

Web Performance

Sergey Chernyshev

March '09 New York Web Standards Meetup

New York, NY

March 19th, 2009

About presenter

- Doing web stuff since 1995
- Director, Web Systems and Applications at [truTV](#)
- Personal projects
 - [TechPresentations.org](#)
 - [MediaWikiWidgets.org](#)
 - [SharingButtons.org](#)
 - [HowWebWorks.com](#) (*in works*)

What is this talk about?

- Performance is how fast your site works
- **NOT** how many users it can serve ([Scalability](#))
- **NOT** how often it's down ([Reliability](#))

Why do we need to care?

- User experience

- Money!!!

- **+100ms** → **-1% sales (Amazon)** *

- **+400ms** → **-5-9% full-page traffic (Yahoo! Autos front page)** **

- **+500ms** → **-20% searches (Google)** ***

* [Make Data Useful](#), Greg Linden at Stanford Data Mining class, fall 2006

** [YSlow 2.0 early preview in China](#), Stoyan Stefanov, December 6th, 2008

*** [Scaling Google for Every User](#), Marissa Mayer at Google Seattle Conference on Scalability 2007

How web works?



Backend

- Usually planned for along with scalability and reliability
- Usual tree-tier
 - *Browser (front-end)*
 - *Web/app server*
 - *Relational database*

DataBase performance

- use INDEXES!
- Study set theory basics and use JOINS instead of cursors or code-level iterations
- Use correct datatypes to put as much stuff into memory as possible
- Use query cache
- Read documentation for your RDBMS *

* For MySQL, [15 Ways to Kill Your MySQL Application Performance](#) by Jay Pipes at PHP Tek 2007

Compile

- Compiled code is faster than interpreted
- C/C++ web apps are rare
- ASP.NET and Java compile into bytecode
- But most of the web is interpreted, use opcode caches
 - *APC for PHP* *
 - *mod_perl for perl*

* [APC at Facebook](#) by Brian M. Shire and [Facebook Performance Caching](#) by Lucas Nealan at PHP Tek 2008

Cache

- Cache results
 - *Shared memory (APC and EAccelerator for PHP)*
 - *memcached*
- Cache not even data, but chunks of HTML
- Reverse proxies (nginx, Squid)

CDNs or lightweight web servers for static content

- CDNs are closer to the user and remove load your app servers
 - *Akamai (also has Edge Suite - reverse proxy)*
 - *Limelight*
- *mystaticfiles.net* with lightweight servers
 - *nginx (also reverse proxy)*
 - *lighttpd*

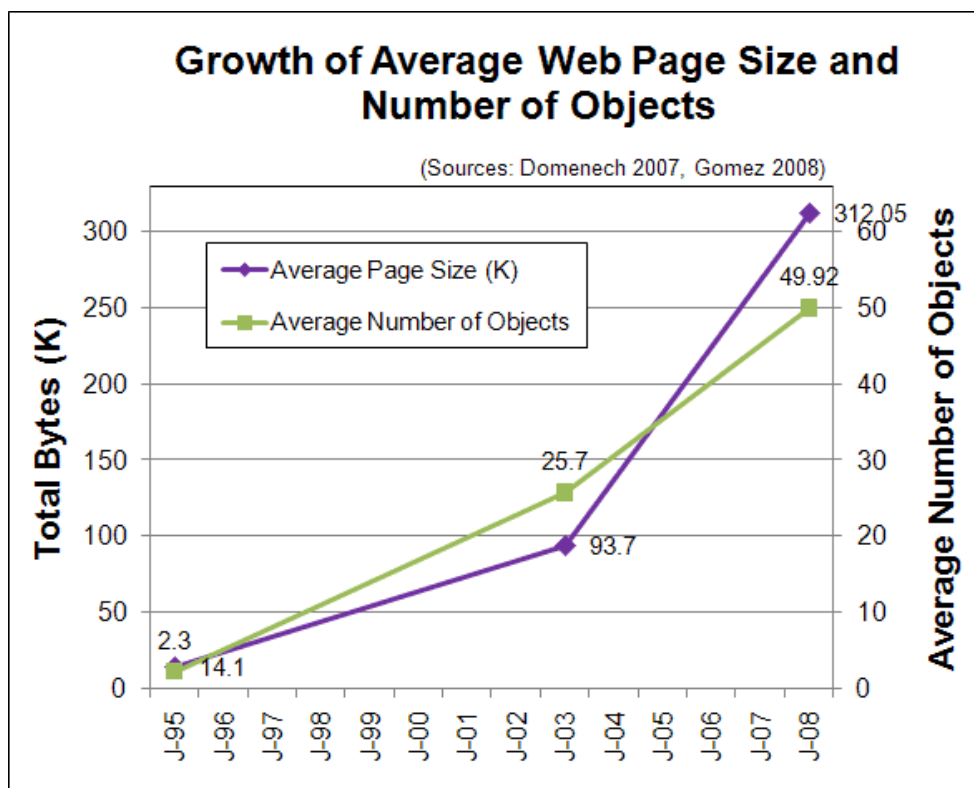
Front-end

- Face of the web
- Heavily influenced by marketing, brand and etc.
- Performance wasn't planned for until Web 2.0
- Slowness became noticable by users

Front-end: Start Here!

Amount of media and requests per page grew exponentially.

1995 - 2008

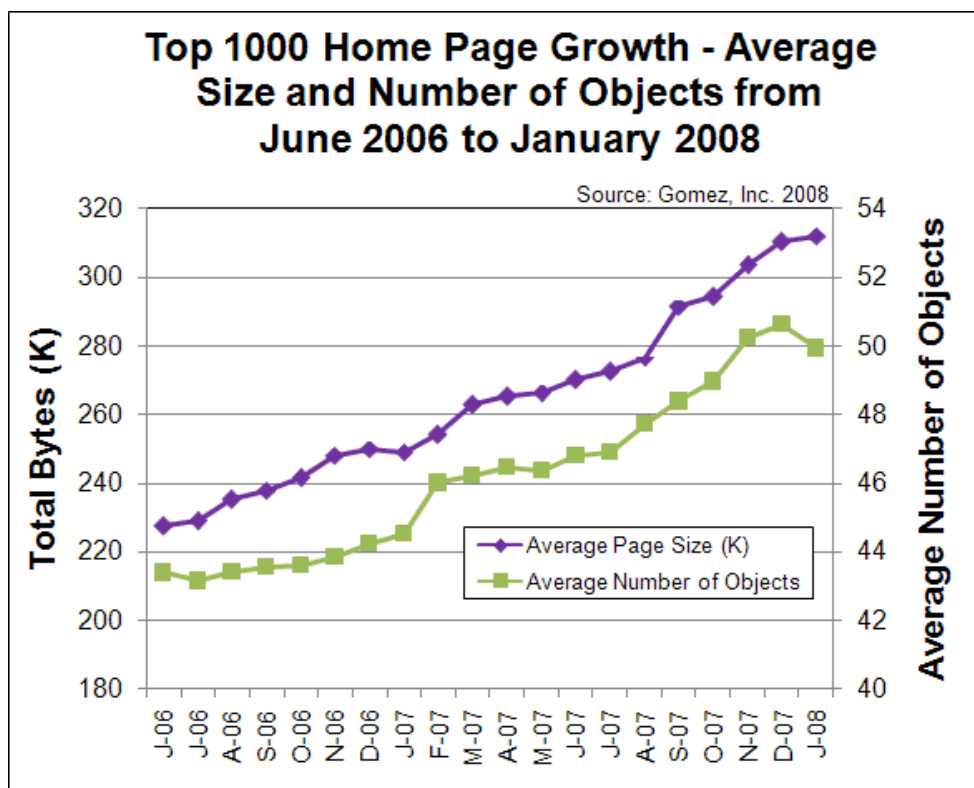


* Average Web Page Size Triples Since 2003, as of end of 2007 (via Nicole Sullivan)

Front-end: Start Here!

Top 1000 Home pages, ~25-30 growth in just 1.5 years.

Jun 2006 - Jan 2008




* Average Web Page Size Triples Since 2003, as of end of 2007 (via Nicole Sullivan)

Amazon Waterfall!

- Total Requests: 88
- Total Time: 6.344 seconds
- Back-end Time: 0.968 seconds = **just 15%**


TRANSFER TIMELINE



CNN Waterfall!

- Total Requests: 174
- Total Time: 4.406 seconds
- Back-end Time: 0.171 seconds = **less than 4%**

TRANSFER TIMELINE



Yahoo!

- Developed YUI design patterns and JavaScript library
- Had to pay attention to performance as part of experience
- **Exceptional Performance Team** at Yahoo!
- **Steve Souders*** (now at Google)
- Developed prioritized **list of best practices**
- Developed **YSlow**** to test performance and promote to business within Yahoo!

* **Steve Souders** @ TechPresentations

** **YSlow** @ TechPresentations

Best Practices (34 already)

- Make Fewer HTTP Requests
- Use a Content Delivery Network
- Add an Expires or a Cache-Control Header
- Gzip Components
- Put Stylesheets at the Top
- Put Scripts at the Bottom
- Avoid CSS Expressions
- Make JavaScript and CSS External
- Reduce DNS Lookups
- Minify JavaScript and CSS
- Avoid Redirects
- Remove Duplicate Scripts
- Configure ETags
- Make Ajax Cacheable
- Flush the Buffer Early
- Use GET for AJAX Requests
- Post-load Components
- Preload Components
- Reduce the Number of DOM Elements
- Split Components Across Domains
- Minimize the Number of iframes
- No 404s
- Reduce Cookie Size
- Use Cookie-free Domains for Components
- Minimize DOM Access
- Develop Smart Event Handlers
- Choose <link> over @import
- Avoid Filters
- Optimize Images
- Optimize CSS Sprites
- Don't Scale Images in HTML
- Make favicon.ico Small and Cacheable
- Keep Components under 25K
- Pack Components into a Multipart Document

Top Best Practices

Most effective, tested by YSlow

- Make Fewer HTTP Requests
- Use a Content Delivery Network
- Add an Expires or a Cache-Control Header
- Gzip Components
- Put Stylesheets at the Top
- Put Scripts at the Bottom

Top Best Practices (cont'd)

- Avoid CSS Expressions
- Make JavaScript and CSS External
- Reduce DNS Lookups
- Minify JavaScript and CSS
- Avoid Redirects
- Remove Duplicate Scripts
- Configure ETags

More Front-end Best Practices

- Flush the Buffer Early
- Post-load Components
- Preload Components
- Reduce the Number of DOM Elements
- Minimize the Number of iframes
- Optimize Images
- Optimize CSS Sprites

Demo

- YSlow (Firefox)
- Firebug's Net panel (Firefox)
- Fiddler 2 Beta (IE)
- AOL Page Test / [WebPageTest.org](http://www.webpagetest.org) (IE)

How do I start?

How do I start? Business people

- use metrics that include site performance, tie it in to your bottom line (e.g. can you afford loosing 20% of your traffic to 500ms slowness? What does it cost you to bridge this gap?)
- incorporate performance testing in QA process (developers should not care more then you do). Use YSlow - it speaks business language.
- but be careful, it might be too expensive to excel on all levels (still cheap to get a lot of improvement)

How do I start? Designers

- watch [Design Fast Websites](#) presentation by Nicole Sullivan (Yahoo!, co-author of Smoosh.it)
- start designing experience, not digital paper
- use consistent styles (reusable CSS, cached)
- don't mandate effects that require heavy lifting on browser side
 - *Rounded corners can still be fast*
 - *Transparent PNGs can still be fast*
- use graceful degradation for less capable browsers (IE6)

How do I start? Front-end developers

- most of it is on your shoulders
- include YSlow grades in your year end review and resume
- include performance into your definition of good craftsmanship (it's a rare case when you can use TDD effectively)
- learn how browsers work, DOM, events, if that trick of yours really makes it faster
- reduce amount of requests, use one CSS file (load it first), one JS file (load as late as possible), CSS sprites
- "smush" images, compress CSS and JS, automate it
- change your code and publishing process to allow infinite expiration of assets (if you can do that for HTML, ask to double your bonus)

How do I start? Backend developers

- you should already know what to do - performance culture is ages old here
- still heavily depends on your apps and connected to scalability and reliability
- use indexes in your databases! Monitor how they perform as you dataset grows
- use caches (RAM is cheap, disks are even cheaper)

How do I start? System administrators

- talk to designers and developers and even business people
- read RFC 2616 (HTTP) and Apache manual
- configure gzipping (mod_deflate) and expires
- install APC, memcached, Squid, nginx

More info

- TechPresentations.org/Performance
- [Yahoo! Exceptional Performance Team](#)

Tools

■ Diagnostic Tools

- *YSlow for Firebug*
- *Fiddler2 (WinINET)*
- *WebPageTest.org and AOL Page Test*
- *HttpWatch (\$\$\$)*

■ Content Compressors

- *JSMin by Douglas Crockford*
- *YUI Compressor - for JS and CSS*
- *SmushIt.com - for images (web interface + API, not Open Source yet)*

■ Server-side

- *APC for PHP*
- *memcached*
- *mod_deflate for Apache2*
- *nginx - lightweight reverse proxy*
- *Squid - caching reverse proxy*
- *Akamai, Limelight - CDNs*

About these slides

These slides use an HTML presentation tool [Slidy](#).

They are [valid XHTML](#) document with [rich metadata](#) embedded using [RDFa](#).



Licensing & Attribution

Web Performance by [Sergey Chernyshev](#)

is licensed under a [Creative Commons Attribution 3.0 United States License](#).

Permissions beyond the scope of this license [available here](#).

