Technical white paper

HP Device Manager

Deployment guide

Table of contents

Executive summary

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

Hardware recommendations

The *HP Device Manager User Guide* and *Release Notes* provide the minimum requirements for server hardware.

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

Table 1. Configuration according to scale

Deployment details	Suggested minimum configuration
Small scale All HPDM components are located on the same server computer This configuration supports up to 5,000 clients	• 4 cores (Intel® Xeon® 5140 or comparable CPU) • 8 GB of RAM • 300 GB of hard disk space for the operating system, HPDM Server, HPDM Gateway, HPDM Master Repository Controller, content of Master Repository, and database files • 1000 Mbps NIC
Medium scale All HPDM components are located on the same server computer This configuration supports up to 10,000 clients	• 8 cores (Intel® Xeon® E5504 or comparable CPU) • 16 GB of RAM • 600 GB of hard disk space for the operating system, HPDM Server, HPDM Gateway, HPDM Master Repository Controller, content of Master Repository, and database files • 1000 Mbps NIC
Large scale All HPDM components except the HPDM Console are located on the site server computer This configuration supports up to 20,000 clients For managing more than 20,000 clients, see the notes following the table.	• 16 cores (Intel® Xeon® L5520 or comparable CPU) • 32 GB of RAM • 1 TB of hard disk space for the operating system, HPDM Server, HPDM Gateway, HPDM Master Repository Controller, content of Master Repository, and database files • 1000 Mbps NIC
Child Repository in a hierarchy A file server allows remote access with chosen protocols No HPDM components on this computer	• 80 GB of hard disk space to for the repository content • 1000 Mbps NIC
HPDM Gateway in a hierarchy Required for PXE Imaging	• 1 core (Pentium® III or greater) • 512 MB of RAM • 512 MB of hard disk space for the HPDM Gateway and all runtime files

NOTE

The RAM needs to be better than DDR3 1333 (PC3-10600). The hard drive needs to be faster than 7200 rpm, and RAID 5 is recommended.

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

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

Deployment recommendations

Topology

The following diagrams show the topology of both a typical and an advanced HPDM deployment.

Figure 1. Topology of typical HPDM deployment

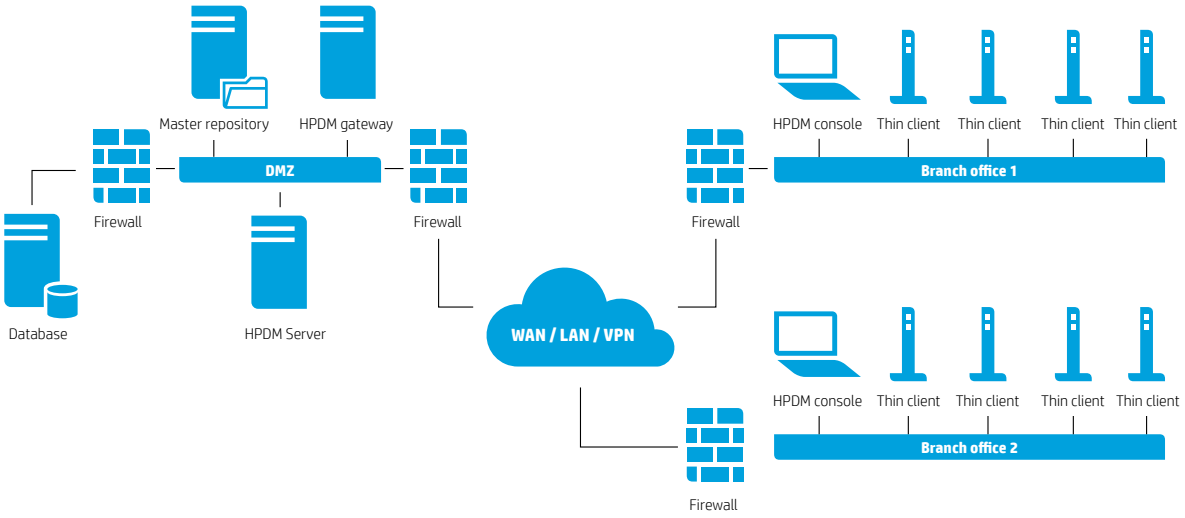

Figure 2. Topology of advanced HPDM deployment

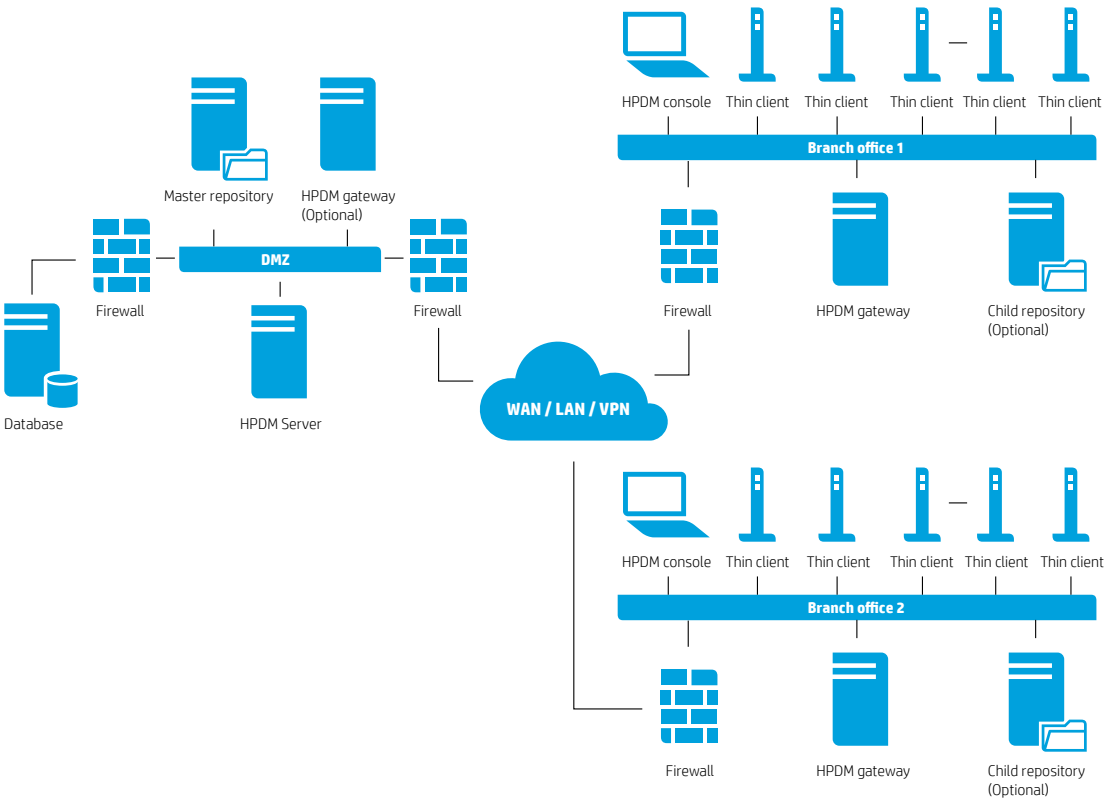

Deployment factors

- Total number of managed thin clients
  - Use more powerful, server-level hardware configuration; for example, RAID.
  - 20,000 is the recommended maximum number in one HPDM system. Lab testing shows that HPDM performs most efficiently up to this number.
  - Current customers successfully use HPDM to manage up to 50,000 devices with a single server with these considerations:
    - Use of networks faster than 1000 MB.
    - 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.
- Volume of transferred files
  - Add more Child Repositories.
- PXE imaging
  - Deploy HPDM Gateway in the same subnet as the thin clients.
  - If the subnet is a NAT subnet, configure port mapping on NAT to make sure that the HPDM Server can talk to the HPDM Gateway directly.

Advantages of additional HPDM Gateways

- Required for PXE imaging tasks.
- Consolidated communication between branch offices and DMZ.
- Accelerated task delivery speed in case of NAT Gateway in branch offices.

Advantages of additional repositories

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

Determining number of repositories

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

For example, you have 20,000 units to be re-imaged. Each image is 1 GB. Therefore, you have 20,000 GB (20 TB) of data to transfer.

Your connection from one repository to thin clients is 100 Mbps. It will take 444.4 hours to transfer all data.

20,000 GB ÷ (100 Mbps ÷ 8 bits per byte ÷ 1000 x 3600 seconds per hour) ≈ 444.4 hours

To reduce it to 48 hours, you need 10 repositories (444.4 ÷ 48). There is some overhead to synchronize from the Master Repository to the Child Repositories. Take that into consideration.

Optimizing large image deployments using Batch Control

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

The continued write speed of current thin client models 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 details about setting batch control, see the Batch tasks section later in this document.

For example, you are deploying a 4 GB image to thin clients. 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 can have the following:

1000 Mbps ÷ 8 bits per byte ÷ 3.5 MBps ≈ 35 concurrent imaging tasks in a batch

To calculate the time it will take:

4000 MB ÷ 3.5 MBps ÷ 60 sec per min ≈ 20 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 leaving a buffer is recommended. The thin client OS type makes no difference; only the image size makes a difference.

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 for each batch and a longer interval between batches to make full use of bandwidth.

Database size

The disk space usage of the database grows with the device and task amount. 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.

Database latency

HP recommends deploying the HPDM Server as close to the database server as possible because the network latency between them has significant impact on the performance. HPDM Console users will experience obvious delays if the network latency between the HPDM Server and database is higher than 30 milliseconds.

Simultaneously connected users

There is no significant increase of memory or CPU usage on the HPDM Server for any additional, connected HPDM Consoles. However, requires more resources for database access and communication.

The HPDM Server spends an additional 3 minutes 30 seconds to notify for 10,000 tasks for each addition HPDM Console. Do not to have too many connected HPDM Consoles.

Console startup

This tip speeds up the startup process of the HPDM Management Console.

Reducing the visible task amount

The HPDM Console takes time to render task summaries when initializing. Reduce the amount of tasks to view for both Manual Tasks and Rule Tasks of each OS tab before closing the HPDM Console by changing the following setting.

Figure 3. Setting number of tasks to view

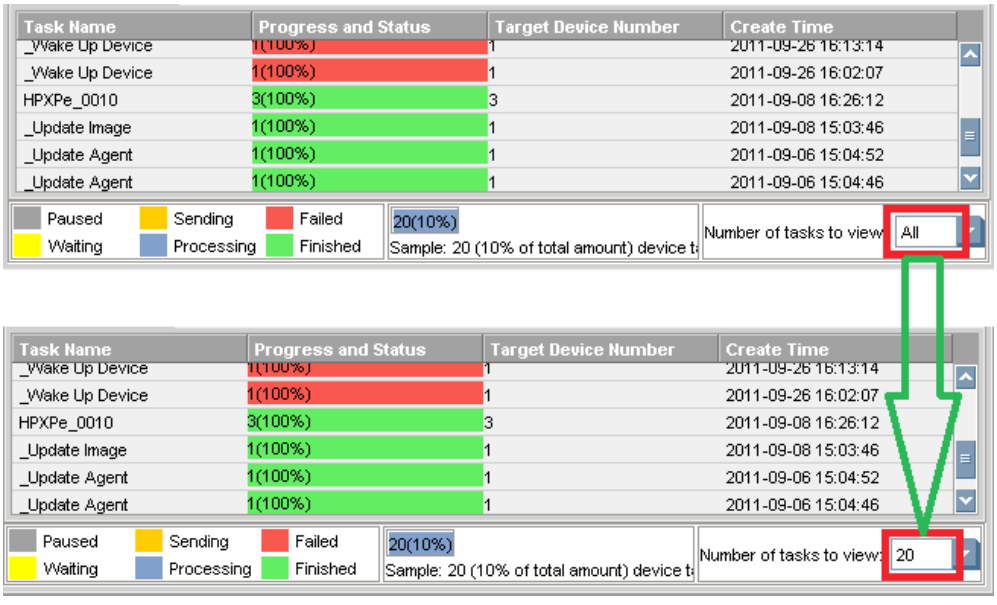

Configurations

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

Batch tasks

There are three kinds of tasks: **PXE**, **FTP**, and **General**. PXE tasks are PXE imaging tasks. FTP tasks are tasks that use file transmission, except for PXE tasks. All other tasks are general tasks. The PXE server and file repositories are bottlenecks for simultaneous task executions. Define default task batch parameters for different types of tasks in the **Configuration Management** dialog.

Figure 4. Setting default batch amount and interval values

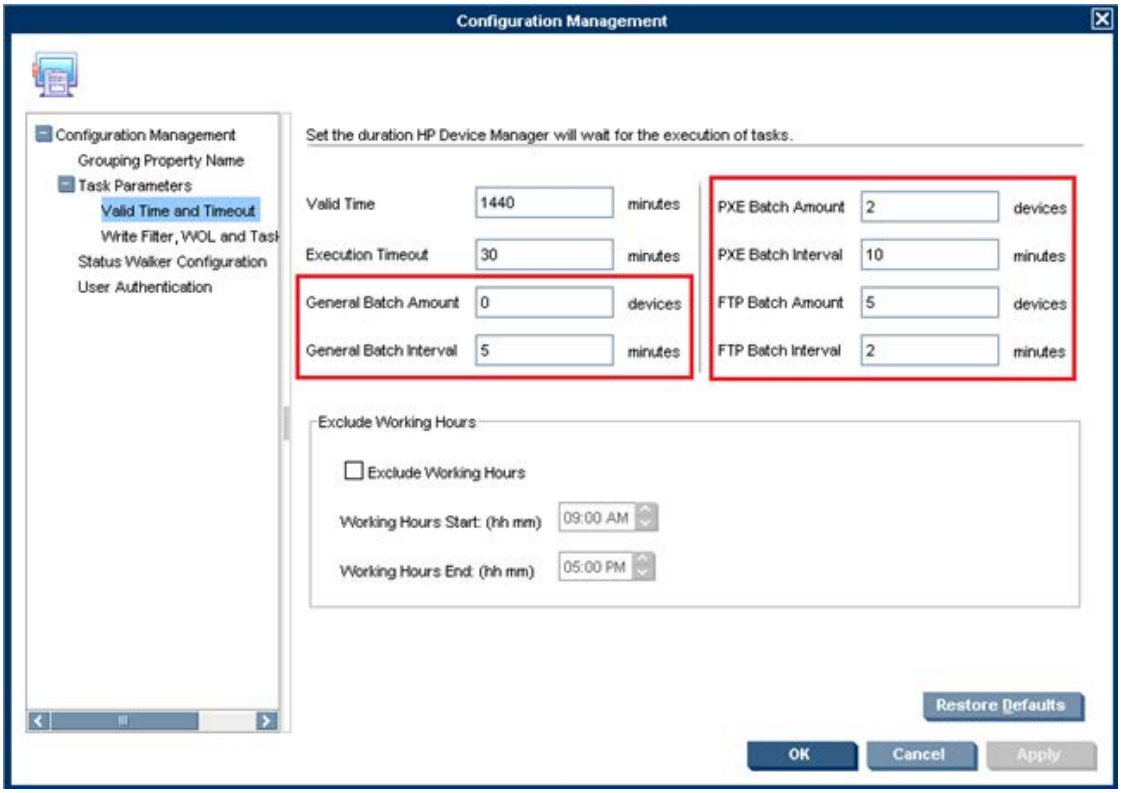

There are two parameters for each type of task: **batch amount** and **batch interval**. The amount is how many devices the server will send in a batch for one task. The interval is how long the server will wait before sending the task to the next batch of devices.

For example, the FTP batch amount is five devices and the interval is two minutes. When an FTP task is sent to 12 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. Two more 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 adjust batch settings for each task before sending them, which supersedes the default settings:

Figure 5. Setting default batch amount and interval values

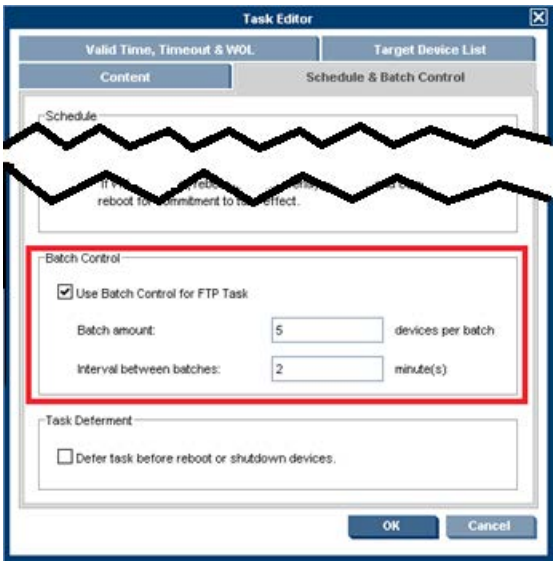

Agent settings

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

Delay Scope

The HPDM Agent sends a startup report at a randomly chosen time between 0 and the "Delay Scope" after it is started. If there are large amounts of devices starting within a short time slot (for example, 9 devices in ± 2 minutes), then configuring a larger Delay Scope prevents those startup reports from reaching the HPDM Server in a short period of time.

Pull Interval

Pull Interval is the frequency the HPDM Agent sends startup reports to the HPDM Gateway to get tasks. There is no need to enable this function for non-NAT environments, because the HPDM Gateway can connect to the HPDM Agent directly to send tasks. If you are managing large amounts of devices and have to enable HPDM Agent pulling, pay attention to the value of the Pull Interval.

For example, you have 10,000 devices with a Pull Interval set to 30 minutes, and the latency between your HPDM Server and database is about 20 ms. The server will receive 10,000 startup reports every 30 minutes. The HPDM Server needs more than 1200 ms to handle one startup report. That means the HPDM Server can handle 1,500 startup reports at most in 30 minutes. Thus, the HPDM Server is flooded with startup reports and barely responds to any HPDM Console operation.

Figure 6. Setting the Delay Scope and Pull Interval on the HPDM Agent

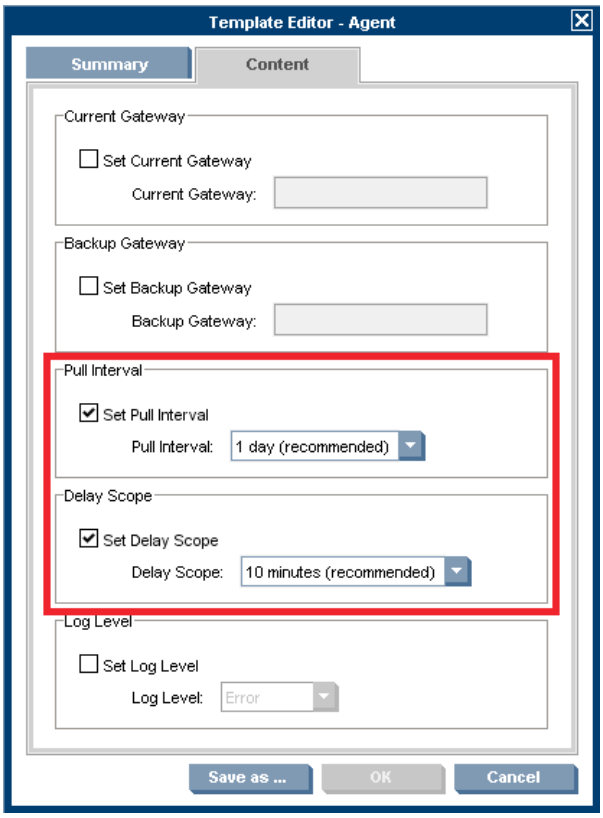

Task automation

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

Using rules

There are three kinds of rules designated by their triggers: **Scheduled**, **First Contact**, and **Startup**.

Use a Scheduled Rule to reduce the traffic from the HPDM Console to the HPDM Server. Sending a task to large amounts of devices from the 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 the HPDM Server generate the target device list instead of sending a huge list from the HPDM Console and flooding the HPDM Server with a large request message.

Without rule tasks, finding devices that meet specific criteria and sending tasks to them could be a major part of the daily work of an administrator. You can define a Startup Rule to make the HPDM Server check whether to send a task to the device when it receives its startup report. Also, you can define a First Contact Rule to make the HPDM Server check whether to send a task to a new device when it is registered in for the first time. This reduces the task report bandwidth occupation for environments where devices start at discrete times.

Backup and restore

Use the HPDM Server Backup and Restore Tool to backup database and HPDM Server files. You can run this backup tool periodically, such as once a month. See the *HP Device Manager User Guide* for detailed steps.

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 to allow 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 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 the HPDM Gateways and repositories, based on existing network architecture, to maximize the use of LAN traffic vs. 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. Once 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 the HPDM Server to use.

Summary

The information in this document provides an idea of how to plan your installation of HPDM. The recommendations given will help to optimize the performance of payload and imaging tasks and the performance of the HPDM Console. The configuration recommendations will also help reduce the amount network traffic generated by HPDM.

For more information

Learn more at
hp.com/go/hpdm

Sign up for updates
hp.com/go/getupdated

Share with colleagues

Rate this document

4AA4-7890ENW, August 2013