



FreeBSD + ASF software &
philosophy + ZFS
==
large \$\$ bonuses from your boss


Why you're here



A record holding TCP/IP stack, *the* cutting edge File System, and most respected Software collection..... What do you get? A much happier system admin and a much better work place. The cloud might be hot stuff, but its not for everyone or every task. How can you take your existing Data Center infrastructure and make it better?



Why you're here



Why you're here



My Past



Job	Title	Tools
Univ. of Maryland	Student, Teacher, Webmaster	FreeBSD, perl, Apache
NEAR Mission	Developer	IDL, Dec Unix
EjournalPress	Sr. Engineer	Windows NT, IIS, Active Perl, Oracle, MSSQL, Apache
Consultant	American Apparel, ISST, others	Misc, Apache
Liquidity Services	Sr. Developer	Redhat, mod_perl, Apache
Ticketmaster	Sr. Software Engineer	Redhat, mod_perl, Apache
RideCharge	Sr. System Admin	FreeBSD, perl, ruby, zfs, Apache
RideCharge	Director Development Operations	FreeBSD, Junipers, Apache



My Current



- Racks of FreeBSD machines
- ~100 ALIX based FreeBSD CDN network
- 10,000 devices in Cabs
- httpd 2.2 + Rails/Passenger
- MySQL/Percona
- FreeBSD+ZFS+Puppet
- Mailing Lists
- An ultra real-time environment
- A budget thats too small
- Sparse staff



My Future



FreeBSD



- since 2.2.8 (1998)
- ports/ Committer
- Licensing
- ZFS
- dtrace
- jails
- source over binaries
- CVE fix rates for PCI DSS Lv 1
- [Tarsnap](#)
- iXSystems



ZFS



- snapshot/clone/promote/rollback
- **send/recv**
- zmirror, raidz[1-3], zil, cache
- data deduplication
- zfs sharenfs



Apache Software



- apr
- httpd
- subversion
- mod_perl
- svnsub
- mod_python
- Infrastructure and RideCharge get to share
 - asf-do.pl / rwsdo.pl
 - ckl / rws
 - wallet.pl



Apache Philosophy



- Release Management
- Mailing Lists
- Meritocracy
- Revision Control
- Transparency
- Agile



These actually happened



The Daily WTF?



Days since last face<->palm: 0



Boss: Upgrade Twiki



- Admin: Are you sure?
- Boss: yes, tomorrow night after hours
- Admin: okay

- svn vendor imports
- zfs snapshot zroot/path/twiki@a.b.c
- do upgrade
- forget to test 1 page that uses one core plugin

- Boss: Undo it now, nobody can work
- Admin: **Sure thing!**

- zfs rollback zroot/path/twiki@a.b.c
- /usr/local/etc/rc.d/apache restart



Boss: Bug in Production



- We're going to blow our allotted SIM card shared bandwidth plan on the 1st of the month (aka Monday).
- That's going to cost us \$10,000 USD
- Can you fix it?
- Without updating the software on all the field devices (POS)

```
$ sysctl net.inet.tcp.fast_finwait2_recycle=1  
$ sysctl net.inet.tcp.finwait2_timeout=30000  
$ sysctl net.inet.ip.fw.dyn_keepalive=1
```

```
$ rwsdo.pl --env prod 'sudo .....
```

- Ops: How about you treat us to a steak dinner?
- Boss: **Sure thing!**



Boss: Deployments



- Fast
- No Downtime
- Close to 1-click
- People other than Ops need to be able to do it

Admin: Well we're an RoR shop, lets use Capistrano

Turns out it sucks, too late its here to stay.

Fear not



Boss: Deployments



- zfs create zroot/path/qa
- zfs snapshot zroot/path/qa@a.b.c
- zfs send zroot/path/qa@a.b.c | \
 - ssh production 'zfs recv zroot/path/prod@a.b.c

Wait this code is still broken!

zfs rollback, or send+recv a different tag

zfs send -i : incremental



Boss: PCI DSS Lv1



- □ 75 page document
- 107 items
 - keep system access data for 1yr
 - keep 90 days hot
 - be able to restore quickly in the event of a breach

Admin: okay this might take some time.
100s of systems, 10s of GB of data / day

Enter rsyslog+relp+tarsnap+EC2/S3

Did you know CERN's LHC uses rsyslog ?



Boss: PCI DSS Lv1



rsyslogd (clients) -> RELP -> rsyslogd (central)

rsyslogd -> logsaggd

real-time aggregation / data warehousing

logsaggd -> MySQL / logstash

some terrific table rotation via percona-toolkit
partitions/merge tables

MySQL -> tarsnap



Boss: PCI DSS Lv1



- Now, our solution is simple
- Daily cron
 - mysqldump table > table.sql
 - tarsnap -C dir -cf table-yyyymmdd table.sql
 - journaled, ref counted, deltas, compression
- Now lets restore
 - create database restores
 - tarsnap -x -f table-yyyymmdd | mysql -D restores
- Takes mere minutes
- SELECT * FROM system_logs WHERE

I don't even need a DATE constraint

Tarsnap+EC2 provide a PCI Lv1 compatible solution for storing offline Encrypted Credit Card Data.



Boss: Verify POS gets correct Content



- Basically we show you advertisements
- Checksum on the server
- Checksum on the PIM
- compare, ==, everyone is happy

WRONG

- CMS+CDN launches, we run out of apache workers, requests take 'forever'
- 0-3GB of data locally on the PIM



Boss: Verify POS gets correct Content



Enter dtrace



Boss: Verify POS gets correct Content



- `dtrace -q -s requestime.d -p `pgrep httpd``
- `md5 file.mpeg ; md5 file.jpg`
 - 3GB of media
 - 10,000 devices making N requests to scan 3GB of data M times / day.

Run the md5 one time in the CMS (internal tool) on file upload, store it in DB. Pass it as part of the manifest sent to the PIM.



Boss: Cloud, Cloud, Cloud



- Cost effective
- Good for the herding problem
- Eases backups, mgmt, scaling
- We should to virtualize



Admin: Facts!



- Cost effective [not for us]
- Good for the herding problem [we don't have one]
- Eases backups, mgmt, scaling [got this licked]
- We should virtualize [enter freebsd jails]



Boss: Jails?



- Yes, Jails
 - man 8 jail
- I can run 100s on my desktop without breaking a sweat
- What should we virtualize?
- DEV, QA, STG, Integration Environments
- Boss: We doubled our Developers in 15 days
- Boss: We need more environments to test in
 - `sudo rsync -rlptgoDvHExh $1/ $2/`
 - egads thats ugly
- `zfs snapshot zroot/jail/A@date`
- `zfs send+recv` it even moves across dungeon masters
- `dev1 --> devN`
 - set the MAC, let puppet run



Boss: Quarterly updates

- Enter FreeBSD Update and FreeBSD Tinderbox



Jails/VMs

- FreeBSD+VirtualBox+Puppet+ZFS
 - vagrant up



Philosophies: Release Mgmt



- Apache Infrastructure is actually agile just without the frills.
- RideCharge actually *is* Agile.
- Iterations require release mgmt
- Dev, then QA, then release
- QA shouldn't block new DEV
- New Dev shouldn't block QA
- Should be able to have emergency single change fixes

- Not taught in colleges
- Branches, tags, merging
- VC agnostic
- Very few people understand what a version triplet means.
- Ruby, Rails, & Rubygems I'm looking at you!
- ~> Really? (tears)

- Apache got this right almost a decade ago

Thanks!



Philosophies: Mailing Lists



- You do send commit diffs to a list right ?
- You do archive it right ?
- You do have other Internal Lists that are searchable?
- Google Apps is quite good at this if you aren't a mailman pro

Thanks Apache!



Philosophies: When they fail?



- 'I' don't like loose ethernet cords
- 'I' plugged it in, the b/c phone wasn't working
- Boss: Call the ISP our Internet is down!
- Juniper firewalls+switches+Spanning Tree

Set up mitigating controls in technology



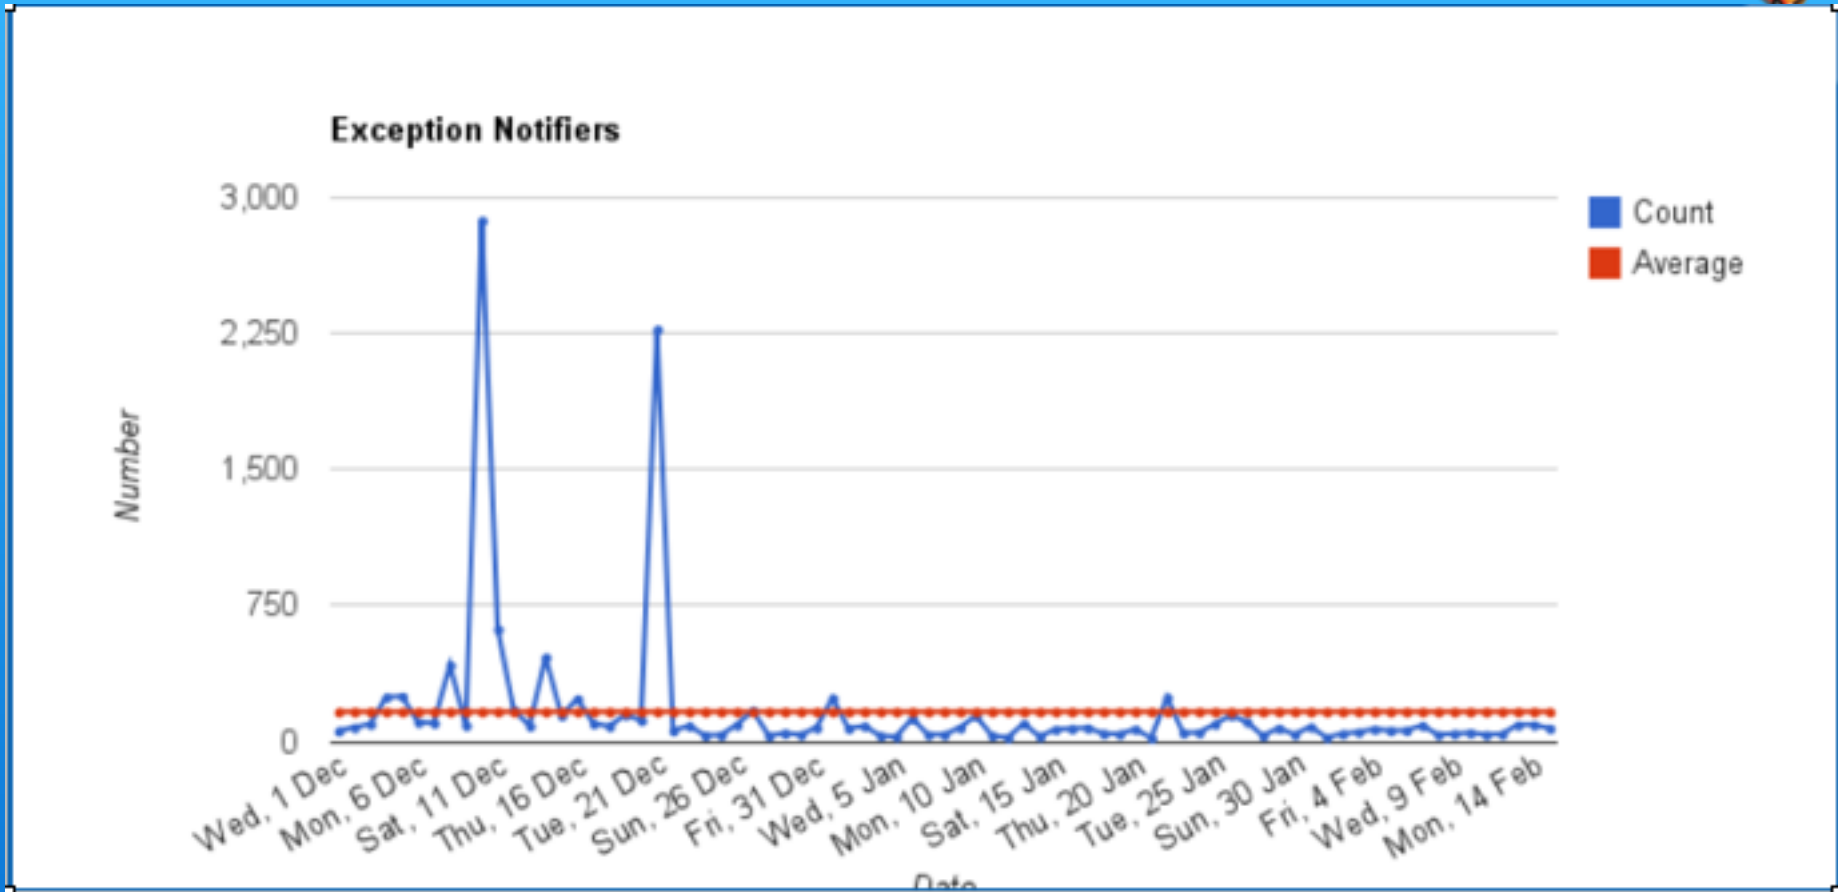
Meritocracy



- Teach Developers how to do common sysadmin tasks
- Enable Developers to do so
- Reward Developers that do (more system access)
- More / Faster support for their pet projects
- Wait those usually end up being projects Ops wants too!



FeedBack



163 ENs/day on avg (0.0001% of requests)
In 1 yr, with 0% EN growth: 1,460/day



When you stray to the dark side



prod.dca1.rc-mock-system	100.000% (100.000%)	0.000% (0.000%)	0.000% (0.000%)	0.000%
Average	100.000% (100.000%)	0.000% (0.000%)	0.000% (0.000%)	0.000%

Servicegroup 'providers' Service State Breakdowns:

Host	Service	% Time OK	% Time Warning	% Time Unknown	% Time Critical	% Time Undetermined
prod.dca1.rc-mock-system	rc-provider-303taxi	97.955% (97.955%)	0.000% (0.000%)	0.000% (0.000%)	2.045% (2.045%)	0.000%
	rc-provider-alexandriayellowcab	98.742% (98.742%)	0.000% (0.000%)	0.000% (0.000%)	1.258% (1.258%)	0.000%
	rc-provider-ann arbor yellow	99.444% (99.444%)	0.000% (0.000%)	0.000% (0.000%)	0.556% (0.556%)	0.000%
	rc-provider-atlanta checker	99.905% (99.905%)	0.000% (0.000%)	0.000% (0.000%)	0.095% (0.095%)	0.000%
	rc-provider-barwoodtaxi	99.616% (99.616%)	0.000% (0.000%)	0.000% (0.000%)	0.384% (0.384%)	0.000%
	rc-provider-blue and white	99.875% (99.875%)	0.000% (0.000%)	0.000% (0.000%)	0.125% (0.125%)	0.000%
	rc-provider-bluediamondtaxi	99.783% (99.783%)	0.000% (0.000%)	0.000% (0.000%)	0.217% (0.217%)	0.000%
	rc-provider-bostoncabdispatch	99.862% (99.862%)	0.000% (0.000%)	0.000% (0.000%)	0.138% (0.138%)	0.000%
	rc-provider-bostonmetrocab	99.880% (99.880%)	0.000% (0.000%)	0.000% (0.000%)	0.120% (0.120%)	0.000%
	rc-provider-broadwaycab	99.716% (99.716%)	0.000% (0.000%)	0.000% (0.000%)	0.284% (0.284%)	0.000%
	rc-provider-californiayellowcab	99.882% (99.882%)	0.000% (0.000%)	0.000% (0.000%)	0.118% (0.118%)	0.000%
	rc-provider-chicago carriage	99.881% (99.881%)	0.000% (0.000%)	0.000% (0.000%)	0.119% (0.119%)	0.000%
	rc-provider-colorado springs yellow	0.000% (0.000%)	0.000% (0.000%)	0.000% (0.000%)	0.000% (0.000%)	100.000%
	rc-provider-columbus yellow	99.842% (99.842%)	0.000% (0.000%)	0.000% (0.000%)	0.158% (0.158%)	0.000%
	rc-provider-dallas golden	99.753% (99.753%)	0.000% (0.000%)	0.000% (0.000%)	0.247% (0.247%)	0.000%
	rc-provider-dallasyellow	99.798% (99.798%)	0.000% (0.000%)	0.000% (0.000%)	0.202% (0.202%)	0.000%
	rc-provider-eastside for hire	99.271% (99.271%)	0.000% (0.000%)	0.000% (0.000%)	0.729% (0.729%)	0.000%
	rc-provider-jacksonville checker	0.000% (0.000%)	0.000% (0.000%)	0.000% (0.000%)	0.000% (0.000%)	100.000%
	rc-provider-laclede	0.000% (0.000%)	0.000% (0.000%)	0.000% (0.000%)	0.000% (0.000%)	100.000%
	rc-provider-layellowcabco	99.683% (99.683%)	0.000% (0.000%)	0.000% (0.000%)	0.317% (0.317%)	0.000%
	rc-provider-loudounyellowcab	99.985% (99.985%)	0.000% (0.000%)	0.000% (0.000%)	0.015% (0.015%)	0.000%
	rc-provider-luxorcabs	99.051% (99.051%)	0.000% (0.000%)	0.000% (0.000%)	0.949% (0.949%)	0.000%
	rc-provider-manhattancorporatededanby	99.181% (99.181%)	0.000% (0.000%)	0.000% (0.000%)	0.819% (0.819%)	0.000%
	rc-provider-mearstransportationgroup	98.699% (98.699%)	0.000% (0.000%)	0.000% (0.000%)	1.301% (1.301%)	0.000%
	rc-provider-memphis yellow checker	97.731% (97.731%)	0.000% (0.000%)	0.000% (0.000%)	2.269% (2.269%)	0.000%
	rc-provider-naples checker cab	99.883% (99.883%)	0.000% (0.000%)	0.000% (0.000%)	0.117% (0.117%)	0.000%
	rc-provider-new orleans carriage	99.862% (99.862%)	0.000% (0.000%)	0.000% (0.000%)	0.138% (0.138%)	0.000%
	rc-provider-oakland friendly	100.000% (100.000%)	0.000% (0.000%)	0.000% (0.000%)	0.000% (0.000%)	0.000%
	rc-provider-orangecabcompany	99.838% (99.838%)	0.000% (0.000%)	0.000% (0.000%)	0.162% (0.162%)	0.000%
	rc-provider-pittsburghyellow	99.875% (99.875%)	0.000% (0.000%)	0.000% (0.000%)	0.125% (0.125%)	0.000%
	rc-provider-pomona yellow	98.378% (98.378%)	0.000% (0.000%)	0.000% (0.000%)	1.622% (1.622%)	0.000%
	rc-provider-portland radio	99.763% (99.763%)	0.000% (0.000%)	0.000% (0.000%)	0.237% (0.237%)	0.000%
	rc-provider-stlouiscounty	99.802% (99.802%)	0.000% (0.000%)	0.000% (0.000%)	0.198% (0.198%)	0.000%
	rc-provider-suburban minneapolis taxi	99.892% (99.892%)	0.000% (0.000%)	0.000% (0.000%)	0.108% (0.108%)	0.000%
	rc-provider-supershuttle	99.904% (99.904%)	0.000% (0.000%)	0.000% (0.000%)	0.096% (0.096%)	0.000%
	rc-provider-the cab	99.247% (99.247%)	0.000% (0.000%)	0.000% (0.000%)	0.753% (0.753%)	0.000%
	rc-provider-totaltransitphoenix	0.000% (0.000%)	0.000% (0.000%)	0.000% (0.000%)	100.000% (100.000%)	0.000%
	rc-provider-washington yellow cab	99.815% (99.815%)	0.000% (0.000%)	0.000% (0.000%)	0.185% (0.185%)	0.000%
	rc-provider-yc san diego	97.276% (97.276%)	0.000% (0.000%)	0.000% (0.000%)	2.724% (2.724%)	0.000%
	rc-provider-yellowcabofbaltimore	99.823% (99.823%)	0.000% (0.000%)	0.000% (0.000%)	0.177% (0.177%)	0.000%
	rc-provider-yellowcabofdenver	99.634% (99.634%)	0.000% (0.000%)	0.000% (0.000%)	0.366% (0.366%)	0.000%
	rc-provider-yellowcabofkanssas	99.764% (99.764%)	0.000% (0.000%)	0.000% (0.000%)	0.236% (0.236%)	0.000%
	rc-provider-yellowcabofsacramento	99.679% (99.679%)	0.000% (0.000%)	0.000% (0.000%)	0.321% (0.321%)	0.000%
	rc-provider-yellowcabofsansaniose	99.570% (99.570%)	0.000% (0.000%)	0.000% (0.000%)	0.430% (0.430%)	0.000%
Average		90.444% (90.444%)	0.000% (0.000%)	0.000% (0.000%)	2.738% (2.738%)	6.818%



Happiness



- Lets review
 - We get bonuses and free steak dinners
 - I'm at a bar, I can call a developer and have them restart something because I trust them
 - I can roll forward or backwards very quickly
 - Real-time business metrics for CEOs

- In short, I sleep better and work less thanks to Apache, FreeBSD, and ZFS



P.S.



- For the love of god, use an issue tracking system for System Administration Tasks.
- Force people to use it to avoid being Bowling Balled or becoming a Pinball.



Questions, Comments



- Slides - <http://goo.gl/0v1DH>

