

Formalising Dynamic Trust Negotiations in Decentralised Collaborative e-Health Systems

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

Access control in decentralised collaborative systems present huge challenges especially where many autonomous entities including organisations, humans, software agents from different security domains seek to access and share resources in a secure and controlled way. Automated trust negotiation (ATN) is one approach that has been proposed for trust discovery and realisation, which enables entities viz. strangers to access resources across autonomous boundaries through iterative exchange of credentials. Various negotiation strategies have been proposed to protect credential disclosure during trust negotiations. However in some domains such as e-Health, not all entities are willing to negotiate credentials or disclose access policies directly to strangers regardless of negotiation strategies and instead prefer to negotiate and disclose sensitive information only to strangers within what we refer to as a circle of trust. In this paper, we introduce a formal model to describe how locally trusted intermediary parties can provide multiple negotiation and delegations hops to protect credentials and access policies. We propose a dynamic trust negotiations (DTN) model that not only protects sensitive information from disclosure but also reduces semantic issues that exist with credentials in decentralised systems. This work is currently being explored and implemented within the e-health domain: specifically in the MRC-funded Virtual Organisation for Trials of Epidemiological Studies (VOTES) project.

Index Terms – Trust Negotiations, Security, P2P

1 Introduction

Access control in decentralised collaborative systems present huge challenges where many autonomous entities such as organisations, humans, software agents from different security domains seek to access and share resources in a secure and controlled way. It is largely

understood how to control access to resources within a given domain, however considerable challenges remain with regards to decentralised access control between collaborating autonomous remote entities. The ideal solution would be a scalable distributed security approach where trust is easily discovered and realised and used to securely extend site autonomy to support collaborative work.

Organisations are often aware that certain resources exist in other organisations but usually will have to negotiate access rights or privileges with those target organisations to access their resources. In some cases, a target organisation publishes those resources they are willing to share with other organisations on a per collaboration basis. In other words an organisation or a user in an organisation will know that certain resources exist and yet lack the means to access them because they do not have appropriate credentials. When this happens, the requesting organisation will initiate a negotiation process with the target organisation for privileges typically given as attributes for resource sharing. These agreements are often difficult to reach because of different organisational security requirements. The agreement challenge is exacerbated when the number of organisations involved in the agreement stage is large and dynamic. A fundamental cause of this is the lack of co-ordination and acceptance of agreements by the wider communities. A common approach for inter-site security policies is pre-exchange of security credentials between organisations. Among the disadvantages to this method are (a) credentials revocation (b) credentials re-distribution and (c) credentials duplication/redundancy.

In most cases when an organisation advertises the availability of some of their resources (for resource discovery), they already have localised access control policies in place that protect it from unauthorised usage. For example an NHS hospital might be willing to make available some statistical data to any organisation within the health services or health-research insti-

tutions. But proving your identity (authentication) as a researcher from a health-research institution is not sufficient to guarantee access without proper privileges (authorisation) because of local security policies that exist at the hospital. In the health domain there is an urgent need to share and collaborate e.g. to identify potential participants, seek consents, recruit participants, collect data from on-going trials at all trial sites and manage on-going studies. All these will require access to electronic health records such as patient demographics, patient medical history, test results, current treatment notes, past and current prescriptions, and so forth. However the lack of access to this geographically and autonomously distributed information may delay a trial or affect the success of a trial.

This paper introduces a formal model of Dynamic Trust Negotiation (DTN) to address the heterogeneous and autonomous federation of credentials and policies. The model describes how locally trusted intermediary parties can provide multiple negotiation and delegations hops to help establish trust between strangers. The model protects credentials and access policies and reduces credential semantic issues that exist in decentralised systems. In Section 2 we review negotiations and *inter*-organisation interactions. Section 3 introduces trust negotiation in decentralised systems. Section 4 presents a trust discovery model using graphs. Section 5 introduces the system used to explore these models within the MRC funded Virtual Organisation for Trials of Epidemiological Studies (VOTES) project, and Section 6 presents our conclusions.

2 Background

Negotiations plays an important role in open pervasive/ubiquitous computing. Typically these forms of negotiation span two major system domains: agent technology and credential-based authorisation systems. Agent technology fits into this problem space as it deals with issues of distribution heterogeneity and autonomy are found in ubiquitous environments such as grid environment. Similarly credential-based authorisation that primarily supports *intra*-organisation access control needs to be extended to support the challenge of *inter*-organisation interactions.

Agent-based negotiation introduces the concept of dialogue between two agents in order to obtain resources [1, 2]. In multi-agent systems, where agents are autonomous as they often represent different individuals or organisations, negotiation is the main form of interaction between them as they cannot perform effectively on their own in order to achieve their objectives. [1, 3, 4] discusses some properties that serve as bases for

negotiation in a multi-agent domain. Properties such as competitive or co-operative behaviours of agents, such as negotiation protocols (auction, heuristic, and argumentation) and negotiation strategies, provide insight into negotiation issues as in trust negotiation. [4] investigates negotiation between partners in loose inter-organisation workflow (IOW) which is similar to the type of cooperation that exists in Grid-based virtual organisations (VOs). Other issues common to both IOW and VO include: finding partners to interact with; contracts specification as in VO agreements between partners, and workflow execution. Negotiation between partners typically follows the discovery of partners in order to realise a service. Although the choice of partners often depends on what is or can be negotiated between them.

Credential-based systems include identity-based systems and property-based systems. Identity-based systems use entity identity or names for authentication and authorisation purposes. These identity-based systems cannot provide the base for negotiating trust between unknown entities (strangers). However, property-based systems also known as attribute-based systems, use attribute certificates and policy assertions to control access to resources in distributed environments. The use of policy assertions enables multiple resource providers to co-exist in the same environment. Thus when an entity tries to access a resource, the entity provides its certificates to the policy enforcement engine that decides if access conditions are satisfied. These systems provide the base for negotiating trust between strangers.

The need to exchange credentials between unknown entities introduces the concept of trust negotiation (TN) otherwise known as automated trust negotiation (ATN)[5]. ATN is an approach for trust establishment between strangers through the exchange of sensitive information such as digital credentials. Trust is established through an iterative but cautious bilateral disclosure of credentials [6, 7]. Digital credentials which are analogous to paper credentials are digital assertions about a credential owner signed by the credential issuer. Currently digital credentials are widely implemented using X.509 certificates [8]. The credential is signed using the issuers private key and the signed credential is verified with the issuers public key. A credential contains attributes that describe properties of the owner asserted by the issuer. Credentials also contain the public key of the credential owner through which the owner can demonstrate its ownership by the corresponding private key. Negotiating these sensitive credentials without any human intervention is the basis of trust establishment [6, 9].

In this paper we explore how credentials can be negotiated as the basis to support collaborative research between autonomous, distributed resources. We note that in sensitive domains such as e-health it is often impossible to deal with strangers owing to the risk involved. This makes it much more difficult to support automated trust negotiations. However when an intermediary party is introduced that is known to both parties (strangers) then the associated risks are reduced since credentials are not perceived to be disclosed to strangers. DTN on the other hand negotiates credentials between known parties who act as mediators on behalf of strangers.

3 Trust Negotiation in Collaborative Systems

In trust negotiations, we regard *credentials also as resources* that can be negotiated. In collaborative research environments such credentials are often used for access control.

3.1 Access Control

Access to resources needs to be controlled and managed especially in ensuring that operations carried out on those resources are adequately authorised. Examples of these resources could be data being provided or services being rendered. Decisions have to be made and enforced in order to protect resources from unauthorised disclosure and alterations while confidentiality and privacy needs to be ensured where needed. Access management systems such as attribute-based access control systems (ABAC)[10, 11, 12] use information contained in policies and credentials to manage access. Desirable properties of an attribute-based access control (ABAC) system include:

1. Decentralised attributes where an entity's attribute can be asserted by another entity.
2. Delegation of attribute authority in which the authority over an attribute can be delegated to another entity.
3. Attribute inference where an attribute can be inferred through another attribute.
4. Attribute intersection in which combination of attributes are used to infer another entity's attribute(s).
5. Attribute fields useful for parameterised attributes such as defining quantities in a credential.

3.2 Modelling Access Control

Key elements of an access control model are: *E*: Environment which defines the context or environment of the access request: $e_0, e_1, \dots, e_k \in E$, e.g. an optional variable representing different virtual organisations (VOs); *S*: Subject, defining the entity such as user, software agent and organisation: $s_0, s_1, \dots, s_k \in S$; *OBJ*: Object, resources or targets more generally: $obj_0, obj_1, \dots, obj_k \in OBJ$; *ACT*: Action, actions on objects: $act_0, act_1, \dots, act_k \in ACT$; *P*: Permission, $P = ACT \times OBJ \times E$ that is $P = \langle act_i, obj_j, e_i \rangle$; *R*: Role is defined as $S.r(h_1, \dots, h_n)$ where r is role name, h_i is parameter for parameterised role's r and S is the issuer of the credential; *PS*: Permissions to role relation, $PS \subseteq P \times R$.

Based on the Role-based Trust-management (RT) language defined for attribute-based access control (ABAC) systems in [11], we define credentials as follows:

- + $S.r(h_1, \dots, h_n) \leftarrow A$ which implies S asserts that A has the attribute $R = r.(h_1, \dots, h_n)$, i.e. A is a member of S .
- + $S.r(h_1, \dots, h_n) \leftarrow A.r_1(l_1, \dots, l_n)$ which implies possibly the same entities with possibly the same role names, that is $R_A \in R_S \mid r_1(l_1, \dots, l_n) \in r(h_1, \dots, h_n)$.
- + $S.R \leftarrow A_1.R_1 \cap \dots \cap A_k.R_k$ which means any entity that has (or is a member of) R_1, R_2, \dots , and R_k is also a member of $S.R$.
- + $S.R \leftarrow S_1.R_1.R_2$ which implies a role-linkage, i.e. entity S asserts that entity A has the attribute R_1 and if A asserts that entity B has the attribute R_2 , then S asserts that B has the attribute R . This is could also be viewed as a form of attribute-based delegation.

3.3 Automated Trust Negotiation

Access control policies (a.k.a policies) and credentials can be defined with languages with well formed semantics and expressed as finite sets of statements[13]. Using propositional logic as in [14] we define a policy for resource D as:

$$P_D \rightarrow F_D(C_1, C_2, \dots, C_k)$$

where C_1, C_2, \dots, C_k are credentials that must be satisfied by the other party; F_D is an expression that entails these credentials, boolean operators \vee and \wedge , and any parenthesis where needed. Access is granted to a resource D when the other party discloses C_i that

satisfies C_k and when $F_D(C_1, C_2, \dots, C_k)$ is evaluated to true. For example Bob wants to access cancer patients records D at hospital X as part of a Cancer clinical trial (XCT). Hospital X’s policy requires the requestor to be an investigator or clinician on the XCT clinical trial before access can be granted. Thus Bob provides credentials such as C_1^{Bob} = “Investigator” or C_2^{Bob} = “Clinician” and C_3^{Bob} = “XCT”, which can be expressed as $P_D \rightarrow (C_1^{Bob} \vee C_2^{Bob}) \wedge C_3^{Bob}$. Similarly, Bob’s access control policies may specify that the requesting target, which in this case is hospital X prove its identity amongst other properties. So for Bob’s credential we have: $P_{C^{Bob}} \rightarrow F_{C^{Bob}}(C_1, C_2, \dots, C_k)$. In a nutshell, the policy of a resource is satisfied when the other party discloses a sequence of credentials for that resource (C_1, C_2, \dots, C_k) . A resource R is said to be unprotected if its access control policy is always satisfied $R \rightarrow true$ or $C \rightarrow true$. A resource is said to have a denial policy if $R \rightarrow false$, that is no credential can satisfy that policy or that resource is not meant for disclosure.

[11, 14] illustrated with examples how trust is established between two peers P_1 and P_2 with each of the peers requesting a series of credentials from each other and how requesting a credential in the series might trigger requests for credentials from the other party. One problem with exchanging credentials this way is that a point of deadlock is reached where both parties wait on each other to disclose the next credential. This credential negotiation deadlock is explained in [14], which occurs whenever there is cyclic credential interdependency: $C_2^X \leftarrow C_2^{Bob}$ and $C_2^{Bob} \leftarrow C_2^X$, where their credential disclosure policies disregard who is on first. [14] also proposed a possible solution to credentials negotiation deadlock. The solution introduces a collaborative peer to the negotiation process called a *locally trusted third party* (LTTP). An LTTP acts as a mediator by disclosing credentials and policy rules to negotiating parties whenever cyclic interdependency occurs to facilitate trust negotiation. A peer P_c is said to be a LTTP for P_a and P_b where P_c has previously successfully exchanged and cached several credentials on more than one occasion at different times with both P_a and P_b . Hence P_a and P_b ask P_c which they both trust to act as their LTTP. P_c then releases missing credentials to both parties, which breaks the cyclic interdependency.

ATN is not all about credential disclosure but also about access policy disclosure. [15, 6, 9, 11] all present models for negotiation strategies to protect the disclosure of sensitive credentials. However for a negotiation to succeed the negotiating peers must operate using the same strategies. [13] discusses the use of interoperable

strategies for credential exchange and why every entity should be free to use whatever strategy they choose before or during negotiation. Two strategies are said to interoperate if trust negotiation succeeds whenever it is possible. A family of disclosure tree strategies was presented which are all mutually interoperable. However we argue that if a trust negotiation succeeds, access policies would have been disclosed. In some context, these access policies are sensitive information that needs to be protected. In this paper we present a model that uses multiple negotiation and delegations hops to protect credentials and access policies.

3.4 Dynamic Trust Negotiation

Dynamic trust negotiation (DTN) also known as dynamic negotiation through delegated trust (DNDDT) is the process of negotiating trust between two non-trusting entities through trusted intermediary entities. The process involves trust delegations through intermediary entities on behalf of these non-trusting entities. Any entity can serve as a negotiator for other entities provided it is trusted by the two non-trusting entities or by their intermediaries. Like ATN, DTN introduces a mediator we call *locally trusted intermediary party* (LTIP) similar to LTTP [14] in ATN. Unlike ATN, an LTIP is just one of the multiple LTIP (many hops) that can exist in a trust negotiation between two peers.

Consider an example of dynamic trust negotiation between two peers P_1 and P_2 , where P_1 is a requestor and P_2 is the domain of the resource R . With the understanding that credentials are also resources, we have two forms of resources in this example: Objects and Credentials. P_1 wants to access an object resource on P_2 . P_1 will have to first negotiate its credential resource¹ for P_2 ’s credential. P_2 has never negotiated with P_1 and its only open for negotiation with peers it has previously negotiated with such as P_3 . We call this a *circle of trust*, shown in Figure 1. Suppose P_2 ’s access policy for R is:

$$\begin{aligned} R_1, R_2 &\subseteq R \\ R_1 &\leftarrow C_1^{P_3} \wedge C_2^{P_3} \\ R_2 &\leftarrow C_3^{P_3} \end{aligned}$$

which means P_2 requires credential C_1 and C_2 from P_3 for resource R_1 while C_3 is required for R_2 . Suppose P_1 belongs to the P_3 circle of trust and that P_3 access

¹From here on our reference to resource is mainly credential.




Figure 1: Circle of Trust

policy with P_1 is:

$$\begin{aligned} C_1^{P_3} &\leftarrow C_1^{P_1} \wedge C_2^{P_1} \\ C_2^{P_3} &\leftarrow C_3^{P_1} \\ C_3^{P_3} &\leftarrow C_3^{P_1} \wedge C_4^{P_1} \end{aligned}$$

To access R_1 , P_1 would have to negotiate with P_3 by making available:

$$\{C_1^{P_1}, C_2^{P_1}, C_3^{P_1}\}$$

while P_3 will negotiate on behalf of P_1 with:

$$\{C_1^{P_3}, C_2^{P_3}\}$$

However if P_1 only makes available:

$$\{C_3^{P_1}, C_4^{P_1}\}$$

then P_3 can only negotiate on behalf of P_1 with:

$$\{C_3^{P_3}\}$$

which will be for P_2 's R_2 and not R_1 .

From the above example it would be seen that, P_3 serves as a link peer known as LTIP for the trust negotiation between P_2 and P_1 .

$$\begin{aligned} P_2 &\leftarrow P_3 \\ P_2 &\leftarrow P_3.P_1 \end{aligned}$$

In a typical trust negotiation where a circle of trust exists, you often find multiple LTIP's involved in the trust negotiation. Each of the involved LTIP acts as a hop or a link to the next LTIP and/or finally to the target peer.

4 Dynamic Trust Negotiation Model

Consider one of the typical scenarios of a DTN model. Rob is a healthcare professional based at Glasgow

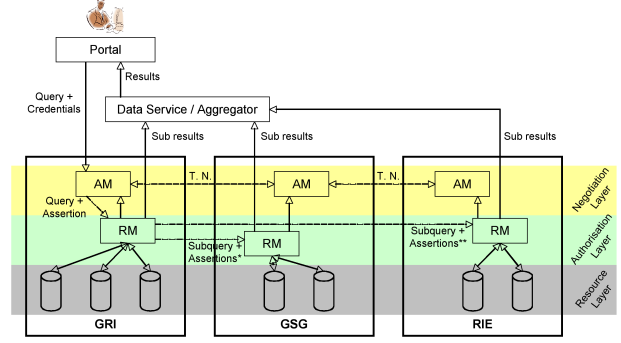


Figure 2: Using Dynamic Trust Negotiation in Clinical Trials Collaborative Environment

Royal Infirmary (GRI) and the principal investigator for the Scottish coronary clinical trial (SCCT) with role Investigator. He logs into the trial portal and his credentials are pulled from his domains credential repository. He decides to query for consented patient records for prospective participants. The portal pushes his credentials and query to GRI Access Manger (GRI-AM). GRI-AM sends a request for data along with Rob's Investigator credentials to peers that are in a static trust relationship with GRI such as the Glasgow Southern General Hospital Access Manager (GSG-AM). Query results are returned if applicable based on Rob's credentials and delegated privileges at GSG. GSG-AM sends a request using credentials it has delegated to Rob through GRI to other peers that GSG is in a static trust relationship with such as Royal Infirmary of Edinburgh Access Manager (RIE-AM). RIE-AM responds with delegated credentials through GSG-AM to GRI-AM. In addition to that RIE-AM also sends a request using credentials it has delegated to Rob through GSG-AM to other peers that RIE is in static trust relationship with. GRI, GSG, RIE are trust-pathways. The request process continues with nodes joining the trust-pathways until all possible trust paths are exploited. These negotiated credentials are forward to GRI-AM, which then makes query request with these credentials on behalf of Rob to each node's Resource Manager (RM). Figure 2 shows a collaborative environment that abstracts a negotiation layer (*DTN*) from an authorisation layer.

In Figure 3, a network is shown and represented as a non-negative, bi-directional, acyclic graph. The network is denoted as $G(V,E)$. The network is an abstraction of negotiation layer shown in Figure 2, which is introduced to augment an authorisation layer.

The Node set V is an abstraction of an autonomous organisation in a network of organisations. A node refers to an end point in a communication chain and




Figure 3: A network of collaborating health organisation

it consists of security credentials, i.e. roles, attributes, rules set or constraints, entities keys that are unique and in use at given end point. The Edge set E represents the direction of trust, which consists of labels, properties and constraints. The Edge also signifies the existence of a Trust-Contract (TC) between two Nodes, which is shown as a bi-directed arc between two Nodes. A TC m is a static agreement between two nodes (u, v) that states the mapping/relationship between two credentials (c^u, c^v) . That is:

$$m = (u, v) \text{ where } u, v \in V$$

That is,

$$f : c^u \rightarrow c^v : m \mapsto f(m)$$

A set of $m \in M$ exists between two nodes when more than one credential mapping is agreed between them, that is:

$$M = (\{u_0, v_0\}, \{u_1, v_1\}, \dots, \{u_k, v_k\})$$

4.1 Trust Discovery

In DTN we use a variant of the link state routing algorithm to discover chains of trusts or trust-pathways, which are necessary before credentials can be negotiated. The algorithm uses the weight w on paths/links, that is the cardinality of trust-contract sets that exist between nodes. In other words, the weight w is the sum of all m 's that exist between two nodes: $w = \sum M_{(u,v)}$.

Since DTN is not about shortest paths to a destination, the algorithm is modified to accommodate the discovery of multiple paths to a destination. Similarly as notification messages are sensitive information, notifications are restricted to trusted peers and messages are encrypted with shared keys or key pairs. Once each node collates routing information, nodes judiciously select appropriate nodes with respect to destination nodes.

Each node u keeps a list of nodes $v_0, v_1, \dots, v_k \in V$ in its *circle of trust (cot)* along with their respective weights, encryption keys and rules-set also known as constraints. When an entity at a node S in figure 3,

makes a request for remote resources², say a list of cancer patients and a discovery service suggests relevant data exists at nodes I, J, D , node S will check for nodes I, J, D in its *cot*. Node S will issue a query request for cancer patients from nodes I, J which are in its *cot*. The request is made with the credentials of the entity making the request. However if trust-contracts that contains those user credentials do not exist at that remote node, a trust-contract request is raised for those credentials. Alternatively, a dynamic trust negotiation (DTN) request through other nodes is explored. Details of a trust-contract request and formation are outside the scope of this paper.

Node D does not exist in S 's *cot*, hence a DTN request is initiated. A DTN request starts off with determining trust-pathways to destination nodes if such information does not already exist in the node's trust-pathways table. The trust-pathways table contains lists of *cot* nodes that act as 'next hops' or links to *non-cot* nodes. The list is prioritised based on the weights of those links. The trust-pathways table is updated as follows.

Link Request Assuming figure 3 where S is the source node and D is the destination node, S sends link request (LREQ) to nodes that exist in its *cot*. Node S contain nodes I, J, K in *cot* and they will all receive a LREQ. A typical LREQ has the following properties: a source IP address, destination IP address, a sequence number and a message authentication code (MAC) which is a message signed using a shared key³.

$$LREQ : \{DN_s, DN_d, Seqnum\} + K_{12}(MAC)$$

Where DN_s and DN_d are distinguished names for source and destination nodes respectively. The sequence number $Seqnum$ is a unique number maintained by the source node for each request made for a destination node. $K_{12}(MAC)$ is a message authentication code signed using K_{12} , which is a key shared between two nodes i.e. sender node and receiver node. Shared keys are created when two nodes first negotiate for a trust-contract. Alternatively a public-private key pair could be used for signing the MAC.

Each node implements a LREQ table that stores link requests it receives from source nodes and intermediary nodes. When a non-destination node receives a LREQ,

²We assume that a discovery service exists that returns a list of target nodes providing certain resources.

³We assume that at initial trust-contract negotiation, nodes exchange keys for message authentication and those keys are periodically maintained. Alternatively a public-private key pair can be used for message authentication.

it checks its *cot* to see if the sender is a trusted node, reads the MAC using their shared key and forwards the message to other trusted nodes re-computing the MAC using keys it shares with those nodes. It then stores that node as ‘next hop’ to the source node in its trust-pathways table. Similarly, if it receives the same LREQ from multiple nodes, it stores those nodes as ‘next hop’ to the source node in its trust-pathways table.

Link Reply The destination node DN_d , on receiving an LREQ, creates a link reply (LREP) which is sent to all nodes it receives a LREQ from.

$$LREP : \{DN_d, DN_s, Seqnum\} + K_{12}(MAC)$$

Where DN_d and DN_s are distinguished names for destination and source nodes respectively. The sequence number *Seqnum* is a unique number maintained by the destination node for each reply made to a source node. $K_{12}(MAC)$ is a message authentication code signed using K_{12} , which is a key shared between two nodes i.e. sender node and receiver node.

When a non-source node receives a LREP, it checks its *cot* for the sender, checks for the corresponding LREQ in its LREQ table and if valid, verifies the MAC using its key. For every valid MAC, a node updates its trust-pathways table registering the LREP sender node as the ‘next-hop’ to the destination node. The non-source node also re-broadcast the LREP to other nodes it receives a LREQ from but replaces the MAC with a MAC signed using keys it shares with those nodes.

Link Update A node may revoke its trust contracts with other nodes and thus renders some links in the trust chains invalid. Similarly new nodes may be added at any time and new trust relations added. Thus when a node detects broken links it sends error messages to other nodes that are in its *cot*. An error message will contain a MAC, which is signed using shared keys. When a node receives link errors, it authenticates the sender and verifies the MAC using their shared key. If the MAC is valid, it updates its trust-pathways table, and if other links to a destination do not exist, it re-sends error messages to other nodes in its *cot*.

Negotiations Once link information exists in either the *cot* or the trust-pathways tables, nodes can prioritise ‘next hop’ nodes based on link weights. The more the TCs/weights the more the chances of having a successful negotiation though not necessarily a guarantee that negotiations will be successful. Similarly,

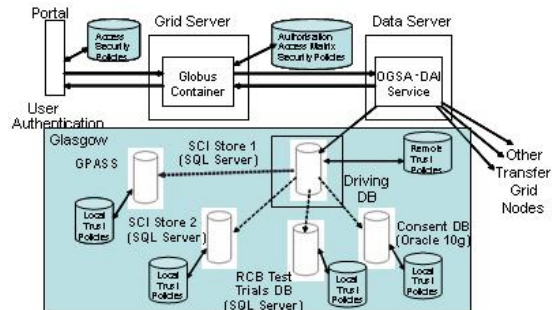


Figure 4: VOTES Architecture

each intermediary node that acts as a hop in the negotiation process prioritised their links with other nodes based on weights for their negotiations.

5 VOTES

The VOTES project[16] is focused upon development of clinical trials infrastructures using Grid technologies. A beta version architecture and implementation is in place for exploring scenarios of federating data across the clinical domain. Figure 4 shows the architecture, with various Grid technologies and a pool of databases housing representative clinical data in actual NHS service resources. The authorization security implemented is composed of local resources (i.e. the database pool) providing local access control policies, and a larger-scale VO-wide access matrix, which is an aggregation of these local policies.

This access control matrix lists individual parameters within each database and matches them against the various roles that have been defined for a particular trial. In the prototype constructed so far, this access matrix model is essentially a centralised model for a given node (for instance, Glasgow). But in the wider scope of the VOTES project, the vision is that nodes at other sites (such as Oxford or Manchester) will join the virtual organisation and share seamless, but secure, access to their resources with the Glasgow node⁴.

This peer-to-peer model brings up the issue of how the access matrix authorization policy between two sites will be shared. For virtual organisations to be truly flexible, nodes requesting to join the VO need not necessarily be known or trusted by all other nodes.

The DTN provides the mechanism by which two initial strangers can establish trust through the exchange of security credentials. If a node that is not known to

⁴Based on the definition of a ‘Grid’ model, these nodes are equally dispensable, with the overall system relying on no single entity.

all VO nodes, requests to collaborate *across* the VO require the node to have a credential-based trust relationship with at least one of the nodes in the VO (i.e. a mutual friend). Using DTN a trust path can then be established between the new node and the rest of the VO, allowing their incorporation into the VO.

By extension, when a new node, with a set of new databases is added to the VO (after establishing trust using the DTN model), the access matrix within all the nodes will be able to query the meta-data of these new databases and populate itself with the necessary parameters and descriptions. The resources within this new node will then be available for querying as part of the VO.

6 Conclusion

In this paper we have introduced a dynamic trust negotiation (DTN) model for the discovery and realisation of trust between strangers, which makes distributed authorisation possible in a collaborative environment. This model differs from ATN models with the introduction of locally trusted intermediary parties that provide multiple negotiation and delegations hops, which protects credentials and access policies in collaborative environments. DTN explores trust relationships that exist between collaborators to discover trust pathways through which trust can be negotiated.

References

- [1] F. Sadri, F. Toni, and P. Torroni, "Logic Agents, Dialogues and Negotiation: An Abductive Approach," in *In Proceedings of the Symposium on Information Agents for ECommerce, AISB'01, March 2001*, 2001.
- [2] C. Bartolini, C. Priest, and N. R. Jennings, "A software framework for automated negotiation," in *Software Engineering for Multi-Agent Systems III: Research Issues and Practical Applications*, pp. 213–235, 2005.
- [3] N. Vulkan and N. R. Jennings, "Efficient mechanisms for the supply of services in multi-agent environments," in *ICE '98: Proceedings of the first international conference on Information and computation economies*, (New York, NY, USA), pp. 1–10, ACM Press, 1998.
- [4] E. Andonoff and L. Bouzguenda, "Agent-Based Negotiation between Partners in Loose Inter-Organizational Workflow," *iat*, vol. 0, pp. 619–625, 2005.
- [5] E. Bertino, E. Ferrari, and A. Squicciarini, "Trust Negotiations: Concepts, Systems, and Languages," in *IEEE Computational Science and Engineering*, pp. 27–34, 2004.
- [6] W. H. Winsborough, K. E. Seamons, and V. E. Jones, "Automated Trust Negotiation," *DARPA] Information Survivability Conference and Exposition (DISCEX)*, vol. 01, p. 0088, 2000.
- [7] W. Winsborough and J. Jacobs, "Automated Trust Negotiation Technology with Attribute-based Access Control," in *In Proceedings of DARPA Information Survivability Conference and Exposition, 2003*, vol. 02, pp. 60–62, 22–24, Apr. 2003.
- [8] "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.
- [9] W. Winsborough and L. Ninghui, "Safety in Automated Trust Negotiation," in *In Proceedings of IEEE Symposium on Security and Privacy, 2004*, pp. 147–160, 2004.
- [10] L. Ninghui, J. Mitchell, and W. Winsborough, "Design of a Role-based Trust-management Framework," in *Proceedings of the 2002 IEEE Symposium on Security and Privacy*, 2002.
- [11] W. H. Winsborough and N. Li, "Towards Practical Automated Trust Negotiation," in *Proceedings of the Third International Workshop on Policies for Distributed Systems and Networks (Policy 2002)*, pp. 92–103, IEEE Computer Society Press, June 2002.
- [12] Y. Zhang, X. Li, J. Huai, and Y. Liu, "Access Control in Peer-to-Peer Collaborative Systems," in *Proceedings of the 25th IEEE International Conference on Distributed Computing Systems Workshops (ICDCSW'05, IEEE, 2005*.
- [13] 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*, (New York, NY, USA), pp. 146–155, ACM Press, 2001.
- [14] 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)*, (Washington, DC, USA), pp. 108–115, IEEE Computer Society, 2004.
- [15] K. Seamons, M. Winslett, and T. Yu, "Limiting the disclosure of access control policies during automated trust negotiation," in *Proc. Network and Distributed System Security Symposium, San Diego, CA, Apr. 2001*.
- [16] R. Sinnott, O. Ajayi, and A. Stell, "Development of Grid Framework for Clinical Trials and Epidemiological Studies," in *HealthGrid 2006 Conference, Valencia, Spain, June 2006*, 2006.