

BANDWIDTH MANAGEMENT GUIDELINES

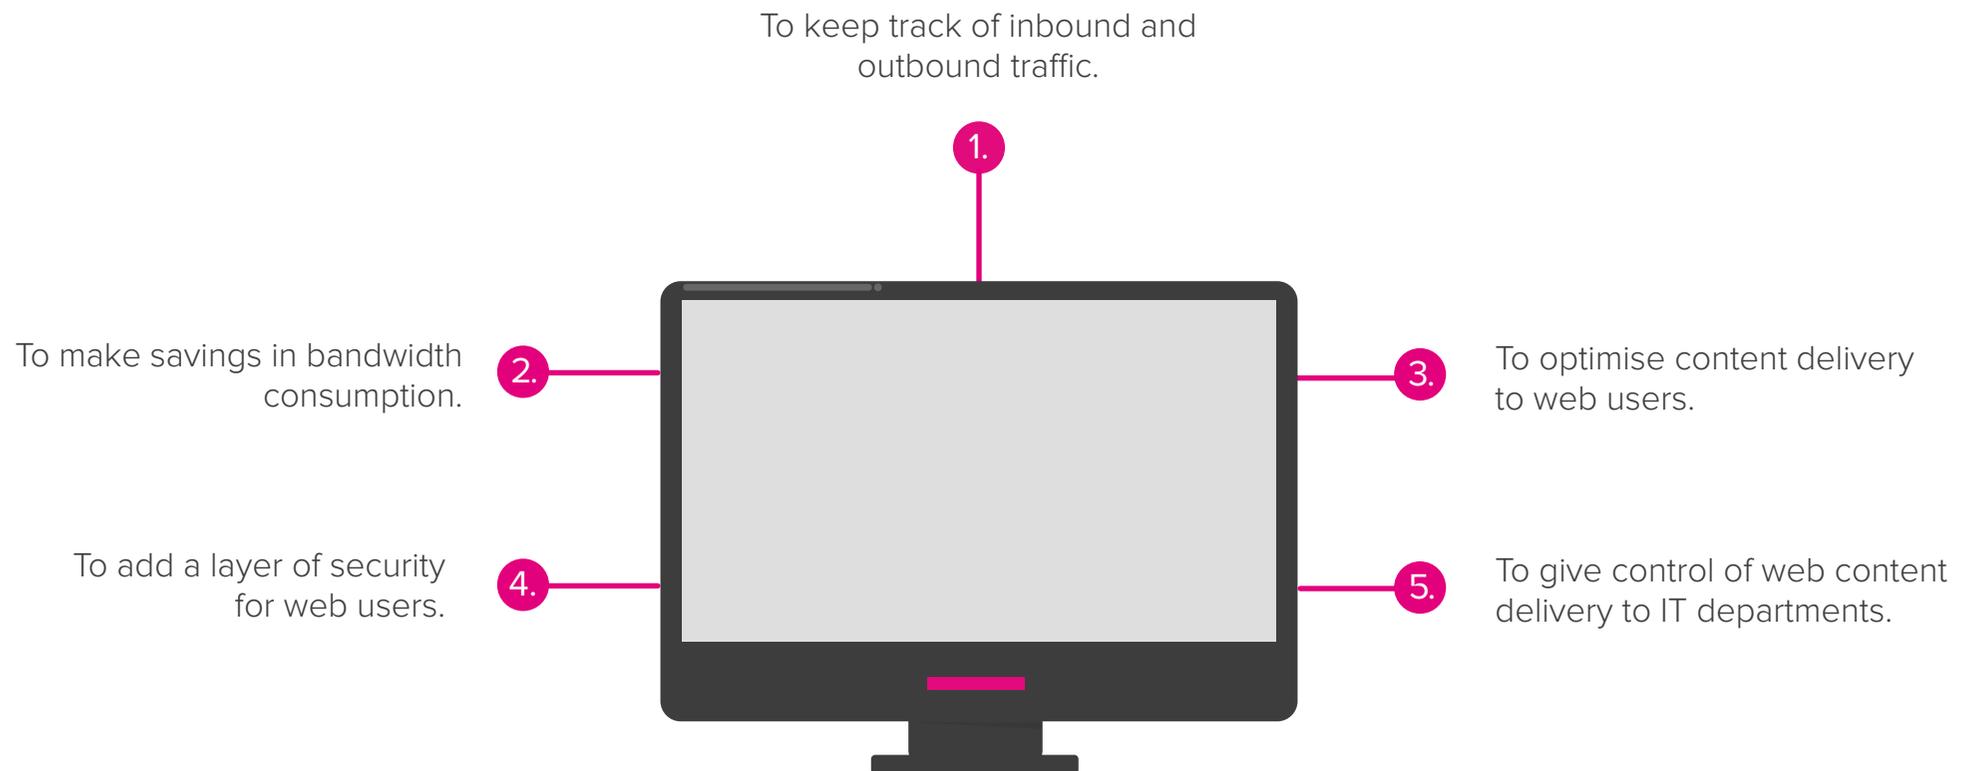
CONTENTS

3	Scope of this document
4	Why use a caching proxy server as part of an internet gateway?
5	Hardware guidelines for a squid proxy cache <ul style="list-style-type: none">- CPU- RAM- HDD- INTERFACES
6 - 9	Squid Caching Based Products <ul style="list-style-type: none">- Configuration Options- Description of refresh patterns- <i>Options</i>
10	Websense Content Gateway Proxy Product <ul style="list-style-type: none">- Forcing object caching
11	Deployment Options
12	Maintenance and Monitoring
13	References and further reading

SCOPE OF THIS DOCUMENT

1. — A guide to understanding caching proxies and their role in optimising Lobster Ink media delivery on a Local Area Network (LAN).
2. — A basic outline of how to cache dynamic streaming content and deploy the proxy server on a LAN.
3. — Advice on products which can provide a simplified caching proxy setup.
4. — This document does not give an exhaustive list of all the configurations available for proxy caching and LAN gateway services.

WHY USE A CACHING PROXY SERVER AS PART OF AN INTERNET GATEWAY?



HARDWARE GUIDELINES FOR A SQUID PROXY CACHE

For an organisation with thousands of users it is advisable to have quite a powerful proxy, or even multiple proxy servers, load balanced.

For a high load scenario the following server grade specifications would be advisable:

CPU

- QuadCore Xeon processor such as the intel E6540.

RAM

- 8 to 16 GB RAM.

HDD

- 250GB 7200 rpm SATA disc for operating system, binaries, reports and logs.
- x2 1TB drives in a hardware RAID 0+1 configuration for on disc cache storage.

INTERFACES

- X2 Gigabit ethernet interfaces.

SQUID CACHING BASED PRODUCTS

These are mainly appliance solutions which combine different technologies to make bandwidth management easier for organisations:

<http://www.unveiltech.com/>

<http://www.bivio.net/products/al/>

<http://www.clearcenter.com/Solution/distributed-enterprise.html>

Configuration Options

The basic configuration that comes with a fresh squid install is a bare minimum, the capabilities of a squid binary installation are limited by the compile options of the actual binary.

It is advisable in extremely large environments (over 50 000 clients) to use peer caches and compile the squid binaries with options and optimisations for specific hardware. In most organisational environments however the binaries will work well.

The possibilities for configuring a proxy range from simply caching common mime types to adblocking, url redirection, virus scanning and delay pooling.

The default algorithm for removal of cached content with Squid is LRU (Least recently used) and is calculated dynamically based on the amount of disc available to the cache.

The configuration we are interested in for enabling caching of streaming media from Lobster Ink services is the refresh patterns, cache dir and acl directives explained below.

SQUID CACHING BASED PRODUCTS

Description of refresh patterns

'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 object's 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** enforces min age even if the server sent an 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 were 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.

SQUID CACHING BASED PRODUCTS

Basically a cached object is: (the order is changed from 1.1.X)

- FRESH if expires
- FRESH if lm-factor
- 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.

acl manager proto cache_object
acl localhost src 127.0.0.1/32 ::1
acl to_localhost dst 127.0.0.0/8 0.0.0.0/32 ::1
acl SSL_ports port 443
acl Safe_ports port 80 # http
acl Safe_ports port 21 # ftp
acl Safe_ports port 443 # https
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 CONNECT method CONNECT
http_access allow manager localhost
http_access deny manager
http_access deny !Safe_ports
http_access deny CONNECT !SSL_ports
http_access allow localhost
For basic authentication
auth_param basic program /usr/lib/squid3/ncsa_auth /etc/squid3/passwd
acl tester proxy_auth lobsternet
http_access allow tester

SQUID CACHING BASED PRODUCTS

Deny all clients which don't authenticate – not necessary for a LAN – make an acl for the LAN
http_access deny all
http_port 3128
coredump_dir /var/spool/squid3
refresh_pattern ^ftp: 1440 20% 10080
refresh_pattern ^gopher: 1440 0% 1440
refresh_pattern -i (/cgi-bin/ \?) 0 0% 0
refresh_pattern (ReleasePackages(.gz)*)\$ 0 20% 2880
refresh_pattern . 0 20% 4320
Video Caching for cloud front
Enable queries for training videos
acl QUERY2 urlpath_regex video\? cloudfront\?
cache allow QUERY2
acl lobsterink dstdomain .cloudfront.net
cache allow lobster

Store training videos forever
refresh_pattern ^ftp: 1440 20% 10080
refresh_pattern -i \.flv\$ 10080 90% 999999 ignore-no-cache override-expire ignore-private
refresh_pattern -i \.mp4\$ 10080 90% 999999 ignore-no-cache override-expire ignore-private
refresh_pattern -i (/cgi-bin/ \?) 0 0% 0
refresh_pattern . 0 0% 4320
Enables skipping of very small redirect files and complete caching of large files
quick_abort_min -1 KB
maximum_object_size 1 GB
minimum_object_size 512 bytes
CACHE DIRECTORY File System Type, Path, Size in Mbytes, 1st level folders, 2nd level folders
cache_dir ufs /var/spool/squid3/ 500000 16 256

WEBSense CONTENT GATEWAY PROXY PRODUCT

Forcing object caching

You can force Content Gateway to cache specific URLs (including dynamic URLs) for a specified duration regardless of **Cache-Control** response headers.

1. Navigate to Configure > Protocols > HTTP > Cacheability.
2. Click **Edit File** at the end of the page to display the configuration file editor for the **cache.config** file.
3. In the fields provided, supply the following information:
 - From the **Rule Type** drop-down box, select **tll-in-cache**.
 - From the Primary Destination Type drop-down box, select url_regex.
 - In the **Primary Destination Value** field, specify the URL you want to force cache.
 - In the **Time Period** field, specify the amount of time that the proxy can serve the URL from the cache.
 - In addition, you can add secondary specifiers (such as **Prefix** and **Suffix**) to the rule. All the fields are described in HTTP.
4. Click **Add**, and then click **Apply**.
5. Click **Close**.

(See http://www.websense.com/content/support/library/web/v75/wcg_help/forcec.aspx)

DEPLOYMENT OPTIONS

There are many ways to deploy a proxy to a LAN, the most simplified examples are outlined below:

- **Non-transparent proxy** – A proxy that applications need to be aware of in order to use. This is probably the safest system to use as it enables easier load-balancing, failover and can be deployed to specific subset's of users rather than the entire network.
- **Transparent proxies** – This is when users requests are pushed through a proxy via the firewall and the traffic is all relayed through the proxy. This setup is most commonly used by ISPs for bandwidth optimising.

The simplest method would be to make use of a pac/wpad file served on the LAN allocated to hosts by DHCP. This is achieved by generating a pac file and serving it on the LAN over http. The DHCP will then update all clients connecting to the network with the location of the pac file.

A proxy pac file is simple javascript to inform clients about how to use the proxy.

MAINTENANCE AND MONITORING

It is important to make sure the cache is healthy, a Squid cache directory is created by running `squid -z` cached can be cleared by removal at file system level and regenerated should they become corrupted. Basic system monitoring is important to keep track of load, disc and interface resources.

There are some very powerful open source tools that can be used to get statistics:

- Squid analyzer (<http://squidalyzer.darold.net/>)
- Munin (<http://munin-monitoring.org/>)
- Monit (<http://mmonit.com/monit/>)

For the most part the cache will look after itself but as with all IT infrastructure it is important to have health checks in place.

REFERENCES AND FURTHER READING

These contain a more exhaustive list of explanations and implementation of specific technologies:

- Complete configuration options for squid proxy servers:

<http://www.squid-cache.org/>

<http://www.vicomsoft.com/learning-center/lan-based-web-caching/>

http://wiki.clug.org.za/wiki/Stephen_Murcott

<http://artemisa.unicauca.edu.co/~dabravo/squid24s1/contents.htm>

- To understand how to implement wpad and pac files on a LAN:

http://www.websense.com/content/support/library/web/v76/pac_file_best_practices/PAC_best_pract.aspx

<http://technet.microsoft.com/en-us/library/dd361887.aspx>

- To generate the proxy pac file:

<http://sourceforge.net/projects/pactory/?source=dlp>

<http://wiki.safesquid.com/hardware-requirements>

<http://ws.edu.isoc.org/data/2006/1156601184448218a95833f/squidhardwareconfiguration.htm>

- For the Websense Content Gateway product:

http://www.websense.com/content/support/library/web/v75/wcg_deploy/first.aspx

http://www.websense.com/content/support/library/web/v75/wcg_help/forcec.aspx