An Internet Storage Architecture Guide

Optimally deploying
DAS and NAS
in strategic large
scale Internet sites

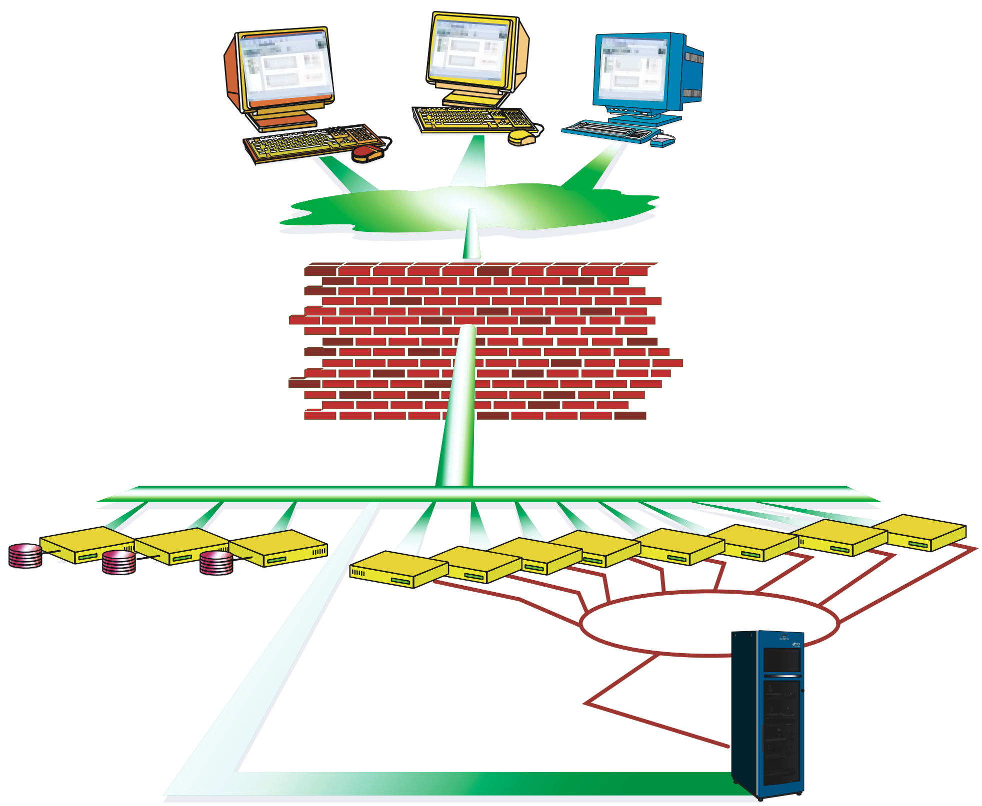

Auspex Technical Report

AUSPEX

An Internet Storage Architecture Guide

Optimally deploying
DAS and NAS
in strategic large
scale Internet sites

How the front-end Web browser works

Connecting to the Internet

The Internet itself is a loose collection of Networks

Anatomy of a Web Site

Components of an effective Web Site Architecture

Appropriate World Wide Web applications for an intelligent web cache

The Auspex NetServer 2000 deployed in internet sites

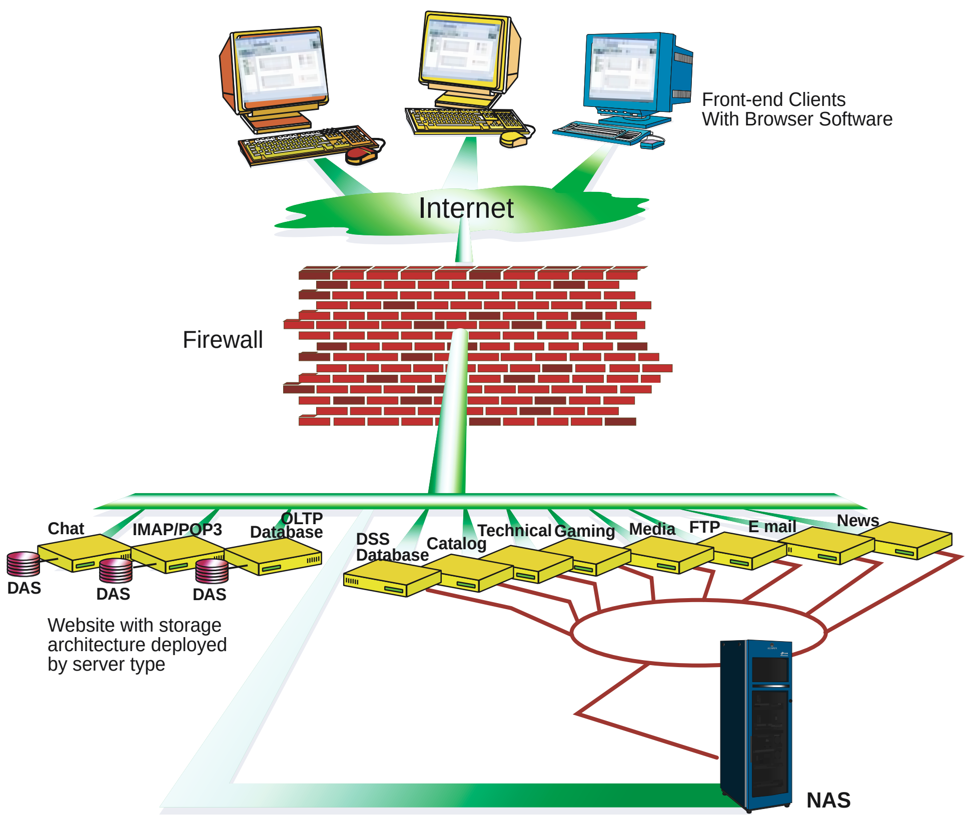

Website with storage architecture deployed by server type

AUSPEX

Table of Contents

Table of Contents ... v
Table of Figures ... viii
Tables ... ix

Chapter 1 – Introduction ... 1
The importance of optimally designing Web Site architectures ... 1
Challenges faced in optimizing Web Site architectures ... 1
A Multi-level functionally specialized Web Site architecture is proposed ... 1
Auspex pioneered functional specialization in computer storage design ... 3
The importance of understanding Internet technology and terminology ... 3
A standard way of viewing the Front-End to Back-End Internet Model ... 4
Other Information sources ... 4

Chapter 2- How the front - end Web browser works ... 5
How URLs work ... 6
Domain Names ... 6
Domain Name Servers ... 7
How cookies work ... 7

Chapter 3- Connecting to the Internet ... 9
POTS ... 9
Web TV ... 9
Cable Modems ... 10
ASDL ... 10
SDSL ... 10
XDSL ... 10
Wireless Internet Connections ... 11
ISDN ... 11
Other network connections ... 11

Chapter 4 - The Internet itself is a loose collection of Networks ... 13
Backbones ... 13
Regional networks ... 14
Routers ... 14
The Internet Society ... 14

Internet Protocols 15
OSI Model 15
TCP/IP 16
HTTP 16
UDP/IP 16
FTP 16
SLIP and PPP 16
CGI 16
File types on the Internet 18

Chapter 5 - Anatomy of a Web Site ________ 19

Border Routers 20
Gateways 20
Firewall 20
Bastion Host 20
Bridges and Hubs 20
Load Balancers 20
FTP Servers 21
HTTP Servers 21
Chat Servers 21
Email Servers 21
Database Servers 21

Chapter 6 – Components of an effective Web Site Architecture ________ 23

A flat clustered Web Site architecture 23
Frequently reported problems with a flat clustered Web Site design 24
A multi-level functionally specific hierarchical architecture 24
The secrets to designing an optimized hierarchical Web Site 25
Deciding when to deploy a server with DAS or NAS 26
Server types appropriate for DAS deployment 26
Server types appropriate for NAS deployment 26
Front-end redundant routers, e-commerce servers, and back-end network paths 27
Scaling 27
Assigning and Balancing Web Loads – Load Balancers 27
Fault Tolerance 29
Tuning individual servers in an Internet site 29

Chapter 7 – The Auspex NetServer 2000 deployed in Internet sites ______________ 31

Network Attached Storage (NAS) has migrated to the Internet 31

Auspex NS2000 Servers in Internet Sites 31

The most advanced design available in Network Attached Storage 32

The highest level of expertise available among NAS vendors 33

Parallel hardware and software design of the Auspex NetServer 2000 33

I/O Node Building Blocks 34

Auspex NS2000 Functional Multiprocessing (FMP) Parallel Architecture 35

Auspex NetServer 2000 Host Node and System Management 35

Why Auspex chose an FMP over an SMP architecture 36

Chapter 8 – Appropriate World Wide Web applications for an Intelligent Web cache __ 37

How Web Caches work 37

Client Side Web Caching 37

Traditional "Client Side" Web Caches 37

Pipelining: Fast Content Retrieval, The First Time 37

Server Side Web Caching 38

Overloaded Servers and Impatient Consumers 39

Pinpointing E-commerce Bottlenecks 40

Intelligent Web Caches reduce the risk of backend devices becoming points of failure and causing unplanned outages 41

"Server Side" Intelligent Web Caches reduce many risks to e-commerce sites 41

Chapter 9 - Conclusion ______________ 43

Summary of the Benefits of a Hierarchical Web Site Architecture 43

Flexibility (Scalability) 43

7x24 (Availability) 44

Speed (Performance or Fast Response Time to Users) 44

Manageability (Ease of Administration) 44

Cost Effectiveness 45

Native FTP Support 45

Glossary of Terms ______________ 47

Table of Figures

Figure 1 A multi-level functionally specialized hierarchical Web Site architecture provides an optimal balance of 24x7, flexibility, speed and manageability. 2

Figure 2 The basic three part Internet model. 4

Figure 3 How a browser retrieves a Web page from a Web server. 5

Figure 4 Web connections can be accomplished through POTS, Web TV, ADSL, SDSL, XDSL, wireless, ISDN and other networks. 9

Figure 5 The Internet is a loose organization of public and private networks comprised of backbones, routers, hubs, switches, software protocols and naming conventions. 14

Figure 6 How CGI programs enable database searches over the Web. 17

Figure 7 A Web Site must be secure, available 24x7, fast and easily manageable. 19

Figure 8 A flat clustered Web Server architecture results in scattered critical data and scattered hot and cold data on the same disks. 23

Figure 9 A multi-level functionally specialized hierarchical Web Site architecture provides an optimal balance of 24x7, flexibility, speed and manageability. 25

Figure 10 The Auspex NS2000 product line scales from workgroup to data center. 33

Figure 11 System block diagram of the Auspex NS2000 parallel processing architecture. 34

Figure 12 I/O Node block diagram illustrating the Auspex patented Functional Multiprocessing (FMP) design. 35

Figure 13 Intelligent “Client side” Web Caches retrieve objects in parallel instead of serially. 38

Figure 14 A typical E-commerce Web site where all requests must pass through the firewall and L4 switch. 40

Figure 15 Server Side Caching, with an Intelligent Web Cache, services frequent requests before the firewall. 42

Tables

Table 1 Common Top Level Domain (TLD) names. 6

Table 2 OSI Reference Model structures communication software in layers. 15

Table 3 Popular file formats sent over the Internet. 18

Table 4 The NS2000 distributes processing functions not only among processors within nodes but also among processors between nodes. SMP computers perform all functions in one node. 36

Table 5 Many users cancel transactions if a page takes longer than 8 seconds to retrieve. 38

Introduction

1

The importance of optimally designing Web Site architectures

The computing architecture of ISPs, ASPs, and corporate Web Sites[1] have grown up as quickly as the Internet to support the booming Business to Business (B2B) and Business to Consumer (B2C) applications that comprise E-commerce. Now is the time to optimize Web Site architectures for the increasingly critical requirements of 24x7, flexibility, speed and manageability. Recent well-publicized security attacks and system outages due to server failures have underscored this need. Network Attached Storage (NAS) can play an important role in meeting these objectives and should be strategically placed in the Internet site along with other Web Site components in order to ensure the best results.

Challenges faced in optimizing Web Site architectures

There are many Web Site architecture design challenges faced by B2B and B2C Internet technology enterprises. When a "fat O/S" model Direct Attached Storage (DAS) is used for large scale Web server farms in a flat cluster – like topology, the same problems experienced by corporations with many distributed general purpose servers are encountered. This is because an individual Web server is asked to do more than it is designed to do.

Servers are designed to be compute intensive, yet this flat two-tier architecture asks the servers to process both IP requests, file system requests and manage storage for whatever domain names are associated with that particular server.

Reliability is decreased as each server's working set increases to perform both IP requests and filesystem processing.

There is no single point of disk backup and a network based backup scheme is often used. Disk utilization is inefficient since disk space cannot be allocated as needed. Load balancing techniques are less easily applied since there is no clean way to divorce filesystem processing from IP-request processing.

Flat two-tier architecture asks the servers to process both IP requests, file system requests and manage storage for whatever domain names are associated with that particular server.

A Multilevel functionally specialized Web Site architecture is proposed

Implementing a multilevel functionally specialized Web Site architecture with devices optimized to do only one job well can solve these problems. This architecture essentially balances computing tasks in a way that enables an optimal balance of 24x7, flexibility, speed and manageability in a cost effective manner not achieved by other designs. It is shown in **Figure 1**:

Implementing a multi-level functionally specialized Web Site architecture with devices optimized to do only one job well can solve these problems.

[1]See the Glossary of Terms for definitions of all acronyms used in this Guide.

Figure 1 – A multi-level functionally specialized hierarchical Web Site architecture provides an optimal balance of 24x7, flexibility, speed and manageability .

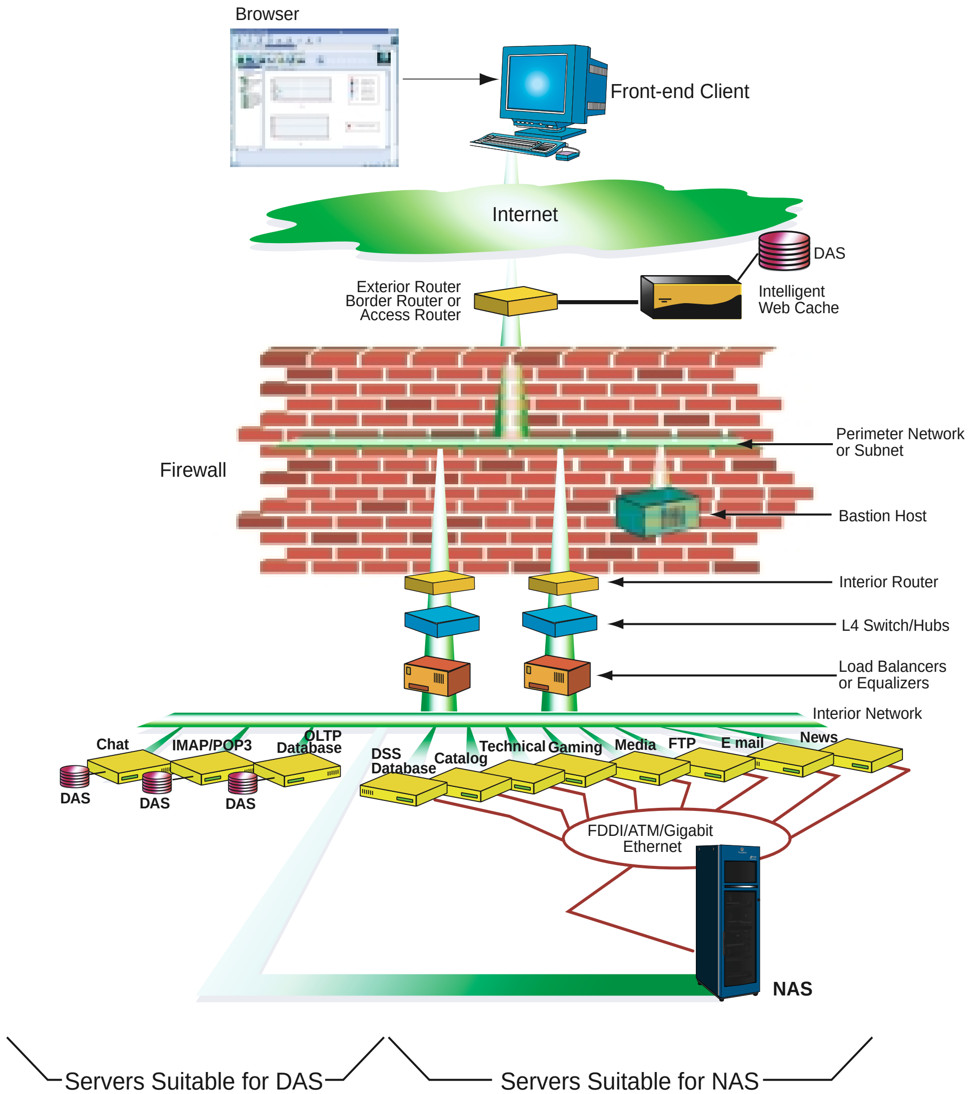

The back-end NAS servers are optimized for moving raw data between disks and networks, and removes the filesystems workload from the front-end servers to functionally balance overall Web Site performance.

- Border Routers route traffic to and from the Internet.
- Specialized front-end Intelligent Web Caches, such as the product series of Auspex partner Cache Flow Inc., handle many IP requests in front of the firewall and relieve unnecessary I/O activity from the firewall, L4 switches, servers and storage on the back-end networks.
- Routers and bastion hosts on the subnet in the firewall provide security.
- Switches hubs and load balancers provide redundancy and route traffic.
- Servers with Direct Attached Storage (DAS) may support heavily transaction oriented I/O workloads particularly those having a large percentage of writes.
- Optimized “thin O/S” Web servers improve reliability and availability by running only the minimum code necessary to accomplish the specific function of Web serving for non-transaction oriented workloads where I/Os tend to be large and data is referenced repeatedly.
- Back-end NAS servers consolidate data, improve availability, reduce backup windows and improve network and system manageability. The back-end NAS servers are optimized for moving raw data between disks and networks, and removes the filesystems workload from the front-end servers to functionally balance overall Web Site performance. The back-end NAS server provides a common access point and disk pool for Web server applications such as HTML pages, CGI scripts, mailspool files, and news.

With devices optimized only for a specific task and running only the minimum code necessary to accomplish these tasks, the overall Web Site architecture is not subject to uneven performance problems or unnecessary outages as is encountered in other Web Site architecture designs.

Auspex pioneered functional specialization in computer storage design

The *multilevel functionally* specialized Web Site architecture is a natural extension of the Auspex concept of ***functional specialization*** implicit in its Functional Multiprocessing (FMP) architecture. As an Auspex NS2000 separates and assigns processing resources to specific tasks within the server, a multilevel hierarchical Web server architecture provides many advantages over the traditional two level flat hierarchical architecture for all critical Web Site computing requirements:

A multi-level hierarchical Web server architecture provides many advantages over the traditional two level flat hierarchical architecture.

- ***24x7*** – The NAS server improves reliability/data availability, flexibility, speed and manageability by insulating file service requests from the general purpose operating system and will allow files to be served even in the event of an O/S failure.
- ***Flexibility*** – Storage can be easily added as the site grows. In addition storage can be dynamically allocated to where it is needed within the NAS system.
- ***Speed*** – The “Web cache” relieves many HTTP I/O requests from the back-end network and NAS server. Working together the “Web cache” and NAS server improve overall user response time for all requests.
- ***Manageability*** – The unique FMP design of the NS2000 improves manageability by offering both conventional Unix / NT system level tools and unique device specific tools.
- ***Cost effectiveness*** – Consolidating servers storage on NAS devices reduces capital acquisition and operating costs. Increased uptime results in fewer lost customers.

The importance of understanding Internet technology and terminology.

Those readers who already understand the basic terms and concepts of how the Internet works can skip Chapters 2, 3 and 4. The Internet is unique in the field of computer architecture since it includes a broad range of technologies, ranging from browsers, PCs, client server models, Local Area Networks (LANs), Wide Area Networks (WANs), many types of servers, and computer storage architecture. However it was felt that it was important to include an overviews of how the Internet works, since the Internet is still new to many consumers and computer professionals alike. Since 95% of all Web Sites today are being hosted on single computers and major growth is expected to come from firms under 100 employees and Web entrepreneurs, the vast majority of Web Site architecture designs are yet to come. More and more small companies and individuals will therefore be making Web Site architecture and Web Site design decisions in the next few years. Many of these persons are unfamiliar with Internet concepts, and incorrect Web Site architecture decisions can be very costly to a new or existing E-business ventures. Therefore, Chapters 2, 3, and 4 are recommended for all readers except Web professionals. Further definitions of all terms used in this chapter can be found in the Glossary.

A standard way of viewing the Front-End to Back-End Internet Model

For consistency throughout this paper we will use the concepts of front-end and back-end of the Internet. The front-end is also known as the “client-side” and the back-end is known as the “server-side.” Additionally all illustrations will show the front-end at the top of an illustration and the back-end at the bottom of the illustration.

The front-end is also known as the “client-side” and the back-end is known as the “server-side”.

Front-end in computer architecture traditionally refers to the part of a computer system that is closest to the user. In the case of the Internet the front-end refers to the PC or client computer with the browser that is usually either Microsoft’s Internet Explorer or Netscape’s

Navigator. Browser is short for Web browser, a software application used to locate and display Web pages. Both of these are graphical browsers, which means that they can display graphics as well as text. In addition, most modern browsers can present multimedia information, including sound and video, though they require plug-ins for some formats.

At the very back-end of the Internet is where the data resides that the front-end browser has requested. In the case of the hierarchical model presented in this paper the back-end is the Network Attached Storage Server and it's disks upon which the desired data resides.

The basic Internet model used in this report is divided into three parts as shown in **Figure 2**.

Figure 2 – The basic three part Internet model.

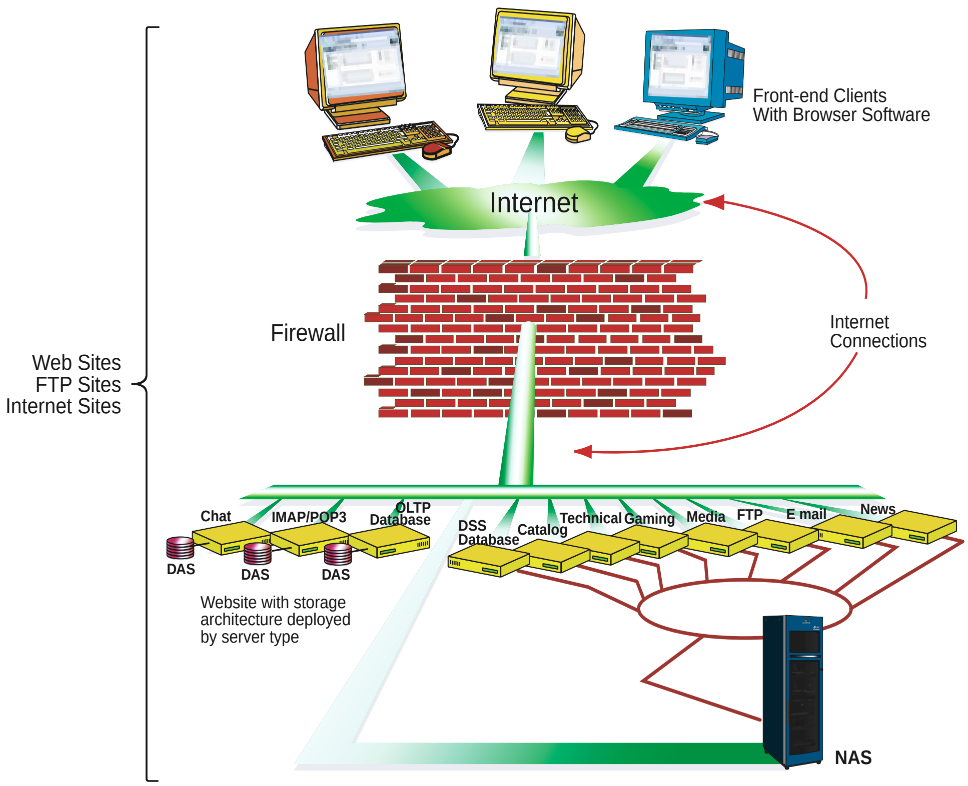

Other Information sources

This paper augments the concepts discussed in the Auspex report titled *A Storage Architecture Guide* which is available on the Auspex home page and discusses the four alternatives available in NAS products today. A thorough review of the Auspex Net Server 2000 Series is available in the Auspex report titled ***NS2000 Product Guide*** which can also be referenced from the Auspex home page at www.auspex.com. This report also contains a complete Glossary of Terms and many Figures and Tables that clearly explain the technical topics involved in large scale Internet site design.

How the front-end Web browser works

2

The World Wide Web is the most recent and exciting part of the Internet. It allows the viewing of multimedia Web pages comprised of graphics text, sound and video. The Web allows a client to jump from one place to another through the use of hypertext linked Web pages, called hyperlinks, with the click of a mouse. These links are embedded in a Web page. Hypertext Markup Language or HTML is the programming language that allows a client to view Web pages on the browser program of the computer.

The process of retrieving a Web page from the front-end to the back-end of the Internet is shown in **Figure 3.**

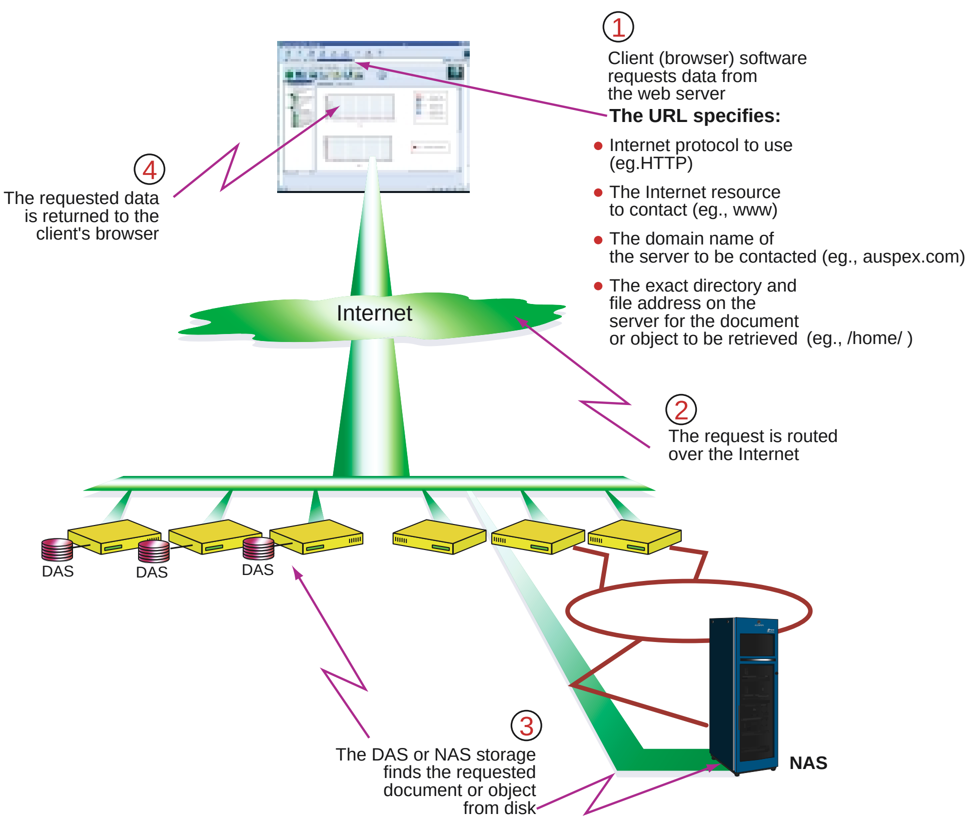

Figure 3 – How a browser retrieves a Web page from a Web server.

The Web works on a client/server architecture model in which browser software runs on a local computer known as the client. The browser software is at the very front-end of the Internet because this program controls what appears on the screen. Browser is short for Web browser software, which is used to locate and display Web pages. The two most popular browsers are Netscape Navigator and Microsoft Internet Explorer. Most modern browsers can present multimedia information, including sound and video, though they require plug-ins for some formats.

How URLs work

In a browser you type the URL or Uniform Resource Locator and click on the link that will send you to the desired location. The first part of a URL indicates which Internet protocol to use e.g., HTTP, FTP etc. The second part indicates what Internet resource to contact e.g. www is most common. The third part specifies the domain name of the Internet resource to be contacted and the fourth part contains the exact directory on the server, document or other Internet object to be retrieved. For example http://www.auspex.com will direct the browser program on a client computer to the World Wide Web resource of the Internet using the HTTP protocol, and then to the Auspex home page which contains hyperlinks to other information hosted on that Web Site.

The request is then routed over the Internet where Internet routers determine where to send the request. When the Web server receives the HTTP request, it finds the requested home page document or object and sends it back to the requesting Web browser client.

Domain Names

Domain names (e.g. www.auspex.com) are assigned and registered by firms such as InterNIC or ICANN which is a collaborative project between AT&T and Network Solutions (NSI) that is paid for by the National Science Foundation (NSF). Companies, organizations and individuals that register domain names pay an annual fee for the service. Since a domain name can be very important to frequency of site access, and an enterprise branding strategy, some individuals register domain names and then try to sell them to likely candidates, usually for a large profit.

Domain names are used in URLs to identify particular Web pages

Domain names are used in URLs to identify particular Web pages. For example, in the URL http://www.auspex.com/index.html, the domain name is auspex.com. Every domain name has a suffix that indicates which top-level domain name (TLD) it belongs to. There are only a limited number of such domains and all are preceded by dots as in .com, .net etc. **Table 1** shows some of the most common TLDs.

Table 1 – Common Top Level Domain (TLD) names.

TLD (Top Level Domain Name)	Domain Category
.gov	Government agencies
.edu	Educational institutions
.org	Organizations (nonprofit)
.mil	Military organizations
.com	Commercial business
.net	Network organizations
.ca	Canadian sites
.usps	United States Post Office
.vi	Virgin Islands

Domain Name Servers

DNS is an abbreviation for Domain Name System (or Service), an Internet service that translates domain names into IP addresses. A DNS server provides this service on the Internet. Because domain names are alphabetic, they're easier to remember. The Internet however, is really based on IP addresses. Every time you use a domain name, therefore, a DNS service must translate the name into the corresponding IP address. For example, the domain name www.example.com might translate to 198.105.232.4. The DNS system is, in fact, its own network. If one DNS server doesn't know how to translate a particular domain name, it asks another one, and so on, until the correct IP address is returned.

A DNS service must translate the name into the corresponding IP address.

How cookies work

Cookies are a feature optionally supported by some Web Browsers and they are standard with others. Cookies are on the forefront of Internet privacy discussions versus the Internet's ability to customize information for users. A cookie is a message given to a Web browser by a Web server that the user visits. The browser stores the message in a text file called cookie.txt. The message is then sent back to the server each time the browser requests a page from the server. The main purpose of cookies is to identify users and possibly prepare customized Web pages for them. When you enter a Web Site using cookies, you may be asked to fill out a form providing such information as your name and interests. This information is packaged into a cookie and sent to your Web browser, which stores it for later use. The next time you go to the same Web Site, your browser will send the cookie to the Web server. The server can use this information to present you with custom Web pages. So, for example, instead of seeing just a generic welcome page you might see a welcome page with your name on it. The name cookie derives from UNIX objects called magic cookies. These are tokens that are attached to a user or program and change depending on the areas entered by the user or program. Cookies are also sometimes called persistent cookies because they typically stay in the browser for long periods of time.

The name cookie derives from UNIX objects called magic cookies.

Connecting to the Internet

3

There are many ways for a client browser to connect to the Internet. The most common is for the browser to connect using plain old telephone service or POTS and a 56K modem. **Figure 4** illustrates how POTS and other Web connections work.

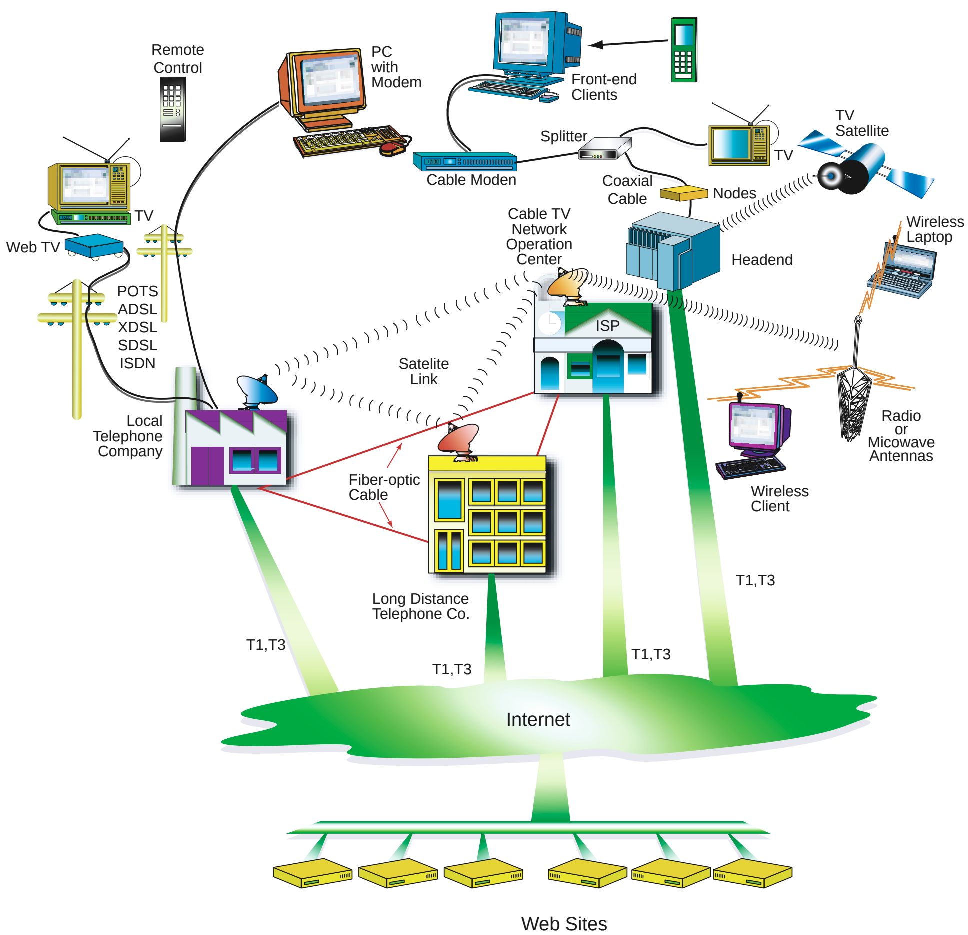

Figure 4 – Web connections can be accomplished through POTS, Web TV, ADSL, SDSL, XDSL, wireless, ISDN and other networks.

POTS

POTS or plain old telephone service refers to the standard telephone service that most homes use. In contrast, telephone services based on high-speed, digital communications lines, such as ISDN and FDDI, are not POTS. The main distinctions between POTS and non-POTS services are speed and bandwidth. POTS is generally restricted to a maximum of about 52 Kbps (52,000 bits per second).

Web TV

Web TV is a general term for a whole category of products and technologies that enable you to surf the Web on your TV set. Most WebTV products today consist of a small box that connects to your telephone line and television. It makes a connection to the Internet via your telephone service and then converts the downloaded Web pages to a format that can be displayed on your TV. These products also come with a remote control device so that you can navigate through the Web. A future class of WebTV products will not require telephone connections at all, but instead will access the Internet directly through the cable TV lines.

Cable Modems

A cable modem is a modem designed to operate over cable TV lines. Because the coaxial cable used by cable TV provides much greater bandwidth than telephone lines, a cable modem can be used to achieve extremely fast access to the World Wide Web. This, combined with the fact that millions of homes are already wired for cable TV, has made the cable modem something of a Holy Grail for Internet and cable TV companies. To install Internet access through a cable modem, a splitter divides the bandwidth from the coaxial cable coming from a cable supplier into the signal going to the high-speed cable modem and the client computer and the signal going to the TV set. This is shown in **Figure 4**.

There are a number of technical difficulties with cable modems, however.

There are a number of technical difficulties with cable modems, however. One is that the cable TV infrastructure is designed to broadcast TV signals in just one direction—from the cable TV company to people's homes. The Internet, however, is a two-way system where data also needs to flow from the client to the server. In addition, it is still unknown whether the cable TV networks can handle the traffic that would ensue if millions of users began using the system for Internet access. Despite these problems, cable modems that offer speeds up to 2 Mbps are already available in many areas.

ASDL

ADSL is short for asymmetric digital subscriber line; a new technology that allows more data to be sent over existing copper telephone lines (POTS). ADSL supports data rates of from 1.5 to 9 Mbps when receiving data (known as the downstream rate) and from 16 to 640 Kbps when sending data (known as the upstream rate). ADSL requires a special ADSL modem. ADSL is growing in popularity as more areas around the world gain access.

SDSL

SDSL is short for symmetric digital subscriber line, a new technology that allows more data to be sent over existing copper telephone lines (POTS). SDSL supports data rates up to 3 Mbps. SDSL works by sending digital pulses in the high-frequency area of telephone wires. Since these high frequencies are not used by normal voice communications, SDSL can operate simultaneously with voice connections over the same wires. SDSL requires a special SDSL modem. SDSL is called symmetric because it supports the same data rates for upstream and downstream traffic. A similar technology that supports different data rates for upstream and downstream data is called asymmetric digital subscriber line (ADSL). ADSL is more popular in North America, whereas SDSL is being developed primarily in Europe.

XDSL

Compared to POTS however, XDSL offers much higher speeds–up to 32 Mbps for downstream traffic, and from 32 Kbps to over 1 Mbps for upstream traffic.

XDSL refers collectively to all types of digital subscriber lines, the two main categories being ADSL and SDSL. Two other types of XDSL technologies are High-data-rate DSL (HDSL) and Single-line DSL (SDSL). DSL technologies use sophisticated modulation schemes to pack data onto copper wires. They are sometimes referred to as last-mile technologies because they are used only for connections from a telephone switching station to a home or office, not between switching stations. XDSL is similar to ISDN inasmuch as both operate over existing copper telephone lines (POTS) and both require "short runs" to a central telephone office usually less than 20,000 feet. Compared to POTS however, XDSL offers much higher speeds—up to 32 Mbps for downstream traffic, and from 32 Kbps to over 1 Mbps for upstream traffic.

Wireless Internet Connections

There are many types of Internet connections that are accomplished through cellular phone technology, radio technology and other wireless communications. Current transmission rates are approaching 128Kbps, although most wireless modems currently transfer at only 28Kbps. The client computer requires a wireless modem and antenna and the connection is made much like the client dialing out on a regular 56K [POTS] modem.

ISDN

ISDN or integrated services digital network, is an international communications standard for sending voice, video, and data over digital telephone lines or normal telephone wires. ISDN supports data transfer rates of 64 Kbps (64,000 bits per second). Most ISDN lines offered by telephone companies give you two lines at once, called B channels. One line is used for voice and the other for data, or both lines can be used for data to provide data rates of 128 Kbps, three times the data rate provided by today's fastest modems. The original version of ISDN uses baseband transmission. A modern version of ISDN, called B-ISDN, uses broadband transmission and is able to support transmission rates of 1.5 Mbps. B-ISDN however, requires fiber optic cables and is not widely available.

Other network connections

Networks in offices and enterprises are connected to the Internet by a variety of technologies including some of those discussed in this chapter. Other connection links are T1 and T3 lines which are more often used in the high speed Internet infrastructure itself. T1 and T3 lines are the most frequent type of attachment to the Internet for large scale private and public sites.

The Internet itself is a loose collection of Networks 4

The Internet is a loose organization of networks paid for by many organizations, government agencies and private organizations. All of these networks work together in a loosely organized alliance with members ranging from small private networks to commercial online services such as AOL and CompuServe and private Internet Service Providers (ISPs) that provide the ability to connect to the Internet for a fee.

A particularly important characteristic of the Internet is its ability to support Virtual Private Networks or VPNs. A VPN is often referred to as an Intranet or a Web Site that restricts access only to authorized employees or individuals. A VPN is a network that is constructed like any other Web Site by using public wires to connect nodes. There are a number of systems that enable you to create VPNs using the Internet as the medium for transporting data. These systems use encryption and other security mechanisms to ensure that only authorized users can access the network and that the data cannot be intercepted.

A VPN is a network that is constructed like any other Web Site by using public wires to connect nodes.

Backbones

As shown in **Figure 5,** the main network connections composing the Internet are known as backbones. Backbones are connected by routers which examine IP envelope addresses and send them on their way to the correct destination. Some backbones are paid for by the government agencies such as the National Science Foundation (NSF) or the National Aeronautics and Space Administration (NASA) that links sites across the country. Backbones connect Network Access Points (NAPS) throughout the world and many are based on ATM OC-12 at 622Mbps. A future higher speed backbone known as vBNS (very high-speed Backbone Network Service) which is in limited deployment at this time and runs at 2.5 Gbps. Local networks and Web Sites are connected to the Internet in a variety of ways. As discussed in Chapter 3, a residential telephone line can transmit data at 56Kbps (kilobits per second), and higher bandwidth leased telephone lines can carry data at 1.544Mbps (megabits per second) for T1 links and 44.746Mbps for T3 links.

Backbones connect Network Access Points (NAPS) throughout the world.

Figure 5 – The Internet is a loose organization of public and private networks comprised of backbones, routers, hubs, switches, software protocols and naming conventions.

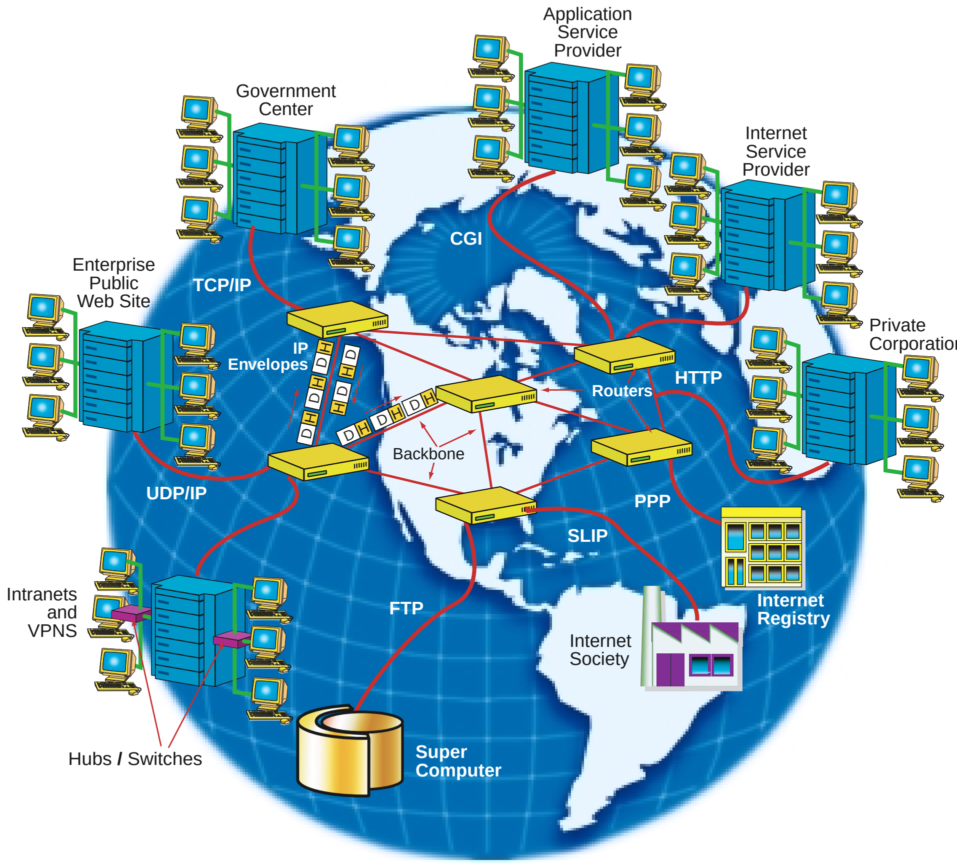

Internet

Regional networks

Regional networks maintain and provide Internet access within a certain geographic area, with a worldwide presence.

Routers

Routers connect backbones and do most of the work directing traffic over the Internet.

Routers connect backbones and do most of the work directing traffic over the Internet. When TCP creates a packet it adds a header to the data which defines how the packets are to be reassembled at the other end. It also adds a checksum to determine if any errors occur during transmission. The packets are then placed into an IP envelope, which adds another type of header that contains information about the address to which the packet is to be sent. Using standard protocols such as Routing Information Protocol (RIP) for interior routers and Exterior Gateway Protocol (EGP) for exterior routers, the router examines the header for each IP envelope to determine the exact IP address to which the envelope is to be sent.

The Internet Society

The Internet Society is a nonprofit organization that guides the direction and growth of the Internet and manages its architecture by recommending how Internet protocols should work on the Internet.

Internet Protocols

The Internet supports a wide variety of standard protocols for sharing data and messages of different types between the millions of nodes and networks comprising the Internet. The Internet is known as a packet switched network. This means that all information is broken down into pieces called packets or datagrams and sent across the Internet from one computer to another. Packets are very small and less than 1500 characters so a large message can involve the transmission of many packets.

OSI Model

This process is governed by the OSI model, which defines the different layers through which a message passes from one computer to another. OSI stands for Open System Interconnection, an ISO standard for worldwide communications that defines a networking framework for implementing protocols in seven layers. Control is passed from one layer to the next, starting at the application layer in one station, proceeding to the bottom layer, over the channel to the next station and back up the hierarchy. Although the OSI model was never fully implemented as a standard, it defines the functionality through which all messages pass in the communications process. Most of the functionality in the OSI model exists in all communications systems, although two or three OSI layers may be incorporated into one. **Table 2** shows the seven layers of the OSI model

In the OSI model, control is passed from one layer to the next.

Level or Layer	Name	Function
7	Application Layer	Program-to-program communication.
6	Presentation Layer	Manages data representation conversions. For example, the Presentation Layer would be responsible for converting data from EBCDIC (IBM mainframes) to ASCII (UNIX and Windows NT).
5	Session Layer	Responsible for establishing and maintaining communications channels. In practice, this layer is often combined with the Transport Layer.
4	Transport Layer	Responsible for end-to-end integrity of data transmission.
3	Network Layer	Routes data from one node to another.
2	Data Link Layer	Responsible for physically passing data from one node to another.
1	Physical Layer	Manages putting data onto the network media and taking the data off.

Table 2 – OSI Reference Model structures communication software in layers.

TCP and IP are the two most important communications protocols on the Internet.

TCP/IP

Transmission Control Protocol (TCP) breaks down and reassembles packets that are sent over the Internet, and Internet Protocol (IP) makes sure that packets are sent to the right destination. TCP and IP therefore are the two most important communications protocols on the Internet. They are frequently referred to as TCP/IP. Packets can travel over the Internet on separate routes and sometimes can be lost if they arrive at an Internet router when its queue is full or if the "hop count" exceeds the maximum. When this happens TCP will keep re-sending packets until they are received by the destination computer. "Packet loss" is one of the biggest reasons for delays in message transmission. The software that is required for a computer to understand and execute TCP/IP protocol is known as a socket or TCP/IP stack. This software acts as an intermediary between the Internet and a computer.

HTTP is an abbreviation for HyperText Transfer Protocol, the underlying protocol used by the World Wide Web. HTTP defines how messages are formatted and transmitted.

HTTP

HTTP is an abbreviation for HyperText Transfer Protocol, the underlying protocol used by the World Wide Web. HTTP defines how messages are formatted and transmitted, and what action Web servers and browsers should take in response to various commands. For example, when you enter a URL in your Web browser, this actually sends an HTTP command to the Web server directing it to fetch and transmit the requested Web page.

UDP/IP

User Datagram Protocol (UDP) is used instead of TCP for certain types of data such as live video and audio data. UDP will not resend packets if they are lost since the packets would be out of synchronization at the other end with the audio or video being received at the destination computer.

FTP

The Internet's File Transfer Protocol or FTP is the most common protocol used for downloading files to your computer or uploading files to another computer. The files can vary from programs that can run on your computer to graphics you can view or sounds and music you can hear. Special software is required to connect to an FTP server, which is available from many places on the Internet.

SLIP and PPP

Serial Line Internet Protocol (SLIP) or Point-to-point (PPP) protocol are methods of connecting a computer to the Internet via a serial port and a modem when a network card is not used.

When a visitor to the Web Site clicks on the URL the CGI program is launched.

CGI

Common Gateway Interface or CGI is a communications protocol that enables a Web server to communicate with other applications especially those that trigger access to databases by a Web browser. When visiting a Web Site, certain links offer users the opportunity to search a database for specific information. These types of databases are called search engines. Examples are Yahoo, Excite, or mapping applications with various clickable parts and interactive forms. As shown in **Figure 6**, CGI applications are written at the Web Site and posted to a special directory on the Web server with a link to it embedded in a URL on a Web page. When a visitor to the Web Site clicks on the URL the CGI program is launched. The program sends the requesting browser a Web page in HTML format that is then filled out detailing what the client wants to find. When the data is completely filled in, it is sent to the CGI program which activates the database program to retrieve the requested information. The database program then retrieves the requested data and passes it to the CGI program which in turn passes it to the requesting client in HTML format. The CGI program is therefore like a fancy hyperlink triggering a database search or other application instead of returning a Web page or file to the client. Leading edge E-commerce firms such as Verity, Inc. utilize "fuzzy" search and retrieval technology to increase the browser-to-buyer conversion ratio for catalog searches in many industries.

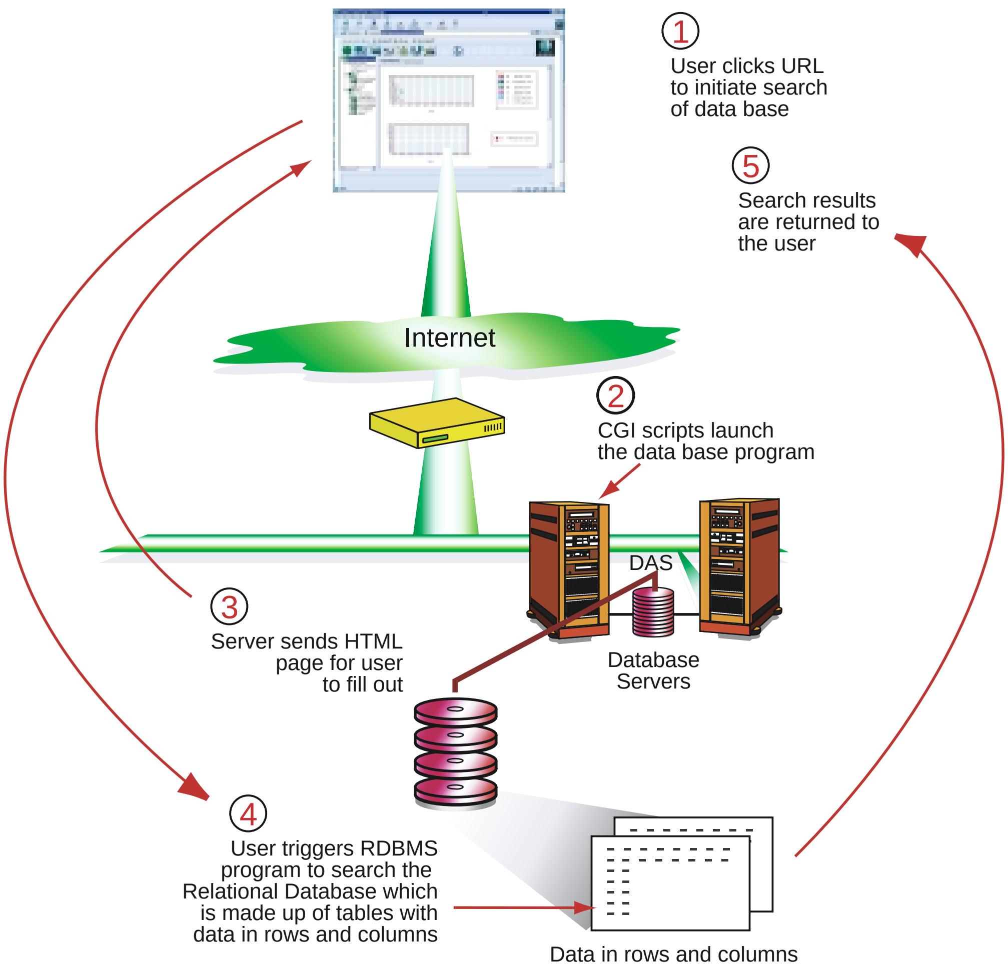

Figure 6 – How CGI programs enable database searches over the Web.

File types on the Internet

There are two types of files commonly sent over the Internet, ASCII and binary files. ASCII is a code system for representing characters as numbers. Binary files represent data using the unique digits 0 through 9, combining these ten digits can represent all other numbers.

In addition files with specific formatting and layout parameters can be packaged in IP envelopes and sent over the Internet to other computers. The other computer must have a program called a “reader” that allows the particular file format to be displayed on the screen or to activate video cards, sound cards or other devices.

Table 3 below lists popular file types that can be sent over the Internet.

Table 3 – Popular file formats sent over the Internet.

File type	File extension	Definition
BMP	.bmp	Standard bitmap Windows graphics format
Excel	.xls	Microsoft Excel file format
EXE	.exe	Executes programs
GIF	.gif	Standard graphics file for the World Wide Web
JPEG	.jpg	Compression standard for color images
MIME	.mim	A standard for non-ASCII messages that allows them to be sent over the Internet
PCX	.pcx	Graphics file format for PCs, scanners, fax
PDF	.pdf	Adobe’s Portable Document Format
PNG	.png	A license free graphics format similar to GIF
PKZIP	.zip	Popular shareware compression technique
Powerpoint	.ppt	Microsoft PowerPoint file format
TIFF	.tif	A bitmap graphics format for Apple and PCs
WAV	.wav	Microsoft standard for sound files
Word	.doc	Microsoft Word file format

A comprehensive online dictionary of file formats can be found at www.stack.com.

Anatomy of a Web Site

5

Whether they are publicly accessible or privately accessible, Web Sites need to be protected from security attacks from the Internet. Security therefore is the primary concern of the front-end of the Web Site. Since Web Sites are connected to internal corporate or government networks and accessible through the local area network (LAN), they are critical to the ongoing business of the enterprise. They also need to exhibit the highest levels of availability (7x24) and therefore fault tolerance in design is a major consideration. The Anatomy of a Web Site is shown in **Figure** 7 including the many different types of WebServers each running different protocols with different I/O workloads. These differences are very important in optimizing the performance and reliability of a Site and are discussed further in this report.

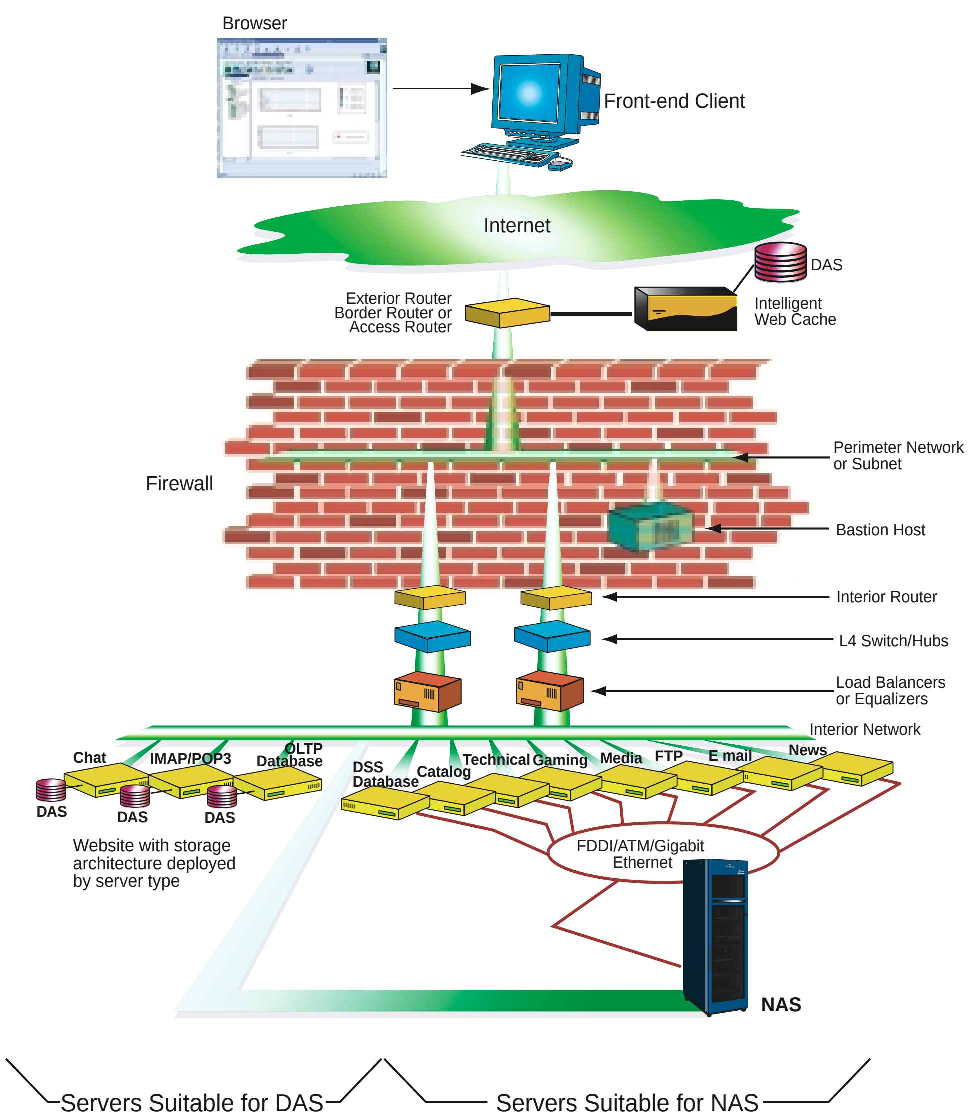

Figure 7 – A Web Site must be secure, available 24x7, fast and easily manageable.

Border Routers

A specialized type of router that operates at the Network Layer (**Table 2**) of the OSI model and is configured at the border of a network and the Internet is called a Border Router, Exterior Router or Access Router. This router screens packets between the Internet and the perimeter network. It adds redundancy to Web Site security by screening incoming packets from the Internet using the same rules as the internal screening router and protecting the network even if the internal router fails.

Gateways

A gateway is a combination of hardware and software that links two different types of networks. An email gateway for example allows users on two different types of email systems to exchange emails.

Firewall

A firewall is a collection of computers and routers on a special network between an organization's network and the Internet that shields the internal network from attack from the Internet. On the Internet boundary, an exterior screening router screens packets between the Internet and a perimeter network known as the subnet which is the first line of defense against attacks from the Internet. The exterior router uses the same rules as the interior router, or choke router in the event the interior router fails.

Bastion Host

The choke router is the first line of defense for packets going out from the internal network and the last line of defense for packets coming into the internal network.

A bastion host in the firewall is the primary point of contact for connections coming in from the Internet for services such as email and receiving data from the internal FTP site. The Bastion host is a heavily protected server with many layers of security software built in and it is the only contact point for incoming Internet requests. Bastion hosts can also be set up as proxy servers. A proxy server processes any requests from the internal corporate network to the Internet, such as browsing the Web or downloading files via FTP. The first line of defense for packets going out from the internal network and the last line of defense for packets coming into the internal network is the interior router, or choke router. This router examines all packets traveling in both directions and decides which packets can be sent to and from specific Internet locations if those locations have been found to be suspicious. System and network administrators set the rules for determining which packets to allow and which to block in either direction.

Bridges and Hubs

A bridge is a device that operates at the Data Link Layer of the OSI model (**Table 2**), and connects two local-area networks (LANs), or two segments of the same LAN. The two LANs being connected can be alike or dissimilar. For example, a bridge can connect an Ethernet with a Token-Ring network. Unlike routers, bridges are protocol independent. They simply forward packets without analyzing and rerouting messages. Consequently, they're faster than routers, but also less versatile.

Hubs are commonly used to connect segments of a LAN.

A hub is a common connection point for devices in a network. Hubs are commonly used to connect segments of a LAN and contain multiple ports. When a packet arrives at one port, it is copied to the other ports so that all segments of the LAN can see all packets.

Load Balancers

Load balancers can prevent servers from being overloaded by distributing large numbers of access requests across multiple Internet servers. The requests are directed to the most-available server based on how busy that server is. Other load balancers take the form of special software that runs on a dedicated server, a proprietary black box, or a capability in an advanced LAN switch. Not only can site administrators configure their load balancers to route requests based on server use, but requests can also be routed based on server content such as video, audio, FTP, text etc.

FTP Servers

FTP servers are usually optimized to send large files via the FTP protocol. This could take the form of financial documents, articles, publications or other files.

HTTP Servers

HTTP servers are the work horses of many Web Sites since there are more HTTP requests on the Internet than any other protocol type. These servers form the backbone of the E-commerce Web Sites and most product presentations are accessed in HTTP. Also most corporate Web Sites receive HTTP requests for pages from their site.

Chat Servers

Chat servers use the IRC format for message exchange. IRC stands for Internet Relay Chat, a chat system developed by Jarkko Oikarinen in Finland in the late 1980s. IRC has become very popular as more users get connected to the Internet because it enables users connected anywhere on the Internet to join in live discussions. Unlike older chat systems, IRC is not limited to just two participants. To join an IRC discussion, you need an IRC client and Internet access. The IRC client is a program that runs on your computer and sends and receives messages to and from an IRC server. The IRC server, in turn, is responsible for making sure that all messages are broadcast to everyone participating in a discussion. There can be many discussions going on at once; each one is assigned a unique channel.

IRC is not limited to just two participants.

Email Servers

An email server runs the Simple Mail Transfer Protocol (SMTP) and is used exclusively for email traffic. Email servers experience very high I/O rates of varying I/O sizes because of the many pictures and files being attached to email messages. For this reason special algorithms are often applied to emails to delay the processing of an email message if it is very large. Some email servers compress attachments or send them at off-peak hours when more bandwidth is available. Email servers use Post Office Protocol (POP) or Internet Message Access Protocol (IMAP) to check addresses.

Special algorithms are often applied to emails to delay the processing of an email message if it is very large.

Database Servers

A database server[2] enables a client to communicate with database applications via a Web browser. Based on the CGI discussed in Chapter 4, database servers or search engines offer users the opportunity to search a database for specific information. Examples of these types of databases are the search engines such as Yahoo or Excite, mapping applications with

[2]Database servers typically run one of the major Relational Database Management Programs (RDBMS) such as are available from Oracle, Informix, Microsoft (SQL Server), and others.

various clickable parts and interactive forms. At the Web Site, applications are written and posted to a special directory on the Web server and a link to it is embedded in a URL on a Web page. When a visitor to the Web Site clicks on the URL the program in launched. The program sends the requesting browser a Web page in HTML format that is then filled out detailing what the client wants to find. When the data is completely filled in, it is sent to the CGI program which activates the database program to retrieve the requested information. The database program then retrieves the requested data and passes it to the CGI program which in turn passes it to the requesting client in HTML format. The CGI program is therefore like a fancy hyperlink triggering a database search or other application instead of returning a Web page or file to the client.

Components of an effective Web Site Architecture 6

A flat clustered Web Site architecture

There are two approaches to scaling an Internet Site. The first is a flat clustered architecture (**Figure 8**), and the second is a functionally specific hierarchical architecture (**Figure 9**). The flat clustered architecture has been systematically rejected by those who have tried to make serious use of it, and it is discussed here to illustrate the reasons for this rejection.

Since Web hyperlinks shield the user from cryptic-looking names and addresses, designers have suggested that when a Web Site grows beyond a single machine's capacity to serve its pages, a second, third, nth similar server be added. This results in the flat clustered architecture as shown in **Figure 8.** Flat topology has "n" essentially independent servers each with data critical to its application. One variation on the theme calls for some servers to replicate Web pages of other servers within the set, for the sake of additional request-handling capacity or fault tolerance. Another variation gives the illusion of a pool of shared disk data; it calls for the servers to cross-mount each other's filesystems, so that they can process the same set of requests without actually duplicating the data. In both variations, critical data and hot data is scattered to the many Direct Attached Storage (DAS) disks of the Web servers in the server farm.

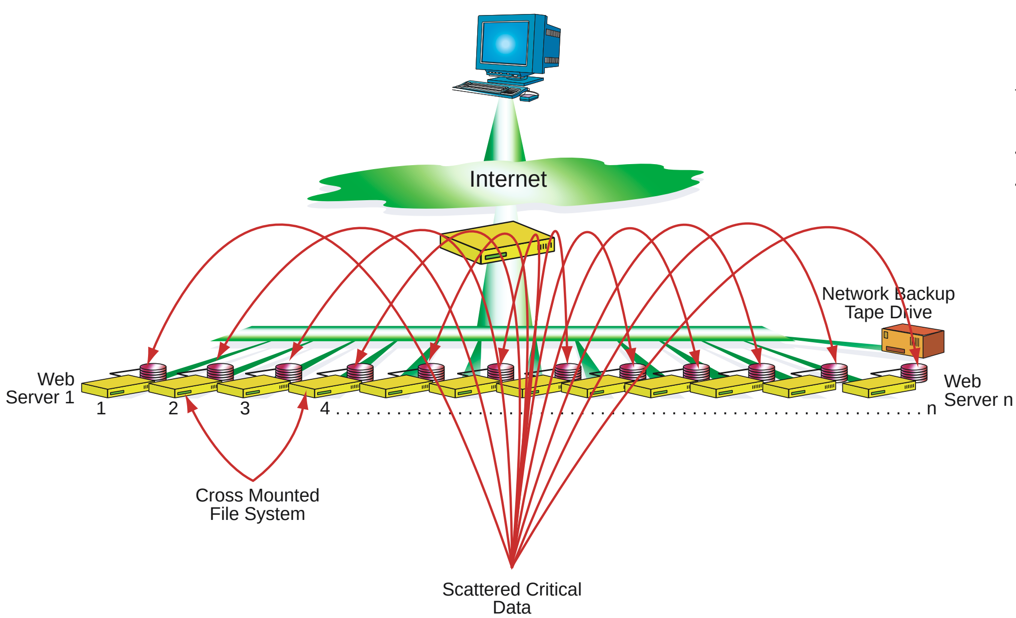

Figure 8 – A flat clustered Web Server architecture results in scattered critical data and scattered hot and cold data on the same disks.

Frequently reported problems with a flat clustered Web Site design

To get a preview of some of the pitfalls of the flat clustered architecture, consider just some of the most often reported problems with a flat clustered design. A basic reasons for these problems, which are discussed below, is that individual servers are asked to do more than they are designed to do. The design goal for workstations is usually fast computation. This results in a net increase in the total count of systems beyond what would be required in a hierarchical approach where front-end servers with DAS run appropriate DAS applications and servers with NAS run applications optimized for NAS. Frequent problems with this approach are as follows:

Reliability is decreased as each server's working set increases to service both IP request and filesystem processing. Even worse, if cross-mounted file systems are used, system interdependencies increase and with that, the number of failure modes goes up.

There is no single point of disk backup. Multiple tape drives or a network-based backup scheme must be used for a flat clustered architecture. Network-based backup can be especially onerous to a World Wide Web Site. For many global businesses running applications twenty-four hours a day, some time zone is always in prime time. Therefore it's hard to pick a good time to steal network cycles away from webserving to backup data on a tape subsystem.

Disk utilization is inefficient, since there is no truly shareable pool.

Performance suffers since load-balancing techniques are less easily applied. This is because there is no clean way to divorce filesystem processing from IP-request processing.

A multilevel functionally specific hierarchical architecture

A multilevel functionally specific Web Site architecture is shown in **Figure 9** and is in service today at multiple Auspex sites. It derives its name from the fact that IP requests received from clients are satisfied by a hierarchy of servers each optimized and load balanced for an appropriate protocol with DAS or NAS storage tailored to what is best for the application. Sites that have adopted this architecture typically cite five key benefits from the hierarchical approach:

Flexibility (scalability)

7x24 (availability)

Speed (performance or fast response time to users)

Manageability (and effective backup)

Cost effectiveness.

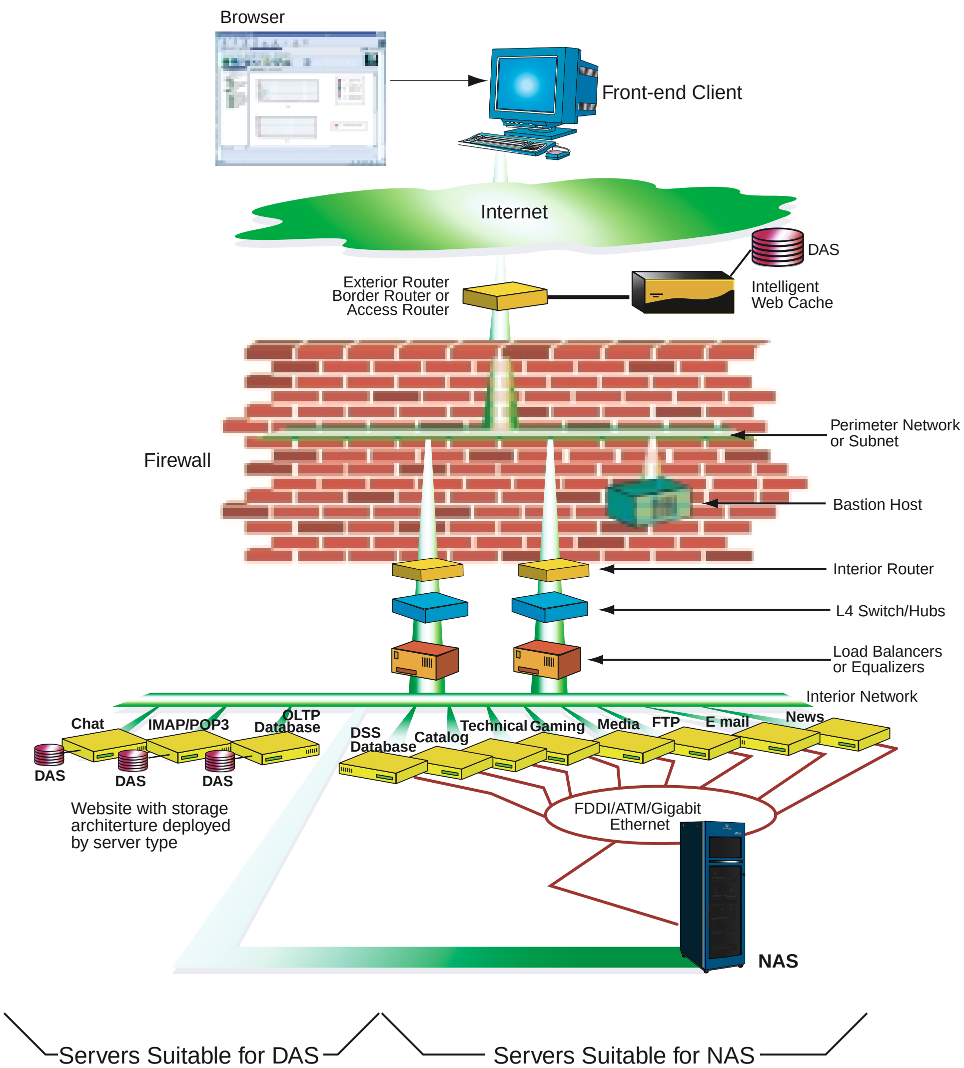

Figure 9 – A multi-level functionally specialized hierarchical Web Site architecture provides an optimal balance of 24x7, flexibility, speed and manageability.

The secrets to designing an optimized hierarchical Web Site

The trick is to know when to use DAS and when to use NAS.

The trick to designing a Web Site with many servers and many types of servers is to know which servers to deploy with DAS (Direct Attached Storage) and which servers to deploy with NAS (Network Attached Storage) servers such as the Auspex NS2000. When servers are deployed with NAS, thin O/S Web Servers, and Intelligent Web Caches (such as those provided by Auspex partner Cache Flow) are recommended for certain configurations.

With devices optimized only for a specific task and running only the minimum code necessary to accomplish these tasks, the overall Web Site architecture is not subject to uneven performance problems or unnecessary outages as is encountered in other Web Site architecture designs.

- Border Routers route traffic to and from the Internet and provide the first level of security as discussed in Chapter 5.
- Specialized front-end Intelligent Web Caches, such as the product series of Auspex partner Cache Flow Inc., handle many IP requests in front of the firewall and relieve unnecessary I/O activity from the firewall, L4 switches, servers and storage on the back-end networks.
- Routers and bastion hosts on the subnet in the firewall provide security.

- Switches, hubs and load balancers provide redundancy and route traffic so as prevent bottlenecks.
- Servers with Direct Attached Storage (DAS) support transaction oriented small I/O workloads such as mail, news and chat where data is not referenced repeatedly.
- Optimized "thin O/S" Web servers improve reliability and availability by running only the minimum code necessary to accomplish the specific function of Web serving for non-transaction oriented workloads where I/Os tend to be large and data is referenced repeatedly.
- Back-end NAS servers consolidate data, improve availability, reduce backup windows and improve network and system manageability. The back-end NAS servers are optimized for moving raw data between disks and networks, and removes the filesystems workload from the front-end servers to functionally balance overall Web Site performance. The back-end NAS server provides a common access point and disk pool for example HTML pages, CGI scripts, mailspool files, and news.

Back-end NAS servers consolidate data, improve availability, reduce backup windows and improve network and system manageability.

With devices optimized only for a specific task and running only the minimum code necessary to accomplish these tasks, the overall Web Site architecture is not subject to uneven performance problems or unnecessary outages as is encountered in other Web Site architecture designs.

Deciding when to deploy a server with DAS or NAS

In **Figure 9** certain types of servers are shown with Direct Attached Storage (DAS) storage attached and certain types of servers are shown with NAS back-end servers. There is a clear division of responsibility between the front-end servers and the back-end server.

Server types appropriate for DAS deployment

In general Direct Attached Storage (DAS) is appropriate for server deployment where the server I/O workload is transaction oriented with little rereading of the same files as in database OLTP (Online Transaction Processing) applications.

- E-commerce On-line Transaction Servers
- Chat servers
- IMAP/POP3 servers
- OLTP Database servers

The I/O workload pattern from these servers to their storage subsystems tends to be characterized by high I/O rates of small sized I/Os with frequent writes to storage.

The I/O workload pattern from these servers to their storage subsystems tends to be characterized by high I/O rates of small sized I/Os with frequent writes to storage.

Server types appropriate for NAS deployment

In contrast to workloads appropriate to DAS, those that are appropriate for deployment on NAS tend to be characterized by large file sizes, low write activity, and rereading of the same files. In a separate report titled ***A Storage Architecture Guide*** appropriate applications for DAS and NAS are discussed when deploying NAS in an enterprise data center. This report is available at www.auspex.com.

Many of the same applications have become geographically distributed over Intranets and the Internet, and map directly into the type of applications appropriate for NAS in a large scale Web Site architecture. They are:

- DSS database servers
- E-commerce product catalog servers
- Email servers
- Regionally distributed technical and scientific applications such as MCAD and ECAD
- News servers

- Streaming audio and video media
- Banner and graphics servers
- Oil and gas development applications
- Unix and NT file sharing applications
- Online Gaming
- FTP Servers
- Software development
- Mapping databases

Workloads that are appropriate for deployment on NAS tend to be characterized by large file sizes, low write activity, and many re-reading of the same files.

Additionally, if a front-end server must run some non-Internet ancillary application (e.g., doing user accounting), site administrators may likewise choose to store the data on a NAS back-end.

Front-end redundant routers, E-commerce servers, and back-end network paths

For redundancy and extra bandwidth, at least *two* routers should connect to the Internet as shown in Figure 1. A dedicated secure commerce server is redundantly configured separately from the more generic "Web servers. This will provide increased availability (7x24) and security.

A *redundant* path to the NS2000 should be available. Front-end servers normally take Path A to the NetServer to access their data. However, Path B into the NetServer guarantees all parties can communicate if a failure develops on the back-end network.

Though these redundant embellishments may be considered overkill in certain environments, no Web Site managers who have lived through an outage would argue with justifying the extra measures of redundancy for 7x24 service.

Scaling

Separation of IP-request processing from filesystem processing significantly contributes to the scalability of the hierarchical architecture by reducing the work that the Web Server has to perform. NAS servers such as the Auspex NS2000 are ideal for scalability and the reader is referred to the ***NS2000 Product Guide*** which is available from www.auspex.com.

Assigning and Balancing Web Loads – Load Balancers

The hierarchical Web Site architecture typically focuses all file system loads on one consolidated server (perhaps more, if multiple fail-over back-ends are used or storage capacity requirements are extraordinarily high). The architecture requires some strategy by which to route incoming IP-request workloads to one of the front-end servers. This section discusses the pros and cons of several techniques.

- **Static Assignment—** Although this method is simple, it has gained a following, especially among ISPs. Via the domain names made known to requesting clients, user request loads are directed to particular front-ends. In other words, a static assignment of function to server is made. However, there are shortfalls to this approach: Although it balances access for Web Pages access because the webmaster has easy access to all HTML code, the technique doesn't translate to services like mail and news.
- **Dynamic Top-Level Webpage Technique—** If clients bookmark pages subordinate to the top-level "switching" page, when they subsequently access the site via the bookmarks, the resultant bypass of the top-level page undermines the round-robin effect and provides no recourse in the event of a down front-end server.
- **HTTP Redirection—** HTTP redirection is similar to the dynamic top-level webpage technique described in the previous section, but it is implemented differently. Rather

than return a dynamically crafted webpage of URLs to the client, the master front-end server takes advantage of a feature of the HTTP protocol. Specifically, it responds to a request with a status code commanding the client to automatically retry the request (which must be a GET or HEAD) to a different front-end. With the cooperation of the client's browser, the second operation lands on the front-end server destined to actually fulfill the request. Some have compared this user-transparent behavior to the call forwarding offered by a phone company.

By choice of status code returned, the master server can indicate that the resource requested has been permanently or temporarily moved to another front-end server, the location of which needs obviously to be returned by the server as part of the initial redirect response. The receipt of a "moved permanently" redirection status code has the advantage of allowing the client to bypass redirection in subsequent accesses, with the attendant disadvantage that the particular client's load has now been statically assigned to a particular resource server. The receipt of a "moved temporary" status code affords more granular dynamic load balancing, since clients keep returning to the master server to learn the current location of the desired resource.

- **Client-based Round-Robining—**Consider the case where a webmaster has a) complete control of which browser is used by all clients (as one might in an Intranet environment), b) has access to that browser's source, and c) can predict the addresses of a particular set of front-end servers. In this case the webmaster might choose client-based round-robining for load balancing. It is accomplished by coding the browser to accept something like http://www.specialsite.com from the user and then surreptitiously rewrite the URL to http://www.specialsiteN.com, where N varies randomly over some range. This technique only load-balances from clients that agree to use a particular browser to access a particular supersite.

  Netscape "popularized" this technique when it employed it to balance load on a very popular public Web Site (www.netscape.com). Its effectiveness hinged on the obvious—a large percentage of Internet clients could be counted upon to run Netscape Navigator! Unfortunately Navigator offers no hook by which another set of front-end servers may be specified and most sites are not ready to undertake custom-client writing.
- **DNS Round-Robining**—This technique is very commonly cited in load-balancing discussions. A discussed in Chapter 2, the Internet's Domain Name System's (DNS) name servers translate text server names (e.g., www.companyname.com) into numeric Internet addresses (e.g., 198.95.224.2). This dot-delimited address enables the underlying IP software and hardware to route a client's packets to their rightful destination. The challenge is to have clients transparently direct IP requests to one location (e.g., www.companyname.com), and yet have the request satisfied by one of the front-end servers, each having its own unique numeric IP address. The "round-robin" features of the Berkeley Internet Name Domain (BIND) software accomplishes this.
- **How Round-Robining Works—** On the site's name server(s), the administrator defines www.companyname.com with multiple "A records," each "A record" defines the IP address of a different front-end server (these all have access to identical WebPages, mounted off the Auspex NS2000 File Server). BIND returns all the IP addresses, not just one, to the client requesting name resolution. One benefit in receiving a list is that an intelligent client can choose an alternative IP address if the first is down or fails over time. As long as all the front-end servers are healthy, websurfing clients will target the first server on the list. When the administrator turns on BIND's round-robin feature, with each query to the name server, the order of the addresses is shifted, so that a different one from the set is always "on top." Since it's based purely on the average number of outstanding requests, the mechanism is both simple (reliable) and efficient. An additional advantage is that scaling is easy: just add another IP address ("A" record) to the name server's list.

Positive experiences have been reported by system administrators, who uses DNS round-robining for POP3 mail, news and Web front-end servers. Extensive metering (of all kinds, not just performance) indicates that the front-ends are largely well balanced. If DNS caching is perturbing the balance, it is at most on the order of 5-10%.

Fault Tolerance

Fault tolerance is often a design criterion for the strategic Web Site. Redundancies are built-in where desirable, feasible, and cost-justifiable. An important first step to fault tolerance is redundant communications—alternate paths into the site (via two T3 lines or other high-speed connection) and alternative paths to the back-end NS2000.

With multiple front-end servers configured, it's natural to think about the possibility of surviving on N-1 front-ends when one fails (or needs to be upgraded or replaced). There are many load balancers that can eliminate the need for special fault-tolerant programming on the part of the local administrators. As long as one server in a load balancers group remains up, and when a load balancer detects an "infinite response time" from a member server, the load balancer will automatically reapportion the load that would otherwise have gone to the non-responding front-end server to the surviving server(s). Other load-balancing approaches require custom programming for fault tolerance.

Tuning individual servers in an Internet site

Because IIS is tightly integrated with the operating system, it is relatively easy to administer.

There are many software products available for configuring a server as a Web Server or Internet Server. An example is Microsoft's Internet Information Services (IIS) 5.0 which runs on Windows 2000 Server. Because IIS is tightly integrated with the operating system, it is relatively easy to administer. However, IIS is currently available only for the Windows NT platform, whereas Netscape's Web servers run on all major platforms, including Windows NT, OS/2 and UNIX.

The product guides and manuals for products such as IIS cover many tips and tricks from the vendor as to how to optimally tune an individual server in terms of all hardware and Operations system settings. There are usually robust technical papers accompanying a Web site of either software or hardware vendors for the Internet Web serving applications. An example of one such paper for Microsoft's IIS can be downloaded in pdf format from: http://www.microsoft.com/windows2000/library/operations/Web/tuning.asp

The Auspex NetServer 2000 deployed in Internet sites

7

Network Attached Storage (NAS) has migrated to the Internet

Network Attached Storage (NAS) evolved from the networking industry and was first pioneered by Auspex. The NAS concept migrated from LAN networks within enterprises to Web based WAN networks such as Intranets, Extranets and the Internet itself. In these sites, NAS servers such as the Auspex NS2000 coexist with Direct Attached Storage (DAS) as back-end storage for Web servers attached to the Internet.

With networking, there are strong standards for connectivity, data security and load balancing. In the Auspex NS2000 Series, the file system resides in the NS2000 as opposed to Web Servers with DAS where the file system competes for resources with the application server's CPU. Both Web Servers with DAS or NAS are appropriate for an enterprise depending on the application being supported. NAS is the best choice for UNIX and Windows NT data sharing applications, audio and video streaming video applications, on-line gaming applications, consolidated file serving applications, technical and scientific applications, Internet and Intranet application and certain types of Decision Support (DSS) applications. In particular the NS2000 has been shown to be highly effective for FTP serving since the protocol is native to the server. See also Chapter 6 (Page 27) for a discussion of other types of appropriate applications for NAS servers such as the Auspex NS2000.

Analysts predict major growth for NAS over the next five years due to its major benefits in the areas of consistent availability, performance, scalability, and manageability.

Analysts predict major growth for NAS over the next five years due to its major benefits in the areas of consistent availability, performance, scalability, and manageability compared to DAS and SAN. A thorough discussion of NAS, its benefits compared to DAS and SAN and specific decision criteria for an enterprise to use in making appropriate deployment choices for a particular application can be found in a companion Auspex report titled ***A Storage Architecture Guide***, which should be read first, and can be found at: www.auspex.com.

Auspex NS2000 Servers in Internet Sites

Ever since its formation in 1987, Auspex has been delivering network server and storage solutions to companies with mission-critical profiles similar to today's Internet environment: 24x7 uptime, high-performance while handling thousands of operations per second, and multiple-terabyte capacities. With the NS2000 Enterprise Storage Server, introduced in January 1999, Auspex has increased its value to mission-critical customers, particularly those building Internet businesses.

Auspex has increased its value to mission-critical customers, particularly those building Internet businesses.

Auspex has taken the basic Network Attached Storage model, and added significant capabilities. First, the NS2000 is based on the patented DataXpress™ architecture. DataXpress is unique in that it is the only system of any kind that can specifically scale Internet data throughput. Also, the NS2000's system software, called NetOS is tuned for DataXpress and scales gracefully as the system is expanded. With DataXpress and NetOS, and as Internet load inevitably increases, networks, processors, controllers, and disks can be added to meet the higher demands. Without these Auspex technologies, additional storage units would need to be acquired and installed. Beyond cost issues, the benefits of having a centralized Network Attached Storage system with a single, master copy of data would be diminished. Since the Internet and data/load growth are synonymous, the combination of DataXpress and NetOS allows a technology manager to deploy the right solution for today's needs with confidence. The NS2000 can also be easily expanded to handle tomorrow's needs as well.

Second, the Internet is primarily a "read" environment. This is why Auspex Network Attached Storage is optimized for reads. The NS2000's FastFLO file system utilizes a technique called contiguous block allocation. This means that when a file is written to disk, the file system ensures that the blocks are organized contiguously, or side-by-side. On read operations, disk head efficiency is greatly optimized since the head can move in one direction versus hopping around the spindle to access the different blocks. Using this approach, the NS2000 can generate industry-leading read performance – 75MB/sec to a Network Attached Storage server.

Third, all Internet data is not created equally. Data sizes will also differ between sites and within sites. Then NS2000 provides "fragment support" for small files, and avoids wasting disk space.

A typical E-commerce site might include product images, text descriptions, and video/audio-based marketing content – each file a very different size. The NS2000 accounts for these differences through technology in the FastFLO file system. The standard block size within FastFLO is 4KBytes. With large files, disk I/Os are minimized and data throughput is increased since FastFLO can access data in 4KB chunks. But I/O optimization isn't the issue with very small files; it's disk utilization. To minimize wasted storage, FastFLO supports fragments, or pieces of a block. Specifically, a 4KB block can be sliced into (8) 512 byte fragments. With fragment support, a 500 byte file would occupy only one 512 byte fragment on disk. Without fragments, a 500 byte file would consume a full 4KB block – wasting 3.5KB of space. For a file system with thousands or even millions of small files, significant disk space would be rendered useless. Through its support of fragmentation, the NS2000 with FastFLO maximizes storage efficiency and minimizes disk costs. The FastFLO file system's dynamic read algorithms and robust block and fragment model helps an Internet site increase throughput, improve disk utilization and manage costs effectively.

The most advanced design available in Network Attached Storage

The Auspex NS2000 architecture is considered the most advanced design available for the specific task of serving files to the Internet with market leading performance, consistent data availability and robust security.

1. It is a ***modern parallel architecture*** introduced in 1999 and is the only choice in parallel hardware and software design among network file serving alternatives.
2. ***Very high availability*** is provided by a robust product design with customers experiencing 99.99+% availability for systems currently in production. This results in average NetServer unscheduled downtime of less than 30 minutes of unscheduled downtime per year. Full environmental monitoring and hot swap capability is also provided.
3. ***Network support*** options include 10/100BaseT Ethernet, FDDI, ATM and Gigabit Ethernet.
4. It provides ***full function data sharing*** between UNIX and Windows and any other NFS or CIFS client.
5. It offers ***scalable performance*** through a multiple CPU architecture, hardware RAID, real time operating system (kernel) design, contiguous file system and superior network and system management tools.
6. ***Backup windows and disaster recovery are optimized*** with the Auspex NDMP Turbo Replicator and efficient parallelism. Data can be replicated in parallel on each of three I/O Nodes. This provides for a total data replication performance of 200 GB+ per hour on a fully configured system. Data can also be replicated to remote locations for disaster protection, using the Auspex TurboCopy feature.
7. The ***full suite of UNIX network and system management tools*** are available, in addition to Auspex custom ***Control Point***™ ***Software.***
8. Auspex offers ***robust factory dial-in capability*** that provides world class functionality for remote diagnosis of problems when they occur
9. It offers ***easily scaleable capacity*** from the smallest department to the largest enterprise as shown in **Figure 10.**

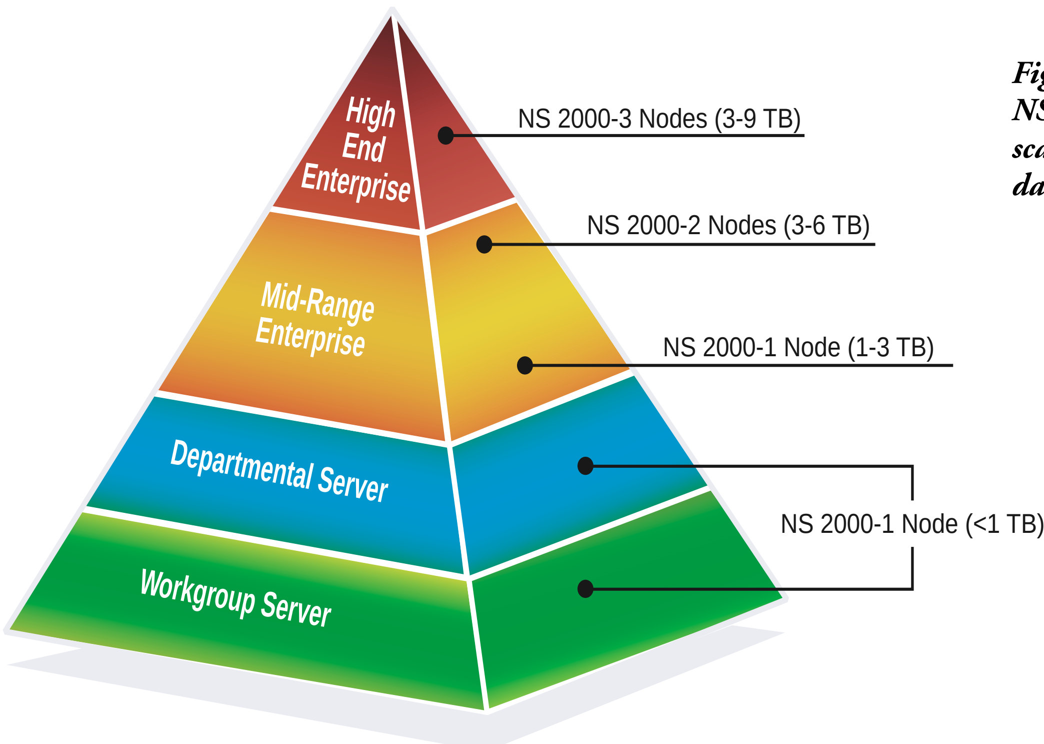

Figure 10 – The Auspex NS2000 product line scales from workgroup to data center.

The highest level of expertise available among NAS vendors

Being the originator of Network Attached Storage (NAS), Auspex is widely considered by customers and analysts alike to have the highest level of expertise in both NAS and networking technology. Since the topic of NAS is new to many customers, Auspex is committed to provide the best public information available on optimizing the flow of accurate information and support on both a pre- and post-sales basis. The Auspex sales and system engineering teams have at their disposal experts in each of the important topics discussed in this report if needed. As with any IT architecture decision, probably the most important issue is the selection of a vendor/partner with the best "total" solution. This means not only choosing a vendor who remains at the forefront of technology with the most advanced parallel architecture, but also making sure the vendor can supply the most knowledgeable professional services, consulting services and support personnel.

Auspex is widely considered by customers and analysts alike to have the highest level of expertise in both NAS and networking technology.

Parallel hardware and software design of the Auspex NetServer 2000

The Auspex architecture is based on a parallel processing design (See **Figure 11**). The Functional Multiprocessing (FMP) architecture used in Auspex NetServer 2000 is unique to Auspex. FMP assigns a specific processor to handle each of the following processes: network processing, storage processing, file system processing. In comparison, Network Appliance architecture uses only one processor for all these processes. The FMP building block makes Auspex architecture highly scalable, easier to expand, and better performing than other alternatives.

The FMP building block makes Auspex architecture highly scalable, easier to expand, and better performing than other alternatives.

Figure 11 – System block diagram of the Auspex NS2000 parallel processing architecture.

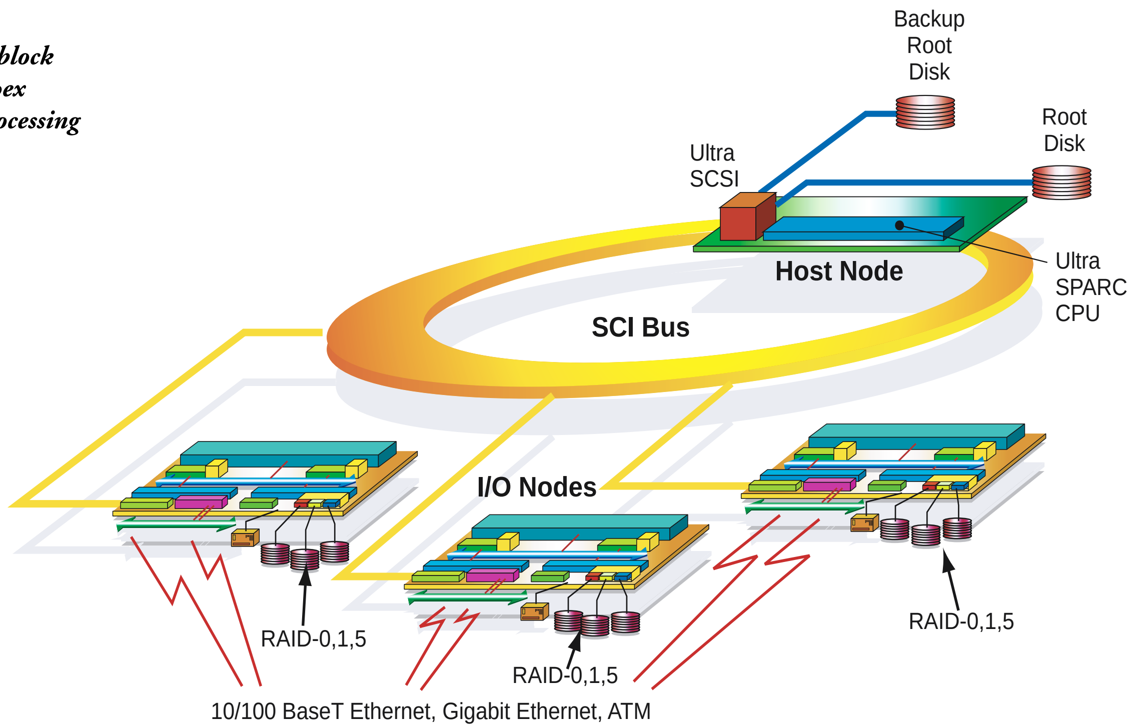

The product line supports up to three nodes, allowing efficient centralized storage and network management.

A Host Node serves as the control center of Auspex NetServer 2000 and manages intelligent I/O Node modules that are connected via a high speed link based on the Scaleable Coherent Interface (SCI) standard designed especially for parallel processing designs. The product line supports up to three nodes, allowing efficient centralized storage and network management. For example, typical high-end applications that require multiple Network Appliance F800 Series file servers would require one Auspex NetServer 2000. This translates to lower resource allocation, easier maintenance and administration, and smaller footprint.

I/O Node Building Blocks

Each Auspex NetServer 2000 I/O node has two Intel Pentium processors. The Control Node (or Host Node) has one UltraSparc processor for system management. This means a total of seven processors (six for I/O processing and one for host management) for a fully configured system. Each I/O node manages three shelves of four drawers, each with seven drives per drawer or 3 X 4 X 7 = 84 drives per node distributed to three RAID controllers. In the case of 36GB drives this equals 3TB per node or 9TB for a fully configured system. This results in a level of scalability that is not available in other enterprise storage options.

The internal hardware architecture of each I/O node is based on dual Pentium processors that reside on the bus with 1GB of DRAM memory that is used for read cache and program files. About 95% or more of the DRAM memory (950MB) is available for read cache. The NVM write cache is battery-backed SRAM that serves as a cache for fast writes.

Writes execute at memory speed for write back cache applications. An option exists for write-through caching where the write goes to disk instead of cache as required by Oracle and other database applications. Each node also contains a PCI add-on board with another 128MB of NVM that serves as file systems cache for journaling logs and fast restore. This results in a level of system performance and reliability that is not available in other enterprise storage options.

Auspex NS2000 Functional Multiprocessing (FMP) Parallel Architecture

Although all hardware is based on the industry-standard Intel high-volume architecture, each processor is assigned a specific function instead of operating symmetrically as in a typical Intel-based Symmetric Multi Processor (SMP) system. Each I/O node consists of industry-standard dual Intel processors, dual PCI busses, associated PCI cards and ECC memory (see **Figure 12**).

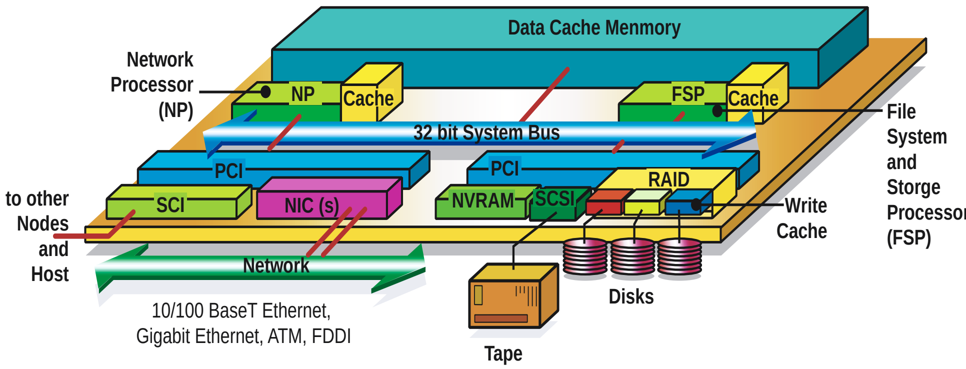

Figure 12 – I/O Node block diagram illustrating the Auspex patented Functional Multiprocessing (FMP) design.

One I/O node processor is called the Network Processor (NP) and manages highly reliable customized software that controls all network protocol and caching functions. The other I/O node processor is called the File and Storage Processor, (FSP) which also executes highly reliable, customized software that handles file system processing and storage processing. The customized software on both processors works closely together in a configuration called the DataXpress kernel. FMP system software consists of a unique custom messaging system that enables efficient network and storage processing on the I/O nodes and efficient system and data management on the host node. It allows the nodes to each efficiently perform their assigned functions. In addition, the FMP architecture improves system availability by isolating the I/O nodes from unplanned outages of the general purpose Solaris OS; I/O processing can continue even when the host node is down. For more information about the Auspex functionally parallel architecture, see the NS2000 Product Guide at: http://www.auspex.com.

The customized software on both processors works closely together.

Auspex NetServer 2000 Host Node and System Management

The Auspex NetServer 2000 Host node processor runs the standard Solaris OS. The Host node enables all management and control functions typically expected in a data center UNIX environment. Auspex Control Point™ is management control software that provides additional features like non-disruptive local backup on each I/O node. Control Point™ is a Java-based Graphical User Interface (GUI) program that runs in standard Web browsers and allows simple and effective remote control of Auspex NetServer 2000.

Auspex Control Point™ is proprietary management control software that provides additional features like non-disruptive local backup on each I/O node.

Why Auspex chose an FMP over an SMP architecture

Symmetric Multiprocessing (SMP) is a type of computer architecture that provides fast performance by making multiple processors available to simultaneously execute multiple software programs. SMP systems are suited for compute-intensive applications. In SMP, any processor executes any program or process. A variety of specialized operating systems are available to support SMP architectures.

In Auspex's patented FMP architecture, each processor executes a predefined set of programs or processes. In a highly predictable environment, like I/O and network processing, this architecture can provide superior performance and scalability characteristics.

As an SMP system gets busier, the processors must spend more time in scheduling work and less time performing work.

Auspex NetServer 2000 links multiple I/O nodes with a host node for data and system management. This architecture distributes the work to many processors working in parallel. The efficient Scaleable Coherent Interface (SCI) interconnect allows the multiple nodes of an Auspex NetServer 2000 system to act as one. The SCSI bus provides a superfast network for message exchange between computer nodes of the Auspex system.

SMP Multiprocessing systems are much more complicated than single-process operating systems because the operating system must allocate resources to competing programs, or processes, in a reasonable manner. The more processes an operating system must support, the more complex the scheduling algorithms necessary to accomplish this, and the more time the processors spend task-switching and running scheduling programs to determine what to do next. Also, SMP machine performance degrades more quickly in performance than an FMP design. This is because as an SMP system gets busier, the processors must spend more time in scheduling work and less time performing work as in the FMP design of Auspex NetServer 2000. This makes an 8-processor SMP system more expensive, while yielding only marginally improved performance than a 4 processor SMP for heavy workloads. Not only does the NS2000 distribute different functions to the multiple CPUs in each I/O node, but also additional functional specificity is achieved by distributing different functions to the host node compared to the I/O node. In addition, the NS2000 design allows multiple I/O nodes to function in parallel. This efficient parallel distribution of work illustrates the advantages of the NS2000 parallel architecture compared to SMP architecture (See **Table 4**).

This NS2000 distribution of work, both within nodes and between nodes, results in higher performance compared to SMP machines especially at higher I/O workloads. At the I/O node level, this is due to reduced task switching and reduced time spent running scheduling routines. By taking advantage of parallel processing, Auspex NetServer 2000 avoids the bottlenecks that result from scheduling complexities with heavy workloads in an SMP environment (see **Table 4**). The Auspex design also provides greater predictability and consistency in file service performance compared to an SMP design. This arises from the greater predictability of the time to complete each program on each node and each processor because of the greatly reduced task switching and scheduling overhead compared to SMP.

Table 4 – The NS2000 distributes processing functions not only among processors witin nodes but also among processors between nodes. SMP computers perform all functions in one node.

Architecture	NS2000 (FMP)							SMP
Processing Nodes	One Host Node	Three FMP I/O Nodes						One Node
		Network Processors			File System and Storage Processors			
		NP1	NP2	NP3	FSP 1	FSP 2	FSP 3	
Network Processing	No	Yes	Yes	Yes	No	No	No	Yes
File System Processing	No	No	No	No	Yes	Yes	Yes	Yes
Storage Processing	No	No	No	No	Yes	Yes	Yes	Yes
Management software	Yes	No	No	No	No	No	No	Yes
Peripheral management	Yes	No	No	No	No	No	No	Yes

Appropriate World Wide Web applications for an Intelligent Web cache

8

How Web Caches work

Web Caches are deployed in a variety of places on the "client side" and on the "server side."

In this Chapter we will discuss the reasons how Intelligent Web Caches work and the benefits of deploying them to prevent unnecessary outgoing and incoming Internet traffic to an Internet site. As discussed in Chapter 6, Intelligent Web Caches are often deployed when a Network Attached Storage (NAS) server is configured as back-end storage for appropriate applications. Web Caches are deployed in a variety of places on the "client side" and on the "server side." They are sometimes known as proxy servers since the cache acts on behalf of the origin server when serving data to a client. In this Chapter we will discuss two of the more frequently encountered configurations of Web Caches although there are many more uses for them in the enterprise and the Internet. For more information, see also www.cacheflow.com.

Client Side Web Caching

As end-user access speeds increase dramatically, the true cause of the Internet's inherent slowness is becoming more visible: latency. Multiple factors contribute to the total latency of the Internet. Congested networks, congested servers and packet loss are well-known causes. In addition to the causes above, there is another critical basis for the Internet's latency: the speed of light. Internet protocols are two-way, requiring hundreds of round trips between the client and the Web server housing the information being requested.

Improving the response time of the Internet requires technology that eliminates these latencies. It is not enough to simply cache Web pages. Accelerating the Internet requires facilitating faster access to the Web server content that exists at the moment the request is made.

Traditional "Client Side" Web Caches

In the past, vendors have offered products that are simple caches. These historical products were originally designed to save bandwidth; they do little or nothing to reduce latency. They simply cache Web objects and deliver those same objects to users over time. If the content on the origin Web server changes, the cache does not intelligently keep itself up to date; it likely will keep delivering the old ("stale") content until a cache timeout period is reached.

Pipelining: Fast Content Retrieval, The First Time

When a browser requests content, dozens of round trips must take place between the browser and the distant Web server. This is because a Web page is typically composed of dozens of objects that are serially retrieved. For each object there typically must first be a TCP session setup followed by an HTTP "get" request.

This serial retrieval of objects presents a major delay for the end user. With Auspex partner, CacheFlow's intelligent web cache deployed on the client side, a large portion of this delay is eliminated. The client connection terminates at the CacheFlow device. One of the CacheFlow algorithms is called Pipeline Retrieval. Instead of retrieving objects serially, this patent-pending algorithm opens as many simultaneous TCP connections as the origin server will allow and retrieves objects in parallel. The objects are then delivered from the appliance straight to the user's desktop as fast as the browser can request them **(Figure 13)**.

Figure 13 – Intelligent "Client side" Web Caches retrieve objects in parallel instead of serially.

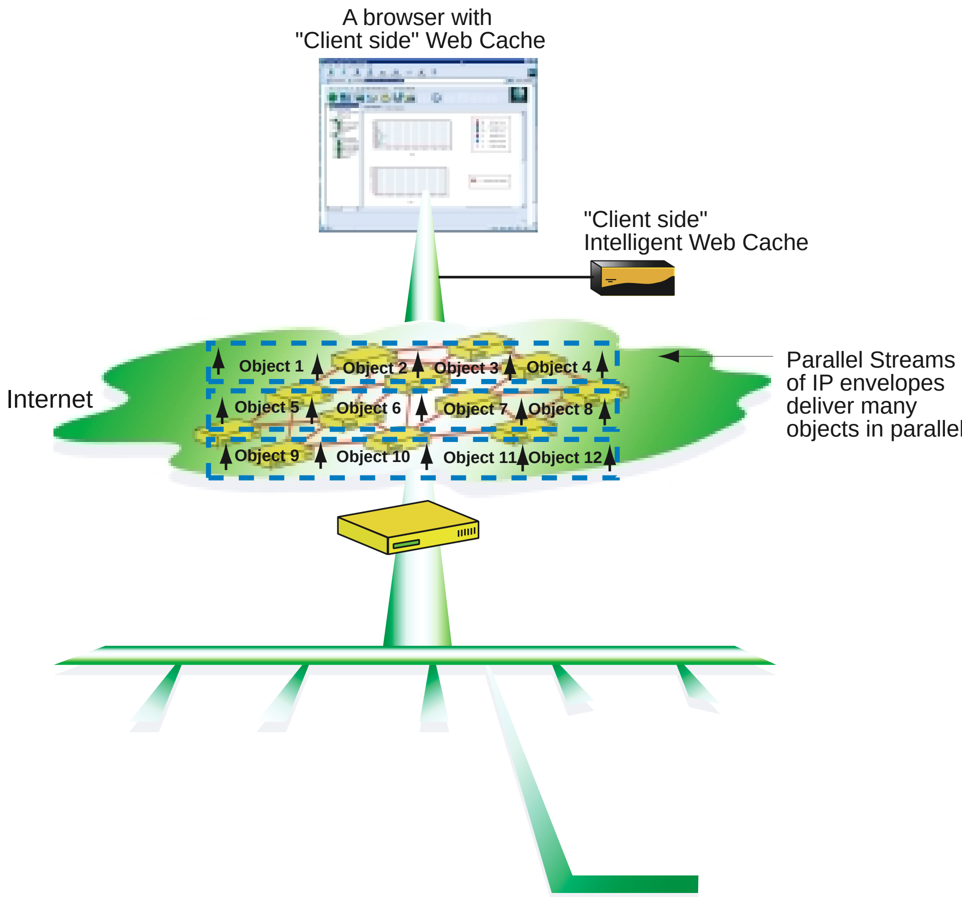

In some cases it has been reported that up to 80% of unnecessary Internet traffic can be eliminated by a client side Web cache.

Using certain rules that examine the frequency with which the elements in a particular page changes in content, it is possible for the WebCache to update only those elements that are likely to change thereby eliminating many unnecessary HTTP requests from being sent over the Internet. In some cases it has been reported that up to 80% of unnecessary Internet traffic can be eliminated by a client side Web cache.

Server Side Web Caching

"Server side caching" can deliver immediate and powerful benefits to many type of B2B and B2C E-commerce business models. Intelligent Web Caches are unparalleled in delivering E-commerce acceleration that keeps customers coming back to transact business.

Table 5 – Many users cancel transactions if a page takes longer than 8 seconds to retrieve.

Time to Load Page	Customer Exit Rate
< 7 seconds	7%
8 seconds	30%
>12 seconds	70%

Overloaded Servers and Impatient Consumers

Electronic commerce (E-commerce) is complementing and in some cases surpassing traditional retail sales for many businesses. Unfortunately, consumers expect much more from Web-based stores than they do traditional retail facilities. While people are accustomed to waiting in lines at stores, they are less tolerant of delay while waiting for a Web page to load. As shown in **Table 5**, recent research concluded that if a page takes longer than 8 seconds to load, many users will cancel the transaction or exit the site.

To eliminate the lost revenue potential, and satisfy E-commerce consumers, there are four key requirements an E-commerce solution must address the following:

- Throughput. The number of transactions that an E-commerce solution is capable of handling.
- Speed. Delivering Web pages and fulfilling E-commerce transactions in an acceptable amount of time.
- Reliability & Availability. The ability to consistently serve requests and complete transactions with no service interruptions.
- Scalability. Building a solution that easily grows as the business grows, and is capable of handling traffic spikes.

Network administrators are challenged to build solutions that are capable of delivering the scalability, load capacity, and throughput that is needed to meet customer expectations and to capitalize on the enormous revenue potential in E-commerce.

Figure 14 – A typical E-commerce Web site where all requests must pass through the firewall and L4 switch.

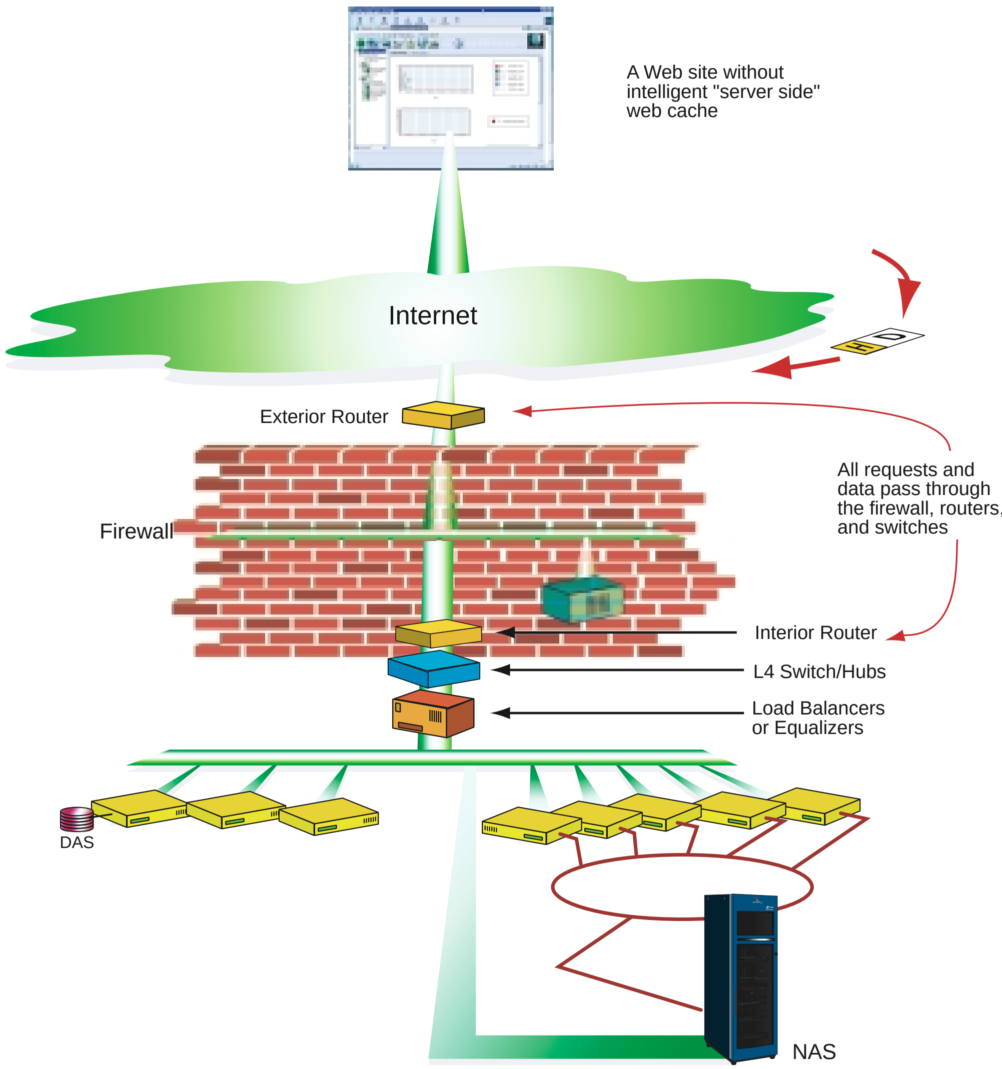

Pinpointing E-commerce Bottlenecks

There are many factors limiting E-commerce performance. Within an E-commerce site and across the Internet, everything from link speed to routers and server bottlenecks add delay to an end-user request. **Figure 14** shows a typical E-commerce Web Site in which requests for Web pages must go through the router, firewall, and Layer 4 switch (L-4 switch) before they reach the Web servers, are processed, and sent back out through the same network components. Forcing user requests through these devices, each with finite throughput, can easily expose network scalability constraints and performance bottlenecks.

In addition, as more and more users try to access the same content, the redundant load on the firewalls and servers for the same Web objects increases significantly. This is particularly problematic in connection with promotions or major events. The most common sources of E-commerce slowdowns are:

- Redundant connections required to serve the same objects.
- Traffic surges due to promotions or major events
- Overburdened firewalls
- Overburdened servers

Intelligent Web Caches reduce the risk of back-end devices becoming points of failure and causing unplanned outages

In an E-commerce network, a large percentage of Web objects are static. Examining the elements of a Web page shows that objects can generally be grouped into three categories: objects that rarely change (logos), objects that change weekly, hourly, or daily (news items, promotions), and items that are dynamically generated and unique for each site visitor. As much as 90% or more of Web objects can be static. Delivering these objects over and over from the Web servers unnecessarily force requests through the internal network. This model causes delays and exposes each network device as a potential bottleneck or point of failure.

"Server Side" Intelligent Web Caches reduce many risks to E-commerce sites

In an E-commerce environment, caches can be deployed at the "server-side" and provide the following benefits:

- Offer site visitors an optimal Web site experience.
- Offload overburdened servers.
- Offload overtaxed firewalls.
- Scale the network to handle more customer transactions.
- Scale the network to handle large traffic spikes.
- Reduce capital and operating costs.

In the example displayed in **Figure 15**, customer HTTP (port 80) requests coming in from the Internet are redirected through the router to the cache. The cache then serves the Web objects that it has stored directly back to the client. Requests for dynamic content or secure (SSL) transactions are passed through to the origin servers for processing. It is not unusual for cache hit rates to exceed 90% in E-commerce environments.

Figure 15 – Server Side Caching, with an intelligent Web Cache, service frequent requests before the firewall.

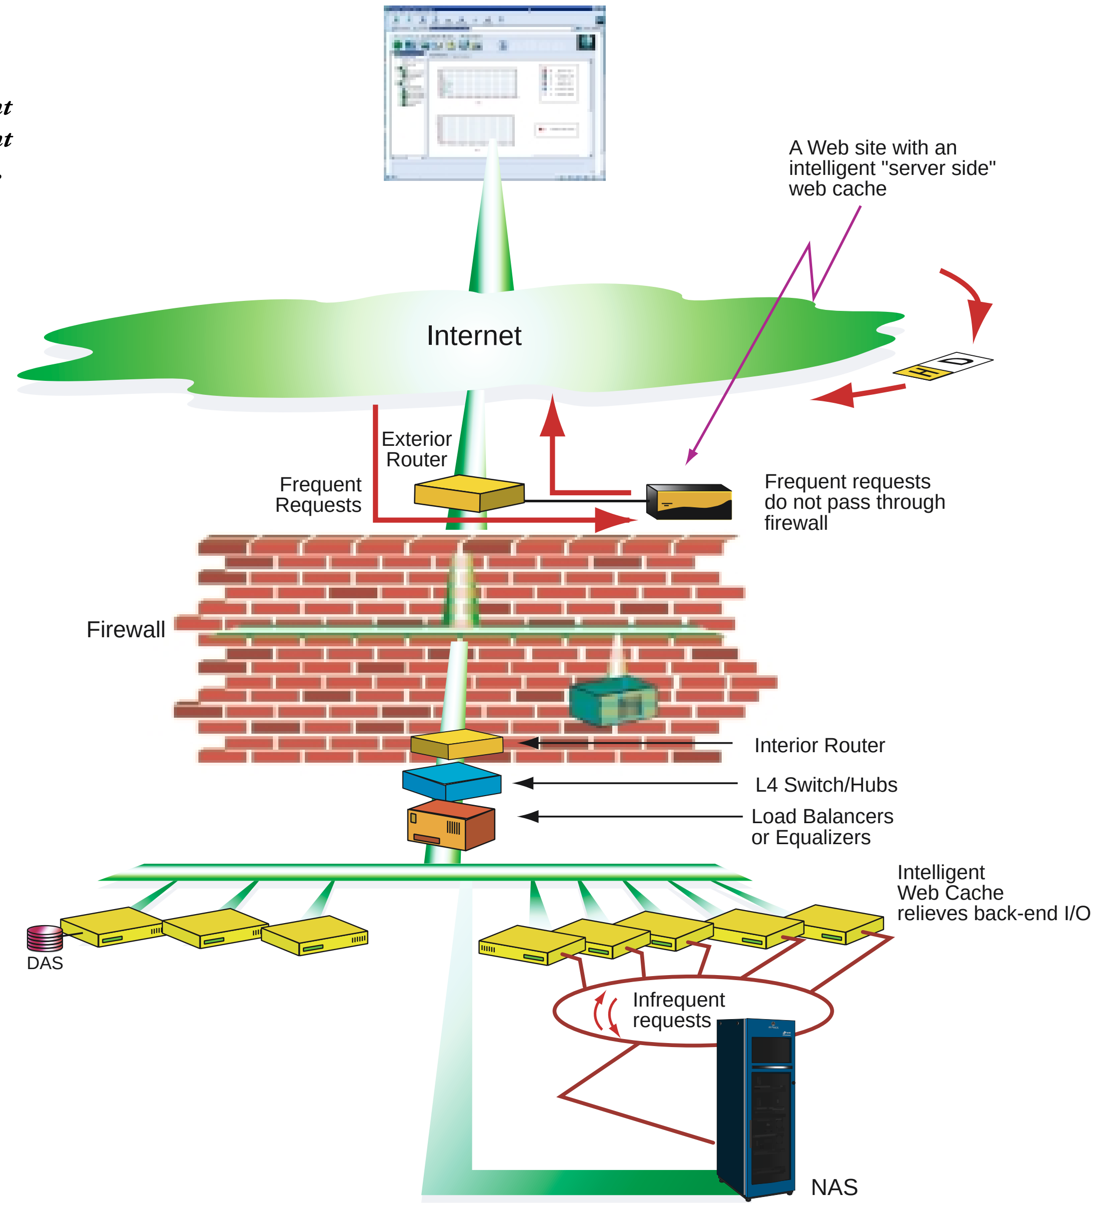

Intranet users who are dependent on the firewall for their network-based access receive improved performance.

Other than the obvious load reduction on firewalls and servers, serving a very high number of Web requests outside of the firewall has other benefits. Having the cache handle most of the Web requests reduces the security risks of users directly accessing servers that are inside the firewall. Further, Intranet users who are dependent on the firewall for their network-based access receive improved performance.

Because as much as 90% or more of Web objects can be cached in a typical E-commerce environment, delivering these objects from the edge of the network, outside the firewall, and in front of the Web servers can deliver dramatically better response time and allow the site to handle more transactions.

Conclusion

9

This report has explained the inner workings of the Internet and provided recommendations in optimizing the design of Web Site architectures in order to meet the increasingly critical requirements of 24x7, flexibility, speed, manageability and cost effectiveness. We have reviewed how large scale Web Sites function in terms of border routers, firewalls, Intelligent Web Caches, switches, servers, and DAS and NAS storage. The drawbacks of a flat clustered storage architecture with all Web Servers configured with DAS were discussed. A multilevel functionally specific hierarchical Web Site architecture has been proposed. Appropriate applications for DAS and NAS Web Server configurations were recommended based on known workload characteristics of the various protocols involved in Internet servers. Sites wishing to optimally configure Web Server storage for an Internet site can use the Infrastructure and Planning Service (IPS) offered by Auspex.

The drawbacks of a flat clustered storage architecture with all Web Servers configured with DAS were discussed.

Summary of the Benefits of a Hierarchical Web Site Architecture

A hierarchical functionally specific Web Site architecture is a completely congruent and logical extension of Auspex's FMP's concept of functional specialization, Cache Flow's concept of a dedicated appliance, and Network Engine's optimized "thin O/S" Web Server. A hierarchical design is a variation of the client-server model applied to computer devices in a Web Site. In a functionally specialized hierarchical design each device in the Internet site is optimally designed and configured to do a specific task in an efficient manner. The functions of interior and exterior routers, firewalls, Intelligent Web Caches, switches and hubs, load balancers, different types of Internet servers, and DAS and NAS storage have been discussed relative to deploying them in an optimized hierarchical manner.

The functions of interior and exterior routers, firewalls, Intelligent Web Caches, switches and hubs, load balancers, different types of Internet servers, and DAS and NAS storage have been discussed.

Many of the same benefits accrue for Web Sites deploying a mixture of DAS and NAS as in other client server network architectures deploying DAS and NAS. These benefits are 24x7, flexibility, speed, manageability and cost. The benefits of using Auspex, Cache Flow, and other functionally specialized products in a Web Site design are summarized below:

Flexibility (Scalability)

The hierarchical approach lets designers scale and tune front-end Web request servers HTTP, FTP, news, mail separately from the back-end Web page server(s). The scalability of the overall architecture and the NAS servers themselves has been singled out as an essential benefit by corporate sites like Fujitsu and ISPs such as Easynet (Frames), CERFnet and DIGEX.

The hierarchical approach permits additional request processors to efficiently access that storage.

Increasing volumes of HTTP requests do not always equate to increasing storage requirements. However, in many environments, the increased request load is focused on the same storage, but the hierarchical approach permits additional request processors to efficiently access that storage.

A given hardware vendor's workstation offering may be very attractive from a processing standpoint, but lag in disk subsystem performance or lack a robust feature set Auspex NS2000 NetServers perform bit-level disk checking, a feature not supported by all vendors. Segregating the CPU function from the storage function gives designers freedom to choose the best components.

The hierarchical approach gives the designer freedom to optimize, and re-optimize.

A particular implementation of webserver code (Netscape, Microsoft, Apache, OpenMarket, CERN, etc.) may run better/faster/cheaper on one OS-CPU combination than another. Indeed, this performance advantage can flip-flop over time from release to release. The hierarchical approach gives the designer freedom to optimize, and re-optimize. The need for one kind of tuning (e.g., disk capacity allocated per front-end server) actually disappears altogether for NAS hosted applications. Once configured, a dataless request server's disk complement rarely needs to be reconsidered. It simply relies on increased capacity on the back-end page server. In ISP or Intranet webpage hosting environments, the customer can start small, utilize freeware (e.g., CERN HTTPD) and work within a modest budget. If performance and functional demands outstrip the original configuration, the front-end server can easily be swapped out for a more powerful one, where no disk data needs to be relocated.

7x24 (Availability)

When Fujitsu was using a single-server approach, rebooting the server when it went down took a long time. In fact, Fujitsu has indicated that, prior to moving to the hierarchical architecture, their HTTP service was down about a third of the time. The high availability that the hierarchical approach provides is essential to strategic Web Sites, whether internal or external.

The architecture is more <u>resilient to total front-end server failures,</u> since N-1 surviving peers can carry on, all the while mounted to the shared NFS data.

Compartmentalization of function <u>isolates failures</u> and makes problems easier to identify and faster to resolve. For example, UNIX reboots are faster when there's no disk on the front-end servers to fsck.

The <u>elimination of UNIX</u> or any complex general-purpose operating system, for that matter from the net-to-disk access path makes file service more stable.

Narrower, <u>more specialized work is given to the front-ends.</u> They have a smaller working set and thus are statistically less likely to trip over bugs.

The <u>ability to scale disk storage</u> on-line reduces the need for planned outages. Use of a NetServer makes the following an irrelevant question: "Do my HTTP servers support hot-pluggable drives?"

Speed (Performance or Fast Response Time to Users)

NAS servers with "thin O/S" optimized application servers in combination with Intelligent WebCaches increase the user response time.

Configuring NAS servers with "thin O/S" optimized application servers in combination with Intelligent WebCaches increase the user response time (speed) or the Web Site. The "thin O/S provides improved availability since there is less operating system code to cause outages. The Intelligent WebCache when mounted on the front-end of the Web Site (sometimes before the firewall) takes the majority of HTTP requests from the back-end network allowing for more efficient use of processing resources throughout the site.

Manageability (Ease of Administration)

All data comes out of a pool; small chunks of free disk space are no longer scattered across numerous Web request servers. Backup is consolidated onto a central repository. SCSI-attached stackers and tape libraries are well suited for the NS2000's SCSI-attached disks. There is neither a need to consume network bandwidth to back up individual disk-rich stand-alone servers, nor must each be equipped with its own backup device.

The configuration of front-end servers is simplified and rendered more generic. CERFnet has taken this to the limit by restricting its thirteen front-end servers to Solaris platforms. In their "cookie cutter" approach, one reference system image is available for new or rebooted front-ends.

Loosely-coupled scaling means each configured front-end server can continue to be productive for a longer period of time, contributing to serving the application load as best it can. This, in turn, reduces pressure on the administrator to continually evaluate (and reevaluate) all the workstation vendor offerings as they furiously leapfrog one another other's price-performance.

Loosely-coupled scaling means each configured front-end server can continue to be productive for a longer period of time.

Cost Effectiveness

The hierarchical approach gives the Internet Site architect the freedom to save money on request processing by buying commodity-priced "thin O/S" optimized front-end servers. Thus, CPUs may be cost-effectively upgraded as fast as the technology cycle allows.

Though ISPs tend to have a bias toward UNIX variants, and have more in-house UNIX than Windows expertise, with the advent of some very good software platforms like Microsoft Internet Information Server (IIS is bundled with NT) and O'Reilly Web Site, may gain some ground with the ISPs.

The advent of some very good software platforms like Microsoft Internet Information Server (IIS is bundled with NT) and "O'Reilly Web Site," may gain some ground with the ISPs.

Another cost benefit is that the site can amortize the proportionally higher investment in a quality back-end NAS server. The price can be justified given the value of the contents and the features delivered such as redundant power supplies, mirroring, striping, hot-plugability. The price/performance of a NAS device rides a much shallower curve than the commodity desktops which are routinely obsoleted every 12-18 months. The administrator may recycle outmoded desktop machines.

Native FTP Support

Another important benefit of NAS deployment for selected applications is the native support of FTP in products such as the Auspex NS2000. By streaming FTP data directly to requesting clients, the need to have a UNIX host intervene in the actual data transfer is eliminated saving in complexity and increasing speed of transfer FTP files to the Internet.[3] Since many ISPs and Intranet sites currently provide FTP services as well as HTTP services, and until WebNFS becomes ubiquitous, FTP will remain the choice for those wanting to transfer medium to large-size files. The use of this unique feature of NAS servers like the Auspex NS2000 will provide better data throughput and greater scalability for sites, which provide FTP services.

Another important benefit of NAS deployment for selected applications is the native support of FTP in products such as the Auspex NS2000.

[3]FTP uses two connections—a control connection and a data connection. In the case of the Auspex NS2000, the NP handles the data connection.

Glossary of Terms

A

10BaseT
Ethernet with a data transfer rate of 10 Mbits/sec.

100BaseT
Also known as Fast Ethernet with a data transfer rate of 100 Mbits/sec.

ASCII
Pronounced ask-ee, ASCII is a code for representing English characters as numbers, with each letter assigned a number from 0 to 127. There are several larger character sets, also generally called ASCII, that use 8 bits, which gives them 128 additional characters. The extra characters are used to represent non-English characters, graphics symbols, and mathematical symbols.

ASDL
Asymmetric Digital Subscriber Line; a technology that allows more data to be sent over existing copper telephone lines. ADSL supports data rates of from 1.5 to 9 Mbps when receiving data (known as the downstream rate) and from 16 to 640 Kbps when sending data (known as the upstream rate). ADSL requires a special ADSL modem. See also XDSL and SDSL.

ASP
Application Service Providers are companies that manage and distribute software-based services and solutions to customers across a wide area network from a central data center. According to http://www.aspnews.com, ASPs are broken down into five subcategories:

1. Enterprise ASPs — deliver high-end business applications.
2. Local/Regional ASPs — supply wide variety of application services for smaller businesses in a local area.
3. Specialist ASPs — provide applications for a specific need, such as Web Site services or human resources.
4. Vertical Market ASPs — provide support to a specific industry such as healthcare.
5. Volume Business ASPs — supply general small/medium-sized businesses with prepackaged application services in volume.

ASP
Active Server Pages is a specification for a dynamically created Web page that utilizes ActiveX scripting — usually VB Script or Jscript code.

ATM
Asynchronous Transfer Mode. ATM is a suite of protocols intended to provide the switching and multiplexing services necessary to carry voice, data and video and multimedia traffic using fixed 53-byte cells. Standards are being defined to allow ATM to emulate traditional LANs (LANE).

b
Abbreviation for "bit" where 8 "bits" comprise a byte.

B
Abbreviation for byte or the equivalent of one character in text.

B2B
Business to Business usually refers to an enterprises inclusion of customers and vendors on an Extranet portion (password required) of their internal network.

B2C
Business to Consumer refers to the portion of E-commerce where goods and services are transacted between businesses and consumers over the Internet instead of through traditional retail channels.

Backbone
A backbone is often used to describe the main network connections composing the Internet.

Bastion Host
A bastion host in the firewall is the primary point of contact for connections coming in from the Internet for services such as email and receiving data from the internal FTP site. The Bastion host is a heavily protected server with many layers of security software built in and it is the only contact point for incoming Internet requests. Bastion hosts can also be set up as proxy servers - servers that process any requests from the internal corporate network to the Internet, such as browsing the Web or downloading files via FTP.

BGP
BGP is short for Border Gateway Protocol, an Internet protocol that enables groups of routers to share routing information so that efficient, loop-free routes can be established. BGP is commonly used within and between Internet Service Providers (ISPs). The protocol is defined in RFC 1771.

BIND
Berkeley Internet Name Domain

BMP
BMP is the standard bit-mapped graphics format used in the Windows environment. Files having a .BMP extension.

Border Router
A specialized type of router on the border of a network and the Internet is called a Border Router or Exterior Router. This router screens packets between the Internet and the perimeter network. It adds redundancy to Web Site security by screening incoming packets from the Internet using the same rules as the internal screening router and protecting the network even if the internal router fails.

Bridge
A bridge is a device that connects two local-area networks (LANs), or two segments of the same LAN. The two LANs being connected can be alike or dissimilar. For example, a bridge can connect an Ethernet with a Token-Ring network. Unlike routers, bridges are protocol -independent. They simply forward packets without analyzing and rerouting messages. Consequently, they're faster than routers, but also less versatile.

Browser
Browser is short for Web browser, a software application used to locate and display Web pages. The two most popular browsers are Netscape Navigator and Microsoft Internet Explorer.

Cable Modem

A cable modem is a modem designed to operate over cable TV lines. Because the coaxial cable used by cable TV provides much greater bandwidth than telephone lines, a cable modem can be used to achieve extremely fast access to the World Wide Web. There are a number of technical difficulties, however. One is that the cable TV infrastructure is designed to broadcast TV signals in just one direction - from the cable TV Company to people's homes. The Internet, however, is a two-way system where data also needs to flow from the client to the server. In addition, it is still unknown whether the cable TV networks can handle the traffic that would ensue if millions of users began using the system for Internet access. Despite these problems, cable modems that offer speeds up to 2 Mbps are already available in many areas.

Cache

Cache (pronounced cash) can be either a reserved section of main memory or an independent high-speed disk storage device. Two types of caching are commonly used in personal computers: memory caching and disk caching. Disk caching can dramatically improve the performance of applications, because accessing a byte of data in RAM can be thousands of times faster than accessing a byte on a hard disk. When data is found in the cache, it is called a cache hit, and the effectiveness of a cache is judged by its hit rate. The strategies for determining which information should be kept in the cache constitute some of the more interesting problems in computer science. Caching is applied to the Internet whenever a proxy server or Web cache is placed at the Internet boundary. See also Proxy servers.

Chat server

Chat servers use the IRC format for message exchange. IRC stands for Internet Relay Chat. To join an IRC discussion, you need an IRC client and Internet access. The IRC client is a program that runs on your computer and sends and receives messages to and from an IRC server. The IRC server, in turn, is responsible for making sure that all messages are broadcast to everyone participating in a discussion. There can be many discussions going on at once; each one is assigned a unique channel.

Checksum

Checksum is a simple error-detection scheme in which each transmitted message is accompanied by a numerical value based on the number of set bits in the message. The receiving station then applies the same formula to the message and checks to make sure the accompanying numerical value is the same. If not, the receiver can assume that the message has been garbled.

CIFS

Common Internet File System. A connection-oriented, network file-sharing protocol developed by IBM and Microsoft as part of LAN Manager. CIFS is the native file sharing protocol for systems running Windows for Workgroups, Windows95/98 and Windows NT. Sometimes referred to as SMB.

Client Server

Client Server Architecture is a network architecture in which each computer or process on the network is either a client or a server. Servers are powerful computers or processes dedicated to managing disk drives (Network Attached Storage Servers or, generically, file servers), printers (print servers), or network traffic (network servers). Clients are PCs or workstations on which users run applications. Clients rely on servers for resources, such as files, devices, and even processing power. Another type of network architecture is known as a peer-to-peer architecture because each node has equivalent responsibilities. Client-server architectures are sometimes called two-tier architectures.

Control Point™

Auspex's customized management control software.

Cookie

A cookie is a message given to a Web browser by a Web server which is then sent back to the server each time the browser requests a page from the server. The main purpose of cookies is to identify users and possibly prepare customized Web pages for them.

CPU

Central Processing Unit. Can refer to either a processor chip such as Sun's SPARC or Intel's Pentium, or to a processor chip or chips and support circuitry on a CPU board.

DataXpress

Communication among the NS2000's multiple hardware processors and software processes are handled by DataXpress, a low-overhead message-passing kernel executing on each processor.

DNS Server

DNS is an abbreviation for Domain Name System (or Service), an Internet service that translates domain names into IP addresses. A DNS server provides this service on the Internet. Because domain names are alphabetic, they're easier to remember. The Internet however, is really based on IP addresses. Every time you use a domain name, therefore, a DNS service must translate the name into the corresponding IP address. For example, the domain name www.example.com might translate to 198.105.232.4. The DNS system is, in fact, its own network. If one DNS server doesn't know how to translate a particular domain name, it asks another one, and so on, until the correct IP address is returned.

Domain Name

A domain name is a name that identifies one or more IP addresses. For example, the domain name microsoft.com represents about a dozen IP addresses. Domain names are used in URLs to identify particular Web pages. For example, in the URL http://www.auspex.com/index.html, the domain name is auspex.com. Every domain name has a suffix that indicates which top-level (TLD) domain it belongs to.
There are only a limited number of such domains. For example:

- gov - Government agencies
- edu - Educational institutions
- org - Organizations (nonprofit)
- mil - Military
- com - commercial business
- net - Network organizations
- ca - Canada
- th - Thailand

Because the Internet is based on IP addresses, not domain names, every Web server requires a Domain Name System (DNS) server to translate domain names into IP addresses.

EBCDIC

Abbreviation of Extended Binary-Coded Decimal Interchange Code, it is an IBM code for representing characters as numbers.

ECAD

Electrical Computer Aided Design

Email

Email (sometimes email) is short for electronic mail, the transmission of messages over communications networks.

Email Server

An email server is a Web Server used exclusively for email traffic. Email servers experience very high I/O rates of varying I/O sizes because of the many pictures and files being attached to email messages. For this reason special algorithms are often applied to emails that delay the processing of an email message if it is very large. Some email servers compress attachments or send them at off peak hours when more bandwidth is available. Email servers use Post Office Protocol or POP. See also POP.

Encryption

Encryption refers to the translation of data into a secret code. To read an encrypted file, you must have access to a secret key or password that enables you to decrypt it. Unencrypted data is called plain text; encrypted data is referred to as cipher text. There are two main types of encryption: asymmetric encryption (also called public-key encryption) and symmetric encryption.

EM-Net

The NS2000 Environmental Monitoring Network that connects to all chassis in an NS2000 system and reports a variety of control information to the Host Node.

Ethernet

A Local Area Network (LAN) protocol developed by Xerox in cooperation with Digital Equipment and Intel in 1976. Ethernet supports a star or bus topology and supports a data transfer rate of 10 megabits per second or 10 Mbps. The Ethernet specification formed the basis of the IEEE 802.3 standard, which specifies the physical and lower software layers. Ethernet uses the CSMA/CD access method for handling simultaneous demands and is one of the most widely implemented LAN standards.

Extranet

Extranet refers to an Intranet that is partially accessible to authorized outsiders. Whereas an Intranet resides behind a firewall and is accessible only to people who are members of the same company or organization, an Extranet provides various levels of accessibility to outsiders.

Fast Ethernet or 100BaseT

Defined by the IEEE 802.3 committee, provides a 100 Mbps standard that is compatible with existing 10BaseT installations, preserving the CSMA/CD media access control (MAC) protocol.

FastFLO

The NS2000 proprietary file system that is optimized for providing high performance and consistent file services.

FDDI

Fiber Distributed Data Interface. A standard for local area networks that typically uses fiber-optic media capable of data rates up to 100 megabits/second over distances up to 100 km. An FDDI network is a token-based logical ring, and is often constructed as a pair of counter-rotating redundant rings (called dual-attachment mode) for reliability. Ethernet, in contrast, is a bus-based, non-token, 10-megabits/second network standard.

Firewall

A firewall is a system designed to prevent unauthorized access to or from a private network. Firewalls can be implemented in both hardware and software, or a combination of both. Firewalls are frequently used to prevent unauthorized Internet users from accessing private networks connected to the Internet, especially Intranets. All messages entering or leaving the Intranet pass through the firewall, which examines each message and blocks those that do not meet the specified security criteria.

There are several types of firewall techniques:

- Packet filter: Looks at each packet entering or leaving the network and accepts or rejects it based on user-defined rules. Packet filtering is fairly effective and transparent to users, but it is difficult to configure. In addition, it is susceptible to IP spoofing.
- Application gateway: Applies security mechanisms to specific applications, such as FTP and Telnet servers. This is very effective, but can impose performance degradation.
- Circuit-level gateway: Applies security mechanisms when a TCP or UDP connection is established. Once the connection has been made, packets can flow between the hosts without further checking.
- Proxy server: Intercepts all messages entering and leaving the network. The proxy server effectively hides the true network addresses.

In practice, many firewalls use two or more of these techniques in concert. A firewall is considered a first line of defense in protecting private information. For greater security, data can be encrypted.

FMP

Functional Multiprocessing (FMP) is the term Auspex uses for its patented distributed parallel processing NS2000 architecture. Each NS2000 I/O Node is based on an Asymmetric Multiprocessing design with two processors running Auspex's real time kernel (DataXpress™) to simultaneously and efficiently execute on each processor different functions in the network file serving process. One processor handles network processing and the other processor handles File and Storage Processing. A Host Node based on the traditional general-purpose single CPU computer and OS is used primarily for system management activity. Up to three I/O Nodes and one Host Node are connected by a Scalable Coherent Interface (SCI). System software consists of a unique custom messaging system that enables efficient network and storage processing on the I/O Nodes and efficient system and data management on the Host Node. The FMP architecture improves system availability compared to other approaches by isolating the I/O Nodes from unplanned outages of the general purpose OS, and I/O processing can continue even in the event that the Host Node is down. See also SMP, parallel processing, SCI.

FSP

The File System and Storage Processor refers to one of the two Intel Pentium processors on an I/O Node of an NS2000 system. This processor runs highly optimized software that manages all file system and storage processing of the I/O Node and communicates with other I/O Nodes and the Host Node. See also Network Processor (NP).

FTP

FTP is an abbreviation of File Transfer Protocol, the protocol used on the Internet for sending files.

Gateway

A gateway is a combination of hardware and software that links two different types of networks. An email gateway for example allows users on two different types of email systems to exchange emails.

GIF

GIF is pronounced jiff or giff (hard g) stands for graphics interchange format, a bit-mapped graphics file format used by the World Wide Web, CompuServe and others. GIF supports color and various resolutions. It also includes data compression, making it especially effective for scanned photos. A patent owned by Unisys, Inc. covers the format, and certain applications may require a license.

Gigabit Ethernet

Gigabit Ethernet is a standard of the IEEE committee that provides a mechanism for conveying Ethernet formatted packets at Gigabyte/second speeds.

Gigabyte

1000 Megabytes.

GUI

An acronym referring to a Graphical User Interface that is the screen presented to a user in any computer application.

HDDA

A term that refers to a High-density Disk Array shelf in an NS2000 containing 28 drives arranged in four drawers of 7 drives each. A maximum three I/O Node NS2000 system contains nine HDDAs or 9x28=252 disk drives.

Head end

A head end receives television transmissions from satellites and has Internet access via high-speed links to the Internet. These feeds provide cable programming and Internet access to cable customers.

HTML

Short for HyperText Markup Language, the authoring language used to create documents on the World Wide Web. HTML defines the structure and layout of a Web document by using a variety of tags and attributes.

HTTP

HTTP is an abbreviation for HyperText Transfer Protocol, the underlying protocol used by the World Wide Web. HTTP defines how messages are formatted and transmitted, and what action Web servers and browsers should take in response to various commands. For example, when you enter a URL in your browser, this actually sends an HTTP command to the Web server directing it to fetch and transmit the requested Web page. HTTP is called a stateless protocol because each command is executed independently, without any knowledge of the commands that came before it.

Hub

A hub is a common connection point for devices in a network. Hubs are commonly used to connect segments of a LAN. A hub contains multiple ports. When a packet arrives at one port, it is copied to the other ports so that all segments of the LAN can see all packets. A passive hub serves simply as a conduit for the data, enabling it to go from one device (or segment) to another. So-called intelligent hubs include additional features that enable an administrator to monitor the traffic passing through the hub and to configure each port in the hub. Intelligent hubs are also called manageable hubs. A third type of hub, called a switching hub, actually reads the destination address of each packet and then forwards the packet to the correct port.

ICMP

ICMP is short for Internet Control Message Protocol, an extension to the Internet Protocol (IP) defined by RFC 792. ICMP supports packets containing error, control, and informational messages. The PING command, for example, uses ICMP to test an Internet connection. See also PING.

IIS

IIS Short for Internet Information Server, Microsoft's Web server that runs on Windows NT platforms.

IMAP

IMAP Short for Internet Message Access Protocol, a protocol for retrieving email messages. The latest version, IMAP4, is similar to POP3 but supports some additional features. For example, with IMAP4, you can search through your email messages for keywords while the messages are still on mail server. You can then choose which messages to download to your machine. Like POP, IMAP uses SMTP for communication between the email client and server. IMAP was developed at Stanford University in 1986.

Inode

In UNIX, an inode is an index to files.

Internet

Internet refers to a global network connecting millions of computers. As of 1999, the Internet has more than 200 million users worldwide, and that number is growing rapidly. More than 100 countries are linked into exchanges of data, news and opinions.

Intranet

A network based on TCP/IP protocols (an Internet) belonging to an organization, usually a corporation, accessible only by the organization's members, employees, or others with authorization. An Intranet's Web Sites look and act just like any other Web Sites, but the firewall surrounding an Intranet fends off unauthorized access. Like the Internet itself, Intranets are used to share information. Secure Intranets are now the fastest-growing segment of the Internet because they are much less expensive to build and manage than private networks based on proprietary protocols.

InterNIC

The InterNIC is a collaborative project between AT&T and Network Solutions (NSI) and is supported by the National Science Foundation (NSF). It currently provides four services:

1. Registration Services – The InterNIC is responsible for maintaining the domains registered through registries. It also tracks the connections between Internet addresses, such as 128.33.5.21 and domain names such as www.auspex.com.
2. Directory and Database Services – This services provides an online white pages service for telephone numbers and a directory of publicly accessible online databases managed by AT&T.
3. Support Services – Outreach, education and Information services for the Internet community managed by NSI.
4. Net Scout Services – Managed by NSI this services provides access to online publications that summarize recent happenings of interest to the Internet community.

IP

The IP (Internet Protocol) is the underlying protocol for routing packets on the Internet and other TCP/IP-based networks. IP is an internetwork protocol that provides a communication standard that works across different types of linked networks for example Ethernet, FDDI or ATM. In an internetwork, the individual networks that are joined are called subnetworks or subnets. IP provides a universal way of packaging information for delivery across heterogeneous subnet boundaries. See also TCP Transmission Control Protocol.

IRC

IRC is short for Internet Relay Chat. IRC may have multiple participants. To join an IRC discussion, you need an IRC client and Internet access. The IRC client is a program that runs on your computer and sends and receives messages to and from an IRC server. The IRC server, in turn, is responsible for making sure that all messages are broadcast to everyone participating in a discussion.

ISDN

ISDN or integrated services digital network, is an international communications standard for sending voice, video, and data over digital telephone lines or normal telephone wires. ISDN supports data transfer rates of 64 Kbps (64,000 bits per second). Most ISDN lines offered by telephone companies give you two lines at once, called B channels. You can use one line for voice and the other for data, or you can use both lines for data to give you data rates of 128 Kbps, much faster than the data rate provided by today's fastest modems. The original version of ISDN uses baseband transmission. A modern version of ISDN, called B-ISDN, uses broadband transmission and is able to support transmission rates of 1.5 Mbps. B-ISDN however, requires fiber optic cables and is not widely available.

ISP

An ISP or Internet Service Provider is a firm that provides individuals and companies with access to the Internet for a monthly fee. The ISP provides a software package, username, and password and access phone numbers. With a dial in modem and individual can logon to the Internet and browse the World Wide Web and Usenet, and send and receive emails. ISPs serve large companies by providing direct connections from the company's computer networks to the Internet. ISPs are connected to each other through Network Access Points (NAPs). ISPs are also called Internet Access Providers or IAPs.

Java

Java is a high-level programming language developed by Sun Microsystems. Java is an object-oriented language similar to C++, but simplified to eliminate language features that cause common programming errors. Java source code files (files with a .java extension) are compiled into a format called bytecode (files with a .class extension), which can then be executed by a Java interpreter. Compiled Java code can run on most computers known as Java virtual computers. Java interpreters and runtime environments exist for most operating systems, including UNIX, the Macintosh OS, and Windows. Bytecode can also be converted directly into machine language instructions by a just-in-time compiler (JIT). Java is a general purpose programming language with a number of features that make the language well suited for use on the World Wide Web. Small Java applications are called Java applets and can be downloaded from a Web server and run on your computer by a Java-compatible Web browser, such as Netscape Navigator or Microsoft Internet Explorer.

Journaling

A journaling file system keeps track of all changes to files as transactions occur in real time. In the event of unexpected system problems, the file system can be restored to a consistent state by updating a prior copy of the file system for the changes made from the point in time that the copy was made.

JPEG

JPEG is short for Joint Photographic Experts Group, and pronounced jay-peg. JPEG is a compression technique for color images. Although it can reduce files sizes to about 5% of their normal size, some detail is lost in the compression.

Kilobyte

1,000 bytes. Sometimes, particularly in main memory discussions, or among Apple Computer aficionados it can refer to 1,024 bytes.

L-4 Switch

Although most switches operate at level two of the OSI reference model an advanced type of switch is called an L4 switch which operates at the fourth layer (Transport Layer) of the OSI Reference Model and is responsible for the integrity of data transmissions between LAN segments. See also Switch.

LADDIS

An acronym formed by names of the group (Legato, Auspex, Data General, Digital Equipment Corporation, Interphase, and Sun) that developed and popularized SPEC's vendor-neutral NFS server benchmark of the same name. See SPEC.

LAN

Local area networks or LANs are networks of computers that are geographically close together; this usually means within the same building. Most LANs are confined to a single building or group of buildings. However, one LAN can be connected to other LANs over any distance via telephone lines and radio waves. A system of LANs connected in this way is called a wide-area network (WAN). Most LANs connect workstations and personal computers. Each node (individual computer) in a LAN has its own CPU with which it executes programs, but it is also able to access data and devices anywhere on the LAN. This means that many users can share expensive devices, such as laser printers, as well as data. Users can also use the LAN to communicate with each other, by sending email or engaging in chat sessions. There are many different types of LANs and Ethernet LANs is the most common for PCs. Most Apple Macintosh networks are based on Apple's AppleTalk network system, which is built into Macintosh computers. LANs are capable of transmitting data at very fast rates, much faster than data can be transmitted over a telephone line; but the distances are limited, and there is also a limit on the number of computers that can be attached to a single LAN.

LDAP

LDAP stands for Lightweight Directory Access Protocol, a set of protocols for accessing information directories. LDAP should eventually make it possible for almost any application running on virtually any computer platform to obtain directory information, such as email addresses and public keys. Because LDAP is an open protocol, applications need not worry about the type of server hosting the directory.

LFS

Local File System. A file system type developed by Auspex and used in the NetServer for filesystem communication between the network processors and the file processor and between the host processor and the file processor. LFS provides local file operations similar to NFS remote operations, but without the protocol processing overhead.

Load Balancer

Load balancers can prevent servers from being overloaded by distributing large numbers of access requests across multiple Internet servers. The requests are directed to the most-available server based on how busy that server is and other considerations such as storage configuration, server power etc. Load balancers can take the form of special software that runs on a dedicated server, a proprietary black box, or a capability in an advanced LAN switch.

MAC

Media Access Controls or MACs are the rules defined within a specific network type that determines how each station accesses the network cable. Using a token-passing method, a carrier sensing and collision detection method or a demand priority method prevents simultaneous access to the cable. The MAC used for 100BaseT Ethernet as implemented in the Auspex EtherBand™ feature is based on the "demand priority" access method in which the central hub scans all its ports in a round-robin fashion to detect stations that want to transmit a frame. Higher priorities can be requested by ports to transmit real-time information like video or audio.

Megabyte
1,000,000 bytes.

MCAD
Mechanical Computer Aided Design

MIB
Management Information Base is a set of standards for detailed system information that is reported to a control console for SNMP compliance. It's intent is to provide common metrics for heterogeneous computer systems.

MTBF
Mean Time Between Failure. A key component of the availability equation, AVAILABILITY = (MTBF – MTTR) ÷ MTBF. For example, a server that on average fails once every 5,000 hours and on average takes 2 hours to diagnose, replace faulty components and reboot would have an availability rating of (5,000 – 2) ÷ 5,000 = 99.96%.

MTTR
Mean Time To Repair. Includes the time taken to diagnose the failure, replace or repair faulty component(s) and reboot the system. See MTBF.

NAP
A NAP or Short for Network Access Point, a public network exchange facility where Internet Service Providers (ISPs) can connect with one another in peering arrangements. The NAPs are a key component of the Internet backbone because the connections within them determine how traffic is routed. They are also the points of most Internet congestion.

NAS
Network Attached Storage.

NDMP
NDMP is a standard protocol for network-based backup of network-attached storage. NDMP hides the unique interfaces from third party backup software which allows this software to execute on any NDMP compliant system on the network (such as the NS2000 Host Node), and control backups on the NS2000 using standard commands.

NeTservices
The Auspex software product that provides for consistent high performance UNIX and NT file services and makes the NS2000 a true bilingual file server.

Newsgroup servers
A newsgroup server is a server that servers as a forum for an on-line discussion group. On the Internet, there are literally thousands of newsgroups covering every conceivable interest. To view and post messages to a newsgroup, you need a newsreader, a program that runs on your computer and connects you to a news server on the Internet. USENET is a worldwide bulletin board system that can be accessed through the Internet or through many online services. The USENET contains more than 14,000 forums, called newsgroups that cover every imaginable interest group. It is used daily by millions of people around the world. See also USENET.

NIC
Network Interface Cards (or NICs) in the NS2000 support 10/100BaseT Ethernet, Gigabit Ethernet, or ATM. There are from one to three on each I/O Node.

NFS

Network File System. NFS is an ONC application-layer protocol for peer-to-peer, distributed, file system communication. NFS allows a remote file system (often located on a file server) to be mounted transparently by client workstations. The client cannot perceive any functional difference in service between remote and local file systems (with trivial exceptions). NFS is the most popular ONC service, has been licensed to over 300 computer system vendors, runs on an estimated 10 million nodes and is a de facto UNIX standard. See also VFS, ONC, and NFSv3.

NFSv3

NFS version 3. References to NFS generally imply NFS version 2 protocol. NFS version 3 is an update to the NFS protocol. Significant among the many changes made for NFSv3 are the adoption of a safe asynchronous write protocol and the use of block sizes up to 64 KB. Other protocol changes are intended to improve the overall network and client efficiency and provide improved support for client-side caching.

NFS ops/s

NFS operations per second. Typical NFS operations include: lookup, read, write, getattr, readlink, readdir, create, remove, setattr, and statfs.

Node

See Fibre Channel.

NP

The Network Processor (NP) refers to one of the two Intel Pentium processors on an I/O Node of an NS2000 system. This processor runs highly optimized microcode that manages all network processing of the NS2000 I/O Node and communicates with other I/O Nodes and the Host Node. See also File System and Storage Processor (FSP).

NTFS

Windows NT File System.

NVM

Non volatile memory is a term used to refer to battery backed up DRAM so that data will not be lost in the event of power failure.

NVRAM

Non-volatile random access memory such as static RAM will not lose data in the event that power is lost to the chip.

ONC

Open Network Computing is the trade name for the suite of standard IP-based network services, including RPC, XDR and NFS. ONC was promulgated by Sun Microsystems.

Online Services

An online service such as Microsoft Network (MSN), America Online (AOL) or CompuServe is a business that provides its subscribers with a wide variety of data transmitted over telecommunications lines, including web access. Subscribers can get almost any information that has been put in electronic form. Of course, accessing all this data carries a price and an Online Service's monthly fees are usually more expensive than for simple Internet access from an ISP.

Operating System

The Operating System (OS) performs basic tasks on a computer such as recognizing input from a keyboard, sending output to the display screen, keeping track of files and directories on the disk and controlling peripheral devices such as disk drive and printers or a mouse. The OS acts as a traffic cop and schedules the various programs that the computer executes. The OS is also responsible for security, ensuring that unauthorized users do not access the system. Operating systems can be classified as follows: 1) Multi-user – allows two or more users to run programs at the same time. 2) Multiprocessing – supports running a program on more than one CPU. 3) Multi-tasking – allows more than one program to run concurrently. 4) Multi-threading – allows different parts of a single program to run concurrently. 5) Real Time – Usually a stripped down OS that responds to input instantly.

OSI

OSI is an abbreviation for Open System Interconnection, an ISO standard for worldwide communications that defines a networking framework for implementing protocols in seven layers. Control is passed from one layer to the next, starting at the application layer in one station, proceeding to the bottom layer, over the channel to the next station and back up the hierarchy. At one time, most vendors agreed to support OSI in one form or another, but OSI was too loosely defined and proprietary standards were too entrenched. Except for the OSI-compliant X.400 and X.500 email and directory standards, which are widely used, what was once thought to become the universal communications standard now serves as the teaching model for all other protocols. Most of the functionality in the OSI model exists in all communications systems, although two or three OSI layers may be incorporated into one.

Packet

A packet is a piece of a message transmitted over a packet-switching network. One of the key features of a packet is that it contains the destination address in addition to the data. In IP networks, packets are often called datagrams. Parallel processing refers to when a single computer simultaneously uses more than one CPU to execute a program. Ideally parallel processing makes a program run faster because there are more CPUs running it. In practice, it is often difficult to divide a program so that separate CPUs can execute different portions without interfering with each other. Among NAS vendors, only the Auspex NS2000 effectively overcomes this problem by designing each I/O Node with two processors each performing separate portions of the network file-serving task. In addition the NS2000 links multiple I/O Nodes together by a highly efficient Scaleable Coherent Interface (SCI) interconnect which allows the multiple nodes to act as one system. See also Functional Multiprocessing (FMP).

PCI

The Peripheral Channel Interconnect is an ANSI standard for an I/O bus used predominantly in PC design.

PCX

PCX is a graphics file format for many graphics programs running on PCs. Most optical scanners, fax programs, and desktop publishing systems support it. Files in the PCX format end with a ".pcx" (pronounced dot -p-c-x) extension. Two other common bit map formats are BMP and TIFF.

PDA

Personal Digital Assistant, a handheld device that combines computing, telephone/fax, and networking features.

PDC

One server running Windows NT Server acts as the Primary Domain Controller (PDC), that maintains the centralized security databases for the domain. Other computers running Windows NT Server in the domain function as Backup Domain Controllers (BDC) and can authenticate logon requests. The PDC or BDC authenticates users of a Windows NT Domain.

PDF

PDF stands for Portable Document Format, a file format developed by Adobe Systems. PDF captures formatting information from a variety of desktop publishing applications, making it possible to send formatted documents and have them appear on the recipient's monitor or printer as they were intended. To view a file in PDF format, you need Adobe Acrobat Reader, a free application distributed by Adobe Systems.

PDU

Power Distribution Unit or Power Shelf in the NS2000. A cabinet model contains from three to seven power supplies and a stack model contains up to three power supplies. Both types of PDUs can be N+1 redundant.

PKZIP

PKZIP is one of the most widely used file compression methods. PKZIP was developed by PKWARE, Inc. in 1989 and distributed as shareware. Files that have been compressed using PKWARE are said to be zipped. Decompressing them is called unzipping.

Plug In

A hardware or software module that adds a specific feature or service to a larger system. For example, there are number of plug-ins for the Netscape Navigator browser that enable it to display different types of audio or video messages. Navigator plug-ins are based on MIME file types.

PNG

PNG is short for Portable Network Graphics, and pronounced ping, a new bit-mapped graphics format similar to GIF. In fact, PNG was approved as a standard by the World Wide Web consortium to replace GIF because GIFF uses a patented data compression algorithm called LZW. In contrast, PNG is completely patent- and license-free. The most recent versions of Netscape Navigator and Microsoft Internet Explorer now support PNG.

Port / Port ID

See Fibre Channel.

POP

POP (1) Given a stack of items, popping one of the items means to pull it off the stack. The opposite of pop is push, which means to move an object onto a stack. (2) Short for Point of Presence. (3) Short for Post Office Protocol.

POTS

POTS or plain old telephone service refers to the standard telephone service that most homes use. In contrast, telephone services based on high-speed, digital communications lines, such as ISDN and FDDI, are not POTS. The main distinctions between POTS and non-POTS services is bandwidth. POTS is generally restricted to a maximum of 52 Kbps (52,000 bits per second). The POTS network is also called the public switched telephone network (PSTN).

PPP

PPP stands for Point-to-Point Protocol, a method of connecting a computer to the Internet. PPP is more stable than the older SLIP protocol and provides error-checking features. See also SLIP.

Proxy Server

Also known as an intelligent Web cache, a Proxy Server is a server that sits between a client application, such as a Web browser, and a real server. It intercepts all requests to the real server to see if it can fulfill the requests itself. If not, it forwards the request to the real server. Proxy servers have two main purposes: 1) Improve Performance - Proxy servers can dramatically improve performance for groups of users. This is because it saves the results of all requests for a certain amount of time. 2) Filter Requests - Proxy servers can also be used to filter requests. For example, a company might use a proxy server to prevent its employees from accessing a specific set of Web Sites.

RPC

Remote Procedure Call. An RPC is an (almost) transparent subroutine call between two computers in a distributed system. ONC RPC is a Sun-defined session-layer protocol for peer-to-peer RPC communication between ONC hosts. ONC RPC underlies NFS.

QoS

Quality of Service

RAID

Redundant Array of Independent Disks. RAID is used to increase the reliability of disk arrays by providing redundancy either through complete duplication of the data (RAID 1, i.e., mirroring) or through construction of parity data for each data stripe in the array (RAID 3, 4, 5). RAID 5, which distributes parity information across all disks in an array, is among the most popular means of providing parity RAID since it avoids the bottlenecks of a single parity disk. RAID Controllers The NS2000 RAID controllers provide a highly optimized scheme for securely managing RAID configurations on NS2000 systems. The Auspex RAID controllers allow RAID arrays to be expanded online, and support conversion of an array from one RAID level to another.

Reader

A file reader is a software program that reads all the file layout characteristics of a particular format so that it can be displayed on the computer screen.

Router

A Router is a device that connects any number of LANs. Routers use headers and a forwarding table to determine where packets go, and they use ICMP to communicate with each other and configure the best route between any two hosts. Very little filtering of data is done through routers. Routers do not care about the type of data they handle. See also ICMP.

Search Engine

A program that searches documents for specified keywords and returns a list of the documents where the keywords were found. Although search engine is really a general class of programs, the term is often used to specifically describe systems like Google and Yahoo that enable users to search for documents on the World Wide Web and USENET newsgroups. Typically, a search engine works by sending out a spider to fetch as many documents as possible. Another program, called an indexer, then reads these documents and creates an index based on the words contained in each document. Each search engine uses a proprietary algorithm to create its indices such that, ideally, only meaningful results are returned for each query.

SCI

Scaleable Coherent Interface is an ANSI standard (#1596-1992) that is the modern equivalent of a processor-memory-I/O bus and a Local Area Network combined and made parallel to support distributed multiprocessing. The SCI interconnect has very high bandwidth, very low latency and a scaleable architecture. This allows building large high performance systems and is used by Convex/HP supercomputers, Sun Clusters, Sequent, Auspex and others. Network latency has been measured at 150 times less than previous network connections for efficient and fast communication between computer nodes.

SCSI

Small Computer System Interface. An intelligent bus-level interface that defines a standard I/O bus and a set of high-level I/O commands. The SCSI busses in the NS2000 are used to connect multiple peripheral devices such as disk drives tape drives. Each SCSI device has an intelligent SCSI controller built into it. There are currently many flavors of SCSI defined by different bus widths and clock speeds. The seven major variations of SCSI are SCSI 1, SCSI 2 (Fast / Narrow), SCSI 2 (Fast / Wide), Ultra SCSI (Fast / Narrow), Ultra SCSI (Fast / Wide) – also called SCSI 3, Ultra 2 SCSI (Narrow), Ultra 2 SCSI Wide. Single ended SCSI is used when the peripheral device is close to the point of attachment as in the NS2000 method of attaching disk drives. Differential SCSI provides for reliable operation over greater distances and is used in the NS2000 for tape drive connections.

SDSL

SDSL is short for symmetric digital subscriber line, a new technology that allows more data to be sent over existing copper telephone lines (POTS). SDSL supports data rates up to 3 Mbps. SDSL works by sending digital pulses in the high-frequency area of telephone wires. Since these high frequencies are not used by normal voice communications, SDSL can operate simultaneously with voice connections over the same wires. SDSL requires a special SDSL modem. SDSL is called symmetric because it supports the same data rates for upstream and downstream traffic. A similar technology that supports different data rates for upstream and downstream data is called asymmetric digital subscriber line (ADSL). ADSL is more popular in North America, whereas SDSL is being developed primarily in Europe.

SE

System Engineer(s) perform a variety of technical pre and post sales services for customers and prospects.

SID

An NT term meaning System Identification (SID).

SLIP

SLIP stands for Serial Line Internet Protocol, a method of connecting to the Internet. Another more common method is PPP (Point-to-Point Protocol). SLIP is an older and simpler protocol, but from a practical perspective, there's not much difference between connecting to the Internet via SLIP or PPP. In general, service providers offer only one protocol although some support both protocols. See also PPP.

SMB

Server Message Block protocol. See CIFS.

SMTP

SMTP or Simple Mail Transfer Protocol, is the predominant Internet protocol for sending email messages between servers. Most email systems that send mail over the Internet use SMTP to send messages from one server to another; the messages can then be retrieved with an email client using either POP or IMAP. In addition, SMTP is generally used to send messages from a mail client to a mail server. This is why you need to specify both the POP or IMAP server and the SMTP server when you configure your email application. See also POP and IMAP.

SMP

Symmetric Multi-Processing. A computer architecture in which processing tasks are executed in parallel on multiple, identical, general-purpose CPUs that share a common memory. SMP computer systems usually have modified operating systems that can themselves execute concurrently. The SMP architecture offers high computational throughput, but not necessarily high I/O throughput. See FMP.

Snapshot

A term that refers to a copy of a file system at a certain point in time. Snapshots are used for backup and recovery.

SNMP

Simple Network Management Protocol. SNMP is a protocol used for communication between simple, server-resident SNMP agents that respond to network administration requests from simple-to-sophisticated SNMP manager tools running on remote workstations.

Socket

A socket is the software that is required for a computer to understand and execute TCP/IP protocol is known as a socket or TCP/IP stack. This software acts as an intermediary between the Internet and a computer.

Solaris

Sun's UNIX operating system.

Spam

Spam or electronic junk mail or junk newsgroup postings. Some people define Spam even more generally as any unsolicited email. However, if a long-lost brother finds your email address and sends you a message, this could hardly be called Spam, even though it's unsolicited. Real Spam is generally email advertising for some product sent to a mailing list or newsgroup.

SPARC

Scalable Processor Architecture. SPARC International's specification for the Reduced Instruction-Set-Computer (RISC) CPUs found in systems sold by Sun Microsystems, Auspex, etc.

Splitter

A splitter divides the bandwidth from the coaxial cable coming from a cable supplier into the signal going to the high-speed cable modem and the signal going to the TV set.

Stack

1) In networking, short for protocol stack. (2) In programming, a special type of data structure in which items are removed in the reverse order from that in which they are added, so the most recently added item is the first one removed. Adding an item to a stack is called pushing. Removing an item from a stack is called popping.

Switch

In networking terminology, a switch is a computing device that filters and forwards packets between Local Area Network (LAN) segments. Switches operate at the data link layer (layer 2) of the OSI Reference Model and therefore support any packet protocol. A special type of switch call an L4 switch operates at the fourth layer (Transport Layer) of the OSI Reference Model and is responsible for the integrity of data transmissions between LAN segements. LANs that use switches to join segments are called switched LANs or, in the case of Ethernet networks, switched Ethernet LANs.

T1

T1 refers to a dedicated phone connection supporting data rates of 1.544Mbits per second. A T-1 line actually consists of 24 individual channels, each of which supports 64Kbits per second. Each 64Kbit/second channel can be configured to carry voice or data traffic. Most telephone companies allow you to buy just some of these individual channels, known as fractional T-1 access. T-1 lines are a popular leased line option for businesses connecting to the Internet and for Internet Service Providers (ISPs) connecting to the Internet backbone. The Internet backbone itself consists of faster T-3 connections. T-1 lines are sometimes referred to as DS1 lines.

T3

T3 A dedicated phone connection supporting data rates of about 43 Mbps. A T-3 line actually consists of 672 individual channels, each of which supports 64 Kbps. T-3 lines are used mainly by Internet Service Providers (ISPs) for connecting to the Internet backbone and for the backbone itself. T-3 lines are sometimes referred to as DS3 lines.

TB

A Terabyte (TB) equals 1,000 Gigabytes.

TCP

Transmission Control Protocol or TCP is a transport layer component of the Internet's TCP/IP protocol suite. It sits above IP in the protocol stack and provides reliable data delivery services over connection-oriented links. TCP uses IP to deliver information across a network and makes up for the deficiency of IP providing a guarantee of reliable delivery services that IP does not. TCP messages and data are encapsulated into IP datagrams and IP delivers them across the network.

TIFF

Tagged Image File Format, one of the most widely supported file formats for storing bit-mapped images on personal computers (both PCs and Macintosh computers). TIFF graphics can be any resolution, and they can be black and white, gray-scaled, or color. Files in TIFF format often end with a .tif extension.

TLD

Top level domain refers to the suffix attached to the Internet Domain names such as .com, .gov, mil etc. New TLDs are being added as the Internet grows. For example the Post Office recently changed from a .gov TLD to its own TLD .usps.

URL

URL is an abbreviation of Uniform Resource Locator, the global address of documents and other resources on the World Wide Web. The first part of the address indicates what protocol to use, and the second part specifies the IP address or the domain name where the resource is located.

USENET

USENET is a worldwide bulletin board system that can be accessed through the Internet or through many online services. The USENET contains more than 14,000 forums, called newsgroups that cover every imaginable interest group. It is used daily by millions of people around the world.

Uuencode

Uuencode is a set of algorithms for converting files into a series of 7-bit ASCII characters that can be transmitted over the Internet. Originally, uuencode stood for Unix-to-Unix encode, but it has since become a universal protocol used to transfer files between different platforms such as Unix, Windows, and Macintosh. UUencoding is especially popular for sending email attachments. Nearly all email applications support UUencoding for sending attachments and uuecoding for receiving attachments.

vBNS

vBNS is an abbreviation for very high-speed Backbone Network Service, an experimental wide-area network backbone sponsored by the National Science Foundation (NSF) and implemented by MCI. vNBS has replaced NSFnet and is designed to serve as a platform for testing new, high-speed Internet technologies and protocols. It currently links several Supercomputer Centers (SCCs) and Network Access Points (NAPs) at OC-12 speeds (622 Mbps). By 1998, it will support data, voice and video traffic at 2.5 Gbps.

VPN

VPN is an abbreviation for virtual private network, a network that is constructed by using public wires to connect nodes. For example, there are a number of systems that enable you to create networks using the Internet as the medium for transporting data. These systems use encryption and other security mechanisms to ensure that only authorized users can access the network and that the data cannot be intercepted.

WAN

Wide Area Networks or WANs are networks of computers that are geographically dispersed and connected by radio waves, telephone lines or satellites.

WAV

The defacto standard for storing sound in files that was developed by IBM and Microsoft.

Web

A system of Internet servers that support specially formatted documents. The documents are formatted in a language called HTML (HyperText Markup Language) that supports links to other documents, as well as graphics, audio, and video files. This means you can jump from one document to another simply by clicking on hot spots. Not all Internet servers are part of the World Wide Web. There are several applications called Web browsers that make it easy to access the World Wide Web; Two of the most popular being Netscape Navigator and Microsoft's Internet Explorer.

Web Site

A site (location) on the World Wide Web. Each Web site contains a home page, which is the first document users see when they enter the site. The site might also contain additional documents and files. Each site is owned and managed by an individual, company or organization.

Webmaster

An individual designated at a Web Site to be in control of many of the operational aspects of the site.

Web TV

Web TV is a general term for a whole category of products and technologies that enable you to connect to the web through your TV set.

Workload

An I/O workload refers to the pattern of I/O activity that a disk RAID array is subjected to by a certain application software program or mix of application software programs.

xDSL

xDSL refers collectively to all types of digital subscriber lines, the two main categories being ADSL and SDSL. Two other types of xDSL technologies are High-data-rate DSL (HDSL) and Single-line DSL (SDSL). DSL technologies use sophisticated modulation schemes to pack data onto copper wires. They are sometimes referred to as last-mile technologies because they are used only for connections from a telephone switching station to a home or office, not between switching stations. xDSL is similar to ISDN inasmuch as both operate over existing copper telephone lines (POTS) and both require the short runs to a central telephone office (usually less than 20,000 feet). However, xDSL offers much higher speeds - up to 32 Mbps for downstream traffic, and from 32 Kbps to over 1 Mbps for upstream traffic.

$12.95

Corporate Headquarters
2800 Scott Boulevard
Santa Clara, CA 95050
USA
408.566.2000 **phone**
800.735.3177 **fax**

http://www.auspex.com **Web**

Japan
Auspex K.K.
ATT Shinkan, 6K
2-11-7 Akasaka
Minato-ku
Tokyo 107-0052
Japan
+81 3 3586 1751 **phone**
+81 3 3586 1778 **fax**

United Kingdom
Auspex Systems, Ltd
Portland House
Aldermaston Park
Church Road
ALDERMASTON
Reading RG7 4 HR
United Kingdom
+44 1189407300 **phone**
+44 1189407333 **fax**

France
Auspex Systems, SA
40, Boulevard Henri Cedex
France
+33 1 42 04 9150 **phone**
+33 1 42 04 9151 **fax**

Germany
Auspex Systems, GmbH
Ingolstaedter Stasse 22
80807 Muenchen
Germany
+49 89 350151-0 **phone**
+49 89 350151-50 **fax**

Auspex Asia Pacific
#06-02 International Building
360 Orchard Road
Singapore
238869
+65 732 9770 **phone**

Auspex Worldwide

www.auspex.com

AUSPEX