Compaq ActiveAnswers

Technical Guide

April 1999
ECG557/0499

Prepared by:
Internet and E-Commerce
Solutions Business Unit

Enterprise Solutions Division

Compaq Computer Corporation

# Open Source ISP Solutions on Linux

**Contents**

***Abstract:***

This guide provides an overview of Internet Service Provider (ISP) solutions based on the software available from the Open Source community. The focus of this paper is on the Linux operating system and selected Open Source applications running on Linux platforms.

# Notice

The information in this publication is subject to change without notice and is provided “AS IS” WITHOUT WARRANTY OF ANY KIND. THE ENTIRE RISK ARISING OUT OF THE USE OF THIS INFORMATION REMAINS WITH RECIPIENT. IN NO EVENT SHALL COMPAQ BE LIABLE FOR ANY DIRECT, CONSEQUENTIAL, INCIDENTAL, SPECIAL, PUNITIVE OR OTHER DAMAGES WHATSOEVER (INCLUDING WITHOUT LIMITATION, DAMAGES FOR LOSS OF BUSINESS PROFITS, BUSINESS INTERRUPTION OR LOSS OF BUSINESS INFORMATION), EVEN IF COMPAQ HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

The limited warranties for Compaq products are exclusively set forth in the documentation accompanying such products. Nothing herein should be construed as constituting a further or additional warranty.

This publication does not constitute an endorsement of the product or products that were tested. The configuration or configurations tested or described may or may not be the only available solution. This test is not a determination of product quality or correctness, nor does it ensure compliance with any federal, state or local requirements.

Compaq, Deskpro, Compaq Insight Manager, Systempro, Systempro/LT, ProLiant, ROMPaq, QVision, SmartStart, NetFlex, QuickFind, PaqFax, and Prosignia are registered with the United States Patent and Trademark Office.

ActiveAnswers, Netelligent, Systempro/XL, SoftPaq, Fastart, QuickBlank, QuickLock are trademarks and/or service marks of Compaq Computer Corporation.

Microsoft, Windows and Windows NT are trademarks and/or registered trademarks of Microsoft Corporation.

Linux is a registered trademark of Linus Torvalds

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Ltd.

Intel, Pentium and Xeon are trademarks and/or registered trademarks of Intel Corporation.

Other product names mentioned herein may be trademarks and/or registered trademarks of their respective companies.

Open Source ISP Solutions on Linux
Technical Guide prepared by Internet and E-Commerce Solutions Business Unit

Enterprise Solutions Division

First Edition (April 1999)
Document Number ECG557/0499

# Table of Contents

# 1 Introduction

This guide offers a basic introduction to the steps involved in establishing an Internet Service Provider (ISP). It introduces Internet applications typically provided by ISPs, and describes bandwidth requirements, equipment needed, registration processes for Internet Protocol (IP) address space, domain names, and security needs. Most importantly, it will point you to further sources of information for specific details on particular Open Source Solutions utilized by ISPs on Linux platforms.

The past few years have seen explosive growth in the Open Source community, in both the quality and quantity of applications available. The general public is beginning to understand what the ISP community has known for several years. Open source software provides qualities essential to an ISP. In layman’s terms, the software is offered at no cost, source included, and includes the right to modify the software.

This lowers the entry barriers for a new ISP, and most importantly gives control to an ISP. Having access to the source allows ISPs to understand the inner workings of an application, and provides the ability to diagnose and fix a problem in a timely manner.

# 2 Getting Started

So, you want to be an ISP. Have you decided what type of services you'd like to offer? Would you like to offer consumer dial-up accounts with unlimited access, for a fixed fee? Will you include free email and web storage? Would you like to offer business customers Web page design, DNS registration, host business sites, maintain business applications/equipment? Will you offer electronic commerce solutions to your customers? Hint: the money is in the business side.

The following figure gives an architectural overview of an ISP.

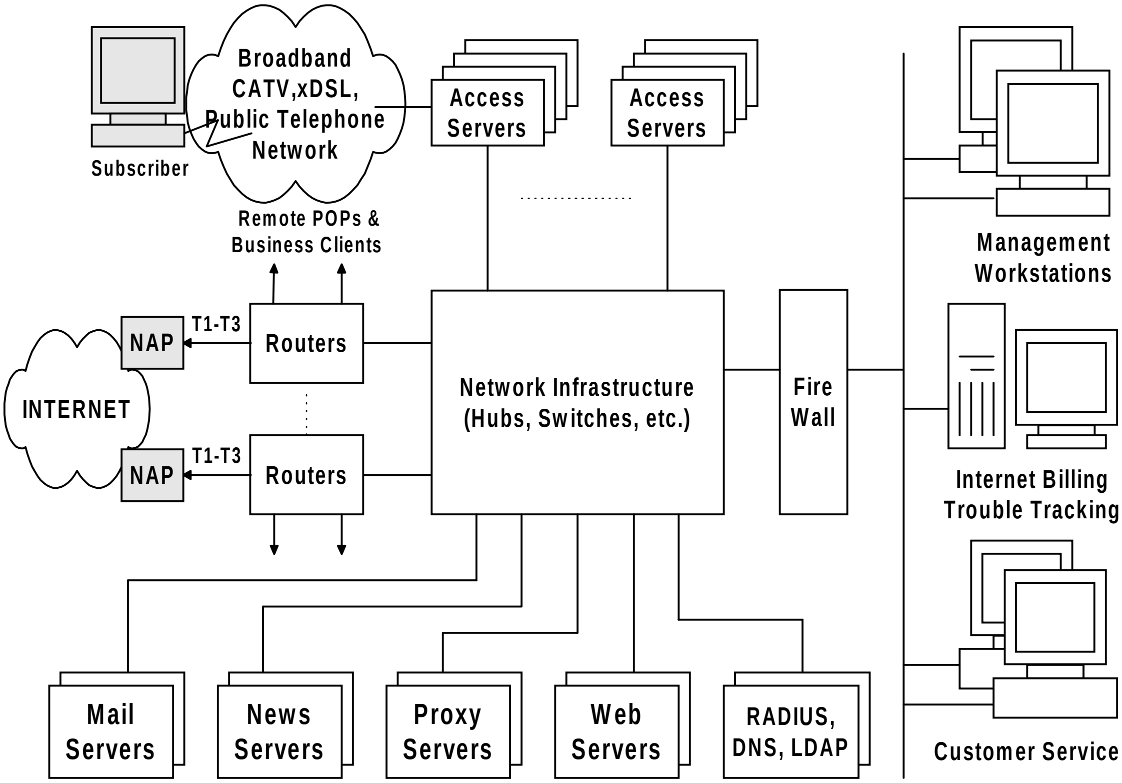

First, identify the applications that will be used to provide the services. The most popular Open Source solutions are listed later in this brief, along with a description of what they do. This is the time to choose a domain name (such as ispconnect.com). This is the name that will uniquely identify an ISP business on the Internet.

Next, decide on how much bandwidth you need. Perhaps even more importantly you'll need to decide how much bandwidth you can afford for your connection to the Internet. Will you need a fractional T1[1] or E1, a full T1 or E1, multiple T1s or E1s, T3, or E3 lines? Upstream Internet access is available from a variety of sources. It can be offered by backbone providers, national providers, or regional providers. In the diagram above they've all been listed under a generic name of Network Access Providers, NAP. In general, the higher up the food chain, the more expensive the connectivity and the better the Internet topology. Interpret this as a major factor in a customer's web site response times. Decide on the equipment required (servers, disk storage arrays, backup equipment, routers, switches, hubs, dial-up access servers, modems). Compaq offers qualified and tested equipment in all these areas, either directly or via our solutions partners.

After deciding Internet upstream connectivity, you'll need to get a block of IP addresses. You'll use these for your own equipment, and potentially offer addresses to business and maybe even consumer customers. Register your ISP's domain name. Allotment of an address pool and even

[1] The T-carrier system, introduced by the Bell System in the U.S. in the 1960s, was the first successful system that supported digitized voice transmission. The original transmission rate (1.544 Mbps) in the T-1 line is in common use today in Internet service provider (ISP) connections to the Internet. Another level, the T-3 line, providing 44.736 Mbps, is also commonly used by ISPs. Another commonly installed service is a fractional T-1 line, which is the rental of some portion of the 24 channels in a T-1 line, with the other channels going unused.
The T-carrier system is entirely digital, using pulse code modulation and time-division multiplexing. This definition is provided by the http://www.whatis.com site.

registration of a domain name can be handled by your upstream provider. Even better, go directly to the address assignment, and the network naming authorities for your country. They are listed below under the IP Addressing, and Domain Name Registration Resources section. In the U.S., it is IANA, and the Internic. Before doing this, establish a Domain Name Service (DNS), and have assigned primary and secondary name servers to provide this service.

The Linux Documentation project is a great source of information on how to set up networking, DNS, create Linux routers, Linux firewalls etc. Information is available in the form of HOW-TO articles.

ISPs also need to begin thinking about security needs – how to protect systems and critical customer data (such as credit card information and customer lists).

As a final point, an ISP solution is still not complete until you also take into account the elimination of single points of failure, and the need for redundant equipment.

There is detailed information at numerous external sources that are referenced in this brief, and Compaq recommends these sources where additional software detail is needed. Information about the software and organizations used in this solution is provided below. They are categorized as Linux Information Resources, Internet Protocol Addressing and Domain Name Registration Resources, and General Open Source Software Resources.

# 3 Linux Information Resources

A large amount of information is available on the worldwide web. The following web sites provide information on Linux, Linux distributors, organizations controlling domain name registration, and IP address assignment:

- http://www.linux.org – Linux Online.
- http://www.li.org – Linux International.
- http://www.tux.org – East coast (U.S.) Linux Users groups.
- http://www.redhat.com –Red Hat Linux distribution.
- http://www.calderasystems.com – Caldera Linux distribution.
- http://www.suse.com – SuSE Linux distribution.
- http://www.pht.com – Pacific HiTech Linux distribution.
- http://www.debian.org – Debian Linux distribution.
- http://www.slackware.org – Slackware Linux distribution.
- http://metalab.unc.edu/LDP – The Linux Documentation project.

# 4 IP Addressing and Domain Name Registration Resources

The following web sites provide IP addressing, and domain name registration resources:

- http://www.iana.org – The Internet Assigned Numbers Authority.
- http://www.internic.net – Assigns network names in top-level domains (i.e. COM, NET, ORG).
- http://www.apnic.net – Asia-Pacific Network Information Center.
- http://www.arin.net – American Registry for Internet Numbers.
- http://www.ripe.net – Reseau IP Europeens.

# 5 General Open Source Software Resources

These applications are typically available in package form with the Linux distributions listed above.

- http://www.opensource.org – Open source software site.
- http://www.gnu.org – The GNU (GNU's Not Unix) project.
- http://www.fsf.org – The Free Software Foundation (GNU).
- http://www.freshmeat.net – Great site for open source solutions.

# 6 Compaq ActiveAnswers Solution Resources

The *Compaq ActiveAnswers* site contains a growing list of documents that address the planning, deployment, and operations of Linux solutions on Compaq platforms.

- http://www.compaq.com/activeanswers

Some examples of documents contained within the *ActiveAnswers* Linux solutions are:

- Installation and Configuration Guide for Linux and Apache Web Server on Intel
- Preventing Unsolicited Bulk Email on Linux
- FrontPage Server Extensions using the Apache Web Server on Linux
- Virtual Web Hosting using the Apache Web Server on Linux

# 7 Open Source Software Commonly Used by ISPs

This section contains a list of web sites and other information sources that are relevant to Open Source solutions typically used by ISP's in providing their services. Most of the applications discussed here are available in package form with the Linux distributions listed above. The goal of this section is to give an overview of each software component. It also provides a definitive source for additional information on that component. The Open Source solutions are grouped under Domain Name Service Software, Worldwide Web Server (WWW) Software, Mail Software, News Software, File Transfer Protocol Software, Firewall Software, System Management Software, Routing Software, Dial-up Authentication Software, Directory Service Software, Chat Servers, Shared File/Print Services Software, Database Software, Development Tools, and finally pointers to Client Side Application Software.

## 7.1 Domain Name Service Software (DNS)

This is the software that provides the domain name services for the Internet. DNS is a critical component for all of the other services an ISP provides (Web, Mail, News, E-Commerce, etc.). The DNS is the way that Internet domain names are located and translated into IP addresses. The DNS is a distributed database, that takes care of mapping machine names (fully qualified domain name, like mail.ispconnect.com) to/from their Internet IP addresses (for example, 192.255.16.1). The lists of domain names and IP addresses are distributed throughout the Internet in a hierarchy of authority. The original top level domains are:

- .com (for commercial organizations, like Compaq, compaq.com)
- .edu (for educational organizations, like Harvard, harvard.edu)
- .gov (for Government organizations, like NASA, nasa.gov)
- .mil (for military organizations, like the U.S. Army, army.mil)
- .net (for networking organizations, like uunet.net, or ispconnect.net)
- .org (for noncommercial organizations, like National Public Radio, npr.org)
- .int (for international organizations, like NATO, nato.int).
- There are also country-wide top-level domains, like .us (for the United States), .jp (for Japan), .de (Germany) etc.

ISP’s typically utilize the .net domain for their own infrastructure (like ispconnect.net) and utilize the .com domain for their business operations (like ispconnect.com). ISPs need to run at least two DNS servers. One will function as a primary (definitive name/address information for the domain). The second will function as a secondary (backup for name/address information, in case the primary is unavailable. These servers also function as the central query location for all Internet domain information. They will resolve information not in their domain space, by querying the appropriate DNS root name servers. It’s a good idea to run these on systems that are not on the same LAN.

### 7.1.1 BIND

http://www.isc.org/

The Internet Software Consortium is in the final stages of becoming a nonprofit corporation dedicated to production-quality software engineering for key Internet standards. They provide a DNS software package known as BIND (the Berkeley Internet Name Domain). It is the most popular software for providing DNS services, and is currently running the DNS services for the DNS root servers (definitive DNS information for the top-level domains). DNS and BIND by Paul Albitz & Cricket Liu (ISBN: 1-56592-236-0) is an excellent source of information.

## 7.2 Web Server Software

These are the servers supporting the Hypertext Transfer Protocol (HTTP). The list below covers both WWW Servers and WWW Proxy/Cache Servers. Also included are useful publishing and log analysis tools.

### 7.2.1 Apache

http://www.apache.org/

The Apache Web Server Project homepage contains versions of the Apache Web Server that you can download and documentation on each version. The January 1999 WWW server site survey by Netcraft found that over 53% of the web sites on the Internet are using Apache (58% if Apache derivatives are included), thus making it more widely used than all other web servers combined.

The Apache project is an effort to develop and maintain an open-source HTTP server for various modern desktop and server operating systems, such as UNIX and Windows NT. The goal of this project is to provide a secure, efficient and extensible server that provides HTTP services in sync with the current HTTP standards.

The Apache Web Server provides some very nice tools supporting virtual web hosting, server side extensions, as well as bandwidth limiting capabilities for the virtual web sites you are hosting. Detailed information on virtual web hosting can be found in the Compaq ActiveAnswer document.

http://www.apacheweek.com

The ApacheWeek site has excellent information on Apache configurations, and modules available for the server. Go to this site for a definitive list of Apache modules. You’ll find features for blocking access, counters, server-side scripting, limiting bandwidth, url rewriting, etc.

Popular Apache modules shipped with Apache can be found at:

http://www.apache.org/docs/mod/index.html

### 7.2.2 FrontPage 98 Server Extensions

http://www.microsoft.com/frontpage

The FrontPage 98 Server Extensions Resource Kit facilitates document publishing on a Linux Web server from remote PC clients running Microsoft FrontPage 98. More detailed information can be found in the *Front Page Server Extensions on the Apache Web Server on Linux* document, located in the *ActiveAnswers* for Apache Web Server on Linux and *ActiveAnswers* for ISP

Infrastructure on Linux solutions (http://www.compaq.com/activeanswers). These extensions are not open source, but add nice publishing capabilities to your open source web server solution.

## 7.3 Web Log Analysis

This section describes web log analysis tools.

### 7.3.1 Analog

http://www.statslab.cam.ac.uk/~sret1/analog/

Analog is a tool for analyzing log files. This site provides a description of Analog, sample reports that can be generated by analog, and access to sites for downloading the latest version of the software.

### 7.3.2 Webalizer

http://www.mrunix.net/webalizer

Webalizer provides highly detailed, easily configurable usage reports in HTML format.

### 7.3.3 Squid Proxy Server

http://squid.nlanr.net/Squid/

Squid is a high-performance proxy/caching server. It supports FTP, Gopher, and HTTP data objects. A proxy server can save you on upstream bandwidth costs. It can also be utilized to implement HTTP screening capabilities for blocking "objectionable" material from customers, or a subset of your customer base. Squid tends to be faster than other proxy/caching servers because it maintains linked objects in RAM, caches DNS lookups, and supports non-blocking I/O, negative caching of objects, and DNS lookups. Squid documentation and FAQs are available from the Squid homepage.

## 7.4 Mail Software

This is the software that supports the SMTP, POP-3, and IMAP-4[2] protocols. Also listed in this section are tools for managing mailing lists, and setting up mail filtering and processing for subscriber vacation messages. Included at the end is a resource for understanding the problems of SPAM mail, and how to deal with them.

### 7.4.1 Sendmail

http://www.sendmail.org

Sendmail is the BSD Mail Transport Agent supporting mail transport by means of TCP/IP using Simple Mail Transfer Protocol (SMTP). Sendmail V8 is the most popular mail transport agent

[2] SMTP (Simple Mail Transfer Protocol) is a TCP/IP protocol used in sending and receiving e-mail. However, since it's limited in its ability to queue messages at the receiving end, it's usually used with one of two other protocols, POP3 or IMAP, that let the user save messages in a server mailbox and download them periodically from the server. In other words, users typically use a program that uses SMTP for sending e-mail and either POP3 or IMAP for receiving messages that have been received for them at their local server. Most mail programs such as Eudora let you specify both an SMTP server and a POP server. On UNIX-based systems, sendmail is the most widely-used SMTP server for e-mail. SMTP usually is implemented to operate over TCP port 25. The details of SMTP are in RFC 821 of the Internet Engineering Task Force (IETF). This definition is provided by the http://www.whatis.com site.

being used on the Internet today. Sendmail has support for mail masquerading (the ability to make all outgoing mail appear to come from one mail domain), and for virtual mail domains (supporting many mail domains on a single system). It also has nice anti-SPAM features for rejecting inbound mail from known spammers, validating domains, keywords and pattern recognition ("make money fast"), and limiting the number of recipients in outgoing/incoming mail.

Note: Sendmail is heavily dependent on DNS. It is in DNS that you define mail exchange records. These records define the system or systems that handle mail for a specific domain.

### 7.4.2 Qmail

http://www.qmail.org

Qmail is considered by many as a replacement for sendmail. It also supports SMTP, and provides performance features not found in sendmail.

### 7.4.3 POP-3

ftp://ftp.qualcomm.com/eudora/servers/unix/popper/

http://eudora.qualcomm.com/free/servers.html

POP (Post Office Protocol) allows single-user hosts to read electronic mail from a server. These sites provide information on the Qualcomm POP server (qpopper).

### 7.4.4 IMAP

http://www.washington.edu/imap/

* University of Washington IMAP server

http://andrew2.andrew.cmu.edu/cyrus/imapd/

* Cyrus IMAP server by Carnegie Melon University

The Internet Message Access Protocol (IMAP) allows a client to access and manipulate electronic mail messages on a server. The current version of the protocol is 4 and is described in RFC 1730.

### 7.4.5 procmail

http://www.procmail.org

Procmail is the mail-processing utility language written by Stephen van den Berg of Germany. Using procmail, you can filter hundreds or thousands of incoming mail messages per day according to a predefined set of rules. Since the procmail language understands details about most UNIX mail transport and delivery agents, it is the tool of choice for writing custom mail filtering scripts. The procmail filtering engine can be invoked by sendmail or by a user's .forward file. Procmail is also the basis of high performance mailing list software like Smartlist.

See procmail(1), procmailrc(5), and procmailsc(5) for more information on using procmail. Extensive examples are provided in procmailex(5).

You can also visit the following Web site to learn more about using procmail:

The Procmail tips site at ( ftp://cs.uta.fi/pub/ssjaaa/pm-tips.html ) contains pointers to documentation, examples, and FAQs, and a link to the FTP site where you can obtain the procmail kit.

### 7.4.6 Fetchmail

http://www.tuxedo.org/~esr/fetchmail/

Fetchmail is a full-featured, robust, well-documented remote-mail retrieval and forwarding utility.

### 7.4.7 Majordomo

ftp://ftp.greatcircle.com/pub/majordomo/

Majordomo is a set of programs that automate operation and maintenance of Internet mailing lists.

### 7.4.8 MAPS

http://maps.vix.com/rbl/

The Mail Abuse Protection System (MAPS) Realtime Blackhole List (RBL) is a system for creating intentional network outages for the purpose of limiting the transmission of known unsolicited bulk e-mail (also known as SPAM). This site also explains the SPAM problem, how to test your mail system for vulnerability to unsolicited bulk e-mail, and how to prevent your mail system from being used for unsolicited bulk emailing. The consequences of SPAM, and how to prevent it are addressed in the Active Answers, Preventing Unsolicited Bulk Email on Linux document.

## 7.5 News Software

This software is used for handling Usenet. Usenet is a worldwide distributed discussion system that consists of “newsgroups” classified hierarchically by subject. Articles" or "messages" are "posted" to these newsgroups, and the articles are then fed to other interconnected computer systems via a wide variety of networks. Architecturally it looks like this:

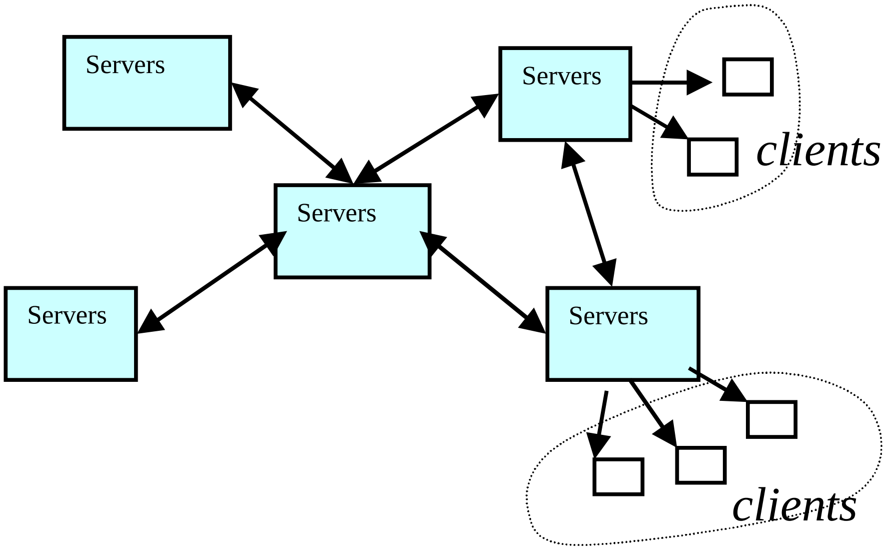

### 7.5.1 INN

http://www.isc.org/inn.html

The Internet Software Consortium homepage for InterNetNews (INN) provides release notes and access to the latest kit.  For more information on newsfeeds and the InterNetNews server, see the Usenet and InterNetNews document by Rich Salz, and the set of FAQs (Frequently Asked Questions).  Both of these are available at the above URL.

For information on how to enable authentication of Usenet group changes using PGPverify, visit the following URL:

ftp://ftp.uu.net/networking/news/misc/pgpcontrol/README.html

### 7.5.2 Diablo

http://www.backplane.com/diablo

DIABLO is a backbone news transit system, designed to replace INND on backbone machines. The transit part of diablo is well established and should work flawlessly. The transit server portion of diablo cannot be run on a machine that needs to accept nntp post commands or newsreading-related nntp commands. You cannot point news reader clients to a host running the transit portion of Diablo.

### 7.5.3 NNTPCache

http://www.nntpcache.org

NNTPCache (efficiently) executes on the local host masquerading as an NNRP news reading server. In fact, what it does is pass certain NNTP commands through to real (remote and possibly local) news-servers based on various pattern matching rules. NNTPCache then takes the output from those servers.  That output is then cached and indexed. The next time information is asked for, or other information which can be logically inferred from the previously collated information, it is sent directly from the cache, without consulting the remote servers.

## 7.6 File Transfer Protocol Software

http://wugate.wustl.edu/ftp/

Washington University, St. Louis, Missouri.

File Transfer Protocol (FTP) is a client/server protocol that allows a user on one computer to transfer files to and from another computer over a TCP/IP network.  An anonymous FTP account on your system allows any remote user to log in to your system using an FTP or anonymous user name.  Once logged in, that user has access to a special directory hierarchy that contains publicly accessible files and to transfer those files to another system using FTP.

## 7.7 Firewall and system access control software

This section covers IP Masquerading, Packet Filtering and Proxy Firewalls, and additional system access control software.

Masquerading is a nice facility for translating network traffic from one network to another by hiding one or more addresses behind a single visible address. It accomplishes this by using a form of network address translation (NAT).

There are two basic types of firewalls, IP or Packet Filtering Firewalls, and Proxy Firewalls. Packet filtering firewalls block all but selected network traffic. All traffic through a TCP/IP network is sent in the form of packets. The start of each packet says where it's going, where it came from, the type of the packet, and other administrative details. This start of the packet is called the header. The rest of the packet contains the actual data being transmitted and is called the body.

A packet filter is a piece of software that looks at the header of packets as they pass through, and decides the fate of the entire packet. It might decide to deny the packet (ie. discard the packet as if it had never received it), accept the packet (ie., let the packet go through), or reject the packet (like deny, but tell the source of the packet that it has done so). Under Linux, packet filtering is built into the kernel.

Proxy firewalls make the network connections on your behalf. They are designed to control the flow of packets based on the source, destination, port and packet type information contained in each packet.

### 7.7.1 Linux with IP Masquerading

http://www.linux.org

http://ipmasq.cjb.net

IP masquerading can be implemented with Linux, and standard networking software. IP Masquerading is a form of network address translation that many routers already support. The idea behind this implementation is that people running Linux can install the IP masquerading features and get the features of the high priced routers and NAT boxes without paying the high prices.

IP masquerading lets you use a single Internet-connected computer running Linux with a real IP address as a gateway for non-connected machines with unregistered (not routed out to the internet, like the class A network 10 ) IP addresses. The Linux box with a real address handles mapping packets from your intranet out to the Internet, and when responses come back, it maps them back to your intranet.

### 7.7.2 Ipfilter

http://coombs.anu.edu.au/~avalon/ip-filter.html

IP Filter is a TCP/IP packet filter, suitable for use in a firewall environment. It can either be used as a loadable kernel module or incorporated into your Linux kernel. IP filter can explicitly deny/permit any packet from passing through, distinguish between various interfaces, filter by IP networks or hosts, selectively filter any IP protocol, and act as a Network Address Translator (NAT) .

If you are using a pre-2.2 linux kernel, you’ll need to use Ipfilter. However if you are using a 2.2 or later linux kernel, you should use Ipchains.

### 7.7.3 Ipchains

http://www.rustcorp.com/linux/ipchains

Linux ipchains is a rewrite of the IP filter code. If you are using a 2.2 or later linux kernel, you should use ipchains.

### 7.7.4 TIS

ftp://ftp.tis.com

http://www.ssc.com/lj/issue25/1204.html

The Trusted Information System toolkit. The Linux Journal has a great article on creating a firewall using the TIS toolkit.

### 7.7.5 TCP Wrapper

ftp://ftp.win.tue.nl/pub/security/

TCP Wrapper intercepts an incoming network connection and verifies that the connection is allowed before passing the connection to the network daemon. TCP Wrapper is configured through the /etc/hosts.allow file. The FTP archive listed above is for the Mathematics and Computing Science Department at Eindhoven University of Technology (the Netherlands) and contains TCP Wrapper kits.

## 7.8 System Security Sites

Information on system security is available from the sites described in this section.

### 7.8.1 CERT

http://www.cert.org/

The Computer Emergency Response Team (CERT) is a clearinghouse for security-related events that occur in the Internet community. If you are an administrator, subscribe to the CERT mailing list and frequently check the CERT advisories.

According to its charter, CERT works with the Internet community to facilitate the community's response to security events involving hosts, takes proactive steps to improve the community's awareness of security issues, and conducts research aimed at improving the security of existing systems. CERT services include a 24-hour hotline for responding to security incidents, product vulnerability assistance, and technical documentation and tutorials.

### 7.8.2 CIAC

http://ciac.llnl.gov/

The Computer Incident Advisory Capability (CIAC) site is maintained by the U.S. Department of Energy. Their Web site offers computer security information, as well as workshops, consulting, and security incident-handling information.

### 7.8.3 CSRC

http://csrc.ncsl.nist.gov/

The Computer Security Resource Clearinghouse (CSRC) is a U.S. government archive on security information and contacts maintained by the National Institute on Standards and Technology (NIST).

### 7.8.4 BugTRAQ

http://www.geek-girl.com/bugtraq

BugTRAQ is a full-disclosure UNIX security mailing list. This is a detailed discussion of UNIX security holes: what they are, how to exploit, and what to do to fix them.

## 7.9 System Management Software

### 7.9.1 VNC

http://www.uk.research.att.com/vnc

Virtual Network Computing (AT&T labs Cambridge) develops remote access software, that allows you to bring up a remote display of the system that is being managed. VNC is a client/server solution. It is composed of vncserver, vncviewer, and a java viewer for browser access. Unlike X based tools, no state is stored in the client side viewer. Linux/Unix/NT/W95/W98/Macintosh systems can be remotely managed with this software.

## 7.10 Routing Software

Gated

http://www.gated.org

Merit GateD Consortium. Routed and Gated are the most popular software for routing. You can use your Linux box as a router. With the Linux 2.2 kernel you can optimize the system to be used as a router. There is also support for routing with Linux for T1-interface network cards.

## 7.11 Dial-up Authentication Software

### 7.11.1 RADIUS Server (Basic Merit AAA Server)

http://www.merit.edu/aaa/

The Basic Merit AAA Server supports standard Remote Authentication Dial In User Service (RADIUS) authentication. The Merit server extends standard RADIUS authentication to include features like distributed authentication and authorization schemes. The Basic Merit AAA Server can also act as a proxy server that relays authentication and authorization requests to other RADIUS servers, and receives and relays the responses. The Basic Merit AAA Server is not Open Source, but it is freely available and the source is provided. You can find detailed information on Merit Radius solutions at the Compaq ActiveAnswer web site.

## 7.12 Directory Service Software

Directory Services provide a way to centrally administer subscriber information in a highly scalable manner. The most common way to do this is through the use of a directory server that supports the LDAP (Lightweight Directory Access Protocol). The Lightweight Directory Access Protocol is an Internet standard directory service protocol that runs over TCP/IP. It can be used to provide a standalone directory service or to provide lightweight access to the X.500 directory. You can use these servers to centralize authentication information, message store locations, vacation messages, access privileges etc. Perl integration with LDAP is available through the PerlLDAP libraries.

### 7.12.1 LDAP

http://www.umich.edu/~dirsvcs/ldap/

One of the most popular LDAP servers is provided by the University of Michigan. Documentation, patches, kits and archives for the University of Michigan LDAP server are available at this Web site.

### 7.12.2 OpenLDAP

http://www.openldap.org

The OpenLDAP Project is a collaborative effort to provide a robust, commercial-grade, fully-featured and open source suite of LDAP applications and development tools. The project is managed by a worldwide community of volunteers who use the Internet to communicate, plan, and develop the OpenLDAP suite and its related documentation.

## 7.13 Chat Software

This section describes chat software.

### 7.13.1 IRC (Internet Relay Chat)

http://www.irchelp.org/

Internet Relay Chat (IRC) allows users to communicate with each other in real time across a network of Internet servers. IRC is an internet-based system that allows people from all over the world to have live real-time textual "conversations." It is based around a network of interconnected servers for which the user runs an IRC client program to connect. Each user on an IRC network must have a unique nickname, and can join one or more channels. They can also send and receive private messages and even transfer files. The help archive URL above provides general information and pointers to many FAQs and clients and servers.

## 7.14 Shared File and Print Services

This section describes file and print servers.

### 7.14.1 Samba File and Print Server

http://samba.anu.edu.au/samba/

news:comp.protocols.smb

The Samba File and Print Server for Windows provides file and print services to SMB clients, such as Windows for Workgroups, Windows NT, or LanManager. It also provides Netbios name serving and browsing support. At the Samba site, you can obtain documentation on the Samba server, report problems and download the latest software. Visit the comp.protocols.smb news group for information on the Samba server protocol.

## 7.15 Development Tools

This section describes development tools commonly used by ISP's.

### 7.15.1 Perl

http://www.perl.org

The Practical Extraction and Report Language (Perl) is an interpreted language distributed over Usenet. See this web site for more information on Perl, including FAQs, documentation, newsgroups and software updates. Perl is a favorite tool of system administrators, and is often referred to as the "duct tape of the Internet".

### 7.15.2 CPAN

http://www.cpan.org

Comprehensive Perl Archive Network contains resources for using Perl as building blocks, and integration tools for your web, mail, news, ldap, etc. solutions.

### 7.15.3 GNU development tools

http://www.gnu.org

The GNU site lists a variety of development software and tools at its site above. It is at this location you can find the GNU C and C++ compilers, version control software, and build utilities. If you are porting a software package, and need to build the source, there is a good chance that it will need to be built using the GNU tools.

### 7.15.4 Tcl

http://www.scriptics.com/

ftp://ftp.scriptics.com/pub/tcl/

Tool Command Language (Tcl) is a string-processing language for issuing commands to interactive programs. The Web site at Scriptics Corporation provides information on topics related to Tcl. Their FTP site contains a collection of Tcl kits that are available for downloading.

### 7.15.5 expect

http://expect.nist.gov/

Expect is a tool for automating and testing interactive applications, such as telnet, FTP, passwd, fsck, rlogin, tip, to name a few. Exploring Expect: A Tcl-Based Toolkit for Automating Interactive Applications (ISBN 1-56592-090-2), written by Don Libes and published by O'Reilly & Associates, is an excellent source of information. The expect homepage provides access to FAQs, examples, contributed scripts and software.

### 7.15.6 Tcl_cgi

http://ruulst.let.ruu.nl:2000/tcl-cgi.html

Tcl_cgi provides a secure WWW interface to Tcl applications. The Tcl_cgi homepage provides pointers to examples of applications that use Tcl_cgi and contains a link to download the source code.

### 7.15.7 TclTk

http://www.scriptics.com/resource/download/tcltk/

TclTk is a programming toolkit developed by John Ousterhout at the University of California, Berkeley. The TclTk toolkit is similar to other GUI toolkits (such as Xlib, Xview, and Motif). Unlike those toolkits, TclTk does not require that you use C or C++ in order to manipulate the widgets, and you can build useful applications very rapidly after you gain some expertise.

The TclTk web site at Scriptics Corporation provides pointers to TclTk kits. The Tcl WWW Info page at SCO provides FAQs, introductory and reference material, and sample applications.

### 7.15.8 TclX

http://www.NeoSoft.com/tclx/

TclX (Extended Tcl), a superset of standard Tcl, was created by Karl Lehenbauer (karl@neosoft.com) and Mark Diekhans (markd@grizzly.com), and can be freely distributed for any use without license or fee. Available over the Internet since 1989, TclX adds capabilities to Tcl and is the source of many of the capabilities of the baseline Tcl release, including arrays, files, sockets, file events, and date and time handling. TclX has three basic functional areas: a set of new commands, a Tcl shell (UNIX shell-style command line and interactive environment), and a user-extensible library of useful Tcl procedures.

The TclX homepage provides pointers to documentation and contains a link to download the source code.

## 7.16 Client Application Software

This section describes client application software.

### 7.16.1 Applications

http://www.tucows.com/

The Ultimate Collection of Winsock Software web site contains access to numerous Internet applications for personal computer users.

### 7.16.2 Lynx

ftp://ftp2.cc.ukans.edu/pub/lynx/

Lynx is a fully featured Web browser for users connected to a system via cursor-addressable, character-cell terminals or terminal emulators. Lynx is a product of the Distributed Computing Group within Academic Computing Services of the University of Kansas.

### 7.16.3 Pine and Pico

http://www.washington.edu/pine/

Pine is an IMAP e-mail client for terminals or terminal emulators. Pine is easy to learn; command choices are presented at the bottom of each screen. Pico is the editor that Pine uses. This web site at the University of Washington provides complete documentation, FAQs, and access to the software.

### 7.16.4 TIN

http://www.tin.org

TIN is a full-screen news reader for terminals or terminal emulators on UNIX systems. The TIN home page allows you to download the latest version of TIN, access TIN mailing lists, and submit bug reports.

## 7.17 Database Software

http://www.mysql.org

Multi-user, multi-threaded SQL database server. An excellent tool for centrally managing your solutions. There are also nice modules for the Apache Web Server that allows you to utilize your SQL database.

http://www.postgresql.org

PostgreSQL is an object-relational DBMS. It supports most SQL constructs, including subselects, transactions, and user-defined types and functions.

# 8 Summary

This is just the tip of the Open Source Iceberg. Every tool you need is either available, or can be constructed from an existing Open Source tool. Open Source gives you, the ISP, control over your solution.

# 9 Glossary

**Anonymous FTP**

Using the Internet's File Transfer Protocol (FTP), anonymous FTP is a method for giving users access to files so that they don't need to identify themselves to the server. Using an FTP program or the FTP command interface, the user enters "anonymous" as a user ID. Usually, the password is defaulted or furnished by the FTP server. Anonymous FTP is a common way to get access to a server in order to view or download files that are publicly available.

**.forward file**

On a Unix system, this file (unique to each user and located in their home directory) is used to reroute mail or perform some processing on that mail message.

**IMAP**

Internet Message Access Protocol. It is a protocol for retrieving e-mail messages. The latest version, IMAP4, is similar to POP3 but supports some additional features. You can store and retrieve messages from folders that you maintain on the mail server. With IMAP4, you can search through your e-mail messages for keywords while the messages are still on mail server. You can then choose which messages to download to your machine.

**INN**

The InterNetNews package.

**Innd**

The main process that implements the network news transport protocol (NNTP) for InterNetNews. The ‘d’ notation at the end of ‘inn’ denotes daemon, which means it is a standalone process a process that runs in the background and performs a specified operation at predefined times or in response to certain events.

**Internic**

The global registrar for the top-level domains of .com, .net, and org.

**NNTP**

Network News Transport Protocol. This is the protocol used for Usenet News. It consists of a set of client-server and server-server commands that are utilized to send/receive usenet news messages across the Internet. For example, the NNTP POST command is used by client programs wishing to post a Usenet news article to a particular discussion group.

**POP**

Short for Post Office Protocol. It is a protocol used to retrieve e-mail from a mail server. Most e-mail applications (sometimes called an e-mail client) use the POP protocol.

**Procmail**

Autonomous mail processor. Usually invoked via the .forward file mechanism, or from within the mail server.

**procmailrc**

procmail resource file, which contains a mixture of environment variable assignments, and recipes for processing mail.

### Procmailsc

procmail weighted scoring technique. In addition to the traditional true or false you can specify in a recipe, you can use a weighted scoring technique to decide if a certain recipe matches or not.

### Procmailex

Example procmailrc files.

### RFC

Request for comments. More information on RFC's, can be found at

http://www.rfc-editor.org/

### Root Name Servers

The root name servers know where name servers authoritative for all the top-level domains are located. Given a query about any domain name, the root name servers can provide the names and addresses of the name servers authoritative for the top-level domain of that domain name.

### Server side extensions

In the context of the Worldwide Web client-server model, server side extensions are extensions to the web server that allow for additional processing/functionality on the server side.

### SMTP

Short for Simple Mail Transfer Protocol, a protocol for sending e-mail messages between servers. Most e-mail systems that send mail over the Internet use SMTP to send messages from one server to another.

### SPAM

A term used to describe unsolicited junk e-mail. They typically consist of commercial advertising, get rich quick schemes, and the sale of dubious products.

### Unsolicited Bulk E-mail (UBE)

Unsolicited Bulk E-mail is the targeting of large numbers of Internet mail users with unsolicited direct mail messages. Commonly referred to as SPAM.

### X.500

An ISO and ITU standard that defines how global directories should be structured. X.500 directories are hierarchical with different levels for each category of information, such as country, state, and city.