Date: 13 February 2002
Author: David Ferguson
Version: 1.1

Zeus Technology

Comparing High Performance Web Servers

Zeus Technology
Zeus House
Cowley Road
Cambridge
England
CB4 0ZT

Phone: +44 1223 525000
Fax: +44 1223 525100
Email: info@zeus.com
Web: http://www.zeus.com/

zeus

Table of contents

1 Overview

In the last few months, a number of vendors have released software billed as high performance web servers or web server accelerators and have submitted SPECweb99 benchmark results using them. There has however, not yet been a real "apples for apples" comparison carried out among these products.

SPECweb99, the second-generation web server benchmark from SPEC (http://www.spec.org), is the industry standard for web server benchmarking. All the major industry players regularly submit results, including: Sun, IBM, HP, Compaq and Dell. Zeus is a member of SPEC and through PEPP (http://www.zeus.com/about/pepp) we work closely with leading vendors on benchmarking activities. Zeus is also involved in the design of upcoming updates to the SPECweb tests.

In this report, we separate out the different elements of a SPECweb99 benchmark run and test each of the web servers with each element.

These tests show that, even with the introduction of accelerators such as Tux or X15, the Zeus Web Server still retains a significant performance advantage. Compared to the most serious contender, Redhat Tux, Zeus shows throughput gains of approximately 50% for small cached files and approximately 30% for large non-cached files.

2 Background

2.1 The Industry

2.1.1 SPEC

"The System Performance Evaluation Cooperative (SPEC) was founded in 1988 by a small number of workstation vendors who realized that the marketplace was in desperate need of realistic, standardized performance tests. The key realization was that an ounce of honest data was worth more than a pound of marketing hype.

SPEC has grown to become one of the more successful performance standardization bodies with more than 60 member companies. SPEC publishes several hundred different performance results each quarter spanning across a variety of system performance disciplines."

http://www.spec.org/spec/

2.1.2 SPECweb99

"SPECweb99 is the next-generation SPEC benchmark for evaluating the performance of World Wide Web Servers. As the successor to SPECweb96, SPECweb99 continues the SPEC tradition of giving Web users the most objective and representative benchmark for measuring a system's ability to act as a web server. In response to rapidly advancing Web technology, the SPECweb99 benchmark includes many sophisticated and state-of-the-art enhancements to meet the modern demands of Web users of today and tomorrow:

- *Standardized workload, agreed to by major players in WWW market*
- *Full disclosures available on this web site*
- *Stable implementation with no incomparable versions*
- *Measurement of simultaneous connections rather than HTTP operations*
- *Simulation of connections at a limited line speed*
- *Dynamic GETs, as well as static GETs; POST operations.*
- *Keepalives (HTTP 1.0) and persistent connections (HTTP 1.1).*
- *Dynamic ad rotation using cookies and table lookups.*
- *File accesses more closely matching today's real-world web server access patterns.*
- *An automated installation program for Microsoft Windows NT as well as Unix installation scripts.*
- *Inter-client communication using sockets."*

http://www.spec.org/osg/web99/

Since SPECweb99 was ratified, all the major industry players have submitted benchmark results, including: IBM, Sun, HP, Compaq and Dell. The Zeus Web Server has been the server software of choice for UNIX based submissions throughout SPECweb99 and its predecessor (SPECweb96).

2.1.3 Zeus and SPEC

Zeus has worked closely with SPEC participants since 1996. In 2000, we set up the PEPP (http://www.zeus.com/about/pepp) initiative to further our relationships with leading vendors. Current members include IBM, Sun and HP. Through this relationship we work with vendor engineering teams to evaluate and support new hardware and operating system features.

This close relationship proves the scalability, performance and stability of the Zeus Web Server and the vendor product. The demanding nature of web serving and the proven stability of Zeus products means that Zeus is used in wider circles as part of CPU, system or platform verification testing.

2.2 The Benchmark Lifecycle

The industry is moving into the final stages of life of SPECweb99. Typically, each benchmark moves through the following stages:

Definition

A group of companies come together to produce a standard method by which their various products can be compared. During this time, each vendor tries to implement the emerging standard in order to ensure that the path that is being taken doesn't represent a competitive risk.

Initial Results

Sooner after ratification of the benchmark, vendors start to make their first tentative submissions. This is a chance to really get to grips with the benchmark and to see how competitors are approaching the benchmark

Serious Results

This is the longest part of the benchmark lifecycle. The vendors ramp up their submissions until they are testing their very latest, fastest hardware. Every 6 to 12 months, one vendor leapfrogs the others by setting a score 50%-100% higher than previously seen. Soon afterwards, the others respond with similar scores.

Optimisation

At this stage, the vendors start looking for optimisations that can be made to the hardware, operating system and software that will further boost performance. Benchmark rules are carefully phrased so that optimisations must benefit all workloads (and not just that used by the benchmark). This phase drives considerable innovation, particularly in software. The results of which are high-performance hardware, operating systems and servers for the whole market.

During the history of SPECweb (SPECweb96 and SPECweb99), kernel accelerators have both appeared during this phase. There are claims that they provide substantial performance gains but commercial adoption is very slow. This document evaluates the actual improvement.

At this stage, the vendors realise that the benchmark is in need of revision and start developing the next generation.

2.3 Kernel Accelerators

2.3.1 Usage

Kernel Accelerators have featured in the lifetime of both SPECweb96 and SPECweb99. SPECweb99 has seen the introduction of the following accelerators:

- RedHat Tux
- Sun NCA

- IBM FRCA (Fast Response Cache Accelerator)
- Microsoft SWC (Scalable Web Cache)

Kernel accelerators work by attempting to cache and serve as much content as possible without entering a user level process. In the case of SPECweb99, they generally handle the serving of static content but pass requests for dynamic content onto a user level web server.

The Zeus Web Server supports all of these accelerators.

Limitations

Most kernel accelerators are not full blown web servers. Today, they are generally only capable of directly serving unencrypted static content.

Market Acceptance

There is little evidence of commercial usage of kernel accelerators. According to the latest Netcraft Survey, the most widely known, RedHat's Tux, runs only 49 sites[1].

Technical Issues

One of the most significant reasons why kernel accelerators have not, and probably will not, be used significantly in commercial environments is that by their very nature, they present a risk to the stability of the operating system kernel. The enviable stability of Linux and UNIX kernels comes from the fact that the kernel is kept small, evolves slowly, and implements only the most necessary functions. By adding commercially unproven components, that will require very large amounts of memory, corporates and operators place the stability of their business critical web servers at risk.

2.4 The Web Servers

2.4.1 Zeus

"Zeus Web Server is the most scalable, high-performance web server software available, underpinning business-critical solutions for the world's leading web hosting, content provider and secure e-commerce companies. Flexible web-based management, extensive integration capabilities, and the most comprehensive range of features available from commercial server software, combine to provide fully extensible and future-proof web-based solutions."

http://www.zeus.com/products/zws/

Zeus Web Server has always and continues to dominate web server benchmarks. We work closely with IBM, HP and Sun engineers to continuously improve the performance of our core technology. Regular submissions of SPECweb99 results on Zeus prove the stability and scalability of our software.

Recent work with HP on Itanium based servers highlights the benefits that Zeus software can provide in e-Commerce environments, see: http://www.zeus.com/library/technical

Zeus Web Server 4.0 was used for these tests.

[1] http://www.netcraft.com/Survey/Reports/0109/byserver/TUX/index.html

2.4.2 Apache 1

"Apache has been the most popular web server on the Internet since April of 1996. The February 2001 Netcraft Web Server Survey found that 60% of the web sites on the Internet are using Apache (around 62% if Apache derivatives are included), thus making it more widely used than all other web servers combined.

The Apache project is an effort to develop and maintain an open-source HTTP server for various modern desktop and server operating systems, such as UNIX and Windows NT. The goal of this project is to provide a secure, efficient and extensible server which provides HTTP services in sync with the current HTTP standards."

http://httpd.apache.org/

Apache is notoriously unrepresented in web server benchmark results. This is due to a fundamental limitation in the Apache architecture that limits the number of simultaneous connections that the Apache web server can handle. For information on Zeus and Apache scalability testing, see: http://support.zeus.com/doc/tech/perf.pdf

Apache 1.3.17 was used for these tests.

2.4.3 Apache 2

At the time of writing, it appears that Apache 2 is about to leave the beta phase. Covalent have announced that Apache 2 based products will ship in December 2001.

Apache 2 supports a number of concurrency models. In these tests we carried out tests using the preforked, multi-threaded and worker models.

Apache 2.0.28 was used for these tests.

2.4.4 Tux

"TUX is a kernel-based, threaded, extremely high performance HTTP server. It is able to efficiently and safely serve both static and dynamic data. TUX moves the HTTP protocol stack to the kernel, and can handle requests for data with both kernel-space and user-space modules.

The TUX 2.0 release is an incremental upgrade to TUX 1.0 and keeps source-code level compatibility with user-space modules."

http://www.redhat.com/docs/manuals/tux/

As yet there is no evidence of commercial acceptance of the Tux solution. The September 2001 Netcraft survey[2] shows only 49 sites running Tux.

Tux 2.0 was used for these tests.

2.4.5 X15

"Chromium X15 WebServer is the world's fastest Linux web server. X15 makes a clean break from web servers of the past, featuring a new "from the ground up" server architecture. Designed to exploit the latest Linux kernel advances, X15 has been optimized for high performance on SMP and new "high density" server hardware. X15 does not tradeoff resource efficiency for

[2] http://www.netcraft.com/Survey/Reports/0109/byserver/TUX/index.html

speed; with a small memory footprint and using only a small number of threads to achieve it's performance. Using industry standard benchmarks, X15 routinely outperforms Apache by a factor of five."

http://www.chromium.com/products.html

X15 was due for release in August 2001. There is no publicly available evidence to support the "factor of five" claim, but given the architecture of Apache, it is not unlikely that a "factor of five" gain could easily be obtained under a SPECweb99 load.

A pre-release copy of X15 was used for these tests.

2.4.6 Web Server Comparison

The following table summarises the architecture, availability, usage and features of the above web servers:

Feature	Zeus	Apache[3]	Tux	X15
Kernel/User Space	User	User	Kernel	User
First Release	July 1995	April 1995[4]	September 2000[5]	August 2001[6]
Sites powered[7]	783,261	19,279,109	49	0
Highest published SPECweb99 score	15,000	No submissions	6,248	2,800
Dynamic Content APIs	CGI, FastCGI, ISAPI, NSAPI, ZDAC, Jserv	CGI, FastCGI, Jserv	CGI, forward to secondary server	CGI, forward to secondary server
Native SSL Support	Yes	No – but readily available	No	No
Web Based User Interface	Yes	No	No	No
Platforms Supported	HPUX (PA-RISC, IA 64) Solaris (IA32, Sparc), AIX Linux (IA32, IA64, Alpha, MIPS, PPC) IRIX Tru64 FreeBSD OpenBSD SCO MAC OS X BSDi	All major Linux and Unix Netware Windows	Linux (IA32)	Linux (IA 32)
Native clustering support	Yes	No	No	No

[3] Apache 2 has not yet been released and there have been no benchmark submissions using it. Therefore we just cover Apache 1 here.
[4] http://httpd.apache.org/ABOUT_APACHE.html
[5] http://boudicca.tux.org/hypermail/linux-kernel/2000week36/0780.html
[6] http://www.chromium.com/x15.html
[7] From the September 2001 Netcraft survey, see http://www.netcraft.com/survey/ - total sites surveyed 32,398,046

3 Methodology

3.1 Test Hardware

Tests were carried out using a single 1U Linux client and server. Each machine has a 700MHz Pentium III processor, 256MB RAM and a 100Mbps Ethernet connection.

Note: whilst the design of this test does result in client and network throughput bottlenecks, it does represent test setups that can be readily recreated on commonly used web serving hardware.

3.2 Test Software

- Apachebench – the benchmarking tool supplied with the Apache distribution was used to generate a load on the servers.

The following web servers were used - in each case "out of the box" configurations were used:

- Zeus 4.0r1
- Apache 1.3.17, 2.0.28
- Tux 2.0
- X15 pre-release 1.0

3.3 Test Workloads

A SPECweb99 benchmark primarily tests the web server's ability to serve three different types of content:

- Serving in-memory cached content. On most recent HP, IBM and Sun results, the machines have been equipped with enough RAM to hold the document root for the entire test.

- For IA32 based Linux systems, the 4GB memory limit means that only a subset of the document root can be cached in memory. In this case, the performance of the disk system and the efficiency of the operating system in streaming content are vital. While technologies like PAE (Page Address Extension) can break through the 4GB limit, there is no evidence of vendors using this feature.

- The SPECweb99 benchmark also involves a substantial number of requests for dynamically generated pages.

The disadvantage of a SPECweb99 result is that it only provides a single point number. In these tests, the SPECweb99 test has been "unbundled" into the three different types listed above at a set of different content lengths. This information can then be used to estimate the actual performance differences for different workloads.

The workloads were defined as follows:

- Serving in-memory cached content
  - These tests used a range of content lengths from 1 to 10,000,000 bytes (SPECweb99 uses content lengths from 100 bytes to 900,000 bytes)

- Serving non-cached content from disk
  - These tests used sample files from the SPECweb99 document root ranging from 100 to 800,000 bytes

- Serving dynamic content
  - These tests used the vendor supplied (Zeus or Tux) implementations of the SPECweb99 dynamic components. These generated a range of content lengths from 345 to 205,050 bytes. The sources for these components can be found at:

    - http://www.spec.org/osg/web99/results/api-src/HP-20010917-API.tar.gz (Zeus)

    - http://www.spec.org/osg/web99/results/api-src/Compaq-20010709-ml370-tux.tar.gz (Tux)

All content was stored on local hard disks.

3.4 The SPECweb99 Content Distribution

Each test consists of testing the performance of the web servers across a range of content lengths. While in some cases, performance at one end or another is of particular interest, it is useful to be able to condense these to a single value. We define the "weighted throughput" as being the sum of the throughput at a particular content-length multiplied the percentage of SPECweb99 requests that would be made at that throughput.

The following table shows the SPECweb99 distribution of content lengths. Within each class, the distribution is assumed to be uniform.

Class	File Size	%
0	< 1k	35
1	< 10k	50
2	< 100k	14
3	< 1000k	1

4 Results

4.1 Cached Content

In theory, serving of cached content should provide the highest throughput for each server. Comparing these results with the charts below, it is clear that there is a substantial performance improvement to be obtained from caching content.

This graph shows that while there are less significant differences above the 10,000 byte point, there are very substantial performance differences in the 50 to 1,000 byte range. This band is significant for two very frequent request types:

- Requests for small GIF and PNG files
- Verification by the client that content has not changed (304 Not Modified responses)

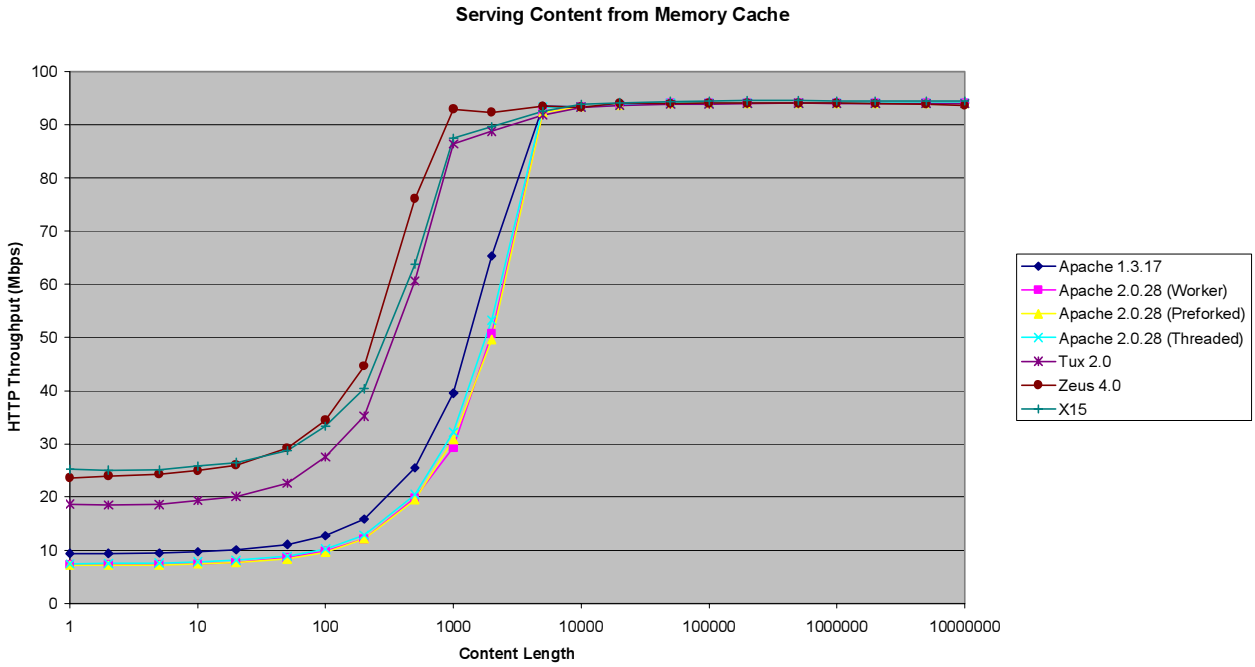

For these tests, the server only has a single 100Mbps Ethernet interface and this limits the HTTP throughput to around 95Mbps.

Within the 50 – 1000 byte range, Zeus provides a 150% to 200% increase in raw throughput over Apache.

This demonstrates the advantage that Zeus gives in raw throughput over all of the competition. However, in real-world situations there are certain types of requests that are far more common than others. The *de facto* SPECweb99 benchmark reflects this reality, and weights different content lengths based upon how important it is for most websites.

By applying the same weightings as used by SPECweb (see Appendix A) we can demonstrate that Zeus is significantly faster than Apache for the content that your website is likely to serve. This means that your customers will get faster responses when requesting pages, and your hardware will be capable of delivering this higher performance to more customers simultaneously.

Server	Weighted Throughput (Mbps)	Percentage (of Zeus)	Percentage (of Apache 1)
Zeus 4.0	87.65	100.00	125.56
X15	85.13	97.13	121.95
Tux 2.0	83.84	95.65	120.10
Apache 1.3.17	69.81	79.65	100.00
Apache 2.0.28 (threaded)	67.12	76.57	96.15
Apache 2.0.28 (preforked)	66.35	75.70	95.04
Apache 2.0.28 (worker)	66.21	75.54	94.83

This table shows that when the SPECweb99 weightings are applied to the traffic that is being served, Zeus still has a significant advantage over all versions of Apache.

Notes:

1. In all cases, higher throughput figures show better test scores.

2. Further tests have shown that at the low end (1 byte content), the client can become the bottleneck. The addition of a second client can lift the Zeus figures by around 50%, and will raise the X15 figures by a smaller amount. In the case of Tux and Apache, no benefit is observed.

4.2 Non-Cached Content

This test represents the most demanding tests for these systems. Here, the web server is forced to visit the disk for each and every request. Recent SPECweb99 Tux 2.0 submissions show that very large and fast disk systems have been necessary in order to deliver the necessary throughput. These results show that while performance is similar below the 10,000 byte zone. As content lengths continue to increase, the gap widens between Zeus and Tux.

Note: during these tests each disk on file was requested only one time to ensure that neither the web server of the operating system could cache information.

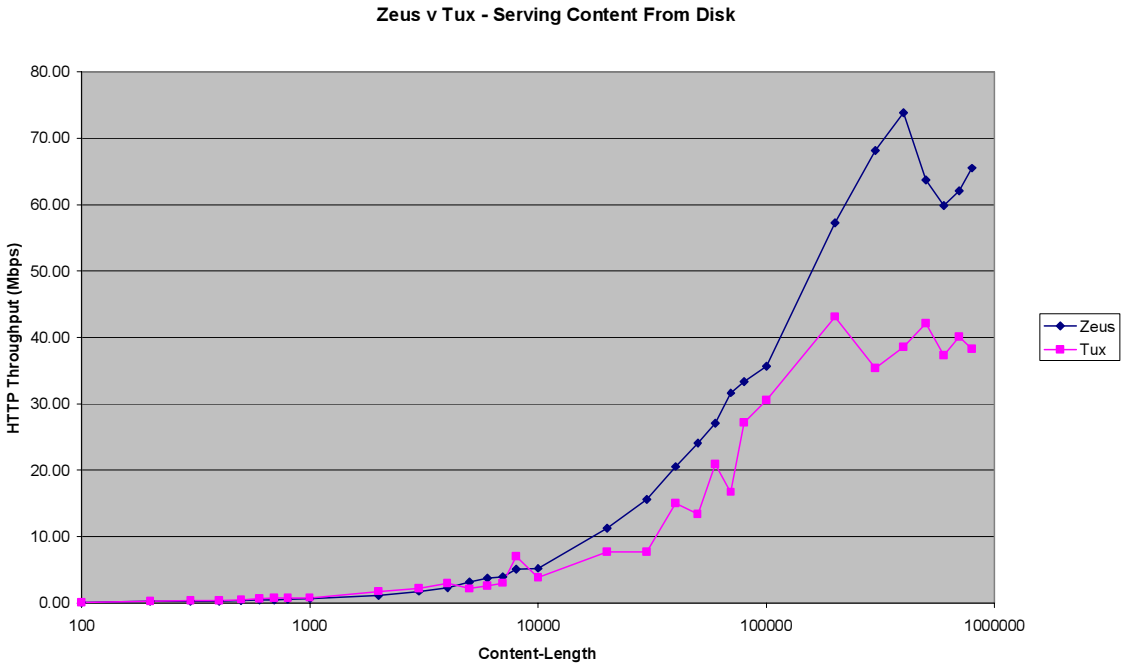

Server	Weighted Throughput (Mbps)	Index (Zeus=1.00)
Zeus	6.06	1.0000
Tux	4.77	0.7878

4.3 Dynamic Content

The final set of tests examined the performance of the Zeus and Tux SPECweb99 dynamic components. SPECweb99 uses 4 different types of dynamic request; this test uses the Dynamic GET request, as this is the most common request in the suite. These tests show marginal difference in performance, indicating perhaps that this test has been more network I/O bound (they reach the 95Mbps limit around 2K content) rather than CPU bound.

The SPECweb99 dynamic components always generate content in excess of 300 bytes, therefore it is not possible to show results below this length.

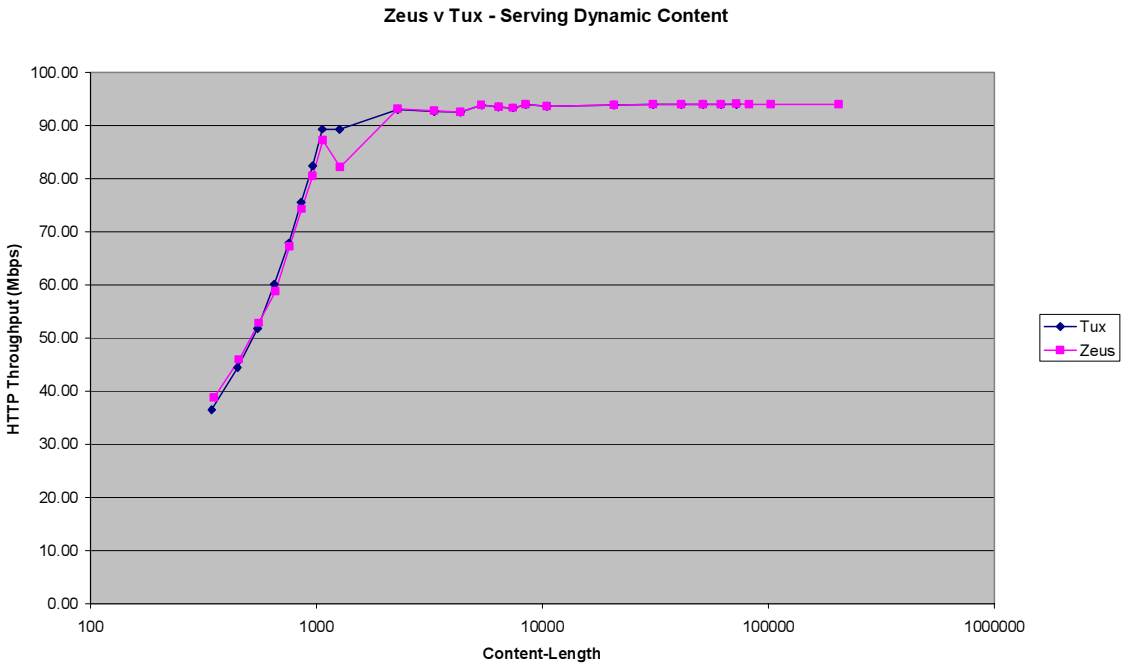

Server	Weighted Throughput (Mbps)	Index (Zeus=1.00)
Zeus	80.47	1.0000
Tux	79.86	0.9925

5 Conclusions

By unbundling the various elements of the SPECweb99 benchmark, this report shows that Zeus outperforms all other web servers in each component of the benchmark and shows therefore that were Zeus to be used on exactly the same platform as a Tux or X15 submission was made, that higher results would be obtained by Zeus.

Apache 2 has been promoted as a faster more scalable web server. These tests show that Apache 2 cannot yet achieve the performance of Apache 1.

As a user-space application, the Zeus Web Server does not present the risk of kernel accelerators such as Tux and, with an established customer base, the Zeus Web Server is commercially proven with customers such as eBay and UUNET.

If the business requirement is for rock solid high performance scalable web serving, with Zeus it isn't necessary to trade off stability against performance or features.

6 Contact Information

6.1 Zeus Products and Services

For more information on the Zeus Web Server, visit:

http://www.zeus.com/products/zws/

To see what our customers say about the Zeus Web Server, visit:

http://www.zeus.com/customers/testimonials.html

For more information about SPECweb99, visit:

http://www.specbench.org/osg/web99/

For more information about our products, please visit:

http://www.zeus.com/products/

To find out more about our consultancy, technical support and training services, please visit:

http://www.zeus.com/services/

6.2 Office Contact Details

Zeus Technology Ltd.
Zeus House
Cowley Road
Cambridge
CB4 0ZT
UK

Tel.: +44 (0)1223 525000
Fax: +44 (0)1223 525100

Zeus Technology Inc.
5201 Great America Parkway
Suite 340
Santa Clara
CA 95054
USA

Tel.: +1 408-350-9400
Fax: +1 408-350-9408

Appendix A – SPECweb99 Benchmark

As benchmark results are not usually typical of real life performance, it is useful to interpret the data to answer the question "how well will my system perform in real life?"

We calculated the servers "quality factor", Q as being the weighted sum of all the data points on the throughput graph: if you choose the weights to reflect real life page requests then this will give a fair indication of the normal real life throughput.

The weightings used for our calculations are those which are used in the SPECweb99 benchmark:

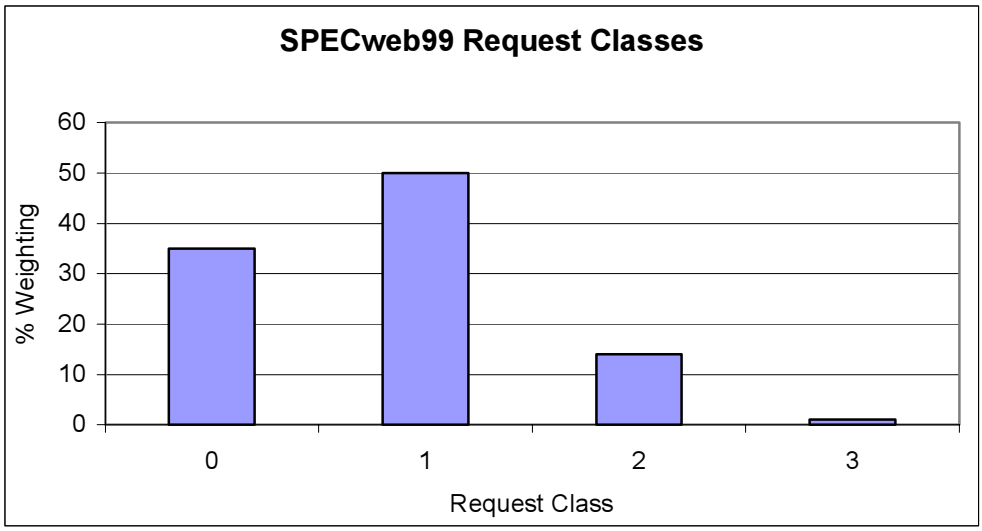

Class	File size	% Weighting
0	< 1k	35
1	< 10k	50
2	< 100k	14
3	< 1000k	1