

MediaWiki Core Review

October 2014

What we said in July

HHVM/Performance



Ori



Tim



Aaron



Rob

SUL



Chris



Dan



Bryan

Data & Developer Hub



Brad



Sumana

Search

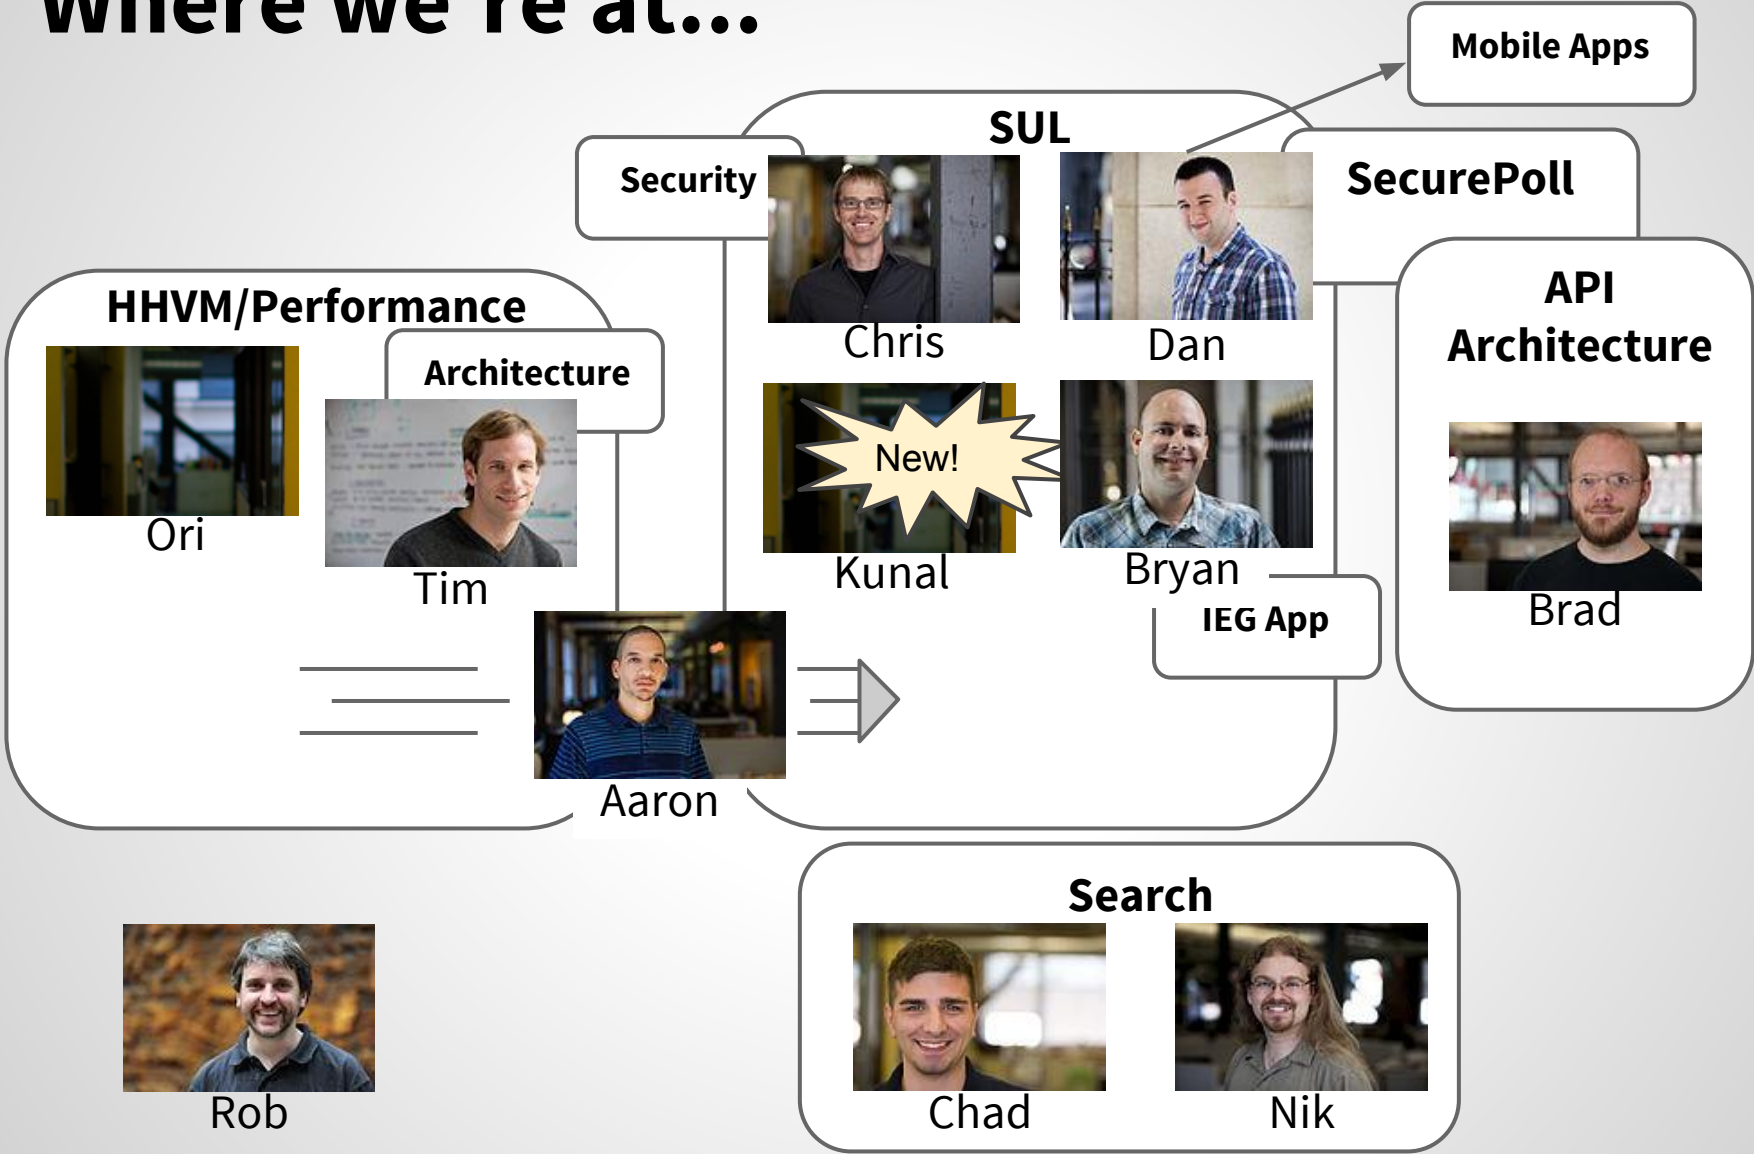


Chad



Nik

Where we're at...



Where we're going...



Previous quarter bigger projects

- HHVM
- Single user login finalisation
- Search

HHVM

- Team:
 - Core: Ori, Tim, Aaron
 - Ops: Giuseppe, Filippo
 - ..plus Max, Brett, Chad, Bryan, Erik B
 - Product Owner: Ori Livneh

- High-level goals:
 - Cut the time it takes to save an edit in **half**
 - **Eliminate** performance penalty for logged-in users
 - Modernize our infrastructure in the process

HHVM - Plan

- Upgraded “HAT stack” - HHVM, Apache 2.4, Trusty (Ubuntu)
- Steps
 - Each step involves deploying to beta cluster prior to production
 - Limited deployment to a few job runners in production
 - Limited deployment to small number of application servers
 - Ramp up deployment to more application servers until most servers use HHVM
 - Deploy to remainder of services

Single User Login Finalisation (SUL)

- Team:
 - Engineers: Bryan Davis, Kunal Mehta, Chris Steipp, Aaron Schulz
 - Community: Keegan Peterzell, Rachel diCerbo
 - Design: Brandon Harris
 - Product Owner: Dan Garry

SUL finalisation work required

- Big rename script (i.e. the finalisation proper)
- Merge “easy” accounts (DONE)
- Request a rename
- Log in with old credentials
- Global account merge
- Global rename user (DONE)

CirrusSearch

- Team
 - MWCore: Nik Everett, ½ of Chad Horohoe
 - Ops: Andrew Otto, Filippo Giunchedi
- Progress
 - Serve all requests for all but four wikis
 - Sadly, those four wikis are ~75% our search traffic
 - Other wikis are available as well received BetaFeature
- Hardware
 - 2x faster disks and 2x more nodes coming
- Once we migrate all the wikis we'll work more to improve results

Bigger projects proposed for next quarter

- Library infrastructure for MediaWiki

Show a path to decomposition of MediaWiki into smaller pieces.

- Editing performance

What probably doesn't get done as a result (among other things)

- Generalized A/B testing infrastructure
- Generalized performance measurement infrastructure and KPIs
- Code review/management for site JS
- Data center failover work on MediaWiki

Future guidance (Q3 and beyond)

- What ideas should we flesh out for future consideration?
 - Currently in the 2014-15 plan of record:
 - Help with Wikidata's complex query language
 - Tool for category intersection and/or Wikidata tag intersection
 - Template styling
 - Work with Operations on SPDY (see bug 33890)
 - Service-oriented architecture authentication