

SNMP: the Simple Network Management Protocol

The SNMP protocol has currently three versions. They are all in use.

RFC 2570 defines SNMPv3

RFC 1157, or STD 15 defines SNMPv2.

The protocol is intended to manage components. There are many types of management:

- Availability management
- Capacity management
- Resource management
- Change management

Information model

SNMP provides a simple model to address these. All the information that can be managed is defined in a tree structure. That tree is represented in the Management Information Base (MIB). For each device such a tree can be defined and published. The nodes in a tree are objects. All SNMP objects have an object identifier. They are globally unique.

A MIB is created by the organization and described formally using Abstract Syntax Notation One (ASN.1). This is the de-facto standard before XML.

Each vendor of hardware or software that wants to support SNMP must make himself part of the Vendor subtree.

Every vendor can request an object identifier for its organization. The names are extended to the right to form new object identifiers: kuleuven.5.2.1 etc. The organization can structure its MIB any way it wants.

Object identifier is hierarchical, it resembles tree walking. The name “iso org dod internet mgmt mib system sysDescr” gets translated into the object identifier “1 3 6 1 2 1 1 1”.

Information access

The data in a MIB can be read and written. The SNMP protocol supports these functions, which are called set/get methods on these objects, much like Java Bean properties.

Example ASN.1 from the RFC

```
-- 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
```

Agents

On the managed system, there is an agent that will respond to SNMP requests.

It receives requests and sends answers The agent defines the actions it will take for a get or set. Just like with Java bean methods, the set and get may involve a complex computation rather than a simple reading or writing of a storage cell. Writing to an object and reading that value back may not give the expected result.

Simple objects may give the number of bytes sent, the mean value of the CPU load during the last 5 minutes, free disk space, etc.

Authentication

Access to the SNMP MIB in set or get mode can have a serious impact on the system. In many cases, such actions can only be permitted by authorized users. SNMP has been slow in providing a secure solution. The basic authentication of requests is via a community string: a secret that can be considered a combination of userID and password. The protocol is UDP, hence network level authentication is very weak. There is no encryption, hence interception of the community string is easy.

Management stations

At the server side the management uses management stations. They vary from simple GUI to complex systems with autodiscovery etc.

Some well known solutions are: HP openview, SUN Solstice, IBM Tivoli (Netview), CA Unicenter.

To experiment with SNMP, one can use very simple SNMP tree browsers.