



Science & Technology  
Facilities Council

# ICAT in Production

**Kevin Phipps (RAL/STFC)**

CLF and ISIS Service Manager  
& ICAT Developer

October 2016

# Overview

- Development/Production differences and similarities
- Splitting the ICAT system up
- Certificates
- Monitoring – is my ICAT happy?
- Informing users
- Off-site access
- Log files



# Development/Production differences

- What is production? Level of service?
- More people involved in production
- Off-site access
- Backups and monitoring required
- Change control
- Keeping users informed
- Scheduling downtimes

This service will be unavailable from ×  
09:00 to 11:00 on Friday 14th  
October due to a software upgrade.  
Apologies for any inconvenience.



# Development/Production similarities

- Make dev/preprod systems copies of production
  - not always possible
  - database and file store size
- Keep data in sync as much as possible
- Be clear which systems are for dev/preprod
  - choose a sensible machine name
  - put a banner/message on dev/preprod systems
- Beware of dev systems being treated as production!



# Splitting the ICAT system up

- For “small” systems all components can be run on one machine
- For anything larger plan to have ICAT and the IDS on separate machines
- ICAT – good connection to database
- IDS – good connection to storage
- Use remote managed ICAT database
- Ingestion also on a separate machine
- TopCAT 2 server lighter



# Certificates

- Avoid self-signed certificates
  - Browser warnings!
  - PKIX exceptions!
  - The Java keystore!
  - Python now stricter!
- Get genuine SSL certificates
- letsencrypt.org, sslforfree.com – a non-profit CA

## Your connection is not secure

The owner of topcat-dev.esc.rl.ac.uk has configured their website improperly. To protect your information from being stolen, Firefox has not connected to this website.

[Learn more...](#)

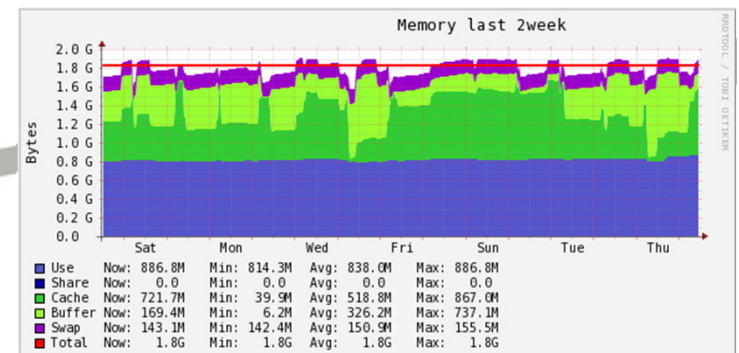
Go Back

Advanced

Report errors like this to help Mozilla identify and block malicious sites

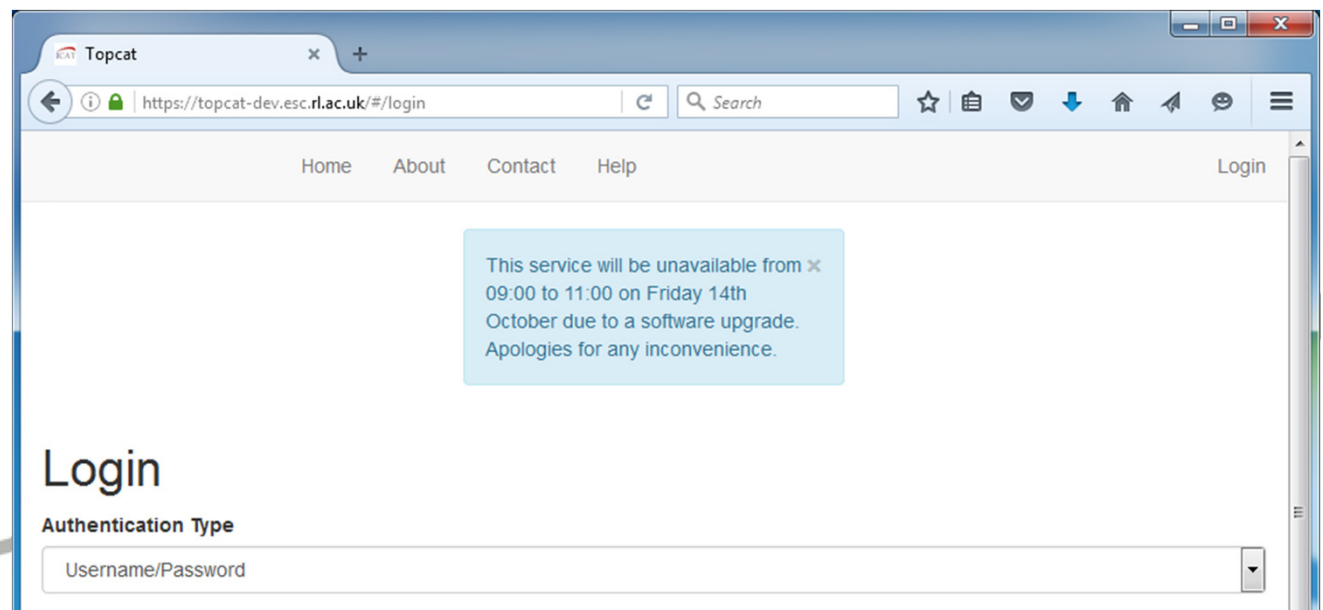
# Monitoring – is my ICAT happy?

- Email alerts essential (eg. nagios)
  - don't rely on users informing you!
  - basic checks – HTTP, disk space etc.
  - GF/ICAT/ingestion – process running, errors
  - `-XX:OnOutOfMemoryError="kill -9 %p"`
  - functional test – can users access it?
- Add status reporting to ingestion process
- Set up system monitoring (eg. ganglia)
  - useful for finding when a problem started
  - spotting patterns eg. backups
- Swap usage kills performance



# Informing users

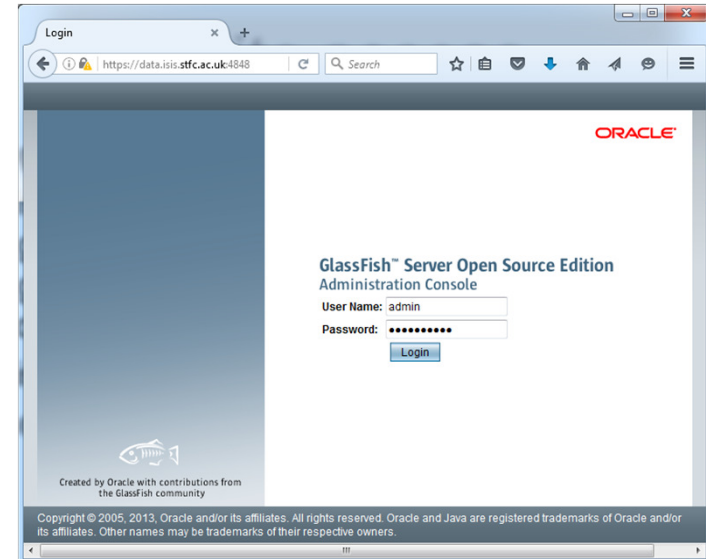
- Advanced warning of downtimes
- Use the TopCAT status message
- Use Apache web server during TopCAT downtimes?
- TopCAT admin page or Dashboard to check for current users?





# Off-site access

- Only open port 443 through site firewall
- Glassfish
  - change Admin password
  - change keystore password
  - restrict port 4848 access
- Be more vigilant with patching!
- Organisation branded/style TopCAT?



# Log files

- Set up a log partition
- Don't set the logging level too high
- Zip up non-current log files
- Useful information in there somewhere
- Difficult to follow because of threads
- Advanced grepping and regular expressions required!



# Conclusions

- Prepare for increased scrutiny of your service
- Keep dev and production as similar as possible
- Split components up for higher loads
- Use genuine SSL certificates
- Consider how your actions affect users
- Keep externally available services secure
- Only use the log files when necessary
- If in doubt, try restarting Glassfish!





Science & Technology  
Facilities Council

**Thanks for listening.**

**Questions ?**