



<http://atlas.ch>



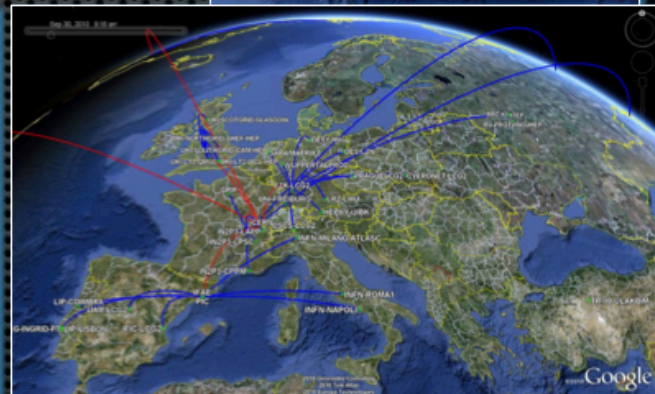

ATLAS Frontier Update

ATLAS Software & Computing Week
April 2011
John DeStefano
Brookhaven National Laboratory

Summary (from ASC 12/2010)

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.


Summary (from ASC 12/2010)

Site deployment details

- Common ATLAS Frontier site deployment



- Recommended ATLAS Frontier site deployment



Site Information

- “Resilient” Frontier server sites with redundant Frontier instances and independent local site cache node(s): BNL, CERN, RAL
- Some sites with single Frontier instance
- Some sites without any local site Squids

Why Are Site Caches Needed?

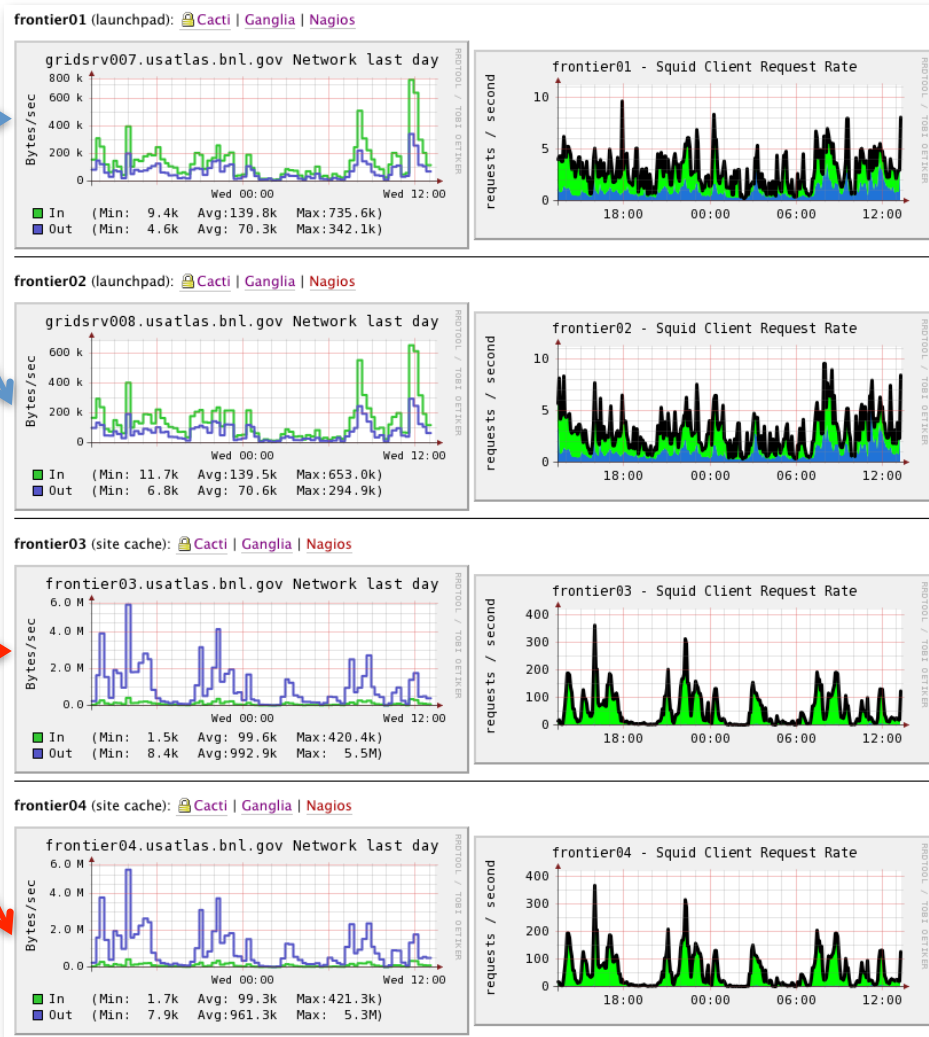
(don't Frontier launchpads have a Squid, too?)

- Redundancy and resiliency for Frontier launchpads
- Protect launchpads by reducing request load on launchpad Squids by multiple factors
- Safeguard local jobs against known problems at Frontier launchpad sites without cache nodes
- ATLAS recommendation for *all* sites: two Squid caches per $500 < 2000$ user analysis job slots; one cache per each 1000 slots above 2000

Why Are Site Caches Needed?

Frontier launchpads

Max Output:
375 KB/sec



Site caches

Max Output:
6 MB/sec

Max Requests:
10/sec

Max Requests:
375/sec

Site Reports

- BNL: Running well; services up-to-date; plans to upgrade hardware in 2011
- CERN: Local configuration problems; out-of-date software; no local resource to maintain/monitor site
 - Currently being trouble-shot and remedied from off-site
- KIT: Running well; services up-to-date
- LYON: (Unconfirmed) Services behind a few versions
- PIC: Running well; services up-to-date; no local site cache
- RAL: Running well; services up-to-date
- TRIUMF: Running well; services up-to-date; no local site cache

Monitoring: AWStats


- <http://frontier.cern.ch/awstatsatlas.html>
- Service provided as courtesy by CMS
- Not working completely; not all sites are reporting

Squid Cache Hits				
Access Status - Full list	Pages	Hits	Bandwidth	Last visit
TCP_REFRESH_HIT:FIRST_UP_PARENT	180243	180243	773.39 MB	04 Apr 2011 - 08:01
TCP_IMS_HIT:NONE	127297	127297	0	04 Apr 2011 - 08:00
TCP_REFRESH_MISS:FIRST_UP_PARENT	87998	87998	3.57 GB	04 Apr 2011 - 08:00
TCP_DENIED:NONE	57315	57315	78.22 MB	04 Apr 2011 - 08:01
TCP_MEM_HIT:NONE	19440	19440	199.31 MB	04 Apr 2011 - 08:00
TCP_MISS:FIRST_UP_PARENT	3965	3968	252.98 MB	04 Apr 2011 - 07:56
TCP_HIT:NONE	733	733	290.06 MB	04 Apr 2011 - 07:51
TCP_REFRESH_HIT:NONE	177	177	142.15 KB	04 Apr 2011 - 04:35
TCP_REFRESH_MISS:NONE	72	72	2.71 MB	04 Apr 2011 - 04:35
TCP_CLIENT_REFRESH_MISS:FIRST_UP_PARENT	58	58	454.95 KB	04 Apr 2011 - 04:53
TCP_IMS_HIT:FIRST_UP_PARENT	17	17	0	03 Apr 2011 - 13:11
TCP_SWAPFAIL_MISS:FIRST_UP_PARENT	1	1	18.45 KB	03 Apr 2011 - 12:55
Total	477316	477319	5.13 GB	

Monitoring: MRTG

- <http://frontier.cern.ch/squidstats/indexatlas.html>
- Service provided as courtesy by CMS
- Not working completely, but updates have been made; most sites now reporting most statistics

`Daily' Graph (5 Minute Average)



	Max	Average	Current
HTTP reqs	378.0 req/min	93.0 req/min	204.0 req/min
HTTP fetches	257.0 req/min	52.0 req/min	123.0 req/min

Monitoring: SAM

- <https://lcg-sam.cern.ch:8443/sam/sam.py>
- Tests most fail-over scenarios for EGEE, but does not work for OSG sites
- SAM framework recently discontinued in favor of Nagios-based testing
 - Possibility to reactivate manual, daily tests while waiting for port to Nagios


No	RegionName	SiteName	NodeName	Status	atlas		
					None	None	None
1	UKI	RAL-LCG2	lcgce06.gridpp.rl.ac.uk	OK	ok	ok	ok
2	UKI	UKI-LT2-Brunel	dc2-grid-65.brunel.ac.uk	OK	ok	ok	ok
3	UKI	UKI-LT2-Brunel	dgc-grid-40.brunel.ac.uk	OK	ok	ok	ok
4	UKI	UKI-LT2-Brunel	dgc-grid-44.brunel.ac.uk	OK	ok	ok	ok
5	UKI	UKI-LT2-QMUL	ce02.esc.qmul.ac.uk	MAINT	na	error	na
6	UKI	UKI-LT2-RHUL	ce2.ppgrid1.rhul.ac.uk	OK	ok	ok	ok
7	UKI	UKI-LT2-UCL-HEP	lcg-ce01.hep.ucl.ac.uk	OK	ok	ok	ok
8	UKI	UKI-NORTHGRID-LANCS-HEP	fal-pygrid-44.lancs.ac.uk	OK	ok	ok	ok
9	UKI	UKI-NORTHGRID-LIV-HEP	hepgrid5.ph.liv.ac.uk	OK	ok	ok	ok
10	UKI	UKI-NORTHGRID-SHEF-HEP	lcgce0.shef.ac.uk	OK	ok	ok	ok
11	UKI	UKI-SCOTGRID-DURHAM	ce01.dur.scotgrid.ac.uk	OK	ok	ok	ok
12	UKI	UKI-SCOTGRID-DURHAM	ce02.dur.scotgrid.ac.uk	OK	ok	ok	ok
13	UKI	UKI-SCOTGRID-ECDF	ce2.glite.ecdf.ed.ac.uk	OK	ok	ok	ok
14	UKI	UKI-SCOTGRID-ECDF	mw05.ecdf.ed.ac.uk	OK	ok	ok	ok
15	UKI	UKI-SOUTHGRID-BHAM-HEP	epgr02.ph.bham.ac.uk	OK	ok	na	ok
16	UKI	UKI-SOUTHGRID-BHAM-HEP	epgr04.ph.bham.ac.uk	OK	ok	na	ok
17	UKI	UKI-SOUTHGRID-CAM-HEP	serv07.hep.phy.cam.ac.uk	ERROR	ok	error	ok
18	UKI	UKI-SOUTHGRID-OX-HEP	t2ce05.physics.ox.ac.uk	OK	ok	ok	ok
19	UKI	UKI-SOUTHGRID-RALPP	heplnx206.pp.rl.ac.uk	OK	ok	ok	ok
20	UKI	UKI-SOUTHGRID-RALPP	heplnx207.pp.rl.ac.uk	OK	ok	ok	ok
21	UKI	csTCDie	gridgate.cs.tcd.ie	OK	ok	ok	ok

- Mentioned in *ATLAS SSB monitoring* session talks

Monitoring: SLS

- <https://sls.cern.ch/sls/service.php?id=ATLAS-Frontier>
- Checks service up/down, but includes false positives; reported errors are not robust enough to aid troubleshooting








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

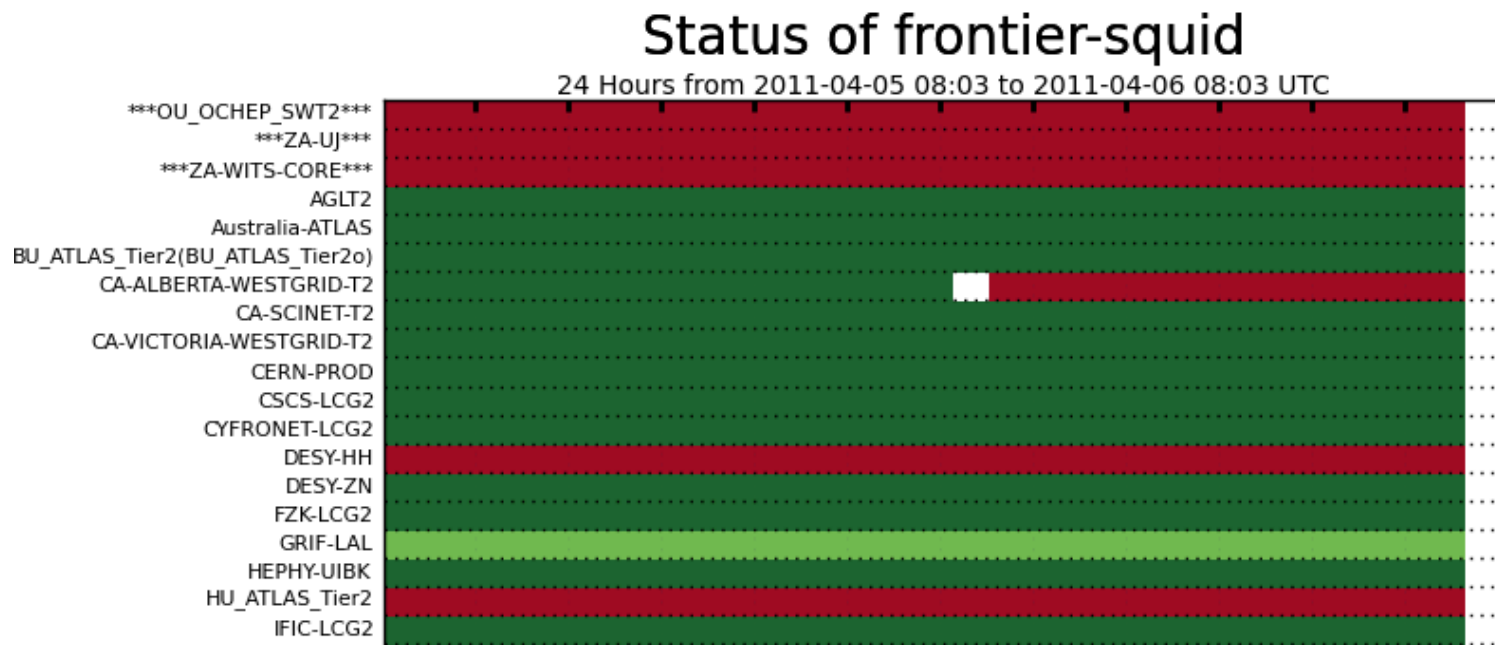
From: Sls.Administrator@cern.ch
Subject: SLS Notification (ATLAS-BNL-Frontier)
Date: March 4, 2011 8:40:49 AM GMT+01:00
To: John DeStefano , SLS.Administrator@cern.ch
Reply-To: Sls.Administrator@cern.ch

The service: ATLAS-BNL-Frontier is NOT AVAILABLE in SLS, please investigate the cause of this disruption.

You can see the service page in SLS:
<http://sls.cern.ch/sls/service.php?id=ATLAS-BNL-Frontier>

Monitoring: SSB

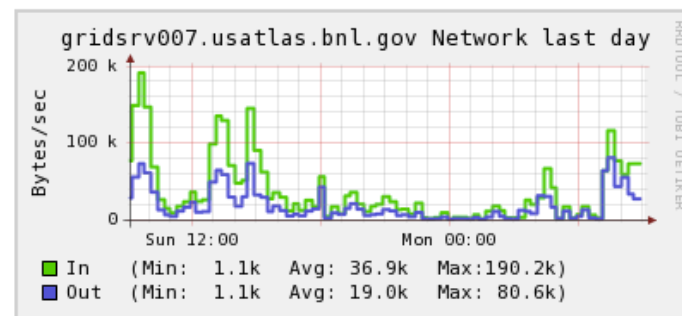
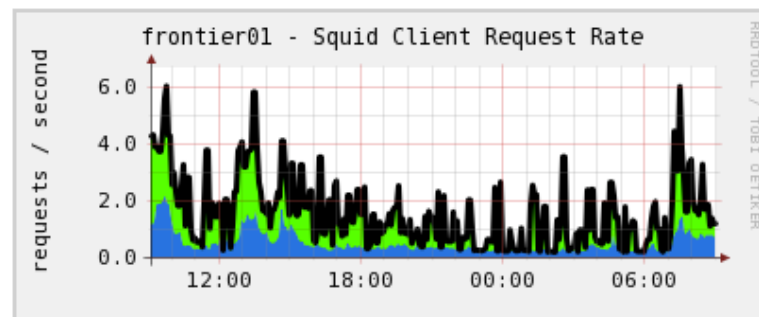
- <http://dashb-atlas-ssb.cern.ch/dashboard/request.py/siteviewhistory?columnid=47>
- Most T2 sites reporting, but Squid stats not completely accurate



- See: *ATLAS SSB monitoring* session talks (J. Schovancova et al)

Monitoring: Misc.

- Independent site tools:
 - Cacti
 - Ganglia
 - Custom log mining scripts
 - Etc.
- *None* of these tools are monitored by shifters



Savannah Activity

- 72011: Jobs failed at FZK
 - Solution: Site cache node added at FZK
 - See: 78607
- 78014: Configuration issue with CERN Squids
 - Solution: Cache expiration times at CERN reconfigured to match ATLAS recommendations
- 78607: Segmentation fault at Frontier sites without site cache nodes:
 - Solution: Frontier client recoded to handle scenario in v.2.8.0; not yet included in ATLAS software

Concerns

- Site Squids
 - Frontier server sites are encouraged to deploy one or more site Squid instances to cache request data, off-load Frontier servers, provide resiliency, redundancy, fail-over
 - Savannah bugs reported against Frontier due to lack of site Squids and subsequent failure of server fail-over (#78992 et al)
- CERN site maintenance
 - Squid and Frontier at CERN currently running precariously without local administrative resources, and with misconfigurations that were recently discovered and addressed manually
- Monitoring
 - In need of centralized, robust ATLAS solution as part of ADC monitoring
 - See: *ADCoS/Comp@P1/Squad/AMOD/DAST Roles* (J. Schovancova)
- Down-time
 - Currently no method to declare Frontier or Squid down-time in ATLAS, or inherit service down-time from related services (Conditions database, site down-time, etc.)
 - Possible future solution via AGIS integration, but nothing currently
- Squid sharing
 - Desire to share Squids among services or VOs to reduce administrative burden
 - Example Squids at RAL: 2 CMS Frontier, 2 ATLAS Frontier, 2 CVMFS, 1 LHCb test

News and Plans

- Geometry data
 - G. Dimitrov, V. Tsulaia: ATLAS currently testing Frontier to cache Geometry data for user analysis jobs
 - Impact on Frontier servers seems minimal thus far
 - See: *DB releases and Geometry DB in Frontier* (V. Tsulaia)
- ATLAS Software Packages
 - D. Dykstra, D. Front: Optimization in progress for RPMs for Squid, Tomcat, and Frontier servlet
 - Good time for comments, recommendations for packages
- AGIS integration
 - A. Anisyonkov, A. Di Girolamo: Technical specification being reviewed and discussed, work to begin soon
 - See: *AGIS News* (A. Anisyonkov, A. Di Girolamo)
- CVMFS
 - Frontier Squid package used in CVMFS cache/replica test bed at BNL
 - See: *CVMFS Status and Plans* session (J. Blomer, I. Collier, A. De Salvo, S. Traylen, A. Undrus)
- HI reprocessing
 - Frontier used for some heavy ion reprocessing queries
 - See: *Reprocessing of Heavy Ion and Muon Data* (A. Vaniachine)