Squid, Web Proxy Cache

Sergio González González

Universidad de León, España

sergio.gonzalez@hispalinux.es

Guía de instalación y configuración de Squid, con las opciones de aceleración para servidores HTTP, cacheo transparente, jerarquías de caches, delay pools para el control del ancho de banda y control de acceso a la caché, entre otras. La guia está pensada para la distribución Debian GNU/Linux y para un caso concreto, por lo que habrá de adaptarse a las necesidades personales de cada lector.

Introducción

Squid es un proxy cache para objetos de Internet, ya que no se limita sólo al tráfico web. La elección de este proxy es por las características casi ilimitadas que posee. Algunas de las más importantes son:

- actuazión como proxy y cache para HTTP, FTP y Gopher
- soporte para SSL
- jerarquía de caches
- ICP, HTCP, CARP, Cache Digests
- cacheo transparente
- WCCP - Web Cache Coordination Protocol (en las versiones iguales o superiores a la 2.3)
- listas de control de acceso
- aceleración de un servidor HTTP
- SNMP
- cacheo de peticiones DNS

La elección de este proxy ha sido por que se adaptaba a todas las características que queríamos implementar en la red: integración con el cortafuegos utilizado[1] para hacer de proxy transparente, control de acceso al servidor proxy, control del ancho de banda, "caché rápida", etc.

Instalación

Sólo necesitamos instalar un paquete para tener Squid listo para ser configurado. El paquete en cuestión es *squid*. La forma de instalarlo será la siguiente:

```
# apt-get install squid
```

Una vez que se ha instalado correctamente Squid, pasamos a configurarlo.

Configuración

La configuración de Squid se lleva a cabo con el siguiente fichero:

- `/etc/squid.conf`

Antes de pasar a especificar las directivas de configuración más importantes, mostraré como está configurada la jerarquia de caches para comprender mejor la configuración. Actualmente existen 3 ordenadores con Squid: *Cancerbero*, *Hidrogeno* y *Sanpedro*. La función de cada uno puede verse a continuación:

- *Hidrogeno* tiene una caché con control de ancho de banda. *Cancerbero* es su padre y *Sanpedro* su hermano.
- *Sanpedro* tiene una caché con control de ancho de banda. *Cancerbero* es su padre e *Hidrogeno* su hermano.
- *Cancerbero* tiene una caché sin control de ancho de banda. *Hidrógeno* y *Sanpedro* son sus hermanos.

En la imágen queda reflejada la comunicación entre cachés:

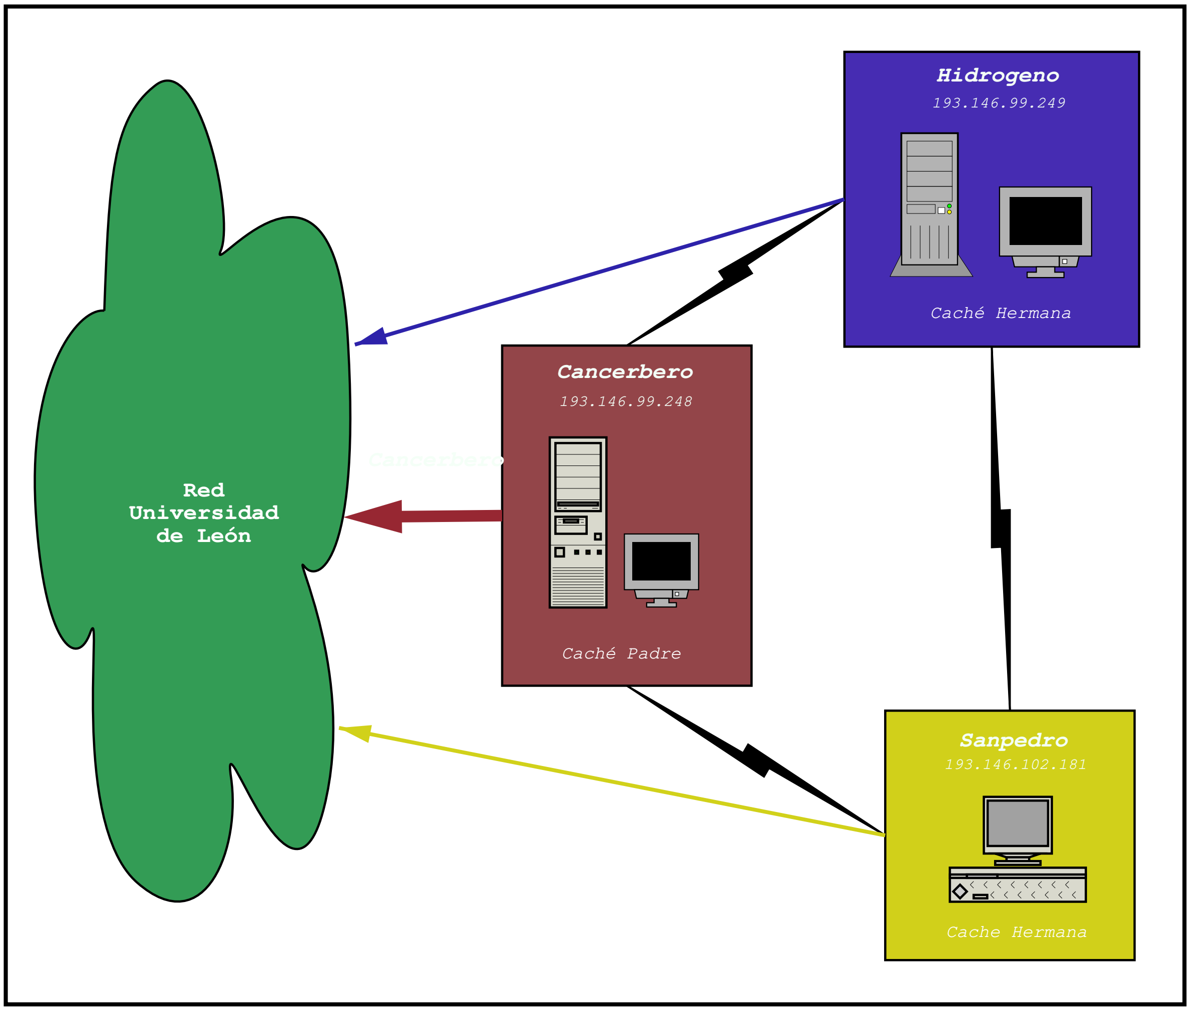

Hay cuatro puntos importantes para llegar ha configurar la jerarquía de cachés que está funcionando actualmente:

- definición de la jerarquia a utilizar
- control de acceso a las cachés
- control del ancho de banda
- cacheo transparente de las peticiones

En los siguientes apartados se mostrarán las opciones necesarias para llevar a cabo estas funciones.

Definición de la jerarquía a utilizar

Antes de definir la jerarquía, debemos definir los puertos por los cuales se van a comunicar las cachés, esto se hace desde la sección *NETWORK OPTIONS* del archivo de configuración. Las directivas quedarían:

```
http_port 3128
```

```
icp_port 3130
htcp_port 4827
```

En la sección *OPTIONS WHICH AFFECT THE NEIGHBOR SELECTION ALGORITHM* del archivo de configuración de Squid definiremos las cachés vecinas de la actual y el parentesco que tienen con esta. Si tomamos por ejemplo a Hidrógeno, sus vecinos y parentescos quedarían configurados de la siguiente forma:

```
cache_peer 193.146.99.248 parent 3128 3130 no-digest default
cache_peer 193.146.102.181 sibling 3128 3130 no-delay
```

Control de acceso a las cachés

Apartado en el que veremos la forma que tiene Squid para controlar el acceso al servicio que proporciona. Si nos desplazamos hasta la sección *ACCESS CONTROLS* del archivo de configuración, veremos la forma de llevar a cabo nuestro cometido. El ejemplo que veremos a continuación, es una parte de la lista de control de acceso de *Hidrogeno*. Este ha de permitir el acceso a la red local a la que proporciona el servicio de chaché así como a sus cachés vecinas. Veamos como se hace:

- Definimos las redes y los host a los que vamos a proporcionar acceso:

  ```
  acl red_local src 192.168.2.0/24
  acl cancerbero src 193.146.99.248/255.255.255.255
  acl sanpedro src 193.146.102.181/255.255.255.255
  ```

- Les permitimos acceder a la caché:

  ```
  http_access allow red_local
  http_access allow cancerbero
  http_access allow sanpedro
  ```

- Permitimos a las cachés vecinas intercambiar mensajes de control por los puertos definidos para tal efecto:

  ```
  icp_access allow cancerbero
  icp_access allow sanpedro
  ```

Control del ancho de banda

Este control lo llevan a cabo las *Delay pools*. Vamos a configurar nuestro proxy de forma que determinados archivos se bajen de Internet más lentamente que otros a determinadas horas. Con esto conseguimos que la navegación web en horas puntas sea fluída, y no sea entorpecida por “bajadores compulsivos” ;-)

Si seguimos los siguientes pasos, habremos logrado nuestro cometido:

- No queremos limitar las descargas en nuestra red local:

```
acl magic_words1 url_regex -i 192.168
```

- Queremos limitar la descarga de este tipo de archivos:

```
acl magic_words2 url_regex -i ftp .exe .mp3 .vqf .tar.gz .gz .rpm .zip .rar .avi .mpeg .mp
```

- Queremos limitar el ancho de banda durante el día, permitiendo el ancho de banda completo durante la noche:

```
acl day time 09:00-23:59
```

- Tenemos dos delay_pools diferentes:

```
delay_pools 2
```

- *Primer delay pool* - no queremos retrasar nuestro tráfico local:

```
delay_class 1 2
delay_parameters 1 -1/-1 -1/-1
delay_access 1 allow magic_words1
```

- *Segundo delay pool* - queremos retrasar la descarga de los archivos mencionados en “magic_words2”. Limitaremos el ancho de banda de bajada de este tipo de archivos a 5 Kbytes/s:

```
delay_class 2 2
delay_parameters 2 5000/150000 5000/120000
delay_access 2 allow day
delay_access 2 deny !day
delay_access 2 allow magic_words2
```

Cacheo transparente de las peticiones

Para evitar la configuración de cada cliente de la red a la que sirve el proxy caché, tenemos la opción de configurar Squid para que actue de forma transparente. Esto quiere decir que todas las peticines que los clientes realicen al puerto 80 (si nos referimos a la navegación web), pasarán a través de Squid automáticamente[2].

Las siguientes opciones son necesarias para que Squid trabaje de forma transparente:

```
httpd_accel_host virtual
httpd_accel_port 80
httpd_accel_with_proxy on
httpd_accel_uses_host_header on
```

Ejemplo de fichero de configuración de Squid - Cancerbero

Este archivo de configuración está pensado para *Cancerbero*, servidor de la Unidad de Imagen. Este servidor actua como padre de una jerarquía de cachés Squid y no tiene control de ancho de banda.

```
# WELCOME TO SQUID 2
# ------------------
#
# This is the default Squid configuration file. You may wish
# to look at the Squid home page (http://www.squid-cache.org/)
# 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.
#
#Default:
 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.
#
#Default:
 icp_port 3130


#  TAG: htcp_port
# The port number where Squid sends and receives HTCP queries to
# and from neighbor caches.  To turn it on you want to set it 4827.
# By default it is set to "0" (disabled).
#
# To enable this option, you must use --enable-htcp with the
# configure script.
#
#Default:
 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://www.squid-cache.org/FAQ/).
#
# Usage: mcast_groups 239.128.16.128 224.0.1.20
#
# By default, Squid doesn't listen on any multicast groups.
#
#Default:
# none

#  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.
#
# A *_incoming_address value of 0.0.0.0 indicates that Squid should
# listen on all available interfaces.
#
# If udp_outgoing_address is set to 255.255.255.255 (the default)
# then it will use the same socket as udp_incoming_address. Only
# change this if you want to have ICP queries sent using another
# address than where this Squid listens for ICP queries from other
# caches.
#
# NOTE, udp_incoming_address and udp_outgoing_address can not
# have the same value since they both use port 3130.
#
# NOTE, tcp_incoming_address has been removed.  You can now
# specify IP addresses on the 'http_port' line.
#
#Default:
# tcp_outgoing_address 255.255.255.255
# udp_incoming_address 0.0.0.0
# udp_outgoing_address 255.255.255.255


# OPTIONS WHICH AFFECT THE NEIGHBOR SELECTION ALGORITHM
# -----------------------------------------------------------------------------

#  TAG: cache_peer
# To specify other caches in a hierarchy, use the format:
#
# cache_peer hostname type http_port icp_port
#
# For example,
#
# #                                        proxy  icp
# #          hostname             type     port   port  options
# #          -------------------- -------- ----- -----  -----------
# cache_peer parent.foo.net       parent    3128  3130  [proxy-only]
# cache_peer sib1.foo.net         sibling   3128  3130  [proxy-only]
# cache_peer sib2.foo.net         sibling   3128  3130  [proxy-only]
#
#       type:  either 'parent', 'sibling', or 'multicast'.
#
# proxy_port:  The port number where the cache listens for proxy
#      requests.
#
#   icp_port:  Used for querying neighbor caches about
#      objects.  To have a non-ICP neighbor
#      specify '7' for the ICP port and make sure the
```

```
#       neighbor machine has the UDP echo port
#       enabled in its /etc/inetd.conf file.
#
#      options: proxy-only
#       weight=n
#       ttl=n
#       no-query
#       default
#       round-robin
#       multicast-responder
#       closest-only
#       no-digest
#       no-netdb-exchange
#       no-delay
#       login=user:password
#       connect-timeout=nn
#       digest-url=url
#       allow-miss
#
#       use 'proxy-only' to specify that objects fetched
#       from this cache should not be saved locally.
#
#       use 'weight=n' to specify a weighted parent.
#       The weight must be an integer.  The default weight
#       is 1, larger weights are favored more.
#
#       use 'ttl=n' to specify a IP multicast TTL to use
#       when sending an ICP queries to this address.
#       Only useful when sending to a multicast group.
#       Because we don't accept ICP replies from random
#       hosts, you must configure other group members as
#       peers with the 'multicast-responder' option below.
#
#       use 'no-query' to NOT send ICP queries to this
#       neighbor.
#
#       use 'default' if this is a parent cache which can
#       be used as a "last-resort." You should probably
#       only use 'default' in situations where you cannot
#       use ICP with your parent cache(s).
#
#       use 'round-robin' to define a set of parents which
#       should be used in a round-robin fashion in the
#       absence of any ICP queries.
#
#       'multicast-responder' indicates that the named peer
#       is a member of a multicast group.  ICP queries will
#       not be sent directly to the peer, but ICP replies
#       will be accepted from it.
#
#       'closest-only' indicates that, for ICP_OP_MISS
#       replies, we'll only forward CLOSEST_PARENT_MISSes
#       and never FIRST_PARENT_MISSes.
```

```
#
#       use 'no-digest' to NOT request cache digests from
#       this neighbor.
#
#       'no-netdb-exchange' disables requesting ICMP
#       RTT database (NetDB) from the neighbor.
#
#       use 'no-delay' to prevent access to this neighbor
#       from influencing the delay pools.
#
#       use 'login=user:password' if this is a personal/workgroup
#       proxy and your parent requires proxy authentication.
#
#       use 'connect-timeout=nn' to specify a peer
#       specific connect timeout (also see the
#       peer_connect_timeout directive)
#
#       use 'digest-url=url' to tell Squid to fetch the cache
#       digest (if digests are enabled) for this host from
#       the specified URL rather than the Squid default
#       location.
#
#       use 'allow-miss' to disable Squid's use of only-if-cached
#       when forwarding requests to siblings. This is primarily
#       useful when icp_hit_stale is used by the sibling. To
#       extensive use of this option may result in forwarding
#       loops, and you should avoid having two-way peerings
#       with this option. (for example to deny peer usage on
#       requests from peer by denying cache_peer_access if the
#       source is a peer)
#
# NOTE: non-ICP neighbors must be specified as 'parent'.
#
#Default:
# cache_peer 192.168.1.2 sibling 3128 3130 weight=5 no-delay
# cache_peer hidrogeno.unileon.es sibling 3128 3130 no-delay

#  TAG: cache_peer_domain
# Use to limit the domains for which a neighbor cache will be
# queried.  Usage:
#
# cache_peer_domain cache-host domain [domain ...]
# cache_peer_domain cache-host !domain
#
# For example, specifying
#
# cache_peer_domain parent.foo.net .edu
#
# has the effect such that UDP query packets are sent to
# 'bigserver' only when the requested object exists on a
# server in the .edu domain.  Prefixing the domainname
# with '!' means that the cache will be queried for objects
# NOT in that domain.
```

```
#
# NOTE: * Any number of domains may be given for a cache-host,
#   either on the same or separate lines.
# * When multiple domains are given for a particular
#   cache-host, the first matched domain is applied.
# * Cache hosts with no domain restrictions are queried
#   for all requests.
# * There are no defaults.
# * There is also a 'cache_peer_access' tag in the ACL
#   section.
#
#Default:
# none

#  TAG: neighbor_type_domain
# usage: neighbor_type_domain parent|sibling domain domain ...
#
# Modifying the neighbor type for specific domains is now
# possible.  You can treat some domains differently than the the
# default neighbor type specified on the 'cache_peer' line.
# Normally it should only be necessary to list domains which
# should be treated differently because the default neighbor type
# applies for hostnames which do not match domains listed here.
#
#EXAMPLE:
# cache_peer  parent cache.foo.org 3128 3130
# neighbor_type_domain cache.foo.org sibling .com .net
# neighbor_type_domain cache.foo.org sibling .au .de
#
#Default:
# none

#  TAG: icp_query_timeout (msec)
# Normally Squid will automatically determine an optimal ICP
# query timeout value based on the round-trip-time of recent ICP
# queries.  If you want to override the value determined by
# Squid, set this 'icp_query_timeout' to a non-zero value.  This
# value is specified in MILLISECONDS, so, to use a 2-second
# timeout (the old default), you would write:
#
# icp_query_timeout 2000
#
#Default:
# icp_query_timeout 0

#  TAG: maximum_icp_query_timeout (msec)
# Normally the ICP query timeout is determined dynamically.  But
# sometimes it can lead to very large values (say 5 seconds).
# Use this option to put an upper limit on the dynamic timeout
# value.  Do NOT use this option to always use a fixed (instead
# of a dynamic) timeout value. To set a fixed timeout see the
# 'icp_query_timeout' directive.
#
```

```
#Default:
# maximum_icp_query_timeout 2000

#  TAG: mcast_icp_query_timeout (msec)
# For Multicast peers, Squid regularly sends out ICP "probes" to
# count how many other peers are listening on the given multicast
# address.  This value specifies how long Squid should wait to
# count all the replies.  The default is 2000 msec, or 2
# seconds.
#
#Default:
# mcast_icp_query_timeout 2000

#  TAG: dead_peer_timeout (seconds)
# This controls how long Squid waits to declare a peer cache
# as "dead."  If there are no ICP replies received in this
# amount of time, Squid will declare the peer dead and not
# expect to receive any further ICP replies.  However, it
# continues to send ICP queries, and will mark the peer as
# alive upon receipt of the first subsequent ICP reply.
#
# This timeout also affects when Squid expects to receive ICP
# replies from peers.  If more than 'dead_peer' seconds have
# passed since the last ICP reply was received, Squid will not
# expect to receive an ICP reply on the next query.  Thus, if
# your time between requests is greater than this timeout, you
# will see a lot of requests sent DIRECT to origin servers
# instead of to your parents.
#
#Default:
# dead_peer_timeout 10 seconds

#  TAG: hierarchy_stoplist
# A list of words which, if found in a URL, cause the object to
# be handled directly by this cache.  In other words, use this
# to not query neighbor caches for certain objects.  You may
# list this option multiple times.
#
#We recommend you to use at least the following line.
hierarchy_stoplist cgi-bin jsp asp miguelez ?

#  TAG: no_cache
# A list of ACL elements which, if matched, cause the reply to
# immediately removed from the cache.  In other words, use this
# to force certain objects to never be cached.
#
# You must use the word 'DENY' to indicate the ACL names which should
# NOT be cached.
#
#We recommend you to use the following two lines.
acl QUERY urlpath_regex cgi-bin \?
acl PAGINA_CANCERBERO dst 193.146.99.249/255.255.255.255
acl PAGINA_CANCERBERO_LOCAL dst 192.168.1.1/255.255.255.255
```

```
no_cache deny PAGINA_CANCERBERO
no_cache deny PAGINA_CANCERBERO_LOCAL
no_cache deny QUERY


# 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.
#
#Default:
# 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.
#
#Default:
 cache_swap_low 90
 cache_swap_high 93

#  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.
#
#Default:
 maximum_object_size 10240 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.
#
#Default:
 minimum_object_size 0 KB

#  TAG: maximum_object_size_in_memory (bytes)
#        Objects greater than this size will not be attempted to kept in
#        the memory cache. This should be set high enough to keep objects
#        accessed frequently in memory to improve performance whilst low
#        enough to keep larger objects from hoarding cache_mem .
#
#Default:
 maximum_object_size_in_memory 16 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.
#
#Default:
 ipcache_size 1024
 ipcache_low 90
 ipcache_high 93

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

```
#Default:
 fqdncache_size 1024

#  TAG: cache_replacement_policy
# The cache replacement policy parameter determines which
# objects are evicted (replaced) when disk space is needed.
#
#     lru       : Squid's original list based LRU policy
#     heap GDSF : Greedy-Dual Size Frequency
#     heap LFUDA: Least Frequently Used with Dynamic Aging
#     heap LRU  : LRU policy implemented using a heap
#
# Applies to any cache_dir lines listed below this.
#
# The LRU policies keeps recently referenced objects.
#
# The heap 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 heap 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 the GDSF and LFUDA 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.
#
#Default:
 cache_replacement_policy heap LFUDA

#  TAG: memory_replacement_policy
# The memory replacement policy parameter determines which
# objects are purged from memory when memory space is needed.
#
# See cache_replacement_policy for details.
#
#Default:
 memory_replacement_policy lru


# LOGFILE PATHNAMES AND CACHE DIRECTORIES
# -----------------------------------------------------------------------------
```

```
#  TAG: cache_dir
# Usage:
#
# cache_dir Type Directory-Name Fs-specific-data [options]
#
# 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 "aufs" 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.
#
# The ufs store type:
#
# "ufs" is the old well-known Squid storage format that has always
# been there.
#
# cache_dir ufs Directory-Name Mbytes L1 L2 [options]
#
# '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.
#
# The aufs store type:
#
# "aufs" uses the same storage format as "ufs", utilizing
# POSIX-threads to avoid blocking the main Squid process on
# disk-I/O. This was formerly known in Squid as async-io.
#
# cache_dir aufs Directory-Name Mbytes L1 L2 [options]
#
# see argument descriptions under ufs above
#
# The diskd store type:
#
# "diskd" uses the same storage format as "ufs", utilizing a
# separate process to avoid blocking the main Squid process on
```

```
# disk-I/O.
#
# cache_dir diskd Directory-Name Mbytes L1 L2 [options] [Q1=n] [Q2=n]
#
# see argument descriptions under ufs above
#
# Q1 specifies the number of unacknowledged I/O requests when Squid
# stops opening new files. If this many messages are in the queues,
# Squid won't open new files. Default is 64
#
# Q2 specifies the number of unacknowledged messages when Squid
# starts blocking.  If this many messages are in the queues,
# Squid blocks until it recevies some replies. Default is 72
#
# Common options:
#
# read-only, this cache_dir is read only.
#
# max-size=n, refers to the max object size this storedir supports.
# It is used to initially choose the storedir to dump the object.
# Note: To make optimal use of the max-size limits you should order
# the cache_dir lines with the smallest max-size value first and the
# ones with no max-size specification last.
#
#Default:
 cache_dir ufs /var/spool/squid 100 16 256

#  TAG: cache_access_log
# Logs the client request activity.  Contains an entry for
# every HTTP and ICP queries received.
#
#Default:
 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.
#
#Default:
 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.
#
#Default:
 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 each
# '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 %s can be used in the file name then it will be replaced with a
# a representation of the cache_dir name where each / is replaced
# with '.'. This is needed to allow adding/removing cache_dir
# lines when cache_swap_log is being used.
#
# If have more than one 'cache_dir', and %s is not used in the name
# then 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.
#
#Default:
# none

#  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.
#
#Default:
 emulate_httpd_log off

#  TAG: log_ip_on_direct on|off
# Log the destination IP address in the hierarchy log tag when going
# direct. Earlier Squid versions logged the hostname here. If you
# prefer the old way set this to off.
#
#Default:
 log_ip_on_direct 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.
```

```
#
#Default:
 mime_table /usr/lib/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'.
#
#Default:
 log_mime_hdrs off

#  TAG: useragent_log
# Squid will write the User-Agent field from HTTP requests
# to the filename specified here.  By default useragent_log
# is disabled.
#
#Default:
 useragent_log /var/log/squid/useragent.log

#  TAG: referer_log
# Squid will write the Referer field from HTTP requests to the
# filename specified here.  By default referer_log is disabled.
#
#Default:
# none

#  TAG: pid_filename
# A filename to write the process-id to.  To disable, enter "none".
#
#Default:
 pid_filename /var/run/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".
#
#Default:
 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.
#
```

```
#Default:
 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'.
#
#Default:
 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).
#
#Default:
 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.
#
#Default:
 ftp_list_width 32

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

#  TAG: cache_dns_program
# Note: This option is only available if Squid is rebuilt with the
#       --disable-internal-dns option
#
# Specify the location of the executable for dnslookup process.
#
#Default:
# cache_dns_program /usr/lib/squid/
```

```
# cache_dns_program none


#  TAG: dns_children
# Note: This option is only available if Squid is rebuilt with the
#       --disable-internal-dns option
#
# 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.
#
#Default:
# dns_children 10

#  TAG: dns_retransmit_interval
# Initial retransmit interval for DNS queries. The interval is
# doubled each time all configured DNS servers have been tried.
#
#
#Default:
 dns_retransmit_interval 5 seconds

#  TAG: dns_timeout
# DNS Query timeout. If no response is received to a DNS query
# within this time then all DNS servers for the queried domain
# is assumed to be unavailable.
#
#Default:
 dns_timeout 5 minutes

#  TAG: dns_defnames on|off
# Note: This option is only available if Squid is rebuilt with the
#       --disable-internal-dns option
#
# 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.
#
#Default:
# 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
#
```

```
#Default:
# none

#  TAG: diskd_program
# Specify the location of the diskd executable.
# Note that this is only useful if you have compiled in
# diskd as one of the store io modules.
#
#Default:
 diskd_program /usr/lib/squid/diskd

#  TAG: unlinkd_program
# Specify the location of the executable for file deletion process.
#
#Default:
 unlinkd_program /usr/lib/squid/unlinkd

#  TAG: pinger_program
# Note: This option is only available if Squid is rebuilt with the
#       --enable-icmp option
#
# 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.
#
#Default:
# pinger_program /usr/lib/squid/

#  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.
#
#Default:
# 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.
#
#Default:
 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.
#
#Default:
# 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.
#
#Default:
# none

#  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
#
#Default:
# 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.
#
#Default:
 authenticate_children 7

#  TAG: authenticate_ttl
# The time a checked username/password combination remains cached.
# If a wrong password is given for a cached user, the user gets
# removed from the username/password cache forcing a revalidation.
#
#Default:
 authenticate_ttl 1 hour

#  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 to allow
# the user to redial without hassle. If all your users are
# stationary then higher values may be used.
#
# See also authenticate_ip_ttl_is_strict
#
#Default:
# authenticate_ip_ttl 0 seconds

#  TAG: authenticate_ip_ttl_is_strict
# This option makes authenticate_ip_ttl a bit stricted. With this
# enabled authenticate_ip_ttl will deny all access from other IP
# addresses until the TTL has expired, and the IP address "owning"
# the userid will not be forced to reauthenticate.
#
#Default:
# authenticate_ip_ttl_is_strict on


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

#  TAG: wais_relay_host
#  TAG: wais_relay_port
# Relay WAIS request to host (1st arg) at port (2 arg).
#
#Default:
# wais_relay_port 0

#  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.
#
#Default:
 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.
#
#Default:
```

```
 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.
#
#Default:
 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.
#
# Note, you must uncomment all the default lines if you want
# to change one. The default setting is only active if none is
# used.
#
#Default:
# refresh_pattern ^ftp: 1440 20% 10080
# refresh_pattern ^gopher: 1440 0% 1440
# refresh_pattern . 0 20% 4320

#  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.
#
#Default:
 reference_age 3 months

#  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.
#
#Default:
 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.
#
#Default:
 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.
#
#Default:
 positive_dns_ttl 6 hours
```

```
#  TAG: negative_dns_ttl time-units
# Time-to-Live (TTL) for negative caching of failed DNS lookups.
#
#Default:
 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 requested. (default)
#
#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).
#
#Default:
 connect_timeout 2 minutes

#  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.
#
#Default:
 peer_connect_timeout 30 seconds

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

```
 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.
#
#Default:
 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.
#
#Default:
 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.
#
#Default:
 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."
#
#Default:
 half_closed_clients off

#  TAG: pconn_timeout
```

```
# Timeout for idle persistent connections to servers and other
# proxies.
#
#Default:
 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.
#
#Default:
 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.
#
#Default:
 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 ...
# acl aclname ident_regex [-i] pattern ...
#   # 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 ...
# acl aclname proxy_auth_regex [-i] pattern ...
#   # 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.
#
# acl req_mime_type mime-type1 ...
#   # regex match agains the mime type of the request generated
#   # by the client. Can be used to detect file upload or some
#   # types HTTP tunelling requests.
#   # NOTE: This does NOT match the reply. You cannot use this
#   # to match the returned file type.
#
#Examples:
#acl myexample dst_as 1241
#acl password proxy_auth REQUIRED
#acl fileupload req_mime_type -i ^multipart/form-data$
#
#Recommended minimum configuration:
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 red_local src 192.168.1.0/24
acl hidrogeno src 193.146.99.249/255.255.255.255
acl sanpedro src 192.168.1.2/255.255.255.255
acl SSL_ports port 443 563
acl Safe_ports port 443
acl Safe_ports port 80 # http
acl Safe_ports port 21 # ftp
acl Safe_ports port 443 563 # https, snews
acl Safe_ports port 70 # gopher
acl Safe_ports port 210 # wais
acl Safe_ports port 1025-65535 # unregistered ports
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 Safe_ports port 901 # SWAT
acl purge method PURGE
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 ...
#
# NOTE on default values:
#
# If there are no "access" lines present, the default is to deny
# 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:

#http_access deny all
#
#Recommended minimum configuration:
#
# Only allow cachemgr access from localhost
http_access allow manager localhost
http_access deny manager
# Only allow purge requests from localhost
http_access allow purge localhost
http_access deny purge
# Deny requests to unknown ports
http_access deny !Safe_ports
# Deny CONNECT to other than SSL ports
http_access deny CONNECT !SSL_ports
#
# INSERT YOUR OWN RULE(S) HERE TO ALLOW ACCESS FROM YOUR CLIENTS
#
http_access allow red_local
http_access allow hidrogeno
http_access allow sanpedro
http_access allow localhost

# And finally deny all other access to this proxy
http_access deny all


#  TAG: icp_access
# Allowing or Denying access to the ICP port based on defined
# access lists
#
# icp_access  allow|deny [!]aclname ...
#
# See http_access for details
#
```

```
#Default:
 icp_access allow sanpedro
 icp_access allow hidrogeno
 icp_access deny all
#
#Allow ICP queries from eveyone
# 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.
#
#Default setting:
miss_access allow hidrogeno
miss_access allow sanpedro
miss_access allow red_local
miss_access allow localhost
miss_access deny 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://www.squid-cache.org/FAQ/FAQ-10.html).
#
#Default:
# none

#  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).
#
#Default:
# 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.
#
#Default:
 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."
#
#Default:
 cache_mgr webmaster

#  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 proxy and GID to proxy.
#
# 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.
#
#Default:
 cache_effective_user proxy
 cache_effective_group proxy

#  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.
#
#Default:
# none
```

```
#  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.
#
#Default:
# none

#  TAG: hostname_aliases
# A list of other DNS names that your cache has.
#
#Default:
# none


# 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://www.ircache.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.
#
#Default:
 announce_period 0
#
#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.
#
#Default:
# 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.
#
#Default:
httpd_accel_host virtual
httpd_accel_port 80

#  TAG: httpd_accel_single_host on|off
# If you are running Squid as a accelerator and have a single backend
# server then set this to on. This causes Squid to forward the request
# to this server irregardles of what any redirectors or Host headers
# says.
#
# Leave this at off if you have multiple backend servers, and use a
# redirector (or host table or private DNS) to map the requests to the
# appropriate backend servers. Note that the mapping needs to be a
# 1-1 mapping between requested and backend (from redirector) domain
# names or caching will fail, as cacing is performed using the
# URL returned from the redirector.
#
# See also redirect_rewrites_host_header.
#
#Default:
httpd_accel_single_host off
```

```
#  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'. Note however that your
# proxy users may have trouble to reach the accelerated domains
# unless their browsers are configured not to use this proxy for
# those domains (for example via the no_proxy browser configuration
# setting)
#
#Default:
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.
#
#Default:
httpd_accel_uses_host_header on


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

#  TAG: dns_testnames
# The DNS tests exit as soon as the first site is successfully looked up
#
# This test can be disabled with the -D command line option.
#
#Default:
# 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>'.
```

```
#
# Note2, for Debian/Linux the default of logfile_rotate is
# zero, since it includes external logfile-rotation methods.
#
#Default:
# logfile_rotate 0

#  TAG: append_domain
# Appends local domain name to hostnames without any dots in
# them.  append_domain must begin with a period.
#
#Example:
# append_domain .yourdomain.com
#
#Default:
# none

#  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.
#
#Default:
# 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.
#
#Default:
# none

#  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.
#
#Default:
# none
```

```
#  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.
#
#Default:
 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.
#
#Default:
 memory_pools_limit 30 MB

#  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
#
#Default:
 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.
#
#Default:
 log_icp_queries off

#  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'.
#
#Default:
 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.
#
#Default:
 minimum_direct_hops 4

#  TAG: minimum_direct_rtt
# If using the ICMP pinging stuff, do direct fetches for sites
# which are no more than this many rtt milliseconds away.
#
#Default:
 minimum_direct_rtt 400

#  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.
#
#Example:
# cachemgr_passwd secret shutdown
# cachemgr_passwd lesssssssecret info stats/objects
# cachemgr_passwd disable all
#
#Default:
# none

#  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.
#
#Default:
 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.
#
#Default:
# store_objects_per_bucket 20

#  TAG: client_db on|off
```

```
# If you want to disable collecting per-client statistics, then
# turn off client_db here.
#
#Default:
 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.
#
#Default:
# 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.
#
#Default:
 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.
#
#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.
#
#Default:
 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).
#
#Default:
 buffered_logs on

#  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.
#
#Default:
 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.
#
#Default:
# none
```

```
#  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.
#
#Default:
# none

#  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).
#
#Default:
# none

#  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!)
#
#Default:
# none

#  TAG: icon_directory
# Where the icons are stored. These are normally kept in
# /usr/lib/squid/icons
#
#Default:
# icon_directory /usr/lib/squid/icons

#  TAG: error_directory
# If you wish to create your own versions of the default
# (English) 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.
#
#Default:
 error_directory /usr/lib/squid/errors/Spanish

#  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.
#
#Default:
# 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.
#
#Default:
 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: on Debian/Linux, the default is zero - you need to
# set it to 3401 to enable it.
#
# NOTE: SNMP support requires use the --enable-snmp configure
# command line option.
#
#Default:
 snmp_port 0

#  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
#
#Default:
 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 snmp_incoming_address (0.0.0.0) is to listen on all
# available network interfaces.
#
# If snmp_outgoing_address is set to 255.255.255.255 (the default)
# then it will use the same socket as snmp_incoming_address. Only
# change this if you want to have SNMP replies sent using another
# address than where this Squid listens for SNMP queries.
#
# NOTE, snmp_incoming_address and snmp_outgoing_address can not have
# the same value since they both use port 3401.
#
#Default:
# snmp_incoming_address 0.0.0.0
# snmp_outgoing_address 255.255.255.255

#  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.
#
#Default:
# as_whois_server whois.ra.net
# as_whois_server whois.ra.net

#  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.
#
#Default:
 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.
#
#Default:
# 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.
#
#Default:
# wccp_incoming_address 0.0.0.0
# wccp_outgoing_address 255.255.255.255


# 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.
#
#Default:
# 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:
#
#Example:
# 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"
#
#Default:
# none

#  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:
#
#Example:
# delay_access 1 allow some_big_clients
# delay_access 1 deny all
# delay_access 2 allow lotsa_little_clients
# delay_access 2 deny all
#
#Default:
# none

#  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.
#
#Default:
# none

#  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).
#
#Default:
 delay_initial_bucket_level 50

#  TAG: incoming_icp_average
#  TAG: incoming_http_average
#  TAG: incoming_dns_average
#  TAG: min_icp_poll_cnt
#  TAG: min_dns_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!
#
#Default:
# incoming_icp_average 6
# incoming_http_average 4
# incoming_dns_average 4
# min_icp_poll_cnt 8
# min_dns_poll_cnt 8
# min_http_poll_cnt 8

#  TAG: max_open_disk_fds
# To avoid having disk as the I/O bottleneck Squid can optionally
# bypass the on-disk cache if more than this amount of disk file
# descriptors are open.
#
# A value of 0 indicates no limit.
#
#Default:
# max_open_disk_fds 0

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

#  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.
#
#Default:
# 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.
#
#Example:
# acl buggy_server url_regex ^http://....
# broken_posts allow buggy_server
#
#Default:
# none

#  TAG: mcast_miss_addr
# Note: This option is only available if Squid is rebuilt with the
#       -DMULTICAST_MISS_STREAM option
#
# 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.
#
#Default:
# mcast_miss_addr 255.255.255.255

#  TAG: mcast_miss_ttl
# Note: This option is only available if Squid is rebuilt with the
#       -DMULTICAST_MISS_TTL option
#
# 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.
#
#Default:
# mcast_miss_ttl 16

#  TAG: mcast_miss_port
# Note: This option is only available if Squid is rebuilt with the
#       -DMULTICAST_MISS_STREAM option
#
# This is the port number to be used in conjunction with
# 'mcast_miss_addr'.
#
#Default:
```

```
# mcast_miss_port 3135

#  TAG: mcast_miss_encode_key
# Note: This option is only available if Squid is rebuilt with the
#       -DMULTICAST_MISS_STREAM option
#
# The URLs that are sent in the multicast miss stream are
# encrypted.  This is the encryption key.
#
#Default:
# mcast_miss_encode_key XXXXXXXXXXXXXXXX

#  TAG: nonhierarchical_direct
# By default, Squid will send any non-hierarchical requests
# (matching hierarchy_stoplist or not cachable request type) direct
# to origin servers.
#
# If you set this to off, then Squid will prefer to send these
# requests to parents.
#
# Note that in most configurations, by turning this off you will only
# add latency to these request without any improvement in global hit
# ratio.
#
# If you are inside an firewall then see never_direct instead of
# this directive.
#
#Default:
# nonhierarchical_direct on

#  TAG: prefer_direct
# Normally Squid tries to use parents for most requests. If you by some
# reason like it to first try going direct and only use a parent if
# going direct fails then set this to off.
#
# By combining nonhierarchical_direct off and prefer_direct on you
# can set up Squid to use a parent as a backup path if going direct
# fails.
#
#Default:
# prefer_direct off

#  TAG: strip_query_terms
# By default, Squid strips query terms from requested URLs before
# logging.  This protects your user's privacy.
#
#Default:
# 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.
#
#Default:
# none

#  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.
#
#Default:
# redirector_bypass off

#  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'.
#
#Default:
# 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.
#
#Default:
# 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.
#
#Default:
# digest_bits_per_entry 5

#  TAG: digest_rebuild_period (seconds)
# This is the number of seconds between Cache Digest rebuilds.
#
#Default:
# digest_rebuild_period 1 hour

#  TAG: digest_rewrite_period (seconds)
# This is the number of seconds between Cache Digest writes to
```

```
# disk.
#
#Default:
# 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.
#
#Default:
# 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.
#
#Default:
# 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.
#
#Default:
# none

#  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.
#
#Default:
 client_persistent_connections on
 server_persistent_connections on

#  TAG: pipeline_prefetch
# To boost the performance of pipelined requests to closer
# match that of a non-proxied environment Squid tries to fetch
# up to two requests in parallell from a pipeline.
#
#Default:
# pipeline_prefetch on

#  TAG: extension_methods
# Squid only knows about standardized HTTP request methods.
# You can add up to 20 additional "extension" methods here.
#
```

```
#Default:
# none

#  TAG: high_response_time_warning (msec)
# If the one-minute median response time exceeds this value,
# Squid prints a WARNING with debug level 0 to get the
# administrators attention.  The value is in milliseconds.
#
#Default:
# high_response_time_warning 0

#  TAG: high_page_fault_warning
# If the one-minute average page fault rate exceeds this
# value, Squid prints a WARNING with debug level 0 to get
# the administrators attention.  The value is in page faults
# per second.
#
#Default:
# high_page_fault_warning 0

#  TAG: high_memory_warning
# If the memory usage (as determined by mallinfo) exceeds
# value, Squid prints a WARNING with debug level 0 to get
# the administrators attention.
#
#Default:
# high_memory_warning 0

#  TAG: store_dir_select_algorithm
# Set this to 'round-robin' as an alternative.
#
#Default:
# store_dir_select_algorithm least-load

#  TAG: forward_log
# Note: This option is only available if Squid is rebuilt with the
#       -DWIP_FWD_LOG option
#
# Logs the server-side requests.
#
# This is currently work in progress.
#
#Default:
# none

#  TAG: ie_refresh on|off
# Microsoft Internet Explorer up until version 5.5 Service
# Pack 1 has an issue with transparent proxies, wherein it
# is impossible to force a refresh.  Turning this on provides
# a partial fix to the problem, by causing all IMS-REFRESH
# requests from older IE versions to check the origin server
# for fresh content.  This reduces hit ratio by some amount
# (~10% in my experience), but allows users to actually get
```

```
# fresh content when they want it.  Note that because Squid
# cannot tell if the user is using 5.5 or 5.5SP1, the behavior
# of 5.5 is unchanged from old versions of Squid (i.e. a
# forced refresh is impossible).  Newer versions of IE will,
# hopefully, continue to have the new behavior and will be
# handled based on that assumption.  This option defaults to
# the old Squid behavior, which is better for hit ratios but
# worse for clients using IE, if they need to be able to
# force fresh content.
#
#Default:
 ie_refresh on
```

Ejemplo de fichero de configuración de Squid - Hidrogeno

Este archivo de configuración está pensado para *Hidrogeno*, servidor del Laboratorio F1. Este servidor actua como hermano de una jerarquía de cachés Squid y tiene control de ancho de banda.

```
# WELCOME TO SQUID 2
# ------------------
#
# This is the default Squid configuration file. You may wish
# to look at the Squid home page (http://www.squid-cache.org/)
# 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.
#
#Default:
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.
#
#Default:
icp_port 3130

#  TAG: htcp_port
# The port number where Squid sends and receives HTCP queries to
# and from neighbor caches.  To turn it on you want to set it 4827.
# By default it is set to "0" (disabled).
#
# To enable this option, you must use --enable-htcp with the
# configure script.
#
#Default:
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://www.squid-cache.org/FAQ/).
#
# Usage: mcast_groups 239.128.16.128 224.0.1.20
```

```
#
# By default, Squid doesn't listen on any multicast groups.
#
#Default:
# none

#  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.
#
# A *_incoming_address value of 0.0.0.0 indicates that Squid should
# listen on all available interfaces.
#
# If udp_outgoing_address is set to 255.255.255.255 (the default)
# then it will use the same socket as udp_incoming_address. Only
# change this if you want to have ICP queries sent using another
# address than where this Squid listens for ICP queries from other
# caches.
#
# NOTE, udp_incoming_address and udp_outgoing_address can not
# have the same value since they both use port 3130.
#
# NOTE, tcp_incoming_address has been removed.  You can now
# specify IP addresses on the 'http_port' line.
#
#Default:
# tcp_outgoing_address 255.255.255.255
# udp_incoming_address 0.0.0.0
# udp_outgoing_address 255.255.255.255


# OPTIONS WHICH AFFECT THE NEIGHBOR SELECTION ALGORITHM
# -----------------------------------------------------------------------------

#  TAG: cache_peer
# To specify other caches in a hierarchy, use the format:
#
# cache_peer hostname type http_port icp_port
#
# For example,
#
# #                                        proxy  icp
```

```
# #          hostname             type     port   port  options
# #          -------------------- -------- ----- -----  -----------
# cache_peer parent.foo.net       parent    3128  3130  [proxy-only]
# cache_peer sib1.foo.net         sibling   3128  3130  [proxy-only]
# cache_peer sib2.foo.net         sibling   3128  3130  [proxy-only]
#
#       type:  either 'parent', 'sibling', or 'multicast'.
#
# proxy_port:  The port number where the cache listens for proxy
#      requests.
#
#   icp_port:  Used for querying neighbor caches about
#      objects.  To have a non-ICP neighbor
#      specify '7' for the ICP port and make sure the
#      neighbor machine has the UDP echo port
#      enabled in its /etc/inetd.conf file.
#
#     options: proxy-only
#      weight=n
#      ttl=n
#      no-query
#      default
#      round-robin
#      multicast-responder
#      closest-only
#      no-digest
#      no-netdb-exchange
#      no-delay
#      login=user:password
#      connect-timeout=nn
#      digest-url=url
#      allow-miss
#
#      use 'proxy-only' to specify that objects fetched
#      from this cache should not be saved locally.
#
#      use 'weight=n' to specify a weighted parent.
#      The weight must be an integer.  The default weight
#      is 1, larger weights are favored more.
#
#      use 'ttl=n' to specify a IP multicast TTL to use
#      when sending an ICP queries to this address.
#      Only useful when sending to a multicast group.
#      Because we don't accept ICP replies from random
#      hosts, you must configure other group members as
#      peers with the 'multicast-responder' option below.
#
#      use 'no-query' to NOT send ICP queries to this
#      neighbor.
#
#      use 'default' if this is a parent cache which can
#      be used as a "last-resort." You should probably
#      only use 'default' in situations where you cannot
```

```
#       use ICP with your parent cache(s).
#
#       use 'round-robin' to define a set of parents which
#       should be used in a round-robin fashion in the
#       absence of any ICP queries.
#
#       'multicast-responder' indicates that the named peer
#       is a member of a multicast group.  ICP queries will
#       not be sent directly to the peer, but ICP replies
#       will be accepted from it.
#
#       'closest-only' indicates that, for ICP_OP_MISS
#       replies, we'll only forward CLOSEST_PARENT_MISSes
#       and never FIRST_PARENT_MISSes.
#
#       use 'no-digest' to NOT request cache digests from
#       this neighbor.
#
#       'no-netdb-exchange' disables requesting ICMP
#       RTT database (NetDB) from the neighbor.
#
#       use 'no-delay' to prevent access to this neighbor
#       from influencing the delay pools.
#
#       use 'login=user:password' if this is a personal/workgroup
#       proxy and your parent requires proxy authentication.
#
#       use 'connect-timeout=nn' to specify a peer
#       specific connect timeout (also see the
#       peer_connect_timeout directive)
#
#       use 'digest-url=url' to tell Squid to fetch the cache
#       digest (if digests are enabled) for this host from
#       the specified URL rather than the Squid default
#       location.
#
#       use 'allow-miss' to disable Squid's use of only-if-cached
#       when forwarding requests to siblings. This is primarily
#       useful when icp_hit_stale is used by the sibling. To
#       extensive use of this option may result in forwarding
#       loops, and you should avoid having two-way peerings
#       with this option. (for example to deny peer usage on
#       requests from peer by denying cache_peer_access if the
#       source is a peer)
#
# NOTE: non-ICP neighbors must be specified as 'parent'.
#
#Default:
cache_peer 193.146.99.248 parent 3128 3130 no-digest default
cache_peer 193.146.102.181 sibling 3128 3130 no-delay
cache_peer_domain 193.146.99.248 !unileon.es !192.168.2.2

#  TAG: cache_peer_domain
```

```
# Use to limit the domains for which a neighbor cache will be
# queried.  Usage:
#
# cache_peer_domain cache-host domain [domain ...]
# cache_peer_domain cache-host !domain
#
# For example, specifying
#
# cache_peer_domain parent.foo.net .edu
#
# has the effect such that UDP query packets are sent to
# 'bigserver' only when the requested object exists on a
# server in the .edu domain.  Prefixing the domainname
# with '!' means that the cache will be queried for objects
# NOT in that domain.
#
# NOTE: * Any number of domains may be given for a cache-host,
#   either on the same or separate lines.
# * When multiple domains are given for a particular
#   cache-host, the first matched domain is applied.
# * Cache hosts with no domain restrictions are queried
#   for all requests.
# * There are no defaults.
# * There is also a 'cache_peer_access' tag in the ACL
#   section.
#
#Default:
# none

#  TAG: neighbor_type_domain
# usage: neighbor_type_domain parent|sibling domain domain ...
#
# Modifying the neighbor type for specific domains is now
# possible.  You can treat some domains differently than the the
# default neighbor type specified on the 'cache_peer' line.
# Normally it should only be necessary to list domains which
# should be treated differently because the default neighbor type
# applies for hostnames which do not match domains listed here.
#
#EXAMPLE:
# cache_peer  parent cache.foo.org 3128 3130
# neighbor_type_domain cache.foo.org sibling .com .net
# neighbor_type_domain cache.foo.org sibling .au .de
#
#Default:
# none

#  TAG: icp_query_timeout (msec)
# Normally Squid will automatically determine an optimal ICP
# query timeout value based on the round-trip-time of recent ICP
# queries.  If you want to override the value determined by
# Squid, set this 'icp_query_timeout' to a non-zero value.  This
# value is specified in MILLISECONDS, so, to use a 2-second
```

```
# timeout (the old default), you would write:
#
# icp_query_timeout 2000
#
#Default:
icp_query_timeout 0

#  TAG: maximum_icp_query_timeout (msec)
# Normally the ICP query timeout is determined dynamically.  But
# sometimes it can lead to very large values (say 5 seconds).
# Use this option to put an upper limit on the dynamic timeout
# value.  Do NOT use this option to always use a fixed (instead
# of a dynamic) timeout value. To set a fixed timeout see the
# 'icp_query_timeout' directive.
#
#Default:
# maximum_icp_query_timeout 2000

#  TAG: mcast_icp_query_timeout (msec)
# For Multicast peers, Squid regularly sends out ICP "probes" to
# count how many other peers are listening on the given multicast
# address.  This value specifies how long Squid should wait to
# count all the replies.  The default is 2000 msec, or 2
# seconds.
#
#Default:
# mcast_icp_query_timeout 2000

#  TAG: dead_peer_timeout (seconds)
# This controls how long Squid waits to declare a peer cache
# as "dead."  If there are no ICP replies received in this
# amount of time, Squid will declare the peer dead and not
# expect to receive any further ICP replies.  However, it
# continues to send ICP queries, and will mark the peer as
# alive upon receipt of the first subsequent ICP reply.
#
# This timeout also affects when Squid expects to receive ICP
# replies from peers.  If more than 'dead_peer' seconds have
# passed since the last ICP reply was received, Squid will not
# expect to receive an ICP reply on the next query.  Thus, if
# your time between requests is greater than this timeout, you
# will see a lot of requests sent DIRECT to origin servers
# instead of to your parents.
#
#Default:
# dead_peer_timeout 10 seconds

#  TAG: hierarchy_stoplist
# A list of words which, if found in a URL, cause the object to
# be handled directly by this cache.  In other words, use this
# to not query neighbor caches for certain objects.  You may
# list this option multiple times.
#
```

```
#We recommend you to use at least the following line.
hierarchy_stoplist cgi-bin jsp asp ?

#  TAG: no_cache
# A list of ACL elements which, if matched, cause the reply to
# immediately removed from the cache.  In other words, use this
# to force certain objects to never be cached.
#
# You must use the word 'DENY' to indicate the ACL names which should
# NOT be cached.
#
#We recommend you to use the following two lines.
#acl PAGINA_CANCERBERO dst 193.146.99.248/255.255.248.0
acl PAGINA_HIDROGENO dst 193.146.99.249/255.255.255.255
acl QUERY urlpath_regex cgi-bin \?
#no_cache deny PAGINA_CANCERBERO
no_cache deny PAGINA_HIDROGENO
no_cache deny QUERY


# 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.
#
#Default:
cache_mem 30 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.
#
#Default:
 cache_swap_low 90
 cache_swap_high 93

#  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.
#
#Default:
 maximum_object_size 300000 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.
#
#Default:
 minimum_object_size 0 KB

#  TAG: maximum_object_size_in_memory (bytes)
#        Objects greater than this size will not be attempted to kept in
#        the memory cache. This should be set high enough to keep objects
#        accessed frequently in memory to improve performance whilst low
#        enough to keep larger objects from hoarding cache_mem .
#
#Default:
```

```
 maximum_object_size_in_memory 16 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.
#
#Default:
 ipcache_size 1024
 ipcache_low 90
 ipcache_high 93

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

#  TAG: cache_replacement_policy
# The cache replacement policy parameter determines which
# objects are evicted (replaced) when disk space is needed.
#
#     lru       : Squid's original list based LRU policy
#     heap GDSF : Greedy-Dual Size Frequency
#     heap LFUDA: Least Frequently Used with Dynamic Aging
#     heap LRU  : LRU policy implemented using a heap
#
# Applies to any cache_dir lines listed below this.
#
# The LRU policies keeps recently referenced objects.
#
# The heap 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 heap 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 the GDSF and LFUDA 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.
#
```

```
#Default:
cache_replacement_policy heap GDSF

#  TAG: memory_replacement_policy
# The memory replacement policy parameter determines which
# objects are purged from memory when memory space is needed.
#
# See cache_replacement_policy for details.
#
#Default:
 memory_replacement_policy lru


# LOGFILE PATHNAMES AND CACHE DIRECTORIES
# -----------------------------------------------------------------------------

#  TAG: cache_dir
# Usage:
#
# cache_dir Type Directory-Name Fs-specific-data [options]
#
# 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 "aufs" 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.
#
# The ufs store type:
#
# "ufs" is the old well-known Squid storage format that has always
# been there.
#
# cache_dir ufs Directory-Name Mbytes L1 L2 [options]
#
# '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.
```

```
#
# The aufs store type:
#
# "aufs" uses the same storage format as "ufs", utilizing
# POSIX-threads to avoid blocking the main Squid process on
# disk-I/O. This was formerly known in Squid as async-io.
#
# cache_dir aufs Directory-Name Mbytes L1 L2 [options]
#
# see argument descriptions under ufs above
#
# The diskd store type:
#
# "diskd" uses the same storage format as "ufs", utilizing a
# separate process to avoid blocking the main Squid process on
# disk-I/O.
#
# cache_dir diskd Directory-Name Mbytes L1 L2 [options] [Q1=n] [Q2=n]
#
# see argument descriptions under ufs above
#
# Q1 specifies the number of unacknowledged I/O requests when Squid
# stops opening new files. If this many messages are in the queues,
# Squid won't open new files. Default is 64
#
# Q2 specifies the number of unacknowledged messages when Squid
# starts blocking.  If this many messages are in the queues,
# Squid blocks until it recevies some replies. Default is 72
#
# Common options:
#
# read-only, this cache_dir is read only.
#
# max-size=n, refers to the max object size this storedir supports.
# It is used to initially choose the storedir to dump the object.
# Note: To make optimal use of the max-size limits you should order
# the cache_dir lines with the smallest max-size value first and the
# ones with no max-size specification last.
#
#Default:
 cache_dir ufs /var/spool/squid 1000 16 256

#  TAG: cache_access_log
# Logs the client request activity.  Contains an entry for
# every HTTP and ICP queries received.
#
#Default:
 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.
```

```
#
#Default:
 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.
#
#Default:
 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 each
# '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 %s can be used in the file name then it will be replaced with a
# a representation of the cache_dir name where each / is replaced
# with '.'. This is needed to allow adding/removing cache_dir
# lines when cache_swap_log is being used.
#
# If have more than one 'cache_dir', and %s is not used in the name
# then 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.
#
#Default:
# none

#  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.
#
```

```
#Default:
 emulate_httpd_log off

#  TAG: log_ip_on_direct on|off
# Log the destination IP address in the hierarchy log tag when going
# direct. Earlier Squid versions logged the hostname here. If you
# prefer the old way set this to off.
#
#Default:
 log_ip_on_direct 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.
#
#Default:
 mime_table /usr/lib/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'.
#
#Default:
 log_mime_hdrs off

#  TAG: useragent_log
# Squid will write the User-Agent field from HTTP requests
# to the filename specified here.  By default useragent_log
# is disabled.
#
#Default:
 useragent_log /var/log/squid/useragent.log

#  TAG: referer_log
# Squid will write the Referer field from HTTP requests to the
# filename specified here.  By default referer_log is disabled.
#
#Default:
# none

#  TAG: pid_filename
# A filename to write the process-id to.  To disable, enter "none".
#
#Default:
 pid_filename /var/run/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".
#
#Default:
 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.
#
#Default:
 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'.
#
#Default:
 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).
#
#Default:
 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.
#
#Default:
 ftp_list_width 32
```

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

#  TAG: cache_dns_program
# Note: This option is only available if Squid is rebuilt with the
#       --disable-internal-dns option
#
# Specify the location of the executable for dnslookup process.
#
#Default:
# cache_dns_program /usr/lib/squid/
# cache_dns_program none


#  TAG: dns_children
# Note: This option is only available if Squid is rebuilt with the
#       --disable-internal-dns option
#
# 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.
#
#Default:
# dns_children 10

#  TAG: dns_retransmit_interval
# Initial retransmit interval for DNS queries. The interval is
# doubled each time all configured DNS servers have been tried.
#
#
#Default:
 dns_retransmit_interval 5 seconds

#  TAG: dns_timeout
# DNS Query timeout. If no response is received to a DNS query
# within this time then all DNS servers for the queried domain
# is assumed to be unavailable.
#
#Default:
 dns_timeout 5 minutes

#  TAG: dns_defnames on|off
# Note: This option is only available if Squid is rebuilt with the
#       --disable-internal-dns option
#
# 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.
#
#Default:
# 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
#
#Default:
# none

#  TAG: diskd_program
# Specify the location of the diskd executable.
# Note that this is only useful if you have compiled in
# diskd as one of the store io modules.
#
#Default:
 diskd_program /usr/lib/squid/diskd

#  TAG: unlinkd_program
# Specify the location of the executable for file deletion process.
#
#Default:
 unlinkd_program /usr/lib/squid/unlinkd

#  TAG: pinger_program
# Note: This option is only available if Squid is rebuilt with the
#       --enable-icmp option
#
# 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.
#
#Default:
# pinger_program /usr/lib/squid/

#  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.
#
#Default:
# 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.
#
#Default:
 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.
#
#Default:
redirect_rewrites_host_header off

#  TAG: redirector_access
# If defined, this access list specifies which requests are
# sent to the redirector processes.  By default all requests
# are sent.
#
#Default:
# none

#  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
#
#Default:
# 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.
#
#Default:
 authenticate_children 7
```

```
#  TAG: authenticate_ttl
# The time a checked username/password combination remains cached.
# If a wrong password is given for a cached user, the user gets
# removed from the username/password cache forcing a revalidation.
#
#Default:
 authenticate_ttl 1 hour

#  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 to allow
# the user to redial without hassle. If all your users are
# stationary then higher values may be used.
#
# See also authenticate_ip_ttl_is_strict
#
#Default:
# authenticate_ip_ttl 0 seconds

#  TAG: authenticate_ip_ttl_is_strict
# This option makes authenticate_ip_ttl a bit stricted. With this
# enabled authenticate_ip_ttl will deny all access from other IP
# addresses until the TTL has expired, and the IP address "owning"
# the userid will not be forced to reauthenticate.
#
#Default:
# authenticate_ip_ttl_is_strict on


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

#  TAG: wais_relay_host
#  TAG: wais_relay_port
# Relay WAIS request to host (1st arg) at port (2 arg).
#
#Default:
# wais_relay_port 0

#  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.
#
#Default:
 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.
#
#Default:
 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.
#
#Default:
 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.
#
# Note, you must uncomment all the default lines if you want
# to change one. The default setting is only active if none is
# used.
#
#Default:
# refresh_pattern ^ftp: 1440 20% 10080
# refresh_pattern ^gopher: 1440 0% 1440
# refresh_pattern . 0 20% 4320

#  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.
#
#Default:
 reference_age 3 months

#  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.
#
#Default:
 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.
#
#Default:
 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.
#
#Default:
 positive_dns_ttl 6 hours

#  TAG: negative_dns_ttl time-units
# Time-to-Live (TTL) for negative caching of failed DNS lookups.
#
#Default:
 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 requested. (default)
#
#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).
#
#Default:
```

```
 connect_timeout 2 minutes

#  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.
#
#Default:
 peer_connect_timeout 30 seconds

#  TAG: siteselect_timeout time-units
# For URN to multiple URL's URL selection
#
#Default:
 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.
#
#Default:
 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.
#
#Default:
 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.
#
#Default:
 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."
#
#Default:
 half_closed_clients off

#  TAG: pconn_timeout
# Timeout for idle persistent connections to servers and other
# proxies.
#
#Default:
 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.
#
#Default:
 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.
#
#Default:
 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 ...
# acl aclname ident_regex [-i] pattern ...
#   # 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 ...
# acl aclname proxy_auth_regex [-i] pattern ...
#   # 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.
#
# acl req_mime_type mime-type1 ...
#   # regex match agains the mime type of the request generated
#   # by the client. Can be used to detect file upload or some
#   # types HTTP tunelling requests.
#   # NOTE: This does NOT match the reply. You cannot use this
#   # to match the returned file type.
#
#Examples:
#acl myexample dst_as 1241
#acl password proxy_auth REQUIRED
#acl fileupload req_mime_type -i ^multipart/form-data$
#
#Recommended minimum configuration:
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 red_local src 192.168.2.0/24
acl cancerbero src 193.146.99.248/255.255.255.255
acl sanpedro src 193.146.102.181/255.255.255.255
acl SSL_ports port 443 563
acl Safe_ports port 80 # http
```

```
acl Safe_ports port 21 # ftp
acl Safe_ports port 443 563 # https, snews
acl Safe_ports port 70 # gopher
acl Safe_ports port 210 # wais
acl Safe_ports port 1025-65535 # unregistered ports
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 Safe_ports port 901 # SWAT
acl purge method PURGE
acl CONNECT method CONNECT
#
# Limitando la velocidad de algunos archivos
#
# información obtenida de:
#
# http://www.tldp.org/HOWTO/Bandwidth-Limiting-HOWTO/index.html
#

#
# No queremos limitar las descargas en nuestra red local.
#
acl magic_words1 url_regex -i 192.168

#
# Queremos limitar la descarga de este tipo de archivos. No bloqueamos
# .html, .gif, .jpg y archivos similares porque por lo general
# no consumen demasiado ancho de banda.
#
# Ponga todo esto en una única línea
#
acl magic_words2 url_regex -i ftp .exe .EXE .mp3 .MP3 .vqf .VQF .tar.gz .TAR.GZ .tar .TAR .g

#
# Queremos limitar el ancho de banda durante el día permitiendo
# el ancho de banda completo durante la noche.
# ¡Cuidado! con el acl de abajo sus descargas se interrumpirán
# a las 23:59. Lea la FAQ si quiere envitarlo.
#
acl day time 09:00-23:59


#  TAG: http_access
# Allowing or Denying access based on defined access lists
#
# Access to the HTTP port:
# http_access allow|deny [!]aclname ...
#
# NOTE on default values:
#
# If there are no "access" lines present, the default is to deny
# 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:

#http_access deny all
#
#Recommended minimum configuration:
#
# Only allow cachemgr access from localhost
http_access allow manager localhost
http_access deny manager
# Only allow purge requests from localhost
http_access allow purge localhost
http_access deny purge
# Deny requests to unknown ports
http_access deny !Safe_ports
# Deny CONNECT to other than SSL ports
http_access deny CONNECT !SSL_ports
#
# INSERT YOUR OWN RULE(S) HERE TO ALLOW ACCESS FROM YOUR CLIENTS
#
http_access allow red_local
http_access allow cancerbero
http_access allow sanpedro
# And finally deny all other access to this proxy
http_access deny all


#  TAG: icp_access
# Allowing or Denying access to the ICP port based on defined
# access lists
#
# icp_access  allow|deny [!]aclname ...
#
# See http_access for details
#
#Default:
icp_access allow cancerbero
icp_access allow sanpedro
icp_access deny all
#
#Allow ICP queries from eveyone
# 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.
#
#Default setting:
miss_access allow cancerbero
miss_access allow sanpedro
miss_access allow red_local
miss_access deny 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://www.squid-cache.org/FAQ/FAQ-10.html).
#
#Default:
# none

#  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).
#
#Default:
# 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.
#
#Default:
 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."
#
#Default:
 cache_mgr webmaster

#  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 proxy and GID to proxy.
#
# 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.
#
#Default:
 cache_effective_user proxy
 cache_effective_group proxy

#  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.
#
#Default:
# none

#  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.
#
#Default:
# none

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

```
#Default:
# none


# 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://www.ircache.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.
#
#Default:
 announce_period 0
#
#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.
#
#Default:
# 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.
#
#Default:
# httpd_accel_host virtual
httpd_accel_host litio.sistemasop.ui
httpd_accel_port 80

#  TAG: httpd_accel_single_host on|off
# If you are running Squid as a accelerator and have a single backend
# server then set this to on. This causes Squid to forward the request
# to this server irregardles of what any redirectors or Host headers
# says.
#
# Leave this at off if you have multiple backend servers, and use a
# redirector (or host table or private DNS) to map the requests to the
# appropriate backend servers. Note that the mapping needs to be a
# 1-1 mapping between requested and backend (from redirector) domain
# names or caching will fail, as cacing is performed using the
# URL returned from the redirector.
#
# See also redirect_rewrites_host_header.
#
#Default:
httpd_accel_single_host off

#  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'. Note however that your
# proxy users may have trouble to reach the accelerated domains
# unless their browsers are configured not to use this proxy for
# those domains (for example via the no_proxy browser configuration
# setting)
#
#Default:
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.
#
#Default:
httpd_accel_uses_host_header on


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

#  TAG: dns_testnames
# The DNS tests exit as soon as the first site is successfully looked up
#
# This test can be disabled with the -D command line option.
#
#Default:
# 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>'.
#
# Note2, for Debian/Linux the default of logfile_rotate is
# zero, since it includes external logfile-rotation methods.
#
#Default:
# logfile_rotate 0

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

```
#Example:
# append_domain .yourdomain.com
#
#Default:
# none

#  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.
#
#Default:
# 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.
#
#Default:
# none

#  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.
#
#Default:
# none

#  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.
#
#Default:
 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.
#
#Default:
 memory_pools_limit 128 MB

#  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
#
#Default:
 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.
#
#Default:
 log_icp_queries off

#  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'.
```

```
#
#Default:
 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.
#
#Default:
 minimum_direct_hops 4

#  TAG: minimum_direct_rtt
# If using the ICMP pinging stuff, do direct fetches for sites
# which are no more than this many rtt milliseconds away.
#
#Default:
 minimum_direct_rtt 400

#  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.
#
#Example:
# cachemgr_passwd secret shutdown
# cachemgr_passwd lesssssssecret info stats/objects
# cachemgr_passwd disable all
#
#Default:
# none

#  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.
#
#Default:
 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.
#
#Default:
# store_objects_per_bucket 20

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

#  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.
#
#Default:
 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.
#
#Default:
# none

#  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.
#
#Default:
# none

#  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).
#
#Default:
# none
anonymize_headers deny User-Agent

#  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!)
#
#Default:
# none
fake_user_agent Mozilla/5.0 (X11; U; Linux i686; es-ES; rv:1.0rc2) Gecko/20020520

#  TAG: icon_directory
# Where the icons are stored. These are normally kept in
# /usr/lib/squid/icons
#
#Default:
# icon_directory /usr/lib/squid/icons

#  TAG: error_directory
# If you wish to create your own versions of the default
# (English) 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.
#
#Default:
 error_directory /usr/lib/squid/errors/Spanish

#  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.
#
```

```
#Default:
# 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.
#
#Default:
 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: on Debian/Linux, the default is zero - you need to
# set it to 3401 to enable it.
#
# NOTE: SNMP support requires use the --enable-snmp configure
# command line option.
#
#Default:
 snmp_port 0

#  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
#
#Default:
 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 snmp_incoming_address (0.0.0.0) is to listen on all
# available network interfaces.
#
# If snmp_outgoing_address is set to 255.255.255.255 (the default)
# then it will use the same socket as snmp_incoming_address. Only
# change this if you want to have SNMP replies sent using another
# address than where this Squid listens for SNMP queries.
#
# NOTE, snmp_incoming_address and snmp_outgoing_address can not have
# the same value since they both use port 3401.
#
#Default:
# snmp_incoming_address 0.0.0.0
# snmp_outgoing_address 255.255.255.255

#  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.
#
#Default:
# as_whois_server whois.ra.net
# as_whois_server whois.ra.net

#  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.
#
#Default:
 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.
#
#Default:
# 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.
#
#Default:
# wccp_incoming_address 0.0.0.0
# wccp_outgoing_address 255.255.255.255


# 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.
#
#Default:
# delay_pools 0

# Tenemos dos delay_pools diferentes
# Acuda a la documentación de Squid para familiarizarse
# con delay_pools y delay_class.
#
delay_pools 2


#  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:
#
#Example:
# 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"
#
#Default:
# none

# Primer delay pool
# No queremos retrasar nuestro tráfico local
# Hay tres cases de pools; aquí sólo hablaremos de la segunda.
# Primera clase de retraso (1) de segundo tipo (2).
#
delay_class 1 2


# Segundo delay pool.
# Queremos retrasar la descarga de los archivos mencionados en magic_words2.
# Segunda clase de retraso (2) de segundo tipo (2).
#
delay_class 2 2


#  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:
#
#Example:
# delay_access 1 allow some_big_clients
# delay_access 1 deny all
# delay_access 2 allow lotsa_little_clients
# delay_access 2 deny all
#
#Default:
# none

# magic_words1: 192.168 que ya hemos puesto antes
#
delay_access 1 allow magic_words1

# Los números siguientes son valores en bytes;
# Debemos recordar que Squid no tiene en cuenta los bits de inicio/parada
# 5000/150000 son valores para la red al completo
```

```
# 5000/120000 son valores para la IP independiente
# una vez los archivos descargados exceden los 150000 bytes,
# (o el doble o el triple)
# las descargas proseguirán a 5000 bytes/s
#
delay_parameters 2 5000/150000 5000/120000

# Ya hemos configurado antes el día de 09:00 a 23:59.
#
delay_access 2 allow day
delay_access 2 deny !day
delay_access 2 allow magic_words2


#  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.
#
#Default:
# none

# -1/-1 significa que no hay límites.
#
delay_parameters 1 -1/-1 -1/-1


#  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).
#
#Default:
 delay_initial_bucket_level 50

#  TAG: incoming_icp_average
#  TAG: incoming_http_average
#  TAG: incoming_dns_average
#  TAG: min_icp_poll_cnt
#  TAG: min_dns_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!
#
#Default:
# incoming_icp_average 6
# incoming_http_average 4
# incoming_dns_average 4
# min_icp_poll_cnt 8
# min_dns_poll_cnt 8
# min_http_poll_cnt 8

#  TAG: max_open_disk_fds
```

```
# To avoid having disk as the I/O bottleneck Squid can optionally
# bypass the on-disk cache if more than this amount of disk file
# descriptors are open.
#
# A value of 0 indicates no limit.
#
#Default:
# max_open_disk_fds 0

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

#  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.
#
#Default:
# 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.
#
```

```
#Example:
# acl buggy_server url_regex ^http://....
# broken_posts allow buggy_server
#
#Default:
# none

#  TAG: mcast_miss_addr
# Note: This option is only available if Squid is rebuilt with the
#       -DMULTICAST_MISS_STREAM option
#
# 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.
#
#Default:
# mcast_miss_addr 255.255.255.255

#  TAG: mcast_miss_ttl
# Note: This option is only available if Squid is rebuilt with the
#       -DMULTICAST_MISS_TTL option
#
# 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.
#
#Default:
# mcast_miss_ttl 16

#  TAG: mcast_miss_port
# Note: This option is only available if Squid is rebuilt with the
#       -DMULTICAST_MISS_STREAM option
#
# This is the port number to be used in conjunction with
# 'mcast_miss_addr'.
#
#Default:
# mcast_miss_port 3135

#  TAG: mcast_miss_encode_key
# Note: This option is only available if Squid is rebuilt with the
#       -DMULTICAST_MISS_STREAM option
#
# The URLs that are sent in the multicast miss stream are
# encrypted.  This is the encryption key.
#
#Default:
# mcast_miss_encode_key XXXXXXXXXXXXXXXX

#  TAG: nonhierarchical_direct
# By default, Squid will send any non-hierarchical requests
```

```
# (matching hierarchy_stoplist or not cachable request type) direct
# to origin servers.
#
# If you set this to off, then Squid will prefer to send these
# requests to parents.
#
# Note that in most configurations, by turning this off you will only
# add latency to these request without any improvement in global hit
# ratio.
#
# If you are inside an firewall then see never_direct instead of
# this directive.
#
#Default:
# nonhierarchical_direct on

#  TAG: prefer_direct
# Normally Squid tries to use parents for most requests. If you by some
# reason like it to first try going direct and only use a parent if
# going direct fails then set this to off.
#
# By combining nonhierarchical_direct off and prefer_direct on you
# can set up Squid to use a parent as a backup path if going direct
# fails.
#
#Default:
# prefer_direct off

#  TAG: strip_query_terms
# By default, Squid strips query terms from requested URLs before
# logging.  This protects your user's privacy.
#
#Default:
# 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.
#
#Default:
# none

#  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.
#
#Default:
# redirector_bypass off

#  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'.
#
#Default:
# 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.
#
#Default:
# 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.
#
#Default:
# digest_bits_per_entry 5

#  TAG: digest_rebuild_period (seconds)
# This is the number of seconds between Cache Digest rebuilds.
#
#Default:
# digest_rebuild_period 1 hour

#  TAG: digest_rewrite_period (seconds)
# This is the number of seconds between Cache Digest writes to
# disk.
#
#Default:
# 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.
#
#Default:
# 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.
#
#Default:
# 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.
#
#Default:
# none

#  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.
#
#Default:
 client_persistent_connections on
 server_persistent_connections on

#  TAG: pipeline_prefetch
# To boost the performance of pipelined requests to closer
# match that of a non-proxied environment Squid tries to fetch
# up to two requests in parallell from a pipeline.
#
#Default:
# pipeline_prefetch on

#  TAG: extension_methods
# Squid only knows about standardized HTTP request methods.
# You can add up to 20 additional "extension" methods here.
#
#Default:
# none

#  TAG: high_response_time_warning (msec)
# If the one-minute median response time exceeds this value,
# Squid prints a WARNING with debug level 0 to get the
# administrators attention.  The value is in milliseconds.
#
#Default:
# high_response_time_warning 0

#  TAG: high_page_fault_warning
# If the one-minute average page fault rate exceeds this
# value, Squid prints a WARNING with debug level 0 to get
```

```
# the administrators attention.  The value is in page faults
# per second.
#
#Default:
# high_page_fault_warning 0

#  TAG: high_memory_warning
# If the memory usage (as determined by mallinfo) exceeds
# value, Squid prints a WARNING with debug level 0 to get
# the administrators attention.
#
#Default:
# high_memory_warning 0

#  TAG: store_dir_select_algorithm
# Set this to 'round-robin' as an alternative.
#
#Default:
# store_dir_select_algorithm least-load

#  TAG: forward_log
# Note: This option is only available if Squid is rebuilt with the
#       -DWIP_FWD_LOG option
#
# Logs the server-side requests.
#
# This is currently work in progress.
#
#Default:
# none

#  TAG: ie_refresh on|off
# Microsoft Internet Explorer up until version 5.5 Service
# Pack 1 has an issue with transparent proxies, wherein it
# is impossible to force a refresh.  Turning this on provides
# a partial fix to the problem, by causing all IMS-REFRESH
# requests from older IE versions to check the origin server
# for fresh content.  This reduces hit ratio by some amount
# (~10% in my experience), but allows users to actually get
# fresh content when they want it.  Note that because Squid
# cannot tell if the user is using 5.5 or 5.5SP1, the behavior
# of 5.5 is unchanged from old versions of Squid (i.e. a
# forced refresh is impossible).  Newer versions of IE will,
# hopefully, continue to have the new behavior and will be
# handled based on that assumption.  This option defaults to
# the old Squid behavior, which is better for hit ratios but
# worse for clients using IE, if they need to be able to
# force fresh content.
#
#Default:
 ie_refresh on
```

Prueba

Una vez adaptado el archivo de configuración de Squid a nuestras necesidades, estamos en disposición de ejecutar el programa. Para ello tecleamos:

```
# /etc/init.d/squid restart
```

Si por cualquier razón, una vez arrancado Squid, se modifica el archivo de configuración, no hace falta reiniciar Squid, simplemente le decimos que vuelva a leer el archivo de configuración:

```
# /etc/init.d/squid reload
```

Si a partir de este momento notas que la velocidad de navegación ha incrementado, Squid está haciendo bien su trabajo ;-)

Más información

Para saber todas las características y posibilidades que ofrece Squid, así como modificar o añadir opciones al mismo, se recomienda la lectura de las páginas del manual que acompañan al programa así como su documentación. Las FAQs (http://www.squid-cache.org/Doc/FAQ/FAQ.html) y la Guía de configuración (http://squid.visolve.com/squid24s1/contents.htm) son de obligada lectura.

De todas formas, la mejor documentación es el propio archivo de configuración como se ha podido observar.

Otros enlaces de interés pueden ser:

- Página principal de Squid (http://www.squid-cache.org)
- Limitar el ancho de banda COMO (http://mural.uv.es/~joferna/doc/Limitar-ancho-de-banda-COMO/html/)
- Enhancement and Validation of Squid's Cache Replacement Policy (http://www.hpl.hp.com/techreports/1999/HPL-1999-69.html)

Sobre este documento

Este documento ha sido escrito en formato XML utilizando la DTD de DocBook (http://www.docbook.org). Mediante este sistema, puede ser fácilmente transformado a múltiples formatos (HTML, TXT, PDF, PostScript, LaTeX, DVI, ...). Se recomienda su utilización como herramienta de documentación potente y libre.

Notas

1. Netfilter (http://www.netfilter.org)
2. Para que esto funcione realmente, es necesario configurar el cortafuegos de una manera determinada. Se verá en la documentación del cortafuegos