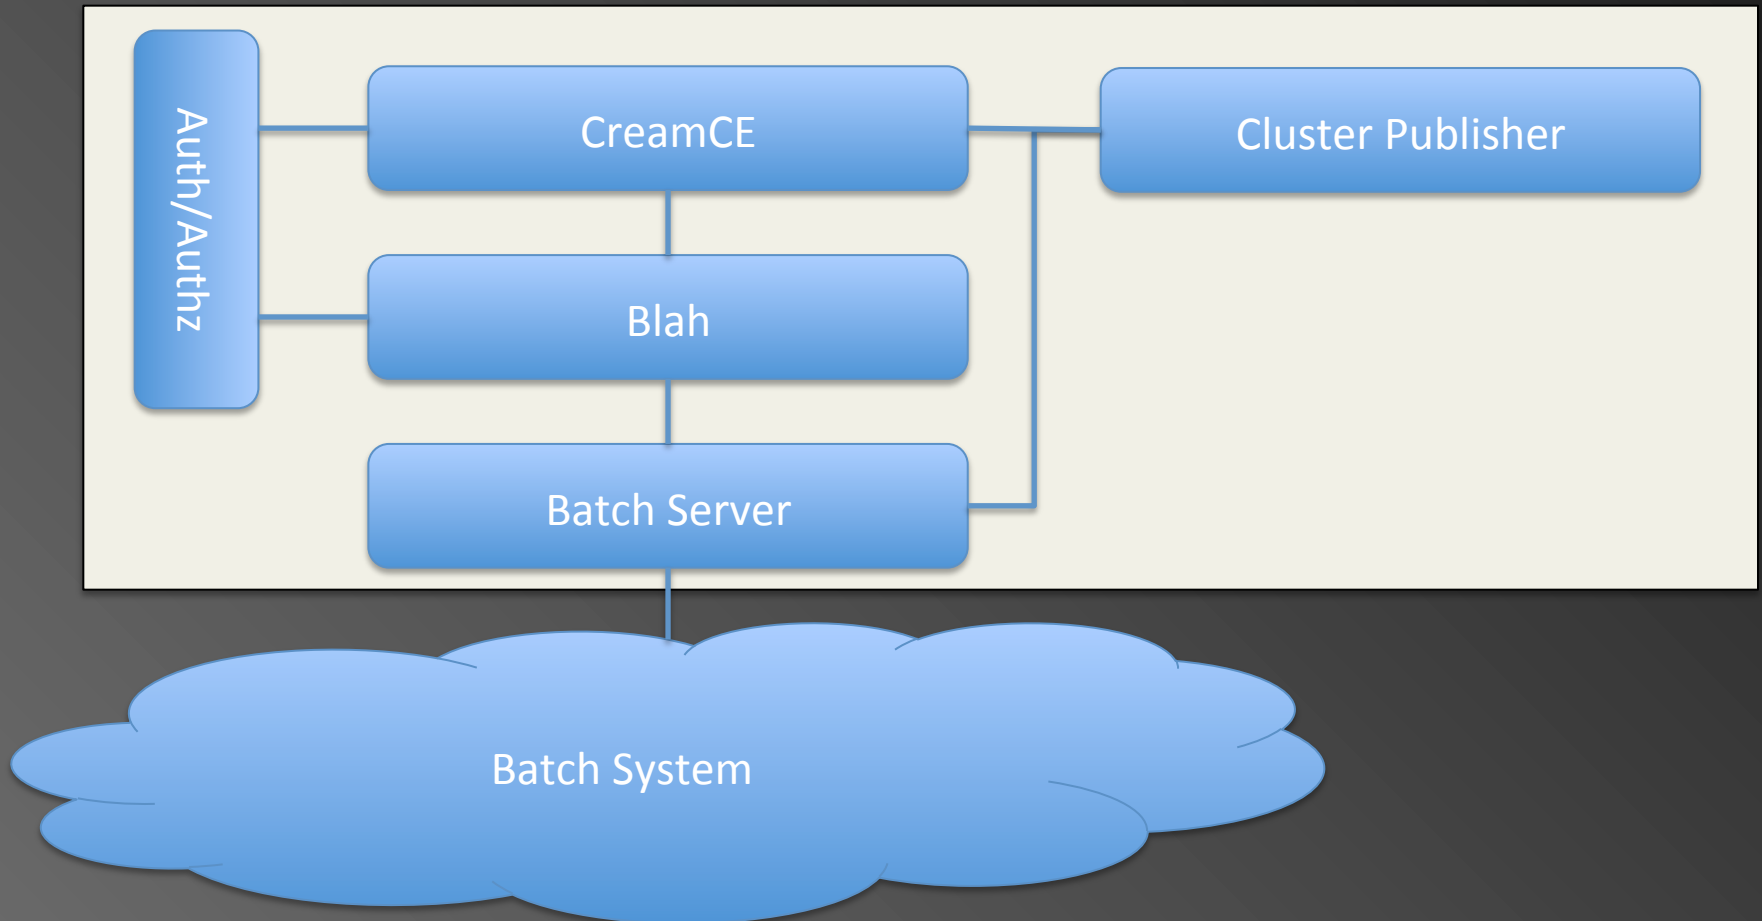


UMD CreamCE Experiences at RALPP

Chris Brew
Rob Harper

CE Components



Timeline

- 03/08 – heplnx206 reinstalled as CreamCE
- 23/08 – heplnx207 starts draining
- 24/08 – heplnx206 started crashing under load
- 24/08 – We start looking for CreamCE tuning
- 01/09 – heplnx207 returns as CreamCE
- 02/09 – heplnx207 starts having load issues
- Early Sept – Apply every CreamCE tweak we can
- Mid Sept – CEs become stable under load but get repeatedly blacklisted
- End Sept/Early Oct – Migrate to new batch server
deploy heplnx208 as third CreamCE – same issues

Database tweaks

- CreamCE/YAIM does not configure the my.cnf file in any way
- innodb_buffer_pool_size needs increasing for a production CreamCE
- We also ran tunemysql.pl and ended up with:

```
skip-bdb  
thread_cache_size=8  
query_cache_size=512M  
query_cache_type=1  
query_cache_limit=64M  
innodb_buffer_pool_size=1024M  
innodb_open_files=2710  
innodb_max_dirty_pages_pct=10  
innodb_flush_method=O_DIRECT  
table_cache=1200  
tmp_table_size=256M  
max_heap_table_size=256M
```

Other Tweaks

- Reduce the Blah purge interval
- Increase the tomcat process memory limits
- Increase the system nofile ulimit for root, tomcat and mysql
- Increase the limit on the number of gridftp connections
- Put `/var/log/cream` and `/var/cream/sandbox` on their own (largish) partitions

blacklisting

- CreamCEs would get randomly blacklisted by CERN WMSes
 - Docs talk of 30min blacklisting but seemed to continue until we rebooted the node
 - Restarting Cream components did not seem to fix the problem
 - Did not appear to be load related
- Appears to have been fixed by adding wildcard VOMS groups to the Argus server

CreamCE and Argus

- CreamCE can use Argus for Authz
- YAIM mostly configures this
 - AFAICT It configures the tomcat/CreamCE part to use Argus but the Blah/glexec part uses the gridmapdir
 - Installing the Argus lcms plugins and copying the config from a glexec WN appears to work.

Cluster Publisher

- Finally separates the cluster from CE
- Only need to publish worker nodes in one place
- Also handles software tags
 - Runs GridFTP server, needs a certificate
 - Hackable to use Argus as Authz

Other notes

- UMD doesn't chkconfig the BDII on

What We've ended up with

