A Network Management Framework for Multi-Layered Network Survivability: An Overview *

D. Medhi, S. Jain[†], D. Shenoy Ramam[§], S. R. Thirumalasetty[§], M. Saddi[‡], F. Summa[¶]
Computer Science Telecommunication ⟨www.cstp.umkc.edu⟩
University of Missouri–Kansas City, USA

Abstract

We consider an interconnected network environment where one network can act as a 'provider' or 'service' network to the 'user' network or a collection of 'user' networks. In such an environment, a major failure in a provider network can affect the user networks. Very often, failure management is addressed in each individual network domain independently. In this work, we present an overview of a loosely-coupled network management framework through the development of a multi-layered network manager of managers concept for correlated management, for failure cases that can *not* be addressed by each component network individually.

I Introduction and Motivation

In this work, we consider a multi-layered view of providing different *network* services. For simplicity, consider providing the network service such as the Internet or the public switched telephone network (PSTN). While such networks have their own routers/switches connected by links, they may in turn use the transport network service (e.g. SONET) from another network. That is, the links for the "upper" network where voice and Internet are provided are actually mapped to transmission links and routes on a "lower" network.

In such a multi-layered network environment, each layer network can have independent policies regarding routing and resource management. Thus, some important issues arise; how a failure at the lower resource layer will affect the overall network, what is the best way for overall networks to respond to a failure, and what functionalities are needed for the failure management.

The administration and management of such a resource-directed multi-layered communication network involve some additional complexities in the routing and resource management strategies during link/node failure situations. Although traditional management systems, which are designed to manage the network of single administrative domain or of homogeneous technology, do deal with issues related to resource management and survivability, the scope of the management information available to the domain-specific management systems is localized. For example, the management system of a virtual private network may request an additional capacity on an overflowing point-to-point link. Similarly, a link failure at the physical layer may affect some of the

* Supported by DARPA and Air Force Research Lab, agreement No. F30602-97-1-0257.
† Now with Transportation.com. § Now with Ciena Corporation.
‡ Now with Calix Networks. ¶ Now with Informix Corporation.

virtual links of service provider networks at the layers above the physical network. Currently most of these inter-domain issues are handled by human managers.

In this work, we propose a loosely-coupled hierarchical network management framework that can facilitate maximal survivability of services for various failure situations in multi-layered networks with different administrative domains. We assume that user networks and the provider network are *not* all completely controlled by the same network administration. In each of the user network or the provider network, we assume that there is at least a domain-specific manager.

Although a reader might be tempted to equate the user network concept to customer network management [7], they are not the same. For one, while customer network management often refers to end-users at the customer premise equipment (CPE) end, ours is primarily on providing *network* services. Further, CPE-end-based customer network management is often provided flexibility in terms of billing capabilities but is given limited network capabilities; in this sense, current customer network management can be classified as a passive mechanism as far as dynamics of network services is concerned.

Throughout the last decade, the design of distributed, both hierarchical and non-hierarchical, management systems as well as customer network management has attracted the attention of many researchers, e.g. [1, 2, 3, 8]. While the need for inter-relating between different networks have received some attention [4, 5], there is little work in the dynamic multi-layered network management architecture in a multi-ownership environment. Our attempt here is to address the management framework in this area, specifically in the context of survivability.

II Architectural Framework

In a more detailed report [6], we have illustrated examples that show the insufficiency of the domain-specific network managers in a multi-layered network to address a certain level of coordination needed for maximal possible benefits. Now imagine a network environment where there are multiple user networks (each with it own administrative domain) that may require transport network service from a provider network (or networks). In such an environment, the level of survivability needed for each user network could be different as agreed through service level agreements.

Our approach here is to address coordination between user networks and provider networks through an integrated management framework for multi-layered, multi-domain networks. However, with the increase in size of the network and in number of domains, the amount of management information increases exponentially, hence implementing this integrated management system as a distributed/hierarchical system rather than a single monolithic system is preferable. The choice of a loosely coupled hierarchical distributed architecture is preferred over a flat distributed architecture in order to provide flexibility to each administrative domain. It should be noted that in a fully-distributed architecture, each domain specific manager should understand the abstraction of the management information of all other domains, whereas this is not required in a loose-coupled management architecture.

We propose a loosely coupled hierarchical management system (see Figure 1) where there is an intermediary that interacts with both the upper layer and the lower layer domain-specific managers. We name such an intermediary component as the *Across Layer Manager of Managers (ALMoM)*. The loosely defined framework allows the possibility that ALMoM may not have control over the internal workings (such as routing) within a specific user network. The capabilities between the user network and the provider network can be negotiated at the time of network manager registration through service level agreements. With ALMoM as the coordinating point, each of the domain-specific managers exchange management information, eliminating the need for the domain specific managers to understand the abstraction of management information of all other domains. MIMIC Agents are the middleware between the agents on the Managed Nodes and the

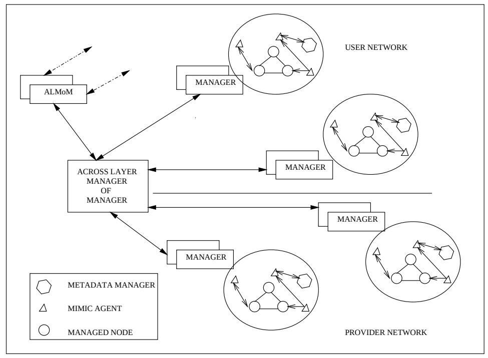

Figure 1: Loosely-Coupled Management System for Multi-Layered Networks

MIMIC Manager of the network domain. The MIMIC Metadata Manager facilitates the offering and the discovery of services of particular types.

The overall architecture of the hierarchical management system for the multi-layered network consists of three levels of management components. At the lowest level, technology-specific agents act as proxies between the network nodes and the domain-specific managers at the middle level. The domain-specific managers perform the core of the management functionalities within its administrative domain. The Across Layer Manager of Managers (ALMoM) in the upper-most level is responsible for the inter-domain management functionalities such as survivability and resource management. The exchange of management information between ALMoM and the domain-specific manager is also limited by these functionalities. The role of ALMoM, in this context, can be compared to that of resource trader or broker among the various domains. The interface of communication with management components of the lower level can be of more than one type depending on lower level components. This interface is referred to as the *vertical interface*. Similarly, the interface used for communicating with the management component of the upper level is referred to as the *horizontal interface*. Each management component will interact with more than one management components of the immediately lower level and with only one component in the upper level. The management components in this architecture can also be seen as service providers to the upper level components and as users of service provided by the lower level components. The name *horizontal interface* is used to mean that whichever component in the upper level wants to access the services provided by this component, say 'M1', it needs to use the horizontal interface of 'M1'. Similarly, when a management component 'M1' accesses the service offered by a component 'M2' in the lower level, the component 'M1' uses the vertical interface specific to the component 'M2'. In other words, the horizontal interface of component 'M2' will be the vertical interface of another component 'M1' with respect to the component 'M2'.

Specific to the survivability mechanisms, the definition of the interface also consists of a set of messages exchanged by ALMoM and the domain-specific managers. For example, when a physical link fails in the provider network, the manager of the provider network sends a message to ALMoM. Then ALMoM would determine the set of logical links affected by this failure in various user networks at the upper layers of the network. Depending on the specific implementation of ALMoM, ALMoM may request the current status of the provider network from its manager, while determining the set of affected components. The ALMoM would inform the managers of administrative domains of user networks whose logical links are affected. This will then initiate a string of message exchanges between ALMoM and domain managers for the restoration of affected areas. Similarly, when one of the administrative domains wants the allocated bandwidth of one of its logical links to be increased, the domain manager will send a message to ALMoM to that effect. ALMoM would then determine whether the request can be served. Depending on this determination, ALMoM will initiate a string of message exchanges with domain-specific managers.

Another important aspect of this architecture is that the level of relationship and interaction with ALMoM can be different from one domain-specific manager to another. We envision that each user network may require different levels of survivability requirements. For example, user network-A (UN-A) may require full restoration while user network-B (UN-B) may require partial restoration. What is desirable should be negotiated at the time of registration of the user network domain with ALMoM through a service level agreement (SLA). In such an environment, to address for a failure, ALMoM may be required to correlate notification of failures from different user networks to prioritize what restoration algorithms to activate in the provider network, such that the user network with the higher priority is restored first, when the bandwidth is limited in the provider network. We have also developed algorithms for this type of environment [9]; however, this is outside the scope of the present paper.

Finally, as shown in Figure 1, there can be multiple instances of ALMoM for scalability purposes which do distributed networking among themselves, while outside entities such as the user network manager and the provider network manager are transparent to the internal view of ALMoM. In other words, physically there's no reason to have one ALMoM component. A good analogy to this concept is the domain name system (DNS) used on the Internet. While DNS is used by every host in the Internet, its internal working (between various DNS servers) are transparent to the host.

More details about the framework and a proof-of-concept implementation can be found in [6].

References

[1] A. Anerousis and A. Lazar, "An Architecutre for Managing Virtual Circuit and Virtual Path Services in ATM Networks," *Jnl. Network & Systems Mgnt.*, Vol. 4, pp. 425-454, 1996.

[2] L. Bjerring, D. Lewis and I.H. Thararensen, "Inter-domain Service Management of Broadband Virtual Networks," *Jnl. Network & Systems Mgnt.*, Vol. 4, pp. 355-373, 1996.

[3] R. Diaz-Caldera *et. al.*, "An Approach to the Cooperative Management of Multitechnology Networks," *IEEE Comm. Mag.*, Vol. 37, No. 5, pp. 119-125, May 1999.

[4] P. Demeester *et. al.*, "Resilience in Multilayer Networks," *IEEE Comm. Mag.*, Vol. 37, No. 8, pp. 70-76, August 1999.

[5] D. Medhi, "A Unified Approach to Network Survivability for Teletraffic Networks: Models, Algorithms and Analysis," *IEEE Trans. Comm.*, Vol. 42, pp. 534-548, 1994. Also see: D. Medhi and R. Khurana, "Optimization and Performance of Network Restoration Schemes for Wide-Area Teletraffic Networks," *Jnl. Ntwk. Sys. Mgnt.*, Vol. 3, pp. 265-294, 1995.

[6] D. Medhi, S. Jain, D. S. Ramam, S. R. Thirumalasetty, M. Saddi, F. Summa, "Network Management for Multi-Layered Network Survivability: Management Framework and Implementation," CST Technical Report, University of Missouri-Kansas City, 2000.

[7] B. Basset, F. Heinrich, D. Kuhl, and A. Shadman, "Customer Network Management: A Service Provider's View," *IEEE Comm. Mag.*, Vol. 28, No. 3, pp. 31-34, March 1990.

[8] M.-A. Mountzia and G. D. Rodosek, "Using the Concept of Intelligent Agents in Fault Management of Distributed Services," *Jnl. Network & Systems Mgnt.*, Vol. 7, No. 4, 1999.

[9] S. R. Thirumalasetty and D. Medhi, "MPLS Traffic Engineering for Surivable Book-Ahead Guaranteed Services," CST Technical Report, University of Missouri-Kansas City, 2000.