

# Drupal Performance for the Rest of Us

Eric Peterson (iamEAP)

# Goals

1. Understand caching and how Drupal does it.
2. Get comfortable configuring Drupal for performance.
3. Learn a subset of performance modules:
  - What they do and when to use them
  - How to configure them
  - See the impact

# Why trust this guy?

- My day job is to work on a site that supports over 1.5 million pageviews per month.
- My company's IT department has enough on its plate besides this "Drupal" thing.
- As a result, we've been forced to make our Drupal installation scale and perform without the help of Varnish or Memcache.
- You can always trust a handsome man.

# So what exactly is caching?




Caching is when you save off data that was expensive to obtain so that subsequent calls for that data are less expensive.

There are essentially three ways to configure:

1. What are you caching?
2. Where are you storing it?
3. How long are you storing it?

# The obligatory contrived example

Your in-laws are staying for the month and you're out of groceries. They're also on different and incredibly stupid diets.




**YOU'RE ON MY SHIT LIST**

OFFENDER: ☐ YOU

VIOLATION: ☐ BULLSHIT

SEVERITY: 1 2 3 4 5 6 7 8 9 10




**YOU'RE ON MY SHIT LIST**

OFFENDER: ☐ YOU

VIOLATION: ☐ BULLSHIT

SEVERITY: 1 2 3 4 5 6 7 8 9 10



# Everything is a cache within a cache

- DNS cache
  - Browser cache
    - Page cache
      - Block cache
      - Views cache
        - Query result cache
        - Output cache
          - MySQL query cache
          - PHP static caching
            - Opcode caching



**CACHE**  
**INCEPTION**

# What's your Drupal cache strategy?

## Primarily Anonymous Traffic

Focus on ways to cache each page as a whole:

- Page Cache
- External Cache (for browsers)
- Boost

## Primarily Authenticated Traffic

Focus on caching individual components of a page:

- Blocks
- Views
- Panels
- Entities

# Drupal performance configuration

## CLEAR CACHE

Admin > Config > Development > Performance

Clear all caches

## CACHING

☐ Cache pages for anonymous users

☐ Cache blocks

Block caching is inactive because you have enabled modules defining content access restrictions.

### Minimum cache lifetime

<none> ▾

Cached pages will not be re-created until at least this much time has elapsed.

### Expiration of cached pages

<none> ▾

The maximum time an external cache can use an old version of a page.

## BANDWIDTH OPTIMIZATION

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

☐ Aggregate and compress CSS files.

☐ Aggregate JavaScript files.



# Cache pages for anonymous users

This checkbox does exactly what it claims with only a couple of sensible caveats:

- If a user just POSTed form data, the page is not cached.
- If a module stores data for anonymous users in the session, pages will not be cached.

**Bonus:** you can disable page cache hook invocations by setting the following in your settings.php file.

```
$conf['page_cache_invoke_hooks'] = FALSE;
```

**Caution:** Some modules depend on this functionality. Configure with care. Those familiar with D6: this is the same as "aggressive" caching.

# Minimum cache lifetime

The most misunderstood configuration, maybe in all of Drupal.

*"Cached pages will not be re-created until at least this much time has elapsed."*

This is something of a half truth:

- True, this is not an explicit expiration time; it is a *minimum* time period during which the cache will not be deleted.
- However, it does not guarantee that any individual cache entry lives as long as you select. We'll explain this in a bit.


# Expiration of cached pages

This setting relates only to how systems *external* to Drupal may cache your page. Don't confuse this with minimum cache lifetime.

Think of this as the length of time a user's browser is allowed to hang on to a copy of your page before requesting another copy from your server.

As your scalability needs increase, this setting is also what you use to let Varnish and other reverse proxies know how long it should hold on to the same page.

# Clear all caches



This button does exactly what it claims. Use with caution!

We'll discuss alternatives and augmentations later.

# Cache garbage collection


If you specify a *minimum* cache lifetime instead of an explicit expiration time, how do those cache entries ever expire?

**Answer:** Cron!

- Cache entries that adhere to your minimum lifetime configuration are always marked *temporary*.
- Cron clears all *temporary* cache entries no more than once every minimum cache lifetime interval.

# Caveat emptor: saving content

Anytime you save a node, Drupal performs the same garbage collection procedure that cron does (except just on page and block caches).



# Unexpected consequences

## **When minimum cache lifetime = none**

- Every time cron is run, your entire cache is emptied.
- Anytime anyone adds or updates any piece of content, your entire page and block caches are emptied.


*For the love of god, never under any circumstances keep your minimum cache lifetime set to "none."*

## **When minimum cache lifetime > cron interval**

- In practice, your minimum cache lifetime is a maximum.

# Unexpected consequences (cont.)

Because cache garbage collection is bin-based rather than entry-based, garbage collection is expensive.






# Recommendations

- If at all possible, run cron less often.
- If that's not possible, increase minimum cache lifetime as much as is reasonable.
- What's "reasonable" depends on your site.
  - The more often content is updated or added, the smaller you want your minimum lifetime (for editor UX).
  - The same applies for comments if you're using the core commenting system.
- The above not good enough?
  - Alternative Database Cache
  - Finer control over cache expiration

# Alternative Database Cache (ADBC)

<http://www.drupal.org/project/adbc>

The ADBC module attempts to make garbage collection less traumatic by clearing cache bins on an entry-by-entry basis, rather than all at once.



# Installing & configuring ADBC

1. Enable the module like any other module.
2. Add the following code to settings.php:

```
$conf['cache_backends'][] = 'sites/all/modules/adbc/adbc.cache.inc';  
$conf['cache_default_class'] = 'AlternativeDrupalDatabaseCache';
```

No additional configuration is necessary beyond how you would normally configure Drupal for performance.

The result will be increased overall cache hit rate.

# Finer control over cache expiration

Another strategy is to increase cache minimum lifetime abnormally high (12-24 hours), and create a smarter cache invalidation system.

Two modules are well suited for this particular use-case in different ways:

- Cache Actions ([http://www.drupal.org/project/cache\\_actions](http://www.drupal.org/project/cache_actions))
- Flush Page Cache ([http://www.drupal.org/project/flush\\_page\\_cache](http://www.drupal.org/project/flush_page_cache))

# Cache Actions

Cache Actions integrates with the Rules module by introducing cache-related actions that can be triggered on any regular Rules event.

Some common use-cases:

- Clear the page cache for the current page when an editor makes content modifications.
- Clear the page cache for the current blog post when a user adds a comment.
- Clear the views cache for a view after re-ordered by an editor.

# Configuring Cache Actions

[Home](#) » [Administration](#) » [Configuration](#) » [Workflow](#) » [Rules](#)

## Events

### EVENT

After updating existing content

[+ Add event](#)

## Conditions

### ELEMENTS

None

[+ Add condition](#) [+ Add or](#) [+ Add and](#)

## Actions

### ELEMENTS

[+ Clear a specific cache cid](#)

Parameter: *Cache bin:* cache\_page, *Cache key:* [site:current-page], *Use wil*

[+ Add action](#) [+ Add loop](#)

[Home](#) » [Administration](#) » [Configuration](#) » [Workflow](#) » [Rules](#) » [Editing reaction rule "test"](#)

### CACHE BIN

The cache table where the cid is

Value \*

cache\_page

### CACHE KEY

The key to clear

Value \*

[site:current-page]

[▶ PHP EVALUATION](#)

[▶ REPLACEMENT PATTERNS](#)

[Switch to data selection](#)

# On-demand cache expiration

While Cache Actions is great for general cache clear rules, it quickly becomes cumbersome and difficult as events become more obscure and cache items become more specific.

One way to address these situations is with the Flush Page Cache module.

**Flush Page Cache** gives you and your editors a button on each page that invalidates every cache loaded on the page where it was clicked.

# Configuring Flush Page Cache

1. Once the module is installed, add the following to your settings.php:

```
$conf['cache_backends'][] = 'sites/all/modules/flush_page_cache/flush_page_cache.cache.inc';  
$conf['cache_default_class'] = FlushPageCacheDrupalDatabaseCache;
```


2. You'll need to place the button the module generates on every page and limit it to trusted roles. This can be done by placing a block or enabling the footer link the module optionally provides.
3. You may also want to style the button to a fixed position on the page.

Flush page cache



# Caching more

Now that we've covered Drupal's cache idiosyncrasies and smarter expiration techniques, let's dive into specifics...



# Module: Block Cache Alter

Drupal 7 introduced an API for different cache levels.

Block Cache alter provides a GUI to manage those cache levels on a block-by-block basis.

## ▼ CACHE SETTINGS

### Cache setting

✓ Do not cache

Cache once for everything (global)

Per page

Per role

Per role per page

Per user

Per user per page

for this block.

# Caveat emptor: Blocks via Context

If you use context to place your blocks, they're not necessarily run through the normal Block API.

There's a module for D6, but the D7 port is still in dev.

## **Context Block Cache Alter**

[http://drupal.org/project/context\\_blockcache\\_alter](http://drupal.org/project/context_blockcache_alter)

# Caching Views

**Page: Customer Stories Main Table: Caching options**

For

**Query results**  
  
The length of time raw query results should be cached.

**Rendered output**  
  
The length of time rendered HTML output should be cached.

# Caching Panels

Cache settings for this display

✖ Close Window

Lifetime

1 hour ▾

Granularity

None ▾

If "arguments" are selected, this content will be cached per individual argument to the entire display; if "contexts" are selected, this content will be cached per unique context in the pane or display; if "neither" there will be only one cache for this pane.

Save

# Caching Entities

Nodes! Comments! Taxonomies! Files! Users! Oh my!

Occasionally, we need to load these things.

The **Entity Cache** module provides caching for the above core entities.

Installation is extremely simple. Just download and install.

In practice, only really beneficial on frequently accessed pages that load large numbers of entities.

# Storing cache entries as files

Performing a database query requires a certain amount of overhead.

Reading a file is an extremely fast operation. Why not store cache entries as files rather than in the database?

The **File Cache** module opens up this possibility.

Install the module, then add the following to settings.php:

```
$conf['cache_backends'][] = 'sites/all/modules/filecache/filecache.inc';  
$conf['cache_default_class'] = DrupalFileCache;
```

# Caching pages to HTML (Boost)

How is this different from File Cache?

File cache switches all (or select) cache storage to be file-backed. Boost saves the final rendered page markup to an HTML file, then subsequently serves the HTML file.

File cache will still boot PHP and Drupal on each request. Boost avoids PHP completely once the page is saved off.



# Installing & configuring Boost

Somewhat more involved than most module installations.

1. Install and enable the module as usual.
2. Disable core page cache.
3. Ensure the cache directory is writeable by the web server.
4. Configure Apache with rules generated by Boost (.htaccess).

Details available at <http://drupal.org/node/1459690>

# Interactive Performance Comparisons

Jump over to Tableau Public!

<http://public.tableausoftware.com/shared/G76NWKNK7>

# Questions



# Thanks!


## Eric Peterson

**E-mail**                      [epeterson@tableausoftware.com](mailto:epeterson@tableausoftware.com)


**Handle**                      iamEAP

**WWW**                        <http://www.asmallwebfirm.net>


# Benchmarking Core block cache



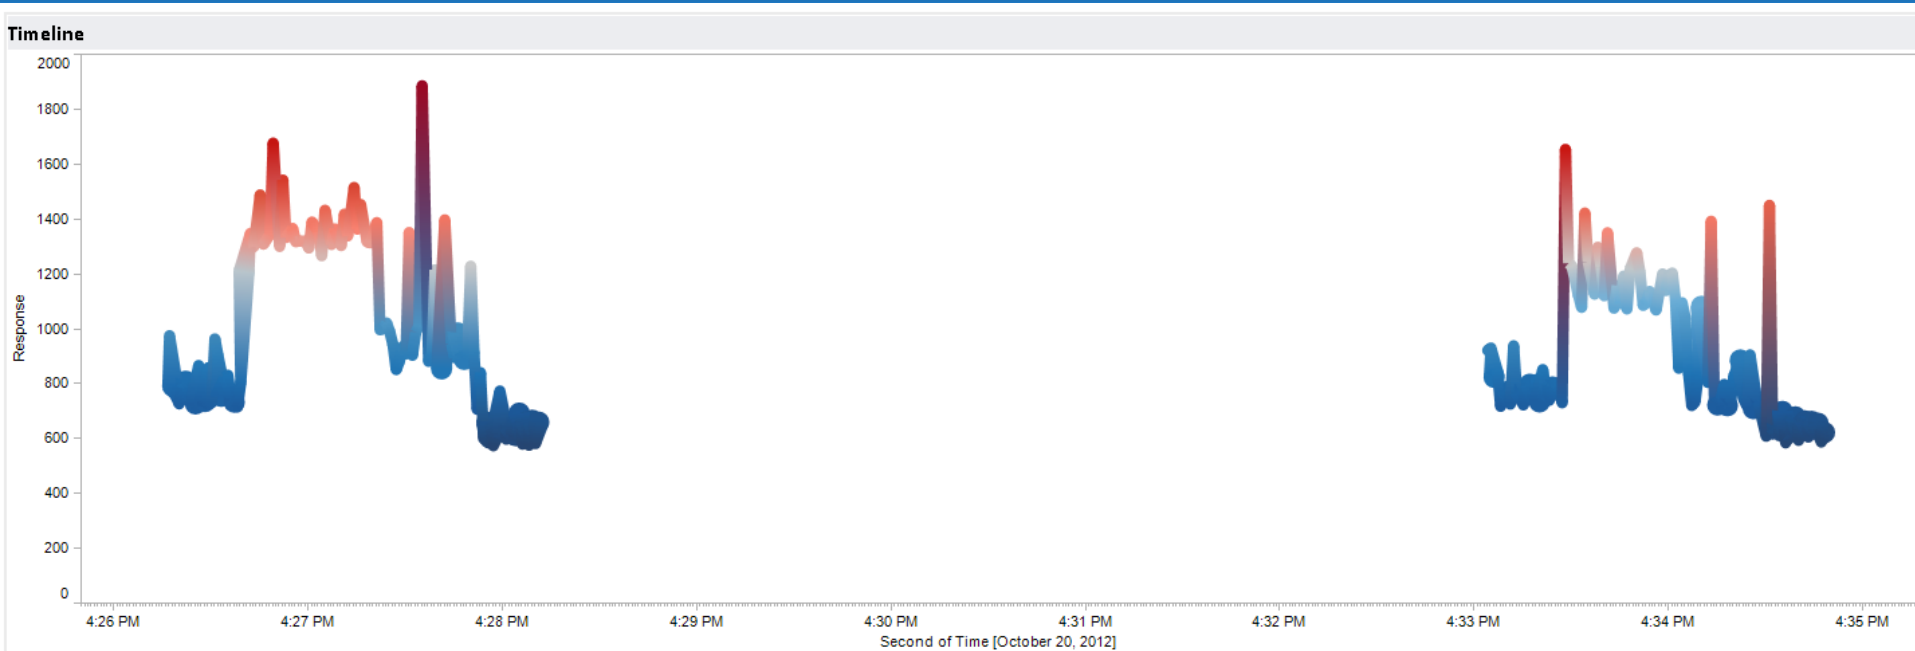
## Page Type Comparison by Cache




## Overall Average Response




# Benchmarking Views cache (output)



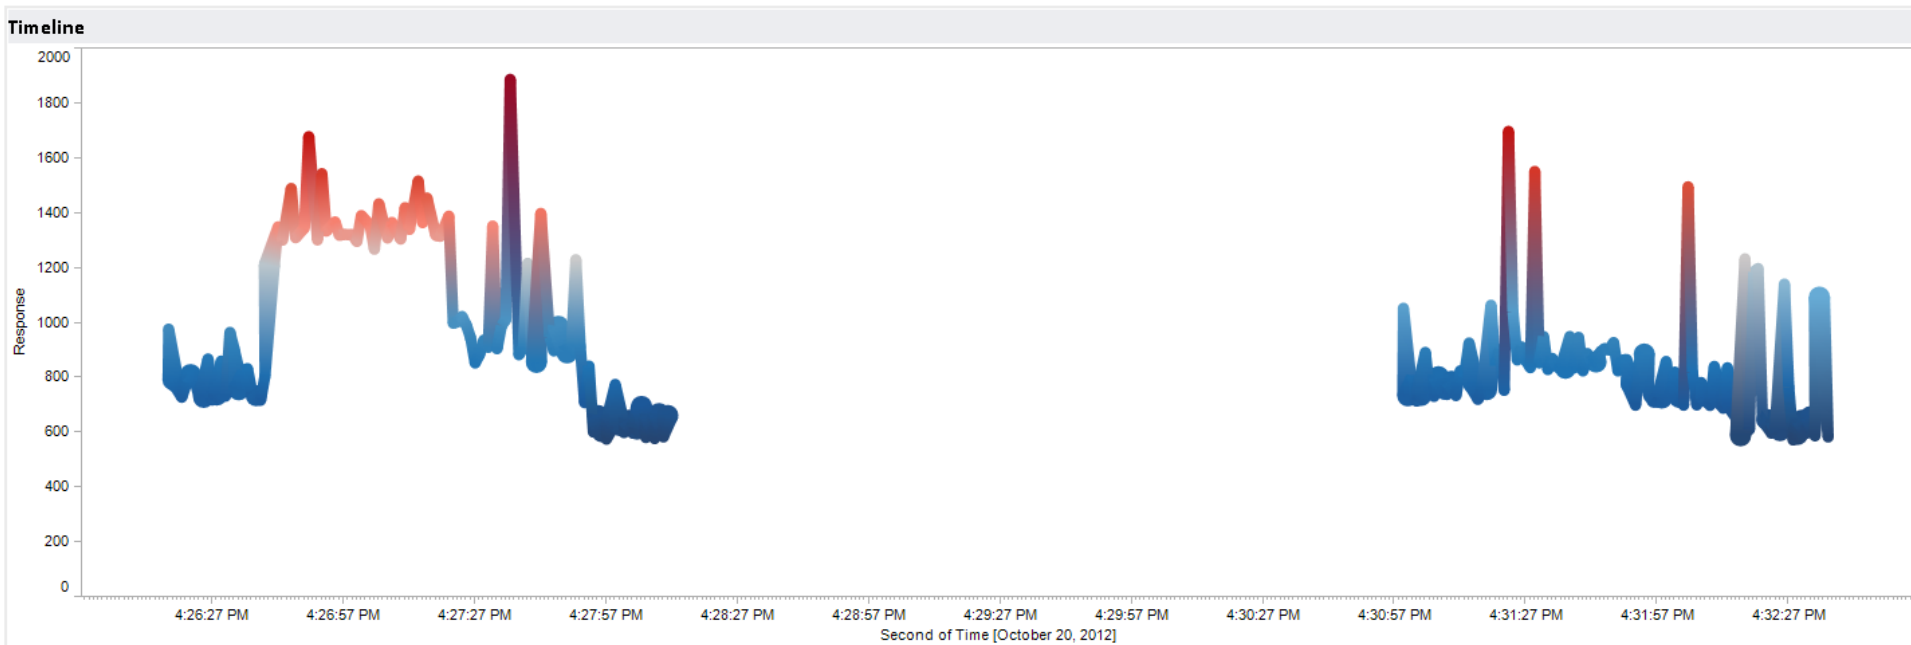
## Page Type Comparison by Cache




## Overall Average Response




# Benchmarking Views cache (results)



## Page Type Comparison by Cache




## Overall Average Response




# Benchmarking Views cache (both)


Timeline



Page Type Comparison by Cache




Overall Average Response






# Benchmarking Panels cache


Timeline



Page Type Comparison by Cache

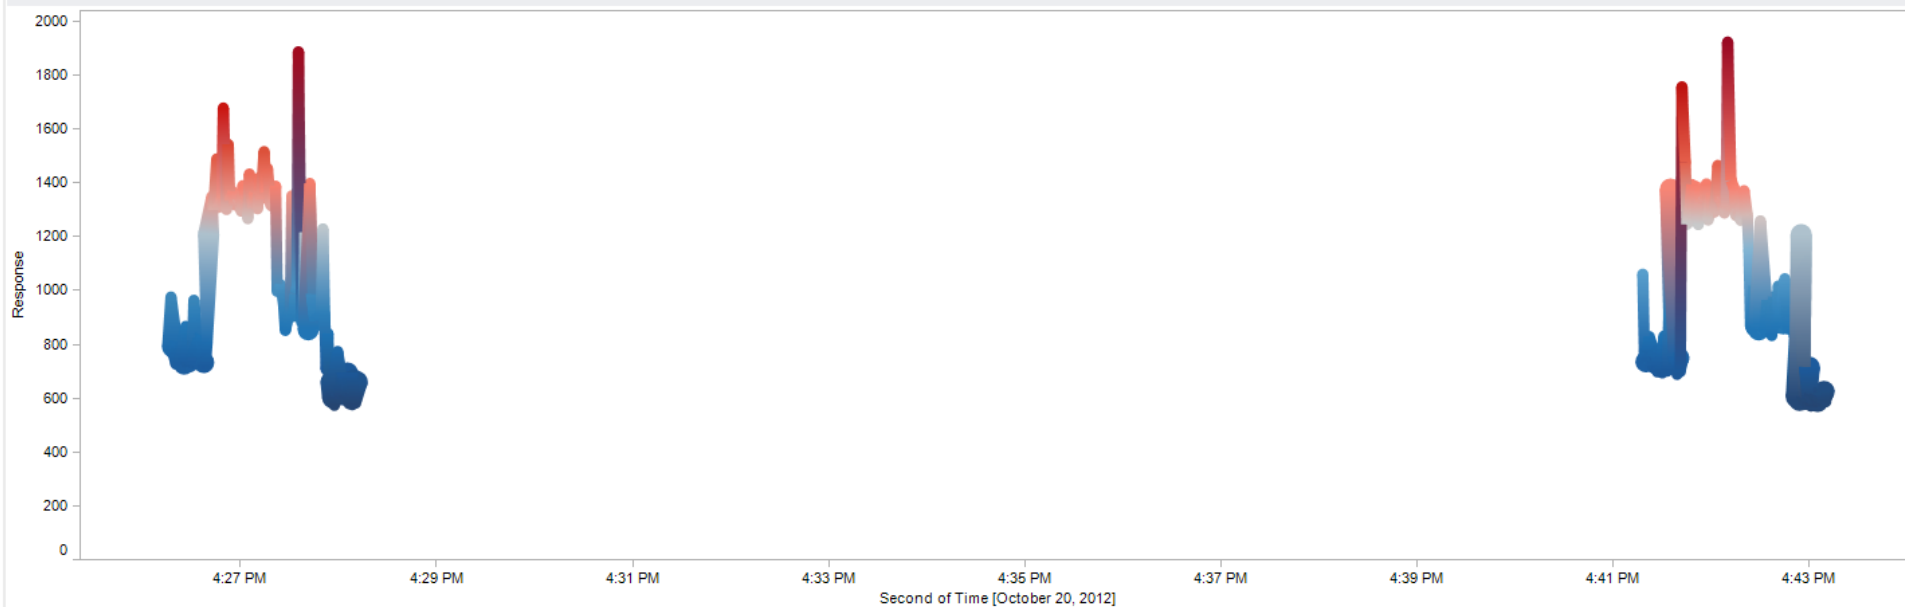


Overall Average Response




# Benchmarking Entity Cache


Timeline



Page Type Comparison by Cache

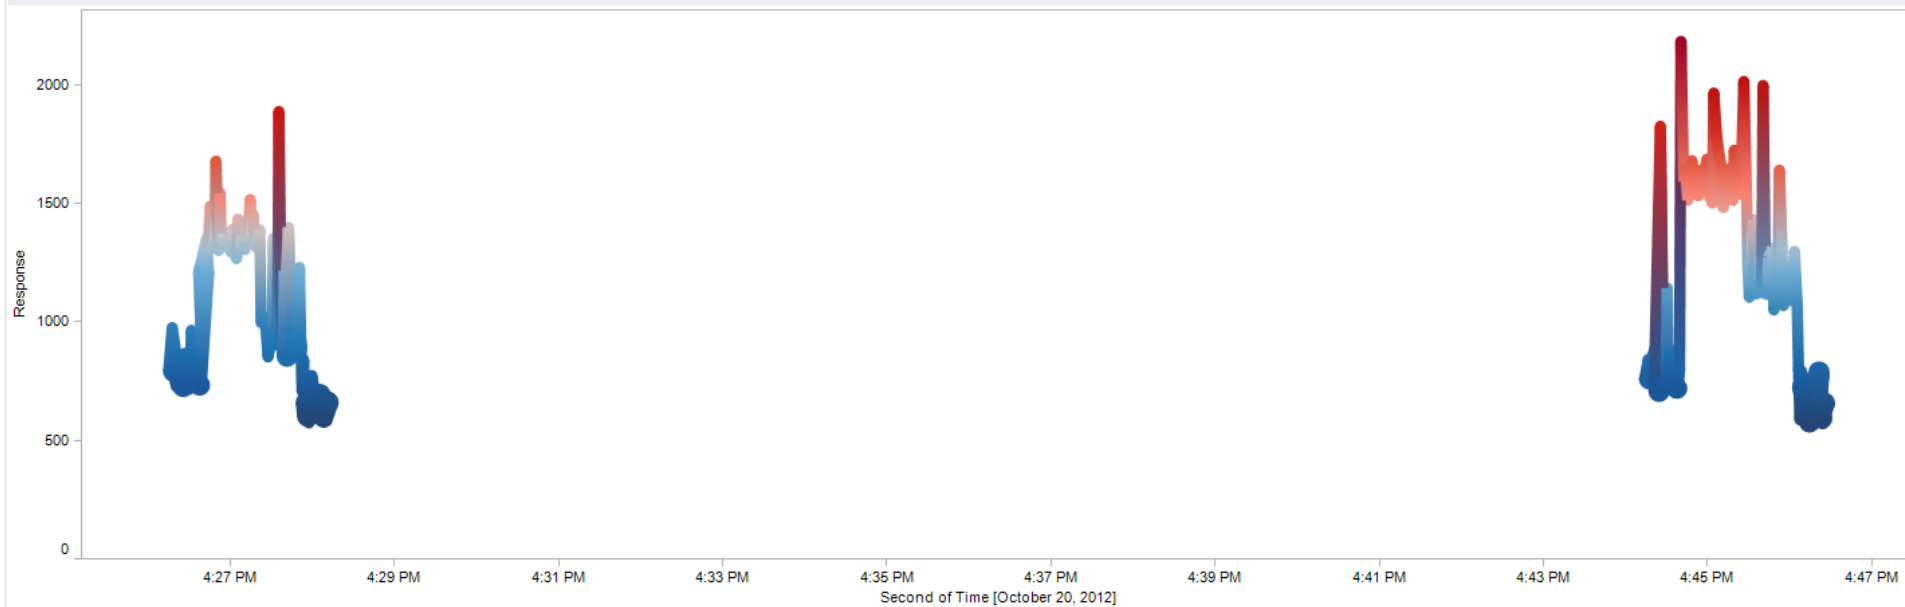


Overall Average Response




# Benchmarking File Cache


Timeline



Page Type Comparison by Cache




Overall Average Response




# Benchmarking Core page cache


Timeline



Page Type Comparison by Cache




Overall Average Response




# Page cache with File Cache enabled


Timeline



Page Type Comparison by Cache




Overall Average Response




# Benchmarking Boost


Timeline



Page Type Comparison by Cache




Overall Average Response




# Boost vs. Page Cache

Timeline



Page Type Comparison by Cache



Overall Average Response

