DAME
Engineering Grid solutions

One of the central challenges posed to the Distributed Aircraft Maintenance Environment (DAME) project, as an e-Science pilot project, was to address the question of the scalability, robustness and maturity of the emerging Grid for complex industrial engineering applications. The team has addressed this challenge within an industrial deployment scenario provided by Rolls-Royce and Data Systems & Solutions in the domain of remote health monitoring for civil aerospace engines.

The engineering scenario provided an ideal framework for investigating and developing Grid capability. Fault diagnosis and prognosis systems on this scale require access to high-volume, remote sensor data and benefit from on demand computing resources for data analysis. They involve complex and distributed interactions between different stakeholders and must provide supporting evidence for the diagnosis or prognosis offered. Commercial issues such as quality of service (QoS), dependability and security are also critical to successful deployment. The Grid computing paradigm offers one potential framework on which to build and manage systems to meet all these diverse and complex requirements.

The DAME demonstrator

Rolls-Royce has collaborated with Oxford University in the development of an on-wing monitoring system, called QUICK, for future generations of civil aeroengines. QUICK performs engine analysis on data derived from continuous monitoring of broadband engine vibration. This is complex signal data, and a single flight can produce up to 1 gigabyte, which when downloaded and archived amounts to several terabytes a year for an operational fleet. Therein lies one of the primary technical challenges and deliverables for DAME: producing a Grid-based data management architecture that permits distributed, seamless access to this data and which provides a secure, virtual diagnostic environment for its analysis.

The DAME proof of concept demonstrator was based on the idea of a diagnostic workbench environment, hosted within a secure Grid portal. This workbench is virtual, in the sense that the diagnostic and data management services are distributed across the Grid (the demonstrator services are deployed across the White Rose Grid). The portal brings together the advanced data-mining and data-management services that have been developed, in addition to a range of novel and Grid-based tools and diagnostic services, including visualisation and modelling tools, decision support systems, workflow management services and a workflow escalation system.

The Data-Mining service consists of the AURA pattern match system (developed by the University of York and now marketed through Cybula) which allows engine health monitor data to be searched for features and can cope with tera-scale data sets. This service facilitates the rapid searching of large, distributed engine data archives through the use of sophisticated pattern matching techniques. The data-mining architecture itself is inherently scalable, and is based around a deployment of the Storage Request Broker (SRB) system from the San Diego Supercomputer Center. A brokerage service allows search queries to be sent out to the data archives, via a call to the SRB, which allows abstraction away from the physical locations of the data repositories. The architecture allows the AURA data-matching capabilities to be deployed across substantially greater datasets than has previously been possible. As such it will be one of the major generic deliverables from DAME to the e-Science community, as it can be redeployed across any domain requiring a high-performance, high-volume pattern matching capability (for example, in medical or high energy physics applications).

Decision support capabilities are provided through two mechanisms; Case Based Reasoning (CBR) tools and the Engine Modelling Service, both of which have been developed by the University of Sheffield. DAME's CBR tools have contributed novel methods for capturing and managing workflow and diagnosis processes within a Grid-based system. The CBR service manages a knowledge base and captures the DAME fault diagnosis methods in a procedural way. Hence, previous expertise in diagnosing fault conditions is captured and used via a workflow advisor system to prompt engineers with procedural information and automatic workflow scripts to deal with known problems. These knowledge capture methods build on provenance and meta-data schemes that have been developed by the University of Leeds, and for which Grid service architectures are well suited.

A second application of CBR is to build and maintain the DAME knowledge base that correlates observed QUICK engine anomalies with the results of root cause investigations by the various engine maintenance, repair and overhaul (MR&O) organisations (for example, from flight-line maintenance technicians to engine overhaul contractors). The learning achieved with CBR tools supports continuous improvement of the diagnosis and prognosis application. Again, these CBR methods are generic in nature and may be re-deployed in any health monitoring domains, for example medical.

The Engine Modelling Service that has been built provides an ideal application for demonstrating the Grid's ability to provide computing resources on demand. The complex, compute intensive models can be launched from the portal, accessed from any so-called thin client hosting a browsing tool, giving data analysts instant access to powerful modelling capabilities. The services take parameters from flight data and run thermodynamic models of the engine and so can infer its current state. The system can also simulate different flight conditions such as idle, take-off, climb and cruise. The team has also developed security models for deployment of commercially sensitive applications.

In addition to diagnostic services, the portal has had to address important commercial drivers such as QoS, security and role and access management. To this end, the University of Leeds has been developing advanced workflow management and brokerage systems for the portal. A novel 'three-phase commit' broker has been implemented that allows optimal assignment of tasks to the available Grid computing resources, and in such a way that Service Level Agreements (SLA) can be assigned to job schedules. SLA's permit engineers to specify the commercial constraints associated with a deployed task. Some tasks may have greater urgency or criticality than others, and the Leeds SLA system allows these parameters to be specified and acted upon by the broker. These services provide important and novel capabilities to investigate the commercial realisation and deployment of Grid computing. The portal also hosts a fully deployed Grid security infrastructure, based upon X509 digital certificates. Detailed dependability analysis and asset risk analysis have underpinned the definition of this security infrastructure.

In summary, DAME has been addressing many important issues relating to the deployment and scalability of Grid computing in complex engineering domains. Novel and robust methods for large-volume data-management and mining have been developed, in addition to commercially realisable methods for workflow management, QoS management and decision support tools for domain knowledge capture. All of these techniques are inherently generic in their deployment and could be applied to many other domains: medicine to engineering and transport to aerospace. The proof of principle provided by the DAME test-bed has been an important step towards understanding how the Grid computing paradigm may shape the future of system developments for engineering and health-monitoring applications.

DAME collaborators

– Universities of York, Leeds, Sheffield and Oxford
– Rolls Royce plc
– Data Systems and Solutions LLC
– Cybula Ltd

Website

www.cs.york.ac.uk/dame

Contacts

Professor Jim Austin (austin@cs.york.ac.uk)
Dr Tom Jackson (tom.jackson@cs.york.ac.uk)