

Fault Management based on peer-to-peer paradigms

A case study report from the CELTIC project Madeira

Markus Leitner, Philipp Leitner, Martin Zach

PSE – Program and System Engineering
Siemens AG Austria
Vienna, Austria

{markus.leitner, philipp.leitner, martin.zach}@siemens.com

Sandra Collins

Ericsson R&D Ireland
Network Management Research Centre
Dublin, Ireland

sandra.collins@ericsson.com

Claire Fahy

TSSG – Telecommunications Software & Systems Group
Waterford Institute of Technology
Waterford, Ireland
cfahy@tssg.org

Abstract — We present an approach to fault management based on an architecture for distributed and collaborative network management as developed in the CELTIC project Madeira. It uses peer-to-peer communication facilities and a logical overlay network facilitating decentralized and iterative alarm processing and correlation. We argue that such an approach might help to overcome key challenges that are posed by NGN scenarios to traditional centralized network management systems. Its feasibility is demonstrated by means of a case study from the area of wireless mesh networks, where an application prototype has been developed.

Keywords – *fault management, decentralized and distributed management, policy based management, overlay networks*

I. INTRODUCTION

Most architectures of currently deployed network management systems (NMS) for telecommunication networks can be characterized as *centralized* and *hierarchical*: While marking a great achievement in making operation more efficient, such NMS solutions require powerful machines because of the complex logic and large amount of management information to be processed; furthermore they involve costly redundancy mechanisms in order to avoid single points of failure. Hierarchy additionally introduces different levels of abstraction – from element management of individual network elements (NEs) up to business management at the top of the Telecommunication Management Network (TMN) pyramid [1]. Management information flow in both directions – up and down this *static* hierarchy – is cascaded, with information mapping performed at each layer. From a functional point of view, the five *FCAPS* disciplines are rather separated. Interactions between these disciplines typically happen at a higher management layer, or even through a human operator.

This approach to manage communication networks is very well understood and works well for *classic* telecommunication networks; plenty of technically mature systems that incorporate

that approach have been on the market for years. The last couple of years, however, have revealed several emerging technologies like Voice over IP (VoIP) and IP-TV, ubiquitous and pervasive computing, new wireless technologies, and various implications of ad-hoc and peer-to-peer (P2P) networks. This results in interesting Next Generation Network (NGN) scenarios that have already been (or are about to be) realized by network operators. Even if these NGN scenarios cover a wide spectrum of use cases and technologies, they have a few characteristics in common: (1) *Large scale* – up to 10^6 possibly small nodes, (2) *Heterogeneity* - different HW and SW platforms, different vendors, different access and communication protocols, and (3) *Dynamics* - nodes might appear and disappear regularly, and topology changes will be the rule rather than the exception.

In a typical state-of-the-art fault management system the main fault processing logic resides in a powerful correlation engine at the top of the processing chain where a lot of alarms from each network node are received. Using sophisticated statistical analysis this engine might find out the most probable root cause for a given sequence of alarms. Further, in order to be meaningful to an operator, the information contained in this large amount of alarms has to be enriched by correlation against data bases containing topology and other configuration related information, stored on the central server. If the number of emitting nodes and alarms exceeds a certain critical value such a centralized correlation engine will become a *bottleneck*. Moreover, the *heterogeneity* of NEs increases the complexity of the system, and finally – maybe the most critical issue – if *static* information on network topology is used for fault analysis, the system will inevitably suffer from inconsistencies as soon as this topology exhibits dynamic aspects.

Motivated by these issues we will present a conceptually different fault management design that is based on a peer-to-peer (P2P) approach to network management suggested in the CELTIC research project Madeira [4]. The remaining part of this paper is organized as follows: Chapter II briefly

summarizes related work, chapter III presents the overall Madeira solution and in particular the design of the Fault Management application. Chapter IV contains a case study report explaining how certain scenarios in a wireless mesh network have been realized using our approach. Chapter V concludes the paper and outlines further ongoing work.

II. RELATED WORK

Several approaches to distributed network management have been made since Goldszmidt et al. presented their concept of *Management by Delegation* in 1991 [5]. A recent overview on distributed paradigms in this field has been published by Martin-Flatin et al. [8]. With respect to their taxonomy, our Madeira approach might be classified as following the *strongly distributed hierarchical paradigm* with the additional advantage of using a completely *dynamic* hierarchy.

The concept of Management by Delegation tries to release the central management entities of management tasks that can be done locally, delegating them on demand using local agents and mobile code. Similar approaches to network management use Intelligent [3] or Mobile Agents [6]. Agent-based approaches typically delegate management tasks dynamically on demand. In contrast to these, we follow the paradigm of performing a management task as *soon* as possible, i.e. whenever at least partial information is available. This leads to probably partial results generated by multiple nodes which are combined to final results at higher hierarchical levels.

Vertical delegation (delegation by domain) of network management tasks has also been implemented in centralized approaches using mid level managers [8]. The Madeira approach enhances them by avoiding the central bottleneck due to the *dynamic* and *self-organising hierarchy*, delegating management tasks even down to the lowest possible level (i.e. the *network elements* themselves) and using the *P2P paradigm*.

A P2P approach to network management has been presented by Binzenhöfer et al. [2]. Their approach of *Distributed Network Agents* (DNA) is on a high level conceptually similar to the Madeira system, but there are some noteworthy differences to Madeira: DNA uses the P2P features primarily to carry out *distributed tests* (e.g. measure the throughput on a subnetwork, ping a remote host, etc.) on the network it manages. Other functionality (such as Configuration and Fault Management, as investigated within Madeira) is not considered.

Finally, a critical aspect of distributing network management responsibilities is *decentralized alarm correlation*. Such models have a number of advantages over the more classical alarm correlation on a central entity: (1) No entity in the network has to have complete topology information of the network. (2) No central entity is responsible for doing all the correlation work of the whole network, which allows for a more scalable design. (3) Local alarm correlation is often much simpler than global correlation. Pencolé et al. published analytical work on decentralized alarm correlation and have also done some experiments [9], which evince the general feasibility of the decentralized alarm correlation idea.

III. THE MADEIRA SOLUTION TO FAULT MANAGEMENT

This chapter briefly introduces the architecture of Madeira by describing its layering and services, discusses how the logical overlay network is constructed and finally describes the functional components and achievements of the Madeira Fault Management (FM) application. A more detailed description of the architecture and the overlay can be found in [1].

A. Architecture

The Madeira software, which is deployed on each network element, consists of several loosely coupled components arranged in a layered structure as illustrated in Fig. 1.

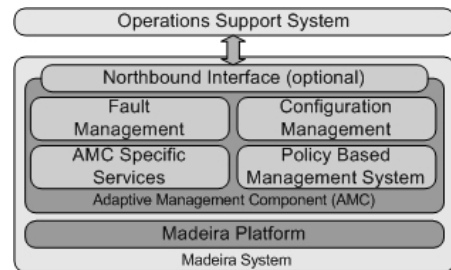


Figure 1. Architecture of Madeira

The base layer is the *Madeira Platform*, a P2P middleware implementing basic lifecycle management and communication services needed for network management in a peer-to-peer environment. Mostly relevant for this paper are the *Directory Service* which keeps track of the one-hop neighbours of a peer and the *Grouping Service* responsible for dynamically forming and monitoring the logical clustering. The *Adaptive Management Component* (AMC) encapsulating basic management functions constitutes the next layer. From an FM point of view three AMC services are worth mentioning: On the one hand the *Policy Based Management System* (PBMS) used for dynamically configuring fault formatting and fault correlation rules, on the other hand the *Notification Correlation Service* providing the generic logic used for fault correlation, and finally the *Network Element Adapter* responsible for interactions with the physical network element layer. Beyond these basic services arbitrary application modules may be implemented as a part of the AMC. Finally, on top of these Application Modules resides the *Northbound Interface* (NBI), representing the link between the Madeira Management System and any Operations Support System (OSS). The NBI is currently implemented by means of Web Services and Web Service Notifications. As an optional module the NBI is not required on all nodes.

B. The Madeira Distributed Overlay Network

To overcome the drawbacks of first generation peer-to-peer networks [1] the Madeira Grouping Service establishes a distributed overlay network by logically partitioning the network into management clusters with a cluster head (super peer) assigned to each of them. In opposite to typical hybrid P2P networks these cluster heads at level 1 are grouped into (level 1) clusters again and elect their level 2 cluster heads. This process is repeated until a single top level cluster head

remains at the highest level. This overlay network offers an efficient method of message passing and management task delegation down to the lowest cluster able to carry out the delegated task, reducing message load and therefore increasing scalability. Note that our approach to hierarchy establishment is completely different to the usual hierarchical network management as the overlay network of Madeira is dynamic and self-organizing, and supports automatic reconfiguration in case of failures. Therefore it increases robustness and avoids single points of failure. Further details on the logical overlay are discussed for two example faults in chapter IV.

C. The Madeira Fault Management Application

The simple but nevertheless challenging goals for a fault management application are (1) automatic (and autonomic) resolution of a problem as far as possible and (2) reporting of the identified fault – with the unique root cause – to the operator in a simple and meaningful way.

The basic idea of our decentralized approach is that on each node in the network the FM application tries to obtain a consolidated view of a possible problem before forwarding its information to the next level (its cluster head) using the logical overlay network of Madeira. This procedure is iteratively repeated up to the top level cluster head.

The functional components of our FM application are:

- *Fault Formatting*: Faults can be detected (and corresponding alarms generated) on any layer of the Madeira architecture. These different types of alarms are transformed into a common alarm data format by means of the PBMS, using specialised fault formatting policies.
- *Notification Correlation Service (NCS)*: The aim of the NCS is to determine related alarms for each newly generated Madeira alarm notification, and to merge those into a new one containing the same or more information as the source alarms, using simple correlation rules specified by means of policies.
- *Fault Analysis*: According to the result from the NCS and possibly from any other impact analysis that has been carried out the fault analysis module decides whether it is necessary to forward this information to the next level. Interaction with other distributed NMS applications (e.g. trigger immediate reconfiguration due to a fault) can easily be integrated here.
- *Fault Reporting*: Finally the application has to query the next level cluster head to which the alarm is forwarded; at the top level the final alarm containing the determined root cause is transformed into a X.733 compliant alarm which is forwarded to the NBI for northbound reporting.

We expect, amongst others, following achievements of the Madeira approach to Fault Management:

- *Scalability*: The delegation of alarm processing to lower cluster levels reduces load of single cluster heads (especially the top level cluster head). Also, only one or a few alarm notifications per fault have to be forwarded through the whole network, as most of them will typically be correlated into single consolidated alarms on low levels.
- *Availability and robustness*: The fact that even the tasks of higher level cluster heads are not different to that of normal

nodes and do not need much more computational power allows for machines to dynamically hand over responsibilities (automatic reconfiguration of the logical overlay). This increases robustness of the whole management network and avoids single points of failure.

- *Consistency in dynamic networks*: The fact that alarm analysis and correlation is performed locally ensures that the information on topology and other relevant network data used for this process is up-to-date. This is most important in highly dynamic network topologies, where it would otherwise be almost impossible to distinguish between network configuration changes and real faults.
- *Flexibility through policy-based approach*: Fault formatting, alarm correlation and fault analysis have been implemented based on policies [7] which can be updated dynamically. Further, by virtue of the distributed PBMS approach, policies inserted or updated by an operator through the NBI are immediately distributed to all Madeira nodes, ensuring consistency.
- *Simple fault analysis and correlation rules*: One can assume that the set of correlation rules required on a single node (i.e. locally) will always be rather simple, due to the limited scope of network and topology knowledge. The emerging rules for global alarm correlation might then get arbitrary complex in large networks. Unlike in traditional centralized alarm correlation approaches, however, knowledge of these complex global correlation rules is not required at all.

IV. EXAMPLE FAULTS IN A WIRELESS MESH NETWORK SCENARIO

We briefly discuss the handling of two simple fault types (node outage, link outage) taken from the scenario of managing a wireless meshed network which has been used for investigating the Madeira solution [1]. Determining the root cause for more complex faults is possible using similarly simple rules as for these basic ones.

In the following examples *Directory Service* (DS) alarms which are generated at a node whenever a one-hop neighbour is lost or a new one appears and *Grouping Service* (GS) alarms that are issued at a cluster head whenever a node joins or leaves its cluster are relevant as primary alarms. They are assigned a *level* attribute according to their position in the logical overlay (only GS alarms may be generated at levels unequal to zero).

1) Outage of a Base Station

There is a failure of a single base station (BS01) which is physically depicted in Fig. 2 (level 0 clusters are encircled by a line). We assume BS02 (and BS09) to be the level 2 cluster heads of BS00 and BS02 (BS06 and BS09) and BS09 being the top level cluster head at level 3.

Outage of BS01 now leads to the generation of DS alarms on all one-hop neighbours (i.e. BS00, BS02, BS03, BS04, BS05, and BS08) and a GS alarm on its cluster head (BS00). As only one alarm is generated on each node at level 0, no correlation is possible on this level. Therefore all alarms are forwarded to the corresponding cluster head (level 1). Each cluster head that received multiple alarms at level 1 (BS00, BS02) is now able to identify them as related (by comparing the alarm attribute containing the node identifier) and merge

them into a single alarm. Following this concept, the amount of alarms is reduced step by step to one single alarm which is then forwarded until the top level cluster head is reached. Fig. 3 summarizes this alarm flow from level 0 to the top level.

Note that in real world networks consisting of much more nodes than in our example, all alarms would most probably still be correlated into a single one at a low level while higher level cluster heads will only need to forward this final alarm. Additionally, this example shows how load is distributed across the network. Out of seven alarms being initially generated, no cluster head has to consider more than three alarms simultaneously.

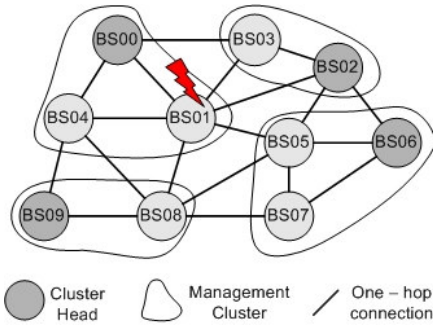


Figure 2. Outage of BS01

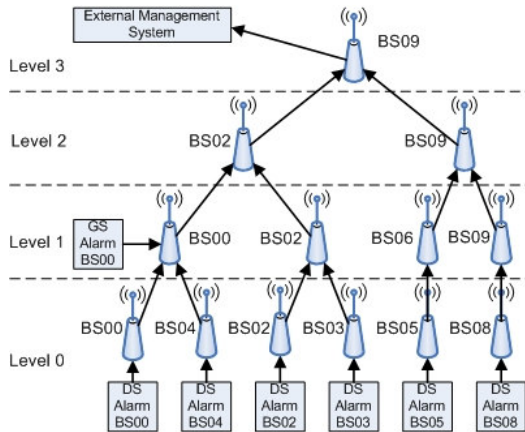


Figure 3. Alarm Flow after outage of BS01

2) Outage of a Link between two Base Stations

Our second example assumes a physical link outage between two base stations (BS01 and BS03) belonging to different clusters. (We assume both physical topology and logical overlay to be equal to the previous example).

Here, two DS alarms are generated, one at BS01 and one at BS03. Due to the lack of additional knowledge both are forwarded to the corresponding level 1 cluster heads (BS00 and BS02) which also can only report to their common level 2 cluster head (BS02). BS02 at level 2 can now easily relate these two alarms to each other by a simple rule (node instances of reporting and affected node are permuted) and merge both alarms into a new one with the correct root cause.

Similar correlation rules have been applied within our case study to identify further issues like controlled shutdown of nodes, interface outages of gateway nodes, outages of cluster heads, and finally also clearances of all these problems.

V. CONCLUSION AND OUTLOOK

We have pinpointed some key challenges for management of NGNs and presented a distributed network management solution based on peer-to-peer paradigms designed to master these challenges. We focused on the area of fault management, explaining our FM application prototype and discussing possible benefits of our decentralized solution, supported by a case study from the area of wireless mesh networking. Our approach promises to exhibit functional benefits – a simple set of correlation rules proved to be able to manage quite complex fault management scenarios – as well as to address key non-functional requirements like scalability and robustness. We also saw crucial benefits from having a flexible interaction between the functional areas of CM and FM, that facilitates more responsive and pre-emptive network management of highly dynamic, distributed networks.

In order to fully evaluate the extent, however, to which the predicted scalability behaviour is realised, a more detailed verification using complementary approaches such as theoretical evaluations, simulations, and experiments on large-scale real test beds is required and is currently in progress.

REFERENCES

- [1] P. Arozarena et. al., "Madeira: A peer-to-peer approach to network management," Proceedings of the Wireless World Research Forum Meeting 16, Shanghai, China, 2006
- [2] A. Binzenhöfer, K. Tutschku, B. Graben, M. Fiedler, and P. Carlsson, "A P2P-based framework for distributed network management," Report No. 351, Institute of Computer Science, University of Würzburg, 2005
- [3] R. Boutaba and J. Xiao, "Network management: state of the art," Proceedings of IFIP World Computer Congress (WCC'02) TC6 Stream on Communication Systems: The State of the Art, Montreal, Canada, pp. 127-146, 2002
- [4] Celtic-Madeira project, <http://www.celtic-madeira.org>
- [5] G. Goldszmidt, S. Yemini, and Y. Yemini, "Network Management by Delegation: The MAD Approach," In Proc. 1991 CAS Conf., Toronto, Canada, pp- 347-359, 1991
- [6] T. Magedanz and T. Eckardt, "Mobile Software Agents: A new Paradigm for Telecommunications Management," In Proc. of NOMS'96, pp. 260-269, Kyoto, Japan, 1996
- [7] R. Marin et. al., "A distributed policy based solution in a fault management scenario, In Proc. of Globecom 2006, San Francisco, USA
- [8] J. Martin-Flratin, S. Znaty, and J. Hubaux, "A Survey of Distributed Enterprise Network and Systems Management Paradigms," Journal of Networks and Systems Management, 7(1), pp.9-26, 1999
- [9] Y. Pencolé, M. Cordier, and L. Rozé, "A decentralized model-based diagnostic tool for complex systems," 13th IEEE International Conference on Tools with Artificial Intelligence (ICTAI'01), pp. 95-102, Dallas, USA, 2001
- [10] A. Pras, B. Beijnum, and R. Sprenkels, "Introduction to TMN," CTIT Technical Report 99-09, University of Twente, 1999.
- [11] R. Steinmetz and K. Wehrle (Eds.), "Peer-to-peer systems and applications," Lecture Notes in Computer Science 3485, Springer Berlin, 2005