The Semantic Grid: Myth Busting and Bridge Building

Carole Goble† and **David De Roure**‡

Abstract. The Semantic Grid is an extension of the current Grid in which information and services are given well-defined meaning, better enabling computers and people to work in cooperation. The full richness of the Grid ambition depends upon realizing the Semantic Grid, but it is still, to many, a mysterious hybrid of the Semantic Web and the Grid, both of which are subject to myths and misunderstandings. In this paper we explain the changing landscape of the Grid, and describe the bridge-building and myth busting needed to achieve the Semantic Grid.

1 INTRODUCTION

In 2001 a group of researchers recognized that the vision of the Grid as "coordinated resource sharing and problem solving in dynamic, multi-institutional virtual organizations" [1] is closely related to that of the Semantic Web – which is also, fundamentally, about joining things up. The value of applying Semantic Web technologies to the information and knowledge in Grid applications was immediately apparent. The group anticipated the new service-oriented Grid and saw the potential of the Semantic Web in working with and for service-based Grid middleware[2]. Semantic Web technologies would not only be useful managing the knowledge resources of the applications *using* the Grid, but also *within* it [3]

In May 2004, Google gives 4140 hits for "Semantic Grid" and 10,300 for "Knowledge Grid". A web survey reveals Semantic Grid projects in every continent; tens of workshops, conferences and journal special issues; industry "briefing days"; and government funding agency programmes. The European Commission's FP6 Complex Problem Solving on the Grid programme lists at least five significant projects with "semantic" components. Global Grid Forum[1], the heartland of mainstream Grid Computing, casually refers to ontologies, and has chartered a Research Group on Semantic Grid.

The Semantic Grid vision brings a challenging research agenda and a promise of immediate practical benefit through deployment of available Semantic Web technologies in current Grid applications However, the Semantic Grid is still, to many, a mysterious hybrid of the Semantic Web and the Grid, both of which are also subject to myths and misunderstandings.

Ultimately, the Semantic Grid, or more accurately a "Grid with semantics", is an exercise in bridge building: across gaps in technologies; gaps between the Grid vision and current reality; and the chasm between the Grid and Knowledge communities and cultures. For community practitioners outside the Semantic Web and within the Grid it is important to understand what can be achieved immediately and what is a research issue. For practitioners outside the Grid and within the Semantic Web it is important to realize, as well as understanding the problems that the Grid is trying to address, that Grid technologies have direct practical application to their own development of scalable, distributed, secure and robust knowledge technologies for the Web.

There are bridges under construction to move towards the vision – with established communities such as Knowledge Management and Agent Based Computing and also newer communities such as Semantic Web Services and Ubiquitous Computing. We are seeing increasing interest in Semantic Grid from the broadening range of disciplines turning their attention to Grid computing, notably Arts, Humanities and Social Sciences. But there are issues to address.

- **The communities to bridge.** Forrester recently surveyed 149 large companies in North America to learn about their knowledge and use of grid technologies [4]. Firms are confused about what "Grid" means, and they are not alone. This confusion is unhelpful. We need to bust some myths by clarifying the Grid vision, current Grid practice, and the state of the art in semantic technologies. The ground on both sides of the bridge is unstable. In particular the Grid is undergoing a significant change which should aid the absorption of semantic technologies.

- **The bridge.** What is the Semantic Grid? Where does it fit in the roadmap of the Grid? Put simply, the Semantic Grid is an extension of the current Grid in which information and services are given well-defined meaning, better enabling computers and people to work in cooperation. It directly enables the next generation Grid infrastructure as well as the use of higher-level knowledge in Grid applications, a layer characterized as the Knowledge Grid.

- **Two-way travel.** To fulfill its vision Grid needs the "intelligent" technologies that the Semantic Web and larger A.I. community has to offer. For the Semantic Web community to develop large scale, robust, distributed solutions, as opposed to interesting prototypes, they should look to the Grid's portfolio of software specifications, middleware components and practical deployed applications.

† Information Management Group, Department of Computer Science, University of Manchester, Oxford Road, Manchester M13 9PL, UK. carole@cs.man.ac.uk

‡ Grid and Pervasive Computing, School of Electronics and Computer Science, University of Southampton, Southampton SO17 1BJ, UK. dder@ecs.soton.ac.uk

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

- Building from both ends: Semantic Grids need commitment from both communities. But there is distrust and suspicion on both sides. The Semantic Web (and larger A.I.) community commonly see the Grid one as scruffy engineers, drawn from the application communities so not even real computer scientists, who will deliver ungainly solutions just to make their applications work. Many Grid workers seem to see large parts of the knowledge community as a bunch of "crazy aesthetics-lovers," good at talk and theory, but with no practical use; more interested in prototyping (complicated) toys than building real (compromised?) solutions. In Grid, a running, usable implementation is the badge of credibility. The "A.I. Winter" still penetrates to this day. Even though there are friendly overtures on both sides, these deep-rooted tendencies are there and can get in the way of fruitful collaboration.

This paper addresses these points, arguing that the Next Generation Grids will need to be semantic and that the AI community can benefit from Grid computing, and finally we identify some challenges to overcome in our bridge building.

2 THE GRID AND ITS MYTHOLOGY

Analogous to an electricity power grid, Grid computing views computing, storage, data sets, expensive scientific instruments and so on as utilities to be delivered over the Internet seamlessly, transparently and dynamically as and when needed, by the virtualization of these resources [1]. It was originally conceived as a means of securely and transparently sharing (super)computing resources on demand using high bandwidth networks – a view now caricatured as 'big iron and fat pipes'. The application pull arose from global large-scale scientific collaborations that transcended organisational and geographic boundaries, and the need to make the best use of expensive or scarce resources as and when needed.

Large-scale science is achieved through the interaction of people, applications, computing resources, information systems, and instruments, despite the heterogeneity of their respective policies, platforms and technologies, and their geographical and organizational dispersal.

Scientists cluster together "virtually" to solve problems. So Grid computing is about enabling unanticipated *collaboration* between resources to achieve something that was not possible before. It is about the temporary *pooling* of resources to solve a problem, for example harnessing compute cycles on thousands of machines to enable high throughput screening of 10,000's of drug compounds in an hour instead of a year, as in the SmallPox Grid[2]. It is about the *democratization of science* so that neuroscientists in San Diego, USA can remotely collect data from an ultra high voltage electron microscopy in Osaka, Japan [5], and astronomers steer remote telescopes from their offices. The Sloan Sky Observatory[3] and AstroGrid[4] take available vast digital sky surveys to all astronomers not just a lucky few, and provide the computing, storage and software resources needed to analyse the data.

As the Grid's motivation is to facilitate the routine interactions of resources, since the mid-1990s its vision has evolved to now mean middleware for "**flexible, secure, coordinated resource sharing among dynamic collections of individuals, institutions, and resources - what we refer to as virtual organizations (VO)"** [1].

The resource configurations are *transient*, *dynamic and volatile* as a consortium of services (databases, sensors, compute servers) participating in a complex analysis may be switched in and out as they become, or cease to be, available. They are anticipated to be *ad-hoc* as service consortia have no central location, no central control, and no existing trust relationships. The configurations and the services that make them up may be large, with potentially hundreds of services orchestrated at any time, and *long-lived*, for example a protein folding simulation could take weeks. The *scale* of data (petabytes) and compute resources (tens of thousands) are large and the quality of service and performance criteria are severe. The platform must be scalable, be able to evolve and hence be future proof, and be fault-tolerant, robust, persistent and reliable. It should work *seamlessly*, and *transparently* – the user won't necessarily know or care where their calculation is done using how many machines, or where data is actually held.

The underlying principles of the Grid are *virtualization* and *dynamic provisioning*. For example, just as an Internet user views a unified instance of content via the Web, a compute grid user essentially sees a single, large virtual computer. In implementing "one image from many resources", Grid should support: virtualization at every layer of the computing stack; provisioning of work and resources based on policies and dynamic requirements and the pooling of resources to increase utilization. To manage "many resources as one resource" requires self-adaptation for tuning and failover, and unified management and provisioning [6]. The reality of this is truthfully some way off. Current Grid infrastructure is fully exposed to programmers and is unreliable enough to require continual nursing by administrators; the software is difficult to use, deploy and debug. The next generation of Grids must be simplified, to lower the barrier of take-up, and provide appropriate abstraction mechanisms controlled by metadata and enabled by agent-based computing [7].

2.1 Meanings and Myths

The term "Grid" is used to mean many things. To avoid confusion, the UK funding agency named their programme "e-Science" - the use of the Grid, the Web, Web Services, the Semantic Web and other software infrastructure to benefit Science.

To clarify the situation, here are some definitions: **"The" Grid** refers to the concept and the vision, and as such does not exist as an artifact. **"A" Grid** is a particular virtual organisation of heterogeneous distributed resources that will collaborate to solve a problem. A VO of machines linked, by a high performance network, to form a virtual machine, such as the TeraGrid[5], is a Computational Grid. Other VOs reflect resources (a Data Grid forms a virtual database), geography (National Grid Service of UK resources), fields (Mouse Genome Grid), disciplines (BioGrid), or problems (protein folding simulation). There are *many* interoperating Grids. **Production Grids** are Grids of resources used routinely, commonly using previous generation middleware. They currently solve hetereogeneity and dynamic resource problems by enforcing rigorous conditions on participation, and using hardwired middleware. **Grid middleware infrastructure** is the software services stack, policies, protocols, standards and APIs that make Grids. Reference implementations include Condor[6], Globus[7] and Unicore[8]. **Grid tools**, such as resource heartbeat monitors and portals, enable the management and use the Grid infrastructure. Finally, **Grid applications** use a Grid by means of its middleware to solve a problem, often a scientific one. Now let's bust some Grid myths.

[2] http://www.grid.org/projects/smallpox/about.htm
[3] http://www.sdss.org
[4] http://www.astrogrid.org
[5] http://www.teragrid.org
[6] http://www.cs.wisc.edu/condor/
[7] http://www.globus.org
[8] http://www.unicore.org

- **The Physics-only Myth** True, most production Grids are computational Grids for Physics, like the EGEE[9] in preparation for the Large Hadron Collider experiment. The Particle Physicists typically have a single distributed team working on a single experiment that was planned from the outset. However, Life Sciences will be the science to dominate the Grid with its large numbers of collaborating teams, large scale data integration, and series of experiments that were never planned to work together. Grid activity cuts across an increasing number of similarly characterized disciplines earth sciences, astronomy, e-Social Science, arts and humanities. The Grid will create new disciplines, where people come together in ways that were not possible before.

- **The Academics-only Myth.** All the commercial vendors are making serious investment. IBM DB2 and Oracle 10g will be Grid-compliant. Forrester reports [4] that of their surveyed companies 37% reported that they're piloting, rolling out, or have implemented some form of grid technology and another 30% are considering grid technology. The Grid is a major initiative, with buy-in from industry as well as nations, funding agencies, and researchers. Commercial deployments are offered from Avaki, Fujitsu, HP, IBM, Oracle, Platform, Sun Microsystems, United Devices and so on.

- **The Computational Grid Myth** The Grid is proposed as a *generic* mechanism for forming, managing and disbanding dynamic federations of services. Application integration and collaboration are the keys. Data-intensive problem solving is the "killer application" of Grid computing. Scientific and non-scientific applications have vast amounts of difficult data of different kinds to be integrated, aggregated, analyzed and archived. *Information Grids* provide data access and manipulation mechanisms that can take account of the structure or content (semantics) of the data. The scale, complexity, and fluid behaviour of the Grid and its applications challenge databases: multi-database query processing, schema integration and evolution, data transport, distributed transaction management, data mining and hybrid processing of semi-structured and structured data, change notification, archiving and security. Some Grid solutions use virtual "wrap & mediate" databases to integrate data-generating resources; most orchestrate services using workflows. Workflow discovery, composition, data streaming for large volumes of data, suspension and resumption of workflows and failure compensation is an essential high level Grid service. Another is the accurate logging and exploitation of provenance (who, what, when, how, why) metadata. To support VOs, Grids also provide support for *collaboration* – both bringing collaborators together and supporting their successful engagement.

- **The No-Architecture Myth** In 2002 first generation Grid middleware suffered from: no standards (the Globus toolkit was de facto); monolithic design making it difficult to integrate existing components and little synergy with existing middleware standards and tools. Not anymore. The second generation moves from a bag of protocols to a service-oriented architecture. The Open Grid Service Architecture (OGSA) [9] gives a middleware stack built on an object model specified in terms of industry standard web services, making the middleware more open and extensible, and more palatable to commercial concerns. The Web Service Resource Framework [10] (the successor to the Open Grid Service Infrastructure) defines a distributed object architecture over Web Services. It specifies abstract descriptions of Grid capabilities and standard service interfaces for creation, soft-state management, lifetime management, introspection, and grouping, separating resources from services to aid virtualisation. Along with essential low level "plumbing" such as GridFTP, a whole middleware stack of higher-level services are currently being specified and built: security, registries, policy management, data access and integration, service management, and workflow. Figure 1 gives a simplified OGSA Architecture.

[9] http://public.eu-egee.org

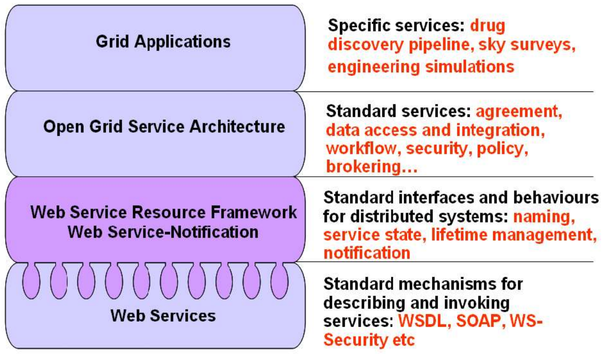

Fig. 1 The Open Grid Services Architecture., adapted from [8]

3 THE SEMANTIC GRID

Metadata pervades the Grid middleware stack; describing the environment, the services available and the ways they can be combined and exploited, keeping audit trails of resource use and configurations and so on. The operation of the Grid middleware itself is dependent on the capturing and harnessing of metadata, as well as the metadata required for the Grid applications themselves. Coupled with the move to a service oriented architecture is a move to the explicit representation, sharing and linking of metadata [3].

The Semantic Web as defined by the W3C is "…an extension of the current Web in which information is given well-defined meaning, better enabling computers and people to work in cooperation. It is the idea of having data on the Web defined and linked in a way that it can be used for more effective discovery, automation, integration, and reuse across various applications […] where data can be shared and processed by automated tools as well as by people." [11]. The ambition is of an environment where software agents are able to dynamically discover, interrogate and interoperate resources, building and disbanding virtual problem solving environments discovering new facts, and performing sophisticated tasks on behalf of humans. This is not so far from the Grid ambition. To achieve this, work has been undertaken to assert the meaning of Web resources in a common data model (RDF) using consensually agreed ontologies expressed in a common language (OWL) so we can share the metadata and add in background knowledge. From this basis we should be able to query, filter, integrate and aggregate the metadata, reason over it to infer more metadata using rules, and attribute trust to the metadata.

Efforts have concentrated on the languages and software infrastructure needed for the metadata and ontologies, and these technologies are ready to be adopted. But we are some ways off the full "distributed agents" vision. Many commentators advocate that the power of the semantic web's techniques and technologies will work best in closed controlled domains (like Enterprises or Grids) that need large scale metadata integration and querying (like Grids) where a little semantics will go a long way[10] (which may well be the case with Grids). Simple metadata and simple queries give a small but not insignificant improvement in information integration.

Both the Grid and the Semantic Web are fundamentally exercises in **interoperability**. By analogy with Semantic Web, the Semantic Grid is an extension of the current Grid in which information and services are given well-defined meaning, better enabling computers and people to work in cooperation. Table 1 gives a comparison of character of the Grid and the Semantic Web. The Semantic Grid is a recognition of the potential of their partnership.

Table 1. The Grid *vs* The Semantic Web.

The Grid	The Semantic Web
On demand transparently constructed multi-organisational federations of distributed services	An automatically processable, machine understandable web
Distributed computing middleware	Distributed knowledge and information management
Programmatic integration, originally based on protocols & toolkits	Information integration, based on metadata, ontologies and reasoning
Information & compute power *as* a utility	Information & knowledge *is* the new utility
Application pull: pioneers are application scientists with large scale collaboration problems, originally computationally-oriented.	Technology push: pioneers are primarily from the knowledge, agent and A.I. communities.

As higher-level knowledge plays an important role in the future Grid applications, issues related to the representation, discovery, and integration of domain knowledge, data mining, knowledge discovery, text and multimedia content analysis, semantic information extraction and integration, etc are relevant. We distinguish between a **Knowledge Grid** which is a *Grid of semantics* based on knowledge generated by applications using and mining the Grid and a **Semantic Grid** which manages semantics for the Grid to manage and execute its architectural components [3, 12] (Fig. 2).

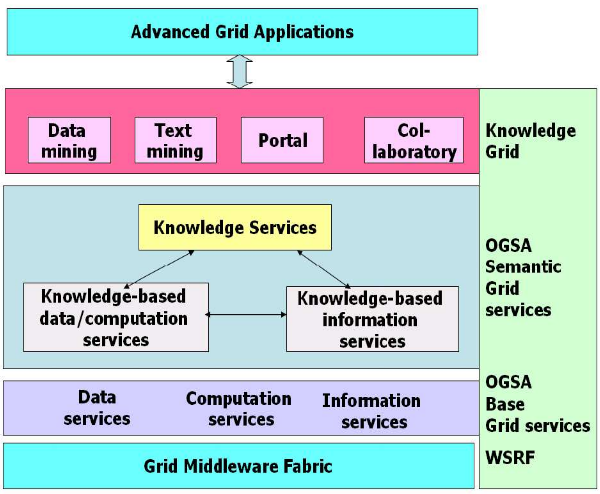

Fig. 2 The Semantic/Knowledge Grid Stack

[10] The Hendler Principle

3.1 A Bridge from Grid to Semantics

There is a big gap between where the Grid is today and the Grid vision of a high degree of easy-to-use and seamless automation and in which there are flexible collaborations and computations on a global scale. Current solutions are commonly bespoke; the middleware is difficult to install, manage and update; configuration management and fail-over are not straightforward; even acquiring authorization certificates is cumbersome. The Next Generation Grids will need to allow for much easier life cycle management and smooth evolution; subsidiarity of control of services; self-organisation to minimize failure; transparency and straightforward trouble-free administration [6]. The levels of abstraction for programming the Grid are currently wrong, and we need standards for common information models for cross domain and cross Grid terminology. It is difficult to reuse Grid services because so much of their semantics is implicit and embedded. Bluntly, the Grid is too low-level, too in the plumbing, and the OGSA high level services are missing. The Grid community is aware of this and trying to fix it. There is, for example, an activity to develop a Common Management Model for Grid Resources.

To manage a Grid beyond hardwired "hackery" requires "Grid intelligence" to interpret extensive quantities of (dynamic) knowledge about the state and properties of Grid components and their configurations for solving problems [3]. The dynamic discovery, formation, and disbanding of ad hoc virtual organizations of (third-party) resources requires that Grid middleware be able to use and process knowledge about the availability of services; their purpose; the way they can be combined and configured or substituted; and how they are discovered, are invoked, and evolve. Figure 3 gives some examples. Here we describe a few.

- **Knowledge management** A Semantic Grid depends upon making knowledge *explicit* and *machine processable* to be used in decision making, all the way down. The Grid Interoperability Project used ontologies deep in the heart of Grid middleware to support mediation between the resource brokers of Unicore and Globus [13]. myGrid uses RDF to represent and integrate the provenance records of workflows [14]. The GriPhyN projects used RDF-based ontologies and reasoning to improve resource matchmaking [15]. GEON and SEEK use RDF/OWL for semantic data integration and mediation [16]. Geodise uses past community knowledge to drive problem solving selection to advise engineers on the configuration and monitoring of long-running engineering simulations [17]. Other knowledge opportunities include: policies, performance metrics, job control; event notification topics, provenance trails; role based security and access rights, personal profiles and security groupings; charging infrastructure; quality of service metrics; intelligent portals, debugging, monitoring etc. WSRF specifies a state model for Grid services but the role of semantics and state is yet to be explored in this context.

- **Semantic Grid Services** Semantics is the key to negotiation, discovery and workflow composition to match policy models in a service-oriented architecture. The Semantic Grid leverages work on Semantic Web Services to provide declarative specification of services and workflows and their requirements, the classification of computational and data resources and the typing of service inputs and outputs. Geodise and myGrid use RDF/OWL to discovery services and workflows [17, 18].

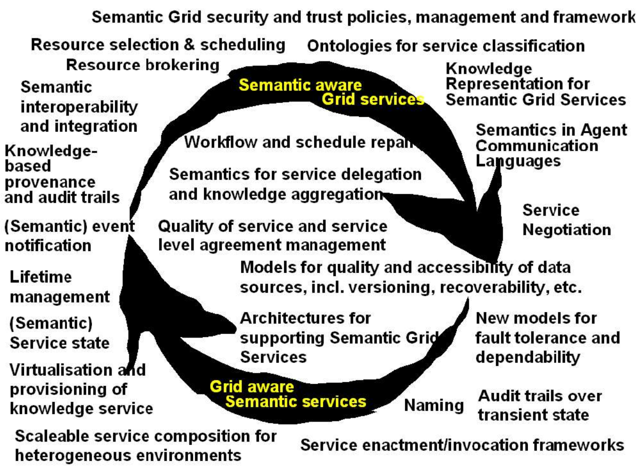

Fig. 3 Opportunities and active areas of mutual benefit

- **Multi-agent systems** A Semantic Grid is characterized as an open system in which users, software components and computational resources (all owned by different stakeholders) come and go on a continual basis. The services are offered under contract, and can be accepted by any number of consumers in a marketplace to form agreements and coalitions. The exploitation of techniques and methodologies from intelligent software agents seems obvious, including peer-to-peer communication, dynamic decision-making, decentralization, coordination, and autonomous / autonomic behaviour [19]. A Grid does not need to be based on an agent framework to profit from notions of agency and the fruits of extensive research in this community – agent architectures address the same problem of flexible action (balancing proactive and reactive behaviour) in dynamic and uncertain environments. The CoABS Grid[11]enables the dynamic interoperability of distributed, heterogeneous, objects, services, and multi-agent systems. Effective resource discovery, allocation and job scheduling to meet stringent quality of service criteria can benefit from economic principles developed within agent systems to influence the GGF Grid Economic Scheduling Architecture. Agent based software engineering [20] is immediately applicable to grid computing, proposing strategies of decomposition, abstraction and organization, drawing on the agents paradigm to gain advantage in using these strategies. By focusing on the objectives of agents and on emergent behaviour, not all interactions need to be specified in advance.

- **Planning and Scheduling.** Grids are still the domain of a few highly trained programmers. Capturing knowledge and heuristics about resource selection and using that to generate automatically executable job workflows for a grid can use A.I. planning techniques. Resource brokering, resource scheduling and workflow / schedule repair are ripe for planning [21].

3.2 A Bridge from Semantics to Grid

For Grid and Semantic Web solutions, knowledge services and data sources are likely to be distributed, heterogeneous, and dynamic. We need to go beyond centralized knowledge service provision, and develop effective open, distributed, knowledge-based solutions. The next generation of Grid Computing offers a practical architecture in which to deploy knowledge services. For knowledge services to become Grid services (rather than the other way about) they become **Grid-compliant** by adopting the lifetime, notification and state management characteristics of virtualized WSRF Grid services; and **Grid-aware**, adopting Grid approaches to security, large scale secure file transfer, data replication protocols, distributed query processing, workflow repair, provisioning, self-repair, data access, event notification, registration etc. Imagine a reasoner able to harness a compute grid. Imagine a distributed RDF store as a virtualized store. Gil reveals that deploying agents as Grid services using the state modeling interfaces greatly enhances their robustness and durability[12]. A great deal of effort is being spent on engineering high quality, reliable Grid infrastructure that it seems obvious the Semantic Web should leverage. This direction over the bridge is almost entirely unexplored.

4 BRIDGE BUILDING

The Grid community is international, interdisciplinary, vibrant and open. This is a genuinely application-driven community – to build a successful Grid solution requires commitment: the scientists who hold the application; the service providers who maintain the Grid fabric and the resources on it, and the software professionals who develop and configure the middleware (Fig. 4).

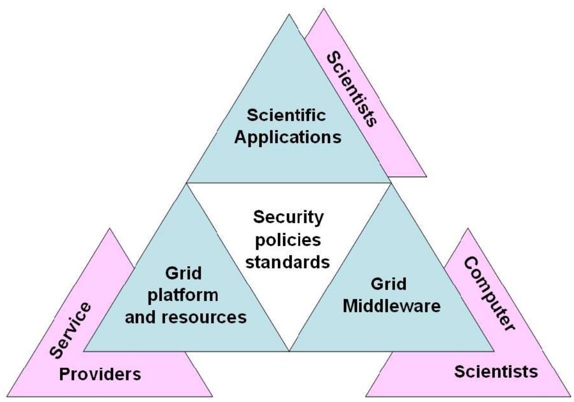

Fig 4. Stakeholders to be bridged for a Grid.

The active participation of application communities is important. Their range of engagement reflects the maturity of the Grid: Particle Physics needs robust and sophisticated production grids to support ambitious application scenarios, whereas Life Sciences look to using at least some grid technologies. The Grid community has done a good job to build these bridges, but it has been tough. The desire to build software that is interesting rather then useful is strong; the tendency to build middleware that service providers cannot deploy or is too insecure is tempting. At the moment Grid middleware is difficult to administer and difficult to use; it's the service providers that need the most convincing. On the other hand the early adopters are prepared to try out new things and are remarkably tolerant of untidiness and instability, so long as it moves them along. Can we do as well with the stakeholders of the Semantic Grid (Fig. 5)?

- **Overcoming divisions**. Grid computing is all about plumbing, connecting up the fat iron and big pipes; AI is more focused on understanding tomorrow's intelligent machines than delivering solutions today. These are typical suspicions,

[11] http://www.darpa.mil/ipto/programs/coabs/

[12] Personal communication, October 2003

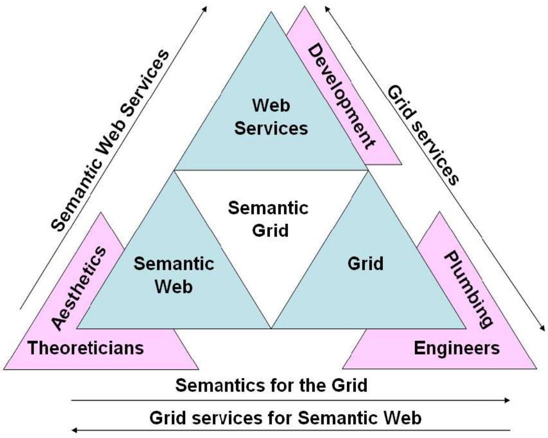

Fig 5. Stakeholders to be bridged for a Semantic Grid.

and people need to be persuaded to take another look to see the value. While the Grid community is a little suspicious of the "crazy A.I. people" of the Semantic Web, airing concerns regarding performance, pragmatism and effectiveness, they are prepared to be convinced. In fact it is not essential to buy into the full vision in order to derive benefits across the bridge – the first steps may be more selfishly motivated. For example, Semantic Web gives the Grid a standard ontology language for interchange, while the Grid gives the Semantic Web a middleware platform for enterprise data integration. Sadly these divisions are often directly reflected in the funding agencies, promoting a silo-mentality in the community rather than defeating it.

- **Growing pains.** Despite very substantial progress, neither Semantic Web nor Grid can really be said to be reality – and Semantic Grid involves both! Recall the three-stage project lifecycle: first the observers say 'it will never work' then 'it could be useful' and thirdly 'it was my idea all along.' We are in phase 2. Many of the benefits of Semantic Grid will be apparent where reuse is occurring over time and/or across large collaborations, and this will happen in the medium to long term. Meanwhile the Grid middleware is immature and unstable, and there are few production examples. We need proof that the Semantic Web can scale, and deliver the QoS and performance that Grid demands. Unfortunately, deployment, research, development, applications, standardisation and commercialisation are all happening at once for both communities. We are experiencing standards swirl, with multiple standards and standards bodies. This maelstrom makes it difficult to build a bridge which is solid and remains so.

- **Making it easier, not harder.** Both the Grid and Semantic Web middleware stacks are pretty complex (in part due to the growing pains we described above). Putting them together runs the risk of a baroque architecture that no-one will understand let alone use. We must strive for simplification. Simplification leads to repeatability - and we want developers and applications to be able to easily pick up and repeat the techniques and technologies of the Semantic Grid.

- **The network effect.** The magic behind the World Wide Web is the added value of the 'network effect' gained by allowing any page to point to any other page, allowing any user to link to, or annotate, any statements made on the Web by another. The scaling of the Semantic Web depends on a similar network effect being reached in "information space"-allowing the sharing and linking of machine readable content, and gaining power by linking to, extending, or even disagreeing with that specified in another Semantic Web document [22]. The diverse, distributed metadata created by various users about the same entities is effectively interlinked by the objects it describes. This, for example, enables us to ask new kinds of questions that draw on the aggregated knowledge. Semantic Grid applications need to adopt shared, naming schemes and agreed metadata schema. Although this can be a challenge in the general case, this is achievable on the scale of a project or a 'closed' community of users – for example, the Life Sciences Identifier [23].

- **Return on investment** The benefits of using Semantics in Grid come when there is heterogeneity and reuse involved; heterogeneity in the data, resources, cross grid and cross organisation – and across time. It enables resources to be reused in ways which were not anticipated when they were first created. Hence it provides a potential saving for organizations, and the promise of a competitive edge. Whether this is effective, and worthy of investment, will depend on the business model - why invest extra now for longer term benefit if that stakeholder will not benefit directly? This also means that one small semantic grid exercise will not show the full benefits to the organisation – it requires an overarching strategy.

- **Applications needed.** The most effective way of overcoming suspicions, demonstrating the potential for Return of Investment, and overcoming growing pains, is surely to provide the right exemplar applications – beacons located mid-bridge to attract the attention of passers-by at each end. The nature of these applications is that they should score easy benefit from Semantic Grid technologies, which means that they must exhibit scale and heterogeneity but at the same time enjoy circumstances where community-based agreement mechanisms are already in place. Hence it is no surprise that we see our 'beacon applications' in communities such as Life Sciences.

5. CONCLUSION

Established communities sometimes don't take new initiatives seriously, wearily dismissing them as another re-run of old news; patronising them if they fly in the face of orthodoxy (the dismissal of the Web by the Hypertext community is a famous case in point); or fearing them as a diversion or a potential threat to the "purity" of their core subject. It is easy to get into a rut.

The new community may not recognize their problems as having old solutions, albeit in disguised clothes, or may not want to believe that approaches already exist. Research findings (in particular failures) of ten years previous, particularly from outside the immediate field of expertise, are those that are doomed to be rediscovered over and over again. Established communities hold a corporate memory of what doesn't work as much as what does.

For researchers with interests in agents, ontologies, scheduling, search, knowledge representation, reasoning, data mining and autonomic systems – to name just a few topics – the Grid evolution (revolution?) provides an exciting opportunity for research. It's a challenging application that can immediately benefit from past experience and current expertise. It's a sand-box with real users for large scale experimentation on extensive grid deployments. The Semantic Grid is a particularly exciting prospect. The Grid community may be a little suspicious of parts of the AI community but is prepared to be convinced. The A.I. and knowledge communities embraced the Semantic Web with zeal; many of these same pioneers see the Grid as an application of Semantic Web approaches.

To get intercommunity activities going we need (a) between-community travelers who will spread the news, and transfer technologies and broker partnerships – in other words, building bridges – and (b) early pioneers, open to new opportunities and different ways of doing things, prepared to cross them.

In this paper we have defined the Semantic Grid: an extension of the current Grid in which information and services are given well-defined meaning, better enabling computers and people to work in cooperation. We have shown that the full richness of the Grid ambition depends upon realizing the Semantic Grid, and indeed the Semantic Web ambition benefits too. We explored the art and craft of bridge construction between unstable lands. Success will be achieved when the bridge becomes invisible because the Semantic Grid is just part of the infrastructure for Grids and the Semantic Web.

For further information on the Semantic Grid, see http://www.semanticgrid.org

ACKNOWLEDGEMENTS

We should like to acknowledge our colleagues Sean Bechhofer, Marije Geldof, Yolanda Gil, Frank van Harmelen, Jim Hendler, Nick Jennings, Bertram Ludaescher, Luc Moreau, Norman Paton, Nigel Shadbolt and Chris Wroe for their insightful discussions on Semantic Grids.

REFERENCES

1. Foster I. and Kesselman C., *The Grid: Blueprint for a New Computing Infrastructure*, Morgan Kaufmann, 1999.
2. De Roure, D., Jennings, N.R. and Shadbolt, N.R. *Research Agenda for the Semantic Grid: A Future e-Science Infrastructure*, UK e-Science Technical Report Series UKeS-2002-02,, National e-Science Centre, Edinburgh, UK, 2001.
3. Goble C.A., De Roure D., Shadbolt N.R. and Fernandes A.A.A. *Enhancing Services and Applications with Knowledge and Semantics* in The Grid: Blueprint for a New Computing Infrastructure Second Edition (eds. I Foster and C Kesselman), Morgan Kaufman 2004.
4. Gillett FE, *Grid Gets Big, But The Term Is Confusing*, Forrester, 180504 http://www.forrester.com/Research/ 2004.
5. Takaoka A., Yoshida K., Mori H., Hayashi S., Young S.J., Ellisman M.H. *International telemicroscopy with a 3 MV ultrahigh voltage electron microscope* Ultramicroscopy. May;83(1-2):93-101 2000.
6. IBM Grid Computing http://www-1.ibm.com/grid/ May 2004
7. Next Generation Grid(s) European Grid Research 2005 – 2010, EU Expert Group Report, June 2003
8. Parastatidis S., Webber J., Watson P. and Rischbeck T.A *Grid Application Framework based on Web Services Specifications and Practices*, CS-TR: 825, School of Computing Science, University of Newcastle, Jan 2004
9. Foster I, Kesselman C, Nick J, Tuecke S. *Grid Services for Distributed System Integration*. Computer, 35(6), 2002.
10. *The WS-Resource Framework*, 180504 http://www.globus.org/wsrf/ 2004.
11. Goble C.A. and De Roure D. *The Grid: an Application of the Semantic Web*. ACM SIGMOD Record 31(4): 65-70. 2002
12. Cannataro, M. and Talia, D. *Semantics and Knowledge Grids: Building the Next-Generation Grid.* IEEE Intelligent Systems, 19 (1). 56-63 2004.
13. Brooke J., Fellows D., Garwood K. and Goble C.A., *Semantic matching of Grid Resource Descriptions* in Proceedings of The Second European Across Grids Conference, Nicosia, Cyprus, 2004
14. Stevens R., Tipney H.J., Wroe C., Oinn T., Senger M., Lord P., Goble C.A., Brass A. and Tassabehji M. *Exploring Williams-Beuren Syndrome Using myGrid* to appear in Proc of 12th ISMB, 31st Jul-4th Aug 2004, Glasgow, UK 2004.
15. Tangmunarunkit H., Decker S. and Kesselman C., *Ontology-based Resource Matching* in Proceedings 2nd International Semantic Web Conference (ISWC2003), Florida, USA, Oct, 2003.
16. Ludaescher B., Gupta A., and Martone M.E., *A Model-Based Mediator System for Scientific Data Management*, in T. Critchlow and Z. Lacroix (eds), Bioinformatics: Managing Scientific Data Morgan Kaufmann, 2003.
17. Chen L., Shadbolt N.R., Goble C.A., Tao F., Cox S.J., Puleston C., Smart P.R. *Towards a Knowledge-based Approach to Semantic Service Composition* in 2nd International Semantic Web Conference, 2003 (ISWC2003), Florida, USA, October 2003.
18. Wroe C., Goble C.A., Greenwood M., Lord P., Miles S., Papay J., Payne T., Moreau L. *Automating Experiments Using Semantic Data on a Bioinformatics Grid* in IEEE Intelligent Systems special issue on e-Science Jan/Feb 2004.
19. Moreau L. *Agents for the Grid: A Comparison with Web Services (Part 1: the transport layer)*. In Henri E. Bal, Klaus-Peter Lohr, and Alexander Reinefeld, editors, *Second IEEE/ACM International Symposium on Cluster Computing and the Grid (CCGRID 2002)*, pages 220-228, Berlin, Germany, May 2002. IEEE Computer Society.
20. Jennings N.R. *On Agent-Based Software Engineering* Artificial Intelligence, 117 (2) 277-296 2000.
21. Blythe J., Deelman E., Gil Y., Kesselman C., Agarwal A., Mehta G., Vahi K. *The Role of Planning in Grid Computing*, In Proceedings of the 13th International Conference on Automated Planning and Scheduling (ICAPS), June 9-13, 2003, Trento, Italy, 2003.
22. De Roure, D. and Hendler, J.A. *E-Science: the Grid and the Semantic Web*. IEEE Intelligent Systems, 19(1), 65-71 2004.
23. Clark T. Liefeld T., and Martin S. *Globally Distributed Object Identification for Biological Knowledgebases.* Briefings in Bioinformatics, 5(1):59–70 2004.