The Simple Times™

THE BI-MONTHLY NEWSLETTER OF SNMP TECHNOLOGY, COMMENT, AND EVENTS℠
VOLUME 1, NUMBER 4 SEPTEMBER/OCTOBER, 1992

The Simple Times is an openly-available publication devoted to the promotion of the Simple Network Management Protocol (SNMP). In each issue, ***The Simple Times*** presents: a refereed technical article, an industry comment, and several featured columns. In addition, some issues include brief announcements, summaries of recent publications, and an activities calendar. For information on submissions, see page 20.

In this Issue:

Technology and Commentary

Featured Columns

Miscellany

The Simple Times is available via both electronic-mail and hard-copy. For information on subscriptions, see page 20.

Technical Article

Vikram Saksena and Tibor Schonfeld,
AT&T Bell Laboratories

In this issue: *Customer Network Management of the InterSpan Frame Relay Service*

Frame relay is rapidly emerging as an attractive wide area networking solution for the efficient transport of bursty data traffic generated by LAN-based applications. Alternate solutions for wide area LAN interconnection have relied upon the use of leased private lines. While private line networks can offer very good performance for some applications, the highly bursty nature of inter-LAN traffic leads to low average utilization levels and inefficient use of dedicated wide area bandwidth. By using statistical multiplexing and streamlined protocol processing, frame relay can deliver the performance benefits of private line networks while achieving significant bandwidth efficiencies.

The adoption of frame relay into a corporate enterprise network requires good network management. One key barrier has been the use of proprietary management protocols by frame relay components — these do not integrate easily with the SNMP-based infrastructure that is being put in place to manage hosts, LANs and internetworking components, such as bridges and routers.

AT&T’s InterSpan Frame Relay Customer Network Management (CNM) Service addresses this fundamental management problem. The CNM Service provides customers with a management “window” into their portion of the InterSpan frame relay network from standards-compliant management stations located on their premises. The InterSpan CNM architecture combines the use of sophisticated management applications with partitioning and proxy agent capabilities to provide customers with SNMP network management and processed reports accessed via FTP. While SNMP access provides real-time management support, the FTP-based report services feature provides summarized and thresholded information that can be used for bandwidth engineering and tuning of application parameters for efficient use of the InterSpan Frame Relay Service.

The InterSpan Frame Relay Service

AT&T's InterSpan Frame Relay Service is an enhanced public frame relay service, that has been available to customers since January, 1992. Customers with standards-compliant frame relay DTEs, access the service at speeds that range from 56 kbps to T1. To provide service ubiquity, AT&T has established several access points across the country. Connectivity between customer locations is provided using *Permanent Virtual Circuits* (PVCs) and protocol conversion capabilities. (Protocol conversion aspects of the InterSpan Frame Relay Service are not germane to this article.) A PVC is a permanent communications path that is administratively established through the frame relay network. PVCs eliminate point-to-point leased lines because circuits are not physically dedicated to a single destination site. The PVCs are engineered to provide a throughput, called the *Committed Information Rate* (CIR), ranging from 32 kbps to 1024 kbps. Unlike fixed-rate private line connections, InterSpan frame relay PVCs give users increased flexibility, enabling them to send bursty data applications at the access line speed while achieving a throughput of CIR.

The InterSpan backbone network consists of StrataCom's IPX fast packet switches and protocol conversion devices that are densely interconnected by a mesh of T1 facilities (moving to T3) to provide a reliable and fault-tolerant network. Sophisticated rate control and bandwidth allocation mechanisms are used to support customer CIRs in a fair and efficient manner. Dynamic routing strategies employed by the nodes ensure that PVCs are rerouted rapidly in the event of internal failures with imperceptible effect on end-user performance. The backbone network is engineered to meet customer expectations of high performance and reliability.

The InterSpan Frame Relay Service Network Operations Center (FR-NOC) provides 7x24–hour support for the service and supports all activities related to operations, administration, maintenance, engineering, and customer network management. The IPX nodes are managed from the FR-NOC using the StrataView Plus element management system in conjunction with systems developed by AT&T Bell Laboratories for the FR-NOC operations. The StrataView Plus directly interfaces with the network nodes and collects fault and configuration information and traffic statistics. A relational database within the StrataView Plus system stores this information and makes it available to other systems for downstream processing. A graphical user interface on the StrataView Plus allows the FR-NOC operators to view network topology, configuration, and fault information, and to perform administrative tasks related to the management of the IPX nodes.

The InterSpan CNM Architecture

The goal of InterSpan Customer Network Management is to provide customers management information concerning their portion of the InterSpan frame relay network. In this model, the objects that are associated with a particular customer are the frame relay ports that terminate the customer's access lines and PVCs that provide connectivity between the customer's locations. The backbone network is a single entity as far as the customer is concerned and its presence is only perceived through the effect it has on the performance of the PVCs.

Customer needs for network management span a broad spectrum. At a very basic level, customers need real-time access to fault and performance information that allows them to do problem diagnosis, isolation, and resolution. Customer network managers use real-time network status information to deal with end-user trouble-reporting issues, e.g., slow response time, host unreachability, and so on. On a longer-term basis, customers need access to processed information that allows them to do pro-active capacity management and planning. Trended and thresholded reports on a periodic basis (weekly/monthly) identify for customers potential capacity exhaust situations before they become service-effecting. Based on this processed information, InterSpan Frame Relay Service customers can take planned action for capacity management rather than being in a continually reactive mode. AT&T's CNM architecture is rich and flexible enough to serve such a broad spectrum of customer needs.

There are two main components of the CNM architecture: the repository of CNM information and the customer premises workstation. The customer premises workstation initiates requests for transferring network management information. The repository of customer network management information, henceforth referred to as the CNM server, resides in the FR-NOC. The customer's applications access the CNM server through a frame relay PVC established between the customer location and the FR-NOC.

The CNM server supports the following capabilities:

- an SNMP agent that supports the Internet-standard MIB-II and the InterSpan Frame Relay Service MIB for access to current traffic, performance and configuration information stored in a partitioned database; and
- FTP access to processed management reports stored in a partitioned file system.

Partitioning allows for individual customers to view management information on their portion of the InterSpan Frame Relay Network in a manner that ensures privacy.

Management Information

SNMP access is provided to customers for interactively monitoring information useful for day-to-day problem management. SNMP support for the InterSpan Frame Relay Service allows customers to use standard SNMP management systems to collect and process data specific to their use of the InterSpan Frame Relay Service and integrate it with the data collected from other components of the enterprise network. The SNMP management framework described here is a first for frame relay networks.

Two MIB modules are supported by the CNM Service: MIB-II and the Interspan Frame Relay Service MIB. From the former MIB module, the system, interfaces, and snmp groups are supported. These are well-known groups and require no explanation. However, although it is meaningful to speak of interfaces with respect to a frame relay DTE, the concept of an interface does not relate to a frame relay service. As such, support for the interfaces group of MIB-II is minimal. (`ifNumber` is zero-valued.)

The second MIB module, the InterSpan Frame Relay Service MIB, is defined in AT&T's enterprise-specific space, and contains three groups: the *Administrative* group, the *Statistics* group, and, the *Calculated* group.

The Administrative group provides static configuration information about a customer's service parameters. This includes the port and PVC configuration information.

The Statistics group presents raw data on the usage and performance of customer ports and PVCs. While port statistics are maintained on a local basis, PVC statistics are maintained on an end-to-end basis that capture the network effects. The information provided by the Statistics group includes frame discards, frames transmitted and received, and frames sent with the congestion indicators set.

The Calculated group presents processed information on the usage and performance of customer ports and PVCs. It includes a statistical analysis of port utilization, PVC usage-to-CIR ratios, frame rates, frames with CRC errors, frame discards, and frames sent with congestion bits set. All variables in the Statistics group and the Calculated group contain information on a rolling 24-hour basis.

SNMP Agent

The SNMP agent consists of a front-end process, the *Proxy Multiplexer* (PMUX), that receives requests from all customers and forwards them to one of several processes that parse the requests, retrieve information, and form the responses. The responses are sent back to the front-end process which then forwards them to the requesting customer. The back-end processes are called the *proxy agents*. The functionality split between PMUX and the proxy agents ensures proper load-sharing and fairness in dealing with customer requests.

The PMUX adds two fields to an SNMP request received from a customer before forwarding it to a proxy agent. One field contains the source IP-address of the request and the other contains the source UDP-port. The method of communication between the PMUX and the proxy agent is UDP-based.

The PMUX needs to know the UDP socket addresses of the proxy agents to forward requests to them. This is accomplished by having each proxy agent register itself with the PMUX when the proxy agent first starts running. The PMUX fans out customer requests to the proxy agents based on source IP-address information. In turn, each proxy agent performs all the tasks of a distinct SNMP agent, and then sends back an SNMP response to the PMUX.

Since the functions of the SNMP agent are spread among multiple processes, AT&T Bell Laboratories developed a *Process Management System* (PMS) for managing and monitoring the PMUX and proxy agents, and reconfiguring the PMUX when necessary. When customers are added or deleted, or when a customer's configuration changes, the PMS initiates reconfiguration of the tables in the PMUX, creating new proxy agents while terminating others, as needed.

At start-up, the assignment of a customer to a proxy agent is made by the PMUX which then passes the Customer-ID to that proxy agent. Once the proxy agent is initialized, it waits on a UDP-port for an incoming SNMP request. Once it receives a request, it validates the request. If the request passes the authentication test, then the proxy agent retrieves the customer profile associated with the community string and continues processing. Otherwise, the proxy agent discards the request and logs an appropriate security violation message.

The next step taken by the proxy agent is to parse the PDU portion of the SNMP request. The proxy agent examines each element in the variable-bindings list and attempts to match its OBJECT IDENTIFIER (OID) against a list of OIDs known to the agent. The retrieval is enacted by calling a function associated with

MIB objects. The OIDs have two parts: the first part identifies the object type, while the second part identifies the specific instance of that object.

The retrieval part of the proxy agent is designed such that four functions are defined for retrieving MIB object values. Each of these functions extracts the instance-identifiers from the end of the OID and stores them in memory. After extracting the instance-identifiers, each function then makes a database query for the requested information. This process is repeated for each variable-binding. Once all the objects have been successfully retrieved, the proxy agent builds a response packet, encodes it and sends it back to the PMUX, and then waits for another request.

Processed Management Reports

In addition to SNMP management services, customers are allowed access to summarized management information contained in a variety of reports. These reports can be used for pro-active engineering and capacity management functions related to their use of the InterSpan Frame Relay Service. Given the bursty nature of inter-LAN traffic, managing capacity and performance issues over wide area connections is important to users. In the context of frame relay networks, the proper selection of CIRs, port speeds, host protocol window sizes and timer settings is important for optimizing an application's end-to-end performance.

The FTP-based report services rely upon the use of systems developed by AT&T Bell Laboratories for measuring, reporting, analyzing and troubleshooting network performance and capacity problems. These systems are used internally by the InterSpan network engineers to pro-actively plan network growth and respond promptly to changing traffic and new end-user needs. Through the FTP-based report services, results generated by these systems are made available to customers for engineering of their own network configuration. Parameters that capture the service's usage, performance, and configuration from a customer's perspective are reported on a weekly and monthly basis. These include port utilization, frames transmitted and received, usage-to-CIR ratio for PVCs, frame/byte discard rates, frames with CRC errors, frames sent with congestion bits set, and the customer's network configuration. This information is provided through several types of ASCII and graphical reports.

Customer Premises Workstation Applications

Customers can access CNM services from any management workstation that supports the following minimal set of capabilities:

- an SNMP-based NMS and an FTP client; and
- support of MIB-II and the ability to incorporate the InterSpan Frame Relay Service MIB.

The FTP application is routinely supported on any UNIX workstation. Therefore, it is straightforward for customers to access the CNM server and transfer reports that are of interest to them. These reports can be locally printed. A PostScript printer is required for printing the graphical reports.

For SNMP access, support of MIB-II and the ability to incorporate vendor-specific MIBs are routinely supported by most SNMP management stations. However, applications support for processing and display of MIB information varies significantly from vendor to vendor, and the notion of managing a shared public service, rather than devices, is relatively new. As such, the InterSpan Frame Relay Service MIB is sufficiently different in its structure and scope from traditional MIB modules, that changes and enhancements to existing applications may be warranted.

AT&T has developed enhanced premises applications on the NCR StarSentry management platform for both the SNMP management services and the FTP-based report services components of the CNM Service. These applications, collectively called the Frame Relay Data Manager, provide a graphical view of the customer's network topology and configuration information as it relates to the InterSpan Frame Relay Service. The topology display can be created automatically by examining and retrieving objects in the Administrative group. Status information on customer-specific objects, e.g., ports and PVCs, can be easily obtained via simple point-and-click operations. Traffic and performance information available from the Statistics group and the Calculated group can be polled and graphically displayed. Customers have the option of setting up the polling mechanism and automatically updating this information on their workstation.

Current Status

The Customer Network Management Service for AT&T's InterSpan Frame Relay Service is operational and the response has been very positive. Work continues on partnering with additional management platform vendors to integrate the AT&T InterSpan Frame Relay Service MIB and provide premises applications support on other management platforms. Extensions to this service are being considered through customer feedback. The realization of the CNM architecture has been made possible due to many team efforts; we thank them all for their valuable contribution to this project.

Industry Comment

Marshall T. Rose

Welcome to the fourth issue of ***The Simple Times***.

In this issue, the *Industry Comment* discusses one of the dangers facing the network management market. But first, an administrative announcement.

The Simple Times is available in a new format: *richtext*. Richtext is a simple page description language which is defined in RFC 1341 (the MIME RFC). Depending on the capabilities of your output device, with a richtext previewer you can see text presented in different typefaces, styles, and sizes. Richtext is text-oriented, rather than image-oriented, so although richtext is not nearly as expressive as PostScript, it is a lot easier to use text-handling tools on a richtext document. For example, on a UNIX system, you can use programs such as `grep` to search for a particular string.

At present, 88% of subscribers to ***The Simple Times*** receive the PostScript edition, which requires that the New Century Schoolbook and Courier typefaces are available. At the other end of the spectrum, the MIME edition, which is primarily ASCII with some modest structuring information, can be reasonably viewed on a dumb terminal or line printer. In contrast, the richtext edition is something of an intermediary between the PostScript and MIME editions: although it is textual in nature, it requires both a MIME-compliant mail-reader and a richtext previewer. There are many different kinds of previewers; some run on dumb terminals, whilst others use the *X Window System*. In fact, this is one of the great features about richtext — it never needs more capabilities than those supplied by your output device! In contrast, to render PostScript, you need a bit-mapped device, whilst having a fancy display or printer will never make ASCII look any better.

In order to encourage subscriptions to this new format, the automatic subscription instructions now contain information on various implementations of MIME-compliant user agents and richtext previewers. I'm interested in listing other implementations (regardless of whether they're openly-available). So, if you know of other software packages for MIME — particularly those that include a richtext previewer, please drop me a note. (Contact information is at the end of this issue.)

The Danger of Dreams

The last five years have seen a lot of interest in so-called *standardized management*. Some of this interest has resulted in products, and some of these products have resulted in solutions. This reflects the natural evolution of a market: a technology is defined, implemented, marketed, and undergoes constant incremental revision.

For our purposes, that technology is the Internet-standard Network Management Framework, best known as SNMP-based management. In addition to providing some solutions, the introduction of standardized management has provided hope to many that the industry can find automated solutions to the network management problem. However, this hope must be tempered with a determined pragmatism, lest we fall victim to the danger of dreams.

Dreams sell well. They make excellent dogma and even better ad-copy. In times of great stress, they are particularly seductive. Dreams are the stuff of great marketing opportunity. However, dreams, like controlled substances, may be harmful. Although vision is important to progress, *dreaming* often gets in the way of *doing*. In our case, the computer-communications industry spends a lot more time dreaming than coding.

Dreams don't alway come true, however. For example, Einar A. Stefferud has masterfully pointed out that:

> "OSI is a beautiful dream, and TCP/IP is living it."

What he means, of course, is that while everyone supports the concept of open systems, it takes engineering discipline and numerous production-quality implementations in order to realize a competitive, robust open systems market.

One might further observe that the "US GOSIP dream" has garnered a lot of hype, many problems, few products, and little credibility. US GOSIP is *checklist procurement* at its finest: a US federal agency specifies GOSIP and TCP/IP, makes sure that each bid has a GOSIP component (possibly to be supplied later), and then proceeds to make technical evaluations based on the TCP/IP products offered. Of course, including GOSIP components (which will probably never be taken out of shrink-wrap) does drive up the price, But that's what happens when dreams meet reality. Unfortunately, the US government never learned the lesson of US GOSIP, so it seems likely that US GOSIP's successor, IGOSS, will meet the same fate.

The network management area is not immune from dreaming — far from it. Originally, in the network management area, there was the CMIP/CMOT/CMOL dream, which history has discredited. CMIP and its ilk are a sad testament to a tragically skewed thrust/payload ratio. Today, the latest dream is called GNMP, released last month by the US NIST and heartily embraced by the NM/Forum. In terms of its disregard for cost-effectiveness and feasibility, GNMP is the king of dreams. To sum up: GNMP is CMIP, with FTAM for bulk transfer of information, MHS for store-and-forward transfer of information, and VT for remote terminal access to information, along with a few other kitchen sinks. (I'm not making this up; GNMP is some *serious* dreaming.)

The first public announcement for GNMP was in April 1991. When the presenter was asked about time-frames, the response was that GNMP was formulated to meet the *immediate* needs of US federal agencies. When asked about availability of GNMP-related products, the presenter was confident that industry would rush to meet this challenge to provide cost-effective solutions, even though there were no GNMP-like products available. When asked about SNMP, the response was that this was for further study. As it turned out, Dr. SNMP was in the audience asking these questions. (The editor was also in the audience, but was laughing so hard that he couldn't catch his breath to ask the same questions.)

After receiving these amazingly naive answers, the good Doctor related the following story: Operation Desert Storm had just ended and there were lessons to be learned. In particular, it was rumored that the military was looking to procure a new weapon: one that could sink a ship, shoot down an airplane, stop a tank, be carried by a single infantry man, be edible (in case the soldier ran out of rations), and be cheap to procure. Although such a weapon does not exist today, Dr. SNMP was confident that industry would rush to meet this challenge to provide cost-effective solutions.

Dr. SNMP was reminding us that:

> "The problems of the real world are remarkably resilient to administrative fiat."

That is, it's one thing to dream. It is another to do. It is one thing to base one's dreams on a proven technology (SNMP), and quite another to pay lip-service to things that work whilst acting like a child in a candy store, pinning one's hopes on a largely-discredited paper standard (CMIP).

The whole situation is reminiscent of an amusing exchange in which someone was boasting about his department's large budget for travel to network management standards meetings. We are talking *serious* money here. In response, Martin L. Schoffstall, a co-inventor of the SNMP, observed that the original research grant for SNMP was in the amount of US$10K. Marty speculated that if the money spent on producing standards were instead to be spent on developing workable network management technology, then we'd probably be able to manage every molecule in the universe. This is a sobering thought: dreams are seductive, but without a pragmatic perspective, they exact a terrible price.

My hope is that we open our eyes before these dreams turn into nightmares.

Applications and Directions

Steven L. Waldbusser

In this issue: *Applications stand to benefit from SMP*

The Simple Management Protocol (SMP) Framework promises to add many features to SNMP and also to correct some deficiencies. In addition to the better known ways in which SMP extends the SNMP framework, there are some interesting ways in which SMP makes life easier for network management application developers. SMP should make applications smaller, faster, and less complex. It is particularly important to make application development easier due to the critical need that network managers have for better applications. In this issue, we'll find out how SMP helps.

The Awesome Get-Bulk Operator

Although the "awesome" get-bulk operator is highlighted for its blazing speed, it also has an effect of reducing the size of SMP applications. The get-bulk operation fills up a packet with variable-bindings (in lexicographical order) until the packet is full. There is no danger of returning a `tooBig` error with the get-bulk operator. This obviates the need for code in the application which, upon receipt of such a reply, splits up the request packet into smaller chunks. In addition, there is no longer any need to have the management station dynamically discover (by trial and error) the optimal number of variable-bindings to put in the request packet to get the most data in the reply packet, without going over the threshold and getting back a reply with a `tooBig` error and no data.

MIB modules with large tables, such as the RMON MIB, sometimes provide mechanisms for bulk downloading of specific tables, in addition to the usual approaches. When applications are written for these MIBs, they often are more complex because they have to support the additional MIB objects that provide faster retrieval. The speed improvements of the get-bulk operator make it unnecessary for MIB modules and the corresponding applications to have special case support for bulk data retrieval.

Exceptional Responses

Another improvement with SMP is that some situations no longer cause the rejection of the whole packet. The most significant case is when a variable-binding in a get operation does not refer to a variable that exists in the relevant MIB view. With SNMP, this would result in a response being sent back with a `nosuchName` error indicating the variable-binding in question. Instead, with SMP, an exception is returned for that variable-binding, but valid data is returned for other bindings in the packet. Because such situations don't happen often, this has negligible effects on performance, but it has significant benefits in the form of reduced application complexity.

With SNMP, if an application were to be robust enough to withstand `noSuchName` errors (typically due to unimplemented objects), it would have to implement a state machine that would respond to an error by taking the offending variable-binding out of the request and re-sending it. When a successful reply was received, it would then have to match up the responses with the initial requests. SMP does not need the complexity of the state machine that would drive this multi-packet transaction. In fact, due to the complexity of this code in SNMP, many management stations did not implement such an algorithm and would not interoperate with agents that had unimplemented variables. This gave rise to the practice among agent vendors of returning an arbitrary value for objects that were not implemented, rather than the correct response of the `noSuchName` error. This practice gave misleading data to network managers, potentially causing confusion and mistakes. SMP's use of exceptions promises to eliminate the need for agent developers to choose between interoperability and correctness.

Set Requests

Another area of significance for applications is SMP's set operation. New error codes have been defined so that a management station can easily pinpoint the cause of a failed request and take direct action to solve the problem. With new errors such as `wrongType`, `wrongLength`, `inconsistentValue` and `resourceUnavailable`, an SMP application can take the appropriate action to solve the problem or alert the user to the specific cause of the problem. Further, the application can now determine if the error reflects a transient or a permanent condition.

In addition, the row creation mechanism allows applications to learn the agent's notion of appropriate default values for the new row and to use or modify those values. Similarly, the application can discover columns that are not implemented so that it may ignore them and continue to interoperate.

An often overlooked need, when remotely changing device configuration, is the capability of specifying whether to change the currently running configuration or the "bootup" configuration that will be loaded at the next restart of the system. SMP provides the `restartDomain`, which accesses a MIB view that contains values of managed objects to be used when the device restarts. Thus, by modifying the `ipAddrTable` in the `restartDomain`, one can change the IP address that the system will use when it next reboots, which can be accomplished without defining new MIB objects for those configuration objects that have already been defined. This avoids the need for standardization activity for data that has already been effectively modeled in MIBs, and allows our current configuration applications to be easily modified to administer the "bootup" configuration.

Systems and application management have extensive configuration management requirements and make heavy use of set operations. For example, the remote administration of users on a file server would make use of these improvements to the set operation. The above enhancements to set operations, coupled with the addition of SNMP Security, allow sets across the network and, therefore, system management to become more widespread.

Table Extensibility

An addition to the OBJECT-TYPE macro, the AUGMENTS clause also plays a substantial role for management stations. This clause allows a table to be specified as an extension of a previously defined table. With this mechanism, if a vendor specifies vendor-specific objects as extensions to a standard MIB, it will be easier for applications to automatically take advantage of the value-added objects. This will also promise to aid the standardization process as well — if a vendor can expect to get meaningful use out of vendor-specific objects across a variety of management stations, the vendor is likely to lessen the intensity of the fight for inclusion of those objects in a standard MIB. This will allow standard MIBs to remain simple and easy to implement.

In Closing

These features in SMP are not the ones that get the most attention; nonetheless, they are likely to play a significant role in improving applications in the future. SMP applications can be expected to be smaller and less complex, while being faster and more interoperable. SMP will pave the way for a new generation in network management technology — a generation in which useful applications are fielded and network management extends to system management.

Ask Dr. SNMP

Jeffrey D. Case

Dear *Dr. SNMP*,
In the May/June issue of ***The Simple Times***, you made a reference to research you have been conducting including comparative studies of SNMP and the OSI management protocol, CMIP. I am looking for CMIP implementations, especially for PCs under DOS. Can you help?
— *Misguided in Massachusetts*

Dear *Misguided in Massachusetts*,
Back on the farm, we have a saying:

> "You can't put ten pounds of manure in a five pound bag."

What this means is that, while we have worked to implement a CMIP agent for UNIX and manager station software for VAX/VMS and UNIX as a part of our research, we have not done any work to implement CMIP on resource limited platforms, such as DOS. We found that in order to conduct comparative studies of SNMP and CMIP, one really has to have an implementation of CMIP. Since none were readily available when we began this work, we set about creating them. We do not know of anybody who has completed a CMIP implementation for DOS, nor do we know of anyone who plans to attempt such an effort, which appears to be an impractical shoehorn job.

Dear *Dr. SNMP*,
I am trying to determine what the reasoning is that the get-next and get-bulk operators are around. Why was it necessary to use get-next for one item at a time, rather than pulling back either a row, the whole table, or whatever. Has anyone ever thought or discussed any other schemes?
— *Without a clue in West Newbury*

Dear *Without a clue in West Newbury*,
No, it's just something we threw together without any thought — **NOT**. The *powerful* get-next operator is one of the most often discussed issues of the Internet-Standard Network Management Framework. Much of the discussion of these issues has occurred on the SNMP mailing list and is available to those interested in the archives.

Back on the farm, we have a saying:

> "You can lead a horse to water but you can't make him drink."

What this means is that if you want to know how SNMP got to be the way that it is, you should read the rich history found in the archives.

Dear *Dr. SNMP*,
When the SNMP Research staff added the bathroom MIB extensions to the SMP agent, did they use SMUX?
— *Distended in Denver*

Dear *Distended in Denver*,
First, some background. Some of the programming staff down on the farm were recently frustrated by repeatedly finding the restroom (a one-holer) busy when they walked over to use it. Being network management experts, they decided to make it manageable via SMP/SNMP version 2. Consequently, there is now a magnetic switch of the burglar alarm variety on the restroom door. It is interfaced to a system via an RS-232 port.

To directly address your question, SMUX was not used to implement the bathroom MIB. The SMUX protocol and MIB may be used to construct modular agents via a master agent and potentially many subagents. Various subagents support different MIBs. For example, a SMUX subagent could be used to support the bathroom MIB, another to support MIB-2, and another to support other MIBs.

Back on the farm, we have a saying:

> "Don't ever try to teach a pig to sing, it wastes your time and it annoys the pig."

What this means is that Dr. SNMP is not particularly fond of the SMUX protocol and MIB, even though he was one of the eight who designed it one afternoon. SMUX has several shortcomings. First, it is resource intensive in several ways. It duplicates many functions of the master agent in each subagent. This leads to large programs, and our experience shows that programs with more lines of code tend to take longer to execute than smaller, simpler programs. It also requires the developers of subagents to be very skilled and knowledgeable about SNMP. Many subagent developers are specialists in FDDI, X.25, writing word processing applications, or using the bathroom and don't want to have to take the time to gain this expertise in order to implement their MIB objects. Second, it is infeasible to implement SMUX on many platforms, such as DOS. Consequently, SMUX is not able to provide a consistent management interface for applications developers across multiple platforms. Many applications packages are now available on both DOS and UNIX platforms, and they will become manageable more rapidly if such a consistent management interface is available.

For these reasons and others, Dr. SNMP agrees with Dr. Marshall T. Rose, who recently said:

> "SMUX was a successful experiment — the outcome of the experiment was that it was demonstrated that SMUX wasn't a good way

> to do things. The successful part is that the outcome was pretty clear.”

There are two ways to tackle this problem: a proxy approach and an API approach. The proxy approach is described in the new SNMP security RFCs. An implementation of this proposed standard replaced SMUX in the 4BSD/ISODE SMP implementation. The other approach is to use an application programming interface (API). This is what was used to implement agent support for the bathroom MIB.

The management station application is a modified version of the *X Window System* tool `xbiff` with a suitably designed “outhouse” icon. The door on the outhouse reflects the current state of the bathroom’s availability. Some manager stations also produce a squeaky door sound on the speaker as the door closes and a flushing sound as it opens.

Of course, the entire system uses SMP/SNMP version 2, as all new management applications should. The system utilizes SMP’s security and privacy, because privacy is important in bathroom applications.

The developers tell me they are looking forward to flushing out the remaining problems and bowling over visitors with demonstrations. Tanks for your question.

Dear *Dr. SNMP*,
I see lots of marketing literature for products which claim compliance with various MIBs. Do any really comply?
— *Caveat Emptor in Caledonia*

Dear *Caveat Emptor in Caledonia*,
Dr. SNMP shares your concern about truth in advertising. All too often it seems that the only difference between a network management salesman and a used car salesman is that the used car salesman knows he’s lying. The only exception, of course, is network management sales personnel who read ***The Simple Times***.

Back on the farm, we have a saying:

> “They’re as rare as hen’s teeth."

What this means is that there are a few agents which are fully compliant with the relevant specifications. It turns out that it is often very difficult to fully implement a MIB without making modifications to the source code of the protocol stack in order to add the relevant MIB instrumentation. This is difficult, if not impossible, in some situations.

The SMP/SNMP version 2 Structure of Management Information (SMI) has addressed this problem by defining a new ASN.1 macro, called AGENT-CAPABILITIES. This macro can be used to describe an agent’s implementation of a particular set of MIB variables. It provides a concise method for detailing how well an agent supports the MIB modules that it claims to implement. It, therefore, can show what object groups are not implemented, which objects are not implemented within groups which are implemented, ranges of supported values, and details about row creation requirements. This technology is a refinement of similar reports first described in RFC 1303. Very few agents have AGENT-CAPABILITIES reports available today, but many will in the future. Yours is one of the many problems addressed by the new SMP/SNMP version 2.

Security and Protocols

Keith McCloghrie

Other columns in this issue of ***The Simple Times*** describe some of the benefits found in the Simple Management Protocol (SMP) Framework, which is to be the basis of SNMP version 2. In this column, we’ll discuss the details of the SMP’s new and updated features. In this article, we’ll start by looking at the changes in the way management information is defined. In subsequent articles, we’ll look at: the changes in the protocol data units (PDUs), including the addition of the “awesome” get-bulk and the inform operations, the new ways to specify compliance and conformance, and the new SMP-specific MIBs.

Object Syntaxes

The SMP extends the set of data types available for defining management information. One new type is Counter64, a 64–bit counter for use when 32–bit counters overflow more rapidly than they can be retrieved by a management station. Such counters, although more complex to implement, are becoming necessary in today’s environment of higher-bandwidth networking. However, SMP limits the use of Counter64 to objects for which the extra complexity is needed, i.e., counters for which a 32–bit quantity could wrap in less than one hour. Another new type is the enumerated form of the ASN.1 BIT STRING, for representing a data item as a set of bits, each with a different meaning, and where any number of the bits may be set. In SNMP, such data items could only be represented via a clumsy definition using INTEGERs. NsapAddress is a new data type for representing the internetwork-layer addresses used in OSI networks. All other SNMP data types are retained, although a loophole in the definition of signed integers is closed by specifically restricting them to be, at most, 31 bits of data plus a sign bit. The Opaque type is retained solely for backward-compatibility, with its use in new MIBs being prohibited.

Object Types

In SNMP, MIB objects are defined using the OBJECT-TYPE macro. The SMP extensions to this macro include several new clauses.

A new UNITS clause is available to define a textual label that can be applied to the object's value to make a management station's user-interface more user-friendly.

A new NUM-ENTRIES clause is available for use with MIB tables to specify the name of another object whose value indicates how many entries are contained in the table. Experience in defining MIBs has shown that objects indicating the number of entries in a particular table are often defined. This clause provides information on this relationship to a management station which may be useful, especially when using the new get-bulk operation to retrieve the contents of an entire table.

Another new clause, AUGMENTS, recognizes the fact that management information from several MIBs is often related even though the definitions may be split into several documents, and thus their names (OBJECT IDENTIFIERs) may be unrelated. In particular, a well-established standard MIB may contain a table for which additional columns need to be defined in a new MIB, either a potential future standard or an enterprise-specific MIB. The AUGMENTS clause records the fact that the new columns logically exist as additional columns in some other MIB table, even though they are not defined in the same document. In particular, the AUGMENTS clause replaces the INDEX clause to indicate that the conceptual rows in the new table are identified in the same way as (and exist under the same circumstances as) the conceptual rows in the original table.

The ACCESS clause has been renamed to MAX-ACCESS in order to clarify that it specifies the maximum access which makes "protocol sense", and the values are ordered, from least to greatest, as follows: "not-accessible", "read-only", "read-write", "read-create". The "read-create" value is used for write-able objects in a conceptual row for which new instances can be created via network management. Another change is the recommended use of "not-accessible" for auxiliary objects (those objects defined in a table solely for use in identifying a conceptual row).

There are two other OBJECT-TYPE changes of note: the use of INDEX (or AUGMENTS) and DESCRIPTION are now mandatory (they were only optional previously so that the now obsolete MIB-I was not prematurely made non-compliant); and, optionality is no longer specified in the STATUS clause. Instead, SMP provides the more powerful alternatives of object groups and module compliance definitions.

Object Groups

Object groups specify the basic units of conformance. In SNMP, groups were defined only through ASN.1 commentary text, and were constrained to be aligned with the naming structure of the MIB (i.e., its OBJECT IDENTIFIER subtree hierarchy). SMP provides a new macro, OBJECT-GROUP, to define groups in a more formal manner, and to allow the contained objects to be specified explicitly, no matter where they are positioned within the naming structure of a MIB. We'll save discussion of module compliance for a future article.

Textual Conventions

In addition to the data types defined in the SMI, SMP allows the definition of *textual conventions*, where a name is given to a specific data type, possibly with some restriction on range/size, together with a particular meaning of the value. These names are then available for use as additional data types in MIB definitions, although they retain the encoding (in PDUs) of their underlying data type. SMP provides a formal means of defining these via the new TEXTUAL-CONVENTION macro.

A useful part of the macro is the DISPLAY-HINTS clause which provides a hint as to how the value may be displayed by a management station, in accordance with the particular meaning of the textual convention. The hint for a MacAddress, for example, suggests that each octet should be printed in hexadecimal, separated by a colon.

Several of the generally useful textual conventions defined in Internet-standard MIBs are included in the SMP specifications, including DisplayString, PhysAddress, MacAddress, TruthValue, AutonomousType, and InstancePointer. Of the new textual conventions, two are noteworthy: TestAndIncr and RowStatus.

The TestAndIncr textual convention provides for atomic, or sequenced, operations. When an object defined with this syntax is contained in an SMP set operation, the new value must match the current value; after a successful set, the value increments. Such objects can be used as local and/or global advisory locks on data structures to arbitrate between interactions from multiple management applications; potentially on multiple distributed management stations. They can also be used to provide assurance that multiple concurrent set operations are only actioned in the desired order, even if re-ordered in transmission.

The RowStatus textual convention is an evolution of a convention first used in the RMON MIB for row creation, modification, and deletion. One of its values provides a means for a management station to create new instances in a conceptual row, guaranteeing that the operation will succeed only if the instances do not already exist. Other values allow a conceptual row to be taken out of use while it is modified, and allow conceptual rows to be deleted.

Compatibility

As can be seen, SMP provides an evolution in the capabilities available to define management information in MIBs, which pays particular attention to backward-compatibility. This evolution is such that all current SNMP MIBs, while not compliant to the SMP SMI, are readily useable with SMP, with no changes in object-specific implementations in agents or management stations being required. The few straight-forward rules for upgrading an existing MIB to compliance with the SMP SMI are detailed as part of the SMP specifications.

Standards

David T. Perkins

In July and August, the IP Forwarding Table MIB and the three documents defining SNMP security were published; in addition, at the July meeting of the IETF a proposal was presented as the next version of SNMP. In this proposal, called the Simple Management Protocol (SMP) Framework, were changes to the SNMP security documents. The result, unfortunately, is that the long process of defining security for SNMP was re-opened two weeks after it was thought to be completed. (In fairness however, the SNMP security documents had laid dormant for over six months.) As such, the just recently published RFCs now need to be re-evaluated to consider the changes made in the SMP proposal. This suggests that it would be unwise to ship product based on these SNMP security RFCs and, instead, to be prepared for changes.

Several items on the standards track are being considered for advancement. These include: the IEEE 802.4 Token Bus Interface Type MIB; the IEEE 802.5 Token Ring Interface Type MIB; the Extensions to the generic-interface MIB; and, the Ether-like Interface Type MIB.

In addition, the following items are under review for addition to the standards track: SNMP over AppleTalk, SNMP over OSI, and the IEEE 802.3 Repeater MIB.

Recently Published RFCs

RFC 1321 - The MD5 Message Digest Algorithm (Informational)

This document describes the algorithm used for computing a “fingerprint” or *message digest* used by SNMP Security. The message digest is used to ensure that a message has not been modified. The algorithm is a more “conservative” design of an earlier algorithm called MD4 that was being considered for use in SNMP security.

RFC 1351 - SNMP Administrative Model (Proposed Standard)

In this RFC, the SNMP administrative model (as defined in RFC 1157, the SNMP protocol specification) is expanded and clarified for a new security approach that includes authentication and message-integrity, privacy, access control, and proxies. The new model uses distinct identifiers for peers exchanging SNMP messages instead of a common identifier or community string approach used in the SNMP specification. The document introduces a new term called *Party*. Security attributes are associated with a Party. This document gives an overview of the information contained in the Security Protocols and Party MIB documents.

RFC 1352 - SNMP Security Protocols (Proposed Standard)

The details of the SNMP security protocol are specified in this document. The security threats and the goals along with constraints of the proposed solution are first presented. This is followed by the detailed exposition of SNMP security, including Parties, the digest authentication protocol, how to generate a message, how to receive a message, clock and secret distribution, and usage considerations.

RFC 1353 - SNMP Party MIB (Proposed Standard)

This document defines the MIB objects needed for implementation of SNMP security. These objects are organized into four tables. They are: the *Party Public* database, the *Party Secrets* database, the *Access Control* database, and the *View* database. The Party Public database contains addressing information, public key information, party clocks, and party maximum message sizes. The Party Secrets database contains private authentication and encryption keys. The Access Control and View databases together determine the permitted operations to collections of MIB objects.

RFC 1354 - IP Forwarding Table MIB (Proposed Standard)

This MIB defines a replacement for MIB-II’s routing table, `ipRouteTable`. The change fixes the indexing problem found in the routing table as well as adding the next hop Autonomous System number and policy information. The result is that multiple paths to a destination can now be managed.

The IEEE Standards Process

The Institute of Electrical and Electronics Engineers, Inc. (IEEE) is a trans-national organization whose standardization activities include local area network technologies. These include: ethernet-like LANs (IEEE 802.3); token ring (IEEE 802.5); token bus (IEEE 802.4); and, bridging (IEEE 802.1d), which are the base documents for IETF-developed MIBs. A previous issue of this column briefly described the process to develop IETF standards. In this issue, we look at the IEEE standards process with emphasis on the 802 group.

In general, the IEEE standards process is characterized by many forms and paperwork that must be completed. (This is due to a combination of two concerns: first, to eliminate possible problems with antitrust violations; and, second, to smooth the transition of standards from IEEE to ANSI, ISO, and other standards organizations.) The output from the process is a standard that has a lifetime of five years. At the end of five years it must be reviewed for revision or reaffirmation.

Pictorially, the process is:

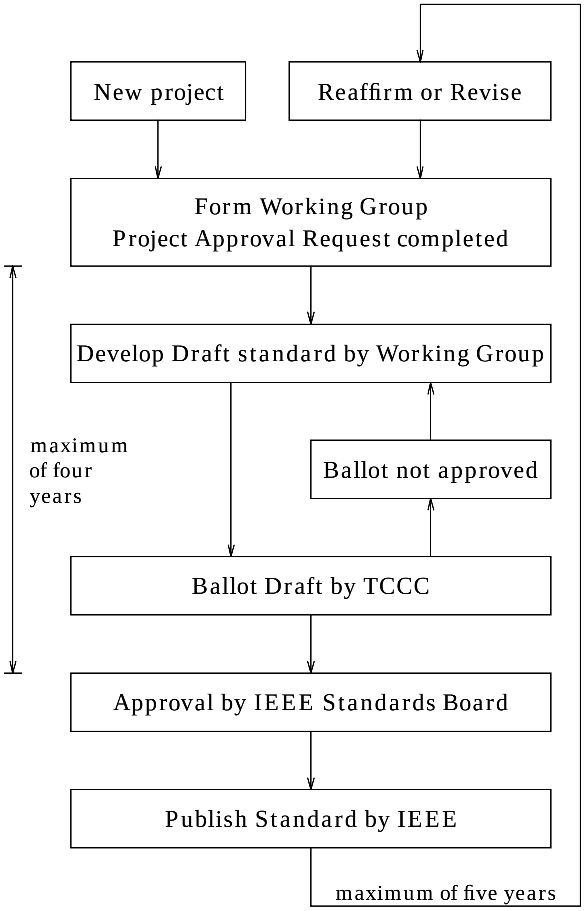

The work on a standard begins by identifying the forum in which it should take place. There are a number of standing committees in IEEE. If there isn't an appropriate existing committee, then one must be formed and its sponsor identified. A *Working Group* (WG) is then organized to decide on the scope of the standard. The first output from the WG is a formal (and legal) document called a *Project Authorization Request* (PAR) which describes:

- the type of project (a standard, guide, or recommended practice);
- indication of intent for project to form the basis of an international standard;
- project title;
- scope of the project;
- purpose of the project;
- name of working group;
- a list of other groups to review and coordinate with the output from the project; and,
- legal odds-and-ends such as copyright statements.

After approval of the PAR, the working group is expected to write a draft of the standard. Projects (i.e., working groups) may run for up to four years. Drafts are written and voted on by voting members of a working group. For WG members to gain and retain voting privileges, the rules for IEEE 802 projects require them to attend at least 75% of the sessions at any two of the last four most recent 802 meetings; state their intent to be a voting member; and, respond to at least one out of the last three WG ballots. Votes at 802 WG meetings require a 75% approval of the total of the APPROVE or DISAPPROVE votes to be carried. Votes by letter ballot (i.e., a letter sent via the Postal Service) require at least 50% to be returned with at least 75% approval of the APPROVE or DISAPPROVE votes to be carried. (The ABSTAIN votes count in the 50% return requirement, but not in the 75% approval requirement.) There will often be several draft versions before a WG decides to send out a letter ballot. If the subject matter is contentious, it may take more than one letter ballot to achieve approval.

The next level of approval is through the sponsor ballot. In 802, a letter of interest and intent to vote is sent to qualified members. This is done to improve the probability that the required percentage of ballots will be returned. The approved draft from the WG is sent out for review to the sponsor group members whose interest has been verified. In 802 this is called "going out for TCCC ballot". (The TCCC is the *Technical Committee for Computer Communications* of IEEE.) For the ballot to be completed (closed), at least 75% of the ballots must be returned. The ballots must indicate ABSTAIN, APPROVE with optional comments, or DISAPPROVE with required "specific reasons in sufficient detail so that the specific wording of the changes that will cause the vote to be changed to APPROVE can be readily determined." For all DISAPPROVE votes from the ballot, the WG must attempt to resolve them.

The revised draft with the list of unresolved negative comments is sent out as a continuation ballot to the balloting group to check that the 75% approval is maintained after the changes. If so, the draft, results from the latest ballot, and the evidence of an attempt to resolve each outstanding negative comment is forwarded to IEEE *Standards Review Committee* (RevCom), who check to make sure that all procedures have been followed and that all paperwork is complete.

The final step is forwarding to IEEE Standards Board for approval and then publishing as an IEEE standard. After approval as an IEEE standard, the document may be forwarded to ANSI and other standards organization for adoption.

IEEE is an organization that has been in existence for over 100 years. The original philosophy was to create standards that codify existing practices. This has changed over the years so that now the 802 group leads in the definition of new technologies. This can (and does) result in "paper standards" that lead the development and deployment of the technology. This accelerates the creation of new standards, but can result in standards that are never implemented due to business or technical barriers.

The current process relies on the Postal Service for notification of business issues and distribution of drafts for voting. This is costly and slow, compared to electronic methods used by the IETF. On the other hand, the requirements that WG members stay current to be able to vote, and that document reviewers list problems with specific recommended changes with negative votes, seem like good additions to the IETF standards process. Additionally, the requirement that the WG show evidence that an effort was made to resolve all negative comments seems like a very constructive way to balance out the wishes of the majority with the concerns of the minorities in votes.

In summary, the IEEE standards process uses a different model than that used in the IETF. Both the models and processes would benefit from using some of the other group's ideas. Some of the friction that has resulted in the past could be eliminated with cross-pollination of culture from each organization.

This issue's column has highlighted the IEEE standards process. The next issue will focus on the process that will be used to evolve the SNMP framework.

Summary of Standards

Full Standards:

- 1155 - Structure of Management Information (SMI);
- 1157 - Simple Network Management Protocol (SNMP); and,
- 1213 - Management Information Base (MIB-II).

Draft Standards:

- 1212 - Concise MIB definitions.

Proposed Standards:

- 1229 - Extensions to the generic-interface MIB;
- 1230 - IEEE 802.4 Token Bus Interface Type MIB;
- 1231 - IEEE 802.5 Token Ring Interface Type MIB;
- 1232 - DS1 Interface Type MIB;
- 1233 - DS3 Interface Type MIB;
- 1239 - Reassignment of experimental MIBs to standard MIBs;
- 1243 - AppleTalk MIB;
- 1253 - OSPF version 2 MIB;
- 1269 - BGP version 3 MIB;
- 1271 - Remote LAN Monitoring MIB;
- 1284 - Ether-Like Interface Type MIB;
- 1285 - FDDI Interface Type MIB;
- 1286 - Bridge MIB;
- 1289 - DECnet phase IV MIB;
- 1304 - SMDS Interface Protocol (SIP) Interface Type MIB;
- 1315 - Frame Relay DTE Interface Type MIB;
- 1316 - Character Device MIB;
- 1317 - RS-232 Interface Type MIB;
- 1318 - Parallel Printer Interface Type MIB;
- 1351 - SNMP Administrative Model;
- 1352 - SNMP Security Protocols;
- 1353 - SNMP Party MIB; and,
- 1354 - SNMP IP Forwarding Table MIB.

Experimental:

- 1187 - Bulk table retrieval with the SNMP;
- 1224 - Techniques for managing asynchronously generated alerts;
- 1227 - SNMP MUX protocol and MIB;
- 1228 - SNMP Distributed Program Interface (SNMP-DPI);
- 1238 - CLNS MIB;
- 1283 - SNMP over OSI; and,
- 1298 - SNMP over IPX.

Informational:

- 1147 - A network management tool catalog;
- 1215 - A convention for defining traps for use with the SNMP;
- 1303 - A convention for describing SNMP-based agents; and,
- 1321 - MD5 message-digest algorithm.

Historical:

- 1156 - Management Information Base (MIB-I).

Working Group Synopses

Robert L. Stewart

The time span covered in this issue's column is about three months, due to a delayed publishing date in deference to summer vacations. During this 50% increase in the usual time, I have almost a 150% increase in the amount of discussion to synopsize. Nearly half of the traffic was on the SNMP mailing list. Speaking of which, we'll start with an announcement.

There is a new working group for SNMP version 2 formed:

> "to consider all technical contributions to the SNMP evolution process and to produce a single recommendation as to which contributions (or combinations or modifications thereof) should define the next generation SNMP network management framework."

The charter calls for the group to complete its work and submit its recommendation to the IESG after the March 1993 IETF meeting.

With this in mind, future columns will have separate synopses for SNMP general discussion on the previous mailing list and the new SNMP-2 Working Group.

SNMP General Discussion

Some specific questions on formatting SNMP requests and responses received a detailed answer and a reference to *The Simple Book*.

Multiple people sought ongoing MIB work for ISDN. Various responses pointed to Israel and CMU. A suggestion to start a working group and mailing list drew a volunteer to coordinate. There is now a mailing list, accessible through `imib-request@hobbit.netcs.com`

A question whether an index integer is 16 or 32 bits drew the architectural answer that it is a sub-identifier in an OBJECT IDENTIFIER (OID) and thus is potentially unbounded. There was also an implementor's response that some implementations had incorrectly limited them to 16 bits and that it should be 32 bits or more, without losing a bit for the sign.

The suggestion that a Network Management System (NMS) could determine if an object is read-only or read-write if MIBs had such an indicator for each object drew a reference to RFC 1303, which offers a way to describe specific agent implementations.

"It's better to use minimal IP/ARP" was the response to a suggestion to use the source IP and Ethernet addresses when formatting an SNMP response via UDP.

Concern arose over the Simple Management Protocol (SMP), first announced in the press as a proposal for SNMP version 2. The complaint that there was no time to read the documents and prepare for the BOF to discuss it in Boston got the reply that the BOF would include a presentation; meetings aren't the only way to make progress; and, this shouldn't be delayed until Fall. The SMP authors made eight documents available from private sources (later re-released as Internet Drafts), and also released four independent, interoperable implementations.

An inquiry for an SDLC MIB or interest in developing one received no response.

An inquiry for a gateway between SNMP for LANs and CMIP for management domains for the Network Management Forum's CONCERT architecture got the response that, although SNMP and CMIP both have managers and agents and both use ASN.1, they have little else in common, as their models for entity, information, naming, protocol operations, and transport are all different.

The question as to whether upcoming public domain SMP implementations will be based on SNMP security or SMP or both, brought the response that the 4BSD/ISODE package will be based on SMP and will proxy RFC 1157 SNMP. Among many concerns was one that it's scary for SMP to drop features of SNMP security as "unnecessary". The response was that the dropped function, protection from message reordering, had negative impact with insufficient justification and referred to reasons in the SMP introductory document. A later inquiry about SNMP security availability got the response that SNMP security is not separate but is an intrinsic part of SMP, which specifies a few variations.

Questions about the future plan for SMP included: is this the mailing list to discuss it; will it become SNMP version 2 with a proper working group; and concern that it would not receive proper review. Multiple responses indicated that this is the proper list for now; SMP was called that only because the authors did not want to presume the name SNMP; SMP is a proposal not a pronouncement, and complete review will be expected from the entire community in an open process.

Many technical questions about SMP included: why did 64-bit counters suddenly become acceptable; would unrestricted BIT STRINGs and removal of other ASN.1 restrictions cause major interoperability problems; surprise that ASN.1 remained; the complaint that RMON row creation formalisms lack experience; the complaint that the rule of a set acting “as if simultaneous” was weakened; the question if sets now had an undo; and the assertion that bulk transfer is a band-aid. Responses to those stated that 64–bit integers remain expensive and Counter64 is restricted to information that would wrap 32–bits in less than an hour; BIT STRINGs are limited to primitive encoding; restrictions on BER have been clarified and tightened; dropping ASN.1 would violate goal of evolution and preservation of investment; RMON interoperability test showed no problems with row creation nor did four SMP implementations; the “as if simultaneous” rule is not weakened; the undo in set processing is a recognition of reality; and the bulk transfer “band-aid” has racing stripes, showing an order of magnitude increase in speed of data retrieval.

Public domain SMP code was announced from ISODE with no privacy, but with instance-level access control, cursory recognition of all ASN.1 macros, and intense scrutiny of the OBJECT-TYPE macro. CMU announced a similar public domain package.

The subject of creating and deleting rows consumed many messages. Proponents of adding create and delete functions to SNMP pointed out that: CMIP has them; SMP’s textual conventions help but don’t go far enough; they are easy and cheap to implement and ease interoperability with NMSs; we should encourage discussion and analysis of benefits, and approaches such as the “RMON polka” for finding a new index and controlling row state leave problems with complexity and collisions. Opponents of such an addition pointed out that: SNMP has considered creation and deletion as included since its inception; existing mechanisms are sufficient; every discussion among the SNMP and SMP authors has concluded that distinct functions are not worth the change as it adds complexity for all MIB objects; it requires more “method” routines for all object leaves, more PDU types, more error handling, and more SNMP parties for security. Some proponents were working on an experimental implementation to be reported later.

A question about migrating to the SNMP security proxy mechanism turned into a discussion of SMUX with some people insisting that the IETF should standardize mechanisms for integrating sub-agents because it is a big problem that needs a solution, and others insisting that this is an issue of implementation internals and application programming interfaces and is thus not in the traditional scope of IETF standards.

There was a suggestion that SMP get-bulk could be further improved by about 26%, with a scheme to compress OIDs. The response was that 26% was a modest return for considerable complexity, particularly when the get-bulk proposal improves performance by orders of magnitude with little increase in complexity.

The observation that agents and managers should use different UDP ports so they will not conflict on the same system received the response that this is not a problem: managers send to port 161, listen on what they specified for responses, and listen to port 162 for traps, while agents listen to port 161, respond as instructed, and send traps to port 162. Another message pointed out that one vendor’s manager sends from port 161, although the vendor should have known better.

The statement that SMP is clearly an improvement over SNMP security brought the response that the SNMP security working group was not chartered to make other improvements. A related question asked if it was necessary for SNMP security to have an incompatible format rather than reformatting the community field. The response was that such a change to the community string was considered but was a mess and didn’t provide needed features that come with the SNMP security message wrapper. Furthermore, SMP changes only the SNMP security inner PDU, not the wrapper.

A request for review of the SMP BOF at the Boston IETF meeting got the response the consensus was: SMP is a good technical proposal, the community doesn’t want two major changes (SNMP security and SMP), and thus desires SMP to advance very quickly. The BOF was over three hours, with a one and a half hour presentation from the SMP authors and the remainder discussion. The conclusion of the BOF was that SMP is to be submitted to the standards track as a proposal and a working group is to be formed, to work closely with the SNMP security working group to resolve differences.

Several people agreed that SMP seems to change gauges so they no longer lock permanently at their maximum value. A response speculated that they were not supposed to lock but were to hold there until the measured value came back down. The author agreed. An ensuing discussion over word definitions, acceptable interpretations, and reasonable implementations stopped just short of personal name calling and didn't reach a resolution other than the suggestion that leaving it unclear is not acceptable.

To the question if vendors can use MIBs from the enterprise repository on `venera.isi.edu` in their own products, one MIB owner granted permission and another person pointed out that the purpose of that repository is to make the MIBs available for network managers.

Several discussions got into high-level issues about the definition and goals of network management and whether SNMP is the right solution. Attempts to keep the discussion at a low level had little effect other than some bitterness over suppression of ideas. Those in favor of the low-level discussion agreed that SNMP is just part of an overall network management solution and insisted that trying to go much beyond it will flounder as it always has in the past. One such discussion was over the functions allowed to agents and what constitutes a manager. Another was regarding traps, their place in network management and whether they could or should be made reliable. A third considered auto-topology and auto-discovery. Yet another questioned the place of SNMP in managing systems. The volume, both in terms of quantity and voracity, was too much to reproduce here.

An inquiry for a mailing list for SNMP users rather than designers and implementors brought the response that this is the right list and such questions will get answers. This got the objection that the discussions have been too theoretical and a separate group is needed. A respondent suggested subscribing through `net-ops-request@decwrl.dec.com`, but bemoaned the lack of traffic there. The final statement was an encouragement to ask any question, with the warning that many of the questions do not have clear answers.

There were many more topics discussed, both long and short, but the space here does not permit even listing them.

SNMP-2 Working Group

The first order of business was a discussion of the need for meetings, an interim meeting and time at the Washington IETF. The suggestion of a meeting at the University of Tennessee at Knoxville (UTK) engendered discussion over whether the community would view this as part of an SMP "railroad job". No one indicated such, and as a result a meeting was scheduled for October 5-6 at UTK, and several meeting slots were requested for the Washington IETF.

The chairman stated a rumor that the SNMP-2 mailing list would not be open was entirely untrue.

A question about editor and rules for changes to SNMP-2 documents brought the chairman's response that the group would start with the SMP Internet Drafts (the only proposal received); that Marshall T. Rose (one of the SMP authors) would be editor; and, a preference not to start with strict rules, although all changes are to be discussed publicly.

Request Address:
`snmp2-request@thumper.bellcore.com`
Archive:
`pub/davin/snmp2-archive@faline.bellcore.com`

Character MIB WG

An inquiry for implementation experience for the RS-232 synchronous port table indicated that certain hardware doesn't report frames interrupted due to modem line status changes. The response was that perhaps this was not a good or useful counter and indicated that the definition of the object for the number of hardware signals qualifies for hardware capability and was intended to cover only useful information such as slower changing modem signals. Another person agreed that deprecating the counter may be in order, as well as clarifying the intent on modem signals in general. These messages will be kept as the basis for changes at consideration for promotion to Draft Standard.

A question about the value for `charPortHardware` when the hardware is RS-232 got the response that AutonomousType indicates the value is the OID for the entire RS-232 MIB subtree.

Another question, about the relationship between the `charPortAdminStatus` and `charPortOperStatus` objects, received the response that the former is advisory and setting it to "off" should result in the termination of connections, and the latter (which reflects reality) should go to "down".

Chassis MIB WG

Although little had happened in this group before and during the Boston IETF, interest and discussion escalated afterwards.

The question whether proxied entities with the same Internet address should share a MIB-II SNMP group or have their own brought the response that each proxy community (or party) should appear self-consistent and complete, and a top-level agent should count the total

SNMP traffic. The follow-up question whether a minimal implementation of MIB-II has only the system and SNMP groups got the response that this is not wrong, some areas will remain gray; the Chassis MIB won't help, and it comes down to implementor's choice.

There were two submissions of power supply MIBs as input for the group. The submittor of one liked the descriptor and health objects from the other but questioned the "exotic" N+1 boolean for indicating redundancy and asked if power supply slots were like other slots. The response was that power supply slots are typically different and that N+1 is similar to a "dormant" state but implies smooth failover and requires knowledge of power consumption. Another message indicated that power supplies would be in the MIB's slot table only if they occupied regular slots.

On the question of retaining community strings and IP addresses now that SNMP security is in place, the consensus was that they are needed for backward compatibility over a long transition period.

A question about the need to resolve *field replaceable units* (FRUs) in the Chassis MIB brought the response that vendor implementations vary too greatly.

A submission of an informal proposal for the Power Supply MIB brought some specific suggestions for improvements and additions and a general discussion of where power supply and environmental information belonged. The result was the decision to put them into the Chassis MIB as that was the limited intent of the charter. Appropriate ASN.1 is to be submitted as optional groups parallel to the Chassis MIB group.

A long discussion of an interface table to show the mapping from interfaces to backplane segments ultimately resulted in the decision to add a type and index to the configuration table, as other suggestions did not sufficiently consider the variations among such systems as routers, bridges, and repeaters.

The assertion that a chassis needs multiple agents for robustness brought the response that one is sufficient but multiple are allowed.

A message with several questions about correct implementation brought the realization that the model in the document had too many implementation-dependent assumptions and required clarification, which are to be added. An example was the attempt to define "sparse" and "dense" which concluded that "sparse" meant leaving out entries rather than using a wide-spaced numbering system.

DECnet MIB WG

The indication of a problem with the indexing of the adjacency table resulted in a long discussion about what the MIB meant, what was appropriate, and what was an acceptable change. The outcome was that semantics of a Proposed Standard MIB cannot be changed, therefore, the old table is to be marked obsolete and duplicated as a new table with different indexing, namely, a circuit index and node address, consistent with the rest of the MIB and DECnet itself.

The response to how do you zero counters for an individual circuit was that DECnet does not provide that capability, therefore, neither does the MIB.

A question about DECnet-style locking counters got the response that they are there only for DECnet compatibility and are not intended for implementations with SNMP-style wraparound counters.

The statement that the MIB is missing a way to handle multiple equal cost routes brought the reply that DECnet Phase IV also omits that information from its management interface, and a warning that circuits in DECnet and SNMP MIBs must always appear in the same order as the effects of order can extend all the way into the routing algorithms.

Domain Name Service WG

A new MIB draft was presented and accepted with a few comments in Boston. Edits were made and submitted as a new Internet Draft.

Ethernet MIB WG

A discussion over the objects that were marked as deprecated during the draft stages resulted in the decision to remove them completely for the Draft Standard.

An Internet Draft was submitted and recommended as a Draft Standard by the IESG.

FDDI MIB WG

Much discussion over SMT version 7.2 resulted in a decision at the Boston IETF to align with SMT rather than advance to Draft Standard, particularly considering that there are few implementations of the MIB.

A long discussion over a path configuration table considered the difference between manageability with a generic NMS and compatibility with ANSI standards and resulted in unconvinced proponents of each side but no change.

Host MIB WG

The group received several submissions of MIBs from various systems, including examples of one MIB's information from UNIX, IBM PC, and Apple Macintosh systems. Discussion on the list was mostly over detailed values down to the level of the number of pixels on a screen and path names to storage devices, with arguments over what was necessary, useful, and available, including an argument over values read from the hardware versus values supplied by an operator. This brought comments that the MIB scope needed better definition. Discussion was largely ended by the statements that the charter is relatively narrow, limited to common OS characteristics, existing MIB objects, and a quick finish and that, although all objects have some value, they must be chosen for the most value without exceeding a cost budget, which includes coding, standardization, interoperability, and market acceptance as indicated by implemented examples from vendors.

Multiport Repeater MIB WG

An Internet Draft has been forwarded and recommended as Proposed Standard to the IAB by the IESG.

OSPF WG

The MIB was updated and a detailed list of changes sent to the mailing list, as was a new traps document.

Remote Monitoring MIB WG

To the question would it be wise for an implementor to generalize the use of event entries to include other events such as `coldStart`, the response was that RMON can set alarms on outside objects, so it seems reasonable to extend `eventEntries` as well.

One query included the question of whether an agent can see its own traffic and complained about the lack of ability to AND two filters. The consensus of responses was that counting its own traffic is necessary, even if it's a software hack, as two probes on the same wire should see the same traffic. On the issue of an AND, you can do any boolean combination with a sum of products.

There was considerable detailed discussion of the new token ring extensions for RMON, much of it centered around what stations are visible on a token ring and how they are represented in the MIB.

An inquiry for a public-domain implementation of an RMON agent received the reply that none are known, but some are expected eventually.

Some suggested solutions to row-dependency problems that appeared at an interoperability test session were discussed, with the conclusion that an agent must always be allowed to reject requests it considers wrong, but multiple managers will always present race conditions.

The question what is the "RMON polka" brought the short response that it is a detailed algorithm in the RMON MIB for using status variables to protect against interference among managers when creating rows. A longer response stated it is a derogatory term for RMON's method for choosing row indexes and protecting creation, which is, in fact, not inefficient and has a very small collision window where one manager would fail, and that window can be further reduced by picking random indexes from a large space.

An expression of disappointment that the list does not discuss the practical use of RMON brought the response the engineers are good at determining what can be implemented rather than what is useful, and it would be good to see a standard that starts with uses then defines the technology to support them. This got a response that it is effective to design, implement, then see uses appear, as a marketing survey is too difficult in a technology-driven market.

SNMP Security WG

The documents are now Proposed Standards, RFCs 1351, 1352, and 1353.

An objection that the source party is out of order in the security message and a suggestion to fix this by requiring source and destination parties to use the same authentication brought the response that the order is based on the architectural order of dependencies, that ASN.1 is best handled by parsing the entire message before applying authentication, and that binding between source and destination parties was a problem that was fixed from the previous architecture.

A question about using transport addresses in party identifiers got the response that party identification is any legal OID, and although the RFC establishes the convention of a space with the agent's IP address, this is not to be taken as the transport address or part of it.

A report that the National Security Agency (NSA) is allowing export of RC2 or RC4 encryption with key sizes less than 40 bits drew the comment that this implies NSA's rejection of export relaxation for other algorithms and key sizes.

Uninterruptable Power Supply MIB WG

An exploratory BOF resulted in a decision to form a working group. MIBs were submitted to the list and a survey circulated. Detailed discussion of the MIB proposals ensued.

X.25 MIB WG

After some rounds of minor changes, the Internet Drafts for the LAPB and Packet MIBs stabilized and have been recommended as Proposed Standards to the IAB by the IESG.

Announcements

SMIC compiler

A compiler for SNMP MIBs, called SMIC, is now available. It is a tool to help develop and use SNMP MIBs. This compiler is being made "freely available" by SynOptics Communications, Inc.

The first version of this package which includes the compiler, associated documents, MIBs, along with a MIB stripper is available via anonymous FTP access to host `sweetwater.synoptics.com` in directory `pub/mibcompiler`. The file `readme.txt` contains information about the components of the package.

The package will run on MS-DOS and many UNIX implementations, and includes source code for a MIB compiler and a MIB stripper, along with "corrected" versions of MIBs from various RFCs. Compiler features include:

- multiple input files;
- concise MIB format (RFC 1212);
- concise trap format (RFC 1215);
- multiple MIB modules;
- items in IMPORTS;
- textual conventions;
- alias assignments for modules and object names;
- selective checking of MIB constructs;
- extensive MIB syntax checking and continuation of syntax checking after syntax errors ;
- extensive checking of MIB consistency;
- multiple output options (including `mosy` compatible output); and,
- environment variable to locate "included" files.

The many people who helped getting SMIC complete and tested are listed in the SMIC User's Guide. Special thanks go to Fred Baker for the MIB stripper program and Dean Throop for help on the YACC error productions.

This package is "freely available" but it is not public domain. The following is the copyright and rights to use message for SMIC:

> © 1992 SynOptics Communications, Inc. All Rights Reserved. SynOptics grants a non-exclusive license to use, copy, modify, and distribute this software for any purpose and without fee, provided that this copyright notice and license appear on all copies and supporting documentation. SynOptics makes no representations about the suitability of this software for any particular purpose. The software is supplied "AS IS", and SynOptics makes no warranty, either express or implied, as to the use, operation, condition, or performance of the software. SynOptics retains all title and ownership in the software.

For further information, contact:

`dperkins@synoptics.com`

Recent Publications

ConneXions. ISSN 0890–8044

- How to Win the Battle for Network Management (July, 1992)

Activities Calendar

- Interop Fall 92

  October 26–30, San Francisco, CA

  For information: +1 415–941–3399 x2502

- 25th Meeting of the IETF

  November 16–20, Washington, DC

  For information: +1 703–620–8990

- 3rd IFIP/IEEE Symposium on Integrated Network Management

  April 18–23, 1993, San Francisco, CA

  For information: kzm@hls.com

Publication Information

The Simple Times is published with a lot of help from the SNMP community.

Publication Staff

Coordinating Editor:	
Dr. Marshall T. Rose	Dover Beach Consulting, Inc.
Featured Columnists:	
Dr. Jeffrey D. Case	SNMP Research, Inc. University of Tennessee
Keith McCloghrie	Hughes LAN Systems, Inc.
David T. Perkins	SynOptics Communications, Inc.
Robert L. Stewart	Xyplex, Inc.
Steven L. Waldbusser	Carnegie Mellon University

Contact Information

Postal: ***The Simple Times***
c/o Dover Beach Consulting, Inc.
420 Whisman Court
Mountain View, CA 94043–2112

Tel: +1 415–968–1052
Fax: +1 415–968–2510
E-mail: `st-editorial@simple-times.org`
ISSN: 1060–6068

Submissions

The Simple Times solicits high-quality articles of technology and comment. Technical articles are refereed to ensure that the content is marketing-free. By definition, commentaries reflect opinion and, as such, are reviewed only to the extent required to ensure commonly-accepted publication norms.

The Simple Times also solicits terse announcements of products and services, publications, and events. These contributions are reviewed only to the extent required to ensure commonly-accepted publication norms.

Submissions are accepted only in electronic form. A submission consists of ASCII text. (Technical articles are also allowed to reference encapsulated PostScript figures.) Submissions may be sent to the contact address above, either via electronic-mail or via magnetic media (using either 8mm `tar` tape, 1/4in `tar` cartridge-tape, or 3-1/2in MS-DOS floppy-diskette).

Each submission must include the author's full name, title, affiliation, postal and electronic mail addresses, telephone, and fax numbers. Note that by initiating this process, the submitting party agrees to place the contribution into the public domain.

Subscriptions

The Simple Times is available via electronic-mail in three editions: *PostScript*, *MIME* (the multi-media 822 mail format), and *richtext* (a simple page description language). For more information, send a message to

```
st-subscriptions@simple-times.org
```

with a `Subject` line of

```
help
```

In addition, ***The Simple Times*** has numerous hard-copy distribution outlets. Contact your favorite SNMP vendor and see if they carry it. If not, contact the publisher and ask for a list. (Communications via e-mail or fax are preferred).