

Dynamic Trust Negotiation For Flexible e-Health Collaborations

Oluwafemi Ajayi, Richard Sinnott, Anthony Stell
National e-Science Centre
University of Glasgow
G12 8QQ, United Kingdom
{o.ajayi, r.sinnott, a.stell}@nesc.gla.ac.uk

ABSTRACT

Security issues have always limited the way we do things. In an organisation we provide security by granting privileges to either identities or roles. However this becomes more challenging when the objective is collaboration across organisational boundaries. Numerous access control approaches exist today to address the cross-boundary control issues. However an optimal approach would be to fold remote security credentials into local security credentials, thereby bridging the gap that makes decentralised security policies for multi-domain collaboration difficult. In this paper, dynamic trust negotiation is presented as a possible optimal approach that provides support for decentralised access control. We show how trust pathways can be established and how remote security credentials could be folded to local security credentials through trust contracts.

Categories and Subject Descriptors

K.6.5 [Management of Computing and Information Systems]: Security and Protection; D.4.6 [Operating Systems]: Security and Protection—Access controls, authentication; C.2.3 [Computer-Communication Networks]: Network Operations

General Terms

Security, Management

Keywords

Access control, trust negotiation, e-health

1. INTRODUCTION

Security issues have always limited the way we do things. In an organisation we provide security by granting privileges to either identities or roles. However this becomes more challenging when the objective is collaboration across organisational boundaries. Today, identity-based management systems

and attribute/role-based systems are two widely used access control management systems.

Identity-based management systems use entity identity to manage access to resources for a domain. These systems require a foreknowledge of identities that need access to particular resources. In a large organisation this is often difficult to maintain or properly control. These systems in most cases are designed based on access control lists (ACL) and they inherit manageability and control issues that are known to ACL. Examples of identity-based systems are grid systems that use a grid-mapfile[14, 19].

Attribute/Role-based access control systems manage access to resources by granting privileges to attributes/roles. These attributes/roles are assigned to entities, which eliminate some of the identity-based issues. Typically these attributes take the form of a digital attribute certificate [4]. These attribute certificates form part of an entity credential that they present at the point of resource request. Attributes together with policy assertions enable multiple resource providers to control access in a distributed environment. The use of policy assertions enables multiple resource providers to co-exist in the same environment to share resources.

Today many distributed access controls approaches are based on these two management methods. One of these approaches is the use of a single access control policy that governs authorisation across domains. This is achieved when all collaborating domains pre-negotiate and agree on privileges amongst other things on access to shared resources[17]. The implication of this approach includes detailed knowledge and agreements of global policies comprising lots of local policies, global policy maintenance, initial integration effort and the static relationship between security attributes/ credentials. Another variation of this approach is delegating authority to remote entities to assign privileges to their (remote) users [18].

Another approach to the single policy paradigm is based on a shared ontology. In this approach collaborating parties agree on a security ontology that describe roles or privileges that could be used for collaboration. Each collaborating party then map their local security ontology to the shared ontology in order to access remote resources. Remote parties in turn map shared security ontology to their local security ontology in order to grant or deny access to their resources [12, 13]. Associated implications of this approach are high global ontology maintenance and moderate initial integration effort.

Ideally an optimal approach would be to fold remote secu-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Mardi Gras Conference '08, Jan 31 – Feb 2, Baton Rouge, Louisiana, USA
Copyright 2008 ACM X-XXXXX-XX-X/XX/XX ...\$5.00.

rity credentials into local security credentials, thereby bridging the gap that makes decentralised security policies for multi-domain collaboration difficult. This approach should have very low knowledge of global policies, low maintenance of global policies and allow feasible relationships between security attributes/credentials.

In this paper we present Dynamic Trust Negotiation (DTN) as an optimal approach that provides support for decentralised access controls, where *direct* trust negotiation with non-trusted entities is unacceptable. We show how remote security credentials can be folded into local security credentials through trust contracts. In section 2 we review trust negotiation and introduce DTN. Section 3 expands on *circle of trust* and *trust contracts* while section 4 presents the DTN negotiation architecture. In section 5 we describe the DTN implementation in Virtual Organisation for Trials of Epidemiological Studies (VOTES), its performance and its performance similarities with the Border Gateway Protocol. We conclude in section 6.

2. BACKGROUND

Trust is the underlying phenomenon of any security system. Most security systems are designed using security policies, which define and describe what is trusted, how it is trusted and where it is trusted. Trust is built on the concept of limiting expected behaviour [9]. It is associated with an assurance measurement. That is, the level of confidence in limiting behaviour within a security policy determines the level of assurance.

Automated trust negotiation (ATN) [21] is the process of establishing trust between strangers through the exchange of digital credentials. These credentials are sensitive information and are often protected through the use of disclosure policies. These disclosure policies inevitably require negotiation strategies as each entity tries to protect what credentials are released. However for a negotiation to succeed entities are expected to operate using the same family of disclosure strategies[23].

Different trust negotiations approaches have been proposed to support access control policies in open decentralised environments [20, 10, 22, 16, 11]. Some approaches are based on a trust negotiation framework in the context of a peer-to-peer environment. [22] introduces a *locally trusted third party* (LTTP) which acts like a cache and mediator between two entities for the purpose of successful trust negotiations in peer-to-peer systems. Similarly [10] introduces a *sequence prediction module* that caches and manages used credential sequences from previously trust negotiations. While [16] proposes a trust chain based negotiation strategy (TRANS), which dynamically constructs trust relationships using a *trust proxy* that can cache common credentials or partial trust chain information from previous negotiations.

Various ATN systems have been developed, they include Trust-X [10] and TrustBuilder [24]. Trust-X is a framework that provides an XML-based language that is used to encode policies and certificates for trust negotiations. It also provides a peer-to-peer architecture used for negotiation management. TrustBuilder is an architecture that focuses on negotiation strategies. The architecture verifies credentials and checks policy compliance. Other systems like Traust [15] have been developed to augment TrustBuilder to provide interaction between applications or systems that offer trust negotiation services.

2.1 Dynamic Trust Negotiation

Dynamic trust negotiation (DTN) formalised in [6], is the process of realising trust between strangers or non-trusting entities (e.g. domains), through locally trusted intermediary entities. Trust is realised when an entity delegates its digital credentials to trusted intermediary entities through which it can interact with non-trusted entities. These intermediary entities can in turn delegate to other intermediary entities resulting in what we call *n-tier* delegation hops. The trust negotiation process involves trust delegations through intermediary trusted entities on behalf of non-trusting entities, where *direct* trust negotiation with non-trusted entities is unacceptable. Any entity can serve as a negotiator for other entities provided it is trusted by the two non-trusting entities or by their intermediaries.

DTN explores how credentials can be negotiated as the basis to support collaborative research between autonomous, distributed resources. It addresses the heterogeneous and autonomous issues of trust management like credentials and policies in multi-domain environments. DTN negotiates credentials between trusted parties also known as a *circle of trust* (COT) [6], who act as mediators on behalf of strangers and thus bridge trust gaps. This bridge also reduces the risk associated with disclosing policies to strangers.

As an example of *circle of trust*, consider the following scenario. Alice from the Glasgow Royal Infirmary hospital - hereafter referred to as domain GRI - is an investigator on a cancer clinical trial. She wants to recruit patients onto specific trials and in doing so needs to query patient consented health records in Scotland. To achieve this, she logs in to the trial portal and her credentials (privileges/attributes/roles...) are pulled from her domain, e.g. through Shibboleth pull or push from the portal service provider or GRI identity provider respectively. The trial portal initiates a credential negotiation request with all other domains that GRI trust such as Southern General Glasgow hospital (SGG). SGG returns patient records that satisfies GRI request based on Alice's credentials and delegated privileges at SGG. SGG also negotiates with other domains it trust such as Royal Infirmary Edinburgh (RIE) using Alice's SGG delegated privileges. Similarly, RIE negotiates with other domains it trust using SGG's RIE delegated privileges. Thus GRI, SGG, RIE are *trust pathways*. The request process continues with nodes joining the trust pathways until all possible trust paths are exploited. These negotiated credentials such as *RIE.investigator* are forward to GRI, which then makes a query request with these credentials on behalf of Alice.

$$\begin{aligned}
 GRI.investigator &\leftarrow Alice \\
 GRI.circleOfTrust &\leftarrow SGG \cup SGH \cup GRH \\
 SGG.circleOfTrust &\leftarrow RIE \cup IRH \\
 GRI.investigator &\leftarrow SGG.delegatedInvestigator \\
 &\quad \cap RIE.investigator
 \end{aligned}$$

where Southern General Hospital is referred to as SGH; Gartnavel Royal Hospital as GRH; and Inverclyde Royal Hospital as IRH.

DTN differs from ATN in that (1) negotiation occurs between trusted parties and not between strangers even though the goal is to realise trust between strangers. (2) it introduces multiple hops and delegation into the trust negotiation, which resolves some heterogeneity issues. (3) it limits

disclosure of access control policies, which reduces the need for disclosure strategies. However ATN can be used between trusted parties to negotiate for new trust contracts.

3. TRUST NEGOTIATIONS

In this section we describe two features of our trust negotiation system: *circle of trust* and *trust contract*.

3.1 Circle of Trust

In the formalised model of DTN, the concept of *circle of trust (COT)* [6, 7] for trust negotiation introduced. Figure 1 describes a *COT*, which is a network of locally trusted intermediary peers that a peer (or entity) trust and collaborates with through one or more trust-contracts between each peer. A trust contract *TC* is an agreement that exists between two entities. This network of trusted peers enable interactions between peered and non-peered domains.

Through overlapping *COT*s a trust-pathway (chain of trust) can be discovered. Consider two peers P_1 and P_5 , where P_1 is a requester and P_5 is a resource provider in another domain. P_1 and P_2 has $\{P_3, P_4, P_6, P_7\}$ and $\{P_3, P_4, P_5\}$ in their *COT* respectively. For P_1 to access P_5 resources, they will need to be trusted by P_2 . In addition, P_2 will need to understand and trust credentials from P_1 . Since P_1 has trust relationships with $\{P_3, P_4\}$, which are also in trust relationship with P_2 , P_1 will initiate a trust negotiation with P_2 through $\{P_3, P_4\}$. Similarly, P_2 will initiate a trust negotiation with P_5 . Thus $\{P_3, P_2\}$, $\{P_4, P_2\}$ are trust-pathways between P_1 and P_5 . Hence trust is realised by exploring overlapping *COT*s between P_1 and P_5 .

$$P_1 \leftarrow (P_3 \vee P_4) \leftarrow P_2 \leftarrow P_5$$

That is trust is realised between P_i and P_j when:

$$P_i \leftarrow P_j : \\ COT(P_i) \cap COT(P_{i+1}) \dots \cap COT(P_j) \neq \{\}$$

COT improves the likelihood of successful negotiations as peers can cache trust chains from previous negotiations, which will reduce the likelihood of future negotiations failing. The cache can also speed up future trust negotiations. However, this additional benefit of *COT* is yet to be explored in our current implementation.

The advantages of having *COT* are quickly overshadowed as the number of overlapping *COT* increases. This is because the more hops you have, the less likely peers will be delegating privileges in open decentralised collaborative systems.

Despite this limitation, *COT* provides an additional benefit. Overlapping *COT*s can help to abstract virtual organisations through which trust can be discovered and realised dynamically. In virtual organisations, relational hierarchies often exist, which can be modelled over the underlying *COT*.

3.2 Trust Contract

The presence of multiple domain authorities and policy enforcement introduces a policy semantics divide between domains, i.e. knowing that *org1.investigator* = *org2.investigator*. Trust contracts *TC* [6] are static agreements between two trusting peers to map credentials between their domains. These agreements cover key management and identity management (authentication) issues.




Figure 1: Circle of Trust




Figure 2: A network of collaborating health organisation

Trust contracts provide one mechanism to overcome the semantic issue of what a credential from one domain means (or should mean) in another domain. However trust contracts require that overlapping *COT*s exist.

Figure 2 shows an abstract network of a collaborative environment. The network, which is a non-negative, acyclic graph is denoted as $G(V,E)$. The Node set V represents autonomous organisations. A node refers to an end point in a communication chain and consists of security credentials. The Edge set E represents the direction of trust, which consists of policies and constraints. Edges have weights, which represents the cardinality of Trust-Contracts (*TC*) sets between two Nodes. A *tc*, $tc \in TC$ is an agreement between two nodes (u, v) that states the mapping/relationship between two credentials (c^u, c^v) . *TC* exists between two nodes when one or more credential mappings are agreed between them, that is:

$$TC = (\{c^{u_0}, c^{v_0}\}, \{c^{u_1}, c^{v_1}\}, \dots, \{c^{u_k}, c^{v_k}\})$$

Relationships between credentials are based on *credential equivalence rules*. A *tc* stems from these rules, which are modelled by function *tc*:

Let c^u and c^v be the set of credential in domain u and v respectively.

$$[c^u, c^v]$$

$$\left| \begin{array}{l} tc : c^u \leftrightarrow c^v \\ \exists x : c^u; y : c^v \bullet tc(x) = y \end{array} \right.$$

Trust contracts provide one solution to credential equivalence problem that exist between autonomous organisations by using equivalence rules. Credential equivalence rules define the relation that exist between credentials. These relations are used in the folding of one credential to another be-

tween different organisations. Some of the credential equivalence rules modelled by trust contracts are as follows:

1. Transitive membership rule:

$$R \leftarrow R_1 \text{ and } R_1 \leftarrow R_2 \Rightarrow R \leftarrow R_2$$

This rule means that R_1 is a member of R and R_2 is a member of R_1 , then R_2 is a member of R . As an example,

$$\begin{aligned} org1.investigator &\leftarrow org2.healthpractitioner \\ org2.healthpractitioner &\leftarrow org3.specialist \\ \Rightarrow org1.investigator &\leftarrow org3.specialist \end{aligned}$$

2. Linking delegation rule: $R \leftarrow R_1 \cdot R_2$

This rule means an entity that has R_2 can act as R if the entity also has R_1 . Requires at least two dependent roles and the order of dependency matters. As an example,

$$\begin{aligned} votes.investigator &\leftarrow org1.generalpractitioner.investigator \\ org1.generalpractitioner &\leftarrow org2.gp \\ generalpractitioner.investigator &\leftarrow gp.investigator \\ org2.gp.investigator &\leftarrow org3.investigator \\ \Rightarrow votes.investigator &\leftarrow org3.investigator \end{aligned}$$

3. Intersection rule: $R \leftarrow R_1 \cap \dots \cap R_k$ implies an entity that has R_1, R_2, \dots , and R_k is a member of R . In most cases the least upper bound of the intersections is inferred. An intersection of two independent roles must be a non-empty set.

4. ARCHITECTURE

Two systems make up our DTN architecture, discovery system and negotiation system. Figure 3 shows the architecture of the negotiation system. In this section we describe the main components that make up the architecture. This architecture is based upon Security Assertion Markup Language (SAML) [2] as the underlying framework. The protocol of the discovery service (not shown in the figure) that we referred to in the architecture is described in [6].

Negotiation Service.

This service is the point of interaction between domains. It provides a secure interface through which domains can request and exchange attribute assertions. The service encapsulates the negotiation protocol that is used for interpreting messages and carrying out the process of checking *trust contracts*. The service interacts with the negotiation agent and enforces the decision of the agent. Similarly (not shown in the figure), the service interacts with the discovery service in order to identify other trusted domains (next-hops), when it is acting as an intermediary domain.

Trust Enforcement Point.

This component interacts with the negotiation agent and enforces the decision of the agent by communicating responses or by interacting with the SAML module. It communicates responses through the negotiation service to a requester or to an intermediary domain. Based on the negotiation agent decision, it initiates a negotiation request with other next-hop domains on behalf of the requester.

Negotiation Agent.

An agent must understand the protocol used for trust negotiation as in [15] and manages the negotiation session. An agent validates a negotiation request and checks that access and release policies are not violated. The attribute assertions received are validated against access policies and checked against trust contracts that exist between domains. Depending on the negotiation strategy in use, further requests can be made for more attribute assertions from the request domain. The agent checks the release policies upon validating the accept policies. If any release policies are satisfied, attribute assertions are issued for further negotiations with other intermediary domains or for interaction with the SAML module.

SAML+.

The SAML *plus* module are triggered when a domain is the targeted domain. It is called SAML *plus* because it extends SAML by using a Negotiated Attribute Store (NAS). The store is populated with attribute assertions that are issued based on the domain's release policy. The SAML extended module is described in Section 5.1.

5. IMPLEMENTATION

In the Virtual Organisations for Trials and Epidemiological Studies (VOTES) project we are investigating the application of Grid technologies in the clinical trials domain. The project is a collaboration involving several UK universities - Glasgow, Oxford, Imperial, Manchester, Nottingham and Leicester. The project goal is to set up data grids that further enhance data quality and support clinical trials and epidemiological studies. This has been achieved by federating clinical data across the regional and national health board boundaries that exist in the UK. Three key areas of trials and studies has been supported: Patient Recruitment, Data Collection and Study Management.

The sensitive nature of clinical data makes security a high priority and any method of federating this data must adhere rigorously to the local security policies that protects this data. In addition, health providers like the NHS are only willing to interact with parties they have explicit contracts with [1]. However, the flexibility required for implementing cross-boundary data queries must also be maintained; hence a grid solution to provision data and operation security must be developed for collaboration to succeed. To this end, prototype portals have been developed as part of the VOTES project, designed with modular role-based access control and supporting fine-grained access control.

Access control is applied at two levels in the portal:

- At the local resource level, a local manager denotes what roles are allowed to access data in local databases.
- At the VO level, an access matrix is available that denotes what roles can access what data, before the query can be executed. The construction of this matrix is based on agreements reached between the local managers.

Unifying these two security policies adequately is a great challenge. The ideal situation is when ultimate control rests with the local resource providers - and the VO policy simply acknowledges their autonomy. In order for this to be achievable, these roles (or credentials) must be negotiated




Figure 3: Negotiation Service Architecture

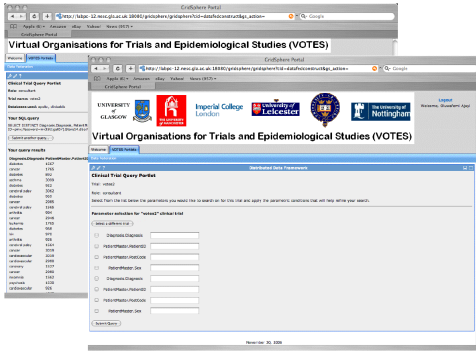


Figure 4: A VOTES Portal

and exchanged between resource providers (nodes) in a flexible and secure way. DTN facilitates this by introducing a negotiation layer, where the local trust policies are managed by Resource Managers (RM) which grant or deny access to their resources based on negotiated assertions. It should be noted that data providers are unwilling to negotiate with trusted parties directly.

In figure 4, a view of the portal is shown. This view is driven by the roles a user has across all federated resources. Based on various trials, different nodes have different roles implemented in their local policies such as nurse, investigator, consultant, administrator, GP, neurologist, psychiatrist, etc. These roles are defined by each organisation based on existing privileges within their respective organisations. For instance a nurse role in Org1 cannot necessary act as a nurse in Org2 even if a nurse role exists in Org2, except if a trust contract already exists asserting that fact between Org1 and Org2. Similarly Org3 might not have a nurse role and so Org2 will have to negotiate based on trust contracts that exist between them where possible.

5.1 SAML-DTN

The negotiation layer was implemented as GT4-based – negotiation and discovery services. In order to integrate with VOTES, we have an Identity Provider (IdP) connector that initiates credential negotiation via the negotiation service. The discovery service is used to realise trust-pathways, which must be invoked whenever a new path needs to be discovered or when existing paths need rediscovering.

When a user tries to access a remote data resource that is protected by a Service Provider (SP), they are redirected to their home Identity Provider (IdP) through the WAYF service [3]. The user authenticates at the IdP e.g. using a LDAP repository. The IdP sends the SP a SAML [2] response that contains an authentication assertion. This assertion is forwarded to the SP’s assertion consumer service, which validates the assertion. This authentication assertion includes a temporary pseudonym for the user (the handle) that the SP can use to reference the user. After validating the authentication assertion, the SP creates an attribute-token, which is sent to the user’s IdP’s attribute authority along with the user’s handle in a SAML attribute query message. The attribute-token and the user’s handle are linked together by the SP to provide the SAML-DTN support.

On receiving a SAML attribute query message, the IdP initiates the home negotiation service using the SP’s attribute-token to negotiate the user attributes. The home negotiation service negotiates the user attributes with nodes (organisations/sites) that serves as next-hop nodes to the target node (SP). Each negotiation hop includes the passing of attribute-token from node to node, which uniquely links negotiated attributes to an attribute-token. Attribute-token and attributes negotiated between the last-hop node and the target node are stored by the target node in it’s Negotiated Attributes Store (NAS).

When an IdP receives negotiation responses from its next-hop nodes, it returns a SAML attribute assertion message to the SP. If negotiations were successful, the assertion notifies the SP to collect negotiated attributes from it’s NAS. The SP uses a user’s handle to retrieve a user’s attribute-token, used to query a NAS for negotiated attributes. The negotiated attributes are used to make authorisation decisions as to what the user can access. If negotiations fail, the assertion includes null attributes. These null attributes are invalid at the SP and hence cannot be used for authorisation decisions by the SP.

5.2 Performance

We tested our DTN implementation across four *COTs*. The network (trust) topology used where arbitrary but were constrained to four *COTs*. The reason for this constraint was based on our simulated experiment [8], which indicated an exponential fall in performance as the number of *COTs* involved in trust negotiation increases.




Figure 5: SAML-DTN model view

We tested the performance of our implementation using several scenarios. Each of the scenarios were tested over several runs comprising the averages of 10 runs on similar network (trust) topologies, each with a total of 8 nodes. Each node was a 2.2Ghz Celeron with 512MB RAM running Linux. All nodes had Grid middleware installed hosting both the discovery and negotiation services. In all our scenarios we identified a node as a source node and another node as a target node. In the first scenario, a source node requests the discovery of a target node by sending route messages to nodes that exists in it's *COT*. This scenario executed an average of 59 seconds. In the second scenario, a source node initiates attribute negotiation with nodes that serves as intermediaries for a target resource. The last negotiation feedback took 15 seconds.

We are carrying out further tests as network (trust) topologies size, number of hops and size of *trust contracts* were seen to affect our results. However the results were as expected for our test environment. The combined effect of number of hops and size of *trust contracts* had the most effect in our tests.

5.3 Similarities with BGP

Border Gateway Protocol (BGP) is an inter-autonomous system routing protocol [5]. BGP is used to exchange routing or reachability information with other BGP systems. Internet service providers (ISP) use this protocol to determine a route to a destination. Since multiple paths to a destination could exist, BGP uses various attributes (or properties) to determine the best route. BGP attributes include: weights, local preference, origin and next-hop [5].

As with BGP, *circle of trust*, number of hops and size of *trust contracts* are attributes that affect the way DTN performs. The combined effect of these attributes exhibits characteristics that are similar to BGP. DTN similarities with BGP include the following:

- Weights (number of *trust contracts* (*TC*), trust levels, level of risks): in this case the route with the highest weights will be preferred. Since negotiation success is based on the satisfiability of one or more *TC* at each hop, more *TCs* at each node improve the negotiation success rate.
- Local preferences: if there are multiple paths to a destination, the local preference attribute is used to select the next node for a particular destination. Local preferences are based on the level of trust, which are often based on past negotiations with a particular node. For example if a previous negotiation was successful through a particular node, that node would be preferred in subsequent negotiations.

- Next-hop attribute: as nodes create trust-pathways [6] to a target node, intermediary nodes that exist in the local *COT* are identified as next-hops. These intermediary nodes are then considered when a target node is to be reached.

6. CONCLUSION

In this paper we presented dynamic trust negotiation as a possible optimal approach for decentralised access control within the e-Health Domain. This approach enables remote security credentials to be folded into local security credentials through trust contracts. We described the design and architecture of DTN services, discussed our implementation and reviewed its performance. Lastly, we discussed similarities between DTN attributes (e.g. *circle of trust* and *trust contract*) and the border gateway protocol (BGP).

In the future we intend to leverage knowledge from similar protocols. We expect that improvements such as peer clustering, a method similar to route reflector in BGP will enhance our trust negotiation model. In addition, we plan to explore the effect of attribute hierarchies on trust contracts as this could enhance trust negotiations.

Acknowledgments

This work is supported by a grant from the Medical Research Council (MRC) UK and is undertaken as part of the Virtual Organisation for Trials of Epidemiological Studies (VOTES) 3-year project.

7. REFERENCES

- [1] Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data. http://ec.europa.eu/justice_home/fsj/privacy/law/index_en.htm.
- [2] Organization for the Advancement of Structured Information Standards (OASIS). Assertions and Protocol for the OASIS Security Assertion Markup Language (SAML) V2.0, OASIS Standard, 15 March 2005.
- [3] Shibboleth Architecture Protocols and Profiles. <http://shibboleth.internet2.edu/docs/draft-mace-shibboleth-arch-protocols-latest.pdf>.
- [4] ITU-T Recommendation X.509 | ISO/IEC 9594-8: Information Technology Open Systems Interconnection the Directory: Public-key and Attribute Certificate Frameworks, 3, May 2001.
- [5] RFC 4271, A Border Gateway Protocol 4 (BGP-4). <http://tools.ietf.org/html/rfc4271>, Jan. 2006.
- [6] O. Ajayi, R. Sinnott, and A. Stell. Formalising Dynamic Trust Negotiations in Decentralised Collaborative e-Health Systems. In *Proceedings of the 2nd International Conference on Availability, Reliability and Security, (ARES07), Vienna, Austria*. IEEE Computer Society, Apr. 2007.
- [7] O. Ajayi, R. Sinnott, and A. Stell. Trust Realisation in Collaborative Clinical Trials Systems. In *HealthCare Computing Conference HC2007, Harrogate, England, Mar. 2007*.

- [8] O. Ajayi, R. Sinnott, and A. Stell. Trust Realisation in Multi-domain Collaborative Environments. *To Appear in Proceedings of 6th IEEE International Conference on Computer and Information Science, ICIS'07*, July 2007.
- [9] M. Benantar. *Access Control Systems: Security, Identity Management and Trust Models*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2005.
- [10] E. Bertino, E. Ferrari, and A. C. Squicciarini. Trust-X: A Peer-to-Peer Framework for Trust Establishment. *IEEE Transactions on Knowledge and Data Engineering*, 16(7):827–842, 2004.
- [11] V. Bharadwaj and J. Baras. Towards Automated Negotiation of Access Control Policies. In *Proceedings of the Fourth International Workshop on Policies for Distributed Systems and Networks (Policy 2003)*. IEEE Computer Society Press, 2003.
- [12] M. Boniface and P. Wilken. *ARTEMIS: Towards a Secure Interoperability Infrastructure for Healthcare Information Systems*, pages 181–189. From Grid to Healthgrid. IOS Press, 2005.
- [13] M. Ehrig and Y. Sure. Ontology Mapping - An Integrated Approach. In *Proceedings of the First European Semantic Web Symposium*, volume 3053 of *Lecture Notes in Computer Science*, pages 76–91. Springer Verlag, MAY 2004.
- [14] I. T. Foster, C. Kesselman, G. Tsudik, and S. Tuecke. A Security Architecture for Computational Grids. In *ACM Conference on Computer and Communications Security*, pages 83–92, 1998.
- [15] A. J. Lee, M. Winslett, J. Basney, and V. Welch. Traust: A Trust Negotiation-based Authorization Service for Open Systems. In *SACMAT '06: Proceedings of the eleventh ACM symposium on Access control models and technologies*, pages 39–48, New York, NY, USA, 2006. ACM.
- [16] J. Li, J. Huai, J. Xu, Y. Zhu, and W. Xue. TOWER: Practical Trust Negotiation Framework for Grids. *2nd IEEE International Conference on e-Science and Grid Computing*, Dec. 2006.
- [17] L. Pearlman, V. Welch, I. Foster, C. Kesselman, and S. Tuecke. A Community Authorization Service for Group Collaboration. In *POLICY '02: Proceedings of the 3rd International Workshop on Policies for Distributed Systems and Networks (POLICY'02)*, page 50, Washington, DC, USA, 2002. IEEE Computer Society.
- [18] R. Sinnott, J. Watt, J. Koetsier, D. Chadwick, O. Otenko, and T. Nguyen. Supporting decentralized, security focused dynamic virtual organizations across the grid. In *Proceedings of 2nd IEEE International Conference on e-Science and Grid Computing, Amsterdam, December 2006*, 2006.
- [19] V. Welch, F. Siebenlist, I. Foster, J. Bresnahan, K. Czajkowski, J. Gawor, C. Kesselman, S. Meder, L. Pearlman, and S. Tuecke. Security for Grid Services. In *Proceedings of 12th IEEE International Symposium on High Performance Distributed Computing*, pages 48–57, June 2003.
- [20] W. Winsborough and L. Ninghui. Safety in Automated Trust Negotiation. In *Proceedings of IEEE Symposium on Security and Privacy, 2004*, pages 147–160, 2004.
- [21] W. H. Winsborough, K. E. Seamons, and V. E. Jones. Automated Trust Negotiation. *DARPA Information Survivability Conference and Exposition (DISCEX)*, 01:0088, 2000.
- [22] S. Ye, F. Makedon, and J. Ford. Collaborative Automated Trust Negotiation in Peer-to-Peer Systems. In *P2P '04: Proceedings of the Fourth International Conference on Peer-to-Peer Computing (P2P'04)*, pages 108–115, Washington, DC, USA, 2004. IEEE Computer Society.
- [23] T. Yu, M. Winslett, and K. E. Seamons. Interoperable Strategies in Automated Trust Negotiation. In *CCS '01: Proceedings of the 8th ACM conference on Computer and Communications Security*, pages 146–155, New York, NY, USA, 2001. ACM Press.
- [24] T. Yu, M. Winslett, and K. E. Seamons. Supporting Structured Credentials and Sensitive Policies through Interoperable Strategies for Automated Trust Negotiation. *ACM Trans. Inf. Syst. Secur.*, 6(1):1–42, 2003.