# Performance at Mass.gov

**Rob Bayliss**

CTO/Last Call Media

**Moshe Weitzman**

Drupal Architect

# A word about Mass.gov

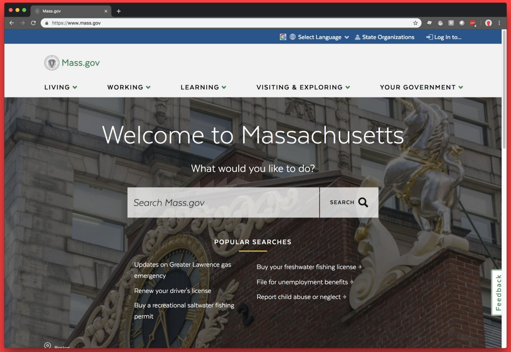

Welcome to Massachusetts

What would you like to do?

POPULAR SEARCHES

- Updates on Greater Lawrence gas emergency
- Renew your driver's license
- Buy a recreational saltwater fishing permit
- Buy your freshwater fishing license →
- File for unemployment benefits →
- Report child abuse or neglect →

Serves the Commonwealth of Massachusetts

Stakeholders are:

- Constituents (as visitors)
- State Organizations (as visitors and publishers)

Receives ~15M pageviews/month

Changes are released twice weekly by a team of developers.

## Backend Performance:

State employees need to be able to publish content in **near real time**.

Constituents need **fast access** to information with **no disruptions**.

The DevOps team needs to balance **freshness** with **reliability**.

# The Traffic Profile

- Visitors come in all day (heaviest 8AM - 8PM)
- Editors work between business hours (8AM - 5PM)
- Releases happen twice weekly (after 8PM)

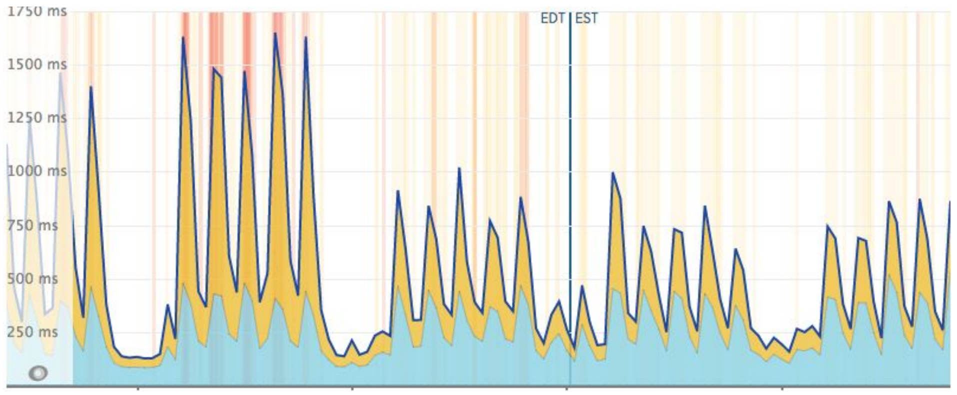

## The Operating Environment

- CDN caches everything for anonymous traffic
- Varnish caches everything for anonymous traffic
- Webs are only hit when there is a miss on CDN and Varnish

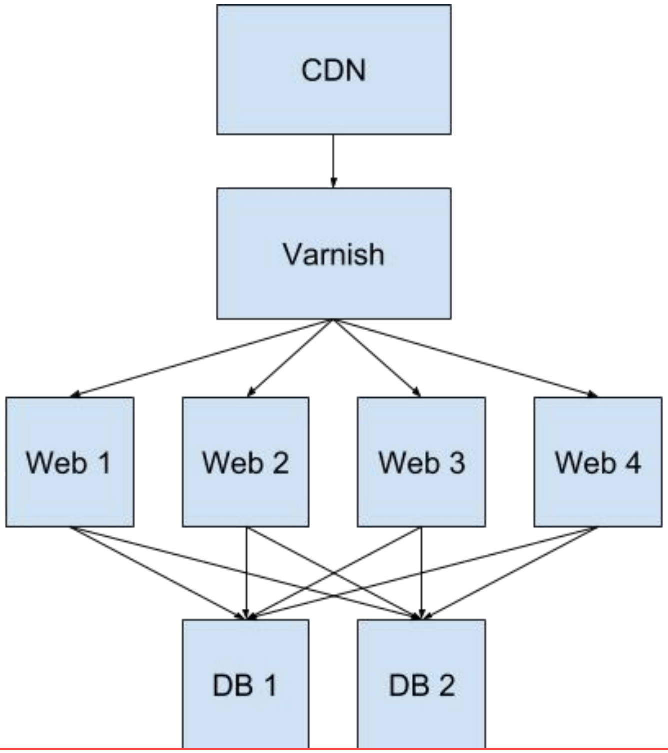

# September, 2017

Site “Launch”

DNS is switched to change the URL of the Drupal site from *pilot.mass.gov* to *mass.gov*

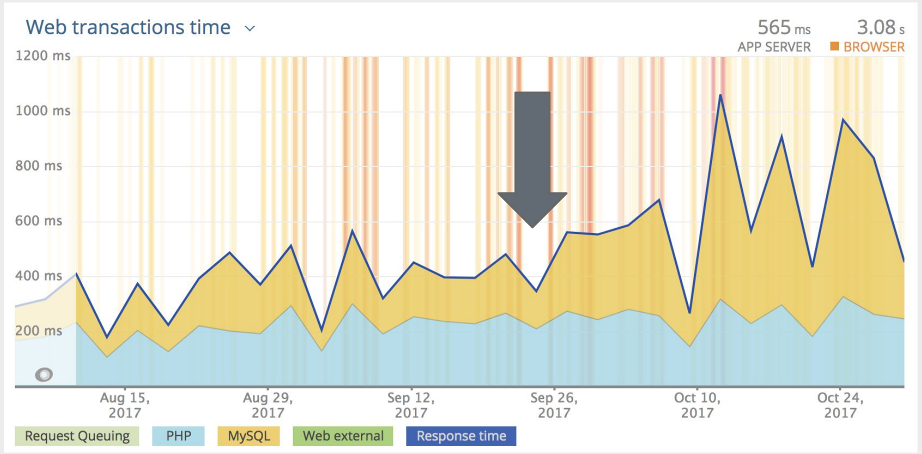

# November 30, 2017

Purge module is rolled out

Varnish cleared on content update

Editors see changes in < 1 hour

Freshness improves, but performance take a big hit.

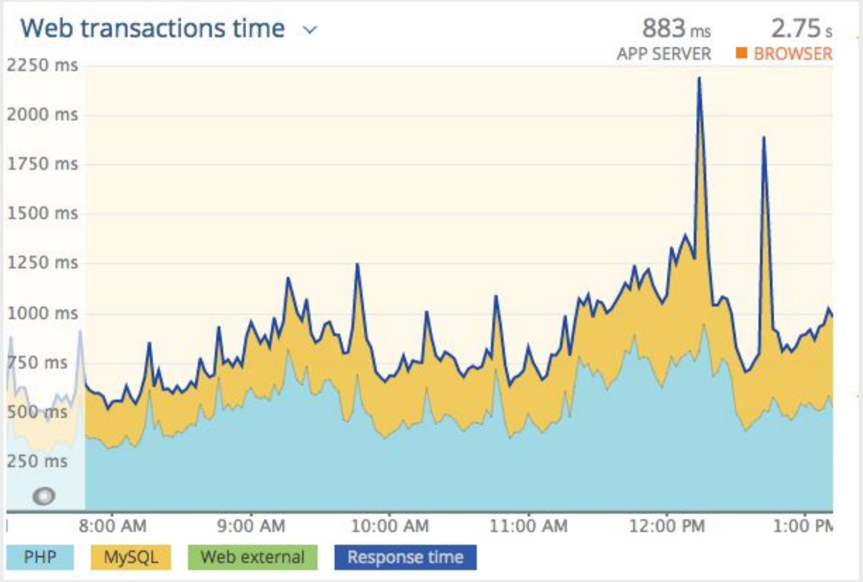

# December, 2017

Many missing cache tags are fixed.

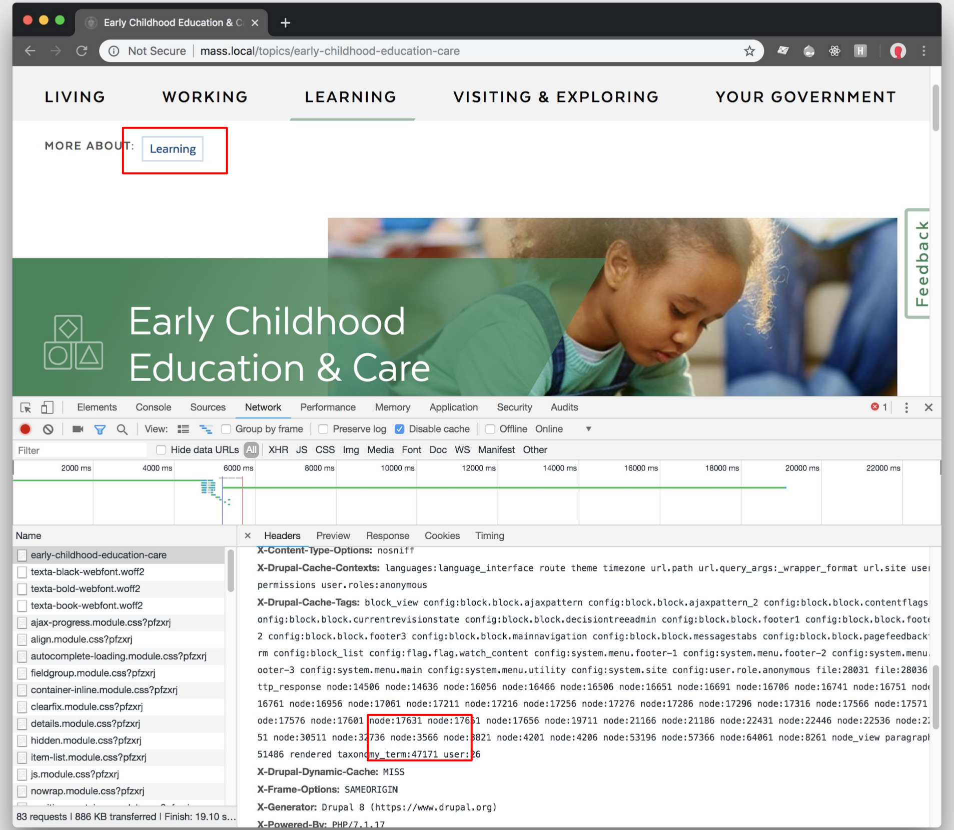

Early Childhood Education & C
Not Secure | mass.local/topics/early-childhood-education-care
LIVING
WORKING
LEARNING
VISITING & EXPLORING
YOUR GOVERNMENT
MORE ABOUT: Learning
Early Childhood Education & Care
Feedback
Elements Console Sources Network Performance Memory Application Security Audits
Name
Headers Preview Response Cookies Timing

```
early-childhood-education-care
texta-black-webfont.woff2
texta-bold-webfont.woff2
texta-book-webfont.woff2
ajax-progress.module.css?pfzxrj
align.module.css?pfzxrj
autocomplete-loading.module.css?pfzxrj
fieldgroup.module.css?pfzxrj
container-inline.module.css?pfzxrj
clearfix.module.css?pfzxrj
details.module.css?pfzxrj
hidden.module.css?pfzxrj
item-list.module.css?pfzxrj
js.module.css?pfzxrj
nowrap.module.css?pfzxrj
83 requests | 886 KB transferred | Finish: 19.10 s...
```

```
X-Content-Type-Options: nosniff
X-Drupal-Cache-Contexts: languages:language_interface route theme timezone url.path url.query_args:_wrapper_format url.site user
permissions user.roles:anonymous
X-Drupal-Cache-Tags: block_view config:block.block.ajaxpattern config:block.block.ajaxpattern_2 config:block.block.contentflags
onfig:block.block.currentrevisionstate config:block.block.decisiontreeadmin config:block.block.footer1 config:block.block.foote
2 config:block.block.footer3 config:block.block.mainnavigation config:block.block.messagestabs config:block.block.pagefeedback
rm config:block_list config:flag.flag.watch_content config:system.menu.footer-1 config:system.menu.footer-2 config:system.menu.
ooter-3 config:system.menu.main config:system.menu.utility config:system.site config:user.role.anonymous file:28031 file:28036
ttp_response node:14506 node:14636 node:16056 node:16466 node:16506 node:16651 node:16691 node:16706 node:16741 node:16751 node
16761 node:16956 node:17061 node:17211 node:17216 node:17256 node:17276 node:17286 node:17296 node:17316 node:17566 node:17571
ode:17576 node:17601 node:17631 node:17651 node:17656 node:19711 node:21166 node:21186 node:22431 node:22446 node:22536 node:2
51 node:30511 node:32736 node:3566 node:3821 node:4201 node:4206 node:53196 node:57366 node:64061 node:8261 node_view paragrap
51486 rendered taxonomy_term:47171 user:26
X-Drupal-Dynamic-Cache: MISS
X-Frame-Options: SAMEORIGIN
X-Generator: Drupal 8 (https://www.drupal.org)
X-Powered-By: PHP/7.1.17
```

LAST CALL MEDIA

# December, 2017

Drupal cache lifetime increased to 3 hours

Attempt to take advantage of the purge work by holding pages that *haven't changed* for up to 3 hours.

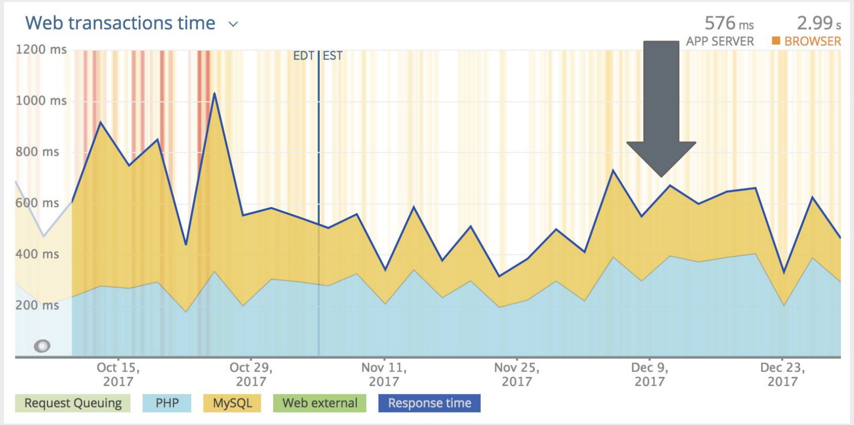

## January, 2018

Backend disruption due to accidental self-DDOS

Work on the emergency alerts system disrupts the site after it is deployed with a cache-busting query parameters in an AJAX request.

Takeaways:

- Every PR needs to be considered for performance impact.
- Having a system in place to run some level of load test or benchmarking on changes before production is key if uptime matters.
- Even front end work can affect backend performance.

# January, 2018

Strip out the *node_list* cache tag from most pages, replacing with “relationship” tag clearing.

The absence of the *node_list* tag is backed up by Behat tests, and relationship clearing is also tested.

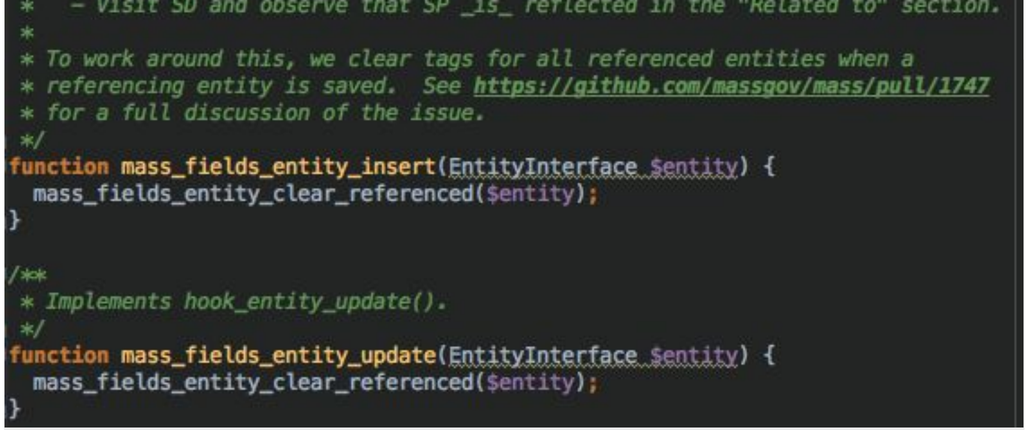

```
 *   - Visit SD and observe that SP _is_ reflected in the "Related to" section.
 *
 * To work around this, we clear tags for all referenced entities when a
 * referencing entity is saved.  See https://github.com/massgov/mass/pull/1747
 * for a full discussion of the issue.
 */
function mass_fields_entity_insert(EntityInterface $entity) {
  mass_fields_entity_clear_referenced($entity);
}

/**
 * Implements hook_entity_update().
 */
function mass_fields_entity_update(EntityInterface $entity) {
  mass_fields_entity_clear_referenced($entity);
}
```

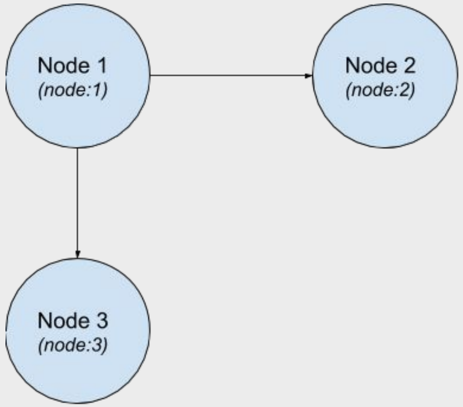

# January, 2018

Strip out the *node_list* cache tag from most pages, replacing with “relationship” tag clearing.

The absence of the *node_list* tag is backed up by Behat tests, and relationship clearing is also tested.

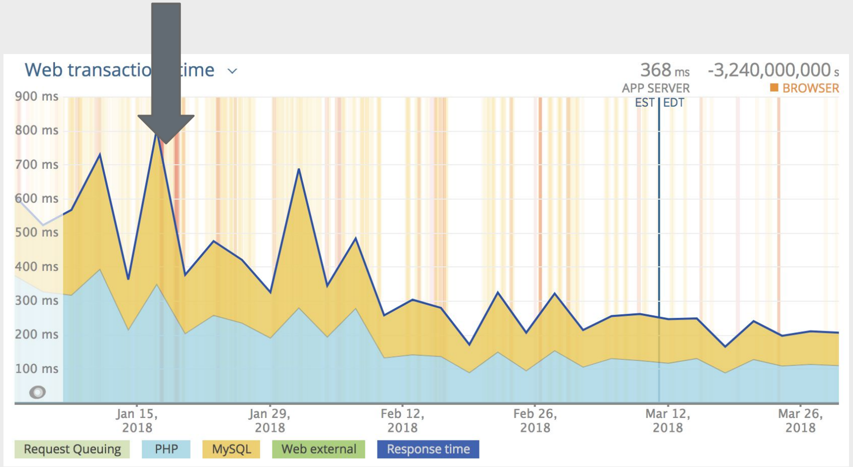

# January, 2018

Metatags module is patched to prevent 2x token generation.

This performance enhancement was submitted and accepted on D.O.

Takeaways:

- Contrib modules aren't immune to performance problems.

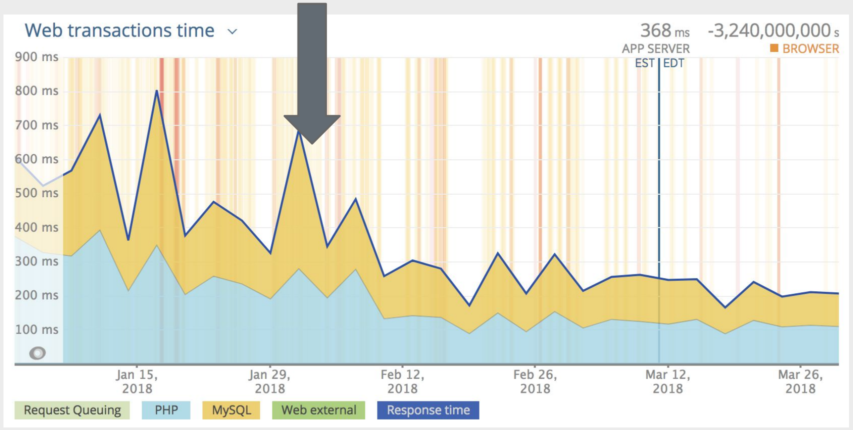

# February, 2018

Backend disruption due to heavy traffic on admin views

Symptoms: Users unable to log in, slowness on uncached pages of public site.

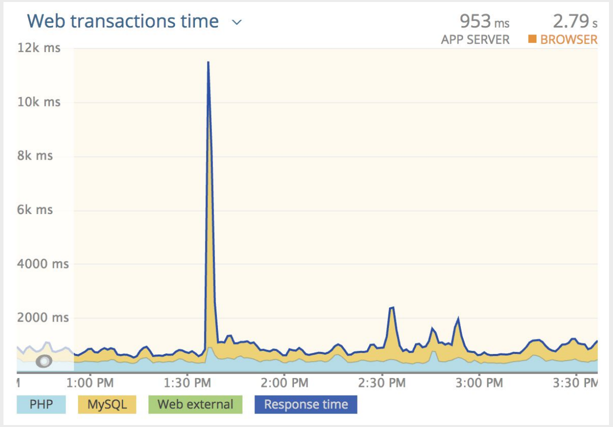

# February, 2018

Backend disruption due to heavy traffic on admin views

Symptoms: Users unable to log in, slowness on some pages of public site.

Takeaways:

1. Your bottlenecks will change as traffic patterns change.
2. Be proactive about fixing small issues before they become big issues - we spotted this issue the month prior, but didn't prioritize.
3. Full pagers are evil because they require count queries that are always as bad or worse than the views query.

# February, 2018

Upgrade to PHP 7

“Free” performance win celebrated by everyone.

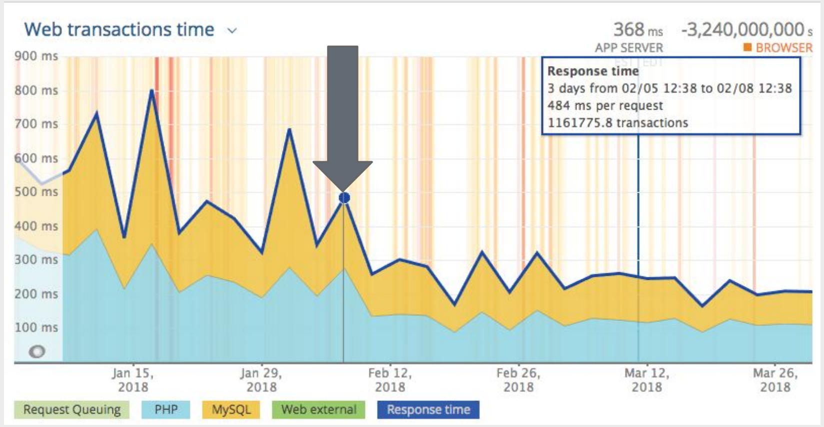

**Response time**
3 days from 02/05 12:38 to 02/08 12:38
484 ms per request
1161775.8 transactions

## February, 2018

New Relic Monitoring Customizations.

You can't improve what you can't measure

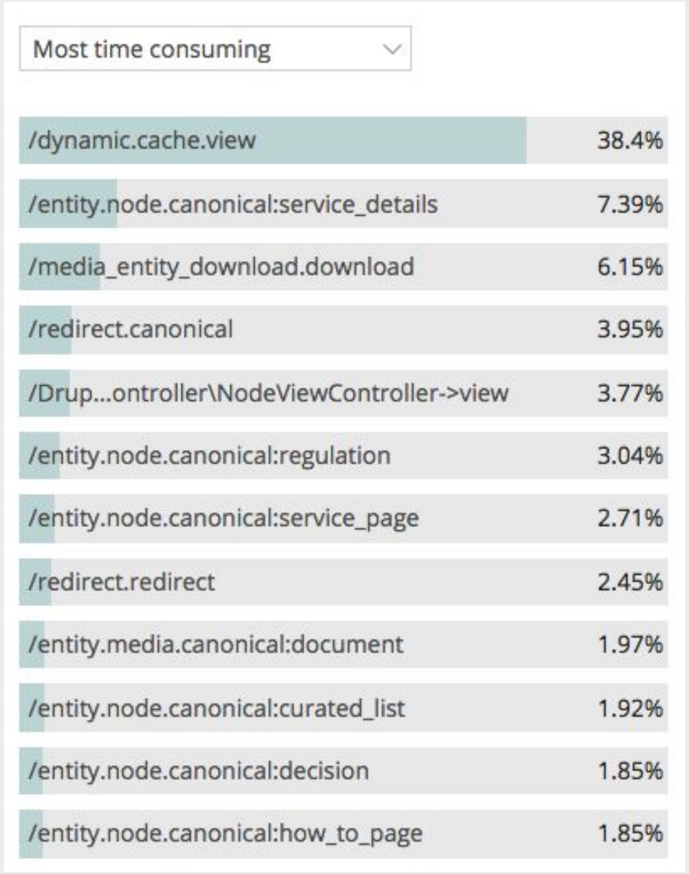

| | |
|---|---|
| /dynamic.cache.view | 38.4% |
| /entity.node.canonical:service_details | 7.39% |
| /media_entity_download.download | 6.15% |
| /redirect.canonical | 3.95% |
| /Drup...ontroller\NodeViewController->view | 3.77% |
| /entity.node.canonical:regulation | 3.04% |
| /entity.node.canonical:service_page | 2.71% |
| /redirect.redirect | 2.45% |
| /entity.media.canonical:document | 1.97% |
| /entity.node.canonical:curated_list | 1.92% |
| /entity.node.canonical:decision | 1.85% |
| /entity.node.canonical:how_to_page | 1.85% |

# February, 2018

New Relic Monitoring Customizations.

You can't improve what you can't measure

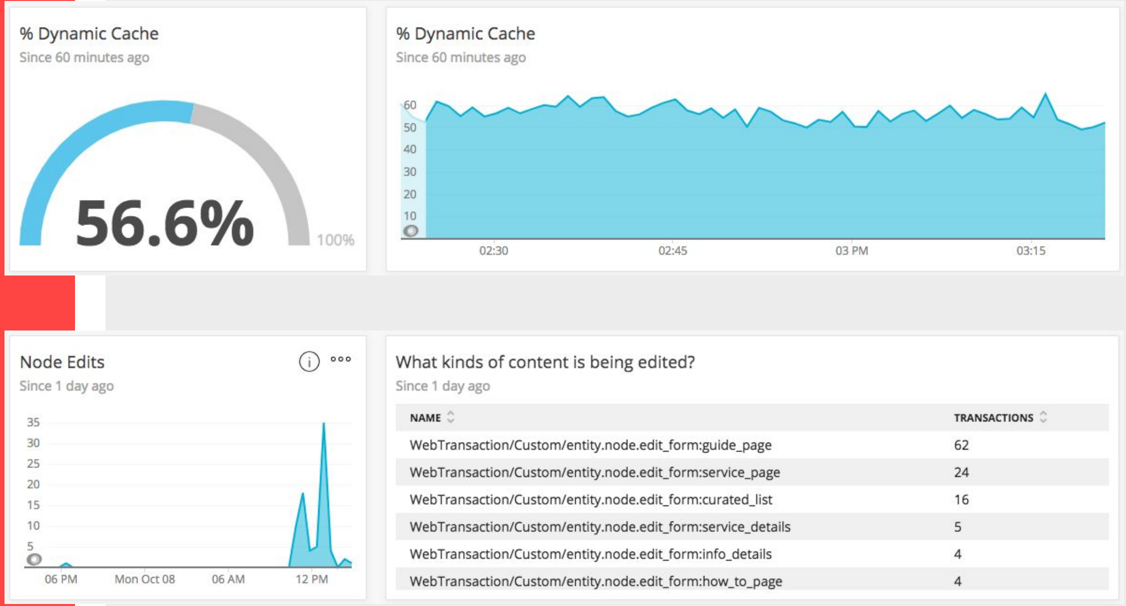

What kinds of content is being edited?

Since 1 day ago

| NAME | TRANSACTIONS |
|---|---|
| WebTransaction/Custom/entity.node.edit_form:guide_page | 62 |
| WebTransaction/Custom/entity.node.edit_form:service_page | 24 |
| WebTransaction/Custom/entity.node.edit_form:curated_list | 16 |
| WebTransaction/Custom/entity.node.edit_form:service_details | 5 |
| WebTransaction/Custom/entity.node.edit_form:info_details | 4 |
| WebTransaction/Custom/entity.node.edit_form:how_to_page | 4 |

# February, 2018

Audit and remediate cache hogging items

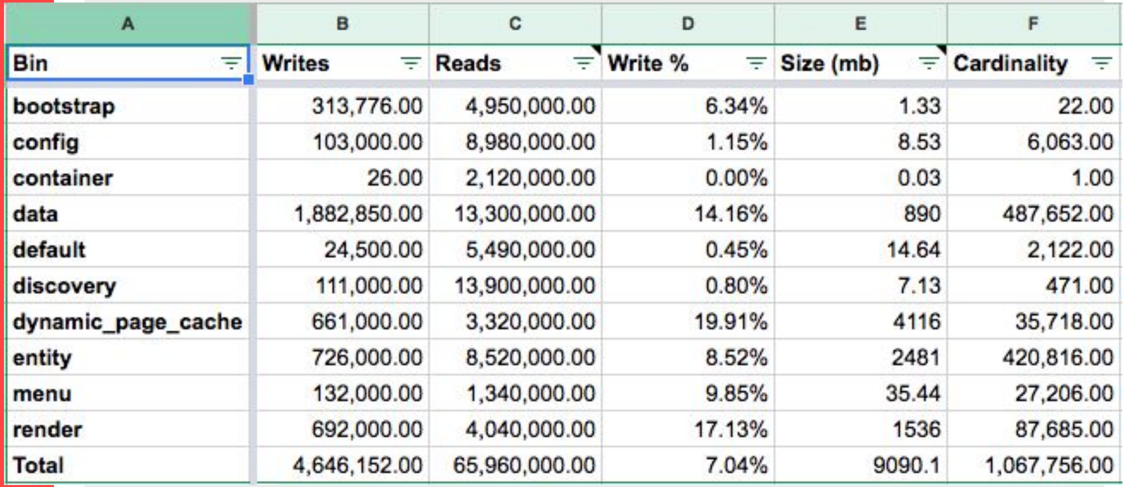

| Bin | Writes | Reads | Write % | Size (mb) | Cardinality |
|---|---|---|---|---|---|
| bootstrap | 313,776.00 | 4,950,000.00 | 6.34% | 1.33 | 22.00 |
| config | 103,000.00 | 8,980,000.00 | 1.15% | 8.53 | 6,063.00 |
| container | 26.00 | 2,120,000.00 | 0.00% | 0.03 | 1.00 |
| data | 1,882,850.00 | 13,300,000.00 | 14.16% | 890 | 487,652.00 |
| default | 24,500.00 | 5,490,000.00 | 0.45% | 14.64 | 2,122.00 |
| discovery | 111,000.00 | 13,900,000.00 | 0.80% | 7.13 | 471.00 |
| dynamic_page_cache | 661,000.00 | 3,320,000.00 | 19.91% | 4116 | 35,718.00 |
| entity | 726,000.00 | 8,520,000.00 | 8.52% | 2481 | 420,816.00 |
| menu | 132,000.00 | 1,340,000.00 | 9.85% | 35.44 | 27,206.00 |
| render | 692,000.00 | 4,040,000.00 | 17.13% | 1536 | 87,685.00 |
| Total | 4,646,152.00 | 65,960,000.00 | 7.04% | 9090.1 | 1,067,756.00 |

# February, 2018

## Audit and remediate cache hogging items

Anatomy of a cache item:

```
response:[languages:language_interface]=en:[request
_format]=html:[route]=mass_map.map_page5a04d4631e65
8a97aee299dca5490f20606afa3b1d74727d2789c61256bf616
f:[theme]=mass_theme:[url.path]=/visit-massachusett
s-state-QeOOjmfLF-hoOOWMXuD9XMrqNVUDXcOsEZrAm7KYfas
```

This cache item uses the following “contexts”:

- *languages:language_interface*
- *request_format*
- *route*
- *theme*
- *url.path*

## February, 2018

Audit and remediate cache hogging items

A simple example:

```
response:[route]=entity.node.canonical5a04:[url.path]
=/visit-massachusetts-state-parks
```

This represents the same content as:

```
response:[route]=entity.node.canonical5a04:[url.path]
=/node/123
```

You will create a new cache entry every time a context you care about changes. Be a nihilist… care about less stuff!

## February, 2018

Audit and remediate cache hogging items

Problematic contexts:

1. *url.path* - prefer *route* wherever you can
2. *url.query_args* - prefer *url.query_args.X* instead
3. *headers.** - there aren't many valid use cases for this
4. *cookies.** - Don't use cookies in your backend code.

## February, 2018

Audit and remediate cache hogging items

**Rule of thumb**: cache global elements with the broadest/least contexts possible, using a #lazy_builder if you need a single, simple element to change based on some conditions.

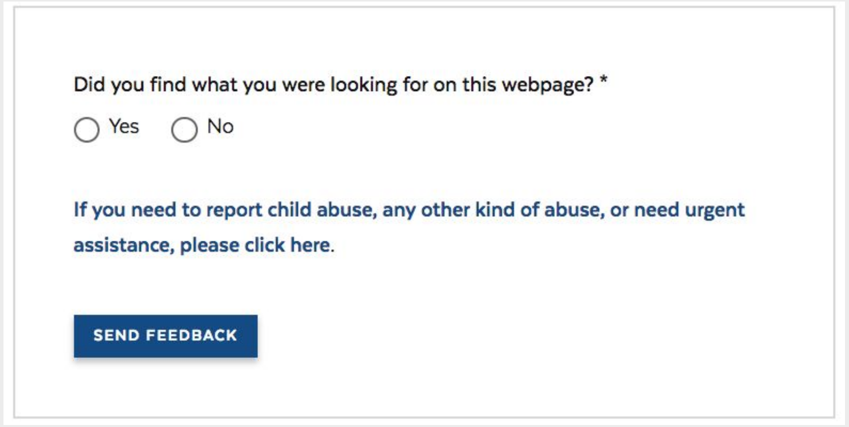

Did you find what you were looking for on this webpage? *

◯ Yes ◯ No

If you need to report child abuse, any other kind of abuse, or need urgent assistance, please click here.

SEND FEEDBACK

# February, 2018

Audit and remediate cache hogging items

(a word about the dynamic page cache)

**Rule of thumb:** If you have a reverse proxy cache in front of your site, use *dynamic_page_cache* instead of *page_cache*. This will save a lot of cache space due to url variations that don't matter to your site (eg: */?utm_source=Twitter*).

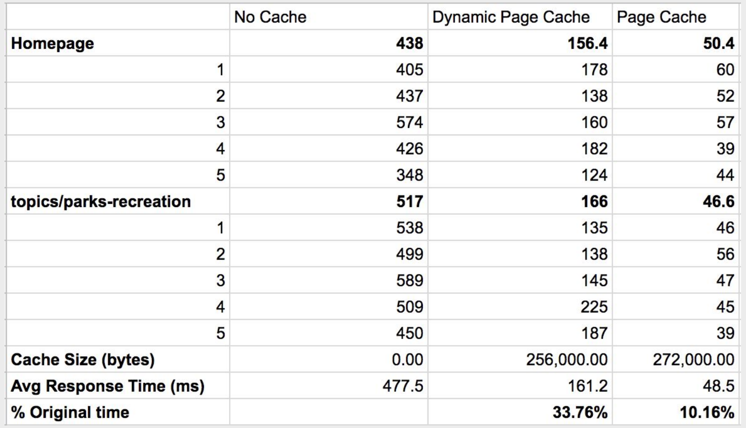

| | No Cache | Dynamic Page Cache | Page Cache |
|---|---|---|---|
| **Homepage** | **438** | **156.4** | **50.4** |
| 1 | 405 | 178 | 60 |
| 2 | 437 | 138 | 52 |
| 3 | 574 | 160 | 57 |
| 4 | 426 | 182 | 39 |
| 5 | 348 | 124 | 44 |
| **topics/parks-recreation** | **517** | **166** | **46.6** |
| 1 | 538 | 135 | 46 |
| 2 | 499 | 138 | 56 |
| 3 | 589 | 145 | 47 |
| 4 | 509 | 225 | 45 |
| 5 | 450 | 187 | 39 |
| **Cache Size (bytes)** | 0.00 | 256,000.00 | 272,000.00 |
| **Avg Response Time (ms)** | 477.5 | 161.2 | 48.5 |
| **% Original time** | | **33.76%** | **10.16%** |

## February, 2018

Begin Crawling the *CD* environment nightly.

We intended this as a nightly error check, but it provides a pretty good daily performance audit for pre-production code, too!

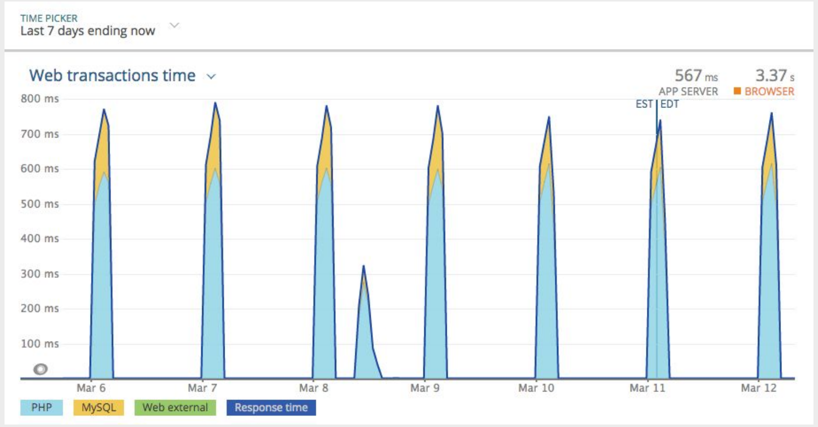

# February, 2018

CDN Cache Lifetime goes from 60 minutes to 30 minutes

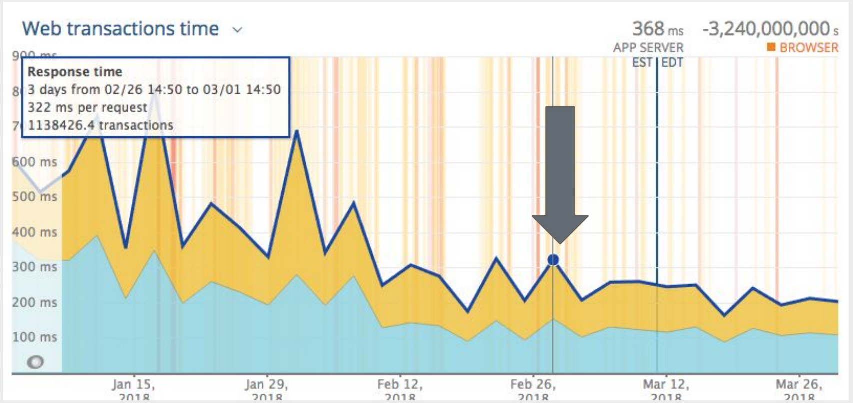

Response time
3 days from 02/26 14:50 to 03/01 14:50
322 ms per request
1138426.4 transactions

# Overall Effect

Time to publish is cut in half, with average response time cut by over 50%

Nov 1: 526ms

Mar 29: 227ms

This has resulted in a 0 disruptions due to performance issues since we completed our work.

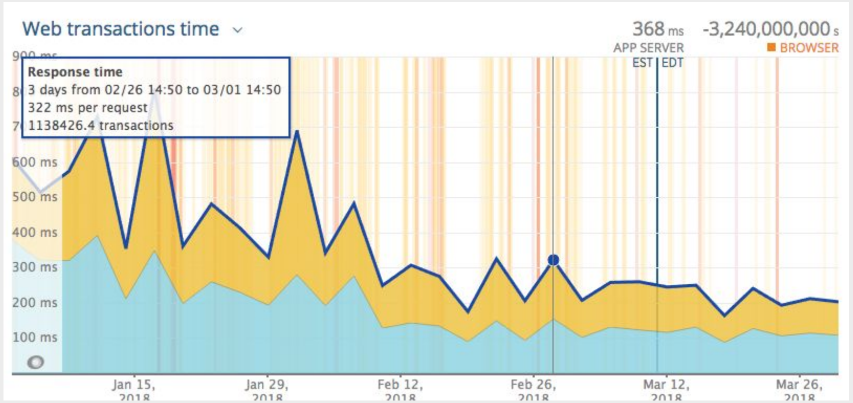

Response time
3 days from 02/26 14:50 to 03/01 14:50
322 ms per request
1138426.4 transactions

## Frontend Performance:

Constituents need **fast access** to information across many devices and varying network connections.

Our research:

0-1s load time: 1.6% bounce rate

1-3s load time: 10.4% bounce rate

3-7s load time: 22.7% bounce rate

7-13s load time: 24.7% bounce rate

# April, 2018

Enable HTTP/2

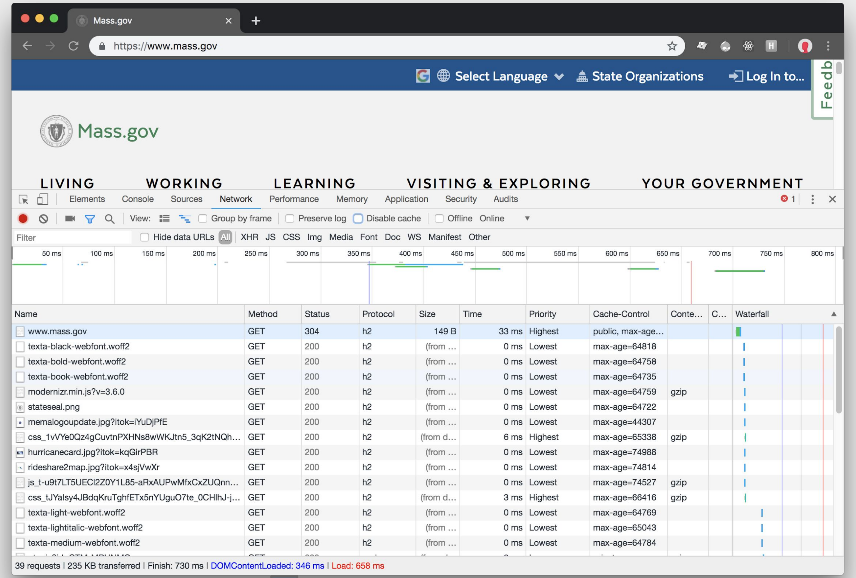

| Name | Method | Status | Protocol | Size | Time | Priority | Cache-Control | Conte… | C… | Waterfall |
|---|---|---|---|---|---|---|---|---|---|---|
| www.mass.gov | GET | 304 | h2 | 149 B | 33 ms | Highest | public, max-age… | | | |
| texta-black-webfont.woff2 | GET | 200 | h2 | (from … | 0 ms | Lowest | max-age=64818 | | | |
| texta-bold-webfont.woff2 | GET | 200 | h2 | (from … | 0 ms | Lowest | max-age=64758 | | | |
| texta-book-webfont.woff2 | GET | 200 | h2 | (from … | 0 ms | Lowest | max-age=64735 | | | |
| modernizr.min.js?v=3.6.0 | GET | 200 | h2 | (from … | 0 ms | Lowest | max-age=64759 | gzip | | |
| stateseal.png | GET | 200 | h2 | (from … | 0 ms | Lowest | max-age=64722 | | | |
| memalogoupdate.jpg?itok=iYuDjPfE | GET | 200 | h2 | (from … | 0 ms | Lowest | max-age=44307 | | | |
| css_1vVYe0Qz4gCuvtnPXHNs8wWKJtn5_3qK2tNQh… | GET | 200 | h2 | (from d… | 6 ms | Highest | max-age=65338 | gzip | | |
| hurricanecard.jpg?itok=kqGirPBR | GET | 200 | h2 | (from … | 0 ms | Lowest | max-age=74988 | | | |
| rideshare2map.jpg?itok=x4sjVwXr | GET | 200 | h2 | (from … | 0 ms | Lowest | max-age=74814 | | | |
| js_t-u9t7LT5UECl2Z0Y1L85-aRxAUPwMfxCxZUQnn… | GET | 200 | h2 | (from … | 0 ms | Lowest | max-age=74527 | gzip | | |
| css_tJYalsy4JBdqKruTghfETx5nYUguO7te_0CHlhJ-j… | GET | 200 | h2 | (from d… | 3 ms | Highest | max-age=66416 | gzip | | |
| texta-light-webfont.woff2 | GET | 200 | h2 | (from … | 0 ms | Lowest | max-age=64769 | | | |
| texta-lightitalic-webfont.woff2 | GET | 200 | h2 | (from … | 0 ms | Lowest | max-age=65043 | | | |
| texta-medium-webfont.woff2 | GET | 200 | h2 | (from … | 0 ms | Lowest | max-age=64784 | | | |

39 requests | 235 KB transferred | Finish: 730 ms | DOMContentLoaded: 346 ms | Load: 658 ms

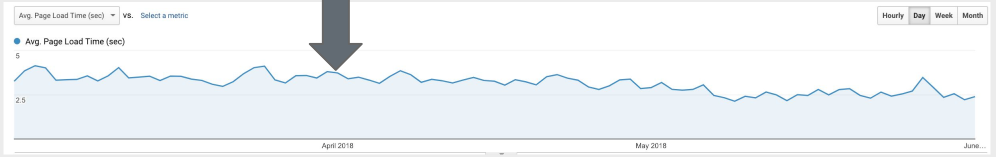

## April, 2018

Switch from directly inlining SVG to embedding with a *use* statement and inlining each used SVG 1x per page.

We discovered SVG was taking up a large % of the initial HTML response due to duplicated icons on the same page.

We came up with a solution to embed each icon that was used on a page 1 time in the header, and reference it multiple times with an SVG *use* statement.

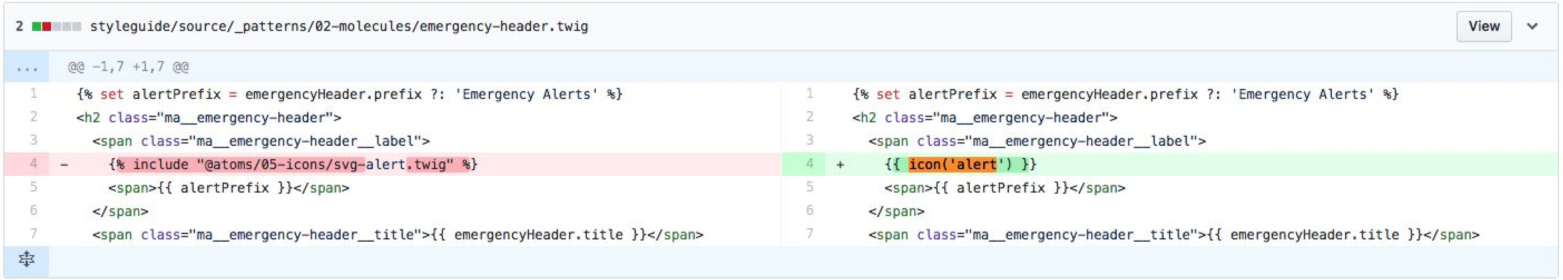

2 styleguide/source/_patterns/02-molecules/emergency-header.twig

```diff
@@ -1,7 +1,7 @@
   {% set alertPrefix = emergencyHeader.prefix ?: 'Emergency Alerts' %}
   <h2 class="ma__emergency-header">
     <span class="ma__emergency-header__label">
-      {% include "@atoms/05-icons/svg-alert.twig" %}
+      {{ icon('alert') }}
       <span>{{ alertPrefix }}</span>
     </span>
     <span class="ma__emergency-header__title">{{ emergencyHeader.title }}</span>
```

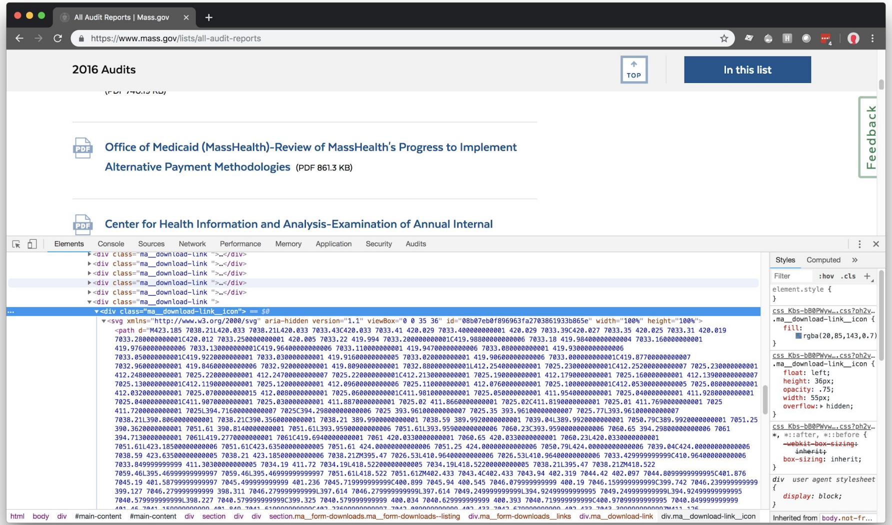

All Audit Reports | Mass.gov

https://www.mass.gov/lists/all-audit-reports

2016 Audits

In this list

Office of Medicaid (MassHealth)-Review of MassHealth's Progress to Implement Alternative Payment Methodologies (PDF 861.3 KB)

Center for Health Information and Analysis-Examination of Annual Internal

```
▸<div class="ma__download-link ">…</div>
▸<div class="ma__download-link ">…</div>
▸<div class="ma__download-link ">…</div>
▸<div class="ma__download-link ">…</div>
▸<div class="ma__download-link ">…</div>
▾<div class="ma__download-link ">
  ▾<div class="ma__download-link__icon"> == $0
    ▾<svg xmlns="http://www.w3.org/2000/svg" aria-hidden version="1.1" viewBox="0 0 35 36" id="08b07eb0f896963fa2703861933b865e" width="100%" height="100%">
       <path d="M423.185 7038.21L420.033 7038.21L420.033 7033.43C420.033 7033.41 420.029 7033.400000000001 420.029 7033.39C420.027 7033.35 420.025 7033.31 420.019
       7033.280000000001C420.012 7033.250000000001 420.005 7033.22 419.994 7033.200000000001C419.98800000000006 7033.18 419.98400000000004 7033.160000000001
       419.97600000000006 7033.130000000001C419.96400000000006 7033.110000000001 419.94700000000006 7033.080000000001 419.93000000000006
       7033.050000000001C419.92200000000001 7033.030000000001 419.91600000000005 7033.020000000001 419.90600000000006 7033.000000000001C419.87700000000007
       7032.960000000001 419.84600000000006 7032.920000000001 419.80900000000001 7032.880000000001L412.25400000000001 7025.230000000001C412.25200000000007 7025.230000000001
       412.248000000001 7025.220000000001 412.24700000000007 7025.220000000001C412.21300000000001 7025.190000000001 412.179000000001 7025.160000000001 412.13900000000007
       7025.130000000001C412.11900000000001 7025.120000000001 412.09600000000006 7025.110000000001 412.076000000001 7025.100000000001C412.05300000000005 7025.080000000001
       412.032000000001 7025.0700000000015 412.008000000001 7025.060000000001C411.98100000000001 7025.050000000001 411.954000000001 7025.040000000001 411.928000000001
       7025.040000000001C411.907000000001 7025.030000000001 411.887000000001 7025.02 411.866000000001 7025.02C411.819000000001 7025.01 411.769000000001 7025
       411.720000000001 7025L394.71600000000007 7025C394.29800000000006 7025 393.96100000000007 7025.35 393.96100000000007 7025.77L393.96100000000007
       7038.21L390.806000000001 7038.21C390.356000000001 7038.21 389.990000000001 7038.59 389.992000000001 7039.04L389.992000000001 7050.79C389.992000000001 7051.25
       390.362000000001 7051.61 390.814000000001 7051.61L393.95900000000006 7051.61L393.95900000000006 7060.23C393.95900000000006 7060.65 394.29800000000006 7061
       394.713000000001 7061L419.277000000001 7061C419.694000000001 7061 420.033000000001 7060.65 420.033000000001 7060.23L420.033000000001
       7051.61L423.18500000000006 7051.61C423.63500000000005 7051.61 424.00000000000006 7051.25 424.00000000000006 7050.79L424.00000000000006 7039.04C424.00000000000006
       7038.59 423.63500000000005 7038.21 423.18500000000006 7038.21ZM395.47 7026.53L410.96400000000006 7026.53L410.96400000000006 7033.42999999999C410.96400000000006
       7033.849999999999 411.30300000000005 7034.19 411.72 7034.19L418.52200000000005 7034.19L418.52200000000005 7038.21L395.47 7038.21ZM418.522
       7059.46L395.46999999999997 7059.46L395.46999999999997 7051.61L418.522 7051.61ZM402.433 7043.4C402.433 7043.94 402.319 7044.42 402.097 7044.809999999995C401.876
       7045.19 401.58799999999997 7045.499999999999 401.236 7045.719999999999C400.899 7045.94 400.545 7046.07999999999 400.19 7046.159999999999C399.742 7046.239999999999
       399.127 7046.279999999999 398.311 7046.279999999999L397.614 7046.279999999999L397.614 7049.249999999999L394.92499999999995 7049.249999999999L394.92499999999995
       7040.579999999999L398.227 7040.579999999999C399.325 7040.579999999999 400.034 7040.629999999999 400.393 7040.719999999999C400.97099999999995 7040.849999999999
```

html body div #main-content #main-content div section div div section.ma__form-downloads.ma__form-downloads--listing div.ma__form-downloads__links div.ma__download-link div.ma__download-link__icon

```
element.style {
}
css_Kbs-bB0PWyw….css?ph2v…
.ma__download-link__icon {
  fill:
    ■rgba(20,85,143,0.7)
}
css_Kbs-bB0PWyw….css?ph2v…
.ma__download-link__icon {
  float: left;
  height: 36px;
  opacity: .75;
  width: 55px;
  overflow: ▸ hidden;
}
css_Kbs-bB0PWyw….css?ph2v…
*, *::after, *::before {
  ~~-webkit-box-sizing:
    inherit;~~
  box-sizing: inherit;
}
div  user agent stylesheet
{
  display: block;
}
Inherited from body.not-fr…
```

Size: > 1MB, Load: > 30s

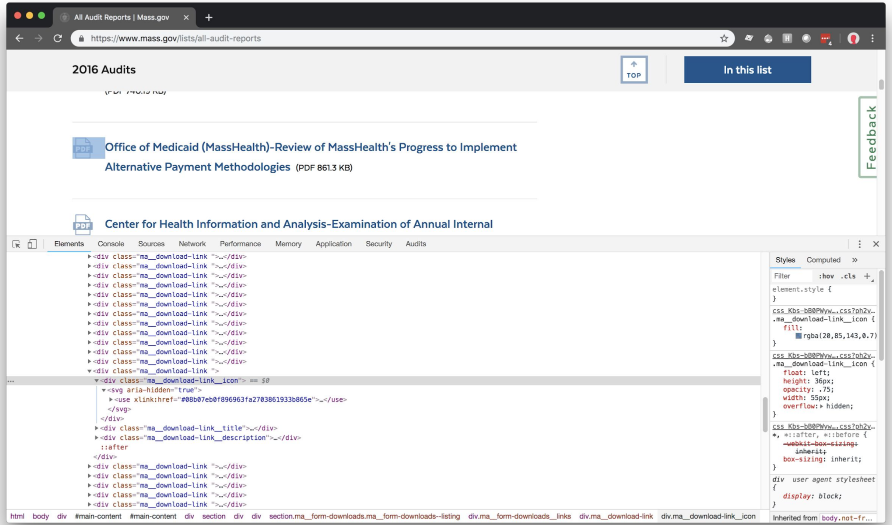

All Audit Reports | Mass.gov

https://www.mass.gov/lists/all-audit-reports

2016 Audits

In this list

Office of Medicaid (MassHealth)-Review of MassHealth's Progress to Implement Alternative Payment Methodologies (PDF 861.3 KB)

Center for Health Information and Analysis-Examination of Annual Internal

```
▸<div class="ma__download-link ">…</div>
▸<div class="ma__download-link ">…</div>
▸<div class="ma__download-link ">…</div>
▸<div class="ma__download-link ">…</div>
▸<div class="ma__download-link ">…</div>
▸<div class="ma__download-link ">…</div>
▸<div class="ma__download-link ">…</div>
▸<div class="ma__download-link ">…</div>
▸<div class="ma__download-link ">…</div>
▸<div class="ma__download-link ">…</div>
▸<div class="ma__download-link ">…</div>
▸<div class="ma__download-link ">…</div>
▾<div class="ma__download-link ">
  ▾<div class="ma__download-link__icon"> == $0
    ▾<svg aria-hidden="true">
      ▸<use xlink:href="#08b07eb0f896963fa2703861933b865e">…</use>
     </svg>
   </div>
  ▸<div class="ma__download-link__title">…</div>
  ▸<div class="ma__download-link__description">…</div>
   ::after
 </div>
▸<div class="ma__download-link ">…</div>
▸<div class="ma__download-link ">…</div>
▸<div class="ma__download-link ">…</div>
▸<div class="ma__download-link ">…</div>
▸<div class="ma__download-link ">…</div>
```

html body div #main-content #main-content div section div div section.ma__form-downloads.ma__form-downloads--listing div.ma__form-downloads__links div.ma__download-link div.ma__download-link__icon

HTML 78Kb, Load: 2.6s

LAST CALL MEDIA

# April, 2018

Switch from directly inlining SVG to embedding with a *use* statement and inlining each used SVG 1x per page.

Backend cache sizes also dropped tremendously, since we were no longer storing all of that inline SVG in *cache_render,* and *cache_dynamic_page_cache*.

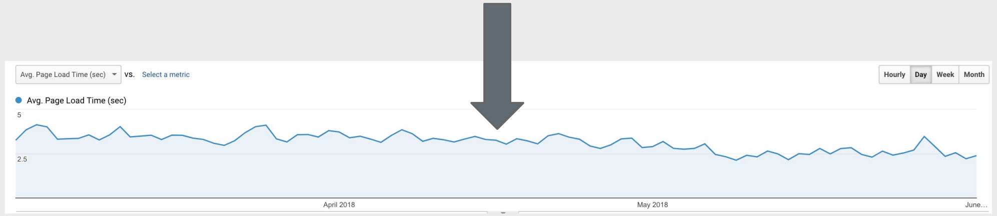

## April, 2018

Audit, deduplicate, and clean up javascript

No longer loaded:

- 1 extra copy of Modernizr
- 1 extra copy of Handlebars
- Several deprecated custom JS files
- Google Maps (only loaded where needed)
- Google CSE (lazy loaded where needed)

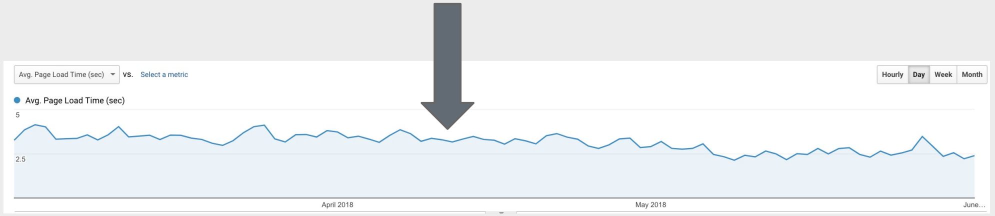

# April, 2018

Reformulate JSONAPI AJAX request made on every page to be more specific about what's it's asking for

## Before (781k before gzip):

```
https://www.mass.gov/jsonapi/node/alert?include=field_target_pag
es_para_ref,field_alert&filter[status][value]=1
```

## After (16k before gzip):

```
https://www.mass.gov/jsonapi/node/alert?page[limit]=250&sort=-ch
anged&include=field_target_pages_para_ref,field_alert&filter[st
atus][value]=1&fields[node--alert]=title,changed,entity_url,fie
ld_alert_severity,field_alert,field_target_pages_para_ref,field
_alert_display&fields[paragraph--emergency_alert]=id,changed,fi
eld_emergency_alert_timestamp,field_emergency_alert_message,fie
ld_emergency_alert_link,field_emergency_alert_content&fields[pa
ragraph--target_pages]=field_target_content_ref
```

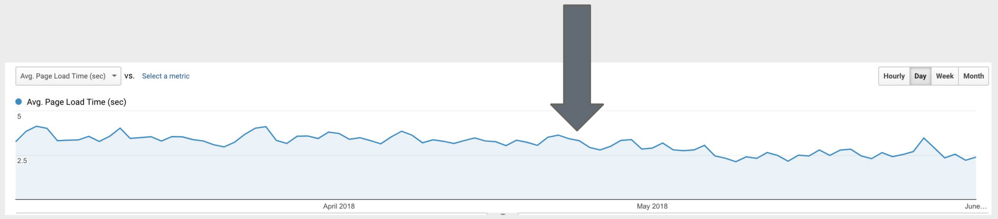

# April, 2018

Add WOFF2 webfont variants and use preloading to shorten the critical asset chain.

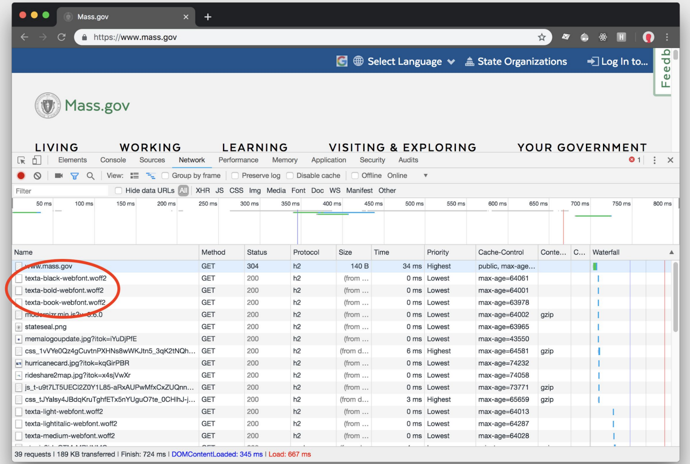

| Name | Method | Status | Protocol | Size | Time | Priority | Cache-Control | Conte… | C… | Waterfall |
|---|---|---|---|---|---|---|---|---|---|---|
| www.mass.gov | GET | 304 | h2 | 140 B | 34 ms | Highest | public, max-age… | | | |
| texta-black-webfont.woff2 | GET | 200 | h2 | (from … | 0 ms | Lowest | max-age=64061 | | | |
| texta-bold-webfont.woff2 | GET | 200 | h2 | (from … | 0 ms | Lowest | max-age=64001 | | | |
| texta-book-webfont.woff2 | GET | 200 | h2 | (from … | 0 ms | Lowest | max-age=63978 | | | |
| modernizr.min.js?v=3.6.0 | GET | 200 | h2 | (from … | 0 ms | Lowest | max-age=64002 | gzip | | |
| stateseal.png | GET | 200 | h2 | (from … | 0 ms | Lowest | max-age=63965 | | | |
| memalogoupdate.jpg?itok=iYuDjPfE | GET | 200 | h2 | (from … | 0 ms | Lowest | max-age=43550 | | | |
| css_1vVYe0Qz4gCuvtnPXHNs8wWKJtn5_3qK2tNQh… | GET | 200 | h2 | (from d… | 6 ms | Highest | max-age=64581 | gzip | | |
| hurricanecard.jpg?itok=kqGirPBR | GET | 200 | h2 | (from … | 0 ms | Lowest | max-age=74232 | | | |
| rideshare2map.jpg?itok=x4sjVwXr | GET | 200 | h2 | (from … | 0 ms | Lowest | max-age=74058 | | | |
| js_t-u9t7LT5UECl2Z0Y1L85-aRxAUPwMfxCxZUQnn… | GET | 200 | h2 | (from … | 0 ms | Lowest | max-age=73771 | gzip | | |
| css_tJYalsy4JBdqKruTghfETx5nYUguO7te_0CHlhJ-j… | GET | 200 | h2 | (from d… | 3 ms | Highest | max-age=65659 | gzip | | |
| texta-light-webfont.woff2 | GET | 200 | h2 | (from … | 0 ms | Lowest | max-age=64013 | | | |
| texta-lightitalic-webfont.woff2 | GET | 200 | h2 | (from … | 0 ms | Lowest | max-age=64287 | | | |
| texta-medium-webfont.woff2 | GET | 200 | h2 | (from … | 0 ms | Lowest | max-age=64028 | | | |

39 requests | 189 KB transferred | Finish: 724 ms | DOMContentLoaded: 345 ms | Load: 667 ms

# April, 2018

Add WOFF2 webfont variants and use preloading to shorten the critical asset chain.

- WOFF2 is about 25% smaller.
- Preloading allows the browser to initiate fetching the fonts before it even requests the CSS file that declares their usage. Fonts are needed for the first render.
- All browsers that support WOFF2 also support preloading.

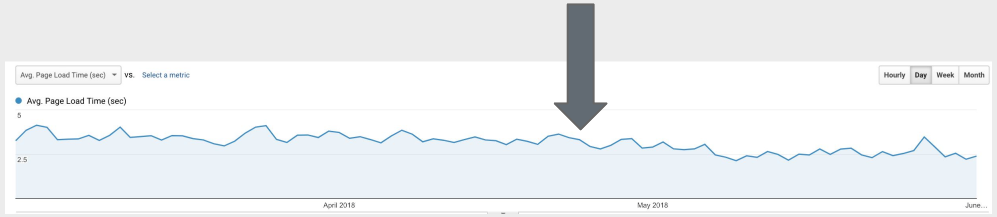

# May, 2018

Use imagick for “on the fly” image optimization

This simple toggle cut some of our larger hero images by over 100K!

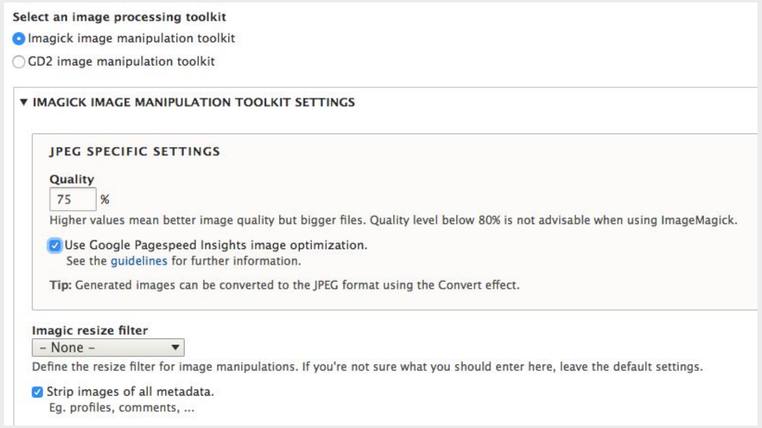

**Select an image processing toolkit**

- [✓ Imagick image manipulation toolkit]
- GD2 image manipulation toolkit

▼ **IMAGICK IMAGE MANIPULATION TOOLKIT SETTINGS**

**JPEG SPECIFIC SETTINGS**

**Quality**

75 %

Higher values mean better image quality but bigger files. Quality level below 80% is not advisable when using ImageMagick.

[✓] Use Google Pagespeed Insights image optimization.
See the guidelines for further information.

**Tip:** Generated images can be converted to the JPEG format using the Convert effect.

**Imagic resize filter**

– None –

Define the resize filter for image manipulations. If you're not sure what you should enter here, leave the default settings.

[✓] Strip images of all metadata.
Eg. profiles, comments, ...

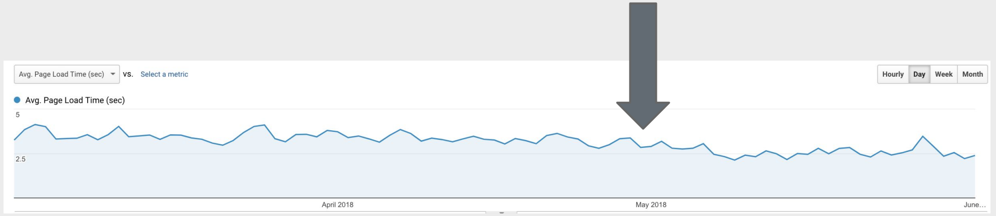

## May, 2018

Use a static feedback form rather than a javascript embed

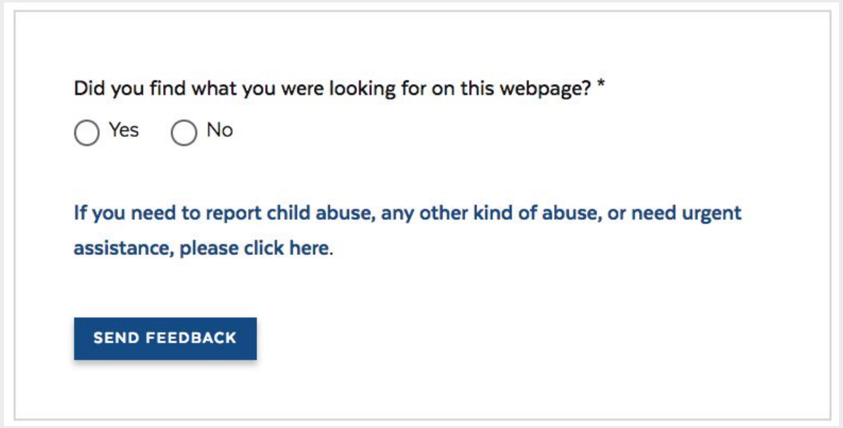
Did you find what you were looking for on this webpage? *

Yes No

If you need to report child abuse, any other kind of abuse, or need urgent assistance, please click here.

SEND FEEDBACK

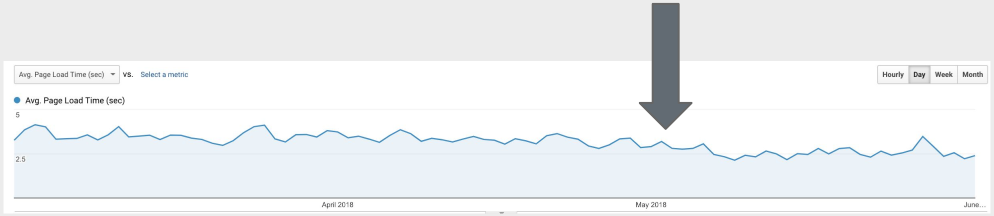

# May, 2018

Optimize the state seal and use it for both the header and footer.

- The "logo" is really the seal + some text, whereas it used to be an image.
- The footer now uses the same seal.
- Cut ~100k from every uncached pageview.

Mass.gov

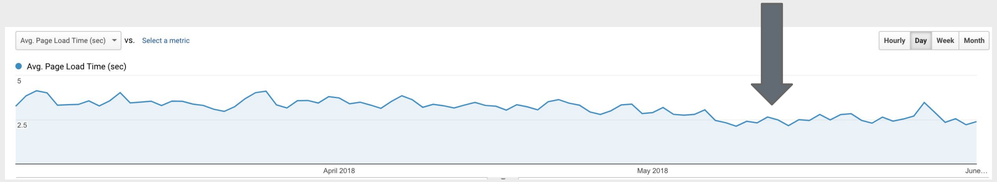

# May, 2018

Speedcurve diff of individual "Curated List" page

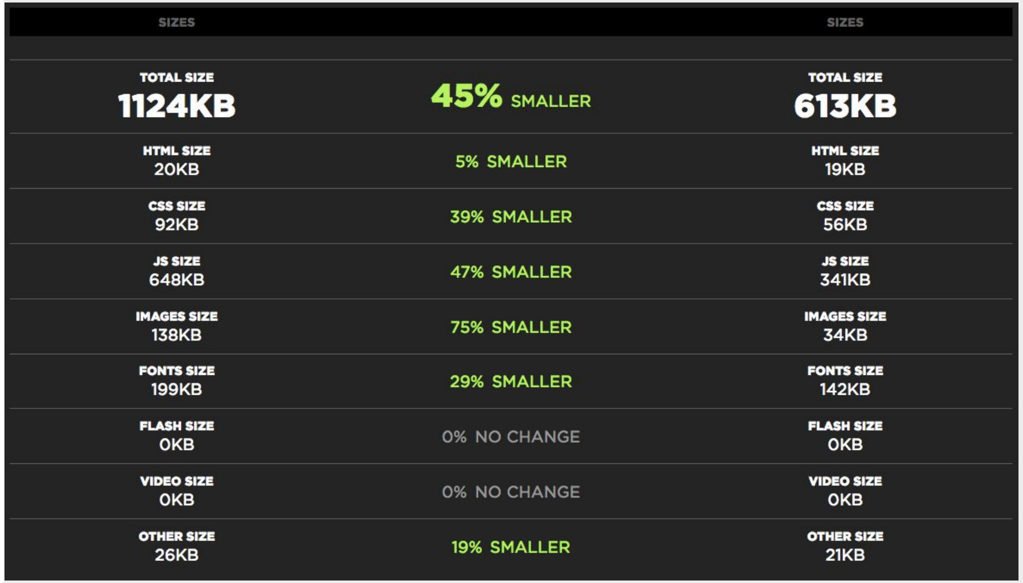

| SIZES | | SIZES |
|---|---|---|
| TOTAL SIZE 1124KB | 45% SMALLER | TOTAL SIZE 613KB |
| HTML SIZE 20KB | 5% SMALLER | HTML SIZE 19KB |
| CSS SIZE 92KB | 39% SMALLER | CSS SIZE 56KB |
| JS SIZE 648KB | 47% SMALLER | JS SIZE 341KB |
| IMAGES SIZE 138KB | 75% SMALLER | IMAGES SIZE 34KB |
| FONTS SIZE 199KB | 29% SMALLER | FONTS SIZE 142KB |
| FLASH SIZE 0KB | 0% NO CHANGE | FLASH SIZE 0KB |
| VIDEO SIZE 0KB | 0% NO CHANGE | VIDEO SIZE 0KB |
| OTHER SIZE 26KB | 19% SMALLER | OTHER SIZE 21KB |

## May, 2018

SpeedCurve visualization of homepage load

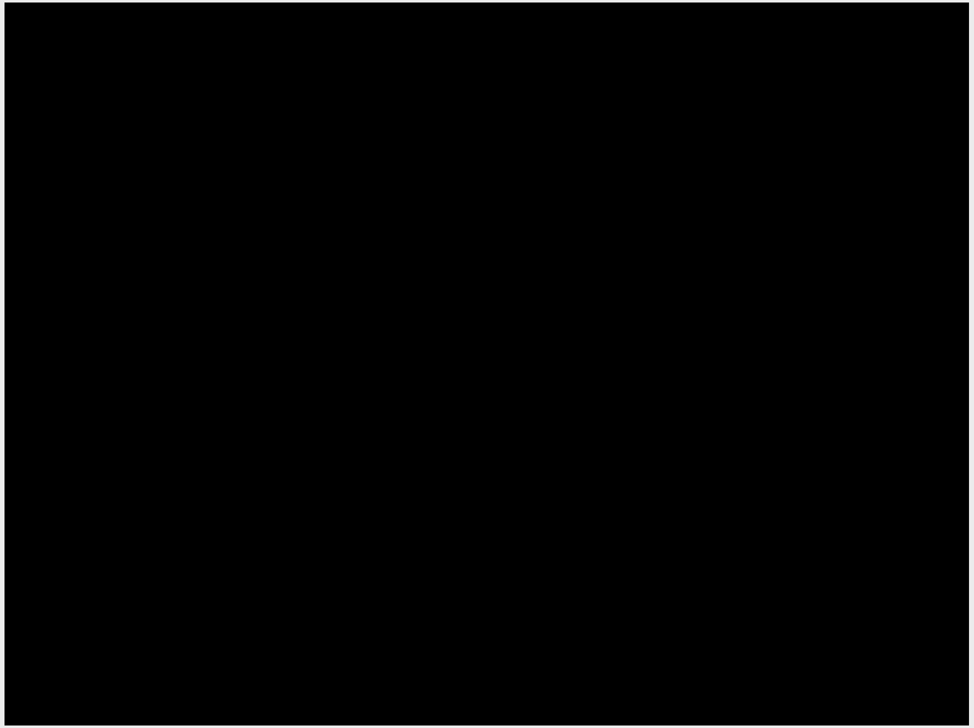

## May, 2018

SpeedCurve visualization of "REALID" guide

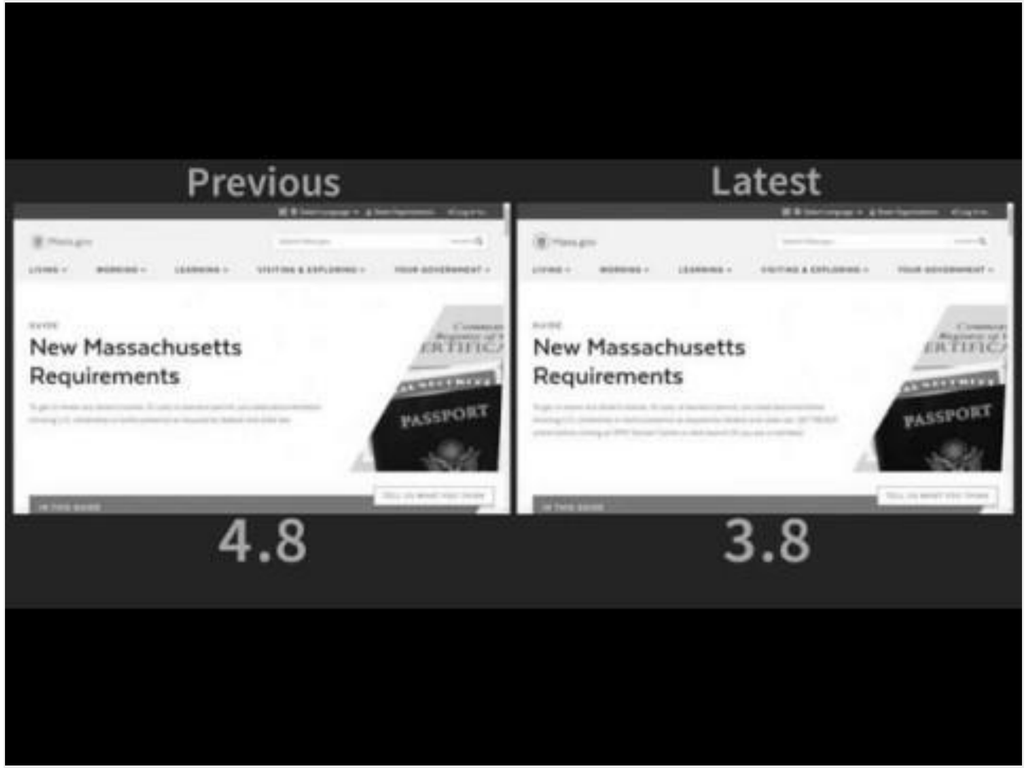
Previous
Latest
New Massachusetts Requirements
New Massachusetts Requirements
PASSPORT
PASSPORT
4.8
3.8

# Overall Effect

35% improvement (3.6s vs 2.4s)

(as per Google Analytics)

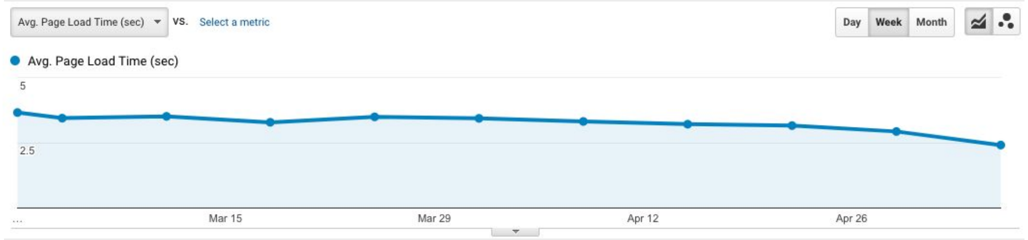

## Questions?

Psst… Like what you see here? Massachusetts Digital Services is hiring!
https://www.mass.gov/digital-services/jobs
Visit @MassGovDigital on Twitter or Medium

## December, 2017

Performance Documentation is added to the repository

Allows us to **share** a common set of performance standards with the rest of the team. Provides a **living standard** for how we think about performance on the site.

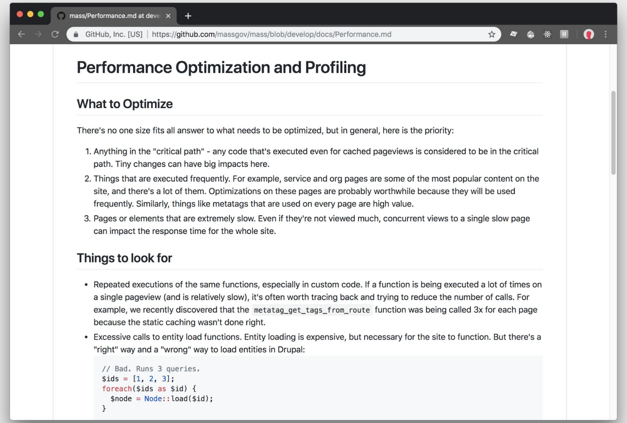

mass/Performance.md at deve

GitHub, Inc. [US] | https://github.com/massgov/mass/blob/develop/docs/Performance.md

### Performance Optimization and Profiling

#### What to Optimize

There's no one size fits all answer to what needs to be optimized, but in general, here is the priority:

1. Anything in the "critical path" - any code that's executed even for cached pageviews is considered to be in the critical path. Tiny changes can have big impacts here.
2. Things that are executed frequently. For example, service and org pages are some of the most popular content on the site, and there's a lot of them. Optimizations on these pages are probably worthwhile because they will be used frequently. Similarly, things like metatags that are used on every page are high value.
3. Pages or elements that are extremely slow. Even if they're not viewed much, concurrent views to a single slow page can impact the response time for the whole site.

#### Things to look for

- Repeated executions of the same functions, especially in custom code. If a function is being executed a lot of times on a single pageview (and is relatively slow), it's often worth tracing back and trying to reduce the number of calls. For example, we recently discovered that the `metatag_get_tags_from_route` function was being called 3x for each page because the static caching wasn't done right.
- Excessive calls to entity load functions. Entity loading is expensive, but necessary for the site to function. But there's a "right" way and a "wrong" way to load entities in Drupal:

```
// Bad. Runs 3 queries.
$ids = [1, 2, 3];
foreach($ids as $id) {
  $node = Node::load($id);
}
```

# February, 2018

Other, minor work

Batch load entities wherever you can:

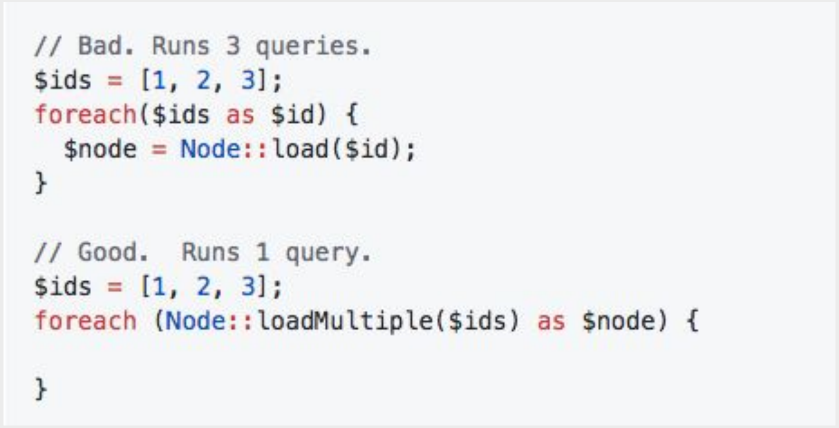

```
// Bad. Runs 3 queries.
$ids = [1, 2, 3];
foreach($ids as $id) {
  $node = Node::load($id);
}

// Good.  Runs 1 query.
$ids = [1, 2, 3];
foreach (Node::loadMultiple($ids) as $node) {

}
```

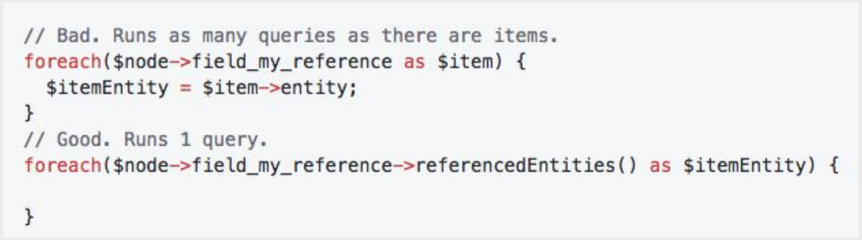

```
// Bad. Runs as many queries as there are items.
foreach($node->field_my_reference as $item) {
  $itemEntity = $item->entity;
}
// Good. Runs 1 query.
foreach($node->field_my_reference->referencedEntities() as $itemEntity) {

}
```

## February, 2018

Other, minor work

Fix Expensive Queries, including:

- Anything more complicated than a simple equality check or LIKE (avoid substring matching).
- COUNT queries
- Anything that's showing up in your slow query logs