

# Robust Infrastructure

*services for a global user base*




**[brong@opera.com](mailto:brong@opera.com)**

# *Introduction*

Who Am I?

Who is FastMail?

Who is Opera?



Email is dead

- 1 Seamless Messaging
- 2 Conversation History
- 3 Social Inb...

# *Email is Dead?*

Compatibility

Un-changeable

Privacy

Business / Orders / Receipts

## *Our Customers*

Individuals, Businesses, “White Label”

Services: Webmail / IMAP / “extras”

24hr Operations

1/2 million users / 3 part time admins

## *When I started – the bad bits*

Redhat 7.3

Manual “install script”

No replication

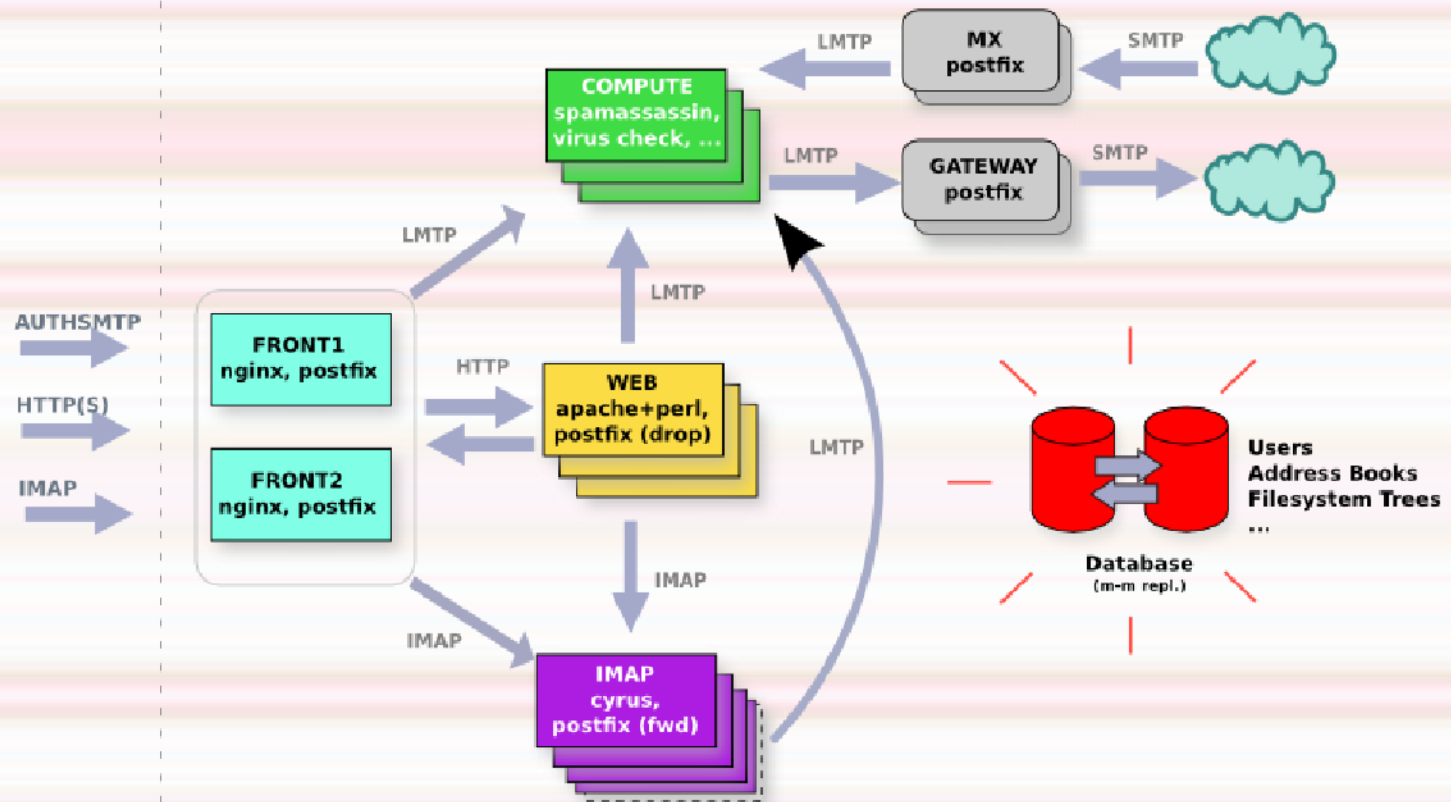
## *When I started – the good bits*

End to end monitoring

Too good! I kept getting woken up

Version control (CVS => SVN => Git)

# fastmail infrastructure



<http://www.fastmail.fm>  
<http://mail.messagingengine.com>

Last update: 19 Jan 2011



## ***Case study: RAID6 Failure***

3 disks failed in 12 hours

Over a week to get everyone back online


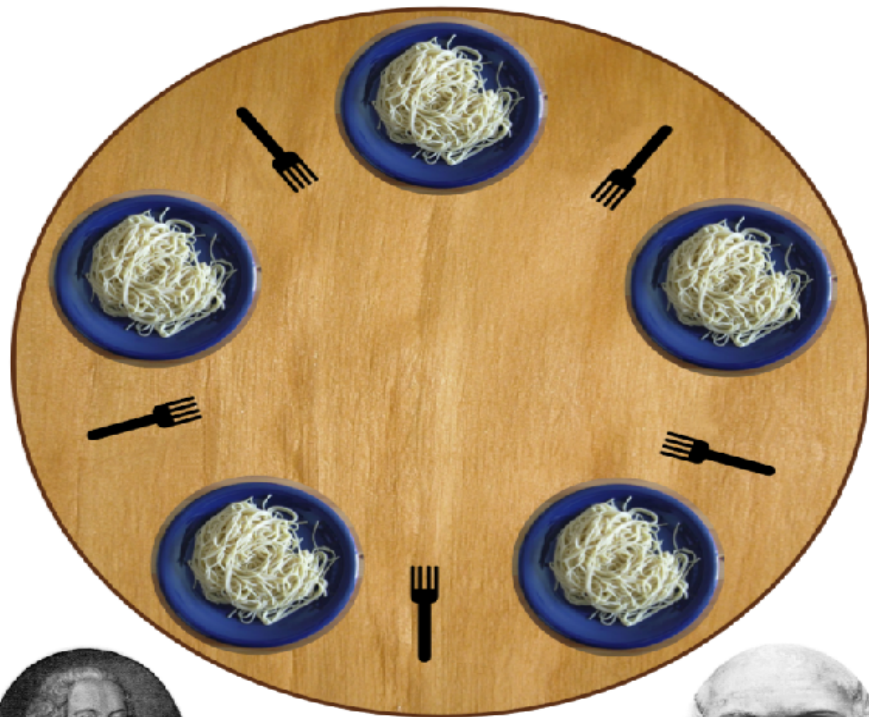
Filesystem check alone took days

## *Replication: choices*

Application native vs low-level (drbd)


Minimise single points of failure

Cyrus – something already existed



# Dining Philosophers Problem

**SOLVED**



## ***Side track – philosophy of data***

3 types ( creator / primary / cache )

Backups

Integrity

## *Data Integrity*

Never trust any layer

Check at the level where you care

Fail-safe checks

## *Linux-HA / Automated Failover*

More outages from heartbeat than real!

Cascading failures

Kept the tools, not the “brain”

## ***Control your core tools***

Cyrus: from patches to maintainer

Postfix: still 8 patches

Fix the right place, don't “work around”



## *Repeatable Automated Installs*

~10 min to reinstall any machine

“production.dat” microformat

Templated, version controlled config

## *Redundant Everything*

Can shut down any machine in 5 min.

Different replication per data type

“Never” need scheduled downtime

## *Slots and Stores*

Every machine works – no “hot spare”

Spread load on failure

300 => 500Gb “units of storage”

Hundreds of them!

## *Summary*

Reliability is an ongoing process

Monitor, fix what fails for next time

Simple, low interdependence is best

Make everything repeatable and version controlled

# *mail.opera.com*

The screenshot displays the My Opera email client interface. At the top left is the 'My Opera' logo. The main header area includes the 'E-post' and 'Kontakter' tabs, the user email 'brongondwana@myopera.com', and links for 'Innstillinger' and 'Logg ut'. Below this is a toolbar with icons for 'Arkiver', 'Slett', 'Rapporter søppelpost', 'Flytt til', and 'Marker', along with a search bar labeled 'Søk Innboks'. The left sidebar contains navigation options: 'Skriv', 'Innboks' (with 26645 items), 'Fastmail 1' (with 178 items), 'Arkiv', and 'Sendt'. The main inbox area shows a list of emails with columns for checkboxes, sender names, subject lines, and dates. The top of the list is filtered to 'Alle e-post'.

**My Opera**

brongondwana@myopera.com | [Innstillinger](#) | [Logg ut](#)

E-post    Kontakter

Arkiver    Slett    Rapporter søppelpost    Flytt til    Marker

Søk Innboks

Skriv

**Innboks** 26645

Fastmail 1 178

Arkiv

Sendt

Alle e-post    Uleste    Personlige    E-postlister    Sorter: nyeste først

<input type="checkbox"/>		<b>brongondwana (2)</b>	To Myself	14 okt
<input type="checkbox"/>		So hi there -- brongondwana@myopera.com		
<input type="checkbox"/>		<b>Bron Gondwana</b>	Testing stuff	10 sep
<input type="checkbox"/>		Hello again		
<input type="checkbox"/>		<b>Bron Gondwana</b>	testing stuff	10 sep
<input type="checkbox"/>		hi		
<input type="checkbox"/>		<b>Bron Gondwana</b>	Testing stuff again	10 sep