



Inside the ASF Infrastructure Team

Philip M. Gollucci <pgollucci@apache.org>

Who Am I? (in 15s)



- VP Apache Infrastructure
- root@ member
- HTTP PMC [yes, I just broke it yesterday]
- APR committer
- mod_perl PMC
- FreeBSD ports/ committer
- Director of Development Operations @ Taxi Magic
- Freelance Consultant
- Reader of 261 mailing lists @apache.org

In short, I'm over committed!



What is the ASF Infra Team?



A group of people (with copious free time) that enjoy making things go smoothly.



What is the ASF Infra Team?



A group of people (with copious free time) that enjoy beer.



What is the ASF Infra Team?



A group of people (with~~out~~ copious free time) that enjoy
dealing with end users.



What is the ASF Infra Team?



- Its not a TLP
- Its a President's Committee
- The VP is appointed by the President
- The VP can gather others to help him/her



Who's in the know



Username	Name	Location	Years	DayJob	Speciality
jim	Jim Jagielski	Maryland, USA	10	CTO	President
noirin	Noirin Plunkett	Ireland	5	Tech writer	EVP, English
pgollucci	Philip Gollucci	Maryland, USA	6	Director Ops	FreeBSD, Perl
joes	Joe Schaefer	Florida, USA	7	FT SysAdmin	C, Perl, Linux
gmcdonald	Gavin McDonald	Brisbane, Australia	3	FT SysAdmin	Builds, Java
pctony	Tony Stevenson	Cambridge, UK	3	X PT SysAdmin	VMs, Windows, Nagios
danielsh	Daniel Shahaf	Givatayim, Israel	1	PT SysAdmin	SVN, Jack of All trades
markt	Mark Thomas	Hampshire, UK	4	Consultant	Bugzilla, Tomcat, Jira
pquerna	Paul Querna	California, USA	6	X VP Infra, CTO	Python, httpd, Lua
arrayder	Chris Rhodes	Iowa, USA	1.5	Engineer	Networking, Perl
norman	Norman Maurer	Germany	3	System Admin	SunOS, zfs, qmail



Who's in the know



Username	Name	Location	Years	DayJob	Speciality
jim	Jim Jagielski	Maryland, USA	10	CTO	Instigator
noirin	Noirin Plunkett	Ireland	5	Tech writer	Red Headed Devil
pgollucci	Philip Gollucci	Maryland, USA	6	Director Ops	Long Winded
joes	Joe Schaefer	Florida, USA	7	FT SysAdmin	Difficult
gmcdonald	Gavin McDonald	Brisbane, Australia	3	FT SysAdmin	Damn Aussie
pctony	Tony Stevenson	Cambridge, UK	3	X PT SysAdmin	Clueless
danielsh	Daniel Shahaf	Givatayim, Israel	1	PT SysAdmin	Paranoid
markt	Mark Thomas	Hampshire, UK	4	Consultant	Insane
pquerna	Paul Querna	California, USA	6	X VP Infra, CTO	Lazy
arreyder	Chris Rhodes	Iowa, USA	1.5	Engineer	Insomniac
norman	Norman Maurer	Germany	3	System Admin	Deluded



Lets See it on a Map



~550 of ~3,200 committers shown



First, Some Trivia



- apache.org shared the same physical server with:

mcdonalds.com

alt.drugs

burning man



More Trivia



- One of the 1st and fastest mirrors of apache.org was

mirrors.microsoft.com

mirrors.playboy.com

mirrors.ibm.com



More Trivia



- Some of the earliest servers are now memorialized where?

museum

unplugged in some rack in some data center

brian's mantle in his house



How did we do all this?



- 1994: NSCAAd
- 1996: daedalus & icarus by Brian
- □ 1996: 1st commit to Apache Group CVS
- 1999: The Apache Software Foundation
- 1999: Apache Board elected yearly
- 2001: United Layer (San Francisco)
- 2002: `make buildworld` takes 2 months
- 2002: Infrastructure Committee created
- 2002: Apache Incubator created
- 2003: 200 Paul DC started (San Francisco)
- 2003: CVS -> SVN



How did we do all this?



- 2004: Sara (.NL) DC started hosted by Surfnet
- 2005: Sun Hardware Kit
- 2005: CVS -> SVN finishes finally!
- 2006: 1st paid SysAdmin
- 2006: The great datacenter move
- 2006: OSUOSL (Oregon) DC started (moved)
- 2006: minotaur dies in transit
- 2007: Dell Hardware Kit
- 2008: VP ASF Infrastructure created



How did we do all this?



- 2008: Traci.net (Florida) DC started
- 2009: Formal budgets
- 2009: 2nd paid SysAdmin
- 2009: Git Read only support
- 2010: Rsync -> SvnPubSub and in-house CMS
- 2010: Git Read/Write support
- 2011: 3rd paid SysAdmin (P/T)
- 2011: PT Sysadmin swap
- 2011: iXSystems Kit pondered
- 2011: VP ASF Infrastructure moved to EVP
- 2012: dist.apache.org, SvnPubSub mandated



What do we actually do?



- Manage / Coordinate our Data Centers
 - OSUOSL (Portland, Oregon) 3 Racks
 - Traci (Miami, Florida) 1 half Rack
 - Sara (Netherlands) 1 Rack
 - Bytemark (London) VPS



What do we actually do?



- img1
- img2



What do we actually do?



- Buy Stuff (Servers, Switches, VPN, SSL Certs, Licenses)
 - Dell
 - Sun/Oracle
 - Microsoft
 - Sonicwall
 - Dotster domains
 - Intervision
 - Godaddy SSL
 - Wells Fargo
 - Silicon Mechanics

Yes, that means we make a budget!



What do we actually do?



```
== Infrastructure ==

Staffing                                :$212,624.00
  System Admin (FT)                    :$186,904.00*
  System Admin (FT)                    :$179,560.00*
  System Admin (PT)                    :$139,000.00
  Bonus Pool                            : $7,160.00 (4260/3900)

Hosting                                  : $8,216.00
  Cloud                                  : $2,000.00
  OSU OSL (OR, USA)                     : $0.00
  SURFnet (AMS, NL)                     : $0.00
  Traci.net (FL, USA)                   : $6,216.00

Hardware                                  : $94,000.00
  Replacement                           : $45,000.00
  Expansion                              : $20,000.00
  Build farm                             : $24,000.00
  Disks/RAM                              : 5,000.00

Service contracts                        : $7,000.00
  Dell                                   $4,500.00
  Silicon Mechanics                      $1,500.00
  Sun                                     $500.00
  Other                                  $500.00

Misc                                      : $1,800.00
  Spamhaus RBL                          : $300.00
  SSL / DNS renewals                    : $700.00
  3rd Party Thank you                   : $800.00

Travel                                   : $17,400.00
  3 Staff+VP @ ACNA2011                 : $12,000.00
  1 infra mtg                           : $5,000.00
  ACNA2011 food 1 mtg                   : $400.00

Category Total                          :$341,040.00 (+28%)
```



What do we actually do?



```
== Infrastructure ==

Staffing                               :$212,624.00
  System Admin (FT)                   :$ 86,904.00*
  System Admin (FT)                   :$ 79,560.00*
  System Admin (PT)                   :$ 39,000.00
  Bonus Pool                           : $7,160.00 (4260/3900)

Hosting                                 : $8,216.00
  Cloud                                 : $2,000.00
  OSU OSL (OR, USA)                   : $0.00
  SURFnet (AMS, NL)                   : $0.00
  Traci.net (FL, USA)                 : $6,216.00

Hardware                                : $94,000.00
  Replacement                          : $45,000.00
  Expansion                             : $20,000.00
  Build farm                           : $24,000.00
  Disks/RAM                            : 5,000.00

Service contracts                       : $7,000.00
  Dell                                  $4,500.00
  Silicon Mechanics                    $1,500.00
  Sun                                   $500.00
  Other                                 $500.00

Misc                                     : $1,800.00
  Spamhaus RBL                         : $300.00
  SSL / DNS renewals                   : $700.00
  3rd Party Thank you                  : $800.00

Travel                                  : $17,400.00
  3 Staff+VP @ ACNA2011               : $12,000.00
  1 infra mtg                          : $5,000.00
  ACNA2011 food 1 mtg                  : $400.00

Category Total                         :$341,040.00 (+28%)
```



What do we actually do?



- □ Install, Configure, **Secure**, **Maintain**, **Debug** OSes
- FreeBSD
- Ubuntu
- Solaris
- Windows
- Jails
- Zones
- VMs



What do we actually do?



- Services (internal and external)
- VC (cvs->subversion, git)
- WWW (100+ tlps, foundation site, others, httpd)
- Mailing Lists (300+, archives, searches, spam)
- Shell/SSH
- Release Distribution and Mirror Network
- Dev (Jira, Bugzilla, Confluence, moinmoin, mediawekia, Roller, custom stuff)
- Dev (continuum, gump, nexus, sonar, url shortener, reviewboard)
- Monitoring (nagios, nagvis, irc, pagers, twitter)
- LDAP
- VPN
- NFS
- DNS (domains, geo-balancing of services)
- User Account Mgmt (Self Serve) All 3,318 of them
- FBSD Tinderbox
- FBSD Update Server
- Backups
- Log Stats
- CI (Jenkins, Buildbot, Hudson)
- MySQL, PostgreSQL
- Voting System
- And many more.....



What do we actually do?



- Committers: 3318
- Members (active): 370
- iCLAs on file: 4158
- Projects: 94
- Incubating: 57
- Labs: 32
- Attic: 16
- SVN Revisions: 1.3million+ (just the public asf one)
 - asf: 119 GB
 - infra: 13 GB
 - private: 21 GB
 - tck: 30 GB
- Web stats: <http://people.apache.org/~henkp/analog/>
- Mailing lists and SVN commits graph: <http://www.apache.org/dev/stats/>



How do we do all this?



- Remember only 2.5 of us are paid, the rest of us are volunteers (that come and go) without offices and in sporadic time zones!
- With volunteers you might be waiting Y-E-A-Rs for something to get done
- Volunteers are not always experts at what they are contributing



How do we do all this?



- The Apache Way (meritocracy)
- Version control
- Mailing lists, archives
- IRC
- Recruiting new members by involving the communities
- Saying NO!
- Admitting mistakes
 - ssh via sourceforge
 - ssh key weakness via apachecon vps
 - jira xss via slicehost
- Not repeating Mistakes
- Random meetings at conferences

Oversight, **Trust**, & Transparency!



How do we do all this?



- Tools of the Trade
 - asf-do.pl
 - apue & pw-driver.pl
 - opie
 - ckl
 - ssh & ssh keys!
 - svnpubsub



ASF-DO.PL



```
$ apue/pty -d apue/pw-driver.pl screen
OPIE Secret Password (^D aborts):
LDAP Password (^D aborts):

$ ./asf-do.pl --terminal --unbuffered --serial --host minotaur,hermes 'sudo date'
minotaur.apache.org:
minotaur.apache.org:          otp-md5 365 mi3675 ext
minotaur.apache.org:          Password:
minotaur.apache.org:
minotaur.apache.org:          Fri Nov  4 23:23:37 UTC 2011
minotaur.apache.org:          Connection to minotaur.apache.org closed.
hermes.apache.org:
hermes.apache.org:          otp-md5 465 he6315 ext
hermes.apache.org:          Password:
hermes.apache.org:
hermes.apache.org:          Fri Nov  4 23:23:42 UTC 2011
hermes.apache.org:          Connection to hermes.apache.org closed.
```



Cloud Kick Logger(CKL)



```
ckl list
( 85) 2011-09-16 18:04:58 UTC by danielsh on eris.apache.org
      apply svnmailer-1.0.8 patch from Andre Malo (nd@)
( 18) 2009-12-31 07:19:32 UTC by pquerna on eris.apache.org
      Banning 216,129,119,0/24 from SVN
(   3) 2009-12-24 17:15:42 UTC by joes on eris.apache.org
      installed on eris
[11:42 PM](0)pgollucci@eris.apache.org j0 h53 c553
~
ckl cat -i 18 -U
http://monitoring.apache.org/ckl//display/?id=18;
```

[ckl web](http://monitoring.apache.org/ckl//display/?id=18;)

```
$ ckl save

commands and output here
# changelog entry:
#   host:   eris.apache.org
#   user:   pgollucci
#   endp:   http://monitoring.apache.org/ckl/
# (lines starting with # are ignored)

OK to continue (y/n)? y
entryid=89

$
```

- Very useful for distributed System Admin teams
- Already used by RackSpace (CloudKick), Mozilla, ASF, \$work
- And now you!



SvnPubSub & CMS



- I want to edit a webpage
- I want to preview it
- I want to stage my changes
- I want to see diffs
- I want an e-mail about it
- I want a webpage that uses more than just static html
- I want to push them live to production
- I want it to be '1-click'
- I want it to take seconds

Crap, Production is on both sides of the planet

No, I didn't let you down it is possible.

We developed a subscriber/push system for SVN and custom in-house CMS on top of it with ZFS.



How do we do all this?



- For most people failure **IS** the exception not the **NORM**
- We don't have hundreds or thousands of machines
- Do you know the serial numbers for support or parts if the machine is down because of a hardware failure?

- We do, we keep it SVN



How do we do all this?



- Do you have a port-map ?
- We do, we keep it in SVN



How do we do all this?



- Do you have the documentation about your PDUs, switches, console servers, and other obsolete hardware your vendors don't support anymore ?
- We do, we keep it in SVN



How do we do all this?



- Do you version your machine config ?
- Do you version your application config ?

- We do, we keep it in SVN



How do we do all this?



- UTC=EDT+4 (Florida, USA)
- UTC=EST-10 (Brisbane, Australia)
- UTC=IST-2 (Jerusalem, Israel)

Thats almost 24/7 coverage. 4 Would be better. We'll get there eventually.

In the mean-time, what if I'm a volunteer with Karma and something breaks? I'm even motivated to fix it but I've never messed with this service before?

RUN Books



Run Books



- Start
- Stop
- Log File
- Where is it installed
- Where is its config
- How to install/upgrade
- Not just services, but processes too



Run Books



```
/home/pgollucci/repos/asf/infra/infrastructure/trunk/docs/services
%
ll
./          ldap.txt
../        mail.txt
README     mirrors.txt
archiva.txt modules.a.o.txt
backup_topology.txt mx.txt
bugzilla.txt nexus.txt
buildbot.txt ooo-forums.txt
cms.txt     projects_ci_server_matrix.txt
confluence.txt projects_ci_server_matrix.xls
continuum.txt gmail-ezmlm.txt
dns-updates.txt roller.txt
dns.txt     sonar.txt
fail2ban.txt sshd.txt
git.txt     svn-dist.txt
gump.txt    svn-setup.txt
http.txt    svn.txt
issuetracking.txt svnpubsub.txt
jira.txt    upgradeconfluence.txt
ldap-lpk.txt wiki.txt
ldap-restore.txt
```



Communications



- Ever had a user e-mail you to say that e-mail doesn't work ?
- We sure have
- How do you communicate schedule maintenance?
- IRC, [Twitter](#), Facebook, [Nagios](#)



Mailing Lists



List	Purpose	Subscribers
infrastructure	Tickets, svn, git	329
infrastructure-issues	Jira Tickets	111
infrastructure-private	People doing the work	72
infrastructure-dev	Large project public discussions	131

You do archive you mailing lists right ?



IRC



- Whats a good motivator for your friendly System Admin ?
- B-E-E-R!
- Checkout our [Infrabot](#)
- If my talk was Wednesday
- We're taking bribes **every** night this week
- Oh and we accept deliveries the other 51 weeks too
- You could also do a lightening talk about how great we are



Funding

- 100% Donations 501c(3) charity
- Thanks to our Sponsors!



Questions, Comments



- Slides - <http://goo.gl/3c2Hn>

