

wikia

Artur Bergman
Surge 2010

Artur Bergman

- CTO Wikia
 - Operations
 - Long term technology
 - Performance
- Old time perl core hacker (perl threading is my fault)
- Varnish contributor
- Against network neutrality

Not afraid to say the
truth

artur@wikia-inc.com
@crucially

Wikia

- Hosts 150 000 wikis
- Mediawiki
- 5 million content articles
- ~ 950 mm pageviews / month
- 36 mm unique users

Mediawiki

- Legacy code base
- Can't change it too much
- More like an enterprise shop in some regards due to this limitation

World of Warcraft univ guide - WoWiki

http://www.wowwiki.com/Portal_Main?noitemals=1

Wikia ENTERTAINMENT

Create a new wiki

Connect Log In Create an account

Edit this page History Follow Share this article Portal Discussion

beta is running. Post away! 3.3.5: Ruby Sanctum Hate ads? Log in, it's free! Report bad ads

Welcome to WoWiki
the free Warcraft universe info source anyone can edit

WoWiki content portals

World of Warcraft Wrath of the Lich King Cataclysm

WARCRAFT UNIVERSE

World of Warcraft info WoWiki community

Selected links from across WoWiki

World of Warcraft (WoW)

- Classes
- Factions
- Guides
- In-game events
- Instances by level
- Mounts
- Patches
- Races

Wrath of the Lich King

- Icecrown Citadel
- Argent Tournament
- Ornyla
- Ul'uar raid
- Naxxramas
- Death knights
- Northrend
- Patch 3.1.0, 3.2.0, 3.3.0

Warcraft universe

- Warcraft games
- Books
- Azerath
- Outland

Community

- Guides
- Feats
- WoWiki forums
- Other fan sites
- Other wikis

Featured article

The red dragonflight, ruled by the Dragonqueen, Alexstrasza, is a noble and honorable group of creatures. They consider themselves the protectors of all life, and in many ways they are. More often than not, they are found protecting sacred areas and items, seeking to keep them from lesser beings who might be hurt by their energies. Their breath is a fearsome stream of fire, and they have been known to swallow enemies whole and slowly digest them over the course of a day. Their most hated enemies are the black dragons ruled by the great, dark leviathan, Deathwing.

WoWiki stats

86,629 pages on this wiki

Tichondrius 13 minutes ago by Gruntjacked

Landou's Homed Greathelm 13 minutes ago by CravusArgule

Ace (AddOns) 13 minutes ago by Sycmaster

Aca2 15 minutes ago by Sycmaster

Crok Scourgebane 41 minutes ago by Kilgord

The Swedish Chef - Muppet Wiki

http://muppet.wikia.com/wiki/Swedish_Chef?noitemals=1

wikia Start a wiki Entertainment Gaming Lifestyle

Connect Log In Create an account

Muppet Wiki

Random Muppet character The Muppet Show Sesame Street Community 3CRandom Page Wiki Activity

The Swedish Chef Log in to edit

(Redirected from Swedish Chef)

Edited 19 hours ago by BradFraggle

Read more: Muppet Characters, The Muppets Characters

29 Search this wiki

22,229 PAGES ON THIS WIKI Create a Page

Episode 4215 2 hours ago by Oceanfan

Sign Language Moment of the Day 4 hours ago by Loriel-man

Episode 114: All Connected 4 hours ago by CDD32 See more >

61,385 PHOTOS ON THIS WIKI Add a Photo

Wikia Spotlights

Performer: Jim Henson

The Swedish Chef is the incomprehensible preparer of foodstuffs from *The Muppet Show*. A rather literal variation of the Live-Hand Muppet concept, the Swedish Chef is a humanoid character, with human hands rather than gloves. On *The Muppet Show*, Frank Oz performed the character's hands. The chef first appeared in *The Muppet Show: Sex and Violence* in 1975, with Chinese subtitles for his dialogue, and Richard Hunt performing the hands. Only in poster puppets used specifically for photo shoots does the Swedish Chef have standard Muppet hands.

According to Brian Henson in one of his introductions for *The Muppet Show*, "Jim Henson had this tape that he used to play which was "How to Speak Mock Swedish". And he used to drive to work and I used to ride with him a lot. And he would drive to work trying to make a chicken sandwich in Mock Swedish or make a turkey casserole in Mock Swedish. It was the most ridiculous thing you had ever seen. And people at traffic lights used to stop and sort of look at him a little crazy. But that was the roots of the character that would eventually become the Swedish Chef."

Nearly all Swedish Chef sketches begin with him in a kitchen.

Share Follow

True Blood Wiki

http://trueblood.wikia.com/wiki/True_Blood_Wiki?noitemals=1

wikia ENTERTAINMENT

Create a new wiki

Connect Log In Create an account

View source History Follow Share Article Discussion

TRUE BLOOD

Home Page True Blood Top Content Characters Character Galleries Cast Cast Galleries Locations Episodes Seasons TB Top 10 Lists Random Article Community Take A Bite I Wiki Leaderboard

Upload a new photo

Help Recent changes Multiple upload Special pages

Image Showcase

facebook twitter

True Blood Wiki on Facebook

312 people like True Blood Wiki

EPISODES CHARACTERS ACTORS

About True Blood Wiki

True Blood is a collaborative Wiki dedicated to the HBO series "True Blood" created by Alan Ball, based on Charlaine Harris's *Southern Vampire Mysteries*. Follow us on Facebook and Twitter! We don't bite...much.

Season 3 Finale

True Blood Season 3 Episode 12 "Evil Is Going O..."

Featured image

King of the Hill

Wookieepedia, the Star Wars Wiki

http://starwars.wikia.com/wiki/Main_Page?noitemals=1

wikia ENTERTAINMENT

Create a new wiki

Connect Log In Create an account

View source History Follow Share this article Main Page Discussion

WOOKEEPEDIA

The Star Wars encyclopedia that anyone can edit.

Warning: this wiki contains spoilers.

80,136 articles since March 4, 2005

click for article navigation

overview index categories browse discussion faq help things to do article requests languages copyrights vip interviews

Quote of the Day

"Obviously, the bad guys have more fashion sense in the Star Wars universe, and Lando didn't have enough screen time to, well, bring Balance to the Fashion."
--Aaron Alston

Featured article

The Battle of Koros Major was a clash fought during the Great Hyperspace War between the army of the united Koros system and the forces of the Sith Empire. In 5,000 BBY, Dark Lord of the Sith Naga Sadow gathered an invasion force and launched an attack on the Galactic Republic. A part of the Sith fleet led by Sadow's Sith apprentice and Koros Major native Gav Daragon, was sent to conquer the Koros system in a two-pronged attack focused on both his homeworld and the planet Kirok. The Koros military, commanded by Empress Teta and Vantage, one of her top advisors and generals, had a diminished fleet due to a recent war but nevertheless believed that they could provide a good accounting of themselves. During the planet-side engagement on Koros Major, Gav Daragon took a dropship down to the main area of battle in the city of Cinnagar to find his sister, Jori, whom he had not seen in a long time. While looking for her with his company of Massassi guards, Daragon had an altercation with a former acquaintance, Aarba Hutt, which resulted in the Hutt's death. Jori Daragon arrived at that moment and confronted her brother about the Hutt's death. Overcome with emotion, Gav fled and abandoned the battle soon afterward, deciding to attack Sadow and end the war. With Gav Daragon no longer present to direct them, the Sith lost the battle. (Read more...)

Recently featured: Gav Daragon - Blockade of the Hydian Way - Bulter Swan

Archive - Good articles - How can I write a Featured article?

Random articles

Random Featured Article: Bertoff Heiss

Random Good Article: Artrix Day group

In the news

- CC7567 becomes Wookieepedia's 43rd administrator.
- The computer game *Star Wars: Clone Wars Adventures* enters open beta testing.
- A new *Wrath Squadron* novel by Aaron Alston, *The Old Republic* novel by Drew Karpyshyn, and update of *The Essential Guide to Characters* by Daniel Wallace are announced at Celebration V.
- Crimson Empire III* is set for release in March 2011.
- Celebration V begins in Orlando, Florida.
- Darth Plagueis by James Luceno is announced for a 2012 release, after its prior cancellation in 2007.

Wookieepedia NewsNet

- "Floyd" newest Wookieepedia administrator (07 Aug)
- Grand Moff Tranner named Wookieepedia administrator (19 Apr)

On this day...

September 30:

- 1907 - Filming for *The Phantom Menace* finishes.

Wikia

- 1000 cores
- 5 data centers
- 3 Gbps
- 3.5 ops guys

Metrics

- Google Analytics
- Pingdom
- Browsermob
- Websitepulse

Metrics

- Love em
- Too cheap to spend too much time
- Abuse tools
- Google Analytics to the rescue
- So these are what our numbers are based on

End user performance

```
<script type="text/javascript">/*! [CDATA[*/  
var wgNow = new Date();  
/*]]>*/</script>
```


```
<!-- content -->
```

```
if (typeof window.wgNow == "object"){  
    var now = new Date();  
    var ms = (now.getTime() - window.wgNow.getTime()) / 1000;  
        var pageTime = Math.floor(ms * 10) / 10; // Round to 1  
decimal  
    var slashtime = "/lean_monaco/" + pageTime.toString().replace(/\./,  
"/");  
    _uff=0;  
    _uacct="UA-XXXX-XX";  
    urchinTracker(slashtime);  
}
```









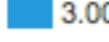

Page	↑ <input type="button" value="None"/> ↓	Pageviews	Unique Pageviews	Avg. Time on Page	Bounce Rate	% Exit
/lean_monaco/0				00:00:48	35.48%	8.80%
/lean_monaco/0/1				00:00:56	36.62%	9.58%
/lean_monaco/0/2				00:01:04	38.76%	11.03%
/lean_monaco/0/3				00:01:11	39.94%	12.60%
/lean_monaco/0/4				00:01:16	40.36%	13.93%
/lean_monaco/0/5				00:01:19	40.26%	15.06%
/lean_monaco/0/6				00:01:22	39.89%	16.02%
/lean_monaco/0/7				00:01:24	39.57%	16.79%
/lean_monaco/0/8				00:01:26	39.18%	17.50%
/lean_monaco/0/9				00:01:28	38.87%	18.00%
/lean_monaco/1				00:01:29	38.81%	18.50%
/lean_monaco/1/1				00:01:31	38.58%	18.91%
/lean_monaco/1/2				00:01:32	38.49%	19.29%
/lean_monaco/1/3				00:01:33	38.32%	19.65%
/lean_monaco/1/4				00:01:35	38.24%	19.95%
/lean_monaco/1/5				00:01:35	38.14%	20.16%
/lean_monaco/1/6				00:01:37	38.05%	20.43%
/lean_monaco/1/7				00:01:38	37.94%	20.61%
/lean_monaco/1/8				00:01:39	37.92%	20.86%
/lean_monaco/1/9				00:01:40	37.87%	21.04%
/lean_monaco/10				00:02:14	44.10%	27.92%
/lean_monaco/11				00:02:18	44.28%	28.30%
/lean_monaco/12				00:02:17	45.01%	28.88%
/lean_monaco/13				00:02:19	45.43%	29.08%
/lean_monaco/14				00:02:23	46.24%	29.26%


Page	↑ <input type="button" value="None"/> ↓	Pageviews	Unique Pageviews	Avg. Time on Page	Bounce Rate	% Exit
/lean_monaco/0				00:00:48	35.48%	8.80%
/lean_monaco/0/1				00:00:56	36.62%	9.58%
/lean_monaco/0/2				00:01:04	38.76%	11.03%
/lean_monaco/0/3				00:01:11	39.94%	12.60%
/lean_monaco/0/4				00:01:16	40.36%	13.93%
/lean_monaco/0/5				00:01:19	40.26%	15.06%
/lean_monaco/0/6				00:01:22	39.89%	16.02%
/lean_monaco/0/7				00:01:24	39.57%	16.79%
/lean_monaco/0/8				00:01:26	39.18%	17.50%
/lean_monaco/0/9				00:01:28	38.87%	18.00%
/lean_monaco/1				00:01:29	38.81%	18.50%
/lean_monaco/1/1				00:01:31	38.58%	18.91%
/lean_monaco/1/2				00:01:32	38.49%	19.29%
/lean_monaco/1/3				00:01:33	38.32%	19.65%
/lean_monaco/1/4				00:01:35	38.24%	19.95%
/lean_monaco/1/5				00:01:35	38.14%	20.16%
/lean_monaco/1/6				00:01:37	38.05%	20.43%
/lean_monaco/1/7				00:01:38	37.94%	20.61%
/lean_monaco/1/8				00:01:39	37.92%	20.86%
/lean_monaco/1/9				00:01:40	37.87%	21.04%
/lean_monaco/10				00:02:14	44.10%	27.92%
/lean_monaco/11				00:02:18	44.28%	28.30%
/lean_monaco/12				00:02:17	45.01%	28.88%
/lean_monaco/13				00:02:19	45.43%	29.08%
/lean_monaco/14				00:02:23	46.24%	29.26%

Exit rate per load time in seconds

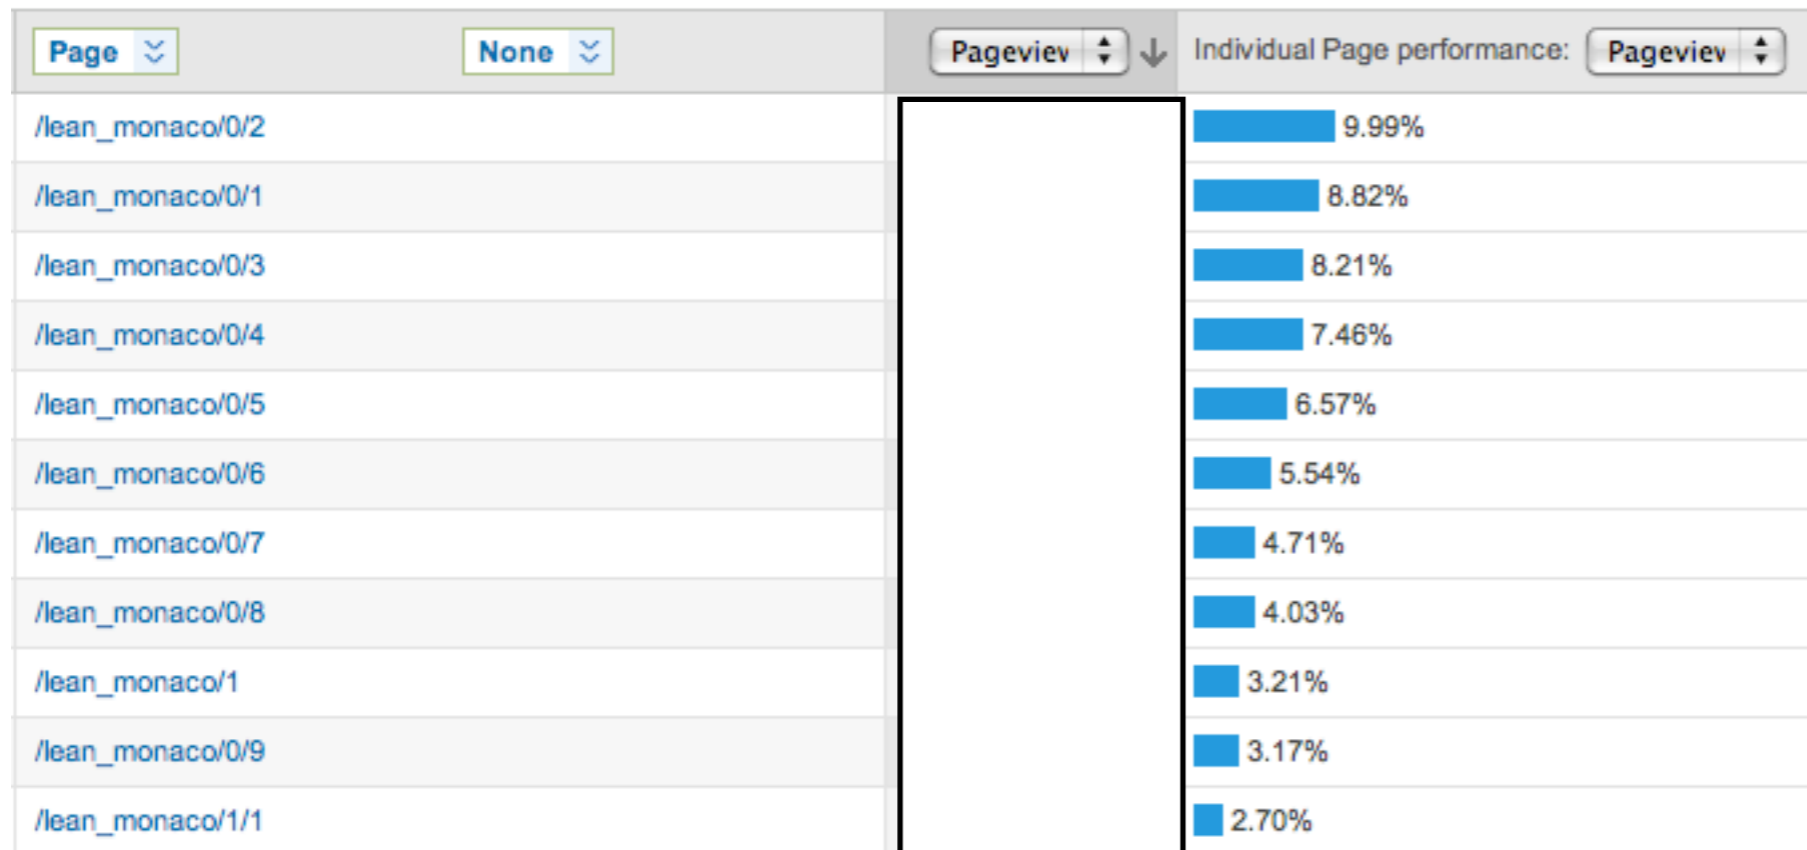


Distribution of load time

	Page None	Pageview ↓	Individual Page performance: Pageview
1.	/lean_monaco/0/1		 17.67%
2.	/lean_monaco/0/2		 14.17%
3.	/lean_monaco/0/3		 9.71%
4.	/lean_monaco/0		 8.34%
5.	/lean_monaco/0/4		 7.44%
6.	/lean_monaco/0/5		 5.81%
7.	/lean_monaco/0/6		 4.54%
8.	/lean_monaco/0/7		 3.78%
9.	/lean_monaco/0/8		 3.00%
10.	/lean_monaco/0/9		 2.52%



Distribution California



Distribution Brasil

Page	Browser	Individual Page pe Pageview
/lean_monaco/0/1	Internet Explorer	15.66%
/lean_monaco/0/2	Internet Explorer	14.48%
/lean_monaco/0/3	Internet Explorer	10.27%
/lean_monaco/0/4	Internet Explorer	8.10%
/lean_monaco/0/5	Internet Explorer	6.41%
/lean_monaco/0/6	Internet Explorer	4.85%
/lean_monaco/0/7	Internet Explorer	4.37%
/lean_monaco/0/8	Internet Explorer	3.28%
/lean_monaco/0/9	Internet Explorer	2.80%
/lean_monaco/0	Internet Explorer	2.79%

Page	Browser	Individual Page pe Pageview
/lean_monaco/0/1	Chrome	21.28%
/lean_monaco/0	Chrome	19.02%
/lean_monaco/0/2	Chrome	12.79%
/lean_monaco/0/3	Chrome	9.00%
/lean_monaco/0/4	Chrome	6.73%
/lean_monaco/0/5	Chrome	4.95%
/lean_monaco/0/6	Chrome	3.80%
/lean_monaco/0/7	Chrome	2.96%
/lean_monaco/0/8	Chrome	2.31%
/lean_monaco/0/9	Chrome	1.91%

Page	Browser	Individual Page performan Pageview
/lean_monaco/0/1	Firefox	17.81%
/lean_monaco/0/2	Firefox	15.52%
/lean_monaco/0/3	Firefox	10.27%
/lean_monaco/0/4	Firefox	7.87%
/lean_monaco/0/5	Firefox	6.15%
/lean_monaco/0/6	Firefox	5.02%
/lean_monaco/0	Firefox	4.60%
/lean_monaco/0/7	Firefox	4.02%
/lean_monaco/0/8	Firefox	3.28%
/lean_monaco/0/9	Firefox	2.71%

Page	Browser	Individual Page performance: Pageview
/lean_monaco/0	Safari	26.08%
/lean_monaco/0/1	Safari	18.99%
/lean_monaco/0/2	Safari	10.89%
/lean_monaco/0/3	Safari	6.28%
/lean_monaco/0/4	Safari	4.60%
/lean_monaco/0/5	Safari	3.32%
/lean_monaco/0/6	Safari	2.75%
/lean_monaco/0/7	Safari	2.34%
/lean_monaco/0/8	Safari	1.91%
/lean_monaco/0/9	Safari	1.67%

GA

- Per country
- Per browser
- Per ISP
- First pageview (uncached)
- For free
- Just do it


Errors

- Track all user errors
- Whenever they see one
- Too cheap for servers
- Google analytics to the rescue

```
sub vcl_error {
    synthetic {"
<html>
<head>
<title> "} req.http.host req.url {"</title>
<script src="http://www.google-analytics.com/urchin.js" type="text/
javascript">
    </script>
    <script type="text/javascript">
        _uacct = "UA-XXXXXX-YY";
        urchinTracker("/errmsg/") server.identity {"/" } obj.status {"/" }
req.backend {"");
    </script>
        </head><body></body></html>
    "};
    return(deliver);
}
```

```
sub vcl_error {
    synthetic {"
<html>
<head>
<title> "} req.http.host req.url {"</title>
<script src="http://www.google-analytics.com/urchin.js" type="text/
javascript">
    </script>
    <script type="text/javascript">
        _uacct = "UA-XXXXXX-YY";
        urchinTracker("/errmsg/") server.identity {"/" } obj.status {"/" }
req.backend {"");
    </script>
        </head><body></body></html>
    "};
    return(deliver);
}
```

Errors served YTD



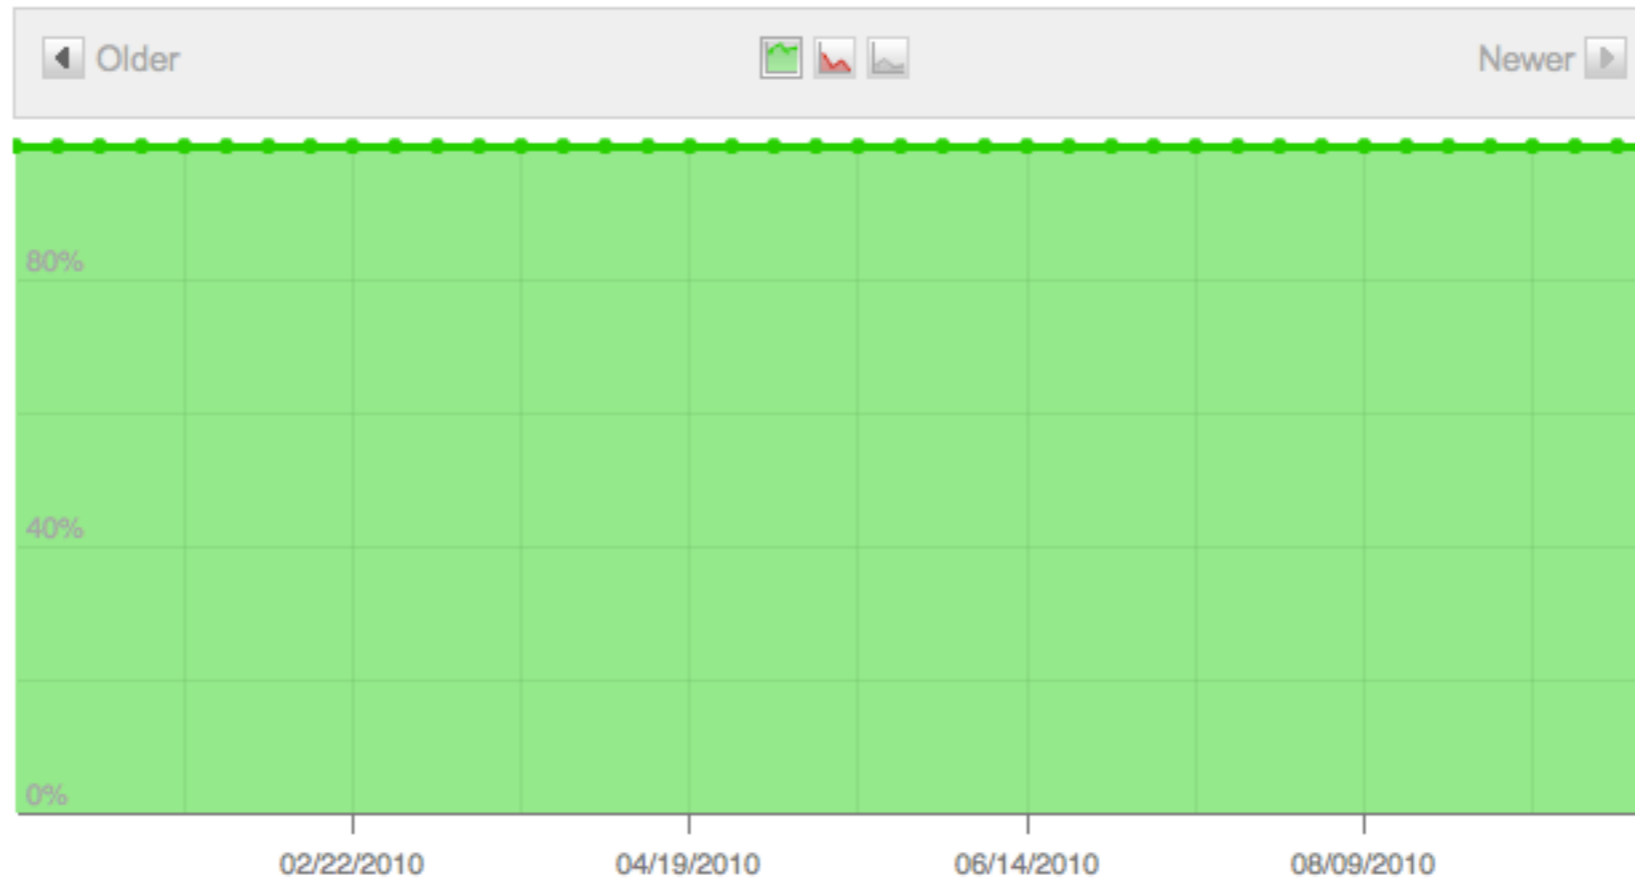
275 pages were viewed a total of 1,173,697 times
 Filtered for pages excluding "(spotlights|liftium)"

Content Performance		Views: Table Chart Grid Filter Settings					
Pageviews 1,173,697 % of Site Total: 5.05%	Unique Pageviews 676,080 % of Site Total: 3.19%	Avg. Time on Page 00:01:00 Site Avg: 00:00:47 (28.32%)	Bounce Rate 55.84% Site Avg: 20.77% (168.83%)	% Exit 40.47% Site Avg: 10.52% (284.57%)	\$ Index \$0.00 Site Avg: \$0.00 (0.00%)		
Page	Pageviews ↓	Unique Pageviews	Avg. Time on Page	Bounce Rate	% Exit	\$ Index	
1. /lvs/sjc/varnish	230,733	82,437	00:01:04	51.67%	35.32%	\$0.00	
2. /varnish/varnish8-LON/503/sjc_varnish	196,843	92,496	00:00:26	22.67%	12.69%	\$0.00	
3. /varnish/varnish9/503	80,916	36,356	00:00:45	57.29%	44.41%	\$0.00	
4. /varnish/varnish-i6/503	66,891	34,892	00:00:19	65.68%	51.55%	\$0.00	
5. /varnish/varnish8-LON/503/iowa_apache	60,151	39,196	00:01:00	47.40%	37.79%	\$0.00	
6. /varnish/varnish-s1-SJC/503/iowa_apache	42,545	31,481	00:01:35	58.94%	49.90%	\$0.00	
7. /varnish/varnish-i7/503	39,963	16,834	00:00:19	51.91%	41.09%	\$0.00	
8. /varnish/varnish9/503/sjc_varnish	38,049	29,493	00:00:49	37.86%	32.04%	\$0.00	
9. /varnish/varnish-s1-SJC/500/iowa_apache	26,173	21,411	00:02:18	66.80%	56.55%	\$0.00	
10. /varnish/varnish-s3-SJC/500/iowa_apache	26,066	20,193	00:02:17	63.25%	50.59%	\$0.00	

Filter Page: (spotlights|liftium) [Advanced Filter](#)

Go to: Show rows: 1 - 10 of 275

ping wikia.com YTD




Uptime
100.00%

Downtime
2m
The average downtime length is 1m

Number of downtimes
2
The longest downtime was 1m on 01/06/2010 10:36:37PM and the shortest was 1m on 01/06/2010 10:36:37PM

http health check wikia.com YTD



Uptime
99.99%

Downtime
42m 54s

Number of downtimes
30

The average downtime length is 1m 25s

The longest downtime was 3m 38s on 08/16/2010 10:31:34PM and the shortest was 58s on 05/24/2010 1:26:36PM

Uptime

- Network - 99.995%
- HTTP - 99.989%
- No error delivered - 99.982%

How

- Ops driving engineering focus
- Geographically distributed (servers and team)
- Fail fast - revert fast
- Internal redundancy in code and infrastructure

Magic in a box



Varnish

- Our entire architecture revolves around it
- High speed HTTP cache and load balancer
- Also HTTP routing

Detour through hardware

- No virtualization
- No cloud
- Too slow
- Too expensive
- Too little control

Amazing control

- Spec your hardware to your app
- Write your app for the hardware
- Realise fantastic gains in efficiency

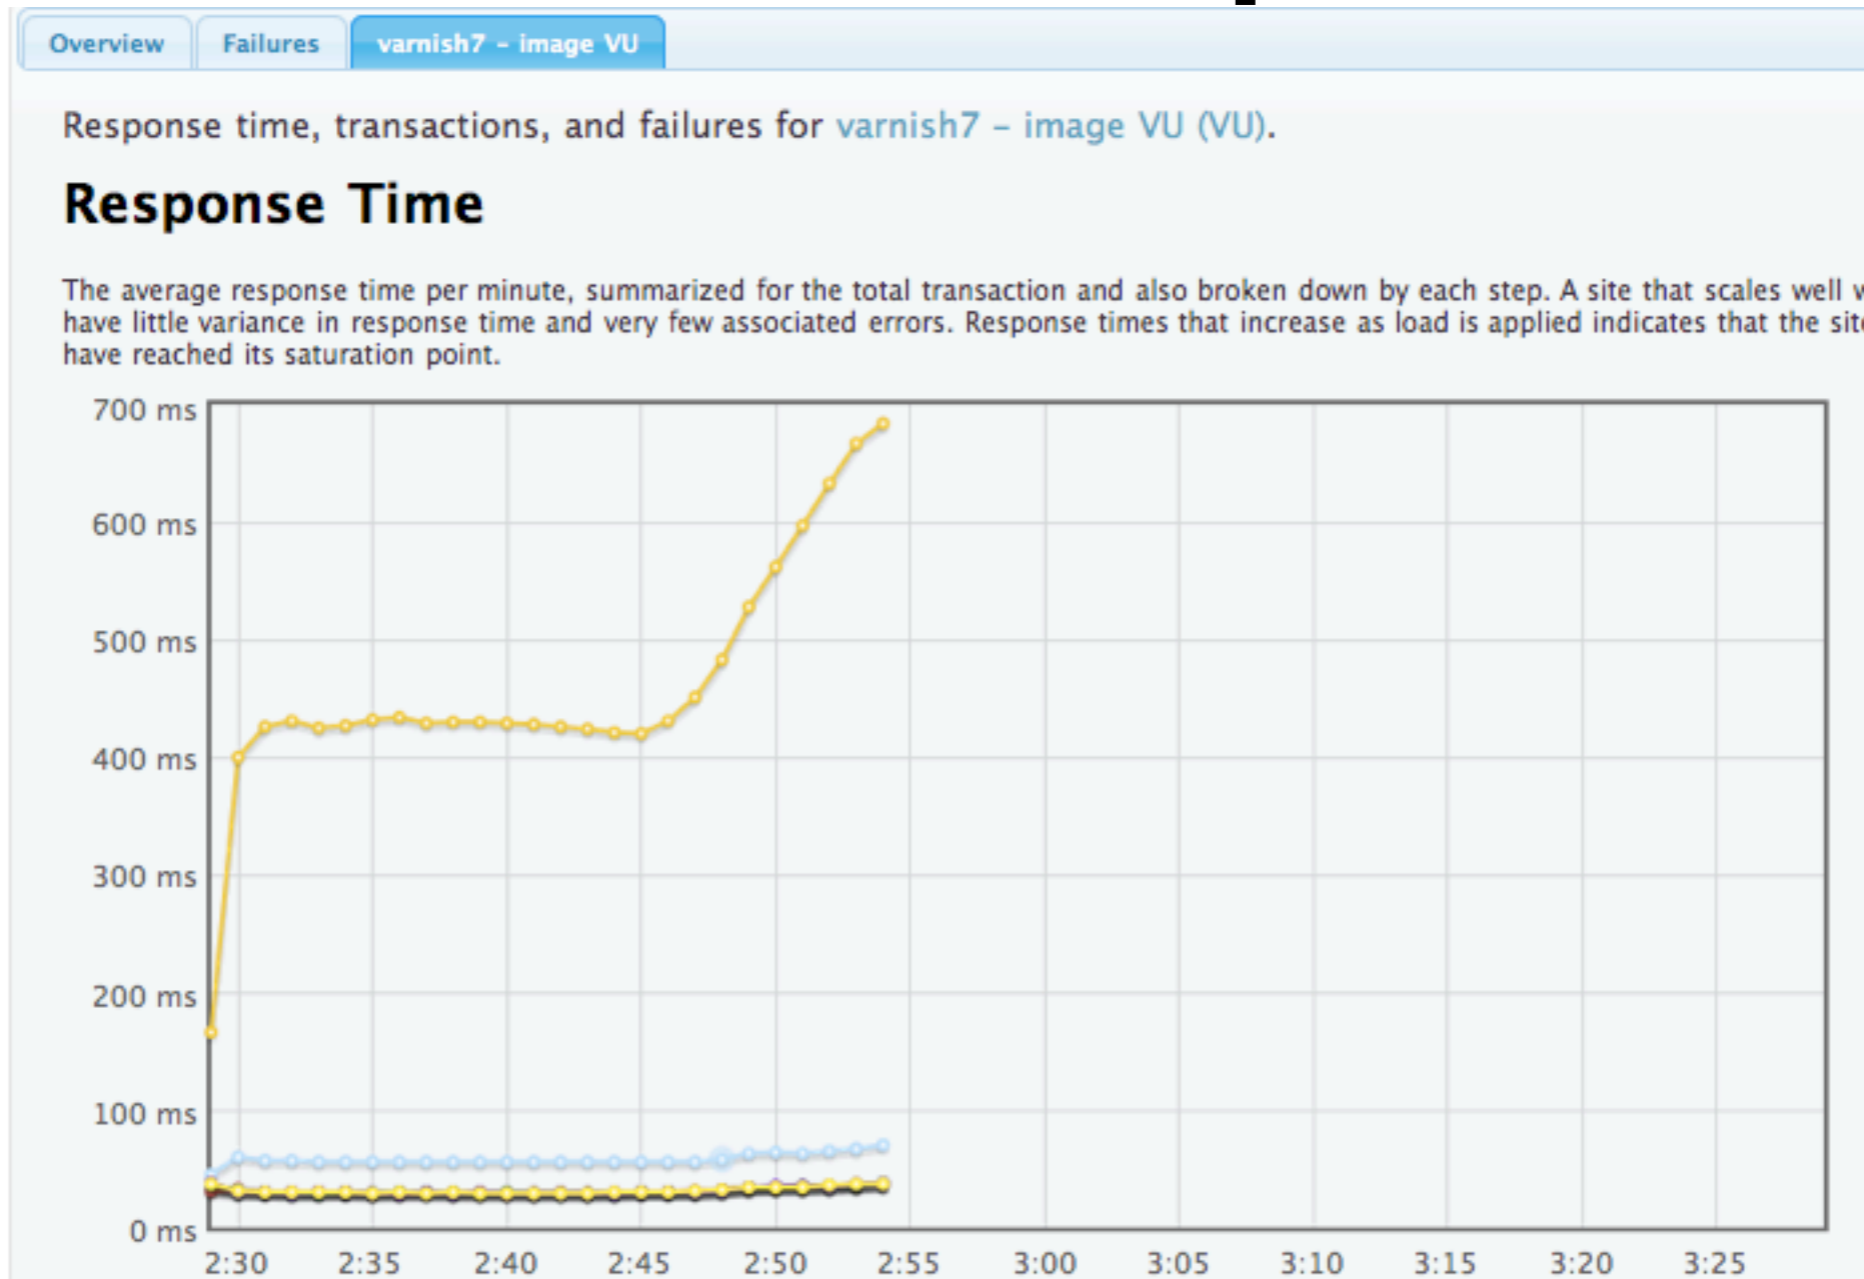
SSDs

- If you don't use them, you are an backwards idiot
- Less power per device
- Less spindles needed
- Way less power when you do $\text{power} * \text{spindles}$

Varnish/Database SKU

- 2 x 6 core westmere x5670
- 48 GB of RAM
- 2 spinning media (transaction logs)
- 6x Intel X25M 160 GB SSD
 - connect directly to motherboard and software raid -- no raid card can keep up
- ~350 watts in normal use for us

106k req/s



<http://browsermob.com/load/test/share/18d4849fae524f44b76148077bbccc4f>

joins

- SSDs to the rescue
- Not really a problem
- Semantic Mediawiki, I am looking at you
- 500 MB/sec of random reads with average latency of 0.2 ms for ~2000 USD
- 220 MB/sec of random writes

Wikia CDN

- Built our own
- DNS distribution based on Dynect
- Cache nodes in key internet locations
 - (Frankfurt in DE-CIX)

CDN Node

- Runs
 - Varnish
 - Quagga
 - Bind
 - Scribe based purging daemon

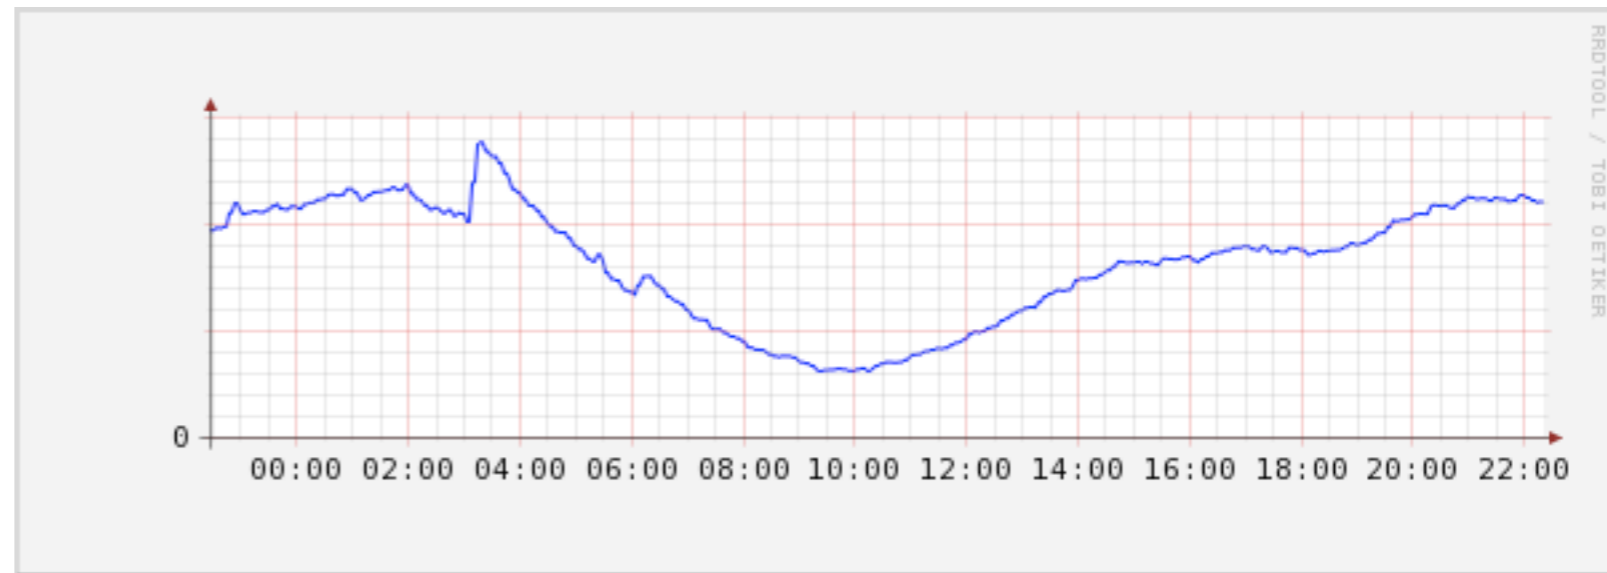
Why?!?

- Instant purging
- Fine grained control
- No commercial offering
- We still use akamai for static non dynamic content

Dynamic Content

- Just cache it
- Invalidate
- Or short TTLs
- Is ttl to save your backend

Load spikes



- Short ttl
- Prevents overloads
- Protects against DOS
- Drastically increases end user performance

Linking varnish

- Frankfurt talks to London
- London talks to Iowa
- Iowa talks to San Jose
- San Jose talks to local backend
- Frankfurt retries directly with Iowa
- London retries directly to San Jose
- Iowa retries with local backend
- San Jose retries with Iowa backend

Geographical Redundancy

- Master datacenter
- Backup datacenter
- Rendering datacenter
- Caching datacenter

Nuclear Bunker!

- <http://ganglia.wikia.net/iowa/>
- Live replication of site
- 20 megatons at 3 miles :)



Current state

- Automatic failover
- Replicated mysql
- Readonly mode

Expensive

- 100% cost overhead
- no benefit to end user in normal operations

Goal of active-active

- Cookie sniffing
- Send to closest DC unless editing
- Faster response times

Active-active-active

- Smaller datacenters
- More datacenters
- Way less overhead
- Happier users

Problems

- Memcache synchronization
- Session synchronization
- File synchronization

Memcache

- Facebook technique
- Insert memcache deletes into mysql blackhole
- Read replication log
- Remote invalidation

Riak

- key-value store
- HTTP interface
- scales across many machines
- N-Master inter-dc replication

Riak

- key-value store
- HTTP interface
- scales across many machines
- N-Master inter-dc replication

Files - RiakFS

- Fuse layer
- Separates metadata and data into two different buckets
 - SSDs for metadata
- MW uses filesystem
- Serve files directly from Riak with Varnish in front

```
if (req.http.host == "xxxx") {  
  if(req.url ~ "^/..") {  
    set req.url = regsub(req.url, "^/(.)(.)", "\1\1\2\1\2");  
  } else {  
    set req.url = regsub(req.url, "^/(.)", "\1\1\1");  
  }  
  set req.url = regsuball(req.url, "/", {"%2F"});  
  set req.url = regsuball(req.url, "\.", {"%2E"});  
  set req.url = regsuball(req.url, "\_", {"%5F"});  
  set req.url = "/riak/bucket/" req.url;  
  set req.backend = riak;  
  return(lookup);  
}
```

RiakFS

- Uses vector clocks to detect conflicts to metadata
- Last-write wins on the file
- Caches locally
- **Experimental**
- Will soon support new Riak large file store

Sessions

- Need to stay in sync
- Otherwise people appear logged out when failover
- Riak as a session store
- N-Master replication

Blob store

- Mediawiki store
- enormous blob tables in mysql
- horizontally scalable instead

Active-Active-Active

- Going live in January
- With a new node on the east coast

Routing requests

- VPN tunnels
 - Mesh
- OSPF
 - Set OSPF cost to latency
 - Faster than the internet

Varnish

- More content is cacheable than you think
- Non cacheable content is also faster
- For small objects, 2x speed increase!

syn-ack

- Long running connections
- London to San Jose
 - syn 100ms
 - ack 100ms
 - request 100ms
 - response 100ms

syn-ack (cdn)

- Long running connections
- London to London to San Jose
 - syn 15ms
 - ack 15ms
 - request 100ms
 - response 100ms

Slow start

- Persistent backend connections
- No syn-ack
- No slow start
- Big content comes down way faster

VCL

- Varnish Configuration Language
- Converted to C code
- Compiled
- Inline into the execution pass
- Wikia
 - 3500 statements

Purging

- Scribe based (used to be Rabbit) stream of purges
- Normalised URLs to purge all vary

```
if( req.http.X-Transformed-URL != "I"
  && ( req.url !~"\.php"
    || req.url ~ "title=."+ )
) {

# make sure we dont do this processing twice
set req.http.X-Origurl = req.url;
set req.http.X-Transformed-URL = "I";
set req.http.X-Set-Vary-URL = "I";

if (req.url ~ "^/index.php") {
  set req.url = regsub(req.url, ".*title=([^&]*).*", "/index.php?title=\1");
} else if (req.url ~ "^/wiki/") {
  set req.url = regsub(req.url, "/wiki/([^?]*).*", "/index.php?title=\1");
} else {
  set req.url = regsub(req.url, "/([^?]*).*", "/index.php?title=\1");
}
log "Tansformed: " req.http.X-Origurl " to: " req.url;
}
```

```
if( req.http.X-Transformed-URL != "I"  
  && ( req.url !~"\.php"  
    || req.url ~ "title=."+ )  
  ) {
```

```
# make sure we dont do this processing twice
```

```
set req.http.X-Origurl = req.url;  
set req.http.X-Transformed-URL = "I";  
set req.http.X-Set-Vary-URL = "I";
```

```
if (req.url ~ "^/index.php") {  
  set req.url = regsub(req.url, ".*title=([^&]*).*", "/index.php?title=\1");
```

```
} else if (req.url ~ "^/wiki/") {  
  set req.url = regsub(req.url, "/wiki/([^?]*).*", "/index.php?title=\1");
```

```
} else {  
  set req.url = regsub(req.url, "/([^?]*).*", "/index.php?title=\1");  
}
```

```
log "Tansformed: " req.http.X-Origurl " to: " req.url;
```

```
}
```

```
if( req.http.X-Transformed-URL != "I"
  && ( req.url !~"\.php"
    || req.url ~ "title=."+ )
) {

# make sure we dont do this processing twice
set req.http.X-Origurl = req.url;
set req.http.X-Transformed-URL = "I";
set req.http.X-Set-Vary-URL = "I";

if (req.url ~ "^/index.php") {
  set req.url = regsub(req.url, ".*title=([^&]*).*", "/index.php?title=\1");
} else if (req.url ~ "^/wiki/") {
  set req.url = regsub(req.url, "/wiki/([^?]*).*", "/index.php?title=\1");
} else {
  set req.url = regsub(req.url, "/([^?]*).*", "/index.php?title=\1");
}
log "Tansformed: " req.http.X-Origurl " to: " req.url;
}
```

```
# restore the cookie
if (req.http.X-Orig-Cookie) {
    set bereq.http.Cookie = req.http.X-Orig-Cookie;
}
if (req.http.X-Set-Vary-URL) {
    set bereq.url = req.http.X-Origurl;
}
```

```
unset bereq.http.X-Origurl;
unset bereq.http.X-Set-Vary-URL;
unset bereq.http.X-Orig-Cookie;
```

```
sub vcl_hit {
    if (req.request == "PURGE") {
        C{VRT_purge(sp, 0, 0);}C
        error 200 "Purged.";
    }
}
```

```
sub vcl_miss {

    if (req.request == "PURGE") {
        C{VRT_purge(sp, 0, 0);}C
        error 404 "Not purged";
    }
}
```

Loop detection

- Very long loops between varnish machines
- 20-30 hops
- very fast, no one noticed

```
if (req.http.X-Forwarded-From) {  
    C{  
        const char *identity = VRT_r_server_identity(sp);  
        const char *forwarded = VRT_GetHdr(sp, HDR_REQ, "\021X-  
Forwarded-From:");  
        char *substr = strstr(forwarded, identity);  
        if (substr != NULL) {  
            VRT_error(sp, 800, "Loop detected\n");  
            VRT_done(sp, VCL_RET_ERROR);  
        }  
    }C #;  
}
```

```
if (req.http.X-Forwarded-From) {  
    set req.http.X-Forwarded-From = req.http.X-Forwarded-From ", "  
server.identity;  
} else {  
    set req.http.X-Forwarded-From = server.identity;  
}
```

```
if (req.http.X-Forwarded-From) {  
    C{  
        const char *identity = VRT_r_server_identity(sp);  
        const char *forwarded = VRT_GetHdr(sp, HDR_REQ, "\021X-  
Forwarded-From:");  
        char *substr = strstr(forwarded, identity);  
        if (substr != NULL) {  
            VRT_error(sp, 800, "Loop detected\n");  
            VRT_done(sp, VCL_RET_ERROR);  
        }  
    }C #;  
}
```

```
if (req.http.X-Forwarded-From) {  
    set req.http.X-Forwarded-From = req.http.X-Forwarded-From ", "  
server.identity;  
} else {  
    set req.http.X-Forwarded-From = server.identity;  
}
```

```
sub vcl_error {
  if (obj.status == 800) {
    set obj.status = 503;
    set obj.http.Error-Reason = "loop detected";
    set obj.http.X-Forwarded-From =
      req.http.X-Forwarded-From;
    set req.http.X-Debug-Headers = "1";
    set obj.response = "Internal Server Error";
    synthetic {" Loop occured: "} req.http.X-Forwarded-
From;
    return(deliver);
  }
}
```

Geotargeting

- Geo cookie on requests
- Used for ads
- Set-Cookie: Geo =
{"city":"Baltimore","country":"US","lat":"39.3
18100","lon":"-76.613998"}; path=/

```

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\"};
path=/",
            record->city,
            record->country_code,
            record->latitude,
            record->longitude
            );
    VRT_AddHdr(sp, HDR_RESP, "\013Set-Cookie:", json, vrt_magic_string_end);
}
pthread_mutex_unlock(&geoip_mutex);
}C

```

A/B testing

- Everything is static
- Use cookie to select people

```
// it is 30
int choice = random() % 30;

if(choice == 0) {
    snprintf(cookie, 255, "wikia-test=xxxx; expires=Fri, 14-Sep-2012
15:44:53 GMT; domain=.wikia.com; path=/");
} else if(choice == 1) {
    snprintf(cookie, 255, "wikia-test=yyyy; expires=Fri, 14-Sep-2012
15:44:53 GMT; domain=.wikia.com; path=/");
} else {
    snprintf(cookie, 255, "wikia-test=zzzz; expires=Thu, 30-Sep-2010
23:59:59; domain=.wikia.com; path=/");
}
VRT_AddHdr(sp, HDR_RESP, "\013Set-Cookie:", cookie,
vrt_magic_string_end);
```

X-Served-By: varnish-s3-SJC, varnish-i6-IOWA

X-Cache: HIT, HIT

X-Cache-Hits: 6, 5

X-Age: 395

X-Varnish-Config: \$Revision: 19021 \$

```
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;
```

```
if(req.http.Cookie ~ "wikicities-wikia-colocation=iowa") {  
    set req.http.X-DC-Requested = "IOWA";  
}
```

```
if (req.http.X-DC-Requested == "IOWA") {  
    if(req.backend == sjc_apache) {  
        set req.backend = iowa_apache;  
    }  
    if(req.backend == sjc_apache_chute) {  
        set req.backend = iowa_apache_chute;  
    }  
    if(req.backend == imgvarnish_sl) {  
        set req.backend = imgvarnish_il;  
    }  
}
```

```
if (!req.backend.healthy || req.restarts > 2) {  
    if(req.backend == sjc_apache) {  
        set req.backend = iowa_apache;  
    } elseif (req.backend == sjc_apache_chute) {  
        set req.backend = iowa_apache_chute;  
    } elseif(req.backend == imgvarnish_sl) {  
        set req.backend = imgvarnish_il;  
    } elseif (req.backend == imgvarnish_il) {  
        set req.backend = imgvarnish_sl;  
    } elseif(req.backend == iowa_varnish) {  
        set req.backend = sjc_varnish;  
    }  
}
```

```
if(req.http.X-Origurl && req.http.User-Agent ~
"(Windows CE|Kindle|Opera Mini|IEMobile|Android|
Symbian|Danger|BlackBerry|iPod|iPhone)" && re\
q.http.host !~ "answers.wikia.com" && req.http.X-Origurl !
~ "(\.php|\?.)") {
    set req.http.X-Origurl = regsub(req.http.X-Origurl, "$",
"?useskin=wikiaphone");
}
```

```
if(server.ip ~ SJC || server.ip == "127.0.0.1" ||
server.hostname ~ "dev-") {
    if(req.http.host ~ "origin-images.wikia.com") {
        set req.backend = imgvarnish_sl;
    } else {
        set req.backend = sjc_apache;
    }
} elseif (server.ip ~ FRA) {
    if (req.restarts > 1) {
        set req.backend = iowa_varnish;
    } else {
        set req.backend = london_varnish;
    }
}
```

```
} elsif (server.ip ~ LON) {  
  if (req.restarts > 1) {  
    set req.backend = sjc_varnish;  
  } else {  
    set req.backend = iowa_varnish;  
  }  
}  
}  
} elsif (server.ip ~IOWA) {  
  if (req.http.host == "origin-images.wikia.com") {  
    set req.backend = imgvarnish_il;  
  } else {  
    set req.backend = sjc_varnish;  
  }  
}  
} else {  
  set req.backend = sjc_varnish;  
}
```

```
if (!req.backend.healthy || req.restarts > 2) {  
  if(req.backend == sjc_apache) {  
    set req.backend = iowa_apache;  
  } elseif (req.backend == sjc_apache_chute) {  
    set req.backend = iowa_apache_chute;  
  } elseif(req.backend == imgvarnish_sl) {  
    set req.backend = imgvarnish_il;  
  } elseif (req.backend == imgvarnish_il) {  
    set req.backend = imgvarnish_sl;  
  } elseif(req.backend == iowa_varnish) {  
    set req.backend = sjc_varnish;  
  }  
}
```

Thank you

- Questions?