



Project no. FP6-033636

## NESSI-Grid

Networked European Software and Services Initiative - Grid

Instrument: Specific Support Action (SSA)  
FP6 – Priority IST-2.5.4 Advanced grid technologies, systems and services

### Deliverables D.1.1 & D.1.3 Grid Vision and Strategic Research Agenda

Due date of deliverable: 31 October 2006  
Actual submission date: 17.11.2006

Start date of project: 1st May 2006

Duration: 30 months

Lead contractor for this deliverable: SAP AG  
Revision: V.1.0 - 17.11.2006

Project co-funded by the European Commission within the Sixth Framework Programme (2002-2006)	
Dissemination Level	
PU Public	Yes

Keywords:

**Business Grids, SOKU, Agility, Adaptive IT, Virtualization, SLA enforcement**

**Table of content**

**Executive Summary ..... 3**

**1 Introduction ..... 5**

1.1 Vision ..... 5

1.2 Technical scope ..... 6

1.3 Non-goals ..... 7

1.4 Methodology ..... 8

**2 Basic Model ..... 9**

2.1 Services ..... 9

2.2 Enterprise grids ..... 10

2.3 Departmental grids ..... 11

2.4 Virtual organisations ..... 13

**3 Business Scenarios ..... 14**

3.1 Enterprise ..... 14

3.2 Hierarchical enterprise ..... 16

3.3 Hosting: a special case of a virtual organization ..... 17

3.4 Extended enterprise ..... 20

3.5 Real-time enterprise ..... 20

3.6 Mergers & acquisitions ..... 21

3.7 Business value networks ..... 22

**4 State of the art ..... 22**

4.1 Scientific Grids ..... 23

4.2 Enterprise Grids ..... 24

4.3 Virtual machine based grids ..... 25

**5 Research challenges ..... 25**

5.1 Functional & Commercial Issues ..... 25

5.2 Dependability ..... 26

5.3 Security ..... 26

5.4 Performance ..... 27

5.5 Interoperability ..... 27

5.6 Manageability ..... 27

5.7 Governance ..... 28

5.8 Flexibility ..... 28

5.9 Overarching challenges ..... 29

**6 Conclusions and Outlook ..... 29**

**7 References ..... 30**

## Executive Summary

The Networked European Software and Services Initiative (NESSI) aims to create a strategic research agenda for European research in services and their foundations. NESSI-Grid forms part of that activity by defining a vision and strategic research agenda (SRA) for grid infrastructures used in business environments and in particular in NESSI scenarios. A first version of both, vision and SRA, focussing on more short-term challenges is presented in this document.

### The Driver

The main driver behind this SRA is the fact that businesses in future service-oriented economies need to act in a more agile fashion than ever before. This will be possible only if high-level business requirements can be translated into lower level ICT requirements with a high level of automation, so that ultimately the ICT environment adapts automatically to changing business needs.

### The Vision

We envision *Business Grids as the adaptive service-oriented utility infrastructure for business applications*. They will become the general ICT backbone in future economies, thus achieving a profound economic impact. The adoption of Business Grids is expected to happen in 3 steps: first as the ICT backbone within enterprises, second as a basis for hosting scenarios and ultimately as the general ICT-infrastructure for service-oriented economies. In this way, Business Grids will eventually support the emergence of new types of application.

### Business Scenarios

The industry-driven approach of this SRA is evident in its identification of a few business scenarios and associated business requirements. Three major scenarios are considered to be most relevant for the short-term perspective: basic enterprise, hierarchical enterprise (organized via a hierarchy of departments) and hosting (a very basic form of a virtual organization). These and more long scenarios are described in Section 3.

### Research Challenges

The following 3 key challenges for Business Grids capture the major requirements derived from the short-term business scenarios:

**KC1: Harmonizing traditional grid and virtualization technologies.**

This allows building truly interoperable core infrastructures that combine the best of both worlds, in particular by providing strong support for typical business applications while inherently reflecting different organizational structures.

**KC2: Automated translation of high-level requirements and policies** of all kinds to the infrastructure level and aggregation of infrastructure behaviour to higher-level information.

This allows the business value of requirements, policies and capabilities to be directly related to properties of the infrastructure and eventually supports managing IT as a business.

***KC3: Automated SLA and policy enforcement at infrastructure level.***

This allows reliable (resilience to higher level violations) and reproducible (due to well-defined linkage to higher level business goals) enforcement while automation contributes to cost reduction and responsiveness.

A detailed description of the identified research challenges is set out in Section 5.

# 1 Introduction

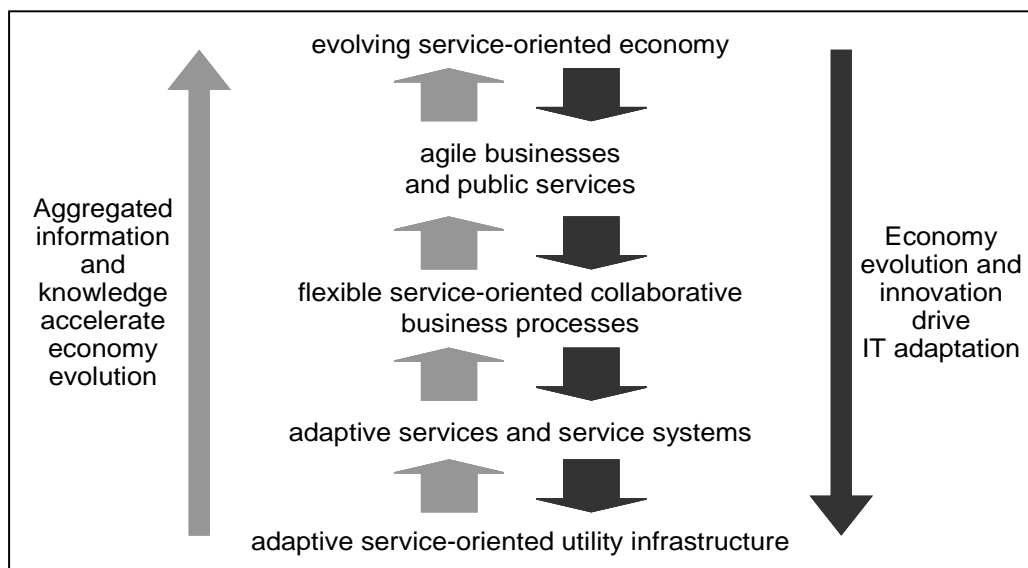
The Networked European Software and Services Initiative (NESSI) aims to create a strategic research agenda for European research in services and their foundations. NESSI-Grid forms part of that activity by defining a vision and strategic research agenda (SRA) for grid infrastructures used in business environments and in particular in NESSI scenarios.

This document presents a first version of both, a business-driven vision for grid infrastructures as well as a strategic research agenda focussing on more short-term research challenges. Rather than describing vision and SRA in two separate documents, we decided to combine both deliverables into one document as they have a significant overlap. The vision as such is explained on a high level in the first section but also refined through various business scenarios, presented in Section 3. On the other hand, the business requirements used to characterize those scenarios in more detail as well as the order of the scenarios is already a first step towards the SRA. The SRA is then refined by concrete research challenges in Section 5.

## 1.1 Vision

The main driver behind this SRA is the relationship between high-level characteristics of service-oriented economies [1] and supporting Information and Communication Technology (ICT) infrastructures. This is depicted in Figure 1.

Businesses in future service-oriented economies need to act in a more agile fashion than ever before. This will be possible only if high-level business requirements can be translated into lower level ICT requirements with a high level of automation, so that ultimately the ICT environment adapts automatically to changing business needs. Figure 1 shows the different stages of this translation process.



**Figure 1: Agile businesses drive adaptive IT**

In detail, an evolving service-oriented economy will lead to new requirements in the sense that existing businesses and services must react more quickly to changing circumstances. This means that they must adopt more flexible service-oriented business processes and be able to exploit collaborative relationships. Software technologies enabling the implementation and composition of flexible services will be important in achieving this and are now being seriously addressed by the software industry. This can be seen in the widespread trend towards service-oriented architectures and business processes built upon this architecture. At the infrastructure level, an adaptive service-oriented utility infrastructure will allow the dynamic, on-demand allocation and assembly of resources needed to support service components and assemblies. A crucial prerequisite for this layered adaptability is a proper information flow translating higher-level business goals into lower-level technical goals and aggregating lower level characteristics to higher-level capabilities.

We call this envisaged adaptive service-oriented utility infrastructure “*Business Grids*”. Business Grids will become the general ICT backbone in future economies, thus achieving a profound economic impact. In this respect they differ significantly from traditional e-Science grids. These focus on specific application areas often aiming to provide shared access to specialised high performance computing resources or datasets. Typically they deal with independently executing, stateless batch jobs which can be easily moved around a network and executed based on some given input files. In contrast to this, Business Grids are a general purpose infrastructure for arbitrary business applications, primarily aimed at providing business flexibility efficiently. Typically, business applications do not exist as independent executables but involve a complex technology stack, containing for example application servers, databases, other middleware components and other co-existing and cooperating business applications. In addition, usage patterns are often interactive or session-based.

We anticipate that Business Grids will become the major ICT infrastructure via the following scenarios:

1. Business Grids will be the ICT backbone for enterprise solutions.
2. Business Grids will support hosting scenarios for small and medium sized enterprises.

Based on these two initial scenarios and as the interworking between administrative domains becomes commonly accepted and is supported as part of the infrastructure this will evolve to the following:

3. Business Grids will be the ICT infrastructure to support service-oriented economies and eventually support the emergence of new types of applications.

It is, of course, important to note that these developments will require changes in business culture which will undoubtedly influence the timescales for change.

## **1.2 Technical scope**

Figure 2 shows the main ICT layers of existing business solutions.

Applications (e.g. collaborative business processes)
Business logic (e.g. business web services)
Middleware (e.g. application server)
Infrastructure (e.g. OS, hardware)

**Figure 2: ICT layers of business solutions.**

Broadly speaking, the infrastructure layer is about the provisioning of any kind of hardware resources (compute power, storage, network, sensors, actuators) together with associated low-level software components (e.g. operating systems). Middleware covers generic software functionality such as application servers (e.g. J2EE), databases, messaging services, portal frameworks and also serves for hiding technical complexity (distribution, heterogeneity) from application developers. The business logic layer includes specific business functionality (e.g. processing a purchase order) while the application layer combines business logic in concrete usage scenarios (e.g. a procurement system). It is important to note that these layers do not have well-defined boundaries but are presented as an illustration of different technical areas which are relevant within business solutions.

Business Grids primarily target the infrastructure layer and partially also the middleware layer. They follow the paradigm of a Service Oriented Knowledge Utility (SOKU) [2]. The overall SOKU vision goes far beyond infrastructure problems as it “identifies a flexible, powerful and cost-efficient way of building, operating and evolving IT intensive solutions for use by businesses, science and society”. Applying the SOKU paradigm to Business Grids means that Business Grids offer infrastructure resources to higher levels according to the following main principles:

- service-oriented (i.e. dynamic allocation and assembly of resources via infrastructure services)
- knowledge-assisted (i.e. translating high-level business requirements to infrastructure requirements and infrastructure capabilities vice versa)
- utility (i.e. immediately available, dependable usage, predictable).

Business Grids need to be able to participate in the execution of the business models of the applications they support in an accountable fashion, addressing issues including auditing, billing, and linkage of resource consumption to business goals. In this context, Business Grid solutions will span several ICT layers.

### **1.3 Non-goals**

Following the described scope it becomes clear that research questions related to higher layers of ICT solutions are out of scope for research in Business Grids. In particular this is the case for business models, human aspects, and business services.

- *Business Grids will not address isolated low-level grid business models.* Several trends in hardware, software, and the relative cost of IT operations conspire to make it nearly impossible to construct viable business models just

around low-level grid services in isolation from the supported applications<sup>1</sup>. The added value stems from the actual services offered on top of the grid which may include provision of resources as managed services with assured performance levels.

- *Business Grids will not address human aspects* such as political, organizational or cultural issues.

Business Grids present themselves as a utility; they are invisible to the user and even largely to the developer of applications, business logic, and middleware.

- *Business Grids will not address issues of business service engineering and service provisioning (other than provisioning of infrastructure services)*. This will be addressed in other parts of the overall NESSI-SRA.

## **1.4 Methodology**

Business Grids are about the provisioning of infrastructure resources to business services as infrastructure services. As resources are always bound to an administrative domain (there is always an entity with ultimate authority over the management of a resource), the approach of this document is based on the distinction between the following two cases: provisioning of resources within one domain and provisioning of resources between several domains.

Consequently this document is structured as follows. Section 2 provides a basic model for describing Grid scenarios with respect to administrative domains and associates them with basic business characteristics. Section 3 details some important business scenarios that can be expressed by means of the basic model. Three of the business scenarios which are assumed to have most short-term relevance are described by fairly detailed business requirements. Section 4 reflects the state of the art, both in academia and industry. In Section 5, research issues corresponding to the short-term business scenarios are derived before Section 6 ends with a brief conclusion.


---

<sup>1</sup> Compute cycle, network and storage costs will continue to decrease for the next few years. Purchased network capacity, software and higher-level service costs, including infrastructure management will remain relatively unchanged. Thus the share of pure hardware costs in overall IT operation costs becomes more and more irrelevant, inhibiting companies from basing their business models on commodity resource provisioning (argumentation taken from [3]).

## 2 Basic Model

### 2.1 Services

At the core of the NESSI vision are services. Consider the diagram below:



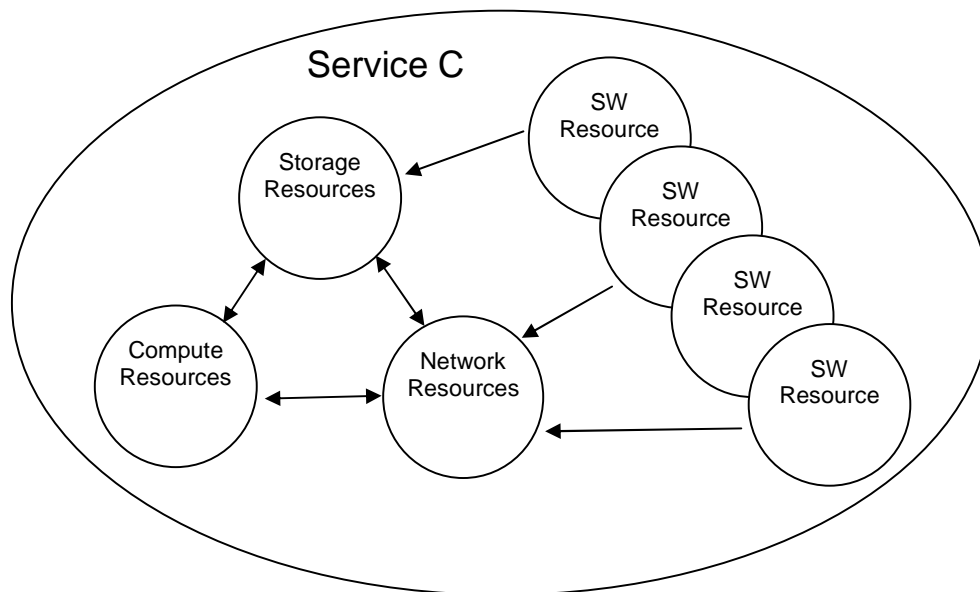
A service provides a given capability. It may be “simple”; e.g. service A, in that it provides all of the required functionality contained within itself. Alternatively, a service may be “complex”; e.g. service B, in that it may itself be made up from other services combined in some fashion. Note that the diagrams above are not intended to show the “ownership” of any given entity. It is perfectly possible that Services X, Y, Z and B are sourced from different organisations.

Services are provided and consumed. This is the case whether they are “simple” or “complex”. Relationships between organisations can be expressed in terms of this provision or consumption of services.

- A service may be provided to one organisation for its own consumption.
- A service may be provided to an organisation which consumes it to in turn provide another (complex) service.
  - Services thus provided may be provided “as-is”
  - Services thus provided may be augmented or extended in some (value-added) fashion

Services may have explicit or implicit relationships with one another. These relationships may in turn be expressed or maintained by other services (e.g. by a brokering service, a lookup service, an identification service etc.). Relationships may be static or dynamic. These service relationships allow common business models to be supported.

A service (whether simple or complex) implicitly uses resources to deliver its capability. An example of this is shown in the diagram below:



Service C uses a number of underlying hardware (e.g. compute, storage, networking etc.) and software (e.g. operating system, applications, licences, etc.) resources. Note that there may also be other types of resources used by some services. These resources are implicitly inter-related. As in the case of the services described earlier, there is no implied “ownership” of any resource in this diagram. An organisation may use resources provided by another to deliver a service. Resources; as with services; may be provided or consumed.


By treating these underlying resources as just another case of a service, then we can see that we can use the same models (implicit / explicit and dynamic / static relationships) to describe the relationships of both the underlying resources to one another as well as the relationships of these resources to the services they support. These relationships between services as well as those between services and their underlying resources need to be well integrated and optimised. Whilst this simple model implies that the relationships between different services and the relationships between services and resources may be expressed in a similar fashion, this does not imply that they can be implemented in a similar fashion. Currently, relationships between different services lack the efficiencies and optimisations that are found between resources and the executable software they support.

## 2.2 Enterprise grids

An “enterprise grid” is the starting point for this discussion. This is a grid within a single enterprise. The main factor identifying an enterprise grid is that it exists within a single administrative domain. An administrative domain in this context is defined by the existence of a common set of management policies. It is likely to also include a specific set of identifiers to refer to relevant entities. The enterprise grid is made up of the lowest level resources: computing, network, storage, and basic services that make these resources available for use.

### 2.3 Departmental grids

Enterprise grids naturally include those grids implemented across different departments within a wider organisation. A “departmental grid” is implicitly a subset of an enterprise grid. The diagram below shows this relationship:




The enterprise grid can be regarded as a “grid of grids” i.e. a specific instance of a “system of systems”. It could also be regarded as demonstrating some other hierarchical relationship between the departmental grids that make it up. The depth of this relationship will define the capabilities of the enterprise grid and the attributes that it exhibits. A very basic relationship may mean that the ‘enterprise grid’ is nothing more than a loose collection of departmental grids and as such adds little value or capability. At the other end of the spectrum is a more tightly integrated vision that delivers a synergistic value to the entire enterprise through the harnessing of all resources in an enterprise. The ability of a given departmental grid to harness resources within and external to itself will vary, based upon the degree of integration. Some enterprise IT organisations are already setting up internal ‘utilities’ to provide services and resources to other parts of the enterprise. In general, these enterprises have a higher degree of sharing / utilisation than others that are less integrated.

Within an enterprise, resources or services will be owned (in the sense of ultimate authority over usage) by someone. These may be part of an identified 'grid' within a given department or might be owned by 'corporate IT' in some other fashion. In any case, they ought to be able to be utilised (subject to the appropriate permissions etc.) without necessarily formally being part of a departmental grid. Using the 'internal utility' analogy, a resource/service within such a utility can be used by a departmental grid without being part of it, though implicitly if the internal utility is built using grid-like primitives then there is sharing between grids demonstrated here too.

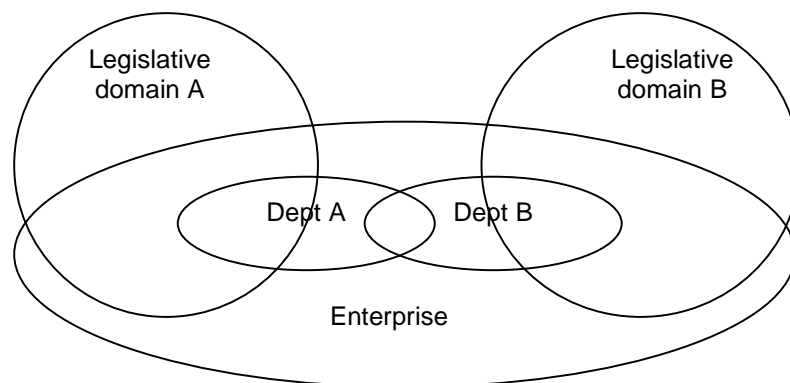
In general, any departmental grids within the same enterprise will come under the same set of enterprise administrative policy rules. For instance, they are both likely to have to conform to similar security or data protection requirements set by the enterprise as a whole. This is not always going to be the case however. Any enterprise that spans multiple countries (and hence multiple legislative domains) may have different local legal requirements (say around the protection of personal information) which need to be complied with. Similarly, some enterprises do not mandate rules to be adopted (or some

departments choose to ignore them!). In any case, there are some problems to be solved in terms of what resources can/cannot be used by whom. Consider the following diagram:



In the diagram above, we see two different departments (A and B) within a single enterprise. Each department is subject to a different legislative domain (A & B respectively). It can be seen that if there is NO overlap between the two legislative domains then there is a fundamental limit in terms of what can be shared between them. Note that this does not necessarily preclude the sharing of computational resources, but may well completely prohibit the sharing of data. This will have implications in terms of how isolation between different services and resources can be achieved (and proven to the satisfaction of the respective legislator).


The reality of the situation is usually such that any given enterprise will have a consistent subset of administrative rules that comply with all of the legislative domains within which they operate. This will allow for the sharing of resources without major restrictions thus:



It is not the purpose of this SRA to define mechanisms to overcome shortfalls in enterprise (or departmental) administrative policy setting or in circumventing legislative restrictions. Departments that are inherently forbidden to share resources with others will be unable to do so regardless of any NESSI framework. The NESSI framework should allow such restrictions to be expressed and enforced. However, in those environments where there are consistent legal rules, then the NESSI framework should allow for the appropriate sharing of services and resources according to these rules.

## 2.4 Virtual organisations

Virtual organisations form when two or more enterprise grids intersect; more correctly when the two administrative domains (and hence their enterprise grids) overlap and share resources. The reality of a VO is that only a subset of the overall grid within an enterprise is likely to be “contributed” to this virtual organisation. The majority of the systems within a given enterprise grid are likely to remain outside the VO. The diagram below shows this VO relationship:



For the most part, contracts in the business world are between two parties. We believe that the abstractions used to describe the interaction of two administrative domains are good enough in general to be a means to describe the interaction of multiple domains (e.g. companies). Consider a tri-party agreement. This may be formed between the parties A, B & C in such a way that individual bi-party relationships describe the interactions, i.e.

- A – B
- B – C
- C – A

An alternative might be the case where the three parties form a consortium D, in which case, the relationships are between the parties and the consortium itself, i.e.

- A – D
- B – D
- C – D

Two kinds of inter-organisation relationship are typical within the business world. The first is a “master / slave” relationship. This is an asymmetric relationship where one party takes on the rules and policies laid down by the other. A typical example of this might be a sub-contractor interaction with a large manufacturer. All of the sub-contractors must meet the rules of the large manufacturer as a condition of doing business with it. One administrative domain has priority over another, such that the policies of the VO describing the relationship may be indistinguishable from the policies of the master.

The second relationship is “peer-to-peer”. This is a symmetric relationship between the two parties. Most inter-organisation relationships fall into this category. This differs from the “master / slave” relationship in that neither administrative domain has priority over the other. The policies of the VO describing the relationship are likely to be some combination of the policies of the contributing organisations.

### 3 Business Scenarios

This section describes a number of common / emerging business scenarios. We will use the basic model of the enterprise grid to show how these may manifest themselves and detail the major business characteristics for them. Scenarios differ in when we expect them to become relevant in the context of Business Grids. The first three business scenarios (enterprise, hierarchical enterprise, hosting) are assumed to have most short-term relevance. Therefore, they are further described by fairly detailed business requirements out of which research challenges are later derived in Section 5.

Business requirements are discussed according to the following major categories:

- *Functional & commercial issues*, i.e. issues stemming from functional requirements typical for enterprises that want to use business grids. Especially important in this area is the commercial context under which business grids are operated. This covers issues such as service level agreements, cost and revenue agreements etc.
- *Dependability*, i.e. the trustworthiness of a computing system in terms of availability, reliability and safety. Dependability covers issues such as redundancy, autonomy, quality of service etc.
- *Security* which we list as separate issue from dependability because of its paramount importance in business interactions. Security covers issues such as confidentiality, authentication, integrity, authorization etc.
- *Performance*, i.e. the system performs and uses resources in an efficient, predictable and accountable way.
- *Interoperability*, i.e. the inter-working with other systems including legacy systems.
- *Manageability*, i.e. the easy, transparent and low-cost management of systems in particular including their maintenance.
- *Governance*, i.e. the specification and assurance of requirements and policies at various levels.
- *Flexibility*, i.e. the capability to react on/implement changed business requirements as fast as possible. Specific issues in here are about scalability etc.

These categories are neither complete nor completely independent from each other. However, they serve as our major perspective under which we present the following requirements as well as the various research challenges discussed in Section 5. Except for the first one, all categories focus on non-functional requirements and as such are often hard to express, measure and quantify in current business systems.

#### 3.1 Enterprise

Basic enterprises consist of one homogeneous administrative domain. Following Section 2, they can be supported by basic enterprise grids which should provide general-purpose infrastructure services within one domain and which should ultimately act as the ICT backbone for the complete IT infrastructure. Consequently, the basic enterprise scenario imposes the following major business requirements.

## Functional & Commercial

- To support the **reliable and secure management of business data**. This is vital to any business as all of its internal activities eventually materialize in changed business data.
- To allow for **operating IT infrastructures as a business**. This includes complete transparency on how system parameters or setup decisions affect the total cost of ownership.

## Dependability

- To allow for **high availability** of (possibly virtualized) infrastructure resources in case of failures or maintenance activities. This contributes to overall high availability of business solutions facing the situation that already short downtimes in service-economies can lead to significant loss of money or even a company's bankruptcy.
- To support the **balancing of availability levels with economic costs**. This allows enterprises to tune the actual availability to the most profitable level.
- Enterprise grids should be **autonomic** in the sense that they should be able to independently repair and recover in the event of errors, including complete failure of system components.

## Security

- To reflect **security policies on infrastructure level** thus guaranteeing integrity and confidentiality of business data at the lowest possible level.
- To support application domain specific end-to-end (infrastructure to user) security demands.

## Performance

- The non-functional behaviour and in particular the performance of applications running on an enterprise grid should be **predictable and accountable**. Thus, users can know in advance what they can get and the resources used can be accounted for properly in order to have a transparent overview of who used which resources. This is especially important in service-oriented systems where services from different providers are composed to make higher-level services.
- To support **massive enterprise job scheduling** by pre-emptive and/or planned allocation/booking of the necessary infrastructure resources.

## Interoperability

- Enterprise Grids can be implemented **transparently to arbitrary business applications**. Thus, the benefits of enterprise grids can be applied to complex existing IT stacks in place today.
- They should be **standards based and offer effective interoperability** with other standards based services within and outside the enterprise.

## Manageability

- To allow for **homogeneous, low-cost, secure, easy and transparent management** of arbitrary infrastructure resources. This leads to maximum exploitation of available infrastructure resources with the lowest possible effort.

- To remain fully-functional while undergoing **incident based or planned maintenance**.
- To enable **on-demand provisioning and termination of services**.

### **Governance**

- To support the **transparent translation of business requirements and policies** to infrastructure capabilities and vice versa thus supporting their reflection and enforcement at infrastructure level.
- To reflect **enterprise policies on infrastructure level** via appropriate logging, tracking and auditing of resources and services. Such infrastructure support for compliance offers most robust and reliable implementation of governance policies even providing some resilience to badly behaving software, unforeseen security holes or erroneous sizing.

### **Flexibility**

- To support **flexible changes of business processes and applications** which impact on the underlying infrastructure, such as dynamic allocation of additional resources. This supports the overarching goal of business agility, a major prerequisite for successful interaction in rapidly changing markets.
- To **dynamically adapt to meet demand** at any point in time in its operations to support sudden ‘bursty’ or unpredictable peak workloads on demand.

## **3.2 Hierarchical enterprise**

Larger enterprises are typically organized via a hierarchy of departments, each of which following the general enterprise policies while potentially deviating from them e.g. due to different legislative constraints. Hierarchical enterprise grids can be built upon a hierarchy of departmental grids all of them tied together in one enterprise grid. Business grids should be able to reflect the different structures of different enterprises. Adopting a business grid should not force a business to change the way that it is organised; rather it should enable businesses to change their organisational structure to meet internal and external demands. Thus, in addition to basic enterprise scenarios, hierarchical enterprises impose the following major business requirements.

### **Functional & Commercial**

- To support both **central and local service level and cost level agreements** as different departments might be operated as almost separated business entities (e.g. different cost centres).

### **Security**

- To implement both **central and local security policies at the infrastructure level** and to support them at higher service levels.

### **Manageability**

- To respect **departmental boundaries in resource management** as departments might be willing to share some resources while others are used exclusively.

- To support a **consistent approach to management from central or local resource and service delivery platforms** including autonomic conflict management and resolution.

### Governance

- To **reflect enterprise policy hierarchies on infrastructure level**, either following a hub-spoke model (where the degree of centralised control and policy consistency with central directives is high) or a distributed model (where the degree of centralised control and policy consistency with central directives is low because individual policy controllers retain significant freedom within a loosely defined Enterprise IT policy framework).

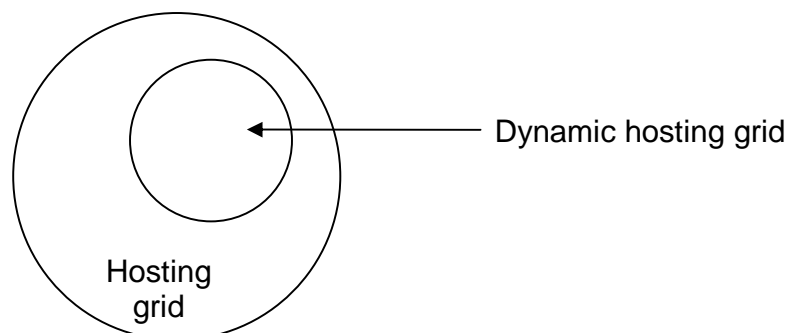
### Flexibility

- To **propagate changes in the departmental structure** of the enterprise automatically to the IT infrastructure.

### 3.3 Hosting: a special case of a virtual organization


An enterprise that provides hosting facilities can also be described using the terminology of an enterprise grid. Today, most hosting can be described to be ‘static’; in that the resources and services are pre-allocated by the hosting company to the using company well in advance of them being used and are typically allocated for a long period of time (weeks, months, years). The relationship between the hosting and using companies is well-defined and also typically long-term. Static hosting relationships are quite predictable and hence can be billed in a well understood fashion. In the future, these static relationships will likely become more ‘dynamic’. Dynamic hosting is typically much more short-term in nature. Hosting resources and services are allocated on demand and in as near to real time as possible. The length of time that a resource or service is allocated to the using company may potentially only be for a few seconds or minutes. Relationships between companies may similarly only exist for the length of time that the resources or services are being used. Dynamic hosting is much more unpredictable and is likely to require new billing and accounting mechanisms.

A hosting company (an enterprise that perhaps provides both ‘static’ and ‘dynamic’ hosting facilities) is illustrated below:




A hosting relationship is thus a variant of the Virtual Organisation, in that it is simply a special relationship between two administrative domains. Here we see a simple

environment where the hosting company (represented by the enterprise on the right) provides resources to the consuming enterprise on the left.



Inter-relationships that are more complex can also exist; for example the provision of dynamic resources by a hosting company to a departmental grid within an enterprise. Such an inter-relationship can be shown thus:



Some complexity starts to arise here. There is a two-level relationship between the two enterprises. The former is the VO that is a high-level interaction between their two administrative domains (a company-to-company agreement). The relationship between the departmental grid and the dynamic hosting environment is then a special case relationship within, albeit still encompassed by the overlap of the two administrative domains. This special relationship is likely to be a subset of the VO relationship. These problems are relatively easy to solve in the case of a static hosting situation. As the hosting needs of the using company become more dynamic, the challenges of building (and maintaining) relationships, of creating the necessary security and interactive linkages and of managing the interaction between the two administrative domains become significantly harder problems to solve.

Apart from the characteristics discussed above, basically all the requirements listed for (hierarchical) enterprise scenarios apply in hosting scenarios but to an even more rigorous extent. These and some further specific requirements are detailed below.

### Functional and Commercial Requirements

- To support **reliable and secure accounting and billing, especially for dynamic, short-term hosting relationships**. In particular, they must support it for different business (e.g. software as a service), financial (e.g. pay as you go, monthly contract, pay in advance, pay in arrears, volume discounts) and licensing models.

- To support **SLA monitoring and penalties management** in order to deal with violated SLAs.
- To support **decisions when to break SLAs**. A provider may service different clients and it is possible that, at some times there are difficulties in complying with all SLAs. The provider may decide that it is more cost-effective to break one or more SLAs in order to be able to comply with the rest (e.g. to prioritise a large or important customer over a smaller one).

### **Dependability requirements**

- To support **flexible, customer-specified dependability levels** for hosted solutions which are automatically mapped to the infrastructure setup and which can even be changed over time and per requested service. Thus, customers can choose the required availability level and balance that with the resulting costs.

### **Security requirements**

- To guarantee **strongest security and isolation** of hosting resources between different customers. As hosting customers put highly sensitive data into a hosting grid they need complete confidence in their secure and isolated management.
- To provide strongest protection against external threats, e.g. from the Internet.

### **Performance requirements**

- To support **highly efficient operations**, e.g. by intelligent sharing of resources among hosting customers. This is of special importance in mass hosting scenarios
- To support **prediction and enforcement of performance characteristics** (response time and resource usage) and to derive them from high-level requirements. This allows the provider to offer appropriate SLAs and plan resource deployment to deliver against them efficiently.

### **Manageability requirements**

- To support **remote infrastructure management**. Though a hosting environment may belong to the administrative domain of the provider, it may be physically collocated with the client's data centre. Such setups require strong support for the remote management of the hosting environment, in particular the hosting infrastructure.
- To support **mass automation** in order to allow management effort to scale to very large numbers of hosting customers.

### **Governance requirements**

- To support **governance and compliance for both providers and customers**.

### **Flexibility requirements**

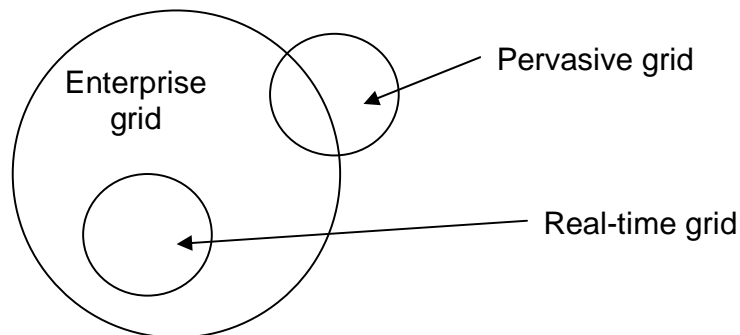
- To provide **hosting services in a highly dynamic and on-demand way** in order to satisfy spontaneous customer needs in extremely short time frames (seconds or minutes).
- To support **dynamic migration of business solutions to/from a hosting environment**. Businesses cannot afford to interrupt their operations just to migrate some IT solutions into a hosting environment or back into their own data centre.

### 3.4 Extended enterprise

Enterprises may intend to extend the scope of their IT-managed activities beyond the traditional boundaries of data centres and desktop machines. For example, pervasive scenarios aim for the integration of pervasive devices (PDAs, mobile phones, RFID systems etc.) into the ICT backbone of a company. *Real-time scenarios aim for the integration of real world entities via sensors/actuators (e.g. in shop floor integration or asset tracking) thus exhibiting more real-time or event-driven characteristics.* Technically, these scenarios might be supported by special-purpose grids, e.g. pervasive grids or real-time grids. Extended enterprises will then seek to integrate those kinds of special-purpose grids into their overall enterprise grid.

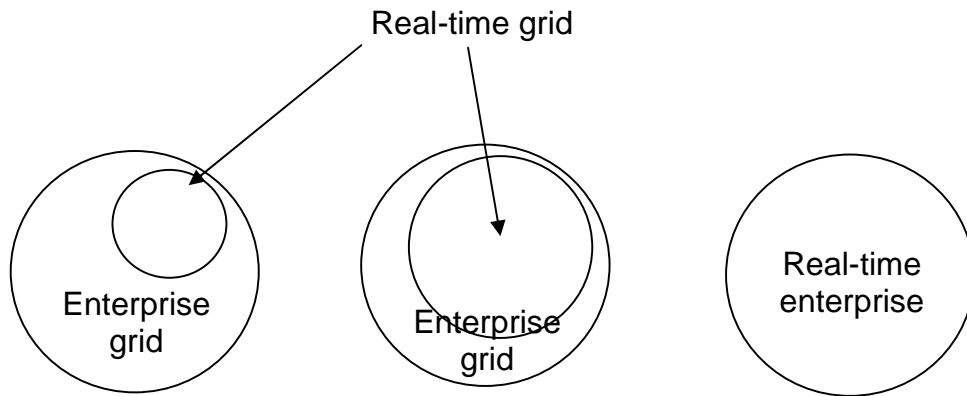
Special-purpose grids could be realized in different ways: within the administrative domain of an enterprise as a specific technical subset, as departmental grids with some specific policies, or as separate administrative domains interlinked with an enterprise grid via mechanisms of virtual organisations.

The following diagram shows possible relationships of these entities to the enterprise grid.




### 3.5 Real-time enterprise

Real-time enterprises are very dynamic, event-driven entities. Typically, real-time capabilities only make up a very small part of an overall enterprise. There is a growing body of evidence to suggest that this event-driven nature will be the key to many enterprises becoming more dynamic. The challenge is how to achieve this. The following diagram shows an initial small subset of an enterprise grid with real time capabilities. Over time, more parts of the enterprise adopt or are equipped with these capabilities and the relative size of the subset grows. Finally, when all components of the grid have some real-time or event-driven capabilities, the enterprise grid will be able to support a real-time enterprise. It may also be possible, should the nature of the real-time operation allow it, for dynamically hosted environments to play a role here. Obviously, if the time it takes to “use” (the whole activate, use, deactivate process) a resource or service provided by a third-party organisation is incompatible with the real-time requirements of the organisation, then this is not going to be worth pursuing.



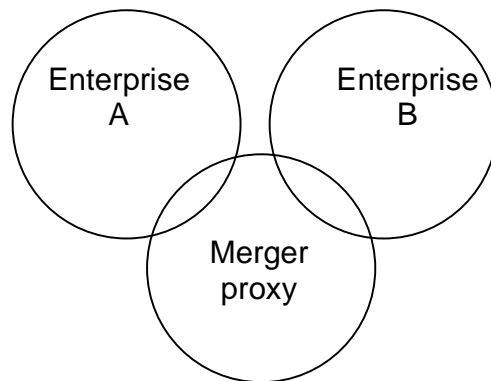
### 3.6 Mergers & acquisitions

Merging with, or acquiring another company is an increasingly common method for companies to grow. The traditional means by which companies merge typically takes a long time to fully complete. This can be envisaged using the enterprise grid model in one of two ways. The first occurs when two enterprise grids merge over time. The diagram below shows this:



The two enterprise grids that are initially separate form a virtual organisation that links their two grids together in a limited fashion. Over time, as the grids coalesce, the virtual organisation encompasses more and more of each individual enterprise until a single merged enterprise remains. In the case of a merger, the relationship is likely to be more of a symmetric one. In the case of a takeover or acquisition, the relationship between the enterprises is more likely to be asymmetric.


One other form of merger / acquisition mechanism is also possible here. This introduces the concept of a merger proxy, as shown below:



The merger proxy is an external entity that provides a link between the two enterprises. Both enterprises form virtual organisations with this proxy and use the services of the proxy to merge logically. Such a proxy might be used to facilitate a new range of business services in this arena.

### 3.7 Business value networks

Business value networks are ways in which organisations interact with each other to drive increased business value. There are many of these in common existence today; the supply chain being the most familiar. The application of enterprise grids to this environment will allow more exotic business value networks to be formed and dissolved as business conditions change. The following shows a business value network around the joint development of a product:



In this example, Enterprises A and D are collaborating on the development of a new product. Enterprise A has sub-contractors or partners B & C, whilst Enterprise D has similar relationships with E & F. The network (shown here with arrows) consists of a series of virtual organisation relationships in both master-slave and peer-to-peer between the respective companies. By pooling resources, information, insight & capabilities a greater business value can be achieved than by a single organisation acting alone.

## 4 State of the art

This section presents a high-level overview on the state of the art of technology relevant for the Business Grid vision. A more detailed review of technologies with respect to the introduced business scenarios and requirements is given within Section 5. The overview here is roughly grouped in the current technology usage in the academic community (across administrative domains), within enterprises (within one domain) and the emerging virtualization technologies usually not associated with grid as such that provide important mechanisms for implementing the requirements of Business Grids.

## 4.1 Scientific Grids

In this section, we introduce some of the major scientific grid approaches. In detail, we will concentrate on standards (OGSA), middleware (Globus, gLite, GRIA) and a collection of selected research projects.

Originally, grid computing was an approach for using a network for sharing computers for compute-intense scientific applications [4]. The term Data Grid was then introduced to refer to an architecture primarily focusing on the distributed management and analysis of large scientific datasets [5]. Nowadays, it is a way of organizing computing resources such that they can be flexibly and dynamically allocated and accessed, often to solve problems requiring many organizations' resources. The Open Grid Forum [6] is the main standards development organization devoted to grid computing, and has resource sharing as a general goal. The Open Grid Services Architecture [7] defines relevant services such as logging, reservation, workflow, etc. In this evolution of Grid to a service infrastructure the Web Service Resource Framework (WSRF) was created as a set of specifications designed to modelling and managing state in a Web Service context.

The open source *Gobus Toolkit* [8] is a widely used technology to establish scientific grids. Globus is not a grid system by itself but provides several tools that can be combined to form an overall grid like system. In detail, basic functionalities are provided that enable the users to share computing power, data, and other tools online across corporate, institutional, and geographic boundaries. To this end, several services like resource monitoring, discovery, and basic management functionalities are implemented. Furthermore, the toolkit consists of components related to security, information infrastructure, data management, communication, and fault detection. Globus users need to combine and extend those basic services in order to build the required grid functionalities. This process further incorporates the adaptation of existing software products to be able to use the given basic Globus grid services. This adaptation is often very time consuming and in some cases simply not possible. Furthermore, the different services provided by Globus only provide basic functionalities. More sophisticated customer oriented solutions must be generated individually by extending or re-implementing the Globus-provided services. In some cases this proves to be difficult, as an extended version of some Globus grid services requires broader interfaces with more functionality of the other Globus grid services.

Another grid computing alternative that is used by many research projects is the *Lightweight Middleware for grid computing (gLite)* [9]. gLite was developed as part of the Enabling grids for E-science (EGEE) project [10]. gLite consist of several grid services that follow a service oriented architecture. This architecture ensures interoperability between the different grid services. Furthermore, upcoming standards can easily be integrated. The services of gLite are grouped in five main categories, namely access services, security services, information & monitoring services, data services, and job management services. The access services simple enable users to contact to the grid system. Security services are responsible for authentication, authorization, and auditing. Furthermore, data confidentiality is ensured. The information and monitoring services are capable to publish and consume information and to use them for monitoring purposes.

They include functionalities like job monitoring, service discovery, and network performance monitoring. The data service consists of the metadata catalogue, the storage element, the file and replica catalogue, and the general data management. The job management service enables the users to allocate resources for their jobs. Furthermore, the resource usage and the corresponding accounting are done by using this service. Overall, gLite is a kind of toolbox that can be used to build grid systems. However, existing applications still need to be adapted and certain services extended to reach the user demand.

GRIA [11] is Grid middleware particularly aimed at business users. It supports use of the Grid in a secure, interoperable and flexible manner by making use of business models, processes and semantics to allow resource owners and users to discover each other and negotiate terms for access to high-value resources. GRIA implements an overall business process to find, procure and utilise resources capable of carrying out high-value, expert-assisted computations. Services from different providers can be combined together via an API in order to create applications.

Beside the introduced middleware systems, several other projects that are sometimes reusing or enhancing elements of Globus, gLite or GRIA should be mentioned. These projects are for example NextGrid [12], Cactus-G [13], NetSolve [14], UNICORE [15], GrADS [16], and the EU DataGrid [17].

## **4.2 Enterprise Grids**

Enterprise grid deployments generally rely on the same base technologies as the academic/scientific grids mentioned above, but due to the lack of adequate security mechanisms (or at least their perceived lack or difficulty of use) are usually restricted to a single administrative domain within an enterprise. It is not uncommon to find deployments that span several administrative domains within an enterprise, i.e. between different arms of the internal IT organization (such as engineering, business and general desktop divisions), but there is no technical support for supporting interactions between these administrative groups in the technical infrastructure deployed. Current enterprise grids rarely support transactional applications. Usually they are used for a few dedicated, compute-intensive applications such as in life-sciences, engineering simulations and financial analysis. There is a dearth of data-management solutions for grids. Managing information life-cycle and data movement is left to individual applications.

Commercially available solutions for these environments feature policy based scheduling and workload management on heterogeneous infrastructures made out of desktops, servers and clusters. These systems contain basic resource control and mechanisms for fault tolerance as well as analysis tools for performance and debugging. Due to the lack of standardization in this space these solutions typically support a variety of *de facto* standards and translate them to a solution-specific format. Finally, these solutions often contain their own billing and user management solutions, partially integrating with common security infrastructures prevalent in enterprises.

Examples of such commercially available solutions (without weighing their market share or importance) are IBM Grid Computing solutions [18], SUN's "N1 Grid Engine" [19] Platform Computings "Enterprise Grid Coordinator" [20], DataSynapse [21], Univa Globus Enterprise [22] or Cluster Resources's "Moab" suite [23].

### **4.3 Virtual machine based grids**

A variety of IT vendors, both large and small, are offering solutions that can transparently manage arbitrary collections of applications in execution containers that can be dynamically deployed onto a homogeneous pool of physical hosts. These containers then are treated like jobs in a "normal" grid but can be dynamically controlled and the resources allocated to them adjusted.

Typically these containers are implemented using virtual machine monitors, and hence contain an operating system together with the deployed application. Examples of this technology are XEN [24], VMWare [25], Virtual Iron [26], SUN's Solaris Container [27] or the (planned for 2007) System Center Manager from Microsoft [28]. Containers can also be implemented without hypervisors inside the operating system (i.e. SWSOft's Virtuozzo [29] or IBM's virtualization suite [30]).

Features of these solutions are similar in capability to what enterprise grid solutions offer, but add "steady state" resource management, additional security mechanisms by encapsulating applications in the containers and by adding control of network partitions through setting up of virtual VPNs.

## **5 Research challenges**

In the following, we describe the resulting main research challenges of the described business scenarios. To this end, we follow the categories introduced in Section 3. For each category the state of the art is summarized. Then, the research challenges for the scenarios enterprise, hierarchical enterprise, and hosting are derived accordingly.

### **5.1 Functional & Commercial Issues**

Traditional grid systems focus on isolated batch jobs while virtual-machine based grids address complete systems to be submitted. However, Business Grids focus more on interactive applications with a low response time. Furthermore, the individual applications are not completely decoupled in the sense, that they are executed as part of business processes or coexist with other collaborating applications.

Though many grid systems provide basic means for accounting, this typically does not allow associating accounting data with related business activities or decisions. This significantly restricts the cost transparency of IT operations. Furthermore SLAs typically relate to technical properties rather than being expressed at business level. Consequently, business properties are not properly translated to technical properties and vice versa.

The major research challenge in this area is how to provide grid systems that support typical business applications and consequently the transactional, interactive management of business data. Special issues are efficient infrastructure support for data management, integration and mining. These challenges require research on harmonizing existing grid

approaches, namely traditional and virtual machine based grid systems, towards Business Grids that leverage the best capabilities of both areas into a common environment. In particular, research is needed on how users can specify business processes and how business processes can be implemented so that they can be efficiently mapped on available resources.

Another research area is how to provide precise accounting across all layers of an IT infrastructure so that ultimately precise TCO (total cost of ownership) calculations or cost level agreements can be done. Likewise, translation of flexibly defined SLAs across all layers and SLA enforcement on infrastructure layer including an automated penalty management is a major research area. These commercial research challenges are most important for hosting scenarios but also increasingly important for the (hierarchical) enterprise scenario.

## **5.2 Dependability**

In today's systems dependability characteristics are typically implemented in a static manner and most often only at the level of infrastructure resources (e.g. the number of redundant nodes assigned for a certain activity). Characteristics are not derived from application level requirements and consequently their economic costs are hard to qualify.

Business Grids aim for dependability requirements specified individually by customers. Consequently, the most important research challenge in this area is the automatic mapping of the customer requirements to the hardware level and the aggregation of resulting costs back to the customer. While cost awareness is highly relevant for all business scenarios, complete automation of the requirement mapping is of outstanding relevance for hosting scenarios.

Another research challenge is on harmonizing traditional and virtual-machine based grid technologies for supporting autonomic behaviour, i.e. system management with a minimum of human interference, via so-called self\* mechanisms (e.g. self-configuring, self-healing, self optimizing) which aim at cutting the costs of system administration, but also ensuring maximum service availability and performance., and automate the management of complex systems.

## **5.3 Security**

At the moment, grid security mechanisms target traditional infrastructure resources (e.g. compute nodes, files, devices, .etc.) in a rather static way and do not consider the overlaying organizational structure. Enterprise security mechanisms are focused on organizational and user roles but make assumptions about the underlying infrastructure topology which is subject to change in grid environments.

Business Grids need to provide security mechanisms for possibly virtualized infrastructure resources (e.g. a virtual node or a distributed grid job) with a granularity and flexibility that allows the reflection of higher-level security zones such as a process group, a business process or organizational entities including departments, customers and partners. Consequently, Business Grids require advanced security models and policies that can accommodate both organizational structures and infrastructure topologies. This is most important for hosting scenarios, where systems process applications of

competitors within one operational domain. In this context, complete isolation on all levels (application to infrastructure) is necessary. These challenges are synchronized with the ones stated by the European Security Forum for Web Services, Software and Systems (ESFORS) [31].

A second research challenge is how to enhance security at infrastructure level so that external threats (e.g. from the Internet) are optimally addressed. Such deeply integrated mechanisms may support traditional techniques such as network-level firewalls or might turn them even obsolete.

## **5.4 Performance**

Performance issues are poorly addressed in current systems. Thus, specific performance metrics of the system cannot be reliably predicted in advance. This results, in some sense, in unpredictable costs as the final resource consumption of certain applications can only be estimated very roughly. Furthermore, scheduling mechanisms are typically restricted to one virtualization layer and not harmonized across layers, resulting in sub-optimal balancing of work loads.

More precise performance characterisation techniques are necessary for business processes. This would enable the system to predict certain characteristics of the resources and also enable the provider to have more accurate activity-based accounting.

Furthermore, specific performance estimation is necessary in order to apply advanced scheduling techniques. Advanced scheduling systems are necessary in order to allocate and reserve resources in advance (for co-allocation problems) and to enable the system to effectively apply pre-emption. Such advanced scheduling has to incorporate the logical boundaries in hierarchical enterprises and in hosting scenarios. Finally, scheduling algorithms have to be researched that take possibly different virtualization layers into account.

A specific research challenge is on resource sharing in mass hosting environments, i.e. research in sharing compute, network and storage resources among hosting customers while still keeping them effectively isolated.

## **5.5 Interoperability**

In today's grid-like systems, interoperability is restricted in two dimensions: First and in contrast to virtual-machine based grids, traditional grid systems require the specific engineering of applications to be "gridified". Second, different grid systems cannot easily be connected as they are based on different standards and use different interfaces.

Consequently, Business Grid research has to address the harmonization of traditional and virtual machine based grid capabilities to allow for executing existing business applications without further modifications. A second research area is on open standards that support the interoperability of different Business Grid flavours. This is especially important for the hierarchical enterprise scenario.

## **5.6 Manageability**

Management tasks are still a major cost driver in today's business systems. In particular, many management tasks are not yet extensively automated. Grid systems provide some

support for automated management of infrastructure resources but typically follow a narrow perspective that simply covers the management of resources as assets rather than looking at the overall business solution stack.

Flexible manageability as envisioned by Business Grids requires research in virtualization technologies across a number of layers (hardware, operating systems, middleware, applications, etc.). Virtualization allows for decoupling of management decisions, thus giving more flexibility. Moreover, research in analysing and presenting the business impact of management decisions (such as SLA violations, costs etc.) is needed.

These challenges apply to all of our business scenarios. A specific research challenge for hierarchical enterprises involves concepts for mediation and conflict management in distributed environments. Hosting scenarios require special attention on remote management and mass automation support (e.g. cloning of customer systems).

## **5.7 Governance**

Today, enforcement of governance policies is poorly supported and if at all requires manual intervention by administrators. In particular, support for complex organizational setups (such as hierarchical enterprise) is almost non-existent.

An important research challenge is therefore the enforcement and tracking of policies at the infrastructure level. To this end, certain techniques for logging, tracking, and auditing of resources are necessary. Second, the automatic translation of business policies to policies at the infrastructure level is to be explored. A third issue is on automated mechanisms for enforcement and conflict resolution.

These challenges apply to all business scenarios. However, the hierarchical enterprise scenario requires additional research on how to derive models and techniques to specify and reflect a hierarchical structure of policies. The corresponding models must be very flexible such that they can easily be adapted for the different enterprises.

Finally, hosting scenarios require governance support for both providers and customers.

## **5.8 Flexibility**

Today, business decisions are often implemented in a fixed way, e.g. a fixed landscape design reflecting security zones or a static deployment of applications on resources.

Flexibility as envisioned by Business Grids requires research in two main areas: the automation of information flows and management processes from business process to infrastructure, and the dynamic adjustment of infrastructure resources. As automation issues are already discussed in other sections we focus on resource topics.

All of our scenarios require research in dynamic allocation and release of resources. This also includes migration mechanisms for supporting load balancing and scalability requirements. In contrast to eScience applications, the migration of business applications may also include migration of related application pools and data sources. Finally, research needs to investigate the interplay of migration and scheduling mechanisms.

## 5.9 Overarching challenges

Finally, there are some overarching research challenges which cannot be directly derived from the listed business scenarios and business requirements. Those research challengers cover:

- **Reduction of the complexity** of grid technology. Grid technologies are complex and will grow in complexity with the introduction of automation and self-\* characteristics. For this reason, installing, configuring and developing applications on this environment currently requires specialised staff. With the aim of reducing complexity of the grid infrastructure new abstraction layers and user friendly interfaces will be created which will facilitate the work of administrators and developers.
- **Enabling enterprise business goals through architecture-driven development.** Domain-specific enterprise grid scenarios place particular requirements for architectural frameworks that create a fabric for domain-specific processes, methods and components to interwork in a coherent and consistent manner. Service grid architectures address these requirements, which are typically related to trust, security, dependability, performance and service coordination aspects. Distributed service grid architectural frameworks specific to the domains addressed, will therefore need to be created as a crucial step to enabling domain-specific enterprise grid scenarios.

## 6 Conclusions and Outlook

Summarizing the presented research challenges we can identify the following 3 key challenges for Business Grids following the paradigm of a service-oriented knowledge utility:

**KC1: Harmonizing traditional grid and virtualization technologies.**

This allows building truly interoperable core infrastructures that combine the best of both worlds, in particular by providing strong support for typical business applications while inherently reflecting different organizational structures.

**KC2: Automated translation of high-level requirements and policies** of all kinds to the infrastructure level and aggregation of infrastructure behaviour to higher-level information.

This allows the business value of requirements, policies and capabilities to be directly related to properties of the infrastructure and eventually supports managing IT as a business.

**KC3: Automated SLA and policy enforcement at infrastructure level.**

This allows reliable (resilience to higher level violations) and reproducible (due to well-defined linkage to higher level business goals) enforcement while automation contributes to cost reduction and responsiveness.

These challenges apply to all the short-term business scenarios analysed so far, namely enterprise, hierarchical enterprise and hosting.

Future versions of this document will focus on the in-depth analysis of more advanced and mid- to long-term scenarios such as the briefly-sketched extended enterprise, real-time enterprise, merger & acquisition and business value networks.

## 7 References

- [1] NESSI: *SRA – Volume 1 – Framing the future of the service oriented economy*. Jan.2006. [http://www.nessi-europe.com/documents/NESSI\\_SRA\\_VOL\\_1\\_20060213.pdf](http://www.nessi-europe.com/documents/NESSI_SRA_VOL_1_20060213.pdf)
- [2] NGG3 Expert Group Report: *Strategic Future for European Grids: Next Generation GRIDs based on SOKU - A new paradigm for service delivery and software infrastructure*. December 2005
- [3] Jim Gray: *Distributed Computing Economics*. Technical Report, MSR-TR-2003-24, Microsoft Research, March 2003.  
[http://research.microsoft.com/research/pubs/view.aspx?tr\\_id=655](http://research.microsoft.com/research/pubs/view.aspx?tr_id=655)
- [4] I. Foster, C. Kesselman, S. Tuecke: *The Anatomy of the Grid: Enabling Scalable Virtual Organizations*. International J. Supercomputer Applications, 15(3), 2001.
- [5] The Data Grid: *Towards an Architecture for the Distributed Management and Analysis of Large Scientific Datasets*.  
<http://www.globus.org/alliance/publications/papers/JNCApaper.pdf>
- [6] Open Grid Forum. <http://www.ogf.org/>
- [7] S. Tuecke et al.: *Grid service specification*. June 02. <http://www.gridforum.org/ogsi-wg>
- [8] Globus. <http://www.globus.org>
- [9] gLite. <https://edms.cern.ch/file/476451/1.0/architecture.pdf>
- [10] Egee. <http://egee-intranet.web.cern.ch/egee-intranet/gateway.html>
- [11] GRIA - Grid Resources for Industrial Applications. <http://www.gria.org/>
- [12] NextGRID: *NextGRID Architecture White Paper*. December 2005.  
[http://www.nextgrid.org/white\\_paper.htm](http://www.nextgrid.org/white_paper.htm)
- [13] G. Allen et al.: *Cactus-G: Enabling High-Performance Simulation in Heterogeneous Distributed Computing Environments*. <http://edoc.mpg.de/3477>
- [14] NetSolve. <http://icl.cs.utk.edu/netsolve/>
- [15] Unicore Forum. <http://www.unicore.org/>
- [16] GrADS - Grid Application Development Software Project.  
<http://hipersoft.rice.edu/grads/>
- [17] Datagrid. <http://eu-datagrid.web.cern.ch/eu-datagrid>
- [18] IBM: IBM Grid Computing. <http://www-1.ibm.com/grid/>
- [19] SUN: Sun N1 Grid Engine. <http://www.sun.com/software/gridware/>
- [20] Platform Computing: *Enterprise Grid Coordinator*.  
<http://www.platform.com/Products/Platform.Enterprise.Grid.Orchestrator/>
- [21] DataSynapse products. <http://www.datasynapse.com/products/default.asp>
- [22] Univa Globus Enterprise. <http://www.univa.com/products/>
- [23] Cluster Resources: *Moab*.  
<http://www.clusterresources.com/pages/products/moab-cluster-suite.php>
- [24] University of Cambridge: *XEN*. <http://www.cl.cam.ac.uk/research/srg/netos/xen/>
- [25] VMWare: *Virtual Infrastructure 3*. <http://www.vmware.com/products/vi/>
- [26] Virtual Iron: *Virtual Iron*. <http://www.virtualiron.com/>
- [27] SUN: *Solaris Container*. <http://www.sun.com/software/solaris/utilization.jsp>
- [28] Microsoft: *System Center Manager*.  
<http://www.microsoft.com/systemcenter/scvmm/default.mspx>
- [29] SWSoft: *Virtuozzo*. <http://www.swsoft.com/products/virtuozzo/>
- [30] IBM: *Virtualization Suite*. <http://www.ibm.com/virtualization>
- [31] ESFORS: ESFORS Position Statement. Version 2, June 2006. <http://www.esfors.org>