Internet Service Provider Configuration Guidelines

Version 2.5

February, 1999

Sun microsystems

901 San Antonio Road, Palo Alto, CA 94303 U.S.A.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

THIS PUBLICATION IS PROVIDED “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT.

THIS PUBLICATION COULD INCLUDE TECHNICAL INACCURACIES OR TYPOGRAPHICAL ERRORS. CHANGES ARE PERIODICALLY ADDED TO THE INFORMATION HEREIN; THESE CHANGES WILL BE INCORPORATED IN NEW EDITIONS OF THE PUBLICATION. SUN MICROSYSTEMS, INC. MAY MAKE IMPROVEMENTS AND/OR CHANGES IN THE PRODUCT(S) AND/OR THE PROGRAM(S) DESCRIBED IN THIS PUBLICATION AT ANY TIME.

This paper was made possible through the contributions of Steve Gaede, James Baty, Vincent Costa, Abe Schoenfeld, Jaap Romers, Ellen Siegel, James Simon, Michael Wallace, and Dave Williams.

BRUCE BAIKIE, Febuary 1999

Contents

1. Introduction 1

2. The Nature of ISPs 7

Basic ISP Functionality 8

Local, Regional, and National ISPs 10

3. ISP Infrastructure 11

Electronic Mail 11

Netnews 17

World Wide Web Services 17

Domain Name Service 20

Security with Routers and Firewalls 22

ISP Infrastructure 23

Summary 29

4. Architectural Principles 31

Scalability 31

Functional Decomposition 33

Reliability and Availability 37
Architectures for Security 39
Summary 41
5. Ensuring ISP Security 43
Paradoxical Security Requirements 43
Policy and Implementation 45
A Range of Security Options 47
Network Design 48
Limiting Access 50
Active Firewalling 51
The Role of Encryption 54
Summary 59
6. High-Availability Solutions For ISPs 61
Fault Tolerance vs. High Availability 61
High Availability and Load Balancing 63
Symmetric and Asymmetric Configurations 66
Replicating, Partitioning, and Sharing Data 67
Three-Tier Approach to HA 69
Summary 73
7. Value-Added Web Hosting 75
Static vs. Dynamic Content 75
Dedicated vs. Shared Hosting 79
Site Management 83

Security Issues 89

Summary 92

8. Local ISP Configurations 93

Characteristics of Local ISPs 93

Architectural Overview 94

Access Network 94

Services Network 97

Internal Network 101

Summary 101

9. Regional and National ISPs 103

Characteristics of Regional ISPs 103

Characteristics of National ISPs 104

Architectural Overview 105

Access Network 107

Demilitarized Zone 108

Services Network 109

Storage Network 110

Internal Network 110

Modifications for the Telco Environment 110

Summary 111

10. Conclusion 113

11. References 115

Introduction 1

The Internet Service Provider (ISP) market continues to defy prediction. As recently as two years ago, analysts were expecting a dramatic consolidation that would eliminate all but the largest providers. Indeed, some consolidation has occurred — however many local and regional ISPs in the U.S. have solidified their market niche by focusing on maintaining close, service-based relationships with their customers. The consolidation that has occurred in the U.S. has more than been offset by a rapid increase in the number of ISPs in Europe and Asia resulting from deregulation in these regions.

The registry at *thelist.internet.com* lists more than 5,600 providers as of January 1999, indicating that the overall number of ISPs is increasing at a moderate pace — however in traffic and number of subscribers, the market continues to explode. Internet traffic exceeded voice traffic for the first time in 1998, *surpassing in less than a decade* the telephony market that has been *growing for nearly a century!* A 1998 survey by Infonetics Research finds that local and regional ISPs expect to nearly double their customer base, number of access ports, and number of up-stream connections in the course of the year.

Emerging Markets

Only a few years ago, the market consisted of a large number of ISPs and a handful of Internet backbone providers. Today, the dramatic increase in the number of Internet subscribers and the traffic they generate has fueled the emergence of three distinct players: Internet Service Providers, Network Service Providers, and Application Service Providers. Each of these providers are developing their specific market while making inroads into complementary areas. Those companies which provide services in all three areas are known as Full Service Providers (Figure 1):

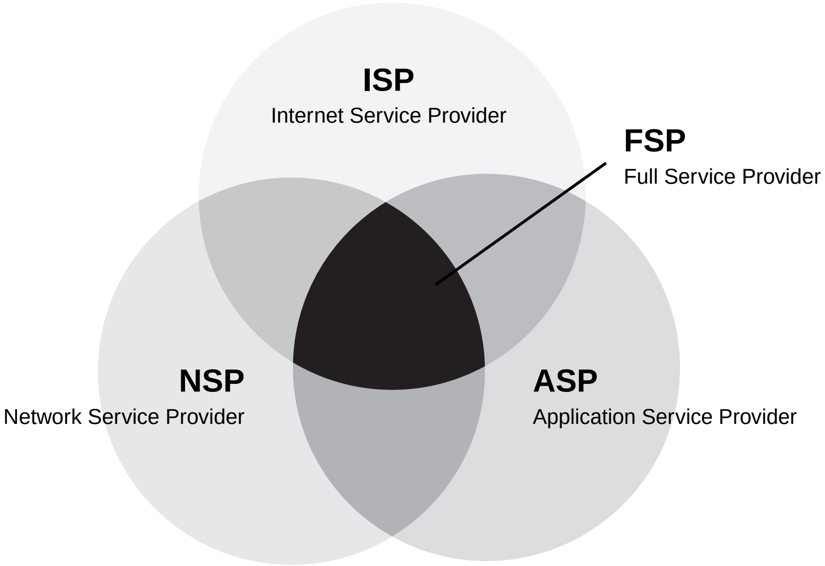

Figure 1 The ISP market is currently segmented into ISPs, ASPs, and NSPs, with some Full Service Providers addressing the needs of each segment.

- *Internet Service Providers (ISPs)*

  Internet Service Providers continue to provide millions of customers with on-ramp access and a wide range of services including e-mail, Web site hosting, local content provision, corporate firewalls, and virtual private networks. Outside of the U.S., many ISPs have been regulated entities affiliated with government postal services and telephone companies, but the current climate of deregulation has fostered an international market of competing ISPs. And where once each foreign ISP had their own separate — and costly — high-speed connections to the United States, now these ISPs are free to aggregate service using international and global network service providers, resulting in lower costs and increased business.

- *Application Service Providers (ASPs)*

  Recognizing the capability of the Internet to deliver application services to businesses in a secure manner, a growing number of ISPs and application vendors are acting in the capacity of Application Service Providers. In the U.S., some telephone company ISPs are beginning to enable access to small business accounting services over the Internet. As a way to meet the needs of small and medium-sized companies which cannot justify the outlay for large computing systems and costly licenses, application vendors offer complete packaged services like Oracle Financials and Enterprise Resource and Planning applications.

- *Network Service Providers (NSPs)*

  The enormous increase in Internet traffic has stimulated significant growth in providers making up the Internet backbone. National providers like BBN and UUNET have reached out into international markets. There is a whole range of “new generation” providers like Qwest, which uses existing pipeline and railroad right-of-ways for the installation of extensive fiber optic networks. Meanwhile, regional NSPs like Merit continue to provide geographically-limited services.

  The growth in NSPs has resulted in a large increase in network Points-Of-Presence (POPs), where ISPs obtain their bandwidth; and network exchanges, where traffic crosses between NSPs. These providers have begun to make inroads into both the ISP and ASP arenas by providing Web site hosting and co-location services where servers are located close on the backbone. This method results in faster access, which leads to increased traffic.

Capitalizing on the tremendous opportunities for growth and innovation, many ISPs are already crossing over into ASP and NSP areas. Likewise many Network and Application Service Providers are crossing over into complimentary markets. Indeed, some providers — like Full Service Provider GTE — already offer all three service areas to their customers.

Regardless of market niche, all types of providers must position themselves for growth and agility to handle increasing numbers of subscribers, additional services, and more challenging workloads. System architectures that meet these demands are critical to success.

Meeting the Challenge with Sun

Sun Microsystems is uniquely positioned at the forefront of the hardware and software technologies that support service providers of all types. Sun has been committed to the Internet since 1982, when its first freely-distributed software was made available on FTP sites through the SunSITE program. Sun also has been a leader in intranet technology, with its own internal network spanning 44 countries and hosting more than 1,800 Web servers. Meanwhile, Sun's Java™ technology — freely-available for almost every computing system — is speeding the maturation of the Internet from an informational utility to a true electronic marketplace.

The innovations that have made Sun's own internal network so successful are now available to customers through a wide range of workstation, server, and software products. Sun's Ultra™ 5, 10, 60, and 450 workstations provide high-performance desktop computing capabilities with up to four 300 MHz UltraSPARC™ processors and 4 GB of main memory. Sun's server products — ranging from the single-processor Ultra 5S to the 64-processor Enterprise 10000 server — provide the computing power and headroom required by ISPs, NSPs, and especially compute-hungry ASPs. For ISPs requiring rugged systems with NEBS Level 3 certification, Sun's Netra™ t 1120 and 1125 servers provide single- and dual-processing capabilities, respectively.

The Solaris™ operating environment provides the cornerstone of Sun's line of ISP-ready software products, delivering the reliable, mature networking capabilities that built the Internet. Solaris delivers today what the competition only promises for the future, and is the basis for Sun's ISP-specific products including Solaris ISP Server™, Sun™ Internet Mail Server™, and the Solaris™ Bandwidth Manager. For proven network security, SunScreen™ SPF-200 and SunScreen EFS™ software can be hosted on any of Sun's server platforms. SunScreen SPF-200 provides a hardened version of Solaris with packet-filtering software that is impervious to intrusion and undetectable on the network. All of these are reasons why 89 percent of telephone company ISPs are based on Sun hardware and software products.

ISP products and services are more than just low-cost high-performance scalable servers — Sun understands that the way these hardware solutions are deployed and integrated into an ISP platform is just as important as the servers themselves. Sun and its integration affiliates know that all ISP installations are unique and require the utmost care in design and deployment.

Well-Kept Secrets

Most of the configuration details that make a successful ISP installation are well-kept secrets — however there are some basic guidelines and principles that can be applied to many ISP architectures. This document presents some of the approaches that can position ISP infrastructures for performance, scalability, security, and flexibility, focusing on the areas which are of common interest to ISPs, ASPs, and NSPs.

Chapter 2 discusses the nature of ISP services today, and chapter 3 presents some of the basic software components needed by Internet Service Providers. The effectiveness of software is only as good as its deployment onto a hardware platform, and chapter 4 discusses some of the architectural principles involved. Chapters 5 and 6 examine the areas of security and high availability, respectively, in more detail. Chapter 7 discusses the issue of application hosting, in particular how to configure value-added Web hosting services. Chapter 8 presents an architecture for an ISP that can support between 5,000 and 10,000 subscribers. Finally, chapter 9 presents some of the architectural considerations that must be made to accommodate regional and national ISPs with customer bases of more than 50,000 subscribers.

This document cannot — indeed no document can — make the reader an expert in ISP configuration. Innovations in both hardware and software happen so quickly that it is sometimes difficult to get the most current information. But the stakes are high, and up-front investments in well-planned architectures will pay off in scalable, high-performance, secure, reliable networks that withstand the test of time. This is why there is no substitute for having the most up-to-date architecture possible by working with Sun's integration affiliates or with Sun Professional Services^{SM} from the beginning. Sun Professional Services has been a key facilitator for the highest-performance and largest-scale ISPs in the world.

The Nature of ISPs 2

The services that customers expect of Internet Service Providers have changed dramatically in the last several years. When fixed, high-speed connections to the Internet were the rule, only businesses, educational institutions, and government agencies participated in the world-wide interconnection of networks. Dial-up services provided only terminal emulation. This changed in the late 1980’s as protocols were developed to support Internet Protocol (IP) connections over serial links such as those established over dial-up telephone lines using modems. These well-known protocols — Serial Line Internet Protocol (SLIP), and later Point-to-Point Protocol (PPP) — have enabled millions of individual users and small businesses to enjoy low-cost entry into the world of the Internet. Dial-up services today allow direct connections to the Internet.

While SLIP and PPP enabled casual Internet connections, few individuals were attracted to Internet Service Providers for the ability to use applications such as FTP and telnet. The invention of the World Wide Web — along with the Mosaic Web browser — provided a graphical user interface for accessing and displaying information provided by other hosts. The combination of serial IP protocols and the “killer application” of the Internet provided the synergy that has caused use to skyrocket.

The rapidly-growing ISP market is born of these two technologies, and now Internet access is demanded by millions of new users every year. Whereas once ISPs had only to provide high-speed connectivity to a small number of highly-technical customers, they now have to provide network-based applications, content, and value-added network services to attract and keep the new wave of business and residential users.

Basic ISP Functionality

As the Internet market continues to mature, the nature of what it takes to be an ISP has changed dramatically. Internet Service Providers must now deliver attractive packages that include the ability to store and retrieve mail; to access huge numbers of netnews articles; to view and host Web sites; to view streaming video clips; to interact in chat rooms and to communicate using voice over IP. In order to provide these value-added services, companies entering the ISP marketplace must ensure that their infrastructure includes the full range of basic functionality:

- *ISP Applications*

  The base of applications that ISPs are expected to provide to their users includes electronic mail, netnews, World Wide Web access, and Web page hosting. More sophisticated users may also demand access to other Internet services such as Virtual Private Networks (VPNs), chat, phone servers, FTP, telnet, gopher, and archie. Decisions must be made regarding which services will be provided to ISP customers — and in large part these choices are what differentiate one ISP from another.

- *System Integrity*

  ISPs must take the utmost care to ensure both the availability and integrity of customer data. Therefore the network through which an ISP provides services must be protected against breaches of security and against hardware and software failures. Firewalls are a necessity for protecting an ISP's core services from intrusion from both the Internet and its customers. High-availability (HA) for ISP applications can be achieved by deploying multiple-server clusters like Sun Clusters, which provide integrated configurations for both NFS™ and many Internet server applications. System integrity must also be ensured backups.

- *Full Range of Connectivity*

  At minimum, today's ISPs provide dial-in service using the public switched telephone network (PSTN). Cable-based ISPs provide access through cable modems and routers at the head-end site. High-bandwidth connections using ISDN, ADSL, leased lines, ATM, and Frame Relay are needed to support more sophisticated residential and home office customers, as well as most commercial subscribers.

Points of Presence (POPs) extend an ISP’s home service area, reaching a broader customer base, and allowing customers to access ISP services while travelling. The more successful ISPs are those who establish roaming arrangements with other service providers, allowing national- and world-wide access to a uniform set of services.

- *System and User Management*

  Service Providers must have well thought-out strategies for managing their own network and the potentially explosive growth that results from providing a high level of service to customers. Basic facilities such as Domain Name Service (DNS) must be established to provide addressibilty to ISP services and customer-owned domain names. Lightweight Directory Access Protocol (LDAP) is used to coordinate subscriber information across a variety of services and operating system platforms. Network management tools like Sun’s Solstice Enterprise Manager™ helps to manage routers, firewalls, and the core servers themselves. Effectively managing users and creating new accounts becomes increasingly critical as an ISP grows. Sun’s affiliates observe that user administration costs tend to peak at around 20,000 subscribers, after which it is essential to establish automatic subscription mechanisms and credit-card billing.

Local, Regional, and National ISPs

Internet Service Providers tend to be clustered in three general size categories (Table 1). Local ISPs begin with configurations for 5,000 to 10,000 subscribers and scale upwards from this range. Regional ISPs begin in the range of 50,000 to 75,000 subscribers. National ISPs have 150,000 to 200,000 subscribers, and sometimes more than one million subscribers. Each category of provider has its own characteristics and challenges, many of which are discussed in this document.

ISP Categories	Number of Subscribers
Local	5,000 to 10,000
Regional	50,000 to 75,000
National	150,000 to 200,000 and more

Table 1 ISPs categorized by size

ISP Infrastructure 3

This chapter discusses the major Internet Service Provider components and infrastructure from a software standpoint. Core components that ISPs deploy include electronic mail, netnews, the World Wide Web, domain name service, routers, and firewalls. Just as important as the components providing the ISP services is the ISP infrastructure itself, including operating environments, high-availability components, billing, and network management.

Electronic Mail

Electronic mail is one of the most important services an ISP provides to subscribers. Establishing a mail service requires making choices on mail server software, storage of user mail files, assigning mail accounts, and selecting a mail reader for customers:

- *Server Software*

  There are many mail server software solutions available today. The choice of server software depends on the size of the ISP and the requirements for scaling to a greater number of users. ISPs must choose server software carefully because of increasing demands to support a greater number of different mail access protocols beyond the Internet standard SMTP, POP, and IMAP protocols. Anti-spam features are a must.

- *Storage Space*

  A rule of thumb is that users will consume, on average, up to 1 MB of disk space each for queued mail messages, depending somewhat on the mechanisms used to store the messages. Some mail systems allow flexible storage limits, and some support fixed limits. The ability of a mail system to bill for actual disk space used is one way to manage users with large mail files. In estimating disk space for mail users, additional space must be estimated for incoming and outgoing system-level queues.

- *Assigning Accounts*

  Assigning mail accounts can require a large amount of administration time, so whichever registration mechanisms the ISP uses, the automatic creation of mail accounts and repositories is a necessity. Mail accounts should be automatically deleted when user accounts are terminated. In general, there should be mechanisms to ensure consistency between the ISP’s billing information and the existence of user accounts. Lightweight Directory Access Protocol (LDAP) is becoming an industry-standard means of maintaining user information across multiple services and operating environments.

- *Mail Readers*

  Providing a bundled Web browser and mail reader to clients with pre-configured addresses is one way to reduce support costs. Some ISPs configure products like Netscape Communicator for these purposes. Bundled packages simplify the process of bringing new customers on-line, and the reduced administration overhead results in a net savings to the ISP. With many customers wishing to access their e-mail using Web browser interfaces, mail server software that provides a Web-based subscriber interface is a plus.

Mail Server Software

The most critical decision in establishing mail service is the choice of mail server software. There are many mail server software packages on the market today, however only a few are in wide use by ISPs. This section discusses four of the most commonly-used packages.

Different mail server packages allow for varying numbers of users and scalability. Standard Solaris utilities support local ISPs up to approximately 10,000 users; Software.com’s *Post.Office* is appropriate for regional ISPs in the

50,000 user range; Sun™ *Internet Mail Server*™ and Netscape's *Messaging Server Hosting Edition* can each host up more than a million users; finally, Software.com's *InterMail* is specifically targeted for National ISPs (Table 2).

Mail Software	Local	Regional	National
Standard Solaris Utilities	✔		
Post.Office		✔	
Sun Internet Mail Server		✔	✔
Netscape Messaging Server Hosting Edition		✔	✔
InterMail			✔

Table 2 Mail software for various sizes of ISPs

Standard Utilities

Small ISPs are often configured with standard Solaris sendmail for outgoing mail, and public-domain Post Office Protocol (POP) and Internet Message Access Protocol (IMAP) servers for providing incoming mail to subscribers.

Sendmail accepts outgoing mail from the customer's mail client software using Simple Mail Transfer Protocol (SMTP). Sendmail also transfers and queues mail messages to SMTP agents at other sites.

POP and IMAP servers are used to provide access for mail clients to read, save, and delete messages from their mailboxes. IMAP enables clients to compose, delete, and send mail while disconnected from the server — synchronizing changes once the connection is re-established. POP and IMAP protocols are used for mail retrieval only.

Standard sendmail does not scale well beyond approximately 10,000 subscribers for two reasons:

1. Sendmail's local delivery agent stores all incoming mail messages in a single directory (*/var/mail*), making performance and access using normal UNIX® tools quite difficult.

2. Both sendmail and standard POP/IMAP servers require subscribers to be registered through the standard UNIX mechanisms, either */etc/passwd*, NIS, or NIS+. These mechanisms all use sequential searches — which do not scale well — to locate records in their respective databases.

Various ISPs have customized sendmail and POP servers to meet their specific needs. One ISP has modified sendmail's delivery agent to deposit mail in a directory structure where the full path name of a mailbox can be generated by the subscriber's account. Another approach is to store mailboxes in large indexed files. Using off-the-shelf databases is another solution for overcoming the limitations of */etc/passwd* and NIS.

Post.Office

Software.com's Post.Office is a drop-in replacement for standard sendmail and POP, and scales from 30-50,000 subscribers per mail host machine. It supports mail accounts for users without login ids, eliminating a major restriction of the standard utilities. The most important advantage of Post.Office is that it provides increased functionality and simplified administration.

Account management is handled by a mail or Web-based forms interface. Post.Office does not run with root permissions, which makes it significantly more secure than standard sendmail. It supports size limits on individual mail messages, mailboxes, and for the entire mail system. Its ability to handle mail for multiple domains makes it easy to support commercial and small business customers with their own domain names. Performance is enhanced over standard sendmail by the fact that Post.Office's servers are multi-threaded, allowing multiple simultaneous connections. One limitation of Post.Office is that it is designed to run on a limited number of configurations, making it more difficult to deploy across a national ISP architecture; both SIMS and InterMail are designed to overcome limitations like these.

Sun Internet Mail Server

Sun Internet Mail Server (SIMS) is used by ISPs of all sizes, and is particularly appropriate for regional and national ISPs because it can scale easily to large numbers of subscribers. In one test performed by Shiloh Consulting (1998), SIMS supported 120,000 concurrent active POP and IMAP users on a single multiprocessor Sun Enterprise 6000 server — which projects to more than a million subscribers given a standard estimate of 10 percent subscriber utilization. SIMS provides the features that large ISPs need, including support for multiple mail domains and mail quotas.

The multi-threaded nature of SIMS provides excellent vertical scalability; SIMS also supports multiple servers with the messaging proxy option, enabling horizontal scalability as an option. Multiple hosts can be configured to create high-availability mail delivery systems, and the individual server processes that make up SIMS can be established on separate physical servers — flexibility that allows an ISP's architectural choices to be made independent of the mail system's architecture. This flexibility is critical in designing architectures based on the principles of functional decomposition and security discussed in the following chapter.

Management in SIMS is accomplished using a Java™ technology Web browser interface. SIMS utilizes LDAP for user management, enabling the seamless integration with other LDAP-enabled ISP components, eliminating the performance-limitations of linear searches through UNIX directory name spaces, and reducing overall administrative costs. Per-user hardware costs are minimized through the highly-efficient SIMS Message Store that minimizes disk space and hardware resources necessary to support large subscriber bases.

The SIMS Message Store reflects a design choice that was made specifically for the high-performance support of ISP installations. Storing messages as individual files results in a high degree of file system overhead, limiting performance. On the other end of the spectrum, storing messages in a full-scale database adds latencies and requires large amounts of computing power for message storage and retrieval. The SIMS two-level Message Store strikes an optimal performance level between these two extremes. It is a custom-designed repository with a simple structure optimized for the storage of mail messages; use of disk space is minimized by storing messages addressed to multiple recipients only once. The SIMS Message Store supports mail message access through POP, IMAP, as well as Microsoft MAPI, IBM Profs, and Solaris Mailtool client access through */var/mail.*

Netscape Messaging Server Hosting Edition

Netscape Messaging Server Hosting Edition is a high-performance messaging solution that delivers scalability, performance, and ease of administration. Netscape Messaging Server can support 500,000 active users on mid-range UltraSPARC servers, with advanced feature sets for prevention of spam and relaying. Concurrent access to a centrally managed, scalable universal message store — accessible via IMAP, HTTP, or POP — makes it possible for ISPs to offer messaging applications such as location-independent access via HTTP and unified messaging.

Integration with Netscape Directory Server, an LDAP-based server, allows service providers to centrally manage and store user and account information. User management is simplified by storing all information in the Netscape Directory Server, thus allowing administrators one-stop access to user, group, services, and shared data management — all through a browser interface. The modular architecture of Netscape hosting allows each server to be independent of the Directory Server, thus enabling multiple servers to share a single directory. A shared-directory architecture reduces both the number of directories and the amount of replication among directories. Support for SNMP agents enables integration of SNMP-based systems and network management products.

To simplify user management of configurations that require multiple messaging servers, Netscape has developed Mail Multiplexor (MMP). With MMP, each user is assigned to a specific messaging server but requests mail from the MMP server. MMP transparently directs incoming requests based on user entries in the LDAP directory. All users connect to a single server domain (for example, mail.isp.com), making it easy to pre-configure mail clients.

InterMail

InterMail is Software.com’s high-end mail server software, and is able to support thousands of concurrent connections to more than one million mailboxes. It has management features not found in Post.Office, and is designed to be deployed across a number of servers, a necessity in national ISP architectures. InterMail provides a highly-scalable solution through three basic features:

- It is multi-threaded and can exploit multi-processing servers.
- It can be configured to coordinate mail delivery with multiple peers, allowing additional servers to share heavy workloads.
- InterMail integrates with relational databases as the back-end storage retrieval mechanism for messages. InterMail supports standard SMTP, POP, and IMAP protocols for client mail access and delivery.

Netnews

Hosting a netnews service is a high-bandwidth proposition. Newsgroup traffic is increasing rapidly, with news volumes increasing significantly every couple of months. Each ISP has one or more news feeds from which it obtains new articles and newsgroups. New articles are distributed using a flood model, where the ISP receiving new articles is responsible for weeding out duplicates. Many ISPs also host local news groups containing information and discussions of local interest.

As of December 1998, a full news feed exchanged 500,000 messages per day, using more than 18 GB of disk space to store each day's news. Sufficient Internet bandwidth must be allocated to handle this data rate, and ISPs must also consider how much bandwidth their commercial customers will be requiring if they establish private news feeds. The amount of disk space required by netnews is enormous, however it can be adjusted to accommodate the disk space available by carefully tuning expiration times for messages.

Almost all ISPs of reasonable size use netnews server software based on Inter-Network News (INN). An ISP can choose to use the public-domain version, often modifying it to suit the particular news server architecture. Another option is to use Sun™ Internet News Server™, which is a supported, robust, and easy-to-configure part of Solaris ISP Server™ that was originally developed using INN as a base.

INN's main daemon process handles all of the news feed connections and configuration commands. A separate news reader daemon is spawned for each client connection request, and transfers news articles to the client using Netnews Transfer Protocol (NNTP). Some modifications to public-domain INN are necessary in order to handle the demands of large ISPs. These modifications have to do with multi-threading the daemon to improve news throughput handling, and multi-threading to handle client processes more efficiently than through heavyweight UNIX forks. These enhancements — critical for ISP performance and scalability — are standard features of Sun Internet News Server.

World Wide Web Services

The reason that has compelled the most residential and commercial subscribers to sign up with an ISP is the World Wide Web. The Web is a proven information-gathering resource, with the volume of Web-based electronic commerce rising rapidly. Informational Web pages range from movies and personal health to company product information. Since its reliability for informational use has been proven, its use in electronic commerce is the current forefront, with Web sites offering consumer items ranging from books and CDs to commercial products that enable even the purchase of Sun workstations. Transactional activities are expected to increase, ranging from mom-and-pop grocery stores speeding their soft drink orders via thin-client network computers to micro-transactions where fractions of a cent are charged for database searches — with the Internet Service Provider serving as the clearing house.

Forward-looking ISPs are preparing for these changes by establishing high-performance Web services to their clients using architectures that will scale as Web use expands and changes. Providing a Web service requires sufficient Internet bandwidth for low-latency browsing, and it requires hosting local Web pages for residential and commercial subscribers. A rough rule of thumb is that approximately 10 percent of residential subscribers will choose to host their own Web pages, and when they do they require from 2-3 MB of disk space per hosted site. These numbers can vary dramatically — Sun's affiliates have observed that there is a significant difference between urban and rural points-of-presence. Urban subscribers tend to have greater Internet bandwidth requirements, and rural subscribers tend to access the Internet less, and more frequently access Web sites of local interest hosted by their ISP.

The major ISPs predict that future revenue growth will come from content services. The Internet is becoming a source for pay-per-use video games, personalized on-line news services, on-line shopping services, and other sources of entertainment. As content providers themselves, cable ISPs have already recognized this potential, and they tend to build ISP configurations geared to providing interesting and entertaining Web sites. It is possible with sophisticated configuration techniques to open different avenues of Web page access to each subscriber depending on the level of service purchased.

Web page design and custom Web page hosting services are important revenue-generating areas for ISPs. Some ISPs offer varying levels of Web server performance appropriate for businesses wishing a high-performance Web presence on the ISP's side of their costly leased-line connections. The cost model for content provision is likely to change to the point where ISPs provide network access at a net loss, with their major revenue accruing from the content that they provide.

Web Server Software

Popular packages among ISPs for supporting Web hosting services and electronic commerce transactions include:

- *Netscape Web Servers.* Netscape has been at the forefront of Web server design, and has led the evolution from first-generation tools. Early Web servers forked separate processes for each CGI script to be executed, but today, multi-threaded processes allow high-performance multi-processor systems to host Web sites capable of responding to millions of hits per day.
- *Sun™ WebServer™.* Sun's Web server implementation delivers the best SPECWeb96 performance numbers on record as of January 1999, and incorporates a very strong caching solution. Sun WebServer supports Java servlets, which are one of the most popular mechanisms for integrating with back-end server functionality.
- *Apache.* Another popular Web server is the public-domain Apache Web server. Apache is available in free and commercially-supported versions, and is currently ported to most UNIX-based platforms as well as Microsoft Windows NT.

Other Considerations

Web server software is just one part of an effective Web service. Other considerations that an ISP must make include:

- *Proxy Servers.* Proxy servers act as an intermediary between the ISP's subscribers and the Internet. HTTP requests issued from client Web browsers are directed to the proxy server which, in turn fetches the requested data from the Internet site. By acting as an intermediary, proxy servers limit the impact of hijacking attacks on TCP sessions.
- *Content Filtering.* Although used primarily for security, proxy servers also provide the capability to filter content, allowing an ISP to provide subscriber-specific packages that limit access to sites having content inappropriate for children.
- *Caching.* The larger the ISP, the more likely it is that Internet bandwidth will be used to request the same Web pages on behalf of many different users. A caching server makes more efficient and cost-effective use of Internet resources while hiding network latency from subscribers. Caching has been shown to save ISPs up to 40 percent on network infrastructure costs while at the same time providing superior quality of service to

subscribers. Clearly all Web pages cannot be cached, with dynamic documents generated by CGI scripts being a common example. When used in conjunction with caching servers, transparent proxies like the Altheon Switch can allow ISPs to reap the benefits of caching without requiring subscribers configure proxy servers in their browsers.

These three services can be deployed on one or many physical servers, with the architecture dependent on the ISP's size, growth curve, and target markets. The Inktomi Proxy Server and the public-domain SQUID packages provide mechanisms for implementing proxies, content filters, and caches.

Domain Name Service

Domain name service (DNS) maps domain names to host addresses for business and residential customers as well as for clients on the Internet. In order for a client to access any Internet service, an address for its server must be obtained through DNS. One of the most important functions of DNS in an ISP network is to provide a limited view of the ISP network so that clients are only allowed to access a small set of servers. In fact, many large ISPs have a client DNS that reveals only a small number of addresses, while an internal DNS provides mappings that make all of the ISP's systems accessible to their staff.

Separate domain names are typically defined on a *per-service* basis. For example, a news service might be accessed through the domain name *news.isp.net*. A Web service might be accessed through *www.isp.net*. These domain names provide a layer of abstraction that an ISP can use for ease of configuration, load balancing, and fail-over. For example, a small ISP might have the names *news.isp.net* and *www.isp.net* map to the same physical server. As the ISP grows, these two services may be hosted on separate servers and with only a change in DNS, customers will correctly access the new configuration.

Domain Name Services	Local	Regional	National
Standard Solaris Utilities	✔		
Berkeley BIND		✔	✔

Table 3 DNS software for various sizes of ISPs

Local ISPs typically use the standard domain name service supplied with the Solaris operating environment (Table 3). Regional and National ISPs often begin with the public-domain Berkeley Internet Name Daemon (BIND) and modify it to suit their particular needs. Some of the variations in using DNS, and some common modifications include:

- *Static Load Balancing* can be achieved with both versions of DNS. When multiple servers are used to host a single service, DNS can be configured to provide different addresses for the same domain name on a round-robin basis. The effect of round-robin DNS is to statically balance the workload across multiple servers.
- *Dynamic Load Balancing* can be achieved by modifying BIND to monitor the load on various servers and map addresses on the basis of measured server loads. Some ISPs have created dynamic load-balancing schemes using the public-domain *lbnamed* package, while an increasing number of ISPs use load-balancing routers. Load-balancing routers are simple to configure, and avoid many of the drawbacks of DNS modifications. These approaches are discussed in detail in chapter 6.
- *Fail-Over* can be supported by modifying BIND to interact with fail-over software on the hosts that provide a particular service. If, for example, mail is hosted on two servers, and one fails, DNS would stop providing address mappings for the failed server and route all traffic to the second server — which could be either a hot spare or part of a load-balanced set of servers.

  One problem with DNS approaches to fail-over is that many Web browsers cache DNS responses as long as they are running. In the event of a server failure, new addresses are not obtained and a failed-over Web site continues to appear dead to the user. Load-balancing routers and IP address fail-over are two solutions which avoid this issue. The latter approach is used in Sun Clusters.
- *Client-Sensitive Mappings* can be established, for example, to provide access to different services for different subscribers depending on the package of services they have purchased. This can be done based on the IP address of the entity requesting a name mapping, and it requires modifying BIND.

One of the primary shortcomings of the current versions of DNS is the lack of intuitive user interfaces to simplify configuration — making DNS administration a tricky and error-prone activity. One solution to this problem is to integrate the name mapping tables with a DBMS. A graphical user interface can be developed for the DBMS to allow entries to be added, deleted, and modified; and back-end software could periodically generate new configurations for the DNS service and re-start the daemons. This solution greatly simplifies administration overhead, and is especially useful for ISPs with customer domain names that are frequently added and modified.

Security with Routers and Firewalls

It is essential for ISPs to establish mechanisms and policies that minimize the occurrences and the effects of intrusion into the ISP network. ISPs must control access to each server in their network, and packet-filtering routers and firewalls are two of the most important tools at their disposal:

- *Packet filtering routers* are the first line of defense, as they allow packets to be routed based on source and destination IP addresses, and source and destination TCP or UDP port numbers. This is the basic mechanism by which an ISP can ensure, for example, that only HTTP requests can be made of a Web server. Packet filtering routers are necessary, but not sufficient, for establishing a secure ISP network. A primary shortcoming is that routers do not provide a logging facility that can be used to detect and track intrusion attempts. Also, rule sets can be quite complex and prone to error. Finally, since routers are also stateless, they cannot perform complex analysis of transactions with internal hosts.
- *Firewalls*, when deployed in conjunction with packet filtering routers, can be used to establish multi-layer, secure gateways throughout the ISP network. With the router configured to block all connections except those appropriate to specific servers, the firewall can be used to perform more fine-grained filtering of traffic. Firewalls can perform stateful inspection of packets using knowledge of the specific application protocols being used. As a result, firewalls can allow easily-spoofed protocols — for example FTP and most UDP-based protocols — to pass safely through the firewall while dropping suspicious packets which are not received in the correct context.

  Firewalls can be used to perform detailed logging of traffic to internal hosts, which is important for detecting any intrusion attempts. Encryption-enabled firewalls can be used to set up virtual private networks (VPN) which can be used in electronic commerce applications and for interconnecting remote corporate customers. Finally, Network Address Translation (NAT) features can be used to hide internal network addresses, enhancing security and enabling an ISP to allocate more private IP addresses than they actually have available. Although firewalls can host proxy services, the most secure networks are protected by servers hosting only the firewall software.

There are three firewall options that are commonly used in ISP configurations:

- *CheckPoint FireWall-1* can be hosted on any Sun server, and it provides a good balance between cost and security. FireWall-1 addresses all of the security concerns discussed above, and also contains proxies which can be used both for ISP firewalls and for protecting corporate networks.
- *SunScreen SPF-200* is a high-performance software solution that provides a significantly higher level of security by using a dedicated, hardened version of the Solaris operating environment. The most significant advantage to SunScreen SPF-200 is that it establishes a transparent network device with no IP address. Since it is invisible to intruders, it is highly impervious to intrusion. SunScreen SPF-200 can be installed on a wide range of Sun servers, allowing ISPs to configure firewalls to meet desired performance levels.
- *SunScreen EFS* has several benefits that are particularly important to ISPs. It utilizes the same enhanced packet filtering engine as SunScreen SPF-200 software. SunScreen EFS It is multi-threaded, which allows it to fully-utilize multiprocessor servers from Sun. It can be administered from a secure Web browser interface. Finally, SunScreen EFS is a low-cost approach that can be used to add an additional layer of security to each server within the ISP's firewall-protected sub-networks.

Routers and firewalls are used to provide multi-layer protection between the various sub-networks that make up an ISP installation. They are integral to the architecture of an ISP, and will be discussed in the context of ISP architectures in chapter 4, in the context of ISP security in chapter 5, and in the context of each of the designs that are discussed in chapters 8 and 9.

ISP Infrastructure

There are many choices that an ISP must make when designing the infrastructure to support the services provided to customers. Choice in operating environments, high-availability strategies, billing, and network management have significant impact on an ISP's ability to deliver services and scale with a rapidly-growing subscriber base.

Operating Environment

Choice of operating environment is one of the most important decisions facing Internet Service Providers. Depending on the choice, ISPs can constrain their architectures to a limited number of possibilities, or maintain the greatest amount of flexibility for future growth. Sun's Solaris operating environment is designed from the ground up to provide reliable, scalable, multi-user, multi-platform network computing — one of the reasons why Dataquest has found that 10 of the top 12 worldwide Internet Service Providers use Sun servers.

- Scalability is key for successful ISPs, whose subscriber bases can grow as by orders of magnitude per year. With Solaris, ISP services can be scaled to utilize as many as 64 processors in a single server, providing multi-threaded services like Web servers the utmost in scalability. Using Sun Clusters, ISPs can deploy services in a high-availability environment with as many as eight servers sharing a single IP address and sharing the workload equally. As the following chapter will discuss, both horizontal and vertical scalability is important in ISP architectures.
- Multi-platform environments are the norm in ISP organizations today, Solaris delivers more network computing power on Intel platforms than Microsoft Windows NT. When Intel unveiled its Pentium II Xeon processor — designed specifically for high-performance workstations and multiprocessor servers — Sun announced Web server performance on an Intel-based platform running Solaris that exceeded by 146 percent the performance of an Intel-based Unisys platform running Microsoft Windows NT.
- Reliable systems are a requirement for ISPs, and Solaris delivers maturity that can only be achieved by years of use by many different users in many different environments — backed up by the hard work of eliminating bugs over time. Hardware failures are inevitable facts of life, and the dynamic configuration features of the Solaris operating environment support hot-pluggable components, management of individual processor sets, and separate operating environment domains on supported Sun and Intel platforms.
- Support for standard Internet protocols is the basis for an ISP's operations. Sun invested in Internet standards like TCP/IP beginning more than fifteen years ago, and the result is a protocol stack whose reliability, functionality, and interoperability has been refined for many years. In fact, the 1998

release of Solaris 7 includes a multi-threaded TCP/IP stack that enables Solaris-based servers to leverage symmetric multiprocessing right down to the networking level.

- Multi-user systems result in more cost-effective computing because the resources of networked systems can be shared remotely. Authorized users can log in and access services as needed. Client systems can access a wide range of facilities on servers — including file, print, and name services, Web pages, and database services. Administrators can manage networked systems remotely, resulting in significantly lower cost of ownership. Solaris servers are easily configured to provide multiple services at the same time — such as mail, file, and Web services — to many simultaneous users. Features such as these are not available on systems where multi-user support is an add-on.
- Ease of administration is required for ISPs whose architectures are increasingly being characterized by distributed points-of-presence; with network-based administration provided by the Solaris operating environment, ISPs can retain central control of resources and reduce the expense of administrators having to perform on-site maintenance tasks.
- Security in ISP networks is a fundamental requirement, and Solaris provides facilities that enable servers to be locked down and made impervious to attack. Between the inherent security provided by the operating environment and the additional security that can be provided with the SunScreen product line, Solaris-based firewalls provide the best protection available.

Solaris ISP Server

Sun’s Solaris ISP Server™ leverages the features of the Solaris operating environment with customized extensions designed specifically for ISPs. It includes a range of essential ISP platform services that simplify administration and provide facilities to deliver, monitor, and manage customer service level guarantees. Solaris ISP Server includes a service bus into which core service-delivery products can be plugged. The newest release of Solaris ISP Server includes Sun™ WebServer™, Sun™ Internet News Server™, and Sun™ Internet FTP Server™. In addition to providing services which are key to managing the ISP infrastructure and core ISP services, Solaris ISP Server also hardens the base Solaris installation to increase its level of security The product is organized into two collections of software: ISP Solaris Platform Extensions, and ISP Services.

- *ISP Platform Extensions*

  The platform extensions are specifically designed to deliver facilities needed by ISPs, including repeatable installation for multiple hosts, configuration management, server process monitoring, and log file management. The platform extensions also provide a hardened environment with intrusion detection, administrator authentication, access control, central auditing, privacy, and integrity for all network traffic. Specific features include:

  - Sun™ Internet Administrator™, a unified system management console which simplifies management and administration tasks by providing a single location from which all ISP services can be managed securely.
  - Sun™ Internet Services Monitor™, a core services performance monitor that enables automatic monitoring of a customer's service level by collecting historical information and providing notice of changes in service levels.
  - Sun Directory Services, a multi-protocol directory server. It currently supports Lightweight Directory Access Protocol (LDAP), NIS, and RADIUS, providing a shared repository for administrator and user authentication and privilege information, as well as the user-specific configuration for services such as e-mail server location.
  - Host Configuration, a robust, Web-based installation tool that enables the rapid installation and replication of desired Solaris ISP server components. The Host Configuration component also administers and configures services on local and remote servers.
  - SunScreen™ SKIP, which enables secure transmission between ISP servers using encrypted virtual private networks.

- *ISP Services*

  The ISP services are layered into the Solaris operating environment to deliver integrated, scalable, high-performance core services.

  - Sun Internet FTP Server, a scalable, high-performance FTP server that is fully compliant with Internet Engineering Task Force (IETF) standards.
  - Sun Internet News Server, an efficient, manageable, and cost-effective NNTP server that scales to more than 1,000 simultaneous clients.
  - Sun WebServer, a high-performance, highly reliable, secure, standards-based HTTP engine.

Additional third-party packages can be integrated into the administration environment, and can provide authentication using information maintained by the Sun Directory Services server. Optional services from Sun include Sun Internet Mail Server, and Sun Internet Calendar Server.

An additional tool useful for managing ISP networks is Sun's Bandwidth Manager. This facility enables ISPs to allocate specific network bandwidth levels that are to be maintained for specific services. It can be used, for example, to maintain a minimum T1 level of service for one customer's Web site while allowing another site a maximum 56 Kbps of bandwidth.

High Availability

High Availability (HA) is becoming more of a necessity in ISP networks, especially in telco and cable ISPs where subscribers have the expectation that Internet services will be as available as normal telephone and cable services. For ISPs, high availability is usually achieved by implementing fail-over mechanisms that allow backup servers to take on the load for a failed server, with more sophisticated solutions allowing multiple servers in a cluster configuration to share the workload equally. HA solutions typically have a small time lag from the time a failure is recognized to the point at which the service is restored on other servers.

- Sun Clusters are designed for environments requiring high-availability data, file, and application services. When services are managed in the Sun Cluster environment, they are provided with rapid detection and recovery from any hardware, network, operating system, or application software failure.
- Sun Clusters support pre-packaged software that utilizes its high-availability features, most importantly NFS, database, and Internet services including mail, news, and Web servers. Sun Clusters also provide an HA toolkit with which ISPs can create their own HA solutions.

Solutions for high availability in ISP environments are discussed in detail in chapter 6.

Rugged Telco-Grade Servers

Sun's Netra™ t servers leverage the latest UltraSPARC technology from Sun into a Bellcore NEBS Level 3 certified package for mission-critical applications. It provides a compact, rack-mountable carrier-grade solution for AC power environments with alarm capability conforming to central office alarm and system management requirements. The Netra t 1125 is packaged with single or dual 300 MHz UltraSPARC-II processors, with the ability to incorporate standard PCI peripherals on an internal 66 MHz PCI bus. Combined with the power of the Solaris operating environment, the Netra t enables telco and cable ISPs to more easily extend Internet services to their customers without compromising their high environmental standards.

Network Management

Network management functions are typically handled using Sun's Solstice™ Enterprise™ Manager, Sun™ Enterprise SyMON™, or HP's OpenView, with components from other vendors — such as router vendors — providing specific management modules for their products. Network management can sometimes be a challenge in telco environments, where both IP and OSI protocols are prevalent. For these environments, Solstice Enterprise Manager supports the specific needs of telco ISPs.

Administration

There are a variety of other administration issues with fairly straightforward solutions. Backups and customer care can be managed with third-party software packages. The most important part of these solutions — the process — is one that should be designed with the help of Sun Professional Services or one of Sun's affiliates.

Authentication and Billing

Authentication and billing are critical functions for ISPs. Authentication services ensure that only valid subscribers can gain access to the network, that subscribers' access to the network is controlled, and that reliable details on the customer's use is provided to the billing system. Billing systems vary widely in complexity from those which simply print invoices, to those which integrate with credit card companies and with legacy telephone company billing systems. Since there is such a wide range of options in this area, Sun refers these issues to a set of affiliates who are well-qualified to handle them.

Summary

Many software components must be assembled to provide a full range of ISP services, and the deployment of each service can vary substantially depending on the subscriber base, workload, and expectations of growth. Solutions — such as standard UNIX utilities — that are appropriate for small Internet Service Providers may not scale acceptably into mid- and large-sized configurations. Given that scalability is key, it is often best to make up-front investments in tools and infrastructure that will carry an ISP well into the future. Just as important as the services themselves are the architectures onto which they are deployed — this is the topic of the following chapter.

Architectural Principles 4

The architectures utilized by Internet Service Providers generally consist of a set of sub-networks having different functional responsibilities — for example user access, services, and administration. These sub-nets are connected via routers and/or firewalls that are used to control access from one sub-net to another. There are many ways in which the services discussed in the previous chapter can be deployed onto a physical network of machines. The choice of which architectures to use depends on a large number of factors including the number of subscribers, expected workloads, services to be provided, desired performance level, expectations for growth, and security concerns.

The factors influencing ISP architectures vary so much from one installation to another that no two architectures are the same. Architectures that are appropriate at one point in time are often superseded by those which exploit new developments in hardware and software. However, there are some underlying principles used in the design of ISP architectures, and this chapter discusses how these principles can guide their development.

Scalability

Scalability is the most important issue facing new and existing Internet Service Providers. Given the dramatic growth in Internet use, it is not uncommon for an ISP's subscriber base to grow by an order of magnitude in only one year. Beginning with a base of 10,000 subscribers, the requirement for an ISP to quickly adapt to handling 100,000 subscribers will put architectural decisions made early-on to the test.

Growing from 10,000 to 100,000 subscribers quickly takes an ISP out of the range where Solaris utilities such as NIS and standard sendmail can provide adequate solutions. When growth is expected, architectures must be designed to handle it, which often means making a heavy initial investment in a scalable architecture, and the hardware and software infrastructure to support it.

Horizontal and Vertical Scalability

Sun believes that ISP architectures should be designed to grow with both horizontal and vertical scalability.

Vertical Scalability

Vertical scalability is the ability to increase the processing power of a single server. This can be done by adding processors and memory to an existing server, or upgrading from one platform to a more powerful one. Sun has an extremely scalable product line from the uni-processor Ultra 5S server to the 64-processor Sun Enterprise 10000 (Starfire™) server. Many ISP applications — netnews for example — are not inherently multi-threaded and do not scale well with additional CPUs, so another dimension of scalability must be considered.

Horizontal Scalability

Horizontal scalability provides the ability to add more servers to a particular service area, such as mail, Web, or news services. Since scaling an ISP installation inevitably means adding more machines, an architecture that provides for horizontal growth is superior to one that depends on vertical scalability alone. In fact, horizontal scaling is typically the first approach to growing an ISP installation.

It is common for ISPs to deploy a set of services across a set of Sun Enterprise 5S servers. Their high-performance and low per-unit cost makes them an appealing starting point. When the ability to scale processing power is important, the Sun Ultra 2 server is a platform that enables a doubling of performance through vertical scaling.

Functional Decomposition

The services for an entire medium-sized ISP can be based on a single multi-processor Sun Enterprise server. Indeed, the standard approach for hosting enterprise-wide databases is to deploy one or two fully-configured Sun Enterprise 10000 servers. This ‘big box’ approach is not the most effective in the Internet realm, however. Decomposing an ISP’s services onto a set of smaller servers has several advantages:

- *Scalability*

  The use of multiple, UltraSPARC™ processor-powered servers from Sun Microsystems supports both vertical and horizontal scaling. ISP architectures that support growth begin with software architectures that allow a single service to be spread across more than one server. It is more difficult at first, for example, to configure a mail service to allow all subscribers to access their mailboxes from any mail server. Once this work is done, however, the architecture will sustain growth well into the future. For those ISP applications which do not scale well on multi-processor systems, horizontal scaling will be necessary at some point in the growth cycle, and it is more cost-effective to address these issues early on. In contrast, the “big box” approach limits scaling to the performance that can be achieved from a single large server.

- *Performance*

  Different services require different server configurations to provide optimum performance. By hosting each service on separate servers, CPU, memory, I/O bandwidth, and operating system parameters may be tuned for the task at hand. It is sometimes even appropriate to split different aspects of a single service onto separate servers — for example Web servers providing dynamic content can respond more quickly to user requests when they farm out specific requests to servers specially tuned for that purpose.

- *Reliability*

  With multiple small servers, reliability is enhanced because the number of single points-of-failure are reduced. If a news server fails for some reason, chances are that the loss will not affect the mail service. With one service down, the ISP can still run in a degraded mode. With more sophisticated architectures, each service can have multiple servers, increasing reliability even further.

- *Security*

  As much as packet-filtering routers and firewalls protect ISP networks from intrusion, it is always possible for some previously unknown attack to succeed in compromising security. An ISP configuration having separate servers for each service can limit the effect of such intrusions in two ways. First, if security is compromised on one server, the intruder may not necessarily gain access to any other server. Second, if firewalls limit network traffic to only the protocol for which the server is responsible, penetration attempts are only possible using that protocol. For example, guessing passwords through a telnet connection is not possible if that protocol is restricted from news, Web, and mail servers.

- *Flexibility*

  Finally, an ISP where services are spread across multiple servers is inherently more flexible. With a domain name service configured to dynamically allocate users to different servers, additional platforms may be configured, tested, and then turned on for subscribers by simply changing DNS configuration. This is difficult to accomplish with an ISP having only a small number of large servers.

Large servers certainly have their place in an ISP so long as they fit into a plan that supports both horizontal and vertical scaling. Larger servers are needed as vertical scaling is exploited to gain higher performance. For example, a Web service that uses multi-threaded Web server software and which has been vertically scaled may be hosted on a large Sun Enterprise 10000 server.

It is important to maintain the capability to scale in both dimensions, and the key to having this capability is in decomposing ISP services into separate components that can operate — and scale — independently. There are several ways to implement functional decomposition — by services, task layer, and by special function. All of these techniques can be used separately or together.

Partitioning by Service

An obvious way to functionally decompose an ISP's services is to partition each service onto a separate machine (Figure 2). This example illustrates independent mail, news, and Web servers. Each server can be configured and tuned for the service it is to host, yielding the best possible performance for each service. Because each service is independent, management is easier and the ISP can flexibly reconfigure each server — for example increasing main memory — while affecting only one service at a time.

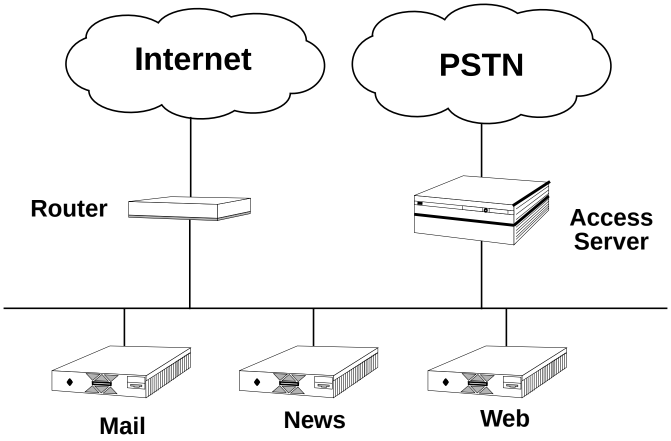

Figure 2 Partitioning by service

Partitioning by Task Layer

A less obvious way to decompose ISP services is by task layer (Figure 3). Each of the three basic services — mail, news and Web — have functions that can be layered into a three-tier architecture.

An access server provides the interfaces with which clients interact. For mail, this would be a POP/IMAP server that enables users to read their mail. For news, this is where an NNTP server resides. For Web services, the HTTP daemon would reside on the access server.

A content server holds the data accessed by each service, and may host a database server to store user mailboxes and to provide dynamic Web content, and an NFS server for news articles and static Web pages.

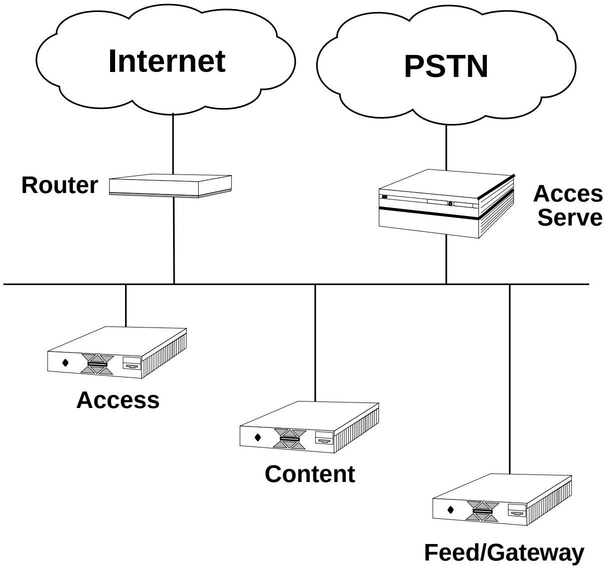

Figure 3 Partitioning by task layer

The feed / gateway server would provide the interface between the ISP and the Internet. It would accept incoming mail via SMTP and deliver mail to user mailboxes on the storage server. It would handle a news feed from an upstream site and place articles into a spool area. For Web services, a proxy server might be used instead of a feed / gateway server.

This example is fairly academic; a more realistic implementation (Figure 4) would partition by service and by task layer. This architecture has several advantages. Performance tuning can be done by the service and by the function within the service that requires optimization. Security is improved because penetration of one host does not necessarily yield access to the entire ISP network. This architecture is both vertically and horizontally scalable, and offers the ISP increased flexibility in configuring and managing the network.

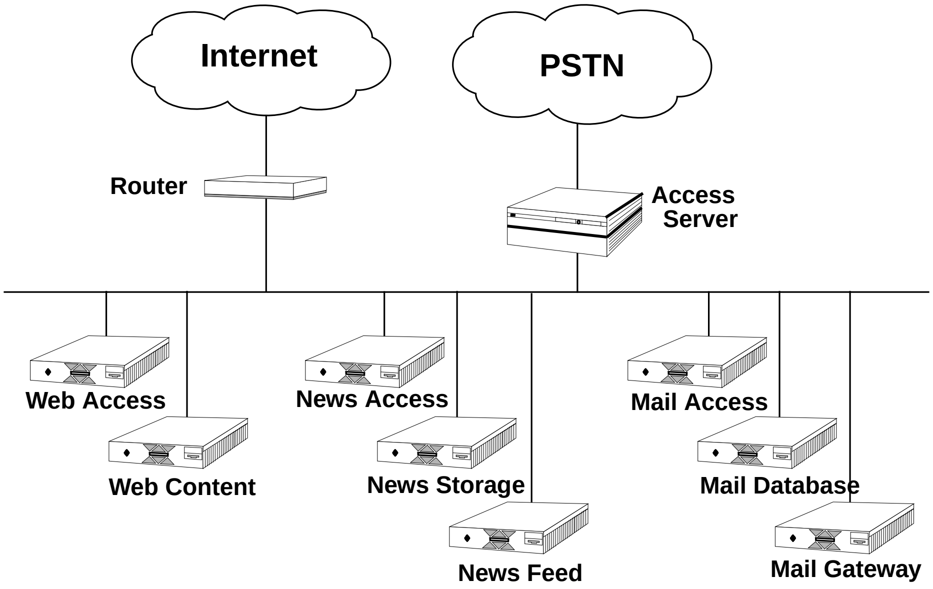

Figure 4 Partitioning by service and task layer

Partitioning by Special Function

The third way to partition ISP services is by special function. For example, splitting a Web service between the HTTP server and dynamic content delivery allows each function to be allocated the computing resources they need as they grow and change over time. Hosting DNS, network management, authentication, and billing functions on separate servers are also examples of partitioning by special function.

Reliability and Availability

Reliability and availability can be improved by adding redundancy into the ISP architecture (Figure 5). Multiple Internet connections can be established to different external networks (Sprint and MCI, for example) to ensure that, if one link fails, the other link can carry traffic to and from the Internet (with reduced performance). Multiple connections may also be partitioned by function. For example, a news feed could be supported on a dedicated connection so that the fluctuations in news traffic do not affect performance of other services. Of course the ISP network requires multiple routers in order to accrue the reliability benefits of multiple Internet connections.

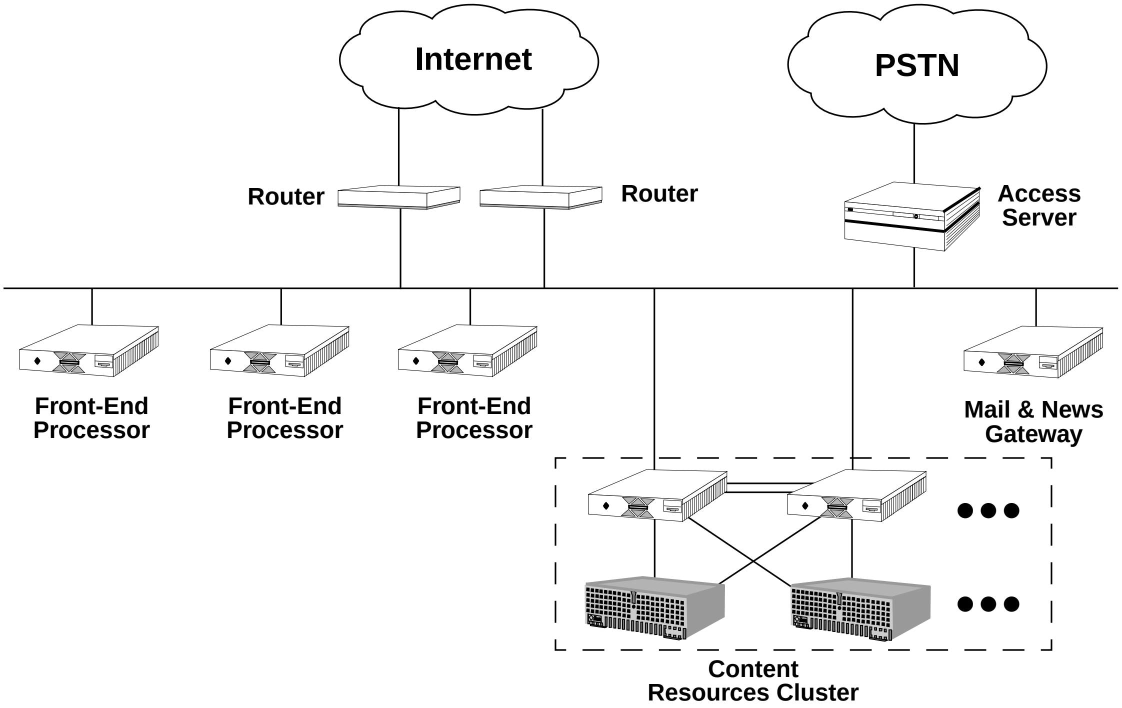

Figure 5 Architectures for reliability and availability

Multiple front-end processors provide access to each ISP service, for example mail, news, and Web services. If three systems were deployed as in the example, there are two ways to configure them, each having a unique set of trade-offs:

- One service could be allocated to each front-end processor, which would not significantly improve availability over the example shown in Figure 4. The failure of one front-end processor would cause the service that it hosts to be unavailable.

- Every service could be provided on each front-end processor, which would make performance and security somewhat more difficult to manage. In this case, however, round-robin DNS could be used to provide static load-balancing between front-end processors.

Note that the weak link in this architecture is the fact that only one mail and news gateway system is configured. In the event of a failure, incoming mail and news would not be handled — however, existing data in all services could still be accessed by clients.

Content storage can be made more reliable by employing Sun Clusters hosting high-availability data services. In this example, two Ultra 2 servers are clustered using high-availability NFS software to provide high-availability network file services. Each server is configured with two interfaces, one to each Sun™ StorEdge™ A1000 disk array. Using RAID level 5 or mirroring, each array is impervious to single disk failures. In the case of catastrophic disk or controller failures, mirrored content can be quickly accessible to each server.

Sun has found that the most effective solution for reliability and availability is to configure clusters to handle each service. These configurations will be discussed in forthcoming chapters.

Architectures for Security

Architectural support for security is usually provided by creating multiple sub-networks that are isolated by firewalls. A typical ISP configuration (Figure 6) consists of an access network, demilitarized zone (DMZ), and services network. One of the advantages of partitioning services by task layer is that firewalls can be interposed between layers — in this example a firewall controls access between the front-end processors and the content storage.

The access network provides the connection by which services may be used by both Internet and dial-up users. The packet-filtering router allows Internet traffic to proceed only to the access servers and the DMZ.

The DMZ creates a moat between the Internet and the services network. This example shows a mail and news gateway in the DMZ. All SMTP and NNTP traffic from the Internet is examined by the firewall and routed only to the gateway server. Once the mail and/or news is deposited on the gateway server, it proceeds to transfer it to the storage server. The DMZ is appropriately named because the indirect transfer of data from the Internet to the storage network requires any intruder to “drop their weapons” in the DMZ — making direct penetration into the services or content storage network more difficult. A fully-configured DMZ might contain separate news feed servers, mail gateways, proxy caching servers, DNS, and authentication servers.

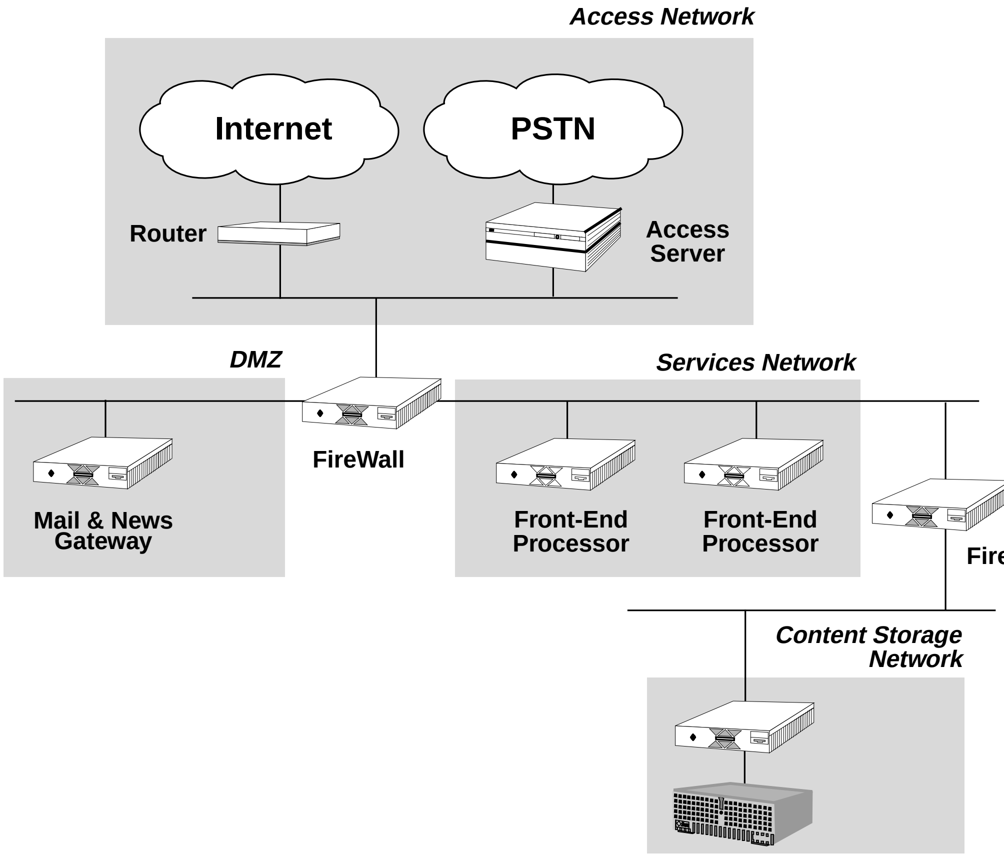

Figure 6 Architectures for security

The services network provides access to mail, news, and Web services. The combination of router and firewall could be configured to allow only subscribers connected via the access server to access the mail and news servers. Internet access is usually provided to Web servers, and in this example the firewall could allow only HTTP service to a specific front-end processor. This limitation of protocols to specific hosts makes intrusion via multiple protocols difficult to achieve.

Because the content storage network must satisfy requests from both the DMZ and the services network, a separate firewall is configured so that performance is not limited. This firewall, interposed between the access and the storage aspects of each service, limits traffic to NFS requests from the authorized front-end processors.

Security issues will be covered in detail in chapter 5, and example of a local ISP configuration utilizing a DMZ will be discussed in chapter 8.

Summary

Every ISP architecture is unique, and is determined by multiple variables, including the number of subscribers, the services to be offered, the workload expected, and the degree of scalability, availability, and security that is needed. Scalability is the most important aspect of an ISP configuration, and it can be achieved by architectures that support both vertical and horizontal scaling. The benefits of functional decomposition — by service, task layer, and special function — accrue by deploying a number of smaller servers rather than a single large server. The examples of decomposition, enhancements for reliability and availability, and security concerns are fairly academic. In actual use, many of these approaches are superposed, realizing ISP architectures that utilize the best of all of these techniques.

Ensuring ISP Security 5

Security is an increasingly important concern for Internet Service Providers. News media reports indicate that Internet security breaches are on the rise, and ISPs are right in the middle of the action — with potential intrusions launched from customer connections, from the Internet, and sometimes even from internal sources.

Paradoxical Security Requirements

The security requirements of ISPs are somewhat paradoxical. An ISP's internal network — containing billing, network management, security logging, and customer service functions — must be locked down as tightly as any major corporation's networks. Indeed, the "corporate jewels" that any ISP needs to protect the most are the data that controls customer access and bills for services. In contrast, the front-end of the ISP network — providing network access and services to customers — must be relatively open. Many ISPs allow unrestricted traffic between dial-up customers and the Internet, which gives them freedom to utilize whatever protocols they wish in accessing services over the network. ISPs also must allow access to customer Web pages by both dial-up users and those accessing services from the Internet. These paradoxical security concerns — open at the front, closed at the back — must be addressed in the ISP architecture through a set of carefully-constructed and access-controlled sub-networks.

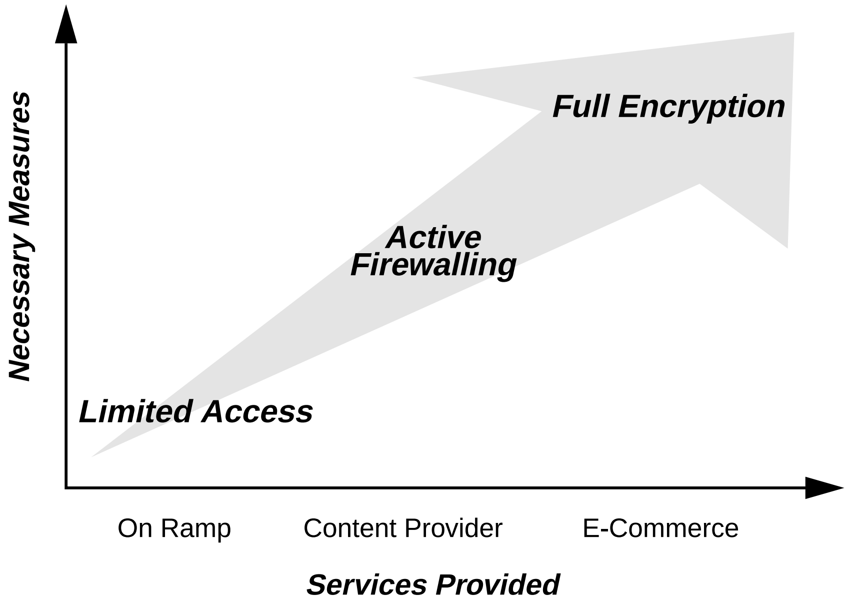

Figure 7 Changing security needs over time

Complicating the open/closed nature of ISP networks is the fact that the security requirements of ISPs change over time. The industry as a whole has been undergoing a transition in the services offered. Initially, ISPs acted as simple “on ramps,” offering little or no value-added services to customers. Most ISPs today are providing content in some capacity — offering various levels of service from hosting customer Web pages to providing multimedia interactive games and entertainment-based sites. Increasing numbers of ISPs are acting as clearing houses for electronic commerce transactions. This evolution demands increasing security measures (Figure 7):

- On-ramp ISPs must ensure the security of user authentication and billing information, and can do so with relatively simple security measures.
- Content provider ISPs must establish multiple sub-networks having carefully-controlled access from one security layer to the next. These ISPs make extensive use of firewall technology to implement their security policies.

- ISPs that support electronic commerce must provide the highest level of security for financial transactions — requiring extensive use of encryption from Secure Socket Layer (SSL) connections between the ISP and the user, to Virtual Private Networks (VPNs) that establish encrypted IP tunnels between layers in the ISP network.

Policy and Implementation

The first step in determining what security measures are necessary in an ISP network is development of a security policy — an unambiguous statement of what the ISP is attempting to accomplish. This statement should declare a general policy, as well as provide precise statements regarding the points between which network traffic is allowed to flow. This may seem elementary, but without taking this step, it is impossible to determine which security measures to deploy. As well, it is impossible to assess whether the various components of the security policy (routers and firewalls, for example) are properly configured to implement the policy.

The two most common security policies are to *allow access to all services unless expressly denied*, or to *deny access to all services unless expressly permitted*. The first policy, allowing all services unless expressly denied, allows the widest range of services with the least amount of intervention — and security mechanisms. This policy may be appropriate for ISPs that provide connectivity to the Internet with no value-added services. The problem with this policy is that, as new protocols and services become available, ISPs are vulnerable until explicit action is taken to prevent them from being exploited. This is why most ISPs — especially those providing more than just Internet connectivity — choose to implement a security policy that involves denying all services unless they are specifically allowed. As new services and protocols become available, the ISP will need to assess whether to allow them and what measures need to be taken to prevent unauthorized use or intrusions.

A good security policy should take into account cost and convenience trade-offs. For example, an ISP needs to consider whether customers are to be protected from intrusion originating from the Internet, or whether this is more appropriately a value-added service provided at extra cost — particularly to corporate customers, who require this level of security. Similarly, ISPs need to consider the cost / benefit of electronic mail security. Password-protection of customer mailboxes is clearly a benefit worth the cost, however is the cost of providing completely encrypted SMTP, POP, and IMAP mail connections and mail storage worth the benefit to the few customers who need this level of security?

Implementation

Once a good security policy is written, network designers can begin implementing it. Sun believes in establishing security measures at two distinct levels:

1. *Establish access control to each network element.* This is accomplished by controlling traffic to each of the ISP's sub-networks with packet-filtering routers and firewalls. This step ensures that, for example, only HTTP requests reach a Web server, limiting the reliance on the security of each individual host.

2. *Secure each network element itself.* This can involve simple measures such as removing support for unapproved protocols — for example removing telnet from a Web server. The most sophisticated approaches involve restricting traffic to a small set of authenticated and encrypted connections between servers. This kind of "lock-down" requires state-of-the-art encryption technology discussed later in this chapter.

It is important to consider where to begin to implement a security policy. The logical sequence is for an ISP to first ensure the security of the internal network, then the services network, and finally the subscribers.

Security Audits

It is easy for designers to be so close to the problem of security that some aspect invariably gets overlooked. As part of an ISP's security policy, it is important to incorporate plans for ongoing, targeted audits of network security. This is an area where third-party perspectives can be extremely valuable, and specialized security consultants can be quite helpful. Sun Professional Services is equipped for both post-design consultation and ongoing security audits of ISP networks. Public-domain packages such as SATAN can be utilized by ISPs for their own security audits, as well.

A Range of Security Options

Chapter 4 introduced the concept of separate firewall-protected sub-networks as an architecture for security. This is known as *perimeter defense*. Perimeter defense effectively creates a series of concentric layers of security — like moats around a castle — making the services network *most* accessible and the internal network the *least* vulnerable to intrusion. These practices are almost always used in ISP installations, however they are not the only set of security measures possible. There is a wide range of security options for ISPs to consider, and perimeter defense sits in the middle of the range from the simplest to the most complex approaches.

In considering a security policy and the range of security measures that an ISP could deploy, it's important to keep in mind the four basic kinds of attacks that need to be guarded against:

- *Denial of Service Attacks* are becoming more common, and are aimed at making parts of the network unavailable by flooding it with superfluous requests, for example the ICMP requests that result from pinging hosts on the network. Other denial of service attacks include gaining root access to a server and making its services unavailable.
- *Confidentiality Attacks* compromise the privacy of user or ISP data by gaining unauthorized access to servers on the network. At the user level, these attacks could result in an intruder gaining valuable information from a subscriber's mail messages. At the core of the ISP internal network, confidentiality attacks could result in an intruder gaining access to subscribers' credit card numbers.
- *Integrity Attacks* occur when an intruder gains access to unauthorized data and then modifies the data. For example, a subscriber's Web page might be compromised and modified — a common ploy of hackers. More serious attacks on ISP data might invalidate credit card information or eliminate billing data, resulting in a loss of revenue.
- *Attacks on Authenticity* occur when an intruder replaces a standard part of the system with one that compromises security in some way. For example, a Trojan Horse login program might accept user passwords and pass them through a covert channel to an intruder who could use them to gain unauthorized access to the ISP network.

There are several basic areas that an ISP can exploit to reduce the likelihood that the above categories of attacks will be successful. The steps range from straightforward — but often overlooked — network design principles to the application of encryption technology that is beginning to play a more important role in ISP networks.

Network Design

Three fairly simple and very effective ways to increase security of ISP networks are based on the principle of isolating network traffic so that it can pass — and hence can be snooped — only between a limited number of hosts.

Router Access Control Lists

Even without active firewalls, routers should be configured with access control lists that allow only necessary protocols to be routed to particular hosts. For example, packet-filtering routers can be used to limit only HTTP traffic to a specific Web server, and to deny all other access from the Internet. This is an important level of protection that can be accomplished by ISPs, however as access control lists become more complex, router performance is significantly impacted. This is why most ISPs use combinations of routers and firewalls to control network traffic.

Switched Ethernet

Ethernet switches can be used to limit traffic on each ISP network segment to only the packets that are directed to, or have originated from, the host on that segment. Consider the benefits of having switched Ethernet in a services network that contains a Web server, a mail front-end server, and a news server. In the case of a network configured with shared Ethernet hubs, an intruder who gains access to the news server might put the host's network interface into promiscuous mode and snoop mail and Web server traffic — compromising the privacy of user mail messages, and ultimately obtaining passwords that can result in penetration of the mail and Web servers. Using switched Ethernet, mail and Web server traffic is not visible on the network segment containing the news server and therefore snooping cannot yield the same information to the intruder. In addition to its security benefits, switched Ethernet significantly increases network throughput by partitioning traffic between network segments.

Separate Administration Networks

When administration traffic is allowed to co-exist on the same physical network as user traffic, intrusions can be accomplished by snooping administration traffic — including passwords — and then masquerading as an administrator using the compromised information. When root logins and administrative access are allowed only over a separate network, the ability to obtain and use administrative privileges is significantly curtailed.

The use of a separate network for administration is analogous to telephone companies using out-of-band signalling to direct long-distance calls. When in-band signals were used by telephone companies, it was possible for "blue boxes" to reproduce the control signals, resulting in network intrusions and significant losses of revenue. Today's telephone network uses out-of-band signalling exclusively, directing long-distance calls over a separate, administrative channel. Making ISP administration commands out-of-band to normal users limits the pathways by which security can be compromised. This can be done in two ways:

- *Console Network*

  An effective but surprisingly low-technology approach is to utilize the serial port of each server in the ISP network. Each serial port is wired to a terminal concentrator (Figure 8) which is configured onto the ISP's internal network. From any host in the internal network, a telnet session can be established to the appropriate port on the terminal concentrator to obtain access to one of the server's serial ports. Sun has found this to be a useful approach in many installations where access to the consoles of a large number of servers must be provided — including Sun's own engineering design automation laboratories. Using this technique in an ISP network involves allowing root access — and hence administrative access — only to logins on the serial port, limiting a set of intrusion paths from the Internet.

- *Separate Ethernet Network*

  The main drawback of the serial port approach is that TCP/IP-based administration protocols cannot easily be utilized over the serial port network. For this reason, most ISPs with separate administration networks use a 100 Mbps Ethernet segment for ISP network traffic, and a separate, low-cost 10 Mbps Ethernet segment for administration traffic. This approach has to be more carefully implemented than the serial port approach — for example by restricting some services to respond only to requests from the administration network. The end result is significantly more flexible than the serial port approach. When implementing a separate Ethernet segment, it's important to limit the extent of the network so that it cannot, for example, be used by an intruder to gain access to all of the ISP sub-networks at once.

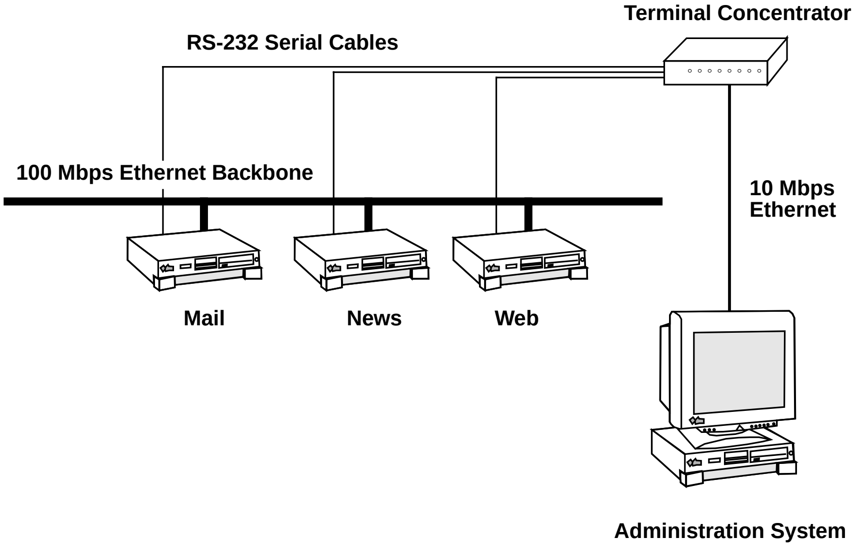

Figure 8 Console network using terminal concentrator

Limiting Access

Access control lists, switched Ethernet, and separate administration networks are ways to control network access to each host on the ISP network. There are three important ways to control the functions that each host on the network will provide:

- *Access Control*

  Each server should only allow connections appropriate for the service that it is to provide. For example, telnet and FTP access should be turned off on a Web server so that any weaknesses in these protocols cannot be exploited to gain access to the server.

- *Removing Non-Essential Applications*

  Further control can be achieved by actually removing the software used to implement non-essential services. It is one thing to rely on the portmapper to limit connections to unauthorized services; it is another to ensure that these services are not even loaded onto the server. This is a relatively simple security measure, but it is often overlooked.

- *Non-Standard Ports*

  Services used within the ISP network should use non-standard port numbers. In this approach, standard port numbers must be used for services — for example Web and mail — provided to subscribers and Internet clients. Non-standard port numbers can be used where traffic using standard protocols is generated within the ISP network. For example, where SMTP is used to transfer mail from a mail gateway server to a mail storage server, using a non-standard port number makes it more difficult for intruders to gain access to lower layers of the ISP network from a higher-level machine that has been compromised.

Sun’s SunScreen™ EFS™ firewall product can help with the process of controlling access to servers. Individual services can be configured or denied with ease, and stateful packet filtering can further limit access to the server by adding an additional level at which packets are inspected before being accepted by the host.

Active Firewalling

When most Internet Service Providers consider security issues, firewalls are usually one of the first measures that are considered. Firewalls are most effectively deployed once the first two measures — network design and limiting access — are implemented. Active firewalls usually include a combination of packet-filtering routers and firewall systems.

Packet Filtering Routers

Packet filtering routers are the first line of defense for an ISP, as they allow packets to be routed based on their source and destination IP addresses and TCP or UDP port numbers. This is the basic mechanism by which, for example, an ISP can ensure that only HTTP requests can be made of a Web server. Packet filtering routers are necessary, but not sufficient, measures for establishing a secure ISP network. A primary shortcoming is that routers are stateless, which means that they do not inspect the contents of each packet in a protocol stream to ensure that they are valid and consistent over time. Another shortcoming is that routers generally do not provide a logging facility that can aid in the detection and tracking of intrusion attempts. Finally, complex rule sets dramatically decrease the performance of routers and can severely limit the available bandwidth over costly leased-line connections. Sun has found that the most effective way to use packet-filtering routers is to use the simplest access control lists possible, and leave the more complex filtering activities to the firewall.

Firewalls

Firewalls can be used to establish highly-secure gateways in and out of the ISP network, and to control access from one level of the ISP installation to the next — from the DMZ to the services domain, for example. With routers taking the first line of defense by dropping packets that have invalid source addresses or whose destinations are not allowed by the ISP, firewalls can handle a more fine-grained filtering of traffic. In fact, since firewalls inspect application-specific data inside packets, they can be viewed as a form of access control lists for the application domain. Some of the services provided by firewalls that are key to ISPs include:

- *Stateful Inspection*

  Firewalls typically inspect the contents of packets and validate the connection type, address, protocol, and port numbers. Firewalls recommended by Sun perform stateful inspection of packets, which means that they can ensure that each packet is valid within the context of the protocol stream in which it is contained. Consider, for example, filtering of FTP requests. When an outgoing FTP request is made, a firewall notes that it is expecting a response from the remote FTP server. When the response is received, the firewall inspects it and then passes it on to the client. Incoming FTP packets that are not responses to outgoing requests are dropped, eliminating the possibility of FTP spoofing. Since firewalls typically contain knowledge of application protocols, they can also filter UDP requests based on the packet contents. For example, requests to the portmapper can be filtered by protocol, adding an additional layer of assurance that disabled applications cannot be accessed from improper sources. This stateful handling of protocols is simply not possible with packet filtering routers alone.

- *Logging and Detection*

  Just as important as preventing breaches in security is detecting, logging, and sounding alarms as intrusion attempts occur. This allows the source of intrusion attempts to be localized, for the intruder's methods to be monitored, and for the firewall to be hardened against future attempts using similar techniques. A solid architecture for ISPs to use is for the firewall logs to be stored on a server in the most highly-protected internal network. Firewalls configured with SunScreen SPF-200 software requires a dedicated administration server to which log data is transferred over a secure channel.

- *Address Translation*

  Some ISPs use firewalls to translate private, internal addresses to public, external addresses at the firewall. Private addressing helps to obscure the ISP's network configuration to any would-be intruders; it can be used to ease the shortage of IP addresses by mapping a large number of private addresses to a smaller set of external addresses; and it can be used in server load-balancing strategies that are discussed in chapter 6.

- *Encryption*

  All of Sun's firewall products support encrypted communication with remote firewalls, which is necessary for the creation of virtual private networks. The SunScreen SKIP package allows encrypted communication with remote clients as well. Encryption technology is quickly becoming a key differentiator between ISPs, and its implications are discussed more fully in the next section.

Firewall Security Options

The market offers a variety of security options, each of which has an important place in implementing ISP security policies. Sun's SunScreen family of products provide the foundation for securing ISP networks as well as positioning them to support value-added corporate virtual private networks and electronic commerce applications. There are many add-on firewall packages available from other vendors, CheckPoint FireWall-1 being one of the most popular in ISP networks.

- *SunScreen SPF-200* is a software product that transforms any Sun server into an integrated single-purpose hardware/software firewall, providing a higher level of security than add-on firewall software. The most significant advantage of servers configured with SunScreen SPF-200 software is that they become transparent network devices with no IP address. This stealth design makes them invisible to intruders and therefore impervious to intrusion. SunScreen SPF-200 runs on a hardened version of the Solaris operating environment, making it less susceptible to attacks that exploit known weaknesses in other UNIX-based systems. Because of the degree of security that it provides, SunScreen SPF-200 is most often deployed at the gateway between the ISP network and the Internet.

- *SunScreen EFS* brings stateful packet filtering and SKIP encryption to a wide variety of Solaris-based SPARC hardware platforms. Because its packet-filtering engine is multi-threaded, SunScreen EFS is a solution that can scale with the underlying hardware platform, increasing firewall throughput as ISP networks grow. SunScreen EFS is designed to support high-availability environments where multiple firewalls are configured so that one can take over the other systems's filtering functions in the event of a failure. SunScreen EFS is easily-configured with a convenient Web browser interface that is managed through a SKIP-enabled virtual private network.

  SunScreen EFS can be used to control access to ISP sub-networks and to individual servers. SunScreen EFS can be installed on any of the servers in the ISP network, giving an additional layer of packet inspection and security monitoring. For example, when deployed on each server within a DMZ, SunScreen EFS can help limit the successful penetration of one server from affecting the other severs within the sub-network.

- *CheckPoint FireWall-1* is the most popular firewall software in use by ISPs today. CheckPoint FireWall-1 provides an excellent balance between cost and security. FireWall-1 addresses most of the security concerns important to ISPs, and includes both packet filtering engines and proxies for a wide range of protocols. In contrast to FireWall-1, SunScreen EFS provides the additional advantages of Web-based administration, vertical scalability, and SKIP-based encryption. With these advantages, Sun expects SunScreen EFS to quickly become popular with Internet Service Providers.

The Role of Encryption

Once a secure set of sub-networks is established, Internet Service Providers can then begin to utilize encryption technology to move network security from perimeter-based defense to security that is distributed throughout the ISP network. The level of security afforded by encryption technology is important for both the ISP and its corporate and individual subscribers.

Corporate Virtual Private Networks

ISPs can deploy secure virtual private networks (VPN) as value-added services for corporate subscribers. As shown in a simplified network diagram (Figure 9), an ISP can use a firewall to encrypt all traffic between a local company office and a remote office connected via the Internet. This is sometimes known as an encrypted IP tunnel.

Traffic within each company office is un-encrypted, but once a packet addressed to a remote office passes through the firewall, it is encrypted and can cross through the Internet in complete privacy. The corporate customer can enjoy the same security benefits as an expensive leased-line connection between remote offices at a fraction of the cost by using a VPN over the Internet.

Sun's SunScreen SKIP is an emerging IETF and ANSI standard for key management and IP encryption. SKIP is the core encryption and key management technology used in the entire SunScreen product line, and it can authenticate and encrypt all IP-based protocol streams.

VPN software is available for client systems, which means that individual users can join their company's VPN and enjoy the same level of privacy that remote offices have. These clients can access the Internet through the same ISP providing the virtual private network, or via dial-up connections through other ISPs. The end result is that companies have increased flexibility in supporting the computing resources of telecommuters and travelling employees.

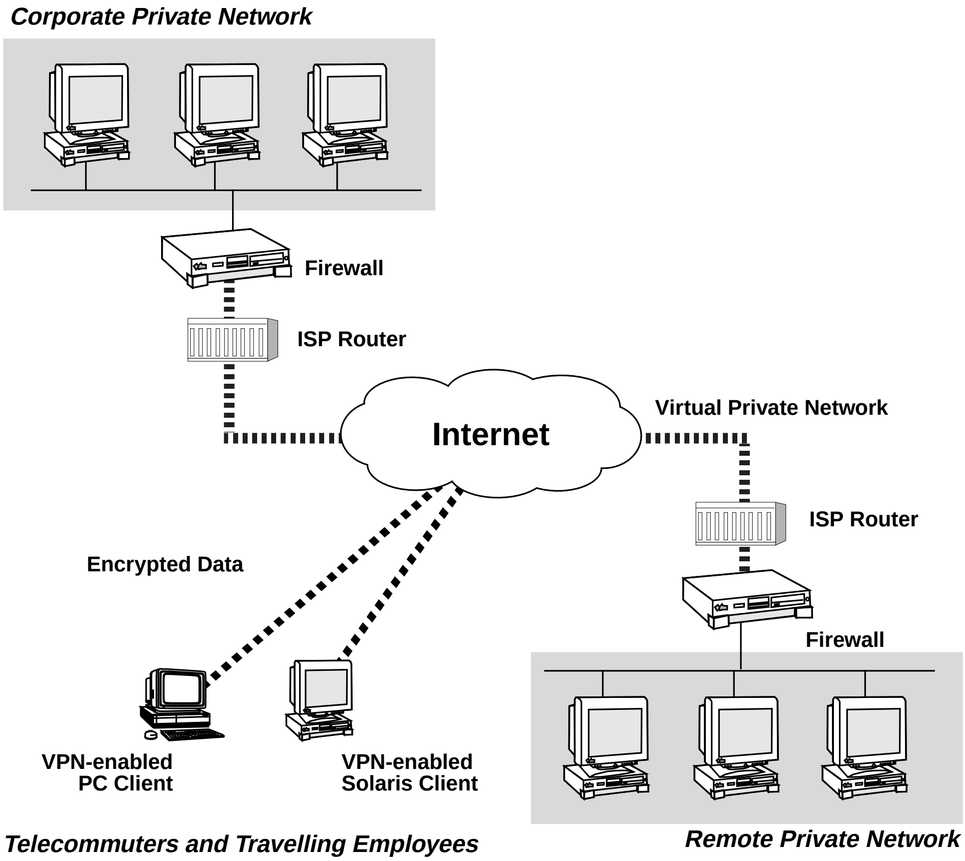

Figure 9 Corporate Virtual private networks using SunScreen SPF-200 software and SKIP-enabled clients

Virtual Private Networks Within ISPs

Virtual private networks have an important role within the ISP networks because they can be used to ensure authenticated, secure traffic between sub-networks and between hosts in the ISP network itself. Consider the benefits of secure communication between a Web server and systems on the ISP's internal network (Figure 10):

- Administration of the Web server could be accomplished over a Virtual Private Network, limiting the ability of an intruder to hijack the encrypted administration channel. Indeed, Sun uses SKIP-based VPNs to administer SunScreen EFS firewalls so that only authenticated, secure channels can be used to modify firewall parameters.
- Billing information such as credit card data could be passed into the secure internal network with little risk of an intruder using the channel to gain access to enterprise-critical data stored at this level. Indeed, many corporations use VPNs in their internal networks to ensure secure transmission of sensitive data between departments — for example corporate financial data or payroll records. The benefits for corporate intranets and for ISPs are the same — even internal users cannot view data that is meant to be private.
- Virtual private networks are an effective way to implement a security policy in which all access is denied unless specifically allowed. By "hard-wiring" the allowed connections between servers, a successful intrusion into one server is stopped short of affecting other servers in the same or more secure sub-networks, effecting a highly-secure "lock-down" of the ISP's servers. This can be used to secure all servers at a particular layer, for example the DMZ or the services network. It can also be used between layers to provide secure communication between the central ISP location and a remote point-of-presence.

Support for Electronic Commerce

An ISP network that is properly wired for secure virtual private networks is one that is ready to support electronic commerce transactions. Because any loss in security in an electronic commerce environment has direct and immediate financial implications, it's important to make the deepest data more secure than in any other application. Encryption between network elements — for example between a Web server taking credit card numbers and a credit card clearing server — becomes absolutely essential.

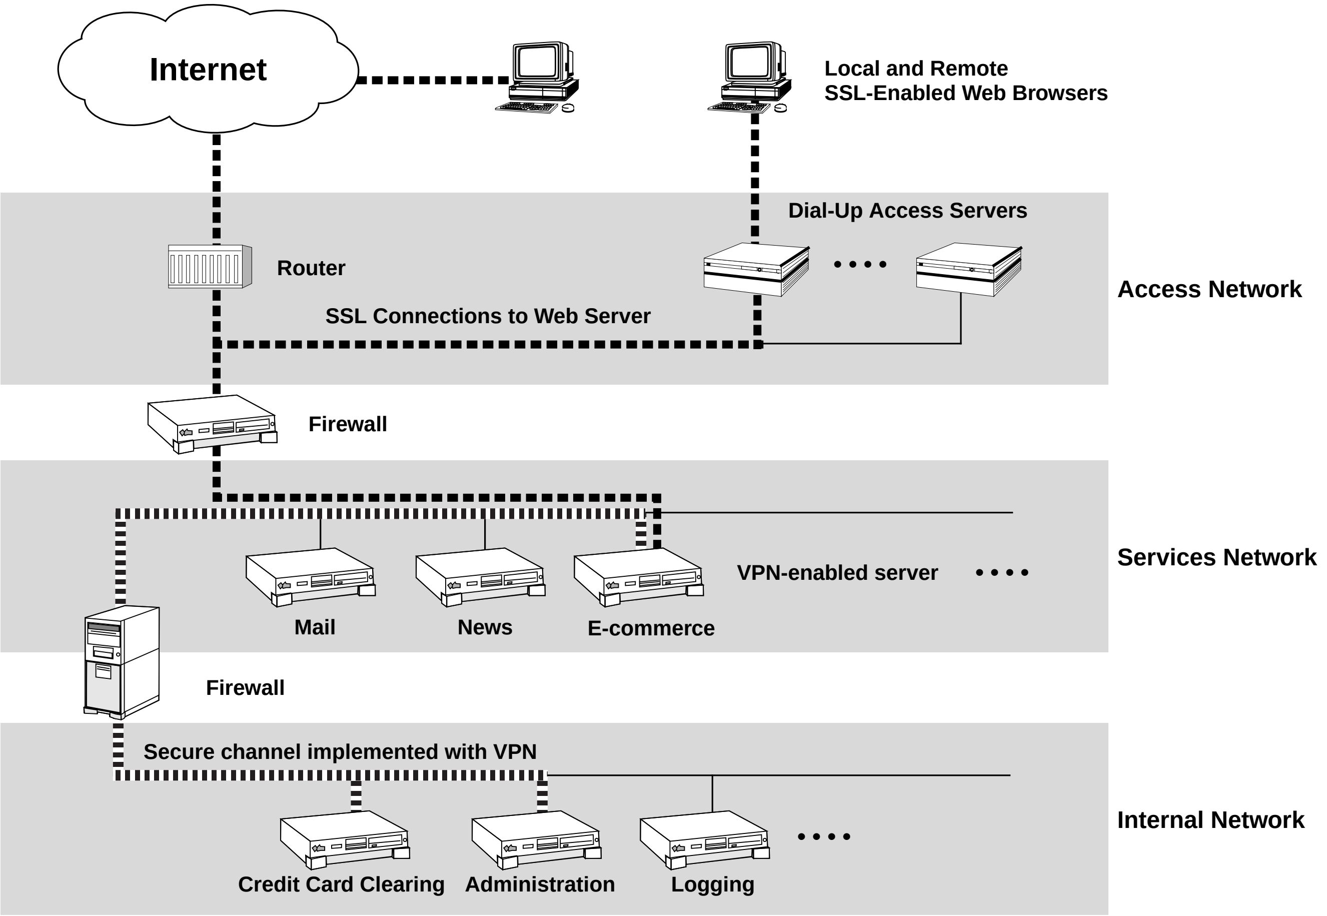

Figure 10 Encrypted channels can be used for increased security within the ISP network. Secure Socket Layer connections are often used for secure communication between Web browsers and Web servers.

There are a variety of electronic commerce servers available from different vendors, and they generally provide services such as catalog management, search engines, automatic generation of product pages from catalog databases, sales analysis, and automated shipping and sales tax calculation. With respect to security, commerce servers typically provide secure ordering and payment methods, secure payment processing methods, and additional restrictions on access to the commerce server itself.

There are two commonly-used strategies for passing encrypted information between Web browsers and Web servers:

- *Secure Socket Layer*

  Most browsers include integrated support for SSL, the secure socket layer. Because of this support, SSL is often used to communicate private information such as credit card numbers and purchase orders between Web browsers and Web server. (SSL connections are illustrated in Figure 10). As with SKIP, SSL verifies that a client is communicating with the intended server, and encrypts the data to prevent its use by an unauthorized party.

- *Secure-HTTP*

  Secure-HTTP is an application layer encryption mechanism, which means that it provides encryption one layer above the TCP/IP layer. Although it offers similar encryption mechanisms to SSL, it is not as flexible because it is specifically built only for use with HTTP.

There are other protocols specifically designed for passing payment information between browsers and banks. MasterCard and Visa's SET protocol is one example. Having support from both Netscape and Microsoft, an advantage of SET is that only the card holder and the bank are able to see the actual credit card numbers. The merchant never sees the plain-text credit card numbers, providing a high degree of security for these transactions.

Summary

In the past, Internet Service Providers have been on-ramps to the Internet. Today's ISPs provide value-added services such as Web site hosting and other content-related facilities. Advanced ISPs are supporting encryption-based services such as corporate virtual private networks, and they are gearing up to support the marketplaces of the future. As the sophistication of ISPs increases, so do their security needs. Once an ISP has established a written security policy, implementation should evolve along a path that begins with solid network design principles and techniques to deny unauthorized access, to active firewalling, and finally to the deployment of encryption technology for internal and external VPNs and electronic commerce transactions. Sun believes in securing both the networks and the individual servers within the ISP installations, and is ready with security products to meet these needs.

High-Availability Solutions For ISPs 6

Today’s Internet users have high expectations about the availability of the services they purchase. Service interruptions and lagging response times make national news, and there is intense competitive pressure amongst ISPs to provide high-availability services and to stay out of the limelight. Indeed, as ISPs make the change from being on-ramp services to being content providers, the availability of customer data — from Web sites to electronic mail — becomes even more important. Providing highly-available services is important to all ISPs, but it is especially so to telephone company ISPs, where customers expect Internet services to be as reliable as their dial tones.

Fault Tolerance vs. High Availability

The notion of high availability is often confused with fault tolerance:

- *Fault tolerant* computers are specifically designed to provide uninterrupted service even after catastrophic system or environmental failures that would completely shut down other configurations. Fault tolerant systems typically employ specialized hardware with processors running in lock-step execution. Applications requiring fault tolerance, such as telephone switching and air traffic control, cannot sustain *any* interruption in service.
- *High availability* (HA), is used when the uninterruptability of fault tolerance systems is not needed, but a much higher degree of service is required than is normally expected from a single system. The processors in HA systems

run asynchronously, are loosely-coupled, and execute separate copies of the operating system. HA environments provide full hardware and software redundancy, and recovery from failures takes only seconds or minutes.

While fault tolerant systems provide a very high level of hardware availability, they provide no coverage for application software failures. Because the servers in HA configurations run asynchronously, and execute their own copies of the operating system, a software problem in one host — such as an application failure — is unlikely to affect the other host in exactly the same way. Indeed, since the majority of system failures are due to software faults, servers which can quickly detect and recover from hardware *and* software failures are preferable to fault tolerant systems.

Sun's high-availability systems are built with low-cost, scalable, off-the-shelf servers which are often only one-third the cost of equivalent fault tolerant hardware. Sun Clusters, with their fast recovery from failure, can provide nearly the same levels of total availability as fault tolerant systems, but at a much lower cost — and without compromising the flexibility and scalability that is so important to Internet Service Providers (Figure 11).

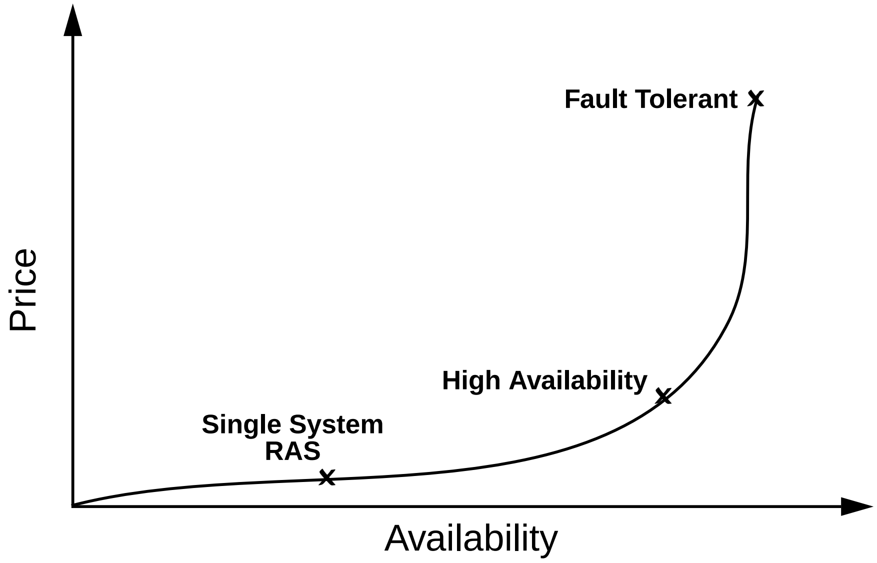

Figure 11 High-availability configurations provide the best price-performance trade-off for Internet Service Providers.

High Availability and Load Balancing

Among the options for providing high-availability services, the issues of load balancing and HA are tightly coupled. ISPs know that horizontal scaling is the preferred way to grow any one service, because the failure of one server does not necessarily mean the loss of the entire service. These ISPs know that the architecture best poised for growth is one that has services already split across a set of servers. ISPs with horizontally-scaled services must have mechanisms for balancing the workload between servers — and good load-balancing systems are one way to provide services with greater availability.

Experienced ISPs typically combine horizontal scaling and clever addressing schemes to provide both high performance and high availability. Consider a simple example of an ISP that wishes to provide both high performance and high-availability Web service. A simplified network diagram illustrates only an access network and a services network (Figure 12). The access network consists of a router with an Internet connection and a set of access servers. The services network contains a DNS server and three Web servers, each containing exact replicas of the static Web page data that they serve. When more than one server is deployed to handle a single service, the set of servers is known as a *service group*. Traffic between the access network and the services network is controlled with a firewall system. (True high-availability services require redundant Internet connections, routers, firewalls, DNS servers, and backbone network; however the purpose of this example is to focus on the Web service).

In order to provide the best performance possible, an ISP must balance incoming Web server requests across the service group. Techniques used to balance workloads can also be used to improve the availability of the service because a failed server is functionally the same as one that is overloaded — neither should allocated any more client requests.

There are several approaches to load balancing across a service group. The simplest is to use round-robin DNS to parcel out different addresses to different users. A somewhat more complex variation is to use a load-balancing DNS. The most popular approach today — which overcomes the shortcomings of the DNS approaches — is to use routing techniques that dynamically balance workloads using network address translation.

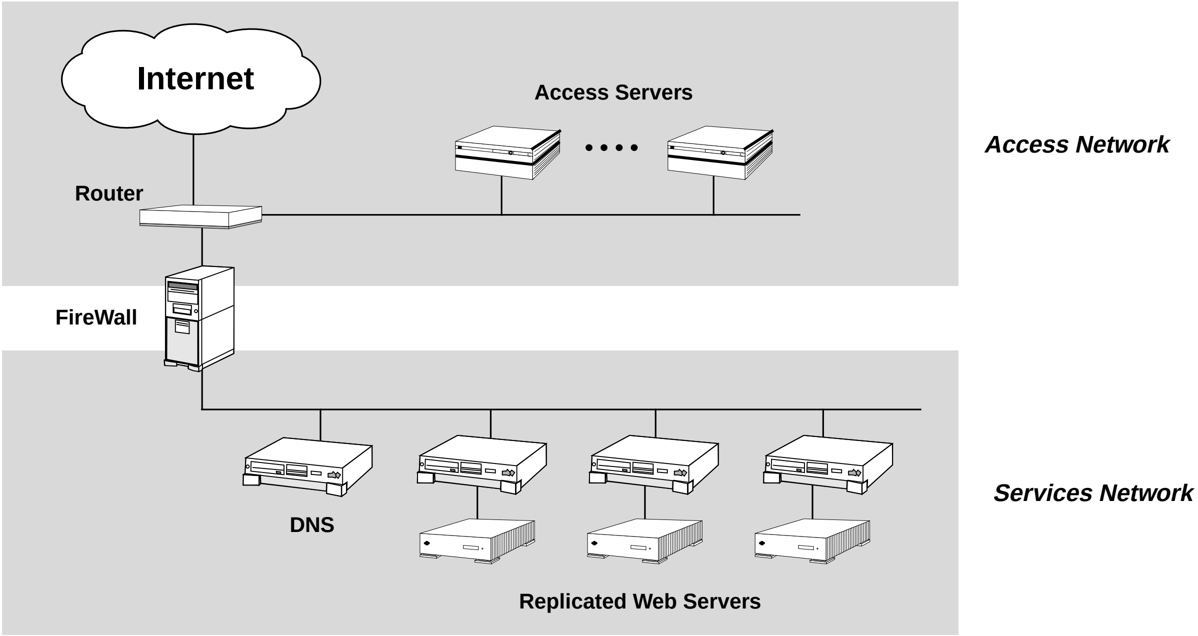

Figure 12 Example configuration for a high-availability Web service

Round-Robin DNS

Standard Solaris DNS and Berkeley BIND can be configured to resolve a single name used for the service group — for example *www.isp.net* — to one of the three Web server addresses. This is known as Round-Robin DNS, and it involves simply configuring IP addresses for the same domain name. DNS rotates through the three addresses each time a name resolution is performed, and the result is that each Web server is allocated every third client. The address records provided by DNS have their time-to-live (TTL) set to zero, so DNS clients will not cache the address records. This is static load balancing, which is independent of the actual workload on the Web servers, and is also independent of whether the servers are even available.

The usefulness of this scheme for high availability is limited, because it depends on an administrator manually modifying the DNS tables in the event of a server failure. During the time that the server is down and the DNS configuration is unchanged, one third of the clients will receive an address for the failed server, and will find the Web service to be unavailable. This window can be narrowed by modifying BIND to periodically check on server availability and to automatically modify its address resolution tables to account for failed servers. The usefulness of this scheme is still limited, as some Web browsers cache resolved addresses independently of the TTL in the DNS address records. Users experiencing a failure on their HTTP connection will not receive a new, valid, IP address unless they completely quit and restart their Web browser.

Lbnamed

Some ISPs have modified the public-domain *lbnamed* software and used it in place of BIND. Lbnamed stands for Load-Balancing Name Daemon, and it heuristically allocates addresses for the service group depending on each server's workload. The DNS server hosts two programs, *lbnamed* and a *poller*. The poller contacts a client daemon running on each server in the service group. Each time that it completes polling the client daemons, it creates a new configuration file and sends a signal to lbnamed that notifies it to read the new configuration. If the poller does not receive a response from any client, it simply removes its entry from the configuration given to lbnamed. The configuration given to lbnamed includes the one-minute load average on each server. Lbnamed uses this information to calculate a weight for each server, and it changes the server's weight each time a request is assigned to it.

This mechanism achieves a basic form of dynamic load balancing, and has the advantage that it can quickly and automatically remove a failed server from its configuration list. This approach shares the same disadvantage of round-robin DNS, namely that some users will experience a disruption in service when a server fails.

Address Translation

The most commonly-used approach in use today is to advertise only a single public address for the service group and use network address translation mechanisms to route network traffic. This approach ensures that a given IP address, once obtained by a client, will always point to an operational server, overcoming the limitations of the previous two schemes. Load-balancing address translation mechanisms can be implemented using load-balancing routers, or with software approaches bundled with cluster solutions:

- Load-balancing routers measure each server’s latency in handling requests and use response times to calculate a performance index which is used to decide how to route incoming requests. Routers vary in the sophistication of their load-balancing mechanisms. Some make their decisions based purely on routed traffic, while others also factor in data obtained by agents running on the servers themselves. Cisco’s LocalDirector can be configured to provide simple round-robin load balancing that distributes the load to the next available server regardless of load; it can allocate requests to the server servicing the least number of connections, and it can allocate requests to the server with the lowest response times.

  Load-balancing routers can be used to create high-availability service groups by not routing requests to servers which are down; some routers can be configured in HA pairs so that service continues uninterrupted through any single router failure. Other load-balancing routers are available from F5 Labs, HydraWeb Technologies, and RND Networks. The routers vary in load-balancing capabilities, administration capabilities, network interface speeds, and their ability to handle protocols beyond HTTP.
- Service groups can be hosted on cluster configurations that handle load balancing and routing decisions internally. Sun Clusters provide a general purpose cluster computing environment where existing applications can run unchanged; additionally, packages are available that provide HA-enabled Web, mail, and news servers. Sun Clusters appear as a single server to the outside world, while internal load-balancing mechanisms route requests to appropriate servers. Sun Clusters provide a unique solution that combines load-balancing and HA services; they support up to eight servers today in both symmetric and asymmetric configurations.

Symmetric and Asymmetric Configurations

Many approaches to providing high-availability services make the assumption that one server can potentially take on its workload plus the entire workload of a failed server. This results in a dramatic decrease in overall performance in the event of a failure. Fortunately, there are two ways to deploy Internet services in high-availability configurations:

- *Asymmetric Configurations* use one server as a hot standby which has access to, or contains a copy of, the data that is needed in order to take over the operation of a failed server. Asymmetric configurations provide consistent performance in the event of a server failure, as the workload of the failed server is assumed by a hot standby server with no other users to support. This approach, clearly, provides performance at the cost of having an idle server to handle fail-over.

- *Symmetric Configurations* use all servers to handle user requests which all continue to share the workload in the event of a server failure. In symmetric configurations, each server may, in fact, be master of a different set of data, and even support different services. Traditional HA solutions deploy only two servers, however modern cluster approaches to HA enable the workload to be allocated across *N* servers, with *N-1* servers available in the event of a single server failure. The larger the number of servers, the less the impact of a single server failure. This provides ISPs with the greatest leverage of their hardware investment, and enables them to bring servers down for maintenance without dramatic performance loss for their customers.

Most Internet Service Providers use symmetric configurations in their networks, as they believe that degraded performance in the event of a server failure is preferable to the additional cost of a hot standby server.

Replicating, Partitioning, and Sharing Data

The Web server example discussed so far hosts identical, replicated, Web pages. This configuration is artificially simplified — in actual use, careful thought must be given to the deployment of server data across high-availability servers. These data management issues are common regardless of whether simple or full-featured high-availability solutions are implemented. When multiple servers are used to provide a single service, their data must be either replicated, partitioned, or shared.

Replicated Data

In cases where the data used by the application does not change much over time, and does not require a large number of disk volumes, replication makes HA simpler to implement. DNS is a prime example of a service where a small amount of data can easily be replicated so that one server can take over without having dual-ported access to the underlying data. News servers can use replicated databases by having news feed servers provide incoming articles to each server in a service group simultaneously. If a server goes down, the standard mechanisms provided by INN will ensure that all articles will be propagated to it when it comes back on-line. If Web server data does not change frequently, replicating content is a reasonable approach.

Partitioned Data

When the volume of data is large, and the changes are frequent, partitioning is a preferred approach. User mailboxes are frequently partitioned by ISPs because it is far too costly to maintain replicated copies of user mailboxes on each server in a mail service group. Instead, user mailboxes are typically allocated to individual servers, each of which is required to handle storing and retrieving mail for a partition of the ISP's subscriber base. In mail servers like Sun's Solaris Internet Mail Server™ (SIMS), LDAP databases are used to manage the allocation of users to mail servers.

One approach to partitioning the mailboxes is to deploy server clusters. The remaining servers in a cluster take over in the event of a failure (Figure 13). For a simple, two-server cluster, one could deploy a pair of Sun™ StorEdge™ A1000 arrays. Each array is configured with two sets of disks — one set on each array for both mail partitions. Each server is master of its own set of disks, each of which is mirrored on the other array. In the event that one server fails, the other server can immediately take over the mail partition on the failed server's disk array. In the event that an array fails, each server can access all of its data on the remaining array. Managing data consistency between multiple mirrors is a complex problem, and is handled by mechanisms provided in Sun Clusters.

Shared Data

When there is a large amount of data to be provided by the service group, and file locking mechanisms are implemented to prevent inconsistent data from being used by members of a service group, shared data solutions can be deployed. The majority of Web traffic today is dynamic, requiring a back-end set of servers that provide dynamic content to all of front-end processors. Dynamic content is often provided by database management systems. In fact, a growing number of mail servers also utilize this architecture, enabling a very large number of user mailboxes to be stored on back-end servers, reducing the number of partitions required. Another approach popular among ISPs is to store user mailboxes as files on a back-end set of NFS servers.

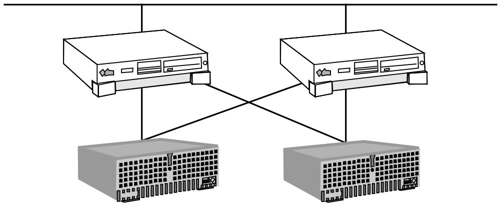

Sun StorEdge
A1000 arrays

Figure 13 Partitioned data can be deployed using clusters. Shown is a two-server cluster with dual, mirrored, Sun StorEdge A1000 arrays.

Three-Tier Approach to HA

With the majority of Web server traffic provided by dynamic back-end content servers, a three-tier approach is a natural choice for many ISPs today. This architecture is based on shared data on a back-end content servers, a set of front-end processors, and gateway servers in the DMZ (Figure 14).

Gateway Servers

The gateway servers are configured in the DMZ and route incoming data between the Internet and the content servers. Mail gateway servers store incoming mail with Secure Mail Access Protocol (SMAP) and store them briefly on disk before standard sendmail safely passes mail to the storage servers using SMTP. News gateway servers store incoming articles and likewise pass them on to the storage servers using NNTP.

Web proxy servers, if configured, can work in both directions. Access is allowed to proxy servers in the DMZ from Internet users and from ISP subscribers on the access network. The proxy servers, in turn, can access Web sites located within the ISP's services network and those on the Internet. All of the gateway server approaches serve to isolate vulnerable systems from attack by requiring all requests to go through one level of indirection before they are handled by a server in the services network.

Although gateway servers are often deployed in pairs, they do not necessarily require a high-availability configuration because the failure of a gateway server simply delays the arrival of mail or news rather than making a service completely unavailable to subscribers. The workload of incoming mail can be statically- or dynamically-balanced to the mail gateway servers with load-balancing routers or other techniques described in the previous section.

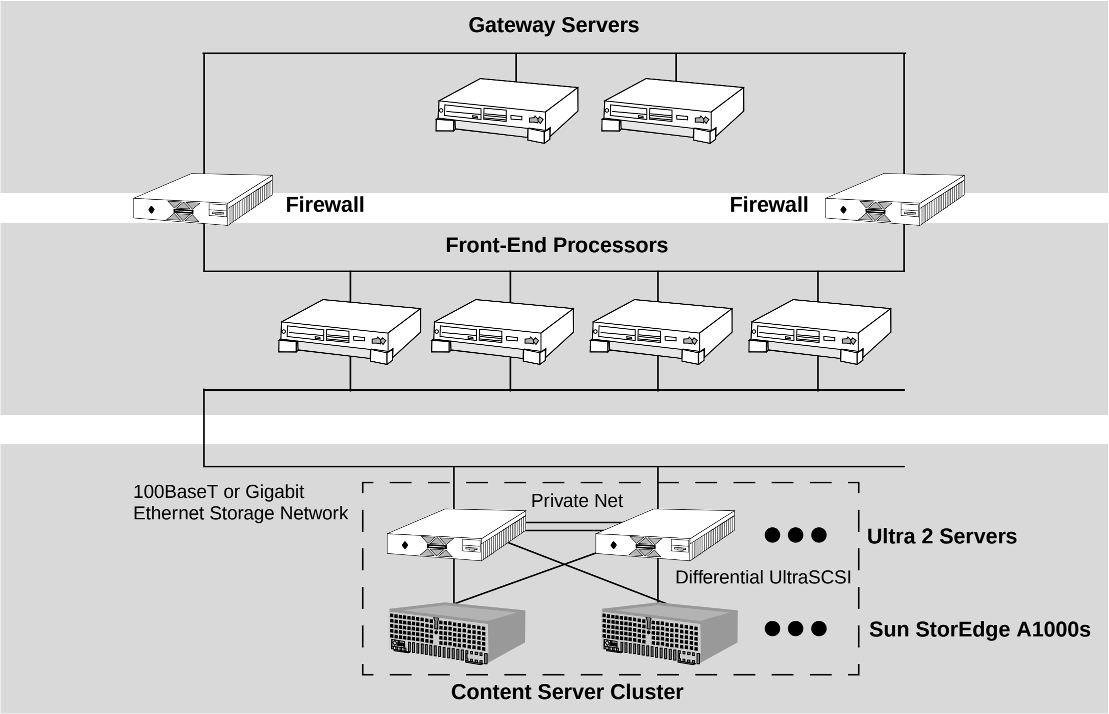

Figure 14 Three-tier solution with content provided using a content server cluster

For incoming news, separate feeds from different sources are often statically partitioned to specific news gateway servers. In fact, news servers often utilize separate leased-line connections so that bursts of incoming news do not interfere with the Internet bandwidth needed by other ISP services. The Solaris Bandwidth Manager can also be used to partition bandwidth so that news is allotted only its share of bandwidth.

Front-End Processors

A set of front-end processors (FEP) handles user requests for data. For a mail service, the FEPs provide mail to users with SMTP, POP3, and IMAP4 protocols, accessing user mail files from the back-end NFS or database servers. Similarly, Web and news front-end processors interact with the ISP subscribers and access dynamically-created content from back-end database servers.

The benefits of using front-end processors for interacting with users is that many inexpensive systems can be deployed, and the service can be expanded by horizontal scaling as the ISP requires. Large ISPs with many racks of front-end processors might consider using the multiple-domain features of the Sun Enterprise 10000 platform in order to reap the benefits of horizontal scaling in an easy-to-maintain reliable, and cost-effective package. Each of the eight domains in an Enterprise 10000 sever hosts a separate copy of the operating system, leveraging the advantages of functional decomposition with the convenience large servers.

Regardless of platform, a range of high-availability solutions can be deployed with front-end processors, ranging from the use of load-balancing routers to the use of an FEP clusters. Because all of the FEPs share the same back-end content, the complexity of failing-over partitioned data is tremendously reduced.

The front-end processors can be connected to the content network over the ISP's backbone network. The high-performance solution illustrated in Figure 14 utilizes a separate network in order to dedicate as much bandwidth as possible to the content servers. ISPs that are particularly security-conscious will use a pair of firewalls between the two network layers in order to ensure that only content requests come in to the content network, and that responses are provided only to the authorized front-end processors.

Content Servers

The content server is a Sun Cluster configured with HA database or NFS software. Clustered architectures are well suited for providing high-availability solutions in ISPs. Sun Clusters feature redundant paths between all systems, between all disk subsystems, and to all external networks (redundant external networks are not shown in Figure 14). No single point of failure — hardware, software, or network — can bring a cluster down. Fully integrated fault management software in the cluster detects failures and manages the recovery process without operator intervention, allowing failed components to be replaced on-line, without impacting availability. These configurations provide high levels of service without requiring costly, proprietary technology. The architectural approach taken in Sun Clusters includes:

- *Redundant Disk Systems*

  Sun Clusters utilize multiple redundant, multi-homed disk systems, which are often mirrored to allow uninterrupted operation in the event that one of them fails. With its ability to support Differential UltraSCSI or Fibre Channel connections, the StorEdge array is the preferred storage platform for Sun clusters. In cost-sensitive installations, Sun's Multi-Disk Packs can also be used to reduce storage costs.

- *Redundant Servers*

  Sun Clusters can be deployed with two to eight servers. (Two Sun Enterprise 2 servers are illustrated, and can be vertically scaled using multiprocessor Sun Enterprise 3500, 6500 and up to Sun Enterprise 10000 servers). Most ISPs deploy clusters in symmetric configurations, which allows all servers to satisfy content requests from a partitioned data set. Note that, even though the content provided by the back-end servers is partitioned, each front-end processor has access to all content, so partitioning of mail users (for example) to FEPs is unnecessary. The symmetric configuration used by most ISPs combines the performance of *N* servers for normal operation, and degrades to the performance of *N-1* servers in the event of a failure.

- *Private Redundant Network*

  The Sun Cluster utilizes redundant private interconnects to monitor the status of each server.

- *Sun Cluster Software*

  The capabilities of Sun Clusters — once a separate software package — are gradually being integrated as standard features of the Solaris operating environment. Sun Clusters provide the intelligence by which the hardware, operating environment, and applications are monitored, as well as the mechanisms by which failures are detected and resolved. Sun Clusters provide a global operating environment in which any ISP service can be executed; Sun provides a range of services which are specifically designed to take advantage of the HA features of Sun Clusters, including NFS, database, and many ISP functions like Netscape news, mail, and Web services.

High availability NFS and database software, for example, is designed to be re-started in the event of a software failure, and to migrate to different servers in the event of hardware failures or scheduled maintenance activities. With Sun Clusters, many services can continue uninterrupted despite failures without having to re-establish connections. Sun's HA software monitors the health of services by acting as clients, constantly evaluating availability and performance, initiating the re-routing of requests, re-starting of the service, or migrating or initiating migration to a different service as specified by the Sun Cluster configuration.

Deploying Three-Tier Configurations

The three-tier solution is popular with ISPs because it is highly scalable. In fact, some ISPs even use multiple content server clusters in their content network, further increasing their storage capacity. This approach is convenient for ISPs because they do not have to enter into the complications of creating static partitions of user data such as mailboxes, requiring different sets of users to attach to different mail servers to retrieve their mail.

The three-tier approach is flexibly scalable because different layers and separate functions can be scaled independently, affording ISPs the greatest amount of freedom in their configurations. The services network can be horizontally scaled by adding more front-end processors. When content servers become saturated, new clusters can be added to absorb the workload. Using the three-tier approach to deploy separate Web and content servers allows the greatest degree of functional scalability; HTTP and content servers can be scaled independently, giving ISPs hosting large amounts of dynamic content a high degree of flexibility.

Summary

Providing high-availability Internet services requires coordinating interactions between DNS, load-balancing mechanisms, and horizontally scaled server configurations. A three-tier approach using gateway servers, front-end processors, and content servers is the most flexible approach for ISPs because scaling of different functions — such as HTTP and dynamic content support, front-end services, and storage — can be handled independently. High-availability services for ISPs can be delivered with the increasing sophistication of Sun Clusters, which tightly integrate load-balancing mechanisms with HA services. Sun is committed to clustering approaches in all its products, so the current state-of-the-art in high-availability services for ISPs is likely to advance quite rapidly.

Value-Added Web Hosting 7

Beyond the basic ISP core services that include mail, news, Web, and Internet access, ISPs are increasingly providing other value-added services to enable competitive distinction and supply additional revenue. The most widely-adopted additional service is Web page/site hosting, with virtually 100% of ISPs offering some form of hosting. Web hosting services can be categorized into three types: non-commercial, or ‘vanity’ pages, light commercial, and commercial. Each of these three categories have different requirements for service and management.

As the popularity of hosting has increased, so has the complexity. The industry has evolved from hosting basic static content to delivering dynamic content, advanced applications, and electronic commerce functionality. Accompanying this increased complexity are issues of performance, availability, and security surrounding proprietary data and electronic transactions.

Static vs. Dynamic Content

One of the major factors influencing the design of Web hosting services is whether the sites deliver static or dynamically-generated content. The most straightforward Web hosting service to offer is static content hosting. Static content does not change, or is only changed manually by the Web site creator. This includes HTML files, text files, pdf files, graphical images, and other media. Each customer is allocated disk space according to their service level, and FTP or direct access to their own document root directory.

Dynamically-generated Web pages, in contrast, use some programmatic method to create content. These pages can contain dynamic content such as stock quotes and customer-specific information such as order status. A good example of a dynamically-generated page is an application where the viewer can query the Web site and cause the material to be assembled using some template mechanism. Dynamic pages are generated using a wide range of facilities including basic Common Gateway Interface (CGI) and Web server APIs. The most common approach is to use an application or process, separate from the Web server, to process a request and perform some function. This is a fundamental difference from static pages, for which the Web server simply delivers the file content.

Resource Requirements

Static content presents little infrastructure demand on the ISP. Static content is generally easily hosted in a shared environment. In addition there are relatively few security concerns.

Dynamic content, on the other hand, has special security considerations, both on the part of the ISP and its customers. Given that various programming methods are used to provide dynamic content, there is the possibility that these programming methods may cause problems within the site. The source of these problems may be either programming errors generated by the customers, or attacks mounted on the site by hackers.

For these reasons it is often desirable for true commercial Web sites to use dedicated resources like servers and disk systems for hosting each customer’s dynamic content. In addition, dedicated resources provide specific hardware methods for ensuring quality-of-service levels for each customer. They also ensure that each customer’s programs and data do not interfere with, or physically intermingle with those of other customers.

In addition, different Web server software might be used for different application functionality. For example some sites might inter-mix the use of commercial servers like Netscape’s for hosting and electronic commerce functions — with the use of high-performance servers such as Zeus or Apache for GIF delivery. Many ISPs are finding that servers like Apache are suitable for light commercial use where static pages such as catalogs are hosted.

Application Architecture Complexity

There are several application architectures to choose from. These are presented in order of increasing complexity, where the more complex architectures offer the greatest benefits in terms of scalability, security, performance, and support for dynamic content generation.

- *Static Content using CGI*

  This is a simple form of dynamic content generation that includes Web pages created by CGI processes that access static file content. While this architecture does provide dynamic content, it does not maximize performance, generality, and maintainability. There are several alternatives which provide additional functionality and performance, but which also increase the application architecture complexity.

- *Database Access Using JDBC™ and SQL*

  The level of sophistication is to utilize an interface from the page generation mechanism to access a generic database system. This can be done, for example, via standard SQL interfaces or a JDBC interface. A major advantage of this model is that the database content is readily linkable (with other applications) and easily maintainable. This allows for an increase in the sophistication of the underlying data resources, but does not result in significantly-increased application functionality.

- *Application Servers*

  Application servers — like those from Net Dynamics and Netscape — can provide dramatic improvements in application functionality. There are drawbacks to using a proprietary programming model; however, they are outweighed by the potential improvements in programmer productivity resulting from using these servers in Web applications. In addition, the use of application servers provides improved performance and scalability.

- *Transaction Servers*

  In complex Web applications, the Web server may interact with several back-end resources and could involve complex, high-value transactions. In large-scale Web applications, transaction servers like BEA's Tuxedo or M3 can improve the performance of complex transactions and can accommodate multi-phase commit. The additional transaction processing monitor improves performance and reliability.

- *Multi-Tier Architectures*

  While each of the above application architectures could be deployed on a single, monolithic server, there are significant advantages to multi-tier models. Each layer of the architecture — Web server, application server, transaction server, and database server — provides specific server resources that can be tuned and scaled.

  While the resulting multi-tier application architectures may seem quite complex, a multi-tier environment offers superior scalability, system performance, and availability. Each tier can use the scalability approach best suited to the function: Web servers may be made horizontally redundant through the use of multiple servers, while database servers can be made highly available through the use of clusters. In addition, the multi-tier model naturally supports geographic distribution and multiple-site operations.

Dynamic Hosting Performance

The mechanism used to provide dynamic content dramatically affects the system performance.

- *CGI Programming*

  CGI is the most basic dynamic content mechanism, offering rapid and easy-to-code functionality. The use of C, Perl, and shell programming, however, tends to create the greatest performance impact. The CGI mechanism spawns additional processes which are constrained by basic operating system resources. A Web server engine that waits on CGI functionality may deliver only one-eighth to one-tenth of the Web server performance (in hits per second) of a similar, static Web page server.

- *Compiled code*

  One alternative for improved performance is the use of a Web server that allows compiling of the application code into the Web server base. For example, Apache allows the compiling of application code (e.g., C, modPerl) into the Web server itself. This results in a two- to four-fold increase in performance. On the other hand, while Apache offers good performance, it is a public-domain server and does not have the commercial support that large ISPs desire.

- *Server APIs*

  Perhaps the best combination of performance and commercial support is the use of proprietary Web server APIs like the Netscape API *NSAPI*. Use of proprietary APIs results in code for dynamic content or other application functionality that can deliver performance up to eight times that of CGI — nearly that of static Web page serving. Another way to flexibly manage services is through the public-domain Servlet API, which is being supported by an increasing number of Web servers.

Dedicated vs. Shared Hosting

The decision on whether to offer static or dynamic content is accompanied by a decision regarding how to allocate hardware resources to each hosting customer. Application hosting has evolved to include both dedicated and shared-hardware resource hosting models.

Vanity hosting is generally provided via a shared, static content model. A user authoring tool may be supported and pages can be uploaded using simple FTP or even e-mail. Light commercial hosting may include either static or dynamically-generated content. This level of service is typically provided in a shared environment, although some ISPs may offer dedicated light hosting on small server platforms. Heavy commercial hosting generally includes dynamic page generation and uses a dedicated platform which provides both isolation of load impact and higher quality of service.

In both the shared and dedicated models, the platform operator determines system hardware and software configurations. Customers may select components but cannot interpose their own configurations or products. It is imperative that the operator be familiar with all components of the platform. For this reason, pre-defined, standard configurations should be prepared and tested.

Dedicated Hosting

Dedicated hosting provides dedicated hardware for each customer application (Figure 15). Dedicated hosting automatically provides a mechanism for physically outsourcing Web sites and for offering independent quality-of-service levels to each customer. Any service interruptions — whether due to system failure, maintenance, Web site data or programming — affect only the individual customer.

Management of these dedicated Web sites is straightforward, but the administration of a very large number of systems does result in some inefficiencies, for example, the time it takes for upgrading the OS on multiple systems.

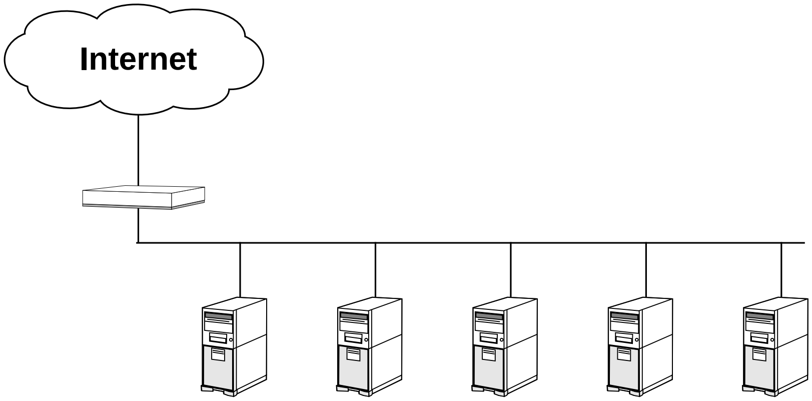

Figure 15 Dedicated hosting provides specific, independent hardware resources for each Web site that is hosted.

Advantages of Dedicated Hosting

- Ease of administering individual sites — each customer is an independent unit with its own configuration. No additional management framework is required beyond normal system administration and customer provisioning.
- No cross-contamination between customers — all event, configuration, and performance impacts are confined to individual customers. Compromises in security, the overloading of an application, or application failure will only impact a single site.
- Isolated quality of service — the dedication of hardware enables specific quality-of-service levels to be offered to individual customers according to a pre-set fee structure.

Disadvantages of Dedicated Hosting

- Cost of dedicated hardware — The most significant disadvantage to dedicated hosting is the entry-level cost to support a customer because each site must be allocated a server and disk resources.
- Administration overhead of large server farms — while the administration of dedicated hosting is straightforward and does not require a specialized framework, it also becomes more difficult to perform basic tasks (e.g., software upgrades) over a large number of servers.
- No flexible allocation of resources — dedicated hosting results in a lack of flexibility in re-allocating resources between customers. For example, in addition to the core servers for a customer, any redundant system components must also be allocated to each customer. This means sophisticated additional services may not be offered because they are not cost effective for individual customers.

Shared Hosting

Many early hosting demands were for static commercial sites, with early Web servers used for this purpose offering little in the way of virtual Web site hosting capabilities. Today, individual users require Web sites hosted with their own domain names, and most Web servers have added the capability to provide many virtual sites from one server. Shared hosting provides an improved outsourcing model with additional economies of scale, lower cost of entry, and potential interoperability. Larger, more cost-effective servers and disk arrays can be utilized (Figure 16).

Advantages of Shared Hosting

- Lower cost — the greatest advantage of shared hosting is the resulting lower cost of supporting individual customers.
- Improved leverage of resources — inherent in the shared hosting model is the ability to flexibly allocate hardware and other system resources across customers. Adding more disk to a customer site is as easy as changing a configuration setting for a pooled disk resource, rather than the actual installation of additional hardware.

- More advanced services — In addition to lower cost for basic core hosting services, the economies of shared hosting may make it possible to offer advanced service functionality. This includes both point technologies, such as electronic commerce transactions or disk caching, as well as the ability to offer superior, more functional application architectures.

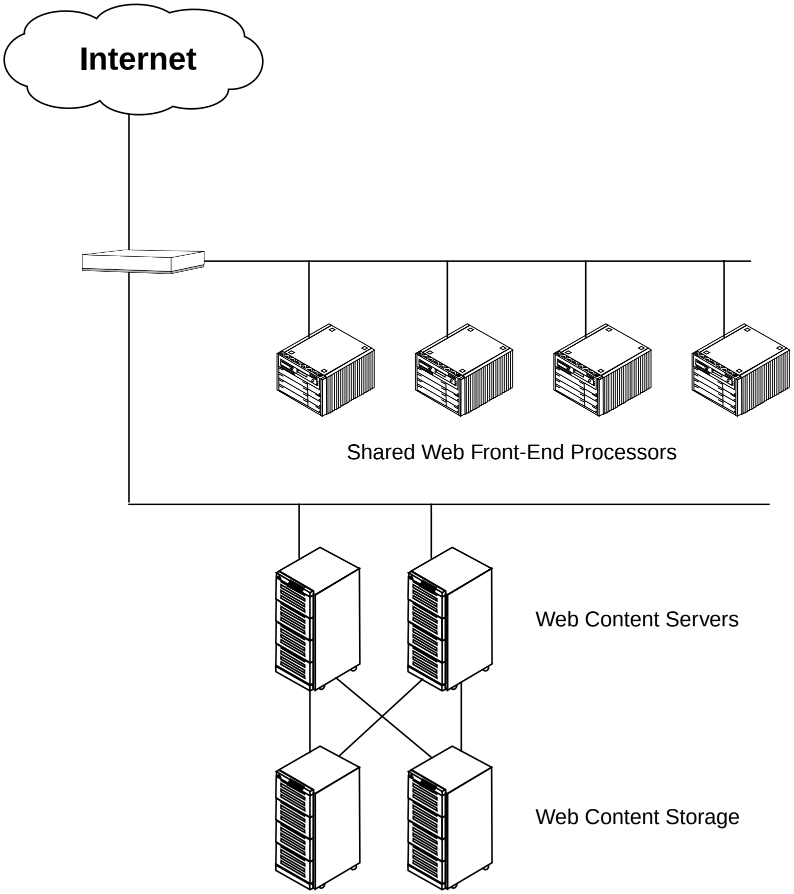

Figure 16 Shared hosting environments can exploit the economies of scale and reduced cost of ownership of larger servers and disk arrays.

Disadvantages of Shared Hosting

- Requires special administrative framework — in general, both server hardware and application software are managed for individual customer utilization — not for shared hosting. For example, shared Web hosting requires the ability to run multiple virtual Web servers answering to different names all on one system.
- Possible cross-contamination — while dedicated hosting means that no customer's problems will affect others, shared hosting means that a security compromise or performance problem affecting one customer has the potential to affect others.

Site Management

Site management of Web hosting facilities is an important consideration for ISPs. Issues include provisioning, content development, and facilities for meeting quality-of-service guarantees.

Provisioning

Provisioning includes all of the activities to create and configure a hosted customer and their environment. Dedicated hosting environments involve significant intervention by systems support staff. Shared hosting environments can, on the other hand, often be provisioned directly by the customer.

In addition to the configuration of hardware, provisioning includes the configuration of appropriate identity and naming conventions (e.g., DNS entries, and user profiles), and the configuration of system services (e.g., billing).

Provisioning by the ISP

Web hosting is usually provisioned directly by the ISP after the customer has selected the desired level of service and functionality. This involves the configuration of any dedicated system resources — including the server itself — or simply the configuration of file space on a shared server for a shared hosting customer. This process may be augmented by an on-line service request form. Personal home page hosting may be implemented automatically, while more significant commercial Web site hosting may require the answers to many questions asked by a sales agent.

Automated Provisioning

While basic provisioning may be accomplished using manual procedures and independent service applications, service providers may want to implement automated provisioning systems to improve this process:

- Automated provisioning improves the speed at which the platform operator may implement services and reduces or eliminates the possibility of configuration errors.
- Use of automated provisioning provides an audit trail of provisioning and administration activities.
- The delegation of certain customer care activities to the customer itself reduces the platform operator's workload and enables the customer to effect and review service changes more easily.

Typically, automatic provisioning has been accomplished with custom scripts developed by ISPs. More recently, products are emerging that provide additional support for easy, automatic addition of hosted customers.

Self Provisioning

Self provisioning involves a framework and application that enables the delegation of specific provisioning capabilities to not just the ISP staff, but to customers themselves. Generally, each delegated administrative authority has the ability to perform specific functions by means of a Web interface. For example, customers may change passwords without involving the sales agent or the system administrator. However, the system administrator must be engaged in order to create a new site.

Content Development

The type of content development support offered usually depends on the nature of the hosting being supported. Heavy commercial customers requiring dedicated hosting of dynamic content will most often provide their own content creation platform. Light commercial users are frequently attracted to content authoring support if it is offered.Vanity users, on the other hand, will not have a sophisticated content creation capability and are often offered content creation support including packages such as Home Site. For these users, ISPs may support server-side extensions like those for Microsoft FrontPage, however extensions inevitably have support, security, and performance implications.

Staging

Staging involves the placement of Web content into a staging area, prior to it actually being placed on-line. This may involve the testing or analysis of the content.

Log Analysis and Viewing

In addition to content creation tools, the ISP may offer log analysis capability to its commercial customers.

Quality of Service

Quality of service is a critical differentiator between Web hosting services. Consideration must be given to controlling the quality of service offered to Web hosting clients, monitoring service levels, and managing resources.

Quality of Service Controls

As a value-added service for which there may be additional service charges, ISPs may offer quality-of-service controls. The platform must include mechanisms to ensure that customers get the performance for which they contract. Before the platform operator may offer quality-of-service guarantees to customers it is necessary for the following to exist:

- Controls that limit the consumption of systems and network resources
- A measurement system that demonstrates the utilization of these resources

The consumption of system resources can be controlled in a number of ways: First, by utilizing a multiple-tier architecture, and by partitioning services according to their relative impact on the underlying resources (i.e., resources are built to isolate load impact) Second, by applying system and network resource management within various systems (i.e., monitoring is used to control load impacts). These methods are discussed in the following sections.

Another critical component of quality of service is reporting to the customer and to the platform operator's management. Metrics and statistics should be delivered to the customer in order to demonstrate compliance with contracted quality-of-service levels. Management reporting enables the customer to determine which quality of service best suits the application while allowing the platform operator to tailor services to the customer's needs. For example, a customer who is consistently reaching their service limit should be offered an increased level of service and be charged accordingly.

Monitoring

To gather sufficient information to check quality-of-service levels, it is necessary to augment basic system monitoring with monitoring of the Web services themselves.

- *System and Network Monitoring*

  For simple monitoring, ISPs can take advantage of basic system and network monitoring services. These include SNMP and native Solaris and server performance monitoring capabilities.

- *Service Monitoring*

  Server hardware may continue to operate and fail to report alert conditions even though an application has hung or failed. In addition to basic server-level monitoring, the applications themselves must be monitored. Sun™ Internet Services Monitor™ software — a component of Solaris ISP Server — monitors the way ISP Internet services perform and appear to customers. It measures the performance of six Internet protocols — HTTP, IMAP4, POP3, SMTP, NNTP, LDAP (v2 and v3) — and reports on performance metrics. Sun Internet Services Monitor gives the information the administrator needs in order to take action before a service fails or degrades to an unacceptable level.

  For example, at user-defined intervals, Sun Internet Services Monitor software can be configured to attempt an e-mail or a specific URL request. It then reports the elapsed time for response, or if the service is even running (service failures may generate an e-mail alarm to the administrator).

Additionally, Sun Internet Services Monitor traps response-time data and helps the administrator to identify trends in service performance. For example, a plot of mail server performance showing an upward trend in response time may indicate the need to add additional server resources.

Through a browser interface, Sun Internet Services Monitor clients are started manually or automatically at remote locations. ISPs determine which services to test, which data to send and get, and how frequently the tests should be made. After the information is collected, the results are stored in files to allow loading into other graphical packages. As a result, Sun Internet Services Monitor software increases the ability to meet customer-specific service-level agreements by giving information about network services before problems arise, and ensuring that a contracted quality of service is delivered.

Resource Controls

Delivering specific service levels is achieved by partitioning services, carefully managing server resources, and allocating specific levels of bandwidth to customers.

- *Service Partitioning*

  In service partitioning, services are separated according to their logic boundaries, such as Web server, application server, and database server, and then arranged in a multi-tier configuration on the available hardware. For example, there are several different types of application execution platforms including CGI, JavaScript, Java servlet, and application server. It may be necessary to separate these platforms for reasons of performance, security, and feature set. It is likely, for example, that the CGI execution platform will be separated from other application servers since this type of Web application presents specific concerns with respect to security. Therefore, CGI programs will be executed on their own server or a cluster of servers dedicated to this purpose.

- *Resource Management*

  Controlling resources within a shared server is necessary in order to guarantee quality of service. The Solaris™ Resource Manager™ provides the ability to control and allocate CPU time, processes, virtual memory, connect time, and logins. This control can be applied on a fine-grained, hierarchical basis — making it possible to define a resource pool for each customer Web site. Each pool is allocated a pre-defined portion of the available system resources. The Resource Manager ensures, for example, that applications running on behalf of one customer Web site will not consume more than the pre-determined amount of CPU and memory. This mechanism allows applications belonging to other Web sites to run without being starved of resources.

Another capability of the Solaris Resource Manager is to define hierarchical relationships between resource pools. Resources may be assigned to a *parent* pool that contains a number of subordinate pools, or *children*. The parent is assigned resources as usual. The child pools are assigned resources *from* the parent pool. This technique is useful for creating classes of service. The parent represents the overall class of service and the children represent the individual customers that belong to the class.

The Solaris Resource Manager ensures a fair sharing of the available system resources over hosted applications. With this system component, ISPs can control costs and server resource consumption using methods similar to those found on costly mainframes. Multiple applications, groups, and individuals users can be guaranteed a consistent level of service on a single server. By dynamically allocating unused CPU capacity and virtual memory, resource utilization is increased. Systems are easier to manage because of the ability to set and enforce policies that control how system resources are utilized, ensuring that customers will always receive the assigned service level within a shared resource environment.

Another method used to provide quality service is to deploy a network caching infrastructure. Many ISPs are placing caches in their distributed Web hosting centers to provide both improved response time and as a way to level out peaks in user demand cycles.

- *Network Bandwidth Allocation*

  Quality of service starts with a reliable, scalable, high-performance network operating environment that can quickly and safely deliver required services, and enable customers to pro-actively manage their ISP-provided bandwidth and service-level agreements. System resource allocation provided by the Solaris Resource Manager does not strictly control the use of network resources; however without some control over the consumption of bandwidth, an application may consume excessive network resources.

The Solaris™ Bandwidth Manager™ performs a function quite similar to the Resource Manager, but it is specifically designed to allocate network bandwidth, ensuring a fair sharing of the network bandwidth available to each Web server. This provides ISPs with the means to deliver upon service-level agreements set with customers.

Bandwidth allocation is most critical in the Web server front-end, since the bulk of network resources are consumed by the delivery of static content. Users typically consume static content at a greater rate than they consume content resulting from the execution of applications. The Web server front-end is far more efficient at delivering content than applications in the middle tier and will therefore be the chief consumer of Internet and intranet bandwidth. The amount of network bandwidth consumed by each customer's virtual Web server will be limited by the Solaris Bandwidth Manager in order to affect fair sharing. For networks requiring regulation of internal bandwidth, Solaris Bandwidth Manager also provides an effective solution.

Security Issues

With security attacks that sometimes leave Web sites completely disabled or contaminated with hacker-provided content, security is a paramount concern of ISPs providing Web hosting services.

Security Controls

Because it involves customer access to content servers, the ISP’s Web hosting environment should be part of a comprehensive security strategy. This strategy includes both the implementation of security features including firewalls, as well as a security policy.

The implementation of the security policy is facilitated over the ISP infrastructure in general, and the Web-hosting environment in particular, by dividing the systems architecture into *security zones* (Figure 17). A security zone is a collection of system, content, and processes that require the same, or similar, level of protection. The security zones can be used to compartmentalize access to systems by users authorized to update content.

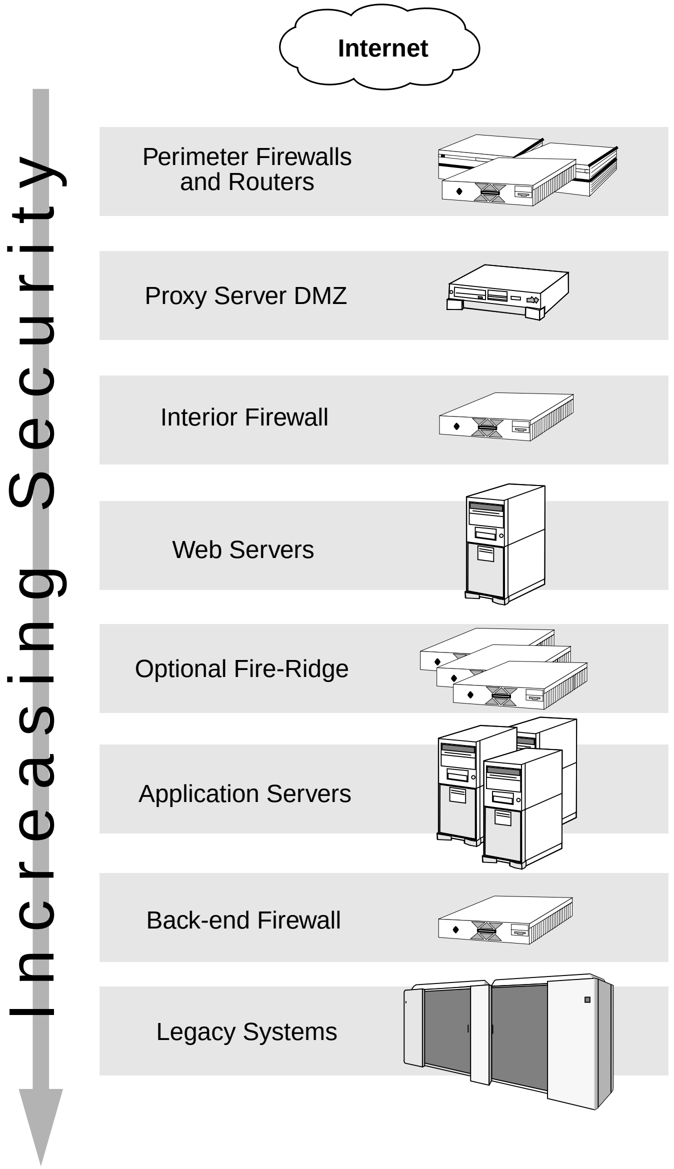

Figure 17 Security zones are separated by firewalls, with careful consideration given to the accesses allowed from one layer to the next.

These security zones can be separated by firewalls that filter traffic entering and exiting the platform. Only authorized protocols, and source and destination addresses are accepted. All other attempted connections will be dropped and the firewall software will generate log messages.

Although the Internet is generally considered the least secure layer in the security model, and the legacy systems as the most secure, the model becomes divided when the legacy systems represent services such as stock ticker feeds, which are susceptible to outside intrusion. In this case, the Web hosting infrastructure needs to be protected from the legacy systems, with the back-end firewall and optional fire-ridge providing the needed protection.

In addition to the careful deployment of firewalls, a number of additional precautions should be taken, including:

- *Server Hardening*

  Each server should be pruned of any unnecessary programs, services or configuration options. In addition, all security-related operating system and application patches should be installed.

- *Cryptographic File System Integrity Testing*

  Automated processes can be used to test the integrity of files and institute permissions. Several times a day a cryptographic (MD5) signature is computed for each critical system or application file. These signatures are compared to a set of “known good” signatures stored in a read-only file system.

- *Event Logging*

  System and application event messages should be forwarded to a centralized log facility.

- *Threat Detection and Risk Mitigation*

  Several public-domain programs can test systems periodically for known security weaknesses. Such programs should be run on a frequent basis.

- *Penetration Testing*

  Security policies and their implementation are most effective when augmented by periodic penetration testing by an outside service.

Secure User Authentication

When Web applications require authentication and encryption, centralized LDAP directory servers can be used to authorize users and applications. This directory service can allow for traditional password challenges as well as verification by way of digital certificate. In order to support authentication via X.509 certificate, the client browser and/or the server must obtain a digital certificate from a certificate authority (CA). There are various commercial and non-commercial organizations that generate and provide custody for digital certificates. In other cases, the ISP may wish to establish their own certificate authority in order to maximize control over the issuance, custody, and revocation of certificates.

Data Encryption Using SSL

Encrypted Web site sessions using the Secure Socket Layer (SSL) are possible, with client-side and server-side certificate-based authentication (i.e. X.509) as necessary. Secure communication consists of three basic components: authentication, data concealment, and integrity. SSL offers authentication based on a secure exchange of digital certificates and session encryption.

Users selecting a secure page or site are directed to a Web server that supports SSL. Since the encryption process associated with supporting an SSL connection can be a large drain on the Web server in terms of resources, it may be wise to dedicate separate server systems to handle SSL sites. These front-end servers may be outfitted with additional CPUs or even hardware-based encryption accelerators.

Summary

Many ISPs provide Web-hosting capabilities to casual, light commercial, and heavy commercial users. Web hosting is a specific case of general application hosting, and it requires ISPs to consider whether to provide dedicated or shared hosting. Management facilities — ranging from site management to automatic provisioning — must be provided to ensure effective administration of services. Finally, security against attack on the application services and the user data itself must be established.

Local ISP Configurations 8

Local Internet Service Providers typically handle between 5,000 and 10,000 subscribers. With a concurrency rate of 10 percent, this requires accommodating a maximum 500 to 1,000 simultaneous users. This chapter examines a typical local Internet Service Provider architecture. The example provided in this chapter is a good prototype for those considering entering the ISP market with approximately 5,000 subscribers.

Characteristics of Local ISPs

Local Internet Service Providers can use standard tools, deploy small networks — often with entry-level servers. ALthough performance and reliability are important to these ISPs, they can tolerate sine compromises:

- *Standard Tools.* Local ISPs can use many standard UNIX tools and off-the-shelf software. The choices range from standard utilities provided with the Solaris operating environment to sophisticated packages that can ensure scalability long into the future, including Solaris ISP Server and Sun Internet Mail Server.
- *Small Configurations.* Local providers tend to use networks with a small number of servers to achieve maximum functionality at low cost.
- *Entry-Level Servers.* Local ISPs often deploy low-cost, entry-level servers while maintaining the ability to upgrade to more high-performance servers as needed.

- *Single Points-of-Failure.* Compromises are often made in these ISP architectures, leaving open areas in which a single failure can bring down some or all of the ISP's operations.

The architecture described in this chapter reflects a choice to provide a reasonable level of service at a low cost. Its capabilities can be improved depending on performance, reliability, flexibility, and security needs, and it can be scaled down if it is necessary to further reduce costs or to handle fewer subscribers.

Architectural Overview

Like Internet Service Providers of all sizes, the basic configuration (Figure 18) consists of a set of sub-networks isolated by firewalls. The cornerstone of this local ISP is an Enterprise 5S server that is configured with SunScreen EFS software and a low-cost, Quad FastEthernet™ card that provides security and isolation between three basic networks:

- The access network provides connectivity to the ISP from the Internet and from dial-up users, as well as routing between dial-up users and the Internet.
- The services network provides mail, netnews, and Web access. Additionally, the domain name service visible to subscribers and the outside world is hosted in this network.
- The internal network provides authentication, billing, and network management services — and is the most carefully-protected sub-network in the ISP design.

Access Network

The core of the access network is a router that handles network traffic to and from the Internet over a redundant pair of T1 connections. This Internet connectivity provides an aggregate bandwidth of approximately 3 Mbps, and the two connections ideally attach to the Internet through two different providers. This reduces the chance that a single network failure would cut off all outside access.

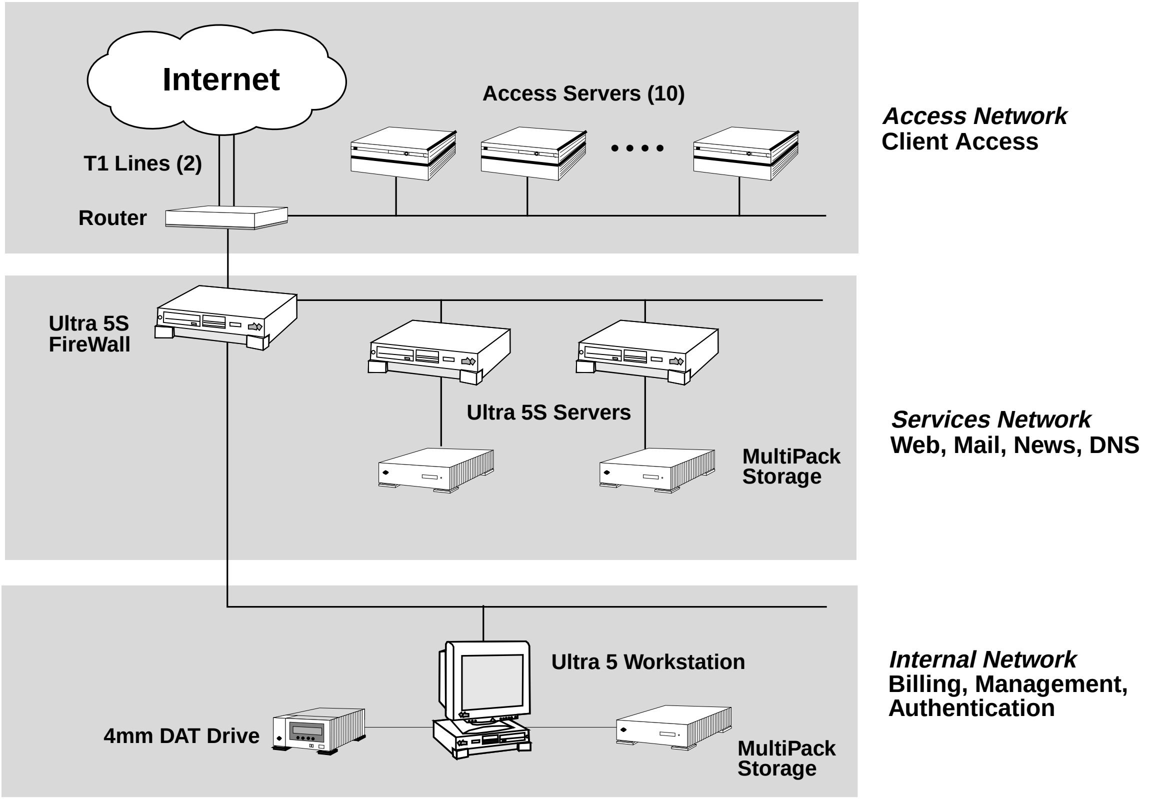

Figure 18 Example local ISP architecture

Dial-up access is provided by a set of 10 access servers which provide a total of 480 dial-up connections — sufficient for 4,800 dial-up subscribers at a 10 percent concurrency rate. Note that the access network enters the ISP facility through the core router. This enables packet filtering between both dial-up and Internet users and the services network.

The access network can be scaled to support more dial-up customers by adding more access servers. Reliability can be increased by configuring a redundant pair of routers so that a failure at this point does not curtail Internet access.

Alternate Head Ends

Although this network is configured for dial-up users, the set of access servers could be replaced with access devices for a number of different media — including wireless, satellite, and cable. Choice of media affects ISP sizing, and must be considered if different head-ends are configured. For example, a dial-up network such as that illustrated in this chapter has a built-in load limit — it is physically impossible for more than 480 subscribers to be connected at one time.

In contrast, cable company ISP subscribers are always attached through cable modems. For cable providers, it is more difficult to limit peak workloads — and experience shows that cable ISP subscribers use ISP services more heavily than dial-up subscribers. Sun estimates that a dial-up ISP can serve three times more subscribers than a cable ISP with the same infrastructure.

Address Translation

Because of the shortage of Class B and Class C addresses, some ISPs configure dial-up users with private IP addresses, translating them to public addresses as requests are passed to the Internet. This allows sharing of a large number of private addresses on a small set of official, public IP addresses. This functionality is available with SunScreen EFS software, and could be added to this network by connecting the collection of access servers to a separate subnet attached to the firewall.

Hubs vs. Switched Ethernet

Note that the network diagram is a logical perspective that does not include the physical wiring details of the ISP. An Ethernet hub or switch is needed for each of the three independent sub-networks. Small ISPs typically use Ethernet hubs in order to minimize costs. There are performance and security issues with hubs, however, making switched Ethernet connectivity preferable. Because Ethernet switches transfer packets only between source and destination interfaces, collisions are minimized and throughput is enhanced.

This has an important security implication — in switched Ethernet configurations each host sees only packets directed to it — not to any other hosts. If an intruder manages to compromise security on one host and attempts to snoop the network using promiscuous mode on the interface, packets destined for other hosts cannot be observed.

Although this network could be based on 10 Mbps Ethernet, Sun advises the use of 100 Mbps Ethernet backbones for ease of scalability. With many servers — like those from Sun — coming with 10/100 Mbps Ethernet as standard equipment, the incremental cost of hubs and router interfaces is worthwhile.

Services Network

The services network provides all of the familiar ISP services to subscribers, including mail, netnews, and Web services. It is populated with two Ultra 5S servers, disk storage, and one internal 4mm DAT drive for backups using third-party backup software:

- The news server is an Ultra 5S server configured with 128 MB memory, and 25 GB of disk space in an external storage pack. News software is public-domain Inter-Network News (INN) or commercially-supported products like Sun™ Internet News Server™. This host acts as a primary DNS server.
- The mail and Web server is an Ultra 5S server configured with 128 MB memory and 25 GB of external disk space. This disk space allows for more than 2 MB per user for mailboxes, and 4 MB for each hosted Web page, assuming 10 percent of users host Web pages. Services on this host can be provided with Sun™ WebServer™ and Sun™ Internet Mail Server™. This host acts as a secondary DNS server.

The allocation of services to these two servers is somewhat arbitrary. The ISP's strategic business goals and expected workload should be used to determine the actual allocation at deployment time. This starting point was chosen because netnews has a heavy, constant bandwidth requirement that can limit the ability of the mail and Web services to handle the peak demands of their services. The 4mm backup drive is configured on the mail/Web server because backups for mail are more important than for netnews. The Multi-Disk packs for each of these servers are configured on their own Fast/Wide SCSI chain for performance and ease of configuration. At additional cost, a high-capacity DLT drive could be substituted for the DAT device. With a capacity of up to 20 GB, the DLT drive would reduce the amount of operator involvement necessary for backups.

Advantages of Solaris ISP Server

Because local Internet Service Providers typically do not require custom software solutions, Solaris ISP Server along with Sun Internet Mail Server can satisfy most of a local ISP’s service needs with LDAP-based integration — speeding installation, simplifying initial configuration, and providing easy-to-use management interfaces:

- Mail services are provided by Sun Internet Mail Server, which provides SMTP, POP3, and IMAP4 servers, all running on an Internet-tuned version of the Solaris operating environment and utilizing a high-performance two-level message store.
- News service is provided by Sun™ Internet News Server™, which is an integrated feature of Solaris ISP Server™.
- World Wide Web services are provided with the Sun WebServer, a high-performance, secure Web server that provides support for creating and maintaining the diverse forms of data stored in sophisticated Web sites.
- Domain name service is provided with hardened services from Solaris ISP Server that can act in primary, secondary, or caching mode.

Scalability

The services network can be scaled up or down depending on the expected number of subscribers and the workload the ISP expects to encounter.

Horizontal Scalability

Horizontal scaling is probably the first possibility to consider in the services network. An additional server could be added to allow mail, news, and Web services to be hosted on separate systems. It is possible to scale down to a single server, however the existence of at least two hosts ensures that the failure of one server does not bring down all of the ISP’s services.

Vertical Scalability

Vertical scaling is secondary to horizontal scaling in such a small configuration. A more flexible starting point for this network would be to use two Sun Enterprise 2 servers. As performance requirements dictate, one additional 300 Mhz UltraSPARC processor could be added to each server as needed.

Reliability

Reliability of the services network could be improved by the addition of more servers, more reliable storage subsystems, and high-availability clusters.

Additional Servers

Adding a third server would allow each service to be independently partitioned onto its own host. This step would ensure that the failure of a single server would cause the ISP to lose only that function, and users could still access the remainder of the ISP's services.

Storage System Reliability

The storage systems specified in the services network are low-cost multi-disk packs where the failure of a single disk has the potential to make the server unavailable. An upgrade to this storage solution would be to use the RAID capabilities of a Sun StorEdge A1000 unit to eliminate the disk farm as a single point-of-failure. There are two ways in which this could be handled:

1. A single Sun StorEdge A1000 array could be deployed and configured with separate sets of disks reserved for each host. Using the dual-porting feature of the array, one set of disks could be accessed by each server. This is an initial, low-cost way to bring RAID technology to a local ISP (Figure 19).

2. Each server could be configured with its own array, eliminating the possibility that a single controller failure on an array could bring down both servers. This configuration also provides for significant growth in disk space without changing the nature of the storage architecture.

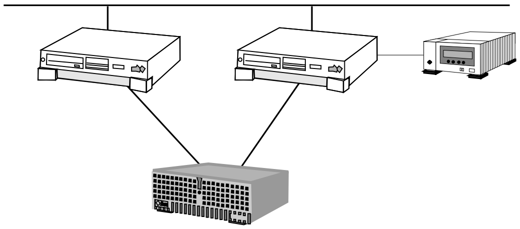

Figure 19 Reliability with a dual-ported SPARCstorage Array

High-availability Clusters

The most sophisticated — and easy-to-implement — approach is to deploy Sun Cluster HA-ready Internet software (Figure 20). The Sun Cluster can be configured with two Ultra 2 servers and a pair of cross-connected StorEdge A1000 disk arrays. A private pair of network connections allows each server to reliably monitor the status of the other server.

Each array is configured with two sets of disks, one for each connected server. Each set is configured as a RAID device, and mirroring is used to keep a consistent image of each partition on both arrays.

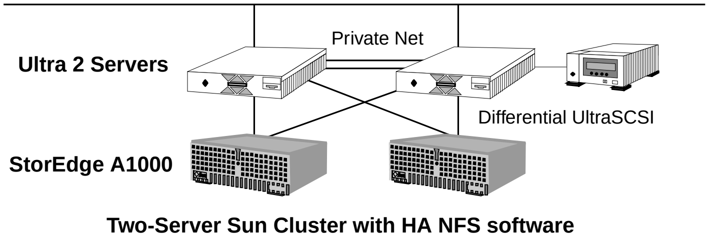

Figure 20 Reliability with a Sun Cluster based on two Sun Ultra 2 Servers configured with high-availability NFS software.

Internal Network

The internal network is where the business operations for the ISP take place. In a local ISP installation, one Ultra 5 workstation can be used for authentication, billing, network management, and problem tracking. Because this network contains mission-critical data for the ISP, it must be protected accordingly. Some ISPs have additional authentication servers outside of the internal network so that network traffic can be restricted to pass between the internal network and the services network only.

Authentication and Billing

Software for authentication and billing is beyond the scope of this document, however there are many excellent solutions available from Sun's integration affiliates.

Network Management, Backup, Customer Service

There are many choices for network management, backup, and customer service functions. Network management can be provided with tools like Sun Enterprise SyMON and HP OpenView. Other third-party tools can be used for backups and for streamlining help desk support.

Scaling

The administration system is an Ultra 5 workstation with 25 GB of external storage for billing software and data, and a 4mm DAT drive for backups. As the ISP subscriber base grows, it is easy to add more workstations to the internal network, each hosting separate functions.

Summary

Local Internet Service Providers can use off-the-shelf products — like Solaris ISP Server and Sun Internet Mail Server — to configure services and quickly come on-line. Compromises in areas such as performance and security can be made to reduce costs, However even with low-cost configurations like the one used in the example, additional servers and more powerful CPUs can be added as needed. While scalability is an important consideration for local ISPs, it becomes of paramount concern for regional and national ISPs.

Regional and National ISPs 9

Regional Internet Service Providers handle between 50,000 and 75,000 subscribers. National ISPs begin with a base of 150,000 subscribers, and it is not uncommon for them to scale to millions of subscribers. Competition between large ISPs is fierce, and their configurations are usually closely-held secrets. As a result, actual regional or national ISP architectures cannot be presented here. As a rule, these ISPs utilize all of the architectural techniques discussed in chapter 4 — and the way in which these techniques are applied to their networks is discussed here.

Characteristics of Regional ISPs

Regional ISPs begin to have larger, more complex networks for handling between 5,000 and 7,500 concurrent users with reasonable performance and a high degree of security. These networks are always custom-designed, and vary significantly depending on the kinds of services to be provided. Because of the custom nature of these ISP installations, consultation with Sun Professional Services or Sun's integration affiliates is essential. Some of the characteristics of these ISPs include:

- *Custom Tools.* Most standard UNIX utilities have limited capabilities for handling more than 50,000 subscribers, so specialized packages and home-grown software solutions are the rule.
- *Scalability.* A growth rate of even a few percent can result in an additional thousand subscribers, so scalability is especially key for regional ISPs. These networks must support fast and easy configuration of additional servers to

maintain performance at a target level. Ethernet using 10BaseT is no longer sufficient; 100BaseT and gigabit Ethernet are the rule throughout these networks.

- *Reliability.* Reliability is a priority for regional ISPs. These ISPs are often owned by large organizations — quite often telephone companies — who do not want the publicity of network failures to tarnish their reputations in other service areas.
- *Importance of Processes.* Sun's affiliate Solect observes that as the number of subscribers at an ISP reaches 20,000, administration costs peak. Beyond this point — in the realm of the regional ISPs — it is essential to establish automated operating procedures, including subscriber sign-up and pre-configured tools like Netscape Communicator with settings for their network. The savings in support costs make it possible to provide these packages to subscribers at no cost. Automated help desk software helps to track customer problems and ensure timely problem resolution.
- *Points-of-Presence.* Regional ISPs must extend their services across many local calling areas, making it essential to have POPs in a wide range of geographical locations. These ISPs usually prefer remote routers and modem banks that connect through high-speed lines (e.g., T1, T3, and ATM) to a single core service center. These POPs must operate in unattended, "lights-out" mode, and require telco-grade equipment.

Characteristics of National ISPs

National ISPs serving more than 150,000 subscribers inherit all of the issues of regional ISPs, and their large size tends to raise additional issues:

- *Scalability.* Still key to the ISP's success, national ISPs can provide scalable architectures in two ways. Centralized architectures can be designed with unattended POPs. Alternatively, centralized architectures can replicate regional-sized networks across the country with central control of administration and billing. Network traffic is such that even 100 Mbps Ethernet cannot handle all of the ISP's traffic, and sub-networks must be carefully designed to ensure that no single segment of 100 Mbps Ethernet becomes overloaded.

- *Reliability.* As failures in these ISP networks frequently receive national news coverage, reliability is a matter of public credibility. National ISPs must have high-availability networks that can dynamically reconfigure to work around outages in parts of the network, much like the reliability that telephone companies have learned to provide.
- *Roaming.* In order to reach more subscribers, national ISPs are beginning to form alliances with other ISPs, allowing subscribers to dial-in using another ISP's points of presence. Using an integrated authentication network, a roaming user can gain access through leased lines back to the national ISP services.
- *Proxy/Caching Servers.* Because of the huge volume of Internet traffic that can be generated by even a small percentage of subscribers accessing the same resources on the World Wide Web, a layer of caching servers is frequently deployed in regional and national ISP networks. These servers cache frequently-accessed Internet data — including that transmitted with the HTTP, FTP, RSTP, and NNTP protocols — and respond to subscriber requests with cached information. Caching hides network latency and can reduce Internet bandwidth requirements by up to 40 percent.

Architectural Overview

The architecture discussed in this chapter illustrates a fully-configured regional Internet Service Provider network along with some of the features that are found in national ISP configurations (Figure 21). This architecture features mail, news, and Web services decomposed by service, task layer, and by special function. This approach supports scaling up to the national ISP range. Security is enhanced by inserting firewalls between task layers, namely between front-end processors and storage systems. This ISP configuration consists of five separate sub-networks with inter-sub-net traffic carefully controlled by firewalls:

- The access network connects the ISP to the Internet as well as to residential and business subscribers. There are some local modem pools and connections to points-of-presence which are unattended installations of modem banks, routers, and high-speed connections to the core installation.

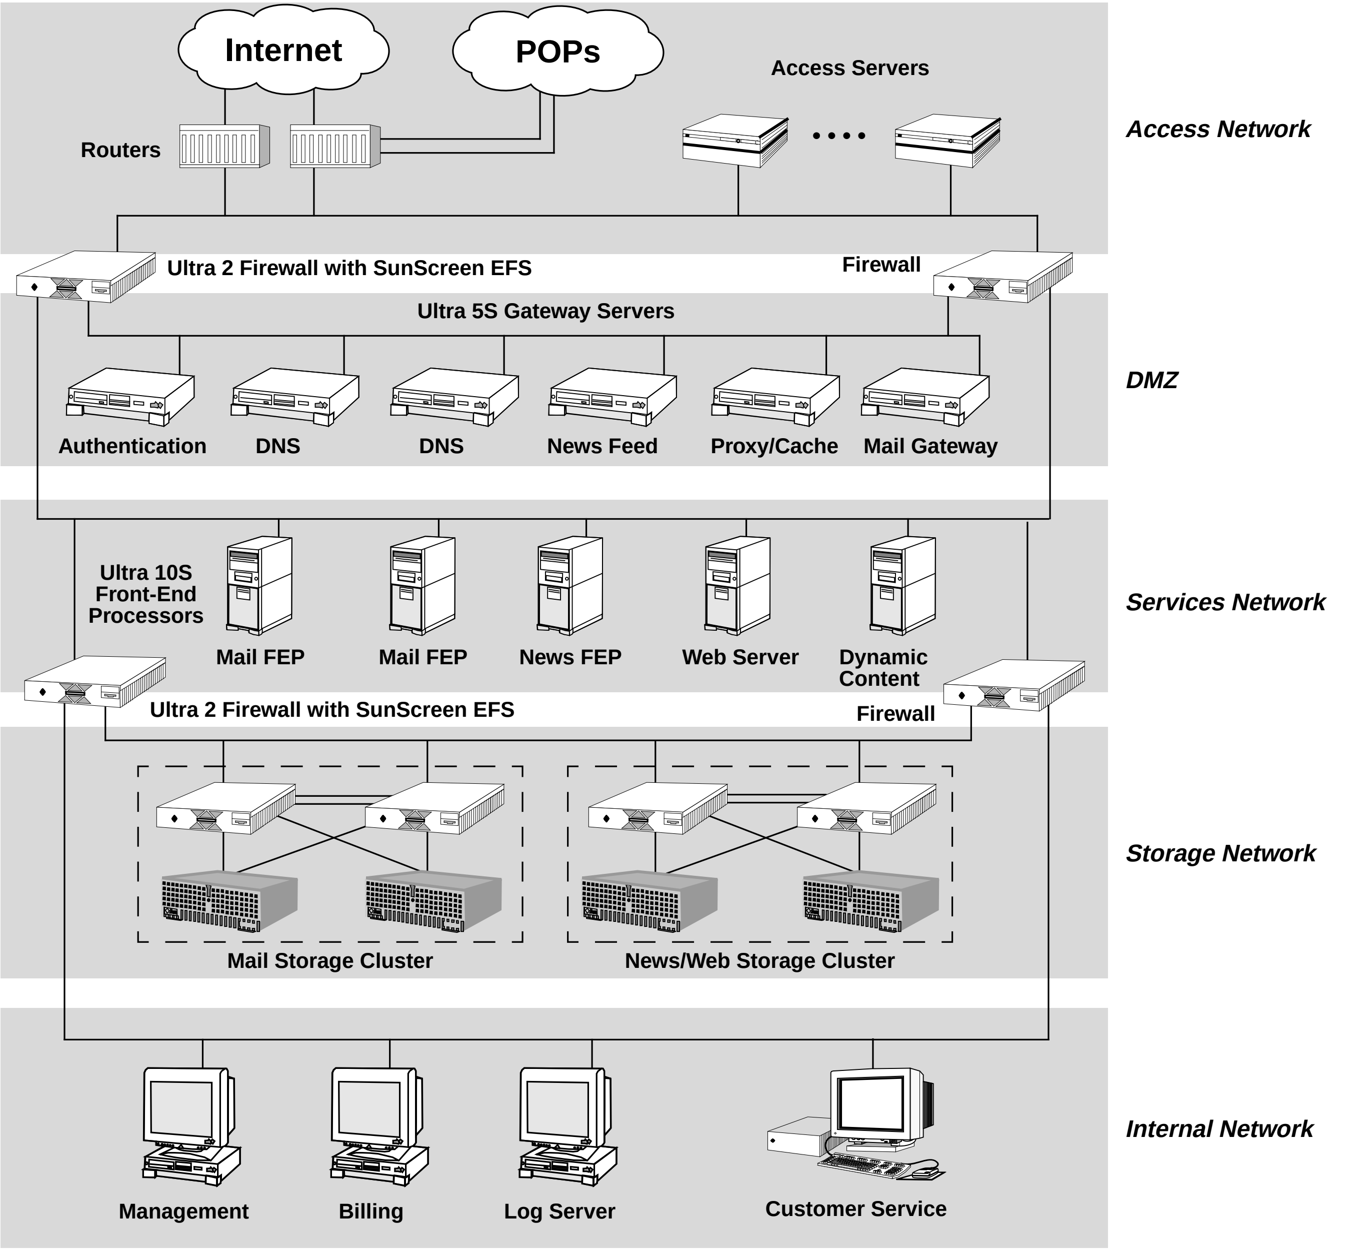

Figure 21 Regional ISP network architecture

- The access network is attached to the demilitarized zone (DMZ) through a redundant pair of firewalls providing both reliable and secure connectivity. The firewalls are single-processor Sun Ultra 2 servers hosting SunScreen EFS for optimum scalability — the Ultra 2 servers can be upgraded with an additional processor as bandwidth requirements dictate. The DMZ contains feed servers that handle incoming mail and news, a proxy / cache server that caches requested Web pages and acts as a proxy for outgoing Internet requests, and a pair of DNS servers, one primary one secondary.
- The services network contains a set of front-end processors, each supporting a single service. Each service can be hosted on a single or multiple hosts with load-balancing routers to distribute the load. Decomposition by special function is used: the dynamic content generation function is split from the Web front-end processor onto a specially-tuned server.
- The storage network is connected to the services network by a pair of redundant firewalls. This network contains two servers running high-availability content-management software that provides data to the front-end processors. Two firewalls removed from the Internet, this back-end network is highly secure.
- The internal network contains several hosts for operations, management, and customer service. A logging server collects log information from each of the firewalls, ensuring that any intrusion attempts are traceable. A Sun JavaStation™ network computer provides access to customer service functions, providing low-cost access to Java technology-enabled applications.

Access Network

The access network contains a redundant pair of routers which manage T3 connections to the Internet and high-speed connections to the POPs. Access servers manage modem pools and handle user authentication with the authentication server in the DMZ. This, and all other networks in this ISP architecture, are supported with 100BaseT Ethernet. The access network can be scaled by adding more of any of its basic components.

Demilitarized Zone

The DMZ provides a first line of defense against intrusion and gives low-latency access to services. All of the hosts in this network are Ultra5S servers with standard 10/100 Mbps Ethernet interfaces. The workload on these systems is not as critical as the decomposition of the different services onto separate hosts.

Domain Name Service

The two DNS servers provide load balancing to the services network through round-robin DNS. This capability can be enhanced by deploying load-balancing routers to allocate each request to the most lightly-loaded server, ensuring uniform response times for DNS requests.

Authentication

The authentication server communicates with local and POP-based access servers to provide user access to the network. Failure of this single server can bring the ISP down, so an important upgrade to the DMZ would be to integrate multiple servers using a cluster configuration.

Mail Gateway

The mail gateway server uses Secure Mail Access Protocol (SMAP) to receive mail and store it to a file for later transfer to the mail storage server using SMTP. Initial acceptance of mail using SMAP is preferable to using the sendmail daemon because the SMAP daemon does not actually deliver the mail, making the mail service more resistant to penetration. The mail gateway can be scaled by adding a second server and using load-balancing techniques, or by using Sun Internet Mail Server, which can partition the user workload using configuration information provided by an LDAP server.

News Gateway

The news server receives articles from upstream news feeds and delivers them to the news storage server; additional servers can be added as necessary. Since failure of this server only causes a delay in articles due to queuing on the upstream side, this server is probably the least important one to duplicate.

Proxy/Cache Server

The proxy/cache server filters outgoing Web requests and caches incoming data so that multiple subscribers accessing the same Web page receive the latest information from the cache — significantly reducing the demand for Internet bandwidth while simultaneously improving quality of service for subscribers. The Inktomi Traffic Server and the public-domain SQUID packages are popular among ISPs.

Services Network

The services network consists of a set of front-end processors that provide user access to mail, news, and Web services. These processors are configured to run their assigned services, however they contain no user data — they access mail, news, and Web pages from the storage network. This services network contains a dynamic content server to aid in the Web front-end processor's ability to quickly generate dynamic Web pages. Sun Ultra 10S servers are sufficient for these applications.

Scalability

The beauty of this architecture is that the front-end processors do not hold any user data, making horizontal scaling simple. Servers can be added by re-configuring DNS to include them in a service group, or by configuring load-balancing routers. Using load-balancing routers, a collection of mail front-end processors are addressed with a single address like *mail.isp.net*, with the router translating the address to the least heavily-loaded server each time that a request is made.

As regional and national ISPs begin to grapple with the problems of deploying large numbers of servers in each sub-network, the Sun Enterprise 10000 (Starfire) server presents an interesting option. This high-end server offers reliability, availability, and serviceability (RAS) features such as system-wide environmental sensing, resource management, dynamic reconfiguration, and hot-swap capability of all system components. These features are all critical for ISPs with large physical installations. The feature that sets the Starfire server apart from other large servers is the ability to configure it as a set of independent domains, each acting as an electrically-isolated server. This allows the ISP to exploit the RAS features of a large system while at the same time reaping the benefits of horizontal scalability — all in the same server configuration. Dynamic reconfiguration in the Starfire server allows ISPs to respond to workload fluctuations in real time by re-allocating processors, memory, and even I/O devices without re-booting any of the domains.

Storage Network

The storage network contains two sets of Sun Enterprise 2 servers configured as Sun Clusters running high-availability NFS software. Each cluster consists of a pair of single-processor Enterprise 2 servers dual-attached to StorEdge A1000 disk storage units. Configured with RAID and mirroring, these clusters can quickly detect and recover from the failure of any hardware, network, operating system, or application software component.

Scalability

The storage network is populated with two high-availability content clusters — one for mail, and the other for news and Web page storage. The host systems are Sun Enterprise 2 servers, each with the capability to accommodate additional processors; more vertical scaling can be achieved by upgrading to multi-processor capable servers. Horizontal scaling can be accomplished with the addition of separate clusters for each service, and by partitioning some services (such as mail) across multiple clusters.

Internal Network

The internal network for this ISP contains a minimum of three Ultra 5 desktop workstations. One is used for management functions, one for billing, and one for logging activities on the four firewalls.

With Java technology-based customer service and billing software already available from Sun’s affiliates, low-cost no-administration JavaStation network computers can be used to support customer service representatives.

Modifications for the Telco Environment

Many regional and national ISPs are telephone companies which require servers to conform to strict central office configuration and environmental requirements. In telco environments, all of the servers described in this chapter could be upgraded to single- and dual-processor Netra t servers. Sun’s Netra t systems leverage UltraSPARC technology from Sun into a Bellcore NEBS Level 3 certified package that is tailored for telco ISPs. It provides a compact, rack-mountable carrier-grade solution for AC power environments with alarm capability that conforms to central office alarm and system management requirements. The Netra t 1125 is packaged with single or dual 300 MHz UltraSPARC-II processors, with the ability to incorporate standard PCI peripherals on an internal 66 MHz PCI bus. Combined with the power of the Solaris operating environment, the Netra t enables telco and cable ISPs to more easily extend Internet services to their customers without compromising their high reliability standards.

Summary

Regional and national Internet Service Providers are faced with using highly-scalable software packages, and developing custom software and unique architectural solutions to the problems of scaling to large numbers of subscribers. Performance, reliability, and security are paramount, and ISPs of this class cannot accept compromises. These configurations are always pushing the state-of-the-art, and utilize combinations of all of the architectural techniques that were discussed in chapter 4.

Conclusion 10

As the Internet Service Provider market matures, a growing number of companies are providing specialized services. ISPs continue to be the on-ramp of the Internet for most subscribers; Network Service Providers specialize in high-bandwidth connections and supporting Internet backbones; Application Service Providers host services ranging from sophisticated Web sites that conduct electronic commerce transactions to Enterprise Resource and Planning services; and finally, many ISPs are covering all niches and acting as Full Service Providers.

Regardless of how the ISP market is segmented, its growth continues unabated, and one immutable fact remains — that it will take ever more computing power to satisfy the growing number of commercial and residential subscribers and their increasingly sophisticated demands for services. Independent of which market segment they serve, ISPs, NSPs, ASPs, and FSPs will have astonishing needs for scalable architectures to support enormous growth.

With the scalable hardware and software technologies available to support ISPs of all sizes, Sun and its affiliates are well-positioned to support the range of services that all providers must deliver. The constantly-changing demands for services and the quickly-changing set of ISP technologies results in every ISP configuration being different. A network designed for one set of customers today may look quite different if designed six months in the future — as the state-of-the-art advances, ISPs can exploit new configuration options to remain competitive in this ever-changing market.

This document has presented an overview of some of the services that ISPs are expected to deliver and the principles behind the network architectures that support them. Examining these principles in actual ISP networks reveals insights into how to configure new ISPs, and how to scale existing ones. This document has drawn on the experiences of Sun Professional Services and Sun's integration affiliates, namely Global One, Omnes, Solect, and Technology Applications, Inc. An ISP can only be as good as the architecture on which it is built, so partnering with an organization that understands the latest technologies — and how to implement them in ISP configurations — is key to the success of Internet Service Providers.

References 11

Sun Microsystems posts complete information on Sun's hardware and software products and service offerings in the form of data sheets, specifications, and white papers on its Internet Web page at *http://www.sun.com/*.

For more information on this white paper or on Sun in the ISP industry, please contact:

- **Sun Telco**
  901 San Antonio Road
  Palo Alto, CA 94303

  Contact:
  Bruce Baikie
  bruce.baikie@eng.sun.com

This document was prepared with the assistance of Sun Professional Services and Sun's integration affiliates. Contact information for Sun Professional Services is:

- **Sun Professional Services**
  901 San Antonio Road
  Palo Alto, CA 94303

  Contact:
  Mark Bauhaus
  mark.bauhaus@sun.com

  Web sites:
  http://www.sun.com/service/sunps/

Sun Microsystems, Inc.
901 San Antonio Road
Palo Alto, CA 94303 USA
650 960-1300
FAX 650 969-9131
http://www.sun.com

Sales Offices

Argentina: +54-1-311-0700
Australia: +61-2-9844-5000
Austria: +43-1-60563-0
Belgium: +32-2-716-7911
Brazil: +55-11-524-8988
Canada: +905-477-6745
Chile: +56-2-638-6364
Colombia: +571-622-1717
Commonwealth of Independent States:
+7-502-935-8411
Czech/Slovak Republics:
+42-2-205-102-33
Denmark: +45-44-89-49-89
Estonia: +372-6-308-900
Finland: +358-0-525-561
France: +33-01-30-67-50-00
Germany: +49-89-46008-0
Greece: +30-1-680-6676
Hong Kong: +852-2802-4188
Hungary: +36-1-202-4415
Iceland: +354-563-3010
India: +91-80-559-9595
Ireland: +353-1-8055-666
Israel: +972-9-956-9250
Italy: +39-39-60551
Japan: +81-3-5717-5000
Korea: +822-3469-0114
Latin America/Caribbean:
+1-415-688-9464
Latvia: +371-755-11-33
Lithuania: +370-729-8468
Luxembourg: +352-491-1331
Malaysia: +603-264-9988
Mexico: +52-5-258-6100
Netherlands: +31-33-450-1234
New Zealand: +64-4-499-2344
Norway: +47-2218-5800
People's Republic of China:
Beijing: +86-10-6849-2828
Chengdu: +86-28-678-0121
Guangzhou: +86-20-8777-9913
Shanghai: +86-21-6247-4068
Poland: +48-22-658-4535
Portugal: +351-1-412-7710
Russia: +7-502-935-8411
Singapore: +65-224-3388
South Africa: +2711-805-4305
Spain: +34-1-596-9900
Sweden: +46-8-623-90-00
Switzerland: +41-1-825-7111
Taiwan: +886-2-514-0567
Thailand: +662-636-1555
Turkey: +90-212-236-3300
United Arab Emirates:
+971-4-366-333
United Kingdom: +44-1-276-20444
United States: +1-800-821-4643
Venezuela: +58-2-286-1044
Worldwide Headquarters:
+1-415-960-1300

Printed in USA