CHAPTER 3

Wholesale Dial NMS Case Study

The chapter presents a case study illustrating a network management system designed to meet the requirements of a wholesale dial network. The design presented here uses components intended to ensure that each element of the FCAPS model is addressed. In addition, the NMS implementation featured in this case study addresses the requirement for a high level of redundancy and availability.

This case study illustrates the application of several popular industry NMS software packages as well as Cisco element management systems. These tools combine to provide a complete NMS architecture. When managing large networks (as in this case study), it can be advantageous to construct a distributed NMS solution—instead of a centralized solution. Distributing functions allows SNMP polling to be more local to each managed device, spreads system resources (thereby reducing the threat of system overload), and makes the overall NMS environment more resilient.

The remainder of this chapter consists of the following sections:

- The Wholesale Dial NMS Architecture
- Applying the FCAPS Model
- Server Recommendations and Configurations
- System Fail-over / Backup Configuration for Collection Stations

The Wholesale Dial NMS Architecture

This case study presents a wholesale dial NMS environment in the context of the following key elements:

- Management Stations and Collection Stations
- Distributed Network Management Systems

The sections that following summarize each of these elements and general considerations for managing a wholesale dial environment.

Management Stations and Collection Stations

A distributed NMS environment consists of two basic systems:

- Management Station (MS)—A management station relies on the collection station to manage a specific group of nodes. This group of nodes is called a *management domain* (MD). A management station provides an end-to-end view of the network status.

- Collection Station (CS)—A collection station monitors its portion of the network and relays the status to its management station. The collection station performs status monitoring, event handling (traps), and health checking.

Each management station shares its information with the other management stations, providing a complete view of the wholesale dial network.

Note A management station can also be a collection station depending how the architecture is designed. However, the details of such design considerations are beyond the scope of this document.

A management station receives information from collection stations that are actively receiving traps and performing status polling.

Based on the network design, each collection station is responsible for managing a set of access servers and support equipment. The number of access servers and support equipment associated with a collection station generally should not exceed 600 devices, but can based upon the resources of the collection station. The hierarchical placement of a collection station's associated management station can be based on geographical region or organizational lines. For this case study, the hierarchy of management stations is organized based on geographical location. If you need to add another set of managed devices, the only change necessary is the addition of another collection station in order to extend the NMS.

Distributed Network Management Systems

The fault management system design presented in this case study is based on HP OpenView's Network Node Manager (NNM) Distributed Internet Discovery and Monitoring (DIDM) model. The DIDM model is the mechanism for providing a redundant distributed NMS architecture. Using the DIDM model, helps form a conceptual foundation for integrating the rest of the NMS components.

Figure 3-1 illustrates an example HP OpenView-based NMS environment featuring a redundant management scheme supporting multiple management domains. The section that follows ("Applying the FCAPS Model") explores how this environment can be used to implement an FCAPS-based network management system.

CiscoWorks 2000 (CW2000) is used for configuration management in this case study. In this environment, each management station runs CW2000. This gives each Network Operation Center (NOC) the ability to perform configuration management tasks. CW2000 provides configuration and inventory management for Cisco NASs, switches, and routers.

The configuration and maintenance of the SC2200 is handled by Cisco Media Gateway Controller Manager (CMM). CMM is an element management system designed specifically for the SC2200. CMM handles all the configuring and provisioning the SC2200s in the network. CMM also can reside on the management station platforms—permitting each NOC to perform configuration management on SC2200s.

Figure 3-1 Overview of Example NMS Case Study Environment

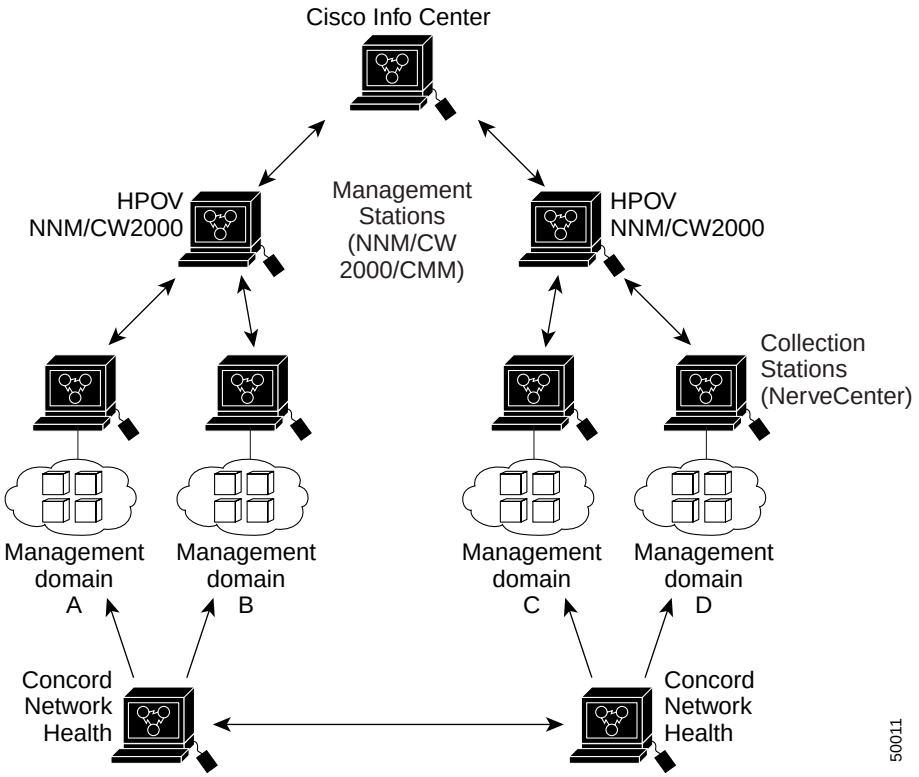

Applying the FCAPS Model

In the context of a wholesale dial environment, the HP OpenView DIDM model and the OSI NMS FCAPS model can be applied to the following network management areas:

- Fault and Configuration Management
- Event Management
- Performance Management
- Security Integration

Fault and Configuration Management

The DIDM model allows for more than one NOC to have a fully operational fault management system (as illustrated in Figure 3-1). This allows for the network to be separated into multiple management domains—each monitored by its own fault manager. Thus, two primary management domains exist in this solution. Each fault manager shares real-time information with the other fault manager. Furthermore, each of the management domains has a sub-domain, containing a Veritas NerveCenter. This distributed approach promotes scalability and reduces traffic caused by polling because it pushes polling closer to the managed devices.

Note This environment could have implemented HP OpenView for collection stations; however, NerveCenter was selected because of its support of event correlation, trap filtering, and its flexible polling engine.

An enterprise-class UNIX server should be deployed for each fault management system. This facilitates growth while allowing for some element managers to be co-located on these servers. A Sun E4000 server would support a medium to large network. HPOV Network Node Manager Version 6.0 resides on this system, along with CiscoWorks 2000. The E4000 class server is powerful enough to handle the fault and configuration management load for the entire wholesale dial network if one of the servers fails.

Note For related information in the document about element managers, see the "Configuration Management Implementation" section on page 2-13.

Event Management

Event management systems provide mechanisms that automate actions to resolve issues, email/pager notifications, and create trouble tickets. The Cisco Info Center (CIC) supports these functions. A CIC system can reside at the same location as a HPOV NNM system. While the CIC in this case would reside at one of the NOCs and would be associated with one of the HPOV NNM management stations, both management stations would have the same level of access. Figure 3-1 illustrates the CIC as situated logically above both management stations. A CIC `trapd` converter is loaded on the HPOV server, allowing SNMP on HPOV NNM to be converted and sent to the CIC.

Performance Management

In selecting and implementing a performance management system, first consider the usability of the application, followed by scalability and reliability. Based on these concerns, this case study calls for the implementation of Concord's *Network Health*.

Network Health is a web-based performance management tool that provides predefined graphs and graphing templates. Network Health can be deployed in a distributed architecture, in very much the same manner that HPOV NNM is deployed in this case study.

For this case study we use a SUN Microsystems 420R containing Concord's Network Health with one located in each NOC and splitting the wholesale dial network into two performance management domains. Each Concord Network Health shares information with its peers so that each NOC has a complete view of the network as it relates to performance.

Security Integration

Access Registrar is implemented for access security in this case study. Access Registrar's configuration consists of two servers strategically placed in the network. Each server includes an underlying Oracle database. The databases are mirrored to allow the Access Registrar servers to share AAA information. This approach provides a backup system if one server is unavailable and distributes the AAA load.

Access Registrar does not integrate with other elements of the NMS solution, except that it should be configured to send a server daemon status to CIC, alerting the network operator of server problems.

The wholesale dial service provider should implement its own access server for staff accounts and maintenance accounts. However, each retailer (ISP) should be responsible for its own account security. This administrative separation insulates the wholesale dial provider from supporting individual dial users.

Server Recommendations and Configurations

The following list of recommended servers, followed by the model and UNIX platform specification, summarizes the hardware and software requirements for each NMS component introduced in this case study. Full deployment of this solution requires a large number of Sun UNIX servers. Cisco recommends that this network use Network Information Services Plus (NIS+) servers to reduce the need for system administration.

Note The selection of appropriate NMS platforms can vary greatly and depends on the specifics of a given network architecture. Recommendations presented here are designed for this case study. The purpose of presenting these recommendations is to provide a starting point for determining appropriate hardware platforms for real-world implementations.

Recommended hardware platform for each HPOV server:

- Sun Enterprise 4000 running Solaris 2.6
  - 4 GBytes RAM
  - 32 GBytes HD
  - Two 400MHz CPUs

Recommended hardware platform for each CIC server:

- Sun Enterprise 420R running Solaris 2.6
  - 2 GBytes of RAM
  - 34 GBytes HD
  - Two 400 MHz CPUs

Recommended hardware platform for each Veritas NerveCenter/Concord Network Health instance:

- Sun Enterprise 420R running Solaris 2.6 (unless deployed in the central office, then Sun Netra 1450)
  - 1 GBytes of RAM
  - 18 GBytes HD
  - Two 400MHz CPUs

Recommended hardware platform for each Access Registrar:

- Sun Enterprise 420R running Solaris 2.6
  - 1 GBytes of RAM
  - 18 GBytes HD
  - Two 400MHz CPUs

System Fail-over / Backup Configuration for Collection Stations

In order for one collection station to take over if another fails, the *standby* collection station must be aware that there is a problem with the failed collection station and have a working knowledge of all the collections stations devices. This is accomplished by distributing a *node list*.

The node list of each NerveCenter collection station includes nodes from its managed domain and the configured neighboring domain. The managed domain nodes are configured to forward traps to two destinations:

- The primary collection station
- The stand-by collection station

The managed domain nodes belong to the *active property group* while the neighboring domain nodes belong to the *inactive property group*. In considering the network illustrated in Figure 3-1, Nerve Center Collection Station (NC) 1, NC2 and NC3 manage only the *active nodes* (nodes in the active property group).

1. NC2 polls NC1, NC3 polls NC2, and NC3 polls NC1 on two-minute intervals.
2. When any collection station detects the loss of its neighbor, the property group of the neighboring nodes is changed to *active*.
3. When that collection station detects the re-establishment of its neighbor, the property group of the neighboring nodes is changed back to *inactive*.
4. The neighbors are configured in a round robin fashion, (for example, NC1 backs up NC2, NC2 backs up NC3, NC3 backs up NC1).

The same process applies to maintain the redundant management station configurations. Assume that management station (MS) 1 is the active server for the network and managed domain. MS2 remains passive, maintaining a one-minute heart beat poll with MS1. If MS1 is unavailable, MS2 becomes the active server.

While in the passive state, MS2 hosts the user sessions and required element managers. Both management stations, MS1 and MS2, contain a copy of the element managers. The element manager's data is synchronized between the two management stations on a nightly basis.

Note For related information in the document about element managers, see the "Configuration Management Implementation" section on page 2-13.