Apple Computer, Inc.

Apple Network Server Benchmarks

- **Web Server Performance**
- **File Server Performance**

SUMMARY

Apple Network Server performance has been evaluated based on industry-adopted benchmark-testing parameters. These benchmarks include WebStone tests developed by Silicon Graphics, Inc., which evaluate web server performance, and the Mac NetBench test developed by ZD Labs, to evaluate file server performance.

The Apple Network Server demonstrated higher performance across all three server application categories than competitive UNIX solutions from Sun Microsystems, Inc., and IBM Corp., as well as NT solutions from Compaq Computer Corp. and Digital Equipment Corp. Test methodology and parameters, and detailed results are provided by server application category.

September 9, 1996

I. WEB SERVER PERFORMANCE

The Network Servers were tested according to WebStone 1.1 and WebStone 2.0.1 benchmarks.* Two of the three best results are attributed to the Network Servers. To put this into perspective, the Network Server 700/200 is capable of saturating up to 47 T1 lines at its maximum connection rate.

WebStone 1.1 Benchmark

Netscape 1.12 Software

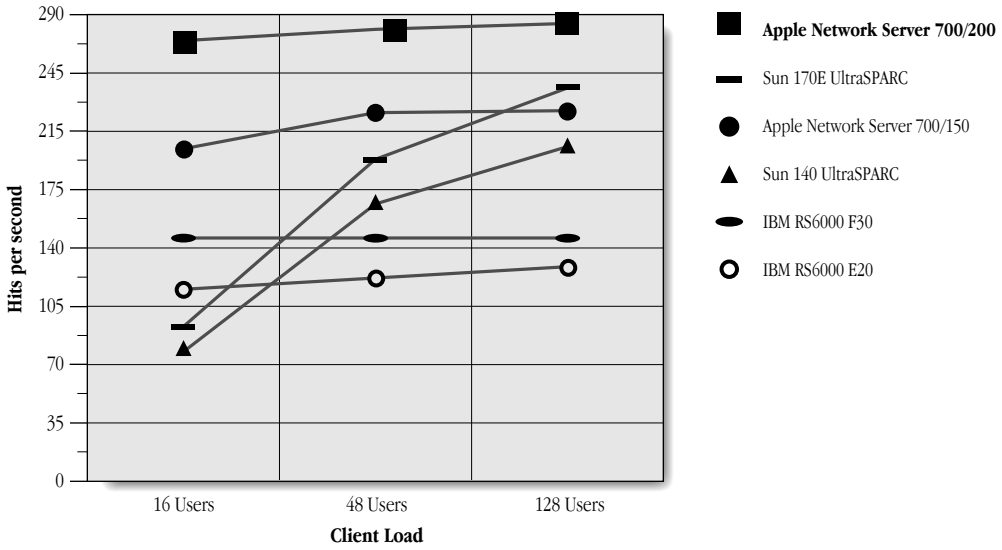

***Connection Rate**. Connections per second are also referred to as hits per second. Unless the size of the HTTP response is also known, it has little meaning. The average size of the HTTP response for the Webstone 1.1 test is approximately 7K. If the average response size had been larger, the connection rate would have been lower. Notice that even at 16 clients, the Apple Network Servers are delivering a very high maximum connect rate. This translates into excellent response times for individual clients. More connections per second is better.*

* *Tests were conducted according to both WebStone test versions to provide a basis for competitive evaluation. Not all vendors have yet completed tests using the WebStone 2.0.1 benchmarks.*

WebStone 1.1 Tested Server Configurations

	Apple	*IBM*	*IBM*	*Sun*
Tested Server	Network Server 700	RS6000/E20	RS6000/F30	Sun 140, 170
Processor	PPC 604: 150, 200	PPC 604/100	PPC 604/133	UltraSPARC 143,167
Main Memory	64MB	48MB	64MB	64MB
OS	AIX 4.1	AIX 4.1	AIX 4.1	Solaris 2.5
Internet SW	Netscape Comm Server 1.12	Netscape Comm Server 1.12	Netscape Comm Server 1.12	Netscape Comm Server 1.12
Test Suite	WebStone 1.1	WebStone 1.1	WebStone 1.1	WebStone 1.1
Who Ran Test?	Apple	Apple	IBM	Apple

WebStone 2.0.1 Benchmark

Zeus Software

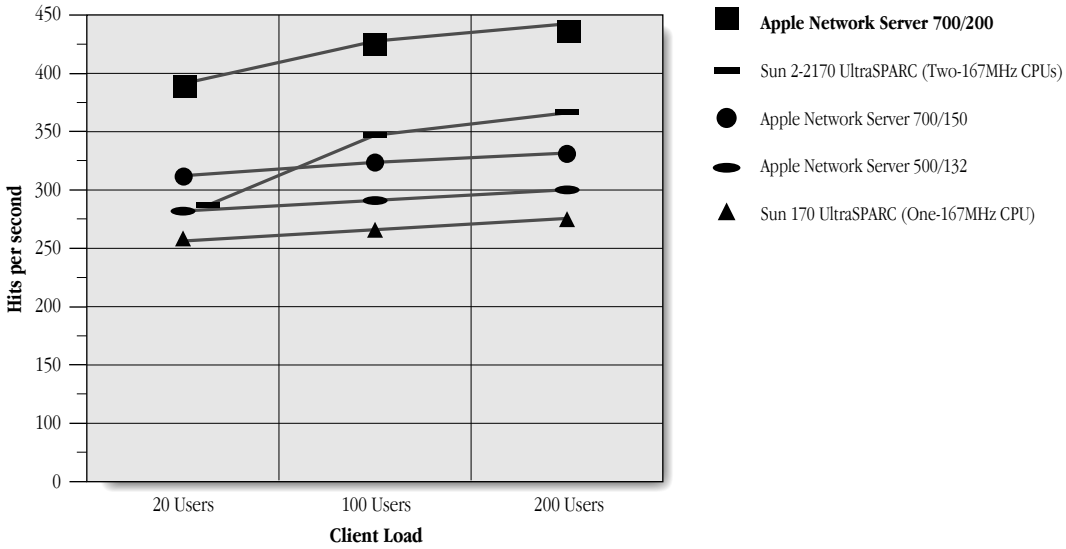

Connection Rate. *Connections per second are also referred to as hits per second. Unless the size of the HTTP response is also known, it has little meaning. The average size of the HTTP response for the WebStone 2.0.1 test is approximately 19K. If the average response size had been larger, the connection rate would have been lower. Notice that at only 20 clients, the difference in the Apple Network Servers Ethernet rate from the highest number of clients is within 10 percent. This translates into excellent response times for individual clients. More connections per second is better.*

WebStone 2.0.1 Tested Server Configurations

	Apple	*Apple*	*Sun*	*Sun*
Tested Server	Network Server 500, Network Server 700	Network Server 700/200	Sun 140	Sun 2-2170
Processor	PPC 604: 132, 150	PPC 604e/200	UltraSPARC 143	Two UltraSPARC 167s
Main Memory	64MB	64MB	128MB	128MB
OS	AIX 4.1	AIX 4.1	Solaris 2.5	Solaris 2.5
Internet SW	Zeus	Zeus	Zeus	Zeus
Test Suite	WebStone 2.0.1	WebStone 2.0.1	WebStone 2.0.1	WebStone 2.0.1
Who Ran Test?	Apple	Apple	Sun	Sun

II. FILE SERVER PERFORMANCE

File server performance was benchmarked according to Mac NetBench 4.0.1. The Network Servers delivered two of the best three performance results, and demonstrated more than four times the performance of NT servers from Digital and Compaq when serving files to Macintosh desktops.

Mac NetBench 4.0.1

File Server Benchmark (Sequential Reads)

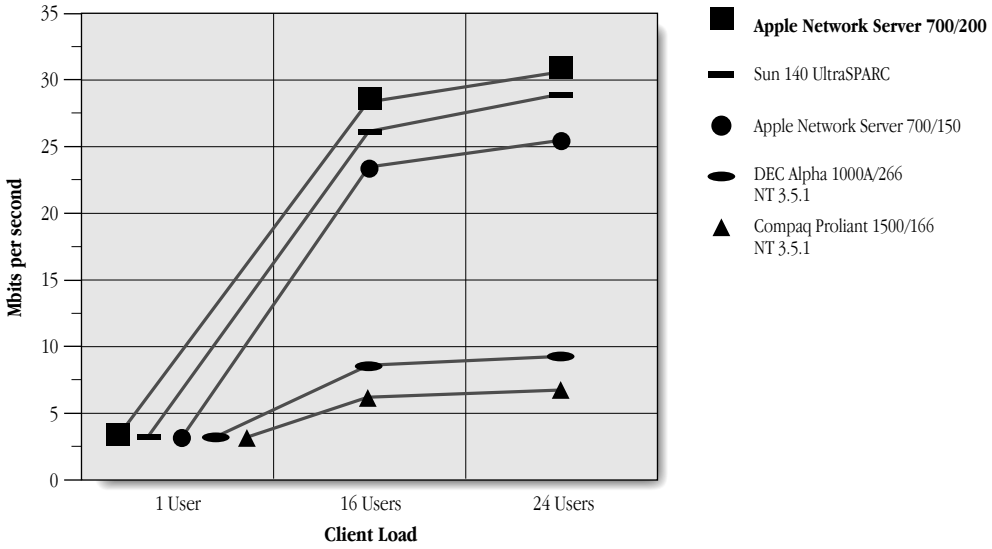

Sequential Reads Test. *The sequential reads benchmark measures the server's ability to read large sequential files in/out of the server. This activity is typical in a professional publishing environment.*

NetBench 4.0.1 Tested Server Configurations

	Apple	*Sun*	*Digital*	*Compaq*
Tested Server	Network Server 700	Sun 140	AlphaServer 1000A	ProLiant 1500
Processor	PPC 604: 150, 200	UltraSPARC 143	21064 266MHz	Pentium 166MHz
Main Memory	64MB	64MB	64MB	64MB
Disk Drives	2x2GB	2x2GB	2x2GB	2x2GB
Fast Ethernet	Yes	Yes	Yes	Yes
OS	AIX 4.1	Solaris 2.5	NT 3.5.1	NT 3.5.1
File Server SW	IPT uShare	IPT uShare	NT AFP	NT AFP
Test Suite	NetBench 4.0.1	NetBench 4.0.1	NetBench 4.0.1	NetBench 4.0.1
Who Ran Test?	Apple	Apple	Apple	Apple

APPENDIX: TEST METHODOLOGY & PARAMETERS

A brief description of each performance benchmarking method and test configuration is provided. Detailed benchmark testing documents can be obtained from their respective authors.

WEBSTONE 1.1 & 2.0.1 BENCHMARK

WEBSTONE 1.1

Overview

The WebStone 1.1 benchmark was conducted using the Apple Network Server 500/132, 700/150 running AIX 4.1.4, and 700/200 running AIX 4.1.4.1 . The report may be used to compare the relative performance of these platforms against other platforms reporting WebStone 2.0.1 numbers.

Note: Web server benchmarking is extremely sensitive to software utilization. Current web server benchmarks are fast-moving targets and care should be taken to insure that comparisons between reported benchmarks are meaningful. Competitive test results for the WebStone 1.1 benchmark must be based on Netscape Corp. Communications Server 1.12 or Netscape Commerce Server 1.12 to be valid.

Testbed Configuration

The tests were conducted over a 100Mbit private network (100baseTX Fast Ethernet). All systems (client and server) were running AIX 4.1.4 with APAR IX56968 (general kernel TCP/IP performance enhancements) installed. (Note: APAR IX56968 is included in AIX 4.1.4.1). Two client machines were used to generate the load.

The Clients

Client 1	*Client 2*
Apple Network Server 700	Apple Network Server 700
150MHz 604	200MHz 604e
1MB L2 Cache	1MB L2 Cache
64MB memory	64MB memory
PCI - 100Mbit Ethernet board	PCI - 100Mbit Ethernet board
AIX 4.1.4 + APAR IX56968	AIX 4.1.4.1

The Servers

Server 1	*Server 2*	*Server 3*
Apple Network Server 500	Apple Network Server 700	Apple Network Server 700
132MHz 604	150MHz 604	200MHz 604e
512K L2 Cache	1MB L2 Cache	1MB L2 Cache
64MB memory	64MB memory	64MB memory
PCI - 100Mbit Ethernet board	PCI - 100Mbit Ethernet board	PCI - 100Mbit Ethernet board
AIX 4.1.4 + APAR IX56968	AIX 4.1.4 + APAR IX56968	AIX 4.1.4.1

Testbed Parameter Tuning

Web Server Parameter Tuning		*System Parameter Tuning*		*Network Interface Tuning*	
Access logging	off	Max. processes per user	1000	Transmit queue size	256
DNS lookups	off			Receive buffer pool size	128
MinProcs	16			Transmit FIFO threshold	128
MaxProcs	64				

Workload Filelist

The original filelist that accompanies the WebStone 1.1 benchmark was used for performance testing. This filelist was created by Silicon Graphics, Inc. (SGI), using data from one of its own web sites.

#Silicon Surf model: pages and files to be tested for

```
8
40 2
/file2k.html
/file3k.html
25 2
/file1k.html
/file5k.html
15 2
/file4k.html
/file6k.html
5 1
/file7k.html
4 4
/file8k.html
/file9k.html
/file10k.html
/file11k.html
4 5
/file12k.html
/file14k.html
/file15k.html
/file17k.html
/file18k.html
6 1
/file33k.html
1 1
/file200k.html
```

Testing Methodology

The standard workload filelist from the SGI web site was used to generate the requests to the server. The average response size of each HTTP request is 7K using this load. Two client stations were used to insure that the server was completely saturated with requests; this was verified through the use of the Vmstat system performance utility.

Testing began at 16 virtual clients, incrementing by 16 up to a maximum of 128 clients. The duration of each test run was 5 minutes. Test cycles as long as 60 minutes were run for selected cases to insure that the 5-minute runs were yielding representative results. There was no significant difference in the reported metrics between runs of 5 minutes and 60 minutes at the same client load.

The tests were conducted over Fast Ethernet (use of 10Mbit Ethernet does not allow server testing without network bottlenecks). There was greater than 95 percent utilization of 10Mbit Ethernet at 140 connections per second with more than 40 percent of the processor still available as reported by the Vmstat system performance utility.

WebStone 1.1 Test Results

Test Results for the Apple Network Server 700/200

Clients	*Conn/second*	*Errors/second*	*Latency*	*Little's Law*	*Throughput (Mbits/sec)*
16	285.12	0.0000	0.0557	15.88	15.24
32	284.52	0.0000	0.1120	31.87	15.22
48	285.79	0.0000	0.1671	47.75	15.35
64	284.24	0.0000	0.2199	62.50	15.23
80	286.96	0.0167	0.2758	79.14	15.16
96	284.84	0.0000	0.3360	95.72	15.21
112	286.51	0.0000	0.3891	111.49	15.21
128	283.37	0.0000	0.4509	127.77	15.21

Test Results for the Apple Network Server 700/150

Clients	*Conn/second*	*Errors/second*	*Latency*	*Little's Law*	*Throughput (Mbits/sec)*
16	215.60	0.0000	0.0742	15.82	11.30
32	219.05	0.0000	0.1456	31.89	11.72
48	216.58	0.0000	0.2206	47.77	11.73
64	217.93	0.0000	0.2915	63.52	11.78
80	219.10	0.0000	0.3646	79.87	11.65
96	217.63	0.0000	0.4342	94.50	11.81
112	213.33	0.0000	0.5195	110.81	11.56
128	217.33	0.0000	0.5799	126.02	11.58

Test Results for the Apple Network Server 500/132

Clients	*Conn/second*	*Errors/second*	*Latency*	*Little's Law*	*Throughput (Mbits/sec)*
16	189.76	0.0000	0.0839	15.91	9.99
32	189.39	0.0000	0.1676	31.75	10.16
48	190.72	0.0000	0.2511	47.89	10.40
64	191.25	0.0000	0.3342	63.90	10.46
80	194.37	0.0000	0.4109	79.86	10.39
96	193.66	0.0000	0.4915	95.19	10.46
112	192.89	0.0000	0.5791	111.71	10.27
128	191.08	0.0000	0.6685	127.74	10.28

Performance Metrics

Connection Rate. Connections per second are also referred to as hits per second. Unless the size of the HTTP response is also known, it has little meaning. The average size of the HTTP response for the tests conducted is approximately 7K. If the average response size had been larger, the connection rate would have been lower. Notice that even at 16 clients, the Apple Network Servers are delivering a very high maximum connect rate. This translates into excellent response times for individual clients. More connections per second is better.

Error Rate. As the load to the server increases (represented by the number of simultaneous connection requests), the server may begin to generate "connection refused" messages. This occurs when a server is not fast enough to process the volume of incoming requests in real time and must queue them for future attention. When this queue becomes full, the server starts to turn away new requestors. For the tests conducted, both Apple Network Servers completed with no refusals. This indicates an ability to deal with peak load conditions (that is, a large number of clients make an HTTP request at almost the same instant). Fewer errors per second is better.

Latency. Latency represents the average response time that a user perceives. As the number of simultaneous requests to the server increases, the response time for a particular user will degrade. This happens because once the server is saturated, it must begin to queue additional requests until it is finished processing the current requests. This metric is only valid when the error rate is at or near zero. Lower latency times are better.

Little's Load Factor. Little's Law is based on queuing theory. Little's Load Factor is a ratio of time spent talking to the web server versus time waiting for the server to respond. In the WebStone 1.1 tests, this metric should be just under or at the number of client processes at each of the sample points. Thus the best a server could achieve would be a ratio of 1 (for example, at 32 clients Little's Load Factor would equal 32).

Throughput. WebStone 1.1 includes only the HTTP header and actual response data as part of this metric. The actual number of bits being transferred on the wire is about 10 percent higher for these particular tests. Thus each reported 1.4Mbits of throughput translates roughly into one T1 line. The Apple Network Server 700 would be able to saturate eight T1 lines at its maximum connection rate given an average response size of 7K per HTTP request. This metric is directly proportional to the connection rate when the response size is held constant. Higher throughput is better.

WEBSTONE 2.0.1

Overview

The WebStone 2.0.1 benchmark was conducted using the Apple Network Server 500/132, 700/150, and 700/200 running AIX 4.1.4.1 and the Zeus Server. The report may be used to compare the relative performance of these platforms against other platforms reporting WebStone 2.0.1 numbers.

Note: Web server benchmarking is extremely sensitive to software utilization. Current web server benchmarks are fast-moving targets and care should be taken to insure that comparisons between reported benchmarks are meaningful. Competitive test results must be based against the Zeus Server with the Keep-Alive feature of HTTP 1.1 not enabled. As distributed from Silicon Graphics, Inc., the WebStone 2.0.1 benchmark suite does not utilize the Keep-Alive feature of HTTP 1.1. These benchmarks were conducted without this feature enabled. When comparing results to those published by other vendors it is important to determine whether the benchmark suite was modified to exploit the Keep-Alive feature, since the use of this feature can result in significantly better results.

Testbed Configuration

The tests were conducted over a 100Mbit private network (100base TX Fast Ethernet). All systems (client and server) were running AIX 4.1.4 with APAR IX56968 (general kernel TCP/IP performance enhancements) installed. Two client machines were used to generate the load.

The Clients

Client 1	*Client 2*
Apple Network Server 700	Apple Network Server 700
150MHz 604	200MHz 604e
1MB L2 Cache	1MB L2 Cache
64MB memory	64MB memory
PCI - 100Mbit Ethernet board	PCI - 100Mbit Ethernet board
AIX 4.1.4 + APAR IX56968	AIX 4.1.4.1

The Servers

Server 1	*Server 2*	*Server 3*
Apple Network Server 500	Apple Network Server 700	Apple Network Server 700
132MHz 604	150MHz 604	200MHz 604e
512K L2 Cache	1MB L2 Cache	1MB L2 Cache
64MB memory	64MB memory	64MB memory
PCI - 100Mbit Ethernet board	PCI - 100Mbit Ethernet board	PCI - 100Mbit Ethernet board
AIX 4.1.4.1	AIX 4.1.4.1	AIX 4.1.4.1

Testbed Parameter Tuning

Web Server Parameter Tuning		*System Parameter Tuning*		*Network Interface Tuning*	
Access logging	off	Max. processes per user	1000	Transmit queue size	256
DNS lookups	off			Receive buffer pool size	128
Number of processes	1			Transmit FIFO threshold	512
Idle timeout	5 seconds				

Workload Filelist

The original filelist that accompanies the WebStone 2.0.1 benchmark was used for the performance testing. This filelist was created by SGI, using data from one of its own web sites.

#The standard filelist distributed with WebStone 2.0.1

```
/file500.html       350    #500
/file5k.html        500    #5125
/file50k.html       140    #51250
/file500k.html      9      #512500
/file5m.html        1      #5248000
```

Testing Methodology

The standard workload filelist from the SGI web site was used to generate the requests to the server (see above). The average response size of each HTTP request ends up being about 19K using this load. Two client stations were used to insure that the server was completely saturated with requests. This was verified through the use of the Vmstat system performance utility.

Testing began at 10 virtual clients, incrementing by 10 up to a maximum of 200 clients. The duration of each test run was 10 minutes. Test cycles as long as 60 minutes were run for selected cases to insure that the 10-minute runs were yielding representative results. There was no significant difference in the reported metrics between runs of 10 minutes and 60 minutes at the same client load.

The tests were conducted over Fast Ethernet (use of 10Mbit Ethernet does not allow server testing without network bottlenecks). There was greater than 95 percent utilization of a 10Mbit Ethernet at 140 connections per second with more than 40 percent of the processor still available as reported by the Vmstat system performance utility.

WebStone 2.0.1 Test Results

Test Results for the Apple Network Server 700/200

Clients	*Conn/second*	*Errors/second*	*Latency*	*Little's Law*	*Throughput (Mbits/sec)*
10	373.63	0.0000	0.0264	9.85	57.02
20	394.57	0.0000	0.0501	19.75	61.62
30	417.30	0.0000	0.0712	29.70	62.28
40	419.07	0.0000	0.0939	39.33	65.28
50	430.71	0.0067	0.1140	49.10	64.86
60	439.66	0.0000	0.1353	59.48	65.14
70	429.54	0.0000	0.1618	69.49	65.71
80	447.06	0.0000	0.1772	79.24	64.58
90	431.38	0.0000	0.2054	88.58	66.58
100	438.90	0.0000	0.2257	99.08	66.68
110	437.11	0.0000	0.2486	108.64	65.07
120	443.57	0.0000	0.2675	118.66	66.12
130	441.58	0.0000	0.2925	129.14	66.14
140	434.46	0.0617	0.3094	134.42	65.86
150	441.97	0.0167	0.3331	147.22	64.93
160	444.82	0.0000	0.3527	156.90	65.98
170	447.18	0.0000	0.3779	169.00	65.93
180	449.16	0.0000	0.3979	178.71	65.32
190	448.00	0.0000	0.4213	188.75	65.87
200	444.85	0.0167	0.4427	196.93	65.61

Test Results for the Apple Network Server 700/150

Clients	*Conn/second*	*Errors/second*	*Latency*	*Little's Law*	*Throughput (Mbits/sec)*
10	290.52	0.0000	0.0341	9.91	43.85
20	302.07	0.0000	0.0659	19.90	47.51
30	314.20	0.0000	0.0943	29.64	48.23
40	311.30	0.0017	0.1269	39.49	47.47
50	313.90	0.0000	0.1584	49.73	48.61
60	325.47	0.0017	0.1827	59.48	49.17
70	329.52	0.0000	0.2118	69.79	49.72
80	329.28	0.0000	0.2406	79.23	49.78
90	331.65	0.0000	0.2692	89.28	49.04
100	337.54	0.0000	0.2944	99.37	48.93
110	328.84	0.0000	0.3314	108.99	49.33
120	333.24	0.0150	0.3466	115.51	49.92
130	330.98	0.0000	0.3897	128.98	50.45
140	330.62	0.0600	0.4067	134.48	50.08
150	332.72	0.0000	0.4466	148.61	50.09
160	326.71	0.0317	0.4803	156.93	50.51
170	327.10	0.0500	0.5060	165.52	50.57
180	330.05	0.0300	0.5249	173.25	50.03
190	324.40	0.0167	0.5757	186.76	50.71
200	335.86	0.0000	0.5897	198.05	50.05

Test Results for the Apple Network Server 500/132

Clients	*Conn/second*	*Errors/second*	*Latency*	*Little's Law*	*Throughput (Mbits/sec)*
10	269.45	0.0000	0.0368	9.92	40.48
20	279.98	0.0000	0.0707	19.79	41.82
30	283.90	0.0000	0.1047	29.71	42.94
40	279.57	0.0000	0.1419	39.68	43.38
50	291.78	0.0050	0.1654	48.27	48.27
60	287.74	0.0000	0.2062	59.34	44.08
70	298.80	0.0000	0.2329	69.58	43.80
80	290.54	0.0483	0.2616	76.01	43.80
90	292.06	0.0000	0.3057	89.30	43.47
100	290.38	0.0400	0.3322	96.48	43.74
110	296.96	0.0167	0.3649	108.37	43.40
120	296.65	0.0450	0.3863	114.60	43.89
130	293.38	0.0100	0.4320	126.73	43.82
140	291.23	0.0883	0.4564	132.91	43.80
150	293.29	0.0267	0.5031	147.55	43.79
160	293.90	0.0517	0.5261	154.61	44.39
170	289.22	0.0700	0.5612	162.30	44.66
180	294.10	0.0517	0.5908	173.75	44.08
190	290.48	0.0000	0.6400	185.92	44.75
200	298.31	0.2083	0.6030	179.88	43.84

Performance Metrics

Connection Rate. Connections per second are also referred to as hits per second. Unless the size of the HTTP response is also known, it has little meaning. The average size of the HTTP response for the tests conducted is approximately 19K. If the average response size had been larger, the connection rate would have been lower. Notice that at only 30 clients, the difference in the Apple Network Servers Ethernet rate from the highest number of clients is within 10 percent. This translates into excellent response times for individual clients. More connections per second is better.

Error Rate. As the load to a server increases (represented by the number of simultaneous connection requests), the server may begin to generate "connection refused" messages. This occurs when a server is not fast enough to process the volume of incoming requests in real time and must queue them for future attention. When this queue becomes full, the server starts to turn away new requestors. For the tests conducted, both Apple Network Servers completed with no refusals. This indicates an ability to deal with peak load conditions (that is, a large number of clients make an HTTP request at almost the same instant). Fewer errors per second is better.

Latency. Latency represents the average response time that a user perceives. As the number of simultaneous requests to the server increases, the response time for a particular user will also degrade. This happens because once the server is saturated, it must begin to queue additional requests until it is finished processing the current requests. This metric is only valid when the error rate is at or near zero. Lower latency times are better.

Little's Load Factor. Little's Law is based on queuing theory. Little's Load Factor is a ratio of time spent talking to the web server versus time waiting for the server to respond. In the WebStone 2.0.1 tests, this metric should be just under or at the number of client processes at each of the sample points. The best a server could achieve would be a ratio of 1 (for example, at 32 clients Little's Load Factor would equal 32).

Throughput. WebStone 2.0.1 includes only the HTTP header and actual response data as part of this metric. The actual number of bits being transferred on the wire is about 10 percent higher for these particular tests. Thus each reported 1.4Mbits of throughput translates roughly into one T1 line. The Apple Network Server 700/200 would be able to saturate 47 T1 lines, or 1.5 T3 at its maximum connection rate given an average response size of 19K per HTTP request. This metric is directly proportional to the connection rate when the response size is held constant. Higher throughput is better.

MAC NETBENCH 4.0.1 BENCHMARK

MAC NETBENCH 4.0.1

Overview

ZD Labs designed a file server performance test suite to compare file server throughput for various LAN server products. PC NetBench 4.0.1 measures PC file server performance to Windows clients, while Mac NetBench 4.0.1 measures AFP file server throughput between the server and Macintosh desktops.

Sequential Reads Test

The sequential reads benchmark measures the server's ability to read large sequential files in/out of the server. This activity is typical in a professional publishing environment.

Apple Computer, Inc.
1 Infinite Loop
Cupertino, CA 95014

© 1996 Apple Computer, Inc. All rights reserved.
Apple, the Apple logo, and AppleTalk are trademarks of Apple Computer, Inc., registered in the U.S.A. and other countries. AIX is a registered trademark and PowerPC is a trademark of International Business Machines Corporation, used under license therefrom. Windows is a trademark of Microsoft Corporation. Mention of third-party products is for informational purposes only and constitutes neither an endorsement nor a recommendation. Apple assumes no responsibility with regard to the selection, performance, or use of these products. All understandings, agreements or warranties, if any, take place directly between the vendors and the prospective users.