
Telecommunications Management Network: Vision vs. Reality

Although very few organizations can claim today that TMN has improved their ability to manage their telecommunications network, there is no doubt that the TMN vision — with its inherent promise — will gradually become a reality.

Roch H. Glitho and Stephen Hayes

Since the early 1980s, the standards bodies have been specifying the Telecommunications Management Network (TMN) principles. Millions of dollars have been spent. The TMN principles aim at being applicable across telecommunications technologies. They recommend the use of independent management networks to manage telecommunications networks, elements in the telecommunications networks (managed networks), and managing systems (in managing networks), communicating via well defined, standardized interfaces.

The standards bodies envisioned TMN as a possible solution to the complex problem of telecommunications networks and services Operation, Administration, Maintenance & Provisioning (OAM&P) in today's open, multivendor environment. However, the vision stumbles against the reality. Various factors still hinder the implementation of TMN-based OAM&P systems. This article provides a tutorial on TMN by contrasting the vision and the reality.

TMN: The Vision

In order to understand how TMN was envisioned, it is necessary to first grasp the issues that led to the development of TMN. Another prerequisite is an understanding of Open Systems Interconnection (OSI) systems management, one of the key technologies upon which TMN is based.

Motivations

Management of telecommunications networks used to be simpler. In the days before the deregulation and privatization of the telephone industry, there were fewer issues to deal with, and generally less competitive pressure. Inefficient and work-intensive operations and management practices were more acceptable. In general, the network was composed of equipment from fewer vendors, thus there were fewer multivendor management issues. Also, the introduction and integration of new technologies and services proceeded at a slower pace.

It was apparent even then to service providers

and telecommunications equipment suppliers that this situation could not last. The wave of the future required increased automation of operations and maintenance tasks, the management of multivendor networks, and the rapid integration of new technologies.

The need for automation required that machine-to-machine interfaces be developed to replace many of the manual functions. The need for managing heterogeneous equipment required that some form of standardization be implemented. Finally, the need to support rapid technological evolution required that the interfaces that were developed be both general and flexible. Furthermore, to ensure that the interfaces had sufficient consistency to allow some level of integrated management, it was necessary to develop a set of guiding principles. That set of guiding principles is the TMN vision.

The overall vision was of a network of management systems linked together and to the various telecommunications networks. This set of systems and the links between them comprised TMN. It constantly monitored and tuned telecommunications networks and, in general, removed the need for human intervention, except for exceptional circumstances or activities that required physical intervention (such as replacing circuit boards). The interfaces were standardized so that introducing equipment from new vendors occurred smoothly (at least as far as OAM&P is concerned). New technologies can be introduced with a minimum of adaptations so that operational procedures may be changed via evolution and not revolution.

Few people (except perhaps those craftspersons put out of work) would fault the TMN overall vision. Its promise has motivated the expenditures. TMN has made considerable progress over the past several years, but few organizations would claim that it has kept the promise implicit in the vision. This does not imply that TMN is a failure or that it is all hype; there is considerable substance behind TMN. However, for various reasons, much of its promise has not yet been realized. In order to understand why, it is first necessary to understand the principles and the technology chosen as the pillar for the interface development.

ROCH H. GLITHO is a principal engineer at Ericsson, while pursuing a Ph.D. at the Royal Institute of Technology of Stockholm in Sweden.

STEPHEN HAYES is a principal engineer at Ericsson, while pursuing a Ph.D. at the University of Texas at Dallas.

There is considerable substance behind TMN. However for various reasons, much of its promise has not yet been realized.

Principles

TMN stands for "Telecommunications Management Network." The term "network" is a key concept. It was envisioned that management would be performed by a cooperating set of systems rather than a single monolithic manager. Even if an administration were to buy a huge monolithic supercomputer and entrust it to control its network, this would be inadequate. It would still need to communicate with the management systems in other administrations in order to resolve faults that spanned jurisdictions (e.g., faults on trunks between the jurisdictions). Thus, the management is envisioned as being done by a network of systems and not a single system.

One of the first needs in defining the TMN principles was to specify the architecture of TMN. This includes the identification of the different types of nodes and the interfaces between them. One of the pivotal documents in TMN is Recommendation M.3010 [1], which deals with the TMN principles, including the architecture. We could try to specify the architecture only in terms of physical nodes and communications interfaces. The various physical components of the TMN are specified in column 1 of Table 1.

In doing so, we immediately run into a problem. How do we classify a node where the vendor has provided both network element (NE) functionality and operation system (OS) functionality? How do we refer to the information being transferred between the functional components since there is no longer a physical interface? To overcome these issues, TMN uses the concept of functional architecture, which is defined using function blocks and reference points.

Function blocks are logical entities that can be implemented in a variety of physical configurations. Table 1 shows the function block, which is mandatory, for each TMN node. Reference points represent the exchange of information between two function blocks. The correspondence between interfaces and reference points can be seen in Table 2. Interfaces are designated in upper case; reference points are designated in lower case.

In earlier years, the concepts of functional components and reference points were simply useful abstractions. However, as distributed computing evolves, these concepts may well become more tangible. So far, the specification of Application Program Interfaces (APIs) and object interfaces has remained outside of the scope of TMN.

This subsection discusses the TMN nodes and TMN interfaces, and introduces the TMN management services and TMN interface specification methodology — two other important concepts of TMN.

TMN Nodes — Figure 1 (taken from Fig. 16/M.3010 [1]) shows a simplified example of a physical TMN architecture. It is used to review the nodes and, subsequently, the interfaces. In this and in the subsequent discussion we shall refer to the nodal and interface designation. Readers need to be aware that for each of these an equivalent functional component or reference point exists.

The OS represents the supervisory or control systems in TMN. Although Fig. 1 does not explicitly show it, OSs can be interconnected. Thus, OSs can form management hierarchies or other structures.

TMN node	Mandatory function blocks
Operation system (OS)	OSF
Mediation device (MD)	MF
Q-adaptor (QA)	QAF
Work station (WS)	WSF
Network element (NE)	NEF

■ Table 1. TMN nodes and their mandatory function blocks.

TMN interfaces	TMN reference points
Q3	q3
Qx	qx
F	f
X	x

■ Table 2. TMN interfaces and reference points.

OS functionality can also be layered using the Logical Layered Architecture (LLA) concept.

There is debate on the actual number of layers. However, a proposal that is not formally part of TMN has gained considerable popularity. It clusters OS functionality into the following layers: element management layer, network management layer, service management layer, and business management layer.

Mediation devices (MD) are probably the most vague component of TMN. They may provide storage, adaptation, filtering, thresholding, or condensing operation on data received from subtending equipment. Since the concept of MD is a nebulous one, it is questionable whether any MDs have been developed to date. Consequently, it is important to point out that what is often referred to as MD in the industry is actually a Q-Adaptor (QA).

The QA is a concession to reality. Its mission is to connect a TMN system to a non-TMN system. Q-adaptors are the great hope for integrating existing networks into TMN. In reality they have been difficult to develop due to problems in mapping between the TMN interfaces and the pre-existing interfaces.

The NE is the only node actually residing in the managed network, the telecommunications network. Its primary job is to handle traffic and not management. It is, however, the ultimate origin or destination of the management supervision and control.

The work station (WS) is where the human sits. It provides the presentation function to the user. It should be noted that WS as a TMN node does not convey the same notion as the workstation of the computer world.

The nodes identified above communicate through the Data Communication Network (DCN), which is the transportation means used in the TMN world. Initially, DCN was assumed to be independent from the telecommunications network, but this restriction has been relaxed due to the costs associated with the maintenance of a distinct physical network.

TMN Interfaces — Although the architecture discusses nodes, functional components, interfaces, and reference points, the bulk of the standards deal with interfaces. The manner in which management systems interact is governed by the interfaces. The functional components and reference points are abstractions and thus are not subject to standardization.

There is a reluctance to standardize the functionality of nodes, as this would constrain product offerings. Standardizing interfaces is sufficient to allow the nodes to interwork as long as the protocols are specified to a level that allows applications to interact.

The Q3 interface is the flagship interface of TMN. It is the one for which the specifications are fairly complete. It connects an OS and an NE, or an OS and a QA, or an OS and a MD, or two OSs that belong to the same TMN.

The Qx is the Q3 interface's underdeveloped brother. It is like a Q3 but with less functionality. It was intended to be used when cost or efficiency issues precluded a fully functional Q3 interface. The problem is that there has been no agreement on what can be dropped from the Q3, and there is no resolution in sight.

The X interface is used for communicating between OSs belonging to different TMNs, or between a TMN OS and a non-TMN OS that supports a TMN-like interface. There is considerable interest in X interfaces. However, very few have been specified up to now, due to their complexity.

The F interface is used for communicating between the WS and the other nodes. Little effort has been made so far on the F interface.

TMN Management Services and TMN Interface Specification Methodology — A management service can be defined as an offering that fulfills a TMN user's specific telecommunications management need. Many management services have been identified; examples are customer administration and traffic management. For an exhaustive list of the TMN management services identified so far, see Recommendation M.3200 [2].

TMN offers a methodology for the specification of the various interfaces identified earlier in this article. The methodology is described in Recommendation M.3020 [3]. However, it has never been applied successfully in its entirety to any known real sub-network, network, or service. It promotes a top-down approach, while in most real cases a bottom-up approach or a mixed approach is used.

OSI Systems Management


As previously stated, one of the criteria imposed on TMN is the ability to accommodate the management of diverse technologies. This requires that the TMN interfaces must be both general and flexible. In addition, the requirement for consistency has motivated the use of standardized protocol suites. A very powerful technology was needed to meet the requirements.

After some debate, the OSI systems management technology was selected as the basis for the TMN interfaces. Although not intrinsically part of TMN, the concepts of OSI management have become so intimately associated with TMN that it is impossible to understand TMN without a basic understanding of OSI management.

Although OSI systems management and TMN have evolved together, they are quite different. OSI systems management is a set of standards developed jointly by the International Standards Organization (ISO) and the International Telecommunications Union (ITU). These standards were oriented primarily toward managing data networks. In reality, their evolution has been considerably influenced by the TMN requirements.

The following subsection describes OSI systems management from the perspective of TMN. It reviews the OSI systems management concepts, presents the organization of the TMN interface standards, and discusses the benefits of using the OSI systems management as the basis for the specification of the TMN interfaces.

OSI Systems Management Overview — Figure 2 illustrates the key concepts of OSI systems management. It depicts a local area network (LAN) card that is managed using OSI systems management. The card resources include the communication chip that implements the LAN protocol. The protocol implemented by the chip is supposed to be Ethernet. We review below the OSI systems management concepts using the elements depicted by Fig. 2 as concrete examples. By necessity, this overview is cursory. Readers are referred to [4] for a more comprehensive treatment of OSI systems management.



■ Figure 1. A simplified TMN physical architecture.

The concepts of OSI management have become so intimately associated with TMN that it is impossible to understand TMN without a basic understanding of OSI management.

A managed object (MO) is the conceptual view of a resource (physical or logical) that needs to be monitored and controlled in order to avoid failures and performance degradation in a network. Ether_Chip in Fig. 2 is an example of MO. It is the abstract view of the Ethernet chip that is on the LAN card.

MOs with the same properties are instances of an MO class. Although not shown in Fig. 2, an example of MO class is LAN_Chip. It groups all instances of chips that implement a LAN protocol, including Ether_Chip in Fig. 2.

The Management Information Base (MIB) is the conceptual repository of the MOs instances. The MIB in Fig. 2 contains among other MO instances the Ether_Chip.

A MO class is defined by the attributes, the management operations, the behavior, and the notifications.

- The attributes are data elements and values that characterize the MO class. The attributes of LAN_Chip include the protocol the chip implements (Ethernet, Token Ring, other), the serial number, and the manufacturer identifier.
- The management operations are operations that can be applied to MO instances. Examples of operations that can be applied to Ether_Chip are the various tests.
- The behavior exhibited by a MO instance is based on the resource the MO class represents. The potential outcome of the various tests are part of the Ether_Chip behavior.
- The notifications are messages that MO instances emit spontaneously. The notifications the communication chip emits include "packet received." This is emitted whenever a packet is received by the node.

There are two roles defined in OSI management, the manager and agent roles.

- The manager is the specific entity in the managing system that exerts the control, the coordination, and the monitoring. It issues the requests to perform operations against the agent. It also receives the notifications emitted by the MOs and sent by the agent.
- The agent is the specific entity in the managed system to which the control, the coordination and the monitoring are directed. It receives and executes the requests sent by the manager, and sends the notifications to the manager.

Manager and agent may communicate using a 7-layer OSI protocol suite. A key element of the suite is the Common Management Information Service Element (CMISE), which is one of the building blocks used at the application layer. CMISE consists of a service definition, the Common Management Information Service (CMIS); and a protocol specification, the Common Management Information Protocol (CMIP). For an overview of CMIS/CMIP, refer to [5].

Thanks to the use of CMISE, all messages exchanged between the manager and the agent have a basic form of either requesting something of one or more object or an object informing another system of some event. The requests may be as simple as returning the value of a parameter, or as complicated as asking the NE to reconfigure itself.

The agent receiving the message is responsible for carrying out the request(s). It maps the request(s) on the MO(s) into request(s) on real

resources. However, the mechanisms used for the mapping are implementation-specific and not subject to standardization.

Using the above concepts, the resources are modeled so that the manager and the agent have a common view. This specification of object-oriented information is called information modeling. The majority of the effort spent in defining TMN interfaces has gone into the development of these information models.

Organization of the TMN Interface Standards — The OSI systems management standards provide power and flexibility in defining interface standards, but they are not in themselves the TMN interface standards. The TMN interface standards comprise generic standards and technology-dependent standards.

The generic standards are intended to be applicable across all telecommunications technologies and services. A classic example is Recommendation M.3100 [6], which contains MOs that are generic enough to describe information exchanged across all TMN interfaces, independently of the telecommunication technology. The objects specified in the technology-specific standards are often imported from the generic standards or are subclasses of generic objects.


Inheritance (also called subclassing) is the procedure of specifying a new object class based upon a previously defined object class. Thus, the new object class has all the characteristics of the base object class (superclass) with some new characteristics. This policy of deriving technology-specific object classes from base generic object classes ensures a level of similarity between different technology-specific information models.

Allomorphy is a capability that may be used to manage the telecommunications technologies in a generic manner. It is the procedure of specifying a subclass that masquerades as a superclass. One use of this is to allow a technology-specific object to be treated as a more generic object. Thus, a technology-specific object can be managed as a generic object.

The disadvantage of the above approach is that technology-specific management capabilities are inaccessible. A related use of allomorphy is to provide a generic set of management capabilities in certain situations, while providing vendor-specific enhancements in other situations. In reality, there has been insufficient use of TMN standards to determine if allomorphy is truly a useful concept.

Inheritance and allomorphy, along with the concept of generic and technology-specific standards, are the mechanisms for providing the generality and consistency desirable in TMN interfaces. An overview of TMN standardization activities is found in [7]. Other articles in this issue summarize the status of specific technology-specific standards.

Benefits — The protocol suite to be used at the Q3 interface is the OSI protocol suite. In adopting the OSI system management protocol suite, TMN has gained, among other things, reliable and robust communications capabilities, and a wealth of application-layer building blocks. The application layer building blocks include the Association Control Service Element (ACSE) and the CMISE.



■ **Figure 2.** A LAN card managed using OSI systems management.

The ACSE provides a means for establishing associations and negotiating application protocol capabilities.

In addition to the gains linked to the OSI protocol suite, various other benefits are worth mentioning:

- A semi-formal specification technique (templates) for defining the information model, including object classes, attributes, actions, notifications, etc.
- The use of object-oriented techniques such as inheritance and allomorphism.
- Naming rules that facilitate the structuring of objects in a database.
- A large set of objects already defined for such things as routing alarms (event forwarding discriminators), logging data (logs), report generation (scanners), etc.
- A data-specification language (ASN.1) for defining data structures in an abstract (machine-independent) notation.
- A method for encoding and decoding application-layer data (BER and presentation layer), independent of machine-specific representation.

TMN: The Reality

The TMN vision does not give any insight into how far TMN has gone in the real world. Although many administrations and companies have voiced their support of TMN, it has only rarely been deployed in the field. In this section we discuss the reasons for the current state of affairs and make some predictions about the fate of TMN.

The Complexity of TMN — Ironically, the principles specially developed for turning the TMN vision into reality have been among the many stumbling blocks hindering its implementation. For instance, the adoption of the OSI system management as the basis for the TMN interface specification has not been without penalty. Two of them are mentioned below.

The price to pay for the power and the flexibility of OSI systems management is that the task of specifying TMN interfaces is dauntingly complex. The pool of individuals versed in the specifications tools and capable of actually developing those specifications is quite small.

Another penalty associated with the use of the OSI systems management is that the OSI systems

management standards were not stable when the development of TMN principles was initiated. This has caused delays in the development of the TMN standards.

The requirements on the TMN information models are actually very difficult to meet. The models should be robust enough to accommodate both existing and future technologies. They should not restrict excessively architectural or implementation approaches of either existing or future products. At the same time they should support the management procedures of a diverse set of administrations. The challenge has been to provide models that meet the above criteria while still being useful in helping to solve concrete OAM&P problems.

The Persistence of Legacy Interfaces — Existing technologies, such as POTS, have been a hindrance for both economic and technical reasons. It has been difficult to justify the expenses of migrating to new interfaces. The development of TMN interfaces in general carries a high initial price tag.

This has led to a chicken and the egg situation. Although most companies voice support of the TMN standards, OS developers have been reluctant to develop interfaces for which there is no NE support, and NE developers do not want to develop interfaces for which there is no OS support.

There are several NEs already deployed that do not support a Q3 interface. Integrating these NEs into any TMN environment requires the development of QAs. Additional work is needed to map to the data in the NEs and to identify the functions to be done in the QAs.

Alternative Management Protocols — As mentioned earlier, TMN is tightly coupled with OSI management. This implies an alignment with OSI in the OSI vs. Transmission Control Protocol/Internet Protocol (TCP/IP) wars. TCP/IP has its own management protocol, the Simple Network Management Protocol (SNMP). Due to the prevalence of this protocol in data communications networks, there is pressure to use SNMP in many of the TMN applications. Due to space limitations, it is not possible to describe the SNMP protocol here. Interested readers are urged to consult [4].

Although many administrations and companies have voiced their support of TMN, it has only rarely been deployed in the field.

**The future
will see
more and
more TMN
systems as
confidence
in the TMN
vision grows.
The vision
will then
become the
reality.**

SNMP is simpler and less powerful than OSI management. It provides services similar to the CMISE services, but does not cater to MIBs with complex structures. An enhanced version of SNMP (referred to as SNMPv2 or SNMP) has been specified. On the one hand it improves the initial functionality, but on the other hand it loses part of the initial simplicity. SNMP is currently making inroads in areas such as Customer Network Management (CNM).

A Motivated Prognosis

After years of incubation, TMN is finally hatching. More and more TMN systems will be deployed in the field. This section addresses the motivations for the anticipated growth of TMN systems.

A couple of years ago, a valid reason for not implementing TMN systems was that TMN standards did not exist. The long waiting period for tangible output from the standards bodies led to a widespread perception that TMN is a dream without substance. But this is no longer true. The choices made for TMN did make the standards development process an uphill task, but did not ultimately make it impossible.

While the coverage of TMN standards is not complete, there now exist concrete and stable TMN interface standards for many areas. To complement the existing standards, industry forums such as the Network Management Forum (NMF) are also developing implementation guidelines. Those guidelines ease the implementation of standards.

The complexity of TMN is no longer such a problem. OSI toolkits (including CMISE toolkits) now exist and are available both commercially and as freeware. It is therefore possible to use these products to greatly simplify the development of TMN interfaces. Much of the intricacies of OSI management can be avoided in this manner. In addition, the pool of people familiar with the concepts of OSI management has been slowly growing.

Even the existing interfaces are beginning to be supplanted by TMN interfaces. There are plans underway to replace many of the existing switch interfaces for data collection and traffic management interfaces with an interface based upon OSI management. It is unreasonable to expect that the existing infrastructure for POTS will be transformed in the near future to a TMN-based system. However, as TMN is introduced for new technology, it will become increasingly attractive to develop QAs for managing the legacy systems.

The TMN vision is also being deployed in OS-to-OS interconnection (X interfaces). In a desire to automate the activities that occur between administrations, TMN interfaces for trouble administration are now being deployed.

The TMN vision is a reality for new technologies such as the ones discussed in this issue. TMN-compliant systems are being developed and deployed. The TMN systems are not yet common because these technologies are only beginning to be applied. It is interesting to note that one of the key features of OSI management is specification reuse. Due to this, it is generally easier to specify a TMN standard for a new technology than to develop a new management interface. This is due to the

fact that a TMN standard for a new technology can build upon generic TMN standards.

While it may make economic sense to use SNMP in certain situations, this does not spell the end of TMN. In fact, it may be desirable to expand TMN to include SNMP. In general, however, the power of OSI management will be preferable due to the complexity of the telecommunications equipment being managed.

Conclusion

The TMN principles have been developed to address several of the fundamental problems facing telecommunications networks management. TMN provides a structure for categorizing the management network according to physical or functional entities, and according to interfaces and reference points. It provides for the structuring of the various management services, and it offers an interface-specification methodology that is currently closely coupled to OSI system management.

The deployment of TMN has been slow due primarily to its complexity and the inertia of legacy systems. As the telecommunications environment changes, these roadblocks are giving way. The future will see more and more TMN systems as confidence in the TMN vision grows. The vision will then become the reality.

Acknowledgments

The authors would like to thank the following people: Carsten Gyrn, Teledanmark, Denmark; Masahiko Matsushita, NTT, Japan; Lakshmi Raman, Bellcore, United States; John Wilber, GTE, United States. Their comments have helped improve the quality of this article. We, however, accept full responsibility for all the controversial statements.

References

- [1] CCITT Recommendation M.3010 "Principles for a Telecommunications Management Network," Geneva, 1992.
- [2] CCITT Recommendation M.3200 "TMN Management Services: Overview," Geneva, 1992.
- [3] CCITT Recommendation M.3020 "TMN Interface Specification Methodology," Geneva, 1992.
- [4] Y. Yemini, "A Critical Survey Of Network Management Protocols and Standards," In S. Aidarous and T. Plevyak, eds., *Telecommunications Management into the 21st Century*, IEEE Press, 1994.
- [5] L. Raman "CMISE Functions and Services," *IEEE Commun. Mag.*, May 1993.
- [6] CCITT Recommendation M.3100 "Generic Network Information Model," Geneva, 1992.
- [7] D. Sidor "Managing Telecommunications Networks Using TMN Interface Standards," *IEEE Commun. Mag.*, in this issue.

Biographies

ROCH H. GUTHO received M.Sc. degrees in computer science (1984) and mathematics (1985) from the University of Geneva (Switzerland), and an M.Sc. in business economics (1990) from the University of Grenoble (France). He is currently pursuing a Ph.D. at the Royal Institute of Technology of Stockholm (Sweden). He is a principal engineer at Ericsson, Montreal, Canada, where he works in the area of OAM&P. He joined Ericsson in 1990 after having worked five years for a computer manufacturer in Oslo (Norway). He is active in many standardization bodies, including ITU-T. He serves as a technical editor on the *IEEE Communications Magazine* editorial board. His email address is: lmcrogl@mc.ericsson.se.

STEPHEN HAYES received a B.S. degree in computer science and mathematics from Texas Christian University in 1977. He received an M.S. in computer science from the University of Texas in 1980, and is currently pursuing a Ph.D. at the University of Texas. He is a principal engineer at Ericsson, Montreal, Canada working in the area of OAM&P. He is active in T1M1.5, where he is involved in modeling activities. Past responsibilities include leading work on standards for performance monitoring. He is currently chairing the T1M1.5 Ad-Hoc group on PCS management. His email address is: exush@exu.ericsson.se.