

DECLARATION OF SANDY GINOZA FOR IETF  
RFC 1157: A Simple Network Management Protocol (SNMP)

I, Sandy Ginoza, hereby declare that all statements made herein are of my own knowledge and are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code:

1. I am an employee of Association Management Solutions, LLC (AMS), which acts under contract to the IETF Administration LLC (IETF) as the operator of the RFC Production Center. The RFC Production Center is part of the "RFC Editor" function, which prepares documents for publication and places files in an online repository for the authoritative Request for Comments (RFC) series of documents (RFC Series), and preserves records relating to these documents. The RFC Series includes, among other things, the series of Internet standards developed by the IETF. I hold the position of Director of the RFC Production Center. I began employment with AMS in this capacity on 6 January 2010.

2. My responsibilities as Director of the RFC Production Center include acting as the custodian of records relating to the RFC Series, and I am familiar with the record keeping practices relating to the RFC Series, including the creation and maintenance of such records.

3. From June 1999 to 5 January 2010, I was an employee of the Information Sciences Institute at University of Southern California (ISI). I held various position titles with the RFC Editor project at ISI, ending with Senior Editor.

4. The RFC Editor function was conducted by ISI under contract to the United States government prior to 1998. In 1998, the Internet Society (ISOC), in furtherance of its IETF

activity, entered into the first in a series of contracts with ISI providing for ISI's performance of the RFC Editor function. Beginning in 2010, certain aspects of the RFC Editor function were assumed by the RFC Production Center operation of AMS under contract to ISOC (acting through its IETF function and, in particular, the IETF Administrative Oversight Committee (now the IETF Administration LLC). The business records of the RFC Editor function, as it was conducted by ISI, are currently housed with a cloud vendor under contract with IETF Administration LLC.

5. I make this declaration based on my personal knowledge and information contained in the business records of the RFC Editor as they are currently housed at AMS, or confirmation with other responsible RFC Editor personnel with such knowledge.

6. Prior to 1998, the RFC Editor's regular practice was to publish RFCs, making them available from a repository via FTP. When a new RFC was published, an announcement of its publication, with information on how to access the RFC, would be typically sent out within 24 hours of the publication.

7. Since 1998, the RFC Editor's regular practice was to publish RFCs, making them available on the RFC Editor website or via FTP. When a new RFC was published, an announcement of its publication, with information on how to access the RFC, would be typically sent out within 24 hours of the publication. The announcement would go out to all subscribers and a contemporaneous electronic record of the announcement is kept in the IETF mail archive that is available online.

8. Beginning in 1998, any RFC published on the RFC Editor website or via FTP was reasonably accessible to the public and was disseminated or otherwise available to the extent that persons interested and ordinarily skilled in the subject matter or art exercising reasonable

diligence could have located it. In particular, the RFCs were indexed and placed in a public repository.

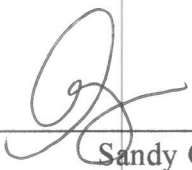
9. The RFCs are kept in an online repository in the course of the RFC Editor's regularly conducted activity and ordinary course of business. The records are made pursuant to established procedures and are relied upon by the RFC Editor in the performance of its functions.

10. It is the regular practice of the RFC Editor to make and keep the RFC records.

11. Based on the business records for the RFC Editor and the RFC Editor's course of conduct in publishing RFCs, I have determined that the publication date of RFC 1157 was no later than May 1990, at which time it was reasonably accessible to the public either on the RFC Editor website or via FTP from a repository. An announcement of its publication also would have been sent out to subscribers within 24 hours of its publication. A copy of that RFC is attached to this declaration as Exhibit 1.

Pursuant to Section 1746 of Title 28 of United States Code, I declare under penalty of perjury under the laws of the United States of America that the foregoing is true and correct and that the foregoing is based upon personal knowledge and information and is believed to be true.

Date: 30 October 2024

By:   
Sandy Ginoza

Network Working Group  
Request for Comments: 1157  
Obsoletes: RFC 1098

J. Case  
SNMP Research  
M. Fedor  
Performance Systems International  
M. Schoffstall  
Performance Systems International  
J. Davin  
MIT Laboratory for Computer Science  
May 1990

## A Simple Network Management Protocol (SNMP)

### Table of Contents

1. Status of this Memo .....	2
2. Introduction .....	2
3. The SNMP Architecture .....	5
3.1 Goals of the Architecture .....	5
3.2 Elements of the Architecture .....	5
3.2.1 Scope of Management Information .....	6
3.2.2 Representation of Management Information .....	6
3.2.3 Operations Supported on Management Information .....	7
3.2.4 Form and Meaning of Protocol Exchanges .....	8
3.2.5 Definition of Administrative Relationships .....	8
3.2.6 Form and Meaning of References to Managed Objects ..	12
3.2.6.1 Resolution of Ambiguous MIB References .....	12
3.2.6.2 Resolution of References across MIB Versions.....	12
3.2.6.3 Identification of Object Instances .....	12
3.2.6.3.1 ifTable Object Type Names .....	13
3.2.6.3.2 atTable Object Type Names .....	13
3.2.6.3.3 ipAddrTable Object Type Names .....	14
3.2.6.3.4 ipRoutingTable Object Type Names .....	14
3.2.6.3.5 tcpConnTable Object Type Names .....	14
3.2.6.3.6 egpNeighTable Object Type Names .....	15
4. Protocol Specification .....	16
4.1 Elements of Procedure .....	17
4.1.1 Common Constructs .....	19
4.1.2 The GetRequest-PDU .....	20
4.1.3 The GetNextRequest-PDU .....	21
4.1.3.1 Example of Table Traversal .....	23
4.1.4 The GetResponse-PDU .....	24
4.1.5 The SetRequest-PDU .....	25
4.1.6 The Trap-PDU .....	27
4.1.6.1 The coldStart Trap .....	28
4.1.6.2 The warmStart Trap .....	28
4.1.6.3 The linkDown Trap .....	28
4.1.6.4 The linkUp Trap .....	28

4.1.6.5 The authenticationFailure Trap .....	28
4.1.6.6 The egpNeighborLoss Trap .....	28
4.1.6.7 The enterpriseSpecific Trap .....	29
5. Definitions .....	30
6. Acknowledgements .....	33
7. References .....	34
8. Security Considerations.....	35
9. Authors' Addresses.....	35

## 1. Status of this Memo

This RFC is a re-release of RFC 1098, with a changed "Status of this Memo" section plus a few minor typographical corrections. This memo defines a simple protocol by which management information for a network element may be inspected or altered by logically remote users. In particular, together with its companion memos which describe the structure of management information along with the management information base, these documents provide a simple, workable architecture and system for managing TCP/IP-based internets and in particular the Internet.

The Internet Activities Board recommends that all IP and TCP implementations be network manageable. This implies implementation of the Internet MIB (RFC-1156) and at least one of the two recommended management protocols SNMP (RFC-1157) or CMOT (RFC-1095). It should be noted that, at this time, SNMP is a full Internet standard and CMOT is a draft standard. See also the Host and Gateway Requirements RFCs for more specific information on the applicability of this standard.

Please refer to the latest edition of the "IAB Official Protocol Standards" RFC for current information on the state and status of standard Internet protocols.

Distribution of this memo is unlimited.

## 2. Introduction

As reported in RFC 1052, IAB Recommendations for the Development of Internet Network Management Standards [1], a two-prong strategy for network management of TCP/IP-based internets was undertaken. In the short-term, the Simple Network Management Protocol (SNMP) was to be used to manage nodes in the Internet community. In the long-term, the use of the OSI network management framework was to be examined. Two documents were produced to define the management information: RFC 1065, which defined the Structure of Management Information (SMI) [2], and RFC 1066, which defined the Management Information Base (MIB) [3]. Both of these documents were designed so as to be

compatible with both the SNMP and the OSI network management framework.

This strategy was quite successful in the short-term: Internet-based network management technology was fielded, by both the research and commercial communities, within a few months. As a result of this, portions of the Internet community became network manageable in a timely fashion.

As reported in RFC 1109, Report of the Second Ad Hoc Network Management Review Group [4], the requirements of the SNMP and the OSI network management frameworks were more different than anticipated. As such, the requirement for compatibility between the SMI/MIB and both frameworks was suspended. This action permitted the operational network management framework, the SNMP, to respond to new operational needs in the Internet community by producing documents defining new MIB items.

The IAB has designated the SNMP, SMI, and the initial Internet MIB to be full "Standard Protocols" with "Recommended" status. By this action, the IAB recommends that all IP and TCP implementations be network manageable and that the implementations that are network manageable are expected to adopt and implement the SMI, MIB, and SNMP.

As such, the current network management framework for TCP/IP-based internets consists of: Structure and Identification of Management Information for TCP/IP-based Internets, which describes how managed objects contained in the MIB are defined as set forth in RFC 1155 [5]; Management Information Base for Network Management of TCP/IP-based Internets, which describes the managed objects contained in the MIB as set forth in RFC 1156 [6]; and, the Simple Network Management Protocol, which defines the protocol used to manage these objects, as set forth in this memo.

As reported in RFC 1052, IAB Recommendations for the Development of Internet Network Management Standards [1], the Internet Activities Board has directed the Internet Engineering Task Force (IETF) to create two new working groups in the area of network management. One group was charged with the further specification and definition of elements to be included in the Management Information Base (MIB). The other was charged with defining the modifications to the Simple Network Management Protocol (SNMP) to accommodate the short-term needs of the network vendor and operations communities, and to align with the output of the MIB working group.

The MIB working group produced two memos, one which defines a Structure for Management Information (SMI) [2] for use by the managed

objects contained in the MIB. A second memo [3] defines the list of managed objects.

The output of the SNMP Extensions working group is this memo, which incorporates changes to the initial SNMP definition [7] required to attain alignment with the output of the MIB working group. The changes should be minimal in order to be consistent with the IAB's directive that the working groups be "extremely sensitive to the need to keep the SNMP simple." Although considerable care and debate has gone into the changes to the SNMP which are reflected in this memo, the resulting protocol is not backwardly-compatible with its predecessor, the Simple Gateway Monitoring Protocol (SGMP) [8]. Although the syntax of the protocol has been altered, the original philosophy, design decisions, and architecture remain intact. In order to avoid confusion, new UDP ports have been allocated for use by the protocol described in this memo.

### 3. The SNMP Architecture

Implicit in the SNMP architectural model is a collection of network management stations and network elements. Network management stations execute management applications which monitor and control network elements. Network elements are devices such as hosts, gateways, terminal servers, and the like, which have management agents responsible for performing the network management functions requested by the network management stations. The Simple Network Management Protocol (SNMP) is used to communicate management information between the network management stations and the agents in the network elements.

#### 3.1. Goals of the Architecture

The SNMP explicitly minimizes the number and complexity of management functions realized by the management agent itself. This goal is attractive in at least four respects:

- (1) The development cost for management agent software necessary to support the protocol is accordingly reduced.
- (2) The degree of management function that is remotely supported is accordingly increased, thereby admitting fullest use of internet resources in the management task.
- (3) The degree of management function that is remotely supported is accordingly increased, thereby imposing the fewest possible restrictions on the form and sophistication of management tools.
- (4) Simplified sets of management functions are easily understood and used by developers of network management tools.

A second goal of the protocol is that the functional paradigm for monitoring and control be sufficiently extensible to accommodate additional, possibly unanticipated aspects of network operation and management.

A third goal is that the architecture be, as much as possible, independent of the architecture and mechanisms of particular hosts or particular gateways.

#### 3.2. Elements of the Architecture

The SNMP architecture articulates a solution to the network management problem in terms of:



- (1) the scope of the management information communicated by the protocol,
- (2) the representation of the management information communicated by the protocol,
- (3) operations on management information supported by the protocol,
- (4) the form and meaning of exchanges among management entities,
- (5) the definition of administrative relationships among management entities, and
- (6) the form and meaning of references to management information.

#### 3.2.1. Scope of Management Information

The scope of the management information communicated by operation of the SNMP is exactly that represented by instances of all non-aggregate object types either defined in Internet-standard MIB or defined elsewhere according to the conventions set forth in Internet-standard SMI [5].

Support for aggregate object types in the MIB is neither required for conformance with the SMI nor realized by the SNMP.

#### 3.2.2. Representation of Management Information

Management information communicated by operation of the SNMP is represented according to the subset of the ASN.1 language [9] that is specified for the definition of non-aggregate types in the SMI.

The SGMP adopted the convention of using a well-defined subset of the ASN.1 language [9]. The SNMP continues and extends this tradition by utilizing a moderately more complex subset of ASN.1 for describing managed objects and for describing the protocol data units used for managing those objects. In addition, the desire to ease eventual transition to OSI-based network management protocols led to the definition in the ASN.1 language of an Internet-standard Structure of Management Information (SMI) [5] and Management Information Base (MIB) [6]. The use of the ASN.1 language, was, in part, encouraged by the successful use of ASN.1 in earlier efforts, in particular, the SGMP. The restrictions on the use of ASN.1 that are part of the SMI contribute to the simplicity espoused and validated by experience with the SGMP.

Also for the sake of simplicity, the SNMP uses only a subset of the basic encoding rules of ASN.1 [10]. Namely, all encodings use the definite-length form. Further, whenever permissible, non-constructor encodings are used rather than constructor encodings. This restriction applies to all aspects of ASN.1 encoding, both for the top-level protocol data units and the data objects they contain.

### 3.2.3. Operations Supported on Management Information

The SNMP models all management agent functions as alterations or inspections of variables. Thus, a protocol entity on a logically remote host (possibly the network element itself) interacts with the management agent resident on the network element in order to retrieve (get) or alter (set) variables. This strategy has at least two positive consequences:

- (1) It has the effect of limiting the number of essential management functions realized by the management agent to two: one operation to assign a value to a specified configuration or other parameter and another to retrieve such a value.
- (2) A second effect of this decision is to avoid introducing into the protocol definition support for imperative management commands: the number of such commands is in practice ever-increasing, and the semantics of such commands are in general arbitrarily complex.

The strategy implicit in the SNMP is that the monitoring of network state at any significant level of detail is accomplished primarily by polling for appropriate information on the part of the monitoring center(s). A limited number of unsolicited messages (traps) guide the timing and focus of the polling. Limiting the number of unsolicited messages is consistent with the goal of simplicity and minimizing the amount of traffic generated by the network management function.

The exclusion of imperative commands from the set of explicitly supported management functions is unlikely to preclude any desirable management agent operation. Currently, most commands are requests either to set the value of some parameter or to retrieve such a value, and the function of the few imperative commands currently supported is easily accommodated in an asynchronous mode by this management model. In this scheme, an imperative command might be realized as the setting of a parameter value that subsequently triggers the desired action. For example, rather than implementing a "reboot command," this action might be invoked by simply setting a parameter indicating the number of seconds until system reboot.

#### 3.2.4. Form and Meaning of Protocol Exchanges

The communication of management information among management entities is realized in the SNMP through the exchange of protocol messages. The form and meaning of those messages is defined below in Section 4.

Consistent with the goal of minimizing complexity of the management agent, the exchange of SNMP messages requires only an unreliable datagram service, and every message is entirely and independently represented by a single transport datagram. While this document specifies the exchange of messages via the UDP protocol [11], the mechanisms of the SNMP are generally suitable for use with a wide variety of transport services.

#### 3.2.5. Definition of Administrative Relationships

The SNMP architecture admits a variety of administrative relationships among entities that participate in the protocol. The entities residing at management stations and network elements which communicate with one another using the SNMP are termed SNMP application entities. The peer processes which implement the SNMP, and thus support the SNMP application entities, are termed protocol entities.

A pairing of an SNMP agent with some arbitrary set of SNMP application entities is called an SNMP community. Each SNMP community is named by a string of octets, that is called the community name for said community.

An SNMP message originated by an SNMP application entity that in fact belongs to the SNMP community named by the community component of said message is called an authentic SNMP message. The set of rules by which an SNMP message is identified as an authentic SNMP message for a particular SNMP community is called an authentication scheme. An implementation of a function that identifies authentic SNMP messages according to one or more authentication schemes is called an authentication service.

Clearly, effective management of administrative relationships among SNMP application entities requires authentication services that (by the use of encryption or other techniques) are able to identify authentic SNMP messages with a high degree of certainty. Some SNMP implementations may wish to support only a trivial authentication service that identifies all SNMP messages as authentic SNMP messages.

For any network element, a subset of objects in the MIB that pertain to that element is called a SNMP MIB view. Note that the names of the object types represented in a SNMP MIB view need not belong to a

single sub-tree of the object type name space.

An element of the set { READ-ONLY, READ-WRITE } is called an SNMP access mode.

A pairing of a SNMP access mode with a SNMP MIB view is called an SNMP community profile. A SNMP community profile represents specified access privileges to variables in a specified MIB view. For every variable in the MIB view in a given SNMP community profile, access to that variable is represented by the profile according to the following conventions:

- (1) if said variable is defined in the MIB with "Access:" of "none," it is unavailable as an operand for any operator;
- (2) if said variable is defined in the MIB with "Access:" of "read-write" or "write-only" and the access mode of the given profile is READ-WRITE, that variable is available as an operand for the get, set, and trap operations;
- (3) otherwise, the variable is available as an operand for the get and trap operations.
- (4) In those cases where a "write-only" variable is an operand used for the get or trap operations, the value given for the variable is implementation-specific.

A pairing of a SNMP community with a SNMP community profile is called a SNMP access policy. An access policy represents a specified community profile afforded by the SNMP agent of a specified SNMP community to other members of that community. All administrative relationships among SNMP application entities are architecturally defined in terms of SNMP access policies.


For every SNMP access policy, if the network element on which the SNMP agent for the specified SNMP community resides is not that to which the MIB view for the specified profile pertains, then that policy is called a SNMP proxy access policy. The SNMP agent associated with a proxy access policy is called a SNMP proxy agent. While careless definition of proxy access policies can result in management loops, prudent definition of proxy policies is useful in at least two ways:

- (1) It permits the monitoring and control of network elements which are otherwise not addressable using the management protocol and the transport protocol. That is, a proxy agent may provide a protocol conversion function allowing a management station to apply a consistent management

framework to all network elements, including devices such as modems, multiplexors, and other devices which support different management frameworks.

- (2) It potentially shields network elements from elaborate access control policies. For example, a proxy agent may implement sophisticated access control whereby diverse subsets of variables within the MIB are made accessible to different management stations without increasing the complexity of the network element.

By way of example, Figure 1 illustrates the relationship between management stations, proxy agents, and management agents. In this example, the proxy agent is envisioned to be a normal Internet Network Operations Center (INOC) of some administrative domain which has a standard managerial relationship with a set of management agents.



Domain: the administrative domain of the element  
 PCommunity: the name of a community utilizing a proxy agent  
 DCommunity: the name of a direct community

Figure 1  
 Example Network Management Configuration

### 3.2.6. Form and Meaning of References to Managed Objects

The SMI requires that the definition of a conformant management protocol address:

- (1) the resolution of ambiguous MIB references,
- (2) the resolution of MIB references in the presence multiple MIB versions, and
- (3) the identification of particular instances of object types defined in the MIB.

#### 3.2.6.1. Resolution of Ambiguous MIB References

Because the scope of any SNMP operation is conceptually confined to objects relevant to a single network element, and because all SNMP references to MIB objects are (implicitly or explicitly) by unique variable names, there is no possibility that any SNMP reference to any object type defined in the MIB could resolve to multiple instances of that type.

#### 3.2.6.2. Resolution of References across MIB Versions

The object instance referred to by any SNMP operation is exactly that specified as part of the operation request or (in the case of a get-next operation) its immediate successor in the MIB as a whole. In particular, a reference to an object as part of some version of the Internet-standard MIB does not resolve to any object that is not part of said version of the Internet-standard MIB, except in the case that the requested operation is get-next and the specified object name is lexicographically last among the names of all objects presented as part of said version of the Internet-Standard MIB.

#### 3.2.6.3. Identification of Object Instances

The names for all object types in the MIB are defined explicitly either in the Internet-standard MIB or in other documents which conform to the naming conventions of the SMI. The SMI requires that conformant management protocols define mechanisms for identifying individual instances of those object types for a particular network element.

Each instance of any object type defined in the MIB is identified in SNMP operations by a unique name called its "variable name." In general, the name of an SNMP variable is an OBJECT IDENTIFIER of the form x.y, where x is the name of a non-aggregate object type defined in the MIB and y is an OBJECT IDENTIFIER fragment that, in a way

specific to the named object type, identifies the desired instance.

This naming strategy admits the fullest exploitation of the semantics of the GetNextRequest-PDU (see Section 4), because it assigns names for related variables so as to be contiguous in the lexicographical ordering of all variable names known in the MIB.

The type-specific naming of object instances is defined below for a number of classes of object types. Instances of an object type to which none of the following naming conventions are applicable are named by OBJECT IDENTIFIERS of the form x.0, where x is the name of said object type in the MIB definition.

For example, suppose one wanted to identify an instance of the variable sysDescr. The object class for sysDescr is:

```
iso org dod internet mgmt mib system sysDescr
 1  3  6      1      2    1    1      1
```

Hence, the object type, x, would be 1.3.6.1.2.1.1.1 to which is appended an instance sub-identifier of 0. That is, 1.3.6.1.2.1.1.1.0 identifies the one and only instance of sysDescr.

#### 3.2.6.3.1. ifTable Object Type Names

The name of a subnet interface, s, is the OBJECT IDENTIFIER value of the form i, where i has the value of that instance of the ifIndex object type associated with s.

For each object type, t, for which the defined name, n, has a prefix of ifEntry, an instance, i, of t is named by an OBJECT IDENTIFIER of the form n.s, where s is the name of the subnet interface about which i represents information.

For example, suppose one wanted to identify the instance of the variable ifType associated with interface 2. Accordingly, ifType.2 would identify the desired instance.

#### 3.2.6.3.2. atTable Object Type Names

The name of an AT-cached network address, x, is an OBJECT IDENTIFIER of the form l.a.b.c.d, where a.b.c.d is the value (in the familiar "dot" notation) of the atNetAddress object type associated with x.

The name of an address translation equivalence e is an OBJECT IDENTIFIER value of the form s.w, such that s is the value of that instance of the atIndex object type associated with e and such that w is the name of the AT-cached network address associated with e.



For each object type, *t*, for which the defined name, *n*, has a prefix of *atEntry*, an instance, *i*, of *t* is named by an OBJECT IDENTIFIER of the form *n.y*, where *y* is the name of the address translation equivalence about which *i* represents information.

For example, suppose one wanted to find the physical address of an entry in the address translation table (ARP cache) associated with an IP address of 89.1.1.42 and interface 3. Accordingly, *atPhysAddress.3.1.89.1.1.42* would identify the desired instance.

#### 3.2.6.3.3. *ipAddrTable* Object Type Names

The name of an IP-addressable network element, *x*, is the OBJECT IDENTIFIER of the form *a.b.c.d* such that *a.b.c.d* is the value (in the familiar "dot" notation) of that instance of the *ipAdEntAddr* object type associated with *x*.

For each object type, *t*, for which the defined name, *n*, has a prefix of *ipAddrEntry*, an instance, *i*, of *t* is named by an OBJECT IDENTIFIER of the form *n.y*, where *y* is the name of the IP-addressable network element about which *i* represents information.

For example, suppose one wanted to find the network mask of an entry in the IP interface table associated with an IP address of 89.1.1.42. Accordingly, *ipAdEntNetMask.89.1.1.42* would identify the desired instance.

#### 3.2.6.3.4. *ipRoutingTable* Object Type Names

The name of an IP route, *x*, is the OBJECT IDENTIFIER of the form *a.b.c.d* such that *a.b.c.d* is the value (in the familiar "dot" notation) of that instance of the *ipRouteDest* object type associated with *x*.

For each object type, *t*, for which the defined name, *n*, has a prefix of *ipRoutingEntry*, an instance, *i*, of *t* is named by an OBJECT IDENTIFIER of the form *n.y*, where *y* is the name of the IP route about which *i* represents information.

For example, suppose one wanted to find the next hop of an entry in the IP routing table associated with the destination of 89.1.1.42. Accordingly, *ipRouteNextHop.89.1.1.42* would identify the desired instance.

#### 3.2.6.3.5. *tcpConnTable* Object Type Names

The name of a TCP connection, *x*, is the OBJECT IDENTIFIER of the form *a.b.c.d.e.f.g.h.i.j* such that *a.b.c.d* is the value (in the familiar

"dot" notation) of that instance of the tcpConnLocalAddress object type associated with x and such that f.g.h.i is the value (in the familiar "dot" notation) of that instance of the tcpConnRemoteAddress object type associated with x and such that e is the value of that instance of the tcpConnLocalPort object type associated with x and such that j is the value of that instance of the tcpConnRemotePort object type associated with x.

For each object type, t, for which the defined name, n, has a prefix of tcpConnEntry, an instance, i, of t is named by an OBJECT IDENTIFIER of the form n.y, where y is the name of the TCP connection about which i represents information.

For example, suppose one wanted to find the state of a TCP connection between the local address of 89.1.1.42 on TCP port 21 and the remote address of 10.0.0.51 on TCP port 2059. Accordingly, tcpConnState.89.1.1.42.21.10.0.0.51.2059 would identify the desired instance.

#### 3.2.6.3.6. egpNeighTable Object Type Names

The name of an EGP neighbor, x, is the OBJECT IDENTIFIER of the form a.b.c.d such that a.b.c.d is the value (in the familiar "dot" notation) of that instance of the egpNeighAddr object type associated with x.

For each object type, t, for which the defined name, n, has a prefix of egpNeighEntry, an instance, i, of t is named by an OBJECT IDENTIFIER of the form n.y, where y is the name of the EGP neighbor about which i represents information.

For example, suppose one wanted to find the neighbor state for the IP address of 89.1.1.42. Accordingly, egpNeighState.89.1.1.42 would identify the desired instance.

#### 4. Protocol Specification

The network management protocol is an application protocol by which the variables of an agent's MIB may be inspected or altered.

Communication among protocol entities is accomplished by the exchange of messages, each of which is entirely and independently represented within a single UDP datagram using the basic encoding rules of ASN.1 (as discussed in Section 3.2.2). A message consists of a version identifier, an SNMP community name, and a protocol data unit (PDU). A protocol entity receives messages at UDP port 161 on the host with which it is associated for all messages except for those which report traps (i.e., all messages except those which contain the Trap-PDU). Messages which report traps should be received on UDP port 162 for further processing. An implementation of this protocol need not accept messages whose length exceeds 484 octets. However, it is recommended that implementations support larger datagrams whenever feasible.

It is mandatory that all implementations of the SNMP support the five PDUs: GetRequest-PDU, GetNextRequest-PDU, GetResponse-PDU, SetRequest-PDU, and Trap-PDU.

```
RFC1157-SNMP DEFINITIONS ::= BEGIN
```

```
IMPORTS
```

```
    ObjectName, ObjectSyntax, NetworkAddress, IpAddress, TimeTicks
    FROM RFC1155-SMI;
```

```
-- top-level message
```

```
Message ::=
    SEQUENCE {
        version          -- version-1 for this RFC
            INTEGER {
                version-1(0)
            },
        community        -- community name
            OCTET STRING,
        data              -- e.g., PDUs if trivial
            ANY           -- authentication is being used
    }
```

```
-- protocol data units

    PDUs ::=
        CHOICE {
            get-request
                GetRequest-PDU,

            get-next-request
                GetNextRequest-PDU,

            get-response
                GetResponse-PDU,

            set-request
                SetRequest-PDU,

            trap
                Trap-PDU
        }

-- the individual PDUs and commonly used
-- data types will be defined later

END
```

#### 4.1. Elements of Procedure

This section describes the actions of a protocol entity implementing the SNMP. Note, however, that it is not intended to constrain the internal architecture of any conformant implementation.

In the text that follows, the term transport address is used. In the case of the UDP, a transport address consists of an IP address along with a UDP port. Other transport services may be used to support the SNMP. In these cases, the definition of a transport address should be made accordingly.

The top-level actions of a protocol entity which generates a message are as follows:

- (1) It first constructs the appropriate PDU, e.g., the GetRequest-PDU, as an ASN.1 object.
- (2) It then passes this ASN.1 object along with a community name its source transport address and the destination transport address, to the service which implements the desired authentication scheme. This authentication

service returns another ASN.1 object.

- (3) The protocol entity then constructs an ASN.1 Message object, using the community name and the resulting ASN.1 object.
- (4) This new ASN.1 object is then serialized, using the basic encoding rules of ASN.1, and then sent using a transport service to the peer protocol entity.

Similarly, the top-level actions of a protocol entity which receives a message are as follows:

- (1) It performs a rudimentary parse of the incoming datagram to build an ASN.1 object corresponding to an ASN.1 Message object. If the parse fails, it discards the datagram and performs no further actions.
- (2) It then verifies the version number of the SNMP message. If there is a mismatch, it discards the datagram and performs no further actions.
- (3) The protocol entity then passes the community name and user data found in the ASN.1 Message object, along with the datagram's source and destination transport addresses to the service which implements the desired authentication scheme. This entity returns another ASN.1 object, or signals an authentication failure. In the latter case, the protocol entity notes this failure, (possibly) generates a trap, and discards the datagram and performs no further actions.
- (4) The protocol entity then performs a rudimentary parse on the ASN.1 object returned from the authentication service to build an ASN.1 object corresponding to an ASN.1 PDU object. If the parse fails, it discards the datagram and performs no further actions. Otherwise, using the named SNMP community, the appropriate profile is selected, and the PDU is processed accordingly. If, as a result of this processing, a message is returned then the source transport address that the response message is sent from shall be identical to the destination transport address that the original request message was sent to.

## 4.1.1. Common Constructs

Before introducing the six PDU types of the protocol, it is appropriate to consider some of the ASN.1 constructs used frequently:

```
-- request/response information
```

```
RequestID ::=
    INTEGER
```

```
ErrorStatus ::=
    INTEGER {
        noError(0),
        tooBig(1),
        noSuchName(2),
        badValue(3),
        readOnly(4),
        genErr(5)
    }
```

```
ErrorIndex ::=
    INTEGER
```

```
-- variable bindings
```

```
VarBind ::=
    SEQUENCE {
        name
            ObjectName,

        value
            ObjectSyntax
    }
```

```
VarBindList ::=
    SEQUENCE OF
        VarBind
```

RequestIDs are used to distinguish among outstanding requests. By use of the RequestID, an SNMP application entity can correlate incoming responses with outstanding requests. In cases where an unreliable datagram service is being used, the RequestID also provides a simple means of identifying messages duplicated by the network.

A non-zero instance of ErrorStatus is used to indicate that an

exception occurred while processing a request. In these cases, `ErrorIndex` may provide additional information by indicating which variable in a list caused the exception.

The term variable refers to an instance of a managed object. A variable binding, or `VarBind`, refers to the pairing of the name of a variable to the variable's value. A `VarBindList` is a simple list of variable names and corresponding values. Some PDUs are concerned only with the name of a variable and not its value (e.g., the `GetRequest-PDU`). In this case, the value portion of the binding is ignored by the protocol entity. However, the value portion must still have valid ASN.1 syntax and encoding. It is recommended that the ASN.1 value `NULL` be used for the value portion of such bindings.

#### 4.1.2. The `GetRequest-PDU`

The form of the `GetRequest-PDU` is:

```
GetRequest-PDU ::=
  [0]
    IMPLICIT SEQUENCE {
      request-id
        RequestID,

      error-status      -- always 0
        ErrorStatus,

      error-index      -- always 0
        ErrorIndex,

      variable-bindings
        VarBindList
    }
```

The `GetRequest-PDU` is generated by a protocol entity only at the request of its SNMP application entity.

Upon receipt of the `GetRequest-PDU`, the receiving protocol entity responds according to any applicable rule in the list below:

- (1) If, for any object named in the `variable-bindings` field, the object's name does not exactly match the name of some object available for get operations in the relevant MIB view, then the receiving entity sends to the originator of the received message the `GetResponse-PDU` of identical form, except that the value of the `error-status` field is `noSuchName`, and the value of the `error-index` field is the index of said object name component in the received

message.

- (2) If, for any object named in the variable-bindings field, the object is an aggregate type (as defined in the SMI), then the receiving entity sends to the originator of the received message the GetResponse-PDU of identical form, except that the value of the error-status field is noSuchName, and the value of the error-index field is the index of said object name component in the received message.
- (3) If the size of the GetResponse-PDU generated as described below would exceed a local limitation, then the receiving entity sends to the originator of the received message the GetResponse-PDU of identical form, except that the value of the error-status field is tooBig, and the value of the error-index field is zero.
- (4) If, for any object named in the variable-bindings field, the value of the object cannot be retrieved for reasons not covered by any of the foregoing rules, then the receiving entity sends to the originator of the received message the GetResponse-PDU of identical form, except that the value of the error-status field is genErr and the value of the error-index field is the index of said object name component in the received message.

If none of the foregoing rules apply, then the receiving protocol entity sends to the originator of the received message the GetResponse-PDU such that, for each object named in the variable-bindings field of the received message, the corresponding component of the GetResponse-PDU represents the name and value of that variable. The value of the error-status field of the GetResponse-PDU is noError and the value of the error-index field is zero. The value of the request-id field of the GetResponse-PDU is that of the received message.

#### 4.1.3. The GetNextRequest-PDU

The form of the GetNextRequest-PDU is identical to that of the GetRequest-PDU except for the indication of the PDU type. In the ASN.1 language:

```

GetNextRequest-PDU ::=
    [1]
        IMPLICIT SEQUENCE {
            request-id
            RequestID,

```



```
        error-status      -- always 0
            ErrorStatus,

        error-index       -- always 0
            ErrorIndex,

        variable-bindings
            VarBindList
    }
```

The GetNextRequest-PDU is generated by a protocol entity only at the request of its SNMP application entity.

Upon receipt of the GetNextRequest-PDU, the receiving protocol entity responds according to any applicable rule in the list below:

- (1) If, for any object name in the variable-bindings field, that name does not lexicographically precede the name of some object available for get operations in the relevant MIB view, then the receiving entity sends to the originator of the received message the GetResponse-PDU of identical form, except that the value of the error-status field is noSuchName, and the value of the error-index field is the index of said object name component in the received message.
- (2) If the size of the GetResponse-PDU generated as described below would exceed a local limitation, then the receiving entity sends to the originator of the received message the GetResponse-PDU of identical form, except that the value of the error-status field is tooBig, and the value of the error-index field is zero.
- (3) If, for any object named in the variable-bindings field, the value of the lexicographical successor to the named object cannot be retrieved for reasons not covered by any of the foregoing rules, then the receiving entity sends to the originator of the received message the GetResponse-PDU of identical form, except that the value of the error-status field is genErr and the value of the error-index field is the index of said object name component in the received message.

If none of the foregoing rules apply, then the receiving protocol entity sends to the originator of the received message the GetResponse-PDU such that, for each name in the variable-bindings field of the received message, the corresponding component of the

GetResponse-PDU represents the name and value of that object whose name is, in the lexicographical ordering of the names of all objects available for get operations in the relevant MIB view, together with the value of the name field of the given component, the immediate successor to that value. The value of the error-status field of the GetResponse-PDU is noError and the value of the errorindex field is zero. The value of the request-id field of the GetResponse-PDU is that of the received message.

#### 4.1.3.1. Example of Table Traversal

One important use of the GetNextRequest-PDU is the traversal of conceptual tables of information within the MIB. The semantics of this type of SNMP message, together with the protocol-specific mechanisms for identifying individual instances of object types in the MIB, affords access to related objects in the MIB as if they enjoyed a tabular organization.

By the SNMP exchange sketched below, an SNMP application entity might extract the destination address and next hop gateway for each entry in the routing table of a particular network element. Suppose that this routing table has three entries:

Destination	NextHop	Metric
10.0.0.99	89.1.1.42	5
9.1.2.3	99.0.0.3	3
10.0.0.51	89.1.1.42	5

The management station sends to the SNMP agent a GetNextRequest-PDU containing the indicated OBJECT IDENTIFIER values as the requested variable names:

```
GetNextRequest ( ipRouteDest, ipRouteNextHop, ipRouteMetric1 )
```

The SNMP agent responds with a GetResponse-PDU:

```
GetResponse (( ipRouteDest.9.1.2.3 = "9.1.2.3" ),
              ( ipRouteNextHop.9.1.2.3 = "99.0.0.3" ),
              ( ipRouteMetric1.9.1.2.3 = 3 ))
```

The management station continues with:

```
GetNextRequest ( ipRouteDest.9.1.2.3,
                 ipRouteNextHop.9.1.2.3,
```

```
ipRouteMetric1.9.1.2.3 )
```

The SNMP agent responds:

```
GetResponse (( ipRouteDest.10.0.0.51 = "10.0.0.51" ),
             ( ipRouteNextHop.10.0.0.51 = "89.1.1.42" ),
             ( ipRouteMetric1.10.0.0.51 = 5 ))
```

The management station continues with:

```
GetNextRequest ( ipRouteDest.10.0.0.51,
                 ipRouteNextHop.10.0.0.51,
                 ipRouteMetric1.10.0.0.51 )
```

The SNMP agent responds:

```
GetResponse (( ipRouteDest.10.0.0.99 = "10.0.0.99" ),
             ( ipRouteNextHop.10.0.0.99 = "89.1.1.42" ),
             ( ipRouteMetric1.10.0.0.99 = 5 ))
```

The management station continues with:

```
GetNextRequest ( ipRouteDest.10.0.0.99,
                 ipRouteNextHop.10.0.0.99,
                 ipRouteMetric1.10.0.0.99 )
```

As there are no further entries in the table, the SNMP agent returns those objects that are next in the lexicographical ordering of the known object names. This response signals the end of the routing table to the management station.

#### 4.1.4. The GetResponse-PDU

The form of the GetResponse-PDU is identical to that of the GetRequest-PDU except for the indication of the PDU type. In the ASN.1 language:

```
GetResponse-PDU ::=
  [2]
    IMPLICIT SEQUENCE {
      request-id
      RequestID,
```

```

        error-status
            ErrorStatus,

        error-index
            ErrorIndex,

        variable-bindings
            VarBindList
    }

```

The GetResponse-PDU is generated by a protocol entity only upon receipt of the GetRequest-PDU, GetNextRequest-PDU, or SetRequest-PDU, as described elsewhere in this document.

Upon receipt of the GetResponse-PDU, the receiving protocol entity presents its contents to its SNMP application entity.

#### 4.1.5. The SetRequest-PDU

The form of the SetRequest-PDU is identical to that of the GetRequest-PDU except for the indication of the PDU type. In the ASN.1 language:

```

SetRequest-PDU ::=
    [3]
        IMPLICIT SEQUENCE {
            request-id
                RequestID,

            error-status          -- always 0
                ErrorStatus,

            error-index          -- always 0
                ErrorIndex,

            variable-bindings
                VarBindList
        }

```

The SetRequest-PDU is generated by a protocol entity only at the request of its SNMP application entity.

Upon receipt of the SetRequest-PDU, the receiving entity responds according to any applicable rule in the list below:

- (1) If, for any object named in the variable-bindings field,

the object is not available for set operations in the relevant MIB view, then the receiving entity sends to the originator of the received message the GetResponse-PDU of identical form, except that the value of the error-status field is noSuchName, and the value of the error-index field is the index of said object name component in the received message.

- (2) If, for any object named in the variable-bindings field, the contents of the value field does not, according to the ASN.1 language, manifest a type, length, and value that is consistent with that required for the variable, then the receiving entity sends to the originator of the received message the GetResponse-PDU of identical form, except that the value of the error-status field is badValue, and the value of the error-index field is the index of said object name in the received message.
- (3) If the size of the Get Response type message generated as described below would exceed a local limitation, then the receiving entity sends to the originator of the received message the GetResponse-PDU of identical form, except that the value of the error-status field is tooBig, and the value of the error-index field is zero.
- (4) If, for any object named in the variable-bindings field, the value of the named object cannot be altered for reasons not covered by any of the foregoing rules, then the receiving entity sends to the originator of the received message the GetResponse-PDU of identical form, except that the value of the error-status field is genErr and the value of the error-index field is the index of said object name component in the received message.

If none of the foregoing rules apply, then for each object named in the variable-bindings field of the received message, the corresponding value is assigned to the variable. Each variable assignment specified by the SetRequest-PDU should be effected as if simultaneously set with respect to all other assignments specified in the same message.

The receiving entity then sends to the originator of the received message the GetResponse-PDU of identical form except that the value of the error-status field of the generated message is noError and the value of the error-index field is zero.

## 4.1.6. The Trap-PDU

The form of the Trap-PDU is:

```

Trap-PDU ::=
  [4]
    IMPLICIT SEQUENCE {
      enterprise          -- type of object generating
                          -- trap, see sysObjectID in [5]
      OBJECT IDENTIFIER,
      agent-addr          -- address of object generating
      NetworkAddress,    -- trap
      generic-trap        -- generic trap type
      INTEGER {
        coldStart(0),
        warmStart(1),
        linkDown(2),
        linkUp(3),
        authenticationFailure(4),
        egpNeighborLoss(5),
        enterpriseSpecific(6)
      },
      specific-trap      -- specific code, present even
      INTEGER,           -- if generic-trap is not
                          -- enterpriseSpecific
      time-stamp          -- time elapsed between the last
      TimeTicks,         -- (re)initialization of the network
                          -- entity and the generation of the
                          trap
      variable-bindings  -- "interesting" information
      VarBindList
    }

```

The Trap-PDU is generated by a protocol entity only at the request of the SNMP application entity. The means by which an SNMP application entity selects the destination addresses of the SNMP application entities is implementation-specific.

Upon receipt of the Trap-PDU, the receiving protocol entity presents its contents to its SNMP application entity.

The significance of the variable-bindings component of the Trap-PDU is implementation-specific.

Interpretations of the value of the generic-trap field are:

#### 4.1.6.1. The coldStart Trap

A coldStart(0) trap signifies that the sending protocol entity is reinitializing itself such that the agent's configuration or the protocol entity implementation may be altered.

#### 4.1.6.2. The warmStart Trap

A warmStart(1) trap signifies that the sending protocol entity is reinitializing itself such that neither the agent configuration nor the protocol entity implementation is altered.

#### 4.1.6.3. The linkDown Trap

A linkDown(2) trap signifies that the sending protocol entity recognizes a failure in one of the communication links represented in the agent's configuration.

The Trap-PDU of type linkDown contains as the first element of its variable-bindings, the name and value of the ifIndex instance for the affected interface.

#### 4.1.6.4. The linkUp Trap

A linkUp(3) trap signifies that the sending protocol entity recognizes that one of the communication links represented in the agent's configuration has come up.

The Trap-PDU of type linkUp contains as the first element of its variable-bindings, the name and value of the ifIndex instance for the affected interface.

#### 4.1.6.5. The authenticationFailure Trap

An authenticationFailure(4) trap signifies that the sending protocol entity is the addressee of a protocol message that is not properly authenticated. While implementations of the SNMP must be capable of generating this trap, they must also be capable of suppressing the emission of such traps via an implementation-specific mechanism.

#### 4.1.6.6. The egpNeighborLoss Trap

An egpNeighborLoss(5) trap signifies that an EGP neighbor for whom

the sending protocol entity was an EGP peer has been marked down and the peer relationship no longer obtains.

The Trap-PDU of type `egpNeighborLoss` contains as the first element of its variable-bindings, the name and value of the `egpNeighAddr` instance for the affected neighbor.

#### 4.1.6.7. The `enterpriseSpecific` Trap

A `enterpriseSpecific(6)` trap signifies that the sending protocol entity recognizes that some enterprise-specific event has occurred. The `specific-trap` field identifies the particular trap which occurred.



## 5. Definitions

```
RFC1157-SNMP DEFINITIONS ::= BEGIN

IMPORTS
    ObjectName, ObjectSyntax, NetworkAddress, IPAddress, TimeTicks
        FROM RFC1155-SMI;

-- top-level message

Message ::=
    SEQUENCE {
        version          -- version-1 for this RFC
            INTEGER {
                version-1(0)
            },

        community       -- community name
            OCTET STRING,

        data             -- e.g., PDUs if trivial
            ANY          -- authentication is being used
    }

-- protocol data units

PDUs ::=
    CHOICE {
        get-request
            GetRequest-PDU,

        get-next-request
            GetNextRequest-PDU,

        get-response
            GetResponse-PDU,

        set-request
            SetRequest-PDU,

        trap
            Trap-PDU
    }
```

```

-- PDUs

GetRequest-PDU ::=
    [0]
        IMPLICIT PDU

GetNextRequest-PDU ::=
    [1]
        IMPLICIT PDU

GetResponse-PDU ::=
    [2]
        IMPLICIT PDU

SetRequest-PDU ::=
    [3]
        IMPLICIT PDU

PDU ::=
    SEQUENCE {
        request-id
            INTEGER,

        error-status      -- sometimes ignored
            INTEGER {
                noError(0),
                tooBig(1),
                noSuchName(2),
                badValue(3),
                readOnly(4),
                genErr(5)
            },

        error-index      -- sometimes ignored
            INTEGER,

        variable-bindings -- values are sometimes ignored
            VarBindList
    }

Trap-PDU ::=
    [4]
        IMPLICIT SEQUENCE {
            enterprise      -- type of object generating
                            -- trap, see sysObjectID in [5]

            OBJECT IDENTIFIER,

```

```

agent-addr      -- address of object generating
                 NetworkAddress, -- trap

generic-trap    -- generic trap type
  INTEGER {
    coldStart(0),
    warmStart(1),
    linkDown(2),
    linkUp(3),
    authenticationFailure(4),
    egpNeighborLoss(5),
    enterpriseSpecific(6)
  },

specific-trap   -- specific code, present even
  INTEGER,      -- if generic-trap is not
                -- enterpriseSpecific

time-stamp      -- time elapsed between the last
  TimeTicks,    -- (re)initialization of the
                -- network
                -- entity and the generation of the
                -- trap

variable-bindings -- "interesting" information
  VarBindList
}

-- variable bindings

VarBind ::=
  SEQUENCE {
    name
      ObjectName,

    value
      ObjectSyntax
  }

VarBindList ::=
  SEQUENCE OF
  VarBind

END

```

## 6. Acknowledgements

This memo was influenced by the IETF SNMP Extensions working group:

Karl Auerbach, Epilogue Technology  
K. Ramesh Babu, Excelan  
Amatzia Ben-Artzi, 3Com/Bridge  
Lawrence Besaw, Hewlett-Packard  
Jeffrey D. Case, University of Tennessee at Knoxville  
Anthony Chung, Sytek  
James Davidson, The Wollongong Group  
James R. Davin, MIT Laboratory for Computer Science  
Mark S. Fedor, NYSERNet  
Phill Gross, The MITRE Corporation  
Satish Joshi, ACC  
Dan Lynch, Advanced Computing Environments  
Keith McCloghrie, The Wollongong Group  
Marshall T. Rose, The Wollongong Group (chair)  
Greg Satz, cisco  
Martin Lee Schoffstall, Rensselaer Polytechnic Institute  
Wengyik Yeong, NYSERNet

## 7. References

- [1] Cerf, V., "IAB Recommendations for the Development of Internet Network Management Standards", RFC 1052, IAB, April 1988.
- [2] Rose, M., and K. McCloghrie, "Structure and Identification of Management Information for TCP/IP-based internets", RFC 1065, TWG, August 1988.
- [3] McCloghrie, K., and M. Rose, "Management Information Base for Network Management of TCP/IP-based internets", RFC 1066, TWG, August 1988.
- [4] Cerf, V., "Report of the Second Ad Hoc Network Management Review Group", RFC 1109, IAB, August 1989.
- [5] Rose, M., and K. McCloghrie, "Structure and Identification of Management Information for TCP/IP-based Internets", RFC 1155, Performance Systems International and Hughes LAN Systems, May 1990.
- [6] McCloghrie, K., and M. Rose, "Management Information Base for Network Management of TCP/IP-based Internets", RFC 1156, Hughes LAN Systems and Performance Systems International, May 1990.
- [7] Case, J., M. Fedor, M. Schoffstall, and J. Davin, "A Simple Network Management Protocol", Internet Engineering Task Force working note, Network Information Center, SRI International, Menlo Park, California, March 1988.
- [8] Davin, J., J. Case, M. Fedor, and M. Schoffstall, "A Simple Gateway Monitoring Protocol", RFC 1028, Proteon, University of Tennessee at Knoxville, Cornell University, and Rensselaer Polytechnic Institute, November 1987.
- [9] Information processing systems - Open Systems Interconnection, "Specification of Abstract Syntax Notation One (ASN.1)", International Organization for Standardization, International Standard 8824, December 1987.
- [10] Information processing systems - Open Systems Interconnection, "Specification of Basic Encoding Rules for Abstract Notation One (ASN.1)", International

Organization for Standardization, International Standard  
8825, December 1987.

[11] Postel, J., "User Datagram Protocol", RFC 768,  
USC/Information Sciences Institute, November 1980.

#### Security Considerations

Security issues are not discussed in this memo.

#### Authors' Addresses

Jeffrey D. Case  
SNMP Research  
P.O. Box 8593  
Knoxville, TN 37996-4800

Phone: (615) 573-1434

Email: case@CS.UTK.EDU

Mark Fedor  
Performance Systems International  
Rensselaer Technology Park  
125 Jordan Road  
Troy, NY 12180

Phone: (518) 283-8860

Email: fedor@patton.NYSER.NET

Martin Lee Schoffstall  
Performance Systems International  
Rensselaer Technology Park  
165 Jordan Road  
Troy, NY 12180

Phone: (518) 283-8860

Email: schoff@NISC.NYSER.NET

James R. Davin  
MIT Laboratory for Computer Science, NE43-507  
545 Technology Square  
Cambridge, MA 02139

Phone: (617) 253-6020

E-mail: [jrd@ptt.lcs.mit.edu](mailto:jrd@ptt.lcs.mit.edu)