Service-Oriented Event Correlation - Workflow and Information Modeling Approached

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

Abstract

The paradigm shift from device-oriented to service-oriented management has also implications to the area of event correlation. Today's event correlation mainly addresses the correlation of events as reported from management tools. However, a correlation of user trouble reports concerning services should also be performed. This is necessary to improve the resolution time and to reduce the effort for keeping the service agreements. We refer to such a type of correlation as service-oriented event correlation.

For introducing service-oriented event correlation for an IT service provider, an appropriate modeling of the workflow and of the information is necessary. Therefore, we examine the process management frameworks ITIL and eTOM for their contribution to the workflow modeling in this area. The MNM Service Model, which is a generic model for IT service management proposed by the MNM Team, is used to derive an appropriate information modeling. The different kinds of dependencies that we find in our general scenario are used to develop a workflow for the service-oriented event correlation.

1. Introduction

In huge networks a single fault can cause a burst of failure events. To handle the flood of events and to identify the root cause of a fault, event correlation approaches like rule-based reasoning, case-based reasoning or the codebook approach have been developed. The main idea of correlation is to condense and structure events to retrieve meaningful information. Until now, these approaches address primarily the correlation of events as reported from management tools or devices. We call these approaches resource-oriented (for an overview of the state-of-the-art see [1]).

As in today's IT environments the offering of services with an agreed service quality becomes more and more important, this change also affects the event correlation. To avoid service level agreement (SLA) violations, it is especially important for service providers to identify the root cause of a fault in a very short time, when trouble reports are received from customers or the provider's own service surveillance. We call the kind of event correlation for such a scenario *service-oriented* as it uses knowledge about services, service provisioning and SLAs. As we showed in [1] the following reasons for service-oriented event correlation can be identified:

Resolution time minimization: The time interval between the first symptom (recognized either by provider, network management tools, or customers) that a service does not perform properly and the verified fault repair needs to be minimized. This is especially needed with respect to SLAs.

Effort reduction: If several trouble reports are symptoms of the same fault, the fault processing should be performed only once and not several times. If e.g. the fault has been repaired, all affected customers should be informed about that automatically.

Impact analysis: In case of a fault in a resource, its influence on associated services and affected customers can be determined. This analysis can be performed for short term (when there is currently a resource failure) or long term (e.g. network optimization) considerations.

To receive the benefits of the service-oriented event correlation, it is necessary to have an appropriate information modeling, e.g. with respect to the dependencies from services to subservices and resources. The workflow also needs to be modeled to show which steps are necessary during the event correlation process.

The rest of the paper is organized as follows. In Section 2 we present the management models ITIL and eTOM and examine their contribution to the area of fault management and especially to event correlation and show that the MNM Service Model is useful as basis for the information modeling of service-oriented event correlation. Our workflow design for this kind of correlation as well as the derived information modeling for the service events are presented in Section 3. The last section concludes the paper and presents future work.

2 Usability of Existing Models for Service-Oriented Event Correlation

In the following we examine the established IT process management frameworks ITIL and eTOM. The aim is find out where event correlation can be found in the process structure and how detailed the frameworks currently are. This is helpful to model the workflow for the service-oriented event correlation.

2.1 ITIL

The British Office of Government Commerce (OGC) and the IT Service Management Forum (itSMF) [2] provide a collection of best practices for IT processes in the area of IT service management. The collection is called "IT Infrastructure Library (ITIL) [3]". The service management is described by 11 modules which are grouped into Service Support Set (provider internal processes) and Service Delivery Set (processes at the customer-provider interface). Each module describes processes, functions, roles and responsibilities as well as necessary databases and interfaces. In general, ITIL describes contents, processes, and aims at a high abstraction level and contains no information about management architectures and tools.

The fault management is divided into Incident Management process and Problem Management process.

Incident Management: The Incident Management contains a service desk as interface to the customer (e.g. receives reports about service problems). In case of severe errors structured queries are transferred to the Problem Management.

Problem Management: The Problem Management's tasks are to solve problems, take care of keeping priorities, minimize the reoccurrence of problems, and to provide management information. After receiving requests from the Incident Management the problem has to be identified and information about necessary countermeasures is transferred to the Change Management.

The ITIL processes describe only what has to be done, but contain no information how this can be actually performed. As a consequence, event correlation is not part of the modeling. The ITIL incidents could be regarded as input for the service-oriented event correlation, while the output could be used as a query to the ITIL Problem Management.

2.2 TOM/eTOM

The TeleManagement Forum (TMF) [4] is an international non-profit organization from service providers and suppliers in the area of telecommunications services. Similar to ITIL a process-oriented framework has been developed at first, but the framework was designed for a narrower focus, i.e. the market of information and communications service providers. A horizontal grouping into processes for customer care, service development & operations, network & systems management, and partner/supplier is performed. The vertical grouping (fulfillment, assurance, billing) reflects the service life cycle.

In the area of fault management three processes have been defined along the horizontal process grouping.

Problem Handling: The purpose of this process is to receive trouble reports from customers and to solve them by using the Service Problem Management. The aim is also to keep the customer informed about the current status of the trouble report processing as well as about the general network status (e.g. planned maintenance). It is also a task of this process to inform the QoS/SLA management about the impact of current errors on the SLAs.

Service Problem Management: In this process reports about customer-affecting service failures are received and transformed. Their root causes are identified and a problem solution or a temporary workaround is established. The task of the "Diagnose Problem" subprocess is to find the root cause of the problem by performing appropriate tests. Nothing is said how this can be done (e.g. no event correlation is mentioned).

Resource Trouble Management: A subprocess of the Resource Trouble Management is responsible for resource failure event analysis, alarm correlation & filtering, and failure event detection & reporting. Another subprocess is used to execute different tests to find a resource failure. There is also another subprocess which keeps track about the status of the trouble report processing. This subprocess is similar to the functionality of a trouble ticket system.

The process description in eTOM is not very detailed. It is useful to have a check list which aspects for these processes have to be taken into account, but there is no detailed modeling of the relationships and no methodology for applying the framework. Event correlation is only mentioned in the resource management, but it is not used in the service level.

2.3 MNM Service Model

The MNM Service Model [5], which was developed by the Munich Network Management Team, is a generic model for service modeling. It distinguishes between *customer side* and *provider side*. The customer side contains the basic roles *customer* and *user*, while the provider side contains the role *provider*. The provider makes the service available to the customer side. The service as a whole is divided into usage which is accessed by the role user and management which is used by the role customer.

The model consists of two main views. The *Service View* (see Fig. 1) shows a common perspective of the service for customer and provider. Everything that is only important for the realization of the service is not contained in this view. For these details another perspective, the *Realization View*, is defined (see Fig. 2).

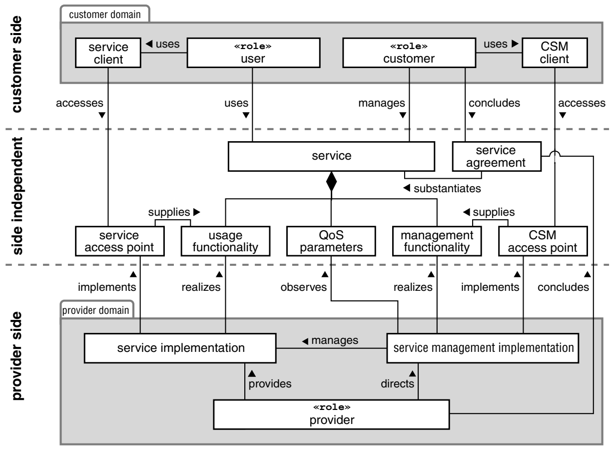

Figure 1. Service View

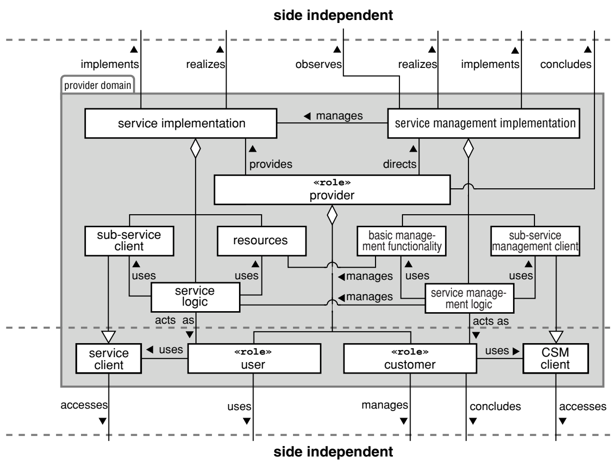

Figure 2. Realization View

The Service View contains the *service* for which the functionality is defined for management as well as for usage. There are two access points (service access point and CSM access point) where user and customer can access the usage and management functionality, respectively. Associated to each service is a list of QoS parameters which have to be met by the service at the service access point. The QoS surveillance is performed by the management.

In the Realization View the service implementation and the service management implementation are described in detail. For both there are provider-internal resources and subservices. For the service implementation a service logic uses internal resources (devices, knowledge, staff) and external subservices to provide the service. Analogous, the service management implementation includes a service management logic using basic management functionalities [6] and external management subservices.

The MNM Service Model can be used for a similar modeling of the used subservices, i.e. the model can be applied recursively.

As the service-oriented event correlation has to use dependencies of a service from subservices and resources the model is used in Subsection 3.4 to derive the needed information for service events.

3 Workflow and Information Modeling for Service-Oriented Event Correlation

Fig. 3 shows a general service scenario which we will use as basis for the workflow modeling for the service-oriented event correlation. The provider offers different services which depend on other services called subservices (service dependency). Another kind of dependency exists between services/subservices and resources. These dependencies are called resource dependencies. These two kinds of dependencies are in most cases not used for the event correlation performed today. This resource-oriented event correlation deals only with relationships on the resource level (e.g. network topology).

As both ITIL and eTOM contain no description how event correlation and especially service-oriented event correlation should actually be performed, we propose the following design for such a workflow (see Fig. 4). The workflow is divided into the three phases fault detection, fault diagnosis, and fault recovery. In general, we have two kinds of events: Resource events, which contain information about failures in resources, and service events, which contain information about service problems.

In the fault detection phase these events can be generated from different sources. The resource events are issued during the use of a resource, e.g. via SNMP traps. The service events are originated from customer trouble reports, which are reported via the Customer Service Management (see below) access point. In addition to these two "passive" ways

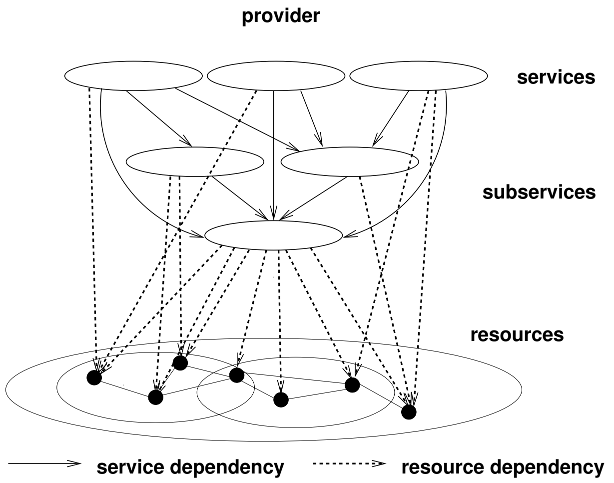

Figure 3. Different kinds of dependencies for the service-oriented event correlation

to get the events, a provider can also perform active tests. These tests can either deal with the resources (resource active probing) or can assume the role of a virtual customer and test a service or one of its subservices by performing interactions at the service access points (service active probing).

An important part of the fault diagnosis phase is the event correlation. The correlation contains the resource event correlator which can be regarded as the event correlator in today's commercial systems. Therefore, it deals only with resource events. The service event correlator does a correlation of the service events, while the aggregate event correlator performs a correlation of both resource and service events. If the correlation result in one of the correlation steps shall be improved, it is possible to go back to the fault detection phase and start the active probing to get additional events. These events can be helpful to confirm a correlation result or to reduce the list of possible root causes.

After the event correlation an ordered list of possible root causes is checked by the resource management. When the root cause is found, the failure repair begins. This last step is performed in the fault recovery phase.

The next subsections present different elements of the event correlation process.

3.1 Customer Service Management and Intelligent Assistant

The MNM Service Model contains a Customer Service Management (CSM) access point as a single interface between customer and provider. Its functionality is to provide information to the customer about his subscribed services, e.g. reports about the fulfillment of agreed SLAs. It can also

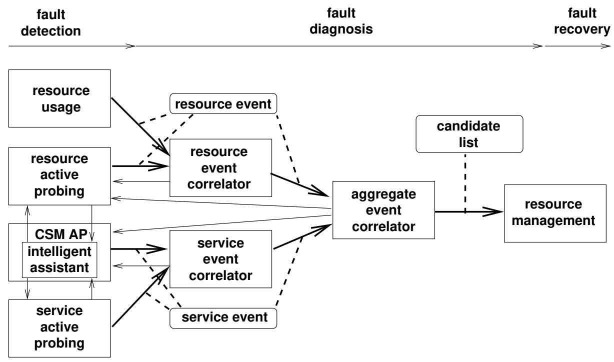

Figure 4. Event correlation workflow

be used to subscribe services or to allow the customer to manage his services in a restricted way. Reports about problems with a service can be sent to the customer via CSM.

To reduce the effort for the provider's first level support, an Intelligent Assistant can be added to the CSM. The Intelligent Assistant structures the customer's information about a service trouble. The information which is needed for a preclassification of the problem is gathered from a list of questions to the customer. The list is not static as the current question depends on the answers to prior questions or from the result of specific tests. A decision tree is used to structure the questions and tests. The tests allow the customer to gain a controlled access to the provider's management. At the LRZ a customer of the E-Mail Service can e.g. use the Intelligent Assistant to start a "ping" request to the mail server. But also more complex requests could be possible, e.g. requests of a combination of SNMP variables.

3.2 Active Probing

Active probing (see Fig. 5) is useful for the provider to check his offered services. The aim is to identify and react to problems before a customer notices them. The probing can be done from a customer point of view or by testing the resources which are part of the services. It can also be useful to perform tests of subservices (own subservices or subservices offered by suppliers).

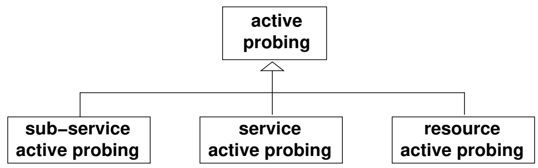

Figure 5. Active Probing

Different schedules are possible to perform the active probing. The provider could select to test important services and resources in regular time intervals. Other tests could be initiated by a user who traverses the decision tree of the Intelligent Assistant including active tests. Another possibility for the use of active probing is a request from the event correlator, if the current correlation result needs to be improved. The results of active probing are reported via service or resource events to the event correlator (or if the test was demanded by the Intelligent Assistant the result is reported to it, too). While the events that are received from management tools and customers denote negative events (something does not work), the events from active probing should also contain positive events for a better discrimination.

3.3 Event Correlator

Because we have to deal with two types of events (resource events and service events) in the service-oriented scenario, the event correlation should be performed in different steps. The reason for this are the different characteristics of the dependencies (see Fig. 3).

On the resource level there are only relationships between resources, e.g. caused by the network topology. An example for this could be a switch linking separate LANs. If the switch is down, events are reported that other network components which are behind the switch are also not reachable. When correlating these events it can be figured out that the switch is the likely error cause. At this stage, the integration of service events does not seem to be helpful. The result of this step is a list of resources which could be the problem's root cause. The resource event correlator is used to perform this step.

In the service-oriented scenario there are also service and resource dependencies. As next step in the event correlation process the service events should be correlated with each other using the service dependencies. The result of this step which is performed by the service event correlator is a list of services/subservices which could contain a failure in a resource. If e.g. there are service events from customers that the video conference service and e-mail service do not work and both services depend on a common subservice (in this case e.g. the DNS), it seems more likely that the resource failure can be found inside the subservice.

In the last step the aggregate event correlator matches the lists from resource event correlator and service event correlator to find the problems possible root cause. This is done by using the resource dependencies.

Fig. 6 shows the different event correlators.

3.4 Resource Events and Service Events

Today's event correlation deals mainly with events which are originated from resources. Beside a resource identifier these events contain information about the resource status,

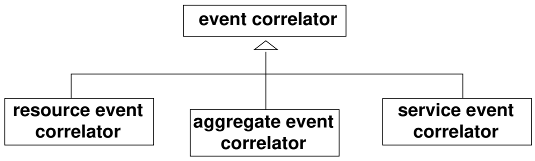

Figure 6. Event Correlators

e.g. SNMP variables. To perform a service-oriented event correlation it is necessary to define events which are related to services. These events can be generated from the provider's own service surveillance or from customer reports at the CSM interface. They contain information about the problems with the agreed QoS. In our information modeling we define an event superclass which contains common attributes e.g. time stamp. Resource event and service event inherit from this superclass (see Fig. 7).

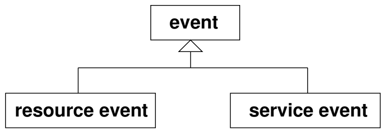

Figure 7. Events

Derived from the MNM Service Model we can define the information which is necessary for a service event.

Event description: This field has to contain a description of the problem. Depending on the interactions at the service access point (Service View) a classification of the problem into different categories should be defined. It should be possible to add an informal description of the problem.

Issuer's identification: This field can either contain an identification of the customer who reported the problem, an identification of a service provider's employee (in case the failure has been detected by the provider's own service active probing) or a link to a parent service event (see below). The identification is needed, if there are ambiguities in the service event or the issuer should be informed (e.g. that the service is available again). The possible issuers refer to the basic roles (customer, provider) in the Service Model.

Dates: This field contains key dates in the processing of the service event such as initial date, problem identification date, resolution date. These dates are important to keep track how quick the problems have been solved.

Status: This field represents the service events actual status (e.g. active, suspended, solved).

Priority: The priority shows which importance the service event has from the provider's perspective. The importance is derived from the service agreement, especially the agreed QoS parameters (Service View).

Assignee: To keep track of the processing the name and address of the provider's employee who is solving or solved the problem is also noted. This is a specialization of the provider role in the Service Model.

Service: As a service event shall represent the problems of a single service, a unique identification of the affected service is contained here.

QoS parameters: For each service QoS parameters (Service View) are defined between the provider and the customer. This field represents a list of these QoS parameters and agreed service levels. The list can help the provider to set the priority of a problem with respect to the service levels agreed.

Resource list: This list contains the resources (Realization View) which are needed to provide the service. This list is used by the provider to check if one of these resources causes the problem.

Subservice service event identification: In the service hierarchy (Realization View) the service for which this service event has been issued may depend on subservices. If there is a suspicion that one of these subservices causes the problem, child service events are issued from this service event for the subservices. In such a case this field contains links to the corresponding events.

Other event identifications: In the event correlation process the service event can be correlated with other service events or with resource events. This field then contains links to other events. This is useful to, e.g., send a common message to all affected customers when their subscribed services are available again.

The fields date, status, and other events are not derived directly from the Service Model, but are necessary for the service event correlation process.

4 Conclusion and Future Work

In our paper we showed the need for a service-oriented event correlation. For an IT service provider this new kind of event correlation makes it possible to automatically map problems with the current service quality onto resource failures. This helps to find the failure's root cause earlier and to reduce costs for SLA violations. In addition, customer reports can be linked together and therefore the processing effort can be reduced.

To receive these benefits we presented our approach for performing the service-oriented event correlation as well as a modeling of the necessary correlation information. In the future we are going to apply our workflow and information modeling for different services offered by the Leibniz Supercomputing Center.

Several issues have not been treated in detail so far, e.g. the consequences for the service-oriented event correlation if a subservice is offered by another provider. If a service of the provider does not work, it has to be determined whether this is caused by the provider himself or by the subservice. Another issue is the use of active probing in the event correlation process which can improve the result, but can also lead to delay in the correlation.

Acknowledgments

The authors wish to thank the members of the Munich Network Management (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 and the Leibniz Supercomputing Center of the Bavarian Academy of Sciences. Its Web server is located at `wwwmnmteam.informatik.uni-muenchen.de`.

References

[1] Hanemann, A. and Schmitz, D.: Why is Service-Orientation Necessary for Event Correlation? Proceedings of the DAIS/FMOODS PhD Student Workshop, Paris, France, November 2003.

[2] Websites: www.itil.co.uk and www.itsmf.com

[3] ITIL Service Delivery, ISBN 0113300174
ITIL Service Support, ISBN 0113300158
ITIL Planning to Implement Service Management, ISBN 011330014X

[4] Website of the TeleManagement Forum: www.tmforum.org

[5] Garschhammer, M., Hauck, R., Hegering, H.-G., Kempter, B., Langer, M., Nerb, M., Radisic, I., Rölle, H., and Schmidt, H.: Towards generic Service Management Concepts - A Service Model Based Approach. In: Pavlou, G., Anerousis, N., and Liotta, A. (eds.): Integrated Network Management, VII, pages 719-732, IEEE/IFIP, May 2001.

[6] Hegering, H.-G., Abeck, S., and Neumair, B.: Integrated Management of Networked Systems - Concepts, Architectures and their Operational Application. Morgan Kaufmann Publishers, ISBN 1-55860-571-1, 1999.