A Framework for Failure Impact Analysis and Recovery with Respect to Service Level Agreements

Andreas Hanemann, David Schmitz
Munich Network Management Team
Leibniz Supercomputing Center
Barer Str. 21, D-80333 Munich, Germany
{hanemann,schmitz}@lrz.de

Martin Sailer
Munich Network Management Team
University of Munich (LMU)
Oettingenstr. 67, D-80538 Munich, Germany
sailer@informatik.uni-muenchen.de

Abstract

In today's IT service market customers urge providers to grant guarantees for quality of service (QoS) which are laid down in Service Level Agreements (SLAs). To satisfy customers and to avoid penalties, service providers have to ensure that the agreed SLAs are met. Therefore, it is necessary to be able to effectively deal with resource failures which could endanger the SLAs by affecting the provided services. The effort for recovering from failures should be selected corresponding to the expected SLA violation costs.

In this paper we present a framework to automatically determine the impact of resource failures with respect to services and service level agreements. We achieve this by monitoring the service quality from inside and outside the service provider and also by incorporating information about the current and expected future service usage. The expected costs of the resource failures are assessed to select an appropriate recovery alternative. Besides this short term perspective the impact analysis can also be employed to identify critical resources and to improve the service provisioning.

1 Introduction

In today's IT environments provisioning of services with guaranteed QoS has become a crucial point for IT service providers. Thus, it is an important problem for a provider to ensure that the agreements with their customers are met. As the Mean Time between Failures (MTBF) and the Mean Time to Repair (MTTR) are often part of the Service Level Agreements (SLAs), it has become vital for a provider to react quickly and accurately when detecting a failure in a resource.

To be able to react in this way, a holistic view of the inner structure of service provisioning is required. This means to have knowledge about the dependencies of the offered services on subservices and resources as well as the customers' SLAs, their QoS parameters, and the current service usage. Appropriate recovery actions can then be chosen by applying this knowledge.

Today, service providers often achieve this by relying on the experience of their employees, which has several drawbacks. Important influence factors are likely to be left out in the decision making process or employees can leave the company which will lead to an information loss.

Therefore, we posit a modeling framework that formalizes the mentioned coherences and automates the decision procedure. This framework shall be applicable with respect to different time horizons. In a short-term perspective, single or multiple failures which are currently present in the network shall be treated. It has to be determined which services and SLAs based on these services are affected by the failures. The impact has to be evaluated especially with respect to possible SLA violations. As a consequence, recovery alternatives have to be identified and a decision support should be given to the operation staff. For mid-term considerations the service provider may want to simulate what would happen in case of resource failures. This is useful to identify critical resources and therefore optimize the way services are provisioned. On the long run, the service provider may like to change his offerings in order to react to customer demands. The framework should in this case be helpful to know whether the current IT infrastructure is suitable for those newly introduced services. It might also be interesting to think about a different pricing model to e.g. avoid utilization peaks.

In this paper we define a *service* as a set of *functionalities* which are offered by a *provider* to a *customer* at a *customer provider interface*. The definition of a "service" is therefore more general than the definition of a "Web Service", but a "Web Service" is included in this "service" definition. As a consequence, the results are applicable for Web Services as well as for other kinds of services. An *SLA* is defined as a contract between customer and provider about guaranteed service performance.

This paper is organized as follows. In Section 2 we present a scenario for the provisioning of an IT service. It is used to motivate the requirements for the framework to perform the impact analysis and fault recovery. Related work covering different aspects of the requirements is outlined in Section 3. Our modeling framework is presented in Section 4. It includes a workflow corresponding to the framework which is illustrated using the scenario. The last section concludes the paper and presents future work.

2 E-Mail Service Scenario and Requirement Analysis

The E-Mail Service offered by the Leibniz Supercomputing Center (LRZ) serves as an example IT service. The LRZ, which is the computing center for the Munich universities and runs the scientific network in Munich, offers e-mail access for students and staff of the universities and the LRZ itself. In this E-Mail Service Scenario the impact of resource failures for different kinds of QoS parameters shall be examined.

Besides more general QoS parameters such as the availability, there also exist parameters specific to this service. An example is the delivery delay which could be defined for incoming and outgoing mail. Such delays may be caused by long mail queues occurring at the LRZ mail server. It is obvious that an agreement for this parameter will only cover the LRZ domain and will therefore not be end-to-end in the first place.

The E-Mail Service Scenario is shown in Figure 1. The service is offered by the LRZ (provider) to its customers at the customer provider interface. A customer can allow several users to use the service by granting them to create their personal e-mail accounts. The quality and cost issues of the service are laid down in SLAs. On the provider side, the E-Mail Service is provided using subservices. In the scenario these subservices are DNS, proxy service, connectivity service (IP), and storage service. Both services and subservices depend on resources which they are provisioned upon. These are e.g., network components, network links, an end system's main memory, or processes running on a server. As depicted in Figure 1 a service can depend on more than one resource and a resource can be used by one or more services.

As the LRZ would like to ensure that the agreed SLAs are met, it is necessary to react accurately to faults occurring in one or more resources. To achieve this, an impact analysis has to be performed where different kinds of dependencies are used to identify affected services and corresponding SLAs. Furthermore, a methodology has to be defined to decide which recovery steps should be carried out in order to deal with faults.

The following issues need to be addressed which are motivated by the scenario. The requirement number refers to locations in the Figure. Our previous work [16] contains details about these requirements.

A workflow (requirement 1) has to be defined to identify steps needed during failure impact analysis and recovery. This workflow e.g. has to comprise the mapping of resource problems onto services and the information of customers about the current service status.

A service model (requirement 2) is needed covering the service features relevant for the impact analysis, especially the QoS parameters. The QoS modeling should be independent from the provider's service implementation. This is a requirement made by customers in order to be able to compare the offers of different providers. In addition to the service modeling, an appropriate resource modeling (e.g., failure states, possible recovery measures) is also required.

In the scenario, there are three kinds of dependencies, i.e., dependencies between different services, dependencies between services and resources, and dependencies on the resource level. It is important to identify the characteristics of these dependencies and their necessary attributes (requirement 3). An example for this is an appropriate modeling of resource redundancies.

As the impact analysis is performed with respect to SLAs, an SLA modeling is needed (requirement 4) based on the QoS modeling mentioned above. A monitoring infrastructure appropriate for the SLA/QoS definition has to be in place.

After the impact of a resource failure has been determined, recovery measures have to be performed. The ways a provider can react to a certain situation have to be modeled depending on the kind of resource failure (requirement 5).

Besides these requirements, the impact analysis should be integrated with the service-oriented event correlation, which we proposed in [15, 14]. Event correlation as used today typically deals with events on the resource level (e.g., link up/down, authentication process crashed). In our approach we extend the correlation by integrating events about service problems. These events can either be generated from customer reports about service malfunctions or by the provider's own service monitoring. The dependency modeling developed for the impact analysis should be designed in a way that it can easily be used or adapted for the service-oriented event correlation. The output of the service-oriented event correlation, i.e. resource failures, can be utilized as input for the impact analysis. Therefore, the modeling of resources (and especially the possible resource failures) should be done with respect to service-oriented event correlation. In addition, it could also make sense to

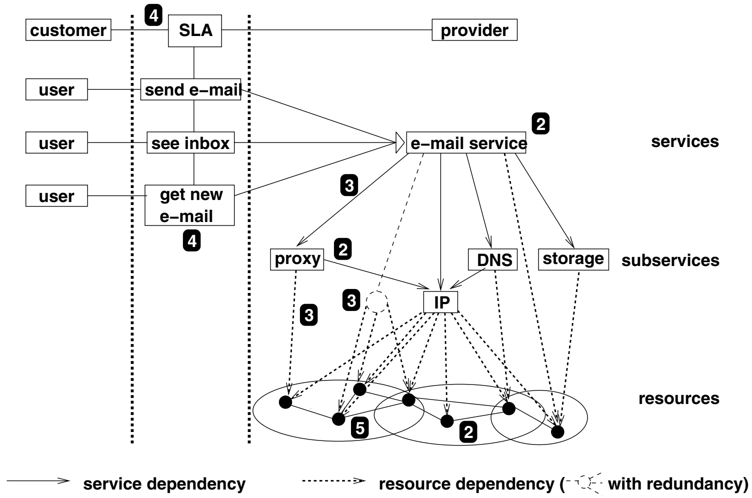

Figure 1. E-Mail Service scenario

transfer other information gained during the event correlation (e.g., other affected resources, affected services) to the impact analysis.

3 Related Work

To our knowledge, no approaches that can be found in the literature cover all parts of the requirements that were derived in the previous section. In the following the approaches are grouped according to these challenges. At first, work related to process frameworks and the modeling of services and resources is discussed. Approaches which either deal with the modeling and/or finding of dependencies or with SLA monitoring and management are presented subsequently. The Management by Contract approach which covers the decision making between different recovery actions is also reviewed. Finally, a conclusion is drawn to identify approaches which serve as input for the framework proposed in Section 4.

3.1 Process Frameworks

The IT Infrastructure Library (ITIL) [22] is a continuously evolving collection of best practice documents with regard to the service management of an IT service provider. It defines (among others) the process sets of service support and service delivery. Process descriptions are derived from expert knowledge in a particular field and written more or less in prose, so that ITIL can be regarded to be a bottom-up approach. As a consequence, the workflow modeling for our framework cannot be directly derived from these process descriptions.

The enhanced Telecom Operations Map (eTOM) [11] published by the TeleManagement Forum in 2002 is a business process framework for the telecommunications industry. eTOM is customer-centric and covers a broad range of important processes including processes for strategy, infrastructure, product, and operations. The Service Problem Management (SM&O-A) process deals with the diagnosis and resolution of service problems including assessment of the impact on customers. However, this process is not described in a formal way, and neither input and output parameters, nor the linking of processes are described explicitly.

3.2 Service and Resource Modeling

In order to establish a common understanding of the term "service" a generic service management model (MNM Service Model) has been proposed by our research group [9]. Therein a service is defined as a set of functionalities that are offered by a provider to a customer with an agreed quality of service. The customer can allow users to access the service at the service access point. An important feature of the model is that service management is integrated into the model. Similar to the service access point a reference point is defined for the exchange of management information between customer and provider (e.g., for ordering of new services, provisioning of service quality reports, information about service failures). It is called Customer Service Management (CSM, [20]) access point.

A detailed presentation and classification of QoS approaches can be found in [10]. The QoS definition and measurement methodology proposed by Garschhammer [8] addresses the issue of implementation independent QoS specification. Therefore, the measurement operations are performed at the service access point independent of the inner structure of the service provisioning.

The Common Information Model (CIM,[5]) introduces a management information model that aims at integrating information models of existing management architectures. CIM acts as an umbrella that allows to exchange management information in an unrestricted and loss-free way. This umbrella architecture was chosen to achieve vendor independence. As a consequence, CIM could be used for the resource modeling.

3.3 Dependency Modeling and Finding

In the literature the term "dependency" is often used without a precise definition. In the dependency graph which was proposed by Gruschke [12] for event correlation a dependency is a relationship between different entities. As the approach was designed to be as generic as possible, it can be used for different abstraction levels (e.g., service level, system level, network level). The term "dependency" was not further specified. A differentiation is performed between absent, weak, medium, and strong dependencies by Bagchi et al. [1], but no methodology how to assign such values is given. In the work performed by Kaiser [23] the availability of a service is defined in detail and a corresponding measurement methodology is proposed. A distinction is made between "and" and "or" dependencies, but their characteristics are not specified any further.

Caswell and Ramanathan [4] describe dependencies for services especially for Internet Service Providers. They define five kinds of dependencies:

Execution dependency: Performance of an application server process depends on the status of the host.

Link dependency: Performance of a service depends on the link status.

Component dependency: In case of a web service that is provided on different front-end servers which are selected by a round-robin DNS scheduling the performance depends on the currently selected server.

Inter-service dependency: This type of dependency occurs between services, e.g. e-mail service depends on an authentication service and on an NFS service.

Organizational dependency: Services and/or server may be mapped to different domains of responsibility.

The model has similarities with our scenario. We employ this model's understanding of inter-service dependency when we refer to service dependency. Also, the link and component dependencies as well as the execution dependency are called resource dependencies in our terms. The organizational dependency is not part of our model as a dependency, but it could be covered by the definition of subservices.

While our starting point in this paper is that the dependencies are given, the issue of finding them has also to be addressed. This knowledge is gathered from experts or configuration databases or log files, etc. As changes in the service provisioning are quite frequent, approaches to automatically detect dependencies have been proposed. Ensel [6] proposed the use of neural networks for the dependency detection, while the approach of Gupta et al. [13] analyzes temporal relationships of interactions to derive dependencies.

3.4 Service Level Monitoring and Management

Service Level Monitoring approaches and tools are used to monitor whether an SLA is met, but they do not deal with the treatment of faults. In the SoLOMon framework [7] a language is defined to specify metrics in an expressive way. The metrics are therefore user-oriented and independent of a specific application. A run-time system was implemented for those metrics which especially aims at achieving scalability. Several commercial tools like Infovista[18] are not only able to monitor the network and systems performance, but can also be used to monitor the service performance.

An overview of the issues related to Service Level Management can be found in [21]. Problems arising when dealing with SLAs across domain borders are addressed in [3]. This publication also contains a language to define SLAs. A specification of SLAs for Web Services can be found in [19].

3.5 Failure Recovery Modeling

Salle and Bartolini [25, 2] approached the management of SLAs from a business perspective (called "Management by Contract"). While other approaches have in focus how to meet service level agreements, the possibility to break an SLA voluntarily is seen as a viable option. A modeling of SLAs and an algorithm to decide which effort should be applied to meet an endangered agreement is presented. A formalization of the cost of violating the agreement is needed as input.

While this approach seems to be appropriate for the modeling of SLAs and to select one of different recovery measures, other issues have not been addressed in a general manner so far. These issues are e.g., the modeling of services and resources, the way the impact analysis of a resource failure should be performed, methods to find possible recovery measures, and how to define costs for not meeting SLAs (it is often not sufficient to consider only the costs defined in the SLAs, but also long-term effects like unsatisfied customers terminating their contracts have to be taken into account).

3.6 Assessment of the Approaches

The analysis of the existing approaches revealed that no solution covers all requirements that have been identified in Section 2. Nevertheless, the MNM Service Model and CIM can be used for service and resource modeling. For QoS definition and measurement the approach from Garschhammer yields the abstraction level appropriate for our framework. In terms of dependency modeling none of the models provides the desired features. Such a modeling is going to be contained in a representation of information necessary for services called "Service MIB" which we are currently working on [24]. The SLA definition and algorithm to select a recovery measure from the Management by Contract approach are suitable for our framework.

4 Approach for Impact Analysis

In this section our approach for performing the impact analysis is presented. Driven by the identified requirements, subtasks which need to be performed during the impact analysis are defined and appropriate components are introduced. For some components there are already existing approaches as mentioned before, while others will have to be addressed in detail in future work. A workflow (*requirement 1*) is derived for the interaction of the components.

4.1 Impact Analysis and Recovery Framework

The framework that we propose for the impact analysis is depicted in Figure 2. The main components (gray boxes) communicate with each other following the workflow described in 4.2. Additionally, these components access databases or repositories to retrieve information they require to perform their tasks (as indicated by dashed arrows).

Network and systems management: The network and systems management component stands for a management system like HP OpenView or IBM Tivoli. It has access to the dependencies between resources which are stored in the network topology database and in the systems configurations database. These dependencies are traversed to identify other affected resources in case of one or more resource failures (*requirement 2, 3*).

Service management: The service management component has access to the dependencies on the service level and between services and resources. This information is contained in a repository called "Service MIB". By traversing the dependencies it is possible to identify services which are affected by resource failures. At this stage it would be possible to draw conclusions regarding the service quality, but this QoS would not be implementation independent as demanded in the requirements. Therefore, a QoS measurement component is introduced that provides the desired functionality (*requirement 2, 3*).

QoS measurement: In this component the QoS measurement is performed by intercepting interactions at the service access point [8] and can therefore be regarded as being implementation independent and customer-oriented (*requirement 2*).

SLA verificator: The SLA verificator receives information about services whose quality is endangered. To verify whether an SLA violation is likely to happen, knowledge about the customers' SLAs for these services is needed which is in turn stored in an SLA database. This database contains the agreements between provider and customers and also information about the current SLA status. In addition, the QoS monitoring information is used to determine whether the agreed service quality is currently met. To calculate the expected costs for SLA violations the current service usage can also be taken into account. The output of the SLA verification, i.e. the expected costs for SLA violations are transferred to the recovery management. The CSM also obtains information to notify affected customers about the current service status and SLA violations (*requirement 4*).

Service usage monitoring and prediction: This component monitors the current service usage and forwards this information to the SLA verificator. Depending on the service and the provider's experience with respect to the customers' behavior, prediction models could be applied to forecast the service usage during the expected repair time. Therefore, a service usage history is stored in the service usage database and analyzed for the prediction (*requirement 4*).

Recovery management: The recovery management component receives the expected cost for the current resource failures. Depending on the kind of resource failures different possibilities to recover from the fault may be known. The information about possible recovery measures is stored in a recovery actions repository. The recovery measure with the expected minimum cost calculated as sum of repair costs and SLA violation costs is selected (*requirement 5*).

Customer service management: CSM receives information about the service and SLA status from the SLA verificator and information about the recovery measures performed by the recovery management. De-

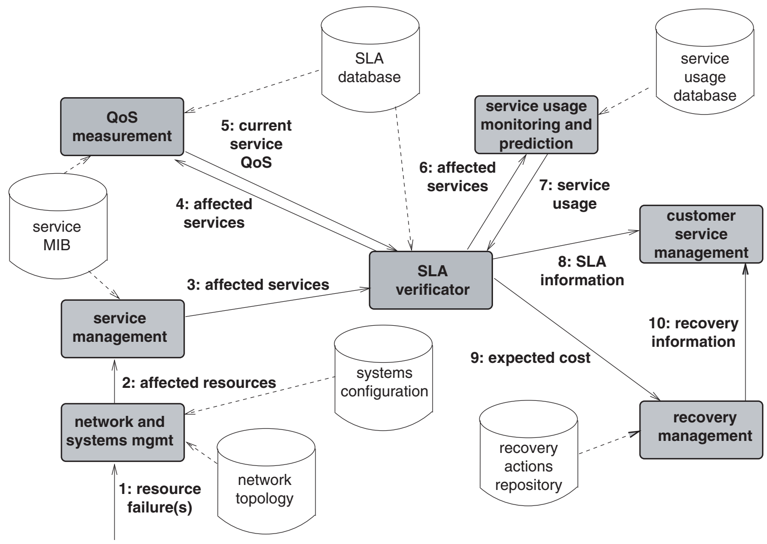

Figure 2. Impact analysis framework and workflow

pending on the provider's customer relationship management policy, this information is provided to the customers via CSM (*requirement 1*).

4.2 Impact Analysis Workflow

The solid arrows in Figure 2 indicate our proposal for a general workflow to perform the impact analysis and fault recovery. At first, the network and systems management receives one or more resource failures as input (*step 1*).

Figure 3 shows the resources which are used for the provisioning of the E-Mail Service. For load sharing and redundancy reasons the dispatching of mails is performed by two different mail relays controlled by a load balancer. Mails received can be accessed from different incoming mail servers depending on the user group (mail for LRZ employees itself or mail for staff/students off the Munich universities LMU and TU which are the customers of the LRZ). The E-Mail Service can also be accessed by using a dedicated web mail server.

Examples of resource failures causing QoS degradations for the E-Mail Service might be:

- Failure of a hard disk storing the e-mail inboxes on one of the incoming mail servers
- Mail dispatching at the mail relays is very slow because of too many mails waiting in the mail queue (possibly caused by a lot of spam mails)
- The load balancer is not working properly (causing high delay and packet loss) because of wrong routing tables

Using the relationships on the resource level which are contained in the network topology database and in the systems configuration other resources which are affected by the failure can be identified. For the hard disk failure in the e-mail server, it is possible that processes running on this system will not work properly anymore. If the load balancer is not working properly the mail dispatching of the mail relays will be affected. Information about affected resources is transferred to the service management (*step 2*).

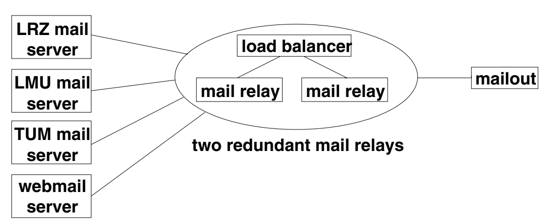

Figure 3. E-Mail Service at the LRZ

In the service management, the services which use the malfunctioning resources are identified traversing the dependencies between services and resources. The severity of the impact is also derived. It may be low if a service is provided using two redundant servers like the E-Mail Service and only one of these servers is currently not reachable. In addition, the dependencies between services are used to retrieve other affected services. The information about both types of dependencies is contained in the Service MIB. The list of all affected services including the expected QoS degradation is transferred to the SLA verificator (*step*

3). In case of the hard disk failure it is important to determine if there is any kind of data backup (e.g., a mirroring second drive or a RAID system using several disks).

While the QoS is derived by the service management in a provider-oriented way, the quality a user receives should also be taken into account in the impact analysis. This customer-oriented quality has to be measured in any case as it is used for the definition of service level agreements. Here, this kind of measurement can be regarded as a control procedure for the provider-internal derivation result. Therefore, the list of affected services is sent to the QoS measurement (*step 4*) and information about the severity of the service quality degradation is transferred back to the SLA verificator (*step 5*).
For the E-Mail Service an example SLA is defined as follows:

Availability: 99.9% during business hours, weekly basis

Delay: Sending of mail (to next mail domain) takes less than 25 minutes in 99% of the cases, maximum size 10 MB, maximum 500 mails per user a day

Penalty: 10.000 $ per month, immediate possibility to change the provider in case of violation

To determine the expected costs for not correctly providing the service, the current service usage by customers (and their users) is taken into account. If e.g., a service is not working properly, but it is only used by few customers whose SLAs do not contain severe penalties, then the impact can be classified as low. Prediction models can be used to get an expected service usage for future time intervals. To get such usage information, the affected services are sent to the service usage measurement and prediction (*step 6*). The result is received by the SLA verificator in *step 7*.

To keep the customers informed about the status of the services with respect to the SLAs, the information gathered so far is transferred to the CSM (*step 8*). From the collected information the SLA verificator can now determine an expected cost function over time for not repairing the resource failure(s). This information together with the resource failure(s) and corresponding repair possibilities are reported to the recovery management (*step 9*). The recovery management decides which recovery steps should be performed and tracks the recovery progress. The customers are kept informed by transferring information to the CSM (*step 10*).

In case of the E-Mail Service, there might be different options to efficiently react to faults:

Do nothing: In some cases a resource failure may cause no or only minor impact because of redundancy. In this case it is possible to do nothing in the current situation.

Maintain by regular staff: The employees which are responsible for this service try to fix the problem themselves.

Call additional staff from other groups: Other employees are requested to help to deal with the failure, if they have appropriate knowledge.

Call other company: A company could be called which is either specialized in this specific kind of problem or is the vendor of a malfunctioning component.

Outsource the service: Another company is hired to provide the E-Mail Service if the LRZ cannot do this anymore because of severe problems. This can be seen as last resort to avoid SLA penalties. Maybe it could be thought of trying to renegotiate some contracts or to terminate them.

5 Conclusion and Future Work

When a resource failure occurs at a provider today, the recovery action is often chosen by relying on the experience of employees. As the requirement analysis has shown many factors (dependencies, SLA definition and penalties, current service usage, costs for recovery alternatives, etc) need to be considered to react reasonably onto the resource failure. The framework and corresponding workflow presented in this paper formalize these coherences and define a detailed decision procedure which can be automated. By making sure not to neglect important influence factors, the accuracy of the recovery decisions can be improved using the workflow.

In the future, some issues of the requirement analysis need to be addressed in more detail. A search procedure and appropriate data structure has to be found to quickly identify the affected services. This is part of the "Service MIB", a generic representation of information needed for service provisioning which is currently developed by our research group. Methodologies for monitoring and prediction of service usage need to be examined as well.

Acknowledgments

The authors wish to thank the members of the MNM Team for helpful discussions and valuable comments on previous versions of the paper. The MNM Team, directed by Prof. Dr. Heinz-Gerd Hegering, is a group of researchers of the Munich Universities, the University of Armed Forces in Munich, and the Leibniz Supercomputing Center of the Bavarian Academy of Sciences. Its webserver is located at `http://wwwmnmteam.informatik.uni-muenchen.de`.

References

[1] S. Bagchi, G. Kar, and J. Hellerstein. Dependency analysis in distributed systems using fault injections: Application to problem determination in an e-commerce environment. In *Proceedings of the 12th International IFIP/IEEE Workshop on Distributed Systems: Operations and Management (DSOM 2001)*, Nancy, France, October 2001.

[2] C. Bartolini and M. Salle. Business driven prioritization of service incidents. In *Proceedings of the 15th International Workshop on Distributed Systems: Operations and Management (DSOM 2004)*, pages 64–75, Davis, California, USA, November 2004. IFIP.

[3] P. Bhoj, S. Chutani, and S. Singhal. Sla management in federated environments. In *Proceedings of the 6th IFIP/IEEE International Symposium on Integrated Network Management (IM 1999)*, pages 293–308, Boston, Massassuchets, USA, May 1999. IFIP/IEEE.

[4] D. Caswell and S. Ramanathan. Using service models for management of internet services. In *HP Technical Report HPL-1999-43, HP Laboratories*, Palo Alto, California, USA, March 1999.

[5] CIM Standards. http://www.dmtf.org/standards/cim.

[6] C. Ensel. New approach for automated generation of service dependency models. In *Network Management as a Strategy for Evolution and Development; Second Latin American Network Operation and Management Symposium (LANOMS 2001)*, Belo Horizonte, Brazil, August 2001. IEEE.

[7] S. Frolund, M. Jain, and J. Pruyne. Solomon: Monitoring end-user service levels. In *Proceedings of the 6th IFIP/IEEE International Symposium on Integrated Network Management (IM 1999)*, pages 261–274, Boston, Massassuchets, USA, May 1999. IFIP/IEEE.

[8] M. Garschhammer. Dienstguetebehandlung im Dienstlebenszyklus: von der formalen Spezifikation zur rechnergestuetzten Umsetzung - in German. PhD thesis, University of Munich, Department of Computer Science, August 2004.

[9] M. Garschhammer, R. Hauck, H.-G. Hegering, B. Kempter, M. Langer, M. Nerb, I. Radisic, H. Roelle, and H. Schmidt. Towards generic service management concepts - a service model based approach. In *Proceedings of the Seventh IFIP/IEEE International Symposium on Integrated Network Management*, pages 719–732. IFIP/IEEE, May 2001.

[10] M. Garschhammer and H. Roelle. Requirements on quality specification posed by service orientation. In *Proceedings of the 15th International Workshop on Distributed Systems: Operations and Management (DSOM 2004)*, November 2004.

[11] enhanced Telecom Operations Map (eTOM), The Business Process Framework for the Information and Communications Services Industry. Technical Report GB 921 Approved Version 3.0, TeleManagement Forum, June 2002.

[12] B. Gruschke. Integrated event management: Event correlation using dependency graphs. In *Proceedings of the 9th IFIP/IEEE International Workshop on Distributed Systems: Operations & Management (DSOM 98)*, pages 130–141. IEEE/IFIP, October 1998.

[13] M. Gupta, A. Neogi, M. Agarwal, and G. Kar. Discovering dynamic dependencies in enterprise environments for problem determination. In *Proceedings of the 14th IFIP/IEEE Workshop on Distributed Sytems: Operations and Management*. IFIP/IEEE, October 2003.

[14] A. Hanemann and M. Sailer. Towards a framework for service-oriented event correlation. In *Proceedings of the Service Assurance with Partial and Intermittent Resources Conference (SAPIR 2005)*, Lisbon, Portugal, July 2005. IARIA/IEEE. to appear.

[15] A. Hanemann, M. Sailer, and D. Schmitz. Assured service quality by improved fault management - service-oriented event correlation. In *Proceedings of the 2nd International Conference on Service-Oriented Computing (ICSOC04)*, New York City, New York, USA, November 2004. ACM.

[16] A. Hanemann, M. Sailer, and D. Schmitz. Towards a framework for failure impact analysis and recovery with respect to service level agreements. In *Proceedings of the 9th IFIP/IEEE International Conference on Integrated Network Management (IM 2005)*, Nice, France, May 2005. IFIP/IEEE. to appear.

[17] H.-G. Hegering, S. Abeck, and B. Neumair. *Integrated Management of Networked Systems - Concepts, Architectures and their Operational Application*. Morgan Kaufmann Publishers, 1999.

[18] Infovista Corporation. http://www.infovista.com.

[19] A. Keller and H. Ludwig. The wsla framework: Specifying and monitoring service level agreements for web services. In *IBM Research Report, RC22456 (W0205-171)*, Yorktown Heights, New York, USA, May 2002.

[20] M. Langer, S. Loidl, and M. Nerb. Customer service management: A more transparent view to your subscribed services. In *Proceedings of the 9th IFIP/IEEE International Workshop on Distributed Systems: Operations & Management (DSOM 98)*, pages 195–206, Newark, DE, USA, October 1998. IFIP/IEEE.

[21] L. Lewis. *Service Level Management for Enterprise Networks*. Artech House, Inc., 1999.

[22] Office of Government Commerce (OGC), editor. *The Business Perspective*. IT Infrastructure Library (ITIL). The Stationary Office, Norwich, UK, 2003.

[23] G. D. Rodosek and T. Kaiser. Determining the availability of distributed applications. In *Proceedings of the Fifth IFIP/IEEE International Symposium on Integrated Network Management (IM 1997)*, pages 207–218, San Diego, California, USA, May 1997. IFIP/IEEE.

[24] M. Sailer, V. Danciu, and A. Hanemann. Towards a service management information base. In *submitted to International Conference on Networks and Services (ICNS2005)*, San Jose, California, USA, September 2005. IARIA/IEEE.

[25] M. Salle and C. Bartolini. Management by contract. In *Proceedings of the 9th IEEE/IFIP International Network Management and Operations Symposium (NOMS 2004)*, pages 787–800, Seoul, Korea, April 2004. IFIP/IEEE.