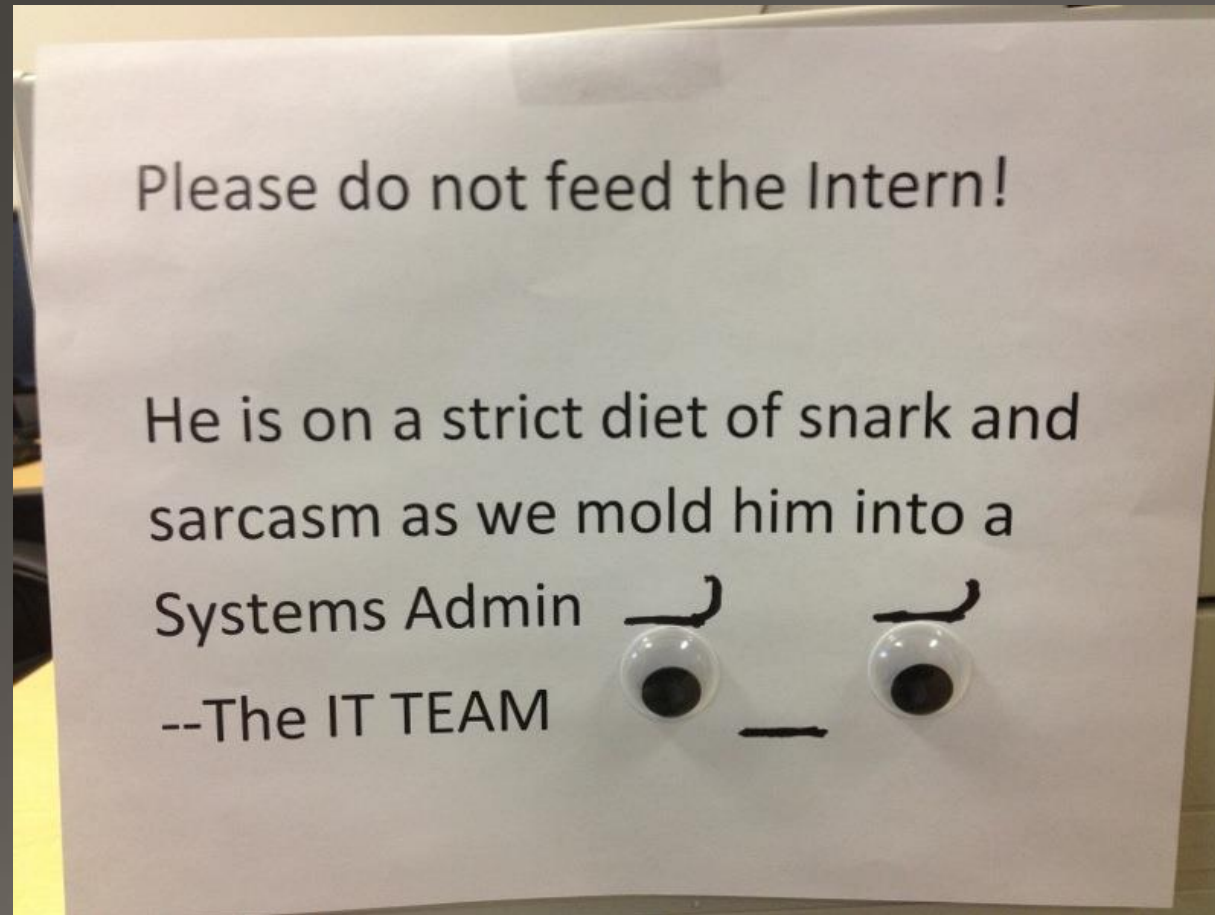


What can Infrastructure do for you today?

Daniel "Humbedooh" Gruno

Infrastructure Architect, The Apache Software Foundation



What is infrastructure?

What is infrastructure?

- The Apache Infrastructure Committee (henceforth 'Infrastructure') is the steward of code and development provenance.
- Infrastructure manages all the machines and services that tie together the ASF.
- Infrastructure grants and oversees the resources project teams need to be able to collaborate on the software that makes the ASF interesting.
- Infrastructure facilitates the common community resources that allows people to communicate and make the ASF awesome.

Who is infrastructure?

Infrastructure is a vast group of people and hardware

- 90 members of the Infrastructure LDAP group
- 42 people in the supporting infrastructure-interest LDAP group
- 16 people in the presidents committee (root@)
- 6 paid staffers (4 full-time, 2 part-time)

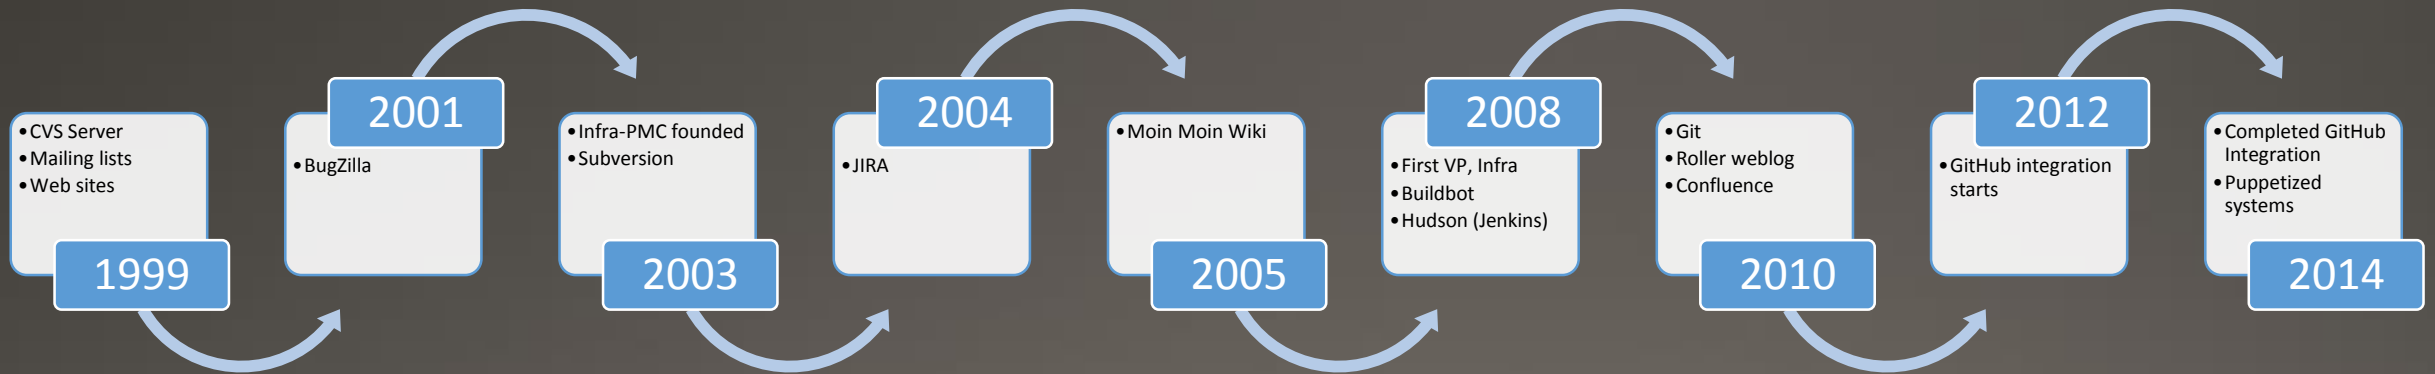
- 35 bare-metal machines
- 100+ VMs and jails

A historical look at infrastructure

- The Infrastructure team was informally founded in 1999
- In 2002, a resolution was made to form the Infrastructure Committee as a board committee. IT WAS DECLINED!!!¹¹one
- Somehow, (without anyone apparently knowing when), the Infrastructure Committee was formed as a President's Committee somewhere between October, 2002 and February, 2003.
- First actual Vice President was in 2008 (Paul Querna)
- Originally tasked with handling email, website and subversion repo
- Started out as an all-volunteer group of committers

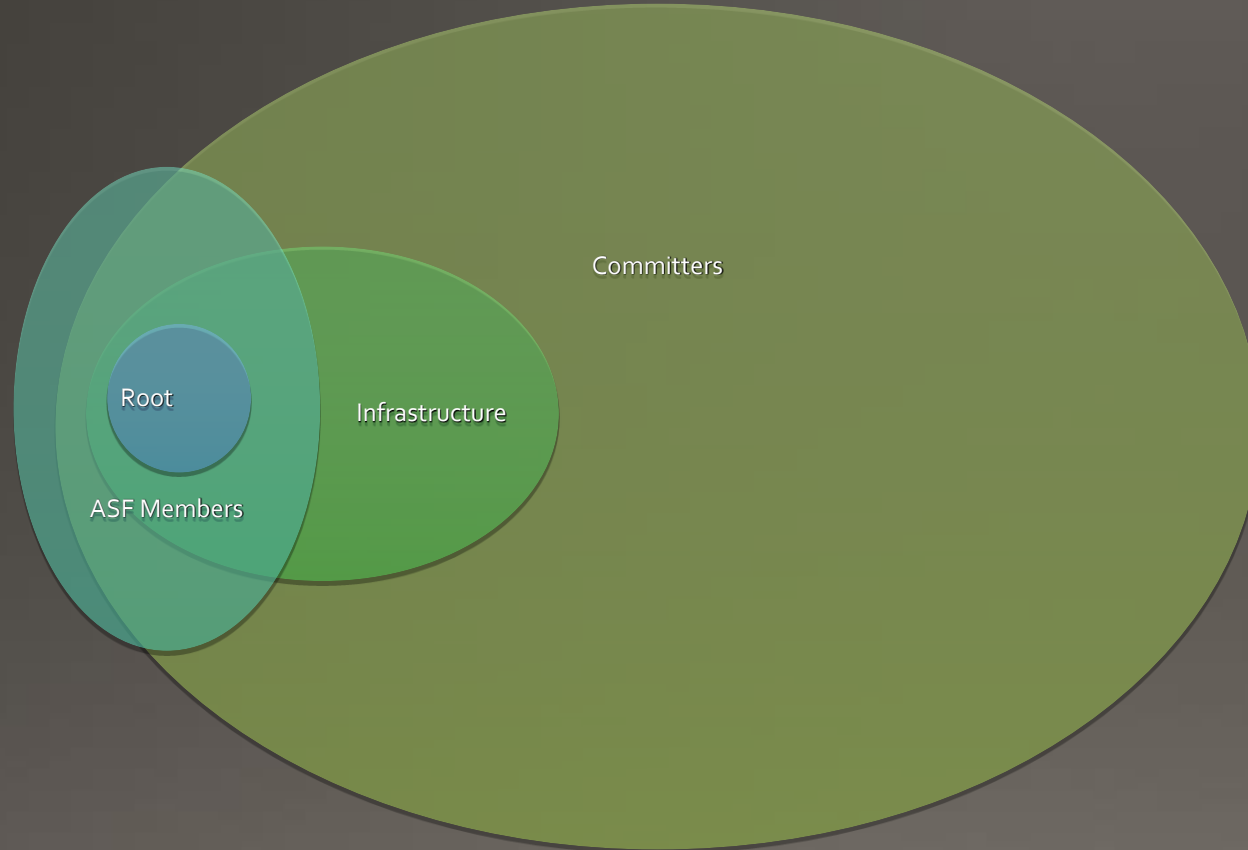
Infrastructure service timeline

A list of some of the main services as they appeared in the ASF:



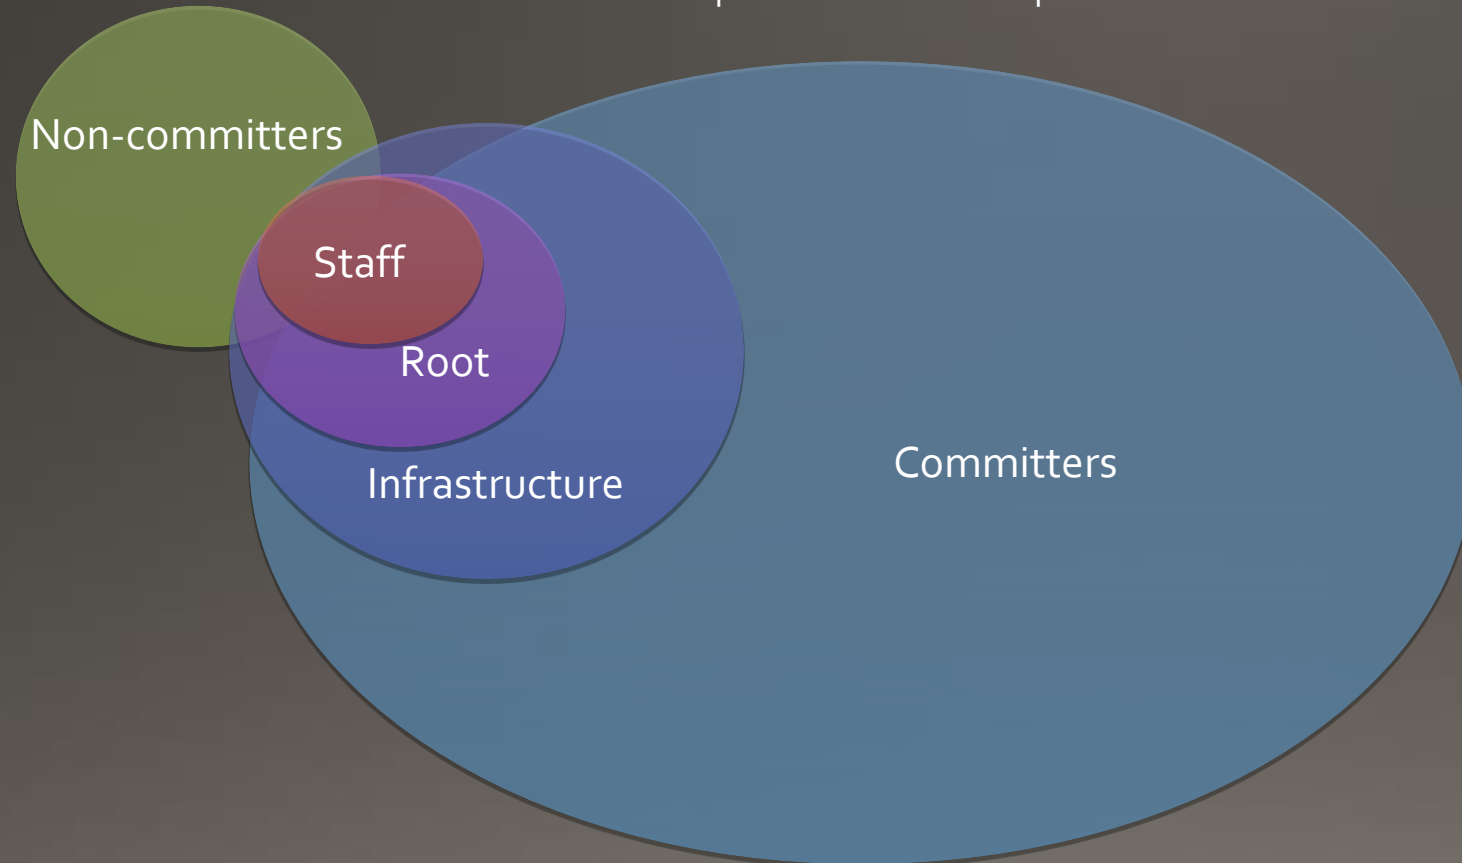
How was infrastructure comprised back in the day?

- All volunteer based
- Infrastructure members were picked from committers
- Root was by merit and ASF Members only



How is Infrastructure comprised today?

- All-volunteer model did not scale
- Staff was hired to deal with the growth of the ASF
- Root picked from infrastructure members or hired staff
- Root does not require ASF membership



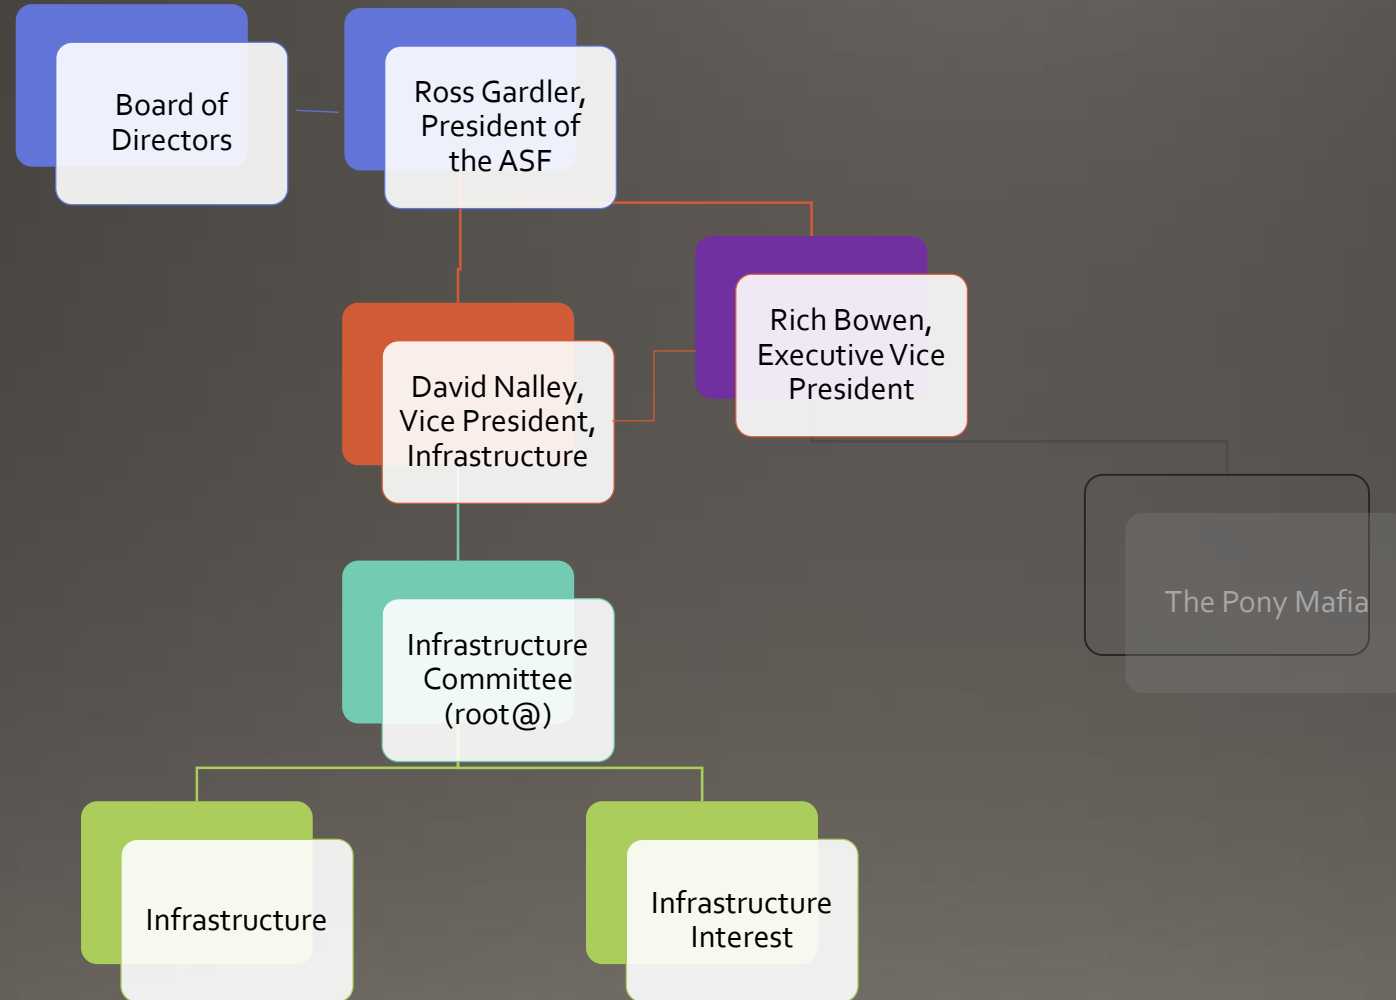
Chain of Command

Infrastructure is a President's Committee. It is comprised of 16 members, including a Vice President of Infrastructure (Currently David Nalley since April 2014).

Infrastructure reports to the president of the ASF (or the EVP in his/her absence).

Unlike Top Level Projects, that reports to the board every quarter, Infrastructure is required to report to the president every month, who in turn reports to the board at the monthly board meeting.

Chain of Command



What does Infrastructure report on?


- Infrastructure reports on the general activity and future of the infrastructure at the ASF:
 - General activity
 - Significant events the past month (CVEs, faulty h/w, maintenance, upgrades)
 - Overall uptime statistics (see next slide)
 - Ongoing changes to the infrastructure
 - Future development plans
 - Post mortem on failures and incidents

Service level agreements and reality

- The Infrastructure team is bound by an SLA for a select group of services:
 - **Critical** services (mail, web sites, svn, git) must have 99.50% uptime
 - **Core** services (BugZilla, JIRA, CI, Whimsy, SSL Frontends) must have 99.00% uptime
 - **Standard** services (Weblogs, Wikis, Pootle, ReviewBoard etc) must have 95.00% uptime.
 - Overall, services must have an average 98.00% uptime.

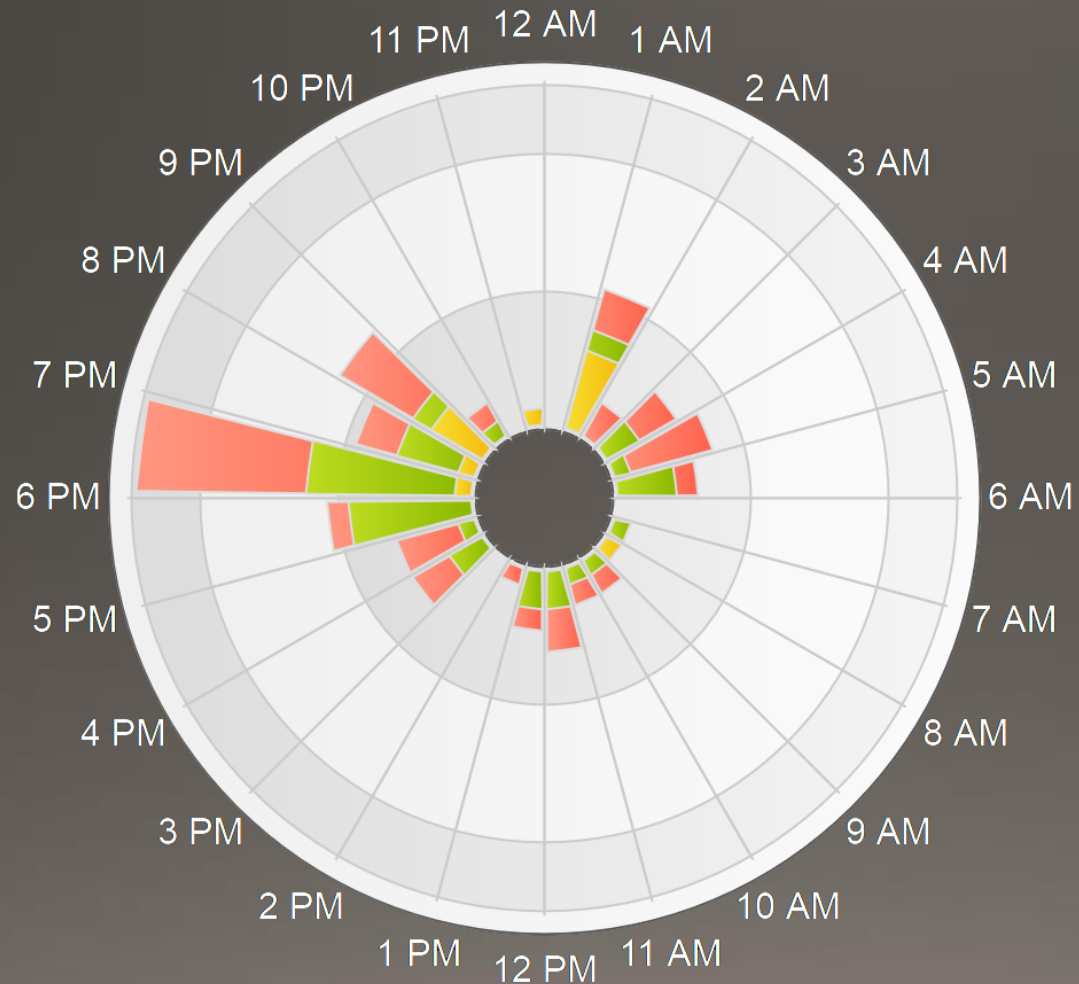
Service level agreements and reality

- Current uptime statistics for the past 6 months:
 - Critical services: 99.96% (0.46% above target)
 - Core services: 99.75% (0.75% above target)
 - Standard services: 97.98% (2.98% above target)
- Overall uptime: 99.31% (1.31% above target)
- Source: <http://s.apache.org/uptime>



Service level agreements and reality

Each week, we have between 5 and 10 service failures




Failures and service restoration relative to (my) time of day, past month

What does infrastructure do?


What does infrastructure do today?

- Infrastructure manages the 40+ unique services used by today's developers and users.
- These services include:


- Mailing lists
- Mail archives
- Committer email accounts

Email 


- Subversion repositories
- Git repositories
- Nexus repository

Code repositories 


- Main apache.org site
- Project web sites
- Paste, comment, pad etc

Web sites 


- JIRA
- Bugzilla
- ReviewBoard
- Allura

Issue and bug tracking 


- Moin Moin Wiki
- Confluence Wiki

Wiki services 


- Dist repository
- Release archive
- RSYNC

Release distribution 

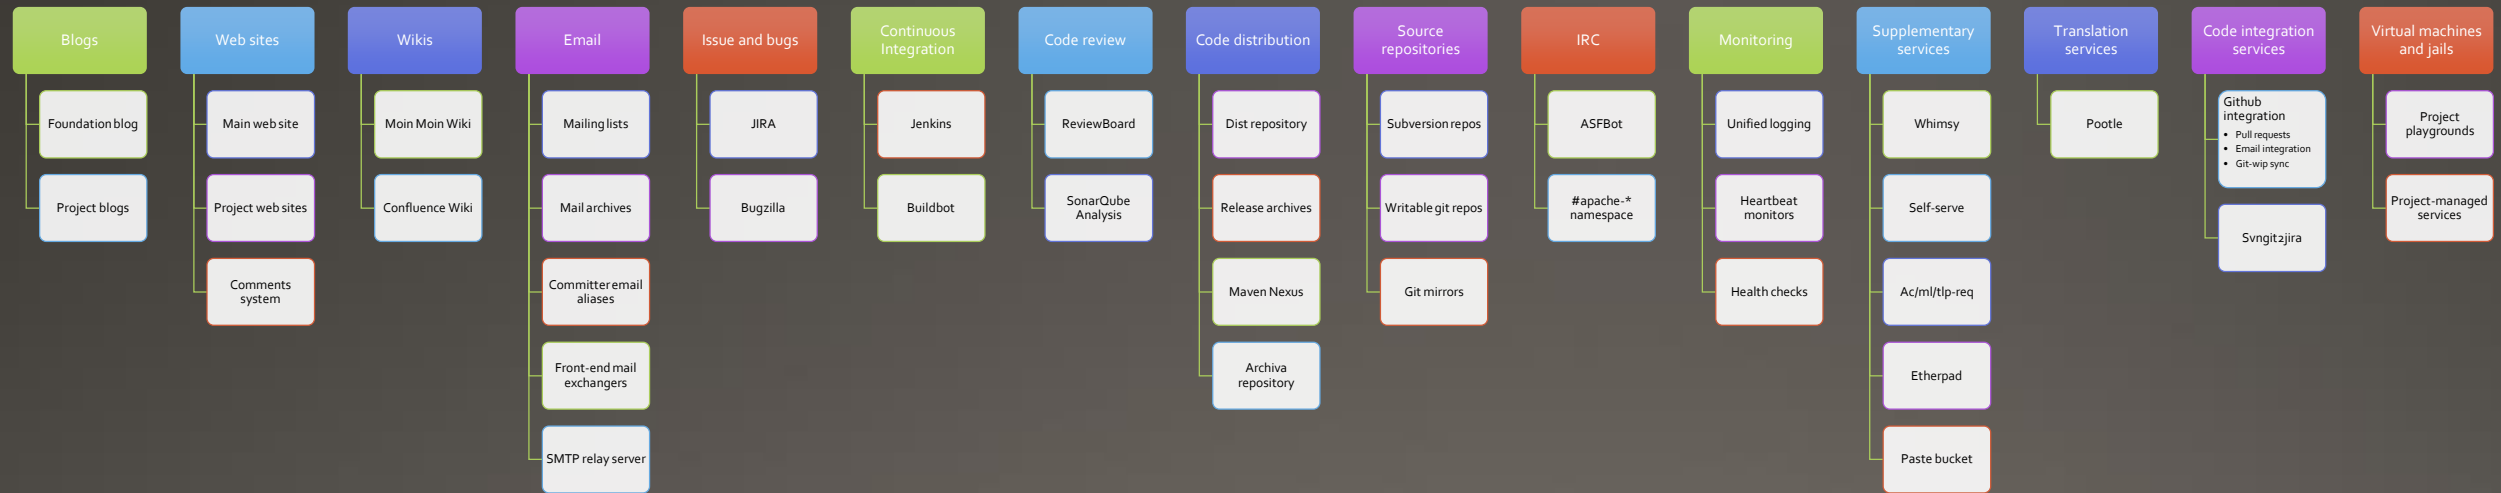
- Logging
- Heartbeats
- Health checks

Monitoring 

- Buildbot
- Jenkins

Continuous Integration 

Heck, let's list all services!*



*These are the ones I could think of

How does infra work on a daily basis?

- Most direct day-to-day communication happen on HipChat
 - <http://s.apache.org/infrachat>
- Important decisions/discussions happen via the mailing lists
- Weekly operational team meetings take place on Google Hangout
 - *And everyone is invited, see our HipChat room topic.*
- Weekly and monthly reports are shared with the ASF operational group

Who does what on a normal day?

- The bulk of infra handles open tickets and ongoing projects
- On-call staff handles immediate queries/alerts and account creations on a rotating week-by-week basis
- Escalation plans are in place that delegate tasks to staff based on response times and severity of incidents

Typical questions we get

1. This doesn't work! (no details provided)
 - Please always provide enough information to replicate the error/bug
2. I can't commit anything to the repo!
 - Make sure you use https instead of http! Make sure you're not banned!
3. I can't log onto JIRA/BugZilla/whatever using my LDAP creds!
 - We don't use LDAP for everything (yet!), some services require local accts.
4. Your \$project software is ruining my life, fix it!
 - Yeah, I'm gonna need you to come in on Sunday and work late...
5. Unsubscribe me!!!
 - Read the footer in the ML emails you get, it has a link to unsubscribe you

Contacting Infrastructure

- Canonical contact list: www.apache.org/dev/infra-contact.html
- We are no longer on IRC – use HipChat: <http://s.apache.org/infrachat>
- Via email: infrastructure@apache.org
- Or, you can break something and we'll notice immediately.

That's all, folks!

For inquiries, comments, snark:

- humbedooh@apache.org
- Twitter: [@humbedooh](https://twitter.com/humbedooh)