Scholar@UC - getting there

Linda Newman, University of Cincinnati
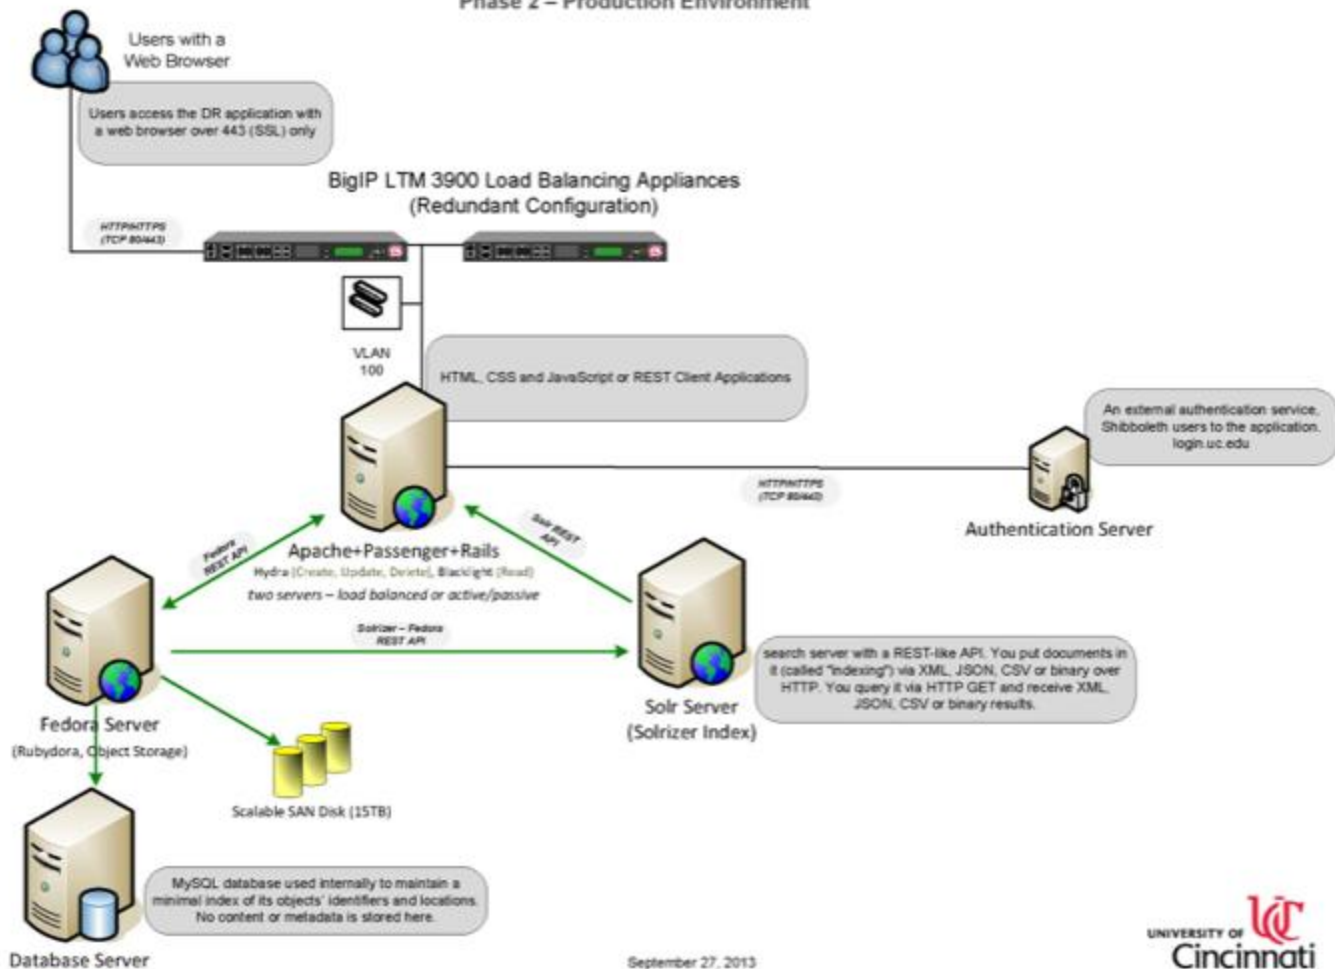linda.newman@uc.edu

Preliminaries...

- Fall 2013 -- joined Hydra as a partner, attended Hydra camp, established sandbox and installed sufia, avalon, curate

- January 2014 - joined the shared Curate development project with Northwestern, Notre Dame, Indiana University, University of Virginia. Learned agile, gitflow, coding and testing local app on laptops, sandbox before committing to project github. Kept our sandbox app up to date with "vanilla curate"-- the sandbox app of the shared project.
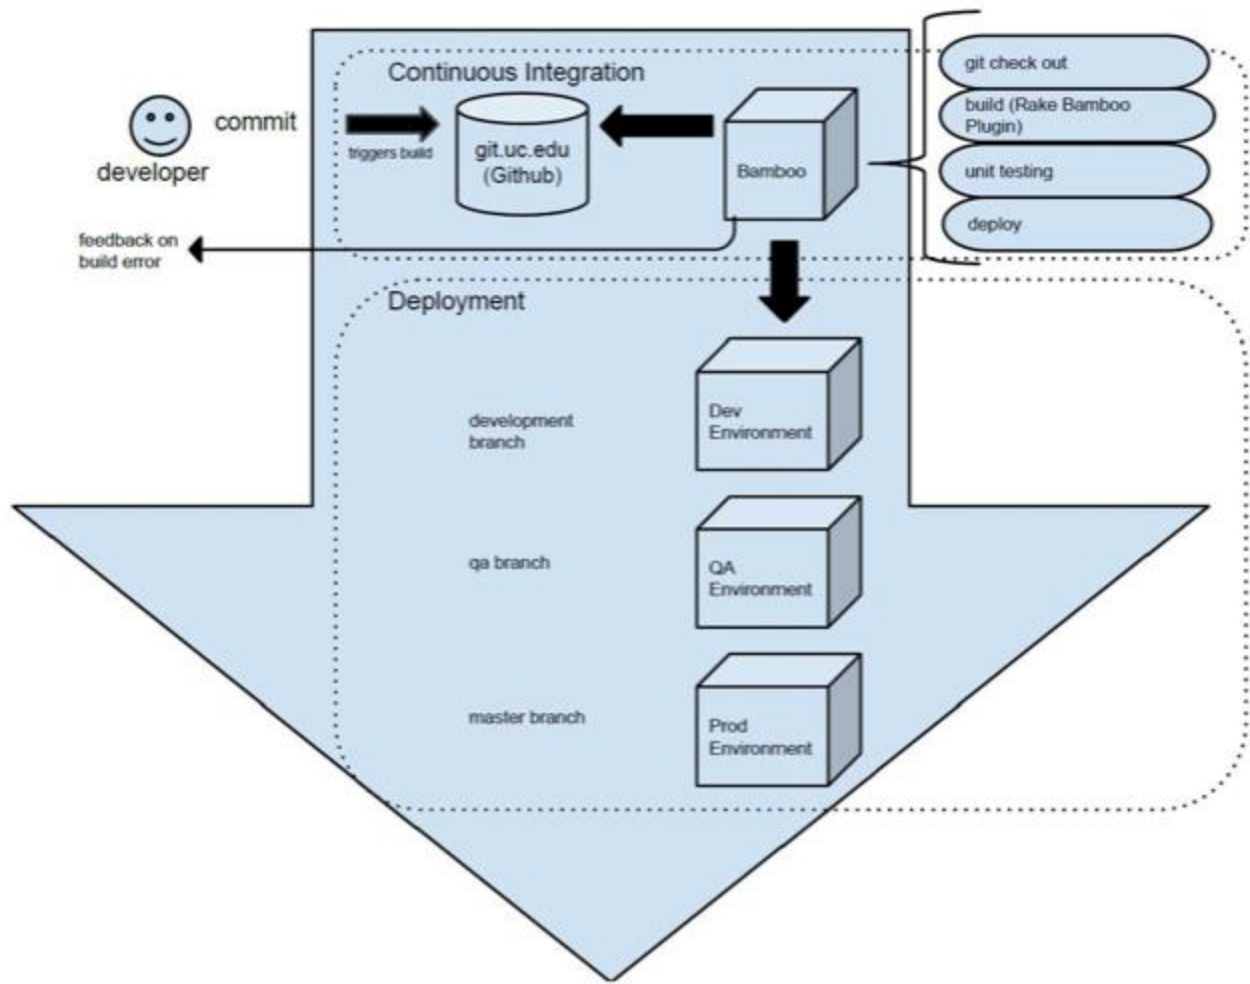
Assembly...

- March 2014 -- work begins -- a software developer and devops person working together -- to establish dev, qa, and production server environment, and create a deploy strategy using bamboo and puppet.

UC Digital Repository

Phase 2 – Production Environment

Users with a
Web Browser

Users access the DR application with
a web browser over 443 (SSL) only

BigIP LTM 3900 Load Balancing Appliances
(Redundant Configuration)

HTTP/HTTPS
(TCP 80/443)

VLAN
100

HTML, CSS and JavaScript or REST Client Applications

An external authentication service.
Shibboleth users to the application.
login.uc.edu

HTTP/HTTPS
(TCP 80/443)

Authentication Server

Fedora
REST API

Solr REST
API

Apache+Passenger+Rails
Hydra (Create, Update, Delete), Blacklight (Read)
two servers – load balanced or active/passive

Solrizer – Fedora
REST API

Fedora Server
(Rubydora, Object Storage)

Solr Server
(Solrizer Index)

search server with a REST-like API. You put documents in
it (called "indexing") via XML, JSON, CSV or binary over
HTTP. You query it via HTTP GET and receive XML,
JSON, CSV or binary results.

Scalable SAN Disk (15TB)

MySQL database used internally to maintain a
minimal index of its objects' identifiers and locations.
No content or metadata is stored here.

Database Server

September 27, 2013

Shared environment re-calibrates...

- May 2014 -- shared development of Curate is suspended and Notre Dame begins development of fundamental redesign (Hydramata)

- May 2014 -- Notre Dame has deployed Curate to Early Adopters with a campus wide rollout planned for Fall. We decide to push for a deploy to Early Adopters as well, dividing work on remaining bugs between Cincinnati and Notre Dame.

Getting close...

- August -- name and URL chosen:

     scholar@uc   --   scholar.uc.edu

- September 5 - Dev, QA, and production environments operational, puppet scripts on bamboo server working, re-deploy tested. [Our puppet scripts are available at https://github.com/uclibs/scholar-puppet.]

- September 10-12 -- Fedora 3 XACML policies in place, our application (based on a fork of Curate, with some modifications) is deployed.

- September 19 -- UC IT Office of Information Security begins 'Hailstorm' security scan.

Deploying our app...

- September 25 - Branch for Metadata changes and support for DOI merged with our fork, testing checks out, 'must fixes' done -- we are ready to deploy.

- September 26 - We deploy to sandbox and dev

- September 29 - We deploy to QA

- Goal of opening up to Early Adopter content on Oct. 6th is in sight.

Goals in sight but then...

But -- backwards motion

- September 29 - OIS says we failed Hailstorm scan and are vulnerable to cross-site scripting and cannot run a production system.

Forward motion again

- September 30 -- OIS agrees with us that we are not vulnerable and reverses decision.

- September 30 - We deploy to production.

- https://github.com/uclibs/curate_uc

https://scholar.uc.edu