

# Service-based Collaborative Workflow for DAME

Duncan Russell     Peter Dew     Karim Djemame  
*Informatics Institute, School of Computing*  
*University of Leeds*  
*LS2 9JT, United Kingdom*  
*{duncanr, dew, karim}@comp.leeds.ac.uk*

## Abstract

*The paper reports on research to design and implement a service-based, secure collaborative workflow system for Distributed Aircraft Maintenance Environment (DAME). This has been developed by the Universities of York, Sheffield, Oxford and Leeds, and industrial partners, Rolls-Royce and Data Systems and Solutions (DS&S). DAME is a prototype system to support aircraft engine maintenance and predictive services. It is an example of a virtual organisation with grid-based services running at the four universities sites. To meet the industrial requirements strong security is implemented to protect commercially sensitive services and data.*

*Application services have been coordinated using the DAME workflow management system to automate business processes and control collaborative access. A dynamic workflow-team policy is used to authorise user access to workflow instances and corresponding service instances.*

*The paper includes aspects of an initial evaluation with the industrial participants and illustrates the feasibility of using DAME in an industrial environment.*

## 1. Introduction

Service oriented architecture (SOA) provides the ability to build distributed services, integrate them through a workflow service that supports the business processes, and deliver it to the users via a grid portal. The services can be managed by different organisations using a virtual organisation (VO). The purpose of this paper is to describe a service-based, secure collaborative workflow services for Distributed Aircraft Maintenance Environment (DAME). This has been developed by the Universities of York, Sheffield, Oxford and Leeds, and with our industrial partners,

Rolls-Royce and its informatics partner DS&S. DAME is an e-Science test bed project [1].

The purpose of DAME is to prototype a grid-based environment to support aircraft engine maintenance and predictive services. The three main application services are: (1) analysis of engine vibration data recorded from an aircraft engine during flight to detect possible faults at Oxford, (2) pattern match across the engine fleet at York and (3) Case Based Reasoning at the Sheffield [2]. At Leeds, the services are coordinated using a workflow model and delivered to the user via a grid portal. Because the production version of DAME is to operate in an industrial environment, it must protect commercially sensitive services and data against both external attacks and internal misuse. Further, the access rights must operate across organisational bounds (e.g. VO) [3].

The DAME workflow management system has been designed to automate the business processes [4]. The workflow service coordinates the application services at the four sites. A workflow creates the context for the services; it also controls the collaborative access to the services. The workflow context (workflow instance) has an access policy, which is known as the team policy. There is one team policy per workflow instance. It is created from a template policy defined from the workflow definition. It is a dynamic version of the template policy, resolving user to role mappings, and defining access to service instances for fine-grained access control. The team policy establishes the authorisation of users' access to workflow instances and the corresponding services instances that are managed and deployed by different organisations. It is assumed that identity management (authentication) is already handled by agreement between VO members, for example using SAML [5] assertions from authentication systems such as Shibboleth [6]. One important fact is the workflow instance is trusted, such that it guarantees to carry out the team policy. This means that the access control to the service instance only needs to validate the workflow context. A dynamic team policy specifies

users, their role mapping, and the services instances the users can access. This is used in the collaborative environment to share the service instances among the workflow participants. In DAME, the modelled workflow is the analysis of engine data using a team model to control collaboration on diagnosis and access to the service instances that execute diagnostics tools.

The workflow technology BPEL [7] is rapidly dominating commercial workflow systems, commercial web service workflow systems, emerging above other workflow languages, such as XPDL, XLANG, WSFL, BPML and WSCI. However, BPEL lacks support for access control. Additionally, the DAME system uses the Globus Toolkit 3.2 OGSA model [8] for grid services. It defines the service factory and creates stateful service instances with a unique URL, which is also not supported by BPEL. Therefore, the proposed workflow and access control model is not dependent of these emerging and currently unstable technologies.

The paper is structured in the following way. In section 2 a summary of relevant research is given. Next, the DAME system is presented. In section 4, the details of the DAME collaborative workflow system are described. This has been written in a general manner, so that it can be used in other applications. Section 5 provides implementation details. Some initial feedback from the industrialists is reported and the lessons learnt are given in Section 6. Further evaluations will be undertaken in the follow-on project BROADEN (Business Resource Optimisation for After market and design On Engineering Networks).

## 2. Background

Workflow activities can be automated in computer systems. By implementing these activities as services, then a service oriented architecture can provide loose coupling between business process definition and workflow implementation. Distributed services can be replaced or modified to improve performance or quality without the need to change the business process. This allows the workflow implementation to change more often than the business process definition and for services to be outsourced if required. Outsourcing services leads to a business model that includes service suppliers and even commodity computing from compute resource suppliers. The combination of these creates the model for grid computing and VOs [3].

Methods from the web service and grid community have aimed to tackle this. Shibboleth [6] provides a user authentication system that crosses organisational boundaries and can include attribute assertions, such as role. The Globus implementations of grid middleware in Globus Toolkit 2.4 and 3.2 [9] have required user

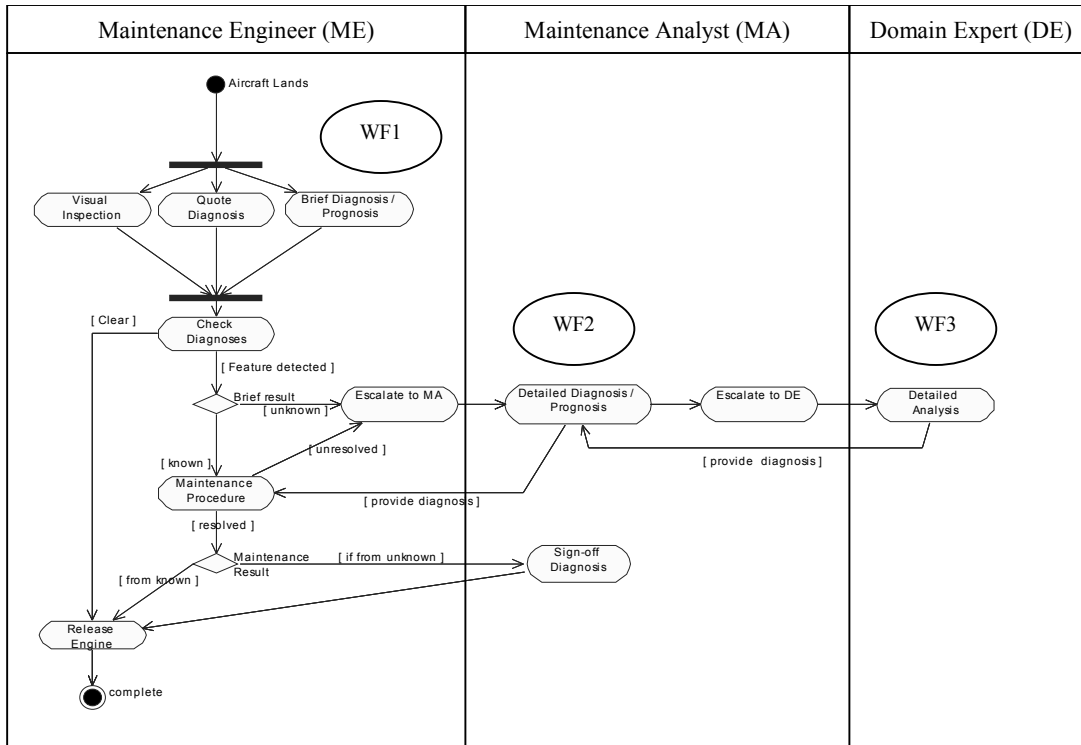
identities to be entered on each resource to enable authentication and authorisation that comes from local policies. Extensions to the grid systems attempt to combine authentication with authorisation assertions. Examples are CAS [10] and VOMS [11] which define the members of a VO and their permissions to services and resources in the system. The VO permissions are designed to be slow changing and stored in static policies, which does not allow for service instances. Policy decision engines such as Akenti [12] and PERMIS [13] can retrieve policy documents on each decision.

RBAC [14] provides an intermediary entity to separate access permissions of users to objects. The NIST standard [15] extends RBAC by using dynamic activation of user to role mapping activated on a session basis and are dependant on the task requirements. An extension to NIST could be to specialise the session as an instance of a workflow. However, other methods of access control use workflow to define the context. Task-Based Authorization Controls (TBAC) [16] shows how active access control can automate the fine-grained association of subject to object security. The TBAC model demonstrates how to apply roles (i.e. RBAC) to the dynamic policies, but does not include a workflow context.

Similarly, Team-based Access Control (TMAC) [17, 18] uses roles in teams for access control in collaborative environments. This presents a link between role-based permissions across object types, and provides fine-grained, identity based control on individual users to individual object instances. The team here is defined as a project team [17], a subset of a VO rather than a group collaboration on fulfilling a specific goal. In TMAC2004 [18], the team is instance based, but not linked to workflow. Therefore, a different approach is proposed to provide access control for teams that enact workflows.

## 3. DAME

The case study for secure collaborative workflows is DAME [1, 3]. The DAME system provides aircraft engine support, determining maintenance requirements by diagnosis and predictive services. This section presents the DAME business environment, describing the main collaborative workflow and the VO partners. This example workflow is the analysis of engine vibration data that is recorded from an aircraft engine during flight.



**Figure 1 Diagnostics business process**

### 3.1. Workflow Example

The main diagnostics business process is a result of Use Cases analysis [19] and the workflow is shown in Figure 1. The people involved in the workflow are shown as the roles that can be instantiated by relevant personnel. These roles represent different skill sets and their job is summarised as:

- **Maintenance Engineer (ME):** carries out inspection, diagnosis and maintenance of aircraft engines; Employed by the Airline and is based at the airport;
- **Maintenance Analyst (MA):** provides technical advice and coordinates analysis; Employed by the Fleet Maintenance Management company (DS&S) and is based at the diagnostics support centre where the airline's aircraft maintenance contracts are managed;
- **Domain Expert (DE):** acts as a repository of knowledge and will provide expert diagnostic advice on unidentified features; Employed by the Aircraft Engine Manufacturer (Rolls-Royce), and is based at the engine manufacturer's design centre and is an experienced aircraft engine designer.

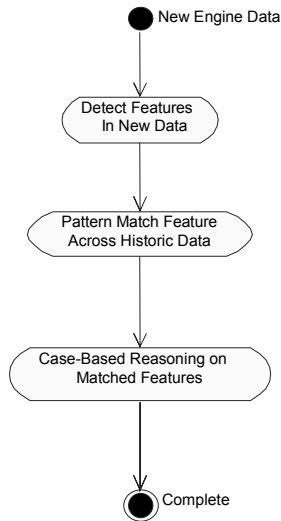
The mechanism to control team membership is provided by an escalation process. The process is the task that requests assistance with the diagnosis from more specialised personnel. This builds the team. The

ME escalates the problem to the MA when the initial automated process (WF1) cannot recognise out of parameter vibration signals. The MA investigates the problem with a range of tools (WF2) and if needed escalates the problem to a DE. The DE uses further investigation tools (WF3).

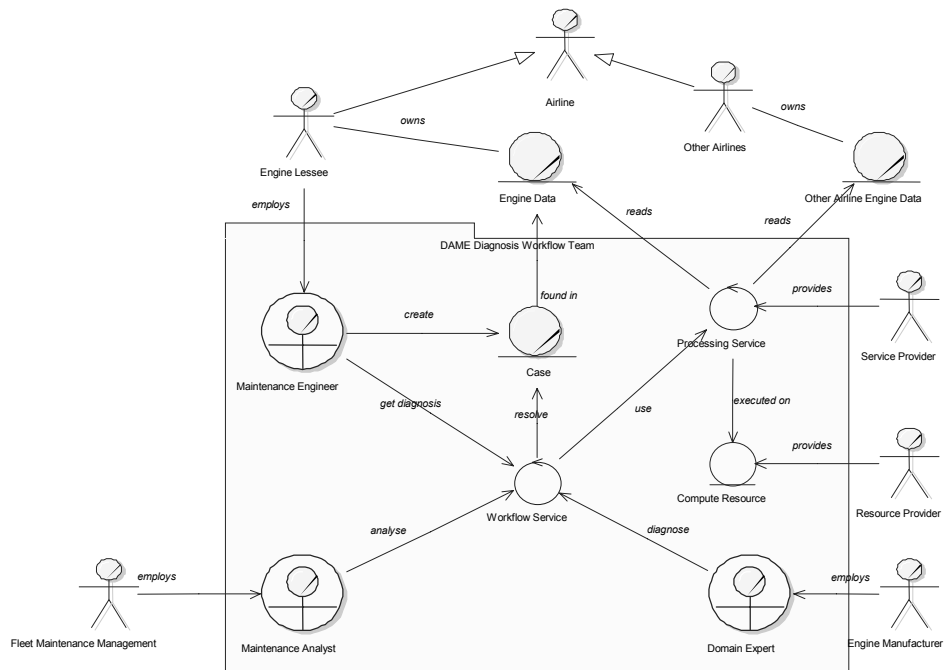
The workflow in Figure 1 is a simplified view for illustration. It assumes that each case solves the problem during this workflow. Of course, they may be occasions of longer diagnosis, where the ME is removed from the process, but the case continues to be diagnosed by the MA and DE.

### 3.2. Service based workflow

The activities shown in Figure 1 are high-level business activities. There are three high-level activities, indicated as WF1, WF2 and WF3, which contain workflows in themselves. The first of these, WF1 shown in Figure 2, is an automated workflow triggered by the arrival of new engine data from the on-wing system. It uses a chain of data processing tools to produce a most likely prognosis. These processing tools are part of on-going development in diagnosis and in order to support the inclusion of new versions and different implementations each tool is implemented as a service component, executing on grid compute resources. The workflow management system coordinates the different services in a SOA.



**Figure 2 WF1 - Brief Diagnosis / Prognosis**



**Figure 3 DAME Virtual Organisation, showing the diagnosis workflow-team**

### 3.3. Workflow across Virtual Organisations

When workflows are a collaboration of users and services from different organisations then we can say it is a Virtual Organisation (VO). The VO model is extended to include supply of compute power to the services, from grid computing [3].

The DAME workflow is executed by a VO including users from the Airline, the Fleet Maintenance Management and the Engine Manufacturer, and the Service and Resource Providers. Figure 3 shows the VO along with the user roles and their respective organisations involved in the diagnostics workflow. The Workflow Service uses Processing Services, which in turn use Compute Resources. These are supplied by Service and Resource Providers. The users, services and resources involved in a particular Case are linked by the Workflow-Team context.

A Workflow Service, executed by a workflow engine is instantiated for each Case, when it is created by the Maintenance Engineer. The Workflow Service instantiates Processing Services. It includes or precludes other users as required by the workflow case, controlling the membership of the workflow-team.

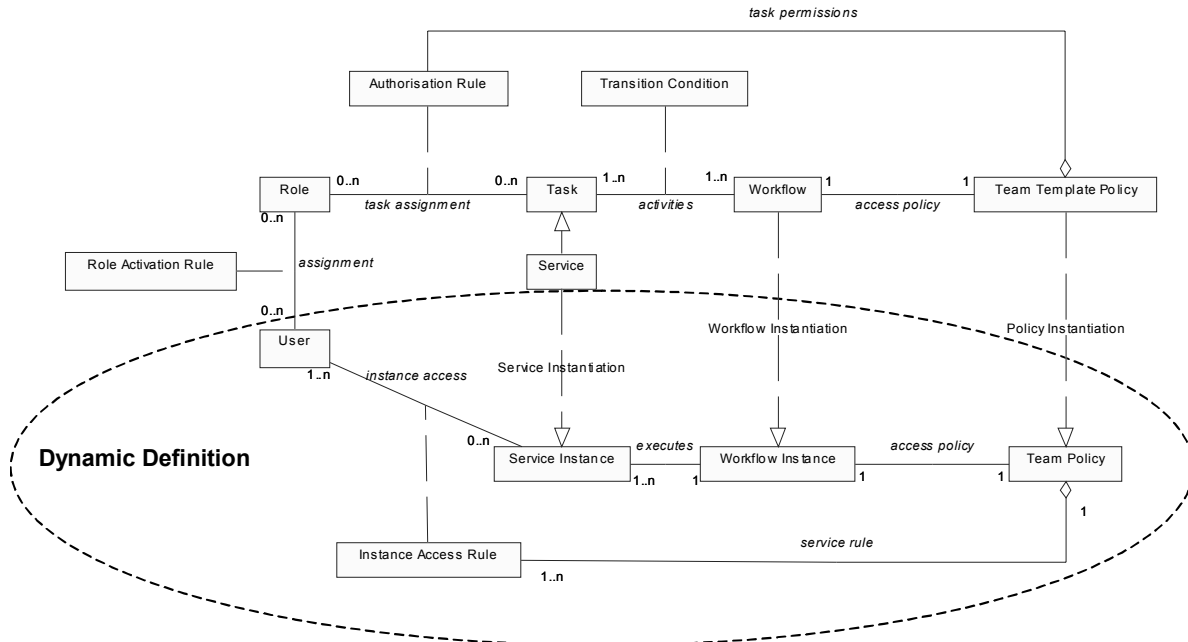
### 3.4. DAME Security

In the DAME example, there are strict requirements for security to protect commercial assets.

These assets include the engine vibration data, resultant diagnoses, processing and modelling algorithms and information about the business processes. The security study [20] identified assets at risk. This study is aimed at protecting the confidentiality of those assets that are commercially sensitive.

Security issues include:

- Requirement for strong security, due to commercially sensitive data and processing;
- In some cases, the engine data is owned by the airline operating the aircraft;
- Processing services use historic engine data and diagnosis results, so there may be restrictions on detailed access to diagnosis output, depending on role and organisation attributes of the user;
- Access policy mechanisms should not impede the workflow time frame when authorising security assertions.



**Figure 4 Workflow-Team policy architecture**

## 4. Workflow-Team Policy Architecture

The DAME scenario shows a VO with security requirements. Grid security mechanisms, such as CAS [10] & VOMS [11], can provide control for predetermined teams using static services and data in the VO. However, a workflow has a subset of users from the VO and includes service instances created for the workflow, which are not detailed in the VO.

The workflow creates the context for the use of services; it also controls the collaborative access to the services. Therefore, the workflow context (workflow instance) has an access policy, which is called the team policy. There is one team policy per workflow instance. It is created from a template policy defined from the workflow definition and is a dynamic version of the template policy, resolving user to role mappings and defining access to service instances for fine-grained access control.

The Team policy is concerned with the authorisation of users' access to workflow instances and the corresponding services instances that are managed and deployed by different organisations. One important fact is that the Workflow Instance is trusted, such that it guarantees to carry out team policy. Therefore, service instance access control only needs to validate the workflow context.

### 4.1. Workflow-Team Architecture

Figure 4 shows the Workflow definition and the Team Template Policy. These are static definitions,

meaning they change according to the business needs. The Workflow is made of Tasks that are linked to Transition Conditions. The other static types defined in Figure 4 are:

- Role activation rule, which when conditions are true permits a User to assume the identity and permissions of a Role
- Task: an activity in a workflow.
- Authorisation Rule, provides conditions for the link between Roles and Tasks. Such as, role must belong to organisation.
- Transition Condition, defines transition between Tasks in a Workflow. Transition types include Sequence, OR-split, AND-join [4]
- Team Template Policy: This static description is linked to a Workflow description. It constrains how Users enact a Workflow by containing the Authorisation Rules for Roles to enact Tasks in the Workflow.

When the workflow is enacted, the definition of the workflow is instantiated, becoming a Workflow Instance. The Team Template Policy becomes a Team Policy linked to the Workflow Instance. Tasks from the workflow definition that enact services become Service Instances, and Roles that were permitted to access those services via the Authorisation Rule are entered in the Team Policy as Users, by their corresponding Instance Access Rule. In Figure 4, the Dynamic Definition is the architecture used in the enactment of the workflow instance. Types are defined as:

- Service: A grid service (factory) that implements a Task.

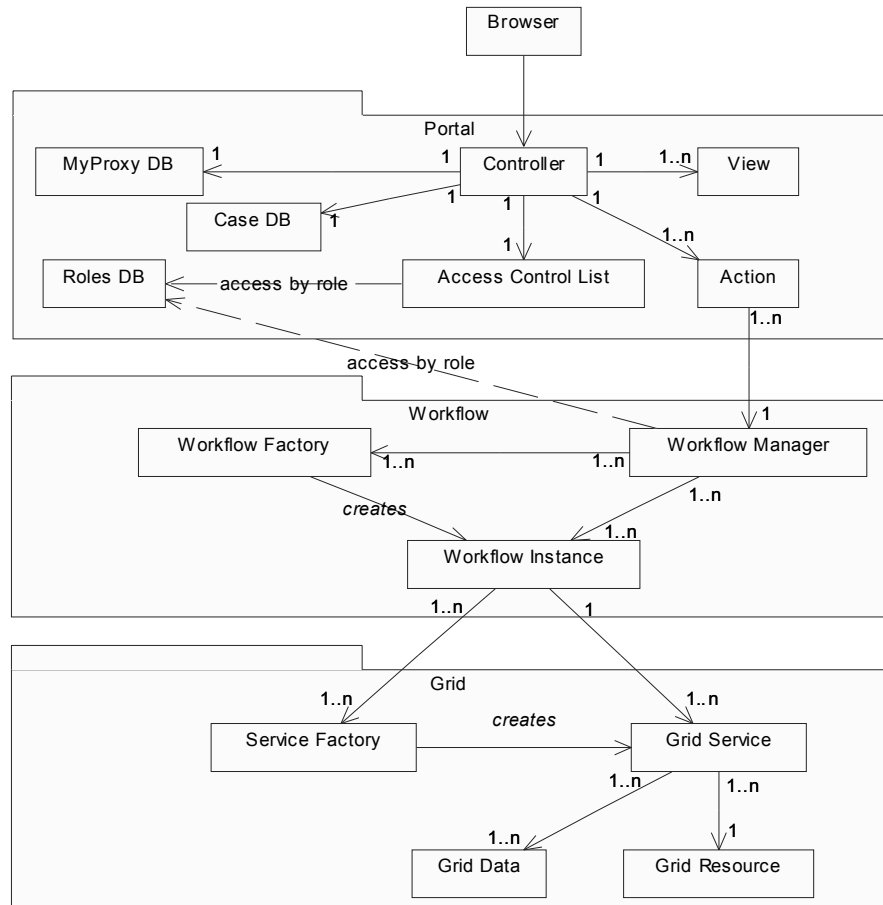
- Workflow Instance: The runtime state associated with the Workflow (definition).
- Service Instance: A stateful grid service instance created by the service factory.
- Instance Access Rule: Derived from both the Role Activation Rule and the Authorisation Rule, this links the User to the Service Instance, stating the Role of the User and retaining constraints from the rules.
- Team Policy: the policy instance created from the Team Template Policy. This is the dynamic constraint on the executing Workflow Instance, providing fine-grained control as to Users permissions for actions on Service Instances.

This dynamic policy architecture can be used across a VO, so long as the workflow engine is trusted by VO partners. However, the architecture can be easily implemented to provide team based access control within a single organisation, and prevent unauthorised access within an organisation.

## 5. DAME Collaborative Workflow Implementation

The DAME workflows are managed in a workflow system connected to the DAME portal, shown in Figure 5. The workflows execute secure grid services that are factory based, implemented using GT3.2 along with GT2.4 across the White Rose Grid [21]. The portal implementation uses the Apache Struts framework to provide some important features for the DAME collaborations:

- Secure Login using sessions over https. Users are identified by username and password, authenticated against their X.509 certificate from myProxy.
- Role Based Access. The user is mapped to a role, from Roles DB, which is used by Struts to restrict the actions and views.



**Figure 5 Implementation Architecture**

- Business Process. The Struts framework allows actions and views configured by user and role.

Each portal session connects to a workflow manager by role. The role restricts access to the workflow types. The main diagnostics business process starts workflows automatically on a user's behalf, assigned to the relevant ME. The ME can login to view their active workflows and annotate the case if escalation is required.

Each active process is stored in the case, which holds the workflow instance and user instance identities. The escalation process is controlled by team definitions; these specify the next user in an escalation chain (i.e. from ME to MA to DE). These team definitions provide a template, and the case record the instances involved (the active team), such that when a user leaves the process, the next correct user receives the case.

The current DAME implementation uses a single corporate certificate to access services and resources. This certificate provides unrestricted access to the DAME services and service instances. Access control

takes place in higher-level components providing workflow services and portal login.

## 6. Feedback

The collaborative workflow-team model presented has been based on DAME documents and discussions between project partners [19, 22]. Feedback to this model has been obtained from presentations both internal and external to DAME [23, 24] and analysis from the DAME [20]. Furthermore, a structured interview technique has been used with representatives from DS&S and Rolls-Royce [25, 26].

The interview results show that the industrial partners agreed the DAME demonstrator reflects the secure role-based collaborations and proves the feasibility of integrating grid-services from different organisations. The demonstrator also found that role alone is not enough to restrict access to the workflow. A user's organisation would need to be used to ensure Airlines are restricted to viewing their own data. This was reinforced in the interviews, and Dependability and Security study, but had not been identified in the original use case requirements.

Another point to arise from the evaluation was that this system provides the opportunity to separate the business critical workflow definition from the service implementation using SOA and grid-services. However, the industrialists stated that the demonstrator does not support collaborations that may require longer-term diagnostics or collaborations between more than one user in the same role.

The portal and workflow manager have been demonstrated with user authentication using Single Sign-On (SSO) across all organisations, which is required for integration of commercial services. In a deployed system, the identity management would have to be standard across all DAME partners, or use an underlying assertion system like Shibboleth. DS&S have expressed that user-role mappings rules could include attributes of suitability, such as a user's qualifications. Access control for services could judge suitability by requiring it to correspond to certain standards, such as z-mod for data exchange (a Rolls-Royce data format) and the MIMOSA information standards [27].

A point noted from Rolls-Royce: *"The problem of the grid is that it provides rapid access to...potentially confidential data. What will probably happen is that there will be a piecemeal approach. In that, access to individual resources, data and information will be made available under specific circumstances ... on a one by one basis."* Despite this, it is important that access control mechanisms do not prevent collaborations that would benefit the business.

Airlines have expressed interest in the DAME capabilities and wish to develop analysis tools that integrate into DAME workflows. They also wish to integrate access to DAME in their own corporate. The integration of services reinforces the VO model for DAME, but also shows that it is important to use access control to support the export of workflows and services into VOs outside DAME.

## 7. Implications for e-Science

UK e-Science projects, such as myGrid and Triana [28], show that service and workflow definitions can be shared between users. DAME advances this by showing how secure access and collaborative workflow can be implemented within the DAME VO. This enables us to securely integrate the application services in a collaborative workflow, which can be re-used in other projects. Further, the team policy can integrate other services and utility compute resources. By providing a team access control layer on the workflow, service access control is simplified by basing authorisation on the identity of VO. This, however, is only one method to address where authorisation responsibility should be located in a collaborative problem-solving system.

The workflow-team policy focuses on authorisation for grid-based collaborations, which compliments the identification of security protocols by the UK e-Science Security Task Force.

## 8. Conclusions

The workflow management in DAME addresses the need to control access to commercially sensitive services and data in an aircraft engine diagnostics environment. The dynamic workflow-team policy architecture enables a secure collaboration between users, services and compute resources in different organisations. This has been successfully demonstrated to the industrial partners, who have provided feedback for the ongoing work in the follow-on project BROADEN (Business Resource Optimisation for After market and design On Engineering Networks).

## 9. Acknowledgements

This research is funded by the Engineering and Physical Sciences Research Council (EPSRC), e-Science Programme, Contract No. GR/R67668/01. Thanks go to Georges Honore and Martyn Fletcher and for the time provided by Graham Hesketh at Rolls-Royce and Charlie Dibsedale at DS&S.

## 10. References

- [1] Austin, J. and et al., *Distributed Aircraft Maintenance Environment DAME: A GRID e-Science Full Proposal*, DAME Project, 28/06/2001.
- [2] Russell, D., *DAME: A Distributed Diagnostics Environment for Maintenance*, NeSC, 2003. <http://www.nesc.ac.uk/talks/303/DAMEWorkflowNESC41203.pdf>.
- [3] Foster, I. and C. Kesselman, *The grid 2: blueprint for a new computing infrastructure*. 2nd ed. San Francisco, Calif.: Morgan Kaufmann. 2004.
- [4] Aalst, W.M.P.v.d. and K.v. Hee, *Workflow Management: Models, Methods, and Systems (Cooperative Information Systems S.)*: The MIT Press. 2004, 384.
- [5] OASIS, *OASIS Security Services (SAML) TC*, OASIS, 2004. [http://www.oasis-open.org/committees/tc\\_home.php?wg\\_abbrev=security](http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=security).
- [6] Cantor, S., *Shibboleth Architecture*, 2004. <http://shibboleth.internet2.edu/docs/draft-mace-shibboleth-arch-protocols-05.pdf>.
- [7] Andrews, T., et al., *BPEL4WS v1.1*, BEA, IBM, Microsoft, SAP AG and Siebel Systems, 2003. <http://www-106.ibm.com/developerworks/library/ws-bpel/>.
- [8] Foster, I., et al., *The Physiology of the Grid*, 2002. <http://www.globus.org/research/papers/ogsa.pdf>.
- [9] *Globus*, The Globus Project, 2003. <http://www.globus.org>.
- [10] Foster, I., et al. The Community Authorization Service: Status and future. in *CHEP 03 2003*. La Jolla, California.
- [11] Alfieri, R., et al., *VOMS, an Authorization System for Virtual Organizations*, DataGrid Project, 2003. <http://grid-auth.infn.it/docs/VOMS-Santiago.pdf>.
- [12] Thompson, M., A. Essiari, and S. Mudumbai, Certificate-based Authorization Policy in a PKI Environment. *ACM Transactions on Information and System Security*, 2003. **6**(4): p. 566 - 588.
- [13] Chadwick, D.W. and O. Otenko. The PERMIS X.509 Role Based Privilege Management Infrastructure. in *Proc.7th ACM symp. on Access control models and technologies 2002*. Monterey, California: ACM Press.
- [14] Ferraiolo, D.F., D.R. Kuhn, and R. Chandramouli, *Role-Based Access Control*. Norwood, MA: Artech House. 2003.
- [15] Ferraiolo, D.F., et al., Proposed NIST Standard for Role-Based Access Control. *ACM Transactions on Information and System Security*, 2001. **4**(3): p. 224–274.
- [16] Thomas, R.K. and R.S. Sandhu, *Task-Based Authorization Controls (TBAC): A Family of Models for Active and Enterprise-Oriented Authorization Management*, in *Proc. 11th Int. Conf. on Database Security XI: Status and Prospects*. 1998, Chapman & Hall, Ltd. p. 166-181.
- [17] Thomas, R.K. Team-based access control (TMAC): a primitive for applying role-based access controls in collaborative environments. in *Proc. 2nd ACM workshop on Role-based access control 1997*. Fairfax, Virginia: ACM Press.
- [18] Alotaiby, F.T. and J.X. Chen. A model for team-based access control (TMAC 2004). in *Proc. ITCC 2004. Int. Conf. on IT: Coding and Computing, 2004*. 2004: IEEE.
- [19] Fletcher, M., *DAME Requirements: Use Cases*, DAME Technical Report, DAME/York/TR/02.001, DAME Project,
- [20] Fletcher, M., *DAME Dependability and Security Study: Final Report*, version 1, Technical Report, DAME/York/TR/04.007, DAME Project, 27/10/2004.
- [21] Dew, P.M., et al. The White Rose Grid: practice and experience. in *UK eScience - All Hands Meeting 2003*. Nottingham, UK.
- [22] Fletcher, M. and DAME Architecture Working Group, *DAME Service Definitions and Descriptions*, 1k, DAME Technical Report, DAME/York/TR/02.005, DAME Project, 01/05/2003.
- [23] Russell, D., P. Dew, and K. Djemame. Self Securing Dynamic Virtual Organisations. in *Workshop on Grid Security Practice and Experience*. H. Chivers and A. Martin, (eds). 2004. Oxford, UK: University of York.
- [24] Russell, D., P.M. Dew, and K. Djemame. Access Control for Dynamic Virtual Organisations. in *UK e-Science All Hands Meeting 2004*. S.J. Cox, (ed) 2004. Nottingham, UK: EPSRC.
- [25] Russell, D., *DS&S Interview Results 04/11/2004*, DAME, 2004.
- [26] Russell, D., *Rolls-Royce Interview Results 12/11/2004*, DAME, 2004.
- [27] *MIMOSA - Machinery Information Management Open Systems Alliance*, MIMOSA, 2004. <http://www.mimosa.org/>.
- [28] *UK e-Science Programme*, Research Councils UK, 2005. <http://www.rcuk.ac.uk/escience/>.