



<3

- Artur Bergman
- artur@crucially.net
- perl hacker
- varnish hacker
- Operations & Engineering at Wikia

Wikia

- Hosts wiki communities
- www.wowwiki.com (second largest)
- starwars.wikia.com (wookieepedia)
- uncyclopedia.wikia.com

- 30 000 wikis or so

wikia ≠ wikipedia

- Same software stack
- Apart from Varnish

Mediawiki

- PHP
- Mysql
- Memcache
- Varnish
- RabbitMQ

varnish!

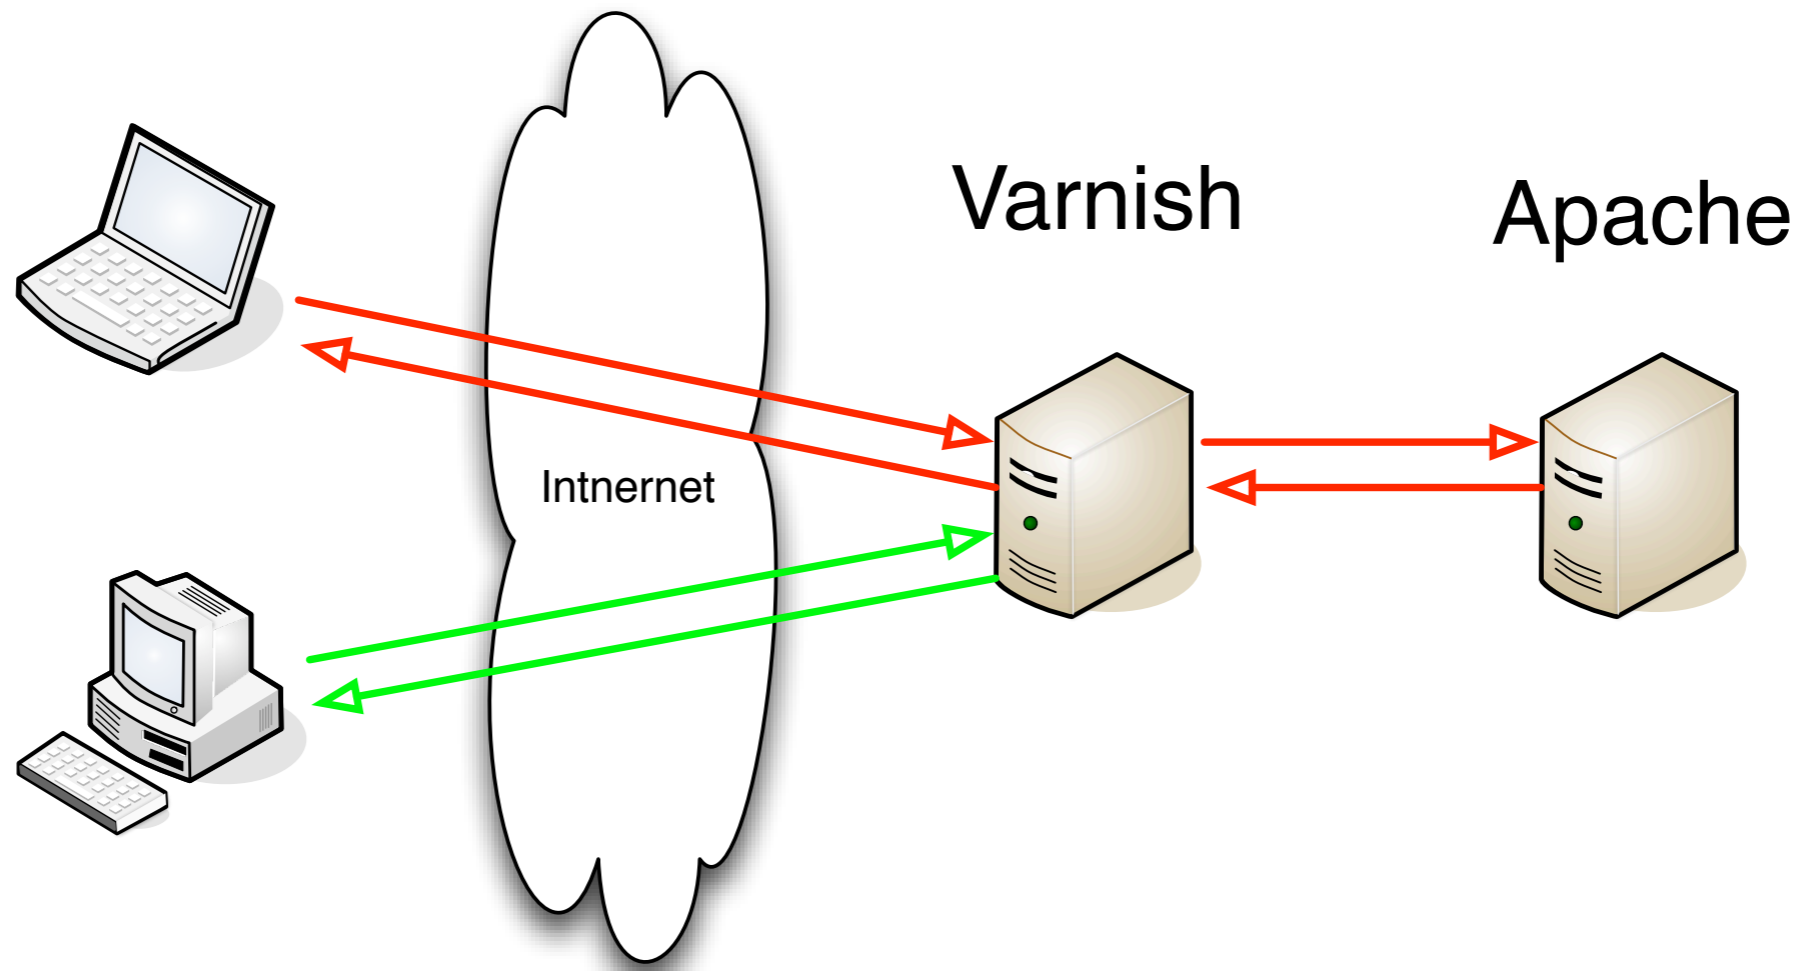
- HTTP accelerator
 - reverse proxy
 - cache
- very fast
 - very much faster than squid
- open source

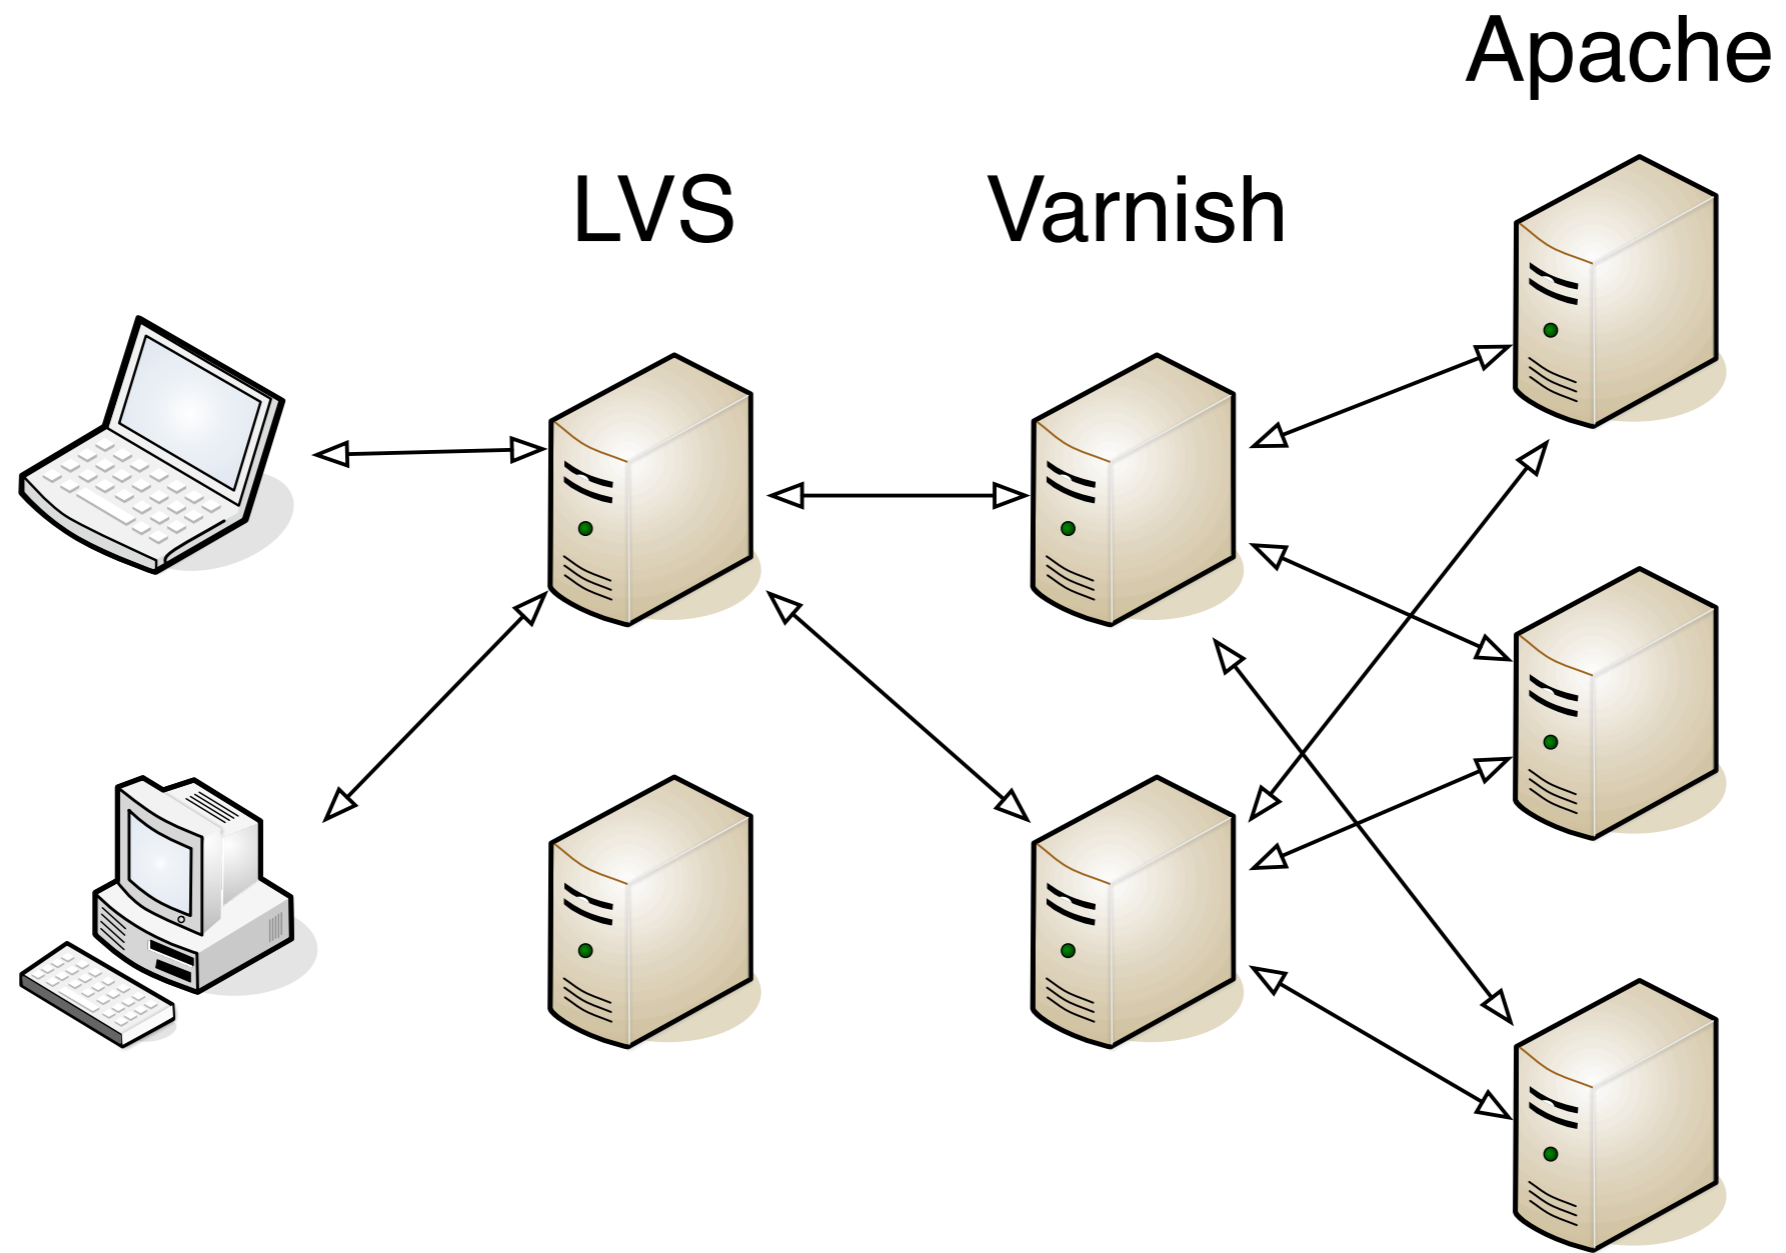
varnish!

- what is it
- architecture
- configuration options
- VCL
- stats/log
- performance numbers

reverse proxy / http accelerator

- squid replacement
- no forward caching
- sits between end user and backend servers
- caches content
- directs requests





Apache

LVS

Varnish

architect(ure)

- Poul-Henning Kamp (phk)
- FreeBSD
- <http://varnish.projects.linpro.no/wiki/ArchitectNotes>

architecture

- mmap
- threads
 - massive amounts
- event driven
- compiled configuration
- 7 syscalls (last count?)

mmap

- maps the store into memory
- (or alternatively use jemalloc)
 - (makes kswapd use 100% cpu -- bad linux)
- `madvise(MADV_RANDOM)`
- `writew` directly from mapped memory

store

- not persistent between restarts
 - (support in trunk)
- works really well with SSDs
- page in to evict :(

logfile

- all headers
- no syscall!
- ring buffer
 - just follow along

workspaces

- allocated up front
- no malloc in normal execution path (except ESI)
- aborts request on overflow
 - (or panic)
- request + response headers + metadata

configuration

/etc/default/varnish

Maximum number of open files (for ulimit -n)

NFILES=131072

Locked shared memory (for ulimit -l)

Default log size is 82MB + header

MEMLOCK=90000

configuration

/etc/default/varnish

Maximum number of open files (for ulimit -n)

NFILES=131072

Locked shared memory (for ulimit -l)

Default log size is 82MB + header

MEMLOCK=90000

```
DAEMON_OPTS="-a :80 \  
-T localhost:6082 \  
-s file,/var/lib/varnish,140GB \  
-f /etc/varnish/default.vcl \  
-u varnish \  
-g users \  
-p obj_workspace=4096 \  
-p sess_workspace=131072 \  
-p listen_depth=8192 \  
-p ping_interval=2 \  
-p log_hashstring=off \  
-h classic,250007 \  
-p thread_pool_max=8000 \  
-p lru_interval=60 \  
-p esi_syntax=0x00000003 \  
-p sess_timeout=10 \  
-p thread_pools=8 \  
-p thread_pool_min=500 \  
-p shm_workspace=32768 \  
-p srcaddr_ttl=0 \  
-p thread_pool_add_delay=1"
```

```
DAEMON_OPTS="-a :80 \  
-T localhost:6082 \  
-s file,/var/lib/varnish,140GB \  
-f /etc/varnish/default.vcl \  
-u varnish \  
-g users \  
-p obj_workspace=4096 \  
-p sess_workspace=131072 \  
-p listen_depth=8192 \  
-p ping_interval=2 \  
-p log_hashstring=off \  
-h classic,250007 \  
-p thread_pool_max=8000 \  
-p lru_interval=60 \  
-p esi_syntax=0x00000003 \  
-p sess_timeout=10 \  
-p thread_pools=8 \  
-p thread_pool_min=500 \  
-p shm_workspace=32768 \  
-p srcaddr_ttl=0 \  
-p thread_pool_add_delay=1"
```

storage type

- **-s file, /var/lib/varnish, 140GB**
 - mmap file
- **-s malloc, 140GB**
 - malloc
 - linux kswapd 100%
 - preferable for in memory workloads

```
DAEMON_OPTS="-a :80 \  
-T localhost:6082 \  
-s file,/var/lib/varnish,140GB \  
-f /etc/varnish/default.vcl \  
-u varnish \  
-g users \  
-p obj_workspace=4096 \  
-p sess_workspace=131072 \  
-p shm_workspace=32768 \  
-p listen_depth=8192 \  
-p ping_interval=2 \  
-p log_hashstring=off \  
-h classic,250007 \  
-p thread_pool_max=8000 \  
-p lru_interval=60 \  
-p esi_syntax=0x00000003 \  
-p sess_timeout=10 \  
-p thread_pools=8 \  
-p thread_pool_min=500 \  
-p srcaddr_ttl=0 \  
-p thread_pool_add_delay=1"
```

obj_workspace=4096

- overhead per object
 - scales with number of objects
- keep small
- can overflow from backend
 - very large cookies
 - very big headers

sess_workspace=131072

- overhead per thread
- scratchpad for VCL work
- can overflow
- only panics if you do excessive copying in VCL and then run out of space to compute the hash
- probably want something smaller than 128k

shm_workspace=32768

- overhead per thread
- temporary storage for logs
- tune to decrease shmlog mutex
- `varnishstat -f shm_cont`

```
DAEMON_OPTS="-a :80 \  
-T localhost:6082 \  
-s file,/var/lib/varnish,140GB \  
-f /etc/varnish/default.vcl \  
-u varnish \  
-g users \  
-p obj_workspace=4096 \  
-p sess_workspace=131072 \  
-p shm_workspace=32768 \  
-p listen_depth=8192 \  
-p ping_interval=2 \  
-p log_hashstring=off \  
-h classic,250007 \  
-p thread_pool_max=8000 \  
-p lru_interval=60 \  
-p esi_syntax=0x00000003 \  
-p sess_timeout=10 \  
-p thread_pools=8 \  
-p thread_pool_min=500 \  
-p srcaddr_ttl=0 \  
-p thread_pool_add_delay=1"
```

Argument to listen()

```
DAEMON_OPTS="-a :80 \  
-T localhost:6082 \  
-s file,/var/lib/varnish,140GB \  
-f /etc/varnish/default.vcl \  
-u varnish \  
-g users \  
-p obj_workspace=4096 \  
-p sess_workspace=131072 \  
-p shm_workspace=32768 \  
-p listen_depth=8192 \  
-p ping_interval=2 \  
-p log_hashstring=off \  
-h classic,250007 \  
-p thread_pool_max=8000 \  
-p lru_interval=60 \  
-p esi_syntax=0x00000003 \  
-p sess_timeout=10 \  
-p thread_pools=8 \  
-p thread_pool_min=500 \  
-p srcaddr_ttl=0 \  
-p thread_pool_add_delay=1"
```

varnish
restarts if no response

```
DAEMON_OPTS="-a :80 \  
-T localhost:6082 \  
-s file,/var/lib/varnish,140GB \  
-f /etc/varnish/default.vcl \  
-u varnish \  
-g users \  
-p obj_workspace=4096 \  
-p sess_workspace=131072 \  
-p shm_workspace=32768 \  
-p listen_depth=8192 \  
-p ping_interval=2 \  
-p log_hashstring=off \  
-h classic,250007 \  
-p thread_pool_max=8000 \  
-p lru_interval=60 \  
-p esi_syntax=0x00000003 \  
-p sess_timeout=10 \  
-p thread_pools=8 \  
-p thread_pool_min=500 \  
-p srcaddr_ttl=0 \  
-p thread_pool_add_delay=1"
```

number of buckets
1/10th of objects

```
DAEMON_OPTS="-a :80 \  
-T localhost:6082 \  
-s file,/var/lib/varnish,140GB \  
-f /etc/varnish/default.vcl \  
-u varnish \  
-g users \  
-p obj_workspace=4096 \  
-p sess_workspace=131072 \  
-p shm_workspace=32768 \  
-p listen_depth=8192 \  
-p ping_interval=2 \  
-p log_hashstring=off \  
-h classic,250007 \  
-p thread_pool_max=8000 \  
-p lru_interval=60 \  
-p esi_syntax=0x00000003 \  
-p sess_timeout=10 \  
-p thread_pools=8 \  
-p thread_pool_min=500 \  
-p srcaddr_ttl=0 \  
-p thread_pool_add_delay=1"
```

total number of max
threads

careful to not let
threads run high in io
pressure situations

```
DAEMON_OPTS="-a :80 \  
-T localhost:6082 \  
-s file,/var/lib/varnish,140GB \  
-f /etc/varnish/default.vcl \  
-u varnish \  
-g users \  
-p obj_workspace=4096 \  
-p sess_workspace=131072 \  
-p shm_workspace=32768 \  
-p listen_depth=8192 \  
-p ping_interval=2 \  
-p log_hashstring=off \  
-h classic,250007 \  
-p thread_pool_max=8000 \  
-p lru_interval=60 \  
-p esi_syntax=0x00000003 \  
-p sess_timeout=10 \  
-p thread_pools=8 \  
-p thread_pool_min=500 \  
-p srcaddr_ttl=0 \  
-p thread_pool_add_delay=1"
```

one threadpool per
CPU

we want to force create
threads on startup so
min of 500*8

```
DAEMON_OPTS="-a :80 \  
-T localhost:6082 \  
-s file,/var/lib/varnish,140GB \  
-f /etc/varnish/default.vcl \  
-u varnish \  
-g users \  
-p obj_workspace=4096 \  
-p sess_workspace=131072 \  
-p shm_workspace=32768 \  
-p listen_depth=8192 \  
-p ping_interval=2 \  
-p log_hashstring=off \  
-h classic,250007 \  
-p thread_pool_max=8000 \  
-p lru_interval=60 \  
-p esi_syntax=0x00000003 \  
-p sess_timeout=10 \  
-p thread_pools=8 \  
-p thread_pool_min=500 \  
-p srcaddr_ttl=0 \  
-p thread_pool_add_delay=1"
```

disable srcaddr_ttl!

slows things down with
no benefit

removed from trunk

```
DAEMON_OPTS="-a :80 \  
-T localhost:6082 \  
-s file,/var/lib/varnish,140GB \  
-f /etc/varnish/default.vcl \  
-u varnish \  
-g users \  
-p obj_workspace=4096 \  
-p sess_workspace=131072 \  
-p shm_workspace=32768 \  
-p listen_depth=8192 \  
-p ping_interval=2 \  
-p log_hashstring=off \  
-h classic,250007 \  
-p thread_pool_max=8000 \  
-p lru_interval=60 \  
-p esi_syntax=0x00000003 \  
-p sess_timeout=10 \  
-p thread_pools=8 \  
-p thread_pool_min=500 \  
-p srcaddr_ttl=0 \  
-p thread_pool_add_delay=1"
```

how long to wait
between threads

default is 20ms

far too long -- makes
startup cause failures


```
-p 'cc_command=exec cc -fpic -shared -Wl,-x -L/usr/local/  
lib/ -lGeoIP -o %o %s'
```

shmlog on tmpfs

- shmlog is mlocked()
- still written to disk
 - dirty buffers
- IO \neq requests per second

```
tmpfs          /var/lib/varnish/ tmpfs  
noatime,defaults,size=150M 0 0
```

vcl

- domain specific language
- translated into C
- compiled
- dynamically loaded and executed
- [https://svn.wikia-code.com/utils/
varnishhtcpd/wikia.vcl](https://svn.wikia-code.com/utils/varnishhtcpd/wikia.vcl)

vcl_recv

- first entry point
- results in
 - pipe
 - pass
 - lookup

```
sub vcl_recv {  
  
    # normalize Accept-Encoding to reduce vary  
    if (req.http.Accept-Encoding) {  
        if (req.http.User-Agent ~ "MSIE 6") {  
            unset req.http.Accept-Encoding;  
        } elseif (req.http.Accept-Encoding ~ "gzip") {  
            set req.http.Accept-Encoding = "gzip";  
        } elseif (req.http.Accept-Encoding ~ "deflate") {  
            set req.http.Accept-Encoding = "deflate";  
        } else {  
            unset req.http.Accept-Encoding;  
        }  
    }  
}
```

(I hate browsers)

```
sub vcl_recv {  
  
    # normalize Accept-Encoding to reduce vary  
    if (req.http.Accept-Encoding) {  
        if (req.http.User-Agent ~ "MSIE 6") {  
            unset req.http.Accept-Encoding;  
        } elseif (req.http.Accept-Encoding ~ "gzip") {  
            set req.http.Accept-Encoding = "gzip";  
        } elseif (req.http.Accept-Encoding ~ "deflate") {  
            set req.http.Accept-Encoding = "deflate";  
        } else {  
            unset req.http.Accept-Encoding;  
        }  
    }  
}
```

(I hate browsers)

```
# clean out requests sent via curls -X mode and LWP
if (req.url ~ "^http://") {
    set req.url = regsub(req.url, "http://[^/]*", "");
}

# lvs check
if (req.url == "/svccheck.html") {
    error 200 "OK";
}

if (req.url == "/__ervername") {
    error 200 "OK";
}
```

```
# save the cookie for later use
set req.http.X-Orig-Cookie = req.http.Cookie;
if(req.http.Cookie ~ "(session|UserID|UserName|Token|LoggedOut)") {
    # dont do anything, the user is logged in
} else {
# dont care about any other cookies
# for vary purposes
    unset req.http.Cookie;
}
```



```
# pipe post
if (req.request != "GET" && req.request != "HEAD" && req.request !=
"PURGE") {
    pipe;
}

# dont cache Authenticate calls
# we dont use those?
if (req.http.Authenticate) {
    pass;
}
set req.grace = 3600s;
lookup;
}
```

vcl_pipe

- pipe switches to byte transfer mode
- no further work is done on the connection

```
sub vcl_pipe {
# do the right XFF processing
# we chain XFF correctly
  set bereq.http.X-Forwarded-For = req.http.X-Forwarded-For;
  set bereq.http.X-Forwarded-For = regsub(bereq.http.X-Forwarded-For, "$", ",");
  set bereq.http.X-Forwarded-For = regsub(bereq.http.X-Forwarded-For, "$", client.ip);
# restore cookie
  set bereq.http.Cookie = req.http.X-Orig-Cookie;
# we don't want any more requests on this connection
# or XFF won't work
  set bereq.http.connection = "close";
}
```

varnish default XFF support is broken

vcl_hit

- called on hit
- be careful
 - **DO NOT MODIFY THE OBJECT**
 - (except TTL)
- way to implement purging
 - (remember to purge all Vary versions)

```
sub vcl_hit {  
    if (req.request == "PURGE") {  
        set obj.ttl = 0s;  
        error 200 "Purged.";  
    }  
}
```

vcl_miss

- called on a miss
- just before fetch from backend
- can change the bereq object

```
sub vcl_miss {  
# tell the client the purge failed  
if (req.request == "PURGE") {  
    error 404 "Not purged";  
}  
  
set bereq.http.X-Forwarded-For = req.http.X-Forwarded-For;  
set bereq.http.X-Forwarded-For = regsub(bereq.http.X-Forwarded-For, "$", ",");  
set bereq.http.X-Forwarded-For = regsub(bereq.http.X-Forwarded-For, "$", client.ip);  
  
# reset the cookie to what it was originally  
set bereq.http.Cookie = req.http.X-Orig-Cookie;  
  
}
```

vcl_fetch

- just after an object has been fetched
- request object
- cached object


```
sub vcl_fetch {
    if(req.url == "/robots.txt") {
        set obj.http.X-Pass-Cache-Control = "max-age=86400";
        set obj.ttl = 86400s;
    }
    if (!obj.cacheable) {
        set obj.http.X-Cacheable = "NO:Not-Cacheable";
        pass;
    }
    if (obj.http.Cache-Control ~ "private") {
        if(req.http.Cookie ~ "(UserID|_session)") {
            set obj.http.X-Cacheable = "NO:Got Session";
        } else {
            set obj.http.X-Cacheable = "NO:Cache-Control=private";
        }
        pass;
    }
    if (obj.http.Set-Cookie ~ "(UserID|_session)") {
        set obj.http.X-Cacheable = "NO:Set-Cookie";
        pass;
    }
}
```

```
sub vcl_fetch {
    if(req.url == "/robots.txt") {
        set obj.http.X-Pass-Cache-Control = "max-age=86400";
        set obj.ttl = 86400s;
    }
    if (!obj.cacheable) {
        set obj.http.X-Cacheable = "NO:Not-Cacheable";
        pass;
    }
    if (obj.http.Cache-Control ~ "private") {
        if(req.http.Cookie ~ "(UserID|_session)") {
            set obj.http.X-Cacheable = "NO:Got Session";
        } else {
            set obj.http.X-Cacheable = "NO:Cache-Control=private";
        }
        pass;
    }
    if (obj.http.Set-Cookie ~ "(UserID|_session)") {
        set obj.http.X-Cacheable = "NO:Set-Cookie";
        pass;
    }
}
```

```
sub vcl_fetch {
    if(req.url == "/robots.txt") {
        set obj.http.X-Pass-Cache-Control = "max-age=86400";
        set obj.ttl = 86400s;
    }
    if (!obj.cacheable) {
        set obj.http.X-Cacheable = "NO:Not-Cacheable";
        pass;
    }
    if (obj.http.Cache-Control ~ "private") {
        if(req.http.Cookie ~ "(UserID|_session)") {
            set obj.http.X-Cacheable = "NO:Got Session";
        } else {
            set obj.http.X-Cacheable = "NO:Cache-Control=private";
        }
        pass;
    }
    if (obj.http.Set-Cookie ~ "(UserID|_session)") {
        set obj.http.X-Cacheable = "NO:Set-Cookie";
        pass;
    }
}
```

```
if ( obj.http.X-Pass-Cache-Control ) {
    set obj.http.X-Internal-Pass-Cache-Control = obj.http.X-Pass-Cache-Control;
} elsif ( obj.status == 304 ) {
# no headers on if-modified since

} elsif ( req.url ~ ".*\/index\.php.*(css|js)"
|| req.url ~ "raw" ) {
# dont touch it let mediawiki decide

} elsif ( req.http.Host ~ "images.wikia.com" ) {
# lighttpd knows what it is doing

} else {
    set obj.http.X-Internal-Pass-Cache-Control = "private, s-maxage=0, max-age=0,
must-revalidate";
}
```

Seperate from cache-control since external cache-control
is not what we want varnish to follow

```
if (obj.ttl < 1s) {  
    set obj.ttl = 5s;  
    set obj.grace = 5s;  
    set obj.http.X-Cacheable = "YES - FORCED";  
    deliver;  
} else {  
    set obj.http.X-Cacheable = "YES";  
  
    if (obj.ttl < 600s) {  
        set obj.grace = 5s;  
    } else {  
        set obj.grace = 3600s;  
    }  
}
```

grace

- Serve stale object
- Fetch new object from background
- If backend dead serve stale
- Avoids thread pileups on invalidations

URL coalescing

- Multiple front end requests
- One backend request
- Unlike Squid
- Wikipedia suffered when Michael Jackson died because of cache storms

```
if(obj.status == 404) {  
    set obj.http.Cache-Control = "max-age=10";  
    set obj.ttl = 10s;  
    set obj.grace = 10s;  
}  
deliver;  
}
```


vcl_deliver

- modify response object
- don't modify the cached object
- no access to the request object
- (changed in next major version)

```
#add or append Served By
if(!resp.http.X-Served-By) {
  set resp.http.X-Served-By = server.identity;
  if (obj.hits > 0) {
    set resp.http.X-Cache = "HIT";
  } else {
    set resp.http.X-Cache = "MISS";
  }
  set resp.http.X-Cache-Hits = obj.hits;
} else {
# append current data
  set resp.http.X-Served-By = regsub(resp.http.X-Served-By, "$", ", ");
  set resp.http.X-Served-By = regsub(resp.http.X-Served-By, "$", server.identity);
  if (obj.hits > 0) {
    set resp.http.X-Cache = regsub(resp.http.X-Cache, "$", ", HIT");
  } else {
    set resp.http.X-Cache = regsub(resp.http.X-Cache, "$", ", MISS");
  }
  set resp.http.X-Cache-Hits = regsub(resp.http.X-Cache-Hits, "$", ", ");
  set resp.http.X-Cache-Hits = regsub(resp.http.X-Cache-Hits, "$", obj.hits);
}
```

```
252:~ sky$ curl -I http://www.wowwiki.com/Portal:Main -x varnish8.wikia.net:80
HTTP/1.1 200 OK
Server: Apache
Content-language: en
Vary: Accept-Encoding, Cookie
Last-Modified: Sun, 19 Jul 2009 05:35:33 GMT
Content-Type: text/html; charset=utf-8
Content-Length: 64672
X-Cacheable: YES
Date: Thu, 23 Jul 2009 07:12:31 GMT
Connection: keep-alive
X-Served-By: varnish1, r9-8-23, varnish8
X-Cache: HIT, HIT, HIT
X-Cache-Hits: 3, 979, 4877
X-Age: 27498
Cache-Control: private, s-maxage=0, max-age=0, must-revalidate
```

```
#add or append Served By
if(!resp.http.X-Served-By) {
  set resp.http.X-Served-By = server.identity;
  if (obj.hits > 0) {
    set resp.http.X-Cache = "HIT";
  } else {
    set resp.http.X-Cache = "MISS";
  }
  set resp.http.X-Cache-Hits = obj.hits;
} else {
# append current data
  set resp.http.X-Served-By = regsub(resp.http.X-Served-By, "$", ", ");
  set resp.http.X-Served-By = regsub(resp.http.X-Served-By, "$", server.identity);
  if (obj.hits > 0) {
    set resp.http.X-Cache = regsub(resp.http.X-Cache, "$", ", HIT");
  } else {
    set resp.http.X-Cache = regsub(resp.http.X-Cache, "$", ", MISS");
  }
  set resp.http.X-Cache-Hits = regsub(resp.http.X-Cache-Hits, "$", ", ");
  set resp.http.X-Cache-Hits = regsub(resp.http.X-Cache-Hits, "$", obj.hits);
}
```

```
#don't confused caches
```

```
set resp.http.X-Age = resp.http.Age;
```

```
#allow overrides of Cache-Control header
```

```
if (resp.http.X-Internal-Pass-Cache-Control) {
```

```
set resp.http.Cache-Control = resp.http.X-Internal-Pass-Cache-Control;
```

```
unset resp.http.X-Internal-Pass-Cache-Control;
```

```
}
```

```
unset resp.http.Age;
```

```
unset resp.http.X-Varnish;
```

```
unset resp.http.Via;
```

```
unset resp.http.X-Vary-Options;
```

```
unset resp.http.X-Powered-By;
```

```
deliver;
```

```
}
```

vcl_error

- used for synthetic responses
- errors or just generated

```
sub vcl_error {
  if (req.url ~ "/__servername") {
    synthetic server.identity;
    deliver;
  }
}
```

```
if(req.url ~ "svccheck.html") {
  synthetic {"varnish is okay"};
  deliver;
}
```

```
synthetic {"
<script src="http://www.google-analytics.com/urchin.js" type="text/javascript">
</script>
<script type="text/javascript">
try {
_uacct = "UA-xxxx-xxx";
urchinTracker("/varnish/" server.identity {"/" } obj.status {"");
} catch(err) {}</script>
"};
  deliver;}
```

```
sub vcl_error {  
    if (req.url ~ "/__servername") {  
        synthetic server.identity;  
        deliver;  
    }  
}
```

```
if(req.url ~ "svccheck.html") {  
    synthetic {"varnish is okay"};  
    deliver;  
}
```

```
synthetic {"  
<script src="http://www.google-analytics.com/urchin.js" type="text/javascript">  
</script>  
<script type="text/javascript">  
try {  
_uacct = "UA-xxxx-xxx";  
urchinTracker("/varnish/") server.identity {"/" } obj.status {""};  
} catch(err) {}</script>  
"};  
    deliver;}
```



```
sub vcl_error {
  if (req.url ~ "/__servername") {
    synthetic server.identity;
    deliver;
  }
}
```

```
if(req.url ~ "svccheck.html") {
  synthetic {"varnish is okay"};
  deliver;
}
```

```
synthetic {"
<script src="http://www.google-analytics.com/urchin.js" type="text/javascript">
</script>
<script type="text/javascript">
try {
_uacct = "UA-xxxx-xxx";
urchinTracker("/varnish/" server.identity {"/" } obj.status {"");
} catch(err) {}</script>
"};
  deliver;}
```

C code!

- Embed C code in the config
- Quite useful for
 - Cookie inspection
 - Generating Expire header
 - Geoip generator
- `varnishd -C -f` to see generated code

```
C{
#include <string.h>
double TIM_real(void);
void TIM_format(double t, char *p);
}C
```

```
C{
#include <dlfcn.h>
#include <stdlib.h>
#include <stdio.h>
#include <string.h>
#include <GeolPCity.h>
#include <pthread.h>
pthread_mutex_t geoip_mutex = PTHREAD_MUTEX_INITIALIZER;
GeolP* gi;
void geo_init () {
    if(!gi) {
        gi =
GeolP_open_type(GEOIP_CITY_EDITION_REV1,GEOIP_MEMORY_CACHE);
    }
}
}C
```

```

# if there isnt an expiry
if (!resp.status == 304) {
    C{
        char *cache = VRT_GetHdr(sp, HDR_REQ, "\016cache-control:");
        char date[40];
        int max_age;
        int want_equals = 0;
        if(cache) {
            while(*cache != '\0') {
                if (want_equals && *cache == '=') {
                    cache++;
                    max_age = strtoul(cache, 0, 0);
                    break;
                }

                if (*cache == 'm' && !memcmp(cache, "max-age", 7)) {
                    cache += 7;
                    want_equals = 1;
                    continue;
                }
                cache++;
            }
            if (max_age) {
                TIM_format(TIM_real() + max_age, date);
                VRT_SetHdr(sp, HDR_RESP, "\010Expires:", date, vrt_magic_string_end);
            }
        }
    }C
    #;
}

```

```

C{
char *ip = VRT_IP_string(sp,VRT_r_client_ip(sp));
char date[40];
char json[255];

pthread_mutex_lock(&geoip_mutex);

if(!gi) { geo_init(); }

GeoIPRecord *record = GeoIP_record_by_addr(gi, ip);
if(record) {
    snprintf(json, 255, "Geo = {\"city\": \"%s\", \"country\": \"%s\", \"lat\": \"%f\", \"lon\": \"%f\", \"classC\": \"%s\", \"netmask\": \"%d\"}",
        record->city,
        record->country_code,
        record->latitude,
        record->longitude,
        ip,
        GeoIP_last_netmask(gi)
    );
    pthread_mutex_unlock(&geoip_mutex);
    VRT_synth_page(sp, 0, json, vrt_magic_string_end);
} else {
    pthread_mutex_unlock(&geoip_mutex);
    VRT_synth_page(sp, 0, "Geo = {}", vrt_magic_string_end);
}

TIM_format(TIM_real(), date);
VRT_SetHdr(sp, HDR_OBJ, "\016Last-Modified:", date, vrt_magic_string_end);
}C
deliver;
}

```

```
Geo = {"city":"White Plains","country":"US","lat":"41.029099","lon":"-73.758003","classC":"209.133.114.31","netmask":"23"}
```

geoiplookup.wikia.com

varnishlog

```
4725 SessionOpen c xxx.xxx.xxx.xxx 1441 :80
4774 ReqEnd      - 0 1245712664.794090033 1245712664.794090033 0.003499746 0.000000000 0.000000000
4774 StatSess   - xxx.xxx.xxx.xxx 1442 0 1 0 0 0 0 0
4749 SessionOpen c xxx.xxx.xxx.xxx 2748 :80
10216 ReqStart  c xxx.xxx.xxx.xxx 51324 1570384079
10216 RxRequest  c GET
10216 RxURL      c /runescape/images/4/4c/Defence_cape.gif
10216 RxProtocol c HTTP/1.1
10216 RxHeader   c Accept: */*
10216 RxHeader   c Referer: http://runescape.wikia.com/wiki/Defence_cape
10216 RxHeader   c Accept-Language: en-gb
10216 RxHeader   c UA-CPU: x86
10216 RxHeader   c Accept-Encoding: gzip, deflate
10216 RxHeader   c User-Agent: Mozilla/4.0 (compatible; MSIE 7.0; Windows NT 6.0; FunWebProducts; GTB6; SLCC1; .NET CLR 2.0.50727; Media Center PC 5.0; .NET CLR 3.5.30729; .NET CLR 3.0.30618; OfficeLiveConnector.1.3; OfficeLivePatch.0.0)
10216 RxHeader   c Host: images3.wikia.nocookie.net
10216 RxHeader   c Connection: Keep-Alive
10216 VCL_call    c recv
10216 VCL_acl     c NO_MATCH SJC
10216 VCL_acl     c MATCH LON xxx.xxx.xxx.xxx
10216 VCL_return  c lookup
10216 VCL_call    c hash
10216 VCL_return  c hash
10216 Hit        c 1216457642
10216 VCL_call    c hit
10216 VCL_return  c deliver
10216 Length     c 1851
10216 VCL_call    c deliver
10216 VCL_acl     c NO_MATCH LON
10216 VCL_acl     c NO_MATCH SJC
10216 VCL_acl     c NO_MATCH IOWA
10216 VCL_return  c deliver
10216 TxProtocol  c HTTP/1.1
10216 TxStatus    c 200
10216 TxResponse  c OK
10216 TxHeader    c Cache-Control: max-age=30
10216 TxHeader    c Content-Type: image/gif
10216 TxHeader    c ETag: "209654623"
10216 TxHeader    c Last-Modified: Thu, 12 Mar 2009 04:58:56 GMT
10216 TxHeader    c Server: lighttpd/1.4.18
10216 TxHeader    c Content-Length: 1851
```

- `varnishlog -o -c RxURL part_of_url`
- `varnishlog -b -i TxURL | head -l 1000 |
cut -c 22- | sort | uniq -c | sort -rn |
head -20`
- `varnishlog -o -c ReqStart | 27.0.0.1`

varnishncsa

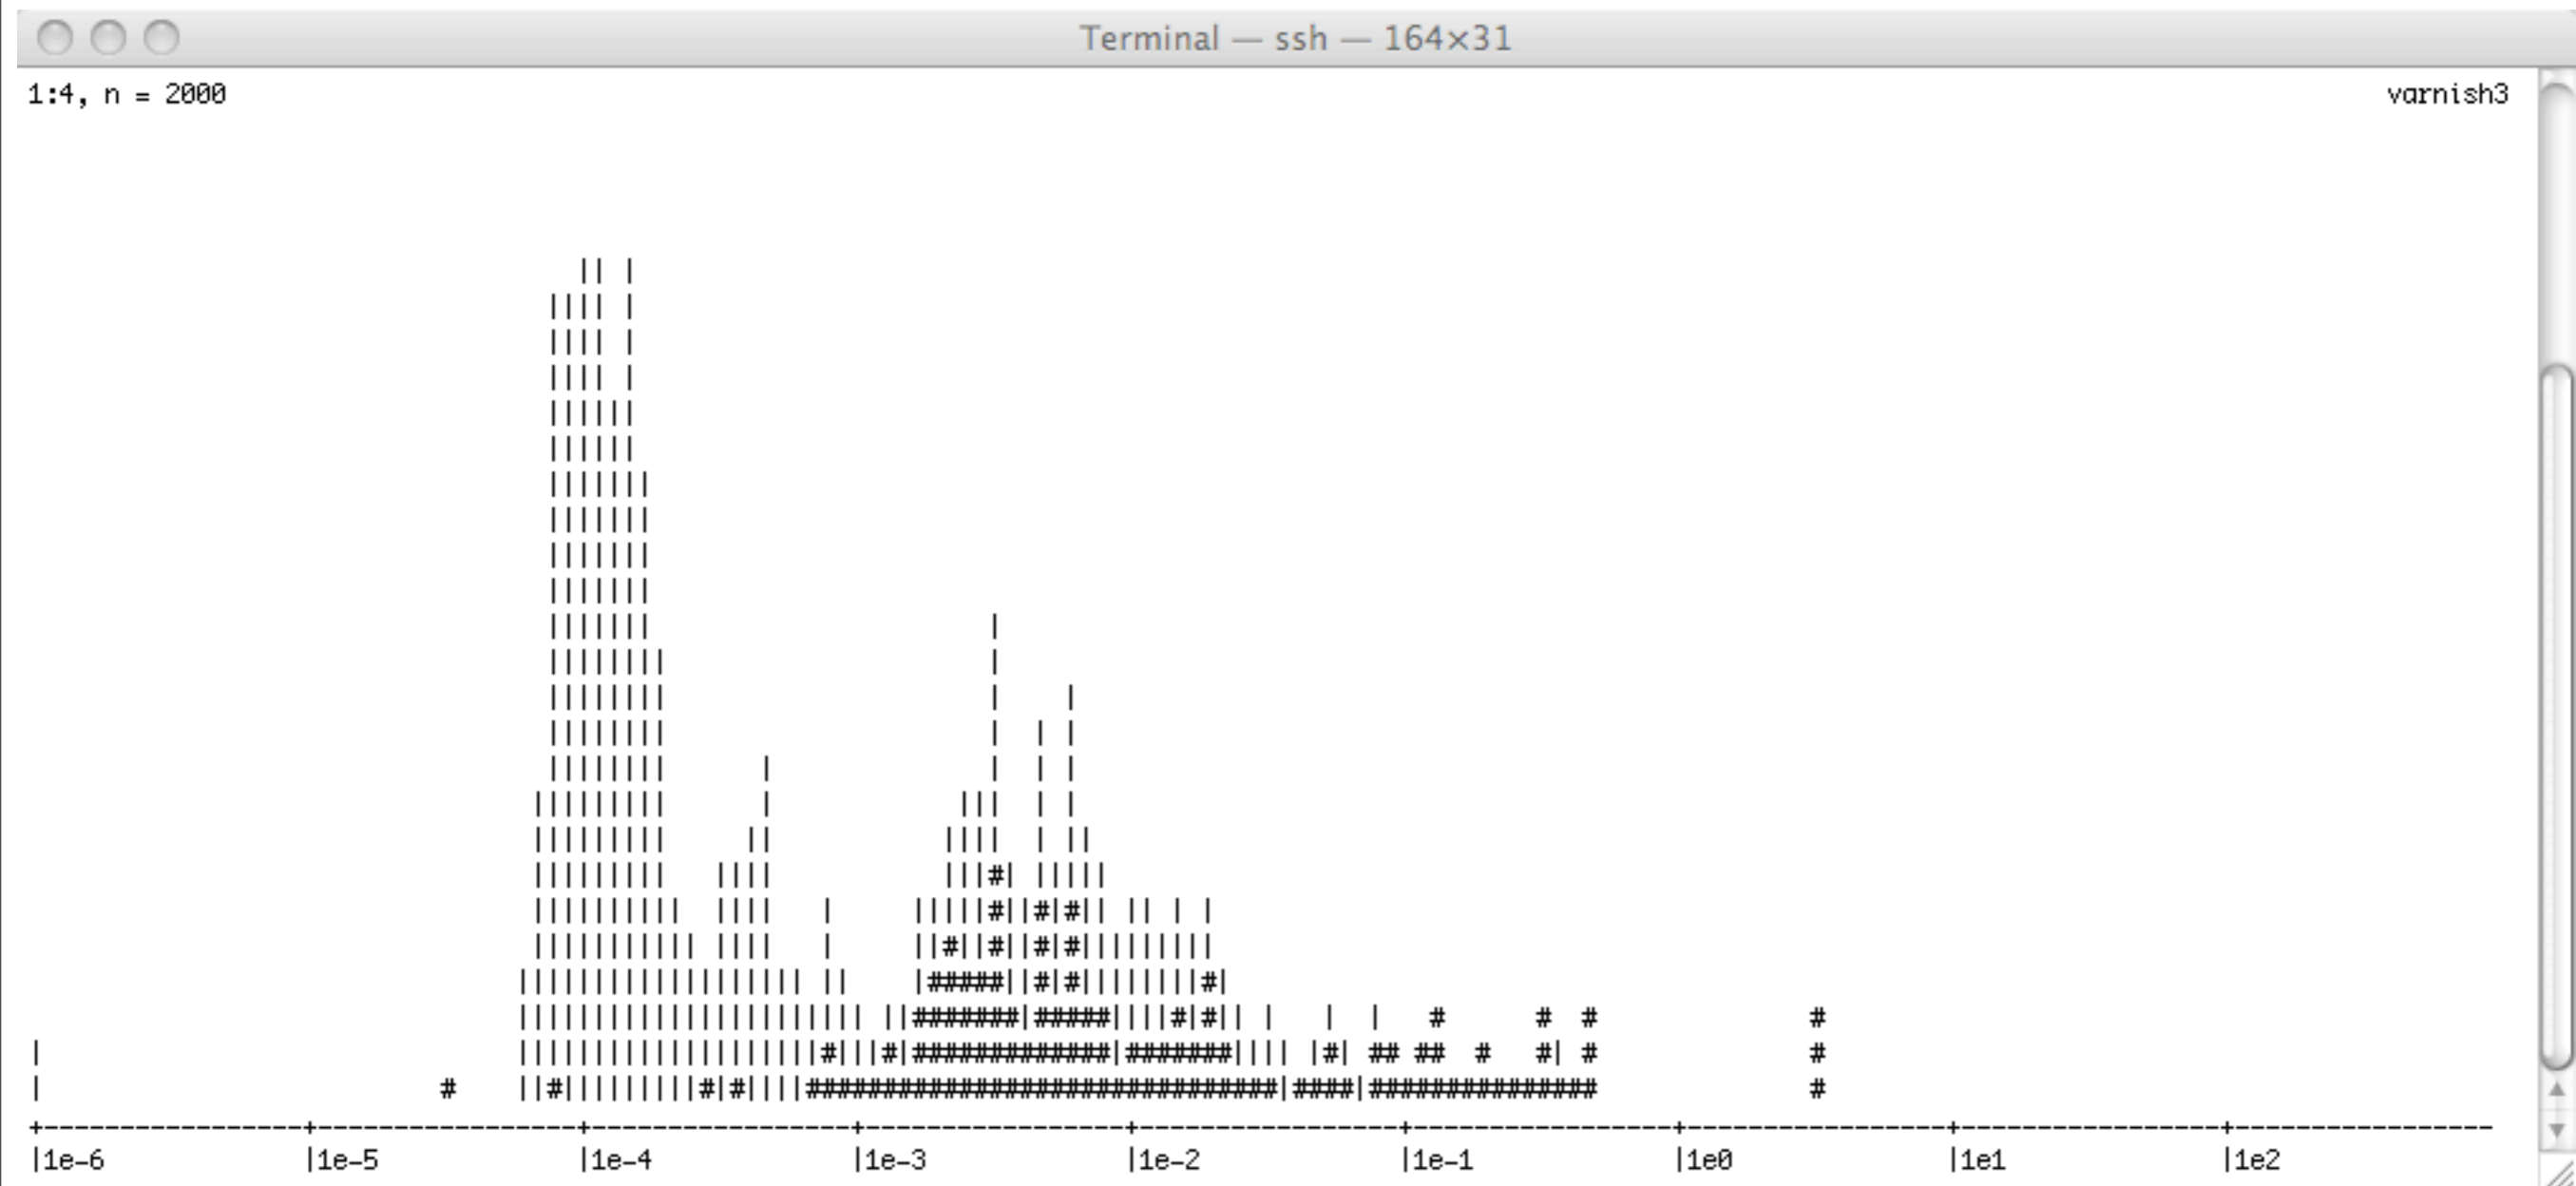
xxx.xxx.xxx.xxx - - [23/Jul/2009:05:49:55 +0000] "GET http://gijoe.wikia.com/extensions/wikia/StaticChute/?type=css&packages=monaco_css&checksum=a5dc1f8a009ce63aea7661b1ba330a8 HTTP/1.1" 200 16570 "http://gijoe.wikia.com/wiki/Duke_(Movie)" "Mozilla/4.0 (compatible; MSIE 7.0; Windows NT 6.0; GTB6; SLCCI; .NET CLR 2.0.50727; Media Center PC 5.0; .NET CLR 3.0.04506; InfoPath.1)"

xxx.xxx.xxx.xxx - - [23/Jul/2009:05:49:55 +0000] "GET <http://www.wowwiki.com/api.php?action=parse&prop=text&text={{:He%20Feeds%20On%20Your%20Tears|mode=home}}&format=json> HTTP/1.1" 200 928 "http://www.wowwiki.com/Algalon_the_Observer" "Mozilla/5.0 (Windows; U; Windows NT 6.0; en-US; rv:1.9.0.12) Gecko/2009070611 Firefox/3.0.12"

xxx.xxx.xxx.xxx - - [23/Jul/2009:05:49:55 +0000] "GET <http://images1.wikia.nocookie.net/uncyclopedia/images/thumb/b/bb/Wotm.jpg/70px-Wotm.jpg> HTTP/1.1" 304 0 "http://uncyclopedia.wikia.com/wiki/Main_Page" "Mozilla/5.0 (Windows; U; Windows NT 6.0; en-US) AppleWebKit/530.5 (KHTML, like Gecko) Chrome/2.0.172.37 Safari/530.5"

xxx.xxx.xxx.xxx - - [23/Jul/2009:05:49:55 +0000] "GET http://banjokazooie.wikia.com/wiki/Jiggy_Switch HTTP/1.1" 200 11041 "-" "Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.0;.NET CLR 1.0.3705; ContextAd Bot 1.0)"

varnishhist



| cache hit
cache miss

61+09:08:16

varnish9

Hirate ratio: 10 16 16
Hirate avg: 0.9274 0.9286 0.9286


1554133248	377.00	293.05	Client connections accepted
3978072764	867.00	750.11	Client requests received
3614483333	801.00	681.55	Cache hits
11639361	0.00	2.19	Cache hits for pass
325788549	61.00	61.43	Cache misses
182821195	19.00	34.47	Backend connections success
25954	0.00	0.00	Backend connections failures
175872686	19.00	33.16	Backend connections reuses
176615269	14.00	33.30	Backend connections recycles
35452	.	.	N struct sess_mem
52444	.	.	N struct sess
2605151	.	.	N struct object
2532375	.	.	N struct objecthead
5293878	.	.	N struct smf
48126	.	.	N small free smf
33357	.	.	N large free smf
93	.	.	N struct vbe_conn
1427	.	.	N struct bereq
2000	.	.	N worker threads
2000	0.00	0.00	N worker threads created
6447	0.00	0.00	N overflowed work requests
13496	0.00	0.00	N dropped work requests
19	.	.	N backends
163660934	.	.	N expired objects
1101311441	.	.	N LRU moved objects
2034	0.00	0.00	HTTP header overflows
2296939595	419.00	433.12	Objects sent with write
1554124359	368.00	293.05	Total Sessions
3978485680	863.00	750.19	Total Requests

varnishstat

Performance

- *Very fast*
- 6 varnish machines handle all of Wikia
- 3 locations -- each location needs 1 varnish
- close to 800 mbit

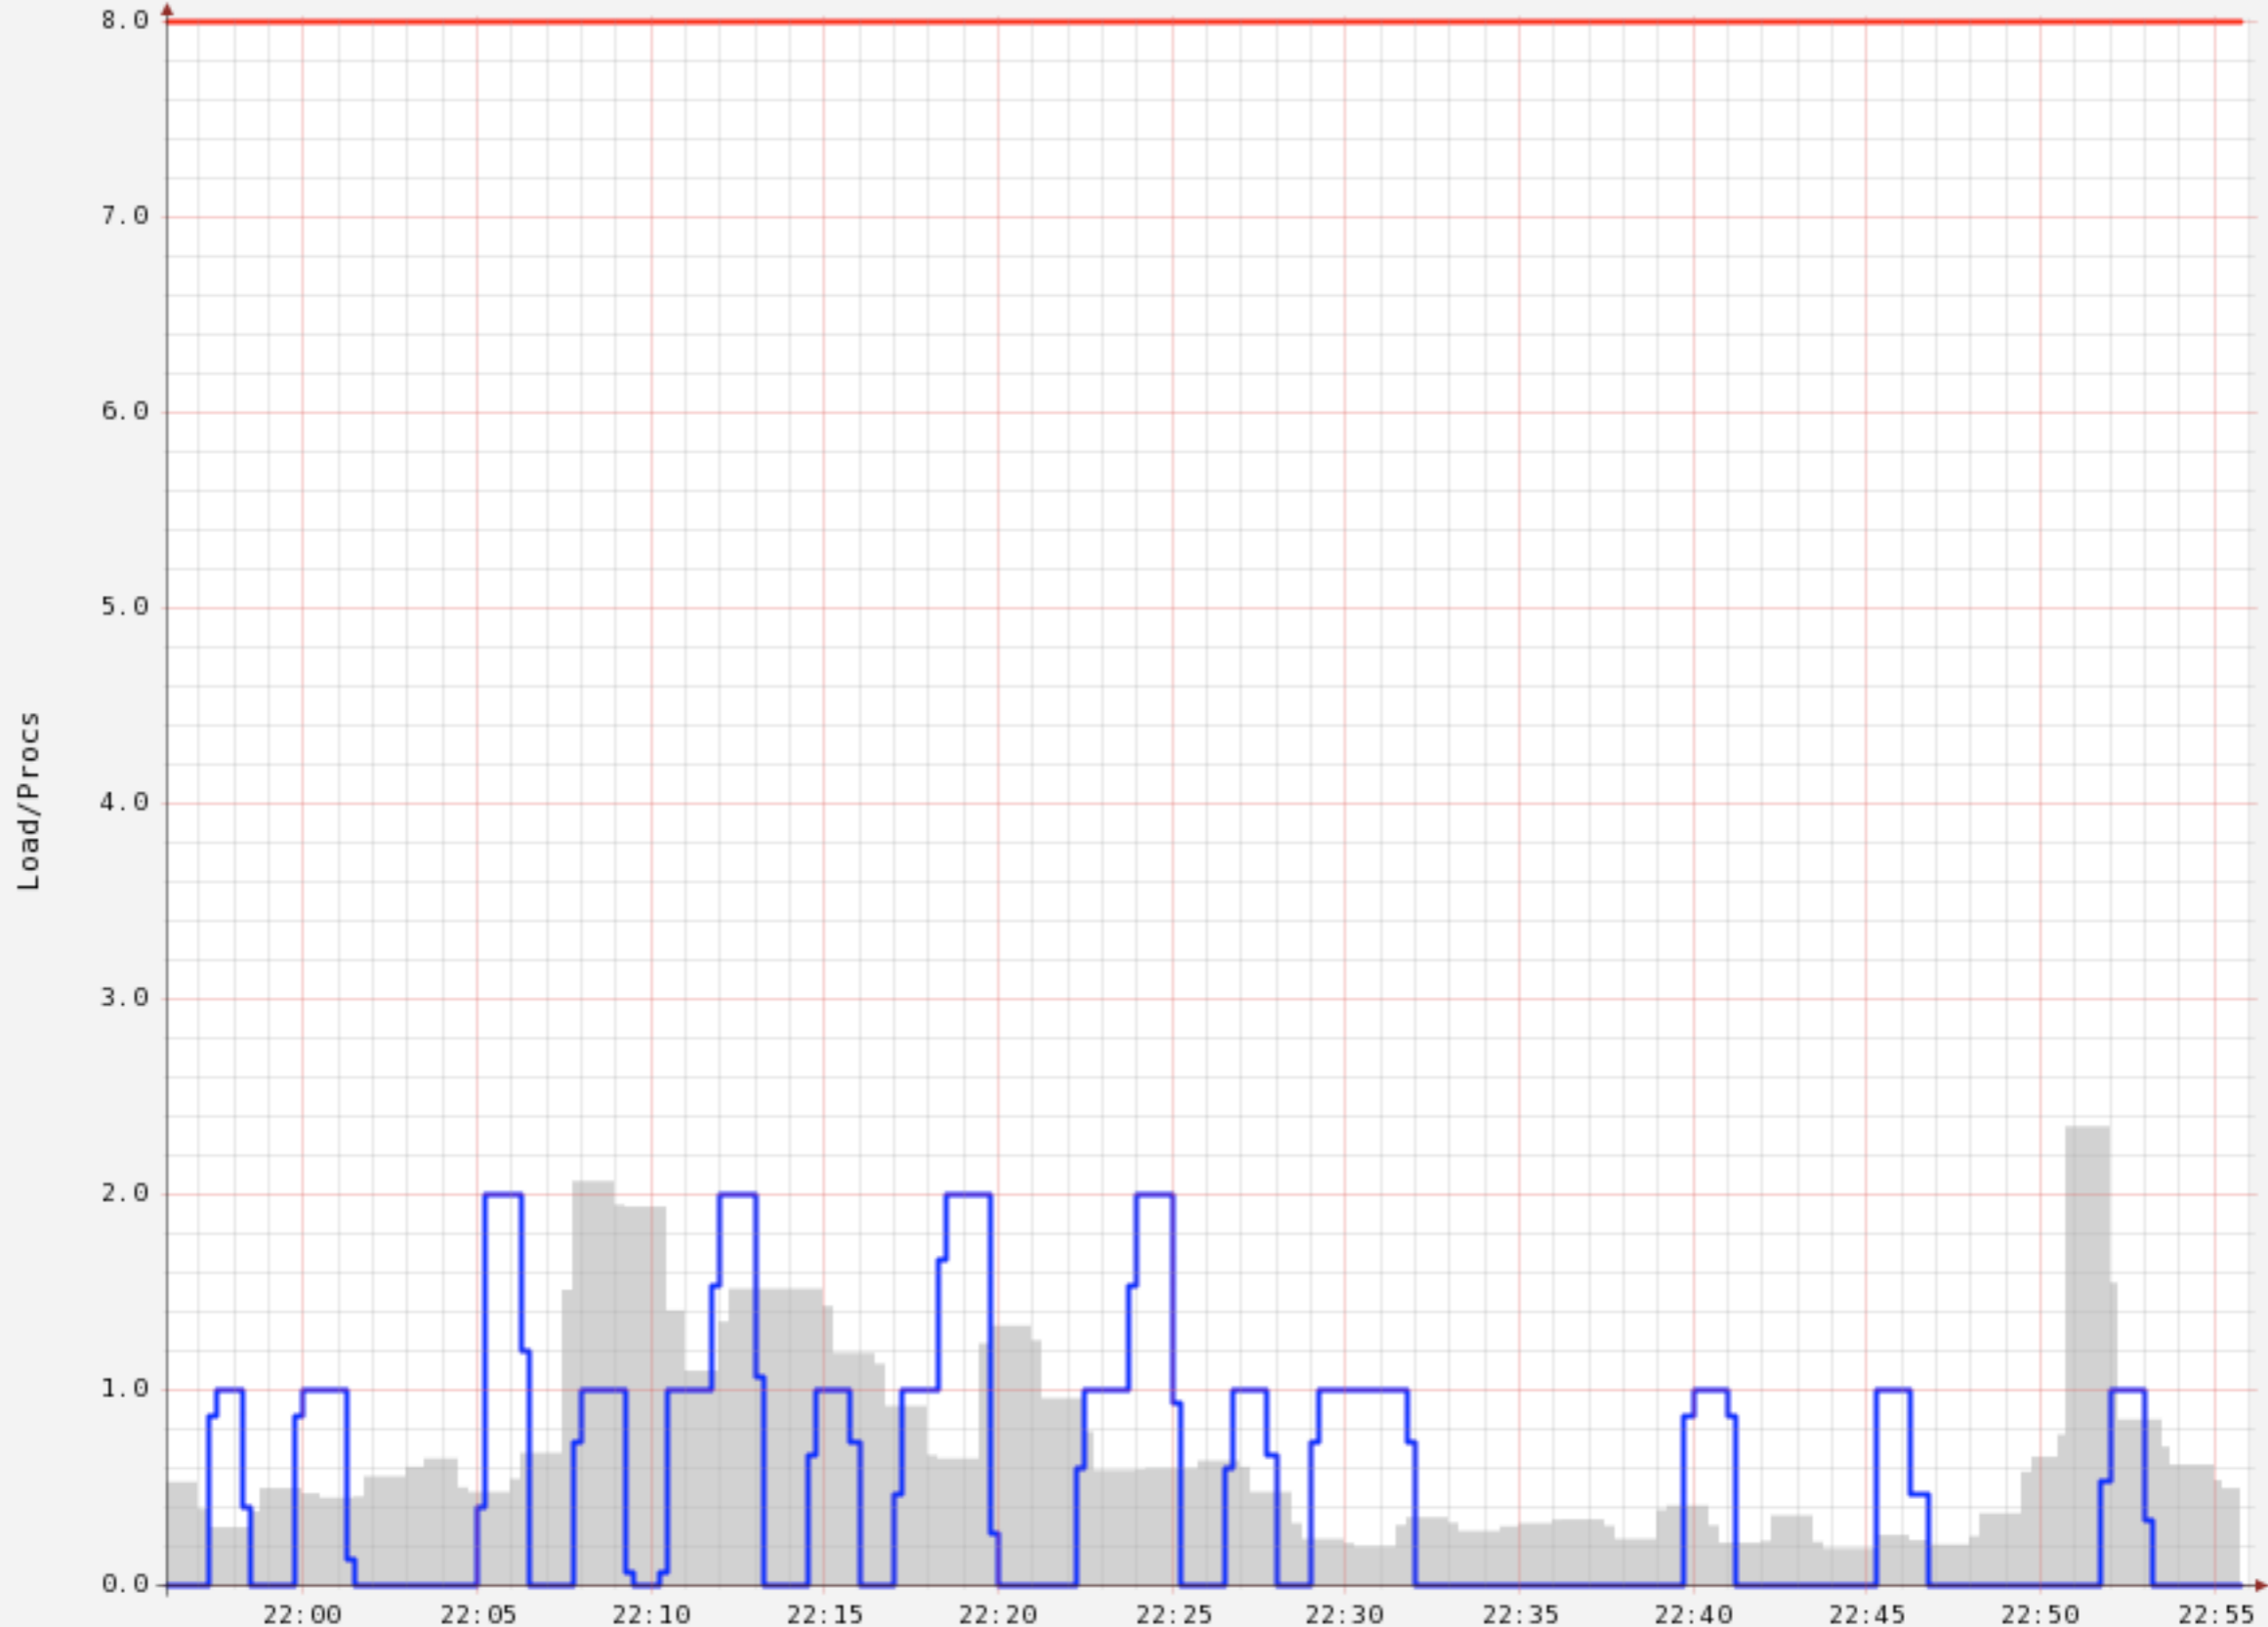
varnish9.wikia.net Network last hour



RRDTOOL / TOBI OETIKER

varnish9.wikia.net Load last hour


RRDTOOL / TOBI OETIKER




1-min Load CPUs Running Processes

Client requests received

RRDTOOL / TOBI OETIKER

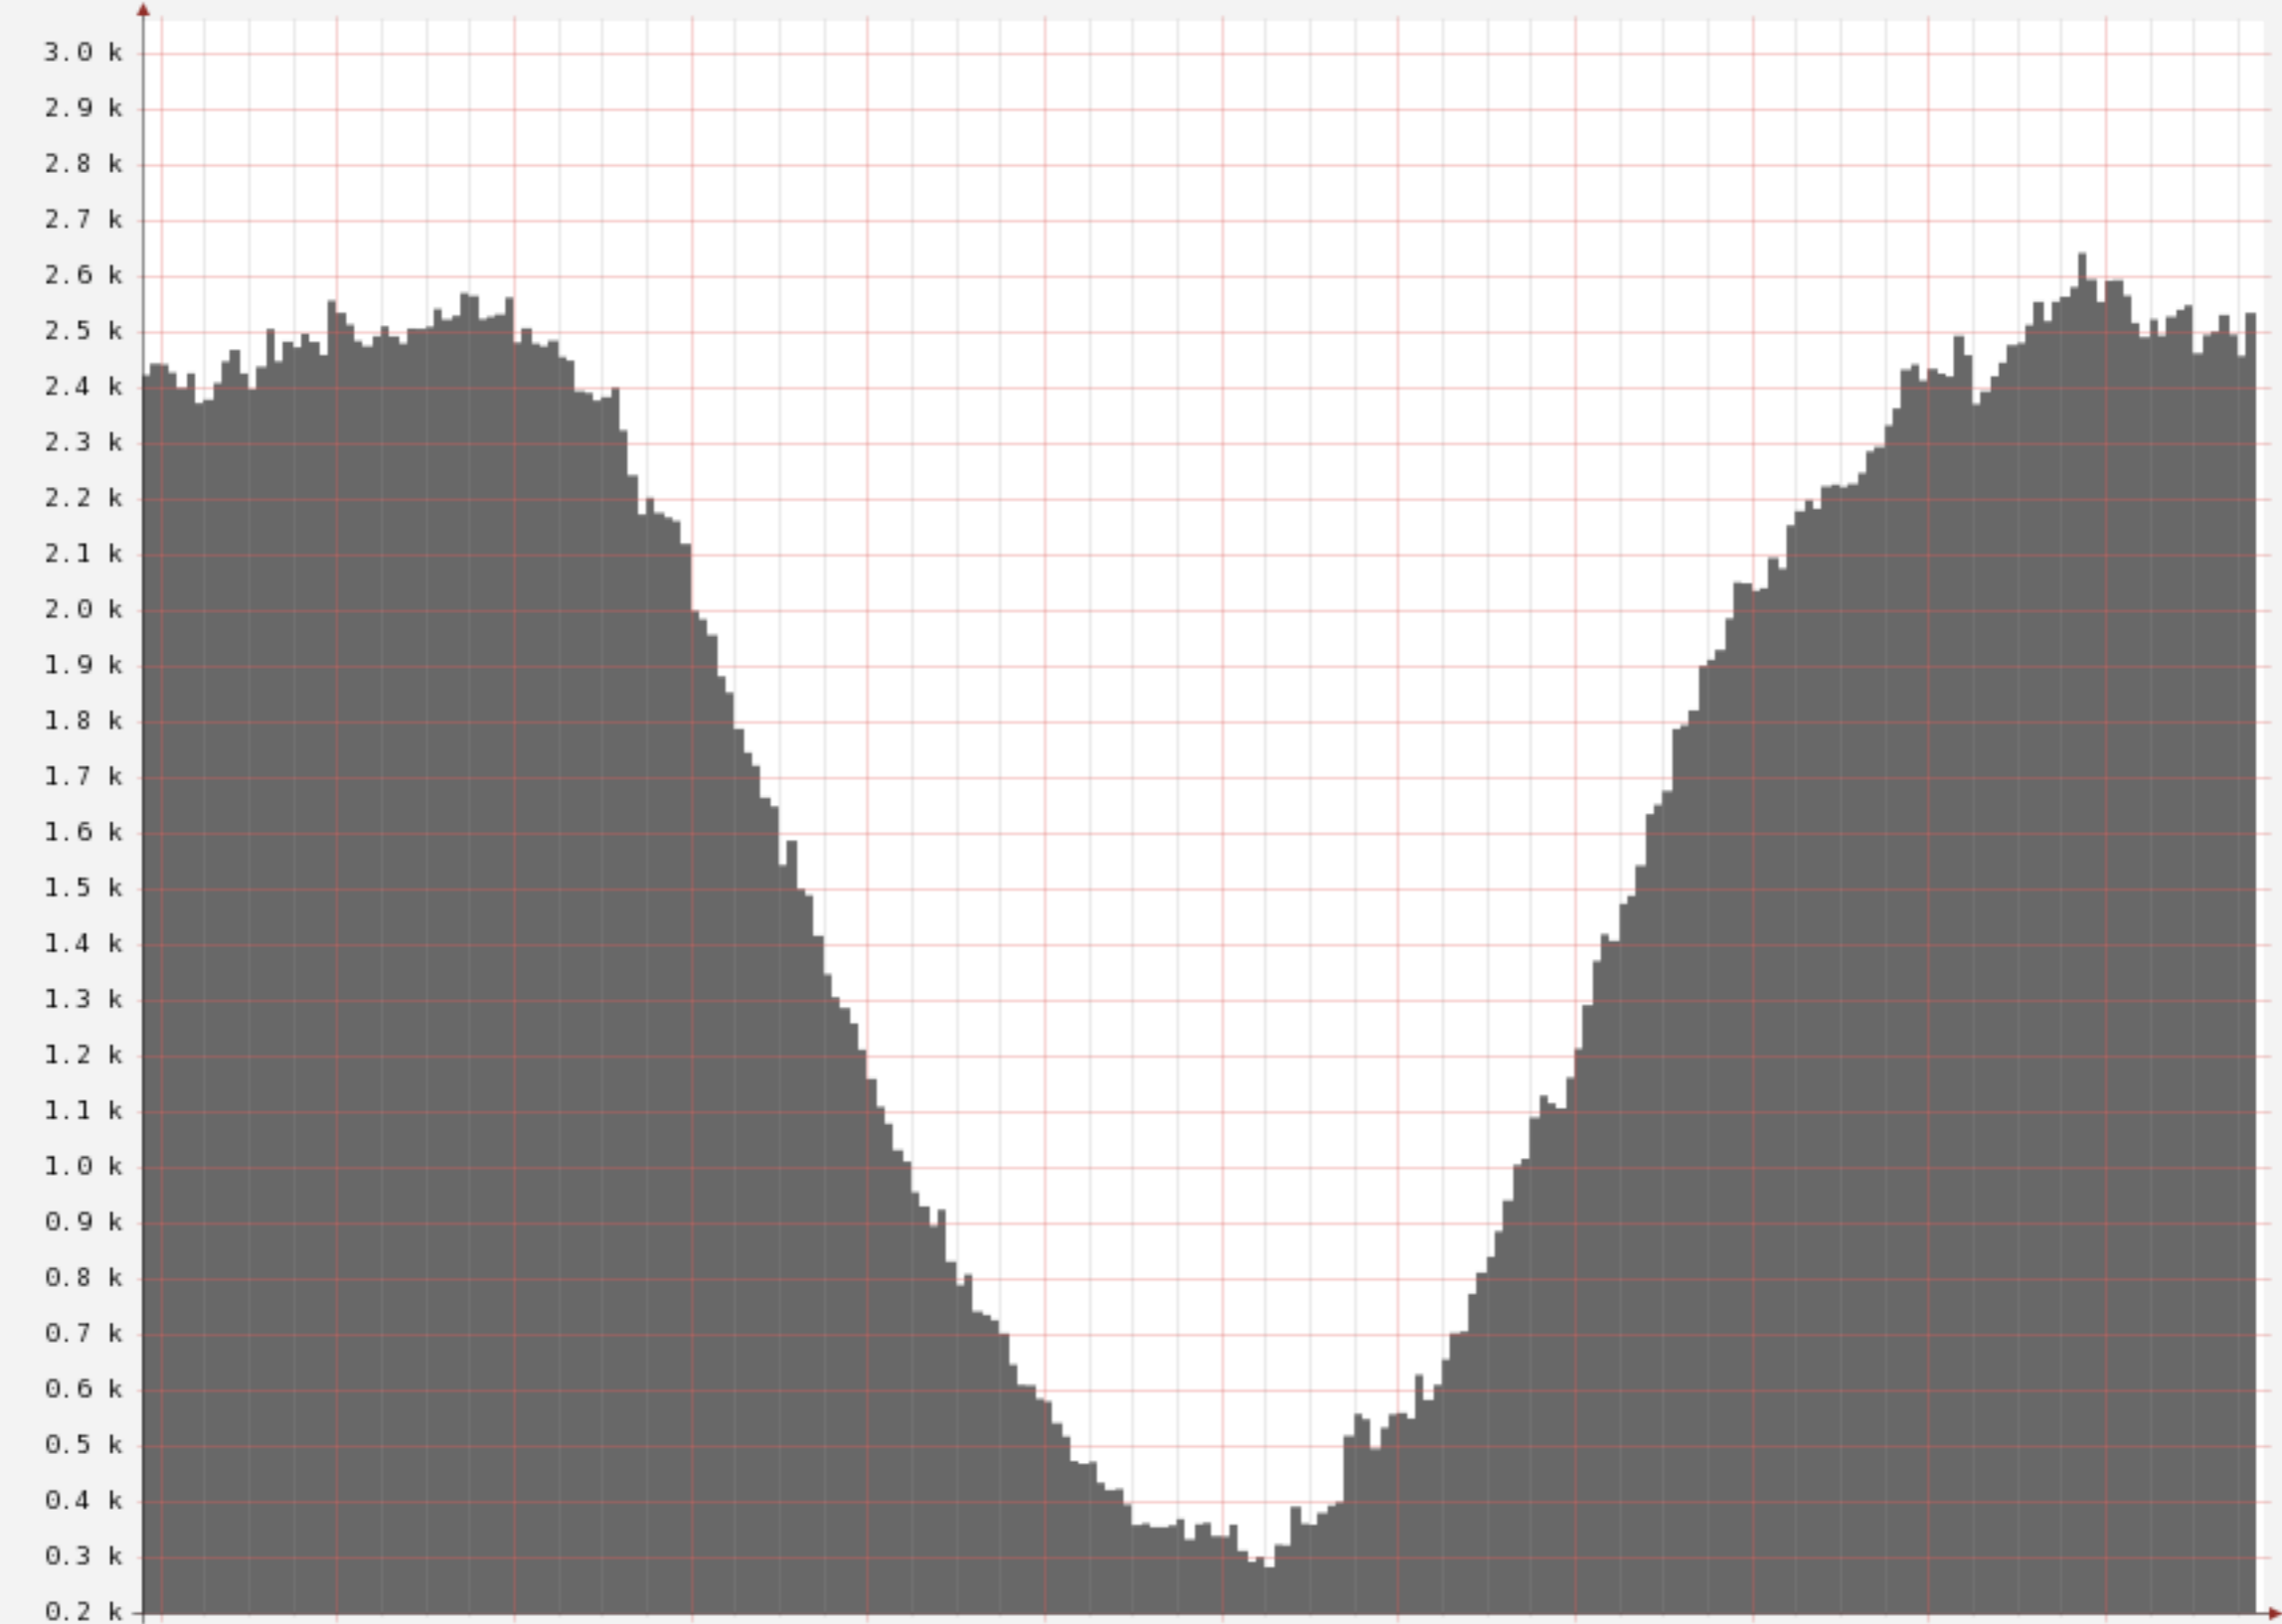


■ varnish9.wikia.net last hour (now 713.00)

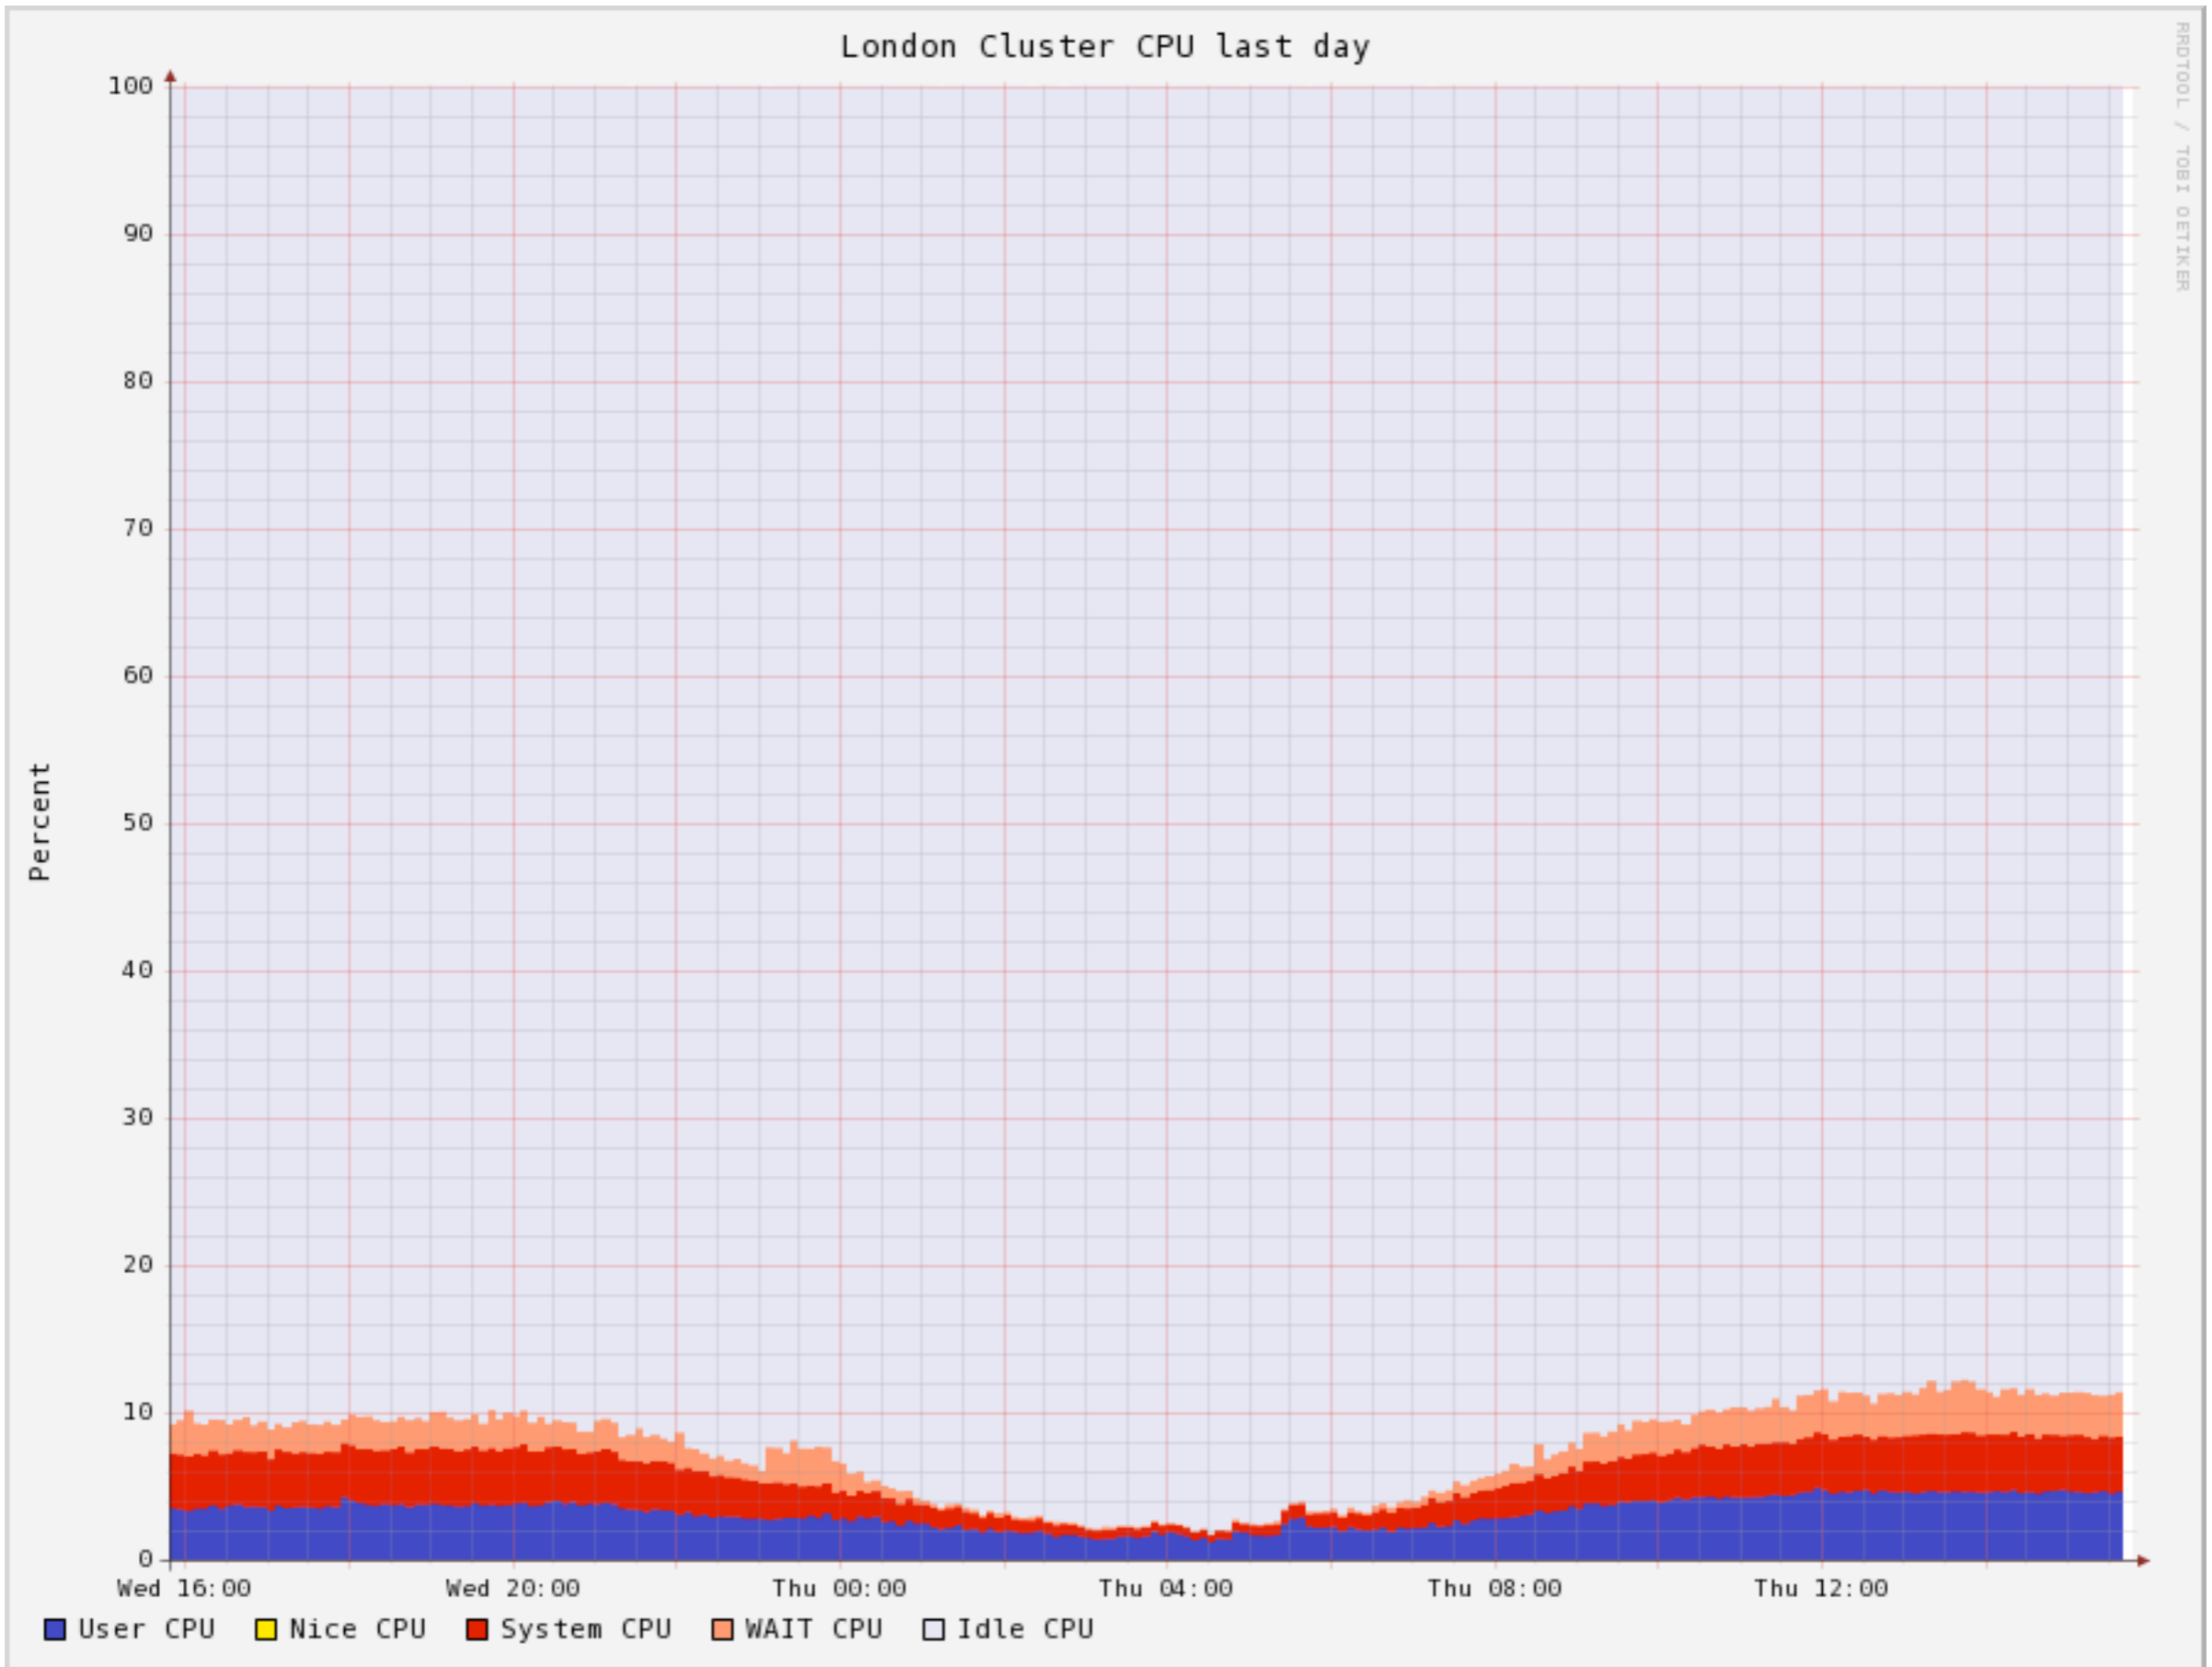


Client requests received


RRDTOOL / TOBI OETIKER



■ `__SummaryInfo__` last day (now 2,500)



Navigator Graph



Median Ping RTT (103.3 ms avg) 0 1/20 2/20 3/20 4/20 10/20 19/20


Packet Loss: 1.29 % average 11.04 % maximum 10.62 % current

Probe: 20 HTTP pings using echoping(1) every 300 seconds created on Wed Sep 20 01:21:23 2006



andersb@vg.no

Navigator Graph



Median Ping RTT (57.8 ms avg) ■ 0 ■ 1/20 ■ 2/20 ■ 3/20 ■ 4/20 ■ 10/20 ■ 19/20
Packet Loss: 0.72 % average 19.59 % maximum 0.43 % current
Probe: 20 HTTP pings using echoping(1) every 300 seconds created on Wed Sep 20 12:26:27 2006



andersb@vg.no

synthetic benchmarks

- 1.8 gbit/s 15% CPU (2500 requests / second)
- max 64000 requests per second 70% cpu

Our own CDN

- Cache HTML
- Fine grained access control
- Cheaper
- London node good example

Hardware

- 8 cores
- Intel(R) Xeon(R) CPU E5420 @ 2.50GHz
- 16/32 GB RAM
- 2 x Intel X25-M SSD

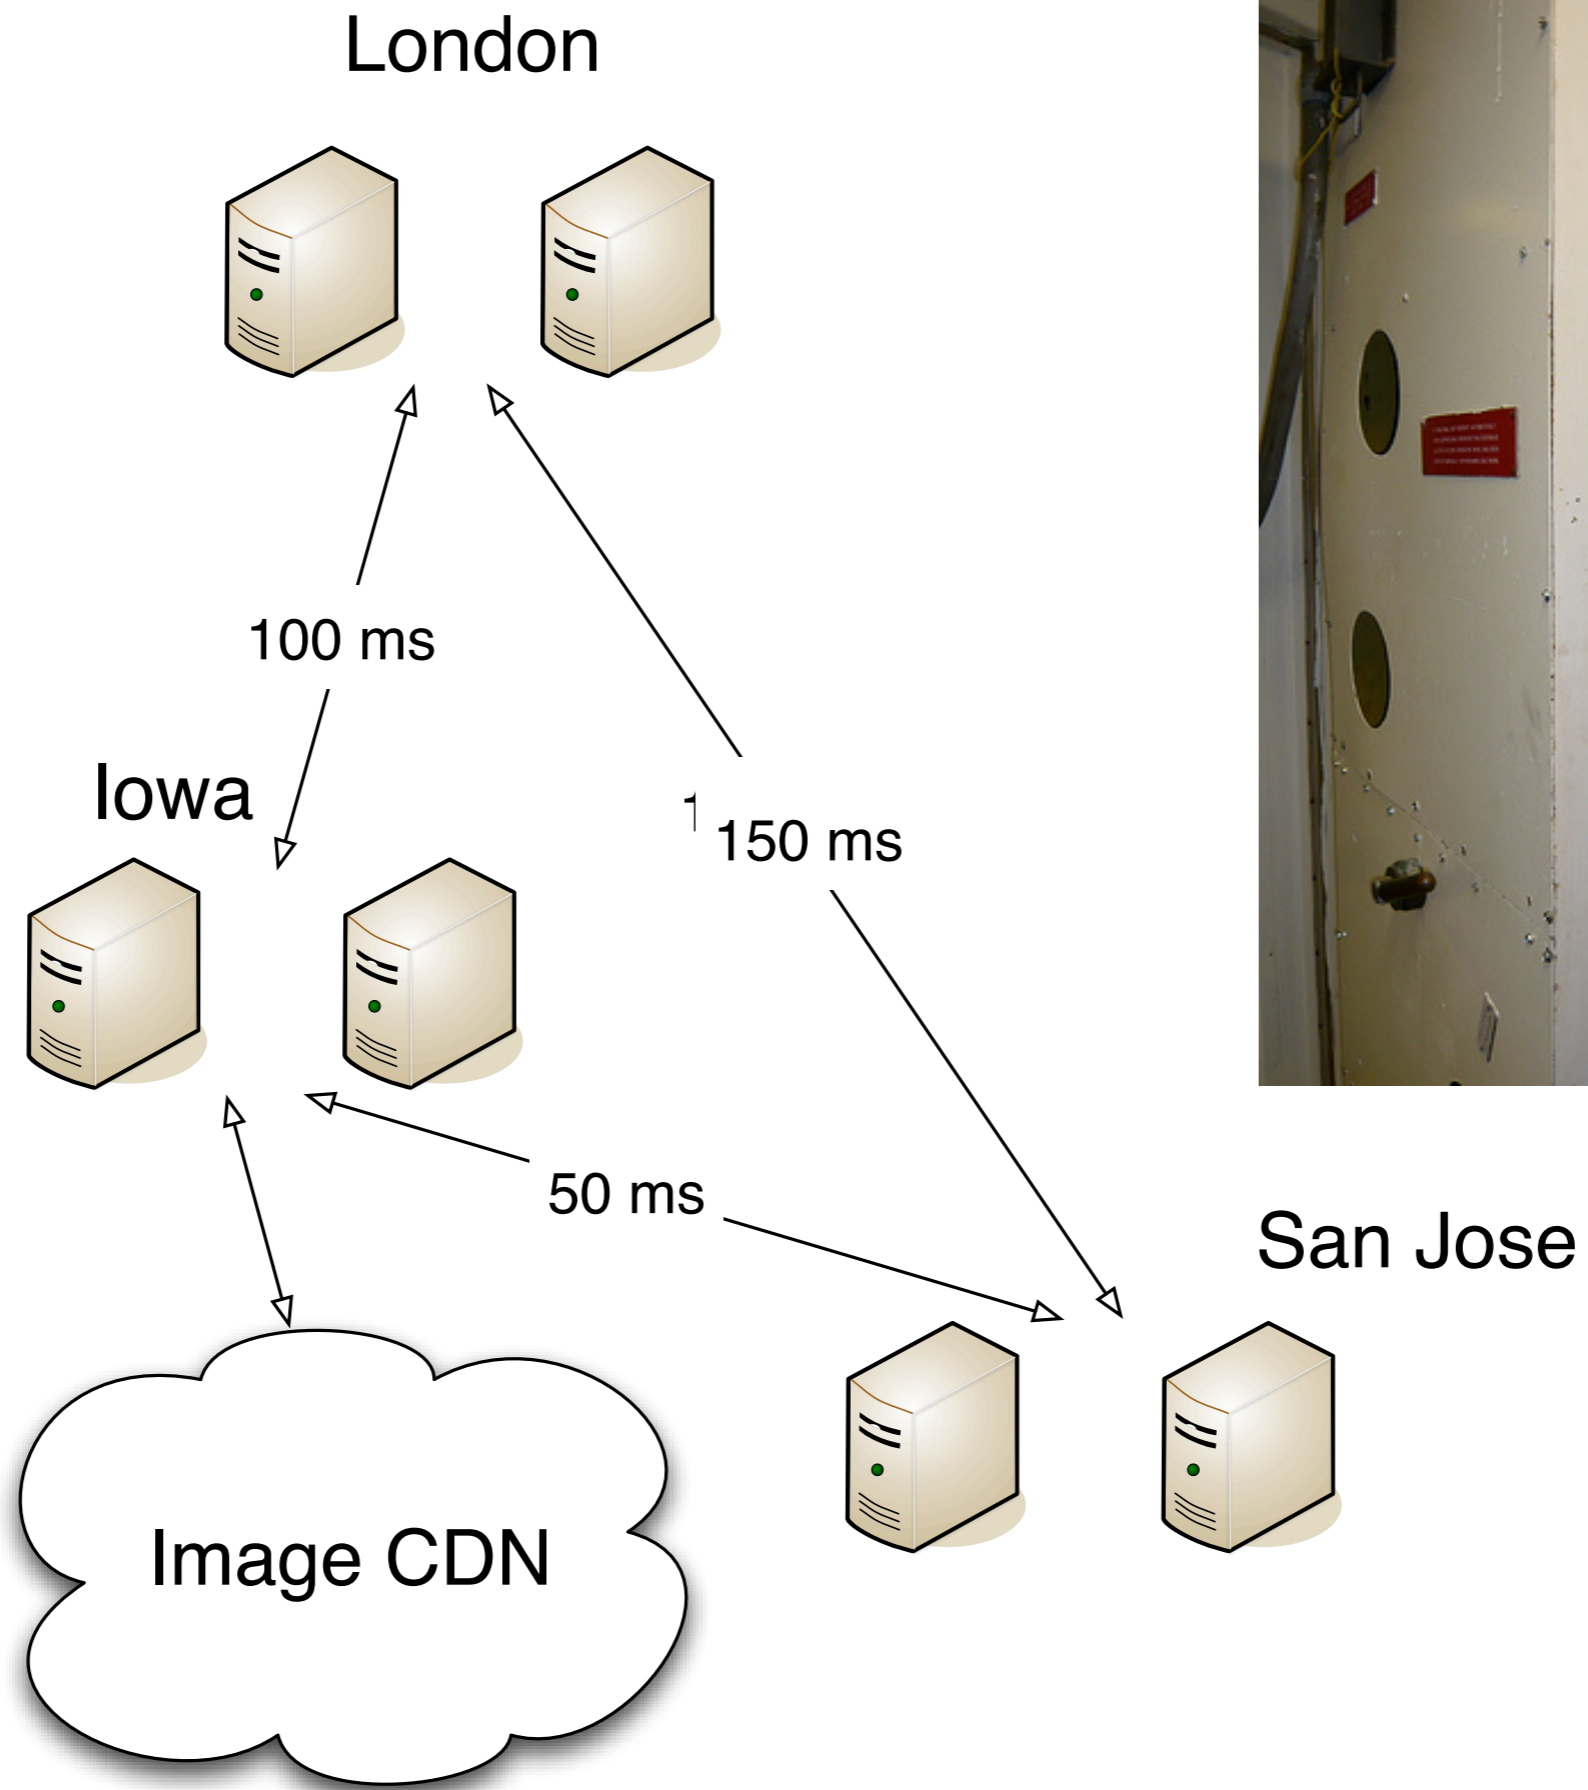
Software

- Ubuntu
- Linux varnish3 2.6.30-wikia #1 SMP
- varnish 2.0.4 + patches
 - defer accept
 - mincore stats
- quagga
- DNS + Dynect

- San Jose
- Iowa
- London

- 2 servers each





Cache hierarchy

- < X-Served-By: varnish3, varnish6, varnish9
- < X-Cache: HIT, MISS, HIT
- < X-Cache-Hits: 5, 0, 137

SSD Love

- Very cost effective
- I love them
- Random IO out of this world
- Varnish uses lots of random IO

SSD optimisation

- noatime
- elevator doesn't seem to matter
- turn off journal
- turn off readahead using hdparm
- turn off all readahead you can
 - cut IO read rate by 10x 80 MB/sec > 8 MB/sec
- don't use a RAID card

Device:	rrqm/s	wrqm/s	r/s	w/s	rMB/s	wMB/s	avgrq-sz	avgqu-sz	await	svctm	
%util											
sdb	0.00	0.00	333.00	0.00	1.30	0.00	8.00	0.07	0.21	0.19	6.32
sdc	0.20	0.00	358.40	0.00	1.40	0.00	8.00	0.07	0.19	0.18	6.56
md0	0.00	0.00	691.60	0.00	2.70	0.00	8.00	0.00	0.00	0.00	0.00

- off peak
- full cache
- not many writes

ESI

- Edge side includes
- Akamai standard
- Currently doesn't support
 - gzip
 - if-modified-since



Edit this page History Delete Move Change protection Watch Article Discussion

WotLK HQ · Patch 3.1.0 / 3.2.0 Ulduar · Argent Tournament View the BlizzCon ticket contest winners! [dismiss]

Class

For the classes in the Warcraft RPG, see RPG Classes. For a short description on class lore, see Class lore.

A class is the primary adventuring style of a player character which determines the type of weapons and armor it can use, as well as what abilities, powers, skills, and spells it will gain throughout its adventures.

Be aware that the choice of which class to play is constrained by the choice of race; each race has a different group of available classes to choose from. Some classes (Shaman, Paladin, and Druid) are available to only one race within a faction. Draenei are the only Alliance race that can play as Shamans; Blood Elves are the only Horde race that can play as Paladins. Night Elves (Alliance) and Tauren (Horde) are the only races that can play as Druids for their respective factions.

Each class has three unique talent trees. As a character progresses, they can specialize their skills by assigning class talent points in any of the trees for their class. There are three talent trees for each class. For

Classes:										
Class races:	Dk	Dr	Hu	Ma	Pa	Pr	Ro	Sh	Wi	Wr
Quests:	Dk	Dr	Hu	Ma	Pa	Pr	Ro	Sh	Wi	Wr
Abilities:	Dk	Dr	Hu	Ma	Pa	Pr	Ro	Sh	Wi	Wr
Trainers:	Dk	Dr	Hu	Ma	Pa	Pr	Ro	Sh	Wi	Wr
Talents:	Dk	Dr	Hu	Ma	Pa	Pr	Ro	Sh	Wi	Wr
Glyphs:	Dk	Dr	Hu	Ma	Pa	Pr	Ro	Sh	Wi	Wr
Talent builds:	Dk	Dr	Hu	Ma	Pa	Pr	Ro	Sh	Wi	Wr
Tactics:	Dk	Dr	Hu	Ma	Pa	Pr	Ro	Sh	Wi	Wr
Armor sets:	Dk	Dr	Hu	Ma	Pa	Pr	Ro	Sh	Wi	Wr
Starting a:	Dk	Dr	Hu	Ma	Pa	Pr	Ro	Sh	Wi	Wr
PvE guide:	Dk	Dr	Hu	Ma	Pa	Pr	Ro	Sh	Wi	Wr
PvP guide:	Dk	Dr	Hu	Ma	Pa	Pr	Ro	Sh	Wi	Wr
Useful macros:	Dk	Dr	Hu	Ma	Pa	Pr	Ro	Sh	Wi	Wr

```
<esi:include src="/header" />
```

Collecting class sets is a way players can improve a character when they have reached the level cap. Most dungeon and raid sets are made up of armor items designed specifically for the class; however, faction rewards sets also include trinkets and neck pieces. When players collect more pieces of the set, they can receive set bonuses.

Always remember that it doesn't matter what class you chose as long as you are willing to stand by it and don't give up. If your class gets nerfed (or made less powerful by Blizzard) always remember that nerfs and buffs are like a pendulum, sooner or later the pendulum is going to swing to your advantage.

Contents [hide]

- 1 Brief summary of each class
 - 1.1 Original classes
 - 1.2 Hero classes
- 2 Class roles
 - 2.1 Legend
 - 2.2 Ratings
- 3 Class colors
- 4 See also
- 5 External links

Brief summary of each class [edit]

There are currently ten playable classes in-game.

Search WoWWiki

- WoWWiki portals
- World of Warcraft
- Interface customization
- Warcraft universe
- Community
- Popular content
- Help
- Wiki upgrades and skin info
- Write an Article
- Add Image
- Advanced search
- Google search
- Recent changes
- Random article
- Help
- Special pages
- What links here

75,852 articles on this wiki

Welcome back, Crucially
My page | My talk | Widgets

Latest Activity

- Clockwork Assistant by Benbold14, 54 seconds ago
- Jepetto Joybuzz by Benbold14, one minute ago

more...

What's Hot Growth in readership

- 700% BRAWL HACKS
- 105% OBLIVION

ESI future

- Chain walking for if-modified-since
- Synthetic parts
- Gzip support

Common problems

- Cache headers
 - s-maxage = for varnish
 - maxage = for clients
- Warning s-maxage also for other caches
- Treat varnish as part of your application
- X-Pass-Cache-Control hack

Common problems 2

- Incorrect cache headers
 - Cache minimum 1 sec
 - prevents DOS to backend
- Cachebusters
 - I hate you jquery!
 - ?randomnumber f**k you

Thank you

- artur@crucially.net
- varnish.projects.linpro.no
- [#varnish irc.linode.com](irc://irc.linode.com/#varnish)

```

0:00:27:22
Hitrate ratio: 10 100 1000
Hitrate avg: 0.6130 0.6550 0.6065

5123 3.00 3.12 Client connections accepted
17937 7.99 18.92 Client requests received
8896 5.99 4.93 Cache hits
0 0.00 0.00 Cache hits for pass
4440 1.00 2.71 Cache misses
4712 2.00 2.87 Backend connections success
0 0.00 0.00 Backend connections not attempted
0 0.00 0.00 Backend connections too many
1 0.00 0.00 Backend connections failures
2621 2.00 1.60 Backend connections reuses
4439 1.00 2.70 Backend connections recycles
0 0.00 0.00 Backend connections unused
1257 . . N struct sroaddr
13 . . N active struct sroaddr
35 . . N struct sess_wm
14 . . N struct sess
2412 . . N struct object
2334 . . N struct objecthead
0 . . N struct saf
0 . . N small free saf
0 . . N large free saf
2 . . N struct vbe_conn
6 . . N struct bereq
0 . . N worker threads
1800 0.00 0.61 N worker threads created
0 0.00 0.00 N worker threads not created
4436 0.00 2.70 N worker threads limited
0 0.00 0.00 N queued work requests
0 0.00 0.00 N overflowed work requests
0 0.00 0.00 N dropped work requests
5 . . N backends
2840 . . N expired objects
0 . . N LRU nuked objects
0 . . N LRU saved objects
6394 . . N LRU saved objects
0 . . N objects on deathrow
11 0.00 0.01 HTTP header overflows
0 0.00 0.00 Objects sent with sendfile
13385 5.99 0.13 Objects sent with write
0 0.00 0.00 Objects overflowing workspace
5120 2.00 3.12 Total Sessions
17935 6.99 18.92 Total Requests
263 0.00 0.16 Total pipe
0 0.00 0.00 Total pass
4447 1.00 2.71 Total fetch
8768551 3889.00 5335.29 Total header bytes
69174335 18694.01 42128.10 Total body bytes
1249 1.00 0.76 Session Closed
0 0.00 0.00 Session Pipeline
2 0.00 0.00 Session Read Ahead
0 0.00 0.00 Session Linger
16761 5.99 18.21 Session herd
977779 369.62 595.40 SHM records
53467 23.90 32.56 SHM writes
0 0.00 0.00 SHM flushes due to overflow
2 0.00 0.00 SHM HTX contention
0 0.00 0.00 SHM cycles through buffer
0 0.00 0.00 allocator requests
0 . . outstanding allocations
0 . . bytes allocated
0 . . bytes free
13673 3.00 0.33 SMA allocator requests
4540 . . SMA outstanding allocations
112449985 . . SMA outstanding bytes
529843279 . . SMA bytes allocated
417393374 . . SMA bytes free
5129 0.00 3.12 SMS allocator requests
0 . . SMS outstanding allocations
0 . . SMS outstanding bytes
2128588 . . SMS bytes allocated
2128588 . . SMS bytes free
4447 1.00 2.71 Backend requests wade
1 0.00 0.00 N vcl total
1 0.00 0.00 N vcl available
0 0.00 0.00 N vcl discarded
4781 . . N total active purges
4750 0.00 2.89 N new purges added
49 0.00 0.83 N old purges deleted

```

```

0:04:59:49
Hitrate ratio: 10 100 1000
Hitrate avg: 0.8389 0.8337 0.8261

3582969 273.99 172.49 Client connections accepted
13838847 1367.97 768.85 Client requests received
11693498 1118.98 658.84 Cache hits
5487 3.00 0.30 Cache hits for pass
1947657 234.00 188.27 Cache misses
2881343 248.99 115.78 Backend connections success
0 0.00 0.00 Backend connections not attempted
0 0.00 0.00 Backend connections too many
175 0.00 0.81 Backend connections failures
1356518 172.00 75.41 Backend connections reuses
1958875 231.00 188.45 Backend connections recycles
0 0.00 0.00 Backend connections unused
5984 . . N struct sroaddr
4855 . . N active struct sroaddr
14893 . . N struct sess_wm
12388 . . N struct sess
144419 . . N struct object
136855 . . N struct objecthead
368784 . . N struct saf
75632 . . N small free saf
8511 . . N large free saf
128 . . N struct vbe_conn
359 . . N struct bereq
0 . . N worker threads
1800 0.00 0.86 N worker threads created
0 0.00 0.00 N worker threads not created
2298 0.00 0.13 N worker threads limited
0 0.00 0.00 N queued work requests
0 0.00 0.00 N overflowed work requests
0 0.00 0.00 N dropped work requests
5 . . N backends
1585221 . . N expired objects
218258 . . N LRU nuked objects
0 . . N LRU saved objects
2586124 . . N LRU saved objects
0 . . N objects on deathrow
496 0.00 0.83 HTTP header overflows
0 0.00 0.00 Objects sent with sendfile
8388729 963.98 461.71 Objects sent with write
0 0.00 0.00 Objects overflowing workspace
3182832 269.99 172.48 Total Sessions
13832421 1361.97 768.94 Total Requests
128392 11.00 7.14 Total pipe
5487 3.00 0.30 Total pass
1952818 235.00 188.56 Total fetch
7835388998 732496.63 391893.98 Total header bytes
61113661292 6662657.21 3397279.52 Total body bytes
561869 66.00 31.19 Session Closed
897 0.00 0.85 Session Pipeline
1471 0.00 0.88 Session Read Ahead
0 0.00 0.00 Session Linger
13354747 1384.97 742.38 Session herd
738737881 75874.42 48621.37 SHM records
37788974 3462.93 2896.86 SHM writes
138789 18.00 7.71 SHM flushes due to overflow
183677 25.00 5.76 SHM HTX contention
322 0.00 0.82 SHM cycles through buffer
4137627 478.99 238.81 allocator requests
284561 . . outstanding allocations
8189378568 . . bytes allocated
6843886976 . . bytes free
0 0.00 0.00 SMA allocator requests
0 . . SMA outstanding allocations
0 . . SMA outstanding bytes
0 . . SMA bytes allocated
0 . . SMA bytes free
56965 3.00 3.17 SMS allocator requests
0 . . SMS outstanding allocations
0 . . SMS outstanding bytes
24759923 . . SMS bytes allocated
24759923 . . SMS bytes free
1952983 237.00 188.56 Backend requests wade
1 0.00 0.00 N vcl total
1 0.00 0.00 N vcl available
0 0.00 0.00 N vcl discarded
5694 . . N total active purges
18618 2.89 0.59 N new purges added
4924 0.83 0.27 N old purges deleted

```

```

Every 1.0s: echo 'show slave status'G'lay... Wed Sep 18 19:25:24 2008
Seconds_Behind_Master: 0

Every 1.0s: echo 'show slave status'G'lay... Wed Sep 18 19:25:24 2008
Seconds_Behind_Master: 0

Every 1.0s: echo 'show slave status'G'lay... Wed Sep 18 19:25:25 2008
Seconds_Behind_Master: 0

Every 1.0s: echo 'show slave status'G'lay... Wed Sep 18 19:25:25 2008
Seconds_Behind_Master: 0

Every 1.0s: echo 'show slave status'G'lay... Wed Sep 18 19:25:25 2008
Seconds_Behind_Master: 0

Every 2.0s: sudo ipvsadm -L | grep -A 2 216.22... Wed Sep 18 19:25:23 2008
-> http Route 18 86 291
-> http Route 1000 28942 27657

```

```

top - 19:25:25 up 16 days, 4:59, 5 users, load average: 1.56, 1.75, 1.55
Tasks: 151 total, 1 running, 150 sleeping, 0 stopped, 0 zombie
Cpu(s): 12.6%us, 1.9%sy, 0.0%ni, 82.2%id, 1.7%wa, 0.1%st, 0.0%ni, 0.0%ot
Mem: 16432292k total, 1634588k used, 86488k free, 188678k buffers
Swap: 1852248k total, 0k used, 1852248k free, 1458798k cached

PID USER PR NI VIRT RES SHR S NCPU MNEM TIME+ COMMAND
31147 varnish 18 0 24.9g 11g 18g S 162 71.5 85:28.65 lt-varnishd
4751 root 15 0 73584 17k 1548 S 19 0.1 0:29.23 iftop

top - 19:25:25 up 27 days, 16:48, 3 users, load average: 0.81, 0.83, 0.88
Tasks: 156 total, 1 running, 155 sleeping, 0 stopped, 0 zombie
Cpu(s): 0.8%us, 0.8%sy, 0.0%ni, 99.8%id, 0.1%wa, 0.0%st, 0.0%ni, 0.0%ot
Mem: 16432292k total, 628648k used, 15883644k free, 82288k buffers
Swap: 1852248k total, 78864k used, 973384k free, 176252k cached

PID USER PR NI VIRT RES SHR S NCPU MNEM TIME+ COMMAND
1 root 15 0 18312 84 48 S 0 0.0 0:38.59 init
2 root RT 0 0 0 0 S 0 0.0 0:00.00 migration/0

19.1%b 38.1%b 57.2%b 76.3%b 95.4%b
216.224.121.143 <=> 1.39%b 858kb 564kb
216.224.121.143 <=> 1.80%b 798kb 714kb
216.224.121.143 <=> 2.38%b 737kb 104kb
216.224.121.143 <=> 441kb 498kb 288kb
216.224.121.143 <=> 464kb 442kb 111kb
216.224.121.143 <=> 285kb 432kb 125kb
216.224.121.143 <=> 588kb 387kb 383kb
216.224.121.143 <=> 86.5kb 383kb 97.0kb
216.224.121.143 <=> 588kb 376kb 153kb
216.224.121.143 <=> 479kb 371kb 289kb

TX: cum: 500MB peak: 50.0MB rates: 53.0MB 49.0MB 44.1MB
RX: cum: 16.3MB peak: 18.0MB rates: 9.94MB 9.70MB 9.35MB
TOTAL 96.3MB 62.9MB 62.9MB 58.7MB 53.5MB

```

Slow site == bad

