Scope and Methods for Impact and Recovery Analysis

Andreas Hanemann[1], David Schmitz[2], and Claudio Bartolini[3]

[1] German Research Network (DFN)
Stresemannstr. 78, 10963 Berlin, Germany
hanemann@dfn.de
[2] Munich Network Management Team
Leibniz Supercomputing Center
Boltzmannstr. 1, 85748 Garching, Germany
schmitz@lrz.de
[3] Hewlett Packard Labs
1501 Page Mill Road, Palo Alto, CA 94304-1126, USA
claudio.bartolini@hp.com

Abstract. Impact and recovery analysis is an important issue for service providers in the context of service fault management. However, the scope of impact and recovery analysis is often unclear for which criteria are proposed in this paper. Furthermore, the scope is compared to existing approaches and it is shown how additional benefits related to the methods can be achieved.

1 Problem Statement

IT resource faults or performance degradations can have critical consequence to IT service provisioning so that the handling of faults is an important subject to service providers. Besides of fault isolation, it is also necessary to make decisions about fault treatment. Usually, little information is available to the decision making process and, as a consequence, such decisions are made on the basis of staff experience in many organizations which is likely to lead to inaccurate priorities. Therefore, decision support tools are required which need to consider failure impact and recovery possibilities.

The rest of the paper is organized as follows. In order to specify the demanded capabilities of such tools in a systematic manner there needs to be a clear specification of the scope of impact analysis which is discussed in Section 2. It helps to identify which kinds of issues can be addressed by a common set of methods. In Section 3 a combination of previous methods is presented in detail which is set in context of the potential scope. A comparison with the state-of-the-art is performed in Section 4. Section 5 concludes the paper.

2 Scope of Impact Analysis

For the scope of impact analysis a set of aspects needs to be distinguished which is explained in the following. The basic scenario is that an organization offers IT

services based on its internal resources and needs to have a method for calculating the impact of actual or assumed faults onto the offered services.

Amount of failures The amount of failures to be treated in parallel in an organization can range from one to multiple or even many in disaster situations.

For a single resource an impact analysis method should be able to detect what other resources and services (including the relation to customer SLAs) are affected. For some resources it can be distinguished between different kinds of failures (e.g. complete failure of an IP link or high packet loss on it). This analysis should be quickly executable in a current situation to determine the effort to be invested for resolving the issue.

In case of multiple failures the impact analysis should assist in selecting a priority to address the issues in addition to the considerations for the single failure situation.

The impact method should not only be useful for the fault management domain, but should also be applicable to disaster situations when a lot of resources are damaged or destroyed.

Time horizon Impact analysis can be beneficial with respect to different time dimensions.

On the short term, the impact of currently unavailable resources has to be determined in order to select appropriate recovery measures.

For medium term network planning purposes, the estimation of a hypothetical failure impact is helpful. It can quantify the criticality of resources and can be the basis for improving the network design.

On the long term, the pricing of resources can be based on the impact calculation so that the use of critical resources is charged with a higher price.

Management domain In addition to the fault management domain, impact analysis is also an important means for change management. It enables change management to estimate an impact in case a change could fail and therefore enables to make more accurate decisions on changes.

Furthermore, both faults and security incidents have an impact on the organization. For an important part of security incidents such as denial-of-service attacks similar methods for impact calculation can be applied.

Recovery linking Finally, impact analysis is not made for its own purpose, but has to be seen in context with the drawing of conclusions (e.g. which changes should be made). The linking to the decision making is therefore another criterion.

3 Combined Approach

Based on this scope it is investigated which potential methods can be applied. The idea is to combine two previous approaches: The impact analysis framework proposed by the MNM Team [6] and HP labs' methods for impact analysis which are known under the term "Management by Business Objectives" (MBO) [1, 2].

Framework overview Due to the complexity of impact analysis the framework has been developed in a set of steps. It starts with a relatively simple basic modeling of a workflow and of information related to the workflow which is refined over time by dropping more and more of initially made assumptions.

In general, the work is split-up into three main tasks. *Impact Analysis* refers to the pure analysis of the current situation without considering the recovery. The recovery is divided into *Recovery Analysis* and *Recovery Tracking*. Recovery Analysis refers to the selection of recovery alternatives based on the impact. Recovery Tracking deals with checking of decisions which have been made over time and allows to incorporate new developments in the decision making.

Impact Analysis The Impact Analysis is mainly composed of two steps, i.e. service impact analysis and business impact analysis (see Fig. 1). The service impact analysis starts with a list of resources for which degradations in their performance exist. This means that complete failures as well as partial degradations are possible. These are mapped to the offered services so that consequently a list of affected services is received. The business impact analysis takes the impact onto services as input and calculates the consequences on the business by integrating contractual data such as SLA conditions.

In Fig. 1 three kinds of information artifacts have been defined accordingly: resource degradations, service degradations and business degradations. For all three of them the scope/value aspects are detailed (what is affected, how is it affected, to what degree is it affected) as well as time aspects (permanent, random or known time conditions of the degradation).

Recovery Analysis The aim of the Recovery Analysis is to have a decision aid on what should be done to reduce the business impact calculated in the previous step. The idea is to (semi-)automatically retrieve a set of recovery alternatives and to calculate the consequences of choosing some of them over time. Important aspects are the time that is expected for a recovery plan to succeed and the costs involved. Concerning the costs it has to be differentiated between the costs that relate to impact calculated earlier and the costs of recovery alternatives. A further aspect that needs to be considered is that some recovery alternatives may only be partial in nature. This means that they may solve the problem only temporarily or only for a part of the functionality. Some may also have side effects, for example reducing the security level which may not be acceptable for a longer period of time.

Recovery Tracking The Recovery Tracking has been added in addition to the two previous steps. It is different from the two other steps in the sense that it happens over time while the two other steps usually happen in a short time interval (how short depends on the degree of automation achieved, etc). The reason why Recovery Tracking is explicitly mentioned and not considered as part of the normal service and business operation is that information in context with recovery is sometimes neglected. For instance, dependencies may have changed after the recovery and it has to be checked whether a safe and secure state of the

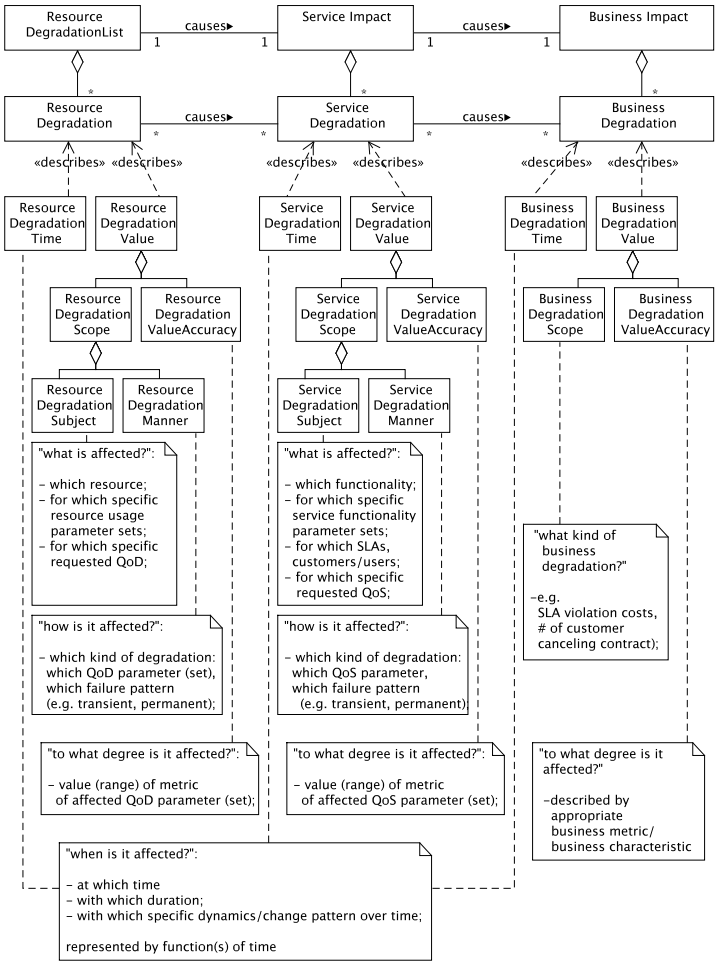

Fig. 1. Basic kinds of information artifacts involved in impact analysis

business has again been reached. This is required by future impact analysis runs in order to have consistently updated model information for performing impact analysis of future potential degradations.

Methods The framework description has been focused on what has to be done and in particular on information involved, but has not mentioned the methods yet. The Impact Analysis is addressed by the use of rules which are used to encode the existing dependencies among resources, services and business. In order to make use of previous impact an organization of knowledge as deductive memory (deductive database) is done. The language Frame Logic (F-Logic, [8]), which is often used for deductive databases, has turned out to be useful for modeling this knowledge. For Recovery Analyis F-Logic is not suitable since it can only

deal with existing facts and draw conclusions in contrast to the hypothetical application of alternatives needed here. For this situation the Transaction Logic [4] can be applied. For the Resolution Tracking again F-Logic is useful.

The MBO work has been focused on capturing the business relationships by quantitative methods. In [1] the SLA terms have been used to determine how much effort should be invested in a current fault situation, while incident management has been integrated into the considerations in [2]. In the latter work it is investigated which impact incidents may have onto the business and how they should be treated.

Joint impact analysis scope The approaches are complementary in the sense that the quantitative methodologies provided by MBO can be regarded as a refinement within the steps of the impact analysis framework.

The starting point of both approaches has been the short term analysis of one or few faults. Concerning the amount of faults there are scalability issues when many faults are occurring in parallel. For this case heuristics should be designed to put individual faults into larger groups and also to summarize the dependencies (e.g. in case of fire in a large data center it is not reasonable anymore to try to cope with individual up/down messages for single links). The methods to be applied on grouped faults and grouped dependencies can then be similar. The time horizon can be extended to mid-term planning where it is already possible to calculate hypothetical costs for individual failures. This can be input to optimization when redundancy alternatives are modeled appropriately. The long-term extension in the general case is more difficult since this is dependent on the business environment of the organization. With respect to the management domain the impact can also be estimated for changes and security incidents without much change in the information modeling. The linking to recovery management is already a feature of the provided framework.

4 Competitive Approaches

Impact analysis is contained in the management frameworks ITIL (IT Infrastructure Library, [9]) and eTOM (enhanced Telecom Operations Map, [10]). These frameworks provide a terminology for important terms (such as impact, urgency, priority) and describe in prose what needs to be done. However, recommendations for methods to implement impact analysis are only briefly referred to.

The literature contains a set of methods for impact analysis. A quantitative method comparable to MBO has been presented in [5]. A policy-based method has been presented in [3] based on the idea to prescribe and enforce logic-based policies.

Further related work is often useful for certain aspects of impact analysis. For instance, there are approaches for scheduling changes (e.g. the CHAMPS system, [7]) which is useful for managing the recovery.

Finally, commercial products are available such as HP's Service Navigator. These tools share common limitations, in particular with respect to the SLA terms they assume.

5 Current Status and Next Steps

The work is currently focused on joining the methods for impact analysis and to extend the scope towards the refinement of the recovery method. An implementation is under way at the Leibniz Supercomputing Center using the F-Logic language. In the future the modification of methods for extending the scope of impact and recovery analysis are examined in details.

References

1. C. Bartolini and M. Salle. Business Driven Prioritization of Service Incidents. In *Proceedings of the 15th IFIP International Workshop on Distributed Systems: Operations and Management (DSOM 2004)*, pages 64–75, Davis, California, USA, November 2004.
2. C. Bartolini, M. Salle, and D. Trastour. IT Service Management Driven by Business Objectives - An Application to Incident Management. In *Proceedings of the 10th IEEE/IFIP International Network Management and Operations Symposium (NOMS 2006)*, Vancouver, British Columbia, Canada, April 2006.
3. K. Bhaskaran, H. Chang, and J. Jeng. Policy Driven Business Performance Management. In *Proceedings of the 15th IFIP International Workshop on Distributed Systems: Operations and Management (DSOM 2004)*, Davis, California, USA, November 2004.
4. A. Bonner and M. Kifer. Concurrency and communication in transaction logic. In *Joint International Conference and Symposium on Logic Programming*, pages 142–156, Bonn, Germany, 1996.
5. M. Buco, R. Chang, L. Luan, C. Ward, J. Wolf, and P. Yu. Utility Computing: SLA Management Based upon Business Objectives. *IBM Systems Journal*, 2004.
6. A. Hanemann, M. Sailer, and D. Schmitz. A Framework for Failure Impact Analysis and Recovery with Respect to Service Level Agreements. In *Proceedings of the IEEE International Conference on Services Computing (SCC 2005, volume II)*, pages 49–56, Orlando, Florida, USA, July 2005. IEEE press.
7. J. Hellerstein, A. Keller, V. Krishnan, J. Wolf, and K. Wu. The CHAMPS System: Change Management with Planning and Scheduling. In *Proceedings of the 9th IFIP/IEEE International Network Management and Operations Symposium (NOMS 2004)*, pages 395–408, Seoul, Korea, April 2004.
8. M. Kifer, G. Lausen, and J. Wu. Logical foundations of object-oriented and frame-based languages. *ACM Journal*, 42(4):741–843, May 1995.
9. OGC (Office of Government Commerce), editor. *Service Operation Book*. IT Infrastructure Library (ITIL). The Stationary Office, Norwich, UK, 2007.
10. enhanced Telecom Operations Map (eTOM), The Business Process Framework for the Information and Communications Services Industry, TeleManagement Forum. GB 921 Release 5.0, April 2005.