

Optimizing a News Website for Speed and Stability

Julian Dunn
SecondMarket, Inc.

Blake Crosby
Canadian Broadcasting Corp.

Friday, December 14, 12

- Who has worked for a news website
- Who has configured a CDN - familiar with a CDN?

What's This Talk About?

Friday, December 14, 12

What's This Talk About?

- **How to build & operate a large media website**

- **without spending millions of \$**
- **without having it go down every 5 minutes**

Some technical things, but set the stage with non-technical things



How to build and operate a large media website...
... without spending millions of dollars

Friday, December 14, 12

```
eax: 00000000
esi: c2c24100
es: 0018
ebx: c2c24100
edi: c2c24100
es: 0018
ss: 0018
destroy_inode [kernel] 0x21 (2.4.20-37.7.1)
ecx: 00000000
ebp: 00000270
```

... and without your site going down every 5 minutes.

```
Stack: c0311254 c015dd8c c2c24100 00217600
f4a43880 c2c24100 c015afc1 c2c24100
c36bff84 c36bff84 00000000 00000000 303eee06 00000000
[<c015dd8c>] iput [kernel] 0x2ac (0xc36bff84)
dentry_iput [nfs] 0x5a (0xc36bff84)
[kernel] 0xe1 (0xc36bff84)
[kernel] 0x84 (0xc36bff84)
[kernel] 0x20 (0xc36bff84)
[kernel] 0x13 (0xc36bff84)
```

Who Am I?

Friday, December 14, 12

Systems engineer, worked for CBC as digital operations team lead for many years
Now working for alternative markets startup in New York City doing cloud infrastructure & web operations.



Friday, December 14, 12

CBC – Canada's National Public Broadcaster

– Similar to NPR + PBS – combined radio and television station

IN THE NEWS

- Gift gaffes
- Photos: Deadly typhoon
- Next Hillary Clinton?
- NHL lockout quiz

- Home
 - World
 - Canada
 - Politics
 - Business
 - Health
 - Arts & Entertainment
 - Technology & Science
 - Community
 - Weather
 - Video
- Home Photo Galleries The National CBC News Network

DEVELOPING Coroner's jury wants all Ontario daycares registered » < || >

NHL rejects latest CBA proposal from players' association



The NHL labour talks have hit another roadblock — and this one appears to be massive. After two optimistic days that suggested a deal could be imminent, negotiations between the league and the NHL Players' Association broke off dramatically Thursday. 9:52 PM ET

- Statements from Ron Burkle, other NHL owners
- Social Media: NHL players react to latest rejected offer
- WATCH: Theo Fleury confident in Sidney Crosby's role in ending lockout
- NHL, NHLPA still need to bridge 'critical open issues' 10:07 AM ET

NEW | F-35 deal not cancelled, Tories insist



The Harper government says it has not made a decision on the F-35 as a replacement for Canada's CF-18 fighter

More Headlines

Pregnant Kate leaves hospital feeling 'much better'

EXCLUSIVE | Costa Concordia wreck an accident, not crime, captain says

No plans to end restrictions on prohibited guns, Harper says

UPDATED | Man arrested in Surrey, B.C., SUV hit and run

Accused in NYC subway death says man grabbed him first

Neutrality on Mideast supported by half of Canadians, poll suggests

Stay Connected with CBC News

- Mobile
- Facebook
- Podcasts
- Twitter
- Alerts
- Newsletter

Must Watch



The National for December 06, 2012
1:00:02
Welcome to The National, the flagship nightly newscast of CBC News,



Political punch
1:34
Definitely not democracy; the top political brawls from around the world



'End of the world is coming'



RAW: Winter arrives out West

Friday, December 14, 12

cbc.ca

- Most popular news website in Canada - 1 million unique visitors per day
- News, sports, radio & TV streaming (live and on demand), regional news, arts & entertainment, etc. etc.

2005



Quality
● 240p
Auto

0:19 / 0:19

YouTube

Friday, December 14, 12

Very first YouTube video was posted in 2005. Jawed Karim co-founder right here in San Diego at the zoo. (which I highly recommend -- it's \$44 but worth the money. one of the best zoos)



Friday, December 14, 12

2005 was also the year of a provincial election in British Columbia



Main

News

Overall Election Results

Party	Elected	Leading	Total	Vote Share
LIB	46	0	46	46.03%
NDP	33	0	33	41.27%
GRN	0	0	0	9.11%
DR	0	0	0	.84%
OTH	0	0	0	2.76%
Last Update: June 1, 5:11:20 PM PDT				79 seats

[an error occurred while processing this directive]




B.C. Premier Gordon Campbell greets his supporters during a rally in Vancouver, B.C., after he and his Liberal party won a second term in office, Tuesday, May 17, 2005.

[an error occurred while processing this directive]

District Profiles

Select a district

Highlights



Your View
If you have something to say about the results, this is the place say it.
[more »](#)

B.C. Votes 2005

- Main
- Features
- Parties & Leaders
- District Profiles
- Analysis & Commentary
- News Archive
- Voters Toolkit
- Your View
- Radio & TV Coverage

cbc.ca/bc



Friday, December 14, 12

We put up this website -- unbelievably it's still up. I can also talk about data retention issues on a news website but that's a whole other talk.

Aw Jeez, not this shit again!



Error 500 - Internal Server Error


I've got some bad news for you sunshine : the server made a boo-boo. Please go take a coffee and retry in a few minutes.

If you feel that contacting me will somehow improve your experience on my site, go ahead : mail ozh at planetozh dot com.

Friday, December 14, 12

Polls closed at 9 p.m. Pacific Time

Shortly after the site was down for over two hours



50 requests/sec per server!

MAKE GIFS AT GIFSOUP.COM

Friday, December 14, 12

(300 origin req/sec over 6 servers) WTF, right? Why?



Bad Assumptions

Friday, December 14, 12



“Make the news stories
appear as fast as possible”

Friday, December 14, 12

Journalists say, “I want the stories to be as fresh as possible” – we’re a news site after all



Friday, December 14, 12

Developers say, ok that's great, we'll put all the stories in an Oracle database



+



Friday, December 14, 12

And we'll render them using a J2EE web application.



Architecture was wrong

Friday, December 14, 12

But this is actually not a great architecture for a news website.

To understand why this is a bad idea, you have to look at news site access patterns


Site Content Delivered, in Mbits per Second

Total bandwidth includes all request traffic plus edge, midgress, and origin response traffic.

Total Volume: 8.8 TB

▼ Peak: 13,960.78 Mbits/sec at 08:25PM

◀ Latest: 692.98 Mbits/sec at 11:55PM



Eastern Daylight Time

	Total Volume	% Volume
Edge Traffic	7.8 TB	88.1%
Midgress Traffic	654.7 GB	7.4%
Origin Traffic	397.4 GB	4.5%

[Show Upload/download breakdown](#)


Origin Bandwidth Offload 94.9% of bytes offloaded to Akamai

Friday, December 14, 12

News of Whitney Houston's death Feb 11, 2012


From < 1 Gbps to 13 Gbps virtually instantaneously

Number of Story Updates Per Hour



Friday, December 14, 12

Story is updated very infrequently, though. Only about 5-6 times per hour even when news is breaking, and over a story's lifetime perhaps one or two dozen times.



Friday, December 14, 12

Senseless to use J2EE application to render stories.

Wastes a lot of CPU doing the same thing over and over.

Can even read about this in the outage post-mortem report from BC Votes 2005.

- “Election news stories were not cached on disk.”
- “All stories were generated dynamically which requires more computer processing, hence resulting in more load on the servers.”

Real, Conflicting Requirements

- Stability and reliability
- Content freshness
- System complexity
- Cost

Site Architecture: SSI

- Primitive, but fast/secure
- Audit trail
- Easy to turn off site features
- Details of “wrapper” mechanism in paper

Friday, December 14, 12

- “Cache cache cache” -- we took this to the extreme and made the backend as close to a cache as possible.
- News portion: no dynamic features outside SSI
 - SSI is primitive, but fast/secure
 - Audit trail - nothing gets into production w/o going through CMS
 - Turn off features easily by modifying wrappers to discard site elements

Delivery Parameters



Set Global Site TTL

Friday, December 14, 12

Originally all kinds of TTL rules for different types of content, path names, etc. Too complicated; journalists constantly calling helpdesk to “purge things from cache”

If you have to “purge things from cache” you’re doing it wrong.

Setting a global site TTL (of 20 seconds) covers most cases and is “fresh enough” (ran some experiments with journalists and they seem not to notice that publishing is not actually instant)

If-Modified-Since:



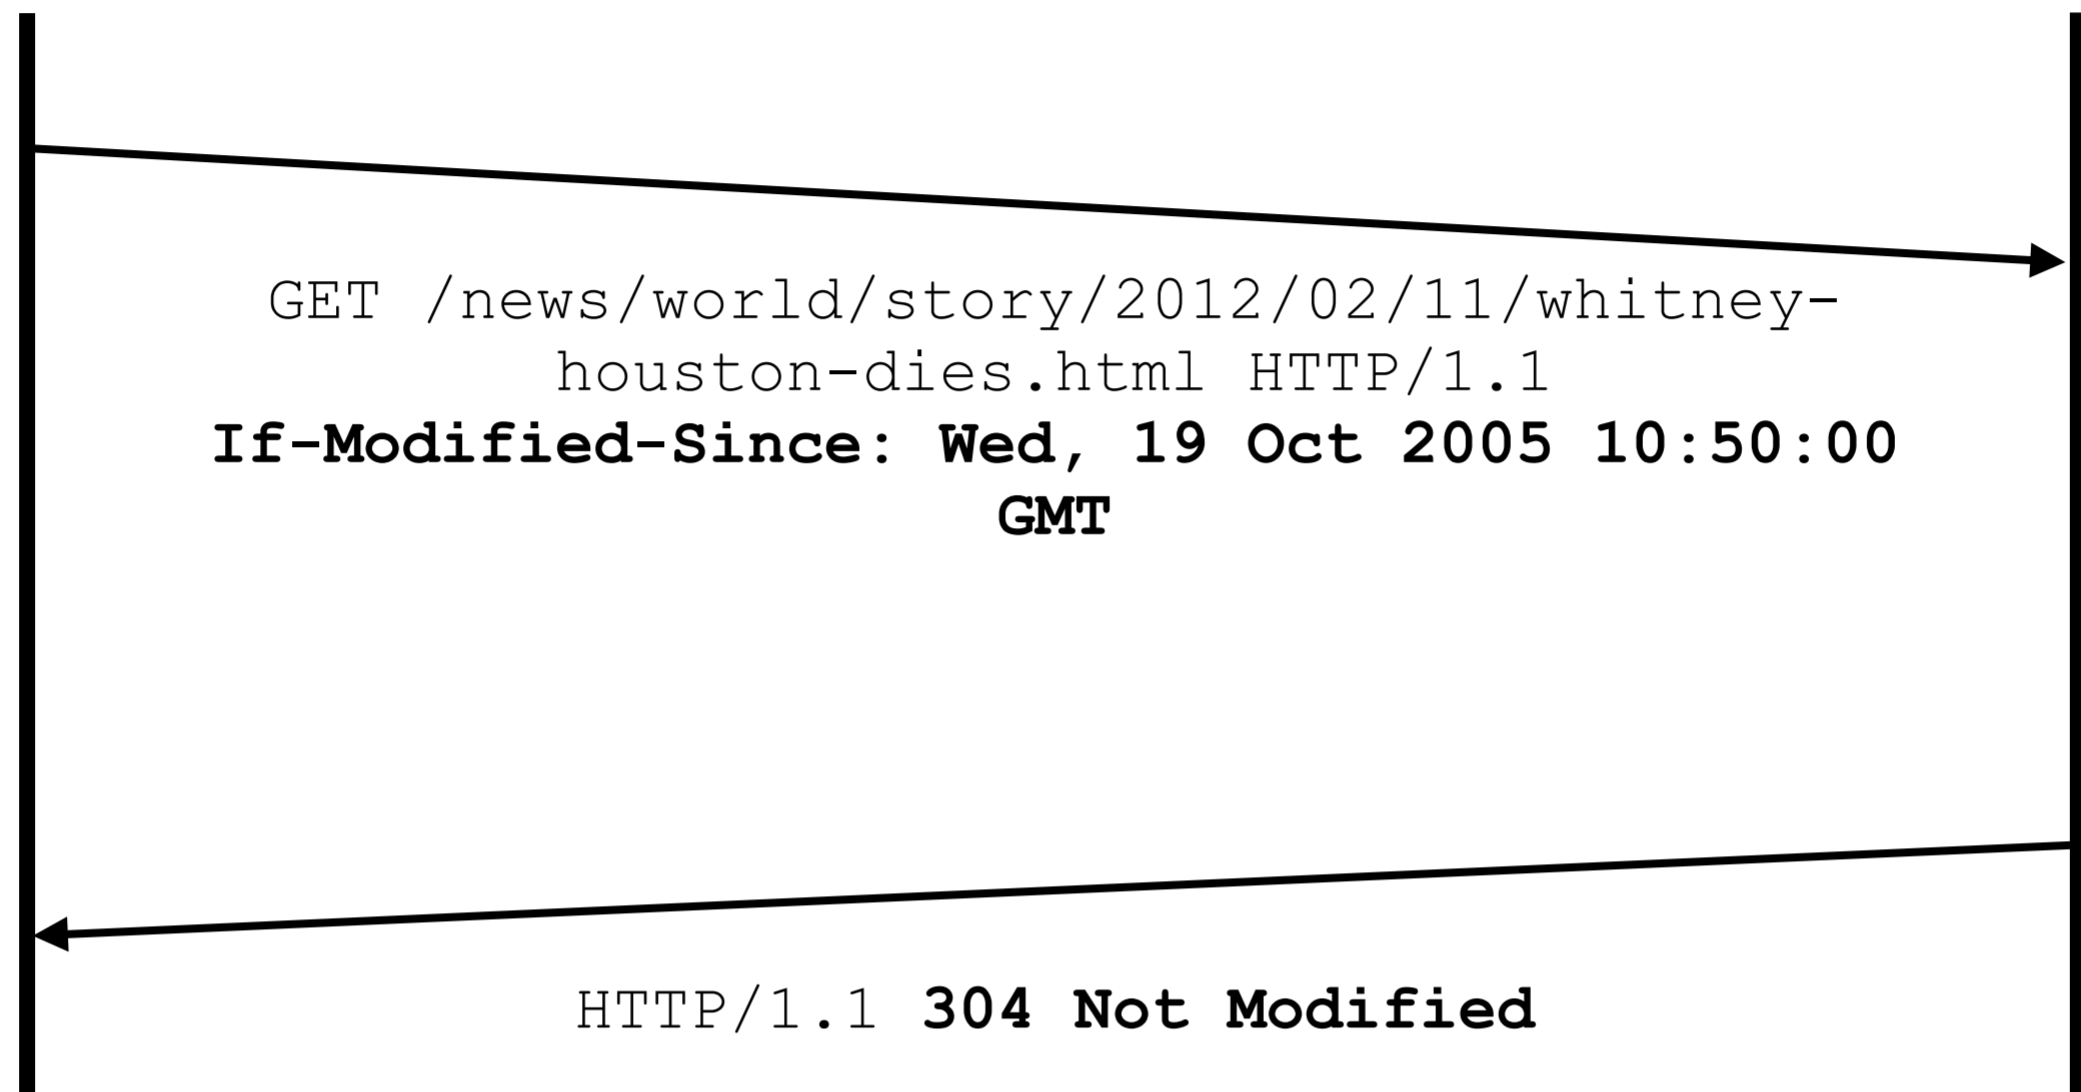
Friday, December 14, 12

This is the regular HTTP transaction that we all know and love.

If-Modified-Since:

CDN

Origin




Friday, December 14, 12

If IMS is configured properly and the CDN has the object in cache, this is what the interaction looks like. Very small origin response, no need to send content body in most cases.

Origin Response Codes, in Requests per Second

Origin Requests: 2.1 Billion

▼ Peak: 1,623.2 requests/sec at 11:30AM
 ▲ Latest: 1,051 requests/sec at 10:00AM



Eastern Daylight Time

Estimated data begins at 5:00 AM

	NUMBER OF REQUESTS	% OF REQUESTS
■ OK: 200, 201, 204, 206, 210	547,813,196	26.2%
■ Not-Modified: 304	1,457,472,639	69.8%
■ Redirect: 301, 302	22,993,403	1.1%
■ Not-Found: 404	29,173,266	1.4%
■ Permission: 401, 403, 415	782,026	0%
■ Server Error: 500, 501, 502, 503, 504	107,952	0%
■ Other: (all other status codes)	28,530,176	1.4%

Friday, December 14, 12

“Most cases” actually turns out to be a lot... over two thirds of requests can be served out of CDN cache



Organize filesystem by desired TTL

Friday, December 14, 12

Allows you to override TTLs by path, e.g. long-lived objects can go in one top-level directory, medium-lived objects go in another, and so on. (examples: /favicon.ico, global site JS and CSS, etc.)

Last mile acceleration, origin compression



Friday, December 14, 12

origin compression – make sure to talk gzip to the CDN

LMA – enable CDN to pass along the same gzipped content to the end-user



Persistent connections

Appropriate timeouts

Friday, December 14, 12

* Set pconns with the CDN

* Origin KeepAliveTimeout value synchronized with CDN (CDN, max 300s for pconn, origin KeepAliveTimeout 301s) so that origin doesn't kill the pconn before the CDN does

General Lessons Learned

- Keep cache rules simple
- Keep tuning knobs at origin if you can
- Organize & categorize content
- **Understand meaning of “TTL”**

Friday, December 14, 12

TTL is not like a DNS TTL. It just means the time after which the CDN will check the origin for object freshness. If still fresh, then it'll keep serving the object out of cache.

Outcomes

Outcomes

- 2003-2010 - no need to grow origin
- 2010-today - 9 origin webserver

- Average 40% CPU utilization


Friday, December 14, 12

7 years – upgraded from HP/UX to SUSE Linux to RHEL – 3 generations of servers but didn't change the quantity


2010 – decided to buy ourselves a bit more runway in the event of high traffic situations

Entire origin farm to deliver content costs around \$15K to buy, astonishing for a large news website.

Origin Bandwidth Offload, in percentage



Origin Request Offload, in percentage



Friday, December 14, 12

Another set of graphs showing that the strategy pays off; CDN serves 80–90% of traffic (both bandwidth & requests)

@julian_dunn
jdunn@aquezada.com
www.juliandunn.net
ICQ: 1326831
(that's what was cool in 2005)

Friday, December 14, 12

My wife's ICQ # is 1290758 so she brags that she's cooler than me because her number is lower.