

Trust and Security in Grids: A State of the Art

*Marcim Adamski¹, Alvaro Arenas², Angelos Bilas^{3,4},
Paraskevi Fragopoulou^{3,5}, Vasil Georgiev⁶, Alejandro Hevia⁷,
Gracjan Jankowski¹, Brian Matthews², Norbert Meyer¹,
Jorg Platte⁸, Michael Wilson²*

adamski@man.poznan.pl, A.E.Arenas@rl.ac.uk, bilas@ics.forth.gr,
fragopou@ics.forth.gr, v.georgiev@fmi.uni-sofia.bg, ahevia@dcc.uchile.cl,
gracjan@man.poznan.pl, B.M.Matthews@rl.ac.uk, meyer@man.poznan.pl,
Joerg.Platte@udo.edu, M.D.Wilson@rl.ac.uk

¹ Poznan Supercomputing and Networking Center, Poland

² E-Science Centre, STFC Rutherford Appleton Laboratory, UK

³ Foundation for Research and Technology-Hellas (FORTH), Greece

⁴ University of Crete, Greece

⁵ Technological Educational Institute of Crete, Greece

⁶ Institute on Parallel Processing, Bulgarian Academy of Sciences (IPP-BAS), Bulgaria

⁷ University of Chile, Chile

⁸ University of Dortmund, Germany



CoreGRID White Paper
Number WHP-0001
May 26, 2008

Contributions from all institutes

CoreGRID - Network of Excellence
URL: <http://www.coregrid.net>

Trust and Security in Grids: A State of the Art

Marcim Adamski¹, Alvaro Arenas², Angelos Bilas^{3,4},
Paraskevi Fragopoulou^{3,5}, Vasil Georgiev⁶, Alejandro Hevia⁷,
Gracjan Jankowski¹, Brian Matthews², Norbert Meyer¹,
Jorg Platte⁸, Michael Wilson²

adamski@man.poznan.pl, A.E.Arenas@rl.ac.uk, bilas@ics.forth.gr,
fragopou@ics.forth.gr, v.georgiev@fmi.uni-sofia.bg, ahevia@dcc.uchile.cl,
gracjan@man.poznan.pl, B.M.Matthews@rl.ac.uk, meyer@man.poznan.pl,
Joerg.Platte@udo.edu, M.D.Wilson@rl.ac.uk

¹Poznan Supercomputing and Networking Center, Poland

²E-Science Centre, STFC Rutherford Appleton Laboratory, UK

³ Foundation for Research and Technology-Hellas (FORTH), Greece

⁴ University of Crete, Greece

⁵ Technological Educational Institute of Crete, Greece

⁶Institute on Parallel Processing, Bulgarian Academy of Sciences (IPP-BAS), Bulgaria

⁷ University of Chile, Chile

⁸ University of Dortmund, Germany

CoreGRID WHP-0001

May 26, 2008

Abstract

The Trust and Security activity in CoreGRID runs as a horizontal integration activity related to all the research areas, making the Network participants aware of the use of the technologies associated with trust and security. This paper presents an overview of the different concepts and technologies relevant to trust and security in Grid systems. It analyses the relation between trust and security, describes trust and security challenges in the Grid, and introduces the existing mechanisms for managing trust and security. The core of the document is the trust and security requirements across the CoreGRID Institutes, and the description of the work being carried out to meet such requirements.

1 Introduction

The CoreGRID Network of Excellence aims at strengthening and advancing scientific and technological excellence in the area of Grid and Peer-to-Peer technologies. This objective is achieved through a joint programme of activity structured around six research areas: knowledge and data management; programming models; system architecture; Grid information and monitoring services; resource management and scheduling; problem solving environments, tools and Grid systems. The Trust and Security activity in CoreGRID runs as a horizontal integration activity related to all the research areas, making the Network participants aware of the use of the technologies associated with trust and security.

This document presents an overview of the different concepts and technologies relevant to trust and security in Grid systems. Their content is based on surveys on Grid security carried out previously within the Network [DIA03,

This research work is carried out under the FP6 Network of Excellence CoreGRID funded by the European Commission (Contract IST-2002-004265).

DIA16]. The document is organised as follows. The next section analyses the relation between trust and security, and prepares the content of the rest of the paper. Then, section 3 presents trust and security challenges in the Grid, and describes the existing mechanisms for managing trust and security. Section 4 discusses the impact of trust and security in the areas tackled by CoreGRID Institutes. Finally, Section 5 concludes the paper by summarizing the work and highlighting future work.

2 Trust and Security

In the Internet world, trust has been recognised as an important aspect of decision making for electronic commerce [Gra00, Jos07]. Customers must trust that sellers will provide the services they advertise, and will not disclose private customer information (name, address, credit card details, etc). Trust in the supplier's competence and honesty will influence the customer's decision as to which supplier to use. Sellers must trust that the buyer is able to pay for goods or services, is authorised to make purchases on behalf of an organisation or is not underage for accessing service or purchasing certain goods.

How is the situation in the Grid? Fundamental to the Grid definition is the idea of *resource sharing* [Fos01]. The Grid was initiated as a way of supporting scientific collaboration, where many of the participants knew each other. In this case, there is an implicit trust relation, all partners have a common objective -for instance to realise a scientific experiment- and it is assumed that resources would be provided and used within some defined and respected boundaries. However, when the Grid is intended to be used for business purposes, it is necessary to share resources with unknown parties. Such interactions may involve some degree of risk since the resource user cannot distinguish between high and low quality resource providers on the Grid. The inefficiency resulting from this asymmetry of information can be mitigated through trust mechanisms.

This section analyses the concept of trust and its relation with security. There is a vast source of information on the theory and application of trust, For instance [Cas00, Wai02, Nix03, Jen04, Her05]. Here we visit the main definitions of trust and study the relation between trust and security.

2.1 Trust Definitions

This report focuses on trust in the context of networked and distributed computing systems. In this context, the remote system needs to be trusted, as well as interactions over underlying services such as communication services. As expressed by Grandison and Sloman [Gra00], the significance of incorporating trust in distributed systems is that trust is an enabling technology. Its inclusion will enable secure electronic transactions.

There is not consensus in the literature on what trust is [McK96]; it is recognised as an important and complex subject relating honesty, truthfulness, competence, reliability, etc. of the trusted person or service.

One of the influential works towards a practical definition of trust is given by Gambetta [Gam88b]: *"When we say we trust someone or that someone is trustworthy, we implicitly mean that the probability that he will perform an action that is beneficial or at least not detrimental to us is high enough for us to consider engaging in some form of cooperation with him. Correspondingly, when we say that someone is untrustworthy, we imply that that probability is low enough for us to refrain from doing so"*. Gambetta's definition stresses that trust is fundamentally a belief or estimation, which has inspired the use of subjective logic as a way of measuring trust [Jos99]. Castelfranchi and Falcone [Cas98] extend Gambetta's definition to include the notion of competence along with predictability.

Kini and Choobineh [Kin98] examine trust from the perspectives of personality theory, sociology, economics and social psychology. They highlight the implications of these definitions and combine their results to create their definition of trust in a system. They define trust as: *"a belief that is influenced by the individual's opinion about certain critical system features"*. Their analysis covers various aspects of human trust in computer dependent systems but they do not address the issue of trust between parties (humans or processes) involved in e-commerce transactions.

In the Trust-EC project of the European Commission Joint Research Centre (ECJRC), Jones [Jon99] defines trust as *"the property of a business relationship, such that reliance can be placed on the business partners and the business transactions developed with them"*. Jones states as relevant issues such as the identification and reliability of business partners; the confidentiality of sensitive information; the integrity of valuable information; the prevention of unauthorised copying and use of information; the guaranteed quality of digital goods; the availability of critical information; the management of risks to critical information; and the dependability of computer services and systems.

Grandison and Sloman [Gra00] survey various definitions of trust. Following a brief analysis of these definitions, they build their own one as *"the firm belief in the competence of an entity to act dependably, securely and reliably within a specified context"*. They argue that trust is a composition of many different attributes - reliability, dependability, honesty, truthfulness, security, competence and timeliness - which may have to be considered and defined depending on the environment in which trust is being specified.

Dimitrakos [Dim01] has defined trust as follows: *"Trust of a party A in a party B for a service X is the measurable belief of A in B behaving dependably for a specified period within a specified context in relation to X"*. In his definition, a party can be an individual entity, a collective of humans or processes, or a system; the term service is used in a deliberately broad sense to include transactions, recommendations, issuing certificates, underwriting, etc; dependability is used broadly to include security, safety, reliability, timeliness, and maintainability; a period may be the duration of the service, refers to the past, future (a scheduled or forecasted critical time slot), or always; finally, the term context refers to the relevant service agreements, service history, technology infrastructure, legislative and regulatory frameworks that may apply.

Josang, Ismail and Boyd [Jos07] define trust as *"the extent to which one party is willing to depend on something or somebody in a given situation with a feeling of a relative security, even though negative consequences are possible"*. They argue that their definition includes aspects such as dependence on the trusted entity or party; the reliability of the trusted entity or party; utility in the sense that positive utility will result from a positive outcome, and negative utility will result from a negative outcome; and a certain risk attitude in the sense that the trusting party is willing to accept the situational risk resulting from the previous elements.

Some aspects of these definitions are common, other are complementary. For example, [Gam88b] emphasises that trust is in part subjective, a characteristic present in other definitions such as [Gra00], [Dim01] and [Jos07]. [Gra00] underlines that trust is a belief in the competence of an entity within a specified context, while [Kin98] lay stress on that the entity that manifests trust (the "trustor") is the human - not the system. The definition in [Jon99] focuses on the aspect that in commerce trust is relative to a business relationship. One entity may trust another entity for one specific business and not in general. Such business relationship can be seen as the context of [Gra00] definition. Finally, the definition in [Dim01] highlights an important point, trust evolves in time and is measurable.

We do not intent to provide a definition of trust, rather to show the diversity of definitions and those points in common: subjective, context and evolution in time, among others. In the next part we analyse how trust is related to security, for the case of distributed systems.

2.2 Trust and Security

In general, the purpose of security mechanisms is to provide protection against malicious parties. Traditional security mechanisms typically protect resources from malicious users by restricting access to only authorised users. However, in many situations within distributed applications one has to protect oneself from those who offer resources so that the problem is in fact reversed. For instance, a resource providing information can act deceitfully by providing false or misleading information, and traditional security mechanisms are unable to protect against this type of threat. As noted in [Jos07], trust systems can provide protection against such threats. The difference between these two approaches to security was first described by Rasmusson and Janssen in [Ras96] who used the term hard security for traditional mechanisms like authentication and access control, and soft security for what they called social control mechanisms, of which trust is an example.

Grandison and Sloman [Gra00] have defined a trust classification as a useful way of categorising the literature relating to trust in Internet services. We have found such taxonomy helpful in linking trust and security for the purpose of this work. Trust is specified in terms of a relation between a trustor, the subject that trusts a target entity, and a trustee, the entity that is trusted. [Gra00] defines the following classes of trust.

- **Service Provision Trust** describes the relying party's trust in a service or resource provider. The trustor trusts the trustee to provide a service that does not involve access to the trustor's resources. This type of trust is essential for Grids, and can be seen as a minimal trust requirement in dynamic Virtual Organisations (VOs). Many Grid applications assume this type of trust implicitly; a partner in a VO presupposes a service provision trust as a result of participating in VO, although the VO does not provide mechanisms to enforce it.

In general, service provision trust is related to the reliability or the integrity of the trustee. For instance, in e-banking the customer trusts the bank to support mechanisms that will ensure that passwords are not divulged, and to maintain the privacy of any information such as name, address and credit card number. The Liberty

Alliance Project uses the term "business trust" to describe a provision trust, a mutual trust between companies emerging from contract agreements that regulate interaction between them [Lin04]. Mobile code and mobile agent-based applications also include service provision trust; the mobile code trusts the execution environment provided by the remote system [Dan07].

- **Resource Access Trust** describes trust in principals for the purpose of accessing resources owned by the relying party. A trustor trusts a trustee to use resources that he own or controls. Resource access trust has been the focus of security research for many decades [Abr95], particularly on mechanisms supporting access control. Generally, resource access trust forms the basis for specifying authorisation policies, which then are implemented using access control mechanisms, firewall rules, etc.

[Gra00] highlights the distinction between trusting an entity to read or write a file on your server and trusting an entity to execute code within your workstation. Simple file access requires that the trustee will follow the correct protocol, will not divulge information read, and will write only correct data, etc. Allowing an entity to execute code on your workstation implies much higher level of trust. The code is expected not to damage the trustor's resources, to terminate within reasonable finite time and not to exceed some defined resource limits with respect to memory, processor time, local file space, etc. [Sur02] has also drawn the attention to the case of trusting an entity to execute remote code in Grids; it shows practical examples of the possible consequences how to minimise dangers.

- **Delegation Trust** denotes the case when a trustor trusts a trustee to make decisions on his behalf, with respect to a resource or service that the trustor owns or controls.

Although delegation is conceptually simple, designing and deploying it within a Grid environment has proved to introduce problems regarding security. Such security implications have been analysed by Broadfoot and Lowe in [Bro03a], work carried out in the context of the EU DataGrid project. A point that is addressed is the level of trust assumed when delegation is employed, in particular the effect of having onwards delegation. They also investigate all the security implications for two delegation mechanisms widely used in Grids: delegation chaining [Gas90] and call-back delegation [Fos98].

- **Certification Trust** is based on the certification of the trustworthiness of the trustee by a third party, so trust would be based on a criteria relating to the set of certificates presented by the trustee to the trustor.

Trust systems that derive certification trust are typically authentication schemes such as X.509 and PGP [Zim95]. This class of trust is called "authentication trust" in Liberty Alliance [Lin04] and "identity trust" in [Jos07]. Grandison [Gra00] views certification trust as a special form of service provision trust, since the certification authority is in fact providing a trust certification service; however Josang [Jos07] views certification trust and service provision trust as two layers on top of each other, where provision trust normally cannot exist without certification trust; in the absence of certification trust, it is only possible to have a baseline provision trust in an entity.

Certification trust has played an important role in Grid environments; it is present with the inclusion of certification authorities, which play a central role in the Grid Security Infrastructure [Nag03] and have been exploited in production Grids [Joh03].

- **Context Trust** describes the extend to which the relying party believes that the necessary systems and institutions are in place in order to support the transaction and provide a safety net in case something should go wrong. It refers to the base context that the trustor must trust. This type of trust is called infrastructure trust in [Gra00], here we prefer to use the broader term of context trust used by [Jos07], which also involves social and legal factors such as insurance and legal system and law enforcement.

The main motivation of Grandison and Sloman's classification is to define classes of high-level trust specifications, which may be refined to low-level implementation policies, such as policies about access control, authentication and encryption [Gra03]. Gambetta [Gam88a] has highlighted that to make a society prosper, one needs rules (both written and unwritten), understanding of good and bad behaviour with its consequences and accountabilities, initial trust and earned trust, identification of the risks associated with transactions, and so on. As mentioned in [Sie05], a similar view should be taken if we want to achieve a secure Grid society. Many of the rules of the secure Grid society can be expressed in the form of trust specifications, which can consequently be refined into policies.

2.3 Dealing with Privacy

“Privacy is the claim of individuals, groups or institutions to determine for themselves when, how, and to what extent information about them is communicated to others” [Wes67]. Privacy is not a purely technical issue; it also involves aspects of legislation, corporate policy, and social norms. Further, privacy is a flexible concept in practice, based on individual perceptions of risk and benefit. There are several studies of privacy from different fields. In this part we analyse some definitions and relate them to the management of privacy for Grids.

In the social sciences, it is notably the work by Margulis [Mar03], who introduces privacy as a social issue and a behavioural concept, reviewing the benefits of obtaining privacy and the cost of failing to achieve and of losing privacy. From the human-computer-interaction view, Grudin describes several privacy issues that one must be dealt with when developing context-aware systems [Gru01]. He notes that “when people see benefits that outweigh risks, they voluntarily adjust their comfort levels regarding privacy”, citing surveillance cameras as a prime example. However, he also notes that a more fundamental problem is that technology is transforming what it means to be situated because we no longer control access to anything we disclose. Grudin’s essay emphasizes the potential dangers of personal data gathered in the past affecting a person in the future, leading to the need for limited data retention.

The work by Palen and Dourish [Pal03] argues that privacy is not simply a problem of setting rules and enforcing them, but rather an ongoing and organic process of negotiating boundaries of disclosure, identity, and time. They suggest *genres of disclosure* for managing interpersonal privacy, which are socially-constructed patterns of privacy management, as a sort of design pattern approach to support the development of privacy-sensitive applications. Examples include creating and managing accounts at shopping Web sites, taking appropriate photographs at social events, exchanging contact information with a new acquaintance, and the kinds of information one reveals to strangers. A person fulfils a role under a genre of disclosure through his performance of his expected role in that genre, and the degree to which a system does not align with that genre is the degree to which it fails to support the user’s and the genre’s privacy regulation process.

There is a vast literature on privacy in ubiquitous computing, since it is undisputed that a future world full of smart and cooperating artifacts will pose great risks to our personal privacy. The work in [Lan02] presents a privacy awareness system that allows data collectors to both announce and implement data usage policies, as well as providing data subjects with technical means to keep track of their personal information as it is stored, used, and possibly removed from the system. The interplay between privacy and risk has been recognised in [Hon04], which introduces privacy risk models for ubiquitous computing as a way of refining privacy from an abstract concept into a set of concrete concerns for a specific domain and community of users. Their privacy risk model consists of two parts: a *privacy risk analysis* which poses a series of questions to help designers think about the social and organizational context in which an application will be used, the technology used to implement that application, and control and feedback mechanisms that end-users will use. The second part looks at *privacy risk management*, and is a cost-benefit analysis intended to help designers prioritize privacy risks and develop architectures, interaction techniques, and strategies for managing those risks.

In Grid computing, most work related to privacy has been done for the case of Health Grids. For instance, the EU FP5 project GEMSS¹ analysed the legal aspects related to the privacy of patients’ data when using Grids [Mid04]. The HealthGrid white paper [HG04] also includes legal issues in relation to privacy when using Grids for health applications.

Work describe above shows different views of privacy. All of them highlight privacy as way of controlling information; not simply *limiting* what others know about one, but also stressing the need for one to *control* such information.

Fair information practices are a set of guidelines to help large organizations, such as corporations and governments, manage people’s personal information in a responsible manner [Wes67]. They include concepts such as notice, choice, security, and recourse. The fair information practices influenced the end-user privacy needs by describing the strong need for control and feedback (especially notice and consent), as well as limited data retention. The idea that personal information should be collected only for express purposes and that people should be able to access and amend their personal information.

2.3.1 A Note on Privacy Laws and Regulation

The existence of a legal and regulatory framework for privacy aims at increasing users’ trust when using information and communication technologies. Privacy laws and regulations vary widely throughout the world. Technology

¹<http://www.it.neclab.eu/gemss/>

developers need to know and follow such legislation.

The European Union Directive on Data Protection ² is the most comprehensive of set of data privacy laws currently in existence. In many respects, this directive closely follows the fair information practices described above, and includes several information privacy principles such as data quality (e.g., data is collected for specified purposes only, is not collected excessively, is accurate), legitimate processing (e.g., consent, notification of purpose and sharing, etc), adequate security, and so on.

The OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data ³ help to harmonise national privacy legislation and prevent interruptions in international flows of data. These guidelines represent a consensus on basic principles that can be built into existing national legislation, or serve as a basis for legislation in those countries which do not yet have it.

There is also a set of voluntary privacy guidelines such as the Online Privacy Alliance ⁴, the Networking Advertising Initiative Principles ⁵, and the Best Practices and Guidelines for Location Based Services of the Wireless Association ⁶.

3 Trust and Security in a Grid Environment

3.1 Trust and Security Requirements in the Grid

The Virtual Organisation (VO) is a key concept in the Grid community. A VO can be seen as a temporary or permanent coalition of geographically dispersed individuals, groups, organisational units or entire organisations that pool resources, capabilities and information to achieve common objectives. Depending on the context, dynamic ensembles of the resources, services, and people that comprise a scientific or business VO can be small or large, short- or long-lived, single- or multi-institutional, and homogeneous or heterogeneous. Trust and security challenges within the Grid environment are driven by the need to support scalable, dynamic distributed VO [Fos01].

The OGF has initiated the definition of the next-generation of Grid middleware by extending the emerging Web services technology that is currently being developed across the IT industry, under the umbrella of the Open Grid Services Architecture (OGSA). Trust and security requirements can be analysed from different perspectives. This section analyses requirements as defined by the OGF OGSA Security Workgroup, as well as according to the VO topology and through the different phases of the VO lifecycle.

3.1.1 Security Challenges According to OGF

The OGF OGSA Working Group has submitted a memo proposing a strategy for addressing security with OGSA [Nag03]. According to the group, the security challenges faced in a Grid environment can be grouped into three categories:

- integration solutions where existing services needs to be used, and interfaces should be abstracted to provide an extensible architecture;
- interoperability solutions so that services hosted in different virtual organizations that have different security mechanisms and policies will be able to invoke each other; and
- solutions to define, manage and enforce trust policies within a dynamic Grid environment.

A solution within a given category will often depend on a solution in another category. For example, any solution for federating credentials to achieve interoperability will be dependent on the trust models defined within the participating domains and the level of integration of the services within a domain. Defining a trust model is the basis for interoperability but trust model is independent of interoperability characteristics. Similarly level of integration implies a level of trust as well as a bearing on interoperability.

²http://ec.europa.eu/information_society/policy/ecommm/current/consumer_rights/data_protection/index_en.htm

³http://www.oecd.org/document/18/0,2340,es_2649_34255_1815186_L1_L1_00.html

⁴<http://www.privacyalliance.org>

⁵<http://www.networkadvertising.org/>

⁶<http://www.ctia.org/content/index.cfm/AID/11300>

In a Grid environment, where identities are organized in VOs that transcend normal organizational boundaries, security threats are not easily divided by such boundaries. Identities may act as members of the same VO at one moment and as members of different VOs the next, depending on the tasks they perform at a given time. Thus, while the security threats to OGSA fall into the usual categories (snooping, man-in-the-middle, intrusion, denial of service, theft of service, viruses and Trojan horses, etc.) the malicious entity could be anyone. An additional risk is introduced, when multiple VOs share a virtualized resource (such as a server or storage system) where each of participating VOs may not trust each other and therefore, may not be able to validate the usage and integrity of the shared resource.

- **The Integration Challenge**

For both technical and pragmatic reasons, it is unreasonable to expect that a single security technology can be defined that will both address all Grid security challenges and be adopted in every hosting environment. Existing security infrastructures cannot be replaced overnight. For example, each domain in a Grid environment is likely to have one or more registries in which user accounts are maintained (e.g., LDAP directories); such registries are unlikely to be shared with other organizations or domains. Similarly, authentication mechanisms deployed in an existing environment that is reputed secure and reliable will continue to be used. Each domain typically has its own authorization infrastructure that is deployed, managed and supported. It will not typically be acceptable to replace any of these technologies in favour of a single model or mechanism.

- **The Interoperability Challenge**

Services that traverse multiple domains and hosting environments need to be able to interact with each other, thus introducing the need for interoperability at multiple levels. At the protocol level, it is required mechanisms that allow domains to exchange messages; this can be achieved, for instance, via SOAP/HTTP. At the policy level, secure interoperability requires that each party be able to specify any policy it may wish in order to engage in a secure conversation-and that policies expressed by different parties can be made mutually comprehensible. Only then can the parties attempt to establish a secure communication channel and security context upon mutual authentication, trust relationships, and adherence to each other's policy. At the identity level, it is required mechanisms for identifying a user from one domain in another domain.

- **The Trust Relationship Challenge**

The VOs that underlie collaborative work within Grids may form quickly, evolve over time and span organisations; as discussed before, their effective operation depends on trust. In the simple case, personal knowledge between parties in the VO allows policies to be derived from identifiable trust "anchors" (parties vouching for other parties). An example in current Grid systems is the use of certificate authorities to root certificate-based identity mechanisms. For these to work, one must "know" about the trustworthiness of the certificate authority used to establish the identity of a party in order to bind it to specific usage policies. However, personal knowledge does not scale for the case on non-trivial VOs, which are most of the VOs, and it is required other technologies such as reputation management [Res00] to create and monitor relationships.

3.1.2 Security Issues According to the VO Topology

There have been several proposals to characterise VOs in topologies [Bur99, Kat00]. For instance, Burn et al [Bur99] defines six types of VOs, ranging from organisations providing services in the web (such as web shops or newspapers on the web) which does not control any user of the service to dynamic networks of entities collaborating to meet market opportunities. Arguably, some of the types of organisations do not comply with the way VOs are seen within Grid computing.

We focus here on three topologies that were introduced initially by Katzy et al. [Kat00], based on network topologies. These are shown in Figure 1 below.

From an organizational perspective, these VO types describe the main coordination structure that governs information and material flows as well as the power relationships and decision making within the network and projects.

- **Supply-Chain VOs**

A supply chain in general may be defined as a coordinated system of organizations, people, processes and resources that moves information or services from one end called the producer (or supplier) to another end called the consumer (or customer). Supply-chain VOs consist of several organisations that are collaborating to achieve the goal of supplying the consumer with the end product.



Figure 1: VO Topologies.

VOs, which adopt a supply chain topology, often use Supply Chain Management and Efficient Consumer Response (ECR) to improve inter-organisational co-ordination and control. Integration of information flow (e.g. EDI) and material flow creates transparency in the entire value chain and reduces waste and doubles effort in the virtual enterprise.

Apart from the usual security issues related to supply chains, such as the authentication of participants, maintaining the integrity of the moving product and its desirable properties, ensuring the security of the product while in transit and auditing at each stage, Kang et al. [Kan01] identify three other main security issues related to inter-organisational workflows, which may be regarded as the generic form of the VO supply chain topology. These are:

- Separation of workflow-level security requirements from organisation-level security. This is necessary since organisation may leave/enter a supply chain at any time or may merge/split from one another. Therefore, it is important that supply-chain-level security requirements be considered independently from local security at each organisation.
- Fine-grained and context-based access control is required since an inter-organisational supply-chain may consist of a large number of small tasks and each of these implies different usage contexts. Therefore, granting access at the level of organisations (or even users or processes) would be coarse-grained.
- Supporting dynamic constraints in order maintain desirable properties such as the separation of duties.

The UCON model, proposed by Park and Sandhu in [Park04] and adapted for Grids by Martinelli and Mori [Mar07], provides a means for solving the last two issues, since access is granted based on continuous monitoring of the behaviour of the entity requesting the access to some object and dynamic context-based constraints form a part of this model. CoreGRID work within the Institute of Grid Information, Resource and Workflow Monitoring Services is also tackling some of these issues, as shown in [Hoh06].

● **Hub and Spoke VOs (star or main contractor VOs)**

In a star topology, partners interact with one central hub or strategic centre. This type of VOs corresponds to a coordinated network of interconnected members, where each member provides key functionalities, and distinguished member plays the role of a leading actor (star), coordinating the whole operation of the VO.

Despite the decline in popularity of hub and spoke topologies in networks, they are still considered to be a good solution in VOs for large scale enterprise application integration problems. For example, having a central management unit facilitates the control of VO membership, even though this centralised management may become a performance bottleneck and a point of failure.

Traditional Grid security solutions, as those advocated in GSI [Nag03], have been focused on this type of topology. The CoreGRID Institute on Data and Knowledge Management has been analysing the security for this type of topologies for the case of Grid storage systems, as reported in [Lun07].

● **Peer-to-Peer VOs**

Peer-to-peer topologies are characterised by the lack of hierarchy where any peer may interact directly with any other peer. Their management is usually based on self-organisation.

Wallach [Wal02] highlights a few security concerns in peer-to-peer systems. These include:

- Secure routing, which ensures that when messages sent by non-faulty peers arrive at their non-faulty destinations without any compromises to their secrecy and/or integrity.
- Secure storage, where a node maintains the integrity of the data it stores and the data cannot harm the node (e.g. because the data contains a virus or a worm).
- Distributed auditing, which is useful in resource usage monitoring and control. This could be related to other issues outside of security, as in load-balancing.
- Trust and reputation, which are important factors in the security of peer-to-peer systems, in particular in identifying peers and evaluating their past behaviour. Such an issue arises due to the lack of hierarchy in the topology.

In CoreGRID, the Institute of Architectural Issues is investigating security issues in P2P VOs. For instance, a study on the application of reputation techniques in Grids is presented in [Sil07a]. Further, [Dom07] presents trust techniques for sabotage tolerance in P2P Desktop Grids and [Sil07b] describes how to tackle collusion threats in P2P Desktop Grids.

3.1.3 Requirement Analysis through the VO Lifecycle

The VO Roadmap project [Cam03] developed a VO lifecycle including phases such as identification, formation, operation/evolution and dissolution. The identification phase is dealing with setting up the VO; this includes selection of potential business partners by using search engines or looking up registries. VO formation deals with partnership formation, including the VO configuration distributing information such as policies, agreements, etc, and the binding of the selected candidate partners into the actual VO. After the formation phase, the VO can be considered to be ready to enter the operation phase where the identified and properly configured VO members perform accordingly to their role. Membership and structure of VOs may evolve over time in response to changes of objectives or to adapt to new opportunities in the business environment. Finally, the dissolution phase is initiated when the objectives of the VO has been fulfilled.

The TrustCoM project has derived security and trust requirement by analysing the lifecycle of a VO [Are05, Wes05]. Here we summarise such requirements.

- **VO Identification**

The identification phase addresses setting up the VO - this includes selection of potential business partners from the network of enterprises by using search engines or looking up registries. Generally, relevant identification information contains service descriptions, security grades, trust and reputation ratings, etc. Depending on the resource types, the search process may consist in a simple matching (e.g., in the case of computational resources, processor type, available memory and respective data may be considered search parameters with clear cut matches) or in a more complex process, which involves adaptive, context-sensitive parameters. For an example, the availability of a simulation program may be restricted to specific user groups or only for certain data types, like less confidential data, etc. The process may also involve metadata such as security policies or Service Level Agreement (SLA) templates with ranges of possible values and/or dependencies between them, such as bandwidth depending on the applied encryption algorithm. The identification phase ends with a list of candidates that potentially could perform the roles needed for the current VO.

After this initial step from the potentially large list of candidates, the most suitable ones are selected and turned into VO members, depending on additional aspects that may further reduce the set of candidates. Such additional aspects cover negotiation of actual Quality of Service (QoS) parameters, availability of the service, "willingness" of the candidate to participate, etc. It should be noted that though an exhaustive list of candidates may have been gathered during the identification phase, this does not necessarily mean that a VO can be realised - consider the case where a service provider may not be able to keep the promised SLA at a specific date due to other obligations.

In principle, the intended identification may fail due to at least two reasons: (a) no provider (or not enough providers) is able to fulfil all given requirements comes to SLA, security, etc. or (b) providers are not (fully) available at the specified time. In order to circumvent these problems, either the requirements may be reduced ("choose the best available") or the actual process may be delayed to be re-launched at a more suitable time. Obviously there may be the case, where a general restructuring of the requirements led to a repetition of the identification phase.

- **VO Formation**

At the end of the (successful) identification phase the initial set of candidates will have been reduced to a set of VO members. In order to allow these members to perform accordingly their anticipated role in the VO they need to be configured appropriately. During the formation phase a central component such as a VO Manager distributes the VO level configuration information, such as policies, SLAs, etc. to all identified members. These VO level policies need to be mapped on local policies. This might include changes in the security settings (e.g. open access through a firewall for certain IP addresses, create users on machines on the fly, etc.) to allow secure communication or simply translation of XML documents expressing SLAs or Obligations to a product specific format used internally.

- **VO Operation**

The operational phase could be considered the main life-cycle phase of a VO. During this phase the identified services and resources contribute to the actual execution of the VOs task(s) by executing pre-defined business processes (e.g. a workflow of simulation processes and pre- and post processing steps). A lot of additional issues related to management and supervision are involved in this phase in order to ensure smooth operation of the actual task(s). Such issues cover carrying out financial arrangements (accounting, metering), recording of and reacting to participants' performance, updating and changing roles and therefore access rights of participants according to the current status of the executed workflow, etc. In certain environments persistent information of all operations performed may be required to allow for later examination e.g. to identify fault-sources.

Throughout the operation of the VO, service performance will be monitored. This will be used as evidence when constructing the reputation of the service providers. Any violation -e.g. an unauthorised access detected by the access control systems- and security threats -e.g. an event detected by an intrusion detection system- need to be notified to other members in order to take appropriate actions. Unusual behaviours may lead to both a trust re-assessment and a contract adaptation. VO members will also need to enforce security at their local site. For example, providing access to services and adapting to changes and the violations.

Evolution is actually part of the operational phase: as participants in every distributed application may fail completely or behave inappropriately, the need arises to dynamically change the VO structure and replace such partners. This involves identifying new, alternative business partner(s) and service(s), as well as re-negotiating terms and providing configuration information as during identification, respectively formation phase. Obviously one of the main problems involved with evolution consists in re-configuring the existing VO structure so as to seamlessly integrate the new partner, possibly even unnoticed by other participants. Ideally, one would like the new service to take over the replaced partners' task at the point of its leaving without interruption and without having to reset the state of operation. However, it is considered application-specific to decide whether the VO remains the same despite the arrival/leaving of participants. There may other reasons for participants joining or leaving the VO, mostly related to the overall business process, which might require specific services only for a limited period of time - since it is not sensible to provide an unused, yet particularly configured service to the VO for its whole lifetime, the partner may request to enter or leave the VO when not needed. Provenance techniques as presented in [Mor08] can be used to register VO-evolution changes.

- **VO Dissolution**

During the dissolution phase, the VO structure is dissolved and final operations are performed to annul all contractual binding of the partners. This involves the billing process for used services and an assessment of the respective participants' (or more specifically their resources) performances, like amount of SLA violations and the like. The latter may of particular interest for further interactions respectively for other potential customers. Additionally it is required to revoke all security tokens, access rights, etc. in order to avoid that a participant may (mis)use its particular privileges. Generally the inverse actions of the formation phase have to be performed during Termination. Obviously partial termination operations are performed during evolution steps of the VO's operation phase.

3.2 Security Technologies in the Grid

This section presents the traditional security areas that play an important role in defining security for the Grids and the associated technologies. We build this analysis on top of previous surveys and text books on security for the Grids [Sur02, Bro03b, Cha07].

3.2.1 Authentication

Authentication deals with verification of the identity of an entity within a network. An entity may be a user, a resource or a service provided as part of the Grid. Authentication is one of the mechanisms helpful in implementing certification trust.

One of the technologies playing a central role in authentication is Public Key Infrastructure (PKI), which defines message formats and protocols that allow entities to securely communicate claims and statements. The most used assertions are those that bind identity and attributes statements to keys. The most popular PKI is defined by the IETF's PKIX working group, which defines a security system used for identifying entities (users and resources) through the use of X.509 identity certificates. In this PKI, highly trusted entities known as certificate authorities (CA) issue X.509 certificates where essentially a unique identity name and the public key of an entity are bound through the digital signature of that CA.

Revocation is vital for authentication, for example when a key is compromised or when a user's project ends. PKI relies upon the periodic distribution of Certificate Revocation Lists (CRLs) in order to allow those relying upon certificate to gain confidence in their present validity. The use of CRLs needs careful management, particularly in relation to the frequency of updates.

An important issue in authentication is the storage of credentials. Credential-storage systems take the responsibility of storing credentials securely, so that users can get credentials anytime on demand. The most popular credential-storage system for Grids is MyProxy [Bas05], which is an implementation of the Virtual Soft Token proposed by Sandhu in [San02], where X.509 proxy certificates are used to store and retrieve user credentials without having to expose the private key.

One of the challenges encountered in key management include the need of users of having different credentials, since users may play different roles or be part of several projects which have elected to trust different CAs. While PKI could handle this situation by signing the same public key into several different certificates, in practice the user may end up with numerous key pairs to manage. To link these different identities, the notion of federated identities has been developed, as shown in the Liberty Alliance project [Lin04]. Shibboleth [Sca05] is one of the most-used federated identity management systems. It is essentially a transport mechanism built on top of an institution's existing architecture that allows organisations to exchange information about their users in a secure and privacy-preserving manner. The basic idea explored by Shibboleth is the federated administration by the alliance of trusted home sites for users. Users are registered only at their home site rather than at each resource provider. Resource providers rely on users' home sites to provide identity information as well as attributes about users. In Shibboleth an opaque handle is returned as a response for an authentication request, which may not carry any user privacy information such as his identity. A user is allowed to select a subset of his attributes to present to providers.

3.2.2 Authorisation

Authorisation deals with the verification of an action that an entity can perform after authentication was performed successfully. In a grid, resource owners will require the ability to grant or deny access based on identity, membership of groups or virtual organisations, and other dynamic considerations. Thus policies must be established that determine the capabilities of allowed actions. Authorisation is closely related to access control trust. A good description of the current state of authorisation in Grid computing appears in [Cha05].

There are several architectural proposals for handling authorisation in Grids. One of the earliest attempts at providing authorisation in VOs was in the form of the Globus Toolkit Gridmap file. This file simply holds a list of the authenticated distinguished names of the Grid users and the equivalent local user account names that they are to be mapped into. Access control to a resource is then left up to the local operating system and application access control mechanisms. As can be seen, this neither allows the local resource administrator to set a policy for who is allowed to do what, nor does it minimise his/her workload. The Community Authorisation Service (CAS) [Pea02] was the next attempt by the Globus team to improve upon the manageability of user authorisation. CAS allows a resource owner to grant access to a portion of his/her resource to a VO (or community hence the name CAS), and then let the community determine who can use this allocation. The resource owner thus partially delegates the allocation of authorisation rights to the community. This is achieved by having a CAS server, which acts as a trusted intermediary between VO users and resources. Users first contact the CAS asking for permission to use a Grid resource. The CAS consults its policy (which specifies who has permission to do what on which resources) and if granted, returns a digitally self-signed capability to the user optionally containing policy details about what the user is allowed to do. The user then

contacts the resource and presents this capability. The resource checks that the capability is signed by a known and trusted CAS and if so maps the CAS's distinguished name into a local user account name via the Gridmap file.

The EU DataGrid and DataTAG projects developed the Virtual Organisation Membership Service (VOMS) [Alf03] as a way of delegating the authorisation of users to managers in the VO. VOMS is a system to attach a set of membership information (VOname, group and/or roles) to a user's own credentials, therefore allowing services to authorise users on the basis of these attributes instead of having to specifically list all users. VOMS has gone through a number of iterations in its development. While initially used only as an alternative means to generate gridmap files, services have quickly evolved to take full advantage of its capabilities, ranging from software like LCAS (Local Centre Authorization Service) and LCMAPS (Local Centre MAPPING Service) [Ste03] for the computing and storage services, to the recent integration with a policy system (G-PBox) capable of defining and evaluating policies based on group and role membership.

Other authorisation systems are Akenti and PERMIS. Akenti [Tho99] is an authorization infrastructure developed at Lawrence Berkeley National Laboratory. It is designed to address complex authorisation problems involving multiple administrative domains and multiple stakeholders (a stakeholder is an X.509 "source of authority"). For authentication, Akenti relies on X.509 certificates and the SSL/TLS protocol to securely authenticate a user, like most PKI-based systems. For authorisation, Akenti uses a pure pull model. When a resource request comes, the Akenti policy engine collects all relevant certificates from both the user and the resource, and derives the user rights from them. It is the server side that contact all authorities once the user gets authenticated. One potential problem within Akenti is its policies are expressed using a proprietary XML format rather than X.509 based. With many common features with Akenti, PERMIS [Cha02] is another policy-based authorization system, from the EC funded Privilege and Role Management Infrastructure Standards validation (PERMIS) projects. Unlike Akenti's distributed and hierarchical policies, the policy in PERMIS is a single attribute certificate stored in a LDAP directory. It supports the role based access control (RBAC) paradigm, which means PERMIS infers the access right (roles and attributes) according to the given user's DN, a resource and an action. It supports classical hierarchical RBAC in which superior roles inherit the privileges of subordinate roles in the hierarchy.

Grid developers have advocated for the need of more fine-grained authorisation mechanisms. One of the recent research efforts in this area comes from the EU GridTrust project, where Martinelli and Mori [Mar07] are adapting the UCON model for Grid system. The UCON model, proposed by Park and Sandhu in [Park04], is a new access control model that addresses the problems of modern distributed environments. In UCON, the existence of a right for a subject is not static, but it depends on dynamic factors. This is possible because, while the standard access control model is based on authorisations only, UCON extends this model with other two factors that are evaluated to decide whether to grant the requested right: obligations and conditions. Moreover, this model introduces mutable attributes paired with subjects and objects and, consequently, introduces the continuity of policy enforcement. Mutable attributes can be updated before (preUpdate), during (onUpdate), or after (postUpdate) that the action is performed. If the attribute is updated before the action, the new value could be exploited to evaluate the authorisation predicate and to determine the right to execute this action, while if the attribute is updated after the execution of the action, the new value will be exploited for the next actions.

3.2.3 Confidentiality

The data being processed in a Grid may be subject to considerable confidentiality constraints, either due to privacy concerns or issues of intellectual property. For instance, grid applications may involve medical data [Bra03], bioinformatics and genomic databases [Cro05] and industrial design information [Wes05].

As mentioned in [Bro03b], confidentiality is usually associated with the encryption of data only, however there are other aspects to be considered for the case of Grids. The use of Grids implies that confidential data is stored in online accessible databases. Access to their interfaces must be carefully controlled, both to allow access only to appropriate users, and also to allow queries and simulations to run over these highly confidential data without that data being compromised or revealed. If the database is to be shared in a Grid, it might need to be operated by a trusted third party. A further novelty of Grid applications is that they may entail running confidential code or using confidential data on a remote resource; running a job on a dynamically-selected cluster according to load may be good resource management, but the data owner may know nothing about the trust status of the cluster selected by the grid software.

3.2.4 Privacy

Confidentiality also extends to the privacy requirements of the actual users and resources. Users are protected under privacy laws and these must be adhered by all components of proposed Grid technology.

Privacy-Enhancing Technologies (PET) have been defined as *"a coherent system of ICT measures that protects privacy by eliminating or reducing personal data or by preventing unnecessary and/or undesired processing of personal data; all without losing the functionality of the data system"* [Bor01]. PET could increase privacy of Grid applications. Adaption and implementation of these technologies for Grids are areas of open research.

There are several principles in consideration when building PET: transparency, being aware of what personal data is transmitted/promoted and how it is processed; data minimisation, reduction of processed personal data by anonymity and pseudonymity procedures, minimising the linkability between a person and the personal data; system integration, privacy protection built into the system; user empowering, privacy self-protection for users; and multilateral security, realisation that only minimal trust in other parties is required.

Examples of existing PETs include: communication anonymisers hiding the real online identity (email address, IP address, etc.) and replacing it with a non-traceable identity (disposable / one-time email address, random IP address of hosts participating in an anonymising network, pseudonym, etc.); wallets of multiple virtual identities; which allow the efficient and easy creation, management and usage of unlinkable virtual identities; and, anonymous credentials, asserted properties/attributes or rights of the holder of the credential that don't reveal the real identity of the holder and that only reveal so much information as the holder of the credential is willing to disclose.

The W3C Platform for Privacy Preferences Project (P3P) specification describes the encoding of privacy policies into machine-readable XML [Cra06], allowing automated processes to read such policies and take actions on them. Using machine-readable preference language such as APPEL [Cra02], users can express personal preferences over all aspects of privacy policies and have automated processes judge the acceptability of any such policy, or prompt for a decision instead. Since it might be cumbersome to manually create such preferences from scratch, a trusted third party (e.g., a consumer interest group) could provide pre-configured preference specifications that would then be downloaded and individually adjusted by each user. Current work of the W3C Policy Language Interest Group (PLING) aims at integrating different policies languages such as P3P, OASIS XACML and IETF Common Policy.

3.3 Emerging Trust and Security Technologies

Service-oriented architectures provide the shared organising principles that underpin the collaborative operation of services in open dynamic distributed systems. In this section we review the main Web Services Security standards, proposed by standardisation bodies such as W3C and OASIS. Then, we review OGSA Security Model.

3.3.1 Web Services Security

Web services offer an interoperable framework for stateless, message-based and loosely coupled interaction between software entities. These entities can be spread across different companies and organisations, can be implemented on different platforms, and can reside in different computing infrastructures. Web services expose functionality via XML messages, which are exchanged through the SOAP protocol. The interface of a Web service is described in detail in an XML document using the "Web Service Description Language" (WSDL).

In order to provide security, reliability, transaction abilities and other features, additional specifications exist on top of the XML/SOAP stack. The creation of the specifications is a cross-industry effort, with the participation of standardisation bodies such as W3C and OASIS. A key element in the Web services specifications is the so-called combinability. Web services specifications are being created in such a way that they are mostly independent of each other, however they can be combined to achieve more powerful and complex solutions. In this section we describe some individual specifications, specifically focusing on those dealing with secure and reliable transactions. A complete description of the specifications and its usage is presented in [Geu05].

- **Reliability**

The WS-ReliableMessaging specification describes a protocol for reliable delivery of SOAP messages in the presence of system or network failures. To do so, the initial sender retrieves a unique sequence identifier from the ultimate receiver of the sequence to be sent. Each message in the sequence is uniquely bound to that identifier, together with a sequence number. The receiver of the sequence acknowledges the sender what messages have already been received, thus enabling the sender to determine based on the sequence number which messages

have to be retransmitted. WS-ReliableMessaging should be used in conjunction with WS-Security, WS-SecureConversation and WS-Trust in order to provide security against attackers at the network layer.

- **Policies**

The Web Services Policy Framework, WS-Policy, provides a general-purpose model to describe web service related policies. WS-Policy by itself only provides a framework to describe logical relationships between policy assertions, without specifying any assertion. WS-PolicyAttachment attaches policies to different subjects. A policy can be attached to an XML element by embedding the policy itself or a link to the policy inside the element or by linking from the policy to the subject that is described by the policy. WS-PolicyAttachment also defines how policies can be referenced from WSDL documents and how policies can be attached to UDDI entities and stored inside a UDDI repository. WS-MetadataExchange defines protocols to retrieve metadata associated with a particular web services endpoint. For example, a WS-Policy document can be retrieved from a SOAP node using WS-Metadata. WS-PolicyAssertions specifies some common WS-Policy assertions, related to text encoding, required SOAP protocol version and so-called 'MessagePredicate' assertions that can be used to enforce that a particular header combination exists in a given SOAP message.

- **Security**

WS-SecurityPolicy defines certain security-related assertions that fit into the WS-Policy framework. These assertions are utilised by WS-Security, WS-Trust and WS-SecureConversation. Integrity and confidentiality assertions identify the message parts that have to be protected and it defines what algorithms are permitted. For instance, the 'SecurityToken' assertion tells a requestor what security tokens are required to call a given Web service. Visibility assertions identify what particular message parts have to remain unencrypted in order to let SOAP nodes along the message path being able to operate on these parts. The 'MessageAge' assertion enables entities to constrain after what time a message is to be treated as expired.

The WS-Security specification defines mechanisms for integrity and confidentiality protection, and data origin authentication for SOAP messages and selected parts thereof. The cryptographic mechanisms are utilized by describing how XML Signature and XML Encryption are applied to parts of a SOAP message. That includes processing rules so that a SOAP node (intermediaries and ultimate receivers) can determine the order in which parts of the message have to be validated or decrypted. These cryptographic properties are described using a specific header field, the `<wsse:Security>` header. This header provides a mechanism for attaching security-related information to a SOAP message, whereas multiple `<wsse:Security>` header may exist inside a single message. Each of these headers is intended for consumption by a different SOAP intermediary. This property enables intermediaries to encrypt or decrypt specific parts of a message before forwarding it or enforces that certain parts of the message must be validated before the message is processed further.

Besides the cryptographic processing rules for handling a message, WS-Security defines a generic mechanism for associating security tokens with the message. Tokens generally are either identification or cryptographic material or it may be expressions of capabilities (e.g. signed authorization statements). WS-Security 1.0 does only define a simple user name token, a container for arbitrary binary tokens (base64 encoded) and a container for XML-formatted tokens. Additional specifications define various 'token profiles' that introduce special token formats. For instance, the 'WS-Security X.509 Certificate Token profile' defines how X.509 certificates, certificate chains or PKCS#7 certificate revocation lists may be used in conjunction with WS-Security.

The WS-Trust specification introduces the concept of 'security token services' (STS). A security token service is a Web service that can issue and validate security tokens. For instance, a Kerberos ticket granting server would be an STS in the non-XML world. A security token service offers functionality to issue new security tokens, to re-new existing tokens that are expiring and to check the validity of existing tokens. Additionally, a security token service can convert one security token into a different security token, thus brokering trust between two trust domains. WS-Trust defines protocols including challenge-and-response protocols to obtain the requested security tokens, thus enabling the mitigation of man-in-the-middle and message replay attacks. The WS-Trust specification also permits that a requestor may need a security token to implement some delegation of rights to a third party. For instance, a requestor could request an authorization token for a colleague that may be valid for a given time interval.

WS-Trust utilises WS-Security for signing and encrypting parts of SOAP messages as well as WS-Policy/SecurityPolicy to express and determine what particular security tokens may be consumed by a given Web service. WS-Trust

is a basic building block that can be used to rebuild many of the already existing security protocols and make them fit directly in the web services world by using Web service protocols and data structures.

WS-Federation introduces mechanisms to manage and broker trust relationships in a heterogeneous and federated environment. This includes support for federated identities, attributes and pseudonyms. 'Federation' refers to the concept that two or more security domains agree to interact with each other, specifically letting users of the other security domain accessing services in the own security domain. For instance, two companies that have a collaboration agreement may decide that employees from the other company may invoke specific web services. These scenarios with access across security boundaries are called 'federated environments' or 'federations'. Each security domain has its own security token service(s), and each service inside these domains may have individual security policies. WS-Federation uses the WS-Security, WS-SecurityPolicy and WS-Trust specifications to specify scenarios to allow requesters from the one domain to obtain security tokens in the other domain, thus subsequently getting access to the services in the other domain.

- **Web Services Specification in Implementing the VO Lifecycle**

Some of the requirements presented in the analysis of requirement through the VO lifecycle can be met by application of Web services specification, as shown in [Are05].

The identification phase includes defining VO wide policies as well as selecting potential business partners who are both capable of providing the required services and of fulfilling the trustworthiness requirements of the VO. The selection of potential business partners involves looking at repositories, which can realize. The usual Web service technology to be applied is WSDL/UDDI, WSDL describes messages and operations while UDDI offers a discovery mechanism. To include the provision of SLA, "Web Service Level Agreements" (WSLA) has been developed, a XML language for specifying and monitoring SLA for Web Services, which is complementary to WSDL. Determining the required service providers and a proper negotiation requires secure communication. The WS-Security specification and data origin authentication for SOAP messages can be used between the entities to secure the communication.

The realisation of the VO requires the creation of federations, where two or more security domains agree to interact with each other, specifically letting users of the other security domain accessing services in the own security domain. The WS-Federation specification deals with federations by providing mechanism to manage and broker trust relationships in a heterogeneous and federated environment. This includes making use of WS-Trust to support for federated identities, attributes and pseudonyms. The dissemination of configuration information requires secure communication as provided by the WS-Security specification.

Throughout the operation of the VO, service performance will be monitored. This will be used as evidence when constructing the reputation of the service providers. Any violation -e.g. an unauthorised access detected by the access control systems- and security threats -e.g. an event detected by an intrusion detection system- need to be notified to other members in order to take appropriate actions. VO members will also need to enforce security at their local site. For example, providing access to services and adapting to changes and the violations. Monitoring can be supported by event management and notification mechanisms using the WS-Eventing and WS-Notification specifications. This allows the monitoring service partner to receive messages when events occur in other partners. A mechanism for registering interest is needed because the set of Web services interested in receiving such messages is often unknown in advance or will change over time.

3.3.2 OGSA Security

To address the Grid specific security requirements of OGSA, the OGSA Security Group has proposed an architecture leveraging as much as possible from the Web Services Security specifications [Nag03].

As we mentioned previously, secure operation in a Grid environment requires that applications and services be able to support a variety of security functionalities, such as authentication, authorization, credential conversion, auditing and delegation. These functionalities are based on mechanisms that may evolve over time as new devices are developed or policies change. As suggested in [Sie03], Grid applications must avoid embedding security mechanisms statically.

Exposing security functionalities as services (i.e., with a WSDL definition) achieves a level of abstraction that helps provide an integrated, secure Grid environment. An OGSA infrastructure may use a set of primitive security functions in the form of services themselves. [Nag03] suggest the following security services:

- **Authentication Service**

An authentication service is concerned with verifying proof of an asserted identity. One example is the evaluation of a User ID and password combination, in which a service requestor supplies the appropriate password for an asserted user ID. Another example involves a service requestor authenticating through a Kerberos mechanism, and a ticket being passed to the service provider's hosting environment, which determines the authenticity of the ticket before the service is instantiated.

- **Identity Mapping Service**

The identity mapping service provides the capability of transforming an identity that exists in one identity domain into an identity within another identity domain. The identity mapping service is not concerned with the authentication of the service requestor; rather it is strictly a policy driven name mapping service. Authorization service: The authorization service is concerned with resolving a policy based access control decision. The authorization service consumes as input a credential that embodies the identity of an authenticated service requestor and for the resource that the service requestor requests, resolves based on policy, whether or not the service requestor is authorized to access the resource. It is expected that the hosting environment for OGSA compliant services will provide access control functions, and it is appropriate to further expose an abstract authorization service depending on the granularity of the access control policy that is being enforced.

- **VO Policy Service**

The VO policy service is concerned with the management of policies. The aggregation of the policies contained within and managed by the policy service comprises a VO's policy set. The policy service may be thought of as another primitive service, which is used by the authorization, audit, identity mapping and other services as needed.

- **Credential Conversion Service**

The credential conversion service provides credential conversion between one type of credential to another type or form of credential. This may include such tasks as reconciling group membership, privileges, attributes and assertions associated with entities (service requestors and service providers). For example, the credential conversion service may convert a Kerberos credential to a form that is required by the authorization service. The policy driven credential conversion service facilitates the interoperability of differing credential types, which may be consumed by services. It is expected that the credential conversion service would use the identity mapping service. WS-Trust defines such a service.

- **Audit Service**

The audit service similarly to the identity mapping and authorization services is policy driven. The audit service is responsible for producing records, which track security relevant events. The resulting audit records may be reduced and examined as to determine if the desired security policy is being enforced. Auditing and subsequently reduction tooling are used by the security administrators within a VO to determine the VO's adherence to the stated access control and authentication policies.

- **Profile Service**

The profile service is concerned with managing service requestor's preferences and data which may not be directly consumed by the authorization service. This may be service requestor specific personalization data, which for example can be used to tailor or customize the service requestor's experience (if incorporated into an application which interfaces with end-users.) It is expected that primarily this data will be used by applications that interface with a person.

- **Privacy Service**

The privacy service is primarily concerned with the policy driven classification of personally identifiable information (PII). Service providers and service requestors may store personally identifiable information using the Privacy Service. Such a service can be used to articulate and enforce a VO's privacy policy.

3.3.3 Grid Security Infrastructure

The Grid Security Infrastructure (GSI) is a specific implementation of an OGSA-based Grid security architecture that include as part of the Globus Toolkit Version 3 (GT3) [Wel03]. Given the prominent use of Globus within the Grid community, let us briefly revise such implementation.

- **Authentication**

GSI defines a credential format based on X.509 identity certification. An X.509 certificate, in conjunction with an associated private key, forms a unique credential set that a Grid entity (requestor or service provider) uses to authenticate itself to other Grid entities (e.g., through a challenge-response protocol such as TLS).

- **Identity Federation**

GSI uses gateways to translate between X.509-based identity credentials and other mechanisms. For example, the Kerberos Certificate Authority (CKA) and SSLK5/PKNIT provide translation from Kerberos to GSI and vice versa, respectively. These mechanisms allow a site with an existing Kerberos infrastructure to convert credentials between Kerberos and GSI as needed.

- **Dynamic Entities and Delegation**

GSI introduces X.509 proxy certificates, an extension to X.509 identity certificates that allows a user to assign dynamically a new X.509 identity to an entity and then delegate some subset of their rights to that identity.

- **Message Level Security**

Globus Toolkit Version 3 (GT3) uses the Web Services Security specifications to allow security messages and secured messages to be transported, understood and manipulated by standard Web services tools and software.

In relation to stateful and secured communication, GSI supports the establishment of a security context that authenticates two parties to each other and allows for the exchange of secured messages between the two parties. GT3 achieves security context establishment by implementing preliminary versions of WS-SecurityConversation and WS-Trust specifications. Once the security context is established, GSI implements message protection using the Web Services standards for secured messages XML-Signature and XML-Encryption.

To allow for communication without the initial establishment of a security context, GT3 offers the ability to sign messages independent of any established security context, by using XML-Signature specification.

- **Trust Domains**

The requirement for overlaid trust domains to establish VOs is satisfied by using both proxy certificates and security services such as CAS. GSI has an implicit policy that any two entities bearing proxy certificates issued by the same user will inherently trust each other. This policy allows users to create trust domains dynamically by issuing proxy certificates to any services that they want to interoperate.

4 Trust and Security in CoreGRID

This section shows how Grid trust and security technologies are impacting the different research areas studied in CoreGRID.

4.1 Trust and Security in the Institute of Knowledge and Data Management

Trust and security are fundamental aspects of work in this Institute, as knowledge and data management are at the core of storing and retrieving information in Grids as well as building new services on top of stored information. The Institute follows a layered approach in addressing data and knowledge related issues in Grid systems.

- The lowest layer (Task 2.1) deals with systems-level, distributed storage management issues.
- The middle layer (Task 2.2) explores techniques that will turn storage systems into knowledge representation systems.
- Finally, the top-most layer (Task 2.3) addresses issues in automatic mining and resource discovery techniques.

Each of these tasks has several objectives as outlined in [Bil05]. Most of these objectives have a component that is related to trust and security. More specifically, for each task, the objectives that are related to trust and security are defined below.

4.1.1 Distributed Data Management

Future storage systems will contain critical user information for various applications and purposes (e.g. health, financial). In fact, more and more information is becoming available in digital form with the goal of keeping all available content online. How to guarantee that storage systems are not compromised and used for other purposes, or information does not leak, is modified or even destroyed without users consent, are main issues in building and managing all available storage. The first part of our research on Grid storage system's security consists on analyzing the security of state-of-the-art technologies using a framework originally proposed for generic storage systems which we have extended to Grid-specific configurations. The results of this analysis have been published in [Lun07]. Another area of research related to security within this task is the definition of security requirements for distributed data management systems, as reported in [Naq07, Naq06].

The main scientific challenges in distributed data management include dealing with data security and privacy at the different Grid layers (starting from the code security, secure network technologies, finishing at the safe rooms for the storage devices) as a way to provide a comprehensive solution able to deal not only with the security gaps that must be covered, but also to optimize existing mechanisms that might be providing redundant security services (i.e. authorization decisions taken by both, local storage systems and Virtual Organizations). Security and privacy issues are vital features in wide area distributed systems. Basic Grid functionality (e.g., Globus security infrastructure - GSI) must be exploited to support secure client-server interactions without impacting on the usability and performance of the Grid infrastructure and services. For instance, there are currently efforts by WP2 partners in building "National Data Stores" which low-level security can be based the following two approaches:

- Encryption: Secure data at rest, but managing the encryption keys in such a way that untrusted storage elements and/or compromised clients can not decrypt them (i.e. when colluded). Performance is also an issue in this approach.
- Trust level: This is mostly a preventive approach, where each storage element is identified by an advertised security policy that can be quantitatively associated to a security level. Clients are provisioned of storage resources based on a security criterion that considers this level.

Thus, all three objectives of task 2.1 have a dimension related to trust and security:

- Storage Infrastructure: Studying the replacement of existing high-end scalable storage systems with commodity physical storage devices, controllers, and interconnects within Grids and examining how current storage systems can migrate to this new architecture.
- Providing Management Mechanisms: Providing techniques for automatically managing storage resources in the Grid and providing "high-quality" storage at low cost to users.
- Specifying Management Policies: Examining the different classes of storage services that could/should be offered to users and description methods and techniques for specifying service classes and management policies.

4.1.2 Information and Knowledge Management

A number of aspects of this task are related to trust and security. For instance, from the Semantic Web point of view, it is necessary to develop semantics for privacy, security and access-rights as well as dealing with dynamic information, state, QoS and states. As another example, from the Grid point of view, it is important to move from fixed-pipeline processes to dynamic compositions. In forming Virtual Organisations, security management must become autonomic and adaptation must occur automatically in real-time, rather than through human intervention. Furthermore, autonomic security management will have to be complemented by extensible and machine processable standards for negotiating, validating and amending collaboration agreements, encoded by means of electronic contracts, which can be autonomically enacted by the platform. Such extensible and machine processable standards require the development of common vocabularies and negotiation protocols.

The objectives of Task 2.2 that are related to trust and security are:

- Semantic Modelling: Developing metadata for Grid service discovery and information management and the design of knowledge-oriented Grid services

- **Semantic Representation:** Exploiting the use of Semantic Web technologies for sharing machine readable Semantic Grid models and techniques for knowledge intensive applications
- **Agent Infrastructure:** Analyzing the use of agent technologies to exploit semantic representation of users and resources to support workflow and knowledge management across distributed virtual organizations in science and business.
- **Standardization and Integration:** Extending and standardizing the existing OGSA middleware for knowledge-based Grid services.

4.1.3 Data Mining and Knowledge Discovery

The dimensions of this task that are related to trust and security are:

- **Policy Publication and Enforcement:** Service providers will publish policies for their use, detailing the obligations, privileges and expected levels of service, which a user should accept before using the service. Some initial efforts in the use of Semantic Web representations for basic security applications (authentication, access control, data integrity, and encryption) have begun to bear fruit. For example, Denker et al. [Den03] have integrated a set of ontologies (credentials, security mechanisms) and security extensions for Web Service profiles with the CMU Semantic Matchmaker. Kagal et al. [Kag03] are also developing Rei, a Semantic Web based policy language. Furthermore, KAoS services and tools allow for the specification, management, conflict resolution, and enforcement of policies within the specific contexts established by complex organizational structures represented as domains [Bra03a]. A comparison of KAoS, Rei, and more traditional policy approaches such as Ponder can be found in [Ton03]. KAoS provides a powerful tool-set that appears to be capable to address publication and deployment of complex policies for Semantic Web Services. However the incorporation of trust metrics and a distributed enforcement and performance assessment scheme remain the main challenges, in addition to the production of a critical mass of domain/application-specific ontologies to allow its uptake and validation in large scale systems. With respect to the latter there is an ongoing effort to adapt KAoS for use in Grid Computing environments in conjunction to OGSA [Joh03a].
- **Monitoring and Policy Enforcement:** During the execution of the service, which may be over a long period, its progress is monitored. The experience of the quality of the service may modify the relationship between the parties. For instance, if usage is long-lived the experience of the parties during the interaction may modify their behavior for its remainder; Good experience may result in the loosening of restrictions and a higher-level of trust, changing the valuations in internal "trustbases", and reducing the policy enforcement overhead.

Thus, the high-level objectives of Task 2.3 that are related to trust and security are:

- **Semantic Mapping:** Studies about the representation and mining of relationships between different Grid entities and resources,
- **Distributed Grid Services:** Design of services, and tools for distributed data mining and knowledge discovery on Grids, with Grid-aware highly adaptive data mining algorithms, considering data integrity and privacy.
- **Monitoring services:** Services providing accurate estimates of the cost of data mining tasks on Grids. Knowledge-based OGSA Grid services - how knowledge-discovery and knowledge-based services can be implemented by using the OGSA model?

A recent line of research is the development of security frameworks for data curation. A security framework for data curation should be able to preserve data as well as metadata, security of the data (policies) and data integration over time and technological changes.

Overall, with increasing amount of information being stored and managed in digital form and with more services being deployed that require access to private and personal information, there is increasing concerns about loss of privacy and wrongful use of information. These concerns affect all tasks (layers) of WP2.

4.2 Trust and Security in the Institute on Programming Models

Security issues have to be considered a primary concern when investigating suitable programming models for the grid. In particular, typical grid configurations including processing elements from different virtual organizations and interconnected by means of "public" network segments should be addressed in such a way data and code confidentiality, their integrity as well as proper user authentication can be guaranteed.

When using existing, traditional programming tools to implement grid applications, the "securing duty" is usually completely the responsibility of the application programmers. In addition to the burden related to programming the functional aspects of the application (i.e. what the application computes) and the non functional aspects related to distribution and performance (i.e. how the application is computed), the user/programmers has to deal with choosing proper security mechanisms as well as security policy implementations. As a result, the application functional code may end up completely intermingled with the non functional code, in particular with the non functional code dealing with security aspects. This code is difficult to debug, fine tune and maintain. In the perspective of "raising the level of abstraction" of programming models, as predicated in the third report of the Next Generation Grids Expert Group document⁷ security concerns should instead be included in proper ways in the tools provided to grid application programmers to implement their applications.

Following this hint, the approach taken in the Programming Model Institute to deal with security of grid applications is a high level approach. As detailed below, security issues are mainly investigated in the context of Task 3.3 "Advanced component models" to verify how security issues are handled directly in the programming environment tools (static and dynamic) following qualitative user hints mainly identifying sensible data and code. The investigation of security aspects was formerly discussed with respect to a lightweight security research group. Actually, the Programming Model Institute decided at the end of the second year of activity to maintain the original, internal division in three tasks rather than to split further the activities in research groups. As a consequence, the lightweight security research group was terminated and the activities related to security were subsumed under Task 3.3.

4.2.1 Advanced Component Models

The main problems related to security in the framework of GCM (the Grid Component Model currently under development at the Programming Model Institute) are secure component deployment, on the one side, and secure component access and data transfer, on the other side. Both require two specific actions: i) proper mechanism choice (i.e. identifying the more suitable mechanisms to be used among those available) and ii) programming proper security policies according to the mechanisms available/chosen (i.e. inserting into the functional code the non functional code needed to deal with security policy implementation). The approach taken exploits GCM composite components as well as GCM component autonomic managers as follows:

- High level components are provided that model common grid programming patterns, in the style of what already happens in HOC and ASSIST, respectively from WWU Muenster and UNIPI
- Programmers are required to annotate the code and data parameters used to instantiate the high level components with security "hints/requirements" (e.g. which data/code has to kept confidential)
- Autonomic managers in the composite components take care of performing (sub)component deployment or component invocation using proper secure mechanisms. This includes dynamically customizing cryptographic tools and options to optimize computation performance without compromising specified security requirements.

Following this approach activities have been set up and are currently carried out by Programming Model Institute partners related to different research topics, and in particular:

- Evaluation of the impact of introducing security mechanism in the code of typical parallel grid applications (responsible partner UNIPI)
- Analysis of the available security mechanisms, tools and policies required to implement the desired security goals. In particular, by identifying concrete but more relaxed security requirements for common scenarios which are satisfiable by lightweight cryptographic tools, and designing efficient tools to distribute useful computation units while preserving meaningful privacy requirements (responsible partner UCHILE)

⁷<http://cordis.europa.eu/ist/grids/ngg.htm>

- Design of semi formal techniques to support proper exploitation of user supplied security information (metadata denoting sensible data and code) (responsible partner QUB)
- Design of proper security policies in the GCM component managers suitable to secure component deployment (code confidentiality) as well as component invocation (data confidentiality) (responsible partner UNIPI)

Currently, these activities already led to some publications. In particular, [Kil07, Ald07a] shortly outline how security related metadata provided by the programmer can be exploited to optimize the security mechanisms usage in grid applications, and the paper [Ald07b] carefully evaluates the cost of securing communications in a task/data parallel computation patterns on the grid.

4.3 Trust and Security in the Institute on Architectural Issues: Scalability, Adaptability, Dependability

The scale, dynamism, and openness of the Grid, together with demands on trust and security, on reliability, and on manageability, poses new, unique architectural challenges. The CoreGRID Institute on Architectural Issues performs a significant improvement of architectural designs of future Grids by focusing on three key aspects: scalability, adaptability, and dependability. Each of these key aspects raises various trust and security issues:

- Scalable Grid Services and Resource Discovery. A fully decentralized model for Grid Architecture should be deployed in order to be able to manage Grids composed of thousands of nodes. As the scalability increases security becomes an even more critical issue, since tools should be designed that automatically react to security threats in resource access and service provision.
- Adaptive Management of Systems and Resources. Adaptability for resource management should provide mechanisms for automated adaptation and reconfiguration of the Grid on all levels of the hierarchy. Adaptation mechanisms that involve resource sharing and process migration should pay special attention to trust and security.
- Dependability and Fault-tolerance in Grids. Grid middleware should be instrumented with fault-tolerance techniques in order to assure the resiliency of applications and the high-availability of crucial Grid services. The principle of the dependable Grid should involve security mechanisms.

For Grid systems to be used for business applications, it is necessary to share the resources between unknown, un-trusted parties. As the scope of Grid enlarges to ubiquitous and pervasive computing, there is need to assess and maintain the reputation of entities. Reputation is one of the tools the research community will have to supply if the use of Grid expands beyond the organization boundaries. In [Sil07a] the main reputation-based trust management systems and their applicability to Grid systems are reviewed.

More specifically, the objectives of the CoreGRID Institute on Architectural Issues that are related to trust and security are the following:

- **Scalable Grid Services and Resource Discovery**

As Grid architectures keep on growing, middleware tools will have to adapt to the large number of resources. Currently, tools providing Grid resources access do not address as much as they should trust and security issues. These issues are of vital importance on large scale architectures. As Grid services could be potentially used by unauthorized users, methods should be devised so that the system as a whole is not compromised and mechanisms have to be provided to automatically react to these events. Grid services themselves should adapt to the larger scale and must offer valid support to systems to cope with user authentication, security and privacy of data.

Currently, much of the Grid research has been concerned with providing security by building a wall around the Grid, and by trying to keep malicious nodes outside. But as the Grid scales, this will be insufficient by itself. It is not safe to assume that nodes already within the walls will never be compromised. An interesting axis of research in this direction is to deploy gossiping techniques which disseminate information in a decentralized and scalable fashion whenever a malicious action is detected. The issue of scalability needs to be considered at the level of Grid services.

- **Adaptive Management of Systems and Resources**

There are several aspects that need to be addressed to make Grid systems adaptable to internal and external changes. Adaptability for resource management should provide mechanisms for automated adaptation and reconfiguration of the Grids on all hierarchy levels.

One of the main targets of this task is creation of automated tools based on atomic actions that allows an application to adapt its workload based on the available resources. This is directly related to resource sharing and process migration and raises several issues related to trust and security. Moreover, access to some information can be restricted due to security. The resource sharing and processes migration tasks make the systems vulnerable to attacks.

The incorporation into the description of the precondition and effects of the atomic actions based on which the workflows result of special mechanisms for anomaly detection (based on the CIM framework) that check the workflows. Prediction and reconfiguration should take into consideration the security levels of the resources.

- **Dependability and Fault-tolerance in Grids**

The issue of Grid dependability is closely connected to trust and security. Grid middleware should be instrumented with the support for fault-tolerance techniques to assure the resilience of applications and the high-availability of crucial Grid Services.

An interesting topic of research that merges the fields of fault-tolerance and trust and security is the development of distributed protocols for sabotage tolerance. These protocols are particularly relevant in Global computing environments since the existing schemes still present some restrictions that should be solved. So, there is need for development of more effective techniques for sabotage tolerance and trust in open environments and more scalable protocols for resilient task distribution. Initial work in this area by Institute's member is reported in [Dom07].

Two other axes of research that raise trust and security issues and which there is intensive research going on in this CoreGRID Institute are the following:

- **Peer-to-Peer Systems**

A topic of research for which there is currently intensive collaboration taking place from the partners involved in this Institute is Peer-to-Peer (P2P) computing, and the use of P2P paradigm for Grid resource discovery. For the P2P paradigm to be established as a successful paradigm not only for the sharing of information but also for business oriented applications, it should exhibit significant trust and security requirements. Currently P2P systems are self-organized entities that are not used for critical applications. In order to overcome this drawback, they will have to provide a trustable environment. P2P is investigated as a tool to build reputation lists and to achieve trust management. P2P can be used for massive computation and their notion is closely related to this of Internet desktop Grids.

- **Desktop Grids**

An important axis of research of this CoreGRID Institute is Internet Desktop Grids, an environment that becomes increasingly popular for demanding applications that make use of large amounts of data. Recently P2P techniques for the efficient distribution of large chunks of data have been investigated using either unstructured techniques (BitTorrent) or structured alternatives (Chord DHT). The validation of desktop Grid computing is heavily depended on the replication of computation. In order to avoid a large degree of replication, reputation-based trust management techniques have been proposed in [Sil07b]. The master builds reputation lists of the slaves based on the validity of past results. Reputation information can be incorporated in the result-validation process by using a weighted voting algorithm in order to reduce significantly the replication factor. Using similar trust management systems collusions of slave organizing themselves into P2P systems can be drastically reduced. The P2P paradigm has also been investigated as the means (the tool) to build reputation lists and to achieve trust management.

4.4 Trust and Security in the Institute on Grid Information, Resource and Workflow Monitoring Services

In modern IT infrastructure the Information has to be gathered and provided in secure way and the resource and the workflow management activities have to be reliable. Therefore the IRWM institute is obligated to consider the

adequate security issues. The next four subsections represent the individual tasks' approach to these issues.

4.4.1 Network Monitoring Security Issues

The work of the Network Monitoring task follows the guidelines indicated in its roadmap [Mey06], and aims at the operational definition of the Grid Infrastructure Monitoring Services and the Network Monitoring Element. The infrastructure monitoring part is based on the C-GMA (Capability-based Grid Monitoring Architecture), which is an extension to the GMA (Grid Monitoring Architecture), defined by the Global Grid Forum.

The C-GMA concept requires two levels of security and trust management. The interaction between producers and consumers and the mediator requires an authentication scheme based on trusted third parties (such as X.509 certificates). As the mediator could not restrict new producers or consumers to join and publish data, the certificate based authentication is used just to check the appropriateness (the element does possess a valid certificate verifying its authenticity with respect to trusted certification authorities) but not to impose restrictions in terms of access control (i.e. no strict authorization in a general framework is required). The second level deals with the actual interaction between selected producer and consumer, where the C-GMA concept on purpose does not dictate any explicit security mechanism and deliberately leaves on the partners (the producer and consumer) willing to communicate to select the authentication and authorization method best suited to their need (low to none for non sensitive data, very high and complex authentication and authorization for highly sensitive data). However, the C-GMA metadata layers offer means for specifying authentication and authorization mechanisms (using capabilities and attributes) so that these non-functional requirements may also be a subject to the described matchmaking process.

4.4.2 Grid Checkpointing Architecture Security Issues

The services making up the GCA are encountering the security issues, as there is need to access to the users' job metadata and ensuring security while handling the images of the applications. As the images might be considered as a special files belonging to owner of the application, the GCA components will rely on standard Grid mechanisms managing access to user-owned resources. Encryption of the files might be a feature increasing the security of the files, but by default the GCA trusts the storage services to be secure.

As regards other security issues, we have taken an assumption that the implementation of Grid Checkpointing Architecture will be performed with help of technologies that allow to employ the build-in security mechanisms. The most recent Web Services based technologies (i.e. those available in GT4) ensure privacy, integrity and proper authorization and authentication based on widely accepted world wide standards (X.509 certificates, proxy certificates, digital signs, TSL transport protocol). However, the security issues as such are not within the area of interest of task 5.2. Therefore the services designed within task 5.2 to provide an adequate trust and security level will take advantage of mechanisms offered by other services and the proof-of-concept implementation environment.

4.4.3 Workflow Services Security Issues

The composition and enactment of workflows in a Grid environment gives additional requirements to the Grid security infrastructure, which are mostly not fulfilled in established Grid middleware environments that focus on atomic or parallel job execution, rather than on complex workflow enactment. Conventional Grid security architectures, such as GSI (Globus Toolkit), focus on the service provider's perspective and do not cover all concerns of the service user. For instance the management of user credentials is delegated to services which are not under full control of the user (e.g. MyProxy). Another drawback is that common Grid security systems do not cover the fine-grained authorization of services, taking into account the methods, their parameters, and the message flow in deciding whether a user or another service is authorized to access the service. This is particularly important for Grid workflow systems, which have additional requirements, such as fine-grained and role-based security mechanisms in combination with restricted delegation of privileges, as presented in [Hoh06].

The security issues identified within the research groups of the WP5.3 are as follows:

- **Workflow description languages using high-level Petri nets for GRID workflows**

The goal of this workgroup is to develop a platform-independent workflow language and platform-specific tools.

The tools developed in this research group have to be able to take into account and make use of the trust and security mechanisms provided by the underlying Grid platforms, such as e.g. the widely used Grid Security

Infrastructure. In order to achieve this goal, workflow managers must implement the security requirements of these infrastructures.

- **Fault tolerance in Grid workflow execution**

The research group is using the EGEE and UK NGS production Grid environments for its work. Both of them are using the Grid Security Infrastructure (GSI). The research group does not foresee any demand for stronger security or trust management mechanisms.

- **Workflow-oriented Grid infrastructure for biomedical purposes**

Medical communities raise serious concerns about the security solutions of current production Grids. As it has already been expressed by the biomedical user community of the EGEE Grid, the Grid Security Infrastructure - the security infrastructure used by all the main production Grids worldwide - does not provide an acceptable level of security and flexibility for several application areas. Consequently, the EGEE developer community is already working on a set of new solutions to fulfil the needs of such communities. Because these systems are in a prototype (or even more initial) stage, our research group is only following their status but has not started their adoption. As soon as they are available, our group will adapt and integrate them into the targeted medical workflow environment.

4.4.4 Accounting and User Management Services security issues

The main aim of user account management systems is to provide controlled and secure access to grid resources. Security requires authentication of the user and authorization based on combined security policy from the resource provider and the virtual organization of the user. Users must be allowed to use the resources to the extent allowed by the user roles and the policy, while resources must be secured against unintentional as well as malicious policy breaks. The issue of authentication is well addressed by existing standards and solutions. There are some solutions in authorization area, but the subject is still being investigated also within this task.

The second important thing is the possibility of logging user activities for accounting and auditing (security reasons) and then gathering these data or their aggregates both by the resource provider and virtual organization of the user. The logging features are closely related to one of the main problems of user management: mapping the global user id to the local one, because they require identification of the user who performed some action. So that, we try to propose a solution that simplify the user management and will be scalable, but still allow for user identification. The access to the accounting and logging data must be properly limited depending on the users' roles: consumers of the resources, administrators of resources and managers of the virtual organizations. Thus, the accounting services must be secured and the access to the data must be controlled.

In the proposed framework for Virtual Environments authentication and confidentiality issues are addressed by the existing Globus Grid Security Infrastructure. Fine-grained and flexible authorization is achieved by the Globus Toolkit v.4 authorization framework and a set of authorization plugins, implementing different authorization methods and reusing existing services like e.g. VOMS. Authorization enforcement and job isolation is assured by mapping the user to properly configured VE. Control-over-user actions (audit and accounting) are possible thanks to the VE Database which contains data from the local logs stored in the context of the global user identity and his/her Virtual Organization. Results achieved by the institute in this topic include [Den06].

4.5 Trust and Security in the Institute on Resource Management and Scheduling

As it was mentioned in previous sections, Trust and Security are one of the main challenges in open network grid environments. In case when the Resource Management and Scheduling are analysed in the grid environment, Trust and Security are very important issues. Definition of security requirements has to be considered from two perspectives:

- from the internal perspective, how particular institute areas/tasks fit to Trust and Security;
- from the whole system point of view, how security systems can protect some of the subsystems developed within the Institute on Resource Management and Scheduling.

Main effort in this chapter will be put on the internal security issues within the institute. There are several areas/tasks defined for the Resource Management and Scheduling where security aspects have to be discussed: analyzing scheduling architectures; multi-Level scheduling strategies; workflow scheduling strategies; evaluation and

benchmarking scheduling system; mapping and scheduling of HPC applications; coordination of GRID scheduling and data management; and performance prediction. An example of the results achieved by the Institute in relation to trust and security is [Far07], which extends UNICORE authorisation capabilities.

- **Analyzing Scheduling Architectures**

For all analysis coming from this task, it would be crucial to consider it from a security perspective. Probably most significant issues for this task are collaboration between services, access rights and security infrastructures. It is important to review those aspects for all architecture components because one gap can make the whole system insecure. Besides security aspects ensuring privacy is an important concern especially for commercial environments when analyzing scheduling architectures. Up to now, most scheduling architectures provide access to user data like the user names, the number of submitted jobs or their runtime. All this data can be used to build user profiles or to predict the current status of products (e.g. short simulations may indicate parameter variations at the beginning of the development of a new product whereas longer simulation times may indicate final evaluations before releasing the product).

Due to the growing demand for virtualization technologies and native support for security mechanisms in recent processors, scheduling architectures should support virtualization and hardware security solutions. Providing support for trusted hardware or trusted virtual machines like AMD's virtualization technology in conjunction with Trusted Platform Modules could increase the trust in grids especially for security aware applications. By using trusted hardware and software scheduling architectures can be certified and their compliance of well defined security policies can be remotely verified. Therefore, security and privacy aspects have to be taken into account when the common CoreGrid scheduling architecture will be designed. It is suggested to elaborate the architecture together with a security model.

- **Multi-Level Scheduling Strategies**

T & S in the 'Multi-Level Scheduling Strategies' is important, starting from level of collaborations between services and resources and finishing on access rights to files and processes on local operation system. This area has identified some security needs such as access to resources or service-quality - one of the criteria can be level of trust. During 'Multi-Level Scheduling Strategies' analysis, it is important to check security aspects of system. It is a need to check the level of security that is offered by service providers. This measure has to be established based on service concepts, technologies and used systems and components.

Running multiple jobs on the local system provide special requirements to such system. The main need concerns on problem of privacy of user's data (tasks and jobs have to run in separate space, on the specific account and have to be fully cleaned after computations). Many local operation systems provide extended security mechanisms, which additionally protect the user data against attackers. It is good manner to use such extensions, which increase level of security of the local system.

- **Workflow Scheduling Strategies**

This task is connected to T & S at abstract level. It is important for the whole system to have security not only at a low level (hardware, single resources) but at abstract layers too. For this task, there is a need to define the security policies for workflow. The need to use an authorization service, where such security policy (e.g. access rights to workflows or parts of them) can be stored, may arise. The security policy has to fulfill possible user security requirements from one side and service requirements from the other side. Based on these requirements coherent policy for workflow has to be delivered. The security policy language has to be chosen, which allows specifying all requirements in logical and distinct form. Further, the integration between different mechanisms like the workflows mechanism and an existing Grid portal solution, can make both above systems unsafe. Such integration needs to be well prepared in the context of security.

- **Evaluation and Benchmarking Scheduling System**

The benchmarking models can be created based on the security requirements, which come from other tasks. Since today's grid infrastructure utilizes virtualization technologies, they have to be taken into account for benchmarking models. Beside memory footprints and available processing power the overhead for trusted virtual machines has to be measured. Especially secure bootstrapping using a secure loader and additional memory and software verification mechanisms may decrease the performance of virtual machines. Once the requirements are well defined, they can be used to create benchmarking models.

- **Mapping and Scheduling of HPC Applications**

The needs presented in this task: "Fault-tolerance and reliability are important subject for HPC scenarios" are quite connected to T & S. Generally speaking, task 6.5 seems to be one of the most important task for the CoreGrid. This task collects other tasks' results. So the security aspects like services collaboration, access rights, quality of service etc. become needs in the term of T & S. Additionally, problem of "...corruption (...) due to benign or malicious attacks" is presented in this task. One of solution for these problems can be "Certification of global computations" like it is mentioned in the task description.

Before starting computations, each HPC application has to be analysed in context of security aspects. Analysis needs to fulfill both user and services requirements. User requirements can be connected to user data protection problems, services requirements point to applications problems like testing if the application code, which is running on service/resource, is secure enough. From service/resource perspective, it will be good to define "quality of application" measure, which tells how good application is secured from service/resource perspective; although, definition of such measure can be very complicated. These two requirements can be conflicted in many cases. Good practice could be need to define some general rules of the above analysis which will be accepted by all sides.

- **Coordination of GRID Scheduling and Data Management**

The coordination between scheduling and the data management concerns the service collaboration and access rights checking. As it is mentioned in the description of this task: "job scheduling will need to communicate with data and network management" and in such case "the derivation of necessary information on data existence, location and access rights will be an issue for examination". Therefore, the need to verify access rights (maybe via an authorization service) and secure communications between services and resources are important parts of this task. Access rights could not be treated as a simple problem. From single user point of view, it is enough to define access control list, which define user rights to specific data. In real use, defining security policy is a need for collaboration group of people. In the real world, group contains several people responsible for different areas of work. It has to be possible to define access rights for them (for example: in context of role in group). The security policy rules define interrelation between members of group.

- **Performance Prediction**

Performance prediction models and strategies can be based on a data coming from previous computations. This task may define a need of cooperation with an accounting system (such system may cooperate with an authorization service - there are many security aspects concerning accounting issue). For the performance prediction, collaboration between services and quality of services are important in the context of T & S. In addition, problem of accessing private data occurs and need to be analysed in this context. Data coming from the accounting system describes how users use system and based on it, presents their preferences. It has to be established if users are agreed to use their data in performance prediction process.

As a conclusion, all above tasks touch several security problems:

- Problem of privacy of personal and company data
- Quality of Service measure from the security perspective
- Quality of Application measure from the security perspective
- Secure communication between entities of system (users, services, resources, others)
- The security policy and access rights,
- Security services and technologies AAA, (authentication, authorization, accounting)
- Level of security for local operation systems

Only deep analysis of described problems (and maybe others which are not presented here) in context of above tasks can help in creation of the security model for the CoreGRID which will fulfill all Resource Management and Scheduling requirements in context of T & S.

4.5.1 Trust and Security in the Institute on Grid Systems, Tools and Environments

One of the main objectives of the Institute on Grid Systems, Tools, and Environments is to design a generic grid platform, based on extensible component technology. Within the Institute there is one research group Security in Grid Platforms tackling issues of trust and security. The aim of the research group is to address the security issues in the institute research groups, most specifically in applications and middleware components. Hence, we are focused on trust and security requirements for services provided in generic grid platforms.

We assume an advanced grid platform should support the following security related requirements and functionalities:

- User Management user registration and user authentication featuring single sign-in
- Authorization for access to resources
- Auditing of user access and resources usage
- Delegation of rights for resource sharing

The main concepts we consider of key importance for implementing the desired minimal but sufficient security model are:

- Context-oriented layering of the security services infrastructure recognizing two distinguished levels:
 - trusted network (intranet, cluster) - bypass security for performance
 - collaborative network (Internet) - minimal performance concern, maximal security
- Pluggable support for any extra security feature.

Our research on trust and security issues is focused in formulating of a Security model consisting of a threat model based on usage scenario spaces, and security domains infrastructure for hierarchically structured grids. Initial results are reported in [Kir05].

The **threat model** is the foundation for building a secure system. It organizes threats and vulnerabilities into general cases so that they can be addressed with certain protection technologies. In our model the threats are categorized according to specific behaviors, this classification being considered more robust than technological categorization, because technologies may change very rapidly. This also encourages a broader view when evaluating security risks.

In our generic grid platform we adopt a role based security model where a set of roles are defined in the grid and the access rights are associated with roles. This corresponds with subdividing the usage scenarios (and consequently - the user/administration accounts) into three spaces (with further subdivision of these spaces into defined roles according to the recruitment for minimal sufficiency).

- service deployers (users)
- resource owners (administrative roles)
- service producers (administrative roles)

The security models include possible relationship assumptions between actors from the different spaces. The assumptions have graph-like presentation by links of type trust/no-trust between the actors of the same or different spaces.

We consider a generic grid platform consisting of multiple security domains following physical or trust-based boundaries. According to the hierarchical concept any large domain is divided into multiple sub domains, thus forming a layered structure within the domain. The division into sub domains may be performed according to security requirements or level of devices. The model includes definition of the interfaces between the sub domains. Specifically the trustworthiness of services and resources (in sense of service provision trust and resource access trust) can be quantified and used further for developing of scheduling and workflow control systems as formulated in [Anc07].

The trust and security measures are very likely to consume more resources, thus hindering the provision of certain QoS. A specific requirement of our security model is achieving balance between security and QoS.

5 Conclusions

The purpose of this update survey is to summarise current technologies for managing trust and security in Grids so that CoreGRID participants could make use of them in their developments. Section 2 starts by studying the relation between trust and security in distributed system. From the technological point of view, section 3 is central to the document: it describes general trust and security requirements, and the existing technologies for achieving such requirements. The work in section 4 presents trust and security requirements for the CoreGRID Institutes and summarises current security developments within the project tackling such requirements:

- The security work within the Institute on Data and Knowledge Management has concentrated on analysing the technologies for securing storage systems. Two approaches have been followed: one is the analysis of the current technologies for storage security using a framework originally proposed for generic storage systems, which we have been extended to Grid-specific configurations [Lun07]; the other approach has adapted established requirements-engineering methodologies for expressing security requirements of Grid-based data management systems and derive systematically its security policies [Naq06, Naq07]. It is worth mentioning that the systematic derivation of policies from requirements, as initiated in CoreGRID and being further extended in the EU FP6 GridTrust project, has been identified as one of the current challenges by the NESSI-Grid project.
- The security work within the Institute on Programming Models has focused on the evaluation of the impact of introducing security mechanisms in the code of typical parallel Grid applications, and the design of security policies in the GCM component managers suitable to secure component deployment and component invocation, among others. Initial results presented in [Kil07, Ald07a] outline how security-related metadata provided by the programmer can be exploited to optimize the security mechanisms usage in grid applications. The security work carried out in this Institute was created as response to the need of raising the level of abstraction of programming models, as expressed in the Third Report of the Next Generation Grids Expert Group.
- The security work within the Institute on Architectural Issues has identified security issues in the areas of interest for the Institute: scalable grid services and resource discovery, adaptive management of systems, and dependability. Examples of security work in this Institute include a study on reputation techniques and their applicability to Grids [Sil07a], the use of trust to develop more effective sabotage-tolerance techniques in Desktop Grids [Dom07], and investigation on collusion threats in P2P Desktop Grids [Sil07b].
- The Institute on Grid Information, Resource and Workflow Monitoring Services has identified several security requirements in network monitoring, checkpointing architecture, workflow as well as accounting and user management. Fine-grained and role-based security mechanisms for workflow systems have been studied in [Hoh06]. Other promising results are presented in [Den06], which proposes a framework for virtual environments authentication and confidentiality.
- The Institute on Resource Management and Scheduling has carried out an exhaustive security requirements exercise, identifying security issues when analysing scheduling architectures, evaluating and benchmarking scheduling system, and mapping and scheduling of HPC applications, among others. An example of the security results achieved by the Institute is shown in [Far07], which extends UNICORE authorisation capabilities.
- The security work in the Institute on Grid Systems, Tools and Environments is focusing on security for lightweight Grids. They have developed security models based on usage scenario spaces, and security domains infrastructure for hierarchically structured Grids, as shown in [Kir05].

References

- [Abr95] M.D. Abrams, M.V. Joyce. Trusted Computing Update. *Computers and Security*, 14(1): 57-68. 1995.
- [Ald07a] M. Aldinucci, M. Danelutto, P. Kilpatrick. Adding Metadata to ORC to Support Reasoning about Grid Programs. *Proceedings of the CoreGRID Symposium 2007*
- [Ald07b] M. Aldinucci, M. Danelutto. The Cost of Security in Skeletal Systems. *Proceedings of PDP07, 2007*.
- [Alf03] R. Alfieri et al. VOMS: An Authorization System for Virtual Organizations. In *Proceedings of 1st European Across Grids Conference, Santiago de Compostela, 2003*. Available from: <http://grid-auth.infn.it/docs/VOMS-Santiago.pdf>.

- [Anc07] A. Anciaux-Sedrakian, R. M. Badia, J. M. Perez, R. Sirvent, T. Kielmann, A. Merzky. Reliability and Trust Based Workflows' Job Mapping on the Grid. CoreGRID Technical Report TR-0069, 2007
- [Are05] A.E. Arenas, I. Djordjevic, T. Dimitrakos, L. Titkov, J. Claessens, C. Geuer-Pollman, E.C. Lupu, N. Tuptuk, S. Wesner, L. Schubert. Towards Web Services Profiles for Trust and Security in Virtual Organisations. IFIP Working Conference on Virtual Enterprises - PRO-VE'05, Valencia, Spain. 2005.
- [Bas05] J. Basney, M. Humphrey, V. Welch. The MyProxy Online Credential Repository. IEEE Software Practice and Experience, vol 35, issue 9, pp. 801-816, 2005.
- [Bil05] A. Bilas (ed). Roadmap Version 1 on Knowledge and Data Management. CoreGRID Deliverable D.KDM.01, 2005.
- [Bor01] J.J. Borking, C. D. Raab. Laws, PETs and other Technologies for Privacy Protection. Journal of Information, Law and Technology (JILT), Issue 1, 2001.
- [Bra03] M. Brady, D. Gavaghan et al. eDiamond: A Grid-Enabled Federated Database for Annotated Mammograms. In F. Berman, G. Fox, T. Hey (editors), Grid Computing: Making the Global Infrastructure a Reality, Wiley, 2003
- [Bra03a] J. Bradshaw, A. Uszok, R. Jeffers, et al. Representation and reasoning about DAML-based policy and domain services in KAoS. In Proc. of The 2nd Int. Joint Conf. on Autonomous Agents and Multi Agent Systems (AAMAS2003). 2003.
- [Bro03a] P.J. Broadfoot, G. Lowe. Architectures for Secure Delegation within Grids. Oxford University Computing Laboratory Technical Report, PRG-RR-03-19, 2003.
- [Bro03b] P.J. Broadfoot, A.P. Martin. A Critical Survey of Grid Security Requirements and Technologies. Oxford University Computing Laboratory Technical Report, PRG-RR-03-15, 2003.
- [Bur99] JM Burn, P Marshall, M Wild. Managing Changes in the Virtual Organisation. Proceedings of the Seventh European Conference on Information Systems 40-54, Copenhagen Business School, Copenhagen, 1999.
- [Cam03] L.M. Camarinha-Matos, H. Afsarmanesh. A Roadmap for Strategic Research on Virtual Organizations. Proceedings of IFIP Working Conference on Virtual Enterprises - PRO-VE'03, Lugano, Switzerland, pages 33-46, 2003.
- [Cas98] C. Castelfranchi, R. Falcone. Principles of Trust for MAS: Cognitive Anatomy, Social Importance, and Quantification. In Y. Demazeau (editor), Proceedings of the Third International Conference on MultiAgent Systems. IEEE C.S., Los Alamitos, 1998.
- [Cas00] C. Castelfranchi, R. Falcone, B. Sadighi, Y-H Tain. Guest Editorial. Applied Artificial Intelligence, 14(9), Taylor & Frances, 2000.
- [Cha07] A. Chakrabarti. Grid Computing Security. Springer, 2007.
- [Cha05] D. Chadwick. Authorisation in Grid Computing. . Information Security Technical Report, Elsevier, 10(1)33:40, 2005.
- [Cha02] D. Chadwick, S. Otenko. The PERMIS X.509 Role Based Privilege Management Infrastructure. In Proceedings of 7th ACM Symposium on Access Control Models and Technologies (SACMAT 2002), 2002.
- [Cra06] L. Cranor et al. The Platform for Privacy Preference 1.1 (P3P1.1) Specification. W3C Working Group Note, 2006. Available at <http://www.w3.org/TR/P3P11/>.
- [Cra02] L. Cranor M. Langheinrich, M. Marchiori. A P3P Preference Exchange Language 1.0 (APPEL 1.0). W3C Working Draft, 2002. Available at <http://www.w3.org/TR/P3P-preferences/>.
- [Cro05] S. Crompton, B. Matthews, A. Gray, A. Jones, R. White. Data Integration in Bioinformatics using OGSA-DAI. In Proceedings of Fourth All Hands Meeting, AHM2005, UK, 2005.
- [Dan07] F. D'andria, J. Martrat, T. Kirkham, S. Naqvi, J. Gallop, A.E. Arenas. The Evolving Use of sService Level Agreements and the Influence of Trust within the Support and Development of Grids to Enable a Next Generation of Business Models. International Workshop on Service Oriented Computing: a Look at the Inside (SOC@Inside'07), Austria, 2007
- [Den06] J. Denemark, M. Jankowski, K. Ales, L. Matyska, N. Meyer, M. Ruda, P. Wolniewicz. Best Practices of User Account Management with Virtual Organization Based Access to Grid. In Parallel Processing and Applied Mathematics, Vol. 3911:633-642 of LNCS, Springer-Verlag, 2006.
- [Den03] G. Denker, L. Kagal, T. Finin, M. Paolucci and K. Sycara. Security for DAML Web Services: Annotation and Matchmaking. In D. Fensel, K. Sycara, J. Mylopoulos (Ed.), The Semantic Web-ISWC 2003. Proceedings of the 2nd International Semantic Web Conference, Sanibel Island, Florida, USA, October 2003, LNCS 2870, 2003.
- [DIA03] A.E. Arenas (editor). Survey Material on Trust and Security. CoreGRID Deliverable D.IA.03, 2005.
- [DIA016] A.E. Arenas (editor). Update of the Survey Material on Trust and Security. CoreGRID Deliverable D.IA.16, 2007.

- [Dim01] T. Dimitrakos. System Models, e-Risk and e-Trust. Towards Bridging the Gap? in *Towards the E-Society: E-Business, E-Commerce, and E-Government*, eds. B. Schmid, K. Stanoevska-Slabeva, V. Tschammer. Kluwer Academic Publishers, 2001.
- [Dom07] P. Domingues, B. Sousa, L. M. Silva. Sabotage-Tolerance and Trust Management in Desktop Grid Computing. *Future Generation Computer Systems*, In Press, 2007.
- [Far07] A. Faroughi, R. Faroughi, W. Ziegler, P. Wieder. Attributes and VOs: Extending the UNICORE Authorisation Capabilities. *CoreGRID Technical Report TR-0097*, 2007.
- [Fos98] I. Foster, C. Kesselman, G. Tsudki, S. Tuecke. A Security Architecture for Computational Grids. In *Proceedings of 5th ACM Conference on Computer and Communication Security*, 1998.
- [Fos01] I. Foster, C Kesselman, S. Tuecke. The Anatomy of the Grid: Enabling Scalable Virtual Organizations. *International Journal of Supercomputing Applications* 15(3), 200-222, 2001.
- [Fos03] I. Foster, C. Kesselman. *The Grid: Blue Print for a New Computing Infrastructure*. Morgan Kauffmann, 2003.
- [Gam88a] D. Gambetta (editor). *Trust: Making and Breaking Cooperative Relations*. Department of Sociology, University of Oxford, 1988. Available at <http://www.sociology.ox.ac.uk/papers/trustbook.html>.
- [Gam88b] D. Gambetta. Can We Trust Trust? In [Gam88a], chapter 13, 1988.
- [Gas90] M. Gasser, E. McDermott. An Architecture for Practical Delegation in a Distributed System. *IEEE Symposium on Research in Security and Privacy*, 1990.
- [Gra00] T. Grandison, M. Sloman. A Survey of Trust in Internet Applications. *IEEE Communications Survey and Tutorials*, 3, 2000.
- [Gra03] T. Grandison, M. Sloman. Trust Management Tools for Internet Applications. In [Nix03], 2003.
- [Gru01] J. Grudin. Desituating Action: Digital Representation of Context. *Human-Computer Interaction (HCI) Journal*, 2001.
- [Gue05] C. Geuer-Pollmann, J. Claessens. Web Services and Web Service Security Standards. *Information Security Technical Report*, Elsevier, 10(1)15:24, 2005.
- [Her05] P. Hermann, V. Issarny, S. Shue (editors). *Third International Conference on Trust Management*. Lecture Notes in Computer Science, vol. 3477, Springer, 2005.
- [HG04] The HealthGrid White Paper. Available at <http://initiative.healthgrid.org/the-initiative/healthgrids-concept/white-paper.html>, 2004.
- [Hoh06] A. Hoheisel, S. Mueller, B. Schnor. Fine-Grained Security Management in a Service-Oriented Grid Architecture. In *Proceedings of the Cracow Grid Workshop '06*. Cracow, Poland, 2006.
- [Hon04] J.I. Hong, J. Ng, J.A. Landay. Privacy Risk Models for Designing Privacy-Sensitive Ubiquitous Computing Systems. In *Proceedings of Designing Interactive Systems (DIS2004)*. Boston, MA, pp. 91-100, 2004.
- [Jen04]. C.D. Jensen, S. Poslad, T. Dimitrakos (editors). *Second International Conference on Trust Management*. Lecture Notes in Computer Science, vol. 2995, Springer, 2004.
- [Joh03] W.E. Johnston, J.M. Brooke, R. Butler, D. Foster and M. Mazzucato. Production Deployment: Experiences and Recommendations. In [Fos03], 2003.
- [Joh03a] M. Johnson, P. Chang, R. Jeffers et al. KAoS semantic policy and domain services: An application of DAML to Web services-based grid architectures. *Proceedings of the AAMAS 03 Workshop on Web Services and Agent-Based Engineering*. Melbourne, Australia, 2003.
- [Jon99] S. Jones. TRUST-EC: Requirements for Trust and Confidence in E-Commerce. *European Commission Joint Research Centre*, 1999.
- [Jos99] A. Josang. An Algebra for Assessing Trust in Certification Chains. In J. Kochmar (editor), *Proceedings of the Network and Distributed Systems Security Symposium (NDSS'99)*. The Internet Society, 1999.
- [Jos07] A. Josang, R. Ismail, C. Boyd. A Survey of Trust and Reputation Systems for Online Service Provision. *Decision Support Systems*, 43(2), pp 618-644, 2007.
- [Kag03] L. Kagal, T. Finin, J. Anupam. A Logical Policy Language for a Pervasive Computing Environment., 4th *IEEE Int. Workshop on Policies for Distributed Systems and Networks*, Lake Como, 4-6 June, 2003.
- [Kan01] M. H. Kang, J. S. Park, J. N. Froscher. Access Control Mechanisms for Inter-organisational Workflow. In *Proceedings of the sixth ACM symposium on Access control models and technologies (SACMAT'01)*, ACM Press, 66-74, 2001.
- [Kat00] BR Katzy, C Zhang, H loeh. Reference Models for Virtual Organizations. Working Paper No 2704, Working Paper Series, CeTIM, 2000.
- [Kil07] P. Kilpatrick, M. Aldinucci, M. Danelutto. Deriving Grid Applications from Abstract Models. *CoreGRID Technical Report TR-0085*, 2007.

- [Kin98] A. Kini, J. Choobineh. Trust in Electronic Commerce: Definition and Theoretical Consideration. Proceedings of 31st International Conference on System Sciences, IEEE, 1998.
- [Kir05] L. Kirchev, M. Blyantov, V. Georgiev, K. Boyanov, M. Malawski, M. Bubak, S. Isaiadis, V. Getov. Security Models for Lightweight Grid Architectures. CoreGRID Technical Report TR-0023, 2005.
- [Lan02] M. Langheinrich. A Privacy Awareness System for Ubiquitous Computing Environments. UbiComp 2002: Ubiquitous Computing, Lecture Notes in Computer Science, vol. 2489, 2002.
- [Lin04] J. Linn (ed). Liberty Trust Models Guidelines. Liberty Alliance Project, version 1.0, 2004.
- [Lun07] J. Luna et. al. An Analysis of Security Services in Grid Storage Systems. In Proceedings of the CoreGRID Workshop on Grid Middleware 2007. Dresden, Germany. June, 2007.
- [Mar03] S. T. Margulis. Privacy as a Social Issue and Behavioral Concept. Journal of Social Issues, 59(2):243-262, 2003.
- [Mar07] F. Martinelli, P. Mori. A Model for Usage Control in GRID Systems. In Proceedings of Grid-STP 2007, International Workshop on Security, Trust and Privacy in Grid Systems, IEEE, 2007.
- [McK96] D.H. McKnight, N.L. Chervany. The Meaning of Trust. Technical Report MISRC Working Paper Series 96-04, University of Minnesota. Management Information Systems Research Center, 1996.
- [Mey06] N. Meyer. Roadmap Version 2 on Grid Information, Resource and Workflow Monitoring Services. Core-GRID Deliverable D.IRWM.03, 2006.
- [Mid04] S. Middleton et al. GEMSS: Privacy and security for a Medical Grid. In Proceedings of HealthGrid 2004.
- [Mor08] L. Moreau et al. The Provenance of Electronic Data. Communication of the ACM, 51(4):52-58, 2008.
- [Nag03] N. Nagaratnam, P. Janson, J. Dayka, A. Nadalin, F. Siebenlist, V. Welch, S. Tuecke, I. Foster. Security Architecture for Open Grid Services. Available at <http://forge.gridforum.org/projects/ogsa-sec-wg/>.
- [Naq07] S. Naqvi, C. Ponsard, P. Massonet, A.E. Arenas. Security Requirements Elaborations for Grid Data Management Systems. International Journal of Systems of Systems Engineering. In Press, 2007.
- [Naq06] S. Naqvi, A.E. Arenas, P. Massonet. Deriving Policies from Grid Security Requirements Model. In Proceedings of the 2nd CoreGRID Integration Workshop, Cracow, 2006.
- [Nix03] P. Nixon, S. Terzis (editors). First International Conference on Trust Management. Lecture Notes in Computer Science, vol. 2692, Springer, 2003.
- [Pal03] L. Palen, P. Dourish, Unpacking "Privacy" for a Networked World. CHI Letters, 2003. 5(1): pp. 129-136.
- [Park04] J. Park, R. Sandhu: The UCON Usage Control Model. ACM Transactions on Information and System Security, 7(1) 2004, 128-174
- [Pea02] L. Pearlman, V. Welch, I. Foster, C. Kesselman, S. Tuecke. A Community Authorization Service for Group Collaboration. In Proceedings of the IEEE 3rd International Workshop on Policies for Distributed Systems and Networks, 2002.
- [Ras96] L. Rasmusson, S. Janssen. Simulated Social Control for Secure Internet Commerce. In C. Meadows, editor, Proceedings of the 1996 New Security Paradigms Workshop. ACM, 1996.
- [Res00] P. Resnick, P. Zeckhauser, R. Friedman, K. Kuwabara. Reputation Systems. Communications of the ACM, 43(12):45-48, December 2000.
- [San02] R. Sandhu, M. Bellare, R. Ganesan. Password-Enable PKI: Virtual Smart-Cards Versus Virtual Soft Tokens. In 1st Annual PKI Workshop, pp 89-96, 2002.
- [Sca05] T. Scavo, S. Cantor (eds). Shibboleth Architecture: Technical Overview. Working Draft 0.2, Internet2, 2005. Available at <http://shibboleth.internet2.edu/docs/draft-mace-shibboleth-tech-overview-latest.pdf>.
- [Sie03] F. Siebenlist, N. Nagaratnam, V. Welch, C. Neuman. Security for Virtual Organizations: Federating Trust and Policy Domains. In [Fos03].
- [Sil07a] G. C. Silaghi, A. E. Arenas, L. M. Silva. Reputation-Based Trust Management Systems and their Applicability to Grids. CoreGRID Technical Report TR-0064, 2007.
- [Sil07b] G.C. Silaghi, L.M. Silva, P. Dominques, A.E. Arenas. Tackling the Collusion Threat in P2P-Enhanced Internet Desktop Grids. CoreGRID Workshop on Grid Programming Model, Grid and P2P System Architecture, Grid Systems, Tools and Environments, Heraklion-Crete, Greece, 2007.
- [Ste03] M. Steenbakkens. Guide to LCAS v.1.1.16, September 2003. Available at <http://www.dutchgrid.nl/DataGrid/wp4/lcas/edg-lcas-1.1>.
- [Sur02] M. Surrige. A Rough Guide to Grid Security. Technical Report, IT Innovation Centre, V1.1a, 2002.
- [Sur05] M. Surrige, J. Claessens. TG6 Trust and Security - White Paper on State of the Art and Planned Developments in the Context of FP6 Grid Projects.

- [Tho99] M. Thompson, W. Johnston, S. Mudumbai, G. Hoo, K. Jackson, A. Essiari. Certificate-Based Access Control for Widely Distributed Resources. In Proceedings of the Eighth USENIX Security Symposium (Security '99), pages 215-228, 1999.
- [Ton03] G. Tonti, J. Bradshaw et al. (2003). Semantic Web languages for policy representation and reasoning: A comparison of KAoS, Rei, and Ponder. In D. Fensel, K. Sycara, J. Mylopoulos (Eds.), The Semantic Web-ISWC 2003. Proc. of the 2nd Int. Semantic Web Conf., Sanibel Island, Florida, USA, October 2003, LNCS 2870
- [Wai02] M. Waidner (editor). Ercim News, Special Theme: Information Security. No 49, 2002.
- [Wal02] D. S. Wallach. A Survey of Peer-to-Peer Security Issues. In Proceedings of the International Symposium in on Software Security - Theories and Systems, Tokyo, Japan. Springer Lecture Notes in Computer Science 2609, 42-57, 2002.
- [Wel03] V. Welch, F. Siebenlist, I. Foster, J. Bresnahan, K. Czajkowski et al. Security for Grid Services. In Proceedings of 12th IEEE International Symposium on High Performace Distributed Computing. IEEE Computer Society Press, 2003.
- [Wes05] S. Wesner, L. Schubert, T. Dimitrakos. Dynamic Virtual Organizations in Engineering. In Proceedings of German-Russian Workshop, 2005.
- [Zim95] P.R. Zimmermann. The Official PGP User's Guide. MIT Press, 1995.