

# *Plone Deployment Patterns*

Plone Konferenz München 2012

Munich

2/22/2012

# Overview

- Technical focused talk
- Explain various deployment configurations
- Benefits and drawbacks found in practice
- Many common sense tips
- Technical overview of common features
- Inspired by JimF and MaurizioD

# Configuration

## Proxy/ZServer

- Most common
- $\leq 2$  threads/process
- # process = # cores
- supervisord useful
- LB is external
- Before/after request requires app hooks

## WSGI

- Not yet common
- # process = # cores
- “advance” mgmt
- mod\_wsgi parent
- LB occurs in Apache
- Before/after request WSGI middleware
- MySQL makes LAMP

# WSGI

## The Bad

- No xmlrpc... yet?
- Reload/restart recycles all clients
- ZPublisher and all existing patches need to be handled
- Not widely deployed?

## The Good

- Widely deployed in Python community
- A “Standard”
- Middleware makes before/after Plone easy
  - enfold.static in SVN
- Less process mgmt

# *Segmenting*

- Use Proxy to identify key for audience
  - Anon, Auth, Spiders, Mobile, etc.
- Varnish identifies due to VCL flexibility
- Downstream proxy or Varnish to route to appropriate “pool”

# *Storage*

- ZODB, R/W and R/O
- FileStorage and NetworkStorage (ZEO)
- relstorage, store ZODB records in RDBMS
- “blob” storage
  - Shared using NFS or CIFS
- demostorage
- zc.beforestorage
- zc.zlibstorage

# Storage II

- ZEO is most common (ships with Plone)
- blob, demo, before, zlibstorage are all “storage wrappers”
  - Will work with ZEO or relstorage
- demostorage vs. R/O
  - demo allows changes but doesn't persist
  - R/O will raise exception; most apps/components WRITE in subtle ways

# *relstorage*



- Popular configuration, tested, hardened
- MySQL, Postgres, & Oracle
- Supports memcached for “shared” cache
- MASTER/SLAVE < 10 minutes (MySQL)
- File system BLOBs optional
  - Makes REPL more awkward
- Only downside is running MySQL server / client libraries – is that a downside?



# *relstorage II*

- Supports “read replicas”
- Separated PRE-PACK and PACK
- PACK operations LOCK tables
- PACK faster, release lock occasionally
- Must be executed on MASTER
  
- Systems administrators can own RDBMS

# *Build*

- Have build machine (stage?)
- Hudson can generate build (ContDeploy)
- .tgz of executed buildout works
- uncompress and run; no need to build
- ZC uses RPM; Experience w/ .deb & RPM
  - <http://package.enfoldsystems.com/docs/>
- Plone will never be in Fedora (sys python)

# *Build II*

- We do not run buildout in PROD
- -o is reasonably safe

## Windows builds

- <https://svn.enfoldsystems.com/trac/public/browser/enfold.recipe.innosetup>
- <http://bit.ly/wD72t2> - Google Docs

Whether it's a .exe, .deb, or .tgz its "frozen"

# *Debugging*

- HAProxy can be parsed to generate what URLs take the longest time
- `collective.stats IN:event.log OUT:CSV` file (also HTTPResponse Headers)
- deadlockdebugger & variants are invaluable
- WSGI middleware could help

# *Application*

- Default Plone Mode
  - Content in root of Plone
  - Maybe content “staged” to a container
- ContentMirror
  - Very similar to how Alfresco works these days. It “publishes” to web application/site.
  - In production. Not widely used due to use case (Plone \*not\* serving requests)
- Static deployment useful but edge case

# *Performance*

- Chameleon 15-30%
- Default Navtree too complex; simplify 15%
- “green bar” actions menu complex; ~ 10%
  - cmsui should address this problem; too many adapter lookups computing menu
- Indexing FTS out-of-process required if uploading multiple or large files

# Monitoring

- munin.zope very cool
- PULL or PARSE
- munin.zope/Nagios/Zenoss “PULLs”
- collective.stats / Big O logger “LOGS”
- “PULL” @ interval or “PARSE” logs
- RAM before/after HTTP operation
- Time in publisher vs. traverse/commit
  - Are you queueing requests? Add more clients

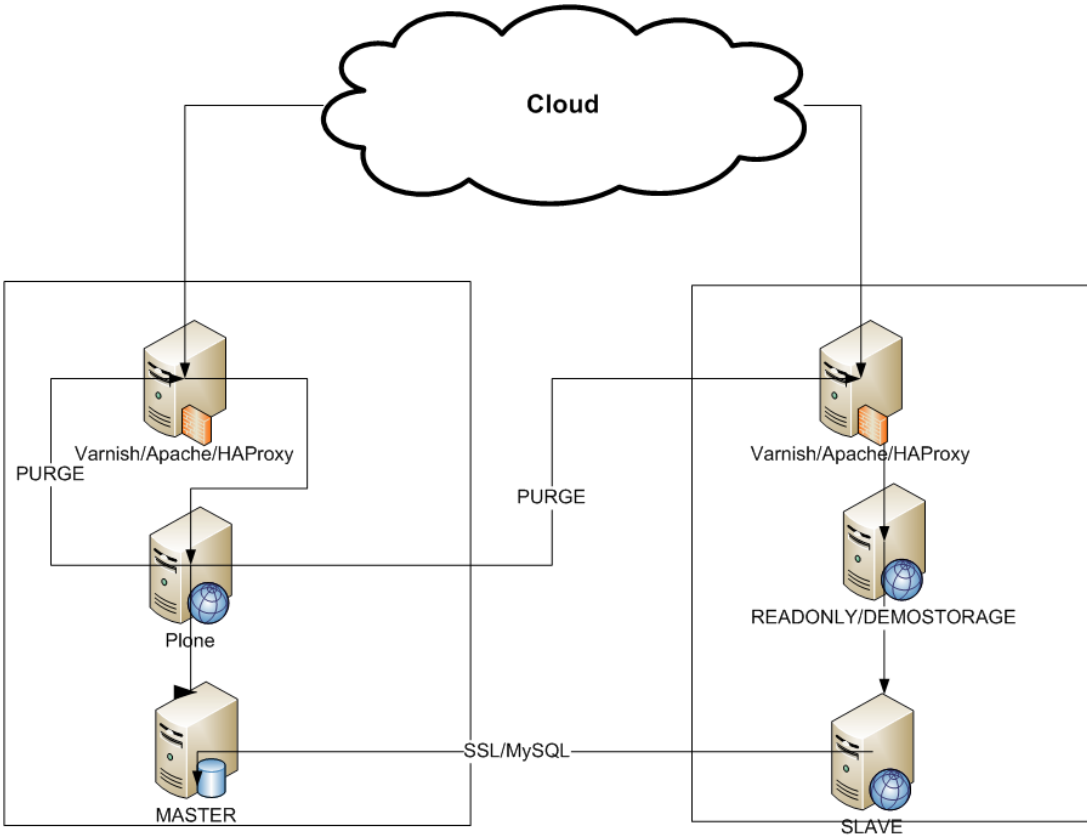
# *Unresolved*

- Content Extraction/Indexing
  - iFilter best text extraction (only Windows)
    - enfold.plone.ifilter in our SVN (fork it!)
  - Apache Tika is fantastic
    - Would be great to have Tika Plone converters
- Messaging / Tasks
  - plone.app.async is “native” but big PRO/CON
  - Carrot / Celery as good as Python gets





# Author vs. Public



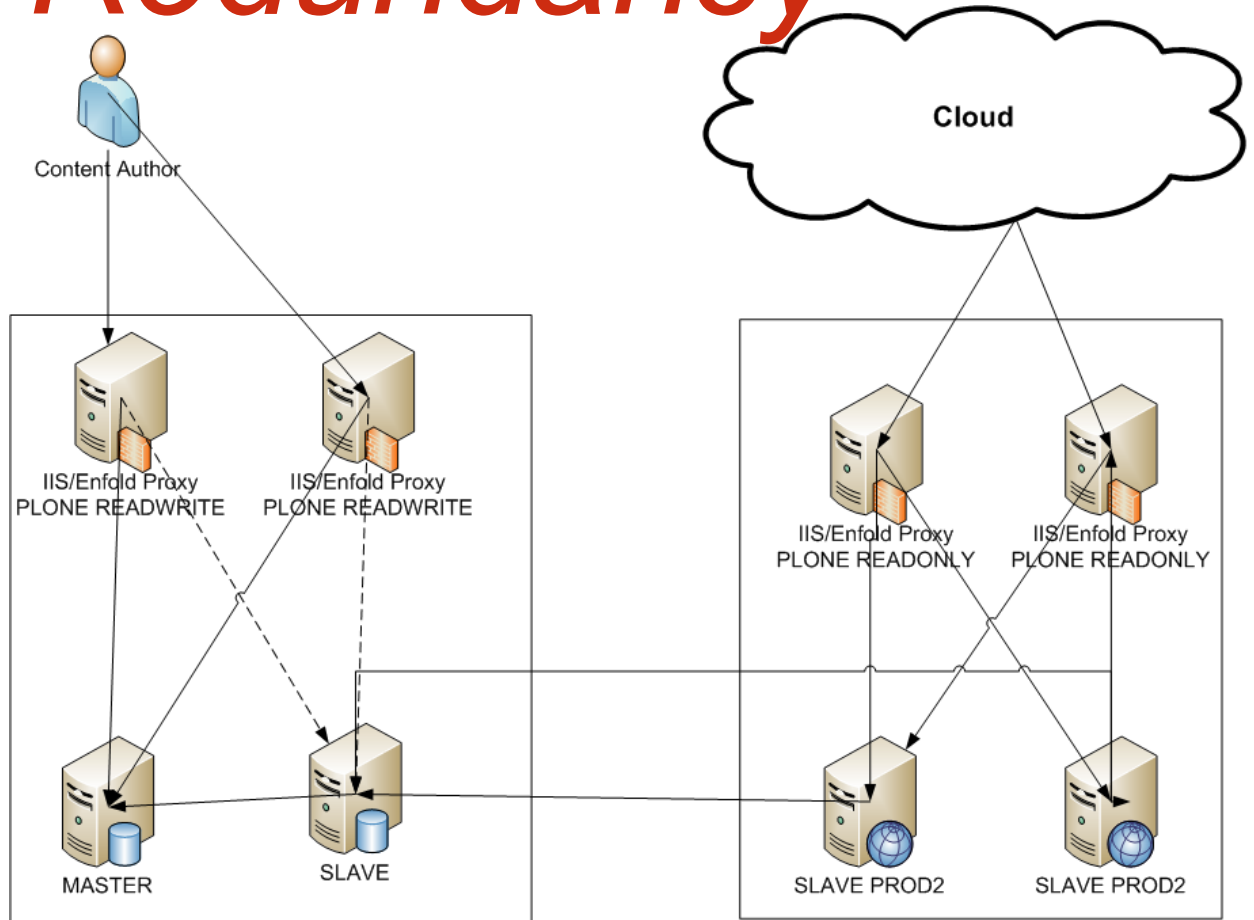
AUTHOR ACCESS;  
WRITES TO MASTER

PUBLIC ACCESS; WRITES  
NOT IN CMS DATABASE

# Redundancy



enfoldsystems



**AUTHOR ACCESS;  
WRITES TO MASTER**

**PUBLIC ACCESS; WRITES  
NOT IN CMS DATABASE**

# *Ploud*



- Varnish/ mod\_wsgi / Apache
  - WSGIPublisher with minor changes
  - ploud.relstorage “multi-tenant” storage in 1 db
    - Postgresql; fills new tables with template data
    - enfold.static
  - Hostname routes to particular client(s)
- Web application (sign up, control panel)
  - Pyramid, Ptah, SQLAlchemy

# *Enfold svn*

- Free Plone site in < 10 seconds
  - <http://ploud.com/>
- Fork any software in our public repo
  - <https://svn.enfoldsystems.com/trac/public>
  - <https://svn.enfoldsystems.com/public>

# Questions

- Any questions/comments send me email
  - [alan@enfoldsystems.com](mailto:alan@enfoldsystems.com)
- I am on #irc
  - runyaga
- Twitter, facebook and all that but email is best.