



Enabling Grids for E-science

Performance Analysis of Existing and Emerging Computing Interfaces for Grid Computing

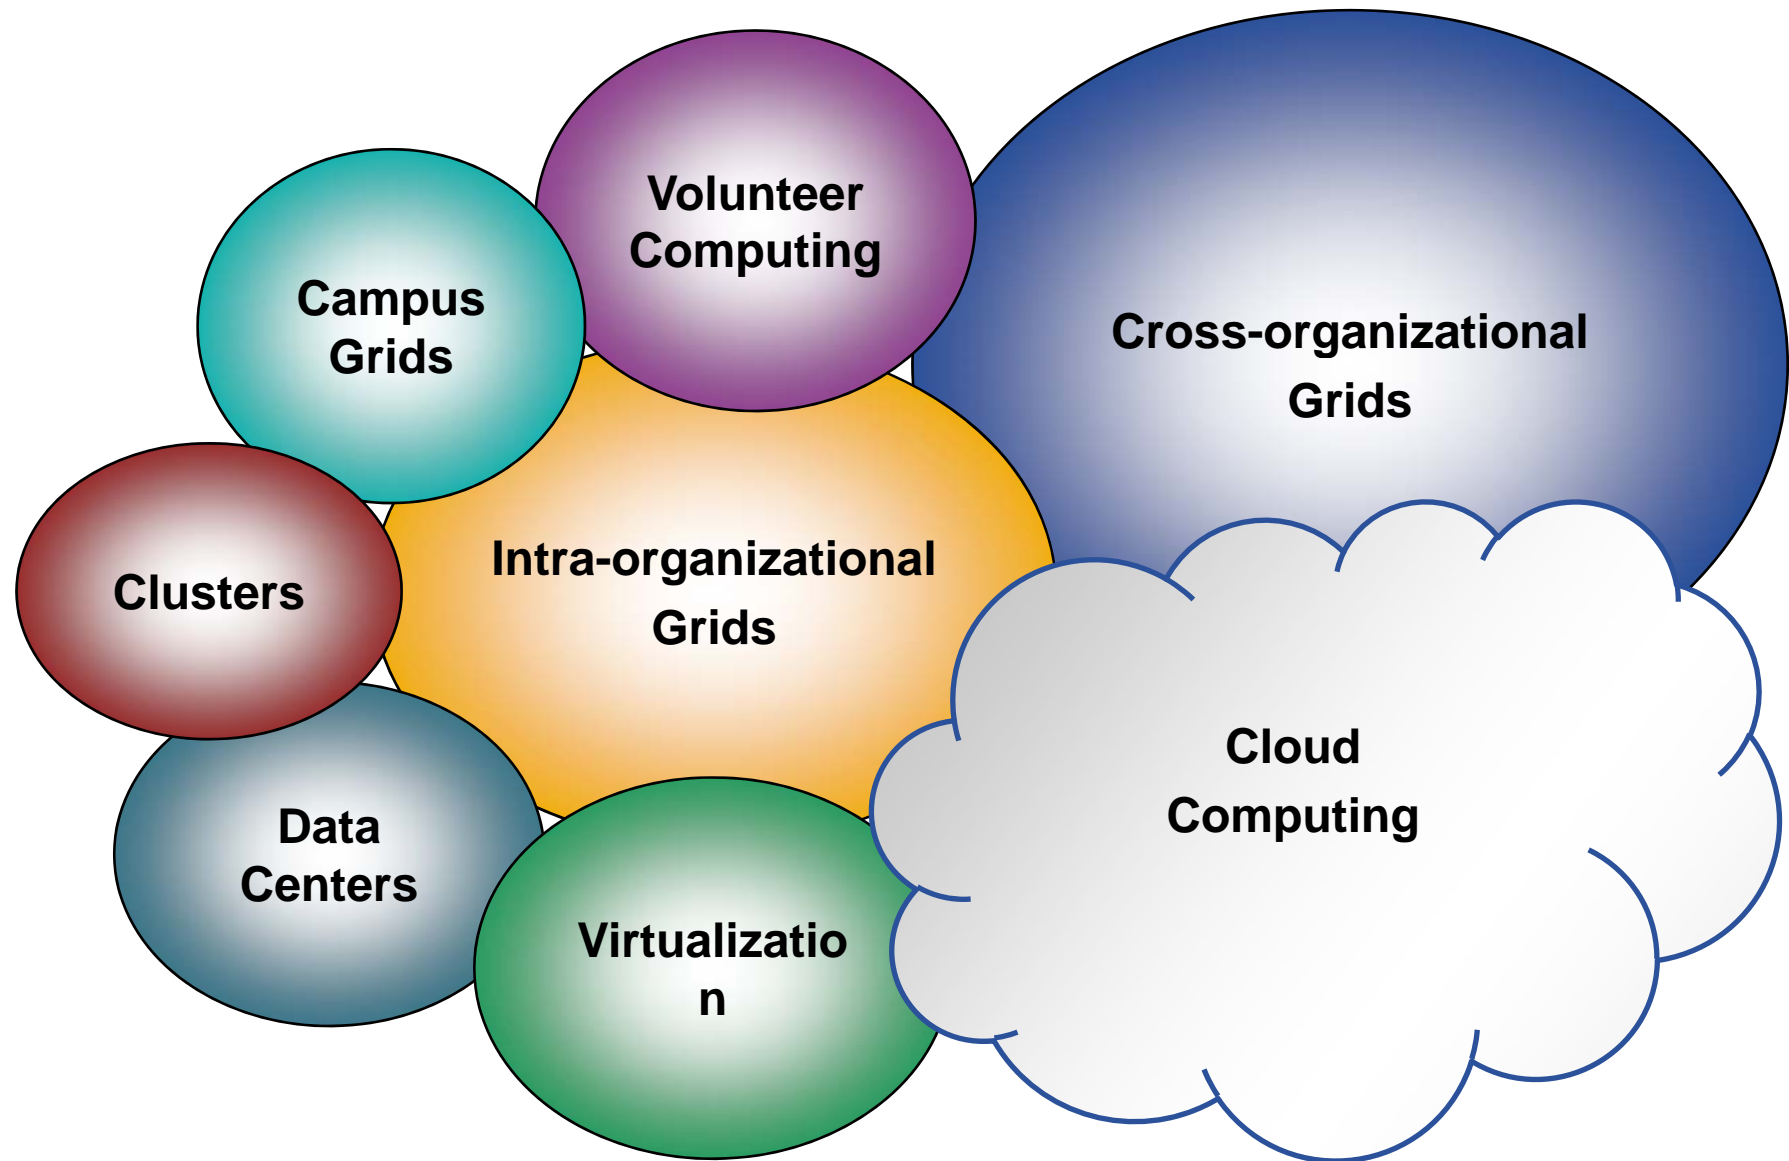
Laurence Field

OGF 25



www.eu-egee.org




- **Introduction to the LSF BES interface**
- **The test environment and methodology**
- **Test results for direct submission**
- **Test results for submission via the WMS**
- **Evaluation of the BES Interface**



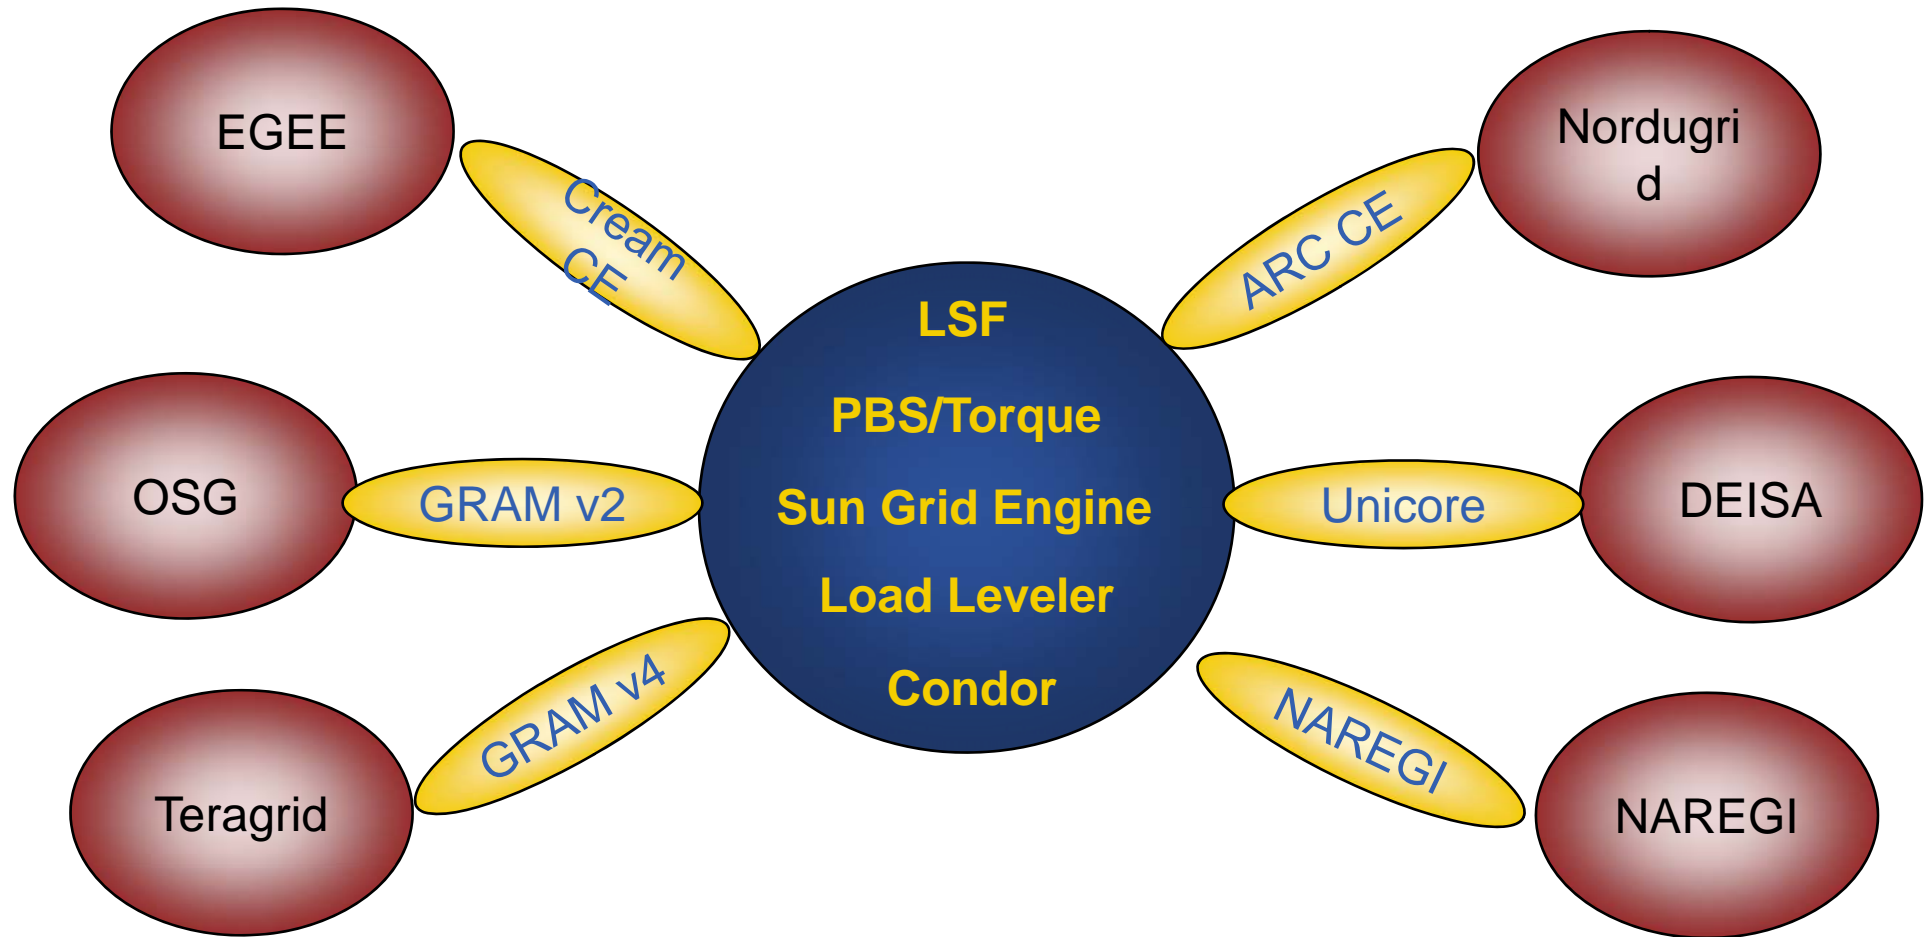
What is the problem?



- **Organization A and B are administrative domains**
 - Independent policies, systems and authentication mechanisms
- **Users have local access to their local system using local methods**
- **Users from A wish to collaborate with users from B**
 - Pool the resources
 - Split tasks by specialty
 - Share common frameworks



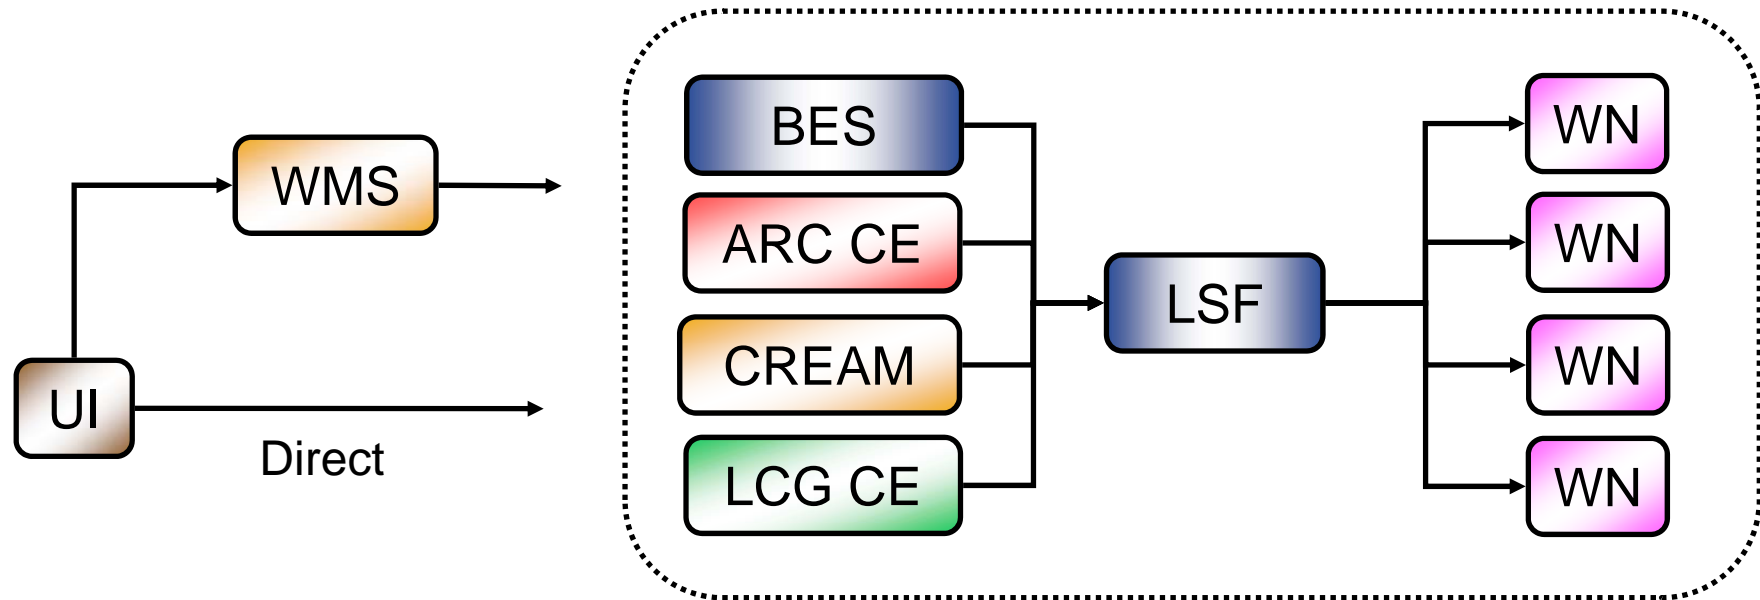
- **The Users from A and B create a Virtual Organization**
 - Users have a unique identify but also the identity of the VO
- **Organizations A and B support the Virtual Organization**
 - Place “grid” interfaces at the organizational boundary
 - These map the generic “grid” functions/information/credentials
 - To the local security functions/information/credentials
- **Multi-institutional e-Science Infrastructures**



• There are as many Computing Interfaces as Batch Systems!

- Basic Execution Service (BES)
 - “A service to which clients can send requests to initiate, monitor, and manage computational activities.”
 - OGF Grid Final Documents (GFD) 108
- **Load Sharing Facility (LSF)**
 - A commercial job scheduler from Platform Computing.
 - BES Interface is available in LSF version 7.0

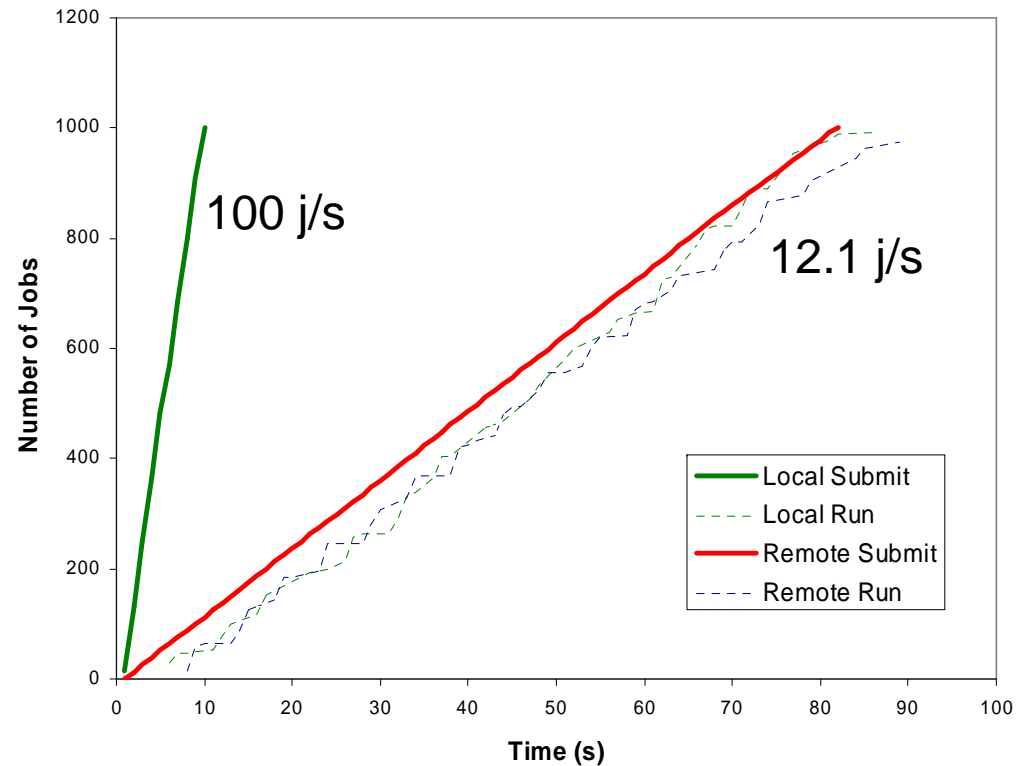
- **Outlined by the Technical Coordination Group (TCG)**
- **Performance**
 - 5000 simultaneous jobs per CE node
 - 50 user/role/submission combinations supported on a single CE
- **Reliability**
 - Job failure rates in normal operations due to the CE <0.5%
 - Job failures due to restart of CE services or reboot <0.5%
 - 5 days unattended running with no performance degradation



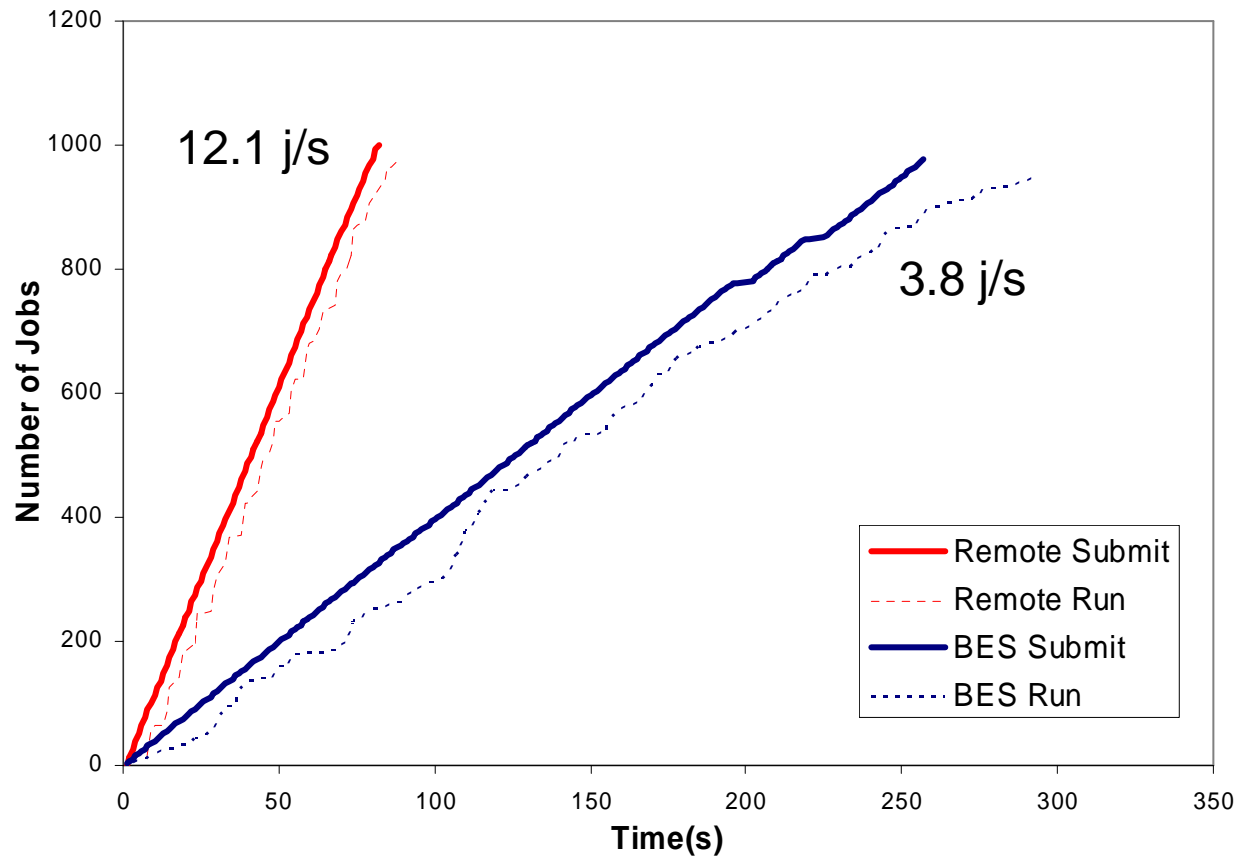
- **Hardware; Intel Xeon 2.33GHz, 8 cores, 16 Gig**
- **Worker Nodes (in VMs) configured with 40 jobs slots each**
- **All CE Interfaces co-hosted on the same physical machine**
- **UI and WMS on separate machines**

- **Direct Submission**
 - Job: Simple bash script running the “hostname” command
 - Submitting 1000 jobs in a loop as a single user
 - Monitoring the LSF log file for timestamps
- **WMS Submission**
 - Job: Simple parametric jobs running the “hostname” command
 - From four separate users
 - Monitoring the Logging and Bookkeeping service for timestamps
- **Metric: $j/s = \text{jobs per second}$**

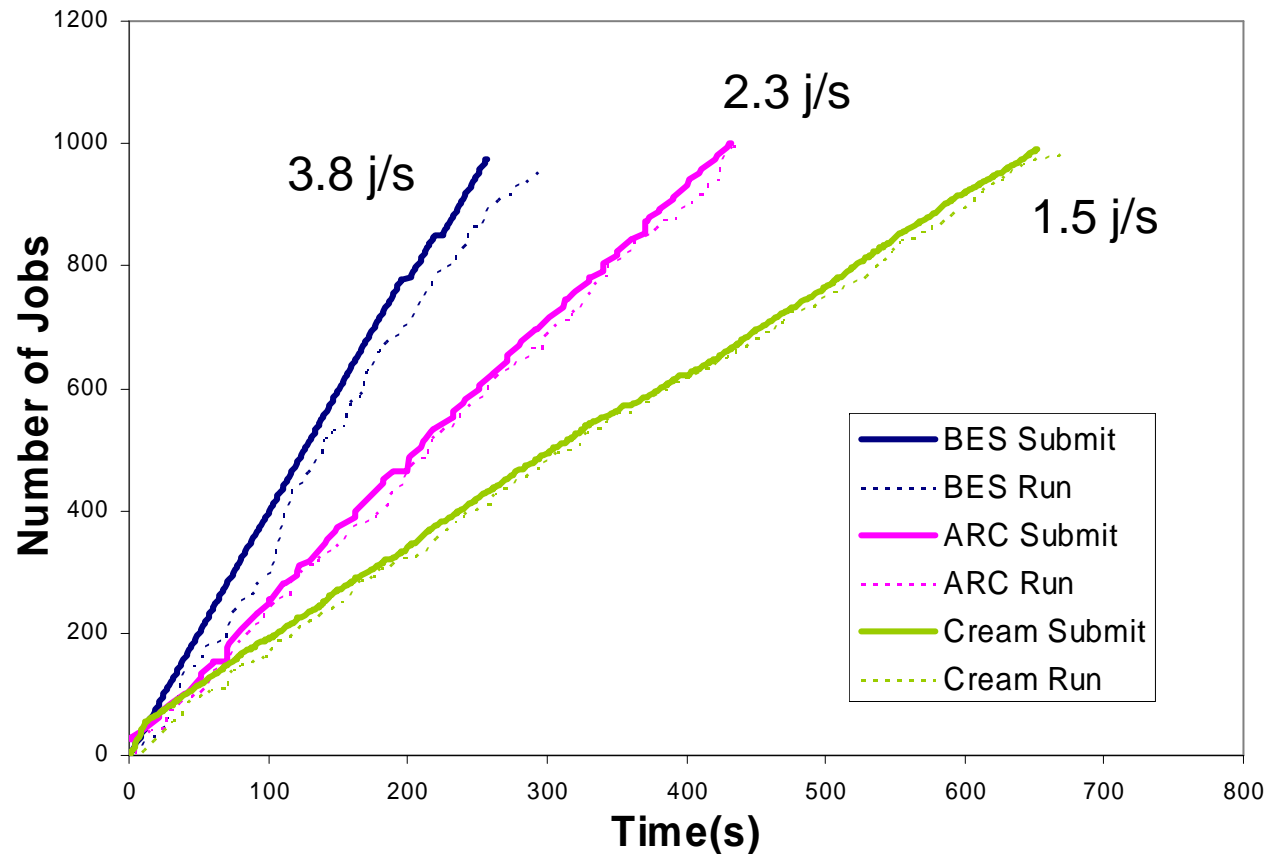
Comparison between submitting jobs to LSF using the bsub command locally and remotely




Comparison between using the LSF bsub command and the LSF BES client.



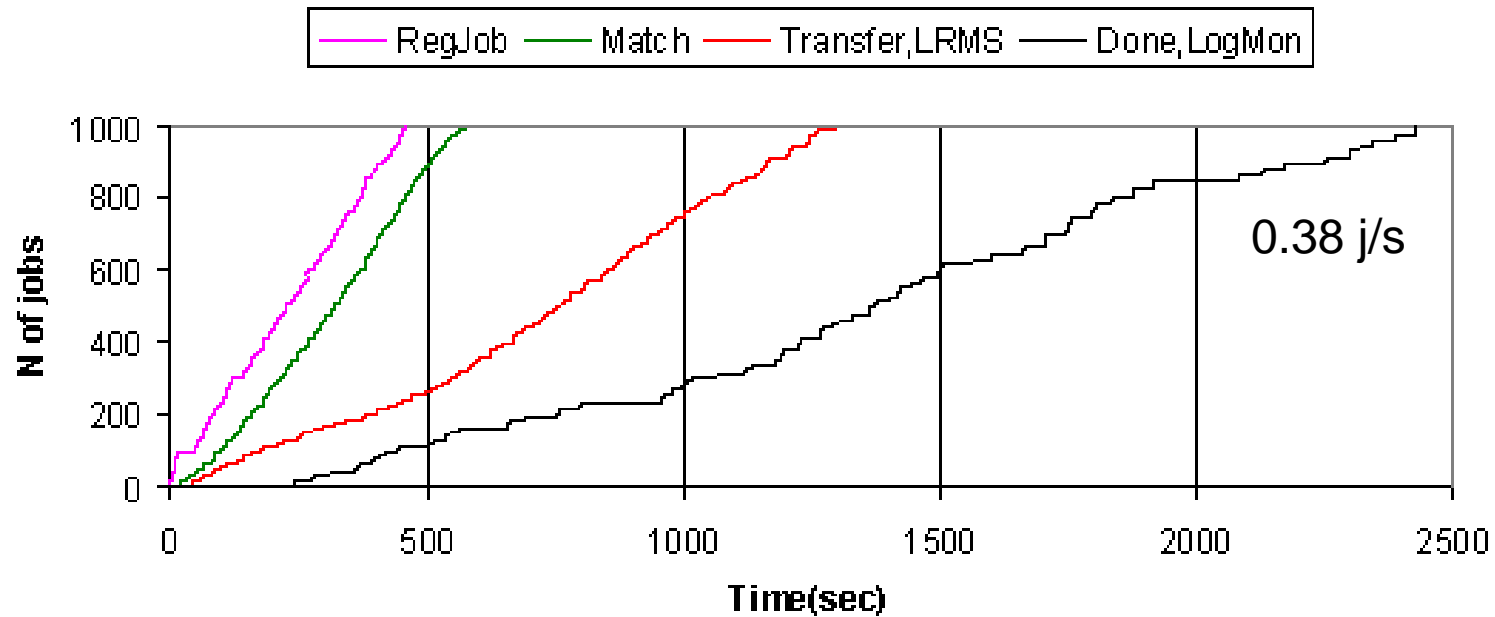
Comparison between the BES, ARC and Cream Interfaces



Job Throughput for the Difference Interfaces

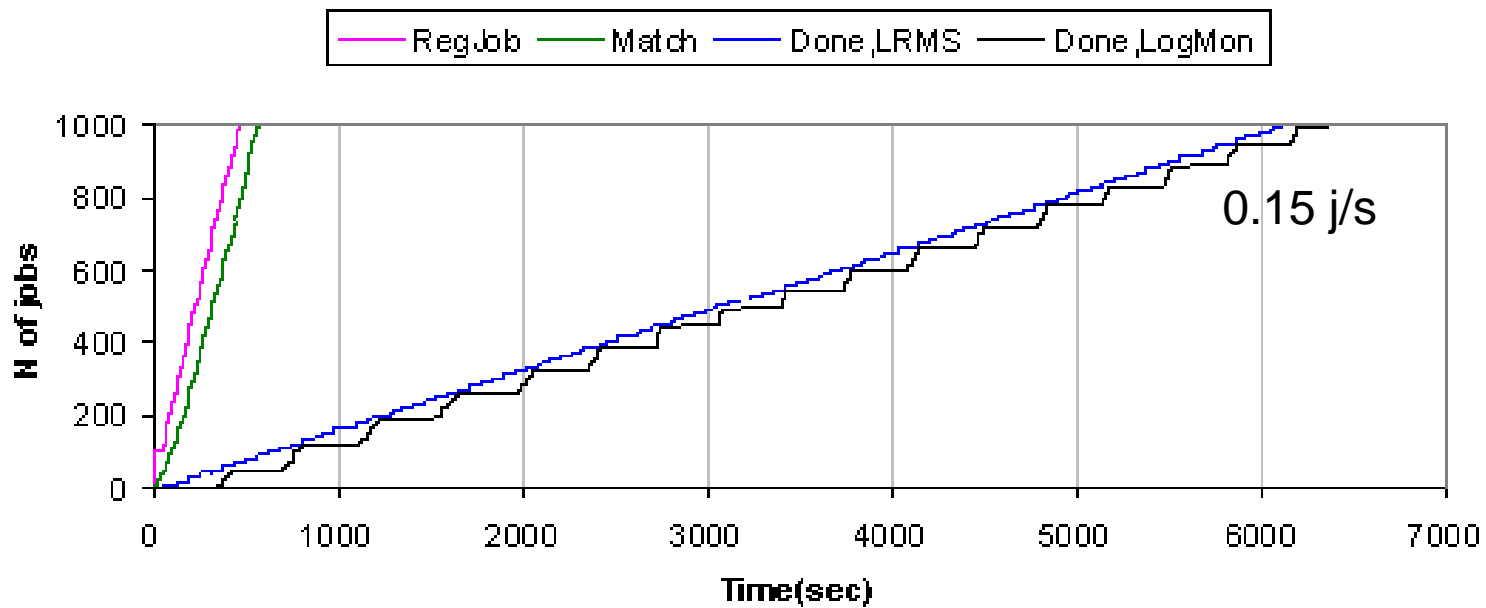


1000 jobs from 10 users (lcg-CE)




Total test time: 0 hour 43 min 2 sec

1000 jobs from 10 users (arc-CE)




Total test time: 1 hour 48 min 17 sec

1000 jobs from 10 users (cream-CE)



Total test time: 0 hour 38 min 20 sec

Job Throughput for the Difference Interfaces



- **Other factors affect performance of real use cases**
 - Size of the sandbox transfers (jobs and input data)
 - Number of parallel submission
 - From different users
 - Bulk job submissions
- **In-depth testing is required for a real comparison**
 - Simple throughput tests only offer rough indications
 - Mainly on the technology used
 - Doesn't highlight different advantages

- **Why in HPC are low performance interfaces used?**
 - Both Web Services and the Security Model affect performance
 - Can not function without the security model
- **Performance of BES could be acceptable for EGEE**
 - Improvements can be made (bulk job transfers etc.)
- **The current implementation can not be directly used**
 - Due to missing functionality (security model not supported etc.)
- **This BES specification is missing these pieces**
 - Could be solved by defining a profile
- **Production Grid Infrastructures PGI-WG**
 - Aims to produce this missing profile
 - So that BES can be adopted by production infrastructures.

- **Different interfaces to LSF have been tested**
 - Via direct submission and via the WMS
- **Results mainly show the affect of the technology used**
 - The performance overhead of Web Services
 - Is similar to that of the security model
- **The LSF BES performance interface is acceptable**
 - However, it is missing required functionality
- **There are missing pieces in the BES specification**
 - These need to be defined if it is to be adopted by EGEE
 - The PGI-WG aims to address this by producing a profile.

- **Many thanks to the following people**
 - Chris Smith (Platform Computing)
 - Nikolai Klopov (PNPI)
 - Andrey Kiryanov (PNPI)
 - Kalle Happonen(HIP)
 - Di Qing (ASGC)