

Delivering a Grid enabled Distributed Aircraft Maintenance Environment (DAME)

Tom Jackson, Jim Austin, Martyn Fletcher, Mark Jessop
Advanced Computer Architectures Group (ACAG), Dept of Computer Science,
University of York, York, YO10 5DD.

[tomj, austin, fletcher, mark.jessop] @cs.york.ac.uk

Abstract

We provide an overview of the DAME project, and a discussion of the progress made to date on the development of a distributed aeroengine diagnosis environment as a proof of concept demonstration for Grid computing. We discuss the challenges faced by the DAME project in meeting the requirements of this complex, data intensive, diagnosis system that must be operated as a distributed ‘virtual organisation’. We describe the core diagnosis applications that have been implemented as Grid services, and explain how these services are being deployed within the overall diagnosis process. We conclude with some remarks on our experiences in developing these Grid services under Globus Toolkit 3, and with an overview of the future development plans for the DAME demonstrator.

1 INTRODUCTION

The Distributed Aircraft Maintenance Environment (DAME) is a pilot project funded by the EPSRC e-Science programme. DAME is undertaken in partnership with Rolls-Royce plc who have provided the aeroengine data for the diagnostic system, Data Systems and Solutions, a company jointly owned by Rolls-Royce and SAIC, who are Rolls-Royce data systems providers and currently deliver commercial aero engine health monitoring services, and Cybula, a York University spin-off managing the data storage technology in the project. The academic partners in the project are the Universities of York, Leeds, Sheffield and Oxford.

The theme of the DAME project is the design and implementation of a fault diagnosis and prognosis system based on the Grid computing paradigm and the deployment of Grid services. In particular, DAME focuses on developing an improved computer-based fault diagnosis and prognostic capability and integrating that capability with a predictive maintenance system in the context of aeroengine maintenance. Here, we use the term “predictive maintenance” to mean that there is sufficient time interval between the detection of a behaviour that departs from normal (via a fault threshold) and the actual occurrence of a failure. The DAME system will deploy Grid services within this time window to develop a diagnosis of why an engine has deviated from normal behaviour, to provide

prognosis (understanding what will happen) and to plan remedial actions that may be taken a safe and convenient point when the impact of maintenance is minimized. We begin with an outline of the characteristics of the problem domain and illustrate the reasons why it is amenable to a Grid based solution.

Fault diagnosis techniques are applicable across many diverse IT domains, for example, medicine, engineering, transport, and aerospace. However, regardless of the application domain, fault diagnosis and prognosis systems share a number of operating and design characteristics:

- These systems are data centric. Monitoring and analysis of sensor data and domain specific knowledge is critical to diagnostic process;
- They typically require complex interactions among multiple agents or stakeholders;
- They are often distributed;
- They need to provide supporting or qualifying evidence for the diagnosis or prognosis offered;
- They can be business critical, and typically have stringent dependability requirements.

The emerging Grid computing paradigm [1] appears to offer an inherently practical framework

in which to build and manage systems to meet these requirements.

2 THE DEMONSTRATION CONTEXT

The context for the demonstrator being developed within the DAME project is an aeroengine diagnosis and prognosis problem. Modern aero-engines operate in highly demanding operational environments and do so with extremely high reliability. To achieve this, the engines combine advanced mechanical engineering systems with tightly coupled electronic control systems. As one would expect, such critical systems are fitted with extensive sensing and monitoring capabilities for performance analysis. To facilitate engine fleet management, engine sensor data are routinely analyzed using the COMPASS health monitoring application developed by Rolls-Royce and prognostic applications employed by Data Systems and Solutions. The resulting commercial services are subscribed to by many commercial airlines.

The basis of monitoring is to detect the earliest signs of deviation from normal operating behaviour. COMPASS achieves this by comparing snap-shots of engine sensor data against ideal engine models. The relatively small data sets may be transmitted in flight or downloaded once on the ground. There is scope to increase the effectiveness of monitoring by looking at more data in greater detail. For this reason, Rolls-Royce has collaborated with Oxford University and has developed an advanced on-wing monitoring system called QUOTE [2]. QUOTE performs engine analysis on data derived from continuous monitoring of broadband engine vibration. The analysis is achieved through data fusion of the vibration data with instantaneous performance measurements. QUOTE does not store data from many flights or cross-reference data from the rest of the fleet. However, a ground based system can maintain fleet wide databases of flight data and other maintenance related information and can use this additional data to perform various analyses. This analysis will enable unknown anomalies to be correlated to root causes and appropriate remedial actions taken.

The DAME project offers the prospect of combining the bandwidth rich QUOTE approach with the sophisticated time series, fleet-wide repositories available to COMPASS with the goal of an enhanced ability to anticipate maintenance requirements.

Developing Grid-based diagnostic systems to facilitate the processing of data in a ground based system presents the DAME project with three principal challenges:

- The type of data captured by QUOTE involves real valued variables monitored over time. An example of this is shown in Figure 1. This plot is typical of the data stored on the ground and utilised to support the diagnostic process. Each flight can produce up to 1 Gigabyte of data, which, if scaled to the fleet level, implies many Terabytes of data per fleet per year. The storage of this data will require vast data repositories which will be distributed across many geographic and operational boundaries, but which must be accessible for health monitoring services;
- Advanced pattern matching and data mining methods must be developed to detect features and analyse the type of data produced by the engine. These methods must be able to operate on Terabytes of data and give a response time that meets operational demands;
- The diagnostic processes require collaboration among a number of diverse actors within the stakeholder organizations, who may need to deploy a range of different engineering and computational tools to analyse the problem. Thus, any Grid-based solution must allow a Virtual Organisation (VO) to support the services, individuals and systems involved.

In the following section we describe the core processing technologies being developed to address these challenges as well as discussing the implications on the Grid middleware required to support the system integration.

3 DAME GRID SERVICES

To support the DAME requirements the system must provide a range of diagnostic and prognostic services, and associated data and processing services. The environment is highly dynamic, and requires dependable and timely operation. The DAME project, now half way through its three-year funding period, is building a proof of concept demonstrator to determine the effectiveness of Grid middleware to support this task.

This proof of concept demonstrator is addressing several issues, including:

- The capability of Grid to support complex processing systems distributed across virtual organisations;
- Robustness of the Grid, including issues of security and availability;
- Performance of the Grid for large volume data management.

The proof of concept demonstrator is based upon 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). A functional view of the service architecture is provided in Figure 2.

DAME has been an early adopter of the Globus Toolkit 3 (GT3) and the OGS/OGSA Grid services model. Hence, after some early development work deploying web services under GT2, the DAME applications have now been ported and deployed as Grid services, running under GT3, release alpha 3.0. The primary motivation for making the early move to GT3 was the requirement for more rigour in the security mechanisms (DAME is working with commercially sensitive data). The early adoption of GT3 has caused some implementation problems, but these are discussed separately in a later section.

Below we describe the core services that have currently been implemented and deployed.

3.1 DAME Engine Data Service

This consists of the QUOTE system, sensors and associated ground links. The Engine Data Service controls the interactions between the on-engine monitoring system (QUOTE), and the communications to a ground station, which establishes the link to Grid data repositories. Since aircraft land in many parts of the world, there will be many replications of this service, each of which will be transient in nature, only existing for long enough for the flight and monitoring data to be transferred from the aircraft to a ground station. A simulated data download application has been built that permits the DAME project to simulate the real-time arrival of flight data and it being archived in the Grid based data repositories (hosted across the White Rose Grid, and at a remote node within Oxford University). The data repositories are currently implemented under MySQL, although in the long term DAME may move to Oracle database technology to improve integration with the commercial partners existing systems.

3.2 Data Mining Service

The Data Mining Service consists of the AURA pattern match system[3] which allows engine health monitor data to be searched for features. This service provides the capability to rapidly search the large, distributed engine data archives through the use of sophisticated pattern matching techniques. The AURA system itself is available in both hardware or software implementations. DAME is currently using a software implementation, built upon the latest AURA software libraries, and for which a Grid service interface has been developed. The use of the software implementation of AURA allows the DAME project to explore the notion of distributed data mining, by permitting multiple instances of the AURA service to be deployed across the available Grid computing resource.

The AURA methods provide search techniques for tera-scale datasets. In the DAME demonstrator context, the AURA system returns vibration data that best matches the anomaly conditions found on the engine. If similar abnormal events are found in the data archive then any supporting information relating to these events, such as maintenance steps or remedial corrective action, can be recovered from the appropriate operations databases that are visible on the Grid. These databases are made available to the Grid through OGS/OGSA-GS (and potentially through an OGS/OGSA-DAI implementation). As the AURA search may return thousands of potential data items, the use of distributed resources is vital to achieve a timely delivery of information.

3.3 Engine Modelling Service

The Engine Modelling Service, developed by Sheffield University, provides the ability to take parameters from flight data and run models of the engine. The aim is to infer the current state of the engine – in effect to perform model based data fusion. This is intended to be an improvement on the engine modelling service currently employed in COMPASS.

A performance model of a Rolls-Royce aero gas turbine engine has been ported into a Grid enabled application which allows the model to be run on Grid computing resource as a Grid Service. Based on GT3, this engine simulation Grid Service can be invoked simultaneously in different “Virtual Organisations” for different applications. The deployed model has a number of other benefits:

- It will be used to assist the experienced maintenance engineers in the fault diagnosis process via on-line simulation;

- It provides GSI enabled secure engine performance simulation for different flight operational conditions and requirements, e.g. Idle, Take-off, Climb and Cruise;
- The Factory Service can generate a group of engine simulation instances for different client requirements
- Both Transport Level and Message Level Security are implemented to protect the secure sensitive engine model and user data.

3.4 Case-Based-Reasoning (CBR) Service

DAME incorporates Case Based Reasoning tools for decision support in the diagnosis process. The CBR system refines the knowledge base and captures the DAME fault diagnosis methods in a procedural way. One potential procedural application is to manage workflows associated with diagnostic operations. A second application is to build and maintain the DAME knowledge base that correlates observed QUOTE engine anomalies with the results of root cause investigations by the various engine maintenance, repair and overhaul (MR&O) organizations (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.

The CBR application has been developed by Sheffield University using an established commercial package, for which a Grid Service/JAVA interface has been developed. Aeroengine maintenance personnel are able to access the service via a secure Grid portal using a web browser on any computer connected to the Internet. The maintenance personnel have access to stores of accumulated maintenance history and operational data as well as large computing resources to support the fault analysis and the decision-making process. A Service Factory supports the creation of multiple instances of the CBR service on the Grid, allowing many CBR processes to be executed in parallel from a single service access point. The advantages are of this approach are:

- Allows search of an extensive casebase of historical maintenance incidents across a fleet of engines, as opposed to just the local available data;
- Secure environment in which to store sensitive data;

- High performance computing resources available to perform analysis across large datasets;
- The ability to integrate performance modelling and simulation with fault and maintenance data for more accurate fault analysis;
- Easier to support and update the system.

The future developments of the CBR services will incorporate workflow support. It is intended to capture information that helps validate or refine the output from the diagnosis processes. This will partly be achieved through the capture of workflow steps used to resolve anomalous engine events. This link to maintenance, repair, and overhaul stakeholders is essential for “closing the loop”, capturing lessons learned, and driving continuous improvement.

3.5 Data Analysis Services

In addition to the decision support services described above, the workbench provides detailed data analysis services that will be deployed by domain experts in tracing the root cause of anomalous events. The current basis of the data analysis tools in DAME is a ground-based version of the QUOTE system. This service, ported and hosted by the University of Leeds, allows fine-grained, interactive analysis of the vibration data, using diverse versions of the feature analysis algorithms developed for the QUOTE system. The concept of a ground based data analysis system allows the option for feature detection algorithms that are more compute intensive to be run, or for Monte-Carlo type simulations to be carried out on fault data. In addition to these Grid enabled analysis functions, the workbench environment also provides data visualisation tools for interactive assessment of the fault data. These services provide domain experts with a deeper insight into the fault symptoms. Through the use of collaborative working tools in the DAME portal, different actors involved in the diagnosis process can share in the results of the processing, permitting remote experts to participate concurrently in the fault analysis.

4 RELATED RESEARCH ACTIVITIES

In addition to development of the core diagnosis services described above, the DAME project is also exploring other important strategic issues relating to the commercial deployment of a Grid based system such as DAME. Two of the most important of these are Quality of Service (QoS) and Security.

Leeds University are leading studies into network performance measurement as the basis for QoS predictions. A local test-Grid has been built to support these activities. QoS assessment is not simply an abstract theoretical issue for the DAME system, but is a fundamental requirement for underpinning commercial deployment of Grid services, where Service Level Agreements (SLAs) are likely to be a necessity. This work is reported in detail elsewhere in these proceedings.

Security is another vital deployment concern for the DAME system, which has stringent confidentiality issues associated with the data and processing being managed across the Grid. DAME therefore has a workgroup, led by the University of York, focussed on the issues of security analysis. This group is working closely with the commercial partners to develop risk and threat models, and to characterise assets in the DAME system. The studies are taking a broad remit, drawing from established software dependability techniques and seeking to establish a methodology for dependability analysis in distributed Grid systems. The longer-term objective is to identify the main threats and risks in deployed systems such as DAME, and to assess the efficacy of the OGSA/OGSI security models for managing these risks. The results of these studies will be reported through the UK e-Science Security Task Force working group.

5 EXPERIENCES OF GT 3

It is useful to highlight some of the experiences gained with the DAME project in implementing the core services under GT3. The DAME project initially looked at GT3 with technology preview 1, following the OGSA demo at GGF4, March 2002. Serious work with Grid services did not really commence until the release of preview 4, working through the alpha releases up to alpha 3.

One of the initial frustrations and hindrances in working with the alpha releases was the lack of supporting documentation for the toolkit. Hence a considerable amount of time and energy was spent getting basic GT3 services up and running.

The OGSI specs are relatively accessible and make Grid service technology seem reasonably straightforward. However, this apparent simplicity did not follow through to service implementation. The documentation that was provided was adequate for building simple services and getting security up and running. Once services were required that implemented interfaces other than the basic Grid Service interface, e.g. registries, factories and working

with Service Data Elements, life got quite difficult. Much time was spent taking apart example code, trawling through and being active on the Globus-developer news group, and a not too small amount of guess work. The bulk of the development work required converting each web service to a Grid service (i.e. porting to the appropriate new CoG libraries in the toolkit) along with implementing registries and factories and the associated service infrastructure. Having invested this effort in the learning process, the DAME developers are now, however, well placed to exploit the full scope of the OGSA/OGSI concepts.

It is worth noting that one major implementation frustration has been the significant change in the specification for factories and registries between the GT3 alpha release to the beta and initial release. The changes in the ServiceGroup/ServiceGroupEntry and the ServiceGroup/Registration model are significant and require extensive reworking of the DAME services. As a result, DAME has continued with the alpha 3 release for the demonstrator produced for this All Hands Meeting.

6 FUTURE PLANS AND DIRECTIONS

The DAME system is developing technology at a number of levels. The base services consisting of AURA-G, QUOTE, CBR and data storage provide universal components that can be used in many other applications. In particular Cybula is committed to providing Grid-enabled versions of the AURA search engine for applications as diverse as biometrics to engineering. Discussions have already commenced with the e-Diamond project to explore the use of AURA for data mining tasks in the domain of medical mammogram analysis.

The composition of the core services into a Diagnostic / Prognostic Workbench or Portal is another development level. The portal deploys Grid middleware services for the management of workflows through the diagnostic process in the demonstrator domain. Having developed and deployed the core diagnosis services there will now be increased focus on the workflow process itself. A basic workflow engine has already been incorporated into the DAME portal, which provides task brokerage and management of Grid security protocols (and in particular the handling of certificate proxies for authentication and validation of access to services). Future work will require the inclusion of a more flexible workflow configuration tool within the portal environment. DAME will be considering the

output of other e-Science projects (e.g. MyGrid) for software to address this requirement.

At a more general level, DAME is concerned with the notion of proof of concept. That is, using the Globus tool kits and other emerging Grid service technologies to develop a demonstration system, which will provide input to requirements capture process for developing standards in areas such as data mining and core Grid services.

Future directions also include more detailed investigations of the timeliness and dependability requirements of Grid based distributed diagnostic systems. The DAME demonstrator provides an excellent example of an advisory system that has soft real-time constraints: the value of the diagnosis/prognosis is dependent on the time at which it is delivered. A diagnosis/prognosis is of most value when there is time for the maintenance crew to perform further checks or minor maintenance without affecting the scheduled departure of the aircraft. Beyond this time, the prognosis is still of value as this work can be scheduled at a subsequent location. The DAME project will investigate the use of contract and anytime algorithms as well as appropriate process scheduling policies to ensure that diagnoses are delivered in time. In addition, a deployed DAME system requires high availability and reliability as well as

communications that ensure data integrity and confidentiality. A fully deployed system also needs to scale to tens if not hundreds of thousands of aircraft. The DAME project will investigate how these requirements can be met with a Grid based system.

7 ACKNOWLEDGEMENTS

This research was supported by Grant Number GR/R67668/01 from the Engineering and Physical Research Council in the UK and through contributions from Rolls-Royce plc and Data Systems and Solutions, LLC.

8 REFERENCES

- [1] I Foster, C Kesselman, J M Nick, S Tuecke, 2002, The Physiology of the Grid, An Open Grid Services Architecture for Distributed Systems Integration
- [2] Nairac A, Townsend N, Carr R, King S, Cowley P, Tarassenko L. A system for the analysis of jet engine vibration data. Integrated Computer-Aided Engineering, 53-65, 1999.
- [3] www.cs.york.ac.uk/aura , The AURA and AURA-G web site, Advanced Computer Architectures Group, University of York, UK.

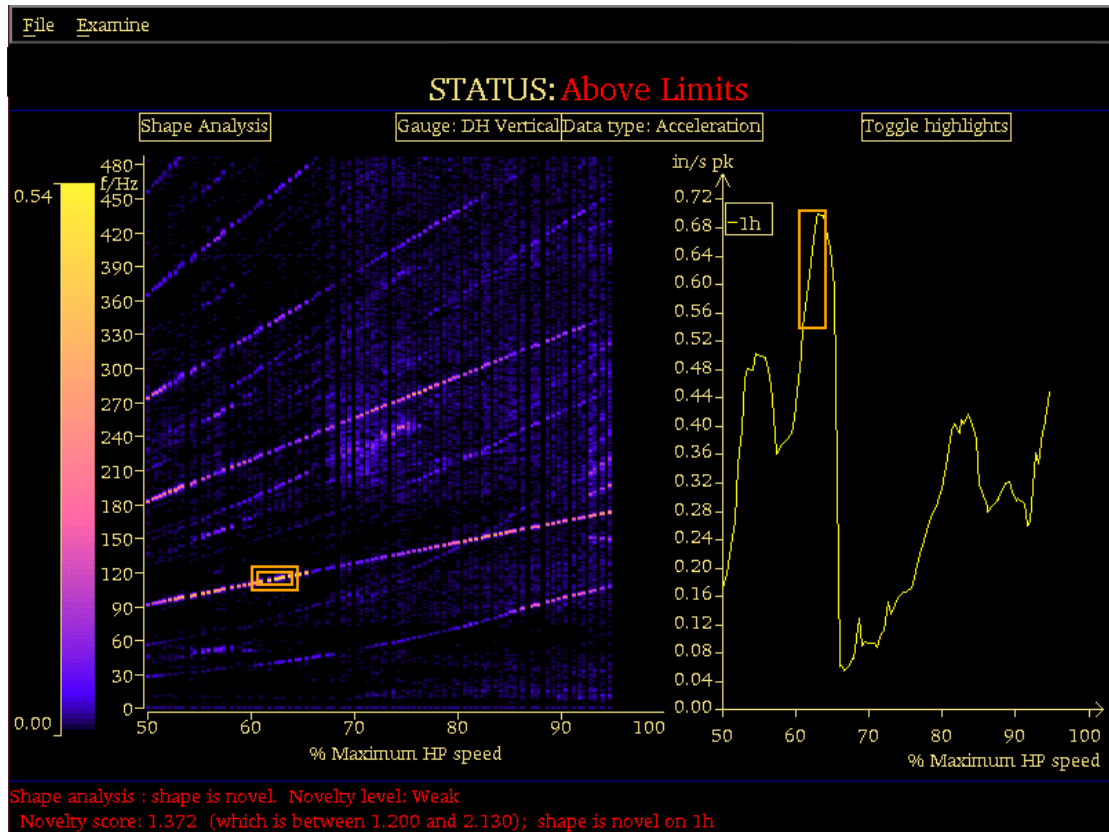


Figure 1. Typical engine vibration data captured by the QUOTE system. On the left is a plot of vibration against shaft speed. The graph shows a box where there is an anomaly, expanded on the right. The right plot shows one ‘tracked order’ and the same anomaly as indicated on the left.

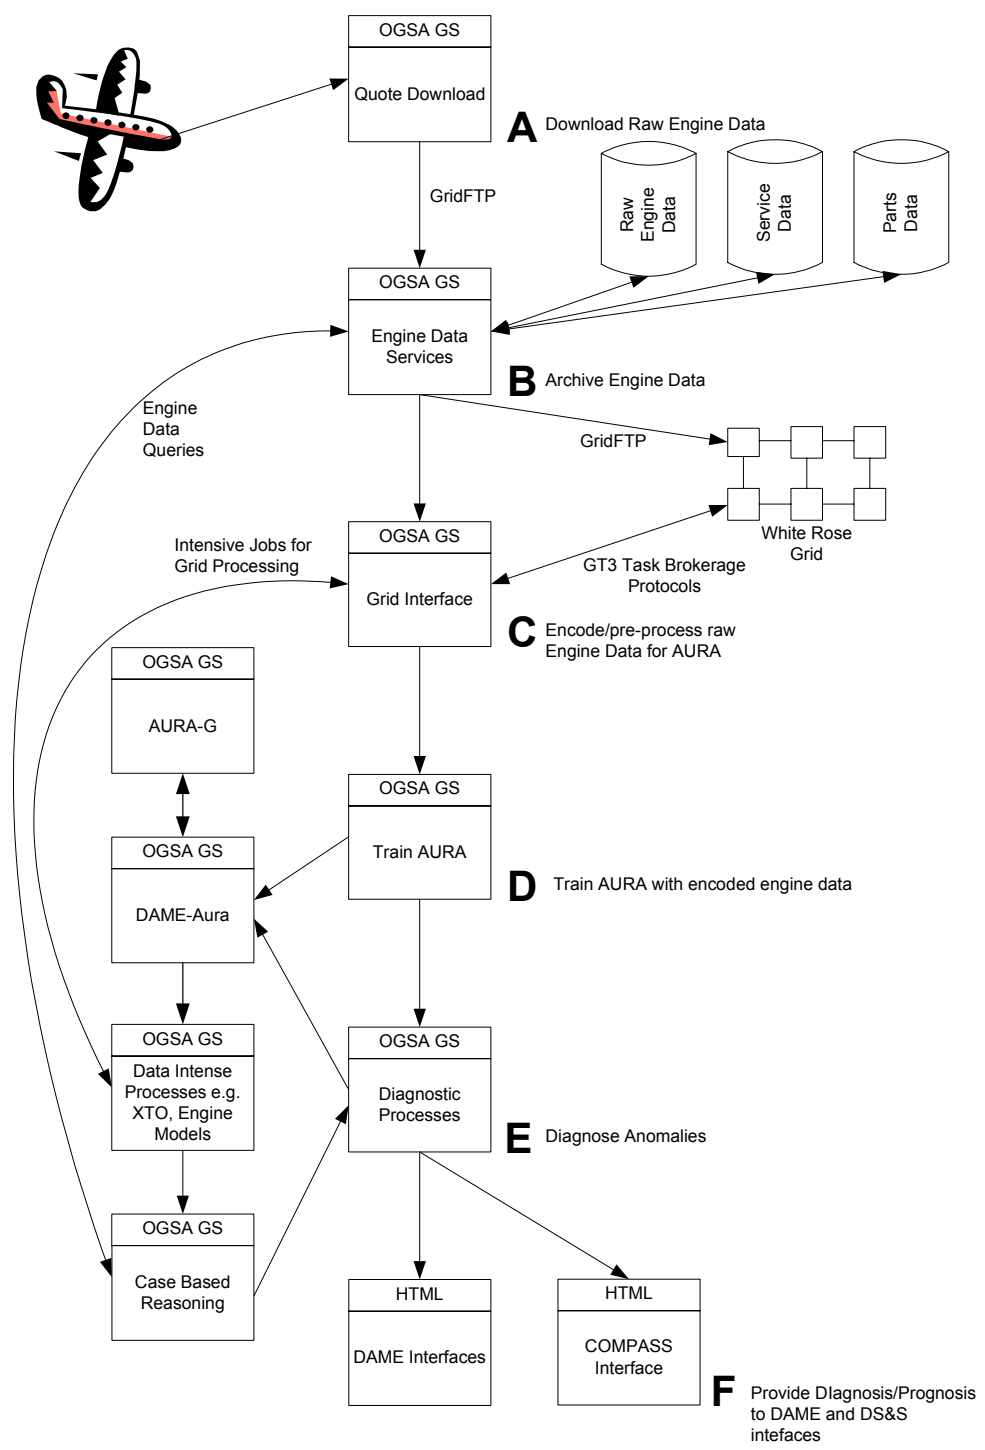


Figure 2. Functional architecture and operation of the diagnostic scenario within the DAME demonstrator.