



ATLAS Operations


Ueda I
GDB 2011 Sep

ATLAS activities since March 2011



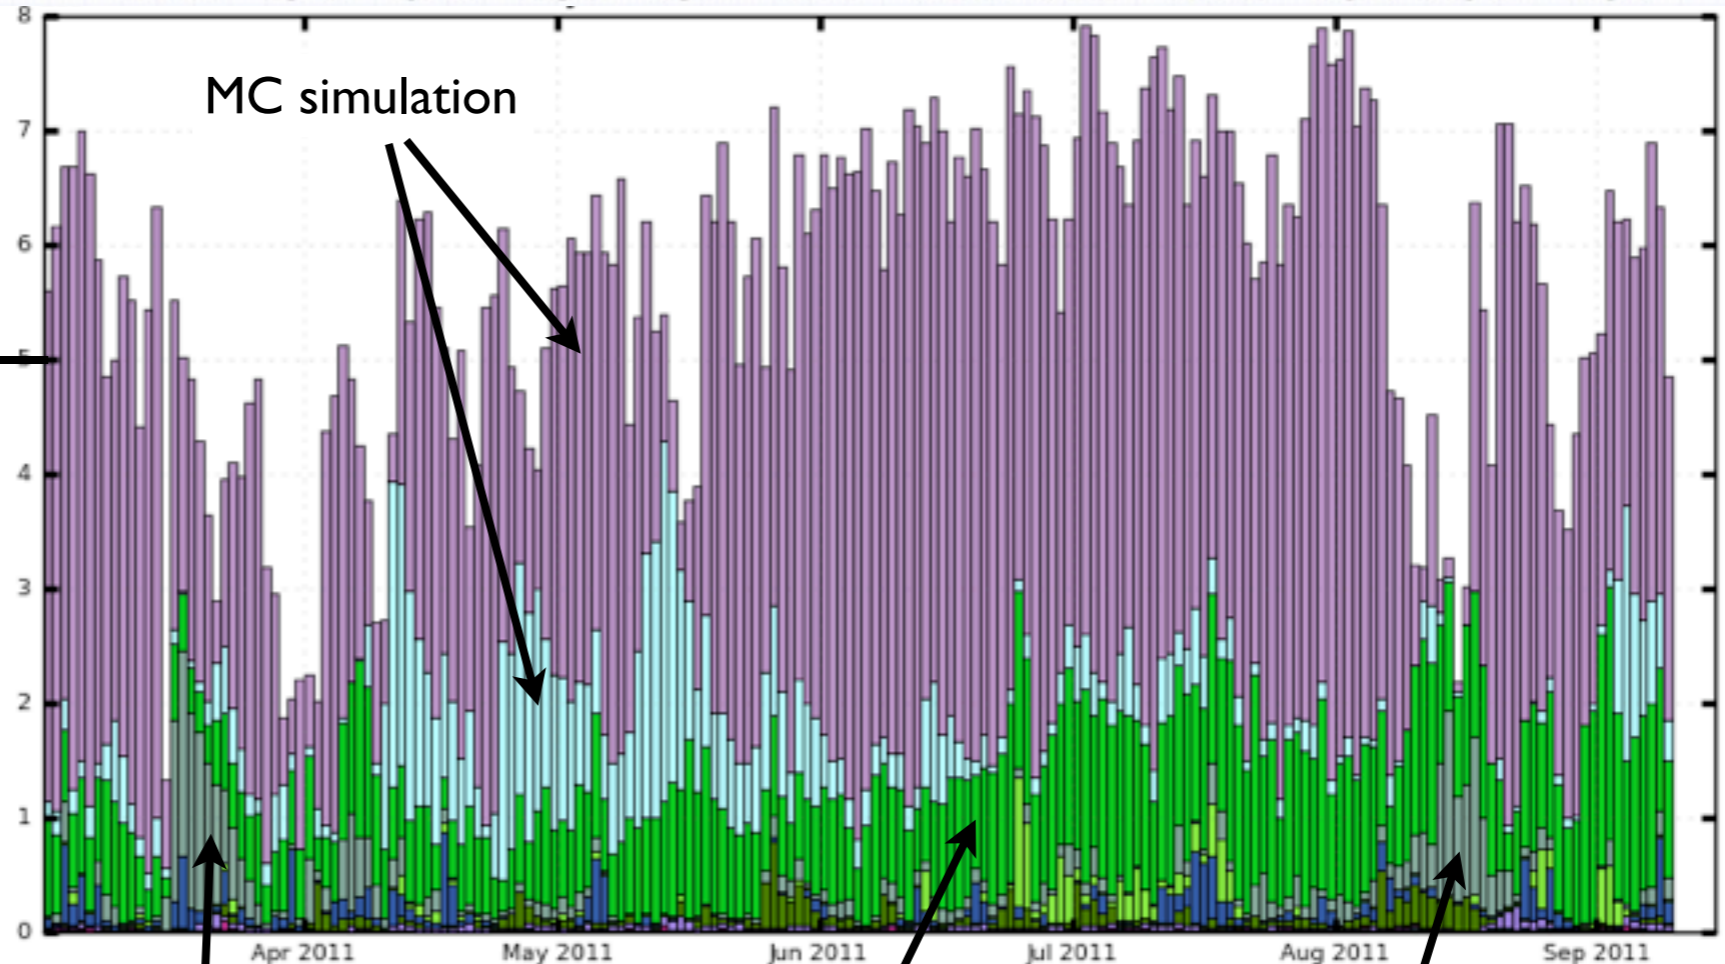
ATLAS DAQ write RAW into T0

300 MB/s



ATLAS grid jobs wall clock consumption


5×10^9 sec



reprocessing

User analysis

reprocessing




RAW data compression at Tier-0 since July



ATLAS Tier-0 write
RAW into Castor

300 MB/s




- Note: not much compression rate expected on TAPE

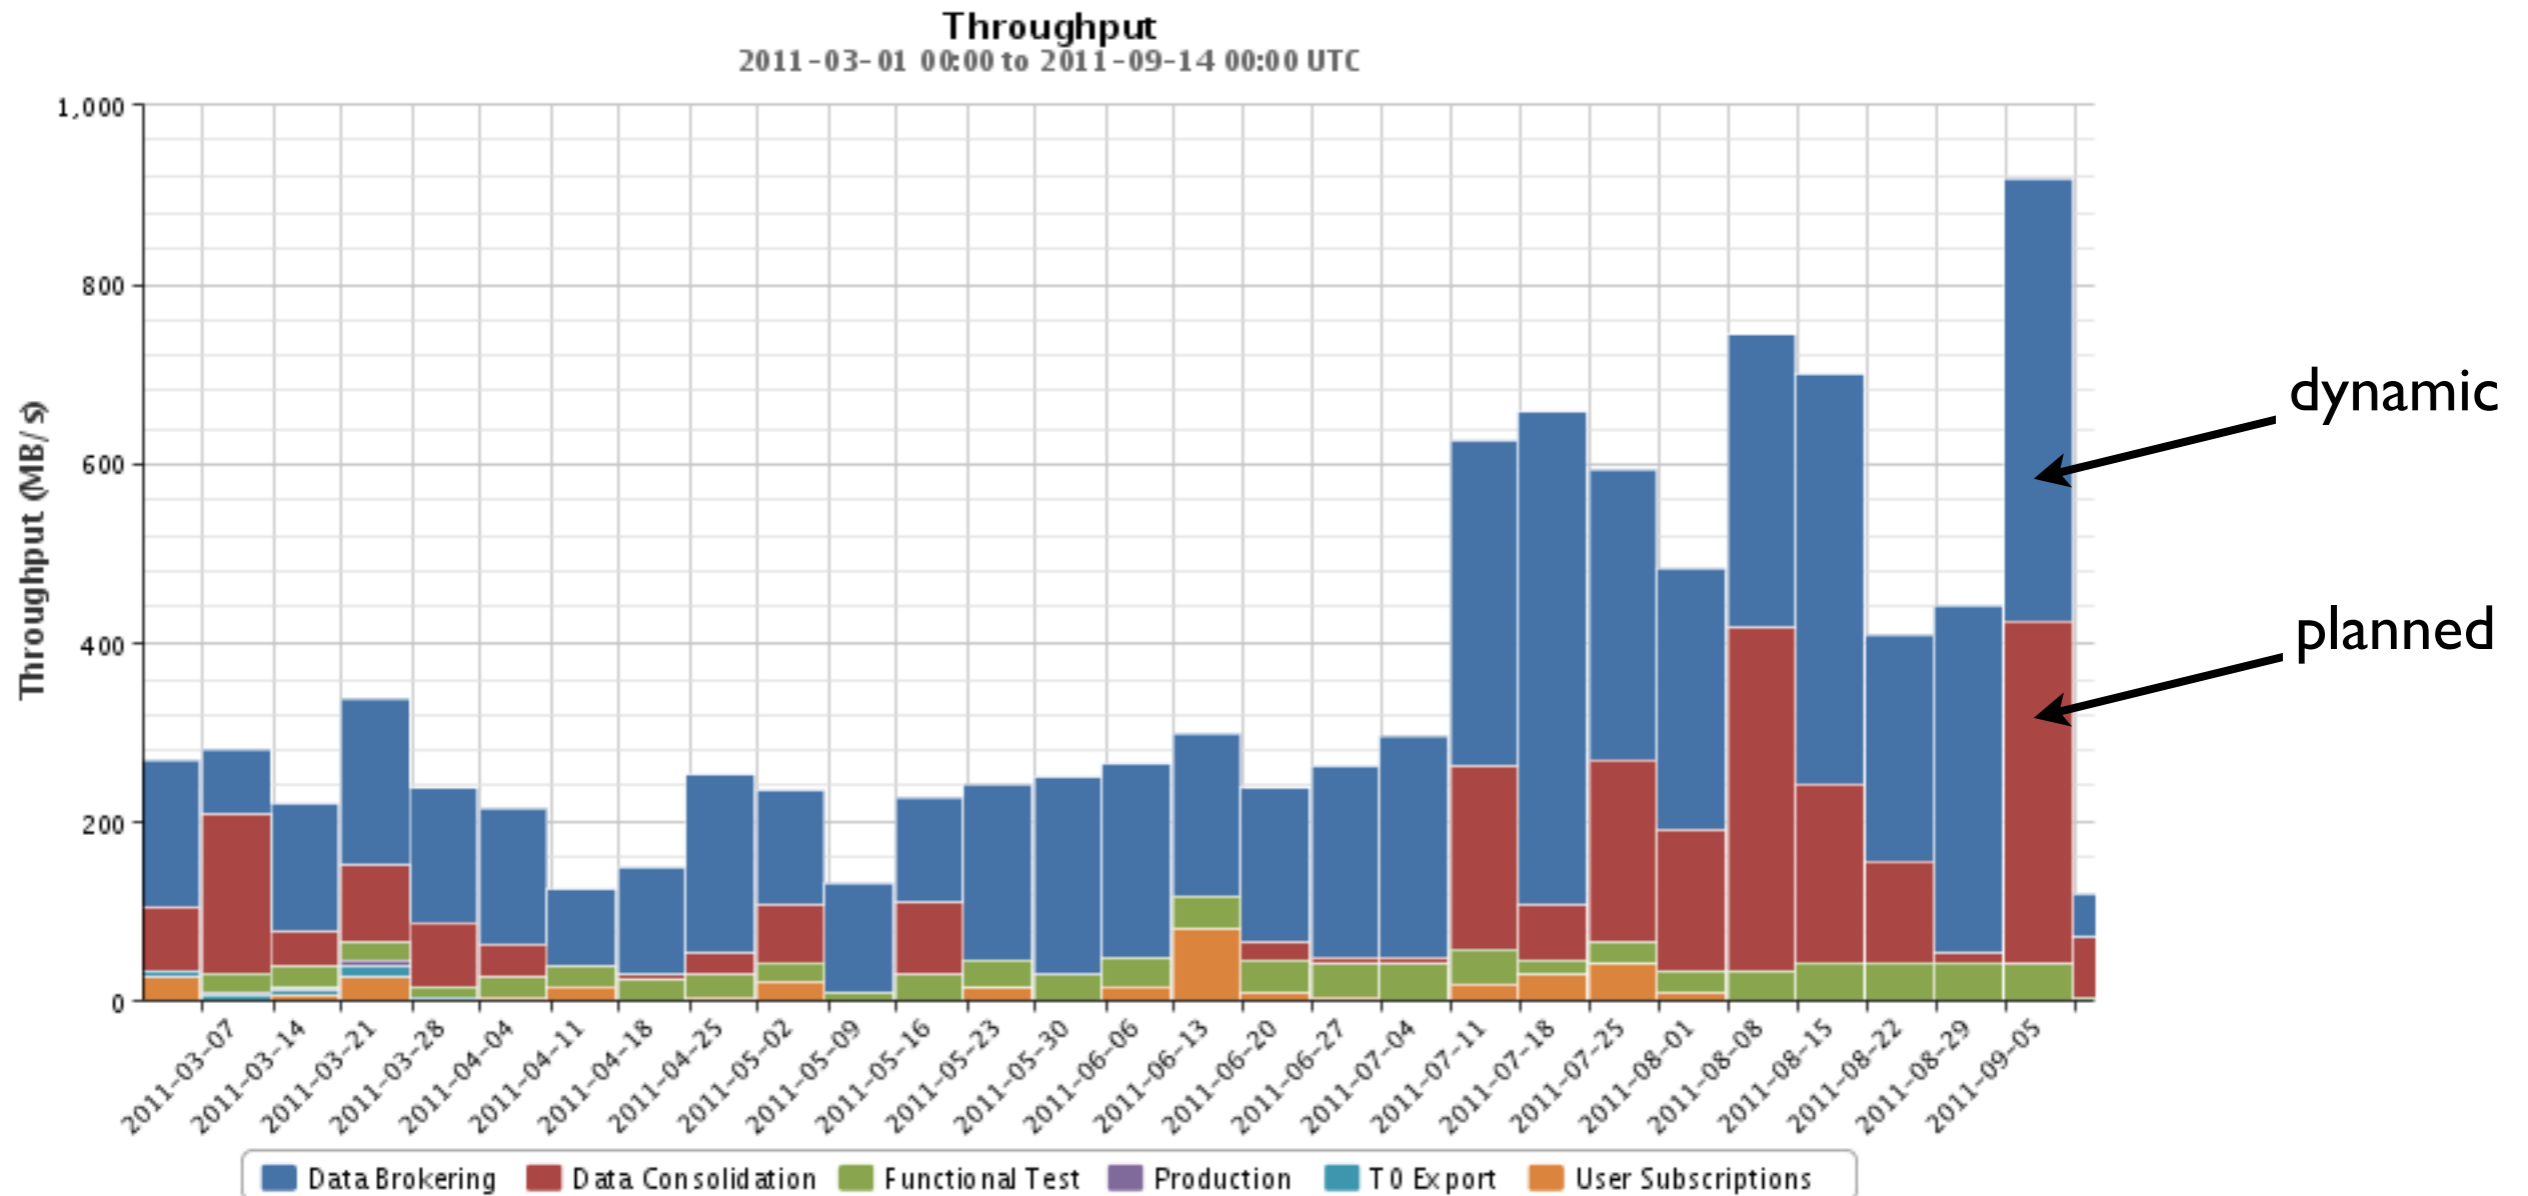
Still no reduction in overall data throughput (ESD are larger)

ATLAS Tier-0 export
to T1s and calibT2s

500 MB/s



Tier-2 data distribution being revised



Planned distribution of NTUPLE and DAOD with pre-defined share (since Jul)

Planned distribution of AOD with pre-defined share (since Sep)

Dynamic data placement with job brokerage algorithm (since last year, algorithm revised in Apr)

Dynamic data placement with pre-defined share (since Jul)

T2Ds

Tier-2 sites that can directly transfer any size of files with the other T1s than the one they are associated to

- First list defined in March
 - ▶ T1SC Mar 17 (<http://indico.cern.ch/contributionDisplay.py?contribId=1&confId=131670>)
- Revised in August
 - ▶ T1SC Sep 01 (<http://indico.cern.ch/contributionDisplay.py?contribId=10&confId=153062>)

T1s have been requested to configure the FTS channels accordingly

- https://twiki.cern.ch/twiki/bin/view/Atlas/DDMOperationsFTS#T2Ds_channels

T2Ds are candidates for

- multi-cloud production sites
- primary replica repository sites (T2PRR)

LHCONE

- Sites joining LHCONE may change in their network connectivity
- Sites are requested to pre-notify ATLAS before joining LHCONE
 - ▶ so that we can do test transfers to compare the situation before and after

LFC consolidation



ATLAS LFC at CERN migrated a new system (Aug 29)

- backend DB from LCGR to ADCR – modified schema for partitioning
- frontends prod-lfc-atlas.cern.ch (3 hosts in load balancing)

Next target: **SARA** LFC contents to be merged into CERN LFC

IN2P3-CC second merge candidate (merge has been tested)

INFN-T1 prefers as early merge as possible

Sites are welcome to express their preference in planning the merge

LFC consolidation at CERN

- Goal :
 - All LFCs aggregated in a single LFC at CERN
 - Read-only replica in another site (probably BNL)
- Reason :
 - ATLAS experienced LFC downtime over few weeks (Summer 2010)
 - Current LFC model: single point of failure
 - all stored data within cloud can be inaccessible
- Status :
 - Discussion between ATLAS and WLCG/CERN to validate the merging procedure
 - Identify possible inconsistencies between catalogs before merging
- Timescale :
 - One LFC migration at a time (should be done within days)
 - Expected to be done during spring/summer 2011

10 9 March 2011

Reprocessing

mid Mar – end Mar : Heavy Ion Reprocessing Campaign

- HI data taken in 2010

late Mar – mid Apr : Muon Reprocessing Campaign

- Reconstruct the 2010 data with muons with the latest software for the muon performance paper

early Apr – late Apr : 2011 Fast Reprocessing Campaign

- Reprocess the first data of 2011 with updated conditions (after looking at the data). The idea that the updates will be ported to the Tier-0 before more data is being taken.

mid July – end Aug : 2011 Reprocessing Phase-I

- for the data until June 30

early Sep – ongoing : 2011 Reprocessing Phase-II

- for the data Jul 1 – Aug 24

Issues during the recent reprocessing



Staging from tape

- GGUS:73490 – INFN-T1
- GGUS:73645 – TRIUMF
- GGUS:73678 – FZK
- GGUS:73680 – PIC
- GGUS:73683 – IN2P3-CC

High failure rate at IN2P3-CC

- DCCP error (*pathological* connections)
- Overload on the worker nodes
- **workaround deployed / planned**
- **express reprocessing** prior to the phase-II bulk was done with a **reduced share** (to finish the tasks in time)
- need to watch the on-going reprocessing

StoRM



ATLAS suffered from many incidents recently

- **INFN-T1**: **GGUS:68243** (2011-03-03 – 04), **GGUS:73054** (2011-07-29 – 31), **GGUS:73068** (2011-07-31), **GGUS:73236** (2011-08-06), **GGUS:73422** (2011-08-12), **GGUS:73452** (2011-08-13 – 14), **GGUS:73456** (2011-08-14 – 15), **GGUS:73619** (2011-08-18 - 24)
- **INFN-MILANO**: **GGUS:68829** (2011-03-21), **GGUS:71124** (2011-05-31), **GGUS:71622** (2011-06-17), **GGUS:72275** (2011-07-06), **GGUS:72879** (2011-07-22 – 23), **GGUS:72986** (2011-07-27), **GGUS:73059** (2011-07-29 – 30), **GGUS:73254** (2011-08-08)
- **INFN-GENOVA**: **GGUS:72937** (2011-07-26 – 27),
- **IFIC**: **GGUS:72891** (2011-07-24)
- **LIP-Coimbra**: **GGUS:73243** (2011-08-07 – 10)
- **QMUL**: **GGUS:71418** (2011-06-10), **GGUS:72075** (2011-06-29 – 07-26), **GGUS:72604** (2011-07-13 – 25)
- **TECHNION**: **GGUS:72998** (2011-07-28 – 08-16)
- **WEIZMANN**: **GGUS:72472** (2011-07-08), **GGUS:72584** (2011-07-13), **GGUS:72783** (2011-07-19 – 21),

T1SC (Sep 01, <http://indico.cern.ch/conferenceDisplay.py?confId=153062>)

- INFN-T1 has been running reliably before the upgrade to StoRM 1.7 and has been very unstable after
- QMUL has been very unstable before the upgrade to 1.7 and is very reliable now

GGUS Alarm



Several cases where ALARM did not work, especially conversion from TEAM to ALARM did not work (alarm mails did not reach the final destination)

- CERN – GGUS:71471 (2011-06-12)
- CERN – GGUS:72890 (2011-07-24)
 - ▶ temporarily limited access to the Alarm-mail interface in GGUS. **fixed**
- SARA – GGUS:72713 (2011-07-17)
 - ▶ Using a script monitoring for ALARM tickets – but contact address sara-matrix was not picked up. Also the list of authorised alarmers is not up to date and site would welcome an update. (WLCGDailyMeetingsWeek110718) — **is this done?**
- INFN-T1 – GGUS:73054 (2011-07-29) GGUS:73236 (2011-08-06)
 - ▶ **solution** GGUS:73317 (2011-08-08).

ALARMers

- ATLAS alarmers are the members of **/atlas/alarm**
- Is the list consistently propagated everywhere? (voms, ggus, T1s)

GGUS alarm tests — test the real use cases?

- conversion from TEAM to ALARM
 - ▶ /atlas/team member (non-Alarmmer) submit a TEAM ticket
 - ▶ /atlas/alarm member convert it to ALARM
- how to make sure every /atlas/alarm member can send, or convert to, ALARM?

Certificate not updated



We sometimes observe problems in certs update at some services

- usually failures in automated update, i.e. glitches.
- but below may need some attention

CA certificate GGUS:73776, GGUS:73794

- affected services
 - ▶ GGUS
 - ▶ SRM for CERN castor (srm-atlas.cern.ch)
 - ▶ LFC (prod-lfc-atlas-local.cern.ch)
 - ▶ BNL VOMS (vo.racf.bnl.gov)
- Why the CA cert renewal did not propagate to so many places?

VOMS certificate (WLCGDailyMeetingsWeek110425)

- ATLAS requested that sites follow the forthcoming instructions so that update of the voms certs rpm would be no longer required in future.
VOMSLSCfileConfiguration
- Has this been followed up?

Security



SSC5 – May

- A user certificate was compromised,
 - ▶ Revocation of this certificate was requested.
 - ▶ The certificate was de-authorized in ATLAS VOMS.
- The pilot proxy was used to execute the malicious payload on multiple sites.
 - ▶ The full proxy has not been compromised.
 - ▶ The limited proxy on sites could have up to 96 hours of validity.
 - ▶ To limit any possible damage, this certificate was revoked

Security incident – July

- One of the ATLAS Tier-2s has possibly been compromised.
- Investigation together with site, WLCG and CERN security contacts.
- As precaution 2 ATLAS certificates used for pilot have been revoked and re-requested.
- Only short living voms-proxy of those certificates have been exposed to the site, certificates themselves are safe

glexec

ATLAS has presented a document to MB

- ▶ <https://indico.cern.ch/materialDisplay.py?contribId=8&materialId=0&confId=115406>
- Different Athena release versions behave differently in the presence of gLexec
 - ▶ it has been rather complicated to debug
- While gLexec covers some vulnerabilities, it creates very serious ones in case of compromised credentials.
 - ▶ The gLexec mechanism does not provide a machinery to propagate user proxies to the worker nodes — ATLAS would need to use the MyProxy server
 - ▶ A stolen pilot credential would therefore be able to fetch ATLAS users proxies
 - ▶ **the number of revocations will be much larger than just the generic pilot proxy in non-gLexec scenario**
- Several tests run at various sites have shown an unstable behavior of the sites coupled with the usage of gLexec
 - ▶ Invoking gLexec might work at a given period of time and not work anymore at a later stage

The conclusion is that ATLAS does not intend to invoke gLExec on the worker node and proposes the following scenario:

- Any Grid site should ban the ATLAS pilot DN or if necessary the entire ATLAS VO in case there is a suspicion of compromised credentials or illegal usage of resources at the site.
 - ▶ The banning mechanism should eventually happen via ARGUS, so that ATLAS services can cooperate in the banning of the illegal activity from their side.
- ATLAS will offer all needed forensics in case of a security violation even if suspect.
 - ▶ In general ATLAS will offer forensics about user activity at sites at the request of the site.
 - ▶ Tools and documentation can be provided so that the investigation can be carried on by the site itself.
 - ▶ **The traceability of ATLAS jobs through the Panda system has been demonstrated during the Security Challenges 4 and 5.**

Miscellanea



VOMS server migration (gLite 3.2 to EMI 1) affected ATLAS production

- Problem traced back to a bug in GridSite ([GGUS:71190](#), 2011-06-03)
- ATLAS had to install lcg-vomscerts-6.4.0-1.sl4 to the service machines

Downtime on castorpublic — srm-atlas.cern.ch marked as affected

- ATLAS does not use castorpublic
- “dependency comes from automatic testing, it has been like this for a couple of years already. **Plan to take this out completely at the beginning of September**, after the many interventions at the end of August” — WLCGDailyMeetingsWeek110815

CVMFS

- ATLAS will largely rely on cvmfs for the sw distribution and will need it well supported
- will prepare formal request for support

SSB



Getting into shape

- ▶ <http://dashb-atlas-ssb.cern.ch/dashboard/request.py/siteview?view=Shifter%20view>

ATLAS dashboard

Index Expanded Table

Show 200 entries Copy Print Save view: Shifter view Search...

| Site Name | Site Info | | | DataManagement | | | Panda Efficiency | | | | | Activity status now: Included-Excluded sites | | | | | |
|-------------------------|-----------|------------|------------|----------------|----------------|-----------|-------------------------|----------------------|-------------------|--------------------------|-----------------------|--|-----------------------|------------------------|---------------|---------------|---------------|
| | Downtime | Tier Level | Cloud Info | DDM 4h [%] | SRM SAM 12 [%] | AFT Panda | Prod Efficiency 12h [%] | Prod Failed Jobs 12h | Prod All Jobs 12h | Analy Efficiency 12h [%] | Analy Failed Jobs 12h | Analy All Jobs 12h | panda prod status NEW | panda analy status NEW | DDM DA Status | DDM DP Status | DDM DT Status |
| BNL-ATLAS | ACTIVE | Tier-1 | US | 99 | 100 | 100 | 100 | 31 | 11873 | 98 | 481 | 25801 | online | online | online | online | online |
| IN2P3-CPPM | ACTIVE | Tier-2 | FR | 100 | 65 | 100 | 100 | 3 | 1613 | 96 | 40 | 1034 | online | online | online | online | online |
| IN2P3-LAPP | ACTIVE | Tier-2 | FR | 100 | 81 | 100 | 100 | 0 | 378 | 98 | 58 | 1335 | online | online | online | online | online |
| INFN-FRASCATI | ACTIVE | Tier-2 | IT | no.data | 100 | n/a | 100 | 1 | 748 | 99 | 3 | 311 | online | online | online | online | online |
| INFN-MILANO-ATLASC | ACTIVE | Tier-2 | IT | 100 | 100 | 100 | 100 | 0 | 1079 | 96 | 84 | 1953 | online | online | online | online | online |
| prague/og2 | ACTIVE | Tier-2 | DE | 91 | 100 | 100 | 100 | 0 | 751 | 98 | 71 | 6167 | online | online | online | online | online |
| RO-07-NIPNE | ACTIVE | Tier-2 | FR | no.data | 100 | n/a | 100 | 1 | 496 | 98 | 3 | 179 | online | online | online | online | online |
| Taiwan-LCG2 | ACTIVE | Tier-1 | TW | 100 | 100 | no-test | 98 | 65 | 4334 | 98 | 22 | 1033 | online | NoQueue | online | online | online |
| TR-10-ULAKBIM | ACTIVE | Tier-2 | NL | 100 | 100 | 100 | 100 | 4 | 667 | 97 | 21 | 608 | online | online | online | online | online |
| UKI-NORTHGRID-LANCS-HEP | ACTIVE | Tier-2 | UK | 100 | 100 | 100 | 100 | 1 | 2376 | 97 | 59 | 2001 | online | online | online | online | online |
| UKI-NORTHGRID-LIV-HEP | ACTIVE | Tier-2 | UK | 100 | 100 | 100 | 98 | 5 | 812 | 100 | 2 | 476 | online | online | online | online | online |
| UKI-NORTHGRID-SHEF-HEP | ACTIVE | Tier-2 | UK | 100 | 100 | 100 | 100 | 0 | 814 | 97 | 81 | 1818 | online | online | online | online | online |
| wuppertalprod | ACTIVE | Tier-2 | DE | 72 | 100 | 100 | 99 | 5 | 953 | 97 | 65 | 2182 | online | online | online | online | online |

to monitor the “site status”

to check the site classification metrics (for T2D, T2PRR, ...)