



CMS FroNTier Report

Lee Lueking, Barry Blumenfeld, Dave Dykstra,
Eric Wicklund

CERN 3D Workshop
Sept. 13-14, 2006



Contents

- Recent software updates
- Deployment details and status
- Testing
- Monitoring and Operations
- Conclusion

Additional information on the Wiki:

<https://uimon.cern.ch/twiki/bin/view/CMS/DatabaseAccess>



What's new in FroNTier lately

- Client can request the data be zipped by the server (compression levels 0-9)
- Keep alive signals sent to client when database is busy, avoids timeouts
- Ported to 64-bit Linux
- Parameters can come in long parenthesized connect string instead of environment vars
- Can define logical name in long string so pool file catalog can use short name



FroNTier connect string example

```
frontier://(logicalserverurl=http://cms_conditions_data)
(proxyurl=http://cmsrv11.fnal.gov:3128)
(proxyurl=http://cmsrv11.fnal.gov:3128)
(serverurl=http://cmsfrontier.cern.ch:8000/Frontier)
(serverurl=http://cmsfrontier1.cern.ch:8000/Frontier)
(serverurl=http://cmsfrontier2.cern.ch:8000/Frontier)
(serverurl=http://cmsfrontier3.cern.ch:8000/Frontier)
(retrieve-ziplevel=5)/CMS_COND_ECAL
```

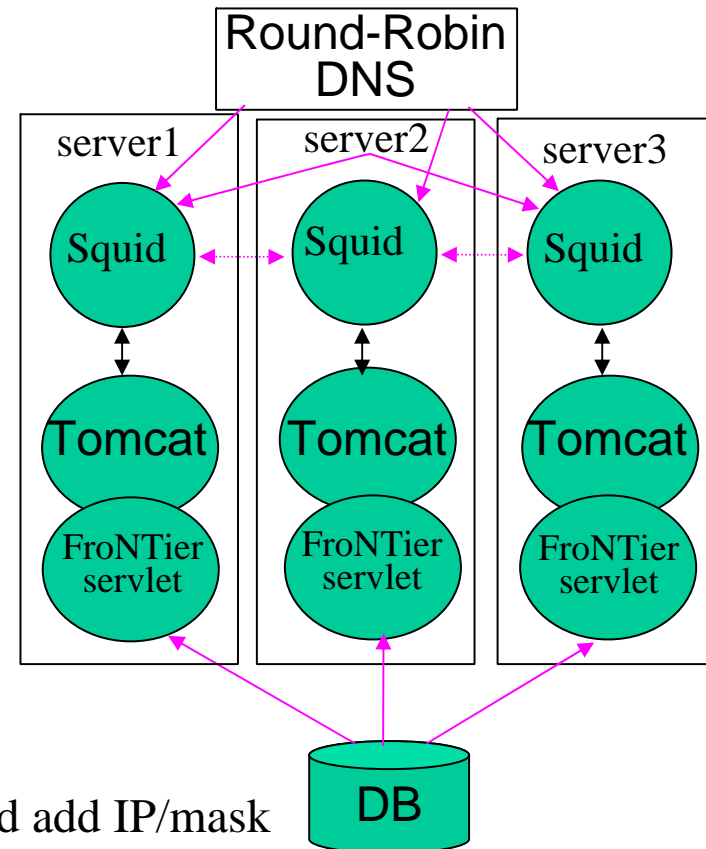
POOL file catalog:

```
frontier://cms_conditions_data/CMS_COND_ECAL
```



FroNTier “Launchpad” software

- Squid caching proxy
 - Load shared with Round-Robin DNS
 - Configured in “accelerator mode”
 - Peer-to-peer caching
 - “Wide open frontier”*
- Tomcat - standard
- FroNTier servlet
 - Distributed as “war” file
 - Unpack in Tomcat webapps dir
 - Change 2 files if name is different
 - One xml file describes DB connection



*In the past, we required the registration so we could add IP/mask to our Access Control List (ACL) at CERN. Recently decided to run in “wide-open” mode so installations can be tested w/o registration.



“Site” Squid details

- Hardware requirements
 - Minimum specs: 1GHz CPU, 1GByte mem, GBit network, 100 GB disk.
 - Needs to be well connected (network-wise) to worker nodes, and have access to WAN and LAN if on Private Network.
 - Having 2 machines for failover is a requirement for T-0/T-1, and a useful option for T-2. Inexpensive insurance for reliability.
- Software installation
 - Squid server and configuration
 - Site-local-config file All Details:
<https://twiki.cern.ch/twiki/bin/view/CMS/CMSSquidDeployment>



Testing Launchpad Failover

Server	Test 1	Test 2	Test 3	Test 4	Test 5
cmsfrontier1	UP	DOWN	DOWN	DOWN	UP
cmsfrontier2	DOWN	UP	DOWN	DOWN	UP
cmsfrontier3	DOWN	DOWN	UP	DOWN	UP
Success rate (10 tries)	100%	100%	100%	0%	100%



Vertical Slice Testing

- The client software is the CMS framework used for most processing (CMSSW/cmsRun). The versions used were CMSSW_0_8_1, CORAL_1_5_1, frontier_client_2.4.5_cms
- The servlet version was frontier servlet 3.2
- Data was in the development DB (devdb10)
- ECAL calibration was used. Each “object” is ~2.8 MBytes in size.
- 8 clients were run simultaneously at each site. The squid cache was pre-loaded prior to the testing.
- In the Grid environment, it is not possible to control variables like the performance of the client node and how many things the client node is doing during the test. Therefore, results are quite variable.



Testing Sites

- ASGC
- Belgium
- CALTECH
- CERN
- CIEMAT
- CSCS
- DESY
- Estonia
- Florida
- FNAL
- GRIDKA
- Legnaro
- MIT
- Nebraska
- PIC
- Pisa
- Purdue
- RAL
- RWTH
- UCSD
- Wisconsin
- 21 sites were successfully tested
- 4 additional sites had various problems w/ squid and/or software.
- Additional 10 sites identified for install soon.



Testing Results (some examples)

Zipping turned off

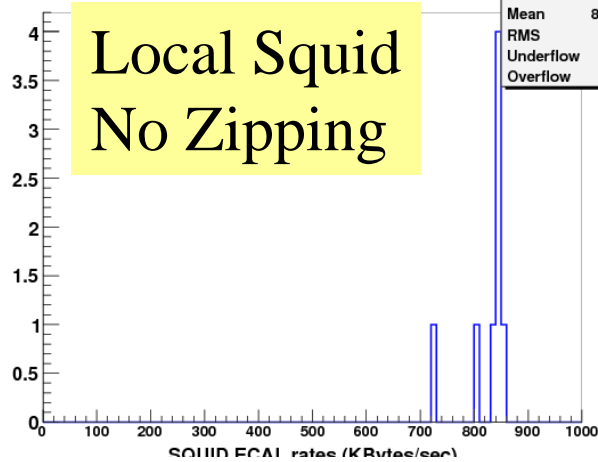
Site	Local Squid kBytes/sec	Frontier@CERN kBytes/sec
DESY	708+/-107	11.4+/-0.2
Estonia	152+/-63	6.3+/-0.8
Florida	612+/-13	40.3+/-0.4
FNAL	482+/-116	18+/-1.4
GridKa	565+/-151	23.1+/-2.3
Performance ranges	110 (Legnaro) to 826 (Belgium)	5.5 (ASGC) to 99 (Belgium)

<https://twiki.cern.ch/twiki/bin/view/CMS/CMS-Frontier-test-integration>



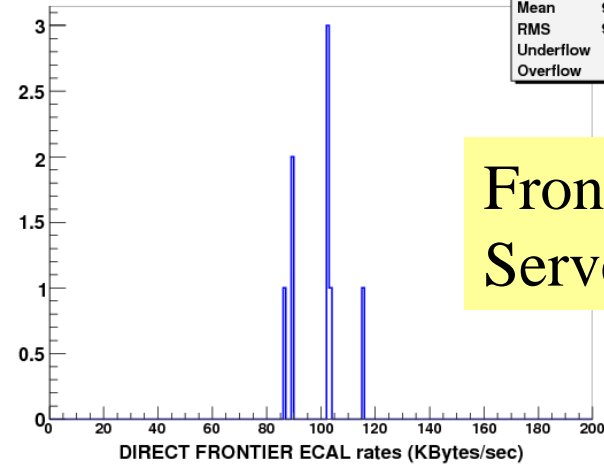
Results at “Representative” Site: Belgium

SQUID ECAL rates



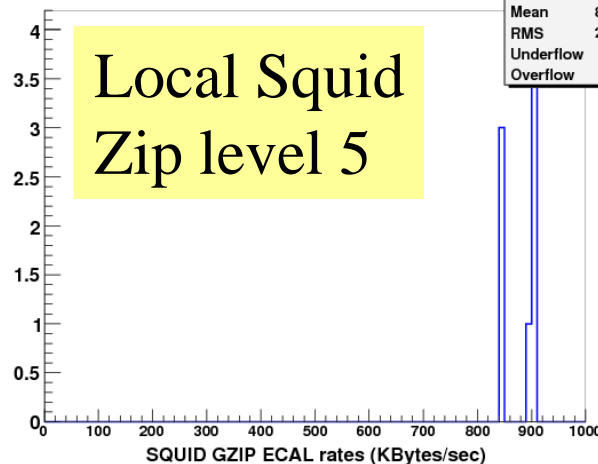
Local Squid
No Zipping

DIRECT FRONTIER ECAL rates



Frontier
Server @ CERN

SQUID GZIP ECAL rates




Local Squid
Zip level 5

- Local squid provides factor of 8 improvement in throughput.
- Zipping seems to help, but will help more for larger number of simultaneous clients.




What's Next?

- Hundreds of simultaneous clients at one site, and across multiple sites.
- Explore regions of operation for which zipping is important.
- Production usage in Computing, Software and Analysis (CSA06) to deliver Calibration and Alignment for Monte Carlo studies. (starting end of September 06)




Squid Monitoring

- All the squids are monitored using the industry standard package MRTG
- This package talks to each squid in turn via SNMP every 5 minutes
- MRTG creates web accessible histograms which can be viewed from anywhere
- In our implementation, MRTG produces histograms of:
 - Requests/Fetches
 - Number of Bytes In/Out
 - Number of Objects in the Cache
- The monitoring computer is a temporary installation at Fermilab <http://cdfdbfrontier4.fnal.gov:8888/indexcms>
- There is also a very simple test script that looks at the histograms every half hour and if the squid is not responding sends e-mail to the cms-frontier-support@cern.ch mailing list



Monitoring Examples


'Daily' Graph (5 Minute Average)



Max HTTP reqs 6924.0 req/min Average HTTP reqs 772.0 req/min Current HTTP reqs 4911.0 req/min
Max HTTP fetches 365.0 req/min Average HTTP fetches 105.0 req/min Current HTTP fetches 635.0 req/min

The green histogram shows the number of requests to the squid. The blue line shows how often the request was not in the cache.

'Daily' Graph (5 Minute Average)



Max Total 16.9 MB/s Average Total 2500.0 kB/s Current Total 11.1 MB/s
Max Fetches 1529.0 kB/s Average Fetches 237.0 kB/s Current Fetches 948.0 kB/s

The green histogram shows how many bytes were delivered by the squid. The blue line shows how many bytes had to be retrieved from the source.



Conclusion

- Several important additions have been made to the FroNTier software.
- We installed squids and tested FroNTier data access at 21 sites; results meet our expectations.
- Additional testing is ongoing to further push the performance limits for 100's of concurrent clients.
- A straightforward monitoring system is in place.
- The system will be used for production in CSA06 starting later in September.



Finish