~DRAFT~

Chapter 11. Network Management

Table of Contents

1. **Context** 1
2. **Introduction** 2
3. **Objectives** 2
4. **What is network management?** 2
5. **Examples of network management activities** 3
6. **ISO network management model** 4
7. **Problem management** 5
8. **Implementing problem management** 5
9. **Network management: an illustrative example** 5
10. **An architecture for network management** 6
    - 10.1. **The hardware point of view** 6
    - 10.2. **The software point of view** 7
11. **Simple Network Management Protocol (SNMP)** 8
    - 11.1. **Components of SNMP** 9
    - 11.2. **SNMP Management Information Base (MIB)** 10
    - 11.3. **SNMP Protocol** 11
12. **Directory-Enabled Networking management** 12
    - 12.1. **Directory** 13
    - 12.2. **Directory service** 13
    - 12.3. **Benefits of Using Directory-Enabled Networking Management** 13
13. **Activities** 14
14. **Review Questions** 14
15. **Answers and Comments** 16

1. Context

Networks of all the types that have been dealt with so far in this module need to be managed if they are to operate efficiently. In this way, this unit relates to all the previous units. In addition, network management is needed to ensure that networked applications such as e-mail and social networks sites make best use of the networks over which they operate.

2. Introduction

A general account of the overall architecture of a network management system is described. Network monitoring, the collection of information on the status of the network, the consolidation and interpretation of this information, and acting on the resulting interpretation are distinguished as the major activities of a network management system. The way in which these activities are realised, the co-ordination of the activities, and their embedding in one coherent network management system are all treated. Various examples are given of the type of network problem that a network management system is responsible for. The way in which a generic management system deals with these problems is then illustrated.

The international Organization for Standardization (ISO) network management framework for classification of network management activities into five broad problem areas, that is, fault management, performance management, configuration management, accounting management and security management, is presented.

The idea of an 'alert' is then introduced, and it is then shown how generic problem management spanning all these categories may be achieved both manually and automatically on the basis of a process that is driven by the use of alerts.

Finally, the two main standards for network management are then described, and the ways in which they relate to this framework can be seen. The standards described are the Internet's Simple Network Management Protocol (SNMP), and the Directory-Enabled Networking management (DENM).

3. Objectives

At the end of this module, you should be able to:

- rehearse the concepts and concerns of network management;
- classify network management activities;
- provide examples of network management activity;
- outline the principles underlying the implementation of network management systems.
- introduce and contrast two views of an architecture for network management;
- describe and contrast the two main standards for network management.

4. What is network management?

In parallel with this topic, you should read Chapter 9 of James F. Kurose and Keith W. Ross, " Networking: a top-down approach", (6th edn.).

When a network is seen as a single entity, that is, as one collection of resources, two views of network management can be offered.

The first construes network management as the monitoring of the activities of a network, following which the information collected is consolidated to present an overview of the state of the network to a human operator who can then take any actions necessary to improve the operation and performance of the network.

The second view understands network management as a completely automated process in which a network is monitored, and the information obtained is consolidated and then presented to an entity

capable of automatically determining from it the actions necessary to adjust and adapt the network so that it maintains the required levels of operation and performance.

These views can be illustrated as follows:

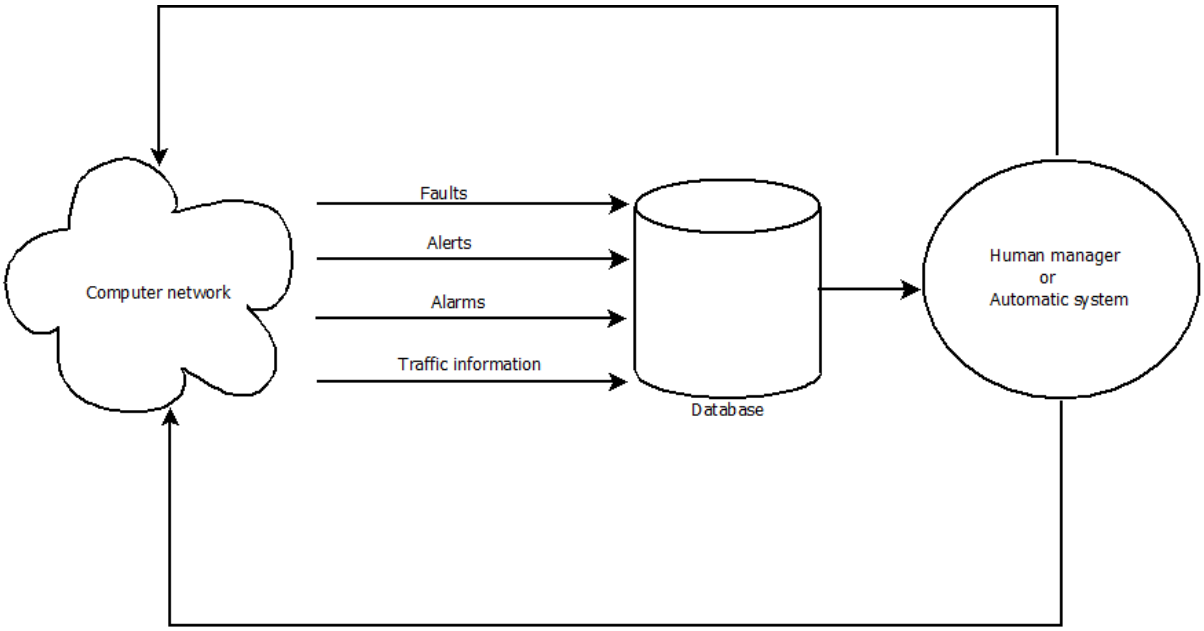

5. Examples of network management activities

The following list give some examples of activities that are typical of those undertaken by a network administrator.

1. Monitoring the computers and links of a network to detect any faults that may develop in them.
2. Monitoring the message queues at each computer to ensure efficient use of network resources and prevent message loss.

   Example: The illustrated network segment contains two printers. When one of them is busy, which is apparent when its queue of waiting jobs becomes full, computers with a job to be printed will be directed to send it to the other printer

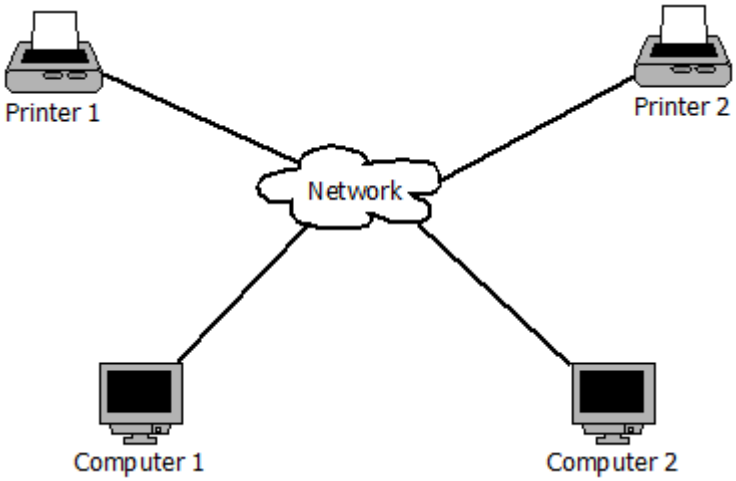

3. Controlling the routing of messages through the network. This is a way of routing round a fault detected in the network. It also provides a way of avoiding the busy parts of a network where a message would be delayed.

   Example: When computer 4 is busy, the routing table at computer 2 can be changed to route messages for computer 3 via computer 1.

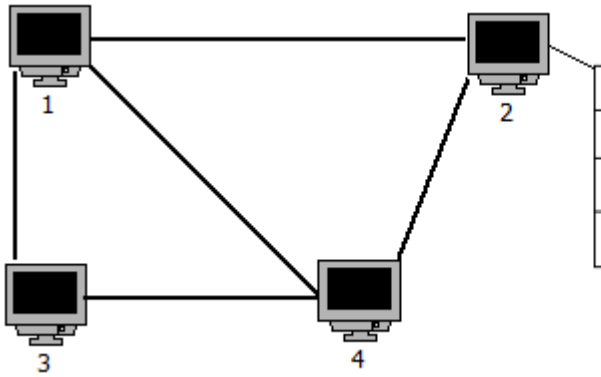

dest.	send to
1	1
3	4
4	4

4. Detecting network intrusion. The network administrator may want to be alerted on suspicious network traffic.

6. ISO network management model

We have touched on some network management activities in the previous section. The international Organization for Standardization (ISO) has created a network management model that assigns the network management activities to one of the following problem areas:

Fault management: The concerns of fault management are the detection of faults, and covering up for them until they can be repaired, after which the network can be returned to its original state. Faults can be covered up by working around them or by providing some alternative that is functionally equivalent to the failed unit.

Configuration management: The configuration of the network is the arrangement of its computers and links, that is, its topology. Configuration management deals with changes to the configuration of the network caused by the addition and removal of computers and links.

Performance management: Managing the performance of a network involves maintaining an acceptable quality of service for all its users. This usually involves the delivery of messages within some specified time. Occasionally, acceptable levels of service may be expressed in terms of levels of reliability. Acceptable delivery times can be achieved by managing the flows of traffic across the network and, as far as possible, preventing the build-up of congestion. Reliability can be ensured by invoking the necessary measures against message impairment and loss.

Accounting management: The responsibilities of accounting management are to keep track of network usage and, correspondingly, to generate bills for the users of the network.

Security management: The concern of security management is to ensure security both for the network itself and for the users of the network. The security of the network can be ensured by allowing access only to authorized users, and by ensuring that those with access do not use it improperly. The users of the network can be given the levels of security they need by providing the appropriate security services.

7. Problem management

The types of problem management listed above can be managed by adopting the following procedure:

1. **Determine that there is a problem**. Much of the network monitoring activity is intended to determine whether network components are operating correctly. Components can send an alert if they detect a problem in their own functioning.
2. **Diagnose the problem**. Find out exactly what the problem is either directly if it is possible or by collecting evidence and deducing what the problem is.
3. **By-pass the problem**. Work round the problem to ensure that the rest of the network is not affected by it.
4. **Resolve the problem**. Fix the problem, and then return the network to its original state.
5. **Keep a record**. Store an account of the problem, so that there is a record. The complete log of problems may well contain valuable information about the running of the network, revealing the presence of persistent problems, recording ways of handling problems, and so on.

8. Implementing problem management

The problems in a network can be determined in two basic styles. The network can be actively monitored by making regular observations of all its components, or the components of the network can send a report when they detect a problem either in their own working or in the working of another component. Given the size and complexity of many networks, it is not surprising that the latter approach is the more common.

Problem management, then, is often done with the use of 'alerts'. An alert is an unsolicited message generated within the network, and sent to the Network Management Centre to alert it to something needing its attention. An alert usually takes the following form:

Typical alerts could take the form:

alert, computer_1, line-6

(Computer 1 complaining about one of its links.)

alert, computer_25, neighbouring_computer

(Computer 25 complaining about a neighbouring computer.)

Note that alerts will usually arrive at the Network Management Centre in clusters. In the first example above, line_6 has computer_1 at one end, but it will have another computer at the other end and, if the line really is faulty, the other line will detect it and will also send an alert sooner or later. In the second example, the faulty neighbour of computer_25 will also cause a cluster of alerts to be sent in due course.

9. Network management: an illustrative example

This section aims to present a simple network management scenario that brings together what we have learnt so far.

Part of a network that is under management is shown in the diagram below.

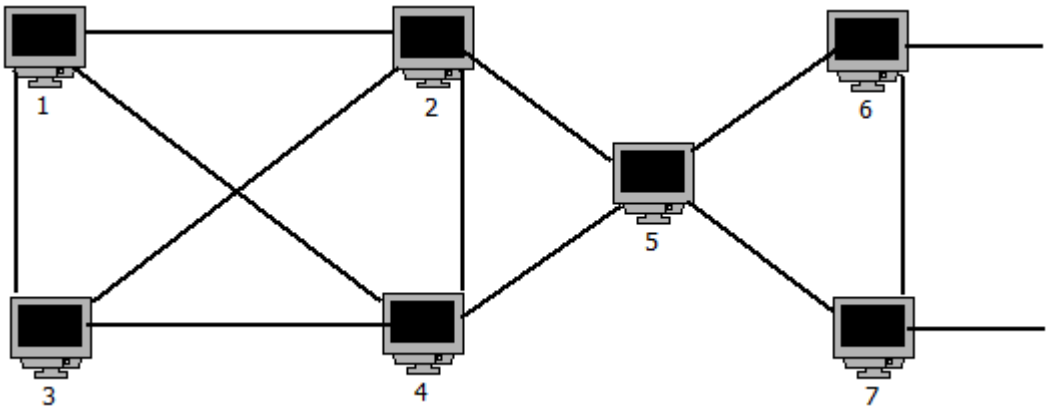

The basis of its management is the use of alerts, which are sent to the Network Management Centre (NMC) and take the form:

alert, <sending_computer>, <symptom>

These alerts do not pinpoint the precise nature of the problem, but provide a symptom or an indication of the problem. The Network Management Centre must diagnose the actual problem by making deductions from the information provide by a cluster of alerts.

The matter that concerns us here is that a certain number of errors inevitably occur on the links of a network. An increase in the error rate of a link is often a sign that the link is about to fail. For this reason nodes often monitor the error rates on their incoming links. On this network, if the error rate on a link reaches a dangerously high level, any computer attached to that link is required to send an alert. When the sending computer is computer_3, the alert takes the form:

alert, computer_3, high_error_rate

Stage 1: The alert arrives, indicating that some problem exists. Suppose it is the alert given above. But we know that alerts arrive in clusters and, assuming that it is the line between computer_3 and computer_2 that is causing the problem, the next alert to arrive will be:

alert, computer_2, high_error_rate

Stage 2: This stage is now reached. From the cluster of two alerts, the NMC can deduce that the link between computer_3 and computer_2 may be about to fail.

Stage 3: The management system must now work around this link and also send a technician to fix or replace the link.

Stage 4: When the link has been repaired, the management system can restore the network to its original state.

Stage 5: An account of the problem and the problem management activity is recorded in the log.

10. An architecture for network management

10.1. The hardware point of view

An architectural conception of network management can be based on the separation of the network hardware into a network under management and a 'shadow' network for management. The

separation is based on the nature of the traffic, with one network carrying the information-bearing traffic and the other carrying the management traffic. The situation is illustrated below:

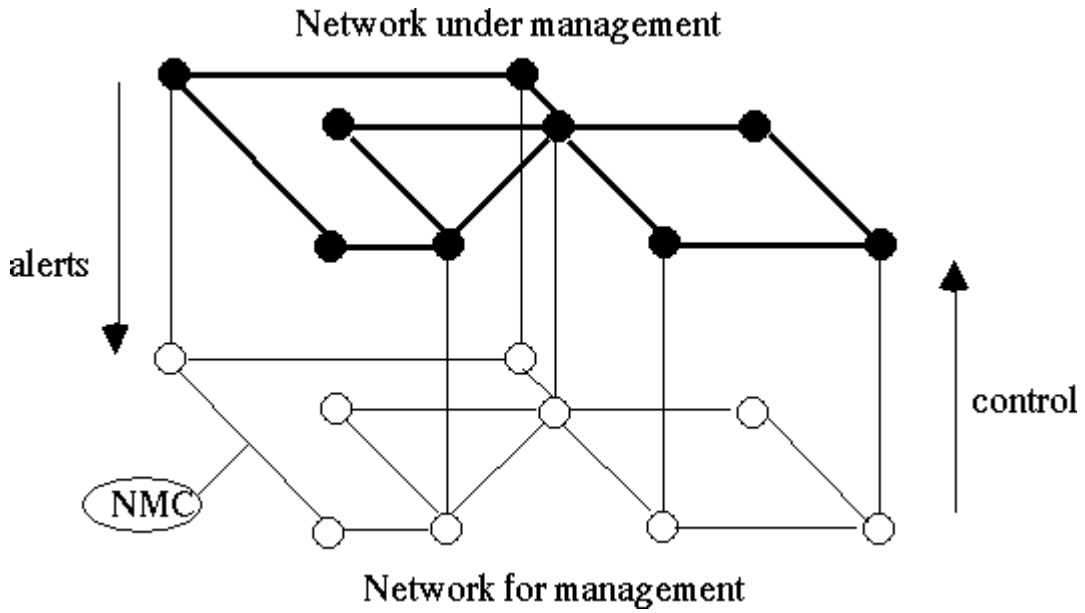

The network under management is 'shadowed' by a network for management with identical structure.

Alerts are raised by components on the network under management, but flow directly to the shadow network, where they are routed to the Network Management Centre. At the Network Management Centre, the information in the alerts is consolidated and, in one way or another, is used to determine the necessary control actions. Return messages carrying this information are routed through the shadow network to the shadow computer with a link to the network under management nearest to the intended destination computer. In this way, the management traffic is retained by the network for management for as long as possible.

10.2. The software point of view

Now consider the software-driven management traffic. The alerts that make up this traffic can be metaphorically considered as an example of 'chatter' in that the network components chatter about their own state, the state of their neighbours or that of some other part of the network, essentially for the overall benefit of the network in much the same way as the people in a community chatter about themselves, their neighbours and others in the community for the general benefit of the community.

The network chatter is all addressed to the Network Management Centre. It collects and consolidates the information conveyed by the chatter, and in this way builds up a picture of the overall state of the network. This overall picture should enable it to determine any problems that need to be dealt with so that the network will run smoothly in all circumstances, and provide essential services to both its users and its operator.

The generation and transmission of alerts as they are needed is the basis of one style of chatter to provide the information needed for problem management. Other styles of chatter are possible, as may be seen by considering the issues that arise with any style of chatter. They include:

- what to chatter about,
- when to chatter,

- how to chatter, and
- who to chatter to.

10.2.1. What to chatter about?

The purpose of chatter is to provide the Network Management Centre with sufficient information about the state of the network to be able to manage the network. With this in mind, we can think in more detail about the issues listed above for each of the categories of problem management. With fault management, for example, if the network computers chatter about the state of their incoming links and neighbouring computers, including whether they are faulty, then the Network Management Centre will receive enough information to identify any faults that have occurred in the network.

We have seen that sending alerts whenever a fault is detected is one way of implementing this scheme

10.2.2. When to chatter?

One way of answering the question of when to chatter is, as we have seen, to use alerts, and alerts are sent only when it is necessary to do so. In fact, up to this point, we have relied entirely on the use of alerts.

When alerts are used to draw attention to the existence of a problem, they carry information about something that has already happened. For this reason, they lend themselves to a style of management that is necessarily reactive in that it can only operate by reacting to events that have already occurred. This style of management can be replaced by one that is rather less concerned with 'fire-fighting', and a little more capable of preventing problems from developing, by allowing alerts to draw attention to the symptoms of a problem rather than to the problem itself. It is, however, also possible to develop a quite different style of management simply by giving a different answer to the question of when to chatter.

One way to find a different answer to the issue of when to chatter is to go back to the analogy between the components of a network and the members of a community, and to note that some of the people in a community chatter all the time (whereas the more taciturn chatter only when they need to). So, one alternative answer to the question is that network components can chatter all the time.

10.2.3. How to chatter?

Chatter is, like any other form of information exchange on a network, carried out by sending messages. When alerts are used, the message is formatted to identify the sending computer and a problem that is known to exist or is suspected to exist. When chatter goes on all the time, rather than only when needed, the message must provide a status report. The report should contain information to assist each type of problem management, so that it ought to be formatted into five main sections.

11. Simple Network Management Protocol (SNMP)

For more information on SNMP, read notes from
http://docwiki.cisco.com/wiki/Simple_Network_Management_Protocol

Simple Network Management Protocol (SNMP) is a protocol that is used to exchange management information between network devices. It allows network administrators to monitor network performance. Two versions of the protocol exists, namely: SNMP version 1 (SNMPv1) and SNMP version 2 (SNMP v2).

The diagram shows a network managed by the SNMP protocol.

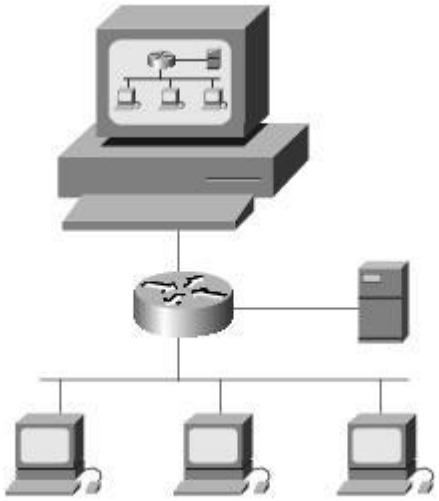

11.1. Components of SNMP

SNMP has three components which are known as:

1. **Managed device**: These are typically bridges, routers, host computers and other network components that can send alerts when necessary and respond to enquiries from the Network Management Station.
2. **Agent**: An agent is a network-management software module that resides in a managed device. This contains information about the network components and their status, recorded in a consistent way compatible with SNMP.
3. **The Network Management Systems (NMSs):** This monitors and controls the managed devices which are the network components that can send alerts and respond to enquiries about their status. It also maintains the Management Information Base (read 'SNMP Network Management Base' section below for more information).

The diagram below shows the relationship between the three components of SNMP.

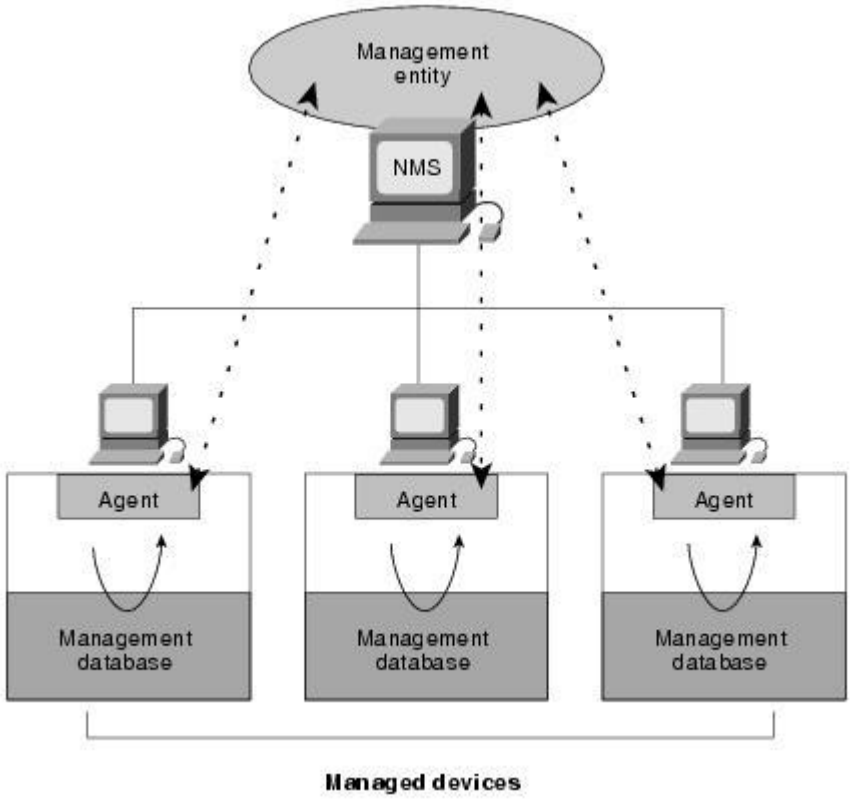

11.2. SNMP Management Information Base (MIB)

Management information base is a database that contains the information obtained from the Network Agents and can be queried by the network manager. The information comprise of managed objects and are identified by object identifiers. A managed object is one of any number of specific characteristics of a managed device. These characteristics may include, basic identification data about a particular piece of hardware, management information about the device's network interfaces and protocols. An object identifier uniquely identifies a managed object. The values of managed objects collectively reflect the current "state" of the network.

International Organization for Standardization (ISO) has proposed a hierarchical framework for naming every possible managed object in MIB (See the diagram below). Note that each branch point in the tree has both a name and a number (shown in parentheses); any point in the tree is thus identifiable by the sequence of names or numbers that specify the path from the root to that point in the identifier tree.

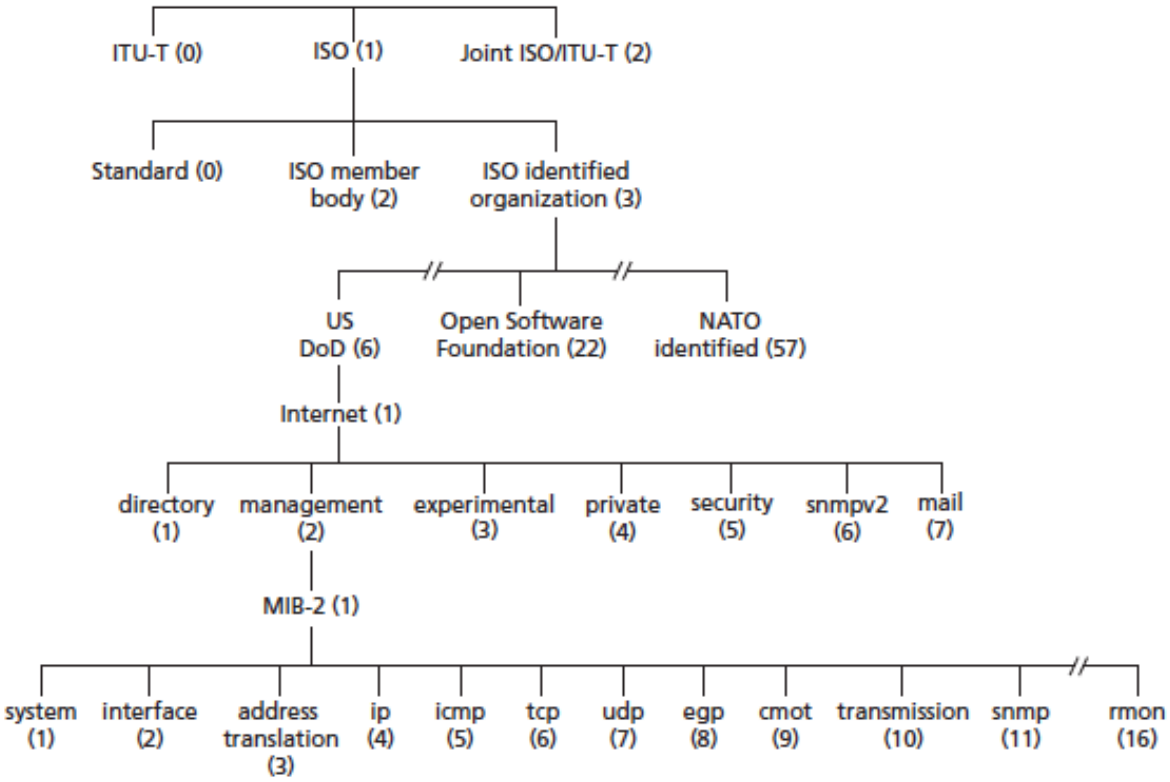

For example: Imagine we have a managed object udpInDatagrams which is a child branch of udp (7). Starting from the root branch, udpInDatagrams can be identified as shown below.

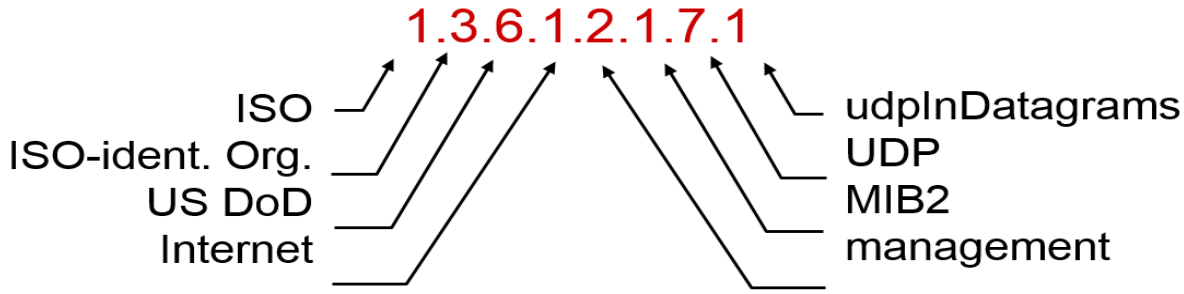

11.3. SNMP Protocol

SNMP supports three commands used to monitor and manage network devices, namely:

1. **Read command:** used to monitor managed devices.
2. **Write command:** used to control managed devices.
3. **Trap command:** used by managed devices to asynchronously report events to the Network Management Systems.

SNMP must account for and adjust to incompatibilities between managed devices. Different computers use different data representation techniques, which can compromise the capability of SNMP to exchange information between managed devices. SNMP uses a subset of Abstract Syntax Notation One (ASN.1) to define message format used to exchange information between diverse systems. SNMPv2 messages consist of a header and a Protocol Data Unit (PDU).

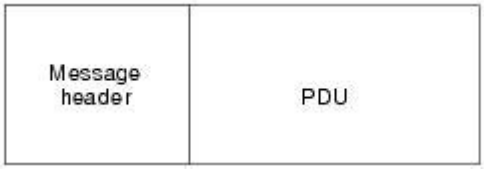

SNMPv2 Message Header

SNMPv2 message headers contain two fields: Version Number and Community Name.

Version number: Specifies the version of SNMP that is being used.

Community name: Defines an access environment for a group of NMSs. NMSs within the community are said to exist within the same administrative domain. Community names serve as a weak form of authentication because devices that do not know the proper community name are precluded from SNMP operations.

SNMPv2 Protocol Data Unit (PDU)

SNMPv2 specifies two PDU formats, depending on the SNMP protocol operation. SNMPv2 PDU fields are variable in length, as prescribed by Abstract Syntax Notation One (ASN.1).

The following descriptions summarize the fields illustrated in Figure below.

PDU type: Identifies the type of PDU transmitted (Get, GetNext, Inform, Response, Set, or Trap).

Request ID: Associates SNMP requests with responses.

Error status: Indicates one of a number of errors and error types. Only the response operation sets this field. Other operations set this field to zero.

Error index: Associates an error with a particular object instance. Only the response operation sets this field. Other operations set this field to zero.

Variable bindings: Serves as the data field of the SNMPv2 PDU. Each variable binding associates a particular object instance with its current value (with the exception of Get and GetNext requests, for which the value is ignored).

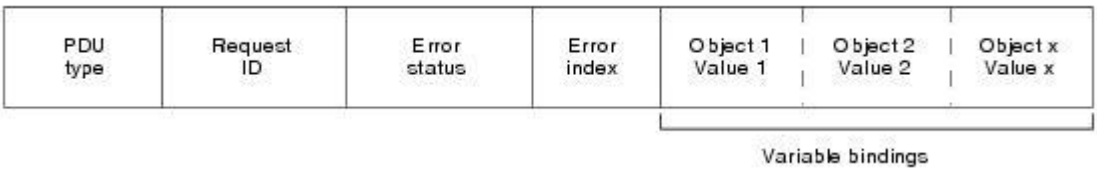

12. Directory-Enabled Networking management

For more information on Directory-Enabled Networking management, read notes from http://docwiki.cisco.com/wiki/Directory-Enabled_Networking#Figure:_Sampling_of_Important_Base_Classes_of_the_DEN_Schema.

Network-management protocols such as SNMP described in the previous section, describe and store the dynamic state of network elements. Apart from the dynamic state, network elements also have the persistent state. The Directory-Enabled Networking (DEN) specification define a standard schema for storing persistent state of network elements and an information model for describing the relationships among objects representing users, applications, network elements, and network

services. SNMP is used to talk to the network elements. DEN is used to talk about network elements and services.

12.1. Directory

Directory-Enabled Networking management uses a directory to manage the entire network environment. A directory is used to record information about users, applications, and network resources such as file servers and printers.

The objects stored in a directory are organized in a hierarchical fashion as shown in the diagram below.

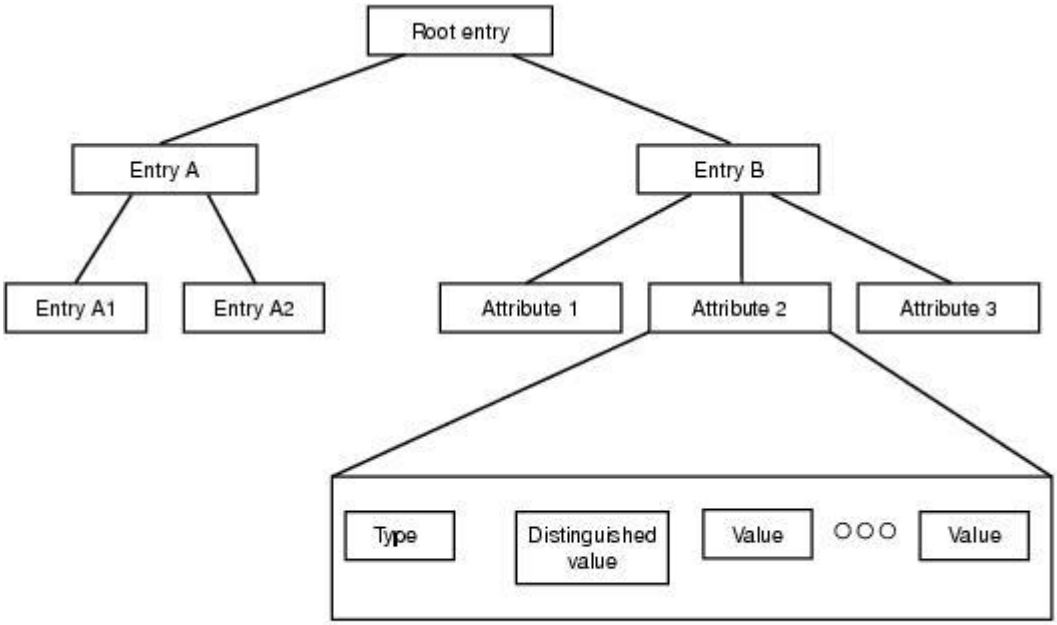

12.2. Directory service

A directory service is a physically distributed, logically centralized repository of infrequently changing data that is used to manage the entire environment. A directory service stores and retrieves information from the directory on behalf of one or more authorized users.

The integration of the network infrastructure with the directory service allows the applications and users to discover the existence of devices and relationships by querying the directory service rather than contacting the individual devices and aggregating the results. Exposing network elements in the directory enhances manageability and usability while reducing the load on the network. The end-user and administrator experience is enhanced because there is a single, authoritative place to obtain the information of interest.

12.3. Benefits of Using Directory-Enabled Networking Management

1. Simplify device configuration. Network devices offers more functionality resulting in increasingly complex device configurations.
2. Control the management and provisioning of network devices through the use of policies. DEN map service-level agreements and business rules to a common set of policies. These policies control the allocation of network resources based on user, subnet, time-of-day, or other appropriate factors.
3. Define a means to make applications more network-aware and to make the network more application-aware.

13. Activities

You can find a discussion of these activities at the end of the chapter.

Activity 1 - problem log

Investigate the kinds of information that could be mined from a network's problem log.

Activity 2 - routing tables

When a performance management strategy of spreading traffic evenly across the network is adopted, the routing tables need to be continually changed to adapt to the state of the traffic in the network. Describe how the information that is needed to determine the appropriate changes to the routing tables may be obtained.

Activity 3 - computer failure

The other concern of the network management system is to deal with the problem of computers that either fail or are disconnected from the network without authorisation. One way to detect such an occurrence is for the computers on a network to monitor their neighbours and note the time since a message was last received from each neighbouring computer. It might then be reasonably assumed that a computer has failed if no message has been received from it after some sufficiently long time, T. When this happens, a monitoring computer sends an alert to the NMC. If computer_2 has received no messages from one of its neighbours in the assigned time, the alert would take the form:

alert, computer_2, no_messages.

Devise a means by which the NMC may detect failed computers. Again, describe the complete management process.

14. Review Questions

You can find an answer/comment for this review questions at the end of the chapter.

Review Question 1

The network management scenario can be described as the network-under-management connected to a Network Management Centre. Information flows from the network to the Network Management Centre. How can this information be characterised in its totality?

Review Question 2

Following on from Question 1, during network management what flows from the Network

Management Centre to the network?

Review Question 3

Given that the information flowing from the network to the Network Management Centre is consolidated in a database, what is the essential difference in what happens during manual network management and automatic network management?

Review Question 4

How can a network management system deal with the situation that arises when a link in a network develops a fault?

Review Question 5

What would happen if a message were sent to a computer with a full input queue? How could this be prevented?

Review Question 6

Devise a routing table for computer 3 in the little network illustrated in ‘Examples of network management activity’ section. Change the table in a way that allows it to respond to the situation when computer acting as the intermediary between computers 3 and 2 is busy.

Review Question 7

How could the fault management system deal with a faulty link in such a way that the fault is not apparent to the users of the network? How could it deal with a failure to a computer providing a specific resource?

Review Question 8

What would be a suitable performance management strategy for a network that aimed to treat all its users equally, and to deliver the messages offered to it as quickly as possible at all times?

Review Question 9

What kind of thing could happen to a network that did not take care of its own security because it failed to provide security management?

Review Question 10

To clarify that the classes of network management correspond to problem management, cast them all in a form such as: 'fault management is the management of the problems that arise in diagnosing and dealing with faults'.

Review Question 11

What is an alert? Why is one sent? Where is it sent to?

Review Question 12

Refer to the ‘Implementing problem management’ section above. After the alert:

alert, computer_25, neighbouring_computer

has been sent, a cluster of associated alerts will also be sent. From where will they originate?

Review Question 13

A network under management could be shadowed by a separate network for management, but this would be an expensive way to support network management. More usually there is a single network which can, conceptually, be separated into two. Describe the way in which a single (actual) network can be conceptually separated into a (virtual) network under management and a (virtual) network for management.

Review Question 14

What do network components chatter about? When do they chatter? How do they chatter? Who do they chatter to?

Review Question 15

Identify the components of SNMP with the items described during the general treatment of the architecture for network management systems.

Review Question 16

Name three of the seven fields of the SNMP v2 GETBULK.

Review Question 17

Name some of the important benefits of DEN.

15. Answers and Comments

Activity 1

It is possible to locate problems that occur consistently, either at a particular time or in a particular place. It is also possible to detect patterns of problems: this may show that some problem is not, in fact, a problem in itself but the consequence of a causal chain of other problems.

Activity 2

The network must be monitored to determine the pattern of the traffic within it. The changes to the routing tables can then be found by determining how they must act to change the current state to the desired state.

Activity 3

The stages are as follows:

Stage 1 is that an alert arrives, indicating that some problem exists.

Suppose it is the alert given above. It will be followed by a cluster of other alerts. These will come from the other neighbours of the problem computer. The other alerts might be:

alert, computer_1, no_messages.

alert, computer_4, no_messages.

Stage 2 is now reached. From the cluster of alerts, the NMC can find the computer that has as its neighbours computer_2, computer_1 and computer_4, and deduce that the problem is with computer_3. (The NMC will have a description of the topology of the network, so that it is in a position to make this deduction.)

Stage 3. The management system must now work around this computer, perhaps find another computer offering the same resources, and also send a technician to fix or reconnect it.

Stage 4. When the computer has been fixed and reconnected, the management system can restore the network to its original state.

Stage 5. An account of the problem and the problem management activity is recorded in the log.

Review Question 1

The information flowing from the network to the Network Management Centre is information about the state of the network.

Review Question 2

Control signals, that is, signals capable of changing or correcting the state of the network as required.

Review Question 3

With a manual system, the person acting as network manager will interrogate the database, while with an automatic system an expert system or some similar item of software will take actions based on the contents of the database.

Review Question 4

By ensuring that no traffic is sent down that link, and by finding an alternative route for any traffic that would have otherwise been directed over that link.

Review Question 5

If a message is sent to a computer with nowhere to store it, the message will inevitably be lost. This can be prevented by monitoring the state of the queues at each computer so as to prevent messages being sent to any computer with a full queue.

Review Question 6

The routing table could be:

Destination	Forward to
1	1
2	1
4	4

The amended routing table would be

Destination	Forward to
1	1
2	4
4	4

Review Question 7

A link failure can be covered up by finding and using an alternative route to any route that involves the failed link. If a computer providing a resource fails, this can be covered up by locating another computer providing the same resource and directing resource demands to it.

Review Question 8

One strategy would be for the network to continually try to spread the traffic evenly across the network, so that there were no congested regions and no inequitable delays to messages held up by congestion.

Review Question 9

The network could damage itself to the point that it could no longer operate by, for example, allowing the messages it carried to cause it to overwrite its own operating software.

Review Question 10

Configuration management is the management of the problems that arise as a result of changing the configuration of the network.

Performance management is the management of the problems that arise as a result of trying to maintain the performance of the network at some assigned level.

Accounting management is the management of the problems that need to be tackled in order to monitor customers' network usage and generate their bills.

Security management is the management of the problems that arise in the course of maintaining the security of the network.

Review Question 11

An alert is a way of drawing attention to a problem. It is sent because a problem has occurred. It is sent to the Network Management Centre.

Review Question 12

Alerts will come from the other neighbours of the computer that computer_25 is complaining about.

Review Question 13

The actual network carrying its user-presented traffic can be regarded as the network under management, while the actual network carrying the traffic for management can be regarded as the network for management. Thus the separation is in terms of the function of the traffic carried. The structure of the two virtual networks is clearly the same.

Review Question 14

They chatter about the state of the network. As described so far in this and the previous unit, they chatter when they need to; they chatter by sending alerts; they chatter to the Network Management Centre.

Review Question 15

The Network Management Station is the computer dedicated to network management at the Network Management Centre. It is used by a person. It contains or controls the MIB, which is the database in which the information recording the status of the network resides. The Network Agents are the network components capable of sending an alert.

Review Question 16

PDU Type, Request ID, Nonrepeaters, Max Repetitions, Variable Bindings (the variable bindings consists of variable object fields that make up the three remaining fields).

Review Question 17

First and foremost, DEN is an object-oriented information model that describes different components of a managed environment in a common way. This enables a close relationship to be

established between classes that define network elements, and services and classes that define other objects. This is the primary mechanism used to define which network services a client needs.

Second, DEN is object-oriented, so it is inherently extensible. This means that concepts not yet defined in DEN can be easily modeled and added to the DEN standard.

Third, DEN enables the application developer as well as the network designer to think of the network as a provider of intelligent services. This enables application developers to describe the functions and treatment that the traffic of their applications requires in terms that the network can represent directly. Thus, if a certain application has specific jitter and latency requirements, DEN can be used to define the set of services that together meet these requirements.

Fourth, and closely related, DEN enables businesses to prioritize the treatment of different applications that are vying for network resources. This enables a business administrator to write a policy that says that SAP and PeopleSoft applications should get preferential treatment over FTP traffic. This enables the network to be designed to treat the applications that a business runs according to the business rules of that organization.

A final example benefit of DEN (although there are more) is that DEN is a standard. This means that it can be used by network vendors, system integrators, and others to define a common framework to describe, define, share, and reuse data.