



WLCG operations

ALICE T1-T2 Workshop
Derby
April 17-19, 2018

Maarten Litmaath
CERN-IT

v1.0

Central services



- Annual power cut tests on Sat Jan 6 had fallout
 - A UPS main power cable melted and shorted out a fuse
 - A few hosts lost power, 1 API server could not be booted
 - A very large number of jobs failed during ~3 hours
- Incident Jan 10 afternoon
 - Slowness and file system corruption due to a bad disk
 - Recovery of the main impacted services took ~20 hours
 - Thanks to heroic efforts by Costin and Miguel !
 - Much less grid activity and most jobs failed during that time
 - Further work behind the scenes afterwards
- AliEn improvements → see Miguel's talk

Site evolution



- New site: **SARFTI** (Sarov, Russia)
 - Thanks to Andrey Zarochentsev for setting it up!
- University of **Derby**: almost there!
- New EOS instances at various sites
 - A second EOS workshop was held Feb 5-6
 - CERN: new *eosalicedaq* SE for 2018 data taking
 - Data copied to CASTOR through ALICE workflow
- Growth – 2018 pledges
 - More disk space is welcome in particular!

CNAF incident & recovery




- Down Nov 9 until late Jan due to a huge flooding incident
 - More details in INFN sites presentation
- First jobs this year ran late Jan
- 68 M disk SE files were unavailable until mid Feb
 - None were lost !
- Recovery of the wet tapes in progress, looking good so far (not critical)
- *We thank and congratulate CNAF with the success of their careful efforts!*




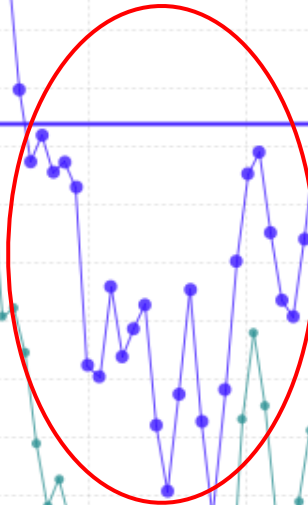
High activity !

Running jobs per user

Dec 25: 175k



Temporary lack of MC



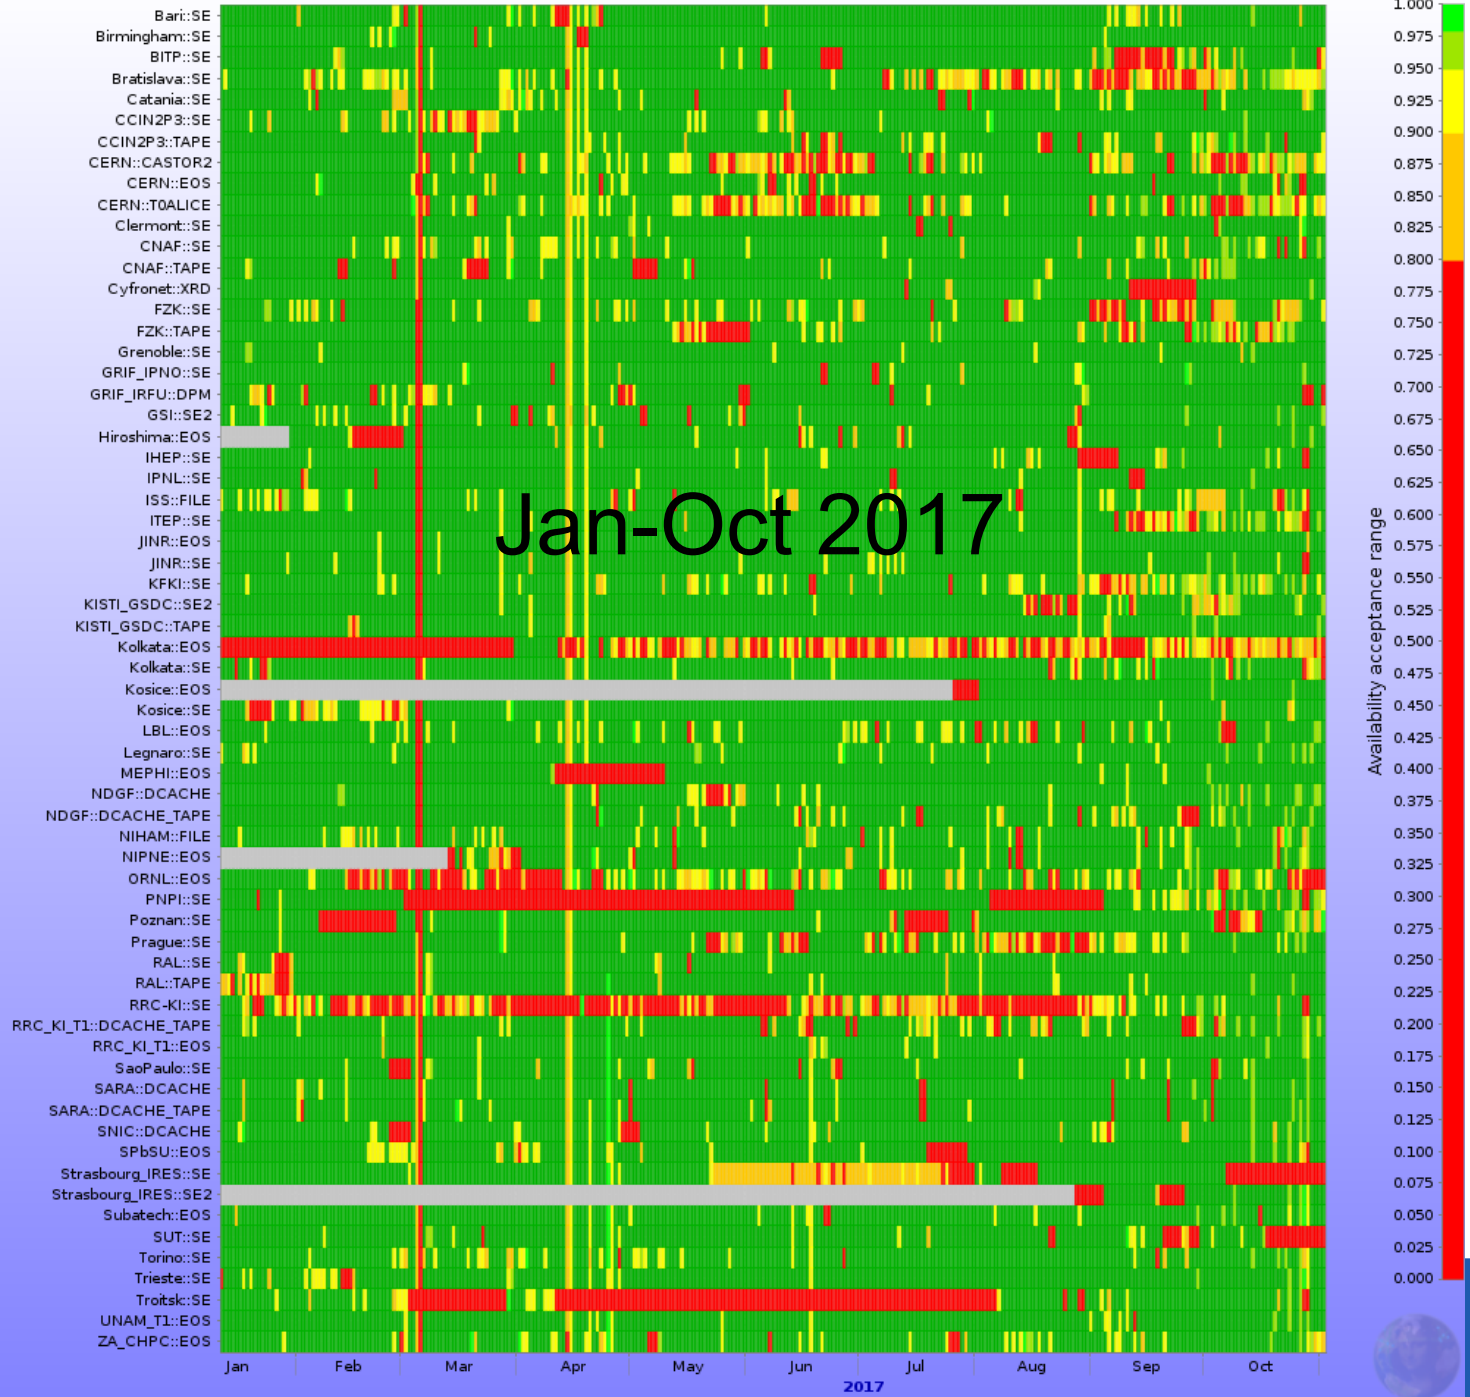
High analysis activity in recent weeks in preparation for QM, May 13-19

Storage



- CERN
 - CASTOR, EOS: mostly stable, intermittent issues
 - The CASTOR default pool was enlarged with disk servers borrowed from EOS
 - After the 2017 data taking CASTOR was reconfigured for reco up to ~50k jobs
- Many SEs too often below 95% availability
 - Target for good job efficiencies & success rates
 - Please keep an eye on your SE, thanks!

AliEn SEs availability for reading

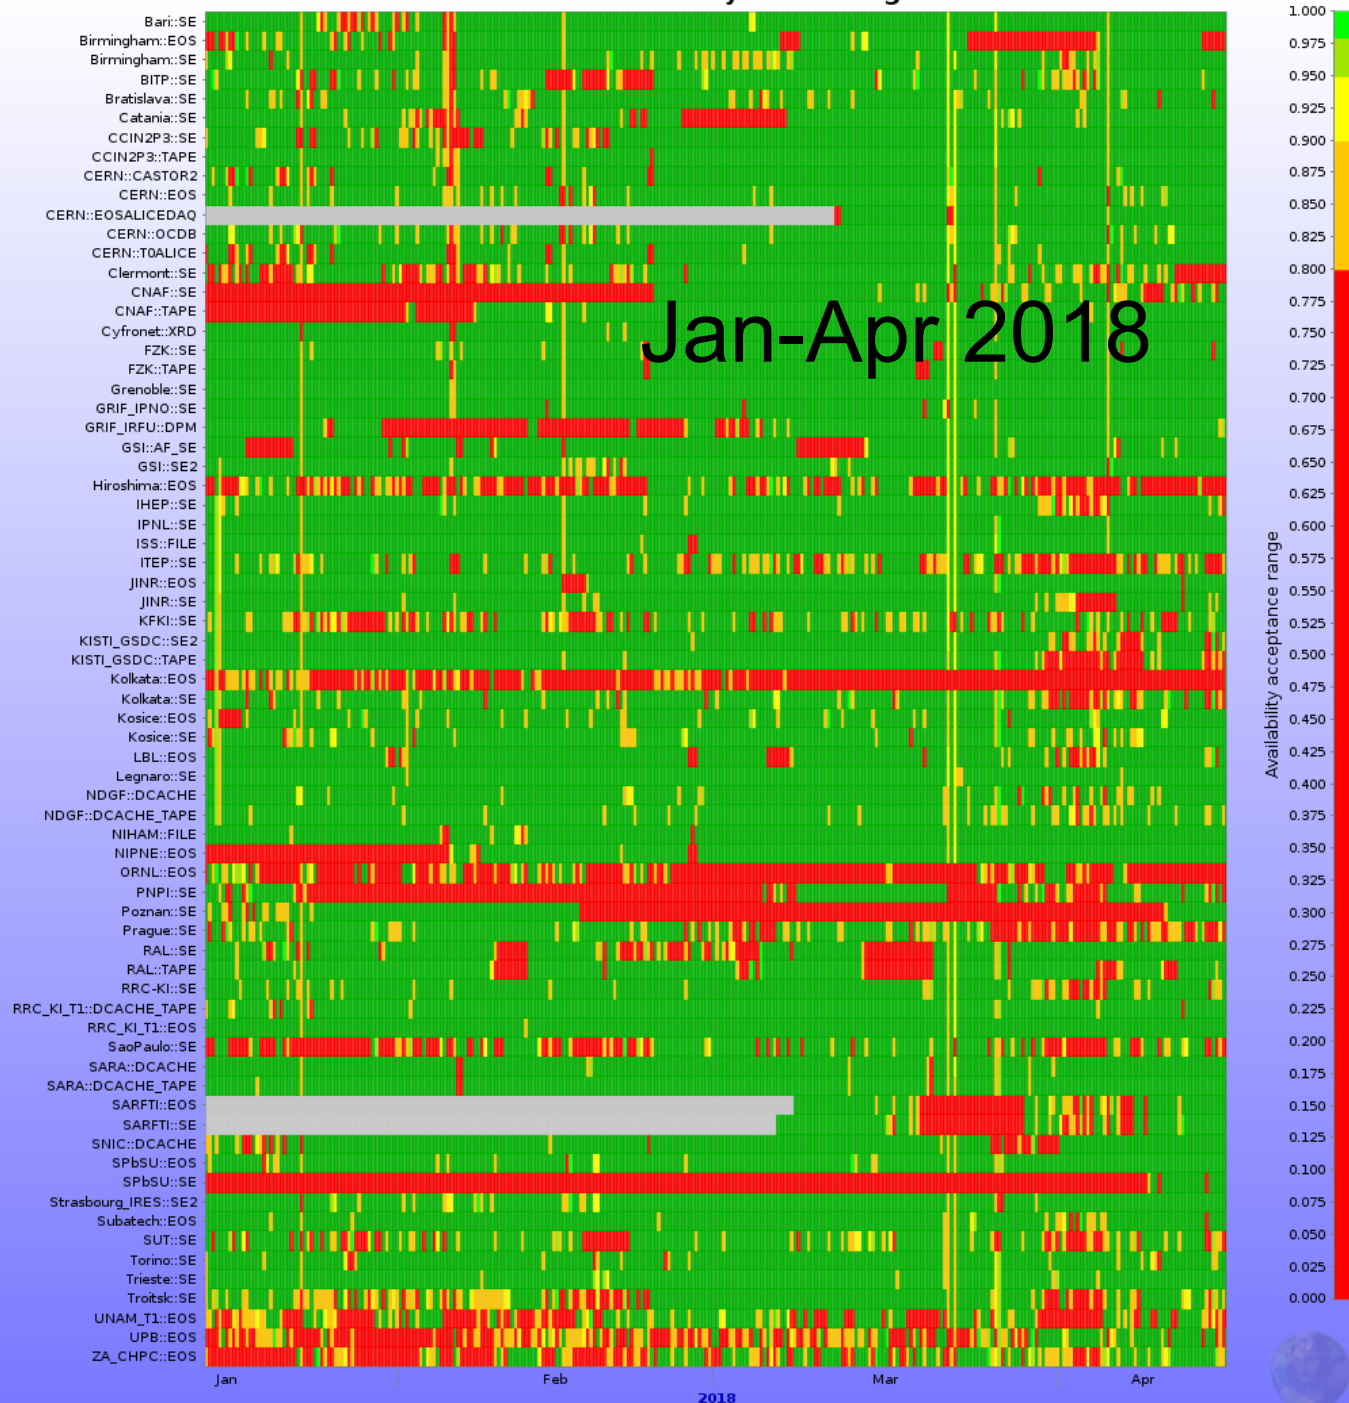


Jan-Oct 2017

Availability acceptance range



AliEn SEs availability for reading



Issues at sites or with jobs (1)



- CERN: further debugging of issues affecting HTCondor cluster as it kept growing
 - 60 → 100% of our resources
 - CE configurations were further improved for robustness and scalability
 - More CEs were added as well
 - Some bugs were fixed
- AliEn HTCondor module overhauled for various improvements
- HTCondor job status monitoring frequency lowered x 10
- Large reductions in I/O and CPU loads to the benefit of the VOBOXes, CEs and HTCondor batch system

Issues at sites or with jobs (2)



- RRC-KI-T1 failing many analysis trains between end of Dec and March 20
 - Despite many investigations not understood (?)
 - Network problem somewhere...
- ORNL EOS instabilities
 - More in the US sites report
- *Meltdown* and *Spectre* vulnerabilities
 - Mitigations released for Linux and other OS
 - Performance losses amounted to a few % at most

Middleware



- **CentOS/EL7** on the rise
 - Most service types available in UMD-4
 - CREAM and UI since Dec (WN was there earlier)
 - *Note: CREAM needs to be configured with Puppet and does **not** provide support for configuring a batch system*
 - WLCG VOBOX available since Feb
- The 4 LHC experiments can all run on CentOS/EL7
 - Various sites already upgraded their WN
 - It is also desirable for *Singularity*
 - Containers (or VMs) could still provide SL6
- UMD-4 only has officially supported products
 - Products like *Maui* remain available in the Community Repository
- UMD-3 to be phased out soon

Containers (1)



- New tool to launch containers gaining momentum in our community and beyond: Singularity
 - CMS require it at ~all their sites
 - Provide SL6 environment on CentOS/EL7 WN
 - And **isolate** each user payload from other processes → the plan is to let our Job Agents use it where available
 - No firm timeline yet; linked to **jAliEn** deployment
 - By default it needs to be *setuid root* because the kernel requires root for *mount namespace* operations
 - Since EL7.4 *user* mount namespaces are configurable through a boot option + *sysctl* setting
 - Mind that RedHat consider that feature a technology preview and recommend not to use it in production...

Containers (2)

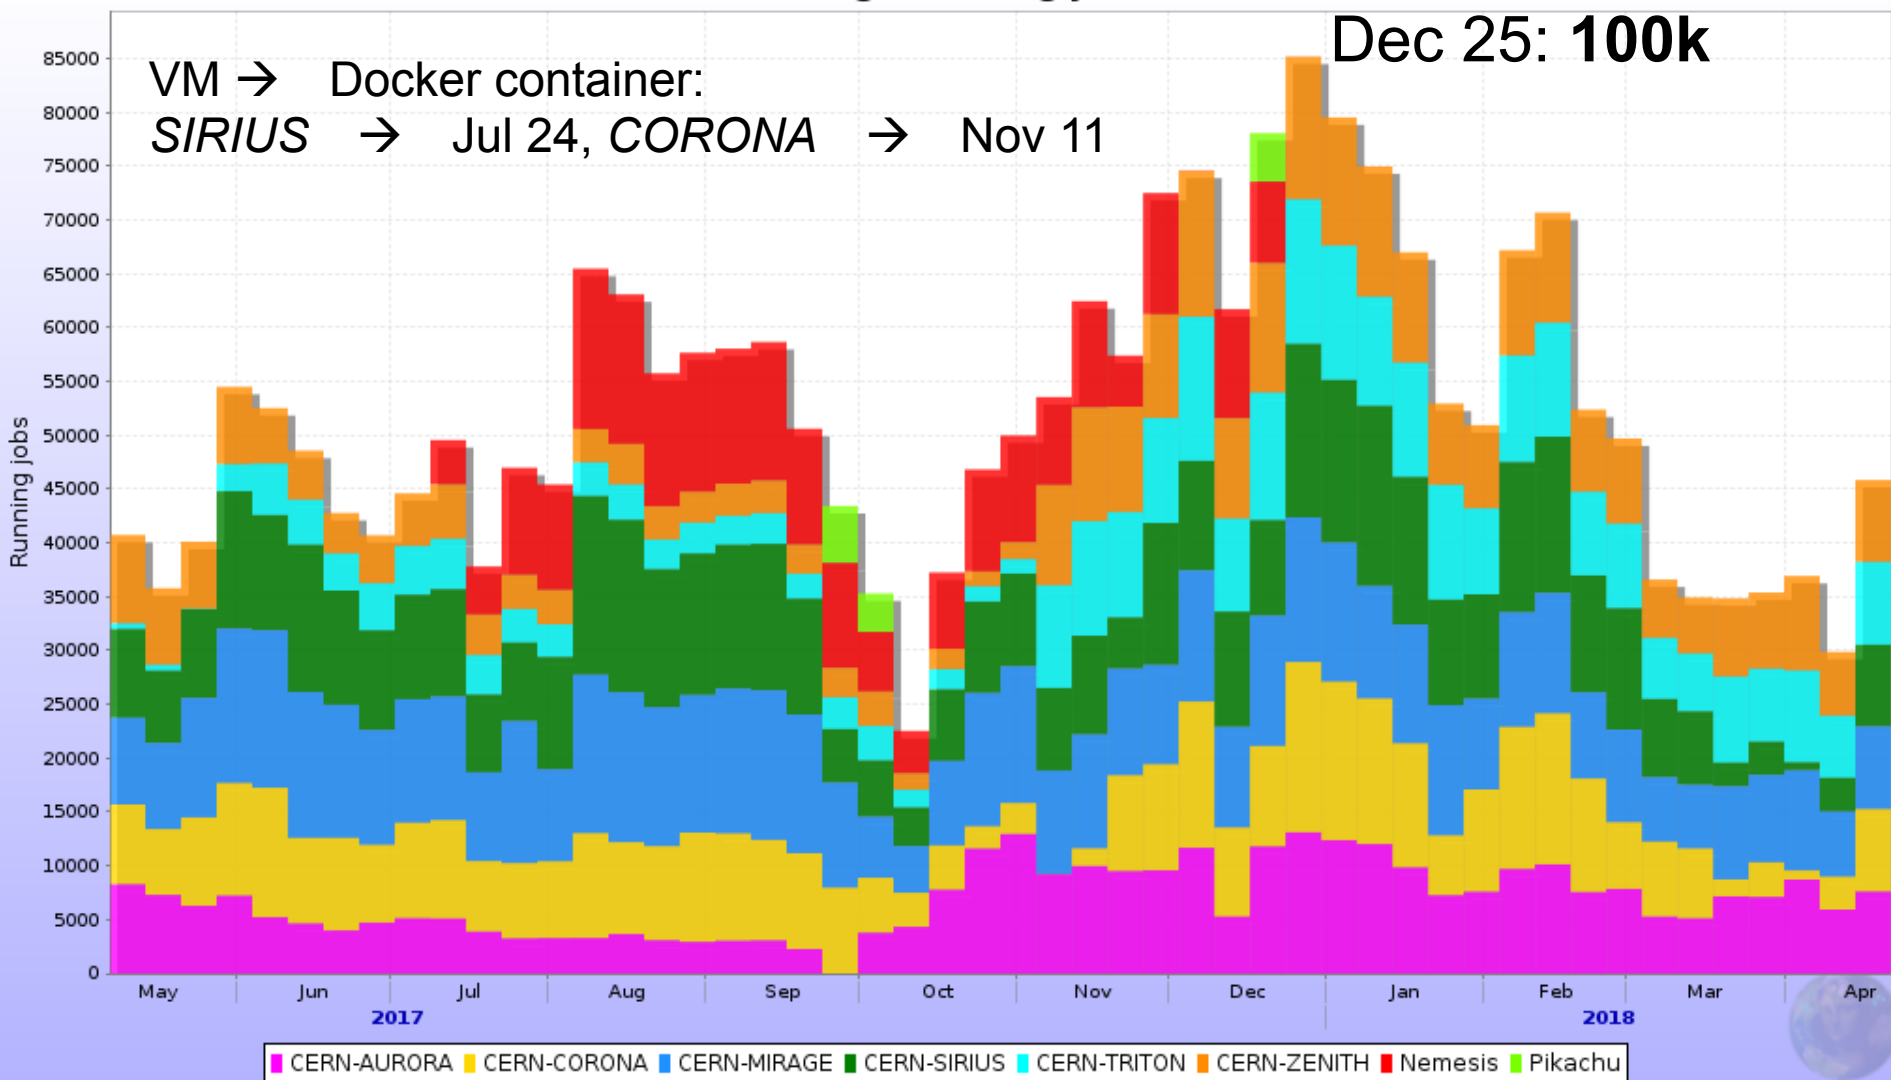


- They can also be used for easy deployment of pre-packaged services
 - Development hosts
 - Build and validation services
 - WLCG VOBOX
 - ...
- A Docker container for a WLCG VOBOX was set up by [Maxim Storetvedt](#) (Western Norway Univ. of Applied Sciences)
 - As an alternative to VMs that were suffering instabilities... 😞
 - See his [talk](#) tomorrow
- It was then configured for a new test “site” called [Nemesis](#) 😊
- After stress tests (in production!) and various improvements we now run 2 VOBOX containers in production on separate hosts
 - Alongside 4 VMs evenly spread over 2 other hosts

CERN virtual sites



Average running jobs



CREAM+LSF → HTCondor: *AURORA* → Oct 1, *TRITON+ZENITH* → Nov 1



CERN cloud extensions



- Oracle Bare Metal Cloud project
 - ~10k cores available to the 4 experiments for ~3 months
 - We used the virtual site *Phoenix* to drive them
 - Up to 9k concurrent jobs between Sep 9 and Dec 4
- Helix Nebula Science Cloud EU project
 - Up to ~20k cores available to the 4 experiments *plus* other participating projects until Dec this year
 - Virtual sites *Capella* and *Regulus*
 - Started mid March, low activity for now
- Non-trivial to configure/debug such resources
 - Network and other issues
 - Default easy use case: MC production

SAM



- **Reminder: SE test failures reduce the A / R!**
- Corrections are applied as needed
 - New policy defining thresholds as of this month, to reduce the number of recalculations
- Test job submission to the HTCondor CE has been added as of May last year

Xrootd reminder



- Sites should continue upgrades to **Xrootd** **>= 4.1**
 - Most sites have done that already, thanks!
 - Required for **IPv6** support
 - EOS → upgrade to v4 a.k.a. *Citrine*
- Communication via LCG Task Force list as usual for expert advice
- ALICE add-ons are available through **rpms**
 - <http://linuxsoft.cern.ch/wlcg/>
 - **Thanks** to Adrian Sevcenco!

Networks



- WLCG IPv6 deployment plans
 - CERN + T1 sites should have all production storage dual-stack by April this year
 - Most sites *ought* to have their production storage dual-stack by the end of Run 2 → early 2019
 - There will be a tail due to site constraints...
 - ALICE needs ~all storage to be dual-stack before *IPv6-only* computing can be used
- 3rd Asia Tier Center Forum held Oct 11–13
 - Hosted by KISTI – thanks!
 - Fostering network evolution in Asia – well done!

Tips for sites – thanks!



- Possible **issues** on VOBOX, CE, WN
 - CVMFS problem, CE not ready for jobs, myproxy running low, myproxy type wrong, ...
 - Absence of “system” library
 - HEP_OSlibs rpm helps avoid that
- Jobs may fail due to SE **problems**
- **Admins please check site issues page**
 - <http://alimonitor.cern.ch/siteinfo/issues.jsp>
- **Subscribe to relevant notifications**
 - <http://alimonitor.cern.ch/xml.jsp>