Chapter 6.8
Grid Service Level Agreements Using Financial Risk Analysis Techniques

Bin Li
University of Surrey, UK

Lee Gillam
University of Surrey, UK

ABSTRACT

Grid computing continues to hold promise for the high-availability of a wide range of computational systems and techniques. It is suggested that Grids will attain greater acceptance by a larger audience of commercial end-users if binding Service Level Agreements (SLAs) are provided. We discuss Grid commoditization, and in particular the use of Grid technologies for certain kinds of financial risk analysis where both data and computation requirements can be substantial. The nature of such analysis, and the need for it to run to completion, suggests the need to guarantee availability and capability in the underlying Grid infrastructure. This further suggests that it is necessary to be able to evaluate the infrastructure in relation both to historic analysis and to the needs of a specific analysis. Our aim, then, becomes one of predicting availability and capability, essentially introducing risk analysis for Grids. Prediction, quantification of risk, and consideration of liability in case of failure, are considered essential for the future provision of Grid Economics – specifically, relating to the provision of SLAs through resource brokers, and comparable to markets in other commodities – but perhaps also more widely applicable to the configuration and management of related architectures such as those of Peer-to-Peer (P2P) and Cloud Computing systems. The authors explore and evaluate some of the factors required for the automatic construction of SLAs, with broad consideration for Financial Risk and the potential formulation of a Grid Economy as a commodity market, which may in future involve the trading and hedging of risk, options, futures and structured products.

DOI: 10.4018/978-1-4666-0879-5.ch6.8

INTRODUCTION

Grid computing has emerged through consideration of secure networked availability of high end computer systems for relatively complex applications, and is reportedly in use in a wide variety of scientific and industrial pursuits. Commercial Grid systems, utility computing, and computing "in the Cloud" are all infrastructural considerations for businesses, and perhaps universities and other academic endeavours, seeking to reduce the costs of their infrastructures. These Grids, Utilities and Clouds can reportedly help organizations to offset environmental impacts and legal requirements relating to the manufacture, delivery, real estate, power and cooling, and subsequent recycling or disposal of the hardware itself. Furthermore, costs relating to ongoing maintenance and software licensing may also be mitigated to some extent. Certain organizations and individuals are variously willing to pay for the use of different forms of commoditized computer systems variously described as kinds of service, for example Software as a Service (SaaS), Platform as a Service (PaaS), Infrastructure as a Service (IaaS) and so on. A number of major IT providers have infrastructural support for SaaS, PaaS and their brethren, and allocate processor hours and storage in managed facilities at fixed prices. The notion, then, is that this enables buyers to focus on the outcomes not the infrastructure, and allows them to rapidly construct applications without entering into procurement and related installation activities from the outset.

To name a few, companies such as HP, Amazon, Google, Sun, Microsoft and IBM already offer or are beginning to promise some form of commoditized servers and pricing models for their Grids, Utilities, and Clouds. These can be used by organizations: (i) directly, for example for running various business services or simulations; (ii) as part of the external-facing business activity, such as supporting a website; (iii) repackaged in some form where a service is sold on to a customer, for example with a relatively simple business model in which some SaaS is built over an infrastructure with a price to the customer that encompasses the cost of the infrastructure. Effective price-performance is going to be a key decision factor in the use of such resources. And yet, commoditization does not necessarily enable the end user to undertake a specific analysis with particular requirements at a given time, to obtain the best price or margin against this analysis, or to manage the risk of the analysis failing and being able to re-run this within limited time. For some applications, the result of analysis may be required at a specific time, beyond which the opportunity to be gained from the analysis is lost. Lacking such assurances of availability, reliability, and, perhaps, liability, is likely to limit numbers of potential end-users. Indeed, a number of Grid Economics researchers have previously suggested that Grids will only be able to attain greater acceptance by a larger audience of commercial end-users if binding Service Level Agreements (SLAs) are provided. This could well be true for Utilities, Clouds, and any other future labelling of computing provision.

The financial sector appears to be one of the principal potential users of substantial commoditized computation, particularly for pricing models and portfolio risk management. However, the limited commercial adoption reported of commoditized infrastructures by the financial sector, preferring instead to maintain bespoke infrastructure with a variety of concomitant costs and related overheads, may suggest that, at minimum, the lack of assurances of availability, reliability and liability may contribute to the apparent lack of adoption of these infrastructures for such use.

Our aim is to support the adoption of such infrastructures by predicting availability and providing for liability through the consideration of risk analysis for Grids. Prediction, quantification of risk, and consideration of liability in case of failure, are considered essential for the future provision of Grid Economics – specifically, relating to the provision of SLAs through resource brokers, and comparable

Related Content