69 - SQUID - Proxy Server

SQUID2 - Proxy Server
Users Manual

Table of Contents (see SQUID_Proxy_Server_Users_Manual.sdw for the details)

Notes:

- When Squid starts:
  - Squid reads the /etc/host.conf to find the sequence of name resolve
  - If `oder hosts bind` then it takes a snapshot of the `/etc/hosts`
  - it then checks if the directive dns_nammeservers exist in `/etc/squid.conf`
    - if yes then it uses these DNS Servers for name resolution
    - otherwise it uses the `/etc/resolv.conf` if it exists.
    - It then starts some DNS-Serving-Daemons called [dnsserver]

- **How can I purge an object from my cache?**

  Squid does not allow you to purge objects unless it is configured with access controls in `squid.conf`. First you must add something like

```
acl PURGE method purge
acl localhost src 127.0.0.1
http_access allow purge localhost
http_access deny purge
```

  The above only allows purge requests which come from the local host and denies all other purge requests.

  To purge an object, you can use the client program:

```
client -m PURGE http://www.miscreant.com/
```

  If the purge was successful, you will see a ``200 OK'' response:

```
HTTP/1.0 200 OK
Date: Thu, 17 Jul 1997 16:03:32 GMT
Server: Squid/1.1.14
```

  If the object was not found in the cache, you will see a ``404 Not Found'' response:

```
HTTP/1.0 404 Not Found
Date: Thu, 17 Jul 1997 16:03:22 GMT
Server: Squid/1.1.14
```

- **How can I assign different bandwidth to different squid clients?**

Squid has the possibility to limit the `http/https/ftp` bandwidth via 'delay pools'.
The procedure is as follows:

- Assignment of acl aliases to groups of hosts.
- Create 'delay pools' and assign a class type to each pool.
- Assign bandwidth parameters to each pool
- Assign each acl group to a 'delay pool'.

Example No1:
We want to assign two bandwidth:

- 512kbit/s(64000 Bytes/s) for the office hosts
- 128kbit/s(16000 Bytes/s) for hosts of rest of the company

Configuration in `/etc/squid.conf`:

```
acl office_hosts src 192.168.71.10 192.168.71.44 192.168.71.242
acl intranet src 192.168.71.0/255.255.255.0

# Create 2 delay pools
delay_pools 2
# Assign a class type to each delay pool
#  Class 1 has class type 2
delay_class 1 2
#  Class 2 has class type 2
delay_class 2 2
# Assign bandwidth to each pool
# Pool 1 get no overall limit(-1/-1) but each host is limited to max 64KBytes/sec.
delay_parameters 1 -1/-1 64000/64000
# Pool 2 get no overall limit(-1/-1) but each host is limited to max 16KBytes/sec.
delay_parameters 2 -1/-1 16000/16000
# Assign each pool to the hosts groups
delay_access 1 allow office_hosts
delay_access 2 allow intranet
```

Example No2:
We want to assign a single bandwidth limit:

- 8KBytes/s to a all hosts in 2 subnets (192.168.70.0/24 and 192.168.71.0/24)
- No bandwidth limit for some special hosts in those 2 subnets.

Configuration in `/etc/squid.conf`:

```
# Assign acl to host groups
acl intranets src 192.168.70.0/255.255.254.0
acl special_hosts src 192.168.71.10 192.168.71.42 192.168.71.44\
     192.168.71.130 192.168.70.10 192.168.70.130 192.168.71.201
# Create 1 delay pool
delay_pools 1
# Assign the class 2 to pool 1
delay_class 1 2
# Assign bandwidth parameters to pool 1
# Pool 1 gets no overall limit(-1/-1) but each host is limited to 8KBytes/s
delay_parameters 1 -1/-1 8000/8000
# Assign pool 1 to hosts group(acls)
```

Notice that `special_hosts` are simply denied access to the pool1 therefore getting no bandwidth limit. To achieve this, the deny rule must be declared first.

```
delay_access 1 deny special_hosts
delay_access 1 allow intranets
```

This above examples are some of the most useful configurations I've found so far.
The following squid configuration file of SuSE 9.x explains in more details the different possibilities.

```
# 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.
#                       (Individual hosts in a class C network)
#
#     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
#                       (Overall limit for each Class C subnet)
#                    and a "individual" bucket chosen
#                       from bits 17 through 32 of the IP address.
#                       (Individual limit for each host in a class C network
#
#     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
#                 (Overall bandwidth for the sum all hosts of the pool)
#                 (class 1, 2, 3).
#
#     network     the "delay parameters" for the network buckets
#                 (class 3 only).
#
#     individual  the "delay parameters" for the individual buckets
#                 (Bandwidth for each individual host)
#                 (class 2 and 3 only).
#
#     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,
#                  - 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
```