

NAVAL POSTGRADUATE SCHOOL MONTEREY, CALIFORNIA



THESIS

INFRASTRUCTURE CONSIDERATIONS FOR WORLD WIDE WEB SERVERS

by

John T. Casler

June, 1996

Co-Advisors:

Hemant Bhargava
Suresh Sridhar

Approved for public release; distribution is unlimited

19960805 015

DTIC QUALITY INSPECTED 1

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave Blank)		2. REPORT DATE June, 1996	3. REPORT TYPE AND DATES COVERED Master's Thesis	
4. TITLE AND SUBTITLE INFRASTRUCTURE CONSIDERATIONS FOR WORLD WIDE WEB SERVERS			5. FUNDING NUMBERS	
6. AUTHOR(S) Casler, John T.				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey, CA 93943-5000			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)			10. SPONSORING / MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government.				
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release, distribution is unlimited			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) This thesis explores issues associated with defining and selecting infrastructure requirements for World Wide Web sites. The explosive growth the WWW has made it the largest single service on the Internet. With this growth comes a need for guidance to organizations or individuals desiring to establish new Web sites. This thesis provides the guidance needed to define a potential site's requirements and select the infrastructure necessary to fulfill those requirements. A combination of literature review of current books and periodicals, as well as surveys of WWW sites was used to obtain information. This information was used to develop the framework for defining requirements. A rule based heuristic was also adopted from the literature and subsequently validated. It is used to select the computing hardware needed for a site. A key lesson learned is that most organizations do not conduct any initial requirements analysis to determine a site's infrastructure needs. The reasons range from oversight to indifference. The potential penalty for not conducting proper assessment of requirements is the same as for any venture, a substandard product and poorly leveraged investment.				
14. SUBJECT TERMS World Wide Web Site Requirements and Definition Server Hardware Hueristic			15. NUMBER OF PAGES 143	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UL	

Approved for public release; distribution is unlimited.

INFRASTRUCTURE CONSIDERATIONS FOR WORLD WIDE WEB SERVERS

John T. Casler
Lieutenant Commander, United States Naval Reserve
B.S., University of South Carolina, 1979


Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN INFORMATION TECHNOLOGY MANAGEMENT

from the


NAVAL POSTGRADUATE SCHOOL
June, 1996


Author: _____


John T. Casler

Approved by: _____


Hemant Bhargava, Co-Advisor


Suresh Sridhar, Co-Advisor


Reuben Harris, Chairman
Department of Systems Management

ABSTRACT

This thesis explores issues associated with defining and selecting infrastructure requirements for World Wide Web sites. The explosive growth the WWW has made it the largest single service on the Internet. With this growth comes a need for guidance to organizations or individuals desiring to establish new Web sites. This thesis provides the guidance needed to define a potential site's requirements and select the infrastructure necessary to fulfill those requirements.

A combination of literature review of current books and periodicals, as well as surveys of WWW sites was used to obtain information. This information was used to develop the framework for defining requirements. A rule based heuristic was also adopted from the literature and subsequently validated. It is used to select the computing hardware needed for a site.

A key lesson learned is that most organizations do not conduct any initial requirements analysis to determine a site's infrastructure needs. The reasons range from oversight to indifference. The potential penalty for not conducting proper assessment of requirements is the same as for any venture, a substandard product and poorly leveraged investment.

TABLE OF CONTENTS

I. INTRODUCTION	1
A. ENVIRONMENT	1
B. OBJECTIVES	2
C. SCOPE	2
D. METHODOLOGY	3
E. ORGANIZATION	4
II. WWW TECHNICAL ASPECTS	5
A. TCP/IP	5
B. CLIENT SERVER MODEL	6
C. WORLD WIDE WEB	7
1. HyperText Transfer Protocol (HTTP)	8
a. Stateless	8
b. Connectionless	9
2. HyperText Markup Language (HTML)	10
3. Uniform Resource Locators (URLs)	10
4. Common Gateway Interface (CGI)	11
a. Scripts	12
D. CONCLUSION	13
III. PROBLEM DESCRIPTION	15
A. SITE SURVEY RESULTS	15
B. LITERATURE REVIEW RESULTS	18

C. COST	20
D. CONCLUSION	21
IV. INFRASTRUCTURE REQUIREMENTS.....	23
A. ANTICIPATING TRAFFIC	24
B. PURPOSE OF THE SITE.....	26
C. CONNECTION.....	28
1. Line Options	28
2. Bandwidth Requirement	29
3. Additional Requirements	33
D. HARDWARE.....	34
1. Processing power	34
2. RAM.....	35
3. Disk Space	36
E. SOFTWARE.....	36
1. Operating Systems.....	37
2. Server Software.....	37
F. CONCLUSION.....	39
V. HARDWARE SELECTION.....	41
A. HARDWARE HEURISTIC.....	41
B. HEURISTIC BENCHMARK	43
1. SPEC	43
a. "Base" Measurement	44
(1) Speed.....	44
(2) Throughput	44
b. "Peak" Measurement	45
(1) Speed.....	45

(2) Throughput	45
2. Heuristic.....	46
C. HEURISTIC VALIDATION	50
D. CONCLUSION	52
VI. SYSTEM RELIABILITY	53
A. DISK STORAGE SUBSYSTEM	54
1. RAID 0	55
2. RAID 1	55
3. RAID 2	55
4. RAID 3	56
5. RAID 4	56
6. RAID 5	56
B. SYSTEM REDUNDANCY	57
1. Clustering.....	57
a. Hyperlinked Computers.....	58
b. Commercial Solutions.....	59
c. NCSA scalable Server Approach.....	59
2. Super Servers	62
C. CONCLUSION	63
VII. CONCLUSION	65
A. THESIS SUMMERY	65
B. CONTRIBUTIONS	65
C. AREAS OF ADDITIONAL RESEARCH.....	66
1. Operating Systems.....	66
2. AFS	67
D. CONCLUSIONS AND RECOMMENDATIONS	67
APPENDIX A. WEB SERVERS SURVEY	69

APPENDIX B. SPEC REFERENCE TABLES	77
APPENDIX C. HEURISTIC BENCHMARK	103
APPENDIX D. WEB SITE SURVEY	107
LIST OF REFERENCES.....	127
INITIAL DISTRIBUTION LIST.....	131

I. INTRODUCTION

A. ENVIRONMENT

The growth of the Internet in general and the World Wide Web (WWW or Web) in particular is well documented and publicized. It has been estimated that between July 1993 and July 1995 the number of Internet hosts worldwide increased from 1.78 million to 6.64 million (Network Wizards, 1995). During this same general period the WWW accounted for 130 sites in June 1993 and 23,500 sites in June of 1995. The June 1995 estimates equate to one Web server per 270 hosts (Gray, 1995). Just six months later, in January 1996, Web sites numbered 90,000, which equated to an estimated one Web server per 100 hosts (Gray, 1995)!

Due to this growth the WWW is now the largest single service on the Internet. Because the WWW is to the Internet as Windows is to the personal computer it seems natural that the majority of the growth in the Internet has been and will continue to be in this area. With this growth comes a need for guidance to organizations or individuals desiring to establish new Web sites.

To the 'uninitiated' the computer industry is shrouded in jargon and meticulous technical issues. The success of the WWW as an information and entertainment source is thrusting it into the lives and businesses of those uninitiated. To establish a new Web site these individuals/organizations often turn to computer industry professionals for assistance. Depending on the credentials and motivations of these professionals the resulting Web site may, or may not, accurately reflect what the client needs.

To the computer and Internet 'literate' the subject is more tractable. However, because the requirements for specific sites can vary greatly depending on the function of that site, it is not uncommon for important nuances to be overlooked or the complexity of the task to be underestimated. The result is often an investment in Web site infrastructure that is insufficient or inappropriate to meet the demands of the site.

Guidance for establishing a Web site is needed to steer both the 'uninitiated' and the 'literate' in determining their specific requirements and then assist them in choosing the infrastructure to fulfill those requirements.

B. OBJECTIVES

The goal of this thesis is to provide guidance for individuals and organizations so they may define and fulfill their WWW requirements. It will address the basic infrastructure issues that must be answered and will provide a rule based heuristic to facilitate the selection of the Web site infrastructure based on the identified requirements.

In this thesis 'infrastructure' is used to represent all the necessary components of a Web site. This will include the size of the connection or 'pipe' (56 Kbps for instance) required to carry the electronic information to and from the Web site. It also includes all the hardware associated with the site such as the Central Processing Unit (CPU) capacity, Random Access Memory (RAM) and hard disk storage capacity, as well as the operating system used and the server software chosen to perform the Web server functions.

In those situations where it is necessary to refer to all site requirements *except the size of the connection*, the term 'Platform' will be used. Similarly, the term 'Hardware' will be used to reference the CPU, RAM and hard disk (everything except the size of the pipe and the operating system and server software).

C. SCOPE

The basic premise of this thesis is that the decision has been made by an organization to acquire the infrastructure necessary to establish a WWW presence. The thesis does not deal with an organizations strategic decision to purchase and employ information technology. Although often impulsively made, this decision, as Clemons (1991) points out, is neither easy or obvious and can be far more difficult than defining and selecting hardware requirements. For a useful discussion on this elemental topic see

Clemon's article "Evaluation of STRATEGIC INVESTMENTS in Information Technology".¹

Similarly, this thesis will not deal with defining which Internet services are most appropriate for a particular organization to offer. It is assumed that reasonable consideration on this issue has been made. An excellent reference to facilitate this decision is *Managing INTERNET Information Services*.²

Security is another area that will not be addressed. The subject is complicated and it is a topic that demands dedicated study, not a cursory mention.

Finally, HTML authoring will not be discussed. Web site content is, however, the most fundamental issue associated with creating a WWW presence. The success or failure of the site can depend on this issue. One of the many excellent published or on-line references available in the creation of HTML documents and Web authorship should be consulted.

D. METHODOLOGY

The initial research approach for this thesis was a combination of literature review of current books and periodicals, and site surveys of existing WWW sites to obtain information for both requirements guidance and the rule based heuristic. This information was to be used to help develop the framework for defining requirements as well as for developing a rule based heuristic that would be used to pick hardware. This approach worked well for establishing guidance for requirements. However, it was necessary to modify this approach for the rule based heuristic. Due to a lack of information it became necessary to adapt, instead of develop, an existing heuristic. This heuristic was then

¹E. Clemons, "Evaluation of STRATEGIC INVESTMENTS in Information Technology". *Communications of the ACM*, 34(1):22-3, 1991.

²C. Liu, J. Peek, R. Jones, B. Buus, and A. Nye. *Managing INTERNET Information Services*, O'Reilly and Associates, Inc., Sebastopol, California, 1994.

validated against existing Web sites. The method for this approach will be amplified in Chapter V.

E. ORGANIZATION

This thesis contains seven chapters and five appendices. Chapter I contains the introduction and overview of the thesis. Chapter II covers the pertinent technical characteristics of the Internet and the World Wide Web. Chapter III discusses research results. Chapter IV details infrastructure requirements. Chapter V presents the heuristic used for selecting hardware. Chapter VI examines redundancy and reliability issues. Chapter VII covers the conclusion, recommendations and suggested areas for further research.

Appendix A is a usage survey of World Wide Web servers. Appendix B supplies SPEC Reference Tables to be used when comparing computer hardware to the heuristic levels. Appendix C lists benchmark values assigned to each level of the heuristic. Appendix D contains surveys which were conducted to determine the validity of the hardware heuristic.

II. WWW TECHNICAL ASPECTS

"The Web is a collection of protocols and standards for accessing information on the Internet, and the Internet is the physical medium used to transport the data." (Net.Genesis and Hall, 1995)

As this thesis is intended to assist the 'uninitiated' as well as those with a broader understanding of the subject, a brief introduction to some of the pertinent technical aspects of the Internet and the WWW will be useful.

A. TCP/IP

Since the Internet is "...the physical medium used to transport the data.", a basic understanding of the two primary protocols used by the Internet is necessary. As mentioned, the Internet is a 'network of networks' all communicating with the common software protocols of TCP/IP (Transmission Control Protocol and Internet Protocol).

Basically, IP has one function, it acts like an envelope to 'carry' (encapsulate) a message to an 'address' (computer on the Internet). IP does **not** guarantee delivery. However, it is fast and easy to implement. (Liu, et al., 1994)

TCP provides three functions that IP does not: serialization of data, guaranteed delivery, and a port number. 'Serialization' (consecutively number) of the data packets ensures they are reassembled in the correct order when received. In this way the delivery of the data can also be guaranteed (if a sequence number is missing that data packet can be retransmitted). Port numbers identify individual services (such as Gopher) or applications within a destination computer that are being requested. (Liu, et al., 1994)

One of the issues with regard to WWW servers is how 'efficient' the server operating system is at handling TCP/IP. This will be discussed in Chapter IV.

B. CLIENT SERVER MODEL

The WWW are based on the 'Client-Server' model (see Figure 1). This concept refers to the relationship between two (or more) computers. One computer - the 'Client' - establishes a connection (via the Internet and a hypertext link) and requests information or services from another computer - the 'Server' - which processes the request and then returns (via the Internet) the information or services. The client server model, in principle, is the same as going to a store and asking to be served. You, the client, request an item and the store clerk, the server, gives that item to you. (Net.Genesis and Hall, 1995)




Figure 1: Client-Server Model

For this model to work the client and server must be using the same communications protocols. If you go to the store and request service in a foreign language you are not likely to obtain the desired response from the clerk. TCP/IP are the underlying protocols required for client/server functionality on the Internet. Additional protocols such as FTP or Gopher are required depending on the particular services you wish to provide or access. A distinct advantage of the WWW is that it was designed to support previous Internet protocols. Therefore, a Web client running Web browser software such as Mosaic

has 'backward' computability with most of the various services (FTP, Gopher, etc.) on the Internet. (Net.Genesis and Hall, 1995)

Web servers must be running a 'server' software package which provides all the functionality required for the job. As indicated above, a Web client computer must be running 'browser' software (such as Mosaic). The functions of client and server can reside on the same computer, however, it is more common for these functions to run on separate computers. The client/server model is what makes it possible for computers connected on the Internet to provide services to a multitude of other computers. (Liu, et al., 1994)

C. WORLD WIDE WEB

The unique aspect of the WEB that differentiates it from other Internet services is its ability to 'hyperlink'. Hyperlink means that a document has 'pointers' in the form of HyperText (the bold text in Figure 1) to related documents or other forms of information such as multimedia files (graphics, sound, video, etc.). The marriage of the two concepts hyperlink and multimedia has given rise to the concept of 'Hypermedia'. (Net.Genesis and Hall, 1995)

Because thought and communication patterns are associative we commonly link words, sounds and images during our intercourse. Associations can be direct or supportive (providing amplification) as well as tangential. The ability to provide a means within documents to hyperlink, at will, to associative 'hypermedia' is a very powerful tool and a prime reason for the phenomenal success of the WWW. (Net.Genesis and Hall, 1995)

(Note: Gopher also allows linking to distributed servers to retrieve information. However, the link is via the text menu only and not within the body of a document. Gopher also only supports text files and not the multitude of information formats (graphics, etc.) that the WEB does.)

Also, as mentioned previously, the WWW integrates most of the other services available on the Internet. Separate client applications are no longer needed to access FTP,

Gopher or WAIS. Web browsers provide the inter-connectivity needed to accomplish this 'transparency' to the user. (Liu, et al., 1994)

It is the WWW standards and protocols that give it these abilities. The Web is primarily defined by four protocols - HyperText Transfer Protocol (HTTP), HyperText Markup Language (HTML), Uniform Resource Locators (URLs), and Common Gateway Interface (CGI). Clients and servers on the WWW use these protocols to locate, access and display information. (Net.Genesis and Hall, 1995)

1. HyperText Transfer Protocol (HTTP)

HTTP is the supporting protocol for the WWW. It is "...a protocol for transferring information with the efficiency necessary for making hypertext jumps. The data transferred may be plain text, hypertext, images, or anything else." (Berners-Lee, et al, 1994). HTTP is a client-server model protocol similar to the FTP protocol, however unlike that protocol it is 'stateless' and 'connectionless' (Net.Genesis and Hall, 1995).

a. Stateless

A stateless protocol is one in which there is no record, or 'memory' of a connection from one request for information to the next. Each request is a 'new' request with no reference to any previous requests.

To compare, FTP maintains state. For example, when you log onto an FTP server it 'remembers' information such as what directory you are in or what file transfer mode you have selected. The next time you make a request it responds in accordance with this information. With HTTP there would be no memory of the directory or the transfer settings.

The advantage of a stateless protocol is that the protocol can run faster because it does not have to maintain extra information. On the other hand more information must be transferred with each connection to report necessary data from prior transactions. (Net.Genesis and Hall, 1995)

b. Connectionless

A 'connectionless' protocol is one which does not maintain a connection between requests. After a client has made its request and the information is transferred, the connection is broken. Prior to each new request for information a new connection must be established.

To again use FTP for comparison, when an FTP request is made two connections are established, one for controlling the connection and one for transferring data. The first connection is maintained as long as the user is logged on. The second connection is activated only during data transfers. (Liu, et al., 1994)

The advantage of a connectionless protocol is that it is efficient. Since servers can only have a finite number of connections open at one time, any connection that is idle (if for example the user is reading or away from the computer) is a waste of resources. (Net.Genesis and Hall, 1995)

A disadvantage is that it can take time to continually reconnect if , for example, a client is downloading a Web document with numerous in-line graphics. To speed the rate of data transfer some browser's (such as Netscape Navigator) open multiple connections and receive the data in parallel. This results in a faster download. This approach can be a problem if there is insufficient bandwidth resulting in a bottleneck. (Net.Genesis and Hall, 1995)

The multitude of connections or 'hits' at a heavily used Web site can seriously affect that server's performance. This is a central issue when building a site.

In addition to being stateless and connectionless, HTTP is also a very simple protocol. Few 'operations' are necessary to carry out a request transaction. The advantage of this is that it can handle a large number of requests efficiently and HTTP servers software can be small and simple. (Net.Genesis and Hall, 1995)

There are four parts to an HTTP transaction:

- 1) The client establishes a connection to the server.

2) The client sends a request to the server. The request includes all the information required by the server to carry out the transaction. Among other things this information lets the server know which Internet services the client can accept (FTP, Gopher, WWW, etc.).

3) The server sends a status response to the client (indicating it's ability to comply with the request) along with the requested information if available.

4) The connection is broken by either the client or the server.

2. HyperText Markup Language (HTML)

Documents on the WWW are 'written' in a HTML. As explained by L. Aronson in *HTML Manual Of Style*, HTML specifies the grammar and markup tags that, when inserted into a text documents, tell Web browsers how to present the documents. "The term markup came from the publishing industry, where it refers to the coded typesetting instructions inserted into a manuscript by an editor." HTML is an example of Standard Generalized Markup Language (SGML) which originated at IBM in the late 1960's as an attempt to solve the problem of moving documents between different computer systems.

Because of it's heritage, HTML works on the same principle used in word processing programs. For instance, 'marks' were inserted into this page to instruct the word processor how to present or format the document. These marks dictate the spaces between words, paragraph indention, bold print, etc.

Additionally, markup tags are also used to hold 'address' information for resources. In order to retrieve a hypertext resource the location of that resource must be known. Markup tags format hypertext so that when activated (clicked on) the resource is located and retrieved using Uniform Resource Locators (URLs).

3. Uniform Resource Locators (URLs)

URLs are central to the WWW architecture. To easily access sources of information anywhere on the Internet it is essential to have an addressing scheme that

scales easily and is independent of any particular network configuration. URLs provide that function. (Berners-Lee, et. al., 1994)

The URL naming scheme provides four basic pieces of information needed to retrieve information: The protocol used to reach the target server (HTTP, FTP, etc.) , the address of the target server, the directory path to the information within the target server, and the name of the information desired. (Liu, et al., 1994)

For example, the URL of the resource called 'Explore the Internet' at the Library of Congress is *http://lcweb.loc.gov/global/explore.html*. The first part of the URL - *http* - tells us that HTTP is being used (this could be FTP, Gopher, etc. depending on the protocol in use). The next section - *lcweb.loc.gov* -specifies the Internet address of the server at the Library of Congress. *Global* is the path within the server to the document, and *explore.html* is the name of the document.

One of the biggest advantages of URLs is that they provide a single, uniform system for identifying any resource on the Internet (such as Telnet, FTP, etc.). Future services will also be accessible. For this reason WWW browsers are considered to be a universal Internet access tool.

4. Common Gateway Interface (CGI)

Most, but not all, of the information on the WWW is in the form of 'static' HTML documents. These files are created prior to being used and are then placed on a servers hard drive from which they can be retrieved and then displayed. They are considered static because the content does not change unless physically updated by the author, site administrator, etc. (Net.Genesis and Hall, 1995)

The alternative to a static HTML document is one that has been generated on demand ('on the fly') based on the request of a client. Reasons for generating 'dynamic' HTML documents include services such as conducting database searches, ordering merchandise, personalizing documents and providing feedback. Additionally, although Web browsers can directly access most Internet services there are a few such as Archie,

and in some cases WAIS, that they cannot. To enable browsers to access the information in these services a 'translation' method is needed (Liu, et al., 1994).

The Common Gateway Interface (CGI) provides the means by which HTML documents can be generated to provide any of the above services. CGI is a standard for allowing a server to interface with custom computer programs that generate HTML documents.

a. Scripts

The custom programs are known as CGI 'scripts' (or just 'scripts') or alternately as Gateways. If the program is written to generate HTML documents based on input from a client it is referred to as a CGI script. Alternatively, if the program is written to make inaccessible services (such as Finger or WAIS) available to a Web browser it is referred to as a Gateway. In any regard, both work the same way and provide the client browser with an HTML document. (Liu, et al., 1994)

Script programs (and Gateways) can be written in 'scripting languages', such as Perl, or they can be written in a regular programming language such as C++ or Visual Basic. These programs are written to specifically enable some feature such as conducting database searches.

Input to CGI scripts (and Gateways) are commonly collected with a 'form' or via 'queries'. Forms are Web documents that 'capture' information entered by a user on a client browser. (Typically Web forms resemble paper forms and so are intuitive to the user.) Queries do not necessarily resemble a form but capture the data in the same manner. Forms are generally used to collect larger amounts of information, whereas a query may only collect one piece of information such as a search string.

The advantage of Scripts is that they allow for a truly interactive Web site. The ability to write custom programs enables site administrators to become very creative with how services are displayed, accessed, etc.

However, a potential disadvantage is that scripts put an additional computational load on the Web server. Given enough traffic this can significantly affect turnaround time for the information and may dictate additional infrastructure. Additionally, if not written carefully they can introduce 'security holes' into the Web server (Liu, et al., 1994).

D. CONCLUSION

The client server model is the foundation for the WWW. It's functionality directly affects an issue that is prime concern when establishing a Web site - the number of hits received by a site. The number of connections a WWW server must handle is a major factor in determining the infrastructure required. As we will see this can be one of the most difficult requirements to gauge.

Additionally, how the requests are handled within the server can also play a major role in defining site requirements. If dynamic HTML documents are served there will most likely be an increased need for additional processing power. This will dictate a greater infrastructure requirement.

III. PROBLEM DESCRIPTION

The issues associated with selecting the infrastructure for a Web site revolve around two questions - the anticipated traffic (number of connections or hits) the site will receive and the purpose of the site (Tabibian, 1995). Both these issues taken together will dictate infrastructure requirements. (Note: Cost is also an obvious issue. However, aside from the comments in the Cost Section below and general comments in various other locations it will not be directly addressed as an issue.)

The advantages to carefully determining what infrastructure is needed, which includes planning for reasonable growth, is a properly running Web site and a cost effective investment. The alternative can be an expensive investment that is insufficient or inappropriate to meet the demands of the site.

Prior to discussing the details of infrastructure requirements in Chapter IV, the issue of how certain Web sites and published literature handle infrastructure requirements definition will be examined.

A. SITE SURVEY RESULTS

During the research for this thesis informal surveys of several existing WWW sites indicated that requirement definition was largely absent. For instance, The Naval Command, Control & Ocean Surveillance Center (NCCOSC) in San Diego, CA. runs a Web site called the PLANET EARTH HOME PAGE.³ This site is an excellent source for locating resources on the WWW. It receives many thousands of hits each day. When interviewed, the Web master for this site indicated that the site initially ran on existing equipment and - "As traffic increased, the system just totally bogged down." (Evans, 1995). The equipment in use and the amount of traffic at the site are secondary to the

³ PLANET EARTH HOME PAGE at http://www.nosc.mil/planet_earth/info.html.

point of the current discussion - that a detailed requirements analysis was not conducted to determine site requirements.

A similar response was produced by an E-mail interview with the manager of the network operations center at Massachusetts Institute of Technology (MIT) when he was asked if a requirements analysis was done for their Web site: "Nope. We just took some hardware we had laying around (DEC 5000's) and went to work. As the workload increased, we went looking for better software and hardware." (Schiller, 1995)

Both of the two organizations above had existing equipment available which was 'pressed' into service to establish a site. Common sense tells us that in a situation such as this it is usually easier to justify creation of the site with equipment on hand than it is to request additional funding for new equipment. (Additionally, as MIT is an educational institution there is benefit for them in learning by just doing.) If undertaken with the correct mind-set this approach can provide benefits, as will be discussed later. Nonetheless, these examples illustrate the response received from all sites surveyed. There was a distinct lack of requirement analysis with regard to establishing the sites.

As stated by Tom Littlejohn at the Library of Congress (LOC), one of the reasons organizations do not conduct a detailed requirements analysis is - "Because it (a Web site) is not viewed as a strategic investment. That is why older spare equipment is used." This comment was offered during a LOC site survey to determine if a requirements analysis was conducted prior to LOC offering their first Internet (Gopher) server. As indicated by his comment, they did not conduct an analysis and also used existing equipment to come on-line. However, he agreed that it quickly did become strategic. In their case they were forced to upgraded within six months.

Another apparent reason is that estimating the anticipated traffic the site will receive is subjective and can be difficult to predict. Because of this it is viewed as easier to 'just do it' and handle problems as they occur. Part of this mind set may be the result of a cavalier attitude on the part of some information systems (IS) professionals, and ignorance of the need on the part of neophytes.

Defining the other issue - the purpose of the site - is much easier. However, determining exactly why the site is needed and what will be offered does require deliberate consideration. If the above mind-sets are prevalent, this side of the requirements definition issue may also receive perfunctory treatment.

Another observation resulting from the surveys pertains to consultants. While consulting firms do undertake a requirements analysis of a customer's needs they may or may not provide a client with an *optimal* solution. The solution may have more to do with a given product line (theirs!) than an honest appraisal of needs. It is likely that the solution will be acceptable, but it may not be the most appropriate or cost effective investment for that particular organization. The best solution may be a competitors infrastructure solution. Also the requirements analysis itself may be a mental iterative process vice a formal evaluation, and therefore not subject to easy review or scrutiny. (Hunter, 1995)

The most interesting response to the site surveys was a philosophy articulated by Dave Norman at the Naval Postgraduate School. Mr. Norman is the Director of Computing Services and has been involved with the Internet since very early in its inception. When asked if he had any 'rules-of-thumb' for determining his hardware needs he responded - "Figure out what you can afford and buy one step higher." His point is that instead of purchasing, for instance, a 'loaded' 486, buy a 'stripped down' Pentium. It will then be possible to expand the Pentium as the need arises. It is also *much* easier to justify additional funds to augment an existing capability than it is to totally junk a new, but inadequate system and start over.

This is a useful, realistic approach for an existing system as it builds upon intimate knowledge of a current site's infrastructure, load and expanding needs. Unfortunately, it does not assist with defining *initial* requirements.

B. LITERATURE REVIEW RESULTS

Literature review consisted of surveying computer industry periodicals, recently published books and on-line material which dealt with the subject of establishing Web sites. Because of the newness of the subject there are a limited number of books available dedicated to establishing Web sites. On the other hand, because of the popularity of the subject this is rapidly changing. Overall the literature review was more useful for assessing and evaluating site needs than the site surveys. However, no single source reviewed provided a comprehensive guide to determining site infrastructure requirements. It was necessary to review several sources to properly address the issues.

In general, the books reviewed covered establishing a web site from a broad perspective. The periodicals reviewed fell into two rough categories, one dealing with the general topic of establishing a web site, and the other dealing with a specific subset issues such as picking a connection provider or RAID (Redundant Array of Independent Disks) technology. The on-line material reviewed also dealt with specific subset issues.

The first category of periodicals and the majority of books provided general background information for establishing a site, providing information on issues such as server software configuration, selecting a Internet access provider, and HTML authorship. As would be expected due to their length, the books provided much more depth. Both provide some valuable "rules-of-thumb" buried within the material. However, they did a much better job at categorizing *what was available* than in assisting with determining *what is required* for a specific site.

An example is *Running a Perfect Web Site* by David Chandler. Overall this is a good book which provides useful discussion and information on the entire range of subjects needed to understand what goes into establishing and maintaining a Web site. However, with regard to the issue of *defining* the needs of an individual site it gave very little guidance. To illustrate, during the discussion on leased lines (connection speeds) it states - "If you're setting up your own server, you will want a connection fast enough to handle your anticipated traffic...". The discussion continues on the next page with regard

to 56Kbps (connection speed) leased lines - "If your site contains large files or graphics, delays in loading pages will be noticeable and multiple simultaneous connections will slow it to a crawl." Nowhere in the book does it address *how* to anticipate a sites potential traffic load or what would be considered heavy or light traffic. Neither does it define "large files" or quantify "multiple simultaneous connections". Because of this an organization desiring to establish a site (and having no prior experience with the Web of course) could not determine what size connection would be sufficient for their needs.

Similarly when dealing with hardware selection the book states - "How much hardware you need for a Web Server depends entirely on your application." In a subsequent chapter it says - "...the Web server hardware itself is the single most important factor in determining performance." And - "...the single most important factor in determining how long it takes a Web sever to respond to a request is the processor speed." Nowhere does it quantify which hardware would be suited for which uses. It does, however, state that a site with "...several hundred users." could be served by a 486/33 with a fast hard drive or an equivalent Mac, and that for "...several thousand users..." a Unix workstation such as an HP 715/50 would be needed. Although this provides some guidance it is very little to work on. It would be difficult to accurately define an organizations infrastructure requirements from these two sentences.

Two other (book) examples - *Marketing on the Internet* by Michael Mathiesen, and *Building Business Web Sites* by Adam Blum - produced similar results.

The second category of periodicals (specific subset issues) and the material retrieved from on-line was useful for investigating individual issues in detail, such as determining connection (bandwidth) requirements. This information, together with the "rules-of-thumb" taken from the first category of periodicals and the books, will be used where applicable to answer Web site requirements issues.

The single most useful reference was the book *Build A Web Site* by Net.Genesis and Devra Hall (which has already been quoted in the Chapter II). Besides the information it provides, this book is noteworthy because one of its authors (the "Chief Technologist" for Net.Genesis) is Matthew Gray who founded the original MIT Web site and created the

World Wide Web Wanderer (the first WWW cataloging 'robot'). Net.Genesis itself is a very successful company creating commercial Web sites for such organizations such as ESPN, DEC and IBM. Additionally, Tim Berners-Lee has written the foreword which lends an air of credibility to the entire work.

Primarily written as a programmers guide to "creating, building and maintaining a Web presence" it provides programming code and tips for Web sites. More importantly it gives advice on determining a sites potential traffic as well as providing the only rule based heuristic for selecting hardware that was found during research for this thesis. The heuristic is the subject of Chapter V.

With the exception of *Build A Web Site* and some of the issue specific material, most of the literature reviewed again illustrates the tendency toward insufficient requirements definition with regard to establishing WWW sites.

C. COST

As in most aspects of life, cost is an obvious constraint. If it were not we would all possess the 'best' of everything. Web sites are no exception. Once requirements have been defined if the resulting infrastructure exceeds the budget constraints, then the 'perceived' requirements are clearly out of line with fiscal reality, and a re-examination of the site's function and scope is necessary.

An important point to realize is that people cost more than equipment. Jeff Schiller (network operations center manager at MIT) is quoted as saying in a LAN TIMES article - "The most expensive part of having your own Web server is the [technical] expertise..." "Computers and hardware are cheap compared to the cost of hiring an expert." (Armstrong, 1995). Due to this some firms find that it is more cost effective to outsource the entire enterprise (Wilder, 1995). However, a company that possesses its own Web site has greater control over document management (Armstrong, 1995).

Another implication of this is that if an organization relies on an outside contractor to provide the expertise required to establish and/or run a Web site it can become very

expensive if forced to upgrade as a result of inappropriate or insufficient requirements analysis (a point of view the contractor may not necessarily share).

Depending on the anticipated traffic and the purpose of the site, the cost of the infrastructure requirement (not including personnel) can range from \$2,000 to \$100,000 or more (Tabibian, 1995). Where feasible general costs will be listed to provide insight into the issue. A detailed cost analysis would obviously be needed for any given infrastructure solution.

D. CONCLUSION

The biggest lesson learned from the site surveys is that initial requirement evaluations are not being conducted. This may be the result of oversight, ignorance or indifference.

The literature review revealed that most material gave varying, but generally acceptable, descriptions of *what was available* but not *how to define what was needed*. They manifested a lack of guidance on actually defining needs and then selecting infrastructure based on those needs. This lack of emphases on defining requirements may 'feed' attitudes 'in the field'. Part of the problem may be that due to the relative newness of the subject there is a limited number of books available on the topic. The popularity of the WWW is rapidly changing this situation. Hopefully, as more material becomes available some will address the issue.

The potential penalty for not conducting a proper assessment of requirements is the same as for any venture, a substandard product and poorly leveraged investment.

IV. INFRASTRUCTURE REQUIREMENTS

To determine infrastructure requirements for a Web site the following questions must be answered:

- 1) How much connection bandwidth (size of the 'pipe')?
- 2) How much CPU capacity?
- 3) How much memory?
- 4) How much hard disk storage space?
- 5) Which operating system?
- 6) Which server software (NCSA, CERN, NT, etc.)?
- 7) How to provide system and stored data integrity/redundancy?

By providing answers to the two driving issues, anticipated traffic and purpose of the site, solutions for the first four questions can be obtained.

A solution to the first question - how much connection bandwidth - can be calculated using a formula presented in this chapter. The second and third questions - how much CPU capacity and how much memory - can be answered by employing the hardware heuristic in Chapter V. Similarly, a reasonable estimate can be made to determine a solution to question four - how much hard disk storage space.

The answer to question five - which operating system - may be driven by the hardware solution obtained. However this can be a subjective issue which can take on religious tones! Question six - which server software - can also be subjective. Some server software packages are better suited for particular jobs than others. More will be said about these later in this chapter.

The question of integrity and redundancy (question seven) is not actually required to set up a site. It is, however, a very important issue and should be an integral part of the planning for any site as with the other requirements, the answers to the two primary issues will drive the level of integrity and redundancy needed. (Note: Security is also a very fundamental question and must be taken into consideration. However, as mentioned in Chapter I, it will not be addressed in this thesis.)

The rest of this chapter will discuss how to obtain answers to the two basic questions of traffic and purpose. Issues associated with the size of the connection (question one), hardware requirements (questions two - four), and software requirements (questions five and six) will also be covered. Integrity and redundancy (question seven) will be discussed in Chapter VI.

A. ANTICIPATING TRAFFIC

Perhaps the single most important issue to consider is how much traffic the site will receive - the connection rate. How much will it be browsed by clients? This issue (along with the purpose of the site) drives connection, and platform (hardware and software) requirements. Therefore, it needs to be considered carefully. Different sites can experience vastly different loads, ranging from a handful of hits a day to hundreds of thousands. There are a variety of ways to approximate the potential hit rate. (Net.Genesis and Hall, 1995)

It is essential that an analysis be carried out to determine who the clientele (audience) are and what will they be served. This question is directly related to the next section - Purpose of the Site; what will it provide and to whom? If you are providing arcane information to a very select group the usage at the site will be light. If on the other hand you are providing access to valuable information and popular services, as the LOC site is, usage could be extremely heavy. Also, how 'unique' is the information? If the site will be one of a handful to offer this information it will likely receive heavier use (Net.Genesis and Hall, 1995).

One useful approach to determining potential traffic is to study USENET newsgroups that cover topics similar to the content intended for the new site. One group - news.lists - provides estimates as to how heavily any particular newsgroup is read. Also some newsgroups maintain archives, by examining these and the FAQ (frequently asked questions) file, references can be found to mailing lists and interest levels for particular information can be estimated. By analyzing this information you can not only gauge

potential traffic but also gain insight that can be used for designing the site. (Net.Genesis and Hall, 1995)

Another technique is to scrutinize similar Web sites. In many cases the site will list its traffic level, if not, most site administrators will be willing to provide this information upon request. As with newsgroups, this can also provide valuable insight into what information is in demand. (Net.Genesis and Hall, 1995)

How heavily the site is publicized can also make a substantial difference. "It is quite clear that advertising a site on important lists like the NCSA 'What's New Page' and Scott Yanoff's 'List of Internet Services' has a very direct and immediate impact on how many people use a site." Listing on one of these services can cause the site to receive thousands of connections per week during the month the announcement is made. Other sources of publicity include posting to newsgroups as well as other Web sites that are willing to list you. (Net.Genesis and Hall, 1995)

In general, if the site contains limited information (such as a simple home page) and/or is not well publicized it will likely receive very light traffic - a few hundred hits a day. If it has more useful information and is well advertised it can receive thousands of hits per day. Most sites fall in this range. Few sites receive hundreds of thousands or millions of hits a day. These are usually main players in the WWW such as Netscape, NCSA (National Center for Supercomputing Applications), etc. (Net.Genesis and Hall, 1995)

Although the above methods would prove useful for *estimating* potential traffic, the most accurate method would be to actually measure usage of a site. If approached from the correct perspective, this is where the use of existing equipment can be employed as an effective requirements analysis tool. Instead of trying to estimate needs this equipment can be used to accurately measure the new sites requirements. Since the equipment is a 'sunk cost' this approach can be a cost effective method for determining a sites true infrastructure needs. The danger with this approach is that it may be viewed as a permanent solution instead of an interim arrangement resulting in a lack of financial commitment toward upgrading to the sites real requirements.

If an organization does not have spare equipment on hand an alternative solution would be to create the Web site contents and then rent space on a providers WWW server for the first six months or so of operation. During this period the actual load on the site can be measured. Based on this information, along with a growth analysis, the requirements for the connection, and hardware can be determined as well as which operating system and server software suits would best satisfy the requirements.

Finally, this thesis makes the assumption that the intent of the proposed Web site is to serve information to the Internet. However, many organizations find that Web servers and HTML are an excellent means of distributing information *within* the organization. If this is the case it is an easy matter to estimate the usage, not only is the employee count known, but the traffic level on the organization's LAN (local area network) will also be known and is an actual measure of use.

B. PURPOSE OF THE SITE

The best way to decide on the appropriate platform and software is to decide the purpose of your server. In our testing, we found that there are currently two ways you can utilize a Web server. You can include basic documents that convey information and provide links to other sites. Or you can set up a more complicated Web server that integrates search engines and forms. In the future expect - perhaps most impressive - a third alternative, which will add security to Web servers so they can conduct financial transactions on the Internet. (Tabibian, 1995)

It is absolutely necessary to define the site purpose. The use relates directly to anticipated content of responses (and therefore data transfer rates), reliability issues and platform (hardware and software) issues. Again, a market analysis approach should be used. Potential uses are:

- 1) Commercial (product and/or order information).
- 2) Corporate and government (organizational information, product and services information, and/or public relations).
- 3) WWW/Internet Service Provider (such as Netscape or Yahoo)
- 4) Education and/or research.

- 5) Internal (internal organizational use)
- 6) Private (homepage)

(note: Depending on the mission, military organizations would fall under one of above categories.)

The purpose of the site can be considered to be a content and audience issue - what size responses (data files) will the site provide and to whom?

Questions that may assist in determining this are:

- 1) Will the site provide static HTML text documents or dynamic documents?
- 2) Will the site provide hypermedia (audio, video, and movies)?
- 3) Will gifs (pictures) be imbedded in the documents? If so, what size and how many?
- 4) Will the site act as a database front-end?
- 5) What are the likely technical limitations of the intended audience?
- 6) What is the 'complexity' of documents that the audience can accommodate?
- 7) Is there a need for 24 hour, seven days a week availability?

For example, if the intended audience is technically sophisticated corporations, the assumption can be made (or perhaps a definitive answer obtained) that they will have large bandwidth capabilities (such as a T-1), and therefore, the size of the files presented (data transferred) will be less of an issue. On the other hand, if the intended purpose of the site is for commercial advertising targeting individual homes, the documents and embedded Gifs must be kept to a reasonable size because of the bandwidth limitations of home modems (up to 28.8Kbps at present). To illustrate, on a 14.4 Kbps modem a ten second sound file can take several minutes to download and a one minute movie file may take an hour (Chandler, 1995). It is therefore desirable to keep the files delivered pertinent to the intended audience.

As stated in Chapter I, the question of Web site content is the most fundamental issue associated with creating a WWW presence, and the success or failure of the site can ride on this issue. Due to this, and because the infrastructure requirements are linked directly to what the site is for, careful deliberation must be given as to its purpose and scope. Additional guidance for determining what is appropriate for a given use may be gleaned by studying existing sites.

C. CONNECTION

- *How much connection bandwidth (size of the 'pipe')?*

1. Line Options

In order to be accessible a Web site must have a connection to the Internet. Connections may be via a switched line or a leased (dedicated) line. A switched line is similar to the telephone service provided to a house. Each time a connection is required the call is routed over available circuits. Because the line is shared with other customers, the quality (transmission error rate and bandwidth) and availability of the line are not guaranteed (Blumenfeld, et al., 1995).

Leased lines are dedicated connections that are rented from a service provider. They provide 24 hour availability for the site as well as delivering a consistent level of performance. The recurring cost for a leased line is more than that of a switched line and is based on the level of performance and the rates of the provider. Generally to provide an acceptable level of performance for a dedicated Web server and give it 24 hour availability leased lines are the only viable alternative.

Table 1 lists a selection of Web site connection options and associated data transfer rates. A 56K line is usually the most economical and may suit a smaller sites needs. If however, the site provides large files (due to graphics, video, etc.) or experiences heavy traffic this speed will not be sufficient (Chandler, 1995). Cost for a 56K line range from \$300 to \$400 a month plus approximately \$500 to start-up (Net.Genesis and Hall, 1995).

ISDN is not a leased line, it is a dial-up connection. However, due to its functionality it may provide a viable alternative to leased lines. ISDN's Basic Access Service is composed of two 64K 'B' channels and one 16K 'D' channel. The two B channels are used for all data transfer (voice and/or digital) while the D channel is used for control data. Cost for ISDN can vary widely.

TYPE	Data Rate (Bits Per Sec)
56K	5,600
ISDN	128,000
T-1	1,544,000
T-3	44,376,000
T-4	274,176,000

Table 1: Connection Options

T-1's are a common option. Costs can range from \$2000 to \$5000 per month plus an estimated \$3000 to \$8000 for installation (Net.Genesis and Hall, 1995). A T-1 can be subdivided into 24, 64Kbps individual channels, which is referred to as a *fractional* T-1. One or more of these channels can be leased for a corresponding reduction in monthly rates.

It should be noted that even if a 56K line is sufficient to currently handle the anticipated loads of a site it may be prudent to obtain one 64K channel of a fractional T-1. The reason for this is that if upgrading is required at a latter date the cost to upgrade to an additional T-1 fraction is relatively little. However, moving from a 56K line to a fractional T-1 is expensive because the installation cost must again be paid. Also, to upgrade to additional T-1 fractions can be done in a day or so, while upgrading from a 56K to a fractional T-1 could take weeks or months (Net.Genesis and Hall, 1995).

Beyond T-1 the options become too expensive for most individual Web sites. T-2's are the next level up in capacity, however, they are used within the 'phone system' and are therefore not available. T-3's and T-4's are generally used by Internet providers and as major Internet backbones. (Chandler, 1995)

2. Bandwidth Requirement

The level of service chosen depends on the amount of anticipated traffic (connection rate) and the purpose/content of the site (data transfer rate). The issue is how much *bandwidth* is required to satisfy the sites connection rate and data transfer rate.

Bandwidth refers to the speed or capacity a line has for transferring data. For instance, a T-1 has a bandwidth of 1,544,000 bits per second, and a 56K line has a bandwidth of 56,000 bits per second (see Table 1). It is fairly easy to estimate the required bandwidth given correct estimation of the traffic, the size of the files that a site will transfer (the 'content'), and the desired *latency*.

Recall the client server model, the client sends a request (query) and the server sends back a response. Latency is the round trip time for the client query to get to a server, be processed and then for the response to be sent and received by the client (Net.Genesis and Hall, 1995). As to the problem of determining what level of connection (amount of bandwidth) is required, latency can be thought of as how long the client has to wait to receive the requested data *after the request has been received*. This will be referred to as the file transfer time.

The reasoning for this is that queries are relatively small (Liu, et al., 1994). Therefore, although the time it takes for the request to be received can definitely matter to a client (especially if the target server is so busy it takes a long time for the request to be accepted) if sufficient bandwidth exists to return the request to the client in a 'reasonable' amount of time, then ample bandwidth should exist to cover the relatively small client queries.

Based on computer command line studies it was determined that five seconds was the amount of time people would wait before becoming impatient with the system (Meyer, 1995). This figure serves as a useful reference but can be altered to provide a more reasonable goal, especially if large hypermedia files are being served. The actual target time depends on the level of service desired for the site.

Equation 1 was adopted from *Business Data Communication* by Jerry Fitzgerald and can be used to estimate the bandwidth required by a WWW site. This formula does not account for control characters transmitted or retransmissions caused by errors or delays. To account for this ten percent can be added to the estimation. (Also, it does not account for any internal LAN delays or delays resulting from the service provider.)

$$\text{File transfer time (to return requested information to Client)} = \frac{\text{Number of Records} \times \text{Number of Bytes per Record} \times \text{Number of Bits per Byte}}{\text{Bits per Second Transmission Speed}}$$

Equation 1: File Transfer

In the calculations below Equation 1 has been rearranged to find the transmission speed. This is based on the assumption that the file transfer time (latency) will be a 'given' - selected by the site administrators to provide a desired level of performance (such as 5 or fifteen seconds).

Additional rules-of-thumb that will assist in estimating bandwidth are (Gray, 1995):

- 1) The average size of an HTML file is 10K.
- 2) Peak traffic hit rates are roughly double the daily average.

To illustrate: a potential Web site serving 10K static HTML documents has estimated the traffic (hit rate) to be an average of 360 connections per hour and wants to provide a response time of five seconds.

$$\text{Transmission Speed} = \frac{\text{one record} \times 10,000 \times 8 \text{ bits per byte}}{4.5 \text{ seconds file transfer time}}$$

To calculate this the problem was set up as follows:

- 1) 360 connections per hour = a peak of 720 hits per hour.
- 2) 720 hits per hour divided by 60 minutes in an hour = 12 hits per minute.
- 3) 12 hits (peak) per minute = one hit every five seconds.
- 4) Static HTML documents = 10,000 bytes.
- 5) One byte = eight bits.
- 6) Five seconds minus 10% (for transmission error, etc.) = four and a half seconds file transfer time.

As the estimated bandwidth required is 17,778bps, this site can easily be served by a 56K connection.

To illustrate another example: a site serving 50K dynamic HTML documents (via scripts), has an estimated average hit rate of 480 connections per hour and wants to provide a response time of 15 seconds. As can be seen, this site would require two fractional T-1 channels (or an ISDN connection).

$$\text{Transmission Speed} = \frac{\text{four records} \times 50,000 \times 8 \text{ bits per byte}}{13.5 \text{ seconds file transfer time}}$$

To calculate this the problem was set up as follows:

- 1) 480 connections per hour = a peak of 960 hits per hour.
- 2) 960 hits per hour divided by 60 minutes in an hour = 16 hits per minute.
- 3) 16 hits (peak) per minute = four hits every 15 seconds.
- 4) Dynamic HTML document = 50,000 bytes (in this example)
- 5) One byte = eight bits.
- 6) 15 seconds minus 10% (for transmission error, etc.) = 13.5 seconds file transfer time.

It is important to realize that if a client in the above example is using a 14.4K modem (bandwidth = 14,400 bps) the response time as perceived by the client would be over 30 seconds (one record $\times$ 50,000 $\times$ 8 / 14,400 = 28 plus 10%). This is why it is essential to keep in mind who the potential audience will be.

If the site's actual document sizes are known (perhaps the content has already been created) then a better estimate can be obtained based on the actual size, instead of the estimating the file size. Also, do not forget to factor in the size of in-line graphics, this can significantly increase the size of a file (Liu, et al., 1994).

Finally, if the site will be serving a selection of documents (such as static documents and script generated information), then a determination must be made as to the ratio that these documents will be retrieved (Liu, et al., 1994). For example, from the previous example, if on average every 15 seconds one static HTML document (10K) and

three script driven documents (50K each) are delivered, then the total document sizes must be added together:

$$\text{Transmission Speed} = \frac{94,815 \text{ bps} = \{(1 \times 10,000) + (3 \times 50,000)\} \times 8 \text{ bits per byte}}{13.5 \text{ seconds}}$$

As discussed in Section A, above, studying other Web sites may assist in estimating this.

3. Additional Requirements

In association with obtaining a connection to the Internet, an IP address and various hardware will also be required.

An IP address is the address of the site on the Internet. To obtain an address a request must be submitted to the Internet Networking Information Center (InterNIC). It can take two or three weeks to process the request (Blumenfeld, et al., 1995). It is considered desirable to obtain a domain name, such as microsoft.*com*, because it is easier to locate such a site as only the company name must be known (Tabibian, 1995).

Equipment needed for the connection include routers and DSU/CSU's. Routers act as the interface that allows the Web site and the Internet to communicate by controlling traffic flow. Among other things they maintain the address routings tables for routing messages into and out of the site. The router must at least be able to handle the speed of the sites connection (Chandler, 1995). Prices for Internet routers range around \$2500 (Chandler, 1995).

A DSU/CSU (Data Service Unit/Channel Service Unit) basically serves a function similar to a home modem. It can 'condition' digital signals to, among other things, reduce noise, distortion and errors (Fitzgerald, 1993). A DSU/CSU is installed between a router and the connection to the Internet. Prices range from \$250 to \$3000 for a 56K unit, to \$1200 to \$2500 for a T-1 DSU/CSU (Chandler, 1995).

IP addresses and the hardware above are mentioned to provide further insight as to what is required to establish a connection. They will not be covered in any greater detail.

D. HARDWARE

The level of hardware required will be driven by the anticipated traffic (connection rate) and the purpose/content of the site (data transfer rate), and will be determined using the heuristic in the next chapter. This section will provide amplifying information concerning the hardware.

1. Processing power - *How much CPU capacity?*

"Figuring out how much computational power a given server will use is even more of a guessing game" (than estimating traffic). This quote from *Managing INTERNET Information Services* expresses the position of much of the literature reviewed and is probably one of the reasons so few guidelines are available. The book goes on to say - "...WWW servers consume CPU in proportion to the number of queries they receive and the size of the files they process." This statement is echoed by the following quote from Jeff Schiller, "The amount of CPU power and RAM requirements depend on what type of data you will be transmitting and how many people will be hitting the Web server at one time." (Armstrong, 1995).

This again emphasizes why it is essential to estimate traffic and to determine the use of a site. The greater the number of processor-intensive functions a site serves the greater the platform requirements will be. However, as long as the server can keep pace with the speed of the connection, CPU induced bottlenecks should not be a problem (Armstrong, 1995). Examples of potential processor intensive functions exclude:

- 1) Forms
- 2) Image maps
- 3) Searches
- 4) Computations
- 5) Scripts

Until relatively recently the debate over which type of CPU - RISC or CISC - was more appropriate for a server was a on-going debate. The issue is now largely moot.

The debate centered around which CPU design approach was 'best' (best generally meant faster). CISC (complex instruction set chip) are typified by the Intel designs (386/486) and use a large set (complex) of CPU internal instruction to enable the CPU to carry out job. RISC (reduced instruction set chip) technology on the other hand is typified by use in UNIX machines by companies such as Sun Microsystems and Digital Equipment Corp. As the name implies RISC uses a smaller instruction set and therefore the CPU can perform in a job faster because it has a smaller pool of instructions to search through.

As pointed out in "RISC/CISC Debate Over: Customers Win" by Damian Rinaldi, a bigger impact on system performance arises from issues such as memory, operating system, disk and I/O subsystem, application mix and transaction loading. There are no assurances that if RISC manufacturers actually achieves a measurable performance edge that the user will experience an overall improvement in throughput. Also, with the current generation of Pentium chips and its follow on, Intel has already incorporated significant RISC like features and functionality. Currently then, the most important questions for the customer is not which type of chip is used, but can the system run the desired applications. (Rinaldi, 1995)

2. RAM - How much memory?

"The single factor that buys you the most speed is RAM, so get as much as you can afford." (Chandler, 1995)

"The more memory you have, regardless of the platform, the better your performance and the server's response time will be." (Tabibian, 1995)

RAM is like money - no matter how much you have, you can always use more. The reason more RAM is better has to do with the way a server functions. Each time a client sends a query, the server responds by creating a copy of itself to handle the request. This is called *forking*. The larger the hit rate the more copies are required to be open. Without sufficient RAM the server will use available hard disk space to temporarily act as storage for system memory. This is called *swapping* and is undesirable because it takes 1,000 times longer to access information on a hard drive than it does to access RAM. The

result is that the system performance will greatly slow down. As a rule the system should never have to perform swapping, except perhaps during peak loads. (Net.Genesis and Hall, 1995)

The heuristic in the next chapter lists recommended level of RAM associated with each level of computer. These RAM amounts should be considered a minimum level. Also, guidance is usually provided by the manufactures of the hardware, server software and operating systems.

3. Disk Space - *How much hard disk storage space?*

Hard disk space is not addressed directly by the heuristic in the next chapter, however the amount of disk space needed is driven by the size and complexity of the document served.

At the very least there is an obvious need to have enough hard disk space to accommodate all the files that will be offered. Usage logs, kept for the site will require 10 to 20 megabytes of storage per month (Chandler, 1995). Also, based on the RAM swapping issue, there should be several megabytes of spare disk space to allow swapping during very high usage. Beyond that there is very little guidance as to how much spare hard disk space to obtain. Tripling or quadrupling the space actually required to hold existing files should provide enough room for growth and any system use required.

E. SOFTWARE

As stated at the beginning of the chapter, the answer to which operating system and which server software package to use can be subjective issues. In many cases they will be driven by the hardware solution obtained. The following two sub-sections will present general information to assist in the decision.

1. Operating Systems - *Which operating system?*

Next to additional RAM the next biggest difference to system performance is the operating system (Gray, 1995). This fact was brought out during several of the site surveys conducted and was also mentioned in a paper obtained from the National Center for Supercomputing Applications (NCSA) at the University of Illinois (McGrath and Yeager, 1995). Evidently, some implementations of Unix handle the TCP/IP stack (programs) more efficiently than others. Unfortunately, no other literature was found which critically compared operating systems. This area is a good candidate for further research.

The usual debate about operation systems is whether Windows NT is a suitable platform or whether Unix is the only viable alternative for a robust server. NT has made in-roads and is now considered by some to be as viable as Unix (Blumenfeld, 1995). However, the current conventional wisdom (or perhaps prejudice) is still that for a extremely stable and generally more secure system some version of Unix is the answer (Campbell, et al., 1996). Other operating systems, such as Macintosh and Windows 3.11, are available but generally not considered viable for very demanding environments (Tabibian, 1995). Because of this if the site requirements call for a heavy duty machine, some version of Unix will probably be needed.

A useful Internet site for additional information on operating systems is, "Operating Systems on the Web", run by RWTH Aachen University of Technology, Aachen, Germany, found at: <http://www.lfbs.rwth-aachen.de/~sven/OS-Projects/>. This site provides an extensive list of links to worldwide sources of information concerning all aspects of operating systems. Another useful site is Yahoo, Operating Systems at: http://www.yahoo.com/Computers_and_Internet/Operating_Systems/. This provides a searchable index of a wide range of operating systems listing specifications and features.

2. Server Software - *Which server software?*

Web server software (often referred to as Web servers, HTTP servers or just 'servers') gives a server all the functionality required to operate as a Web site. In the past

year a plethora of Web server software packages have been made available. There are basically two approaches that can be taken to acquire server software, one avenue is to download a free version from the Internet, and the other is to obtain a commercial package.

Until last year downloading from the Internet was the primary means of obtaining server software. Two of the original 'servers' available were the NCSA 'server' and CERN 'server'. These two continue to be very popular with NCSA taking 41% and CERN taking 11% of the market in a recent survey (Hoffman, 1996).

The alternative route, purchasing a commercial offering, has shown a significant increase in the last year. The 'server' offered by Netscape, now has 13% of the market and WebSite from O'Reilly & Associates, Inc. has 4%. (Hoffman, 1996)

The above statistics were obtained from the WWW site "Web Servers Survey", by Paul Hoffman of Proper Publishing, at: <http://www.proper.com/www/servers-survey.html>.

Appendix A. Web Servers Survey, is a copy of this document and is provided as a ready reference on the current market use of server software packages. This information can be useful because more numerous servers will have 'longer track records' and thus problems and shortcomings will more likely be known. Also, it may be easier to find additional information (such as configuration information) and support for the more popular software.

According to this survey, Unix 'servers' still dominate the market. Some of the reasons Unix systems predominate are the same reason Unix operating systems are still seen as the superior choice - they are very stable, and provide superior security. However, Windows NT servers (such as WebSite) are making inroads, just as the NT operating system is. And just as the NT operating system is now considered viable by some, so are NT servers. One of NT's strong draws is its graphical user interface, and resulting relative ease of configuration and administration.

As stated, NCSA is the most widely used server software. Due to its 'speed' it is considered to be a good choice if there is a need to handle a lot of connections. The CERN software, on the other hand, is considered to make a good *proxy* server. (A proxy

server is one that acts as an intermediary between the client and the actual server the requested information is on. Proxy servers provide an extra degree of security for networks.) The CERN software also supports document caching, which means that information from a remote server or network is copied to the proxy servers own disk and retained for a set period of time. This provides faster response should that document be retrieved within the set period. (Net.Genesis and Hall, 1995)

Although free, one of the draw-backs of the Unix based software available on-line is that in-house Unix expertise is required to configure and operate the server. For 'turn-key' (ease of installation) solutions and continued customer support a commercial version may be a better option. Also, as financial transactions on the Internet increase newer commercial server software will incorporate security features such as authentication and encryption (Smith, 1995). The cost of commercial server packages ranges from \$100 to \$25,000 (Smith, 1995).

For a detailed look at a large number of free and commercial 'servers' another useful document, again authored by Paul Hoffman, and can be viewed on-line at: <http://www.proper.com/www/servers-chart.html>. This site provides a good overview of available servers and their features and also lists other useful links to additional Web sites.

Two other Web sites worth visiting are both listed in the "Web Server Comparison" Web page. The first is "World Wide Web Server Software" by the World Wide Web Consortium (W3C) at: <http://www.w3.org/hypertext/WWW/Servers.html>. The second is a searchable index of Web server software from Yahoo, at: http://www.yahoo.com/Computers_and_Internet/Internet/World_Wide_Web/HTTP/Servers/. Both sites review various server software.

F. CONCLUSION

The issue of how much traffic the site will receive along with the purpose of the site drives all the major requirements. It is absolutely necessary to conduct a deliberate study of these two issues to properly determine connection speed, and platform (hardware

and software) requirements. If the assessment of these two parameters is not reasonably accurate the resulting site infrastructure will likely be inappropriate to a corresponding degree.

With regard to traffic, it is very difficult to accurately estimate the connection hit rate. The techniques list should provide a reasonable, educated estimate of how heavy the load will be. However, if possible the traffic should be measured.

Determining the purpose of the site can be accurately determined. The type and size of content (files) that will be provided should be known. The intended audience should also be known especially if any sort of market analysis was conducted.

Given answers to these two questions it should be reasonably straightforward to calculate the required connection speed as well as the level of hardware required.

Choosing software is less objective. It will be driven to some extent by the site and hardware requirements. Beyond that it is largely a matter of preference.

V. HARDWARE SELECTION

As explained previously, the initial research approach was to build a heuristic from 'rules-of-thumb' learned through site surveys and literature review. Unfortunately little guidance was found, and the emphasis shifted from creating a heuristic to validating the one heuristic the research did turn up. This chapter will present the heuristic, describe how it was benchmarked, and discuss the validation method and results.

A. HARDWARE HEURISTIC

The heuristic in Figure 2 is taken from the book *Build A Web Site* by Net.Genesis and Devra Hall. It has been slightly edited and reformatted.

Using the requirements information determined in IV, calculate the 'hardware level' by following steps one through five. This 'hardware level' number is applied to the "Quickie Server levels" table to determine the level of hardware (computer and memory) needed.

To illustrate, an organization wants a Web server and intends to serve primarily static HTML documents with a few small (less than 25K) gif images. They have estimated an average traffic load of about 400 hits per hour and have calculated that a fractional T-1 will provide sufficient connection bandwidth. Given these conditions the table is entered with a one, a one is added for the files being served, a one is subtracted because they have less than half a T-1, and another one is added because they will be handling more than 150 connections per hour. This gives them a total of two and indicates that a high-end PC or Mac would provide the minimum hardware to satisfy the sites needs.

The authors of this heuristic are careful to point out that this is "... not a tried-and-true law, but a useful place to start." It is important to keep this in mind as the heuristic is evaluated.

Steps for Determining Hardware Level

1. Start with a score of 1 if you want a Web server.
2. Add 1 if you will be transferring a typical balance of images and HTML documents (average document size about 10K);

or add 2 if you will be transferring unusually large files (average above 25K), such as audio, graphics, or video.
3. Subtract 1 if your connection is greater than a half T-1.
4. Add 1 if you will be serving a substantial number of processor-intensive functions (such as data-base searches).
5. Subtract 1 if, on average, you will be handling fewer than 150 connections per hour (with peak usage at about 300 connections per hour);
or add 1 if you will be handling more than 500 connections per hour (with peak usage at about 1,000 connections per hour);
or add 2 if you will be handling more than 1,500 connections per hour (with peak usage at about 3,000 connections per hour);
or add 3 if you will be handling more than 4,000 connections per hour (with peak usage at about 8,000 connections per hour);
or add 4 if you will be handling more than 10,000 connections per hour (with peak usage at about 20,000 connections per hour).

Quickie Server levels, Based on Machines and Memory

<i>Level</i>	<i>Machine</i>	<i>Memory</i>
1	Mid-level PCs (486 up to about 50MHz) or mid-level Macs	4-8 MB
2	High-end PCs (high-end 486 or Pentium) or high-end Macs	8-16 MB
3	Mid-level Unix workstations (DS5000, SparcIPX, etc.)	8-24 MB
4	Higher-end Unix workstations (Sparc5, Pentium/486 UNIX box)	16-40 MB
5	Very powerful UNIX workstations (Sparc20, DEC Alpha, HP9000)	32-64 MB
6	Parallel processing workstations (multiprocessor machine, or multiple machines)	40-80 MB

Notes:

- 1) It is reasonable to assume that peak usage will be roughly double average usage.
- 2) If a system has more memory than is shown for its level, consider the system to be in a category between one-half and one level higher. For example, a DEC DS5000 (level-3) becomes a level-3.5 machine if it has 40 MB of memory.

Figure 2: Hardware Heuristic.

B. HEURISTIC BENCHMARK

1. SPEC

In order to determine if this heuristic was valid, some method was needed to compare dissimilar computer systems. Because of widespread computer industry acceptance and use, the benchmarks provided by the Standard Performance Evaluation Corporation (SPEC) were selected. SPEC was founded in 1988 and is a non-profit organization "...devoted to establishing, maintaining and endorsing a standard of relevant benchmarks that can be applied to the newest generation of high-performance computers." (Reilly, 1995)

SPEC has issued several benchmark suites to measure various aspects of computer performance. After contacting SPEC it was determined that the SPEC92 and SPEC95 benchmarks would be the most practical suites to use (Carlton, 1996). This decision was a compromise because the most appropriate suite to use would be the System-Level File Server (SFS) suite which runs the "LADDIS" benchmark for testing NFS (Network File System Protocol) file server performance. However, because there is little use of this suite (and therefore limited data) it was decided that the SPEC92 and SPEC95 benchmarks would provide a reasonable platform for comparing computers for Web sites. (A soon to be realized Web Server benchmark is similar to SFS and will specifically test Web server performance.)

The SPEC95 and SPEC92 benchmarks are designed to provide a comparable measure of performance for systems executing known computer-intensive workload. As the name implies, SPEC92 was released in 1992. SPEC95 was released in August of 1995 and is in the process of replacing SPEC92. Much of this discussion will address SPEC95, however, unless specified it also applies to SPEC92.

Like SPEC92, SPEC95 is a component-level as opposed to a system-level benchmark. Specifically, performance of the CPU, memory system including cache, and compiler code generation are tested. The benchmarks were not designed to measure graphics, networking, I/O, or operating systems (Dixit and Reilly, 1995). The benchmarks

are normally provided as un-compiled UNIX code which is compiled and run on the target systems. Other operating systems can also be employed, and results for Windows NT would be similar to those of UNIX (Carlton, 1996).

The SPEC95 suite consists of two sub-suites, CINT95 and CFP95. CINT95 (the “C” stands for “component-level” benchmark) contains eight individual benchmarks designed to measure CPU performance by performing *integer* computations. CFP95 contains ten individual benchmarks designed to measure CPU performance by performing *floating-point* computations. Because of a difference in units, it is not possible to directly compare results between CINT95 and CFP95 (SPEC, 1995).

In addition to the results from each of the individual benchmarks, “aggregate” values are also provided within each sub-suite. These “aggregate” configurations for CINT95 are listed below. For simplicity CFP95 will not be shown, however it also has an identical breakdown.

a. “Base” Measurement

SPECint_base: A “base” measurement, obtained with “conservative” (specified) optimization of the compiler/linker options. Unlike SPEC92, the base measurement is required for all reported results under SPEC95.

(1) Speed. *SPECint_base95*: The “geometric mean” of the eight benchmarks testing for speed “when compiled with *conservative optimization* for each benchmark”. It is expressed as a ratio of how long it takes the benchmarks to execute compared to a fixed “SPEC reference time”. A Sun Microsystems SPARCstation 10/40 (40MHz) is used as the reference machine for SPEC95. By definition the benchmark value of this system for both *SPECint_base95* and *SPECfp_base95* is “1”.

(2) Throughput (Rate). *SPECint_rate_base95*: The “geometric mean” of the eight benchmarks testing throughput “when compiled with *conservative optimization* for each benchmark”. This indicates the systems performance while executing

several copies of a particular benchmark in a given period. This test is best suited for multi-processor systems.

b. "Peak" Measurement

SPECint: A "Peak" measurement obtained with "aggressive" (tailored) optimization of the compiler/linker options. This is the benchmark most manufactures report because performance values are normally enhanced.

(1) Speed. *SPECint95*: This is the "geometric mean" of the eight benchmarks testing for speed "when compiled with *aggressive optimization* for each benchmark". It is expressed as a ratio of how long it takes the benchmarks to execute compared to a fixed "SPEC reference time". Because optimization is allowed, this value as measured on a Sun Microsystems SPARCstation 10/40 will be greater than the SPEC95 reference value of "1".

(2) Throughput (Rate). *SPECint_rate95*: The "geometric mean" of the eight benchmarks testing throughput "when compiled with *aggressive optimization* for each benchmark". This also measures throughput of the systems performance while executing several copies of a particular benchmark in a given period. As with all the rate measurements, this test is best suited for multi-processor systems.

Based on conversations with SPEC it was determined that since CFP95 is more suitable for 'numeric-scientific' (floating point intensive) applications, it would be more appropriate to use the CINT (integer intensive) values. CFP values are therefore not employed.

It was also determined that within the CINT suite only the *speed* computation benchmarks (*SPECint_base* and *SPECint*) should be used. The rationale behind this decision was that since the computers being benchmarked in levels one through five were single processor units it was most appropriate to use the *speed* values. As Level 6 (multi-

processor or multiple machines) is the highest level, *speed* values exceeding Level 5 will generally indicate computers that will satisfy Level 6 requirements. (Note: For those solutions requiring multi-processors (Level 6), it would be instructive to refer to *rate* values (*SPECint_rate_base* or *SPECint_rate*), in addition to the *speed* values, when comparing candidate computers.)

Finally, because CINT95 is new, many systems have not been tested with it. It was therefore necessary to include the more abundant CINT92 data along with the CINT95 data. Within both CINT95 and CINT92, the *SPECint_base* information represents un-optimized values and therefore would provide a better reference. However, because SPEC92 did not require the submission of 'base' (*SPECint_base92*) values, much of the CINT92 data is in the form of the "aggressively" optimized 'peak' (*SPECint92*) tests. Because of this, it was necessary to use both 'base' (*SPECint_base92* and *SPECint_base95*) and 'peak' (*SPECint92* and *SPECint95*) values.

2. Heuristic

Using the fairly generic "machines" listed in Figure 2, and based on consultation with the authors of *Build A Web Site* and objective judgment, more specific models were identified and used to assign benchmark values to each level of the heuristic. Appendix C, HEURISTIC BENCHMARK, contains a list of these computers, their SPEC benchmark values and average SPEC values for each level. This list is representative of computers within each level. It is not inclusive of all possible computers, listing only those with reported SPEC values.

Table 2 is a summary of Appendix C. *SPECint95* values have not been included in Table 2 because in most cases they are identical to the *SPECint_base95* values, and because *SPECint_base95* represents a better reference statistic.

Four good sources were found for SPEC data. The first is SPEC itself. Their Web site, at <http://www.specbench.org>, lists CINT95 data that has been submitted to them. Because SPEC92 data is more difficult to format, they do not currently have it posted.

<i>Level</i>	<i>Machine</i>	<i>SPECint _base95</i>	<i>SPECint _base92</i>	<i>SPECint92</i>
1 SPEC92 Avg.: 30	486 (33-50 MHz)			18.2-30.1
Averages	'Mac' (50 MHz)			<u>40-41.7</u> 30
2 SPEC92 Avg.: 60	486DX2 (66 MHz)		36.7	32.2-39.6
Averages	Pentium (60-66 MHz)		60.4-74.0	50.0-78.0
3 SPEC92 Avg.: 30	'Mac' (66 MHz)		<u>50.7-62.1</u>	<u>40.0-76.0</u>
Averages	DEC Station 5000 (40-50 MHz)		58	61
4 SPEC92 Avg.: 85 SPEC95 Avg.: 2.18	Sun SPARC IPX (40 MHz)			27.3-43.2
Averages	Sun SPARC 5 (70-110 MHz)	1.37	49.8-68.7	<u>21.8</u> 30
5 SPEC92 Avg.: 129 SPEC95 Avg.: 3.80	Pentium (70-90 MHz)	<u>2.31-2.74</u>	<u>79.0-104.3</u>	<u>57.0-78.6</u> <u>83.8-110.1</u>
Averages	Sun SPARK 20 (75-125 MHz)	2.14	83	86
6	DEC Alpha (100-275 MHz)	2.82	94.0-122.4	104.5-131.2
	HP 9000 (80-125 MHz)	1.48-6.43	68.6-257.1	74.6-289.0
	(Also see <i>SPECint_rate_base</i> or <i>SPECint_rate</i> data)	<u>2.89-4.04</u>	<u>74.5-138.5</u>	<u>82.0-149.4</u>
		3.83	125	131
		> ~6.5	> ~260	> ~300

Table 2: SPEC Benchmark Reference for Hardware Heuristic

The second source is the "PDS: The Performance Database Server", provided by SPEC and the University of Tennessee at URL <http://performance.netlib.org/performance/html/spec.html>. This site provides a listing of SPEC92 statistics. It is also possible to conduct various database searches at the site.

Both the SPEC site and the SPEC/University of Tennessee site provide additional data for each entry. It is possible to obtain all measurement values for the individual benchmarks tested, not just the aggregate value's.

The most comprehensive listing of aggregate benchmark statistics is John Dimarco's University of Toronto site at the <ftp://ftp.cdf.toronto.edu/pub/spectable>. This site provides both SPEC92 and SPEC95 statistic. The contents of this site have been made available in Appendix B as a reference for obtaining computer SPEC values for use with

the heuristic. The drawback to this site is that additional information is not available for the entry's. However, the comprehensive nature of the listings makes this site very valuable.

Finally, the site at Berkeley, URL <http://infopad.eecs.berkeley.edu/CIC/summary/local>, is mentioned because it was the only site where some of the earlier systems could be found. However, the focus of this site is the CPU and system data is limited.

When comparing the values in Table 2 (or Appendix C) to other values within the table, or to SPEC values for other computers (such as from Appendix B), several points should be kept in mind. First, as *SPECint92* values represent "optimized" test results, these values will be greater than *SPECint_base92* values for the same system. For example a DEC AXPpci 33 tested under *SPECint_base92* was rated 69.4 while under *SPECint92* it was rated 76.0 (DiMarco, 1996).

Second, there is no direct mathematical method available to convert between SPEC92 and SPEC95 values (Dixit and Reilly, 1995). However, by obtaining the SPEC92 values for the SPEC95 reference machine (a SPARC 10/40) the two different suites can be roughly equated. Recall that the *SPECint_base95* value for a SPARC 10 is "1". The *SPECint92* value for SPARC 10/40 is 50.2 (DiMarco, 1996). Therefore any computer possessing a *SPECint92* value of less than the low fifties (upper forties for *SPECint_base92*) can be considered to perform worse than a SPARC 10/40 and have a *SPECint_base95* equivalent value of less than "1".

The most important point to keep in mind is that the values should be used as a general relative indication of a computer's performance and not as a precise indication of performance or absolute ranking. Many factors come into play during these tests and the fact that comparisons are also being made across different tests further dilutes the precision with which any exact comparisons could be made.

For example, although the heuristic in Figure 2 lists specific values for RAM at each level, most of the SPEC tests were conducted with 64MB of RAM. Also, the results of the SPEC tests are dependent on the amount of cache a system is tested with, which is not accounted for in the heuristic (or listed in Appendix C/Table 2). Furthermore, because

SPEC tests for Mac's are not available the 'Mac' values listed in level one and two are based on IBM and Motorola machines which use the same CPU as Mac'.

All of the above factors contribute to make Appendix C/Table 2, and therefore the heuristic, very 'coarse'. Based on careful observation of the SPEC statistics and given the factors discussed, it's reasonable to consider SPEC92 values within approximately 10 'points' (for example 20 and 28) of each other to be roughly equivalent. Within SPEC95, values of several tenths (for example 1.5 and 1.8) can be considered to be roughly equivalent.

As for the amount of RAM required, as pointed out in Chapter IV, the RAM amounts recommended in Figure 2 should be considered a minimum level, and guidance provided by the manufactures of the hardware, server software and operating systems should also be heeded. The fact that the SPEC tests were conducted with considerably more RAM than is specified by the heuristic will not affect the *relative* comparison of various candidate computers because, as mentioned, all systems (there were very few exceptions) were tested with a uniform 64MB.

Another significant point to consider is that the benchmark values for Level 1 equals that of Level 3, and those of Level 2 approximate Level 4 values. The reason for the duplication is that Levels 1 and 2 represent non-UNIX based operating systems. The authors of *Build A Web Site* considered Macs and Windows based computers "... good for handling light loads, but not recommended for heavy loads." (Load in this case is defined as the number of processes the computer is performing at one time.) The reasons for this position are the same as those pointed out in Chapter IV, UNIX is viewed as tried and true (i.e. more stable and secure) than newer operating systems, such as Macintosh and Windows.

The validity of this argument will not be debated here. It must be pointed out however, that only UNIX based SPEC statistics were available to benchmark the heuristic. This is one more factor contributing to the 'coarse' granularity of Appendix C/Table 2. It is obvious that the *equipment* represented in these two levels would be capable of handling the server loads of the corresponding higher levels. Only the fact that these levels

represent non-UNIX based operating systems causes them to be ranked as the two lowest levels.

This leads to another point, which is the rapid increase of hardware performance. As predicted by the so called "Moore's Law", named for Intel co-founder Gordon Moore, microprocessor performance is doubling every 18 months (Cohen, 1996). This progression of CPU performance became apparent when researching the benchmarks statistic. At the time the book *Build A Web Site* was written 90 and 100 MHz systems were very powerful machines. Computer systems of this performance, especially the Pentium machines, are now the *minimum* that most organizations would consider. This means that today most organizations or individuals buy *entry* level equipment that automatically places them at Level 4 in the heuristic (ignoring operating systems issues).

If this advance continues, and there is not a corresponding increase in the requirement for more performance (such as new or more CPU intensive scripting or database searches) hardware could cease to be a limiting factor for most sites. In the interim a heuristic that allows hardware selection will continue to be useful, not only for new equipment purchase but for legacy equipment employment.

C. HEURISTIC VALIDATION

To validate the heuristic a number of Web sites were contacted to find out what equipment was being used, what their connection was, what sort of documents they were serving and how heavy their traffic flow was. This information was then used to calculate a *recommended* hardware level using the heuristic. SPEC benchmark values were then found for the actual hardware being used. These values were compared to Appendix C/Table 2 to determine which level of hardware the site was *actually* employing. The two figures were then compared to determine how accurate the calculated recommendation reflected actual hardware.

A total of 29 Web sites were contacted. Of those, 19 sites provided sufficient information. Commercial sites were very difficult to obtain information from. There were

two reasons for this. The first is that many did not maintain their own equipment, employing a provider instead. The second reason was that they were reluctant to give out information. Educational institutions were also difficult to obtain information from because it was very hard to identify individuals who knew the pertinent facts. These problems were not experienced with government organizations, so most of the surveys obtained came from these sites. The surveys are presented in Appendix D.

The results of the surveys reveal that at six of the 19 sites the actual hardware level matched the calculated levels. Of the remaining 13 sites all used hardware at a level which exceeded the calculated level. In many of these cases where the level of actual hardware used *grossly* exceeded the calculated level the site administrators conceded that the equipment was more powerful than the site currently required.

Additionally, based on the level of hardware *actually in use* at the site (*not calculated*), the results show that five of the 19 sites were using RAM amounts that matched the RAM recommended for that hardware level. Of the remaining 14 sites, 12 used RAM which exceeded the amount recommended and two used less RAM than was recommended for the level of hardware actually in use.

To provide a statistical analysis of the survey results, each result can be viewed as a Bernoulli variable with probability, P , that the result will equal or exceed that calculated in the heuristic. The total number of successful trials (S) has a Binomial distribution with parameters N and P . N is the number of independent Bernoulli trials. P is a measure of the accuracy of the heuristic and if, for example, $P \geq .89$ then the heuristic will correctly predict the equipment required at a site 85% of the time. (Woods, 1996)

In this case, where the number of trials ($N=19$) equals the number of successes ($S=19$), using a 90% lower confidence limit for P yields a probability of at least 0.89 that the heuristic will reliably calculate the level of computer needed for a site. If an 80% lower confidence limit is used, the probability increases to at least 0.92 that the heuristic will reliably calculate the level of computer needed for a site. (Lloyd and Lipow, 1984)

When a similar analysis is conducted on the RAM results, where the total number of trials is again 19 ($N=19$) and the number of successes is 17 ($S=17$), a 90% lower

confidence limit for P yields a probability of at least 0.74 that the heuristic will reliably predict the RAM needed for a particular hardware level. If an 80% lower confidence limit is used, the probability is at least 0.79 that the heuristic will reliably predict the RAM needed for a particular hardware level.

Although the validation sample is small, the results demonstrate that the heuristic is valid. However, because most sites were using hardware that exceeded the calculated level, it seems reasonable to use the heuristic as an indication of the *minimum* level of hardware to be employed at a Web site. This reasoning should also be applied to the amount of RAM recommendation in the heuristic.

D. CONCLUSION

This chapter has presented a heuristic adopted for calculating hardware infrastructure for a Web site. The methods of benchmarking and validating the heuristic were also covered.

Although, the heuristic was demonstrated to be valid, as noted previously by the authors of the heuristic, it is not "...a tried-and-true law but a useful place to start." Therefore, it is recommended that the heuristic be viewed as a rough indication of plateaus of computing power needed for Web sites and should be used to determining the *minimum* levels of hardware and RAM necessary.

Another important point brought out in this chapter is that of the rapid increase of hardware performance as 'predicted' by "Moore's Law". Because microprocessor performance is doubling every 18 months machines that were considered very powerful a year ago are entry level platforms today. The implications of this are that if this advance continues hardware could cease to be a limiting factor for most sites.

In the interim however, the heuristic will be useful, not only for new equipment procurement but for legacy equipment employment.

VI. SYSTEM RELIABILITY

- How to provide system and stored data integrity and redundancy?

Finally, we come to the last question with regard to determining infrastructure requirements - system reliability. Depending on how 'mission critical' a Web site is to an organization, various measures can be taken to insure that data is protected and to provide degrees of 'robustness' for the site. The basic question that must be answered is does it matter to the organization if the site goes 'off-line' for periods of time because of equipment problems. In other words - how much fault tolerance does the site require.

A generally accepted rule is that a single (stand-alone) UNIX system provides 99.5% uptime. Adding a RAID (Redundant Array of Independent Disks) subsystem can increase this to 99.9%, and "clustering" (two or more coupled systems with a shared disk subsystem) can provide 99.99% reliability. To put this in perspective, for a 24 hour a day, seven day a week operation (7/24) 99.5% reliability yields over 43 *hours* of un-planned downtime per year, while 99.99% equals 52 *minutes* of downtime each year. (Simpson, 1995)

Due to the global nature of the Internet it is reasonable to expect traffic at all hours of the day. Therefore, any downtime can potentially cost organizations millions of dollars a year. If it is considered critical (due to cost or other 'mission' factors) to an organization to maintain a 7/24 site, then it will need to be designed with a high degree of fault tolerance.

As with any computer enterprise, the most basic precaution to be taken is to implement and religiously adhere to a data backup scheme - this does not, however, introduce any additional fault tolerance into a system. Two primary approaches that facilitate data integrity and a site's fault tolerance will be introduced in this chapter.

A. DISK STORAGE SUBSYSTEM

Studies by Intel have determined that hard disks account for 50% of all system component failures and disk controllers for 4% (system power supplies account for another 28%) (Milne,1995). RAID (Redundant Array of Independent Disks) provides fault tolerance for the failure prone disk storage subsystem.

The basic principle of RAID is to combine two or more hard disks into an 'array' with data copied or distributed across all the disks. Should a hard disk failure occur, data can be recovered or reconstructed from other disks in the array. Most RAID systems go much further than this in that they also provide redundant controller cards, cooling fans, power supplies, cables, etc., thus minimizing single point failure within the storage subsystem.

The approach is similar to that provided by some operating systems such as Novell Netware which provides disk mirroring or duplexing. Mirroring uses back-up hard disk(s) which have the same data written to them (mirrored) as that being written to the primary hard disk(s). If the primary disk(s) fails, data is recovered from the back-up disk(s). Duplexing takes mirroring a step further by providing redundant controller cards, cables, etc., thus it also minimizing single point failure within the storage subsystem. However, with both disk mirroring and duplexing, the system (Web site) must be brought off-line to effect the recovery. (Lin, 1996)

A distinct advantage of RAID is the capability to 'hot swap'. This is the ability to recover from a hard disk or other component failure by replacing the failed component without bringing the system off-line. Some RAID systems also include an 'on-line' spare feature in which a normally idle spare component automatically 'kicks in' when another component fails. Although system performance may suffer after a component failure, these RAID features greatly enhance a site's ability to maintain data integrity and provide a reasonable level of fault tolerance.

RAID can be 'hardware' or 'software' controlled. 'Hardware' controlled is a misnomer because software actually controls the RAID subsystems. However, it is

running on the RAID hardware and opposed to the host system (Levin, 1996). Although it is more expensive, 'hardware' controlled is preferable because it is usually more reliable and faster than a 'software' based system. Unlike 'software' controlled, it does not affect host system CPU performance. Software controlled versions can affect CPU loading during a data rebuild which can lead to problems if the CPU is already loaded. Additionally, hardware based has the advantage of providing more flexibility as to which operation systems are supported (Milne, 1995).

RAID is categorized into several levels each using different methods for data recovery.

1. RAID 0

RAID 0 uses 'data striping' to distribute *blocks* of data evenly across multiple disks (minimum of two) making a single volume (Lin, 1996). This level is best suited for transferring large blocks of data such as large data bases and sequential files, and where read and write performance is important. RAID 0 is very fast but provides no fault tolerance because if a disk fails there is no way to rebuild the data (Milne, 1995).

2. RAID 1

This level uses mirroring to copy blocks of data to spare disks. RAID 1 provides fault tolerance via the backup disks - should the primary disk(s) fail the back-up(s) comes on line. This level is also fast and is suited for applications requiring the transfer of large blocks of data. (Milne, 1995). However, this level is considered to be the least cost efficient because as two complete sets of data are being maintained, only half the total disk storage capacity is available for general use (Levin, 1996).

3. RAID 2

Level 2 functions in a similar manner as RAID 3 below, with the exception that data is 'striped' at the *bite* level as opposed to *bytes* as is done in RAID 3 (or *blocks* as in

RAID 0). RAID 2 is not widely used because it is slow and, as with all levels that dedicate a disk to parity, expensive.

4. RAID 3

Level 3 uses an approach which is a combination of levels 0 and 1. Data is 'stripped' at the *byte* level and distributed across two (or more) drives as in RAID 0. Another drive is used to store a parity bit from each byte of information. If a disk fails lost data can be rebuilt from any two remaining drives thereby providing fault tolerance. Similar to the situation in RAID 1, the disk that is used to store the parity bite is not available for general use by the system. Although not as bad as RAID 1, this approach is also less cost efficient. This level is also best suited for applications requiring the transfer of large blocks of data. (Levin, 1996)

5. RAID 4

Level 4 functions in a similar manner as RAID 2 and RAID 3, with the exception that data is 'stripped' at the *block* level as opposed to *bites* (RAID 2) or *bytes* (RAID 3). This level is best suited for file and print servers with small files and where write performance is not critical (Milne, 1995).

6. RAID 5

Like RAID 4, RAID 5 'strippes' data *blocks* across multiple disks. However, RAID 5 uses all disks (minimum of three) to store both data and parity bits. RAID 5 has excellent read but poor write performance. Therefore (as with RAID 4) this level is best suited to applications for file and print servers with small files, and where write performance is not critical. In the event of a disk failure both read and write performance will be severely affected, because the systems must read from all surviving drives to re-construct the missing data (Milne, 1995).

RAID 5 is increasingly becoming the most popular RAID implementation because it overcomes the parity disk shortcomings of other RAID levels. However, different situations require different solutions and RAID 1 or 3 may be more appropriate for applications requiring the writing of large amounts of data. In practice RAID level 2 and 4 are seldom used because other levels provide better performance and/or functionality. When compared with other RAID levels, RAID 5 becomes attractive when storage capacity approaches 4-5 gigabytes or above, based on cost verses storage capacity (Milne,1995).

Newer RAID systems allow multiple RAID levels. For instance Hewlett-Packard's new AutoRAID system automatically and transparently configures for RAID 1 or RAID 5 as needed (Carr,1995). Although this approach is relatively new, it should become quite common.

B. SYSTEM REDUNDANCY

RAID works well for providing data integrity and disk storage subsystem redundancy, but to provide overall 'system' redundancy another approach is required. Two approaches for providing system redundancy are 'clustering' and 'superservers'.

1. Clustering

Clustering can be defined as "...two or more loosely coupled systems with a shared-disk subsystem and software that handles failure in the case of a node failure." As mentioned previously, clustering can be employed to supply the high degree of fault tolerance required for a 7/24 site by providing 99.99% system reliability. (Simpson, 1995)

Ideally, clustering provides several desirable features - ease of system management, hot swapping of nodes, routine servicing for individual nodes without any interruption in site availability, a single unified view of the file system, and scalability.

The primary advantage of clustering is scalability which allows the addition of multiple nodes. There are two advantage to this. The first is that as a site grows and more

hardware resources are required, additional nodes can simply be added to the site. The second reason is that multiple nodes introduces redundancy into the system. The number of possible nodes depends on the implementation, with two to eight being common for commercial offerings.

Scalability generally does not result in a linear performance increase with each additional node. Depending on factors such as the efficiency of the operation system and whether an application has been suitably optimized, actual performance gains can be expected to be 1.6x to 1.8x (80-90%) when going from one to two node, 2.5x to 3x when going from one to four nodes, and 5x when going from one to eight node. (Simpson, 1995)

A potential drawback to clustering is that a high degree of network message handling (i.e. network bandwidth) can be required for certain shared file system implementations. (Simpson, 1995)

Three approaches to clustering will be discussed.

a. Hyperlinked Computers

Technically this approach (which has been named 'Hyperlinked Computers' for lack of a formal title) may not pass the definition of 'clustering', because it does not use a shared file system and lacks other clustering features. However, it is mentioned here because it does offer an inexpensive solution to providing redundancy and 'computing power' to a site.

This approach involves nothing more than placing different functional areas of a site on separate computers and interlinking them via hypertext. With a little extra effort the contents of each server can be duplicated on the other server(s) (mirroring). In the event one of the platforms fail, reconfiguring the links on the good machine(s) will bring the entire site back up. (Powell, 1994)

An advantage to this is that it is hardware and operating system independent - any spare machine capable of running a Web server can be used in the configuration. It is also inexpensive in that no clustering software need be purchased.

However, as manual configuration and monitoring are required this approach would not be very practical for large or 'mission' critical sites.

b. Commercial Solutions

The most expensive, but perhaps most transparent solutions (for administrators as well as users) are commercial packages. These range in price from \$70,000-95,000 for a complete two node system, to \$2,500-20,000, per node, to add clustering software to existing equipment. These systems offer transparent hardware and software failover, and although some performance degradation may be experienced, the best provide failover times of 15 to 30 seconds. (Simpson, 1995)

As mentioned, a potential disadvantage of clustering is the large amount of network bandwidth that can be required for implementations that use a shared file system. Commercial vendors are developing *proprietary* 'connection' solutions to reduce this traffic overhead. Additionally, because the clustering packages are designed by computer vendors they can be fairly restrictive as to which equipment, operating systems and networks are supported. As to be expected UNIX is widely supported, however, Windows NT systems are not presently available.

Currently, the DEC VMScluster system is the standard by which other commercial systems are judged. This is due to the high level of functionality of its clustered file system and system management software.

c. NCSA scalable Server Approach

Technically, this is also a commercial approach because the key component, the Andrew File System (AFS), is a commercial product. However, because NCSA (National Center for Supercomputing Applications) has used AFS at their site to solve a scalability problem, it will be discussed within that context.

In response to rapid growth in the traffic on their Web servers, NCSA at the University of Illinois researched and implemented a "Scalable HTTP Server" solution to their problem. (Katz, et al., 1994)

An initial solution to the problem was to migrate the site to a faster more capable machine, however, this computer was also overwhelmed (a common occurrence!). They subsequently determined that some form of distributed multi-server configuration would be required. Two approaches to this problem were considered. The first was to divide the document tree among several computers with each responding to a unique host name and each serving a portion of the documents (thus distributing different site functions or process on separate machines). This was considered prohibitively complex for both Web site administration and user access.

The second approach involved sharing the document set among a group ('cluster') of servers each answering to the same host name. This solution was successfully adopted. The key to this approach is the use of a "load distributor" that maps multiple machine IP addresses to a single URL and the Andrew File System (AFS).

The "load distributor" That NCSA used is a version of the Berkeley Internet Name Domain (BIND) which allows a "round-robin" mapping of IP addresses to the site's URL. This arrangement is 'stateless' in that no knowledge of a particular server's loading is maintained. Instead, a time limit is set (currently 15 minutes) after which a new IP address is mapped to the URL (McGrath, 1996). (Problems can, and do, develop with this 'stateless' approach if a client continues to use an old IP address. Some applications employ a 'state-full' approach, however due to problems involving uneven loading resulting from time-lag, this approach was not used.)

AFS, the heart of the NCSA scalable server solution, is based on a distributed file system originally developed at the Information Technology Center at Carnegie-Mellon University. AFS is currently marketed by Transarc Corporation, Pittsburgh, Pa. from whom NCSA obtained their license.

A key feature of AFS that distinguishes it from other distributed file systems, such as NFS, is that it uses a local caching scheme that allows repeatedly accessed documents to be stored at the individual Web server(s) (nodes). This minimizes the common drawback of shared file system clusters - the large amount of network bandwidth that is generated by constant and repeated hits to the file server. Another

distinct advantage of this approach is that it allows much faster document retrieval for users. It is important to note that this approach is ideally suited for a site that experiences repeated requests for the same document sets. If a site does not experience this type of traffic this approach may not be appropriate.

Very briefly, AFS functions by maintaining a complete set of data on a file server which has read and write authority. A complete set of data is also maintained on the individual Web servers (nodes of the cluster) in read only or 'replicated' disks. The data across the system is "consistent and synchronized" (Katz, et al., 1994). Periodically (every hour) the 'master' data is written to the 'replicated' drives to bring that information up to date. When an Internet client connects to the site they are connected to the web server which is currently being mapped to the URL. After retrieving the information from the read only 'replicated' disks (and passing it on to the Internet client) the Web server stashes it in its local cache. Future requests for the documents are served from cache. The information in cache is compared to the information on the read-only disc and flushed and replaced as necessary to maintain currency.

Another key advantage to AFS is the scalability it allows. Unlike many commercial cluster systems, AFS is platform independent - it allows many hardware platforms to be used and intermixed in the system (generally, as long as the computer will accommodate an AFS client, which means a UNIX machine). Because the individual servers do not know about each other, nodes (client servers) can be 'hung' or removed from the cluster without affecting the site. AFS also allows geographically dispersed Web sites sharing the same file content (Houston, 199).

The recommended limit to the numbers of 'nodes' is a ratio of one file server to every 50 client Web servers (50:1). However an "architectural goal" was a ratio of 200:1, which has been successfully achieved at some sites. "AFS cells can range from the small (1 server/client) to the massive (with tens of servers and thousands of clients)." (Transarc Corporation, 1996)

Other prime features include the use of Kerberos for user authentication, and Access Control Lists (ACL) for file and directory access. These features provide a very flexible and secure basis for configuring both local and remote user access.

As with other clustering systems, the NCSA approach also eliminates single point of failure for system components as well as disk subsystem.

AFS demands more attention than can be delivered here. It would make a good area for future study.

2. Super Servers

Finally, another potential solution to the problem of providing fault tolerance is another commercial offering - 'Commodity Superservers'. These units, which can cost as little as several *thousand* to as much as several *hundred thousand* dollars, provide fault tolerance as well as improved performance and potential growth by incorporating multiple components into their design. (Milne, 1995)

By using multiple processors, superservers achieve much of the same performance improvements as clustered systems do, as well as providing fault tolerance via multiple processors. Similarly, by incorporating RAID technology into their disk subsystems these servers offer that level of fault tolerance as well. Additionally, many include redundant components such as power supplies, cooling fans, cabling, etc.

Another area that superservers support is growth for an expanding site. By allowing large RAM upgrades as well as support for additional processors they can be scaled up to meet increasing demand. As with clustered systems, adding processors does not result in a linear performance increase with each additional CPU.

C. CONCLUSION

There are several ways to provide fault tolerance for Web sites. RAID systems can be added to supply redundancy to disk subsystems, systems can be 'clustered' together, or superservers can be purchased which incorporate many individual advantages into one unit.

Because of its expandability, hardware independence, and security as well as its unique caching scheme, the most interesting and flexible approach covered seems to be that of employing AFS.

VII. CONCLUSION

A. THESIS SUMMERY

This thesis explores the issues of defining infrastructure requirements for WWW sites and provides guidance for the selection of that infrastructure based on the requirements identified.

Due to the rapid growth of the Internet in general and the World Wide Web in particular there is a need for guidance to organizations and individuals desiring to establish new Web sites. The requirements for a site's infrastructure varies depending on the function of that site, and it is not uncommon for important nuances to be overlooked or the complexity of the task to be underestimated. The result can be an investment in Web site infrastructure that is insufficient or inappropriate to meet the demands of the site.

A combination of literature review and site surveys of existing WWW sites was used to obtain information. This information was used to identify and define the issues, and to develop the framework for evaluating a site's infrastructure requirements. Additionally, a rule based hardware heuristic was adopted from the literature and subsequently validated.

Taken together, the material in this thesis provides the information necessary to identify and select the infrastructure needed for a site.

B. CONTRIBUTIONS

One of the contributions this thesis has made is to highlight the lack of literature available providing guidance on actually *defining* needs and then *selecting* infrastructure based on those needs. Most material gave varying, but generally acceptable, descriptions of *what was available* but not *how to define or select what was needed*.

Another contribution is the revelation that most organizations do not conduct any initial requirements analysis to determine a site's infrastructure needs. The reasons range from oversight to indifference, however, the potential penalty for not conducting proper assessment of requirements is the same as for any venture, a substandard product and poorly leveraged investment.

The most significant contribution this thesis has made is to provide the material needed to correct the short-coming of most of the literature reviewed. To this end the information necessary to identify and select the infrastructure needed for a WWW site is provided.

Finally, a key contribution is the revelation that hardware could cease to be a limiting factor for most sites. The fact that microprocessor performance is doubling approximately every 18 months (as predicted by "Moore's Law") is fairly well known. However, the effect that has on WWW sites may not be so obvious. It became apparent during the validation of the heuristic that the 'entry level' for computers has significantly increased in the two years since the heuristic was written. If this trend continues without a corresponding increase in computational requirements, 'entry level' computers will soon be able to handle all but the most demanding sites.

C. AREAS OF ADDITIONAL RESEARCH

1. Operating Systems

The operating system used on a server can make a substantial difference in performance. Apparently the "efficiency" with which the TCP/IP stack is handled can be instrumental in how stable and robust the server is. Although this seems to be common wisdom among Web site administrators, no detailed information about the issue was available. This would be an excellent topic for further research. It could add significant insight into server performance.

2. AFS

Many administrators are struggling with questions of scalability and security, as well as how to configure or distribute functional areas of a site. Because of its potential to greatly facilitate a site's functionality, AFS would be very useful to investigate. Although it is an established product, it is unknown among the majority of Web site administrators surveyed in this thesis. Its obscurity suggests that it would make a good candidate for further study and site experimentation.

D. CONCLUSIONS AND RECOMMENDATIONS

As mentioned previously, one of biggest lessons learned in this thesis is that initial requirement analysis to determine a site's infrastructure needs is not being conducted. Contributing to the problem is a lack of literature providing guidance on actually defining needs and then selecting infrastructure based on those needs. In part this lack of literature is probably due to the relative newness of the subject. In any event the potential penalty for not conducting proper assessment of requirements is the same as for any venture, a substandard product and poorly leveraged investment.

It is absolutely necessary to conduct a deliberate study of how much traffic the site will receive as well as defining the purpose of the site. These two issues drive all the major infrastructure requirements.

With regard to traffic, it is very difficult to accurately estimate the connection hit rate. Several techniques were outlined which will provide reasonable, educated estimates of how heavy the load will be. However, the only way to *accurately* determine the traffic for a site is to actually measure it. Two methods for this are available - using existing equipment or rent server space from a provider.

If approached from the correct perspective, employing existing surplus equipment can be used as an effective requirements analysis tool by actually measuring the traffic at a site. Since the equipment is a 'sunk cost' this approach can be a cost effective method for determining a sites true infrastructure needs. However, the danger in this is that the

surplus equipment may be viewed as a permanent solution instead of an interim arrangement resulting in a lack of financial commitment toward upgrading to the real requirements of the site.

For those sites which are being started from 'scratch' it is *strongly recommended* that space be rented on a provider's WWW server for at least the first six months of operation. During this period the actual load on the site can be determined and infrastructure requirements accurately determined.

Determining the purpose of the site is a much more direct issue. It must however, be given deliberate consideration. It is essential that the intended audience be identified as well as the content that will be provided.

Once the traffic and purpose of the site are known it is relatively straightforward to identify the bandwidth, software and hardware which will satisfy those requirements. Chapters IV and V provide the information required to do this.

Finally, with regard to the heuristic in Chapter V, it was demonstrated to be valid. However, it is recommended that it be viewed as a rough indication of the levels of computing power needed for a Web site and should be used to determining the *minimum* levels of hardware and RAM necessary.

APPENDIX A. WEB SERVERS SURVEY

This information was obtained from the WWW site, Web Servers Survey, by Paul Hoffman of Proper Publishing. It can be viewed at: <http://www.proper.com/www/servers-survey.html>. This document is provided as a ready reference to the current market use of server software packages.

Web Servers Survey

Version 2.0, January 1996

by Paul E. Hoffman, Proper Publishing

Many people ask "Which Web server software is the most popular?" The best way to find out is to directly survey the thousands of Web sites using HTTP commands. This document is the result of an extensive scientific survey of this type. This is the second survey I've conducted; the first was done in mid-September 1995, and the relative differences in the results are described below.

If you are more concerned with the features of particular Web servers, not their popularity, please take a look at the Web Servers Comparison, which I also maintain. That chart also has pointers to where you can get more information on over 40 Web server software packages.

Executive Summary

By far, the most popular Web server software remains the free Unix-based servers from NCSA and CERN, as well as Apache, a spin-off from the NCSA server. The next most popular category is commercial software: Netscape's Unix-based software, Mac-based WebSTAR, and PC-based WebSite. There were over two dozen other server packages found, each of which had only a tiny percentage of the server market.

The differences from this survey and the one done four months earlier are also important. Netscape has increased its share from 8% to 13% in just four months. Many people have switched from the NSCA and CERN servers to Apache, and the market share of the combination of NCSA and Apache remain around 60%. WebSite has greatly increased its share of the market, and other Windows-based Web servers are becoming more popular as well.

How The Survey Was Taken

In order to get reasonable results, I polled a random sample from a large database of known Web addresses. Other surveys in the past have used less scientific methods, such as relying on server maintainers to respond to a questionnaire, or by only choosing domain names that start with the string "www".

The database was kindly provided by Yahoo, who has one of the largest and best-cataloged index of Web sites anywhere in the world. Yahoo provided a list of names for over 45,000 unique hosts on the Web, taken from the beginning of January 1996.

Note that these are unique domain names of hosts, not Web pages; the Yahoo database is much, much larger than this list because many hosts have multiple Web pages that appear in the index, and the Yahoo database also has tens of thousands of other resources, such as Gopher sites, FTP archives, Usenet news groups, and Z39.50 (WAIS) databases.

From this large dataset, I selected a random subset of 2000 sites to poll. A Perl script sent an HTTP "HEAD" request to each domain name in the subset, and stored the responses to the requests. All information returned was kept, and all errors were logged.

The polling program encountered the typical errors that Web users do: connection failed, bad host name, and host too busy. To get as complete results as possible, I waited about 36 hours and queried again all the hosts for which errors were received during the first run.

During the first run, 1734 of the 2000 Web servers responded; during the second run on the remaining 266, an additional 58 Web servers responded. In all, 1792 servers responded, approximately 90% of the 2000 polled.

Another survey, run by Netcraft in the UK, came up with similar results as this survey. Their survey is run on a different (and much larger) data set that was acquired using a robot. They also have some fun tools, like the ability to find what server software a given site is running in real time.

Most Popular Servers

The following lists the most popular server software packages that have 1% or more of the market. The table shows the number of queries out of the 1792 each returned and the percentage of the total market. Note that these are not raw data: different versions of each package have been combined into a single number.

Server	Count	Pct
NCSA	732	41%
Apache	305	17%
		70

Netscape	237	13%
CERN	198	11%
WebSTAR / MacHTTP	101	6%
WebSite	73	4%
BESTWWWD (best.com)	37	2%
OSU (Region 6)	14	1%
Purveyor	12	1%

The full set of raw data used to generate this summary table was:

732	NCSA
295	Apache
198	CERN
111	Netscape-Commerce
101	Netscape-Communications
74	WebSTAR
73	WebSite
37	BESTWWWD
17	MacHTTP
14	OSU
12	Netsite-communication
12	Purveyor
9	Netsite-Communications
8	VApache
7	WinHttpd
7	MacHTTP 2.0
6	GN
5	Spinner
5	I-Site Web Server v1.1 w
5	HTTP-Srv-Beta2
5	Alibaba
4	Netsite-Commerce
4	plexus
3	GoServe
3	MacHTTP 2.0.1
3	WN
3	Commerce-Builder
2	WebServer Version 1.0
2	NFIC MultiHost CERN
2	NaviServer
2	IBM Internet Connection Server
2	Worldgroup
2	Apache-SSL
2	FTPd

2	WindsorWeb
1	Prezemyslaw-serv 038H
1	NEIC Superserver 2.19
1	Open Market WebServer
1	NDC Port Redirector
1	IBM Internet Connection Secure Server
1	HTTPS
1	SySNET Route 1.0
1	Hyper-G WWWMaster
1	Internet-Office-Web-Server
1	Marquette Web Server
1	PSIWeb
1	Open-Market-Secure-WebServer
1	Mosaic-Netsite
1	Webshare
1	FolkWeb
1	CMSHTTPD
1	SpiderWEB - WWW Server (MSWindows)
1	Amdahl
1	Branch_Internet_Image
1	INOS_NT
1	Delta's Very pache
1	ECN psudo www redirector

Note that some of the servers in the raw list have names that contain spaces. This is not allowed by the HTTP specification, and most current versions of servers only display names with no spaces.

Differences from the First Survey

The NCSA and CERN servers both lost significant market share in the four months between surveys, but the Apache server, a free Unix-based server based on the NCSA code, made a large increase. The total market for these three servers went from 78% to 69% in the four months, indicating that people are using Unix-based freeware less, but that it is still the vast majority of the Web servers market.

Netscape made an impressive climb from 8% to 13% of the market, and WebSite made an even more impressive climb from 1% to 4%. Both these servers are commercial, and it is likely that the trend toward commercial servers will increase.

Server	1/96	9/95
NCSA	41%	54%
Apache	17%	7%
Netscape	13%	8%
CERN	11%	17%
WebSTAR / MacHTTP	6%	5%
WebSite	4%	1%
BESTWWWD (best.com)	2%	<1%
OSU (Region 6)	1%	<1%
Purveyor	1%	<1%

It is interesting to note that BESTWWWD made it to the list in both rounds. This is a multi-homed Web server created by BEST Internet Communications and used in-house for their Web leasing. Making it onto the list of most popular servers indicates that BEST must be host to a very large number of Web sites.

About the Dataset

The Yahoo dataset started off with 45,494 unique host names. Of these, 790 (about 2%) were hosts specified by IP address only, not by domain name. Most of the Web sites polled were in the US.

The top domains were:

Count	Dom.	Location
20786	com	Commercial and personal sites, mostly US
7515	edu	Educational institutions, mostly US
2706	net	Network providers, mostly US
2004	org	Organizations and non-profit corporations, mostly US
1420	uk	United Kingdom
1414	ca	Canada
879	au	Australia
849	gov	US, government sites (non-military)
837	us	US, sites identified by geographic location
770	de	Germany
502	jp	Japan
446	se	Sweden
437	it	Italy
403	nl	Netherlands
353	fr	France
280	mil	US, military sites
242	ch	Switzerland

211	no	Norway
204	fi	Finland
196	at	Austria

Clearly, the top domains in this dataset are all from countries whose primary language is English. This is due both to the English-centric nature of the Internet and to the fact that the Yahoo database is based in the United States.

Different datasets would yield different counts for the domains, which would certainly change the results of which server software was most popular. For example, servers whose documentation had been translated into different languages would probably be much more popular in countries whose dominant language is not English.

As many people commented after the first survey, it is inaccurate to say that all sites in the US-centric domains are in fact in the US. There are two major reasons why, for example, a domain name that ends in "somecompany.com" might be outside the US:

- The domain name "somecompany.com" might have been given to a non-US company before restrictions were put on the country of origin for the "com" domain.
- The company may be based in the US, but the office hosting the Web server might be located in a different country. For example, the domain name "www.jp.somecompany.com" might indicate a Web site in Japan.

An interesting tidbit from the dataset: 1047 sites (about 2%) used TCP ports other than the standard Web port of 80. This number is significantly lower than in the previous survey, indicating that the use of non-standard ports for new sites is definitely becoming less common. This is good, since using non-standard port numbers makes typing in URLs by hand more prone to error.

Future Surveys

The Web server market is expanding rapidly, although it is not clear whether current Web sites will respond to these new choices by changing server software. There is a great deal of inertia in the market: once you have selected a server, you are hesitant to change even for one that has many new features.

For example, consider GN and WN, two free Unix-based servers written by John Franks. GN has been in use for a year longer than WN. GN is now updated infrequently and has only a few Web-specific features, while WN is actively supported and has a robust and growing set of features for serving Web documents. Yet, there are still more than twice as many GN servers as WN servers, according to the survey.

It will be interesting to see how well new servers with better support and more features fare against the entrenched servers such as NCSA, Apache, and CERN. Will the commercial market,

with well-recognized companies like Netscape, IBM, and Microsoft, be able to grow in the face of many free servers? Will the PC-based servers such as WebSTAR and WebSite thrive as alternatives to Unix-based systems?

The next iteration of this survey will certainly have different results, although it unclear in what way they will change. It is likely that the percentage of sites using newer servers will increase. Also, as Web commerce becomes more pervasive, servers that offer higher security will possibly also increase faster than those with minimal security. At the same time, free Unix servers that are better supported than NCSA and CERN might also increase their share of the market.

If you have comments or suggestions for future surveys, please send them to www-servers@proper.com.

APPENDIX B. SPEC REFERENCE TABLES

The information in this appendix was obtained from John Dimarco's Web site at the University of Toronto - <ftp://ftp.cdf.toronto.edu/pub/spectable>. It contains a comprehensive listing of both SPEC92 and SPEC95 aggregate benchmark statistics. This information is presented as a reference for obtaining SPEC benchmark values for computer hardware. These values can then be used, in conjunction with the heuristic and associated information in Chapter V and Appendix C, to evaluate the hardware for its suitability to satisfy Web site requirements.

What this is: A file of reported SPEC CINT/CFP benchmark results (means only) for various machines. These figures are generally taken from numbers published on the net, or in manufacturer press releases or reports.

This file is organized into eight tables, the first reporting SPECint_base95 and SPECfp_base95, the second reporting SPECint_rate_base95 and SPECfp_rate_base95, the third reporting SPECint95 and SPECfp95, the fourth reporting SPECint_rate95 and SPECfp_rate95, the fifth reporting SPECint_base92 and SPECfp_base92, the sixth reporting SPECint_rate_base92 and SPECfp_rate_base92, the seventh reporting SPECint92 and SPECfp92, and the eighth reporting SPECint_rate92 and SPECfp_rate92. SPECmark89 (obsolete) is no longer reported.

There are no chip-only entries (as opposed to systems with those chips in them); SPEC CINT95/CINT92 CFP95/CFP92 are suites of component-level benchmarks that measure primarily the performance of a system's processor, memory architecture, operating system and compiler. Reporting SPEC results for a chip alone is misleading.

Some specrate numbers have been computed from reported specint/FP92 numbers for various uniprocessor systems. These are indicated by a trailing "c".

Manufacturer estimates, or estimates of any sort, are not normally reported.

This file is available via anonymous ftp from <ftp.cdf.toronto.edu> in the file `/pub/spectable`.

A SPEC FAQ describing the SPEC benchmark suite and the SPEC consortium is periodically posted to `comp.benchmarks`, and can be found on the WWW at "<http://www.specbench.org/spec/specfaq.html>". I strongly recommend reading that document before using these numbers.

More SPEC-related information is available at the SPEC WWW site, "<http://www.specbench.org>", and at the Performance Database Web site, "<http://performance.netlib.org/performance/html/spec.html>".

Note carefully: benchmark results depend not only on processor type, speed, and cache size, but compiler, OS and other machine characteristics that are not reported here. In particular, the compiler can have a significant effect.

Quote:

" While no one benchmark can fully characterize overall system performance, the results of a variety of realistic benchmarks can give valuable insight into expected real performance. "
- SPEC newsletter.

Disclaimer: These numbers have not been verified. Nobody guarantees their correctness, and there is no guarantee that they accurately reflect the true performance of these systems. Furthermore, this is not a publication of the SPEC consortium and is not endorsed by the SPEC consortium in any way.

Please send all corrections, updates, and new entries to jdd@cdf.toronto.edu.

John DiMarco <jdd@cdf.toronto.edu>
Computing Disciplines Facility Systems Manager
University of Toronto
<http://www.cdf.toronto.edu/personal/jdd/jdd.html>

Office: EA201B
Phone: 416-978-1928
Fax: 416-978-1931

***** Legend *****

Guide to Vendor Acronyms:

DEC: Digital Equipment Corporation
DG: Data General
HP: Hewlett-Packard
IBM: International Business Machines
RT: Ross Technology
SGI: Silicon Graphics Inc.
SNI/Py: Siemens-Nixdorf Inc./Pyramid Technology Corp.

Guide to processor families:

88000: 88100
68000: 68040
ALPHA: A21064, A21064A, A21066, A21164
HP PA: PA7100, PA7100LC, PA1.1
i86: 80487SX, 80486DX, 80486DX2, 80486DX4, Pentium
MIPS: R2000, R3000, R4000, R4400, R4600, R6000, R8000
POWER: POWER, POWER2, MPC601(PowerPC), RSC3308, RSC4608

SPARC: SuperSPARC, SuperSPARC-II, HyperSPARC, MicroSPARC, MicroSPARC-II,
 FJMB86902(LSI64911), FJMB86903, RT601, RT605, Weitek PowerUp
 VAX: REX520, SOC, KA46, NVAX, KA650, KA660, KA680

***** TABLE 1: SPECint_base95, SPECfp_base95 *****

Notes:

- SPECint_base95 is derived from the results of eight integer benchmarks compiled with conservative optimization. It is the geometric mean of eight normalized ratios (one for each integer benchmark).
- SPECfp_base95 is derived from the results of ten floating-point benchmarks compiled with conservative optimization. It is the geometric mean of ten normalized ratios (one for each integer benchmark).

System Name	CPU (NUMx)Type	ClkMHz ext/in	Cache Ext+I/D	SPECint base95	SPECfp base95	Info Date	Source Obtained
Sun SS10/40	SuprSP	40	20/16	1.00	1.00	Aug95	SPEC Refc
DEC 250/4/266	A21064A	??/266	2M+16/16	4.18	5.78	Apr95	www.dec
DEC 600/5/266	A21164	38/266	4M+96+8/8	6.43	10.64	Sep95	Digital
DEC 600/5/300	A21164	75/300	4M+96+8/8	7.33	11.59	Sep95	Digital
DEC 600/5/333	A21164	83/333	4M+96+8/8	8.42	12.4	Feb96	Digital
DEC 3000/500	A21064	30/150	512+8/8	2.15	3.65	Sep95	Digital
DEC 3000/700	A21064A	38/225	2M+16/16	3.66	5.71	Sep95	Digital
DEC 3000/900	A21064A	39/275	2M+16/16	4.24	6.29	Sep95	Digital
DEC 2[01]00/5/250	A21164	35/250	4M+96+8/8	5.96	8.39	Feb96	Digital
DEC 2[01]00/5/300	A21164	42/300	4M+96+8/8	7.03	9.26	Feb96	Digital
DEC 2[01]00/5/300	2xA21164	42/300	4M+96+8/8	?	9.0	Feb96	Digital
DEC 2100/5/300	4xA21164	42/300	4M+96+8/8	?	9.0	Feb96	Digital
DEC 8[24]00/5/300	A21164	5/300	4M+96+8/8	7.43	11.7	Feb96	Digital
DEC 8[24]00/5/300	2xA21164	75/300	4M+96+8/8	?	11.9	Feb96	Digital
DEC 8[24]00/5/300	4xA21164	75/300	4M+96+8/8	?	11.7	Feb96	Digital
DEC 8[24]00/5/300	6xA21164	75/300	4M+96+8/8	?	11.8	Feb96	Digital
DEC 8400/5/300	8xA21164	75/300	4M+96+8/8	?	11.8	Feb96	Digital
DEC 8[24]00/5/350	A21164	88/350	4M+96+8/8	8.82	13.2	Feb96	Digital
DEC 8400/5/350	8xA21164	88/350	4M+96+8/8	?	28.9	Feb96	Digital
Dell DimensionXPS	Pentium	66/100	512+8/8	3.16	2.09	Jan96	www.intel
Dell DimensionXPS	Pentium	60/120	512+8/8	3.53	2.26	Jan96	www.intel
Dell DimensionXPS	Pentium	66/133	512+8/8	3.90	2.48	Jan96	www.intel
Dell Optiplex	Pentium	60/120	512+8/8	3.51	2.16	Jan96	www.intel
Dell Optiplex	Pentium	66/133	512+8/8	3.90	2.32	Jan96	www.intel
Gateway P5-75	Pentium	50/75	256+8/8	2.31	1.50	Jan96	www.intel
Gateway P5-90	Pentium	60/90	256+8/8	2.74	1.86	Jan96	www.intel
Gateway P5-100	Pentium	66/100	256+8/8	3.05	2.07	Jan96	www.intel
HP C100	PA7200	100	256/256	3.67	6.20	Dec95	www.hp
HP C110	PA7200	120	256/256	4.41	7.45	Dec95	www.hp
HP 9000/735	PA7100	99	256/256	3.13	3.97	Sep95	SPEC
HP 9000/735	PA7100	125	256/256	3.88	4.54	Sep95	SPEC
HP 9000/J200	PA7200	100	256/256	3.27	6.22	Sep95	SPEC
HP 9000/J210	PA7200	120	256/256	3.93	7.51	Sep95	SPEC

HP 9000/J210	2xPA7200	120	256/256	?	9.91	Sep95	SPEC
IBM C10	MPC601	80	1M+32	2.06	2.94	Sep95	SPEC
IBM C20	MPC604	120	1M+16/16	3.38	3.48	Sep95	SPEC
IBM E20	MPC604	100	512+16/16	3.47	3.11	Oct95	www.ibm
IBM 43P	MPC604	133	512+16/16	4.07	3.27	Sep95	SPEC
IBM 39H/3CT	POWER2	66.7	2M+32/128	2.88	9.28	Sep95	SPEC
Intel XXpress	Pentium	66/100	1M+8/8		2.06	Jan96	www.intel
Intel XXpress	Pentium	60/120	1M+8/8	3.72	2.24	Jan96	www.intel
Intel XXpress	Pentium	66/133	1M+8/8	4.14	2.48	Jan96	www.intel
Intel XXpress	Pentium	50/150	1M+8/8	4.27	3.04	Jan96	www.intel
Intel XXpress	Pentium	55/166	1M+8/8	4.76	3.37	Jan96	www.intel
Intel Alder	PentiumPro	150	256+8/8	6.08	4.76	Jan96	www.intel
Intel Alder	PentiumPro	166	512+8/8	7.11	5.47	Jan96	www.intel
Intel Alder	PentiumPro	180	256+8/8	7.29	5.40	Jan96	www.intel
Intel Alder	PentiumPro	200	256+8/8	8.09	5.99	Jan96	www.intel
Intel Aurora	PentiumPro	150	256+8/8		4.22	Jan96	www.intel
Intel Aurora	PentiumPro	166	256+8/8	?	4.72	Jan96	www.intel
SNI/Pyr 2-225	R4600	133	?+16/16	2.31	?	Sep95	c.bmarks
SNI/Pyr 4-630	R4400	100/200	?M+16/16	3.79	?	Sep95	c.bmarks
SNI/Pyr RM200-C20	R4600	133	16/16	2.53	?	Dec95	c.bmarks
SNI/Pyr RM300-C20	R4600	133	16/16	2.53	?	Dec95	c.bmarks
SNI/Pyr RM300-C60	R4400	100/200	1M+16/16	3.41	?	Dec95	c.bmarks
SNI/Pyr RM400-C70	R4400	100/200	1M+16/16	3.72	?	Dec95	c.bmarks
Sun SS10/40	SuprSP	40	20/16	1.06	1.13	Mar96	c.bmarks
Sun SS[45]/110	MicroSP2	110	16/8	1.37	1.88	Mar96	c.bmarks
Sun SS20/71	SuprSP2	50/75	1M+20/16	2.82	2.96	Mar96	c.bmarks
Sun SS20/151	HyperSP	50/150	512+8/0	3.77	4.73	Mar96	c.bmarks
Sun Ultra1/140	UltraSP	143	512+16/16	4.52	7.73	Mar96	c.bmarks
Sun Ultra1/170	UltraSP	167	512+16/16	5.26	8.45	Mar96	c.bmarks
Sun Ultra2/2200	2xUltraSP	200	M+16/16	6.41	11.6	Mar96	c.bmarks

***** TABLE 2: SPECint_rate_base95, SPECfp_rate_base95 *****

Notes:

- SPECint_rate_base95 is derived from the results of eight integer benchmarks compiled with conservative optimization. It is the geometric mean of eight normalized throughput ratios (one for each integer benchmark).
- SPECfp_rate_base95 is derived from the results of ten floating-point benchmarks compiled with conservative optimization. It is the geometric mean of ten normalized throughput ratios (one for each integer benchmark).

System Name	CPU (NUMx)Type	ClkMHz ext/in	Cache Ext+I/D	SPECint rt_bs95	SPECfp rt_bs95	Info Date	Source Obtained
DEC 3000/500	21064	150	512+8/8	19.4	32.9	Sep95	SPEC
DEC 3000/700	21064A	225	2M+16/16	32.9	52.2	Sep95	SPEC
DEC 3000/900	21064A	275	2M+16/16	38.2	56.5	Sep95	SPEC
DEC 2[01]00/5/250	A21164	35/250	4M+96+8/8	53.6	75.5	Feb96	Digital

DEC 2[01]00/5/250	2xA21164	35/250	4M+96+8/8	108.0	139.0	Feb96	Digital
DEC 2100/5/250	4xA21164	35/250	4M+96+8/8	210.0	216.0	Feb96	Digital
DEC 2[01]00/5/300	A21164	42/300	4M+96+8/8	63.3	83.4	Feb96	Digital
DEC 2[01]00/5/300	2xA21164	42/300	4M+96+8/8	125.0	155.0	Feb96	Digital
DEC 2100/5/300	4xA21164	42/300	4M+96+8/8	246.0	246.0	Feb96	Digital
DEC 8[24]00/5/300	A21164	75/300	4M+96+8/8	64.2	104.0	Feb96	Digital
DEC 8[24]00/5/300	2xA21164	75/300	4M+96+8/8	131.0	205.0	Feb96	Digital
DEC 8[24]00/5/300	4xA21164	75/300	4M+96+8/8	261.0	400.0	Feb96	Digital
DEC 8[24]00/5/300	6xA21164	75/300	4M+96+8/8	388.0	587.0	Feb96	Digital
DEC 8400/5/300	8xA21164	75/300	4M+96+8/8	525.0	753.0	Feb96	Digital
DEC 8400/5/300	10xA21164	75/300	4M+96+8/8	642.0	797.0	Feb96	Digital
DEC 8400/5/300	12xA21164	75/300	4M+96+8/8	767.0	904.0	Feb96	Digital
DEC 8[24]00/5/350	6xA21164	88/350	4M+96+8/8	449	493	Feb96	Digital
DEC 8400/5/350	12xA21164	88/350	4M+96+8/8	890	1025	Feb96	Digital
SNI/Pyt 2-225	R4600	133	?+16/16	20.8		Sep95	c.bmarks
SNI/Pyt 4-730	2xR4400	100/200	?M+16/16	66.4	?	Sep95	c.bmarks
SNI/Pyt 6-3[24]0	8xR4400	100/200	4M+16/16	234	?	Sep95	c.bmarks
SNI/Pyt 6-620	24xR4400	100/200	4M+16/16	658	?	Sep95	c.bmarks
SNI/Pyt RM200-C20	R4600	133	16/16	22.8	?	Dec95	c.bmarks
SNI/Pyt RM300-C20	R4600	133	16/16	22.8	?	Dec95	c.bmarks
SNI/Pyt RM300-C62	2xR4400	100/200	1M+16/16	64.4	?	Dec95	c.bmarks
SNI/Pyt RM400-C70	2xR4400	100/200	4M+16/16	65.0	?	Dec95	c.bmarks
SNI/Pyt RM400-C70	4xR4400	100/200	4M+16/16	131	?	Dec95	c.bmarks

***** TABLE 3: SPECint95, SPECfp95 *****

Notes:

- SPECint95 is derived from the results of eight integer benchmarks compiled with aggressive optimization. It is the geometric mean of eight normalized ratios (one for each integer benchmark).
- SPECfp95 is derived from the results of ten floating-point benchmarks compiled with aggressive optimization. It is the geometric mean of ten normalized ratios (one for each integer benchmark).
- Note that the level of optimization is not mandated. While highly aggressive optimization is permitted, results derived from benchmarks compiled with conservative optimization (as in SPECbase) can be submitted.

System Name	CPU (NUMx)Type	ClkMHz ext/in	Cache Ext+I/D	SPECint 95	SPECfp 95	Info Date	Source Obtained
DEC 250/4/266	A21064A	??/266	2M+16/16	4.18	5.78	Apr95	www.dec
DEC 600/5/266	A21164	38/266	4M+96+8/8	6.43	11.18	Sep95	Digital
DEC 600/5/300	A21164	75/300	4M+96+8/8	7.33	12.16	Sep95	Digital
DEC 600/5/333	A21164	83/333	4M+96+8/8	9.23	13.2	Feb96	Digital
DEC 3000/500	A21064	30/150	512+8/8	2.15	3.65	Sep95	Digital
DEC 3000/700	A21064A	38/225	2M+16/16	3.66	5.71	Sep95	Digital
DEC 3000/900	A21064A	39/275	2M+16/16	4.24	6.29	Sep95	Digital
DEC 2[01]00/5/250	A21164	35/250	4M+96+8/8	5.96	8.39	Feb96	Digital
DEC 2[01]00/5/300	A21164	42/300	4M+96+8/8	7.03	9.64	Feb96	Digital
DEC 2[01]00/5/300	2xA21164	42/300	4M+96+8/8	?	14.0	Feb96	Digital

DEC 2100/5/300	4xA21164	42/300	4M+96+8/8	?	19.2	Feb96	Digital
DEC 8[24]00/5/300	A21164	75/300	4M+96+8/8	7.43	12.4	Feb96	Digital
DEC 8[24]00/5/300	2xA21164	75/300	4M+96+8/8	?	18.1	Feb96	Digital
DEC 8[24]00/5/300	4xA21164	75/300	4M+96+8/8	?	25.9	Feb96	Digital
DEC 8[24]00/5/300	6xA21164	75/300	4M+96+8/8	?	30.1	Feb96	Digital
DEC 8400/5/300	8xA21164	75/300	4M+96+8/8	?	33.5	Feb96	Digital
DEC 8[24]00/5/350	A21164	88/350	4M+96+8/8	10.10	14.2	Feb96	Digital
DEC 8400/5/350	8xA21164	88/350	4M+96+8/8	?	38.5	Feb96	Digital
Dell DimensionXPS	Pentium	66/100	512+8/8	3.16	2.75	Jan96	www.intel
Dell DimensionXPS	Pentium	60/120	512+8/8	3.53	2.92	Jan96	www.intel
Dell DimensionXPS	Pentium	66/133	512+8/8	3.90	3.28	Jan96	www.intel
Dell Optiplex	Pentium	60/120	512+8/8	3.51	2.80	Jan96	www.intel
Dell Optiplex	Pentium	60/133	512+8/8	3.90	2.99	Jan96	www.intel
Gateway P5-75	Pentium	50/75	256+8/8	2.31	2.02	Jan96	www.intel
Gateway P5-90	Pentium	60/90	256+8/8	2.74	2.39	Jan96	www.intel
Gateway P5-100	Pentium	66/100	256+8/8	3.05	2.72	Jan96	www.intel
HAL 330	SPARC64	100	128/128	4.2	7.73	Feb96	www.hal
HAL 350	SPARC64	118	128/128	4.9	9.03	Feb96	www.hal
HP 9000/735	PA7100	99	256/256	3.22	4.06	Sep95	SPEC
HP 9000/735	PA7100	125	256/256	3.97	4.61	Sep95	SPEC
HP 9000/J200	PA7200	100	256/256	3.52	6.32	Sep95	SPEC
HP 9000/J210	PA7200	120	256/256	4.21	7.60	Sep95	SPEC
HP 9000/J210	2xPA7200	120	256/256	?	10.10	Sep95	SPEC
HP 9000/K420	PA7200	120	1M/1M	4.61	8.24	Feb96	www.hp
Intel XXpress	Pentium	66/100	1M+8/8	3.30	2.59	Jan96	www.intel
Intel XXpress	Pentium	60/120	1M+8/8	3.72	2.81	Nov95	www.intel
Intel XXpress	Pentium	66/133	1M+8/8	4.14	3.12	Nov95	www.intel
Intel XXpress	Pentium	50/150	1M+8/8	4.27	3.04	Jan96	www.intel
Intel XXpress	Pentium	55/166	1M+8/8	4.76	3.37	Jan96	www.intel
Intel Alder	PentiumPro	150	256+8/8	6.08	5.42	Jan96	www.intel
Intel Alder	PentiumPro	166	512+8/8	7.11	6.21	Jan96	www.intel
Intel Alder	PentiumPro	180	256+8/8	7.29	6.08	Jan96	www.intel
Intel Alder	PentiumPro	200	256+8/8	8.09	6.75	Jan96	www.intel
Intel Aurora	PentiumPro	150	256+8/8	?	4.71	Jan96	www.intel
Intel Aurora	PentiumPro	166	256+8/8	?	5.20	Jan96	www.intel
SNI/Pyr 2-225	R4600	133	?+16/16	2.41	?	Sep95	c.bmarks
SNI/Pyr 4-630	R4400	100/200	?M+16/16	3.95	?	Sep95	c.bmarks
SNI/Pyr RM200-C20	R4600	133	16/16	2.64	?	Dec95	c.bmarks
SNI/Pyr RM300-C20	R4600	133	16/16	2.64	?	Dec95	c.bmarks
SNI/Pyr RM300-C60	R4400	100/200	1M+16/16	3.55	?	Dec95	c.bmarks
SNI/Pyr RM400-C70	R4400	100/200	4M+16/16	3.92	?	Dec95	c.bmarks
Sun SS10/40	SuprSP	40	20/16	1.13	1.38	Mar96	c.bmarks
Sun SS[45]/110	MicroSP2	110	16/8	1.59	1.99	Mar96	c.bmarks
Sun SS20/71	SuprSP2	50/75	1M+20/16	3.11	3.10	Mar96	c.bmarks
Sun SS20/151	HyperSP	50/150	512+8/0	4.02	4.73	Mar96	c.bmarks
Sun Ultra1/140	UltraSP	143	512+16/16	4.66	7.90	Mar96	c.bmarks
Sun Ultra1/170	UltraSP	167	512+16/16	5.56	9.06	Mar96	c.bmarks
Sun Ultra2/2200	2xUltraSP	200	1M+16/16	6.85	12.9	Mar96	c.bmarks

***** TABLE 4: SPECint_rate95, SPECfp_rate95 *****

Notes:

- SPECint_rate95 is derived from the results of eight integer benchmarks compiled with aggressive optimization. It is the geometric mean of eight normalized throughput ratios (one for each integer benchmark).
- SPECfp_rate95 is derived from the results of ten floating-point benchmarks compiled with aggressive optimization. It is the geometric mean of ten normalized throughput ratios (one for each integer benchmark).

System Name	CPU (NUMx)Type	ClkMHz ext/in	Cache Ext+I/D	SPECint_rate95	SPECfp_rate95	Info Date	Source Obtained
SNI/Pyr 2-225	R4600	133	?+16/16	20.8	?	Sep95	c.bmarks
SNI/Pyr 4-730	2xR4400	100/200	?M+16/16	69.0	?	Sep95	c.bmarks
SNI/Pyr 6-3[24]0	8xR4400	100/200	4M+16/16	247	?	Sep95	c.bmarks
SNI/Pyr 6-620	24xR4400	100/200	4M+16/16	658	?	Sep95	c.bmarks
SNI/Pyr RM200-C20	R4600	133	16/16	23.7	?	Dec95	c.bmarks
SNI/Pyr RM300-C20	R4600	133	16/16	23.7	?	Dec95	c.bmarks
SNI/Pyr RM300-C62	2xR4400	100/200	1M+16/16	65.2	?	Dec95	c.bmarks
SNI/Pyr RM400-C70	2xR4400	100/200	4M+16/16	67.2	?	Dec95	c.bmarks
SNI/Pyr RM400-C70	4xR4400	100/200	4M+16/16	131	?	Dec95	c.bmarks
DEC 2[01]00/5/250	A21164	35/250	4M+96+8/8	53.6	75.5	Feb96	Digital
DEC 2[01]00/5/250	2xA21164	35/250	4M+96+8/8	108.0	139.0	Feb96	Digital
DEC 2100/5/250	4xA21164	35/250	4M+96+8/8	210.0	216.0	Feb96	Digital
DEC 2[01]00/5/300	A21164	42/300	4M+96+8/8	63.3	86.7	Feb96	Digital
DEC 2[01]00/5/300	2xA21164	42/300	4M+96+8/8	125.0	161.0	Feb96	Digital
DEC 2100/5/300	4xA21164	42/300	4M+96+8/8	246.0	251.0	Feb96	Digital
DEC 8[24]00/5/300	A21164	75/300	4M+96+8/8	64.2	109.0	Feb96	Digital
DEC 8[24]00/5/300	2xA21164	75/300	4M+96+8/8	131.0	215.0	Feb96	Digital
DEC 8[24]00/5/300	4xA21164	75/300	4M+96+8/8	261.0	420.0	Feb96	Digital
DEC 8[24]00/5/300	6xA21164	75/300	4M+96+8/8	388.0	601.0	Feb96	Digital
DEC 8400/5/300	8xA21164	75/300	4M+96+8/8	525.0	789.0	Feb96	Digital
DEC 8400/5/300	10xA21164	75/300	4M+96+8/8	642.0	817.0	Feb96	Digital
DEC 8400/5/300	12xA21164	75/300	4M+96+8/8	767.0	919.0	Feb96	Digital
DEC 8[24]00/5/350	6xA21164	88/350	4M+96+8/8	506	505	Feb96	Digital
DEC 8400/5/350	12xA21164	88/350	4M+96+8/8	1004	1039	Feb96	Digital

***** TABLE 5: SPECbase92, SPECbaseFP92 *****

Notes:

- SPECint92 is derived from the results of a set of integer benchmarks, and can be used to estimate a machine's single-tasking performance on integer code. SPECbase92 is a variant of SPECint92 that reports "baseline" results, using stricter run rules.
- SPECfp92 is derived from the results of a set of floating point benchmarks, and can be used to estimate a machine's single-tasking performance on floating-point code. SPECfp_base92 is a variant of SPECfp92 that reports "baseline" results, using stricter run rules.

System Name	CPU (NUMx)Type	ClkMHz ext/in	Cache Ext+I/D	SPECint base92	SPECfp base92	Info Date	Source Obtained
DEC VAX11/780	VAX	5	2	1.0	1.0	Jan89	SPEC Ref
DEC 3000/900	A21064A	39/275	2M+16/16	178.4	244.6	Jul94	Digital
DEC 7000/710	A21064A	39/275	4M+16/16	180.0	265.8	Aug94	Digital
DEC 200/4/100	A21064	??/100	512+8/8	68.6	90.6	Feb95	Digital
DEC [24]00/4/166	A21064	33/166	512+8/8	100.1	128.4	Jul95	Digital
DEC [24]00/4/233	A21064A	39/233	512+16/16	137.4	174.6	Apr95	Digital
DEC 250/4/266	A21064A	??/266	2M+16/16	182.6	246.8	Apr95	www.dec
DEC 600/5/266	A21164	38/266	2M+96+8/8	257.1	365.0	Jul95	Digital
DEC 600/5/266	A21164	38/266	4M+96+8/8	260.6	386.1	Jul95	Digital
DEC 600/5/300	A21164	42/300	4M+96+8/8	279.8	436.1	Jul95	Digital
DEC 1000/4/200	A21064	40/200	2M+8/8	123.3	165.7	Nov94	Digital
DEC 2[01]00/4/200	A21064	47/190	1M+8/8	117.5	154.3	Nov94	Digital
DEC 2[01]00/4/233	A21064A	38/233	1M+16/16	163.7	192.3	Apr95	Digital
DEC 2[01]00/4/275	A21064A	39/275	4M+16/16	187.8	259.5	Apr95	Digital
DEC 2[01]00/5/250	A21164	35/250	4M+96+8/8	244.7	356.3	Apr95	Digital
DEC 2[01]00/5/300	A21164	42/300	4M+96+8/8	283.6	420.0	Feb96	Digital
DEC 8[24]00/5/300	A21164	75/300	4M+96+8/8	314.4	444.0	Apr95	Digital
DEC 8[24]00/5/350	A21164	88/350	4M+96+8/8	72.2	518.5	Feb96	Digital
HP E45	PA7100LC	80	256	74.5	110.6	Mar95	www.hp
HP E55	PA7100LC	96	1M	96.1	149.9	Mar95	www.hp
IBM C20	MPC604	120	16/16	95.2	106.4	Jun95	www.ibm
IBM C20	MPC604	120	1M+16/16	124.3	137.2	Jun95	www.ibm
IBM E20	MPC604	100	512+16/16	110.9	121.1	Oct95	www.ibm
IBM 42[T W]	MPC604	120	16/16	95.2	106.4	Jun95	www.ibm
IBM 42[T W]	MPC604	120	512+16/16	121.8	133.5	Jun95	www.ibm
IBM 43P	MPC604	100	256+16/16	104.3	104.8	Jun95	www.ibm
IBM 43P	MPC604	120	512+16/16	127.1	129.0	Jun95	www.ibm
IBM 43P	MPC604	133	512+16/16	142.2	146.2	Jun95	www.ibm
IBM 591/R21	POWER2	77	32/256	121.2	268.2	Jul95	www.ibm
SNI/Pyr PC/E5S	Pentium	30/90	256+8/8	82.9	67.9	Jul94	c.bmarks
SNI/Pyr PC/E5S	Pentium	33/100	256+8/8	92.4	75.4	Jul94	c.bmarks
SNI/Pyr PC/D5T	Pentium	30/60	256+8/8	63.9	48.3	Nov94	c.bmarks
SNI/Pyr PC/D5T	Pentium	30/90	256+8/8	83.0	62.4	Nov94	c.bmarks
SNI/Pyr 2-12[05]	R4600	50/100	16/16	71.5	?	Nov94	c.bmarks
SNI/Pyr 4-220	R4400	50/100	512+16/16	63.7	?	Nov94	c.bmarks
SNI/Pyr 4-3[34]0	R4400	50/100	1M+16/16	67.7	?	Nov94	c.bmarks
SNI/Pyr 4-420	R4400	75/150	512+16/16	87.1	?	Nov94	c.bmarks

SNI/Pyr 4-4[34]0	R4400	75/150	1M+16/16	94.0	?	Nov94	c.bmarks
SNI/Pyr 4-5[34]0	R4400	75/150	4M+16/16	101.6	?	Nov94	c.bmarks
SNI/Pyr 6-3[24]0	R4400	100/200	4M+16/16	132.1	?	Jun95	SNI/Pyr
SNI/Pyr RM200-C20	R4600	133	16/16	97.5	?	Dec95	c.bmarks
SNI/Pyr RM300-C20	R4600	133	16/16	97.5	?	Dec95	c.bmarks
SNI/Pyr RM300-C60	R4400	100/200	1M+16/16	127.9	?	Dec95	c.bmarks
SNI/Pyr RM400-C70	R4400	100/200	4M+16/16	138.9	?	Dec95	c.bmarks
SNI/Pyr RM1000	R4400	100/200	4M+16/16	139.1	?	Aug95	SNI/Pyr
Sun SS20/61	SuprSP	50/60	1M+20/16	?	95.8	Jun94	SPEC news
Sun SS20/612	2xSuprSP	50/60	1M+20/16	?	111.0	Sep94	SPEC news
Sun SS20/HS11	HyperSP	50/100	256+8/0	?	117.8	Dec94	SPEC news
Sun SS1000E	8xSuprSP	50/60	1M+20/16	15414	17114	Jan96	SunPromo
Sun SS1000E	8xSuprSP2	50/85	1M+20/16	21758	20851	Jan96	SunPromo
Sun SC2000E	20xSuprSP2	50/60	2M+20/16	38213	44722	Jan96	SunPromo
Sun SC2000E	20xSuprSP2	50/85	2M+20/16	57997	54206	Jan96	SunPromo
Intel XXpress	Pentium	66/100	1M+8/8	126.2	?	Jan96	www.intel
Intel XXpress	Pentium	60/120	1M+8/8	143.6	?	Jan96	www.intel
Intel XXpress	Pentium	66/133	1M+8/8	160.5	?	Jan96	www.intel
Intel XXpress	Pentium	??/150	1M+8/8	165.2	?	Jan96	www.intel
Intel XXpress	Pentium	??/166	1M+8/8	181.6	?	Jan96	www.intel
Intel Alder	PentiumPro	150	256+8/8	228.1	?	Jan96	www.intel
Intel Alder	PentiumPro	180	256+8/8	268.1	?	Jan96	www.intel
Intel Alder	PentiumPro	200	256+8/8	296.5	?	Jan96	www.intel

***** TABLE 6: Integer/FP SPECrate_base92 *****

Notes:

- Integer SPECrate is derived from the results of a set of integer benchmarks run multiple times simultaneously, and can be used to estimate a machine's overall multi-tasking throughput for integer code. It is typically used on MP machines.
- Floating-Point SPECrate is derived from the results of a set of floating-point benchmarks run multiple times simultaneously, and can be used to estimate a machine's overall multi-tasking throughput for FP code. It is typically used on MP machines.
- SPECrate_base is a variant of SPECrate that reports "baseline" results, using stricter run rules.
- Computed SPECrate_base figures are indicated by "c". They're computed from SPECint_base92, fp92 (for uniprocessors) using a scaling factor. This number is usually slightly less than or equal to a measured specbaserate on a uniprocessor. The scaling factor is the number of seconds in a week, divided by the time of the longest-running benchmark on the reference SPEC VAX 11/780, which is 604800/25500, or about 23.7.

System Name	CPU (NUMx)Type	ClkMHz ext/in	Cache Ext+I/D	SPECint rt_bs92	SPECfp rt_bs92	Info Date	Source Obtained
DEC VAX11/780	VAX	5	2	24c	24c	Jan89	SPEC Ref
DEC 3000/700	A21064A	38/225	2M+16/16	3682	5106	Jul94	Digital
DEC 3000/900	A21064A	39/275	2M+16/16	4402	5798	Jul94	Digital
DEC 7000/710	A21064A	39/275	4M+16/16	4222	6159	Aug94	Digital

DEC 7000/720	2xA21064A	39/275	4M+16/16	8550	12344	Aug94	Digital
DEC 7000/740	4xA21064A	39/275	4M+16/16	14922	24711	Aug94	Digital
DEC 7000/760	6xA21064A	39/275	4M+16/16	22267	37273	Aug94	Digital
DEC 200/4/100	A21064	??/100	512+8/8	1626	2133	Feb95	Digital
DEC [24]00/4/166	A21064	33/166	12+8/8	2371	3009	Jul95	Digital
DEC [24]00/4/233	A21064A	39/233	512+16/16	3275	4041	Apr95	Digital
DEC 250/4/266	A21064A	??/266	2M+16/16	4300	5726	Apr95	www.dec
DEC 600/5/266	A21164	38/266	2M+96+8/8	6114	8706	Jul95	Digital
DEC 600/5/266	A21164	38/266	4M+96+8/8	6256	9255	Jul95	Digital
DEC 600/5/300	A21164	75/300	4M+96+8/8	6429	10558	Jul95	Digital
DEC 1000/4/200	A21064	40/200	2M+8/8	2944	3906	Nov94	Digital
DEC 2[01]00/4/200	A21064	47/190	1M+8/8	2786	3594	Nov94	Digital
DEC 2[01]00/4/200	2xA21064	47/190	1M+8/8	5495	6914	Nov94	Digital
DEC 2100/4/200	4xA21064	47/190	1M+8/8	10537	12384	Nov94	Digital
DEC 2[01]00/4/233	A21064A	38/233	1M+16/16	3842	4575	Apr95	Digital
DEC 2[01]00/4/233	2xA21064A	38/233	1M+16/16	7367	8605	Apr95	Digital
DEC 2100/4/233	4xA21064A	38/233	1M+16/16	14494	15741	Apr95	Digital
DEC 2[01]00/4/275	A21064A	39/275	4M+16/16	4423	6182	Apr95	Digital
DEC 2[01]00/4/275	2xA21064A	39/275	4M+16/16	8617	12373	Apr95	Digital
DEC 2100/4/275	4xA21064A	39/275	4M+16/16	16963	24273	Apr95	Digital
DEC 2[01]00/5/250	A21164	35/250	4M+96+8/8	6175	8448	Apr95	Digital
DEC 2[01]00/5/250	2xA21164	35/250	4M+96+8/8	11556	17068	Apr95	Digital
DEC 2100/5/250	4xA21164	35/250	4M+96+8/8	22017	33127	Apr95	Digital
DEC 2[01]00/5/300	A21164	42/300	4M+96+8/8	7148	10125	Feb96	Digital
DEC 2[01]00/5/300	2xA21164	42/300	4M+96+8/8	12559	19665	Feb96	Digital
DEC 2100/5/300	4xA21164	42/300	4M+96+8/8	22202	39198	Feb96	Digital
DEC 8[24]00/5/300	A21164	75/300	4M+96+8/8	7831	10632	Apr95	Digital
DEC 8[24]00/5/300	2xA21164	75/300	4M+96+8/8	15691	21225	Apr95	Digital
DEC 8[24]00/5/300	4xA21164	75/300	4M+96+8/8	30772	42497	Apr95	Digital
DEC 8[24]00/5/300	6xA21164	75/300	4M+96+8/8	46584	63388	Apr95	Digital
DEC 8400/5/300	8xA21164	75/300	4M+96+8/8	59901	83108	Apr95	Digital
DEC 8400/5/300	10xA21164	75/300	4M+96+8/8	74347	102194	Apr95	Digital
DEC 8400/5/300 1	2xA21164	75/300	4M+96+8/8	82663	121155	Apr95	Digital
DEC 8[24]00/5/350	A21164	88/350	4M+96+8/8	8739	12108	Feb96	Digital
DEC 8[24]00/5/350	6xA21164	88/350	4M+96+8/8	51394	73044	Feb96	Digital
DEC 8400/5/350	12xA21164	88/350	4M+96+8/8	98348	146114	Feb96	Digital
HP E45	PA7100LC	80	256	1767c	2623c	Mar95	www.hp
HP E55	PA7100LC	96	1M	2279c	3555c	Mar95	www.hp
IBM C20	MPC604	120	16/16	2258c	2524c	Jun95	www.ibm
IBM C20	MPC604	120	1M+16/16	2948c	3254c	Jun95	www.ibm
IBM E20	MPC604	100	512+16/16	2630c	2872c	Oct95	www.ibm
IBM 42[T W]	MPC604	120	16/16	2258c	2524c	Jun95	www.ibm
IBM 42[T W]	MPC604	120	512+16/16	2889c	3166c	Jun95	www.ibm
IBM 43P	MPC604	100	256+16/16	2474c	2486c	Jun95	www.ibm
IBM 43P	MPC604	120	512+16/16	3015c	3060c	Jun95	www.ibm
IBM 43P	MPC604	133	512+16/16	3373c	3468c	Jun95	www.ibm
IBM 591/R21	POWER2	77	32/256	2875c	6361c	Jul95	www.ibm
IBM R30	2xMPC601	75	1M+32	325	3953	Jul95	www.ibm
IBM R30	4xMPC601	75	1M+32	6354	7808	Jul95	www.ibm
IBM R30	8xMPC601	75	1M+32	10072	14415	Jul95	www.ibm
SNI/Pyr PC/E5S	Pentium	30/90	256+8/8	1966c	1610c	Jul94	c.bmarks
SNI/Pyr PC/E5S	Pentium	33/100	256+8/8	2192c	1788c	Jul94	c.bmarks

SNI/Pyr PC/D5T	Pentium	30/60	256+8/8	1470c	1111c	Nov94	c.bmarks
SNI/Pyr PC/D5T	Pentium	30/90	256+8/8	1909c	1435c	Nov94	c.bmarks
SNI/Pyr 2-12[05]	R4600	50/100	16/16	1645c	?	Nov94	c.bmarks
SNI/Pyr 4-220	R4400	50/100	512+16/16	1465c	?	Nov94	c.bmarks
SNI/Pyr 4-3[34]0	R4400	50/100	1M+16/16	1557c	?	Nov94	c.bmarks
SNI/Pyr 4-420	R4400	75/150	512+16/16	2003c	?	Nov94	c.bmarks
SNI/Pyr 4-4[34]0	R4400	75/150	1M+16/16	2162c	?	Nov94	c.bmarks
SNI/Pyr 4-5[34]0	R4400	75/150	4M+16/16	2337c	?	Nov94	c.bmarks
SNI/Pyr 6-5[34]0	12xR4400	75/150	4M+16/16	22878	?	Nov94	c.bmarks
SNI/Pyr 6-5[34]0	16xR4400	75/150	4M+16/16	29316	?	Nov94	c.bmarks
SNI/Pyr 6-5[34]0	20xR4400	75/150	4M+16/16	35111	?	Nov94	c.bmarks
SNI/Pyr 6-5[34]0	24xR4400	75/150	4M+16/16	39427	?	Nov94	c.bmarks
SNI/Pyr 6-3[24]0	R4400	100/200	4M+16/16	3148	?	Jun95	SNI/Pyr
SNI/Pyr 6-3[24]0	2xR4400	100/200	4M+16/16	6122	?	Jun95	SNI/Pyr
SNI/Pyr 6-3[24]0	4xR4400	100/200	4M+16/16	11836	?	Jun95	SNI/Pyr
SNI/Pyr 6-3[24]0	8xR4400	100/200	4M+16/16	22192	?	Jun95	SNI/Pyr
SNI/Pyr 6-620	12xR4400	100/200	4M+16/16	33780	?	Jun95	SNI/Pyr
SNI/Pyr 6-620	16xR4400	100/200	4M+16/16	42953	?	Jun95	SNI/Pyr
SNI/Pyr 6-620	24xR4400	100/200	4M+16/16	61249	?	Jun95	SNI/Pyr
SNI/Pyr RM200-C20	R4600	133	16/16	2383	?	Dec95	c.bmarks
SNI/Pyr RM300-C20	R4600	133	16/16	2383	?	Dec95	c.bmarks
SNI/Pyr RM300-C60	R4400	100/200	1M+16/16	3088	?	Dec95	c.bmarks
SNI/Pyr RM300-C62	2xR4400	100/200	1M+16/16	6079	?	Dec95	c.bmarks
SNI/Pyr RM400-C70	R4400	100/200	4M+16/16	3294c	?	Dec95	c.bmarks
SNI/Pyr RM400-C70	2xR4400	100/200	4M+16/16	6275	?	Dec95	c.bmarks
SNI/Pyr RM400-C70	4xR4400	100/200	4M+16/16	11997	?	Dec95	c.bmarks
SNI/Pyr RM1000	R4400	100/200	4M+16/16	3299c	?	Aug95	SNI/Pyr
Sun SS1000E	8xSuprSP	50/60	1M+20/16	13423	15572	Oct95	Sunflash
Sun SS1000E	8xSuprSP2	50/85	1M+20/16	20225	18741	Oct95	Sunflash
Sun SC2000E	20xSuprSP2	50/60	2M+20/16	33702	41857	Oct95	Sunflash
Sun SC2000E	20xSuprSP2	50/85	2M+20/16	53714	51489	Oct95	Sunflash
Cray CS6400	48xSuprSP	55/60	2M+20/16	75275	95943	Nov94	Cray
Cray CS6400	56xSuprSP	55/60	2M+20/16	82851	109477	Nov94	Cray
Cray CS6400	64xSuprSP	55/60	2M+20/16	92844	122061	Nov94	Cray

***** TABLE 7: SPECint92, SPECfp92 *****

Notes:

- SPECint92 is derived from the results of a set of integer benchmarks, and can be used to estimate a machine's single-tasking performance on integer code.
- SPECfp92 is derived from the results of a set of floating point benchmarks, and can be used to estimate a machine's single-tasking performance on floating-point code.

System Name	CPU (NUMx)Type	ClkMHz ext/in	Cache Ext+I/D	SPECint 92	SPECfp 92	Info Date	Source Obtained
DEC VAX11/780	VAX	5	2	1.0	1.0	Jan89	SPEC Ref
ALR PowerVEISA	80487SX	20	64+8	10.7	4.9	Mar93	SPEC news

CDC 4330	R3000	33	32/32	24.9	23.9	Sep92	SPEC news
CDC 4360	R3000	33	64/64	24.9	26.7	Sep92	SPEC news
CDC 4680	R6000	66	512+64/16	40.6	45.1	Sep92	SPEC news
Compaq Deskpro	80487SX	16	0+8	9.3	4.3	Mar93	SPEC news
Compaq Deskpro	80487SX	25	64+8	14.2	6.7	Mar93	SPEC news
Compaq Deskpro	80486DX	33	128+8	18.2	8.3	Sep92	SPEC news
Compaq Deskpro	80486DX2	25/50	256+8	25.7	12.2	Mar93	SPEC news
Compaq Deskpro	80486DX2	33/66	256+8	32.2	16.0	Mar93	SPEC news
Compaq DeskproXL	Pentium	33/66	256+8/8	65.1	63.6	Sep93	SPEC news
Mobius P5-60	Pentium	30/60	?	50.0	46.7	Jan94	c.sun.hw
Nekotech MachI	A21066	??/166	1M+16	70	105	Jun94	m.sale.wk
Nekotech MachI	A21066	??/200	1M+16	105	135	Jun94	m.sale.wk
Nekotech MachII	A21066	??/210	2M+16	130	184	Jun94	m.sale.wk
Nekotech MachII	A21066	??/225	2M+16	135	205	Jun94	m.sale.wk
Nekotech MachII	A21066	??/275	2M+16	170	240	Jun94	m.sale.wk
DEC VAX3100/38	?	?	?	3.5	3.8	Mar93	DECinfo
DEC VAX3100/76	REX520	?	128	7.1	6.6	Mar93	DECinfo
DEC VAX4000VLC	SOC	?	25	5.8	6.3	Mar93	DECinfo
DEC VAX4000/60	KA46	22.2	?	11.1	12.6	Mar93	DECinfo
DEC VAX4000/90	NVAX	71	2/8	?	30.2	Sep92	SPEC news
DEC VAX6000/410	KA660	36	128	?	7.1	Feb90	uproc rpt
DEC VAX6000/510	KA650	62	512	?	13.3	Sep92	SPEC news
DEC VAX6000/610	KA680	83	2M	?	39.2	Sep92	SPEC news
DEC 5000/900	R3000	40	64/64	27.3	29.9	Sep92	SPEC news
DEC 5000/20	R3000	20	64/64	13.5	18.4	Jun93	DECinfo
DEC 5000/25	R3000	25	64/64	15.7	21.7	Jun93	DECinfo
DEC 5000/33	R3000	33	64/128	20.9	23.4	Sep92	SPEC news
DEC 5000/50,150	R4000	50/100	1M+8/8	46.7	45.9	Sep93	c.arch
DEC 5000/120	R3000	20	64/64	13.8	18.4	Jun93	DECinfo
DEC 5000/125	R3000	25	64/64	16.1	21.7	Jun93	DECinfo
DEC 5000/133	R3000	33	64/128	20.9	29.1	Jun93	DECinfo
DEC 5000/200	R3000	25	64/64	19.5	26.7	Jun93	DECinfo
DEC 5000/240	R3000	40	64/64	27.9	35.8	Jun93	DECinfo
DEC 5000/260	R4400	60/120	1M+16/16	57.1	54.5	Sep93	c.arch
DEC 5000/280	R4400	60/120	1M+16/16	56.9	55.6	Jun93	DECinfo
DEC 2000/300	A21064	30/150	512+8/8	80.9	110.2	Oct93	c.arch
DEC 3000/300	A21064	30/150	256+8/8	66.2	91.5	Apr93	c.sun.mc
DEC 3000/300L	A21064	20/100	256+8/8	45.9	63.6	Apr93	c.sun.mc
DEC 3000/300LX	A21064	25/125	256+8/8	63.5	75.5	May94	SPEC news
DEC 3000/300X	A21064	35/175	256+8/8	84.4	100.5	May94	SPEC news
DEC 3000/400	A21064	27/133	512+8/8	74.7	112.2	Apr93	c.arch
DEC 3000/500	A21064	30/150	512+8/8	84.4	127.7	Apr93	c.arch
DEC 3000/500X	A21064	40/200	512+8/8	110.9	164.1	Apr93	c.sun.mc
DEC 3000/600S	A21064	35/175	2M+8/8	114.1	162.1	Oct93	c.arch
DEC 3000/700	A21064A	38/225	2M+16/16	162.6	230.6	Jul94	Digital
DEC 3000/800S	A21064	40/200	2M+8/8	138.4	187.6	May94	c.sun.hw
DEC 3000/900	A21064A	39/275	2M+16/16	189.3	264.1	Jul94	Digital
DEC 4000/610	A21064	40/160	1M+8/8	94.6	137.6	Oct93	Digital
DEC 4000/710	A21064	38/190	4M+8/8	122.3	185.4	Oct93	c.arch
DEC 7000/610	A21064	50/200	4M+8/8	132.6	200.1	Oct93	c.arch
DEC 7000/710	A21064A	39/275	4M+16/16	193.8	292.6	Aug94	Digital
DEC 10000/610	A21064	50/200	4M+8/8	116.5	193.6	Oct93	Digital

DEC 200/4/100	A21064	??/100	512+8/8	74.6	95.2	Feb95	Digital
DEC 250/4/266	A21064A	??/266	2M+16/16	198.6	262.5	Apr95	www.dec
DEC [24]00/4/166	A21064	33/166	512+8/8	116.2	134.8	Jul95	Digital
DEC [24]00/4/233	A21064A	39/233	512+16/16	157.7	183.9	Apr95	Digital
DEC 600/5/266	A21164	38/266	2M+96+8/8	289.0	405.0	Jul95	Digital
DEC 600/5/266	A21164	38/266	4M+96+8/8	292.8	433.5	Jul95	Digital
DEC 600/5/300	A21164	75/300	4M+96+8/8	337.8	502.1	Jul95	Digital
DEC 600/5/333	A21164	83/333	4M+96+8/8	412.4	545.2	Jan96	Digital
DEC 1000/4/200	A21064	40/200	2M+8/8	135.8	177.0	Nov94	Digital
DEC 1000/4/233	A21064	??/233	2M+8/8	165.3	222.9	May95	Digital
DEC 2[01]00/4/200	A21064	47/190	1M+8/8	131.8	161.0	Nov94	Digital
DEC 2[01]00/4/233	A21064A	38/233	1M+16/16	177.3	215.0	Apr95	Digital
DEC 2[01]00/4/275	A21064A	39/275	4M+16/16	202.9	292.6	Apr95	Digital
DEC 2[01]00/5/250	A21164	35/250	4M+96+8/8	277.1	410.4	Apr95	Digital
DEC 2[01]00/5/300	A21164	42/300	4M+96+8/8	319.3	477.3	Feb96	Digital
DEC 8[24]00/5/300	A21164	75/300	4M+96+8/8	341.4	512.9	Apr95	Digital
DEC 8[24]00/5/350	A21164	88/350	4M+96+8/8	432.8	602.2	Feb96	Digital
DG 4100	88100	20	16/16	13.1	?	Sep92	SPEC news
DG 4300	88100	25	16/16	17.4	?	Sep92	SPEC news
DG 4600	88100	33	16/16	22.6	?	Sep92	SPEC news
DG 4605	88100	33	64/32	26.1	?	Sep92	SPEC news
DG 5225	2x88100	25	128/128	20.3	12.1	May93	c.sun.hw
DG 5500	88100	40	128/128	32.3	41.4	Oct93	
DG 6240	4x88100	25	256/256	20.1	?	Sep92	SPEC news
HP 425t	68040	25	4/4	12.3	10.3	Jun93	DECinfo
HP 425e	68040	25	4/4	12.2	9.3	Jun93	DECinfo
HP 705	PA1.1	35	32/64	21.9	33.0	Nov92	Sunflash
HP 710	PA1.1	50	32/64	31.6	47.6	Oct92	c.arch
HP 712/60	PA7100LC	60	64/32+1	67.0	85.3	Jun95	www.hp
HP 712/80i	PA7100LC	80	256/128+1	84.1	79.0	Jan94	HP
HP 712/80	PA7100LC	80	256/128+1	97.1	123.3	Jun95	www.hp
HP 712/100	PA7100LC	100	256/128+1	117.2	144.2	Jun95	www.hp
HP 715/33	PA7100	33	64/64	32.5	52.4	Jan94	HP
HP 715/50	PA7100	50	64/64	49.2	78.8	Jan94	HP
HP 715/75	PA7100	75	256/256	82.6	127.2	Jan94	HP
HP 715/64	PA7100LC	64	256	80.6	109.4	Jun95	www.hp
HP 715/80	PA7100LC	80	256	96.3	123.2	Jun95	www.hp
HP 715/100	PA7100LC	100	256	115.1	138.7	Jun95	www.hp
HP 715/100XC	PA7100LC	100	1MB	132.2	184.6	Jun95	www.hp
HP 720	PA1.1	50	128/256	38.5	66.1	Jun93	DECinfo
HP 725	PA7100	50	64/64	37.1	72.8	Apr93	Sunflash
HP 725/75	PA7100	75	256/256	80.3	126.8	May94	HP
HP 730	PA1.1	66	128/256	47.8	75.4	May92	c.sun.hw
HP 7[35]5	PA7100	99	256/256	109.1	167.9	Jan94	HP
HP 7[35]5/125	PA7150	125	256/256	136	201	Apr94	HP
HP 750	PA1.1	66	256/256	48.1	75.0	Oct92	c.arch
HP C100	PA7200	100	256/256	140	224	Dec95	www.hp
HP C110	PA7200	120	256/256	167	269	Dec95	www.hp
HP F10	PA1.1	32	32/64	22.0	36.6	Mar93	SPEC news
HP [F-I]30	PA1.1	48	256/256	37.8	62.4	Mar93	SPEC news
HP [FH]20	PA1.1	48	64/64	33.6	56.1	Mar93	SPEC news
HP [GHI]30	PA1.1	48	256/256	37.8	62.4	Apr94	www.hp

HP [GHI]40	PA1.1	64	256/256	65.2	91.3	Apr94	www.hp
HP [GHI]50	PA7100	96	256/256	100.0	158.5	Apr94	www.hp
HP [GHI]60	PA7100	96	1M/1M	108.8	195.3	Apr94	www.hp
HP E25	PA7100LC	48	64	45.0	66.6	Mar95	www.hp
HP E35	PA7100LC	64	256	65.6	98.5	Mar95	www.hp
HP E45	PA7100LC	80	256	82.1	122.9	Mar95	www.hp
HP E55	PA7100LC	96	1M	108.0	163.4	Mar95	www.hp
HP J200	PA7200	100	256/256	139.4	222.5	Jun95	www.hp
HP J210	PA7200	120	256/256	168.7	269.2	Jun95	www.hp
HP 807	PA1.1	32	64/32	20.2	?	Sep92	SPEC news
HP 827/17	PA1.1	48	64/64	31.4	?	Sep92	SPEC news
HP 847	PA1.1	?	?	34.8	?	Apr93	DECinfo
HP 867	PA1.1	64	256/256	45.6	?	Sep92	SPEC news
HP 877	PA1.1	64	256/256	45.8	?	Sep92	SPEC news
HP 897S	PA7100	96	?	78.3	141.6	Sep92	SPEC news
IBM N40	MPC601	50	32	41.7	51.0	Mar95	www.ibm
IBM [2M]20	RSC3308	33.3	8	20.4	29.1	Sep93	c.arch
IBM 230	RSC4608	45.5	128+8	28.5	39.9	Sep93	c.arch
IBM 250	MPC601	66	32	62.6	72.2	Jul94	www.ibm
IBM 250	MPC601	80	32	77.6	89.4	Jul94	www.ibm
IBM 25T	MPC601	66	32	62.6	78.8	Mar95	www.ibm
IBM 25T	MPC601	80	32	72.2	90.4	Mar95	www.ibm
IBM C10	MPC601	80	32	78.8	90.4	Jul94	www.ibm
IBM C10	MPC601	80	1M+32	90.5	100.8	Jul94	www.ibm
IBM C20	MPC604	120	16/16	118.2	116.5	Jun95	www.ibm
IBM C20	MPC604	120	1M+16/16	155.0	150.2	Jun95	www.ibm
IBM E20	MPC604	100	512+16/16	139.6	131.6	Oct95	www.ibm
IBM 320H	POWER	25	8/64	20.9	39.4	Nov92	Sunflash
IBM 340	POWER	33	8/32	27.7	51.9	Oct92	c.arch
IBM 350	POWER	41.6	8/32	35.4	74.2	Nov92	Digital
IBM 355	POWER	41.6	32/32	48.1	83.3	Sep93	c.arch
IBM 365,570	POWER	50	32/32	57.5	99.2	Sep93	c.arch
IBM 37[05T]	POWER	62.5	32/32	70.3	121.1	Sep93	c.arch
IBM 380	POWER2	59	32/64	99.3	187.2	Mar95	www.ibm
IBM 390	POWER2	67	1M+32/64	114.3	205.3	Mar95	www.ibm
IBM 39H	POWER2	67	2M+32/64	130.2	266.6	Mar95	www.ibm
IBM 3AT	POWER2	59	32/64	99.3	187.2	Feb95	www.ibm
IBM 3BT	POWER2	67	1M+32/64	114.3	205.3	Feb95	www.ibm
IBM 3CT	POWER2	67	32/64	122.2	244.6	May95	c.bmarks
IBM 3CT	POWER2	67	1M+32/64	129.1	260.7	May95	c.bmarks
IBM 3CT	POWER2	67	2M+32/64	130.2	266.6	Mar95	www.ibm
IBM 40P	MPC601	66	32	63.7	67.8	Mar95	www.ibm
IBM 40P	MPC601	66	256+32	75.1	77.0	Mar95	www.ibm
IBM 41[T W]	MPC601	80	512+32	88.1	98.7	Jul94	www.ibm
IBM 41[T W]	MPC601	80	32	78.8	90.4	Jul94	www.ibm
IBM 42[T W]	MPC604	120	16/16	118.2	116.5	Jun95	www.ibm
IBM 42[T W]	MPC604	120	512+16/16	150.2	146.5	Jun95	www.ibm
IBM 43P	MPC604	100	256+16/16	128.1	120.2	Jun95	www.ibm
IBM 43P	MPC604	120	512+16/16	157.9	139.2	Jun95	www.ibm
IBM 43P	MPC604	133	512+16/16	176.4	156.5	Jun95	www.ibm
IBM 520H	POWER	25	8/32	20.9	39.6	May92	c.sun.hw
IBM 530H	POWER	41.6	8/64	28.5	64.6	Mar93	c.sun.hw

IBM 550	POWER	41.6	8/64	35.4	71.7	May92	c.sun.hw
IBM 560	POWER	50	8/64	42.0	85.6	Oct92	c.arch
IBM [59]80	POWER	62.5	32/64	73.3	134.6	Sep93	c.arch
IBM 580H	POWER2	55	32/256	97.6	203.9	Sep93	c.arch
IBM 590	POWER2	66.6	32/256	121.6	259.7	Jul94	c.bmarks
IBM 59H	POWER2	66.6	1M+32/128	122.4	250.7	Mar95	www.ibm
IBM 591/R21	POWER2	77	32/256	143.5	307.9	Jul95	www.ibm
IBM 970B	POWER	50	32/64	58.8	108.9	Sep93	c.arch
IBM 990	POWER2	71.5	32/256	126.0	260.4	Sep93	c.arch
IBM R24	POWER2	71.5	2M+32/128	134.1	273.8	Jul94	c.bmarks
Mips Magnum	R4000	50/100	16	36.8	40.0	Oct92	c.arch
SGI 4D/25	R3000	20	64/32	14.0	11.1	Jun93	DECinfo
SGI 4D/35	R3000	36	64/64	28.0	33.4	Jun93	DECinfo
SGI Challenge	R4400	50/100	1M+16/16	62.4	66.5	Apr93	c.arch
SGI Onyx	R4400	100/200	4M+16/16	142	143.3	Jul95	SGI
SGI Onyx	R4400	???/250	4M+16/16	177.5	180.2	Nov95	SGI Ptabl
SGI PowerChl,Onyx	R8000	75	4M+16/16	108.7	310.6	Jun94	c.arch
SGI PowerChl,Onyx	R8000	90	4M+16/16	132.2	396.1	Aug95	SGI Ptabl
SGI Crimson	R4000	50/100	1M+8/8	61.7	63.4	Oct92	c.arch
SGI Crimson	R4400	75/150	1M+16/16	86.0	93.2	Nov94	SGI Ptabl
SGI Indigo	R3000	33	32/32	22.4	24.2	Nov92	Sunflash
SGI Indigo2	R4600	66/133	512+16/16	94.8	72.0	Nov94	SGI Ptabl
SGI Indigo2	R4400	75/150	1M+16/16	90	87	Apr93	c.bmarks
SGI Indigo2	R4400	100/200	2M+16/16	140	131	Jul95	SGI
SGI Indigo2	R4400	???/250	2M+16/16	176	165	Jul95	SGI
SGI PowerIndigo2	R8000	75	2M+16/16	113	269	Oct95	www.sgi
SGI IndigoR4000	R4000	50/100	1M+8/8	57.6	60.3	Mar93	c.sun.hw
SGI IndyPC	R4000	50/100	8/8	34	35	Jul93	SGI anno
SGI IndyPC	R4600	50/100	16/16	62.8	49.9	May94	SGI anno
SGI IndyPC	R4600	44/133	16/16	84.9	61.0	Feb95	SGI anno
SGI IndySC	R4600	44/133	512+16/16	113.5	73.7	Feb95	SGI anno
SGI IndySC	R4400	50/150	1M+16/16	91.7	97.5	Nov94	SGI Ptabl
SGI IndySC	R4000	50/100	1M+8/8	59	61	Jul93	SGI anno
SGI IndySC	R4400	44/175	1M+16/16	122.6	115.5	Feb95	SGI anno
SGI IndySC	R4400	50/200	1M+16/16	140.2	131.0	Jan96	SGI
Sun SS/ELC	FJMB86903	33	64	18.2	17.9	Nov92	Sunflash
Sun SS/IPC	FJMB86902	25	64	13.8	11.1	Nov92	Sunflash
Sun SS/IPX	FJMB86903	40	64	21.8	21.5	Nov92	Sunflash
Sun SS2	RT601	40	64	21.8	22.8	Oct92	c.arch
Sun SS2/PowerUp	WeitekPwUP	40/80	16/8	32.2	31.1	Jun93	c.sun.an
Sun SS10/20	SuprSP	33	20/16	39.8	46.6	Nov92	Sunflash
Sun SS10/30	SuprSP	36	20/16	45.2	54.0	Apr93	Cockcroft
Sun SS10/40	SuprSP	40	20/16	50.2	60.2	Apr93	Sunflash
Sun SS10/41	SuprSP	40/40.3	1M+20/16	53.2	67.8	Apr93	Cockcroft
Sun SS10/51	SuprSP	40/50	1M+20/16	65.2	83.0	Apr93	Sunflash
Sun Classic,LX	MicroSP	50	4/2	26.4	21.0	Nov92	Sunflash
Sun Voyager	MicroSP2	60	16/8	43.2	36.2	Mar94	Sun
Sun SS4/70	MicroSP2	70	16/8	59.6	46.8	Jan95	Sunflash
Sun SS4/85	MicroSP2	85	16/8	65.3	53.1	May95	SunIntro
Sun SS5/70	MicroSP2	70	16/8	57.0	47.3	Mar94	Sunflash
Sun SS5/85	MicroSP2	85	16/8	65.3	53.1	May95	SunIntro
Sun SS5/110	MicroSP2	110	16/8	78.6	65.3	May95	SunIntro

Sun SS20/50	SuprSP	50	0/16	76.9	80.1	May95	SunIntro
Sun SS20/51	SuprSP	40/50	1M+20/16	81.8	89.0	May95	SunIntro
Sun SS20/61	SuprSP	50/60	1M+20/16	98.2	107.2	May95	SunIntro
Sun SS20/71	SuprSP2	50/75	1M+20/16	125.8	121.2	Jan95	SunIntro
Sun SS20/612	2xSuprSP	50/60	1M+20/16	?	127.1	Sep94	SPEC news
Sun SS20/HS11	HyperSP	50/100	256+8/0	104.5	127.6	Nov94	SunIntro
Sun SS20/HS21	HyperSP	50/125	256+8/0	131.2	153.0	May95	SunIntro
Sun SS20/151	HyperSP	50/150	512+8/0	169.4	208.2	Nov95	SunWorld
Sun Ultra1/140	UltraSP	71/143	512+16/16	215	303	Nov95	SunIntro
Sun Ultra1/170	UltraSP	83/167	512+16/16	252	351	Nov95	SunIntro
Sun Ultra2/2200	2xUltraSP	67/200	1M+16/16	332	505	Nov95	SunIntro
Sun SS1000	SuprSP	40/50	M+20/16	?	79.9	Jan95	Cockcroft
Sun SS1000	2xSuprSP	40/50	1M+20/16	?	92.3	Jan95	Cockcroft
Sun SS1000	4xSuprSP	40/50	1M+20/16	?	112.8	Jan95	Cockcroft
Sun SS1000	8xSuprSP	40/50	1M+20/16	?	123.1	Jan95	Cockcroft
RT 100S-55	HyperSP	40/55	256+8/0	57	74	Aug94	Ross
RT 100S-66	HyperSP	40/66	256+8/0	67	87	Aug94	Ross
RT 100S-72	HyperSP	40/72	256+8/0	75	96	Aug94	Ross
RT 100S-90	HyperSP	40/90	256+8/0	98	116	Aug95	www.ross
RT 100S-110/1024	HyperSP	40/110	1M+8/0	135	165	Aug95	www.ross
RT 100S-125	HyperSP	40/125	256+8/0	126	146	Aug95	www.ross
RT 200S-66	HyperSP	50/66	256+8/0	72	94	Aug94	Ross
RT 200S-72	HyperSP	50/72	256+8/0	80	105	Aug94	Ross
RT 200S-90	HyperSP	50/90	256+8/0	103	120	Aug95	www.ross
RT 200S-110	HyperSP	50/110	256+8/0	122	142	Apr95	Ross
RT 200S-110/1024	HyperSP	50/110	1M+8/0	137	171	Aug95	www.ross
RT 200S-125	HyperSP	50/125	256+8/0	133	154	Aug95	www.ross
RT 200S-125/512	HyperSP	50/125	512+8/0	152	181	Aug95	www.ross
Solbourne 6/901	SuprSP	33	16+1M+20/16	44.0	52.5	Dec92	SPEC news
HAL 330	SPARC64	100	128/128	181	230	Sep95	www.hal
HAL 350	SPARC64	118	128/128	212	271	Sep95	www.hal
SNI/Pyr PC/E5S	Pentium	30/60	256+8/8	60.6	55.1	Sep93	SPEC news
SNI/Pyr PC/E5S	Pentium	33/66	256+8/8	67.4	61.5	Sep93	SPEC news
SNI/Pyr PC/E5S	Pentium	30/90	256+8/8	86.3	72.7	Jul94	c.bmarks
SNI/Pyr PC/E5S	Pentium	33/100	256+8/8	96.2	81.2	Jul94	c.bmarks
SNI/Pyr PC/D5T	Pentium	30/60	256+8/8	65.9	52.4	Nov94	c.bmarks
SNI/Pyr PC/D5T	Pentium	30/90	256+8/8	86.0	68.3	Nov94	c.bmarks
SNI/Pyr 2-12[05]	R4600	50/100	16/16	76.3	?	Nov94	c.bmarks
SNI/Pyr 4-120	R4400	50/100	16/16	45.6	?	Oct93	c.bmarks
SNI/Pyr 4-120	R4400	50/100	128+16/16	49.7	?	Jan94	c.bmarks
SNI/Pyr 4-220	R4400	50/100	512+16/16	68.2	?	Nov94	c.bmarks
SNI/Pyr 4-3[34]0	R4400	50/100	1M+16/16	71.4	?	Nov94	c.bmarks
SNI/Pyr 4-420	R4400	75/150	512+16/16	92.0	?	Nov94	c.bmarks
SNI/Pyr 4-4[34]0	R4400	75/150	1M+16/16	100.4	?	Nov94	c.bmarks
SNI/Pyr 4-5[34]0	R4400	75/150	4M+16/16	108.7	?	Nov94	c.bmarks
SNI/Pyr 6-120	R4400	50/100	1M+16/16	55.8	?	Nov93	Siemens
SNI/Pyr 6-220	R4400	75/150	4M+16/16	94.2	?	Nov93	Siemens
SNI/Pyr 6-3[24]0	R4400	100/200	4M+16/16	143.7	?	Jun95	SNI/Pyr
SNI/Pyr RM200-C20	R4600	133	16/16	104.6	?	Dec95	c.bmarks
SNI/Pyr RM300-C20	R4600	133	16/16	104.6	?	Dec95	c.bmarks
SNI/Pyr RM300-C60	R4400	100/200	1M+16/16	140.9	?	Dec95	c.bmarks
SNI/Pyr RM400-C70	R4400	100/200	4M+16/16	150.7	?	Dec95	c.bmarks

SNI/Pyr RM1000	R4400	100/200	4M+16/16	152.1	?	Aug95	SNI/Pyr
Dell DimensionXPS	Pentium	60/120	512+8/8	160.7	105.4	Nov95	www.intel
Dell DimensionXPS	Pentium	66/133	512+8/8	177.9	116.0	Nov95	www.intel
Micronics M4P	80486DX4	33/100	256+16	51.4	26.6	Mar94	c.arch
Intel 486DX	80486	50	256+8	30.1	14.0	Oct92	c.arch
Intel 486DX2	80486DX2	33/66	0+8	32.4	16.1	Sep92	uproc rpt
Intel Xpress	Pentium	60	256+8/8	70.4	55.1	Mar95	www.intel
Intel Xpress	Pentium	66	256+8/8	78.0	63.6	Mar95	www.intel
Intel Xpress	Pentium	50/75	512+8/8	89.1	68.5	Mar95	www.intel
Intel Xpress	Pentium	60/90	512+8/8	106.5	81.4	Mar95	www.intel
Intel Xpress	Pentium	60/90	1M+8/8	110.1	84.4	Mar95	www.intel
Intel Xpress	Pentium	66/100	512+8/8	118.1	89.9	Mar95	www.intel
Intel Xpress	Pentium	66/100	1M+8/8	121.9	93.2	Mar95	www.intel
Intel Xpress	Pentium	60/120	512+8/8	133.7	99.5	Mar95	www.intel
Intel Xpress	Pentium	60/120	1M+8/8	140.0	103.9	Mar95	www.intel
Intel Xpress	Pentium	66/133	512+8/8	147.5	109.6	Jun95	www.intel
Intel Xpress	Pentium	66/133	1M+8/8	155.5	116.9	Jun95	www.intel
Intel XXpress	Pentium	66/100	1M+8/8	137.7	?	Jan96	www.intel
Intel XXpress	Pentium	60/120	1M+8/8	157.3	108.4	Jan96	www.intel
Intel XXpress	Pentium	66/133	1M+8/8	174.2	120.6	Jan96	www.intel
Intel XXpress	Pentium	??/150	1M+8/8	181.4	?	Jan96	www.intel
Intel XXpress	Pentium	??/166	1M+8/8	197.5	?	Jan96	www.intel
Intel Alder	PentiumPro	150	256+8/8	243.9	220.0	Jan96	www.intel
Intel Alder	PentiumPro	166	512+8/8	327.1	261.3	Nov95	www.intel
Intel Alder	PentiumPro	180	256+8/8	287.1	254.6	Jan96	www.intel
Intel Alder	PentiumPro	200	256+8/8	318.4	283.2	Jan96	www.intel

***** TABLE 8: Integer/FP SPECrate92 *****

Notes:

- Integer SPECrate is derived from the results of a set of integer benchmarks run multiple times simultaneously, and can be used to estimate a machine's overall multi-tasking throughput for integer code. It is typically used on MP machines
- Floating-Point SPECrate is derived from the results of a set of floating-point benchmarks run multiple times simultaneously, and can be used to estimate a machine's overall multi-tasking throughput for FP code. It is typically used on MP machines.
- Computed specrates are indicated by "c". They're computed from SPECint92, SPECfp92 (for uniprocessors) using a scaling factor. This number is usually slightly less than or equal to a measured specrate on a uniprocessor. The scaling factor is the number of seconds in a week, divided by the time of the longest-running benchmark on the reference SPEC VAX 11/780, which is 604800/25500, or about 23.7.

System Name	CPU (NUMx)Type	ClkMHz ext/in	Cache Ext+I/D	SPECint rate92	SPECfp rate92	Info Date	Source Obtained
DEC VAX11/780	VAX	5	2	24c	24c	Jan89	SPEC Ref
ALR PowerVEISA	80487SX	20	64+8	254c	116c	Mar93	SPEC news
CDC 4330	R3000	33	32/32	591c	567	Sep92	SPEC news
CDC 4360	R3000	33	64/64	591c	633	Sep92	SPEC news
CDC 4680	R6000	66	512+64/16	963c	1070c	Sep92	SPEC news
CDC 4680	2xR6000	66	512+64/16	?	2232	Sep92	SPEC news
Compaq Deskpro	80487SX	16	8	221c	102c	Mar93	SPEC news
Compaq Deskpro	80487SX	25	64+8	337c	159c	Mar93	SPEC news
Compaq Deskpro	80486DX	33	128+8	432c	197c	Sep92	SPEC news
Compaq Deskpro	80486DX2	25/50	256+8	610c	289c	Mar93	SPEC news
Compaq Deskpro	80486DX2	33/66	256+8	764c	379c	Mar93	SPEC news
Compaq DeskproXL	Pentium	33/66	256+8/8	1544c	1508c	Sep93	SPEC news
Mobius P5-60	Pentium	30/60	??+8/8	1186c	1108c	Jan94	c.sun.hw
Convex SPP1000	PA7100	100	1M/1M	?	3478	Sep94	Convex
Convex SPP1000	8xPA7100	100	1M/1M	?	27701	Sep94	Convex
Convex SPP1000	32xPA7100	100	1M/1M	?	95108	Sep94	Convex
Nekotech MachI	A21066	??/166	1M+16	1660c	2490c	Jun94	c.sale.wk
Nekotech MachI	A21066	??/200	1M+16	2490c	3202c	Jun94	c.sale.wk
Nekotech MachII	A21066	??/210	2M+16	3083c	4364c	Jun94	c.sale.wk
Nekotech MachII	A21066	??/225	2M+16	3202c	4862c	Jun94	c.sale.wk
Nekotech MachII	A21066	??/275	2M+16	4032c	5692c	Jun94	c.sale.wk
DEC VAX3100/38	?	?	?	83c	90c	Mar93	DECinfo
DEC VAX3100/76	REX520	?	128	168c	157c	Mar93	DECinfo
DEC VAX4000VLC	SOC	25	?	138c	149c	Mar93	DECinfo
DEC VAX4000/60	KA46	22.2	?	263c	299c	Mar93	DECinfo
DEC VAX4000/90	NVAX	71	2/8	?	716c	Sep92	SPEC news
DEC VAX6000/410	KA660	36	128	?	168	Feb90	uproc rep
DEC VAX6000/510	KA650	62	512	?	315	Sep92	SPEC news
DEC VAX6000/610	KA680	83	2M	?	930	Sep92	SPEC news
DEC 5000/900	R3000	40	64/64	646	709	Sep92	SPEC news
DEC 5000/20	R3000	20	64/64	320c	351	Jun93	DECinfo
DEC 5000/25	R3000	25	64/64	372c	415	Jun93	DECinfo
DEC 5000/33	R3000	33	64/128	496c	556c	Sep92	SPEC news
DEC 5000/50,150	R4000	50/100	1M+8/8	1107c	1088c	Sep93	c.arch
DEC 5000/120	R3000	20	64/64	327c	436c	Jun93	DECinfo
DEC 5000/125	R3000	25	64/64	382c	514c	Jun93	DECinfo
DEC 5000/133	R3000	33	64/128	495c	690c	Jun93	DECinfo
DEC 5000/200	R3000	25	64/64	462c	633c	Jun93	DECinfo
DEC 5000/240	R3000	40	64/64	661c	848c	Jun93	DECinfo
DEC 5000/260	R4400	60/120	1M+16/16	1353c	1292c	Sep93	c.arch
DEC 5000/280	R4400	60/120	1M+16/16	1349c	1318c	Jun93	DECinfo
DEC 2000/300	A21064	30/150	512+8/8	1930	2634	Oct93	c.arch
DEC 3000/300	A21064	30/150	256+8/8	1535	2137	Apr93	c.sun.mc
DEC 3000/300L	A21064	20/100	256+8/8	1081	1480	Apr93	c.sun.mc
DEC 3000/300LX	A21064	25/125	256+8/8	1506c	1791c	May94	SPEC news
DEC 3000/300X	A21064	35/175	256+8/8	2002c	2384c	May94	SPEC news
DEC 3000/400	A21064	27/133	512+8/8	1763	2662	Apr93	c.arch
DEC 3000/500	A21064	30/150	512+8/8	1997	3023	Apr93	c.arch

DEC 3000/500X	A21064	40/200	512+8/8	2611	3910	Apr93	c.sun.mc
DEC 3000/600S	A21064	35/175	2M+8/8	2722	3857	Oct93	c.arch
DEC 3000/700	A21064A	38/225	2M+16/16	3944	5482	Jul94	Digital
DEC 3000/800S	A21064	40/200	2M+8/8	3137	4377	Oct93	c.arch
DEC 3000/900	A21064A	39/275	2M+16/16	4702	6293	Jul94	Digital
DEC 4000/610	A21064	40/160	1M+8/8	2198	3247	Oct93	Digital
DEC 4000/620	2xA21064	40/160	1M+8/8	3861	6215	May93	sunflash
DEC 4000/710	A21064	38/190	4M+8/8	2900	4340	Oct93	c.arch
DEC 4000/720	2xA21064	38/190	4M+8/8	5144	8272	Apr94	DECinfo
DEC 7000/610	A21064	50/200	4M+8/8	3250	4701	Apr94	DECinfo
DEC 7000/620	2xA21064	50/200	4M+8/8	6347	9329	Apr94	DECinfo
DEC 7000/640	4xA21064	50/200	4M+8/8	12463	18719	Apr94	DECinfo
DEC 7000/660	6xA21064	50/200	4M+8/8	18956	28157	Apr94	DECinfo
DEC 7000/710	A21064A	39/275	4M+16/16	4522	6680	Aug94	Digital
DEC 7000/720	2xA21064A	39/275	4M+16/16	8621	13395	Aug94	Digital
DEC 7000/740	4xA21064A	39/275	4M+16/16	17450	27008	Aug94	Digital
DEC 7000/760	6xA21064A	39/275	4M+16/16	24735	40103	Aug94	Digital
DEC 10000/610	A21064	50/200	4M+8/8	2761c	4588c	Oct93	Digital
DEC 10000/660	6x21064	50/200	4M+8/8	12865	24748	Nov92	c.arch
DEC 200/4/100	A21064	??/100	512+8/8	1749	2258	Feb95	Digital
DEC [24]00/4/166	A21064	33/166	512+8/8	2779	3160	Apr95	Digital
DEC [24]00/4/233	A21064A	39/233	512+16/16	3772	4415	Apr95	Digital
DEC 250/4/266	A21064A	??/266	2M+16/16	4574	6189	Apr95	www.dec
DEC 600/5/266	A21164	38/266	2M+96+8/8	7001	9741	Jul95	Digital
DEC 600/5/266	A21164	38/266	4M+96+8/8	7132	10247	Jul95	Digital
DEC 600/5/300	A21164	42/300	4M+96+8/8	8384	11812	Jul95	Digital
DEC 1000/4/200	A21064	40/200	2M+8/8	3136	4230	Nov94	Digital
DEC 1000/4/233	A21064	??/233	2M+8/8	3921c	5287c	May95	Digital
DEC 2[01]00/4/200	A21064	47/190	1M+8/8	3123	3835	Nov94	Digital
DEC 2[01]00/4/200	2xA21064	47/190	1M+8/8	6178	7296	Nov94	Digital
DEC 2[01]00/4/233	A21064A	38/233	1M+16/16	4135	5112	Apr95	Digital
DEC 2[01]00/4/233	2xA21064A	38/233	1M+16/16	8284	9676	Apr95	Digital
DEC 2100/4/233	4xA21064A	38/233	1M+16/16	15538	17361	Apr95	Digital
DEC 2[01]00/4/275	A21064A	39/275	4M+16/16	4711	6827	pr95	Digital
DEC 2[01]00/4/275	2xA21064A	39/275	4M+16/16	9423	13242	Apr95	Digital
DEC 2100/4/275	4xA21064A	39/275	4M+16/16	18036	25997	Apr95	Digital
DEC 2[01]00/5/250	A21164	35/250	4M+96+8/8	6551	9795	Apr95	Digital
DEC 2[01]00/5/250	2xA21164	35/250	4M+96+8/8	13112	18802	Apr95	Digital
DEC 2100/5/250	4xA21164	35/250	4M+96+8/8	24996	37928	Apr95	Digital
DEC 2[01]00/5/300	A21164	42/300	4M+96+8/8	7148	10125	Feb96	Digital
DEC 2[01]00/5/300	2xA21164	42/300	4M+96+8/8	12559	19665	Feb96	Digital
DEC 2100/5/300	4xA21164	42/300	4M+96+8/8	22202	39198	Feb96	Digital
DEC 8[24]00/5/300	A21164	75/300	4M+96+8/8	8551	11981	Apr95	Digital
DEC 8[24]00/5/300	2xA21164	75/300	4M+96+8/8	16769	24329	Apr95	Digital
DEC 8[24]00/5/300	4xA21164	75/300	4M+96+8/8	33201	48526	Apr95	Digital
DEC 8[24]00/5/300	6xA21164	75/300	4M+96+8/8	50778	71286	Apr95	Digital
DEC 8400/5/300	8xA21164	75/300	4M+96+8/8	63418	94686	Apr95	Digital
DEC 8400/5/300	10xA21164	75/300	4M+96+8/8	80707	117493	Apr95	Digital
DEC 8400/5/300	12xA21164	75/300	4M+96+8/8	91580	140571	Apr95	Digital
DEC 8[24]00/5/350	A21164	88/350	4M+96+8/8	9908	14309	Feb96	Digital
DEC 8[24]00/5/350	6xA21164	88/350	4M+96+8/8	65842	84561	Feb96	Digital
DEC 8400/5/350	12xA21164	88/350	4M+96+8/8	115878	168159	Feb96	Digital

DG 4100	88100	20	16/16	310c	?	Sep92	SPEC news
DG 4300	88100	25	16/16	412c	?	Sep92	SPEC news
DG 4600	88100	33	16/16	536c	?	Sep92	SPEC news
DG 4605	88100	33	64/32	619c	?	Sep92	SPEC news
DG 5225	2x88100	25	128/128	868	532	May93	c.sun.hw
DG 5500	88100	40	128/128	766c	981c	Oct93	
DG 5240	4x88100	25	64/64	1591	971	May93	c.sun.hw
DG 6240	4x88100	25	256/256	1591	?	Sep92	SPEC news
DG 6280	8x88100	25	512+64/64	3245	?	Sep92	SPEC news
HP 425t	68040	25	4/4	292c	244c	Jun93	DECinfo
HP 425e	68040	25	4/4	289c	220c	Jun93	DECinfo
HP 705	PA1.1	35	32/64	519c	782c	Nov92	Sunflash
HP 710	PA1.1	50	32/64	749c	1128c	Oct92	c.arch
HP 712/60	PA7100LC	60	64/32+1	1589c	2023c	Jun95	www.hp
HP 712/80i	PA7100LC	80	256/128+1	1995c	1874c	Jan94	HP
HP 712/80	PA7100LC	80	256/128+1	2303c	2924c	Jun95	www.hp
HP 712/100	PA7100LC	100	256/128+1	2780c	3420c	Jun95	www.hp
HP 715/33	PA7100	33	64/64	574c	1067c	Mar93	c.sun.hw
HP 715/50	PA7100	50	64/64	866	1710	Apr93	Sunflash
HP 715/75	PA7100	75	256/256	1959c	3017c	Jan94	HP
HP 715/64	PA7100LC	64	256	1498	2281	Aug94	www.hp
HP 715/80	PA7100LC	80	256	1866	2865	Aug94	www.hp
HP 715/100	PA7100LC	100	256	2237	3226	Aug94	www.hp
HP 715/100XC	PA7100LC	100	1MB	3135c	4378c	Jun95	www.hp
HP 720	PA1.1	50	128/256	912	1567c	Jun93	DECinfo
HP 725	PA7100	50	64/64	866	1710	Apr93	Sunflash
HP 725/75	PA7100	75	256/256	1905c	3007c	May94	HP
HP 730	PA1.1	66	128/256	1133c	1787c	May92	c.sun.hw
HP 7[35]5	PA7100	99	56/256	1832	2950	Nov92	c.arch
HP 7[35]5/125	PA7150	125	256/256	3226c	4767c	Apr94	HP
HP 750	PA1.1	66	256/256	1141	1778c	Oct92	c.arch
HP F10	PA1.1	32	32/64	521c	867c	Mar93	SPEC news
HP [F-I]30	PA1.1	48	256/256	896c	1479c	Mar93	SPEC news
HP [FH]20	PA1.1	48	64/64	796c	1330c	Mar93	SPEC news
HP [GHI]30	PA1.1	48	256/256	869c	1435c	Apr94	www.hp
HP [GHI]40	PA1.1	64	256/256	1500c	2100c	Apr94	www.hp
HP [GHI]50	PA7100	96	256/256	2300c	3646c	Apr94	www.hp
HP [GHI]60	PA7100	96	1M/1M	2502c	4492c	Apr94	www.hp
HP E25	PA7100LC	48	64	1067c	1580c	Mar95	www.hp
HP E35	PA7100LC	64	256	1556c	2336c	Mar95	www.hp
HP E45	PA7100LC	80	256	1947c	2915c	Mar95	www.hp
HP E55	PA7100LC	96	1M	2562c	3875c	Mar95	www.hp
HP J200	PA7200	100	256/256	3306c	5277c	Jun95	www.hp
HP J200	2xPA7200	100	256/256	6432	9646	Jun95	www.hp
HP J210	PA7200	120	256/256	4001c	6385c	Jun95	www.hp
HP J210	2xPA7200	120	256/256	7892	11900	Jun95	www.hp
HP 827/17	PA1.1	48	64/64	744c	?	Sep92	SPEC news
HP 847	PA1.1	?	?	825c	?	Apr93	DECinfo
HP 867	PA1.1	64	256/256	1201	?	Sep92	SPEC news
HP 870	2xPA1.1	50	512/512	1515	?	Sep92	SPEC news
HP 870	3xPA1.1	50	512/512	2051	?	Sep92	SPEC news
HP 870	4xPA1.1	50	512/512	2479	?	Sep92	SPEC news

HP 877	PA1.1	64	256/256	1085c	?	Sep92	SPEC news
HP 897S	PA7100	96	?	1857	1937	Sep92	SPEC news
HP 890	PA1.1	60	2M/2M	1215	1180	Sep92	SPEC news
HP 890	2xPA1.1	60	2M/2M	2253	2360	Sep92	SPEC news
HP 890	3xPA1.1	60	2M/2M	3306	3529	Sep92	SPEC news
HP 890	4xPA1.1	60	2M/2M	4301	4685	Sep92	SPEC news
HP T500	4xPA7100	90	1M/1M	9017	15341	Jan94	HP
HP T500	8xPA7100	90	1M/1M	17114	28341	Jan94	HP
HP T500	12xPA7100	90	1M/1M	23717	38780	Jan94	HP
IBM N40	MPC601	50	32	989c	1210c	Mar95	www.ibm
IBM [2M]20	RSC3308	33.3	8	377	543	ep93	c.arch
IBM 230	RSC4608	45.5	128+8	675c	946c	Sep93	c.arch
IBM 250	MPC601	66	32	1485c	1712c	Jul94	www.ibm
IBM 250	MPC601	80	32	1840c	2120c	Jul94	www.ibm
IBM 25T	MPC601	66	32	1485c	1869c	Mar95	www.ibm
IBM 25T	MPC601	80	32	1712c	2144c	Mar95	www.ibm
IBM C10	MPC601	80	32	1869c	2144c	Jul94	www.ibm
IBM C10	MPC601	80	1M+32	2146c	2391c	Jul94	www.ibm
IBM C20	MPC604	120	16/16	2803c	2763c	Jun95	www.ibm
IBM C20	MPC604	120	1M+16/16	3676c	3562c	Jun95	www.ibm
IBM E20	MPC604	100	512+16/16	3311c	3121c	Oct95	www.ibm
IBM 320H	POWER	25	8/64	496	935	Nov92	Sunflash
IBM 340	POWER	33	8/32	657	1231	Oct92	c.arch
IBM 350	POWER	41.6	8/32	821	1542	Nov92	DEC anno
IBM 355	POWER	41.6	32/32	961	1936	Sep93	c.arch
IBM 365,570	POWER	50	32/32	1148	2301	Sep93	c.arch
IBM 37[05T]	POWER	62.5	32/32	1332	2612	Sep93	c.arch
IBM 380	POWER2	59	32/64	2355c	4440c	Mar95	www.ibm
IBM 390	POWER2	67	1M+32/64	2711c	4869c	Mar95	www.ibm
IBM 39H	POWER2	67	2M+32/64	3088c	6323c	Mar95	www.ibm
IBM 3AT	POWER2	59	32/64	2355c	4440c	Feb95	www.ibm
IBM 3BT	POWER2	67	1M+32/64	2711c	4869c	Feb95	www.ibm
IBM 3CT	POWER2	67	32/64	2898c	5801c	May95	c.bmarks
IBM 3CT	POWER2	67	1M+32/64	3062c	6183c	May95	c.bmarks
IBM 3CT	POWER2	67	2M+32/64	3088c	6323c	Mar95	www.ibm
IBM 40P	MPC601	66	32	1511c	1608c	Mar95	www.ibm
IBM 40P	MPC601	66	256+32	1781c	1826c	Mar95	www.ibm
IBM 41[T W]	MPC601	80	512+32	2090c	2341c	Jul94	www.ibm
IBM 41[T W]	MPC601	80	32	1898c	2144c	Jul94	www.ibm
IBM 42[T W]	MPC604	120	16/16	2803c	2763c	Jun95	www.ibm
IBM 42[T W]	MPC604	120	512+16/16	3562c	3475c	Jun95	www.ibm
IBM 43P	MPC604	100	256+16/16	3038c	2851c	Jun95	www.ibm
IBM 43P	MPC604	120	512+16/16	3745c	3301c	Jun95	www.ibm
IBM 43P	MPC604	133	512+16/16	4184c	3712c	Jun95	www.ibm
IBM 520H	POWER	25	8/32	495c	939	May92	c.sun.hw
IBM 530H	POWER	41.6	8/64	669	1364	Mar93	c.sun.hw
IBM 550	POWER	41.6	8/64	840	1701	May92	c.sun.hw
IBM 560	POWER	50	8/64	999	2028	Oct92	c.arch
IBM [59]80	POWER	62.5	32/64	1404	2960	Sep93	c.arch
IBM 580H	POWER2	55	32/256	2313c	4832c	Sep93	c.arch
IBM 590	POWER2	66.6	32/256	2884c	6159c	Jul94	c.bmarks
IBM 59H	POWER2	66.6	1M+32/128	2903c	5946c	Mar95	www.ibm

IBM 591/R21	POWER2	77	32/256	3403c	7303c	Jul95	www.ibm
IBM 970B	POWER	50	32/64	1117	2220	Sep93	c.arch
IBM 990	POWER2	71.5	32/256	2986c	6171c	Sep93	c.arch
IBM R24	POWER2	71.5	2M+32/128	3181c	6494c	Jul94	c.bmarks
IBM R30	2xMPC601	75	1M+32	4267	4492	Jul95	www.ibm
IBM R30	4xMPC601	75	1M+32	8430	8689	Jul95	www.ibm
IBM R30	8xMPC601	75	1M+32	16200	16324	Jul95	www.ibm
Mips Magnum	R4000	50/100	8/8	872c	948c	Oct92	c.arch
SGI 4D/25	R3000	20	64/32	332c	263c	Jun93	DECinfo
SGI 4D/35	R3000	36	64/64	664c	792c	Jun93	DECinfo
SGI Challenge	R4400	50/100	1M+16/16	479c	1576c	Apr93	c.arch
SGI Onyx	R4400	100/200	4M+16/16	3368c	3399c	Jul95	SGI
SGI PowerChl, Onyx	R8000	75	4M+16/16	2578c	7367c	Jun94	c.arch
SGI PowerChl, Onyx	R8000	90	4M+16/16	3135c	9395c	Aug95	SGI Ptabl
SGI Crimson	R4000	50/100	1M+8/8	1383	1459	Oct92	c.arch
SGI Crimson	R4400	75/150	1M+16/16	2040c	2210c	Nov94	SGI Ptabl
SGI Indigo	R3000	33	32/32	531	574	Nov92	Sunflash
SGI Indigo2	R4600	66/133	512+16/16	2248c	1708c	Nov94	SGI Ptabl
SGI Indigo2	R4400	75/150	1M+16/16	2133c	2062c	Apr93	c.bmarks
SGI Indigo2	R4400	100/200	2M+16/16	3320c	3107c	Jul95	SGI
SGI Indigo2	R4400	???/250	2M+16/16	4174c	3913c	Jul95	SGI
SGI PowerIndigo2	R8000	75	2M+16/16	2680c	6380c	Oct95	www.sgi
SGI IndigoR4000	R4000	50/100	1M+8/8	1366	1430	Mar93	c.sun.hw
SGI IndyPC	R4000	50/100	8/8	806c	830c	Jul93	SGI anno
SGI IndyPC	R4600	50/100	16/16	1489c	1184c	May94	SGI anno
SGI IndyPC	R4600	44/133	16/16	2014c	1447c	Feb95	SGI anno
SGI IndySC	R4600	44/133	512+16/16	2692c	1748c	Feb95	SGI anno
SGI IndySC	R4400	50/150	1M+16/16	2175c	2312c	Nov94	SGI Ptabl
SGI IndySC	R4000	50/100	1M+8/8	1398c	1446c	Jul93	SGI anno
SGI IndySC	R4400	44/175	1M+16/16	2908c	2739c	Feb95	SGI anno
SGI IndySC	R4400	50/200	1M+16/16	3325c	3107c	Jan96	SGI
SGI Challenge/L	12xR4400	50/100	1M+16/16	13406	17370	May93	c.sun.hw
SGI Challenge/XL	R4400	75/150	1M+16/16	2221	2306	Oct93	Mashey
SGI Challenge/XL	4xR4400	75/150	1M+16/16	8679	9079	Oct93	Mashey
SGI Challenge/XL	8xR4400	75/150	1M+16/16	16849	17854	Oct93	Mashey
SGI Challenge/XL	12xR4400	75/150	1M+16/16	23696	25171	Oct93	Mashey
SGI Challenge/XL	16xR4400	75/150	1M+16/16	27242	33956	Oct93	Mashey
SGI Challenge/XL	20xR4400	75/150	1M+16/16	31073	40013	Oct93	Mashey
SGI Challenge/XL	24xR4400	75/150	1M+16/16	?	45776	Oct93	Mashey
SGI Challenge/XL	28xR4400	75/150	1M+16/16	?	53796	Oct93	Mashey
SGI Challenge/XL	32xR4400	75/150	1M+16/16	?	56840	Oct93	Mashey
SGI Challenge/XL	28xR4400	100/200	4M+16/16	65793	94218	Jan96	SGI
SGI PowerChl/XL	16xR8000	90	4M+16/16	47131	148900	Jan96	SGI
Sun SS/ELC	FJMB86903	33	64	432	425	Nov92	Sunflash
Sun SS/IPC	FJMB86902	25	64	327	263	Nov92	Sunflash
Sun SS/IPX	FJMB86903	40	64	517	510	Nov92	Sunflash
Sun SS2	RT601	40	64	517	541	Oct92	c.arch
Sun SS2/PowerUp	WeitekPwUP	40/80	16/8	763c	737c	Jun93	c.sun.an
Sun SS10/20	SuprSP	33	20/16	943c	1104c	Nov92	Sunflash
Sun SS10/30	SuprSP	36	20/16	1072	1282	Apr93	Cockcroft
Sun SS10/40	SuprSP	40	20/16	1191	1427	Apr93	Sunflash
Sun SS10/402	2xSuprSP	40	20/16	2112	2378	Apr93	Sunflash

Sun SS10/41	SuprSP	40/40.3	1M+20/16	1264	1607	Apr93	Cockcroft
Sun SS10/412	2xSuprSP	40/40.3	1M+20/16	2411	2854	Apr93	Cockcroft
Sun SS10/51	SuprSP	40/50	1M+20/16	1546c	1969c	Apr93	Sunflash
Sun SS10/512	2xSuprSP	40/50	1M+20/16	2950	3744	Apr93	Sunflash
Sun SS10/514	4xSuprSP	40/50	1M+20/16	5155	5809	Dec93	Sun
Sun Classic,LX	MicroSP	50	4/2	626	498	Nov92	Sunflash
Sun Voyager	MicroSP2	60	16/8	1025c	859c	Mar94	Sun
Sun SS4/70	MicroSP2	70	16/8	1414c	1110c	Jan95	Sunflash
Sun SS4/85	MicroSP2	85	16/8	1549c	1259c	May95	SunIntro
Sun SS5/70	MicroSP2	70	16/8	1352c	1122c	Mar94	Sunflash
Sun SS5/85	MicroSP2	85	16/8	1549c	1259c	May95	SunIntro
Sun SS5/110	MicroSP2	110	16/8	1864c	1549c	May95	SunIntro
Sun SS20/50	SuprSP	50	20/16	1628	1842	Mar94	Sunflash
Sun SS20/51	SuprSP	40/50	1M+20/16	1731	1995	Mar94	Sunflash
Sun SS20/61	SuprSP	50/60	1M+20/16	2092	2418	Mar94	Sunflash
Sun SS20/502	2xSuprSP	50	0/16	3218	3193	May95	SunIntro
Sun SS20/612	2xSuprSP	50/60	1M+20/16	4492	4888	May95	SunIntro
Sun SS20/712	2xSuprSP2	50/75	1M+20/16	5726	5439	Jan95	SunIntro
Sun SS20/514	4xSuprSP	40/50	1M+20/16	7072	7341	May95	SunIntro
Sun SS20/HS11	HyperSP	50/100	256+8/0	2478c	3026c	Nov94	SunIntro
Sun SS20/HS14	4xHyperSP	50/100	256+8/0	8124	8906	May95	SunIntro
Sun SS20/HS21	HyperSP	50/125	256+8/0	3112c	3629c	May95	SunIntro
Sun SS20/HS22	2xHyperSP	50/125	256+8/0	5600	6399	May95	SunIntro
Sun SS20/151	HyperSP	50/150	512+8/0	018c	4938c	Nov95	SunWorld
Sun SS20/152	2xHyperSP	50/150	512+8/0	7310	8758	Nov95	SunWorld
Sun 600-120	2xRT605	40	64	043	1066	Sep92	SPEC news
Sun 600-140	4xRT605	40	64	1847	1930	Sep92	SPEC news
Sun Ultra1/140	UltraSP	71/143	512+16/16	5107	7175	Nov95	SunIntro
Sun Ultra1/170	UltraSP	83/167	512+16/16	5982	8323	Nov95	SunIntro
Sun Ultra2/2200	2xUltraSP	67/200	1M+16/16	14962	18675	Nov95	SunIntro
Sun SS1000	2xSuprSP	40/50	1M+20/16	2730	3681	May93	c.sun.hw
Sun SS1000	4xSuprSP	40/50	1M+20/16	5318	7076	May93	c.sun.hw
Sun SS1000	8xSuprSP	40/50	1M+20/16	10113	12710	May93	c.sun.hw
Sun SS1000E	2xSuprSP	50/60	1M+20/16	3999	4584	Oct94	Sunflash
Sun SS1000E	8xSuprSP	50/60	1M+20/16	15414	17113	Oct94	Sunflash
Sun SS1000E	8xSuprSP2	50/85	1M+20/16	21758	20851	Oct95	Sunflash
Sun SC2000	8xSuprSP	40/40.3	1M+20/16	8047	10600	May93	c.sun.hw
Sun SC2000	16xSuprSP	40/50	2M+20/16	21196	28064	Oct93	sunflash
Sun SC2000E	2xSuprSP	50/60	2M+20/16	4282	4952	Oct94	sunflash
Sun SC2000E	20xSuprSP	50/60	2M+20/16	38213	44722	Oct94	sunflash
Sun SC2000E	20xSuprSP2	50/85	2M+20/16	57997	54206	Oct95	sunflash
Cray CS6400	24xSuprSP	55/60	2M+20/16	41967	55734	Mar94	c.sun.hw
Cray CS6400	32xSuprSP	55/60	2M+20/16	54186	72177	Mar94	c.sun.hw
Cray CS6400	48xSuprSP	55/60	2M+20/16	82522	102235	Nov94	Cray
Cray CS6400	56xSuprSP	55/60	2M+20/16	95262	115802	Nov94	Cray
Cray CS6400	64xSuprSP	55/60	2M+20/16	101969	129843	Nov94	Cray
RT 100S-55	HyperSP	40/55	256+8/0	1352c	1755c	Aug94	Ross
RT 100D-55	2xHyperSP	40/55	256+8/0	2368	2838	Aug94	Ross
RT 100Q-55	4xHyperSP	40/55	256+8/0	4554	5457	Aug94	Ross
RT 100S-66	HyperSP	40/66	256+8/0	1589c	2063c	Aug94	Ross
RT 100D-66	2xHyperSP	40/66	256+8/0	2817	3377	Aug94	Ross
RT 100Q-66	4xHyperSP	40/66	256+8/0	5419	6470	Aug94	Ross

RT 100S-72	HyperSP	40/72	256+8/0	1779c	2277c	Aug94	Ross
RT 100D-72	2xHyperSP	40/72	256+8/0	3073	3684	Aug94	Ross
RT 100Q-72	4xHyperSP	40/72	256+8/0	5912	7058	Aug94	Ross
RT 100S-90	HyperSP	40/90	256+8/0	2324c	2751c	Aug95	www.ross
RT 100D-90	2xHyperSP	40/90	256+8/0	4264	4747	Aug95	www.ross
RT 100Q-90	4xHyperSP	40/90	256+8/0	7142	7310	Aug95	www.ross
RT 100S-110/1024	HyperSP	40/110	1M+8/0	3202c	3913c	Aug95	www.ross
RT 100D-110/1024	2xHyperSP	40/110	1M+8/0	6049	173	Aug95	www.ross
RT 100Q-110/1024	4xHyperSP	40/110	1M+8/0	10586	11477	Aug95	www.ross
RT 100S-125	HyperSP	40/125	256+8/0	2988c	3463c	Aug95	www.ross
RT 100D-125	2xHyperSP	40/125	256+8/0	5437	5848	Aug95	www.ross
RT 100Q-125	4xHyperSP	40/125	256+8/0	8882	8933	Aug95	www.ross
RT 200S-66	HyperSP	50/66	256+8/0	1708c	2229c	Aug94	Ross
RT 200D-66	2xHyperSP	50/66	256+8/0	3042	3647	Aug94	Ross
RT 200Q-66	4xHyperSP	50/66	256+8/0	5853	6988	Aug94	Ross
RT 200S-72	HyperSP	50/72	256+8/0	1897c	2490c	Aug94	Ross
RT 200D-72	2xHyperSP	50/72	256+8/0	3318	3979	Aug94	Ross
RT 200Q-72	4xHyperSP	50/72	256+8/0	6385	7623	Aug94	Ross
RT 200S-90	HyperSP	50/90	256+8/0	2395c	2846c	Apr95	Ross
RT 200D-90	2xHyperSP	50/90	256+8/0	4568	5226	Aug95	www.ross
RT 200Q-90	4xHyperSP	50/90	256+8/0	7785	8107	Aug95	www.ross
RT 200Q-100	4xHyperSP	50/100	256+8/0	9132	11389	Apr95	SunExpert
RT 200S-110	HyperSP	50/110	256+8/0	2894c	3368c	Apr95	Ross
RT 200Q-110	4xHyperSP	50/110	256+8/0	9988	12026	Apr95	Ross
RT 200S-110/1024	HyperSP	50/110	1M+8/0	3249c	4056c	Aug95	www.ross
RT 200D-110/1024	2xHyperSP	50/110	1M+8/0	6185	7697	Aug95	www.ross
RT 200Q-110/1024	4xHyperSP	50/110	1M+8/0	11133	13085	Aug95	www.ross
RT 200S-125	HyperSP	50/125	256+8/0	3154c	3653c	Aug95	www.ross
RT 200D-125	2xHyperSP	50/125	256+8/0	5857	6510	Aug95	www.ross
RT 200Q-125	4xHyperSP	50/125	256+8/0	9539	9726	Aug95	www.ross
RT 200S-125/512	HyperSP	50/125	512+8/0	3605c	4293c	Aug95	www.ross
RT 200D-125/512	2xHyperSP	50/125	512+8/0	6717	7805	Aug95	www.ross
RT 200Q-125/512	4xHyperSP	50/125	512+8/0	11311	12507	Aug95	www.ross
Solbourne 6/901	SuprSP	33	16M+1M+20/1	1043c	1244c	Dec92	SPEC news
HAL 330	SPARC64	100	128/128	4163c	5290c	Sep95	www.hal
HAL 350	SPARC64	118	128/128	4876c	6233c	Sep95	www.hal
Marix DTH802	2xHyperSP	??/80	256+8/0	3684	4613	Jan95	Marix
Marix DSH904	4xHyperSP	??/90	256+8/0	7972	8842	Jan95	Marix
SNI/Pyr PC/E5S	Pentium	30/60	256+8/8	1436c	1306c	Sep93	SPEC news
SNI/Pyr PC/E5S	Pentium	33/66	256+8/8	1597c	1458c	Sep93	SPEC news
SNI/Pyr PC/E5S	Pentium	30/90	256+8/8	2047c	1724c	Jul94	c.bmarks
SNI/Pyr PC/E5S	Pentium	33/100	256+8/8	2282c	1926c	Jul94	c.bmarks
SNI/Pyr PC/D5T	Pentium	30/60	256+8/8	1516c	1205c	Nov94	c.bmarks
SNI/Pyr PC/D5T	Pentium	30/90	256+8/8	1978c	1571c	Nov94	c.bmarks
SNI/Pyr 2-12[05]	R4600	50/100	16/16	1755c	?	Nov94	c.bmarks
SNI/Pyr 4-120	R4400	50/100	16/16	1081	?	Oct93	c.bmarks
SNI/Pyr 4-120	R4400	50/100	128+16/16	1177	?	Jan94	c.bmarks
SNI/Pyr 4-220	R4400	50/100	512+16/16	1569c	?	Nov94	c.bmarks
SNI/Pyr 4-3[34]0	R4400	50/100	1M+16/16	1642c	?	Nov94	c.bmarks
SNI/Pyr 4-4[34]0	R4400	75/150	1M+16/16	2309c	?	Nov94	c.bmarks
SNI/Pyr 4-5[34]0	R4400	75/150	4M+16/16	2500c	?	Nov94	c.bmarks
SNI/Pyr 6-120	R4400	50/100	1M+16/16	1293	?	Nov93	Siemens

SNI/PyT 6-120	2xR4400	50/100	1M+16/16	2486	?	Nov93	Siemens
SNI/PyT 6-120	3xR4400	50/100	1M+16/16	3549	?	Nov93	Siemens
SNI/PyT 6-120	4xR4400	50/100	1M+16/16	4798	?	Nov93	Siemens
SNI/PyT 6-140	8xR4400	50/100	1M+16/16	9352	?	Jan94	c.bmarks
SNI/PyT 6-220	R4400	75/150	4M+16/16	2193	?	Nov93	Siemens
SNI/PyT 6-220	2xR4400	75/150	4M+16/16	4196	?	Nov93	Siemens
SNI/PyT 6-220	3xR4400	75/150	4M+16/16	6218	?	Nov93	Siemens
SNI/PyT 6-220	4xR4400	75/150	4M+16/16	8073	?	Nov93	Siemens
SNI/PyT 6-240	8xR4400	75/150	4M+16/16	15197	?	Jan94	c.bmarks
SNI/PyT 6-5[34]0	12xR4400	75/150	4M+16/16	24759	?	Nov94	c.bmarks
SNI/PyT 6-5[34]0	16xR4400	75/150	4M+16/16	31803	?	Nov94	c.bmarks
SNI/PyT 6-5[34]0	20xR4400	75/150	4M+16/16	36968	?	Nov94	c.bmarks
SNI/PyT 6-5[34]0	24xR4400	75/150	4M+16/16	42536	?	Nov94	c.bmarks
SNI/PyT 6-3[24]0	R4400	100/200	4M+16/16	3470	?	Jun95	SNI/PyT
SNI/PyT 6-3[24]0	2xR4400	100/200	4M+16/16	6786	?	Jun95	SNI/PyT
SNI/PyT 6-3[24]0	4xR4400	100/200	4M+16/16	13094	?	Jun95	SNI/PyT
SNI/PyT 6-3[24]0	8xR4400	100/200	4M+16/16	24242	?	Jun95	SNI/PyT
SNI/PyT 6-620	12xR4400	100/200	4M+16/16	36562	?	Jun95	SNI/PyT
SNI/PyT 6-620	16xR4400	100/200	4M+16/16	47422	?	Jun95	SNI/PyT
SNI/PyT 6-620	24xR4400	100/200	4M+16/16	69361	?	Jun95	SNI/PyT
SNI/PyT RM200-C20	R4600	133	16/16	2499	?	Dec95	c.bmarks
SNI/PyT RM300-C20	R4600	133	16/16	2499	?	Dec95	c.bmarks
SNI/PyT RM300-C60	R4400	100/200	1M+16/16	3348	?	Dec95	c.bmarks
SNI/PyT RM300-C62	2xR4400	100/200	1M+16/16	6487	?	Dec95	c.bmarks
SNI/PyT RM400-C70	R4400	100/200	4M+16/16	3574c	?	Dec95	c.bmarks
SNI/PyT RM400-C70	2xR4400	100/200	4M+16/16	6971	?	Dec95	c.bmarks
SNI/PyT RM400-C70	4xR4400	100/200	4M+16/16	13152	?	Dec95	c.bmarks
SNI/PyT RM1000	R4400	100/200	4M+16/16	3607c	?	Aug95	SNI/PyT
Dell DimensionXPS	Pentium	60/120	512+8/8	3811c	2500c	Nov95	www.intel
Dell DimensionXPS	Pentium	66/133	512+8/8	4219c	2751c	Nov95	www.intel
Micronics M4P	80486DX4	33/100	256+16	1219c	631c	Mar94	c.arch
Intel 486DX	80486	50	256+8	713c	332c	Oct92	c.arch
Intel 486DX2	80486DX2	33/66	0+8	768c	382c	Sep92	uproc rpt
Intel Xpress	Pentium	60	256+8/8	1670c	1307c	Mar95	www.intel
Intel Xpress	Pentium	66	256+8/8	1850c	1508c	Mar95	www.intel
Intel Xpress	Pentium	50/75	512+8/8	2113c	1625c	Mar95	www.intel
Intel Xpress	Pentium	60/90	512+8/8	2526c	1931c	Mar95	www.intel
Intel Xpress	Pentium	60/90	1M+8/8	2611c	2002c	Mar95	www.intel
Intel Xpress	Pentium	66/100	512+8/8	2801c	2132c	Mar95	www.intel
Intel Xpress	Pentium	66/100	1M+8/8	2891c	2210c	Mar95	www.intel
Intel Xpress	Pentium	60/120	512+8/8	3171c	2350c	Mar95	www.intel
Intel Xpress	Pentium	60/120	1M+8/8	3320c	2464c	Mar95	www.intel
Intel Xpress	Pentium	66/133	512+8/8	3498c	2599c	Jun95	www.intel
Intel Xpress	Pentium	66/133	1M+8/8	3688c	2773c	Jun95	www.intel
Intel Alder	PentiumPro	150	256+8/8	6553c	5218c	Nov95	www.intel
Intel Alder	PentiumPro	166	512+8/8	7758c	6197c	Nov95	www.intel
Intel Alder	PentiumPro	180	256+8/8	7765c	6039c	Nov95	www.intel
Intel Alder	PentiumPro	200	256+8/8	8681c	6717c	Nov95	www.intel
Intel XXpress	Pentium	60/120	1M+8/8	4084c	2571c	Nov95	www.intel
Intel XXpress	Pentium	66/133	1M+8/8	4528c	2860c	Nov95	www.intel

APPENDIX C. HEURISTIC BENCHMARK

This appendix contains a table of computers, and their SPEC benchmark values, used to assign benchmark values to each level of the heuristic presented in Chapter V. This list is representative of computers within each level. It is not inclusive of all possible computers, listing only those with reported SPEC values. For those computers not listed in this table refer to Appendix B to obtain a SPEC value. That value can then be compare to those listed here to determine a relative heuristic level.

The table begins on the next page of this appendix.

HEURISTIC BENCHMARK

Note, The information was obtained from the following sources:

[1] SPEC: <http://www.specbench.org>

[2] Dimarco: <ftp://ftp.cdf.toronto.edu/pub/spectable>

[3] Univ. Tenn.: <http://performance.netlib.org/performance/html/spec.html>

[4] Berkeley: <http://infopad.eecs.berkeley.edu/CIC/summary/local>

Reference Computer	System		SPECint-Base95		SPECint95		SPECint-Base92		SPECint92	
		MHz		[Source]		[Source]		[Source]		[Source]
	Sun SPARC 10	40	1.0	[1]	1.13	[2]			50.2	[2]
Level I SPEC92 Avg. 30	Compaq Deskpro [486DX]	33							18.2	[2]
	Compaq Deskpro [486DX2]	50							25.7	[2]
	Intel/ 486DX2								30.1	[2]
	Siemens Nixdorf MX300 Model 75 [486DX2]	50							28.8	[3]
	Siemens Nixdorf MX300 Model 75 [486DX2]	50							30.0	[3]
	Intel [486DX2]	50							27.9	[4]
	IBM N40 ['Mac' 601 chip]	50							41.7	[2]
	"Power PC" ['Mac' 601chip]	50							40.0	[4]
Averages									30	
Level II SPEC92 Avg. 60	Compaq Deskpro [486DX2]	66							32.2	[2]
	Intel [486DX2]	66							32.4	[2]
	Micronic M4P [486DX2]	66					36.7	[3]	39.6	[3]
	Intel [486DX2]	66							32.2	[4]
	Siemens Nixdorf PC/D5T [Pentium]	60					63.9	[2]	66.0	[2]
	Mobius P5-60	60							50.0	[2]
	Siemens Nixdorf PC/E5S [Pentium]	60							60.6	[2]
	Intel Xpress [Pentium]	60							70.4	[2]
	Intel Xpress Desktop [Pentium]	60					66.9	[3]	70.4	[3]
	Intel Express Desktop [Pentium]	60					60.4	[3]	63.3	[3]
	Compaq DeskproXL [Pentium]	66							65.1	[2]
	Siemens Nixdorf PC/E5S [Pentium]	66							67.4	[2]
	Intel Xpress [Pentium]	66							78.0	[2]
	Intel Xpress Desktop Pentium]	66					74.0	[3]	78.0	[3]
	Intel Xpress Desktop [Pentium]	66					67.0	[3]	70.2	[3]
	Intel Xpress MX Deskside [Pentium]	66							64.6	[3]
	Intel [Pentium]	66							78.0	[4]
	IBM 250 ['Mac' 601 chip]	66							62.6	[2]
	IBM 25T ['Mac' 601 chip]	66							62.6	[2]
	IBM 40P ['Mac' 601 chip]	66							63.7	[2]
	IBM 40P ['Mac' 601 chip]	66							75.1	[2]
	IBM 40P ['Mac' 601 chip]	66					52.4	[3]	64.2	[3]
	Motorola PowerStack 603 Series E [603 chip]	66					50.7	[3]	63.7	[3]
	Motorola PowerStack 603 Series E [603 chip]	66					48.0	[3]	60.6	[3]
	IBM 40P ['Mac' 601 chip]	66					62.1	[3]	76.0	[3]
	"Power PC" ['Mac' 602chip]	66							40.0	[4]
Averages							58		61	
Level III SPEC92 Avg. 30	Dec Station 5000/ 900	40							27.3	[2]
	Dec Station 5000/ 240	40							27.9	[2]
	Dec Station 5000/ 50	50							43.2	[3]
	Sun SPARC IPX	40							21.8	[2]
Averages									30	

HEURISTIC BENCHMARK

	System	MHz	SPECint-Base95		SPECint95		SPECint-Base92		SPECint92	
			Base95	[Source]	int95	[Source]	Base92	[Source]	int92	[Source]
Level IV	Sun SPARC 5	70							57.0	[2]
	Sun SPARC 5	70					49.8	[3]	57.0	[3]
	Sun SPARC 5	85							65.3	[2]
	Sun SPARC 5	85					56.3	[3]	64.1	[3]
	Sun SPARC 5	85							64.0	[4]
	Sun SPARC 5	110	1.37	[2]	1.59	[2]			78.6	[2]
	Sun SPARC 5	110					68.7	[3]	78.6	[3]
	Sun SPARC 5	110							76.0	[4]
	SPEC95 Avg. 2.18									
	SPEC92 Avg. 85									
	Gateway P5-75	75	2.31	[2]	2.31	[2]			89.1	[2]
	Intel Xpress [Pentium]	75							89.1	[3]
	Intel Xpress Desktop 610 [Pentium]	75					85.0	[3]	83.8	[3]
	Intel Xpress Desktop 610 [Pentium]	75					79.0	[3]	89.1	[4]
	Intel [Pentium]									
	Gateway P5-90	90	2.74	[2]	2.74	[2]				
	Siemens Nixdorf PC/E5S [Pentium]	90					82.9	[2]	86.3	[2]
	Siemens Nixdorf PC/D5T [Pentium]	90					83.0	[2]	86.0	[2]
	Intel Xpress [Pentium]	90							106.5	[2]
	Intel Xpress [Pentium]	90							110.1	[2]
	Intel Xpress Desktop 735 [Pentium]	90					104.3	[3]	110.0	[3]
	Intel Xpress Desktop 735 [Pentium]	90					101.0	[3]	106.5	[3]
	Intel Xpress Desktop 735 [Pentium]	90					99.5	[3]	104.5	[3]
	Intel Xpress Desktop 735 [Pentium]	90					96.1	[3]	100.9	[3]
	Intel Xpress Desktop 735 [Pentium]	90					85.4	[3]	90.1	[3]
	Intel [Pentium]	90							110.0	[4]
	Averages		2.14		2.21		83		86	
Level V	Sun SPARC 20 Model 71	75			2.46	[1]				
	Sun SPARC 20 Model 71	75	2.82	[2]	3.11	[2]			125.8	[2]
	Sun SPARC 20 Model 71	75					116.4	[3]	125.8	[3]
	Sun SPARC 20 HS11	100							104.5	[2]
	Sun SPARC 20 HS11	100					94.0	[3]	104.5	[3]
	Sun SPARC 20 HS21	125							131.2	[2]
	Sun SPARC 20 HS21	125					122.4	[3]	131.2	[3]
	DEC AlphaStation 200 4/100	100	1.48	[1]	1.48	[1]	68.6	[2]	74.6	[2]
	DEC AXPPci 33	166					69.4	[3]	76.0	[3]
	DEC AlphaStation 200 4/166	166	2.31	[1]	2.31	[1]	100.1	[2]	116.2	[2]
	DEC AlphaStation 400 4/166	166					100.1	[2]	116.2	[2]
	DEC AlphaStation 400 4/166	166					107.5	[3]	116.8	[3]
	DEC AlphaServer 1000 4/200	200					123.3	[2]	135.8	[2]
	DEC AlphaServer 2000 4/200	200					117.5	[2]	131.8	[2]
	DEC AlphaServer 2100 4/200	200					117.5	[2]	131.8	[2]
	DEC AlphaStation 200 4/233	233	3.39	[1]	3.39	[1]	137.4	[2]	157.7	[2]
	DEC AlphaStation 400 4/233	233					137.4	[2]	157.7	[2]
	DEC AlphaStation 400 4/233	233					136.2	[3]	155.2	[3]
	DEC AlphaServer 1000 4/233	233							165.3	[2]
	DEC AlphaServer 2100 4/233	233					163.7	[2]	177.3	[2]
	DEC AlphaServer 2100 5/250	250	5.96	[1]	5.96	[1]	244.7	[2]	277.1	[2]
	[Level V continued]									

HEURISTIC BENCHMARK

	System	MHz	SPECint-		SPEC		SPECint-		SPEC	
			Base95	[Source]	int95	[Source]	Base92	[Source]	int92	[Source]
Level V	DEC AlphaStation 250 4/266	266	4.18	[1]	4.18	[1]	182.6	[2]	198.6	[2]
	DEC AlphaStation 600 5/266	266	6.3	[1]	6.3	[1]				
	DEC AlphaStation 600 5/266	266	6.43	[2]	6.43	[2]	257.1	[2]	289.0	[2]
	DEC AlphaStation 600 5/266	266					256.9	[3]	288.6	[3]
	DEC AlphaServer 2100 4/275	275					187.8	[2]	202.9	[2]
	DEC AlphaServer 2100 4/275	275					176.5	[3]	200.1	[3]
	HP 9000 Series 700 Model 712/80	80							97.1	[2]
	HP 9000 Series 700 Model 712/80	80					76.6	[3]	84.3	[3]
	HP 9000 Series 700 Model 712/80	80					85.2	[3]	93.0	[3]
	HP 9000 Series 700 Model 715/80	80							96.3	[2]
	HP 9000 Series 700 Model 715/80	80					75.0	[3]	83.5	[3]
	HP 9000 Series 700 Model 715/80	80					85.2	[3]	93.0	[3]
	HP 9000 E45	80					74.5	[2]	82.1	[2]
	HP 9000 E45	80					84.1	[3]	92.5	[3]
	SPEC95 Avg. 3.80									
	HP 9000 T 500	90							98.3	[2]
	HP 9000 T 500	90					107.2	[3]	115.1	[3]
	SPEC92 Avg. 129									
	HP 9000 G 50	96							100.3	[3]
	HP 9000 E 55	96					96.1	[2]	108.0	[2]
	HP 9000 E 55	96					108.3	[3]	118.9	[3]
	HP 9000 G/H/I 60	96							108.8	[2]
	HP 9000 G/H/I 60	96							82.0	[3]
	HP 9000 Series 700 Model 735/99	99	3.27	[1]	3.27	[1]				
	HP 9000 Series 700 Model 735/99	99	3.13	[2]	3.22	[2]			109.1	[2]
	HP 9000 Series 700 Model 735/99	99					111.9	[3]	119.7	[3]
	HP 9000 Series 700 Model 712/100	100							117.2	[2]
	HP 9000 Series 700 Model 715/100	100	2.89	[1]	2.89	[1]			115.1	[2]
	HP 9000 Series 700 Model 715/100	100					89.7	[3]	99.6	[3]
	HP 9000 Series 700 Model 715/100	100					100.0	[3]	109.6	[3]
	HP 9000 K 400	100	3.58	[1]	3.58	[1]	113.4	[3]	136.4	[3]
	HP 9000 Series 700 Model 735/125	125	4.04	[1]	4.04	[1]				
	HP 9000 Series 700 Model 735/125	125	3.88	[2]	3.97	[2]			136.0	[2]
	HP 9000 Series 700 Model 735/125	125					123.8	[3]	135.7	[3]
	HP 9000 Series 700 Model 735/125	125					138.5	[3]	149.4	[3]
	Averages		3.83		3.77		125		131	

APPENDIX D. WEB SITE SURVEY

This appendix presents 19 surveys which were conducted to determine the validity of the hardware heuristic detailed in Chapter V. The surveys begins on the next page of this appendix.

Site: Defense Information Systems Agency (DISA) ([http://www.itis.disa.mil /](http://www.itis.disa.mil/))

POC: John Bridger (703) 735-3544

File Size:

- 'Typical' HTML (~10K): yes (80%)
- Video/Sound/etc: no

Connection: T-1

CPU:

- Scripts: yes (10%)
- Database Searches: yes (10%)

Traffic:

- Average Hits per Hour: ~70 (Avg daily files for Dec 95: 1,657)
- Peak Hourly Peak: ~140

Equipment:

- Computer: SPARC 10 Model 30
- Speed (MHz): 50
- RAM: 64Mb
- Cache: unknown
- SPEC Benchmark: SPECint92: 45 (for a 30 MHz model 30; SPECint92: 65.4 for a 50 MHz model 51)
- Heuristic Level: Level Three (or Four if based on Model 51)

Calculated Heuristic Level

Start:	+1
Files:	+1
Connection:	-1
CPU:	+1
Hits:	-1
Total (Level):	2

Note: Adding a SPARC 20 and will places different function on different computrs.

Site: Federal Emergency Management Agency (<http://www.fema.gov/>)

POC: Bill Casti (202) 646-4600

File Size:

- 'Typical' HTML (~10K): yes (85%)
- Video/Sound/etc: Voice message from Director (accessed 100 times a day)

Connection: T-1

CPU:

- Scripts: yes (15%)
- Database Searches: no

Traffic:

- Average Hits per Hour: ~800
- Peak Hourly Peak: ~1,600

Equipment:

- Computer: Dual Processor DEC 3000/400
- Speed (MHz): 133 (each processor)
- RAM: 32Mb
- Cache: unknown
- SPEC Benchmark: SPECint92: 74.7 (single processor)
- Heuristic Level: Level Six (due to dual processors)

Calculated Heuristic Level

Start:	+1
Files:	+2
Connection:	-1
CPU:	0
Hits:	+1
Total (Level):	3

Site: IDC Government, Falls Church, Va. (<http://www.idcg.com>)

POC: Kelly Kavanagh (703) 876-5043

File Size:

- 'Typical' HTML (~10K): yes (66%)
- Video/Sound/etc: no

Connection: 64Kb

CPU:

- Scripts: no
- Database Searches: yes (34%)

Traffic:

- Average Hits per Hour: one or two hits an hour
- Peak Hourly Peak:

Equipment:

- Computer: DEC Pentium
- Speed (MHz): 100
- RAM: 64MB
- Cache: unknown
- SPEC Benchmark: unavailable (estimated SPECint_base95: 3.06 - 3.16;
SPECint_base92: 92-126)
- Heuristic Level: Level Five

Calculated Heuristic Level:

Start:	+1
Files:	+1
Connection:	0
CPU:	+1
Hits:	- 1
Total (Level):	2

Site: Internet Society, Reston Va. (<http://www.isoc.org>)

POC: Jay Whittle (703) 648-9888

File Size:

- 'Typical' HTML (~10K): yes (80%)
- Video/Sound/etc: limited (one, 1MB sound file several times a day)

Connection: T-1

CPU:

- Scripts: yes (20%)
- Database Searches: no

Traffic:

- Average Hits per Hour: ~580 (98,000 a week)
- Peak Hourly Peak: ~1,100

Equipment:

- Computer: Four Processor SPARC 1000
- Speed (MHz): 66 (each processor)
- RAM: 128Mb
- Cache: unknown
- SPEC Benchmark: unknown (very Fast!)
- Heuristic Level: Level Six (due to multi-processors)

Calculated Heuristic Level:

Start:	+1
Files:	+2
Connection:	- 1
CPU:	+1
Hits:	+1
Total (Level):	4

Note: Large increase in script use anticipated in next year due to increasing membership.
Current equipment donated.

Site: Library of Congress (<http://lcweb.loc.gov/>)

POC: Tom Littlejohn (202) 707-9073

File Size:

- 'Typical' HTML (~10K): yes (70%)
- Video/Sound/etc: no

Connection: Ethernet (10mb)

CPU:

- Scripts: yes (30%)
- Database Searches: yes (part of 30%)

Traffic:

- Average Hits per Hour: ~6,000
- Peak Hourly Peak: ~9,500

Equipment:

- Computer: Two IBM/ 6000 980's
- Speed (MHz): ~60
- RAM: 256Mb
- Cache: 64k Data/32k Instruction
- SPEC Benchmark: unavailable (estimated to be less than IBM 990:
~100)
- Heuristic Level: Level Six (due to dual computers)

Calculated Heuristic Level:

Start:	+1
Files:	+1
Connection:	- 1
CPU:	+1
Hits:	+3
Total (Level):	5

Note: Upgrading to IBM R30 (eight processors) and one Gig RAM due to "digital Library project" which will have five mil. Images by year 2000.

Site: Thomas Server - Library of Congress (<http://thomas.loc.gov>)

POC: Tom Littlejohn (202) 707-9073

File Size:

- 'Typical' HTML (~10K): yes (light)
- Video/Sound/etc: no

Connection: Ethernet (10MB)

CPU:

- Scripts: yes (heavy)
- Database Searches: yes (heavy)

Traffic:

- Average Hits per Hour: ~3,400
- Peak Hourly Peak: ~6,000

Equipment:

- Computer: IBM RS/ 6000 990
- Speed (MHz): 71
- RAM: 256Mb
- Cache: 256k Data/32k Instruction
- SPEC Benchmark: SPECint92: 125.9
- Heuristic Level: Level Five

Calculated Heuristic Level:

Start:	+1
Files:	+1
Connection:	- 1
CPU:	+1
<u>Hits:</u>	<u>+3</u>
Total (Level):	5

Note: Upgrading to IBM R30 (eight processors) and one Gig of RAM in preparation for "digital Library Project" which will offer five million on-line images by year 2000.

Site: National Aeronautics and Space Administration (NASA)

(<http://www.hq.nasa.gov>)

POC: Woody Smith (202) 358-1486

File Size:

- 'Typical' HTML (~10K): yes (79%)
- Video/Sound/etc: very little (~1%)

Connection: T-3

CPU:

- Scripts: yes (10%)
- Database Searches: yes (10%)

Traffic:

- Average Hits per Hour: ~10,400
- Peak Hourly Peak: ~20,000

Equipment:

- Computer: Dual Processor SPARC 10
- Speed (MHz): 55
- RAM: 64Mb
- Cache: unknown
- SPEC Benchmark: unavailable
- Heuristic Level: Level Six (due to multi-processors)

Calculated Heuristic Level:

Start:	+1
Files:	+1
Connection:	-1
CPU:	+1
Hits:	+4
Total (Level):	6

Site: NASA Ames (<http://www.arc.nasa.gov>)

POC: Tony (415) 604-4181

File Size:

- 'Typical' HTML (~10K): yes (70%)
- Video/Sound/etc: no

Connection: T-1

CPU:

- Scripts: yes (15%)
- Database Searches: yes (15%)

Traffic:

- Average Hits per Hour: ~800
- Peak Hourly Peak: ~1,600

Equipment:

- Computer: SPARC 2
- Speed (MHz): ~40
- RAM: 64Mb
- Cache: unknown
- SPEC Benchmark: SPECint92: 21.8
- Heuristic Level: Level Three

Calculated Heuristic Level:

Start:	+1
Files:	+1
Connection:	-1
CPU:	+1
Hits:	+1
Total (Level):	3

Site: NASA Jet Propulsion Laboratory (JPL) - PDS (Planetary Data System) Server (<http://stardust.jpl.nasa.gov>)
POC: Steve Mortellaro (818) 306-6029

File Size:

- 'Typical' HTML (~10K): yes (25%)
- Video/Sound/etc: no

Connection: T-1

CPU:

- Scripts: yes (75%)
- Database Searches: no

Traffic:

- Average Hits per Hour: ~400 (66,000 a week)
- Peak Hourly Peak: ~800

Equipment:

- Computer: SPARC 10
- Speed (MHz): 50
- RAM: 32MB
- Cache: unknown
- SPEC Benchmark: SPECint92: 65.2
- Heuristic Level: Level Four

Calculated Heuristic Level:

Start:	+1
Files:	+1
Connection:	- 1
CPU:	+1
Hits:	+1
Total (Level):	3

Site: National Archives, College Park, Md. (<http://www.nara.gov>)

POC: Rick Carrick (301) 713-6895

File Size:

- 'Typical' HTML (~10K): yes (79%)
- Video/Sound/etc: few (~1%)

Connection: T-1

CPU:

- Scripts: yes (10%)
- Database Searches: yes (10%)

Traffic:

- Average Hits per Hour: ~1,600
- Peak Hourly Peak: ~3,200
-

Equipment:

- Computer: SPARC 20
- Speed (MHz): 100
- RAM: 128Mb
- Cache: unknown
- SPEC Benchmark: SPECint92: 104.5
- Heuristic Level: Level Five

Calculated Heuristic Level:

Start:	+1
Files:	+1
Connection:	-1
CPU:	+1
Hits:	+2
Total (Level):	4

Site: National Institute of Standards and Technology (NIST)

(<http://www.nist.gov>)

POC: Mark Williams (301) 975-3160

File Size:

- 'Typical' HTML (~10K): yes (85%)
- Video/Sound/etc: limited (5%)

Connection: T-3

CPU:

- Scripts: no
- Database Searches: some (10%)

Traffic:

- Average Hits per Hour: ~580
- Peak Hourly Peak: ~1,000

Equipment:

- Computer: SPARC 20
- Speed (MHz): 75
- RAM: 96Mb
- Cache: unknown
- SPEC Benchmark: SPECint_base95: 2.82; SPECint92: 125.8
- Heuristic Level: Level Five

Calculated Heuristic Level:

Start:	+1
Files:	+2 (note: because of sound/video & database searches add 2)
Connection:	-1
CPU:	0
<u>Hits:</u>	<u>+1</u>
Total (Level):	3

Note: Upgrading to IBM R6000 (four processors) and 256Mbs RAM due to anticipated load increase.

Site: National Science Foundation (<http://www.nsf.gov>)

POC: Michael Morse (703) 306-1145 x4660

File Size:

- 'Typical' HTML (~10K): yes (80%)
- Video/Sound/etc: some (5%)

Connection: T-3

CPU:

- Scripts: some (5%)
- Database Searches: yes (10%)

Traffic:

- Average Hits per Hour: ~2,600 (63,000 a day)
- Peak Hourly Peak: ~5,200

Equipment:

- Computer: SPARC 20
- Speed (MHz): 50
- RAM: 32Mb
- Cache: unknown
- SPEC Benchmark: SPECint92: 76.9
- Heuristic Level: Level Four

Calculated Heuristic Level:

Start:	+1
Files:	+1
Connection:	-1
CPU:	+1 (note: because of some sound/video & database searches add 1)
Hits:	+2
Total (Level):	4

Site: Naval Medical Information Management Center
(<http://support1.med.navy.mil>)

POC: Dale Edington (301) 295-0807

File Size:

- 'Typical' HTML (~10K): yes (90%)
- Video/Sound/etc: no

Connection: T-1

CPU:

- Scripts: yes(10%)
- Database Searches: very limited

Traffic:

- Average Hits per Hour: ~200
- Peak Hourly Peak: ~400

Equipment:

- Computer: Dual Processor Entigraph P-5
- Speed (MHz): 100 (each processor)
- RAM: 32Mb
- Cache: L1 -512kb
- SPEC Benchmark: unavailable (estimated SPECint_base95: 3.06 - 3.16;
SPECint_base92: 92-126)
- Heuristic Level: Level Six (due to dual computers)

Calculated Heuristic Level:

Start:	+1
Files:	+1
Connection:	- 1
CPU:	0
Hits:	0
Total (Level):	1

Note: Recently upgraded from 486/66 in anticipation of possible increase in sight scope.

Site: Navy Online (www.ncts.navy.mil)

POC: Mike Jenkins (904) 452-3501

File Size:

- 'Typical' HTML (~10K): yes (95%)
- Video/Sound/etc: no

Connection: T-1

CPU:

- Scripts: very little (5%)
- Database Searches: no

Traffic:

- Average Hits per Hour: ~1,900
- Peak Hourly Peak: ~3,800

Equipment:

- Computer: Sun ELC
- Speed (MHz): 33
- RAM: 40Mb
- Cache: unknown
- SPEC Benchmark: SPECint92: 18.2
- Heuristic Level: Level Three

Calculated Heuristic Level:

Start:	+1
Files:	+1
Connection:	-1
CPU:	0
Hits:	+2
Total (Level):	3

Site: U.S. Department of Education (<http://www.ed.gov>)

POC: Robert Thompson (202) 219-1847

File Size:

- 'Typical' HTML (~10K): yes (60%)
- Video/Sound/etc: no

Connection: T-1

CPU:

- Scripts: some -expanding (20%)
- Database Searches: some -expanding (20%)

Traffic:

- Average Hits per Hour: ~2,800
- Peak Hourly Peak: ~5,600

Equipment:

- Computer: Dual Processor SPARC 10
- Speed (MHz): 90
- RAM: 320Mb (!)
- Cache: unknown
- SPEC Benchmark: unavailable
- Heuristic Level: Level Six (due to multi-processors)

Calculated Heuristic Level:

Start:	+1
Files:	+1
Connection:	-1
CPU:	+1
Hits:	+2
Total (Level):	4

Site: U.S. Department of Energy (<http://www.doe.gov>)

POC: Lynn Davis (423) 241-6435

File Size:

- 'Typical' HTML (~10K): yes (85%)
- Video/Sound/etc: very little (will expand)

Connection: T-1

CPU:

- Scripts: yes (5%)
- Database Searches: yes (10%)

Traffic:

- Average Hits per Hour: ~1,200
- Peak Hourly Peak: ~3,000

Equipment:

- Computer: Four Processor SPARC 1000
- Speed (MHz): 66 (each processor)
- RAM: 128Mb
- Cache: unknown
- SPEC Benchmark: unknown (very Fast!)
- Heuristic Level: Level Six (due to multi-processors)

Calculated Heuristic Level:

Start:	+1
Files:	+1
Connection:	-1
CPU:	+1
Hits:	+2
Total (Level):	4

Site: U.S. Department of Labor (<http://www.dol.gov>)

POC: Dave Dickerson

File Size:

- 'Typical' HTML (~10K): yes (95%)
- Video/Sound/etc: no

Connection: T-1

CPU:

- Scripts: yes (5%)
- Database Searches: limited

Traffic:

- Average Hits per Hour: ~750
- Peak Hourly Peak: ~1,500

Equipment:

- Computer: Dual Processor SPARC 2000
- Speed (MHz): 75
- RAM: ~34Mb
- Cache: unknown
- SPEC Benchmark: unavailable
- Heuristic Level: Level Six (due to multi-processors)

Calculated Heuristic Level:

Start:	+1
Files:	+1
Connection:	-1
CPU:	0
<u>Hits:</u>	<u>+1</u>
Total (Level):	2

Note: Moving to NT Information Server.

Site: U.S. Fish and Wildlife Service (<http://www.fws.gov>)

POC: Alan Fisher (303) 275-2320

File Size:

- 'Typical' HTML (~10K): yes (90%)
- Video/Sound/etc: no

Connection: T-1 (probably fractional)

CPU:

- Scripts: yes (5%)
- Database Searches: yes (5%)

Traffic:

- Average Hits per Hour: ~350
- Peak Hourly Peak: ~700

Equipment:

- Computer: SPARC 10
- Speed (MHz): 85
- RAM: 32Mb
- Cache: unknown
- SPEC Benchmark: unavailable (approximated as a 85Mhz SPARC 4 or 5 - SPECint92: 65)
- Heuristic Level: Level Four

Calculated Heuristic Level:

Start:	+1
Files:	+1
Connection:	0
CPU:	0
<u>Hits:</u>	<u>0</u>
Total (Level):	2

Site: U. S. Patent and Trademark Office (<http://www.uspto.gov/>)

POC: John Ridell (703) 308-6873

File Size:

- 'Typical' HTML (~10K): yes (90%)
- Video/Sound/etc: no

Connection: T-1

CPU:

- Scripts: yes (5%)
- Database Searches: yes (5%)

Traffic:

- Average Hits per Hour: ~1,300
- Peak Hourly Peak: ~2,600

Equipment:

- Computer: SPARC 10
- Speed (MHz): 40
- RAM: 64MB
- Cache: unknown
- SPEC Benchmark: SPECint_base95: 1.0; SPECint92: 50
- Heuristic Level: Level 3 ½ (SPEC benchmark falls between levels three and four)

Calculated Heuristic Level:

Start:	+1
Files:	+1
Connection:	- 1
CPU:	0
Hits:	+2
Total (Level):	3

LIST OF REFERENCES

- L. Armstrong. "Setting Up a Web Site", *LAN TIMES*, April 24, 1995.
- L. Aronson, *HTML Manual of Style*, Ziff-Davis Press, Emertville, California, 1994.
- T. Berners-Lee, R. Cailliau, A. Luotonen, H.F. Nielsen, and A. Secret, "The World-WIDE WEB". *Communications of the ACM*, 37(8):76-82, 1994.
- J. R. Blumenfeld, I. Etra, J. Gartner, and W. Gee. "Step-by-Step to a WORLD-CLASS WEB SITE", *WINDOWS MAGAZINE*, July 1995.
- S. S. Cambell, R. J. Kohlhepp, and J. Linder. "Serving Up HTML Documents On The Web", *NETWORK COMPUTING*, January, 15, 1996.
- A. Carlton, Technical Engineer for Standard Performance Evaluation Corporation (SPEC). Phone Interview March 8, 1996.
- E. Carr. "RAID Without Compromise", *Network Computing*, September 1, 1995.
- W. Cohen. "Why the Chip is Still the Economy's Champ", *U.S. News & World Report*, March 25, 1996.
- J. DiMarco. SPEC Benchmark Report File, URL <ftp://ftp.cdf.toronto.edu/pub/spectable>, 1996.
- D. Chandler. *Running a Perfect Web Site*, Que Corporation, Indianapolis, Indiana, 1995.
- K. Dixit and J. Reilly. "Answers to Common Questions About SPEC95 Benchmark Suites", URL <http://www.netlib.org/performance/html/spec-q-a.html>, August 1995.
- J. Evans, Naval Command, Control & Ocean Surveillance Center, San Diego, California. E-mail Interview August 8-15, 1995.
- J. Fitzgerald. *Business Data Communications, Basic Concept, Security, and Design*, John Wiley & Sons, Inc., New York, New York, 1993.
- M. Gray. "Measuring the Growth of the Web, June 1993 to June 1995", URL <http://www.netgen.com/info/growth.html>, 1995.
- M. Gray. Chief Technologist, Net.Genesis, Cambridge, Massachusetts. Phone Interview August 16, 1995.

P. E. Hoffman, Proper Publishing, "Web Servers Survey", URL <http://www.proper.com/www/servers-survey.html>, January 1996.

P. Houston, "Transarc Corporation's World-Wide web Strategy Overview", URL <http://www.transarc.com/afs/transarc.com/public/www/Public/ProdServ/Product/Whitepapers/WebStrategy/webstrat.htm>.

J. Hunter, Net + Effect, San Jose, California. Phone Interview August 30, 1995.

E. Katz, M. Butler, and R. McGrath. "A Scalable HTTP Sever: The NCSA Prototype", URL <http://www.ncsa.uiuc.edu/InformationServers/Conferences/CERNwww94/www94.ncsa.html>, 1994.

R. Levine. "RAID. Wasted Days, Wasted Nights", *SunExpert*, March 1996.

C. Lin. "RAIDers of the Lost Drive", *PC MAGAZINE*, March 26, 1996.

T. Littlejohn, Information Technology Services, Library of Congress, Washington D.C. Phone Interview February 22, 1996.

C. Liu, J. Peek, R. Jones, B. Buus, and A. Nye. *Managing INTERNET Information Services*, O'Reilly and Associates, Inc., Sebastopol, California, 1994.

D. Lloyd and M. Lipow. *Reliability: Management, Methods, and Mathematics*, second edition, American Society For Quality Control, Milwaukee Wisconsin, 1984.

R. McGrath and N. Yeager. "Serving Large Files with the Web: Preliminary Outline of Experiments", URL <http://www.ncsa.uiuc.edu/People/mcgrath/docs.html>, January, 19, 1995.

R. McGrath, National Center for Supercomputing Applications (NCSA), University of Illinois. Phone Interview April 16, 1996.

M. Meyer. "How Fast does My Net Connection Need to Be?", URL <http://www.boutell.com/faq/sband.htm>, 1995.

J. Milne. "RAID Storage Solutions: Which Is Right For You?", *Network Computing*, May 1, 1995.

J. Milne. "Commodity Superservers: Muscle-bound PCs", *Network Computing*, October 15, 1995.

Net.Genesis and D. Hall. *Build a Web Site*, Prima Publishing, Rocklin, California, 1995.

Network Wizards. "Internet Domain Survey, July 1995", URL <http://www.nw.com/zone/WWW/report.html>, July 1995.

D. Norman, Director of Computing Services, Naval Postgraduate School, Monterey, California. Interview November 1995.

J. Powell. "Spinning the Web: Setting up WWW Servers", URL <http://scholar.lib.vt.edu/reports/Servers-web.html>, September 1994.

J. Reilly. "SPEC Discusses the History & Reasoning Behind SPEC95", URL <http://www.specbench.org/osg/cpu95/news/cpu95reason.html>, September 1995.

D. Rinaldi. "RICK/CISC Debate Over: Customers Win", *CLIENT / SERVER COMPUTING*, April 1995.

J. Schiller, Network Operations Center Manager, Massachusetts Institute of Technology, Cambridge, Massachusetts. E-Mail Interview October 15-16, 1995.

D. Simpson. "Can't Tolerate SERVER Downtime? CLUSTER'EM!", *DATAMATION*, August 15, 1995.

B. Smith. "Making The Internet Connection", *BYTE*, January, 1995.

Standard Performance Evaluation Corporation (SPEC), "A SPEC Primer", URL <http://www.specbench.org/spec/specfaq.html#contact>, December, 1995.

SPEC and the University of Tennessee, "PDS: The Performance Database Server", URL <http://performance.netlib.org/performance/html/spec.html>.

O. R. Tabibian. "Electronic Buisness", *PC MAGAZINE*, February 7, 1995.

Transarc Corporation. "AFS Frequently Asked Questions", URL <http://www.transarc.com/afs/transarc.com/public/www/Product/AFS/FAQ/faq.html#sub1.05.c>, March 1996.

C. Wilder. "Plugging Into The Web", *INFORMATIONWEEK*, April 3, 1995.

M. Woods, Professor of Mathematics, Naval Postgraduate School, Monterey, California. Conversation April 1996.

INITIAL DISTRIBUTION LIST

- | | | |
|----|--|---|
| 1. | Defense Technical Information Center
8725 John J. Kingman Rd., STE 0944
Ft. Belvoir, VA 22060-6218 | 2 |
| 2. | Dudley Knox Library
Naval Postgraduate School
411 Dyer Rd.
Monterey, CA 93943-5101 | 2 |
| 3. | Professor Hemant K. Bhargava
Department of Systems Management (Code SM/Bh)
Naval Postgraduate School
Monterey, CA 93943 | 4 |
| 4. | Associate Professor Suresh Sridhar
Department of Systems Management (Code SM/Sr)
Naval Postgraduate School
Monterey, CA 93943 | 2 |
| 5. | Professor Ramayya Krishnan
The Heinz School
Carnegie Mellon University
Pittsburgh, PA 15213 | 1 |
| 6. | Professor Rudolf Muller
Institute for Information Systems
Humboldt\University Berlin
Spandauer Str. 1, 10178
Berlin, Germany | 1 |
| 7. | Kelly Kavanagh
Care of IDC Government
3110 Fairview Park Dr.
Suite 1100
Falls Church, VA. 22042 | 1 |

- | | | |
|------|--|---|
| 8. | Bill Casti, CQA
FEMA
500 C Street, Room 820-H
Washington, DC 20472 | 1 |
|
 | | |
| 9. | Lisa Weber
National Archives At College Park
Policy and Information Resource Management Services
Suite 3200 - PIRM
8601 Adelphi Rd.
College Park, MD 20740-6001 | 1 |