

Replacing Squid with ATS

ApacheCon 2015, Austin TX

Kit Chan (kichan@yahoo-inc.com)
Chiru Jaladi (chiru@yahoo-inc.com)

Before We Begin



This Talk

- 1) Not about how to choose a proxy server
- 2) Does not cover all Squid Use Cases

Agenda

History

Reasons

Details

Use Cases

Learnings

Squid in Yahoo

The year is 2006

- Dawn of SOA/Web Service in Yahoo!
- Squid improves performance through caching
 - Other benefits - routing & ACL

The Yahoo! logo is displayed in white, bold, sans-serif capital letters on a solid purple rectangular background.

Squid in Yahoo

- Mark Nottingham was the Champion of Squid in Yahoo!
- Collapsed Forwarding added to Squid 2.6
- SWR/SIE added to Squid 2.7



Squid in Yahoo

- Squid 3.0 - Rewrite of Squid in C++
 - ESI
 - ICAP
- Squid 3.2 - multiple worker support
- Backward Incompatibilities
 - No Collapse Forwarding till 3.5+
 - No SIE till 3.2+
 - Still no SWR
 - BLOCKERS!!!

ATS in Yahoo



I n k t o m i ®

Inktomi

- TS 2.0 - 1998, 3.0 - 1999, 4.0 - 2000
- Customers - AOL, @Home
- e.g. - Transcoding images to smaller sizes for AOL dialup users

ATS in Yahoo



YTS

- Inktomi Acquired by Yahoo - late 2002 / early 2003
- Renamed to YTS - Efforts resumed around late 2005
- By Feb 2010, served 30 billion objects, 400 terabytes a day for Yahoo

ATS in Yahoo



ASF

- Preparation - 700K lines of code change, 9 Months
- Apache Incubator in July 2009
- TLP in April 21, 2010

Why?

Cost of maintenance

Unresolved Blockers to Upgrade

Performance Limitation in 2.7

Details - Configuration

Squid - squid.conf

```
http_port 3128 vhost http11
```

```
cache_mem 10 GB
```

```
negative_ttl 15 seconds
```

```
connect_timeout 15 seconds
```

```
forward_timeout 10 seconds
```

```
read_timeout 15 seconds
```

```
pconn_timeout 70 seconds
```

```
persistent_request_timeout 65 seconds
```

Details - Configuration

ATS - records.config

```
CONFIG proxy.config.http.server_ports STRING 3128
```

```
# ATS requires disk cache to be set up through storage.config as well
```

```
CONFIG proxy.config.cache.ram_cache.size INT 2147483648
```

```
CONFIG proxy.config.http.negative_caching_enabled INT 1
```

```
CONFIG proxy.config.http.negative_caching_lifetime INT 15
```

Details - Configuration

ATS - records.config (cont)

```
CONFIG proxy.config.http.connect_attempts_timeout INT 15
CONFIG proxy.config.http.keep_alive_no_activity_timeout_in INT 15
CONFIG proxy.config.http.keep_alive_no_activity_timeout_out INT 30
CONFIG proxy.config.http.transaction_no_activity_timeout_in INT 30
CONFIG proxy.config.http.transaction_no_activity_timeout_out INT 30
CONFIG proxy.config.http.transaction_active_timeout_in INT 30
CONFIG proxy.config.http.transaction_active_timeout_out INT 30
CONFIG proxy.config.http.accept_no_activity_timeout INT 12
```

Details - Log

Squid - squid.conf

```
logfile_rotate 240
logformat ysquid_extended %ts.%03tu %6tr %>a %Ss/%03Hs %<st %rm %ru %un
%Sh/%<A %mt
access_log daemon:/usr/local/var/logs/squid/access.log ysquid_extended
cache_log /usr/local/var/logs/squid/cache.log

debug_options ALL,1
```

Details - Log

ATS - records.config

```
CONFIG proxy.config.log.logging_enabled INT 3
CONFIG proxy.config.log.logfile_dir STRING logs/trafficserver
CONFIG proxy.config.log.custom_logs_enabled INT 1
CONFIG proxy.config.log.rolling_enabled INT 1
CONFIG proxy.config.log.rolling_interval_sec INT 3600
CONFIG proxy.config.log.rolling_offset_hr INT 0
CONFIG proxy.config.log.rolling_size_mb INT 6000
CONFIG proxy.config.log.auto_delete_rolled_files INT 1

CONFIG proxy.config.diags.debug.enabled INT 0
CONFIG proxy.config.diags.debug.tags STRING http.*|dns.*
```

Details - Log

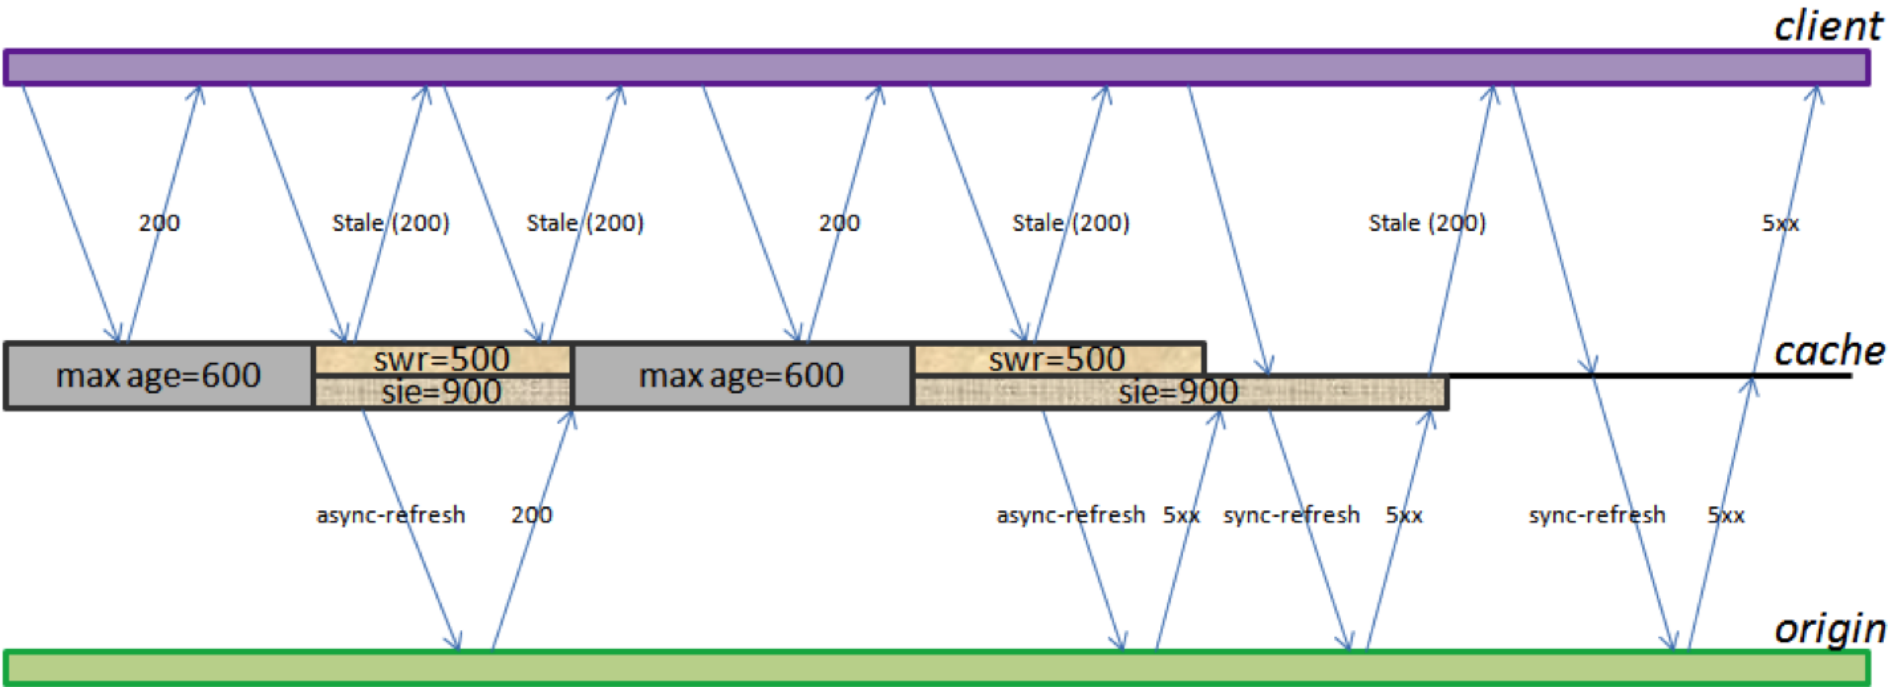
ATS - logs_xml.config

```
<LogFormat>
  <Name = "ats_generic_config"/>
  <Format = "ts=%<cqtq> url=%<cqu> host=%<{Host}cqhq duration=%<ttms>
status=%<pssc> cache=%<crc> ostatus=%<sssc> uurl=%<cquuc> conn=%
<cfsc>/%<pfsc> ip=%<chi> cqhm=%<cqhm> pscl=%<pscl> age=%<{Age}ssh> "/>
</LogFormat>
<LogObject>
  <Format = "ats_generic_config"/>
  <Filename = "mon"/>
  <Mode = "ascii"/>
</LogObject>
```


Details - Metrics

- e.g. cache hit ratio, avg latency, # of reqs
- Squid
 - squidclient mgr:info
 - squidclient mgr:counters
- ATS
 - traffic_line -r
 - https://docs.trafficserver.apache.org/en/latest/reference/commands/traffic_line.en.html

Details - SWR/SIE (RFC 5861)



Details - Collapsed Forwarding

- Multiple client requests as one server request
- Prevent “Thundering Herds” problem

Details - SWR/SIE/Collapsed Forwarding

Squid - squid.conf

```
collapse_forwarding on
```

```
max_stale 8 hours
```

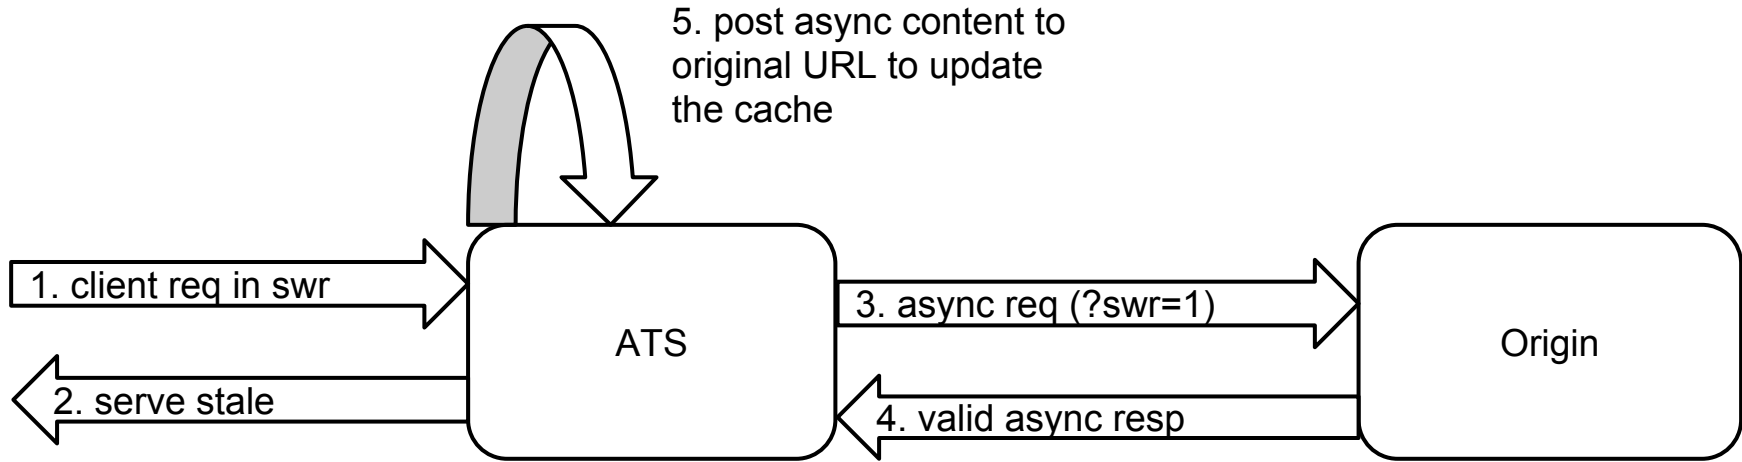
```
refresh_pattern . 5 0 480 stale-while-revalidate=15 ignore-reload override-lastmod
```

Details - SWR/SIE/Collapsed Forwarding

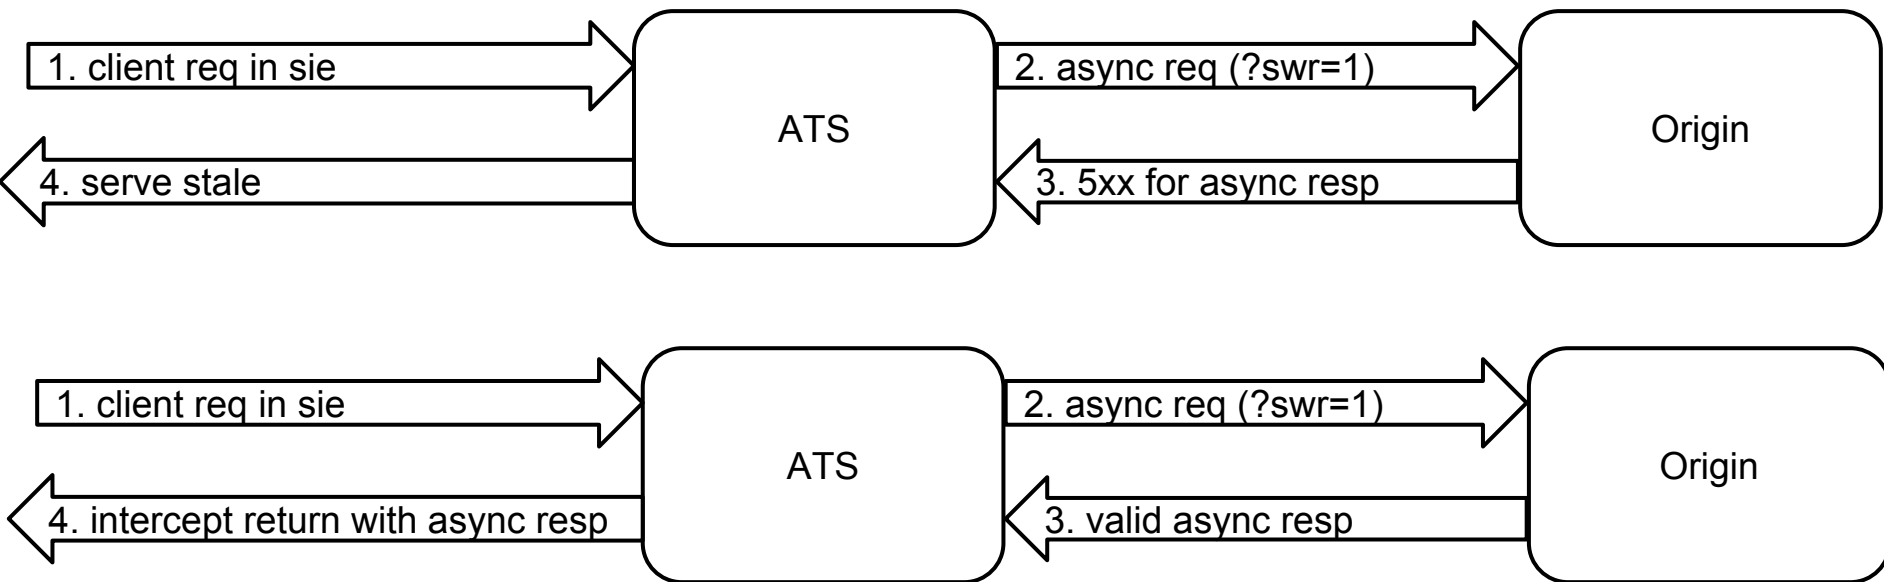
ATS

- Collapse Forwarding - check out [wiki page](#)
- experimental stale_while_revalidate plugin
- TS-1463
 - once the fetch of the object is initiated, the original object is not allowed to be served from cache
- TS-1996
 - Deprecated API - TSHttpTxnNewCacheLookupDo()

Details - Fixing the Plugin (SWR)



Details - Fixing the Plugin (SIE)



Details - ACL control (Squid)

Based on src, dst, time, regex, schema etc.

Squid.conf

```
# Example 1 (give access only to certain clients)
```

```
acl myclients src 172.16.5.0/24
```

```
http_access allow myclients
```

```
http_access deny all
```

```
# Example 2 (disable cache for responses from a domain)
```

```
acl someserver dstdomain .someserver.com
```

```
cache deny someserver
```


Details - ACL control (ATS)

remap.config

```
map http://www.x.com/ http://server.hoster.com/
```

ip_allow.config (Similar to Example 1)

```
src_ip=123.12.3.000-123.12.3.123 action=ip_allow
```

cache.config (Similar to Example 2)

```
dest_domain=mydomain.com action=never-cache
```

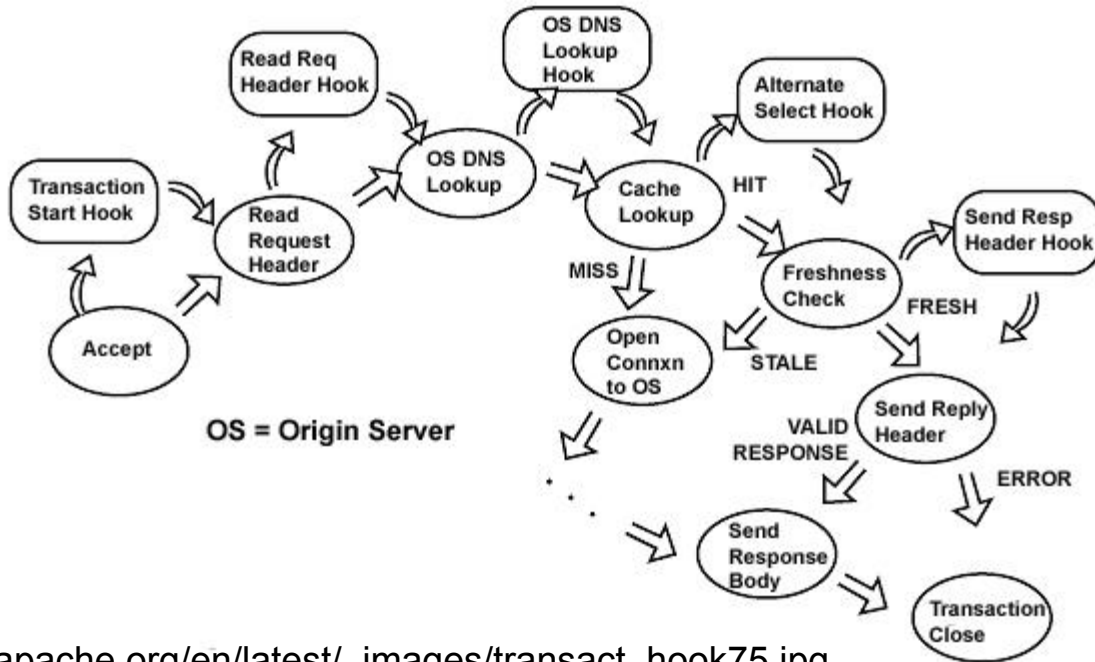
Details - Extensibility (Squid)

- Helper program for ACL, URL Manipulation, DNS lookup
 - Rigid and limited
 - e.g. controlling ACL in squid.conf

```
external_acl_type yca_helper cache=5000 concurrency=1000 children=2 grace=1
%SRC %{App-Auth} %DATA /usr/local/libexec/squid/yc_acl.pl
acl yca external yca_helper
acl yca_appids ext_user REQUIRED
deny_info YCA_AUTH_REQ yca
deny_info YCA_WRONG_APPID yca_appids
```

Details - Extensibility (ATS)

ATS - Plugins (C, C++, Lua)



Details - Peering (Squid)

Type

Parent Selection method

1. sibling (ICP*)
2. parent

1. carp*, round-robin, user-hashing etc.
2. originserver (reverse proxy)

ICP: Internet Cache Protocol
Carp: Cache Array Protocol

Details - Peering (Squid)

Squid.conf

```
# Example 1 (ICP)
```

```
cache_peer 172.16.1.123 sibling 3129 5500 weight=1
```

```
# Example 2 (Reverse Proxy mode)
```

```
cache_peer localhost parent 8080 0 originserver no-query no-digest
```

```
# Example 3 (Carp routing)
```

```
cache_peer server1.com parent 3128 0 carp no-query no-digest monitortimeout=15  
monitorinterval=60 monitorurl=/status.html name=andy
```

```
cache_peer server2.com parent 3128 0 carp no-query no-digest monitortimeout=15  
monitorinterval=60 monitorurl=/status.html name=mandy
```

Details - Peering (ATS)

Type

1. Sibling (ICP*)
2. Parent

* Our Squid ICP peering use case (Example 1) is not needed when we deploy with hierarchical caching + consistent hashing

Details - Peering (ATS)

remap.config (Similar to Squid Example 2)

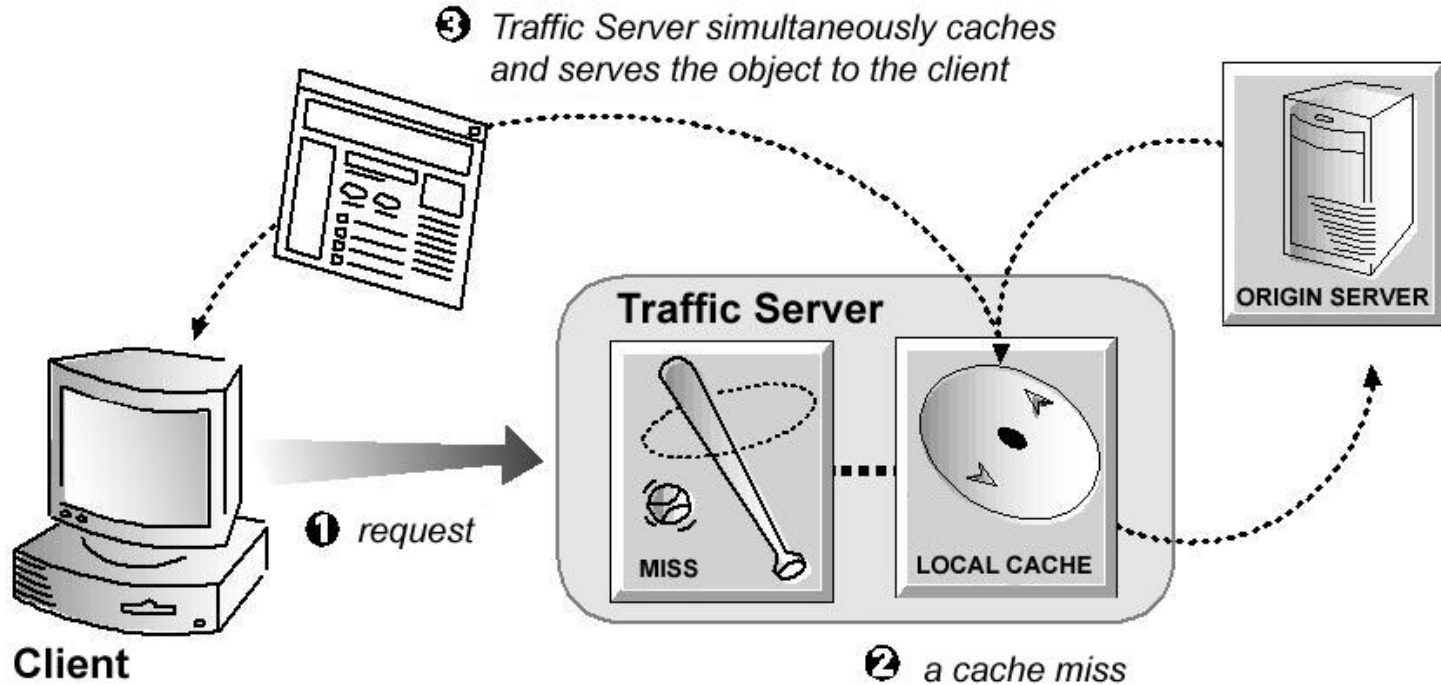
```
map http://www.x.com/ http://server1.com/
```

parent.config (Similar to Squid Example 3)

```
dest_domain=. method=get parent="p1.x.com:8080; p2.y.com:8080"  
round_robin=true
```

```
dest_domain=. method=get parent="p1.x.com:8080|1.0; p2.y.com:8080|2.0"  
round_robin=consistent_hash
```

Use Cases - Forward Proxy



Use Cases - Forward Proxy (Squid)

squid.conf

```
http_port 80 vhost
```

```
# protecting proxy by only allowing clients to connect to port 80
```

```
acl Safe_ports port 80
```

```
http_access deny !Safe_ports
```

```
acl localnet src 10.0.0.0/8 # RFC1918 possible internal network
```

```
http_access allow localhost
```

```
http_access allow localnet
```

Use Cases - Forward Proxy (ATS)

records.config

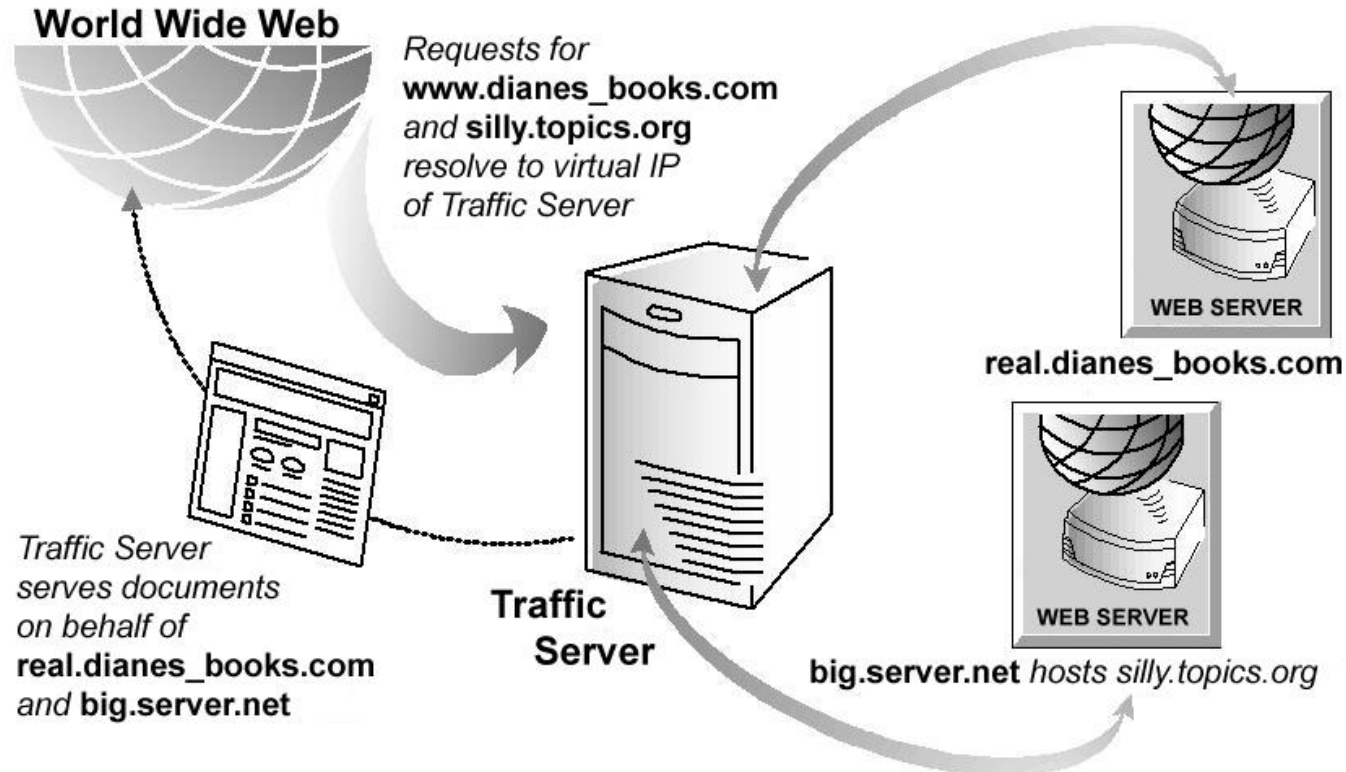
```
CONFIG proxy.config.reverse_proxy.enabled INT 0
CONFIG proxy.config.http.server_ports 80

# for security purpose require remap (Optional)
CONFIG proxy.config.url_remap.remap_required INT 1
```

remap.config

```
# map all origin servers for which we need forward proxy
map http://server1.com/ http://server1.com/
```

Use Cases - Reverse Proxy



Use Cases - Reverse Proxy (Squid)

Squid.conf

```
http_port 80 accel defaultsite=server1.com

# allow access to origin server
cache_peer endpoint.com parent 80 0 no-query originserver name=myAccess

# access controls
acl our_sites dstdomain server1.com server2.com
http_access allow our_sites
cache_peer_access myAccess allow our_sites
cache_peer_access myAccel deny all
```

Use Cases - Reverse Proxy (ATS)

1. records.config

```
CONFIG proxy.config.http.server_ports 80  
CONFIG proxy.config.reverse_proxy.enabled INT 1  
CONFIG proxy.config.url_remap.remap_required INT 1
```

2. remap.config

```
map http://www.server1.com/ http://www.endpoint.com/  
map http://www.server2.com/ http://www.endpoint.com/
```

Learnings

- Managing an open source project
 - Backward compatibility
- Squid configuration
 - config can be confusing when they exist in same file
 - allow/deny pattern confusing
- ATS configuration
 - ATS configuration complicated/disjoint, but clean
 - ATS plugins more flexible
- Improved Performance
 - CPU/Latency/RPS
- Migration
 - (Automated) Testing is important

Automated Testing

- TSQA
 - integration/functional testing framework
 - Not just for ATS, can be used for other proxy server (e.g. squid)
 - To be used heavily during migration

Q&A

Reference

Choosing a Proxy Server -

https://www.slideshare.net/bryan_call/choosing-a-proxy-server-apachecon-2014

Squid-2.6 - <http://www.squid-cache.org/Versions/v2/2.6/squid-2.6.STABLE24-RELEASENOTES.html#s1>

Squid-2.7 - <http://www.squid-cache.org/Versions/v2/2.7/squid-2.7.STABLE9-RELEASENOTES.html#s1>

Squid-3.0 - <http://www.squid-cache.org/Versions/v3/3.0/RELEASENOTES.html#s4>

Squid-3.2 - <http://www.squid-cache.org/Versions/v3/3.2/RELEASENOTES.html#s2>

Squid-3.5 - <http://www.squid-cache.org/Versions/v3/3.5/squid-3.5.3-RELEASENOTES.html#s2>

SIE Squid Bug - http://bugs.squid-cache.org/show_bug.cgi?id=2255

SWR Squid Bug - http://bugs.squid-cache.org/show_bug.cgi?id=2256

Reference (Cont)

ATS records.config - <https://docs.trafficserver.apache.org/en/latest/reference/configuration/records.config.en.html>

ATS logs_xml.config - https://docs.trafficserver.apache.org/en/latest/reference/configuration/logs_xml.config.en.html

ATS traffic_line - https://docs.trafficserver.apache.org/en/latest/reference/commands/traffic_line.en.html

ATS Collapsed Forwarding wiki - <https://cwiki.apache.org/confluence/display/TS/Collapsed+Forwarding>

ATS stale_while_revalidate plugin - https://docs.trafficserver.apache.org/en/latest/reference/plugins/stale_while_revalidate.en.html

ATS Plugin developer guide - <https://docs.trafficserver.apache.org/en/latest/sdk/index.en.html>