

09131 Abstracts Collection
Service Level Agreements in Grids
— Dagstuhl Seminar —

Hans Michael Gerndt¹, Omer F. Rana², Wolfgang Ziegler³ and Gregor von Laszewski⁴

¹ TU München, D

gerndt@in.tum.de

² Cardiff University, GB

³ Fraunhofer Institut - St. Augustin, D

wolfgang.ziegler@scai.fraunhofer.de

⁴ Argonne Nat. Lab., USA

laszewski@gmail.com

Abstract. From 22.03. to 27.03.09, the Dagstuhl Seminar 09131 “Service Level Agreements in Grids ” was held in Schloss Dagstuhl – Leibniz Center for Informatics. During the seminar, several participants presented their current research, and ongoing work and open problems were discussed. Abstracts of the presentations given during the seminar as well as abstracts of seminar results and ideas are put together in this paper. The first section describes the seminar topics and goals in general. Links to extended abstracts or full papers are provided, if available.

Keywords. Service Level Agreement, Grid Computing, SLA Negotiation, SLA Policies, SLA Implementations

09131 Executive Summary – Service Level Agreements in Grids

Grid computing allows virtual organizations to share resources across administrative domains. In its early days, Grid computing was inspired by the need for transparent access to supercomputing resources and by the idea to even couple the resources in a metacomputing environment to create even more powerful computational resources. Currently the focus is on service-oriented architectures (SOA) where a wide variety of services from multiple administrative domains can be accessed by service clients.

One of the most important tasks of current Grid middleware centers on efficient resource management. Resource providers offer their resource to virtual organizations and publish detailed information about the resources. Recent efforts have also focused on exposing computational and data resources as "services" - thereby providing a single abstraction that could be applied at different levels of

software deployment. Based on this information appropriate resources for Grid applications are selected, and jobs are finally submitted to these resources.

Service Level Agreements (SLA) are attracting more and more attention in Grids as a means to guarantee quality of service terms for grid applications and to enable the establishment of novel business models. A wide range of research and development questions have to be addressed in this context. This covers the creation of languages for formulating SLAs that are powerful enough to express the relevant QoS terms, but can also be used to automatically manage the negotiation, execution, and monitoring of SLAs. Brokering systems are required that can select resources for job execution based on the SLA templates offered by the resource owners. Scheduling algorithms that can optimize for different goals in the context of multi-item, multi-attribute, and multi-unit optimization problems are also necessary. Flexible local resource management algorithms are required for provisioning the resources at the provider's side to meet signed SLAs.

The seminar brought together people working on SLAs in the context of grid computing mainly from computer science, but also from information systems and application areas. These researchers come from different areas and bring in a wide range of research work.

Keywords: Service Level Agreement, Grid Computing, SLA Negotiation, SLA Policies, SLA Implementations

Joint work of: Gerndt, Hans Michael; Rana, Omer F.; von Laszewski, Gregor; Ziegler, Wolfgang

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2009/2026>

SLA Management in AssessGrid

Dominic Battré (TU Berlin, DE)

The fundamental goal of the AssessGrid project is to integrate risk awareness into all layers of the grid; provider layer, broker layer, and end-user layer.

In order to provide quality of service guarantees, Service Level Agreements play a crucial role. A central outcome of the AssessGrid project is therefore the Negotiation Manager component, an implementation of the WS-Agreement specification published by the Open Grid Forum. This work presents the architecture and major features of this Negotiation Manager and how these are used to provide quality of service guarantees.

Keywords: SLA, risk assessment, Grid, QoS

Joint work of: Battré, Dominic; Hovestadt, Matthias; Kao, Odej

Towards Negotiation Bootstrapping and Service Mediation

Ivona Brandic (TU Wien, AT)

Nowadays, novel computing paradigms as for example Grid or Cloud Computing are gaining more and more on importance. In case of Cloud Computing users pay for the usage of the computing power provided as a service. Beforehand they can negotiate specific functional and non-functional requirements relevant for the application execution. However, providing computing power as a service bears different research challenges. On the one hand dynamic, versatile, and adaptable services are required, which can cope with system failures and environmental changes. On the other hand, human interaction with the system should be minimized. In this talk we present the first results in establishing adaptable, versatile, and dynamic services considering negotiation bootstrapping and service mediation. We discuss novel meta-negotiation and SLA-mapping solutions for Grid/Cloud services bridging the gap between current QoS models and Grid/Cloud middleware and representing important prerequisites for the establishment of autonomic Grid/Cloud services. We present document models for the specification of meta-negotiations and SLA-mappings. Thereafter, we discuss the sample architecture for the management of meta-negotiations and SLA-mappings.

Negotiation and Management of SLAs in distributed dynamic environments

Frances Brazier (VU University Amsterdam, NL)

Management of large scale distributed autonomous systems in dynamic environments mandates decentralised solutions. Multi-level reflective architectures with local knowledge of dependencies between services at each level and between levels, provide a structure for division of responsibilities, roles, tasks, goals, and communication between systems. Service level agreements (SLAs) are acquired through interaction between autonomous systems: (1) through specific (mediated) negotiation; (2) through auctions, and/or (3) as part of participating in ad-hoc emergent configurations of systems. The type and amount of knowledge needed by each system and/or virtual organisation of systems, differ depending on the domain of application. The IIDS group is currently exploring negotiation and management of SLAs in data centre management, power management, energy management, crisis management, transaction management. The legal and social requirements are explored in close collaboration with experts in these fields.

VO- and Application-Centric Approaches to Service Level Agreement

Marian Bubak (AGH University of Science & Technology - Krakow, PL)

First, we present how the SLA management could be applied to operations of EGEE-like environments. The presented model was designed for the Polish NGI (PL-Grid, however some ideas are currently implemented in EGEE-3 project as a regional solution. The key concept of the model follows. The SLA is negotiated and agreed between VO and resource provider. Each party's decision is taken according to their own resource policy for a specified time frame, and therefore this is done manually, but an appropriate collaboration tool could make it efficient even in case of hundred-to-hundred connections order.

As a proof of the concept, we have deployed Bazaar, a tool which makes the process manageable and transparent, and implemented it in the Central European EGEE-III region. SLA-aware grid systems should enable measurements of QoS metrics defined in SLA, therefore these metrics should be carefully defined for each type of service. We propose metrics for computational and storage services. In a SLA-centric grid, the resources allocation should be limited to those with valid agreements. In other words, the SLA layer should filter out resources provided on the technical level, being able to notify about missing resources at the same time.

In the second part, we demonstrate a Grid application porting methodology using tools which act at a user- and application-level to improve the "user-centric" QoS metrics such as job turnaround time, stability and reliability of the job processing and efficiency of using of the Grid resources. In contrast to the system-level solutions, our approach does not require changes to existing Grid infrastructures or middleware services. However, the limitations of the underlying system apply: for example, the user-centric QoS may be only at best effort if the underlying Grid system does not provide mechanisms for strong system-level QoS. Similarly the access to the resources is influenced by the underlying fair-share mechanisms (or their lack).

An application-aware scheduler is overlaid on a pool of Grid worker nodes. The worker nodes run private pilot-jobs, or worker agents, which communicate with the master server created on demand at the end-user side. There is no strong software coupling between the scheduling, which is implemented in the Master/Worker DIANE framework, and the resource selection and acquisition which is implemented as sending pilot jobs to the Grid via Ganga interface. This allows to customize the system for various usage scenarios, including: high-performance computing (case-study: ITU RRC06 2006), high-throughput computing (case-study: LQCD 2008), interactive feedback (medical simulations), higher-order application logic (DAGs - astrophysics applications with SWIFT, workflows - medical imaging using Moteur and V-Browser, evolutionary algorithms - HEP event reconstruction for COMPASS and parameter-sweep applications).

Joint work of: Bubak, Marian; Moscicki, Jakub; Radecki, Marcin; Szepieniec, Tomasz

SLA Management in an Auction-Based Market Infrastructure

Simon Caton (Cardiff University, GB)

SORMA is a research project investigating the practical application of computational and other related resources in a market infrastructure. A key aspect of this approach is the introduction of potential consumers to fitting resource providers, which is facilitated through asynchronous auctions, where “best” fits are identified and introduced. In its capacity as an impartial third party, the SORMA market place acts as an overseer in the life cycle of the (Service Level Agreements) SLAs that are the product of this process. Here, the life cycle includes: SLA Generation from market matches, enforcement and violation-aware billing. We use WS-Agreement for the representation of an SLA, but have identified several weaknesses in the specification, specifically in the areas of accurately and pragmatically representing market-orientated data and SLA-adherence by its participants. Our solutions to circumvent these issues and how they facilitate more pragmatic market and SLA enforcement strategies is presented.

Keywords: Grid Markets, SLA Management, WS-Agreement, Violation Classification, Penalty Application

Joint work of: Caton, Simon; Borrisov, Nikolay; Rana, Omer

Service Discovery and the Role of SLA Templates

Donal Fellows (University of Manchester, GB)

The BREIN project, which is working to combine outputs from many previous projects to build business-oriented service Grids, uses Service Level Agreements (within overall master contractual contexts) to drive all its inter-organization communications. However, when a participant-space has large numbers of both clients and service providers, it can be difficult to discover the SLA Templates required for setting up these SLAs. This talk outlines a protocol for doing this discovery using best-effort SPARQL queries to produce an ordered list of templates, so permitting implementations that are efficient in terms of the number of messages exchanged and which do not allow the service provider to prevent the client from reaching a decision. It then discusses how this simple protocol can be extended to a whole architecture based on asynchronous distributed merge sorting which discourages malicious subversion of the discovery process through, e.g., cheating or slow answering by participants. This is done by allowing service providers to publish their SLA Templates by multiple simultaneous routes and by enabling a market in discovery services as well as in the services being discovered. This mechanism and architecture are independent of the rest of the negotiation process, services being discovered or term language, and depend on the fact that no commitment is formed until after the completion of the discovery phase.

Keywords: Service discovery sla template sparql architecture protocol

Pattern-based Service Orchestration

Cecilia Gomes (The New University of Lisbon, PT)

This presentation describes an approach based on patterns, and their manipulation through pattern operators, for structured aggregation and orchestration of services.

The approach is discussed in terms of its usage as a basis to define SLA templates which, in turn, can be refined further when specific QoS issues are associated with services.

Grid Accounting and Billing Infrastructure

Houssam Haitof (TU München, DE)

For a commercial entity to entrust the Grid for its business operations either as a consumer or a provider of resources, mechanisms that would guarantee its interests need to be implemented, especially service level agreements (SLA) supported by resource usage accounting and billing. Accounting is essential to SLAs, as it constitute the input to the SLA manager to assert and enforce SLAs as well as against plausible deniability in case something goes wrong. Whereas billing, being an accounting transaction, is typically a consequence to the SLA, and without it, SLA would be just a script language for requirement description or reporting.

We present here our experience with designing and building an accounting and billing system for the Grid in the context of the FinGrid project. Our architecture is based on a service oriented model, we relied on open standards and recommendations for our accounting system and on knowledge representation and reasoning to model our billing infrastructure.

Keywords: Grid Accounting, Grid Billing, Service Oriented Computing, Service Level Agreement, FinGrid

Dynamic Provisioning of Enterprise Systems

Markus Hedwig (Universität Freiburg, DE)

The ecological impact of the computing industry and the cost of operation of enterprise architectures have recently become a recognized problem. This research picks up the aspect of the operation of large enterprise architectures and derives a scheme for a dynamic hardware provisioning. The methodology aims to reduce the energy demand of the system while maintaining the original performance capabilities. This approach consists of three elements. The first part is the evaluation of the performance characteristics of an enterprise system. The next element analyses and forecasts the workload process. Based on this result, an

optimal hardware configuration can be selected in every situation. The different aspects of the model are evaluated with the help of synthetic and generic data. Afterwards the impact of the model is estimated

Usage of Virtualization for enforcing Service Level Agreements

Matthias Hovestadt (TU Berlin, DE)

Numerous research projects focused on SLAs, introducing SLA mechanisms to all levels of Grid computing. Customers and providers are enabled to make an explicit statement about all obligations and expectations in their business relationship. Having an SLA negotiated, customers can depend on getting their results in time and in the required quality. At this, fault tolerance mechanisms at provider side are mandatory for coping with resource outages. We will outline how resource virtualization can be used for increasing the provided level of fault tolerance.

From Service Markets to Service Economies - An infrastructure for protocol-generic SLA negotiations

Sebastian Hudert (Universität Bayreuth, DE)

Visions of 21st century's information systems show highly specialised digital services and resources, which interact continuously and with a global reach. For a broad adoption of this vision in a commercial context it is crucial to have a mechanism in place to guarantee quality of service, even across enterprise boundaries and to decentrally coordinate the involved resources.

Current service infrastructures try to tackle these problems by applying socio-economic mechanisms such as electronic negotiations and service level agreements stating functional and non functional guarantees of service invocations. Such technologies allow for the implementation of electronic service markets in analogy to real-world markets for everyday goods, which are currently investigated in many research projects.

Regarding the negotiation protocol applied in such a market Economic Theory claims that different market situations and negotiated products (i. e. the SLAs under negotiaion) demand different negotiation protocols in order to reach the highest-possible overall efficiency of the system. Thus we argue that the next generation distributed service-systems should not only be based on individual service markets, each applying one particular protocol, but on a global service economy where different protocols and thus service markets are present at any given point in time. We present a novel approach on such an infrastructure based on a structured negotiation protocol description language and a set of applied protocol primitives for agent-based SLA negotiations.

Keywords: SLA negotiations, negotiation protocol, software agents

Service Level Agreements in BREIN and beyond...why we need flexibility

Bastian Koller (Universität Stuttgart, DE)

Research in the area of Service Level Agreements is performed since several years, however the uptake by business has not been yet that good as it could have been. This is an implication of still existing and not addressed gaps in the overall SLA lifecycle, especially in the areas of SLA Creation and Management. Even though the efforts in refining and closing these gaps in SLA Management were increased in past, the creation of SLAs (how business partners find each other and how they come to an agreement) is still immature. The BREIN project addresses maturity by an approach where classical Grid/Web Services are combined/enhanced by Multiagent and Semantics concepts to overcome the limitations of current solutions and to allow for more flexibility.

Keywords: SLA, Discovery, Negotiation, Alternative Approaches, SLA Management, Semantic Annotated SLAs

Grid Resource Management with Service Level Agreements

Tianchao Li (IBM Deutschland - Böblingen, DE)

Based on a commercial application scenario, we discuss the application of service level agreements (SLAs) in distributed service provisioning systems (Grid or cloud) for autonomous resource management. We discuss the major issues including the specification and management of SLAs, monitoring and prediction of application run time, effective scheduling of the execution of provisioned services, as well as the establishment of a global infrastructure that assists the discovery and negotiation of SLAs with candidate services. We also present a demonstration that illustrates our solution with a comprehensive environment to support the development, usage and management of applications for the SLA-managed service provisioning system.

Keywords: Grid Computing, Cloud Computing, Resource Management, Service Level Agreements

Joint work of: Li, Tianchao; Gerndt, Hans Michael

AN ORGANIZATIONAL CONTEXT FOR SLAS

Julian Padget (University of Bath, GB)

The very name “Service Level Agreement” focusses attention on services and the interactions between services, but services are selected in a dynamic operational context that is in turn established by a problem to be solved, that in its turn is derived from the goals of an organization (virtual or otherwise).

In particular, these surrounding contexts are potentially useful in deciding what to do when a given service or workflow fails, since they can provide the normative framework that can help determine the appropriate corrective action, of which the service level is typically necessarily ignorant.

We will outline the three-tier architecture being developed in the EU-funded Alive project (<http://www.ist-alive.eu>) in which organization, coordination and service functions interact with one another to deliver adaptable service-oriented architectures, with illustrations from industrial use cases.

Keywords: Norms, organizational models, workflow execution, adaptation

(More) Negotiation in WS-Agreement

Michael Parkin (Tilburg University, NL)

The agreement protocol defined in the WS-Agreement specification lacks the capability to negotiate agreements through the multi-round exchange of messages within a single context. Current proposals to achieve negotiation have a two-stage approach: an initial quote request-quote phase, then offer-accept. This approach has an inherent problem because if the offer is not accepted in the second stage the protocol terminates and a negotiation must be re-started in a new negotiation context.

This short (and unfinished) piece of work describes a proposal to add negotiation within a single context to the existing WS-Agreement model. It is a backwards compatible (in the sense that the changes are additions and not modifications to the existing specification) and relatively simple approach.

Keywords: WS-Agreement, Negotiation, Proposal

Trust Negotiation and Service Level Agreements

Shamima Paurobally (University of Westminster - London, GB)

Providing effective Grid resource allocation mechanisms is a key factor in realising the Grid vision. In this talk, we present the SLA negotiation component developed in the Ontogrid project. We present the Insurance use case between insurance companies and repair services, where negotiation occurs through a broker service. Various negotiation protocols are implemented to provision a damage report and these are the contract net protocol, bilateral protocol and English auctions with timeouts. Strategies, preferences, and evaluation results are discussed.

In the second part of the talk, we discuss how the existing negotiation service can be extended for trust negotiation between semantic web services. VO authorisation and online trust systems lack the ability to securely negotiate suitable

trust and privacy requirements. Given the limitations of current VO authorisation systems when faced with malicious parties, this talk argues that trust can be built incrementally through trust negotiation, but without actual exchange of credentials until an SLA has been reached of what to exchange specifically. Finally we discuss if trust can help in SLA agreement. For example, less guarantees are needed for more trustworthee partners, or less chance of SLA violation occurring to preserve trust relationships or reputations scores.

Keywords: Negotiation, trust, strategies, protocols

ASKALON: Towards integration of end-user SLA support

Kassian Plankensteiner (Universität Innsbruck, AT)

The ASKALON Grid Workflow Management System implements an integrated approach for the development, planning as well as the execution of scientific workflow applications in computational Grids.

We describe techniques currently in use by the ASKALON system to improve the Quality of Service for the end user: bi-criteria scheduling, advance reservation as well as an approach implementing a continuous double auction model. While these methods improve Quality of Service for the end user on a best effort basis, they fail to provide guarantees concerning QoS parameters specified by the client.

We conclude by presenting an architecture for integrating support for Service Level Agreement-based workflow-QoS guarantees to the end user in environments that do not support the establishment of SLAs between resource providers and consumers, which is a fact in current scientific Grid infrastructures.

Keywords: ASKALON, grid workflow, sla, QoS, grid, HPC

Intelligent Resource Management for Distributed Systems - Concept of a SLA-Based Decision Model -

Tim Püschel (Universität Freiburg, DE)

Many existing applications exhibit strongly varying demand patterns for computing resources. Demand for computing resources is further increased by the introduction of new products and processes. Accommodating an ever increasing and highly fluctuating demand requires that sufficient resources are available.

The more providers offer their services, the more likely it is that they can be accessed in quantity and quality. Therefore it is important to attract more providers. Providers, however, will only offer their services if they can realize sufficient benefit.

To this end a decision model is presented which supports providers in taking decisions about management of resources and capacity, offered services as well as pricing. Different scenarios and stochastic optimization models are introduced. Subsequently mechanisms to solve these models are introduced and evaluated.

Negotiation and Monitoring of SLAs

Thomas Quillinan (VU University Amsterdam, NL)

A Service Level Agreement (SLA) is an agreement between a client and a provider in the context of a particular service provision. SLAs specify Quality of Service (QoS) properties that must be maintained by a provider during service provision – generally defined as a set of Service Level Objectives (SLOs). While significant work exists on how to specify and monitor SLOs, not much work has focused on the consequences of violating these SLOs. Furthermore, the definition of both violations and their associated penalties has not been examined in much detail. While it may seem reasonable to penalise SLA non-compliance, there are a number of concerns when issuing such penalties. For example, determining whether the service provider is the only party that should be penalised, or determining the type of penalty that is applied to each party.

Furthermore, what, if any, penalties are appropriate when a party inadvertently violates an agreement. In the energy domain, producers and consumers negotiate agreements to fulfill specific energy requirements. Typically these agreements represent an aggregation of several producers. For example, if one of these producers utilises wind power and there is no wind, this can technically be seen as a violation. However, determining the equitability of such cases is important towards acceptance of autonomous negotiations. Negotiation of both the terms of service and any penalty clauses in an autonomous fashion is an important step towards the realisation of an economy of services. While significant work in automated negotiation exists, it is primarily aimed at finding exact matches in single shot protocols. Furthermore, SLAs are typically agreed between human actors, for both legal reasons and a basic lack of trust in autonomous interactions.

Extending these negotiations towards multi-round automated agreements is a fundamental challenge that is arguable one of the biggest tasks that must be addressed in both the Web Services and Grid communities.

Keywords: SLA Monitoring, Violations

Issues and Challenges with SLA-Based Job Scheduling

Rizos Sakellariou (University of Manchester, GB)

This talk will present work of a recently completed project which focused on the benefits resulting from the usage of Service Level Agreements in parallel job scheduling for Grid-enabled resources (<http://www.gridscheduling.org>).

The talk will introduce the motivation for this research, highlight some of the key results of the project, and discuss the need for simplicity, which is based on clear engineering objectives, in SLA-based environments.

Keywords: Grid Scheduling

The Role of Service Level Agreements in Distributed Activity Management

Philipp Wieder (TU Dortmund, DE)

When managing distributed resources, fragmented and dispersed information of different provenance is generated. In current distributed systems, this includes information like resource requirements (e.g. hardware, software, or services), resource capabilities, execution states, security and ownership context, accounting, auditing or historic data. The challenge is to track all information related to one activity, as for example a job, a work flow, or a service invocation, and to provide the means to query and manipulate this information.

The approach presented here assumes that once a service consumer and provider agreed upon the service provision through a Service Level Agreement, the activity comes into existence and all related information is captured. The central entity of the approach is the activity document which is managed by an activity provider. Through this, clients can query and update activity-related information. With respect to Service Level Agreements, this enables management services to control SLA-compliance and, in case of violation, notify the relevant parties and steer potential countermeasures.

Keywords: Resource Management, SLA Negotiation, Activity, SLA Violation

WSAG4J: A flexible runtime environment for WS-Agreement based services

Oliver Wälldrich (Fraunhofer SCAI - St. Augustin, DE)

WSAG4J: A flexible runtime environment for WS-Agreement based services

Service Level Agreements (SLA) are used more and in grid environments. WS-Agreement is one of the accepted standards to express SLAs. It defines means to describe services, their constraints and guarantees in a very general way. However, in practice it turned out that this generality of WS-Agreement poses great demands on implementation that want to use SLAs. WSAG4J aims to overcome this problem in lowering the level of entry for services to implement SLAs and to foster best practices in SLA design based on WS-Agreement.

Keywords: Service Level Agreements, Service Oriented Architectures, Service Management

Multi-Layer SLA Management

Ramin Yahyapour (TU Dortmund, DE)

While many projects consider SLA management, most of them assume single consumer-provider relations for an SLA. Typical industrial use cases show that SLA negotiation is often not restricted to two parties.

More often, the consumer of an SLA is not even part of the formal SLA negotiations. At the same time the provider may use several agreements within its own organization or using sub-contractors to provide an overall SLA to a consumer. To this end, we see a clear need to support multi-level SLA management which allows automatic and iterative negotiation beyond several layers. The interaction between layers leads to complex orchestration and planning of SLAs.

We propose a self-* entity to manage individual SLAs. In this way, the use-case specific knowledge and requirements are encapsulated. The so-called SLA management instance does negotiation, planning, initiate provisioning, monitoring, adaption, reporting. We assume different management entities for different layers, services and resources. They can abstract from the use-case specifics and cooperate between each other. Our architecture employs WS-Agreement as a core technology. However, we see need for extensions in the area of (re-)negotiation. Moreover, a common monitoring infrastructure is needed.

The EC project SLA@SOI is addressing this field of research to support business requirements for a service economy.

Keywords: SLA management, QoS, workflow orchestration

Negotiating license terms with WS-Agreement

Wolfgang Ziegler (Fraunhofer SCAI - St. Augustin, DE)

The main conclusion highlighted in The 451 Group's 2005 report on The Impact of Software Licensing in Grid was that "Software licensing practices are limiting the acceleration of Grid adoption". Since 2005 additional paradigms of distributed computing like Clouds and SOA have emerged, however, traditional software licensing practices have not been adapted. Until now, software licenses are provided on the basis of a named user, node-locked host, number of concurrent jobs, or possibly a site license. These models are not sufficiently flexible to use commercial applications in Grids, Clouds or SOA that include access to resources beyond the current administrative domain. Moreover, the rapid emergence of mainstream multi-core processors and virtualization environments is requiring an evolution in software licensing models. We will introduce SmartLM license technology developments and corresponding business models that overcome the limitations of the existing models by rendering software licenses into mobile objects. Service Level Agreements used for creating SmartLM license objects are based on WS-Agreement. We present an extension of WS-Agreement, which is used both for the initial negotiation of license terms required for the execution of an application and a possible re-negotiation of the terms if necessary at a later stage.

Toward Reputable Grids

Gregor von Laszewski (Rochester Institute of Technology, US)

Reputation based framework for Grids:

Toward Reputable Grids G. von Laszewski, B. K. Alunkal and I. Veljkovic, pages 95-106, Volume 6, no. 3 (September 2005), Scalable Computing: Practice and Experience Scientific International Journal for Parallel and Distributed Computing

http://www.scpe.org/vols/vol06/no3/SCPE_6_3_09.pdf

The Grid approach provides a vision to access, use, and manage heterogeneous resources in virtual organizations across multiple domains and organizations. This paper foremost analyses some of the issues related to establishing trust and reputation in a Grid. Integrating reputation into quality management provides a way to reevaluate resource selection and service level agreement mechanisms. We introduce a reputation management framework for Grids to work toward facilitating the complex task of improving the quality of resource selection. Based on community experience we adapt trust and reputation of entities through specialized services.

Simple contextual quality statements are evaluated in order to effect the reputation for a monitored resource. Additionally, we introduce a novel algorithm for evaluating Grid reputation by combining two known concepts using eigenvectors to compute reputation and integrating global trust.

Keywords: Grid, SLA, reputation

Full Paper:

http://www.scpe.org/vols/vol06/no3/SCPE_6_3_09.pdf

See also: Toward Reputable Grids G. von Laszewski, B. K. Alunkal and I. Veljkovic, pages 95-106, Volume 6, no. 3 (September 2005), Scalable Computing: Practice and Experience Scientific International Journal for Parallel and Distributed Computing