

Why your site is slow
DrupalCamp New Jersey
2016-01-30



Hi, I'm Steve Persch



- Agency and Community Engineer
- @stevector and Twitter, Github, D.O
- @DrupalHaikus on Twitter

What we're covering



Key architectural points related to speed:


- Server-side
- Client-side
- The underlying 'why'

What a website does



- Sends responses from a server
- over a network
- to be interpreted by a browser

Clients and servers



What a website does



More abstractly:

- **Converses with a client**
- **Informs the world about an organization**

Server-side



PROBLEM



The website is too slow

- **WHY? - The server is not responding with HTML fast enough**
- **WHY? - (Web 1.0 answer) There is a limit to how fast a .html file is read off a hard disk**

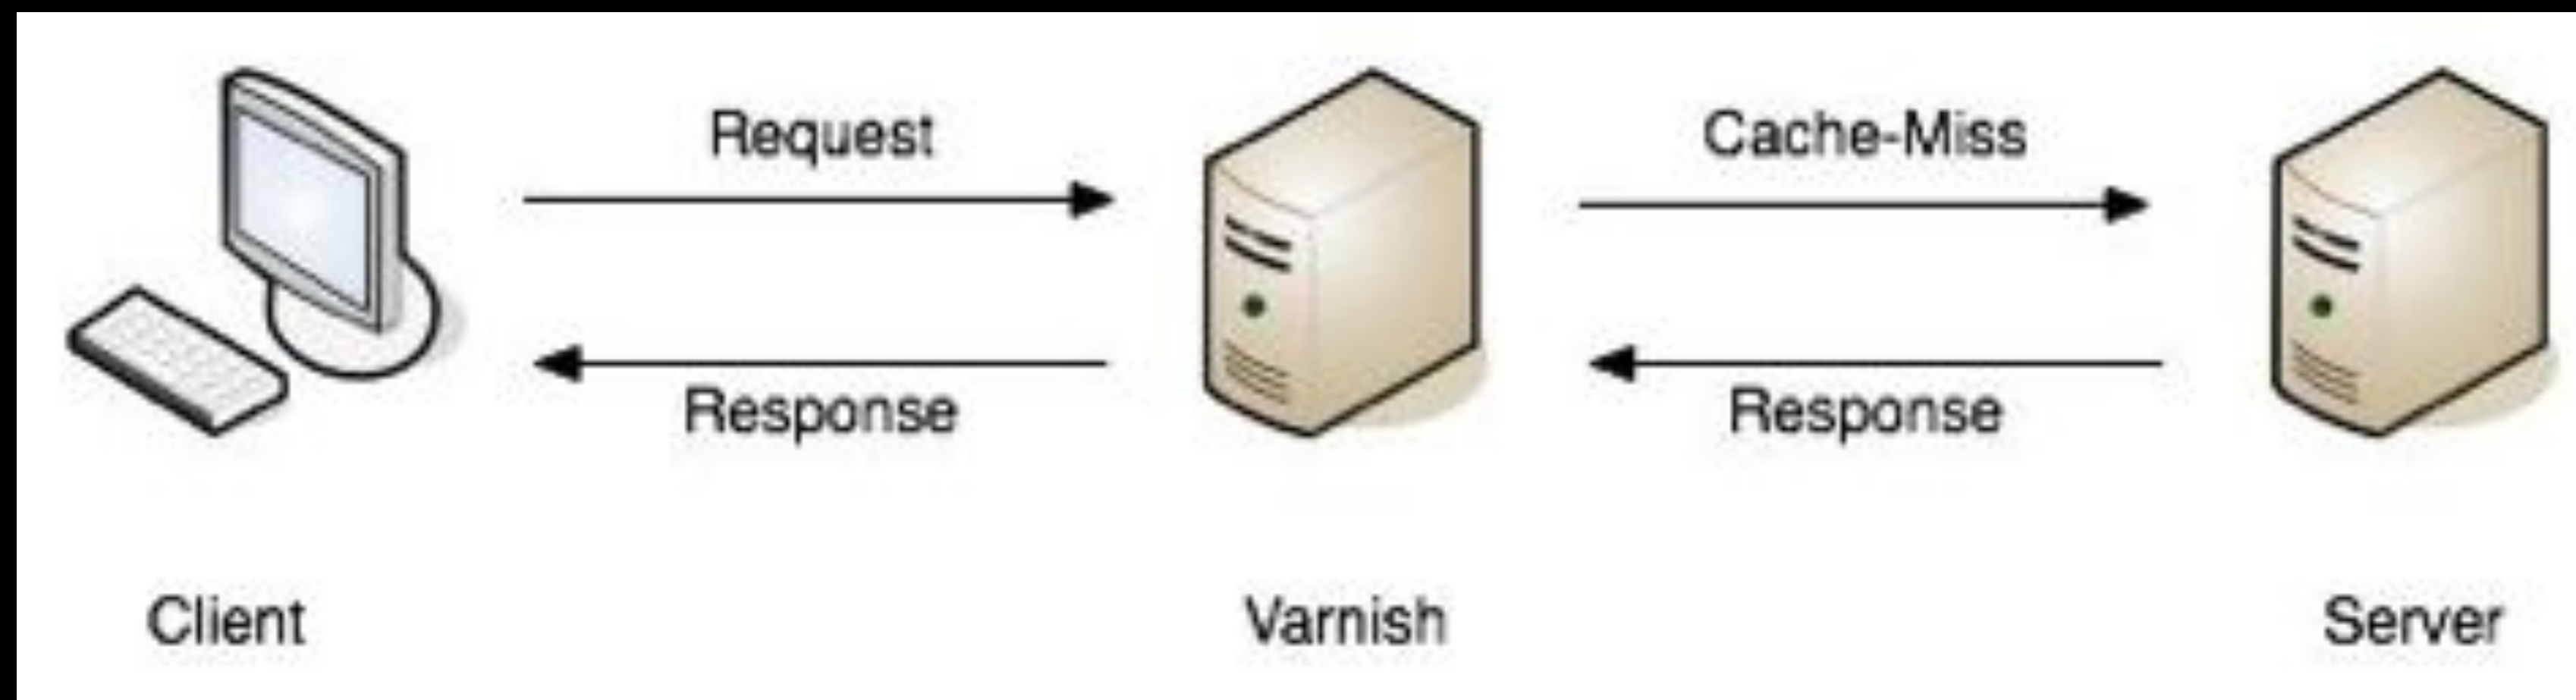
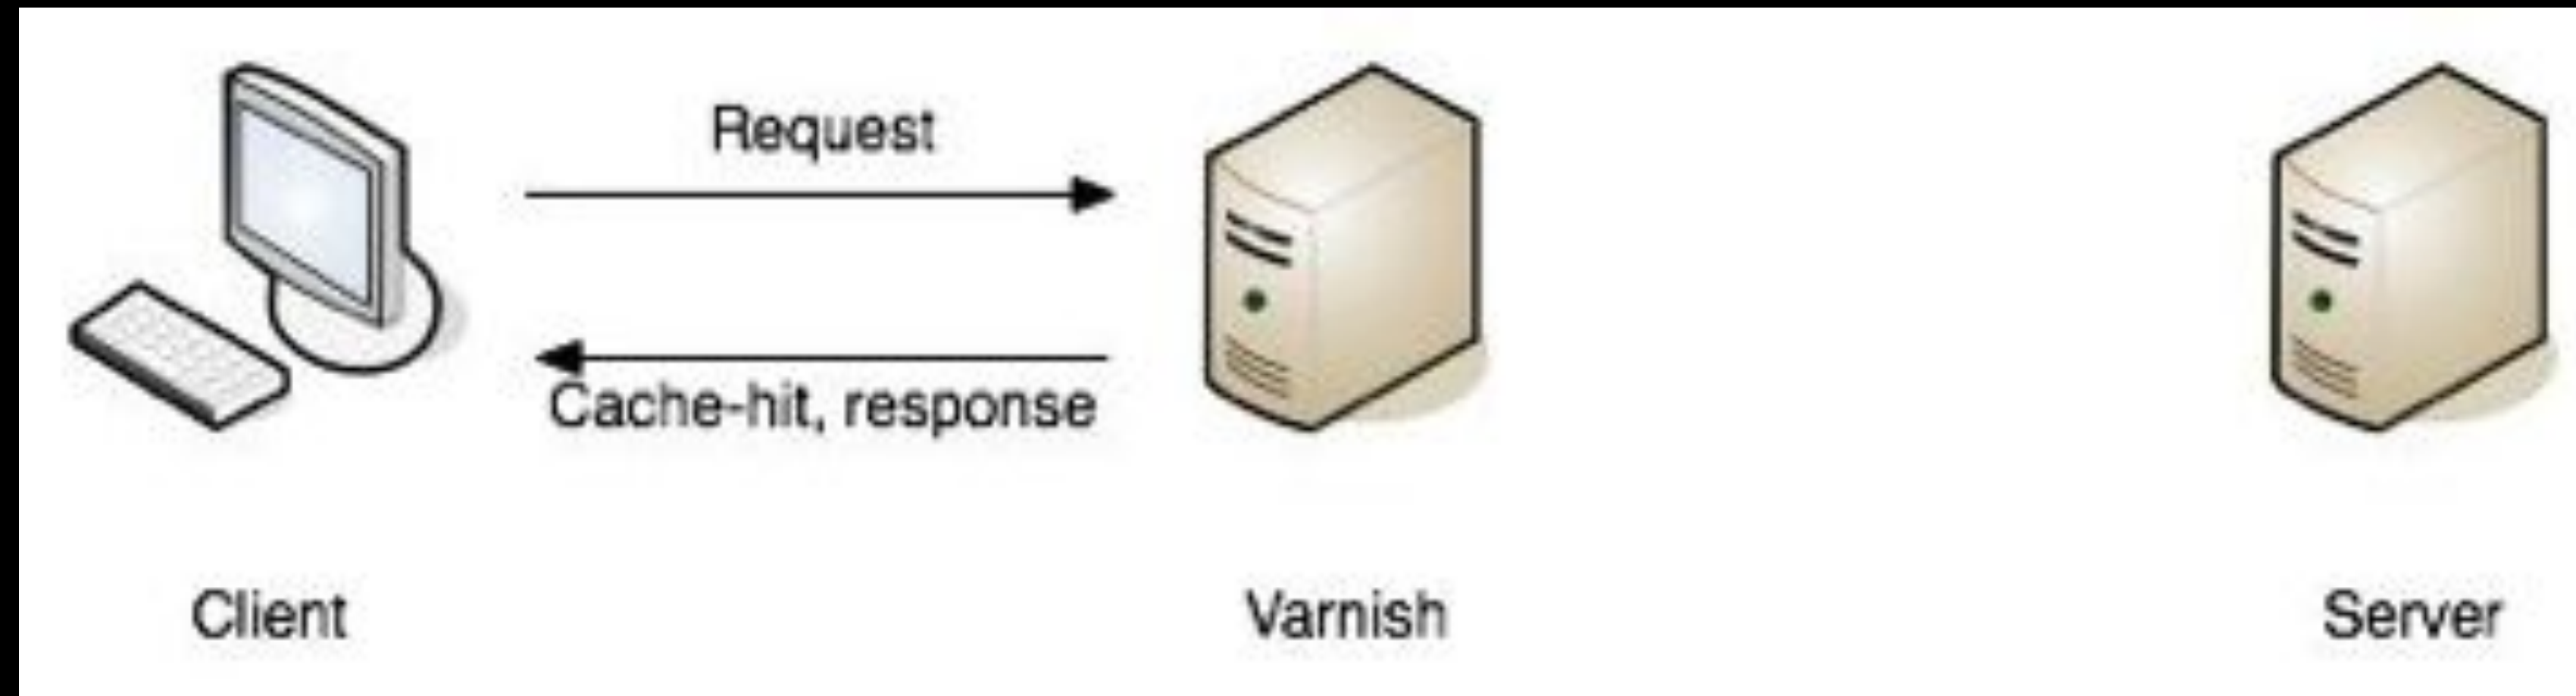
PROBLEM



The HTML response is too slow

- **WHY? - The requested URL is not in Varnish**

What does Varnish do?



<http://book.varnish-software.com/4.0/chapters/HTTP.html>

PROBLEM



The HTML response is too slow

- **WHY? - The requested URL is not in Varnish**
- **WHY? - Bad or limited HTTP headers**

Varnish relies on HTTP headers



verishcheck.getpantheon.com/misionbicycle.com

PANTHEON

Is Varnish working on my Pantheon Site?

Excellent! Varnish is serving cached content for <http://misionbicycle.com> - re-run this report at <http://varnishcheck.getpantheon.com/varnish/misionbicycle.com>

Result details

Header	Value	Details
Via	1.1 varnish-v4	Used by proxies to indicate the intermediate protocol and recipient; in short, the request went through Varnish. This header will always be shown, regardless of whether Varnish served cached content.
Age	6 min 26 sec	The amount of time since the response was generated at the origin server. As the age was greater than 0, it was served from Varnish cache. As Pantheon has multiple Varnish servers and the cached content is not synchronized between them, each Varnish server will have a different age.
X-Varnish	126684301 123212471	This request was served from Varnish cache (both the ID of the current request and ID of the request that populated the cache are shown).
X-Pantheon-Style-Hostname	stys004403c21	Pantheon uses a rotating pool of Varnish servers. Varnish does not have a shared pool or cache, so that means there is a distinct cache for each server. While local DNS typically picks a route and keeps using it, it is possible to access a different Varnish server and experience a cache miss.
HTTP/1.1	200 OK	
Cache-Control	public, max-age=100	

Drupal core settings

Performance ☆

[Home](#) » [Administration](#) » [Configuration](#) » [Development](#)

▼ CLEAR CACHE

Clear all caches

▼ CACHING

Note: Drupal provides an internal page cache module that is recommended for small to medium-sized websites.

Page cache maximum age

15 min ▼

The maximum time a page can be cached by browsers and proxies. This is used as the value for max-age in Cache-Control headers.

▼ BANDWIDTH OPTIMIZATION

External resources can be optimized automatically, which can reduce both the size and number of requests made to your website.

- Aggregate CSS files
- Aggregate JavaScript files

Save configuration

PROBLEM



The HTML response is too slow

- **WHY?** - The requested URL is not in Varnish
- **WHY?** - Bad or limited HTTP headers
- **AND SO** - Varnish is purged frequently
- **WHY?** - We doesn't know where a node is used

PROBLEM



The HTML response is too slow

- **WHY? - The site has no reverse-proxy or CDN**
- **WHY? - It is running on custom servers**
- **AND - No internal expertise**
- **WHY? - internal politics**

Solved problems



- At what number of sites is it cost effective to manage your own Varnish config?
- What decisions will make your site faster and your team faster?

PROBLEM



Drupal responds slowly - why?

- Too many files, too many modules
- Too many queries, slow queries
- Too much memory is used, too many nodes
- Too many function calls, too many calls to the same function

PROBLEM



Drupal responds slowly

- **WHY? - Too many queries**
- **WHY? - We don't know**

Use devel module



Executed 121 queries in 95.11 ms. Queries exceeding 1 ms are highlighted

no	#	status	sql	query	type
38.87	2	Success	CREATE TABLE cache_term (cid, serialized, created, expires, data) VALUES (cid_insert_placeholder_4, cid_insert_placeholder_1, cid_insert_placeholder_2, cid_insert_placeholder_3, cid_insert_placeholder_4)	default	
4.87	3	Success	CREATE TABLE cache_term (cid, serialized, created, expires, data) VALUES (cid_insert_placeholder_4, cid_insert_placeholder_1, cid_insert_placeholder_2, cid_insert_placeholder_3, cid_insert_placeholder_4)	default	
1.38	4	Success	SELECT source, alias FROM url_alias WHERE source IN (:source_0, :source_1, :source_2, :source_3, :source_4, :source_5, :source_6, :source_7, :source_8, :source_9, :source_10, :source_11, :source_12, :source_13, :source_14, :source_15, :source_16, :source_17, :source_18, :source_19, :source_20, :source_21, :source_22, :source_23, :source_24, :source_25, :source_26, :source_27, :source_28, :source_29, :source_30, :source_31, :source_32, :source_33, :source_34, :source_35, :source_36, :source_37, :source_38, :source_39, :source_40, :source_41, :source_42, :source_43, :source_44, :source_45, :source_46, :source_47, :source_48, :source_49, :source_50, :source_51, :source_52, :source_53, :source_54, :source_55, :source_56, :source_57, :source_58, :source_59, :source_60, :source_61, :source_62, :source_63, :source_64, :source_65, :source_66, :source_67, :source_68, :source_69, :source_70, :source_71, :source_72, :source_73, :source_74, :source_75, :source_76, :source_77, :source_78, :source_79, :source_80, :source_81, :source_82, :source_83, :source_84, :source_85, :source_86, :source_87, :source_88, :source_89, :source_90, :source_91, :source_92, :source_93, :source_94, :source_95, :source_96, :source_97, :source_98, :source_99) ORDER BY fit DESC LIMIT 4, 1	default	
1.44	5	Success	SELECT * FROM user_router WHERE path IN (:processors_0, :processors_1, :processors_2) ORDER BY fit DESC LIMIT 4, 1	default	
1.41	6	Success	SELECT alias FROM url_alias WHERE source = :source AND language IN (:language, :language_name) ORDER BY language ASC, cid DESC	default	
1.04	7	Success	SELECT DISTINCT user AND id AND role FROM user_roles WHERE user_id = :user_id AND role_id = :role_id AND status = 1 AND created > :created AND expires > :expires AND serialized = 1 AND data = :data AND cid = :cid AND cid2 = :cid2 AND cid3 = :cid3 AND cid4 = :cid4 AND cid5 = :cid5 AND cid6 = :cid6 AND cid7 = :cid7 AND cid8 = :cid8 AND cid9 = :cid9 AND cid10 = :cid10 AND cid11 = :cid11 AND cid12 = :cid12 AND cid13 = :cid13 AND cid14 = :cid14 AND cid15 = :cid15 AND cid16 = :cid16 AND cid17 = :cid17 AND cid18 = :cid18 AND cid19 = :cid19 AND cid20 = :cid20 AND cid21 = :cid21 AND cid22 = :cid22 AND cid23 = :cid23 AND cid24 = :cid24 AND cid25 = :cid25 AND cid26 = :cid26 AND cid27 = :cid27 AND cid28 = :cid28 AND cid29 = :cid29 AND cid30 = :cid30 AND cid31 = :cid31 AND cid32 = :cid32 AND cid33 = :cid33 AND cid34 = :cid34 AND cid35 = :cid35 AND cid36 = :cid36 AND cid37 = :cid37 AND cid38 = :cid38 AND cid39 = :cid39 AND cid40 = :cid40 AND cid41 = :cid41 AND cid42 = :cid42 AND cid43 = :cid43 AND cid44 = :cid44 AND cid45 = :cid45 AND cid46 = :cid46 AND cid47 = :cid47 AND cid48 = :cid48 AND cid49 = :cid49 AND cid50 = :cid50 AND cid51 = :cid51 AND cid52 = :cid52 AND cid53 = :cid53 AND cid54 = :cid54 AND cid55 = :cid55 AND cid56 = :cid56 AND cid57 = :cid57 AND cid58 = :cid58 AND cid59 = :cid59 AND cid60 = :cid60 AND cid61 = :cid61 AND cid62 = :cid62 AND cid63 = :cid63 AND cid64 = :cid64 AND cid65 = :cid65 AND cid66 = :cid66 AND cid67 = :cid67 AND cid68 = :cid68 AND cid69 = :cid69 AND cid70 = :cid70 AND cid71 = :cid71 AND cid72 = :cid72 AND cid73 = :cid73 AND cid74 = :cid74 AND cid75 = :cid75 AND cid76 = :cid76 AND cid77 = :cid77 AND cid78 = :cid78 AND cid79 = :cid79 AND cid80 = :cid80 AND cid81 = :cid81 AND cid82 = :cid82 AND cid83 = :cid83 AND cid84 = :cid84 AND cid85 = :cid85 AND cid86 = :cid86 AND cid87 = :cid87 AND cid88 = :cid88 AND cid89 = :cid89 AND cid90 = :cid90 AND cid91 = :cid91 AND cid92 = :cid92 AND cid93 = :cid93 AND cid94 = :cid94 AND cid95 = :cid95 AND cid96 = :cid96 AND cid97 = :cid97 AND cid98 = :cid98 AND cid99 = :cid99) ORDER BY user_created DESC LIMIT 10 OFFSET 0	default	
0.84	8	Success	SELECT * FROM user_router WHERE path IN (:processors_0, :processors_1, :processors_2) ORDER BY fit DESC LIMIT 4, 1	default	
0.87	9	Success	SELECT * FROM cache_term WHERE (cid = :cid_insert_placeholder_4)	default	
0.82	10	Success	SELECT source, alias FROM url_alias WHERE source = :source AND language IN (:language, :language_name) ORDER BY language ASC, cid DESC	default	
0.81	11	Success	SELECT alias FROM url_alias WHERE source = :source AND language IN (:language, :language_name) ORDER BY language ASC, cid DESC	default	

PROBLEM



Drupal responds slowly

- **WHY? - Too many queries**
- **WHY? - Too many nodes loaded from the DB**

Caching, caching, caching

Entity cache module +



redis

PROBLEM



Drupal responds slowly

- WHY? - ~~Too many queries~~
- WHY? - One really slow query

Views module: A leaky abstraction



Displays

System [Add](#) Edit View Name/Description

System details

Display name: System View System

TITLE
Title: Content

FORMAT
Format: Table | [Settings](#)

FIELDS [Add](#)

- Bulk operations: Content
- Content: Title (Title)
- Content: Has new content
- Content: Type (Type)
- (author) User: Name (Author)
- Content: Published (Published)
- Content: Updated date (updated)
- Content: Edit link (Operations)
- Content: Delete link

FILTER CRITERIA [Add](#)

- Content: Title (exposed)
- Content: Type (exposed)
- (author) User: Name (exposed)
- Content: Published (exposed)
- (term) Taxonomy terms: Vocabulary (exposed)

SORT CRITERIA [Add](#)

HEADER [Add](#)

FOOTER [Add](#)

PAGER
Use pager: Full | [Paged, 50 items](#)

Advanced

SYSTEM PATH SETTINGS
Path: [admin/content](#)

CONTEXTUAL FILTERS [Add](#)

- Content: Author
- Content: Taxonomy terms on node

RELATIONSHIPS [Add](#)

NO RESULTS BEHAVIOR [Add](#)

- Global: Unfiltered text

EXPOSED FORM
Exposed form in block: No
Exposed form style: Basic | [Settings](#)

OTHER
Machine Name: system_1
Comment: No comment
Use AJAX: Yes
Hide attachments in summary: No
Hide contextual links: Yes
Use aggregation: No
Query settings: [Settings](#)
Field language: Current user's language
Caching: None
CSS class: admin-views-view
Theme: Information

Turn on the timer!



```
Query      SELECT DISTINCT node.nid AS nid, node.title AS node_title, history.timestamp AS history_timestamp, node.created AS node_created, node.changed AS node_changed,
node.type AS node_type, users_node.name AS users_node_name, users_node.uid AS users_node_uid, node.status AS node_status
FROM
(node) node
LEFT JOIN (users) users_node ON node.uid = users_node.uid
LEFT JOIN (taxonomy_index) taxonomy_index ON node.nid = taxonomy_index.nid
LEFT JOIN (taxonomy_term_data) taxonomy_term_data_node ON taxonomy_index.nid = taxonomy_term_data_node.nid
LEFT JOIN (history) history ON node.nid = history.nid AND history.uid = '1'
ORDER BY node_changed DESC
LIMIT 50 OFFSET 0

Title      Content
Path       administration
Query build time 12.65 ms
Query execute time 1.72 ms
View render time 906.91 ms
```


PROBLEM

Drupal responds slowly

- WHY? - ~~Too many queries~~
- WHY? - One really slow query
- WHY? - One really complex requirement



If I had only one hour to save the world, I would spend fifty-five minutes defining the problem, and only five minutes finding the solution.”

— Probably not Albert Einstein

<http://quoteinvestigator.com/2014/05/22/solve/>

PROBLEM



The server responds slowly

- **WHY? - Varnish can't cache pages reliably**
- **WHY? - Poor page-level cache metadata**
- **WHY? - Poor element-level cache metadata**
- **BUT! - Drupal 8 has the solution**

Cache metadata



Watch Wim Leers and Fabian Franz at DrupalCon: <http://wimleers.com/talk/making-drupal-fly-fastest-drupal-ever-here>

Chunks vs. blobs 2.0



Cache metadata

- Cache tags (data dependencies)
- Cache contexts (global dependencies)
- Cache max-age

Cache metadata in render arrays



```
$build[ '#cache' ][ 'tags' ][] = 'node:123' ;  
$build[ '#cache' ][ 'contexts' ][] = 'user.permissions' ;  
$build[ '#cache' ][ 'max-age' ][] = Cache::PERMANENT ;
```



I'm rendering something. That means I must think about cacheability!"

— Wim Leers

<http://wimleers.com/talk-making-drupal-fly-fastest-drupal-ever-here/#/3/5>

PROBLEM



Drupal responds slowly

- **WHY? - Too many modules**

Can you have too many modules?



Filter list	Filter list	Filter list	Filter list
Name	Description		URL
<input checked="" type="checkbox"/> Address permissions (000) (admin, permissions)	Provides per-module-level access control for admin. Used by these filter operations.		
<input type="checkbox"/> Address field (admin, field)	Storage & flexible address field implementation for sites. Module B.		
<input checked="" type="checkbox"/> Admin users	Full support for Drupal admin and managers.		
<input checked="" type="checkbox"/> Administrative Page Management tools (admin, tool)	Administration and debugging functionality for developers and site builders.		
<input checked="" type="checkbox"/> Administrative menu (admin, menu)	Provides a Drupal menu to most administrative tasks and other common instructions for users with the proper permissions.		Permissions, 1
<input checked="" type="checkbox"/> Administrative menu (Admin) style (admin, menu, style)	Admin menu toolbar		
<input type="checkbox"/> Administrative menu (admin, menu)	Provides Drupal admin page management pages in Drupal with rich navigation.		
<input checked="" type="checkbox"/> Aggregator (aggregator)	Aggregated syndicated content (RSS, REX, and more feeds).		
<input checked="" type="checkbox"/> App icon	Provides interface for mobile Apps. Not a Drupal core component.		
<input checked="" type="checkbox"/> Authenticated (authenticated, access)	Implements multiple role support based on the user's Drupal membership.		
<input type="checkbox"/> Backport (backport)	Full support for backport's Not a Drupal B.		
<input checked="" type="checkbox"/> Block (block)	Controls the visual building blocks a page is constructed with. There are pages of content, content items as well, or region of a web page.		Permissions, 1
<input checked="" type="checkbox"/> Blog (blog)	Creates full web blogs.		
<input checked="" type="checkbox"/> Book (book)	Allows content items and region maintenance in an online.		
<input checked="" type="checkbox"/> Bookmarks (bookmarks)	Storage in website.		Permissions, 1
<input checked="" type="checkbox"/> Book report (book, report)	Performs full reporting of data objects from about by these tools.		

PROBLEM



Drupal responds slowly

- **WHY? - Too many modules**
- **WHY? - Overlapping/duplicate modules**
- **WHY? - Lack of clarity/history among team**

Treat “slow” like a bug



- Do you have tests for the module's functionality?
- Can you look through your git history?
- Project management tickets?
- Developer documentation on modules?

PROBLEM



Drupal responds slowly

- **WHY - “slowly” was not defined upfront.**
- **WHY - The team building the site struggles to define many requirements in advance.**
 - **User personas, acceptance criteria, behavioral testing / unit testing strategies.**

For any part of your site



How do you know:

- Is it secure?
- Is it accessible?
- Does it behave as expect?
- Who uses it?
- How fast should it execute?

Client-side



What's on the agenda?



Name	Status	Type	Initiator	Size	Time	Timeline - Start Time	200ms	1.00s	A
checkbox dashboard	200	docum...	Other	20.3 KB	954 ms				
checkbox css__UmjysVANCO9yfbY2cbw...	200	stylesh...	dashboard:11	7.1 KB	117 ms				
checkbox css__QPp1snPjYU_xYr0c_V0IOR...	200	stylesh...	dashboard:12	21.6 KB	481 ms				
checkbox js__79M6UrZjAw3oNCnUjyWip...	200	script	dashboard:23	27.4 KB	1.14 s				
checkbox js__HSAFmq1RrXKsrazzPq_xP...	200	script	dashboard:56	7.5 KB	192 ms				
checkbox js__jY3-3CF_aW2-zV7M-ZJG0...	200	script	dashboard:57	8.6 KB	589 ms				
checkbox js__MU0c1KbcCpTCz8zYuPGZ...	200	script	dashboard:603	12.4 KB	599 ms				
checkbox js__Wnfmj39M4876ad0c3jcc7...	200	script	dashboard:604	8.1 KB	263 ms				
checkbox js__gUyfkFUH9LkUwfrAg-ccCal...	200	script	dashboard:605	14.0 KB	343 ms				
checkbox bg-blue-blueprint.jpg	200	jpeg	dashboard:603	3.3 KB	219 ms				
checkbox sprites.png	200	png	dashboard:603	13.0 KB	295 ms				
checkbox small-icons.png	200	png	dashboard:603	2.7 KB	175 ms				
checkbox js__ZA8k5CE4R6vkLKcvmLJHuP...	200	script	dashboard:606	53.1 KB	1.01 s				
checkbox js__hHUWfbMAa11eJtx5V0ym...	200	script	dashboard:607	1.6 KB	105 ms				
checkbox icon-w-search.svg	200	svg+xml	dashboard:115	1.9 KB	154 ms				
checkbox nr-768.min.js	200	script	dashboard:9	9.0 KB	255 ms				
checkbox e671e252557a=4266189&pi=...	200	script	nr-768.min.js:1	138 B	74 ms				

17 requests | 212 KB transferred | Finish: 3.23 s | DOMContentLoaded: 2.89 s | Load: 2.88 s

Look behind the numbers



sitespeed.io

32 pages analyzed for <http://drupal.org>

Test performed Thu Jan 28 2016 19:18:45 GMT-0500 (EST) with sitespeed.io-desktop rules using a cable connection.

User-Agent: Mozilla/5.0 (Macintosh; Intel Mac OS X 10_9_4) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/47.0.2526.120 Safari/537.36 Viewport: 1280x800

Rule Score
85 (92)

Critical Rendering Path Score
90.0 (90.0)

Number of JS synchronously inside head
0.00 (2.00)

Number of JS files per page
8.00 (12.00)

Number of CSS files per page
2.00 (2.00)

Number of CSS images per page
2.00 (4.00)

Number of font files per page
0.00 (0.00)

Number of images per page
14.00 (28.00)

Number of requests per page
29.5 (47.0)

Requests Without Expires
13.00 (14.00)

Requests without GZip
0.00 (0.00)

Document Weight
25.8 kb (139.3 kb)

JS File Weight Per Page
168.7 kb (315.1 kb)

CSS File Weight Per Page
146.2 kb (158.6 kb)

Total page weight (including all assets)
542.1 kb (1370.7 kb)

Images scaled by the browser
0.0 (3.0)

Number of SPOF per page
0.0 (2.0)

Number of Domains
5.00 (8.00)

Domains with only one request
3.00 (5.00)


Redirects Per Page
19.00 (21.00)

Cache Time
1 day (4 weeks)

PROBLEM

The front-end is too slow

- **WHY?** - The CSS takes too long to load
- **WHY?** - The site has too much CSS
- **UNHELPFUL CONCLUSION** - The website is “fat”



I'd like us to stop talking about "fat sites" not only because that language hurts my feelings, but also because this oversimplification obscures key ways we can make websites faster."

— Marc Drummond

Fast sites do not require insults

<https://the-pastry-box-project.net/marc-drummond/2015-december-21>

PROBLEM

The front-end is too slow

- **WHY? - The CSS takes too long to load**
- **WHY? - The site has too much CSS**
- **WHY? - Unused CSS rules**
- **WHY? - We are afraid to edit our CSS**

Delete CSS with confidence

List of screenshots for shots taken 2015/04/08 07:16:23

Screenshots:

great-vegetables

lovely-vegetables

login

home

great-vegetables

320px



dev



test



diff

0.0 % different

great-vegetables

600px



dev



test




diff

1.83 % different



WRAITH

<https://pantheon.io/blog/test-website-wraith>



Tricks like trying to load a bunch of stuff in parallel, or aggressive caching, might seem like appealing shortcuts, but nothing replaces just auditing the code and making it need less ‘stuff.’”

— Craig Silverstein

Forgo JS Packaging? Not so Fast

<http://calendar.perfplanet.com/2015/forgo-js-packaging-not-so-fast/>

Do you know where your JS is coming from?

```
/**
 * Add JS to every page. Used only on home page.
 */
function custom_js_adder_init() {
  drupal_add_js('has_not_been_audited_since_2011.js');
}
```

Attach libraries to render arrays

```
/**
 * Add JS to every page. Used only on home page.
 */
function custom_preprocess_node(&$variables) {
  &$variables['#attached']['library'][] = 'nameof/lib';
}
```

When you have a lot of JS/CSS files...



- Serve all individually
 - Each cached separately
 - Lots of HTTP request
- Package all CSS/JS per web page
 - 1 HTTP request
 - Poor caching
- Split the difference (Drupal core)

Where is the decision implemented?



PROBLEM



The front-end is too slow!

- **WHY? - Too many third party widgets**



whenever the parser encounters a script it has to stop and execute it before it can continue parsing the HTML... avoid and minimize the use of blocking JavaScript, especially external scripts that must be fetched before they can be executed.”

—Google developer documentation

<https://developers.google.com/speed/docs/insights/BlockingJS>

JavaScript share widgets



Site	Files	Size
Google+	1	15.1KB
Facebook	3	73.3KB
LinkedIn	2	47.7KB
Pinterest	3	12.9KB
Tumblr	1	1.5KB
Twitter	4	52.7KB
Total	14	203.2KB

March 2015: <https://jonsuh.com/blog/social-share-links/>

PROBLEM

The front-end is too slow!

- **WHY? - Too many third party widgets**
- **WHY? - No one feels responsible for them**



Organizations which design systems ... are constrained to produce designs which are copies of the communication structures of these organizations.”

— Melvin Conway (Conway's Law)

PROBLEM

The front-end is too slow!

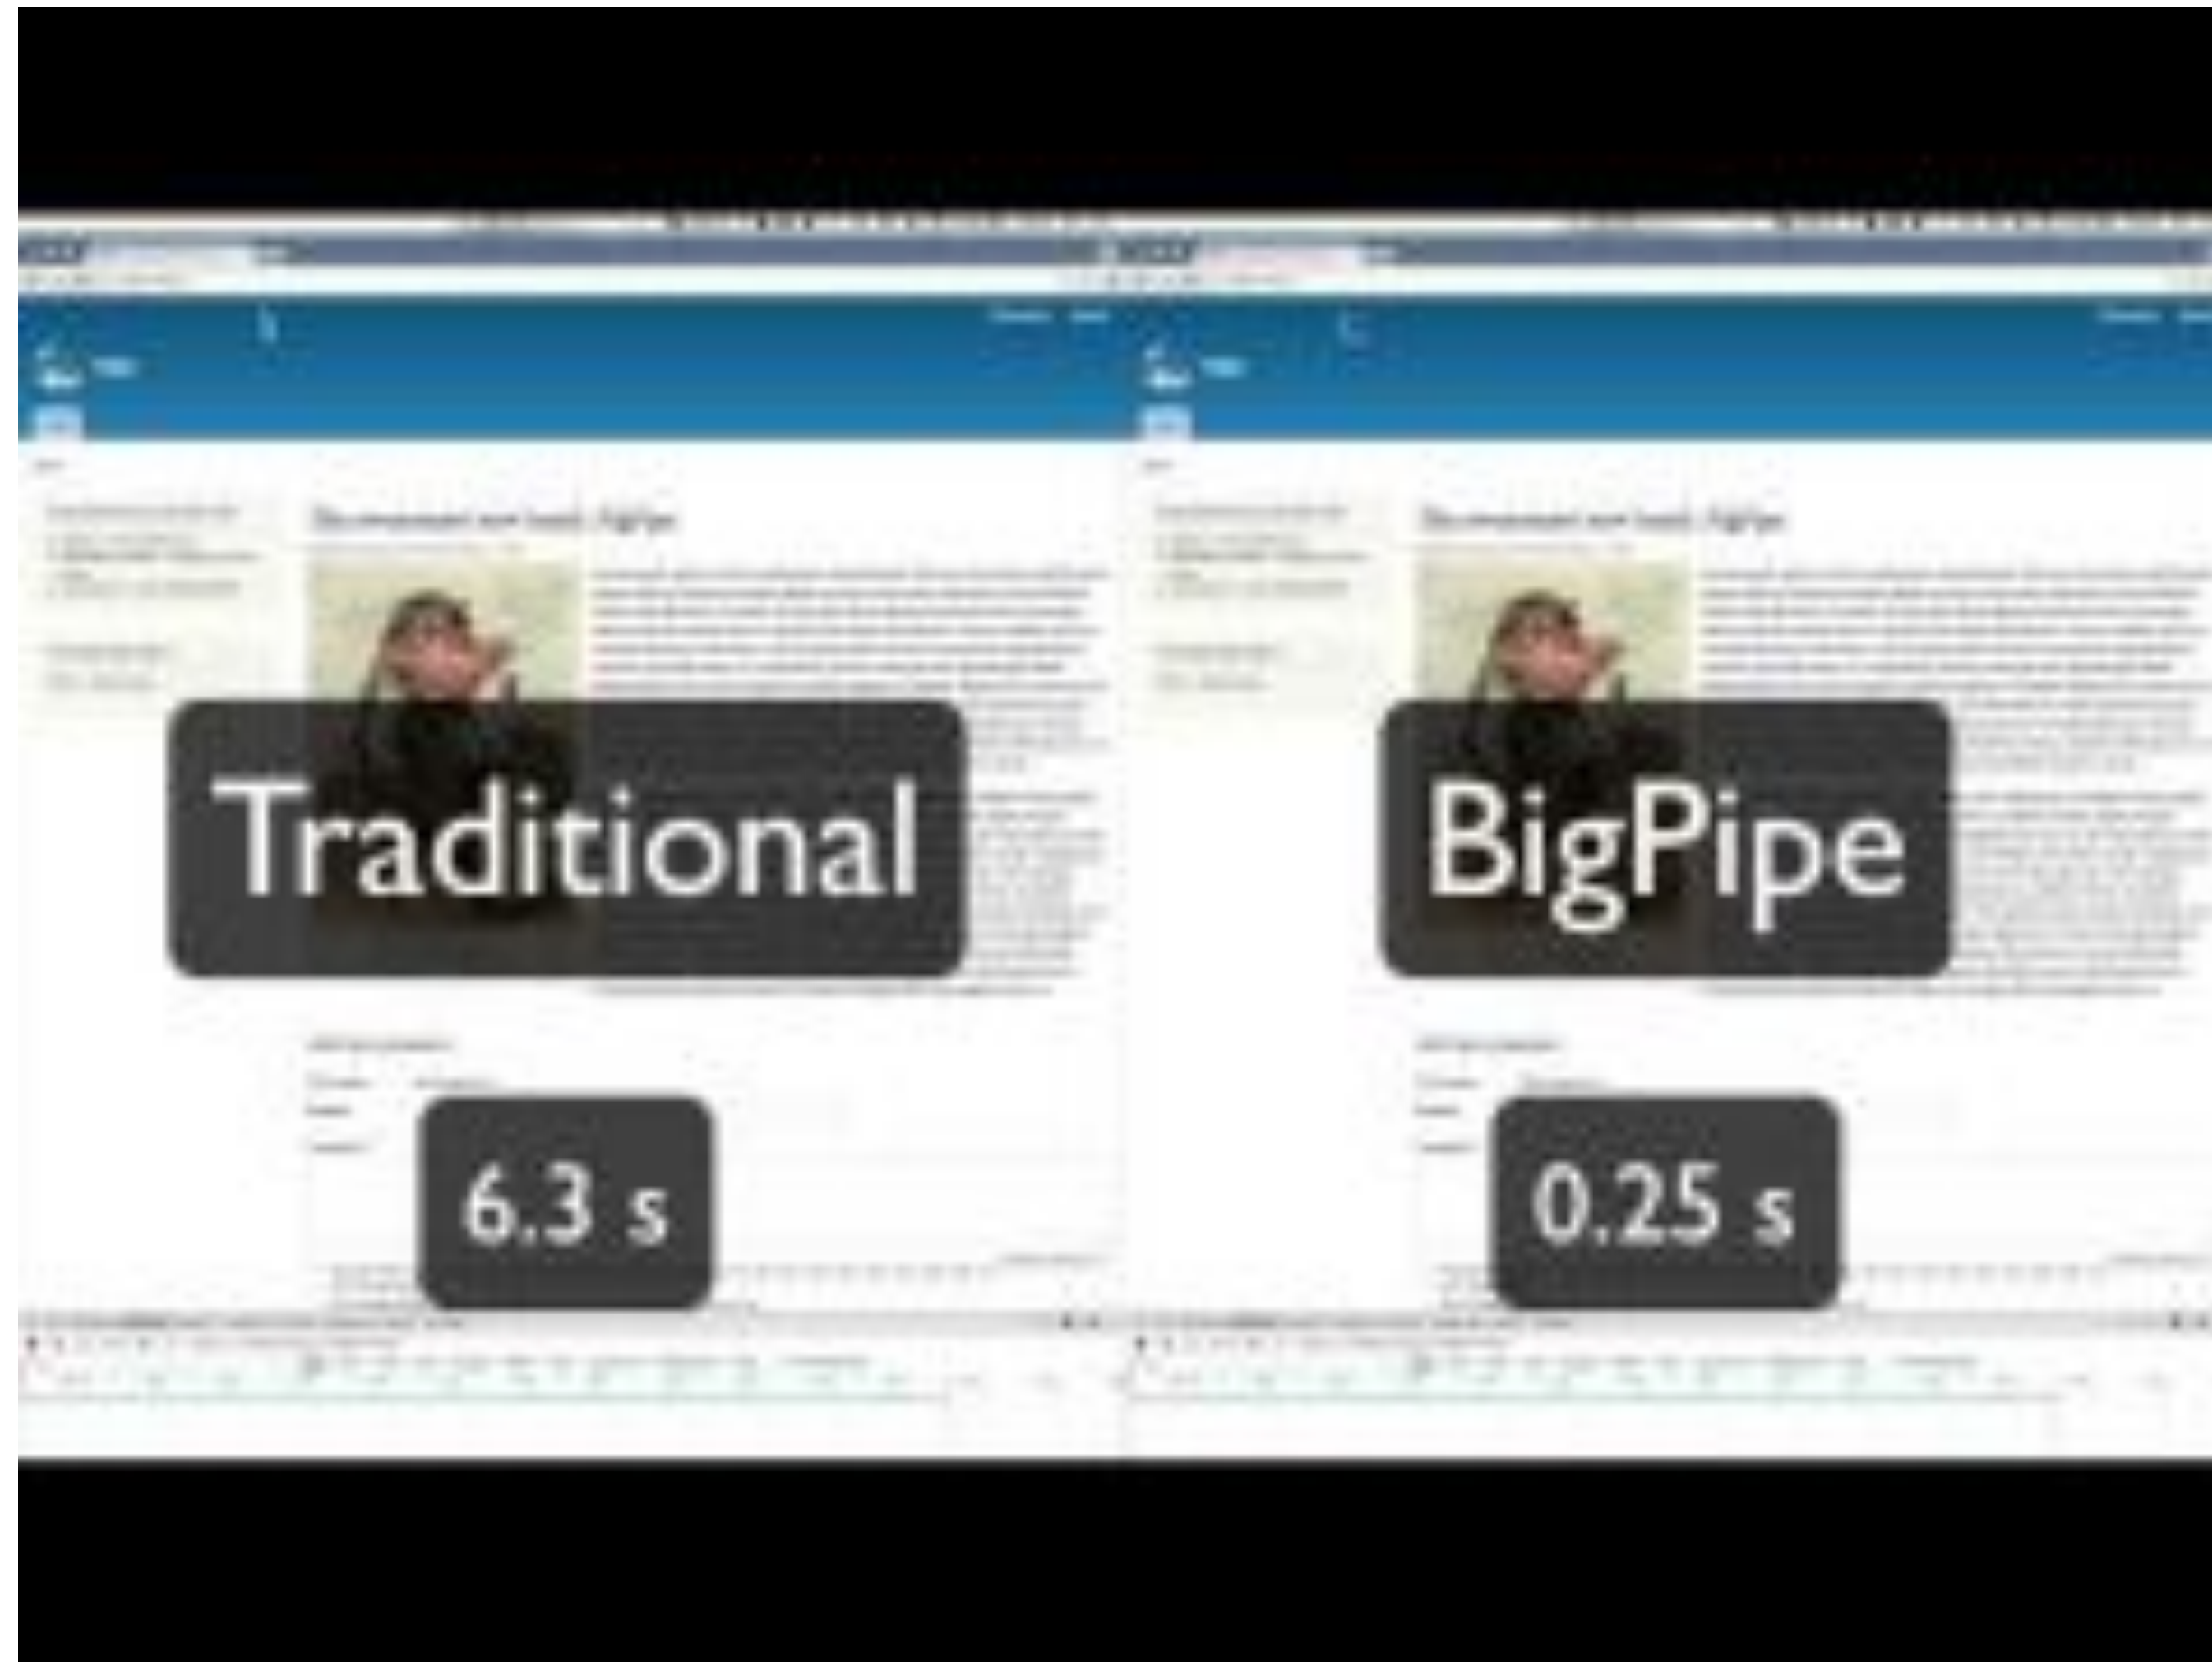
- **WHY? - Too many third party widgets**
- **WHY? - Two sets of analytics tracker codes**
- **WHY? - Two groups of people analyze traffic who don't share any tools.**

When you're in a meeting do you...



- Add agenda items ad-hoc
- Have factions who ignore each other
- Interrupt each other
- Repeat topics from previous meetings
- Start the meeting late

How quickly can you get to a usable site?



You are moving in the right direction when...



- You can describe how the website is built
- You can track changes over time
- You estimate the performance impact of a feature
- You start projects with performance limits
- decisions make your site AND team faster
- When performance is a boring topic

**Thanks,
- @stevector**

