

# Madeira: A peer-to-peer approach to network management

Pablo Arozarena Llopis, Martijn Frints, David Ortega Abad, Javier González Ordás

TELEFÓNICA INVESTIGACIÓN Y DESARROLLO

Liam Fallon

ERICSSON R&D

Martin Zach

SIEMENS AG

Hai Nguyen Thi Van, Joan Serrat Fernández

UNIVERSITAT POLITÈCNICA DE CATALUNYA

Current management approaches like the Telecommunications Management Network are based on a hierarchical structure and use rigid interoperability standards that make them less applicable to dynamic, heterogeneous network environments. This article will give an overview of the innovative Madeira management solution, which is based on distributed, peer-to-peer principles.

Madeira provides novel technologies for a logically meshed Network Management System that facilitates self-management and dynamic behaviour of nodes within the network. It performs network management tasks that are difficult to carry out using conventional methods, or tasks that can be carried out more efficiently using a distributed approach. The focus of Madeira lies on Configuration and Fault Management, and the collaboration between these two functional areas. The management functions are executed in peer-to-peer aware Adaptive Management Components, or AMCs. By using the peer-to-peer interface, these AMCs interact with each other and create an Overlay Management Network in order to perform specific management tasks. The behaviour of the management application is influenced by the use of policies. To ensure the scalability of the network, the nodes are grouped into Management Clusters. Each cluster has a Cluster Head that is responsible for the coordination and topology publishing of its cluster.

Madeira also offers support for a higher layer Operation Support System, or OSS. It can issue alarms containing meaningful information for the operator, publish the complete network topology and receive management commands from this external OSS.

In order to demonstrate the capabilities and strengths of the Madeira management approach, a number of scenarios are included that describe real life problems that occur in WiFi networks, and that are difficult to solve with the traditional management approach.

## INTRODUCCIÓN

Telecommunication networks are getting larger and more heterogeneous. These large, transient next-generation

networks have a great need for automated network management in order to reduce operating expenses (OPEX). The global adoption of IP-based networking solutions since the mid-1990's and the recent emergence of the

Internet as the platform for next generation services has highlighted the need for a powerful management technology and for an overall framework for the automated management of emerging networking infrastructures. Broadband access technologies in particular have started to act as a catalyst for the radical substitution of public switched telephone services with Voice over IP (VoIP), while optical switching technologies have started to be deployed both in the access and core network. Furthermore, wireless communication has been gaining ground in the last few years. Besides the extensive growth of mobile telephony, also WiFi networking is becoming more and more popular. The emerging environments will have new and more sophisticated management requirements, similar to those of telecommunication networks.

In the above context, the use of traditional hierarchical management techniques suffers a number of disadvantages like static architectures and rigid standards, which make them less suitable for these large and dynamic environments. Initial and well established frameworks in the telcos networks, such as the OSI System Management (OSI-SM), using the Common Management Information Protocol (CMIS/CMIP), were management-specific approaches adopting the manager-agent model. This is the foundation of the Telecommunications Management Network (TMN), which was defined by ITU-T as a framework for the management of communications networks in ITU-T M.3000 recommendation series. In the mid-1990's general distributed object technologies (DOTs) emerged, with the Common Object Request Broker Architecture (CORBA) being the most representative one, spurring a lot of research and standardization activities regarding their use for management. Since then, OSI-SM has been confined to telecommunication environments for *network* management purposes, while CORBA and Java-based distributed object technologies have been used for *service* management.

A next generation management technology is required, together with an overall framework for network-wide optimal configuration according to well-defined goals and constraints. Approaches based on network programming, management overlays and peer-to-peer (P2P) computing, as well as distributed aggregation and control schemes, have been recently proposed to engineer management systems with good scalability behaviour and that are robust regarding topology changes and failures.

This paper describes a solution to overcome the drawbacks of present day management approaches: the solution developed in the Madeira<sup>1</sup> project. Madeira is based

<sup>1</sup> Madeira is a two year project of the pan-European CELTIC-Initiative. The consortium consists of Ericsson Ireland, Siemens AG, British Telecom, Telefónica I+D, Telecommunications Software & Systems Group, Universitat Politècnica de Catalunya and Ericsson Research Sweden.

on peer-to-peer technologies for a logically meshed Network Management System that facilitates self-management and dynamic behaviour of nodes within the network. After this introduction, the paper continues with a section devoted to the highlights of current and state of the art solutions for network management systems. This serves as a framework to understand the impact of Madeira's proposal. After a short introduction on the Madeira management approach, the details of the Madeira Architecture are presented. The scenarios that will be used for evaluation during next months follow this. Finally, the paper ends with some concluding remarks.

## CURRENT NETWORK MANAGEMENT APPROACHES AND TRENDS

### Management frameworks

In the past years, TMN has been the dominant network management framework. It promotes a well known centralized approach [16], which has a number of effects on the scalability of the network management application. Managing a large network from a single, central point will increase the load of the central manager and could create bandwidth bottlenecks on links that are close to that central manager [8].

Another disadvantage is the lack of flexibility, since the current generation of management architectures use static topology data, often based on manually generated files. This complicates reconfiguration of a dynamically changing network. Changes in the topology are normally implemented from the top of the management pyramid, and pushed down into the network. Fault Management is often based on the same static topology that is saved in the management system. For example, Alarm Correlation normally assumes a static topology, while in meshed networks this topology is dynamic and time-based.

TMN is divided into five layers, depicted in the pyramid in **Figure 1** [16]:

1. The *Business Management* takes care of the management of the whole enterprise. It can be seen as goal setting rather than goal achieving.
2. The *Service Management* is in charge of managing aspects that are directly observed by the users of the telecommunication network. It addresses topics like customer care, service development and operation, Quality of Service management, accounting, etc.
3. The *Network Management* provides management services that are related to the interaction between multiple pieces of equipment.

4. The *Element Management* takes care of vendor specific management functions, which are hidden from the Network Management layer above. For example, it can detect equipment errors and perform measurements on temperature and resource usage.
5. The *Network Element*. It provides agent services, mapping the physical aspects of the equipment to the TMN framework.

As an evolution of TMN, the TeleManagement Forum, or TMF, proposed the Telecom Operation Map (TOM) [22] and, more recently, the eTOM. These models describe at a high level the processes a telecom operator needs to fulfil to manage its network and services infrastructure. Furthermore, TMF has defined the NGOSS architecture, which is a technology agnostic framework for the construction of management applications. NGOSS fosters a component based architecture with interfaces between components defined as contracts, a shared information model and the separation of implementation from business logic (see [11] and [21]).

An interesting implementation of NGOSS concepts is OSS/J, which pursues the implementation of a set of APIs, based on J2EE technologies, to allow the integration of OSSs [13].

It is also worth mentioning the 3GPP approach to network management, which is based on the IRP (Integration Reference Point) concept. The IRP is analogous to the TMN Q3 reference point, improving it by defining the information models in an implementation independent UML.

An important assumption behind TMN and similar frameworks is that network elements have limited management capabilities, being focused on their communications role. Therefore, management functions are performed externally by dedicated systems; while network elements only provide simple management agents to allow these external systems access and manipulate management data. It was recognised considerable time ago, however, that network devices can do much more than running a simple agent, being even capable of managing themselves [25]. This is the approach taken in IP networks, where nodes are able to perform management tasks for routing, signalling, path provisioning, etc.

Following this trend, the control plane paradigm has emerged in telecom networks, giving more autonomy to network elements for certain tasks. Particularly in optical networks, research is being conducted to create an optical control plane that enables automated multi-vendor network operation. Examples are the ITUs Architecture for Automatically Switched Optical Networks (ASON) and the IETFs Generalized Multi-Protocol Label Switching (GMPLS) [7]. However, decentralized standards are not

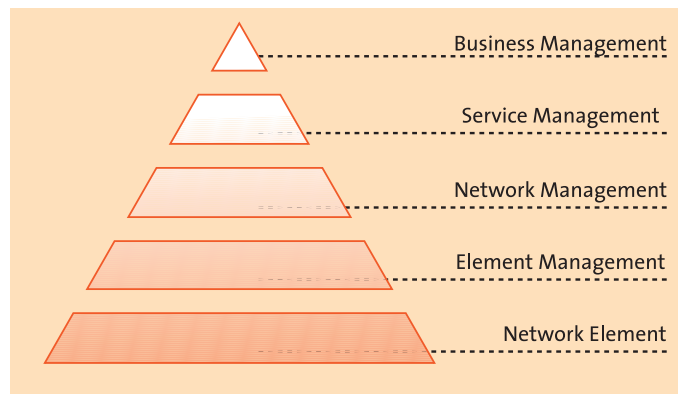


Figure 1. TMN Pyramid

(yet) available for wireless networks.

Note that the existence of a control plane does not mean that the management plane is no longer necessary, but rather that it must adapt to the new scenario, focusing on the tasks for which it is better suited and relying on the control plane for other tasks such as routing and signalling [1].

### Distributed technologies

One of the main problems faced by the network and service management field is the integration of different platforms and applications [9]. Distributed technologies and frameworks have greatly helped to solve this problem, avoiding the need for dedicated or proprietary protocols.

The distributed system management approach with the most impact has undoubtedly been the introduction of the Object Management Group (OMG) Common Object Request Broker Architecture (CORBA). Reasons for using CORBA in TMN environments are that OSI management system was conceived as an object-oriented management technology in the absence of a distributed object-oriented framework. CORBA provides exactly such a framework, with a superior distribution paradigm in which every object could be potentially distributed [3]. Moreover, the performance could also be better than OSI management due to a more lightweight protocol stack than the conventional CMIS/CMIP.

The use of CORBA has been proven valuable both in managing several network technologies and in the challenging area of IP and WDM network integration [6] and has also deserved the attention of the standardization bodies. Nevertheless, the use of CORBA is not exempt of disadvantages. In fact, the OSI object management model is more expressive and the mechanism for accessing objects is superior thanks to the incorporated scoping and filtering concepts. Finally, a distribution mechanism of pre-filtered events makes it appropriate to deal with alarm and event notifications. For these reasons, the idea was to

keep the OSI management model and tools for Element Management and Network Management whilst using CORBA for the upper layers of the TMN pyramid. But the coexistence of both models poses serious compatibility challenges addressed by several authors [14].

Web Services [24] is an emerging Internet-oriented technology that has strong analogies to CORBA. Although in principle not conceived for that purpose it is clear that due to this analogy it could also be used for network management. Currently, research on the use of Web Services for network management is being conducted by organizations as OASIS. Originally, Web Services were stateless entities. The OASIS Web Services Resource Framework [28] solves this problem by defining a framework for stateful resources. Besides, OASIS defined the Web Services Notification standard allowing the usage of a publish-subscribe communication pattern. Among others, these two standards have been used in the Management Using Web Services specification, which is part of the OASIS Web Services Distributed Management Technical Committee [26]. It enables the management of distributed resources. Although these solutions are very promising, it remains important to consider aspects pertinent to Web Services (XML encoding overhead, etc.). In that respect, the fact that Web Services is an XML-based technology makes it very attractive due to its potential easy integration with other applications but this also causes the drawback of higher overheads than CORBA.

Another approach to decentralized network management tasks is the usage of Mobile Agents, which are small software programs that “travel” through the network. These agents decentralize processing and control by remotely performing certain management tasks [2], but the security requirements on the execution environment are extensive, resulting in performance and functional limitations as indicated in [18].

### Peer-to-peer

P2P systems, which have proven to be highly scalable and robust, can be utilised to establish service specific management overlays. Making management information available via such overlay networks has great advantages

since it avoids dedicated management systems and supports fast deployment of new services together with their effective management. According to [4], peer-to-peer systems can be characterized as distributed systems in which all nodes have identical capabilities and responsibilities and all communication is symmetric. In [19], Schollmeier indicated that some define peer-to-peer networks as a collection of heterogeneous distributed resources, which are connected by a network, while others define it simply as the opposite of client/server architectures.

In a P2P system, the nodes have a significant or total degree of autonomy from central servers. As pointed out by [20], P2P systems enable the utilization of previously unused resources as storage, cycles or content for example, by tolerating and working with the variable connectivity of numerous devices.

An overall characteristic of a peer-to-peer network is that the nodes can send and receive information in a way that makes them both servers and clients, or “servents”. In both [19] and [10], a distinction is made between pure peer-to-peer networks and hybrid peer-to-peer networks, in such a way that:

- *Pure P2P architectures are completely decentralized.* There is no central server or router. Each node can issue and respond to requests, or route requests to other nodes.
- *In Hybrid P2P architectures, more types of nodes exist.* The leaf nodes are nodes with an information need or information resource. In other words, they can provide information to or request information from other leaf nodes. Another type of nodes, super peers, have a more “server-like” role in the network. These nodes provide regionally centralized services to the network in order to improve the routing of information requests. In [10], these nodes are called directory nodes or ultrapeers. Each directory node provides directory services for portions of the network and directory nodes work in a cooperative manner to cover the whole network.

Figure 2 gives an example of both types of peer-to-peer networks.




Figure 2. Pure peer-to-peer (left) and hybrid peer-to-peer (right)




Figure 3. Overlay Management Network using AMCs

## THE NETWORK MANAGEMENT APPROACH IN MADEIRA

In an attempt to overcome the shortcomings of the traditional management approaches to face the challenges of next generation telecommunication networks, Madeira aims to develop a new management framework based on peer-to-peer networking concepts. Furthermore, it provides novel technologies for a logically meshed Network Management System that facilitates self-management and dynamic behaviour of nodes within the network. Madeira also takes advantage of the Policy Based Management Paradigm [17] that pursues the separation of management logic from the actual applications. This logic is then specified as a set of rules or policies that can be dynamically fed into the management system allowing a change of its behaviour without the need of changing the application or even restarting it. Besides the innovative architectural framework, the Madeira project will provide interface protocols, standards and a reference software implementation and apply it to a specific network management scenario. Ultimately, by enabling the management of network elements of increasing numbers, heterogeneity and transience, the Madeira approach should reduce the Operational Expenses, or OPEX [29].

Madeira focuses on Fault and Configuration Management functional areas and, especially, on the way they can co-operate to solve management problems. In doing so, it will act as a complement to traditional management systems. There are many management tasks that current network management systems perform well, such as Performance Management. For these tasks, a hierarchical approach is entirely appropriate. Madeira will investigate the feasibility of distributing management responsibilities among peer nodes in order to perform certain

tasks more efficiently. In other words, the Madeira management approach is applied to management tasks that are difficult to carry out using conventional methods, or tasks that can be carried out more efficiently using a distributed approach.

## THE MADEIRA ARCHITECTURE

### AMCs and the Overlay Management Network

The approach of Madeira is completely different than in traditional management systems. It encompasses a much flatter structure that is based on peer-to-peer (P2P) principles. The management functions are executed in P2P aware Adaptive Management Components (AMCs), which correspond directly to the Network Elements (NEs). By using a well-defined east-west (or peer-to-peer) interface, these AMCs can communicate with each other, creating an Overlay Management Network. **Figure 3** gives a graphical representation of this overlay.

Another basic principle in Madeira is to delegate management functions (and corresponding management information) down to the Network Elements as far as possible. It makes no real difference whether management functionality is residing on the Network Elements themselves or on dedicated Network Management nodes (as has been depicted in **Figure 3**). By making use of the east-west interface introduced with the P2P paradigm, both become part of the Overlay Management Network. Such a configuration of AMCs residing on each NE as well as on dedicated Network Management nodes is shown in **Figure 4**.

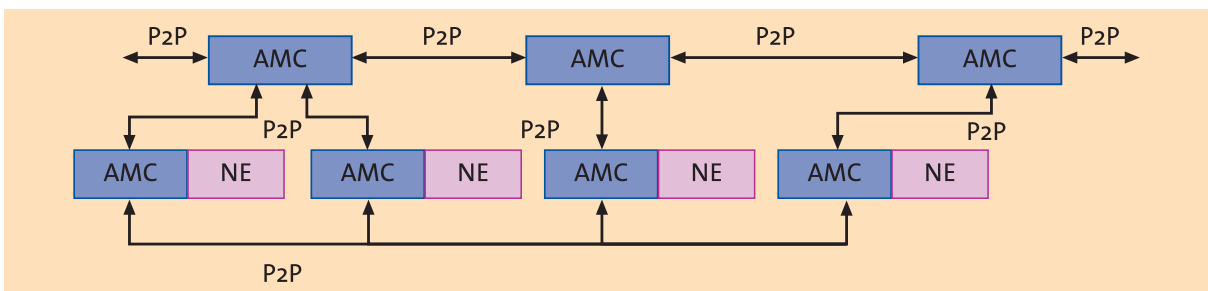


Figure 4. Overlay Management Network including AMCs on NEs

As mentioned before, an Adaptive Management Component (AMC) fulfils the network management functionality of a peer in a P2P network. It can exchange management information between peer management applications in a network element or management node. The AMC has the ability to import (and export) functionality to perform specific tasks. One or more AMCs interact with each other to perform a specific network management application.

In order to perform these tasks, an AMC requires an execution environment (a Container) and a variety of services. The Madeira platform provides both of these. The separation of the Madeira Management System between AMC and platform is depicted in **Figure 5**. The AMC covers the management specific parts for a particular scenario, while the platform provides all the generic functionality required to run tasks in a P2P environment. This separation of functionality enables Madeira to adapt to changing scenarios and requirements in an efficient way.

The same **Figure 5** also shows the high-level layering structure on a Madeira Management Network Element. The following groups of services are available in an AMC:

- The *Configuration Management and Fault Management* contain the specific network management applications. They and provide the ability to set-up the network, react to faults and other FM and CM related tasks. A description of how these tasks are performed will be described in the section dedicated to the scenario.

- The *Northbound Interface* is optional. It offers services that communicate with a higher layer Operation Support System (OSS) via Web Services. The OSS can, for example, retrieve information like network topology, events or alarms. More information on the connection with an external OSS will be given in the section “Connecting to the North”.
- The *AMC Specific Services* offers a base for the Network Management Applications. It mainly provides services to communicate with other AMCs. This can be either publish-subscribe based, or a direct peer-to-peer connection to another AMC.

As mentioned before, the AMC performs all tasks that are directly related to network management. To execute these tasks, the AMC needs services that are offered by the platform. The platform takes care of all additional functionality that is needed in a peer-to-peer environment. Its capabilities are divided into two groups of services:

1. The *Lifecycle Management Services* take care of the management of AMC containers. In more detail, this group offers the following services:
  - The *Lifecycle Service* can perform start/stop/restart operations on all modules loaded by the AMC.
  - The *Code Distribution Service* enables dynamic loading of application logic/data into the AMC. AMCs have a minimum “bootstrap” configuration to function in the network. Additional functionality can be imported as required with this service.

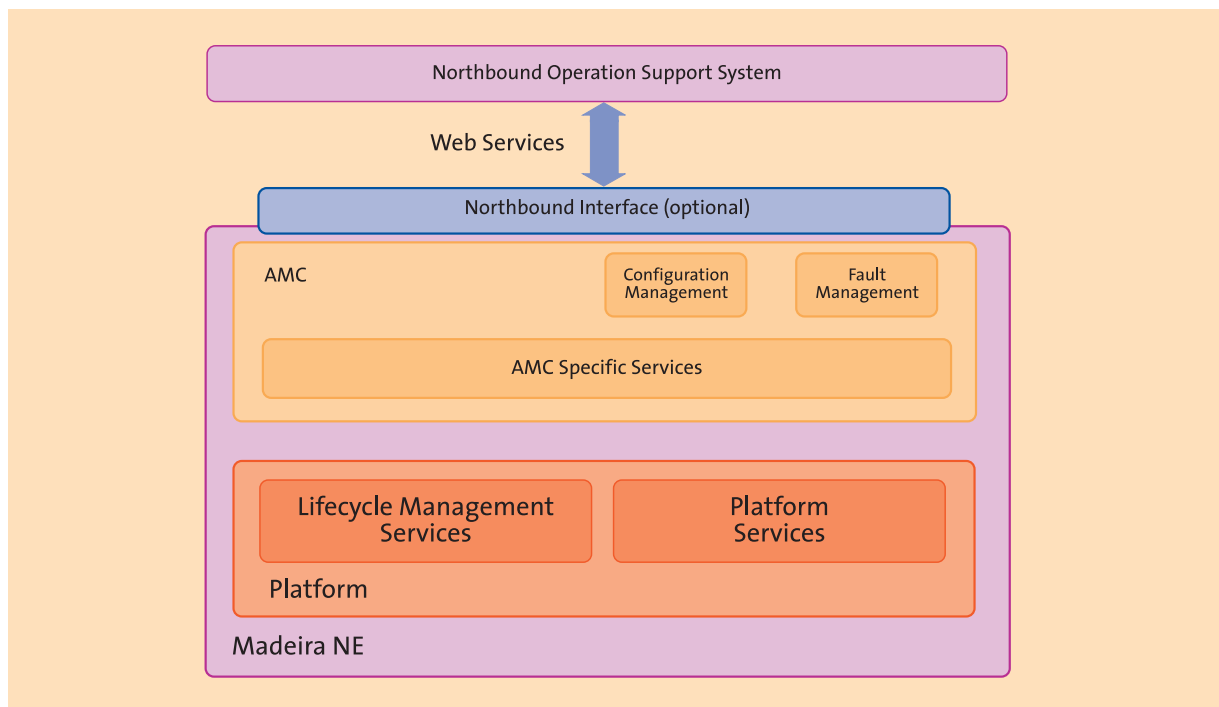


Figure 5. Layering of services on a Madeira Network Element

- The *Security Service* provides all aspects of security and authentication from a platform perspective.
2. The *Platform Services* offers the following services, which are specific for the peer-to-peer environment:
- The *Notification Service* is a basic event notification service, based on a standard publish-subscribe service. It enables AMCs to subscribe to certain event types.
  - The *Directory Service* is a directory of AMCs with their roles and capabilities. It keeps track of the physical one-hop neighbourhood of the NE and enables AMCs to be looked up
  - The *Connectivity Service* provides reliable one to one communication between two AMCs. This point-to-point connection supports multi-hop P2P communication.
  - The *Persistency Service*. With this service, AMC data can be stored for retrieval across restarts. It supports permanent storage of application-defined data.
  - The *Grouping Service* can dynamically form AMC groups for a given management function, It provides application partitioning for AMCs of similar roles or capabilities.

Using the AMC as a basis, four different types of interfaces are specified to facilitate communication with components:

1. The Peer-to-Peer interface, to communicate with other AMCs.
2. The Northbound interface, to communicate with another NMS or OSS.
3. The Southbound interface, to access NE application specific functionality.

4. The Interface(s) to the platform, to communicate with Lifecycle Management and Platform Services.

The relation between the AMC and these interfaces is shown in **Figure 6**.

### Clustering

The Grouping Service provided by the Platform Services offers the ability to create groups of AMC that perform a specific management function. These groups are called Management Clusters. Each Management Cluster contains exactly one network node that acts as the Cluster Head. This Cluster Head is responsible for coordination of and topology publishing for its cluster. The clustering principle makes Madeira a hybrid peer-to-peer solution, where the Cluster Heads can be seen as “super peers”. The other nodes are “normal peers”, as described before.

Different levels of clustering can exist, which leads to a clustering hierarchy. A node in a level “n” cluster is the Cluster Head in a level “n-1” cluster. The top level Cluster Head is responsible for publishing, for example, the topology of the complete network and events or alarms to a higher layer Operation Support System or another Network Management System. In order to perform this task, the top level Cluster Head will contain the Northbound Interface.

The use of this clustering concept is of great importance for the scalability of the management application. It creates the ability to divide a large network into smaller, dynamic groups to perform certain management tasks. The scenario described further on will elaborate more on the use of Management Clusters.

### Policies

In Madeira, policies are used to tune the behaviour of the management application. These policies are introduced to the system by the Northbound Interface and automatically distributed among the appropriate AMCs. Certain

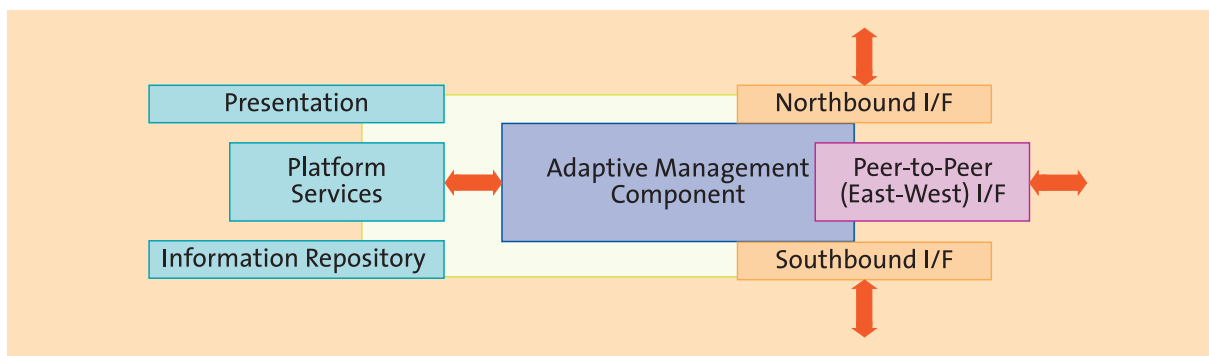


Figure 6. The four different types of interfaces

policies can be present during initialisation. To give an example of this rather abstract principle, the structure of the Management Clusters mentioned above is based on policies. They can, for example, control the number of nodes per cluster, but they can also contain criteria for the election of Cluster Head, such as memory resources, load etc.

### Connecting to the North

The Northbound Interface offers support for a higher layer Operation Support System (OSS) or another Network Management System to access Madeira's network management functionality. From this point on, both the Operation Support System and the Network Management System will be referred to as OSS.

#### Web Services

The communication between Madeira and an external OSS is based on Web Services. The information exchange between Madeira and the OSS will be encapsulated in SOAP messages. It is an XML-based message format and although it can be used over a variety of protocols, the use of SOAP over HTTP is most common.

The Northbound Interface has one corresponding WSDL (Web Services Description Language) specification, which is an XML description of the public interface to the Web Service. It contains information like protocol bindings, message formats and available operations and messages.

The network node that is running the Northbound Interface functionality can dynamically change as a result of network reconfiguration. Due to this, some functionality should be in place so the OSS can discover the new

address. Madeira will use an external UDDI (Universal Description, Discovery and Integration) registry to store the NBI address. It is an XML registry in which information on Web Services can be stored [23]. **Figure 7** shows the relationship between Northbound Interface, OSS and UDDI.

As can be seen in this figure, the Northbound Interface supports request/response as well as notifications as communication types between Madeira and the external OSS. For these notifications, the Northbound Interface acts as a notification producer [27]. The OSS will, naturally, act as a notification consumer. To be able to receive notifications, the OSS listens for SOAP requests. The most suitable solution would be that the OSS also implements a simple Web Service for this functionality.

#### Capabilities

The Northbound Interface propagates instructions from the OSS to the platform, and data from the platform back to the OSS. In more detail, the Northbound Interface offers the following services, divided into five groups:

1. The *OSS Connection Manager* enables the OSS to (re)discover the Northbound Interface node. It publishes, for example, the address of the Northbound Interface in an external UDDI registry. It also responds to polling messages sent by the OSS to check connectivity with Madeira.
2. The *Policy Manager* can be used to add, change or remove policies to the Madeira Management System.
3. The *Events/Notification Manager* receives alarms and events from Madeira, and converts these messages to a format readable by the OSS. Notifications will be sent to every subscribed consumer. This means that

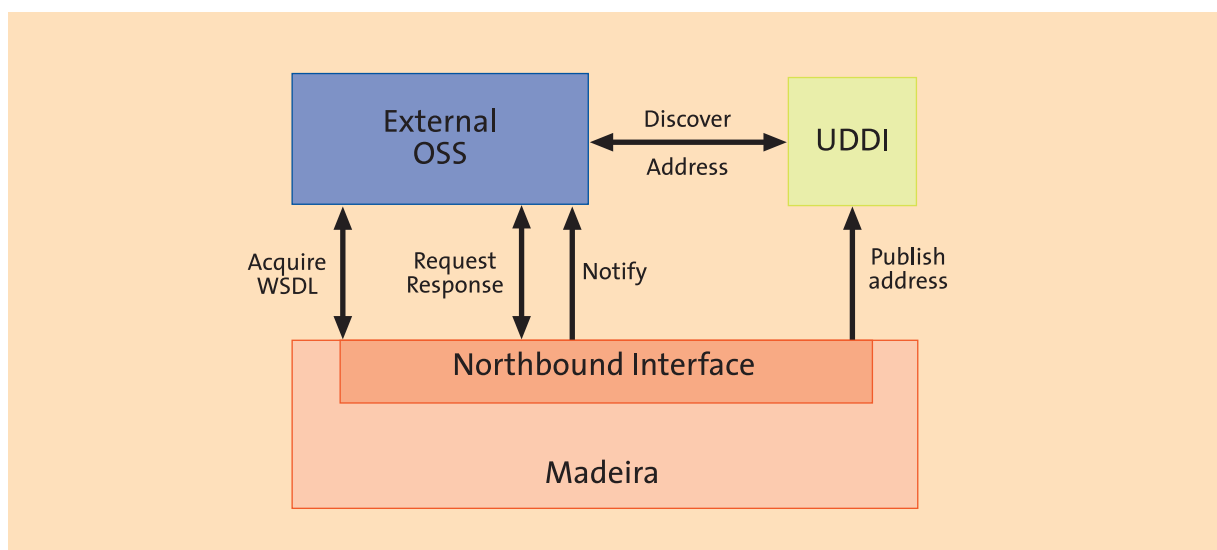


Figure 7. Usage of Web Services for the Northbound Interface

multiple Operation Support Systems can monitor the Madeira Management System.

4. The *Network Manager* provides information about the network topology and Management Clusters.
5. The *Command Injection Manager* receives commands from the external OSS, forwards them to the appropriate Madeira components for further processing and sends back the corresponding response to the OSS. The OSS can use the Command Injection Manager to disable Base Stations or to acquire a list of currently active alarms for example.

## THE MADEIRA SCENARIO

The goal of the scenario is to prove the capabilities of the Madeira approach to deal with real life management problems. It provides a number of challenging tasks to test the management approach. To emphasize the strengths of Madeira, the problems that arise in the scenario are difficult to solve with traditional management approaches, especially with respect to dynamic reconfiguration and changing topologies that occur in WiFi networks.

The scenario focuses on the areas of Configuration Management and Fault Management, with an emphasis on the integration between both of them. In the scenario, a number of wireless base stations are deployed in such a

way that wireless equipment (for example, laptops or PDAs) may have coverage from one or more base stations. Not every base station has a wired connection to the back haul network, as is the case in a traditional wireless network. Base stations directly connected to the backhaul network are called gateways. The rest of the base stations can only use a wireless connection to reach a gateway and thus the backhaul network.

### Configuration management

After deploying the base stations, Madeira automatically sets up the wireless meshed network using OLSR (Optimized Link State Routing protocol) as the routing algorithm. OLSR is a link stated routing protocol that is specifically developed for mobile ad-hoc networks [12].

Based on pre-installed policies, base stations are grouped into a number of clusters by the Grouping Service. These policies can be based on a number of criteria such as, for example, number of nodes per cluster or topological proximity. **Figure 8** depicts an example topology that could be the result of this process. As can be seen, clusters may or may not have direct backhaul connectivity. The network elements in a cluster monitor each other and exchange management information on a peer-to-peer basis.

As mentioned, wireless network equipment that wants to

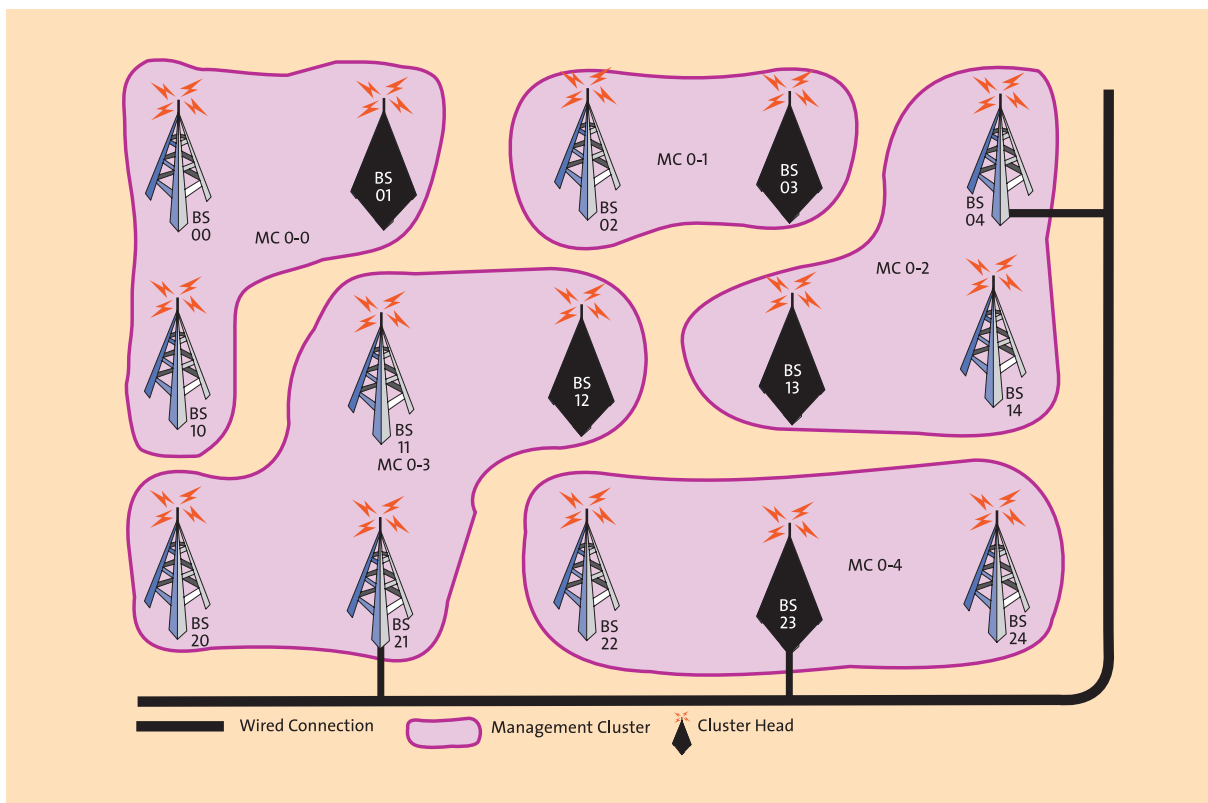


Figure 8. Management Clusters formed in the Wireless Mesh Network

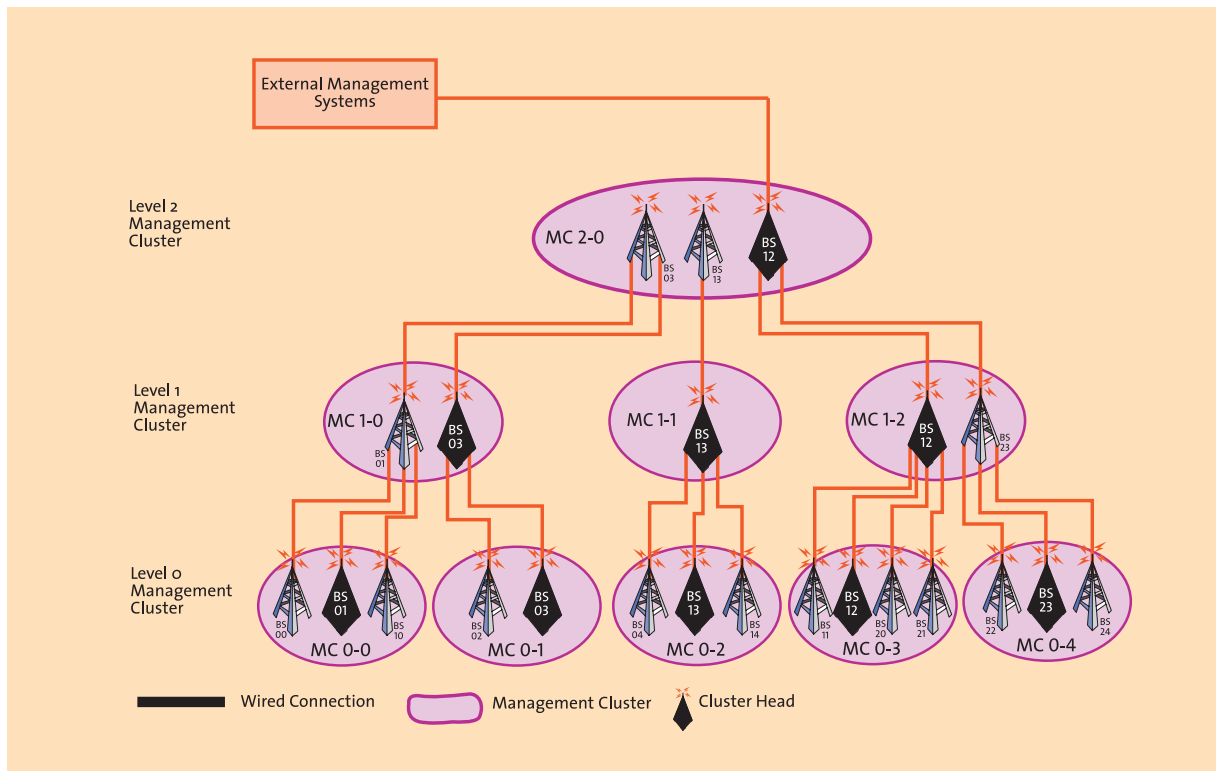


Figure 9. Management Cluster hierarchy

use the network is in range of one or more base stations. If the wireless equipment is in range of more than one base station, it selects one of them as its preferred base station, and uses this connection to use services on the Internet for example. If the wireless equipment is in range of just one base station, it must select that base station as its preferred base station.

Each cluster has exactly one Cluster Head. Policies are used for this election, and can be based on criteria like load, optimal connectivity or robustness. The Cluster Head is responsible for coordination and topology publishing of its cluster.

Different levels of clustering can exist in Madeira. **Figure 9** shows a possible clustering hierarchy that could be constructed in this scenario. The creation of this hierarchy is also based on policies. As mentioned in the previous section dedicated to the architecture, the top level Cluster Head is responsible for publishing the topology of the complete network. This can be done to a higher layer Operation Support System, or another Network Management System for example.

The cluster hierarchy is the basic management overlay that is used by all management functionality in Madeira. It creates a scalable environment for network management. The Madeira Configuration Management application is responsible for the construction, maintenance and viewing of the topology. Other applications, such as Fault Management, use this management overlay to implement

their functionality.

### Fault Management

During usage of the network, it is inevitable that unexpected faults occur. When such a fault occurs, it is important that:

- Appropriate action is undertaken quickly in order to reduce the service impact.
- Meaningful information on the fault is presented to the operator (in particular in those cases where automatic restoration is not or not fully possible).

Besides Configuration Management (CM), Madeira focuses on Fault Management (FM) and how CM actions and events are related to FM faults and alarms. Correlation between CM events and FM faults is an important aspect in order to discover the actual cause of a problem in the network.

Alarms can be generated by two different sources:

1. *Hardware level alarms* are generated by the base station in case of a hardware fault.
2. *Platform level alarms* are generated by either the Directory Service or the CM application. The Directory Service can indicate loss of connection with




Figure 10. Base Station outage

a neighbouring node, and the CM application can indicate changes in the topology (a node leaves or joins a cluster).

When a fault occurs, the FM application will receive one or more alarms. For example, a hardware level problem may also cause a fault on platform level, creating two alarms. These alarms are correlated into a new alarm by FM and sent to the Cluster Head. This Cluster Head also performs correlation of the alarm with alarms originating from other nodes in order to get a clearer picture of the probable cause and possible solution. It can then forward the alarm to a higher hierarchy level. This process is repeated until it reaches the Top Level Cluster Head, which can notify the Northbound Interface in order to produce an alarm for the external OSS.

This paper provides a few example scenarios in order to explain the basic concepts and functionality of the FM application in Madeira. These scenarios describe two faults with similar impact but very different in nature (in the first case a node goes down, while in the second there is just a loss of connectivity between two nodes), and focus on the way Madeira distinguishes these two cases by correlating alarms and CM events at different levels of the management hierarchy.

#### Base Station E outage

**Figure 10** depicts two Madeira Management Clusters, with node A and node G being the Cluster Heads. The solid lines indicate “physical” OLSR links between

nodes. The dotted line represents an inter-cluster connection between node E and node F.

When node E fails, the Directory Service of its one-hop neighbours D and F will notice this. Both nodes will notify their Cluster Heads that the link with node E has failed, and that therefore E might be faulty. Besides receiving this alarm from node D, Cluster Head A also receives a notification from the CM application that node E isn't part of the cluster anymore. It will then send an alarm with this knowledge to a higher hierarchy level. When the Cluster Head G of node F receives the alarm that node E isn't reachable, it will also forward this alarm to a higher hierarchy level.

After receiving the alarms from node A and G, the higher level Cluster Head tries to correlate the information with other alarms. Since both alarms contain the same knowledge (a possible fault of node E) they will be merged to a single alarm with the same content. Afterwards the alarm will be forwarded up the hierarchy pyramid, until the Top Level Cluster Head will notify the NBI, which will inform the external OSS on the failure of node E.

#### Link outage between node D and node E

In this scenario, shown in **Figure 11**, the link between node D and node E fails. The link between E and F remains intact. Both nodes D and E will receive a notification from their Directory Service indicating the neighbouring node is no longer reachable. Node D concludes E might be faulty and forwards this information to its




Figure 11. Link outage

Cluster Head A, who also receives a notification from the CM application that E is no longer part of its cluster. After combining this information, it will send an alarm to the next hierarchy level, identical to the previous scenario. Besides receiving the notification from the Directory Service, node E will recognize that it is no longer connected to its cluster head and it will join another cluster (it is assumed E joins the cluster containing F and G). After this reconfiguration process E will forward the information that D is no longer reachable to its new cluster head G, which will additionally receive a notification from the CM application that E has joined its cluster and forward this information to the next hierarchy level.

The higher level Cluster Head of this next hierarchy level receives the alarm from A, indicating that node D reported that E is unavailable and probably faulty. It also receives a notification from the CM application indicating that node E has joined another cluster, and an alarm from G indicating that node E reported that D is unavailable and probably faulty. After correlating these three notifications, the Cluster Head concludes that a link outage between D and E occurred and either suppresses this alarm (if configured to do so) or forwards this knowledge as a minor alarm up the hierarchy pyramid until the Top Level Cluster Head and NBI is reached, which will then inform the external OSS on the link outage between D and E.

## FUTURE WORK

Currently the project is focusing on the implementation and integration of the different components that encompass the Madeira solution. A prototype management system dealing with specific Configuration and Fault management scenarios of WiFi networks will be developed and tested on top of a test bed provided by the partners. At the time of writing, several aspects have already been implemented. The release of a first prototype is planned for the beginning of this year. A second iteration demonstrating the main aspects of Madeira will be finished at the end of the project, in July 2006.

Although it is not explicitly mentioned as being in the scope of Madeira, some efforts are being undertaken to research security mechanisms for peer-to-peer environ-

ments. In such environments, where there is no centralized security solution, security and trust are important aspects. A possible solution for this could be to use Public Key Cryptography to create a web-of-trust, similar to the PGP [15] approach [5]. This should not be mistaken with the FCAPS Security Management area, which focuses on authorized use of the network, data integrity and confidentiality for example.

## CONCLUSIONS

The Madeira Management framework proposes the use of peer-to-peer techniques to fulfil management tasks. The ability to perform self-management and the usage of Management Clusters solves scalability issues that are present in current hierarchical network management approaches. So-called Adaptive Management Components, or AMCs, that run on network elements, are in charge of performing the various management tasks. By using the peer-to-peer interface, these AMCs can communicate with each other, creating a Management Overlay, in order to execute management tasks and make up network management applications.

Furthermore, because of the peer-to-peer concept, Madeira doesn't have a need for a central server. Besides reducing the operating expenses, this also eliminates the single point of failure that exists in traditional systems.

Madeira also offers support for a higher layer Operation Support System (OSS). Such an OSS can access Madeira via the Northbound Interface. This Web Services based interface enables an OSS to acquire topology information, receive alarms, introduce policies and perform other management tasks. By using a publish/subscribe system for notifications from the network, Madeira can notify multiple external systems about events or alarms at the same time.

In order to prove the feasibility of the Madeira approach, a challenging scenario has been identified, dealing with Configuration and Fault Management of highly dynamical wireless networks. Based on the Madeira framework, a Management System addressing this scenario will be prototyped and tested on a real test bed.

## GLOSSARY OF ACRONYMS

3GPP	3rd Generation Partnership Project, <a href="http://www.3gpp.org">www.3gpp.org</a> .
AMC	Adaptive Management Component. Peer-to-peer aware components that interact with each other in order to perform management tasks.
CM	Configuration Management.
CORBA	Common Object Request Broker Architecture. Language independent distributed object oriented framework.
eTOM	enhanced Telecom Operations Map.
FM	Fault Management
J2EE	Java 2 Platform, Enterprise Edition
NBI	Northbound Interface. Enables an external OSS to acquire management data and configure the Madeira management system.
NE	Network Element.
NGOSS	New Generation Operations Systems and Software.

OLSR	Optimised Link State Routing protocol.
OSS	Operation Support System. System to support the basic operation for service providers.
OSS/J	Operation Support System through Java.
P2P	Peer-to-Peer. Communication structure where nodes can communicate directly to each other without interference of a central entity.
TMF	TeleManagement Forum.
TMN	Telecommunications Management Network. Framework for the management of communications networks, developed by the ITU-T (International Telecommunication Union - Telecommunications sector).
TOM	Telecom Operation Map.
UDDI	Universal Description, Discovery and Integration. Platform-independent XML registry for Web Services.
WSDL	Web Services Description Language.

## REFERENCES

1. Bela Berde, Carolina Pinart, Javier Gonzales Ordas, Piet Demeester and Koen Casier: *An Experience on Implementing Network Management for a GMPLS Network IV Workshop in MPLS/GMPLS networks*. 21-22 April 2005, Gerona, Spain. [bcds.udg.es/gmpls/ws4/papers/wgn4\\_s5\\_p1.pdf](http://bcds.udg.es/gmpls/ws4/papers/wgn4_s5_p1.pdf).
2. A. Bivens, R. Gupta, I. McLean, B. Szymanski and J. White: *Scalability and performance of an agent-based network management middleware*. Int. J. Network Mgmt 2004; 14: 131-146.
3. Object Management Group: *The common object request broker: Architecture and specification*, CORBA. Version 2.0, 1995.
4. P. Druschel and A. Rowstron: *PAST: A large-scale, persistent peer-to-peer storage utility*. HotOS VIII, Schloss Elmau, Germany, May 2001.
5. L. Hyun-rok: *A distributed trust model for peer-to-peer system*. 2001.
6. Lampros Raptis et al: *Integrated Management Of IP Over Optical Transport Networks*. IEEE International Conference on Telecommunications (ICT-01), pp172-177, Bucharest (Romania), June 2001.
7. N. Larcin: *ASON and GMPLS – The battle for the Optical Control Plane*. August 2002.
8. Li Li, Marina Thottan, Bin Yao and Sanjoy Paul: *Distributed Network Monitoring with Bounded Link Utilization in IP Networks*. Bell Labs. IEEE Inforcom 2003.
9. J.A. Lozano y C. de Hita: *Nueva visión en la gestión de redes y servicios*. Comunicaciones de Telefónica I+D, número 18, junio de 2000.
10. Jie Lu and Jamie Callan: *Content-Based Retrieval in Hybrid Peer-to-Peer Networks*. In Proceedings of the Twelfth International Conference on Information and Knowledge Management (CIKM'03). New Orleans, ACM.
11. J.L. Martín y J. A. Paz: *Nuevas tendencias y herramientas OSS para redes IP*. Comunicaciones de Telefónica I+D, número 36, junio de 2005.
12. RFC3626: *Optimized Link State Routing Protocol (OLSR)*. October 2003.
13. OSS through Java Initiative: [www.oosj.org](http://www.oosj.org).
14. George Pavlou: *Using Distributed Object Technologies in Telecommunications Network Management*. IEEE Journal On Selected Areas In Communications, Vol. 18, No. 5, pp. 644-653, May 2000.
15. The International PGP Homepage: [www.pggi.org](http://www.pggi.org) (last visited 2/12/2005).
16. Aiko Pras, Bert-Jan van Beijnum and Ron Sprenkels: *Introduction to TMN*. CTIT Technical Report 99-09. April 99.
17. Willem Romijn et al: *Enhancing Network Management by Applying Policy Management Principles*. Bell Labs Technical Journal, Vol.8, Nº 1, pp.151-156, 2003.
18. D. Schoder and T. Eymann: *The Real Challenges of Mobile Agents*. 2000.
19. Rudiger Schollmeier: *A definition of Peer-to-Peer Networking for the Classification of Peer-to-Peer Architectures and Applications*. 1st International Conference on Peer-to-Peer Computing, August 2001.
20. Clay Shirky: *What is P2P ... and what isn't?* [www.openp2p.com/pub/a/p2p/2000/11/24/shirky1-whatisp2p.html](http://www.openp2p.com/pub/a/p2p/2000/11/24/shirky1-whatisp2p.html) (last visited 5/12/2005).

## REFERENCES

21. TeleManagement Forum (TMF): [www.tmforum.org](http://www.tmforum.org).
22. *Telecom Operations Map (TOM)*. TeleManagement Forum, GB910, issue 2.0. November 1999.
23. OASIS UDDI Specification:  
[www.oasis-open.org/committees/uddi-spec/doc/tcspecs.htm](http://www.oasis-open.org/committees/uddi-spec/doc/tcspecs.htm).
24. W3C Web Services Architecture: [www.w3.org/TR/ws-arch/](http://www.w3.org/TR/ws-arch/). 11 February 2004.
25. Chris Wellens and Karl Auerbach: *Towards Useful Management*. The Simple Times, Volume 4, Number 3, July 1996.
26. OASIS Web Services Distributed Management:  
[www.oasis-open.org/committees/tc\\_home.php?wg\\_abbrev=wsdm](http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=wsdm).
27. OASIS Web Services Notification:  
[www.oasis-open.org/committees/tc\\_home.php?wg\\_abbrev=wsn](http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=wsn).
28. OASIS Web Services Resource Framework:  
[www.oasis-open.org/committees/tc\\_home.php?wg\\_abbrev=wsrf](http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=wsrf).
29. Martin Zach, Daryl Parker, Liam Fallon, Christian Unfried, Miguel Ponce de Leon, Sven van der Meer, Nektarios Georgalas and Johan Nielsen. *CELTIC Initiative Project Madeira: A P2P Approach to Network Management*. Eurescom Summit 2005.