IDIS (2009) 2:171–187
DOI 10.1007/s12394-009-0027-1

A security framework for dynamic collaborative working environments

Matthias Assel · Stefan Wesner · Alexander Kipp

Received: 19 January 2009 / Accepted: 2 September 2009 / Published online: 8 October 2009

Abstract Moving away from simple data sharing within the science community towards cross-organizational collaboration scenarios significantly increased challenges related to security and privacy. They need to be addressed in order to make cross-organizational applications such as collaborative working environments a business proposition within communities such as eHealth, construction or manufacturing. Increasingly distributed scenarios where many different types of services need to be combined in order to implement semantically enriched business processes demand new approaches to security within such dynamic Virtual Organizations. The allocation of access rights need to be possible in an easy and controlled way in order to allow inexperienced users to maintain the information and ensure compliance e.g. with legal and privacy related regulations. In this paper the focus is how security concepts originating from the Grid domain have been applied for collaborative working environments. The chosen scenarios are a Virtual Laboratory for Infectious Diseases (ViroLab) and different collaborative environments from the engineering domain as defined within the CoSpaces project. The requirements from these scenarios are analyzed and a security model enabling such dynamic, secure and trustworthy collaborations is presented.

Keywords Security model · Grid · Virtual Organization · Collaborative working environment · Decentralized identity management · Dynamic policy enforcement

M. Assel (✉) · A. Kipp
Intelligent Service Infrastructures, High Performance Computing Center Stuttgart, Nobelstr. 19, 70569 Stuttgart, Germany
e-mail: assel@hlrs.de

S. Wesner
Applications & Visualization, High Performance Computing Center Stuttgart, Nobelstr. 19, 70569 Stuttgart, Germany

Introduction

As of today most professionals still rely on “local” information, i.e. own expertise, company internal experts or perform manual research in globally available information sources such as libraries, literature or Internet search engines for gathering adequate knowledge for making decisions. With more complex and distributed workflows, such information is easily outdated and difficult to find and access. This can lead to huge delays within entire manufacturing processes for example during design and/or production or for the selection of an appropriate treatment for a patient. In order to speed up such business collaborations, workflows, and processes the necessary information need to be accessible for every employee in a fast, easy and secure way. Additionally, the provision of internal information to collaborators must be possible in a controlled manner.

While cross-organizational data exchange is a daily routine most people do not actually perceive this as a security critical element of their daily work. Sending e-mails to business partners with sensitive information attached to them or transferred via file transfer capabilities of chat tools are common practices. Consequently an environment for data sharing where more control can be maintained about which data is provided to collaboration partners and in which context is necessary.

The concept of a Virtual Organization (VO) is widely used to provide such an environment, namely to make (data/information) resources available dynamically, securely, and on-demand (Schubert et al. 2005). The main purpose of such a concept consists hence in enabling dynamic collaborations with easy access to different resources, respectively a secure sharing of relevant data/information/knowledge, tools/services, and workflows.

In order to achieve such inter-organizational business collaborations that support and enable the sharing of adequate expertise and relevant information, several requirements from both perspectives, the provider’s as well as from the user’s (customer’s) site have to be taken into account and carefully addressed by upcoming systems. Looking from the resource provider’s point of view, different issues and particularly concerns such as security, trustworthiness, and integrity of provided information/data but also legal aspects like copyrights and privacy issues are of great importance before any collaboration could ever be established. The abuse of confidential data pieces-not necessarily personal data but mainly business secrets-is one of the main reasons why current business collaborations are basically limited to uncritical workflows that do not immediately involve any crucial (manufacturing) data. Furthermore, these inter-company liaisons are principally organized together with associated companies in order to take control of each dataflow between corresponding entities in case of extraordinary circumstances.

To overcome these concerns and tackle current problems and limitations, future collaborative working environments have to follow international standards and guarantee a certain level of protection, reusability, interoperability, scalability, and finally trustworthiness. Thus, existing products need to be extended with new security models and technologies that facilitate a smooth and secure exploitation of local resources. Unfortunately, most of today’s solutions such as commonly used identity or database management systems do not actually provide capabilities that can be easily adapted or further extended in order to allow collaborations with external partners in a shorten time and provide access to own resources for foreign users that are explicitly unknown to the locally deployed system(s). For instance, users belonging to a foreign company could not directly use their primary identity information to run remote applications or access resources although being securely connected to the remote network. That is specific problem might result from inconsistent identity tokens that were created by systems using different standards. This lack of flexibility and functionality very often limits enterprises to local applications for data and information sharing but also prevents them from taking advantage of cross-organizational business synergies including a mutual information/knowledge/data exchange that might help improving their own internal business workflows, too.

The Reminder of this article is structured as follows. Starting from the Virtual Organization concept developed within "Service Grid" community the concept of a collaborative working environment as understood here is introduced and further detailed with two application cases. Both are targeting next-generation collaborative working environments in different application domains namely engineering and eHealth. Based on these scenarios their requirements on security are discussed in detail and a new security model overcoming current limitations and reflecting the needs for future collaboration scenarios are presented. The article is concluded with an assessment of the applicability of the proposed model, its limitations and the derived future research challenges.

Modern collaborative working environments

Collaboration models for audio and video communication had been changing from one-to-one communication quite quickly towards one-to-many (e.g. lecture mode) and many-to-many models. The same applies to other sharing applications used on many computing systems such as chat tools or slide sharing applications. The collaboration in such scenarios is started spontaneously in an ad-hoc manner without long preparatory steps. However this inherent flexibility is at the same time a major problem preventing their use beyond informal information exchange. The major drawback is the complete lack of control about the exchanged information. Collaboration within a professional environment needs to be compliant with legal and regulatory constraints as well as companywide policies (e.g. with whom this particular data can be shared).

The collaboration model introduced by the Grid community called Virtual Organizations (VO) for sharing information and resources was initially a coalition of very similar type of providers aiming for a long-term collaboration (e.g several years). This model is designed for applying rigid access control policies on data and resources and is often realized by utilizing a central control instance for user identities and VO memberships (Wesner and Kipp 2007).

However more dynamic collaboration models across organizational boundaries had been developed as extension of the outsourcing concepts already before in the field of economics (see for example Saabeel et al. 2002) proposing a much more dynamic model without any kind of centralized control.

Several research projects such as TrustCoM (Schubert et al. 2005) further developed such dynamic Virtual Organizations enabling collaborative business models. The Akogrimo project (Wesner 2005) realized a mobility aware VO model integrated with the dynamic session concept of multimedia communication.

In order to make this dynamic Virtual Organization a powerful instrument for collaborations with the simplicity of the ad-hoc collaborations of video conferences and chat applications but at the same time compliant with the legal and company policies, the collaboration infrastructure need a tight integration of all kind of shared resources if they are video streams, text windows or IT resources. Additionally, the collaboration need to be more controllable, i.e. that decision processes are documented and how the input data for a collaboration session has been produced and which tools had been used to produce this data.

Hence, in this context we define the term Collaborative Working Environment as follows:

> "A Collaborative Working Environment is a temporary coalition of resources, services and people from different organizations allowing the exchange of information and knowledge in order to work together and achieve a shared goal."

The following two sections outline in more detail two specific instances of such collaborative working environments in the engineering and the eHealth domain.

The CoSpaces collaborative workspaces

Due to an emerging globalization, many engineering companies are decentralizing their operations and working teams as distributed virtual enterprises with the objectives to optimize costs, quality and time. However, the teams working as such virtual enterprises need to come together frequently to assess the project from different engineering points of view and to ensure that the product is meeting the specifications and the predefined set of quality standards. The CoSpaces project aims for developing a distributed engineering environment for remote teams to work together securely and efficiently by reducing the frequency of travelling to a single location. It also aims for supporting both planned and ad-hoc meetings within a distributed virtual workspace (Kipp et al. 2008) with appropriate decision-making and communication tools. The key technical challenges therefore undertaken by the CoSpaces project are:

- Dynamic integration of decision-making and communication tools into the distributed environment;
- Secure cross-domain data sharing;
- Multi-user visualization and interaction metaphor for collaborative working;
- Creation of co-located, distributed and mobile workspaces for multi-functional teams.

Figure 1 gives an overview of involved components within the CoSpaces-Framework. Table 1 shows the technologies being used for the corresponding components as well as collaboration technologies being applied within the different

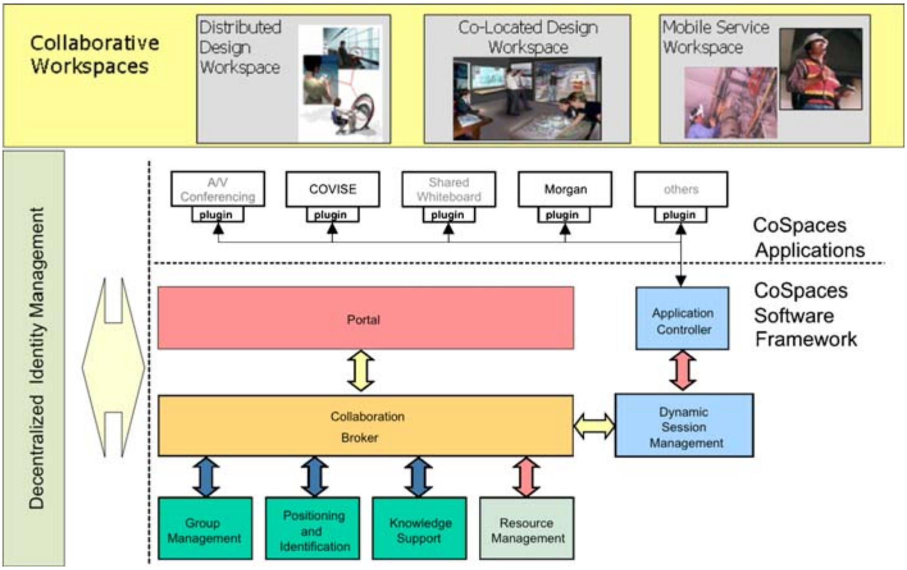

Fig. 1 The CoSpaces framework

workspaces. A detailed description of these components go beyond this paper so interested readers are referred to visit the CoSpaces project website[1].

Distributed workspaces

In the automotive sector, worldwide Original Equipment Manufacturers (OEMs) are supplied by subcontractors from different locations. Similarly, in the aerospace industry several parts of an aircraft are designed and manufactured in geographically distributed locations. The coordination between provider, developer, supplier and customer is costly in staff times and other expenses.

For this reason CoSpaces develops a distributed design workspace for dispersed teams to reduce the need for face-to-face meetings in the overall product lifecycle.

Co-located workspaces

Co-located meetings are frequently organized by industry during the whole product lifecycle. These periodic meetings have the objective to identify issues or potential problems across and between various competencies/skills. Indeed, many industrial studies have shown that identification of problems early in the lifecycle can avoid excessive exponential cost and time overruns.

The efficiency of a project is dependent upon the various geographically and culturally diverse players across the entire supply chain. Traditional tendering strategies give limited possibilities to exchange information and share experience in

[1] http://www.cospaces.org

Table 1 CoSpaces framework components and technologies

Component	Technologies
Decentralized Identity Management	Shibboleth
CoSpaces Portal	MS Silverlight
Knowledge Support	BSCW
Resource Management	Web Service
Collaboration Broker	Web Service
Dynamic Session Management	Web Service
Application Controller	Web Service
Distributed Visualization	Covise
Synchronous A/V conferencing	AccessGrid

a collaborative manner. On the other hand, the new tendering procedures that have been introduced and increasingly used in the industry to enable early and cross-discipline identification of problems require collaborative discussion and decision-making.

Mobile workspaces

The construction industry is faced with a dilemma that as soon as building is in progress some problems might occur having impact on the handover time to the customer.

These frequent unforeseen situations require a decision or action urgently. Such decisions often require a chain of authorization involving a diversity of actors. These situations may fall outside the responsibility of the operative on site, though she might nonetheless initiate a decision or knowledge acquisition process that is likely to engage these actors and other experts. In such situations, there is a need for extensive collaboration between various actors as well as the need for external expertise, which could be human expertise or be presented in form of information or knowledge by accessing the relevant data sources. However, the fragmented nature of construction projects inhibits orchestrated and fast decision-making. Further, the site-based constrains do not allow seamless access nor access across organizational boundaries to the knowledge that is required to support the collaboration and decision making processes.

The ViroLab Virtual Laboratory

The ViroLab Virtual Laboratory is an integrated set of tools and services for accessing and integrating distributed heterogeneous resources (Gubala et al. 2008). Its main purpose is to facilitate medical knowledge discovery and decision support for HIV drug resistance (Sloot et al. 2006) as well as other types of research in the general field of eScience.

Those research studies are carried out in a collaborative working environment using state-of-the-art service-oriented computing technologies and standards and consisting of computing and data resources deployed at various networked sites. As the complexity of interfacing such resources often presents a steep learning curve for application developers, the main goal of the Virtual Laboratory is to present a powerful, flexible and dynamic environment for experiment developers and medical end-users while preserving ease of use and reusability of the proposed solution, thus allowing transparent and secure access to corresponding underlying infrastructures.

Primarily, the virtual workspace is used by medical doctors to review actual HIV drug rankings and recent drug resistance interpretations or by scientists to conduct new experiments and simulations starting from pre-defined process flow templates, which allow an interactive and smooth selection of available bioinformatics applications to be combined into one explicit workflow for studying individual drug resistance susceptibility. Furthermore, the system offers different capabilities such as on-demand expert consultation or real-time data sharing allowing easy collaborations with other medical professionals for discussing previous results and experiments.

In order to support and provide these different functionalities, the overall architecture consists of several components each providing different levels of services. The main driver for building this distributed infrastructure was the fact that data, information, users, and services are dispersed all over Europe, and only a decentralized system could approach this challenge and be sustainable enough to ensure functionality beyond the project's duration. In Fig. 2, an overview of involved modules and their main interactions are schematically shown.

A detailed description of the ViroLab Virtual Laboratory design is outside the scope of this paper but the interested reader should visit the project's website[2] for further information.

Key security issues and requirements

Lots of progress has been made in the past and quite a few solutions have been developed but all these approaches do not explicitly consider the very specific needs in case of dynamic collaboration sessions and distributed data handling respectively.

Nevertheless, the dynamic setup of collaboration sessions and the management of sensible data are considered as a very critical and challenging aspect within the CoSpaces project. Due to the necessity that during industrial collaborations confidential data needs to be shared among all parties involved, the CoSpaces framework provides an infrastructure, which supports the secure distribution of corresponding pieces of data only to participants being foreseen for a specific collaboration. One of the most important issues for industrial partners concerning cross-organizational data exchange constitutes the control of sensible data basically *who is allowed to access or modify which data sets and for which purposes*. In particular, data providers keep the full control on respective data sets and thus, those are not stored at any trusted third party side. These central requirements enforce a decentralized approach ensuring security and trust (Assel and Kipp 2007) among all involved partners within a collaboration.

Similar concerns came up in the context of the ViroLab project since the sensitivity and confidentiality of personal data sets shared within the Virtual

[2] http://www.virolab.org

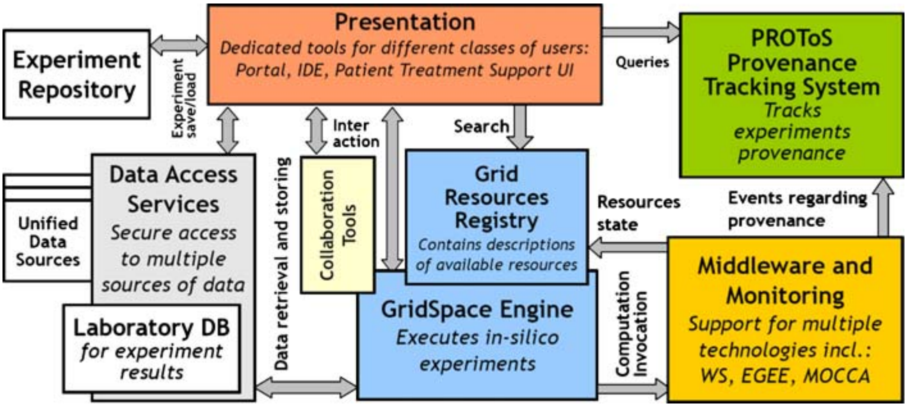
Experiment Repository
Experiment save/load
Presentation
Dedicated tools for different classes of users: Portal, IDE, Patient Treatment Support UI
Queries
PROToS Provenance Tracking System
Tracks experiments provenance
Inter action
Search
Unified Data Sources
Data Access Services
Secure access to multiple sources of data
Data retrieval and storing
Collaboration Tools
Grid Resources Registry
Contains descriptions of available resources
Resources state
Events regarding provenance
Laboratory DB for experiment results
GridSpace Engine
Executes in-silico experiments
Computation Invocation
Middleware and Monitoring
Support for multiple technologies incl.: WS, EGEE, MOCCA

Fig. 2 Simplified architecture of the virtual laboratory

Laboratory require strong and dynamic access control for data resources as well as encryption of all information transferred through the infrastructure. ViroLab overcomes this crucial issue by building highly dynamic and flexible data access services that guarantee security on several levels using established security principles and solutions (Assel and Kalyoncu 2008) in order to protect the confidential information and to keep the patient's privacy appropriately.

In the following, the most important requirements for dynamic collaborative working environments are briefly summarized taking the recent project developments and results into account and combining as well as extending them accordingly.

1) The usage of the collaboration platform should be supported through a decentralized Authentication and Authorization Infrastructure (AAI) based on a Single Sign-On (SSO) procedure in order to ensure flexibility and dynamicity for identity management and distributed access control but also to ease applicability for remote users and service providers respectively;
2) Capabilities for users to control any information (attributes) released to requesting service providers-customization of user profiles according to personal settings and context;
3) Dynamic setup of collaboration partners including the on-demand modifications of firewalls in order to guarantee that only trusted people and particular applications are allowed to execute corresponding operations;
4) Dynamic management and control of attribute-based access policies for creating fine-grain access rules;
5) Context-sensitive access control mechanisms taking the actual user context details into account and thus, only information which is securely accessible in that specific context is exchanged;
6) Keeping the users' privacy and protecting their confidentiality by impersonalizing private data or excluding information such as irrelevant fields and/or any kind of metadata;
7) Secure data transmission using data encryption on several levels (encrypted messages versus secure protocols) and/or context-aware security tokens in order to establish and ensure dynamic trustworthiness and integrity of exchanged information;

8) Additional security mechanisms and policies for long-term and intermediate storage of confidential data sets according to state-of-the-art legal and ethical issues;
9) Recording of relevant user interactions for tracking back the origin of certain data sets (data provenance);
10) Monitoring of critical infrastructure components as well as services or applications to react on suddenly intermittent failures;
11) Flexible system(s) for distributing interesting events, failure reports or predefined topics of interest to foreseen users/components (notification support);
12) Possibilities to immediately judge a provider's performance and reliability based on previous events;
13) Methods for defining and negotiating Service Level Agreements (digital contracts) based on QoS[3] parameters and provider reputation in order to guarantee a certain level of reliability, performance, and scalability for customers as well as providers, and to facilitate the individual accouting of single service or resource capabilities.

A security model for dynamic collaborative working environments

Taking the very specific needs of modern Collaborative Working Environments (CWE) in particular with respect to security into account, a flexible and highly dynamic security model shall be introduced in the following sections.

For this model, we consider automatic and dynamic adaption of particular security policies due to ad-hoc changing security requirements as well as the management of distributed identity information as the two main and at the same time most challenging fields. Prior to investigating the relevant security principles in more detail and exploiting them onto a concrete model, a general architecture for collaborative working environments is being presented and briefly discussed.

As illustrated in Fig. 3, a typical CWE can be split into a four-layer architecture although the borderlines between them are blurred and not clearly defined. Common user interfaces like web-based or standalone applications usually can be found on the top level. Through these "doors", different users are entering the virtual workspace, which enables them easy collaboration and data exchange with various other participants. In order to provide these features, a so-called execution and collaboration framework has to be deployed, which basically controls and allows the setup of particular interactive sessions. Such sessions are usually consisting of several independent workflows (comprising different people, services, and resources), which can more or less dynamically be composed. The framework has for example to ensure that the initiator of a session is able to invite various users for joining a running collaboration. Moreover, the usage and/or sharing of embedded applications also needs to be supported. Sometimes, applications require different remote services for performing specific tasks. This could include, e.g. the access to underlying data resources for querying relevant input data but also the transmission of complicated calculations onto heterogeneous computing resources. For these purposes, the runtime

[3] Quality of Service

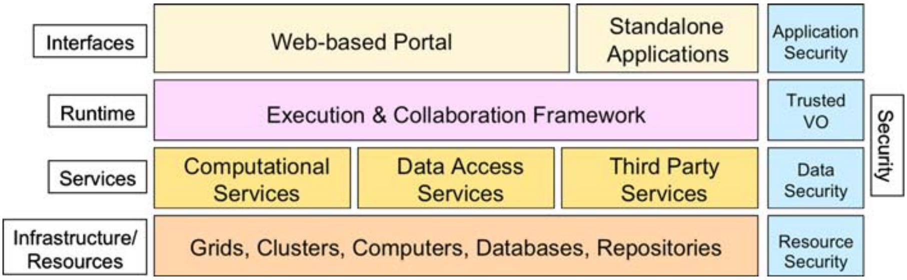

Fig. 3 Simple four-layer architecture

layer has to communicate and interact with distributed (third-party) services by forwarding dedicated tasks (e.g. data queries, complex computations) or requesting additional functionalities (e.g. advanced data visualization). All in all, the entire collaboration infrastructure has to have a flexible, extendable and easy to use methodology that adapts to any dynamic change or requirement almost seamlessly and on demand.

However, security concerns are immediately emerging as soon as any cross-organizational collaboration takes place. In addition, even if all participating entities are located inside the same institutional network, security has to be carefully addressed and appropriately handled. For these reasons, security ranges from the top level devoted to any kind of applications through all lower layers of the overall architectural framework. This involves amongst others the basic execution and collaboration framework, the different interconnected services as well as the underlying infrastructure including single computing nodes, databases and even different file or document repositories. Due to this layered distinction, security duties and responsibilities are distributed throughout the entire stack and not limited to one centralized system and layer respectively. Several components will provide dedicated services being responsible for specific security-related tasks such as user authentication, policy enforcement, data encryption, or distributed access control of integrated resources. In Fig. 4, the aforementioned differentiation of security responsibilities is being depicted in more detail.

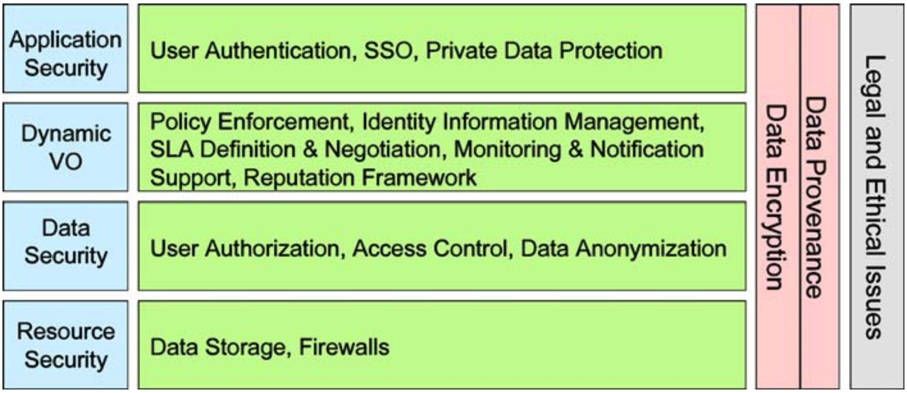
Application Security
User Authentication, SSO, Private Data Protection
Dynamic VO
Policy Enforcement, Identity Information Management, SLA Definition & Negotiation, Monitoring & Notification Support, Reputation Framework
Data Security
User Authorization, Access Control, Data Anonymization
Resource Security
Data Storage, Firewalls
Data Encryption
Data Provenance
Legal and Ethical Issues

Fig. 4 Security stack of collaborative working environments

In order to ensure the challenging issue of building such a decentralized security framework with respect to performance, scalability, reliability and trustworthiness, the architecture of a CWE has to be flexible with a clear focus on the definition of corresponding security policies and duties. To provide this flexibility in terms of management and processing of infrastructure policies, we are adopting the main principles of Virtual Organizations (Wilson et al. 2006) for CWE, too, namely to guarantee trusted collaborations through well-defined security policies and allow their automatic and ad-hoc enforcement and redefinition by dedicated infrastructure components.

In the following, we describe and discuss how those policies and particular security responsibilities and capabilities are being distributed and interacting together throughout the entire security infrastructure.

Establishing dynamic and trusted Virtual Organizations

Identity information management (Requirement 1 and 2)

A security infrastructure should primarily take care of a proper aggregation of relevant identity information from different sides and should allow the exchange among trusted entities only that do require certain user information for any kind of authorization purposes.

Prior to sharing different user-specific characteristics, e.g. personal settings or actual context, such information has to be gathered from different decentralized systems, so-called Identity Provider (IdP) that are integrated with the overall framework and mainly responsible for local (organization-specific) user authentication and distribution of requested identity tokens and user attributes respectively. In order to enable the integration of multiple identity providers and at the same time guarantee a user-friendly way for differently experienced user groups to access foreign resources, we are considering the basic Shibboleth functionalities (Barton et al. 2006) within our security infrastructure. On the one hand, Shibboleth enables Single Sign-On (SSO) capabilities to smoothly access services and resources across or within organizational boundaries and allows the deployment of a decentralized identity management and exchange framework by establishing a trusted federation based on SAML[4] assertions among various identity and service providers (Bhargav-Spantzel et al. 2006). Enhanced privacy functionality is established by allowing the user and their home site to control the attributes released to each application, too.

Policy enforcement (Requirement 3, 7 and 11)

Usually, simple “pre-defined” policies are negotiated and recorded during the setup phase before any collaboration starts. This typically hinders the framework respectively individual components to react on altering circumstances immediately. Hence, policy enforcement should cope with prompt changes and on-the-fly updates of existing policies (Khurana and Gligor 2004) to quickly adapt services and

[4] Security Assertion Markup Language

workflows accordingly and without any notable time shift. For this reason, a dedicated component, the so-called Policy Enforcement Point (PEP), which is responsible for evaluating and processing policy rules while any cross-domain communication takes place (Wasson and Humphrey 2003), needs to be informed (e.g. through event driven dissemination of information (Forestiero et al. 2008)) as soon as, but whenever a policy rule update occurs. Thus, the main role of the PEP is to capture all in- and outgoing message streams and analyze the security headers of relevant messages in order to perform particular actions that have been specified in detail within certain organization-specific security policies. For example, the PEP might enforce users to re-authenticate at the corresponding IdP or automatically tries to request additional security tokens (sometimes required to securely collaborate with external parties and normally provided by a component named Security Token Service (Geuer-Pollmann and Claessens 2005)). Moreover, it is responsible for encrypting or decrypting any kind of information flows as well as checking data integrity of transferred messages. In addition, the PEP also interacts with another central security component, namely the Policy Decision Point (PDP). Prior to forwarding any request to a particular service or resource, the PEP sends an authorization request to the PDP in order to ask for permission(s). The PDP searches its own policy repository and checks for relevant access rules that comply with the actual user's identity information. All these tasks make the PEP an important component that basically builds the connection bridge (gateway) between different endpoints, and allows for secure and trustworthy cross-organizational information exchange.

SLA definition and negotiation (Requirement 13)

Beside the concrete security aspects of the so-called trusted VO, the framework also has to consider legal aspects. For example consumers and the provider of a service need to negotiate electronically the terms of the interaction including potential penalties in case of violation of them. This means that Service Level Agreements (SLA) negotiated must be structured to be legally binding and must go beyond Quality of Service (QoS) terms. Within a dynamic environment also allowing for the involvement of partners not being known at the very beginning, an approach is needed to allow a controlled addition or removal of partners including respective access and user rights for resources (Schubert et al. 2008).

Distributed monitoring and reputation management (Requirement 10, 11 and 12)

Distributed monitoring and reputation management (Yu and Singh 2002) is mainly responsible for benchmarking the individual participants according to their overall performance(s) during the runtime of a collaboration session but also for monitoring different critical components of the entire infrastructure. For that reason, the services interact closely with the SLA infrastructure services, which provide periodical updates on the fulfillment or violation of agreed Quality of Service parameters like reserved versus available storage, free CPU power, availability of services, network bandwidth and a lot more. A continuous observation of all relevant components and immediate notification messages being sent back to the reputation services and the

PEP (to probably enforce a policy update) achieve this. The reputation services then take and convert the measurements into human readable and understandable values to evaluate and determine a partner's reliability, These performance measurements are stored within a so-called "business history" useful to derive someone's trustworthiness in context of signing future (business) contracts or establishing further collaborations with that particular organization. It also allows a general and unreserved view on different partners because such historical values are recorded and managed by third party organizations and not by the business entities themselves, which might provide better results than obviously present.

Ensuring data and resource protection

User authorization and access control (Requirement 1, 4 and 5)

Following the approach of a decentralized AAI, the authorization of users for particular resources constitutes a very important field. Firstly, the final authorization decision should always be taken by the resource's owner and secondly, the management needs to be as easy and flexible as possible. Therefore, we propose to make use of Attribute-Based Access Control (ABAC) policies instead of common Role-Based Access Control (RBAC) policies (compare with (Assel and Kalyoncu 2008)) since the attribute-based solution provides a very detailed specification of corresponding access rules and thus, increases the level of flexibility and granularity. The later issue is extremely important in case of database hierarchies, which facilitate a very fine granular procedure to restrict the access only to a certain number of data sets (one should think here of the database "view" paradigm). Additionally, for improving the applicability and providing more dynamicity to traditional RBAC or ABAC techniques, access policies used within ad-hoc evolving environments such as modern CWEs have to carefully consider user- as well as application-specific context in order to adapt to new customer respectively provider needs seamlessly (Kumar et al. 2002). For evaluating such policies on-the-fly, a dedicated component, a Policy Decision Point is being deployed at each involved entity so that locally created policies can be easily, dynamically and especially securely generated and changed. The PDP is implemented as a web service and directly interacts with the PEP. It makes own decisions based on access control policies stored in its repository by looking up rules that specify whether the user can access that particular resource or service. These rules might contain different sets of attributes as well as further contextual information such as time or location parameters (Demchenkoa et al. 2008). They are expressed in XACML[5] that provides a very general extensible policy language offering lots of flexibility. XACML allows for defining highly detailed policies, which can be used in order to support fine-grained access control down to the level of database tables and single data sets, too (Assel et al. 2009).

[5] eXtensible Access Control Markup Language

Data anonymization and encryption (Requirement 6 and 7)

Cross-organizational data and information exchange respectively demand strong protection mechanisms to securely transfer any kind of data sets, documents or other pieces of information. Traditionally, the owners had to manually anonymize private information (mainly through de-identification of data sets) or extract irrelevant parts as well as had to ensure that all outgoing data is being encrypted accordingly. With an increasing amount of large databases and document repositories as given in many companies, the need for automatic processing of confidential data (not necessarily private data sets but also any kind of business secrets) arises. Hence, we propose the usage of Privacy-Enhancing Techniques (PETs) originating from the eHealth domain (De Moor et al. 2003). PETs are defined as a coherent system of ICT[6] measures that protect 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 information system. Currently, several approaches that present different privacy-preserving methods are in use and under research in the e-Business and e-Commerce areas, too (Smith and Shao 2007). However, in order to create a flexible solution that fits with requirements of dynamic CWEs, PETs should be integrated with the entire infrastructure to allow for their ad-hoc enforcement according to defined security policies. Therefore, the PEP has to determine whether an actual request for data requires specific privacy-preserving actions (an existing policy defines the level of data sensitivity and necessary steps). In addition, this policy should also contain the level of data encryption (on top of the basic encryption standard, e.g. TLS [7]or SSL[8] for network connections) prior to delivering the result set(s).

Data provenance (Requirement 9)

Data Provenance describes the origin of a piece of data, i.e. it answers the question what other pieces of data contributed to a given result (Buneman et al. 2000). However, from the database perspective, the piece of data is a result of a database query while in distributed collaborations or workflows it is a result of multiple queries and other processes respectively. Hence, data provenance is not limited to a single request or action but a sum of various individual and inhomogeneous steps. It plays an important role not only for the administrators but also for end-users to relocate particular bits of data and/or repeat certain previous steps due to verifications of (unexpected) results or even in case of failures. To cope with all these complex requirements, data provenance needs to be vertically integrated along the entire security stack including resources/services on the lowest level and going up to individual end-user applications. Furthermore, it must be centrally controlled by dedicated components of the execution framework in order to provide uniqueness of traced data sets and to ensure their reusability. However, before building a system that allows for tracking any process, data or information flow, a sophisticated

[6] Information Communication Technologies

[7] Transport Layer Security

[8] Secure Sockets Layer

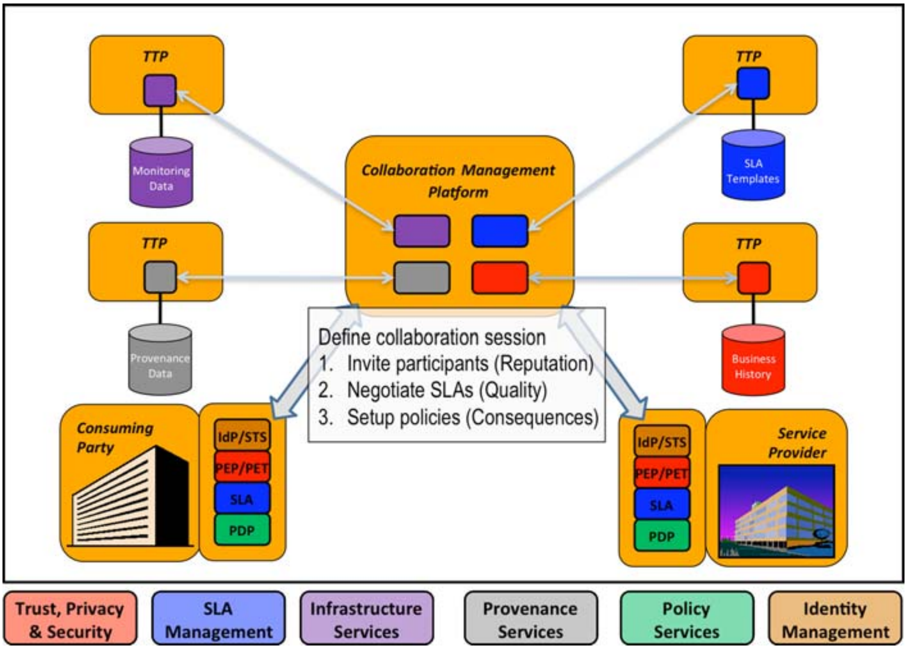

Fig. 5 Architectural snapshot of the CWE security framework

semantic relationship model has to be set up and applied throughout the overall infrastructure. This model must be flexible enough to annotate any piece of data using powerful and extendable domain-specific ontologies.

Exemplary architecture

In order to depict the previous explanations and exploit them onto a concrete exemplary architecture, Figure 5 shows the aforementioned technologies and tries to visualize the main components and basic interactions respectively.

In summary, our model integrates current standards for exchanging identity information between security domains like SAML and WS-Security[9] together with well-established policy languages such as XACML or PDL[10]. This will ensure that the entire framework is fully compliant with existing standards and provides a certain level of sustainability for both users and providers.

Conclusion and outlook

The security framework for the new kind of collaborative working environments with their dynamic and ad-hoc nature presented in this paper is reflecting the need of

[9] Web Services Security
[10] Policy Description Language

cross-organizational collaborations not only for the engineering and eHealth domain as detailed but cover a much wider range of application domains. The model presented here is already covering a broad range of applicability but is not targeting to support scenarios where a very large number of participants are collaborating but is focused on a smaller (in terms of number of concurrent users) business-to-business area where the participants are bound to company policies but at the same time can rely also on an existing IT infrastructure enabling the collaboration.

Future research is needed in order to address more open and larger collaboration scenarios. Larger scenarios would need to empower the individual user even more which information she wants to share in the given context. This would include not only the access right problem on data but also the shared data itself. Parts of data is not necessary to be known to the collaboration partner in the given context and should not be visible or accessible whereas at a later stage the same collaboration partners would need to exchange this information. In this view such models require the security models to be adaptable to the changing roles and context of a collaboration leveraging security from the infrastructure level up to the application level.

Acknowledgements The results presented in this paper are partially funded by the European Commission under contract IST-5-034245 through the project CoSpaces and through the support of the ViroLab Project Grant 027446. The authors want to thank all who contributed to this paper, especially the members of both consortia.

References

Assel A, Kalyoncu O. Dynamic access control management for distributed biomedical data resources. In: Cunningham P, Cunningham M, editors. Collaboration and the knowledge economy: issues, applications, case studies. Amsterdam: IOS; 2008. p. 1593–9.

Assel M, Kalyoncu O, Pan Y. Approaching fine-grain access control for distributed biomedical databases within virtual environments. In: Bubak M, Turala M, Wiatr K, editors. Proceedings of the 8th Cracow grid workshop. Kraków: ACC Cyfronet AGH; 2009. p. 311–9.

Assel M, Kipp A. A secure infrastructure for dynamic collaborative working environments. In: Arabnia HR, editor. Proceedings of the 2007 international conference on grid computing and applications. Las Vegas: CSREA; 2007. p. 212–6.

Barton T, Basney J, Freeman T, Scavo T, Siebenlist F, Welch V, et al. Identity federation and attribute-based authorization through the globus toolkit, Shibboleth, Grid- Shib, and MyProxy. Proceedings of 5th annual PKI R&D workshop, Gaithersburg, USA: 2006.

Bhargav-Spantzel A, Squicciarini AC, Bertino E. Establishing and protecting digital identity in federation systems. J Comput Secur. 2006;14(3):269–300.

Buneman P, Khanna S, Tan WC. Data provenance: some basic issues. In: Kapoor S, Prasad S, editors. Lecture notes in computer science 1974. Berlin: Springer-Verlag; 2000. p. 87–93.

De Moor GJ, Claerhout B, De Meyer F. Privacy enhancing techniques—the key to secure communication and management of clinical and genomic data. Methods Inf Med. 2003;42(2):148–53.

Demchenkoa Y, Mulmob O, Gommansa L, de Laata C, Wana A. Dynamic security context management in grid-based applications. Future Gener Comput Syst. 2008;24(5):434–41.

Forestiero A, Mastroianni C, Spezzano G. Dissemination of information with fair load distribution in self-organizing grids. In: Dorigo M, Birattari M, Blum C, Clerc M, Stützle T, Winfield AF, editors. Lecture notes in computer science 5217. Berlin: Springer-Verlag; 2008. p. 291–8.

Geuer-Pollmann C, Claessens J. Web services and web service security standards. Inf Secur Tech Rep. 2005;10(1):15–24.

Gubala T, Balis B, Malawski M, Kasztelnik M, Nowakowski P, Assel M, et al. Virtual laboratory for development and execution of biomedical collaborative application. In: Proceedings of the 21th IEEE international symposium on Computer-Based Medical Systems (CBMS 2008), Jyväskylä, Finland: 2008. p. 373–8.

Khurana H, Gligor VD. A model for access negotiations in dynamic coalitions. Enabling Technologies, IEEE International Workshops on, 2004. p. 205–10.

Kipp A, Schubert L, Assel M. Supporting dynamism and security in ad-hoc collaborative working environments. In: Callaos N, Lesso W, Zinn CD, Baral J, editors. Proceedings of the 12th world multi-conference on systemics, cybernetics and informatics. Orlando: 2008. p. 259–63.

Kumar A, Karnik N, Chafle G. Context sensitivity in role-based access control. SIGOPS Oper Syst Rev. 2002;36(3):53–66.

Saabeel W, Verduijn TM, Hagdorn L, Kumar K. A model for Virtual Organisation: a structure and process perspective. eJov 2002;4:1–16.

Schubert L, Wesner S, Dimitrakos T. Secure and dynamic Virtual Organizations for business. In: Cunningham P, Cunningham M, editors. Innovation and the knowledge economy: issues, applications, case studies. Amsterdam: IOS; 2005. p. 1201–8.

Schubert L, Kipp A, Koller B. Supporting collaborative engineering using an intelligent web service middleware. Adv Eng Informatics. 2008;22(4):431–7.

Sloot PMA, Altintas I, Bubak M, Boucher CA. From molecule to man: decision support in individualized e-health. IEEE Comput Soc. 2006;39(11):40–6.

Smith R, Shao J. Privacy and e-commerce: a consumer-centric perspective. Electron Commerce Res. 2007;7(2):89–116.

Wasson G, Humphrey M. Policy and enforcement in Virtual Organizations. Grid Computing, IEEE/ACM International Workshop on, 2003:125.

Wesner S. Towards an architecture for the mobile grid. IT-Inf Technol. 2005;47(6):336.

Wesner S, Kipp A. Report on BE classification and recommendations for architecture and interoperability. BEinGrid Whitepaper. 2007.

Wilson MD, Chadwick D, Dimitrakos T, Döser J, Giambiagi P, Golby D, et al. The TrustCoM approach to enforcing agreements between interoperating enterprises. Proceedings of Interoperability for Enterprise Software and Applications (I-ESA'06), Bordeaux, France: 2006.

Yu B, Singh MP. Distributed reputation management for electronic commerce. Comput Intell. 2002;18 (4):535–49.