NORTEL NETWORKS™

Preside Multiservice Data Manager

Engineering Guide

241-6001-101

Document status: Standard
Document version: 15.1RSUP
Document date: August 2004
Product release: 15.1

Contents

Engineering overview 6
MDM engineering philosophy 6
Planning for MDM 7
MDM philosophy of reuse 7
Configuring MDM 7
Monitoring MDM 8

Passport connectivity 9
Connectivity types 9
In-band connectivity 10
Out-of-band connectivity 13
Passport on-switch management protocols 14
FTP with IPSec 14

Distributed architecture 15
Distributed surveillance architecture overview 15
Surveillance servers 16
Generic DCD 17
SMDR server 18
FMDR server 18
System and application management agents 19
GMDR 19
GMDR-GMDR filtering for high cost WAN links 19
Regionalization 19
Engineering recommendation 20
Surveillance redundancy 21
Operator Client architecture 22
MDM User Administration server 23
MDM Operator Clients 23
MDM Server Workstation 23
MDM Toolset 23
Host Group Directory server rules 24
Network data access mediator 25

MDM servers 27
MDM servers to configure 28
Passport groups 28
Groups of Passports for network access 28
Groups of Passports for surveillance access 29
FMDR server redundancy for surveillance access 30
Distribution of servers in large networks 31
Guidelines for deploying servers over multiple workstations 32
NDAM server 32
Component criticality thresholds 33
Component type and regional filtering 34

Choosing a configuration for MDM 36
Types of configurations 36
Stand-alone configurations 37
Stand-alone server model configuration 38
Stand-alone CPU server model configurations 39
Practical limits to stand-alone CPU server configurations 40
Client set/server set configurations 41
Combination configurations 44
Advantages of various server configurations 44
Reduced hardware costs 44
Reduced administration expenses 45
Disadvantages of server configurations 45
Network reliability 45
Response time 45
Complexity 45
Counteracting the disadvantages of server configurations 46
Building in redundancy 46
Server configuration with backup stand-alone workstations 46
Server configuration with redundant servers 47
Minimizing response time 47

MDP deployment options 48
Data protection by RAD 48
RAID 1+0 48
RAID 5 49
UNIX file system limitations 50
Veritas file systems 50
TCP delayed acknowledgement and bandwidth requirements 50

Monitoring MDM 52
System response time 52
Managing workstation resources 53
Disk management 53

Memory management 53
Management of the CPU 54

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 7)
- Configuring MDM (page 7)
- Monitoring MDM (page 8)

MDM engineering philosophy

Deploying Preside Multiservice Data Manager (MDM) is a three-phase, circular process that consists of planning, configuring, and monitoring (see the figure MDM engineering philosophy (page 6)).

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.

MDM engineering philosophy

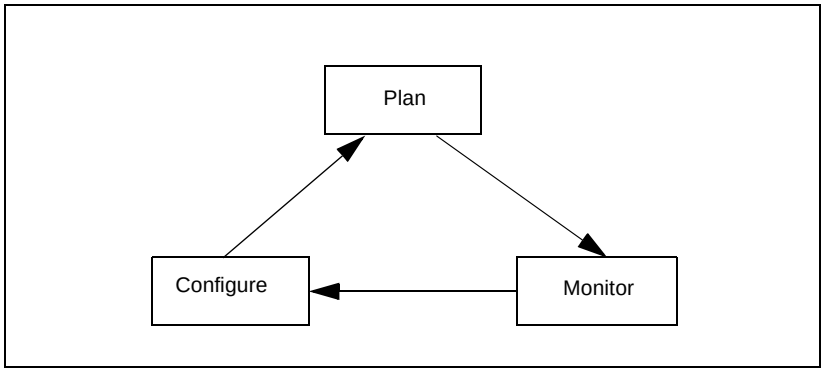

For MDM engineering to succeed, the process is a closed loop. The requirements change continuously as new technologies emerge and the network grows. You must be prepared to change the MDM configuration as the network requirements change. This section provides an overview of each of these phases, and 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 network 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, and for information on determining the specific requirements for your network, see 241-6001-102 *Preside MDM Planning Guide*.

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

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 <insert x-ref to Choosing a configuration for MDM section>.

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.

Passport connectivity

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

- Connectivity types (page 9)
- Passport on-switch management protocols (page 14)

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.

For an example of a Passport network that uses in-band and out-of-band connectivity, see the figure In-band and out-of-band connectivity example (page 10).

In-band and out-of-band connectivity example

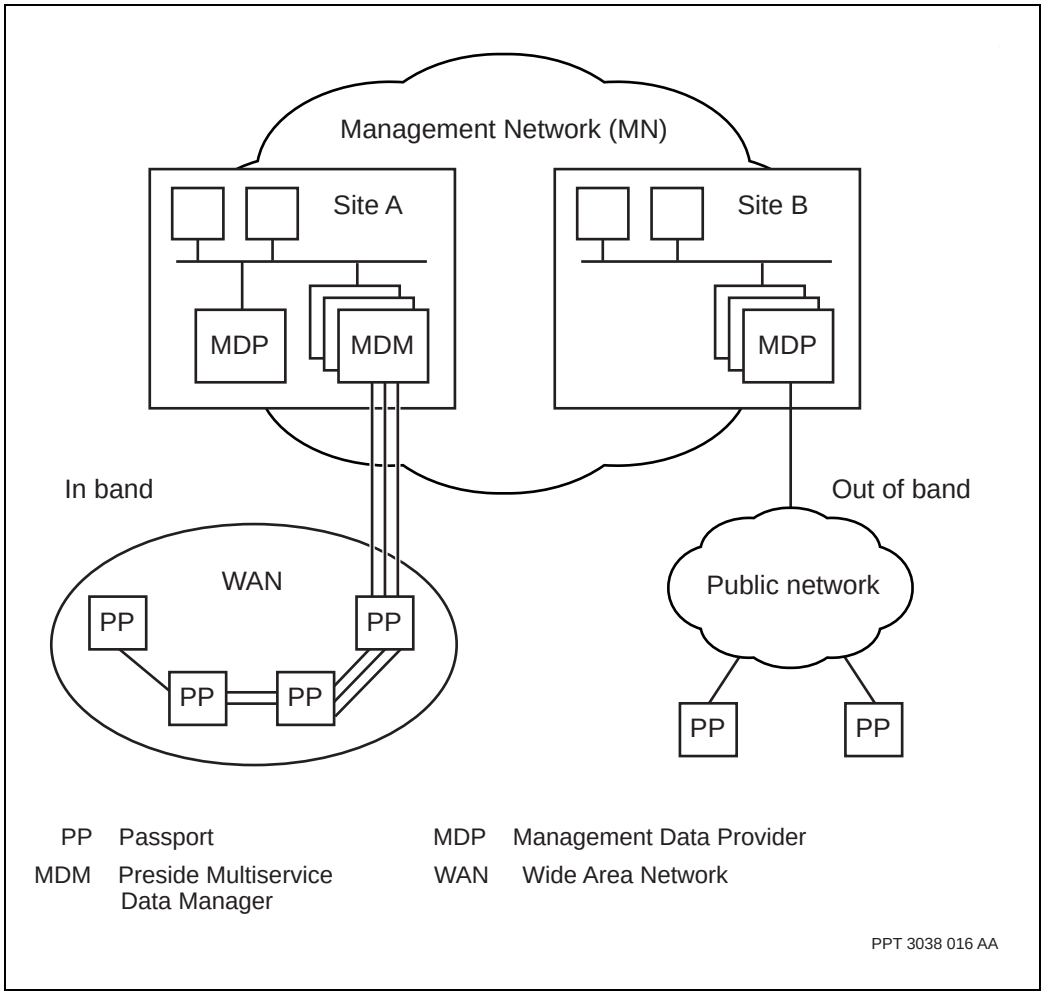

In-band connectivity

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.

For an example of a Passport network using in-band connectivity, see the figure In-band connectivity example using SunATM NIC (page 12). The workstation is connected to the workstation using a SunATM network interface card (NIC).

In-band connectivity example using SunATM NIC

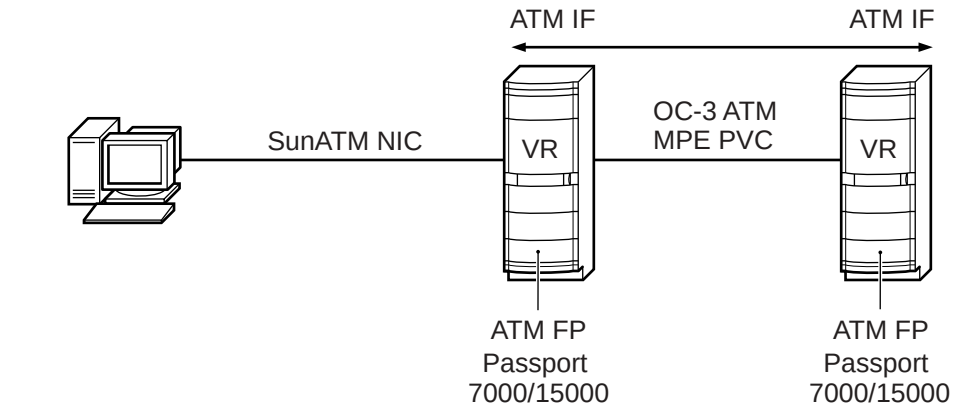

ATM asynchronous transfer mode
FP function processor
IF interface
MPE multi-protocol encapsulation
NIC network interface card
PVC permanent virtual circuit
VR virtual router

PPT 3038 001 AA

For an example of a Passport network using in-band connectivity, see the figure In-band connectivity example using a router (page 12). The workstation is connected to the Passport network by a router.

In-band connectivity example using a router

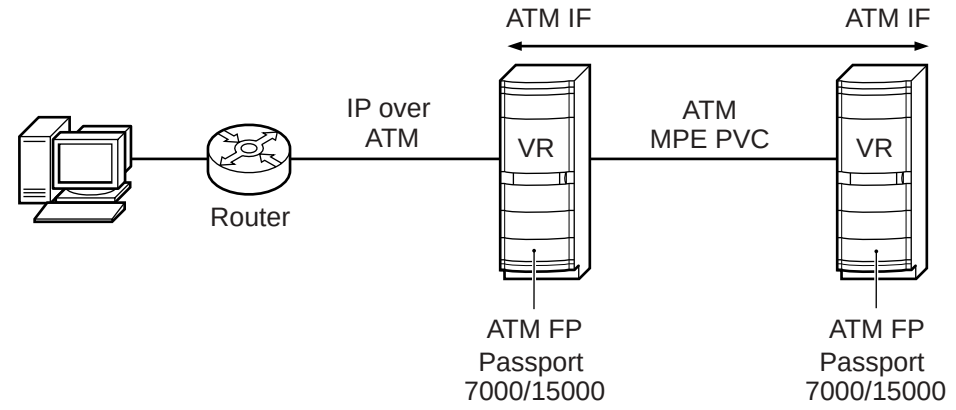

ATM asynchronous transfer mode
FP function processor
IF interface
IP internet protocol
MPE multi-protocol encapsulation
PVC permanent virtual circuit
VR virtual router

PPT 3038 001 AB

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.

For an example of a Passport network using out-of-band connectivity see the figure Out-of-band connectivity example (page 14).

Out-of-band connectivity example

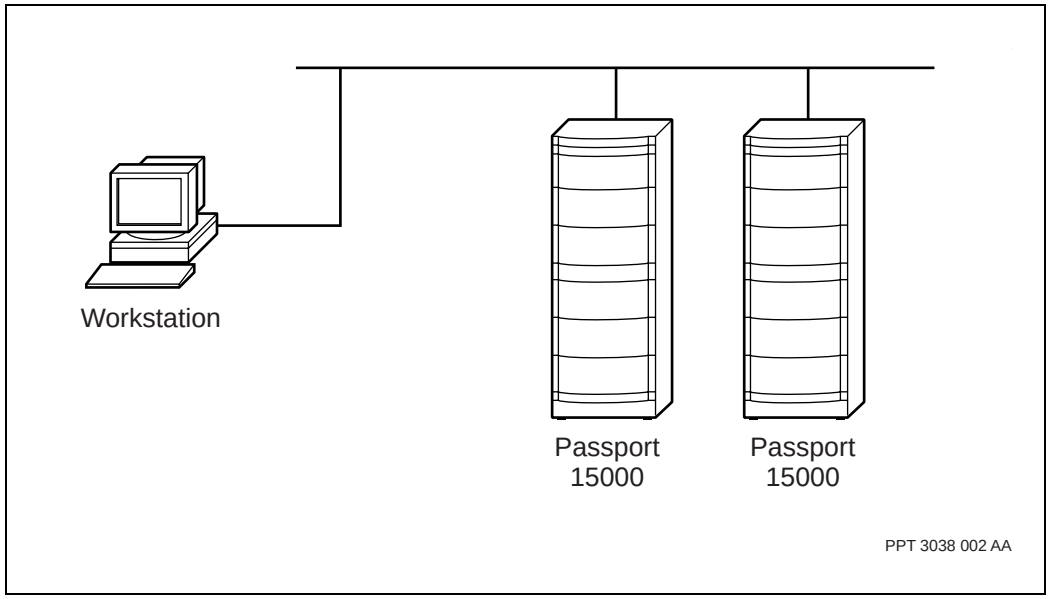

Passport on-switch management protocols

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*.

Distributed architecture

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

- Distributed surveillance architecture overview (page 15)
- Surveillance servers (page 16)
- Regionalization (page 19)
- Surveillance redundancy (page 21)
- Operator Client architecture (page 22)
- Host Group Directory server rules (page 24)
- Network data access mediator (page 25)

Distributed surveillance architecture overview

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.

For an example of stand-alone, management console client, and distributed client server sets for workstation configurations, seesee the figure MDM supported distributed workstation configurations (page 16).

MDM supported distributed workstation configurations

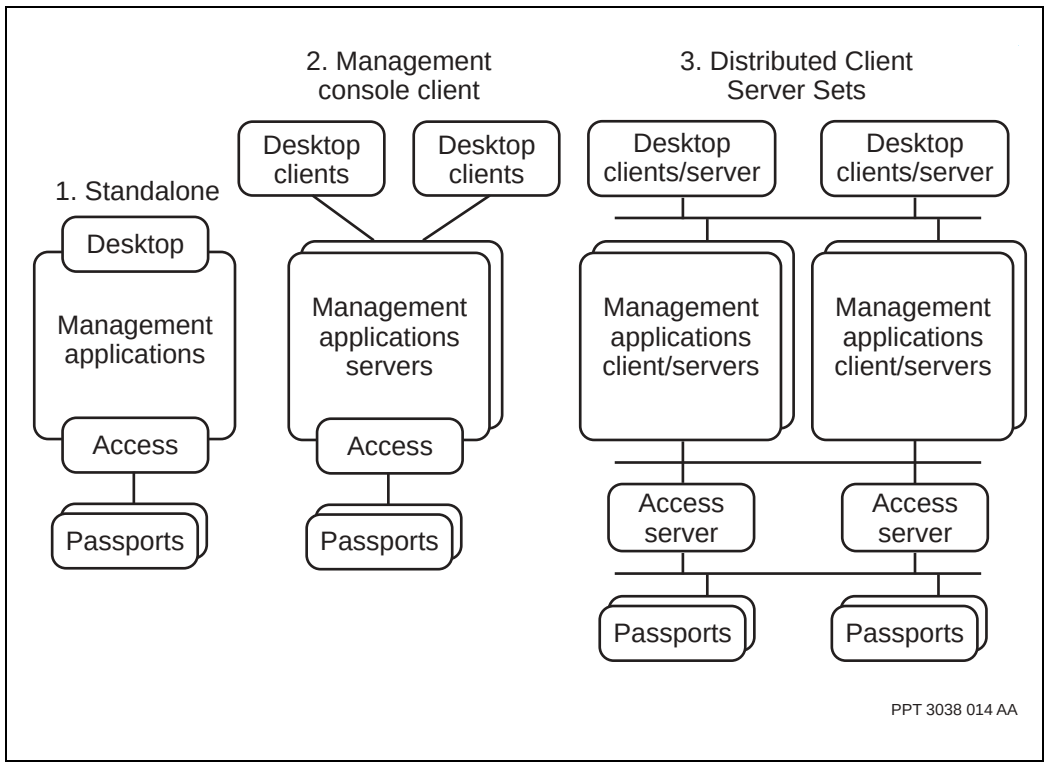

Surveillance servers

This section discusses the scalability of the surveillance servers. The servers can be located on one workstation or on multiple workstations (see the figure Surveillance server architecture (page 17)). The location of the servers depends on the size of the network and the number of network operators. The topics are discussed:

- Generic DCD (page 17)
- SMDR server (page 18)
- FMDR server (page 18)
- System and application management agents (page 19)
- GMDR (page 19)
- GMDR-GMDR filtering for high cost WAN links (page 19)

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

Surveillance server architecture

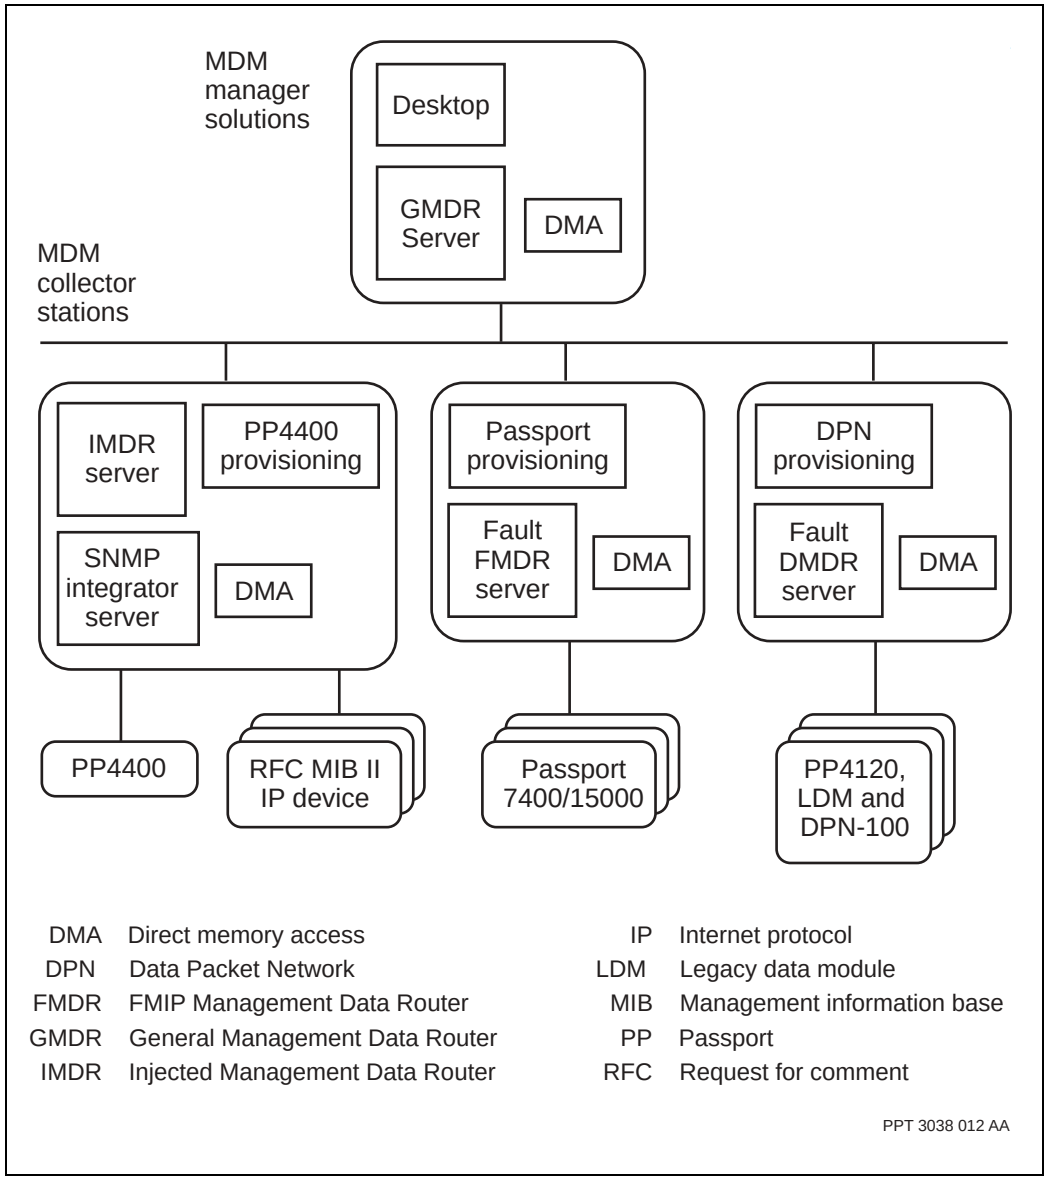

Generic DCD

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.

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

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.

For an example of a WAN configuration with a Passport asynchronous transfer mode (ATM) backbone network being managed from its global headquarters, see the figure Regional surveillance capabilities (page 20). 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.

Regional surveillance capabilities

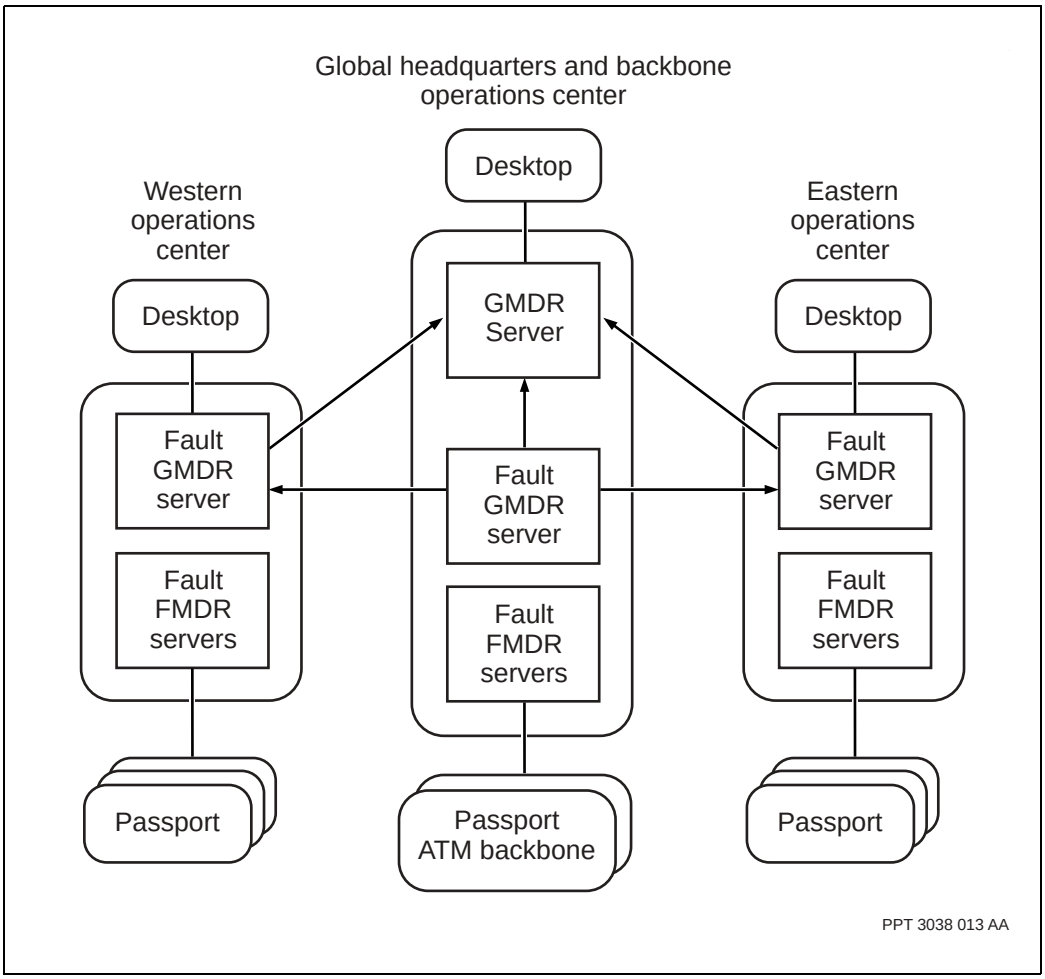

The GMDR server can filter out duplicate surveillance alarms and state events. The GMDR server includes alarm acknowledgement that allows you to know when an operator from any site has acknowledged an alarm. This reduces operations co-ordination efforts and assures an audit trail of actions.

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

For an example of 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, see the figure Redundant surveillance server architecture (page 22)

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.

Redundant surveillance server architecture

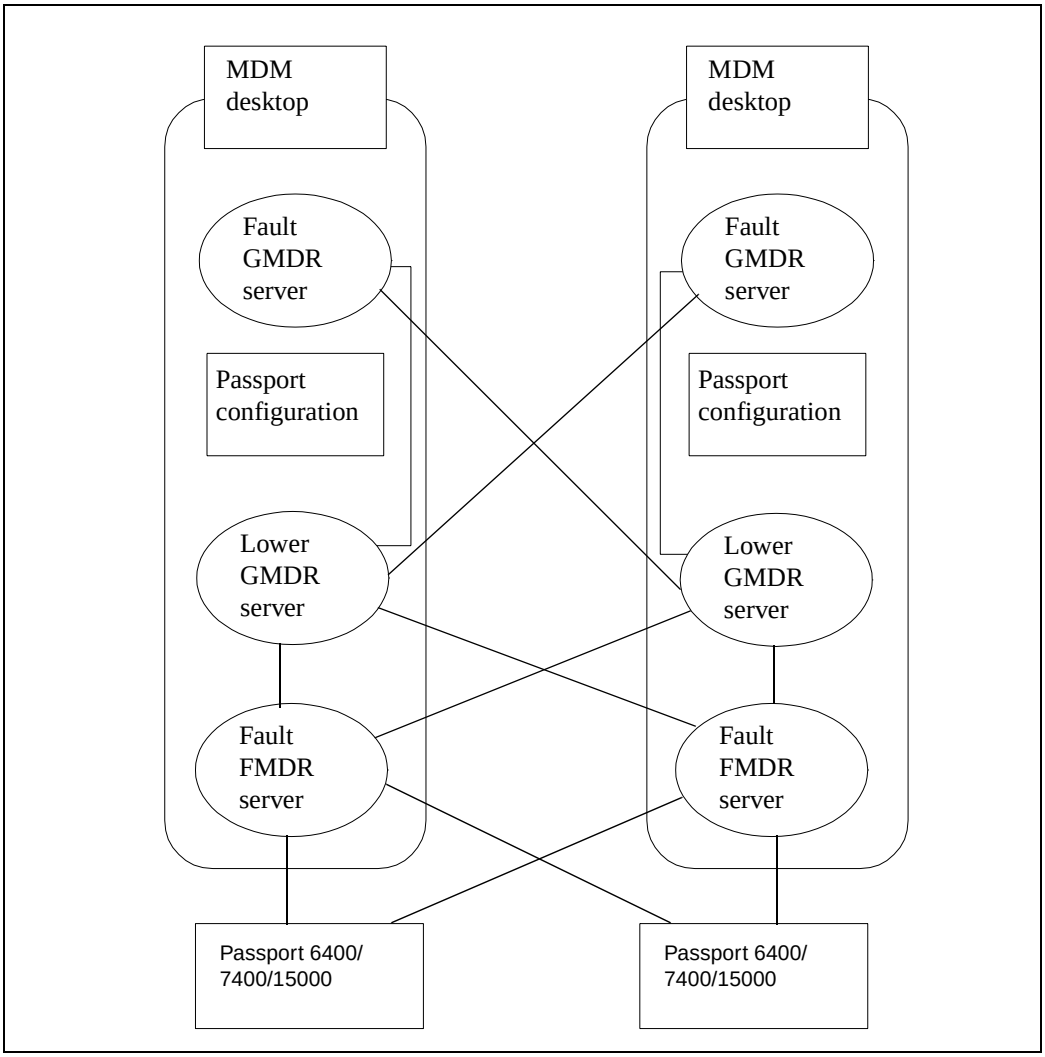

Operator Client architecture

MDM Operator Client provides a desktop environment for Windows-based PCs or Unix-based SUN Workstations. With this architecture one MDM workstation acts as a User Administration server which contains all necessary software and user data. Client PCs or workstations use a web browser to connect to the User Administration server through a URL link. Once connected, the browser downloads the software required to support fault, configuration and performance applications. Operator Client software is never installed on the client PCs and workstations, but is downloaded the first time a user logs in, and is updated whenever the software is upgraded on the User Administration server. User access is controlled through user profiles stored in the User Administration server.

An Operator Client deployment consists of the following four components: the MDM User Administration Server, MDM Operator Clients, MDM Server workstations, and the MDM Toolset. These components can exist on separate workstations or can be combined on one workstation. For a simplified example of these components, see the figure MDM Operator Client components (page 24).

MDM User Administration server

An MDM User Administration Server contains the software that centralizes Java Web Start (JWS) enabling, and Help Server for the Operator Clients. A second User Administration server can be set up for backup purposes. Operator Client users log into this server to download the desktop environment required to run MDM on the client PCs or Workstations.

In addition to supporting Operator Client functions, the User Administration server acts as a central repository for user definitions. By centralizing user definitions, this server provides a single access point for users, controlling how they access Operator Client and Sun UNIX operations, as well as Passport and other network devices.

MDM Operator Clients

Clients can be established on Windows-based PCs or Unix-based SUN workstations. Clients must either locally install Java Run-Time Environment (JRE) and JWS on the client host, or use the auto-install feature to download the JRE from the User Administration Server. All other software required to support MDM on the client is downloaded from the User Administration server when the client establishes an IP connection. Clients can connect to the MDM User Administration server or to an MDM workstation configured as a server.

MDM Server Workstation

At least one or more MDM workstations must be setup with MDM server software. This server provides access to the network, collects data, and supports MDM applications.

MDM Toolset

The MDM toolset refers to the non-operator client version of MDM originally available on UNIX-based SUN workstations.

MDM Operator Client components

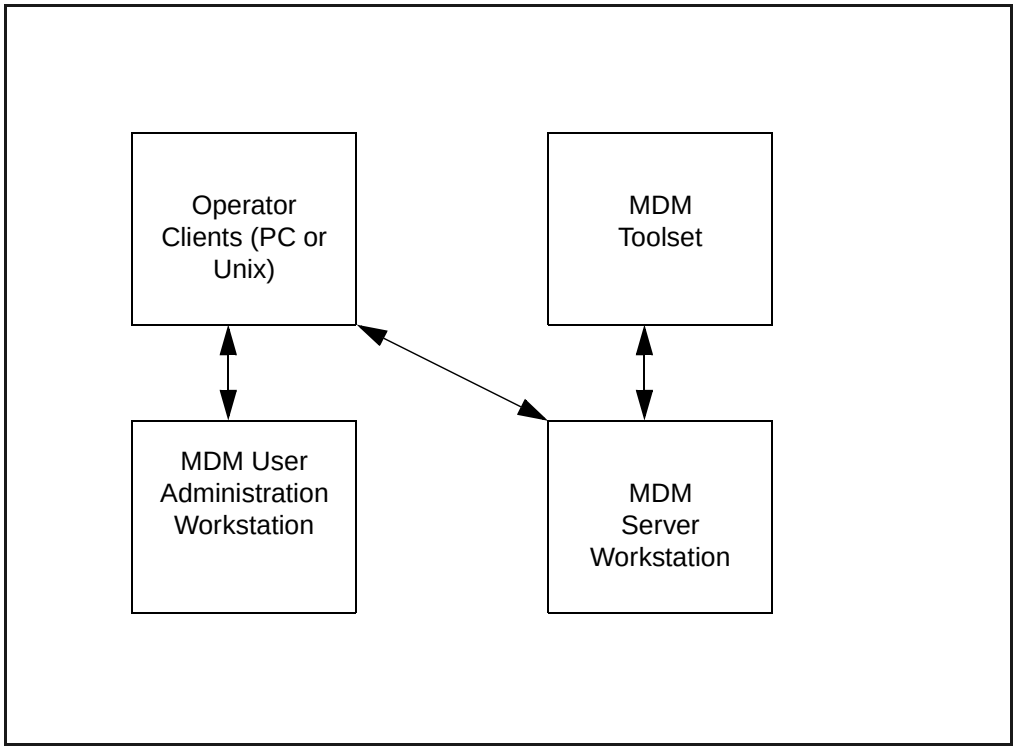

Host Group Directory server rules

For a basic overview of the session server architecture with Preside Multiservice Data Manager (MDM), see the figure HGDS session servers (page 25). 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.

HGDS session servers

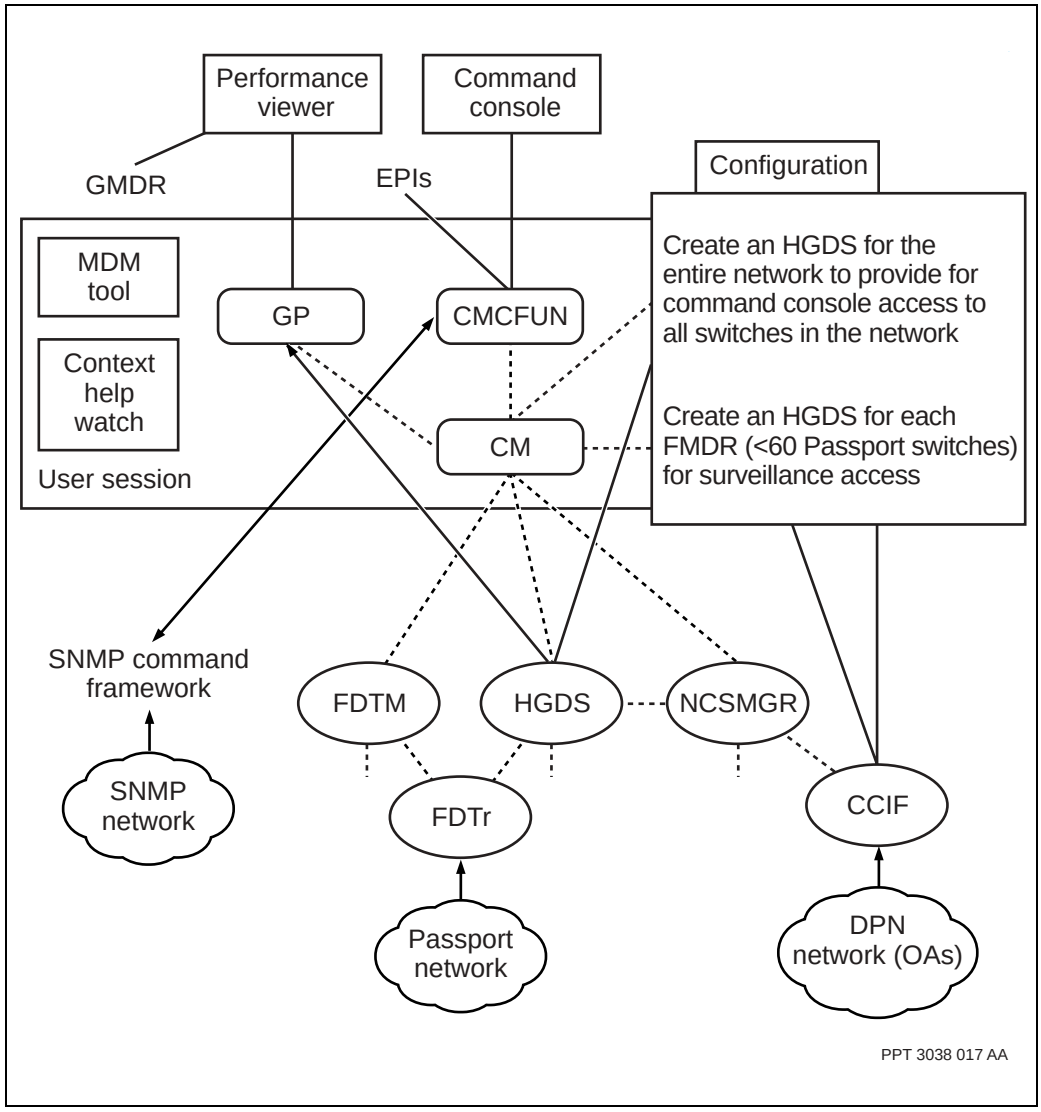

Network data access mediator

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. For an example of the NDAM server within a network, see the figure NDAM server (page 26).

NDAM server

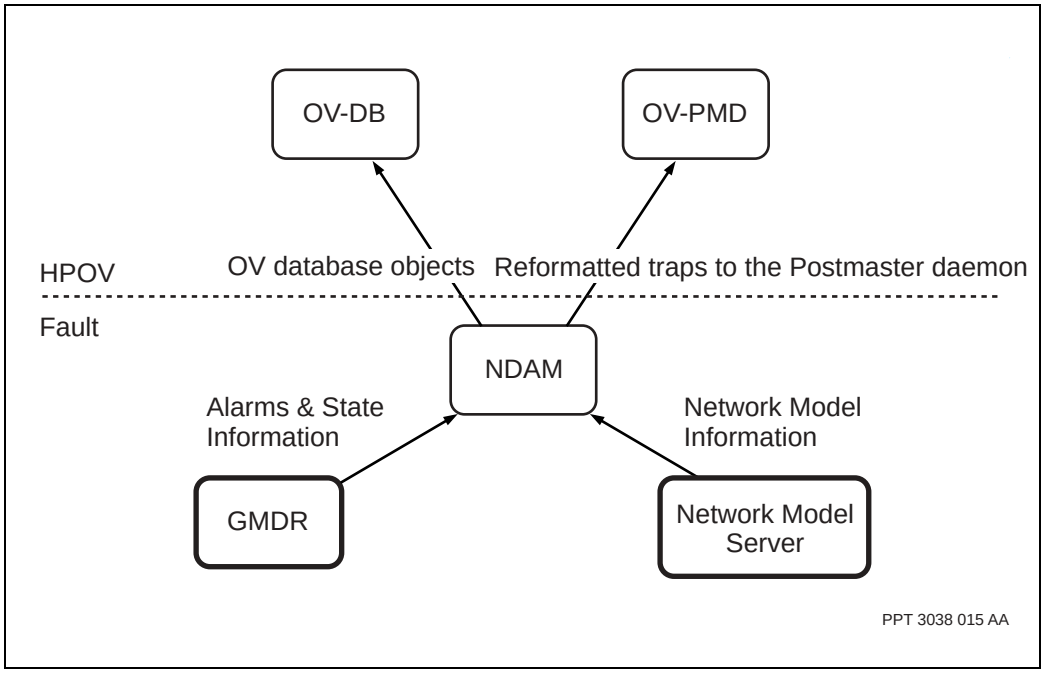

MDM servers

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:

- MDM servers to configure (page 28)
- Passport groups (page 28)
- FMDR server redundancy for surveillance access (page 30)
- Distribution of servers in large networks (page 31)
- NDAM server (page 32)

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.

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.

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.

For an example of a network containing three Passport switches that are monitored by two stand-alone workstations connected by a LAN, see the figure FMDR server redundancy for surveillance access (page 31). 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.

FMDR server redundancy for surveillance access

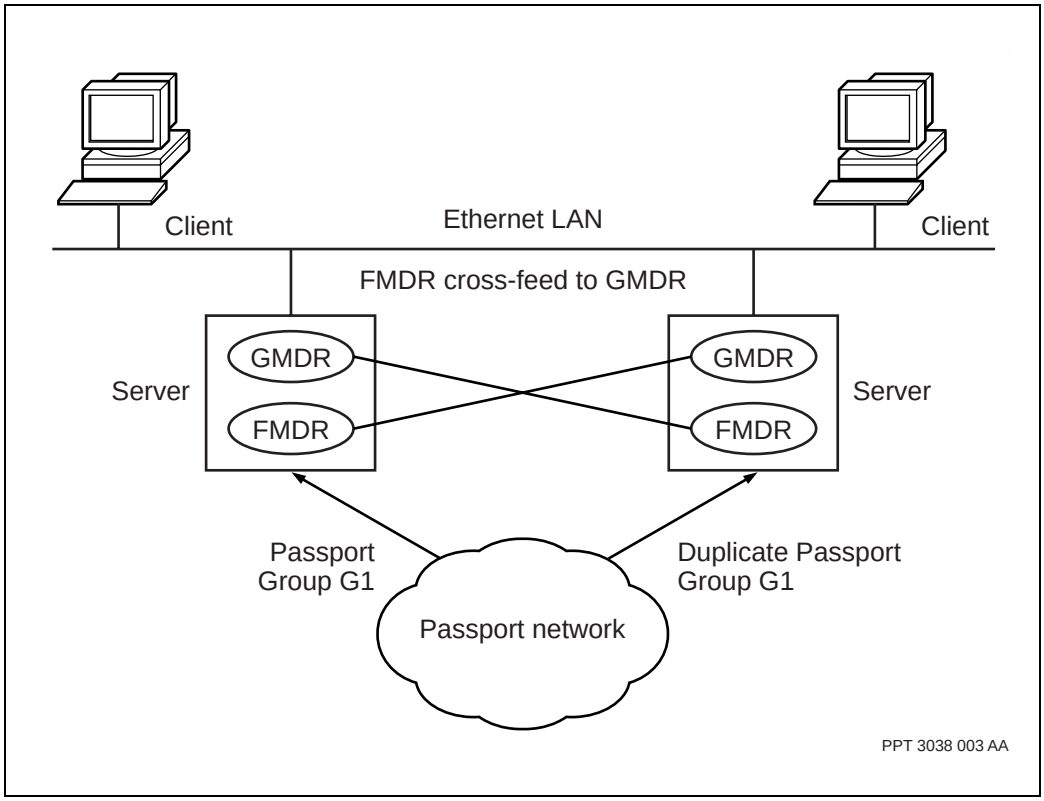

Distribution of servers in large networks

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 33). For more information on component type and regional filtering, see Component type and regional filtering (page 34).

For an example of how filtering is used for regional-central network management, see the figure Filtering based on thresholds (page 33).

Filtering based on thresholds

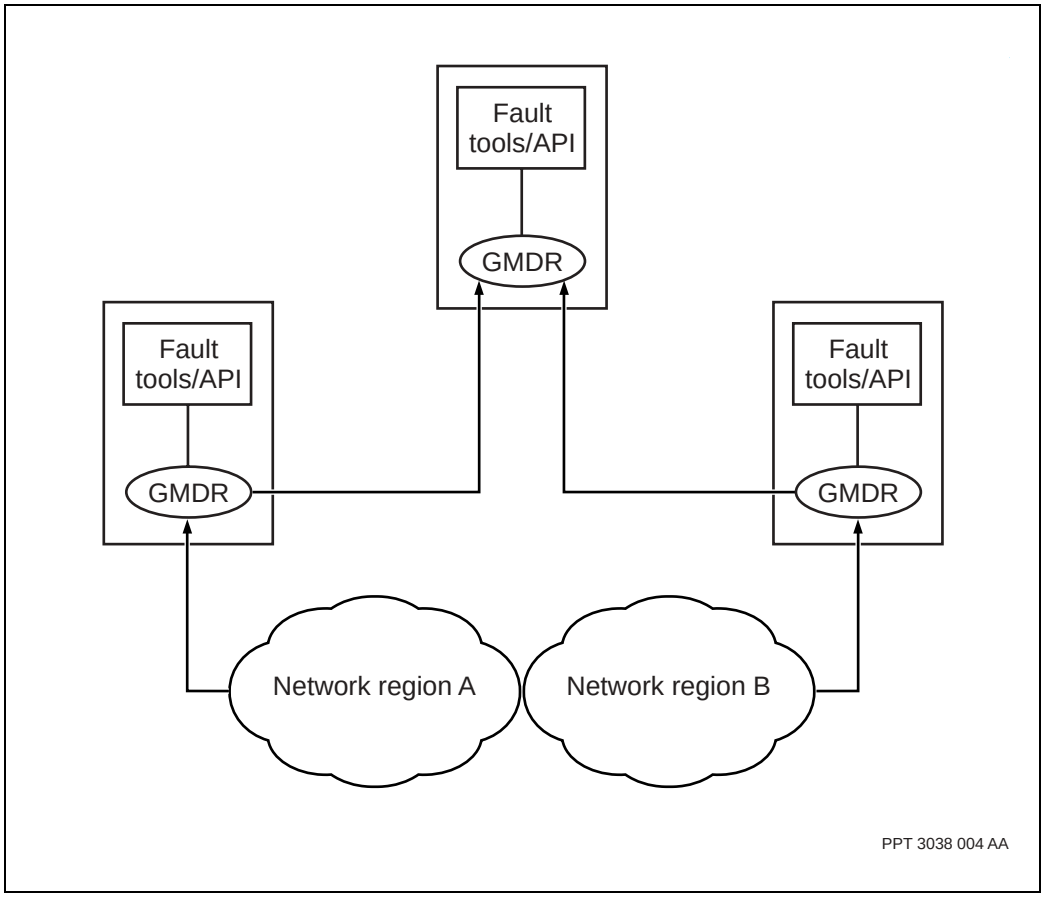

Component criticality thresholds

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 backbone

- 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 (see the figure Service name aliasing (page 35)). 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.

Service name aliasing

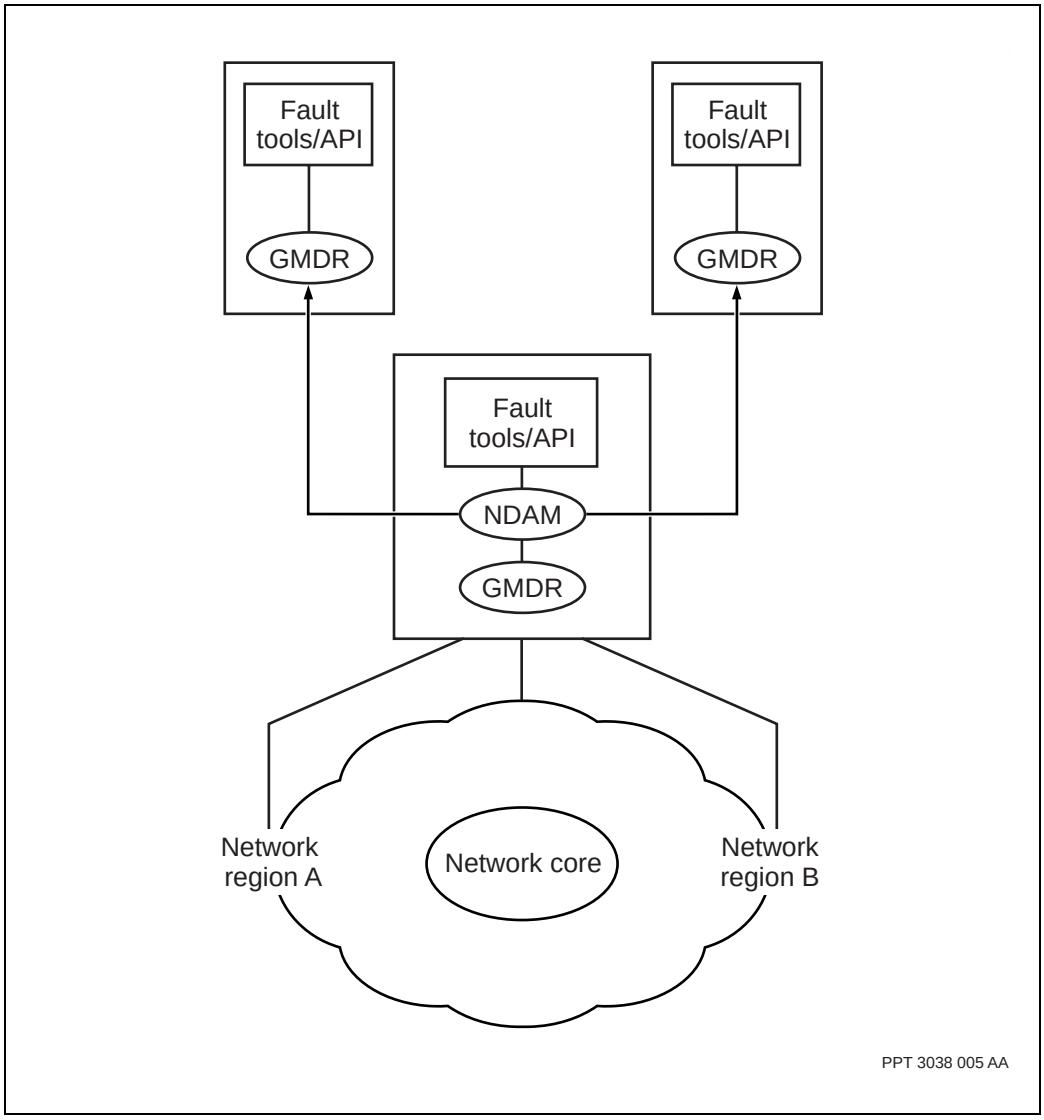

Choosing a configuration for MDM

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 (page 36)
- Stand-alone configurations (page 37)
- Stand-alone server model configuration (page 38)
- Stand-alone CPU server model configurations (page 39)
- Practical limits to stand-alone CPU server configurations (page 40)
- Client set/server set configurations (page 41)
- Combination configurations (page 44)
- Advantages of various server configurations (page 44)
- Disadvantages of server configurations (page 45)
- Counteracting the disadvantages of server configurations (page 46)

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 (Passport only, Data Packet Network (DPN) only, DPN and Passport, simple network management protocol (SNMP) devices, or a combination of these).

For an example of a typical configuration, seesee the figure Typical stand-alone configuration (page 38).

Typical stand-alone configuration

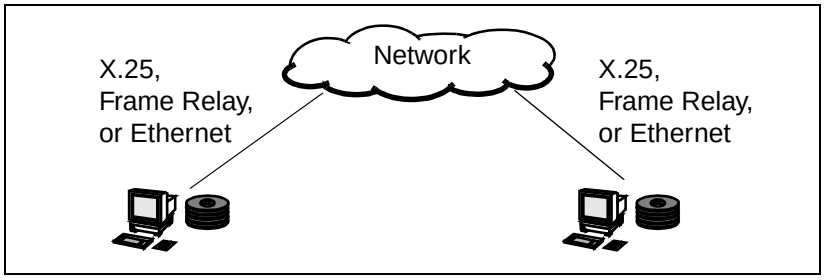

For a new, very small network, only one stand-alone workstation is required to manage the entire network.

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 <insert x-ref to practical limits to stand-alone CPU server configurations>.

Stand-alone server model configuration

This model of a stand-alone server configuration starts with a grouping in a network that contains five Passport switches (see the figure Typical stand-alone server configuration (page 39)). 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.

Typical stand-alone server configuration

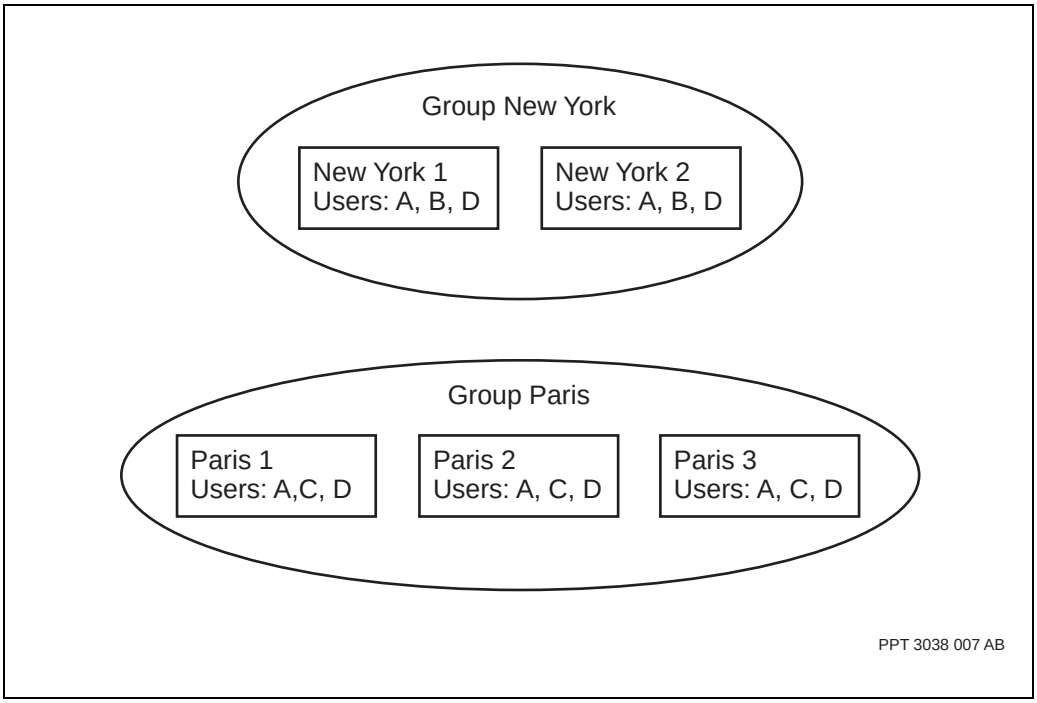

This network has the following administrative requirements:

- 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 model 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. For an example of a typical configuration, see the figure Typical stand-alone CPU server configuration (page 40).

Typical stand-alone CPU server configuration

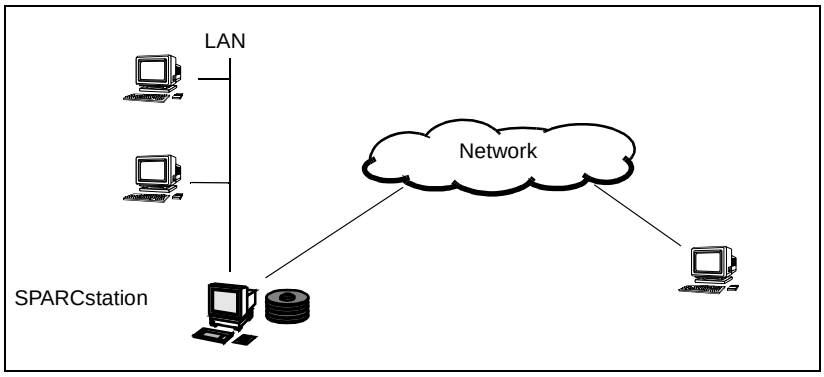

Practical limits to stand-alone CPU server configurations

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
- 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 241-6001-102 *Preside MDM Planning Guide* for guideline for selecting 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 (see the figure Client set and server set workstations (page 41)):

- 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.

Client set and server set workstations

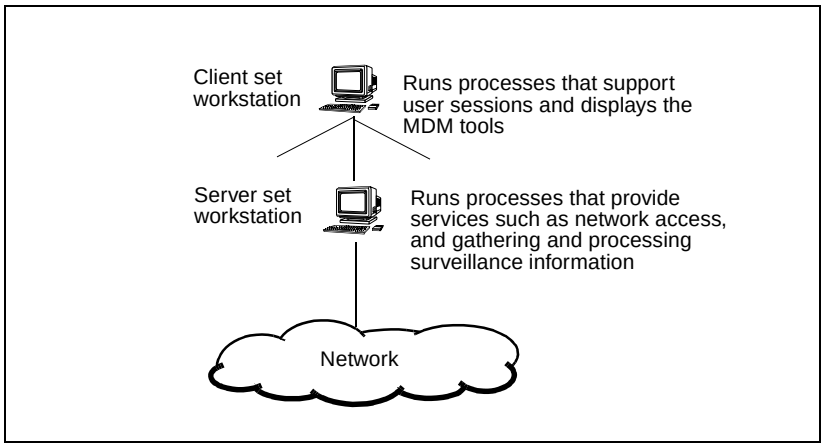

A workstation can be configured to run only the client set processes, only the server set processes, or both of these at the same time (stand-alone workstations).

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 (see the figure Service Selection (page 42)):

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

Service Selection

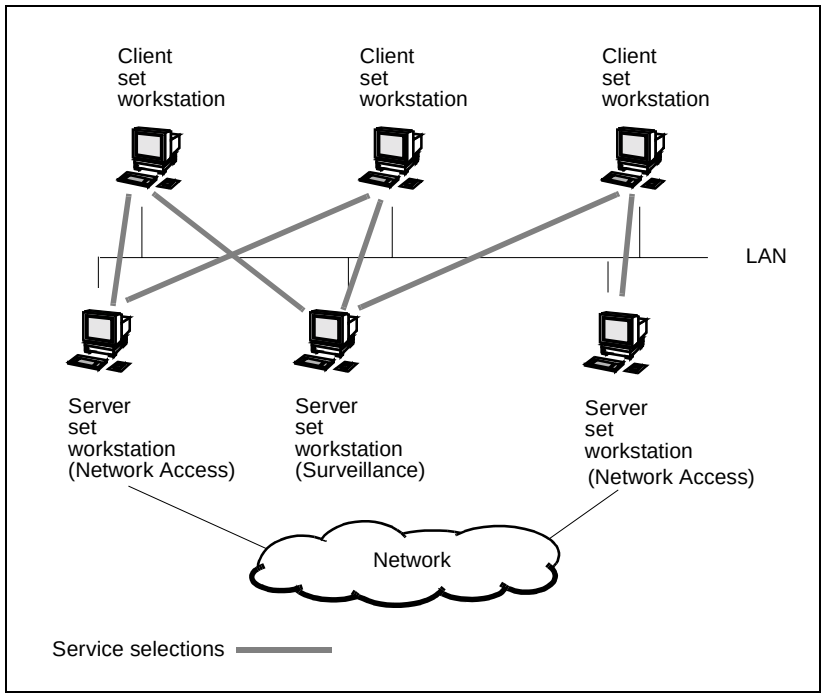

The Service Selection tool allows an administrator to set up the default server set workstations from which a client set workstation obtains support for one of the service areas. An operator can override these defaults when the following happens:

- 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

For an example of a typical server set and client set configuration with two server set workstations and several client set workstations, see the figure Typical client set/server set configuration (page 43). 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.

Typical client set/server set configuration

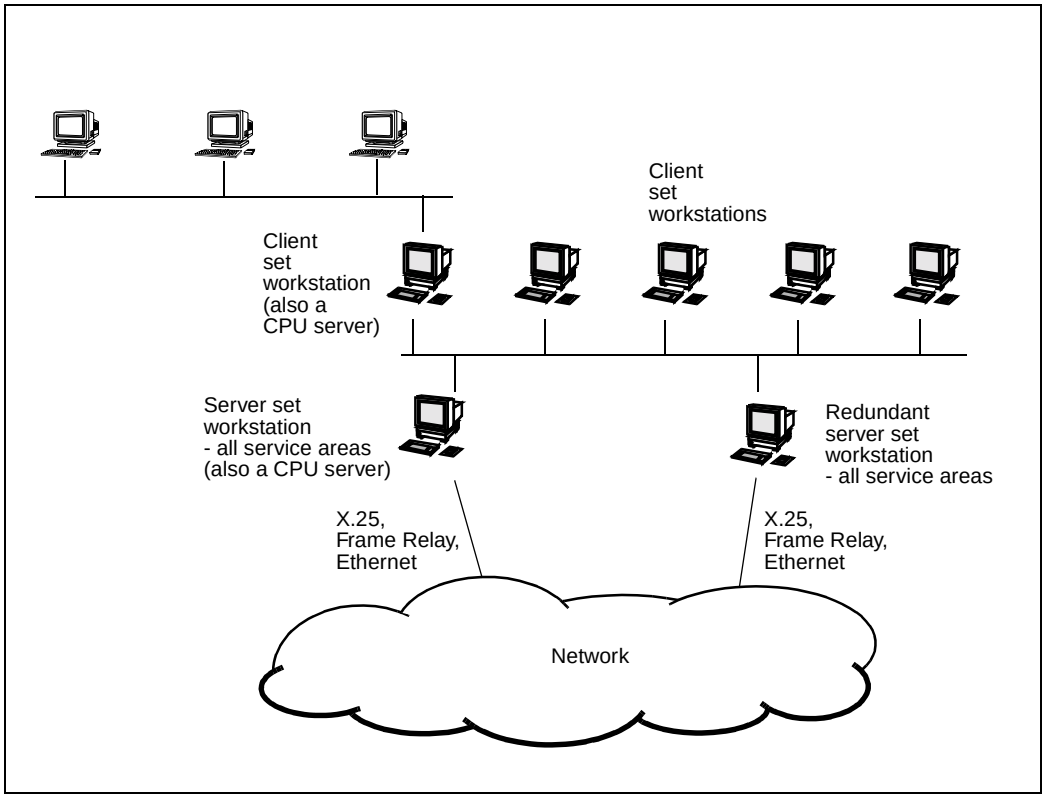

The client set and server set configuration is suited to medium and large networks with more than 200 modules. This configuration provides the following advantages:

- 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

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 of various server configurations

The primary advantage of the CPU server or client set and server set 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 of server configurations

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 52).

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. For an example of a typical configuration, see the figure Typical server configuration with backup stand-alone workstations (page 46).

Typical server configuration with backup stand-alone workstations

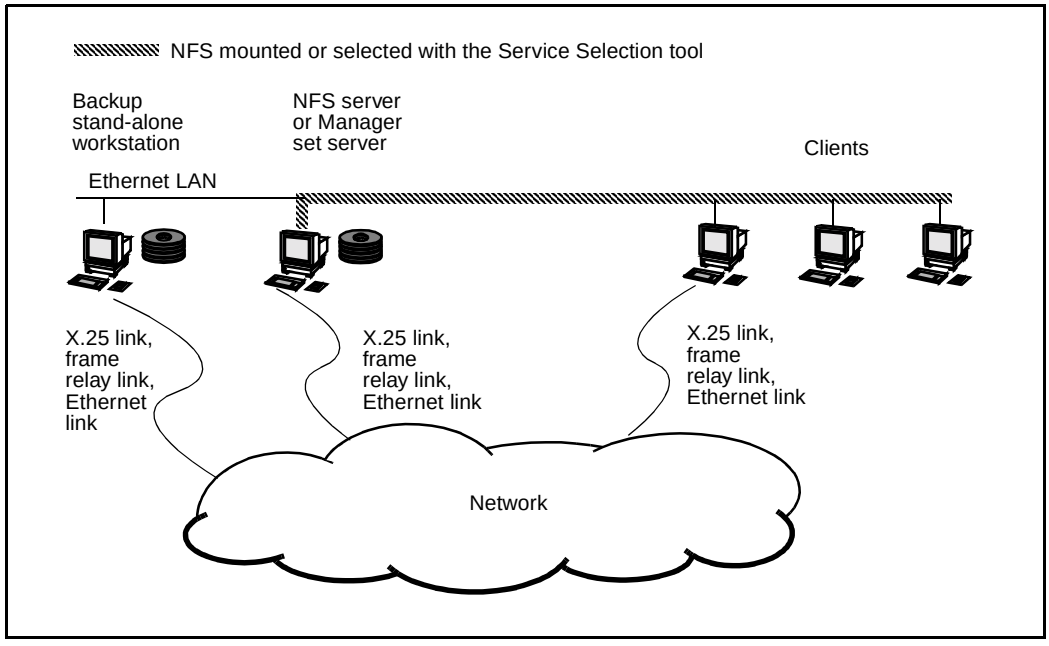

Server configuration with redundant servers

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. For an example of a typical configuration, see the figure Typical server configuration with redundant server workstations (page 47).

Typical server configuration with redundant server workstations

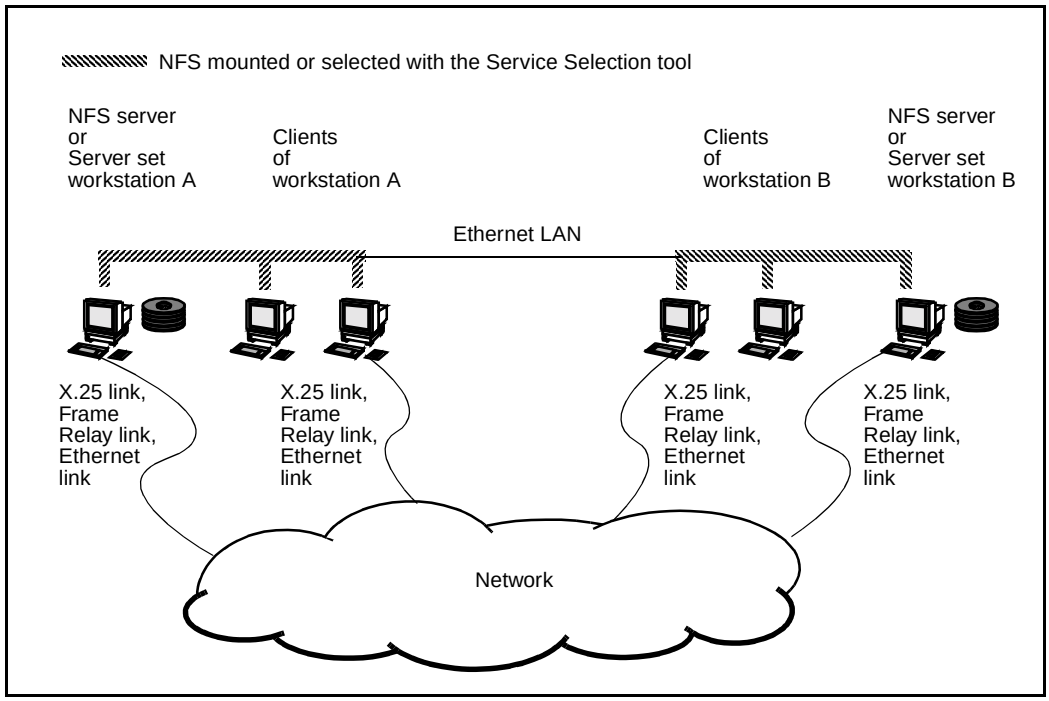

Minimizing response time

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

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. The following options are discussed

- Data protection by RAD (page 48)
- UNIX file system limitations (page 50)
- Veritas file systems (page 50)
- TCP delayed acknowledgement and bandwidth requirements (page 50)

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 for DPN Reference*.

Data protection by RAD

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.

For an example of how RAID 1+0 operates, see the figure RAID 1+0 (page 49). 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.

RAID 1+0

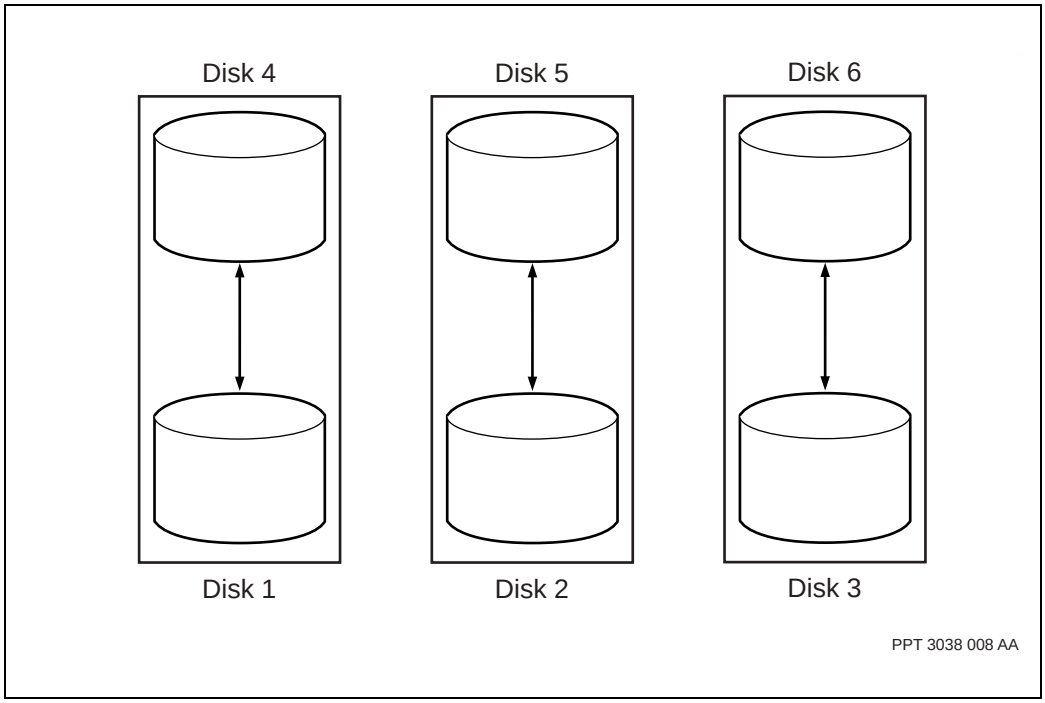

RAID 5

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.

For an example of how RAID 5 works, see the figure RAID 5 (page 50).

RAID 5

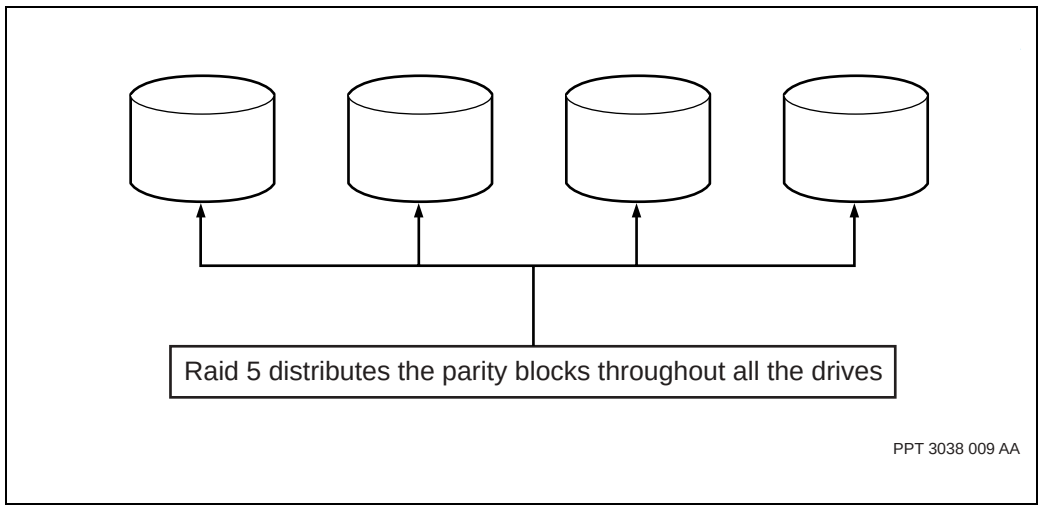

UNIX file system limitations

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 bandwidth 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.

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:

- System response time (page 52)
- Managing workstation resources (page 53)

Before using this section, determine the following:

- the size of the network
- the number of users
- the type of network connectivity
- the configuration you are using

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.

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.

Preside Multiservice Data Manager

Engineering Overview

Publication: 241-6001-101
Document status: Standard
Document version: 15.1RSUP
Document date: August 2004
Product release: 15.1