

Self Protecting Data for De-perimeterised Information Sharing

Pete Burnap and Jeremy Hilton

Cardiff School of Computer Science, Cardiff University
p.burnap@cs.cardiff.ac.uk, jeremy.hilton@cs.cardiff.ac.uk

Abstract - The emergence of high-speed networks, Grid Computing, Service-Oriented Architectures, and an ever increasing ambient connection to mobile Internet has enabled an underpinning infrastructure for the development of dynamically formed, collaborative working groups known as Virtual Organisations (VOs). VOs provide strong motivation for investigation into the infrastructure, and in particular the security necessary to protect the information and resources shared within a VO, both while resident on local machines and when allowed to move beyond the secure boundary of a local organisational network perimeter and into the realm of the distributed VO. Traditional access control systems are perimeter-centric, meaning they apply the controls to both internal and external requests for access to information within or at the perimeter of their information system. This paper presents the initial results of the JISC funded SPIDER project, being led by Cardiff University. Through case based example, the research investigates the limitations to granularity and persistent control over information when using the perimeter-centric approach in a collaborative working environment.

1. Introduction

Collaborative working between multiple organisations will always require some level of information sharing and exchange. A significant amount of information belonging to an organisation will have an associated value for which appropriate protection mechanisms must be put in place in order to prevent the exposure or loss of that information. The emergence of Grid computing and Service-Oriented Architectures have led to the increasing adoption of dynamically formed, collaborative working groups known as Virtual Organisations (VOs). These Virtual Organisations (VOs), as defined in [1] and [2], provide strong motivation for investigation into the infrastructure, and

in particular the security necessary to protect the information and resources shared within a VO, both while resident on local machines and when allowed to move beyond the secure boundary of a local organisational network perimeter and into the realm of the distributed VO.

Much of the previous research in the area of access control approaches to shared information to date such as VOMS [3], PERMIS [4], ShARPE [5] and iRODS [6] has focused on the protection of information resources as an entity within the secure network boundary, for example, an entire classified document; access is either granted or denied using system-level access controls, or Digital Rights Management (DRM) [7] techniques on the entity. This approach often has two major drawbacks:

- It hinders information sharing to some extent due to its limited granularity. That is, information sharing, and as a result collaborative working, is not being allowed to reach its maximum potential because large amounts of information cannot be shared due to small amounts of higher-level sensitive content within the resource raising the overall classification of the resource.
- With current DRM and system-level controls that can control access to information to some extent *after* the information has been allowed to move beyond an organisation's access control perimeter, the access control policy is permanent and cannot be modified by the owner of the information. However, there may be a change in the controls required to protect a resource. For example, the VO working group may disperse or the VO community may be changed, thus wishing to deny access to information previously shared.

The JISC funded SPIDER project at Cardiff University aims to address both drawbacks. Firstly by investigating and modifying approaches to access control which remove the fixed boundaries around the whole information resource, and place boundaries just around the sensitive content within the resource. Thereby putting part of the access control policy within the information itself, and allowing the access

restrictions to apply to not only the entire resource but also the content within. This has the potential to allow the sensitive content to be strictly controlled, while the rest of the information in the resource can be made publicly available. Different views of a resource can be created for varying levels of access control.

Secondly by adapting the access control techniques to include modifiable access control policy, even on resources that have been stored on media outside the control of the system access control perimeter. The proposed work will review and build on current and emerging standards/approaches to Information Security that define policy and place access control restriction criteria with the information resource itself, instead of the common approach which relies on centrally controlled access to information contained within a finite perimeter (e.g. a company network).

Collaborative working arrangements can have a limited and dynamic lifespan. Following the end of collaboration or change in organisational structure, there may be a requirement to discontinue the sharing of information. In fact, sometimes it may be necessary to terminate the sharing arrangement immediately, should a loss of trust or irreparable loss of working relationship occur. In this case, the access control policy used to restrict access and use of the information would ideally be modifiable in order to restrict or remove previously assigned access and usage rights.

Using a risk assessment of information sharing in the UK Probation Service as a case study, this paper presents a model intended to reconfigure and extend current access control techniques and technologies: to allow the ability to enforce restrictions outside of an organisational information system boundary; to drill down into the data level of an information resource in order to apply varying levels of protection to different sections of content within a resource; and finally to enable the modification of access control policy on previously shared, distributed information. The following sections detail current approaches to controlling access and use of information and suggest how these approaches could be modified to better suit collaborative working across distributed, autonomously controlled information systems. At present the work is focussed on the temporal, task-specific structure of VO activity so it should be noted that the information being shared is of a particularly sensitive nature and that there may be additional contractual agreements that need to be enforced within the solution.

2. Background

The development of a collaborative electronic working infrastructure provides the potential for a

more effective and productive working relationship between inter-organisational, cross discipline organisations working on distributed information systems. The United Kingdom criminal justice system is a prime example for such activity when considering the interdisciplinary organisations working towards the same goal. As part of the SPIDER project, the project team have conducted a risk assessment of the information sharing system within the prosecution service for youth offenders. Information is shared between the Courts, the Police, Youth Offender Services and the Parents of the Youth among others. There is a requirement to share information unique to each case such as the details of the offence, the charges being brought, the resulting criminal record, and the identity details of the youth in question. While the information is required to be shared between organisations in support of the collaborative system, it is also clear from the risk assessment that information is at risk of disclosure, loss, corruption and interception during operational activities. It is incredibly important that sensitive content within the shared information is classified according to the requirements for its protection, ensuring that only the relevant people gain access to sensitive information. It is equally as relevant that each organisation maintains control of the information that they are responsible for, in order to reduce the risk to information shared in such a way. This presents some key requirements:

- i. An information classification scheme that is able to accurately define the requirements necessary to protect shared information.
- ii. An access control model that is enforceable by all collaborating parties, taking into account the possibility that information may not remain local to their own information systems and the controls may need to be applied to information that has been distributed to collaborating partner's information systems.
- iii. A fine-grained approach to applying controls within an access control model. Several organisations may be responsible for different sections of content in a collaboratively developed resource. If they contribute some content to a resource, a criminal case report for a youth offender for example, they should be able to define their own protection requirements for that content. Providing controls to the entire resource with a single classification may not be acceptable as each organisation may have a different view on the protection requirements for their information and the content within the resource may vary in classification. Certain sections of content may be classified higher than others and require a greater level restriction. In some cases information may also be protected by

the Data Protection Act, in which case the contributing organisation is legally obliged to protect that information. The definition and application of controls should be able to reflect that.

- iv. As collaborative relationships develop and disband, access control privileges for collaborating partners may need to be modified to add, modify or revoke access privileges. The privileges defined in an access control policy for a resource should be modifiable and should take effect immediately if a change in privileges occurs. This applies to distributed information and information shared in resources held by multiple collaborating parties.

These requirements provide the motivation for the following adaptation of current access control models and techniques to allow more flexible, persistent and accurate protection for information shared in collaborative working environments.

3. Existing Access Control Techniques

Traditional access control systems such as PERMIS and VOMS are perimeter-centric, meaning they apply the controls to both internal and external requests for access to information within or at the perimeter of their information system. There are several issues that arise from the application of perimeter-centric access control to distributed, collaborative information sharing within VOs. Primarily, the use of perimeter-centric security limits the persistent control available after information has been distributed between collaborators. Once the information has been shared, i.e. copied, transferred and stored on another organisation's information system, it can no longer be protected by the same mechanisms as it is no longer stored within the perimeter. One approach to solving this issue would be to provide an access control technique that could continue to be applied to information after it had been copied, transferred and stored on another organisations' information system. Digital Rights Management (DRM) is an emerging technology that aims to provide the remote enforcement of access control policy through the use of proprietary software that controls the access and usage of distributed resources. While this goes some way to a de-perimeterised security model, the controls defined are static and can only currently be applied at the 'entire-resource' level. This highlights the problems of a lack of persistent control and limited granularity. Information resources are often a collection of related pieces of information with varying levels of protection requirements. Certain parts of the document may be restricted to use within the organisation, while other parts may be non-sensitive content that could be shared

with other organisations. It may not be trivial to remove the restricted content before sharing as it may be mixed into paragraphs with the non-sensitive content. As such, the restricted content within a document often means the entire document is not shared, limiting the effectiveness, dynamism, and potential of collaborative working.

There have been efforts in recent years to break down the content of information resources and apply different controls to content with specific protection requirements [8][9]. These approaches rely on the resource being structured so that the content can be fragmented and tagged according to their protection requirements, and the application of encryption techniques and fine-grained key management to provide access control for the resource. Effectively, fragmented content with specific access control requirements are encrypted with a different key to other sections of content within the resource, meaning that keys can be assigned to different users so that they can only decrypt the content that they have privileges to access. This approach has some major drawbacks with a view to maintaining persistent, modifiable access controls for the information. Key management is an incredibly complex task. Issuing keys to a static set of users is hard to manage; issuing them to a dynamic set of users in a collaborative working group is even harder – particularly when the key issuers will be the collaborators themselves and not a single authority. From a scalability viewpoint, the more collaborators you have contributing to distributed information resources between a dynamic set of users, key management becomes incrementally more difficult. Furthermore, once the keys are issued, there is no way of revoking them. Once a user has an information resource in their possession and is able to access it using the keys issued to them, the access control policy cannot be modified and enforced. The same drawback occurs with DRM. This means previously assigned access privileges cannot be modified at the end of a collaboration agreement, something that could be very advantageous show due diligence that all electronic access has been revoked. If this can be achieved, it becomes more likely that future use of the information can be proven to show intent of use outside of the collaborative agreement.

An approach to fragmenting and tagging information content according to its access and usage restrictions is also published by the Creative Commons community [10]. Rather than a technical set of access controls, Creative Commons identifies a set of licensing properties that inform the reader of information resources of their obligations and limitations when sharing and reusing the information contained within the resource, through the use of the

Creative Commons Rights Expression Language (ccREL) to express machine readable copyright licensing terms and related information. The approach follows a similar line of investigation to [8] and [9] by tagging sections of structured content with labels that represent different levels of access and usage restriction. However, within the information resource there is an embedded URL link to the full licence for the content. The document at this location details how the labels should be interpreted by the reader. This is a live URL and, of course, can be updated and modified at any time by creator of the licence. In addition to this, a completely new licence could be created for the resource, in which case the existing link will forward the reader to the latest version. The access control model suggested in this paper considers the possibility of modifying the existing models of fragmenting structured content based on access control requirements, together with a live URL link to an access control policy that can be enforced through a DRM style application at each client location in a collaborative working environment.

4. A De-Perimeterised Approach to Access Control

The approach developed as part of the SPIDER project aims to provide persistent and accurate access control for information, specifically text documents, shared in collaborative working environments.

The proposed solution comprises three components: an information classification scheme that represents the protection requirements available for information shared in collaborative working environments; a fine-grained approach to access control policy definition that allows information resource owners to fragment and classify their information resources, map the classifications onto the information content, and provide a live URL link to the access control policy for the resource; and a proprietary Open Source Java-based tool that can be installed on stand-alone machines or invoked via Web Start [11] both to define and enforce the access control policy.

The resources being used to demonstrate the SPIDER application are proprietary Microsoft Office Word 2003 and 2007 text documents that have been translated through a set of filters to become plain-text, structured XML documents. This process allows the documents to be manipulated using standard XML functionality and can be automated either locally or through a Web Service invocation of the Open Source Docvert Tool [12].

4.1 Information Classification Scheme

The purpose of the information classification scheme is to give information resource owners a set of protection options that can be used to represent the control requirements for their information. SPIDER has produced a set of initial controls which aim to represent controls applicable to information shared in collaborative working environments. These are:

- Community Access
- Restricted Access
- Personal Information
- Organisation Only

These labels allow an access control policy to be derived which can define user access privileges by mapping their identity to Boolean values for each label. The SPIDER application allows sections of content to be highlighted and by clicking an icon that represents the classification label; the user can select the controls applicable to that section of content. The application then transparently embeds the label into the body of the information content, nesting the content within the selected label. An example of the structure of a resource relating to the prosecution case against a youth offender is illustrated in Figure 1.




Figure 1 – Document Structure and Labelling

The vision for the SPIDER application is to act as a plug-in to existing access control techniques and be used across multiple information domains. As such, the Information Classification Scheme can be modified and contain different labels to suit the information protection requirements of information from specific domains such as healthcare, legal and academic research. The method provided within the application for selecting classification labels, mapping them to the information content, and defining access control policy

for the labels is flexible enough to support other Information Classification Schemes.

4.2 Access Control Policy Definition

At present, the access control policy is a centralised database containing usernames and a set of classification labels to which that user has access. This allows the owner of an information resource to maintain the list of users in the collaborative working group and modify the access controls as the working group develops. For example, initially they may allow a collaborator “Community Access” but remove the access after their collaboration has come to an end. Another user may have access to “Personal Information” for a limited period of time, and then have that access removed once they have completed the task for which access was required. This is based on the traditional approach to access control. The SPIDER model extends this control to information that has been distributed and shared with other collaborators making the control over information much more persistent and accurate from the owner’s perspective.

4.3 Policy Enforcement

Once the policy has been defined for the resource, it can be distributed. However, the policy embedded within the body of the resource can only be applied properly if the content remains confidential until the policy is applied. Previous work in the area has used encryption techniques to protect information until the relevant encryption keys are distributed to users and can be used to decrypt the content. SPIDER encrypts the entire resource illustrated in Figure 1 (with the exception of a URL that links to the centralised access control policy and a unique resource identifier) with the same key, reducing the overheads of key management. This renders the resource unreadable until the policy can be applied. The key is centrally stored in a database, along with a unique identifier for the resource.

Upon requests for access, the SPIDER application extracts the URL link from the resource and sends an access request to the URL across a secure connection, which is actually a Web Service resident on the server of the information owner. The request includes the unique identifier for the resource and the identity credentials of the user. The identity details that are sent are previously configured in the client application. The Web Service forwards a query to the access control and key storage databases for the selected resource and returns a security label containing the classification labels to which the user has been granted access, along

with a decryption key for the resource. The client side SPIDER application can then decrypt the resource in memory, parse the resources for the classifications labels that match the security label returned for the user, and generate a dynamic subset of the original resource in unencrypted form for the user to access. The entire policy enforcement process is illustrated in Figure 2.

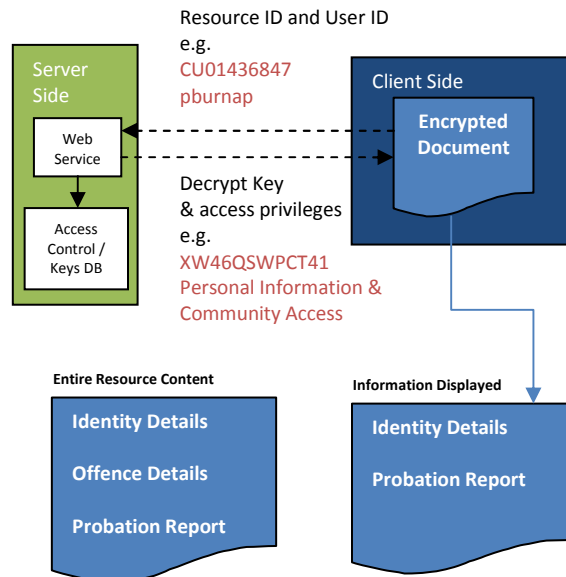


Figure 2 – Policy Enforcement Mechanism

The information is displayed to the user by the SPIDER application so that the user can then edit and add to the content. If they so wish, they can add their own classification labels and another URL link will be added to the resource so that next time a request for access is made, the client side application will query both access control policies before generating the unencrypted document. Once the SPIDER application is closed or the file is saved, the information is returned to encrypted form until the next access request.

The SPIDER text editor is a basic Java editor. The persistent control over information relies on the editing of text being carried out within the proprietary editor in order to enforce the access control policy and restrict access based on the set of embedded information classification labels. This is very much along the same lines as DRM approaches. The SPIDER project assumes that to maintain control in distributed environments there must be an agent that acts on the behalf of the local access control policy to provide control over information. This follows the development of other DRM tools that control access to licensed media such as music and video. The advantage of the SPIDER application is that it is platform independent,

being developed in Java, and it can be invoked over the Web through the Web Start function of Java, meaning that a user does not necessarily need to have any software installed on their local machines to use the editor.

5. Conclusions & Further Work

The SPIDER project is in early stage development. At present, this paper presents the concept of an access control policy that is managed centrally and can be enforced on machines that are under autonomous control and outside of the secure network perimeter of the organisation that owns the information. Also, that information can be fragmented and protected at different levels depending on its protection requirement, with each section of content having modifiable access control policy allowing much more fine-grained, persistent control over information. The rationale comes from the traditional method of labelling paper based information resources with the level of restriction assigned to them. The underlying concepts are developed with some additional functionality still to be researched and developed during the project lifetime.

Identity management is an area that is still under investigation. It is not envisaged that a simple username identity credential will suffice. X.509 certificates are one option under investigation. Attribute certificates allow the potential to incorporate user identity, organisation, VO and location attributes and can be issued by a trusted authority to ensure a greater level of trust in the security of the SPIDER application.

Access control policy could potentially be represented in an actual language. XACML or other standards could provide a more standard, flexible method for representing policy.

Trusted computing is another area of interest. The decryption of the resources in memory is open to attack from memory hacks or leaks, and the addition of a trusted computing module to ensure confidentiality during decryption/encryption is under investigation.

The text editor is quite basic at present with simple text editing capability. It is envisaged that a plug-in will be developed for Open Office in the future so that users gain full text editing functionality. The vision is to share files in proprietary format, translate them to open XML based format if necessary, apply the controls through SPIDER and open the document for editing. All actions being performed transparently to the user as if they were opening a normal text file.

SPIDER research to date has focussed on the restriction of collaboratively developed text documents. However, because of the structured nature

of the information managed by the SPIDER application, the approach could be extended to the control of other structured information such as databases and Web Services. Web Services in particular could be a very interesting target for future development. Web Service description documents (WSDLs) contain information relating the location, all available operations and input/output parameters for Web Services. The SPIDER application could be used to add labels to the available operations in order to classify and define access controls to those operations. For example, a particular Web Service may offer query, modify and delete operations for database access. The ability to persistently control and modify privileges that define who can query, who can modify and who can delete from the database may be particularly advantageous in a collaborative working environment.

6. References

- [1] Foster, I., Kesselman, C., *The Grid 2: blueprint for a new computing infrastructure*, 2nd Ed, San Francisco, Calif.: Morgan Kaufman, 2003
- [2] Foster, I., Kesselman, C. Tuecke, S., *The Anatomy of the Grid, Enabling Scalable Virtual Organisations*. Intl J. Supercomputer Applications, 15(3), 2001
- [3] Alfieri, R., et al. VOMS, an Authorization System for Virtual Organizations . In *Proceedings of the 1st European Across Grids Conference*, Santiago de Compostela 2003
- [4] Chadwick, D. W. and Otenko, A. 2003. The PERMIS X.509 role based privilege management infrastructure. *Future Gener. Comput. Syst.* 19, 2 (Feb. 2003), 277-289.
- [5] ShARPE Project Web Site : <http://www.mams.org.au/confluence/display/SHA/ShARPE;jsessionid=EC9B1D9DE1EFF3E275979B6499BE867B>
- [6] Introduction to iRODS : https://www.irods.org/index.php/Introduction_to_iRODS
- [7] Liu, Q., Safavi-Naini, R., and Sheppard, N. P. 2003. Digital rights management for content distribution. In *Proceedings of the Australasian information Security Workshop Conference on ACSW Frontiers 2003 - Volume 21* (Adelaide, Australia).
- [8] Damiani, E., De Capitani di Vimercati, S., Paraboschi, S., and Samarati, P. 2002. A fine-grained access control system for XML documents. *ACM Trans. Inf. Syst. Secur.* 5, 2 (May. 2002), 169-202. [9] Bertino, E., Castano, S., and Ferrari, E. 2001. On specifying security policies for web documents with an XML-based language. In *Proceedings of the Sixth ACM Symposium on Access Control Models and Technologies* (Chantilly, Virginia, United States).
- [10] Abelson, H et al. ccREL: The Creative Commons Rights Expression Language. Found at: <http://wiki.creativecommons.org/images/d/d6/Ccrel-1.0.pdf>
- [11] Java Web Start. Found at: <http://java.sun.com/javase/technologies/desktop/javawebstart/index.jsp>
- [12] Docvert Tool, available at: <http://holloway.co.nz/docvert/>