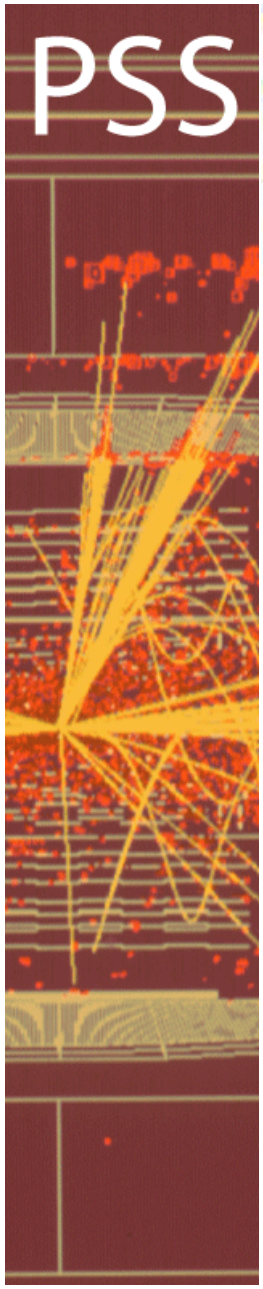

WLCG 3D Workshop Summary

Dirk Düllmann

WLCG Service Challenge Technical Meeting

15 September, CERN

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

PSS

Introduction

CERN IT Department

- Full Production Milestone approaching
- For October all Tier 1 sites will provide Database and Frontier installations
  - According to requirements by experiments and agreed in last Oct GDB/MB
- The 7 sites from the first phase are already available and included in experiment test activities
- 4 Additional sites will need to join now
  - NDGF, NIKHEF/SARA, PIC, TRIUMF

Dirk.Düllmann@cern.ch

Purpose of this workshop

- Summarize the status and remaining problems at all sites
  - Help new sites to join quickly
- Review status of experiment / project replication tests
  - ATLAS, CMS, LHCb and LFC
- Review service procedures and integrate with existing LCG operations infrastructure
- Review/adjust experiment database / frontier resource requests for production in the next 6 month
- Target: Experiments and sites agree on what is (can be) expected during the next 6 month
  - Sites are notified about experiment plans and any requests for database/frontier resource extensions
- Agenda with more detail
  - http://agenda.cern.ch/fullAgenda.php?ida=a063213

Dirk.Düllmann@cern.ch

Since the last workshop

- License Requirements collected and acquired
  - All sites have Oracle s/w and support according to current experiment and project requests
- Database administrator training
  - One week course held with 14 new Oracle administrators from experiments and sites
  - OCP training being setup for November
- Streams through put tests between CERN and T1 sites
  - 10 - 100 MB/min reached (typical 30 MB/min)
  - WAN replication running at ~50% of LAN rates
  - Sufficient for planned use with conditions data
    - Working with Oracle on rate improvement
- Experiments now taking over T1 setups for their replication / client access tests
  - Getting close to full online-offline-T1 chain

Dirk.Düllmann@cern.ch

T0 Production & validation set-up

CERN IT Department

Oracle RAC on Linux project, CERN-IT-PSS
100 CPUs, 200GB RAM, 200TB disk

PRODUCTION

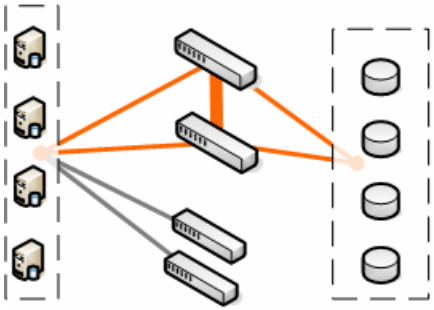

ATLR	ATLAS
CMSR	CMS
LCGR	Grid applications
LHCBR	LHCb
COMPR	Compass

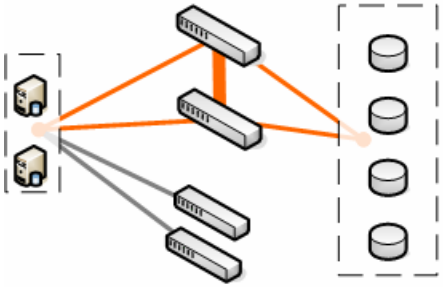

PDBR	Other physics apps
D3R	3D project

VALIDATION

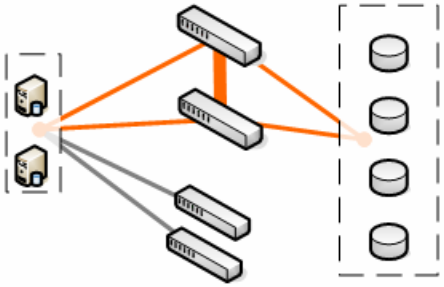

INTR	ATLAS validation
INT2R	CMS validation
INT3R	Grid validation
INT4R	LHCb validation
INT5R	PVSS tests

TEST

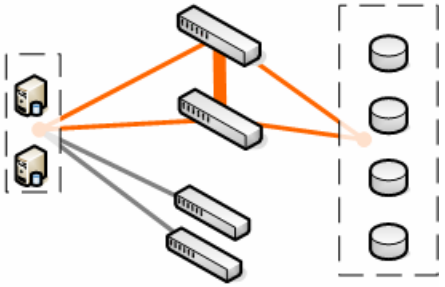

TEST1	Tests

Dirk.Düllmann@cern.ch

T0 FroNtier Production Setup

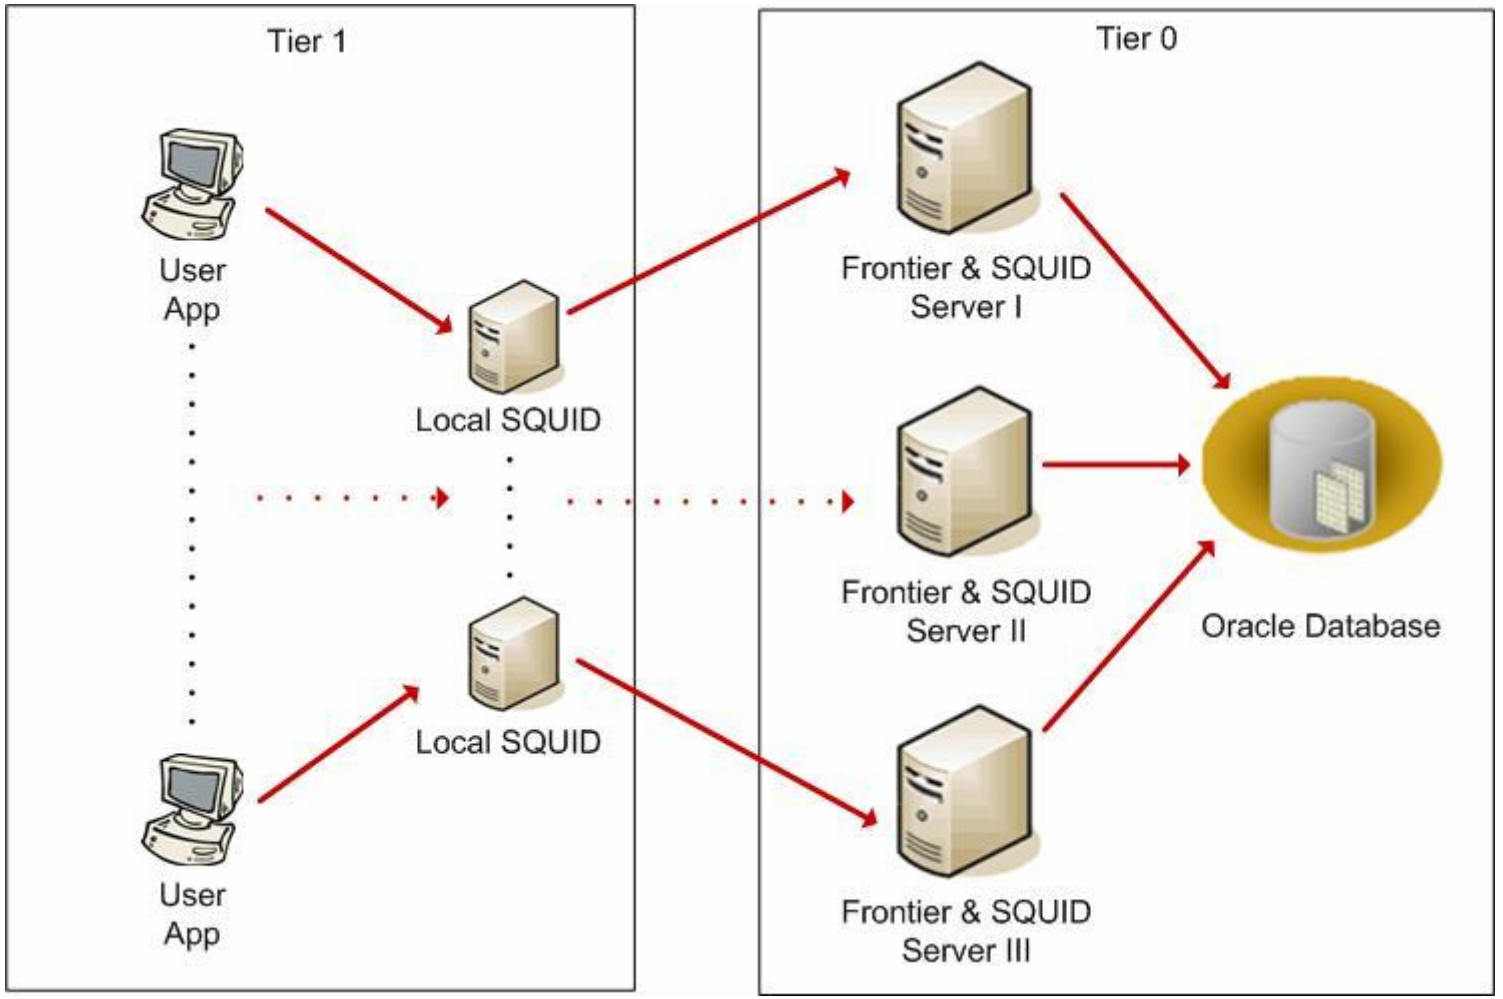

Dirk.Düllmann@cern.ch

T0 Summary

- Service fully based on Oracle RAC/Linux
  - Positive experience
  - Managing many more resources with same team
- The FroNtier servers have been installed and tested
  - Service currently provided by IT-FIO and CMS
- Database service for Tier0 is ready for distributed deployment

Dirk.Düllmann@cern.ch

Site Status

- Phase 1 sites are up and running
  - All sites involved in experiment tests (ASCG is back)
  - Consolidation plans: RAC also grid databases
  - Remaining issues: completion of monitoring and backup setup at some sites
- Attendance still mainly phase 1 sites and TRIUMF
  - NIKHEF/SARA: h/w available, waiting for SAN connection
  - NDG: waiting for h/w arrival, DBA
  - PIC: external company for RAC setup
- New sites need to increase participation in 3D meetings
  - Have to expect delays for those sites wrt October milestone

Dirk.Düllmann@cern.ch

Database & Streams Monitoring

- New web based monitoring pages
  - Show experiment database (online, offline and T1) and replication status in between
  - Database availability, throughput, latency wrt to Tier 0
- Information will be included in ATLAS dashboard
- Integration with SAME availability framework

Dirk.Düllmann@cern.ch

Database setup and maintenance

- Each T1 site is responsible for installation and maintenance of the DB and Frontier servers according to experiment and project request
  - Requests/updates will be collected by 3D and presented to LCG GDB/MB for approval every 6 month
- The sites are responsible for
  - h/w server selection, acquisition, installation and monitoring as well as related network setup
  - s/w installation and upgrade according to the agreed evolution defined
  - regular application of security patches according to site policy
- Tier 0 is responsible for defining the streams setup in agreement with the experiments/grid deployment
- 3D will collect and distribute experience and configuration examples from all sites
  - Including proposals for patch and security upgrades

Dirk.Düllmann@cern.ch

Frontier and Squid Service

- In case of h/w problems of the squid setup sites are required to replace unavailable nodes, but not the cached data.
  - No cache backup is required
- In case the squid cache becomes inconsistent (eg after a power failure) sites may clear the cache (following procedures defined by CMS)
- The Frontier production setup at T0 is operated by FIO (box level) and the FNAL Frontier team (tomcat & squid)
  - Do we need to review this?

Dirk.Düllmann@cern.ch

Database Consistency

- The sites are responsible for recovery from unavailability/inconsistency caused by power or h/w problems
  - Worst case: re-import from T0 and streams resync of one site
  - Need to exercise this at each site!
- The application owners are responsible for recovering from logical corruption caused by their s/w packages
  - Worst case: point-in-time recovery and re-sync on all affected sites
  - Need to schedule a full size recovery test to estimate the unavailability caused by this.

Dirk.Düllmann@cern.ch

Backup and Recovery

- Each site is responsible to setup database backup and recovery infrastructure via Oracle RMAN
  - This may include on-disk backups and should include tape backups and associated media
- Backups should be performed online and with a retention period which is compatible with the time window for point-in-time recovery required by the experiments
  - Eg 1 month or 3 month?
- This is required to allow for a standard recovery procedure including streams re-synchronisation
- The T1 sites are responsible for performing
  - Standard database recovery (eg after media faults)
  - Point-in-time recovery within the agreed time window in case of logical data corruption
- Tier 0 is responsible for the streams re-syncronisation after the site is locally consistent

Dirk.Düllmann@cern.ch

Backup & Recovery

- Database Recovery with Streams
- Collected DB / streams recovery scenarios
  - Recovery after T1 data loss - OK
    - RAL recovered and re-synchronised
    - Replication CENR to CNAF continued unaffected
  - Recovery after T0 data loss - OK
  - Next: coordinated point-in-time recovery
    - Procedure defined, will validate asap
- Service procedure documented on 3D wiki
- Planning local and 3D wide recovery exercise as soon as all sites have backup system in place

Dirk.Düllmann@cern.ch

Problems Resolution & Support Team

- The sites are responsible for staffing their support teams to meet the problem response time and availability numbers defined by the LCG MoU
  - Work split between DBA and other support staff should be organised by each site
- Availability numbers there are only defined indirectly
  - eg for high level application types
  - need to correlate service and database availability monitoring to understand required DB availability
- If application code does not implement DB retry/failover then the DB availability requirements significantly grow.
  - Are there experiment or grid s/w policies ?

Dirk.Düllmann@cern.ch

Service Security

- Security monitoring and patching is entirely site responsibility
  - Each site needs to monitor for security incidents and apply security patches according to site policy
  - This includes emergency interventions like change of compromised admin credentials
- 3D will make information about content and schedule of security patches available and collect experience with their application at the sites
  - But their selection and application is with the sites
- The application of security patches does not require agreement of experiments / software provider
  - It is the responsibility of s/w owners to adapt their applications in case issues with security patches are detected.

Dirk.Düllmann@cern.ch

Problem Reporting

- Should use as much as possible the established reporting channels
- Each site should now join the GGUS support setup
  - Expect that all service users will report problems via this channel
- Each site should pre-announce service changes/outage
  - grid-service-databases@cern.ch
    (new list for any 3D service related discussions)
  - EGEE broadcast and other LCG lists (gmod)
- Integration with operations meetings
  - Either: All sites database team are represented via their grid teams
  - Or: LCG 3D is represented by one person (eg a rotating database manager on duty - dbmod)

Dirk.Düllmann@cern.ch

Experiment Resource Review

- Resource Review for next 6 month (Oct - Mar)
  - Tier 1 database and squid resources are adequate for ATLAS, CMS and LHCb
- Significant increase in medium term expected eg by ATLAS (some 10 TB/y)
- Will review regularly - eg every 6 month
  - Approval of new requests via LCG GDB/MB

Dirk.Düllmann@cern.ch