




VARNISH

Makes Websites Fly


You can cache everything


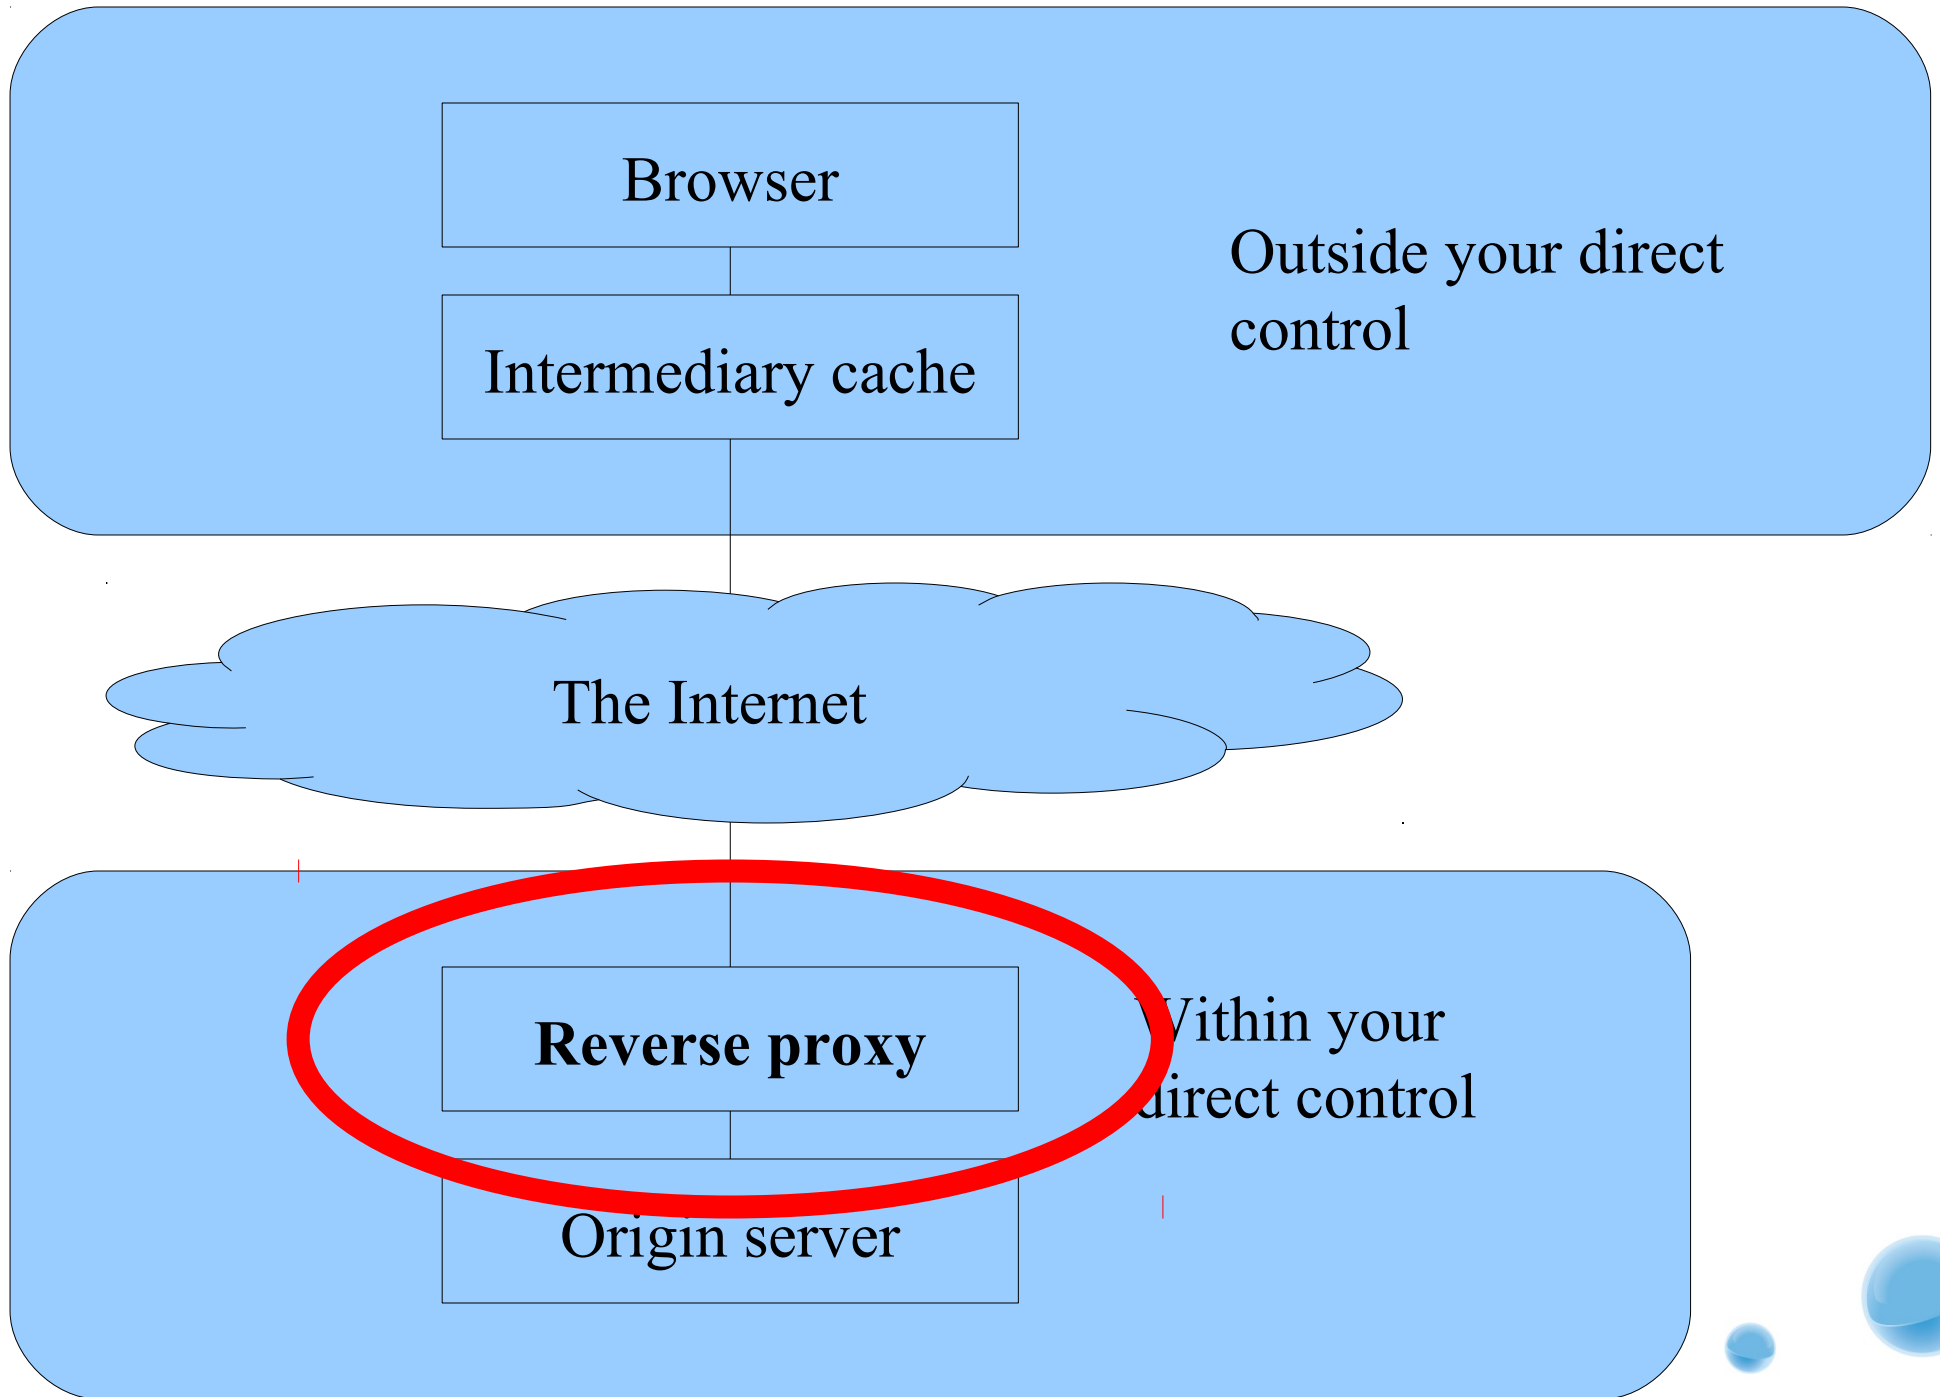
Kristian Lyngstøl
Varnish Software AS
Twitter: @kristianlyng

Montreal, March 2013




What is a reverse proxy anyway?






A reverse proxy...


- Replies in place of your web server
- Fetches content from the web server
- Caches the content
- Is fast
- Is simple






“Can I cache content for less than one second?”


- Audun Ytterdal, Vgnett




(yes, you can)




But WHY?

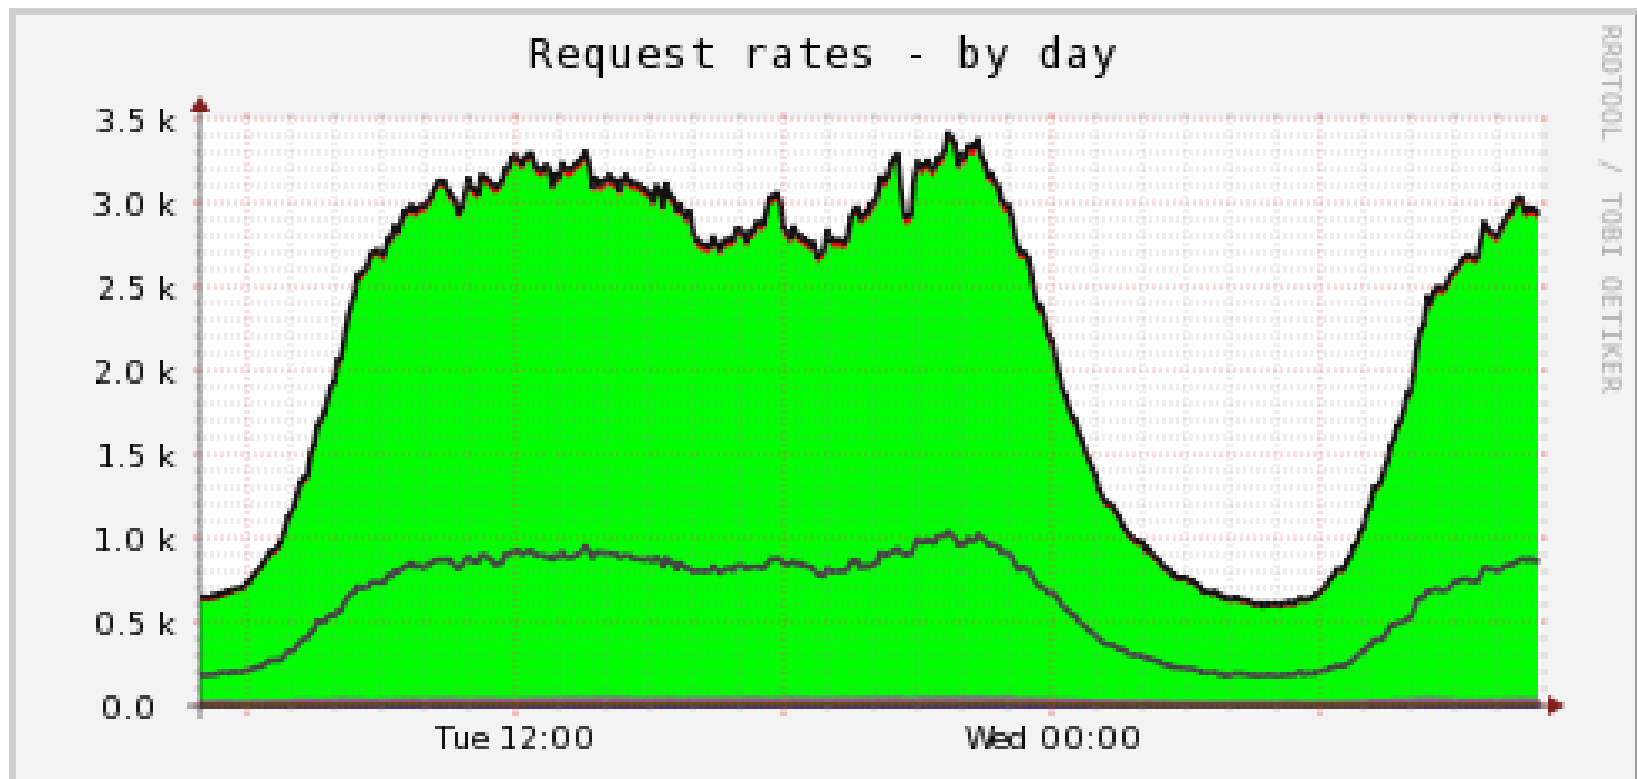


100 req/day
100 req/hour
100 req/minute
100 req/second
4000 req/second




Server load?






Max: 3370 requests per second
Average: 2190 requests per second



CPU usage - by day

RRDTOOL / TOBI OETIKER




- system
- user
- nice
- idle
- iowait
- irq
- softirq
- steal


	Cur:	Min:	Avg:	Max:
system	10.75	2.17	8.06	13.08
user	21.17	5.76	16.28	23.24
nice	0.02	0.00	0.09	1.33
idle	755.33	749.15	766.02	789.34
iowait	0.05	0.04	0.08	0.61
irq	2.87	0.87	2.31	3.17
softirq	9.21	1.52	6.71	10.99
steal	0.00	0.00	0.00	0.00

Last update: Wed Aug 11 10:55:04 2010

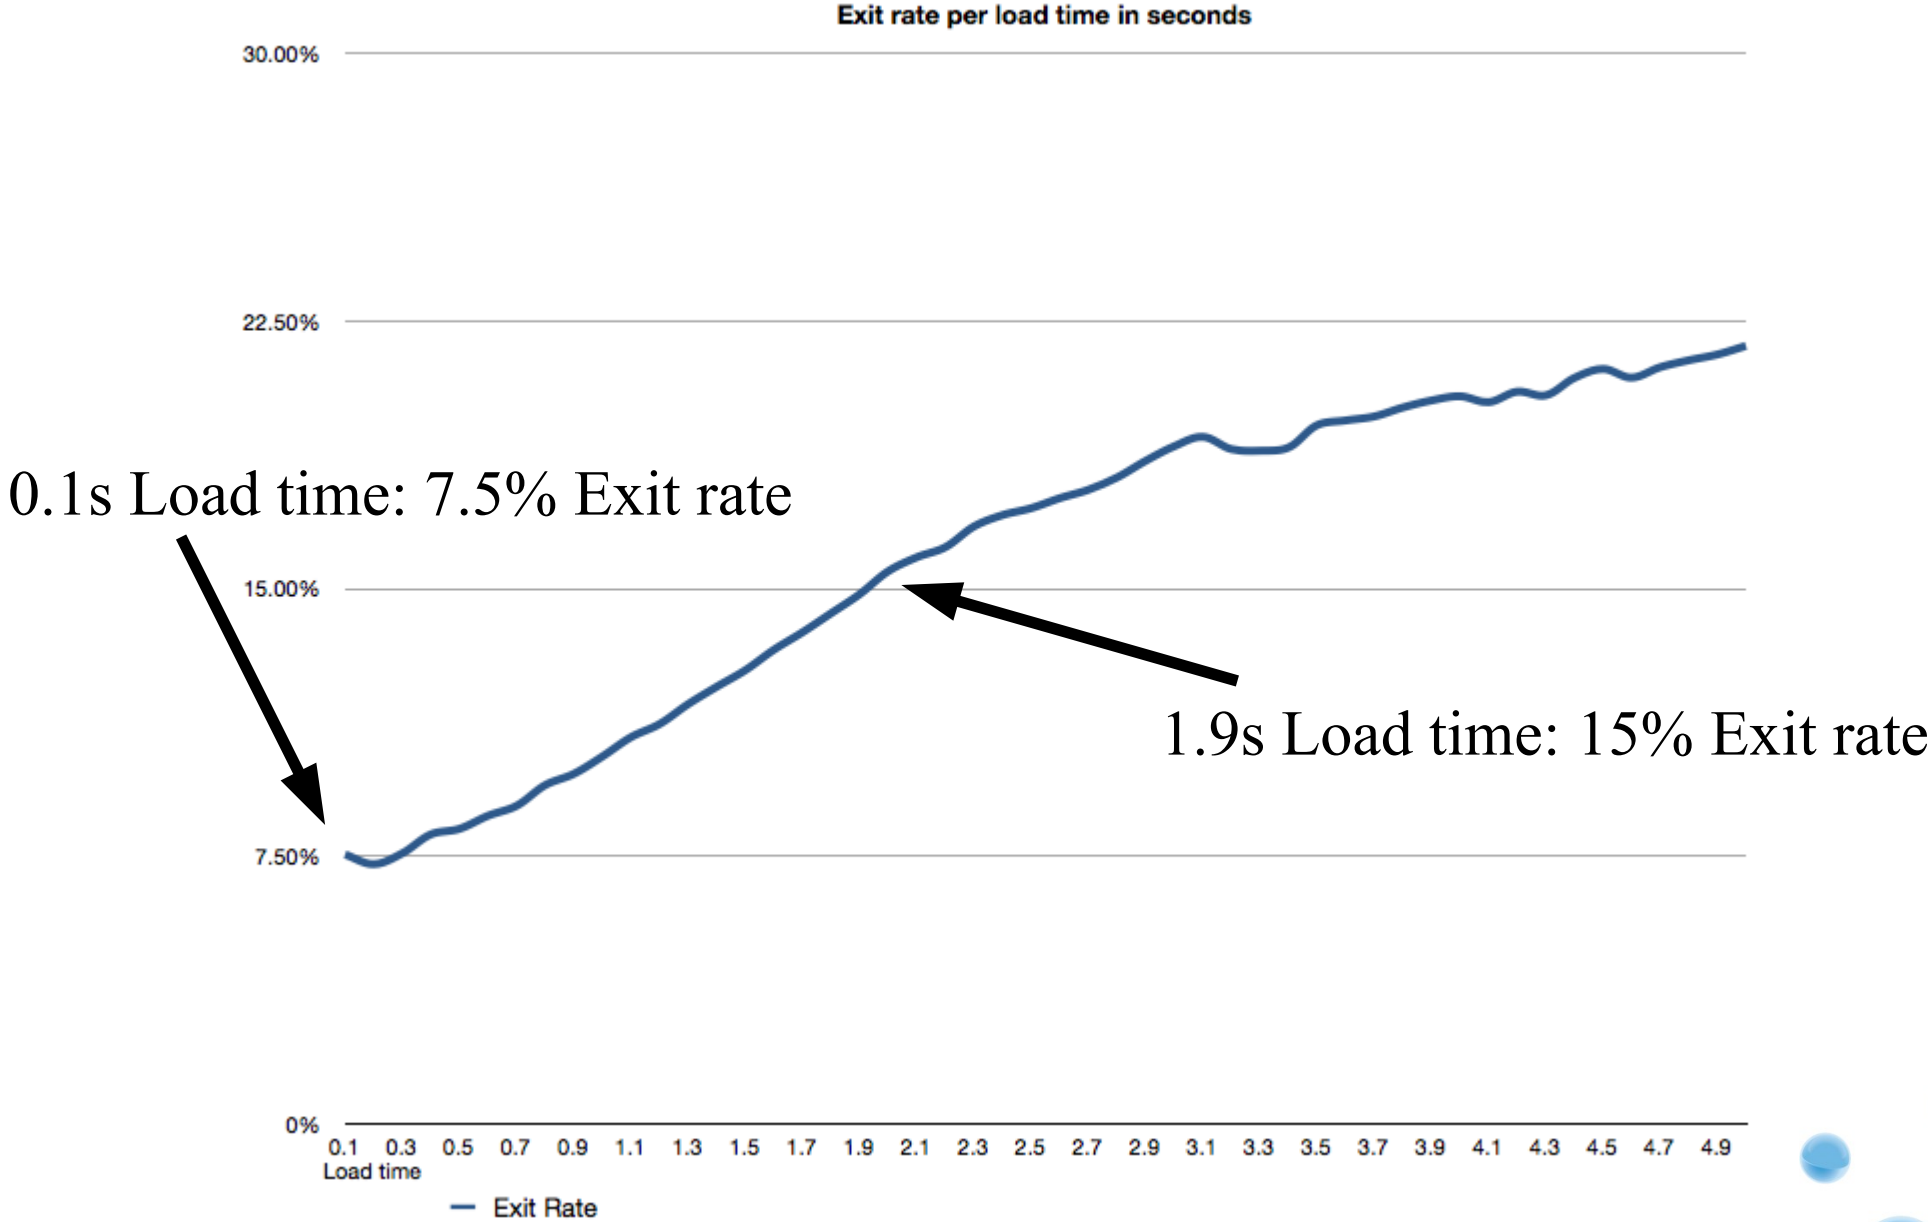
Response time?




Exit rate: How many of your users did not continue browsing the site after this specific page.




Exit rate per load time in seconds


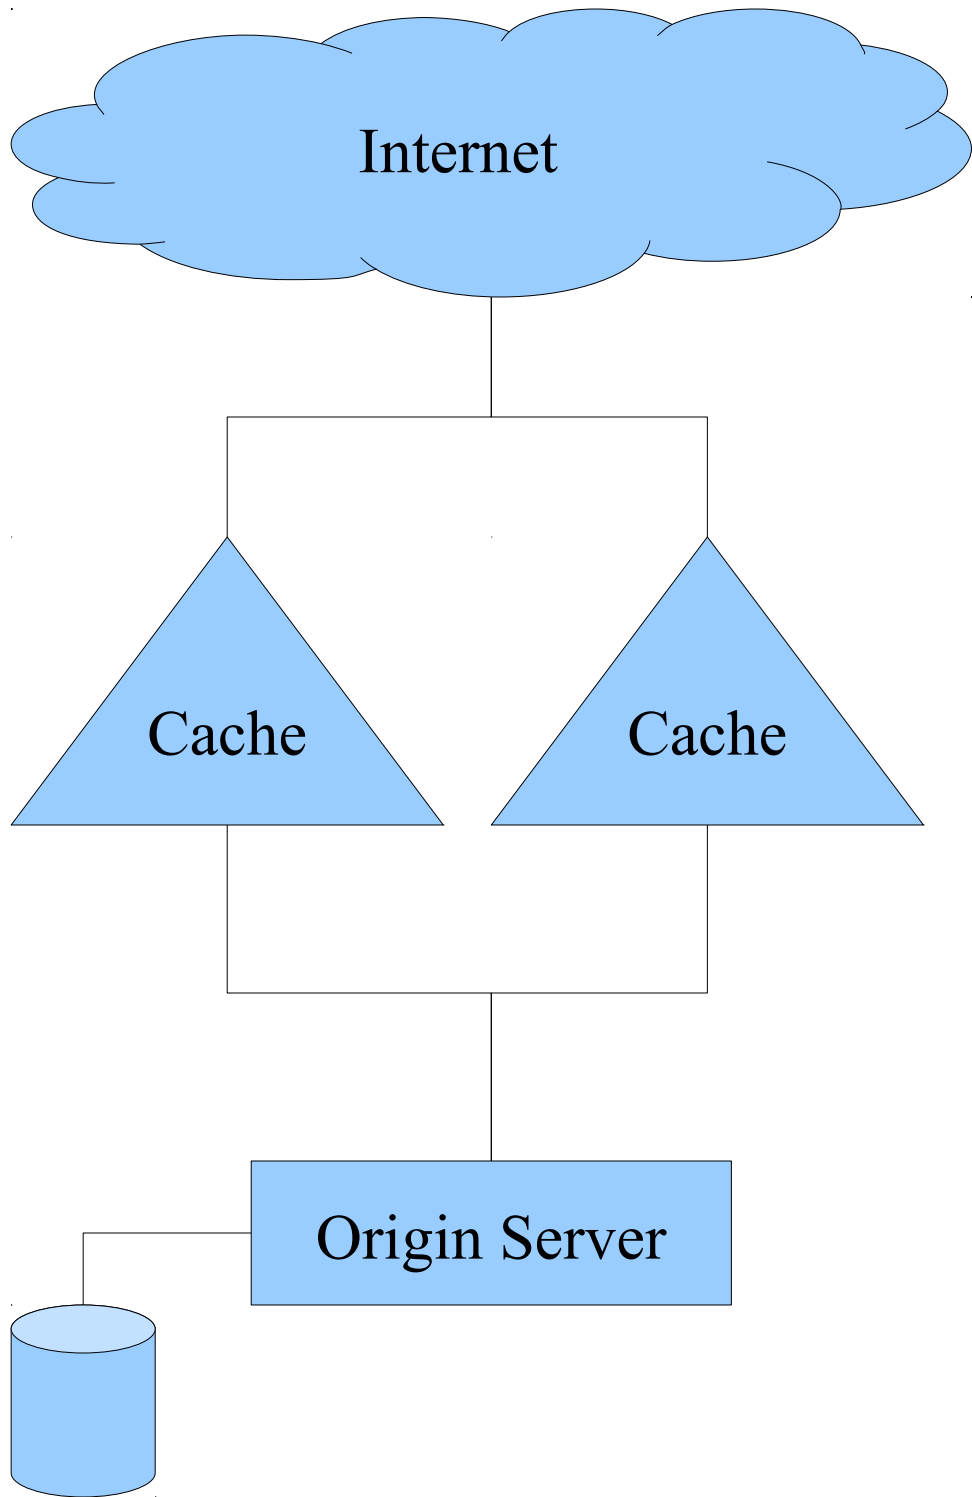


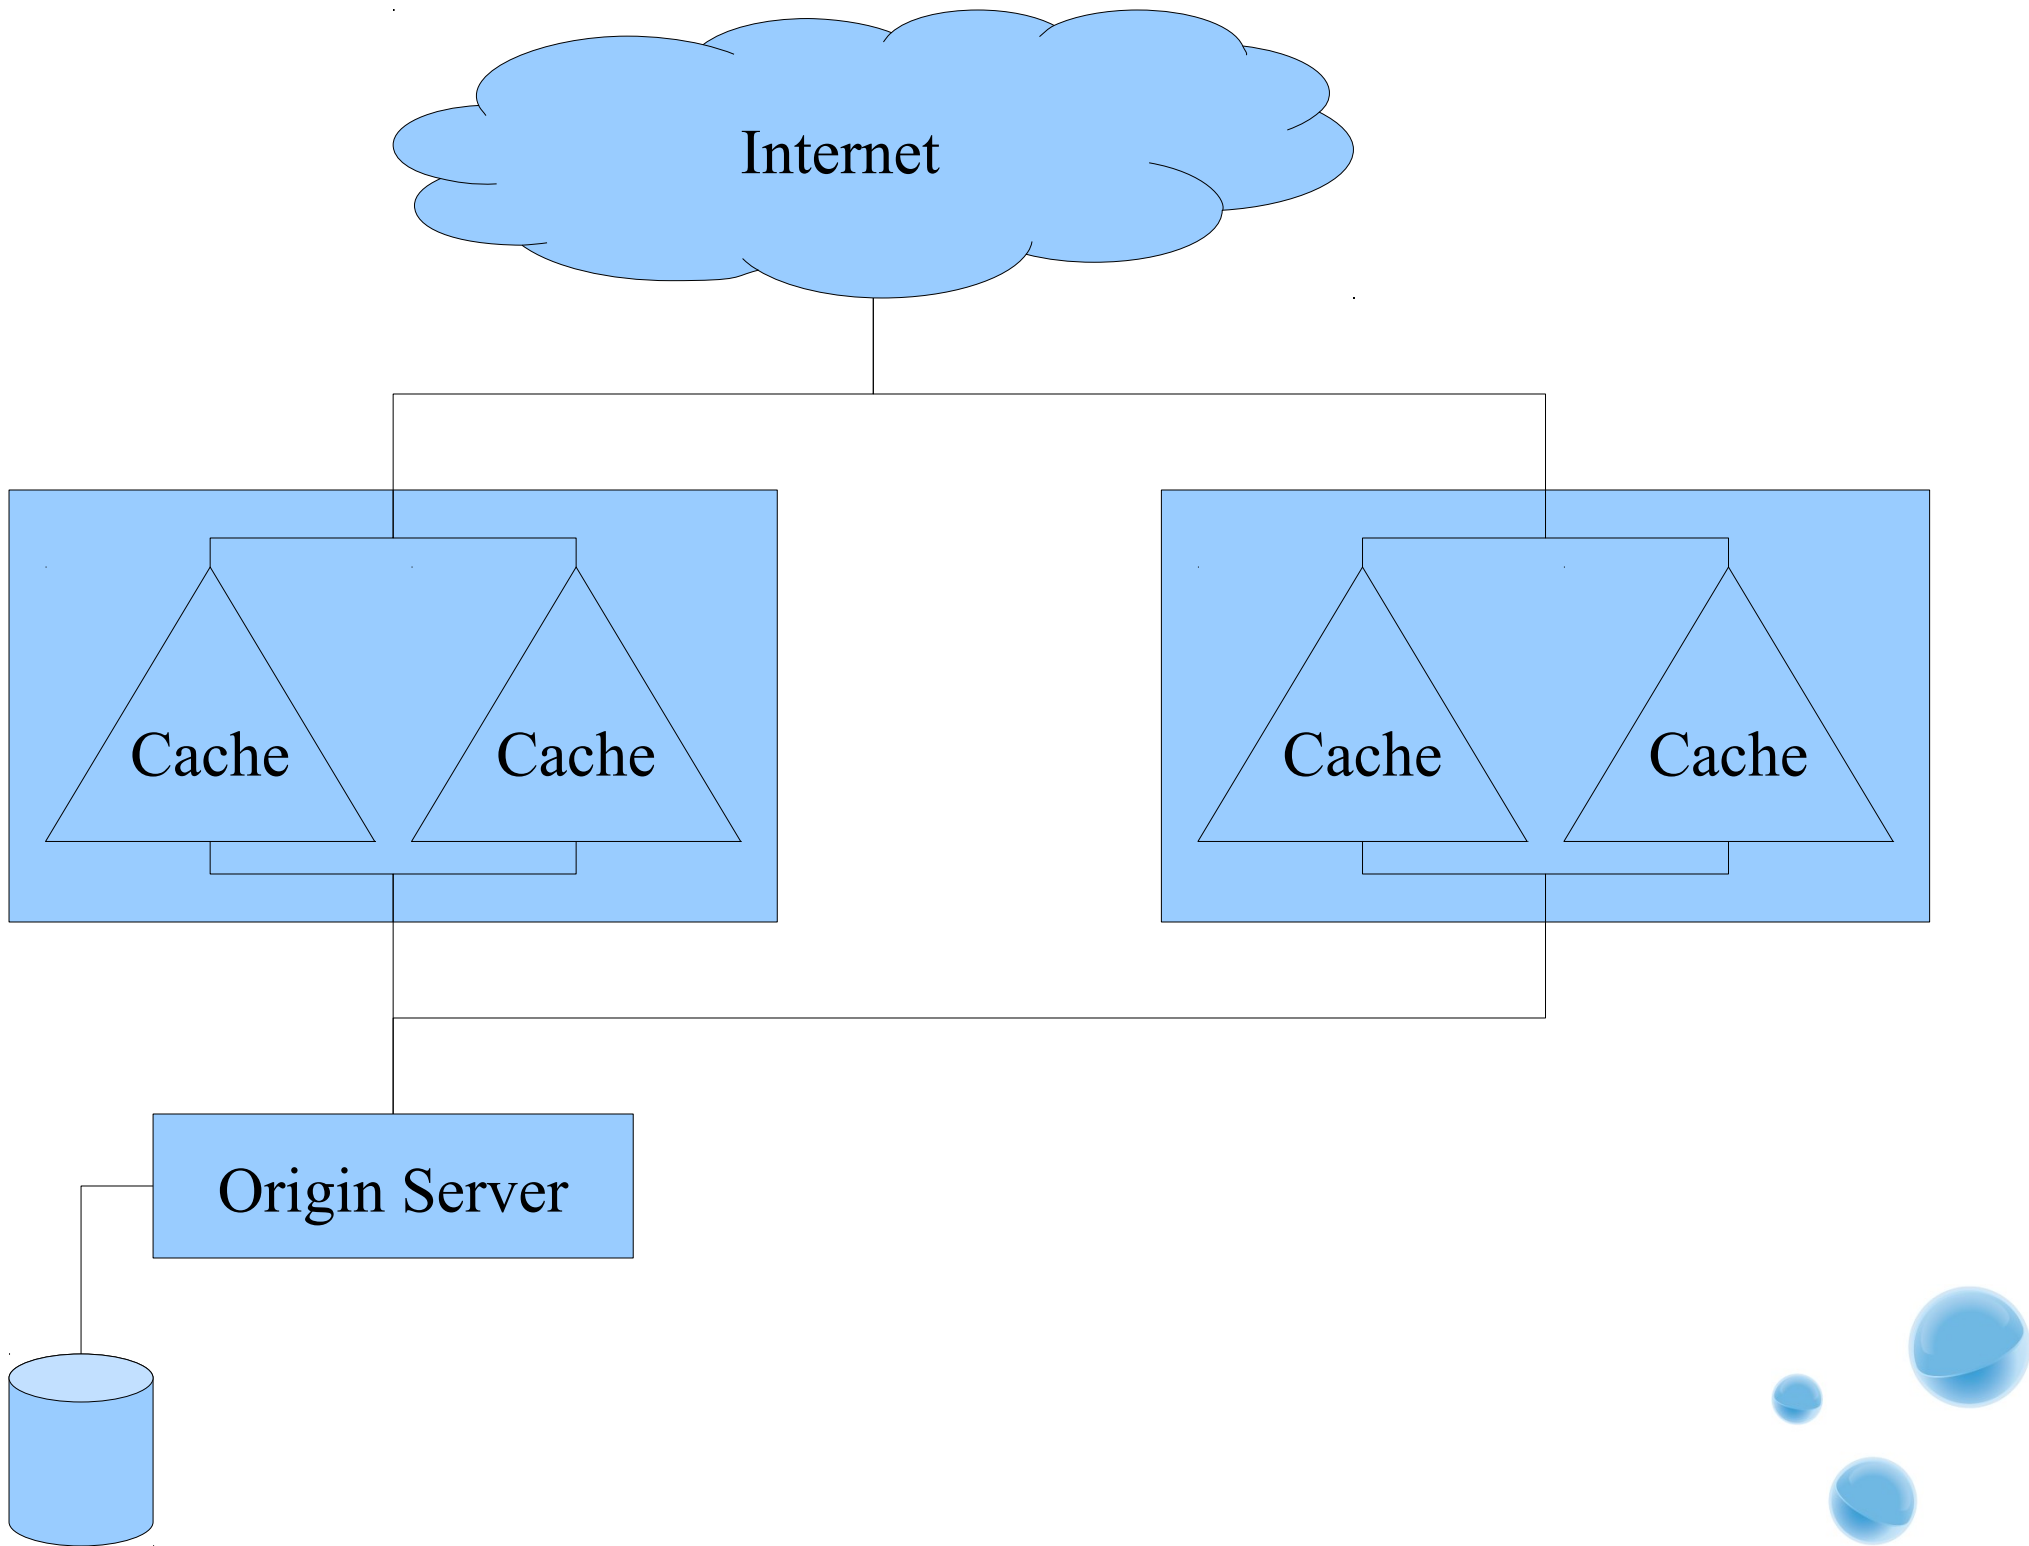
Courtesy of Artur Bergman of Fastly Inc.




Expansion?









How unique is an asset?




How often is an asset requested?



How often does an asset change?




How expensive is it to re-generate an asset?



Easy mode: Static Content


Easy to cache, but relatively small benefit



Anonymous users only


Caching front pages, articles, etc. Fairly easy and a significant benefit.

This is where most people stop.



Caching generic content


Cache even for logged in users, but only content which isn't user specific. Also fairly easy.



Caching content for logged in users

Hard, often not terribly beneficial, requires a larger cache.

BUT POSSIBLE.




Caching content that “changes all the time”

No such thing.

Live streamed video? Cache for 0.5 seconds and see a benefit.


Comments? Relatively easy to cache. Doesn't change very often.

Etc, etc.

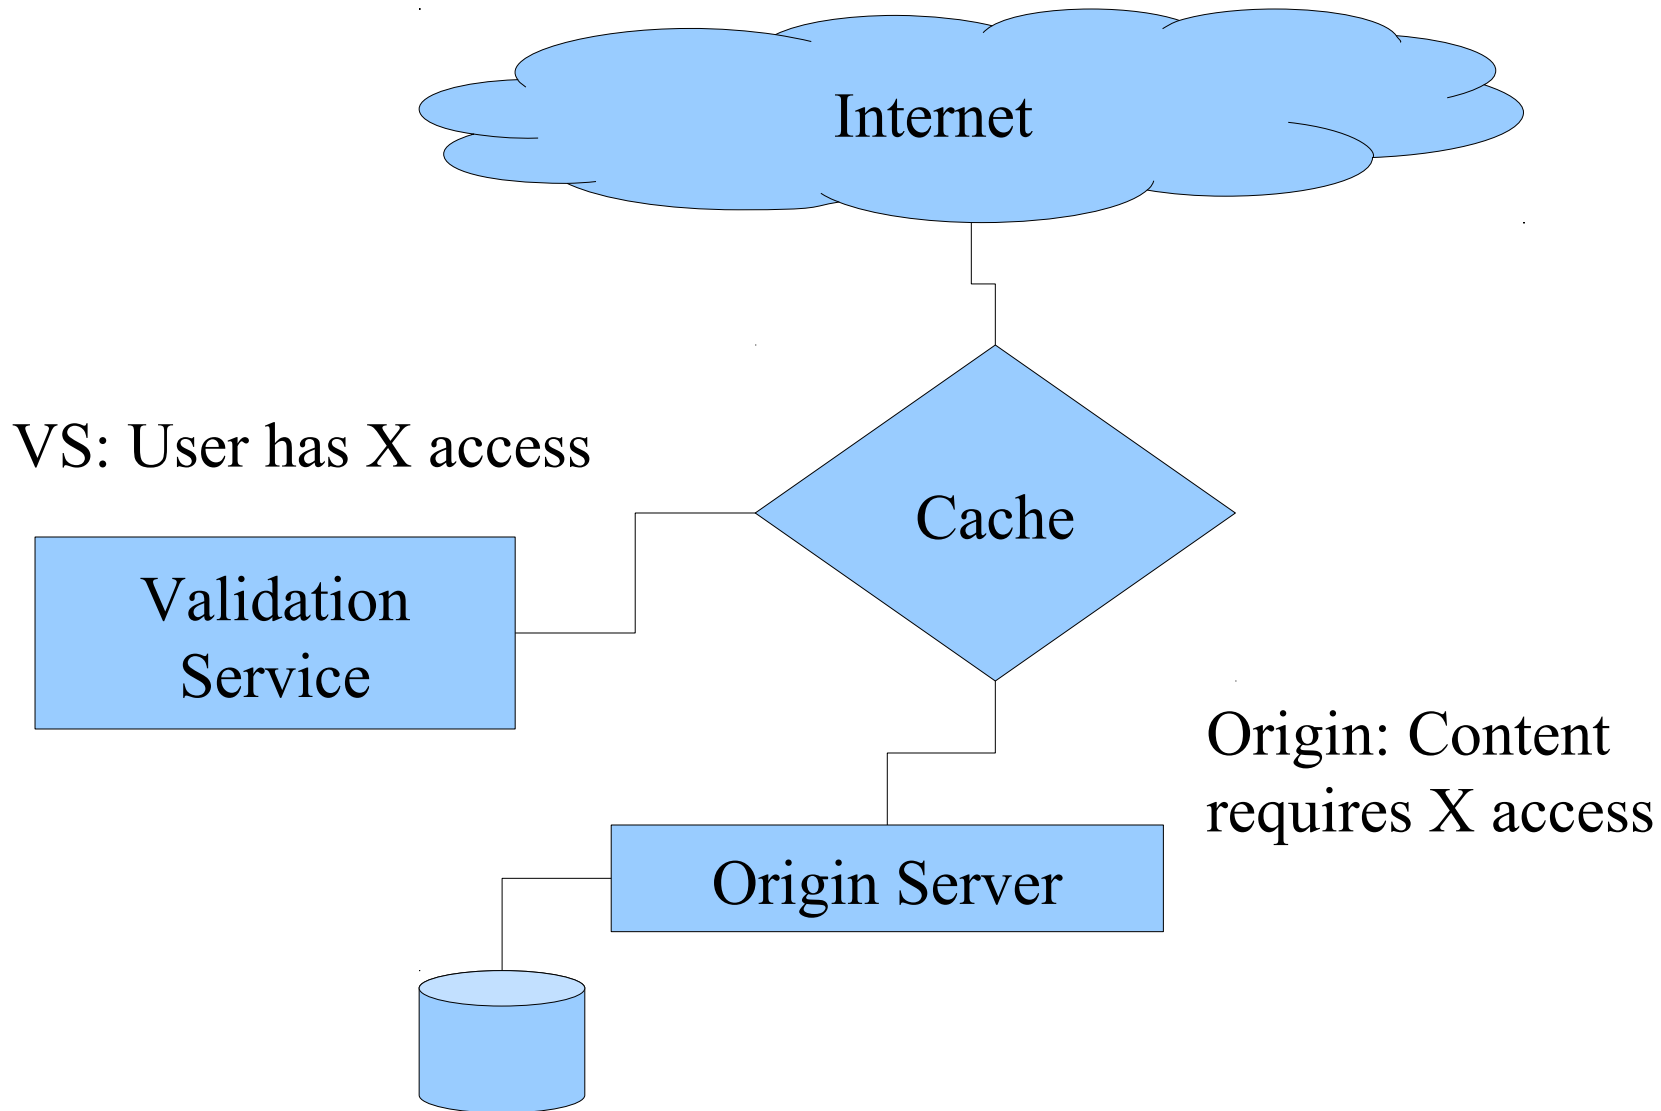


Example: Comment “counters”

- A setups with 20-30 news sites/magazines
- All sites supported comments
- All front pages displayed “X comments on this article”
- These counters were all in a single file – one massive database query.
- “Can't be cached.”
- 0 cache: Site went down.
- 5 second cache: No load at all.




Example: Paywall



Example: Paywall explained


1. User asks for content
2. Cache asks for content
3. Origin server responds with content and “Content requires permission X”
4. Cache asks validation service if user has access
5. Validation service answers with “yes”




Example: Paywall magic

6. Cache creates a cryptographically signed cookie for this user.


Next time the user asks for content, the cache only has to verify the signed cookie, then deliver from cache. No validation service request required.



ESI – Edge Side Includes
(Or: Client side AJAX if you prefer)



```
<title>Hello ESI-World</title>  
<esi:include src="/header">  
<esi:include src="/login-box">  
Content!  
<esi:include src="/user-survey">  
<esi:include src="/footer">
```



Horrible Designed News Site Blah

Welcome, Foobar (logout)


Article
snippet 1

Article
snippet 2

Article
snippet 3

Article
snippet 4

Top 10 popular news items



TTL: 10minutes

TTL: 60min, user specific


TTL: 30min

TTL: 30min

TTL: 30min

TTL: 30min

TTL: 1min



ESI top page

localhost/esi-demo.php

Mon, 19 Nov 2012 11:07:04 +0100

Mon, 19 Nov 2012 11:19:03 +0100

Welcome anonymous user. Who are you?

User page ▾ Purge

ESI top page


localhost/esi-demo.php

Mon, 19 Nov 2012 11:07:04 +0100

Mon, 19 Nov 2012 11:19:03 +0100

Welcome anonymous user. Who are you?

User page ▾ Purge



ESI top page

localhost/esi-demo.php

Mon, 19 Nov 2012 11:07:04 +0100

Mon, 19 Nov 2012 11:19:02 +0100

Welcome "User One".

Log out

User page ▾ Purge

ESI top page

localhost/esi-demo.php


Mon, 19 Nov 2012 11:07:04 +0100

Mon, 19 Nov 2012 11:19:03 +0100

Welcome anonymous user. Who are you?

username Log in

User page ▾ Purge



ESI top page

localhost/esi-demo.php

Mon, 19 Nov 2012 11:07:04 +0100

Mon, 19 Nov 2012 11:19:02 +0100

Welcome "User One".

Log out

User page ▾ Purge

ESI top page

localhost/esi-demo.php


Mon, 19 Nov 2012 11:07:04 +0100

Mon, 19 Nov 2012 11:20:28 +0100

Welcome "Another User".

Log out

User page ▾ Purge



ESI top page

localhost/esi-demo.php

Mon, 19 Nov 2012 11:07:04 +0100

Mon, 19 Nov 2012 11:19:02 +0100

Welcome "User One".

Log out

User page ▾ Purge

ESI top page

localhost/esi-demo.php


Mon, 19 Nov 2012 11:07:04 +0100

Mon, 19 Nov 2012 11:20:28 +0100

Welcome "Another User".

Log out

Top page ▾ Purge



ESI top page

localhost/esi-demo.php

Mon, 19 Nov 2012 11:07:04 +0100

Mon, 19 Nov 2012 11:19:02 +0100

Welcome "User One".

Log out

User page ▾ Purge

localhost/purgearticle.php

Updating the article in the database ...

Purging
http://localhost:80/esi-demo.php

---- Curl debug ----


* About to connect() to localhost port 80 (#0)

* Trying 127.0.0.1... *

connected

* Connected to localhost (127.0.0.1) port 80 (#0)

> PURGE /esi-demo.php HTTP/1.1



ESL top page

localhost/esi-demo.php

Mon, 19 Nov 2012 11:21:20 +0100

Mon, 19 Nov 2012 11:19:02 +0100

Welcome "User One".

Log out

User page ▾ Purge

ESL top page

localhost/esi-demo.php


Mon, 19 Nov 2012 11:21:20 +0100

Mon, 19 Nov 2012 11:20:28 +0100

Welcome "Another User".

Log out


User page ▾ Purge



ESI in Varnish


– `<esi:include ...>` `<esi:remove ...>`

```
sub vcl_fetch {  
    if (req.url ~ "/esi-urls") {  
        set beresp.do_esi = true;  
    }  
}
```




ESI alternatives

- Not loving the use of ESI?
- Use *AJAX* or other magic instead, does the same thing, but on the client side
- Different caching contexts need to be accessed through different URLs.




Where to start?



Caching Strategy

1. Low hanging fruit.
2. Heavy pages cached only for anonymous users.
3. Make a policy on cookies.
4. Cache multiple variants of a page.
5. Cache different elements on a single page.




Cool tricks time!

“Cache this for a long time, but only if it's popular!”


```
sub vcl_fetch {
    set beresp.ttl = 10m;
}

sub vcl_hit {
    if (obj.hits == 500) {
        set obj.ttl = 2d;
    }
}
```



Cookies, only for logged in people

```
sub vcl_recv {  
    if (req.http.cookie !~  
        "logged_in=true") {  
        unset req.http.cookie;  
    }  
}
```




Quick Start?

The Varnish Book!

<http://varnish-software.com/book>

Check out the Varnish Agent 2 for a single tenant REST interface/WebUI for Varnish.

<http://github.com/varnish/vagent2>



Questions?



Kristian Lyngstøl
kristian@bohemians.org
@kristianlyng
<http://kly.no/>
<https://joind.in/7960>

