

ATLAS Frontier Update

Distributed Database Operations Workshop

16 November 2010

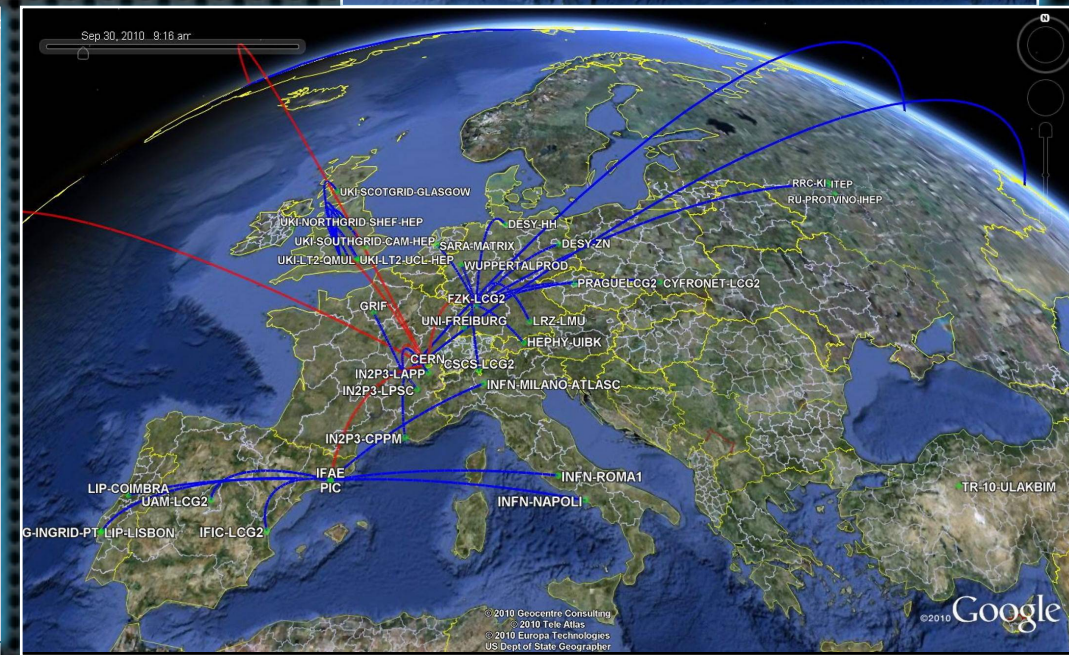
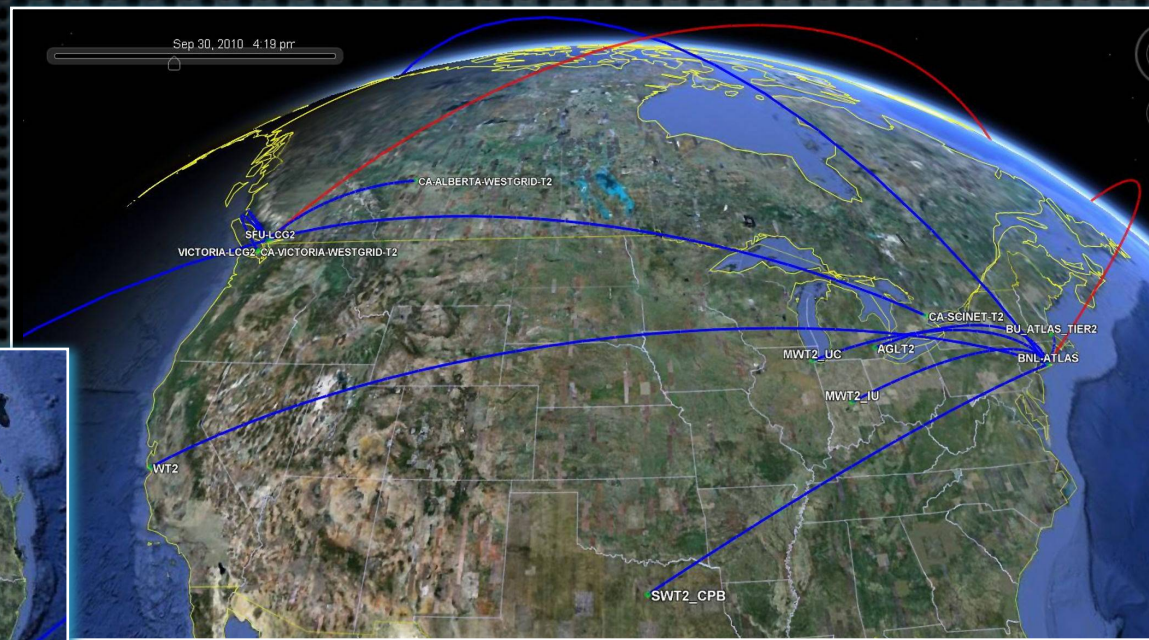

John S. DeStefano Jr.

Brookhaven National Laboratory

jd@bnl.gov

Current deployment

- Frontier servers at seven ATLAS sites:
 - BNL, CERN, IN2P3, KIT, PIC, RAL, TRIUMF
- Squid cache servers at Tier-N sites around the world




Images by Google Earth, from CHEP poster "Frontier in ATLAS", organized for ATLAS by Douglas Smith.

Current deployment

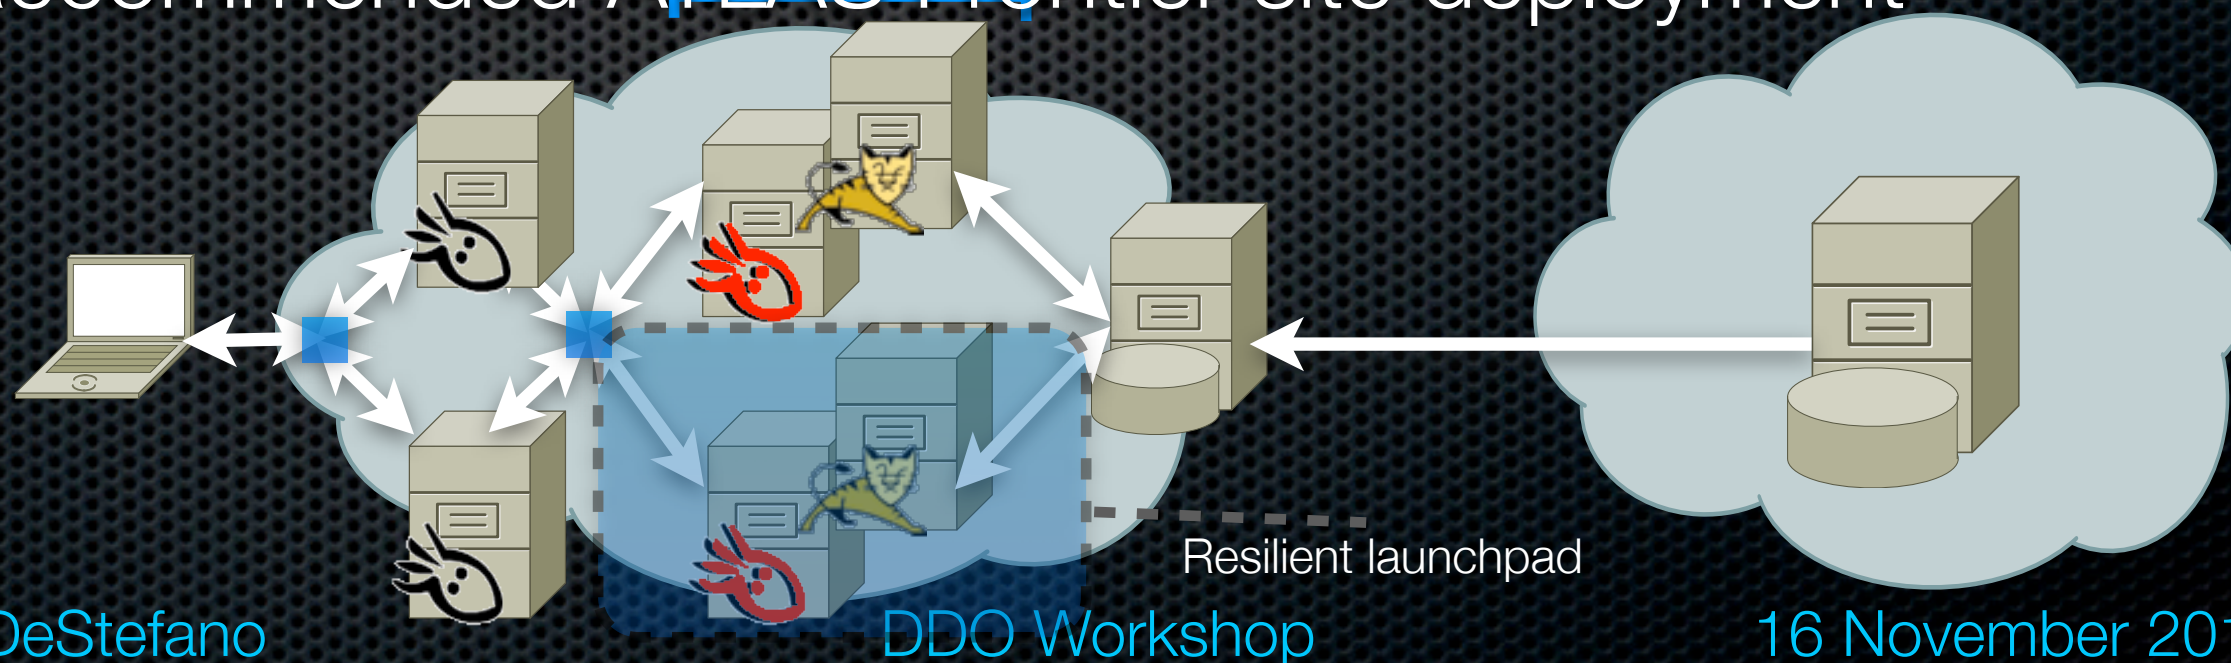
- ✦ Most sites comprised of a single Frontier “launchpad”
 - ✦ Single Tomcat plus single accelerator Squid
- ✦ Many sites do not offer local jobs a local Squid cache

Site deployment details

- Common ATLAS Frontier site deployment



- Recommended ATLAS Frontier site deployment



J. DeStefano

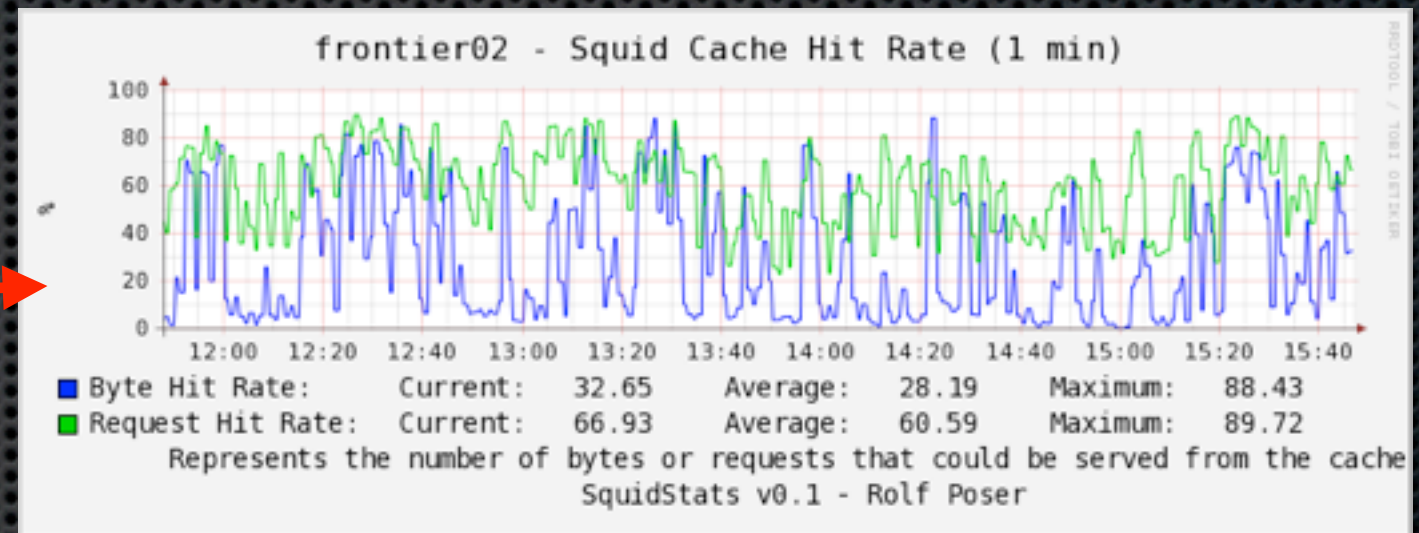
DDO Workshop

16 November 2010


Site deployment details

- Does a separate site Squid cache really make a difference?

BNL Launchpad Squid →



BNL Site Cache Squid →

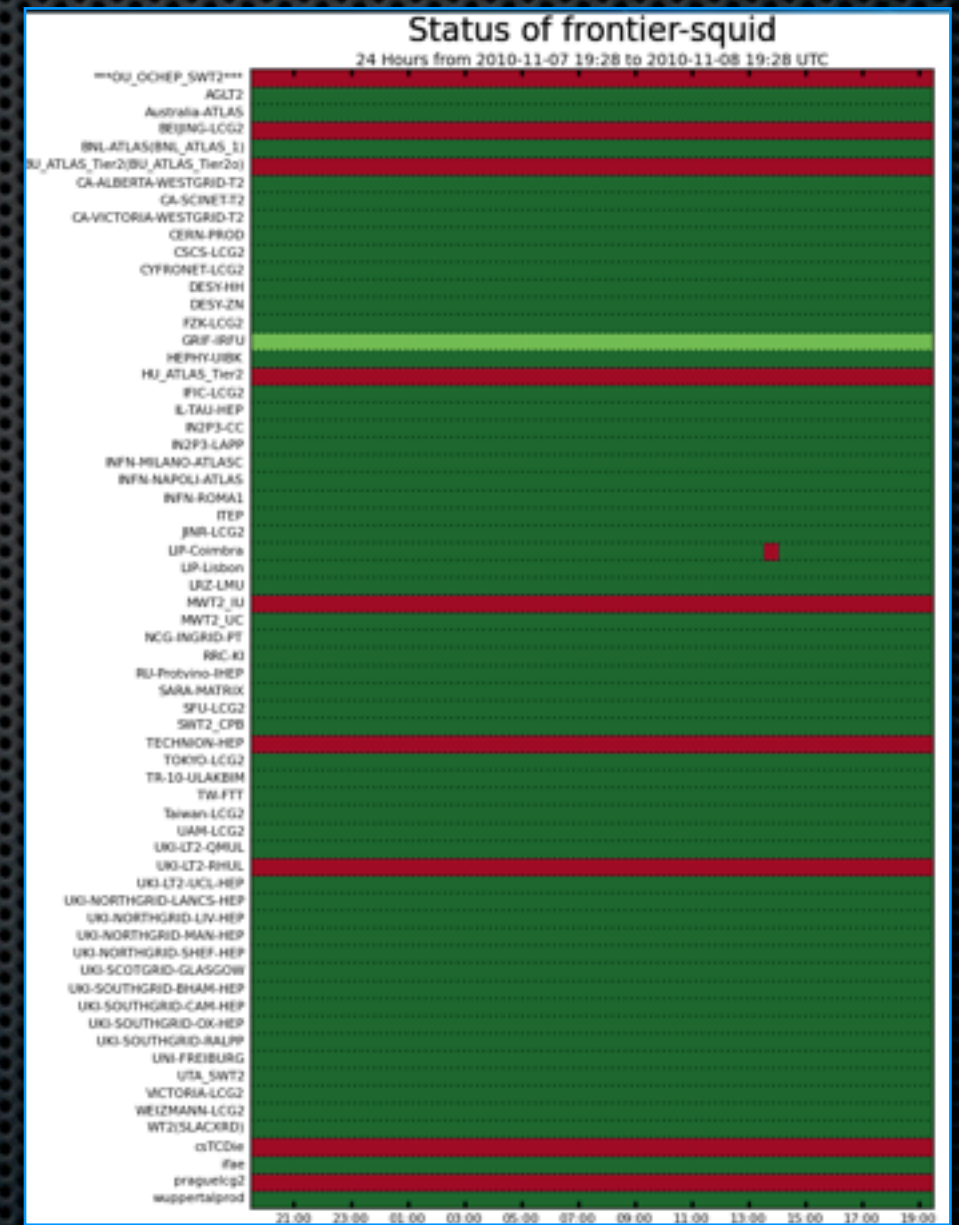
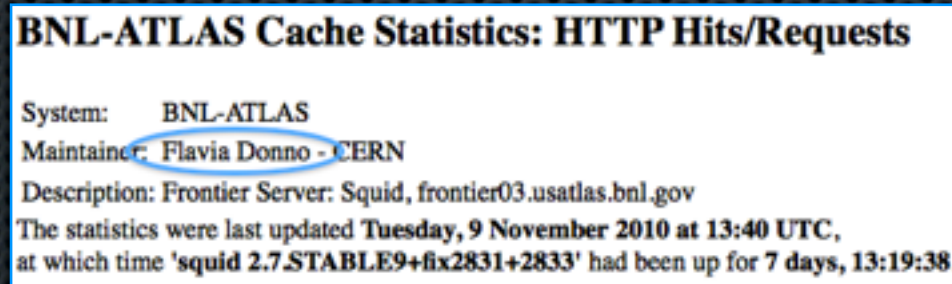


Current deployment

- ✦ Site reports:
 - ✦ BNL: recent site power failure caused brief downtime
 - ✦ KIT: recent Frontier upgrade and reconfiguration
 - ✦ RAL: Frontier working well despite unforeseen traffic and load
 - ✦ TRIUMF: Frontier working well

Production updates

Squid site cache status added to ATLAS dashboard



• <http://dashb-atlas-ssb.cern.ch/dashboard/request.py/siteviewhistory?columnid=47>

Production updates

Frontier site status added to CERN Service Level Status dashboard


Additional information

full name: **ATLAS Frontier Service**
short name: ATLAS-Frontier
group: IT-ES-VOS
site: CERN

email: **Flavia.Donno@cern.ch**
web site: <http://frontier.cern...>

manager: **Flavia Donno**








ATLAS-Frontier

availability: 
(more)

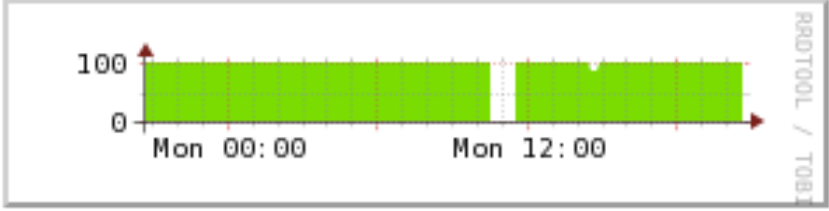
percentage: 100%

status: **available**

this service consists of:

-  ATLAS-CERN-Frontier
-  ATLAS-BNL-Frontier
-  ATLAS-KIT-Frontier
-  ATLAS-LYON-Frontier
-  ATLAS-PIC-Frontier
-  ATLAS-RAL-Frontier
-  ATLAS-TRIUMF-Frontier

availability in the last 24 hours (more):



100
0
Mon 00:00 Mon 12:00

✦ <http://sls.cern.ch/sls/service.php?id=ATLAS-Frontier>

Production updates

Additional ATLAS Frontier monitoring (currently problematic):

- ✦ AWStats for Frontier servers:
<http://frontier.cern.ch/awstatsatlas.html>
- ✦ MRTG for Squids:
<http://frontier.cern.ch/squidstats/indexatlas.html>
- ✦ SAM:
<https://lcg-sam.cern.ch:8443/sam/sam.py>

Production updates

ATLAS RPM packages for Squid, Tomcat, Frontier, and AWStats now available

- ✦ <https://twiki.cern.ch/twiki/bin/view/LCG/WLCGFrontier>
- ✦ <https://twiki.cern.ch/twiki/bin/view/PDBService/SquidRPMsTier1andTier2>

Current available rpms

- <http://grid-deployment.web.cern.ch/grid-deployment/flavia/frontier-awstats-1.0-1.noarch.rpm>
- <http://grid-deployment.web.cern.ch/grid-deployment/flavia/frontier-tomcat-1.0-1.noarch.rpm>
- <http://grid-deployment.web.cern.ch/grid-deployment/flavia/frontier-servlet-1.0-3.24.0.noarch.rpm>

Frontier Savannah reports

- ✦ 72011,73387,73405 (closed): Problems with some Frontier site configurations
- ✦ 73387 (closed): Some sites missing local Squid cache
- ✦ 74441 (closed): Frontier used for reconstruction jobs instead of analysis

Other problems and solutions

- ✦ Site installation and configuration of cache revalidation scripts
 - ✦ Ensure ATLAS_FRONTIER_READER database schema and triggers are installed
 - ✦ https://twiki.cern.ch/twiki/bin/viewauth/Atlas/ScriptsDepotContent#Script_for_creating_DB_user_sche
 - ✦ https://twiki.cern.ch/twiki/bin/viewauth/Atlas/ScriptsDepotContent#Script_for_creating_FRONTIER_rea
 - ✦ Ensure Frontier servlet is properly configured

Other problems and solutions

- Database connection time-outs and disconnects during long-running queries
 - MaxDbExecuteSeconds=60

```
ServerName=frontieratbnl
DataSourceName=jdbc/frontieratbnl
XsdTableName=frontier_descriptors
UseFdoCache=no
ShortCacheExpireSeconds=3600
LongCacheExpireSeconds=3600
LastModifiedTableName=ATLAS_FRONTIER_TRACKMOD.LAST_MODIFIED
ValidateLastModifiedSeconds=300
MaxDbExecuteSeconds=60
```

<https://www.racf.bnl.gov/docs/services/frontier/servlet>

- Athena default time-out (3600 sec) can be disabled with `svcMgr.PoolSvc.ConnectionRetrialTimeOut=0` (R. Hawkings)

Other problems and solutions

- ✦ Errors on missing database tables
 - ✦ Connect with ATLAS_FRONTIER_READER
 - ✦ Enable ATLAS_FRONTIER_TRACKMOD tables

```
driverClassName="oracle.jdbc.driver.OracleDriver"  
maxIdle="2" maxActive="2" maxWait="-1" logAbandoned="true"  
minEvictableIdleTimeMillis="60000" timeBetweenEvictionRunsMillis="20000"  
username="ATLAS_FRONTIER_READER" password="ASecretLocalPassword"
```

```
ServerName=frontieratbnl  
DataSourceName=jdbc/frontieratbnl  
XsdTableName=frontier_descriptors  
UseFdoCache=no  
ShortCacheExpireSeconds=3600  
LongCacheExpireSeconds=3600  
LastModifiedTableName=ATLAS_FRONTIER_TRACKMOD.LAST_MODIFIED  
ValidateLastModifiedSeconds=300
```

<https://www.racf.bnl.gov/docs/services/frontier/servlet>

ATLAS recommendations

- ✦ Squid:
 - ✦ Frontier server sites should maintain at least one site Squid (two for > 500 analysis slots) to offset load on Frontier servers, and for reliability (when server/database is down)
- ✦ Frontier (caveat: ongoing discussion by ATLAS management):
 - ✦ Choices:
 - ✦ Leave existing Frontier sites to operate as they are
 - ✦ Minimize and fortify a subset of existing Frontier server sites

Ongoing issues and discussion

- ✦ Current fail-over scheme: “ordered chaos”
 - ✦ ToA does not currently permit multiple server entries for local cache server definitions

```
fsquidmap={  
# CERN  
'CERN-PROD': {'mysquid': 'http://atlast0fsquid.cern.ch:3128' 'frontiers': [fcer, flyo], 'squids': ['local', '']},  
# CERN-PROD: {'mysquid': 'http://atlast0fsquid.cern.ch:3128' 'frontiers': [fcer, flyo], 'squids': ['local', '']}
```


Ongoing issues and discussion

- ✦ Current fail-over scheme: “ordered chaos”
 - ✦ Frontier server Squid ACLs are not consistent or permissive enough at all T1 Frontier sites to support fail-over from other sites
 - ✦ Tier 2 Squids have fail-over pairs
 - ✦ Fail-overs are not currently tested sufficiently

Ongoing issues and discussion

- Frontier T1 access tests from BNL worker nodes:

Site	Cache	Frontier	Off-site
BNL	OK	OK	OK
CERN	refused	OK	OK
KIT	time-out	OK	forbidden
LYON	refused	OK	OK
PIC	no cache	OK	OK
RAL	no cache	OK	OK
TRIUMF	no cache	OK	time-out

Future

- ✦ Solidify ATLAS Frontier deployment model
- ✦ Modify ToA to permit multiple Frontier server and Squid cache primary and backup entries
 - ✦ Rod, Alessandro working on improvements (^mysquids` multi-element definition to compliment default `mysquid`, etc.)
- ✦ Fix and improve monitoring, including fail-over combinations
- ✦ Implement separate Squid instances for CMS, ATLAS

Documentation

- ✦ <http://frontier.cern.ch>
- ✦ <https://www.racf.bnl.gov/docs/services/frontier>

Miscellaneous

- ✦ “Frontier in ATLAS” poster at CHEP
 - ✦ Organized and presented by D. Smith
- ✦ Frontier weekly meetings: Wed. 4:30 PM CET
 - ✦ Agenda rotates bi-weekly: site admins/project
 - ✦ <https://www.racf.bnl.gov/docs/services/frontier/meetings/>

Acknowledgements

- ✦ Dario Barberis
- ✦ Alessandro De Salvo
- ✦ Alastair Dewhurst
- ✦ Gancho Dimitrov
- ✦ Flavia Donno
- ✦ Dave Dykstra
- ✦ David Front
- ✦ Elizabeth Gallas
- ✦ Fred Luehring
- ✦ Shawn McKee
- ✦ Douglas Smith
- ✦ Rod Walker
- ✦ Andrew Wong
- ✦ Many others