

# Measurement, Analysis and Performance Improvement of the Apache Web Server \*

Yiming Hu<sup>†</sup>, Ashwini Nanda<sup>‡</sup>, and Qing Yang<sup>†</sup>

<sup>†</sup> Dept. of Electrical & Computer Engineering  
University of Rhode Island  
Kingston, RI 02881  
{hu,qyang}@ele.uri.edu

<sup>‡</sup> IBM T.J. Watson Research Center  
P.O.Box 218  
Yorktown Heights, NY 10598  
ashwini@watson.ibm.com

## Abstract

*Performance of Web servers is critical to the success of many corporations and organizations. However, very few results have been published that quantitatively study the server behavior and identify the performance bottlenecks. In this paper we measured and analyzed the behavior of the popular Apache Web server on a uniprocessor system and a 4-CPU SMP (Symmetric Multi-Processor) system running the IBM AIX operating system. Using the AIX built-in tracing facility, we obtained detailed information on kernel events and system activities while running Apache driven by the SPECweb96 and the WebStone benchmarks. After quantitatively identifying the performance bottlenecks, we proposed and implemented 6 techniques that improve the throughput of Apache by 61%. These techniques are general purpose and can be applied to other Web servers as well.*

## 1 Introduction

With the explosive growth of the World Wide Web, more and more corporations and organizations are depending on high performance Web servers for the success of their business. However, there are very few published results addressing the issue of quantitatively characterizing the behavior of Web servers and improving their performance. In this paper, we measured and analyzed the behavior of the Apache Web server running on several different hardware platforms. We show that Apache spends more than 75-80% of its CPU time on OS kernels. We also give detailed measurement results on where the kernel time is spent. Based on our performance analyses, we present 6 techniques to improve the performance of the

Apache server. Together, these techniques increase the throughput of Apache by more than 61% under the SPECweb96 workload.

## 2 Apache, SPECweb96 and WebStone

### 2.1 The Apache Web Server

Apache [1] is a freely available, UNIX-based Web server developed by a team of volunteers. It was originally based on code and ideas found in NCSA httpd 1.3, but has been completely rewritten since then. We choose Apache as the base of our experimental system for the following three reasons. First, Apache is the most popular Web server running today, accounting for more than 53% of all Web domains on the Internet [1]. Second, Apache is a fully featured, high performance Web server, superior than many other UNIX-based Web servers in terms of functionality, efficiency and speed. Third, the source code of Apache is available, enabling us to make changes to the code to improve its performance. We used Apache 1.2.0 during the experiment.

### 2.2 The SPECweb96 Benchmark

SPECweb96 was developed by the Standard Performance Evaluation Corporation (SPEC). It is the first standardized benchmark for measuring the performance of Web servers.

SPECweb96 consists of two parts: a *file set generator* and a *workload generator*. Before testing a Web server, the file set generator must be run in the server machine to populate a test file set consisting of many files of different sizes. The workload generator then runs on one or several computers connected to the Web server machine being tested via a TCP/IP network. It simulates Web clients by sending HTTP "GET" commands to the Web server, requesting for files in the test file set. The workload generator measures the

\*This work is supported in part by NSF under Grants MIP-9505601 and MIP-9714370, and IBM. Part of this work was performed at IBM T. J. Watson Research Center.

Classes	File Sizes	Access Frequencies
Class 0	0 – 1KB	35%
Class 1	1KB – 10KB	50%
Class 2	10KB – 100KB	14%
Class 3	100KB – 1MB	1%

Table 1: File Sizes per Class and Frequency of Access

response time of the Web server, and increases the request rate until the server can not handle them. The maximum HTTP request rate that the Web server can handle is the SPECweb96 value.

Many studies indicate that small files are accessed more frequently than large files in most real world Web servers [2, 3, 4]. Similarly, in SPECweb96, the data files are classified by their sizes into 4 classes, as shown in Table 1. Smaller classes are accessed more frequently by the workload generator than larger classes are. The file sizes and access frequencies are based upon analysis of logs from several popular servers, including NCSA’s site, the home pages for Hewlett-Packard and HAL Computer.

In the SPECweb96 file set, Class 3 files account for 90% of the total file set size. However their chances of being accessed are only about 1%. The remaining files account for only 10% of the total file set size, yet receiving 99% of total accesses. This leads to an important conclusion. That is, a relatively small amount of RAM (say 128 MB) can be used to cache most files in the file set and eliminate most disk accesses. As a result, for such systems, disk activities have little effect on overall server performance.

### 2.3 The WebStone Benchmark

SPECweb96 is a standardized benchmark that generates relatively realistic Web workloads for evaluating the overall system performance. It mixes 4 different classes of files in a way close to the access patterns observed in real world Web servers. However, researchers often need to change the characteristics of the workload, such as the request sizes, to identify the server performance bottleneck under a specific condition. SPECweb96 does not allow users to change the workload. Therefore, when we need to study the behavior of Apache under different request sizes, we will use the WebStone benchmark developed by SGI, which gives users almost complete control over the workload characteristics, including the request sizes and mixtures.

One problem of WebStone is that its default file generator populates the server with only 32 files. If

Server Machine (1 RS/6000)	Model	43P-140
	CPU Type	200 MHz PPC 604e
	RAM Size	128 MB
	Disk Space	2 x 2 GB
	OS	AIX 4.2.1
Client Machines (2 PCs)	Model	Pentium PC
	CPU Type	133 MHz Pentium
	RAM Size	32 MB
	Disk Space	2 GB
	OS	Linux 2.0.30
Network	2 x 100 Mbps Ethernet	

Table 2: Uniprocessor Test System Configuration

we use these files as the server file set, the server will transfer a very small number of files over and over again. This may generate misleading results, since the small number of files will always stay in the file cache and the CPU cache. Fortunately, WebStone allows using any data files. We choose the same SPECweb96 data files, which is large enough (several hundreds megabytes), as the Server data file. We setup WebStone to selectively request the SPECweb96 files on a specific size range.

## 3 Experimental Environments

We studied the behavior of the Apache server on both a uniprocessor system and a SMP multiprocessor system. The uniprocessor experimental system consists of a IBM RS/6000 system as the server and two Pentium PC systems as clients, as shown in Table 2. Each client machine has a dedicated 100 Mbps Ethernet connection to the server. The SMP experimental system consists of two state-of-the-art IBM 4-Way SMP machines with 2 GB of RAM each, one as the server and the other as a client, connected through two dedicated 100 Mbps Ethernets.

We use the AIX built-in trace facility to capture the system activities while Apache is running. The trace facility records kernel events with an extremely fine granularity of details [5]. The collected trace data is post-processed by the UTLD 1.2 program [6] developed by IBM RISC System/6000 Division. UTLD analyzes the trace file and generates a detailed system utilization report including the CPU utilization, locking time, interrupt handling time, as well as CPU times of individual system calls. For better measurement accuracy, we repeat each measurement 3 times and average the results.

## 4 Measurement Results

### 4.1 Detailed CPU Time Breakdown

Figure 1 gives detailed CPU time breakdown for Apache running on the uniprocessor system. Note that because the system has sufficient amount (128 MB) of RAM, most files are cached in the system memory so the disk is not the bottleneck. In fact, the CPU utilization is 100%, meaning the CPU is the bottleneck.

As shown in Figure 1, Apache spends only about 23% of its CPU time on the user code. The rest of the CPU time is spent in kernel. Among the kernel activities, the TCP/IP stack and the low level network handling consume the majority of the CPU resource. The code directly involved with TCP/IP (the dark sectors on the bottom, including the *Ethernet Driver*, the *write*<sup>1</sup>, *select* and other TCP/IP system calls) takes 29% of the total CPU time. The first level interrupt handlers (the light sectors on right, including *I/O Interrupt*, *data page fault* and *other interrupts*) also use a significant amount (24%) of CPU time. Among them, the *I/O Interrupt* portion is mainly for handling Ethernet interrupts. Since most disk files are cached in RAM, there are very few disk activities observed. The TCP/IP and the Ethernet interrupt handlers together take 43% of the total CPU time, which is the main performance bottleneck. File system operations (the light sectors on the left, including the *read*, *open*, *stat* and other system calls) use 17% of the CPU time, which is very wasteful because more than 99% of the active data files are cached in the system RAM. Finally, the “incinterval” system call that is mainly used by the *alarm* and other timer functions uses 3% of the processing time. All other miscellaneous syscalls, which include signal handling and system calls such as *getpid* and *exit*, take only about 2% of the total CPU time together.

We have also run Apache on a 4-CPU SMP system with 2 GB of RAM as a server and SPECweb96 on another SMP system as a client to generate HTTP requests. We found that the behavior of Apache on a SMP system is very similar to that on a uniprocessor one. Interested readers can refer to our technique report [7] for a detailed description of the behavior of Apache on the SMP machine.

### 4.2 Effects of Request Sizes

SPECweb96 generates Web requests with mixed request sizes. While it gives us a balanced view on the

<sup>1</sup>A very small percentage of the write time is actually caused by writing the log file, which is a file system activity. Similarly, a small portion of the *read* time (about 10%) is caused by read requests from the network, although we classify *read* as a file system operation.




Figure 1: Detailed Apache CPU Time Breakdown (128MB)

overall system performance and behavior, it is often desirable to study the system behavior under different request sizes. The studies of ours and others [8] indicate that different request sizes stress the system in very different ways.

To study the behavior of Apache under different request sizes, we conducted another 4 groups of tests on the uniprocessor system, using the WebStone benchmark. For the reasons described in Section 2, we use SPECweb96 data files as the test file set for WebStone. The SPECweb96 files are classified according to their sizes with class 0 being the smallest sizes and class 3 the largest sizes. In each group of the tests we let the WebStone request about 100 different files from a specific SPECweb96 file class. As a results, we are able to measure the behavior and identify the problems of Apache under different file sizes.

Figure 2 summarizes the results of the Apache behavior under different request sizes. This figure clearly shows that when the request sizes are small (class 0 and class 1 files), user time dominates the total CPU time. The file system and the TCP/IP stack are relatively equally stressed. For large requests (class 2 and class 3 files), Apache spends most of its time on the TCP/IP stack. Detailed CPU time breakdowns for different file sizes can be found in our technical report [7].

## 5 Improving the Apache Performance

We have shown that for the SPECweb96 workload, on a uniprocessor system Apache spends about 23% of the total CPU time on user code, 26% on inter-




Figure 2: Apache CPU Time for different file sizes

rupt handlers, 29% on TCP/IP processing and 17% on file system operations. It is difficult to reduce the overheads caused by TCP/IP processing and interrupt handling unless we modify the OS code or change the hardware. On the other hand, it is possible to greatly reduce the file system overhead and the user code overhead, as will be demonstrated in the remaining of this section.

### 5.1 Using the *mmap* function

For our system with 128 MB of RAM, almost all frequently accessed file data are cached in the file system cache. However, Figure 2 clearly indicates that for all 4 different classes of files, Apache still spends a considerable amount (17-19%) of its CPU time on file system activities. The reason to this high overhead is that for almost every HTTP request, Apache has to check the file state using the *stat* system call, *open* the file and *read* the file data from the file system cache into the user buffer, before it can ship the data to the network.

Our first approach is to use the *mmap* function to eliminate the data copying between the file system cache and the user space. When Apache needs to read a file, we let the program open the file, then use the *mmap* function to map the data of the file into the user space. As a result, the server can send the mapped data directly to the network, avoiding the *read* system call altogether. When the entire file content is sent out, the file is unmapped and closed to limit the number of opened files in the system (most OSs pose a limitation on the number of files a process can open).

This approach is easy to implement. However, we found that the effectiveness of this approach is limited for the following two reasons. First, for small

file sizes (class 0 and class 1 files), we found that the *read* system call accounts for only 30-50% of the total file system overhead. Other file system calls that can not be eliminated by *mmap*, such as *open* and *close*, also have significant impacts on the server performance. Second, for each mmapped pages accessed by the Web server, the Virtual Memory system generates a page fault, even the file data are already cached in the kernel buffer. Since the file is unmapped and closed after the data are sent out, a sequence of page faults will occur again when the file is re-mapped into the user space next time. Because of the overhead of page faults, reading data from the mmapped area is about 20-100% slower than from the user memory space in most OSs, as reported by McVoy and Staelin [9]. Our own experiments also confirmed their observation. The *lmbench* benchmark reports that our system has a memory bandwidth of 72 MB/sec for reading data from the user space. The memory bandwidth is 59 MB/sec for reading data from mmapped areas, which is about 22% lower than the former. We will further discuss the performance impact of the *mmap* overhead shortly.

### 5.2 Caching files in the User Space

A better solution is to cache the file data in the Web server user space. A Web server can directly ship the data of cached files to the TCP/IP stack, avoiding all file system calls such as *open*, *read* and *close*. Moreover, as discussed above, reading data cached in the user memory space is faster than from the mmapped area, resulting in a better server performance.

To evaluate the performance benefit of caching in the Web server user space, we conducted the following experiment. we let the WebStone repeatedly retrieve a single 1 MB file from the Apache server. Because of the large file size, the overheads caused by *open* and *close* system calls are negligible. We have already found that in our system, reading data from the user space is about 22% faster than from a mmapped area. Our profiling results show that when transferring large files, Apache spends most of the time on the *write* function, which reads data from the user space or from the mmapped area and writes the data to network buffers. Therefore we expected that the server bandwidth of Apache caching data in the user space should be noticeably higher than that of Apache using *mmap*. We obtained an average server bandwidth of 87.24 Mb/sec by caching the file in the user space, as compared to 78.30 Mb/sec if we use *mmap* to fetch the file data, giving rise to almost 11% improvement.

Our profiling data also indicates that Apache frequently calls the *stat* system call to get the file states,

such as the access permission and create-time information. For small file sizes, the *stat* system call accounts for about 20% of the total file system overhead. It will be beneficial to cache the file state in the Web server user space also. Since the file state information is small, the overhead for caching this information is minimal.

We have implemented a user-level cache for Apache. Because of the space limitation, we will not discuss the implementation details in this paper. Interested readers can refer to our technique report [7] for more information. Basically, the cache stores both the file data and the file states in a user memory region shared by all processes. When the server tries to read the content or state of a file, the cache is searched first. On a cache hit, information stored in the cache is returned immediately. On a cache miss, the cache invokes system calls to obtain data, put the data in the cache, and return them to the server.

### 5.3 Speeding up Logging

Typically a Web server logs every HTTP request into a log file. The logging information includes the client name, the request time, the requested URL, etc. During our experiments we found that the logging process of Apache is a major overhead. For example, on the uniprocessor system with 128 MB of RAM, the SPECweb96 number of Apache with logging is 158 ops/sec. Simply turning off the logging operation results in a SPECweb96 number of 194 ops/sec, which is a 22% improvement. Since logging information is very useful for many Web sites, it is desirable to reduce the logging overhead of Apache.

#### 5.3.1 Caching DNS results

We found that the majority of overheads of logging comes from looking up the host names of clients. Apache uses the *getnamebyaddr* function to perform the DNS (Domain Name Service) lookup using client IP addresses as inputs. Apache calls the function for every HTTP request being logged. This is very wasteful, since users normally send out a sequence of requests to a Web server for multiple objects in a Web page (the HTML file and many small bitmap files). As a result, multiple name lookup operations are performed for the same IP address. The overhead is especially high if there is no name server running at the local machine. In such a case, the name lookup requests must be sent across the network to a name server, causing long delays and extra network traffic. For small file sizes, the logging delays caused by DNS lookups may dominate the HTTP request response time. While it is possible to disable the host name

lookup processes and log only the client IP addresses, most Web administrators would prefer logging host names, because it provides much more information.

We solve the problem by using a simple and effective technique — A DNS cache that caches the host names and IP addresses of clients in the Web server address space. The DNS cache is a small array (about several hundred entries) of records that contain the host names and IP addresses of client machines. The cache entries are indexed by hashing the IP address of machines. When Apache needs to lookup a client's name with an IP address, the DNS cache is checked first. If there is a cache hit, the machine name is returned from the cache. Otherwise the cache calls the *getnamebyaddr* function first to find the machine name, puts the information into the cache, then returns the machine name. Our measurement results show that this simple solution improves the throughput of Apache by more than 14%. This is equivalent to reducing the logging overhead by 63% (14/22), since the logging overhead reduces the server throughput by 22%.

#### 5.3.2 Caching String Results

We ran Apache with the *gprof* profiling tool and found that another noticeable portion of the logging overhead is caused by logging the request time and status code. For each request, Apache has to convert the current system time into a human-readable ASCII string and then write the string into the log file. Similarly, it has to convert the request status code — success or fail, for example — into an ASCII string. Both string conversion operations are expensive.

However, many of the string operations are redundant. For example, the logging time resolution is only 1 second. As a result, in any second, Apache calls the time conversion routines hundreds times, only to convert the same time into the same string again and again. Similarly, the majority (80-90%) of status codes during a normal operation period should be “successful” [2]. It is wasteful to convert the same “successful” status code into the same string over and over again.

The overhead can be significantly reduced by caching the last string conversion result. For example, we let the time-conversion routine store the resulting string in a static array. When Apache calls the conversion routine again, the routine returns the result in the static array if the current time argument is the same as the last call, thus avoiding many redundant string operations. Similar optimizations can be applied to the status code conversion and several other places in the Apache code. This requires changing only several

lines of code. The result is a 20% of reduction on the logging overhead, or a throughput improvement of about 4 percent.

#### 5.4 Caching URI Processing Results

We have shown that Apache spends about 20-23% of its CPU time on user code. Using the *gprof* tool, we found that about 60% of the user time are for processing URIs (Uniform Resource Indicators, the parts of URLs after the colon), e.g., parsing, directory checking, security checking, translating the URI to a file name, etc. The URI parsing overhead becomes a major performance bottleneck after the file system overhead is eliminated by the user-level file cache.

The problem here is very similar to the one faced by physically addressed CPU caches. A physical CPU cache can eliminate the overhead of transferring data between the main memory and the CPU on cache hits. However, the CPU needs to translate virtual addresses to physical addresses before accessing the cache. The address translation process is potentially very expensive, involving accessing page tables in the RAM for each memory access.

Almost all modern CPUs use TLBs (Translation Look-aside Buffers) that significantly reduce the address translation overhead. A TLB caches the translated physical addresses to speed up address translation. We adopted the idea of TLB into the Web server design. The URI parsing process resembles the process of translating a virtual address (a URI) to a physical address (a file name). This similarity leads us to the design and implementation of a URI cache that greatly speeds up the URI processing of Apache. When Apache finishes parsing and checking a URI and obtains a file name corresponding to the URI, the URI/file-name pair is put into the cache. Later when a client send the same URI to the server, the server can obtain the file name corresponding to the request directly from the cache and send the data out, without going through the time-consuming parsing and checking process again. Note that because Apache supports the concept of “Virtual host”<sup>2</sup>, the same URI may result in different file names. Consequently, each virtual host has its own URI cache. Similarly, different URIs may correspond to the same file, although it is easy to handle this situation.

This scheme is especially helpful for dynamically generated Web pages such as directory lists, because it is expensive to generate such pages. Because dynamic pages do not have corresponding physical files,

<sup>2</sup> “Virtual hosts” refers to the ability of a single Web server acting as multiple “virtual” Web servers, each of them has its own name and IP address.

the program assigns a unique “pseudo-file-name” to each page and stores the resulting page data with the pseudo-file-name in the file cache. The URI/pseudo-file-name pair is also cached in the URI cache to speed up accesses.

## 6 Results of the Enhancement Techniques




Figure 3: SPECweb96 Results of the Enhanced Apache

Figure 3 compares the performance of Apache with various performance-enhancement techniques we proposed and implemented. The bar on the far left is the performance of the original Apache. We can see that the DNS cache improves the server throughput by more than 14%. Caching the string conversion brings in another 4% performance increase. We then use the *mmap* function to speed up both small files and large files accesses. This results in a 9% throughput increase. However *mmap* does not eliminate other file system calls such as *open* and *close*, therefore the server still spends a lot of time on file system operations. Next we add the cache of the file states, which eliminates all *stat* calls and results in a 3% improvement. We then add the user-level file cache for small files, and use *mmap* only to speed up accesses for uncached large files. The user-level file cache removes almost all file system calls, improving the performance by about 11%. Finally, we implement the URI cache and obtain another 10% speed up. Together, the 6 techniques we proposed and implemented boost the performance of Apache by more than 61%, as shown by the right most bar in Figure 3.

## 7 Conclusions

In this paper, we present measurement results and performance analyses of the behavior of the Apache Web server on a uniprocessor system and a 4-CPU

No	Technique	Complexity	Effectiveness
1	<i>M</i> mapping large files	0	2
2	Caching small files	2	3
3	Caching file states	1	1
4	Caching DNS	0	3
5	Caching string results	0	1
6	Caching URI processing	2	3

Table 3: Web Server Enhancement Techniques. For implementation complexity, 0 means trivial and 3 means challenging. For effectiveness, 0 means almost useless, while 3 means very effective.

SMP system running the IBM AIX operating system. Using the built-in tracing facility and a trace-analysis tool, we obtained detailed information on OS kernel events and overall system activities while running Apache driven by the SPECweb96 and the WebStone benchmarks. We found that, on average, Apache spends about 20-25% of the total CPU time on user code, 35-50% on kernel system calls and 25-40% on interrupt handling. For systems with small RAM sizes, the Web server performance is limited by the disk bandwidth. For systems with reasonably large RAM sizes, the TCP/IP stack and the network interrupt handler are the major performance bottlenecks. We notice that Apache shows similar behavior on both the uniprocessor and the SMP systems.

After quantitatively identifying the performance bottlenecks, we proposed and implemented 6 techniques to improve the performance of Apache, which are summarized in Table 3. Our experimental results show that these techniques are quite effective. Together they improve the throughput of Apache by 61%. These techniques are general purpose and can be applied to other Web servers as well.

### Acknowledgments

Erich Nahum at IBM Watson read an early draft of this paper and gave many constructive comments and suggestions that greatly improved this paper. We benefited from discussions with Rich Neves at IBM Watson and Pedro Trancoso at University of Illinois. Bulent Abali at Watson helped us on setting up the test network.

### References

[1] The Apache Team, "Apache HTTP server project." <http://www.apache.org/>.

[2] M. F. Arlitt and C. L. Williamson, "Web server workload characterization: The search for invariants," in *Proceedings of the 1996 SIGMETRICS Conference on Measurement and Modeling of Computer Systems*, May 1996.

[3] V. Almeida, A. Bestavros, M. Crovella, and A. de Oliveira, "Characterizing reference locality in the WWW," in *Proceedings of the 1996 IEEE Conference on Parallel and Distributed Information Systems*, Dec. 1996.

[4] M. E. Crovella and A. Bestavros, "Self-similarity in World Wide Web traffic: Evidence and possible causes," in *Proceedings of the 1996 SIGMETRICS Conference on Measurement and Modeling of Computer Systems*, (Philadelphia, PA), May 1996.

[5] D. A. Helly, *AIX/6000 Internals and Architecture*. New York, NY 10020: McGraw-Hill, 1996.

[6] IBM RS/6000 Division, "Utlid 1.2 user's guide."

[7] Y. Hu, A. Nanda, and Q. Yang, "Measurement, analysis and performance improvement of the apache web server," Tech. Rep. 1097-0001, Department of Electrical and Computer Engineering, University of Rhode Island, Oct. 1997.

[8] J. M. Almeida, V. Almeida, and D. J. Yates, "Measuring the behavior of a world-wide web server," in *Seventh Conference on High Performance Networking (HPN)*, (White Plains, NY), pp. 57-72, Apr. 1997.

[9] L. McVoy and C. Staelin, "Imbench: Portable tools for performance analysis," in *Proceedings of the 1996 USENIX Conference*, 1996.

### For More Information

An extended version of this paper is available as a technical report via <http://www.ele.uri.edu/Research/hpcl/network/apache.html>.