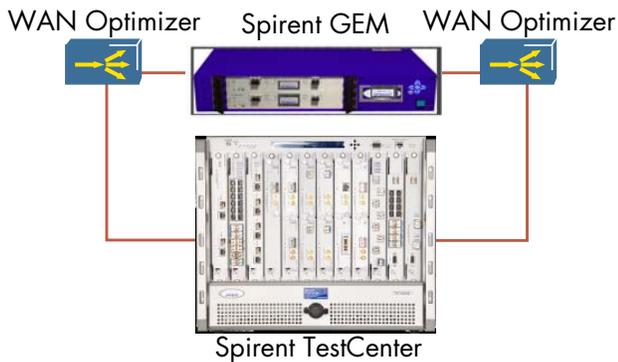
In addition to conducting these hybrid tests (Ethernet/IP and FCoE on a single Ethernet port in the queueput test, and FCoE to Fibre Channel in the above test) with the Spirent TestCenter, we were also able to observe the latency values. We did not set specific expectations for latency, but observed them nevertheless and saw none that would be a cause for alarm.

## WAN Optimization

Dubbed by some as WAN accelerators (which of course do not literally change the speed of your WAN links - unfortunately) Wide Area Network (WAN) optimization hardware aims to increase efficiency of end-to-end traffic in a few ways. Several vendors produce such equipment, each working in slightly different ways. Some analyze data to transmit only changed pieces (why re-upload an entire file when you only had to edit the authors name?) and minimizing unnecessary control data, some simply cache. WAN optimizers’ selling point is not necessarily technical - install one and the time you spend waiting for an upload will be reduced.

Our setup allowed us to test point-to-point connectivity. WAN optimizers need to store a lot of data and therefore consist of pretty heavy duty servers, and a few network interfaces. Two optimizers were used in the test, each positioned on one end of the emulated site. Between each of our two DUTs we inserted the Spirent GEM which adds delay and loss to traffic in an attempt to emulate a WAN. The Spirent GEM can be statically or dynamically configured - the latter allowing the user to create complex configuration by emulating specific real world-like scenarios. These scenarios, and their impairment, have been defined by the ITU-T (G.1050) and TIA (TIA-921), “Network

Model for Evaluating Multimedia Transmission Performance over Internet Protocol". We chose a relatively tough, but realistic profile from the catalogue presented by G.1050 - rate combination 131. The standard builds a series of complex tables in order to define the impairments for over a hundred realistic scenarios. The 131 scenario means that each site has a 20 Mbit/s LAN with 768 and 128 Kbit/s access links to a 13 Mbit/s WAN. We ran tests with and without the WAN optimizer for both 131A (low severity conditions) and 131F (more severe conditions).



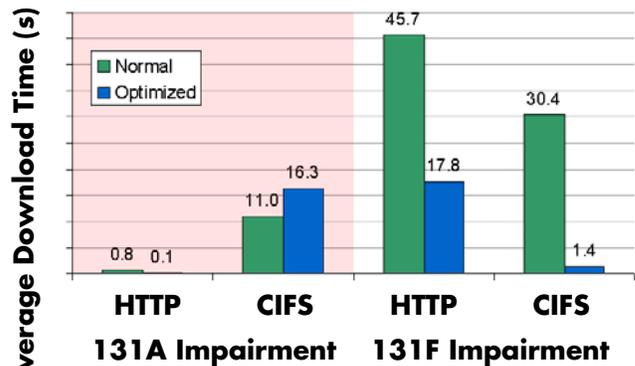
**Figure 5: WAN Optimizers Test Topology**

Two Spirent TestCenter ports running Avalanche software were then connected to the two DUTs, one emulating the data center side, and one user side of the enterprise. A variety of applications were emulated to reflect a realistic user profile, and to see how each application was affected by the delay, and optimization:

- Data - File transfers on a series of real files were executed using both HTTP and CIFS (Microsoft's "Common Internet File System")
- Video - Users requested an H.264 encoded MPEG video stream with RTSP from the server.
- Voice - Emulated users speaking to each other from each location.

We collected the average download time after running the tests for five minutes for each of the two Spirent GEM profiles without the WAN optimizers. After adding the optimizers to the topology we ran the test again to allow them to cache, eliminating cache ramp-up time as a test factor. We then repeated the same test with the WAN optimizer active.

The results were interesting. As expected, the optimizer reduced response times quite a bit for file based data such as HTTP and CIFS. Voice and video data

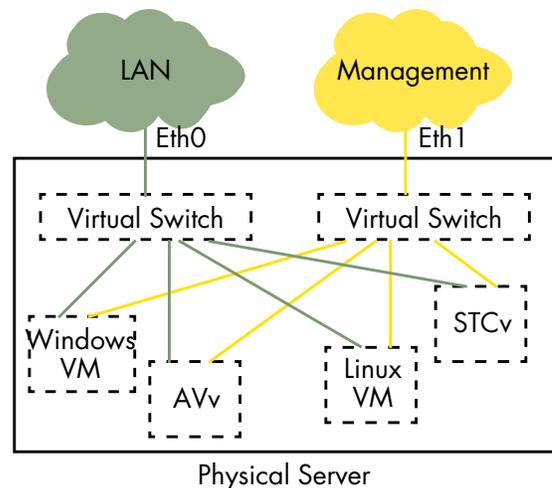


**Figure 6: WAN Optimization Results**

did not show a noticeable difference. However, with the less severe impairment profile the optimizer actually decreased CIFS performance. This repeated throughout our iterations. Without speculating too much into the algorithms of the WAN optimizer, we can still conclude that WAN optimizers will increase performance in most scenarios.

## Going Virtual

Virtualization. The term is everywhere in the tech world these days, and for good reason. The ability to use the resources in a physical server for different virtual machines gives an administrator the power to create different user domains on a machine level, where every user has the experience of using their own server to themselves, and allows him to be more flexible with the resources - Virtual Machines (VMs) can be migrated from one server to another with minimal interruption.



**Figure 7: Basic Virtual Network**

As CPU speeds have increased over the years, and disks became cheaper and networks' capacity grew, the idea to efficiently use said resources cemented itself under the term virtualization. The concept is fairly simple: a server (physical machine with memory, CPU, disk and network interface) spends a large portion of its time idling. Why not use these free resources to offer more services with the same hardware? To facilitate this virtual machines (VMs) vendors offer virtualization. Many VMs can in effect run on the same hardware, allowing efficient hardware resource utilization. What does this mean for networking? As shown in Figure 7 there is a concept of virtual switches and virtual links, but there are also different kinds of switches. How do you architect a virtual network? How do you make it secure? For all of the above - how can this be tested? Spirent has developed virtual versions of their tools to answer the latter question. Just as a VM is a virtual PC or server (running Windows, Linux, etc), Spirent TestCenter virtual (STCv) and Avalanche virtual (AVv) are virtual pieces of Spirent's test equipment. Their idea is to use them to address new testing challenges in the virtual space.

### Virtualization High Availability

One big advantage of virtualization is the ability to move a virtual machine from one physical server to another - while the VM continues to run. In order to do this, with minimal loss, a distributed virtual switch is required. The difference with the distributed virtual switch compared to the standard one is its ability to synchronize across physical platforms. Since this feature allows public clouds to offer highly available services both customers and the cloud service provider

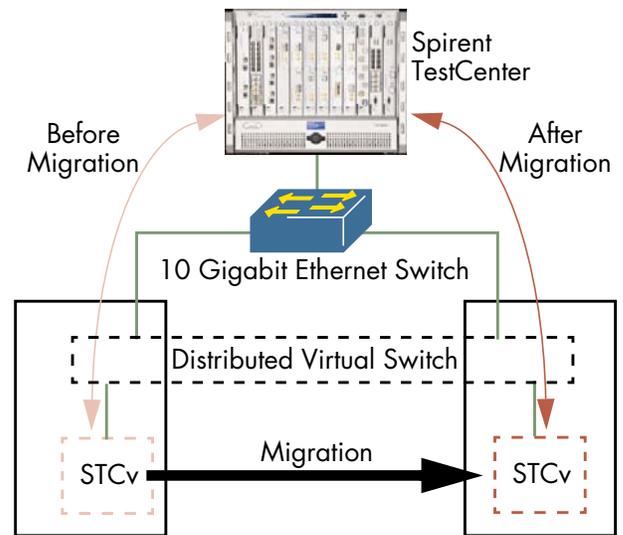


Figure 8: Availability Test Setup

have a vested interest in knowing the impact such virtual machine move has on the service.

To measure the effect a virtual machine move has on the service we transmitted traffic between two 10 Gigabit Ethernet ports- one virtual residing in a virtual machine and one physical within Spirent TestCenter. We transmitted 1280-Byte frames at 1,000 frames per second. We quickly realized then, that the virtual world is not frame-loss free, even before conducting the migration. We constantly observed roughly one lost frame per 5 seconds. This meant that defining out of service time as a summation of the number of lost frames, would lead to inaccurate results - frames were lost after all even when VMs were not being moved around. The solution was to monitor the receive rate on the physical port using Spirent TestCenter's "High



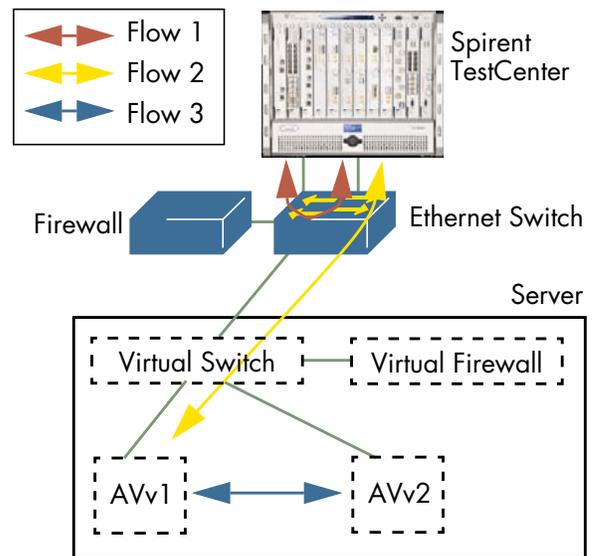
Figure 9: Availability Results

Resolution Sampling". After starting the sampling, and starting traffic, we performed the migration. The system typically took about 10-15 seconds to copy the VM over to the new physical server, but this duration is very much hardware dependant. We conducted the test three times in order to gain confidence in the results. The graph for all three runs looked similar to Figure 9, showing about 1.5 seconds of out of service time on the Ethernet layer. This matched the results collected based on the frame loss observed on that port. The results show that accepting VM High Availability in a service level agreement (SLA) should be taken with a grain of salt. For a web server, losing 1.5 seconds of traffic during migration could have significant consequences. Both HTTP and TCP are robust enough to compensate for the loss. In comparison to video however, which uses UDP as transport, the results will be devastating. This amount of frame loss will lead to the viewer watching a blank or pixilated screen and is likely to result in a very busy call center.

## Security

When it comes to installing and testing security systems and firewalls to protect data centers from attacks and users who should not have access, there are generally two main questions to answer: can my system block what it should, no more no less? and how does the introduction of this system affect the performance of the network? Physical hardware firewalls have tried their best to give impressive answers to these questions for years but this does not help within the virtual world. To take a peek at the status of data center security we ran a test on both a physical firewall and a virtual one.

The traffic profile used was the same as for our WAN optimizer test above consisting of Voice, Video, CIFS data and HTTPS data which filled approximately 90% of the Gigabit Ethernet interface (the protocols do not create a steady traffic rate, hence "approximately"). The physical firewall connected to the Spirent TestCenter with two Gigabit Ethernet interfaces running Avalanche software. We ran a test both with, and without, rules enabled (Flow 1 in Figure 10). The rules blocked a set of 5 IP addresses (of the 250 emulated) from all applications. The good news was that not only did the firewall block what it should have (no more, no less) but the CIFS and HTTPS response times, and the voice and video MOS scores, were not negatively impacted when rules were enabled.



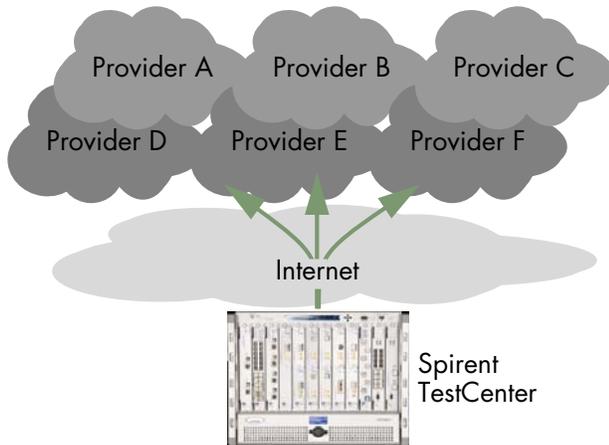
**Figure 10: Security Test Setup**

To test the virtual firewall we used two different scenarios: one where the Spirent TestCenter port was connected directly to the server (through a switch) running Spirent Avalanche Virtual (Flow 2 in Figure 10), and a second where only Spirent Avalanche Virtual was used to emulate user traffic (Flow 3 in Figure 10). The results were similar - each were a bit less consistent since virtual performance is still unpredictable, however we were impressed to find that the virtual firewall correctly blocked only what we configured it to block.

## Public Cloud Computing

As explained earlier in this report, enterprises with giant data centers are realizing that they can sell the unused resources as a service, often called Infrastructure as a Service (IaaS), by modulating them with virtualization. This is an attractive proposition for those customers or small and medium businesses that cannot afford to build their own data center. The price of building a data center is only one argument against rolling out one's own. A second prominent argument is expertise. Companies might rely on large data sets, but are in effect not in the business of managing data centers. Therefore, "outsourcing" data centers and computing power to specialists is gaining more ground. The increase in companies offering such services begs the question - how do they compare? More importantly - how can a user verify that these offers match the needs?

**Test Setup.** We contracted and established services with six different public cloud providers. Due to the terms and conditions of these services, we have left the public clouds nameless. It should be noted that these services are all priced differently, and offer slightly different services (NetworkWorld has published an article specifically on this aspect<sup>1</sup>). Still, they all provided us with the ability to start a VM and load applications as desired.



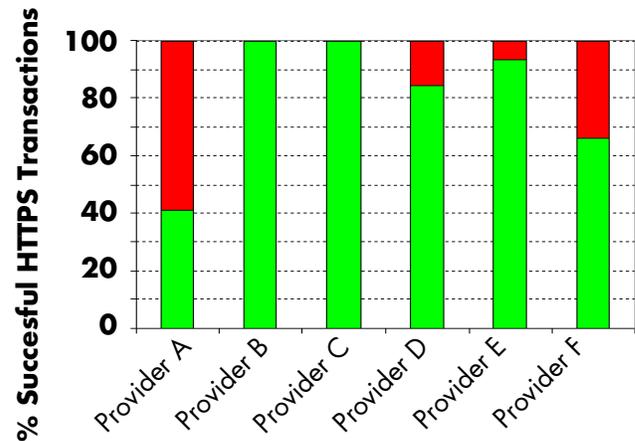
**Figure 11: Public Cloud Setup**

On each cloud provided VM we established the typical three tier architecture: front end web server accepting HTTPS connections, the application server, and an SQL database. We setup one web page that was accessible to anyone and returned a simple string, "It Works!". Another page was established with a search field, which queried the SQL database. Our SQL database was loaded with entries for equipment type objects and fields for their price, so users could search for an equipment type and the price would be returned - simple, but realistic. Finally a third URL was setup to trigger a CPU intensive task. For this we were recommended to use the graphic software POV-Ray to calculate the light reflections on a particular 3D graphic. The point was the use up CPU resources, and see how it affected our web users.

**Scalability.** Many enterprises already rely on public cloud services to host their public web sites expecting some average number of users to access their sites. In our first test we measured each cloud's HTTPS scalability. For each cloud we simulated 410 transactions per section (a user searching through our small under-10-entry SQL database) using a Spirent TestCenter running Avalanche software connected to the Internet.

1. <http://www.networkworld.com/reviews/2010/040510-cloud-computing-test.html?page=1>

We staggered the ramp up of emulated user sessions for one minute, then attempted to hold the 410 total sessions for another minute.

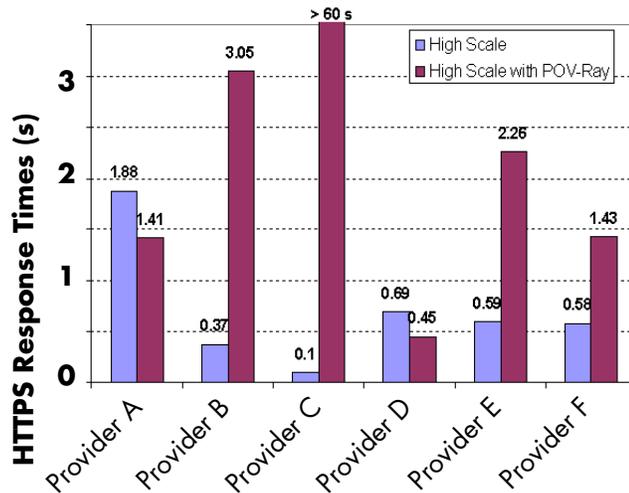


**Figure 12: HTTPS Scalability Results**

The results differed, as expected. Again, each cloud provider is offering different hardware resources, at a different cost. Still, we imagined ourselves in the place of the enterprise user, without detailed knowledge of the cloud resources, looking to host a web site with what is provided. Such a user could use this test methodology and observe that they could not realistically plan to host a web server for hundreds of users using the initial setup for some of these cloud providers. The cases resulting in a significant number of failed HTTPS transactions (see spread in Figure 12) may not necessarily be a consequence of poor resources within that public clouds data center - this may be something configured on purpose by the provider to limit any single user from taking too many resources. Nevertheless, as a user, this is a limitation that one should be aware of so it could be discussed with the provider if it is an issue.

**Performance.** The results for the scalability test showed how many users the cloud could scale up to with HTTPS sessions. We could not be sure, as a user of the cloud, where their network or their CPU could have been a limiting factor. To focus on CPU performance of the clouds, we repeated the test, but configured just 5% of the HTTPS transactions (insignificant bandwidth usage) to call a CPU intensive task - POV-Ray, a graphics program, would recalculate the light reflections across our graphic's surfaces. The same scalability test was run both with and without the POV-Ray task in parallel, and instead of measuring the number of failed sessions, we compared the HTTPS response times for each. This answered two questions; how long it take for the SQL database query to be

returned to the user, when the server was processing the POV-Ray task in parallel?



**Figure 13: HTTPS Performance Results**

The results were surprising. The clouds which excelled in our scalability test measured the longest HTTPS response times when the CPU-intensive task was initiated. To note, the response times are only of the standard SQL query transactions, and do not count the POV-Ray transactions themselves. Therefore, the results show us how heavy CPU load affects all users' response times. In conclusion, each public cloud provider may have its strengths, but there may also be some weaknesses. That said, it may simply be the preference of the majority of those clouds' customers to have CPU intensive tasks prioritized. Our methodology aims to produce a tool for such an observation to be seen by the end user.

## Conclusion

The industry is pretty excited about cloud computing and these latest advances to the data center, and understandably so. Still, this is a developing technological area, and as such there is still a ways to go.

In our private cloud tests we found that performance of data center bridges is not entirely straight forward, but they succeed in convergence. Our WAN optimizer did a great job in most cases, but may not be worth the financial investment for all user scenarios. Firewalls have progressed into the virtual space, yet the unpredictability of virtualized resources proves to still be an area of continued study.

Our public cloud computing tests have only begun to scratch the surface. The services provided us with an effective way to outsource a web server and data base function to the cloud, but what about the performance of these services in the cloud? How does the virtual machine of one user affect the behavior of another virtual machine running on the same physical server in the same public cloud? We look forward to continuing and expanding these tests. We hope to hear from anyone interested in repeating these tests, or providing suggestions for further testing.

## INTEROP Las Vegas - April, 2010



This white paper report has been distributed at the INTEROP conference in Las Vegas in April of 2010 alongside a live presentation of the tests described here. We would like to note that Spirent Avalanche Virtual, which was used in the testing, was awarded "Best of Interop" in the Performance and Optimization category by the conference organizer.



EANTC AG  
European Advanced Networking Test Center

Einsteinufer 17  
10587 Berlin, Germany  
Tel: +49 30 3180595-0  
info@eantc.de  
<http://www.eantc.com>



Spirent Communications

26750 Agoura Road  
Calabasas, California 91302  
Tel: +1-800-SPIRENT / +1-818-676-2683  
<http://www.spirent.com>

This report is copyright © 2010 EANTC AG. While every reasonable effort has been made to ensure accuracy and completeness of this publication, the authors assume no responsibility for the use of any information contained herein.

All brand names and logos mentioned here are registered trademarks of their respective companies in the United States and other countries.

20100505 v2.1