# AssessGrid Strategies for Provider Ranking Mechanisms in Risk–Aware Grid Systems*

Dominic Battré[1], Karim Djemame[2], Iain Gourlay[2], Matthias Hovestadt[1], Odej Kao[1], James Padgett[2], Kerstin Voss[3], and Daniel Warneke[1]

[1] Technische Universität Berlin, Germany
{battre,maho,okao,warneke}@cs.tu-berlin.de
[2] School of Computing
University of Leeds, United Kingdom
{karim,iain,jamesp}@comp.leeds.ac.uk
[3] Paderborn Center for Parallel Computing,
Universität Paderborn, Germany
kerstinv@uni-paderborn.de

**Abstract.** Grid systems are on the verge of attracting the commercial user who requires contractually fixed levels of service quality. Service Level Agreements (SLAs) are powerful instruments for describing all obligations and expectations within such a Grid-based business relationship. Service selection has so far been based on performance and compatibility criteria while neglecting the factor of reliability and risk.

The EC-funded project "AssessGrid" aims at introducing risk assessment and management as a novel decision paradigm into Grid computing. With AssessGrid, providers are able to express the risk associated with an SLA, and broker services are able to judge the trustworthiness of such provider risk statements. This paper focuses on the provider ranking process where a broker or end-user has to decide which provider to choose from, and consequently which SLA to commit to.

**Keywords:** Grid, SLA, Negotiation, Broker, Ranking.

## 1 Introduction

Advances in Grid computing research have in recent years resulted in considerable commercial interest in utilizing Grid infrastructures for application and service provisioning. However, significant developments in the areas of risk and dependability are necessary before widespread commercial adoption can become reality. Specifically, risk management mechanisms need to be incorporated into Grid infrastructures in order to move beyond the best-effort approach that current Grid infrastructures follow to service provision.

* This work has been partially supported by the EU within the 6th Framework Programme under contract IST-031772 "Advanced Risk Assessment and Management for Trustable Grids" (AssessGrid).

J. Altmann, D. Neumann, and T. Fahringer (Eds.): GECON 2008, LNCS 5206, pp. 226–233, 2008.

AssessGrid addresses the key issue of risk by developing a framework to support risk assessment and management for all three Grid actors (end-user, broker, and resource provider) [1]. To integrate risk awareness and support risk management in all Grid layers, new components are introduced: the provider benefits from access to a consultant service that provides statistical information to support both risk assessment and the identification of infrastructure bottlenecks. The broker makes use of a confidence service that provides a reliability measure of a resource provider's risk assessment, based on historical data. In addition, a workflow assessor supports the broker deriving the probability of failure of a workflow from risk estimations of the sub-tasks.

Having risk estimations of single jobs and even workflow jobs available, Grid stakeholders negotiating an SLA have a concrete idea on the risk associated with a particular business activity. Prior to the binding agreement of an SLA, the customer (e.g. the Grid end-user or a Grid broker) usually requests a non-binding SLA quote from one or more providers, which holds all information like price, penalty, or the probability of failure (PoF) of the SLA. This way each party can decide whether or not to accept this risk by committing to a binding SLA. At least if Grid brokers have to map complex workflows to Grid resources, it is common practice to not only request a single SLA quote from a single provider at a time, but from numerous providers in parallel. This way the broker is able to optimize the workflow orchestration according to the particularly available resources at provider side. Such a broadcast request usually results in a large number of non-binding SLA quotes from numerous different providers. Even if the customer is easily able to decide whether the PoF of a particular SLA offer is acceptable, it remains difficult to select the best offer among them.

For supporting the customer in this decision making process, the AssessGrid project will introduce a provider ranking mechanism, which is presented in this paper. After an outlook on related work in section 2, we briefly describe the *getQuote* mechanism in SLA negotiation process in section 3. The main part of this paper focuses in Section 4 on the provider ranking process. A short conclusion ends this paper.

## 2 Related Work

The Grid resource selection from a user's perspective can be separated in two phases. The first phase comprises the discovery of resources that match the user's requirements associated with a job or workflow (necessary condition). The second phase comprises the ranking of these resources so that a user can select the resource with the highest utility/performance.

Resource discovery has been addressed in the past for example with information services like the Globus Monitoring and Discovery Service (MDS) [2]. Performance requirements of the response time are evaluated in [3] and according to data collected by the Grid Index Information Service (GIIS) predications about the response time of queries can be made.

Provider rankings can be based on different utility functions. One obvious criterion is the performance of a provider in terms of resources employed and available. This has been discussed for example in [4] where a framework is described to use Gridbench performance probes to determine resources that are best suited to a user's job. The framework allows the user to individually define filtering, aggregation, and ranking criteria for a custom utility function. A second criterion is of course the price requested by a provider, and projects such as GridEcon [5] define market places for resources.

Given the enormous job failure rates observed in Grids (in DAS-2 more than 10% of all jobs fail [6], in TeraGrid the failure rate is 10–45% [7], and in Grid3 27% jobs fail even with 5–10 retries [8]) it becomes apparent that quality of service and the capability to negotiate SLAs is another key decision factor.

This has been the focus of the AssessGrid project [1] which developed mechanisms to estimate the probability of failure of a job. This estimation can be only be performed from the provider side since it has information about the exact scheduling, planned fault-tolerance mechanisms, and stability of resources which will be used. The confidence service at the broker layer is able to estimate the reliability of the providers' published failure information by setting it into a relation with observed SLA violations [9].

Several papers elaborate on reputation based mechanisms. Elnaffar describes in [10] a ranking mechanism for Grid providers based reputation. The metric employed is a vector of user ranking (a rating entered manually by the user according to the perceived performance), Quality of Service (QoS) conformance (measured discrepancy between asserted and delivered QoS), and fidelity (consistency of delivered performance).

Sonnek and Weissman review in [11] several reputation systems for the Grid and give a quantitative comparison. The reviewed ranking systems comprise the Ebay system. Providers are ranked in a personalized way based on a user's direct experience and other approaches that filter dishonest feedback.

In this paper we define several other criteria that are relevant for the utility of a provider's resources to the user, which are elaborated in Section 4.

## 3 Quote Mechanism in SLA-Negotiation

Grid Service Level management contains QoS descriptions for Web services in the form of SLAs. The Grid community has identified the need for a standard for SLA description and negotiation. This has led to the development of WS-Agreement [12], a language and protocol designed for advertising the capabilities of providers and creating agreements based on initial offers, and for monitoring agreement compliance at runtime. These upcoming standards rely on the Web Services Resource Framework (WSRF).

This WS-Agreement protocol now has been extended to allow flexible SLA negotiation schemes between contractors and service providers. Briefly, modifications consist in the addition of one operation: *getQuote()*. This is only an extension, which allows to change the original single-round acceptance model to a two-phase acceptance model. It introduces the negotiation possibility, in other words the bargaining capability.

The protocol implemented within the EC-funded AssessGrid project is a two-phase commit negotiation. The user first requests a template from the provider, which describes the provider capabilities. The user then specifies his requirements in a quote request. The provider makes an offer by sending a quote based on the request made by the user. The user is then able to accept this quote and sign it. The SLA contract is then signed if the provider accepts the user's signed SLA.

This modification to obtain a flexible and robust SLA negotiation protocol can be seen as a continuation of work within the WS-Agreement specification. The extended protocol answers the requirements where a negotiation before a final agreement is needed.

The agreement mechanism within the WS-Agreement draft specification does not meet the negotiation requirement. The main drawback comes from the single round "offer, accept" agreement mechanism. This has an important consequence: there is no possibility for a service consumer to request offers from different providers so that he can choose the best one among these. In order to do so, he would have to act as the agreement initiator and call *createAgreement()* from several providers to propose some SLA to each. The problem is that he would then be bound to every provider that decides to accept the quote. The concept of "SLA quote" does not exist in the WS-Agreement draft specification: it is not possible for a consumer to simply ask a provider what his terms would be without being committed to the provider by this action. In the real word, a negotiation process usually begins by the initiator asking non-committing questions to the other party.

The solution proposed in AssessGrid is to introduce the concept of "SLA quote" into the agreement mechanism. The *getQuotes()* method offers the end-user the possibility to have a first evaluation of a request for service. Based on this first quote, the user can then decide to accept it using the *createAgreement()* method. If the provider's quote is not satisfactory, a new quote can be requested by entering a new quote request, with slightly different parameters.

## 4 Provider Ranking

The negotiation of an SLA is the first step in the business relationship between a customer and the provider. Even if both parties are interested in a successful execution of an SLA, both are driven by different - often opposing - goals, like high quality service at low cost (customer) vs. maximizing the revenue (provider). In this light, providers may even be tempted to lie regarding their service quality level.

### 4.1 Confidence Service

During the SLA negotiation, the customer may specify the required level of risk in the SLA request. The provider answers this request by publishing the risk level it is able to support. This may either be the risk that the customer demands or lower. At this stage, the customer has to trust the correctness of the provider risk calculation.

For attracting additional Grid jobs and to increase the system utilization and revenue, providers may publish a lower risk than they are actually able to support. For coping with this situation, the AssessGrid broker provides a confidence service, rating the correctness of the provider specified risk value by considering the provider's past reliability [9]. Using appropriate risk models, the broker can deduct the likeliness that a provider performs as announced.

Grid end-users are able to directly request resources from Grid providers. However, the confidence service is a strong incentive for using the services of the Grid broker, because the Grid end-user usually does not have the broker's experience dealing with thousands or millions of SLA-based jobs and workflows and the history behind their specification, management, and outcome, which is mandatory for statistically firm provider ratings. Thus, for the broker the quality of the confidence service is a key argument for attracting customers.

### 4.2 Provider Performance

The information provided by the confidence service is a mandatory key when ranking a list of offerings: all offers having a poor confidence value in the provider specified risk may be filtered, since it is very likely that the actual risk of executing in the context of the SLA in question is not acceptable for the customer. However, the broker's information on provider performance can be further used for not only filtering, but also ranking.

As outlined above both parties are interested in fulfilling an SLA. Hence, AssessGrid will support provider ranking according to the provider's SLA violation rate. Here, the broker is using the floating average of logged provider performance, e.g. focusing on daily or weekly average values. Following this ranking, the Grid end-user is able to select a provider that complies to his risk requirements while showing the best SLA conformance among all offers.

This ranking approach may be enhanced to other parameters. In classic queuing based Resource Management Systems (RMSs), the waiting time of a job indicates how long a job has to wait in the queue until compute resources have been allocated to the job. In general, the smaller the waiting time, the better the service. In case of AssessGrid, the RMS is planning-based and therefore does not use any queues. However, in the case of deadline bound jobs, the provider has a time window ranging from earliest start time (EST) until deadline (DL) for executing the job with a defined runtime (RT). This results in a slack time of $DL - EST - RT$.

In the AssessGrid scenario, the waiting time is reciprocal to the slack time: the end-user is interested in providers offering a small waiting time, so that jobs are started as early as possible after EST. Providers performing with a high waiting time execute the job with only a small slack-time. Thus, the provider executes the job in an SLA-compliant way, returning all results until the specified deadline, but the customer has to anticipate getting results just in time.

### 4.3 Acceptance Rate

As described above, the SLA quote given by the provider has a non-binding character. End-users may request SLA quotes for getting an overview about available resources, mandatory for an SLA-compliant mapping of workflows onto Grid resources. Providers in turn do not have to block resources, answering SLA quotes without any risk or cost.

Providers may exploit this situation by answering SLA quote requests while knowing that offered resources are not available. This is similar to overbooking, where a provider accepts more requests than resources available. Here, the situation is even better for the provider due to the non-binding nature of SLA quotes: the provider may anticipate that resources are available at a later time, when this customer asks for a binding SLA due to the answered SLA quote. Applying a conservative quote policy means the provider would have neither answered the quote request nor got the binding to SLA request, even though resources were available. Hence, providers should be optimistic when answering SLA quotes, but not aggressive, answering quote requests if their fulfillment is unrealistic.

Even if an SLA quote is non-binding and the provider is not obliged to reserve any resources, the requestor should be able to expect that the SLA quote is at least a short term commitment of the provider: if the requestor immediately replies to the request, he may expect the provider to agree to this request.

Using the information on previous negotiations with a specific provider, the AssessGrid broker service is able to determine the ratio of SLA quotes resulting in successful SLA agreements. Moreover, the time between quote and agreement can be considered. Evidently, the ratio decreases with an increasing time span, having multiple customers competing on using a limited set of resources. The more time between non-binding SLA quote and binding SLA request, the more likely it is that resources have been assigned to another request meanwhile. High-quality providers are characterized by a high ratio curve. The higher the probability curve starts and the slower the curve descents, the better for the SLA requestor.

This knowledge is particularly beneficial for brokers mapping workflow tasks on resources, using a two phase procedure. In the first phase an SLA-quote based mapping of tasks to providers is executed, considering time dependencies between consecutive workflow tasks as well as deadlines. If this phase results in a valid mapping, the second phase then focuses on creating binding SLA agreements. SLA offers not resulting in SLA agreements are particularly problematic for the broker at this point, since it impacts the entire workflow mapping, where a single workflow task can no longer be mapped as planned, while other tasks already have binding SLA agreements. In such a case, the broker has to re-enter the first phase for all tasks where no binding SLA has been agreed yet, trying to map the workflow in a different way. As a matter of fact this remapping process may fail, resulting in SLA cancellation costs for the broker.

The AssessGrid broker service is able to use this SLA acceptance ratio curve as ranking or even filtering parameter. Choosing to deal with high-quality providers may be more costly than using low-cost providers, but does pay off at the end due to reduced SLA cancellation costs and an increased service quality level for broker's customers.

### 4.4 Certified and Non-certified Provider Classification

The provider landscape in large scale Grids is extremely heterogeneous, offering all different kinds of resources and services. Also the quality support and administration is very diverse, ranging from high-class compute centers with 24/7 support, over compute resources operated by universities with 12/5 support, up to garage level compute centers with no regular or professional administration at all. Analogously, the quality of hardware resources, the level of redundancy and fault tolerance, or the local security policies also differ significantly.

Obviously the level of support, administration, or other parameters have significant impact on the price of resource usage as well as the provider specified risk value: depending on the acceptable risk, the customer will prefer higher priced SLA offers.

In normal life, classification systems help us in selecting services according to our needs and expectations. In a three star hotel we can expect a color TV in the room, while four star hotels provide 24/7 reception service or a hotel pool. Using data mining methods on the broker information pool, such categories can be established by comparing infrastructure information provided by the provider with the provider's performance data. Obviously, parameters like 24/7 support do have a strong correlation with low risk values and low SLA violation rates. Other parameters like the type of locally used RMS show a strong correlation with acceptance ratio.

Deducting abstract provider classes from these data mining results, may they be specific to a broker or accepted within the entire Grid, help the Grid end-user in ranking and filtering SLA quotes. Similar to the business traveler only looking for business class flights, without really checking the actual services provided to business class travelers, the Grid end-user may select "silver class" providers without checking for detailed services or data.

In this context third-party certificates have focal importance. Even if logfile analysis may reveal contradictions between published data and actual performance (e.g. a provider publishing 24/7 support, only showing 12/5 performance), a provider may lie about other published properties (e.g. policies regarding access security for compute facilities). Here, the provider statements could be certified by a third party. The broker could establish such a certification process for key providers, offering this as additional service for its customers.

## 5 Conclusion

This paper has discussed strategies of the AssessGrid project in relation to ranking mechanisms of SLA offers. The new negotiation process is built on a non-binding SLA request which enables end-users to broadcast an SLA request, and receive and compare SLA offers from a large number of providers. Applying this negotiation mechanism, the end-user requires a ranking mechanism in order to select the best one among them. The history of negotiations recorded by a broker can supplement the ranking mechanisms with several useful metrics such as reliability of estimated failure probabilities, performance (overachieving an SLA) and the acceptance rate of issued non-binding offers. This can help end-users as well as brokers with selecting suitable resources.

## References

1. AssessGrid – Advanced Risk Assessment and Management for Trustable Grids, http://www.assessgrid.eu
2. Czajkowski, K., Kesselman, C., Fitzgerald, S., Foster, I.: Grid information services for distributed resource sharing. High Performance Distributed Computing, 181–194 (2001)
3. Keung, H.N.L.C., Dyson, J.R.D., Jarvis, S.A., Nudd, G.R.: Predicting the Performance of Globus Monitoring and Discovery Service (MDS-2) Queries. In: GRID 2003: Proceedings of the Fourth International Workshop on Grid Computing, Washington, DC, USA, p. 176. IEEE Computer Society, Los Alamitos (2003)
4. Tsouloupas, G., Dikaiakos, M.: Ranking and performance exploration of grid infrastructures: An interactive approach. In: Grid Computing, 7th IEEE/ACM International Conference on, September 2006, pp. 313–314 (2006)
5. Grid Economics and Business Models (GridEcon), http://www.gridecon.eu
6. Iosup, A., Epema, D.: GRENCHMARK: A Framework for Analyzing, Testing, and Comparing Grids. In: CCGRID 2006 Sixth IEEE International Symposium on Cluster Computing and the Grid, pp. 313–320 (2006)
7. Khalili, O., He, J., Olschanowsky, C., Snavely, A., Casanova, H.: Measuring the Performance and Reliability of Production Computational Grids. In: GRID – 7th IEEE/ACM International Conference on Grid Computing (GRID 2006), Barcelona, Spain, September 2006, pp. 293–300. IEEE, Los Alamitos (2006)
8. Dumitrescu, C., Raicu, I., Foster, I.T.: Experiences in Running Workloads over Grid3. In: Zhuge, H., Fox, G.C. (eds.) GCC 2005. LNCS, vol. 3795, pp. 274–286. Springer, Heidelberg (2005)
9. Gourlay, I., Djemame, K., Padgett, J.: Risk and Reliability in Grid Resource Brokering. In: Proceedings of IEEE International Conference on Digital Ecosystems and Technologies 2008 (DEST 2008), Phitsanulok, Thailand (Febuary 2008)
10. Elnaffar, S.: Beyond User Ranking: Expanding the Definition of Reputation in Grid Computing. In: Advances and Innovations in Systems, Computing Sciences and Software Engineering, pp. 381–386. Springer, Netherlands (2007)
11. Sonnek, J.D., Weissman, J.B.: A Quantitative Comparison of Reputation Systems in the Grid. In: GRID 2005: Proceedings of the 6th IEEE/ACM International Workshop on Grid Computing, Washington, DC, USA, pp. 242–249. IEEE Computer Society, Los Alamitos (2005)
12. Andrieux, A., Czajkowski, K., Dan, A., Keahey, K., Ludwig, H., Nakata, T., Pruyne, J., Rofrano, J., Tuecke, S., Xu, M.: Web Services Agreement Specification (WS-Agreement). GRAAP WG - Open Grid Forum (OGF) (March 2007), http://www.ogf.org/documents/GFD.107.pdf