

# CERN openlab II:

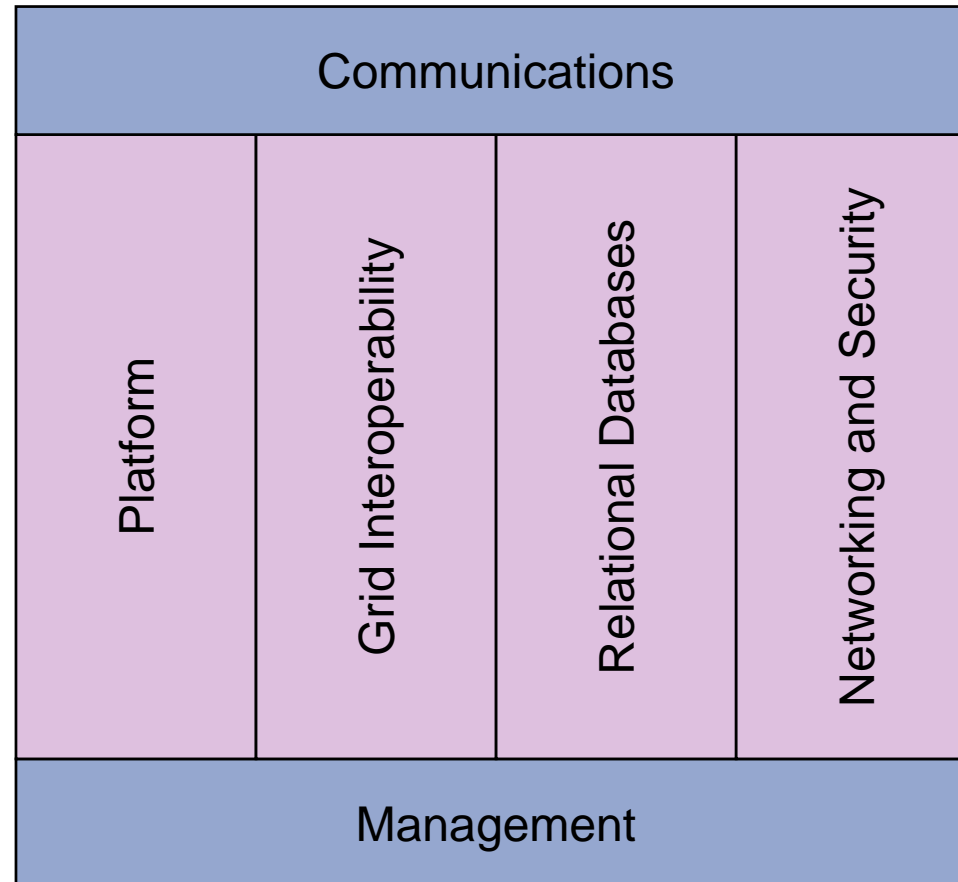
Review of the first 16 months,  
Plans for the future



Sverre Jarp, 26 April 2007

CERN openlab CTO

[sverre.jarp@cern.ch](mailto:sverre.jarp@cern.ch)





# Openlab II success formula

- Dedicated manpower in almost all areas
- Currently 9 fellows/staff recruited
  - 3 Oracle
    - Eva Dafonte Pérez, Anton Topurov, Dawid Wójcik
  - 2 HP (4 as of 1.7.)
    - Jose Dana, Xavier Gréhant
    - Milosz Hulbój, Ryszard Jurga
  - 2 Intel
    - Håvard Bjerke, Andreas Hirstius
  - 2 Marie-Curie (1 as of 1.7.)
    - Andrzej Nowak
- 1 EDS fellow imminent
- Close collaboration with all partners and contributors
- Broad portfolio of relevant projects



# Review of last year



# Platform Competence Centre (PCC)

- Projects:
  - Benchmarking, Thermal testing of servers
    - Intensive testing of Dual-core/Quad-core servers
  - Performance monitoring
    - Working with author of new, universal interface in Linux kernel
  - Compiler project
    - Covers Xeon-64, Itanium (icc and gcc); Article published in Courier
  - Platform virtualization
    - Xen, and review of KVM (new Linux package)
  - CFD studies on opencluster
    - Service continues (also with users from AB)
  - DAPL
    - Investigation of use of openfabrics I/O stack in xrootd
  
- Recent additions:
  - Multithreading
    - Linked to multicore
  - New Computer Centre planning
    - Visits to Intel Centres in the US, dialogue with Intel experts



# Grid interoperability Centre (GIC)

- Projects:
  - TYCOON
    - Started recently; Collaboration w/EGEE, BalticGrid
    - Regular contacts with HP Labs, PA
  - SmartDomains
    - Deployment in EGEE/SA3; Recent acceptance of paper at VHPC (Euro-Par 07)
  - Grid Programming Environment (GPE)
    - Collaboration with Intel, Brühl
  - OSFARM
    - Generation/Retrieval of OS images

- Projects:
  - Streams for database replication in LCG
    - Now in operation between all Tier-1 centres
  - Scalability on RAC
    - 1500x performance improvement for change-archival application
    - Near linear RAC scalability
  - Oracle Data Guard Automatic Failover
    - Now using Data Guard at CERN as a migration tool to minimise server downtime, feedback given to Oracle
  - Benchmarking
    - Recent comparison of Dual-core/Quad-core servers
  - Software regression testing
    - New versions and features tested, issues reported and fixed by Oracle
  - 64-bit DB servers
    - Smooth introduction, production since Q4 2006

- Projects:
  - F-Secure
    - Intrusion detection with emphasis on client security and mail server security evaluated by IS group. Encouraging results in the area of virus detection.
  
  - Stonesoft
    - Network based intrusion detection evaluated for the CERN environment
    - Use cases identified to enhance functionality





# Plans for next 12 months



## Reviews and project documentation

- Maintain Monthly/Quarterly review meetings
  - Monthly:
    - May 22<sup>nd</sup>, June 19<sup>th</sup> (Tuesdays at 2PM)
  - Quarterly:
    - Sept 25<sup>th</sup>
- Workshop on Multithreading
  - May 31<sup>st</sup>/June 1<sup>st</sup> (w/H.Cornelius from Intel)
- Other workshops to follow
  
- Project documentation
  - Continue enhancements on the Web



# Platform Competence Centre (PCC)

- **Maintain:**
  - Benchmarking, Thermal testing of servers
    - New generation Intel processors expected
  - Performance monitoring
    - Large scale deployment w/FIO and physics programmers
  - Compiler project
    - Continue as before
  - Platform virtualization
    - Investigate support of perfmon inside virtual machines
  - CFD studies on opencluster
    - Upgrade to dual-core Montecito processors
  
- **Intensify:**
  - Multithreading
    - Workshop, Collaborate with PH/SFT on applications
  - Computer Centre planning
    - Power optimization inside rack, Continue plan for new CC



# Grid Interoperability Centre (GIC)

- Intensify:
  - TYCOON
    - Aim at real feasibility tests w/EGEE partners
  - SmartDomains
    - Broader deployment
  - Grid Programming Environment (GPE)
    - Integration of hardware control
  - OSFARM
    - Broaden usage
- Initiate:
  - Grid monitoring project w/EDS
    - Detailed plan already being elaborated by Max Böhm (after initial visit)



# Database Projects

- **Maintain:**
  - Streams for database replication in LCG
    - Ramp up work load to full LHC rates and integrate new applications
  - Scalability on RAC
    - PVSS production; New CERN applications and reference document
  - Oracle Data Guard Automatic Failover
    - Deployment in production, emphasis on critical grid services
  - Enterprise Manager
    - Full deployment and better integration for CERN and 3D installations, usage of more features, scale-up and high availability architecture
  - Benchmarking
    - New systems, scalability with more cores per CPU
  - Software regression testing
    - New versions coming, new functionality tests
  - 64-bit DB servers
    - Deployment for CASTOR and other services (CERN and Tier 1 sites), exploit larger memory

- Projects:
  - Procurve (Three years as of 1<sup>st</sup> July)
    - One staff + one fellow
  - Understand behaviour of large computer network in large HPC/campus installations in order to:
    - Detect traffic anomalies in the system
    - Be able to perform trend analysis
    - Take automatic counter measures
    - Provide post-mortem analysis facilities



## Conclusion and outlook

- New, exciting project: HP Procurve
- New, exciting contributor: EDS
- Our broad program is set to continue and expand
  - First aim: Foster cross-project fertilization and collaboration (inside openlab)
  - Second aim: Foster broad collaboration with all relevant IT groups, LCG, and physics communities
- Thanks to our partners and contributors we continue to accelerate!