Adaptive Grid Scheduling and Resource Management

Aleksandar Lazarević
PhD Student
E: a.lazarevic@ee.ucl.ac.uk W: www.ee.ucl.ac.uk/~alazarev/
Tutor: Dr. Lionel Sacks

The Grid

"The Grid" is many things to many people, and is becoming a valuable buzz-word in IT sale brochures worldwide. Behind the hype that is catching momentum stands an enabling technology promising to deliver ubiquitous high performance computing and virtualization of specialized scientific resources across geographical and administrative boundaries. The Grid presents a unified interface to instruments, displays, computational and information resources enabling Grid applications to integrate them in persistent environments

SO - GRM

"Self-organized Grid Resource Management" is an EPSRC funded e-Science project in cooperation with BT ExacT.
SO-GRM draws its motivation from the lack of new-breed, adaptive, autonomous and robust resource management and scheduling frameworks for heterogeneous computing environments.
SO-GRM researches several key areas of Grid computing:

1. High-level policy based VO management and Service Level Agreement (SLA) negotiation
2. Light-weight, resilient and self-organizing resource monitoring and discovery methods
3. Adaptable and dynamic job scheduling based on previously observed performance data and future predictions
4. Run-time process auditing and security monitoring

Framework

Components of SO-GRM framework are built using widely accepted open-source technologies:

1. XML Policies stored in open-source databases such as eXist
2. Integration of Globus Security Infrastructure (GSI) and using Public Key Infrastructure (PKI) based on X.509 certificates
3. Scalable monitoring system based on Ganglia Monitoring Toolkit exchanging XML messages via broadcast/ multicast/unicast
4. Support for various local job managers overlaid with Globus GRAM
5. Wide operating system support

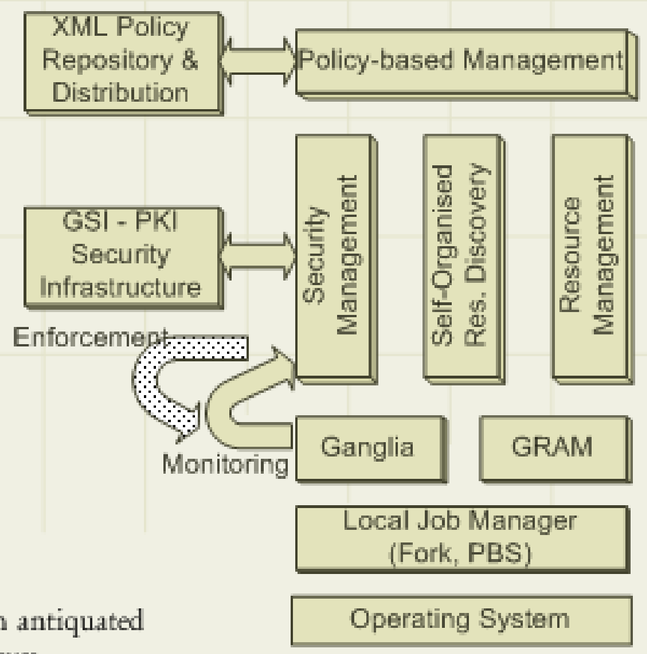

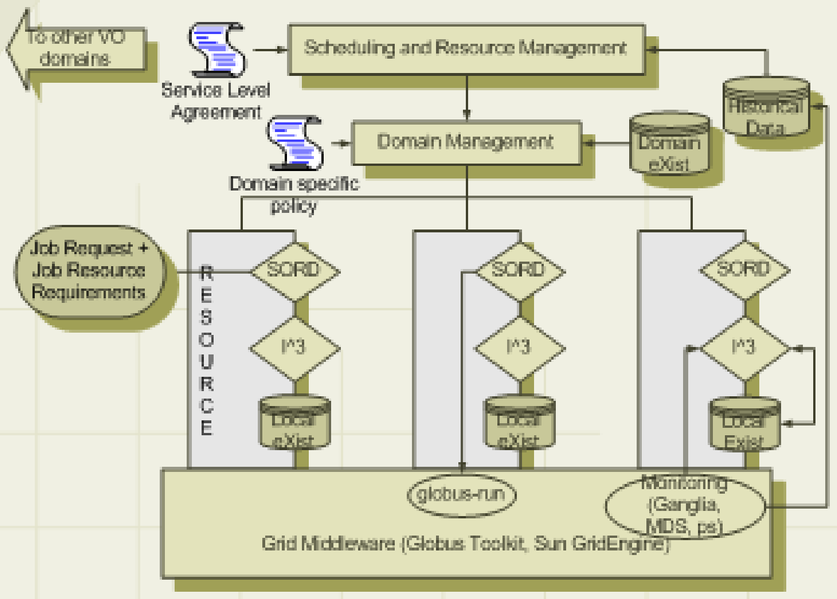

Components

Many current scheduling/queuing approaches are based on antiquated master-worker model. The performance of the master gateway determines the service rates for incoming jobs and, being the only route for job submission, renders it a single point of failure.

SO-GRM approach is to develop a distributed system of independent compute nodes, all able to accept job requests and choose whether to execute or forward them to a more appropriate node. The levels of acceptable loading, quality of service guarantees and SLA enforcement can all be set through policy-enabled management. Multiple job entry points and self-organizing information dissemination based on small worlds topology should significantly increase resilience.

Test-bed

Component integration and testing was done on a test-bed consisting of machines in UCL and BT ExacT at Ad@stral Park. Machines were separated in two distinct administrative domains and connected through a routed IP network LearNET. Services, data and components required for successful job execution were scattered among the machines in both domains to demonstrate integration within a Virtual Organization. Job submission was performed from BT side, while demonstration was followed in UCL.

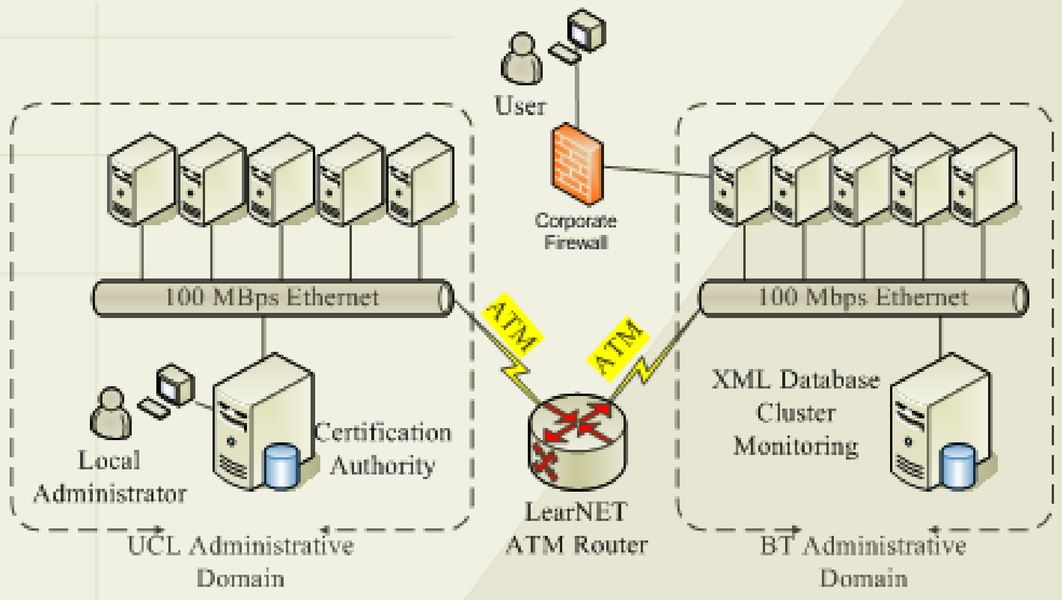

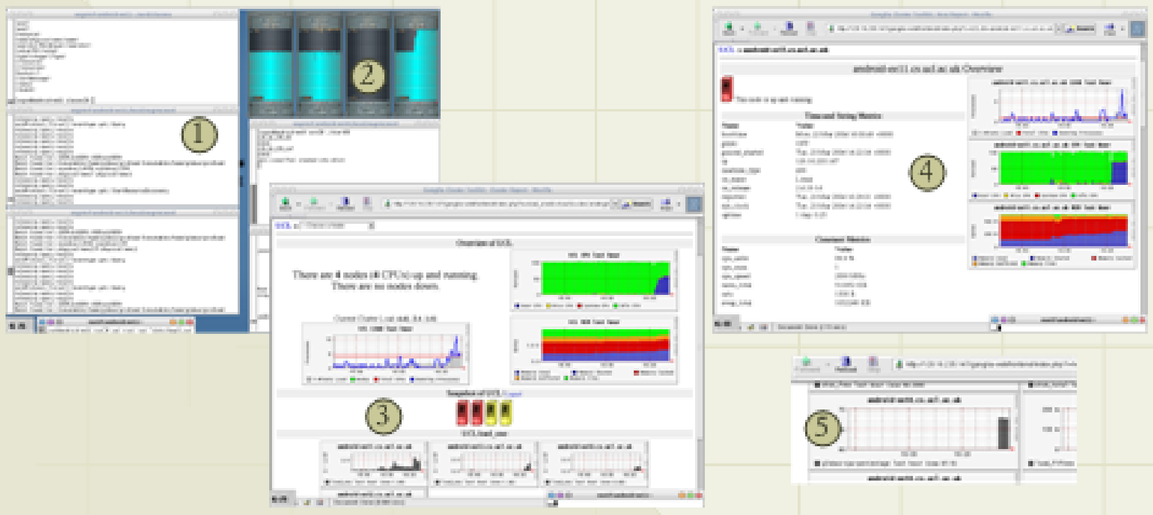

Demo

Thorough monitoring and information capture was essential during test runs, this task being further complicated by the probabilistic nature of SO-GRM's resource discovery and allocation algorithm.

1. Proper functioning of resource discovery is followed through debug messages
2. Machines are being assigned jobs. Load is only one of parameters taken into consideration, thus an already loaded machine may be assigned another job before an idle node.
3. Ganglia monitoring system showing a snapshot of UCL Grid
4. In-depth information on single execution node.
5. Additional information providers have been developed to report CPU utilization of Globus submitted jobs.