NESSI-Grid contributing to NESSI

A NESSI-Grid White Paper

A community-driven Vision on Business Grids Turning Grids into a backbone for Arbitrary Business Applications

Foreword

The NESSI-Grid SRA focuses on opportunities for grids to evolve into major business infrastructures for the more agile and effective delivery of services. The SRA has been developed by the community for the community and details how this evolution can be achieved through a set of research challenges based on business scenarios and current technology trends. The innovative infrastructure is underpinned by the concept of Business Grids. This White paper offers a high-level overview on the SRA.

June 2008 **www.nessi-europe.eu**

NESSI – the Networked European Software and Services Initiative – was launched as a European Technology Platform in September 2005. Coordinated by 22 partners, it unites a community of 300 organisations from industry and academia active in Information and Communication Technologies. It plans to deliver NEXOF, the NESSI Open Service Framework.

NESSI-Grid is a support action launched under the FP6 European Commission research programmes which contributes to NESSI. NESSI-Grid started in May 2006 and will end in October 2008.

Table of contents

Introduction

Evolving Service-Oriented Economies (SOE) bring new requirements for businesses and services, which will need to react more quickly to changing circumstances. This means that businesses must adopt more flexible service-oriented business processes and be able to develop collaborative relationships. This is only possible if high-level business requirements are transformed into lower level ICT requirements with a high level of automation, so that ultimately the ICT environment can automatically cater for changing business needs.

The SRA has been developed through a structured and open community process managed by NESSI-Grid. Special emphasis is placed on the opportunities for Grids to evolve into major business infrastructures for the delivery of services, detailing and prioritizing a set of business scenarios. The vision centres on the shift towards a Service Oriented Economy based on the concept of Business Grid, which is defined as an adaptive service-oriented utility infrastructure for business applications. The SRA explains how this evolution could be achieved through a set of research challenges based on a multi-perspective approach with reference to societal trends, constraints, business scenarios and technology trends.

The methodology of the SRA is based on three steps. Firstly, identifying the requirements that a business grid must fulfil to satisfy current and future business needs. Secondly, evaluating the state of the art to understand whether these requirements can be met with current technology. Thirdly, pinpointing the gaps in the state of the art, in order to develop key research challenges.

The SRA is addressed to a broad spectrum of stakeholders, particularly enterprises, researchers and policy-makers. Industry stakeholders are offered a vision for future IT infrastructures, their related business scenarios, the state of the art for these scenarios, as well as the specific business impact of future research. Researchers are provided with guidance for industry-focused research topics. Policy-makers will gain an understanding of possible areas of influence and expected market impact.

Business Grids & the Shift Towards Service Oriented Economies

Business Grids are an adaptive service-oriented utility infrastructure for business applications and will form the general ICT backbone in the future knowledge and service economies.

***Business Grids** as adaptive service-oriented utility infrastructure*

The Service Oriented Knowledge Utility (SOKU) paradigm identifies a flexible, powerful, and cost-efficient way of building, operating and evolving IT intensive solutions for use by business, science and society. Applying the SOKU paradigm to Business Grids means that they offer infrastructure resources to higher levels based on the following main principles:

- **Service-oriented**: dynamic allocation & assembly of resources via infrastructure services
- **Knowledge-assisted**: translating high-level business requirements to infrastructure requirements & infrastructure capabilities vice versa
- **Utility**: immediately available, dependable usage, predictable

Business Grids will span several ICT layers so that they can participate in the execution of the business models for the applications they support. The aim is to provide transparent accounting and addressing issues such as auditing, billing, and linkage of resource consumption to business goals. At the infrastructure level, an adaptive service-oriented utility infrastructure will enable the dynamic, on-demand allocation and assembly of resources needed to support service components.

Business Grids will become the backbone of the ICT infrastructure through the following steps:

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

Business Grids are the Backbone of ICT infrastructure

As the interworking between administration domains becomes commonly accepted & is supported as part of the infrastructure there will be a shift towards:

3. Business Grids providing the ICT infrastructure to support service-oriented economies (SOE) & ultimately support the emergence of new types of applications.

Business grids denote the specific adaptation of the grid paradigm for the context of business applications. In terms of Foster's definitions, Business Grids follow an open-standard-based approach, key for achieving a utility infrastructure across administrative domains. Business grids also relate to Foster's definition that business applications require non-trivial qualities of service due to their highly networked nature, that is, their co-existence with other applications and services.[1]

1. Ian Foster, *What is the Grid? A Three Point Checklist*, July 2002, http://www-fp.mcs.anl.gov/wwwfp.mcs.anl.gov/~foster/Articles/WhatIsTheGrid.pdf

Beneficiaries of Business Grids

Business Grids have the potential to bring benefits to a wide range of organisations and sectors, as summarised below:

- **Commercial enterprises** where the main adoption drivers are improving performance, bringing significant reductions in computing time, the possibility of doing new things, implementing new services, as well as reducing time to market and costs.

Sectors & Main Uses comprise: **Aerospace & Automotive** (collaborative design & modelling); **Architecture** (engineering & construction); **Electronics** (design & testing); **Energy** (oil & gas for exploration); **Financial Services** (stock analysis & risk management); **Life Sciences** (especially pharmaceuticals & biotech firms); **Manufacturing** (inter/intra-team collaborative design, process management); **Media/Entertainment** (digital animation); **Telecommunications** (processing & analysis of large amounts of data); **Transportation** (modelling & simulation)**; Utilities** (improving efficiency & dealing with peaks & valleys in utilization).

- **Academia** and **research** with emphasis on collaborative work and the utilization of resources.
- **Public Sector** as a major user of high power computing.

The main players comprise: middleware, application, hardware, network, content, service, and payment providers; grid systems integrators; users from the private & public sectors; regulatory bodies (standard bodies & policy-makers); identity providers as certification authorities & other influential groups, such as the media and consultants.

Business Scenarios

The core business scenarios and associated business requirements cover short-term perspectives (basic enterprise, hierarchical enterprise, hosting) as well as mid- to long-term prospects (extended enterprise, merger & acquisition, virtual organizations, dynamic outsourcing, value chains and mega services). Community feedback indicates that Enterprise, Hierarchical Enterprise and Hosting are the most relevant scenarios, though even the lowest ranking (Mergers & Acquisitions) are viewed to be pertinent.

Business Grids are envisioned to support the various scenarios, which can be combined to form more complex, compound scenarios, such as Business Applications; Enterprise; Hierarchical Enterprise; Hosting; Extended Enterprise; Dynamic Outsourcing; Mergers & Acquisitions; Virtual Organizations; Value Networks; and Mega Services.

9 Key Requirements

Each scenario has been analyzed in terms of **nine key requirements**, which are functional & commercial, and non-functional: Dependability, Security, Performance; Interoperability; Manageability; Governance; Flexibility; and Sustainability along with the expected business

impact. A sample of the expected impact of Business Grids on the basic Enterprise, Hierarchical Enterprise & Hosting is provided in the tables below.

Basic enterprises are set up as one homogeneous administrative domain. They should provide general purpose infrastructure services within one domain & ultimately act as the ICT backbone for the entire infrastructure. The SRA identifies ways in which the Enterprise can be extended to include mobile devices for access to grid and sharing mobile device resources, bringing new services.

ENTERPRISE	
VISION	**EXPECTED IMPACT**
Service-oriented; knowledge-assisted; utility. Business Grids as an enabler for agile business.	• Providing IT resources as a service to the entire organization. • Significantly improving flexible IT resource usage; transparency on IT resources to foster further industrialization & adoption. • IT provisioning & management as a utility, bringing significant operational savings.

Larger enterprises are usually organized through a hierarchy of departments with individual departments complying with general enterprise policies while potentially deviating from them because of different legislative constraints.

HIERARCHICAL ENTERPRISE	
VISION	**EXPECTED IMPACT**
Harmonized enterprise-wide infrastructure built on hierarchy of departmental grids. Business grids as enabler for efficient, flexible, & transparent IT operation in complex enterprise set-ups. Enabler for implementing changes in the enterprise organizational structure to meet internal & external demands.	• Enabling the sharing of IT resources in a service-oriented way. • Transparently mirroring the organizational structure of different enterprises. • Supporting automated IT provisioning & management on the enterprise level (where possible) & on the department level (where needed).

Hosting environments provide resources and services, usually by pre-allocating them ("static hosting"). Key challenges are associated with creating dynamic hosting scenarios, that is, between several administrative domains.

HOSTING	
VISION	**EXPECTED IMPACT**
Adding transparency & dependability. Supporting highly dynamic hosting scenarios. Significantly impacting hosting scenarios by supporting new hosting business models: software as a service; on-demand computing & enabling short-term relations. Much more flexible interactions between several organizations.	• Automation support for negotiation & set-up of hosting relationships & Service level Agreements (SLAs). • On-demand (or at least near real-time) allocation of resources & services. • Comprehensive management of complete hosting lifecycle from enactment to decommissioning.

Five Technology Trends Impacting on Business IT Systems

The Research Agenda has identified five Technology Trends that will have a significant impact on future business IT systems and specifically on Business Grids: Storage & Data Management; Processor Technology; Network Connectivity; Mobile Devices; and Sustainability. A number of initiatives are tackling issues, such as the increasing amounts of data generated by businesses and research projects, rapidly growing needs and costs for storage and challenges surrounding data management. Societal concerns also surround increasing energy consumption. More research, targeted ICT policies and coordinated engagement with standard development organizations (SDOs) are needed to support these endeavours, bearing in mind the specific drivers, needs and practices of the business and scientific communities concerned.

The SRA evaluates a number of key research fields: Data Grids; Database Replication; Multi-tier Systems; Application Development; Performance Engineering; Cost & Revenue Management; Autonomic Computing; Virtualization; Security; Interoperability; Manageability, Flexibility & Reconfiguration within the framework of Scientific Grids, Enterprise Grids, Virtual Machine-based Grids, and Mobile Grids. Specific requirements can be matched to top-level challenges and priorities for research, which will ultimately bring benefits to EU's business community.

Match Key Research Fields to top-level challenges

A good case in point is the integration of mobile devices and pervasive scenarios involving embedded processing capabilities into the enterprise grid, which could lead to new opportunities for grid access and the creation of new mobile services. While software developers would need support to produce effective, robust applications that hide the complexity, the shift towards an open infrastructure would ultimately make it easy for developers to write and deploy applications and services. This challenge is closely connected with the creation of new skill sets for programming and scripting languages.

1. STORAGE & DATA MANAGEMENT	
Drivers & Needs	**Challenges & Opportunities**
Compliance with Sarbannes-Oxley & Basel-II is driving businesses towards large-scale storage of all relevant financial records, including email. Data volumes in typical businesses are growing at around 50% p.a. with budgets for storage exceeding other IT hardware budgets.	• Clear need for storage & data management systems so that organizations can keep pace with the rapid growth in volume & complexity of usage. • Storage should provide Business Grid applications with the ability to maintain persistent data & retrieve it as and when required.

2. PROCESSOR TECHNOLOGY	
Drivers & Needs	**Challenges & Opportunities**
Multicore processors offer a way to maintain the rate of improvement in available computing power.	• Benefits include processing power with very low latency & also in power dissipation, an increasingly important factor.

3. NETWORK CONNECTIVITY	
Drivers & Needs	**Challenges & Opportunities**
Access networks use different devices and have different features in terms of bandwidth, latency & reliability with important implications for application performance.	• As Business Grids develop, the role of networks to interconnect resources in different locations will become increasingly important.

4. MOBILE DEVICES	
Drivers & Needs	**Challenges & Opportunities**
PDAs, mobile phones, portable media players, sensors, etc are becoming increasingly important constituents of the global networked ICT ecosystem. Mobile phones are increasing computational capacities using multiple approaches to network connectivity.	• Incorporating mobile devices into Business Grid scenarios means that their context (geographic location, connectivity features) has to be taken into account.

5. SUSTAINABILITY	
Drivers & Needs	**Challenges & Opportunities**
Energy efficiency in data centres for business and the environment is the most important issue facing technology providers and their customers today (The Green Grid Consortium). Energy costs are projected to exceed 50% of the total IT budget over the next few years. The security of supply for large data centres could also become an issue in some locations. Demand for computing capacity continues to grow resulting in increased energy requirements for hardware & associated cooling systems. Important implications for operational costs and for sustainability & climate change.	Energy efficiency optimization is an ever important factor affecting the IT infrastructure. Priorities include: • More efficient hardware & data centre architectures. • Integrated management solutions that cater for ambient temperature & energy consumption. The ability to demonstrate high energy efficiency is expected to serve as a driver in the adoption of Business Grids and an advantage over stand-alone deployments. Regulation & legislation are expected to grow in importance.

Specific Research Challenges to Advance the State of the Art

The research challenges are derived from the nine business scenarios & the five technology trends. The thirteen specific challenges each have a set of related challenges and are ranked in terms of priority, complexity and expected time-lines. **Four functional & commercial requirements** have been identified and comprise Data Management; Applications Development; Network connectivity; and Accounting & SLAs. In addition to the **nine non-functional** requirements, which are Dependability; Security; Performance; Interoperability; Manageability; Governance; Flexibility; Mastering Complex Systems, a core set of **Overarching Challenges** have been defined.

The sample of high priority challenges below illustrates how specific research activities could advance the state of the state of the art for both functional and non-functional requirements. Specific examples are given for functional requirements (data management) and non-functional requirements (dependability; security; performance & interoperability).

Electronic data stored in files & databases is the highest value asset for commercial organizations. Businesses therefore require a data management system with a real potential for improvement in performance, reliability, scalability & integration. To date, there has been limited focus in R&D projects on issues surrounding commercial data management systems. Coupled with this, there has not been a sufficient number of targeted efforts to satisfy the required level of reliability for enterprise applications.

Data Management is a good example of why research in this area is a high priority, in order to meet three key functional and commercial requirements in the Enterprise Scenario*:*

1. *Fast provisioning of systems (time to market) & low costs.*
2. *Reliable & secure management of business data.*
3. *Operating IT infrastructures as a business.*

Data Management Systems allowing improved reliability, scalability, integration & performance.

Three **inter-related challenges** for Data Management are scalable data management; low latency geo data management and autonomic data management. All these challenges are high priority with medium-high complexity and achievable in the medium-long term. A summary of the state of the art serves to highlight the impact of research tackling these related challenges.

Current scale-out approaches achieve a scalability of a few tens of sites. Specific research for **scalable data management** would help develop data-base management systems with high scalability for shared data bases, which are the main requirements for Enterprise Grids managing large amounts of data. Additionally, new approaches to data replication are needed to ensure scalability at hundreds of nodes with low replication overheads.

Low latency is another important goal for the future Business Grid. At present, geographically replicated data shows poor performance. Several business scenarios need geographical distribution of data, tolerance of catastrophic failures due to network connectivity problems, & low latency to distant clients. Achieving **Low Latency Geo Data Management** would enable shared & geographically distributed data to change dynamically and ensure low latency for all clients.

Data management systems need to be able to handle heterogeneous underlying resources & apply them effectively to support business processes & applications. **Autonomic Data Management** is needed to ensure that the systems:

- are self-healing; provide high performance independent of the workload with continuous reconfigurations to maximize performance that is, self-optimizing.
- cater for file systems that can accommodate high & variable latency, complex failure modes & cross-organization operation, including VOs.

From a non-functional perspective, key business requirements have been mapped to specific research challenges to illustrate their impact on the future Business Grid. The examples, which focus on high priority research challenges, include dependability; security; performance and interoperability. The tables below outline the specifics and research challenges for each of the four requirements with reference to current state of the art.

DEPENDABILITY	
Requirements for the Business Grid	**Related Research Challenges**
• High-availability of infrastructure resources. • Balancing availability levels with costs. • Autonomic: automatic repair & recovery in the event of error.	Open Challenge: extending self-healing to complex, multi-tier systems typical of the business infrastructure – high priority. Cost-awareness, which is particularly important for Hosting, and mobile dependability are medium level priorities.

SECURITY	
Requirements for the Business Grid	**Related Research Challenges**
• Security Policies on infrastructure level with guarantees for integrity & confidentiality of business data. • End-to-end (infrastructure to user) security demands. • Supporting several levels of security authentication.	Security considerations are a major roadblock for commercial adoption of Grid. Business Grids need to provide security mechanisms for lower level virtual execution environments with granularity & flexibility enabling a reflection of higher-level security zones: business group, business process or organizational entities – high priority. This approach complies with recommendations of ESFORS (EU Security Forum for Web Services) & SOI formulated by NESSI. Enforcement & brokering of security policies are also high priority.

PERFORMANCE	
Requirements for the Business Grid	**Related Research Challenges**
• Prediction & accounting of non-functional behaviour (applications, services, resources). • Massive Enterprise job scheduling with pre-emptive/planned allocation.	Various virtualization technologies exist but do not indicate predictable impact. Scheduling solutions focus on batch-like jobs. Predictions do not cater for dynamic situations. Resource management is largely homogeneous. Top priorities include: • Develop system & models for virtualization technologies enabling impact analysis of different workloads. • Better understanding of the impact of scheduler decisions across layers & for transactional applications.

INTEROPERABILITY	
Requirements for the Business Grid	**Related Research Challenges**
• As grid solutions become more widely adopted, the need for interoperability & standards increases. Interoperability is crucial for organizations connecting grids within their own & other organizations. • Standards-based with effective interoperability.	There are a few standards that cater for specific enterprise needs (ETSI, OGF). There is a need for open standards ensuring interoperability of grid applications at the infrastructure level and middleware level. Top priority: • Engagement with the standardization process, providing conformance test methods, promoting wide adoption.

Requirements for Manageability; Governance; and Flexibility range from catering for flexible changes of the business process & applications to remaining fully functional during incidents and planned maintenance. The challenges surrounding these requirements are main low-medium priority and achievable in the short to medium term. However, resource management in multiple admin domains and dynamic resource allocation & release are high priorities.

Overarching Challenges: Requirements & Impact

Highlight Business Value & reduce complexity

Overarching challenges entail demonstrating the business value irrespective of the specific technology adopted; reducing complexity & the need for an architecture-driven approach to grid. Reducing the complexity of Grid technology is high priority with medium-high complexity and achievable in the medium to long term.

- The ultimate goal is an "invisible" infrastructure so that businesses can focus on their main concerns. The challenge lies in providing a highly automated configuration & management; different levels of abstraction; user-friendly interfaces; a focus on scalability & interoperability without universal standardization. Raising the level of abstraction is connected with Application Development so there is a need for domain specific languages, tools & development environments that hide complexity & simplify the development of business applications for a specific domain. Understanding software lifecycle issues in open service oriented architecture is also important.

Direct contribution to business value is a medium-high priority with medium-high complexity and achievable in the medium to long-term.

- A business Grid should bring benefits that are clearly linked to the business goals of its users, enabling straightforward value assessments. The challenge is to map high-level business goals & technical functionality. Mapping should be flexible to cope with changes in priorities & business models.

Architecture-driven solutions are a medium-high priority but with high complexity and achievable in the medium-long term.

- The main goal is to define common principles, in order to construct & operate grids that support reproducible & interoperable solutions. The challenge is to build a consensus on best practices & standards for building & operating business grids that cater for a variety of scenarios.

The framework should ensure that the service ecosystem is:

- **Well-governed**: policy & rules driven, monitored by SLAs.
- **Well-balanced** between the need of flexible business flows & mandatory technology processes.
- **Well-orchestrated** at the level of complex end-user services.
- **Well-integrated** at the level of software systems powering core & derived services.
- **Well-instrumented** at the level of the technology stacks powering the software systems.
- **Well-supported** at the level of infrastructure.

Three Key Challenges to Achieve the Business Grid Vision

The extensive set of research challenges identified would be best tackled in combination rather than in isolation, as several of them are inter-dependent and some partly overlap. Joint research programmes would therefore be beneficial to tackle not only top-level but also a number of inter-related challenges more effectively. An assessment of all the challenges has enabled the identification of **three key challenges**, that is, **architectures**; **lifecycle management** and **infrastructure technologies**, where significant progress in the state of the art is needed.

THREE KEY CHALLENGES
1. New system architectures • Harmonizing service architectures (SOA) and infrastructure architectures (SOI). • Advancing the structure of multi-tier, federated & internet scale architectures. • Supporting all kinds of business models, applications and emerging hardware environments. • Providing transparent & integrated access for all relevant stakeholders: architects, engineers, operators, customers.
2. Advanced System lifecycle approaches • Engineering, deployment, composition, provisioning, management & decommissioning phases that support transparent knowledge tracking, feedback loops, prediction & simulation. • Enabling a clear separation of concerns between different stakeholders, e.g. business versus IT, developers and customers. • Supporting the full variety of business scenarios, from traditional data centres to complex service value networks, while complying with overarching sustainability requirements.
3. Advanced infrastructure technologies • Hardware: energy efficient, flexible allocation, virtualization. • Middleware: new multi-tier system design, flexible storage systems, harmonized virtualization on all layers. • Related programming models (parallel programming, multi-core) that meet the required flexibility of the networked economy.

Conclusions

The development of a major business infrastructure designed for future service-oriented economies represents a shift towards the functional view of data transformation services based on the assumption that innovation in businesses and public services drive the evolution of these services down to the technology level. The community-driven NESSI-Grid SRA exemplifies the general Grid vision in a concrete context based on real business needs, drivers, and challenges.

The Research Challenges for Business Grids within the future SOE detail and prioritize specific requirements that focus on real concerns for the commercial sector. The overall aim is to ensure that the infrastructure is developed in compliance with real requirements, standards and ICT policies through industry-research collaborations. The expected impact of Business Grids as a service-oriented infrastructure is to bring tangible benefits for European enterprises in the evolving IT & business landscape.