AssessGrid, Economic Issues Underlying Risk Awareness in Grids*

Kerstin Voss[1], Karim Djemame[2], Iain Gourlay[2], and James Padgett[2]

[1] Paderborn Center for Parallel Computing, University of Paderborn, Germany
kerstinv@uni-paderborn.de
[2] School of Computing, University of Leeds, United Kingdom
{karim,iain,jamesp}@comp.leeds.ac.uk

Abstract. In order to improve the attractiveness and drive the commercial uptake of Grid technologies, the establishment of Service Level Agreements (SLAs) is required. The AssessGrid project contributes to this aim by introducing risk-aware Grid architectural components. Grid service users, brokers and providers benefit from risk assessment functionalities in all phases of service provisioning and utilisation. This paper focuses on the economic issues which result from this new risk-aware approach to Grid computing. Multiple open economic research questions are discussed from the perspective of users, brokers and providers, which point out the potential impact of AssessGrid in this area.

Keywords: Risk Assessment, Risk Management, SLA, AssessGrid.

1 Introduction

The AssessGrid project is motivated by the need to bridge the gap between Service Level Agreements (SLAs) as a concept and as an accepted tool in Grid utilisation and service provisioning. If widespread acceptance is to become a reality, a number of issues need to be addressed that are important to the two main Grid actors: the Grid service providers and the Grid service users. Service providers may be unwilling to agree an SLA since they are aware of the possibility of resource failures. This could lead to an SLA violation and consequently a need to pay a penalty fee. Meanwhile, end-users do not completely trust SLAs since they are also aware of possible resource failures and know that a *one hundred percent* guarantee cannot be given by any provider. To establish the usage of SLAs for Grids, which is essential for commercial Grid utilisation, the aim of the AssessGrid project [1] is to integrate a risk-aware SLA model into current Grid technology. The provision of risk assessment methods brings added benefits for end-users as well as a new potential market for Grid services and resource brokers. The risk information is integrated in SLAs as an additional negotiable parameter in order to notify end-users about the probability that the SLA may

* 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).

D.J. Veit and J. Altmann (Eds.): GECON 2007, LNCS 4685, pp. 170–175, 2007.

be violated. The probability of failure (PoF) published in the SLA will enable end-users to compare different SLA offers. In particular, the PoF will influence the price and the penalty offered by the Grid service provider, e.g. high costs and penalties for job executions which have a high predicted success rate and vice versa. Through this extension end-users obtain a new perspective when selecting an SLA offer since they can individually evaluate the balance between the consequences of an SLA violation and the price they are willing to pay for Grid service usage. The use of PoFs in SLAs opens new possibilities for Grid brokers in the scope of workflows and the evaluation of the reliability of PoF values published by providers.

In order to enable Grid service providers to offer a PoF for an SLA violation, risk assessment methods have to be integrated into the Grid fabric. The usage of risk information is, however, not limited to publishing it in SLA offers. Rather the AssessGrid developments integrate risk management methods for brokers and providers which use the assessed risk as a decisive component in their scheduling and resource management activities.

This paper presents an economic view of the AssessGrid project in order to point out the potential impact of the project. It is structured as follows. In section 2 an overview of the project is given. In section 3 the identified Grid economy research issues are described. General conclusions and future work are presented in section 4.

2 AssessGrid: Risk Management in the Grid

AssessGrid aims to introduce risk assessment and management [2] in a Grid environment. In professional risk management, risk is not only the likelihood of occurrence. Since in the scope of risk management different events have to be compared, it is not sufficient to only consider their likelihood. To develop an accurate risk management process, the consequences of an event also have to be taken into account when comparing different threats. Accordingly, a more complete definition describes risk in terms of the product of the likelihood of an event and the impact of its occurrence. From an economic perspective, it is desirable to express the impact in monetary terms. However, a provider or broker cannot, in general, evaluate the financial impact of an SLA violation on an end-user since the data required to make such an evaluation is not available to them. Consequently, the provider offers a probability of failure (PoF) instead of a risk value. Note that the risk for an SLA violation from the provider's perspective can be determined exactly by the PoF and the agreed penalty.

We have identified three main actors in the scope of Grid service provisioning and utilisation: end-user, broker, and provider. The end-user is the Grid service consumer, who specifies quality of service (QoS) requirements through an SLA. In the computational Grid, end-users run various applications and define SLAs on a per-application basis since each Grid job has different requirements and time constraints. The Grid fabric is comprised of resources that are owned by various Grid resource providers. Providers charge for the utilisation of their resources,

agreeing SLAs that specify the cost, the provider's obligations, and penalty fees to be paid in the event that those obligations are not met. Grid resources can then either be accessed directly or through the use of a broker. A broker is defined as a business role that acts as a third party between organisations that consume Grid resources and Grid providers that offer these services. It can offer additional services to the end-user due to benefiting from the data acquired through its many negotiations with various providers.

Details on these actors as well as the architecture for the risk-aware Grid components developed in the project can be found in [3].

3 Grid Economy Research Issues

The objective of AssessGrid is to contribute to the establishment of commercial adoption of Grid technology. Therefore it is essential to introduce economy-awareness into each role - end-user, broker, and provider. This sections presents aspects of the architecture where an economy model may be exploited. The current scope of AssessGrid is the provision of composite (e.g. SLAs, workflows etc) and computational (e.g. physical Grid infrastructure) service markets.

3.1 End-User

The end-user is provided with a number of abstract applications which make use of Grid services deployed within the Grid fabric layer. SLA requests and offers are exchanged between end-user and broker or provider, in order to agree an SLA which grants permission to invoke a Grid service in the fabric layer. Within each layer, the organisation performing the role of each actor must define a policy statement governing the acceptable bounds of negotiation. This restricts end-users and contractors to request or offer SLAs which fall outside of the organisation's acceptable limits. For example, in addition to specifying budget constraints, there may be a restriction on a provider's penalty conditions to limit the financial loss incurred because of an SLA violation. Taking these policy limits into consideration, an end-user can negotiate an SLA to run a Grid service by defining requirements as well as the requested QoS in an SLA request. During the definition process the end-user evaluates the importance of the job in terms of its urgency and the consequences of delayed results or failure. A further validation of the policy limits must be made against the SLA offers received from the broker or providers.

Where several SLA offers have been negotiated on behalf of the end-user, the broker can return a ranked list - according to price, penalty, and PoF. The challenge for the end-user is to find an SLA offer which offers the best service in terms of price, penalty, and PoF. Here we can apply a mathematical model to help the end-user make the *best* offer selection based on quality criteria. The end-user defines a ranking of the quality criteria (e.g. PoF is more important than price) in order to measure each of the offers according to its closeness to the criteria. A possible approach is the application of an Analytic Hierarchy Process which is based on criteria weights specified by the end-user [4].

From the perspective of the end-user, there are economic issues if the SLA is violated. In this case the end-user needs to evaluate whether the penalty received is sufficient to offset the consequences of the SLA violation. The evaluation results will be useful information for the negotiation of similar SLAs in the future.

3.2 Broker

Within the AssessGrid architecture the broker role facilitates SLA negotiation between entities fulfilling the end-user and resource provider roles. After the negotiation has returned an SLA offer, the broker is responsible for performing reliability checks on the PoFs contained in the SLA offers. Without this check, the end-user has no independent view on the provider's assessment, which cannot be assumed to be impartial. SLA offers that are deemed to be unreliable are subjected to an additional risk assessment by the broker using historical data related to the provider making the offer. Where multiple SLA offers are returned by the SLA negotiation process, the broker can rank these according to a price, penalty, PoF matrix depending on the priorities of the end-user.

The economic benefit of using a broker within the SLA negotiation process effects all three Grid actors and provides the opportunity for an economy model where SLAs for software services are bought and sold based on differentiated classes of service. In the provision of SLA negotiation, the broker offers two classes of service: mediator and runtime-responsible. In the case of mediator service, the broker provides a marketplace for providers to advertise their SLA templates to a wider number of end-users; for end-users it allows selecting a provider and their services from a larger set.

Use of the broker as a runtime-responsible service offers the greatest scope for Grid economy research. In this scenario the broker has the ability to buy SLA offers from providers and resell them to end-users transparently using its own SLA offer. In this way a broker becomes a virtual provider and can offer price, penalty, and PoF combinations unavailable from a single provider. In addition, the broker can orchestrate workflows, which combine multiple single SLA offers and combine these into SLA offers for an entire workflow. The broker can make trade-offs against price, penalty, and PoF between providers in order to maximise the economic benefit for itself. Where a task is executed redundantly, to reduce the PoF, or where it forms part of a workflow orchestration, the broker has additional responsibilities during runtime. Should an SLA governing one of these tasks be violated, the broker must determine whether it is more economical to pay its own penalty fee or to negotiate a new SLA offer with a different provider at a price which minimises its losses.

During post-runtime, the broker is responsible for updating the historical data it holds on each provider registered therewith. When offering it's runtime-responsible service, the broker can easily access the final status of SLA offers as they are agreed between itself and the provider. When acting as a mediator, the broker must persuade the end-user to pass on the same information about the SLA final status. A rebate or bonus payment may be built into the economy

model to encourage end-users to give SLA offer feedback to the broker. As well as the financial benefit, end-users will benefit through up-to-date historical data and greater confidence in the reliability checks.

3.3 Resource Provider

A provider offers access to resources and services through formal SLA offers specifying the requirements as well as PoF, price, and penalty. Providers need well-balanced infrastructures, so that they can maximise the offerable QoS and minimise the number of SLA violations. Such an approach increases the economic benefit and motivation of end-users to outsource their IT tasks.

A number of economic issues have been identified which affect the provider. These issues can be categorised as belonging to the pre-runtime (i.e. during SLA negotiation), run-time and post-runtime phases.

In the pre-runtime phase a risk aware negotiation requires that a provider place an advance reservation for the SLA and calculates the PoF [5]. Based on this, a provider determines the price and penalty fee which will be offered to an end-user. To ensure unsuitable SLA offers are not made, end-users define minimum and maximum limits for price, penalty, and PoF within the SLA request. A provider's decision whether to agree or reject an SLA depends on the fees and the requested PoF in comparison with the current status of its infrastructure.

The publication of the SLA PoF opens further research fields. A provider must not offer the PoF it had assessed during the reservation process since no mechanism can be developed which can coerce it into telling the truth. However, the broker's confidence service is designed to ensure that providers do not lie about published PoFs. Therefore it is the ability of the provider to fulfill SLA offers which marks it out as reliable, rather than its ability to offer SLAs with a low PoF.

For contractors (end-users or brokers), an important provider selection criterion is the price. The SLA template contains pricing information for actions such as data transfer, CPU usage, and storage. Within the AssessGrid model these prices are variable since the price depends on the PoF value specified within the SLA.

The market mechanism will influence the pricing since each provider has only a limited resource set with variable utilisation. Consequently, prices for resource usage will not be fixed but will depend on the economics of supply and demand. Reservations, which were made well in advance, will usually result in a reduced price since there will be access to a greater number of free reservation slots. Equally, immediate resource usage may also result in reduced prices, as providers try to increase their utilisation if demand is low. However, end-users risk resources unavailability if they wait too long before reserving resources. These pricing dynamics are valid only in the scope of the resource costs and do not consider the PoF.

After an SLA has been agreed by the provider and the end-user, the provider has to ensure during runtime that the SLA will not be violated. Accordingly, the provider's risk management activities are controlled by estimating the penalty

payments in the case of an SLA violation. Using the AssessGrid technology enables to initiate precautionary fault tolerance mechanisms in order to prevent SLA violations. The penalty fees, in addition to the PoF (i.e. risk), are the decisive factors in determining which fault tolerance mechanisms are initiated.

In the post-runtime phase the provider has to evaluate the final SLA status to determine whether a penalty fee has to be paid. Even in the case the SLA had been fulfilled the costs for the fulfillment have to be checked since the initiation of a fault tolerance mechanism also consumes resources and therewith results in additional costs. The results of the evaluation process will point out on the one hand whether adjustments in the offer making policies are necessary in order to increase the provider's profit. On the other hand, statistics can be generated which show whether initiated fault tolerance mechanisms had been able to prevent an SLA violation.

4 Conclusion

The integration of risk-awareness into the Grid provides a number of benefits within an economy framework but also gives rise to numerous research problems. This paper presents an overview of the AssessGrid developments and their impact for continuing research in economic models. An unexplored problem is the handling of workflows for a broker. Since the broker is responsible for the SLA fulfillment, it has to react on failures (negotiate with providers for a repeated job execution) in order to prevent paying penalties. Other essential economic issues of AssessGrid are the pricing mechanisms for brokers and providers which must take account of the probability of failure in a risk-aware Grid approach. Further research is required to address these problems in detail.

References

1. AssessGrid : (Advanced risk assessment and management for trustable grids), `http://www.AssessGrid.eu`
2. Koller, G.: Risk Assessment and Decision Making in Business and Industry. CRC Press, Boca Raton (1999)
3. Birkenheuer, G., Hovestadt, M., Voss, K., Kao, O., Djemame, K., Gourlay, I., Padgett, J.: Introducing Risk Management into the Grid. In: Proceedings of the 2nd IEEE International Conference on e-Science and Grid Computing, Amsterdam, The Netherlands (2006)
4. Saaty, T.: How to make a decision: The Analytic Hierarchy Process. European Journal of Operational Research 48, 9–26 (1990)
5. Hovestadt, M., Kao, O., Voß, K.: The First Step of Introducing Risk Management for Prepossessing SLAs. In: IEEE International Conference on Services Computing (SCC), Chicago (2006)