


Grid Middleware

Markus Schulz - LCG Deployment

February 2010, CERN



Overview

- Middleware(s)
- Computing Access
- Workload Management
- MultiUserPilotJob support
- Data Management
- Information System
- Infrastructure Monitoring
- Release Process
- Summary



Focus on:

- Changes since last year
- Issues
- Plans
- Will not cover all components

Middleware(s)

- WLCG depends on three middleware stacks
 - ARC (NDGF)
 - Most sites in northern Europe
 - ~ 10 % of WLCG CPUs
 - OSG
 - Most North American sites
 - > 25 % of WLCG CPUs
 - gLite
 - Used by the EGEE infrastructure



- All based on the same security infrastructure
- All interoperate (via the experiment's frameworks)
- Variety of SRM compliant Storage Systems
 - BestMan, dCache, STORM, DPM, Castor..

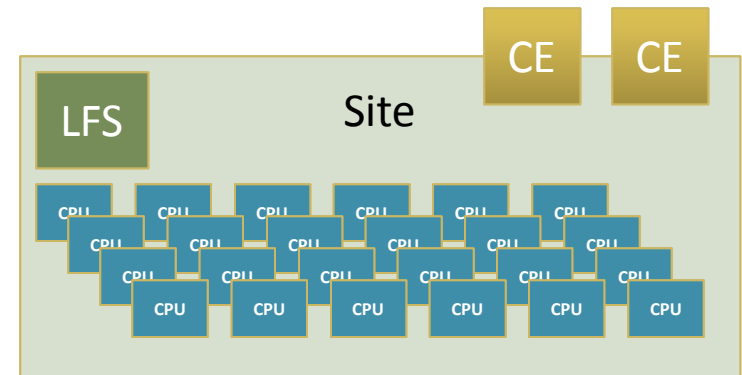
Middleware(s)

- All core components:
 - In production use for several years
 - Evolution based on feedback during challenges
 - And by linking with the LCG Architects Forum
 - Software stabilized significantly during the last year
 - Significant set of shared components:
 - Condor, Globus, MyProxy, GSI OpenSSH, BDII, VOMS, GLUE 1.3 (2) Schema
 - All support at least SL4 and SL5
 - Moved to 64bit on SL5 (RHEL 5), 32bit libraries for compatibility
- Differences
 - gLite strives to support complex workflows directly
 - ARC focuses on simplicity and strong coupling of data and job control
 - OSG (VDT) moves complexity to experiment specific services

Computing Access

- Computing Elements (CE)

- gateways to farms



- EGEE:

- LCG-CE (450 instances)

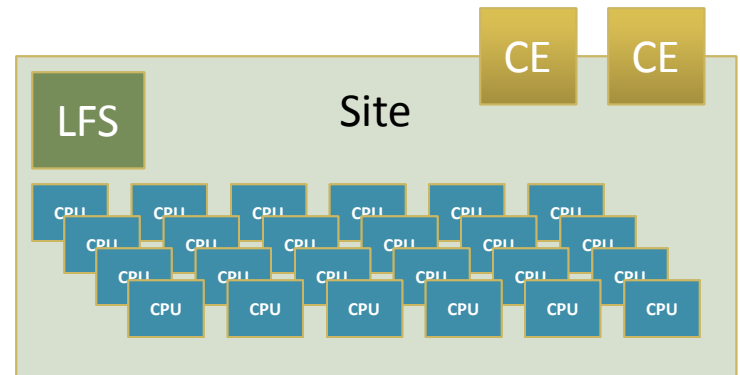
- Minor work on stabilization/scalability (50u/4KJ) , bug fixes
- **LEGACY SERVICE** no port to SL5 planned

- CREAM-CE (69 instances (up from 26))

- Significant investment on production readiness and scalability
- Handles direct submission (pilot job friendly)
 - Production use by ALICE for more than 1 year
 - Tested by all experiments (directly or via WMS)
- SL4/SL5
- BES standard compliant, parameter passing from grid <-> batch
- **Future: gLite Consortium, EMI**
- **Issues: Slow uptake by sites**

Computing Access

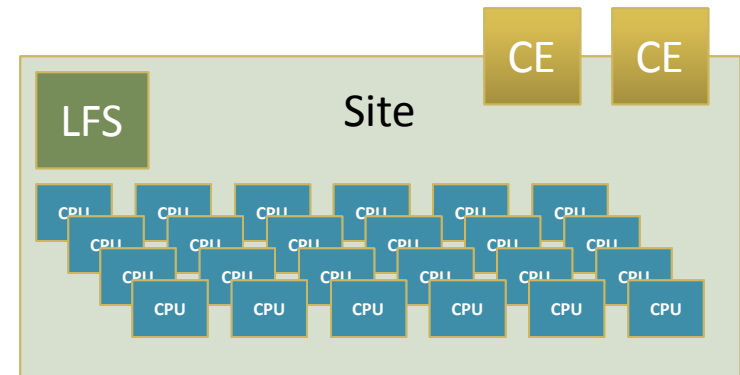
- Computing Elements (CE)
 - gateways to farms




- ARC:
 - ARC-CE (~20 instances)
 - Improved scalability
 - Move to BDII and Glue-1.3
 - KnowArc features included in the release
 - Support for pilot jobs
 - Future: EMI

Computing Access

- Computing Elements (CE)
 - gateways to farms



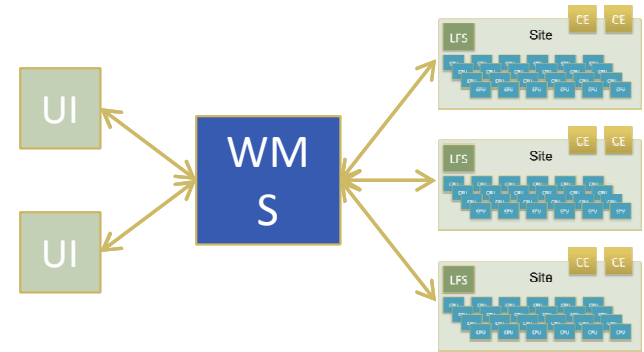
- OSG:
 - OSG-CE (globus) (>50instances)
 - Several sites offer access to resources via Pilot factories
 - Local (automated) submission of Pilot jobs
 - Evaluation of GT-5 gatekeeper (~2Hz, > 2.5k jobs)
 - Integration of CREAM and Condor(-G)
 - Test phase
 - Planning tasks and decisions that lead to deployment
 - Review in mid March
 - Future: OSG/Globus



Workload Management

- EGEE WMS/LB

- Matches resources and requests
 - Including data location
- Handles failures (re-submission)
- Manages complex workflows
- Tracks job status



- EGEE WMS/LB (124 Instances)

- Fully supports LCG-CE and CREAM-CE
 - Early versions had some WMS<->CREAM incompatibilities
- Several updates during the year
 - Much improved stability and performance
- LCG VOs use only a small subset of the functionality
- Future: gLite Consortium /EMI

MultiUserPilotJobs

- Pilot Jobs (Panda, Dirac, Alien...)
 - Framework sends jobs to sites
 - No “physics” workload
 - When active the Pilot contacts the VO’s task-queue
 - The Experiment schedules a suitable job and moves it to the Pilot and executes it
 - This is repeated until the maximum queue time is reached
- MUPJs run workloads from **different** users
 - The batch systems is only aware of the Pilot’s identity
 - Flexibility for the experiment
 - Conflicts with site security policies
 - Lack of traceability
 - “Leaks” between users

MultiUserPilotJobs

- Remedy for this problem:
 - Changing the UID/GID according to the workload
- Implementation:
 - EGEE
 - glexec (setuid code or logging) on the Worker Node
 - SCAS or ARGUS service to handle authorization
 - OSG
 - Glexec / gums
 - In production for several years
- Glexec/SCAS ready for deployment
 - Scalability and stability tests passed
 - Deployed only on a few sites

MultiUserPilotJobs

- Glexec/ARGUS

- ARGUS is the new authorization framework for EGEE
 - Much richer policy management than SCAS
- Certified
- Deployed on a few test sites




- Both solutions have little exposure to production
 - Need some time to fully mature

- Future: glexec/SCAS/ARGUS gLite-Consortium/EMI

Data Management

- Storage Elements (SEs)
 - External interfaces based on SRM 2.2 and gridFTP
 - Local interfaces: POSIX, dcap, secure rfiio, rfiio, xrootd
 - DPM (241)
 - dCache (82)
 - STORM (40)
 - BestMan (26)
 - CASTOR (19)
 - “ClassicSE” (27) → legacy since 2 years....
- Catalogue: LFC (local and global)
- File Transfer Service (FTS)
- Data management clients gfal/LCG-Utils



Data Management

- Common problems:
 - Scalability
 - I/O operations
 - Random I/O (analysis)
 - Bulk operations
 - Synchronization
 - SEs <-> File Catalogues
 - Quotas
 - VO-Admin Interfaces
- All services improved significantly during the year.

Data Management

- Examples:

- DPM

- Several bulk operations added
 - Improved support for checksums
 - RFIO improvements for analysis
 - Improved xrootd support
 - Next release DPM 1.8 (end of April)
 - User banning, VO Admin capacity

- FTS

- Many bug fixes
 - Improved monitoring
 - Checksum support
 - Next Release: 2.3 (end of April)
 - Better handling of downtime and overload of storage elements
 - Move from “channels” to SE representation in DBs
 - Administrative web interface
 - Longer term: Support for small, non-SRM SEs (T3)

Data Management

- Examples:
 - CASTOR
 - Consolidation
 - Castor 2.1.9 deployed
 - Improved monitoring with detailed indicators for stager and SRM performance
 - Next release: SRMv2.9 (February)
 - Addresses SRM instabilities reported during the last run
 - Improved monitoring as requested by the experiments
 - Observation: xroot access to Castor is sufficient for analysis
 - Further improvements:
 - Tuning root client and xroot servers
 - Plan: deploy native xroot instances for analysis
 - Low latency storage
 - Discussion started on dataflow
 - Before summer: disk only
 - After summer: disk + backup

Data Management

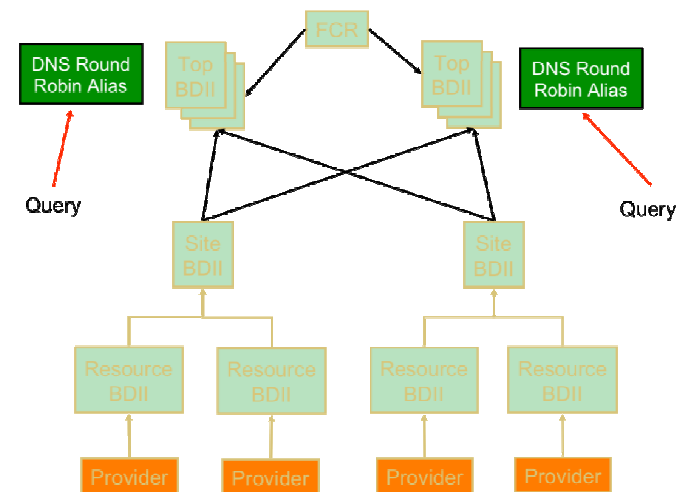
- Examples:
 - dCache
 - Introduced Chimera name space engine
 - Improved scalability
 - Released “Golden Release dCache 1.9.5”
 - Functionality will be stable during first 12 months
 - Bug fix releases as required
 - Plans (12 months):
 - Multiple SRM front ends (improved file open speed)
 - NFS-4.1 (security has to be added)
 - First performance tests are promising
 - WebDav (https)
 - Integration with Argus
 - Information system and monitoring

Data Management

- Examples:
 - STORM
 - Added tape backend
 - SRM-2.2 + WLCG extensions implemented
- Future:
 - dCache, STORM, DPM, FTS, LFC, clients → EMI
 - Castor → CERN
 - BestMan → OSG

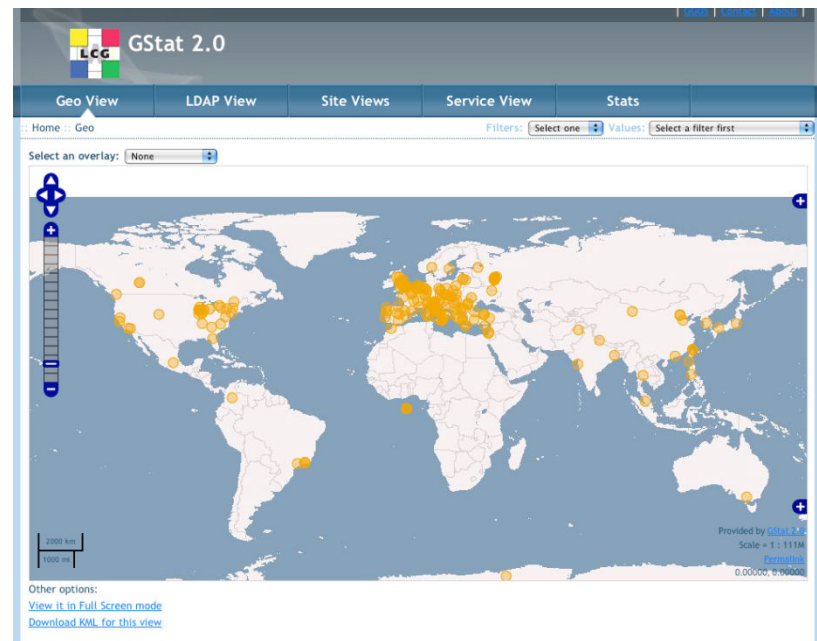
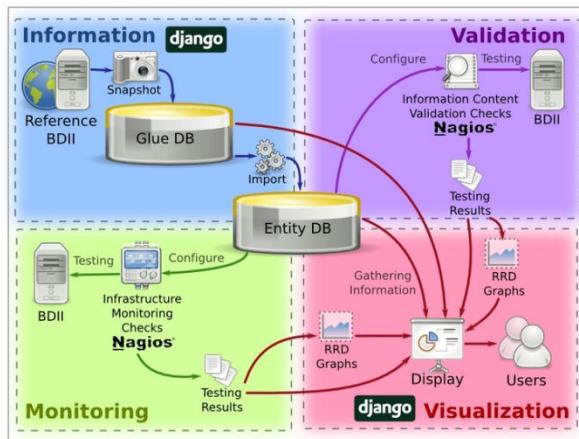
Information System

- BDII
 - Several updates during the year
 - Improved stability and scalability
 - Support for new GLUE-2 schema
 - OGF standard
 - Parallel to 1.3 to allow smooth migration
 - Better separation of “static” and “dynamic” information
 - Opens the door for new strategy towards scalability
 - Issues:
 - Complex schema
 - Wrong data published by sites
 - Bootstrapping
 - Future: gLite Consortium/EMI



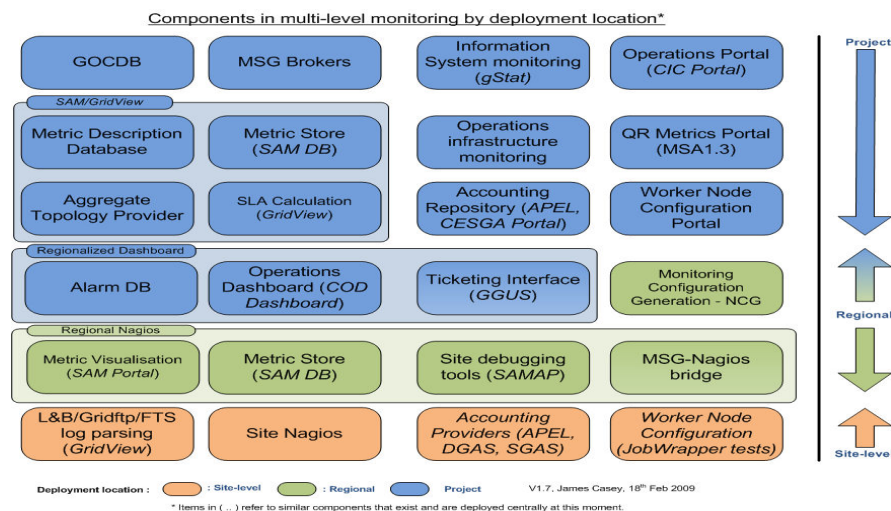
Information System

- Gstat-2.0 <http://gstat-prod.cern.ch/gstat/stats/GRID/ALL>
 - Information system monitor and browser
 - Consistency checks
 - Solid implementation based on standard components
 - CERN/Academia Sinica Taipei



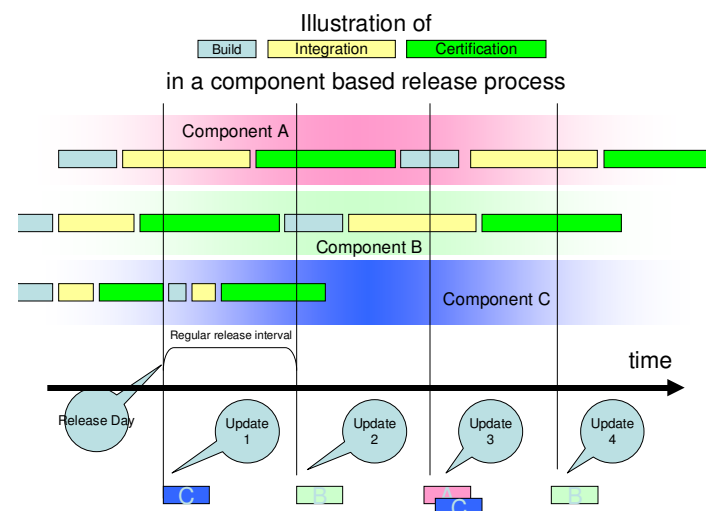
Infrastructure Monitoring

- Distributed system based on standard technology
 - NAGIOS, DJANGO
 - ActiveMQ based messaging infrastructure
 - Integrated existing SAM tests
 - Use MyOSG based visualisation -> MyEGEE
 - Reflects operational structure of EGI
 - Replaces SAM system
 - “Grown” central system



Release Process

- Refined component based release process
 - Frequent releases (2 week intervals)
 - Monitored process
 - Fast rollback
 - Components have reached a high level of quality
 - Synthetic testing is limited
 - Fast rollback limits impact
 - Staged Rollout
 - Validation in production
- Transition to Product Teams
 - Responsible for:
 - Development, Testing, Integration, Certification
 - Based on project policies



General Evolution

- Move to standard building blocks
 - ActiveMQ, Django, Nagios
 - Globus GSI → openssl
- Data Management
 - cluster file systems as building blocks
 - STORM, BestMan, (DPM)
 - Using standard clients NFS-4.1
 - Reducing complexity (FTS)
- Workflow management and direct control by Users
 - Direct submission of Pilots to CREAM-CEs (no WMS)
- Virtualization
 - Fabric/application independence
 - User-controlled environments

Open Issues

- EC funded projects EGI/EMI
 - Not sufficient to continue all activities at current level
 - Change rate can be reduced
 - Some activities can be stopped
 - Middleware support will depend more on community support
 - Build and integration systems will be adapted to support this
- Continuity
 - Significant staff rotation and reduction
- Uptake of new services is very slow
- Development of a long-term vision
 - After 10 years a paradigm change might be due...

Summary

- WLCG Middleware handles core tasks adequately
- Most developments targeted at:
 - Improved control
 - Quotas, security, monitoring, VO-admin interfaces
 - Improved recovery from problems
 - Catalogue/SE resynchronization
 - Simplification
 - Move to standard components
 - Performance improvements
 - Stability
- How stable is the software?

Error rates / Usage

- Bug rate almost flat
- Exponential increase in usage
- Example: gLite

PRODUCTION Normalised CPU time by VO and DATE
LHC VOs (and Other VOS). April 2004 - February 2010



Open vs Closed bugs history

