Towards a Service Management Information Base

Martin Sailer

Munich Network Management Team
University of Munich (LMU)
Oettingenstr. 67, D-80538 Munich, Germany
martin.sailer@mnm-team.org

Abstract. The trend towards service-oriented computing has implications on the area of IT management. While traditionally management of networks and systems has dominated research as well as industrial applications, the focus is now turning to service management. One of the key concepts of integrated management is the management information base (MIB), the composite of management information. While a number of standardized MIBs have been devised in the area of network and systems management, a comprehensive MIB for service management is still missing. The PhD thesis presented in this paper aims at filling this gap by specifying semantic-rich and business-oriented service management information.

1 Introduction

With service-oriented computing becoming widely accepted by the industry, the number of commercial deployments is steadily increasing. The success of these deployments largely depends on the providers' ability to efficiently manage their IT infrastructure. To meet this requirement service providers deploy a wide array of hardware and software to deliver and manage services to their customers. Each management application deals with a certain facet of service management and therefore uses valuable data as key to its operation. The resulting situation is that information about services needs to be widely and electronically shared across departments throughout the provider's enterprise - and with partners and suppliers.

Such information is part of a Management Information Base (MIB), one of the key concepts of integrated management. The MIB contains the definition of the entities that are managed (managed objects), including their properties, their permissible behavior, and the operations that may be performed upon them [1]. Properties of managed objects are explicitly described by attributes, that usually consist of types and values. A MIB enables integration of different management applications in that it allows for exchange of semantically bound management information. For this reason, a large number of standardized managed object (MO) definitions have been devised in the area of network and systems management. However, a comprehensive MIB for service management is still missing. Hence, the realization of interoperability is impeded by the lack of a common information base for service management.

While a number of modeling frameworks for service management information have been devised, they currently show deficiencies regarding description of common service properties. The PhD thesis presented in this paper aims at addressing this gap with the concept of a Service Management Information Base (SMIB). The SMIB approach aims at specifying semantic-rich, business-oriented service management information. This includes common service management attributes as well as association information between services and resources.

In Section 2 a typical service provisioning scenario is presented in order to derive requirements for a Service MIB. Related work in this area is discussed in Section 3 and juxtaposed to the identified requirements. Section 4 introduces the SMIB approach and presents the further steps to be taken in the PhD thesis.

2 Scenario

The Web Hosting Service offered by the Leibniz Supercomputing Center (LRZ) serves as an example IT service (see Figure 1). The LRZ operates the scientific network in Munich and hosts web sites for other research institutions. Figure 1 shows a simplified view of the service incorporating three main components: an array of web servers, a global filesystem (AFS), and a database. In addition, network components need to be considered for management (Loadb routers, Nx-switches). The management systems respectively MIBs employed by the LRZ provide means to manage networks and applications, but do not take into account services or dependencies between services and resources.Yet, that knowledge has to be deduced by systems administrators and maintained separately. It therefore often lacks formalization and is usually not reusable. Thus, a comprehensive service management repository is required to enable consistent management. The challenges stemming from that requirement can be further categorized into specification of service attributes and formalization of service resource dependencies.

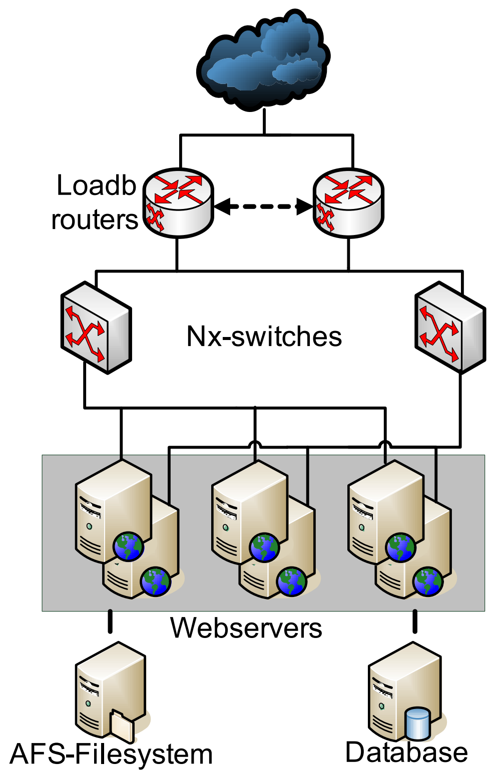

Fig. 1. Web Hosting Service at the LRZ

2.1 Specification of meaningful service attributes

Traditionally, management information has been defined from a vendors' perspective and was therefore often restricted to information that could be easily obtained by instrumenting components [2]. With the shift to service management, that perspective has to be adapted in a way that it reflects the providers' information needs. Furthermore, services are managed along their lifecycle, in different functional management areas, thus motivating the inclusion of these aspects. Additionally, service provisioning – and consequently its management – may cross administrative boundaries. Hence, specification of service management information needs to be tailored towards the following generic requirements:

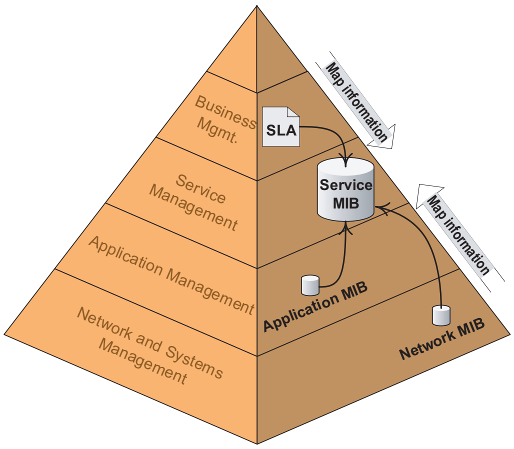

Fig. 2. Management Pyramid

Alignment to providers' information needs The MIBs currently employed contain information pertaining to applications and network elements. Examples for such data are the number of dropped IP packets, the CPU load of a webserver, or http-requests per second. While this may suffice from a network and systems management point of view, service management requires abstracted and enriched data that reflects the providers' information needs. That is, information needed to deliver the service according to the objectives set up in service level agreements (SLAs).This includes information relating to the quality level of the service, to the customers having subscribed to the service, or the impact of service failures on the business.

Administrative domains Application scenarios like dynamic service provisioning or Grid computing require services to be composed in a highly dynamic and flexible manner, thereby crossing organizational boundaries. Accordingly, the ServiceMIB must provide capabilites to encompass different management domains and therefore allow to present different views on service management information according to the needs of the involved organizations.

The functional areas of management, usually classified by the OSI FCAPS, dissect the whole complex of management. Yet, management information is shared between different functional areas, e.g., service usage information is a prerequisite for accounting as well as fault management. Consequently, all of these areas should be included in a ServiceMIB that purports to deal with service management.

The service lifecycle usually consists of a planning, negotiation, provisioning, operation/change, and withdrawal phase. Since management information is propagated and enriched among these phases, the ServiceMIB must accommodate all phases.

2.2 Formalization of dependencies between services and resources

Service management relies on knowledge regarding dependencies between services and resources. In particular, dependencies are the key to service problem resolution in that they provide a foundation for determining causality [3]. Formalization of dependencies has to accommodate different levels (e.g., functional, organizational) as well as provide means to express complex dependencies (e.g., dependencies causing certain service quality degradations). Therefore, information contained in system, network and application MIBs needs to be combined in such a way that it reflects the service. Consequently, a formalization of dependencies must allow for translating/mapping information from other MIBs to leverage an integrated management approach (see Fig. 2).

3 Related Work

In this section prominent representative work in the area of management information modeling is presented and its significance in the focus of this paper is discussed.

Common Information Model (CIM) The Common Information Model [4] introduced by the Distributed Management Task Force is an object-oriented information model that aims at providing a common way to represent information about networks and systems as well as services. The CIM Core Model contains the specification of the class *service*, which acts as a superclass from which specific service classes (e.g. print service) can be derived. Additionally, dependencies between services and other CIM elements (e.g., *ServiceAffectsElement*) are expressed by association classes. However, only few service attributes have been defined so far for the basic service class as well as associations pertaining to that class. CIM constitutes an comprehensive modeling framework that could be extended to meet the requirements delineated in Section 2. Since it accomodates network, application, and service management information, it offers a foundation for modeling dependencies between those areas.

Web Services Distributed Management Two complementary specifications have been presented by the Oasis Web Services Distributed Management technical committee: Management Using Web Services (MUWS) deals with providing a common management interface for distributed IT resources, whereas Management of Web Services (MOWS) introduces the application of the former onto Web Services [5,6]. Service Management information is exposed through manageability capabilities and categorized into common and resource-specific capabilities. At present, only basic capabilites like the *CurrentOperationalState,* the *LastResponseTime* or *NumberOfRequests* have been defined. Moreover, the WSDM approach does not explicitly address dependencies between services and resources.

NGOSS SID As part of the New Generation Operations Systems and Software (NGOSS) program the TeleManagement Forum released a Shared Information/Data (SID) model [7]. It employs an object-oriented modeling approach and defines a framework for defining service management information from a business-oriented point of view. SID offers a promising modeling framework for service management information that builds upon sound design principles. However, it also shows deficits regarding the description of common service attributes.

Internet Information Model The IETF model revolves traditionally around the Simple Network Management Protocol (SNMP) [8]. This is a read/write, state based protocol for altering variables in a management information base (MIB) and is used by the vast majority of commercial products. Despite of newly introduced RFCs that deal with the specification of service management information for WWW services[9], the Internet Information model does not offer object-oriented modeling features, and is therefore not suitable for modeling of service management information.

4 Approach

The Service MIB approach aims at specifying semantic-rich, business-oriented service attributes as well as association information between service and resources. It therefore consists of two main steps, that are aligned with the challenges presented in Section 2:

Specification of business-oriented service attributes In recent time, service management has been influenced by the strong trend towards process-orientation. Process-oriented frameworks contribute to that trend in that they provide guidance for enterprises striving to deploy IT service management processes. A prominent example for such a framework is the IT Infrastructure Library (ITIL). ITIL has been introduced by the British Office of Government Commerce (OGC) and provides a collection of best practices for IT processes in the area of IT service management [10,11]. Service management is described by 11 modules which are grouped into Service Support Set and Service Delivery Set. These modules describe processes along with its goals and activities as well as necessary databases and interfaces.

To facilitate the definition of service attributes we leverage a use-case driven approach. The main idea is to analyze ITIL service management processes in order to derive common service attributes. Since such processes essentially represent a specification for service management tasks, they provide a comprehensive collection of use cases that can be further utilized to derive service management information needs. This approach offers considerable benefits with respect to the requirements denoted in Section 1. Since ITIL constitutes a business-driven approach to IT service management, it already takes lifecyle and interorganizational aspects into account, while at the same time offering a sufficient level of genericity.

Example: The Incident Management Process deals with restoring normal service operation. One of the core activities within this process is the classification of incidents, consisting of e.g., priorization of incidents. To support this activity, service management information relating to the business impact of the incident is required. That is, the SLA associated with the service, the impact on related services, or affected users.

Formalization of service-resource dependencies Formalization of association information regarding services and resources requires high expressive power that cannot be expressed by single attributes. Instead, a more complex formalism needs to be employed that allows to specify associations on a per-attribute basis in a functional manner. This includes means to map information defined for network and systems to ServiceMIB attributes. Therefore, this issue has been approached with a declarative language to specify service-resource dependencies, called Service Information Specification Language (SISL). This concept has been presented in [12], which introduced a monitoring architecture that facilitates the creation of service management information from raw monitoring data according to a SISL specification.

Further Work

Previous work in the context of the incident and problem management processes has analyzed use-cases in the area of service-oriented event correlation and impact analysis [3]. This work has shown applicability of the approach and will be extended towards other ITIL processes. While up to now, basic examples of SISL for expressing the operational state of services have been drafted [12], further work on that topic will target modeling of more complex scenarios.

Acknowledgment

The author wishes to thank the members of the Munich Network Management (MNM) Team for helpful discussions and valuable comments on previous versions of this paper. The MNM Team directed by Prof. Dr. Heinz-Gerd Hegering is a group of researchers of the University of Munich, the Munich University of Technology, and the Leibniz Supercomputing Center of the Bavarian Academy of Sciences. Its web–server is located at http://www.mnm-team.org.

References

1. Black, U.: Network Management Standards - The OSI, SNMP and CMOL Protocols. McGraw-Hill (1992)
2. Hegering, H.G., Abeck, S., Neumair, B.: Integrated Management of Networked Systems – Concepts, Architectures and their Operational Application. Morgan Kaufmann Publishers, ISBN 1-55860-571-1 (1999) 651 p.
3. Hanemann, A., Sailer, M., Schmitz, D.: 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, NY, USA, ACM SIGSOFT and SIGWEB, ACM Press (2004) 183–192
4. Distributed Management Task Force (DMTF): Common Information Model (CIM) Version 2.9. Specification (2005)
5. Organization for the Advancement of Structured Information Standards (OASIS): Web Services Distributed Management: Management Using Web Services (MUWS 1.0) Part 1 (2004)
6. Organization for the Advancement of Structured Information Standards (OASIS): Web Services Distributed Management: Management of Web Services (WSDM-MOWS 1.0) (2004)
7. TeleManagementForum: Shared Information/Data (SID) Model Concepts, Principles, and Domains. Technical Report GB 922 Member Evaluation Version 3.1 (2003)
8. Case, J., McCloghrie, K., Rose, M., Waldbusser, S.: RFC 1902: Structure of management information for version 2 of the simple network management protocol (snmpv2). RFC, IETF (1996)
9. Hazewinkel, H., Kalbfleisch, C., Schoenwaelder, J.: RFC 2594: Definitions of managed objects for www services. RFC, IETF (1999)
10. Office of Government Commerce (OGC), ed.: Service Support. IT Infrastructure Library (ITIL). The Stationary Office, Norwich, UK (2000)
11. Office of Government Commerce (OGC), ed.: Service Delivery. IT Infrastructure Library (ITIL). The Stationary Office, Norwich, UK (2001)
12. Danciu, V., Sailer, M.: A monitoring architecture supporting service management data composition. In: Proceedings of the 12th Annual Workshop of HP OpenView University Association. Number 972–9171–48–3, Porto, Portugal, HP (2005) 393–396