

HPDM Deployment Guide

HP Device Manager 4.7 SP3



Table of contents

Overview.....	3
Typical HPDM Topology	3
Port usage	4
Deployment factors	5
Hardware environment	5
Disk space usage of the database.....	5
Disk space usage of the repository	6
Network environment	6
Number of devices.....	6
HPDM logic.....	7
Multiple HPDM Consoles.....	7
Multiple HPDM Gateways	7
Advantages of additional repositories	7
Determining number of repositories	7
Simultaneously connected users	8
Scenarios	8
Small scale deployment	8
Normal scale deployment.....	9
Huge scale deployments.....	9
Optimizing HPDM configurations.....	10
Reducing the visible task amount	10
Batch tasks	10
HPDM Agent settings	12
Using rules.....	13
Appendix A: Use cases.....	14
Appendix B: Performance testing data.....	15
Appendix C: NAT environment.....	16
Appendix D: DMZ environment.....	18
Selecting a file transfer port	19
Resolving the child repository address.....	19
Using PASV mode on an FTP or FTPS server	19
Limitations.....	20
Ports between networks.....	20

Appendix E: HPDM failover redundancy	21
Redundancy scenario	21
Implementing redundancy	21
Local Traffic Manager	22
Global Traffic Manager	22
Summary	22
For more information	23

Overview

The purpose of this document is to assist customers who are planning the architecture of HP Device Manager (HPDM) components and configuring those components for the number of devices that will be managed. This document focuses on managing larger device deployments using HPDM. This document also includes tips to fine-tune the performance of HPDM.


After the installation of HPDM, it can be divided into the following components:

- HPDM Console
- HPDM Server
- Database (Here it means MS SQL Server, PostgreSQL is taken as an inner part of Server)
- HPDM Gateway
- Master Repository
- Child Repository (not necessary)
- HPDM Agent (pre-installed on device)

Typical HPDM Topology

The following diagrams show the topology of a typical HPDM deployment.

Figure 1. Topology of a typical HPDM deployment

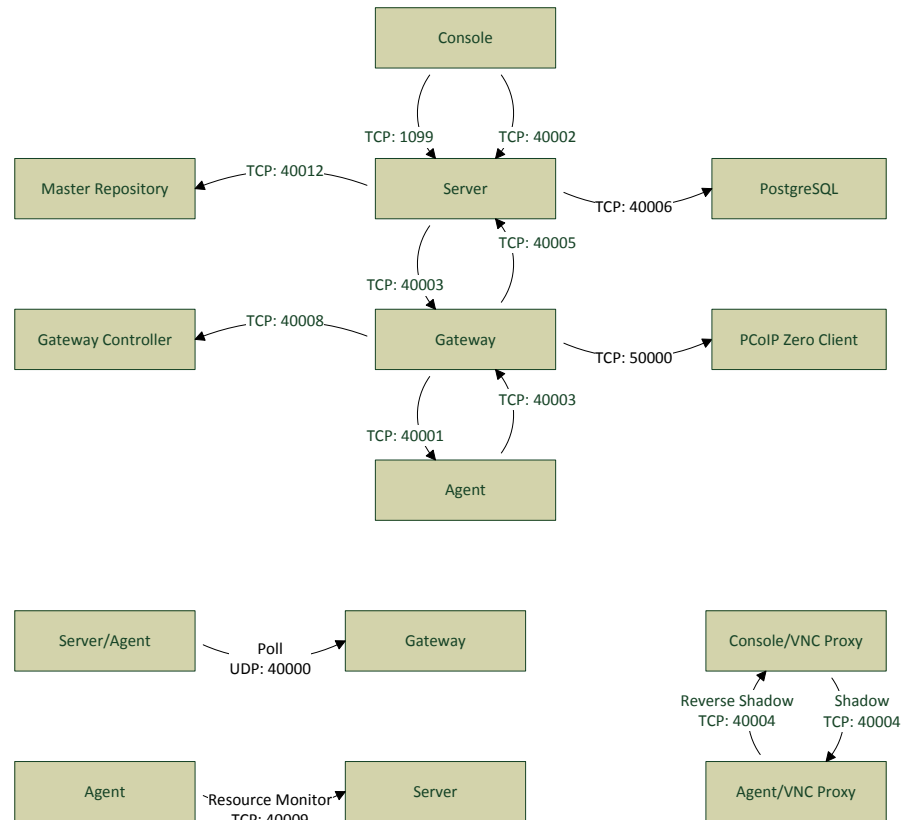


Note

If you want to generate templates through the FTP Software Component Browser, make sure that HPDM Console and Master Repository can reach to HP FTP Server, either through direct connection or an HP FTP Proxy configuration.

Port usage

Figure 2. Port usage in HPDM



Note

This chart lists only the basic ports created by HPDM. For more information about common software port usage in HPDM, see the *Administrator Guide* for HP Device Manager 4.7.

Make sure that the ports are not blocked by a firewall or hijacked by other processes.

Deployment factors

This section lists the main factors that might influence HPDM deployment, and also gives some deployment recommendations. The main factors are as follows:

- Hardware environment
- Network environment
- Number of devices
- HPDM logic

Hardware environment

The following table provides the minimum hardware requirements of HPDM components.

Table 1. System requirements

HPDM component	Operating system	Suggested minimum hardware
HPDM Console	<ul style="list-style-type: none">• Windows® Server 2003• Windows Server 2003 R2• Windows Server 2008• Windows Server 2008 R2• Windows Server 2012• Windows Server 2012 R2• Windows 7• Windows XP Professional	<ul style="list-style-type: none">• Pentium® IV or greater• 1 GB RAM• 1 GB free disk space• 100 Mbps NIC
HPDM Server	<ul style="list-style-type: none">• Windows Server 2003• Windows Server 2003 R2• Windows Server 2008• Windows Server 2008 R2• Windows Server 2012• Windows Server 2012 R2	<ul style="list-style-type: none">• Pentium IV or greater• 1 GB RAM• 2 GB free disk space• 100 Mbps NIC
Database (MS SQL Server)	<ul style="list-style-type: none">• Windows Server 2003• Windows Server 2003 R2• Windows Server 2008• Windows Server 2008 R2• Windows Server 2012• Windows Server 2012 R2	<ul style="list-style-type: none">• Pentium IV or greater• 2 GB RAM• 2 GB free disk space• 100 Mbps NIC
HPDM Gateway	<ul style="list-style-type: none">• Windows Server 2003• Windows Server 2003 R2• Windows Server 2008• Windows Server 2008 R2• Windows Server 2012• Windows Server 2012 R2	<ul style="list-style-type: none">• Pentium IV or greater• 1 GB RAM• 2 GB free disk space• 100 Mbps NIC
Master/Child Repository	<ul style="list-style-type: none">• Windows Server 2003• Windows Server 2003 R2• Windows Server 2008• Windows Server 2008 R2• Windows Server 2012• Windows Server 2012 R2	<ul style="list-style-type: none">• Pentium IV or greater• 1 GB RAM• 2 GB free disk space• 100 Mbps NIC

Disk space usage of the database

The disk space usage of the database grows with the device and task amounts. Calculate the required disk space with the following pattern:

- The initial disk space is less than 50 MB.
- Add an additional 100 MB for every 1,000 devices.

- Add an additional 1 MB for every 100 tasks.

Disk space usage of the repository

The disk space usage of the repositories grows with the payload amount, especially with images of a device operating system. Make sure that the disk space is sufficient to hold all payloads and tools.

Table 2. Recommended size reserved for repository

WES 2009	2 GB
WES 7	4 GB
WES8	5 GB
HP ThinPro	1 GB

Note

All hardware configurations have been tested on virtual machines with comparable virtual settings.

This table is the suggested minimum size environment for each HPDM component on a single machine. Enlarge the parameters if you are deploying multiple HPDM components or other software services on one machine. Other deployment factors might also influence the hardware requirements.

Network environment

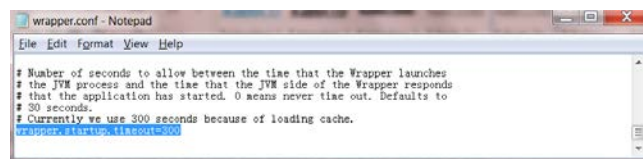
There are many network factors that might influence the deployment of HPDM, such as the network bandwidth or whether HPDM components are deployed on a subnet. Some companies might have different network strategies, such as devices must be in a NAT environment, or devices must be deployed in different regions, or the HPDM components cannot connect to Internet directly. Deploy HPDM according to the situation.

To manage large-scale deployments, HP recommends having HPDM Server and the HPDM Gateway services installed on same subnet as the database. A server-type operating system is needed because of the half-open connection limit on client operating systems.

HP recommends deploying HPDM Server as close to the database server as possible, because the network latency between them has a significant impact on the performance. It causes obvious delays for HPDM Console users if the network latency between HPDM Server and database is higher than 30 milliseconds. Also, HP recommends deploying a Child Repository as close to its related devices as possible.

Note

If latency is unavoidable, go to ...\\Server\\conf\\wrapper.conf, and increase the value of **wrapper.startup.timeout** in wrapper.conf on the HPDM Server side. The default value is 300 seconds. This makes HPDM Server more tolerant to database latency.



Number of devices

As the total number of managed devices increases, use more powerful, server-level hardware configurations such as RAID. 100,000 is the recommended maximum number of devices in one HPDM deployment. Lab testing shows that HPDM performs most efficiently up to this number.

Current customers successfully use HPDM to manage up to 100,000 devices with a single HPDM Server using the following considerations:

- Use of networks faster than 1000 Mbps
- Efficient placement of gateways and repositories to maximize the use of LAN-based communications

- Use of scheduled tasks to execute tasks during non-peak hours of operation

Note

HPDM is a very flexible system and supports the use of any number of HPDM Servers, HPDM Gateways, and repositories to match the customer's existing network architecture.

HPDM logic

There are many factors of HPDM logic that might influence the deployment. See following list:

- HPDM Server does not support clusters. Currently, you can deploy more than HPDM Servers in your environment according to your requirements. Under one HPDM Server, you can deploy the following:
 - Multiple HPDM Consoles
 - One database
 - Multiple HPDM Gateways
 - One Master Repository
 - Multiple Child Repositories
- PXE imaging
 - To use PXE imaging, deploy an HPDM Gateway in the same subnet as the devices.
 - If the subnet is a NAT subnet, configure port mapping on NAT to make sure that HPDM Server can talk to HPDM Gateway directly.

Multiple HPDM Consoles

Currently, HPDM does not limit the number HPDM Consoles connected to an HPDM Server. From the results of extreme performance testing (using 50,000/100,000 devices, one HPDM Gateway, and one HPDM Server), sending tasks from HPDM Consoles to all devices at the same time, HP recommends sending a task from less than five HPDM Consoles at the same time.

Multiple HPDM Gateways

HPDM does not limit the number of HPDM Gateways connected to an HPDM Server. One HPDM Gateway has verified good performance from the performance testing (using 50,000/100,000 devices, one HPDM Gateway, and one HPDM Server), so HP recommends not using too many HPDM Gateways in under one HPDM Server. For some core centers and regions, multiple HPDM Gateways are preferred for the following reasons:

- Required for PXE imaging tasks
- Consolidated communication between branch offices and DMZ
- Accelerated task delivery speed when there is a NAT Gateway in branch offices

Advantages of additional repositories

As the volume of transferred files increases, add more Child Repositories for the following reasons:

- Reduced traffic between branch offices and DMZ
- Faster software updates and image deployment

Determining number of repositories

To determine the number of Child Repositories required, use the following formula:

Number of repositories = (transferred data ÷ bandwidth) ÷ expected time spent

For example, if you have 20,000 units to be re-imaged and each image is 1 GB, you have 20,000 GB (20 TB) of data to transfer. With a connection of 100 Mbps from one repository to a device, it takes 444.4 hours to transfer all data.

$20,000 \text{ GB} \div (100 \text{ Mbps} \div 8 \text{ bits per byte} \div 1000 \times 3600 \text{ seconds per hour}) \approx 444.4 \text{ hours}$

To reduce the data transfer time to 48 hours, you need 10 repositories ($444.4 \div 48$). Keep in mind that there is some overhead to synchronize from the Master Repository to the Child Repositories.

Simultaneously connected users

There is no significant increase of memory or CPU usage on HPDM Server for any additional, connected HPDM Consoles. However, it requires more resources for database access and communication. Do not have too many connected HPDM Consoles.

Scenarios

The following following are some typical scenarios.

Table 3. Minimum requirements for typical scenarios

Devices	HPDM Servers	HPDM Gateways	Database
1–5,000	1	1	PostgreSQL
5,000–20,000	1	1	PostgreSQL or MS SQL Server
20,000–50,000	1	≥1	MS SQL Server
> 50,000	≥ (device number/50,000)	≥3	MS SQL Server

Note


This table introduces the minimum requirements of some typical scenarios, but you must deploy your environment according to your network situation and company strategies, such as whether devices are in a NAT environment or distributed in different places.

Small scale deployment

Device number: <5000

Deployment: 1 HPDM Server, 1 HPDM Gateway, PostgreSQL (or MS SQL Server), 1 Master Repository

Figure 3. Small scale topography



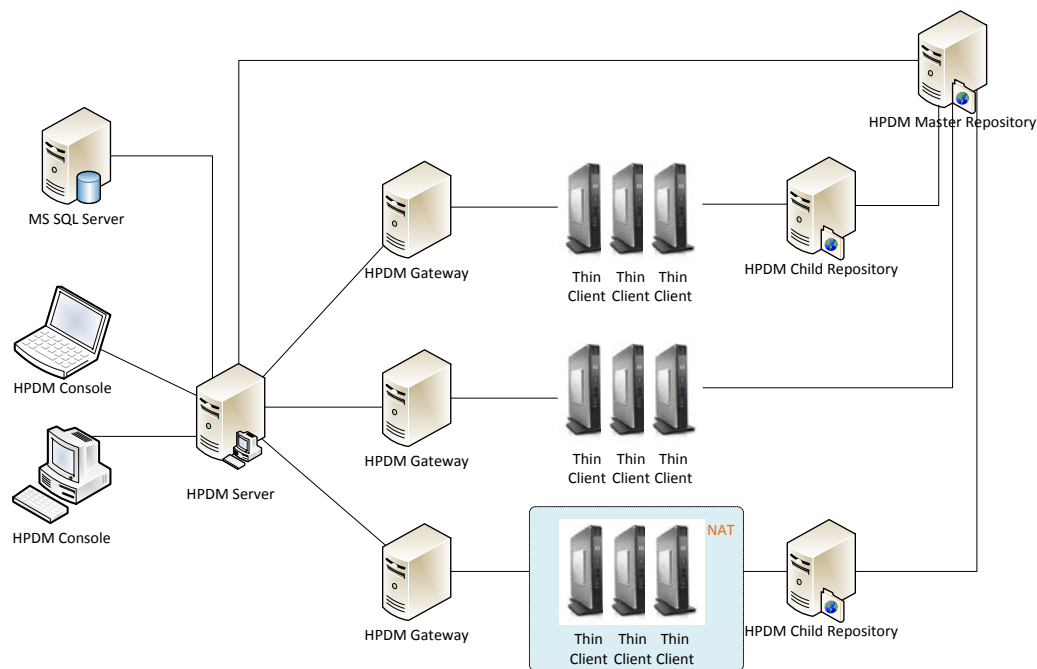
This is a small scale deployment, so the minimum requirement is that you can deploy all HPDM components on one machine.

Normal scale deployment

Device number: 25,000

Deployment: 1 HPDM Server, MS SQL Server, 3 HPDM Gateways, 1 Master Repository, 2 Child Repositories.

Figure 4. Normal scale topography



HP recommends that you deploy each HPDM component on its respective machine. In the case that the Master Repository overloads, there are two Child Repositories to divide the file transmission pressure. There are three HPDM Gateways to separate all devices into three groups. Note that one device group is behind a NAT environment. Use the HPDM poll function to manage those devices.

Note

See Deployment factors for hardware and other requirements.

Huge scale deployments

Device number: >50,000

Deployment: Because one HPDM Server supports up to 50,000 devices with verified performance, deployments of greater than 50,000 devices require multiple instances of HPDM. You might view it as multiple Normal scale deployments to deploy.

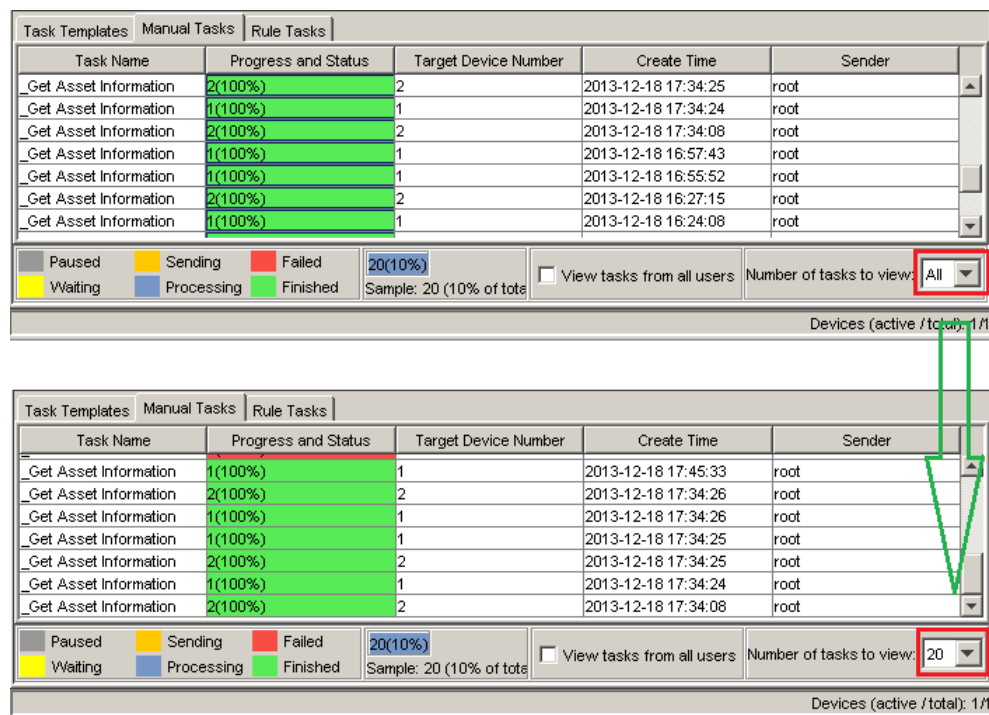
Optimizing HPDM configurations

HPDM manages devices remotely with tasks. Configuring the most suitable parameters for tasks can help improve performance.

Reducing the visible task amount

The following tip increases the speed of the startup process of HPDM Console. HPDM Console takes time to render task summaries when initializing. Before closing HPDM Console, open the **Manual Tasks** tab and select a smaller number under **Number of tasks to view**. Then, change the setting in the **Rule Tasks** tab.

Figure 5. Setting the number of tasks to view



Batch tasks

Batch Control is an optional task parameter that can be used when deploying images or other tasks with a payload to a large number of devices. The batch settings control how many devices are sent the task at a time, which gives you some control over the amount of network traffic that HPDM generates.

The continued write speed of current devices is around 3.5 MBps (megabytes per second). Use that number to decide how many imaging tasks you can send in one batch for each repository and how long the interval between batches should be.

For example, suppose you are deploying a 4 GB image to your devices. The network is a high-quality, 1000 Mbps Ethernet, which means it can provide continuous theoretical bandwidth. Many of the latest 7200 rpm hard drives have higher continuous read speeds than that, and those using RAID even higher. Assuming that the network is the bottleneck, we have the following:

$$1000 \text{ Mbps} \div 8 \text{ bits per byte} \div 3.5 \text{ MBps} \approx 35 \text{ concurrent imaging tasks in a batch}$$

To calculate the time it will take:

$$4000 \text{ MB} \div 3.5 \text{ MBps} \div 60 \text{ sec per min} \approx 20 \text{ minutes to transfer the image file}$$

The actual imaging task takes longer because there are extra steps before and after the file transfer. However, you can start another batch without influencing the previous batch because the extra steps do not need much bandwidth. In this example, the recommended value would be 30 units per batch and 30 minutes between batches, because HP recommends leaving a buffer. The devices' operating system makes no difference; only the image size makes a difference.

There are three kinds of tasks: PXE, FTP, and general. PXE tasks are PXE imaging tasks. FTP tasks are tasks that need file transfer. (PXE tasks use file transmission.) All other tasks are general tasks. The PXE server and file repositories are bottlenecks for simultaneous task executions.

Note

A payload task is different from general task that has a payload using file transmission, HPDM supports FTP, FTPS, SFTP, and SMB protocols for transferring files. Compared to FTP, FTPS (FTP over SSL) is slower, and SFTP (SSH File Transfer Protocol) is the slowest. If you use these protocols for repositories, set a larger amount on units per batch and a longer interval between batches to make full use of your bandwidth.

Define default task batch parameters for different types of tasks in the **Configuration Management** dialog.

Figure 6. Setting default batch amount and interval values

The screenshot shows the 'Configuration Management' dialog box. On the left is a tree view with 'Task Parameters' selected. The main area is titled 'Set the duration HP Device Manager will wait for the execution of tasks.' It contains several input fields: 'Valid Time' (1440 minutes), 'Execution Timeout' (30 minutes), 'General Batch Amount' (0 devices), 'General Batch Interval' (5 minutes), 'PXE Batch Amount' (2 devices), 'PXE Batch Interval' (10 minutes), 'FTP Batch Amount' (5 devices), and 'FTP Batch Interval' (2 minutes). The last four fields are grouped in a red box. Below these is an 'Exclude Working Hours' section with a checkbox and time pickers for 'Working Hours Start' (09:00 AM) and 'Working Hours End' (05:00 PM). At the bottom right are 'OK', 'Cancel', 'Apply', and 'Restore Defaults' buttons.

Parameter	Value	Unit
Valid Time	1440	minutes
Execution Timeout	30	minutes
General Batch Amount	0	devices
General Batch Interval	5	minutes
PXE Batch Amount	2	devices
PXE Batch Interval	10	minutes
FTP Batch Amount	5	devices
FTP Batch Interval	2	minutes

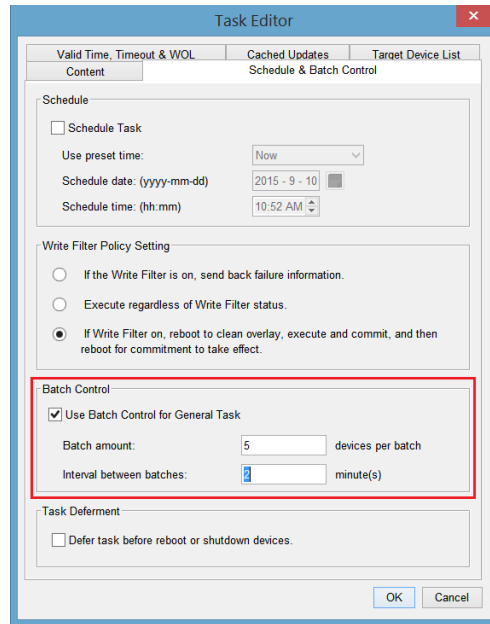
There are two parameters for each type of task: **Batch Amount** and **Batch Interval**. The batch amount is how many devices the server sends in a batch for one task. The interval is how long the server waits before sending the task to the next batch of devices.

For example, the FTP batch amount is five devices and the interval is two minutes. When a payload task is sent to twelve devices, the server sends the task to the first five devices. Two minutes later, the server sends the task to another five devices, regardless of whether the first five devices are finished. Another two minutes later, the server sends the task to the remaining two devices.

Adjust the default batch settings according to your environment. You might have to adjust the FTP server settings to get a better performance on concurrent tasks.

You can also supersede the default settings by adjusting the batch settings for each task before sending it.

Figure 7. Setting default batch amount and interval values



The Task Editor dialog box is shown with the 'Schedule & Batch Control' tab selected. The 'Batch Control' section is highlighted with a red box. It contains the following settings:

- ☒ Use Batch Control for General Task
- Batch amount: 5 devices per batch
- Interval between batches: 2 minute(s)

Other visible settings include:

- ☐ Schedule Task
- Use preset time: Now
- Schedule date: (yyyy-mm-dd) 2015 - 9 - 10
- Schedule time: (hh:mm) 10:52 AM
- Write Filter Policy Setting:
☐ If the Write Filter is on, send back failure information.
☐ Execute regardless of Write Filter status.
☒ If Write Filter on, reboot to clean overlay, execute and commit, and then reboot for commitment to take effect.
- Task Deferment:
☐ Defer task before reboot or shutdown devices.

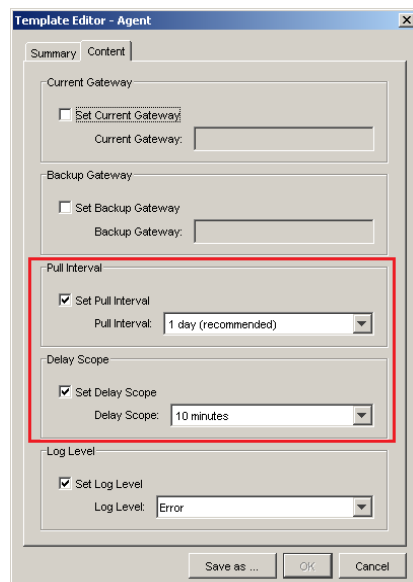
HPDM Agent settings

There are two parameters in the HPDM Agent settings that affect HPDM performance:

- **Delay Scope**—After it is started, HPDM Agent sends a startup report at a randomly chosen time between 0 and the chosen Delay Scope. If there are a large amount of devices starting within a short time slot (for example, 10,000 devices that all start within two minutes), configuring a longer Delay Scope prevents those startup reports from reaching HPDM Server in a short period of time and reduces the network pressure.
- **Pull Interval**—Pull Interval is the frequency at which HPDM Agent sends startup reports to HPDM Gateway to get tasks. There is no need to enable this function for non-NAT environments, because HPDM Gateway can connect to HPDM Agent directly to send tasks. If you are managing a large amount of devices and have to enable HPDM Agent pulling, pay attention to the value of the Pull Interval.

It takes approximately thirty minutes to for 100,000 devices to report to HPDM Server.

Figure 8. Setting the Delay Scope and Pull Interval on HPDM Agent



The Template Editor - Agent dialog box is shown with the 'Content' tab selected. The 'Pull Interval' and 'Delay Scope' sections are highlighted with a red box. It contains the following settings:

- ☒ Set Pull Interval
- Pull Interval: 1 day (recommended)
- ☒ Set Delay Scope
- Delay Scope: 10 minutes

Other visible settings include:

- ☐ Set Current Gateway
- Current Gateway:
- ☐ Set Backup Gateway
- Backup Gateway:
- ☒ Set Log Level
- Log Level: Error

Using rules

You can define rules so that HPDM Server sends tasks to the proper devices automatically. Well-defined rules can reduce not only network traffic from HPDM Console to HPDM Server, but also concurrent task reports.

There are three kinds of rules, designated by their triggers: **Scheduled**, **First Contact**, and **Startup**. Without rule tasks, finding devices that meet specific criteria and sending tasks to them can be a major part of the daily work of an administrator.

Use a Scheduled Rule to reduce the traffic from HPDM Console to HPDM Server. Sending a task to a large amount of devices from HPDM Console manually might cost minutes, or more, to initiate the task. You can define a filter that will find these devices and create a Scheduled Rule to let HPDM Server generate a target device list, which will not flood HPDM Server with a large request message from HPDM Console.

You can define a Startup Rule to make HPDM Server verify whether to send a task to a device when HPDM Server receives the device's startup report. Also, you can define a First Contact Rule to make HPDM Server verify whether to send a task to a new device when the device is registered in for the first time. This reduces the bandwidth used for task reports in environments where devices start at discrete times.

Appendix A: Use cases

HPDM is used by thousands of customers throughout the world in deployments ranging in size from 30 units to 50,000 units. HPDM is modular, which allows customers to structure HPDM to fit into their existing corporate network architecture. Like any tool, the components of HPDM can be optimally configured to make the best use of system and network performance. The use of multiple HPDM Servers is highly recommended to ensure fast testing, validation, and delivery in production environments. HPDM allows you to easily export and import created templates from one system to another to reduce the burden of using multiple HPDM Servers. Strategically place HPDM Gateways and the repositories, based on existing network architecture, to maximize the use of LAN traffic versus WAN traffic. HPDM is included in the cost of HP thin clients, so there is no additional cost or limitation on using multiple installations.

Customer 1

Industry: US retail

Units managed: 50,000 HP ThinPro

Environment: 1 HPDM Server, MS-SQL Database, 10 HPDM Gateways (regional), 1 repository per store (500+)

Usage: Asset management, patch deployment, and distributed administration—entire Help Desk (300+ admins) run through HPDM.

Customer 2

Industry: Global telecom

Units Managed: 45,000 WES, Windows CE, and HP ThinPro

Environment: 1 HPDM Server, MS-SQL Database, 13 HPDM Gateways, 13 repositories

Usage: Asset management, patch deployment, imaging and OS updates, reporting, automation and distributed administration.

Customer 3

Industry: Global banking

Units Managed: 15,000 WES

Environment: 5 HPDM Servers (one per country), MS-SQL Database, 5 HPDM Gateways, 5 repositories

Usage: Asset management, patch deployment, imaging and OS updates, reporting, automation and distributed administration. The customer creates and tests templates in the US instance. After it is validated, the template with payload is exported from the local HPDM Server and sent to remote HPDM installations. The local HPDM administrator receives the template, validates it on the local test HPDM Server, and then imports it into local production for HPDM Server to use.

Appendix B: Performance testing data

HPDM is performance optimized and was subjected to performance testing. The testing results are based on 50,000/100,000 units with a single HPDM Server, single HPDM Gateway, and database is MS SQL Server 2008. All components are distributed in different machines.

Table 1. Testing environment

Items	Sub items	Description
Device	Unit Number	50,000
HPDM Server	Environment	CPU=X5650 @ 2.67GHz 2.66GHz (2 Processors)
		Memory=8G
		Disk=150 G
Database SQL Server 2008	Environment	CPU=X5650 @ 2.67GHz 2.66GHz (2 Processors)
		Memory=8G
		Disk=80 G

Table 2. Testing results

Activity	HPDM 4.5 (50K)	HPDM 4.6 (50K)	HPDM 4.7 (50K)	HPDM 4.7 (100K)
HPDM Server Startup	> 20 minutes	20–30 seconds	20–25 seconds	40–45 seconds
HPDM Console Startup	357 seconds or error: "Cannot connect to Server"	20–40 seconds	15–25 seconds	30–40 seconds
Report In (after boot)	42 minutes	27 minutes	28 minutes	57 minutes
Sent Task	> 20 minutes or error UI hang	5 seconds	9 seconds	14 seconds
Show Task View	60 seconds	5–9 seconds	3–5 seconds	7–9 seconds
Dynamic Grouping	> 20 minutes	1–5 seconds	< 1 seconds	2–3 seconds
Manual Grouping	120 seconds	2–5 seconds	< 1 seconds	2–3 seconds
Delete Task	2–4 minutes	10 seconds	2–6 seconds	2–9 seconds

Appendix C: NAT environment

The physical network in which HPDM is deployed might be complex. When allocating components between two networks, such as a public network and private network, different NAT cases can be separated as in the following model. There is a single HPDM component within the public, or outer, network, and other HPDM components within the a single private, or inner, network. Within this model, it is assumed that without additional configuration, applications on the private, or inner, network can connect to the public, or outer, network and that applications on the public, or outer, network cannot connect to the private, or inner, network.

Figure 9. Topographical model of NAT environment

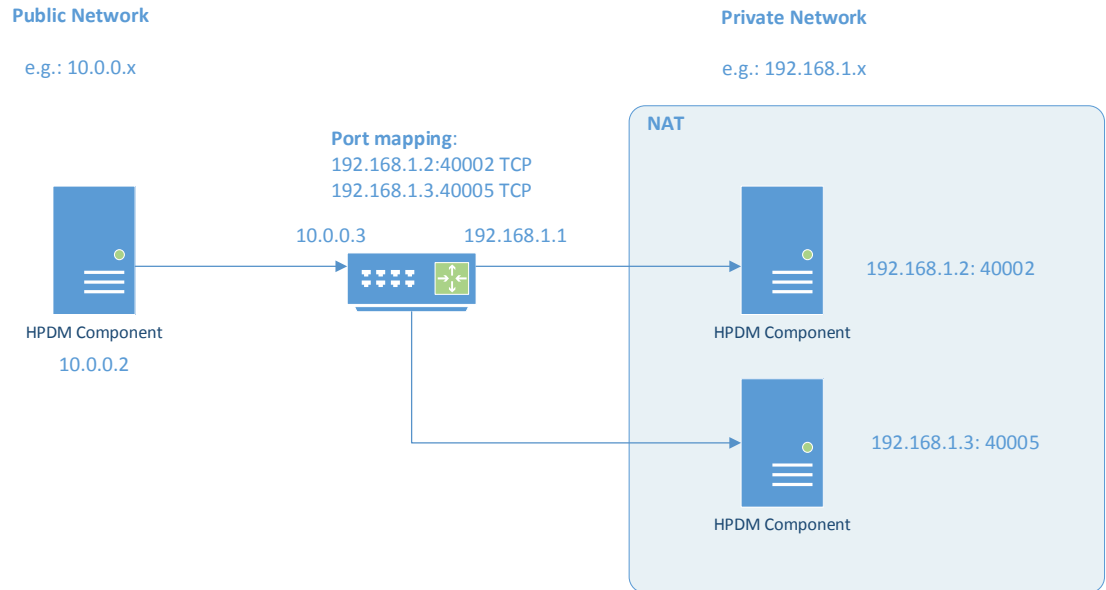


Table 3. Evaluated NAT scenarios

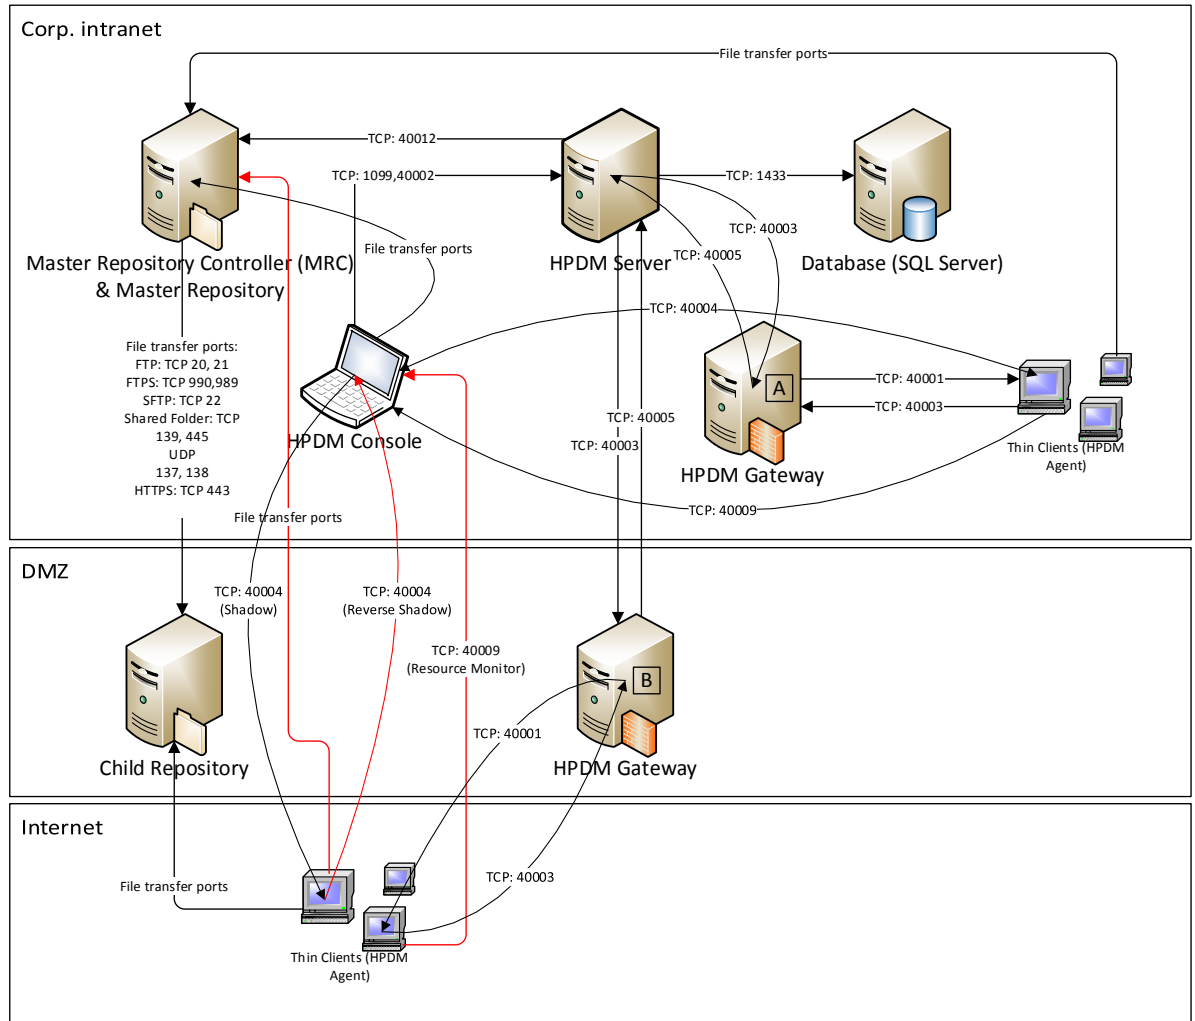
Component (public network)	Component (private network)	Port (private network)	Port type	HPDM connection status ¹		Comments
				With outer router ports mapped to inner ports ⁵	Without outer router ports mapped to inner ports	
HPDM Server	HPDM Gateway	40000	UDP	Passed	N/A ⁴	
	HPDM Gateway	40003	TCP	Passed		
	Master Repository Controller	40012	TCP	Passed	N/A	
HPDM Console	HPDM Server	1099	TCP	Failed ²	N/A	
		40002	TCP			
	HPDM Agent	40004	TCP	Failed ²		
HPDM Gateway	HPDM Agent	40001	TCP	Passed	N/A ³	
	HPDM Server	40005	TCP	Passed	N/A ³	
	PCoIP Zero Client	50000	TCP	Failed	N/A	Only PCoIP-related tasks fail.
HPDM Agent	HPDM Gateway	40000	UDP	Passed	N/A ⁴	
	HPDM Gateway	40003	TCP	Passed	N/A	
	HPDM Console (Reverse Shadow)	40004	TCP	Failed	N/A	Only Reverse Shadow tasks fail.
	HPDM Server (Resource Monitor)	40009	TCP	Failed	N/A	Only Resource Monitor tasks fail.
¹ Passes if HPDM Console can connect to HPDM Server successfully, perform operations, send tasks to devices, and update device status correctly. ² To connect HPDM Console to HPDM Server successfully, do the following on the HPDM Server side: A. Stop HPDM Server. B. Open the following file for editing: <code>\Server\conf\wrapper.conf</code> C. Add the following parameter to the file, where <IPAddress> is the outer IP address of the private network router: <code>wrapper.java.additional.2=-Djava.rmi.server.hostname=<IPAddress></code> D. Restart HPDM Server. ³ HPDM supports a poll mode in which ports 40001 and 40005 can be replaced by port 40000. ⁴ Port 40000 is for HPDM poll mode only, and can be replaced by ports 40001 and 40005. ⁵ In this setup, the router maps the inner (private network) IP address and port to the outer (public network) IP address and port. Based on these results, you can connect from a private network IP address and port from a public network.						

For more detailed information about ports used in HPDM, see the *Administrator Guide* for HP Device Manager 4.7.

Appendix D: DMZ environment

You might want to use HPDM to manage devices located in both the corporate (intranet) network and the Internet. To enable this configuration, you must place an HPDM Gateway and an HPDM repository in the corporate network's DMZ environment to manage the devices outside the corporate network.

Figure 10. Typical HPDM topology in a DMZ environment



Note:

For detailed information about HPDM ports, see the "Port reference" chapter of the *Administrator Guide* for HP Device Manager 4.7.

Figure 10 is an example of how to deploy HPDM in a DMZ environment. In this example, each component is installed on a single machine. You can install more than one component, such as HPDM Gateway A and HPDM Server, on one machine.

If you install all HPDM Server-side components in the DMZ environment it usually makes the topology simpler; however, most companies have strict security policy against this configuration.

Selecting a file transfer port

- If you are using FTP protocol, use ports 20 (in PORT mode) and 21. If you want to use PASV mode, set a port range for PASV mode in the FTP server and ensure that your firewall does not block the selected ports.
- If you are using FTPS protocol, use ports 989 (in PORT mode) and 990. If you want to use PASV mode, set a port range for PASV mode in the FTP server and ensure that your firewall does not block the selected ports.
- If you are using SFTP protocol, use port 22.
- If you are using Shared Folder protocol, use ports TCP 139 and 445 and UDP 137 and 138.
- If you are using HTTPS protocol, use port 443.

Resolving the child repository address

Both the Master Repository Controller located in the corporate intranet and the devices located within the Internet need access the child repository. However, you can only set one address for a repository location from HPDM Console. If you set an intranet address, devices in the Internet cannot access the repository. If you set an outside (Internet) address, the Master Repository Controller might not be able to access the repository.

There are two possible solutions, as follows:

- Configure the network to make sure that the Master Repository Controller can access the outside address.
- On the Master Repository Controller, modify the Windows **HOSTS** file (in %systemroot%\system32\drivers\etc\) to map the outside address (hostname or FQDN) to the internal IP address of the child repository.

For example, the child repository address is `hpdm-dmz.hp.com` and corp. network cannot access it. Add the line `192.168.10.20 hpdm-dmz.hp.com` to the **HOSTS** file of the Master Repository Controller. The Master Repository Controller can then go to 192.168.10.20 to access the child repository.

You can set the outside address as the child repository address in HPDM Console.

Using PASV mode on an FTP or FTPS server

When an FTP or FTPS server receives a PASV command, it replies with an IP address and a port using an `xx,xx,xx,xx,yy,yy` string to the FTP or FTPS client. `xx,xx,xx,xx` is the IP address and `yy,yy` is the port. Then, the client connects to `xx.xx.xx.xx:yyyy`. Both the Master Repository Controller and outside devices need access to the FTP or FTPS server. This is similar to the child repository address issue; however, the FTP or FTPS server only can be set to IP address in PASV mode.

The HPDM 4.6 SP5 or HPDM 4.7 file client library can resolve this issue. The HPDM file client does not use `xx.xx.xx.xx` in reply to a PASV command, but does use the original address for the control socket.

For example, if an HPDM file client connects to `hpdm-dmz.hp.com:21` and sends `PASV`, it receives the reply `192,168,10,20,10,01`. Then, its data socket connects to `hpdm-dmz.hp.com:2561` ($10 \times 256 + 01 = 2561$).

Limitations

In Figure 10, there are three red lines. These lines represent connections that cannot be resolved easily.

- Protocols for file transfer between HPDM Agents and the Master Repository Controller
Capture File and Capture Image tasks always capture files to the Master Repository. If the Master Repository is in the corporate intranet, outside devices cannot connect to the Master Repository.
- Using Reverse Shadow on port TCP 40004 between HPDM Agents and HPDM Console
If HPDM Console is in the corporate intranet, outside devices cannot connect to HPDM Console.
- Using Resource Monitor on port TCP 40009 between HPDM Agents and HPDM Console
- If HPDM Console is in the corporate intranet, outside devices cannot connect to HPDM Console.

If you can move all HPDM Server-side components to the DMZ environment, these issues can be resolved. However, you might not be able to use this configuration based on your security policy.

Ports between networks

Table 4. Ports between networks

Network	Peer	Direction	Type	Port
Corporate intranet	DMZ	Inbound	TCP	20, 989, 40005
		Outbound	TCP	21, 22, 139, 443, 445, 40003
			UDP	137, 138
DMZ	Corporate intranet	Inbound	TCP	21, 22, 139, 443, 445, 40003
			UDP	137, 138
		Outbound	TCP	20, 989, 40005
	Internet	Inbound	TCP	21, 22, 139, 443, 445, 40003
			UDP	137, 138
		Outbound	TCP	40001, 40004
Internet	DMZ	Inbound	TCP	40001, 40004
		Outbound	TCP	21, 22, 139, 443, 445, 40003
			UDP	137, 138

Note:

You do not need allow all file transfer ports in your firewall. For details on valid ports, see [Selecting a file transfer port](#).

Appendix E: HPDM failover redundancy

This appendix provides high-level guidance for implementing failover redundancy of HP Device Manager within F5 Network's BIG-IP infrastructure.

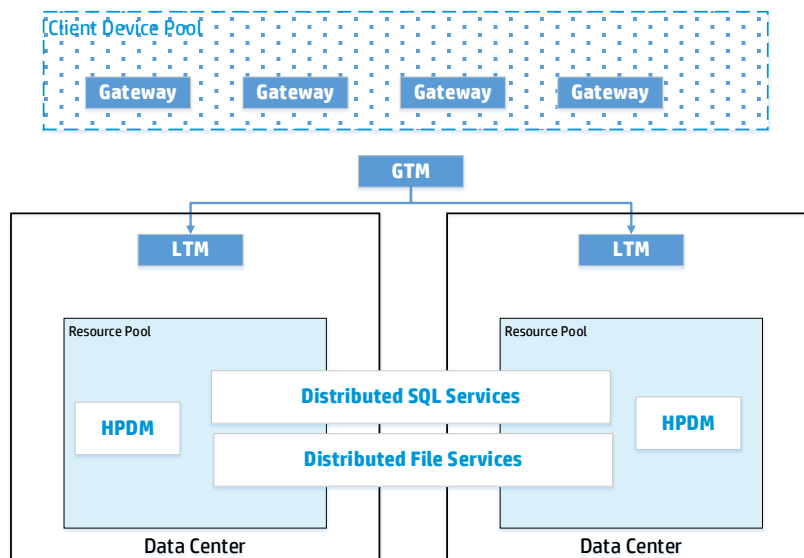
HPDM is a management solution that provides an easy-to-use interface and workflows that are streamlined for the management of HP thin clients. HPDM is a highly scalable management solution for thin clients. HPDM can scale from managing thousands of devices to over a hundred thousand devices, all on a single management server.

Scalability is not enough to cover all unforeseen failures. Diligent planning and built-in redundancy can help ensure that HPDM remains available and resilient despite infrastructure failures or other catastrophic events.

Redundancy scenario

One approach to high availability that covers a wide range of issues is the creation of redundant data centers. You can create mirrors of the resources within your network to make these resources instantly available if the origin service goes offline. This approach is consistent with the current architecture of HPDM, because HPDM is designed to have a single point of operation with gateways and repositories feeding from the single-server service layer in a hub-and-spoke pattern.

The following diagram provides a high-level depiction of an example environment with redundant data centers. The active infrastructure in this model is F5 Network's BIG-IP; however, the concepts demonstrated can also apply to other software-defined networking solutions.



In this model, F5's global traffic manager is used to mask and route traffic between data center implementations, keeping client devices unaware of the redundancy model that has been put in place. In addition, local traffic managers are used to monitor the availability of and provide access to resources within each data center. The local traffic manager does this by managing the resources available, in our case, the HP Device Manager server itself, the HP Device Manager Master Repository Controller (hosted within the same VM as our server), the SQL database service, and the distributed file system service used to store master repository content.

Implementing redundancy

Ideally, your redundant HPDM configuration looks identical to your regular configuration. If you can provide consistent addressing for SQL database services and for master repository file system services, configure the same for the VM housing HPDM Server and the Master Repository Controller agent. This VM can also be used to create local redundancy within each data center by being able to support multiple VM hosts behind Local Traffic Manager. In this example, that level of redundancy is unnecessary. Depending on your own data center configuration, Local Traffic Manager can manage both multiple VMs and multiple VM pools containing the HPDM resources.

Some key things to consider when implementing redundancy of your HPDM Server environment:

- Choose a database solution that provides availability across data centers. HPDM interoperates with a wide selection of databases that can be configured for high availability.

- Be sure that both data centers have access to the same user level; typically, this access can be configured through Active Directory replication. HPDM uses user-level permissions to communicate with the file service layer. Additionally, HPDM uses Active Directory users and groups to provide privilege access controls to various management tasks.
- Do not try to load balance HPDM traffic across multiple VMs. HPDM is not designed to operate in a coordinated effort with multiple servers either servicing the same device or leveraging the same database tables. If you have reached the scalability limits of HPDM within your installation environment, consider segmenting traffic by location and routing that traffic through your available data centers.

Local Traffic Manager

To monitor the availability of HPDM resources, add HPDM as a virtual server resource to manage within each of your data center Local Traffic Managers. Identically configure each Local Traffic Manager per data center.

Inbound traffic from HPDM Gateways to HPDM Server is on TCP port 40005. Outbound traffic to HPDM Gateways is on TCP port 40003.

Inbound connections to HPDM Server from management consoles use TCP ports 1099 and 40002.

Global Traffic Manager

To the client network, there is one and only one instance of HPDM. To implement redundancy, be sure that to the client network still perceives one and only one instance of HPDM.

This example Global Traffic Manager uses a well-known DNS entry HPDM. This address is routed to the appropriate virtual HPDM Server based on availability. Global Traffic Manager asks each Local Traffic Manager that manages virtual server resources for availability information.

In this example, one virtual resource pool (Local Traffic Manager) is designated as the primary HPDM resource. Any other Local Traffic Managers are used as failover sources, if the primary Local Traffic Manager becomes unavailable.

Summary

In this appendix, a simple scenario demonstrated where additional fault tolerance and disaster recovery support for HPDM can be achieved through data center redundancy using F5 Network's BIG-IP infrastructure. While this example does not take into account all environments or customer implementations of BIG-IP or similar software-defined networking infrastructure, it does demonstrates the basic requirements for implementing redundancy with HPDM.

For more information

To read more about HP Device Manager, go to hp.com/go/hpdm.

Sign up for updates
hp.com/go/getupdated

© Copyright 2015, 2016 HP Development Company, L.P.

ARM is a registered trademark of ARM Limited. Java is a registered trademark of Oracle and/or its affiliates. Linux is the registered trademark of Linus Torvalds in the U.S. and other countries. Microsoft, Windows, and Windows Server are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. Pentium is a trademark of Intel Corporation in the U.S. and other countries.

Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Second Edition: August 2016

First Edition: October 2015

