ISTAG Report on

Grids, Distributed Systems and Software Architectures

Report

September 2004

Information Society Technologies

European Commission

*Conception and design of the cover from the Information Society DG,
Unit A3 — Information and communication unit, DTP Dpt.*

IST programme

ISTAG Working Group on

Grids, Distributed Systems and Software Architectures

Information Society
Technologies

http://www.cordis.lu/ist/istag.htm

Europe Direct is a service to help you find answers
to your questions about the European Union

Freephone number:
00 800 6 7 8 9 10 11

A great deal of additional information on the European Union is available on the Internet.
It can be accessed through the Europa server (http://europa.eu.int).

Cataloguing data can be found at the end of this publication.

Luxembourg: Office for Official Publications of the European Communities, 2004

ISBN 92-894-8162-5

Printed in Italy

PRINTED ON WHITE CHLORINE-FREE PAPER

Working group participants:

ISTAG members: Tony Hey (chair), Kiril Boyanov, Manuel Hermenegildo, Andrea Nicolai
Other participants: Guillaume Alleon, Christoph Bussler, Patrick Cousot, Michael Fehse, Burkhard Neidecker-Lutz, Tapio Tallgren, Clemens-August Thole, Ulrich Trottenberg, Rene Ziegler, Alan Gould, Santosh Shrivastava, Tony Storey, David Williams
Rapporteur: Eugene Sweeney

Contents

1 Executive Summary

1.1 Introduction and Background

Working Group 8, "Grids, distributed systems and software architectures" was set up at an ISTAG (Information Society Technology Advisory Group) plenary meeting in December 2003. Its mission was to undertake an in-depth analysis that would allow it to provide recommendations for research orientations that will enable the development and deployment of extremely large complex distributed systems supported by relevant software architectures, tools and platforms.

The Working Group sought to identify the means to address, in a coherent way, systems and software research that are currently following different but nevertheless probably converging tracks. These approaches include Grid middleware, Service Oriented Architectures (SOA), Web Services, the Semantic Web, networked embedded systems as well as robust supporting software development tools and environments. In an increasingly connected world it is no longer sufficient just to consider interconnections though Internet services. The middleware infrastructure must also take into account other components such as mobile, embedded systems and Radio Frequency Identification (RFID).

This Working Group reviewed the research activities supported by the European Commission, and also sought input from experts in academia and industry.

1.2 Technological and Economical Trends

The rapid growth of computing power, in smaller physical packages at lower costs has meant that the global computing and communications systems in the 21st century are now all-pervasive, creating a world where individuals and businesses are able to operate anywhere, anytime. We are moving towards the complete virtualisation of resources and ubiquitous access. As well as fixed network infrastructures, we are increasingly surrounded by wireless and mobile infrastructures. Other devices, such as RFID (Radio Frequency Identification) tags or embedded sensors will also become part of this all-pervasive environment with intelligent interfaces that transmit information as and when required through global persistent and intermittently connected networks.

As integration and interoperability increases between different sectors, the need for semantics also increases as a mediator between the structure and content of the different knowledge bases. There will be a need, not just for semantics to mediate the structure and content, but also for the services themselves.

The crucial component in this interconnected world will be the software, in the form of 'middleware' between the network and the applications, that will implement most functionality and will also ensure the secure and reliable integration and interoperability of mobile, fixed, personal and corporate heterogeneous resources and applications. This will require advanced technology for the development of highly functional and high-quality software, and a truly semantically enabled middleware infrastructure.

It is also clear that social, ethical and legal issues will be of increasing importance in this ubiquitously connected environment, where computing is all-pervasive. Public acceptance of these new technologies will be a pre-requisite for their wide-spread deployment.

Software technologies are already crucially important for the development of advanced products and services in modern economies and their importance will only increase with time. The relative weight of software in the value of these products and services continues to expand to the point where it becomes a strong factor for differentiation and competitiveness. Today, more than 59% of software in Europe is developed internally by enterprises. Software is thus at the heart of European productivity and competitiveness.

There is also shift in the structure of the software industry towards commoditisation of the middleware layer. As commoditisation increases, the high added-value activities will be in the customisation and application of software and the provision of services – i.e. those activities that are closest to businesses, rather than development of the middleware infrastructure. The European software industry is well-placed to benefit from this trend.

1.3 The Opportunity for Europe

Europe already has a strong ICT sector and these inevitable technological advances, coupled with globalisation of the economy should bring many opportunities for Europe. It is of paramount importance to continue to support the development of a strong, agile and high added-value European software industry.

A European software industry that can provide high value-added services and products with high functionality and reliability will be central to the new economy. These new services and products will be based on secure, robust, common open infrastructure standards, including Grid and Web Services middleware, supported by semantics. An efficient software and software service industry will make a significant contribution to the overall future productivity and competitiveness of Europe, the fundamental Lisbon objectives.

With the trend towards the commoditisation of much of the software infrastructure, the economic landscape is shifting towards a greater dependence on software and the associated skills. The value will no longer be in just deploying the middleware infrastructure but in mastering the complexities of applying this infrastructure in real applications – its integration into services and its localisation for businesses. This implies that specialised high value service providers will need to become closer to their customers.

It is vital that the European software base is supported and encouraged to develop the necessary skills and service solutions, and so build the critical mass necessary to compete globally. There is already an exponential expansion of software and a core objective must be to act now to develop the capability to harness the complexity of very large-scale distributed systems and to develop high added-value applications and services.

There is now a key window of opportunity for the European software industry during which action must be taken.

1.4 Research Priorities

The next 5-10 years will see the development of large-scale, resilient, fault-tolerant intermittently connected distributed systems (including mobile systems) – driven by the needs of individuals, business, the public sector and the research community. There will be many challenges to realising the integration and interoperability, managing the complexity and providing value-added services and products based on them. As systems grow larger and become increasingly complex it is important to address scalability issues in order to maintain performance. A strong, agile European software development industry based on secure, robust, common open infrastructure standards, including Grid and Web Services middleware, will be central to achieving this.

Semantics will play an important role in mediation between the structure and content of knowledge bases from different sectors and communities, and in the provision of services, as well as in ensuring the quality and reliability of the intervening software

Addressing trust, security, privacy and ethical issues in a globally interconnected shared resource environment is another core challenge.

Societal, business, and legal issues must also be addressed and closely integrated with implementation plans, since these will impact the eventual take-up and acceptance by industry, government and individuals.

Six priority research themes are identified:

(i) Interoperability and integration
(ii) Management of complex software-intensive systems
(iii) Semantics and knowledge
(iv) Software development methodologies, tools, and standards
(v) Trust, security and privacy
(vi) Societal and commercial acceptance take-up

1.5 Recommendations

1. There should be a clear focus on the research priorities identified for the Framework programmes, since they will have a major impact on European competitiveness and productivity.

   (i) Interoperability and integration
   (ii) Management of complex software-intensive systems
   (iii) Semantics and knowledge
   (iv) Software development methodologies, tools, and standards
   (v) Trust, security and privacy
   (vi) Societal and commercial acceptance take-up

2. Active support should be given to the development of open standards, including Grids and Web Services. This should include support for the development of open source implementations of these emerging international standards.

3. An important contributor in the new service economy will be smaller SME's providing specialised services. New mechanisms should be explored to encourage the involvement and creation of SME's offering specialised services. Such SME's could play an important role for Europe's future competitiveness

4. Societal, ethical and legal issues must be addressed in parallel with the technical agenda to encourage take-up by industry and acceptance by the general public.

5. An important new trend identified along with the move towards Web Services is the 'commoditisation' of much of the middleware infrastructure. This implies that specialised high value service providers will need to be close to their customers and may counteract the trend towards outsourcing of IT development. Since this is now a key window of opportunity for the European software industry, the priorities identified in the WG Report should be reflected not only in FP7 but also in the remaining calls of FP6

6. ISTAG requests a response from the Commission on how Strategic Objectives could be adapted in calls 4 and 5 of Framework Programme 6 to reflect the research priorities identified in this document.

7. The Commission should encourage and support industrial interest in the creation of a Technology Platform in this area.

2 Objectives and Methodology of the Working Group

Working Group 8, *"Grids, distributed systems and software architectures"* was set up at an ISTAG plenary meeting in December 2003. Its mission was to undertake an in-depth analysis that would allow it to provide recommendations for research orientations that will enable the development and deployment of extremely large complex distributed systems supported by relevant software architectures, tools and platforms. These middleware developments will underpin future Ambient Intelligent environments.

The realization of the developing AmI vision and its intersection with other industry visions will require the development of extremely large, complex, heterogeneous, distributed systems. These must be built on flexible middleware platforms capable of utilizing high-performance networking services so as to support delivery of higher-level layers of value-added, functional services to individuals, to industry, and to governments. These heterogeneous middleware platforms will include domain-specific platforms, such as those presently under development in e-Business, in manufacturing, in transport and in the home. They will encompass Grid middleware that is now moving beyond 'computing on demand' and beyond the academic research community, and beginning to develop software infrastructures capable of providing flexible, secure, coordinated resource sharing among dynamic collections of individuals, institutions and resources. The resulting complex systems, which will comprise a myriad of interacting embedded software components, will begin to approach a broader vision of AmI interacting with other industry systems. Such systems will become intelligent, self-configuring, self-healing, self-protective, and self-managed and these requirements will be reflected in the critical middleware infrastructure.

The Working Group has sought to identify the means to address, in a coherent way, systems and software research that are currently following different but nevertheless probably converging tracks. These approaches include Grid middleware, Service Oriented Architectures, Web Services, the Semantic Web, networked embedded systems as well as robust supporting software development tools and environments. As we move to an increasingly connected world, the middleware infrastructure must also now take into account other components such as mobile, embedded systems and RFID.

This Working Group began its analysis by reviewing the research activities supported by the European Commission, and sought input from academia and industry.

It considered the need to have an interconnected vision for ICT futures encompassing not only the personal space, but also encompassing the business community and the public sector, and their applications. The WG attempted to identify the key opportunities for Europe that result from the technological advances and the move to a global knowledge-based economy, and the resulting research and socio-economic challenges.

3 Technology and Economic Trends

3.1 Technology Trends

Technological innovation has always driven economic progress – most recently seen by the impact of semiconductors, the Internet and mobile technologies.

In 1965, Gordon Moore observed an exponential growth in the number of transistors per integrated circuit and predicted that this trend would continue. This became known as "Moore's Law". Since then, through the inevitable advance of technology, the doubling of transistors every couple of years has been maintained and still holds true today. Intel expects that it will continue at least through the end of this decade.

This rapid increase in computing power, in increasingly smaller physical packages at lower costs has meant that the global computing and communications systems in the 21st century are now all-pervasive, creating a world where individuals and businesses are able to operate anywhere, anytime. We are moving towards the complete virtualisation of resources and ubiquitous access. Ideas, products and services can be transferred virtually instantaneously wherever and whenever needed. As well as fixed network infrastructures, we are increasingly surrounded by wireless and mobile infrastructures providing connectivity through, for example, PDA's or mobile phones. Other devices, such as RFID tags or embedded sensors with intelligent interfaces can now transmit information as and when required through these globally connected networks. For example, the software systems used in the aeronautic and automotive industries are already diversified and becoming increasingly interconnected and interdependent - ranging from business software, design, engineering, documentation management, embedded systems and services.

In this pervasive world ubiquitous access is fundamental. The corporate world, once locked within its own well-controlled space, is now being forced through business necessity and customer demand to move towards the pervasive region, where activities are location independent – any time and any place. We are going from a "fixed" world to a "moving" world – where there is no centralised control, and access is ubiquitous. Up to now, these corporate and mobile pervasive worlds are largely separate and do not share information: increasingly, however, these two worlds need to be connected though a middleware layer.

Computing and communications technologies, and the software that drives them, are becoming an intrinsic part of everyday life, be it as an individual, a citizen or an enterprise. Like electricity, modern society could not function in their absence. This trend towards increasingly higher computing capacity (in terms of processing power, storage, and communication bandwidth), brought about by Moore's law, has not been followed by an equivalent improvement in our ability to develop the highly complex software to exploit the available hardware platforms. In fact, despite the obvious progress in programming languages, compilers, and software development environments, significant additional advances are still required in these areas in order to make it feasible, with a reasonable amount of effort and in a reasonable amount of time, to develop the large, reliable, high-quality, interoperable applications that these hardware advances make possible.

Software developments are now moving away from the more tightly linked object-oriented approach towards a more loosely-coupled service oriented approach. The major players in the IT industry are beginning to move towards Web Services based on a service-oriented architecture as the most appropriate technology for large-scale distributed computing. It is an encouraging sign that the major (US) players have agreed on a strategy to develop open, royalty free, interfaces for interoperable web services. This is a new and significant development, which provides an opportunity for the European software industry to build on.

For example, in 2002, the Global Grid Forum agreed to base the new Grid infrastructure on a Service Oriented Architecture similar to Web Services. Working groups at the Global Grid Forum are currently debating the technical specifications of the core services for the proposed Grid middleware architecture, the Open Grid Services Architecture (OGSA). In January 2004 a new Grid middleware specification was proposed by a consortium of major players in the industry, which introduced a new set of Web Services in WS-Resource Framework (WSRF). In this approach, Grid services are just a specific collection of Web Services. This Grid middleware layer will be an important part of the emerging pervasive and corporate middleware infrastructure.

As integration and interoperability of resources between different sectors or communities grow, the need for semantics increases as a mediator between the structure and content of the different knowledge bases. There will be a need, not just for semantics to mediate the structure and content, but also for the services themselves. Whilst semantics helps to provide the interoperability, this only works if the different parties (or software) have the same understanding. This is only possible if the semantics are described explicitly. Although interoperability may be achieved without explicit semantics, this would be less efficient. The use of explicit semantics allows interoperability to become automatic, hence increasing both speed and agility.

These developments, actively supported by industry, will lead to a stable, open middleware infrastructure.

The crucial component in this interconnected world will be the software infrastructure that will implement most functionality and will also ensure the secure and reliable integration and interoperability of mobile, fixed, personal and corporate heterogeneous resources and applications. This will require advanced technology for the development of highly functional and high-quality software, and a truly semantically enabled middleware infrastructure.

3.2 Economic Trends

The European Innovation Scoreboard (EIS), developed at the request of the Lisbon European Council in 2000, focuses on high-tech innovation and provides indicators for tracking the EU's progress towards the Lisbon goal of becoming the most competitive and dynamic knowledge-based economy in the world within the next decade. It contains 17 main indicators, selected to summarize the main drivers and outputs of innovations. In the most recent figures (EIS2003) the indicator-by-indicator analysis shows that large gaps between the US and Japan continue. The EU leads the US for only one of the twelve indicators for which US data are available (S&E graduates). The gap with Japan is comparable to that with the US, and the EU is lagging in all ten indicators that are available for Japan.

The EIS 2003 report does, however, note that

> "*..the only encouraging example of a long lasting catching-up process is in ICT expenditures (gap cut by 50% since 1996). Reaping the full benefits of this positive trend would require acceleration of organisational innovation following investment in ICT hardware.*"

A strong ICT sector is therefore a vital component of EU competitiveness, and Europe has a strong platform from which to build and innovate.

A recent OECD report ("The Sources of Economic Growth in the OECD Countries", 2003) also highlights the important contribution that ICT makes to economic growth and productivity. It comments that

> *"ICT-related technologies continue to have the potential to affect the quality and variety of goods and services, as well as the costs of transactions between many economic agents. In addition, the use of ICT may be increasing the efficiency of innovation, further contributing to long-term growth potential."*

The report notes that there have been significant changes in the OECD economies, brought about by the rapid spread of ICT that is changing work organisation practices, production processes and the relationships between consumers and producers.

In addition to the positive effects of ICT on output and labour productivity, it also generates "spillover" effects in the economy. For example, the economic benefits of improved business-to-business communication via the Internet do not all arise directly from quality improvements in the stock of individual computers, but also from different – and cheaper – ways of organising business processes.

Businesses are increasingly taking greater advantage of better real-time information systems, rationalising inventory stocks and the distribution of their products. Businesses have also started to reduce costs by integrating their suppliers more closely in the design and manufacturing of products, while also using the web to outsource tasks previously carried out internally. With greater information exchange between customers and producers, companies are better able to control fluctuating demand.

Today it is now possible to source raw material from anywhere in the world – wherever it is most cost effective to do so, to transport it to wherever it is best to manufacture, to market and sell products globally and to deliver products directly to the end-user.

Trade liberalisation and the fall in communication and transport costs means it is becoming more difficult for Europe to compete with low wage and well-educated labour forces. Indeed, the EU is overall behind the US and Japan on productivity, but its labour costs are higher. For example, labour costs in Korea are just over half UK levels, and the proportion of graduates in the working age population is almost identical (*Source: US Bureau of Labor Statistics)*.

Businesses are now able to be more flexible over the location of production facilities and resources and to organise their assets and activities on a global basis to reap efficiency gains.

This trend is evidenced today in the software industry by the growing volume of off-shore outsourced software production. However, much outsourcing is focused on localisation activities (e.g. language, local regulations). There is also shift in the structure of the software industry towards commoditisation of software – particularly in the areas of packaged software and core components. Today, as software becomes increasingly pervasive, more and more software value is now contained in non-software companies. As commoditisation of the middleware layer grows, so too does the importance and relevance of open source software. In this environment open source is a market leveller, particularly when supported by industrial players who pool resources for a common interest – i.e. "coopertition". An example of this would be the support now received by Linux from most of the major IT vendors. A further discussion of the issues related to open source software can be found in the ITEA (Information Technology for European Advancement) Report on Open Source Software" published in January 2004. It can be obtained from the ITEA web site (www.itea-office.org).

In this economic landscape, there will be an increase in outsourcing of the commoditisation and localisation processes, but the high added-value will be in the provision of services, the customisation and application of software – i.e. those activities that are closest to businesses, rather than the software infrastructure per se. The European software industry is well placed to benefit from this trend. There is now an opportunity for it to provide cost-effective solutions and concentrate more on high value-added products and services.

In this context it is thus essential to concentrate on developing higher quality software, in less time. This requires much better tools and techniques, as well as better-educated developers, thus raising productivity and competitiveness.

We see the development of an interconnected world with long and complex supply and value chains, which highlight the ever-growing pervasiveness of computing as a general-purpose technology. Even if some enterprises do not yet use these technologies in critical business applications, their suppliers, customers and competitors increasingly do so.

This networked environment will transform businesses and business processes. Agility, flexibility and location independence will be key business drivers. Value-chain relationships will change significantly and will affect every aspect of the work processes. It will be the owners of the higher-level skills and knowledge who will reap the benefits in this new economy, not the owners of the raw materials or production resources.

Social, ethical and legal issues will have an ever increasing importance in this ubiquitously connected environment, where computing is all-pervasive. For example, as information is collected, stored and re-used, social and ethical issues relating to privacy, civil rights, security and access rights must be addressed. As resources are spread across national boundaries, and across national legal systems, workable access and usage rights must be established.

These inevitable technological advances, coupled with the globalisation of the economy, bring not only challenges but also opportunities for Europe to be a major added-value player in the emerging knowledge-based society particularly in the ICT sector.

Software technologies are crucially important for the development of advanced products and services in modern economies. The relative importance of software in the value of these products and services continues to expand to the point where it becomes a strong factor for differentiation and competitiveness. Software is already and will increasingly be the key component in this new economy. Software development is thus at the heart of European productivity and competitiveness, the fundamental Lisbon objectives.

4 Industry Visions

The following scenarios represent views from different industry sectors looking ahead 5 to 10 years at how advances in technology – hardware and software – will be routinely deployed in their particular sector. It is not intended to be a list of research challenges but rather a vision of what the exploitation of imminent technology will mean to the routine capabilities of these business sectors. The visions are from informed industry experts, and illustrate the inevitability of technological progress, the move to pervasiveness, and the opportunities this brings not only for the software industry but within their industry sectors too.

Following the visions, some commonalities and common challenges are identified.

Product Design in the Automotive Industry

Product development in the automotive industry is today already a very distributed process. It involves different groups within the same company as well as external suppliers and consultants, located around the globe. Different teams collaborate closely as product development time decreases as a result of reducing the number of physical prototypes, parallelisation of the development work and the use of sophisticated simulation tools.

Instead of merely sharing compute resources, the sharing of data is the most demanding challenge in this application domain. The developments by the different teams and the numerous variations of prototypes are stored in or referenced by databases. Different development groups traditionally use different problem solving environments, each of which is tailored according to the specific development task. The different databases may express similar meaning in structure and contents in different terms. Therefore semantics and tools like ontologies and related transformers must be available for mediation between the different databases.

These future integrated environments, which will include robust and secure management tools for access to remote data repositories, will create new ways of working in automotive design.

Weather forecast on request

We can envision a system consisting of many components - computing, communication and sensors - which can be connected in an ad-hoc way, as and when required to provide information about, for example, temperature, air-pressure, and humidity. Mobile phones, cars, home or professional weather stations, weather satellites can all be intermittently connected.

A weather information service might process incoming data from networked sensors and weather satellites. Based on the input data weather simulation models will be run to provide fine granular weather forecast for specific regions. We can envision a situation where a consumer requests a weather forecast via mobile phone (or some other device). The weather system / service will localize the customer and identify which device is used. The local weather information will then be automatically adapted and sent to the specific user device (Mobile GSM, SMS, PDA, Notebook etc.). This weather data may also then be used as input to a business process. For example, a co-generation plant might require a local weather forecast to plan the energy consumption of the next day.

Field Service and Support

This scenario describes a system to automatically find and identify members of a field service and support team, who can respond to a problem alert call. A typical problem report could be: “There is a Problem with the Network Colour Printer in HQ, Darmstadt Pfnorstr.” The software system/service must then identify:

- Where is HQ Pfnorstrasse
- Which staff members of field service belong to the area
- Who is in this area right now
- Who is available right now, can deal with printers
- What devices/systems can be used (Mobile phone, PDA, Notebook, Fax etc.)

and to make sure that confidential corporate information gets to the right, authorized technician only.

In this scenario there are many semantic issues which are needed in order to manage the interfaces, for example, different definitions of “HQ” or “printer”. Security will also be vital to ensure sensitive information is properly managed.

This particular business service is currently managed manually in an ad-hoc fashion. As an intelligent middleware infrastructure layer emerges to manage the interfaces and integration of resources, then such business services will become routine.

Health Care

Within the Healthcare sector there are many applications that would benefit from secure, pervasive, ubiquitous and transparent access to information and computing power. However, there are also many challenges to overcome.

A knowledge-based application might require the secure, pervasive, fault tolerant, ubiquitous access to the distributed medical information for data-mining and knowledge extraction from population-level data. Population-level data might already exist in distributed data repositories, or be collected in an ad-hoc way from local (body-area) sensors and be transmitted for processing.

In order to realise such applications there are many important ethical, privacy and security considerations that must be addressed in addition to the issues of standardization of data, federation of databases and content-based knowledge extraction.

Safety-critical software solutions in the Aerospace industry

Software used in the aerospace industry is already quite diversified, encompassing business software, design, engineering, documentation management, embedded systems and service solutions. However, in the aeronautics industry - as well as elsewhere - there is an exponential expansion of the software used within the product itself. In the aerospace industry the nature of the product makes this software safety-critical. Recent catastrophes in various different industries due to software failures serve as graphic reminders that this is a priority that must be addressed as systems become larger, more distributed, and more code is re-used. These "intra-product" distributed software systems are also increasingly connected to external large-scale distributed systems.

A consequence of the growing software code size will be an explosion of the cost to maintain these high-quality objectives. The vision is, therefore, of the availability of the necessary knowledge, methods, technologies and tools to allow software engineering to emerge from its present strong dependence on manual activities towards a more automated and better-tooled environment. The current simulation, audit, review and test technologies, all of which will remain indispensable for the near future, are rapidly reaching their limits and must therefore be complemented by new design, checking, verification and certification tools to achieve the required levels of trust and confidence in specifications and software.

Integrated Product Lifecycle and Workforce Safety Management

Today there are some isolated examples of the management of high-value assets. For example, aircraft engines are managed electronically from design through to production and thereafter through the entire operational lifecycle. Engines "phone home" to have a maintenance crew ready when the plane lands. As this trend is projected into the future, we can expect increasingly more devices and materials to be electronically tracked and managed.

Large enterprise customers are today piloting early experiments to manage and track the entire lifecycle of materials and equipment from suppliers, including their interaction during operation. An advanced example is from oil drilling rigs, where the supply and quality of drilling gear and other equipment from hundreds of different suppliers needs to be managed. Beyond merely tracking the items passively, in these scenarios hazardous materials are tracked with active sensors automatically communicating with personnel trackers to prevent work accidents.

Building and maintaining such a diverse infrastructure poses unique challenges in longevity (non-IT equipment has decades of lifetime), code quality (in safety critical situations) and the number of participating suppliers (in the tens of thousands for aircrafts or oil rigs).

Adaptive supply chains, consumer protection

While supply chains are partially automated today, this trend will dramatically increase in the future. Production, transport and actual retail operations will become increasingly automated. The sheer number of items flowing through consumer supply chains will pose enormous scalability challenges (even today, a single large retailer typically tracks several billion items per year). While adaptive forecasting and replenishment allow suppliers more transparency and enables smooth planning, this will need to be balanced with consumer privacy concerns and government regulations.

There will be a requirement to design identification systems with provable properties of anonymity and privacy while still fulfilling the legitimate needs of enterprises or governments that may wish to override privacy, for example in cases of danger (i.e. in a situations where dangerously defective goods were shipped and criminal activities need tracing).

Knowledge Management for the Individual Citizen

Individuals will have access to information from many different sources – both from the public and private sector. They might have information in different data sources such as hospitals or different businesses from which they purchased products. They will require access to product information to support their purchasing process, or to public sector information when performing their duties as a citizen. The information necessary to accomplish these tasks must be readily accessible and understandable. In general, the information will come from many data sources and therefore requires semantic integration in order to make it useful and manageable.

In order to make the access semantically consistent, a "Semantic Web API" allowing efficient and large-scale access to content, ontologies and annotations will be paramount in order to establish ubiquitous information access. Key aspects that must be supported include mediation, reasoning under uncertainty and versioning. Only these key aspects allow making different data and knowledge sources semantically uniform in their access. In addition, data sources must be both accessible and searchable allowing the user to find the most relevant ones. Once a data source is found it must be searchable, maybe in conjunction with other knowledge sources so that the combination of knowledge is available.

Commonalities in the Industry Visions

These industry visions represent informed views of how technology might be routinely employed in those sectors in 5-10 years. They result from the possible future technological responses to business drivers and user needs.

Even though the applications and services are very different, some commonalities are evident. These are the key technological and societal challenges that must be addressed if these visions are to be achieved.

The realisation of these visions relies on the availability of:

- **Interoperable services built on robust middleware based on open standards to manage the diverse range of technologies and resources.**
- **Semantics to manage the sharing, integration and distribution of content from heterogeneous sources and different communities.**
- **Software linking the mobile and fixed corporate worlds.**
- **Systems for the management of privacy, security, access and use.**
- **Resolution of ethical issues.**
- **Effective management systems for the business models (e.g. charging and accounting).**
- **Software knowledge, methods, technologies and tools to allow software engineering to emerge from its present strong dependence on manual activities towards a more automated and better-tooled environment.**
- **New design, checking, verification and certification approaches to achieve the required levels of trust and confidence in specifications and software.**
- **An emphasis on highly efficient production of reliable, distributed software.**

5 The Opportunity for Europe

In this increasingly interconnected world (both at a personal and enterprise level), software is the key enabler. It is the software that will deliver the functionality, integration and applications for these very large scale, distributed and connected (sometimes intermittently) services which access distributed, remote resources. There is already an exponential expansion of software, and a core objective must be to master the complexity resulting from this expansion. Europe must develop the capability to harness the complexity of very large-scale distributed systems to develop new high added-value applications and services.

Europe already has well-established strengths in critical sectors such as mobile, embedded systems and business process software. As we move from a fixed to a mobile world Europe is well positioned to take advantage of the new opportunities that are presented. Europe also has strong industrial sectors, such as aerospace, automotive and pharmaceutical, which are both developers and users of software. These industries are well aware of the potential benefits to be gained from the development of an intelligent, open standard, interoperable infrastructure.

Open standards, such as GSM, are a well-recognised asset of European ICT industry. The IT industry's move to Web Services, and in particular the move to open standards for interoperability provides Europe with a unique opportunity. There is also an associated trend towards the 'commoditisation' of much of the software infrastructure. As the economic landscape shifts towards a greater dependence on software and the associated skills, the value will no longer be in the middleware infrastructure, but in mastering the complexities of exploiting it in real applications – its integration into services and the localisation for businesses. This implies that specialised high value service providers will need to be close to their customers. This is now a key window of opportunity for the European software industry.

Europe has the opportunity to build not only products and services but also capabilities and skills. The benefits will not just be seen in the IT industry, but in all sectors as IT becomes an integral part of the business process. The ubiquitous and pervasive environment will support and benefit enterprises in all sectors, governments, public service providers as well as individuals. In particular, SME's and smaller "niche" players can benefit from the new business model opportunities provided by the improved access to an enlarged market via the "level playing" field provided by the emerging open standard middleware infrastructure. Europe can take advantage of the opportunities these new developments are creating but there is a limited window of opportunity.

The creation and exploitation of the core enabling "knowledge" and IT skills will also help to ensure that Europe maintains a leading position as both an ICT developer and user. The European research community is strong in the areas of logic and semantics, and these strengths combined with the Web Services infrastructure could provide Europe with key differentiators. The applications of these technologies would be not only in semantic web enabled web services but also in the development of tools to analyse, develop and validate complex safety critical software. Thus there are significant opportunities for Europe to play a leading role in the development and exploitation of a pervasive, intelligent software infrastructure based on open standards. Attention should also be paid to European XML-based standards for All-IP (all services IP compatible) and the development of Wireless All-IP standards supporting multi-channel and vertical roaming.

A strong, agile European software development industry providing high value-added services and products of high functionality and reliability, based on secure, robust, common open infrastructure standards and supported by semantics will be central to the new economy. An efficient software and software service industry will make a significant contribution to the overall future productivity of Europe.

It is vital that the European software development base is encouraged to develop the necessary skills and service solutions, and so build the critical mass necessary to compete globally. This is the challenge that Europe must address to take advantage of the opportunity that presents itself.

6 Research Priorities

The next 5-10 years will see the development of large-scale resilient and fault-tolerant, intermittently connected and mobile distributed systems, driven by the needs of individuals, business and the research community. There will be many challenges to enabling interoperability, managing the complexity and providing value-added services based on them. As systems grow larger and become increasingly complex it is important to address scalability issues in order to maintain performance.

A strong, agile European software development industry providing reliable and high value-added services and products based on this secure, robust, common open standard middleware infrastructure will therefore be central to achieving this.

Semantics will play an important role in mediation between the structure and content of knowledge bases from different sectors and communities. Addressing trust, security and privacy issues in a globally interconnected shared resource environment is another core challenge.

Societal, business, and legal issues must also be addressed and closely integrated with implementation plans, since these will impact the eventual take-up and acceptance by industry, government and individuals.

Six priority research areas are identified:

(i) Interoperability and Integration

As large scale distributed systems grow and become increasingly connected – from thousands to perhaps billions of interconnected systems - we will have access to increasing amounts of information, sometimes intermittently. We are in danger of not taking advantage of the opportunities offered by this imminent "data deluge".

As these heterogeneous resources and services, supported by semantics, come together to deliver complete systems, so too must there be a convergence and interoperability of standards and protocols.

Grid and XML-transactional middleware needs to be developed to manage the interoperability, access and integration of distributed data and resources, and so provide the platform for added value services. European XML-based standards for business and public service processes are urgently needed.

This middleware needs to be both context and location sensitive.

Enterprise interoperability and interoperability with legacy systems must also be addressed.

The development of open standards for grids and web services as well as reference implementations of these standards should be actively supported.

(ii) Management of Complex Software Intensive systems

With the increasing heterogeneity of resources and services the management of systems is important. This creates a new requirement for “autonomic” software management tools. Tools for dynamic system management, for example system performance prediction and for composability of services must be developed as the middleware becomes more autonomic, adaptive and self-healing.

New software systems must be developed to manage this complex environment, which will see an increasing convergence of wire, wireless and mobile. Applications will need to embrace Web Services and Service Oriented Architectures, Grid systems, Wireless, Sensors, RFID, and embedded systems. An active dialogue among the software developers in the different sectors will be necessary. Wireless IP roaming and multi-channel technologies are complex and standards need to be developed in this area.

Specific issues relating to ad hoc networks must be addressed. Resource or service awareness, discovery and identification are important components. Users must be able to automatically find the most appropriate and cost-effective resources and services they require.

Software must be developed to manage service and data updates, to provide tracking tools and to address life-cycle issues and performance improvement.

Technical solutions must be found for managing resource usage and the associated accounting.

Software for the management of security and delegation must be developed. Other issues relating to trust, security, identification and authentication are discussed below.

(iii) Semantics and Knowledge

With the heterogeneity of the resource structure and content in different sectors, a semantic approach will be necessary to deliver effective integrated, interoperable services and solutions. The use of semantics is not confined to just the application sectors. Semantics of services will also be necessary to compose, federate and create complex services. Whenever composition takes place it has to be ‘lossless’ in the sense that any data communicated between two services does not alter its semantics.

Ontologies will need to be developed in relevant applications areas to provide interoperability between different sectors. It will also be necessary to develop strategies, services and tools to mediate between the different resources.

Data from these integrated resources will form an enormous distributed base of information. Mining and extracting knowledge from these distributed repositories will require the development of new algorithms, strategies, tools and services.

(iv) Software Development Methodologies, Tools, and Standards

Extremely large, complex, distributed systems and services involve a myriad of interacting software components. These must conform to open standards and be of an assured quality and be certified in safety critical situations. To achieve the industry visions described above the software will need to become intelligent, self-configuring, self-healing, self-protective, and self-managed. Novel software development methodologies, tools and standards for such heterogeneous distributed systems must therefore be developed. Such methodologies must support rapid cycles of (re-)design, prototyping and evaluation.

The reliability and dependability of such systems must also be addressed. It will be necessary to develop more cost-effective systems to generate trust and confidence in the requirements, specification, design, development, documentation, configuration, modification and quality assurance. Recent catastrophes due to software failures serve as graphic reminders that this is a priority which must be urgently addressed as systems become larger, more distributed, and as more code is re-used. Failures in reliability and dependability will also impact user confidence in the Information Society generally. Quality and robustness of software is therefore critical. Software development environments and tools addressing these aspects must therefore be developed for complex and distributed systems. As software systems grow, the potential for verification tools will increase in importance – particularly as systems are increasingly based on software re-use. Other software re-use challenges must also be addressed, such as certification and "fitness for use". The opportunities for Europe presented by free/open source approaches should also be supported and exploited.

Practical challenges outside the core software development activity must also be addressed as part of the overall process of product or service development. These include requirements capture, managing changing requirements, on-going refinements, life cycle management, self-configuration and affordability. Systems must also be scalable with predictable performance.

(v) Trust, Security and Privacy

Addressing trust, security and privacy issues in a globally interconnected shared resource environment is a core challenge. As enterprise, government and individual activities become increasingly digitised and connected, more personal, public and commercial information will be gathered, stored and possibly disclosed to third parties.

As resources are shared, traded or licensed, transaction authorisation/verification and identity verification over secure (and in many cases non-secure) communication channels must be addressed. Users will need to be assured that their transactions, analyses or resources are not being tampered with, that the results are "trustworthy", and that their identities cannot be stolen.

Robust technical solutions must be developed related to trust and security in this trans-enterprise, trans-national, trans-cultural environment. These solutions must address privacy, anonymity (with need to know), security, confidentiality, identity verification and theft prevention, transaction authorisation and repudiation.

(vi) Societal and commercial acceptance and take-up

Whilst the move to complex, large-scale distributed systems brings with it opportunities to develop new business models, it also raises other non-technical issues. In order to fully engage industry and to clarify commercial viabilities, these issues must be addressed alongside the technical.

These non-technical issues include:

- **The transfer of research** results into business value should be supported through near market support actions and "live" evaluations.

- **Socio-economic factors** such as the development of new skills and complementary services.

- **Societal impacts** such as ecological, entertainment and leisure.

- **Ethical issues** related to privacy, anonymity, identity verification/theft and civil liberties,

- **Legal issues** related to the management of trans-enterprise, trans-border access and usage rights, for example: -

    - Outsourcing of data management (different legal environments and the management of data held in 3rd countries)
    - EU database legislation
    - Security/confidentiality
    - Trans-national legal frameworks
    - Software patent issues
    - IP Harmonisation
    - Legislation for the use of "common" public domain components (e.g. open source, licence free, etc)

- **Education issues** must be addressed if Europe is to have an appropriately skilled workforce to take advantage of these new opportunities.

- **Business Issues.** New business models will need to be developed and validated to encourage and promote take-up - for example, in supply chain management or for application developers and service providers. Whilst the move towards an open standards software infrastructure will assist in generating a more level playing field for Europe in that it is likely to diminish the importance of the behemoth companies, the commercial implications for all of the software industry remain unclear.

  A complex, distributed, trans-national, inter- (or even intra-) enterprise system raises many licensing issues. Workable models for access and usage need to be investigated.

7 Recommendations

1. There should be a clear focus on the research priorities identified for the Framework programmes, since they will have a major impact on European competitiveness and productivity.

   (i) Interoperability and integration
   (ii) Management of complex software-intensive systems
   (iii) Semantics and knowledge
   (iv) Software development methodologies, tools, and standards
   (v) Trust, security and privacy
   (vi) Societal and commercial acceptance take-up

2. Active support should be given to the development of open standards, including Grids and Web Services. This should include support for the development of open source implementations of these emerging international standards.

3. An important contributor in the new service economy will be smaller SME's providing specialised services. New mechanisms should be explored to encourage the involvement and creation of SME's offering specialised services. Such SME's could play an important role for Europe's future competitiveness

4. Societal, ethical and legal issues must be addressed in parallel with the technical agenda to encourage take-up by industry and acceptance by the general public.

5. An important new trend identified along with the move towards Web Services is the 'commoditisation' of much of the middleware infrastructure. This implies that specialised high value service providers will need to be close to their customers and may counteract the trend towards outsourcing of IT development. Since this is now a key window of opportunity for the European software industry the priorities identified in the WG Report should be reflected not only in FP7 but also in the remaining calls of FP6

6. ISTAG requests a response from the Commission on how Strategic Objectives could be adapted in calls 4 and 5 of Framework Programme 6 to reflect the research priorities identified in this document.

7. The Commission should encourage and support industrial interest in the creation of a Technology Platform in this area

Annex 1 – Presentations to the Working Group

European Commission Presentations

These presentations may be downloaded from: -

http://www.cordis.lu/ist/istag.htm#istag2004

Presentation Title	Presenter	Organisation
Software Technologies	Jacques Bus	DG INFSO D3 Software Technologies
Semantic Web/Knowledge Technologies	Roberto Cencioni	DG INFSO E2 Knowledge Management & Content Creation
Grid Technologies	Wolfgang Boch	DG INFSO F2 Grids for Complex Problem Solving
eBusiness and Virtual Organisations	Joël Bacquet	DG INFSO D5 eBusiness
Grids and Research Infrastructures	Mário Campolargo	DG INFSO F3 Research Infrastructure
Embedded Systems	Mercé Griera i Fisa,	DG INFSO C3 Embedded Systems
Health Grid	Sofie Norager	DG INFSO C4

Industry and Academic Presentations

Presentation Title	Presenter	Organisation
European Networking	David Williams	CERN
An Application Vendor's Perspective	Burkhard Neidecker-Lutz	SAP AG
The Engineer's Perspective	Alan Gould	BAE Systems
PharmaGrids	René Ziegler	Novartis Pharma AG
An independent Telco service/systems developer's perspective	Michael Fehse	T-Systems International
Grid computing in the mobile domain	Tapio Tallgren	Nokia Research Centre
Semantically Enabled Web Services	Christoph Bussler	DERI, National University of Ireland
Middleware	Santosh Shrivastava	Newcastle University
Web Services	Tony Storey	IBM

European Commission

ISTAG Report on Grids, Distributed Systems and Software Architectures

Luxembourg: Office for Official Publications of the European Communities

2004 — 26 pp. — 21 x 29.7 cm

ISBN 92-894-8162-5

SALES AND SUBSCRIPTIONS

Publications for sale produced by the Office for Official Publications of the European Communities are available from our sales agents throughout the world.

How do I set about obtaining a publication?
Once you have obtained the list of sales agents, contact the sales agent of your choice and place your order.

How do I obtain the list of sales agents?

- Go to the Publications Office website (http://publications.eu.int/) or
- Apply for a paper copy by fax (352) 29 29-42758

15
KK-61-04-937-EN-C

Publications Office
Publications.eu.int

ISBN 92-894-8162-5
9 789289 481625