

Patricia Morreale et al. "Voice and Data Communications"
The CRC Handbook of Modern Telecommunications
Ed. Patricia Morreale and Kornel Terplan
Boca Raton, CRC Press LLC. 2001

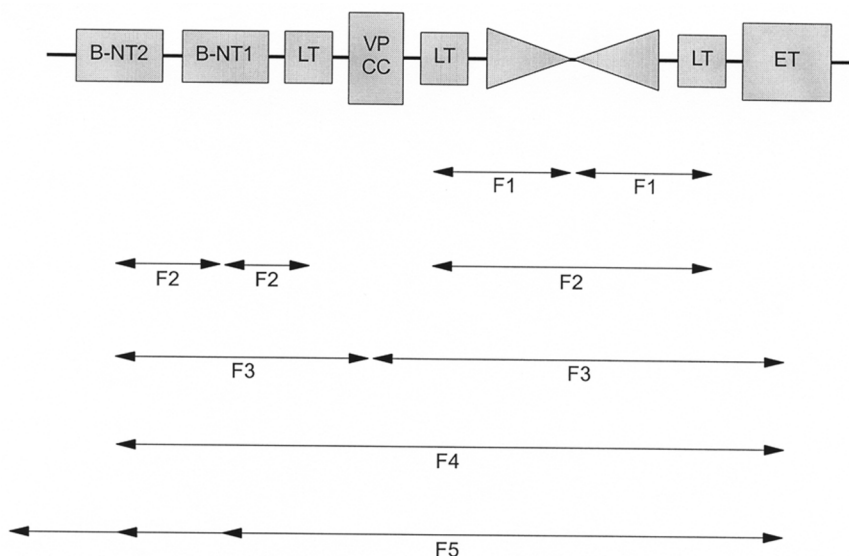


FIGURE 3.2.7 Information flows in B-ISDN.

The five OAM flows are the following:

- F5: OAM information flows between network elements performing VC functions. From the perspective of a BISDN configuration, F5 OAM operations are conducted between B-NT2/B-NT1 endpoints. F5 deals with degraded VC performance, such as late arriving cells, lost cells, cell insertion problems, etc.
- F4: OAM information flows between network elements performing VP functions. From the perspective of a BISDN configuration, F4 OAM flows between B-NT2 and ET. F4 OAM reports on an unavailable path or a virtual path (VP) that cannot be guaranteed.
- F3: OAM information flows between elements that perform the assembling and disassembling of payload, header, and control (HEC) operations, and cell delineation. From the perspective of a BISDN configuration, F3 OAM flows between B-NT2 and VP cross connect and ET.
- F2: OAM information flows between elements that terminate section endpoints. It detects and reports on loss of frame synchronization and degraded error performance. From the perspective of a BISDN configuration, F2 OAM flows between B-NT2, B-NT1, and LT, as well as from LT to LT.
- F1: OAM information flows between regenerator sections. It detects and reports on loss of frame and degraded error performance. From the perspective of a BISDN, F1 OAM flows between LT and regenerators.

3.2.4.6 Mobile and Wireless Communication

The purpose of a mobile communications system is the provision for telecommunications services between mobile stations and fixed land locations, or between two mobile stations. There are two forms of mobile communications: cellular and cordless.

The best approach is to examine their major attributes and compare them to each other. First, a cellular system usually has a completely defined network which includes protocols for setting up and clearing calls and tracking mobile units through a wide geographical area. So, in a sense, it defines a UNI and a NNI. With cordless personal communications, the focus is on access methods in comparison to a closely located transceiver — usually within a building. That is to say, it defines a rather geographically limited UNI.




FIGURE 3.2.8 Mobile management information base.

Cellular systems operate with higher power than do the cordless personal communications systems. Therefore, cellular systems can communicate within large cells with a radius in the km range. In contrast, cordless cellular communication cells are quite small, usually in the order of 100 meters.

Cellular will continue to be a preferred medium for the consumer and business market. PCS will emerge as a driver for wireless. Telecommunications providers will use PCS to help reduce cellular churn, joining the customers closer to the vendor by providing end-to-end service. Wireless data connectivity, driven by lighter and smaller equipment capable of being carried by humans as well as in vehicles, will drive wireless connectivity requirements deeper and further into the network infrastructure.

The technology is not new, but implementation took some time. Cordless systems are undergoing rapid technological changes. Different protocols and standards are used.

MIB Availability

The premise of the Mobile MIB is the following:

- Network managers require more control over remote network workstations.
- To manage these remote machines, the network manager requires data on remote network access.
- Some of this data can be supplied via SNMP from the remote workstation.
- Other important information is available only on the local network at the point where the connection was created.
- This requires, therefore, two important SNMP agents with associated Mobile MIB components, one at the remote workstation and the other at the local network connection point.

The following groups of information have been created (Figure 3.2.8):

- The Mobile Client Group contains information to be relayed from the remote workstation to a network management system on the attached network. This group stores information on mobile client name, description, location, phone number, power management configuration, connection

hardware and software, client type (CPU, RAM, disk, video), system software, date, time, network adapter used, and configuration and statistics.

- The Mobile Server Group contains information to be relayed from the local network server to a network management system on the same network. This group stores information about the network server where remote connections can be originated, such as server's name, remote network hardware, slot and port number, server uptime, connection speed, service type, and traffic statistics.
- The Trap Group describes SNMP trap types for remote workstations, such as mobile computer docked, undocked, suspended, resumed, pcmcia inserted, and a trap table. This is a table of alerts, which can be sent to the specified ipaddress using the specified protocol by setting the value of the mobileTrapTrigger object to the index of an entry in this table.
- Adapter Group contains information about network adapters, including hardware information and type of connection.
- Link Group provides data about mobile network links, such as link status and link performance.

Management is more difficult than with wirelines due to the fact that more information is processed in real-time or near real-time. In most cases, proprietary solutions dominate. The implementation of TMN and other standards is extremely slow.

3.2.4.7 Digital Subscriber Line Technologies

The enabling technology is digital subscriber line (xDSL), a scheme that allows mixing data, voice, and video over phone lines. There are, however, different types of DSL to choose from, each suited for different applications. All DSL technologies run on existing copper phone lines and use special and sophisticated modulation to increase transmission rates (ABER97).

Asymmetric digital subscriber line (ADSL) is the most publicized of the DSL schemes and is commonly used as an ideal transport for linking branch offices and telecommuters in need of high-speed intranet and Internet access. The word asymmetric refers to the fact that it allows more bandwidth downstream (to the consumer) than upstream (from the consumer). Downstream, ADSL supports speeds of 1.5 to 8 Mbit/s, depending on the line quality, distance, and wire gauge. Upstream rates range between 16 and 640 Kbit/s, again depending on line quality, distance, and wire gauge. For up to 18,000 feet, ADSL can move data at T1 using standard 24-gauge wire. At distances of 12,000 feet or less, the maximum speed is 8 Mbit/s.

ADSL delivers a couple of other principal benefits. First, ADSL equipment being installed at carriers' central offices offloads overburdened voice switches by moving data traffic off the public switched telephone network and onto data networks — a critical problem resulting from Internet use. Second, the power for ADSL is sent by the carrier over the copper wire with the result that the line works even when their local power fails. This is an advantage over ISDN, which requires a local power supply and thus a separate phone line for comparable service guarantees. The third benefit is again over ISDN; ADSL furnishes three information channels — two for data and one for voice. Thus, data performance is not impacted by voice calls. Rollout plans are very aggressive with this service. Its widespread availability is expected for the end of this decade.

Rate-adaptive digital subscriber line (RADSL) has the same transmission limits as ADSL. But as its name suggests, it adjusts transmission speed according to the length and quality of the local line. Connection speed is established when the line synchs up or is set by a signal from the central office. RADSL devices poll the line before transmitting; standards bodies are deciding if products will constrain the line speed. RADSL applications are the same as with ADSL and include Internet, intranets, video-on-demand, database access, remote LAN access, and lifeline phone services.

High-bit-rate digital subscriber line (HDSL) technology is symmetric, meaning that it furnishes the same amount of bandwidth both upstream and downstream. The most mature of the xDSL approaches, HDSL has already been implemented in the telco feeder plant — the lines that extend from central office to remote nodes — and also in campus environments. Because of its speed — T1 over two twisted pairs of wiring, and E1 over three — telcos commonly deploy HDSL as an alternative to T1/E1 with repeaters.

At 15,000 feet, HDSL operating distance is shorter than ADSLs, but carriers can install signal repeaters to extend its useful range (typically by 3000 to 4000 feet). Its reliance on two or three wire-pairs makes it ideal for connecting PBXs, interexchange carrier POPs (point of presence), Internet servers, and campus networks. In addition, carriers are starting to offer HDSL to carry digital traffic in the local loop, between two telco central offices and customer premises. HDSL's symmetry makes this an attractive option for high-bandwidth services like multimedia, but availability is still very limited.

Single-line digital subscriber line (SDSL) is essentially the same as HDSL with two exceptions: It uses a single wire pair and has a maximum operating range of 10,000 feet. Since it is symmetric and needs only one twisted pair, SDSL is suitable for applications like video teleconferencing or collaborative computing with identical downstream and upstream speeds. Standards for SDSL are still under development.

Very high-bit rate digital subscriber line (VDSL) is the fastest DSL technology. It delivers downstream rates of 13 to 52 Mbit/s and upstream rates of 1.5 to 2.3 Mbit/s over a single wire pair. But the maximum operating distance is only 1000 to 4000 feet. In addition to supporting the same applications as ADSL, VDSL, with its additional bandwidth, could potentially enable carriers to deliver high-definition television (HDTV). VDSL is still in the definition stage.

A number of critical issues must be resolved before DSL technologies achieve widespread commercial deployment. Standards are still under development. Modulation techniques, such as carrierless amplitude phase (CAP) and discrete multitone (DMT) have been separated by standards bodies. Some other problems include interoperability, security, eliminating interference with ham radio signals, and lowering power systems requirements from the present 8 to 12 watts down to 2 to 3 watts. A nontechnical but important factor will be how well carriers can translate the successes they have realized in their xDSL technology trials into market trials and then to commercial deployments.

MIB Availability

MIB definitions are in the works with standard committees and with vendors. The present definition is not yet complete, but as a guideline is very useful for further steps.

The MIB is presently structured into eight groups:

- `xdslDevIfStats` group provides statistics specific to the xDSL link. Statistics are collected on a per port basis and on specific intervals. Hence, such a table is indexed by the `xdslDevIfStatsIfIndex` and `xdslDevIfStatsInterval`. Also, statistics are grouped into remote and central statistics. "Remote" means that the statistics are collected by the device at the customer premises. "Central" means that the statistics are collected by the device located at the central office. The objects which are not divided into these two groups are related to both ends of the xDSL link.
- `xdslDevIfConfig` group provides configuration information specific to a xDSL device, or system. The table is indexed by an object which corresponds to `ifIndex`. These `ifIndex` entries themselves denote and identify specific xDSL interfaces on the board or module. Also, the configuration parameters are grouped into two broad categories, up and down. Up reflects the upstream direction (from customer premises to the central office). Down reflects the downstream direction (from the central office to the customer premises).
- `xdslRemoteSys` group's description is identical with the `mib-2` system MIB.
- `xdslRemoteDTEStatus` group provides status information about the DTE port of the DSL RTUs.
- `xdslDevMvIfConfig` group provides configuration information specific to a xDSL (MVL) device or system. The table is indexed by an object which corresponds to `ifIndex`. These `ifIndex` entries, themselves, denote and identify specific `xdsl(Mvl)` interfaces on the board or module.
- `xdslDevNAPCustomerAccount` group provides customer accounting information on each DSL port. Network access providers can accurately bill their end station DSL customers by the amount of usage. The table is indexed by `ifIndex` and `xdslDevNAPCustomerAccountInterval`. The `ifIndex` identifies specific xDSL interface on the device and `xdslDevNAPCustomerAccountInterval` specifies the accounting information for the current day or for the previous day. Customer data excludes all traffic used for management purposes.

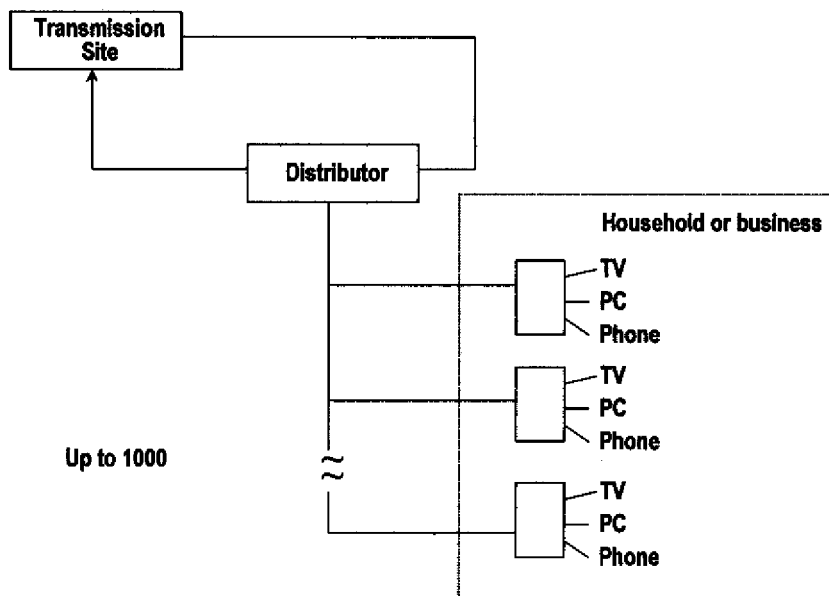


FIGURE 3.2.9 Use of cable technology in subscriber environments.

- xdsLinkUpDownInformation group contains the reason why the DSL Link went down. This information is obtained when the DSL Link comes up.
- xdslRemoteInjection group identifies the processes at the remote site, indicating injection types and various traps.

Management is more difficult than with other technologies due to the fact that more information is processed in real-time or near real-time. In most cases, proprietary solutions dominate. The implementation of TMN and other standards is extremely slow.

3.2.4.8 Cable Technology

Cable service providers can enter the competition for voice and data services. Depending on the country, there are millions of households and businesses with cable television connections. In the majority of cases, cable television is a distribution channel supporting only one-way communication. But, using cable modems, channels could be provided for two-way-communications, allowing consumers to send back data or use the cable for phone conversations.

Figure 3.2.9 shows the structure of such an arrangement. The structure is simple requiring just a few additional networking devices. The components are:

- The equipment at the transmission site generates television signals and houses switches that route phone calls and monitor the network.
- Parallel fiber optic lines spread out over the area. If a line is out, traffic moves almost instantly to another route.
- Electronic nodes convert signals for transmission via coax cables. One node can serve approximately 500 to 1000 homes.
- Coaxial cable runs into a box on the side of the home. Electronics in the box split signals, sending phone traffic over regular phone lines, and television signals over cable lines to the television set.
- Phones would work on the current wiring. Television could be interactive, allowing signals to flow back up the cable line. Personal computers also could be connected through cable-modem, which transmits information up to 1000 times faster than phone lines.

Turning cable television systems into local phone networks will be not easy or cheap. Plans are ambitious to compete with local service providers. A basic cable system starts out as copper coax that carries signals in a line from the head-end — where TV signals are generated — to each consumer. Signals go one way; in the case of cable cuts or other damages, all the consumers from that point are out of service.

A system designed to handle phone calls and interactive TV looks much different. The main trunks of the network are set up in interconnecting rings of fiber-optic lines, which can carry thousands of times more information than a copper line. If a line is cut, traffic moves to another ring in a microsecond and practically no one is losing service. Because phone calls are two-way, a phone-cable system must be two-way and contain the sophisticated switches that send calls to the right places.

About 40% of today's cable TV systems are still copper. The other 60% have some fiber in their trunk lines, but most of them need upgrades. Just a few are set up in rings and have switches.

Power could cause another problem. Telephone lines have just enough electricity running through them to power the phone and the network, independent of the power grid. If a storm knocks out the main electrical grid, most of the time the phone still works. Cable lines don't carry power. If electricity goes out, so does cable. So would a phone hooked to an unpowered cable network. Cable lines, however, are capable of carrying power. But adding power from a node over the coaxial cable could add noise. There is a debate today in the industry about whether to get power from the cable-phone network, or use another solution, like attaching back-up batteries to the side of the households.

Once the pieces of the network come together, cable companies will face other issues. One is number portability — allowing consumers to keep their phone numbers even if they change phone service from the local phone company to the cable company. Right now, if somebody changes carriers, the number must be changed. State or federal regulation is expected to challenge that. Cable-phone systems also will have to prove to a sceptical public that they can be as reliable as current phone systems, which almost never break down. Cable systems also have to overcome a reputation for going out too frequently.

From a technology standpoint, the competition in the high-speed data services market will initially be between a dedicated architecture and a shared architecture (Figure 3.2.2). Both ISDN or other dedicated solutions, like xDSLs and cable-LAN services require infrastructure upgrades, and that means significant investments on the part of service providers. While most performance comparisons focus on peak bandwidth, other aspects of network usage, such as customer density and average session time, can also affect cost and quality of service.

The type of upgrade that is needed to deliver high-speed data over cable networks should be fairly obvious: cable television transmission is typically one way, but a data and phone network must permit two-way traffic. Limited forms of data service are possible by using the telephone line as a return route, but this is not a long-term solution. Some cable networks already have migrated to a hybrid fiber-coax infrastructure, but many operators are still immersed in the upgrading process. Once the networks are tailored for two-way traffic, broadband LAN technology will be incorporated so that digital data can be transmitted over a separate channel. Most of the vendors that supply technology to the cable operators will use an Ethernet-like approach, where the consumer's computer will have to be fitted with a network interface card and a cable modem for accessing the cable LAN. Access speed will depend upon the peak rate of the cable modem and the volume of traffic on the cable LAN. Subscribers will likely experience connection speeds that vary according to such factors as usage.

Cable LANs will operate full duplex with two channels, each channel sending data in a different direction. For customers using Internet or intranet applications, one of these channels would connect to an Internet or intranet POP router, which would then forward data packets from all users on the neighborhood LAN to and from all other systems on the Internet or on intranets. The other channel would receive data from the Internet or intranet service. Cable LANs are unlike ISDN-based services in that they furnish full-time connections, rather than switched or dial-up links. The obvious advantage is that it eliminates the need for dedicated transmission and switching resources. Users can access the cable LAN when needed, and only the cable network makes use of the Internet or Intranet point of presence.

A shared cable LAN requires only a single connection to the Internet or intranet provider. With ISDN, the lines of many individually served subscribers have to be multiplexed and concentrated, and the number of subscribers online at any given time is limited by the number of connections between the provider and the telecommunication network.

In terms of geographical coverage, cable-LANs can be quite large. At the beginning, they are intended to be implemented for residential broadband services. The future target is, however, corporate internetworking.

Currently, there are no standards governing the transmission of data over cable LANs, but that has not slowed down equipment suppliers. They bring proprietary products to the market as rapidly as possible.

The cable modem, which is used to connect the consumer's PC with a coax drop linked to the two-way, broadband cable LAN, is specifically designed for high-speed data communications. When the modem is in operation, cable television programming is still available. But the return path that the cable modem uses is limited and must be shared among all digital services, including interactive television, video on demand, telephony, videoconferencing, and data services. One problem with the cable modem is the immaturity of existing devices. The majority of manufacturers are targeting peak bandwidth between 128 Kbps and 30 Mbps.

Competing with cable technology are, first of all, ISDN and emerging technologies, such as ADSL, RADSL, HDSL, SDSL, and VDSL. They all use existing wires to corporations and to residential customers. Because ISDN delivers both voice and data, it makes use of circuit-switching and packet-switching technology. For high-speed data services, the local telephone plant must be upgraded so that two-way digital transmission — over the existing copper pairs that ISDN also relies upon — is possible. In addition, ISDN-capable digital switches must be installed at the telecommunications central office. ISDN may offer faster data transmission than analog modems, but the dial-up access model is still the same. From the data transmission perspective, this is an inefficient approach because circuit switching requires the service provider to dedicate resources to a customer at all times. In other words, it is not consistent with the bursty nature of data services. The result is wasted bandwidth.

Implementing an ISDN network could end up being more costly than deploying cable LANs. In telecommunications networks, e.g., an assessment of the incremental costs of providing ISDN Internet or intranet access includes the cost of initializing a subscriber's ISDN circuit at the central office; the cost of multiplexers, concentrators, and terminal servers; and possibly the cost of T1 lines from the telco to the Internet or intranet provider.

Cable-LAN access might prove to be less expensive for data and online services. It is assumed that cable television providers are facing less financial risks than telecommunications providers when deploying data services. Lower costs translate into consumer savings. The cost per bit of peak bandwidth (ROGE95) in providing Internet or Intranet access is significantly lower for hybrid fiber-coax networks than it is for ISDN — about 60 cents for a 4 Mbps residential service as opposed to \$16 for a 128 Kbps service. Shared fiber-coax networks also compete well against dedicated ISDNs on average bandwidth and peak bandwidth. A 4 Mbps residential cable service, for example, can provide the same average bandwidth and about 32 times the peak bandwidth of a 128 Kbps ISDN service for 40% less money.

While deployment costs for both technologies continue to decline, ISDN deployment still runs several hundred dollars more per subscriber than cable-LAN deployment. This difference, in addition to the higher performance of cable LANs, will be important factors as Internet, intranet, and online service access become a commodity product. Yet, the cable-LAN approach and the cable services industry have their shortcomings, too. Cable-LAN modems and access products are still proprietary, so once an operator has selected a vendor, it is likely to remain locked in. Also, shared networks function properly only when subscribers' usage habits are well known. The more subscribers deviate from an assumed usage profile, the more performance is likely to deteriorate. Big changes in usage could impact the cost of providing acceptable service levels over neighborhood cable-LANs.

Finally, the cable industry itself is greatly fragmented in many countries. This inhibits coordinated efforts which are vital to the rapid deployment of services. Although telecommunications providers are

in competition against each other, they will work better as a group if they act quickly in deploying the resources needed to make ISDN universally available.

In order to avoid these bottlenecks and other fault- and performance-related problems, cable service providers need a powerful management solution. It requires a complete rethinking of presently used solutions. The management solution should address high volumes of events, alarms, and messages, failsafe operations, traffic control and management, quality assurance of services, and collection of accounting data.

MIB Availability

MIB definitions are in the works with standard committees and with vendors. The present definition is not yet complete, but as a guideline is very useful for further steps.

The MIB is structured at present into six groups:

- The docsDevBase group extends the MIB-II “system” group with objects needed for cable device system management. It includes the device role, date and time, serial number, and reset conditions.
- The docsDevNmAccessGroup provides a minimum level of SNMP access security to the device by network management stations. The access is also constrained by the community strings and any vendor-specific security. The management station should have read-write permission for the cable modems.
- The docsDevSoftware group provides information for network downloadable software upgrades. It includes file identification, administration status, operational status, and current version.
- The docsDevServer group provides information about the progress of the interaction between the CM and CMTS and various provisioning servers. It includes boot state, DHCP, server time, configuration file, and TFTP configuration parameters.
- The docsDevEvent group provides control and logging for event reporting. This group offers very detailed information on control parameters, syslog details, throttle conditions, severity codes, and priorities. It also offers entries for vendor-specific data.
- The docsDevFilter group configures filters at link layer and IP layer for bridges data traffic. This group consists of a link-layer filter table; docsDevFilterLLCTable, which is used to manage the processing and forwarding of non-IP traffic; an IP packet classifier table, docsDevFilterIpTable, which is used to map classes of packets to specific policy actions; and a policy table, docsDevFilterPolicyTable, which maps zero or more policy action tables. At this time, the MIB specifies only one policy action table, docsDevFilterTosTable, which allows the manipulation of the type of services bits in an IP packet based on matching certain criteria. The working group may add additional policy types and action tables in the future, for example, to allow QoS to modem service identifier assignment based on destination.

Management is more difficult than with other technologies due to lack of management capabilities in managed devices. In most cases, proprietary solutions dominate. The implementation of TMN and other standards is extremely slow.

3.2.4.9 Management Solutions

The present status for emerging technologies can be summarized as follows:

- Proprietary solutions dominate: it means that the management protocols selected are controlled by the supplier of the equipment or facilities vendors. Working Groups do not exist, with the exception of the ATM Forum.
- Re-engineered by SNMP: many of the equipment vendors include SNMP agents into their devices to meet the requirements of customers. The SNMP agents provide information for performance management and reporting, but they usually do not change the real-time processing of status data within the devices. Most equipment is extremely intelligent. SNMP is simply too slow for certain decisions.

- The management structures are very heterogeneous: in most cases, domains are managed. Domain boundaries may be defined by the geography or by the family of managed objects. Most of the interfaces to element managers and managed devices are still proprietary. But, at least, the interfaces are going to be defined clearly by TMN.
- TMN has a very low penetration: suppliers have recognized the need for a generic standard, but they are not willing to invest heavily into supporting it. Some of the providers support the Q3-interface. There are no implementation examples for the higher layers of TMN. The network element and network management layers are well understood.
- Definition of MIBs: Each technology requires the definition of public and private MIBs. It is merely a question of time until each of the emerging technologies will be using MIB as the basis of management.
- Operating Support Systems are in the transitioning process: the legacy-type OSSs will give place to the new OSSs that are based on separating operations data from operating, the implementation of very flexible software, and separating network management from service management.

3.2.5 Summary

TMN will change the positions of managing emerged and emerging technologies. Customers will show little interest for the management solutions at lower layers. The primary target is to offer management capabilities at the service and business management layers. It needs significant extensions to existing management products. This segment of the handbook showed the status of management capabilities at element and network management layers only. In order to offer metrics for higher layers, MIB extensions and eventually new tools are required. If TMN layering combined with CNM, customers and providers may share the same management tools.

References

- ABER97 Aber, R.: xDSL Supercharges Copper, *Data Communications*, March 1997, p. 99-105.
- BLAC94 Black, U.: *Emerging Communications Technologies*, Prentice-Hall Series in Advanced Communication Technologies, Englewood Cliffs, NJ, 1994.
- ROGE95 Rogers, C.: Telcos versus Cable TV: The Global View, *Data Communications*, September 1995, p. 75-80.
- TERP98 Terplan, K.: *Telecom Operations Management Solutions*, CRC Press, Boca Raton, 1998.

3.3 Commercial Network and Systems Management Standards

Kornel Terplan

Industry standards help suppliers, vendors, and customers to communicate with each other more efficiently, but finding the common denominator of all interests is not easy. Service creation, provisioning, assurance, delivery, and service management cannot wait for fully completed standards. This chapter summarizes the most important management standards for multimedia communications. Besides TMN, CMIP, SNMP, RMON, OMA, and DMTF, Web-based standards are also introduced that may help to unify and simplify management frameworks and applications. The TeleManagement Forum may help a lot to improve the interoperability between multiple suppliers by providing practical guides, specifications, solution sets, and conformance documents.

3.3.1 Manager-Agent Relationship

Protocols always represent an agreement between the communicating parties. In the area of network management, management-agent relationships have been frequently implemented. [Figure 3.3.1](#) shows

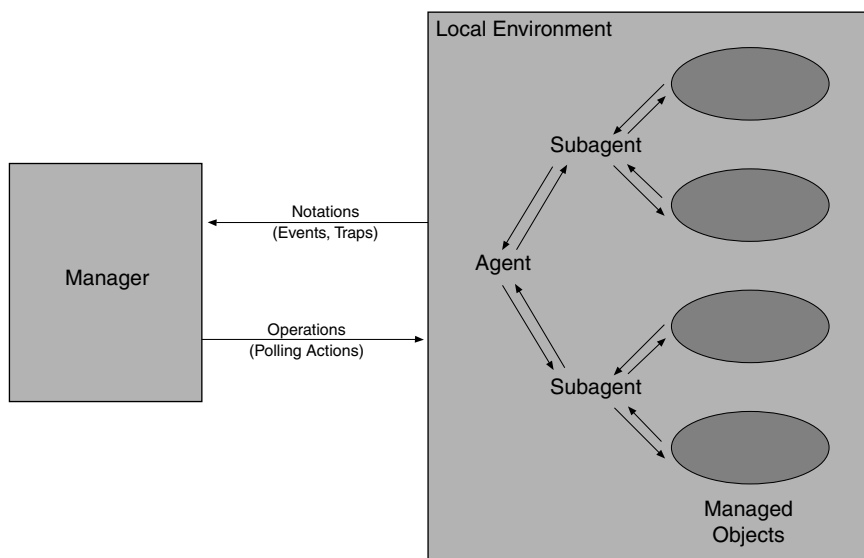


FIGURE 3.3.1 Communications paths between manager, agents, subagents, and managed objects.

this relationship. The agent side may be hierarchical by implementing subagents into specific devices. Depending on the nature of the management protocol, either the manager or the agents starts the dialog. There are always exceptions for high-priority networking events.

In case of CMIP, eventing techniques are used. Assuming that the agents are intelligent enough to capture, interpret, filter, and process events, they will notify the manager about alarming conditions. Of course, the manager is allowed to interrupt, and send inquiry-types of messages to the agents.

In case of SNMP, the manager polls the agents periodically. The agents respond by sending an abundance of information on device status and performance. Usually, the agents wait for the poll unless unusual events occur in the network. For such events, special traps can be defined and implemented.

3.3.2 Common Management Information Protocol (CMIP)

TMN protocols include OSI protocols such as CMIP and FTMP, ISDN, and Signalling System No. 7 protocols. They are organized into protocol suites or profiles for specific TMN interfaces. Functions and protocols support TMN services which include:

- Traffic management
- Customer management
- Switching management
- Management of transport networks
- Management of intelligent networks
- Tariffing and changing

The primary protocol is Common Management Information Protocol (CMIP). The estimated overhead scares both vendors and users away with the exception of the telecommunication industry, where separate channels can be used for management. CMIP is event driven, assuming processing capabilities at the agent level. Once fault and performance thresholds are exceeded, the manager is alarmed by the agent. This is similar to SNMP traps.

Open Systems Interconnected network management follows an object-oriented model; physical and logical real resources are managed through abstractions known as managed objects (MO). Management

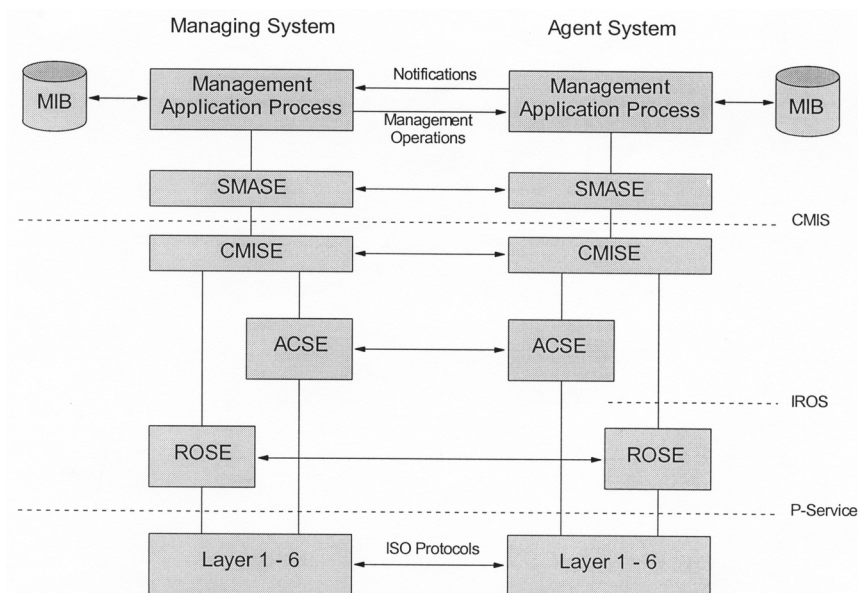


FIGURE 3.3.2 Application layer protocols and services.

systems also need managed objects that do not represent anything but exist for the needs of the management system itself. Most are handled by applications in agent roles and are accessed by applications in manager roles in order to implement management policies. The global collection of management information is consolidated in MIBs; each agent handles a part of it in its Management Information Tree. Information in the manager-agent interaction is conveyed through the management service/protocol common management information services (CMIS) and CMIP. In the context, agent, managed system, and managed node are synonymous; manager, managing application, and management stations are synonymous as well. CMIP is part of the OSI management framework. Its elements are:

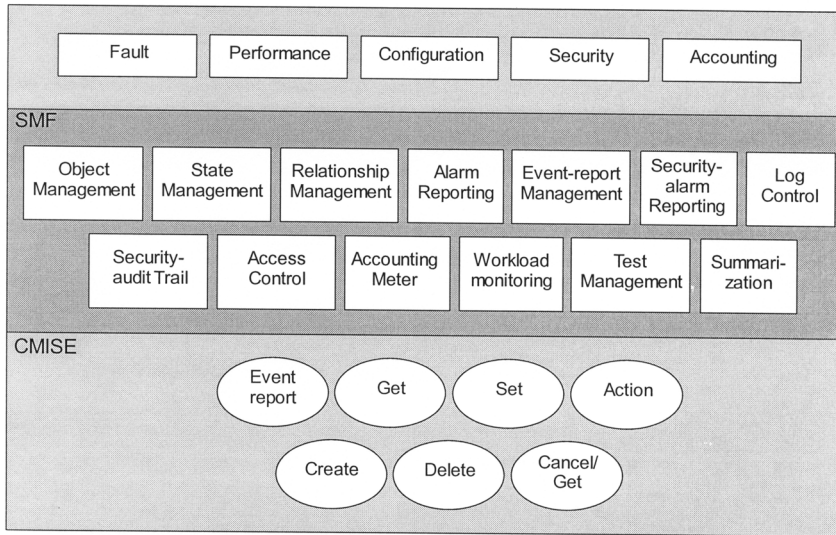
- CMIS/CMIP defines the services provided to management applications and the protocol providing the information exchange capability
- Systems management functions specifying all the functions to be supported by management
- Management information model that defines the MIB
- Layer management that defines management information, service, and functions to specific layers

Based on the seven-layer model, network-management applications are implemented in layer 7. Layers 1 through 6 contribute to network management by offering the standard services to carry network management-related information. Figure 3.3.2 shows the structure for the application layer. In particular, four system management application entities (SMAE) are very useful:

- For generic use: Association control service element (ACSE) and Remote operations service element (ROSE)
- For specific use: Common management information service element (CMISE) and System management application service element (SMASE)

Communication services in the OSI management model are provided by the CMIS. The service is realized through CMIP over a full or lightweight OSI stack. In the OSI world, two mappings are defined:

- a service over a full OSI stack (CMIP) and
- a connectionless one over Logical Link Control 1 (LLC1) (CMOL)



The SMASE provides additional support to the application and the functional areas via a set of SMFs

SMF System Management Functions

FIGURE 3.3.3 Overview of OSI management structure.

In the Internet world, there is a third mapping to provide the service over TCP/UDP using a lightweight presentation protocol (CMOT). CMOT and CMIP applications are portable on each other's stack with the same API, but will not work over CMOL.

Figure 3.3.3 shows the complete structure of CMISEs and SMFs. The overall goal is to support typical network management functions, such as fault, configuration, performance, security, and accounting management. SMFs or a group of SMFs support specific management functions. For the communication between entities, CMISE is used. Figure 3.3.4 summarizes these service elements between SMASE and CMISE. Each of the service elements are defined in Table 3.3.1.

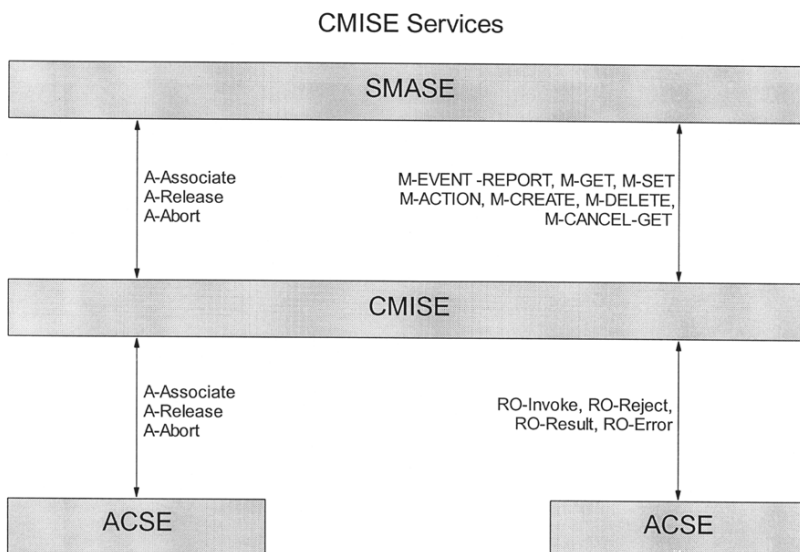


FIGURE 3.3.4 Information exchange using service elements.

TABLE 3.3.1 Definition of Service Elements

Service	Type	Definition
Management Notification Service		
M-event-report	Confirmed/unconfirmed	Reports on an event to a peer CMISE service user
Management Operations Service		
M-get	Confirmed	Requests retrieval of information from peer CMISE user
M-set	Confirmed/unconfirmed	Requests modification of information from peer CMISE user
M-action	Confirmed/unconfirmed	Requests that peer CMISE user perform some action
M-create	Confirmed	Requests that peer CMISE user create an instance of a Managed Object
M-delete	Confirmed	Requests that peer CMISE user delete an instance of a Managed Object
M-cancel-get	Confirmed	Requests that peer CMISE user cancel outstanding invocation of M-get service

Further details can be found in (TERP96) — Effective Mgmt of LANs, McGraw-Hill.
The strengths of CMIP include:

- General and exensible object-oriented approach
- Support from the telecommunication industry
- Support for manager-to-manager communications
- Support for a framework for automation

Weaknesses of CMIP are:

- It is complex and multilayered
- High overhead is the price of many confirmations
- Few CMIP-based management systems are shipping
- Few CMIP-based agents are in use

CMIP assumes the use of the OSI stack for exchanging CMIP protocol data units. In layer 7, there are also other applications which may be combined with CMIP.

3.3.3 Simple Network Management Protocol (SNMPv1, SNMPv2, and SNMPv3)

In the SNMP environment, the manager can obtain information (Figure 3.3.5) from the agent by polling managed objects periodically. Agents can transmit unsolicited event messages, called traps, to the manager. The management data exchanged between managers and agents is called the management information base (MIB). The data definitions outlined in SMI must be understood by both managers and agents.

The manager is a software program residing within the management station. The manager has the ability to query agents using various SNMP commands. The management station is also in charge to interpret MIB data, construct views of the systems and networks, compress data, and maintain data in relational or object-oriented databases. A traditional SNMP manager is shown in Figure 3.3.6 (STAL99).

This figure shows typical functions that are not only valid for managers, but also for managed agents. Figure 3.3.7 shows the typical functional blocks of SNMP agents.

In a traditional manager, the SNMP engine contains a dispatcher, a message processing subsystem, and a security subsystem. The dispatcher is a simple traffic manager. For outgoing PDUs, the dispatcher accepts PDUs from applications and performs the following fuctions. For each PDU, the dispatcher determines the type of message processing required — it may be different for SNMP versions 1, 2, and 3 — and passes the PDU on to the appropriate message processing module in the message processing subsystem. Subsequently, the message processing subsystem returns a message containing that PDU and

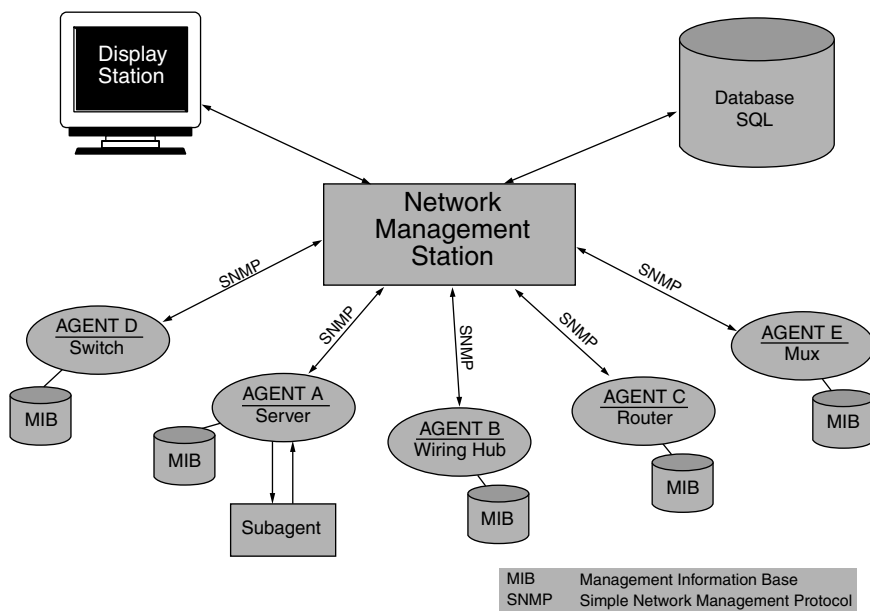


FIGURE 3.3.5 Structure of SNMP-based management services.

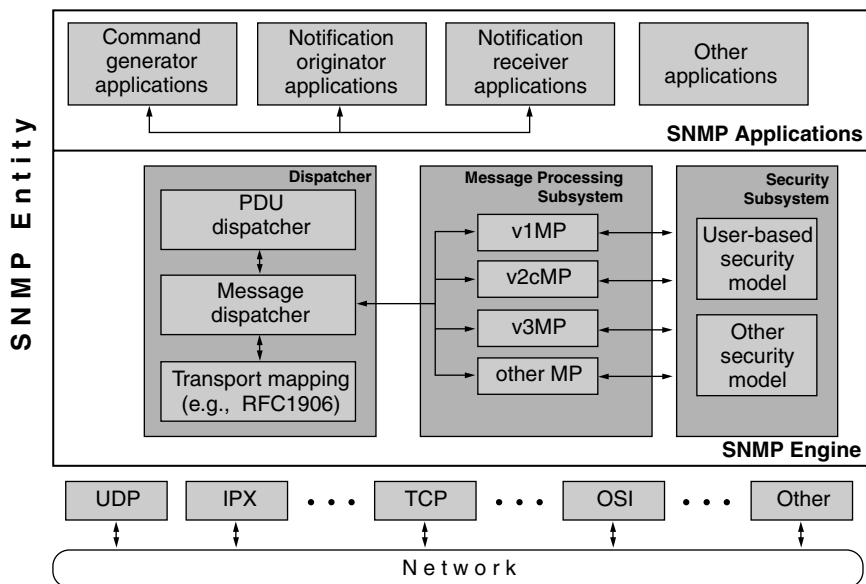


FIGURE 3.3.6 Traditional SNMP manager.

including the appropriate message headers. The dispatcher then maps this message onto a transport layer for transmission. For incoming messages, the dispatcher accepts messages from the transport layer and performs the following functions. The dispatcher routes each message to the appropriate message processing module. Subsequently, the message processing subsystem returns the PDU contained in the message. The dispatcher then passes this PDU to the appropriate application.

The message processing subsystem accepts outgoing PDUs from the dispatcher and prepares these for transmission by wrapping them in the appropriate message header and returning them to the dispatcher.

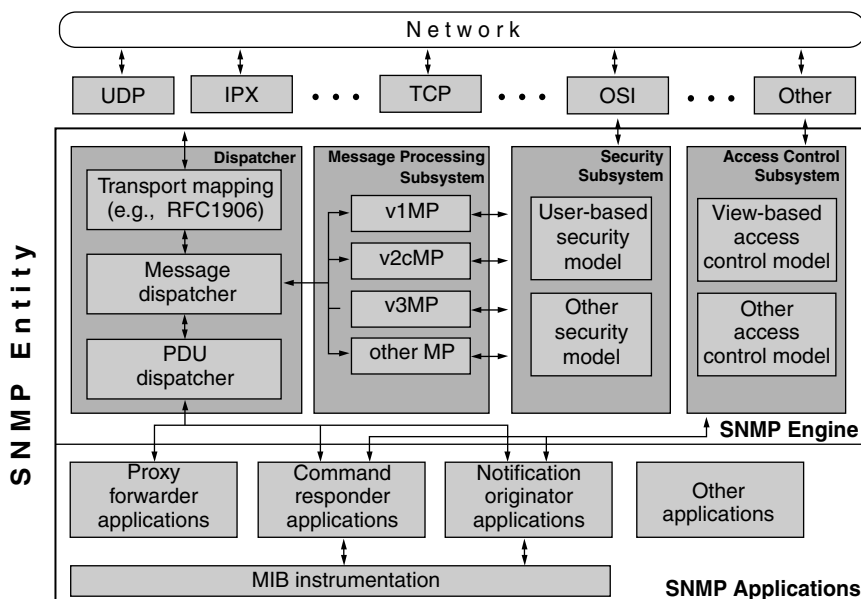


FIGURE 3.3.7 Traditional SNMP agent.

It also accepts incoming messages from the dispatcher, processes each message header, and returns the enclosed PDU to the dispatcher. An implementation of the message processing subsystem may support a single message format corresponding to a single version of SNMP, or may contain a number of modules, each supporting a different version of SNMP.

The security subsystem performs authentication and encryption functions. Each outgoing message is passed to the security subsystem from the message processing subsystem. Depending on the services required, the security subsystem may encrypt the enclosed PDU, and it may generate an authentication code and insert it into the message header. The processed message is then returned to the message processing subsystem. Similarly, each incoming message is passed to the security subsystem from the message processing subsystem. If required, the security subsystem checks the authentication code and performs decryption. An implementation of the security subsystem may support one or more distinct security models.

The SNMP engine for a traditional agent has all the components found in the SNMP engine for a traditional manager, plus an Access Control Subsystem. This subsystem provides services to control access to MIBs for the reading and setting of management objects. These services are performed on the basis of the contents of PDUs. An implementation of the security subsystem may support one or more distinct access control models. The security-related functions are organized into two separate subsystems: security and access control. This is an excellent example of modular design, because the two subsystems perform quite distinct functions and therefore it makes sense to allow standardization of these two areas to proceed independently. The security subsystem is concerned with privacy and authentication, and operates on SNMP messages. The access control subsystem is concerned with authorized access to management information, and operates on SNMP PDUs.

For both SNMP managers and agents, there are a number of components that constitute the management functionalities. These components are listed in [Table 3.3.2](#) (STAL99).

In terms of MIBs, there are continuing changes. In addition to standard MIBs, such as MIB I and II ([Table 3.3.3](#)), the IETF has defined a number of adjunct MIBs covering hosts, routers, bridges, hubs, repeaters, FDDI networks, AppleTalk networks, frame relay networks, switches, ATM nodes, mobile components, and applications.

TABLE 3.3.2 Components of SNMP Entity

Dispatcher	— Allows for concurrent support of multiple versions of SNMP messages in the SNMP engine. It is responsible for (1) accepting PDUs from applications for transmission over the network and delivering incoming PDUs to applications; (2) passing outgoing PDUs to the message processing subsystem to prepare as messages, and passing incoming messages to the message processing subsystem to extract the incoming PDUs; and (3) sending and receiving SNMP messages over the network.
Message Processing Subsystem	— Responsible for preparing messages for sending, and for extracting data from received messages.
Security Subsystem	— Provides security services such as the authentication and privacy of messages. This subsystem potentially contains multiple security models.
Access Control Subsystem	— Provides a set of authorization services that an application can use for checking access rights. Access control can be invoked for retrieval or modification request operations and for notification generation operations.
Command Generator	— Initiates SNMP Get, GetNext, GetBulk, and/or Set request PDUs and processes the response to a request that it has generated.
Command Responder	— Receives SNMP Get, GetNext, GetBulk, and/or set request PDUs destined for the local system as indicated by the fact that the contextEngineID in the received request is equal to that of the local engine through which the request was received. The command responder application will perform the appropriate protocol operation, using access control, and will generate a response message to be sent to the request's originator.
Notification Originator	— Monitors a system for particular events or conditions, and generates Trap and/or Inform messages based on these events or conditions. A notification originator must have a mechanism for determining where to send messages, and which SNMP version and security parameters to use when sending messages.
Notification Receiver	— Listens for notification messages, and generates response messages when a message containing an Inform PDU is received.
Proxy Forwarder	— Forwards SNMP messages. Implementation of a proxy forwarder application is optional.

TABLE 3.3.3 MIB II Structure

11 Categories of Management (2) Subtree	Information in the Category
System (1)	Network device operating system
Interfaces (2)	Network interface specific
Address translation (3)	Address mappings
Ip (4)	Internet protocol specific
Icmp (5)	Internet control message protocol specific
Tcp (6)	Transmission protocol specific
Udp (7)	User datagram protocol specific
Egp (8)	Exterior gateway protocol specific
Cmot (9)	Common management information services on TCP specific
Transmission (10)	Transmission protocol specific
Snmp (11)	Snmp specific

In terms of SNMP, the following trends are expected. SNMP agent-level support will be provided by an even greater number of vendors. SNMP manager-level support will be provided by only a few leading vendors in the form of several widely-accepted platforms. Management platforms provide basic services, leaving customization to vendors and users.

Wider use of intelligent agents is also expected. Intelligent agents are capable of responding to a manager's request for information and performing certain manager-like functions, including testing for thresholds, filtering, and processing management data. Intelligent agents enable localized polling and filtering on servers, workstations, and hubs, for example. Thus, these agents reduce polling overhead and management data traffic, forwarding only the most critical alerts and processed data to the SNMP manager.

RMON MIB will help to bridge the gap between the limited services provided by management platforms and the rich sets of data and statistics provided by traffic monitors and analyzers.

The strengths of SNMP include:

- Agents are widely implemented
- Simple to implement
- Agent-level overhead is minimal
- Polling approach is good for LAN-based managed object
- Robust and extensible
- Offers the best direct manager-agent interface

SNMP met a critical need; it was available and implementable at the right time.

The weaknesses of SNMP include:

- Too simple, does not scale well
- No object-oriented data view
- Unique semantics make integration with other approaches difficult
- High communication overhead due to polling, in particular for WAN-based management objects
- Many implementation-specific (private MIB) extensions
- No standard control definition
- Small agent (one agent per device) may be inappropriate for systems management

SNMP is being continuously improved and extended. SNMPv2 addresses many of the shortcomings of version 1. SNMPv2 can support either a highly centralized management strategy or a distributed one. In the latter case, some systems operate both in the role of manager and agent. In its agent role, such a system will accept commands from a superior manager; these commands may deal with access of information stored locally at the intermediate manager or may require the intermediate manager to provide summary information about subagents. The principal enhancements to SNMPv1 provided by version 2 fall into the following categories (STAL99):

- Structure of management information is being expanded in several ways. The macro used to define object types has been expanded to include several new data types and to enhance the documentation associated with an object. A noticeable change is a new convention has been provided for creating and deleting conceptual rows in a table. The origin of this capability is from remote monitoring (RMON).
- Transport mappings help to use different protocol stacks to transport the SNMP information, including user datagram protocol, OSI connectionless-mode protocol, Novell internetwork (IPX) protocol and Appletalk.
- Protocol operations with the most noticeable change of including two new PDUs. The GetBulkRequest PDU enables the manager to efficiently retrieve large blocks of data. In particular, it is powerful in retrieving multiple rows in a table. The InformRequest PDU enables one manager to send trap-type information to another.
- MIB extensions contain basic traffic information about the operation of the SNMPv2 protocol; this is identical to SNMP MIB II. The SNMPv2 MIB also contains other information related to the configuration of SNMPv2 manager to agent.
- Manager-to-manager capability is specified in a special MIB, called M2M. It provides functionality similar to the RMON MIB. In this case, the M2M MIB may be used to allow an intermediate manager to function as a remote monitor of network media traffic. Also, reporting is supported. Two major groups, Alarm and Event, are supported.
- SNMPv2 security does include a wrapper containing authentication and privacy information as a header to PDUs.

The SNMPv2 framework is derived from the SNMP framework. It is intended that the evolution from SNMP to SNMPv2 be seamless. The easiest way to accomplish this is to upgrade the manager to support SNMPv2 in a way that allows the coexistence of SNMPv2 managers, SNMPv2 agents, and SNMP agents. In order to map commands mutually into the target protocol, proxy agents are used (STAL99). The actual implementation of the proxy agent depends on the vendor; it could be implemented into the agent or into the manager.

The key new feature of SNMPv3 is better security. The design goals for v3 can be summarized as follows (STAL99):

- Use existing work, for which there is implementation experience and some consensus as to its value. As a result, the SNMP architecture and SNMPv3 security features rely heavily on SNMPv2u and SNMPv2*.
- Address the need for secure Set request messages over real-world networks, rectifying the most important deficiency of SNMPv1 and SNMPv2.
- Design a modular architecture that will (1) allow implementation over a wide range of operational environments, some of which need minimal, inexpensive functionality and some of which may support additional features for managing large networks; (2) make it possible to move portions of the architecture forward in the standards track even if consensus has not been reached on all pieces; and (3) accommodate alternative security models.
- Keep SNMP as simple as possible, despite the many necessary and useful extensions.

Based on these design goals, developers have made the following design decisions (STAL99):

- Architecture: An architecture should be defined which identifies the conceptual boundaries between the documents. Subsystems should be defined which describe the abstract services provided by specific portions of a SNMP framework. Abstract service interfaces, as described by service primitives, define the abstract boundaries between documents, and the abstract services that are provided by the conceptual subsystems of a SNMP framework.
- Self-contained documents: Elements of procedure plus the MIB objects that are needed for processing a specific portion of a SNMP framework should be defined in the same document, and as much as possible, should not be referenced in other documents. This allows pieces to be designed and documented as independent and self-contained parts, which are consistent with the general SNMP MIB module approach. As portions of SNMP change over time, the documents describing other portions of SNMP are not directly impacted. This modularity allows, for example, security models, authentication and privacy mechanisms, and message formats to be upgraded and supplemented as the need arises. The self-contained documents can move along the standards track on different time lines.
- Remote configuration: The security and access control subsystems add a whole new set of SNMP configuration parameters. The security subsystem also requires frequent changes of secrets at the various SNMP entities. To make this deployable in a large operational environment, these SNMP parameters must be able to be remotely configured.
- Controlled complexity: It is recognized that manufacturers of simple managed devices want to keep the resources used by SNMP to a minimum. At the same time, there is a need for more complex configurations which can spend more resources for SNMP and thus provide more functionality. The design tries to keep the competing requirements of these two environments in balance and allows the more complex environment to logically extend the simple environment.
- Threats: The security models in the security subsystem should protect against the principal threats, such as modification of information, masquerade, message stream modification, and disclosure. They do not need to protect denial of service and traffic analysis.

SNMPv3 is secure against the following threats:

- **Modification of information:** Any entity could alter an in-transit message generated by an authorized entity in a way that effects unauthorized management operations, including the setting of object values. The essence of this threat is that an unauthorized entity could change any management parameter, including those related to configuration, operations, and accounting.
- **Masquerade:** Management operations that are not authorized for some entity may be attempted by that entity by assuming the identity of an authorized entity.
- **Message stream modification:** SNMP is designed to operate over a connectionless transport protocol. There is a threat that SNMP messages could be reordered, delayed, duplicated, or replayed to effect unauthorized management operations. For example, a message to reboot a device could be copied and replayed later.
- **Disclosure:** An entity could observe exchanges between a manager and an agent and thereby learn the values of managed objects and learn of notifiable events. For example, the observation of a set command that changes passwords would enable an attacker to learn the new passwords.

But SNMPv3 is not intended to secure against the following two threats:

- **Denial of service:** An attacker may prevent exchanges between a manager and an agent.
- **Traffic analysis:** An attacker may observe the general pattern of traffic between managers and agents.

SNMPv3 is still in its initial implementation phase. If fully implemented, the power of SNMP will be significantly improved.

3.3.4 Remote Monitoring (RMON1 and RMON2)

The RMON MIB will help to bridge the gap between the limited services provided by management platforms and the rich sets of data and statistics provided by traffic monitors and analyzers. RMON defines the next generation of network monitoring with more comprehensive network fault diagnosis, planning, and performance tuning features than any current monitoring solution. The design goals for RMON are (STAL99):

- **Off-line operation:** In order to reduce overhead over communication links, it may be necessary to limit or halt polling of a monitor by the manager. In general, the monitor should collect fault, performance, and configuration information continuously, even if it is not being polled by a manager. The monitor simply continues to accumulate statistics that may be retrieved by the manager at a later time. The monitor may also attempt to notify the manager if an exceptional event occurs.
- **Preemptive monitoring:** If the monitor has sufficient resources, and the process is not disruptive, the monitor can continuously run diagnostics and log performance. In the event of a failure somewhere in the network, the monitor may be able to notify the manager and provide useful information for diagnosing the failure.
- **Problem detection and reporting:** Preemptive monitoring involves an active probing of the network and the consumption of network resources to check for error and exception conditions. Alternatively, the monitor can passively — without polling — recognize certain error conditions and other conditions, such as congestions and collisions, on the basis of the traffic that it observes. The monitor can be configured to continuously check for such conditions. When one of these conditions occurs, the monitor can log the condition and notify the manager.
- **Value-added data:** The network monitor can perform analyses specific to the data collected on its subnetworks, thus offloading the manager of this responsibility. The monitor can, for instance,

TABLE 3.3.4 RMON MIB Groups for Ethernet

Statistics Group	Features a table that tracks about 20 different characteristics of traffic on the Ethernet LAN segment, including total octets and packets, oversized packets, and errors.
History Group	Allows a manager to establish the frequency and duration of traffic observation intervals, called “buckets.” The agent can then record the characteristics of traffic according to these bucket intervals.
Alarm Group	Permits the user to establish the criteria and thresholds that will prompt the agent to issue alarms.
Host group	Organizes traffic statistics by each LAN node, based on time intervals set by the manager.
HostTopN Group	Allows the user to set up ordered lists and reports based on the highest statistics generated via the host group.
Matrix Group	Maintains two tables of traffic statistics based on pairs of communicating nodes; one is organized by sending node addresses, the other by receiving node addresses.
Filter Group	Allows a manager to define, by channel, particular characteristics of packets. A filter might instruct the agent, for example, to record packets with a value that indicates they contain DECnet messages.
Packet Capture Group	This group works with the Filter Group and lets the manager specify the memory resources to be used for recording packets that meet the filter criteria.
Event Group	Allows the manager to specify a set of parameters or conditions to be observed by the agent. Whenever these parameters or conditions occur, the agent will record an event into a log.

TABLE 3.3.5 RMON MIB Groups for Token Ring

Statistics Group	This group includes packets, octets, broadcasts, dropped packets, soft errors and packet distribution statistics. Statistics are at two levels: MAC for the protocol level and LLC statistics to measure traffic flow.
History Group	Long-term historical data for segment trend analysis. Histories include both MAC and LLC statistics.
Host Group	Collects information on each host discovered on the segment.
HostTopN Group	Provides sorted statistics that allow reduction of network overhead by looking only at the most active hosts on each segment.
Matrix Group	Reports on traffic errors between any host pair for correlating conversations on the most active nodes.
RingStation Group	Collects general ring information and specific information for each station. General information includes: ring state (normal, beacon, claim token, purge); active monitor; and number of active stations. Ring Station information includes a variety of error counters, station status, insertion time, and last enter/exit time.
Ring Station Order	Maps station MAC addresses to their order in the ring.
Source Routing Statistics	In source-routing bridges, information is provided on the number frames and octets transmitted to and from the local ring. Other data includes broadcasts per route and frame counter per hop.
Alarm Group	Reports changes in network characteristics based on thresholds for any or all MIBs. This allows RMON to be used as a proactive tool.
Event Group	Logging of events on the basis of thresholds. Events may be used to initiate functions such as data capture or instance counts to isolate specific segments of the network.
Filter Group	Definitions of packet matches for selective information capture. These include logical operations (AND, OR, NOT) so network events can be specified for data capture, alarms, and statistics.
Packet Capture Group	Stores packets that match filtering specifications.

observe which station generates the most traffic or errors in network segments. This type of information is not otherwise accessible to the manager that is not directly attached to the segment.

- **Multiple managers:** An internetworking configuration may have more than one manager in order to achieve reliability, perform different functions, and provide management capability to different units within an organization. The monitor can be configured to deal with more than one manager concurrently.

Table 3.3.4 summarizes the RMON MIB groups for Ethernet segments. Table 3.3.5 defines the RMON MIB groups for Token Ring segments. At the present time, here are just a few monitors that can measure both types of segments using the same probe.

RMON is very rich on features and there is the very real risk of overloading the monitor, the communication links, and the manager when all the details are recorded, processed, and reported. The

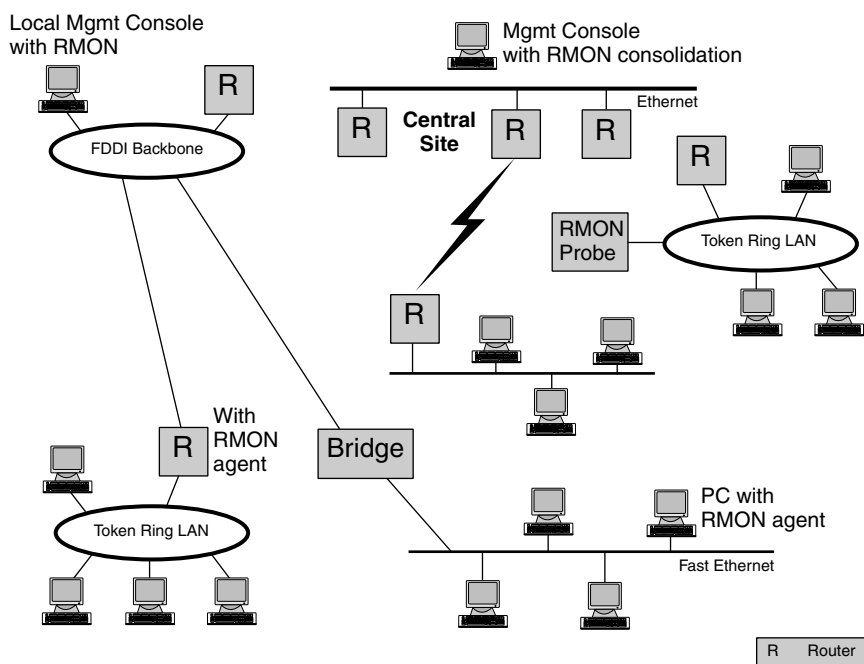


FIGURE 3.3.8 RMON probes in LAN segments.

preferred solution is to do as much of the analysis as possible locally, at the monitor, and send just the aggregated data to the manager. This assumes powerful monitors. In other applications, monitors may be reprogrammed during operations by the managers. This is very useful when diagnosing problems. Even if the manager can define specific RMON requests, it is still necessary to be aware of the trade-offs involved. A complex filter will allow the monitor to capture and report a limited amount of data, thus avoiding overhead on the network. However, complex filters consume processing power at the monitor; if too many filters are implemented, the monitor will become overloaded. This is particularly true if the network segments are busy, which is probably the time when measurement is most valuable.

Figure 3.3.8 shows the RMON probes in the segments. RMON probes can be implemented in three different ways:

- Probe as a standalone monitor
- Probe as a module of hubs, routers, and switches
- Probe as a software module in Unix, NT operating systems, or in PC workstations

Each of these alternatives shows benefits and disadvantages.

1. *Probe as a Standalone Monitor*

The benefits of this alternative are:

- Excellent performance
- Support of all functions
- Availability in various options, such as stackable or rack-mountable

The disadvantages are:

- High costs for an average LAN segment
- Multiple probes are required, when segmentation of LANs is deployed by switches without using probes-ports
- Most advanced LAN technologies might not be supported right away

2. *Probe as a Module of Hubs, Routers, and Switches*

The benefits of this alternative are:

- It represents a very convenient solution because the networking components have already been deployed
- The costs are much lower than in case of standalone probes
- The integration of probes into switches is less expensive than deploying an individual probe in each switched segment

The disadvantages are:

- The networking components need an upgrade or customization to support the probes
- Upgrades are not always economical
- The performance might become a problem
- The conformance to standards may not be met
- Problems with the probe may impact the performance of the networking component
- The RMON functionality may be very much limited; not all RMON indicator groups are supported
- Processing programs are very simple in comparison to standalone probes
- Integration with management platforms is usually incomplete
- RMON modules may be provided by different vendors, which may lead to incompatibility problems

3. *Probe as a Software Module in Unix, NT Operating Systems, or in PC Workstations*

The benefits of this alternative are:

- Much lower costs in comparison to any other alternatives
- Performance is good when run on RISC or Pentium processors
- Scalability and extendability are excellent
- Support of state-of-the technology is easier than with other alternatives
- Combination of supervising Ethernet and Token Ring is possible
- Outband access to probes is possible with proper configuration

The disadvantages are:

- Purchase of adapters and additional workstations may be required
- The user is responsible for the installation

The extension of this alternative might be utilized also for switched LANs. RMON may be installed into the adapter of end-user devices. Usually, just the filter- and packet capture groups are supported at the end-user device level. The other groups are supported by the collector. The overhead is minimal; performance impacts in the switch and in the end-user devices are minimal.

RMON probes are extremely helpful for collecting data on Web site accesses and activities. Independently, how probes are implemented, vendors of probes are expected to work together. Standards are continuously improved offering even more functionality of capturing and processing Web sites-relevant data.

The existing and widely used RMON version is basically a MAC standard. It does not give LAN managers visibility into conversations across the network or connectivity between various network segments. The extended standard is targeting the network layer and higher. It will give visibility across the enterprise. With remote access and distributed workgroups, there is a substantial intersegment traffic. The following functionalities are included:

- Protocol distribution
- Address mapping

- Network layer host table
- Network layer matrix table
- Application layer host table
- Application layer matrix table
- User history collection table
- Protocol configuration.

Protocol distribution and protocol directory table — The issue here was how to provide a mechanism that will support the large number of protocols running on any one network. Current implementations of RMON employ a protocol filter which analyzes only the essential protocols. RMON2, however, will employ a protocol directory system which allow RMON2 applications to define which protocols an agent will employ. The Protocol Directory Table will specify the various protocols a RMON2 probe can interpret.

Address Mapping — This feature matches each network address with a specific port to which the hosts are attached. It also identifies traffic-generating nodes/hosts by MAC, Token Ring, or Ethernet address. It helps identify specific patterns of network traffic. Useful in node discovery and network topology configurations. In addition, the address translation feature adds duplicate IP address detection, resolving a common troubleshooting with network routers and virtual LANs.

Network-layer Host Table — Tracks packets, errors, and bytes for each host according to a network-layer protocol. It permits decoding of packets based on their network layer address. In essence, permitting network managers to look beyond the router at each of the hosts configured on the network.

Network-layer Matrix Table — Tracks the number of packets sent between a pair of hosts by network layer protocol. The network manager can identify network problems quicker by using the matrix table which shows the protocol-specific traffic between communicating pairs of systems.

Application-layer Host Table — Tracks packets, errors, and bytes by host on an application-specific basis (e.g., Lotus Notes, e-Mail, WEB, etc.) Both the application-layer host table and matrix table trace packet activity of a particular application. This feature can be used by network managers to charge back users on the basis of how much network bandwidth was used by their applications.

Application-layer Matrix Table — Tracks packet activity between pairs of hosts by application (e.g., pairs of hosts exchanging internet information).

Probe Configuration — Currently, vendors offer a variety of proprietary means for configuring and controlling their respective probes. This complicates interoperability. The Probe Configuration Specification, based on the Aspen MIB, defines standard parameters for remotely configuring probes — parameters such as network address, SNMP error trap destinations, modern communications with probes, serial line information, and downloading of data to probes. It provides enhanced interoperability between probes by specifying standard parameters for operations, permitting one vendor's RMON application the ability to remotely configure another vendor's RMON probe.

User History Collection Group — The RMON2 history group polls, filters, and stores statistics based on user-defined variables, creating a log of the data for use as a historical tracking tool. This is in contrast to RMON1, where historical data is gathered on a predefined set of statistics.

After implementation, more and more complete information will be available for performance analysis and capacity planning.

3.3.5 Desktop Management Interface (DMI) from Desktop Management Task Force (DMTF)

Basically, SNMP may be utilized to manage systems, assuming system components accommodate SNMP agents. But there are no MIBs yet that describe principal indicators for management purposes. An important emerging standard for desktop management is the Desktop Management Interface (DMI). The Desktop Management Task Force (DMTF) has defined the DMI to accomplish the following goals:

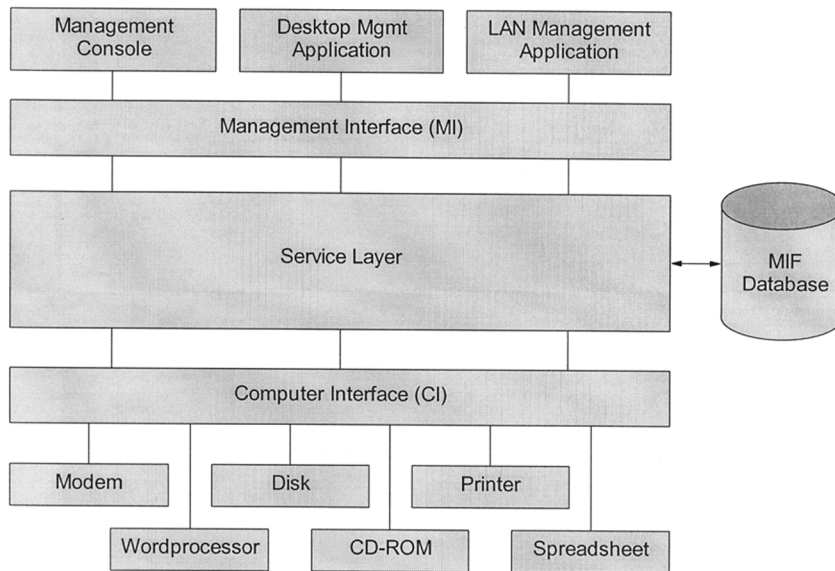


FIGURE 3.3.9 Structure of the desktop management interface.

- Enable and facilitate desktop, local, and network management
- Solve software overlap and storage problems
- Create a standard method for management of hardware and software components using MIFs (management information format)
- Provide a common interface for managing desktop computers and their components
- Provide a simple method to describe, access, and manage desktop components

The scope of management under DMTF includes CPUs, I/Os, mother boards, video cards, network interface cards, faxes, modems, mass storage devices, printers, and applications. Figure 3.3.9 shows the structure of this standard. There is a clear separation between the managed components (CI) and the services offered for management (MI). The commands are similar to SNMP, but they are not identical.

3.3.6 Object Management Architecture (OMA) from OMG

OMA permits the cooperation of distributed objects independently from their locations. OMA has not been designed and structured only for network management. It is a general purpose standard. In particular, it helps to design and implement distributed applications. This architecture has gained interest because the OMA technology is going to be increasingly used for end-user devices.

The organizational model is using a peer-to-peer approach with the result that communicating objects are equivalent in their importance. Figure 3.3.10 shows this architecture.

The communication model is based on CORBA (Common Object Request Broker Architecture). The Object Request Broker is the coordinator between distributed objects. The broker receives messages, inquiries, and results from objects, and routes them to the right destination. If the objects are in a heterogeneous environment, multiple brokers are required. They will talk to each other in the future by a new protocol based on TCP/IP. There is no information model available; no operations are predefined for objects. But an object does exist containing all the necessary interfaces to the object request broker. For the description, the interface definition language (IDL) is being used. There are no detailed MIBs for objects because OMA is not management specific.

The functional model consists of the Object Services Architecture. It delivers the framework for defining objects, services, and functions. Examples for services are instantiation, naming, storing objects'




FIGURE 3.3.10 Open management architecture from OMG.

attributes, and the distribution/receipt of events and notification. CORBA services and facilities represent more generic services; they are expected to occur in multiple applications or they are used in specific applications. The driving force beyond designing common facilities for systems management is the X/Open Systems Management Working Group. The Managed Set Service, defined by this group, encourages grouping of objects in accordance of their management needs, with the result of easier administration. In the future, more services are expected to be defined; the next is an Event Management Service that expands the present Object Event Service by a flexible mechanism of event filtering.

DCOM is the heart of Microsoft's ActiveOSS product suite. Basically, DCOM is an integration infrastructure designed to facilitate communication between software components operating on the same host or with DCOM on multiple networked hosts. It was originally developed to create interoperability between components. It is the most widely deployed component object model. Active OSS acts as a centralized management and translation point for an OSS network. Conceptually, applications ride on top of the framework, but communicate through it. DCOM abstracts various application interfaces into objects, basically mapping the functions of the application into a common model that can be stored in a database. The common model allows the various applications to communicate in a uniform manner within the framework or across multiple networked frameworks.

3.3.7 Standards for Web-based Systems and Network Management

In July 1996, five major vendors announced an initiative to define *de facto* standards for Web-based Enterprise Management (WBEM). This effort, spearheaded by Microsoft, Compaq, Cisco, BMC, and Intel, was publicly endorsed by over 50 other vendors as well. The initial announcement called for defining the following specifications:

- HyperMedia Management Schema (HMMS): an extensible data description for representing the managed environment that was to be further defined by the Desktop Management Task Force (DMTF).
- HyperMedia Object Manager (HMOM): data model consolidating management data from different sources; a C++ reference implementation and specification, defined by Microsoft and Compaq, to be placed in the public domain.
- HyperMedia Management Protocol (HMMP): a communication protocol embodying HMMS, running over HTTP and with interfaces to SNMP and DMI.

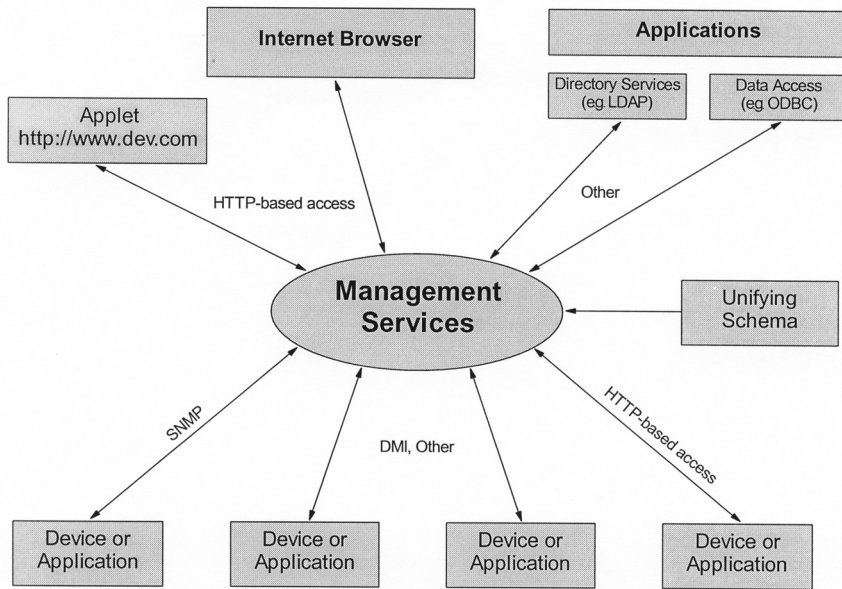


FIGURE 3.3.11 Use of the object manager to coordinate management services.

SunSoft has also announced a programming environment for developing Web-based network and systems management software. This environment, called Solstice Workshop, consists of a Java Management API (JaMAPI), a small footprint database, and a Java programming environment. Solstice Workshop's big drawing card is its extensibility and the popularity of Java's "write once run anywhere" appeal. JaMAPI requires Java, whereas HMMP/HMMS/HMOM specifies HTML/HTTP, although Java is not specifically excluded.

Among these two efforts, the WBEM is certainly the broadest in scope, addressing not only protocol issues, but also data modeling and extensible data description. While JaMAPI includes object class definitions, it does not go as far as data description.

The complete solution is envisioned in Figure 3.3.11 including both major directions of webification.

The initial euphoria over WBEM and JaMAPI is starting to wear off, and it is time for doing the hard work of pounding out specifications and, more importantly, building products. Customer demand will push Web-based management to its limits, but disillusion is sure to set in if a lack of progress becomes obvious on the standards' front. But there are several emerging products that have been developed with an eye for supporting current and future standards, that bring to market a practical approach to take advantage of Web-based management.

The DMTF has developed a common information model (CIM) to take advantage of object-based management tools and provide a common way to describe and share enterprise-wide management information. Using HMMS as an input, the new model can be populated by DMI 2.0 and other management data suppliers, including SNMP and CMIP, and implemented in multiple object-based execution models such as JaMAPI, CORBA, and HMM. CIM will enable applications from different developers on different platforms to describe and share management data, so users have interoperable management tools that span applications, systems, and networks, including the Internet.

CIM is a conceptual information model for describing management that is not bound to a particular implementation. This allows for the interchange of management information between management systems and applications. This can be either "agent to manager" and "manager to manager" communications, which provide for distributed system management.

In a fully CIM-compliant world, it should be possible to build applications such as Service Level Agreement Tracking applications using management data from a variety of sources and different management

systems such as TME, OpenView, ManageWise, SMS, etc. The management data would be collected, stored and analyzed using a common format (CIM) while allowing property extensions providing “value add.”

There are two parts to CIM: the CIM specification and the CIM schema.

The CIM specification describes the language, naming, meta schema, and mapping techniques to other management models such as SNMP MIBs, and DMTF MIFs, etc. The meta schema is a formal definition of the model. It defines the terms used to express the model and their usage and semantics. The elements of the meta schema are classes, properties, and methods. The meta schema also supports indications and associations as types of classes, and references as types of properties.

The CIM schema provides the actual model descriptions. The CIM schema supplies a set of classes with properties and associations that provide a well-understood conceptual framework within which it is possible to organize the available information about the managed environment.

The CIM schema itself is structured into three distinct layers:

- Core schema — an information model that captures notions applicable to all areas of management
- Common schema — an information model that captures notions common to particular management areas but independent of a particular technology or implementation. The common areas are systems, applications, databases, networks, and devices. The information model is specific enough to provide a basis for the development of management applications. This schema provides a set of base classes for extension into the area of technology specific schemas. There are currently four common schemas:
 1. Systems
 2. Applications
 3. Networks (LAN)
 4. Devices
- Extension schemas — Represent technology-specific extensions of the common schema. These schemas are specific to environments, such as operating systems (for example, UNIX or Microsoft Windows).

More schemas are planned for definition in the areas of directory enabled networks (DEN), service level agreements, and distributed application transaction measurement (DATM). Others will follow.

The formal definition of the CIM schema is expressed in a managed object file (MOF) which is an ASCII file that can be used as input into a MOF editor or compiler for use in an application.

The unified modeling language (UML) is used to portray the structure of the schemas. Techniques to develop UML (VISIO) files from MOF files are being developed.

This is the first time in this industry that a common method of describing management information has been agreed upon and followed through with implementation. Other efforts have failed because of the lack of industry support. Because the model is implementation independent, it does not provide sufficient information for product development. It is the specific product areas — applications, system, network, database, and devices — and their product-specific extensions that produce workable solutions.

Status of current CIM-related projects: The CIM TDC is defining rules and categories for an information model that provides a common way to describe and share enterprise-wide management information. The group will define a meta schema or basic modeling language; a core schema, the base set of classes specific to systems, networks, and applications; and a common schema, a base set of platform-neutral, domain-specific extensions of core schema.

CIM will take advantage of emerging object-based management technologies and ensure that the new model can be populated by DMI and other management data suppliers including SNMP and CMIP. The CIM is being designed to enable implementations in multiple, object-based execution models such as CORBA and common objects model (COM), and object-based management technologies such as JaMAPI (Java Management API).