


Atlas Critical Services - CCRC 08 Post Mortem

Birger Koblitz, CERN IT & Atlas ADC






GS

Outline:

- Overview of the Atlas Critical Services
- Monitoring the Central Services
- Experiences in CCRC08
 - **The May 30th power cut**
- Conclusions (and what still has to be done)





GS

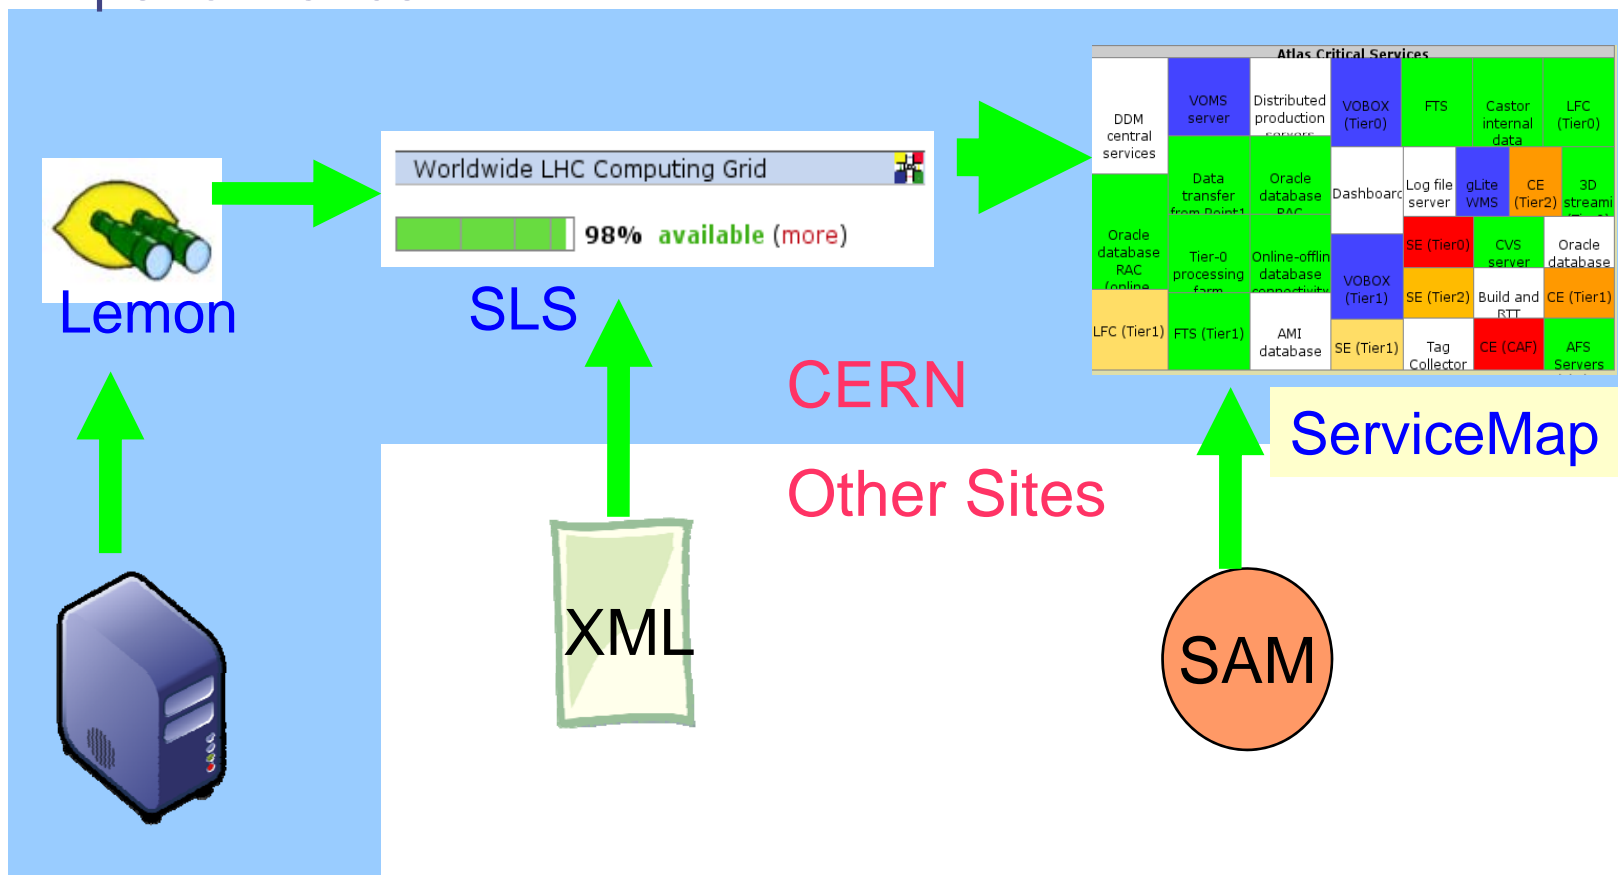
- **Scope: All ADC services needed for the centralized transfer of files, production, reprocessing**
 - Tier 0: Processing of data coming from the pit, Registration on Grid
 - Production System: ProdDB, Panda, Bamboo
 - Atlas Offline DBs
 - Site Services: Transfers out of CERN, within Clouds
 - Central Catalogues: Location of datasets, subscriptions of datasets to sites
 - AMI: Dataset metadata, SiteIndex, Deletion Service, Tracking Service
 - Atlas Dashboard
- **Complete list of all services at:**
<https://twiki.cern.ch/twiki/bin/view/Atlas/AtlasCriticalServices>



Monitoring in Atlas

GS

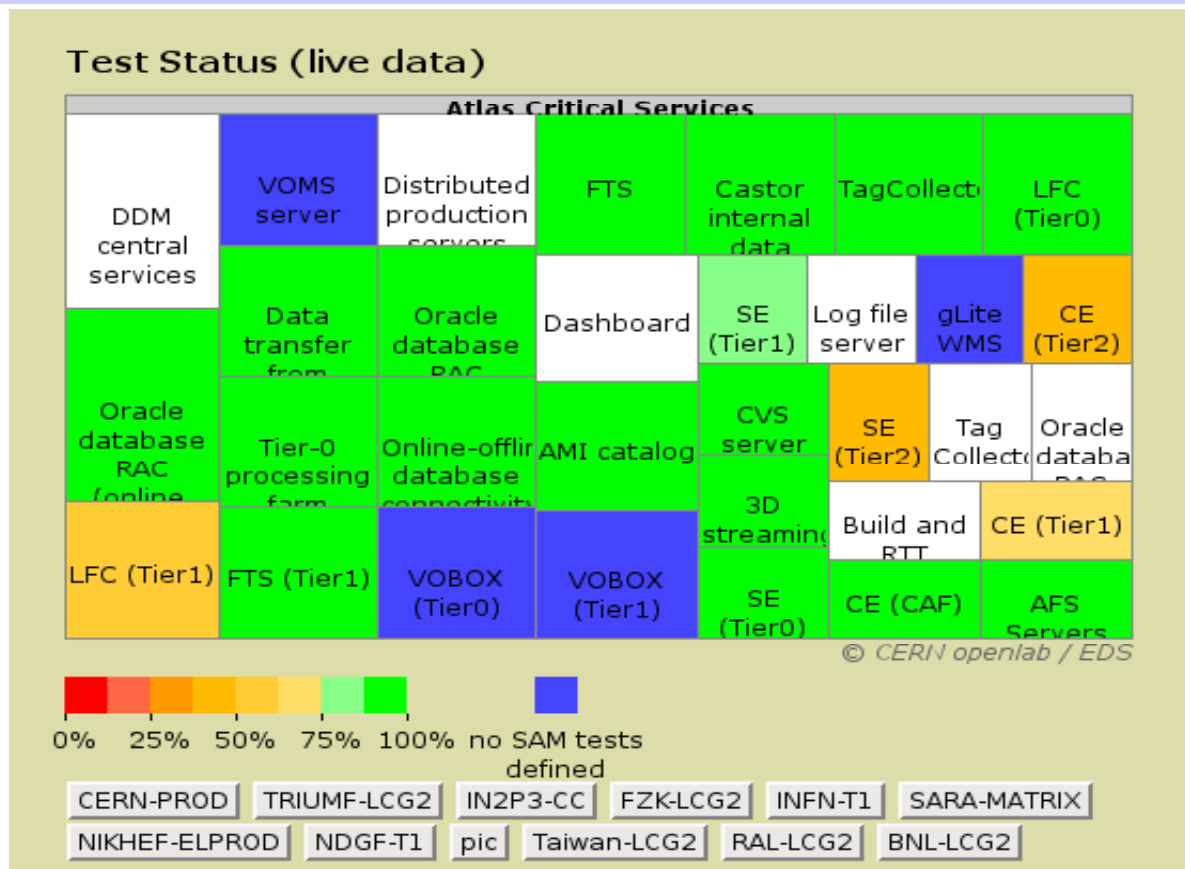
- Use Lemon for sensors, SLS and ServiceMap for monitoring of service status, Dashboard is used to follow service performance





GS

- <http://servicemap.cern.ch/ccrc08/servicemap.html>



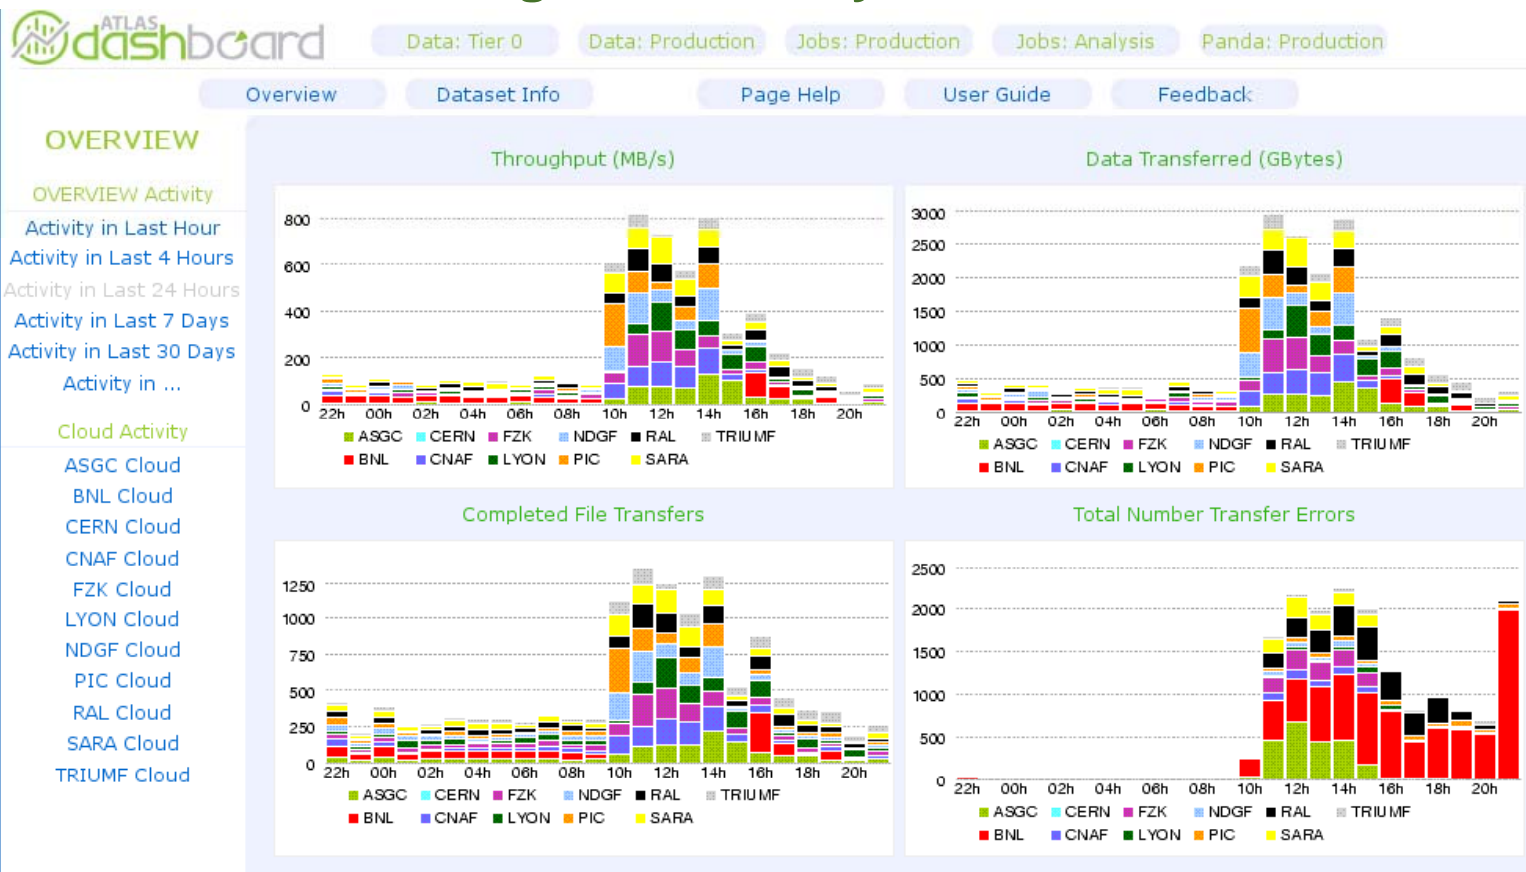
- Now have Lemon sensors and alarms for all critical services at CERN to alert experts
- Still need to connect all sensors to service map



Monitoring

GS

- Atlas Dashboard is main source of information for shifters and central operations
 - Used to monitor MC production and data transfers
 - We will be adding functionality to raise alarms to shifters

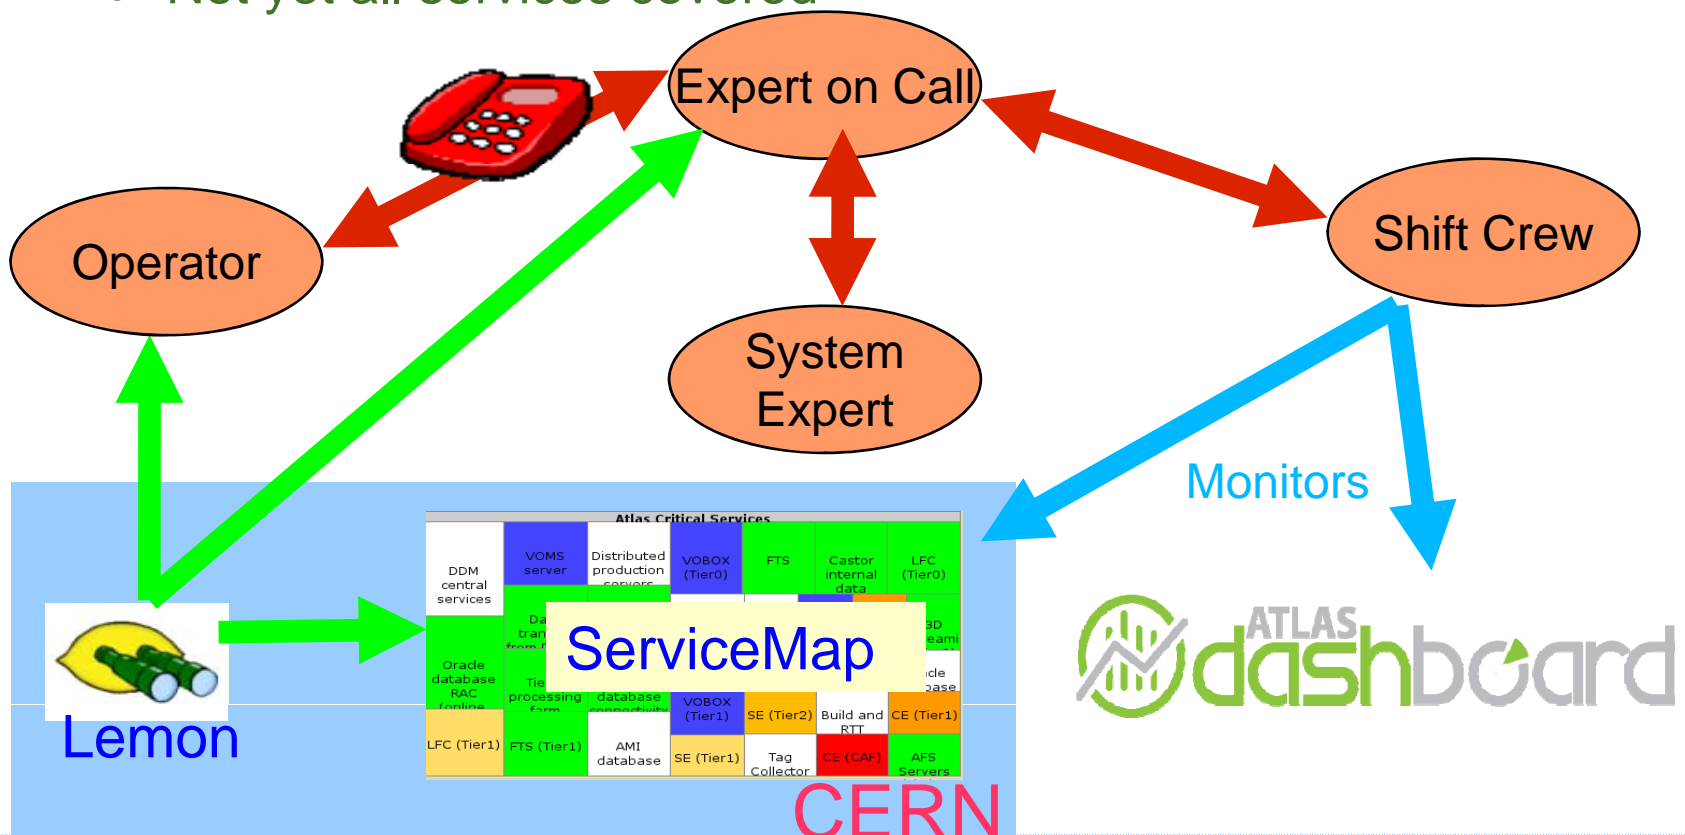




Introduced new Central Services Shifts

GS

- Widened responsibilities of distributed computing shift (16/7, in 2 time-zones) to monitor also data transfers
- Introduced a new Central Services Expert on Call
 - Very good experiences, started training new experts
 - Not yet all services covered





GS

- Since beginning of year Atlas services are moved to production infrastructure
 - Move to modern midrange servers
 - Load balancing for fail-safety for catalogues
 - Quattorizing our services is now mostly done
- Much reduced effort and downtime with quattorized software updates (upgrades Site Services several times)
- During CCRC no noteworthy problems with our central services hardware or software
- However we were severely hit by the the kernel upgrade at CERN, which made our Site Services extremely unstable
 - We were force-upgraded by the power-cut
 - Now have downgraded kernel on all impacted machines
 - In the future we need better protection against forced upgrades and the possibility to test new kernels (and other software updated) for a longer time on our test-infrastructure



A Reality Check

GS

- Power cut on Friday May 30th 6:30 at CERN
 - CASTOR was back ~11:00
 - Until 13:00 all Atlas central data management services were back, the Central Catalogues were prioritized
 - By 14:30 Monitoring and eLog was back
 - Netops reacted immediately after reports of network
- The central operations experts were fortunately all available during the morning and we believe CERN-IT and ADC Operations handled the event very well
 - Lost some files due to bad handshake between pit and Tier-0
- Nevertheless we learned what we can improve:
 - We need procedures, checklists and better coordination
 - We must make sure our communication and monitoring means are available early and possibly have backups outside of CERN



Conclusions

GS

- Atlas Central Services operated very smoothly over CCRC08
- Our investment into quattorized installation, sensors, monitoring and the consolidation of the hardware infrastructure seems to pay off very well
- Procedures and Checklists needed
 - Routing of alarms
 - Organization of shifts
- Higher availability solutions for our communications (eLog) and monitoring means
- Will now run Functional test at low rate continuously
 - Allows to monitor our infrastructure with real activity