HEP Applications Evaluation of the EDG Testbed and Middleware

Stephen Burke (EDG HEP Applications – WP8)

s.burke@rl.ac.uk

http://presentation.address

Talk Outline

- The organisation and mission of EDG/WP8
- Overview of the evolution of the EDG Applications Testbed 2002-3
- Overview of the Task Force activities with HEP experiments and their accomplishments
- Use case analysis mapping HEPCAL to EDG
- Main lessons learned and recommendations for future projects
- A forward look to EDG release 2 and co-working EDG/LCG

Authors

I.Augustin, F.Carminati, J.Closier, E.van.Herwijnen	CERN
J.J.Blaising, D.Boutigny, A.Tsaregorodtsev	CNRS (France)
K.Bos, J.Templon	NIKHEF (Holland)
S.Burke, F.Harris	PPARC (UK)
R.Barbera, P.Capiluppi, P.Cerello, L.Perini, M.Reale, S.Resconi, A.Sciaba, M.Sitta	INFN (Italy)
O.Smirnova	Lund (Sweden)

EDG Structure

- Six middleware work areas:
  - Job submission/control
  - Data management
  - Information and monitoring
  - Fabric management
  - Mass Storage
  - Networking
- Three application groups:
  - HEP (WP8)
  - Earth observation
  - Biomedical
- Testbed operation
- Three-year project, 2001-03
- Current release 1.4, expect 2.0 in May

Mission and Organisation of WP8

- To capture the requirements of the experiments, to assist in interfacing experiment software to EDG middleware, to evaluate functionality and performance, and give feedback to middleware developers and report to the EU
  - Also involved in generic testing, architecture, education etc.
- 5 'Loose Cannons' … full time people helping all experiments (were key members of the Task Forces for ATLAS and CMS)
- 2-3 representatives from each experiment
  - ALICE, ATLAS, CMS, LHCb
  - BaBar, D0 (since Sep 2002)
- Most recent report 'Evaluation of Datagrid Application Testbed during Project Year 2' available now (EU deliverable D8.3)
  - This talk summarises the key points of the report

The EDG Applications Testbed

- The testbed has been running continuously since November 2001
- Five core sites: CERN, CNAF, Lyon, NIKHEF, RAL
  - Now growing rapidly: currently around 15 sites, 900 CPUs, 10 Tb of disk in Storage Elements (plus local storage). Also Mass Storage Systems at CERN, Lyon, RAL and SARA.
- Key dates:
  - Feb 2002, release 1.1 with basic functionality
  - April 2002, release 1.2: first production release, used for ATLAS tests in August
  - Nov 2002 - Feb 2003, release 1.3/1.4: bug fixes incorporating a new Globus version, stability much improved. Used for CMS stress test, and recently ALICE and LHCb production tests.
  - May 2003, release 2.0 expected with major new functionality, application testbed expected to merge with LCG 1 (the LHC Computing Grid testbed)

Middleware in Testbed 1

- Basic Globus services: GridFTP, MDS, Replica Catalog, job manager
- Job submission: submit a job script to a Resource Broker which dispatches the job to a suitable site, using matchmaking information published in MDS. Uses Condor-G.
- Data management: tools to copy files with GridFTP and register them in the Replica Catalogs. Interfaced with the job submission tools to allow jobs to be steered to sites with access to their input files.
- Fabric management: allows automated configuration and update of Grid clusters
- VO management: VO membership lists in LDAP servers, users are mapped to anonymous VO-based "pool accounts"

Resumé of experiment DC use of EDG-see experiment talks elsewhere at CHEP

- ATLAS were first, in August 2002. The aim was to repeat part of the Data Challenge. Found two serious problems which were fixed in 1.3
- CMS stress test production Nov-Dec 2002 – found more problems in area of job submission and RC handling – led to 1.4.x

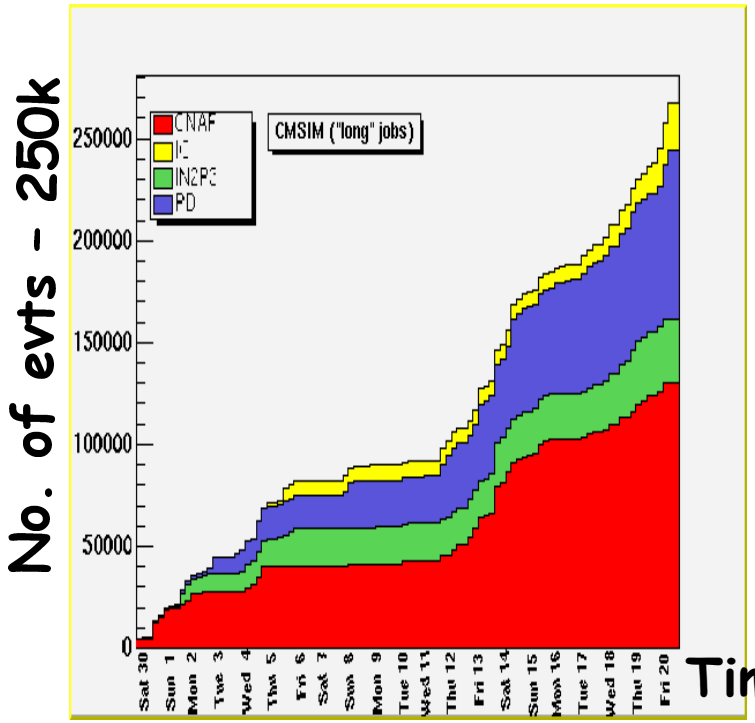

- ALICE started on Mar 4: production of 5,000 central Pb-Pb events - 9 TB; 40,000 output files; 120k CPU hours
  - Progressing with similar efficiency levels to CMS
  - About 5% done by Mar 14
  - "Pull" architecture
- LHCb started mid Feb
  - ~70K events for physics
  - Like ALICE, using a pull architecture
- BaBar/D0
  - Have so far done small scale tests
  - Larger scale planned with EDG 2

Use Case Analysis

- EDG release 1.4 has been evaluated against the HEPCAL Use Cases
- Of the 43 Use Cases:
  - 6 are fully implemented
  - 12 are largely satisfied, but with some restrictions or complications
  - 9 are partially implemented, but have significant missing features
  - 16 are not implemented
- Missing functionality is mainly in:
  - Virtual data (not considered by EDG)
  - Authorisation, job control and optimisation (expected for release 2)
  - Metadata catalogues (some support from middleware, needs discussion between experiments and developers about usage)

Lessons Learnt - General

Many problems and limitations found, but also a lot of progress. We have a very good relationship with the middleware and testbed groups.

- Having real users on an operating testbed on a fairly large scale is vital – many problems emerged which had not been seen in local testing.
- Problems with configuration are at least as important as bugs - integrating the middleware into a working system takes as long as writing it!
- Grids need different ways of thinking by users and system managers. A job must run anywhere it lands. Sites are not uniform so jobs should make as few demands as possible.

Job Submission

- Limitations with the versions of Globus and Condor-G used:
  - 512 concurrent jobs /Resource Broker
  - Max submission rates of 1000 jobs/hr
  - Max 20 concurrent users
- Worked with multiple brokers and within rate limits
- Can be very sensitive to poor or incorrect information from Information Providers, or propagation delays
  - Resource discovery may not work
  - Resource ranking algorithms are error prone
  - Black holes – "all jobs go to RAL"
- The job submission chain is complex and fragile
- Single jobs only, no splitting, dependencies or checkpointing

Information Systems and Monitoring

- It has not been possible to arrive at a stable, hierarchical, dynamic system based on MDS
  - System jammed up with increasing query rate and hence could not give reliable information to clients such as RB
  - Used workarounds (static list of sites, then a fixed database LDAP backend). Works, but not really satisfactory. New R-GMA software due in release 2.
- Monitoring/debugging information is limited so far
- We need to develop a set of monitoring tools for all system aspects

Data Management

- Replica Catalog
  - Jammed with many concurrent accesses
  - With long file names (~100-200 bytes) there was a practical limit of ~2000 entries
  - Hard to use more than one catalogue in the current system
  - No consistency checking against disk content
  - Single point of failure
  - New distributed catalogue (RLS) in release 2
- Replica Management
  - Copying was error prone with long (1-2 Gb) files (~90% efficiency)
  - Fault tolerance is important: error conditions should leave things in a consistent state

Use of Mass Storage (CASTOR, HPSS)

- EDG uses GridFTP for file replication, CASTOR and HPSS use RFIO
- Interim solution:
  - Disk file names have a static mapping to the MSS
  - Replica Management commands can stage files to and from MSS
  - No disk space management
- Good enough for now, but a better long-term solution is needed
  - A GridFTP interface to CASTOR is now available
  - The EDG MSS solution (Storage Element) will be available in release 2

Virtual Organisation (VO) Management

- Current system works fairly well, but has many limitations
- VO servers are a single point of failure
- No authorisation/accounting/auditing/…
  - EDG has no security work package! (but there is a security group)
  - New software (VOMS/LCAS) in release 2
- Experiments will also need to gain experience about how a VO should be run

User View of the Testbed

- Site configuration is very complex, there is usually one way to get it right and many ways to be wrong! LCFG (the fabric management system) is a *big* help in ensuring uniform configuration, but can't be used at all sites.
- Services should fail gracefully when they hit resource limits. The Grid must be robust against failures and misconfiguration.
- Many HEP experiments (and the EDG middleware at the moment) require outbound IP connectivity from worker nodes, but many farms these days have "Internet Free Zones". Various solutions are possible, discussion is needed.

Other Issues

- Documentation – EDG has thousands of pages of documentation, but it can be very hard to find what you want
- Development of user-friendly portals to services
  - Several projects underway, but no standard approach yet
- How do we distribute the experiment application software?
- Interoperability – is it possible to use multiple Grids operated by different organisations?
- Scaling – can we make a working Grid with hundreds of sites?

A Forward Look to EDG 2/LCG 1

- Release 2 will have major new functionality:
  - New Globus/Condor releases via VDT, including GLUE schema
    - Use of VDT + GLUE is a major step forwards in US/Europe inter-operability
  - Resource Broker re-engineered for improved stability and throughput
  - New R-GMA information and monitoring system to replace MDS
  - New Storage Element manager
  - New Replica Manager/Optimiser with a distributed Replica Catalogue

- From July 2003 the EDG Application Testbed will be synonymous with the LCG-1 prototype

- EDG/WP8 are now working together with LCG in several areas – requirements, testing, interfacing experiments to LCG middleware (EDG 2 + VDT)

Summary & Future Work

- The past year has seen major progress in the use by the experiments of EDG middleware for physics production on an expanding testbed - pioneering tests by ATLAS + real production by CMS, and now ALICE and LHCb
  - WP8 has formed the vital bridge between users and middleware, both by generic testing and by involvement in Task Forces
  - And has been a key factor in the move from R/D to a 'production' culture
- There is strong strong interest from running experiments (BaBar and D0) to use EDG middleware
- We have had excellent working relations with the Testbed and Middleware groups in EDG, and this is continuing into LCG
- We foresee intense testing of the LCG middleware combining efforts from LCG, EDG and the experiments, and also in user support and education activities
- **ACKNOWLEDGEMENTS**
  - **Thanks to the the EU and our national funding agencies for their support of this work**