The Journal of Systems and Software 88 (2014) 42–53

Contents lists available at ScienceDirect

The Journal of Systems and Software

journal homepage: www.elsevier.com/locate/jss

Resource failures risk assessment modelling in distributed environments

Raid Alsoghayer[a], Karim Djemame[b,*]

[a] Computer Science Department, King Saud University, Riyadh 11543, Saudi Arabia
[b] School of Computing, University of Leeds, Leeds LS2 9JT, UK

ARTICLE INFO

Article history:
Received 23 February 2013
Received in revised form
10 September 2013
Accepted 11 September 2013
Available online 3 October 2013

Keywords:
Grid computing
Cloud computing
Risk assessment
Quality of Service
Resource failure
Markov Chains

ABSTRACT

Service providers offer access to resources and services in distributed environments such as Grids and Clouds through formal Service level Agreements (SLA), and need well-balanced infrastructures so that they can maximise the Quality of Service (QoS) they offer and minimise the number of SLA violations. We propose a mathematical model to predict the risk of failure of resources in such environments using a discrete-time analytical model driven by reliability functions fitted to observed data. The model relies on the resource historical data so as to predict the risk of failure for a given time interval. The model is evaluated by comparing the predicted risk of failure with the observed risk of failure, and is shown to accurately predict the resources risk of failure, allowing a service provider to selectively choose which SLA request to accept.

1. Introduction

Advances in Grid/Cloud computing research have in recent years resulted in considerable commercial interest in utilising infrastructures such as distributed environments provide to support commercial applications and services (Berman et al., 2003). However, significant developments in the areas of risk and dependability are necessary before widespread commercial adoption can become a reality. Specifically, risk management mechanisms need to be incorporated into Grid/Cloud infrastructures, in order to move beyond the best-effort approach to service provision that current Grid infrastructures follow (Djemame et al., 2006).

Risk management is a discipline that addresses the possibility that future events may cause adverse effects and is defined in The Risk Management Standard (2009) as "the process whereby organisations methodically address the risks attaching to their activities with the goal of achieving sustained benefit within each activity and across the portfolio of all activities."

The importance of risk management in Grid/Cloud computing is a consequence of the need to support various parties involved in making informed decisions regarding contractual agreements. Consider a provider that wishes to offer use of its resources as a pay-per-use service. Interactions between a provider and an end-user (a service consumer or a broker acting on their behalf) can then be governed through a Service Level Agreement (SLA), contractually defining the resource provider's obligations, the price the end-user must pay and the penalty the provider needs to pay in the event that it fails to fulfil its obligations. The use of SLAs to govern such interactions in Grid computing is gaining momentum (Wieder and Yahyapour, 2010; Battré et al., 2007; Waeldrich and Ziegler, 2010). However, such agreements represent a business risk to the parties involved. An SLA violation could be caused by various events such as a node outage or network failure. Consequently a provider may be unwilling to implement such an approach without effective risk assessment.

This paper focuses on a specific aspect of risk management as applied to Grid/Cloud computing: techniques that can be used by a resource provider to assess the risk of failure of resources within its infrastructure. This will enable a provider to identify infrastructure bottlenecks, evaluate the likelihood of an SLA violation and, where appropriate, mitigate potential risk, in some cases by identifying fault-tolerance mechanisms such as job migration to prevent SLA violations. A resource provider's reputation is closely related to the reliability of its product (here risk assessment). The more reliable the provider's risk assessment is, the more likely the provider is to have a favourable reputation.

A mathematical model for the prediction of the resources risk of failure is proposed with the use of a discrete-time analytical model driven by availability functions fitted to observed historical data.

* Corresponding author. Tel.: +44 1133436590; fax: +44 1133435468.
E-mail addresses: ralsoghayer@ksu.edu.sa (R. Alsoghayer), K.Djemame@leeds.ac.uk (K. Djemame).

0164-1212/$ – see front matter.
http://dx.doi.org/10.1016/j.jss.2013.09.017

This research has considered resource failures in Grid computing and the proposed mathematical model can equally be applied in a cloud environment. The main contributions of this paper are:

- A detailed analysis of Grid resource failures using failure data collected from different Grid resources and spanning for three years. The analysis focuses on the statistical properties of the failure data, including the root cause of failures, the mean time between failures, and the mean time to repair.
- A model to describe the time between failures in Grid resources, as well as a model for the time to repair a resource. Modelling failures and repairs are crucial in the design of reliable systems and also when creating realistic benchmarks and test-beds for reliability testing.
- A model to predict the Grid resources risk of failure, which can also be used to rank Grid resources.

The remainder of the paper is organised as follows. Section 2 introduces the risk management discipline. Section 3 explains the vision of risk in Grid computing. Section 4 provides an overview of Grid resources failures data along with the data-collection process and presents an analysis of such data. Section 5 presents the proposed model to predict the risk of failure of a Grid resource using a discrete time analytical approach driven by reliability functions fitted to observed failures data. Section 6 presents some related work and Section 7 ways to further extend this research. In conclusion, Section 8 provides a summary of the research.

2. Risk management

Risk management plays an important role in a wide range of fields, including statistics, economics, systems analysis, biology and operations research. The most central concepts in risk management are the following: an *asset* is something to which a party assigns value and hence for which the party requires protection. An *unwanted incident* is an event that harms or reduces the value of an asset. A *threat* is a potential cause of an unwanted incident whereas a *vulnerability* is a weakness, flaw or deficiency that opens for, or may be exploited by, a threat to cause harm to or reduce the value of an asset. Finally, *risk* is the likelihood of an unwanted incident and its consequence for a specific asset, and *risk level* is the level or value of a risk derived from its likelihood and consequence. For example, a server is an asset, a threat may be a computer virus, the vulnerability a virus protection not up to date, which leads to an unwanted incident: a hacker getting access to this server. The likelihood of the virus creating a back door to the server may be medium, but the integrity of the server (consequence in terms of harm) may be high.

As explained earlier, this paper focuses on a specific aspect of risk management as applied to Grid computing: methods that can be used by a resource provider to evaluate the risk of failure of Grid resources. In this context, a Grid resources is an asset, a threat may be a loss of its connectivity, the vulnerability a faulty hardware, which leads to an unwanted incident: the failure of the resource. The paper only focuses on the *likelihood* (probability) of Grid resource failures, and therefore uses the terms Probability of Failure and Risk of Failure interchangeably.

3. Risk aware Grid computing – the vision

The overall vision is the production of a risk aware decision support system allowing individuals to negotiate and consume Grid resources using Service Level Agreements (SLA). This embraces an extended approach to the utility computing business model, which fits in an open market business model (for example for access to compute power) as used in sectors such as finance, automotive, and energy. This section presents the main actors (end-user and resource provider), an example scenario in which they participate, and the resource provider architectural components for risk assessment.

3.1. Actors

An end-user is a participant from a broad public approaching the Grid in order to perform a task comprising of one or more services. The user must indicate the task and associated requirements formally within an SLA template. Based on this information, the end-user wishes to negotiate access with providers offering these services, in order that the task is completed. The end-user must make informed, risk-aware decisions on the SLA quotes it receives so that the decision is acceptable and balances cost, time and risk.

A provider offers access to resources and services through formal SLAs specifying risk, price and penalty. Providers need well-balanced infrastructures, so they can maximise the Quality of Service (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 prerequisite to this is a provider's trustworthiness and their ability to successfully deliver an agreed SLA. The assessment of risk allows the provider to selectively choose which SLA requests to accept.

Note the possible consideration of a broker in such context, which acts as a matchmaker between end-users and providers, furnishing a risk optimised assignment of SLA requests to SLA quotes (Djemame et al., 2013). It is responsible for matching SLA requests to resources and services, which may be operated by an arbitrary number of providers. The broker's goal is to drive this matchmaking process to a conclusion, when the provider makes an SLA offer.

3.2. Motivating scenario

Considering the situation where a provider wishes to offer use of its resources as a pay-per-use service to potential end-users, and where the use of SLAs govern the interaction between them, a provider may need to implement an effective risk assessment prior to making an SLA offer. In this case, the provider computes the risk of failure for each resource and subsequently allocates the resources to the end-user's job. If the resulted allocation fails to satisfy the end-user's requirements, the resource reservation is revisited; if it does satisfy the end-user's requirements, the resource provider then sends back the SLA offer, updated with cost/penalty fee and pre-commits the resources. The end-user either commits to the SLA or rejects it.

Fig. 1 provides an overview of the interaction of the provider infrastructure components. The user sends an SLA request to the provider specifying the job requirements (1). The provider's *Resource Manager* requests the *Reservation and Allocation* component to reserve the required resources (2). The Reservation and Allocation component reserves the physical resources (3) and passes to the *Risk Assessor* for each reserved resource the time and duration of the reservation (4). The Risk Assessor computes for each resource the risk of failure based on the resource historical information stored in the *Historical Database* (5). The *Monitoring* component is responsible for gathering all necessary runtime information that is collectable by sensors in the infrastructure. The Risk Assessor returns the risk of failure information to the Resource Manager (6). Finally, the Resource Manager sends a response back to the user (7), either in the form of an SLA offer or reject.