SERVICE LEVEL AGREEMENTS AS A SERVICE
Towards Security Risks Aware SLA Management

Katerina Stamou[1], Jean-Henry Morin[1], Benjamin Gateau[2] and Jocelyn Aubert[2]
[1]*Institute of Services Science, University of Geneva - CUI, Geneva, Switzerland*
[2]*Centre de Recherche Public Henri Tudor (CRPHT), Luxembourg City, Luxembourg*

Keywords: Cloud Computing, Service Level Agreement (SLA), Risk Management, Security, SLAaaS, Service Level Objective (SLO).

Abstract: Cloud computing has matured to become a valuable on demand alternative to traditional ownership models for the provisioning of services, platforms and infrastructure. However, this raises many issues for Governance, Risk and Compliance (GRC) and in particular in terms of Information Systems Security Risk Management (ISSRM). Considering such issues lack attention and knowledge, particularly for small and medium sized enterprises (SMEs), and that cloud computing Service Level Agreements (SLA) provide very limited support outside of basic Quality of Service (QoS) parameters, this paper argues that SLAs for cloud computing services should be more customer oriented and aware of security and risk management. A design is proposed where the SLA process, from context initialization to negotiation and agreement is decoupled from the actual cloud service provisioning and itself turned into a Service : SLA as a Service (SLAaaS). This should provide customers with much more customized and fine-grained agreements compared with the ones currently offered.

1 INTRODUCTION

The cloud computing model allows easy, on-demand, internet-based access to computing resources. NIST (Mell and Grance, 2011) defines the following types of cloud deployments: private, community, public, including hybrid combinations of these. Each of them involves different sharing modes (Jansen and Grance, 2011).

Provision of service offerings in cloud computing is instant, requires only network access and minimal client-provider interaction. Service level agreements (SLAs) represent a form of such interaction, as guarantees between a service provider (SP) and a service consumer.

SLAs provide assurances on the expected or mutually agreed service level on behalf of a SP to a service consumer (Dan et al., 2003). They define a multitude of business and technical Service Level Objectives (SLOs), including metrics that reflect such objectives.

Still, current public cloud service offerings lack the use of security risk management (RM), as well as governance and compliance aspects in their provided guarantees. We perceive RM as an important factor to assist the fine-tuning of an SLA according to customer defined objectives.

This paper advocates that SLAs for cloud computing services should be more customer centric and aware of customer risk requirements. We envisage a holistic SLA process from context initialization to negotiation and agreement as a service that can provide a customer with a much more tailored final agreement compared to the ones currently offered.

Section 2 presents the motivation behind this work and provides a high level summary of recent research projects on cloud computing dealing with SLA management. We have identified some representative issues regarding SLAs and their management. We also briefly discuss several derived SLA parameters and management aspects.

Section 3 proposes a high-level SLA decomposition and clause classification. This classification has helped conceptualize our approach and a model in a research project on Cloud Computing SLA Risk Management (CLOVIS, for which a first study was done in (Morin et al., 2012)).

Section 4 discusses the importance of combining RM with SLA management while decoupling it from the actual provisioning of the service. We propose a model and approach where SLA is provided as a Service therefore opening new opportunities for better

Stamou K., Morin J., Gateau B. and Aubert J..
SERVICE LEVEL AGREEMENTS AS A SERVICE - Towards Security Risks Aware SLA Management.
DOI: 10.5220/0003966006630669
In *Proceedings of the 2nd International Conference on Cloud Computing and Services Science* (CloudSecGov-2012), pages 663-669
ISBN: 978-989-8565-05-1

management of SLAs and consequently RM in cloud services. We argue, such an approach may be particularly useful for SMEs having limited resources and knowledge to assess such risks. We conclude with on-going and future work.

2 SLAS IN CLOUD COMPUTING

2.1 Motivation

Cloud computing SLAs do not lack technical objectives in terms of service levels. Undoubtedly, their management lacks several technical attributes such as machine-readability, automation etc. But most importantly, they lack risk management objectives and more generally GRC concerns that can be better addressed with a more customer aware orientation.

Research approaches on SLA management introduce important aspects like automation and dynamicity, yet are deprived of RM objectives that are often crucial to a customer's internal planning. Although commonly found SLA terms semantically relate to GRC management, it is not clear how they relate to SMEs' needs and internal assessments.

Due to lack of experience, many SMEs and start-up companies are not able to evaluate their risks and often do not have a risk management plan. Typically, their size and lack of financial strength prohibits them from doing so. Additionally, many SMEs primarily look into public cloud service offerings, since they cannot afford some other cloud deployment type.

Independently of the service layer, public clouds provide SLAs that typically do not address small organizations' security concerns. Thus, such agreements impose even greater risks for SMEs that do not have internal risk assessments. Moreover, public cloud SLAs do not adjust their pricing according to evaluation of specific risks, but apply massive, one-size-fits-all pricing models.

In (Potoczny-Jones, 2011) the offering of Security Risk Agreements (SRAs) is proposed. Cloud SPs offer such SRAs to SMEs customers as tailored SLAs that include all necessary security information regarding an offered service. An SME can then evaluate if the offered service reflects its security requirements or not. It is suggested that such agreements will increase the level of trust between SPs and customers, and consequently motivate them to adopt public cloud services. Moreover, in (Kaliski et al., 2010) they propose the idea of Risk Assessment as a Service.

There is no doubt that the higher the mutual trust between service provider and service consumer, the easier the adoption of new offered services.

2.2 Related Work

We have conducted a thorough literature review to investigate state of the art approaches related to SLA management for cloud computing services. Table 1 provides a high-level overview of recent research initiatives in cloud computing. Many initiatives are EU FP7 funded projects and, as Table 1 illustrates, are focused on different cloud layers, with the exception of (SLA@SOI, 2011) and (IRMOS, 2009) that follow a more vertical approach across all layers.

Table 1: Research projects focused on different cloud computing layers.

Project/cloud layer	IaaS	PaaS	SaaS
Optimis	x		
Contrail	x	x	
SLA@SOI	x	x	x
IRMOS	x	x	x
Cloud4SOA		x	
mOSAIC		x	
4CaaSt[1]		x	

Most studied SLA frameworks are platform specific, in that every research approach implements its own system architecture and accommodates software modules and tools. Such platforms typically include a multitude of software components to assist fine-grained interoperation of different services like workflow (job) monitoring, resource discovery and accounting. SLA management modules usually include active as well as template SLA repositories and communicate with user and task management components.

Automation of SLA management is a desirable feature. Still, there is the assumption that at some point human intervention is necessary. Depending on the platform, layer and application specifics, some approaches implement a single type of SLA templates, while others combine different types (IRMOS, 2009).

In CLOVIS, the goal is to define, design and develop a service-level framework that integrates RM and SLA capabilities. The focus is to improve governance, risk management and compliance (GRC) aspects for cloud computing services while introducing enhanced levels of managed flexibility through exception management (Morin, 2008) features.

2.3 SLA Issues

Based on our review of related work, we have tried to summarize important problems of currently provided

[1]Morfeo 4CaaSt, http://4caast.morfeo-project.org

SLAs. Undoubtedly, wide adoption of the cloud computing model and increased interest about its offerings by companies worldwide is leading to novel service oriented economic models, that require more efficient, flexible and automated SLA management (Butler et al., 2011).

Traditionally, SLAs are defined exclusively by SPs. The same SLA template is offered to all customers of the same service, regardless of the specific needs of customers. Such SLA approaches tend to ignore that each customer weighs SLA defined metrics differently.

Offered SLAs focus on service level objectives (SLOs) that deal with technical guarantees like uptime, bandwidth availability, etc. Such SLOs are and should be included in any service level agreement, since they constitute the very basics of such an agreement. On the other hand, such SLAs are monolithic in the sense that no customer specific needs are taken into account and their customization depends on manual and time-consuming interventions between the two involved parties.

In public clouds, SLAs do not address properly important security aspects such as data privacy and clearly lack negotiation opportunities between SPs and customers. Provided SLAs are usually in non machine-readable formats and are delivered as static agreements, not considering dynamic demand on service levels.

Moreover, it is commonly accepted that there is a lack of unified standards for public cloud computing services (mOSAIC, 2011) and (Mell and Grance, 2009). The definition of standards can be very helpful for the establishment of common requirements and provide an extra safety layer for SMEs.

Research efforts on SLA management for cloud computing introduce important aspects like automation and dynamicity. They lack however risk management objectives that are crucial to a customer's requirements.

2.4 Important Aspects Regarding Context and Management of SLAs

Our review of related cloud computing initiatives has helped in aggregating important points considering parameters and management of SLAs:

Machine-readable SLAs. Typically, most cloud offerings for public cloud services offer static, non machine-readable SLAs. All research approaches indicate the necessity to offer machine-readable agreements. This feature allows for automated negotiation modules (SLA@SOI, 2011) as well as translation of high-level QoS parameters into low-level ones (IRMOS, 2009).

Digitalization of SLAs construction and management is a mature research topic. IBM's research on utility computing resulted in the definition of the (WSLA, 2003) specification language. The grid computing community has also contributed much into SLAs initiation and management for distributed infrastructures. The definition of the WS-Agreement specification language (Andrieux et al., 2005) as well as the implementation of protocols related to SLA negotiation, resource allocation and monitoring (Czajkowski et al., 2002) designate some grid-oriented accomplishments that are currently reused by many cloud computing activities.

Automation. Many research efforts highlight the need to automate as many SLA management parameters as possible. Such automation is necessary to allow for vertical or multi-layered SLA management (SLA@SOI, 2011). Moreover, it simplifies the interaction of an SLA module with other system components (monitoring, accounting etc) for the better orchestration of the whole SLA life cycle (Cloud4SOA, 2011).

Risk. (Optimis, 2011) refers to risk as a central concept that has to be taken into account throughout the full life cycle of the service. The Cloud Security Alliance provides a GRC Stack toolkit (CSA, 2010) that integrates their initiatives for efficient cloud auditing, security assessments and critical compliance requirements.

Many SLA defined terms and parameters are derived, at least semantically, from RM techniques and objectives.

SLA Management Activity Duration. Proposed SLA frameworks and design approaches suggest that activities related to SLA management should follow a whole service life cycle. This feature enables dynamic management of SLAs by allowing their possible modification during the run time of a service (mOSAIC, 2011). Dynamic deviation of client demands justifies this requirement. Moreover, SLA active duration affects other SLA management parameters like pricing, a customer's risk management planning, trust, etc.

Negotiation. Negotiation is a process that typically takes place prior to the run time activity of the service. (Optimis, 2011) highlights the need to negotiate SLAs with different autonomous providers. (IRMOS, 2009) defines SLA renegotiation as negotiation during the execution time of a service and describes it as a process that takes place in order to solve issues about to cause or having caused SLA violations.

Customization of SLA Related Terms and Parameters. The (mOSAIC, 2011) project advocates the need for user-centric, customer specific SLAs. Moreover, (Contrail, 2011) categorizes SLA related terms into quality of service and quality of protection parameters. Such approaches indicate the need for differentiated SLA offerings, depending on the specific needs of a company, customized to express useful metrics.

Geographic Location, Inter/National Laws. The geographic location of data-centres in relation to the national and/or international laws is an important criterion for many companies that want to invest into cloud based solutions. It is commonly accepted that such options should be included in, if not determine, an SLA. As highlighted by NIST in (Mell and Grance, 2009) and other initiatives, currently there is a lack of unified standards for cloud computing. The establishment of standards may assist in dealing with geographic and legally binding parameters.

Digital Signatures. In (IRMOS, 2009) they refer to the eligibility of signing contracts using digital signatures. They describe how they added signatures as an extension to an agreement's structure, which is based on the WS-Agreement specification (Andrieux et al., 2005).

Regulation of Penalties. In (Contrail, 2011) they investigate a possible differentiation of penalties, based on partial or complete fulfillment of SLA objectives. They discuss penalties or rewards for violating or satisfying an objective. Penalties can be imposed in the form of monetary compensation that is the current case in most public cloud service offerings or by means of reputation track mechanisms (Rana et al., 2008).

3 SLA DECOMPOSITION AND CLAUSE CLASSIFICATION

In SLA Management Handbook Volume 2 (ITU, 2005), the authors define a method for classifying service parameters to be used within an SLA. They classify parameters based on technology specific, service specific and technology/service independent parameters. Motivated by this methodology we have tried to classify SLA clauses and context parameters by decomposing an abstract SLA schema.

As core SLA clauses we have selected to use main risk categories as identified in (ENISA, 2009). Such a decomposition is very helpful to understand where and how the RM factor can possibly be integrated into an SLA. Additionally, it eases the process of recognizing particular needs and requirements that the cloud computing model imposes. Moreover this decomposition is quite practical for the realization of the SLA model in CLOVIS.

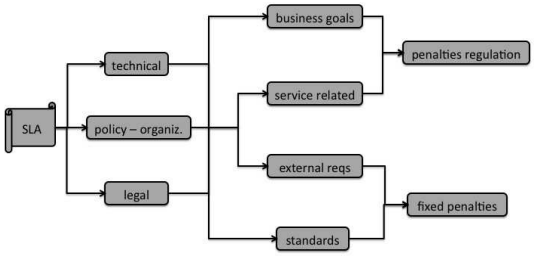

Figure 1: SLA clause classification.

Initially, our conceptual schema organizes SLA clauses into technical, legal and policy organizational elements. All such categories include sub-clauses that are:

- derived from business opportunities
- service related or
- defined by external requirements and standards

A subsequent layer follows where penalties are defined and regulated within an SLA.

This schema allows better mapping of companies identified risks and easier service level updates –if updates are allowed– during an agreement life cycle. The holistic SLA process from construction to initialization and management can be viewed as a modular procedure that introduces more customer oriented characteristics.

The dynamic nature of the cloud computing model substantiates the assumption that there is a tremendous need for more customer oriented, risk management aware SLAs, combined with automated and more flexible characteristics. The proposed decomposition illustrates a high level path towards the first part of such goals.

4 ISSRM COUPLED WITH SLA MANAGEMENT AS A SERVICE

4.1 Proposed Design

A risk management tool coupled with an SLA management framework is likely to provide a valuable solution for both customers and SPs. The SLA clause decomposition and arrangement of Section 3 has inspired an abstract model that is illustrated in Figure 2. Both RM and SLA management are to be considered as provided "as a Service", hence the RMaaS and SLAaaS.

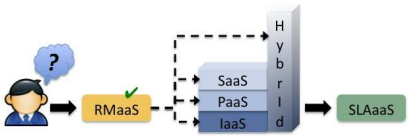

Figure 2: RM and SLA coupling.

Decoupling these components from the actual cloud service provisioning as shown in Figure 3 has many advantages. First and foremost it accounts for a much needed separation of concern between service provisioning and the conditions under which the service is to be provided including monitoring. In doing so, SLA management could be considered as a Trusted Third Party (TTP) between SPs and customers in a similar way Certification Authorities are now well recognized in similar trust roles. Moreover, this may open a whole new area of expertise allowing SLAaaS providers to become new marketplaces for SLA management. One could even imagine industry or service type specific SLA providers with templates tuned to specific needs. For example, a company working in the finance industry and requiring a cloud storage service might be subject to specific rules with respect to where the data physically resides. In such a case, a Financial Services SLAaaS provider could be a good choice for the company and additional peace of mind when contracting cloud based storage services.

Figure 3: SLAaas decoupling from actual cloud service.

Our model assumes the existence of an appropriate RM tool, which can provide a customer with an adequate risk assessment according to a customer's service criteria. Theoretically, such criteria may describe service preferences from a single cloud layer or from hybrid service combinations. The derived assessment can subsequently be fed into a semi-automated process that will map such information into concrete SLOs.

It is not easy to automate a risk assessment process, since every company and potential SLA customer has unique and rather subjective requirements regarding its risks and internal needs. Still, it is possible to use the output of a proper risk assessment process and treat it as input for the orchestration of a customer oriented SLA.

4.2 Service Usage Scenario

Let's now consider the following basic service usage scenario, described in abstract terms:

A customer (e.g., an SME), with or without an internal ISSRM plan, wants to invest into a suitable IaaS solution. The customer is provided with a RM tool that can help towards risk assessment.

The customer submits its risk assessment in raw format along with its service preferences for the desired service level guarantee (what type of service, duration, IaaS common technical requirements). Our SLAaaS framework receives as input the raw format data and produces as output structured information data in an interoperable format.

Such data represent the initial assessment, contextualised into SLA clauses. These clauses are depicted as SLOs. In (ITU, 2005) SLOs are defined as internal forms of SLAs that exist between business functions. SLOs are less formal than SLAs.

Our output is not intended to produce a concrete SLA, but rather context clauses that can directly be integrated in or affect a final SLA output. It is merely an attempt to assist the process of producing RM aware and more customer oriented SLAs.

The processed output may then submitted to an open registry that is accessible by SPs. SPs evaluate the submitted objectives, adjust their own SLOs based on their SLA templates and send their offers. Sent offers are displayed back to the customer. The customer may select or reject among the displayed offers.

4.3 Discussion

The Service usage scenario described above, assumes several technical requirements. First of all, machine readability of all involved information is an essential prerequisite. Information should be processed and utilized in a well-used, interoperable format. A processed SLO output should be able to fit in most SLA management frameworks, be editable and easily updated.

In our current design, we are focused on the IaaS layer. An initial part of our work will involve adapting a selected RM tool into IaaS related risks. The report provided in (ENISA, 2009) will support this effort.

Our hosting platform will include a web interface for the customer (user) - service interaction as well as a negotiation module for a possible customer-provider interaction. Still, the envisioned service has also to

interact with several modules that are running on different platforms and assist services like monitoring, accounting, etc.

Additionally, security is a major concern for our framework since it has to process, move and regularly update context sensitive, private data. Thus, our design will include the orchestration of the necessary mechanisms to ensure data confidentiality and secure service transactions.

Moreover, a generalization of our envisaged service equally applies to services of all cloud layers and any hybrid combinations. We anticipate that the realization of our model at the IaaS layer will generate useful test cases and results that can be extended to any cloud service approaches.

5 FUTURE WORK

Our tentative roadmap attempts to couple RM with SLA management for offered cloud services. We plan to validate our results using real case studies in collaboration with vendors who provide business resilience solutions utilizing holistically the cloud computing model. As mentioned, our produced case studies will initially be applied to the IaaS layer.

The envisioned framework has to provide SLA management capabilities as well as be compatible with and pluggable into external platforms that may integrate their own SLA management framework. Compatibility will allow feeding of existing SLA frameworks with RM parameters and metrics, so that provided SLAs depict more precisely customers' needs.

The produced SLA will be customer centric, taking into account and adjusting service level and pricing parameters to a company's identified risks and security concerns. This may lead to public cloud service offerings that capture more thoroughly customers objectives while also being compliant with standards.

Currently the design process of our prospective framework is initiated. We are in the process of selecting and adjusting an appropriate RM tool to our service scope. The next phase of our work will include the definition of a method that receives as input a risk assessment output and maps it into potential SLA clauses in a semi-automated manner.

Throughout this procedure we will use and refine the SLA decomposition schema of Section 3. Such an arrangement can help with the desired SLA context mapping as indicator of how to integrate and adapt given risk assessment results. Additionally, we will have to ensure that our method will allow for frequent iterations and updates of the generated SLOs, since the RM process is by essence dynamic.

The service we are building could possibly be classified at the SaaS layer, since an end user would ideally interact with the service through a web interface. Still, inevitably the success and viability of our proposed service depends on orchestrating components and modules from lower layers of the cloud computing stack.

6 CONCLUSIONS

We have provided an SLA decomposition and clause classification schema utilizing a method for service parameter classification that is provided in (ITU, 2005) and identified risk categories as identified by (ENISA, 2009). Although abstract, such a classification may contribute as a first step towards a unified conception on SLAs construction, initiation and management for RM aware SLAs. Our assumption is that more customer oriented, risk aware SLAs will assist the integration of SMEs into the emerging cloud computing economy.

Security concerns and often absence or negligence of risk management operations are some of the major obstacles that currently prohibit many SMEs from engaging in and using cloud computing services. Existing SLAs do not address such concerns. Consequently, SMEs cannot estimate precisely their needs regarding service level expectations and service providers' guarantees.

We have described our proposed design for the SLAaaS framework. It will exhibit both risk and SLA management attributes. A service-oriented framework that addresses risk and generally GRC aspects can lead to efficient, more flexible SLAs and be highly valuable for many potential cloud service customers. Finally, it can also be a very useful tool for service providers to better adjust their pricing models and offer negotiable, customer centric services.

ACKNOWLEDGEMENTS

This work is supported by the CLOVIS project jointly funded by the Swiss SNF and Luxembourg FNR Lead Agency agreement; under Swiss National Science Foundation grant number 200021E-136316/1 and Luxembourg National Research Fund (FNR) grant number INTER/SNSF/10/02. K.Stamou would like to thank N.Mayer for external discussions and valuable input in this effort.

REFERENCES

Andrieux, A., Czajkowski, K., Dan, A., Keahey, K., Ludwig, H., Nakata, T., Pruyne, J., Rofrano, J., Tuecke, S., and Xu, M. (2005). Web Services Agreement Specification (WS-Agreement). Available at http://mailman.ogf.org/documents/GFD.107.pdf.

Butler, J., Lambea, J., Nolan, M., Theilmann, W., Torelli, F., Yahyapour, R., Chiasera, A., and Pistore, M. (2011). SLAs Empowering Services in the Future Internet. In *The Future Internet*, volume 6656 of *Lecture Notes in Computer Science*, pages 327–338. Springer Berlin / Heidelberg.

Cloud4SOA (June 2011). D1.3 Reference Architecture. Technical report, EU FP7. Available at http://www.cloud4soa.eu, accessed October 2011.

Contrail (2011). D3.2 SLA Management Services Terms and Initial Architecture. Technical report, EU FP7. Available at http://contrail-project.eu, accessed October 2011.

CSA (2010). GRC Stack, an Integrated Suite of Four CSA Initiatives. Available at https://cloudsecurityalliance.org/research/initiatives/grc-stack, accessed November 2011.

Czajkowski, K., Foster, I., Kesselman, C., Sander, V., and Tuecke, S. (2002). SNAP: A protocol for negotiating service level agreements and coordinating resource management in distributed systems. volume 2537 of *Job scheduling strategies for parallel processing*, pages 153–183. Springer.

Dan, A., Ludwig, H., and Pacifici, G. (2003). Web service differentiation with service level agreements. *White Paper, IBM Corporation*. Available at http://www.ibm.com/developerworks/library/ws-slafram.

ENISA (2009). Cloud computing: Benefits, Risks and Recommendations for Information Security. Technical report. Available at http://www.enisa.europa.eu/act/rm/files/deliverables/cloud-computing-risk-assessment, accessed October 2011.

IRMOS (2009). D7.2.1 Initial version of Path Manager Architecture and Guaranteeing QoS with Dynamic and Automated SLAs in real-time aware SOIs. Technical report, EU FP7. Available at http://www.irmosproject.eu, accessed October 2011.

ITU (2005). *SLA Management Handbook, Concepts and Principles*, volume 2.0. TeleManagement Forum. Release 2.5, GB 917-2.

Jansen, W. and Grance, T. (2011). Guidelines on Security and Privacy in Public Cloud Computing. Available at http://csrc.nist.gov/publications/nistpubs/800-144/ SP800-144.pdf, accessed December 2011.

Kaliski, J., Burton, S., and Pauley, W. (2010). Toward Risk Assessment as a Service in Cloud Environments. In *Proceedings of the 2nd USENIX conference on Hot topics in cloud computing*, HotCloud'10, pages 13–13, Berkeley, CA, USA. USENIX Association.

Mell, P. and Grance, T. (2009). Effectively and Securely Using the Cloud computing Paradigm.

Mell, P. and Grance, T. (2011). The NIST Definition of Cloud Computing (Draft) Recommendations of the National Institute of Standards and Technology. *NIST Special Publication*, 145(6):7.

Morin, J.-H. (2008). Exception Based Enterprise Rights Management: Towards a Paradigm Shift in Information Security and Policy Management. volume 1 of *International Journal On Advances in Systems and Measurements*, pages 40–49.

Morin, J.-H., Aubert, J., and Gateau, B. (2012). Towards Cloud computing SLA Risk Management: Issues and Challenges.

mOSAIC (February 2011). D1.1 Architectural design of mOSAIC's API and platform. Technical report, EU FP7. Available at http://www.mosaic-cloud.eu, accessed October 2011.

Optimis (2011). OPTIMIS SLA Framework and Term Languages for SLAs in Cloud Environment. Technical report, EU FP7. Available at http://www.optimis-project.eu, accessed October 2011.

Potoczny-Jones, I. (2011). Cloud Security Risk Agreements for Small Businesses. Available at http://corp.galois.com/blog/2011/8/23/cloud-security-risk-agreements-for-small-businesses.html, accessed November 2011.

Rana, O. F., Warnier, M., Quillinan, T. B., and Brazier, F. M. T. (2008). Monitoring and Reputation Mechanisms for Service Level Agreements. Grid Economics and Business Models GECON, pages 125–139. Springer.

SLA@SOI (July 2011). D.A1a Reference Architecture for an SLA Management Framework. Technical report, EU FP7. Available at http://sla-at-soi.eu, accessed October 2011.

WSLA (2003). WSLA Language Specification. Technical report, IBM Corporation. Available at www.research.ibm.com/wsla/WSLASpecV1-20030128.pdf, accessed October 2011.