HP Enterprise Discovery

for the Windows® operating system

Software Version: 2.50

Planning Guide

Manufacturing Part Number: None
Document Release Date: October 2007
Software Release Date: October 2007

Legal Notices

Warranty

Restricted Rights Legend

Copyright Notices

Trademark Notices

Support

You can visit the HP Software Support web site at:

www.hp.com/go/hpsoftwaresupport

HP Software online support provides an efficient way to access interactive technical support tools. As a valued support customer, you can benefit by using the support site to:

- Search for knowledge documents of interest
- Submit and track support cases and enhancement requests
- Download software patches
- Manage support contracts
- Look up HP support contacts
- Review information about available services
- Enter into discussions with other software customers
- Research and register for software training

Most of the support areas require that you register as an HP Passport user and sign in. Many also require an active support contract. To find more information about support access levels, go to the following URL:

http://h20230.www2.hp.com/new_access_levels.jsp

To register for an HP Passport ID, go to the following URL:

http://h20229.www2.hp.com/passport-registration.html

Contents

1 An Overview of Enterprise Discovery

Enterprise Discovery^{TM} collects large amounts of data about your network and devices. It can discover devices on its own by working its way through a collection of IP addresses that you provide and can also collect detailed data from devices using configurable Scanners.

From a data-gathering perspective, Enterprise Discovery performs three distinct functions:

- Discovery
- Inventory
- Software Utilization

The following diagram shows the type of data collected by each of these functions.

Device	
Discovery	• Type • IP/MAC • Ports • Name
Inventory	• Disk Configuration • Asset Tag • Monitor Information • Services • Software Licenses
Software Utilization	• Usage information

Further Information

You can find a detailed technical description of how Enterprise Discovery works in the *Reference Guide*.

Discovery

As a starting point, Enterprise Discovery needs to determine what devices are in your network and gather basic information about each of them. This process is referred to as Discovery and allows you to get a good overview of the number and types of devices in your network, as well as a basic set of attributes for each. It also servers as the foundation for the other modules of Enterprise Discovery.

Discovery is based on collections of IP addresses call IP-only device groups. The conditions that define IP-only device groups can be specified using IP ranges, subnets, IP wildcard strings, or individual IP addresses.

For each IP-only device group in your network, Enterprise Discovery can use a variety of methods to discover devices, allowing you to use choose the appropriate settings for different groups. For example, UNIX servers in the data centre may have different requirements for Discovery than laptops in the Finance group.

For an initial discovery setup, it is normally sufficient to define a small set of large IP ranges and then fine-tune the setup later as you discover groups of devices that need to be treated differently in some way.

To deal effectively with very large networks, more than one Enterprise Discovery server can be deployed in an organization, typically organized by geographical location. An additional server can then be designated to aggregate the results from all of the other servers through a process called Aggregation. It would then be possible for example, to run the results from each Enterprise Discovery Server to one central repository database (such as AssetCenter).

Inventory

After discovering a device, Enterprise Discovery can run a Scanner on it to gather detailed hardware, configuration and software license information. This process is referred to as Inventory and makes it possible to drive standardization and compliance initiatives, manage risk, implement chargeback policies, etc.

The Scanners can be launched automatically according to a configurable schedule, allowing complete control over network bandwidth usage and any impact on the end-user.

In order to manage the Scanners, the HP Enterprise Discovery Agent needs to be in place. This is a small program that runs all the time and deals with security and communications. The Agent can be automatically deployed to Windows machines in your network, and must be manually deployed to UNIX machines. Once this is done, Enterprise Discovery can automatically upgrade the Scanners and agents when necessary.

Enterprise Discovery includes Agents and Scanners for most common desktop and server operating systems.

Software Utilization

Enterprise Discovery can gather information about what software is used on the machines on your network. This is referred to as Software Utilization and the information collected is necessary to optimize software license cost, for example by eliminating unused or under-utilized software installations.

The Software Asset Management module of AssetCenter 4.4 or later is ideal for performing the analysis of the data collected by Enterprise Discovery.

2 Preparing Your Network for Discovery

There are several steps you can take to prepare your network for using Enterprise Discovery.

- Turn on SNMP management in all routers and core switches on page 13
- (Optional) Turn on SNMP management in other devices on page 14
- Set DHCP lease time on page 14
- Turn on SNMP management on the DHCP server. on page 14
- About SNMP Configuration on page 14
- Give the Enterprise Discovery server's IP address to all devices using directed community strings on page 15
- (Optional) Adjust bridge aging on page 15
- Plan the device and port to which the Enterprise Discovery server will be attached on page 15
- Server Ports on page 16
- Check Cisco devices on page 18

Turn on SNMP management in all routers and core switches

Depending on the device, this may be a case of enabling an existing SNMP agent or setting up an SNMP agent.

You may also turn on SNMP management in other devices. The more managed devices in your network, the better. However, enable switches and routers first.

➤ If you use HSRP (Hot Standby Routing Protocol) in your network, ensure you turn on SNMP management in all the affected devices.

What if you don't turn on SNMP management in your switches and routers?

- Enterprise Discovery will appear to work, but you'll eventually notice that it is working poorly. Once Enterprise Discovery is up and running, the Exceptions reports can advise you of problems. Much of the information that Enterprise Discovery collects comes from the SNMP MIB of devices in your network, so it is crucial that you enable SNMP management.

How do you turn on SNMP management?

- The exact procedure is different for every device. Consult the documentation that came with your switch or router.

  ➤ When you turn on SNMP management in a device, you often assign a community string (for SNMPv1/v2) or a user (for SNMPv3). If you assign a new string later, be sure you give the community string/user information to Enterprise Discovery. For more information, see About SNMP Configuration on page 14.

(Optional) Turn on SNMP management in other devices

Your decision to turn on SNMP management in your remaining switches, hubs, servers and workstations depends on the results you expect from Enterprise Discovery. For example, in many networks, monitoring the performance of workstations is not important.

Set DHCP lease time

If you use DHCP (Dynamic Host Configuration Protocol) in your network, set the IP address lease time to at least 7 days.

Turn on SNMP management on the DHCP server.

Enable SNMP management on the DHCP server so that Enterprise Discovery can poll the DHCP server ARP cache for the current IP and MAC address pair information of the devices on your network, allowing Enterprise Discovery to keep track of IP addresses as they are reassigned.

About SNMP Configuration

A community string (SNMPv1/v2) or a user (SNMPv3) is like a password. A device uses a community string/user to protect its SNMP MIB—and it's the data from the SNMP MIB that Enterprise Discovery relies on. Enterprise Discovery must know at least one of a device's passwords to collect data from that device. If you do not give Enterprise Discovery a device's community string/user, Enterprise Discovery will behave as though the device does not have SNMP management turned on. Enterprise Discovery will appear to work, but you'll eventually notice that it is working poorly. Once Enterprise Discovery is up and running, the Exceptions reports can advise you of problems.

➤ Community strings are case-sensitive. “Public” and “public” are two different strings.

Directed community strings

Directed community strings give devices another layer of protection: a list of IP addresses of approved devices. When Enterprise Discovery tries to get information from a device with a directed community string, the device asks not only “What’s the password?” but also “Are you on the list?”

Give the Enterprise Discovery server’s IP address to all devices using directed community strings

When directed community strings are used, it is not enough to give Enterprise Discovery access to the device. You must also configure the device to recognize the Enterprise Discovery server. You must put it on the list of approved devices.

What happens if a device with directed community strings is not configured with the IP address of the Enterprise Discovery server?

1. Enterprise Discovery will behave as though the device does not have SNMP management turned on. Enterprise Discovery will appear to work, but you’ll eventually notice that it is working poorly. Once Enterprise Discovery is up and running, the Exceptions reports can advise you of problems.

(Optional) Adjust bridge aging

To improve the reliability and speed of Enterprise Discovery, adjust bridge aging on your bridges, routers, switches, and concentrators. Turn bridge aging on, and set the bridge aging interval to 2-6 hours. Smaller networks can use shorter intervals; larger networks will need longer intervals. Enterprise Discovery’s Exceptions reports can tell you which devices should have their bridge aging adjusted.

Plan the device and port to which the Enterprise Discovery server will be attached

Plan to attach the server:

- behind your corporate firewall
- to an Ethernet port on a device close to the top of your network. Enterprise Discovery works best if the port is SNMP managed.

Server Ports

Firewall Ports

Enabling these firewall ports will allow Enterprise Discovery system to perform through a corporate firewall.

If you have a corporate firewall that could impede Enterprise Discovery, configure the corporate firewall to allow ICMP (ping) to pass through, and enable the following ports:

Table 1 Firewall Ports

Used for	Port	Note	From	To
Echo Reply	0/icmp		device	Enterprise Discovery server
Error Messages	3/icmp		device	Enterprise Discovery server
Echo Request	8/icmp		Enterprise Discovery server	device
TTL Timeout	11/icmp	4	Enterprise Discovery server	device
			device	Enterprise Discovery server
Netmask Request	17/icmp		Enterprise Discovery server	device
Netmask Reply	18/icmp		device	Enterprise Discovery server
SMTP	25/tcp		Enterprise Discovery server	SMTP server
DNS	53/udp		Enterprise Discovery server	DNS server
NetBIOS-n (name server)	137/udp		Enterprise Discovery server	device
NetBIOS-dgm (datagram)	138/udp		management workstation	Enterprise Discovery server
NetBIOS-ssn (session—file and printer sharing)	139/tcp		management workstation	Enterprise Discovery server
SNMP	161/udp		Enterprise Discovery server	device
SNMP traps	162/udp	3	Enterprise Discovery server	external network management server

Table 1 Firewall Ports

HTTPS	443/tcp		management workstation	Enterprise Discovery server
		1	management workstation	device
		1	Enterprise Discovery server	device
		2	Enterprise Discovery remote server	Enterprise Discovery aggregator server
Windows communication (Windows 200x and XP)	445/tcp		Enterprise Discovery server	Windows device where the Enterprise Discovery Agent will be deployed using Windows RPC
HP Enterprise Discovery Agent	2738/tcp		Enterprise Discovery server	device with Enterprise Discovery Agent
Remote Desktop	3389/udp		HP Software Customer Support	Enterprise Discovery server
MySQL ODBC	8108/tcp	1	management workstation	Enterprise Discovery server
Traceroute	33263/udp 33436/udp		Enterprise Discovery server	device

1. Depending on your settings for Server proxy services
2. If you have an Aggregator license
3. If you are using SNMP trap notification
4. TTL Timeout can go in either direction, from the Enterprise Discovery server or to the Enterprise Discovery server.

Other Ports used by the Server

These are additional ports used on the Enterprise Discovery server that do not need to be enabled in the firewall.

Table 2 Other Ports

Service	Port
MIB Browser	8100
Network Map	8101
authd	8109
acs	8110
Java loggers	8112

Table 2 Other Ports

Service	Port
C++/delphi loggers	8113
scheduler	8114
Perl/text based loggers	8115
Tomcat	8116
Apache	8117
Event Manager	8118
Topology Converter	8119
DiscoveryEngine.jay_polls0	8200
DiscoveryEngine.jay_polls1	8201
DiscoveryEngine.jay_polls2	8202
DiscoveryEngine.jay_polls3	8203
DiscoveryEngine.jay_polls4	8204
DiscoveryEngine.jay_idd	8196
DiscoveryEngine.jay_tables	8197
DiscoveryEngine.jay_jaywalks	8198
DiscoveryEngine.jay_debug	8199

Check Cisco devices

It is strongly recommended that firmware/software in your Cisco devices be IOS version 12 or higher. If you want ATM or Frame Relay support, IOS 12 is mandatory in your Cisco devices.

Check Committed Information Rate values

If your network uses Frame Relay, check your Committed Information Rate (CIR) values for your connectivity devices. Make sure you set the CIR on these connections in the Port Manager, so the correct statistics will be calculated.

In Frame Relay networks, a CIR is a bandwidth (expressed in bits per second) associated with a logical connection in a permanent virtual circuit (PVC).

The CIR values for these devices are available from your service provider. Check the appropriate documentation to obtain these values.

If the network activity on any particular PVC goes over normal operating thresholds, you should be aware that the Frame Relay controller may mark some packets to be deleted.

3 Planning Form

This chapter contains a preformatted example of an Enterprise Discovery planning form. Use this form as a starting point and customize it as needed to suit your organization's discovery, inventory, and software utilization needs. You will find information on the following topics:

- The Planning Form Overview on page 19
- Instructions for completing the Planning Form on page 19
- The Enterprise Discovery Planning Form on page 25

The Planning Form Overview

To ensure that Enterprise Discovery knows where to collect data from and how to collect that data, you must do a little preliminary work. You only have to do this once.

By using the planning form in this chapter before implementing Enterprise Discovery you will:

- Ensure the network is ready and prepared for Discovery.
- Determine what Inventory data is to be collected and how it will be used.
- Determine the user asset information that will be recorded.
- Extract the information necessary to plan the logistics of an inventory.

Instructions for completing the Planning Form

This planning form is the first step in defining the requirements for a Discovery and Asset Inventory project. Depending on how the project is to be implemented, further requirements will need to be defined to deal with detailed logistics. These may include, for example, site access for engineers, security clearance, and other miscellaneous issues.

➤ If you want, you can fill in the questionnaire and send it to HP Software customer support. They can review your information and provide feedback on how you set up Enterprise Discovery.

If you need help filling out the questionnaire, contact your HP Software or OEM/VAR (Original Equipment Manufacturer or Value Added Reseller) sales representative.

Discovery

Table 1 Summary of form sections related to discovery

Section	**General instructions**
Network node and subnet setup	Enter information to help determine the scale of your network. Enterprise Discovery defines a node as any network device with at least one MAC address. A managed device is a network device that has an SNMP agent and MIB so it can respond to SNMP requests.
Enterprise Discovery Server network information	Enter the information that you will assign to the Enterprise Discovery Server at startup. If your network uses DHCP, ensure that the IP address for the Enterprise Discovery Server is static.
List IP addresses for Enterprise Discovery to discover	Enterprise Discovery uses IP-only device groups to discover the devices in your network. The discovery process works best when you give it a broad idea of where the devices in your network are—but exclude places where you know there are no devices. While you are making a list of devices in your networks, indicate bridges, routers, switches, and concentrators, so that you can identify them easily. Please list the IP addresses that you want Enterprise Discovery to discover in your network. For example, to discover an entire class C subnet with subnet mask 255.255.255.0 enter an IP range from xxx.xxx.xxx.0 to xxx.xxx.xxx.255 such as 172.17.1.0. to 172.17.1.255. If you require more space, please attach additional sheets as needed. In addition to IP ranges, you can now specify individual IP addresses, subnets, or IP wildcard strings when you create device groups.
List IP addresses for Enterprise Discovery to avoid	If there are subsets of the above IP addresses that you do not want Enterprise Discovery to discover, enter them here.

Table 1 Summary of form sections related to discovery

List the network device community strings (SNMPv1/ v2) and users (SNMPv3)	For an explanation of community strings/users, see the *Installation and Initial Setup Guide*. This is a list of non-directed community strings. Directed Community strings are covered later. Does Enterprise Discovery need to know the Write Community String/User? No. Enterprise Discovery will operate without write strings/users. However, if you do give Enterprise Discovery the write strings/ users, the owner of an Administrator account will be able to manage the device from the Enterprise Discovery interface.
TCP/IP configuration	The Enterprise Discovery server must have its own static IP address, but it can manage devices with either static or dynamic IP addresses. Please enter the following information to show how the devices on your network receive IP addresses.
Unmanaged routers	List the IP addresses of any routers you want Enterprise Discovery to monitor that do not have SNMP management enabled now or will have it in the future. For example, this might be a router controlled by an Internet Service Provider.

Inventory and Software Utilization

Table 2 Summary of form sections related to inventory and software utilization

Section	General instructions
Project staff details	On projects such as these, there are several individuals involved with the core implementation. List those individuals, their roles, and responsibilities as they relate to the project.
Reasons for data collection	Pinpoint the business drivers for conducting the Inventory. Include as little or as much information as you want. This can be used to list the information you seek to gain from the project
Project timing	An inventory project is not an instantaneous event. Sufficient time needs to be given to the development of a deployment plan as well as testing. If the data is required for other parts of a larger project, the sooner the work is started the better. Inventory should progress in parallel with the rest of the project, rather than being left as an afterthought. The business driver should be specified so that interfacing to other parts of a project can be evaluated to see if the time-scale is realistic.

Table 2 Summary of form sections related to inventory and software utilization

Current environment and existing software deployment	What is currently available for asset identification and for deploying software such as Scanners? It is important to identify if any existing facilities are available that can be used in the new project. Unique asset identification is crucial to differentiate machines. If nothing exists, then time and effort needs to be given to considering how this identification is achieved.
Site information and equipment	Knowing the split of locations and the estimate of machines at each site is necessary to establish how much work is required at each site and the amount of space required for storage of the data. This leads to discussion of whether all the data is stored locally at each site and uploaded in bulk or push all the data back to one central repository.
Data storage estimates and speed of network connections	Knowing how often the scan is run will identify traffic flow and when the transfer activities will take place. If there are any current network performance issues, these need to be understood before additional activity is added to the current load.

Table 2 Summary of form sections related to inventory and software utilization

Scanner configuration Hardware and software	The Scanner can be configured to select various combinations of Hardware, Software and Assets. For example, the first scan of an asset may only be concerned with Asset Information and Hardware details. Whereas, some machines may only require a Software scan on subsequent rescans. • Hardware Normally, the default selection of all hardware tests data is sufficient. The tests take very little time to run and unless there is a known problem it is best to leave the settings as they are. • Software The software choices determine how many and what types of files to both scan and store in the resulting scan file. Exclusions may be because of known scanning problems with a particular file. To assist with software recognition, it is useful to have at least one sample set with signatures. The file signature is a calculation on the first 8K bytes of a file. This requires the file to be opened. If the file is opened, then it is also possible to extract the version information (same as properties under Windows). This header information can provide vendor and application details. The Scanner has the ability to store the contents of ZIP files as directories. This allows the display of the file names that have been compressed, but it cannot extract version information or open the individual compressed files. • Asset Information Asset information is data that has been extracted from files and/or Windows registry. The asset data can be read from a previous scan file and loaded into the fields. • Stored Files These files can be embedded in the scan file. They are usually configuration or other data files.
File contents to be collected by the scanner	List files to be collected in addition to the default files.
Software recognition	List the custom software that is to be recognized, as well as the number of applications.
Scan Administration	Checking whether there is any anti-virus or other security software running is important because it can have a profound impact on the speed of the scan. If a file has to be opened for signature, a real-time virus scanner would intercept the request and check the file before releasing it for scanning. To avoid opening a file it may be necessary to configure the scanner to ignore file information. In this instance directory information about the file would still be captured – file name, size, attributes etc. but version information would not be captured.

Computer Population Details

Table 3 Summary of form sections related to population details

Section	General instructions
Total population	Enter the total number of workstations and servers that are to be included in the inventory project. Also include whether this number is believed to be accurate to within 5%, 10%, or greater.
Percentage Networked	Indicate the percentage of machines that are, or are likely to be network connected during the inventory.
Number of Laptops	Indicate the approximate number of laptop within the overall population. Often special arrangements need to be made to ensure that these are on site during the period of the data collection.
Operating Systems (% of populations)	The scanning software can be configured for several operating systems. Indicate the percentage of differing operating systems. From this you can determine the number of different Scanners likely to be needed. See **Help** > **Compatibility Matrix** in the Enterprise Discovery GUI for a complete list of supported operating systems.
Network Types	Include the network operating software and speed of network if available.
Policy on passwords	Indicate whether power on (boot-up) and/or screensaver passwords are used. If possible, consider not only company policy but the likelihood of individual departments/persons making use of them.

Manual Inventories

Table 4 Summary of form sections related to manual inventories

Section	General instructions
Special users • Contractors • Remote • Standalone	• Special users, remote or standalone need to be considered along with any non-working equipment. • You may decide to leave these out or perform the work manually. • Consider whether there might be physical constraints when gaining access to the sites. Are parking facilities available? • Are passes needed for access to locations within the sites. Are swipe cards used? Will they be made available to your project team members?
Non-operational equipment	Indicate if this equipment should be scanned or if data will be captured manually, as well as other considerations.

Table 4 Summary of form sections related to manual inventories

Other miscellaneous considerations - Security Software Installed - Floppy disk locks in place - Asset Labels - Scan schedule times - Access rights to machines	- Indicate whether security software is installed which would prevent an executable file from being run from the floppy drive. If such software is installed, identify. - Indicate whether floppy disk drive locks are fitted. Also consider whether any of the floppy drives might be either disabled or not-connected. - Each scan file needs to be uniquely identified during the data collection. The easiest and most efficient way of doing this is to use an asset number assigned to that equipment. If you do not have a system of asset labelling in place you might like to consider introducing one during the data capture, as the labels can be fitted as each PC is visited. - Check access times for your team. You can often expect that access can only be granted to certain locations/departments outside of these times. This can assist with estimating the time taken to complete each site and hence the overall project. - In order to fully scan some machines, local administrator access rights to the machine are required.

The Enterprise Discovery Planning Form

Client Details and Contacts

Table 5 Client Details and Contacts

Client name	
Address	
Phone	
Project Contact	Name: Phone: Email:
Technical Contact	Name: Phone: Email:

Network node and subnet setup

Table 6 Network node and subnet setup

How many nodes do you believe are active on your network?	
Are there any remote sites to be managed?	Yes □ No □
If yes, approximately how many managed nodes are at remote sites?	
Is your network divided into subnets?	Yes □ No □
If yes, how many subnets does your network contain?	

Enterprise Discovery Server network information

➤ You will give this IPv4 address to new Enterprise Discovery users so they can log in easily.

Table 7 Enterprise Discovery Server network information

Planned IPv4 address for your Enterprise Discovery server	
Subnet mask address	
Default gateway IP address	

List IP addresses for Enterprise Discovery to discover

You can specify a collection of IP addresses by using any of the following formats:

- Single IP—Specify an individual IP address as a string of 4 dotted octets with values 0-255. The first octet must not be 0 (for example, 192.168.56.237).
- IP Range—Specify starting and ending IP addresses to represent a range. All octets may be 0, but wildcards are not permitted. The starting address must be less than or equal to the ending address when considering them as 32 bit integers. (for example, 192.168.56.0-192.168.56.255)
- Subnet—Specify a single IP address and subnet mask. The mask is a 32 bit string of binary 1s and 0s represented as either a dotted octet or a decimal number that indicates the number of binary 1s. For example, the masks 255.255.254.0 and 23 are equivalent. 0 or 0.0.0.0 is not valid as a mask.
- IP Wildcard string—Specify a string of 4 dotted octets with values 0-255. The first octet may not be 0. Octets 3 and/or 4 may be asterisks. (for example, 192.168.56.*).

Be aware of the size of the device groups you will be defining. If you request a large collection of IP addresses to sweep, it can take several hours or days.

Table 8 List of IP addresses for Enterprise Discovery to discover

Device Group Name	IP Address Format	IP Information
Group 1		
Group 2		
Group 3		
Group 4		
Group 5		
Group 6		

List IP addresses for Enterprise Discovery to avoid

You do not need to enter addresses outside the groups you have specified. Enterprise Discovery does not discover IP addresses unless you specify them.

Table 9 List of IP addresses for Enterprise Discovery to avoid

Device Group Name	IP Address Format	IP Information
Group 1		
Group 2		
Group 3		
Group 4		
Group 5		
Group 6		

List the SNMPv1/v2 Community Strings for your network devices

Table 10 List the Community Strings of your network devices

		Rights granted	
Community string	**Associated device or group**	**Read**	**Write**
		□	□
		□	□
		□	□
		□	□
		□	□
		□	□

➤ These should be made available at installation time.

List the SNMPv3 Users for your network devices

Table 11 List the Community Strings of your network devices

				Rights granted	
User Name	**Associated device or group**	**Authentication Algorithm and Pass phrase**	**Encryption Algorithm and Pass phrase**	**Read**	**Write**
				□	□
				□	□
				□	□
				□	□
				□	□
				□	□

➤ These should be made available at installation time.

TCP/IP configuration

Table 12 TCP/IP configuration

Are TCP/IP addresses static or dynamic?	Static □ Dynamic □
If dynamic, enter the following:	
— The IPv4 address(es) of Dynamic Host Configuration Protocol (DHCP) server(s)	
— The DHCP IPv4 address lease time (recommended lease time of at least 7 days.)	
Is SNMP management enabled on the DHCP server?	Yes □ No □
Enable SNMP management on the DHCP server so that Enterprise Discovery can poll the DHCP server ARP cache for the current IP and MAC address pair information of the devices on your network.	

Unmanaged Routers

Table 13 Unmanaged routers

Unmanaged router number 1	
Unmanaged router number 2	
Unmanaged router number 3	

Project Staff Details

Table 14 Project details

Resource	Role	Responsibilities
Name: Phone: Email:	Project manager	
Name: Phone: Email:	Project lead	
Name: Phone: Email:	Security lead	
Name: Phone: Email:	Network lead	
Name: Phone: Email:	Asset management lead	

Reasons for Data Collection

Table 15 Reasons for Data Collection

1	
2	
3	
4	
5	

Project Timing

Table 16 Project Timing

Desired start date for project	
Target implementation date (move to Production) Business driver for target date	

The Current Environment

Table 17 The current environment

Is there an existing asset control system in place with unique identifiers?	Yes □ No □
If yes, is this information stored electronically?	
If no, what identifier is to be used?	

Table 18 Existing software deployment

Is there an existing software deployment system in place?	Yes □ No □
If yes, what system?	
If no, how are applications distributed currently?	

Site Information

Table 19 Site information

Locations	
1	
2	
3	
4	
5	

Table 20 Equipment

Locations	**1**	**2**	**3**	**4**	**5**	**Total**
Workstations						
Windows PCs						
Solaris workstations						
HP/UX workstations						
Linux Kernel						
AIX						
Mac OS						
Servers						

Table 20 Equipment

NT Servers						
Netware Servers						
Site totals						

Data Storage Estimates

If sample files are available from previous scans, use their average size as an indicator.

Table 21 Data storage estimates

Locations	**1**	**2**	**3**	**4**	**5**	**Total**
Workstations						
Mbytes/ Workstation						
Servers						
Mbytes/Server						
Site totals Mbytes						

Table 22 Anti Virus software and high speed network connections

Any Anti Virus/ Security products which may impact audit?	Yes □ No □
Are there high speed network connections between sites?	Yes □ No □

Scanner Configuration

For Automated scanning the Hardware detection is always included in the Scanner.

Table 23 Scanner configuration

Hardware only	Yes □ No □
Hardware, Assets	Yes □ No □
Hardware, Assets, S/W	Yes □ No □
Software only	Yes □ No □

Table 24 Hardware Scanner configuration

Hardware	
Use default setting of all tests? If No, which tests to remove?	Yes □ No □

Table 25 Software Scanner configuration

Software	
All file data to be stored, only Executable files, Selected files If Selected, which file types? Any specific exclusions? If yes, detail files and/or directories	All □ Executables □ Selected □ _____; _____; _____; _____; Yes □ No □
Will file scan include signatures?	Yes □ No □
If Yes, will version information be extracted?	Yes □ No □
Will ZIP files be stored as directories?	Yes □ No □

File Contents to be Collected by the Scanner

Default files collected

Config.sys, Sms.ini, Drvspace.ini, Autoexec.bat, System.ini, Win.ini, Boot.ini, Infrtool.ini, Exclude.fp, Net.cfg, Protocol.ini

For UNIX

fstab, group, hosts, inetd.conf, inittab, profile

Table 26 Files to be collected by the Scanner

File	Location	String for file extract

Software Recognition

Table 27 Software recognition

Custom software to be recognized?	Yes □ No □
Number of applications	
Note: Each application may have many releases. Ensure that you note these.	

Scan Administration

Table 28 Scan administration

What systems will be fed with scan data e.g. AssetCenter, ServiceCenter using [Connect-It]?	

Computer Population Details

Table 29 Computer population details

Total population	Total number of work stations: _ Percentage of accuracy: _ Total number of servers: _ Percentage of accuracy: _ Total number of laptops: _ Percentage of accuracy: _
Percentage Networked	Percentage currently networked: Percentage likely to be networked:
Operating Systems	Percentage of Windows: _ 32 bit: _ 64 bit: _ Percentage of UNIX: _ 32 bit: _ 64 bit: _ Percentage of Linux: _ 32 bit: _ 64 bit: _ Percentage of Mac OS: _ 32 bit: _ 64 bit: _ See **Help** > **Compatibility Matrix** in the Enterprise Discovery GUI for a complete list of supported operating systems.
Network Types	Networking operating software: Speed of Network:
Are power on passwords used?	Yes □ No □
Are screensaver passwords used?	Yes □ No □

Manual Inventories

Table 30 **Special users**

Contractors/Special Users	
Are they to be scanned?	Yes □ No □
If yes, will a restricted Scanner be used?	Yes □ No □
Any other considerations?	

Table 31 **Remote access users**

Remote Access Users	
Are they to be scanned?	Yes □ No □
If yes, will a restricted scanner be used?	Yes □ No □
Any other considerations?	

Table 32 **Standalone users**

Standalone Users	
Are they to be scanned?	Yes □ No □
If yes, will a restricted scanner be used?	Yes □ No □
Any other considerations?	

Table 33 **Non-operational equipment**

Non-operational equipment	
Are they to be scanned?	Yes □ No □
If no, will data be captured manually?	Yes □ No □
If yes, who will make them operational?	
Any other considerations?	

Table 34 Other miscellaneous considerations

Consideration	
Security software installed that prevents an executable file from being run from the floppy drive?	Yes □ No □ If yes, identify: ___________
Floppy disk locks in place?	Yes □ No □
Asset labelling system in place?	Yes □ No □
Scan schedule in place?	Yes □ No □
Correct access rights to scan machines?	Yes □ No □

4 Planning an Enterprise Discovery Deployment

This section provides information on how Enterprise Discovery can be deployed to discover, collect and maintain current inventory data.

You will find information on the following topics:

- What are IT assets
- What is Discovery and Inventory Data
- Purpose of an Enterprise Discovery Deployment
- Planning the Enterprise Discovery Deployment
- Steps for Planning an Enterprise Discovery Deployment

These are key concepts that are important to understand before proceeding with an inventory.

What are IT Assets?

An IT asset is any piece of IT equipment or software that your company owns or leases. The information about these assets should be stored and kept up to date in your IT asset database.

Hardware IT Assets

Examples of your hardware IT assets are:

- Desktops
- Workstations
- Servers
- Portable devices (for example, laptops)
- Printers
- Modems
- Monitors
- Keyboards
- IP Telephones
- Scanners
- Routers
- Bridges
- Switches

Software IT Assets

Your software IT assets are the software and applications that are being run on computers in your organization. Examples of software IT assets are:

- Commercially available software applications
- Proprietary software your company has produced
- Information stored in files

What is Discovery and Inventory Data?

Discovery and Inventory Data is information about all IT Assets in your organization. Enterprise Discovery is used to automatically detect and collect network, hardware and software data for these devices.

The data discovered by Enterprise Discovery can be used in a number of ways, even without integrating it with an asset repository:

- To get an accurate list of devices in your network, including core network devices, IP telephones, computers, and intermittently connected devices such as laptops.
- To get a complete view of software applications deployed,
- To identify top software license requirements by publisher or application,
- To identify devices in your network that do not work as well as they should.

To get the full value of the data discovered by Enterprise Discovery, it should be reconciled with the financial data held in a central asset repository such as AssetCenter. This makes the data valuable in a much wider range of business processes:

- As a valuable company resource for other departments.
- For data management. Keeping track of versions of electronic price lists for consistency, keeping track of databases for business continuity etc.
- To perform spot checks to ensure all changes within an IT asset's life cycle are recorded.
- To manage hardware and software resources efficiently. This results in more effective user support, software purchasing, licensing and better hardware utilization.
- To detect and solve problems such as software piracy, computer pornography and other abuses.
- As a mechanism to ensure compliance with internal standards for software application licensing.
- To drive a range of standardization initiatives, whether hardware- or software-related.

Purpose of Conducting an Enterprise Discovery Deployment

An Enterprise Discovery deployment can be conducted for a number of reasons, including software license compliance, but is also an opportunity to gain beneficial information and effective control of costly IT assets.

It is important to consider various business drivers in order to maximize the benefit. Once Discovery and Inventory data are available within an organization, interested parties begin to request information for their particular needs. Consider the needs of the following functions:

- Technical Support (Helpdesk)
- Asset management
- Software licensing
- Business continuity
- Disk grooming
- Platform upgrading
- Software Licensing
- Virus Impact Assessment

Technical Support (Helpdesk)

Helpdesk inventory data usage can include:

- User and machine information
- Software versions
- Configuration files

Inventory data facilitates Helpdesk by:

- Reducing time spent managing users
- Detecting problems earlier
- Verifying implementation of changes
- Giving users better service

Asset Management

Asset management data usage can include:

- User and machine information
- Software versions
- Software utilization information

This data can be used to:

- Analyze the computer estate
- Reduce software license infringement and down-time
- Optimize software license costs
- Improve asset management and utilization
- Improve supplier leverage

Asset management products often manage only a minimum set of data fields and traditionally have focused on hardware information. To make best use of the data, you should use AssetCenter, which now includes a Software Asset Management module dedicated to managing and optimizing software licenses.

Business Continuity

Data necessary for business continuity is:

- User and machine information
- Software versions

This information is used to:

- Use stored files to aid computer configuration
- Check restored data for versioning

Business continuity is an important aspect of asset management. One concern is recovering machines critical to the organization, after a failure. Another issue is change. As users are re-deployed within an organization, knowing their previous equipment capability makes it easier to assign a machine with the same or better specification. If new equipment is needed, the previous specifications are available.

By comparing the inventory of a rebuilt machine with its baseline, differences in software versions and additional data files can be highlighted. When an inventory is complete, an accurate status of changes in assets, hardware, software or deployment can be maintained. Key files can be embedded into the inventory data so that they are available from an alternative resource to the specific machine.

Disk Grooming

Enterprise Discovery can be configured to collect information that can be used for disk grooming, such as:

- Duplicate file and application installations
- Multiple versions
- Small and empty files
- Unbalanced directory structures
- Mixed data and program directories
- Space available on accessible drives

Disk drives tend to accumulate superfluous files, such as, old versions of programs, more than one copy of a local file and old installations that were incompletely removed. Empty and small files consume at least a few KB of disk space per file, and even the contents of Temp directory of the recycle bins can impact disk space.

Duplicate applications can have licensing implications. Some vendors do not permit more than one version of their software on a computer without an additional license. This can happen when upgrades are installed and a previous version remains on the machine during the migration period.

Applications, particularly graphical and multi-media applications, consume a significant amount of disk space. If applications are to be migrated to a server, then it is vital that odd data files are not mixed with the executable files as they might be given access rights preventing user updates.

By making periodic inventories, disk usage can be tracked and remedial action taken. This will help users organize data more effectively and increase their productivity.

Platform Upgrading

Inventory data facilitates platform upgrading in:

- Determining current configurations
- Comparing current configurations with target requirements
- Controlling rollout programs

When upgrading software or migrating data, it is necessary to know the current and the target configurations. By comparing current and inventory data with target requirements, it can be determined which machines meet specifications, which need to be upgraded and which need to be replaced. Re-inventory of machines after they are upgraded can identify what ‘standard builds’ are compromised and provide a check on the progress of a rollout program.

Software Licensing

Inventory data can be used to:

- Identify applications and versions
- Produce summary counts
- Update software indices for local applications
- Identify redundant software installations that are not used or under-utilized.
- Check license breaches

Due to the number of phases involved in the process, software licensing needs to be handled as a specific project. Since many applications occur as software suites, inventory data is useful for matching software licenses to application counts and can help companies avoid the over-purchasing of licenses.

Establishing ownership of software may require input from sources such as suppliers. Even if unlicensed software is detected, the user should not hesitate to consult suppliers. Software publishers are generally pleased to know that users are taking steps to mitigate the problem.

Virus Impact Assessment

Inventory software is not a replacement for virus detection; however, an important use of inventory data is to check the current version of the anti-virus software deployed. Sometimes the most current version is not fully deployed leaving the computer unprotected against new viruses.

Inventory data provides support information for virus detection. For example:

- If a boot sector virus has been detected and cleaned, it may leave that boot sector inoperable. A standard inventory embeds a copy of the boot sector in the data it collects, and this information can be used to rebuild the damaged sector.
- Some viruses create a time stamp of 62 seconds. The file list can be searched for such an occurrence.
- Some viruses create or rename files. By knowing the file names, a search can be undertaken.
- If a file has been infected, it’s size could have changed. Re-running the scan may result in different data being produced.

Reusing Inventory Data

Once inventory data is collected, it is necessary to keep the data current. Managing upgrades is one re-use of inventory data, others are:

- Producing management reports of product usage
- Cutting out non-essential product evaluations
- Supplier leverage

Planning the Inventory

For an inventory to be successful, front-end analysis should be done to fine-tune the objectives, determine the ultimate use of the data and specify exactly what data is needed. Well conducted front-end analysis coupled with effective project management practices ensures the successful implementation of the IT asset inventory with minimum disruption to your business.

Overcoming the Data Mountain

At the start of the inventory, the exact data requirement may be unknown. Once the user begins looking at the inventory data, there is a temptation to start investigating all sorts of ancillary issues. For example, what about those stored outdated computers? As a result, the people performing the inventory may mistakenly employ the ‘if in doubt, inventory it’ approach.

While there may be a short-term tactical need for specific information, or the inventory might be an opportunity to gather information that is difficult to capture, redundancies can be eliminated by ascertaining exactly what information needs to be extracted.

For example, machines awaiting disposal may not need to be inventoried, but may simply need manual asset recording. Also, consider whether mice and keyboards need to be recorded.

As a general guideline, focus on achieving specific objectives before investigating special interests. If the data collected is not going to be maintained, then its inventory capture should be questioned.

The goals of an IT Asset Inventory

The major goals of an inventory might be:

- Listing all known applications – indicating how many licenses are needed.
- Listing all unknown software – highlights any threats.
- Reporting asset deployment – assists with asset management.
- Listing all computer and server hardware – aids in future upgrade plans.
- Maintaining the Discovery database

Steps for Planning an IT Asset Inventory

The steps outlined here will help bring into focus some of the issues you may face when planning your IT asset inventory.

Step 1: Identifying Your Existing Data Collection Process and Current Environment

Throughout your organization you will find that different departments will be using different methods for the collection of their data.

For example:

- Your HR department may be using a spreadsheet that contains all the data for employees (Employee names, functions, departments, software contracts etc.).
- Your IT department may use an in-house database that contains data on the machines the company owns.

Identifying existing data collection methods is also an important first step in identifying the data needs of your various departments.

The Current Environment

- Have previous scans been undertaken?
- Are there any existing electronic identifiers?
- Are there any facilities for application deployment?
- What percentage of machines are networked?
- How many servers?
- What operating environment do they use?

Step 2: Planning the Collection of This Data

The data collection processes identified in Step 1 will become a source of one-time input into your repository.

Step 3: Designating and Training Members of Staff for the Maintenance of the Data

You must designate tasks to employees who can ensure that the data accuracy and consistency is maintained regularly.

In order to ensure data accuracy and consistency, both across the organization and over time, it is important that ownership of these issues is assigned early in the process.

These tasks must be considered a vital part of the overall asset management effort, as accuracy of the data is an absolute requirement if later analysis is to be of significant value.

It is vital that there are staff nominated as contact points to answer both project and technical questions. There may be more than one person in each category depending on the size of the project.

Step 4: Deciding What Data is Needed and Determining How to Source that Information

You will need to take a detailed look at the data that is required. Bear in mind that the data needs will be different from department to department. Step 1 will have already provided some indication as to what these departmental data needs are.

This can have a significant bearing on how a project is undertaken and the amount of data to be collected. Also consider whether the data will be forwarded to other applications, if so, how is this envisaged.

Further Information

For information on what data scanners collect, see **Help** > **Data Collected by the Scanners**.

Step 5: Configuring the Scanners

Configure a Scanner that collects this data for the various platforms used by the computers in your company using Scanner Generator.

Hardware detection is fast - typically 10-30 seconds. The main areas that need configuring in Scanner Generator are Software and Asset data collection. For almost all purposes, you can use the default Hardware detection settings.

If possible, avoid configuring more than one Scanner for each platform. Running different Scanners for different departments can become a labour intensive exercise and should be avoided if possible.

Further Information

For more information on creating customized Scanners, refer to the Enterprise Discovery *Configuration and Customization Guide*.

Step 6: Deciding How You Will Gather the Data

The Scan Repository

Regardless of whether you carry out a manual or automated scan, the scan files will end up in your repository. The repository should be designed to hold all your scan files and should be cleaned up periodically.

Automatic Inventory

This type of inventory will allow you to collect information about hardware and software assets and pinpoint basic information about where those assets are located, who is logged into them, and what operating systems they are running on etc.

This information can serve as the basis for the initial walk around inventory as it helps define and establish “what is there”.

The Scanners are distributed to individual machines. You can set up a schedule using Enterprise Discovery dictating which machines should be scanned and at what frequency. The retrieved scan files are placed in a central repository.

You can then use the data for use in asset management systems, such as AssetCenter using Connect-It.

Manual Inventory

A walk-round (or manual) inventory captures data about assets that are not connected to your network (i.e. stand-alone).

In these cases a memory stick or floppy disk with the Scanner executable will ensure that all important configuration items about that PC (such as installed software, and hardware), monitors and Asset Tags, etc. are known.

During the course of the inventory you may come across unused and surplus assets. These surplus assets can be re-deployed to other employees or should be properly removed from the inventory.

A walk-round inventory also allows each asset to be physically inspected. The state of the machine can then be stored in an asset field.

In manual deployment, you still need to add the recognized application to the scan file. The XML Enricher performs this task. You must also make some special configurations to accommodate stand alone scan files. They need to be added to the IPv4 range as devices with scan only data. To do this, you need to add the Restricted to scan only attribute to the range. The results of the scan must then be saved to the central repository.

Step 8: Creating and Maintaining a Set of Methodologies

Inventory capture is not a one time event.

Once the initial asset inventory and asset reconciliation is complete, you should create and maintain a set of processes that are designed to keep your inventory up-to-date.

- Any time a machine has had changes made to it, for example, software and hardware upgrades, it has been allocated to another member of staff etc., a scan should be initiated.

  The fields about the user and asset can then be updated and automatically reconciled again against the data in your asset management system for example.

- Spot-check practices should be established that verify automatically-collected data samples for accuracy.
- Non-networked assets must have a process for the manual entry of new data initially and then follow-up inventories must also ensure that these assets are included and the data is kept up-to-date.
- Scans should be run on networked systems at least once monthly to keep the data about these systems up-to-date.

Re-inventory

The regularity of re-inventory depends on a number of factors, including:

- How often assets change condition - moves, upgrades, additions etc.
- How often information needs to be reported. For example, is it necessary to perform a daily inventory check if the results are only reported quarterly?
- How often the people in charge are likely to have the time to look at it.
- What it is used for.
  - If for asset management - then rarely

- If for services/support - then more often.

5 Frequently Asked Questions and Answers

General

What version of Windows 2003 Server is recommended and why?

The recommended system requirements specify a Windows 2003 SP1 or SP2 server. Depending on which document you read, you will also see that Windows XP is supported but not recommended for use in a production environment. There is no distinction between the Windows 2003 Enterprise Edition and the Standard edition. The only real requirement is that the server be dedicated to Enterprise Discovery.

Is there a worldwide customer or referral list?

Product Management is working on compiling a list of customers who are willing to participate in a referral program for Enterprise Discovery. Currently, Product Management is investigating the Americas accounts with EMEA to follow before the end of 2007.

What is the purpose of the Web Asset Questionnaire?

The purpose of the Web Asset Questionnaire (WAQ) is to have a central place for administrators to enter Asset Information about a device. It allows IT staff to fill assignment information eventually on site when physically setting up the system using a simple browser. This questionnaire enables you to associate a person's name, department, phone number, or other personal information that you might want to associate with a given device in the Enterprise Discovery database. This data will be saved with the other data for a specific device (obtained by discovery or scanning) and will appear in the Device Manager.

What version of Perl does Enterprise Discovery use? Why does it cause issues when installing?

Enterprise Discovery uses ActivePerl version 5.8.6. ActivePerl is installed as a 3rd-party utility for many other published applications. Enterprise Discovery will fail to install if it determines that ActivePerl is already installed on the system. In some cases, however, the Enterprise Discovery installer has not been able to identify that ActivePerl is already on the system. Before you begin the Enterprise Discovery installer, run the following command from the command prompt to determine if ActivePerl (or some other version of Perl) is already installed on the system:

```
Perl -v
```

It is also advisable to do a quick search in the Windows registry for the keyword Perl. If any version of Perl is already installed on the system, remove it before installing Enterprise Discovery.

When troubleshooting Enterprise Discovery, is it OK to start Apache manually?

Normally, the HP Enterprise Discovery System Monitor controls the Apache process. There really is no need to start any service manually. If there is an issue with the Apache service, open a support call and discuss this with a technician.

What are the demo requirements?

If you are running Enterprise Discovery VMware demo images, the hardware requirements are at least 2 GB of memory on the host system. If you are planning to run more than just the Enterprise Discovery image, you will need more memory.

Is there an advanced-level forum with developers and product management where experienced field consultants can get involved with the products direction and advancement?

There is a forum that has been established and will be used for most of the above mentioned departments. Of course, there may also be customers in this forum as well, so there will be rules of engagement. For more information, go to the following URL:

http://forums1.itrc.hp.com/service/forums/categoryhome.do?categoryId=682

History

How long has Enterprise Discovery been around?

The history of Enterprise Discovery goes back to companies acquired by Peregrine Systems in the late 1990s and early 2000s. The predecessors of Enterprise Discovery were two stand-alone Peregrine legacy products called Network Discovery and Desktop Inventory. When these two products first came together, the result was known as Enterprise Discovery 1.0. At this point, however, the two pieces were in reality still two separate products. In 2005, the first version of Enterprise Discovery was released in the form of Enterprise Discovery 2.0. This combined the technologies of Network Discovery and Desktop Inventory integrated in a single solution. After a one year development phase, with the Enterprise Discovery 2.1.0 release, the remaining pieces of both products were completed bringing together all the features of the discovery products.

What are the differences between Enterprise Discovery and Desktop Inventory?

See the response for the "How long has Enterprise Discovery been around?" question.

Enterprise Discovery Compared to other Products

How does Enterprise Discovery compare to Mercury Application Management?

Enterprise Discovery and Mercury Application Management (MAM) have complementary features. Together they can be used to populate your Configuration Management Database (CMDB) with detailed information about the hardware and software assets in your organization, as well as the relationships between these assets.

Enterprise Discovery is focused on the discovery and inventory of assets while MAM is focused on providing visibility into the complex business services and application relationships that exist within your IT infrastructure.

However, it is understood that there are still questions around the positioning and integration between these products and so this topic will be addressed in a separate communication that will be released shortly.

Will Enterprise Discovery install software like the Radia Application Manager?

Installing software will remain a function of Configuration Manager (Radia). There is no road map item to allow Enterprise Discovery to perform this task.

Will Enterprise Discovery delete and replace files like the Radia Inventory Manager?

The goal is to continue with the features in Configuration Manager (Radia) while utilizing the data that is collected by Enterprise Discovery.

Does Enterprise Discovery provide the equivalent of Configuration Manager data?

A document is currently being written that maps the tables in the Configuration Manager database to those in the Enterprise Discovery database. It will be made available shortly.

What additional fields does Enterprise Discovery Collect that RIM does not?

Refer to **Help** > **Data Collected by the Scanners** in the GUI of Enterprise Discovery.

Discovery

Is the discovery limit of 50K devices hard coded? Does it have performance impact?

In Enterprise Discovery 2.20, the license schema changed slightly. Discovery and Inventory is one item on the price list. If you purchase a 50,000 device Discovery and Inventory license, you are automatically allowed an extra 10,000 devices for discovery. This is a 20% increase on the number of devices that the license includes. This is hard-coded into the backend of Enterprise Discovery so that you are able to inventory up to 50,000 devices and discover another 10,000 devices that will not be inventoried.

Enterprise Discovery works best when you are allowed to discover the devices that connect the end nodes. More information can be found in the *Configuration and Customization Guide* and the *Release Notes* for Enterprise Discovery 2.20.

What is the IP range-limit?

Enterprise Discovery does not have any hard-coded limits for groups or profiles. As a result of benchmarking configurations, the recommendation is to have no more than 2000 groups. Reasonable estimates are provided in the "Configuring the Discovery Process" chapter of the *Installation and Initial Setup Guide*.

Do you provide better format controls for IP ranges in the current version of Enterprise Discovery?

With the release of Enterprise Discovery 2.20, we have introduced a new Discovery Configuration that uses configuration profiles. Each profile is then added to Device Groups that add in the IP ranges instead of adding in the IP range and assigning a Property Set to that range. This new method of discovery configuration allows for better flexibility in adding ranges and removes the issue of having only 2,000 IP ranges configured.

Additionally, you are able to add in Device Groups that will mange different Devices Types. For example, you may want to scan devices with a hardware only scanner on all Sun servers, while performing a complete scan on all Windows workstations. This is now a configurable item allowing better control over the discovery configurations. Also, see the response for the "Why are there no inherit or global properties in Enterprise Discovery 2.20?" question.

Why are there no inherit or global properties in Enterprise Discovery 2.20?

Enterprise Discovery 2.20 includes a new Discovery Configuration feature that enables you to precisely define what devices in your network it will discover and how these devices will be managed. There are two primary elements:

- Configuration profiles specify how network devices are discovered and managed by Enterprise Discovery.
- Device groups specify what devices are discovered and managed.

You establish device groups by creating one or more conditions that specify a collection of IP addresses, a particular type of device, or both. You then assign configuration profiles to a device group to specify how the devices in that device group should be treated.

For example, previously in Network Configuration, you would use the SNMP Property Group to set all the community strings for use with the IPv4 ranges. This would then be inherited all the way down the IPv4 tree.

Now, you create an SNMP profile with a set of community strings and assign that profile to all the device groups that require this set of community strings.

Because the new Discovery Configuration process provides much greater flexibility and precision than the old Network Configuration process, there is no longer a need to inherit global settings.

What information is available from the OID/MAC? Are there any mismatches?

SNMP Standards are described in Request for Comments (RFC) documents published by the Internet Engineering Task Force (IETF). Standards Topics can generally be categorized into the following:

- Messaging protocols between Managers and Agents (which encompasses security issues)
- MIB syntax standards
- "Standard MIB" definitions

Why are servers, laptops/workstations often reported as NIC card manufacturers?

The MIB and OID information is related to the manufacture of the NIC. The devices identified based on their NIC card are usually unmanaged devices. As a result, the information about the MAC address is not determined by the MIB. The first part of the MAC address (called the OUI) provides some information about the manufacturer. This is the information that Enterprise Discovery can identify. There is no way to really change this. If the NIC is coming up incorrectly and you report it, it is possible to check the MAC assignment by accessing the following URL and typing in the MAC address.

http://standards.ieee.org/regauth/oui/index.shtml

What happens when many components change on a device? How is the device merged into Enterprise Discovery database?

Enterprise Discovery tries to maintain a given device as one device in its database as best it can. It uses many merging criteria to maintain this information. However, it will not be able to merge information on one device if several component on the device change at the same time. For example, if a device gets a new NIC installed and a new IP Address, but the hostname remains the same, Enterprise Discovery can make a reasonable attempt to merge the new information for the correct existing device in the database. If many identifying elements change, for example, new NIC, new IP Address, and new hostname, it is possible that the device will appear as a new device in the database. If the device is not reconciled, a new entry is created for the new device, but the old device will be purged.

Other criteria for merging devices use NetBIOS information, which can maintain the model as well.

How do environment DNS issues affect Enterprise Discovery?

Reverse DNS records are used by IP hosts to resolve an IP address to a DNS name. A DNS server stores these records as follows:

```
52.0.0.10.IN-ADDR.ARPA IN PTR C.EXAMPLE.COM
```

In this case, `52.0.0.10.IN-ADDR.ARPA` (the IP address 10.0.0.52) points to `C.EXAMPLE.COM`.

A reverse lookup for the address 10.0.0.52 in the `EXAMPLE.COM` network resolves to `C.EXAMPLE.COM`. A DNS implementation is incomplete until it has reverse lookup records. IP addresses in reverse DNS records must reference one and only one primary DNS name. DNS servers should never have conflicting reverse lookup records for a given IP address. In networks where a single IP address may have multiple DNS names, choose only a single DNS name for reverse lookup.

Having multiple reverse lookup records will cause the resolved host name for a given IP address to become random and unpredictable. This creates problems in network management systems and in systems that depend on reverse lookups for a measure of security or selection. When this condition exists, Enterprise Discovery may report DNS names for devices incorrectly and inconsistently.

Replace any alias entries in reverse lookups with canonical DNS names. Delete any duplicate entries in reverse lookups. Enable automatic scavenging of stale records in Microsoft Windows 2000 Domain Name System (Refer to Microsoft Knowledge Base Article - 296116). On UNIX systems, use the `host` or `dig` command to examine entries. Avoid using the `nslookup` command since it shows only one reverse DNS name.

What is the best configuration to handle roaming users? DHCP?

Roaming users represent a difficult set of users to control. When VPN users attach to the network, you must configure the device to force ARP table reads, so that you are able to monitor the device as a new IP address is assigned to that device. Internal roaming users are handled with a list of various merging criteria. This helps maintain the device's existence in the network. The key here is to make sure you are forcing ARP table reads to the devices that are giving and assigning IP addresses.

Scanners

What is the typical size of a scanner XSF file?

An XSF file is a compressed XML file, which is created by the scanner when it scans devices for inventory information. The XSF file is compressed to make the size of the file relatively small. The file size does grow with the size of the system being inventoried. An average Windows desktop scan file will be around 240 KB when conducting a targeted scan, while a UNIX workstation scan file will be 1 to 2 MB in size. Servers tend to have a larger XSF file size because of their large file systems. An average size for Windows servers is 2 to 3 MB. An average for UNIX servers is 6 to 8 MB, growing to sometimes as large as 240 MB. Care must be taken on these types of systems so that you do not have a scanner running for a 24 to 48 hour period. Some mounted file systems may not need to be scanned. Customers are advised to exclude those file systems.

Refer to the “Disk Space on Managed Devices“ chapter in the *Installation and Initial Setup Guide* for more detailed information about scan file sizes.

What happens when a file is in use when a scan is running?

If the file is in exclusive use, its file name, size and attributes can still be collected for inventory purposes, but the scanner is not able to either identify the file type (for example, executable file type, archive, and so on) or collect the file’s signature as the file cannot be opened.

What is the load on the CPU with inventory processing? What are various configurations and options?

It is difficult to determine this type of benchmarking because the inventory process is highly configurable. The inventory process has two detection phases, namely, a hardware detection phase and a software detection phase. In the both phases, the scanner will use CPU resources according to the enabled priority settings configured for it.

During the software detection phase, the scanner is collecting file information for software recognition. The scanner can be configured to throttle CPU usage during this phase. Although different methods are used to accomplish throttling on Windows and UNIX, the CPU usage will increase when these settings are used. When the scanner detects that a screensaver is active on Windows, it assumes that the user is not using the system. The inventory scanner will try to use as much CPU as possible at this point. On UNIX, the `nice` command is used. This ensures that the scanner does not take up CPU usage when another intense process is using CPU resources. When throttling is turned off, the scanner will attempt to scan as quickly as possible.

Windows scanners performing a targeted scan have an average running time of 3 to 6 minutes. The UNIX systems can have very large file systems, which causes the scanning process to take longer. In general, an average cannot be determined for UNIX systems. On a small UNIX server, however, the average time is 10 to 15 minutes. Of course, if throttling is turned on, the time increases for both Windows and UNIX systems.

Can publisher information be excluded in the scanner?

No. Publisher information is performed at the recognition level not at the scanner level. The scanner does not know who the application's publisher is at the time of scanning. The scanner is collecting the files for the Recognition Engine. The Recognition Engine then uses this information to identify the publisher.

How many scans can Enterprise Discovery process at any one time?

The Enterprise Discovery server can communicate with up to 80 agents installed on managed computers at any one time. It is the agent's job to upgrade the scanner, execute the scanner, and collect the scan file after scan completion. Each one of these tasks takes up a connection to the device using one of the 80 connections available.

There are also four connections reserved for manual actions. At any time, a step in the scanner execution workflow can be initiated manually. The subsequent steps of the workflow are then executed automatically. For example, if you manually trigger the scanner execution, the scan file retrieval will happen automatically as well.

Although there are only 80 agent connections available, the connections are normally in use for short periods only, and as a result, the Enterprise Discovery server can launch the scans of hundreds or even thousands of managed computers at any one time.

What hardware fields are available for collection for each OS?

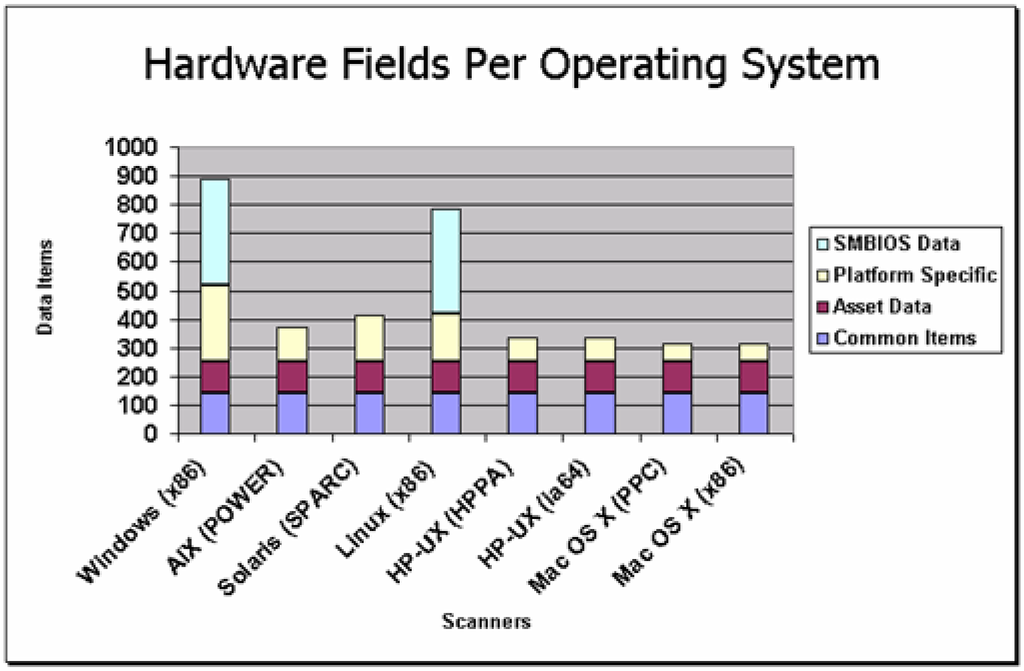

When is scan data available for viewing in the Device Manager?

The Java Viewer option within the Device Manager can take up to 30 minutes to be available after the inventory scan files are collected. This may increase depending on how many scan files are being processed to the database. For immediate viewing you should be able to collect the scan file from the processed directory or in the Device Manager's Diagnostics screen and use the Windows Viewer to display the data.

Will WMI data be collected in future versions of ED?

Installed application information is collected from WMI in Enterprise Discovery 2.50. Details on the future plans for Enterprise Discovery will be communicated as part of the Enterprise Discovery road map.

Scanner configuration can be confusing. How is the best way to get started?

The best way to understand scanner configuration is to start with the pre-defined scanners. They are descriptive enough to get you started. You can then use the Scanner Generator to design a customized scanner. Use the Manual Mode to test scanner generation prior to creating your scanner to better understand the data that can be collected. Refer to the "Scanner Generator" chapter in the *Configuration and Customization Guide*.

Is OS installed application data collected from the registry and how?

Yes, the OS installed application list is collected with the Inventory scanner and is part of Hardware Data collection. This application information can be used for recognition if the recognition setting is set to installed applications. This is only recommended for initial inventory stages. The SAI recognition should be used to obtain a normalized view of applications.

Software Teaching and Recognition

What is the percentage of software recognition?

In a Windows environment, it is typically about 80% of the applications.

In UNIX, it is generally about 30% of the applications.

There are efforts underway to improve UNIX application recognition.

What are the software teaching methods?

There are several techniques for teaching applications. They are based on maintaining a Master library of applications. See the next "How does HP actively maintain the software library?" question. If you want to know more about teaching methods, request the `Teaching_an_Application_HP.pdf` document from HP Support. Also, refer to the "SAI Editor" chapter in the *Scan Data Analysis Guide*.

How does HP actively maintain the software library?

Maintaining a current, accurate software library is very important to HP. HP has a dedicated team of librarians who are continuously adding applications to the Master library, referred to as the Master Software Application Index (SAI). Updates to the library are made available to you monthly through the Support website.

The ability to add applications to the Master SAI is contingent on HP's access to the install media for the applications. HP is working with other software vendors, partners, and customers for such access. See the next "What is the procedure to request software to be added to the Master SAI?" question as to how you can help.

What is the procedure to request software to be added to the Master SAI?

HP provides a Master SAI for out-of-the-box recognition of mainstream software applications. It is expected that customers supplement this with a User SAI. If you want HP to add an application directly to the Master SAI library, file a ticket through Support. Depending on the nature of the application and its suitability for the broader Enterprise Discovery customer base, HP will work with you to add it to the Master SAI as long as access to the install media is provided.

After teaching applications to the User SAI, how can they be added to HP's Master SAI?

HP welcomes the opportunity to update its Master SAI library. See the previous “What is the procedure to request software to be added to the Master SAI?“ question.

What should I do if an application is not recognized even though there is an entry for it in the Master SAI library?

If there is a Master SAI entry for the very same version of the application that is not being recognized, file a ticket with HP Support and HP will investigate the issue.

Will there be an Apple Macintosh SAI? What will the new automatic teaching process be?

The Software Application Index (SAI) is used for application teaching and recognition. At this time, Mac OS X support has just been introduced for the inventory process. The support for Mac OS X in the UNIX master SAI is being considered. However, regular application teaching techniques can be used to create an in-house Mac OS X user SAI.

Does the Master or User SAI take precedence? Is the rating automatic?

The rating system in the recognition engine takes all loaded Software Application Indexes (SAIs) into account. As each device file is passed through the application recognition engine, the file is given a rating based on the matches found within the SAIs. When all of the files have completely loaded into the recognition engine, the rating decides what the best match is.

What is the SAI upgrade process? When is the Master SAI not required?

The current version of Enterprise Discovery has not had any major changes to the SAI. There is a utility in Enterprise Discovery that will upgrade the SAI to the Enterprise Discovery version. Refer to the "SAI Editor" chapter in the *Scan Data Analysis Guide*.

Can a file range be specified in application recognition?

No, there is no option to enable a file range for recognition. This was removed back in Desktop Inventory, and it is unlikely that it will be reintroduced. You can ignore the file size in the teaching process, but all Master libraries have this option turned on, and it is used in the recognition process.

How are browser-based applications recognized?

In most cases, there will be an associated shortcut or something that points to the web application. The shortcut must be configured for collection and then taught to the SAI for recognition.

Why does the Software Application Index file use the zSAI extension?

The SAI extension was used as the extension for all the Master libraries prior to Enterprise Discovery 2.1.0. With this release, the zSAI extension is used. The SAI files are now XML-based and GZIP compressed to reduce the size occupied on disk, which is why the new `.zsai` extension has been introduced.

What is the purpose of the auto.zSAI file?

The purpose of the `auto.zsai` file is to keep track of applications recognized by using recognition rules. It ensures that such applications have consistent IDs associated with them. The rules-based recognition feature helps Enterprise Discovery anticipate new versions of the applications already on your workstations and already saved in your User SAI file. By looking at the current data for an application, you can create a rule (based on a series of regular expressions) that will provide automatic recognition of any new versions of that application.

When teaching the SAI and reprocessing scan files, when do scan results get passed to the Viewer?

In the Windows viewer, you are able to open the XSF scan file as soon as it is returned to the Enterprise Discovery server. In the Java Viewer, there is a processing wait time of about 30 minutes. The Recognition Engine is built into Analysis Workbench, Windows Viewer, SAI editor, and XML Enricher. The purpose of the XML Enricher is to process the scan file information into recognized files, unrecognized files, and partially-recognized files.

The Windows Viewer performs its own recognition, so the recognition results are available immediately, while the Java Viewer shows the recognition information as produced by the XML Enricher and therefore there is a delay, that is, the scan file has to be processed by the XML Enricher and then the resulting data imported into the database.

How can you pre-fill the SAI if application information is not available?

In Analysis Workbench, you are able to specify this data. You can use the SAI Teaching Mode as a way to add Publisher, Application, Release Levels, and Version prior to adding this to the SAI directly. This is known as the Holding area, and you can perform this process by right clicking to add information from the file window. To use this process you need to understand what you are teaching, but it gives you two windows that allow you to show and add the information to the SAI.

Does having more Main files per application give a better recognition result?

Not in all cases, but for larger suite-type applications, it may make sense to teach more Main files.

Is there a published list of SAI contents?

There is a list of published applications. This comes in an HTML format and can be requested from HP Support.

Agents

What agent ports are required?

It is required that one of the two required agent ports be available. All communication goes through this single port. The two ports you can choose from are the following:

- TCP 2738 - Unassigned (legacy)
- TCP 7738 - IANA Assigned

More information is available at the following URL:

http://www.iana.org/

What are the methods for distributing the agent? Can an automated method be provided for UNIX?

There is a UNIX script template that can be configured within Enterprise Discovery to deploy UNIX agents to a device. It requires a UNIX System Administrator to integrate it into the organizational environment. Once the script is configured and able to communicate, you can begin to deploy the agents using a custom deployment method that calls the created script. This script template is supported by HP Support.

Virtualization

What are the virtualization technologies currently available in Enterprise Discovery?

There are two different technologies that are being used in the current version of Enterprise Discovery. With the release of Enterprise Discovery 2.20, virtualization is supported in VMware ESX 3.0 Servers and Solaris Zones (containers). These technologies use different collection methods to understand the virtual environments that are running on the host system.

In VMware, web service calls are used to find the information from the host device about its virtual machines (VMs). Enterprise Discovery agents must be installed on the physical host machine, as well as the hosted VMs. An account is created within the Configuration Profile that accesses the host system and initiates the calls to collect information for the VMs configured on that host system. Refer to the “Configuring the Discovery Process“ chapter in the *Installation and Initial Setup Guide* for more information about Configuration Profiles.

In Solaris, the Solaris Zone (Containers) technology must be configured and active. The host system must be discovered and have an Enterprise Discovery agent installed. Once the host device is discovered and the agent is communicating, a scanner is executed in the global zone only. The collected scan information is then processed by the XML Enricher, which splits the global scan file into individual scan files for each Solaris local container (hosted virtual device).

Enterprise Discovery not only provides the basic discovery, inventory, and software utilization data but also determines the parent/child relationship between the physical host device and the virtual devices hosted on the physical device. As such, you are able to see the host/virtual device information in several places in Enterprise Discovery including the Network Map, the Virtual Devices Window, the Device Manager, and the Virtual Device Reports. Refer to the "Virtualization in Enterprise Discovery" chapter in the *Reference Guide*.

How are Virtual Machines represented on the Network Map?

Virtual machines are still displayed and counted as devices in the database and on the Network Map. The Network Map does not show a connection from a virtual machine to its host. Instead, the host system will display a list of its virtual machines in the Device Manager. If virtual machines appear on the map, they will show a connection, but not to their host. There is also an option to disable sending devices with virtual information to the Topology Engine and, as a result, not display them on the map.

However, it is still possible to create your own package and display this on the Network Map. You can verify which virtual machines belong to the host system by looking at the Device Manager.

Licensing

What are the license changes?

License changes are done as part of the release for Enterprise Discovery 2.20. More information on these changes can be found in the sales portal.

Does a printer usage report require resource managed and/or Alarms license?

Yes, the Alarms license is needed for these types of reports, but this license is not available to new customers.

Security

Is there 168-bit security?

Yes. Agent to server communications are authenticated and encrypted using RSA 1024 and 3DES. The 3DES key length is 168 bits.

Software Utilization

How does the Software Utilization Agent work?

The Application Usage modules collect the information about executable files within the process list. They monitor the executable files and produce usage statistics at the device level. When the inventory scanner runs, it picks up the files that were produced by the Usage monitor. This information is then used on the server to identify the specific applications to which the executed processes belongs. Usage information for an application is not displayed unless it is recognized by the Software Application Index (SAI).

What do the utilization figures in the Viewer represent?

- Days used (last month)
  - The number of days the application was used in the last month.
- Hours used (last month)
  - The number of hours the application was used in the last month.
- Days used (last quarter)
  - The number of days the application was used in the last quarter.
- Hours used (last quarter)
  - The number of hours the application was used in the last quarter.
- Days used (last year)
  - The number of days the application was used in the last year.
- Hours used (last year)
  - The number of hours the application was used in the last year.
- Hours used (average daily)
  - The daily average in hours over the period configured.
- Hours used (peak daily)
  - The highest daily number in hours over the period configured.

SNMP

Does SNMPv3 create further traffic?

The process that is used to generate packets in SNMPv3 is still the same as in SNMPv2 and SNMPv1. There is a little more traffic on the network with SNMPv3, but it is still low so that it does not impact the network's backbone. Here are some of the strengths and weaknesses of the SNMP protocol.

Strengths:

- Widespread popularity
- Many standard MIBs available
- Agents have low impact on monitored system resources
- Well suited to monitoring
- Many products available

Weaknesses:

- Not as comprehensive as some other protocols
- Not bandwidth efficient
- Complicated message encoding rules
- Security has been an on-going concern. SNMPv3 was developed in response to this issue.
- UDP or some other connectionless protocol is used, which creates issues regarding verification of operations

Analysis Workbench

How do you use the load options in Analysis Workbench?

The load options in Analysis Workbench are templates for settings that are used in the Analysis Workbench. The average user would use the default load options. The typical Sales Consultant is familiar with the load options and would most probably customize them. The settings tell Analysis Workbench what to load, how to process information, and what ratings to give.

How do you interpret the license information in Analysis Workbench?

The Applications Details - Licensed by or Licenses information - link back to the Applications window and the Machines window, showing details for that selected application and the machine it is located on.

XML Enricher

How many XML Enrichers should be run? What determines turning on the 2nd enricher?

Generally speaking, one XML Enricher is all you need for a device inventory of about 20,000 devices. A second enricher should be introduced when you have the following conditions:

- A short inventory scanner frequency, under 7 days
- More then 20,000 inventory files being collected on a single server

The hardware specifications for Enterprise Discovery are based on running one XML Enricher. Memory and CPU requirements will need to increase if the second XML Enricher is turned on.

Reports

Can the Windows Vista Readiness reports be modified?

In Enterprise Discovery, the report function for Windows Vista Readiness reports can be exported to a Comma Separated Value (CSV) file. This file can be modified.

When Configuration Manager (CM) 5.0 has been configured with the CM Reporting server, you are able to filter information inside the Windows Vista readiness report. The CM Reporting Server has the capability to display the data collected by Enterprise Discovery.

Do Server CPU, RAM, and Disk Utilization reports require an Alarms license?

Yes, these reports require an Alarms license, but this license is not available to new customers.

Aggregator

Are there port limitations for the aggregator?

The license port limitation only applies to an Enterprise Discovery server that is not an aggregator.

Porting Data

How do you export data from the Enterprise Discovery database?

The Enterprise Discovery database has many components that can be used to export data. The traditional way is to use an HP product called Connect-It. This allows you to connect two different databases together and move the information into the tool of choice. Traditionally, Connect-It transfers the data into Asset Center. Asset Center has workflows and other processes in place to control and take action on the data being transferred. There are three typical workflow cases:

- From the HP solution point of view, importing data into the Asset Center Software Management module is the next step to allow complete tracking of licenses and optimization costs.
- Connect-It can also be used to send Enterprise Discovery data to other third party applications.

- Enterprise Discovery data can also be made available through MySQL and reports.

Using an ODBC connection to the Enterprise Discovery database will also allow you to transfer information into an MS Access database.

You can also use third party database mining tools to access the information in the Enterprise Discovery database. A list of tools can be found at the following URL:

http://www.mysql.com

What import data functionality is available in Enterprise Discovery?

- Introduced in Enterprise Discovery 2.20, there is an option for importing the Discovery Configuration from another server.
- You can perform bulk import of Device/Port characteristics and configuration per device.
- You can restore an Enterprise Discovery server from another server and then clean the discovered data. This provides a mechanism to import the entire server configuration (including accounts and other settings).

Is there an Enterprise Discovery 2.20 Connect-IT scenario available now?

Yes, it is in the Connect IT 3.8 version.

Training

Is Enterprise Discovery training material available?

Some information is available for the previous versions of Enterprise Discovery and can be found at the following URL:

http://www.education.hp.com/openview/ov_enterprise_discovery.htm

New material is being created and will be ready before the last quarter of 2007.

Is Radia delta training available?

Yes, there is Delta training. All past presentations can be found at the HP Software University.

Index