# facebook

facebook

# Scaling Facebook Engineering

David Mortenson

# People on Facebook

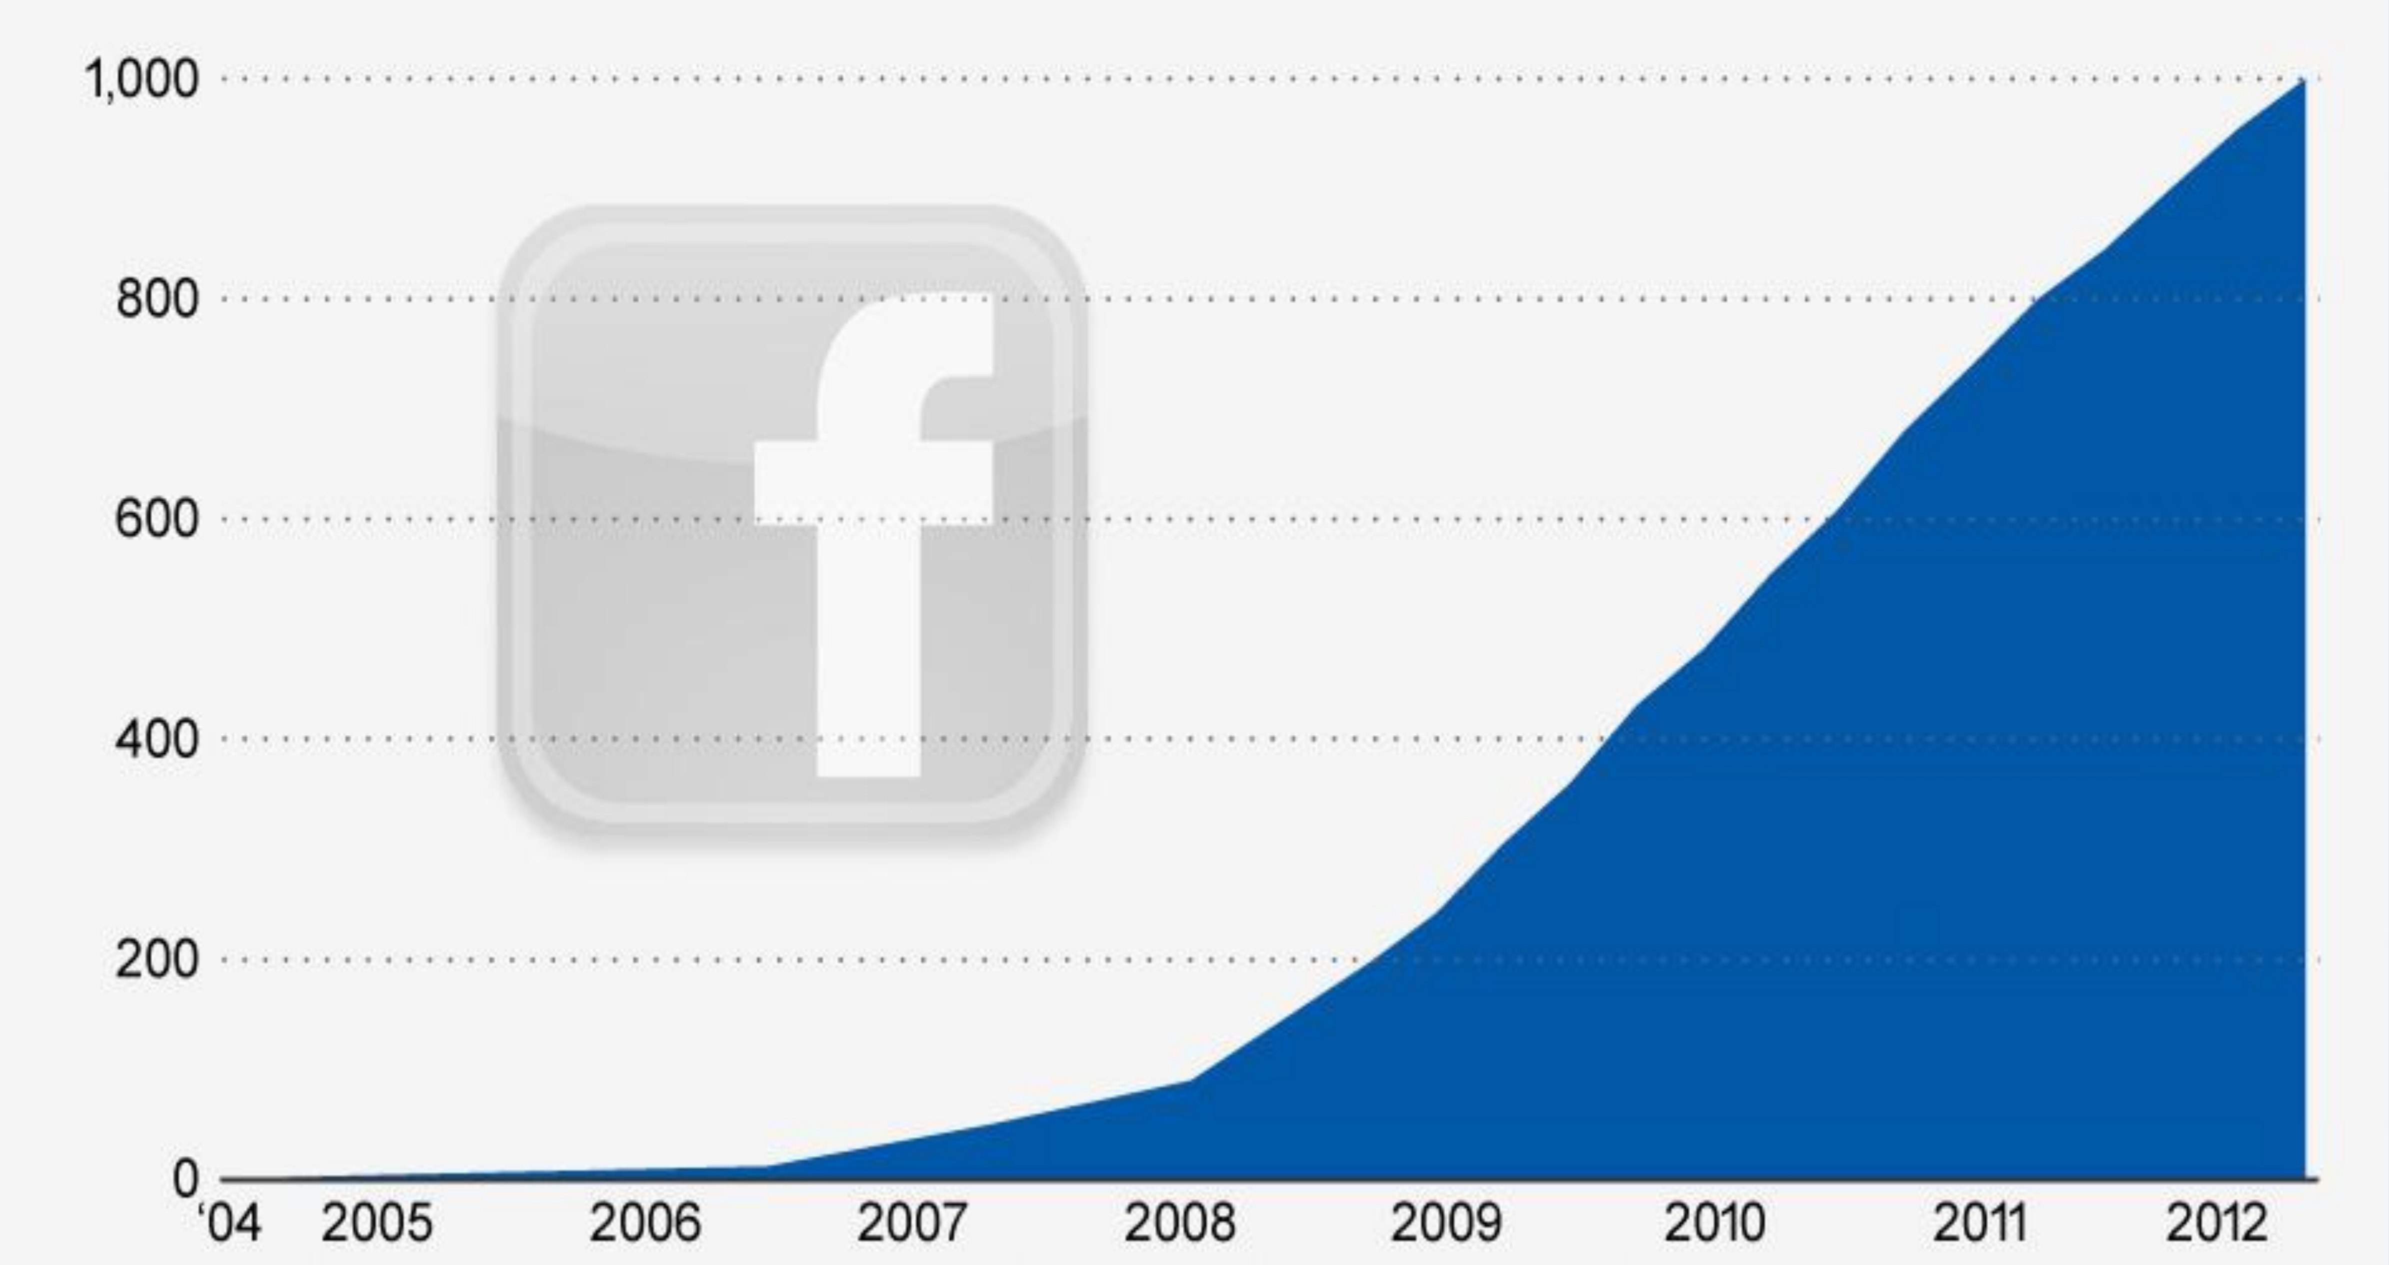

# New Products

- Timeline
- New iOS and Android Apps
- Facebook Messenger
- Ticker
- Open Graph
- Many more

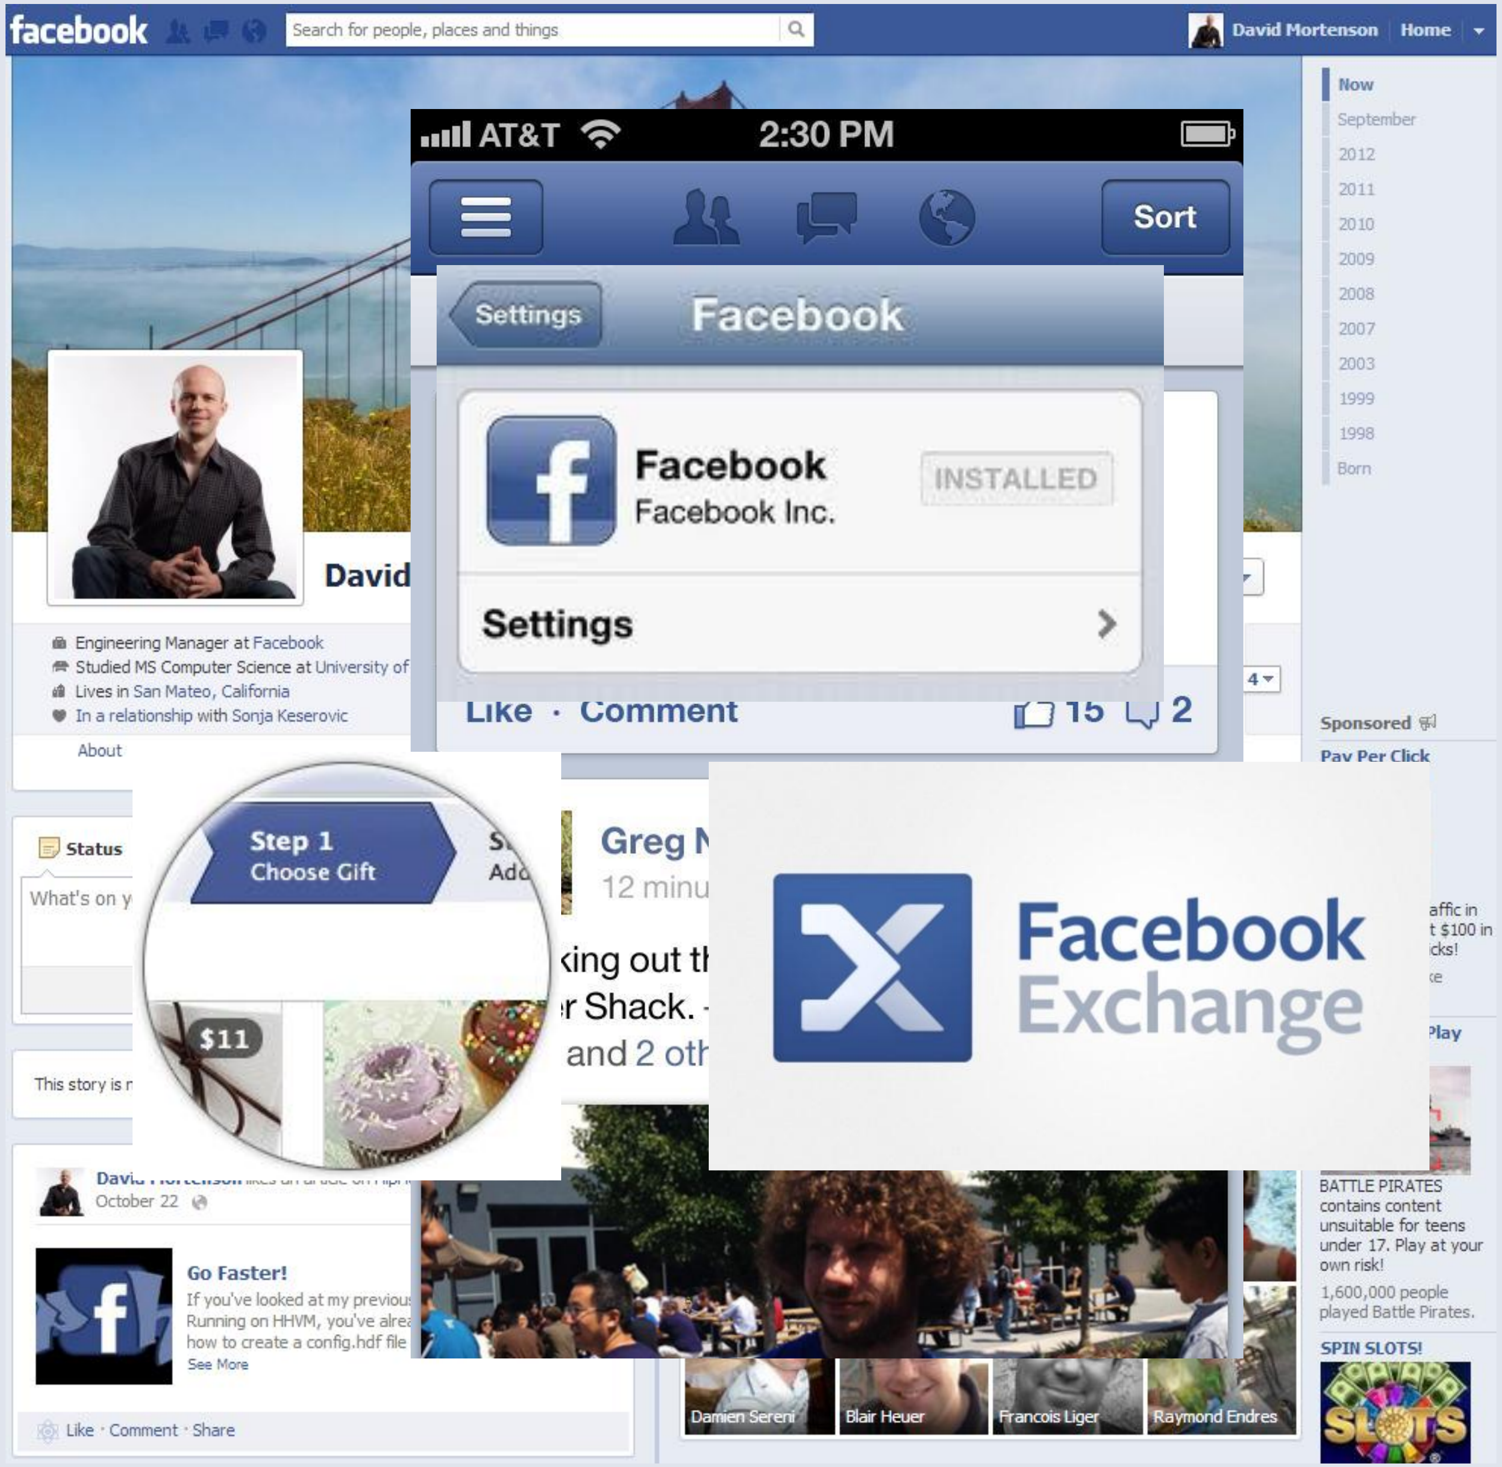
facebook
Search for people, places and things
David Mortenson
Home
Now
September
2012
2011
2010
2009
2008
2007
2003
1999
1998
Born
AT&T
2:30 PM
Sort
Settings
Facebook
Facebook
Facebook Inc.
INSTALLED
Settings
Like · Comment
15
2
David
Engineering Manager at Facebook
Studied MS Computer Science at University of
Lives in San Mateo, California
In a relationship with Sonja Keserovic
About
Status
Step 1
Choose Gift
$11
Greg N
12 minu
Facebook Exchange
Sponsored
Pay Per Click
Go Faster!
If you've looked at my previous
Running on HHVM, you've alrea
how to create a config.hdf file
See More
October 22
Like · Comment · Share
Damien Sereni
Blair Heuer
Francois Liger
Raymond Endres
BATTLE PIRATES contains content unsuitable for teens under 17. Play at your own risk!
1,600,000 people played Battle Pirates.
SPIN SLOTS!

# Facebook Engineering Team

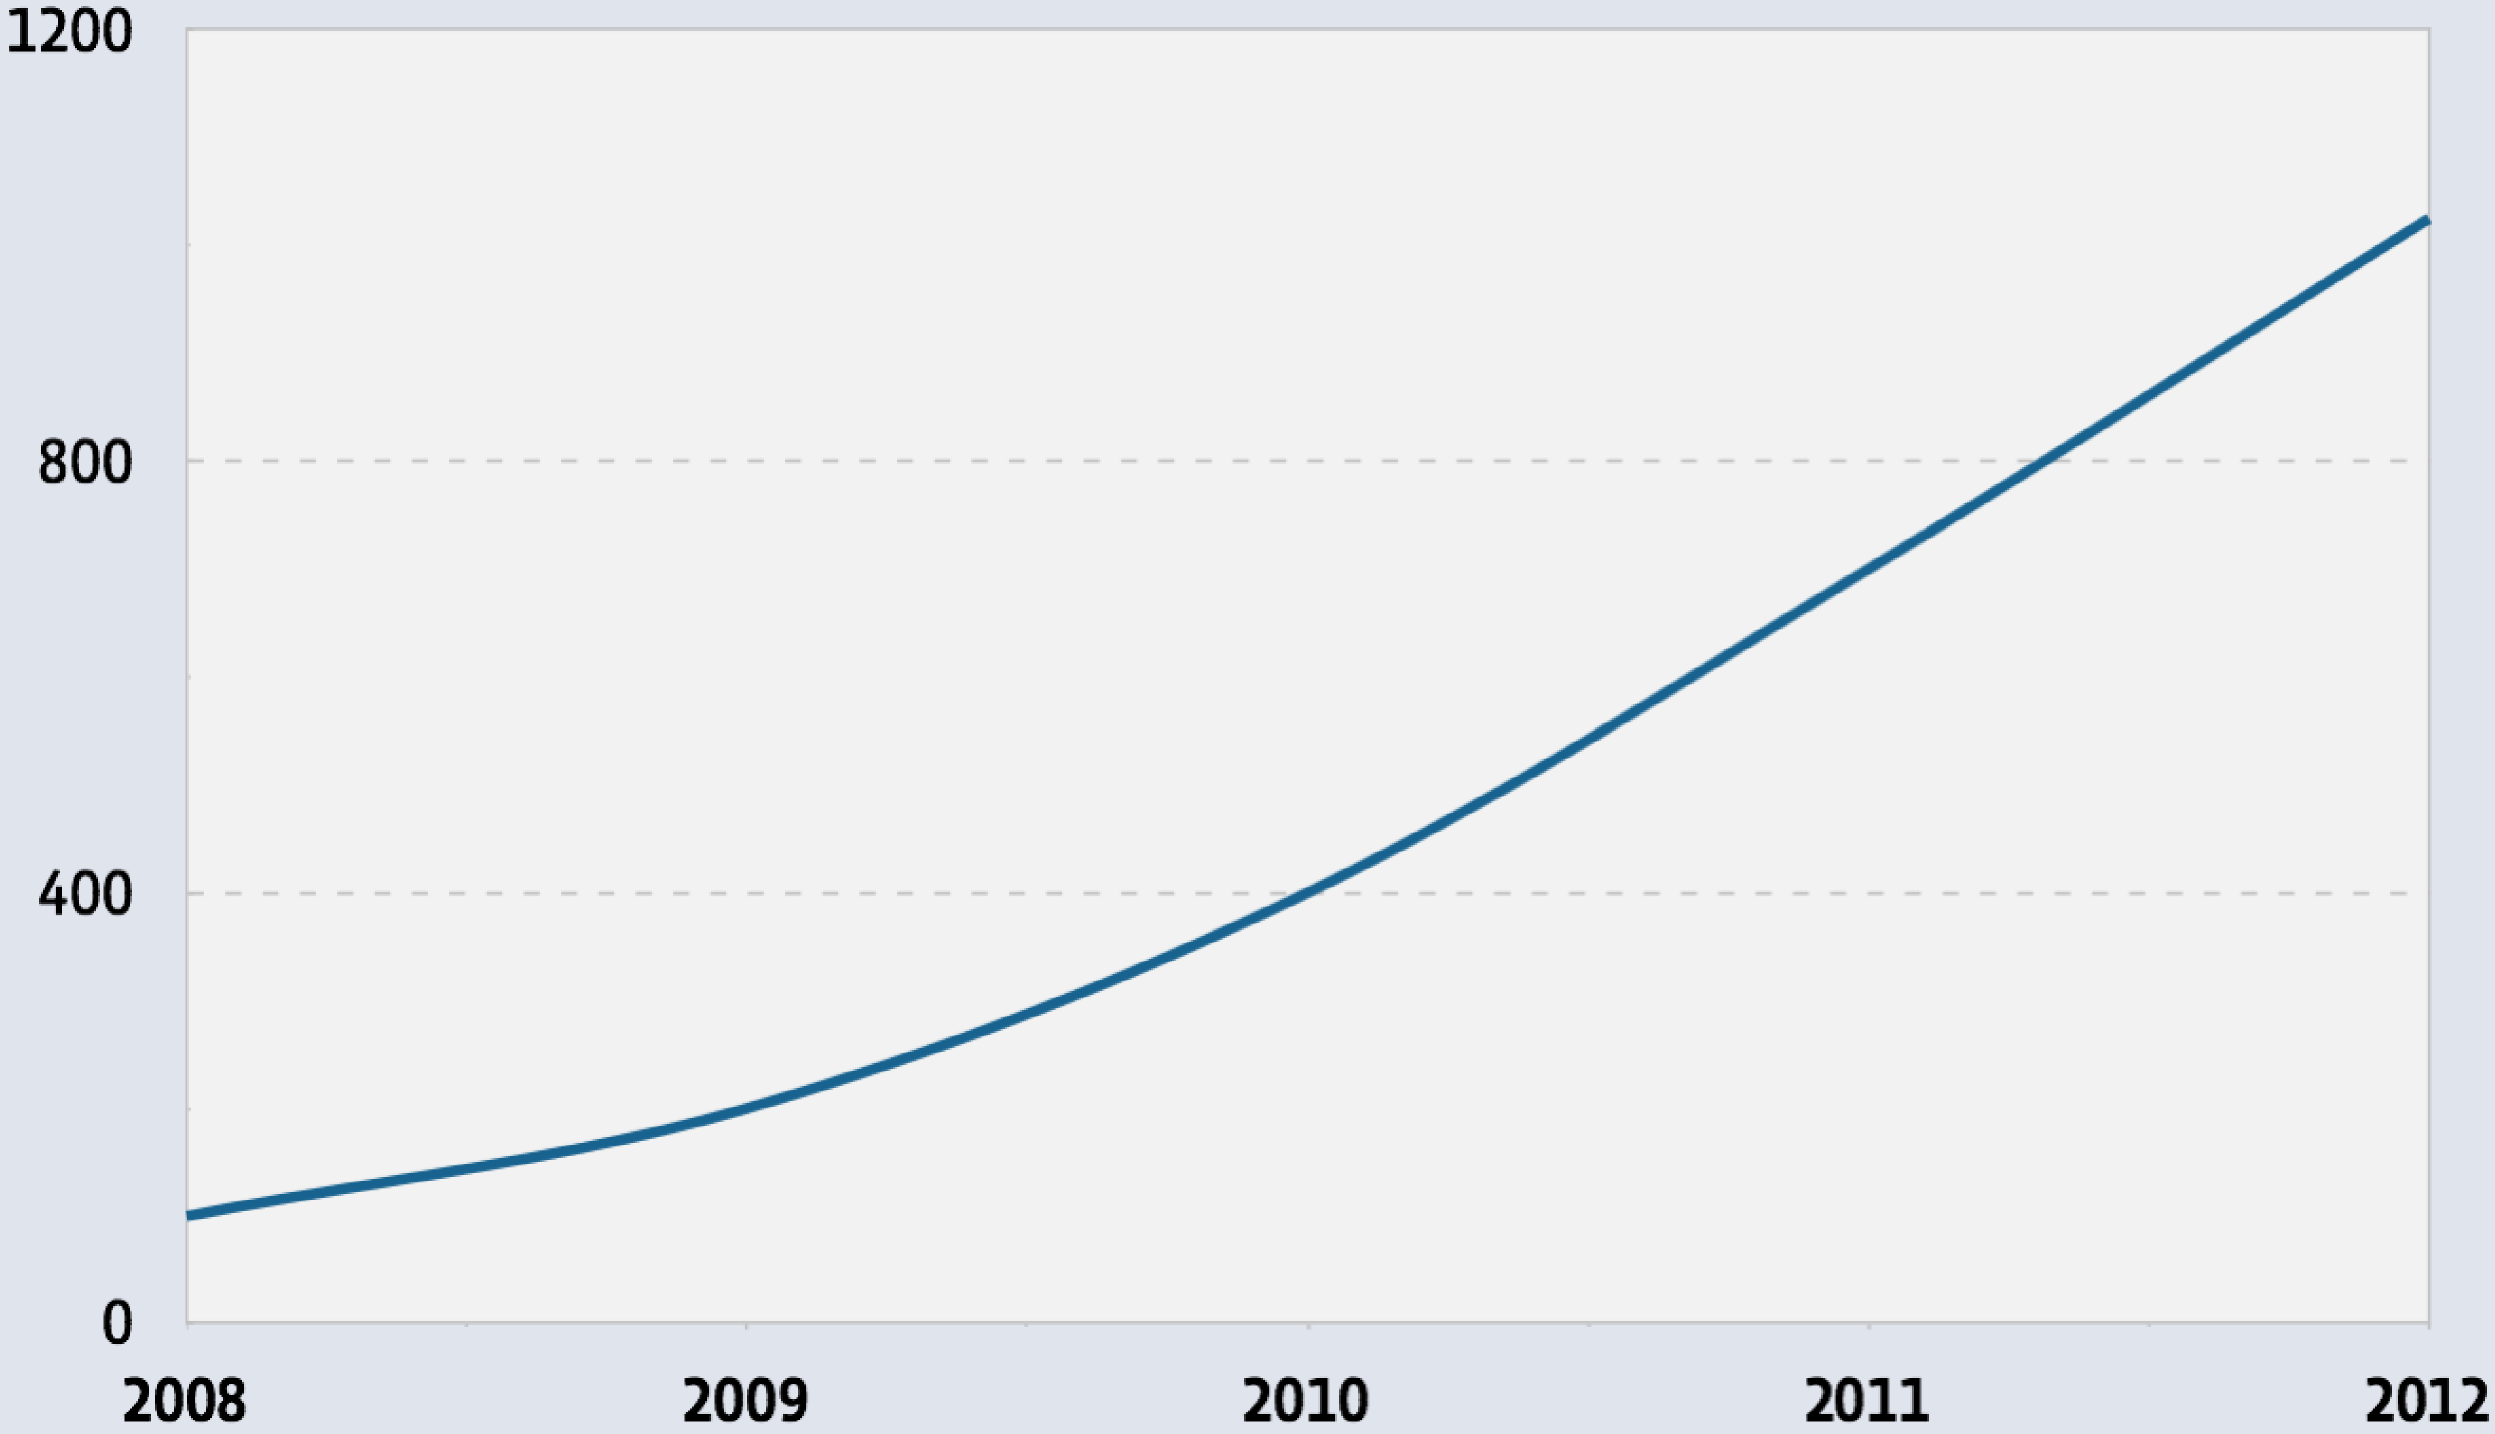

Goal: maintain or increase the efficiency per engineer as we grow

# Challenges

- The n00b time sink
- Keeping development fast
- Unintended consequences

# The n00b Time Sink

# Brief History: 2008

- Hiring 10 engineers / month
- Each team spending lots of time ramping up new engineers
- New engineers taking a while to come up to speed
- About to exceed Dumbar’s number of engineers (~150)

bootcamp

# Goals of Bootcamp

- Get new engineers up to speed technically and culturally
- Give exposure to many different areas of the codebase
- Find a team at the intersection of their passion and impact
- Get new engineers doing useful work as soon as possible

# Bootcamp Day 1: The Basics

- Meet your bootcamp mentor
- Get your dev server set up
- Learn about the core concepts of the FB www codebase
- Get your first bootcamp tasks

# Bootcamp Day 1: First Assignment

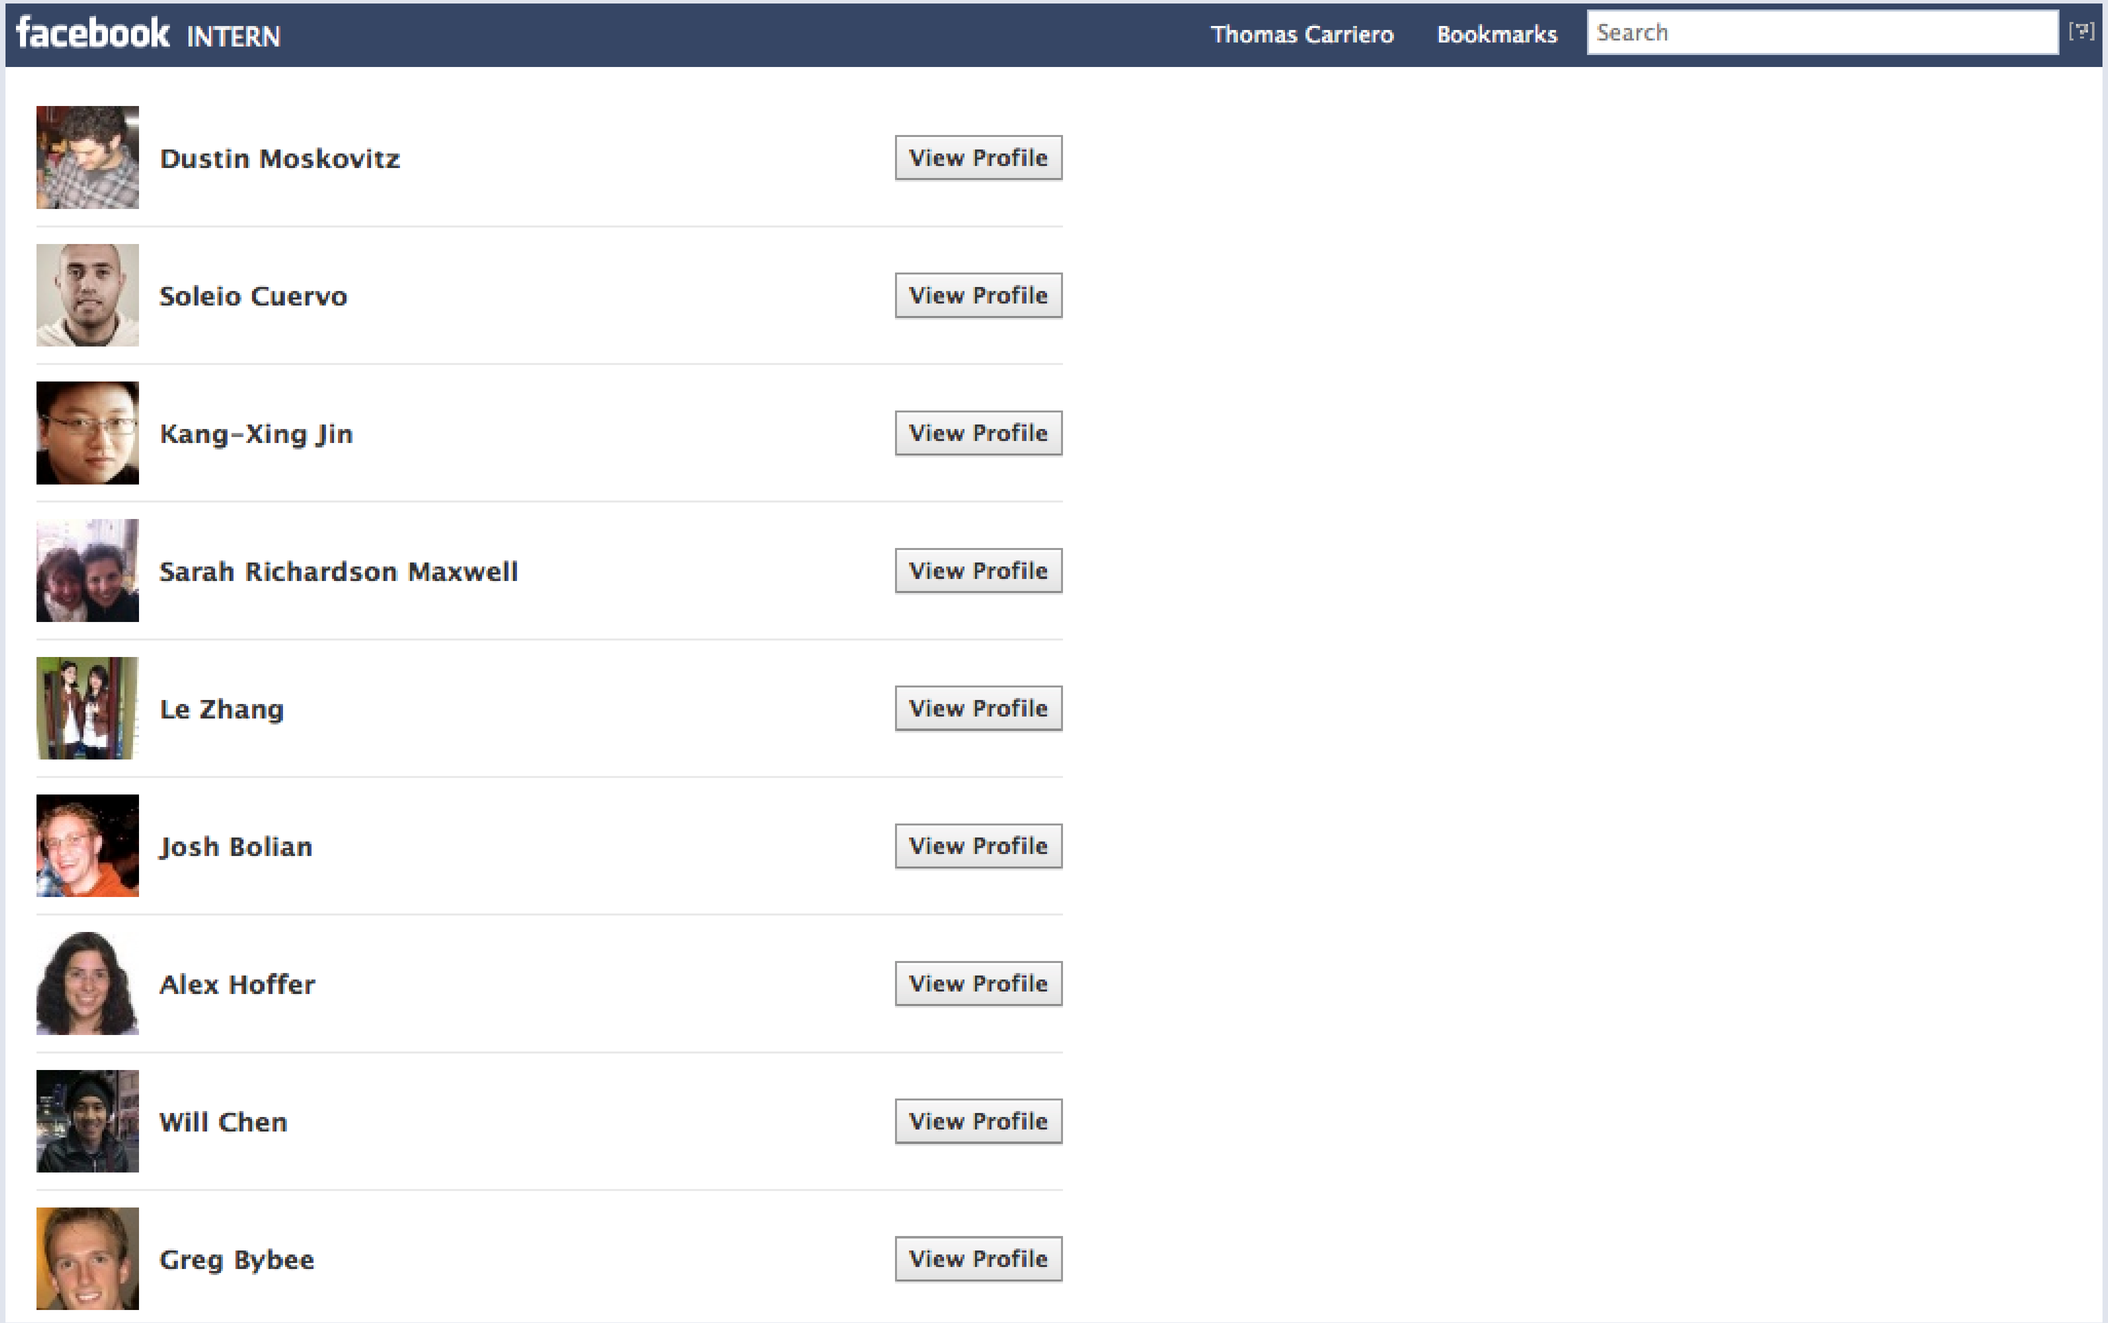

facebook INTERN

Thomas Carriero

Bookmarks

Search

Dustin Moskovitz

View Profile

Soleio Cuervo

View Profile

Kang-Xing Jin

View Profile

Sarah Richardson Maxwell

View Profile

Le Zhang

View Profile

Josh Bolian

View Profile

Alex Hoffer

View Profile

Will Chen

View Profile

Greg Bybee

View Profile

# Bootcamp Week 1: Learning & Hacking

- Intense training sessions on how things work at Facebook
- Customizing your dev environment
- Hang out with your cohort of bootcampers
- Complete & commit your 1st task

# Bootcamp Weeks 2-4

- Your first change goes live to a billion people!
- Training sessions on back-end services, mobile, network, data center design, etc
- Work on bootcamp tasks across FB stack

# Bootcamp Weeks 4-6: Team Selection

- Learn about the many teams that need engineers
- Meet the engineers from teams you are most interested in
- Bootcamp tasks for your top teams
- Pick the team that intersects impact and your passion

# Bootcamp Results

- Almost 1000 engineers graduated
- Feedback from bootcampers is consistently awesome
- New engineers come up to speed much faster
- Greatly reduced cost of ramping up new engineers
- New engineers get established and make friends faster
- Really helps new engineers understand Facebook culture

# Keeping Development Fast

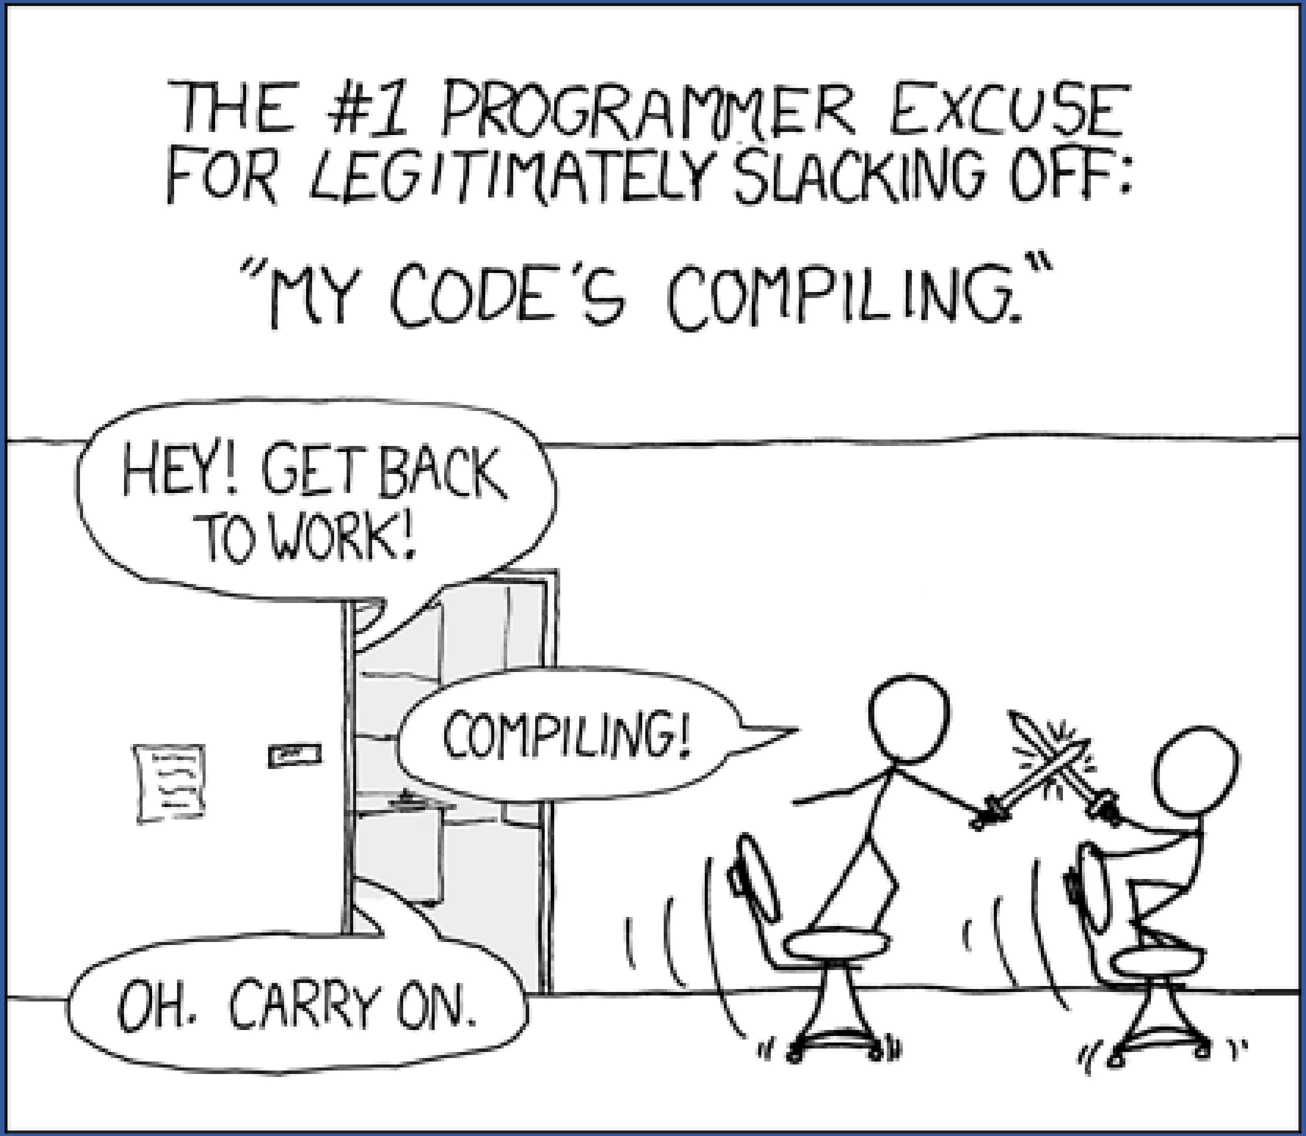

# Facebook Development Environment

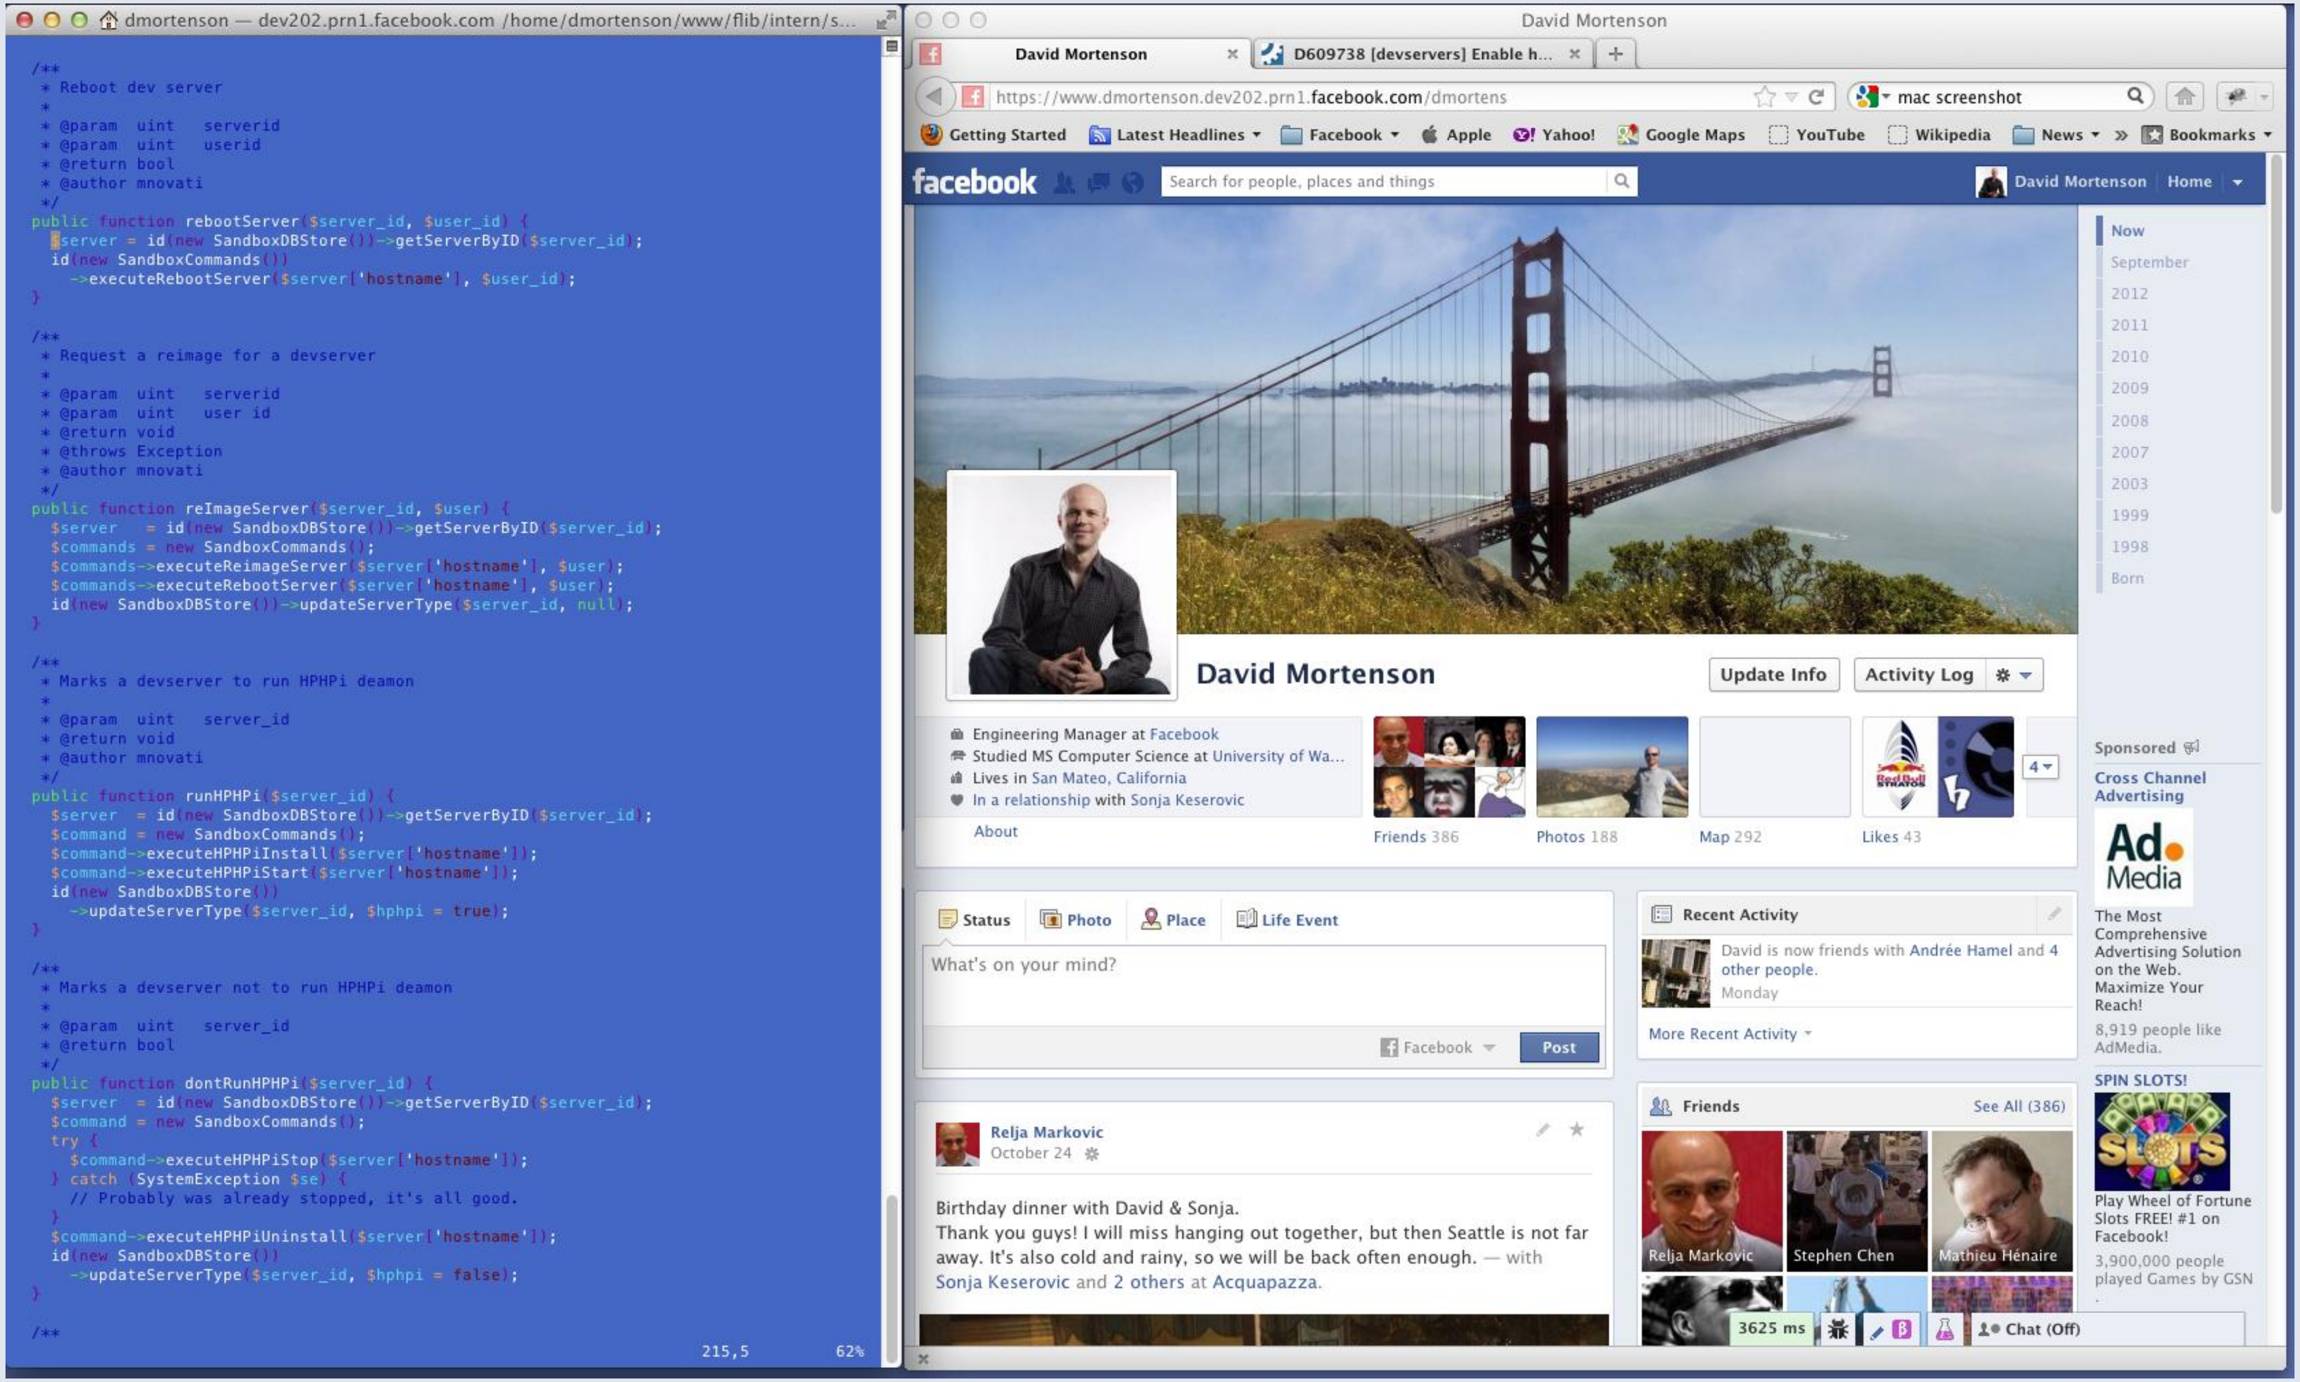

dmortenson — dev202.prn1.facebook.com /home/dmortenson/www/flib/intern/s...

```
/**
 * Reboot dev server
 *
 * @param  uint   serverid
 * @param  uint   userid
 * @return bool
 * @author mnovati
 */
public function rebootServer($server_id, $user_id) {
  $server = id(new SandboxDBStore())->getServerByID($server_id);
  id(new SandboxCommands())
    ->executeRebootServer($server['hostname'], $user_id);
}

/**
 * Request a reimage for a devserver
 *
 * @param  uint   serverid
 * @param  uint   user id
 * @return void
 * @throws Exception
 * @author mnovati
 */
public function reImageServer($server_id, $user) {
  $server   = id(new SandboxDBStore())->getServerByID($server_id);
  $commands = new SandboxCommands();
  $commands->executeReimageServer($server['hostname'], $user);
  $commands->executeRebootServer($server['hostname'], $user);
  id(new SandboxDBStore())->updateServerType($server_id, null);
}

/**
 * Marks a devserver to run HPHPi deamon
 *
 * @param  uint   server_id
 * @return void
 * @author mnovati
 */
public function runHPHPi($server_id) {
  $server  = id(new SandboxDBStore())->getServerByID($server_id);
  $command = new SandboxCommands();
  $command->executeHPHPiInstall($server['hostname']);
  $command->executeHPHPiStart($server['hostname']);
  id(new SandboxDBStore())
    ->updateServerType($server_id, $hphpi = true);
}

/**
 * Marks a devserver not to run HPHPi deamon
 *
 * @param  uint   server_id
 * @return bool
 */
public function dontRunHPHPi($server_id) {
  $server  = id(new SandboxDBStore())->getServerByID($server_id);
  $command = new SandboxCommands();
  try {
    $command->executeHPHPiStop($server['hostname']);
  } catch (SystemException $se) {
    // Probably was already stopped, it's all good.
  }
  $command->executeHPHPiUninstall($server['hostname']);
  id(new SandboxDBStore())
    ->updateServerType($server_id, $hphpi = false);
}

/**
                                                215,5          62%
```

David Mortenson

https://www.dmortenson.dev202.prn1.facebook.com/dmortens

David Mortenson

Engineering Manager at Facebook
Studied MS Computer Science at University of Wa...
Lives in San Mateo, California
In a relationship with Sonja Keserovic

What's on your mind?

David is now friends with Andrée Hamel and 4 other people.
Monday

Relja Markovic
October 24

Birthday dinner with David & Sonja.
Thank you guys! I will miss hanging out together, but then Seattle is not far away. It's also cold and rainy, so we will be back often enough. — with Sonja Keserovic and 2 others at Acquapazza.

Sponsored

Cross Channel Advertising

Ad Media

The Most Comprehensive Advertising Solution on the Web. Maximize Your Reach!

8,919 people like AdMedia.

SPIN SLOTS!

Play Wheel of Fortune Slots FREE! #1 on Facebook!

3,900,000 people played Games by GSN

# Why Fast Tools Matter

- Focus and flow are critical to high productivity
- Any development operation taking over 5s will cause engineers to do a lightweight context switch
- Any operation taking over 2 minutes will cause a heavyweight context switch
- These are productivity killers and need to be avoided at all cost!

# Dev Sandbox Page Load Time

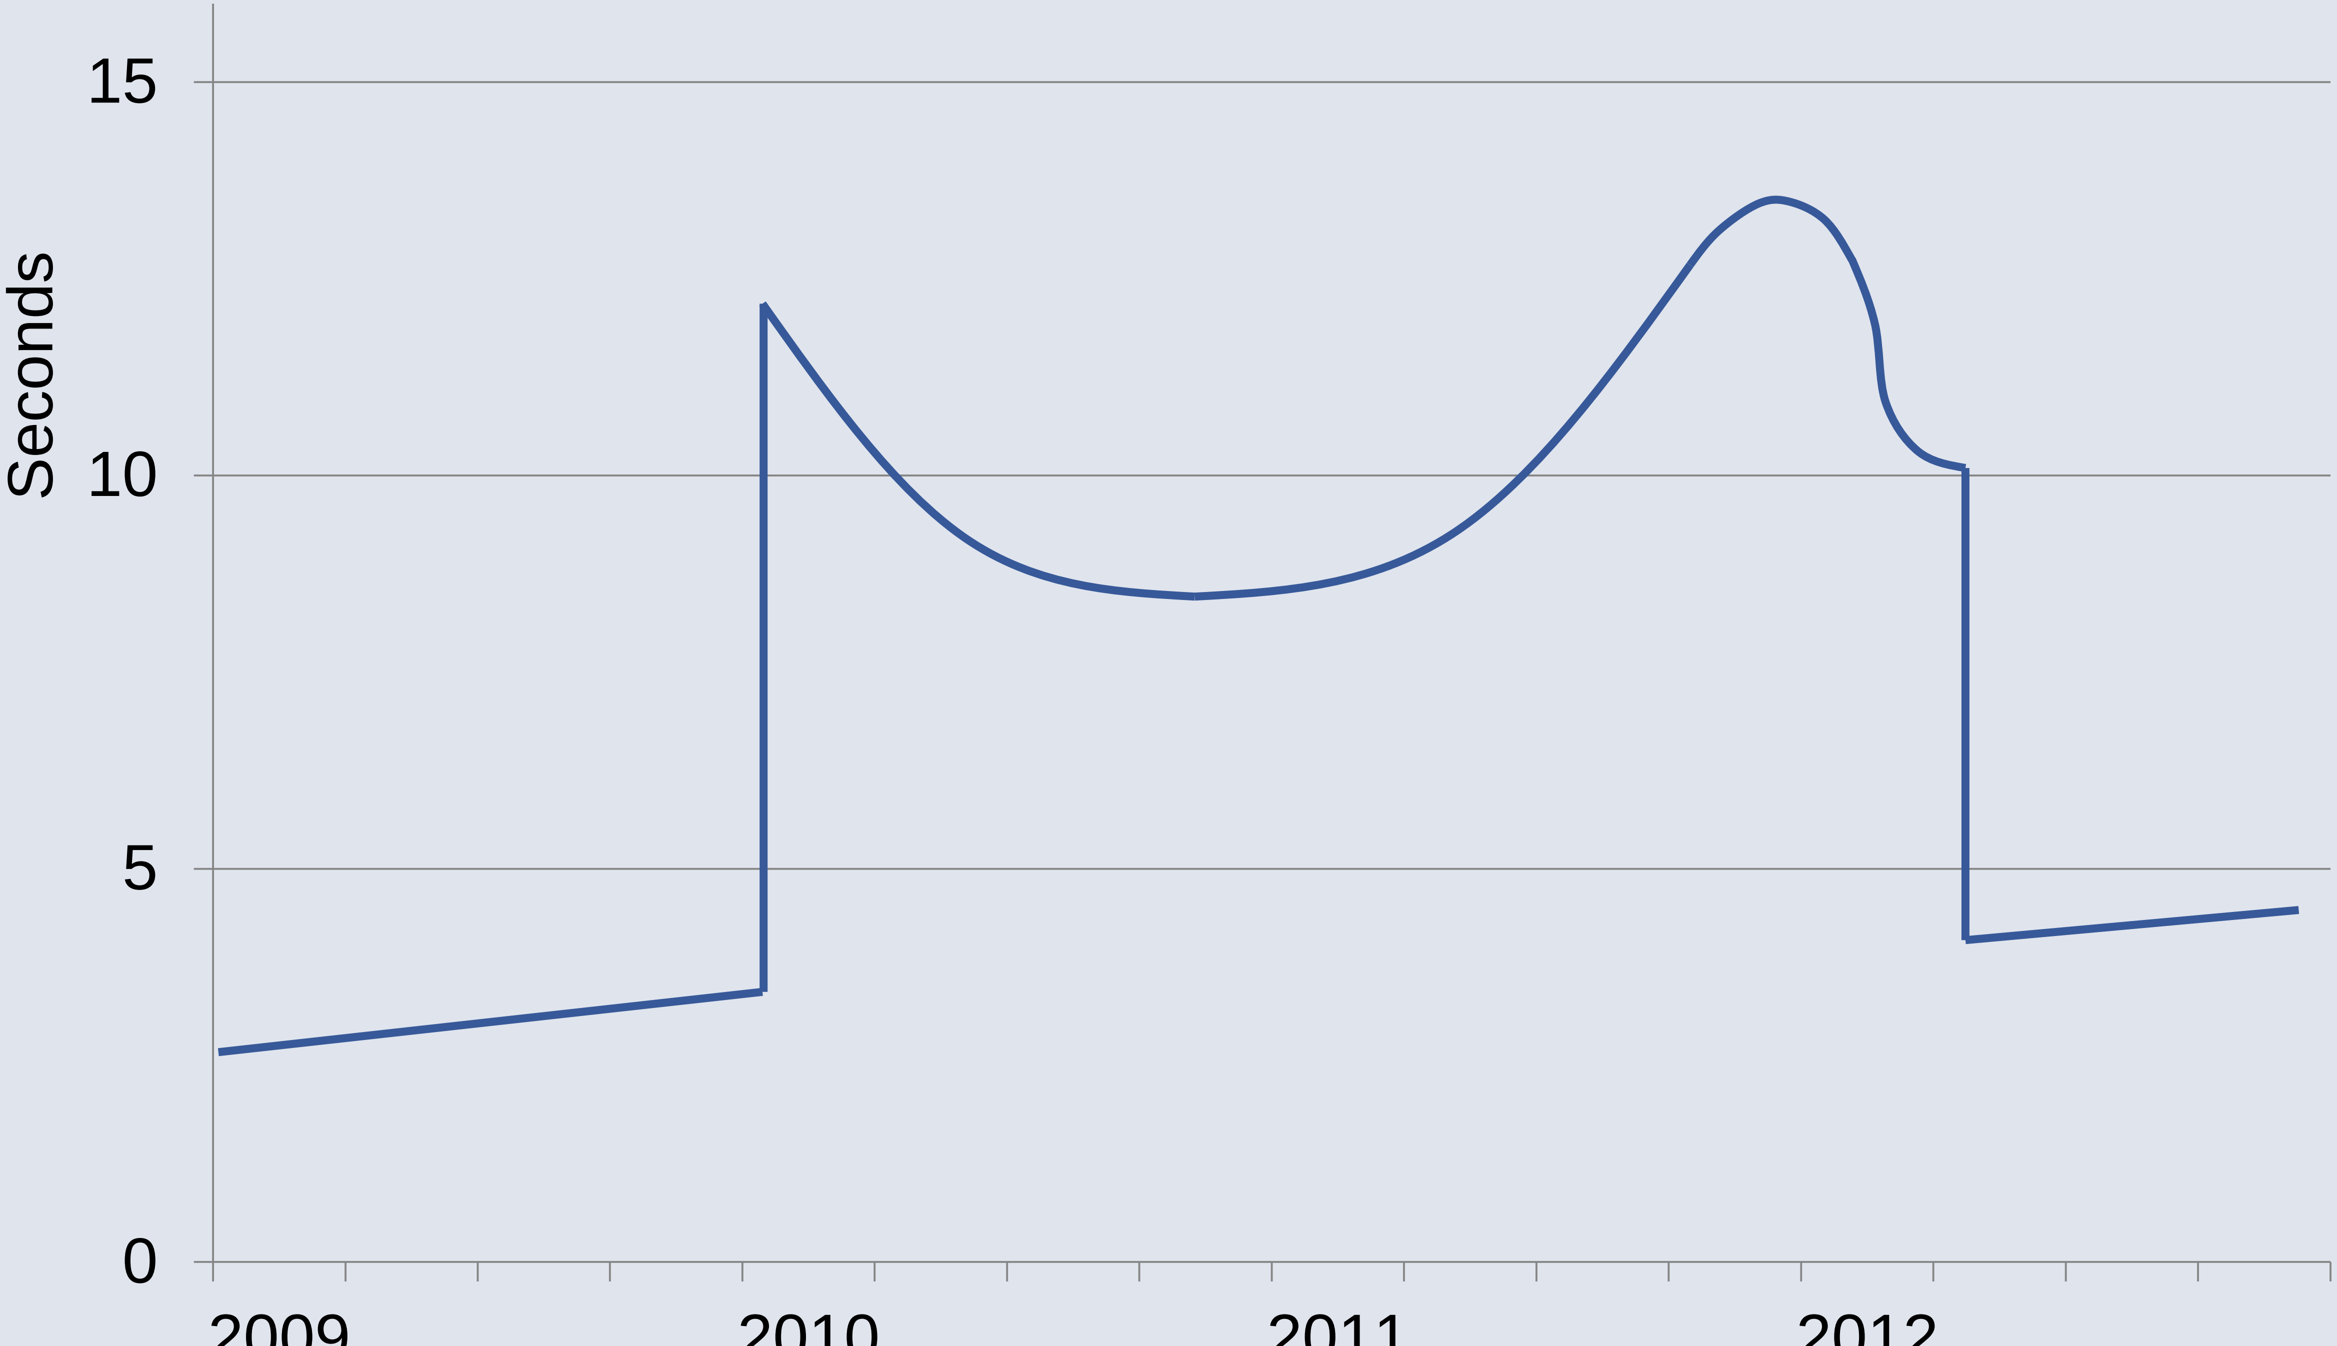

# Other Critical Areas

- Source Control
- Testing
- Static analysis
- Task / bug tracking tool
- Code Review Tools
- Release

# Unintended Consequences

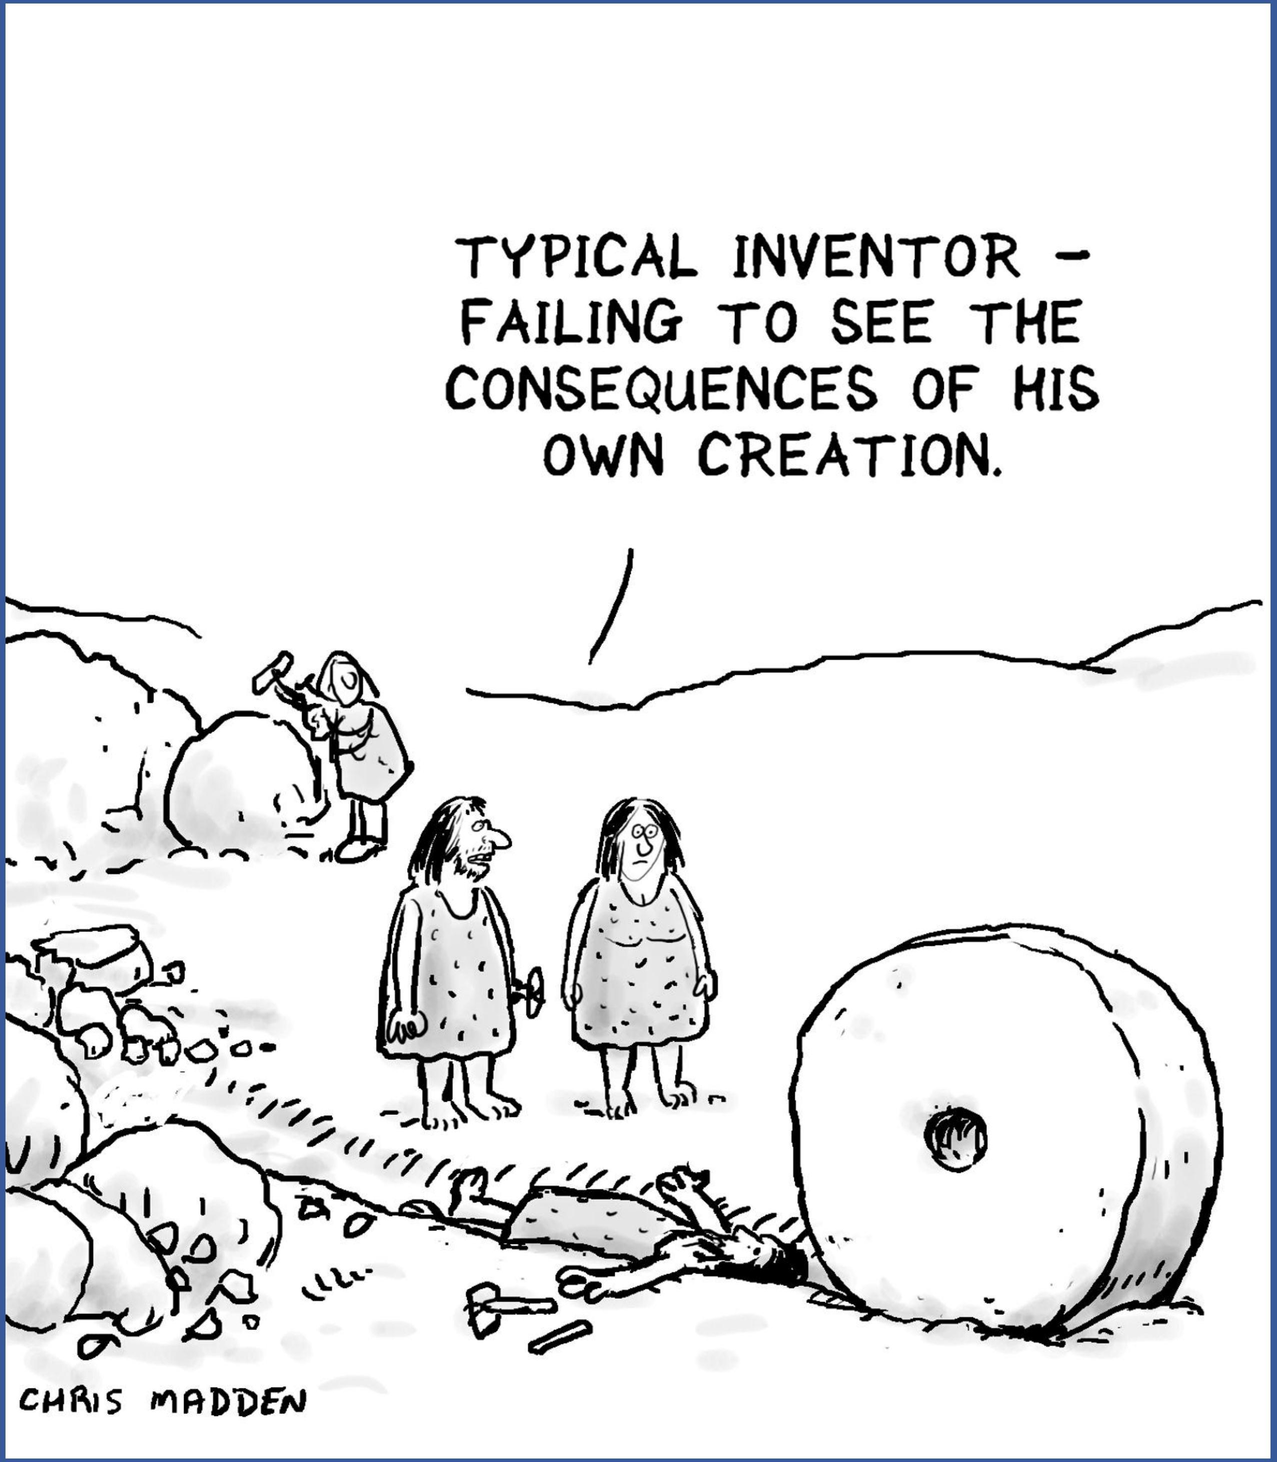
TYPICAL INVENTOR – FAILING TO SEE THE CONSEQUENCES OF HIS OWN CREATION.
CHRIS MADDEN

# Brief History: 2008

- Engineers validate changes by testing them in their sandbox
- Code simple enough & engineering small enough that engineers had a good idea of what a change could break
- Good monitoring in place to catch issues that slipped through

# Brief History: 2009

- Engineering team getting close to 300 people
- Switch to git allowed development of more changes in parallel
- Complexity of the codebase exceeding the point where anyone can keep it all in their head
- Increasing number of breaks that block other engineers
- Increasing number of bugs shipped to production

- Bottom line: the current model wasn’t scaling

# Solution: Automated Testing

# Introducing Testing at Facebook

- Built framework to make it easy to write the tests
- Integrated unit testing in the engineer workflow
- Focused first on writing tests for one key area
- Developed experts / advocates in key teams across FB
- Leveraged incidents to encourage more people to write tests
- Indoctrinated bootcampers

# Testing At Facebook: Initial Results

- ~3000 tests written in 3 years
- Good reduction in breakage in areas with test coverage
- Reduction in number of defects shipped to users

# Testing At Facebook: Curse of Success

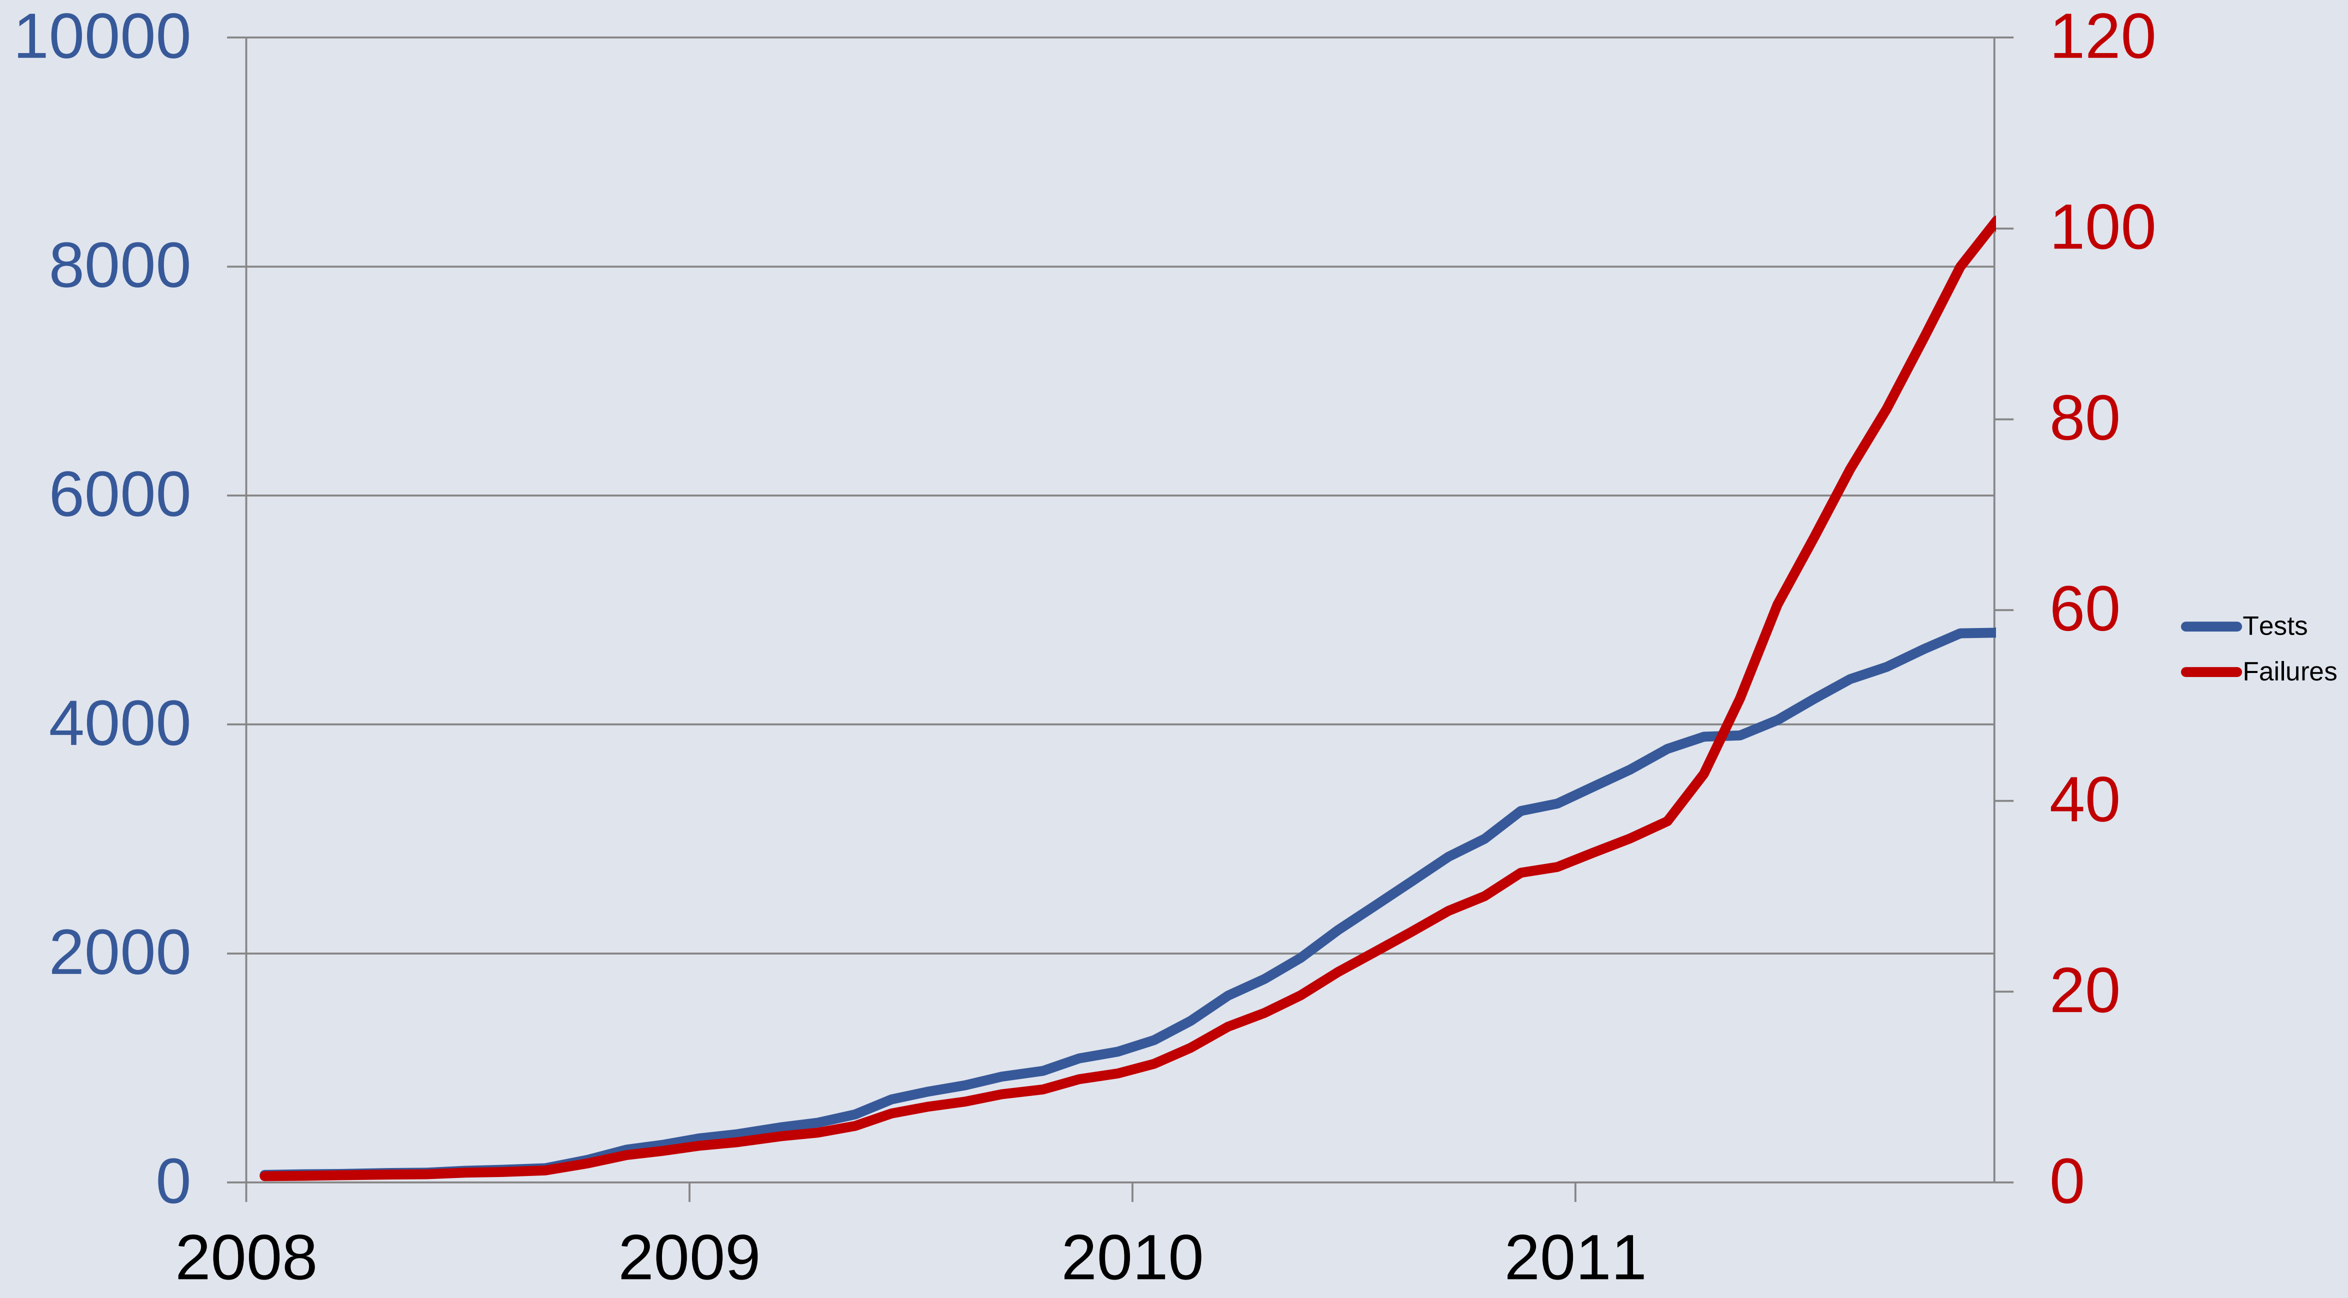

# Cleaning house

- Big effort to get our test failures down to 0
- Killed several hundred low value & commonly failing tests
- Improved the reliability of the testing infrastructure
- Started disabling tests that were failing for > 1 week
- Started disabling intermittently failing tests
- Built a test failure analyzer
- Built a shim for our graph cache to reduce dependencies
- Worked with teams with unreliable tests to improve them

# Game Theory

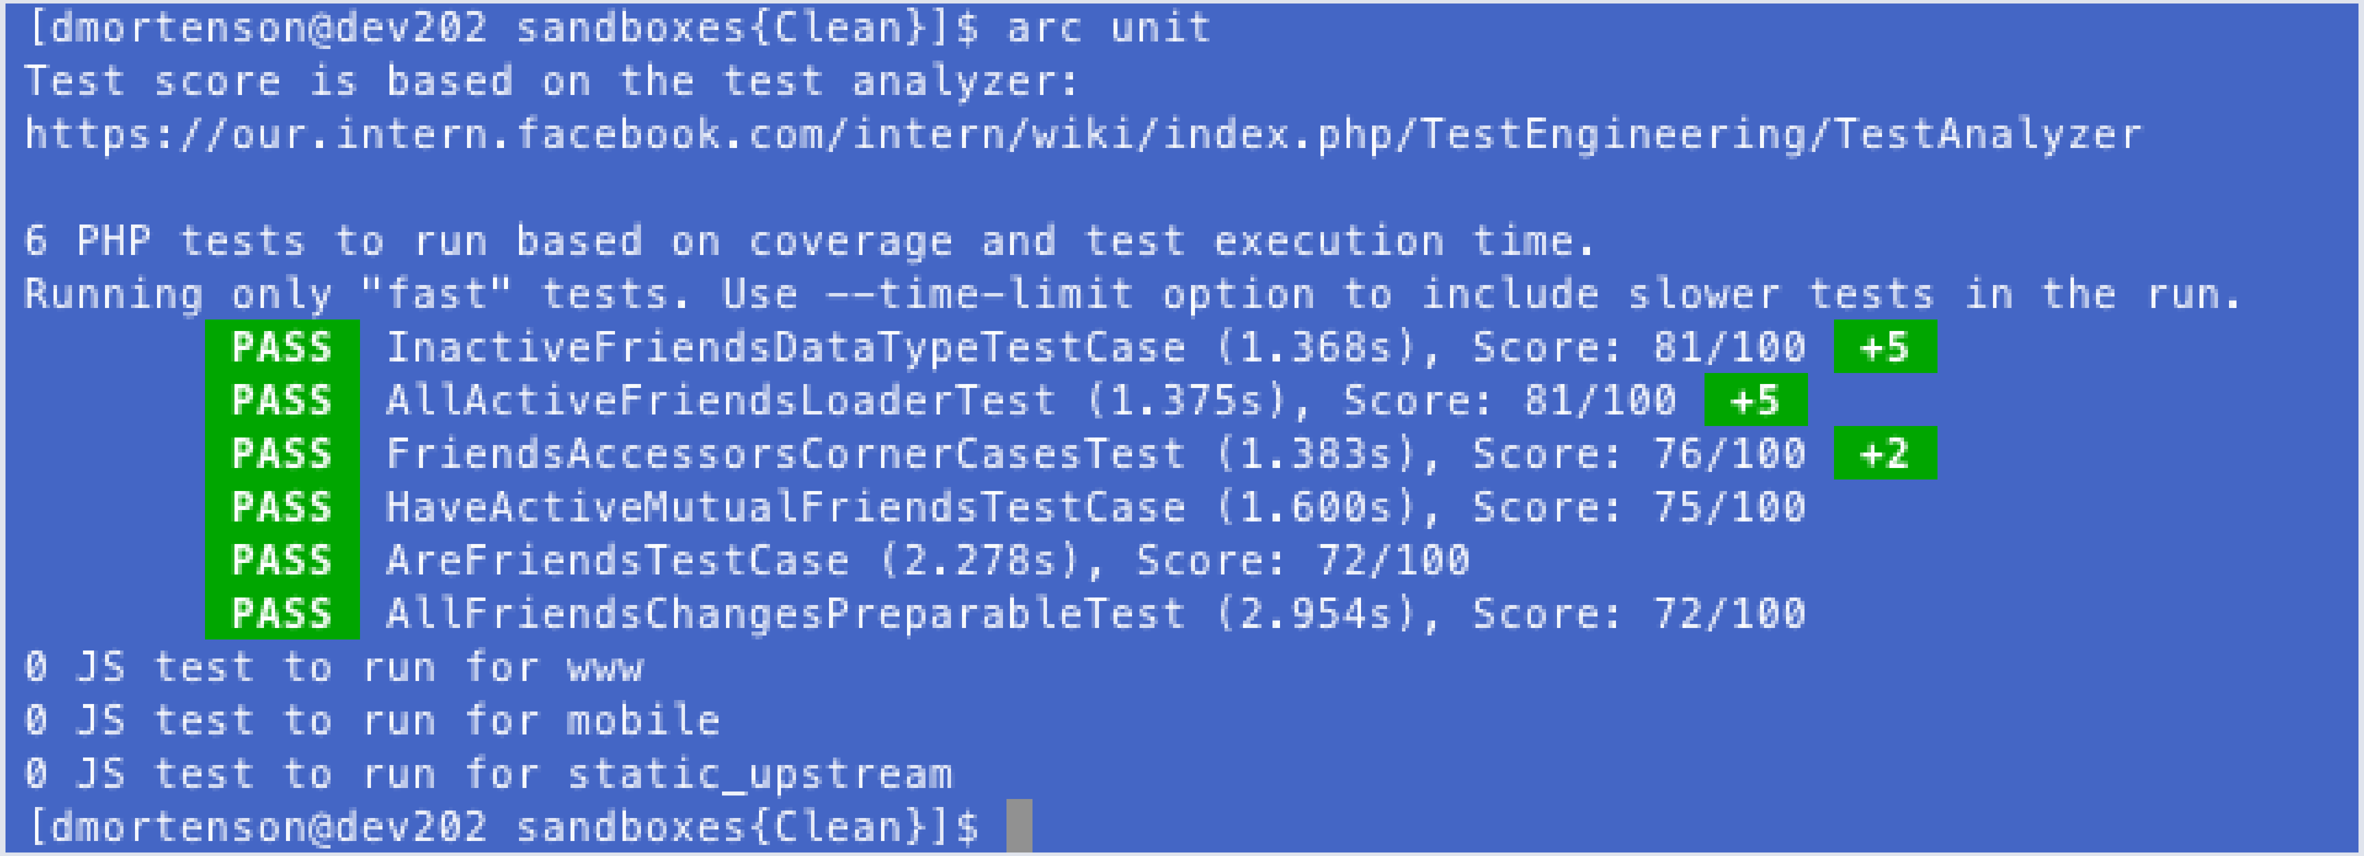

```
[dmortenson@dev202 sandboxes{Clean}]$ arc unit
Test score is based on the test analyzer:
https://our.intern.facebook.com/intern/wiki/index.php/TestEngineering/TestAnalyzer

6 PHP tests to run based on coverage and test execution time.
Running only "fast" tests. Use --time-limit option to include slower tests in the run.
        PASS  InactiveFriendsDataTypeTestCase (1.368s), Score: 81/100  +5
        PASS  AllActiveFriendsLoaderTest (1.375s), Score: 81/100  +5
        PASS  FriendsAccessorsCornerCasesTest (1.383s), Score: 76/100  +2
        PASS  HaveActiveMutualFriendsTestCase (1.600s), Score: 75/100
        PASS  AreFriendsTestCase (2.278s), Score: 72/100
        PASS  AllFriendsChangesPreparableTest (2.954s), Score: 72/100
0 JS test to run for www
0 JS test to run for mobile
0 JS test to run for static_upstream
[dmortenson@dev202 sandboxes{Clean}]$
```

# Results So Far: Better But Not Enough

- Failures still in the 40-50 per run
- Main cause: test failures discovered too late

# Faster Feedback

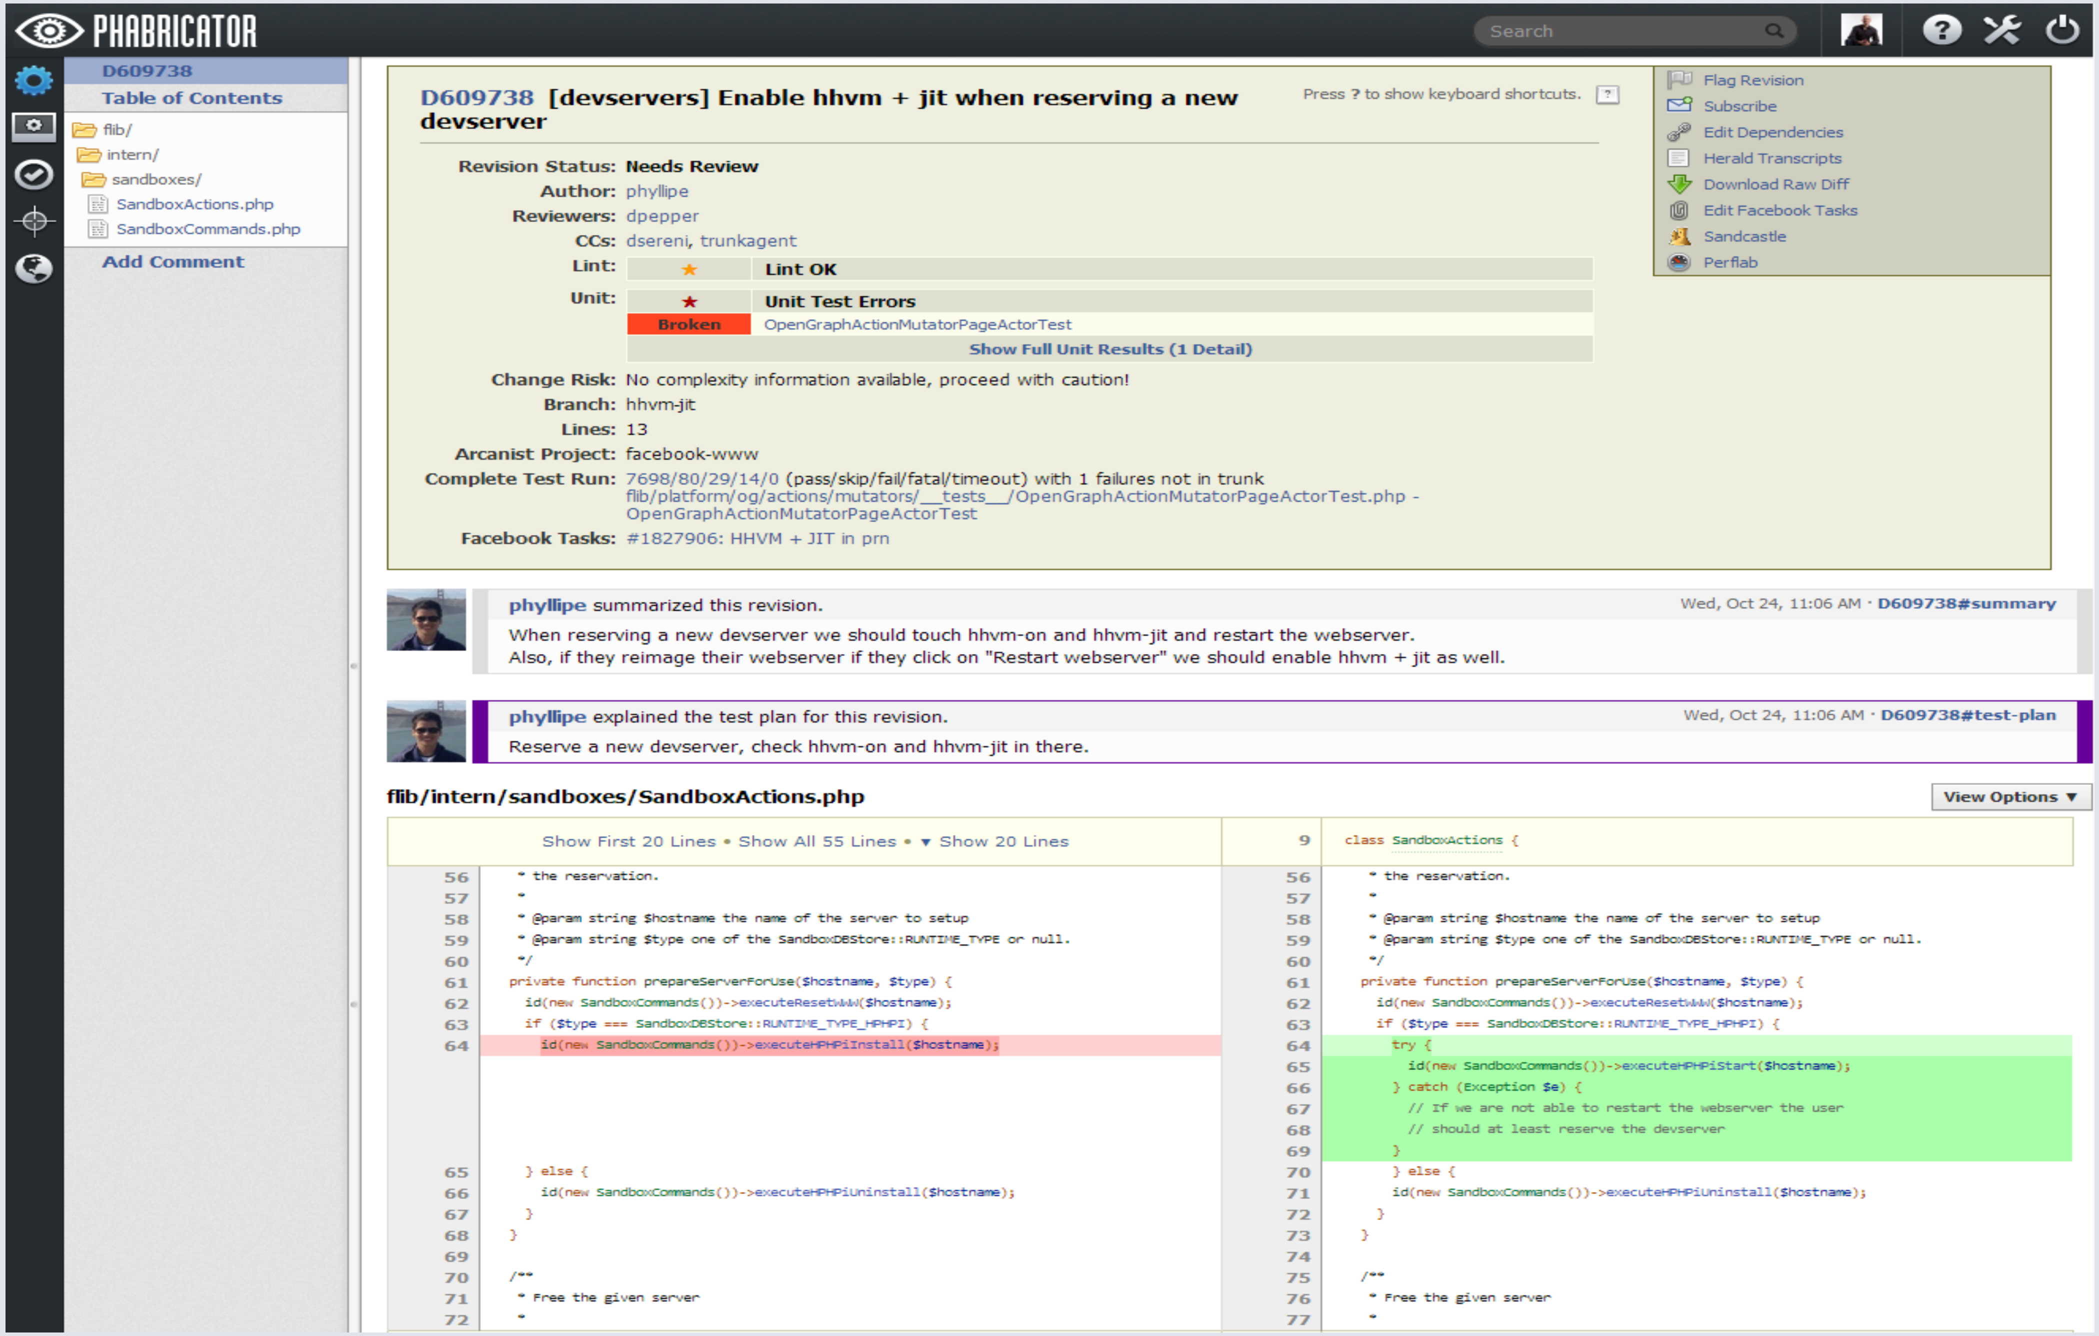

**D609738 [devservers] Enable hhvm + jit when reserving a new devserver**

| | |
|---|---|
| **Revision Status:** | **Needs Review** |
| **Author:** | phyllipe |
| **Reviewers:** | dpepper |
| **CCs:** | dsereni, trunkagent |
| **Lint:** | ★ **Lint OK** |
| **Unit:** | ★ **Unit Test Errors**<br>Broken OpenGraphActionMutatorPageActorTest<br>**Show Full Unit Results (1 Detail)** |
| **Change Risk:** | No complexity information available, proceed with caution! |
| **Branch:** | hhvm-jit |
| **Lines:** | 13 |
| **Arcanist Project:** | facebook-www |
| **Complete Test Run:** | 7698/80/29/14/0 (pass/skip/fail/fatal/timeout) with 1 failures not in trunk flib/platform/og/actions/mutators/__tests__/OpenGraphActionMutatorPageActorTest.php - OpenGraphActionMutatorPageActorTest |
| **Facebook Tasks:** | #1827906: HHVM + JIT in prn |

**phyllipe** summarized this revision. Wed, Oct 24, 11:06 AM · **D609738#summary**

When reserving a new devserver we should touch hhvm-on and hhvm-jit and restart the webserver.
Also, if they reimage their webserver if they click on "Restart webserver" we should enable hhvm + jit as well.

**phyllipe** explained the test plan for this revision. Wed, Oct 24, 11:06 AM · **D609738#test-plan**

Reserve a new devserver, check hhvm-on and hhvm-jit in there.

**flib/intern/sandboxes/SandboxActions.php**

Show First 20 Lines • Show All 55 Lines • ▼ Show 20 Lines

```
class SandboxActions {
```

```
   * the reservation.
   *
   * @param string $hostname the name of the server to setup
   * @param string $type one of the SandboxDBStore::RUNTIME_TYPE or null.
   */
  private function prepareServerForUse($hostname, $type) {
    id(new SandboxCommands())->executeResetWWW($hostname);
    if ($type === SandboxDBStore::RUNTIME_TYPE_HPHPI) {
      id(new SandboxCommands())->executeHPHPiInstall($hostname);
    } else {
      id(new SandboxCommands())->executeHPHPiUninstall($hostname);
    }
  }

  /**
   * Free the given server
   *
```

```
   * the reservation.
   *
   * @param string $hostname the name of the server to setup
   * @param string $type one of the SandboxDBStore::RUNTIME_TYPE or null.
   */
  private function prepareServerForUse($hostname, $type) {
    id(new SandboxCommands())->executeResetWWW($hostname);
    if ($type === SandboxDBStore::RUNTIME_TYPE_HPHPI) {
      try {
        id(new SandboxCommands())->executeHPHPiStart($hostname);
      } catch (Exception $e) {
        // If we are not able to restart the webserver the user
        // should at least reserve the devserver
      }
    } else {
    id(new SandboxCommands())->executeHPHPiUninstall($hostname);
    }
  }

  /**
   * Free the given server
   *
```

# Testing At Facebook: Now

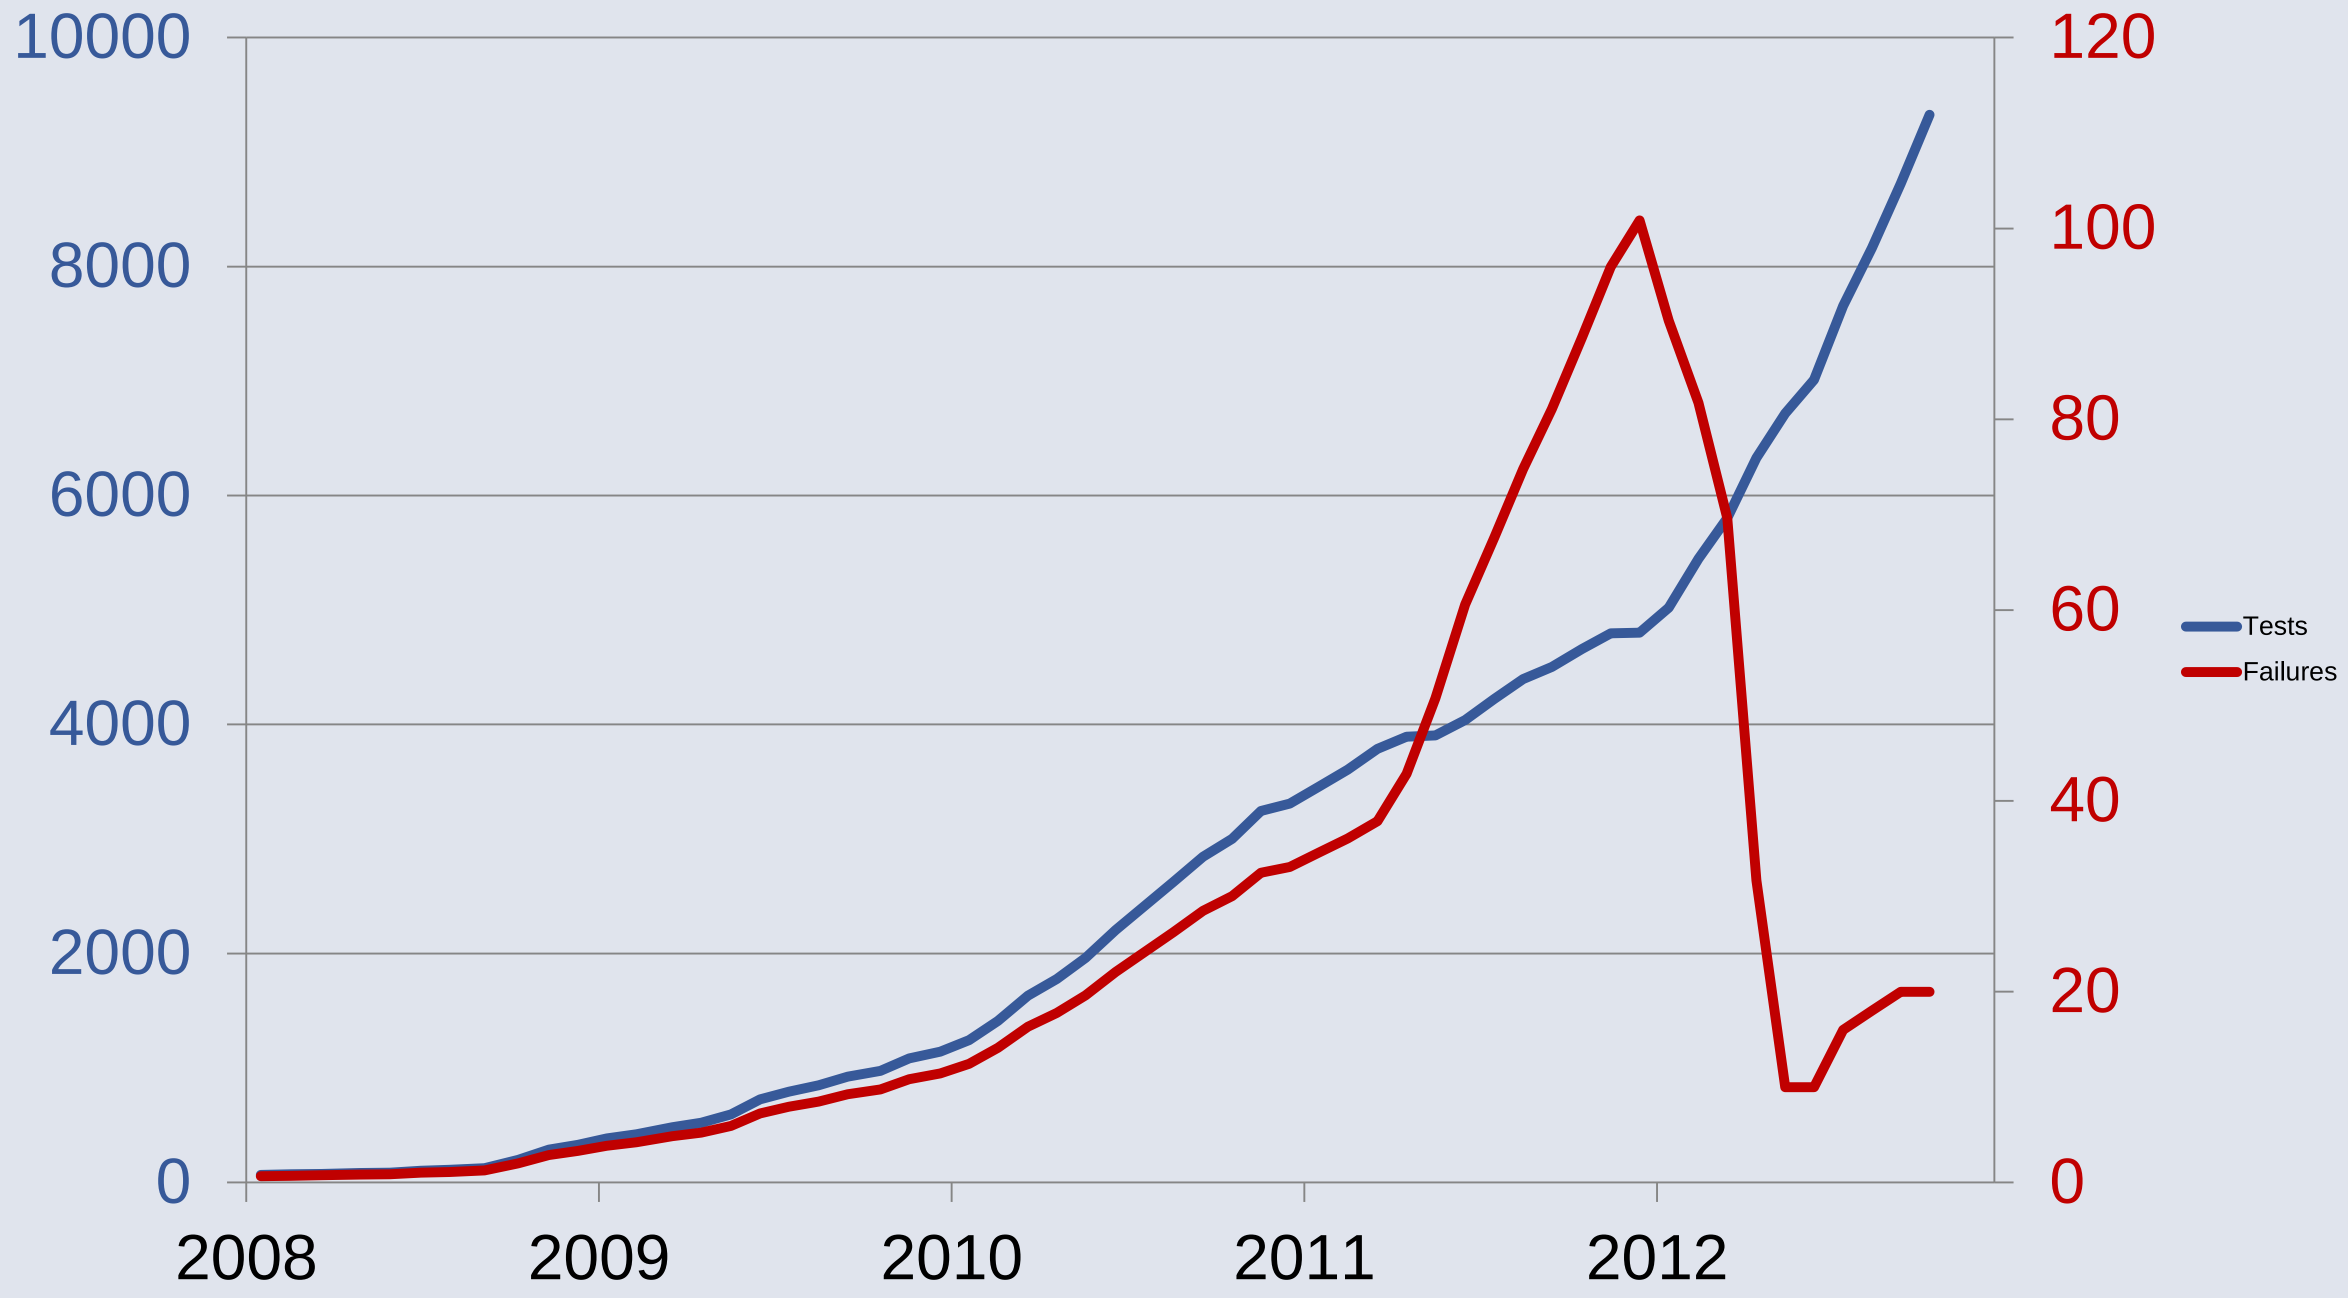

## So: did it work?

# Commits every Month

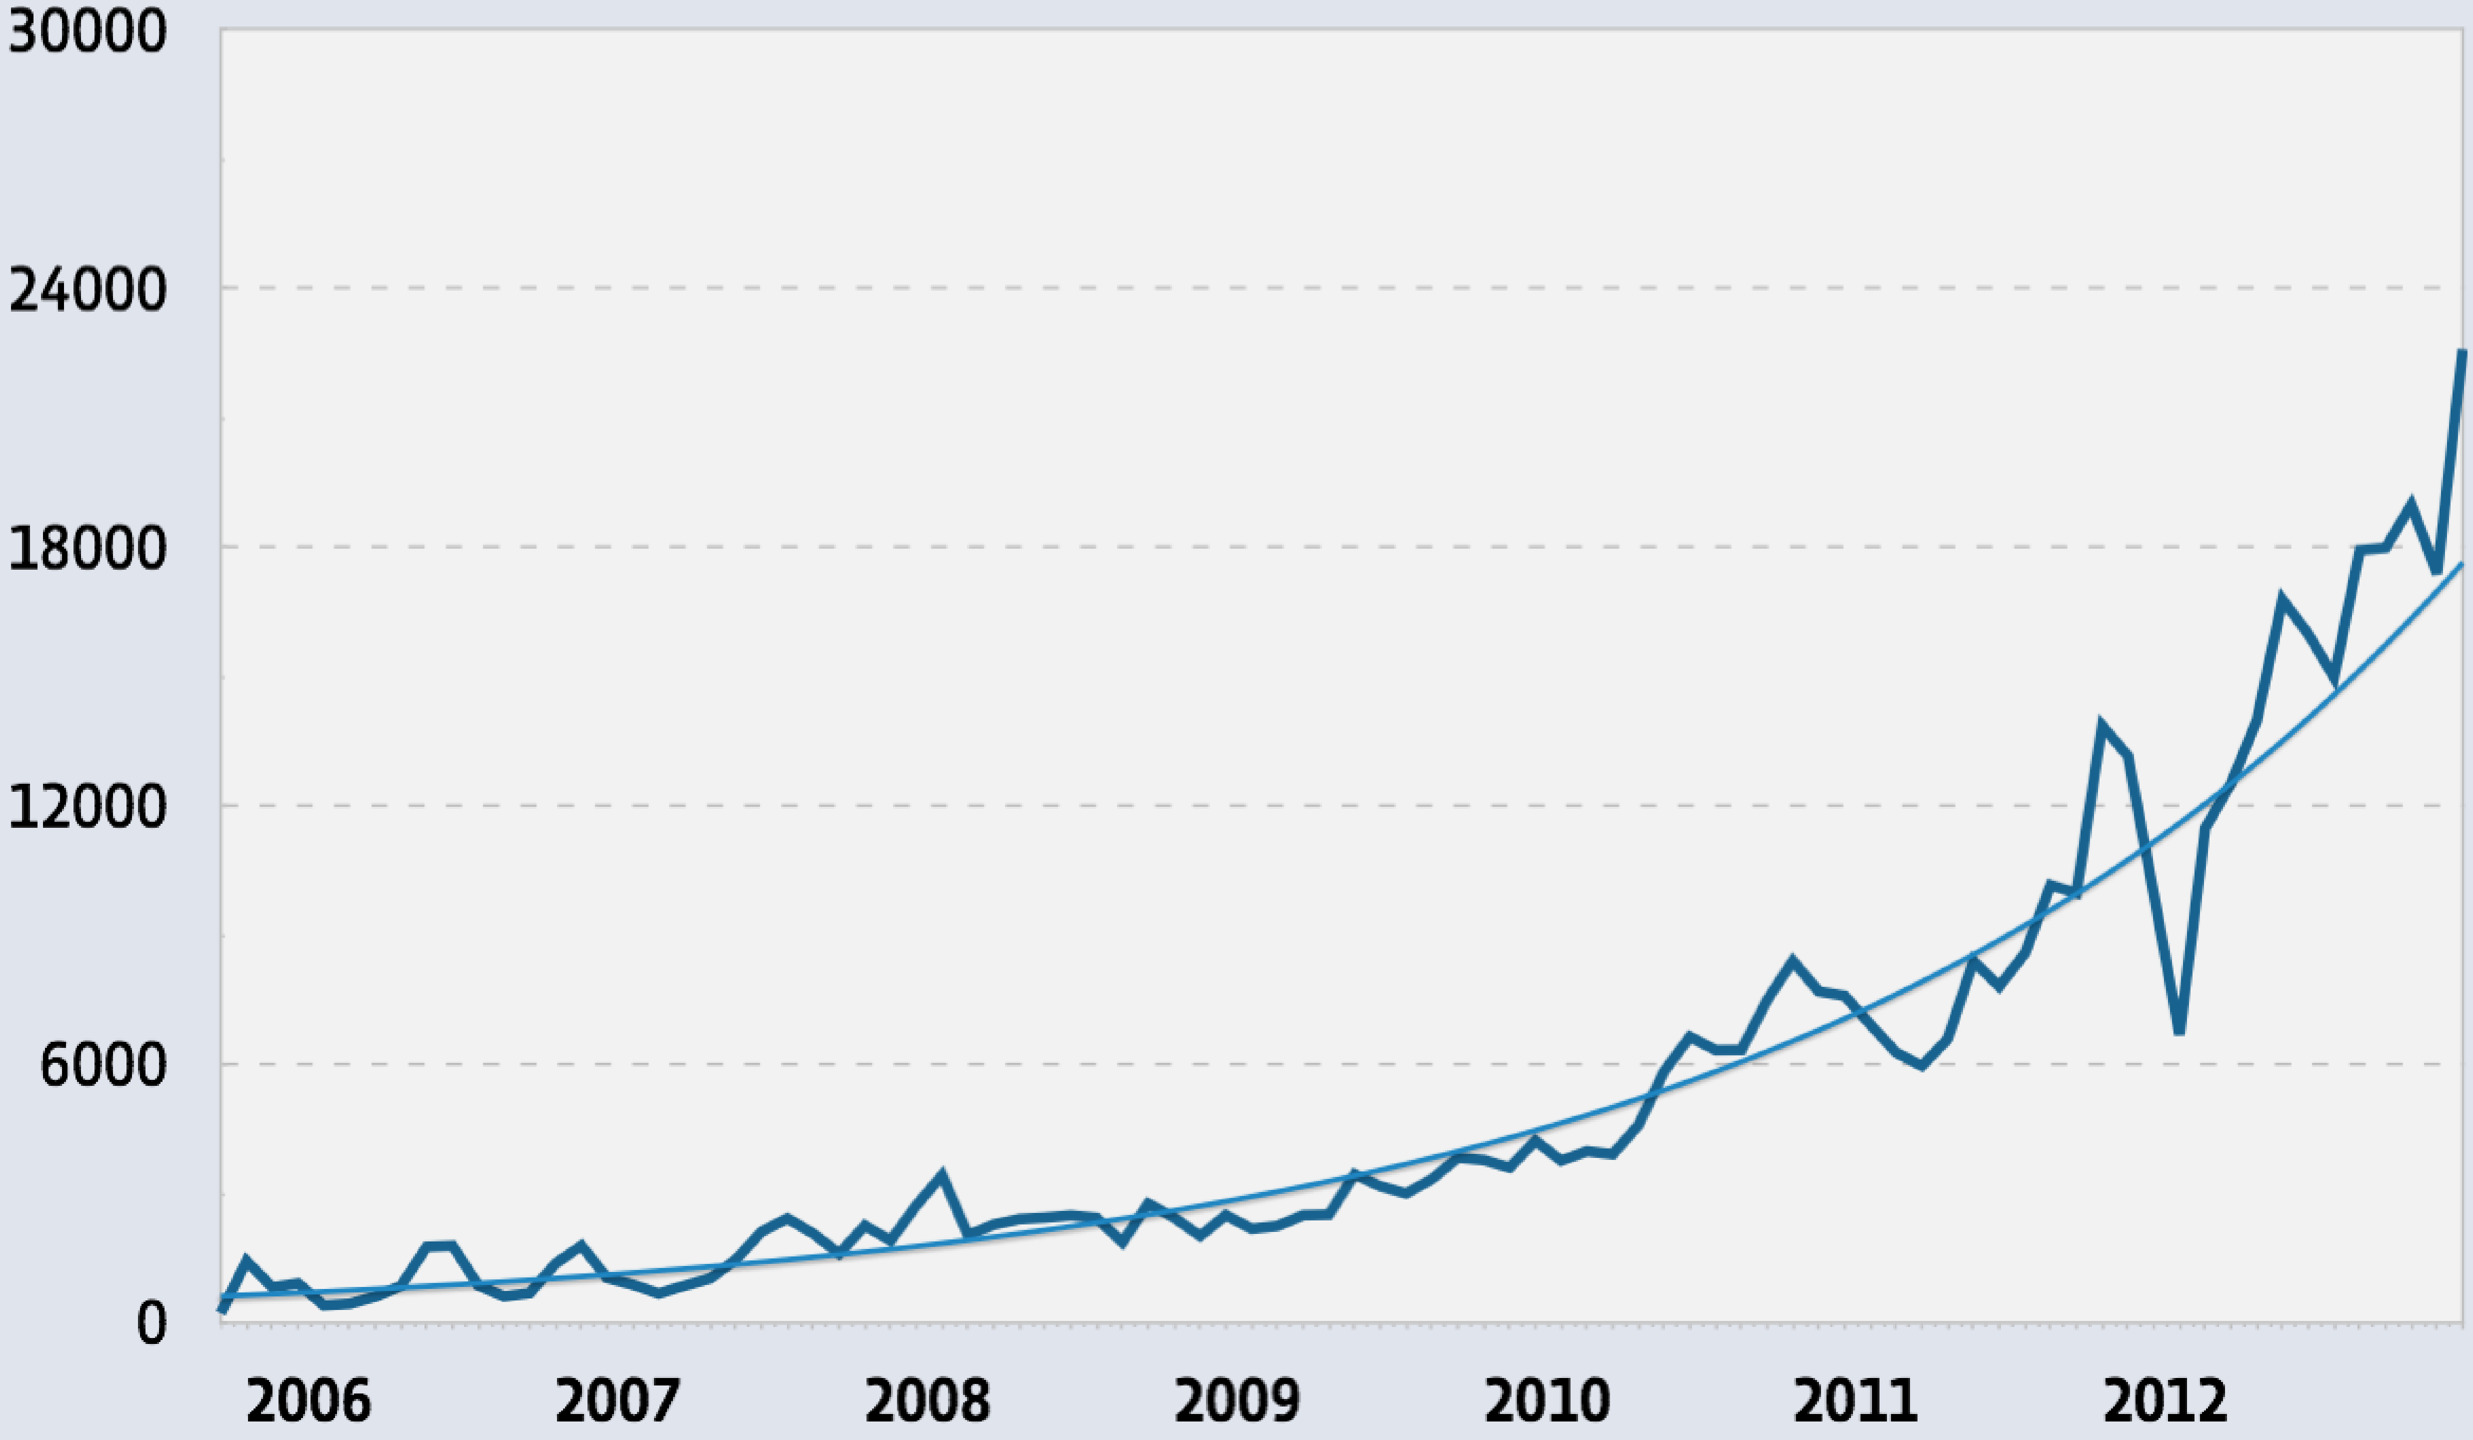

# Lots of Work Left

- Mobile
- Back end services
- Making every engineer more productive as we grow

How are you tackling these challenges?

facebook