

How Zones Served Queen Victoria

Peter Tribble
ProQuest



The speaker...

- Full time sysadmin
 - Solaris, Linux, other unices...
 - Jack of all trades
 - Application focus
- Part time developer/enthusiast
 - OpenSolaris
 - Governing Board Member

ProQuest

- Online Publisher, US parent
- I deal with mostly literature and historical data
 - Early English Books, Early European Books
 - Vogue
 - Literature Online, Film, media, video
 - Periodicals
- Digital scans, indexing, strong search
- Customers – academic, universities, research
 - Not individuals or the public

One day

- “We've got this project...”
- Short term, high traffic
- Very public, on national media
- High profile, fantastic PR
- Potential for egg on face
- Can we build it?

Queen Victoria's Journals

- Princess Victoria started a journal in 1832
- Carried on as Queen until her death in 1901
- 141 volumes, 43765 pages
- Digitally scanned



Initial Evaluation

- Let's assume it's on the BBC News
 - And national newspapers
- And 1 million page views per day
- Given existing similar sites
- System infrastructure will be easy
- Network load might be interesting
 - Maybe 300M/s

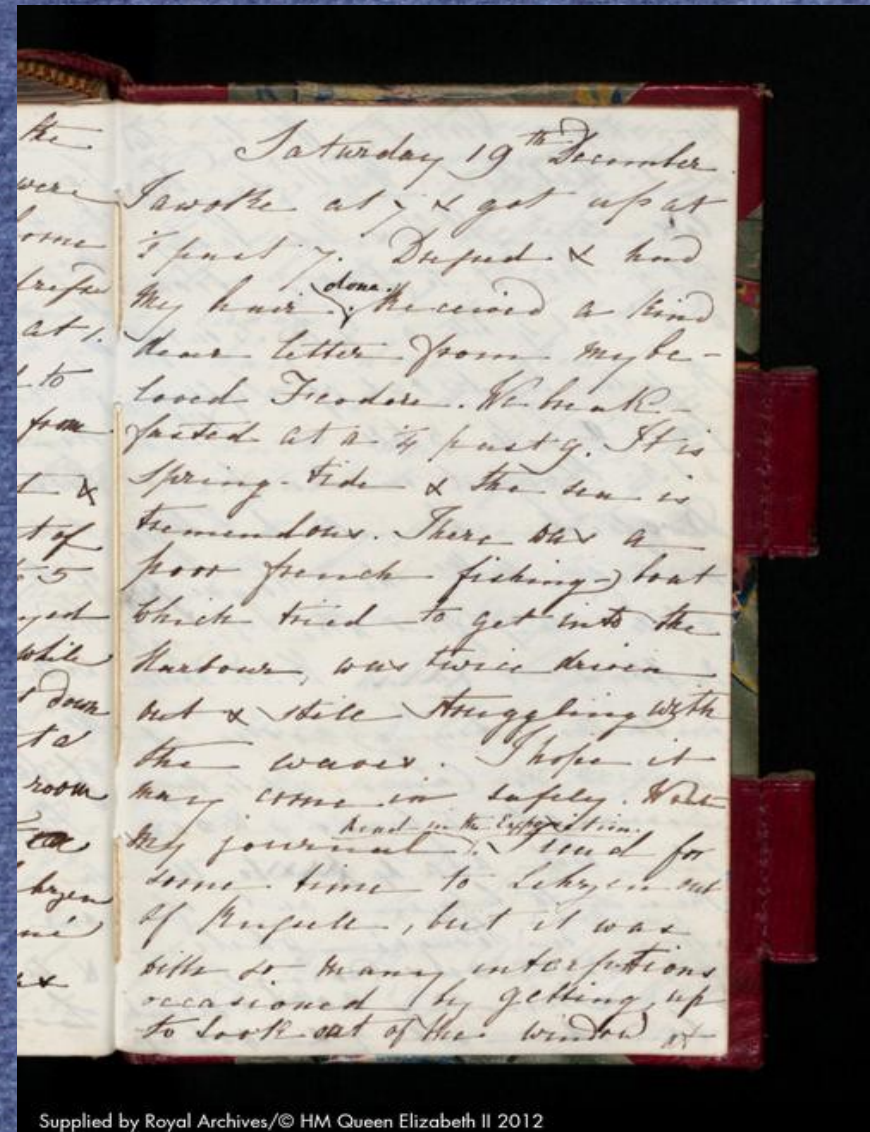
Handling the data

- Scanned/photographed
- Stored on ZFS
- ImageMagick
- Zoomify
- Replicated image servers
- Reprocess until customer happy
- Keyed data for search



Handwriting

- Handwritten texts only exist from the time before Victoria became queen
- Only 13 small purple and marbled volumes survive



Standard Architecture

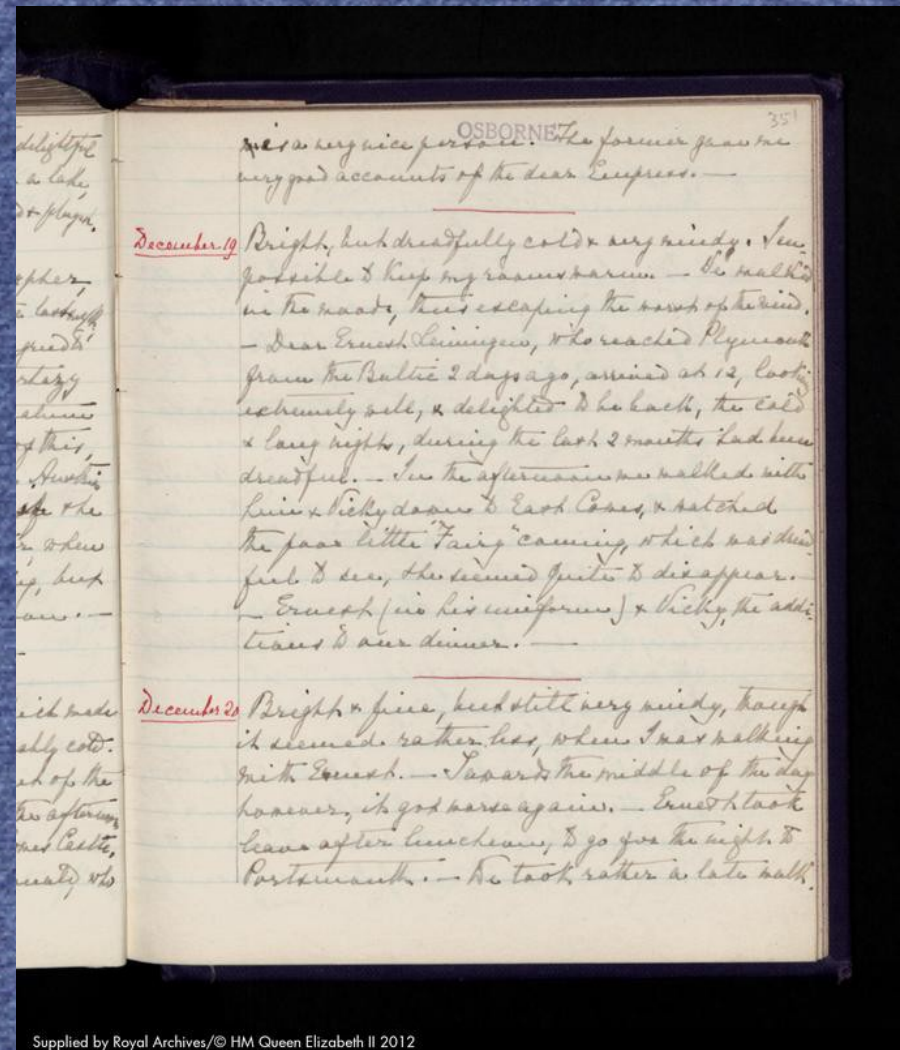
- Use Solaris with Zones for everything
- Sparse root, very lightweight
 - Guarantees every system is identical
- Application lives in one place
 - /opt/proquest
- Private application stack
 - Independence from vendor, patches
- Replicate horizontally
 - Just copy the one directory (+extras)

Application layout

- /opt/proquest
 - Same on every zone
- /opt/proquest/{apache,java,tomcat}
 - Read-only (root -owned)
- /opt/proquest/qvj/{apache,tomcat}
 - Configuration and logs only
- Virtualization at every level
 - System, application, application instance

Princess Beatrice Transcripts

- Manually transcribed after Victoria's death
- On her instructions
- Sanitised and edited
- Originals destroyed
- Sketches and doodles cut and pasted
- Finished in 1940



Build a zone

- `zonecfg -z qvj -f zconfig.qvj`

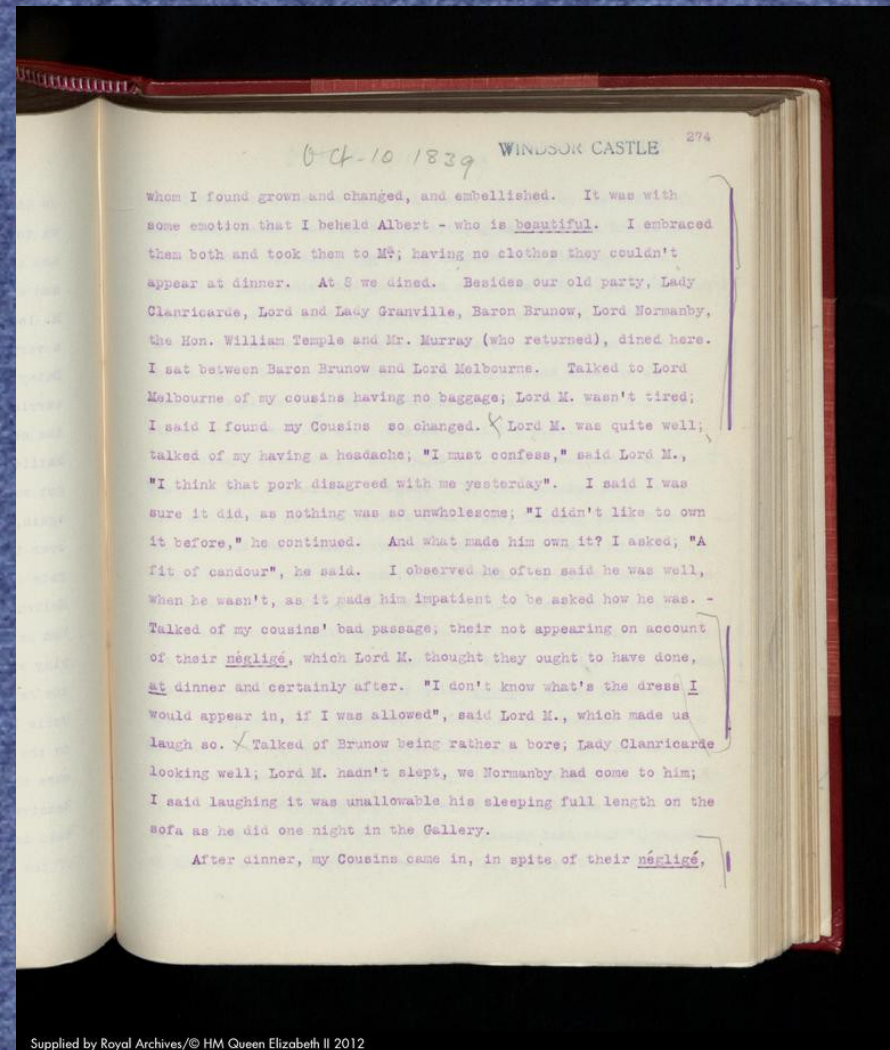
```
create
set zonepath=/storage/zones/qvj
set autoboot=true
add net
set physical=igb0
set address=172.18.1.36/24
end
add fs
set dir=/opt/sfw
set special=/opt/sfw
set type=lofs
set options=ro
end
add fs
set dir=/opt/proquest
set special=/storage/qvj/opt/proquest
set type=lofs
end
verify
commit
```


What we do NOT do

- Manage the zone
 - No packaging, no storage delegation
 - It's all about the application
- Patch zones
 - Migrate to newly patched host, then patch server
- Migrate zones
 - Save app, destroy zone, build new zone, restore

Lord Esher transcripts

- First Keeper of the Royal Archives
- Typewritten
- Unknown to Beatrice
- Unabridged
- 1832-1840



Production Infrastructure

- Load balancer handles individual failures
- N web zones (2, 4, 8) apache/tomcat
- M search zones (Solr in this case)
- Separate url and zones for images
 - Offsite at Everycity during launch
- Our normal site has MySQL – this doesn't
 - Authentication (free to UK)
 - History and personalization (remove)

System Capacity

- Most of our sites use tomcat, so each zone needs approx 2G of RAM
- OS overhead is essentially zero
- Typically low traffic, so memory is the constraint
- Even on busy sites, the application is a bottleneck
- So we basically size for RAM, and add zones until (just before) we start to swap
- Actually need to add lots of swap

The Artist

- Queen Victoria was a keen artist
- Scanned many of her sketches
- Also many illustrations in the journals



Testing, testing...

- Lots of new technologies (in our context)
 - Solr, no mysql, GeoIP
 - And go to current versions in the stack
- Need to work on each of those
 - Easy – create a zone
 - Actually, lots of zones...
- Use our standards to create lots of clones
 - **Clone the app, not the zone**

Performance

- Normally, we are low usage (but high value)
- Critical bottleneck – MySQL – was obvious
 - Removed very early
- Complete environment cloned using zones
 - And a whole bunch of load generators
- Actual testing almost trivial
 - Plain wget + ab
 - Components + whole stack

Fixes

- 1 page was horrifically expensive
 - >2s to render afresh with each request
 - Generate once, save, serve that
- HTML is extremely verbose
 - mod_gzip made a huge difference
- Design images unoptimized
 - 5x improvement possible
- Timeline sent twice for Mobile
 - Detect before sending the data

Queen Victoria's Journals

"My poor old birthday ever returning without him, who made all so bright"

[Home](#) [About](#) [Timeline](#) [Browse Journals](#) [Illustrations](#) [Search](#) [Resources](#)

A message from HM Queen Elizabeth II

In this the year of my Diamond Jubilee, I am delighted to be able to present, for the first time, the complete on-line collection of Queen Victoria's journals from the Royal Archives.

These diaries cover the period from Queen Victoria's childhood days to her Accession to the Throne, marriage to Prince Albert, and later, her Golden and Diamond Jubilees.

Thirteen volumes in Victoria's own hand survive, and the majority of the remaining volumes were transcribed after Queen Victoria's death by her youngest daughter, Princess Beatrice, on her mother's instructions.

It seems fitting that the subject of the first major public release of material from the Royal Archives is Queen Victoria, who was the first Monarch to celebrate a Diamond Jubilee.

It is hoped that this historic collection will make a valuable addition to the unique material already held by the Bodleian Libraries at Oxford University, and will be used to enhance our knowledge and understanding of the past.

Elizabeth II



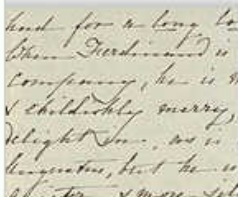
Queen Victoria's coronation portrait. Painted by Franz Xaver Winterhalter 1838
Supplied by Royal Collection Trust / © HM Queen Elizabeth II 2012

Essays: [Queen Victoria's Coronation](#), by **Sir Roy Strong**



From the journals

Browse Journals



Illustrations



Resources



Key events from the timeline

Coronation



Marriage



Diamond Jubilee



Like Tweet 1,242

Share 701

[Diamond Jubilee website](#)

Other tricks - ZFS

- Each zone has its own zfs file system
 - Just 1, so managed as a single unit
 - Automatic snapshots (daily)
 - Compression enabled (great for logs)
- Corrupted zone? Roll back
 - Shut down zone
 - zfs rollback
 - Boot zone

Other Tricks - SMF

- All services managed using SMF
 - Trivial start/stop
 - Clean application environment
 - Dramatic reduction of trivial call-outs
- Developers/app owners have privilege
 - Just to restart, not to modify boot state
- Assign `net_privaddr` if appropriate
 - All applications run on the “official” port
- SMF still has rough edges



Thank You

<http://www.queenvictoriasjournals.org/>