NORTEL NETWORKS

Preside Multiservice Data Manager

Engineering and Planning Guide

241-6001-101

Preside Multiservice Data Manager

Engineering and Planning

Guide

Publication: 241-6001-101
Document status: Standard
Document version: 14.3RSUP
Document date: December 2003

Printed in Canada

Publication History

December 2003

14.3RSUP Standard
Commercial availability

Contents

About this document

This document describes the engineering, configuration and planning activities for the Preside Multiservice Data Manager (MDM). The following topics are discussed in this section:

Who should read this document and why

This document is designed to help you do the following:

- set up your Preside Multiservice Data Manager (MDM) environment to make effective and efficient use of hardware and software
- engineer the MDM environment so that it best suits your needs
- make decisions that deploy MDM effectively into the existing environment
- plan the initial configuration of MDM and plan subsequent reconfigurations to suit system expansions or modifications

This document presents guidelines to configure a workstation (or server) to run the MDM software. Also, the document simplifies the engineering process and minimizes the workstation monitoring effort for the planner. To

achieve these objectives, this document presents conservative, initial workstation configurations, that are supplemented by monitoring rules and software to support configuration optimization and growth.

This document is intended for personnel who plan, install, and monitor MDM systems and applications.

What you need to know

This document assumes familiarity with all aspects of Preside Multiservice Data Manager (MDM) workstations and their applications.

It is also assumed that you are familiar with the following:

- Sun SPARCstation workstation platforms
- the Solaris operating system
- network communications (X.25, frame relay, and Internet Protocol (IP)
- UNIX

How this document is organized

This document contains the following sections:

- "Engineering overview" (page 19) provides an introduction to engineering and planning for Preside Multiservice Data Manager (MDM).
- "Determining your MDM requirements" (page 23) discusses the factors to consider when determining the requirements for an MDM to support your network.
- "Passport connectivity" (page 35) describes connectivity types, management protocols, and bandwidth requirements.
- "Distributed surveillance architecture" (page 49) describes the various components that support network engineering.
- "MDM servers" (page 61) describes MDM servers and how they are deployed in the network. that are used in the MDM network. This section also contains information on Passport groups.
- "Choosing a configuration for MDM" (page 77) describes the MDM configurations you can use to satisfy your MDM requirements.

- "Choosing a workstation to run MDM" (page 95) provides you with information about the minimum workstation configuration required to run MDM.
- "Monitoring MDM" (page 99) describes how to monitor MDM and maintain it at peak operating efficiency.
- "System resource utilization programs" (page 113) provides a listing of the MDM-supplied monitoring scripts.
- "MDM Engineering for SNMP Devices" (page 123) contains engineering calculations and information for SNMP devices such as Baystack 450.
- "MDP requirements" (page 131) describes options for MDP deployment.
- "Upgrading your UNIX server" (page 137) contains information to determine when to upgrade your server. The section also helps you determine which workstation to use as your server.

Text conventions

This document uses the following text conventions:

- `nonproportional spaced plain type`

  Nonproportional spaced plain type represents system generated text or text that appears on your screen.

- **nonproportional spaced bold type**

  Nonproportional spaced bold type represents words that you should type or that you should select on the screen.

- *italics*

  Statements that appear in italics in a procedure explain the results of a particular step and appear immediately following the step.

  Words that appear in italics in text are for naming.

- `[optional_parameter]`

  Words in square brackets represent optional parameters. The command can be entered with or without the words in the square brackets.

- `<general_term>`

  Words in angle brackets represent variables which are to be replaced with specific values.

- UPPERCASE, lowercase

  In Preside Multiservice Data Manager (MDM), uppercase and lowercase letters that appear in UNIX commands and parameters must be matched exactly. The system matches upper and lowercase characters differently.

- |

  This symbol separates items from which you may select one; for example, ON|OFF indicates that you may specify ON or OFF. If you do not make a choice, a default of ON is assumed.

- ...

  Three dots in a command indicate that the parameter may be repeated more than once in succession.

The term *absolute pathname* refers to the full specification of a path starting from the root directory. Absolute pathnames always begin with the slash (/) symbol. A *relative pathname* takes the current directory as its starting point, and starts with any alphanumeric character (other than /).

Related documents

See the following documents for related information:

- 241-6001-100 *Preside MDM Installer Guide*
- 241-6001-303 *Preside MDM Administrator Guide*
- 241-5701-270 *Passport 7400, 15000, 20000 Software Installation Guide*
- 241-6001-309 *Preside MDM Management Data Provider User Guide*

- 241-6001-310 *Preside MDM Server Reference Guide*
- 241-6001-806 *Preside MDM MDP Data Formats Reference Guide*
- 241-6001-118 *Preside MDM SNMP Surveillance Adapter Guide*

Chapter 1
Engineering overview

This section provides an overview of the engineering, configuring and planning functions for Preside Multiservice Data Manager (MDM) on a SUN computing platform. Engineering includes the following information:

- “Planning for MDM” (page 20)
- “Configuring MDM” (page 21)
- “Monitoring MDM” (page 21)

Engineering Preside Multiservice Data Manager (MDM) is a three-phase, circular process that consists of planning, configuring, and monitoring, as shown in the figure “MDM engineering philosophy” (page 20).

The planning phase consists of analyzing the network, and selecting the network topology that best serves your needs. The configuring phase consists of selecting a configuration for your network. The monitoring phase begins once the network is operational. Monitoring the network ensures that it runs at peak efficiency. Changes are required in the monitoring phase when it is not possible to maintain the network at peak efficiency.

Figure 1
MDM engineering philosophy

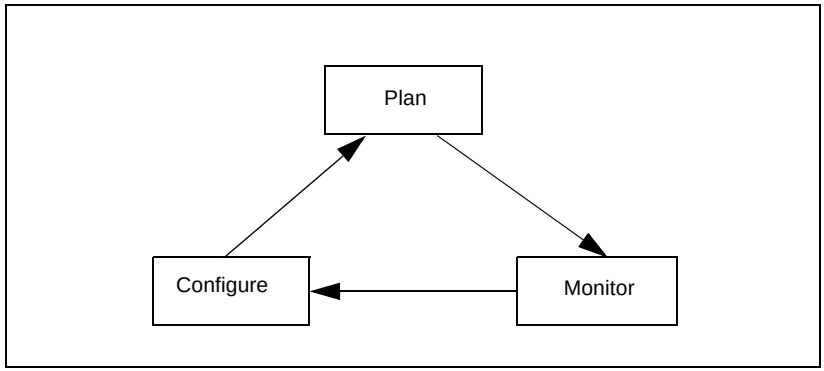

For MDM engineering to succeed, the process is a closed loop. The requirements change continuously as new technologies emerge and the network grows. You must be prepared to change the MDM configuration as the network requirements change. This section provides an overview of each of these phases, and outlines the importance of each phase.

Planning for MDM

Planning for Preside Multiservice Data Manager (MDM) consists of determining the requirements to manage the network. These requirements depend on many factors, such as the size of the network and the number of users.

The requirements dictate that MDM and Data Packet Network (DPN) topology designs be planned at the same time. The size of the network is the most important factor when determining the MDM engineering requirements. The network can be small, medium, or large, based on the number of modules it contains.

For more information on planning your network, see “What is network planning?” (page 23).

MDM philosophy of reuse

If your requirements change due to network growth, it is beneficial to purchase the additional equipment. This can require the following:

- upgrading existing equipment with more memory and additional disks
- redeploying the equipment to different areas
- having the equipment take on different functions

For more information on network planning, see “Determining your MDM requirements” (page 23).

Configuring MDM

After you determine your Preside Multiservice Data Manager (MDM) requirements, select one of the following configurations that meet these requirements:

- stand-alone
- stand-alone CPU server
- client/server
- network file server (NFS)
- combination

For a description of these configurations and information to help you select a configuration, see “Choosing a configuration for MDM” (page 77).

Monitoring MDM

After the Preside Multiservice Data Manager (MDM) topology design is implemented, the network planner regularly monitors each MDM workstation to keep it operating at peak efficiency. This involves monitoring the CPU, memory, and I/O or MDM hardware to achieve a quick response and high MDM throughput.

The network planner must also consider the behavior of the individual workstation and the external networks to which it is attached.

Chapter 2
Determining your MDM requirements

This section includes the following information to help you determine your Preside Multiservice Data Manager (MDM) requirements, and the affects on your network:

What is network planning?

Network planning includes the following tasks:

- estimating the current size of the network that Preside Multiservice Data Manager (MDM) will be deployed in, in addition to
  - — the expected growth rate for the network
  - — the amount of information to be collected and stored for the network

- considering whether a centralized or decentralized management strategy will be implemented
  - — if decentralized, which MDM applications and functions will be managed by each site
- assessing the cost of installing, configuring, and administering:
  - — the frame relay or Internet Protocol (IP) links to Passport 6000, 7000, 15000 and 20000 in the network
  - — the X.25 links from the Preside MDM workstation to DPN-100 in the network
  - — IP connections to devices that are managed through SNMP in the network
- considering personnel access requirements
- determining the number of users and their geographic and functional distribution
- estimating training requirements and costs
- determining the hardware requirements (CPU power, RAM, disk space, SWAP space, and network link bandwidth) for workstations that will run the MDM software
- determining the need for data backups

Estimating the size of your network

Network size is the most important factor in determining the Preside Multiservice Data Manager (MDM) configuration to manage your network. Network size relates to the number of average modules and average SNMP devices in your network.

What is an average module?

The term average module applies to Passport 6000, 7000, 15000, 20000, and DPN-100 switches.

The engineering loads resource modules (RM), access modules (AM), or Passport 6000, 6000, 15000 and 20000 present to Preside Multiservice Data Manager (MDM) are proportional to

- the number of processors surveilled
- the number of lines provisioned
- the MDM tools being used

An average module is a concept used to normalize these factors for various types of modules. See the table "Estimates of the number of average modules" (page 25) for a list of the average module estimates for the modules in your network.

Table 1
Estimates of the number of average modules

Hardware type	Equivalent number of average modules
Single-shelf AM/RM	1
Dual-shelf AM/RM	2
Network Management Module (NMM)	0.5
DPN-100/1	0.16
DPN-100/3	0.25
Passport 4120	0.2
Passport 3 slot	0.12
Passport 5 slot	0.3
Passport 16 slot	1.0
Passport 15000	1.5
Passport 20000	1.5

Calculating the number of average modules in your network

1 Count the number of each type listed in the table "Estimates of the number of average modules" (page 25).

2. Multiply the number of each type by the equivalent number of average modules (from the table “Estimates of the number of average modules” (page 25) to get the number of average modules per type.
3. Add the total to obtain the average number of modules for the network.
4. Use the resulting number, or adjust it proportionately if network monitoring and provisioning functions are to be divided into smaller administrative units. See “Choosing a configuration for MDM” (page 77).
5. If you are installing a new network and you need to calculate the number of modules in your network, estimate the number of modules that your network is likely to contain in two years and add 10%.

What is an average SNMP managed device?

The term average SNMP device applies to devices that communicate with Preside MDM by means of simple network management protocol (SNMP). These devices include

- SNMP devices that provide fault information to Preside MDM through a device integration fault cartridge such as Baystack 450 or Passport 8600
- any device managed through the Preside MDM SNMP Surveillance Adapter described in 241-6001-018, *NMS Engineering Data Reporter User Guide*

The engineering loads presented by SNMP devices to Preside Multiservice Data Manager (MDM) are proportional to

- the number of devices surveilled
- the number of number of components per device
- number of SNMP variables polled per component/interface/device
- polling rate per device
- average alarm/ trap rate per device/interface/component
- the MDM tools being used

The concept of an average SNMP device is used to normalize these factors across the various device types and configurations. See the table “Estimates of the number of SNMP managed device” (page 27) for a list of the average SNMP module estimates for the devices in your network.

The information in this table is intended as a starting point. For details about engineering for SNMP devices, see “MDM Engineering for SNMP Devices” (page 123).

Table 2
Estimates of the number of SNMP managed device

Hardware type	Equivalent number of SNMP device
Passport 8600	0.125
Juniper	0.16
Baystack 450	0.08
Passport 4400	0.04

Calculating the number of average SNMP devices in your network

Use the following procedure to calculate the number of average SNMP devices in your network. To calculate the number of average SNMP devices that are not listed in the table, see “Determining the number of average SNMP devices” (page 126)

1. Count the number of each type listed in the table “Estimates of the number of SNMP managed device” (page 27).
2. Multiply the number of each type by the equivalent number of average modules (from the table “Estimates of the number of SNMP managed device” (page 27) to get the number of average modules per type.
3. Add the total to obtain the number of SNMP managed modules for the network.
4. Use the resulting number, or adjust it proportionately if network monitoring and provisioning functions are to be divided into smaller administrative units. See “Choosing a configuration for MDM” (page 77).
5. If you are installing a new network and you need to calculate the number of devices in your network, estimate the number of devices that your network is likely to contain in two years and add 10%.

Effect of network size

The size of the network can be small, medium, or large, depending on the number of modules it contains.

The performance requirements for workstations increase according to network size. As the network expands, you can add newer, more powerful workstations, and redeploy the less powerful workstations as clients of server workstations or as X-terminals.

A small number of powerful workstations mixed with a large number of less-powerful workstations is the most economical, both in original hardware and ongoing administration costs. In determining your MDM requirements, attempt to extend the service lives of the existing workstations.

The figure "Relationship between network size and MDM deployment" (page 28) shows an evolution path for MDM deployment as the network grows.

Figure 2
Relationship between network size and MDM deployment

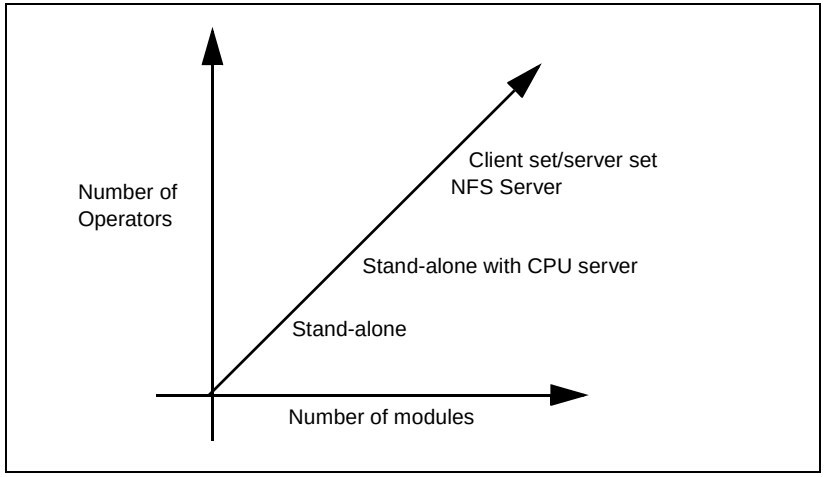

In a small initial network, all MDM functions can be performed on one stand-alone workstation. As the network size increases, the capabilities of some of the stand-alone workstations can extended by using them as CPU servers. Additional network control centers can be used for surveillance and

provisioning. MDM functions can also be distributed among Sun Solaris UltraSPARC machines, which can function as external network file servers or client servers.

Economical incentives to add low-cost MDM access from non-SUN SPARCstation terminals and to use older MDM workstations with small disks will result in the introduction of Sun servers. Finally, in the larger networks, applications can be partitioned among workstations. For example, you can run fault tools on one workstation and configuration tools on another.

The benefits of one configuration can outweigh those of another. Therefore, most medium and large networks use combinations of the different types of servers.

Effect of MDM applications

Individual Preside Multiservice Data Manager (MDM) applications are designed for different types of network behavior. For example, computing requirements depend on the number of modules, processing elements (PE), function processors (FP) and alarms generated by the network. Configuration computing requirements depend on the number of service orders (additions and changes) being processed. In large networks, using the number of average-sized modules is accurate enough for planning purposes and it simplifies the planning function.

Effect of the number of users

Before setting up a network, determine the following:

- how many users are going to use the system
- how they will access the system
- which tools they will access

This information helps you determine the hardware required and how to configure the hardware, especially in server configurations. For example, the number of active users using a server is one of the limiting factors when sizing the workstation.

The tools each user will use, and the user location affects the hardware required and its cost. For example, a provisioning center, where colocated users use configuration tools to perform provisioning, can use the services of a single server set or CPU server running configuration tools by working from less expensive SPARC-compliant or client workstations.

Geographically widespread users can choose between the cost of:

- stand-alone workstations
- additional WAN bandwidth to support IP over WAN connections to remote workstations or X-terminals not accessible over a LAN (if using MDM)

Effect of the projected growth rate

The planning step includes scheduling for changes in the computing environment. Users usually do not want to be making continuous changes to their servers. Continual workstation configuration changes are expensive and reduce system availability. If you foresee rapid network growth, increase initial workstation sizes and link capacities. For example, if you have high-speed interface (HSI) installed for frame relay, use the highest link speed available.

Effect of the organization of administrative staff

Skills required to administer and maintain Preside Multiservice Data Manager (MDM) include a thorough knowledge of UNIX, networking, and MDM. If you select an MDM configuration that minimizes the number of administrators, you can centralize the administrative function at as few sites as possible. Centralizing this function can increase WAN traffic due to the increased need for IP over WAN connections. Centralization can also increase frame relay traffic due to an increased need for IP over frame relay connections.

Effect of network availability

The primary goal of Preside Multiservice Data Manager (MDM) is to monitor the network to ensure that it runs at peak efficiency. To achieve this goal, it is critical that access to the network through MDM is maintained.

When planning your network, anticipate and plan for scheduled and unscheduled outages and failures to ensure that access to the network is always maintained. Ensure network availability by building in redundancy through including stand-alone workstations in your network configuration, or by adding redundant servers when running a server configuration. For more information, see “Building in redundancy” (page 91).

Effect of link capacity

This section contains information about the capacity of following types of links from switches and devices in the network and the Preside MDM workstation:

- “X.25 links to DPN-100 switches in the network” (page 31)
- “IP over X.25 links to Passports in a network that contains DPN-100 and Passports” (page 33)
- “IP over frame relay links and IP over ethernet links to Passports” (page 33)
- “IP over ethernet links to SNMP managed devices” (page 33)

X.25 links to DPN-100 switches in the network

The X.25 connection from Preside Multiservice Data Manager (MDM) to an AM or RM must be configured correctly to ensure successful workstation operation. Because X.25 link utilization will be high, ensure that the trunks or network links supporting the X.25 service have

- sufficient bandwidth, typically twice that of the X.25 link
- available bandwidth

Also optimize SunLink X.25 default parameters. For example, you can raise the packet size to 512 bytes to reduce the amount of bandwidth occupied by transporting packet header information across the X.25 links.

Overall X.25 traffic is the sum of

- status probing
- network alarms
- network commands

- service data provisioning
- software downloads

There is usually some form of file transfer, which also uses up some of the available bandwidth. The table “X.25 bandwidth requirements for DPN networks” (page 32) shows bandwidth and status polling requirements. If the link bandwidth is insufficient for network control system (NCS) status probing intervals, analyze the probing interval for possible reduction. If a status probe interval change cannot be made or provides insufficient additional bandwidth, then increase the link speed.

Table 3
X.25 bandwidth requirements for DPN networks

X.25 link bandwidth (kbit/s)	Maximum recommend bandwidth (65% of actual)	Status polling interval (min): office	Status polling interval (min): subcomponents	Number of modules supported by bandwidth
9.6	6.2	1	1	90
9.6	6.2	1	5	270
19.2	12.4	1	1	180
19.2	12.4	1	5	450
64	41	1	1	600
64	41	1	5	1750
64	41	1	10	2300
128	83	1	1	1200
256	166	1	1	2400

The AM or RM X.25 links to the MDM must be monitored using X.25 link statistics to ensure that utilization levels do not exceed 65%. This reduces link delay and provides sufficient bandwidth for the Configuration tools.

IP over X.25 links to Passports in a network that contains DPN-100 and Passports

Workstations support character mode access and IP over X.25. Character mode access is used in remote network communication system applications. Full graphic access is achieved by using the X.11 protocol by rlogin access to provide Preside Multiservice Data Manager (MDM) window support.

IP over X.25 is required when the following occurs:

- the network contains Passport nodes and DPN nodes
- you are using the *Service Selection tool* to allow several MDM workstations to share the MDM server processes running in a Server set. See 241-6001-303 *Preside MDM Administrator Guide* for a description of the *LAN Selector*.
- you are using the Network Time Protocol to synchronize the network time. See 241-6001-303 *Preside MDM Administrator Guide* for a description of how to use Network Time Protocol to synchronize the network time.
- you are providing rlogin access to remote workstations

IP over frame relay links and IP over ethernet links to Passports

In networks that only contain Passport switches, you can connect the Preside Multiservice Data Manager (MDM) to Passport through an IP over frame relay link. When the Passport switches are configured as an inter-LAN switching (ILS) network, you can connect the MDM to one of the Passport switches through an IP over Ethernet link.

IP over ethernet links to SNMP managed devices

In networks that contain devices that are managed through SNMP, you can connect the Preside Multiservice Data Manager (MDM) to the network through an IP over ethernet link. The link must have sufficient capacity for handling the volume of traps from the managed devices and responses to polling from Preside MDM.

MDM backup

Performing regular backups ensures that the network can be recovered. There are two types of Preside Multiservice Data Manager (MDM) backups:

- file backup uses a 4 mm or 8 mm tape drive to make regular copies of critical files as changes occur. It is important to back up module service data and network files.
- workstation backup involves configuring workstations so that they can assume the functions of failed MDM workstations. Workstations can provide emergency backup support while still performing normal processing functions. Appropriate MDM backup configurations vary greatly according to
  - — the network size
  - — site
  - — application
  - — the projected mean times for repairing failed workstations at different locations

Give priority to providing backups for workstations that function as servers. This document does not contain specific recommendations for backup configurations.

Chapter 3
Passport connectivity

This sectioncontains information on Passport connectivity, management protocols, and bandwidth requirements. The following information is contained in this section:

For more information on Passport connectivity, see 241-5701-270 *Passport 7400, 15000, 20000 Software Installation Guide*.

Connectivity types

Passport networks support in-band and out-of-band connectivity based upon Internet Protocol (IP) management. In a Passport network, you can choose the internal IP (Ipi) or virtual router IP (VrIP) subsystems for connectivity. The subsystem you choose depends on specific circumstances, and whether you have a LAN infrastructure. This section describes in-band and out-of-band connectivity.

The figure "In-band and out-of-band connectivity example" (page 36) shows a Passport network that uses in-band and out-of-band connectivity.

Figure 3
In-band and out-of-band connectivity example

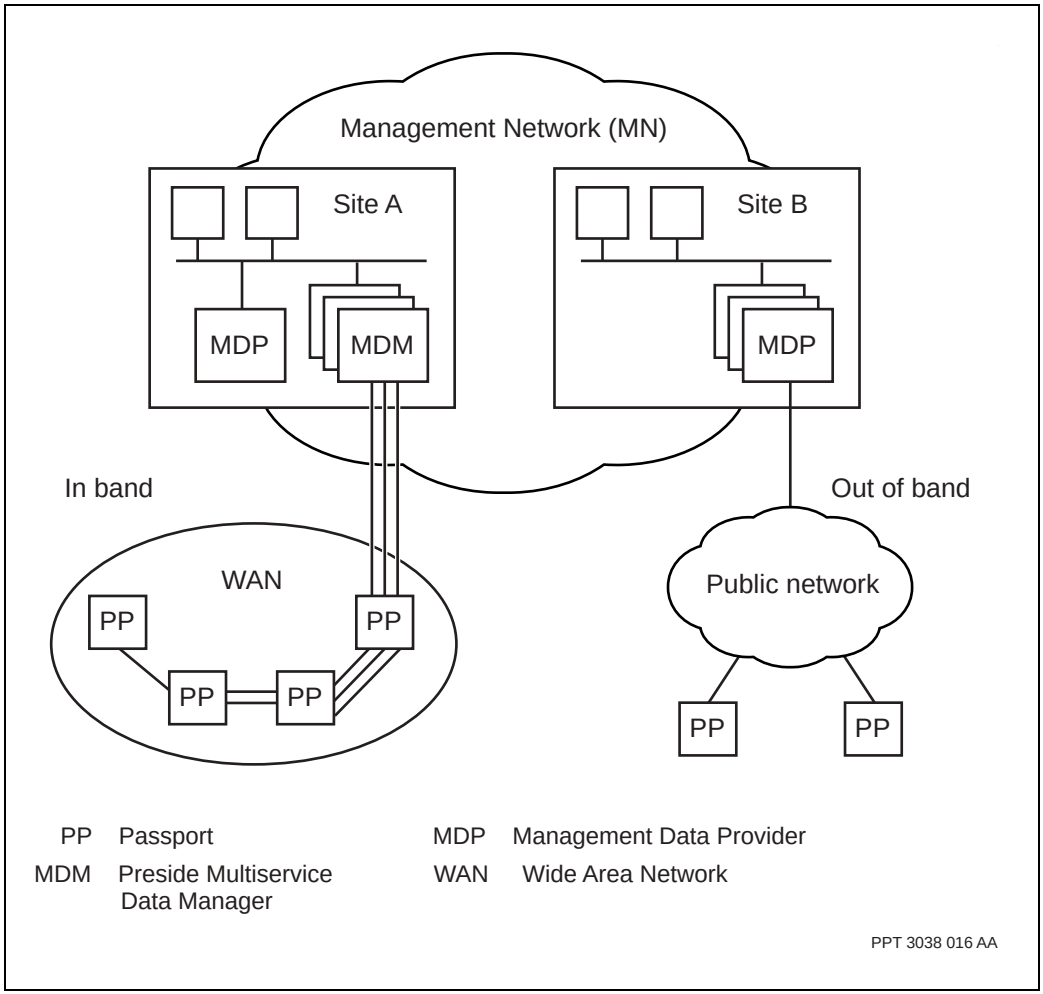

In-band connectivity

The IP stack running on the Passport is called the Ipi subsystem. There are three types of in-band connectivity: IP over virtual circuit (IpiVc), IP over frame relay (IpiFr) and IP over ATM (ATM MPE).

IpiVc is used when the network contains Data Packet Network-100 (DPN-100) and Passport switches. IpiVc routes the X.25 traffic from the DPN-100 to the Passport when Preside Multiservice Data Manager (MDM) is connected by an X.25 link to the DPN-100. The DPN-100 connects to the Passport.

IpiFr is used mainly for a Passport-only network. The system routes all the network management traffic back to the management stations by way of the system backbone network. There are two approaches to connecting the Passport using frame relay:

- Preside MDM is connected to the Passport with a high-speed serial interface (HSSI) card with SunLink frame relay software running on a Solaris workstation.
- An access router is deployed to perform the frame relay to IP conversion. The router acts like a frame relay access device (FRAD). Multiple workstations can share a physical frame relay link, providing the bandwidth can support the management traffic.

IpiFr is the most commonly deployed method because it is proven and low in cost. It is, however, time consuming to get all the virtual circuits up, and requires labor-intensive provisioning.

IP over ATM is only used in a Passport-only network. Preside MDM is connected from an ATM network interface card (NIC) running on a Solaris workstation either through a router or directly to the Passport switch.

Once the initial connection is made to the first Passport, ATM based permanent virtual circuit's can be set up from the first Passport to other Passports to manage them through the same Preside MDM workstation.

The figure,"In-band connectivity example using SunATM NIC" (page 38) show a Passport network using in-band connectivity. The workstation is connected to the workstation using a SunATM network interface card (NIC).

Figure 4
In-band connectivity example using SunATM NIC

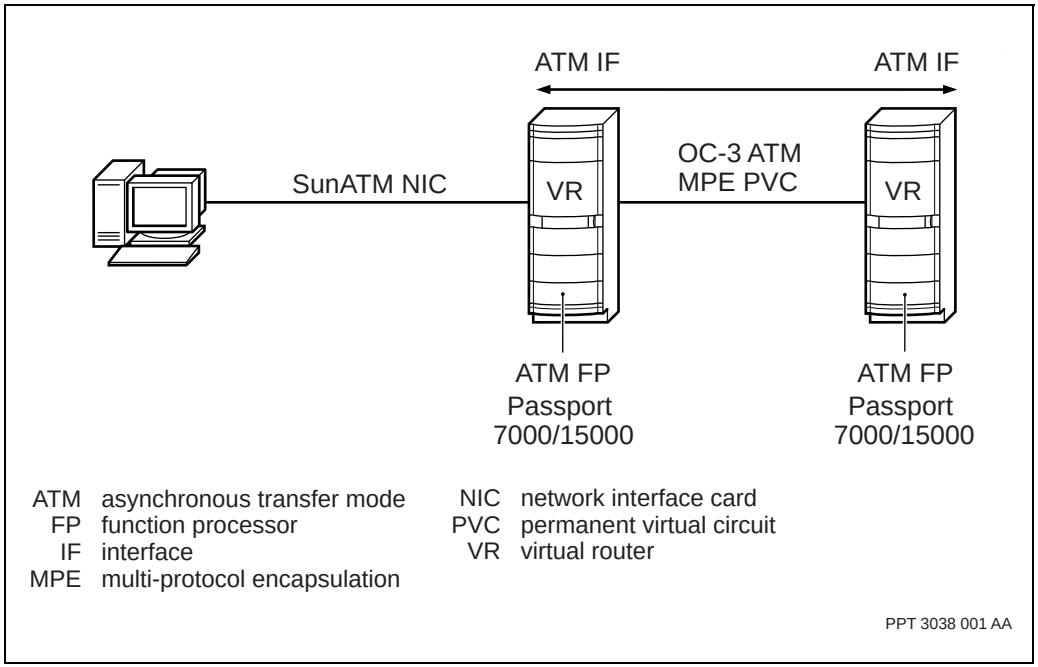

The figure “In-band connectivity example using a router” (page 39) shows a Passport network using in-band connectivity. The workstation is connected to the Passport network by a router.

Figure 5
In-band connectivity example using a router

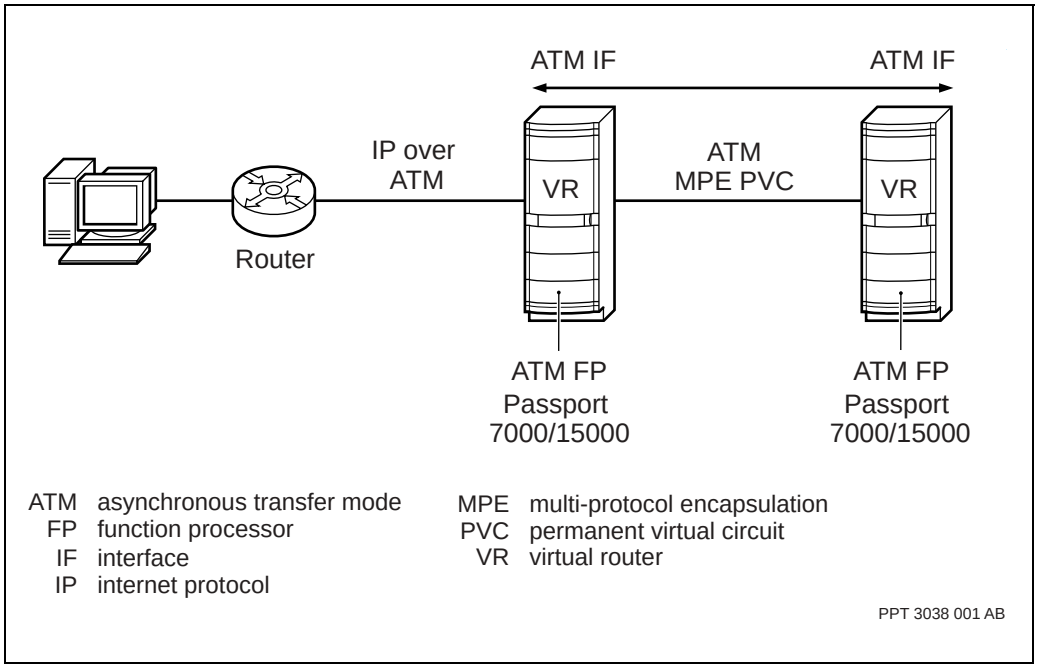

The network needs a Passport 7480 or another device to which the Passport 15000 can connect to. Otherwise, the MDM is connected to the Ethernet port of the control processor (CP) on the Passport 15000. The VrIp subsystem is used to connect to internal Passport 15000 nodes. You cannot have more than 10 internal Passports in this scenario. In-band connectivity by an access device, such as Passport 7000, or out-of-band connectivity by CP Ethernet is recommended.

Out-of-band connectivity

Passport also supports the IP stack, VrIp, which is also known as the inter-lan switching IP (ILS IP). VrIp is used when there is Ethernet or ATM multi-protocol encapsulation (AtmMpe) LAN media.

If you have a LAN infrastructure in place, connect the management stations to the Passport by way of the local Passport Ethernet port to gain network management connectivity. The management traffic is transmitted through the access router to the workstation by way of the router path.

If your Passport network is an ATM network, and you cannot reach your remote sites by an external router, use the Passport to route management traffic back to the local management station. Deploying AtmMpe converts ATM cells back into IP packets. An access router, Accelar 8600, or Passport 7000 perform the AtmMpe function and have the rest of the Passport running ATM bearer services with logical trunks.

A typical deployment consists of using Passport 7000s to serve as access nodes, and Passport 15000s to function as the ATM backbone. You could also use an access router to run the AtmMpe function, as this frees up a Passport shelf and offloads the processing workload onto the router. This leaves the Passport with higher capacity to do other work.

The figure, “Out-of-band connectivity example” (page 41) shows an Passport network using out-of-band connectivity.

Figure 6
Out-of-band connectivity example

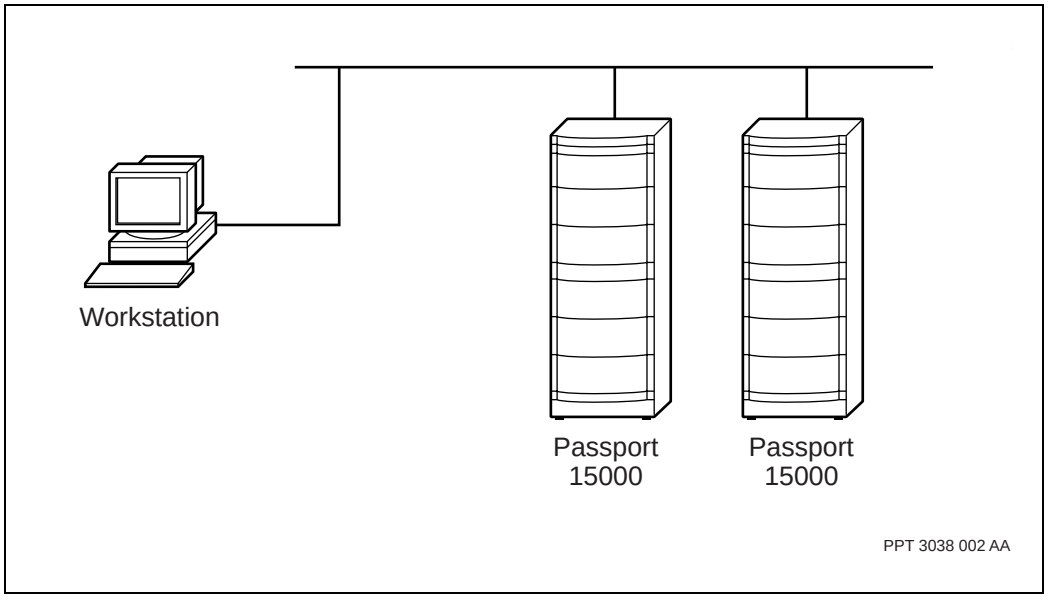

Passport on-switch management protocols

Passport 6000, 7000, 15000 and 20000 support the following management protocols:

- simple network management protocol (SNMP) based management for fault and performance
- file transfer protocol (FTP) based management for software downloaded from a software distribution site (SDS) and data collection to the Management Data Provider (MDP)
- telnet-based management for provisioning using Preside Multiservice Data Manager (MDM)
- network time protocol for network time of day synchronization. Network time protocol is essential for the management of large networks to assure synchronized time stamping for alarms, statistics, and accounting records. Network time protocol can be taken from a local workstation clock, an Internet clock or another external source that is supported on the MDP.

FTP with IPSec

You can apply IPSec to the FTP connection you use for uploading and downloading data between MDM and Passport. For more information about configuring IPSec, see 241-6001-040 *Preside MDM Security User Guide*.

The table “FTP with IPSec” (page 42) lists the average time, in seconds, to FTP a 3.25 Mb file between a Passport and an MDM using various encryption protocols, authentication protocols, or combinations of encryption and authentication protocols.

Table 4
FTP with IPSec

Applied protocol	Time (in seconds) to complete transfer
none	10
DES	27
3DES	58
SHA	14
MD5	13
SHA with DES	32
SHA with 3 DES	1:02
MD5 with DES	30
MD5 with 3DES	1:02

Connectivity bandwidth engineering

The fast management information protocol (FMIP) performs default surveillance of a Passport 7000 or Passport 15000 switch. FMIP is a proprietary protocol for Passport, and runs over transmission control protocol/Internet protocol (TCP/IP).

The minimum bandwidth requirements are as follows:

- Passport alarms: 500 Byte

- Passport open system interconnection (OSI) state change notice (SCN): 150 Byte
- assume 10 SCN per Passport alarm
- assume 2 Passport alarms/minute

Use the following procedure to estimate the FMIP connection supporting surveillance of a single Passport switch.

1. Calculate the number of alarms per minute.
2. Multiple the number of alarms per minute by 0.237 Kbyte/s.

For example, 2 x 0.237 = 0.474 Kbyte/s per Passport 7480, assuming the following:

- Passport 7440 is 0.5 of 0.474 Kbyte/s
- Passport 7420 is 0.25 of 0.474 Kbyte/s
- Passport 15000 is 1.5 of 0.474 Kbyte/s
- Passport 20000 is 1.5 of 0.474 Kbyte/s

On-switch bandwidth requirements

The following are the on-switch bandwidth minimum requirements:

- 256 Kbyte/s available for FTP functions, such as software download and statistic and accounting data uploads
- 8 Kbyte/s available for FMIP functions, such as provisioning and fault analysis per operator. This also addresses network time protocol client requirements.
- 1 Kbyte/s average bandwidth available for daily surveillance operations by the FMDR server
- 16 Kbyte/s available for frame relay committed information rate (CIR) and 256 Kbtye/s available for emitted information rate (EIR) for Passport on-switch access bandwidth

Passport to MDM bandwidth requirements

Assuming an in-band management solution, the average bandwidth requirement between Preside Multiservice Data Manager (MDM) and Passport is a function of an average minimum of:

(1 Kbyte/s * #Passport 7000 or 15000) + (8 Kbyte/s * number of operators * 2) + a minimum of 256 Kbyte/s for adequate SDS FTP downloads and uploads

The Passport to MDM bandwidth requirements for 100 Passports with 10 operators is as follows:

(100 Kbyte/s + 160 Kbyte/s + 256 Kbyte/s) = 516 Kbyte/s

The Passport to MDM bandwidth requirements for 1000 Passports with 100 operators is as follows:

(100 Kbyte/s + 1600 Kbyte/s + 256 * 10 Kbyte/s) = 5160 Kbyte/s

MDP bandwidth requirements

The minimum link bandwidth requirement for MDP is 256 Kbyte/s.

SNMP interfaces on Passport 15000

The Passport 15000 supports several standard SNMP management information bases (MIB) and has its own registered enterprise MIB. This allows third-party applications to gain access to the switch by the SNMP protocol. Check with your third-party vendor to determine if it supports the monitoring of Passport products.

An SNMP session with a third-party application and an FMIP provisioning session with the Preside Multiservice Data Manager (MDM) application can occur at the same time.

Passport on-switch data collection system

Both the Passport 7000 and Passport 15000 contain a hard disk to which accounting and statistics can be spooled. These data files can be moved off switch by using FTP.

The table“DCS spooler guidelines” (page 45) lists the data collection system (DCS) event types, including their default setup.

Table 5
DCS spooler guidelines

DCS type	Spooling	File number	Agent queue size records)
Accounting	True	200	10000
Alarms	True	30	100
Logs	True	10	50
Debug	False	2	0
SCNs	True	10	200
Traps	False	2	50
Stats	True	200	0

You can configure the queue sizes for the data collection agents. When the queue size is set to zero, no collection occurs. When the queue is full, it overflows and discards data.

Data from data collection agents whose files are being spooled to the hard drive on the CP are placed in files for FTP retrieval. A new file is created for spooling when the following occurs:

- the maximum file size for spooling of 500 Kbyte is reached
- midnight or 00:00 hrs is reached daily for accounting
- the active CP resets, restarts, or reboots
- a hitless software migration completes from PCR1.3 release onward
- the file system is locked
- a newFile col/<dcsType> sp operator command is issued

DCS data files are built on Passports in fixed 500-kbyte files. Setting the maximum size to zero means that the number of files is limited to the space available on the disk. Remove files on a regular basis to ensure data is not lost. This is important before the maximum number of files is reached.

When files have been transferred by FTP from the Passport CP to the MDP, the successfully transferred files are automatically removed from the CP hard disk. The loss during CP switchover depends on the size of the file system buffers and on the rate at which data is being spooled to disk. On an 810-mbyte disk, the loss is up to 16 kbyte or up to 300 seconds, whichever is less.

> ***Note:*** DCS reduces the accounting records that are lost in the file system buffers. At each 10-second interval, DCS forces the file system buffer to be flushed to disk if DCS detects that data is flowing, but at a rate of less than 1 kbyte/10 seconds.

All files spooled to the primary CP hard disk are mirrored on the backup CP in case of primary CP failure. When files have been transferred from the Passport CP to the MDP, the successfully transferred files are automatically removed from the CP hard disk.

The table “Passport data collection system summary” (page 47) summarizes the various data collection types.

Table 6
Passport data collection system summary

Data collection Type	Data access methods	Defaults and maximum file size	MDP collected data
Alarm data: essential for real-time surveillance of the network and recommended to be on at all times	Local operator Telnet MDM spooler SNMP agent (traps)	Off 10 records 50 files	Yes
Statistics data: used for mid-term to long-term planning	FTP spooler	On 0 records 200 files	Yes
Accounting data: used for billing purposes	FTP spooler	On 10,000 records 200 files	Yes
State change notification (SCN) data: used to update the network model as opposed to regular polling	Local operator Telnet MDM spooler	Off 200 records 10 files Queue 0	No
Log data: monitor operator commands being issued to the switch	Local operator Telnet spooler	Off 0 records 10 files Queue 0	No
Trap data	SNMP agent	Off 50 records 2 files Queue 0	No
Debug data	Local operator Telnet spooler	Off 2 files 0 records Queue 0	No

Chapter 4
Distributed surveillance architecture

This section describes the Preside Multiservice Data Manager (MDM) distributed surveillance architecture, and includes the following information:

Distributed surveillance architecture provides a scalable surveillance architecture for Passport, other Nortel Networks data products, and third-party simple network management protocol (SNMP) devices.

MDM provides surveillance of the data switch, the connectivity to the data switch and the UNIX workstations, and MDM applications that make up the management network.

The figure "MDM supported distributed workstation configurations" (page 50) shows stand-alone, management console client, and distributed client server sets for workstation configurations.

Figure 7
MDM supported distributed workstation configurations

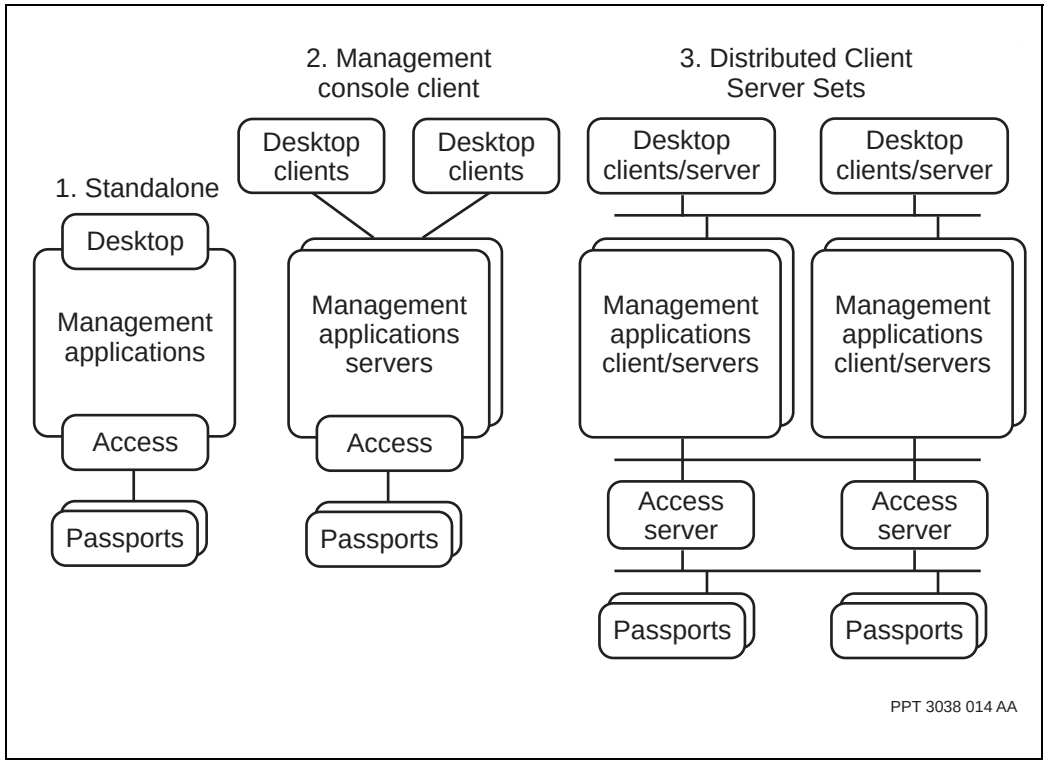

Surveillance servers

This section discusses the scalability of the surveillance servers. The servers shown in figure “Surveillance server architecture” (page 51) can be located on one workstation or on multiple workstations. The location of the servers depends on the size of the network and the number of network operators.

For more information on MDM servers, see 241-6001-310 *Preside MDM Server Reference Guide*.

Figure 8
Surveillance server architecture

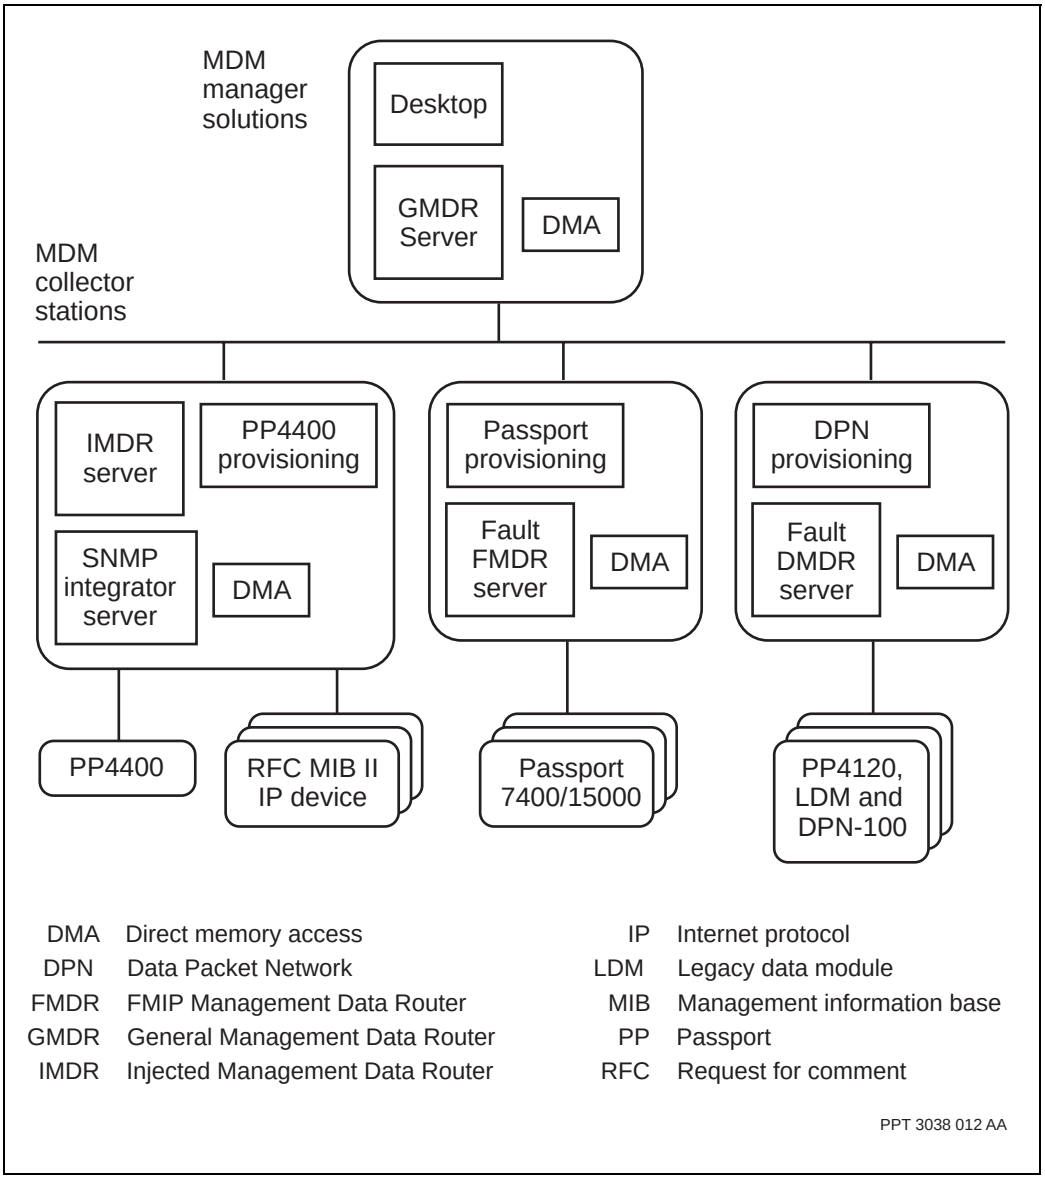

Generic DCD

The SNMP Surveillance Adapter allows the surveillance of SNMP devices, and the auto-discovery of the devices and their configuration. This Adapter is based on a generic data collection daemon (GENDCD). GENDCD can simultaneously monitor devices of different types. Configuration data for each device type, including configuration parameters, polling and response handling configuration, and trap translation rules, is then provided in a separate set of configuration files.

Several GENDCDs can be deployed on the same workstation to monitor different device types or a separate set of devices of the same type (or any combination). GENDCDs can coexist with the SNMP Integrator on the same workstation, but the same device type must not be monitored by both processes.

GENDCDs collect surveillance data by polling devices using the SNMP protocol. GENDCDs make their data available to registered client processes through an MDM server interface. A GENDCD notifies SNMP management data router (SMDR) when there is a change in an object status. Changes in object status include adding a new component, deleting a component and changing the state of a component.

GENDCD also converts traps received from the trap server to the MDM format, based on trap translation rules defined in the GENDCD trap translation rules files. GENDCD forwards these alarms to SMDR. There can be several GENDCDs on the same workstation and each GENDCD can monitor devices of several types.

SMDR server

You can use the SMDR server to merge the SNMP surveillance data obtained from SMDR-based DCDs. The Preside Multiservice Data Manager (MDM) can manage SNMP devices and does not require HP OpenView.

FMDR server

The Passport fast management information protocol (FMIP) management data router (FMDR) collects and routes alarm and state changes from a logical grouping of Passport switches to the general management data router (GMDR) server. A separate and unique FMDR server instance is required for the management of Passport 15000.

Engineering FMDR size

The failure or restart of an FMDR server implies the loss of surveillance to all Passports within the group. When the FMDR restarts, it checks the states of all the Passports. FMDRs must be restarted when Passports are upgraded in the network. The impact of the restarts can be major if the group is too large. Having smaller surveillance groups controls the impact of upgrades. To have smaller surveillance groups, the FMDR size should not exceed 60 Passports. If there are always two active surveillance paths for each node (two FMDRs managing each Passport, and both feeding to the same GMDRs), there is not any impact if one of the FMDRs needs to restart.

It is easier to distribute FMDR feeds into different GMDRs and achieve regional and central management when the FMDRs are smaller. If all data comes from the same process, post-filtering is required, which is inefficient.

System and application management agents

For any large-scale network management solution, workstation and server management is as important as switch management. Fault integration of the applications, servers, and workstations is often referred to as system management. System management is accomplished by the data manager agent (DMA) that forwards server and workstation alarms to a GMDR server.

GMDR server

The GMDR server collects and routes alarms and state changes to various applications supporting the Preside Multiservice Data Manager (MDM) desktop.

The GMDR server takes inputs from all other servers:

- FMDR
- DPN management data router (DMDR)
- SNMP management data router (SMDR)
- OpenView data access mediator (OVDAM)

GMDR provides domain management over a specific grouping of switches. The GMDR can take an input for lower-level, subordinate GMDR servers and provide the capability for regional operation centers using MDM.

GMDR-GMDR filtering for high cost WAN links

The GMDR-GMDR filter restricts the amount of traffic between hierarchical GMDRs, which is useful on high-traffic, low-bandwidth, high-cost links. This filtering is also useful for customizing component criticality and filtering on this criticality.

Regionalization

The GMDR server collects alarms and events from the Preside Multiservice Data Manager (MDM) servers and sends the information to the desktop application. The desktop application provides the network operator with the network-wide surveillance map.

The GMDR server can route to the subordinate GMDR servers, permitting a hierarchy of regions to be created for surveillance operations. This is advantageous for large network management systems.

The figure “Regional surveillance capabilities” (page 55) shows a WAN configuration with a Passport asynchronous transfer mode (ATM) backbone network being managed from its global headquarters. This network also has two regional operations centers in the west and east that receive surveillance information from the backbone network.

The GMDR server routes surveillance information from the backbone to the regional operations centers so that they can monitor the backbone. The global headquarters has a GMDR server that receives surveillance data from the backbone GMDR server, and both regional GMDR servers which it is monitoring.

Figure 9
Regional surveillance capabilities

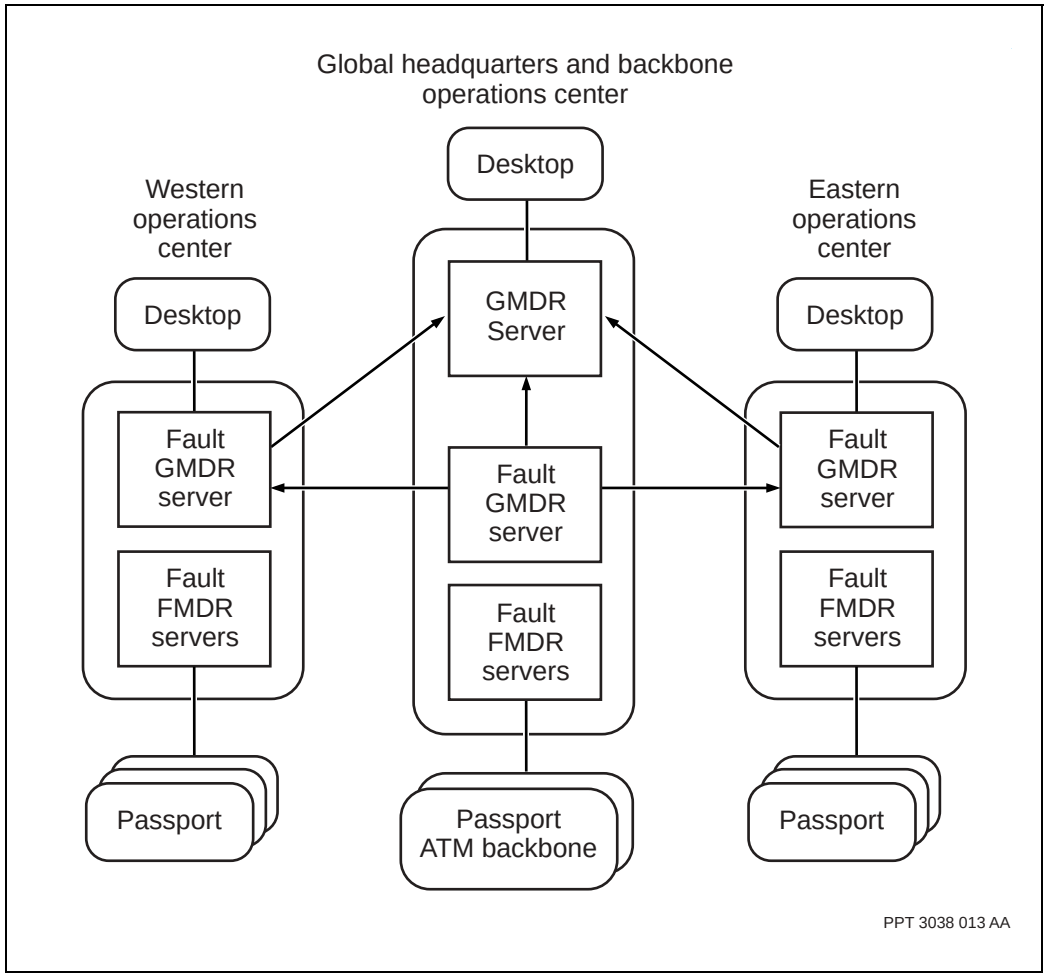

The GMDR server can filter out duplicate surveillance alarms and state events. The GMDR server includes alarm acknowledgement that allows you to know when an operator from any site has acknowledged an alarm. This reduces operations co-ordination efforts and assures an audit trail of actions.

Engineering recommendation

The following rules apply when scaling a large Passport network:

- limit each FMDR server to not have more than 60 Passports
- create each regional GMDR server from less than 15 FMDR servers. These collector FMDRs can then feed into higher-level GMDR regional servers. This preserves the hierarchical nature of MDM. The regional GMDR server can exist on the same workstation as the FMDR servers as long as the workstation has sufficient resources.
- up to 15 regional GMDR servers can then be fed into higher-level regional GMDR servers. There is no GMDR restriction per workstation as long as the CPU and RAM is sufficient.

Surveillance redundancy

The figure “Redundant surveillance server architecture” (page 57) shows a redundant surveillance solution that has a Passport switch send its alarm and state changes to two FMDR servers, each located on a different workstation.

The FMDR servers gather duplicate fault information from the same Passports and pass it to the GMDR server, which performs state calculations and makes the fault information available to network model, and therefore the desktop. The GMDR server discards duplicate fault information.

Duplicated GMDRs are deployed to provide further redundancy.

Figure 10
Redundant surveillance server architecture

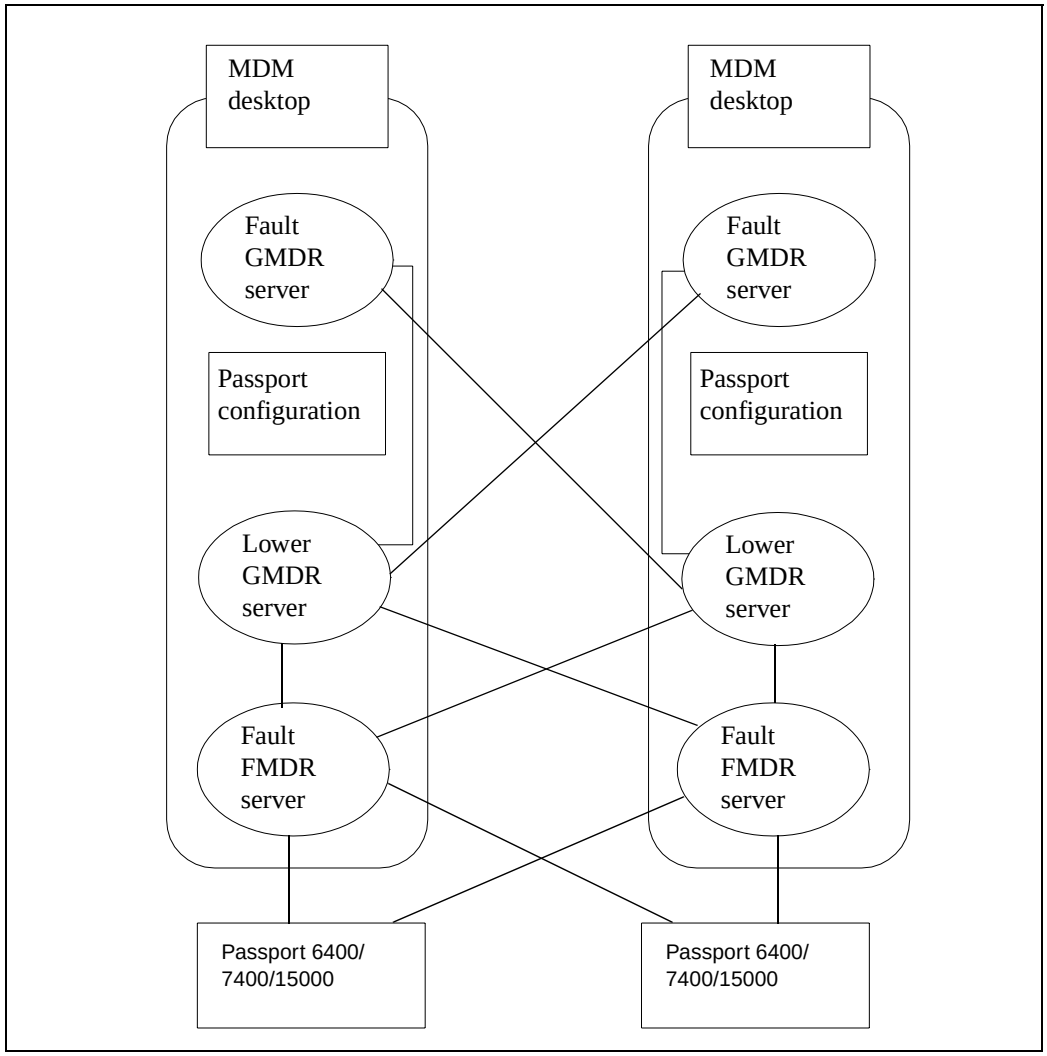

Host Group Directory server rules

The figure “HGDS session servers” (page 59) provides a basic overview of the session server architecture with Preside Multiservice Data Manager (MDM). For improved network performance, a Host Group Directory server (HGDS) containing all network nodes to provide for command console access to all switches in the network. Use a separate HGDS for each FMDR server and the Passport switches supported by that server.

Figure 11
HGDS session servers

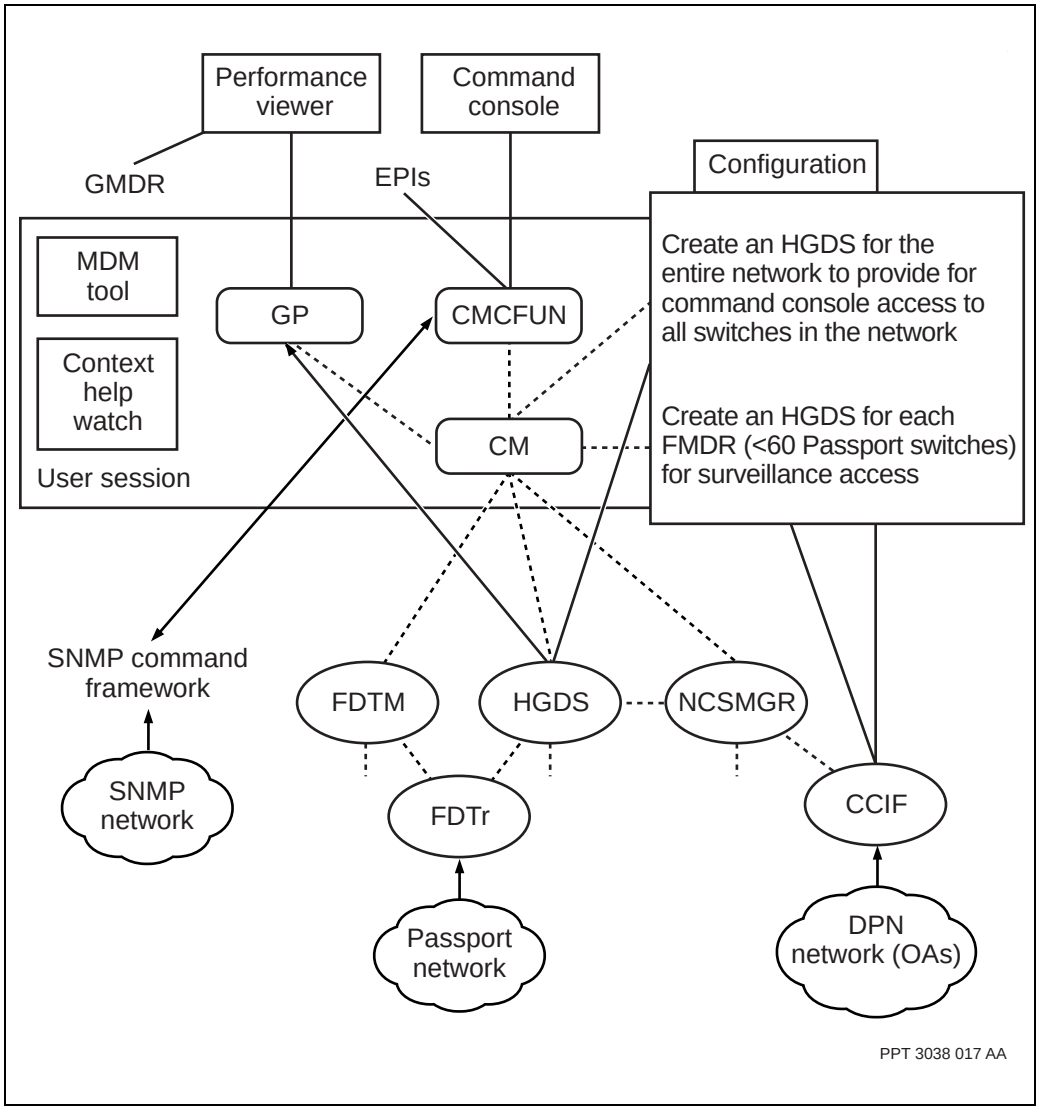

Network data access mediator

The network data access mediator (NDAM) server acts as a distributor of management data to other network management clients, such as the HP OpenView desktop. NDAM also acts as the filter between two GMDR servers. NDAM receives the requests for information from the client applications, and extracts this information from the network model server and GMDR. NDAM then forwards the data to the requesting applications after filtering is applied to them. The figure “NDAM server” (page 60) shows the NDAM server within a network.

Figure 12
NDAM server

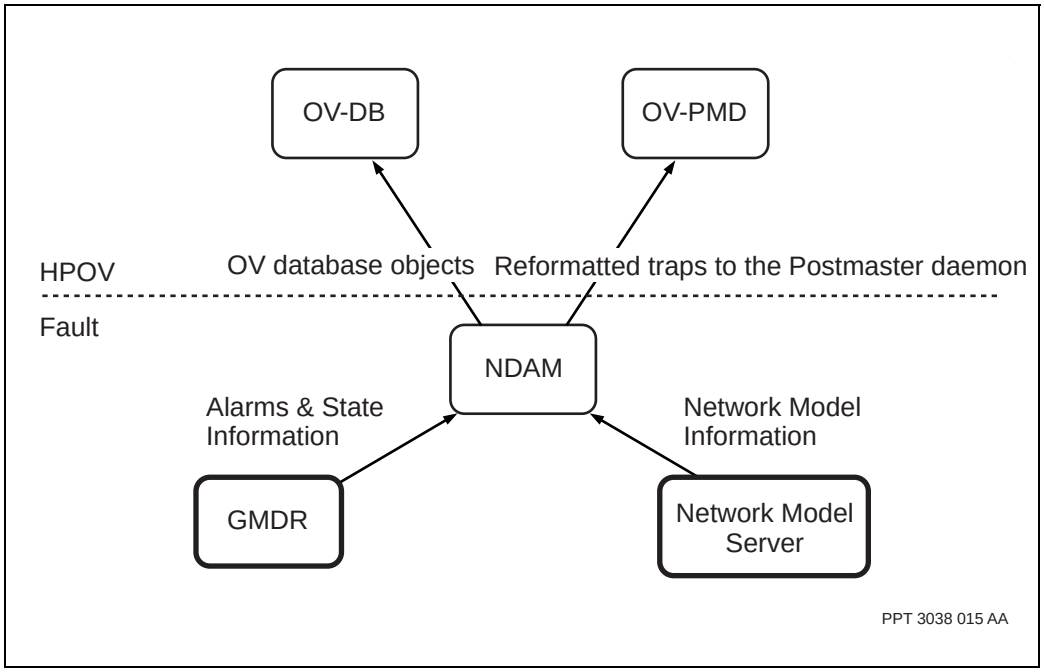

Chapter 5
MDM servers

This section describes the MDM servers that support the three basic functions for networks that contain Passport switches:

- network access lets you log in to Passport switches and perform operations such as provisioning and troubleshooting
- surveillance access lets the MDM software gather surveillance information from Passport switches
- provisioning access lets users configure Passport switches and upload service data descriptions (SDD) to the workstation

This section also contains guidelines for setting up Passport groups, which is part of the server configuration process.

The following information is contained in this section:

For more information on MDM servers, see 241-6001-310 *Preside MDM Server Reference Guide* and 241-6001-303 *Preside MDM Administrator Guide*.

Preside Multiservice Data Manager (MDM) can be used in

- the traditional stand-alone server approach,
- a client-server scenario using the LAN-select mode
- a hierarchical mode configuration to manage a very large network
- a hot-standby mode to provide redundancy and resiliency.

There is no limit to the size of the group, but there is a limit to the number of Passports that a server can connect to.

MDM servers to configure

To perform network, surveillance and provisioning access, you need to configure the following servers:

- Passport communication manager (FDTM)
- Host group directory server (HGDS)
- fast management information protocol (FMIP) management data router (FMDR)
- Data manager agent (DMA)
- General management data router (GMDR)
- Injected management data Router (IMDR)
- Network access mediator (NDAM)
- Network model server (NMSERVER)

Passport groups

A Passport group is a set of Passport switches that share at least one common userID and password for performing network, surveillance, or provisioning access. A Passport group is a group in the configuration files of the Preside Multiservice Data Manager (MDM) software.

Use the Passport group to control access to the administrative functions on a switch. A user logging on with a userID has access to all switches defined in the group, and can perform any administrative function allowed by the userID.

Passport groups control network access for network and surveillance tasks. Network access allows an operator or administrator to log on to a Passport switch with the Command Console or to perform provisioning operations. Surveillance access allows the FMDR server to log on to a group of Passport switches to obtain surveillance information.

A Passport switch can belong to several groups, so that it is accessible by different userIDs for different tasks. For example, a Passport switch can be accessed by an operator for surveillance, and by a network administrator for provisioning.

Groups of Passports for network access

You can define groups that allow users to access all Passport switches in a group, and to perform operations such as provisioning or troubleshooting. Depending on the group ID logged into, a user can access the group of Passports based on the capability defined in the group ID.

Guidelines for grouping Passports for network access

The guidelines for grouping Passport switches to provide network access are as follows:

- You must define at least one common userID and password on all Passport switches in a group for performing network access functions. This common userID and password must authenticate in the same way on all Passport switches in the group. The userID and password must be defined with the same capability on all of the Passport switches, and all of the switches must return the same customer network identifier (CNMID).
- You can define several common userIDs and passwords on the Passport switches in a group and dedicate each to a different function. For example, one userID can have access privileges for performing maintenance functions, however, a common userID and password must authenticate in the same way on all Passport switches in the group.
- The same Passport switch can be used in more than one group.
- There is no limit on the size of the group, but there is a limit to the number of Passports that the server can connect to.
- The default maximum number of Passport switches in a Passport group that is used for network access is 60.

Groups of Passports for surveillance access

This section describes how surveillance information is obtained from the network. To obtain surveillance information, the following sequence occurs:

1 The FMDR server on a Preside Multiservice Data Manager (MDM) server logs in to all of the Passports in a surveillance group with a common userID and password.

2 Each Passport switch authenticates the userID and password, and returns a CNMID.

3 To perform its filtering function, an FMDR server needs to receive surveillance information from all of the devices on all the Passports in the surveillance group. For an FMDR server to receive the information, you must define a common userID and password on all Passport switches. This is required so that the userID and password can obtain the required surveillance information and that it causes all Passports to return a CNMID of 0.

4 Once logged in, the FMDR server is ready to receive alarms and status records automatically from all Passports in the surveillance group.

5 To obtain surveillance information from an FMDR server, a client application, such as the GMDR server, registers with the FMDR server. This registration request is done by a userID and password authentication process.

6 The FMDR server then passes the userID and password obtained in the registration request to one of the Passports in the surveillance group for authentication. If successful, the Passport returns a CNMID to the FMDR server. The FMDR server stores the CNMID for filtering purposes.

7 The FMDR then obtains the states of all components that it surveils. FMDR also obtains information about links that terminate on TRK components and DPNGATE components and sets the initial state of these links to in-service.

8 When the setup is complete, a Passport switch forwards surveillance information to the MDM workstation. The FMDR server filters the surveillance information for the GMDR according to the GMDR's stored CNMID.

9 For GMDR to receive surveillance information from all Passport devices in the surveillance group, the userID and password provided by GMDR must cause the Passports to return a CNMID of 0. For virtual private networks (VPN) in which you only receive information about the devices in the VPN, the userID and password must cause the Passports to return a CNMID other than 0 that is unique to the customer VPN.

Guidelines for grouping Passports for surveillance access

Follow these guidelines for grouping Passports for surveillance access:

- define at least one surveillance group
- use the same Passport switch in more than one group
- ensure that there is one FMDR server for each surveillance group or two for redundancy on two Preside Multiservice Data Manager (MDM) workstations with separate connectivity.
- ensure that the names of surveillance groups are unique on a given workstation. You cannot have two groups with the same name on the same workstation. However, you can duplicate the names of surveillance groups on different workstations.
- do not create groups containing more than 60 Passport switches

FMDR server redundancy for surveillance access

To achieve redundancy, create duplicate surveillance groups on each of the workstations and run a separate FMDR server on each workstation. Then, use the GMDR Administration tool to set up the GMDR server on each workstation to gather surveillance information from the FMDR servers on both workstations.

The GMDR server receives alarms from the FMDR servers on both workstations. The GMDR server only displays the alarms once because it discards duplicate alarm notifications. If one of the FMDR servers fails, the GMDR server continues to receive data from its redundant FMDR servers.

The figure “FMDR server redundancy for surveillance access” (page 65) shows a network containing three Passport switches that are monitored by two stand-alone workstations connected by a LAN. Identical groups called G1 are defined on both workstations. Separate FMDR servers retrieve surveillance data from the groups.

Each GMDR server receives surveillance data from the FMDR server on its own workstation and from the FMDR server on the redundant workstation through the LAN connection.

The GMDR server on workstation A discards duplicate data from the FMDR servers. If FMDR fails on workstation A, the GMDR server on workstation A can obtain the same surveillance information from the redundant FMDR through its LAN connection to workstation B.

Figure 13
FMDR server redundancy for surveillance access

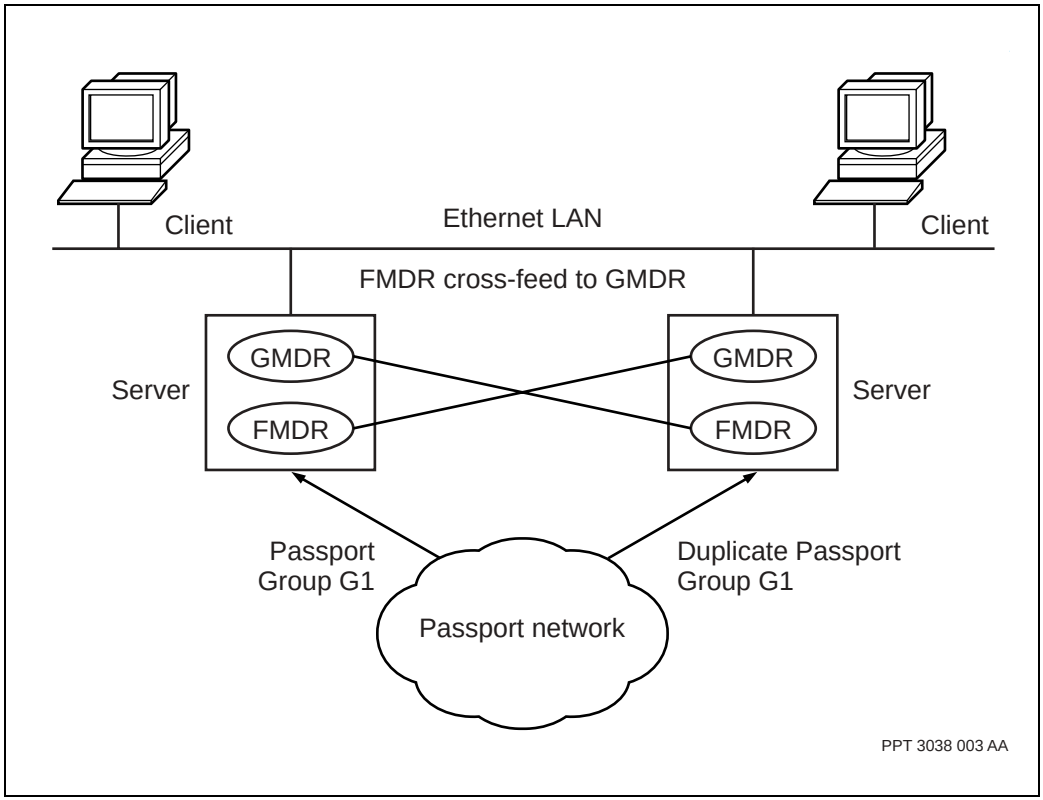

Distribution of servers in large networks

This section describes the distribution of servers among workstations on a LAN in large networks. For small networks, the servers that support Passport network, surveillance, and provisioning access can run on the same workstation.

For medium and large networks, you can deploy servers among workstations connected by the same Ethernet LAN or by a WAN Internet protocol (IP) connection. This is done for several reasons, including the following:

- distribution of the workload over several workstations to improve performance
- effective use of older, less powerful workstations along with newer and more powerful workstations
- redundancy and resiliency for fault management

Guidelines for deploying servers over multiple workstations

The following guidelines apply to deploying the servers for Passport network, surveillance and provisioning access over multiple workstations:

- the HGDS and FDTM servers must run on a workstation that provides network access through an X.25 or frame relay link to the network
- the FMDR server must run on the workstation that provides network access. You can run the server on another workstation as part of the FMDR server start-up command. You must specify the hostname of the workstation that runs the network access server.
- the GMDR server can run on any workstation on the LAN, provided the workstation can handle traffic to the server. To ensure that GMDR receives surveillance information, use the GMDR Administration tool to specify the FMDR server from which the GMDR server obtains the surveillance information.
- configure the IMDR server as a subserver of GMDR

NDAM server

The NDAM server provides clients such as HP OpenView Desktop with access to the Preside Multiservice Data Manager (MDM) surveillance information. NDAM performs filtering according to component type or geographic region for the HP OpenView Desktop and fault tools.

The servers collect, interpret, and concentrate the management data from the network. Clients can access the information collected by the data collection servers from the GMDR server and the NMSERVER servers. This information can be forwarded to clients such as HP OpenView Desktop.

If all of the information is available, users are subjected to excessive amounts of information, which is not always useful. Several methods to reduce the information forwarded to users or to hierarchical GMDR servers, include the following:

- through component criticality thresholds, you can configure the following thresholds:
  - — by specifying thresholds when setting up hierarchical GMDR servers
  - — by supplying parameters in the startup command for the Surveillance Network Model Updater (SURNUP) server
- component type and regional filtering performed by an NDAM server

For more information on component criticality thresholds, see “Component criticality thresholds” (page 69). For more information on component type and regional filtering, see “Component type and regional filtering” (page 70).

The figure “Filtering based on thresholds” (page 69) shows how filtering is used for regional-central network management.

Figure 14
Filtering based on thresholds

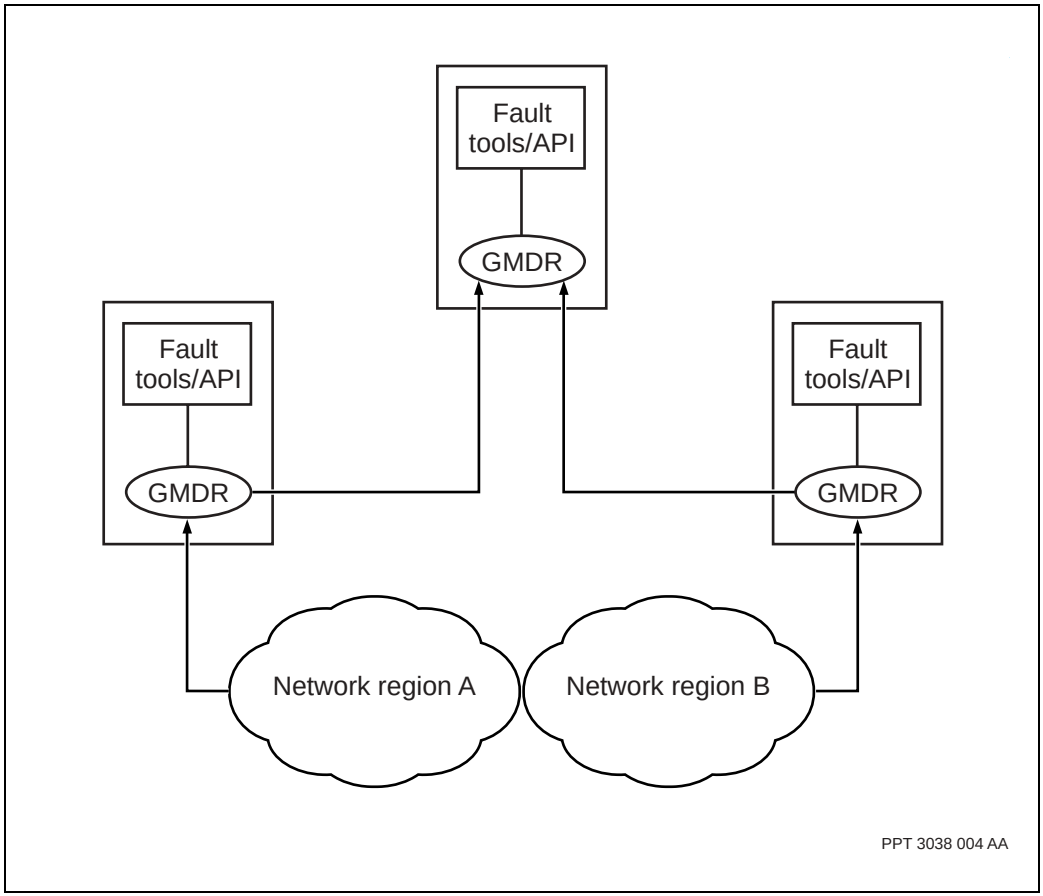

Component criticality thresholds

A subordinate GMDR server assigns a criticality value to all components it manages. When a superior GMDR server connects to a subordinate GMDR server, the superior GMDR server can supply a component criticality threshold value. The subordinate GMDR server only provides management data for components whose faults pass a threshold test. Thresholds allow the deployment of regional management centers that can see all devices in the network. These regional management centers only get information for the most important sub-components that are controlled by the criticality

threshold. You can customize the component criticality assignments by modifying the GMDR criticality schema and by adding exceptional mappings to its criticality overrides configuration file.

Component type and regional filtering

Component type filtering lets you specify the type of module, subcomponent, and link types for which a client can receive management data. Regional filtering lets you subdivide the network into different regions and only supply a client with information from the devices in a region. The NDAM server provides regional filtering through its network data access mediation capabilities.

Filtering can be set up in two ways:

- by specifying a list of type and device filter sets and individual overrides at connection time for HP OpenView Desktop clients
- by forcing authentication and filtering for clients

You can deploy the NDAM server as

- a superior GMDR server
- a subordinate GMDR server
- a proxy GMDR server (in place of a GMDR server)

For HP OpenView Desktop clients, the NDAM server provides combined access to the GMDR database and Network Model information. For clients, the NDAM server provides access to the GMDR server information only.

NDAM supports all the capabilities of GMDR, including filtering. NDAM is always associated with a single GMDR server and its clients. Unlike GMDR, NDAM does not store information in a memory. NDAM passes all queries to its GMDR and Network Model servers, and filters the replies according to the associated filter sets.

NDAM also uses a single notification stream from the GMDR and Network Model servers, filters it, and multiplexes it for NDAM clients. For example, the NDAM server receives an alarm only once from GMDR, but it can forward the alarm to many clients. Like GMDR, NDAM supports component criticality filtering, so both forms of filtering can be combined. This lets you

divide the network into several regions, one of which represents the network backbone. You can then access data for each region with different criticality thresholds as follows:

- connectivity information from the backone
- full regional information from the regional centers
- hardware and connectivity information from the regions
- full backbone information for the central operations center

NDAM also supports service name aliasing as shown in the figure “Service name aliasing” (page 72). Service name aliasing allows a single NDAM server to act as multiple subordinates to the same GMDR server. Each connection has a different filterset and criticality threshold mapping.

Figure 15
Service name aliasing

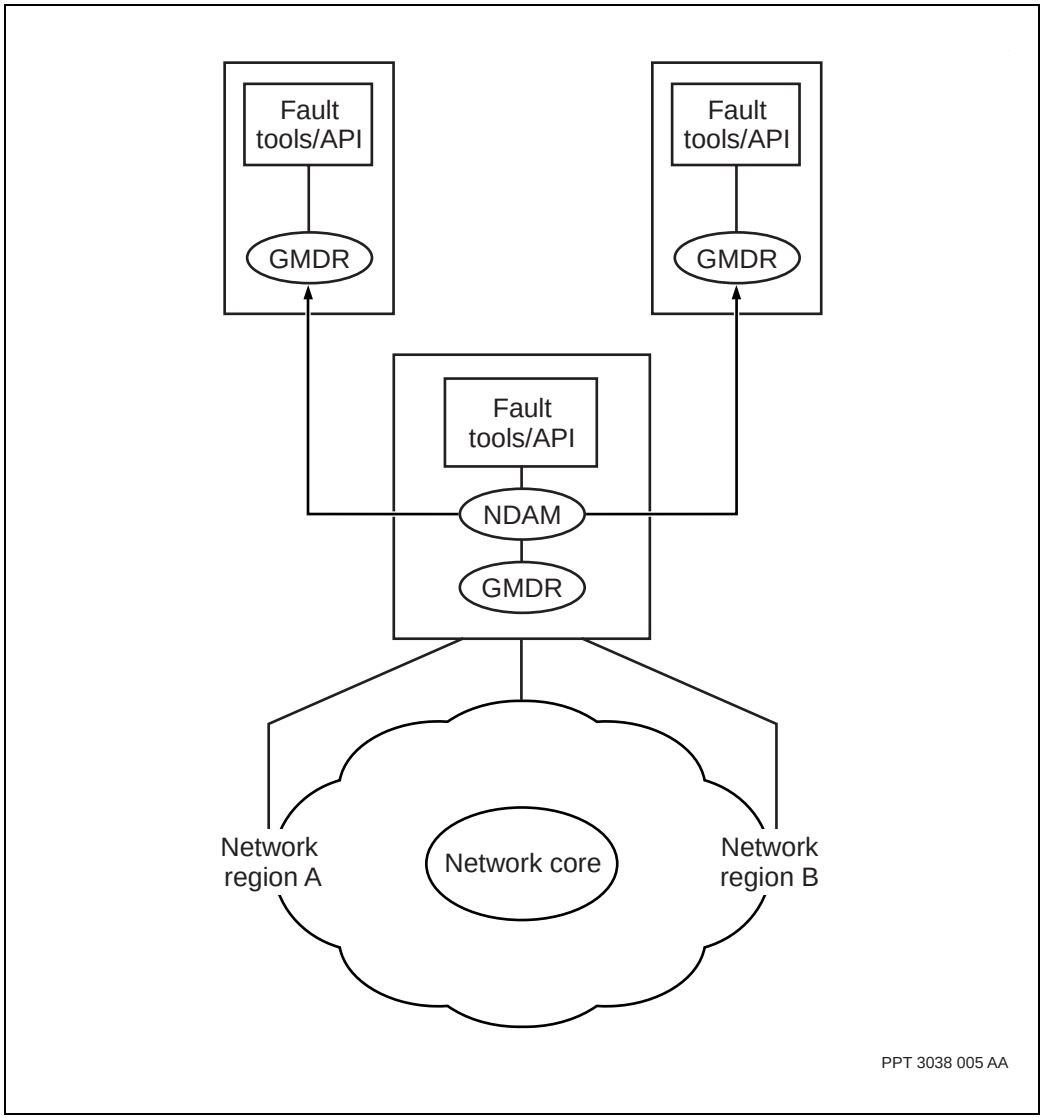

NDAM deployment and configuration

To deploy NDAM for HP OpenView Desktop clients, start the NDAM server with the GMDR server option (-g), and the Network Model option (-m). If multiple NDAM servers need to run on the same workstation, you can configure multiple NDAM servers. You must assign each one a separate service name by using the -n option.

> ***Note:*** More more information on the NDAM startup command, see 241-6001-310 *Preside MDM Server Reference Guide*.

You can configure NDAM as a subordinate GMDR server or as a proxy server for GMDR. For more information, see the following:

- "Configuring NDAM as a subordinate GMDR server" (page 73)
- "Configuring NDAM as a proxy server for GMDR" (page 75)

Configuring NDAM as a subordinate GMDR server

When the NDAM server is used with a subordinate GMDR server, NDAM acts as the intermediary for two hierarchical GMDRs. A superior GMDR server obtains information from a subordinate GMDR. The subordinate GMDR server is capped by an NDAM server which is configured as a subserver to the superior GMDR server.

The figure "Deploying the NDAM server between GMDR servers" (page 74) shows how you can use the GMDR as a subordinate server.

Figure 16
Deploying the NDAM server between GMDR servers

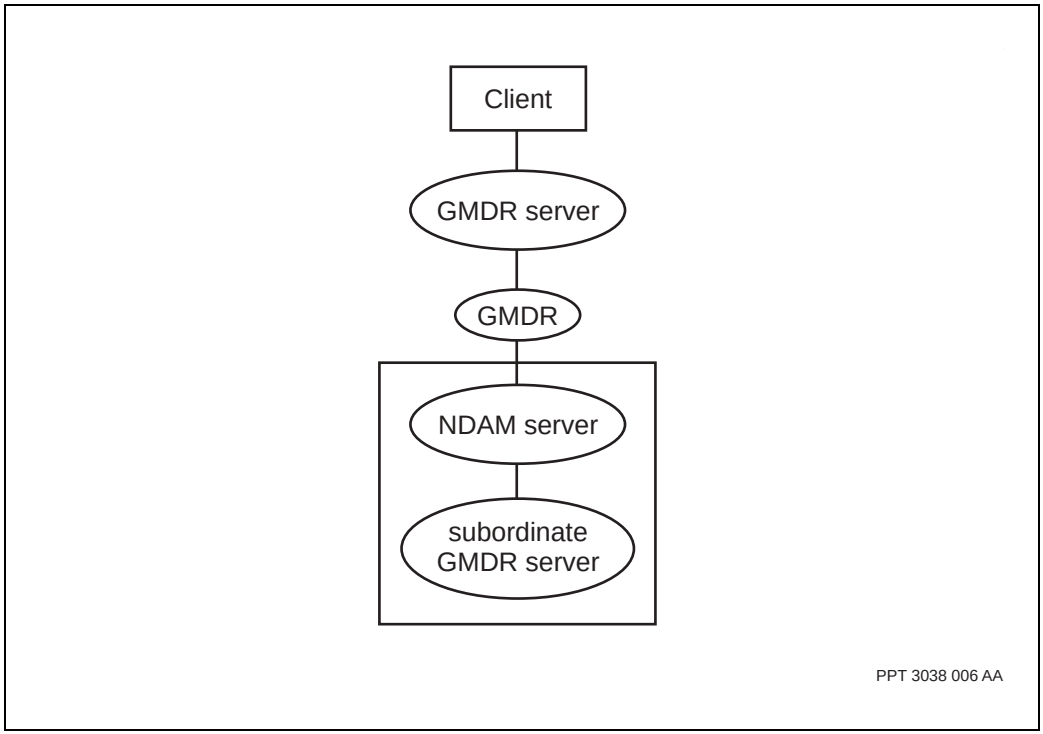

Deploying the NDAM server between GMDR servers allows the superior GMDR server to gather information that is filtered, according to region or component type, from subordinate GMDRs. GMDR servers provide this information to its clients.

To use this configuration, start the NDAM server with the following options:

- GMDR option (-g)
- -m option with a value of "~". You must use the quotation marks. The -m"~" value specifies that no network model information is involved.
- -s option to set up forced authentication

You cannot configure the NDAM authentication information on the superior GMDR server to specify the filter sets explicitly. You must assign them implicitly through the authentication mechanism.

You can configure the NDAM server as a subordinate to the same superior GMDR server. This allows the extraction of different file set/criticality threshold filtering combinations from the same GMDR server. You must specify multiple service name aliases with the -a option, and each alias must include a prefix of GMDR.

Configuring NDAM as a proxy server for GMDR

The NDAM server does not support the GMDR admininistration interface when you use it as a replacement server. You must connect the GMDR Administration tool to the real GMDR server used by the NDAM server. If you use NDAM used as a proxy, GMDR does not support the inbound API capabilities of the alarm and status API.

To inject alarms, you must connect directly to the underlying GMDR server. Some of the alarm and status API attributes are not available because NDAM does not have an object model of its own to which these filters can be applied.

NDAM authentication configuration

You can use NDAM in plain or forced-authentication mode. In plain mode, NDAM requires a REGISTER message. The superior GMDR servers provide the alarm and status API REGISTER message, while the Fault tools implicitly provide this message. Like GMDR, NDAM ignores its contents. In forced-authentication mode, NDAM checks the authentication information against a list of pre-configured authentication, and only accepts matching entries. The authentication information also contains a list of device and type filter sets to be automatically applied to the client connection. This allows the server side of the MDM fault stack to control the filtering applied to specific client connections.

To support tools without deploying an additional GMDR server superior to NDAM, you can use wildcard authentication.

NDAM filter set file configuration

There are type set and device set filter set files. Type set files define the types of components that NDAM reports data on. Device set files contain lists of devices and can be used to divide the network elements into geographic regions. Device set files must list the modules that are part or not part of the region and support standard style pattern matching for more efficiency.

You can specify the type set files by the -f option to NDAM or without the -f option, which is the default option. The filter set files are specified as in the network model API by their highest and lowest categories only.

Chapter 6
Choosing a configuration for MDM

After determining the network requirements, you need to determine the Preside Multiservice Data Manager (MDM) configuration that best suits those requirements. This section describes the ways in which you can configure the workstations for MDM, and provides you with a method to evaluate their merits.

This section contains the following information:

Types of configurations

The main types of configurations are as follows:

- stand-alone
- stand-alone CPU server
- client set/server set
- network file system (NFS)
- combination

The main factors to consider in choosing a configuration are as follows:

- network size

  Be sure to consider both the network and workstation point of view. For the network, consider the number of modules and traffic. For the workstation, consider the system throughput and number of users.

- physical locations and organizational responsibilities of administrative staff

  Review the size and number of different locations and the number of different Preside Multiservice Data Manager (MDM) support staff, particularly in a widely distributed network. If staff at most sites require access to all MDM functions, server configurations can yield workstation administrative savings at the cost of additional links.

- present network operational model

  If provisioning or surveillance, or both, are centralized functions performed by a number of different people, one or a combination of the server configurations can result in workstation cost savings.

Review of these issues determines whether to choose

- a stand-alone configuration
- a stand-alone CPU server configuration
- one or both of the server configurations for MDM
- a combination of configurations

Stand-alone configurations

In a stand-alone configuration, one or more workstations are connected to the network and each workstation runs all of the Preside Multiservice Data Manager (MDM) software processes independently. Each workstation also has a connection to the network for managing switches in the network.

The type of link used for accessing the network depends on the composition of the network (Data Packet Network (DPN) only, DPN, Passport, Passport only, simple network management protocol (SNMP) devices, or a combination of these). For information about the types of network access, see "Types of network access" (page 93).

A typical configuration is shown in the figure "Typical stand-alone configuration" (page 79).

Figure 17
Typical stand-alone configuration

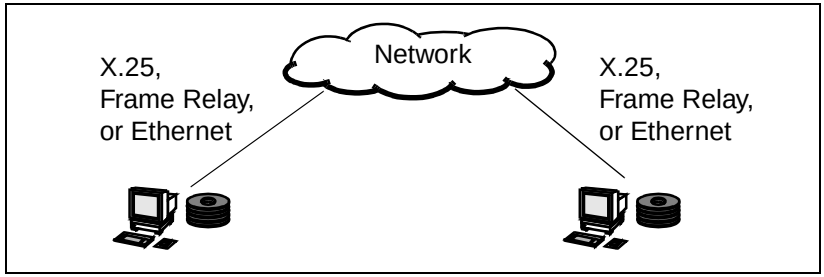

For a new, very small network, only one stand-alone workstation is required to manage the entire network.

As the network expands, additional stand-alone workstations can be used for performing specialized tasks. For example, one stand-alone workstation can be used for surveillance, and another can be used for provisioning. There are practical limits to this scheme, as described in "Practical limits to stand-alone CPU server configurations" (page 82).

Stand-alone server model configuration

The figure“Typical stand-alone server configuration” (page 80) shows a grouping in a network that contains five Passport switches. The switches New York1 and New York2 are located in New York. The switches Paris 1, Paris 2 and Paris 3 are located in Paris.

Figure 18
Typical stand-alone server configuration

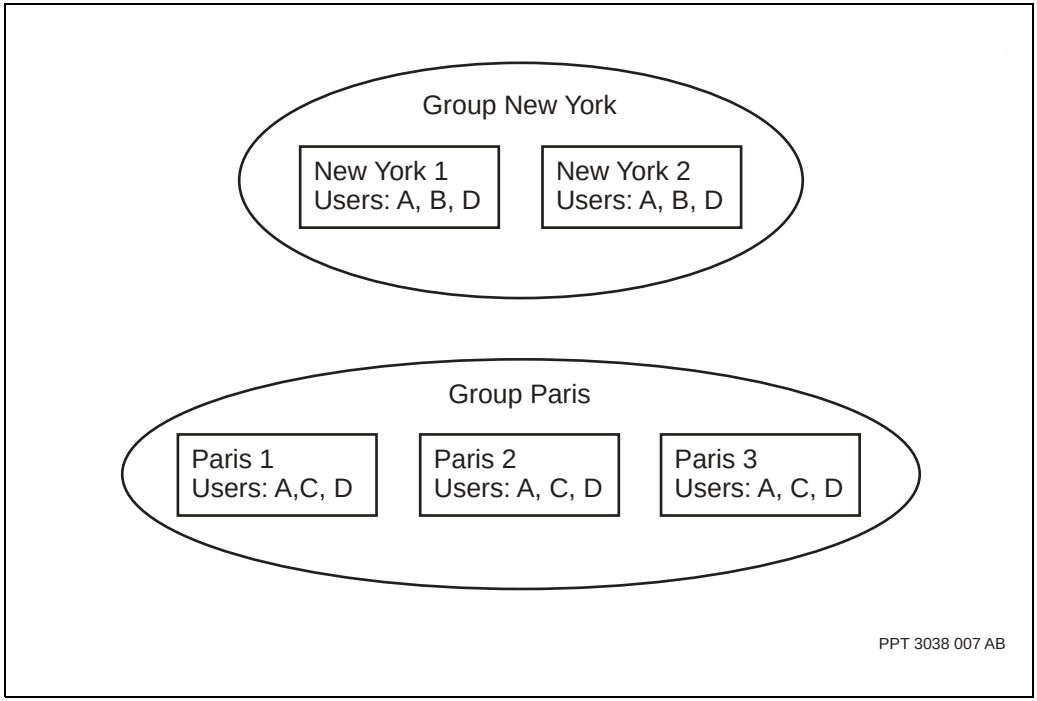

This network has the following administrative requirements:

- User A needs to access all switches in the network for surveillance.
- User B needs to perform provisioning on all of the switches in New York. User C needs to perform provisioning on all of the switches in Paris. Node provisioning is performed locally.
- User D needs to access all switches in the network for network management purposes.

Administering this network requires the following three groups:

- a group that contains all the Passport switches in the network that are accessible by Users A and D
- a group that contains only the New York-based switches that are accessible by User B
- a group that contains all the Paris-based switches that are accessible by User C

Stand-alone CPU server configurations

One or more of the stand-alone workstations is set up to allow users on different platforms to log in and share its CPU to run the Preside Multiservice Data Manager (MDM) software. This configuration allows existing HP-UX or Sun Solaris hosts to be used as terminals to access the workstations by means of an X-11 Release 5 windowing software package. All processing of MDM applications is still performed on the workstation because it is in a plain stand-alone configuration. The clients hosting off the workstation perform window management functions, but do not have a local window manager. A typical configuration is shown in the figure "Typical stand-alone CPU server configuration" (page 81).

Figure 19
Typical stand-alone CPU server configuration

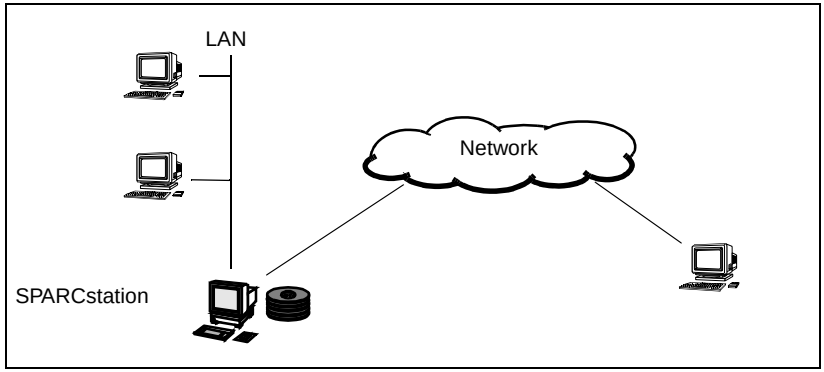

Practical limits to stand-alone CPU server configurations

Stand-alone and stand-alone CPU server configurations become less practical as the number of modules in the network, and the number of Preside Multiservice Data Manager (MDM) workstations, increase. This occurs because of the following reasons:

- the extra effort (and cost) of administering each MDM individually. For example, administering the network model, or updating and maintaining the network link access configuration, requires more effort.
- the additional cost of installing and administering network access software network links, and any high-speed access hardware (such as a high-speed access card) on each workstation. The addition cost results regardless of whether the workstation uses the full bandwidth on the connection to the network.
- the decrease in workstation and link performance as the network size, and the clients running on a stand-alone workstation increase
- the absence of redundancy

  Note: You can compensate for some failures, such as the loss of network access, but not total workstation failure. You can set up two or more stand-alone workstations to belong to the same multi-nodal name server (MNSD) level 2 domain, and use the Service Selection tool to access the software processes that support the service area that has failed (network access, network model, and so on). This scheme does not provide full redundancy.

Stand-alone configurations and stand-alone CPU server configurations are best suited to small networks containing up to 200 modules, in which from one to three MDM users require workstation access.

See “Choosing a workstation to run MDM” (page 95) to select a workstation to suit stand-alone and CPU server configurations.

Client set/server set configurations

The Preside Multiservice Data Manager (MDM) software is structured as a set of processes that communicate by means of transmission control protocol/ Internet protocol (TCP/IP). You can deploy processes among MDM workstations that are connected by an Ethernet LAN (or an IP over WAN connection), and still process interaction.

MDM software processes can be divided into client set processes and server set processes, as shown in the figure "Client set and server set workstations" (page 84):

- Client set processes support user sessions and provide the means to access and display the MDM tools. Client set processes require access to the server set processes. Workstations that run client set processes are referred to as client set workstations.
- Server set processes provide client set processes with services such as the following:
  - — access to the network
  - — gathering, processing, and supplying surveillance information from the network

Workstations that run the server set processes are referred to as server set workstations.

Figure 20
Client set and server set workstations

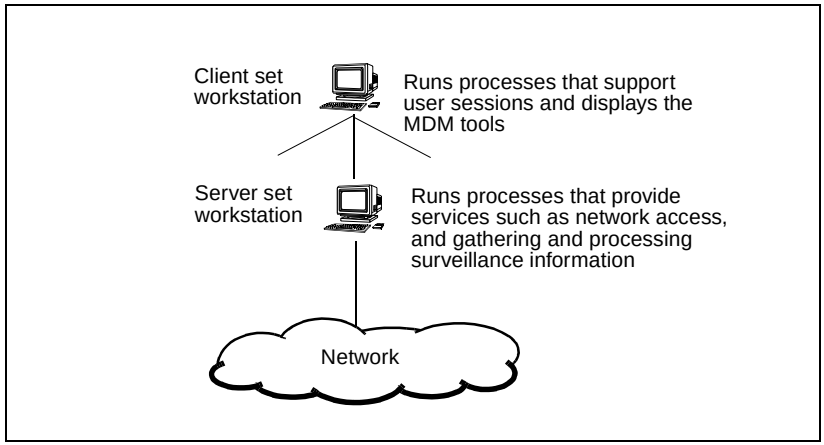

A workstation can be configured to run only the client set processes, only the server set processes, or both of these at the same time (stand-alone workstations).

The Service Selection tool lets an operator or administrator at a client workstation access the server set processes on several server set workstations and server set processes on the workstations to use for one of the following areas of service:

- surveillance
- network model
- DPN configuration
- DPN network access
- Passport network access

See the figure “Service Selection” (page 85).

Figure 21
Service Selection

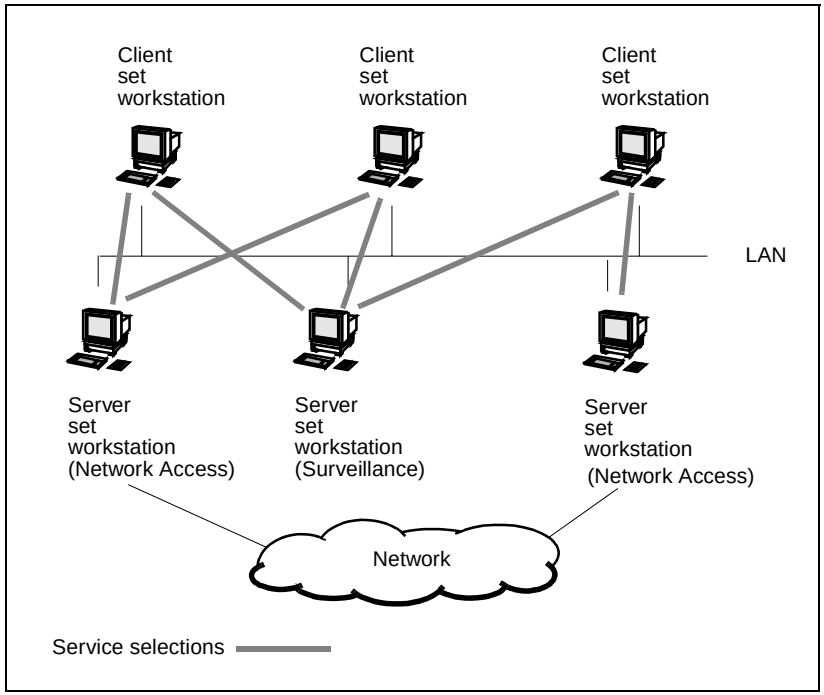

The Service Selection tool allows an administrator to set up the default server set workstations from which a client set workstation obtains support for one of the service areas. An operator can override these defaults when the following happens:

- a server set workstation is no longer able to support an area of service
- the operator is unable to access another part of a regionalized management network

Setting up default server set workstations lets you balance loads on processors and links by specializing server set workstations for different areas of service. For example, you can balance loads on the network access links by using the Service Selection tool to set up the following:

- one server set workstation as the default for providing network access to some client workstations

- another server set workstation as the default for providing network access to other client workstations

A typical server set and client set configuration with two server set workstations and several client set workstations is shown in the figure "Typical client set/server set configuration" (page 86). The client set workstation can also act as a CPU server for X-terminals or PCs. You can set up server set workstations as redundant server set workstations to provide client workstations with an alternate service source.

Figure 22
Typical client set/server set configuration

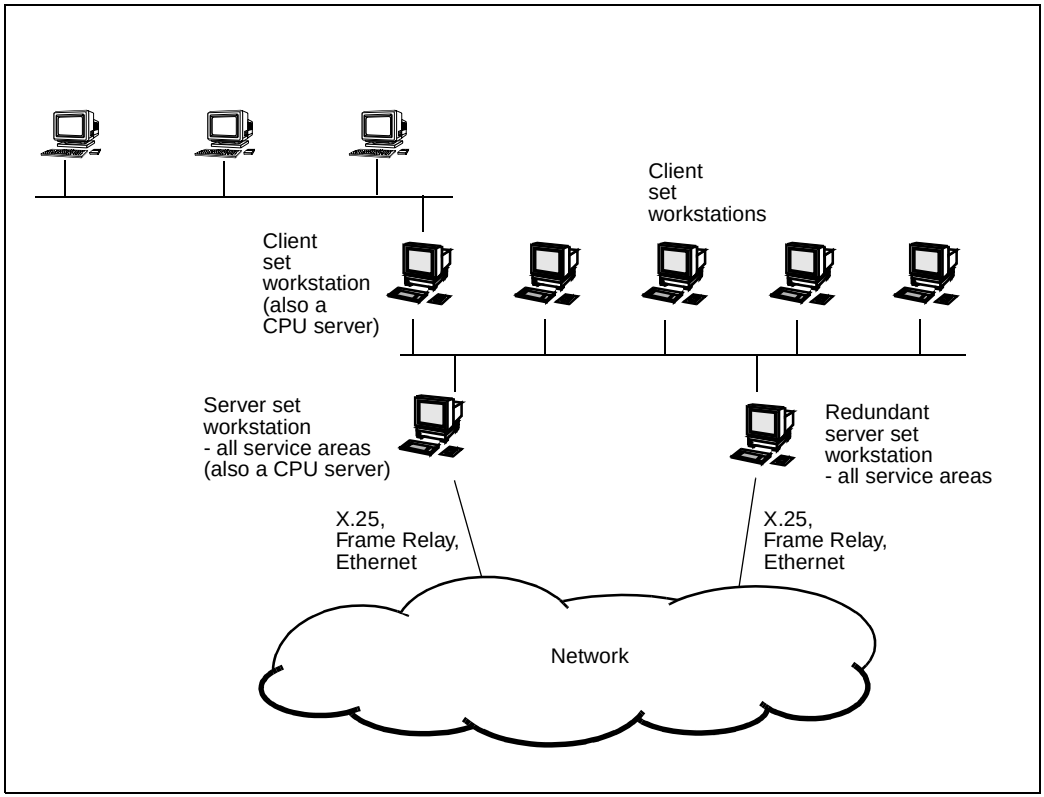

The client set and server set configuration is suited to medium and large networks with more than 200 modules. This configuration provides the following advantages:

- decreases the processing requirements on individual workstations, and improves system performance by distributing processing tasks across workstations
- lets you prolong the usefulness of older, lower-performance workstations by re-deploying them as client set workstations, as the network expands
- reduces costs for purchasing and administering network access hardware and software by requiring it only on the server set workstations
- lets you provide redundancy in your network

NFS server configurations

In a network file system (NFS) server configuration, one or more workstations on a LAN are set up as NFS servers, and the others as NFS clients. NFS allows client workstations to access files on the server in a manner that makes the files appear as if they are being stored locally on the client workstation. With this configuration, client workstations perform Preside Multiservice Data Manager (MDM) application processing, but rely on the NFS server for access to the MDM software. A typical NFS server configuration is shown in the figure “Typical NFS server configuration” (page 88).

In any NFS server configuration the following applies:

- client workstations need to be on the same LAN (or IP over WAN links) as the NFS server. If a WAN link is used, it must have enough bandwidth to provide adequate workstation response.
- for better performance, NFS client workstations need to have local swap space, and enough disk space to hold the operating system, the MDM configuration and data files
- an NFS server can also act as a server set workstation or as a CPU server

The NFS server configuration is suited to medium and large networks of over 200 modules, and lets you do the following:

- perform software upgrades more easily by replacing the software on the server. For NFS client workstations to access new MDM software on the NFS server, you need to
  - — establish a mount to the NFS server
  - — run a script to establish the links to the new software
- perform rollbacks easily by running a script on the NFS client to re-establish links to the old software
- extend the life of older client workstations with small disks

Figure 23
Typical NFS server configuration

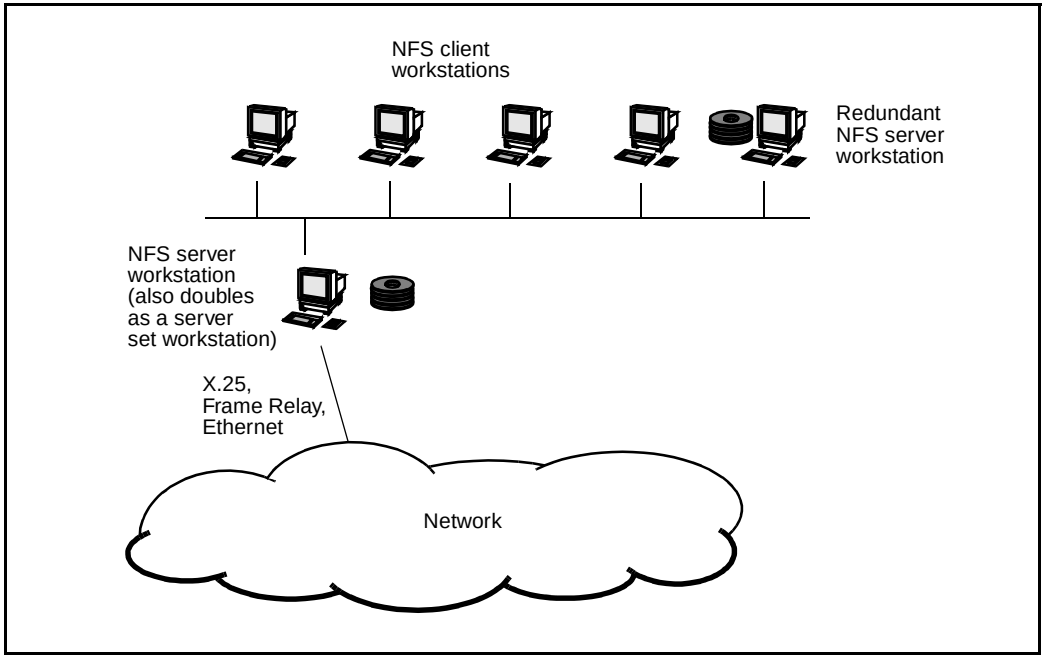

Combination configurations

You can mix stand-alone, CPU server, client set and server set, and NFS server configurations to a certain extent. Most stand-alone, client set, or server set workstations can be used as a CPU server. Ensure that the workstation has sufficient performance to handle the extra X-terminals that are logged in to it.

You can also mix client set and server set configurations with NFS server configurations. A server set workstation can also double as an NFS server provided it has sufficient performance to service the needs of its clients.

Advantages and disadvantages of server configurations

There are advantages and disadvantages in using the CPU server, client set and server set, or NFS server configurations. These are described in the following paragraphs.

Advantages

The primary advantage of the CPU server, client set and server set, or NFS server configurations is reduced cost in initial hardware expenses and in ongoing administration expenses.

Reduced hardware costs

Depending on the server configuration you choose, most or all of the Preside Multiservice Data Manager (MDM) software is stored only on the server workstations. Client workstations do not require disk space to store the MDM software. A file server reduces the overall need for disk space in the network, and reduces the network cost.

CPU and memory savings are possible especially in networks that use the CPU server configuration, the client set and server set configuration, or both of these. The client workstations only run a portion of the MDM software, and require less memory and less powerful CPUs. You can use older, less powerful workstations as MDM client workstations.

Reduced administration expenses

In a CPU or NFS server configuration, only load the software on to the server from tape or compact disk, and not onto all workstations in the network. The time and cost required to install and upgrade MDM software using these configurations is reduced.

You can also centralize the administrative functions. In a widely distributed network, you can use servers and the Sun Administration Tool Suite to centralize your operating system, and reduce the need for system administrators.

Disadvantages

The greatest disadvantages to server configurations are network reliability, response time, and complexity.

Network reliability

The greatest disadvantage of using a server configuration is the dependency of the client workstations on the server for the Preside Multiservice Data Manager (MDM) software. If a CPU server crashes, or is being rebooted, the client workstation does not work because the software it needs is running on the CPU server. This can be offset in client set, server set, and NFS server configurations by providing redundant backup server set or NFS server workstations.

Response time

Because client workstations use a LAN, an X.25 link, or a frame relay link to access Preside Multiservice Data Manager (MDM) software, the response time on client workstations is slower than stand-alone workstations.

If a CPU server configuration is being used, the performance of the CPU server workstation drops as clients log in to it to run the MDM software.

Complexity

Server configurations can require substantial planning and monitoring to ensure that they provide adequate performance. To communicate between workstations, and multi-nodal naming service (MNSD) level 2 domains and using the Service Selection tool, server configurations can require the following:

- LANs
- X.25
- IP over X.25
- frame relay connections

Counteracting the disadvantages of server configurations

With careful consideration of the advantages and disadvantages, you can successfully configure cost-effective, reliable Preside Multiservice Data Manager (MDM) networks using server configurations. For example, you can minimize the impact on network reliability by building a redundant network that includes several file servers and some stand-alone workstations. You can minimize the impact on response time by using high-speed X.25 links to DPN nodes, frame relay or Ethernet links to Passport nodes, and by monitoring your MDM workstation as outlined in “Monitoring MDM” (page 99).

Building in redundancy

Because server failure causes the file server and all clients to fail, you must build in redundancy to conserve network reliability. You can build in redundancy by configuring some of the workstations in the network to be stand-alone, or by configuring redundant servers.

Server configuration with backup stand-alone workstations

In this configuration, some of the workstations are configured in a server configuration. Network redundancy is provided by stand-alone workstations that can be used to maintain the network in case of file server outage. A typical configuration is shown in the figure “Typical server configuration with backup stand-alone workstations” (page 92).

Figure 24
Typical server configuration with backup stand-alone workstations

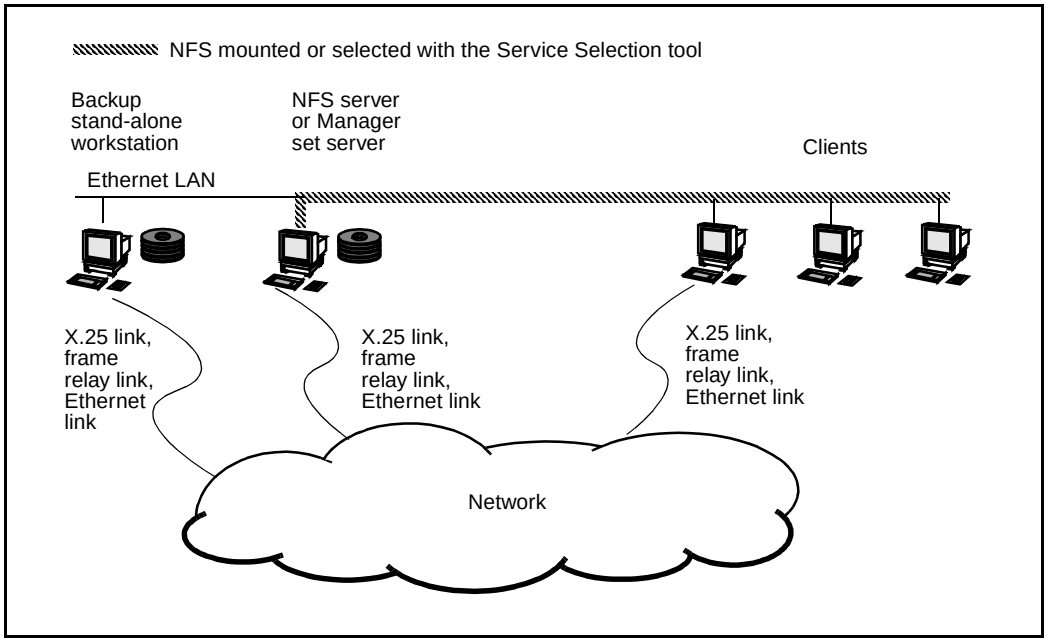

Server configuration with redundant servers

In this configuration, the network is configured with at least two servers. Each server handles a portion of the clients in the network. If one of the servers fails, the clients can be reconfigured to use the Preside Multiservice Data Manager (MDM) software from another server in the network. A typical configuration is shown in the figure “Typical server configuration with redundant server workstations” (page 93).

Figure 25
Typical server configuration with redundant server workstations

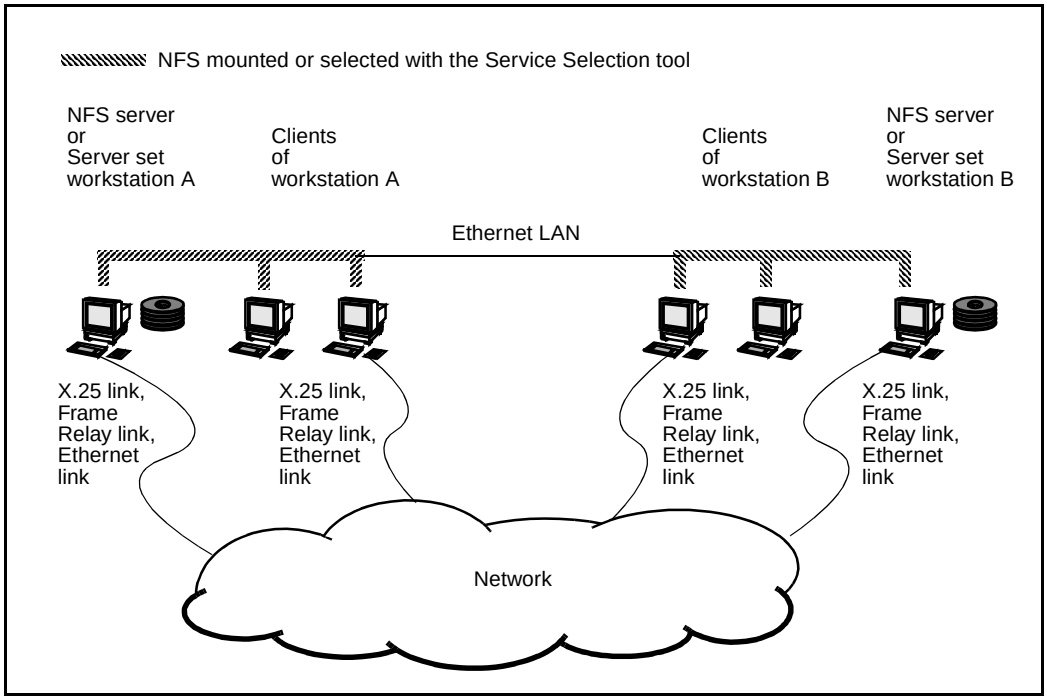

Minimizing response time

Clients to accessing software over the LAN or over X.25 links, IP over X.25 links, or frame relay links can impact the response time for client servers. Minimize this impact by ensuring the following:

- the client workstations have adequate memory
- the LAN or links operate at speeds and traffic levels that allow for rapid data transfer between the client and server workstations

Types of network access

The type of network access for a Preside Multiservice Data Manager (MDM) workstation depends on the composition of the network: DPNs only, DPNs and Passports, or Passports only.

In a DPN-only network, access is provided by an X.25 link between the workstation and a DPN switch that is running a management Operations Agent (OA). The X.25 link connects to a serial port on the workstation or to a port on a high-speed interface (HSI) card installed in the workstation. The serial port is used for data rates up to 19.2 kbits/s and the HSI port is used for data rates above 19.2 kbits/s. On UltraSPARC workstations, the serial port has an extended bandwidth that can be used for data rates up to 64 kbit/s, unless limited by cabling length and construction.

In a network that contains DPN and Passport switches, network access for managing the DPN switches is provided by an X.25 link, as described for DPN-only networks. Network access for managing the Passport switches is provided by an IP connection running on an X.25 link. The DPN switch acts as a gateway to the Passport switches and provides IP connectivity to the Passport switches.

In a Passport-only network, access is provided by one of three methods:

- an IP over frame relay link connected to a port on an HSI card installed in the workstation.
- an IP over Ethernet link connected to the Ethernet port on the workstation. To use the Ethernet link, the Passport switches must be configured as an inter-LAN switching (ILS) network. In an ILS configuration, network management traffic is treated as normal data traffic.
- an IP over ATM link connected to a port on an ATM NIC installed in the workstation. In this approach, the IP traffic is encapsulated on the Passport switch using ATM multi-protocol encapsulation and a permanent virtual circuit is constructed from the workstation to the Passport switch.

Chapter 7
Choosing a workstation to run MDM

This section contains the following guidelines for selecting the hardware to run the Preside Multiservice Data Manager (MDM) software on a workstation in your network:

Prerequisites for choosing a workstation

Before you choose workstation hardware, you must complete the following:

- use the information in "Choosing a configuration for MDM" (page 77) to determine your Preside Multiservice Data Manager (MDM) requirements including:
  - — the number of average modules in your network, and which ones will be managed by this workstation
  - — the number of workstation users, the MDM tools they will use, and the physical location of the users
  - — the projected growth rate for the network
  - — the organization model for the network (centralized, regionalized, decentralized, and so on)
  - — X.25, Internet protocol (IP) over X.25 and frame relay link availability

- Use the information in "Choosing a configuration for MDM" (page 77) to determine the configuration for the workstation:
  - stand-alone
  - stand-alone CPU server
  - client set or server set
  - NFS server
  - NFS client
  - a combination of these

You may need to modify your configuration after consulting the tables in this section.

Minimum workstation configuration needed to run MDM

The minimum hardware configuration to run Preside Multiservice Data Manager (MDM) is a stand-alone workstation with the following specifications:

- CPU performance: 400 MHz
- RAM: 512 MB
- number of disks and capacity: one disk or two disks with a total minimum capacity of 9 GB
- monitor: 21-inch color

This minimum configuration is suitable for a stand-alone workstation as follows:

- only one operator uses the workstation for managing a network containing up to 100 modules (no X-terminals)
- the workstation runs MDM and no other management software, including Management Data Provider (MDP)

  Note: This configuration is the minimum configuration needed to run MDM. When purchasing workstations to run MDM, contact your Nortel Networks technical representative to determine the workstation hardware requirements.

Minimum workstation configuration needed to run MDM with Oracle database

The minimum hardware configuration to run the Oracle database is a stand-alone workstation with the following specifications:

- 2 CPU performance: 400 MHz
- RAM: 1 GB
- number of disks and capacity: one disk or two disk with a total minimum capacity of 18 GB
- monitor: 21-inch color

Manufacturers of platforms to run MDM

The Preside Multiservice Data Manager (MDM) software is designed to run on any platform that meets the following requirements:

- capable of running Solaris as its operating system
- along with its peripheral devices must be certified as SPARC compliant by SPARC International

  For a list of SPARC-compliant workstations, which includes workstations offered by Sun Microsystems and workstations offered by other vendors, contact SPARC International Inc., or access the web site: www.sparc.com/certified.shtml.

- must have the CPU power, the amount of RAM, swap space, disk capacity, and a network communication link bandwidth sufficient for the network to manage. For the engineering information required, contact your Nortel Networks technical representative.

  Note: For optimal results, display the graphical user interface (GUI) on a Sun SPARC workstation.

Obtaining information to select a workstation

For the information to select a workstation configuration that supports anything other than the minimum configuration described in “Minimum workstation configuration needed to run MDM” (page 96), contact your Nortel Networks technical representative.

Chapter 8
Monitoring MDM

This section describes how to monitor Preside Multiservice Data Manager (MDM) workstations to ensure they are running at peak efficiency. This section contains the following information:

Before using this section, determine the following:

- the size of the network
- the number of users
- the required speed of the X.25 link
- the configuration you are going to use

Initial planning must be complete and the hardware must be up and running.

System response time

CPU, memory, and I/O throughput are crucial to ensure quick response time and high Preside Multiservice Data Manager (MDM) throughput. Resolving delay problems can require the following:

- scheduling processing to off-peak hours
- moving some processing to another workstation
- increasing workstation memory and disk sizes
- purchasing a workstation with a faster CPU

Workstation delay is affected by network engineering. If trunks leading to the node supporting a workstation are congested, performance of MDM running on the workstation will be unacceptable. Service selected MDM workstations suffer if the LAN is congested. To meet the engineering objectives, you must consider the behavior of the workstation and the external networks to which it is attached.

Managing workstation resources

One of the most important issues in having a well-engineered Preside Multiservice Data Manager (MDM) environment is the management and administration of system resources. Without this, system resources become overworked and performance decreases. This section describes areas of concern, such as disk, memory, and CPU, and contain suggestions for using UNIX performance monitoring tools.

Knowing how your system runs helps you to determine the workloads that are right for your system. Using some built-in UNIX functions lets you gather enough information to determine what is normal for the system under a variety of conditions. This helps you to spot abnormal behaviors in your system, and helps you to determine what requires adjustment. You can investigate any anomaly, and continue monitoring the system to see if the condition persists.

The configuration file lets you customize the thresholds at which workstation surveillance generates alarms. The parameters in the configuration file are set up to monitor CPU, memory, and disk space. The configuration file also

checks the status of local ports and the ability to reach other IP devices using the ping command. For more information on the threshold configuration file, see 241-6001-303 *Preside MDM Administrator Guide*.

Each network can require different configurations, which depend on the network traffic being monitored and the tools that are running.

Disk management

One of the easiest areas of the system to manage is the disk space. A disk that meets the current requirements can quickly become too small. Give consideration to whether or not the workstation will be used as a server or stand-alone and other software it will store.

Files that accumulate in the /tmp directory, and various log files, are described in 241-6001-303 *Preside MDM Administrator Guide*.

When the system is installed, use the UNIX df command to show actual disk utilization.

Memory management

It is difficult to determine how much real or virtual memory is required due to the changing mix of applications running on Preside Multiservice Data Manager (MDM), and particularly on servers. The most cost-effective strategy is to monitor workstation performance and increase swap space file sizes. Employ this strategy until paging occurs or CPU utilization exceeds the thresholds shown in the table “Tools for monitoring system performance” (page 104).

When paging occurs, CPU utilization levels quickly increase. Paging can be reduced by increasing the number of workstation disk controllers, assuming that three or more disks are connected to one controller. Alternatively, you can increase the size of the real memory configuration.

Management of the CPU

Management of CPU capacity is achieved by the choice of Preside Multiservice Data Manager (MDM) applications configured on the workstation. If workstation performance is CPU limited, you can remove applications or upgrade the workstation to a more powerful workstation. CPU

utilization can be reduced by planning. For example, do not build a network model by running the make configuration data file (MCDF) tool during a high traffic period.

Reducing MDM resource utilization

This section lists some of the applications, and describes how to reduce their impact on resources:

- run Network Viewer (NV), component status display (CSD) at level 2 when possible

  NV can consume Preside Multiservice Data Manager (MDM) system resources. One way of reducing the impact is to run NV at level 2 rather than level 3, whenever possible. For more information on the Multi-nodal Name server, see 241-6001-310 *Preside MDM Server Reference Guide*.

- limit the number of modules displayed by the Network Viewer to 100 modules

  For performance reasons, limit the number of modules displayed (number of icons displayed) to less than 100 modules.

- provision without using Expand All

  Avoid using the Expand All option in component provisioning, especially with large modules. There is no way to stop the lengthy process once started, and it takes a large amount of memory, which is not released until the session is ended.

- maximize system swap space

  To reduce swap time and paging errors, make available as much swap space as your workstation will allow. Use the *swap* command with the -a option without having to shut down or reboot the workstation.

- modify default service selection settings

  The Service Selection tool allows an administrator to set up the default server set workstations that a client set workstation uses to obtain support for various service areas. Move the default settings to balance link and processing loads, and to reduce the workload on the workstation.

Monitoring workstation performance

Monitoring the system means collecting normal and abnormal statistics about the system. Carry out system monitoring over a specified time period, or during intervals over a greater period.

There are several ways to monitor the system. UNIX has several commands to determine system performance of the workstation. These commands include the following:

- ps (displays process status)
- vmstat (provides a virtual memory statistics report)
- iostat (checks on I/O requests)
- netstat (provides the status on networking)
- nfsstat (provides detailed data transfer information)

Each command contains several options to get the relevant information. These can also be incorporated in shell script files or in C programs to perform automated reporting. A script incorporating vmstat is included with the Preside Multiservice Data Manager (MDM) software. See “Sample workstation performance monitoring session” (page 106).

Depending on which part of the system is being monitored, a different time period can be used, as is shown in the table “Tools for monitoring system performance” (page 104). Cover all situations with high and low system load, including high-traffic and low-traffic periods, through the use of the various MDM tools. Investigate these scenarios to determine a sound measurement for the system being engineered.

Table 7
Tools for monitoring system performance

Monitoring tool	Condition to look for...	Recommended sample time
vmstat	Swap (Free) memory utilization > 90% CPU user utilization > 90% Idle CPU utilization < 10% High paging rate > 1/second	several hours
iostat	High disk utilization (> 75%) Uneven disk loads	approximately 2 weeks
netstat	High max. CPU utilization High mbuf utilization	several hours to days
nfsstat	Client bottlenecks High number of time-outs badxid vs. time-outs	several times over monitor period

MDM system performance does not need to be monitored continuously. In a fast-growing network or with initial network startup, monitoring can be performed on a monthly basis. As the network matures, you can monitor the system on a quarterly basis.

Sudden peaks in workstation or network traffic can have a major effect on the system performance in the short term. Incidents such as these can usually be ignored unless they become frequent or a pattern develops. Avoid making adjustments based on what can resemble abnormalities. This can cause the statistics to become skewed and unreliable.

Another useful tool for monitoring workstation performance is SymBEL release 3.0, also known as SE3.0. This tool is built on SymBEL, which is an interpreted language that provides an extensive toolkit for building performance monitoring tools. The tool provides a default set of scripts and rules for monitoring workstation performance. For further information about this tool, and for access to its software, see the following web site: http://www.sun.com/sun-on-net/performance/se3/.

Using the UNIX workstation monitoring utilities

The table “Tools for monitoring system performance” (page 104) lists some of the UNIX utilities you can use to determine the state of a system. This section explains how to use the utility, and how to interpret its output. Although some are specialized in what they report, there is some overlap. For this reason, use vmstat as your basic query utility, and use other specialized utilities when trouble is apparent.

Each of the UNIX utilities focuses on a particular area of the system to help you monitor the system. The problem is usually with the client, the network, or the server. This section helps you to determine which area of the system is causing the trouble, and what to do to reduce or remove the impact.

What to look for on a client

A user at a client workstation can complain that system response seems slow, or that it takes too long to read in files. This can be caused by the way the user path was set up, and solved by reformatting the client path variable. Examine the local disk first, followed by the critical directories on the remote host, and the remaining host files.

Another problem can be the caches on the client machine. Use the UNIX tools netstat, nfsstat, or vmstat to determine if there is a memory problem if the following occurs:

- running mounted executables is slow
- the number of biods:server network buffer overflows. The output from nfsstat is usually in a percentage. The optimal ratio is 4 biods.
- badxid & time-outs from nfsstat -rc; requiring an increase of timeouts on the client

Determining network problems

If the server is not responding, execute the ping command. If a response is received even though the server is not responding to requests from applications, the server can be oversaturated, in a panicky state, or the network is slow. If the server is in these states, or if the network is slow, try to find the server bottleneck to determine the problem. If no response is received from the *ping command*, the server or the network to the server is down.

To help prevent server problems, keep clients on the same subnet as the server.

Invoke the netstat -i command to determine the number of output packets and the number of collisions on your subnet. Then use the following formula to calculate the collision rate for server communications:

high collision rate (if (high collision rate = #collisions/#output packets) > 0.30)

A high collision rate based on the formula is one that exceeds 30%

Checking the server resources

If server performance is repeatedly slow, run the netstat, nfsstat or vmstat tools and look for the

- look for a directory name cache (vmstat -s) hit rate (<85%) and increase the MAXUSERS variable
- page cache - this requires large samples of several hours to determine if there is a high CPU system time, CPU user time, and high paging rate
- if the ionode cache hit rate = (cache hits/cache hits + cache misses) is < 85, try increasing MAXUSERS. A tuned directory name cache usually means this cache is also tuned.
- look for disk over 75% I/O hot spots: iostat -D -l 10 60; look for disk over 75%; load balance disks. The utilization field shows 75% busy.
- use the netstat -m command to look for high Max CPU utilization and high mbuf utilization. Netstat indicates queuing of requests.

Sample workstation performance monitoring session

The example in “Output from the workstation monitor (wsmon) program” (page 108) shows the output from the wsmon program written to monitor workstation performance. This program and its accompanying startup script (/opt/MagellanNMS/bin/wsmon) are included with the Preside Multiservice Data Manager (MDM) software and are documented in “System resource utilization programs” (page 113).

The script generates a report using the output from vmstat. Since the output from the utility tends to produce considerable output, only portions are extracted and presented in the report. The program also uses ps to determine the number of users and processes immediately prior to the time vmstat is run.

The fields that are sent to the report include memory, paging faults, and CPU usage. These fields contain useful information about server activities, and can point to potential bottlenecks.

The table “Legend to output from the wsmonrep (workstation monitor) script” (page 107) gives a brief explanation each field. The vmstat fields that have been omitted from the report indicate the following:

- process running
- processes bound
- processes swapped
- page reclaims
- predicted memory shortfall due to paging
- clocking algorithm scan rate
- disk operations per second

Table 8
Legend to output from the wsmonrep (workstation monitor) script

Heading	Meaning
KB Free Memory	The number of pages (or kbytes) marked as currently free
Paging: In	Pages read from disk on page fault
Paging: Out	Pages written to disk to free core memory
Paging: Freed	Pages which have been marked released by their parent process
CPU Usage: User	Percentage of CPU in use by user applications
CPU Usage: System	Percentage of CPU that is CPU bound by system activities
CPU Usage: Idle	Percentage of the CPU that is not in use by either user or system

A properly running workstation is characterized by the following:

- free memory >20% of total swap space
- minimal paging
- idle CPU usage of 5–20% as measured over a 15-minute period of maximum workstation usage

Running the workstation monitor (wsmon) script

1 Change to the directory that contains the wsmon program.

```
cd /opt/MagellanNMS/bin
```

The report generated by the workstation monitor will be written to this directory.

2 Enter the following command:

```
wsmon <probe_interval> <iterations>
```

where:
`<probe_interval>` is the time between each probe in seconds. Typically, you want to take a snapshot of the workstation performance every 10 to 15 minutes (600 to 900 seconds) during a period of peak usage.

If you notice that a bottleneck occurs at a specific time, you can increase the number of probes during this period by decreasing the probe interval to, for example, every minute (60 seconds).

`<iterations>` is the number of times you want to probe the system. The workstation monitor automatically terminates and generates the report when the number of iterations is reached.

Statistics from the wsmon program are displayed on the screen.

Note: If you want to terminate execution of the script before the number of iterations is reached, press Ctrl-C.

Output from the workstation monitor (wsmon) program

The following example shows some sample output from the wsmon program.

```
WSMON - 3.0 now running...process number 1131. Warning messages will be
displayed to this xterm. To store the output to a file run this
command with a shell redirection or through the tee filter.
To stop program, type ^C or fill process 1131.

1 Workstation Monitor - 3.0                  Mon Feb 20 11:36:06 EST 1995
```

```
  NOTE: Delay time between system queries has been set to 2 seconds.

=============================================================================
|Number| Processes | Memory |      Paging        |      CPU Usage      |
| Users|  (Total)  |  Free  | In    Out   Freed  |  User  System  Idle |
=============================================================================
  7          104     1408    13    9        14     5      6      89
  7          103     1408    0     0        0      30     69     1
  7          103     920     0     0        0      38     62     0
Beginning cleanup routine...
   Killing process number 1131
Killed
```

Output from the wsmonrep program can indicate general workstation problems, and can be used as the first analysis tool. The first line in the following output from the script indicates excessive paging. This can be a result of not enough memory or I/O congestion, and can be corrected by adding additional memory or disk controllers. The second line (from a different workstation and a different execution of the script) indicates that the machine is CPU bound. In this case, the problem is corrected by upgrading to a SPARCstation10/40.

```
=============================================================================
|Number| Number    | KB Free |      Paging        |      CPU Usage      |
| Users| Processes | Memory  | In    Out   Freed  |  User  System  Idle |
=============================================================================
  7        115       10604     324     0     0       24      52     25
  10       122        2400       0     0     0       58      42      0
```

Correcting problems

After the engineered system has stabilized and monitoring has been applied for a specified period of time, analyze the system and make adjustments to optimize performance. This includes the following:

- changing swap space size
- disk utilization
- network alarm rates
- network probing

This can also include adding drives, memory, or upgrading processors if the system has been grossly under-engineered.

Depending on the system usage and load, the measurements that you collected and analyzed can point to problem areas, or can illustrate system abnormalities. This is the main reason that a large sample must be extracted from the system when the system is most likely to be in use.

Even if all seems well after initial system monitoring has taken place, continue the monitoring stage at a reduced rate until it indicates that changes are required to the system. Then increase monitoring.

Most problem areas become clear when the information is reviewed. Typical problem areas are CPU, memory, disk size, and X.25 bandwidth.

CPU

Typically, when a system is CPU bound, it is time to upgrade. This can include acquiring additional workstations, or upgrading the CPU where possible. Alternatively, functions (typically Fault or Configuration) can be removed from the workstation that is CPU bound, and redistributed to other workstations in the network.

Memory expansion

A high number of paging faults being generated indicates that the system is memory bound. This can be caused by the lack of real, or virtual (swap space) memory. A monitoring tool, such as the workstation monitor (wsmon) program, a UNIX utility, or something similar, can be used to help determine whether this condition exists on your server.

Real memory (RAM)

Additional memory can be installed in various increments according to the model of workstation being upgraded. Contact the manufacturer for information about the maximum memory allowed.

Virtual memory or swap space

If the system requires an increase in virtual memory, the only prerequisite to increasing the swap space is disk availability on the local workstation. Swap space cannot be obtained from another workstation.

Swap space can be allocated from a file or from a dedicated disk partition. Initial installations of Preside Multiservice Data Manager (MDM) workstations typically have a dedicated disk partition for swap space. Using a dedicated partition improves performance improvement. The dedicated partition eliminates referring to file system directories when searching for the location of swap files on disk. If the initial partition size is too small, swap space can be increased without having to reboot the system.

To determine how much swap space to add, it is recommended that you allocate three to four times of random access memory (RAM). For larger systems that have more than 2 Gbyte RAM, allocate 1 to 1.5 times of RAM.

Installing additional swap space

1 Log in as root.

2 Create a file in the disk partition in which the swap space file will be located:

`mkfile <swapsize> /<filename>/swap`

where:
`<swapsize>` is the amount of swap space, in megabytes, that you wish to add

`<filename>` is the name of the file where the swap space is added to

Example:

```
mkfile 5 /user1/swap
```

3 Edit the /etc/vfstab file and add the following entry to ensure that the swap space is permanent. Add the following (on a single line) to the vstab file:

device to mount: /<pathname>/<swapfile_name>

device to fsk: -

mount point: -

FS type: swap

fsck pass: -

mount at boot: no

mount options: -

where:
`<pathnames>` is the path name where the device is mounted

`<swapfile_name>` is the name of the swap file

Example:

```
/localdisk/mydirectory/swapfile - - swap - no -
```

4 Activate the swap file without rebooting.

```
swap -a
```

5 Verify the addition of the swap space:

```
swap -l
```

Initially, allocate swap space on a workstation that is roughly three times that of the real memory installed.

Shared memory

This is an area of real memory that is made available to multiple programs, and for that reason, it is usually resident. The use of shared memory saves memory space, rather than having to have a section allocated solely for their need. The Network Model is placed in this area of memory, since many applications require access to model information.

To some extent, the amount of shared memory is proportional to the size of the network. Each access module (AM) or resource module (RM) or Passport requires approximately 25 kbytes of shared memory. A graphical tool is available to help visualize the shared memory. If invoked before loading the network model, you can see how much shared memory is required while the model is loading. The Shared Memory Utilization tool is found in the System -> Utilities menu.

Number of disks

The number of disks is governed by the availability from the manufacturer. It is beneficial to acquire a disk that is larger than you require. The amount of disk space required is determined by

- adding the amount required by the software
- adding the amount needed by each account that will reside on the disk

One or more additional 18 Gbyte or 36 Gbyte disk can be used if the network file system (NFS) is configured on the same server. Assign only one separate partition for each additional NFS.

Appendix A
System resource utilization programs

This appendix contains listings for the wsmon and wsmonrep scripts that are included with the Preside Multiservice Data Manager (MDM) software. You can use these scripts to aid the workstation monitoring function. Topics in this appendix are as follows:

The wsmonrep script is called by the wsmon script. For information on how to use the workstation monitoring scripts see, "Sample workstation performance monitoring session" (page 106).

Workstation monitor startup script

The following macro is included with the Preside Multiservice Data Manager (MDM) software as /opt/MagellanNMS/bin/wsmon.

```
#!/bin/sh

#============================================================
# /opt/MagellanNMS/bin
#           - Startup script for performance monitor.
#             This script contains the startup and
#             cleanup functions for monitoring a
#             workstation.

#

# This function cleans up before exiting

#
```

```
shutdown () {

 list=‘ps -j | grep $pid | grep -v grep | sort -r | awk ’{print $2}’‘
       for child in $list
   do
      n=‘ps $child | wc -l‘
      if [ $n -gt 1 ]

      then
          echo “Beginning cleanup routine...”
          echo “   Killing process number “ $child
          kill -9 $child 2>&1 >/dev/null
       fi
          echo “The following temporary files will now be removed...”
          ls /tmp/wsmon*
          rm -f /tmp/wsmon*
   done
   exit

}

#

# This function prints out a usage blurb

#

#* Make sure an argument for delay (’sleep’) time has be supplied.

useright () {

  echo “Usage: $0 time <iterations>”
  echo “  Where ’time’ is the interval time in seconds,”
  echo “  and ’iterations’ is the number of times to query”
  echo “  the system.”
  echo “  ie.: ’$0 30 10’ would set the interval time to”
  echo “       30 seconds between system stat queries,”
  echo “       for 10 queries.”
  echo “       (Add ’&’ to run as a background job).”
  exit 2
}

 #

# ensure that the user has input 1 parameters

#

if [ $# -lt 1 -o $# -gt 2 ]
```

```
then
   useright
   exit 1
fi

delay=$1
iterate=""
iterate=$2

set -f
set noglob

trap 'shutdown' 0 1 2 3 15
pid=$$
v=2.0

#

# start up a dialog with rncs.

#

clear

echo "WSMMON - $v now running...process number $pid.  Warning messages
 will be"
echo "displayed to this xterm as well as to file.  To stop program, type
^C"
echo "or kill process $pid.    "
echo "                         "

 #

# pipe the raw data from vmstat to 'wsmonrep'

#

#if [ $# = 1 ]
#then
  #vmstat $delay | /iws/usr/bin/wsmonrep &
#else
echo $delay > /tmp/wsmondelay
echo $iterate > /tmp/wsmoniterate
  vmstat $delay $iterate | /iws/usr/bin/wsmonrep &
  #shutdown

#fi
```

```
 while  [ 1 ]
do
  :
done
exit
```

Workstation monitor report generator script

This script uses the output from vmstat to provide system resource utilization data. See “Sample workstation performance monitoring session” (page 106) for more information.

This script is included with the Preside Multiservice Data Manager (MDM) software as /iws/usr/bin/wsmonrep. It is called by the wsmon script.

```
#!/bin/csh -f
#*********************************************************************#
#
# wsmonrep
# General Description:
# This shell script will be used to generate reports on system resource
# utilization on the workstation on which it is run.  It uses the 'vmstat'
# utility for the raw system data, and reformats the output to a report
# format, stripping only the required data.
# The script is launched from a unix xterm, or window, and can be run
# either in the foreground or background.  It requires a single argument,
# which will be used to determine the time (in seconds) between system
# queries. This allows the user to collect as much information as required.
#
#
set activity = $<
set activity = ( $activity )
#
# Initialize the variable for the program.
#
set whatdate = ‘date‘
set pid = $$
set delay = ‘cat /tmp/wsmdelay‘
set interate = ‘cat /tmp/wsmiterate‘
set destaud = $whatdate[2]$whatdate[6].wsmrpt$pid
```

```
set pageno = 1
set maxlines = 55
set maxpages = 100
set linecount = 56
set version = 2.0
set header1 = "1 Workstation Monitor - $version           $whatdate"
set header2 = " "
set header3 = "  NOTE: Delay time between system queries has been set to
     $delay seconds."
set header4 = " "
set header5 = "  |Number| Processes |  Memory  |   Paging  |   CPU Usage   |"
set header6 = "  | Users| (Total) | Free| In  Out  Freed  |
     User   System   Idle |"
set header7 = "
================================================================================
"

set cpuhigh = 0

# Mainline begins here...
# Make sure the destination report exists; if it does not, create it.

while ( $#activity != 0 )

   touch $destaud

# Check to see if a page break is required.

   if ($linecount > $maxlines) then
      echo "$header1" >> $destaud
      echo "$header2" >> $destaud
      echo "$header3" >> $destaud
      echo "$header4" >> $destaud
      echo "$header7" >> $destaud
      echo "$header5" >> $destaud
      echo "$header6" >> $destaud
      echo "$header7" >> $destaud
      set linecount = 8
   endif

# Get the number of users and processes currently running.
   set users = ‘who|wc -l‘
   set procs = ‘ps -aux|wc -l‘
   @ procs--

# Format line for output to report.
```

```
   set activity = ‘echo $activity | egrep -v '(procs|avm)'‘
   if ( $#activity != 0 ) then

# Check to make sure that the raw data is fully readable before trying

# to format.

     switch ( $#activity )

       case 21:
           set activity = ‘echo $activity | awk '{ print $2 "     " $3 "
         " $5 "     " $8 "     " $9 "     " $10 "     " $19 "     " $20 "
         " $21 }'‘

         set bp = $activity[1]
         set sp = $activity[2]
         set fm = $activity[3]
         set pi = $activity[4]
         set po = $activity[5]
         set pf = $activity[6]
         set cu = $activity[7]
         set cs = $activity[8]
         set ci = $activity[9]
       echo "  $users $procs  $fm  $pi  $po  $pf  $cu  $cs  $ci" >> $destaud
         breaksw

       case 22:

         set activity = ‘echo $activity | awk '{ print $2 "     " $3 "
         " $5 "     " $8 "     " $9 "     " $10 "     " $20 "     " $21 "
         " $22 }'‘

         set bp = $activity[1]
         set sp = $activity[2]
         set fm = $activity[3]
         set pi = $activity[4]
         set po = $activity[5]
         set pf = $activity[6]
         set cu = $activity[7]
         set cs = $activity[8]
         set ci = $activity[9]
       echo " $users $procs  $fm  $pi  $po  $pf  $cu  $cs  $ci" >> $destaud
         breaksw
        default:

   echo "   >>>   Problem reading raw data...data for this probe ignored.
<<<" | tee -a $destaud
```

```
         set bp = 0
         set sp = 0
         set fm = 101
         set pi = 0
         set po = 0
         set pf = 0
         set cu = 0
         set cs = 0
         set ci = 0
         breaksw

     endsw

# Lets do some monitoring for troubled areas and report them if found

# Check for CPU utilization above 95%.  If it continues for 5 consecutive
# status cycles, report it.

      if ($cu > 98) then
       @ cpuhigh++
       if ($cpuhigh > 4) then
         echo "       *** WARNING...CPU user processes were above 98% for
          last $cpuhigh cycles. ***" | tee -a $destaud
          echo "               *** Please investigate for possible causes.
          ***" | tee -a $destaud
          echo ""
         echo "       *** The following is a list of high runner processes
          currently running.  ***"
          echo ""
          ps -auxr
          set cpuhigh = 0
          @ linecount = $linecount + 2
       endif
     endif

# See if the Free Memory available has dropped below 100Kb, and if so,

# report it

     if ($fm < 100) then
       @ lowmem++
       if ($lowmem > 4) then
        echo "      *** WARNING...Free Memory available has dropped below
         100 kb over the ***" | tee -a $destaud
        echo "         *** last $lowmem status probs...Please investigate
         possible causes. ***" | tee -a $destaud
```

```
        set lowmem = 0
        @ linecount = $linecount + 2
      endif
    endif

#* See if the Processes are being blocked from resources.

    if ($bp > 0) then
      @ blocked++
      if ($blocked > 4) then
        echo "       *** WARNING...$bp processes have been blocked for
        resources  over the ***" | tee -a $destaud
        echo "          *** last $blocked status probs...May indicate
         memory problems. ***" | tee -a $destaud
        set  blocked = 0
        @ linecount = $linecount + 2
      endif
    endif

#* See if the Processes are being swapped out even though runnable.

    if ($sp > 0) then
      @ swapped++
      if ($swapped > 4) then
        echo "       *** WARNING...$sp processes have been blocked for
        resources  over the ***" | tee -a $destaud
        echo "          *** last $swapped status probs...May indicate
        memory problems. ***" | tee -a $destaud
        set  swapped = 0
        @ linecount = $linecount + 2
      endif
    endif

#* Check to see if we are chewing up disk space ourselves!

    if ($pageno > $maxpages) then
     echo "       **** WARNING...The report has grown to $pageno pages.
      ***" | tee -a destaud
    endif
    @ linecount++

#* Check to see if a page break is required.

    if ($linecount > $maxlines) then
      echo " " >> $destaud
      echo " " >> $destaud
```

```
      echo " " >> $destaud
      echo "                                        page $pageno" >> $destaud
      @ pageno++
      echo "$header1" >> $destaud
      echo "$header2" >> $destaud
      echo "$header3" >> $destaud
      echo "$header4" >> $destaud
      echo "$header7" >> $destaud
      echo "$header5" >> $destaud
      echo "$header6" >> $destaud
      echo "$header7" >> $destaud
      set linecount = 8
    endif
  endif
  set activity = $<
  set activity = ( $activity )
end
```

Appendix B
MDM Engineering for SNMP Devices

This appendix contains performance and configuration information for devices in the network that can be managed with simple network management protocol (SNMP). Examples of SNMP devices to which the information in this appendix applies are Baystack 450 and Passport 8600 devices.

Use the information in the following sections to define your SNMP requirements for maximum performance.

Factors that affect surveillance of SNMP devices

When choosing a computing platform to run Preside Multiservice Data Manager (MDM), you need to select a platform that has sufficient computing resources to support the surveillance activities of the Preside MDM software, and any other applications that run on the same platform.

The computing resources required to run the MDM fault management tools For SNMP devices depend on the following factors:

- the number of polled parameters
- the time interval between two processing cycles. In larger networks, the defaults (typically set between 5 to 7.5 minutes) need to be increased to allow for collection of the polled information before another polling cycle starts.
- the trap rate
- the trap handling:
  - — if the traps cause the software to generate polling requests
  - — if the polling requests that are generated require further processing, then more computing resources are required
- the way in which the configuration files for the device are written. This affects the generation of the proxy alarms that are sent to MDM, and the amount component/subcomponent and active alarm information. This process can involve many computing operations. There may also be no processing, however, minimal processing is done before events are sent to the rest of MDM.
- the number and type of applications running on the MDM server
- the MDM server configuration

Assumptions used in this appendix

This section contains the assumptions we made while creating and testing the engineering information contained in this appendix. See the following sections for information about the assumptions:

- “General assumptions” (page 125)
- “Capacity guidelines” (page 125)
- “Workstation hardware to run Preside MDM” (page 125)

General assumptions

The following assumptions were made when we estimated and calculated the number of SNMP devices on supported by an MDM workstation:

- The MDM server is configured to support SNMP devices only. No Passport or DPN switches are configured.
- A standard or general configuration of Preside MDM is used (no MDMWeb, no exporting of surveillance information to external Operations Support Systems)
- The devices used default device integration cartridge configuration (standard polling rates, variables polled, and so on)
- SNMP v1 is used for base measurements and calculations.
- A small number of operators (less than 5) are using the system. If this number is insufficient you need to consider the number of operators in your calculations.
- There is no load sharing for single device type. That is, a single device is not distributed over many DCD processes.
- Only Preside MDM and no other management software including Management Data Provider (MDP) is running on the platform.

Capacity guidelines

We used the following capacity guidelines when calculating the workstation requirements to support fault surveillance of SNMP devices in the network.

- Maximum CPU for system occupancy is less than 65%, with DCD servers occupying less than 50% total CPU
- Maximum bandwidth consumption for polling variables and traps per workstation is 40% of the link (typically 10/100 mbs) used for device management to the network.
- Maximum number of devices per DCD is 1000.

Workstation hardware to run Preside MDM

We assumed the following common hardware configurations of a standalone workstation in our calculations:

- Small System - Single CPU, ~400 MHz

- Medium System - Dual CPU ~ 400MHz
- Large System Multiple CPU (4 or more)

Determining the number of average SNMP devices

Use the information in this section to calculate the number of average of SNMP devices in your network.

Use the following formula to determine the number of CPUs required in a single workstation to support SNMP. For some samples of the results we obtained using this formula, see the table, "Sample SNMP device information" (page 127).

```
no CPUs =
(no of devices*no of components per device)/
(24000*DPF*TRF)
```

where:

`no of devices` is the number of devices in the network

`no of components per device` is the average number of components per device

`24000` is the maximum number of components per CPU. This figure has been determined based on lab testing. For information about how we arrived at this figure, see "Calculating the number of components per CPU" (page 127).

`DPF` is the DCD process factor. For instructions to calculate the DPF, see "Calculating the DCD process factor" (page 127)

`TRF` is the trap rate factor. For instructions to calculate the TRF, see "Calculating the trap rate factor" (page 128)

Example:

Calculate the number of CPUs required in a workstation for a network containing 200 Passport 8600 devices, each with 60 components:

```
(200 *60)/(24000*1.0*1.0)= 0.5 CPUs
```

Therefore a single CPU system has the ability to support the 200 Passport 8600s.

Table 9
Sample SNMP device information

Device type	Average number of components	Average polling interval (seconds)	Average number of variables polled per component	Estimated number of Passport module equivalents
PP8600	60	300	4	0.125
PP4400	5	20	4	0.04
Juniper M40	75	300	4	0.16
Baystack 450	35	300	4	0.08
Note: The number of components per device may vary due to device configuration. For information, see 241-6001-118 *Preside MDM SNMP Surveillance Adapter Guide*.				

Calculating the number of components per CPU

We estimated the number of components supported by a CPU as follows:

- The equivalent of an average Passport module c onsists of approximately 480 managed components (interfaces/device).
- Based on other studies we estimate that a single CPU will support approximately 50 Passport module equivalents.
- Therefore a single CPU system can support (480*50=) 24000 components.

Calculating the DCD process factor

A single DCD process can support up to 1000 devices. If you wish to support more than 1000 devices, you must configure more than one DCD.

Use the following mathematical representation to calculated the DCD process factor (DPF):

```
if # of devices < 1000
  then DPF =1
  else DPF = 1.25 (Round up (the number of devices/1000)-1)
endif
```

where:

`1.25` is the DPF coefficient derived from analysis and lab testing.

Examples:

For 1500 devices, the DPF = $1.25^{(2-1)}$ = 1.25

For 2200 devices, the DPF = 1.25(2) = 1.56

> ***Note:*** If the overall number of CPUs needed is greater than the number of DCDs, then you need to reconfigure the system to distribute device support over the same number of DCDs as the number of CPUs.

Calculating the trap rate factor

The number of traps generated by a device can greatly affect the number of devices supported. It is important to understand the trap/alarm rate for a device under normal operating conditions. SNMP traps can generate one or more alarms as well as trigger additional polling of SNMP variables form a device. If you anticipate that alarm rates will be higher than normal, you need to estimate the projected rates and include them in the trap rate factor calculations.

For trap rates less than one per device per minute, use a trap rate factor of 1.0. For all other trap rates use the following formula to estimate the trap rate factor (TRF):

```
TRF =
((# components*#traps per minute/device)/#components*
trap rate coefficient)+1
```

The stable network trap rate coefficient if 0.04 (based on lab testing with mathematical analysis)

Example:

For a trap rate of 2 traps/components/minute,the TRF is:
`((60*21)/60*04)+1 = 1.8`

> ***Note:*** This is a general guideline for trap rates in a stable network. As the network becomes unstable and trap rates increase, the factor may change exponentially. For information to engineer Preside MDM for high trap rates, contact Nortel Networks support.

Using the number of average SNMP devices in a mixed network

When estimating the workstation requirements for a Preside MDM workstation that is monitoring a mix of Passport and DPN switches and devices that use the SNMP for surveillance, we recommend that you convert the average SNMP devices into average Passport module equivalents. Here is the equivalent we recommend:

1 Passport average module = 480 SNMP components

Therefore a network with 12 Passport 8600 devices with 60 components per device is the equivalent to 12*60/480=1.5 Passport average modules.

Example:

For a mixed network with 200 Passport 8600 devices, 20 Juniper M40 devices, with 75 components per devices, and 20 Passport 16-slot switches the number of Passport average module equivalents is:

```
(200*60 +20 *75)/480 +20 = 28.25 +20 = 48.25 Passport
average module equivalents
```

This configuration approaches the maximum capacity of a single CPU workstation.

> ***Note:*** For large networks you also need to consider the number of CPUs, the alarrms/trap rates, the number of DCD processes, the number of network operator sessions, and the number of other applications running on the Preside MDM workstation.

To calculate the Passport module equivalent in large networks or when more than 5 simultaneous operator sessions run on the same Preside MDM workstation, contact Nortel Networks support.

Additional recommendations affecting fault management of SNMP devices

The following recommendations focus on response times, efficient CPU usage of the Preside MDM workstation and network bandwidth consumption need to support the SNMP devices:

- The SNMP version supported. We recommend that you use SNMPv2 wherever possible instead of SNMPv1 because SNMP has optimized request and variable packaging.
- Number of DCD servers per device type. We recommend that you configure 1 DCD server per device per CPU on a Preside MDM workstation. For example, on a Preside MDM workstation with 2 CPUs you can configure up to two DCDs with 1000 PP44400 devices per DCD.

Appendix C
MDP requirements

This appendix contains information on Management Data Provider (MDP) requirements. MDP collects and converts Passport accounting and statistical data from the raw data format into a usable format.

This appendix contains the following section: “MDP deployment options” (page 131)

For additional information on MDP, see 241-6001-309 *Preside MDM Management Data Provider User Guide* and 241-6001-806 *Preside MDM MDP Data Formats Reference Guide*.

MDP deployment options

MDP can be set up as a stand-alone primary server and a secondary server. There are also other server options that help to protect MDP data. This section outlines some of the options available in server technology that enhance the MDP data availability, reliability, and integrity.

Data protection by RAID

Redundant array of independent disks (RAID) is the industry standard. Sun offers two types of RAID. For MDP, use RAID 1+0 and RAID 5.

RAID 1+0

RAID 1+0 combines the data stripping and mirroring concept. RAID 1+0 is the most suitable for critical applications where data redundancy and integrity are vital.

The figure “RAID 1+0” (page 132) shows how RAID 1+0 operates. The system is divided into two parts that mirror each other horizontally. Each vertical part is subdivided into a disk pair. Disk 1 and Disk 4 are a pair, Disk 2 and 5 are a pair, and Disk 3 and Disk 6 are a pair.

If Disk 1 fails, you can still access data from Disk 4. If Disk 5 fails, you can access data from Disk 2. This type of arrangement services multiple disk failures.

Figure 26
RAID 1+0

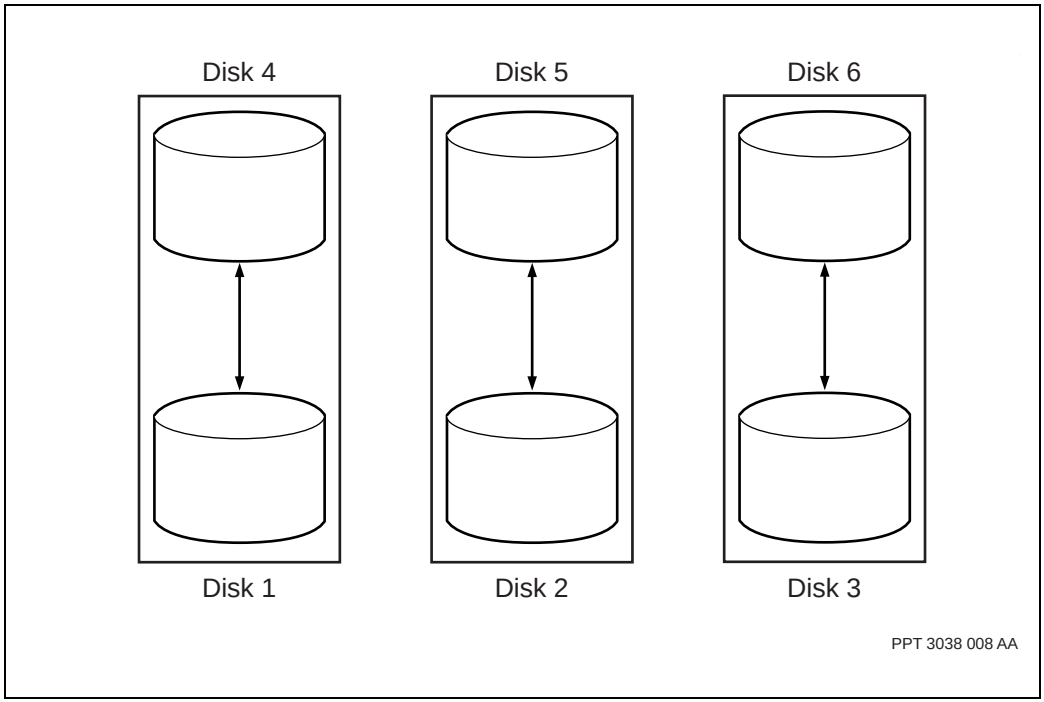

RAID 5

RAID 5 uses a parity check to provide better data integrity, similar to the frame check sequence used for the high-level data link control (HDLC) in communications framing protocols. RAID 5 spreads out the load onto multiple disks to eliminate any performance bottlenecks.

RAID 5 is sufficient for random-read applications, but it suffers performance degradation when there is a disk failure. When it needs to write data to disks, it re-computes the parity check, which slows down the network.

The figure “RAID 5” (page 133) shows how RAID 5 works.

Figure 27
RAID 5

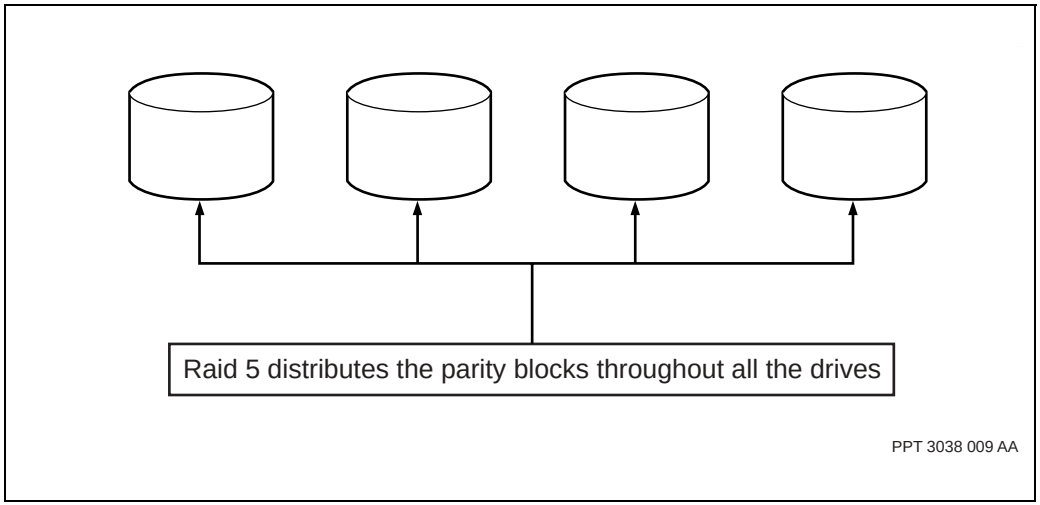

Data availability by server clustering

If an MDP process dies, server clustering lets you access MDP server data. Server clustering consists of having two separate workstations linked together, with one shadowing the other. If one workstation crashes, you can access server data through the other workstation. Without server clustering, you cannot access the server data until the workstation is restored.

The figure “High availability through Sun clustering technology” (page 134) shows how server clustering works.

Figure 28
High availability through Sun clustering technology

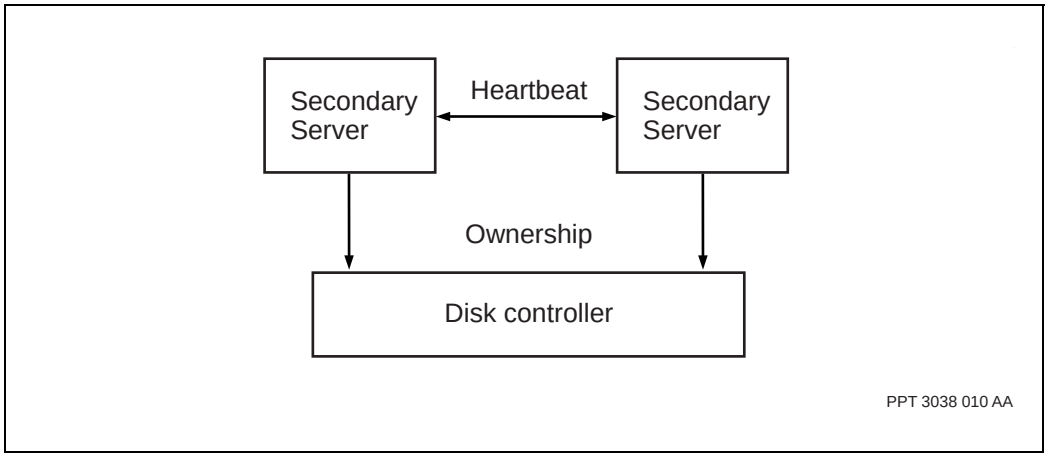

UNIX file system limitations

The original UNIX file system depends on the buffer cache to write all the data into the data block. When the data block is out of synchronization with the superblock, orphan files result. Many other problems occur as a result of the original file system.

Veritas file systems

Veritas file systems work best with high traffic and large volume transactions, where the operating system constantly opens and closes files. Veritas file systems are journalled file systems that keep track of all the transactions by using transaction logs.Vertias file systems use the transaction log to reconstruct the file system at reboot time.

TCP delayed acknowledgement and bandwith requirements

Determining the network bandwidth and disk space to use depends on the amount of data you want to collect and the frequency of the data collection.

A faster transmission link avoids transmission time-outs if there is high-volume traffic with many files to process. Transmission control protocol (TCP) retransmits data packets if the data transfer is unsuccessful, but it causes more network congestion and delayed delivery of packets to the end system.

A dedicated MDP host lets you tune the TCP delayed acknowledgement kernel parameter to improve the file transfer process. This solution only works for batch processes, such as file transfer protocol (FTP), when there is no interactive traffic. Altering the default TCP kernel parameter changes the settings suitable for other traffic types such as telnet. This requires the privileged command called ndd, where you can alter the TCP/Internet protocol (IP) driver while the system is running.

When you are running IP over frame relay on the Passport, you can use another approach. You can set the transmission priority queue on the logical channel number (LCN). This helps smooth out the default interrupt queue, which is meant for Telnet and other traffic types.

Appendix D
Upgrading your UNIX server

This appendix contains information about when to upgrade a UNIX worksation that is used as server, and how to determine which workstation to use. This appendix also contains information about how to prevent server problems in the network. The following information is contained in this section:

When to upgrade to a bigger UNIX server

Your business requirements dictate when to upgrade to a bigger server. When your system capacity is approaching the 65% mark, perform a system engineering study to determine if an upgrade is necessary. When determining whether to upgrade, consider the following factors:

- the importance of network management and its impact to your overall business objectives
- the size of your Passport network and its projected growth rate
- the number of users accessing the server
- CPU, memory, and network bandwidth requirements
- data collection and disk space requirements
- having a redundant or a stand-alone server approach

- deploying a distributed versus centralized Preside Multiservice Data Manager (MDM) server architecture to manage your Passport network
- performance versus cost trade-offs

Workstation sizing

The table “Sun server sizing” (page 139) helps you to determine which workstation to use as your server depending on the application and the number of Passports in your network. For example, you could use an Ultra 10 with 512 Mbyte RAM if you were running Management Data Provider (MDP) in a network of less than 50 Passports.

The equivalent number of Passports used by the table “Sun server sizing” (page 139) is determined by the following formula:

1.5 * (number of Passport 15000s or Passport 20000s)

+ 0.8 * (number of Passport 6400s + number of Passport 7400s)

+ 1.0 * (number of Passport 6480s + number of Passport 7480s + number of DPN access module (AM)/resource module (RM)

+ 0.2 * (number of Passport 4400s)

+0.3 * (number of Passport 4400s + number of Passport 6420s)

The data in the table “Sun server sizing” (page 139) assumes a maximum of five operators for Preside Multiservice Data Manager (MDM). Having more operators requires additional CPUs and memory for an additional Sun workstation. These are non-redundant configurations. The number in brackets refers to the number of CPUs.

Table 10
Sun server sizing

Applications	<50 Passports	100 Passports	200 Passports	400 Passports
MDM (stand-alone with user clients	1 CPU with 1 Gbyte RAM	2 CPU with 1.25 Gbyte RAM	3 CPU with 1.75 Gbyte RAM	4 or 5 CPU with 2.75 Gbyte RAM
MDM (collector only)	1 CPU with 512 Mbyte RAM	1 CPU with 768 Mbyte RAM	2 CPU with 1.25 Gbyte RAM	4 CPU with 2.25 Gbyte RAM
MDP	1 CPU with 512 Mbyte RAM	2 CPU with 1 Gbyte RAM	3 CPU with 1.5 Gbyte RAM	4 CPU with 2 Gbyte RAM
MDM with MDP	2 CPU with 1.5 Gbyte RAM	4 CPU with 2.5 Gbyte RAM		
MDM with Oracle database	2 CPU with 1.0 Gyte RAM	4 CPU with 2 Gbyte RAM	6 CPU with 3 Gbyte RAM	8 CPU with 4 Gbyte RAM

MDM RAM and network connectivity requirements are as follows:

- 256 Mbyte RAM, 64 kbit/s link, if the network contains less than 10 Passports
- 768 Mbyte RAM, 256 kib/s link, if the network contains 50 Passports
- 1.5 GB RAM, 3 x 256 kbit/s link, if the network contains 100 Passports

Note: A 256 kit/s link or higher is preferred for all software downloading to a switch.

MDM/SM requirements, without the Oracle database, can be met with a 9-Gbyte hard drive. Minimum requirements for MDM with the Oracle database is an 18 Gbyte hard drive.

In most cases, five to 10 additional X-terminal users require the deployment of an Ultra 10 management console for MDM or to upgrade the existing server with 1 Gbyte RAM.

For larger networks of 100 to 1000 or more Passports, contact your Nortel Networks representative for sizing requirements.

Using E280(R)

Use E280(R)as your server when the following applies:

- the application is CPU bound, for example, HP OpenView Network Node Manager
- the application does not require a lot of disk storage, for example, Preside Multiservice Data Manager (MDM)
- there is limited floor space (the machine is in the server room)

Using ES 3800/4800

Use ES 3800/4800 as your server when the following applies:

- you are consolidating many applications on a single server (greater than 4 CPUs)
- managing many smaller workstations is not the preferred solution in the primary operations center
- you are looking for high availability and want to minimize service disruption during upgrades

Performance tuning rules for Sun servers

This section contains rules that help you spot potential bottlenecks before they become a problem. Instead of dealing with the problem reactively, these rules help you to actively monitor the system.

Disk rule

If more than 75% of the disk is busy, and there are service times of more than 3050 ms, it is time to upgrade your server. If you also see RAM consumption increasing, you can add more RAM. Adding more RAM can reduce swap space and file system paging activity on the disk subsystem. If there is a lot of write activity, adding non-volatile RAM to the input/output subsystem or enabling fast writes to cache on a SparcStorage Array reduces the disk load.

Network rule

The network rule looks at all the Ethernets in the system at once. If the collision rate is consistently higher than 75%, you need to move the device to another LAN segment.

Swap space rule

If you run out of swap space, the system stops working properly. Use vmstat or sar to monitor the system if there is low swap space.

You can determine the amount of swap space defined in the system by using the command:

```
/usr/sbin/swap -s
```

To determine the amount of available memory installed on the system, use the following command:

```
/usr/bin/dmesg
grep avail mem | uniq
```

If the swap space is less than three times the physical memory, you can add more swap space. For larger systems that have plenty of physical memory (more than 2 Gbyte RAM), consider allocating 1 to 1.5 times the amount of physical memory. Use the following command to get more swap space:

```
mkfile <swapsize> /filename/swap; swap -a /filename/
swap
```

where:
<swapsize> is the amount of swap space, in megabytes, that you wish to add

RAM rule

The virtual memory system indicates that it needs more memory when it looks for idle pages to reclaim for other users. If there is less than 32 Mbyte of reserved memory, the system is low on memory. It is then necessary to add more memory.

To determine if there is a memory shortage, do the following:

- determine if the applications are paging excessively because of a memory shortage

- determine if the system can benefit by making more memory available for the buffering

If the system is paging out excessively, use vmstat to confirm. If vmstat does not confirm this, there is no memory shortage. Excessive paging activity is indicated in the scan-rate(sr) and page-out(po) columns, where these values are constantly non-zero.

CPU rule

Determine if you have adequate an CPU cycle by the run queue statistics reported by the vmstat utility available on UNIX. The sar -q command determines the run queue length. CPU shortage is evident by the lack of idle time, followed by the CPU load average climbing.

Adding more CPUs

Before adding more CPUs, make sure that you can partition the application to run on multiple CPUs. The application must be multi-threaded so that it can use more than one CPU at a time.

Each CPU must have an independent process to run. The optimal number of CPUs for a system is equal to the number of run jobs. For example, if the run queue length is 4 on a system with 3 CPUs, the optimal number of CPUs is 7.

DNLC rule

The directory name look up cache (DNLC) speeds up directory access by caching the file names to let you open files quicker. If the cache is too small, extra disk reads are needed and file system lookup, such as the network file system (NFS) server, response time, and throughput, decreases. If you have less than 256 Mbyte of RAM, set ncssize=5000 in your /etc/system file.

Inode rule

The UNIX file system (UFS) inode cache is related to the DNLC. UFS stores the information about a file in a local UFS file system. There are active entries and inactive entries. Active entries are for open files and must remain in memory while they file is open. Active entries also contain the locations in memory of the pages that cache file data. Inactive entries are for files that are not currently open. The current algorithm limits the number of inactive entries to the value of ufs_inode, which is the same as ncsize, by default.

TCP/IP rule

The transmission control protocol (TCP) stack is monitored using the data that is normally obtained by the netstat command or SNMP. If the collision rate is too high on the Ethernet interface, determine whether you need the latest TCP patches by dividing the number of output packets by the number of errors. If the system is consistently 75% busy because there are too many servers on the same LAN, you may have to move the server to a different LAN segment.

Index

N

P

R

S

U

V

W

X

Preside Multiservice Data Manager

Engineering and Planning

Guide

Release R14.3

Publication: 241-6001-101
Document status: Standard
Document version: 14.3RSUP
Document date: December 2003
Printed in Canada

NORTEL NETWORKS