ANEXO F. SERVIDORES LINUX

TABLA DE CONTENIDO

LISTA DE TABLAS

ANEXO F

F. SERVIDORES LINUX

Antes de iniciar este anexo con la configuración directa de los servidores en el sistema operativo linux es necesario tener un conocimiento general acerca de este s istema, por lo tanto se va a hacer una introducción mirando su historia, evolución y concepto para después dar la caracterización y configuración de los servidores que soportan los diversos servicios que provee la *RTPSTT*.

F.1 CONCEPTO

Linux es un sistema operativo multiusuario y multitarea real compatible con UNIX. Presenta una gestión avanzada de memoria virtual y recursos.

Linux (como UNIX) es un SO seguro, con contraseñas, usuarios y grupos de usuarios que permite una limitación estricta del acceso y el uso del sistema. Está absolutamente orientado al trabajo en redes, especialmente TCP/IP y acceso a Internet SLIP/PPP. Su estabilidad y su resistencia a cargas de trabajo hacen que sea cada vez más utilizado en entornos empresariales.

Linux permite acceder, ya sea en modo de lectura y escritura, ya sea solo lectura a gran cantidad de sistemas de ficheros distintos: FAT16, FAT32, VFAT, OS/2, Mac, NTFS y muchos sistemas UNIX diferentes además emplea su propio sistema de ficheros (ext2) para su instalación.

El núcleo (kernel) de Linux es recompilable y totalmente configurable por el usuario, de manera que se pueden hacer desde configuraciones con total soporte de elementos multimedia como tarjetas de sonido, radio, TV... hasta configuraciones para hard ware modesto o que no precisan gastar recursos en esos elementos.

El sistema se distribuye bajo licencia GNU lo cual quiere decir que se rige por las normas del software libre, concretamente de la licencia GPL.

Presenta dos características que lo diferencian del resto de los sistemas que podemos encontrar en el mercado, la primera, es que es libre, esto significa que no hay que pagar ningún tipo de licencia a ninguna casa desarrolladora de software por el uso del mismo, la segunda, es que el sistema viene acompañado del código fuente. El sistema lo forman el Núcleo del sistema mas un gran número de programas / librerías que hacen posible su utilización.

Día a día, más y más programas / aplicaciones están disponibles para este sistema, y la calidad de los mismos aumenta de versión a versión. La gran mayoría de los mismos vienen acompañados del código fuente y se distribuyen gratuitamente bajo los tér minos de licencia de la GNU Public License.

Las plataformas en las que en un principio se puede utilizar Linux son 386 -, 486-. Pentium, Pentium Pro, Pentium II, Amiga y Atari, también existen versiones para su utilización en otras plataformas, como Alpha , ARM, MIPS, PowerPC y SPARC

F.2 HISTORIA DEL SISTEMA OPERATIVO LINUX

La idea de crear un sistema GNU (General Public License) y el código fuente disponible gratuitamente, surgió en el año de 1991 cuando Linus Trovalds estudiaba la carrera de ciencias informáticas. Trovalds se había interesado en Minix, el único sistema Unix disponible en aquel entonces. Este sistema gratuito había sido creado por Andrew Tanenbaun con el propósito de facilitar a los alumnos de la universidad el estudio y el diseño de sistemas operativos. Minix era un UNIX más, tanto en apariencia como en el kernel, pero distaba mucho de ser comparable a uno de los grandes. A partir de ese momento Trovalds decidió crear un sistema que excediera los estándares de Minix, poniendo en marcha el proyecto Linux.

Torvalds tomó sus primeras clases de C y UNIX en 1990 y en poco tiempo empezó a utilizar el sistema operativo Minix en su nuevo 386. Linux evolucionó del simple

programa **Hello World!** a una terminal. Durante mucho tiempo Torvalds trabajó sólo, hasta la mañana del 3 de Julio de 1991 cuando pidió ayuda a través de Internet. Al principio fueron unos pocos los que le apoyaron, pero al poco tiempo muchos otros internautas se unieron al proyecto. En uno de los primeros e -mails enviado por Torvalds refiriéndose a Linux, informaba sobre su proyecto como si fuera un hobby, nada tan grande no comparable a GNU.

Torvalds se encontró con muchos problemas a lo largo de la programación del kernel (núcleo del sistema). Pero Linux empezó a disponer de con troladores para los dispositivos y un funcionamiento correcto del disco el 3 de Julio, unas horas después de enviar su primer e-mail informando sobre su proyecto. Dos meses más tarde Linux empezaba a funcionar y el código fuente de la primera versión (v0.0 1) ya estaba disponible. La versión 0.01 incluía la bash shell 1.08 y gcc 1.40.

Linux y sus nuevas versiones. También empezó a recibir una avalancha de preguntas sobre su nuevo sistema operativo gratuito. Una de las preguntas más comunes era si Linux se podía portar a otras arquitecturas, cosa que él respondía que no, ya que estaba hecho en gran parte en C y utilizaba 386 MMU. También informó de los dispositivos y programas que Linux podía utilizar, que por aquellos momentos se trataba de gcc, bash shell y la mayoría de utilidades GNU. Uno de los problemas se debía a que los disquetes aún no funcionaban, pero Linux ya empezaba a superar a Minix en algunos aspectos, de tal forma que el proyecto de Trovalds marchaba por buen camino.

Torvalds anunció la versión 0.02 de Linux el 5 de Octubre de 1991. Linux v0.02 ya incorporaba archivos binarios y se podía ejecutar las bash shell, gcc, GNU -make, GNU-sed, compress, etc. Quien estuviera interesado en el código fuente del kernel, algunos binarios (como bash, gcc, etc.) y unos cuantos archivos de ayuda, podían descargarlo de nic.funet.fi.

Linux había progresado de forma considerable en muy poco tiempo, pero aún quedaban muchos arreglos por hacer. Aún no podía funcionar por sí solo, ya que necesitaba el uso de Minix-386, por lo que aún no podía considerarse como un sistema operativo. También necesitaba el uso de un disco duro AT -compatible (IDE funcionaba) y tarjetas

EGA/VGA. Durante el proceso de desarrollo Torvalds comenzó a recibir e -mails con bugs que los usuarios habían encontrado en Linux.

La versión 0.03 pronto apareció y parecía funcionar sin problema alguno. Durante esas semanas el proyecto sufrió un avance muy importante, hasta llegar a la versión 0.11, donde algunos usuarios confirmaban que ya se podía comparar con Minix-386 e incluso indicaban que resultaba superior en algunos aspectos. La versión 0.11 aún no soportaba dispositivos SCSI y tampoco se podía utilizar init/login, de tal forma que se accedía al sistema como usuario root (superusuario, es el usuario que administra el sistema) directamente. La versión 0.11 necesitaba como mínimo 2MB para funcionar pero sin poder compilar y 4MB si se pensaba utilizar binarios GNU.

Durante la Navidad de 1991 apareció la versión 0.12 y Torvalds la definió como un sistema "divertido" de hackear, utilizable en 386/486, sistema libre y mejor que el Minix en muchos aspectos. Linus Torvalds fue recibiendo e -mails de personas que consiguieron hacer funcionar el kernel de Linux en sus ordenadores. El código fuente de Linux, además de sus utilidades ya estaba disponible en otros servidores FTP.

Se creó una nueva lista de correo sólo para Linux, en la que todo usuario interesado en este proyecto podía intercambiar información, ideas, noticias, etc. con el resto de los usuarios. Con v0.12 ya se podía hablar de Linux como un sistema operativo ya que no requería el uso de Minix para configurarlo. Con esta nueva versión, muchos dispositivos de hardware nuevos funcionaban correctamente y no hubo necesidad de implementar ningún patch (pequeño programa escrito normalmente en C que arregla o "tapa" algún bug encontrado) por mucho tiempo. Linux v0-12 tenía las siguientes características:

- Intercambio de data/código entre dos procesos diferentes.
- Mejora en el uso de los disquettes.
- Correcciones de bugs.
- Utilidades: mkfs, fdisk y fsck.
- Teclados US/German/French/Finish
- Com1 y Com2 funcionaban correctamente.
- Init/login finalmente funcionaban.

En muy poco tiempo apareció la siguiente versión de Linux, pero no se denominó v0.13, Torvalds pensó que Linux se merecía estar en los 0.90´s y decidió llamarlo v0.95, pero

contenía gran cantidad de bugs. Todos estos bugs se convirtieron en patchs y gracias a ellos salió la versión v0.96, mucho más estable.

Linux se convirtió en un sistema mucho más fác il de instalar y configurar, y empezó a coger fama en todo el mundo. Al tener en muy poco tiempo miles de usuarios, las nuevas versiones de Linux salían casi semanalmente. En el presente hay millones de usuarios y gracias a ellos y a sus aportaciones, Linu x crece sin respiro alguno. La última versión del Kernel es la 2.4.0 test 4.

Como todos los sistemas operativos, Linux también dispone de un logotipo. Torvalds decidió que la imagen que representaría a Linux sería la de un pingüino. En casi todas las páginas web relacionadas con Linux se puede hallar el logotipo y ya ha sido aceptado como representante de este sistema. No se sabe con certeza cuándo diseñó el logotipo de su sistema, pero se dice que fue al principio del proyecto Linux.

Linux había nacido para ser un sistema operativo del tipo Posix (variante de Unix), totalmente gratuito para el usuario y con libre acceso al código fuente. Estas tres ideas fueron las que lo han convertido al sistema con mejor rendimiento, más fiable, veloz y con más desarrolladores del mundo. Pronto se ha colocado cerca de los grandes sistemas operativos como UNIX en el ámbito de servidores de comunicaciones, especialmente utilizado en empresas proveedoras de acceso a Internet.

Las versiones más recientes de Linux ofrece n la posibilidad de convertir un ordenador personal en una potente estación de trabajo. Puede funcionar como estación de trabajo personal dando la posibilidad de acceder a las prestaciones que ofrece UNIX y cualquier otro sistema operativo. Además, puede c onfigurarse para funcionar como estación de desarrollo, proveedores de acceso a Internet, y muchas otras opciones.

También es interesante convertir un ordenador personal en una estación de desarrollo. Linux dispone de los siguientes lenguajes de programac ión gratuitos: GNU C, GNU C++, GNU Fortran 77, ADA, Pascal, Perl, Modula2 and 3, TCL/Tk, Scheme and Small - Talk/X. Todos ellos vienen con extensas librerías de código fuente.
Muchos proveedores de acceso a Internet emplean Linux como sistema operativo para ofrecer soporte a sus usuarios. Al instalar Linux, también se instalan los siguientes

paquetes, World Wide Web, usenet news, email, ftp, etc. Éstos pueden ser configurados para dar ofrecer a los usuarios los servicios disponibles en Internet.

Linux como sistema operativo gratuito posee unas características que le hacen único. Algunas de las más importantes son: multitarea, memoria virtual, los drivers TCP/IP más rápidos del mundo, librerías compartidas, multiusuario (cientos de usuarios pueden utilizar un ordenador a la vez, ya estén en una ed, Internet, terminales, etc.). Linux trabaja en modo de protección (al contrario de Microsoft Windows) y soporta multitarea a 32 y 64 bits.

Además posee capacidades avanzadas para la interconexión de ordenadores en r edes ya que para crear Linux se hubo de utilizar Internet. El desarrollo de software de interconexión de ordenadores se empezó a implementar en las primeras versiones de Linux y desde entonces ha ido evolucionando a gran velocidad, y más ahora con la gran aceptación de la red.

No hay duda que Linux es uno de los sistemas operativos con más posibilidades, y es el único que se actualiza día a día.

F.3 CARACTERISTICAS DEL SISTEMA OPERATIVO LINUX

Aquí se tiene una lista bastante completa con las caracterísi cas de LINUX

- Multitarea: La palabra multitarea describe la habilidad de ejecutar varios programas al mismo tiempo.
  LINUX utiliza la llamada multitarea preventiva, la cual asegura que todos los programas que se están utilizando en un momento dado serán e jecutados, siendo el sistema operativo el encargado de ceder tiempo de microprocesador a cada programa.
- Multiusuario: Muchos usuarios usando la misma máquina al mismo tiempo.
- Multiplataforma: Las plataformas en las que en un principio se puede utilizar Linux son 386-, 486-. Pentium, Pentium Pro, Pentium II, Amiga y Atari, también existen versiones para su utilización en otras plataformas, como Alpha, ARM,MIPS, PowerPC y SPARC.

- Multiprocesador: Soporte para sistemas con más de un procesador, está disponible para Intel y SPARC.
- Funciona en modo protegido 386.
- Protección de la memoria entre procesos, de manera que uno de ellos no pueda colgar el sistema.
- Carga de ejecutables por demanda: Linux sólo lee del disco aquellas partes de un programa que están siendo usadas actualmente.
- Política de copia en escritura para la compartición de páginas entre ejecutables: Esto significa que varios procesos pueden usar la misma zona de memoria para ejecutarse. Cuando alguno intenta escribir en esa memoria, la página (4Kb de memoria) se copia a otro lugar. Esta política de copia en escritura tiene dos beneficios: aumenta la velocidad y reduce el uso de memoria.
- Memoria virtual usando paginación (sin intercambio de procesos completos) a disco: A una partición o un archivo en el sistema de archivos, o ambos, con la posibilidad de añadir más áreas de intercambio sobre la marcha. Un total de 16 zonas de intercambio de 128Mb de tamaño máximo pueden ser usadas en un momento dado con un límite teórico de 2Gb para intercambio. Este límite se puede aumentar fácilmente con el cambio de unas cuantas líneas en el código fuente.
- La memoria se gestiona como un recurso unificado para los programas de usuario y para el caché de disco, de tal forma que toda la memoria libre puede ser usada para caché y ésta puede a su vez ser reducida cuando se ejecuten grandes programas.
- Librerías compartidas de carga dinámica (DLL's) y librerías estáticas.
- Se realizan volcados de estado (core dumps) para posibilitar los análisis post-mortem, permitiendo el uso de depuradores sobre los programas no sólo en ejecución sino también tras abortar éstos por cualquier motivo.
- Compatible con POSIX, System V y BSD a nivel fuente.
- Emulación de iBCS2, casi completamente compatible con SCO, SVR3 y SVR4 a nivel binario.
- Todo el código fuente está disponible, incluyendo el núcleo completo y todos los drivers, las herramientas de desarrollo y todos los programas de usuario; además todo ello se puede distribuir libremente. Hay algunos programas comerciales que

están siendo ofrecidos para Linux actualmente sin código fuente, pero todo lo que ha sido gratuito sigue siendo gratuito.

- Control de tareas POSIX.
- Pseudo-terminales (pty's).
- Emulación de 387 en el núcleo, de tal forma que los programas no tengan que hacer su propia emulación matemática. Cualquier máquina que ejecute Linux parecerá dotada de coprocesador matemático. Por supuesto, si el ordenador ya tiene una FPU (Unidad de Punto Flotante), esta será usada en lugar de la emulación, pudiendo incluso compilar tu propio kernel sin la emulación matemática y conseguir un pequeño ahorro de memoria.
- Soporte para muchos teclados nacionales o adaptados y es bastante fácil añadir nuevos dinámicamente.
- Consolas virtuales múltiples: varias sesiones de login a través de la consola entre las que se puede cambiar con las combinaciones adecuadas de teclas (totalmente independiente del hardware de video). Se crean dinámicamente y puedes tener hasta 64.
- Soporte para varios sistemas de archivo comunes, incluyendo minix-1, Xenix y todos los sistemas de archivo típicos de System V, y tiene un avanzado sistema de archivos propio con una capacidad de hasta 4 Tb y nombres de archivos de hasta 255 caracteres de longitud.
- Acceso transparente a particiones MS-DOS (o a particiones OS/2 FAT) mediante un sistema de archivos especial: no es necesario ningún comando especial para usar la partición MS-DOS, esta parece un sistema de archivos normal de Unix (excepto por algunas restricciones en los nombres de archivo, permisos, y esas cosas). Las particiones comprimidas de MS-DOS 6 no son accesibles en este momento, y no se espera que lo sean en el futuro. El soporte para VFAT (WNT, Windows 95) ha sido añadido al núcleo de desarrollo y estará en la próxima versión estable.
- Un sistema de archivos especial llamado UMSDOS que permite que Linux sea instalado en un sistema de archivos DOS.
- Soporte en sólo lectura de HPFS-2 del OS/2 2.1
- Sistema de archivos de CD-ROM que lee todos los formatos estándar de CD-ROM.
- TCP/IP, incluyendo ftp, telnet, NFS, etc.

- Appletalk.
- Software cliente y servidor Netware.
- Lan Manager / Windows Native (SMB), software cliente y servidor.
- Diversos protocolos de red incluidos en el kernel: TCP, IPv4, IPv6, AX.25, X.25, IPX, DDP, Netrom, etc.

F.4 SERVIDORES DE LA RTPSTT

Para el funcionamiento de la red de Telemedicina es necesario instalar varios tipos de servidores como son.

- Servidor de acceso remoto, el cual permite la conexión telefónica de los usuarios remotos localizados en aquellos puntos donde no existe el acceso a Internet.
- Servidor Web, el cual permite el manejo de las paginas que servirán como interfaces de usuario para los servicios de Telemedicina.
- Servidor proxy el cual permite la conexión a Internet de los usuarios de la red.
- Servidor DNS que determina el dominio sobre el cual operará la red de Telemedicina.

Y servidores para la prestación de servicios adicionales como son:

- Servidor FTP.
- Servidor de correo electrónico.

F.4.1 SERVIDOR PROXY

Squid es un servidor *proxy* que permite utilizar una sola conexión a Internet para todas las estaciones de los centros de acceso para la navegación por Internet.

El servidor *proxy* además almacena en el disco duro del servidor las páginas más visitadas desde las estaciones de tal manera que se realiza un ahorro significativo del ancho de banda del enlace del centro de acceso cuando se solicita la página nuevamente desde la misma u otra estación. Squid verifica si la página ha cambiado, y de ser así, vuelve a almacenarla localmente. Es posible configurar Squid para definir cuanto tiempo pueden estar almacenadas las páginas en el servidor.
A continuación se muestra un instructivo para su instalación y configuración:

1-. Obtención del programa.

Squid se puede obtener en el CD 2 (Binary CD) de la distribución de RedHat 7 en RPM o copiar la versión que utilizamos en Internet.

También se puede obtener desde la distribución de los programas fuentes. Las fuentes pueden obtenerse de http://www.squid-cache.org/.

2-. Instalación.

2.1-. Desde RPM.

Después de copiar el programa en RPM ejecute:
rpm -iv squid-2.3.STABLE4-1.i386.rpm
Con esto quedará instalado Squid en el sistema. Proceda con la configuración en el apartado 3.

2.2-. Desde las fuentes.

Desde las fuentes, después de obtener el paquete se debe ejecutar:
tar zxvf squid-2-3-STABLE3-src.tgz
cd squid-2.3.STABLE3
./configure
make
make install
En este caso el programa Squid quedará instalado en /usr/local/squid

3-. Configuración de Squid.

Para configurar Squid es necesario modificar el archivo: /etc/squid/squid.conf para la distribución de RPM o /usr/local/squid/etc/squid.conf para las fuentes.

Si realizó la instalación desde RPM teclee:

 # cd /etc/squid/

de lo contrario:

 # cd /usr/local/squid/etc
 # pico squid.conf

El archivo original de configuración trae algunos parámetros con el signo # al comienzo de la línea por lo que es importante quitar el símbolo para activar el parámetro correspondiente. Los archivos de configuración se presentan a continuación: se encuen tra localizado en /etc/squid/squid.conf

```
# WELCOME TO SQUID 2
# ------------------
#
# This is the default Squid configuration file. You may wish
# to look at the Squid home page (http://squid.nlanr.net/)
# for the FAQ and other documentation.
#
# The default Squid config file shows what the defaults for
# various options happen to be. If you don't need to change the
# default, you shouldn't uncomment the line. Doing so may cause
# run-time problems. In some cases "none" refers to no default
# setting at all, while in other cases it refers to a valid
# option - the comments for that keyword indicate if this is the
# case.
#
# NETWORK OPTIONS
# ------------------------------------------------------------------
-------------
#      TAG:  http_port
#                    Usage:    port
#                                   hostname:port
#                                   1.2.3.4:port
#
# The socket addresses where Squid will listen for HTTP client
# requests. You may specify multiple socket addresses.
# There are three forms: port alone, hostname with port, and
# IP address with port. If you specify a hostname or IP
# address, then Squid binds the socket to that specific
# address. This replaces the old 'tcp_incoming_address'
# option. Most likely, you do not need to bind to a specific
# address, so you can use the port number alone.
#
#The default port number is 3128.
#
# If you are running Squid in accelerator mode, then you
# probably want to listen on port 80 also, or instead.
#
```

```
# The -a command line option will override the *first* port
# number listed here. That option will NOT override an IP
# address, however.
#
# You may specify multiple socket addresses on multiple lines.
#
 http_port 3128

# TAG: icp_port
# The port number where Squid sends and receives ICP queries to
# and from neighbor caches. Default is 3130. To disable use
# "0". May be overridden with -u on the command line.
#
#icp_port 3130
icp_port 0

# TAG: htcp_port
# The port number where Squid sends and receives HTCP queries to
# and from neighbor caches. Default is 4827. To disable use
# "0".
#
# To enable this option, you must use --enable-htcp with the
# configure script.
#htcp_port 4827

# TAG: mcast_groups
# This tag specifies a list of multicast groups which your server
# should join to receive multicasted ICP queries.
#
# NOTE! Be very careful what you put here! Be sure you
# understand the difference between an ICP _query_ and an ICP
# _reply_. This option is to be set only if you want to RECEIVE
# multicast queries. Do NOT set this option to SEND multicast
# ICP (use cache_peer for that). ICP replies are always sent via
# unicast, so this option does not affect whether or not you will
# receive replies from multicast group members.
#
# You must be very careful to NOT use a multicast address which
# is already in use by another group of caches.
#
# If you are unsure about multicast, please read the Multicast
# chapter in the Squid FAQ (http://squid.nlanr.net/Squid/FAQ/).
#
# Usage: mcast_groups 239.128.16.128 224.0.1.20
#
# By default, Squid doesn't listen on any multicast groups.
#
#mcast_groups 239.128.16.128

# TAG: tcp_outgoing_address
# TAG: udp_incoming_address
# TAG: udp_outgoing_address
# Usage: tcp_incoming_address 10.20.30.40
# udp_outgoing_address fully.qualified.domain.name
#
# tcp_outgoing_address is used for connections made to remote
# servers and other caches.
```

```
# udp_incoming_address is used for the ICP socket receiving
packets
# from other caches.
# udp_outgoing_address is used for ICP packets sent out to other
# caches.
#
# The default behavior is to not bind to any specific address.
#
# NOTE, udp_incoming_address and udp_outgoing_address can not
# have the same value (unless it is 0.0.0.0) since they both use
# port 3130.
#
# NOTE, tcp_incoming_address has been removed. You can now
# specify IP addresses on the 'http_port' line.
#
#tcp_outgoing_address 0.0.0.0
#udp_incoming_address 0.0.0.0
#udp_outgoing_address 0.0.0.0


# OPTIONS WHICH AFFECT THE CACHE SIZE
# -----------------------------------------------------------------
-------------

# TAG: cache_mem (bytes)
# NOTE: THIS PARAMETER DOES NOT SPECIFY THE MAXIMUM PROCESS
# SIZE. IT PLACES A LIMIT ON ONE ASPECT OF SQUID'S MEMORY
# USAGE. SQUID USES MEMORY FOR OTHER THINGS AS WELL.
# YOUR PROCESS WILL PROBABLY BECOME TWICE OR THREE TIMES
# BIGGER THAN THE VALUE YOU PUT HERE
#
# 'cache_mem' specifies the ideal amount of memory to be used
# for:
# * In-Transit objects
# * Hot Objects
# * Negative-Cached objects
#
# Data for these objects are stored in 4 KB blocks. This
# parameter specifies the ideal upper limit on the total size of
# 4 KB blocks allocated. In-Transit objects take the highest
# priority.
#
# In-transit objects have priority over the others. When
# additional space is needed for incoming data, negative-cached
# and hot objects will be released. In other words, the
# negative-cached and hot objects will fill up any unused space
# not needed for in-transit objects.
#
# If circumstances require, this limit will be exceeded.
# Specifically, if your incoming request rate requires more than
# 'cache_mem' of memory to hold in-transit objects, Squid will
# exceed this limit to satisfy the new requests. When the load
# decreases, blocks will be freed until the high-water mark is
# reached. Thereafter, blocks will be used to store hot
# objects.
#
# The default is 8 Megabytes.
#
```

```
cache_mem 8 MB

# TAG: cache_swap_low (percent, 0-100)
# TAG: cache_swap_high (percent, 0-100)
#
# The low- and high-water marks for cache object replacement.
# Replacement begins when the swap (disk) usage is above the
# low-water mark and attempts to maintain utilization near the
# low-water mark. As swap utilization gets close to high-water
# mark object eviction becomes more aggressive. If utilization is
# close to the low-water mark less replacement is done each time.
#
# Defaults are 90% and 95%. If you have a large cache, 5% could be
# hundreds of MB. If this is the case you may wish to set these
# numbers closer together.
#
cache_swap_low 90
cache_swap_high 95

# TAG: maximum_object_size (bytes)
# Objects larger than this size will NOT be saved on disk. The
# value is specified in kilobytes, and the default is 4MB. If
            # you wish to get a high BYTES hit ratio, you should probably
# increase this (one 32 MB object hit counts for 3200 10KB
# hits). If you wish to increase speed more than your want to
# save bandwidth you should leave this low.
#
# NOTE: if using the LFUDA replacement policy you should increase
# this value to maximize the byte hit rate improvement of LFUDA!
# See replacement_policy below for a discussion of this policy.
#
maximum_object_size 4096 KB

# TAG: minimum_object_size (bytes)
# Objects smaller than this size will NOT be saved on disk. The
# value is specified in kilobytes, and the default is 0 KB, which-
# means there is no minimum.
minimum_object_size 0 KB

# TAG: ipcache_size (number of entries)
# TAG: ipcache_low (percent)
# TAG: ipcache_high (percent)
# The size, low-, and high-water marks for the IP cache.
#
ipcache_size 1024
ipcache_low 90
ipcache_high 95

# TAG: fqdncache_size (number of entries)
# Maximum number of FQDN cache entries.
fqdncache_size 1024


# LOGFILE PATHNAMES AND CACHE DIRECTORIES
# ------------------------------------------------------------------
-------------
# TAG: cache_dir
# Usage:
```

```
#
# cache_dir Type Directory-Name Mbytes Level-1 Level2
#
# You can specify multiple cache_dir lines to spread the
# cache among different disk partitions.
#
# Type specifies the kind of storage system to use. Most
# everyone will want to use "ufs" as the type. If you are using
# Async I/O (--enable async-io) on Linux or Solaris, then you may
# want to try "asyncufs" as the type. Async IO support may be
# buggy, however, so beware.
#
# 'Directory' is a top-level directory where cache swap
# files will be stored. If you want to use an entire disk
# for caching, then this can be the mount-point directory.
# The directory must exist and be writable by the Squid
# process. Squid will NOT create this directory for you.
#
# If no 'cache_dir' lines are specified, the following
# default will be used: /var/spool/squid.
#
# 'Mbytes' is the amount of disk space (MB) to use under this
# directory. The default is 100 MB. Change this to suit your
# configuration.
#
# 'Level-1' is the number of first-level subdirectories which
# will be created under the 'Directory'. The default is 16.
#
# 'Level-2' is the number of second-level subdirectories which
# will be created under each first-level directory. The default
# is 256.
#
#cache_dir ufs /usr/local/squid/cache 100 16 256
cache_dir ufs /var/spool/squid/cache 100 16 256

# TAG: cache_access_log
# Logs the client request activity. Contains an entry for
# every HTTP and ICP queries received.
#
cache_access_log /var/log/squid/access.log

# TAG: cache_log
# Cache logging file. This is where general information about
# your cache's behavior goes. You can increase the amount of data
# logged to this file with the "debug_options" tag below.
#
cache_log /var/log/squid/cache.log

# TAG: cache_store_log
# Logs the activities of the storage manager. Shows which
# objects are ejected from the cache, and which objects are
# saved and for how long. To disable, enter "none". There are
# not really utilities to analyze this data, so you can safely
# disable it.
#
#cache_store_log /var/log/squid/store.log
cache_store_log none
```

```
# TAG: cache_swap_log
# Location for the cache "swap.log." This log file holds the
# metadata of objects saved on disk. It is used to rebuild the
# cache during startup. Normally this file resides in the first
# 'cache_dir' directory, but you may specify an alternate
# pathname here. Note you must give a full filename, not just
# a directory. Since this is the index for the whole object
# list you CANNOT periodically rotate it!
#
# If you have more than one 'cache_dir', these swap logs will
# have names such as:
#
# cache_swap_log.00
# cache_swap_log.01
# cache_swap_log.02
#
# The numbered extension (which is added automatically)
# corresponds to the order of the 'cache_dir' lines in this
# configuration file. If you change the order of the 'cache_dir'
# lines in this file, then these log files will NOT correspond to
# the correct 'cache_dir' entry (unless you manually rename
# them). We recommend that you do NOT use this option. It is
# better to keep these log files in each 'cache_dir' directory.
#
#cache_swap_log

# TAG: emulate_httpd_log on|off
# The Cache can emulate the log file format which many 'httpd'
# programs use. To disable/enable this emulation, set
# emulate_httpd_log to 'off' or 'on'. The default
# is to use the native log format since it includes useful #
information that Squid-specific log analyzers use.
#
emulate_httpd_log on

# TAG: mime_table
# Pathname to Squid's MIME table. You shouldn't need to change
# this, but the default file contains examples and formatting
# information if you do.
#
#mime_table /usr/local/squid/etc/mime.conf
mime_table /etc/squid/mime.conf

# TAG: log_mime_hdrs on|off
# The Cache can record both the request and the response MIME
# headers for each HTTP transaction. The headers are encoded
# safely and will appear as two bracketed fields at the end of
# the access log (for either the native or httpd-emulated log
# formats). To enable this logging set log_mime_hdrs to 'on'.
#
#log_mime_hdrs off

# TAG: useragent_log
# If configured with the "--enable-useragent_log" configure
# option, Squid will write the User-Agent field from HTTP
# requests to the filename specified here. By default
# useragent_log is disabled.
#
```

```
#useragent_log none

# TAG: pid_filename
# A filename to write the process-id to. To disable, enter "none".
#
pid_filename /var/log/squid/squid.pid

# TAG: debug_options
# Logging options are set as section,level where each source file
# is assigned a unique section. Lower levels result in less
# output, Full debugging (level 9) can result in a very large
# log file, so be careful. The magic word "ALL" sets debugging
# levels for all sections. We recommend normally running with
# "ALL,1".
#
debug_options ALL,1

# TAG: log_fqdn on|off
# Turn this on if you wish to log fully qualified domain names
# in the access.log. To do this Squid does a DNS lookup of all
# IP's connecting to it. This can (in some situations) increase
# latency, which makes your cache seem slower for interactive
# browsing.
#
#log_fqdn off

# TAG: client_netmask
# A netmask for client addresses in logfiles and cachemgr output.
# Change this to protect the privacy of your cache clients.
# A netmask of 255.255.255.0 will log all IP's in that range with
# the last digit set to '0'.
#
#client_netmask 255.255.255.0


# OPTIONS FOR EXTERNAL SUPPORT PROGRAMS
# ------------------------------------------------------------
-------------

# TAG: ftp_user
# If you want the anonymous login password to be more informative
# (and enable the use of picky ftp servers), set this to something
# reasonable for your domain, like wwwuser@somewhere.net
#
# The reason why this is domainless by default is that the
# request can be made on the behalf of a user in any domain,
# depending on how the cache is used.
# Some ftp server also validate that the email address is valid
# (for example perl.com).
#
#ftp_user Squid@

# TAG: ftp_list_width
# Sets the width of ftp listings. This should be set to fit in
# the width of a standard browser. Setting this too small
# can cut off long filenames when browsing ftp sites.
#
#ftp_list_width 32
```

```
# TAG: ftp_passive
# If your firewall does not allow Squid to use passive
# connections, then turn off this option.
##ftp_passive on

# TAG: cache_dns_program
# Specify the location of the executable for dnslookup process.
#
#cache_dns_program /usr/lib/squid/dnsserver

# TAG: dns_children
# The number of processes spawn to service DNS name lookups.
# For heavily loaded caches on large servers, you should
# probably increase this value to at least 10. The maximum
# is 32. The default is 5.
#
# You must have at least one dnsserver process.
#
#dns_children 5

# TAG: dns_defnames on|off
# Normally the 'dnsserver' disables the RES_DEFNAMES resolver
# option (see res_init(3)). This prevents caches in a hierarchy
# from interpreting single-component hostnames locally. To allow
# dnsserver to handle single-component names, enable this
# option.
#
#dns_defnames off

# TAG: dns_nameservers
# Use this if you want to specify a list of DNS name servers
# (IP addresses) to use instead of those given in your
# /etc/resolv.conf file.
#
# Example: dns_nameservers 10.0.0.1 192.172.0.4
#
#dns_nameservers none

# TAG: unlinkd_program
# Specify the location of the executable for file deletion
process.
# This isn't needed if you are using async-io since it's handled
by
# a thread.
#
#unlinkd_program /usr/lib/squid/unlinkd

# TAG: pinger_program
# Specify the location of the executable for the pinger process.
# This is only useful if you configured Squid (during compilation)
# with the '--enable-icmp' option.
#
#pinger_program /usr/lib/squid/pinger

# TAG: redirect_program
# Specify the location of the executable for the URL redirector.
```

```
# Since they can perform almost any function there isn't one
included.
# See the Release-Notes for information on how to write one.
# By default, a redirector is not used.
#
#redirect_program none

# TAG: redirect_children
# The number of redirector processes to spawn. If you start
# too few Squid will have to wait for them to process a backlog of
# URLs, slowing it down. If you start too many they will use RAM
# and other system resources.
#
#redirect_children 5

# TAG: redirect_rewrites_host_header
# By default Squid rewrites any Host: header in redirected
# requests. If you are running a accelerator then this may
# not be a wanted effect of a redirector.
redirect_rewrites_host_header on

# TAG: redirector_access
# If defined, this access list specifies which requests are
# sent to the redirector processes. By default all requests
# are sent.

# TAG: authenticate_program
# Specify the command for the external authenticator. Such a
# program reads a line containing "username password" and replies
# "OK" or "ERR" in an endless loop. If you use an authenticator,
# make sure you have 1 acl of type proxy_auth. By default, the
# authenticator_program is not used.
#
# If you want to use the traditional proxy authentication,
# jump over to the ../auth_modules/NCSA directory and
# type:
# % make
# % make install
#
# Then, set this line to something like
#
# authenticate_program /usr/bin/ncsa_auth /usr/etc/passwd
#
#authenticate_program none

# TAG: authenticate_children
# The number of authenticator processes to spawn (default 5). If
you
# start too few Squid will have to wait for them to process a
backlog
# of usercode/password verifications, slowing it down. When
password
# verifications are done via a (slow) network you are likely to
need
# lots of authenticator processes.
#
#authenticate_children 5
# TAG: authenticate_ttl
```

```
# The time a checked username/password combination remains cached
# (default 3600). If a wrong password is given for a cached user,
# the user gets removed from the username/password cache forcing
# a revalidation.
#
#authenticate_ttl 3600

# TAG: authenticate_ip_ttl
# With this option you control how long a proxy authentication
# will be bound to a specific IP address. If a request using
# the same user name is received during this time then access
# will be denied and both users are required to reauthenticate
# them selves. The idea behind this is to make it annoying
# for people to share their password to their friends, but
# yet allow a dialup user to reconnect on a different dialup
# port.
#
# The default is 0 to disable the check. Recommended value
# if you have dialup users are no more than 60 (seconds). If
# all your users are stationary then higher values may be
# used.
#
#authenticate_ip_ttl 0


# OPTIONS FOR TUNING THE CACHE
# ------------------------------------------------------------
-------------

# TAG: wais_relay_host
# TAG: wais_relay_port
# Relay WAIS request to host (1st arg) at port (2 arg).
#
#wais_relay_host localhost
#wais_relay_port 8000

# TAG: request_header_max_size (KB)
# This specifies the maximum size for HTTP headers in a request.
# Request headers are usually relatively small (about 512 bytes).
# Placing a limit on the request header size will catch certain
# bugs (for example with persistent connections) and possibly
# buffer-overflow or denial-of-service attacks.
#request_header_max_size 10 KB

# TAG: request_body_max_size (KB)
# This specifies the maximum size for an HTTP request body.
# In other words, the maximum size of a PUT/POST request.
# A user who attempts to send a request with a body larger
# than this limit receives an "Invalid Request" error message.
# If you set this parameter to a zero, there will be no limit
# imposed.
#request_body_max_size 1 MB

# TAG: reply_body_max_size (KB)
# This option specifies the maximum size of a reply body. It
# can be used to prevent users from downloading very large files,
# such as MP3's and movies. The reply size is checked twice.
# First when we get the reply headers, we check the
```

```
# content-length value. If the content length value exists and
# is larger than this parameter, the request is denied and the
# user receives an error message that says "the request or reply
# is too large." If there is no content-length, and the reply
# size exceeds this limit, the client's connection is just closed
# and they will receive a partial reply.
#
# NOTE: downstream caches probably can not detect a partial reply
# if there is no content-length header, so they will cache
# partial responses and give them out as hits. You should NOT
# use this option if you have downstream caches.
#
# If you set this parameter to zero (the default), there will be
# no limit imposed.
#reply_body_max_size 0

# TAG: refresh_pattern
# usage: refresh_pattern [-i] regex min percent max [options]
#
# By default, regular expressions are CASE-SENSITIVE. To make
# them case-insensitive, use the -i option.
#
# 'Min' is the time (in minutes) an object without an explicit
# expiry time should be considered fresh. The recommended
# value is 0, any higher values may cause dynamic applications
# to be erroneously cached unless the application designer
# has taken the appropriate actions.
#
# 'Percent' is a percentage of the objects age (time since last
# modification age) an object without explicit expiry time
# will be considered fresh.
#
# 'Max' is an upper limit on how long objects without an explicit
# expiry time will be considered fresh.
#
# options: override-expire
# override-lastmod
# reload-into-ims
# ignore-reload
#
# override-expire enforces min age even if the server
# sent a Expires: header. Doing this VIOLATES the HTTP
# standard. Enabling this feature could make you liable
# for problems which it causes.
#
# override-lastmod enforces min age even on objects
# that was modified recently.
#
# reload-into-ims changes client no-cache or ``reload''
# to If-Modified-Since requests. Doing this VIOLATES the
# HTTP standard. Enabling this feature could make you
# liable for problems which it causes.
#
# ignore-reload ignores a client no-cache or ``reload''
# header. Doing this VIOLATES the HTTP standard. Enabling
# this feature could make you liable for problems which
# it causes.
#
```

```
# Please see the file doc/Release-Notes-1.1.txt for a full
# description of Squid's refresh algorithm. Basically a
# cached object is: (the order is changed from 1.1.X)
#
# FRESH if expires < now, else STALE
# STALE if age > max
# FRESH if lm-factor < percent, else STALE
# FRESH if age < min # else STALE
#
# The refresh_pattern lines are checked in the order listed here.
# The first entry which matches is used. If none of the entries
# match, then the default will be used.
#
#Default:
refresh_pattern          ^ftp:                   1440        20%
      10080
refresh_pattern    ^gopher:1440       0%        1440
refresh_pattern .                    0            20%        4320
```

 # TAG: replacement_policy

```
# The cache replacement policy parameter determines which
# objects are evicted (replaced) when disk space is needed.
# Squid used to have only a single replacement policy, LRU.
# But when built with -DHEAP_REPLACEMENT you can choose
# between two new, enhanced policies:
#
# GDSF: Greedy-Dual Size Frequency
# LFUDA: Least Frequently Used with Dynamic Aging
#
# Both of these policies are frequency based rather than recency
# based, and perform better than LRU.
#
# The GDSF policy optimizes object hit rate by keeping smaller
# popular objects in cache so it has a better chance of getting a
# hit. It achieves a lower byte hit rate than LFUDA though since
# it evicts larger (possibly popular) objects.
#
# The LFUDA policy keeps popular objects in cache regardless of
# their size and thus optimizes byte hit rate at the expense of
# hit rate since one large, popular object will prevent many
# smaller, slightly less popular objects from being cached.
#
# Both policies utilize a dynamic aging mechanism that prevents
# cache pollution that can otherwise occur with frequency-based
# replacement policies.
#
# NOTE: if using the LFUDA replacement policy you should increase
# the value of maximum_object_size above its default of 4096 KB to
# to maximize the potential byte hit rate improvement of LFUDA.
#
# For more information about these cache replacement policies see
# http://www.hpl.hp.com/techreports/1999/HPL-1999-69.html and
# http://fog.hpl.external.hp.com/techreports/98/HPL-98-173.html.
#
#replacement_policy LFUDA

# TAG: reference_age
# As a part of normal operation, Squid performs Least Recently
```

```
# Used removal of cached objects. The LRU age for removal is
# computed dynamically, based on the amount of disk space in
# use. The dynamic value can be seen in the Cache Manager 'info'
# output.
#
# The 'reference_age' parameter defines the maximum LRU age. For
# example, setting reference_age to '1 week' will cause objects
# to be removed if they have not been accessed for a week or
# more. The default value is one year.
#
# Specify a number here, followed by units of time. For example:
# 1 week
# 3.5 days
# 4 months
# 2.2 hours
#
# NOTE: this parameter is not used when using the enhanced
# replacement policies, GDSH or LFUDA.
#
reference_age 1 month

# TAG: quick_abort_min (KB)
# TAG: quick_abort_max (KB)
# TAG: quick_abort_pct (percent)
# The cache can be configured to continue downloading aborted
# requests. This may be undesirable on slow (e.g. SLIP) links
# and/or very busy caches. Impatient users may tie up file
# descriptors and bandwidth by repeatedly requesting and
# immediately aborting downloads.
#
# When the user aborts a request, Squid will check the
# quick_abort values to the amount of data transfered until
# then.
#
# If the transfer has less than 'quick_abort_min' KB remaining,
# it will finish the retrieval. Setting 'quick_abort_min' to -1
# will disable the quick_abort feature.
#
# If the transfer has more than 'quick_abort_max' KB remaining,
# it will abort the retrieval.
#
# If more than 'quick_abort_pct' of the transfer has completed,
# it will finish the retrieval.
#
quick_abort_min 16 KB
quick_abort_max 16 KB
quick_abort_pct 95

# TAG: negative_ttl time-units
# Time-to-Live (TTL) for failed requests. Certain types of
# failures (such as "connection refused" and "404 Not Found") are
# negatively-cached for a configurable amount of time. The
# default is 5 minutes. Note that this is different from
# negative caching of DNS lookups.
#
negative_ttl 5 minutes

# TAG: positive_dns_ttl time-units
```

```
# Time-to-Live (TTL) for positive caching of successful DNS
lookups.
# Default is 6 hours (360 minutes). If you want to minimize the
# use of Squid's ipcache, set this to 1, not 0.
#
positive_dns_ttl 6 hours

# TAG: negative_dns_ttl time-units
# Time-to-Live (TTL) for negative caching of failed DNS lookups.
#
negative_dns_ttl 5 minutes

# TAG: range_offset_limit (bytes)
# Sets a upper limit on how far into the the file a Range request
# may be to cause Squid to prefetch the whole file. If beyond this
# limit then Squid forwards the Range request as it is and the
result
# is NOT cached.
#
# This is to stop a far ahead range request (lets say start at
17MB)
# from making Squid fetch the whole object up to that point before
# sending anything to the client.
#
# A value of -1 causes Squid to always fetch the object from the
# beginning so that it may cache the result. (2.0 style)
#
# A value of 0 causes Squid to never fetch more than the client
# client requested. (default)
#
#range_offset_limit 0 KB


# TIMEOUTS
# ------------------------------------------------------------------
-------------

# TAG: connect_timeout time-units
# Some systems (notably Linux) can not be relied upon to properly
# time out connect(2) requests. Therefore the Squid process
# enforces its own timeout on server connections. This parameter
# specifies how long to wait for the connect to complete. The
# default is two minutes (120 seconds).
#
#connect_timeout 120 seconds

# TAG: peer_connect_timeout       time-units
# This parameter specifies how long to wait for a pending TCP
# connection to a peer cache. The default is 30 seconds. You
# may also set different timeout values for individual neighbors
# with the 'connect-timeout' option on a 'cache_peer' line.
#peer_connect_timeout 30 seconds

# TAG: siteselect_timeout time-units
# For URN to multiple URL's URL selection
#
#siteselect_timeout 4 seconds
# TAG: read_timeout time-units
```

```
# The read_timeout is applied on server-side connections. After
# each successful read(), the timeout will be extended by this
# amount. If no data is read again after this amount of time,
# the request is aborted and logged with ERR_READ_TIMEOUT. The
# default is 15 minutes.
#
#read_timeout 15 minutes

# TAG: request_timeout
# How long to wait for an HTTP request after connection
# establishment. For persistent connections, wait this long
# after the previous request completes.
#
#request_timeout 30 seconds

# TAG: client_lifetime time-units
# The maximum amount of time that a client (browser) is allowed to
# remain connected to the cache process. This protects the Cache
# from having a lot of sockets (and hence file descriptors) tied
up
# in a CLOSE_WAIT state from remote clients that go away without
# properly shutting down (either because of a network failure or
# because of a poor client implementation). The default is one
# day, 1440 minutes.
#
# NOTE: The default value is intended to be much larger than any
# client would ever need to be connected to your cache. You
# should probably change client_lifetime only as a last resort.
# If you seem to have many client connections tying up
# filedescriptors, we recommend first tuning the read_timeout,
# request_timeout, pconn_timeout and quick_abort values.
#
#client_lifetime 1 day

# TAG: half_closed_clients
# Some clients may shutdown the sending side of their TCP
# connections, while leaving their receiving sides open.
Sometimes,
# Squid can not tell the difference between a half-closed and a
# fully-closed TCP connection. By default, half-closed client
# connections are kept open until a read(2) or write(2) on the
# socket returns an error. Change this option to 'off' and Squid
# will immediately close client connections when read(2) returns
# "no more data to read."
#
#half_closed_clients on

# TAG: pconn_timeout
# Timeout for idle persistent connections to servers and other
# proxies.
#pconn_timeout 120 seconds

# TAG: ident_timeout
# Maximum time to wait for IDENT requests. If this is too high,
# and you enabled 'ident_lookup', then you might be susceptible
# to denial-of-service by having many ident requests going at
# once.
#
```

```
# Only src type ACL checks are fully supported. A src_domain
# ACL might work at times, but it will not always provide
# the correct result.
#
# This option may be disabled by using --disable-ident with
# the configure script.
#ident_timeout 10 seconds

# TAG: shutdown_lifetime time-units
# When SIGTERM or SIGHUP is received, the cache is put into
# "shutdown pending" mode until all active sockets are closed.
# This value is the lifetime to set for all open descriptors
# during shutdown mode. Any active clients after this many
# seconds will receive a 'timeout' message.
#
#shutdown_lifetime 30 seconds


# ACCESS CONTROLS
# ----------------------------------------------------------------
-------------

# TAG: acl
# Defining an Access List
#
# acl aclname acltype string1 ...
# acl aclname acltype "file" ...
#
# when using "file", the file should contain one item per line
#
# acltype is one of src dst srcdomain dstdomain url_pattern
# urlpath_pattern time port proto method browser user
#
# By default, regular expressions are CASE-SENSITIVE. To make
# them case-insensitive, use the -i option.
#
# acl aclname src       ip-address/netmask ...  (clients IP
address)
# acl aclname src       addr1-addr2/netmask ...       (range of
addresses)
# acl aclname dst       ip-address/netmask ...  (URL host's IP
address)
# acl aclname myip      ip-address/netmask ...  (local socket IP
address)
#
# acl aclname srcdomain       .foo.com ...      # reverse lookup,
client IP
# acl aclname dstdomain       .foo.com ...      # Destination
server from URL
# acl aclname srcdom_regex [-i] xxx ...  # regex matching client
name
# acl aclname dstdom_regex [-i] xxx ...  # regex matching server
#     # For dstdomain and dstdom_regex a reverse lookup is tried
if a IP
#     # based URL is used. The name "none" is used if the reverse
lookup
#     # fails.
#
```

```
# acl       aclname     time  [day-abbrevs] [h1:m1-h2:m2]
# day-abbrevs:
# S - Sunday
# M - Monday # T - Tuesday
# W - Wednesday
# H - Thursday
# F - Friday
# A - Saturday
# h1:m1 must be less than h2:m2
# acl aclname url_regex [-i] ^http:// ...      # regex matching
on whole URL
# acl aclname urlpath_regex [-i] \.gif$ ...    # regex matching
on URL path
# acl aclname port 80 70 21 ...
# acl aclname port 0-1024 ...             # ranges allowed
# acl aclname myport 3128 ...             # (local socket TCP
port)
# acl aclname proto HTTP FTP ...
# acl aclname method GET POST ...
# acl aclname browser [-i] regexp
#      # pattern match on User-Agent header
# acl aclname ident username ...
#     # string match on ident output.
#     # use REQUIRED to accept any non-null ident.
# acl aclname src_as number ...
# acl aclname dst_as number ...
#     # Except for access control, AS numbers can be used for
#     # routing of requests to specific caches. Here's an
#     # example for routing all requests for AS#1241 and only
#     # those to mycache.mydomain.net:
#     # acl asexample dst_as 1241
#     # cache_peer_access mycache.mydomain.net allow asexample
#     # cache_peer_access mycache_mydomain.net deny all
#
# acl aclname proxy_auth username ...
#     # list of valid usernames
#     # use REQUIRED to accept any valid username.
#     #
#     # NOTE: when a Proxy-Authentication header is sent but it is
not
#     # needed during ACL checking the username is NOT logged
#     # in access.log.
#     #
#     # NOTE: proxy_auth requires a EXTERNAL authentication
program
#     # to check username/password combinations (see
#     # authenticate_program).
#     #
#     # WARNING: proxy_auth can't be used in a transparent proxy.
It
#     # collides with any authentication done by origin servers.
It may
#     # seem like it works at first, but it doesn't.
#
# acl aclname snmp_community string ...
#     # A community string to limit access to your SNMP Agent
#     # Example:
#     #
```

```
#     # acl snmppublic snmp_community public
#
# acl aclname maxconn number
#     # This will be matched when the client's IP address has
#     # more than <number> HTTP connections established.
#
#
#Examples:
#acl myexample dst_as 1241
#acl password proxy_auth REQUIRED
#
#Defaults:
acl all src 0.0.0.0/0.0.0.0
acl manager proto cache_object
acl localhost src 127.0.0.1/255.255.255.255
acl mired src 172.16.0.0/255.255.0.0
acl SSL_ports port 443 563
acl Safe_ports port 80 21 443 563 70 210 1025-65535
acl Safe_ports port 280 # http-mgmt
acl Safe_ports port 488 # gss-http
acl Safe_ports port 591 # filemaker
acl Safe_ports port 777 # multiling http
acl CONNECT method CONNECT

# TAG: http_access
# Allowing or Denying access based on defined access lists
#
# Access to the HTTP port:
# http_access allow|deny [!]aclname ...
#
# Access to the ICP port:
# icp_access allow|deny [!]aclname ...
#
# NOTE on default values:
#
# If there are no "access" lines present, the default is to allow
# the request.
#
# If none of the "access" lines cause a match, the default is the
# opposite of the last line in the list. If the last line was
# deny, then the default is allow. Conversely, if the last line
# is allow, the default will be deny. For these reasons, it is a
# good idea to have an "deny all" or "allow all" entry at the end
# of your access lists to avoid potential confusion.
#
#Default configuration:
http_access allow manager localhost
http_access deny manager
#http_access deny !Safe_ports
#http_access deny CONNECT !SSL_ports
http_access allow CONNECT !SSL_ports
#
# INSERT YOUR OWN RULE(S) HERE TO ALLOW ACCESS FROM YOUR CLIENTS
#
#http_access allow localhost
#http_access allow all
 #http_access deny all
http_access allow mired
```

```
http_access deny !mired

# TAG: icp_access
# Reply to all ICP queries we receive
#
icp_access allow all

# TAG: miss_access
# Use to force your neighbors to use you as a sibling instead of
# a parent. For example:
#
# acl localclients src 172.16.0.0/16
# miss_access allow localclients
# miss_access deny !localclients
#
# This means that only your local clients are allowed to fetch
# MISSES and all other clients can only fetch HITS.
#
# By default, allow all clients who passed the http_access rules
# to fetch MISSES from us.
miss_access allow all

# TAG: cache_peer_access
# Similar to 'cache_peer_domain' but provides more flexibility by
# using ACL elements.
#
# cache_peer_access cache-host allow|deny [!]aclname ...
#
# The syntax is identical to 'http_access' and the other lists of
# ACL elements. See the comments for 'http_access' below, or
# the Squid FAQ (http://squid.nlanr.net/Squid/FAQ/FAQ-10.html).

# TAG: proxy_auth_realm
# Specifies the realm name which is to be reported to the client
for
# proxy authentication (part of the text the user will see when
# prompted their username and password).
#
#proxy_auth_realm Squid proxy-caching web server

# TAG: ident_lookup_access
# A list of ACL elements which, if matched, cause an ident
# (RFC 931) lookup to be performed for this request. For
# example, you might choose to always perform ident lookups
# for your main multi-user Unix boxes, but not for your Macs
# and PCs. By default, ident lookups are not performed for
# any requests.
#
# To enable ident lookups for specific client addresses, you
# can follow this example:
#
# acl ident_aware_hosts src 198.168.1.0/255.255.255.0
# ident_lookup_access allow ident_aware_hosts
# ident_lookup_access deny all
#
# This option may be disabled by using --disable-ident with
# the configure script. #ident_lookup_access deny all
```

```
# ADMINISTRATIVE PARAMETERS
# ------------------------------------------------------------------
-------------

# TAG: cache_mgr
# Email-address of local cache manager who will receive
# mail if the cache dies. The default is "webmaster."
#
cache_mgr root

# TAG: cache_effective_user
# TAG: cache_effective_group
#
# If the cache is run as root, it will change its effective/real
# UID/GID to the UID/GID specified below. The default is to
# change to UID to squid and GID to squid.
#
# If Squid is not started as root, the default is to keep the
# current UID/GID. Note that if Squid is not started as root then
# you cannot set http_port to a value lower than 1024.
#
cache_effective_user squid
cache_effective_group squid

# TAG: visible_hostname
# If you want to present a special hostname in error messages,
etc,
# then define this. Otherwise, the return value of gethostname()
# will be used. If you have multiple caches in a cluster and
# get errors about IP-forwarding you must set them to have
individual
# names with this setting.
#
visible_hostname minotauro

# TAG: unique_hostname
# If you want to have multiple machines with the same
# 'visible_hostname' then you must give each machine a different
# 'unique_hostname' so that forwarding loops can be detected.
#
#unique_hostname www-cache1.foo.org

# TAG: hostname_aliases
# A list of other DNS names that your cache has.


# OPTIONS FOR THE CACHE REGISTRATION SERVICE
# ------------------------------------------------------------------
-------------
#
# This section contains parameters for the (optional) cache
# announcement service. This service is provided to help
# cache administrators locate one another in order to join or
# create cache hierarchies.
#
# An 'announcement' message is sent (via UDP) to the registration
# service by Squid. By default, the announcement message is NOT
# SENT unless you enable it with 'announce_period' below.
```

```
#
# The announcement message includes your hostname, plus the
# following information from this configuration file:
#
# http_port
# icp_port
# cache_mgr
#
# All current information is processed regularly and made
# available on the Web at http://ircache.nlanr.net/Cache/Tracker/.

# TAG: announce_period
# This is how frequently to send cache announcements. The
# default is `0' which disables sending the announcement
# messages.
#
# To enable announcing your cache, just uncomment the line
# below.
#
#announce_period 1 day

# TAG: announce_host
# TAG: announce_file
# TAG: announce_port
# announce_host and announce_port set the hostname and port
# number where the registration message will be sent.
#
# Hostname will default to 'tracker.ircache.net' and port will
# default default to 3131. If the 'filename' argument is given,
# the contents of that file will be included in the announce
# message.
#
#announce_host tracker.ircache.net
#announce_port 3131


# HTTPD-ACCELERATOR OPTIONS
# -----------------------------------------------------------------
-------------

# TAG: httpd_accel_host
# TAG: httpd_accel_port
# If you want to run Squid as an httpd accelerator, define the
# host name and port number where the real HTTP server is.
#
# If you want virtual host support then specify the hostname
# as "virtual".
#
# If you want virtual port support then specify the port as "0".
#
# NOTE: enabling httpd_accel_host disables proxy-caching and
# ICP. If you want these features enabled also, then set
# the 'httpd_accel_with_proxy' option.
#
#httpd_accel_host hostname
#httpd_accel_port port
httpd_accel_host virtual
httpd_accel_port 80
```

```
# TAG: httpd_accel_with_proxy on|off
# If you want to use Squid as both a local httpd accelerator
# and as a proxy, change this to 'on'.
#
#httpd_accel_with_proxy off httpd_
accel_with_proxy on


# TAG: httpd_accel_uses_host_header on|off
# HTTP/1.1 requests include a Host: header which is basically the
# hostname from the URL. Squid can be an accelerator for
# different HTTP servers by looking at this header. However,
# Squid does NOT check the value of the Host header, so it opens
# a big security hole. We recommend that this option remain
# disabled unless you are sure of what you are doing.
#
# However, you will need to enable this option if you run Squid
# as a transparent proxy. Otherwise, virtual servers which
# require the Host: header will not be properly cached.
#httpd_accel_uses_host_header off
httpd_accel_uses_host_header on


# MISCELLANEOUS
# -----------------------------------------------------------------
-------------

# TAG: dns_testnames
# The DNS tests exit as soon as the first site is successfully
looked up
#
# If you want to disable DNS tests, do not comment out or delete
this
# list. Instead use the -D command line option
#
dns_testnames netscape.com internic.net nlanr.net microsoft.com

# TAG: logfile_rotate
# Specifies the number of logfile rotations to make when you
# type 'squid -k rotate'. The default is 10, which will rotate
# with extensions 0 through 9. Setting logfile_rotate to 0 will
# disable the rotation, but the logfiles are still closed and
# re-opened. This will enable you to rename the logfiles
# yourself just before sending the rotate signal.
#
# Note, the 'squid -k rotate' command normally sends a USR1
# signal to the running squid process. In certain situations
# (e.g. on Linux with Async I/O), USR1 is used for other
# purposes, so -k rotate uses another signal. It is best to get
# in the habit of using 'squid -k rotate' instead of 'kill -USR1
# <pid>'.
#
#logfile_rotate 0

# TAG: append_domain
# Appends local domain name to hostnames without any dots in
# them. append_domain must begin with a period.
#
```

append_domain .minotauro.telemed.org

```
# TAG: tcp_recv_bufsize (bytes)
# Size of receive buffer to set for TCP sockets. Probably just
# as easy to change your kernel's default. Set to zero to use
# the default buffer size.
#
#tcp_recv_bufsize 0 bytes

# TAG: err_html_text
# HTML text to include in error messages. Make this a "mailto"
# URL to your admin address, or maybe just a link to your
# organizations Web page.
#
# To include this in your error messages, you must rewrite
# the error template files (found in the "errors" directory).
# Wherever you want the 'err_html_text' line to appear,
# insert a %L tag in the error template file.
#err_html_text

# TAG: deny_info
# Usage: deny_info err_page_name acl
# Example: deny_info ERR_CUSTOM_ACCESS_DENIED bad_guys
#
# This can be used to return a ERR_ page for requests which
# do not pass the 'http_access' rules. A single ACL will cause
# the http_access check to fail. If a 'deny_info' line exists
# for that ACL then Squid returns a corresponding error page.
#
# You may use ERR_ pages that come with Squid or create your own
pages
# and put them into the configured errors/ directory.

# TAG: memory_pools on|off
# If set, Squid will keep pools of allocated (but unused) memory
# available for future use. If memory is a premium on your
# system and you believe your malloc library outperforms Squid
# routines, disable this.
#
#memory_pools on

# TAG: memory_pools_limit (bytes)
# Used only with memory_pools on:
# memory_pools_limit 50 MB
#
# If set to a non-zero value, Squid will keep at most the
specified
# limit of allocated (but unused) memory in memory pools. All
free()
# requests that exceed this limit will be handled by your malloc
# library. Squid does not pre-allocate any memory, just safe-keeps
# objects that otherwise would be free()d. Thus, it is safe to set
# memory_pools_limit to a reasonably high value even if your
# configuration will use less memory.
#
# If not set (default) or set to zero, Squid will keep all memory
it
```

```
# can. That is, there will be no limit on the total amount of
memory
# used for safe-keeping.
#
# To disable memory allocation optimization, do not set
# memory_pools_limit to 0. Set memory_pools to "off" instead.
#
# An overhead for maintaining memory pools is not taken into
account
# when the limit is checked. This overhead is close to four bytes
per
# object kept. However, pools may actually _save_ memory because
of
# reduced memory thrashing in your malloc library.

# TAG: forwarded_for    on|off
# If set, Squid will include your system's IP address or name
# in the HTTP requests it forwards. By default it looks like
# this:
#
# X-Forwarded-For: 192.1.2.3
#
# If you disable this, it will appear as
#
# X-Forwarded-For: unknown
#
#forwarded_for on

# TAG: log_icp_queries on|off
# If set, ICP queries are logged to access.log. You may wish
# do disable this if your ICP load is VERY high to speed things
# up or to simplify log analysis.
#
#log_icp_queries on

# TAG: icp_hit_stale on|off
# If you want to return ICP_HIT for stale cache objects, set this
# option to 'on'. If you have sibling relationships with caches
# in other administrative domains, this should be 'off'. If you
only
# have sibling relationships with caches under your control, then
# it is probably okay to set this to 'on'.
#
#icp_hit_stale off

# TAG: minimum_direct_hops
# If using the ICMP pinging stuff, do direct fetches for sites
# which are no more than this many hops away.
#
minimum_direct_hops 4

# TAG: cachemgr_passwd
# Specify passwords for cachemgr operations.
#
# Usage: cachemgr_passwd password action action ...
#
# Some valid actions are (see cache manager menu for a full list):
# 5min
```

```
# 60min
# asndb
# authenticator
# cbdata
# client_list
# comm_incoming
# config *
# counters
# delay
# digest_stats
# dns
# events
# filedescriptors
# fqdncache
# histograms
# http_headers
# info
# io
# ipcache
# mem
# menu
# netdb
# non_peers
# objects
# pconn
# peer_select
# redirector
# refresh
# server_list
# shutdown *
# store_digest
# storedir
# utilization
# via_headers
# vm_objects
#
# * Indicates actions which will not be performed without a
# valid password, others can be performed if not listed here.
#
# To disable an action, set the password to "disable".
# To allow performing an action without a password, set the
# password to "none".
#
# Use the keyword "all" to set the same password for all actions.
#
cachemgr_passwd clave all
#cachemgr_passwd secret shutdown
#cachemgr_passwd lesssssssecret info stats/objects
#cachemgr_passwd disable all

# TAG: store_avg_object_size (kbytes)
# Average object size, used to estimate number of objects your
# cache can hold. See doc/Release-Notes-1.1.txt. The default is
# 13 KB.
#
#store_avg_object_size 13 KB

# TAG: store_objects_per_bucket
```

```
# Target number of objects per bucket in the store hash table.
# Lowering this value increases the total number of buckets and
# also the storage maintenance rate. The default is 50.
#
#store_objects_per_bucket 50

# TAG: client_db on|off
# If you want to disable collecting per-client statistics, then
# turn off client_db here.
#
#client_db on

# TAG: netdb_low
# TAG: netdb_high
# The low and high water marks for the ICMP measurement
# database. These are counts, not percents. The defaults are
# 900 and 1000. When the high water mark is reached, database
# entries will be deleted until the low mark is reached.
#
#netdb_low 900
#netdb_high 1000

# TAG: netdb_ping_period
# The minimum period for measuring a site. There will be at
# least this much delay between successive pings to the same
# network. The default is five minutes.
#
#netdb_ping_period 5 minutes

# TAG: query_icmp on|off
# If you want to ask your peers to include ICMP data in their ICP
# replies, enable this option.
#
# If your peer has configured Squid (during compilation) with
# '--enable-icmp' then that peer will send ICMP pings to origin
server
# sites of the URLs it receives. If you enable this option then
the
# ICP replies from that peer will include the ICMP data (if
available).
# Then, when choosing a parent cache, Squid will choose the parent
with
# the minimal RTT to the origin server. When this happens, the
# hierarchy field of the access.log will be
# "CLOSEST_PARENT_MISS". This option is off by default.
#
#query_icmp off

# TAG: test_reachability     on|off
# When this is 'on', ICP MISS replies will be ICP_MISS_NOFETCH
# instead of ICP_MISS if the target host is NOT in the ICMP
# database, or has a zero RTT.
#
#test_reachability off

# TAG: buffered_logs    on|off
# Some log files (cache.log, useragent.log) are written with
# stdio functions, and as such they can be buffered or
```

```
# unbuffered. By default they will be unbuffered. Buffering them
# can speed up the writing slightly (though you are unlikely to
# need to worry).
#buffered_logs off

# TAG: reload_into_ims  on|off
# When you enable this option, client no-cache or ``reload''
# requests will be changed to If-Modified-Since requests.
# Doing this VIOLATES the HTTP standard. Enabling this
# feature could make you liable for problems which it
# causes.
#
# see also refresh_pattern for a more selective approach.
#
# This option may be disabled by using --disable-http-violations
# with the configure script.
#reload_into_ims off

# TAG: always_direct
# Usage: always_direct allow|deny [!]aclname ...
#
# Here you can use ACL elements to specify requests which should
# ALWAYS be forwarded directly to origin servers. For example,
# to always directly forward requests for local servers use
# something like:
#
# acl local-servers dstdomain my.domain.net
# always_direct allow local-servers
#
# To always forward FTP requests directly, use
#
# acl FTP proto FTP
# always_direct allow FTP
#
# NOTE: There is a similar, but opposite option named
# 'never_direct'. You need to be aware that "always_direct deny
# foo" is NOT the same thing as "never_direct allow foo". You
# may need to use a deny rule to exclude a more-specific case of
# some other rule. Example:
#
# acl local-external dstdomain external.foo.net
# acl local-servers dstdomain foo.net
# always_direct deny local-external
# always_direct allow local-servers
#
# This option replaces some v1.1 options such as local_domain
# and local_ip.

# TAG: never_direct # Usage: never_direct allow|deny [!]aclname
...
#
# never_direct is the opposite of always_direct. Please read
# the description for always_direct if you have not already.
#
# With 'never_direct' you can use ACL elements to specify
# requests which should NEVER be forwarded directly to origin
# servers. For example, to force the use of a proxy for all
# requests, except those in your local domain use something like:
```

```
#
# acl local-servers dstdomain foo.net
# acl all src 0.0.0.0/0.0.0.0
# never_direct deny local-servers
# never_direct allow all
#
# or if squid is inside a firewall and there is local intranet
# servers inside the firewall then use something like:
#
# acl local-intranet dstdomain foo.net
# acl local-external dstdomain external.foo.net
# always_direct deny local-external
# always_direct allow local-intranet
# never_direct allow all
#
# This option replaces some v1.1 options such as inside_firewall
# and firewall_ip.

# TAG: anonymize_headers
# Usage: anonymize_headers allow|deny header_name ...
#
# This option replaces the old 'http_anonymizer' option with
# something that is much more configurable. You may now
# specify exactly which headers are to be allowed, or which
# are to be removed from outgoing requests.
#
# There are two methods of using this option. You may either
# allow specific headers (thus denying all others), or you
# may deny specific headers (thus allowing all others).
#
# For example, to achieve the same behavior as the old
# 'http_anonymizer standard' option, you should use:
#
# anonymize_headers deny From Referer Server
# anonymize_headers deny User-Agent WWW-Authenticate Link
#
# Or, to reproduce the old 'http_anonymizer paranoid' feature
# you should use:
#
# anonymize_headers allow Allow Authorization Cache-Control
# anonymize_headers allow Content-Encoding Content-Length
# anonymize_headers allow Content-Type Date Expires Host
# anonymize_headers allow If-Modified-Since Last-Modified
# anonymize_headers allow Location Pragma Accept
# anonymize_headers allow Accept-Encoding Accept-Language
# anonymize_headers allow Content-Language Mime-Version
# anonymize_headers allow Retry-After Title Connection
# anonymize_headers allow Proxy-Connection
#
# NOTE: You can not mix "allow" and "deny". All
'anonymize_headers'
# lines must have the same second argument.
#
# By default, all headers are allowed (no anonymizing is
# performed).
#
#anonymize_headers
```

```
# TAG: fake_user_agent
# If you filter the User-Agent header with 'anonymize_headers' it
# may cause some Web servers to refuse your request. Use this to
# fake one up. For example:
#
# fake_user_agent Nutscrape/1.0 (CP/M; 8-bit)
# (credit to Paul Southworth pauls@etext.org for this one!)
#
#fake_user_agent none

# TAG: icon_directory
# Where the icons are stored. These are normally kept in
# /usr/lib/squid/icons # TAG: error_directory
# Directory where the error files are read from.
# /usr/lib/squid/errors contains sets of error files
# in different languages. The default error directory
# is /etc/squid/errors, which is a link to one of these
# error sets. # # If you wish to create your own versions of the
error files,
# either to customize them to suit your language or company,
# copy the template English files to another
# directory and point this tag at them.
#
#error_directory /etc/squid/errors

# TAG: minimum_retry_timeout (seconds)
# This specifies the minimum connect timeout, for when the
# connect timeout is reduced to compensate for the availability
# of multiple IP addresses.
#
# When a connection to a host is initiated, and that host has
# several IP addresses, the default connection timeout is reduced
# by dividing it by the number of addresses. So, a site with 15
# addresses would then have a timeout of 8 seconds for each
# address attempted. To avoid having the timeout reduced to the
# point where even a working host would not have a chance to
# respond, this setting is provided. The default, and the
# minimum value, is five seconds, and the maximum value is sixty
# seconds, or half of connect_timeout, whichever is greater and
# less than connect_timeout.
#
#minimum_retry_timeout 5 seconds

# TAG: maximum_single_addr_tries
# This sets the maximum number of connection attempts for a
# host that only has one address (for multiple-address hosts,
# each address is tried once).
#
# The default value is three tries, the (not recommended)
# maximum is 255 tries. A warning message will be generated
# if it is set to a value greater than ten.
#
#maximum_single_addr_tries 3

# TAG: snmp_port
# Squid can now serve statistics and status information via SNMP.
# By default it listens to port 3401 on the machine. If you don't
# wish to use SNMP, set this to "0".
```

```
#
# NOTE: SNMP support requires use the --enable-snmp configure
# command line option.
#snmp_port 3401

# TAG: snmp_access
# Allowing or denying access to the SNMP port.
#
# All access to the agent is denied by default.
# usage:
#
# snmp_access allow|deny [!]aclname ...
#
#Example:
#snmp_access allow snmppublic localhost
#snmp_access deny all

# TAG: snmp_incoming_address
# TAG: snmp_outgoing_address
# Just like 'udp_incoming_address' above, but for the SNMP port.
#
# snmp_incoming_address is used for the SNMP socket receiving
# messages from SNMP agents.
# snmp_outgoing_address is used for SNMP packets returned to SNMP
# agents.
#
# The default behavior is to not bind to any specific address.
#
# NOTE, snmp_incoming_address and snmp_outgoing_address can not
have
# the same value since they both use port 3130.
#
#snmp_incoming_address 0.0.0.0
#snmp_outgoing_address 0.0.0.0

# TAG: as_whois_server
# WHOIS server to query for AS numbers. NOTE: AS numbers are
# queried only when Squid starts up, not for every request.

# TAG: wccp_router
# Use this option to define your WCCP ``home'' router for
# Squid. Setting the 'wccp_router' to 0.0.0.0 (the default)
# disables WCCP.
#wccp_router 0.0.0.0

# TAG: wccp_version
# According to some users, Cisco IOS 11.2 only supports WCCP
# version 3. If you're using that version of IOS, change
# this value to 3. #wccp_version 4

# TAG: wccp_incoming_address
# TAG: wccp_outgoing_address
#     wccp_incoming_address   Use this option if you require WCCP
#                         messages to be received on only one
#                         interface. Do NOT use this option if
#                         you're unsure how many interfaces you
#                         have, or if you know you have only one
#                         interface.
```

```
#
#     wccp_outgoing_address   Use this option if you require WCCP
#                         messages to be sent out on only one
#                         interface. Do NOT use this option if
#                         you're unsure how many interfaces you
#                         have, or if you know you have only one
#                         interface.
#
#     The default behavior is to not bind to any specific address.
#
#     NOTE, wccp_incoming_address and wccp_outgoing_address can
not have
#     the same value since they both use port 2048.
#
#wccp_incoming_address 0.0.0.0
#wccp_outgoing_address 0.0.0.0

# DELAY POOL PARAMETERS (all require DELAY_POOLS compilation
option)
# -----------------------------------------------------------------
-------------

# TAG: delay_pools
# This represents the number of delay pools to be used. For
example,
# if you have one class 2 delay pool and one class 3 delays pool,
you
# have a total of 2 delay pools.
#
# To enable this option, you must use --enable-delay-pools with
the
# configure script.
#delay_pools 0

# TAG: delay_class
# This defines the class of each delay pool. There must be exactly
one
# delay_class line for each delay pool. For example, to define two
# delay pools, one of class 2 and one of class 3, the settings
above
# and here would be:
#
#delay_pools 2          # 2 delay pools
#delay_class 1 2        # pool 1 is a class 2 pool
#delay_class 2 3        # pool 2 is a class 3 pool
#
# The delay pool classes are:
#
#     class 1     Everything is limited by a single aggregate
#           bucket.
#
#     class 2     Everything is limited by a single aggregate
#           bucket as well as an "individual" bucket chosen
#           from bits 25 through 32 of the IP address.
#
#     class 3     Everything is limited by a single aggregate
#           bucket as well as a "network" bucket chosen
#           from bits 17 through 24 of the IP address and a
```

```
#           "individual" bucket chosen from bits 17 through
#           32 of the IP address.
#
# NOTE: If an IP address is a.b.c.d
#     -> bits 25 through 32 are "d"
#      -> bits 17 through 24 are "c"
#      -> bits 17 through 32 are "c * 256 + d"

# TAG: delay_access
# This is used to determine which delay pool a request falls into.
# The first matched delay pool is always used, i.e., if a request
falls
# into delay pool number one, no more delay are checked, otherwise
the
# rest are checked in order of their delay pool number until they
have
# all been checked. For example, if you want some_big_clients in
delay
# pool 1 and lotsa_little_clients in delay pool 2:
#
#delay_access 1 allow some_big_clients
#delay_access 1 deny all
#delay_access 2 allow lotsa_little_clients
#delay_access 2 deny all

# TAG: delay_parameters
# This defines the parameters for a delay pool. Each delay pool
has
# a number of "buckets" associated with it, as explained in the
# description of delay_class. For a class 1 delay pool, the syntax
is:
#
#delay_parameters pool aggregate
#
#     For a class 2 delay pool:
#
#delay_parameters pool aggregate individual
#
#     For a class 3 delay pool:
#
#delay_parameters pool aggregate network individual
#
#     The variables here are:
#
#      pool       a pool number - ie, a number between 1 and the
#           number specified in delay_pools as used in
#           delay_class lines.
#
#     aggregate the "delay parameters" for the aggregate bucket #
(class 1, 2, 3).
#
#     individual the "delay parameters" for the individual #
buckets (class 2, 3).
#
#     network the "delay parameters" for the network buckets #
(class 3).
#
```

```
# A pair of delay parameters is written restore/maximum, where
restore is
# the number of bytes (not bits - modem and network speeds are
usually
# quoted in bits) per second placed into the bucket, and maximum
is the
# maximum number of bytes which can be in the bucket at any time.
#
# For example, if delay pool number 1 is a class 2 delay pool as
in the
# above example, and is being used to strictly limit each host to
64kbps
# (plus overheads), with no overall limit, the line is:
#
#delay_parameters 1 -1/-1 8000/8000
#
# Note that the figure -1 is used to represent "unlimited".
#
# And, if delay pool number 2 is a class 3 delay pool as in the
above
# example, and you want to limit it to a total of 256kbps (strict
limit)
# with each 8-bit network permitted 64kbps (strict limit) and each
# individual host permitted 4800bps with a bucket maximum size of
64kb
# to permit a decent web page to be downloaded at a decent speed
# (if the network is not being limited due to overuse) but slow
down
# large downloads more significantly:
#
#delay_parameters 2 32000/32000 8000/8000 600/64000
#
# There must be one delay_parameters line for each delay pool.

# TAG: delay_initial_bucket_level (percent, 0-100)
# The initial bucket percentage is used to determine how much is
put
# in each bucket when squid starts, is reconfigured, or first
notices
# a host accessing it (in class 2 and class 3, individual hosts
and
# networks only have buckets associated with them once they have
been
# "seen" by squid).
#
#delay_initial_bucket_level 50

# TAG: incoming_icp_average
# TAG: incoming_http_average
# TAG: min_icp_poll_cnt # TAG: min_http_poll_cnt
# Heavy voodoo here. I can't even believe you are reading this.
# Are you crazy? Don't even think about adjusting these unless
# you understand the algorithms in comm_select.c first!
#
#incoming_icp_average 6
#incoming_http_average 4
#min_icp_poll_cnt 8
#min_http_poll_cnt 8
```

```
# TAG: max_open_disk_fds
# TAG: offline_mode
# Enable this option and Squid will never try to validate cached
# objects.

# TAG: uri_whitespace
# What to do with requests that have whitespace characters in the
#     URI.  Options:
#
#     strip:      The whitespace characters are stripped out of
the URL.
#           This is the behavior recommended by RFC2616.
#     deny:       The request is denied. The user receives an
"Invalid
#     Request" message.
#     allow:      The request is allowed and the URI is not
changed. The
#           whitespace characters remain in the URI. Note the
#           whitespace is passed to redirector processes if they
#           are in use.
#     encode: The request is allowed and the whitespace characters
are
#           encoded according to RFC1738. This could be considered
#           a violation of the HTTP/1.1
#           RFC because proxies are not allowed to rewrite URI's.
#     chop:       The request is allowed and the URI is chopped at
the
#           first whitespace. This might also be considered a
#           violation.
#uri_whitespace strip

# TAG: broken_posts
# A list of ACL elements which, if matched, causes Squid to send
# a extra CRLF pair after the body of a PUT/POST request.
#
# Some HTTP servers has broken implementations of PUT/POST,
# and rely on a extra CRLF pair sent by some WWW clients.
#
# Quote from RFC 2068 section 4.1 on this matter:
#
# Note: certain buggy HTTP/1.0 client implementations generate an
# extra CRLF's after a POST request. To restate what is explicitly
# forbidden by the BNF, an HTTP/1.1 client must not preface or
follow
# a request with an extra CRLF.
#
#acl buggy_server url_regex ^http://....
#broken_posts allow buggy_server

# TAG: mcast_miss_addr
# If you enable this option, every "cache miss" URL will
# be sent out on the specified multicast address.
#
# Do not enable this option unless you are are absolutely
# certain you understand what you are doing.
# TAG: mcast_miss_ttl
# This is the time-to-live value for packets multicasted
# when multicasting off cache miss URLs is enabled. By
```

```
# default this is set to 'site scope', i.e. 16.

# TAG: mcast_miss_port
# This is the port number to be used in conjunction with
# 'mcast_miss_addr'.

# TAG: mcast_miss_encode_key
# The URLs that are sent in the multicast miss stream are
# encrypted. This is the encryption key.

# TAG: prefer_direct
# By default, if the ICP, HTCP, Cache Digest, etc. techniques
# do not yield a parent cache, Squid gives higher preference
# to forwarding the request direct to origin servers, rather
# than selecting a parent cache anyway.
#
# If you want Squid to give higher precedence to a parent
# cache, instead of going direct, then turn this option off.
#prefer_direct on

# TAG: strip_query_terms
# By default, Squid strips query terms from requested URLs before
# logging. This protects your user's privacy.
#strip_query_terms on

# TAG: coredump_dir
# By default Squid leaves core files in the first cache_dir
# directory. If you set 'coredump_dir' to a directory
# that exists, Squid will chdir() to that directory at startup
# and coredump files will be left there.

# TAG: redirector_bypass
# When this is 'on', a request will not go through the
# redirector if all redirectors are busy. If this is 'off'
# and the redirector queue grows too large, Squid will exit
# with a FATAL error and ask you to increase the number of
# redirectors. You should only enable this if the redirectors
# are not critical to your caching system. If you use
# redirectors for access control, and you enable this option,
# then users may have access to pages that they should not
# be allowed to request.

# TAG: ignore_unknown_nameservers
# By default Squid checks that DNS responses are received
# from the same IP addresses that they are sent to. If they
# don't match, Squid ignores the response and writes a warning
# message to cache.log. You can allow responses from unknown
# nameservers by setting this option to 'off'.
#ignore_unknown_nameservers on

# TAG: digest_generation
# This controls whether the server will generate a Cache Digest
# of its contents. By default, Cache Digest generation is
# enabled if Squid is compiled with USE_CACHE_DIGESTS defined.
#digest_generation on

# TAG: digest_bits_per_entry
# This is the number of bits of the server's Cache Digest which
```

```
# will be associated with the Digest entry for a given HTTP
# Method and URL (public key) combination. The default is 5.
#digest_bits_per_entry 5

# TAG: digest_rebuild_period  (seconds)
# This is the number of seconds between Cache Digest rebuilds.
# By default the server's Digest is rebuilt every hour.
#digest_rebuild_period 1 hour

# TAG: digest_rewrite_period  (seconds)
# This is the number of seconds between Cache Digest writes to
# disk. By default the server's Digest is written to disk every
# hour. #digest_rewrite_period 1 hour

# TAG: digest_swapout_chunk_size     (bytes)
# This is the number of bytes of the Cache Digest to write to
# disk at a time. It defaults to 4096 bytes (4KB), the Squid
# default swap page.
#digest_swapout_chunk_size 4096 bytes

# TAG: digest_rebuild_chunk_percentage (percent, 0-100)
# This is the percentage of the Cache Digest to be scanned at a
# time. By default it is set to 10% of the Cache Digest.
#digest_rebuild_chunk_percentage 10

# TAG: chroot
# Use this to have Squid do a chroot() while initializing. This
# also causes Squid to fully drop root privileges after
# initializing. This means, for example, that if you use a HTTP
# port less than 1024 and try to reconfigure, you will get an
# error.

# TAG: client_persistent_connections
# TAG: server_persistent_connections
# Persistent connection support for clients and servers. By
# default, Squid uses persistent connections (when allowed)
# with its clients and servers. You can use these options to
# disable persistent connections with clients and/or servers.
#client_persistent_connections on
#server_persistent_connections on
```

A continuación se describen cada uno de los parám etros utilizados en la instalación:

http_port. Este parámetro define en que puerto responderá a las solicitudes S **quid**. Para esta instalación utilizamos el puerto 3128:

http_port 3128

icp_port. Este parámetro define el puerto en que el servidor Squid re cibe solicitudes ICP (Inter-Cache Protocol). Desactívelo asignando el valor de cero al parámetro:

icp_port 0

cache_mem. Memoria utilizada por Squid para ciertos procesos. En esta instalación para un computador de 64MB de memoria utilizamos 8 MB.

cache_mem 8 MB

cache_swap_low. Indica el nivel en porcentaje de capacidad mínima aceptada por Squid, es decir, los objetos se mantendrán en el caché hasta que se cope el límite mínimo:

cache_swap_low 90

cache_swap_high. Parámetro que especifica en porcentaje el límite máximo que utiliza Squid para mantener objetos en el caché. Si el valor asignado es del 95%, Squid comenzara a eliminar los objetos del caché cuando se tope el 95% de la capacidad asignada a Squid:

cache_swap_high 95

maximum_object_size. Este parámetro, especificado en KB, indica el tamaño máximo que se almacena en el caché. Por defecto se utiliza 4MB:

maximum_object_size 4096 KB

cache_dir. Directorio de ubicación del caché, por defecto /usr/local/squid/cache. Este parámetro incluye tres parámetros numéricos adicionales. El primero incluye el número de MB que se utilizarán en este directorio para el caché, por defecto 100MB, el segundo el número de directorios a utilizar en el primer nivel (16 por defecto) y el tercero el número de subdirectorios en el segundo nivel (256 por defecto):

cache_dir ufs /var/spool/squid/cache 100 16 256

cache_access_log. Especifica en que directorio se realizará el registro de accesos al Squid.

cache_access_log /var/log/squid/access.log

cache_log. Define en donde se almacenan los mensajes del sistema:

cache_log /var/log/squid/cache.log

cache_store_log. Este parámetro especifica la ubicación del archivo de registro de objetos sacados del cache. No es necesario activarlo. Es mejor desactivarlo para ahorrar espacio en disco:

cache_store_log none

emulate_httpd_log. Este parámetro define si se desea utilizar emulación de logs del servidor Web (httpd).

emulate_httpd_log on

mime_table. Define la ubicación del archivo mime.conf, se utiliza el valor por defecto:

mime_table /etc/squid/mime.conf

pid_filename. Define la ubicación del archivo squid.pid, se utiliza el valor por defecto:

pid_filename /var/log/squid/squid.pid

debug_options. Opciones de depuración, se utiliza el valor por defecto:

debug_options ALL,1

ALL,1

reference_age. Este parámetro determina cuanto tiempo permanece el objeto en el caché, en este caso se utiliza un mes.

reference_age 1 month

quick_abort. Este parámetro define si un objeto debe almacenarse en el caché cuando el usuario ha interrumpido una solicitud: si el objeto tiene el valor especificado en **min** o falta más del valor especificado en **max** se abortará la transferencia. Si se ha realizado una transferencia mayor del valor en porcentaje especificado en **pct**, no se abortará el almacenamiento del objeto. Se recomienda utilizar los valores por defecto:

quick_abort_min16KB
quick_abort_max16KB
quick_abort_pct 95

negative_ttl. Este parámetro se utiliza para definir cuanto tiempo debe esperar Squid para procesar nuevamente una página que no ha sido encontrada. En este caso se utilizaron 5 minutos.

negative_ttl 5 minutes

positive_dns_ttl. Este parámetro especifica el tiempo que Squid mantendrá la dirección de un sitio visitado exitosamente. El valor por defecto es de 6 horas:

positive_dns_ttl 6 hours

negative_dns_ttl. Especifica el tiempo que espera Squid antes de intentar nuevamente determinar la dirección de un sitio solicitado y que no ha sido encontrado. Por defecto 5 minutos:

negative_dns_ttl 5 minutes

En los acl de la sección ACCESS CONTROLS se adiciona

acl localhost src 127.0.0.1/255.2.55.255.255
acl mired src 172.16.0.0/255.255.0.0

http_access e **icp_access**. Define una serie de permisos para acceso al servidor Squid. ecomendamos utilizar la siguiente configuración:

http_access allow manager localhost
http_access deny manager
http_access deny ¡mired
http_access allow CONNECT !SSL_ports
http_access allow all
http_access allow mi red
icp_access allow all
miss_access allow all

cache_mgr. Definición del administrador del sistema.

cache_mgr webmaster

cache_effective_user y **cache_effective_group**. Estos dos parámetros definen que usuario (user) y grupo (group) ejecuta Squid. La versión de RPM utiliza el usuario Squid y grupo Squid. La versión de las fuentes utiliza el usuario nobo dy y grupo nogroup RPM:

cache_effective_user squid
cache_effective_group squid

Fuentes:
cache_effective_user nobody
cache_effective_group nogroup

Es importante verificar que el grupo **nogroup** exista en su sistema mirando el archivo group ubicado en el directorio /etc. Este archivo debe tener una entrada como se muestra a continuación:

nogroup:x:500:

El número que aparece al final de la línea debe ser único y corresponder al siguiente número disponible en su servidor. Por ejemplo, si la última línea del arc hivo /etc/group es:

slocate:x:21:

Utilice un número para este grupo superior a 21 y cerciórese que no se encuentre ya utilizado por otro grupo en este archivo. Como en el ejemplo anterior recomendamos que utilice el número 500.

El usuario nobody existe en las distribuciones de Linux.

visible_hostname. Este parámetro define el nombre del servidor. Coloque aquí el nombre de su servidor:

visible_hostname minotauro

httpd_accel_uses_host_header. Este parámetro se utiliza para activar el *proxy* transparente, necesario para controlar el acceso a Internet desde las estaciones.

httpd_accel_uses_host_header on

httpd_accel_host y **httpd_accel_port**. Este parámetro también es necesario para activar el *proxy* transparente. Al configurar Squid de esta manera no es necesar io realizar configuración de los navegadores del centro de acceso con servidor proxy lo que se constituye en una ventaja:

httpd_accel_host virtual
httpd_accel_port 80

httpd_accel_with_proxy. Este parámetro es necesario activarlo para el proxy transparente:

httpd_accel_with_proxy on

dns_testnames. Este parámetro define que hosts se utilizan para chequear el servidor de nombres:

dns_testnames netscape.com internic.net nlanr.net microsoft.com

append_domain. Este parámetro indica a Squid que dominio debe a ñadirse a solicitudes que vengan sin dominio completo. Es recomendable colocar el dominio de su Telecentro (nótese que debe comenzar con un punto):

append_domain .minotauro.telemed.org

cachemgr_passwd. Squid es posible ser administrado desde una página W eb una vez instalado (no documentado aquí). Este parámetro define la clave de acceso para tener acceso a esta función:

cachemgr_passwd clave all

4-.Creación del cache.

Antes de utilizar Squid por primera vez es necesario crear el directorio cache ejecut ando para la versión de RPM:
 # /usr/sbin/squid -z

y para la versión de fuentes:

 # /usr/local/squid/bin/squid -z

En la versión de fuentes es importante asignar los permisos del directorio de registro (*logs*) al usuario *nobody* y grupo *nogroup*:

 # cd /usr/local/squid/
 # chmod nobody.nogroup logs

5-. Ejecución del Squid.

5.1-. Inicio del servidor para la instalación de RPM.

Una vez instalado Squid para que se inicie cada vez que se prenda el servidor se puede colocar al final del archivo /etc/rc.d/rc.local la línea de comando /etc/init.d/squid start:

 # cd /etc/rc.d
 # pico rc.local

ir al final del archivo y añadir:
/etc/init.d/squid start
o Usted puede ejecutar este comando manualmente.

5.2-. Inicio del servidor squid para la instalación de las fuentes:

Una vez instalado Squid para que se inicie cada vez que se prenda el servidor se puede colocar al final del archivo /etc/rc.d/rc.local la línea de comando /usr/local/squid/bin/squid:

 # cd /etc/rc.d
 # pico rc.local

ir al final del archivo y añadir:

/usr/local/squid/bin/squid

o usted puede ejecutar este comando manualmente

F.4.2 SERVIDOR APACHE

El servidor Apache es el servicio que se encarga de resolver las peticiones de páginas de Internet de los clientes utilizando el protocolo de Internet http.

A continuación se presentan las instrucciones de instalación del Apache en el servidor Linux. Es muy posible que su instalación de Linux ya haya realizado la instalación del servidor por lo que se presenta un apartado para determinar si Apache ya se encuentr a instalado.
Posteriormente se presentan las instrucciones de instalación para la versión completa y para la versión RPM.

Si Usted dispone del Apache en RPM se recomienda que realice la instalación por este mecanismo.

1-. Comprobar si ya está instalado Apache

Si se instaló una distribución de Linux como RedHat, Conectiva, Corel Linux, Suse, Mandrake, TurboLinux, Slackware, etc, es posible que ya se tenga instalado Apache. Para verificar si está ya instalado utilice el comando whereis así:

whereis httpd

httpd es el programa que ejecuta Apache (http *deamon*) Si el sistema responde:

whereis httpd
httpd:

Apache no está instalado en el servidor.
Si responde:
whereis httpd
httpd: /usr/local/bin/httpd
Apache está instalado en el servidor y el ejecutab le se encuentra en /usr/local/bin/httpd.

También puede utilizarse el siguiente comando para verificar si el servidor está instalado:
httpd -v
Server version: Apache/1.3.12 (Unix) (Red Hat/Linux)
Server built: Aug 23 2000 15:44:50
Si el sistema responde como se presenta en el anterior ejemplo el servidor se encuentra ya instalado. Adicionalmente se muestra la versión y la fecha de instalación.

En este caso busque el archivo **httpd.conf** y añada los parámetros requeridos para la instalación del Programa de Registro.

2-. Instalación de Apache a través de RPM.

RPM es la sigla para El Sistema de Manejo de Paquetes de Red Hat (*Red Hat Package Manager* - RPM).

La instalación a través de RPM es muy sencilla. Simplemente busque en su distribución de Linux el paquete apache y ejecute la instalación con el comando **rpm**.
En Red Hat 7.0 este se encuentra ubicado en el CD-ROM 1 en el archivo RedHat/RPMS/apache-1.3.12-25.i386.rpm.

Si se introduce el CD-ROM en la unidad del servidor este deberá activarse automáticamente, es decir, que deberá poder acceder al CD-ROM y al programa Apache de la siguiente manera:

cd /mnt/cdrom/RedHat/RPMS/
rpm -iv apache-1.3.12-25.i386.rpm

Este comando instalará la versión de Apache en el servidor. El parámetro *i* corresponde a instalar (*install*), el parámetro *v*, corresponde a mostrar información de la instalación (*verbose*).

3-. Instalación de Apache desde las fuentes. Se sugiere bajar la última versión de apache desde http://httpd.apache.org. En esta página se busca el enlace **"Download Apache 1.3"** y posteriormente se selecciona el archivo **apache_1.3.19.tar.gz** para copiar la distribución al computador, o se selecciona la última versión disponible.

Posteriormente se descompacta el archivo con la opción:

tar zxvf apache_1.3.19.tar.gz

Con este comando se creará el directorio
apache_1.3.19

se ingresa al directorio:
cd apache_1.3.19

Y se ejecuta la instalación así (se deberá esperar a que el sistema ejecute las tareas de configuración y compilación):

./configure --enable-module=so
make
make install
Apache quedará instalado en el directorio
/usr/local/apache

4-. Configuración de Apache.

Después de realizar la instalación es necesario configurar el servidor editando el archivo de configuración **httpd.conf**. Este archivo puede encontrarse en uno de los siguientes directorios:
/usr/local/apache/conf
o
/etc/httpd/conf

para este caso se encuentra en /usr/local/apache/conf

En este archivo se definen los parámetros del sistema. Cada parámetro po see un nombre al comienzo de la línea y su valor (o valores separados por espacios **pero en una sola línea!**) en frente separado por al menos un espacio:
nombre_del_parametro valor
nombre_del_parametro valor1 valor2 valor3 valorn

Todas las líneas que comienzan en este archivo con el signo # corresponden a comentarios dentro del archivo de configuración

Se edita el archivo **httpd.conf** con su editor preferido, por ejemplo *vi* o *pico*:

Para instalación **desde las fuentes** seguramente es necesario acceder el siguie nte directorio:

cd /usr/local/apache/conf
o, para instalación por **RPM**:
cd /etc/httpd/conf

pico httpd.conf

El arcivo de configuración es httpd.conf y se encuentra en el fichero /usr/local/apache/conf/httpd.conf. este archivo se muestra a continuació n:

```
##
## httpd.conf -- Apache HTTP server configuration file
##

#
# Based upon the NCSA server configuration files originally by Rob
McCool.
#
# This is the main Apache server configuration file. It contains
the
# configuration directives that give the server its instructions.
# See <URL:http://www.apache.org/docs/> for detailed information
about
# the directives.
#
# Do NOT simply read the instructions in here without
understanding
# what they do. They're here only as hints or reminders. If you
are unsure
# consult the online docs. You have been warned.
#
# After this file is processed, the server will look for and
process
# /usr/local/apache/conf/srm.conf and then
/usr/local/apache/conf/access.conf
# unless you have overridden these with ResourceConfig and/or
# AccessConfig directives here.
#
# The configuration directives are grouped into three basic
sections:
# 1. Directives that control the operation of the Apache server
process as a
# whole (the 'global environment').
# 2. Directives that define the parameters of the 'main' or
'default' server,
# which responds to requests that aren't handled by a virtual
host.
# These directives also provide default values for the settings
# of all virtual hosts.
# 3. Settings for virtual hosts, which allow Web requests to be
sent to
# different IP addresses or hostnames and have them handled by the
# same Apache server process.
#
# Configuration and logfile names: If the filenames you specify
for many
# of the server's control files begin with "/" (or "drive:/" for
Win32), the
# server will use that explicit path. If the filenames do *not*
```

```
begin
# with "/", the value of ServerRoot is prepended -- so
"logs/foo.log"
# with ServerRoot set to "/usr/local/apache" will be interpreted
by the
# server as "/usr/local/apache/logs/foo.log".
#

### Section 1: Global Environment
#
# The directives in this section affect the overall operation of
Apache,
# such as the number of concurrent requests it can handle or where
it
# can find its configuration files.
#

#
# ServerType is either inetd, or standalone. Inetd mode is only
supported on
# Unix platforms.
#
ServerType standalone

#
# ServerRoot: The top of the directory tree under which the
server's
# configuration, error, and log files are kept.
#
# NOTE! If you intend to place this on an NFS (or otherwise
network)
# mounted filesystem then please read the LockFile documentation
# (available at
<URL:http://www.apache.org/docs/mod/core.html#lockfile>);
# you will save yourself a lot of trouble.
#
ServerRoot "/usr/local/apache"

#
# The LockFile directive sets the path to the lockfile used when
Apache
# is compiled with either USE_FCNTL_SERIALIZED_ACCEPT or
# USE_FLOCK_SERIALIZED_ACCEPT. This directive should normally be
left at
# its default value. The main reason for changing it is if the
logs
# directory is NFS mounted, since the lockfile MUST BE STORED ON A
LOCAL
# DISK. The PID of the main server process is automatically
appended to
# the filename.
#
#LockFile /usr/local/apache/logs/httpd.lock

#
# PidFile: The file in which the server should record its process
# identification number when it starts.
#
```

```
PidFile /usr/local/apache/logs/httpd.pid

#
# ScoreBoardFile: File used to store internal server process
information.
# Not all architectures require this. But if yours does (you'll
know because
# this file will be created when you run Apache) then you *must*
ensure that
# no two invocations of Apache share the same scoreboard file.
#
ScoreBoardFile /usr/local/apache/logs/httpd.scoreboard

#
# In the standard configuration, the server will process
httpd.conf (this
# file, specified by the -f command line option), srm.conf, and
access.conf
# in that order. The latter two files are now distributed empty,
as it is
# recommended that all directives be kept in a single file for
simplicity.
# The commented-out values below are the built-in defaults. You
can have the
# server ignore these files altogether by using "/dev/null" (for
Unix) or
# "nul" (for Win32) for the arguments to the directives.
#
#ResourceConfig conf/srm.conf
#AccessConfig conf/access.conf

#
# Timeout: The number of seconds before receives and sends time
out.
#
Timeout 300

#
# KeepAlive: Whether or not to allow persistent connections (more
than
# one request per connection). Set to "Off" to deactivate.
#
KeepAlive On

#
# MaxKeepAliveRequests: The maximum number of requests to allow
# during a persistent connection. Set to 0 to allow an unlimited
amount.
# We recommend you leave this number high, for maximum
performance.
#
MaxKeepAliveRequests 100

#
# KeepAliveTimeout: Number of seconds to wait for the next request
from the
# same client on the same connection.
#
```

```
KeepAliveTimeout 15

#
# Server-pool size regulation. Rather than making you guess how
many
# server processes you need, Apache dynamically adapts to the load
it
# sees --- that is, it tries to maintain enough server processes
to
# handle the current load, plus a few spare servers to handle
transient
# load spikes (e.g., multiple simultaneous requests from a single
# Netscape browser).
#
# It does this by periodically checking how many servers are
waiting
# for a request. If there are fewer than MinSpareServers, it
creates
# a new spare. If there are more than MaxSpareServers, some of the
# spares die off. The default values are probably OK for most
sites.
#
MinSpareServers 5
MaxSpareServers 10

#
# Number of servers to start initially --- should be a reasonable
ballpark
# figure.
#
StartServers 5

#
# Limit on total number of servers running, i.e., limit on the
number
# of clients who can simultaneously connect --- if this limit is
ever
# reached, clients will be LOCKED OUT, so it should NOT BE SET TOO
LOW.
# It is intended mainly as a brake to keep a runaway server from
taking
# the system with it as it spirals down...
#
MaxClients 150

#
# MaxRequestsPerChild: the number of requests each child process
is
# allowed to process before the child dies. The child will exit so
# as to avoid problems after prolonged use when Apache (and maybe
the
# libraries it uses) leak memory or other resources. On most
systems, this
# isn't really needed, but a few (such as Solaris) do have notable
leaks
# in the libraries. For these platforms, set to something like
10000
# or so; a setting of 0 means unlimited.
```

```
#
# NOTE: This value does not include keepalive requests after the
initial
# request per connection. For example, if a child process handles
# an initial request and 10 subsequent "keptalive" requests, it
# would only count as 1 request towards this limit.
#
MaxRequestsPerChild 0

#
# Listen: Allows you to bind Apache to specific IP addresses
and/or
# ports, in addition to the default. See also the <VirtualHost>
# directive.
#
#Listen 3000
#Listen 12.34.56.78:80

#
# BindAddress: You can support virtual hosts with this option.
This directive
# is used to tell the server which IP address to listen to. It can
either
# contain "*", an IP address, or a fully qualified Internet domain
name.
# See also the <VirtualHost> and Listen directives.
#
#BindAddress *

#
# Dynamic Shared Object (DSO) Support
#
# To be able to use the functionality of a module which was built
as a DSO you
# have to place corresponding `LoadModule' lines at this location
so the
# directives contained in it are actually available _before_ they
are used.
# Please read the file README.DSO in the Apache 1.3 distribution
for more
# details about the DSO mechanism and run `httpd -l' for the list
of already
# built-in (statically linked and thus always available) modules
in your httpd
# binary.
#
# Note: The order in which modules are loaded is important. Don't
change
# the order below without expert advice.
#
# Example:
# LoadModule foo_module libexec/mod_foo.so
LoadModule env_module libexec/mod_env.so
LoadModule config_log_module libexec/mod_log_config.so
LoadModule mime_module libexec/mod_mime.so
LoadModule negotiation_module libexec/mod_negotiation.so
LoadModule status_module libexec/mod_status.so
LoadModule includes_module libexec/mod_include.so
```

```
LoadModule autoindex_module libexec/mod_autoindex.so
LoadModule dir_module libexec/mod_dir.so
LoadModule cgi_module libexec/mod_cgi.so
LoadModule asis_module libexec/mod_asis.so
LoadModule imap_module libexec/mod_imap.so
LoadModule action_module libexec/mod_actions.so
LoadModule userdir_module libexec/mod_userdir.so
LoadModule alias_module libexec/mod_alias.so
LoadModule access_module libexec/mod_access.so
LoadModule auth_module libexec/mod_auth.so
LoadModule setenvif_module libexec/mod_setenvif.so
LoadModule php4_module libexec/libphp4.so

# Reconstruction of the complete module list from all available
modules
# (static and shared ones) to achieve correct module execution
order.
# [WHENEVER YOU CHANGE THE LOADMODULE SECTION ABOVE UPDATE THIS,
TOO]
ClearModuleList
AddModule mod_env.c
AddModule mod_log_config.c
AddModule mod_mime.c
AddModule mod_negotiation.c
AddModule mod_status.c
AddModule mod_include.c
AddModule mod_autoindex.c
AddModule mod_dir.c
AddModule mod_cgi.c
AddModule mod_asis.c
AddModule mod_imap.c
AddModule mod_actions.c
AddModule mod_userdir.c
AddModule mod_alias.c
AddModule mod_access.c
AddModule mod_auth.c
AddModule mod_so.c
AddModule mod_setenvif.c
AddModule mod_php4.c

#
# ExtendedStatus controls whether Apache will generate "full"
status
# information (ExtendedStatus On) or just basic information
(ExtendedStatus
# Off) when the "server-status" handler is called. The default is
Off.
#
#ExtendedStatus On

### Section 2: 'Main' server configuration
#
# The directives in this section set up the values used by the
'main'
# server, which responds to any requests that aren't handled by a
# <VirtualHost> definition. These values also provide defaults for
# any <VirtualHost> containers you may define later in the file.
#
```

```
# All of these directives may appear inside <VirtualHost>
containers,
# in which case these default settings will be overridden for the
# virtual host being defined.
#

#
# If your ServerType directive (set earlier in the 'Global
Environment'
# section) is set to "inetd", the next few directives don't have
any
# effect since their settings are defined by the inetd
configuration.
# Skip ahead to the ServerAdmin directive.
#

#
# Port: The port to which the standalone server listens. For
# ports < 1023, you will need httpd to be run as root initially.
#
Port 80

#
# If you wish httpd to run as a different user or group, you must
run
# httpd as root initially and it will switch.
#
# User/Group: The name (or #number) of the user/group to run httpd
as.
# . On SCO (ODT 3) use "User nouser" and "Group nogroup".
# . On HPUX you may not be able to use shared memory as nobody,
and the
# suggested workaround is to create a user www and use that user.
# NOTE that some kernels refuse to setgid(Group) or
semctl(IPC_SET)
# when the value of (unsigned)Group is above 60000;
# don't use Group nobody on these systems!
#
User nobody
Group nobody

#
# ServerAdmin: Your address, where problems with the server should
be
# e-mailed. This address appears on some server-generated pages,
such
# as error documents.
#
ServerAdmin root@minotauro.telemed.org

#
# ServerName allows you to set a host name which is sent back to
clients for
# your server if it's different than the one the program would get
(i.e., use
# "www" instead of the host's real name).
#
# Note: You cannot just invent host names and hope they work. The
```

```
name you
# define here must be a valid DNS name for your host. If you don't
understand
# this, ask your network administrator.
# If your host doesn't have a registered DNS name, enter its IP
address here.
# You will have to access it by its address (e.g.,
http://123.45.67.89/)
# anyway, and this will make redirections work in a sensible way.
#
# 127.0.0.1 is the TCP/IP local loop-back address, often named
localhost. Your
# machine always knows itself by this address. If you use Apache
strictly for
# local testing and development, you may use 127.0.0.1 as the
server name.
#
#ServerName minotauro.telemed.org

#
# DocumentRoot: The directory out of which you will serve your
# documents. By default, all requests are taken from this
directory, but
# symbolic links and aliases may be used to point to other
locations.
#
#DocumentRoot "/usr/local/apache/htdocs"

DocumentRoot "/telemed"

#
# Each directory to which Apache has access, can be configured
with respect
# to which services and features are allowed and/or disabled in
that
# directory (and its subdirectories).
#
# First, we configure the "default" to be a very restrictive set
of
# permissions.
#
<Directory />
Options FollowSymLinks
AllowOverride None
</Directory>

#
# Note that from this point forward you must specifically allow
# particular features to be enabled - so if something's not
working as
# you might expect, make sure that you have specifically enabled
it
# below.
#

#
# This should be changed to whatever you set DocumentRoot to.
#
```

```
<Directory "/usr/local/apache/htdocs">

#
# This may also be "None", "All", or any combination of "Indexes",
# "Includes", "FollowSymLinks", "ExecCGI", or "MultiViews".
#
# Note that "MultiViews" must be named *explicitly* --- "Options
All"
# doesn't give it to you.
#
Options Indexes FollowSymLinks MultiViews

#
# This controls which options the .htaccess files in directories
can
# override. Can also be "All", or any combination of "Options",
"FileInfo",
# "AuthConfig", and "Limit"
#
AllowOverride None

#
# Controls who can get stuff from this server.
#
Order allow,deny
Allow from all
</Directory>

#
# UserDir: The name of the directory which is appended onto a
user's home
# directory if a ~user request is received.
#
<IfModule mod_userdir.c>
UserDir web
</IfModule>

#
# Control access to UserDir directories. The following is an
example
# for a site where these directories are restricted to read-only.
#
#<Directory /home/*/public_html>
# AllowOverride FileInfo AuthConfig Limit
# Options MultiViews Indexes SymLinksIfOwnerMatch IncludesNoExec
# <Limit GET POST OPTIONS PROPFIND>
# Order allow,deny
# Allow from all
# </Limit>
# <LimitExcept GET POST OPTIONS PROPFIND>
# Order deny,allow
# Deny from all
# </LimitExcept>
#</Directory>

#
# DirectoryIndex: Name of the file or files to use as a pre-
written HTML
```

```
# directory index. Separate multiple entries with spaces.
#
<IfModule mod_dir.c>
DirectoryIndex index.html index.htm index.php
</IfModule>

#
# AccessFileName: The name of the file to look for in each
directory
# for access control information.
#
AccessFileName .htaccess

#
# The following lines prevent .htaccess files from being viewed by
# Web clients. Since .htaccess files often contain authorization
# information, access is disallowed for security reasons. Comment
# these lines out if you want Web visitors to see the contents of
# .htaccess files. If you change the AccessFileName directive
above,
# be sure to make the corresponding changes here.
#
# Also, folks tend to use names such as .htpasswd for password
# files, so this will protect those as well.
#
<Files ~ "^\.ht">
Order allow,deny
Deny from all
</Files>

#
# CacheNegotiatedDocs: By default, Apache sends "Pragma: no-cache"
with each
# document that was negotiated on the basis of content. This asks
proxy
# servers not to cache the document. Uncommenting the following
line disables
# this behavior, and proxies will be allowed to cache the
documents.
#
#CacheNegotiatedDocs

#
# UseCanonicalName: (new for 1.3) With this setting turned on,
whenever
# Apache needs to construct a self-referencing URL (a URL that
refers back
# to the server the response is coming from) it will use
ServerName and
# Port to form a "canonical" name. With this setting off, Apache
will
# use the hostname:port that the client supplied, when possible.
This
# also affects SERVER_NAME and SERVER_PORT in CGI scripts.
#
UseCanonicalName On

#
```

```
# TypesConfig describes where the mime.types file (or equivalent)
is
# to be found.
#
<IfModule mod_mime.c>
TypesConfig /usr/local/apache/conf/mime.types
</IfModule>

#
# DefaultType is the default MIME type the server will use for a
document
# if it cannot otherwise determine one, such as from filename
extensions.
# If your server contains mostly text or HTML documents,
"text/plain" is
# a good value. If most of your content is binary, such as
applications
# or images, you may want to use "application/octet-stream"
instead to
# keep browsers from trying to display binary files as though they
are
# text.
#
DefaultType text/plain

#
# The mod_mime_magic module allows the server to use various hints
from the
# contents of the file itself to determine its type. The
MIMEMagicFile
# directive tells the module where the hint definitions are
located.
# mod_mime_magic is not part of the default server (you have to
add
# it yourself with a LoadModule [see the DSO paragraph in the
'Global
# Environment' section], or recompile the server and include
mod_mime_magic
# as part of the configuration), so it's enclosed in an <IfModule>
container.
# This means that the MIMEMagicFile directive will only be
processed if the
# module is part of the server.
#
<IfModule mod_mime_magic.c>
MIMEMagicFile /usr/local/apache/conf/magic
</IfModule>

#
# HostnameLookups: Log the names of clients or just their IP
addresses
# e.g., www.apache.org (on) or 204.62.129.132 (off).
# The default is off because it'd be overall better for the net if
people
# had to knowingly turn this feature on, since enabling it means
that
# each client request will result in AT LEAST one lookup request
to the
```

```
# nameserver.
#
HostnameLookups Off

#
# ErrorLog: The location of the error log file.
# If you do not specify an ErrorLog directive within a
<VirtualHost>
# container, error messages relating to that virtual host will be
# logged here. If you *do* define an error logfile for a
<VirtualHost>
# container, that host's errors will be logged there and not here.
#
ErrorLog /usr/local/apache/logs/error_log

#
# LogLevel: Control the number of messages logged to the
error_log.
# Possible values include: debug, info, notice, warn, error, crit,
# alert, emerg.
#
LogLevel warn

#
# The following directives define some format nicknames for use
with
# a CustomLog directive (see below).
#
LogFormat "%h %l %u %t \"%r\" %>s %b \"%{Referer}i\" \"%{User-
Agent}i\"" combined
LogFormat "%h %l %u %t \"%r\" %>s %b" common
LogFormat "%{Referer}i -> %U" referer
LogFormat "%{User-agent}i" agent

#
# The location and format of the access logfile (Common Logfile
Format).
# If you do not define any access logfiles within a <VirtualHost>
# container, they will be logged here. Contrariwise, if you *do*
# define per-<VirtualHost> access logfiles, transactions will be
# logged therein and *not* in this file.
#
CustomLog /usr/local/apache/logs/access_log common

#
# If you would like to have agent and referer logfiles, uncomment
the
# following directives.
#
#CustomLog /usr/local/apache/logs/referer_log referer
#CustomLog /usr/local/apache/logs/agent_log agent

#
# If you prefer a single logfile with access, agent, and referer
information
# (Combined Logfile Format) you can use the following directive.
#
#CustomLog /usr/local/apache/logs/access_log combined
```

```
#
# Optionally add a line containing the server version and virtual
host
# name to server-generated pages (error documents, FTP directory
listings,
# mod_status and mod_info output etc., but not CGI generated
documents).
# Set to "EMail" to also include a mailto: link to the
ServerAdmin.
# Set to one of: On | Off | EMail
#
ServerSignature On

# EBCDIC configuration:
# (only for mainframes using the EBCDIC codeset, currently one of:
# Fujitsu-Siemens' BS2000/OSD, IBM's OS/390 and IBM's TPF)!!
# The following default configuration assumes that "text files"
# are stored in EBCDIC (so that you can operate on them using the
# normal POSIX tools like grep and sort) while "binary files" are
# stored with identical octets as on an ASCII machine.
#
# The directives are evaluated in configuration file order, with
# the EBCDICConvert directives applied before EBCDICConvertByType.
#
# If you want to have ASCII HTML documents and EBCDIC HTML
documents
# at the same time, you can use the file extension to force
# conversion off for the ASCII documents:
# > AddType text/html .ahtml
# > EBCDICConvert Off=InOut .ahtml
#
# EBCDICConvertByType On=InOut text/* message/* multipart/*
# EBCDICConvertByType On=In application/x-www-form-urlencoded
# EBCDICConvertByType On=InOut application/postscript model/vrml
# EBCDICConvertByType Off=InOut */*


#
# Aliases: Add here as many aliases as you need (with no limit).
The format is
# Alias fakename realname
#
<IfModule mod_alias.c>

#
# Note that if you include a trailing / on fakename then the
server will
# require it to be present in the URL. So "/icons" isn't aliased
in this
# example, only "/icons/". If the fakename is slash-terminated,
then the
# realname must also be slash terminated, and if the fakename
omits the
# trailing slash, the realname must also omit it.
#
Alias /icons/ "/usr/local/apache/icons/"
```

```
<Directory "/usr/local/apache/icons">
Options Indexes MultiViews
AllowOverride None
Order allow,deny
Allow from all
</Directory>

#
# ScriptAlias: This controls which directories contain server
scripts.
# ScriptAliases are essentially the same as Aliases, except that
# documents in the realname directory are treated as applications
and
# run by the server when requested rather than as documents sent
to the client.
# The same rules about trailing "/" apply to ScriptAlias
directives as to
# Alias.
#
ScriptAlias /cgi-bin/ "/usr/local/apache/cgi-bin/"

#
# "/usr/local/apache/cgi-bin" should be changed to whatever your
ScriptAliased
# CGI directory exists, if you have that configured.
#
<Directory "/usr/local/apache/cgi-bin">
AllowOverride None
Options None
Order allow,deny
Allow from all
</Directory>

</IfModule>
# End of aliases.

#
# Redirect allows you to tell clients about documents which used
to exist in
# your server's namespace, but do not anymore. This allows you to
tell the
# clients where to look for the relocated document.
# Format: Redirect old-URI new-URL
#

#
# Directives controlling the display of server-generated directory
listings.
#
<IfModule mod_autoindex.c>

#
# FancyIndexing is whether you want fancy directory indexing or
standard
#
#
# AddIcon* directives tell the server which icon to show for
different
```

```
# files or filename extensions. These are only displayed for
# FancyIndexed directories.
#
AddIconByEncoding (CMP,/icons/compressed.gif) x-compress x-gzip

AddIconByType (TXT,/icons/text.gif) text/*
AddIconByType (IMG,/icons/image2.gif) image/*
AddIconByType (SND,/icons/sound2.gif) audio/*
AddIconByType (VID,/icons/movie.gif) video/*

AddIcon /icons/binary.gif .bin .exe
AddIcon /icons/binhex.gif .hqx
AddIcon /icons/tar.gif .tar
AddIcon /icons/world2.gif .wrl .wrl.gz .vrml .vrm .iv
AddIcon /icons/compressed.gif .Z .z .tgz .gz .zip
AddIcon /icons/a.gif .ps .ai .eps
AddIcon /icons/layout.gif .html .shtml .htm .pdf
AddIcon /icons/text.gif .txt
AddIcon /icons/c.gif .c
AddIcon /icons/p.gif .pl .py
AddIcon /icons/f.gif .for
AddIcon /icons/dvi.gif .dvi
AddIcon /icons/uuencoded.gif .uu
AddIcon /icons/script.gif .conf .sh .shar .csh .ksh .tcl
AddIcon /icons/tex.gif .tex
AddIcon /icons/bomb.gif core

AddIcon /icons/back.gif ..
AddIcon /icons/hand.right.gif README
AddIcon /icons/folder.gif ^^DIRECTORY^^
AddIcon /icons/blank.gif ^^BLANKICON^^

#
# DefaultIcon is which icon to show for files which do not have an
icon
# explicitly set.
#
DefaultIcon /icons/unknown.gif

#
# AddDescription allows you to place a short description after a
file in
# server-generated indexes. These are only displayed for
FancyIndexed
# directories.
# Format: AddDescription "description" filename
#
#AddDescription "GZIP compressed document" .gz
#AddDescription "tar archive" .tar
#AddDescription "GZIP compressed tar archive" .tgz

#
# ReadmeName is the name of the README file the server will look
for by
# default, and append to directory listings.
#
# HeaderName is the name of a file which should be prepended to
# directory indexes.
```

```
#
# If MultiViews are amongst the Options in effect, the server will
# first look for name.html and include it if found. If name.html
# doesn't exist, the server will then look for name.txt and
include
# it as plaintext if found.
#
ReadmeName README
HeaderName HEADER

#
# IndexIgnore is a set of filenames which directory indexing
should ignore
# and not include in the listing. Shell-style wildcarding is
permitted.
#
IndexIgnore .??* *~ *# HEADER* README* RCS CVS *,v *,t

</IfModule>
# End of indexing directives.

#
# Document types.
#
<IfModule mod_mime.c>

#
# AddEncoding allows you to have certain browsers (Mosaic/X 2.1+)
uncompress
# information on the fly. Note: Not all browsers support this.
# Despite the name similarity, the following Add* directives have
nothing
# to do with the FancyIndexing customization directives above.
#
AddEncoding x-compress Z
AddEncoding x-gzip gz tgz

#
# AddLanguage allows you to specify the language of a document.
You can
# then use content negotiation to give a browser a file in a
language
# it can understand.
#
# Note 1: The suffix does not have to be the same as the language
# keyword --- those with documents in Polish (whose net-standard
# language code is pl) may wish to use "AddLanguage pl .po" to
# avoid the ambiguity with the common suffix for perl scripts.
#
# Note 2: The example entries below illustrate that in quite
# some cases the two character 'Language' abbreviation is not
# identical to the two character 'Country' code for its country,
# E.g. 'Danmark/dk' versus 'Danish/da'.
#
# Note 3: In the case of 'ltz' we violate the RFC by using a three
char
# specifier. But there is 'work in progress' to fix this and get
# the reference data for rfc1766 cleaned up.
```

```
#
# Danish (da) - Dutch (nl) - English (en) - Estonian (ee)
# French (fr) - German (de) - Greek-Modern (el)
# Italian (it) - Korean (kr) - Norwegian (no)
# Portugese (pt) - Luxembourgeois* (ltz)
# Spanish (es) - Swedish (sv) - Catalan (ca) - Czech(cz)
# Polish (pl) - Brazilian Portuguese (pt-br) - Japanese (ja)
# Russian (ru)
#
AddLanguage da .dk
AddLanguage nl .nl
AddLanguage en .en
AddLanguage et .ee
AddLanguage fr .fr
AddLanguage de .de
AddLanguage el .el
AddLanguage he .he
AddCharset ISO-8859-8 .iso8859-8
AddLanguage it .it
AddLanguage ja .ja
AddCharset ISO-2022-JP .jis
AddLanguage kr .kr
AddCharset ISO-2022-KR .iso-kr
AddLanguage no .no
AddLanguage pl .po
AddCharset ISO-8859-2 .iso-pl
AddLanguage pt .pt
AddLanguage pt-br .pt-br
AddLanguage ltz .lu
AddLanguage ca .ca
AddLanguage es .es
AddLanguage sv .se
AddLanguage cz .cz
AddLanguage ru .ru
AddLanguage zh-tw .tw
AddLanguage tw .tw
AddCharset Big5 .Big5 .big5
AddCharset WINDOWS-1251 .cp-1251
AddCharset CP866 .cp866
AddCharset ISO-8859-5 .iso-ru
AddCharset KOI8-R .koi8-r
AddCharset UCS-2 .ucs2
AddCharset UCS-4 .ucs4
AddCharset UTF-8 .utf8

# LanguagePriority allows you to give precedence to some languages
# in case of a tie during content negotiation.
#
# Just list the languages in decreasing order of preference. We
have
# more or less alphabetized them here. You probably want to change
this.
#
<IfModule mod_negotiation.c>
LanguagePriority en da nl et fr de el it ja kr no pl pt pt-br ru
ltz ca es sv tw
</IfModule>
```

```
#
# AddType allows you to tweak mime.types without actually editing
it, or to
# make certain files to be certain types.
#
# For example, the PHP 3.x module (not part of the Apache
distribution - see
# http://www.php.net) will typically use:
#
#AddType application/x-httpd-php3 .php3
#AddType application/x-httpd-php3-source .phps
#
# And for PHP 4.x, use:
#
AddType application/x-httpd-php .php .php3
#AddType application/x-httpd-php-source .phps

AddType application/x-tar .tgz

#
# AddHandler allows you to map certain file extensions to
"handlers",
# actions unrelated to filetype. These can be either built into
the server
# or added with the Action command (see below)
#
# If you want to use server side includes, or CGI outside
# ScriptAliased directories, uncomment the following lines.
#
# To use CGI scripts:
#
#AddHandler cgi-script .cgi

#
# To use server-parsed HTML files
#
#AddType text/html .shtml
#AddHandler server-parsed .shtml

#
# Uncomment the following line to enable Apache's send-asis HTTP
file
# feature
#
#AddHandler send-as-is asis

#
# If you wish to use server-parsed imagemap files, use
#
#AddHandler imap-file map

#
# To enable type maps, you might want to use
#
#AddHandler type-map var

</IfModule>
# End of document types.
```

```
#
# Action lets you define media types that will execute a script
whenever
# a matching file is called. This eliminates the need for repeated
URL
# pathnames for oft-used CGI file processors.
# Format: Action media/type /cgi-script/location
# Format: Action handler-name /cgi-script/location
#

#
# MetaDir: specifies the name of the directory in which Apache can
find
# meta information files. These files contain additional HTTP
headers
# to include when sending the document
#
#MetaDir .web

#
# MetaSuffix: specifies the file name suffix for the file
containing the
# meta information.
#
#MetaSuffix .meta

#
# Customizable error response (Apache style)
# these come in three flavors
#
# 1) plain text
#ErrorDocument 500 "The server made a boo boo.
# n.b. the single leading (") marks it as text, it does not get
output
#
# 2) local redirects
#ErrorDocument 404 /missing.html
# to redirect to local URL /missing.html
#ErrorDocument 404 /cgi-bin/missing_handler.pl
# N.B.: You can redirect to a script or a document using server-
side-includes.
#
# 3) external redirects
#ErrorDocument 402
http://some.other_server.com/subscription_info.html
# N.B.: Many of the environment variables associated with the
original
# request will *not* be available to such a script.

#
# Customize behaviour based on the browser
#
<IfModule mod_setenvif.c>

#
# The following directives modify normal HTTP response behavior.
# The first directive disables keepalive for Netscape 2.x and
```

```
browsers that
# spoof it. There are known problems with these browser
implementations.
# The second directive is for Microsoft Internet Explorer 4.0b2
# which has a broken HTTP/1.1 implementation and does not properly
# support keepalive when it is used on 301 or 302 (redirect)
responses.
#
BrowserMatch "Mozilla/2" nokeepalive
BrowserMatch "MSIE 4\.0b2;" nokeepalive downgrade-1.0 force-
response-1.0

#
# The following directive disables HTTP/1.1 responses to browsers
which
# are in violation of the HTTP/1.0 spec by not being able to grok
a
# basic 1.1 response.
#
BrowserMatch "RealPlayer 4\.0" force-response-1.0
BrowserMatch "Java/1\.0" force-response-1.0
BrowserMatch "JDK/1\.0" force-response-1.0

</IfModule>
# End of browser customization directives

#
# Allow server status reports, with the URL of
http://servername/server-status
# Change the ".your_domain.com" to match your domain to enable.
#
#<Location /server-status>
# SetHandler server-status
# Order deny,allow
# Deny from all
# Allow from .your_domain.com
#</Location>

#
# Allow remote server configuration reports, with the URL of
# http://servername/server-info (requires that mod_info.c be
loaded).
# Change the ".your_domain.com" to match your domain to enable.
#
#<Location /server-info>
# SetHandler server-info
# Order deny,allow
# Deny from all
# Allow from .your_domain.com
#</Location>

#
# There have been reports of people trying to abuse an old bug
from pre-1.1
# days. This bug involved a CGI script distributed as a part of
Apache.
# By uncommenting these lines you can redirect these attacks to a
logging
```

```
# script on phf.apache.org. Or, you can record them yourself,
using the script
# support/phf_abuse_log.cgi.
#
#<Location /cgi-bin/phf*>
# Deny from all
# ErrorDocument 403 http://phf.apache.org/phf_abuse_log.cgi
#</Location>

#
# Proxy Server directives. Uncomment the following lines to
# enable the proxy server:
#
#<IfModule mod_proxy.c>
# ProxyRequests On

# <Directory proxy:*>
# Order deny,allow
# Deny from all
# Allow from .your_domain.com
# </Directory>

#
# Enable/disable the handling of HTTP/1.1 "Via:" headers.
# ("Full" adds the server version; "Block" removes all outgoing
Via: headers)
# Set to one of: Off | On | Full | Block
#
# ProxyVia On

#
# To enable the cache as well, edit and uncomment the following
lines:
# (no cacheing without CacheRoot)
#
# CacheRoot "/usr/local/apache/proxy"
# CacheSize 5
# CacheGcInterval 4
# CacheMaxExpire 24
# CacheLastModifiedFactor 0.1
# CacheDefaultExpire 1
# NoCache a_domain.com another_domain.edu joes.garage_sale.com

#</IfModule>
# End of proxy directives.

### Section 3: Virtual Hosts
#
# VirtualHost: If you want to maintain multiple domains/hostnames
on your
# machine you can setup VirtualHost containers for them. Most
configurations
# use only name-based virtual hosts so the server doesn't need to
worry about
# IP addresses. This is indicated by the asterisks in the
directives below.
#
# Please see the documentation at
```

```
<URL:http://www.apache.org/docs/vhosts/>
# for further details before you try to setup virtual hosts.
#
# You may use the command line option '-S' to verify your virtual
host
# configuration.

#
# Use name-based virtual hosting.
#
NameVirtualHost 172.16.130.75
NameVirtualHost 200.21.83.96

#
# VirtualHost example:
# Almost any Apache directive may go into a VirtualHost container.
# The first VirtualHost section is used for requests without a
known
# server name.
#
#<VirtualHost *>
# ServerAdmin webmaster@dummy-host.example.com
# DocumentRoot /www/docs/dummy-host.example.com
# ServerName dummy-host.example.com
# ErrorLog logs/dummy-host.example.com-error_log
# CustomLog logs/dummy-host.example.com-access_log common
#</VirtualHost>
<VirtualHost 172.16.130.75>
ServerAdmin webmaster@dummy-host.example.com
DocumentRoot /home/jmmarti/web/
ServerName www.telemed.org
ErrorLog logs/dummy-host.example.com-error_log
CustomLog logs/dummy-host.example.com-access_log common
</VirtualHost>
<VirtualHost 200.21.83.96>
ServerAdmin webmaster@dummy-host.example.com
DocumentRoot /home/jmmarti/web/
ServerName www.telemed.org
ErrorLog logs/dummy-host.example.com-error_log
CustomLog logs/dummy-host.example.com-access_log common
</VirtualHost>
```

Opciones de índices:

En las opciones que aparecen para el parámetro DirectoryIndex es importante añadir los valores **indindex.htm**, **index.php** . Esta sección debe quedar así:

DirectoryIndex index.html index.htm index.php

Es posible que esta definición se encuentre en la definición de parámetros de la sección **IfModule_mod_dir** (en el caso de la distribución de fuentes de Apache), en este caso deberá quedar así:

<IfModule mod_dir.c>
DirectoryIndex index.html index.htm index.php
</IfModule>

HostnameLookups. Si se dispone de un enlace a Internet permanente, activando este parámetro será posible realizar posteriormente análisis estadísticos que indicarán desde que partes de la red son consultadas las páginas que están alojadas en el servidor.

Esta opción puede hacer que **el acceso a las páginas del Sistema de Registro sea lento** si las estaciones no se encuentran registradas en el DNS de su dominio. Si Usted tiene dudas le recomendamos que deje este parámetro como **off**:

HostnameLookups Off

Salve el archivo de configuración e **inicie el servidor**.
Si la instalación de Apache se realizó desde RPM (o estaba previamente instalada) puede iniciar el servidor así:

/etc/rc.d/init.d/httpd start

Para la instalación desde las fuentes ejecute:

 # /usr/local/apache/bin/apachectl start

El servidor deberá responder:
/usr/local/apache/bin/apachectl start: httpd started

Para que Apache se inicie cada vez que se prenda el servidor se puede colocar al final del archivo **/etc/rc.d/rc.local** la línea de comando que se ejecutó arriba.
Esto no es necesario si se realizó la instalación desde RPM.
Abra un navegador y escriba el nombre de su servidor o la dirección IP que se especificó en el archivo de configuración y deberá poder ver la página de prueba de Apache que indica:

¡Funcionó! ¡El Servidor de Red Apache ha sido instalado en ese sitio!

o en inglés:

Test Page

F.4.3 SERVIDOR DNS

El DNS (Domain Name System)constituye la base de datos que proporciona a sus clientes la información acerca de las direcciones y los nombres de la red. El DNS tiene una orientación cliente servidor y todas las máquinas que qued en bajo él pueden ser consideradas como archivos en una estructura de árbol de directorio. Este nuevo sistema es una base de datos distribuida conocida como BIND (Nombres de Dominio de Internet de Berkeley). El servicio de nombres en linux, es llevado a ca bo por un demonio llamado named, el cual está incluido en la mayoría de las distribuciones de linux y queda ubicado en el directorio /usr/sbin.

Tipos de servidores DNS

Los servidores DNS se dividen en tres grupos: primarios, secundarios y de chaché. Los servidores primarios son a los únicos a los que se les considera autorizados para un dominio en particular. Un servidor autorizado es el único en el cual residen los archivos de configuración del dominio.

Los servidores secundarios trabajan como respaldo y como distribuidores de carga de los servidores de nombres primarios. Los servidores primarios conocen la existencia de los secundarios y les envían periódicamente actualizaciones de sus tablas.

Los servidores de caché no contienen archivos de configur ación de ningún dominio. En su lugar, cuando una maquina cliente realiza una petición a un servidor de caché para que resuelva un nombre, este servidor comprueba su propio caché local primero. Si no lo encuentra buscará un servidor primario y le preguntará . Su respuesta pasará a caché.

Para el caso del servidor de la *RTPSTT* se hizo la configuración de un servidor de nombres de chaché y un servidor primario, además se utilizaròn dos direcciones en el servidor: una dirección real para darle salida a los eq uipos que se coneten al servidor y una dirección de Intranet. Para el primer caso, la dirección fue 200.21.83.96 y para el segundo 172.16.130.75. La configuración del equipo para aceptar esta configuración se hace por medio del *linuxconf* , seleccionando *configuración de red*, *tareas como servidor*, *alias IP para el dispositivi etho, clic sobre etho* y añadir lo siguiente
alias IP para el dispositivo etho: 200.21.83.86
màscara (opc) : 255.255.255.192

Se debe especificar el gateway por defecto a utilizar, e n el *linuxconf*, *tareas como cliente, rutas y gateways* , *especificar valores por defecto ,* el gateway por defecto serà: 200.21.83.126

Después de realizar estos cambios es necesario reiniciar la red para que los cambios surjan efecto. Ejecute entonces la in strucción:

/etc/rc.d/init.d/network restart

si se desea verificar los cambios realizados en la configuración de red se ejecuta la instrucción **ifconfig** y si se desea conocer la ruta la instrucción **route.**

De otra manera, para que al iniciar los servicios se tome la dirección gateway por defecto en etc/rc.local por medio de un editor se añade la siguiente linea:

/etc/rc.d/init.d/network restart

Es necesario también cambiar la configuración básica del equipo por medio del *linuxconf, tareas como cliente, DNS* se añade : IP del servidor de nombres: 172.16.130.75
y dominio de búsqueda : telemed.org

Los archivos de configuración son los siguientes:

Lo primero que se necesita para hacer la configuración del servidor de caché es el archivo llamado /etc/named.boot. Este archivo es leído cuando se inicia named que es el demonio para el DNS. Su configuración es la siguiente:

```
;
; a caching only nameserver config
;
directory                    /var/named
cache              .                root.cache
primary          0.0.127.in-addr.arpa    named.local
primary          130.16.172.in-addr.arpa       130.16.172
primary          telemed.org             telemed.org
```

La línea Directory indica a named dónde buscar los archivos.

A continuación se crea el fichero root.cache que se nombra en boot.named en los directorios /var/named/, es decir que la ruta es /var/named/root.cache, este archivo contiene lo siguiente:

```
.       518400      NS      D.ROOT-SERVERS.NET.
.       518400      NS      E.ROOT-SERVERS.NET.
.       518400      NS      I.ROOT-SERVERS.NET.
.       518400      NS      F.ROOT-SERVERS.NET.
.       518400      NS      G.ROOT-SERVERS.NET.
.       518400      NS      A.ROOT-SERVERS.NET.
.       518400      NS      H.ROOT-SERVERS.NET.
.       518400      NS      B.ROOT-SERVERS.NET.
.       518400      NS      C.ROOT-SERVERS.NET.
;
D.ROOT-SERVERS.NET.     3600000     A     128.8.10.90
E.ROOT-SERVERS.NET.     3600000     A     192.203.230.10
I.ROOT-SERVERS.NET.     3600000     A     192.36.148.17
F.ROOT-SERVERS.NET.     3600000     A     192.5.5.241
G.ROOT-SERVERS.NET.     3600000     A     192.112.36.4
A.ROOT-SERVERS.NET.     3600000     A     198.41.0.4
H.ROOT-SERVERS.NET.     3600000     A     128.63.2.53
B.ROOT-SERVERS.NET.     3600000     A     128.9.0.107
C.ROOT-SERVERS.NET.     3600000     A     192.33.4.12
```

Este archivo contiene los servidores de nombres raíz en el mundo. Este archivo cambiará a lo largo del tiempo y tiene que ser actualizado.

Posteriormente se crea un archivo llamado local host.zone en /var/named

```
$TTL    86400
$ORIGIN localhost.
@                       1D IN SOA       @ root (
                           42; serial (d. adams)
                           3H; refresh
                           15M; retry
                           1W; expiry
                           1D ); minimum

1D              IN              NS
@
1D      IN              A               127.0.0.1
```

A continuación se requiere el archivo /etc/resolv.conf que aparece como:

```
search telemed.org
nameserver 172.16.130.75
```

La línea "search" especifica en que dominios se buscaría para cualquier nombre de máquina a ña que quiera conectar. La línea "nameserver" especifica la dirección del servidor de nombres, ya que es allí donde named se estará ejecutando.

Ahora se requiere mirar el archivo /etc/nsswitch.conf el cual especifica de donde se obtienen las diferentes clases de tipos de datos, y de cual archivo o base de datos. Se busca la línea que comienza por hosts, debe leerse:

```
#
# /etc/nsswitch.conf
#
# An example Name Service Switch config file. This file should be
# sorted with the most-used services at the beginning.
#
# The entry '[NOTFOUND=return]' means that the search for an
# entry should stop if the search in the previous entry turned
# up nothing. Note that if the search failed due to some other
reason
# (like no NIS server responding) then the search continues with
the
# next entry.
#
# Legal entries are:
#
#     nisplus or nis+   Use NIS+ (NIS version 3)
#     nis or yp   Use NIS (NIS version 2), also called YP
#     dns         Use DNS (Domain Name Service)
#     files             Use the local files
#     db          Use the local database (.db) files
#     compat            Use NIS on compat mode
```

```
#     hesiod            Use Hesiod for user lookups
#     [NOTFOUND=return]       Stop searching if not found so far
#
# To use db, put the "db" in front of "files" for entries you want
to be
# looked up first in the databases
#
# Example:
#passwd:    db files nisplus nis
#shadow:    db files nisplus nis
#group:           db files nisplus nis

passwd:           files nisplus nis
shadow:     files nisplus nis
group:            files nisplus nis

#hosts:           db files nisplus nis dns
hosts:            files nisplus nis dns

# Example - obey only what nisplus tells us...
#services: nisplus [NOTFOUND=return] files
#networks: nisplus [NOTFOUND=return] files
#protocols: nisplus [NOTFOUND=return] files
#rpc: nisplus [NOTFOUND=return] files
#ethers: nisplus [NOTFOUND=return] files
#netmasks: nisplus [NOTFOUND=return] files

bootparams: nisplus [NOTFOUND=return] files

ethers: files
netmasks: files
networks: files
protocols: files nisplus nis
rpc: files
services: files nisplus nis

netgroup: files nisplus nis

publickey: nisplus

automount: files nisplus nis
aliases: files nisplus
```

Si no hay una línea que comience por “hosts” es necesario colocarla. Esto indica que los programas deben mirar primero en el fichero /etc/hosts y después comprobar el DNS de acuerdo a resolv.conf.

Ahora se configura el fichero /etc/ host.conf. cuya configuración es:

```
order hosts, bind
```

Si no existe línea order se debe incluir, ya que esta indica a las rutinas de resolución de nombres que busquen primero en /etc/hosts y después en el servidor de nombres.

El fichero /etc/hosts es el siguiente.

```
127.0.0.1       localhost.localdomain   localhost
172.16.130.75   minotauro.telemed.org   minotauro
```

Después de todo esto, se puede iniciar named presionando el comando

```
ndc start
```

se puede comprobar si funciona con el comando nslookup y debe salir esto:

$ nslookup
Default Server: localhost
Address: 127.0.0.1

Ahora se configura el servidor de nombres para el dominio de la RTPSTT:

Inicialmente se configura el archivo /etc/named.conf como se muestra:

```
// generated by named-bootconf.pl

options {
      directory "/var/named";
      /*
       *If there is a firewall between you and nameservers you
want
       *to talk to, you might need to uncomment the query-source
       *directive below. Previous versions of BIND always asked
       *questions using port 53, but BIND 8.1 uses an unprivileged
       *port by default.
       */
      // query-source address * port 53;
};

//
// a caching only nameserver config
//
zone "." {
            type hint;
            file "named.ca";
 };

zone "0.0.127.in-addr.arpa" {
```

```
            type master;
            file "named.local";
            // allow-update { none; };
 };

zone "telemed.org" {
            type master;
            file "telemed.org";
 };

zone "130.16.172.in-addr.arpa" {
            type master;
            file "172.16.130";
                                    };
```

Se va al archivo localhosts localizado en /etc/named/localhosts.zone y se configura como se mostró anteriormente.

Ahora se ve el archivo /var/named/named.ca, este archivo generalmente no se edita.

```
; This file holds the information on root name servers needed to
; initialize cache of Internet domain name servers
; (e.g. reference this file in the "cache . <file>"
; configuration file of BIND domain name servers).
;
; This file is made available by InterNIC registration services
; under anonymous FTP as
; file /domain/named.root
; on server FTP.RS.INTERNIC.NET
; -OR- under Gopher at RS.INTERNIC.NET
; under menu InterNIC Registration Services (NSI)
; submenu InterNIC Registration Archives
; file named.root
;
; last update: Aug 22, 1997
; related version of root zone: 1997082200
;
;
; formerly NS.INTERNIC.NET
;
. 3600000 IN NS A.ROOT-SERVERS.NET.
A.ROOT-SERVERS.NET. 3600000 A 198.41.0.4
;
; formerly NS1.ISI.EDU
;
. 3600000 NS B.ROOT-SERVERS.NET.
B.ROOT-SERVERS.NET. 3600000 A 128.9.0.107
;
; formerly C.PSI.NET
;
. 3600000 NS C.ROOT-SERVERS.NET.
C.ROOT-SERVERS.NET. 3600000 A 192.33.4.12
;
; formerly TERP.UMD.EDU
```

```
;
. 3600000 NS D.ROOT-SERVERS.NET.
D.ROOT-SERVERS.NET. 3600000 A 128.8.10.90
;
; formerly NS.NASA.GOV
;
. 3600000 NS E.ROOT-SERVERS.NET.
E.ROOT-SERVERS.NET. 3600000 A 192.203.230.10
;
; formerly NS.ISC.ORG
;
. 3600000 NS F.ROOT-SERVERS.NET.
F.ROOT-SERVERS.NET. 3600000 A 192.5.5.241
;
; formerly NS.NIC.DDN.MIL
;
. 3600000 NS G.ROOT-SERVERS.NET.
G.ROOT-SERVERS.NET. 3600000 A 192.112.36.4
;
; formerly AOS.ARL.ARMY.MIL
;
. 3600000 NS H.ROOT-SERVERS.NET.
H.ROOT-SERVERS.NET. 3600000 A 128.63.2.53
;
; formerly NIC.NORDU.NET
;
. 3600000 NS I.ROOT-SERVERS.NET.
I.ROOT-SERVERS.NET. 3600000 A 192.36.148.17
;
; temporarily housed at NSI (InterNIC)
;
. 3600000 NS J.ROOT-SERVERS.NET.
J.ROOT-SERVERS.NET. 3600000 A 198.41.0.10
;
; housed in LINX, operated by RIPE NCC
;
. 3600000 NS K.ROOT-SERVERS.NET.
K.ROOT-SERVERS.NET. 3600000 A 193.0.14.129
;
; temporarily housed at ISI (IANA)
;
. 3600000 NS L.ROOT-SERVERS.NET.
L.ROOT-SERVERS.NET. 3600000 A 198.32.64.12
;
; housed in Japan, operated by WIDE
;
. 3600000 NS M.ROOT-SERVERS.NET.
M.ROOT-SERVERS.NET. 3600000 A 202.12.27.33
; End of File
```

Posteriormente se crea el fichero /etc/named/telemed.org a continuación se muestra su contenido.

```
$TTL        86400
@           IN SOA      minotauro.telemed.org
      root.minotauro.telemed.org (
                                        42 ; serial (d. adams)
                                       3H ; refresh
                                       15M ; retry
                                       1W ; expiry
                                       1D ) ; minimum

@              IN             NS     minotauro
@              IN             MX     5     minotauro

minotauro      IN             A      172.16.130.75
www         IN          CNAME       minotauro
telemed           IN           CNAME       minotauro
www            IN       CNAME       minotauro
ftp         IN          CNAME       minotauro
```

Posteriormente se crea el fichero /var/named/130.16.75 como se muestra:

```
$TTL    86400
@   IN SOA     minotauro.telemed.org root.minotauro.telemed.org  (
                                          2001070515 ; Serial
                                          28800      ; Refresh
                                          14400      ; Retry
                                          3600000    ; Expire
                                          86400 )    ; Minimum
@       IN       NS       minotauro

75      IN       PTR      minotauro.telemed.org.
```

Este archivo muestra la zona de resolución inversa la cual se usa pa ra encontrar el nombre de la máquina a partir de su dirección IP.

F.4.4 SERVIDOR DE ACCESO REMOTO – RAS

Un servidor de acceso remoto es un equipo que permite a otro conectarse a él mediante una línea telefónica a través de un módem. Aunque para los efect os del protocolo PPP, los dos extremos de la comunicación son equivalentes, se denomina servidor PPP a aquél equipo que recibe la llamada, y en su caso, valida la contraseña facilitada por el otro, y se denomina cliente al equipo que efectúa la llamada y s olicita el establecimiento de la conexión.

Para poder recibir llamadas, es necesario tener un módem conectado a un puerto serie, soporte en el Kernel, directamente incluido o como módulo, para el protocolo PPP, así como los paquetes pppd y mgetty + sendfa x.

Para que el servidor atienda a otro terminal es necesario informar de ello al sistema. Esto se hace en el fichero /etc/inittab. Este fichero indica que procesos se ponen en marcha durante la operación de arranque del sistema. En este caso el módem se encuentra conectado a la puerta serie ttyS0 (COM1), por lo que se incluye la siguiente línea en el fichero:

S0:2345:respown:/usr/local/sbin/mgetty ttyS0

Aquí se presenta el archivo de configuración del inittab:

```
#
# inittab This file describes how the INIT process should set up
# the system in a certain run-level.
#
# Author: Miquel van Smoorenburg, <miquels@drinkel.nl.mugnet.org>
# Modified for RHS Linux by Marc Ewing and Donnie Barnes

#
S0:2345:respawn:/usr/local/sbin/mgetty ttyS0

# Default runlevel. The runlevels used by RHS are:
# 0 - halt (Do NOT set initdefault to this)
# 1 - Single user mode
# 2 - Multiuser, without NFS (The same as 3, if you do not have
networking)
# 3 - Full multiuser mode
# 4 - unused
# 5 - X11
# 6 - reboot (Do NOT set initdefault to this)
#
id:5:initdefault:

# System initialization.
si::sysinit:/etc/rc.d/rc.sysinit

l0:0:wait:/etc/rc.d/rc 0
l1:1:wait:/etc/rc.d/rc 1
l2:2:wait:/etc/rc.d/rc 2
l3:3:wait:/etc/rc.d/rc 3
l4:4:wait:/etc/rc.d/rc 4
l5:5:wait:/etc/rc.d/rc 5
l6:6:wait:/etc/rc.d/rc 6

# Things to run in every runlevel.
ud::once:/sbin/update
```

```
# Trap CTRL-ALT-DELETE
ca::ctrlaltdel:/sbin/shutdown -t3 -r now

# When our UPS tells us power has failed, assume we have a few
minutes
# of power left. Schedule a shutdown for 2 minutes from now.
# This does, of course, assume you have powerd installed and your
# UPS connected and working correctly.
pf::powerfail:/sbin/shutdown -f -h +2 "Power Failure; System
Shutting Down"

# If power was restored before the shutdown kicked in, cancel it.
pr:12345:powerokwait:/sbin/shutdown -c "Power Restored; Shutdown
Cancelled"


# Run gettys in standard runlevels
1:2345:respawn:/sbin/mingetty tty1
2:2345:respawn:/sbin/mingetty tty2
3:2345:respawn:/sbin/mingetty tty3
4:2345:respawn:/sbin/mingetty tty4
5:2345:respawn:/sbin/mingetty tty5
6:2345:respawn:/sbin/mingetty tty6

# Run xdm in runlevel 5
# xdm is now a separate service
x:5:respawn:/etc/X11/prefdm -nodaemon
```

S0: indica la abreviatura del nombre de dispositivo que usará init.

2345: son los runlevel para lo que se activa l a atención a este terminal.

respawn: indica a inity que la entrada está activa.

/var/local/etc/mgetty+sendfax:”””” es el programa que atenderá al terminal, mgetty puede detectar llamadas de fax o de datos por PPP.

Una vez editado /etc/inittab, ejecute la o rden “kill –1 1” en el terminal para forzar a que el proceso init vuelva a leerlo.

Otro archivo que se debe editar es el archivo loging.config que para este caso esta ubicado en:

/var/local/etc/mgetty+sendfax/loging.config. Este es uno de los ficheros de configuración del programa mgetty. Este programa es el que según indica el /etc/inittab recibirá el control cuando aparezca algo por el puerto serie /dev/ttyS0. Es decir que cuando llegue una llamada mgetty se hará cargo de ella de acuerdo a la configuraci ón de este fichero:

la configuración se presenta a continuación:

```
# login.config
#
# This is a sample "login dispatcher" configuration file for
mgetty
#
# Format:
# username userid utmp_entry login_program [arguments]
#
# Meaning:
# for a "username" entered at mgettys login: prompt, call
# "login_program" with [arguments], with the uid set to "userid",
# and a USER_PROCESS utmp entry with ut_user = "utmp_entry"
#
# username may be prefixed / suffixed by "*" (wildcard)
#
# userid is a valid user name from /etc/passwd, or "-" to not set
# a login user id and keep the uid/euid root (needed for
/bin/login)
#
# utmp_entry is what will appear in the "who" listing. Use "-" to
not
# set an utmp entry (a must for /bin/login), use "@" to set it to
the
# username entered. Maximum length is 8 characters.
#
# login_program is the program that will be exec()ed, with the
arguments
# passed in [arguments]. A "@" in the arguments will be replaced
with the
# username entered. Warning: if no "@" is given, the login_program
has
# no way to know what user name the user entered.
#
#
# SAMPLES:
# Use this one with my Taylor-UUCP and Taylor-UUCP passwd files.
# (Big advantage: tuucp can use the same passwd file for serial
dial-in
# and tcp dial-in [uucico running as in.uucpd]). Works from 1.05
up.
#
#U* uucp @ /usr/lib/uucp/uucico -l -u @
#
# Use this one for fido calls (login name /FIDO/ is handled
specially)
#
# You need Eugene Crosser's "ifmail" package for this to work.
# mgetty has to be compiled with "-DFIDO", otherwise a fido call
won't
# be detected.
#
#/FIDO/ uucp fido /usr/local/lib/fnet/ifcico @
#
# Automatic PPP startup on receipt of LCP configure request
(AutoPPP).
# mgetty has to be compiled with "-DAUTO_PPP" for this to work.
# Warning: Case is significant, AUTOPPP or autoppp won't work!
```

```
# Consult the "pppd" man page to find pppd options that work for
you.
#
# NOTE: for *some* users, the "-detach" option has been necessary,
for
# others, not at all. If your pppd doesn't die after hangup, try
it.
#
# NOTE2: "debug" creates lots of debugging info. LOOK AT IT if
things
# do not work out of the box, most likely it's a ppp problem!
#
# NOTE3: "man pppd" is your friend!
#
# NOTE4: max. 9 arguments allowed.
#
#/AutoPPP/ - a_ppp /usr/sbin/pppd auth -chap +pap login debug
/AutoPPP/ -  -    ppp
#
                                    #
# An example where no login name in the argument list is desired:
# automatically telnetting to machine "smarty" for a given login
name
#
#telnet-smarty gast telnet /usr/bin/telnet -8 smarty
#
# This is the "standard" behaviour - *dont* set a userid or utmp
# entry here, otherwise /bin/login will fail!
# This entry isn't really necessary: if it's missing, the built-in
# default will do exactly this.
#
* - - /bin/login @
```

Posteriormente se configura el fichero mgetty.config localizado en este caso en /var/local/etc/mgetty+sendfax/mgetty,config el cual es un fichero de configuración particular para mgetty. La configuración se presenta a continuación:

```
#
# mgetty configuration file
#
# this is a sample configuration file, see mgetty.info for details
#
# comment lines start with a "#", empty lines are ignored


# ----- global section -----
#
# In this section, you put the global defaults, per-port stuff is
below


# set the global debug level to "4" (default from policy.h)
#debug 4
# set the local fax station id
```

```
fax-id 49 115 xxxxxxxx

# access the modem(s) with 38400 bps
#speed 38400

# use these options to make the /dev/tty-device owned by
"uucp.uucp"
# and mode "rw-rw-r--" (0664). *LEADING ZERO NEEDED!*
#port-owner uucp
#port-group uucp
#port-mode 0664

# use these options to make incoming faxes owned by "root.uucp"
# and mode "rw-r-----" (0640). *LEADING ZERO NEEDED!*
#fax-owner root
#fax-group uucp
#fax-mode 0640


# ----- port specific section -----
#
# Here you can put things that are valid only for one line, not
the others
#

# Zoom V.FX 28.8, connected to ttyS0: don't do fax, less logging
#

port ttyS0
debug 6
speed 115200
data-only yes
#init-chat""\d+++\dAT&FH0 OK
rings 1

# some other Rockwell modem, needs "switchbd 19200" to receive
faxes
# properly (otherwise it will fail with "timeout").
#
#port ttyS1
# speed 38400
# switchbd 19200

# ZyXEL 2864, connected to ttyS2: maximum debugging, grab
statistics
#
#port ttyS2
# debug 8
# init-chat "" \d\d\d+++\d\d\dAT&FS2=255 OK ATN3S0=0S13.2=1 OK
# statistics-chat "" AT OK ATI2 OK
# statistics-file /var/log/statistics.ttyS2
# modem-type cls2

# direct connection of a VT100 terminal which doesn't like DTR
drops
#
#port ttyS3
# direct y
```

```
# speed 19200
# toggle-dtr n
```

donde:

port ttyS0 indica el puerto al que referencia esta configuración.

debug 6es el nivel de depuración en ficheros log.

speed 115200 es la velocidad de conexi ón.

data-only yes indica que solo se reciben datos.

rings 1 es el numero de rings antes de contestar.

La cadena init-chat ””?d+++dAT&FH0 OK esta comentada ya que el módem que se esta utilizando tiene la capacidad para iniciar automáticamente.

Una vez configurado este fichero se procede a hacer la configuración del fichero options localizado en /etc/ppp/options. Este archivo es donde se almacenan las opciones que se pasan al programa pppd para su ejecución, la diferencia más importante entre un servidor y un cliente, es que el servidor, normalmente, entrega una IP al cliente y le pregunta su contraseña para validarla en su base de datos local localizada en /etc/paswwd. Red Hat esta preparado para usar /etc/shadow, pero es posible que en otros sistemas no sea así, por lo que habrá que recompilar las fuentes PPP con la siguiente orden:

make HAS_SHADOW=1

para usar la opción MS_DNS para entregar la dirección IP del servidor DNS a clientes Windows se debe usar la siguiente instrucción:

make USE_MS_DNS=1 HAS_SHADOW = 1

la configuración del fichero options es la siguiente:

```
debug
#
lock
modem
crtscts
/dev/ttyS0
115200
asyncmap 0
#Requiere auntentificación
auth
login
#Dirección IP asignada al ordenador remoto
:172.16.130.10
proxyarp
nodetach
refuse-chap
require-pap
```

```
#
ms-dns 172.16.130.75
```

En esta configuración
auth, obliga al otro extremo a autentificarse antes de permitir el tráfico de paquetes.
login, significa que utilizará la base de datos de usuario para la autentificación (/etc/passwd)
:172.16.130.10, es la dirección IP que entrega al interfaz ppp0 del otro extremo.
proxyarp creará una entrada ARP para la IP del otro extremo asociada a la dirección. Ethernet del servidor, esto hará el efecto del que el otro extremo esta en la misma red que el servidor para otros equipos en la red del servidor. Si se asigna la IP remota dentro del rango de direcciones del segmento de red en que está el servidor, el extremo remoto es como si tuviera una tarjeta de red conectada a esta misma red, lo que simplifica el encaminamiento.
nodetach, no se desconecta del terminal.
refuse-chap, rehusa auntentificación chap.
require-pap, solicita autentificación pap.

Ahora se configura el fichero pap-secrets localizado en /etc/ppp/pap-secrets. Este archivo es donde se almacenan las contraseñas que se entregan cuando se son solicitadas por PPP. El archivo de configuración es el siguiente:

```
*     *     ""    172.16.130.10
# Secrets for authentication using PAP
# client server         secret                  IP addresses
```

Seguidamente se configura el alias global para pppd, esto se realiza en el fichero /etc/bashrc, con el fin de que una orden simple arranque ppp en el servidor una vez que se ha establecido la comunicación. El archivo debe ser como:

Alias ppp= "exec /usr/sbin/pppd –detach"

Lo que esta línea hace es:

- exec: quiere decir que reemplace el programa en ejecución con el programa que está ejecutando.
- pppd – detach: comienza pppd y no se parte en el segundo plano. Esto asegura que cuando pppd salga no queden procesos colgando por ahí.

Por último se requiere configurar pppd , es decir el demonio de ppp, para permitir que los usuarios puedan ejecutarlo, de lo contario los clientes no podrán conectarse al servidor. Para esto se abre la ventana de terminal de Red Hat y se escribe la siguiente instr ucción:

Chmod u+s /usr/sbin/pppd

Para simplificar las cosas a los usuarios del servicio de conexión PPP, se crea un alias global en /etc/bashrc para que una orden simple arranque ppp en el servidor una vez que se han conectado. Esto debe ser:

alias ppp="exect /usr/sbin/pppd –detach"

lo que hace es:

- exec: quiere decir que reemplace el programa en ejecución (en este caso el shell) con el programa que está ejecutando.
- Ppp –detach: comienza pppd y no se parte en el segundo plano. Esto asegura que cuando pppd salga no queden accesos colgando por ahí.

El archivo /etc/bashrc/ se muestra a continuación:

```
# /etc/bashrc
# System wide functions and aliases
# Environment stuff goes in /etc/profile

# are we an interactive shell?
if [ "$PS1" ]; then
      if [ -x /usr/bin/tput ]; then
            if [ "x`tput kbs`" != "x" ]; then # We can't do this
      with "dumb" terminal
               stty erase `tput kbs`
            fi
      fi
      case $TERM in
              xterm*)
                      PROMPT_COMMAND='echo -ne
      "\033]0;${USER}@${HOSTNAME}:   ${PWD}\007"'
              ;;
      *)
              ;;
      esac
      [ "$PS1" = "\\s-\\v\\\$ " ] && PS1="[\u@\h \W]\\$ "

      if [ "x$SHLVL" != "x1" ]; then # We're not a login shell
          for i in /etc/profile.d/*.sh; do
                if [ -x $i ]; then
```

```
                        . $i
                fi
            done
        fi
fi
alias ppp="exec /usr/sbin/pppd -detach"
```

Es importante realizar otro procedimiento para lograr el buen funcionamiento del mgetty y este consiste en :

Crear el grupo módem con el linuxconf de esta manera:
Se va a la ventana de terminal y se coloca

#linuxconf

Aparece la ventana de linuxconf y se va a normal, deficniciones de grupo.

En la parte inferior aparece un botón que indica agregar por medio del cual sale otra ventana. Solo se coloca el nombre del grupo módem que es el que se quiere agregar y se cierra indicando que se activen los cambios.

Posteriormente se coloca en la ventana del terminal la siguiente instrucción para mirar el puerto donde se encuentra conectado el módem:

 # cd /dev/

En este caso especifico el módem se encuentra localizado en ttyS0.

Se hace un enlace simbólico desde ttyS0 al grupo módem de la siguiente manera:

 # ln –s ttyS0 modem

Cuando se lista se ve lo siguiente:

 # ls –l

lrwxr-xr-x 1 root root fecha modem – ttyS0

Se cambian los permisos

```
# chmod 750 modem
```

Si se listan los archivos se ve:

```
# ls –l
lrwxr-x-x  root  root
```

lo que indica que los usuarios no pueden leer y ejecutar modem.
Posteriormente se cambia el grupo root a el grupo MODEM:

```
# chgrp modem modem
```

así se le dan los permisos de ejecución y lectura solo al grupo MODEM.

Luego se adicionan a este grupo los usuarios que pueden acceder al servidor por medio de la línea telefónica.

Es conveniente volver a dar la instrucción:

```
#Chmod u+s /usr/sbin/pppd
```

Para así de esta manera dar el control a los usuarios.

F.4.5 SERVIDOR DE CORREO ELECTRÓNICO

Para la realización de la prueba piloto de la RTPSTT se configuraron dos servicios adicionales como los son el correo electrónico y el FTP.

Un sistema de correo electrónico está compuesto por los siguientes elementos:

- MTA o agente de transferencia de correo: más conocido como servidor de correo, su función principal es almacenar los mensajes e intercambiarlos con las otras MTAs.
- AU o agente de usuario de correo: es el software de usuario final y su función es editar , componer y mostrar los correos (ejemplos: SSH, Outlook, PINE). El AU crea el mensaje y mediante un protocolo AU se lo entrega a una MTA para su transporte u manejo. De la misma forma, el AU se encarga de contactar a su MTA para recibir el correo.
- POP3: protocolo de AU, es el conjunto de comandos que el AU de usuario envía al MTA para recuperar su correo. El más conocido es POP pero también existe el Imap.
- MTP: protocolo utilizado entre servidores de correo (MTAs) para intercambiar los mensajes que tengan.

- WebMail: aplicación web que permite consultar a un servidor de correo a través de uno de los protocolos de AU (pop3, imap) y presentar su correo como páginas HTML. El más conocido es el sendmail.

La configuración del servidor de correo electrónico es bastante sencilla, pues solo se requieren dos paquetes que son conocidos como Imap y Sendmail. Y tener el DNS funcionando. Como primer paso pregunte al sistema si estos paquetes ya están instalados mediante la instrucción:

#rpm –q imap sendmail

Sino los tiene entonces instale los siguientes RPMS desde el CD de instalación de Ret Hat 7.0 :

#rpm –i imap4.7C2-12.rpm
#rpm –i sendmail-8.11.0-8.i386.rpm

Además de ello es necesario que el servidor tenga soporte POP, esto lo consigue entrando al *setup* desde el terminal y en la sección *system services* activando las opciones ipop3, imaps, imap y pop3s.

En los archivos de configuración del DNS es necesario agregar una línea en el archivo que define el nombre del dominio, para nuestro caso el archivo se ha denominado telemed.org. Añada entonces la siguiente línea la cual permite definir el servidor como servidor de correo electrónico.

```
@      IN     MX     5      minotauro
```

El numero 5 nos permite fijar las preferencias en el rango de 1 a 10, entre menor sea el número mayor será la prioridad.

Seguidamente, es necesario editar el archivo de configuración del sendmail ubicado en /etc/sendmail.cf y adicionar o descomentar las siguientes líneas:

CW localhost telemed.org
Dmtelemed.org

En la primera se define el nombre del dominio para el cual va ha recibir correo y la segunda permite que todo el correo saliente tenga como dominio telemed.org, por ejemplo, en nuestro caso existe el usuario prueba, por lo tanto todo el correo saliente será de la siguiente manera: prueba@telemed.org; pero si la línea Dmtelemed.org no estuviera habilitada la dirección sería: prueba@minotauro.telemed.org.

El archivo de configuración es el siguiente:

```
# Copyright (c) 1998-2000 Sendmail, Inc. and its suppliers.
#       All rights reserved.
# Copyright (c) 1983, 1995 Eric P. Allman.  All rights reserved.
# Copyright (c) 1988, 1993
#       The Regents of the University of California.  All rights reserved.
#
# By using this file, you agree to the terms and conditions set
# forth in the LICENSE file which can be found at the top level of
# the sendmail distribution.
#
#

##########################################################################
##########################################################################
#####
#####          SENDMAIL CONFIGURATION FILE
#####
#####
##########################################################################
##########################################################################

#####  $Id: cfhead.m4,v 8.76.4.9 2000/07/11 23:50:30 geir Exp $  #####
#####  $Id: cf.m4,v 8.32 1999/02/07 07:26:14 gshapiro Exp $  #####

#####  linux setup for Red Hat Linux  #####
#####  $Id: linux.m4,v 8.11.16.1 2000/05/09 18:48:58 gshapiro Exp $  #####

#####  $Id: local_procmail.m4,v 8.21 1999/11/18 05:06:23 ca Exp $  #####
#####  $Id: smrsh.m4,v 8.14 1999/11/18 05:06:23 ca Exp $  #####

#####  $Id: mailertable.m4,v 8.18 1999/07/22 17:55:35 gshapiro Exp $  #####

#####  $Id: virtusertable.m4,v 8.16 1999/07/22 17:55:36 gshapiro Exp $
#####

#####  $Id: redirect.m4,v 8.15 1999/08/06 01:47:36 gshapiro Exp $  #####

#####  $Id: always_add_domain.m4,v 8.9 1999/02/07 07:26:08 gshapiro Exp $
#####

#####  $Id: use_cw_file.m4,v 8.9 1999/02/07 07:26:13 gshapiro Exp $  #####


#####  $Id: local_procmail.m4,v 8.21 1999/11/18 05:06:23 ca Exp $  #####
```

```
#####  $Id: access_db.m4,v 8.15 1999/07/22 17:55:34 gshapiro Exp $  #####

#####  $Id: blacklist_recipients.m4,v 8.13 1999/04/02 02:25:13 gshapiro Exp
$  #####

#####  $Id: accept_unresolvable_domains.m4,v 8.10 1999/02/07 07:26:07
gshapiro Exp $  #####


#####  $Id: proto.m4,v 8.446.2.5.2.12 2000/07/19 21:41:19 gshapiro Exp $
#####


# level 9 config file format
V9/Berkeley

# override file safeties - setting this option compromises system security,
# addressing the actual file configuration problem is preferred
# need to set this before any file actions are encountered in the cf file
#O DontBlameSendmail=safe

# default LDAP map specification
# need to set this now before any LDAP maps are defined
#O LDAPDefaultSpec=-h localhost

##################
#   local info   #
##################
```

Cw localhost telemed.org

```
# file containing names of hosts for which we receive email
Fw/etc/mail/local-host-names

# my official domain name
# ... define this only if sendmail cannot automatically determine your domain
#Dj$w.Foo.COM

CP.

# "Smart" relay host (may be null)
DS

# operators that cannot be in local usernames (i.e., network indicators)
CO @ % !
# a class with just dot (for identifying canonical names)
C..

# a class with just a left bracket (for identifying domain literals)
C[[

# access_db acceptance class
C{Accept}OK RELAY
# Hosts that will permit relaying ($=R)
FR-o /etc/mail/relay-domains
```

```
# arithmetic map
Karith arith
# possible values for tls_connect in access map
C{tls}VERIFY ENCR

# who I send unqualified names to (null means deliver locally)
DR

# who gets all local email traffic ($R has precedence for unqualified names)
DH

# dequoting map
Kdequote dequote

# class E: names that should be exposed as from this host, even if we
masquerade
# class L: names that should be delivered locally, even if we have a relay
# class M: domains that should be converted to $M
# class N: domains that should not be converted to $M
CL root

# who I masquerade as (null for no masquerading) (see also $=M)
DMtelemed.org

# my name for error messages
DnMAILER-DAEMON


# Mailer table (overriding domains)
Kmailertable hash -o /etc/mail/mailertable

# Virtual user table (maps incoming users)
Kvirtuser hash -o /etc/mail/virtusertable
CPREDIRECT

# Access list database (for spam stomping)
Kaccess hash /etc/mail/access

# Configuration version number
DZ8.11.0


################
#   Options   #
################

# strip message body to 7 bits on input?
O SevenBitInput=False

# 8-bit data handling
O EightBitMode=pass8

# wait for alias file rebuild (default units: minutes)
O AliasWait=10

# location of alias file
```

```
O AliasFile=/etc/aliases

# minimum number of free blocks on filesystem
O MinFreeBlocks=100

# maximum message size
#O MaxMessageSize=1000000

# substitution for space (blank) characters
O BlankSub=.

# avoid connecting to "expensive" mailers on initial submission?
O HoldExpensive=False
# checkpoint queue runs after every N successful deliveries
#O CheckpointInterval=10

# default delivery mode
O DeliveryMode=background

# automatically rebuild the alias database?
# NOTE: There is a potential for a denial of service attack if this is set.
#       This option is deprecated and will be removed from a future version.
O AutoRebuildAliases

# error message header/file
#O ErrorHeader=/etc/mail/error-header

# error mode
#O ErrorMode=print

# save Unix-style "From_" lines at top of header?
#O SaveFromLine=False

# temporary file mode
O TempFileMode=0600

# match recipients against GECOS field?
#O MatchGECOS=False

# maximum hop count
#O MaxHopCount=17

# location of help file
O HelpFile=/etc/mail/helpfile

# ignore dots as terminators in incoming messages?
#O IgnoreDots=False

# name resolver options
#O ResolverOptions=+AAONLY
# deliver MIME-encapsulated error messages?
O SendMimeErrors=True

# Forward file search path
O ForwardPath=$z/.forward.$w:$z/.forward

# open connection cache size
```

```
O ConnectionCacheSize=2

# open connection cache timeout
O ConnectionCacheTimeout=5m

# persistent host status directory
#O HostStatusDirectory=.hoststat

# single thread deliveries (requires HostStatusDirectory)?
#O SingleThreadDelivery=False

# use Errors-To: header?
O UseErrorsTo=False

# log level
O LogLevel=9

# send to me too, even in an alias expansion?
#O MeToo=True

# verify RHS in newaliases?
O CheckAliases=False

# default messages to old style headers if no special punctuation?
O OldStyleHeaders=True

# SMTP daemon options
O DaemonPortOptions=Name=MTA
O DaemonPortOptions=Port=587, Name=MSA, M=E

# SMTP client options
#O ClientPortOptions=Address=0.0.0.0

# privacy flags
O PrivacyOptions=authwarnings

# who (if anyone) should get extra copies of error messages
#O PostmasterCopy=Postmaster

# slope of queue-only function
#O QueueFactor=600000

# queue directory
O QueueDirectory=/var/spool/mqueue

# timeouts (many of these)
#O Timeout.initial=5m
O Timeout.connect=1m
#O Timeout.iconnect=5m
#O Timeout.helo=5m
#O Timeout.mail=10m
#O Timeout.rcpt=1h
#O Timeout.datainit=5m
#O Timeout.datablock=1h
#O Timeout.datafinal=1h
#O Timeout.rset=5m
#O Timeout.quit=2m
```

```
#O Timeout.misc=2m
#O Timeout.command=1h
#O Timeout.ident=5s
#O Timeout.fileopen=60s
#O Timeout.control=2m
O Timeout.queuereturn=5d
#O Timeout.queuereturn.normal=5d
#O Timeout.queuereturn.urgent=2d
#O Timeout.queuereturn.non-urgent=7d
O Timeout.queuewarn=4h
#O Timeout.queuewarn.normal=4h
#O Timeout.queuewarn.urgent=1h
#O Timeout.queuewarn.non-urgent=12h
#O Timeout.hoststatus=30m
#O Timeout.resolver.retrans=5s
#O Timeout.resolver.retrans.first=5s
#O Timeout.resolver.retrans.normal=5s
#O Timeout.resolver.retry=4
#O Timeout.resolver.retry.first=4
#O Timeout.resolver.retry.normal=4

# should we not prune routes in route-addr syntax addresses?
#O DontPruneRoutes=False

# queue up everything before forking?
O SuperSafe=True

# status file
O StatusFile=/var/log/sendmail.st

# time zone handling:
#  if undefined, use system default
#  if defined but null, use TZ envariable passed in
#  if defined and non-null, use that info
#O TimeZoneSpec=

# default UID (can be username or userid:groupid)
O DefaultUser=8:12

# list of locations of user database file (null means no lookup)
O UserDatabaseSpec=/etc/mail/userdb.db

# fallback MX host
#O FallbackMXhost=fall.back.host.net

# if we are the best MX host for a site, try it directly instead of config
err
O TryNullMXList=true

# load average at which we just queue messages
#O QueueLA=8

# load average at which we refuse connections
#O RefuseLA=12

# maximum number of children we allow at one time
#O MaxDaemonChildren=12
```

```
# maximum number of new connections per second
#O ConnectionRateThrottle=3

# work recipient factor
#O RecipientFactor=30000

# deliver each queued job in a separate process?
#O ForkEachJob=False

# work class factor
#O ClassFactor=1800
# work time factor
#O RetryFactor=90000

# shall we sort the queue by hostname first?
#O QueueSortOrder=priority

# minimum time in queue before retry
#O MinQueueAge=30m

# default character set
#O DefaultCharSet=iso-8859-1

# service switch file (ignored on Solaris, Ultrix, OSF/1, others)
#O ServiceSwitchFile=/etc/mail/service.switch

# hosts file (normally /etc/hosts)
#O HostsFile=/etc/hosts

# dialup line delay on connection failure
#O DialDelay=10s

# action to take if there are no recipients in the message
#O NoRecipientAction=add-to-undisclosed

# chrooted environment for writing to files
#O SafeFileEnvironment=/arch

# are colons OK in addresses?
#O ColonOkInAddr=True

# how many jobs can you process in the queue?
#O MaxQueueRunSize=10000

# shall I avoid expanding CNAMEs (violates protocols)?
#O DontExpandCnames=False

# SMTP initial login message (old $e macro)
O SmtpGreetingMessage=$j Sendmail $v/$Z; $b
# UNIX initial From header format (old $l macro)
O UnixFromLine=From $g $d

# From: lines that have embedded newlines are unwrapped onto one line
#O SingleLineFromHeader=False

# Allow HELO SMTP command that does not include a host name
```

```
#O AllowBogusHELO=False

# Characters to be quoted in a full name phrase (@,;:\()[] are automatic)
#O MustQuoteChars=.

# delimiter (operator) characters (old $o macro)
O OperatorChars=.:%@!^/[]+

# shall I avoid calling initgroups(3) because of high NIS costs?
#O DontInitGroups=False

# are group-writable :include: and .forward files (un)trustworthy?
#O UnsafeGroupWrites=True

# where do errors that occur when sending errors get sent?
#O DoubleBounceAddress=postmaster

# where to save bounces if all else fails
#O DeadLetterDrop=/var/tmp/dead.letter

# what user id do we assume for the majority of the processing?
#O RunAsUser=sendmail

# maximum number of recipients per SMTP envelope
#O MaxRecipientsPerMessage=100

# shall we get local names from our installed interfaces?
O DontProbeInterfaces=true

# Return-Receipt-To: header implies DSN request
#O RrtImpliesDsn=False
# override connection address (for testing)
#O ConnectOnlyTo=0.0.0.0

# Trusted user for file ownership and starting the daemon
#O TrustedUser=root

# Control socket for daemon management
#O ControlSocketName=/var/spool/mqueue/.control

# Maximum MIME header length to protect MUAs
#O MaxMimeHeaderLength=0/0

# Maximum length of the sum of all headers
#O MaxHeadersLength=32768

# Maximum depth of alias recursion
#O MaxAliasRecursion=10
# location of pid file
#O PidFile=/var/run/sendmail.pid

# Prefix string for the process title shown on 'ps' listings
#O ProcessTitlePrefix=prefix

# Data file (df) memory-buffer file maximum size
#O DataFileBufferSize=4096
```

```
# Transcript file (xf) memory-buffer file maximum size
#O XscriptFileBufferSize=4096

# list of authentication mechanisms
#O AuthMechanisms=GSSAPI KERBEROS_V4 DIGEST-MD5 CRAM-MD5

# default authentication information for outgoing connections
#O DefaultAuthInfo=/etc/mail/default-auth-info
# SMTP AUTH flags
#O AuthOptions
# CA directory
#O CACERTPath
# CA file
#O CACERTFile
# Server Cert
#O ServerCertFile
# Server private key
#O ServerKeyFile
# Client Cert
#O ClientCertFile
# Client private key
#O ClientKeyFile
# DHParameters (only required if DSA/DH is used)
#O DHParameters
# Random data source (required for systems without /dev/urandom under
OpenSSL)
#O RandFile



############################
#   Message precedences   #
############################

Pfirst-class=0
Pspecial-delivery=100
Plist=-30
Pbulk=-60
Pjunk=-100

######################
#   Trusted users   #
######################
# this is equivalent to setting class "t"
#Ft/etc/mail/trusted-users
Troot
Tdaemon
Tuucp

#########################
#   Format of headers   #
#########################

H?P?Return-Path: <$g>
HReceived: $?sfrom $s $.$?_($?s$|from $.$_)
        $.$?{auth_type}(authenticated$?{auth_ssf} (${auth_ssf} bits)$.)
        $.by $j ($v/$Z)$?r with $r$. id $i$?{tls_version}
```

```
      (using ${tls_version} with cipher ${cipher} (${cipher_bits} bits)
verified
${verify})$.$?u
      for $u; $|;
$.$b
H?D?Resent-Date: $a
H?D?Date: $a
H?F?Resent-From: $?x$x <$g>$|$g$.
H?F?From: $?x$x <$g>$|$g$.
H?x?Full-Name: $x
# HPosted-Date: $a
# H?l?Received-Date: $b
H?M?Resent-Message-Id: <$t.$i@$j>
H?M?Message-Id: <$t.$i@$j>

#
######################################################################
######################################################################
#####
#####                REWRITING RULES
#####
######################################################################
######################################################################
################################################
###  Ruleset 3 -- Name Canonicalization  ###
################################################
Scanonify=3

# handle null input (translate to <@> special case)
R$@                  $@ <@>

# strip group: syntax (not inside angle brackets!) and trailing semicolon
R$*                  $: $1 <@>                   mark addresses
R$* < $* > $* <@>       $: $1 < $2 > $3             unmark <addr>
R@ $* <@>               $: @ $1                   unmark @host:...
R$* :: $* <@>           $: $1 :: $2               unmark node::addr
R:include: $* <@>       $: :include: $1             unmark :include:...
R$* [ IPv6 $- ] <@>     $: $1 [ IPv6 $2 ]           unmark IPv6 addr
R$* : $* [ $* ]         $: $1 : $2 [ $3 ] <@>       remark if leading colon
R$* : $* <@>            $: $2                     strip colon if marked
R$* <@>                 $: $1                     unmark
R$* ;                   $1                      strip trailing semi
R$* < $* ; >            $1 < $2 >                 bogus bracketed semi

# null input now results from list:; syntax
R$@                  $@ :; <@>

# strip angle brackets -- note RFC733 heuristic to get innermost item
R$*                  $: < $1 >                   housekeeping <>
R$+ < $* >              < $2 >                   strip excess on left
R< $* > $+              < $1 >                   strip excess on right
R<>                  $@ < @ >                    MAIL FROM:<> case
R< $+ >                 $: $1                     remove housekeeping <>

# strip route address <@a,@b,@c:user@d> -> <user@d>
R@ $+ , $+              $2
R@ $+ : $+              $2
```

```
# find focus for list syntax
R $+ : $* ; @ $+        $@ $>Canonify2 $1 : $2 ; < @ $3 >       list syntax
R $+ : $* ;             $@ $1 : $2;                 list syntax

# find focus for @ syntax addresses
R$+ @ $+                $: $1 < @ $2 >              focus on domain
R$+ < $+ @ $+ >         $1 $2 < @ $3 >              move gaze right
R$+ < @ $+ >            $@ $>Canonify2 $1 < @ $2 >      already canonical

# do some sanity checking
R$* < @ $* : $* > $*    $1 < @ $2 $3 > $4           nix colons in addrs

# convert old-style addresses to a domain-based address
R$- ! $+                $@ $>Canonify2 $2 < @ $1 .UUCP >        resolve uucp
names
R$+ . $- ! $+           $@ $>Canonify2 $3 < @ $1 . $2 >         domain uucps
R$+ ! $+                $@ $>Canonify2 $2 < @ $1 .UUCP >        uucp subdomains

# if we have % signs, take the rightmost one
R$* % $*                $1 @ $2                     First make them all @s.
R$* @ $* @ $*           $1 % $2 @ $3                Undo all but the last.
R$* @ $*                $@ $>Canonify2 $1 < @ $2 >      Insert < > and finish

# else we must be a local name
R$*                     $@ $>Canonify2 $1


##################################################
###  Ruleset 96 -- bottom half of ruleset 3  ###
##################################################

SCanonify2=96

# handle special cases for local names
R$* < @ localhost > $*          $: $1 < @ $j . > $2         no domain at
all
R$* < @ localhost . $m > $*     $: $1 < @ $j . > $2         local domain
R$* < @ localhost . UUCP > $*   $: $1 < @ $j . > $2         .UUCP domain

# check for IPv6 domain literal (save quoted form)
R$* < @ [ IPv6 $- ] > $*        $: $2 $| $1 < @@ [ $(dequote $2 $) ] > $3
mark IPv6
addr
R$- $| $* < @@ $=w > $*         $: $2 < @ $j . > $4         self-literal
R$- $| $* < @@ [ $+ ] > $*      $@ $2 < @ [ IPv6 $1 ] > $4      canon IP addr

# check for IPv4 domain literal
R$* < @ [ $+ ] > $*             $: $1 < @@ [ $2 ] > $3          mark [a.b.c.d]
R$* < @@ $=w > $*               $: $1 < @ $j . > $3         self-literal
R$* < @@ $+ > $*                $@ $1 < @ $2 > $3           canon IP addr

# if really UUCP, handle it immediately

# try UUCP traffic as a local address
R$* < @ $+ . UUCP > $*          $: $1 < @ $[ $2 $] . UUCP . > $3
R$* < @ $+ . . UUCP . > $*      $@ $1 < @ $2 . > $3
```

```
# hostnames ending in class P are always canonical
R$* < @ $* $=P > $*            $: $1 < @ $2 $3 . > $4
R$* < @ $* $~P > $*            $: $&{daemon_flags} $| $1 < @ $2 $3 > $4
R$* CC $* $| $*              $: $3
# pass to name server to make hostname canonical
R$* $| $* < @ $* > $*          $: $2 < @ $[ $3 $] > $4
R$* $| $*                    $: $2

# local host aliases and pseudo-domains are always canonical
R$* < @ $=w > $*             $: $1 < @ $2 . > $3
R$* < @ $=M > $*             $: $1 < @ $2 . > $3
R$* < @ $={VirtHost} > $*      $: $1 < @ $2 . > $3
R$* < @ $* . . > $*          $1 < @ $2 . > $3


##################################################
###  Ruleset 4 -- Final Output Post-rewriting  ###
##################################################
Sfinal=4

R$* <@>              $@                      handle <> and list:;

# strip trailing dot off possibly canonical name
R$* < @ $+ . > $*      $1 < @ $2 > $3

# eliminate internal code
R$* < @ *LOCAL* > $*    $1 < @ $j > $2

# externalize local domain info
R$* < $+ > $*          $1 $2 $3                  defocus
R@ $+ : @ $+ : $+      @ $1 , @ $2 : $3          <route-addr> canonical
R@ $*                $@ @ $1                   ... and exit

# UUCP must always be presented in old form
R$+ @ $- . UUCP        $2!$1                     u@h.UUCP => h!u
# fin archivo de configuración
```

Después de realizar estos cambios es aconsejable reiniciar el servicio mediante la instrucción:

/etc/rc.d/init.d/sendmail restart

La configuración de los clientes de correo es la siguiente:

En /etc/mail/acces se definen las máquinas que pueden usar este servidor, como servidor de salida; para nuestro caso se habilitarán todos los equipos de la red 172.16. por los cual se añade entonces la línea :

172.16 RELAY

Finalmente se ejecuta la instrucción makemap hash /etc/mail/acces < /etc/mail/acces

F.4.6 SERVIDOR FTP

FTP (File Transfer Protocol) es usado para conectar el cliente y el servidor con el objetivo de transferir archivos de un extremo al otro. FTP trabaja por que un cliente (software cliente FTP) se conecta al servidor y coloca u obtiene archivos del servidor.

Existen dos tipos de servidores FTP: uno para acceso de usuario y otro para acceso anonymous. El acceso al servidor de tipo usuario requiere un login y password en la máquina que el servidor valida antes de permitir operaciones sobre él. El otro tipo d eservidor permite la conexión sin necesidad de un login específico. Para el caso de la prueba piloto se configuró un servidor anonymous, el más usado es el wu-ftpd.

Entonces para configurar este servicio primero deben instalarse los paquetes wu-ftpd y anon ftp.

rpm –i wu-ftpd-2.6.1-6.i386.rpm
rpm –i anonftp-3.0-9.i386.rpm

Es necesario adicionar una nueva línea en uno de los archivos de configuración del DNS, en var/named/telemed.org se añade:

ftp IN CNAME minotauro

La cual indica que minotauro estará habilitado como servidor FTP.

Después por medio del gestor de archivos en var/ftp es necesario cambiar la configuración de cada una de las carpetas Para ello se debe hacer clic derecho sobre cada una de las carpetas y en propiedades y permisos configurarlas de la forma que se muestra en la tabla E.1:

Tabla E.1. Configuración de las carpetas del FTP

PERMISOS	GRUPO	DUEÑO	DIRECTORIOS
drwxr – xr - x	Root	Root	FTP
d- - x- - x - - x	Root	Root	Bin
d- - x - - - - - -	Root	Root	Etc
drwxr – x r –x	Root	Root	Lib
drwxrwxr - x	ftp	ftp	Pub
- rwxr - - r - -	ftp	ftp	Welcome.msg

Para el caso de la configuración realizada, en la carpeta welcome.msg se escribió el siguiente mensaje de bienvenida :" bienvenido al FTP de Teleled, la información se encuentra en la carpeta PUB". De esta manera, el usuario puede descargar o subir la información en la siguiente dirección: ftp://ftp.telemed.org.

Otro de los aspectos a tener en cuenta, es que si sólo se desea que la carpeta pub sea visible en el FTP se debem seguir los siguientes pasos : en el *linuxconf* en la sección de *cuentas de usuario, ftp* se debe escribir como directorio var/ftp/pub.