



DUBLIN

DRUPALCON





4x High Performance for Drupal - Step by Step

Fabian Franz (Fabianx) - @fabianfranz

Performance Track

4x High Performance for Drupal

Presented by Fabian Franz

Step by Step

Your BOSS is calling!

It happens to the best
of us ...

Especially during
DrupalCon ...

... or during elections.

The site goes down,
the site is slow,
grab a tut',
and make it grow!

But first ...

Lets start with a little
story ...

„Where is the power
of Drupal ... I hate
Drupal!“

„... I hate Drupal!“

„[...] and it always overload the database (mysql) some times the load reach 200 and it never reach it before ... I enabled core cache and the server got down again next day ..."

- groups.drupal.org/node/158244

My site is so slow ...
HEEEEEEELP!

- "Where is the power of Drupal ... I hate Drupal!"
- This is really sad ...



My site is so slow ...
HEEEEEEELP!

- There is a need for High Performance Drupal!
- Faster sites earn more money
- Faster sites get ranked higher by Google

My site is so slow ...
HEEEEEEELP!

- Visitors love fast sites
- Mentioned in the media? What if your server goes down exactly then?

BUT ...

... THE QUESTION IS ...

How do I get a
blazingly fast site?

Doing it Wrong

Doing it Wrong

"Okay, I have now tweaked my Sauerkraut* settings, but the site is still slow. What Sauerkraut settings do I need to tweak so that it is as fast as xyz.com?"

* Sauerkraut was APC in this case, but we come to that later

Doing it Wrong

„I have setup 10 Slave DB Servers, but
once I test the site it is sooo sloooooow!“

Doing it Wrong

„I have setup NGINX with AdvAgg and Varnish combined with Entitycache and Views_Opt_Cache - still the performance remains the same. :-“(

„Have you setup
Memcache?

???

Doing it Wrong

“I have setup static page caching for all the pages. The high traffic day can come! What could possibly go wrong?”

We all wish ...

... we had ...

THE MAGIC PILL!

Optimization is a process

- not a pill

„Just one pill
and the site is
fast!“



Optimization is a process
- not a pill

4 common ways to fail

1. Optimizing one part to death while neglecting all the others:

"If you build the house on one pillar, it'll not hold long ..."

Optimization is a process
- not a pill

4 common ways to fail

2. Optimizing things without knowing
where the pain is:

"Is the bottleneck MySQL, PHP, Apache or
something else?"

Optimization is a process
- not a pill

4 common ways to fail

3. Optimizing things with new methods
without really understanding them:

"Reinvent the wheel or stand on the
shoulder of giants?"

Optimization is a process
- not a pill

4 common ways to fail

4. Optimizing things without testing it'll hold the load:

"You are featured by BigNews.com - Your server goes down."

4 Common Ways to Fail

1. Optimizing one part to death
2. Optimizing just random parts
3. Optimizing parts with <BuzzWord>
4. Optimizing without testing

"Ouch, that are lots of
ways to fail ..."

„That is all so
complicated.“

* sigh *

This is all so
complicated. *sigh*

"Is there nothing I can do
to make this easier and
have a fast site?"

This is all so
complicated. *sigh*

- The easy answer:

Hire a performance consultant.

„This is all so complicated.“ *sigh*

Hire a performance consultant NOW.

CALL NOW IN THIS SECOND:

0800 - DRUPAL PERFORMANCE

and enjoy blazingly fast sites.

The End

Now you know Performance is really difficult to get right and that you should hire a Performance Consultant.

Remember this number:

0800 - DRUPAL PERFORMANCE

QUESTIONS?

JUST KIDDING ...

- Okay, okay. You got me ...
"Hiring a performance consultant can be really useful at times, but even more useful is learning and spreading the knowledge."

JUST KIDDING ...

4x HIGH PERFORMANCE FOR DRUPAL -
STEP BY STEP

*„Lets stand on the shoulders of Giants
and walk the paths of our ancestors!“*

Your Mission

Loading your
mission . . .

The Mission

- Drupal 7 / 8
- Several Performance Problems
- ... real life problems!

Lets meet some
friends ...

... and help them in
their need ...

Your Mission

D. Pages feel
slow, sluggish
and big ...

... and are totally
unhappy :-)

„This is sooooo
heavy load.“



Your Mission

Mrs. MySQL is exhausted and needs a time out

:-(*sigh*):-

„I just need a
SELECT break.“



Your Mission

Mr. Apache is sweating under the load :-((

„I give 100% all the time, but this is just too much.“



Your Mission

Mr. Code is
buggy and a
real trouble
maker ;-)

„Yeah!

He he he!“



Your Mission

- D. Pages feel slow and are unhappy :-((
- Mrs. MySQL is exhausted and needs a time out :-((
- Mr. Apache is sweating under the load :-(((
- Mr. Code is buggy and a real trouble maker ;-)

The Task

- Investigate and Fix!
- Let's go!



1. Server Performance

Measuring Server Performance....

Measuring Server Performance....

- System Load: 4.14
- Page Load Time: 20 sec
- Apache Load: 100%

How to measure Performance on Server?

- *top* command

```
load average: 4.14, 1.40, 0.53
,   1 stopped,   0 zombie
,   0.0 wa,   0.0 hi,   0.0 si,   0.0 st
188836 free,   411892 buffers
999192 free. 29910408 cached Mem
```

R	S	%CPU	%MEM	TIME+	COMMAND
0	R	100.0	0.2	1:04.92	apache2
0	R	100.0	0.2	0:22.09	apache2
4	R	100.0	0.0	0:09.09	apache2
2	R	99.2	0.2	0:20.02	apache2
0	R	99.2	0.2	1:06.60	apache2
8	R	99.2	0.1	0:31.56	apache2

8	B	00.5	0.1	0:31.20	gdbcpes
0	B	00.5	0.5	1:00.00	gdbcpes

How to measure Performance on Server?

- Handy Drush Command for page generation time of any page (Drupal 7):

```
time drush php-eval '  
$path="node/17";  
menu_set_active_item($path);  
menu_execute_active_handler($path,  
TRUE);' > /dev/null
```


Uhm, really?

Yes!

Why would I need
that?

Production
debugging!

Sometimes problems
only show up on
production.

So now we know
there is a problem.

How do we solve
them?

The 4 Shoulders of the Giants

While I have said ...

.... that you should
know your pain
points first ...

... there is a “stack”

.... that many high
performance sites
use.

Pressflow / Good code

APC / Opcache

Memcache / Redis

Varnish / NGINX / CDN

Pressflow / Good code

APC / Opcache

Memcache / Redis

Varnish / NGINX / CDN

PHP-FPM (!)

Hey there, I am new!



Pressflow / Good code


APC / Opcache

Memcache / Redis

Varnish / NGINX / CDN

BigPipe + Streaming

Hey there, I am also new!



Uhm, BigPipe,
you have already
been handled
yesterday!

Pressflow / Good code

APC / Opcache

Memcache / Redis

Varnish / NGINX / CDN

PHP-FPM (!) + PHP ?

Hey there, I am new!



So how can
those help me?

Pressflow

- Only really relevant for Drupal 6 sites
- Drupal 7 already includes most Pressflow patches / approaches
- Drupal 8 has performance best practices all around.

Inofficial Pressflow

- <https://groups.drupal.org/node/210683> (Wiki)
- All collected performance patches that are relevant for Drupal 7.
- Drupal 8: All committed already :D.

PHP 7

- Up to 50% faster
- Drupal 8 supports PHP 7 since 8.0.0
- Drupal 7 supports PHP 7 since 7.50

PHP 7

- Just use it and start transforming your stack to it.

- PHP 5.6

End-of-Life
LTS

31/12/2016
31/12/2018

PHP 7: Is it worth it?

- Seriously: It is really really worth it!

Admin
performance:




Image from:

<https://pantheon.io/blog/measuring-php-7-performance-new-relic-nobsbenchmarks> by Josh Koenig

PHP 7

- Most hosters already support it, because it is just **that** much better.

Fun fact:

Drupal 8 was the reason why PHP 7.0.0 release was delayed by 2 weeks.

APC

- Alternative PHP Cache
- Highly recommended (easy to install)
- PHP PECL Extension
- Speeds up PHP execution by caching pre-compiled PHP objects

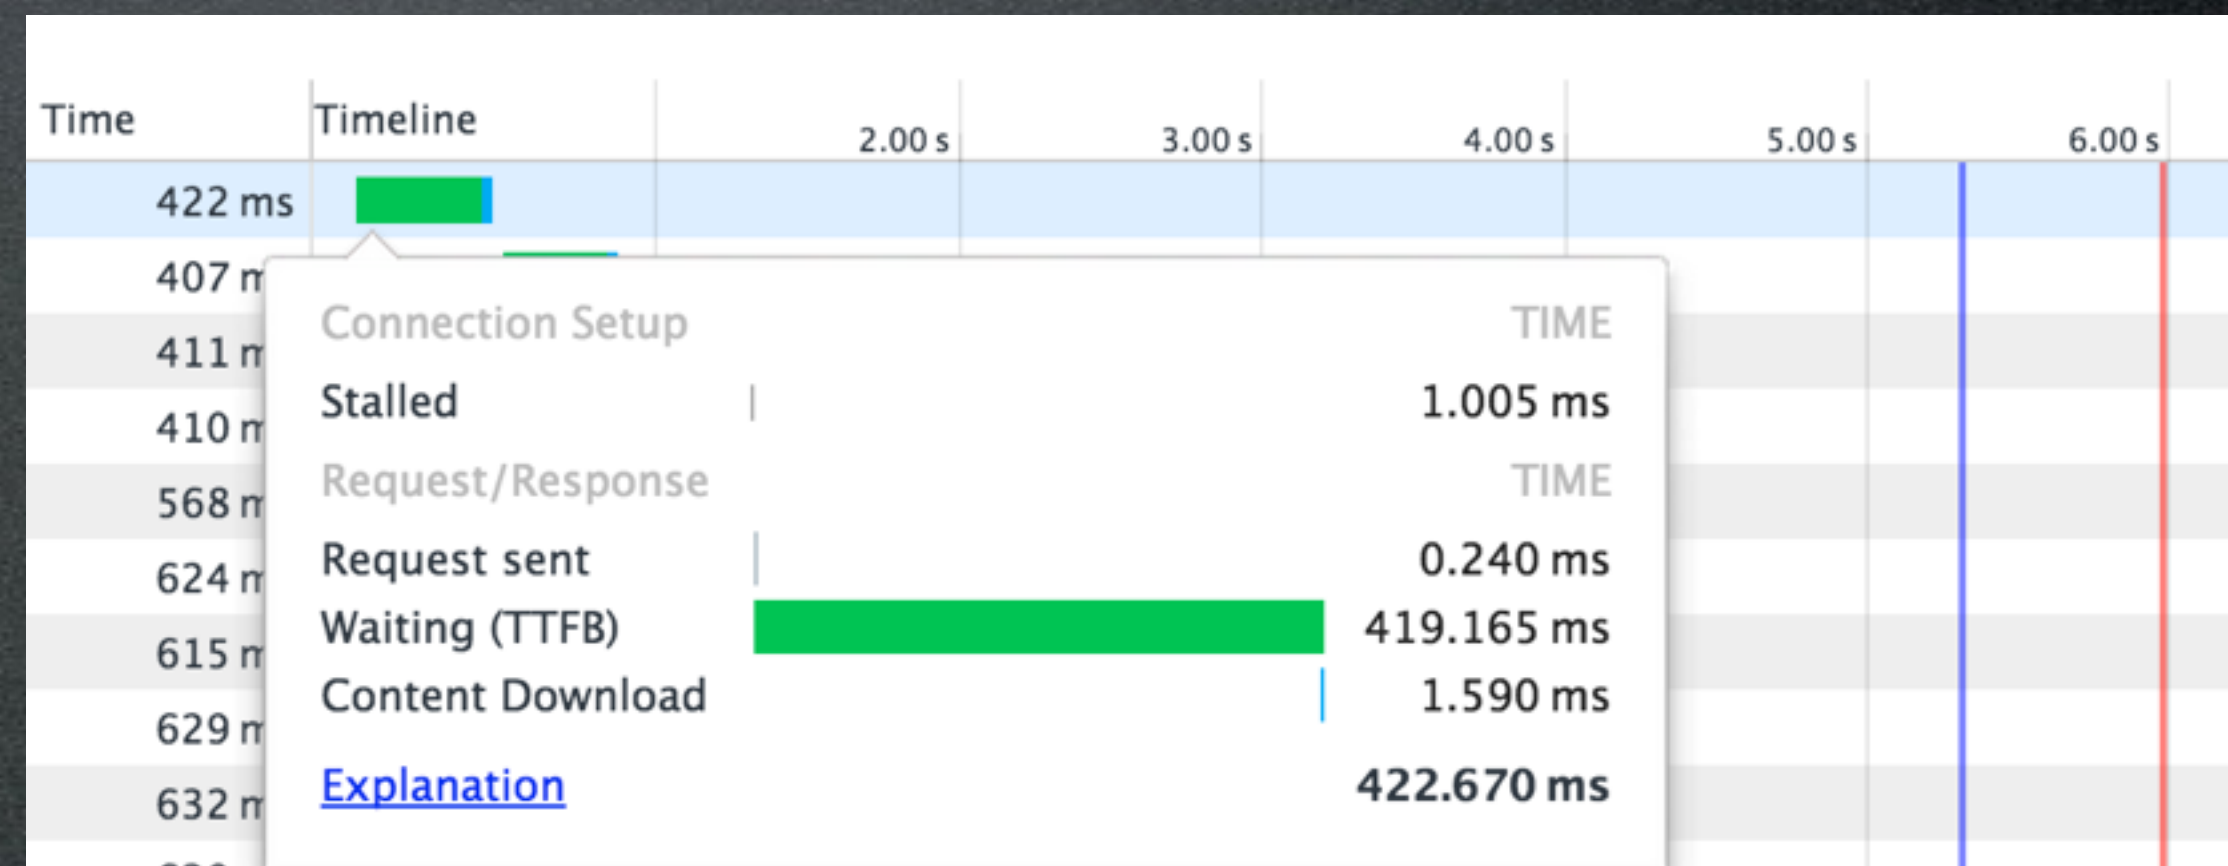
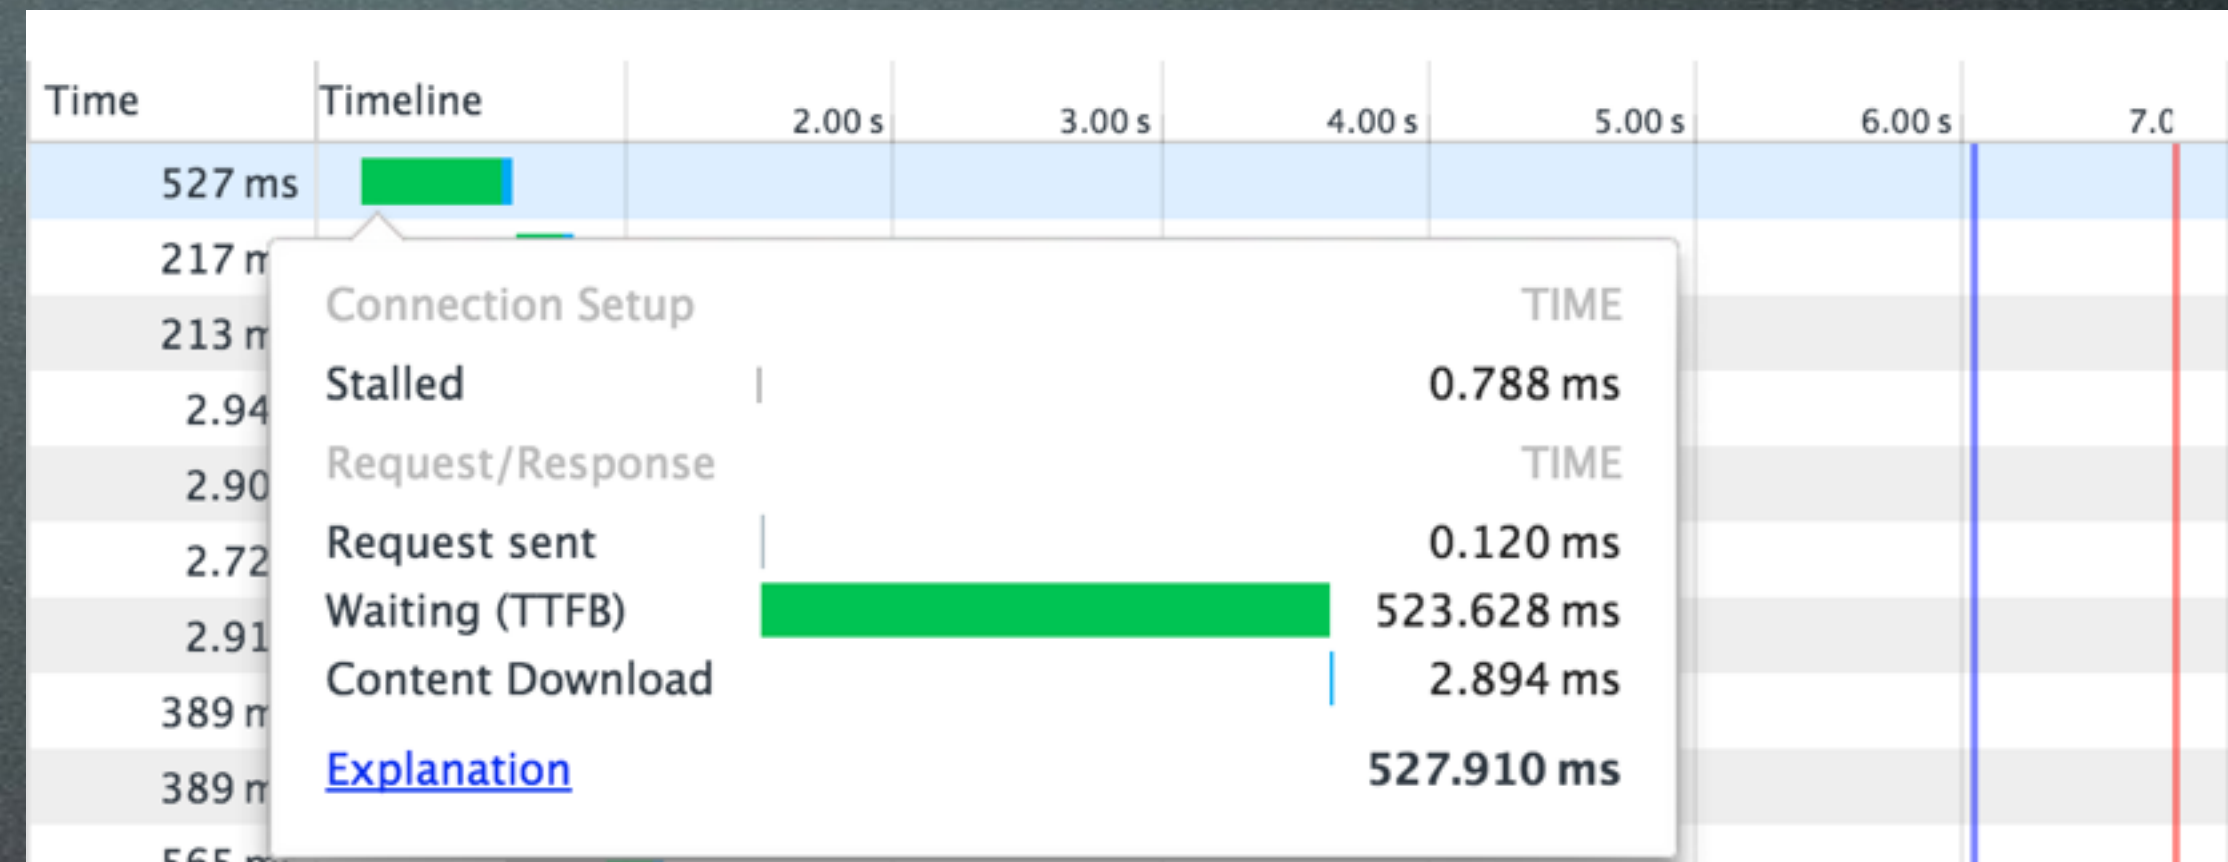
Opcache (PHP \geq 5.5)

- Opcode Cache
- Highly recommended (easy to install)
- In PHP \geq 5.5 by default (Yeah!!!)
- Speeds up PHP execution by caching pre-compiled PHP objects

Opcache: Is it worth
it?

OpCache: Is it worth it?

- Yes
- Saves:
- ~ 100 ms
- ~ 20 MB



OpCache: What can go wrong?

- `opcache_reset()` buggy

=> in all released PHP versions.


- opcache is disabled during restart; can take up to two minutes.

OpCache: What can go wrong?

- Size it properly (if it's full => restart)
- Use php-fpm graceful restart.
- Blog Post about systemd + php-fpm for graceful restart coming soon on tag1consulting.com

OpCache: What can go wrong?

- Look at the stats!
- Opcache going offline can lead to huge spikes!



Memcache/Redis

- Replaces caching in MySQL Database
- Key/Value Store in Main Memory

```
store(Apple, 10)  
get(Apple) == 10
```

- Very Fast!

Memcache/Redis:
What does it get me?

Memcache/Redis: What does it get me?

- Way less load on the database
- Overall faster caches
- Much easier to scale up
(Distributed key-value storage)

LCache (d.go.to/lcache)

- Newest kid on the block (alpha release)
- Key/Value Store within PHP (APCu)

```
store(Apple, 10)  
get(Apple) == 10
```

- Even faster than memcache!

LCache: What does it get me?

- Transaction log + consistent store in one
- No additional service needed (!) - except for APC/APCu extension.
- Presentation online already.

SuperCache

- Similar to lcache / ChainedFast in core
- Especially optimized for write performance and cache tags
- => Caveat: Might have race conditions.

Varnish

- Save whole response to memory
- Serve response from memory
- Like a “shield” for your server
- Varnish is crucial for high throughput.
(unless you use NGINX)

Varnish: That sounds
pretty complicated!

Varnish: That sounds pretty complicated!

- Best practice configurations:
- <https://fourkitchens.atlassian.net/wiki/display/TECH/Configure+Varnish+3+for+Drupal+7>
- <https://www.lullabot.com/blog/article/configuring-varnish-high-availability-multiple-web-servers>

Varnish: What does it get me?

- 50 ms response times!

Your Mission:
UPDATE!

Your Mission: UPDATE!

- Anonymous Pages: „We are blazingly fast, still big, but quite happy“.

Your Mission: UPDATE!

- Mrs. MySQL: „I have less to do now, but if I have it is still too much. Those authenticated users ...“

Your Mission: UPDATE!

- Mr. Apache: „I have much less to do, but when those authenticated users come, I still sweat.“
- „And I hate those anonymous utm_requests!“

Quick Fix for GA Problem

- VCL (Varnish 2.1.x) Rule:

```
# Strip out Google Analytics campaign variables.  
# They are only needed by the javascript  
# running on the page.  
# utm_source, utm_medium, utm_campaign, gclid  
  
if(req.url ~ "(\?|&)(gclid|utm_[a-z]+)=") {  
    set req.url = regsuball(req.url, "(gclid|utm_[a-z]+)=[^  
\&]+&?", "");  
    set req.url = regsub(req.url, "(\?|&)$", "");  
}
```


2. Client Performance

Measuring Client Performance....

Measuring Client Performance....

- Page Load Size: 300 kB
- Page Load Time: 20 sec

How to measure Performance on Client?

- Use Google Chrome Developer Toolbar

➡ Network Tab



How to measure Performance on Client?

- <http://webpagetest.org>



Why are those pages
so big?

Need Compression of
CSS and JS!

Compression of CSS and JS

- Very easy to setup
- In Drupal 7 Core:
 - ▶ Administration Menu
 - ➡ Configuration
 - ➡ Development
 - ➡ Performance

Compression of CSS and JS

- Very easy to setup
- In Drupal 8 Core:
 - ▶ Enabled by default!

Compression of CSS and JS

- Aggregation and Compression
- Do this before Go-Live!
- Users will thank you for it!
- Needs: *mod_rewrite* and *mod_headers*



Compression of CSS and JS (Drupal 6)

- Can be also done in Apache (e.g.)
 - *mod_deflate*
- But: Can put high load on the server!
- ✓ Combine with: Varnish

Minimize CSS and JS
source files!

AdvAgg (Drupal 7)

=> Way less
aggregates

Alternatives to AdvAgg:

- <http://drupal.org/project/agrcache>
 - goes hopefully into core for Drupal 8 (<https://www.drupal.org/node/1014086>)
- +
- https://www.drupal.org/project/simple_aggregation
 - very simple, but effective module to reduce number of aggregates, very low risk

Set proper caching
headers

Set proper caching headers

- Goal: Cache for some time on Client machine
- Drupal 7 sets the headers for you
- Just need to adjust the numbers.



Client Performance: What did we achieve?

- Only 4 HTTP requests.
- Much faster page load time.

Your Mission:
UPDATE!

Your Mission: UPDATE!

- ✓ Anonymous Pages: „We are blazingly fast, really slick, and really really happy :-).“

Your Mission: UPDATE!



- Anonymous Pages: „We are blazingly fast, really slick, and really really happy :-)“.

Really Happy!



Additional techniques

- CDN: Content Delivery Network
- Caches files close to the users location
- Useful for images, CSS / JS files
- <http://drupal.org/project/cdn>



Additional techniques

- CDN: Content Delivery Network
- Fastly
- CloudFlare
- => Put your whole site behind a CDN!

Nicely integrated with tags for Drupal 8!



Additional techniques

- [AJAX/PJAX](#): Only reload the content you need
- Useful for pagers, image galleries
- <http://drupal.org/project/pjax>



Additional techniques

- Refreshless: Only reload the content you need
- Drupal 8 only
- Out of the box!
- <http://drupal.org/project/refreshless>



Quick-Tip: Fix slow JS!

- Unresponsive script error on loading of page?
- Workaround -- Wrap Code in:

```
setTimeout(function() {  
    // Old code  
}, 100);
```


3. Module Performance

Measuring Module Performance....

Measuring Module Performance....

- Drupal Bootstrap: 240 ms
- `menu_execute_active_handler`: 6 sec
- Memory Usage: 104 MB

How to measure Module Performance?

- Use **xhprof** PHP extension
- Integration via **xhprof** module
- ➔ **admin/config/development/devel**



How to measure Module Performance?

- Drupal 8: [webprofiler](#)
- Drupal 7: [block_timer](#)
- Want insights about your caches?

[heisencache](#)



Common Pitfalls

Common Pitfalls

- `variable_set` on each page request

➡ Can bring your DB server to its knees!

Common Pitfalls

- *Generally: Writes* on each page request

➡ Can bring your DB server to its knees!

Common Pitfalls

- Anonymous `$SESSION` set for saving simple data
- Example: low_bw flag
 - ➡ Disables anonymous caching
 - ➡ Solution: Use Javascript to set/receive cookies directly and change page.

Common Pitfalls

- *Having installed way
toooooooooooooooooooooooooooooooooooooo
many modules*
- ➡ Adds a little to the page request every
time. (Drupal 6: worse)

Interesting Pitfalls

- *views* loading 5000 nodes in one page request for *openlayers*

➡ Can easily exceed memory and takes quite some time to attach fields.

➡ Adv. Solution: *openlayers_quick_query_sandbox*

Improve performance

- Use block caching
- Use render caching (render_cache module)



Improve performance

- Use `blockcache_alter` module
- ➡ Lots of more opportunities for caching then by default



Improve performance

- Setup views caching
- Setup panels caching
- => Click those checkboxes and see what is acceptable for your users!



Improve performance

- Drupal 8:
`dynamic_page_cache`
enabled - out of the
box!
- Ensure it is effective!

Pages are still slow =>
something might be
disabling your caching.
- Use a `placeholder`!



Improve performance

- Drupal 7: **authcache**
- Needs to be configured; much easier to use now.
- Don't split too much!



Module Performance:

What did we achieve?

- We found the bad code and removed it!
- Page is much faster now!

Your Mission:
UPDATE!

Your Mission: UPDATE!

- Mrs. MySQL: Almost Happy. She feels still kinda slow sometimes
- ✓ Apache: Really happy!
- ✓ D. Pages: Really Happy!

4. Database Performance

Measuring MySQL Performance....

Measuring MySQL Performance....

- Slow SQL Query: 10 sec

Measuring MySQL Performance....

- Tip: Use Percona's mysql slow query log analyzer
- <https://www.percona.com/doc/percona-toolkit/2.2/pt-query-digest.html>

How to measure MySQL Performance?

- Enable **slow query log**
- **dbtuner** module (6.x only) / **mysqltuner** script for 7.x
- **EXPLAIN** queries



Common Tweaks

Common Tweaks

- Use Inno DB
- ✓ Default: Most current MySQL configurations

Common Tweaks

- Be aware of the barrier!
- `nobarrier=1` for `ext3/ext4` file systems
- Newer Linux kernels / Ubuntu
- Many many threads by Developers

Common Tweaks

- Be aware of the barrier!
- ➡ Attention: Needs to be used with special hardware for production usage.
- ➡ Useful Guide: Red Hat Handbook

Common Tweaks

- Use **XFS** file system
- Good and proven file system for MySQL databases
- Size appropriately to the use-case

Fix Slow Queries

- `EXPLAIN` queries
- ✓ Add `indexes` where necessary
- Run `EXPLAIN` again

MySQL Performance: What did we achieve?

- No more slow queries! (YEAH!)

Your Mission:
UPDATE!

Werbung / Commercials

Recap: Best Practices

Best Practices

- Setup Base Performance

Best Practices

- You want to have your own
„High Performance Stack“

Best Practices

- Having your own stack is not difficult ...
- ... as you have seen

Best Practices

- Analyze Pain Points first

Where is the problem?

- Server based?
- Client based?
- Modules?
- Database?

Best Practices

- Optimize Pain Points

Your Mission:
UPDATE!

Your Mission: UPDATE!

- ✓ Mrs. MySQL: Very Happy!
- ✓ Apache: Really happy!
- ✓ D. Pages: Really Happy!

Mission: Completed!

Wake up, Neo!

QUESTIONS?

FOLLOW ME!

- Fabian Franz
- @Fabianx
- twitter: @fabianfranz
- github: LionsAd



Tag1 Consulting

Performance & Scalability Experts



DUBLIN
DRUPALCON

JOIN US FOR CONTRIBUTION SPRINTS

First Time Sprinter Workshop - 9:00-12:00 - Room Wicklow2A

Mentored Core Sprint - 9:00-18:00 - Wicklow Hall 2B

General Sprints - 9:00 - 18:00 - Wicklow Hall 2A



DUBLIN
DRUPALCON

WHAT DID YOU THINK?

Evaluate This Session

events.drupal.org/dublin2016/schedule

THANK YOU!

