



The Frontier System for the Fr Atlas Cloud

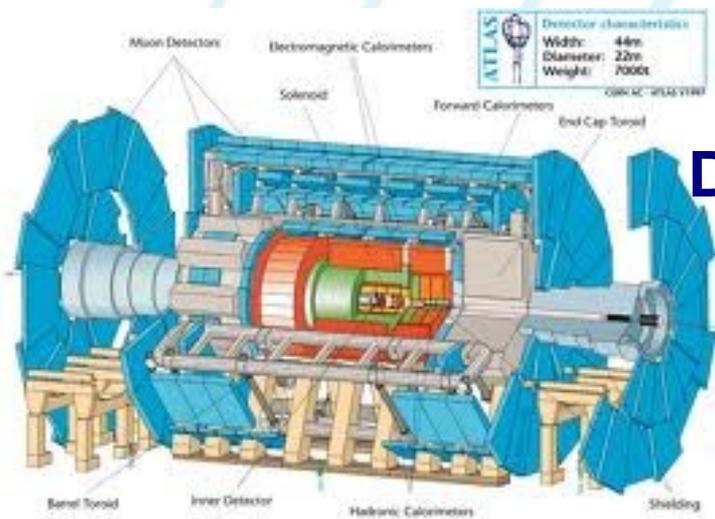
Vamvakopoulos E.
GRIF-Tokyo-Beijing workshop,
15th-19th June,
LPNHE, Paris

Transfer metadata across Frontiers



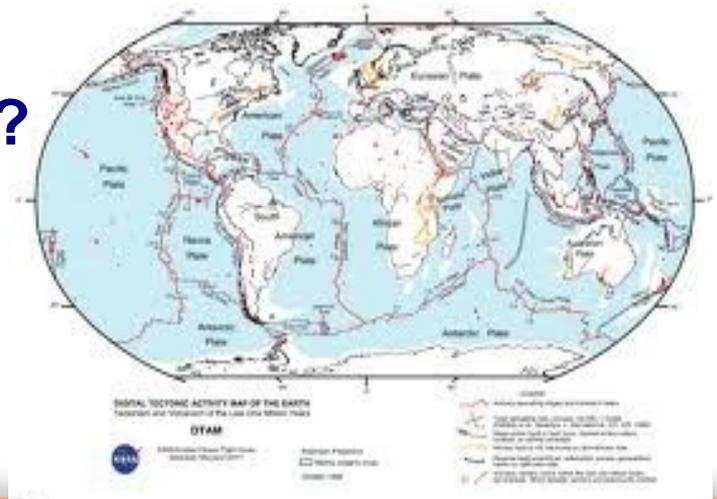
Conditions data refers to nearly all the non-event data produced during the operation of the ATLAS detector, together with that required to perform reconstruction and analysis¹.

Atlas Detector at CERN



Hundreds GRID sites beyond the Frontiers

How
Conditions
Data transfer?



[1]. ATLAS TDR-017, CERN-LHCC-2005-022

▶ *Frontier System :*



A System for Data Distribution from a Central DB to Many Client System²

- Frontier uses HTTP based Protocol (Rest Arch)³
- Developed at FermiLab for the needs CDF
- Adapted from LHC experiments at CERN: **CMS**, **ATLAS** to access "*conditions data*" at all their sites worldwide.

[2] <http://frontier.cern.ch/>

[3] <http://www.ics.uci.edu/~fielding/pubs/dissertation/top.htm>

Transfer metadata across the frontiers ...



In the case of Condition Data from *ATLR* are transferred down to the *ATLAS* payload (jobs) by :

a distrusted network of read-only replicas of the central *ATLR DB* (**CERN**) on selected

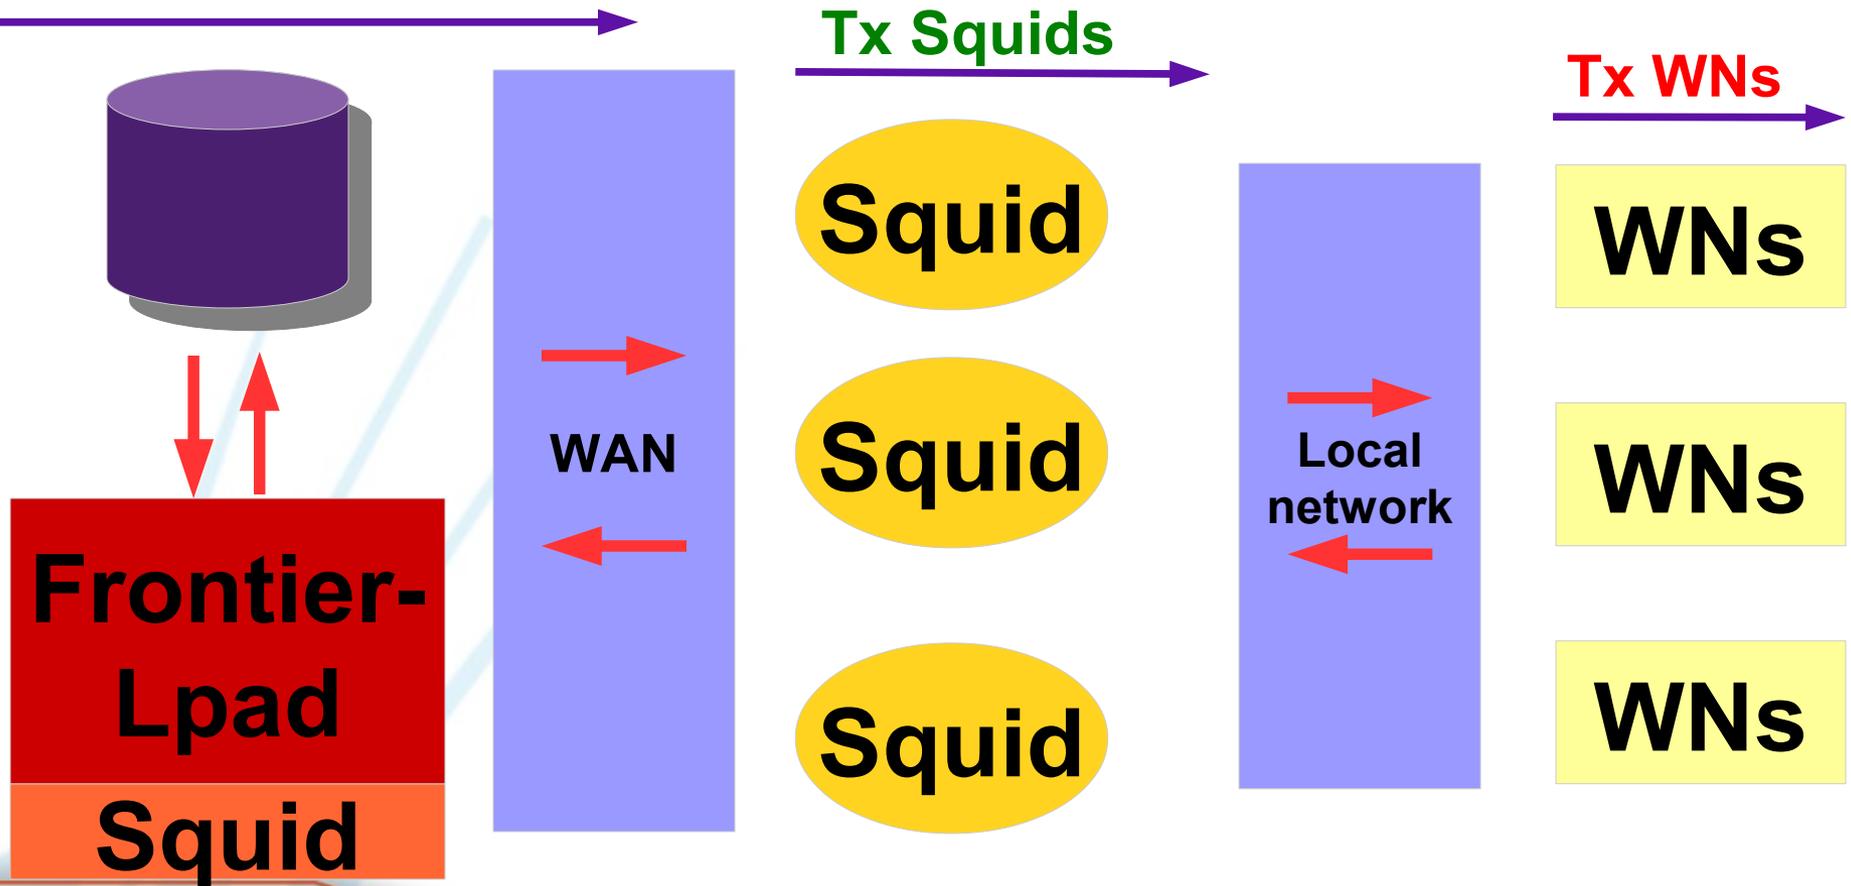
Tier1 sites (**IN2P3-CC**, **RAL**, **KIT**, **BNL** and **TRIUMF**)

and the employment of the **Frontier system**

Atlas Frontier System



RO Replica of the
ATLR Database: T1



▶ Recommendation



ATLAS requests that selected T1 sites install both Frontier (launchpad) and Local Squid servers and that each **Tier 2 site install one local Squid instance per **500 user analysis batch queue slots** hosted at their site**

(ref: ATLAS Trigger and Offline Board [TOB], September 16, 2009).

■ *Two Frontier-Lpad*

- under DNS load balancing
- Alias: **ccfrontier.in2p3.fr**

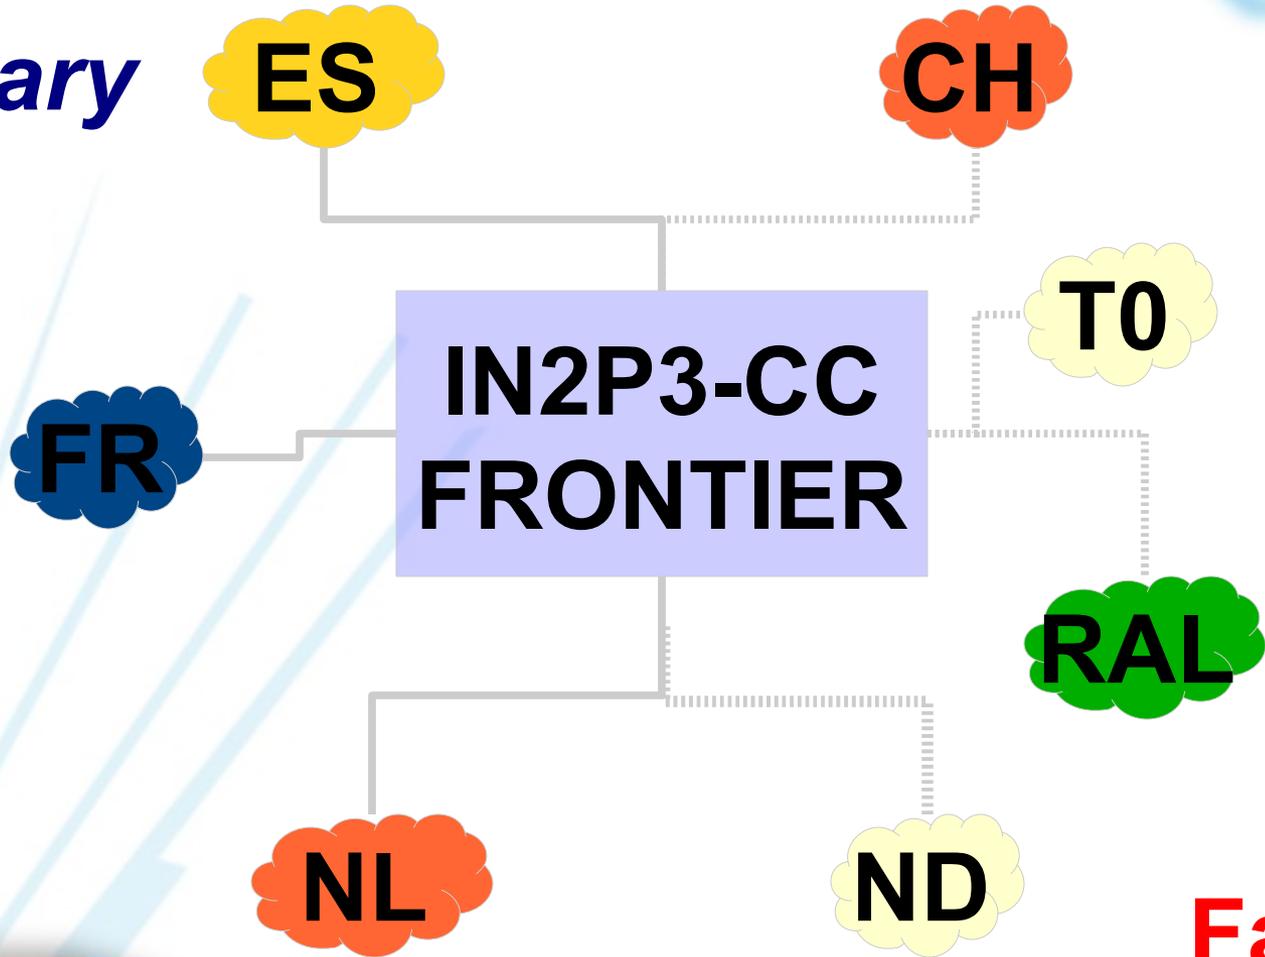
■ *Two Site Squids*

- under DNS load balancing
- Alias: **ccaltassquid.in2p3.fr**
- **only for the Frontier dedicated usage (acl on destination)**

IN2P3-CC FRONTIER



Primary



Failover

Squid Network



- GRIF
 - LAL
 - LPNHE
- IRFU/CEA
- LAPP*
- IN2P3-CC*
 - CPPM
 - LSPC
 - LPC
- BEIJING*
- RO-02-NIPNE*
 - RO-14-ITIM
- RO-07-NIPNE*
- RO-16-UAIC*
- TOKYO*

Frontier endpoint Configuration via AGIS



- *Frontier Configuration for the site: primary and backup Frontier Service endpoint.*
- *Squid Configuration for the site: primary and backup squid service endpoint.*

Atlas installation machinery



add to the atlas local environment setup of ... the, a site variable with a specific format ...

```
FRONTIER_SERVER="(serverurl=http://ccfrontier.in2p3.fr:23128/ccin2p3-AtlasFrontier)
(serverurl=http://lcgft-atlas.gridpp.rl.ac.uk:3128/frontierATLAS)
(proxyurl=http://ccatlassquid.in2p3.fr:3128)"
```

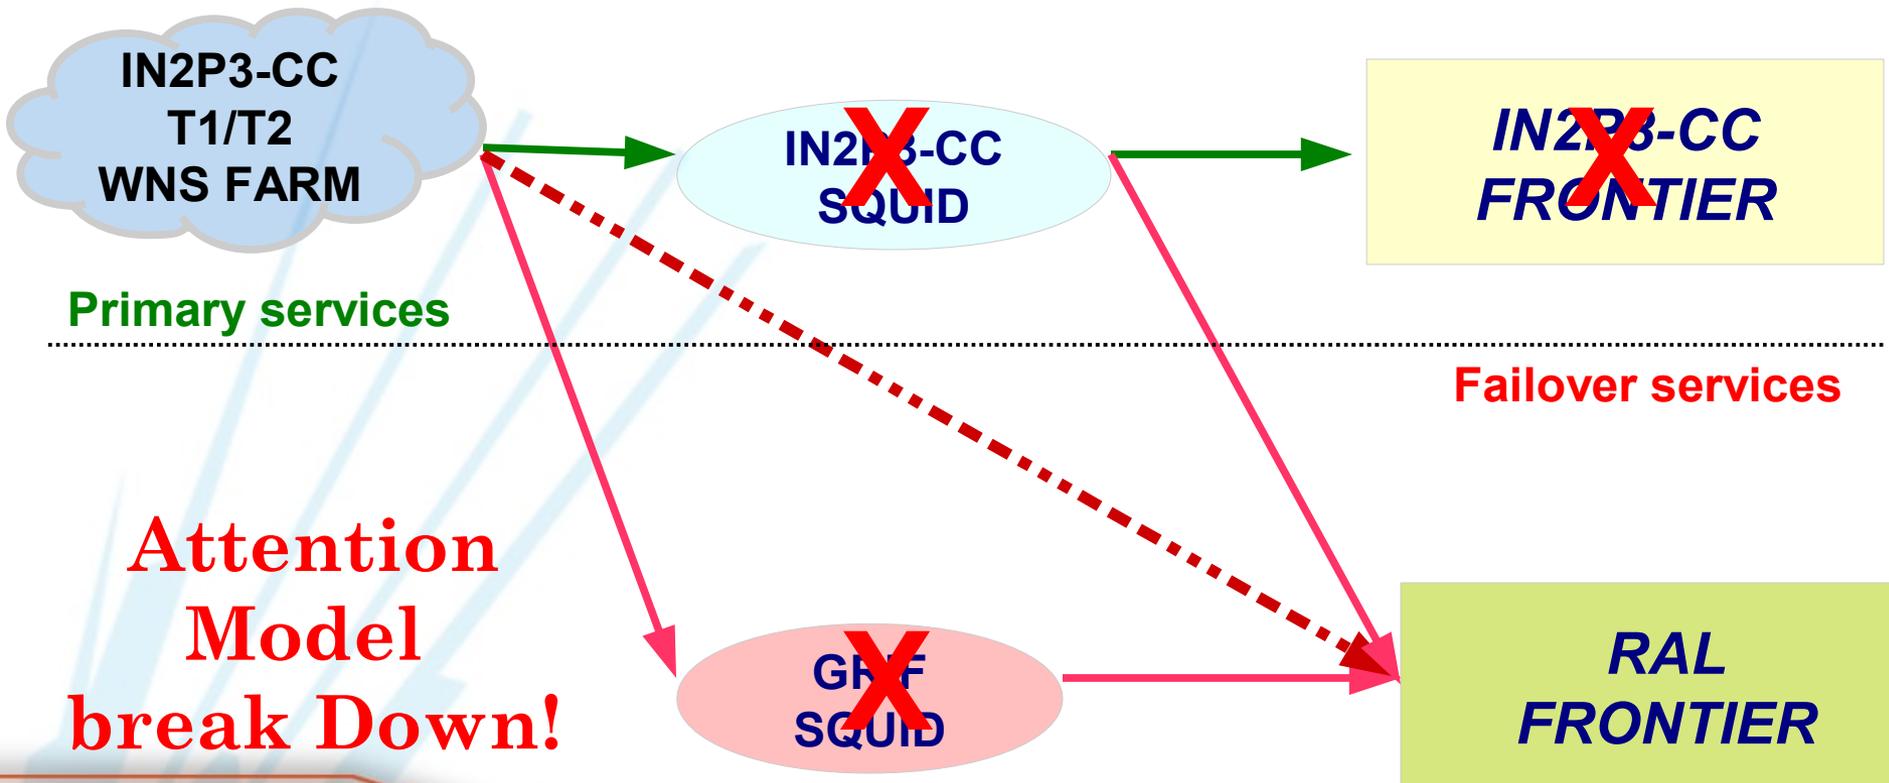
....

Topology & Fail-over



Failover in the level of the client :

if a Frontier server fails then the client (frontier lib) will try other Frontier in order to get their data



*Central Monitoring of
Frontier-Lpad and Squids*

Central Monitoring of Frontier-Lpad and Squid service



- **Service Level Status of ATLAS launchpad servers(SLS)**
 - <https://sls.cern.ch/sls/service.php?id=ATLAS-Frontier>
- **SUM: Site Usability Monitor & SSB dashboard**
 - <http://dashb-atlas-sum.cern.ch>
 - <http://dashb-atlas-ssb.cern.ch/>
- **Awstats Monitoring for Frontier-launchpad** (A nice open source tool that generates advanced squid-proxy,web, ftp ,mail server statistics: Domains/countries, Host, Number of visits, TCP_HIT, TCP_MISS)
 - <http://frontier.cern.ch/awstatsatlas.html>
- **MRTG Monitoring :** (Request/Fetch – in/out Traffic- Count object)
 - <http://frontier.cern.ch/squidstats/indexatlas.html>

Service Level Status at CERN

Service Level Status overview - Mozilla Firefox

File Edit View History Bookmarks Tools Help

AdcMonitoring < Atlas < TW... FroNTier < Atlas < Twiki ATLAS awstats SLS Service Level Status overview SLS Service Level Status overview

cern.ch https://sls.cern.ch/sls/service.php?id=ATLAS-Frontier

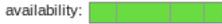
Réunion équipe Op... Software & Comput... Software & Comput... ADC Monitoring Informations about ... Tier-0 Monitoring O... Atlas Running Job... Logbook sur la cam...

Home Search KPIs Tags Admin Documentation Help

ATLAS Frontier Service

13 Jun 2012 Wed 23:45:31

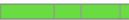
ATLAS-Frontier

availability:  (more)

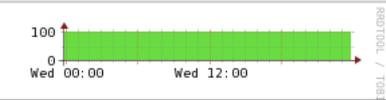
percentage: 100%

status: **available**

this service consists of:

-  ATLAS-CERN-Frontier
-  ATLAS-T0-Frontier
-  ATLAS-BNL-Frontier
-  ATLAS-KIT-Frontier
-  ATLAS-IN2P3-CC-Frontier
-  ATLAS-RAL-Frontier
-  ATLAS-TRIUMF-Frontier

availability in the last 24 hours (more):



availability update

last update: 23:40:39, 13 Jun 2012 (5 minutes ago)

expires after: 60 minutes

 [rss feed with status changes](#)

Part of (subservice of):

-  ADC Central Services

Clusters, subclusters and nodes

none / not declared

Admin

[admin tools](#)

- Reason of Outages
- Schedule downtime of Oracle backend DB
- Problem on back-end storage
- Network connectivity

Service Level Status overview

SUM AVAILABILITY TEST



SUM availability test for Frontier/squid submits a simple job, which tests different combination of the Frontier endpoint and the squid endpoint (normally 4 pair)

File Edit View History Bookmarks Tools Help

SUM Visualization x SUM-Detailed Results x

dashb-atlas-sum.cern.ch/dashboard/request.py/latestresultssmry-sum#profile=ATLAS&group=

Réunion équipe Op... Software & Comput... Software & Comput... ADC Monitoring Informa

Algorithm for calculating the Site and Service Availability

Legends for Metric Result Status

Status:	NA	OK	WARNING	CRITICAL	UNKNOWN	MIS
Legend:	NA	OK	W	C	U	I

Note: brightest colors: test is 0 - 12 hours old, ... lightest colors: test is more that 12 hours old

Legends for Metric Names

Legend	Metric Name	Legend	Metric Name	Legend	Metric Name	Legend	Metric Name
1	org.atlas.WN-FrontierSquid						

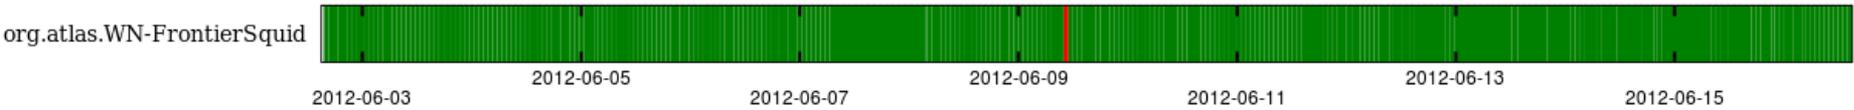
[Link to data](#)

Sitename	Flavour	Hosts	1
BEIJING-LCG2	CE	lcg002.lhep.ac.cn	OK
			1
	CREAM-CE	cce.lhep.ac.cn	OK
GRIF	CE	grid10.lal.in2p3.fr	OK
			1
	CREAM-CE	grid36.lal.in2p3.fr	OK
		lpnhe-cream.in2p3.fr	OK
		node74.datagrid cea.fr	OK
IN2P3-CC	CREAM-CE	cccreamcell05.in2p3.fr	OK
		cccreamcell06.in2p3.fr	OK
IN2P3-LAPP	CE	lapp-ce01.in2p3.fr	OK
			1
	CREAM-CE	lapp-ce02.in2p3.fr	OK
RO-14-ITIM	CREAM-CE	cream.itim-cj.ro	W

SUM test historical view: examples ..

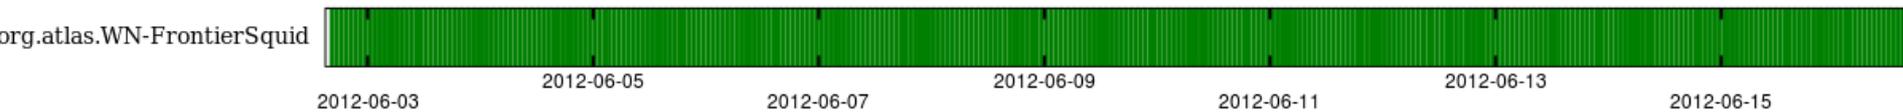
Test history cce.ihep.ac.cn using ATLAS

336 hours from 2012-06-02 17:00 to 2012-06-16 17:00



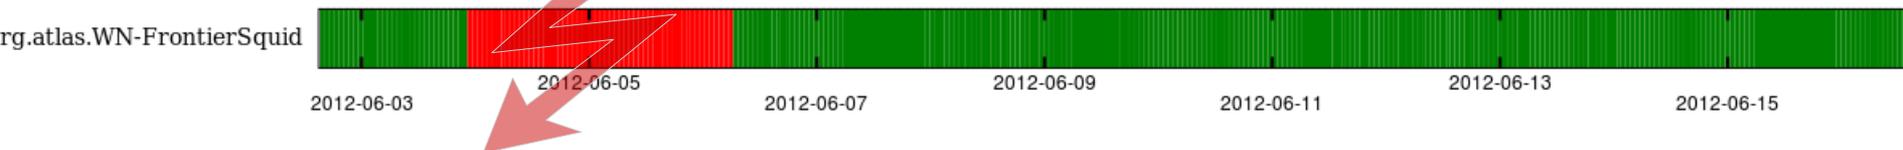
Test history lpnhe-cream.in2p3.fr using ATLAS

336 hours from 2012-06-02 17:00 to 2012-06-16 17:00



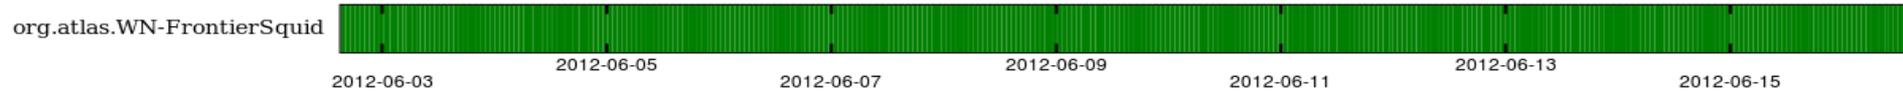
Test history lapp-ce02.in2p3.fr using ATLAS

336 hours from 2012-06-02 17:00 to 2012-06-16 17:00



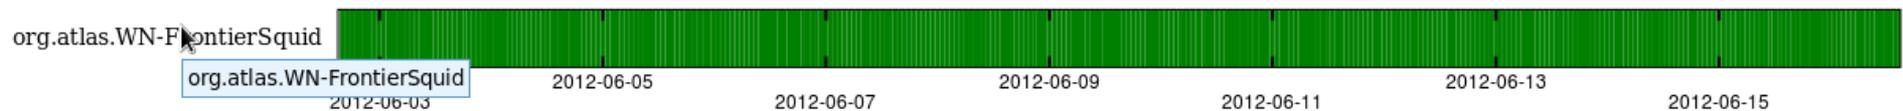
Test history lcg-ce03.icepp.jp using ATLAS

336 hours from 2012-06-02 17:00 to 2012-06-16 17:00



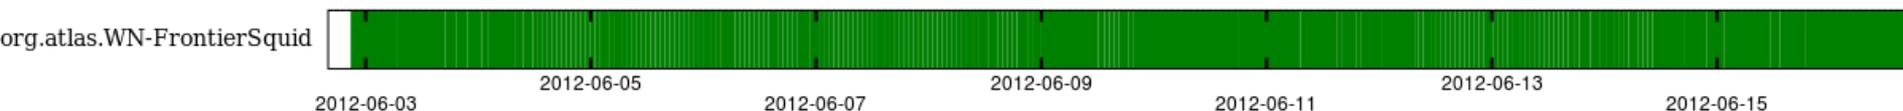
Test history tbat01.nipne.ro using ATLAS

336 hours from 2012-06-02 17:00 to 2012-06-16 17:00



Test history tbit03.nipne.ro using ATLAS

336 hours from 2012-06-02 18:00 to 2012-06-16 18:00



AWSTATS ...



Firefox (3) GNU Image Manipulati... [Talks_paris] [vamvakop@localhost... Paris_15june.odp - Op...]

Applications Places System USA 13 °C Thu Jun 14, 12:29:36 AM Vamvakopoulos Emmanouil

Statistics for ccin2p3-ccsqfatlasli02 (2012-06-13) - main - Mozilla Firefox

frontier.cern.ch/awstats/awstats.pl?config=ccin2p3-ccsqfatlasli02&databasebreak=day&day=13 awstats

Réunion équipe Op... Software & Comput... ADC Monitoring Informations about ... Tier-0 Monitoring O... Atlas Running Job... Logbook sur la cam...

Statistics for: ccin2p3-ccsqfatlasli02

Summary
When:
 Monthly history
 Days of month
 Days of week
 Hours
Who:
 Countries
 Full list
 Hosts
 Full list
 Last visit
 Unresolved IP Address
Navigation:
 Visits duration
 Request size
 Request time
 Viewed
 Full list
 Entry
 Exit
Others:
 HTTP Status codes
 Pages not found
Extra/Marketing:
 Squid Cache Hits

11	41534	41534	230.02 MB	23	33056	33056
----	-------	-------	-----------	----	-------	-------

Visitors domains/countries (Top 10) - Full list

Domains/Countries	Pages	Hits	Bandwidth
France	304713	304713	2.59 GB
Spain	153366	153366	619.49 MB
Israel	108800	108800	430.56 MB
Portugal	83497	83497	240.91 MB
Romania	72988	72988	219.35 MB
Japan	55972	55972	325.41 MB
Netherlands	50804	50804	243.71 MB
China	49725	49725	199.81 MB
Former USSR	47067	47067	208.98 MB
Unknown	30673	30673	24.61 MB

Hosts (Top 10) - Full list - Last visit - Unresolved IP Address

Hosts	Pages	Hits	Bandwidth	Last visit
ccsvli52.in2p3.fr	3341477	3341477	34.24 GB	31 May 2012 - 23:58
ccsvli42.in2p3.fr	3283312	3283312	33.74 GB	31 May 2012 - 23:59
node14.datagrid.cea.fr	3030168	3030168	44.30 GB	30 May 2012 - 16:32
grid15.lal.in2p3.fr	1900470	1900470	14.89 GB	31 May 2012 - 23:59
conddb-px01.icepp.jp	1761651	1761651	13.05 GB	31 May 2012 - 23:51
sq5.ific.uv.es	1752437	1752437	12.62 GB	31 May 2012 - 23:52
servicios.ft.uam.es	1268095	1268095	4.87 GB	31 May 2012 - 23:21
atlassq.ihep.ac.cn	1264793	1264793	6.54 GB	31 May 2012 - 23:58
atrogro08.nipne.ro	1160741	1160741	3.84 GB	31 May 2012 - 23:55
gridgw.nipne.ro	1049794	1049794	4.43 GB	31 May 2012 - 23:50
Others	10748980	10748980	56.22 GB	

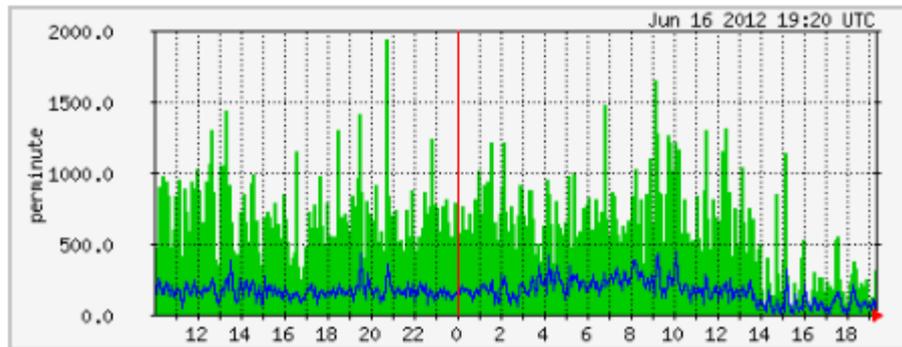
Daily Traffic: Mai 2012

Monthly Traffic: Mai 2012

Frontier-lpad at IN2P3-CC

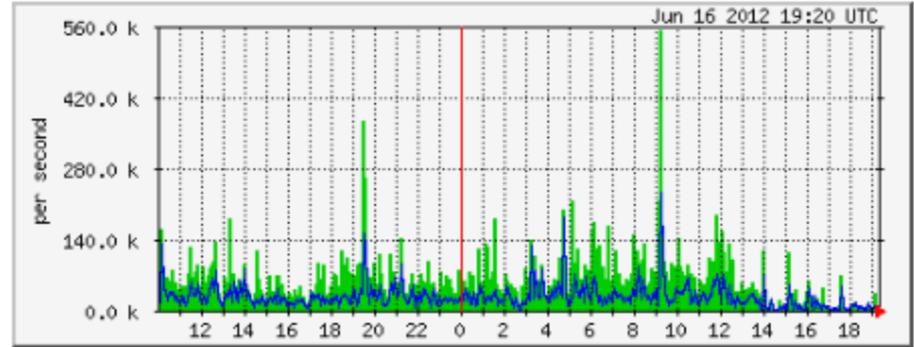


`Daily' Graph (5 Minute Average)



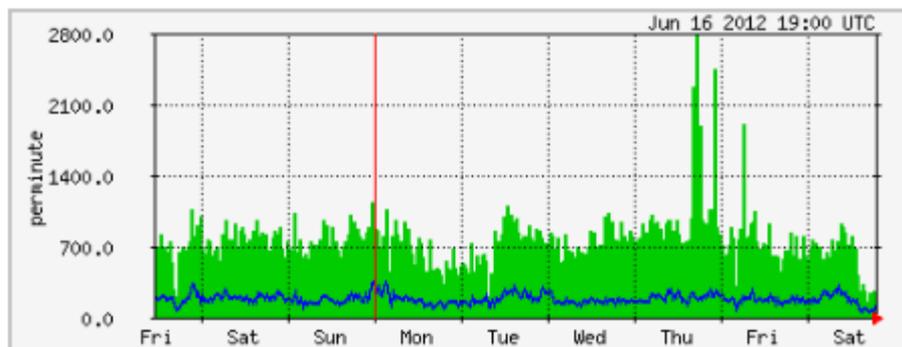
	Max	Average	Current
HTTP reqs	1932.0 req/min	572.0 req/min	306.0 req/min
HTTP fetches	433.0 req/min	157.0 req/min	123.0 req/min

`Daily' Graph (5 Minute Average)



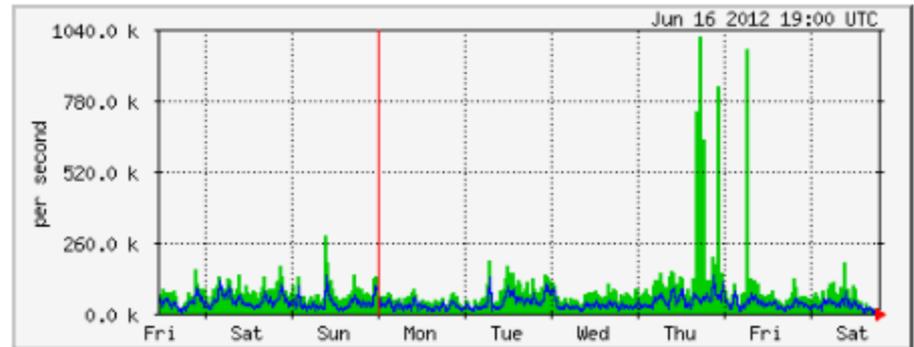
	Max	Average	Current
Total	551.0 kB/s	56.0 kB/s	8.0 kB/s
Fetches	230.0 kB/s	25.0 kB/s	4.0 kB/s

`Weekly' Graph (30 Minute Average)



	Max	Average	Current
HTTP reqs	2780.0 req/min	703.0 req/min	163.0 req/min
HTTP fetches	347.0 req/min	172.0 req/min	59.0 req/min

`Weekly' Graph (30 Minute Average)

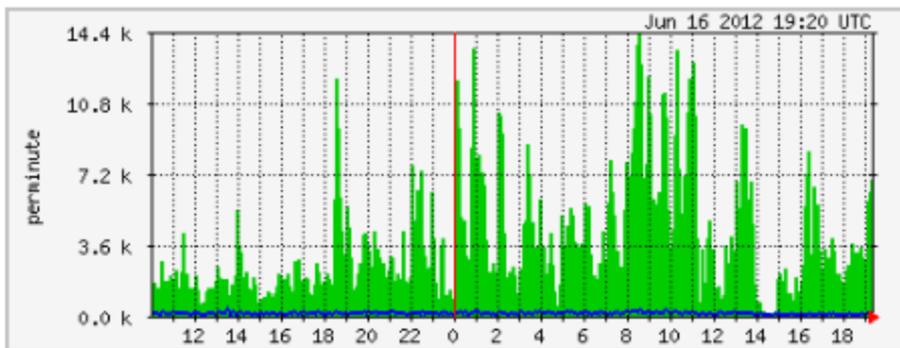


	Max	Average	Current
Total	1009.0 kB/s	81.0 kB/s	6.0 kB/s
Fetches	133.0 kB/s	35.0 kB/s	4.0 kB/s

SQUID at IN2P3-CC

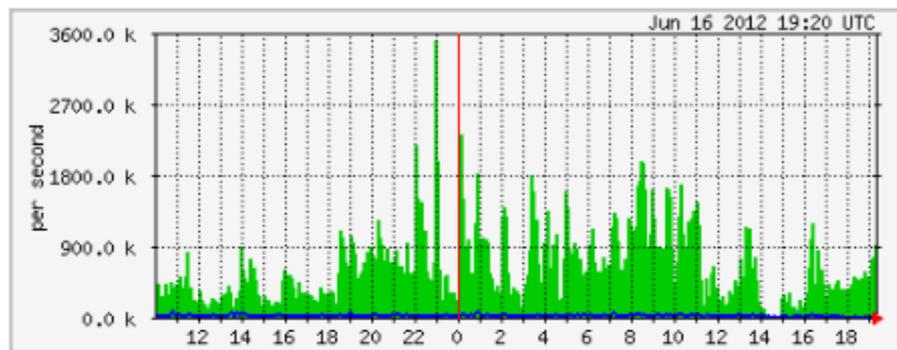


`Daily' Graph (5 Minute Average)



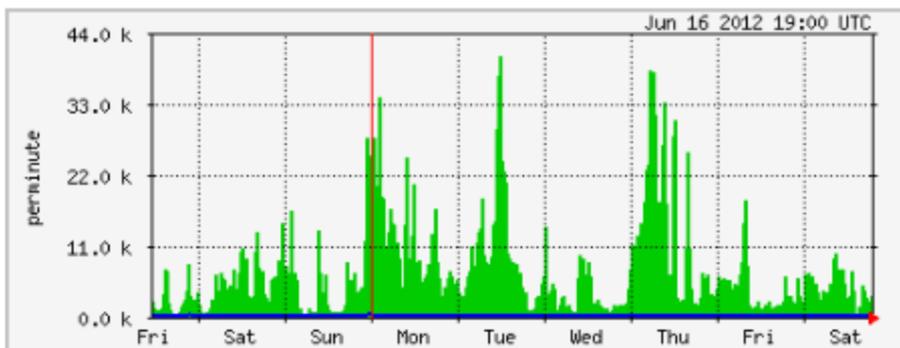
	Max	Average	Current
HTTP reqs	14.3 kreq/min	3346.0 req/min	6928.0 req/min
HTTP fetches	372.0 req/min	112.0 req/min	48.0 req/min

`Daily' Graph (5 Minute Average)



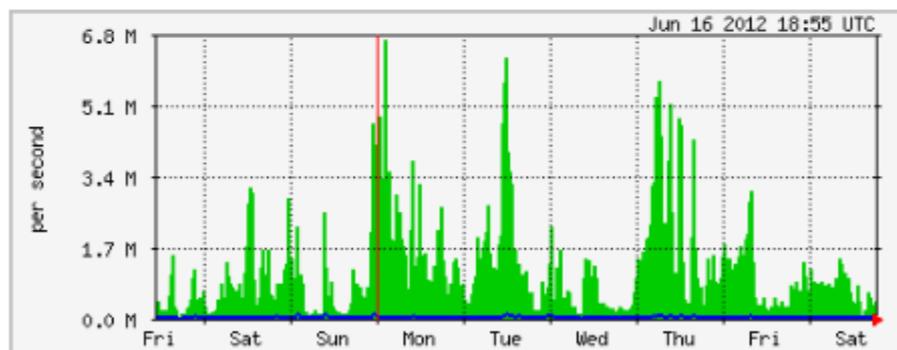
	Max	Average	Current
Total	3506.0 kB/s	578.0 kB/s	918.0 kB/s
Fetches	74.0 kB/s	13.0 kB/s	1.0 kB/s

`Weekly' Graph (30 Minute Average)



	Max	Average	Current
HTTP reqs	40.4 kreq/min	6464.0 req/min	3099.0 req/min
HTTP fetches	305.0 req/min	122.0 req/min	36.0 req/min

`Weekly' Graph (30 Minute Average)

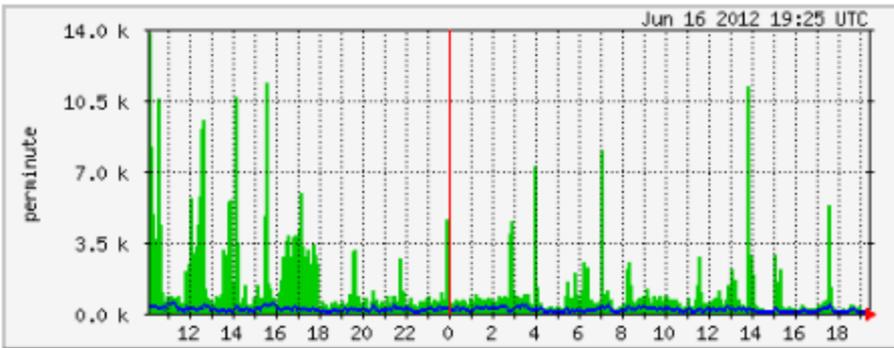


	Max	Average	Current
Total	6799.0 kB/s	1093.0 kB/s	421.0 kB/s
Fetches	120.0 kB/s	18.0 kB/s	0.0 kB/s

TOKYO

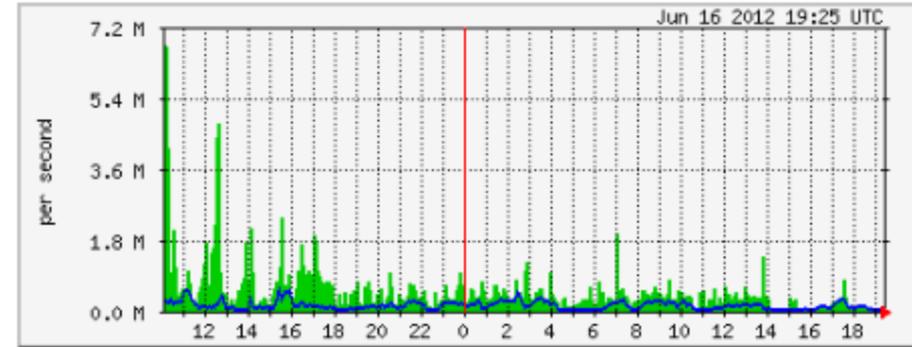


'Daily' Graph (5 Minute Average)



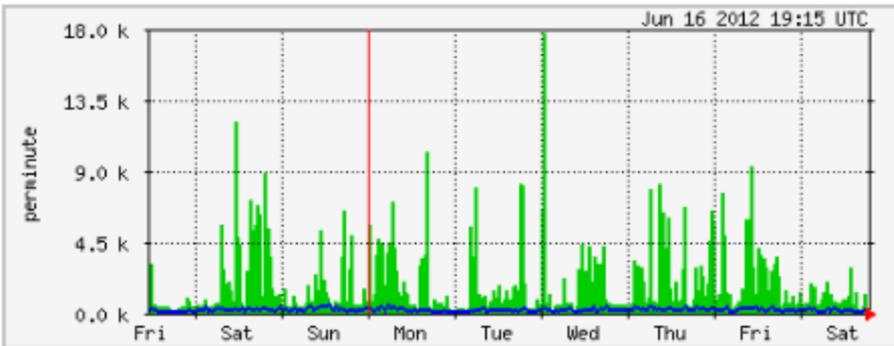
	Max	Average	Current
HTTP reqs	13.8 kreq/min	1156.0 req/min	195.0 req/min
HTTP fetches	502.0 req/min	159.0 req/min	12.0 req/min

'Daily' Graph (5 Minute Average)



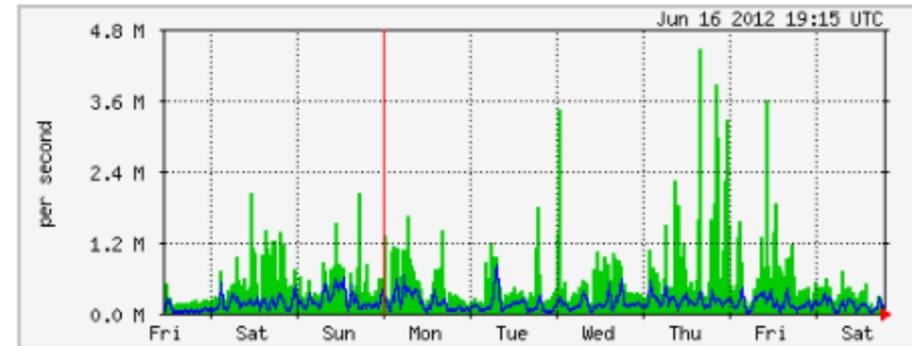
	Max	Average	Current
Total	6901.0 kB/s	470.0 kB/s	28.0 kB/s
Fetches	582.0 kB/s	130.0 kB/s	9.0 kB/s

'Weekly' Graph (30 Minute Average)



	Max	Average	Current
HTTP reqs	17.7 kreq/min	1560.0 req/min	295.0 req/min
HTTP fetches	516.0 req/min	177.0 req/min	91.0 req/min

'Weekly' Graph (30 Minute Average)



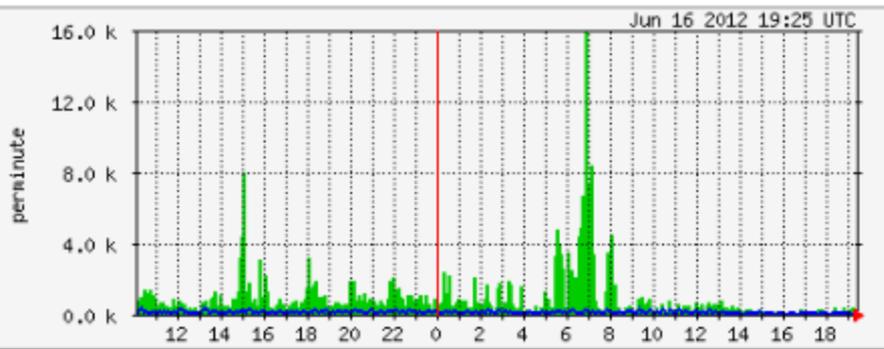
	Max	Average	Current
Total	4560.0 kB/s	552.0 kB/s	68.0 kB/s
Fetches	822.0 kB/s	162.0 kB/s	40.0 kB/s



BEIJING

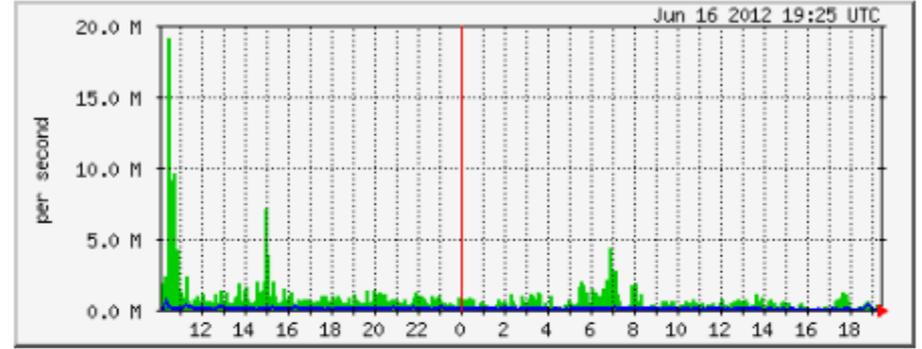


Daily' Graph (5 Minute Average)



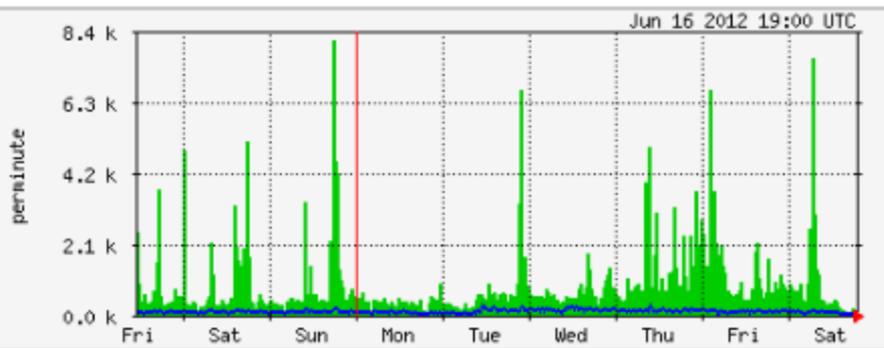
	Max	Average	Current
HTTP reqs	15.8 kreq/min	771.0 req/min	178.0 req/min
HTTP fetches	359.0 req/min	82.0 req/min	11.0 req/min

'Daily' Graph (5 Minute Average)



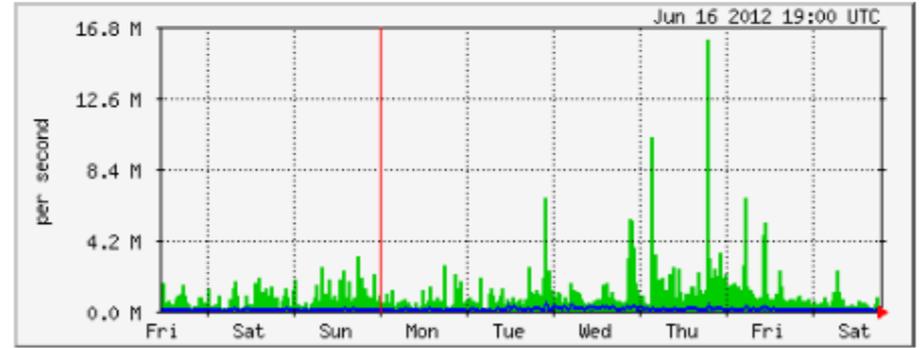
	Max	Average	Current
Total	19.1 MB/s	704.0 kB/s	86.0 kB/s
Fetches	687.0 kB/s	29.0 kB/s	0.0 kB/s

Weekly' Graph (30 Minute Average)



	Max	Average	Current
HTTP reqs	8112.0 req/min	805.0 req/min	125.0 req/min
HTTP fetches	305.0 req/min	97.0 req/min	21.0 req/min

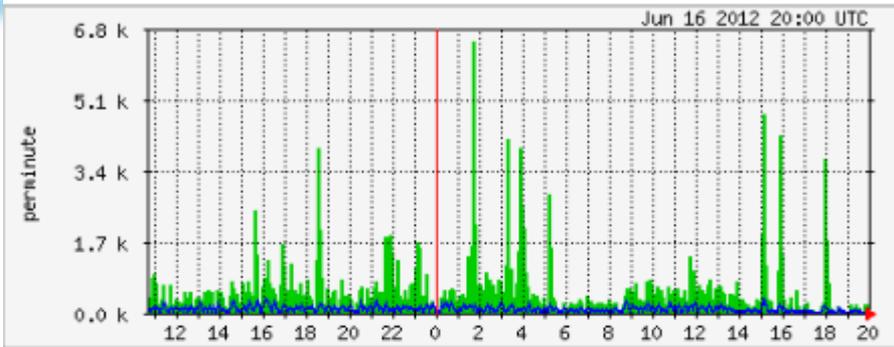
'Weekly' Graph (30 Minute Average)



	Max	Average	Current
Total	16.0 MB/s	978.0 kB/s	261.0 kB/s
Fetches	452.0 kB/s	58.0 kB/s	148.0 kB/s

LAPP

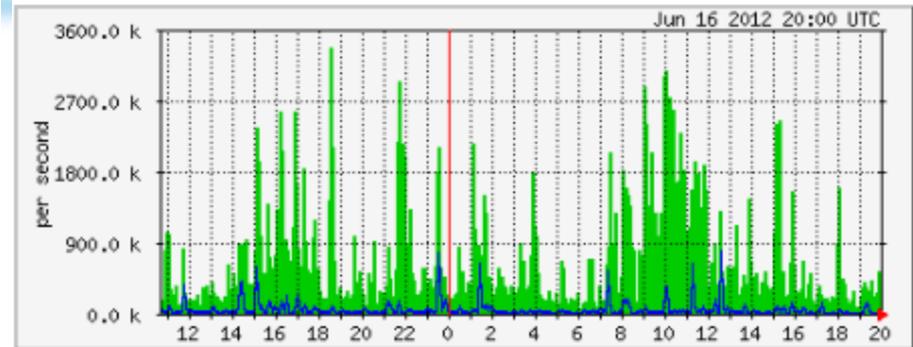
'Daily' Graph (5 Minute Average)



Max **Average** **Current**

HTTP reqs	6474.0 req/min	498.0 req/min	43.0 req/min
HTTP fetches	340.0 req/min	74.0 req/min	0.0 req/min

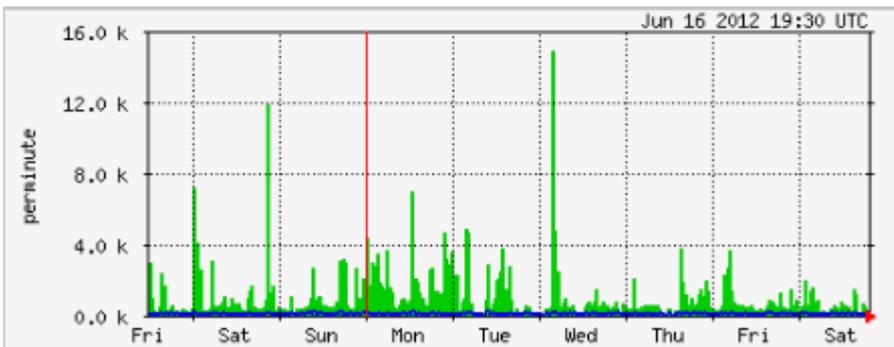
'Daily' Graph (5 Minute Average)



Max **Average** **Current**

Total	3354.0 kB/s	613.0 kB/s	123.0 kB/s
Fetches	766.0 kB/s	37.0 kB/s	0.0 kB/s

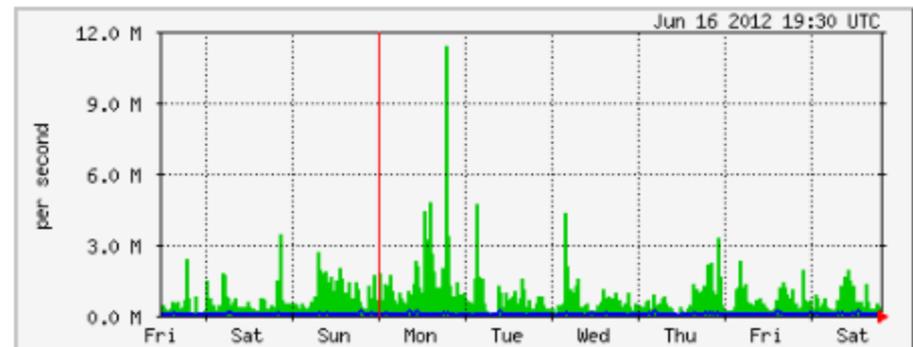
'Weekly' Graph (30 Minute Average)



Max **Average** **Current**

HTTP reqs	14.9 kreq/min	885.0 req/min	111.0 req/min
HTTP fetches	262.0 req/min	77.0 req/min	21.0 req/min

'Weekly' Graph (30 Minute Average)



Max **Average** **Current**

Total	11.3 MB/s	757.0 kB/s	309.0 kB/s
Fetches	174.0 kB/s	29.0 kB/s	49.0 kB/s

■ Service Criticality:

- *Frontier Criticality: **High** Allowed downtime: **Less than 1 day***
- *Frontier_squid Criticality: **Moderate** Allowed downtime: **Less than 2 days***

■ Guideline ADCos (24/7) :

- Frontier SLS : Periodically (2-3 times per shift)
- Squid SSB/SUM: “... Shifters should check the Frontier monitoring once per shift and if problems are observed non-urgent GGUS tickets can be submitted.. “

Operational issues I



- *Frontier System in FR cloud it Stable*
- *Small number of WARNINGS in SUM test for due to acls (access control list)*
- *Recent Problem with Atlas installation software ...*
 - *Frontier service is configured based on ToA DQ2 and not on AGIS (ADC experts informed)*
 - *LAPP case, ROMANIAN sites IN2P3-CC AND NL cloud.*

Further issues



- *Follow the Frontier service at T1*
 - *Improve the local documentation*
 - *Follow the recommendations and rpm updates from ADC*
 - *No GGUS ticket last 7 month*
- *Solve the recent problem of Frontier configuration with ADC experts*
 - *For Romanian sites ,*
 - *LAPP and IN2P3-CC*
- *Improve the squid network for the FR cloud*
 - *Improve the Monitoring for the FR cloud: squid Grif ...*
- *Make a mapping of FR Cloud squids*
 - *Hardware spec ...*
 - *Number of Squid ? / shared with CVMFS or not ?*

ATLAS Frontier Project



- *Alastair Dewhurst (RAL) (cord.)*
- *<https://twiki.cern.ch/twiki/bin/viewauth/Atlas/FroNTier>*
- *atlas-adc-frontier@cern.ch*
- *Regular meeting every second Wed 16:30*

New Mailing list frontier-talk@cern.ch

***!!SPECIAL THANKS!! TO CC-IN2P3
COLLEGUES FOR THEIR CONTRIBUTION
TO THE OPERATION OF THE FRONTIER SERVICE FOR THE ATLAS VO***

***!!SPECIAL THANKS!! TO FR-CLOUD
COLLEGUES FOR THEIR CONTRIBUTION
TO THE CONFIGURATION OF THE FRONTIER SERVICE***

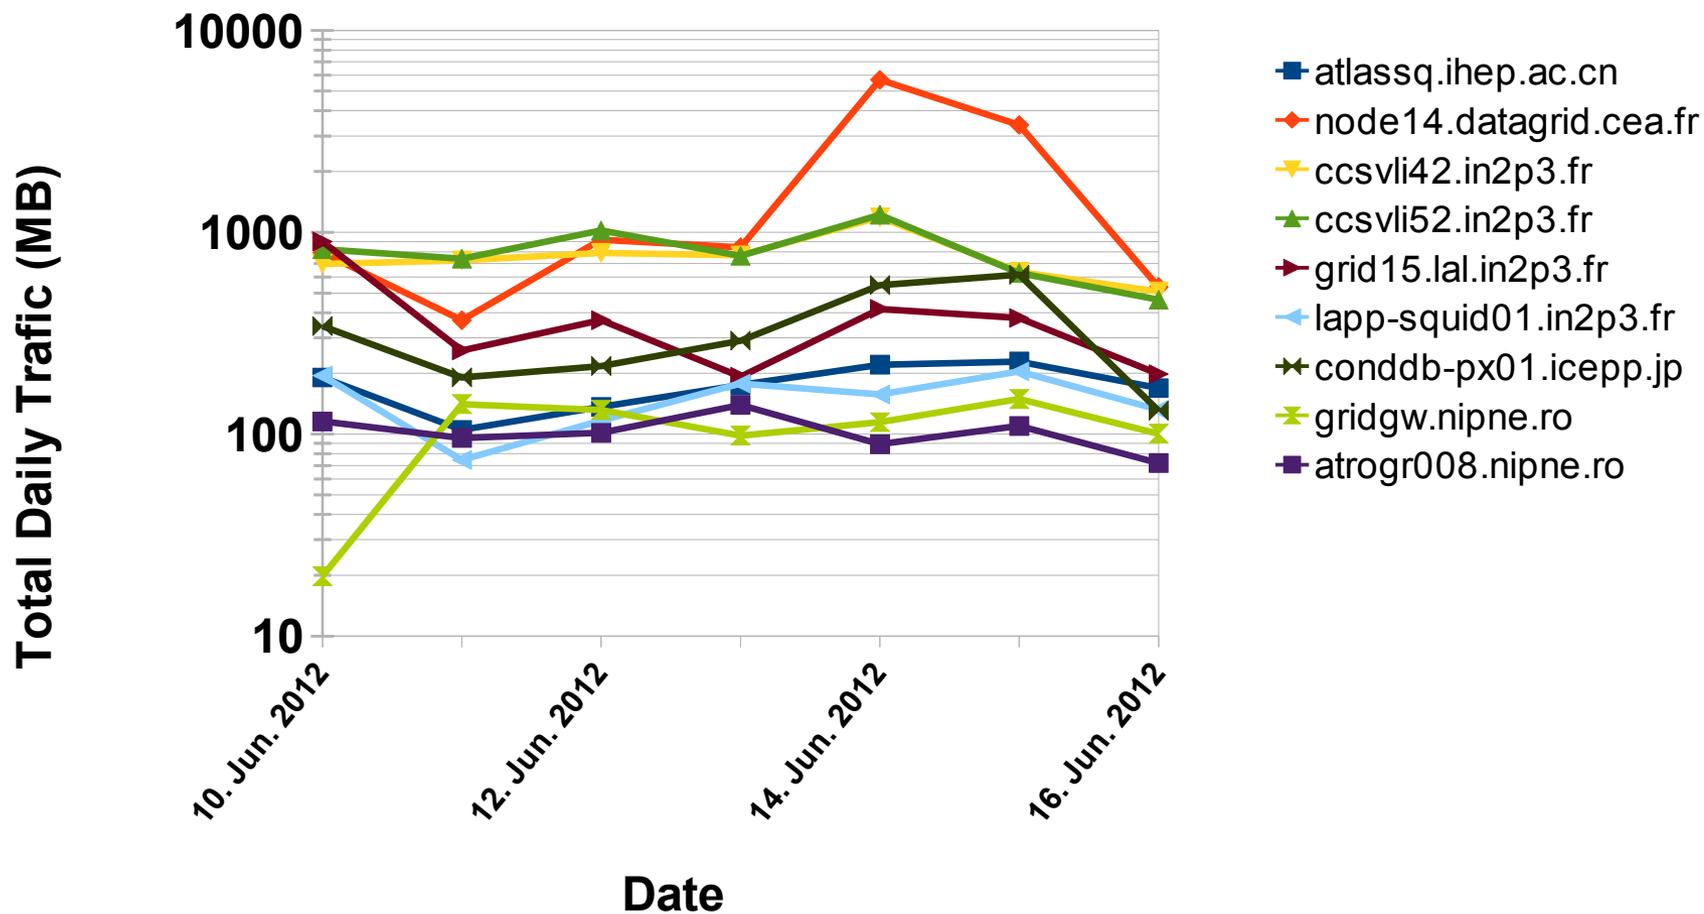
***!!SPECIAL THANKS!! TO ATLAS-FRONTIER
-GROUP FOR THE FRUITFULL DISCUSSION
AND FOR THE HELP***

Thank you for Your Attention



BACKUP SLIDES

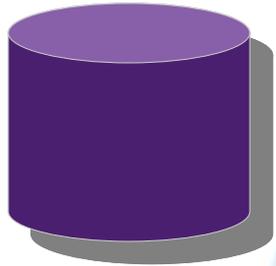
Indicate Daily traffic : for frontier T1 to Fr squid



▶ *Frontier Launchpad*

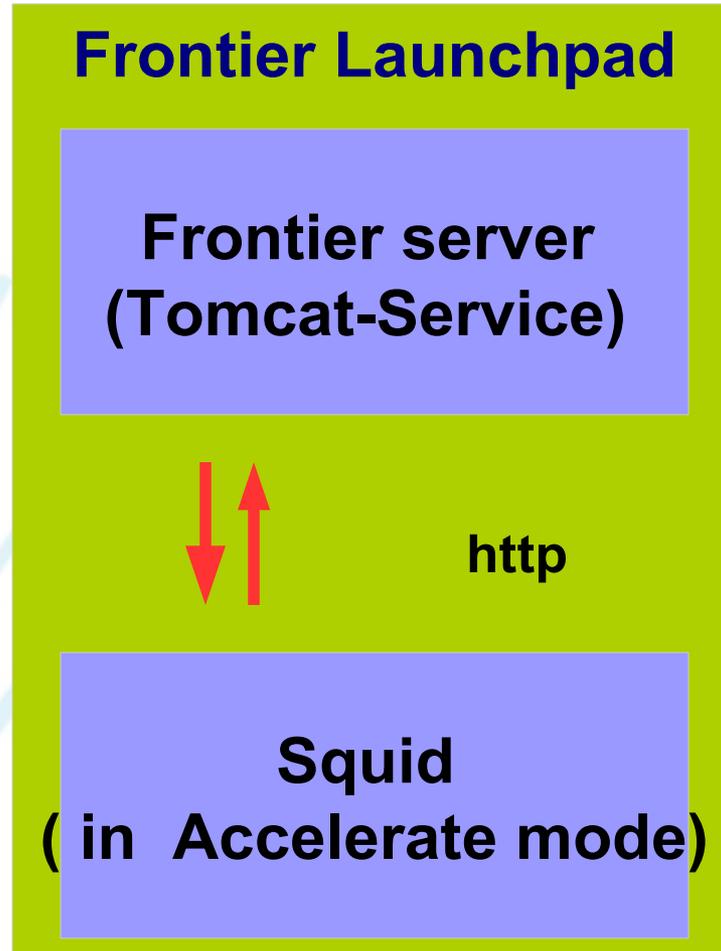


ATLR REPLICA



Database Internal protocol

Squid in Accelerate mode: this squid can cache only one site (frontier-Tomcat Service) and can't be use as a general cache squid



▶ Frontier System^{4 (?)}:

