

CERN openlab II (2006 – 2008)

Grid-related activities

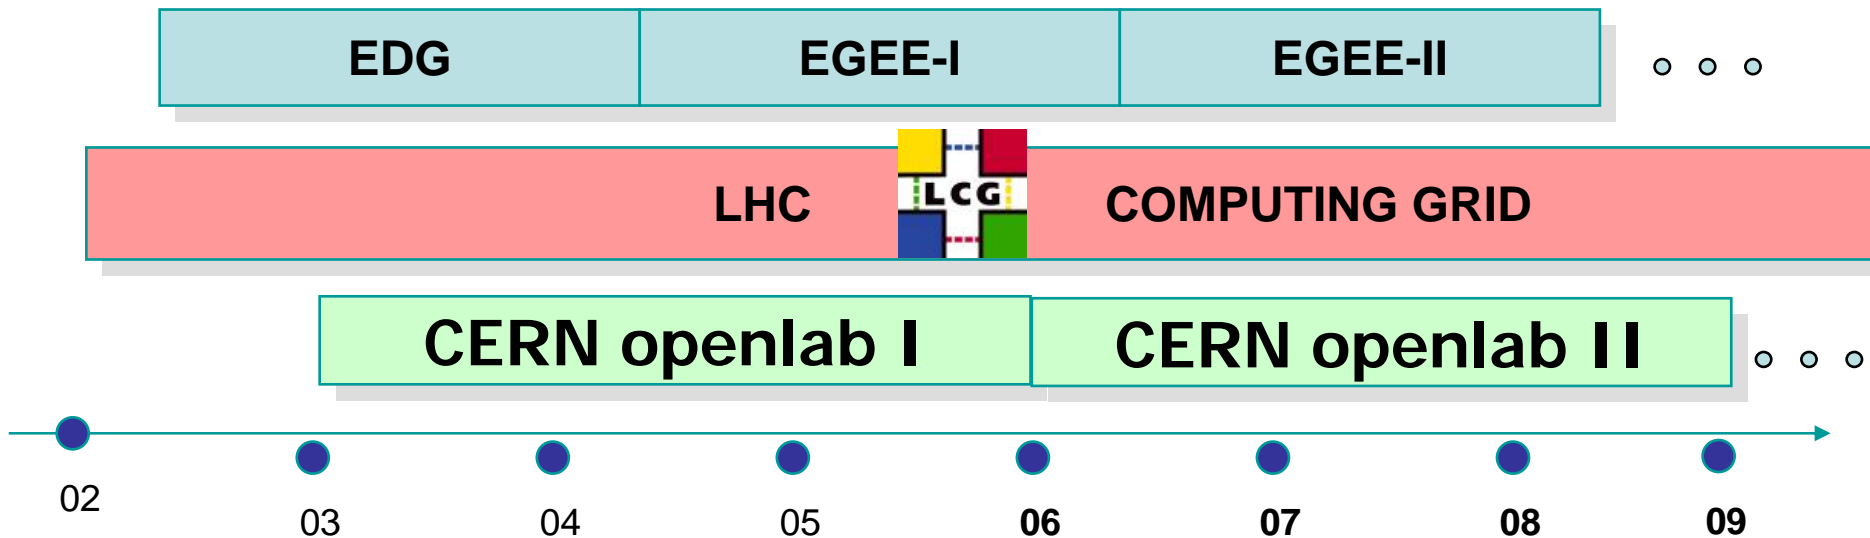
Sverre Jarp
CERN openlab CTO
sverre.jarp@cern.ch





Overview and timeline

- CERN IT Department's main R&D focus
- Framework for collaboration with industry
- Evaluation, integration, validation
 - of cutting-edge technologies that can serve the LHC Computing Grid (LCG)
- Sequence of 3-year agreements
 - 2003 – 2005: the “opencluster” project
 - 2006 – 2008: openlab Phase II with new projects



- Major current focus areas
 - Platform Competence Centre
 - Grid Interoperability Centre
 - Advanced DB Deployment
 - Security
- Being finalised:
 - Network Monitoring



Future partners/contributors. Please see:
http://proj-openlab-datagrid-public.web.cern.ch/proj-openlab-datagrid-public/Guiding_Principles_Partners_and_Contributors-Version_3.pdf

PARTNERS



ORACLE®

CONTRIBUTORS



STONESOFT
Real World Business Security

The “openlab contract”

- Agreed framework for sponsorship
 - Formal agreement at high level (CERN Director General)

- Set of guiding principles
 - 3 year commitment (partners)
 - Agreed level of contribution (500K € per year)
 - 50% in-kind – 50% manpower
 - (Up-front) collaboration agreement
 - Technical annexes
 - Regular Technical Reviews, Workshops
 - Liaison officer to each company
 - Annual Board of sponsors meeting

- A complementary option: Contributor status
 - One year engagement, 1/10 of partner contribution




- Combination of
 - “our key to success”, i.e.
 - Industry-sponsored post-doc positions (CERN fellows)
 - Also: positions for preparing a PhD
 - Marie-Curie EU positions
 - Combined training and contributing positions
 - Strong summer student programme
 - CERN staff
 - Mainly management and secretariat
 - Head: W. von Rüden
 - Manager: F. Fluckiger
 - CTO: S. Jarp
 - Communications: F. Grey
 - Secretary: S. Pizzera

- Important results in multiple fields
 - Large grid compute node (integrated with LCG test bed) – 100 Itanium DP nodes
 - 64-bit application programming
 - 64-bit LCG-2 stack
 - 10-gbit LAN and WAN
 - Infiniband interconnect
 - CFD service for air flow in LHC caverns
 - Initial virtualization experience
 - Working with Xen since Summer 2004
 - Initial SmartFrog experience
 - Becoming familiar with deployment frameworks

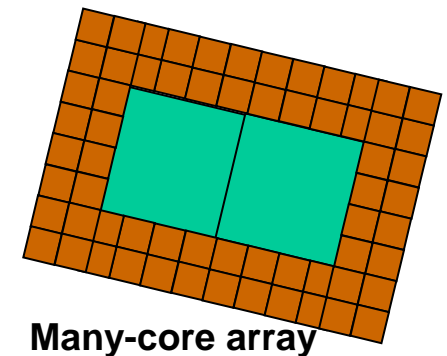
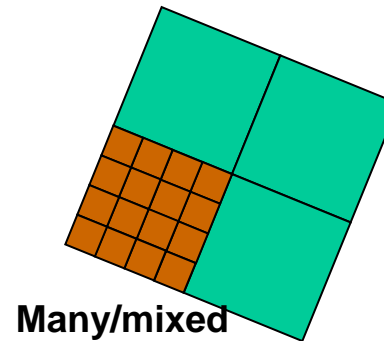



Review of some relevant projects in openlab II

- Fundamental component of future Grid services
 - In the near future, everybody will regard virtualization as “a given”
 - Increased flexibility
 - OS deployment for test and production
 - Increased security
 - Isolation, even with root privileges
 - Increased configurability and manageability
 - Selection of OS at the individual level
 - Manage VM images (across nodes, and grids)
 - Increased flexibility
 - Checkpointing at OS level; migration between servers




- Another foundation for tomorrow's grid computing

- Unequaled opportunity for high-throughput computing
- We are only at the beginning of a long evolution
- In openlab we are active with Intel in benchmarking/throughput testing
 - Also requirements for the future




Evolution

- Oracle technology used in several domains, such as:
 - *Streams* for data replication between LCG sites
 - Distributed Database Deployment (3D)
 - Between CERN (Tier 0) and Tier 1 sites
 - Special emphasis on backup and recovery, as well as monitoring
 - *Data Guard* automatic fail-over
 - Reduce downtime
 - Unexpected failures
 - Planned upgrades
 - Scalability of *RAC* (Real Application Clusters)
 - Vital for subsystems, such as POOL (Pool of persistent objects for LHC)




- Open source framework
 - “Smart FFramework for Object Groups”
 - Developed by HP Labs (Bristol)
 - <http://www.smartfrog.org/>
 - Strong candidate for automated grid management
 - Structured Java programming
 - Emphasis on interrelated set of components and entire lifecycle management
 - Already being used in EGEE/SA3
 - Framework for testing of new Grid components
 - Combined with virtualization




- SmartFrog

- Orchestrating components easily




- Wherever they sit




Courtesy: X.Gréhant/openlab

- Market-based system for managing computing resources in a distributed environment
- Developed by HP Labs, Palo Alto
- <http://tycoon.hpl.hp.com/>
- We want to understand what Tycoon can do in the context of e-science grids:
 - How to deal with swings in demands for computing
 - 1) Provide complementary services?
 - Based on efficient gateways
 - 2) Provide technology that can be integrated into science grids (such as EGEE)?
 - Analyze Tycoon's features in the context of large-scale configurations

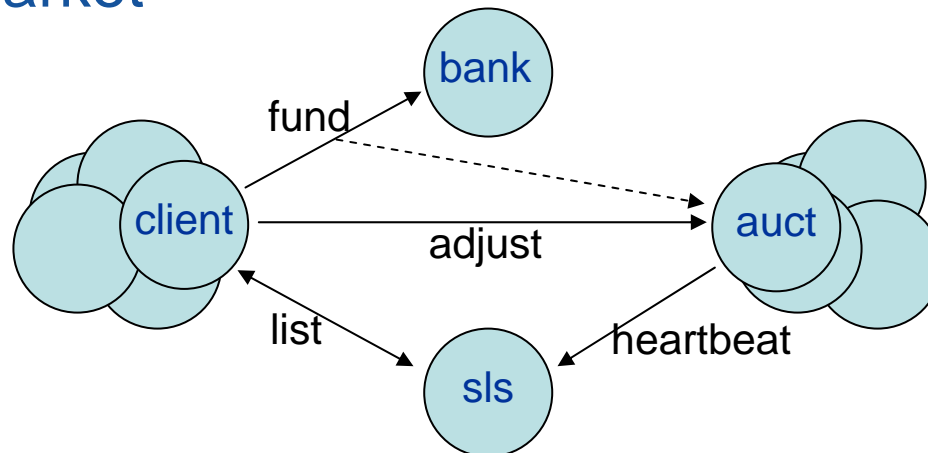


- Architecture

- Auctioneer: Xen VM management with Python



- Market



Courtesy: X.Gréhant/openlab

- CERN openlab II is off to strong start:
 - Solid collaboration with our industrial partners
 - Encouraging results in multiple domains
 - Including Grid Computing
 - Still gathering momentum
 - 2 ¼ years to go
 - We believe partners are getting good “ROI”
 - But only they can really confirm it → so ask them
 - No risk of running short of R&D
 - IT Technology is still moving at an incredible pace
 - What was “R&D” in openlab I is now typically in production
 - The same will happen in a couple of years with the ingredients in openlab II