



CHAPTER 1

Introduction

This guide provides formatted information drawn primarily from the Management Information Base (MIB) file written for the server administrator Simple Network Management Protocol (SNMP) subagent. The MIB filename is **10892.mib**.

For each of the variables defined in the MIB, the following fields are specified:

- Variable name
- Object identifier (OID) or unique identifying number
- Description
- Data type of the variable (for example: integer, string, octet string)
- Whether the variable is accessible, not accessible, read-only, or read-write
- Index or indexes, if applicable

For each MIB group that has unique variable definitions, tables are included in the last section of the chapter to explain the meaning of the terms.

Introduction to SNMP

SNMP is a systems management standard originally designed for network management. SNMP manages much more than networks. Information Technology (IT) professionals use SNMP for monitoring and managing computer systems and the various components and peripherals supported by their systems.

SNMP standards are defined by the Internet Engineering Task Force (IETF). SNMP version 1 was published in August 1988 and is the most commonly supported version of SNMP. SNMP version 2 was first published in May 1993, but has not gained widespread market acceptance. SNMP version 3 was recently completed and has addressed security issues that exist in version 1.

All SNMP systems consist of one or more managed nodes that provide data through an SNMP agent to a management station. The management station provides a user interface to view data from the managed nodes. The management station and managed nodes communicate over a network (typically through User Datagram Protocol/Internet Protocol [UDP/IP]).

The management station and a managed node communicate by means of a common data schema. SNMP Management Information Base (MIB) files define the structure, type, and values of the SNMP data. While MIBs can be standardized or enterprise specific, most operating systems supply SNMP agents for the standard MIB-I and MIB-II schemas. MIB-I defines a base set of standard management information for systems implementing the Internet Protocol (IP) suite. MIB-II defines characteristics of the system, characteristics of network interfaces, and characteristics of components of the IP on the system. In addition to the standard MIBs, many hardware vendors have defined MIBs that provide management data specific to their systems and peripheral devices.

Monitored data can be retrieved through SNMP using the Get command. Typically, this command requires the host name or IP address of the target machine as well as the OID of the data to retrieve. Exact details are dependent on the operating system and the development tools being used to create the management application. The Get command has a variant known as GetNext.

SNMP MIB OIDs

Each data class within an MIB is defined by an object identifier (OID). OIDs are unique across all MIBs. An OID consists of a series of digits separated by periods. The OID functions in a similar fashion to a phone number. The phone number 011-512-471-0000 uniquely identifies a single phone. The phone number can be broken down into a number of components to uniquely identify a phone. The first component, 011, is the country code for the United States. The second component, 512, identifies the area code for central Texas. The third component, 471, is the phone exchange for a large state university in the city of Austin. The final component, 0000, is the main switchboard.

There are two main differences between the phone number example and an actual OID. The first difference is that there are many more components in an OID, up to 128. The combination of these components is called an OID prefix. The second difference is that OIDs support the concept of indexes or keys. The OID prefix specifies the data class but does not specify an instance of the data within the class. Indexes can be used to identify the instances of a data class. These indexes are referred to as the OID suffix.

The assignment of values for each OID prefix component can be illustrated by using a tree structure. The following is an example of an OID assignment:

```
ROOT
  CCITT (0)
  ISO (1)
    ORG (3)
      DOD (6)
        INTERNET (1)
          MGMT (2)
            MIB (1)
              EXPERIMENTAL (3)
                PRIVATE (4)
                  ENTERPRISES (1)
                    DELL (674)
                      SNMPv2 (6)
```

In the above example, the OID prefix for the Dell™ enterprise would be 1.3.6.1.4.1.674.

The numbers in boldface type show the categories and numbers that apply to server administrator. All server administrator-defined OIDs consist of 1.3.6.1.4.1.674 followed by additional component values.

SNMP Security

SNMP version 1 has a very limited security mechanism. SNMP agents support the use of a community string, which is configured at each SNMP agent and is passed as a part of all SNMP request messages. There is no verification that the requester is actually a member of the specified community.

Because most system and network management data is not confidential, this limited security is acceptable for Get types of requests. On the other hand, this security is not acceptable for Set types of operations where an SNMP request could power off a system, reconfigure a redundant array of independent disks (RAID) card, and so on. Some vendors have chosen not to support SNMP Set operations for this reason. Server Administrator is able to support SNMP Set operations because its SNMP agents implement a hash/digest mechanism to prevent unauthorized SNMP Set operations. One limitation of this practice is that only server administrator-developed SNMP management applications have the capability to support the hash/digest mechanism.

Initiating Management Actions

Management actions can be performed using the SNMP Set command. These actions can consist of configuring a phone number for the system's owner, rebooting a system, or changing the asset tag of the system. See the previous section, "SNMP Security," for limitations on Set operations.

SNMP Traps

SNMP is frequently used to monitor systems for fault conditions such as temperature violations, hard drive failures, and so on. Management applications can monitor for these conditions by polling the appropriate OIDs with the Get command and analyzing the returned data. This method has its drawbacks. If it is done frequently, significant amounts of network bandwidth can be consumed. If it is done infrequently, the response to the fault condition may not occur in a timely fashion. SNMP traps avoid these limitations of the polling method.

An SNMP trap is an asynchronous event indicating that something significant has occurred. This is analogous to a pager receiving an important message, except that the SNMP trap frequently contains all the information needed to diagnose a fault.

Two drawbacks to SNMP traps are that they are sent using User Datagram Protocol (UDP), which is not a guaranteed delivery mechanism, and that they are not acknowledged by the receiver.

An SNMP trap message contains the trap's enterprise OID, the agent IP address, a generic trap ID, the specific trap ID, a time stamp, and zero or more variable bindings (varbinds). The combination of an enterprise OID and a specific trap ID uniquely identifies each server administrator-defined trap. A varbind consists of an OID and its value and provides additional information about the trap.

In order for a management station to receive SNMP traps from a managed node, the node must be configured to send traps to the management station. Trap destination configuration is dependent on the operating system. When this configuration is done, a management application on the management station can wait for traps and act on them when received.

For a list of traps supported by the server administrator SNMP subagent, see Chapter 20, "Traps."