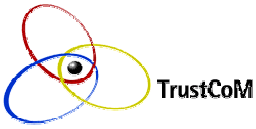

Deliverable

22

TrustCoM Roadmap

Scientific and Technological Assessment of the first phase of the project & Scientific and Technological Roadmap for the remaining of the project

WP12 S&T Roadmap

Theo Dimitrakos

Scientific coordinator, TrustCoM project

BT Group, Chief Technology Office

5 September 2005

Version 2

TrustCoM

A trust and Contract Management framework enabling secure collaborative business processing in on-demand created, self-managed, scalable, and highly dynamic Virtual Organisations

SIXTH FRAMEWORK PROGRAMME

PRIORITY IST-2002-2.3.1.9

Prepared for the European Commission under FP6 Contract No. 01945 as a deliverable form

LEGAL NOTICE

The following organisations are members of the TrustCoM Consortium:

Atos Origin,
Council of the Central Laboratory of the Research Councils,
BAE Systems,
British Telecommunications PLC,
Universitaet Stuttgart,
SAP AktienGesellschaft Systeme Anwendungen Produkte in der Datenverarbeitung,
Swedish Institute of Computer Science AB,
Europaeisches Microsoft Innovations Center GMBH,
Eidgenoessische Technische Hoschschule Zuerich,
Imperial College of Science Technology and Medicine,
King's College London,
Universitetet I Oslo,
Stiftelsen for industriell og Teknisk Forskning ved Norges Tekniske Hoegskole,
Universita degli studi di Milano,
The University of Kent,
International Business Machines Belgium SA .

Deliverable datasheet

Project acronym: TrustCoM

Project full title: *A trust and Contract Management framework enabling secure collaborative business processing in on-demand created, self-managed, scalable, and highly dynamic Virtual Organisations*

Action Line:	**7**
Activity:	**7.1**
Work Package:	**12**
Task:	**12.1-4**

Document title:	**Progress Assessment and S&T Roadmap**
Version:	**1**
Document reference:	
Official delivery date:	**31 July 2005**
Actual publication date:	
File name:	**TrustCoM.D22.SnT.Roadmap.V2**
Type of document:	Report
Nature:	Restricted

Authors:	Theo Dimitrakos (BT) based on input from Alvaro Arenas (CCLRC), Philip Robinson (SAP), Johen Haller (SAP), Lutz Schubert (HLRS), Pablo Giambiagi (SICS), Christian Geuer-Pollmann (EMIC), David Chadwick (U. Kent), Emil Lupu (ICSTM), Babak Sadighi (SICS), Tomas Garcia (ATOS), David Golby (BAE Systems), Joris Claessens (EMIC), Jakka Sairamesh (IBM), Tobias Mahler (NRCCL), Fredik Vraalsen (SINTEF).
Reviewers:	David Golby (BAE Systems), , Emil Lupu (ICSTM), Stefan Wesner (HLRS)
Approved by:	

Version	Date	Sections Affected
0.1	May 2004	First version of the S&T Roadmap.
0.2	August 2004	Minor update
0.3	November 2004	Major restructuring – all sections affected
0.4	February 2005	Major restructuring – all sections affected
0.5	April 2005	Minor update
0.6	July 2005	Major update – “self-assessment” and “recommendations” sections affected (sections 7 & 8 in v0.6).
0.7	August 2005	Minor update / aesthetic improvements
0.8	September2005	Undergone internal reviews and QA procedure
0.9	September2005	Minor update
1.0	September2005	Final version delivered

Table of Content

1 About this document

This document is the main deliverable of the TrustCoM WP12: scientific & technological roadmap. Work started at the beginning of the project (month 1) and will last until month 18, covering the whole duration of this implementation plan. Its main objectives are

1. to periodically revisit, update and refine the research challenges faced by the TrustCoM Consortium, taking into consideration the interests of the TrustCoM Consortium, progress within the project, and the advancements achieved outside of TrustCoM
2. to produce a comprehensive scientific and technological roadmap guiding the research advancements and technological innovations expected during the project's implementation towards meeting the identified research challenges
3. to conduct a self-assessment of the project's progress towards meeting its research and technological objectives
4. to offer specific recommendations in order to improve delivery or re-adjust the ongoing technical work towards converging goals and common research objectives

This deliverable does not have the following objectives that are sometimes attributed to general-purpose roadmap documents.

- It does not attempt to analyse the problem space addressed by the project and scope its objectives in relation to that problem space. Although this objective falls within the same group of activities (AL7), it is addressed by a different deliverable, namely **D1: Problem definition**.
- It does not attempt to sketch the state of the art outside of TrustCoM and define how TrustCoM is positioned against the state of the art. Although this objective falls within the same group of activities (AL7), it is addressed by a different deliverable, namely **D2: State-of-the-art evaluation,** which periodically revisited during the life-time of the project.
- It does not attempt to analyse the relevant Open standards in the areas where TrustCoM is achieving technological advancements, nor does it offer a roadmap towards integrating or extending these standards. This is addressed by deliverable **D24: Roadmap of Technical Standards development**.
- It does not attempt to analyse the exploitation and business opportunities in the areas addressed by TrustCoM or to identify exploitable products within the project and to classify them against identified business opportunities. This is addressed in deliverables **D7: Market Study** and **D25: Outline exploitation plans**.

At the first phase of the project (month 1 to month 18) particular emphasis is placed on achieving the necessary integration that will underpin subsequent scientific advancements and technological innovation. This Roadmap is a "live" document updated as necessary throughout the project and revised at the end of each stage of a project phase.

Some specific questions that are addressed by this deliverable include the following:

- What has changed in the environment since the initiation of the project?
- What is the impact of these changes in the objectives of the overall project?
- What is the recommended reaction of the Consortium to these changes, what can be addressed within the TrustCoM project and what has to be done by other – potentially new – projects?
- For each research challenge initially identified in the technical annex of the project (or earlier versions of this roadmap):
    a. Has the assumption of this challenge changed?
    b. Must it be updated? If so, how?

c. For each change identified: does the change have impact on the project activities targeting this challenge?

d. After the end of the first phase of the TrustCoM project, can it be foreseen that the project will not solve this challenge, but new challenges have appeared?

1.1 Outline of this deliverable

The rest of the deliverable is structured as follows:

Section 2 offers a general introduction where we describe the process of producing this deliverable is explained and the overall impact of the S&T roadmap on the project.

Section 3 provides an overview of the TrustCoM project, the motivation for this research in general and a summary of the high-level research challenges.

Section 4 analyses, decomposes and refines the research challenges summarised the previous section and shows how these are addressed by each specific technical activity of the TrustCoM project.

Section 5 summarises a comparison of the updated research objectives and interim results of the TrustCoM project to the current state of the art in commercial products and applied research

Section 6 offers a graphical representation of the projected timescales of the TrustCoM innovation from conception to market penetration.

Section 7 summarises a self-assessment of the project and places emphasis on the advancement achieved so far compared to the current trends in each relevant area.

Section 8 provides recommendations to each technical activity, in order to ensure convergence and leverage on all relevant interim results in preparation for the second phase of the project.

2 Introduction

The TrustCoM project [http://www.eu-TrustCoM.com/] is developing a framework for trust, security, and contract management for secure, collaborative business processing and resource sharing in dynamically-evolving Virtual Organisations. An overview of the motivation, targeted application domains, and of the scientific and technological objectives of the project is described in chapter **Error! Reference source not found.** of this deliverable. The term "TrustCoM Framework" stands for the principles and paradigms, the processes and functions, and the architecture and the technology that underpin trustworthy, secure, and contract-driven operations of Virtual Organisations.

The purpose of this document is to provide a coherent overview of the scientific and technological objectives of the TrustCoM project, to highlight main research results, and to update research challenges and associated technical goals, and to provide the foundation for research and technological development work to be conducted in the second phase of the project.

2.1 Towards the TrustCoM scientific & technological roadmap

This document is the main deliverable of the TrustCoM WP12: scientific & technological roadmap. Its main objectives are to set the main research challenges of the project, to produce a comprehensive scientific and technological roadmap guiding the research advancements and technological innovations expected during the project's implementation towards meeting the identified research challenges, and to conduct regular progress assessments in order to re-adjust targets and focus of the work. D22 covers an assessment of M1-M18 and offers recommendations for the remaining of the project.

During the first phase of the project (month 1 to month 18), a particular emphasis is placed on achieving the necessary integration that would subsequently result in scientific advancements and technological innovation. The Roadmap is a "live" document, updated as necessary throughout the project and revised at the end of each stage of a project phase.

The process for the development and update of TrustCoM scientific and technological roadmap consisted of the following steps:

1. Validation of the research challenges by the project consortium and the associated communities. This included steering the work conducted in the following tasks:
    a. *Selection of a number of targeted application domains* and analysis of several scenarios in order to identify the main issues relating to the security, trust and contract management across various Virtual Organisation settings. Scenarios in selected areas have been analysed in order to validate the research challenges, on one hand, and to inform the scientific and technological objectives on the other. The analysis was conducted by WP11 during the first quarter of the project and results of the analysis have been documented in deliverable ***D3****: Case study scenarios*. The scenarios pursued in the remaining of the project amalgamate elements that have been identified as critical by this analysis.
    b. *Analysis and evaluation of the state of the art*. One of the intrinsic characteristics of all projects dealing with ICT infrastructures and technologies for Virtual Organisations is their dependence if a large number of potentially diverse enabling technologies. From early on in the TrustCoM project we tried to make sense of this large technology jigsaw and identify what could be leveraged upon, what had to be improved and what had to be developed from scratch in the context of this project or in collaboration with a wider community. The analysis was conducted mainly in the first quarter of the project within WP10 and the results have been documented in ***D2****: State-of-the-art-Evaluation*. This evaluation has been very informative and covered an unprecedented number of the

technologies that have not been analysed before by the same team and in a common context.

c. *Identification and analysis of a set of open standards* that can be used as a foundation for the TrustCoM framework or relate to specific aspects of this framework. A preliminary standardisation roadmap was produced in the first year of the project (D6) and has been refined in ***D24****: Roadmap of technical standards development*. Standards cover areas from service management to identity management and federation, and from messaging to business processing. Similarly to D2, the number and complexity of the standards that have been analysed, classified and experimented with, has been unprecedented for a research project, and the knowledge generated has been particularly useful for understanding how ICT for Virtual Organisations can be designed and developed.

The combined outcome of the work lead into a revision and subsequent refinement of the initial project objectives that was in turn channelled into the two main action lines of the project: AL1 that focuses on conceptual models and architecture and AL2 that focuses on detailed design and reference implementation of key ICT services and components in the areas of security, trust and contract management for dynamic Virtual Organisations.

2. Communication, validation and revision of initial challenges via specific outreach activities. In particular, the TrustCoM project

   a. Ensured that the initial challenges and results of deliverables D2, D3 and D6 have been extensively discussed and accepted among all Consortium partners and in particular software vendors and end-users.

   b. Organised a series of detailed tutorials and panel discussions during the 2nd international conference on Trust Management (iTrust 2004),[1] in April 2004 where the targeted application domains, evaluated technologies and relevant standards were discussed with the community.

   c. Organised a panel discussion for feedback on initial results during a workshop at the 18th IFIP World Computing Congress in August 2004.[2]

   d. Organised two workshops with extensive presentations of research challenges, plans and results of the evaluation during the eChallenges conference in October 2004 [3]

   e. Participated at the DG INFSO Enterprise Interoperability Cluster and the Grid Concertation events.

3. Clarification of the main aspects that may appear in a “blue print” of the TrustCoM framework. This was achieved by steering the work in Action Line 1 of the project that focused during the first year of the project on providing some basic conceptual models and architecture for such a blue print and by relating the interim results with the work in Action Line 2 which focused during that period on experimenting with enabling technologies on diverse platforms (Java and .NET based Web Services) in order to assess the feasibility of prototyping these aspects. The results of this exercise lead into three main results:

---

[1] The first three International Conferences on Trust Management have been supported by the iTrust working group, which was a FP5 Thematic Network funded under the FET programme of IST. The iTrust network continues its operation after the end of the FP5 project and currently brings together over *100 researchers* from over *60 institutes* in Europe, America, Australia and the Far East. For more information on the FP5 project see http://www.itrust.uoc.gr. For information on the TrustCoM events see http://www.trustmanagement.clrc.ac.uk/ and http://www-rocq.inria.fr/arles/events/iTrust2005/ for the 2004 and 2005 events respectively.

[2] See http://www.wcc2004.org and http://www.wcc2004.org/congress/workshops/ws4.htm

[3] See http://www.echallenges.org/2004/ for the eChallenges event. The Workshops were 3e and 4e: “Towards a Trust & Contract Management Framework for Dynamic Virtual Organisations” 1 & 2. See also http://www.echallenges.org/2004/PDF/Workshop_3E.pdf and http://www.echallenges.org/2004/PDF/Workshop_3E.pdf

a. The identification of six subsystems for the TrustCoM framework blueprint. These are summarised in Figure 1 and described in more detailed in section 4.

b. The analysis of the dependencies between the services and the info-sets in each subsystem in order to ensure their best distribution in self-coherent and loosely coupled groups.

c. The recommendation to the TrustCoM project management to radically change the project implementation plan of the main action lines under which research and technological development work is conducted (i.e. AL1 and AL2) in order to achieve a better alignment with the above structure. Following this recommendation the Consortium performed a *major project restructuring* at the end of the first year of the project.

   This restructuring effectively transformed the project from a generic "horizontal" structure that was looking into the "conceptual models", "architecture", "interoperability profiles" and "tools & methods" of the overall TrustCoM Framework, as depicted in Figure 1, into the more specific "aspect-driven" structure depicted in Figure 2 where technical workpackages directly correspond to the subsystems that have been identified. Each of these workpackages focuses on analysing, designing and implementing the key functionalities expected by each subsystem. For convenience, the division between modelling and prototyping activities has been maintained at a high-level although all the modelling and prototyping activities are now closely aligned for each aspect of the TrustCoM framework.

4. At a second phase, we expect to validate our interim results and revised objectives, resulting from the above. This second phase of validation involves

   a. A detailed tutorial at the 3rd international conference on Trust Management (iTrust 2005), in April 2005 where an update of the research challenges in view of the interim project results was discussed.

   b. A detail lecture and tutorial on interim findings and project plans at the FOSAD international post-graduate school on Foundations of Security Analysis and Design, September 2005.

   c. Validation and further input from externals in

      i. DG INFSO Enterprise Interoperability

      ii. The Global Grid Forum

      iii. The project's technical advisory board

      iv. Other selected related initiatives in Europe and worldwide

The purpose of this second phase evaluation is to complete and update the research challenges based on the recent project evolution and experiences. An assessment of the project achievements will be provided in the final update of this deliverable towards the end of the project.

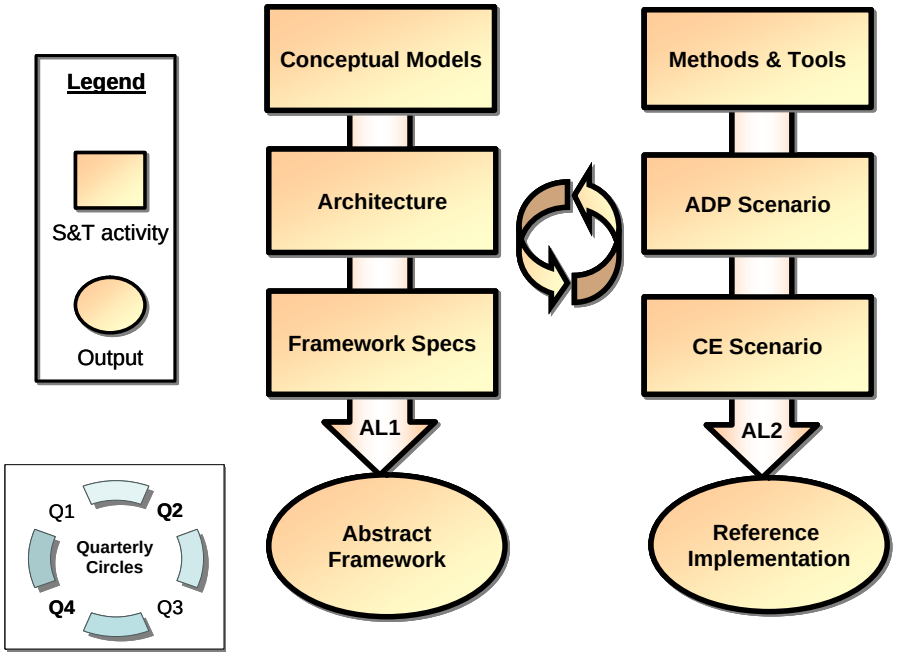

Figure 1: Structure of the project technical core during the first year of the project

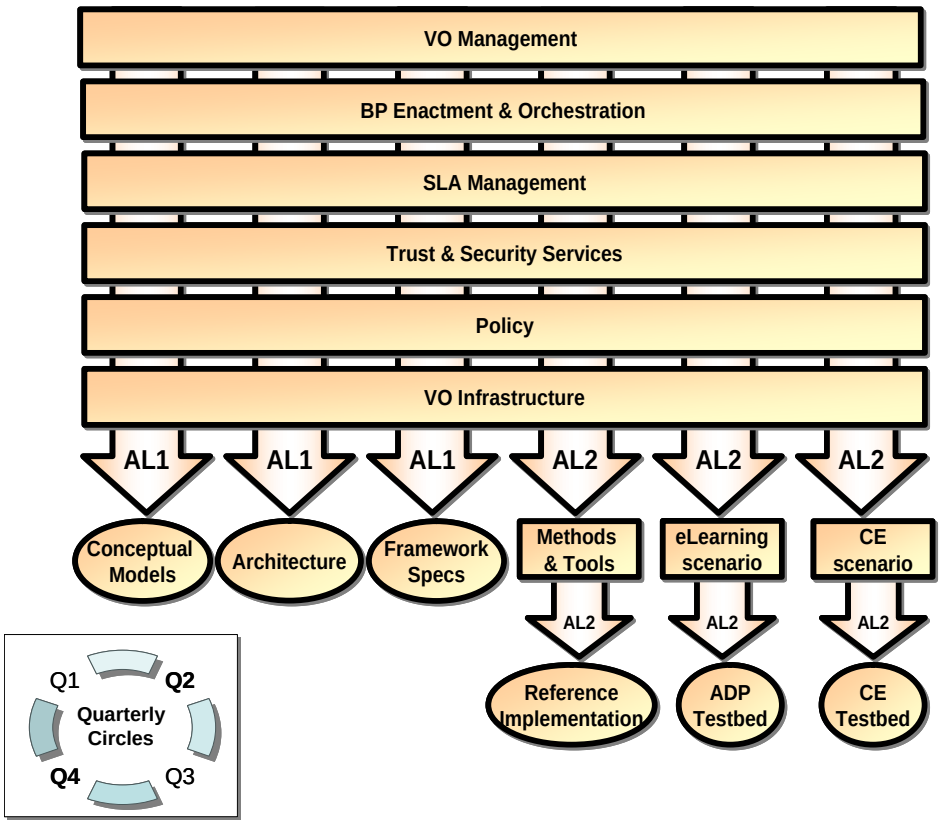

Figure 2: Current structure of the project technical core

3 Main Research Challenges and Project Scope

Recent years have seen an unprecedented acceleration in the evolution of the Internet as the technological vehicle underpinning the expansion of service provision and inter-/intra- enterprise integration in all market sectors. This brings about the prospect of *ad hoc* integration of systems across organisational boundaries to support collaborations that may last for a single transaction or evolve dynamically over many years. This sets new requirements for scalability, responsiveness and adaptability that necessitate the on-demand creation and self-management of dynamically evolving virtual organisations (VO) spanning national and enterprise borders, where the participating entities (enterprises or individuals) pool resources, information and knowledge in order to achieve common objectives. The objectives may be short term - e.g. to deliver an one-off service in response to a specific customer demand - or long-lasting. In the latter case, the VO's structure, business processes and operational infrastructure must adapt as the goals of the collaboration, the participating entities, the business context and the technologies employed, change.

Emerging ICT paradigms such as Autonomic computing, Utility computing and Grid computing are making the formation and operation of virtual organisations easier by providing dynamic management of the distribution of computational processes across available resources. However, the malleability of the digital medium that makes this possible is also a liability: a major limiting factor is a well-founded concern about exposure to fraud or misuse of the technology. Today, concerns about trust and security are acknowledged to be significant barriers to providing access to outsiders. In spite of the major ICT breakthroughs of the last two decades, protecting one's assets while integrating services, processes and resources, remains a major ICT challenge. Overcoming such challenges requires the development of disruptive technology realising innovative ideas over widely acceptable interoperable platforms. The required scalability, responsiveness and adaptability for on-demand created and dynamic virtual organisations, makes the provision of *cost effective* trust and contract management solutions for VO environments, *the* most demanding and timely research challenge in this field. Effective solutions require interdisciplinary approaches integrating tools from law, cognitive and social science in addition to telecommunications and computing. The successful deployment of *secure* and *trusted dynamic VOs* requires converging strategic research at a European level, coupled with mechanisms for integration of existing experimental results and the rapid dissemination, realisation and take-up of new research outputs.

3.1 Main outputs of the TrustCoM project

In response to this challenge, the European Commission and a consortium of end-users, major software vendors and telecom operators, national research institutes and Universities, are implementing the new Integrated Project TrustCoM. TrustCoM conducts multidisciplinary research in order to deliver:

1. *A novel trust and contract management reference architecture* that will enable collaborative work within on-demand created and self-managed dynamic VOs leveraging on the emerging convergence of Web Services and Grid technologies.
2. A set of *conceptual models* explaining the fundamental concepts, principles and methods underpinning the above architecture. Effectively these provide the meta-model of any new architectural constructs that may result from TrustCoM research.
3. A set of *profiles,* that bring together and potentially extend selected Web/Grid Services specifications at specific version levels, along with conventions about how they work together to support potential implementations of the TrustCoM framework.
4. *A reference implementation of the above* integrating and extending already established or emerging interoperability standards for autonomic security, trust and contract management based on Web and Grid services technology.

5. *System and software engineering tools and methods* analysis of the VO life-cycle and offering a library of design patterns and generic software components implementing selected services that offer the core functionalities of the VO.
6. *Testbeds* exhibiting instantiations of the above architecture and reference implementation into two classes of realistic application scenarios, namely collaborative engineering (CE) and provision of ad-hoc aggregated services (ADP).
7. *Selected demonstrators* exhibiting the business value and benefits of the TrustCoM framework in the abovementioned application domains.
8. *Studies analysing selected aspects of the legal and socio-economic context* that underpins such Virtual Organisations.

3.2 Examples of Virtual Organisations

3.2.1 Virtual Organisations in Collaborative Engineering

The development, production and support of modern products such as ships, aircraft etc are highly complex processes that often involve great risk. Principal risks include technical complexity (both in the complexity of products and processes) and changing customer and market requirements. The ability to manage these and other risks is a distinguishing feature of competitive organisations in the engineering sector. A strategy for managing this complexity is to form partnerships or Joint Ventures (JVs) in order to exploit new markets and opportunities through Collaborative Engineering (CE). In a JV partners focus on particular aspects of the product through its lifecycle, enabling more focus on core business capabilities. Emerging technologies such as web and grid computing may facilitate the evolution of JVs into Virtual Organisations (VOs), where organisations quickly come together to share resources without requiring the development of new facilities and systems - a common feature of JVs at present. The CE scenarios described here attempt to cover most of the phases of the product lifecycle within a CE VO through development, production and in-service product upgrade.

In summary, three scenarios have highlighted the importance of effective and flexible security system for building confidence in the extensive and more integrated collaborations that VOs offer over conventional JVs. The security policies should also be correlated both with the collaborative agreements established between partners at the business level and with agreements established within other collaborations as well. The benefits from an effective security and contract management framework are the ability for engineering collaborations to be quickly reconfigured in order to expose the assets that need to be shared to achieve the business goal. Service level agreement monitoring is important for ensuring that suppliers (of components, services etc) perform according to contracts. Benefits here possibly include the automation of processes between clients and suppliers that are usually repetitive. Finally, trust frameworks are required for supporting collaborations. The first of these concerns is for managing the reliability and traceability of engineering data, ensuring that greater confidence can be given to it and that it can be relied upon in the major engineering tasks. The second of these Trust frameworks should facilitate the search for new partners/suppliers of components or services that were previously unknown to the VO. This should include some assessment of the trustworthiness of the security systems and its security policies. It has been recommended that TrustCoM focuses initially on the latter area, as a higher priority, where the technology and methods investigated by the Consortium can have a stronger impact, and then address the former as a lesser priority. See also sections 7.9.1 and 8.9.1 for more information about the current state and plans for our testbed in this area.

3.2.2 Virtual Organisations for Next Generation Service Providers

We are interested here in VOs that are formed through ad hoc aggregation of component services offered by different service providers. Increasingly, enterprises are using web services and related technologies to provide their customers, suppliers and partners with direct access to their services and business processes. Motivations include reducing costs and speeding up processes through

automation. However, the vision behind the web services / service oriented architecture revolution is that distributed applications can be assembled as needed by connecting together pre-existing services. Selection of the services to use takes place through a 'discovery' process. As well as connecting the services together into a supply chain capable of fulfilling a customer order, the business process of the enterprises involved must also be interfaced. Furthermore, contracts need to be agreed establishing the mutual rights obligations of the participating service providers. When connections at these three levels can be established on demand, we can truly say we have an ad-hoc dynamic VO.

We are already seeing services being 'disaggregated', that is, in addition to offering 'complete' services, simpler constituent services are offered separately. Other organisations can then make use of these constituents in combination with their own service elements to offer composite services to their customers. Motivations for disaggregating include regulatory / anti-trust factors, advantages arising from focus on core competences, business agility (ability to launch new services / enter new markets rapidly), a desire of a part of the individual SPs to retain the advantages of small scale (or conversely to avoid the overheads and inertia of large organisations. New services may also be created specifically for use as constituents of larger services offered by other enterprises. This could offer opportunities for specialist start-up companies to enter the market. Benefits of a dynamic aggregation include provision of services that are precisely tailored to a specific customer need. The need to offer a wide range of tailored services could arise from a wide range of preferences or requirements among the targeted customer base, or because the specifics of the service depend on the circumstance of the customer, e.g. current location, the task currently being undertaken, and other context specific variables. The ability to participate in dynamic VOs greatly increases the range of services a provider can offer to its customers, and also the number of end-customers it can reach indirectly via partners.

Five such 'Aggregated Services' (AS) scenarios have been defined and analysed as part of the TrustCoM problem definition activity. In summary, the five scenarios have highlighted that dynamic VOs inevitably incur a management overhead compared to real organisations, and indeed to static VOs (formal consortia). There is a requirement for additional services to provide the glue that enables the VO to function as a viable entity e.g. to provide overall coordination of activities while retaining flexibility. We expect that these services can be defined in such a way that they are basically independent of the particular application domain. Furthermore, there is a requirement for services to replace the trust inherent in operation within an integrated real organisation (trust in colleagues even when not known personally, trust in procedures and processes, etc.), and the trust between customer and an established service provider with a clear legal identity and brand / reputation. This last class of service is a main ingredient of the TrustCoM Framework. Without such a framework, it is likely that enterprises will judge that the risks in participating in dynamic VOs will out-weigh the benefits. Similarly, end-customers will be reluctant to buy from dynamic VOs. It should also be recognised that there are substantial commercial opportunities for enterprises offering the trust, security and contract management services instantiating the TrustCoM framework. The TrustCoM project will prototype the implementations of potentially useful classes of service, drawing on the scenarios mentioned above for the requirements.

Following the analysis of the five aggregated services scenarios, a presentation of alternatives, and advice from the TrustCoM project reviewers, the Consortium decided to select an AS scenario in the area of eLearning. This scenario tackles the full life-cycle of creating communities of eLearning service and content provision and the process-driven integration of these into an aggregate service that follows a personalised learning path.

3.3 Main research Challenges and Anticipated Innovation

TrustCoM aims to develop a coherent framework (architecture, services descriptions, and interaction protocols) that provides means of achieving:

1. *Establishment of trust* relationships by means of digital identities, certification, reputation, and inspection to ensure the security, dependability and competency of the business partners,

2. *Autonomic security*, including the specification, automated management and enforcement of policies controlling fine-grained access to the services and resources contributed by the VO constituents and assuring confidentiality / privacy, integrity, availability and accountability at VO level, while self-adapting to contextual changes within the VO.
3. *Contracting*, focusing on the provision of trusted services to support the management of electronic contracts, the incorporation of guarantees to facilitate trustworthy collaboration, and performance assessment at the enactment of electronic contracts (in particular those related to SLAs).
4. *Business Process Enactment*, focusing on securing the enactment of collaborative business processes invoking services and consuming resources contributed by the VO partners in compliance with their security policies and agreements. Emphasis is also placed on self-adaptation of the business process enactment in response to contextual changes within the VO, including changes to the VO membership, security policy or agreements.

3.3.1 Specific areas of innovation

The development of a coherent trust & contract management framework enabling the on-demand formation and self-management of secure, scalable, highly dynamic, integrated and targeted Virtual Organisations, that share services, resources, information and knowledge across enterprise boundaries is *the overall research challenge* faced by the TrustCoM project. It requires the integration of innovations from trust management, enterprise security management, and contract management into a mutually reinforcing overall solution covering all phases of the VO life cycle, as summarised in Table 1.

VO life cycle	Main areas of innovation		
	Trust & Security	Contract	Collaborative Process
Identification	discovery & justified identification of credible, trusted partners	elicitation of contractual requirements	definition of VO objectives, elicitation of process goals and requirements
Formation	establishment of trust between perspective VO members	identification, negotiation and endorsement of collaboration agreements between VO partners	process definition (overlaying trust information), engagement of collaborators, optimisation of resource utilisation
Operation	maintenance of trust, autonomic security management, adaptive deployment of security policies	contract enforcement, performance monitoring, arbitration & contract amendment	adaptive enactment of collaborative processes, trust-based decision making, secure service orchestration, dynamic service invocation, accounting
Dissolution	termination of trust relationships & maintenance of trust knowledge	nullification of contracts, posterior analysis	resource disengagement, posterior analysis

Table 1: Overview of main areas of innovation classified against phases of the VO life cycle

1. *Advances in contract management* approaches are required to automate the negotiation, the validation and amendment of collaboration agreements (formalised by means of electronic contracts). They will facilitate the operation of electronic contracts by defining the context within which business processes enact, and by providing a description against which any deviations from the expected norm (or non-compliance to a collaboration agreement) can be identified and assessed and by defining counter-actions corresponding to such deviations when appropriate.

2. *Autonomic security management* can facilitate the operation of VOs by providing the adaptability and the required responsiveness to contextual changes in order to ensure that assets remain protected in the volatile environment of a dynamic VO. They can facilitate the dissolution of a VO (or the disengagement of partners from operational VOs) by ensuring that shared resources are released and that access conditions return to their pre-VO norm.
3. *Advances in trust management* are required to enable network entities with or without established relations to inform their choice of collaborators by assessing their trustworthiness (often assisted by mediators who supply information or knowledge about the trustworthiness of an entity in different contexts) for undertaking a specific task or offering a service. Advances also require the development of new trust models that address the complexity of interdependent interactions in dynamic VOs, methods for collecting and propagating evidence and trust-based decision making mechanisms.
4. *Integration*: Although innovation in any of the above areas constitutes in itself a significant contribution to Information Society, the *added value of integration* is enormous, providing:
    a. A balance between the significance of the business process goals, the expected competence of the contributors, the required level of protection of shared assets, and the terms of collaboration, in relation to the VO objectives.
    b. An optimal selection of VO members, based on the goals of the collaborative business processes they will contribute to, their competence for the tasks assigned to them, the policies defining a partner's own terms of involvement, and contracts expressing the mutually accepted context in which collaboration takes place.
    c. A sustainable coherence between the efficiency gained by relying on an entity's competence to perform a delegated task, the need to sufficiently protect one's assets (especially when opening-up to collaborators), the necessity to perform and adapt within the boundaries set by potentially incomplete mutually accepted agreements, and the need to take decisions on-the-fly about which task to assign to whom in order to respond in a timely manner to a business opportunity.
    d. Continuity and sustainable quality in service provision within VO ecosystems, where evolution is characterised by frequent changes of variable force in the organisational context and short period of relative stability. To ensure that such changes do not damage the equilibrium of complex collaborations between potential competitors, one has to ensure rapid responsiveness to sudden changes in trustworthiness, the ability to swiftly renegotiate and amend agreements and to accordingly adjust security policies, and their enforcement mechanisms.

We expect that some of the enabling technologies are being or will be produced by other projects and initiatives. In such cases, TrustCoM focuses on innovation in terms of holistic integration.

Research and technological innovation in the above themes will be informed by analyses investigating the legal and socioeconomic context of VOs:

5. *Socio-economic Context*. Based on an empirical analysis of the market needs, TrustCoM aims to develop new socio-economic models underpinning the establishment of digital economies within which VOs can evolve and generate profit. These will identify methods for creating incentives for engaging in trustworthy electronic collaborations and sharing services, resources information and knowledge within VOs in order to achieve common objectives in a way that multiplies their productivity and allows for the achievement of results that participants could not produce on their own.
6. *Legal Context*. TrustCoM will study selected legal and regulatory issues of collaborative work in VOs, focusing on privacy, data protection, and international issues. Analysis will also assess the expected impact of technological innovation in light of these issues and some legal and regulatory factors that could influence its exploitation.

4 Refinement and decomposition of the main technical research challenges

During the first phase of the TrustCoM project the main technical research challenges summarised in section 3.3 have been revisited and further refined. In this section we summarise the outcome of this process which initiated the overall restructuring of the technical part of the project in its current form.

The first three of following specific research challenges (RC1-RC3) address the overall objective of the TrustCoM project, while the remaining three research challenges (RC4-RC5) correspond to a revision of the specific project objectives in the areas of contract management, trust and security and collaborative processes.

In each case we position the research challenges against emerging solutions that have been identified outside of TrustCoM and, where appropriate, relevant open standards technologies upon which the TrustCoM framework is based.

4.1 RC1: General Objective

The general objective of TrustCoM is to develop a framework for trust, security and contract management in dynamically-evolving virtual organisations. The framework will enable secure collaborative business process management and sharing in an on-demand, self-managed, dynamic value-chains of businesses and governments. The framework will leverage and extend the emerging convergence of open-standards such as Web Services, Grid technologies and protocols for inter-enterprise interactions (using open agent protocols).

In addition to objectives in the specific areas of trust, security and contract management for Virtual Organisations, this general objective entails specific research challenges relating to the development of a common ICT infrastructure underpinning the formation, operation, adaptation and dissolution of Virtual Organisations, as well as a collection of basic services to manage the life-cycle of Virtual Organisations.

In the following subsections we first summarise specific objectives that have been identified to these research challenges and then we proceed to objectives that relate to the more specific areas of trust, security and contract management in this context.

4.2 RC1: Common ICT infrastructure for Virtual Organisations

We have identified and clarified the need for an open standards-based common infrastructure that enables the secure and reliable exposure and integration of the services and resources offered within a Virtual Organisation. This infrastructure may be independent of the assets of the partners who may wish to form virtual organizations (independent in terms of the business function, of the ownership of its assets and of its operational management).

We have identified the following main research challenges in this area.

(i) separation of concerns between

a. the provision and management of business services by the business (in particular SMEs) that may like to participate in Virtual Organisations

b. the provision and operational management of hosting environments and supporting infrastructure services that enable the rapid deployment of application services by different VOs

(ii) developing business models and system deigns that support businesses that would like to take advantage of a network-centric delivery model to reduce the opportunity-cost and the time-to-market for by:

a. *Maximising Return-on-Investment via outsourcing* the development of a dedicated dependable infrastructure and infrastructure services, the cost of which is often prohibiting for a single business that focuses on a vertical market.

b. *Alleviating the operational management cost* of service deployment and hosting through outsourcing hosting and operational management while maintaining overall control of the terms under which their business function is provided within Virtual Organisations.

c. *Reducing the cost of* building *a secure, reliable and accountable capability exposure infrastructure* by enabling the use of a purpose-built infrastructure capabilities for virtualising one's business functions as managed services;

d. *Reducing the risk of exposure to an Open network* by leveraging on the experience of a dedicated infrastructure provider.

(iii) Optimising the time and effort spent for setting-up and dissolving Virtual Organisations and for implementing change during their operation.

4.2.1 Emerging solutions and trends

Our assessment indicated that although a relatively small, but rapidly growing number of research and commercial tools that claim to provide hosting environments or "glue" software for cross-enterprise integration, their maturation timescales are 3-5 years from now (i.e. 2008-2010) and anyhow none of these is targeting supporting the life-cycle of dynamic Virtual Organisations, or is providing advanced security and SLA management features as yet.

Middleware in the first category includes the Globus toolkit version 4[4] which stem out of technological innovation targeting scientific communities with an emphasis on resource hosting and integration. Such products are now evolving into being a significant part of wider scope enterprise infrastructure systems such as IBM's Grid toolkit[5] and products from Platform Computing[6] and United Devices[7].

Software in the second category includes emerging Service-Oriented Architecture (SOA) based Enterprise Service Bus (ESB)[8] products offered by small companies such as Cape Clear, Infravio, Blue Titan and Sonic Software. In this category the products offered by Cape Clear and Sonic Software are representative. The Cape Clear ESB solution focuses mainly on the creation and hosting of standards based (Web) services. The Sonic Software ESB solution focuses more on offering managed capabilities message brokerage, reliable transactions, asynchronous messaging, etc.

In between these two categories lies a recent initiative by the Apache foundation to offer an open source ESB on top of the Apache Axis2 platform. However this initiative was announced during the Summer of 2005 and it is still in an early incubator phase. Similarly to the above, this initiative aims

---

[4] Globus Toolkit v4 can be seen as a Web services based Grid middleware that facilitates the integration of services and resources that have been deployed on multiple hosting environments. See also http://www.globus.org

[5] See also http://www-1.ibm.com/grid/solutions/grid_toolbox.shtml?Open&ca=daw-prod-gridtoolbox

[6] See also http://www.platform.com/Products/

[7] See also http://www.ud.com/solutions/deploy/mp_enterprise.htm

[8] According to Gartner's definition, an ESB is standards-based middleware that uses a Service-Oriented Architecture (SOA) and that has messaging, intelligent routing, and transformation capabilities. In this document we follow other industry experts who validly extend Gardner's definition to include features like orchestration, security federation, and a common service management framework.

at producing a general-purpose ESB and it does not aim at supporting of dynamic virtual organisations.

4.2.2 Open standards and common design patterns

In terms of common design patterns and open standards specifications, a number of specifications partly address some aspects of this objective. In particular:

- The SOAP and WSDL specifications offer a transport independent means for service-to-service interaction by exchanging meta-data (XML) based messages between applications that can be deployed upon different platforms and have been exposed as Web services.
- The WS-Addressing specification offers interoperable constructs that convey address-related information that is typically provided by transport protocols and messaging systems.
- The SOAP interceptor / Handler pattern offers a programming model for network intermediary network points to process message exchanges between services. These intermediary points may be deployed independently of (a.k.a. “Interceptor”), or co-deployed (a.k.a. “Handler”) with, a Web service endpoint.
- The WS Security stack of specifications is delivering a technical foundation for implementing security functions such as integrity and confidentiality in messages implementing higher-level Web services applications.
- The WS-Notification specification is offering a pattern-based approach to allow Web services to disseminate information to one other
- The WSRF/WSDM (or alternatively the competing WS-Transfer/WS-Enumeration/WS-Management) stack of specifications define a Web services architecture for managing distributed resources, including other Web services endpoints.
- The WS-Coordination / WS Transaction stack of specifications (and alternatively the competing WS-CAF) are defining an open framework for supporting coordinated transactional compositions of multiple Web service applications.

Although the Web Services interoperability organisation (www.ws-i.org) has produced a basic interoperability profile and it is finalising a basic security profile, there is no current initiative to define profiles for realising the basic functionalities targeted by this research challenge.

4.3 RC2: VO management capabilities

Tackling the challenges related to a common ICT infrastructure for Virtual Organisations is a necessary prerequisite but in itself it does not suffice for tackling the general objective of TrustCoM. It has to be enhanced with higher-level “VO Management” capabilities that enable the life-cycle management of Virtual Organisations on top of such an ICT infrastructure.

We have identified the following research challenges relating to VO management:

- The means of modelling and representing the structure of VO have to be investigated
- The development of services for maintaining and propagating the state of a Virtual Organisation has to be investigated. Particular emphasis has to be placed on the provision of mechanisms for programmatically activating and de-activating the availability of services and resources that operate in the context of a Virtual Organisation in order to be able to enforce VO-wide life-cycle changes.
- The development of services for managing VO membership, and in particular services for maintaining information about the engagement and disengagement of VO partners, their role and services and resources they contribute to, or may use within, the Virtual Organisation.

- Models of General Virtual Organisation Agreement (GVOA) need to be implemented. This includes describing the roles and relationships between VO members, general VO-wide policies that need to be enforced across the VO as well as monitoring the performance of VO partners in relation to the GVOA. Such a GVOA can be understood as the framework within which specific service policies, assertions and bilateral SLA are interpreted and integrated.

4.3.1 Emerging solutions and trends

Currently there are no concrete results relating to Virtual Organisation Management as understood in TrustCoM. The term is often found in the literature either in relation to service hosting and service management (e.g. as in Grid Computing projects such as www.gridpp.ac.uk), to access control and/or membership management of virtual domains (e.g. as in CAS www.globus.org/security/CAS/ and VOMS http://infnforge.cnaf.infn.it/projects/voms/) or to abstract VO frameworks (see TrustCoM deliverable **D2: State of the art evaluation** for a comprehensive analysis of VO frameworks).

4.3.2 Open standards and common design patterns

There are no open standards technologies addressing VO Management as such. The following technologies address aspects of components that may be contribute to the development of VO Management services

- UDDI technology may be useful for discovering and maintaining information about the services that VO members could offer (or have committed to offer) within a VO
- WS-Notification allows for topic based notifications. The use of simplified ontologies offers an attractive alternative to fixed, pre-existing, explicit and bilateral publish/subscribe agreements and offers considerable advantages when used for disseminating information about VO life-cycle changes.
- WS-Policy could be used as a means of documenting specific assertions or constraints on interactions relating to (Web) services offered by on ore more VO partners.
- IOEDF can be used as a basis for incident report dissemination among VO members – probably implemented as a profile on top of WS-Notification.
- WS-Agreement can offer a scheme that is partly used in a GVOA. However the layers address by WS-Agreement and current the limitations of the WS-Agreement make it inadequate for tackling the complexity of a GVOA in its full extend.
- ebXML registry could offer an alternative solution implementing a “fully fledged” business registry. However our preliminary investigation found certain problems with using ebXML registry as a potential solution. These are explained in section “3.1: VO Management” of TrustCoM deliverable **D24: Standardisation Roadmap v2**.
- ebXML CPP and CPA offer inspiration for defining profiles and agreement between VO partners, however it is not understood as yet how to generate a CPA based on two or more CPPs. See also section “3.1: VO Management” of **D24: Standardisation Roadmap v2**.

4.4 RC3: SLA and Contract Management

In relation to SLA and contract management, the main research challenges that have been identified are the development of models and mechanisms for the specification of contract templates and the negotiation, monitoring and enforcement of collaboration agreements between existing or prospective VO members. Particular emphasis should be placed on ensuring that such agreements are in harmony with the trust and security management policies across a VO and that and they provide a context for the definition and enactment of collaborative business processes across a VO.

In addition to the GVOA, which has been judged as being more relevant to VO management, we have identified research challenges relating to two main types of agreement that applies to VO partners:

- *Category A*: Research challenges relating to providing support for managing legal contracts between organisations and automate part of the process associated with their definition and enforcement.
- *Category B*: Research challenges relating to – typically bilateral – agreements that capture customer-provider relationships and the Quality of Service promise associated with the provision of a (Web-) service.

Following the problem analysis and technology evaluation presented in deliverables **D3** and **D2**, respectively, the following specific research challenges have been identified:

- To develop explicit conceptual model for supporting agreements at both business and service level needs to be developed based on a conjunction of WSLA, WS-Agreement and relevant concepts from the more generic contacts architecture such as the BCA[9] developed at DSTC.
- The development of this conceptual model needs to devolve significant efforts to two aspects: a) the impact and use of trust and reputation relationships in service discovery, SLA negotiation and enforcement phases and b) the handling of SLA violations in a more flexible form that may include the enactment of business processes to implement compensation.
- In conjunction with the legal team in TrustCoM, to identify which elements of contract management are likely to be the most useful within the framework as well as what security controls in terms of confidentiality, integrity and non-repudiation will be necessary.
- To identify which specific design patterns and implement services for the negotiation of SLA templates, the creation of SLA instances, high-level SLA evaluation and infrastructure-level monitoring mechanisms to support the enactment of (Web-) SLAs.

Given the immaturity of solutions to support contracts that fall in *category A*, including the lack of standardised representations and of mechanisms facilitating operational support for such agreements, and taking into consideration the background, commercial interests and expertise of the members of the consortium, we have decided to start tackling research challenges relating to contracts that fall in *category B* before considering the former.

4.4.1 Emerging solutions and trends

The BCA architecture[6] is one example of a comprehensive ICT model for dealing with legal contracts comprising sophisticated means of describing contracts as well as processes for contract arbitration and enforcement. However, such frameworks (including BCA) are rather complex and its implementation status are usually non-existent or uncertain. Most importantly, they have not been used outside relatively restricted research environments. From a conceptual viewpoint, however, such frameworks propose a number of solutions that are worth investigating in conjunction with a legal team.

Work on Service Level Agreements (SLAs) on the other hand is comparatively more mature and better understood. Originally developed as part of the network and systems management community in order to cater for the specification of the Quality of Service (QoS) parameters characterising the provision of network connectivity services, this work has evolved into general frameworks for the characterisation of application level services and more recently business services. Most of the solutions proposed in this area provide the means for: specifying SLAs and associating them with the WSDL services concerned, discovering and locating services based on profiles of QoS that can be delivered for those services, defining simple negotiation protocols for negotiating QoS parameters, and monitoring the compliance with the SLA objectives (including monitoring and metric definition).

[9] See http://www.dstc.edu.au/Research/Projects/Elemental/BCA.htm for a summary of research activities relating to the Business Contracts Architecture (BCA).

However, the extent to which these features are supported varies greatly amongst the different SLA solutions proposed. Probably the most concrete framework that is likely to provide a solid foundation for TrustCoM is WSLA, which in addition to specification and structuring of SLA agreements also provides detailed monitoring aspects including an extensible framework for metric definition. The other framework of particular interest is WS-Agreement. Originating initially from the OGSI framework, and a good example of how Grid platforms evolve towards a more open web service environment, WS-Agreement caters for the discovery of services including SLA retrieval and negotiation and is compliant with the other WSRF specifications. WS-Agreement is however a relatively new specification, which has not been evolving as rapidly as the community had initially anticipated.

ebXML Trading Party Agreement and Collaboration protocol Agreement also offer an attractive alternative baseline. However, their specifications and implementations are tightly coupled to the other ebXML specifications, which predated recent developments in Service Oriented Architectures and often do not integrate well with the more recent web service specifications.

4.4.2 Open standards and common design patterns

The following open standards technologies are related to the research challenges mentioned in this subsection and may offer part of the baseline used by the TrustCoM consortium.

- *WSLA*: The WSLA specifications allow for the definition of QoS service parameters and the relationship between involved partners and so-called supporting parties that may take over monitoring and related functionalities.
- *WS-Agreement*: As opposed to WSLA, WS-Agreement focuses on interaction protocols and provision of templates. WS-Agreement has little or no support for the definition of QoS parameters. Notably, there seems to be a strong interest by IBM (the developer of WSLA) to integrate WS-Agreement and WSLA.
- *ebXML CPPA*: ebXML CPPA is strongly integrated into the ebXML set of specifications, and may hence not be directly used without significant impact on other technologies used in TrustCoM. However some the concepts used in ebXML CPA appear to be very relevant to the objectives of the project. Such concepts will be adopted following adaptation where appropriate.

4.5 RC4: Trust & Adaptive Security

Trust Management models support the supply and collection of evidence or derived information about the trustworthiness of a prospective VO member to perform a specific task towards an objective of the VO, and the assessment of their reputation by other VO members, who will have to rely on that prospective member (or not) for the specific task.

Research challenges in the area of autonomic security management include the development of models and mechanisms that underpin the life-cycle management of federations of security realms of VO partners as well as security management within and across VO partner realms. Particular emphasis has to be placed on adaptation of security policy and mechanisms to changes in the VO context, self-management, and resiliency to faults or misbehaviour within the realm of a VO partner or the realm of its collaborators.

In order to tackle more effectively this extensive area and to identify common functionalities (such enforcement, adaptation policies and reputation), which may be of a more generic nature than specific to security, we have decided to divide these research challenges into the following areas:

- **RC4.1**: Specific research challenges relating to security token services, which includes basic mechanisms that underpin credentials management and offer a foundation of federation security realms of different VO members. These include the development of specific ***"security token services"*** and mechanisms for issuing, validating and exchanging security claims.
- **RC4.2**. Specific research challenges relating to "**trust negotiation"**, including policies that guide the incremental disclosure of credentials that are needed for satisfying a minimal set of

requirements for a particular transaction within a particular VO, and protocols for securely implementing such exchanges of credentials.

- **RC4.3**: Specific research challenges relating auditing. These include the development of an archetype of an "***audit service***" for VOs and mechanisms that underpin the collection and dissemination of evidence about transactions between VO members.
- **RC4.4** Specific research challenges relating to reputation management. These include the development of the archetype of a **"reputation service"** for VOs, of models for meaningfully quantifying reputation and computing reputation values as well as mechanisms for collecting and collating evidence or other information (e.g. recommendations) on the basis of which the performance of a VO partner may be assessed.
- **RC4.5** Specific research challenges for "**access control policies**", including policies concerning the **delegation** of administrative authority. This includes the development of authorisation and delegation policy templates as well as the development of **Policy Decision Points** that have the intelligence to produce decisions at run time based on such policies.
- **RC4.6** Specific research challenges relating to **adaptation**. These include the development of adaptation models, and notations for specifying policies that describing conditions under which the system may automatically adapt its behaviour and of services that implement adaptation actions, i.e. actions that result in adapting system behaviour in reaction to contextual changes.

It has to be noted that following the research conducted during the first half of the TrustCoM project in this area, it became apparent to us that a sufficiently generic form of adaptation policies underpins goals and/or solutions relating to other challenges such as VO management, SLA management and BP enactment. Consequently a set of goals relating to policies has been separated from the set of goals that are specific to trust, secure federation and reputation. The goals relating to the Policy address our specific research challenges relating to adaptation policies, on the one hand, and permission, prohibition, obligation and delegation policies on the other.

4.5.1 Emerging solutions relating to adaptive security and federation

Security aspects of a VO framework span a large number of concerns that broadly divide in the following categories: Secure Federation, Authorisation, and Adaptive Security. These will each be addressed in turn in the following paragraph. Overall security and policy are not only a substantial part of TrustCoM but one where the consortium has considerable expertise.

Access Control Models are well understood within a single administrative domain and new concepts such as Role Based Access Control are increasingly appearing in main stream products. Authorisation policies are used in a number of different frameworks (Ponder, Permis, SPKI, etc) and standards (XACML). Despite apparent differences between the specification languages their functionality is broadly similar. Their enforcement is sometimes different, in particular when applied in distributed environments but the advantages and disadvantages of the various solutions are again well understood. However, distributed access control within environments that cross domain boundaries remains fundamentally an open research problem. Grid environments have attempted to address these issues in a number of platforms (Akenti, VOMS, CAS, etc.) however the assumptions on which these models are based are too restrictive for VO enforcement. In particular, most grid-platforms are concerned with access control to resources by distributed tasks and do not allow for recursively composable VOs in federated structures (i.e., a Grid is not itself a VO that can participate in higher-level VOs). One common characteristic across all platforms is however the increased usage of arbitrary security tokens to convey relevant security information. As domain boundaries are crossed, local identity loses any meaning and access control decisions are made based on properties that the requestor proves he possesses. These properties may include its role, qualifications and other attributes as well as privileges he/she holds or that have been delegated to him/her. This evolution is also evidenced in the more recent web-service standards such as WS-Trust, SAML and WS-Federation. The latter, in particular, focuses on the exchange and use of such tokens across domain boundaries. Authentication, and in particular authentication based on identity, becomes then a particular case of the more general token based framework described above. Recent studies and standards have particularly focussed on Single Sign-On systems such as Liberty

Alliance and Shibboleth. Both of these overlap in scope with WS-Security, WS-Trust, WS-Federation based standards but tend to be less flexible (e.g., lack of support for “active” requestors), focus on identity management alone and rely on SAML for communication of information and SSL as the underlying secure transport protocol.

4.5.2 Emerging solutions and trends on trust management

Trust management remains a significant area of research despite numerous attempts to address this issue. The fundamental paradox of trust management as a research area is that although there is wide spread agreement on the importance of using trust in a variety of contexts including business transactions and although each one of us has an intuitive belief for what/who we trust, there is little agreement on what trust *is* or how to characterise it. Indeed, the various trust management frameworks proposed in the literature differ significantly both in their definition as well as is their computation of trust. The following aspects are by and large agreed in the various studies on trust:

- Trust is intimately linked (or derived from) different elements such as: recommendation, reputation, risk, and evidence of behaviour.
- Trust is linked to a well identified contexts including the activities being performed, the parties engaged in the interaction as well as other contextual elements of the transactions. However, none of the solutions in existence address this adequately.
- Trust may be expressed in relation to different characteristics of the parties involved in a transaction or the activities being performed such as competence, and honesty of the parties, correctness of the execution of the transaction or its result.
- Trust should be quantifiable as otherwise little use could be made of it. However, n consensus has been reached on the desired metrics for its quantification.

The various studies can be broadly divided into two categories, those that focus on trust aspects of a security infrastructure in particular with regards to the authentication of users or disclosure of information and general frameworks for trust management that focus on trust analysis, quantification and trust services. The former are relatively well understood in particular when relating to PKI infrastructures. In addition, there are also a number of emerging studies on trust negotiation i.e., the incremental disclosure of security relevant information such as credentials and requirement for access although further studies are needed in this area. The latter have also been subject of a number of studies but there is little consensus on how to define, manage and compute trust based on an infrastructure of trust services.

4.5.3 Open standards and common design patterns

- *X.509 (PKI), X.509 PKI Profile, WSS X.509Token*: default security token format, particularly for intra-organization use.
- *X.509 PMI, X.509 AC Profile*: used for authorization tokens and enabling delegation of authority.
- *WS-Trust*: web service interface adopted for issuance and validation of security tokens, i.e., interaction between enforcement point and security token service; a specific profile is implemented.
- *WS-Federation, WS-Federation Active Requestor Profile*: the federation model is adopted; while specific features (such as the pseudonym service) are not supported in v1 of the TrustCoM framework, we may adopt more of these in the future.
- *WSS SAMLToken, SAML Token Profile*: a custom token format is implemented for cross-organisation use; the SAML token format would be a good candidate to migrate to for the next version of the framework.

- *WS-Federation Passive Requestor Profile*: the eLearning scenario implements a custom username/password authentication scheme, but may likely adopt the WS-Federation passive requestor profile in the future.
- *SAML*: the SAML token format is a good candidate for cross-organization use (see above); the SAML protocols are currently not adopted, as the WS-Trust protocols have been selected for token (including authorization tokens) interaction between enforcement points and security token services.
- XACML is currently used as the main intermediate-level policy language for defining *attribute-based* Access Control policies that are loaded in an XACML compliant Policy Decision Point (PDP). Also XACML request / response operations in a SOAP envelope are used as a baseline for implementing message exchanges relating to access control policy decisions made by an XACML compliant PDP.
- *WS-PolicyAttachment, WS-MetadataExchange*: in v1 of the framework, the creation of access control policy instances and the configuration of the appropriate PDPs is performed out of band. All related metadata is exchanged out of band; in-band exchange is important to be considered in the next version of the framework.
- XACML profile of SAML can provide an alternative protocol for interacting with XACML-compliant PDPs. SAML protocol is not considered at present as a baseline protocol for authorisation and access control request/response message exchanges.
- *XML Key Management (XKMS)*: as a VO-wide PKI was not a direct objective, this specification is not considered during the conceptual investigation.
- *Liberty, Shibboleth, Web Single Sign-On Interoperability Profile, Web Single Sign-On Metadata Exchange Protocol*: in its second phase the project expects to revisit interoperability with single sign-on systems.
- *Use of SAML for OGSA Authorisation Profile*: this is a relevant ongoing standards initiative, but a specification is not yet available.
- *WSS UsernameToken, WSS KerberosToken*: X.509 certificates are use for intra-organisation communications, but the platforms underpinning the TrustCoM framework allow transparent use of these alternative token formats.

4.6 RC5: Collaborative Business Processes

Business Process Modelling and implementation should proceed based on mature open-standard specifications and any available packages providing adequate implementations. The following are specific challenges that are being tackled by the TrustCoM Consortium in this area.

- **RC5.1**: *Specific challenges about supporting the life-cycle of Business Process instances*. A major challenge is, however, to support the whole life-cycle of instances of VO-wide collaborative processes. Besides consistently modelling such processes, the extraction and distribution to process views to VO partners and the joint enactment of such processes are particularly challenging.
- **RC5.2**: *Specific challenges about correlating Business Processes and Service Level Agreements.* As it is elaborated in TrustCoM deliverable **D2: State-of-the-art evaluation,** few (if any) of the existing studies properly tackle business processes in conjunction with SLA and none in conjunction with trust and reputation information for service selection and composition.
- **RC5.3**: *Specific challenges about supporting adaptation & administrative processes.* Another challenge for the consortium is to find how business processing technology or methods can be used for defining light-weight mechanisms for implementing transactions that support the administrative and possibly adaptation processes that bring added value to the ICT infrastructure upon which VOs evolve.

After an initial phase of defining and implementing the core business process functionality, the efforts should focus on three aspects: integration with the SLA infrastructure, leveraging the availability of trust and reputation for providing enhanced flexibility in the enactment of the processes especially across administrative domains, and offering models and technology for automating common transactions between the infrastructure and supporting services that TrustCoM develops.

4.6.1 Emerging solutions and trends

By comparison, business Processes are probably the best understood and defined technology. Indeed, the issues regarding executable collaborative business processes in the last few years have been more focussed towards standardisation aspects rather than basic research, as many software vendors and business integration consultants are using a wide spectrum of proprietary protocols.

Standardisation allows addressing the problems of executable business process aggregation and collaboration across administrative domains that use proprietary solutions as well as outsource workflow control and implementation to third parties.

A number of specifications have been investigated including:

- WS-Coordination that defines the means to coordinate distributed actions during process runtime including agreement on outcome through the propagation of activity contexts
- WS-Transactions that extends context information to include transactional capabilities for both atomic transactions (WS-AtomicTransactions) and long running business transactions (WS-BusinessActivity),
- WS-CDL that focuses on the choreography of message exchanges starting at design time across multiple parties and
- BPEL4WS that provides the means to describe abstract and executable business processes in terms of their structure, control as well as offered and invoked service interfaces.

BPEL4WS and BPML/WS-CI have overlapping functionality, in particular for the business process specification although from different points of view. Whilst BPEL4WS relies on supporting Web Service standards such as the WS-Coordination model, which relies on the use of a single coordinator entity or a hierarchy of coordinators to control the execution of the workflow, WS-CI advocates a more loosely coupled choreography model with distributed control. Since many of the use-case scenarios established for TrustCoM do not explicitly require the use of a coordinator the latter mode may provide some flexibility. Regrettably, development of the BPML/WS-CI has been abandoned with most of the concepts being integrated in a new specification, WS-CDL. The latter however, is still evolving and is not sufficiently stable to base the TrustCoM development upon it, at least during the first stage of the project. At present WS-CDL is also not adequately catering for a collaborative business process choreography description capturing complex message exchanges across administrative domains, for instance in tendering and quotation processes.

Finally, there are few solutions, if any, that attempt to tackle the problem of relating business processes with the SLA of the services they engage. Furthermore, the co-use of choreography approaches (e.g. WS-CDL based approaches), which naturally fit for describing high-level VO-wide processes and WS-BPL (BPEL), which naturally fit for implementing more dynamic processes within the realm of a specific VO partner, has not been investigated adequately although it has been often discussed.

4.6.2 Open standards and common design patterns

- *WS-CDL:* This process choreography language is used to define a collaboration definition for a VO. Based on this collaboration definition, the public business processes and WSDL interfaces of the VO members are derived. However, WS-CDL is still evolving. Nevertheless, it seems to be most promising for specifying the collaboration definition (business protocol) of a VO. It is not complete to cover all complex business interactions (e.g. multicasting). However, the current version of it can be used to base the TrustCoM development upon it. Future versions of the specification will be monitored for further developments.

- *WSCI* is also relevant but we have noticed that WS-CDL to a large extend covers those aspects WSCI functionality that appear to be more relevant to the TrustCoM goals.
- *WSBPEL* provides a reasonably mature language set for executable business processes. It focuses on the control and orchestration aspects of business processes and leaves business logic to invoked web service implementation. WSBPEL can be used to specify “public” processes (views) of VO members.
- *BPML* is also relevant but it appears that WSBPEL covers the necessary BPML functionality for the needs identified in TrustCoM.
- *WS-Coordination* and *WS-AtomicTransaction* are used as a means of implementing distributed transactions and coordination protocols at the level of VO Infrastructure. Mechanisms developed in this subsystem for Business Process enactment will leverage on the VO Infrastructure capabilities wherever such protocols are required.

5 TrustCoM Framework: an overview

The TrustCoM Framework is divided into six loosely-coupled subsystems, each of which focuses on a complementary aspect of an ICT infrastructure for dynamic Virtual Organisations. In this section we provide an overview of the TrustCoM framework. Refer to deliverables **D16**: *Conceptual Framework* and **D9**: *Reference Architecture* for a detailed description of an abstract architecture proposal, deliverable **D19**: *Reference Implementation* for a description of the detailed designs of the components and services that are currently being developed and to deliverables **D24**: *Standards Roadmap* and **D18**: *Framework Specifications* for a detailed description of the open standards technologies upon which the TrustCoM Framework is based and for an overview of how these standards are extended and integrated into interoperability profiles for each subsystem.

5.1 Enterprise Network versus Virtual Organisation

A starting point for the TrustCoM framework is the requirement for an advanced form of an open distributed and standards based Enterprise Service Bus (ESB) which we have named “Network of Enterprises” or “Enterprise Network” (EN) in order to avoid overloading the ESB term.

In addition to the common ESB characteristics, i.e. being based on Service-Oriented Architecture and having messaging, intelligent routing, and transformation capabilities, we require that EN provides capabilities for brokerage, notification, distributed transactions, security federation, policy enforcement, and a common service management framework, including the ability to programmatically deploy new application capabilities and to create, and manage the life-time of dedicated endpoint instances for virtualising these capabilities in the context of different VOs.

The EN concept also extends the ESB model by incorporating VO agreement, service-level agreement and policy templates that can be instantiated upon request in order to facilitate the rapid formation of VOs. In analogy to the EBS paradigm, there is a clear separation between the EN, where capabilities are exposed and advertised, and the application hosts that simply accommodate application-specific or supporting components that implement the capabilities. Access to the capabilities takes place only in the context of some VO and only via dedicated, managed endpoints. (At a conceptual level the latter are analogous to service instances of a capability that are offered exclusively to a VO and are subject to the agreements and policies of that VO).

The EN can be understood as the infrastructure underpinning a VO ecosystem. Although the EN/VO Infrastructure subsystem of the TrustCoM framework aims to offer key functionalities of the EN concept described above, all other subsystems of the TrustCoM framework focus mainly on what happens within such a VO ecosystem.

5.2 TrustCoM Framework subsystems

5.2.1 Virtual Organisation Management

The VO Management subsystem aims to offer the essential capabilities for managing the state and life-cycle of a Virtual Organisation. In particular, defines and maintains details of each Virtual Organisation which is operating within the Enterprise Network and offers three main modules that are respectively responsible for the lifecycle changes to the VO, the VO membership management, and the General VO Agreement management.

5.2.2 Business Process Enactment and Orchestration

This subsystem aims to offer the essential capabilities for modelling, deployment and execution of collaborative business processes across a Virtual Organisation. In particular it offers services for

producing choreographies from high-level business process models, distributing views of such processes to different VO partners, and for the secure enactment of these processes by services offered by the corresponding VO partners.

5.2.3 SLA Management

This subsystem aims to offer the essential capabilities for managing the life-cycle of (Web services) SLA instances among different VO partners about the services they provide and for monitoring the fulfilment of these agreements. Its ultimate goal is to support the full “lifecycle” of a service level agreement between the service provider and a customer, respectively the virtual organization – this covers provision of SLA-related information about a service, negotiation of SLA terms, configuration of the involved components, enactment of the SLA (monitoring and evaluation), feedback, and finally “unbinding” the service provider at the end of the SLA instance life-time.

5.2.4 Trust & Security Services

This subsystem aims to offer essential capabilities for security credentials management, auditing and reputation in dynamic Virtual Organisations. In particular, this subsystem contains services for issuing, processing, negotiating and validating credentials assigned to services and resources of different VO partners; services that enable auditing message exchanges within a VO; and services for evaluating “reputation” of a VO partner based on evidence about the performance of the services and resources that are contributed to the VO by this provider;

5.2.5 Policy

This subsystems aims to provide the capabilities for defining, managing the life-cycle of, and making decisions at run-time on the basis of, policies that control access to services and resources of VO partners, policies for delegating (under constraints) the authority to administer specific types of access control policy, as well as of policies for dynamically reacting to changes of the VO context.

5.2.6 VO Infrastructure

This subsystem aims to offer the infrastructure upon which the capabilities offered by other TrustCoM subsystems may be deployed. In particular to allow for

- Remotely deploying new business services or TrustCoM capabilities as Web services within an Enterprise Network or an already formed Virtual organisation.
- Creating on demand new VO-specific instances of business services or of already deployed TrustCoM capabilities, and managing of the life-cycle of such service instances through dedicated management services.
- Enforcing specific security, SLA management and transaction actions on VO-specific service instances.
- Dynamically re-configuring at run-time the binding of VO-specific service instances to the trust, security and SLA monitoring components that support their operation without any need for redeployment.
- Dynamically adapting at run-time the enforcement actions applied on VO-specific service instances without a need to redeploy the service.
- Allowing the implementation of secure and reliable message exchange protocols between VO-specific service instances.
- Allowing the implementation of explicitly defined protocols that implement common transactions requiring the dynamic integration of several components from one or more TrustCoM subsystems.

6 Projected timescales

In this section we summarise projections indicating the timescales within which we expect the areas where TrustCoM making research advancements to have an impact. We do this by means of three diagrams: the projected impact of the technologies relating to the TrustCoM subsystems, the projected timescales of the standards adoption relating to the TrustCoM Framework and the projected timescales by which the research advancements tackled in each TrustCoM subsystem are likely to have an impact. Instead of absolute timescales our diagrams have been normalised in relation to the following distribution.[10]

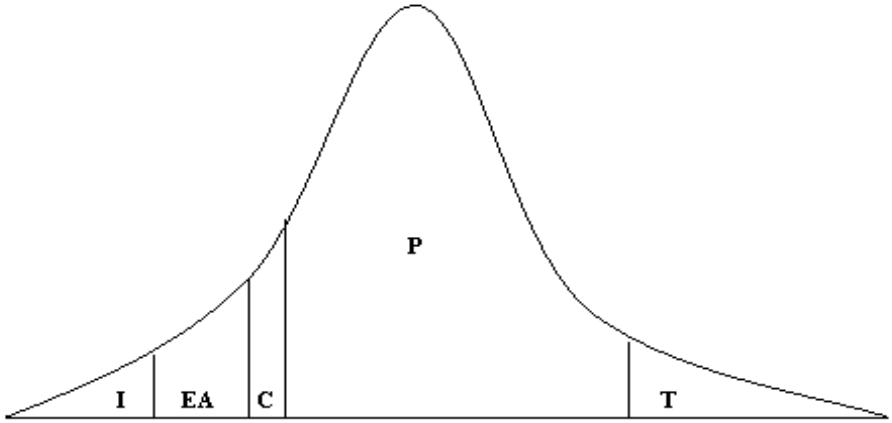

Innovators (I)	The enthusiasts who like technology for its own sake.
Early Adopters (EA)	Those who have the vision to adopt an emerging technology to an opportunity that is important to them.
The Chasm (C)	Time gap in technology adoption, which is between the early adoptors and the pragmatists.
Pragmatists (P) Early Majority	Early majority pragmatists are the solid citizens who do not like to take the risks of pioneering, but are ready to see the advantages of tested technologies. They are the begining of a mass market.
Pragmatists (P) Late Majority	Late majority pragmatists, who represent about one-third of available customers, disklike discontinuous innovations and believe in tradition rather than progress. They buy high-technology products reluctantly and do not expect to like them.
Traditionalists (T)	Traditionalists (laggards) do not engage with high technology products - except to block them. They perform the valuable service of pointing out regularly the discrepancies between the day-to-day reality of the product and the claims made for it.

Figure 3: Overview of Moore's high technology product adoption pattern and adopters' classification

[10] The distribution is based on the elaborate analysis on trends underpinning the introduction of new technology by Geoffrey Moore in *"Crossing the Chasm: marketing and selling high-tech products to mainstream customers"*.

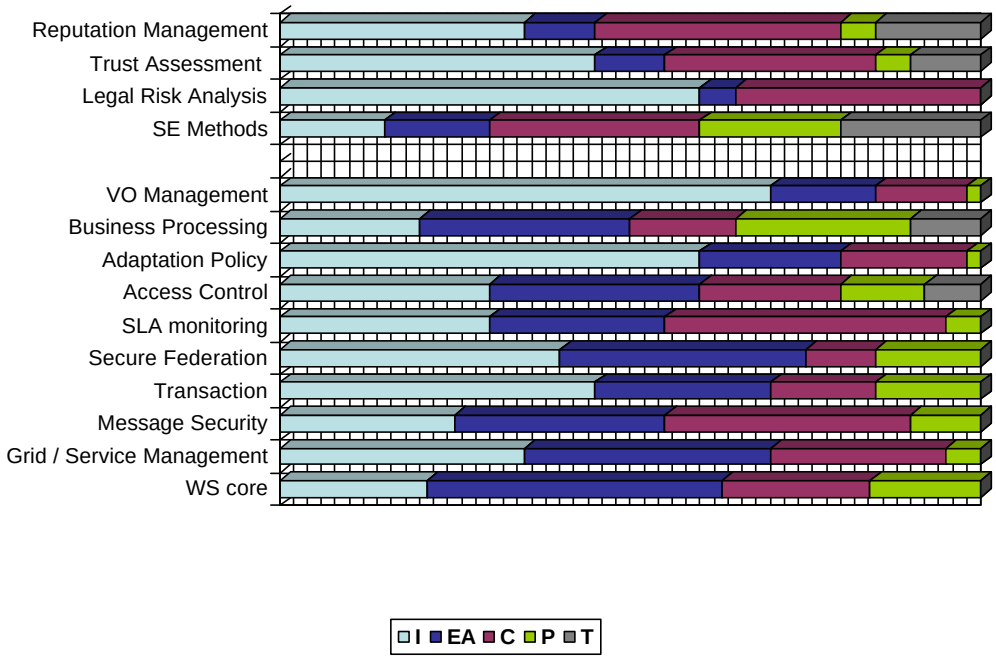

Figure 4: Applied research results uptake normalised over Moore's distribution.

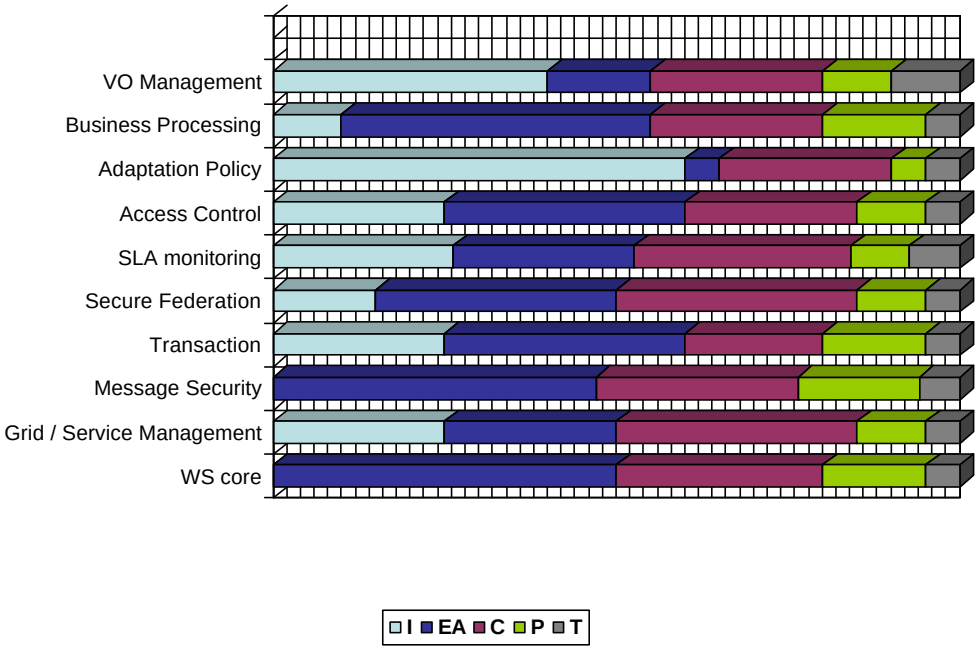

Figure 5: Standards adoption normalised over Moore's distribution.

7 Progress assessment

In this section we summarise the main achievements of the project so far and offer a self-assessment of the progress in the project compared to state-of-affairs outside of TrustCoM. We explain the main achievements in each area and set the context within which the recommendations made in section 8 about the second phase of the project will have to be interpreted.

7.1 VO Management

The VO Management subsystem defines and stores details of each virtual organisation which is operating within the Enterprise Network. It is divided into three main modules responsible for the lifecycle changes to the VO, the VO membership management, and the General VO Agreement management. During the first phase of TrustCoM the conceptualisation and architectural design of the VO Management subsystem was defined and the first version of the VO membership management component, including its fundamental functionalities, was implemented.

The main advancement to the existing state-of-the-art consists in the additional functionalities for VO-level contract management (in the form of General VO Agreements) and VO lifecycle support, taking into consideration trust and security aspects. An innovation is the inclusion of collaboration agreements, the so-called General VO Agreement, which allows the VO membership component to manage membership to the VO in a dynamic way, taking into account trust and security constraints (policies). Another innovation is the existence of a life-cycle management component, which detects events that may change the current state of the VO, such as violation of agreements by a partner.

7.2 Business Processing

The BP Enactment and Orchestration subsystem allows the modelling, deployment and execution of secure collaborative business processes. Business processes consist of a public process which exposes an interface necessary for collaborative interactions of a private executable process. It builds on existing technology standards for the definition of business processes (BPEL) and choreographies (WSCDL) and existing components for the execution of the (private) business process (BPEL engine). During the first phase of the project the conceptual model, architecture and design documentation identified the following services to be implemented: UML2CDL service, CDL2BPEL, and BPM service. Furthermore, the necessary extensions to the standards above for the integration of trust, security and contract management into business processes have been identified.

The design of the BP Enactment and Orchestration subsystem allows for the high-level design of the business process and then automatically derivates the necessary deployment parts in a top-down fashion. The CDL2BPEL does the automatic translation of the choreography into executable business processes using a knowledge base supporting it with process design patterns. This is a new form of design for collaborative business processes bridging the information gap between the coarse grained CDL descriptions capturing the global view of the collaboration and the detailed subjective process model for each collaboration partner. The CDL2BPEL service is capable of generating both, the private and public business processes. Additionally, the concept of TSC extensions has been introduced. TSC extensions (roles and context) are design elements for the business process that allow the inclusion of security, trust and contract management activities and decisions into the business process. They are modelled at the highest level of the choreography, preserved during above translation and deliver private and public process control during their execution, so that they are easy to specify for the business process designer.

7.3 SLA Management

During the first phase of the project, we developed the main services that allow for monitoring and evaluating the quality of service of specific participants in the virtual organization. The setup allows providers to supervise the service performance given that they support either WMI or Ganglia for monitoring the system[11]. The current architecture assumes that evaluation of the performance is carried out at VO-level by a trusted third party, whilst monitoring covers potentially sensitive data and hence information is provided by interceptors co-hosted with the service or resource being monitored. The SLA definition is easily comprehensible and extensible, thus allowing business entities to adapt it to their needs.

The TrustCoM project has advanced SLA Management further to practical usage in the business domain. The SLA Management system takes up automatic monitoring of service performance with respect to a predefined set of quality of service parameters – this allows service providers to simply define the relevant parameters and supervise their services accordingly. The SLA monitoring infrastructure that is currently being prototyped principally allows service providers to apply their own methods for calculating the required parameters, thus potentially enabling the provider to "neutralize" sensitive data. For example instead of reporting the actual clockspeed used for the service, work on the percentage of the speed in relationship to the maximum cpu speed[12].

7.4 Trust & Security Services

7.4.1 Trust & Security: Security Token Services

The main achievement in the area of 'Security Token Services' (STS) is the implementation of STSs that support cross-partner security token validation for virtual organizations. We developed an architecture that helps business owners inside a partner organization to manage VO-membership and assignment for their assets in a secure way. These 'assets' include employees of their organization that work in the VO, as well as services and resources that the organization contributes to the VO.

The novelty in the area of security for cross-partner service invocation is the way of how we ease security management. In our model, we support business people with knowledge about their assets (employees, services, resources) to make the appropriate security decisions. For example, the project owner already knows which people work on the project, which services are assigned to the VO, etc. In many existing security management systems, the business owners have to communicate their business requirements to IT administrators, which implement these requirements by configuring the IT system. In our model, business people manage 'their' part of the system themselves. The system makes the results of the management operations directly visible to the IT administrators, thus there is no loss in control. Our management model does not require project owners to become security experts in order to configure their projects and systems appropriately.

7.4.2 Trust & Security: Reputation Service

The achievement regarding 'trust and reputation' is the implementation of a generic trust and reputation architecture. The term 'generic' means that the system can be modified to support arbitrary actors and relationships, through the configuration of a reputation schema. Different trust metrics can also be built into the system through modifying the algorithms. Therefore, customers can tailor the system to their specific needs and fulfil a wide range of different requirements for trust and reputation management in virtual organizations. For example, customers can implement VO

[11] Note that WMI comes with the windows operating system and that Ganglia (http://ganglia.info/) is a free tool for both the Windows and the Linux platform

[12] This requires a distinction between negotiation and monitoring/enactment – see issues & targets

formation phases using VO partner identification with recommendation-based mechanisms, as well as using evidence collected via the VO infrastructure for rating VO partners during the VO execution phase.

Research conducted so far in this area, advances the state-of-the-art by defining a generic model and schema for reputation systems and showing how this can be applied to existing reputation systems such as e-bay and Slashdot, and by exposing the core functionality via a web services message exchanges so that it can be seamlessly used by multiple different VO entities.

7.4.3 Trust & Security: Secure Audit Web Service

The Secure Audit Web Service (SAWS) allows client applications to securely store important log messages in secure audit trails on which any tampering can be detected. During the first phase of the project, the core functions of SAWS have been implemented, and a Java API interface has been developed for client applications to invoke SAWS. The SAWS component comprises both a SAWS writer and an 'audit trail viewer'. Processes inside the VO, such as other web services or clients, can send log messages in any digital formats to the SAWS writer. In the next phase, apart from the Java API interface, SAWS will further expose its functionality via a web service front end.

SAWS design includes the following distinct features that constitute unique improvements to competing solutions. The design of the secure audit web service allows client applications to store log messages in any digital formats on untrusted machines. Only append mode of access is allowed for the audit trail, so that users or applications cannot rewind the audit file and delete or modify information that has already been stored there. Also only the SAWS writer is authorised to be able to append log records to the audit trail. Though unauthorised applications or attackers may gain access to the audit trail and try to append fake log records to the audit trail, or modify or remove the audit trail, this can be detected by the tamper detection mechanism adopted by SAWS. Since an audit trail may be stored on untrusted machines, the SAWS security mechanism also ensures persistent and resilient storage of the audit trail and detection of tampering of the audit trail – modification, deletion, insertion, truncation, or replacement. Through the SAWS web service interface, SAWS can support multiple simultaneous clients, and it can record any digital content coming from any SAWS client. Since the audit trail may contain sensitive information, the secure audit mechanism can optionally ensure that only authorised applications or people have the privilege to read the audit trail.

7.4.4 Trust & Security: Trust Negotiation Engine

The 'trust negotiation engine' implements a mechanism for the selective disclosure of credentials. The goal is to provide mechanisms for policy exchange between parties, policy match evaluation and credential exchange between parties. These mechanisms, such as the credential negotiation process, are controlled with disclosure policies that regulate which information is disclosed under which circumstances.

In the first phase of the project, we defined a general model of the Negotiation and clarified the diverse policies needed to drive the negotiation process, as well as their specific purposes. Moreover, the internal architecture and the interactions of the Trust negotiation system with other TrustCoM services were defined and a first proof-of-concept prototype was implemented, not yet integrated with other TrustCoM services.

The introduction of a Trust Negotiation Engine to an ICT infrastructure for Virtual Organisations constitutes an improvement to the existing solutions (both research prototypes and commercial products). The availability of a Trust Negotiation capability in a VO setting enhances the VO's flexibility during the formation and evolution phase. Actually, a VO manager can use the Trust Negotiation in order to accept a new VO partner when no reputation information is available about it. In our model, we express negotiation policies as conditions on assertions (credentials); thus, the model allows to associate credential attributes with a property that tells to what extent the *value* of the attribute can be disclosed during the negotiation. This model feature provides us with fine-grained disclose protection for the subject's sensitive information to the counter-party during a negotiation, depending on the sensitivity of the information.

7.5 Policy

During the first phase of the project we have designed and started the implementation of a generic policy model that caters for obligation policies in the form of event-condition-action-rules and access control policies in the form of both authorisation and delegation rules. This model permits policies to be dynamically enabled or disabled, thus adapting the configuration of the VO in response to events such as security violations, SLA violations, failures, and changes in trust or reputation, without interrupting the functioning of the VO. This enables the VO to adapt to changes throughout its life-time including searching for new partners and removing new partners when needed, changing the security parameters used in interactions, enabling or disabling audit and gathering of evidence according to the reputation and performance of the VO participant in the business process.

Authorisation policies permit the specification and enforcement of access control rules that include constraints based on the attributes of the requestor and additional context parameters such as time. Delegation policies enable decentralized and distributed management of access control[13] [14] by making it possible to specify who may administer access control policies, based on the attributes of users, resources and administrators. Delegation policies are also useful in expressing the sharing of resources in a VO since such sharing entails delegation of access control. The policy subsystem further provides the means to group policies in relationships that capture the authorisations and adaptation requirements of particular interactions within the VO that relate to a specific purpose.

The implementation of relationships provides the means to distribute policy enforcement across several policy interpreters in a structured way thus enabling the framework to scale to large VO structures whilst preserving encapsulation and separation of concerns. Policies and relationships are themselves objects that can be managed using other policies. Thus, it is possible to trigger the deployment of a new set of policies without human intervention. For example, when reputation of a participant in the VO drops significantly a new set of access control policies may be deployed in order to withdraw that participant's access to a set of application services. This will also permit to cater for the ad-hoc services scenario where new groups of policies need to be deployed upon service instantiation.

The majority of current VO Infrastructures and Enterprise Integration products lack the ability to adapt the VO's functioning based policies (i.e., declarative rules that can be dynamically replaced) although the principles and suitability of the approach for enterprise systems has been presented before[15] [16]. Thus, the application and use of policy-based management within a VO framework is entirely novel. The policy sub-system in the TrustCoM framework is based upon the lessons learnt from the development and implementation of the Ponder[17] [18] framework, which has been developed over a number of years at Imperial College London. Developed within the context of network and systems management, Ponder has been applied to numerous application areas including network QoS management, mobile systems, mobile agents platforms, enterprise systems and role-based and distributed access control. With respect to Ponder, the policy model proposed in the TrustCoM project presents a number of innovations including:

---

[13] Olav Bandmann, Mads Dam, and B. Sadighi Firozabadi. Constrained Delegations. In *proceedings of 2002 IEEE Symposium on Security and Privacy*, 2002.

[14] B. Sadighi Firozabadi, M. Sergot, and O. Bandmann. Using Authority Certificates to Create Management Structures. In *proceedings of Security Protocols, 9th International Workshop, Cambridge, UK*, April 2001

[15] E. Lupu, M. Sloman, N. Dulay and N. Damianou. Ponder: Realising Enterprise Viewpoint Concepts. Fourth International Enterprise Distributed Object Computing Conference (EDOC 2000), Makuhari, Japan, Sept. 2000.

[16] E. Lupu, Z. Milosevic and M. Sloman Use of Roles and Policies for Specifying, and Managing a Virtual Enterprise. Ninth IEEE International Workshop on Research Issues on Data Engineering: Information Technology for Virtual Enterprises (RIDE-VE'99). March 23-24, 1999, Sydney, Australia.

[17] N. Damianou, N. Dulay, E. Lupu and M. Sloman. The Ponder Policy Specification Language. Policy Workshop 2001, Jan. 2001, Bristol, U.K., Springer-Verlag, LNCS 1995.

[18] N. Dulay, E. Lupu, M. Sloman and N. Damianou. A Policy Deployment Model for the Ponder Language. IFIP/IEEE Symposium on Integrated Network Management, Seattle, USA, 2001, IEEE Press.

- A new policy specification language that emphasises relationships between services.
- A new aggregation model for policies based on nested relationships
- A new enforcement model for obligation policies
- The use of XACML for authorisation policies.
- Extensions to the XACML model in order to specify and enforce delegation policies.

Finally, the extensions of delegation to XACML are novel in that XACML, which is perhaps the most known commercial standard for access control policies, has not had any support for policy administration. We are currently participating in the XACML committee to extend the XACML specification with delegation, which has been well received, and is being worked into the next version of the official standard.

7.6 VO Infrastructure

The focus during the first phase of the project has been on defining the concept and building the core functionalities of an ICT infrastructure upon which VO may evolve. This involved:

- identifying and clarifying the need for the existence of a VO infrastructure that may be independent of the assets of the partners who may wish to form virtual organizations (independent in terms of the business function, of the ownership of its assets and of its operational management)
- eliciting requirements from the scenarios and abstracting away their application domain specific aspects
- providing the overall design of a VO infrastructure
- identifying the core functionalities that are needed by most other VO business or supporting services and do not fall into the realm of some other domain specific subsystem (i.e. business processing, SLA & contract management, trust & security services, policy)
- producing a roadmap towards integrating these functionalities into a whole
- producing detailed designs and early prototypes of some of these functionalities based on widely available (and preferably open source) software

In the following paragraphs we summarise the main areas addressed during the first phase of the project emphasising the key research and technological advancements achieved in each area.

7.6.1 Enforcement & Service Management

Work on enforcement and service management focused mainly on creating service endpoints that used for exposing a (Web) service in the context of a specific VO, which have a potentially limited life-time that is tied to the period during which the service is offered in this VO, and which are “manageable” in the sense that their life-time and configuration can be set and changed programmatically by dedicated clients (e.g. administrator’s GUI) or by management services.

7.6.1.1 Capability deployment

In the first phase of the project we have developed a service that exposes programmable interfaces allowing an administrator to “push” application code into a remote host and deploy it as a web service. Capability deployment offers a means of enriching the capabilities of the VO infrastructure or an application domain by introducing application code to implement new functionalities that are exposed as web services. Particular emphasis has been placed on securely deploying new capabilities in dedicated “sandboxes” therefore isolating the associated execution processes. Prototype implementation of “capability deployment” started from a .NET environment. The functionalities implemented go beyond what is currently available over the .NET platform.

7.6.1.2 Service Management

Work on service management has focused on exposing programmable interfaces that enable:

- Creating for each capability one or more dedicated service endpoints that are dedicated to exposing a specific Virtual Organisation and implement different trust, security and contract management configurations. This offers an inexpensive way of securely exposing a service to different VOs, on the one hand, and making the same service exhibit different behaviours in different VOs, on the other hand.
- Managing the (security, trust, contract management) behaviour of these endpoints through programmable interfaces. Consequently enabling their management via dedicated client interfaces (e.g. Administrator's GUI) or via dedicated management services that implement intelligent controls such as the "adaptation" policies of section 7.5, which allow automating certain adaptation scenarios.

In the first phase of the project we identified, designed and implemented:

1 A novel technique for creating, exposing and operating manageable endpoints that can provide the means via which a service interacts within a virtual organisation.

2 A novel technique for maintaining the (trust, security, monitoring, transaction) configuration meta-data for each of these service endpoints, so that the same "capability" may expose different service behaviours (in terms of trust, security, contract monitoring and transactions) through different endpoints.

3 The archetype of a "Factory" service, which creates manageable endpoints of a service by implementing the techniques described above.

4 A novel technique for updating the configuration of a specific endpoint via an authorised client (e.g. Administrator's GUI) or via a dedicated service.

5 The archetype of a "Configuration Management" service, which loads specific configurations and is able to accordingly reconfigure the endpoints that it is authorised to manage. As yet this does not include advanced intelligence such as the ability to implement general purpose adaptation policies such as those mentioned in section 7.5.

The Factory service is an advanced implementation of a commonly used "creational" design pattern[19]. It is usually associated with an already deployed capability that has been exposed as Web service. The factory is exposed as a separate Web service and can create services instances of a capability at a service host. The main advantage of this approach is that it separates the deployment of a Web service from its exposure in a specific VO and enables the creation of multiple Endpoints, each of which comes with an explicitly described, security, contract and transaction configuration. Although the concept of a factory service is well understood in CBSE and Grid Computing, the use of a Factory service for creating dedicated, manageable and reconfigurable service Endpoints is novel and offers a new perspective on what can be achieved by leveraging on the converging points of Grid and Web services technologies.

The Management service is an advanced implementation of a concept similar to the "manageability client" of the WSDM standard. This service can be explicitly associated with a number of capabilities and implements mechanisms that allow it to change the configuration meta-data contained in service instances therefore adapting their behaviour at run-time. The management service may load different configuration templates from a repository, receive them by another service or by an external client (e.g. GUI). The Management service concept and design addresses functionalities that are currently missing from the vast majority of commercial and research prototypes.

Although the interfaces exposed by the Factory and Configuration Management service are based on emerging OASIS standards. The combination of the functionalities and techniques described in this paragraph bring offer management capabilities that have not been implemented as yet by any other commercial product or any other research prototype available at present (including Globus Toolkit v4).

---

[19] See Gamma, E., Helm, R., Johnson, R., Vlissides, J., Design Patterns © 1995, Addison Wesley.

7.6.1.3 Adaptive Enforcement

The enforcement allows for: remote configuration of security, monitoring and transactional properties of a particular endpoint (a.k.a. “service instance”); implementation of the enforcement configuration documented in an explicit configuration policy; and provision of the appropriate notifications in case of any significant event relating to monitoring or enforcement.

The technique builds on top of the service management capabilities described in section 7.6.1.2 and uses SOAP message interception techniques for enforcing actions on SOAP message exchanges their path to the destination. Actions are enforced by dynamically selecting and chaining handlers (a.k.a. interceptors) based on the contents of the intercepted SOAP message and the statement contained within the enforcement policy. The enforcement policy is assigned to every instance of the resource by the Configuration Manager via Management Interface.

Successful completion of the enforcement assumes that: the appropriate policy is retrieved form the resource property document; tokens are successfully validated; authorisation procedure conducted and results are used; handler chain of a new instance is properly configured; and the appropriate notification are generated and dispatched. In the case of a failure of an action (e.g. token/ signature validation, message part decryption, non-granted authorisation), the communication is terminated and the appropriate notifications generated and dispatched.

Although advanced prototypes such as the early versions of the forthcoming Apache Axis2 product (see http://ws.apache.org/axis2/) and ALLESTA (http://www.allesta.com/about/allesta.html), a California based, Silicon Valley start-up, are attempting to tackle similar targets, neither of these has been using the same techniques as we do, nor are they addressing enforcement for the purpose of dynamic Virtual Organisations.

7.6.2 Messaging

Work towards a Messaging Infrastructure for VOs focused on identifying and designing mechanisms that enable advanced notification, routing and endpoint naming features, as well as policies that are required for supporting standards-based, reliable and flexible message exchanges in an environment where transient service instances interact with each other, often asynchronously, in multiple and potentially diverse contexts.

7.6.2.1 Notification

We have identified the need to offer a standards based event management and notification infrastructure across the VO. Given the dynamic nature of the Virtual Organizations supported by TrustCoM, it is has been judged essential to generalize basic publish/subscribe notification systems to notification systems that can operate based on ontologies of notification topics.

A baseline web services notification subsystem has been prototyped. At present this is used in the context of the VO infrastructure and SLA management subsystems. In the next phase of the project and in addition to improving and integrating the above with a selection of other services that will become available to the TrustCoM project, we anticipate to design and prototype the following:

7.6.2.2 Messaging Service

We have identified the need for a collection of *messaging services* that may be needed in order to facilitate mediation in, or the management of, message exchanges between service instances participating in VO interactions. Such messaging services should enable:

- *message redirection*, i.e., to the ultimate recipient or an intermediary for further processing,
- *name-to-address resolution*, i.e. resolving a globally unique service name to the current address (EPR) of a service, and
- *message protocol implementation*, i.e. includes reliable messaging, evidence gathering, etc.

No detailed design or experimental development has taken place in this area as yet.

7.6.3 Registries and Meta-Data Repositories

We have identified the need for a federation of meta-data repositories to facilitate controlled sharing of information among the VO Infrastructure services, and for one service instance registry that maintains references to metadata relating to the identification, address, hosting and operational state of all the capabilities and active service instances, Neither of these should be confused with higher-level service registries such as UDDI and/or EbXML registries that are designed to maintain and manage information about business application services rather than infrastructure capabilities.

No detailed design or experimental development has taken place in this area as yet.

7.6.4 Common VO infrastructure transactions

In addition to providing basic functionalities for capability exposure, service instance management and adaptive enforcement, we identified a set of transactions that are required in order to implement complex interactions between supporting services. These transactions are realised on top of a collection of services (which we call Coordination infrastructure) that are dedicated to the enactment of Web service transactions that are implemented by means of effectively asynchronous message exchanges.

7.6.4.1 Coordination Infrastructure

This includes a collection of services that enable explicitly coordinated, distributed transactions that are implemented in an effectively asynchronous fashion.

In the first phase of the project we have adapted an open source implementation of the de-facto standards WS-Coordination and WS-AtomicTransaction as well as transaction types for some fundamental service instance life-cycle management transactions, such as those describedi n the following paragraph.

7.6.4.2 Service Instance life-cycle management

Recall that we view “service instances” as the combination of a Web service and a specific endpoint exposed by this service for the purpose of implementing interactions within a VO. We call the Web Service itself a “capability” in order to emphasise the fact that interactions with a VO take place only via an endpoint dedicated to this VO (i.e. via a “Service Instance of the Capability”).

In the first phase of the project we analysed, designed and developed a number of VO infrastructure service archetypes and transaction templates that contribute to the lifecycle management of service instances. These included a Factory service type, a Configuration Management service type, and an Instantiation service type, as well as, an Instantiation transaction.

The Instantiation service provides the means by which one can request the creation of a new service instance for an already deployed capability. The instantiation activates a particular instance of an Instantiation transaction.

The instantiation transaction implements distributed, explicitly coordinated and asynchronous transactions for efficiently managing interactions that bring together a number of different VO infrastructure services in order to implement the process of creating a new service instance. That is, creating a new manageable endpoint, configuring the endpoint with the appropriate enforcement actions, configuring all services that support the operation of the new service instances (e.g. token services, policy decision points, monitoring components, SLA management services, etc.) and configuring the bindings between the service instance and the supporting services. If the creation or any of the configuration steps fail then the instantiation transaction aborts if all succeed then the a reference to the endpoint of the new service instance (EPR) is returned to the requestor. At this stage all support services (security, SLA, processing, etc.) and the service registries will have been updated so as to support the operation of the new service instance.

Service instance destruction is an analogous transaction where all registry entries of the service instance to be destroyed are removed, the configuration of all supporting services is updated to exclude supporting that service instance and finally the endpoint of service instance is destroyed.

In the first phase of the project, prototyping is focused on the implementation of protocol for instantiation and destruction of a service instance. The implementation uses Open source Web services standards implementation on Apache Axis. This implementation of the instantiation process focuses on the configuration of the service instance and does not cover as yet the re-configuration of security token services, or services supporting higher-level policy, SLA monitoring functions.

7.7 Integration

One of the risks that we identified during the first phase of the project has been that of integration. Following the decomposition of the overall research challenges into specific targets and decomposition of the TrustCoM framework into six subsystems (i.e. VO Management, BP Enactment & Orchestration, SLA management, Trust & Security Services, Policy & VO infrastructure). We had to allow the teams addressing each subsystem to be able to focus on producing innovative designs and prototype implementation while maintaining a degree of consistency and convergence so as to alleviate the difficulties of integration when interim results have been produced.

In order to achieve this, we took the following actions within the first part of the project:

1. We created an internal representation (Figure 6) where we maintain information about

- the main services (“capabilities”) provided in the context of each subsystem
- the main interfaces these services expose
- the main dependencies between services – especially dependencies across subsystems
- the main info-sets that characterise information specific to a subsystem or information shared across subsystems, including
    1. message exchange scheme
    2. policy schemes
    3. main transaction templates

2. We have put in place a procedure whereby the organisations responsible for leading design and prototype implementation in each subsystem regularly update the information in the above representation and highlight changes that may have an impact in other subsystems. Major changes that affect dependencies can be implemented only if the directly affected parties endorse.

In the second phase of the project we expect to experiment with integration scenarios both in terms of the reference implementation and also, most importantly, in the context of the TrustCoM testbeds

Matrix of interactions

Read: ROW provides to COLUMN

Agree

Types of relationships:

N - notification

C - conversation (incl. Invoke)

? - Any ideas?

X - important capability for testbed

Please also refer to the accompanying Word document (**SubsystemTable-IvanD_v03.doc**) that contains initial input & the table used as the basis for the matrix below.

Subsystem		EN/VO Infrastructure						VO Mgmt					Business Processing				Contracts & SLA						Trust & Security				Policy			eLearning				CE							
	Service	Service instantiator factory	Service Instance (includes PEP)	Messaging Service	Notification service	Service instance registry/ information repository	Coordination infrastructure	Membership Manager	Context Manager (Role/Task Mgr)	Choreography Support	VO agreement/ partner profiles	Lifecycle manager	Collaborative BP repository manager	CDL++2BPEL service	BP TSC management service	BP engine	SLA template repository	Monitors	Evaluators	Negotiator	SLA Manager	Notary	Token services	Sec context coordinator	Reputation service	Trust Negotiator	Sec/trust broker	ECA PEP	PDP (authorisation)	Policy service	Training Consultant	Learning Resource Providers	Metacampus Discovery service	Metacampus Portal	PDD Service	Design Service	Customer Neg Service	Analysis Service	Project Mgmt Service	Storage Provider	Enterprise Messaging Service	NEC Service
EN/VO Infrastructure	Service instantiator factory		C		N		C	N	N		N				N	N C					N C			N			N	N?	?								X		X		X
	Service Instance (includes PEP)		N C		N		N C			N		N				N		N C						N C	N			C									X		X		X
	Messaging Service	C	C	C	C	C	C	C	C	C	C	C	C	C	C	C	C	C	C	C	C	C	C	C	C		C	C	?	C?					X			X	X	X	X	X
	Notification service	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N		N	N	C	N C		N	N	N		N	N	?	?			X	X	X			X	X	X	X	X
	Service instance registry/ information repository	C	C	C	N	N	C	C	C		C			C	C	C		(C)	(C)		C	(C)	C	C	C		C	C?	C?	C				X				X		X			X
	Coordination infrastructure	C	C	N	C		C		C	C		C				C		C	C	C	C		C	C			C		?			X						X				
VO Mgmt	Membership Manager	N				C,N						N		C	C								C	C	N	N				N	X	X			X	X	X	X	X	X	X	X
	Context Manager (Role/Task Mgr)	N	N									N			C					C	(C)		C	C							X	X			X	X	X	X	X	X	X	X
	Choreography Support	?												? N							(C)										X						X	X				
	VO agreement/ partner profiles		N					C				C,N		C	C		C			C	C									C								X				
	Lifecycle manager	N	N		N,C	N										N																						X				
Business Processing	Collaborative BP repository manager													C																	X	X							X			
	CDL++2BPEL service				N					C	C																				X	X						X				
	BP TSC management service				N				C							C															X	X							X			
	BP engine		?C		N		?C							C		?C			C	[C]			[C]	C							X	X						X	X			
Contracts and SLAs	SLA template repository							(C)				(C)								C		(C)						C											X		X
	Monitors		C		N													N/C	N	C	C							C	C										X			X
	Evaluators				N			N				N				C			(N)		C				N			C	C										X			X
	Negotiator							C				C								C																			X			X
	SLA Manager				C	C			C	(C)	C			(C)						C									?	C									X			X
	Notary																				C																			X			X
Trust and Security	Token services	C	C	C					C									(C)	(C)	C	C	C	C		C	C		N			X	X		X	X	X	X	X	X	X	X	X
	Sec context coordinator															C															X	X		X								
	Reputation service							X	C					C		C							?		N				C			X							X			X
	Trust Negotiator				C			C,N																																		
	Sec/trust broker																																					X			X
Policy	ECA PEP			?															C									N						X	X	X	X	X	X	X	X
	PDP (authorisation)			?C	N			C	C						C	? C		C	C				C		N									X	X	X	X	X	X	X	X
	Policy service	C		?																C	C?						C								X				X	X		X
eLearning	Training Consultant							X	X				X	X	X	X							X	X							X	X	X	X								
	Learning Resource Providers							X	X				X	X	X	X							X	X	X		X				X		X									
	Metacampus Discovery service				X																										X			X								
	Metacampus Portal				X		X																X	X							X											
CE services	PDD Service																																									
	Design Service																																		C							
	Customer Neg Service																																									
	Analysis Service																																		C			C				C
	Project Mgmt Service																																									
	Storage Provider Service																																									
	Enterprise Messaging Service																																									
	NEC Service																																									

Other	Subsystem	Required	
		Functionality	**Provided by**
missing	**VO Management**	AC/PEP	? Policy
	Policy	Policy/SLA compliance assessment mechanism	VO management
		Implementation for all policy actions to be called through ECA	All
		Policies	All
	Business Processing	AC/PEP	? Policy
	eLearning	AC Security policy	? T&S
		SP Federation	EN/VO
requirements/ assumptions	**EN/VO**	WS-core; WS-Notification; WSRF/ WS-Transfer; WS-S; WS-SC	N/A
		Distributed database	N/A
		UDDI server	N/A
		WS-Coordination; WS-Transaction	N/A

Figure 6: Internal representation of TrustCoM Framework, subsystems, services and dependencies

7.8 Open Standards

Standards are a way to promote and achieve interoperability between technologies across different vendors. While businesses need to balance between agreed functionality, competitive advantage, and need for interoperability, interoperability is a key requirement in today's multi-vendor market. Standardisation is an important part of successful exploitation. TrustCoM therefore aims at building upon existing well established and accepted standards and published specifications, where appropriate. If new technology is not compatible with existing standards that are well established in the market, then it may be more difficult to commercialize this into products and services which can interact with products and services provided by others. TrustCoM furthermore intends to contribute to the evolution of, and feed research results into, standards, where and in which way appropriate. The TrustCoM Standardisation Roadmap supports and documents the standardisation activities within the TrustCoM project, and is regularly updated throughout the lifetime of the project.

The first version of the TrustCoM Standardisation Roadmap (deliverable **D6**) established a first baseline for further standardisation activities, identified the standardisation areas which are relevant to the project, and provided an initial assessment of the state of standardisation in each of these areas.

The second version of the Standardisation Roadmap – **D24**, delivered at the end of the first phase of the project – gives a precise positioning status for each relevant standard and published specification, with respect to the first implemented version of the TrustCoM framework. While the list of standards and specifications relevant to TrustCoM has been growing during the last months, we have at the same time positioned the relevance of each of these specs in a fine-grained, qualitative way. Within each of the TrustCoM subsystems this had led to a clear identification of standards and specifications that have been adopted in the first version of the framework (as such, with restrictions, or adapted), standards and specifications that will not be considered, standards and specifications that may become very relevant in the second version of the framework, and standards and specifications that have popped up more recently and are kept within the relevance horizon for further investigation if time permits. This information intends to serve two purposes. We provide feedback to the standards world on the applicability of existing specifications within the TrustCoM framework and on the effective impact of standards on the different subsystems in TrustCoM. Furthermore, we inform the outside world of the standards choices made for the first version of the framework, in order to get feedback and to promote interoperability with products and services as well as research work in other projects.

Based on the positioning of the relevant standards and specifications in each of the subsystems, with respect to the first version of the TrustCoM framework, a forward look for standards impact to/from TrustCoM in each area is formulated, updating the broad standards assessments given in the first version of the roadmap, and concentrating on the envisaged adoption of standards in v2 of the framework (i.e., expected future impact from standards on TrustCoM), and on potential profiles or other specific standards contributions arising in each area – and particularly across areas – from the developments so far (i.e., potential envisaged impact from TrustCoM on standards).

The updated assessment in each of the TrustCoM standards areas allows us to select the areas that offer most potential for a significant contribution and to concentrate our efforts there.

The "Trust, PMI and PKI", "Contracts and SLAs", "Policies & Security", and "Collaborative business processes" standards areas are aligned with the TrustCoM objective of developing a framework for trust, contract, and security management, for collaborative business processing in dynamic VOs. The trust, contract, security, and business processing related components in the first version of the framework leverage specific standards in each of their areas, and during conceptualization, design, and development, potential standards contributions have been identified and refined.

The first version of the TrustCoM framework builds on top of selected, composable specifications in the "Web and Grid Services" area, having chosen (mainly driven by implementation availability) specific standards where overlap exists, and ensuring easy migration to other specifications where needed. These specific choices and restrictions may provide valuable feedback to standards.

TrustCoM will not be using (as not yet sufficiently mature to adopt) or contributing to (as not sufficiently core to the project) standards in the "Semantic technologies" area.

"Model driven security" is a separate area, yet closely linked to the other areas. Instead of addressing the functional and implementation aspects of the TrustCoM framework, this area addresses the design and development aspects that will be needed in order to easily build secure solutions using the TrustCoM framework. The project does not intend to start a standardization effort in this area. However, the security modelling languages developed are expected to provide a useful starting point for a potential future standardization effort. Experiences gained in model driven architecture may be used to validate and give feedback to emerging specifications.

During the first phase of the project TrustCoM has disseminated its objectives and some early results into various, relevant standards initiatives (including IETF, ISO, GGF, OASIS, OMG), and has maintained a liaison with other projects in its area through the Cluster Enterprise Interoperability around the ATHENA project, and the Grid Trust and Security concertation.

7.9 Application domains

Following an analysis of numerous scenarios in the selected application domains, the Consortium selected two scenarios, described in deliverables **D10** and **D21**, respectively. Although a critical mass of TrustCoM Framework services has not been integrated with these application scenarios as yet, the scenarios themselves exhibit distinct elements of innovation, relating either to the application of VO concepts in the respective domains or the re-engineering of their underlying ICT infrastructures so as to adhere to the basic principles of Service Oriented Architectures. In his section we summarise the main advancements to the state of the art that has been already achieved at the early stages of these testbed prototypes.

7.9.1 Collaborative Engineering

The main achievements in the first phase have been the development of an application scenario, the analysis of the scenario into collaborative processes, and the identification of the basic set of application services. A number of these application services have been implemented and are ready for testing in association with TrustCoM services. The services are currently deployed across three partner sites (Figure 7), as it is explained in TrustCoM deliverable **D10**: "Baseline Prototype Infrastructure for the CE Scenario".

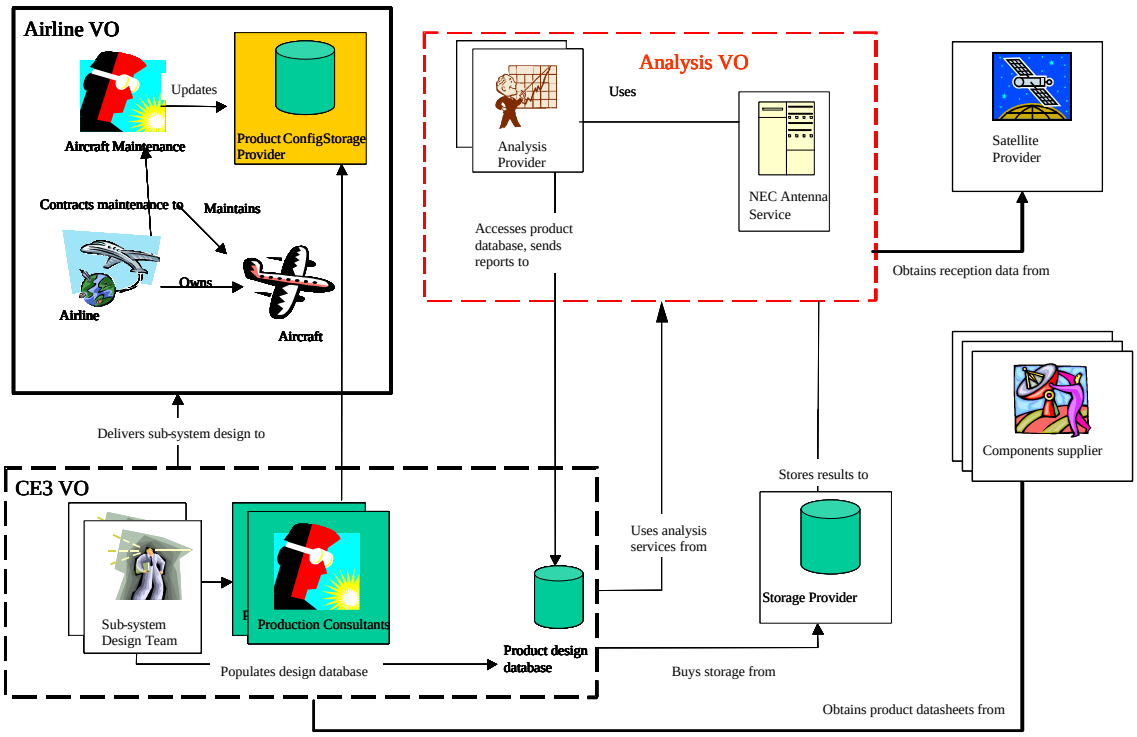

Figure 7: The CE testbed scenario with the main actors and the interactions between them.

At present the CE testbed advances the commercial state of the art in collaborative design by implementing a prototype of an advanced Grid-like system using a mixture of open source and "off-the-self" commodity technology. In research terms, however, it does not make advances to the state of the art in collaborative design. In effect the CE testbed is at present a "text-book" Grid use case but one which has implemented on commodity technology and still lacks the robust business-oriented security, trust and contract management infrastructure that is required in a commercial environment. Our intention is to introduce step-by-step this infrastructure (as described in section 8.9.1) by adapting and optimizing selected outputs of the reference implementation of the TrustCoM framework and therefore gain an insight on the potential impact and commercial value of these outputs.

7.9.2 E-Learning Services (ADP scenario)

The implementation done for the baseline prototype at this first phase has consisted in the distribution of the core modules of the Metacampus marketplace that enable their operation through web service communication. More specifically, a SOA (Service Oriented Architecture) approach has been taken in order to design the future interaction between these modules.

The next figure shows the new distribution of modules:

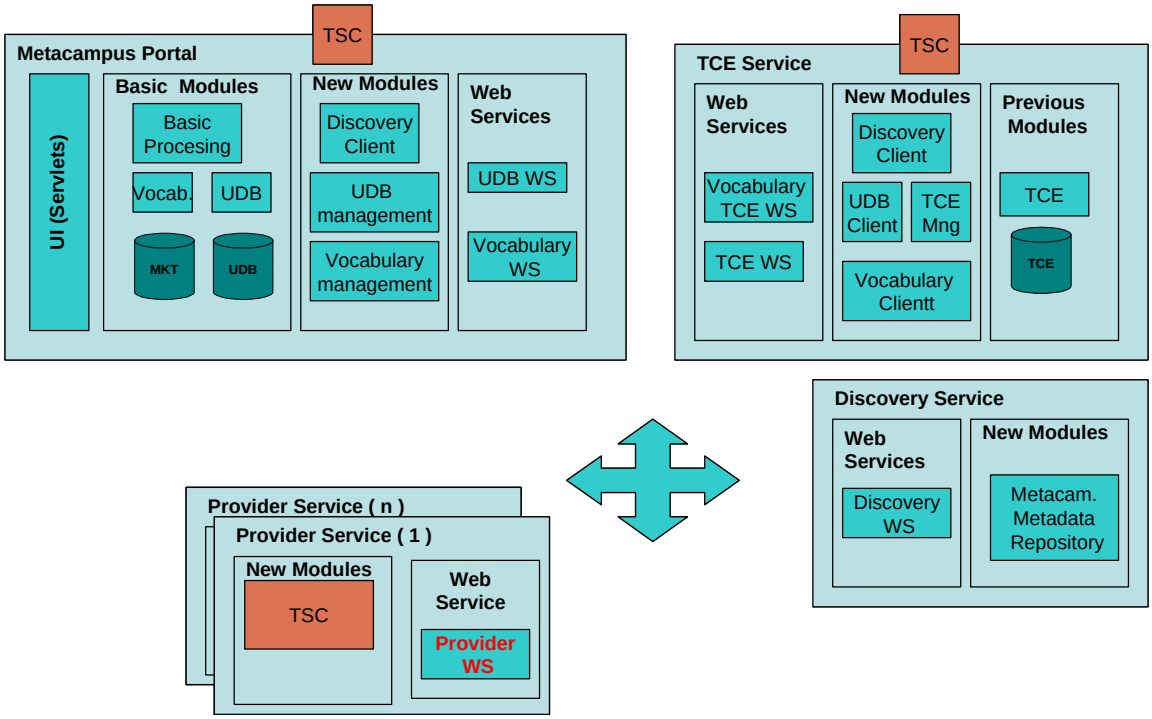

The baseline prototype does not incorporate the TSC module yet. This module will implement the TrustCoM framework and supporting services. Following this new distribution from the Metacampus original platform we now have 3 distributed services:

- The *Metacampus marketplace*, which keeps the User Database and the Vocabulary Database that stores the Metacampus specific metadata related to competences.
- A *Training Consultant service*, that evolves from a simple module within the marketplace to become an independent service. This step will provide strong benefit to the marketplace as many vendors could implement different Training Consultant services working with the marketplace in competition by offering aggregated services.
- A *discovery service* that will substitute the Learning Resource Catalogue, which will store the metadata description of the learning resources published by the service providers with reference to the access web service (stored in a UDDI).

For the baseline prototype, which does not interface any TrustCoM specific subsystem yet, the major enhancement that it has been re-engineered in order to follow Service-Oriented Architecture (SOA) principles and take advantage of a network-centric Application Service Provision business model. The separation of concerns between the main business functions, their encapsulation into services and their exposure by means of programmable Web service interfaces reflects how we expect future network-centric systems to be built and provides a fertile ground for making most out of the use of selected services from the reference implementation of the TrustCoM Framework.

In parallel and implicitly there are other novel service interfaces that need to be created in order to communicate these new distributed services.

All Learning Resource Providers (LRP) will also have to implement a web service to provide the requested learning resource when invoked.

7.10 Contextual objectives

7.10.1 Legal aspects

The objective of TrustCoM's legal activity has been to study selected legal issues in relation to trust, security and contract management for virtual organisations and WP 9's focus is on the legal risks that may arise for participants in VOs. The work has been performed in close collaboration with other TrustCoM partners, in particular in relation to the TrustCoM scenarios on collaborative engineering and e-learning. This work has contributed to the overall TrustCoM framework by defining some legal requirements for trust, security and contract management in VOs.

This activity applies legal risk management as a novel methodology to analyse legal issues related to trust, security and contracts in VOs. The utilization of methods from legal risk analysis has allowed the team to have a proactive approach on legal issues, which can be seen as opposed the reactive perspective inherent in traditional legal methods. Moreover, legal risk analysis has facilitated the integration of the perspectives of trust and security with the focus on legal issues related to virtual organisations.

The legal research performed so far falls into three categories: data protection law, intellectual property law and international issues. Specific legal issues within these categories were selected based on their relevance for the TrustCoM project. The risk analysis results indicate how legal risks can be treated through an integrated solution that joins together contractual elements, trust management and security management. The contractual treatments should consist of an adaptation of a contract template to the specific risks identified in the scenario.

With respect to the collaborative engineering scenario, the legal risk analysis focused on intellectual property rights and confidentiality. The legal analysis of the AS / eLearning scenario focused on legal risks related to international issues, i.e. choice of law and jurisdiction.

The workpackage is also developing methodology and tools for legal risk analysis. In particular, a graphical language is being developed which extends the Unified Modelling Language (UML) with relevant legal concepts, in order to make it more suitable for documentation and communication of legal risks and treatments. The utilisation of UML also ensures compatibility of the legal study with the work of the other TrustCoM partners. Work on formalising the language semantics is undertaken to make it more precise and to facilitate the development of automated tools for processing the graphical models.

7.10.2 Socio-economic aspects

Midway in the first year of the project, competitive game models were applied for VO selection and trust enablement between two parties. The models developed were focussed on individual trust models when compared to the requirements of the complex VO lifecycle management, which involves complex relationships between the VO members (group level trust). The game model was applied on a few attributes of the members and deeper insights into VO management were not revealed. One of the primary reasons was that the VO management framework was in the design stage and not fully conceived. Furthermore, focusing overall the priorities of the Consortium towards areas where a significant impact could be achieved in shorter timescales, meant that the Consortium were more interested on investigating the Business Models that underpin VO configurations where different trust relationships can be identified as well as exploring business contracts, an important area that could not be fully addressed, in the more technical part of the TrustCoM Framework (cf section 4.4).

Following the TrustCoM review done in April, 2005 the socio-economic team re-assessed their plan with the help of the TrustCoM Scientific Coordination and Programme Management, and the Consortium agreed on a set of modified objectives for November 2005 towards models of Reputation, Business models for Trust and others. The final modified objectives are as follows:

1. Explore advanced multi-tier Models of Business Contracts for VO Management

2. Economics of Business Contracts and Terms and Conditions for Reputation
3. Investigate models for Reputation based on metrics defined around contract terms and conditions. Investigate advanced scoring models for Reputation.
4. Investigate Business Models for Trust and Interoperability between VO members and other VO organizations.
5. Investigate Business models for Trust in third-party neutral or dominant group environments.

The socio-economic research team worked closely with AL1 and AL2 sub-projects and teams on Business Contracts, Terms and Conditions, VO management and Business Metrics for Reputation and VO supply-chain models. Work on Business Contracts has contributed knowledge and criteria to the TrustCoM framework subsystems. Currently a working group has been established to investigate role of Business Contracts in VO supply chains and in designing reputation mechanisms. The IBM team also intends to provide advanced knowledge, definitions and mechanisms around Business contracts, which will encompass the SLA work that is being done elsewhere in the project.

The team also provided input on “Generic Reputation Service” which is an important part of the VO lifecycle management and Trust/Security Services. The input has been on supply-chain metrics and contract attributes and management for building an industry driven reputation system. The socio-economic research team has also provided input to VO management (on Reputation scoring) on reputation models and scoring functions for VO members in a VO environment.

During the first phase of the project, the team working on socio-economic aspects developed a model of Business contracts for enabling VO supply chain interactions based on terms and conditions between VO supply chain partners. This is detailed in section 3 of D14: Report on socio-economic models

The team also developed novel reputation model based on Business Contracts and contract-specific terms and conditions. The reputation model is based on monitoring contract terms and conditions over a long-period of time in order to score and rate VO members. The contracts signify the agreed upon terms, which if violated the conditions apply. Business rules can be set by the VO members on the violations to understand and select out the VO members. This has been accompanied by new models for scoring based on contract attributes and functions for VO member reputation. The attributes for reputation are based on rules applied to the terms and conditions. For example, in a VO supply chain if a VO member violated a specific term and condition 5 times then the scoring function will rate the VO member based on the type of term and condition. If multiple terms and conditions are violated the scoring function considers multiple attributes and weighting functions based on the semantics and criticality of the violations.

Economics and business models for Interoperability were developed as a part of the project’s socio-economic research activities. The models for interoperability consider trusted third-party, trusted consortia and trusted group models. The business models were compared and contrasted around various degrees of trust and reputation (this study was presented at a cluster interoperability workshop).

Other contribution includes offering advice on standards for business contracts, models for the Enterprise Interoperability Cluster of FP6 projects.

8 Recommendations

In this section we conclude the self-assessment presented in this deliverable by summarising specific actions that have been recommended to the TrustCoM consortium for the second phase of the project in order to for the project to tackle the research challenges outlined in section 3.3 and the specific scientific and technological development goals set in section 4. Where appropriate we highlight goals that fall outside the programme of the project and will have to be addressed by other research projects.

8.1 VO Management

In relation to VO Management, the TrustCoM consortium should aim at producing three components covering the essential requirements of a VO Management system: lifecycle management, membership management and GVOA management. Such components will enable the management of dynamic VO, taking into account trust and security constraints.

It has been identified that no VO Management product exists today covering these important issues (see Section 4.9 of deliverable **D25**, the exploitation plan). In order to realise this vision, we will implement a VO management subsystem covering the full cycle of a VO, concentrating in particular on the evolution phase of the VO lifecycle.

The relevance of semantic technologies for VO management has been identified. Semantic technologies have been found particularly useful for representing VO structure and dependencies between VO partners as well as offering a common foundation for processing. Current resources and priorities in TrustCoM do not allow an adequate investigation of the use of semantic technologies for VO management and GVOAs. It is recommended that further research relating to the use of semantic technologies for modelling VO management structures and GVOAs is conducted outside of, and where appropriate, in relation to TrustCoM.

8.2 Collaborative Business Processes

The collaborative business process models identified so far should be refined during design and implementation in two directions. First, the operational components, such as services, have to interface with other subsystems. These interfaces need to be defined, reviewed, implemented, evaluated and improved based upon experience and feed-back. Furthermore, a common and simple deployment process for all components including the additions to existing technology in the form of TSC extensions has to be developed. The TSC extensions also have to interface with other subsystems. Again, these interfaces need to be refined. Details of TSC extensions, such as their representation in public and private business processes need to be defined. A final implementation of the BP Enactment and Orchestration subsystem with the identified components and interfaces is supposed to be finished at conclusion of the second phase of the project.

Because of a change of priorities within the consortium and also because of the immaturity of technologies relating to technology for implementing choreographies of Web services, it has not been possible to properly address self-adaptive processing at all layers. It is recommended that further research relating to self-adaptive implementation of choreographies is commenced outside of, and where appropriate in relation to, TrustCoM. In the context of TrustCoM it is recommended that the focus shifts on the use of mechanisms for adaptation policies (such as those presented in section 7.5) and distributed transaction mechanisms (such as those presented in section 7.6.4.1) in order achieve some level of adaptation during the enactment of orchestrated business processes. Further research will also be required in order to correlate business process goals with security policies and SLA obligations. Also in this area the complexity and immaturity of enabling technologies in combination with limitations on resources has necessitated scaling down the ambition of the original project objectives, which however remain relevant research challenges.

8.3 SLA Management

By the end of this project, we want to be able to support the full “lifecycle” of a service level agreement between the service provider and a customer, respectively the virtual organization – this covers provision of SLA-related information about a service, negotiation of SLA terms, configuration of the involved components, enactment of the SLA (monitoring and evaluation) and finally “unbinding” the service provider again.

This will allow service providers and consumers to offer, respectively exploit services that meet a specific quality of service and to ensure that this is maintained during enactment.

We furthermore want to enable service providers to manage their services with respect to a specific quality, but TrustCoM will only deliver the basis for this and not examine the required intelligence of such an autonomous self-management component.

To realize these issues, we pursue the following goals:

- Extend the SLA template schema so as to allow for discovery of services on the basis of a quality of service description and to act as a premise for negotiation
- Implement a template repository that can be queried for discovery
- Develop negotiation services that are capable of agreeing upon SLA parameters using a fairly simple protocol. The discovered SLA templates will serve as a basis for negotiation.
- Support notary services with the ability to plug in various SLA signing protocols
- Integrate the SLA lifecycle with that of the VO, exploiting the relationship with the General VO Agreement (GVOA) and including the discovery and negotiation of SLA Management services in the corresponding VO phases.
- Realize an SLA Manager service that allows for automatic configuration (i.e. binding and “unbinding”) of the involved components based on the agreed upon SLA. An SLA Manager will in fact take the form of a federation of services, operating on different administrative domains.

The relationship between SLA obligations and business rules in terms of common foundations, common semantic representations and common enforcement and monitoring mechanisms will not be fully addressed in this project. Investigating such a relationship has been among the initial objectives of this project and it is important for fully tackling its original research challenges. However the limitation of resources combined with the immaturity of open source technologies for monitoring the execution of simple contracts such as (Web) SLAs, necessitates reducing the ambition of the original research objectives in this area. Nevertheless the Consortium recognizes the importance of this research objective and recommends that community supports further research in this direction. Such research will also have to take the legal implications of technology-driven automation into account.

8.4 Trust & Security

8.4.1 Security Token Services

Our main goal for this project is to implement an intuitive security management model that support humans in controlling *their* part of a virtual organization. In today’s IT systems, it is difficult to achieve a constant level of security. The limiting factor in maintaining consistent system security settings is the complexity of the security management. It is increasingly difficult for users and administrators to understand the implications of their daily security decisions. We already have designed a security management model and have implemented a first proof-of-concept prototype. The configuration of this first prototype is rather static, which shows that the cross-organizational service invocations work. However, it does not yet fully support the security management aspects. In the second phase of the project, we will provide management functionality for security token services and improve the implementation, as outlined below.

- We will use interoperable security token types. It is critical for the success of TrustCoM that we provide security token formats for cross-organizational service invocations, which are usable within a wide range of scenarios. In the first prototypical implementation, we used simple self-defined claims and attributes inside the security tokens. We will replace these claims by already existing claim and token formats, such as SAML/XACML, in order to re-use existing technologies as much as possible.
- The administrative interfaces in the first implementation of the security scenario are purely web service based and only support management through custom applications running on the client. In addition to these web service interfaces for autonomic security management, we will create web applications (HTML interfaces) that support business people with the administration of security-related aspects of their part of the VO.
- The first security scenario STS prototype (as implemented in mid-July 2005) solely performs authorization as part of the STS, i.e. we did not yet integrate the authorization process with the policy decision points (PDP) provided by the policy work package WP32. The next step before the review demonstration in October 2005 is to link these building blocks together. The implementation will support multiple deployment alternatives, i.e., we will support multiple interaction patterns between policy enforcements points, policy decisions points and STSs.
- The STS-internal configuration currently is based on plain text files, i.e. the management implementation currently does not provide sufficient and fine-grained access controls. The internal STS data structures should persist inside a scalable and manageable system, such as a database.

In this area TrustCoM anticipates offering fairly mature and generic technologies for supporting secure federation and the distribution. TrustCoM will not fully tackle however schemes for representing, processing and translating such credentials. Again this is an area where semantic technologies may be able to provide effective solutions. It is recommended that the community supports further research on languages and models for representing credentials in the context of dynamic Virtual Organizations, as well as, on techniques for transforming such credentials across trust realms.

8.4.2 Reputation service

Additional functionality that needs to be included in the system is the creation of an *electoral role*. This is important as it will allow the reputation system to record which actors are allowed to express the reputations of other actors. This will be achieved by linking the VO management and SLA systems into the reputation management system. We also want to *link the reputation system into the authorization system*, so that authorization decisions can be made based upon the changing reputation (or trustworthiness) of the principal. Current authorization systems (also known as trust management systems) only work in binary mode, in that a principal is either always trusted or always not trusted. Linking this to a reputation management system will allow much finer grained authorization to take place based upon dynamic control of trustworthiness.

Finally we want to define the *TrustCoM reputation metrics*, which will calculate a partner's reputation based on how well they have performed in the context of VO transactions. This will require the definition of the attributes, their possible weightings, the algorithms for combining the weightings, and interfaces for setting and retrieving reputations based upon them.

Although the TrustCoM consortium anticipates providing the archetype of a generic reputation system and implement protocols that enable the collection of evidence and other information (e.g. recommendations) that is consumed by the reputation system, the consortium will not conduct in-depth research on the role of reputation in Virtual Organisations and/or on optimal models for computing and representing reputation information. We anticipate making a proposal in this field with the understanding that the further research will have to be conducted in order to validate, revise and improve this proposal.

8.4.3 Secure Audit service

By the end of this project, we want to be able to provide the secure audit web service through all three interfaces: Java API interface, a TLS-protected TCP/IP interface and a web service front end interface. A high logging efficiency is an important factor for the successful application of SAWS in virtual organisations. We hope the minimal logging speed for SAWS can be 500 log records per second.

Currently, symmetric secret keys and private keys are used in SASW and they are protected by a Java KeyStore in the format of JCEKS. Obviously this is a software-based key management method, and it may pose a security weakness in reality. The security of the keys is the most important security issue in SAWS. We will look into the possibility of using a hardware-based key management method based on Trusted Platform Module, so as to increase the security level of the key management in SAWS.

Security enforcement and auditing are two areas where Trusted Computing Platform technology may play a significant role as an enabler of innovation. Although in the context of the project we will conduct preliminary research on this field, we strongly recommend that a further in-depth investigation is conducted outside of the scope of the project. Such an investigation will have to take into account the legal implications of using Trusted Computing Platform as one enabling technology that underpins ICT infrastructures for Virtual Organisations.

8.4.4 Trust negotiation engine

A wider diffusion of trust negotiation mechanisms requires solutions for several issues. The first step is to disclose, as far as possible, the requirements that both negotiation parties are equipped with the same negotiation engine. To this end, it is necessary to separate the disclosure policy-driven negotiation engine from the protocol that the two parties conduct the negotiation conversation through. At the same time it is necessary to express and publish the negotiation protocol in a suitable way, so that a client can automatically generate at least parts of the code needed to implement the negotiation 'stub'. Secondly, we must enhance the flexibility of the trust negotiation, by defining different negotiation strategies. The third issue is to improve the protection of the disclosed policies.

In order to solve these issues, we recommend the following steps:

- To address the first issue, we plan to expose the trust negotiation functionality using standardized protocols, or extensions of them. Candidate protocol standards are WS-Trust, WS-Policy and WS-PolicyAttachement.
- To improve the adaptability of the trust negotiation and the protection of the disclosed policies, we plan to extend the current prototype by introducing new negotiation strategies. In particular, these strategies are '*suspicious'*, '*strongly suspicious'* and '*trusting'* strategies:

  1. If the '*suspicious'* strategy is adopted, the credential proof is always requested during the policy evaluation phase for each of the involved credentials.
  2. The '*strongly suspicious'* strategy is a specific case of the *suspicious* strategy. Under the '*strongly suspicious'* strategy, parties require attribute disclosure as the corresponding policies are satisfied.
  3. The goal of the '*trusting'* strategy is to speed-up the negotiation process whenever possible. This can be done using credential *suggestions*, to be described (and stored) in an additional component of the negotiation policy. The main advantage of this strategy is that, if used for all involved policies, it reduces by half the number of rounds to complete a negotiation.

  Overall, the second phase of the project will focus on the implementation of an engineered version taking into account the above mentioned points.

  Further research to clarify the use of Trust Negotiation in virtual organisations is required. We anticipate making a proposal in this field with the understanding that the further research will have to be conducted in order to validate, revise and improve this proposal.

8.5 Policy

Whilst the first phase of the project has focussed on conceiving the policy framework, its architecture and principles of operation, substantial work remains to be done towards its implementation. Currently, simple obligation policies can be defined and enforced and simple policy groupings can be realised. However, substantial work remains to be done in terms of:

- Distribution, implementation of relationships and interactions across policy interpreters.
- Precise definition of the policy specification language and adaptation of the existing toolset to the new notation.
- Translation between the policy notation and the XACML format used by the policy decision point.
- Working out the details of how dynamic attributes, policy and attribute revocation and organizational dynamics are best modelled and handled in the context of delegation and policy administration.
- Enactment of the two project scenarios in this framework.

Thus, in summary, the second phase of the project will focus on the implementation of the policy sub-system as described in deliverables **D16**, **D9** and **D18**. This may require refinement of the models and architecture where appropriate.

8.6 VO infrastructure

In the first phase of the project a number of infrastructure capabilities have been identified. If these are implemented and integrated, the VO infrastructure developed in TrustCoM will offer substantial technological advancements compared to the majority of existing ESB and service management platforms. The main VO infrastructure capabilities that have are still in early phase of design and prototyping include:

- The messaging infrastructure, including support for advanced message routing and support for implementing different types of reliable message exchanges, based on explicitly defined and extensive messaging policies.
- Mechanisms for providing service registries and federated meta-data repositories that maintain the configuration of and allow seamless access to the management interfaces of potentially large groups of transient service instances.

Further to the above, the existing VO infrastructure capabilities will be refinement. Refinement during the second phase of the project will focus on tackling the following priorities related to the VO infrastructure:

1. Improving the design and implementation of the adaptive enforcement mechanism by:
    a. Analysing the dependencies between different categories of interceptors and reflecting these in the Configuration Policy in order to improve the reliability of the infrastructure.
    b. Integrating the existing mechanism with services that implement higher-level adaptation policies. Such integration is a high priority because, if implemented, it has the potential to realise a major technology breakthrough towards providing a technological foundation for self-adaptation during VO operation.
    c. Selectively including interceptors to assist SLA monitoring and process enactment wherever these are made available by the teams working on the corresponding subsystems.
2. Improving service management mechanisms by exposing advance interfaces for human administrators and by enabling the efficient, automated management of potentially large groups of service instances distributed across multiple hosts.

3. Extending the basic transactions that underpin life-cycle management so that they include configuration of various supporting services, such as registries, token services, monitoring and SLA evaluation services, policy decision points and adaptation policy services.
4. At present the transaction mechanisms can deal mainly with interactions between persistent services (although these can interact for e.g. managing the life-time of transient service instances). Implementing transaction mechanisms that can deal with transient service instances requires a technology breakthrough that is within the scope and the abilities of the TrustCoM consortium. Further research in this direction will be conducted in the second phase of the project.
5. Capability exposure and service instance life-cycle have been implemented at present mainly for application services rather than supporting services, which are assumed to be persistent. Improving these mechanisms so that they can also be used for infrastructure or other supporting services will enable the creation of distinct instances of infrastructure capabilities for different Virtual Organisations. This is a major challenge that is yet to be addressed and an essential technology if the operation of multiple Virtual Organisations over the same ICT infrastructure is to be achieved.

Finally the different capabilities of the VO infrastructure need to be brought together so as to form a coherent platform that offers advanced ESB and service management capabilities.

8.7 Integration

Integration is the outstanding challenge that has yet to be met by the TrustCoM consortium. The following actions are recommended in order to facilitate integration within the scope of the activities relating to the TrustCoM framework and its reference implementation:

1. The team working on the architecture revisits the designs of the services that have been developed so far and defines (in conjunction with the corresponding workpackages) basic transactions that span across the VO infrastructure and (trust, security, SLA) supporting services. Examples of such transactions include:
    a. Transactions that underpin the life-cycle management of service instances; this includes creating instances of policies that apply to the new instance and configuring the necessary policy decision points, identifying the necessary SLA and configuring the corresponding SLA monitoring and evaluation services in order to support the operation of a new instance or deactivate an operational service instance endpoint and implement the graceful destruction of that instance.
    b. Transactions that underpin reconfiguration or update of trust & security services (including enforcement, security token services, policy and reputation) in reaction to an SLA violation
    c. Transactions that underpin adaptation to the change of the level of reputation of a VO partner or of the state of trust relationships between different VO partners.
    d. Transactions that underpin the life-cycle of secure federations across the trust realms of several VO partners
    e. Transactions that underpin the distribution, enactment and adaptation of a collaborative process. Where adaptation comes in response to an SLA violation or a security failure.
    f. Transactions that underpin major changes to the life-cycle of a VO including formation, engagement of a new partner, disengagement of existing partner and dissolution.
2. The teams working on the different subsystems implement these transactions in a bottom-up fashion. This requires selectively integrating services on top of a common ICT infrastructure, implementing transactions for realising complex interactions between these services and adapting their interfaces where appropriate, and proceeding to the integration of a layer above once an adequate level of integration has been achieved at all levels below.

3. Identify selective integration that add value to each application scenario and try to apply them in the context of enhancing the corresponding application scenario testbed.
4. Understand the implications of the integration on the collection of Open standards technologies used as a technological base-line in each case.
5. Understand the implications of the results of the legal and socio-economic research, and where appropriate, implement a selective take-up of these results in whenever they clearly add value;

Meeting these distinct integration milestones is a major outstanding research challenge for the TrustCoM consortium.

8.8 Open Standards

A broad set of selected standards and specifications have impacted the development of the components in the first version of the TrustCoM framework. Potential standards contributions have also been identified as part of these developments. The next steps for pursuing standards contributions include:

- continue dissemination to, and collaboration with, other projects in this area, in order to promote the TrustCoM technology, get feedback, and further align work;
- further integrate and mature the components into application testbeds as to get a more in-depth common understanding and a first validation of the results within the project;
- concentrate on integration profiles, bringing together the isolated subsystem developments; while we have refined the potential standardisation contributions *within* each specific TrustCoM research and development area, the most immediate result of the TrustCoM standardisation activity is expected to be in the integration of existing standards *across* the different areas.

It is important to emphasise that TrustCoM is an *integrated* project addressing trust, security, and contract management, for collaborative business processing, as a whole, focusing on the relationships and interactions between, and integration of, these issues, rather than investigating each of these issues separately and independently. The primary focus of the TrustCoM standardisation activity is expected to be in the creation of profiles that integrate existing standards *across* the different areas. While there are already numerous specifications addressing various issues within most of the identified areas, there are almost no concrete guidelines at all with respect to combining different specifications into a single interoperable framework.

8.9 Application Scenarios

8.9.1 Collaborative Engineering

The main outstanding issues are:

1. The development of additional application services to support more extensive collaboration services.
2. The use of the application services within the TrustCoM framework.

The additional application services will cover other aspects of the design-cycle such as customer negotiation, design, document review and a further decomposition of the Analysis phase into sub-tasks which could be distributed over different partner sites. This would make the testbed much more realistic and relevant to current aerospace procedures. Other services which offer information and data services will also be deployed. These include a satellite reception service and a product information service provided by a component supplier.

The concrete targets for Phase 2 are therefore:
1. implementation of additional application services outlined above, and

2. the initial deployment of TrustCoM services to implement a first version of a TrustCoM Virtual Organisation.

Regarding the latter of the two targets, the following sections indicate how the TrustCoM sub-systems will be used within the CE testbed.

- **EN/VO infrastructure.** We expect that for the CE scenario the EN/VO infrastructure will provide the following capabilities:
  - Service Instantiation: The NEC Antenna service is a good first candidate for demonstrating how a 'virtualisation' of services can lead to a more flexible VO operation that responds to varying customer demand. Issues to be addressed in relation to service instantiation include:
    - Establishing and managing the logical identity of the service instance
    - Offering values to agreed metrics for monitoring the service performance
    - Supporting the lifecycle of the application instance
    - Managing relationships to other service instances
  - The Policy Enforcement Point. The use of PEP will be important in enforcing security policies- a high priority in Industrial applications. The PEP of the service instance will intercept SOAP messages in order to:
    - Monitor invocations of the HPC Service
    - Enforce access control decisions- particularly important for the PDD, where different sections of the PDD will be subject to different access policies
    - Update the message header with new credentials and contextual information, eg, in accessing the PDD service or the product information service offered by a potential new supplier.
  - Notification: A certain number of services have an 'asynchronous' behaviour. A primary example of this includes the Analysis Service where invoking the service activates an internal business process that can be monitored using a 'handle'. This kind of behaviour is typical of many business processes where an incoming order is processed within some internal business process. In these situations, other services in a business process would wish to be notified when the associated task completes.
- **VO Management.** The CE Testbed will make use of the following components within VO Management:
  - Membership Management
  - Context Management (role & task management)
  - Choreography support

  In the current scenario, the assumption is that the CE VO is a pre-existing entity that accommodates new members and releases a number of these when the scenario goal is achieved. The management of this without the support of VO Management services would be expected to be highly labour intensive and time-consuming. Future editions of the Testbed will consider more interesting VO lifecycle use cases with more dynamic membership scenarios. For example, the members who are released at the end of the scenario include the Analysis consultancy and the NEC Antenna provider services who would be expected to receive payment for their services. Other transitory members do not expect payment, and these include the commodity component suppliers that join/leave the VO as the suitability of their offerings are assessed during the collaborative design process.
- **Business Processing.** Business Processing will be investigated in earnest in future versions of the CE testbed. In the present version of the CE testbed the business processes are only formally defined and managed 'out-of-band' of the collaboration. A more extensive design process will be implemented that will attempt to exercise this and the choreography aspects of the framework

- **Contracts and SLA**. The availability of reliable HPC resources are critical for the timely and accurate prediction of candidate designs. The introduction of the SLA will provide a step change over the current Grid model where issues such as SLA are resolved informally or by intervention of operatives. These the main services provided by the SLA Management module that will interface the application services:
  - **Monitors**: Relevant at different levels of the VO, depending on the functionality they provide: *Monitoring Interface* (so-called "Data Provider"), *Host Monitor*, *Domain Monitor***,** *VO wide Monitor* (so-called "Monitoring Aggregator"). In the CE scenario, and for the first iteration of its implementation the ***response time*** metric[20] will be critical for testing its functionality. Notably values for this particular metric can be provided either by explicit SLA monitors or by interceptors embedded within the enforcement mechanisms associated with the VO infrastructure.
  - **SLA Template Repository and Notary services**: both supporting services will store the SLA template associated with each kind of service provided. In the first case, the Service Provider, when publishing the service in the UDDI will also publish the associated SLA template with the relevant metrics and Quality of Service that the service offers. The Notary service will store the signed SLA agreed in the negotiation phase between the NEC Antenna service provider on behalf of the end user.

- **Trust & Security.** The general benefits are expected to be the management of resources within the CE VO that ensures data and service access to selected roles. This is a difficult and challenging topic for businesses that have extremely valuable IP but who wish to participate in different collaborations.

  The following table summarises the most important requirements that the application needs from the security & trust subsystem:

Entity	Requirement
End User & Service Providers	Single Sign-on Authentication Authorization and access control Reputation
Data (User Profile, LRP metadata)	Encryption
Orchestration (BPMS)	May require federated trust management May require data-level encryption and integrity control
Transactions	Digital signatures for non-repudiation (eg, in the case of uses of the NEC Antenna service)

  The issue of suppliers' Reputation will also be a topic for investigation. Reputation is important in establishing the most suitable supplier that could meet a sudden change in demand for a resource in the VO. A good example of this would for HPC resources, where one could have a large number of potential NEC Antenna service providers. Faced with a large number of suppliers who have the same technical capabilities, one would have to consider other factors that may be based on 'trust' or 'reputation' measures. Using Reputation one could have a much more open business environment in which new start up companies can quickly participate.

[20] ***Response Time:*** Response time is the time it takes the service to respond to a specific request. Normally, this is measured as the difference between the time the request was *sent* and the time the response was *received*. However, as this involves third party monitoring which is currently not present, we here simply return the temporal difference between the *ingoing* and the *outgoing* message – likewise, this value represents rather the processing time than the response time, but the metric itself can be easily reused for the first type.

- **Policy.** In addition to the specific targets related to authorisation and access control summarised in the pervious paragraph, the ability to make high-level policies that can be enforced throughout the infrastructure of the CE VO will be another important step change in capability.

8.9.2 eLearning Services

The Metacampus marketplace will be clearly enhanced with the TrustCoM framework services which will provide the current solution with the following new features:

- Adoption of a Service Oriented Architecture. Metacampus services will be deployed in different domains (service aggregator, Training consultant...)
- Learning Resources information distributed over different Service Providers
- Provide the application with trusted and secure services from Content Providers
- a new Discovery service to get up to date information of learning objects at the Service Provider sites
- Provide the Metacampus users with a more reliable and secure access control system
- A more dynamic and evolving service provision. Service providers will be federated as virtual organisations and the aggregation of services will clearly be more dynamic and efficient.
- A more flexible system for federation and registration of services
- A more powerful business process system to manage the provision of aggregated services
- A more user-oriented vision based on the negotiation and agreement of the service provision and quality of service
- Flexible and dynamic association between the different actors (SP, Training Consultants, users)

8.10 Socio-economic aspects

The following three major conclusions have been made through the socio-economic studies so far. Each of them requires a follow-up action.

1. *Economic models play a strong role in enabling trust mechanisms*. D?? describes the various economic and business models for enabling trust in third-party environments. The major result is that trust between parties or players is better when keeping a history of transactions. This was proven in several experiments that were conducted.

   Follow-up action: This result is being taken into consideration by the Trust & Security services subsystem where Audit and Reputation services have been introduced.

2. *Contracts are the life-line of building trust in Business Environments and VO supply chains systems*. Contract terms and conditions provide a tremendous foundation for Trust and Reputation management. The major result from the work is the design of a novel contract driven and attributes based reputation rating of partners in a VO supply chain. The reputation rating models consider attributes and criteria that are semantically driven.

   Follow-up action: the Reputation service considered under the Trust & Security workpackage should take these models into consideration and the SLA monitoring subsystem should offer the necessary information throughout the SLA life-cycle.

3. *Business Metrics based on contract terms and conditions are critical* for evaluating the reputation of VO members, monitoring the contracts terms and ensuring the proper enforcement of the terms. These measurements will provide feedback into the generic reputation system models (section) for performing rating of members and new VO creation. The same applies to the subsystem addressing VO management and GVOA in particular.

   Follow-up action: the team working on the VO Management and SLA Management subsystems should consider the option of explicitly associating business metrics to certain terms and conditions described in the GVOA and SLA templates. The socio-economic team should refine the model and assist the other teams in defining concrete examples of such metrics.

8.11 Legal aspects

During the second phase of the project, the team working on legal aspects will continue to focus on the TrustCoM AS testbed scenarios, and in particular on access-rights management and also on the liability of VO participants in relation to access rights, e.g. how the access to resources in a virtual organisation may be legally classified and the effect this may have on the liability of its members. Legal risk analysis will be used both to identify legal risks and treatments relevant to VOs and to evaluate the suitability of the tools, method and language. Moreover, the work with methods for legal risk analysis for VOs will be finalized.

Regarding the TrustCoM CE scenario, the research performed in phase 1 identified the need for a more detailed legal analysis of confidentiality issues. This research will aim at identifying the risks to information that VO partners will have access to, with respect to (i) illicit access to confidential information by VO members or third parties, (ii) illicit dissemination of confidential information to entities that are not entitled to access the information. For a VO using the TrustCoM technology, the challenge will be to integrate the access based on policies as defined in the TrustCoM framework with the legal protection of confidential information. In this context the legal protection of confidential information in selected statutory laws will be analysed and the need for additional contractual clauses will be assessed. The analysis will include, for example, what is to be understood as confidential, how this information is to be identified in a web-services environment, how long the protection should last, procedures for lifting the limitation and procedures for exceptional circumstances.

9 Conclusion

As mentioned in section 2.1 of this deliverable, a interim progress assessment conducted by the project's scientific coordination and management teams during the first year of the TrustCoM project identified as a major risk the fact that the Consortium could be spending unreasonably large effort and resources getting deeper into a vicious circle of analysing dependencies between the various aspects of the TrustCoM framework at the expense of producing interim results in any of these areas. Consequently it was recommended that the project structure is drastically changed in order to achieve a clear separation of concern between the main aspects of the TrustCoM Framework and focus on producing a first round of tangible results in each area. In turn this brought about a major project restructuring that has been unprecedented for a collaborative project, especially if one takes into account that this restructuring was implemented following an internal project initiative and not an external review.

The restructuring of the project plan and re-focusing of work in specific self-coherent sub-areas has been successful to the extent that the new teams formed focused on delivery within their respective areas of expertise and, in particularly tight timescales, the Consortium produced substantial advancements to the state of the art, and in many cases we managed to place ourselves ahead of our contemporary research trends.

Notwithstanding this success, it is now felt that the TrustCoM consortium needs to find a way to build on this early success by successfully integrating their results instead of falling victims of their success by focusing on perfecting our current partial solutions. For one of our main objectives – one that is particularly difficult to classify in any specific research area and has been a major motivation for bringing this Consortium together – is to produce a comprehensive framework in order to overcome shortcomings of previous attempts which fail at the boarders of the self-coherent albeit partial solutions they offer. In addition to improving the solutions in each subsystem by filling (at least) the identified gaps, the TrustCoM consortium needs to put in practice the so far top-down attempts to maintain cohesion across the TrustCoM subsystems, application scenarios and contextual research. In particular, the TrustCoM consortium needs to:

1. Understand the results of the legal and socio-economic research and selective use these results in those cases that they clearly add value.
2. Design and analyse integration scenarios for the TrustCoM framework subsystems. Then implement them bottom-up, starting by selectively integrating services on top of a common ICT infrastructure, implementing transactions for realising complex interactions between these services and adapting their interfaces where appropriate, and proceeding to the integration of a layer above once an adequate level of integration has been achieved at all levels below.
3. Bring appropriate integrations of the TrustCoM Framework subsystems into the application scenarios and conduct a second of integration while ensuring the selection of services to be integrated is driven by the needs of each scenario
4. Document the findings of these integration attempts, and abstract application domain specific optimisation in order to produce "blue-prints" of the TrustCoM profiles.

Beyond and above the specific recommendations and targets set for the continuation of research and development in each aspect of the TrustCoM Framework, meeting the above integration milestones is a major outstanding research challenge for the TrustCoM consortium.

Bibliography

1. **TrustCoM deliverable D2**:“State-of-the-art evaluation”, Emil Lupu (editor)
2. **TrustCoM deliverable D3**:“Case study scenarios”, Paul Kearney (editor)
3. **TrustCoM deliverable D6:** “Roadmap of technical standards development v1.0”, Joris Claessens (ed.)
4. **TrustCoM deliverable D7**:”Market Study”, Yücel Karabulut, Jakka Sairamesh (editor)
5. **TrustCoM deliverable D9:**”TrustCoM reference architecture, version 1”, Lutz Schulbert (editor)
6. **TrustCoM deliverable D10**:”Baseline Prototype infrastructure for CE scenario”, Dave Golby (editor)
7. **TrustCoM deliverable D11**:”Baseline prototype infrastructure for ADP scenario”, Tomas Garcia (editor)
8. **TrustCoM deliverable D14**:”Report on socio-economic models”, Jakka Sairamesh, Jonathan Sage (eds)
9. **TrustCoM deliverable D15**:”Report on legal issues”, Tobias Machler (editor)
10. **TrustCoM deliverable D16**:”TrustCoM Conceptual Framework - version 1” (draft, July 2005)
11. **TrustCoM deliverable D18**:”TrustCoM Framework Specifications - version 1” (draft, July 2005)
12. **TrustCoM deliverable D19:**“Basic set of TrustCoM methods & support tools” (draft, July 2005).
13. **TrustCoM deliverable D25:**“Exploitation plans of identified innovations” Yücel Karabulut, (ed.)
14. **TrustCoM deliverable D24**:”Roadmap of technical standards development, v2.0”, Joris Claessens (ed.)