Application Discovery Manager User's Guide

vCenter Application Discovery Manager 6.1

This document supports the version of each product listed and supports all subsequent versions until the document is replaced by a new edition. To check for more recent editions of this document, see http://www.vmware.com/support/pubs.

EN-000453-00

vmware®

You can find the most up-to-date technical documentation on the VMware Web site at:

http://www.vmware.com/support/

The VMware Web site also provides the latest product updates.

If you have comments about this documentation, submit your feedback to:

docfeedback@vmware.com

VMware, Inc.
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Contents

About This Book

The VMware vCenter™ Application Discovery Manager (ADM) User's Guide describes the user interface of the ADM. It also provides information that the customers need, to manage the ADM.

Intended Audience

This document is part of the VMware vCenter Application Discovery Manager documentation set, and is intended for use by corporate information technology (IT) personnel who needs to monitor enterprise applications and resources and make decisions about acquiring, allocating, and modifying these resources.

VMware Technical Publications Glossary

VMware® Technical Publications provides a glossary of terms that might be unfamiliar to you. For definitions of terms as they are used in VMware technical documentation go to http://www.vmware.com/support/pubs.

Document Feedback

VMware welcomes your suggestions for improving our documentation. If you have comments, send your feedback to docfeedback@vmware.com.

Technical Support and Education Resources

The following sections describe the technical support resources available to you. To access the current version of this book and other books, go to http://www.vmware.com/support/pubs.

Online and Telephone Support

To use online support to submit technical support requests, view your product and contract information, and register your products, go to http://www.vmware.com/support.

Customers with appropriate support contracts should use telephone support for the fastest response on priority 1 issues. Go to http://www.vmware.com/support/phone_support.

Support Offerings

To find out how VMware support offerings can help meet your business needs, go to http://www.vmware.com/support/services.

VMware Professional Services

VMware Education Services courses offer extensive hands-on labs, case study examples, and course materials designed to be used as on-the-job reference tools. Courses are available onsite, in the classroom, and live online. For onsite pilot programs and implementation best practices, VMware Consulting Services provides offerings to help you assess, plan, build, and manage your virtual environment. To access information about education classes, certification programs, and consulting services, go to http://www.vmware.com/services.

Architecture Overview 1

The VMware vCenter™ Application Discovery Manager (ADM) is an enterprise datacenter management solution that uses agentless discovery and provides continuous dependency mapping of applications. ADM helps you gain an understanding of your service dependencies. ADM also provides automated and real-time application discovery capability across physical and virtual environments.

An accurate application dependency map is essential to virtualize applications, move datacenters, build a site recovery plan, and to move a service to the cloud.

This chapter describes different ADM architecture solutions and also describes how to configure and deploy them. Topics include:

What Does ADM Do?

ADM performs the following high-level functions:

- ADM discovers Configuration Items (CIs). It also discovers the relationships and dependencies between these instances in real time. The discovery data is stored in the Management Data Repository. The map feature in the User Interface (UI) provides a graphical representation of the relationships and dependencies between all the CIs.

  NOTE ADM provides both known and unknown information about an infrastructure. It tells you what you know, and what you do not know. For example, you might think that no laptops are communicating with a database server. Using the map, you can see the hosts that you know are in your infrastructure. However, you might also see that there are hosts or laptops communicating with the database server.

- ADM determines the baseline of usage for these configuration items. Further, the ADM detects inconsistencies from the norm.

  NOTE Baseline is the first 24 hours of activity of a configuration item.

- ADM provides impact and predictive reporting and helps you troubleshoot errors.

ADM helps to accurately answer the following questions:

- What hosts, applications, and connections do I have?
- How are CIs used?
- What are the dependencies among them?

- Where are the hidden optimization opportunities and how can I capitalize on them?
- How will these application changes impact my business?
- What are the risks?
- What are the problems and how can I solve them?

Essential Concepts

Table 1-1 defines essential concepts of ADM.

Table 1-1. Essential Concepts

Concept	Definition
Configuration Item (CI)	A CI is a host (this can also include network devices such as switches or routers), an application (also known as a service), or a network connection. For example, a Linux host, Cisco router, or VMware® ESX™ Server are all host CIs. Oracle is an application CI. HTTP and SSH are network connection CIs.
Discovery	Discovery is a continuous process that creates and maintains a detailed model of your application environment. ADM discovers CIs. Chapter 5 contains more information about discovery.
Management Data Repository (MDR)	The MDR is the database of ADM. When configuration item information is discovered, it is stored in the MDR after reconciliation.
Groups	A group is a built-in or user-defined collection of one or more CIs. The *admin* users can create groups so that they can easily view, analyze, and track the specific subset of data. Group creation allows the *admin* user to classify CIs into logical sets or collections so that actions such as creating reports, policies, or viewing the inventory is applied to specific groups, instead of the entire database of CIs. Chapter 4 contains more information about groups.
Change	A change is any change to a CI. For example, a change to a configuration file on a host is a change.
Change tracking	Change tracking refers to a change in the discovered environment, such as a change in the amount of activity on a host, a change in the demand on a host or service, or a new host or connection that is present in the discovered environment. Change tracking refers to behavioral changes. ADM allows you to create change tracking policies that track behavioral changes in your application environment.
Application pattern	ADM discovers business applications by creating application patterns. The *admin* user can create these from the **Manage** tab, or a VMware Professional Services representative can provide you these. An application pattern definition consists of a set of rules — node rules and connectivity rules — that describe the requested pattern. ADM analyzes the application pattern definition and discovers instances of the application pattern. Node rules are group-like rules that describe the instances in the topology graph. Connectivity rules describe the edges in the topology graph (that is, the connectivity between two instances).
Entity aging	ADM lets you create entity aging policies. For example, a service running in your network is populated in the MDR, and is therefore visible to you from the User Interface (UI). If you uninstall this service at a later time, you should no longer see it as a running service in the UI. Creating an entity aging policy lets you view the most updated state of your network, since it is a constantly changing environment. (An entity is a network element, service or dependency, and any of their derived elements such as a J2EE or database instance.)

ADM Components

VMware provides ADM on one or more appliances. The mode of the appliance determines which component is running. ADM components are described in Table 1-2.

Table 1-2. ADM Components

Component	Description
Active Discovery-Unix	Collects data from the configuration objects in your data center. The following discovery types apply: ■ IP discovery — Detects hosts or other configuration items with a specific IP address when passive discovery fails to discover them. ■ Detail discovery — Extends the information obtained using passive and IP discovery. It uses common network protocols to remotely query servers in the managed network and obtains supplementary information about network hosts that is added to the database.
Active Discovery-Windows	A discovery engine that uses WMI based discovery policies for performing active discovery on Windows machines.
Passive Discovery	Passively observes the network traffic by performing a deep-packet analysis to discover applications and component relationships in physical and virtual environments. Passive discovery also allows you to do the following: ■ Map dependencies. ■ Count the activity of these dependencies. ■ Identify services.
Aggregator	Receives data from the discovery components and reconciles the data before transferring it to the database component. The aggregator also provides the user interface for using ADM and is the integration point for various integrations, for example, ERDB.
Database	An Oracle RDBMS used to store discovered data and ADM configuration.

Active Probe

The active probe is the ADM process used for both Detail and IP discovery. Active probe responds to the policies defined through the management component, discovers the items assigned through the policies, and returns the data to the management component. This data is reconciled, stored, and presented in the console.

To configure the active probe

1 Assign discovery items to a specific Collector.

2 Define the protocols that are supported for discovery.

3 Connect the management component (Aggregator) to the discovery component (Collector).

The *VMware vCenter Application Discovery Manager Administration Guide* includes the steps for configuring the active probe for Detail and IP discovery.

Architecture Solutions

ADM provides a Single-box setup and Distributed setup to meet the needs of different environments. The following sections contain more information about the architecture solutions.

Single-box Setup

In a Single-box setup, all the ADM modes are enabled on a single appliance as shown in Figure 1-1.

Figure 1-1. Single-box Setup Architecture

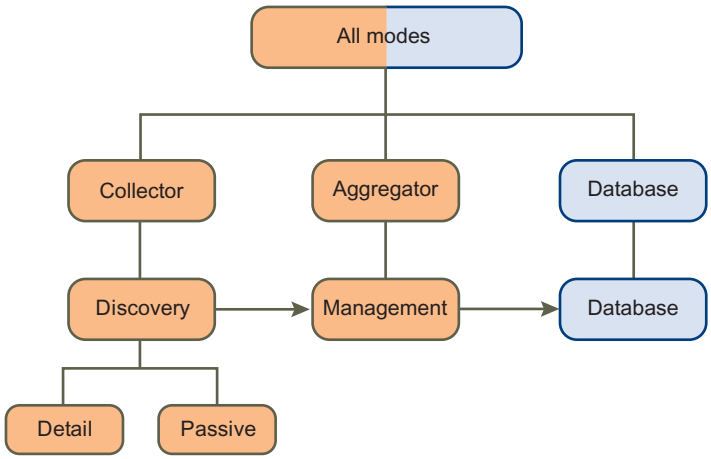

Single-box Appliance Configuration

If you are using a Single-box setup, you need to perform the steps described in *VMware vCenter Application Discovery Manager Administration Guide*, after completing the installation instructions provided in the *VMware vCenter Application Discovery Manager Appliance Platform Installation Quick Reference Guide*.

VMware vCenter Application Discovery Manager Administration Guide also describes how to move an existing Single-box setup to a Distributed setup or Distributed with remote database setup.

Distributed Appliance Solutions

A Distributed appliance solution has the following two options:

- Distributed setup
- Distributed with remote database setup

Distributed Setup

The Distributed setup has at least one designated appliance enabled as a Collector, and another appliance enabled as an Aggregator and Database as shown in Figure 1-2.

Figure 1-2. Distributed Setup Architecture

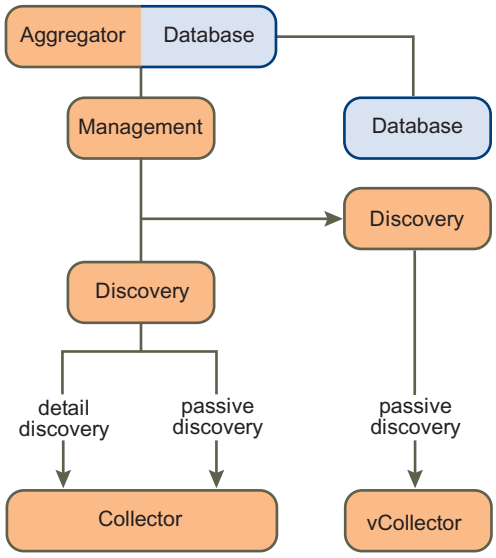

NOTE Passive and Detail Discovery can run on single or multiple Collectors.

Distributed Appliance with Remote Database Solution

In a Distributed with remote database setup, there is a designated appliance to host the database as shown in Figure 1-3.

Figure 1-3. Distributed with Remote Database (Split)

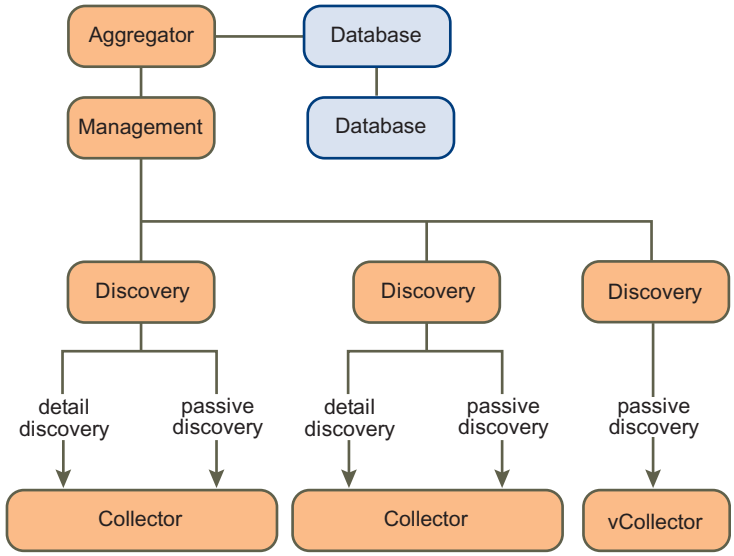

NOTE Passive and Detail Discovery can run on single or multiple Collectors.

The steps to configure distributed setup and distributed with remote database setup are described in *VMware vCenter Application Discovery Manager Administration Guide*.

Getting Started 2

This chapter describes accessing the ADM console. Topics include:

Accessing the ADM Console

The ADM Console is the Graphical User Interface (GUI) to access and manage the appliance processes. You connect to the ADM console using a browser.

NOTE Before running the ADM Console, verify that the Microsoft Windows client meets the system requirements outlined in the *VMware vCenter Application Discovery Manager Appliance Platform Installation Quick Reference Guide*.

Log In to the ADM Console

To log in to the system using the ADM Console

1. Open Microsoft Internet Explorer
2. Type the IP address of the management appliance into the URL and click **Go**. The Welcome screen appears.
3. In the **Username** field, type your user name.
4. In the **Password** field, type your password.
5. Click **Login**.

NOTE The *admin* user can log in by typing **`admin`** in the **Username** field and **`123456`** in the **Password** field.

Managing ADM 3

This chapter describes the **Manage** tab that enables you to create, edit, and delete groups, policies, fingerprints, and users. You can run discovery and manage the system from this tab. Also, it allows you to create and manage application patterns and aging policies. Topics include:

NOTE The **Manage** tab is visible only to *admin* user.

Groups

The **Manage** tab opens to the **Groups** tab by default. A group is a built-in or user-defined collection of one or more CIs. ADM provides with the ability to create groups so that you can easily view, analyze, and track the specific subset of data. Group creation allows you to classify CIs into logical sets or collections so that actions such as creating reports, policies, or viewing the inventory is applied to specific groups, instead of the entire database of CIs.

The *VMware vCenter Application Discovery Manager Online Help* provides step-by-step instructions on how to perform these actions.

IMPORTANT Only *admin* users can create groups from **Manage > Groups**. Chapter 4 provides more information on ADM groups.

Discovery

Discovery is a continuous process that creates and maintains a detailed model of your application environment. ADM discovers CIs.

In ADM, there are three types of discovery: Passive Discovery, IP Discovery, and Detail Discovery.

The *VMware vCenter Application Discovery Manager Online Help* provides step-by-step instructions on how to perform discovery.

NOTE Chapter 5 provides more information on discovery.

Application Patterns

ADM discovers business applications by creating application patterns. An application pattern definition consists of a set of rules — node rules and connectivity rules — that describe the requested pattern.

ADM analyzes the application pattern definition and discovers instances of the application pattern. Node rules are group-like rules that describe the instances in the topology graph. Connectivity rules describe the edges in the topology graph (that is, the connectivity between two instances).

The *VMware vCenter Application Discovery Manager Online Help* contains task-based information on application patterns.

IMPORTANT The *admin* user can create application patterns from the **Manage** tab, or a VMware Professional Services representative can provide these to you. Chapter 6 provides more information on application patterns.

Changes

ADM allows you to create policies that help you track application behavior and alerts you when changes occur in your environment. These changes might be demand changes, such as a higher number of transactions for a service, or behavior changes such as new clients accessing a service or new services added into the application environment.

A change tracking policy is a rule that governs what happens when a specified change occurs among the discovered items in a group. A group can have multiple policies defined for it. Change tracking policies run only after the system is in the monitoring stage, after completing discovery.

Change policies are displayed in the **Change Tracking Policies List**. To access this list, navigate to **Manage** > **Changes**. Table 3-1 lists and describes the information in the Change Tracking Policies List.

Table 3-1. Change Tracking Policies List

Column	Description
Active	This box has a green flag for an active policy and a grey flag for an inactive policy. An active policy generates change notifications.
Name	Name assigned to the change tracking policy when you created it.
Description	Change tracking policy description.
Created By	User name of the person who created this change tracking policy.
Creation Date	Date this change tracking policy was originally created on the ADM appliance where this change tracking policy is defined.
Updated By	User name of the person who last modified this change tracking policy.
Update Date	Date this policy was last modified. The date is for the ADM appliance where this policy is defined.

Click on the top of the column to sort the list by that column.

To view changes in the environment, navigate to **Change Tracking** > **Changes**. The changes that are defined in the **Manage** > **Changes** tab are actually displayed in the **Change Tracking** > **Changes** tab.

Changes are also displayed in the **Dashboard > Changes** pane.

You can perform the following actions that are accessed from the **Actions** menu on the left side of the screen:

- Add Policy — Creates a new change tracking policy.
- Copy Policy — Makes a copy of an existing change tracking policy.
- Edit Policy — Modifies an existing change tracking policy.
- Delete Policy — Removes a change tracking policy.

- Enable Policy — Enables a change tracking policy.
- Disable Policy — Disables a change tracking policy.

The *VMware vCenter Application Discovery Manager Online Help* provides step-by-step instructions on how to perform these actions.

Setting Up Scripts

You can configure a change tracking policy to run a script on the ADM appliance in response to an alert.

The ADM console can run any script that can run on a Linux-based computer. You can write a script for each change tracking policy, or the same script for several change tracking policies. When you specify a script, the change tracking policy automatically runs it whenever the selected change event occurs. For example, if you want to write a script that pages the appropriate person or group to notify them of the change event.

The script:

- Needs to be in the `/home/nlayers/Seneca/custom_scripts` directory.
- Can be any script that can run on a Linux-based computer.
- Runs in a separate process.
- Runs under the ADM user account. This account does not have root permissions. It is just a regular user account.
- Can be any script that can perform all the operations on the network from the interface. This depends on the network structure and permissions, but usually there is no internet access and limited access to other resources on the net.
- Invokes other programs installed on the appliance.

The *VMware vCenter Application Discovery Manager Online Help* provides step-by-step instructions on how to set up a script to automate a change response.

Fingerprints

Fingerprints are the core of the ADM business application discovery. They uniquely identify both packaged and custom developed in-house applications.

Identifying Application by Ports

Fingerprints enable you to identify custom applications by the ports they use. When a custom service or connection that matches a defined fingerprint is discovered, it appears with the service/connection name that you supplied, instead of an unclassified service or connection. When you navigate to **Manage > Fingerprints**, a list of all defined fingerprints appears.

The fingerprints list contains the information for each fingerprint, as shown in Table 3-2.

Table 3-2. Fingerprints List

Column	Description
Port	Port that this service/connection uses.
Transport	Method of transport that the port could use, either TCP or the User Datagram Protocol (UDP).
Protocol	Communications protocol your service/connection uses.
Service	Name of the service that you created to use this port.

You can perform the following actions from the **Actions** pane on the left side of the screen:

- Add Fingerprint — Creates a new fingerprint.
- Edit Fingerprint — Modifies an existing fingerprint.

- Delete Fingerprint — Removes a fingerprint.

The *VMware vCenter Application Discovery Manager Online Help* provides step-by-step instructions on how to perform these actions.

Aging

ADM allows you to create entity aging policies. Aging is the removal of an inactive entity and its owned entities from ADM. An inactive entity is a network element, service, or dependency, and any of their derived elements, such as a J2EE or database instance that have not been verified as live for some predefined amount of time. For example, a service running in your network is populated in the MDR, and is therefore visible to you through the ADM console. If this service is uninstalled at a later time, it should no longer be shown as a running service in the ADM console. Creating an entity aging policy allows you to view the most updated state of your network, since it is a constantly changing environment.

Aging is performed through aging policies. When you create a new aging policy, the default time limit is seven days. You can change this default as described in the *VMware vCenter Application Discovery Manager Online Help*

The Aging Policies list contains general information about existing aging policies. To access the Aging Policies list, select **Manage** > **Aging**. The Aging Policies list is shown below.

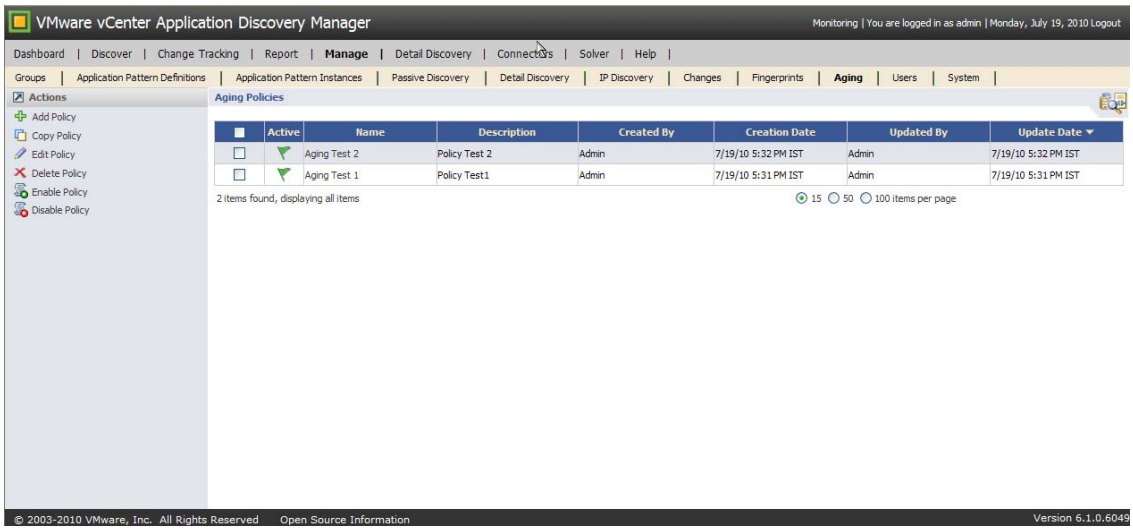

Click on the top of a column on the **Aging Policies** list to sort the list by that column.

You can perform the following actions with aging policies:

- Add Policy — Creates a new aging policy.
- Copy Policy — Makes a copy of an existing aging policy.
- Edit Policy — Modifies an existing aging policy.
- Delete Policy — Removes an aging policy.
- Enable Policy — Enables an aging policy.
- Disable Policy — Disables an aging policy.

The *VMware vCenter Application Discovery Manager Online Help* provides step-by-step instructions on how to perform these actions.

Table 3-3 lists and describes the aging policy information in the Aging Policies list.

Table 3-3. Aging Policy Information

Column	Description
Active	This box has a green flag for an active policy and a grey flag for an inactive policy.
Name	Name assigned to the aging policy when it was created.
Description	Description of the aging policy.
Created By	User name of the person who created this Aging Policy.
Creation Date	Date this aging policy was originally created. The date is for the ADM appliance where this Aging Policy is defined.
Updated By	User name of the person who last modified this Aging Policy.
Update Date	Date this Aging Policy was last modified. The date is for the ADM appliance where this Aging Policy is defined.

Deleting Aging Logs

To delete the aging logs, navigate to **Manage > System** and then click **Delete All Aging Logs** from the **Actions** pane on the left side of the screen.

Users

In ADM, there are two types of users: administrators and operators. If you log in as an operator, you do not have access to the **Manage**, **Detail Discovery**, and **Connectors** tabs. If you log in as an administrator, you have access to all the tabs.

On the **Discover** > **Inventory** page, only an administrator can perform the following actions:

- Delete — Removes a selected item entirely.
- Add to Group — Adds a selected item to a group.
- Remove from Group — Removes a selected item from a group.

Selected items consist of hosts, services, and devices.

Only administrators can add, copy, edit, delete, enable, and disable authorized system users. Navigate to **Manage** > **Users** to perform these actions:

- Add User — Creates a new user definition.
- Copy User — Makes a copy of an existing user definition.
- Edit User — Modifies an existing user definition.
- Delete User — Deletes a user definition.
- Enable User — Enables a user definition.
- Disable User — Disables a user definition.

The *VMware vCenter Application Discovery Manager Online Help* provides step-by-step instructions on how to perform these actions.

Role-Based Access Control

ADM provides role-based access control. This allows you to assign permission to a role instead of directly assigning permission to a user. ADM roles define the basic permission level for operations that users assigned to the role can perform. When there is a large amount of data in your environment, role-based access control helps to discover just the information that you might be interested in.

You specify role-based access control when you add a user. You must select the operator role, select the **Enable Role Based Access Control** check box, and select one or more existing groups. This ensures that the operator account has access only to certain groups. Only those groups are displayed for the operator in ADM.

The *VMware vCenter Application Discovery Manager Online Help* provides step-by-step instructions on how to enable role-based access control.

System

The administrator can perform several management functions on an ADM system. Navigate to **Manage** > **System** to perform the following functions:

- Self Test — Performs internal tests and looks for errors.
- System Reboot — Reboots the ADM appliance.
- System Log — Displays the system log file, which shows the time of each event, log message, and severity. For example, services restart is an event.
- Delete All System Logs — Deletes all the system log files.
- Aging Log — Displays the aging log files, which shows the time of each event, log message, and severity. For example, removal of an entity is an event.
- Delete All Aging Logs — Deletes all the aging log files.
- Support Package List — Displays the list of all created support packages.
- Create Product Support Package — Creates a product support package and adds it to the Support Package List.
- Mail Configuration — Sets your mail server and address.
- Active Probes Configuration — Configures active probes. *VMware vCenter Application Discovery Manager Administration Guide* provides information about adding and configuring an active probe.
- Advanced Configuration — Sets parsing of configuration files. Use only when an IT Compliance Analyzer- Application Edition (ITCA-AE) appliance is connected to ADM.
- Update — Updates to be done through the command line interface using the ADM Appliance Platform. The *VMware vCenter Application Discovery Manager Administration Guide* provides more information about updating ADM using ADM Appliance Platform.
- Restart Discovery — Restarts discovery.
- Licensing — Displays license information and allows you to upload a new license.

The *VMware vCenter Application Discovery Manager Online Help* provides step-by-step instructions on how to perform these actions.

Licensing

You must have an ADM license to initiate discovery process.

IMPORTANT Only the Administrator can view and upload the license from the **Manage** tab.

To view existing licenses

1 From the ADM console, navigate to **Manage** > **System**.

2 In the **Actions** left pane, click **Licensing** to display the **License properties** page.

The following details are displayed for existing licenses:

- **License Feature** — The type of license.
- **Amount Licensed** — The number of servers registered in the license.
- **Amount Used** — The number of servers already discovered.
- **Expiration Date** — The date on which the license is scheduled to expire.
- **Serial Number** — The 25 digit serial number string.

NOTE If the number of discovered servers exceed the number of servers registered in the license, the following notification appears in the **Discover** > **Inventory** page:

Maximum discovered Servers exceeded! Discovery may be incomplete. Please obtain additional licenses from your VMware sales representative.

Upload a License

A newly installed ADM setup is not licensed by default. Before you begin, obtain the license from the VMware sales representative.

To upload a license

1 From the ADM console, navigate to **Manage** > **System**.

2 In the **Actions** left pane, click **Licensing** to display the **License properties** page.

3 Click **Upload a new license**.

4 Type the serial number provided by VMware in the text box and click **Apply**.

NOTE If the serial number string is invalid, an error message is displayed immediately.

Viewing License Details

After you upload a license, you can view the license details, including the license feature, the amount licensed, the amount used, the expiration date, and the serial number from the **License properties** page.

To view license details

1 From the ADM console, navigate to **Manage** > **System**.

2 In the **Actions** left pane, click **Licensing** to display the **License properties** page.

Groups 4

This chapter discusses groups in ADM. Topics include:

Overview

ADM is capable of discovering thousands of configuration items on a single appliance. The *admin* user can create groups so that they can easily view, analyze, and track the specific subset of data. Group creation lets the *admin* user classify configuration items into logical sets or collections so that actions such as creating reports or policies, or viewing the inventory are applied to specific groups, instead of the entire database of configuration items.

A group is a built-in or user-defined collection of one or more hosts, services, J2EE applications, database instances, hypervisors, or virtual hosts. There are three types of groups: View, Business Application, and Cluster.

ADM administrators can perform the following actions in this tab:

- Add Group — Adds a new group.
- Copy Group — Opens a dialog box to copy the selected group.
- Edit Group — Opens a dialog box to modify the selected group.
- Delete Group — Deletes the selected group (after confirming).
- Refresh Groups — Refreshes the selected group against the ADM database.

  An automatic nightly refresh process synchronizes existing groups with information in the ADM database.
- Import — Imports the group (previously saved or exported as an XML file).

  If a group with the same name already appears in the Groups List the words `Copy of` is appended to the beginning of the group name.
- Export — Exports group as an XML file.

The *VMware vCenter Application Discovery Manager Online Help* provides step-by-step instructions on how to perform these actions.

IMPORTANT Only *admin* users can create groups from the **Manage > Groups** page. The **Manage** tab is visible only to *admin* users.

Requirements

Other than being an *admin* user, group creation has no other prerequisites. No passive or detail discovery must be performed prior to group creation, and defining the group does not require that specific CIs have already been discovered.

Group Refresh

If you have not yet started discovery, then no CIs are populated in the MDR and any predefined groups do not contain any services or hosts. After the MDR is populated, however, services or hosts are added to the group.

When you create a group, you can decide whether to automatically refresh the group. The default setting is that the group will refresh automatically.

Built-In Groups

The following groups are built into ADM:

- Microsoft SQL Servers
- MySQL databases
- IIS servers
- VMware ESX Servers
- Jboss servers
- VMware virtual machines
- Routers
- Sybase databases
- Oracle databases
- Switches
- WebLogic servers
- Hosts running Windows
- Apache servers
- Websphere servers
- Tomcat servers
- DB2 databases
- Hosts running Linux

User-Defined Groups

Navigate to **Manage** > **Groups** and click **Add Group** to view the Group Definition screen and create a group.

An ADM group is shown in Figure 4-1. To create a group, select a primary object and then define one or more rules that apply to that primary object. A primary object is a host, service, J2EE application, J2EE module, or database instance. Based on the primary object you select, the fields for defining the rules change. Table 4-1 lists these options.

Figure 4-1. ADM Group

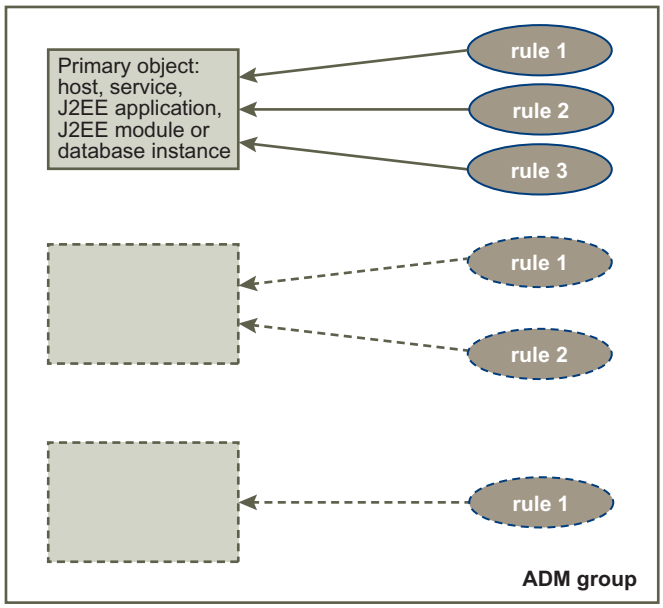

A group consists of one or more primary objects. Each primary object has one or more rules, which specify more details about the primary objects that compose your group. Individual rules can be included or excluded from the group definition.

NOTE Since you have to select a primary object to create a rule, you can create rules for different primary objects and thus create a group that consists of multiple CIs. For example, you can create a group that consists of both Linux hosts and Oracle databases.

Table 4-1. Group Definition Parameters for Each Primary Object

Primary Object	Parameters
Host	Basic: ■ all hosts in scope ■ which are of type ■ running services ■ with incoming protocols connection ■ with outgoing protocols connection ■ having services accessed by URL ■ having host configuration Extensions: ■ with its connected clients ■ with its connected servers ■ with its virtualized virtual machines ■ with its hypervisors
Service	Basic: ■ all services ■ running on hosts in scope ■ with incoming protocols connection ■ with outgoing protocols connection ■ being accessed by URL ■ running on host having configuration Extensions: ■ with its connected clients ■ with its connected servers ■ with its virtualized virtual machines ■ with its hypervisors
J2EE application	Basic: ■ all J2EE applications ■ running on services ■ where services running on hosts in scope ■ with incoming protocols connection ■ with outgoing protocols connection Extensions: ■ with its connected clients ■ with its connected servers ■ with its virtualized virtual machines ■ with its hypervisors

Table 4-1. Group Definition Parameters for Each Primary Object (Continued)

Primary Object	Parameters
J2EE module	Basic: ■ all J2EE modules ■ deployed on J2EE applications ■ running on services ■ where services running on hosts in scope ■ with incoming protocols connection ■ with outgoing protocols connection Extensions: ■ with its connected clients ■ with its connected servers ■ with its virtualized virtual machines ■ with its hypervisors
DB instance	Basic: ■ all DB instances ■ running on services ■ where services running on hosts in scope ■ with incoming protocols connection ■ with outgoing protocols connection ■ having DB tables Extensions: ■ with its connected clients ■ with its connected servers ■ with its virtualized virtual machines ■ with its hypervisors

When you select a parameter, it appears in the Rule Editor and becomes a hyperlink. When you click the hyperlink, you can define the parameter.

The *VMware vCenter Application Discovery Manager Online Help* contains information on using the Group Definition screen.

Importing and Exporting Group Definitions

You can import or export group definitions as XML files from **Manage** > **Groups**. To import a group definition, click **Import**. To export a group, select it from the list of groups and click **Export**.

Discovery 5

This chapter describes the Discovery functionality and configuration. Topics include:

Discovery Types

Discovery is the process of populating ADM's management data repository with CIs and identifying the relationships between them. In ADM, there are three types of discovery: Passive Discovery, IP Discovery, and Detail Discovery. Table 5-1 defines each discovery type.

Table 5-1. Types of Discovery

Discovery Types	Definition
Passive Discovery	Passive Discovery is the process in which network traffic is listened to passively. Passive Discovery is a non-intrusive process where you can specify an IP range to search for hosts and applications.
IP Discovery	IP Discovery is the process that detects hosts or other devices with a specific IP address when passive discovery fails to find them. An ICMP or TCP connection scan detects devices that are active but not a source or destination of network traffic, such as switches or routers.The IP scan discovers the devices and adds them to the inventory list. You can create IP discovery policies and set them to run periodically.
Detail Discovery	Detail Discovery is the process that provides the granular details of hosts and services (such as OS information, the installed software list, disk size, configuration file settings, and so on) that are not available with passive discovery.

All the details — the configuration items as well as all their hardware and software configuration information — that are discovered through these three discovery methods are displayed on the **Discover > Inventory** page. Clicking on a host, service, or device will display its properties.

Discovery Plans

A discovery plan helps you to control the depth of information discovered. For example, software changes frequently and hardware does not change often. Therefore, you might want to discover software information more often than hardware information. A discovery plan allows you to define exactly what will be discovered, and therefore improves performance. ADM lets you include or exclude both the passive and detail discovery of specific hosts, services, and connections with the creation of passive discovery plans and detail discovery plans. "Passive Discovery" on page 32 contains information on passive discovery plans, and "Detail Discovery" on page 33 contains information on detail discovery plans.

Passive Discovery

Passive Discovery detects hosts, services, relationships and dependencies. Further, it detects the use of each relationship, extracts some basic properties (that is, URLs, table names, version numbers), and resolves IPs to host names.

Policies

Start Passive Discovery after creating a Passive Discovery policy in which you specify an IP address range to search for traffic and hosts. Navigate to **Manage > Passive Discovery** to view the Passive Discovery Policy Definition screen and create a policy. The online help contains procedural information for Passive Discovery tasks.

Plans

You can create a Passive Discovery plan from the **Plan** tab within the Passive Discovery Policy Definition screen.

To create a Passive Discovery plan

1 Select the discovery plan rules. You have an option to discover, not to discover and ignore the following:

- Services
- Protocols
- Ports
- Service categories

2 Select or clear **Discover behavior** option.

3 Click **Update.**

Passive Discovery plans are useful when you want to include or exclude the passive discovery of certain hosts, services, or connections. This will also improve performance. Scenarios in which you might want to create a Passive Discovery plan are as follows:

- The SSH protocol is noisy and clutters the MDR without providing any value. Since you are not interested in discovering the SSH protocol passively, you can specify this in the Passive Discovery plan.
- For licensing control, you want to include or exclude only certain discovery services such as BEA OEM.

IP Discovery

IP Discovery is a method for detecting hosts or other configuration items with a specific IP address when passive discovery fails to find them. For example, if a host is powered down or if it is outside the specified IP range, passive discovery fails to find it. In this case, you can use IP discovery with a TCP or ICMP connection to find CIs. IP discovery policies are created from the **Manage > IP Discovery** page.

Policies

To create IP Discovery policy

1 Click **Add Policy** on the IP Discovery Policies page.

2 Type the name of the plan in the **Name** field.

3 Type the description.

4 Select or clear the different options available on the three tabs - **General**, **Scheduling**, and **Scope**. Type the information required in these screens.

5 Click **Create**.

VMware vCenter Application Discovery Manager Online Help provides the detailed steps for creating an IP Discovery policy.

Detail Discovery

Detail Discovery is a method to find more granular details, such as hardware and software configuration information that are not available with passive discovery. Detail Discovery extends the information obtained using ADM's Passive Discovery technology and uses common network protocols to remotely query servers in the managed network. Using these protocols, you can obtain supplementary information about network hosts and add it to the MDR. You can view the additional configuration information found by detail discovery in the properties screen for any host, service, or device through the **Discover > Inventory** page. All other detail discovery information and administrative tasks, such as creating detail discovery policies, are done through the **Detail Discovery** tab as shown below.

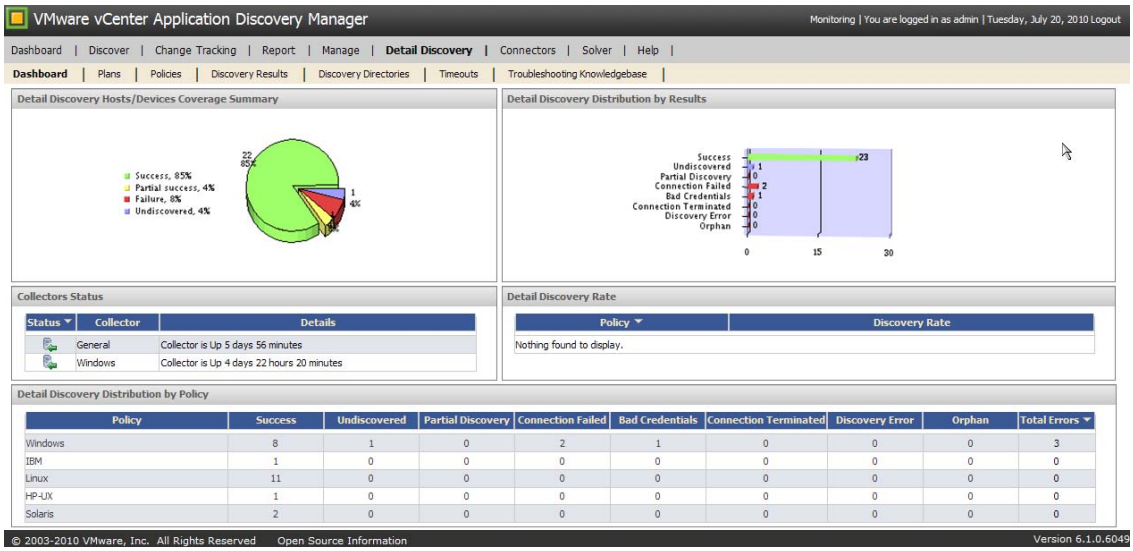

Detail Discovery Tab

Dashboard

The **Detail Discovery > Dashboard** page displays two graphs and three tables. Dashboard is a visual tool that lets you understand detail discovery status, process, and problems to be resolved. You can see the system state and detail discovery policies. Dashboard summarizes the detail discovery status of the system and enables you to analyze detail discovery status (success, failure, and so on) in different views. The following graphs and tables are available:

- Detail Discovery Hosts/Devices Coverage Summary
- Detail Discovery Distribution by Results

- Collectors Status
- Detail Discovery Rate
- Detail Discovery Distribution by Policy

IMPORTANT The **Detail Discovery** tab is visible only to the *admin* user.

Plans

You can create Detail Discovery plans from the **Detail Discovery > Plans** page. This page displays a list of detail discovery plans, including three built-in plans:

- Shallow — Discovers operating systems and network configuration.
- Medium — Discovers everything (OS, hardware, software) except services and network connections.
- Deep — Discovers everything, except network connections.

Detail discovery plans are useful when you want to include or exclude the detail discovery of certain hosts, services, or connections. Navigate to **Detail Discovery > Plans** > **Add Discovery Plan** to create your own plan. Scenarios in which you might want to create a detail discovery plan are as follows:

- The software in your environment changes frequently and the hardware hardly changes. For Detail Discovery, you want software discovered daily, but hardware information is discovered only once a week.
- You want to exclude the Detail Discovery of hosts with specific host names or operating systems.

NOTE Every Detail Discovery policy has a plan associated with it. **Deep** is the default discovery plan for Detail Discovery policies. Changes to a discovery plan are effective the next time the CIs within the policy scope are discovered.

The Discovery Plan Definition screen is shown below.

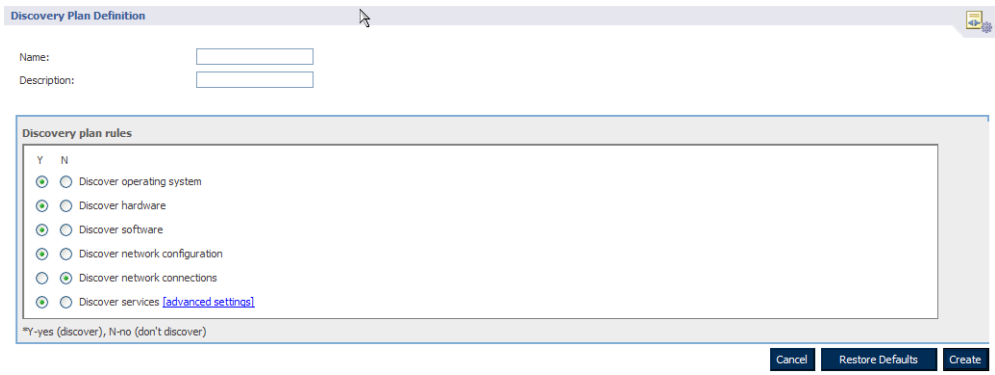

Policies

You can initiate Detail Discovery after creating a Detail Discovery Policy from the **Detail Discovery > Policies** page. To maximize security and minimize the risk of impact on the managed network, only *admin* users control the setup and scheduling of detail discovery policies.

Detail Discovery policies describe a query that runs periodically to collect network information. Each policy defines the following parameters:

- The protocol to use (SNMP, SSH, Telnet, WMI, or VI-SDK), as well as the authentication parameters required for the specified protocol
- A schedule for running the policy
- The scope (a named group or IP address range) of hosts to query

Discovery Results

The **Detail Discovery > Discovery Results** page displays the latest detail discovery results and allows you to create or modify a detail discovery policy, or view the report card for a policy.

Report Card

You can use the report card to view information about discovery, such as the total discovery time, detailed status information, protocols used, discovery stages and their times, and much more. You can also use the report card as a troubleshooting utility. You can view and resolve errors directly from the report card as well as create a support package for the specific report card. The report card is useful in the following scenarios:

- When you want to find out more detailed information about the discovery process such as search strategies and directories, or the time taken to complete the full discovery.
- For troubleshooting discovery errors.
- When discovery is reported as successful, but a CI seems to be missing or has an incorrect value, and the user wants to investigate this further.

Discovery Directories

The **Detail Discovery > Discovery Directories** page allows you to specify or customize the directories that ADM will search when it searches for applications. ADM then searches the directories added through this interface in addition to all the default directories through which it searches.

NOTE The discovery directories specified in this tab apply only to detail discovery.

The following is an example in which you would use this page:

1 The default Oracle installation on UNIX is located at `/opt/oracle`, but in your environment, Oracle is installed on a different location such as `/opt/applications/oracle`.

2 You need to add `/opt/applications/oracle` as a discovery directory and can do so through the **Detail Discovery > Discovery Directories** page.

Timeouts

The **Detail Discovery > Timeouts** page lets you configure the timeout values for different protocol methods. Each protocol used for discovery (that is, SNMP, SSH, Telnet, WMI, or VI-SDK) uses a different method (such as running a remote command in an SSH session or executing a remote query with WMI) in order to acquire information. A timeout is enforced on the execution of each protocol method. This page allows you to modify the timeout values. Typically, this is for advanced users.

The following is an example in which you would configure the timeout values:

1 You perform a discovery and receive a timeout error in the **Discovery Results** tab of the report card. Following error message is displayed:

```
Method “running remote shell command” has timed out. Current timeout is 10 minutes.
```

2 Click the resolution link that redirects you to the **Detail Discovery Timeouts** table in the **Detail Discovery > Timeouts** tab.

3 Change the timeout for “Run a shell command” from 10 minutes to 20 minutes.

Troubleshooting Knowledgebase

Navigate to **Detail Discovery > Troubleshooting Knowledgebase** page for troubleshooting scenarios. The following is an example in which you would use this tab:

1 You encounter detail discovery problems or errors.

2 Your Customer Support Representative cannot resolve the problem. Engineering requires more information and creates a specialized knowledge base file in order to debug the problem.

3 You upload the file from the **Detail Discovery > Troubleshooting Knowledgebase** page.

NOTE Troubleshooting knowledge base applies only to create a detail discovery support package until the file is removed. It does not customize the detail discovery knowledge base permanently.

Deploying Detail Discovery

This section describes how to deploy Detail Discovery.

IMPORTANT ADM can merge data from multiple hosts in clustered environment that were assigned the same IP or secondary IP (such as clusters and balancers) as if only a single host were discovered. You need to inform your Customer Sales Representative if your network contains more than one host using the same IP so that ADM is configured to treat each host separately.

Detail Discovery Configuration

ADM has an extensive detail discovery feature that enhances and complements the results obtained through Passive Discovery. Detail Discovery is performed by communication between hosts in a managed network, using four common network protocols: SNMP, SSH, Telnet, and WMI. Detail discovery is also performed using the discovery protocol VI-SDK, which is a VMware proprietary API used to query model of VMware, control the behavior of virtual infrastructure, and receive notifications on changes in the virtual environment.

Using these protocols, supplementary information about network hosts is obtained and added to the Configuration Management Database (CMDB). To maximize security and minimize the risk of impact on the managed network, the administrative user of the ADM Console can tightly control the setup and scheduling of the Detail Discovery queries to the various network hosts.

How Do Active Probe Configurations Affect Detail Discovery Policies?

For Detail Discovery to be successful in a Distributed setup, the hosts to be discovered, needs to be included in both a Detail Discovery policy and an Active Probe configuration. The Detail Discovery policy includes a list of items to include in the Detail Discovery, while the Active Probe configuration assigns the specific hosts on which each Collector can perform Detail Discovery.

Detail Discovery is not be performed on a host that is included in a Detail Discovery policy if it is not included in an Active Probe configuration. However, the host is still recognized through Passive Discovery and will appear as Orphaned in the Detail Discovery dashboard because the Active Probe configuration has no correlation with Passive Discovery.

Detail Discovery Policies

Detail Discovery is controlled and configured through user-defined Detail Discovery policies. There can be many policies, each describing a *query* that is running periodically and collecting network information. For each policy, the following parameters are defined:

- The discovery scope, that is, the set of network hosts that is queried for information.
- The protocol being used, which is one of the following: SNMP, SSH, Telnet, WMI, or VI-SDK.
- Protocol-specific authentication and communication parameters, such as usernames and passwords to use for authentication.
- Scheduling information — How often to access the network hosts covered by the policy. The scheduling scheme lets ADM users control and balance two conflicting factors: the need to have the most updated picture, and the need to avoid overloading the network with too many Detail Discovery queries.
- Discovery plan — A discovery plan allows you to define exactly what will be discovered, and therefore improve performance. You can create passive or Detail Discovery plans that can include or exclude the discovery of specific hosts, services, and connections.

Configuring Standard OS Agents

To avoid agent installation, ADM performs Detail Discovery using only standard network protocols. Support for these protocols is built into most modern operating systems, and so no special software needs to be installed. However, a moderate amount of configuration work might be needed to make network hosts respond to Detail Discovery queries made by ADM. Whether and how much configuration work is needed depends on the operating systems used, their existing configuration, and the standard network protocol used for Detail Discovery. "Detail Discovery Protocols" on page 37 describes the standard protocols and the necessary configuration.

Dealing with Firewalls

When deploying Detail Discovery, firewalls, when placed between the ADM appliance and the hosts that it is discovering, might block the network protocols used for Detail Discovery, and thus prevent Detail Discovery from taking place. The following solutions deals with the firewall:

- Open the necessary ports in the firewall, if only just for client-side use by the IP address assigned to the ADM appliance. The relevant ports are typically:
  - For SNMP, UDP/161
  - For Telnet, TCP/22 for SSH and TCP/23
  - For WMI access, "WMI Deployment Recommendations" on page 41 and "Firewall Settings " on page 41 include details on firewall settings.
  - For VI-SDK, TCP/443 for HTTPS
- Deploy another Collector appliance on the other side of the firewall. This extra device communicates with the Aggregator appliance. This communication uses standard HTTPS (port 443) or HTTP (port 80) and has to be open in the firewall for connections initiated by the Collector into the Aggregator. The default is HTTPS on port 443. There might be multiple Collectors installed at various locations in the network.

Checking for Results

After Detail Discovery policies have been defined, check the Detail Discovery dashboard for the discovery status of each policy and each host. Click Detail Discovery to view the dashboard.

Detail Discovery Protocols

This section describes the network protocols used for Detail Discovery: SNMP, SSH, Telnet, WMI, and VI-SDK.

The *VMware vCenter Application Discovery Manager Discovery Coverage Spreadsheet* contains a list of systems on which ADM has been tested. ADM supports the systems that are listed in this document.

Discovery data obtained from more than one protocol is reconciled according to the priority order below. More information on detail discovery protocols is provided for:

1 "SSH" on page 38

2 "SNMP" on page 39

3 "WMI" on page 40

4 "Telnet" on page 43

5 "VI-SDK" on page 45

6 Listener (passive)

7 IP Discovery (passive)

SSH

Secure Shell (SSH) is a standard protocol for secure remote access to UNIX-like operating systems. SSH servers are built into most Linux distributions, Mac OS-X, Sun Solaris, OpenBSD, and most other UNIX-like operating systems. SSH servers from various vendors are also available for Windows.

Remote access to a host that runs an SSH server starts by authenticating the client's identity. After the client identity is authenticated, an encrypted communication channel opens. The client can then examine files and run commands on the server host. The privileges and permissions of the client are determined by the server according to its identity. For example, if the client uses a guest account with few privileges, most of the information is not available to this user.

SSH has two versions. Version 2 is normally in use, while version 1 is less recommended. Since SSH clients and servers automatically detect each other's versions and coordinate their communications, no action is required.

Detail Discovery with SSH

ADM uses SSH to access hosts that run SSH servers, and to obtain information about the operating system, hardware, and software installed on the server host.

Both SSH versions 1 and 2 are supported automatically with no user interaction. Authentication is based on specifying a user name and password to use when accessing the managed hosts; these are stored by ADM internally in an *encrypted* form.

SSH Server Deployment Recommendations

Firewall Settings

SSH queries are normally performed on TCP port 22 on the server. If a firewall exists between the ADM appliance and the monitored network, this port needs to be open for connections initiated by the ADM appliance.

SSH Server Settings

Discovery with SSH of servers running the OpenSSH server (sshd) requires that the "PasswordAuthentication" field contain the value "yes" in the server settings file (often, `/etc/ssh/sshd_config`). In some operating systems, such as SuSE, the default is "no" and needs to be changed for the SSH discovery to complete.

Credentials

Detail discovery with SSH is based on accessing the managed host with a predefined user name and password. For more information on necessary privileges, download the document *discovery_coverage.xls* from:

http://downloads.vmware.com/Application Discovery Manager

IMPORTANT It is not recommended to use the user "root" for security reasons.

If ADM is used to discover configuration of services such as application servers, databases, and web servers, this user might need more read privileges if the configuration files of these services are not accessible by ordinary users.

For example, in some sites, the Oracle database server is installed and run with a special "oracle" user belonging to a special "oracle" group. The configuration files for the server might only be readable by users in the "oracle" group. Having ADM use a user in this group would allow it to access these files and retrieve valuable and detailed configuration information that is otherwise unavailable.

A similar scenario might also occur with other types of servers, depending on how they are installed. However, often this is not an issue: for example, in the default installation of the Apache web server under Red Hat Linux, all configuration information is stored in a location that is readable by the general public (under the /etc branch of the file system). In such cases, no group memberships are required for ADM to be able to read this detailed configuration.

SNMP

The Simple Network Management Protocol (SNMP) is a popular and standard protocol for remotely monitoring and managing various types of network nodes. Managed network nodes are often regular servers, but other network devices such as network switches and network printers can also be managed using SNMP.

SNMP is based on the notion of agents running on the managed network nodes. An agent is a software component, installed on the managed node that can answer remote queries about the state of the managed node. The remote component that makes such queries is termed the SNMP manager. A typical SNMP deployment includes many SNMP agents installed on the various managed nodes, and a single SNMP manager that collects information from all of them.

The main differences among the three versions of SNMP (1, 2, and 3) are the security mechanisms. Versions 1 and 2 rely on a fairly primitive mechanism of **community strings**, each defining a different set of SNMP operations that is performed. Version 3 introduces more advanced authentication and privacy mechanisms, based on usernames and passwords. SNMP agents are built into most modern operating systems, but might need to be turned on or configured to be able to provide relevant information. The information that is obtained from the SNMP agent is defined in modules called Management Information Base (MIB). The core MIB, which is available in most agents, is called MIB-2, and it supplies system and hardware information.

Detail Discovery with SNMP

ADM can act as an SNMP manager and collect information from any host that has an SNMP agent running on it. All versions of the SNMP protocol are fully supported. For versions 1 and 2, community strings are used. For version 3, you can select the authentication and privacy modes in compliance with this newer standard.

SNMP Agent Deployment Recommendations

Firewall Settings

By default, SNMP queries are performed on UDP port 161 of the agent, although this can be changed if desired. If there is a firewall between the ADM appliance and the monitored network, this port needs to be open for connections that are initiated by the ADM appliance.

Linux and Net-SNMP

The SNMP agent that is built into Linux distributions is Net-SNMP (http://net-snmp.sourceforge.net). This agent runs as a service called “snmpd” and is located in the services directory `/etc/init.d/`.

The default Net-SNMP configuration allows the use of the **public** community string with SNMP version 2, to query the SNMP agent for particular system configuration items. However, this default configuration only allows access to a portion of the standard MIB-2 information base. Specifically, it does not allow querying the list of network interfaces, which is a very important piece of information.

To allow Net-SNMP to also report this missing information, it is recommended that you modify the Net-SNMP configuration file in /etc/snmp/snmpd.conf. Add the following line to the section of the file that has lines starting with “view”:

```
view    systemview    included   .1.3.6.1.2.1.2
```

Windows

Windows 2000 is usually installed with its own SNMP agent. If it is not, it is quickly installed by selecting: **Control Panel > Add/Remove Programs > Add/Remove Windows Components > Management and Monitoring Tools > Details > Simple Network Management Protocol**.

By default, this server supports the “public” community string for querying system information.

Solaris and HP-UX

Solaris and HP-UX systems do not include built-in SNMP agents. You can download and install Net-SNMP from http://net-snmp.sourceforge.net and configure it as in Linux.

BEA WebLogic Agent

The BEA WebLogic application server comes with its own SNMP agent that is built into its installation.

To enable the BEA WebLogic application

1 In the WebLogic management console, select **Services > SNMP**.

2 Select **Enabled**.

3 Restart the application server.

NOTE If there is another SNMP agent running on the same machine, such as the native agent of the operating system, it is recommended to change the port used by the WebLogic agent. In the same location in the management console, set the port to the desired port.

WMI

Windows Management Instrumentation (WMI) is a proprietary Microsoft technology for modeling, querying, and managing the configuration of Windows hosts. WMI follows a public modeling and management standard known as Common Information Model (CIM), as well as another related standard called Web-Based Enterprise Management (WBEM).

The WMI software component is built into all Server editions of the Windows operating system since the Windows 2000 Server. It might or might not be installed by default as part of Windows XP, but it is easily installed there as an add-on.

WMI is modular and extendable: common information about the host is obtained with the basic built-in WMI module. Additional components called **WMI providers** is installed to model and query in detail the configuration of services such as IIS Server, Active Directory, BizTalk server, and so on.

The WMI component in Windows is based on Microsoft Component Object Model (COM) technology, and is queried both locally and remotely. Remote queries are through RPC access to the WMI component, using the remote access flavor of the COM technology known as Distributed Component Object Model (DCOM).

Detail Discovery with WMI

ADM can perform Detail Discovery using the WMI protocol. When creating a WMI Detail Discovery policy, you need to specify a user name, password, and domain name. These are used by the WMI component to authenticate and authorize access to the host information.

NOTE When using a non domain user to perform WMI discovery, specify "WORKGROUP" in the domain field.

WMI discovery is used for discovering machines that run the following operating systems: Windows 2000 Server, Windows Server 2003, Windows 2008, and Windows XP SP2.

IMPORTANT The following steps are new to ADM 6.0 and later versions as it discovers more information compared to 5.3. If you have already set up your servers for WMI discovery using ADM 5.3, you need to perform the additional steps listed under "Setting Execute Permissions for Used Executables" on page 43.

IMPORTANT The permissions required to complete the same WMI operations might vary between different versions of Windows and different Service Packs installed. Some windows versions such as Windows 2003 Server with SP2 require an account with local administrator permissions in order to successfully complete all queries performed by ADM.

WMI Deployment Recommendations

Creating a User for WMI Detail Discovery

Using WMI to query remote hosts for their configuration details requires appropriate privileges, as described next. To easily manage these privileges, it is recommended to use a separate domain user for this purpose. Therefore, the first step in deploying WMI Detail Discovery is to create a domain user account. This user should not have any special administrative privileges. In fact, there is no reason for it to belong to any groups at all.

In the event that a local administrator user is used instead of a specially created user, it is important that DCOM configuration allows remote access and launch for administrator users. Troubleshooting tips regarding WMI and DCOM permissions is found in the article at:

http://blogs.technet.com/askperf/archive/2007/08/14/wmi-troubleshooting-permissions.aspx

You need to create a profile and temporary folder on all machines where Detail Discovery is to be performed by logging in to those machines.

If a local user is used rather than a domain user, follow the instructions in "Configuring the Windows Telnet server" on page 44 regarding local security policy settings.

Firewall Settings

WMI queries involve the Microsoft RPC network protocol that uses dynamically-assigned ports on the server side, and is therefore quite firewall-unfriendly. To avoid firewall trouble, it is recommended to deploy the Detail Discovery, Collector appliance in the same network as the managed hosts without a firewall between them.

If there must be a firewall between the Management, Aggregator appliance and the Detail Discovery, Collector appliance, it should be configured to allow RPC traffic. This is done in two stages:

1 Configure the managed hosts to use a narrow range of dynamic ports for their RPC. The following URLs provide further information:

```
http://msdn2.microsoft.com/en-us/library/ms809327
http://msdn.microsoft.com/library/default.asp?url=/library/en-us/dndcom/html/msdn_dcomfirewall.asp
```

2 In the firewall settings, open TCP port 135 (the RPC Service Control Manager port), in addition to the full range of RPC ports specified in Step 1, for access by the Detail Discovery, Collector appliance.

Disabling Internal Firewall for Windows XP Service Pack 2

IMPORTANT Windows XP with Service Pack 2 has a built-in internal firewall that might block incoming RPC/DCOM requests.

The internal firewall should be turned off or partially disabled to allow direct connection to the local network.

To change the firewall configuration

1 Go to **Control Panel > Security Center > Windows Firewall**.

2 To fully disable the firewall, in the **General** tab, select **Off**.

3 If you want to leave the firewall enabled but still allow RPC/DCOM communication, select **On** in the **General** tab, and in the **Advanced** tab, clear local network.

Setting DCOM Privileges

In the following steps, it is assumed that the domain name is MYDOMAIN and that the user used for WMI Detail Discovery and that domain is named DOMAINUSER.

Since WMI access to a Windows host involves DCOM technology, the DOMAINUSER needs to be allowed to perform DCOM operations on each managed host. This is already the default setting in most Windows servers (Windows 2000 and 2003 server families), but not in Windows XP or in servers that had their defaults changed.

To enable DCOM for a host running Windows XP or Windows 2003

1 Log in to the host as *admin* user (either a local administrator or a domain administrator).

2 Click **Start** menu, select **Run** and type `DCOMCNFG`. **Component Services** window appears.

3 From the **Console Root** left pane, select **Component Services** > **Computers** > **My Computer**.

4 Right-click **My Computer**, and then click **Properties**. **My Computer Properties** window appears.

5 Select **COM Security** tab.

NOTE In Windows 2000, the DCOM management console is simpler. So skip Step 3 and Step 4.

6 Under **Launch and Activation Permissions**, click **Edit Limits**. **Launch Permission** window appears.

7 If the domain user (DOMAINUSER in this example) is not present in the **Groups or user names** list, Click **Add**, to add the domain user.

8 Select the user you added from the **Groups or user names** list.

9 From the **Permissions for Administrators** list, select **Remote Launch** and **Remote Activation**.

10 Clear all other checkboxes and click **OK** to return to the **My Computer Properties** window.

11 Click A**pply** to save your changes and click **OK**.

Setting WMI Privileges

After WMI requests pass through the DCOM communication layer, Windows runs another authorization check, for specific user permissions at the level of the WMI service. Therefore, to allow the domain user to run WMI queries, the WMI service needs to be configured for every managed host. This is done either locally on that host or remotely.

To configure WMI access on the managed hosts

1 If you are configuring a remote machine, use a domain administrative account. A local administrative account will not work.

2 Log in as *admin* user (either local or remote administrator) to the target host.

3 Click **Start** menu, select **Run** and type `wmimgmt.msc`. The **Windows Management Infrastructure (WMI)** window appears.

4 Right-click **WMI Control (Local)** in the left pane and click **Properties**. **WMI Control (Local) Properties** window appears.

NOTE If you are configuring remote settings for WMI privileges, right-click **WMI Control (Local)** and select **Connect to another computer**. Type the name of the remote host and click **OK**. From this point, configuration changes will be applied to the selected host instead of the local host.

5 Click **Security** tab and navigate to **Root > Security** in the namespace tree. **Security for Root** window appears.

NOTE By selecting the **Security** option in **Root**, you are allowing WMI queries to all available WMI providers. In case Root cannot be used, **CIMV2** should be used; this option is not recommended. Giving the ADM user security permissions to Root provides little to no risk that queries will be nonintrusive.

6 Click **Add**. The **Select Users or Groups** window appears.

7 Type the user name in the following format:

Domain name\user name

For example, `MYDOMAIN\DOMAINUSER`

NOTE Type the user name of the user who will be performing the Detail Discovery.

8 Click **OK** to return to the **Security for Root** window.

9 Ensure that the newly added user name is highlighted in the **Groups or user names** list.

10 Click **Advanced**. The **Advanced Security Settings for Root** (**Root** in some systems) window appears.

11 Select the newly created user name from the list and click **Edit** (**View/Edit** in some systems).

12 Set **Apply onto** to **This namespace and subnamespaces** from the drop-down menu.

13 Select **Remote Enable** from the **Permissions** list and set it to **Allow**.

14 Clear all other checkboxes and click **OK**.

15 Continue to click **OK** until all of the dialog boxes are closed and then close the **Windows Management Infrastructure (WMI)** window.

16 Click **Yes** if you encounter the following message:

Save console settings to winmgmt?

Setting Execute Permissions for Used Executables

ADM 6.0 and later versions discover more information than previous versions. To discover this additional information, for each managed server, on each of these files (cmd.exe, cscript.exe, and netstat.exe located in the system32 folder where Windows is installed), perform the following steps:

1 Right-click the file and click Properties.The **Properties** dialog box appears.

2 Select the **Security** tab.

3 From the **Group or user names** list, select the user who will be performing the Detail Discovery.

4 Select **Read & Execute** and **Read** from the **Permissions for user** list, to grant the necessary permissions.

5 Click **OK** to confirm.

Telnet

The Telnet protocol is one of the oldest and most common protocols for remote shell access. However, in recent years it is replaced in many cases with the SSH protocol, which encrypts its network traffic and is considered more secure. Still, some network devices, such as network routers and switches, support remote access through Telnet exclusively. Additionally, Microsoft Windows has a built-in Telnet server, and does not have a similar SSH server. Therefore, Telnet is used by ADM for Detail Discovery, similarly to the use of SSH.

Detail Discovery with Telnet

In general, Detail Discovery with Telnet is supported by any machine running a Telnet server that:

- Supports the terminal type known as *dumb*.
- Either allows simple command-line authentication or accepts NTLM authentication.

Telnet Server Deployment Recommendations

This section includes Telnet recommendations for deployment.

Firewall Settings

Telnet queries are normally performed to TCP port 23 of the target device. If there is a firewall between the ADM appliance and the monitored network, this port needs to be open for connections initiated by the ADM appliance. Specifically, in Windows XP Pro SP2, the internal firewall must be turned off for Telnet discovery to take place.

Credentials

As with SSH, Detail Discovery with Telnet is based on accessing the managed host with a user name and password that it recognizes. The considerations regarding the choice of user account and privileges are the same as those for SSH, described in "Detail Discovery with SSH" on page 38.

Configuring the Windows Telnet server

Certain operating system settings must apply for a Windows host to be accessible with Telnet. Depending on the specific edition of Windows and on the existing configuration, any of the following configuration modifications might be necessary.

IMPORTANT Ensure you have local administrator permissions on the machines you are performing these procedures on.

To start up the Windows Telnet Services automatically using Windows

1 From the Windows **Start** menu, navigate to **Setting** > **Control Panel** > **Administrative Tools > Services**.

2 Locate and right-click **Telnet** service.

3 Ensure the startup type is **Automatic** and start the service if it is not already started. This change is required in most Windows editions.

To start up the Windows Telnet Services automatically using a command line interface

- On a local machine, type:

  ```
  sc config TlntSvr start= auto && sc start TlntSvr
  ```

- On a remote machine where COMPUTER is the remote computer name or IP address, type:

  ```
  sc \\COMPUTER config TlntSvr start= auto && sc \\COMPUTER start TlntSvr
  ```

To log in to a machine using Telnet, you need to be listed as a member of either the local TelnetClients group on that machine, or as a member of the domain's TelnetClients group.

To modify users and groups using Windows

1 From the Windows **Start** menu, navigate to **Setting** > **Control Panel** > **Administrative Tools** > **Computer Management** > **Local Users and Groups** > **Groups**.

2 If the **TelnetClients** group exists, double-click **TelnetClients**. **TelnetClients Properties** window appears.

3 Click **Add** to add relevant user as its members.

4 If the **TelnetClients** group does not exist, create a new group with the name **TelnetClients**, and then add the user to it.

Creating a Local TelnetClients Group

If you want to add a user to the local TelnetClients group, but no such group exists yet, you can simply create a new group by this name. This operation is automated using VBScript or Jscript. Use the following commands:

```
computer = "COMPUTER"
user="USER"
domain="DOMAIN"
Set objGroup = GetObject("WinNT://" & strComputer & "/TelnetClients")
Set objUser = GetObject("WinNT:// " &  domain & "/" & user)
objGroup.Add(objUser.ADsPath)
```

IMPORTANT In some editions, particularly XP Pro SP1 and later, remote access by local users is always treated as if the "guest" user is involved. This extra security measure might cause Telnet (and also WMI) to fail with local users, but it has no effect on domain users.

To change this behavior for operating as a local user

1 From the Windows **Start** menu, navigate to **Setting** > **Control Panel** > **Administrative Tools** > **Local Security Policy** > **Security Settings** > **Local Policies** > **Security Options**.

2 Locate **Network access: sharing and security model for local accounts** policy from the list.

3 Right-click **Network access: sharing and security model for local accounts** and click **Properties**. **Network access: sharing and security model for local accounts Properties** window appears.

4 Select **Classic – local users authenticate as themselves** from the drop-down menu and click **Apply**.

For security reasons, it is recommended to allow the Telnet server to use only NTLM authentication, and to disable password-based authentication.

NOTE You must have local administrator permissions to use NTLM.

To view the current Telnet server authentication settings

1 From the Windows **Start** menu, select **Run** and type `cmd`. The **Command Prompt** appears.

2 Type `tlntadmn` to view the local host settings.

3 Type `tlntadmn config sec=+ntlm` to turn on NTLM authentication.

4 Type `tlntadmn config sec=-passwd` to turn off password-based authentication.

VI-SDK

VI-SDK is a VMware proprietary API used by third-party applications to query VMware's model, control the virtual infrastructure's behavior, and receive notifications on changes in the virtual infrastructure environment. VI-SDK also allows you to discover, configure, and monitor all aspects of VMware ESX servers and Virtual Center.

VI-SDK on Virtual Center accesses information about the entire Virtual Infrastructure deployment, whereas VI-SDK on an ESX only access information about that ESX.

VI-SDK is implemented by standard web services with a published WSDL and runs over HTTPS by default. A VI-SDK URL is the URL of the VI-SDK server on the target host. This URL is used on all hosts in the policies scope so it does not contain the server part of the URL. For example, if the URL is /sdk:925 and discovery is done against host 1.2.3.4, the actual URL used to connect to the host is https://1.2.3.4/sdk:925. By default, the URL field is initialized to /sdk, which is the default VI-SDK URL.

NOTE The VI-SDK reports the IP address of a virtual machine only if VM-tools is running on the virtual machine's guest OS.

Detail Discovery with VI-SDK

ADM uses VI-SDK to access a Virtual Center to obtain information about ESX Server and the virtual machines that are installed on them.

VI-SDK Deployment Recommendations

For VI-SDK to work properly, the Web Access Option must be enabled in the Virtual Center. VMware recommends that you assign read-only permission to the user account that is created for Detail Discovery.

Firewall Settings

VI-SDK queries are performed over HTTPS. If a firewall exists between the ADM appliance and the monitored network, TCP port 443 must be open for connections that are initiated by the ADM appliance.

Limitations

VI-SDK does not expose some information and therefore ADM cannot discover some CIs. Specifically, the following CIs will not be discovered through VI-SDK:

- Services
- Configuration files
- License files
- Installed software
- Operating system kernel related information (For Example, kernel version)

Discovering Dependencies with Detail Discovery

This section provides the necessary information to decide the method to discover dependencies and to configure Detail Discovery to discover dependencies.

ADM Dependency Discovery Methods

Discovering dependencies is done through ADM using either of the following methods: Passive and Detail Discovery together or only using Detail Discovery.

Passive Discovery assumes that a packet represents a dependency between two hosts, for example, there's most likely a good reason that host A sent a packet to host B. The shortcoming of this assumption in Passive Discovery might not identify the service on the client-side of a connection (client-server communication usually contains much more information on the server side than on the client side) correctly. As a result, ADM might rely on Passive Discovery for discovering dependencies, while leveraging Detail Discovery to correctly identify the service on the client side of the connection.

However, some clients have regulatory, business, or infrastructural barriers that prevent them from using Passive Discovery. To penetrate these clients, ADM allows you to correctly identify dependencies using Detail Discovery only, without relying on Passive Discovery.

The Process of Dependency Discovery Using Only Detail Discovery

To discover dependencies using Detail Discovery, ADM must leverage the capabilities exposed to it by the protocol it uses to connect remotely to the interrogated host. ADM uses a cross-platform and widely supported program called netstat for remote shell protocols (for example, SSH and Telnet) and WMI. Since WMI does not expose any port-related information, netstat is used to run commands on the target Windows machine. With SNMP, ADM interrogates a standard MIB2 table that exposes information about open TCP connections and ports used by applications and services running on the interrogated host. ADM then uses heuristics to guess the protocols used by those connections.

NOTE ADM uses heuristics to guess the protocols only for ports that are lower than 512.

Table 5-2 describes what ADM discovers about dependencies and how.

Table 5-2. Dependency Discovering Methodology

What	How
Discover the open connections that the interrogated host has to other hosts on the network (much like passive).	By querying the services exposed by the ADM protocol, as described previously.
For each of those open connections, discover the service that is using the connections.	ADM finds the process ID (PID) of the services running on the host, and matches up that PID with the PID of each open connection. Note: There is a limitation with some major OS platforms (SunOS, AIX, and HP-UX) for which netstat is unable to retrieve PID information. On these platforms, this capability is disabled.

Table 5-2. Dependency Discovering Methodology (Continued)

What	How
Discover the open ports that the services running on the interrogated host are using to listen for incoming connections ("service endpoints").	Both netstat and SNMP tables expose listening ports that is used to create a service endpoint with that port. An example of this is discovering that an Apache service is listening on ports 80 and 8080, even if no active connection exists at the time of the discovery.
Guess the top-level protocol used by those connections.	To avoid false positives, it is done only on low ports: <512.

Choosing a Method of Dependency Discovery

Table 5-3 outlines the differences in the information that are discovered through either Passive Discovery (PD) or Detail Discovery (DD) to help you determine which type to use in your environment.

Table 5-3. Comparison of Passive and Detail Discovery Information

	Advantage	
Difference	**PD**	**DD**
Detailed discovered dependencies do not include activity, whereas passively discovered dependencies do.	✓	
Protocol identification is by far more accurate with Passive Discovery. This of course, results from Passive Discovery's specialization in protocol analysis.	✓	
Passive Discovery is weak in discovering the source of a connection for reasons explained in "ADM Dependency Discovery Methods" on page 46.		✓
Passive Discovery cannot discover the ports on which a service is listening unless a client sent a packet to it.		✓
Detail Discovery discovers connections that are active at the time of discovery, whereas Passive Discovery samples all communication traffic on the network. This means that ephemeral connections have less of a chance to be discovered through Detail Discovery. Note: However the connections that are active and representing an interaction with a live business application are not likely to be ephemeral.	✓	
Only Detail Discovery discovers *documented dependencies* that are dependencies discovered by looking at the configuration of service, such as in the files and registry.		✓

Note the following when you are using Passive or Detail Discovery for discovering dependencies:

- If the same connection is discovered through both Passive and Detail Discovery, the connection is reconciled to appear as a single connection; for example, if Passive Discovery discovers the protocol, activity and the server-side service, and Detail Discovery discovers the client-side service, the two discoveries would be reconciled to include all the information collected by both discovery types, without redundancy.
- The process of reconciling hosts might take some time.
- ADM uses the same Passive Discovery scope IP filters to filter remote hosts (hosts connected to the interrogated host) discovered during Detail Discovery. This feature avoids the problem of Detail Discovery overriding the IP ranges that were excluded as part of the Passive Discovery scope.
- By default, Detail Discovery policies do not discover network dependencies due to issues surrounding performance. The discovery of network dependencies substantially increases the amount of time it takes to reconcile the discovered results, and since the default deployment of ADM includes Passive Discovery, this default configuration still provides a full view of the network, including network dependencies.
- vCollector support Passive Discovery only.

Discovery Strategy for Using Only Detail Discovery for Dependencies

Here is one suggested strategy to use when creating a Detail Discovery plan for discovering dependencies:

1 Ensure that you set up the Passive Discovery scope with the IP ranges of the hosts with the dependencies you want included or excluded from discovery.

2 Create a Detail Discovery policy (or set of policies, depending on the discovery protocol) with a Shallow plan that will run frequently (for example, all hosts once a day) to quickly scan the network for minimal OS and networking information.

3 Create a Detail Discovery policy (or set of policies, depending on the discovery protocol) with a Deep plan that will run less frequently (for example, discover a host once every few days) than the policy created in Step 1.

4 Once the policies in Step 2 and Step 3 have discovered a substantial part of the network, and the rate of new discovery decreases:

   a Create a custom discovery plan that has the **only Network Connections** enabled.

   NOTE The **only Network Connections** option is disabled by default.

   b Create a new Detail Discovery policy that runs frequently and apply the custom discovery plan created in the previous step to it.

   The frequent SHALLOW scan, in Step 2, serves two purposes. First, hosts with more than one IP are merged to appear as a single host. Second, Detail Discovery policies will be tailored to match the discovery protocol with the OS of the hosts in their scope.

   This less frequent, DEEP policy, created in Step 3, is used to retrieve deep configuration information of the environment.

   The Network Connections policy, defined in Step 4, will discover only network connections, and do so after the new discovery rates have decreased. This is important because Network Connection plans can have a performance penalty, which is the reason that discovery of network dependencies is excluded from the DEEP discovery plan by default.

NOTE This phased approach creates a delay of a few days to discover network dependencies, because option 3 is enabled only after the discovery rate decreases. This approach is used to avoid the performance penalty caused by using Detail Discovery to discover dependencies (which is the reason that discovery of network dependencies is excluded from the DEEP discovery plan by default). By starting to discover network dependencies only *after* much of the environment has been discovered by Detail Discovery, the performance penalty is minimized. Alternatively, if there is an immediate need to see network dependencies sooner rather than later, there is the option of creating a custom discovery plan that includes network dependencies and running it immediately.

VMware Discovery

VMware products such as ESX Server are used to create the virtual machines in the form of a set of configuration and disk files that together perform all the functions of a physical machine. Through the virtualization platform, you run the virtual machines, install operating systems, run applications, and configure the virtual machines. This includes identifying the virtual machine resources, such as a storage device.

VMware Terminology Overview

Virtual Center monitors and manages components of your virtual and physical infrastructure. These components are as follows:

- Virtual machines — A virtualized x86 personal computer environment in which a guest operating system and associated application software can run. Multiple virtual machines can operate on the same managed host machine concurrently.
- Hosts — The physical computers on which the virtualization platform software, such as ESX Server, is installed. They are hosts to the virtual machines.

  NOTE A Virtual Center host is the physical machine on which the virtual machines are running. All virtual machines within the VMware Infrastructure environment are physically on ESX Server hosts. The term "host" in this Help system refers to the ESX Server host that has virtual machines on it.

- Resources — Selected resources belonging to the host and assigned to the virtual machines that are resident on that host. The managed resources are CPU, memory, disk space, and I/O. Virtual Center uses the resources to provide, through the VMware DRS components, various options for monitoring the status of the resources and adjusting or suggesting adjustments to the virtual machines.
- Virtual Infrastructure Client (VIC) — The UI used to connect to the Virtual Center Server.

VMware Discovery in ADM

ADM can discover VMware environment, track changes in the environment, and find dependencies within that environment. ADM can discover the VMware components and CIs in Table 5-4.

Table 5-4. Discovered VMware Components and CIs

Component (CI type)	CI	Discovery Method
Virtual Center (service)	■ Product, vendor, and version. ■ Documented dependencies to all ESX servers attached to that VirtualCenter (only with VI-SDK).	Telnet, WMI, VI-SDK
Virtual Infrastructure Client (service)	Product, vendor, and version.	Telnet, WMI
VMware ESX Server	■ All standard Linux-based hardware and software. ■ For each virtual machine configured on it, the virtual machine configuration. ■ Generic ESX configuration. Virtual infrastructure version.	VI-SDK (partial), Telnet, SSH, SNMP
Virtual machine configuration	Configuration of a virtual machine on an ESX Server: ■ The virtual machine configuration (content of .vmx file, license, and state of the virtual machine (running or stopped)). ■ Whether VM-Tools is installed. ■ Information on the virtual machine hardware as configured on the ESX Server (memory, CPU, physical drives, and so on). ■ Information on the virtual machine OS, as configured on the ESX Server (OS type and family).	Telnet, SSH, VI-SDK
Virtual machines	Standard host discovery.	All methods (excluding VI-SDK).

Table 5-4. Discovered VMware Components and CIs (Continued)

Component (CI type)	CI	Discovery Method
Player (service)	Product, vendor, and version.	Telnet, SSH, WMI
VM-Tools (installed software)	Product, vendor, and version. Note: VM-Tools is an attribute of the guest OS, not the virtual machine configuration.	Telnet, SSH, WMI

Use Case

The following use case describes the typical flow for passive and detail discovery of a VMware environment.

Assumptions

Assumptions for both passive and detail VMware Infrastructure discovery are as follows:

- The VMware Infrastructure components (VMware Infrastructure 3 or higher, VMware Virtual Center and VMware ESX Server) are properly configured and operational. VMware Virtual Center is used for management.
- A web interface or Virtual Infrastructure Client (VIC) is used for interfacing with VMware Virtual Center. (A VIC is a front-end UI client used to administer Virtual Center. It is either a Windows application or a web application.)
- One of the following protocols is allowed in the VMware Infrastructure environment:
  - SSH/Telnet/SNMP on the VMware ESX Server.
  - VI-SDK with either Virtual Center or all of the ESX Server in the VMware Infrastructure environment.
- You have the credentials to log in to the target machine using the SSH, Telnet, SNMP, or VI-SDK protocols, and these credentials allow you to access information about virtual machines.
- For detail discovery of VMware environments, it is assumed that Virtual Center hosts have already been discovered with passive or IP discovery.

VMware Infrastructure Discovery Flow

Table 5-5 describes the typical flow for VMware Infrastructure discovery.

Table 5-5. VMware Infrastructure Discovery Flow

Step		Description	Results
1	Passive discovery (not mandatory)	A passive discovery policy is created from the **Manage > Passive Discovery** page.	Since the VMware Infrastructure is discovered passively, this will be only a partial discovery. You will be able to view the following information: ■ The VMware Infrastructure environment as a whole—the ESX Servers and services (VirtualCenter, Virtual Infrastructure Client). ■ Connections between ESX Servers and VirtualCenter. ■ The clients that are managing VirtualCenter (that is, Virtual Infrastructure Client, web browser). ■ Connections between Virtual Infrastructure Client and VirtualCenter, all labeled with VMware's proprietary protocol, VMDB. ■ Third-party software integrated with VMware Infrastructure. ■ Virtual machines.
2	Detail discovery	A detail discovery policy is created from the **Detail Discovery > Policies** page. When entering the protocol type during detail discovery policy creation, use any or all of the following protocols: ■ SSH/SNMP on the ESX Server included in the VMware Infrastructure deployment. ■ WMI on the host running VirtualCenter. ■ VI-SDK on any or all of the VMware ESX Servers and hosts running VirtualCenter.	The VMware Infrastructure is discovered fully. You will be able to view additional information that was discovered after passive discovery: ■ Detailed configuration information of all ESX Servers and services (VirtualCenter), including configuration information and license information. ■ For each ESX Server—all the virtual machines defined on it, including the detailed hardware configuration information for each virtual machine.
3	Application pattern creation (optional)	After performing passive or detail discovery, you might decide that you want ADM to automatically detect instances of VMware Infrastructure environments in your network. To do this, create an application pattern from the **Manage > Application Pattern Definitions** page.	ADM automatically detects instances of VMware Infrastructure environments in your network and will populate the MDR with the relevant data.
4	Virtual Collector/vCollector deployment (optional)	Virtual machines running on the same physical host communicate with each other without going out to the network. Since ADM listens to network traffic, it will not have visibility into this communication. In this case, the vCollector must be deployed.	Details on vCollector deployment is found in the *VMware vCenter Application Discovery Manager Version 6.0 vCollector Installation and Configuration Guide.*

Using VI-SDK for Detail Discovery

ADM includes the VI-SDK protocol as an option when creating Detail Discovery policies. VI-SDK is a VMware proprietary API used by third-party applications to query a VMware model, control VMware Infrastructure's behavior, and receive notifications of changes in the VMware Infrastructure environment. It also allows its user to discover, configure, and monitor all aspects of VMware ESX Server and Virtual Center. The VI-SDK option is shown below.

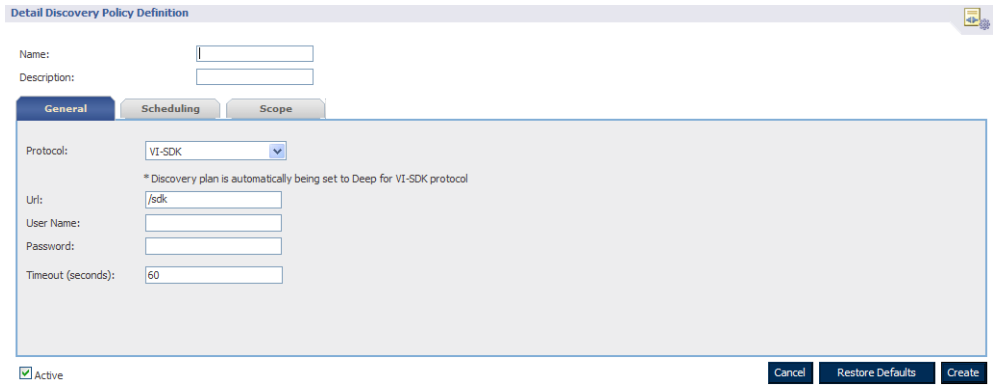

Capabilities

Once a VMware environment is discovered, you can perform several functions.

Viewing Virtualization Dependencies

You can view dependencies in your VMware Infrastructure environment by selecting the relevant items from the **Discover > Inventory** page and clicking **Virtualization** in the **Dependencies** panel. This option allows you to view dependencies for the selected object. For example, you can select a virtual machine from the inventory and then click **Virtualization** to view its containers. You can also view dependencies in the map.

Virtualization-Related Search

After discovering your VMware Infrastructure environment, you can perform searches on it. For example, you can search for virtual machines or VMware ESX Server. To perform a search, go to the **Discover > Inventory** page and click **Search**. This option is shown below.

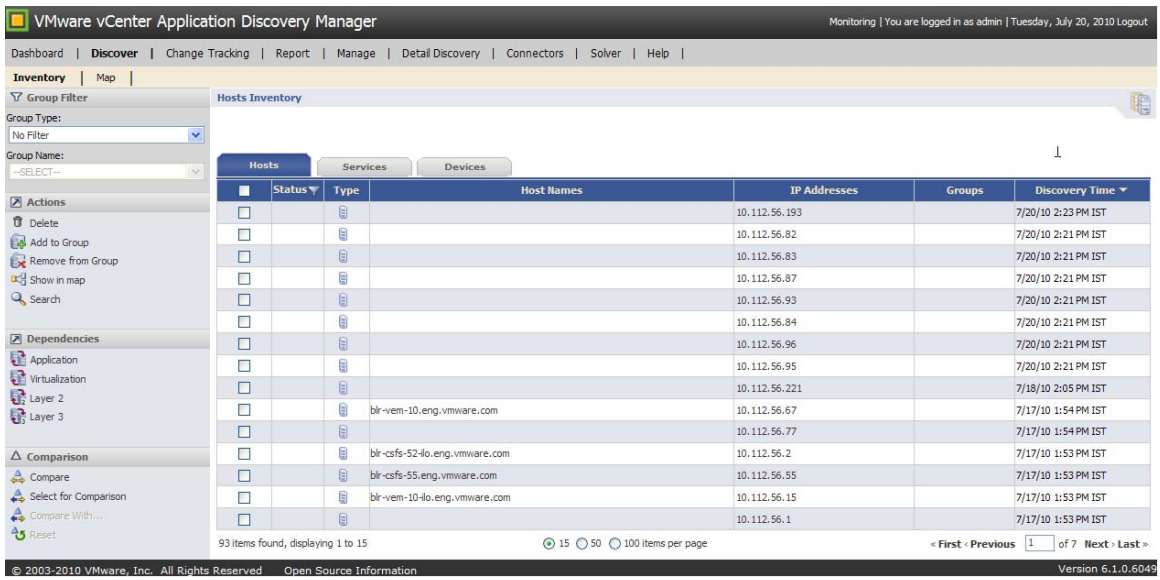

Finding Dependencies Between a VMware Infrastructure and Virtualized Business Applications

Assumptions

You have completed active discovery of the VMware Infrastructure environment at least once, and discovery of the virtual machines in the VMware Infrastructure at least once.

Goal

After detecting VMware Infrastructure environment instances in your network, you want to find out which business applications (for example, PeopleSoft) are on these instances.

Flow 1: Foundation to Virtualized Environment

To view your VMware environment

1 Create a group that contains the VMware ESX Server, Virtual Center, and the Virtual Infrastructure Clients.

2 To view all the virtual machines in the VMware Infrastructure environment, do either of the following:

- Extend the VMware Infrastructure environment group to include the environment.
- View the environment through the map (**Discover > Map**) or the inventory (**Discover > Inventory**).

Chapter 4 contains more information on groups.

Flow 2: Virtualized Environment to Foundation

To view your VMware environment

1 Create a group that contains all the virtual machines in the VMware Infrastructure environment.

2 To view the VMware Infrastructure environment, do any of the following:

- Automatically extend the group you created to include its Virtual Infrastructure environment.
- View the environment through the map (**Discover > Map**) or the inventory (**Discover > Inventory**).

Application Patterns

6

This chapter describes how to create application pattern definitions and view the results as application pattern instances. Topics include:

Overview

Data centers run business applications that handle the core business and operational data of the organization. These business applications typically consist of several hosts running databases, application servers, file servers, and various other components. ADM provides application patterns that allows you to identify and group together entities that comprise the various instances of a business application.

Creating application patterns helps you to easily follow any changes in a specific business application. The instance is updated automatically if there are server changes or a change in services in the application, thus relieving you from having to manually track changes.

Application Pattern Definitions

To use Application Patterns, you must first create an Application Pattern definition. An Application Pattern definition is a topology defined by a set of endpoints (such as hosts, services, and databases) and the relations (connections) between them. Figure 6-1 demonstrates an example of an application pattern that is a Jboss Server, connected to an Oracle database through a Java Database Connectivity (JDBC) connection.

Figure 6-1. Application Pattern Example

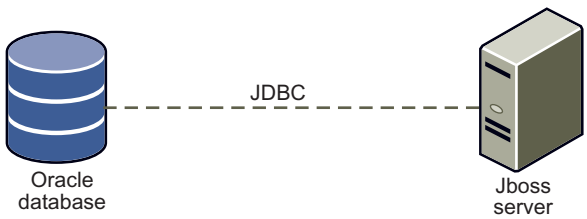

ADM provides the following types of rules for identifying Application Patterns:

- Node rules
- Connectivity rules

Node rules identify the endpoints of the application pattern. There are two types of node rules:

- "Mandatory Node Rules" on page 56
- "Unifying Node Rules" on page 57

Connectivity rules identify the connections between the nodes. Connectivity rules also assign each node as a source or target of the application pattern definition.

NOTE ADM requires that you define at least two node rules, one for the source and one for the target, plus one connectivity rule for each application pattern definition.

Node Rules

In Figure 6-2, there are three endpoints of the application pattern:

- Oracle database
- Jboss Server
- Apache Server.

Figure 6-2. Application Pattern Endpoints

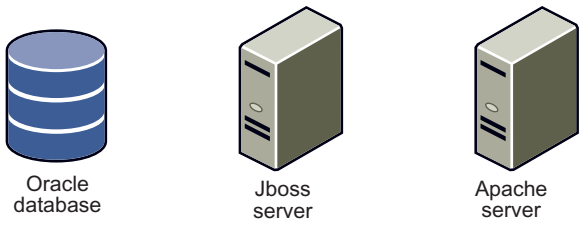

This application pattern definition requires three node rules to identify each endpoint.

Connectivity Rules

Connectivity rules identify the connections between the endpoints and are used to assign an endpoint as a source or target of the application pattern.

Figure 6-3 demonstrates a connectivity rule that includes JDBC and HTTP connections, as well as assigns the Apache Server as a source to the Jboss Server (target) and the Jboss Server as a source to the Oracle database (target).

Figure 6-3. Connections Between Endpoints

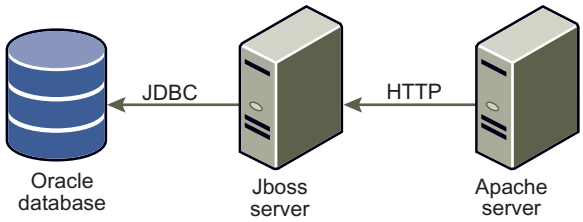

Mandatory Node Rules

Application pattern definitions also require that you define a node rule either as:

- Mandatory
- Optional

Mandatory elements must exist in the environment to be identified as an instance. They are core elements of the application pattern definition.

Optional elements are not core and, if they exist, they are included in the discovered instance.

Figure 6-4 shows an application pattern definition that contains both mandatory and optional elements.

Figure 6-4. Mandatory and Optional Elements in an Application Pattern Definition

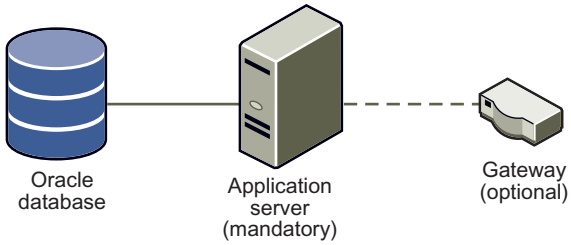

Unifying Node Rules

The endpoints of an application pattern instance might be shared by other entities. Unification allows you to identify all entities that share a resource as a single application pattern instance.

To help identify all similar instances that share the same components, you can specify a node-based rule as a unification rule. When the rule is applied, it results in one instance being detected, instead of several that contain all the same components.

For example, you might not know how many ADM Collector components exist in an instance of an ADM application. You could end up with multiple instances instead of just one, as shown in Figure 6-5.

Figure 6-5. Multiple Instances Sharing the Same Components Identified Without Unification

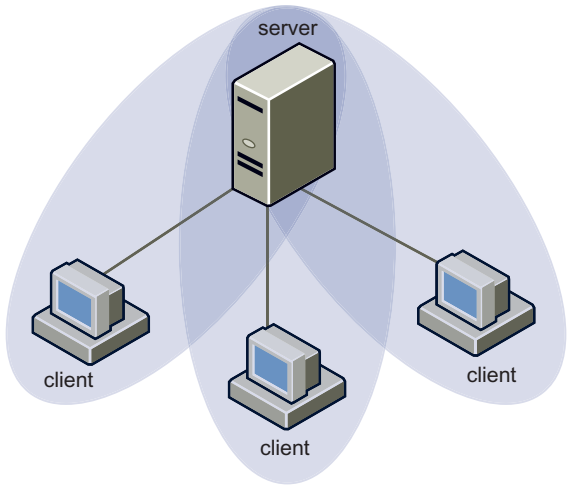

When you specify a node rule as a unification rule, the application pattern definition unites all aggregator-collector instances that share the same ADM aggregator into one instance as shown in Figure 6-6.

Figure 6-6. Similar Instances Are Identified as One Instance With Unification

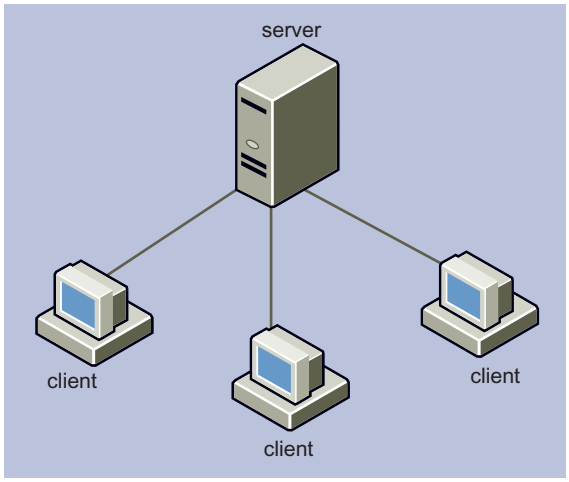

Application Pattern Instances

ADM analyzes newly created application pattern definition and discovers instances of the application pattern. An application pattern instance is a set of components (hosts, services, and so on) and their relations that answer an application pattern definition. This definition is applied against the repository thus resulting in a display of all application pattern instances that match that definition.

Viewing Application Pattern Definitions and Instances

The application pattern definitions and instances are viewed and managed through the **Manage** tab. The *VMware vCenter Application Discovery Manager Online Help* contains field descriptions and information about using the interface.

Application Pattern Definitions

The Application Patterns List contains general information about existing application pattern definitions. To access the Application Patterns List, navigate to **Manage** > **Application Pattern Definitions**. The Application Patterns List is shown below.

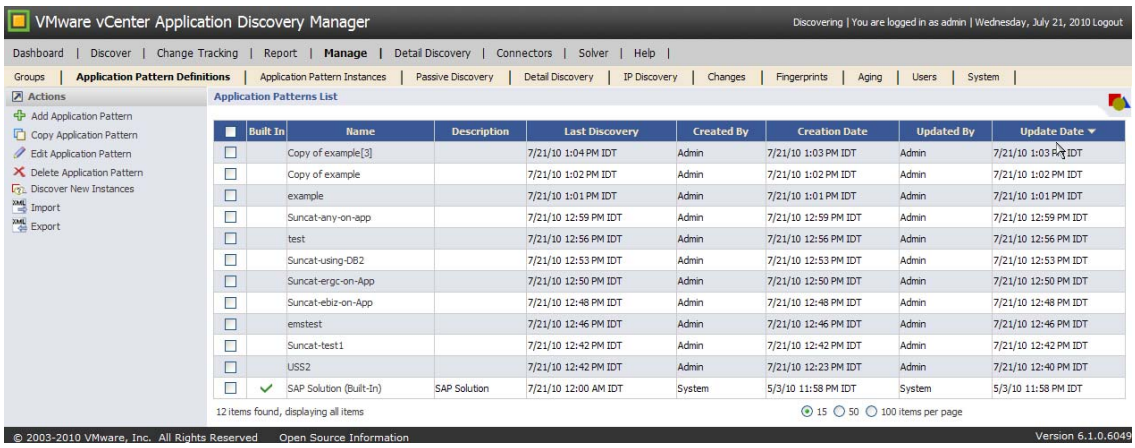

Table 6-1 describes the columns that appear in the Application Patterns List.

Table 6-1. Application Patterns Definition Column Descriptions

Column	Description
Built In	ADM has a set of default Application Pattern Definitions that you can use as is or as templates to create custom Application Pattern Definitions. The **Copy** option allows you to copy a built-in Application Pattern Definition and customize the copy to create a new group. Discovery of the built-in Application Pattern Definitions is performed once a day by default. The flag automatically discovers new instances once a day and is cleared by the user.
Name	Name given to the Application Pattern when it was created.
Description	Description of the Application Pattern Definition (optional).
Last Discovery	The last time the ADM searched for the Application Pattern. The first time an Application Pattern Definition is created, ADM searches the ADM database for the CIs that meet the criteria specified in the Application Pattern Definition. Application Pattern Instances are offered for each discovered instance of the defined pattern. Select the instances you want the ADM to save. These instances will now be displayed in the **Application Pattern Instances** tab. If the **Automatically discover** option is selected in the definition, the ADM will automatically search for new instances once a day and update the existing instances.
Created By	User name of the person who created this Application Pattern Definition.
Creation Date	Date the Application Pattern Definition was created. The date is for the ADM appliance where this definition is defined.
Updated By	User name of the person who last modified the Application Pattern Definition.
Update Date	Date the Application Pattern Definition was last modified. The date is for the ADM appliance where this definition is defined.

Click the column heading to sort the list by that column.

You can perform the following actions with Application Pattern Definitions:

- Add Application Pattern — Creates a new definition.
- Copy Application Pattern — Makes a copy of an existing definition.
- Edit Application Pattern — Modifies an existing definition.
- Delete Application Pattern — Removes an existing definition.

  NOTE You cannot delete an Application Pattern Definition if it is built-in or if one or more Application Pattern Instances reference that definition.

- Discover New Instances — Discovers new instances of an application pattern definition.
- Import — Imports application pattern definitions from other ADMs.
- Export — Exports application pattern definitions from other ADMs.

The *VMware vCenter Application Discovery Manager Online Help* provides procedures on how to perform these actions.

Application Pattern Instances

The Application Pattern Instances List contains the application pattern instance created as a result of discovery of instances of associated application pattern definitions.

To access the **Application Pattern Instances List**, select **Manage** > **Application Pattern Instances**. The Application Patterns Instances List is shown below.

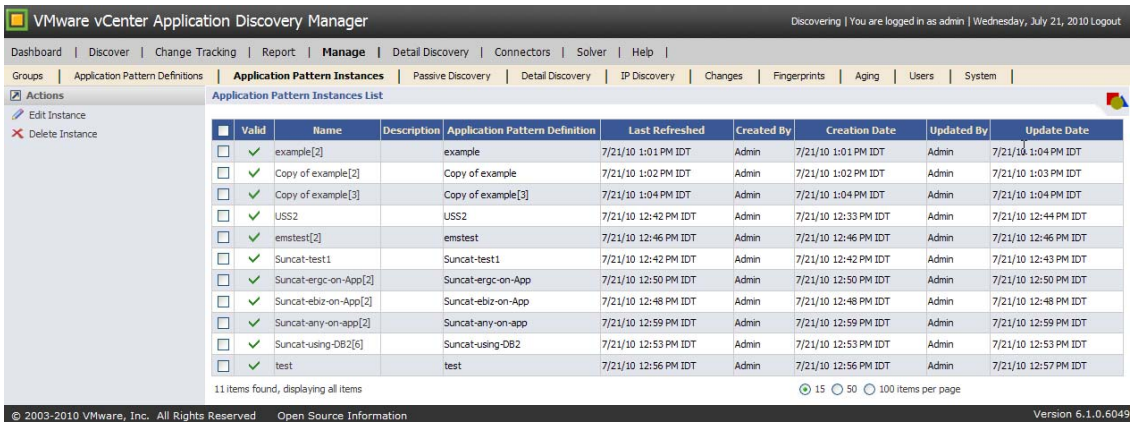

Table 6-2 describes the columns that appear in the Application Patterns Instances List.

Table 6-2. Application Patterns Instances Column Descriptions

Column	Description
Valid	A green check mark means the application pattern instance was valid as of the last **Refresh**.
Name	The name of the Application Pattern Instance. ADM provides a default name, but you can provide a custom name using the **Edit** action.
Description	The instance description is an optional field and might not have been defined. Use the **Edit** action to enter a description for the instance.
Last Refreshed	The last time the ADM database was checked for application pattern instance validity.
Created By	User name of the person who created this application pattern definition.
Creation Date	The time this application pattern definition was originally created on the ADM appliance where this pattern was defined.
Updated By	User name of the person who last modified the name or description of this application pattern instance.
Update Date	The time this application pattern definition was last modified on the ADM appliance where this pattern was defined.

Click the top of a column to sort the list by that column.

You can perform the following actions on Application Pattern Instances:

- Edit Instance — Modifies Name or Description of an existing instance.
- Delete Instance — Removes an existing instance.

NOTE You cannot delete an Application Pattern Instance if the instance is part of the scope of either Active probing or Aging policies, or part of a user's configuration.

The *VMware vCenter Application Discovery Manager Online Help* provides step-by-step instructions on how to perform these actions.

Application Discovery Process

Figure 6-7 shows a high-level overview of the process to create application patterns and view the results in the ADM console. The callouts correspond to the "Use Case: Creating Definitions and Viewing the Resulting Instances" on page 61.

Figure 6-7. Overview of the Application Discovery Process

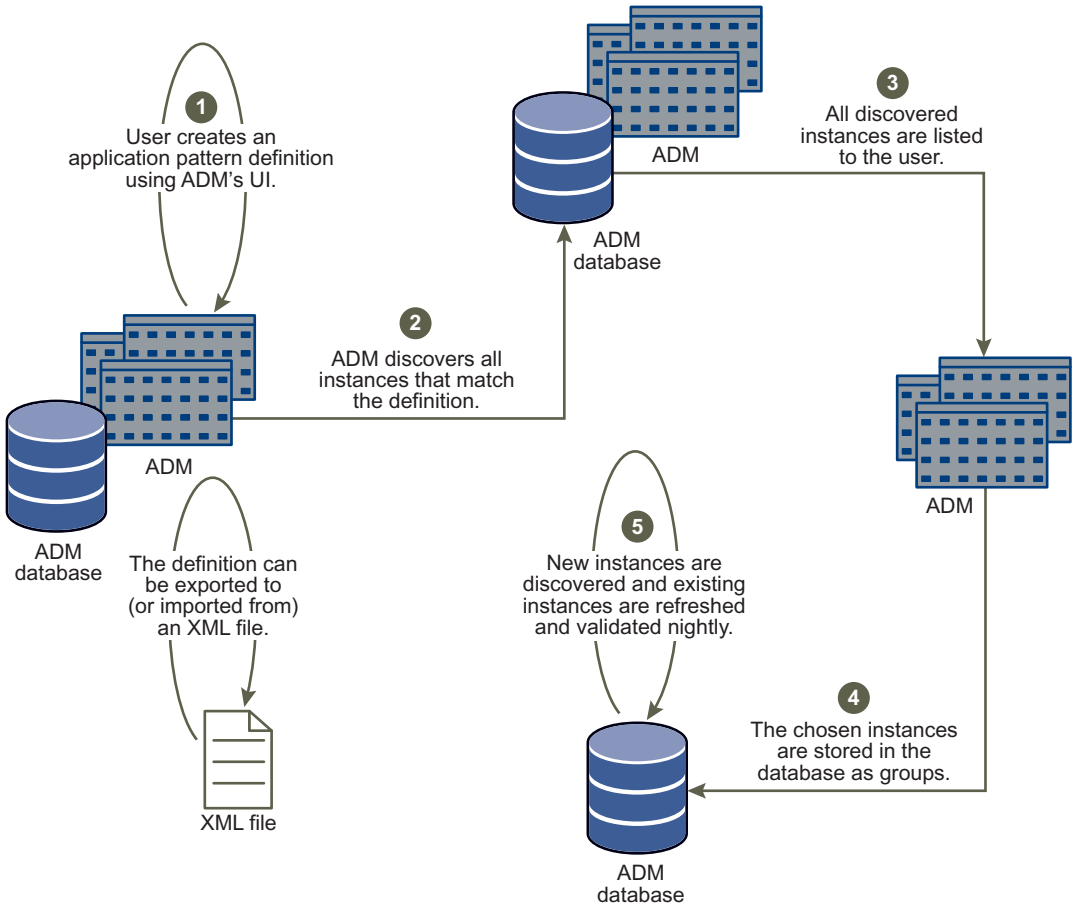

Use Case: Creating Definitions and Viewing the Resulting Instances

This use case provides an overview on how to create an application pattern definition and view the resulting instances. The *VMware vCenter Application Discovery Manager Online Help* describes the fields in the display.

Step 1: Create an application pattern definition

To create an application pattern definition

1 Navigate to **Manage** > **Application Pattern Definitions**.

2 Click **Add Application Pattern** from the **Actions** pane on the left side of the screen.

3 Create the node rules for each endpoint of the application pattern from the **Node Rules** tab.

IMPORTANT At least one node rule is mandatory.

---

4 Use the **Connectivity Rules** tab to define the type of connections to include in the instance and to define the nodes as a source or target of the application pattern instance.

Step 2: Discover All Instances That Match the Definition

The discovery process runs the first time an application pattern definition is created. ADM searches the database for the CIs that meet the criteria specified in the application pattern definition. If the criteria is met, then an application pattern instance is created.

Alternatively, the discovery process is triggered manually at any time as follows:

1 Navigate to **Manage** > **Application Pattern Definitions**.

2 Select the appropriate application pattern definition.

3 Click **Discover New Instances** from the **Actions** pane on the left side of the screen.

Monitoring of the background process is done through the **Last Discovery** column.When the discovery process is complete, the number of newly detected application pattern instances appear in the status bar of the window.

4 In the **Last Discovery** column, click **click here** to view the application pattern instances that were discovered.

Step 3: Viewing All Discovered Instances

After completing "Step 2: Discover All Instances That Match the Definition" on page 61, the Discovered Application Pattern Instances page appears displaying the scope of these instances.

Select the application pattern instances that you would like to store and click **Create**.

Step 4: Storing Selected Instances as Groups

The application pattern instances are stored as groups and are viewed and managed from the **Manage** > **Application Pattern Instances** tab.

Step 5: Discovering New Instances Automatically

You have an option to enable or disable the automatic discovery of new instances.

If the option Automatically discover new instances once a day is selected in the Application Pattern Definition, the ADM database is searched for new Application Pattern instances. Newly discovered instances are displayed and are saved manually as shown in "Step 3: Viewing All Discovered Instances" on page 62.

If the option is not selected, no new instances are discovered but a nightly refresh process synchronizes existing instances with information in the ADM database. If changes in the Application Pattern Instance (for example, relevant CIs were no longer discovered) render that instance irrelevant, the green check mark disappears in the Valid column. You can delete such an instance manually, provided the instance is not part of the scope of either Active Probing or Aging policies, nor part of a user's configuration.

Report 7

This chapter describes the **Report** tab in ADM. Topics include:

Report types

Table 7-1 lists all of the reports you can create in the **Report** tab.

Table 7-1. Available Reports in the Report Tab

Report Type	Definition	Choices For Each Report Type
Inventory reports	Inventory reports show what hardware and software are installed, the versions, and dependencies related to your business application.	■ Host Inventory Report ■ Device Inventory Report ■ Service Inventory Report ■ Connection Inventory Report ■ Host Operating System Breakdown Report ■ Host CPU Breakdown Report ■ Host Kernel Version Breakdown Report ■ Host Physical Memory Breakdown Report ■ Service Inventory Breakdown Report ■ Installed Software Inventory Report ■ Isolated Hosts Report ■ Isolated Services Report ■ Abandoned Services Report ■ Unclassified Connections Report ■ Unclassified Services Report ■ Changes Report ■ Host Configuration Report
Dependency reports	Dependency reports show what objects are dependent upon your hosts, services, and applications.	■ Host Dependency Report ■ Service Dependency Report ■ Application Dependency Report

Table 7-1. Available Reports in the Report Tab (Continued)

Report Type	Definition	Choices For Each Report Type
Demand reports	Demand reports show use information for hosts.	■ Host Baseline Behavior Report ■ Service Baseline Behavior Report ■ Host Baseline Comparison Report ■ Service Baseline Comparison Report ■ Most Used Hosts Report ■ Least Used Hosts Report ■ Most Active Host Users Report ■ Least Active Host Users Report ■ Host Demand Analysis Report ■ Service Demand Analysis Report ■ Most Used Services Report ■ Least Used Services Report ■ Most Used Connections Report ■ Least Used Connections Report ■ Most Active Application Users Report ■ Most Active Service Users Report ■ Host Demand Trend ■ Service Demand Trend ■ Host Activity Breakdown Report ■ Current Changes Report
Configuration reports	Configuration reports show either hosts with few or no connections, or a list of alerts.	■ Group Report ■ Change Policy Report

Exporting and Printing Reports

Once a report is generated, you can export and print it. The export choices are as follows:

- Excel format — Enables you to download the report through your browser as a Microsoft Excel spreadsheet.
- Rich Text Format (RTF) — Enables you to save the report in RTF that is opened in Microsoft Word.
- Portable Document Format (PDF) — Enables you to save the report in PDF.
- Print — Opens the standard **Print** dialog box and allows you to print the report.

Connectors

8

This chapter describes the **Connectors** tab that enables you to integrate ADM with other applications. Topics include:

NOTE The integration software requires a license to work with ADM. Contact your Customer Sales Representative for information on purchasing a license.

Connectors Overview

The **Connectors** tab enables you to integrate ADM with other applications, if you have them installed. You must also have a license for them. Integration between ADM and other applications allows detailed information to be discovered and populated into the other application.

Information about hosts, routers, switches, services, and connections are transferred between applications. For example, you might want to use the network devices from EMC Smarts Service Assurance Manager (SAM) with the applications from ADM. Information about ADM could be collected and displayed in the SAM.

To view the integration screen from the ADM Console, click **Connectors** tab. Depending on the applications you have installed and licensed, you see the following tabs:

- **EMC Smarts**
- **Reports**

The *VMware vCenter Application Discovery Manager Online Help* provides specific details and steps.

EMC Smarts Integration

Integration between the ADM and the SAM allows the detailed information collected by the ADM to be populated into the SAM.

NOTE Before ADM data can appear in SAM, you must verify that the `ics.conf` file has been edited to specify that the data is to flow from the SAM Adapter Platform to the SAM. The section "Defining Domain Parameters" in the *EMC Smarts Service Assurance Manager Configuration Guide* provides detailed information on editing the `ics.conf` file. Once all integration steps are complete, the ADM data is viewed using the Smarts, Launch in Context feature.

To view the ADM and SAM Integration screens

1 From the ADM Console, click **Connectors** tab.

2 Click **EMC Smarts** to view the **Smarts Integration Status** as shown below.

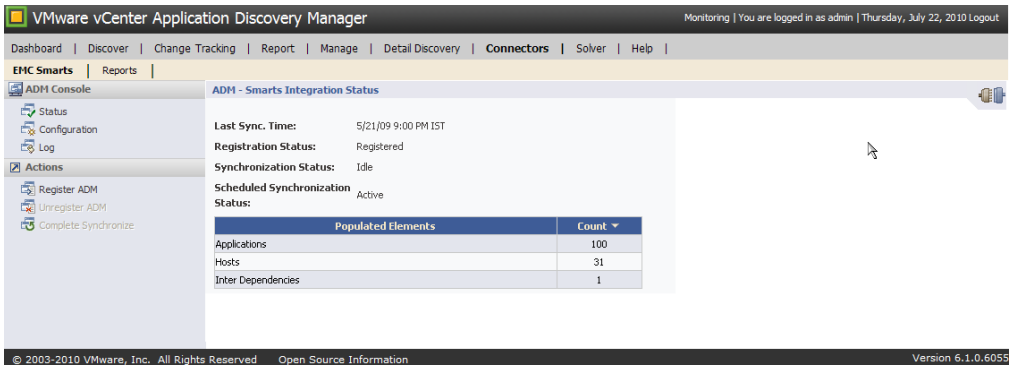

The ADM-Smarts Integration Status screen displays a summary of the ADM Smarts integration status.

The following menu items are available in the **ADM Console** left pane:

- Status
- Configuration
- Log

The following actions items are available in the **Actions** left pane:

- Register ADM
- Unregister ADM
- Complete Synchronize

These menu items are described in the following sections.

Status

Table 8-1 describes the various fields of the ADM-Smarts Integration Status screen.

Table 8-1. ADM-Smarts Integration Status Screen Information

Field	Description
Last Sync. Time	Displays the time of the last successful ADM-SAM synchronization.
Registration Status	Displays the current status of the ADM registration in the SAM. Available values include: ■ Not configured — The integration has not been configured. ■ Unregistered — The ADM device is not registered in the SAM. ■ Registered — The ADM device is registered in the SAM.
Synchronization Status	Displays the current status of the ADM device and the SAM synchronization. Available values include: ■ Idle — Indicates that no synchronization between the ADM device and the SAM is occurring. ■ Complete in Progress — A complete synchronization is currently in progress. ■ Incremental in Progress — An incremental synchronization is currently in progress.

Table 8-1. ADM-Smarts Integration Status Screen Information (Continued)

Field	Description
Scheduled Synchronization Status	Displays the status of the scheduled ADM device and the SAM synchronization. Available values include: ■ Active — Automatic synchronization scheduling has been activated. ■ Not Active — Automatic synchronization scheduling has not been activated.
Populated Elements/Count	Displays the type and the number of elements populated to the SAM.

Use your Web browser **Refresh** option to update the ADM-Smarts Integration Status screen information.

Click **Status** in the left pane to view the ADM-Smarts Integration status.

Configuration

The ADM-Smarts Integration Configuration screen includes the following tabs:

- **Publisher** — Configures the SAM server connecting to an ADM device.
- **Scheduling** — Configures automatic scheduling of the integration.

  NOTE When scheduling an integration, the schedule becomes active at 12:00 A.M. the next day.

- **Scope** — Defines the entities that are populated to the SAM.

To configure the EMC Smarts Connector

1 Configure the SAM server connecting to the ADM device.

2 Schedule the synchronization between the ADM device and the SAM.

  The following synchronization options are available:

  - **Incremental** — The incremental update option populates all objects that were created or modified since the last synchronization.

    Since the Incremental update option transfers deltas of data between two points in time, it is used more frequently.

  - **Complete** — The complete update option populates all objects that are included in the configured ADM to the SAM scope.

    Since this option populates all objects, do not use it frequently. It is recommended that you perform a complete update to correct any synchronization issues that have occurred over time.

3 Define the resources that are populated from the ADM device to the SAM.

  The **Scope** tab defines the resources that are populated from the ADM to the SAM.

  Since the ADM discovery process can result in a relatively large set of server and client resources, you must carefully plan on the scope of integration. The scope of integration should be limited to resources that are of interest within the SAM context. A broad scope can result in slow synchronization and a large set of entities in the SAM.

4 Register the ADM in the SAM.

  Before transferring any information from the ADM device to the SAM, ADM needs to be registered in the SAM.

  Only one ADM device is registered in the SAM at any given time. If you register a new ADM device by supplying a different name in the **ADM Name** field, the existing ADM device is unregistered. If you register a new ADM device with a name similar to an existing ADM device, the integration assumes that this is a replacement ADM device and attempts to synchronize the data of the ADM device and the SAM.

The *VMware vCenter Application Discovery Manager Online Help* provides the complete procedures for these steps.

Log

Click **Log** in the **ADM Console** left pane to view the logs for the ADM-Smarts integration. The log files are displayed from the most recent to the least recent. Click **Time** or **Message** header to reverse the order of the log files displayed. The **Time** header will toggle between the most current to the least current. Sorting by the **Message** column will first sort alphanumerical A–Z. Click **Message** again to reverse this order.

Unregister ADM

Click **Unregister ADM** in the **Actions** left pane to unregister ADM and disables any synchronization between ADM and the SAM. Once the ADM device has been unregistered, the **Registration Status** appears as **Unregistered** and all objects discovered by the ADM device are deleted.

Unregistering an ADM device when the SAM is not available displays a **Force Unregister** message. Click **Yes** to unregister the ADM device from the SAM without notifying the SAM. You must manually unregister the ADM device from the SAM using the SAM console.

Complete Synchronize

After you have configured the ADM and SAM integration, the system must be synchronized for the SAM to retrieve the data. The following two options are available for synchronizing:

- Schedule the synchronization between the ADM device and the SAM.
- Perform a complete synchronization now — To perform a synchronization now, click **Complete Synchronization** in the **Actions** left pane.

A complete synchronization is performed anytime after the two systems have been configured for integration.

The *VMware vCenter Application Discovery Manager Online Help* provides instructions for synchronizing the integration.

Displaying ADM data in SAM

Before ADM data can appear in SAM, you must verify that the `ics.conf` file has been edited to specify that the data is to flow from the SAM Adapter Platform to the SAM.

NOTE The **Defining Domain Parameters** section in the *EMC Smarts Service Assurance Manager Configuration Guide* provides detailed information on editing the `ics.conf` file.

To display discovered and populated ADM data in the SAM:

1. From the Notification Log Console, navigate to **InCharge Manager** > **Attach**. This displays the **Attach InCharge Manager** dialog box.
2. Select **INCHARGE-SA** from the **InCharge Manager** list box.
3. Navigate to **File** > **New** > **Map Console**. The **Map Console** appears.

4 From the left pane of the **Map Console** open the `ADMSoftwareRequest` folder.

5 Select an item. The graphical representation of the software request with the related software services appears in the right pane of the console as shown below.

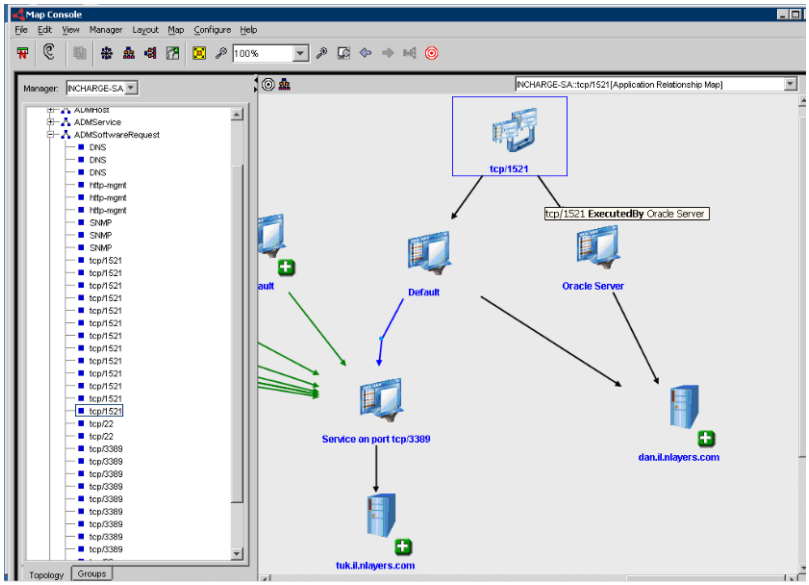

Launch in Context

To open an ADM console in context using the Smarts Launch in Context functionality, right-click an object and navigate to **Client Tools** > **Launch ADM Web Console** as shown below.

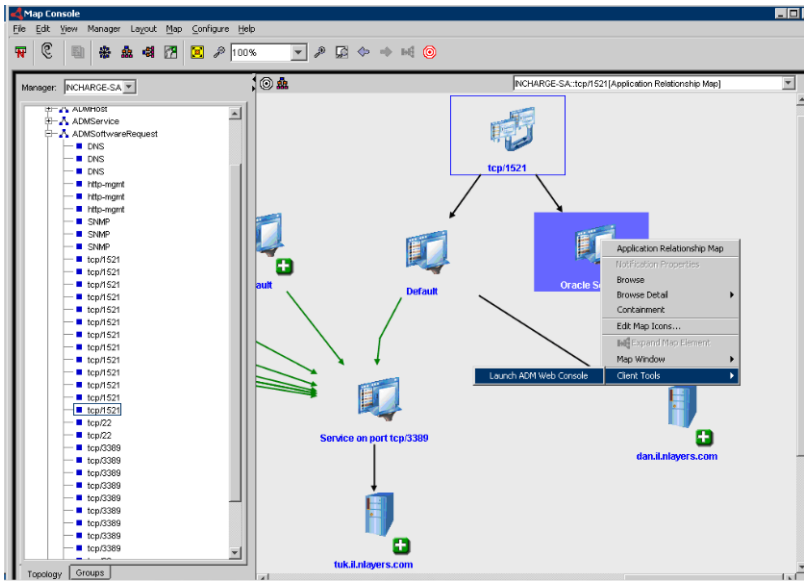

As a result of integration, the ADM Console displays detailed information about the selected object populated into the SAM.

NOTE On a host running Solaris 9 or 10, for the **Smarts Launch in context** functionality to open the **ADM Console in context**, you must edit the `AMDLIC.sh` file to configure the location of your Mozilla browser.

Editing the ADMLIC.sh File

In Solaris 9 and Solaris 10, the Mozilla browser replaced the Netscape browser. Because of this change, the `ADMLIC.sh` file must be edited to configure the location of the browser of your choice.

To edit the ADMLIC.sh file

1 Open the `ADMLIC.sh` file using a text editor.

2 Locate the following section:

```
# ==========| Customize only until next dashed line |==============
# ----------------------------------------

### Required:
# Some typical browser locations
# /usr/dt/bin/netscape    (Solaris)
# /opt/netscape/netscape  (HPUX)
# /usr/bin/mozilla        (Linux)
BROWSER=/usr/sfw/bin/mozilla
# --------------------
# ====================| End Customizations |======================
# ----------------------
```

3 Edit the path for the browser running on a Solaris host.

For example, BROWSER=/usr/sfw/bin/mozilla

4 Save the file to the `BASEDIR/smarts/actions/client/ADM` directory.

Custom Reports

In addition to standard reports, ADM provides functionality for creating custom reports. To create custom reports, an external database must be installed and configured. Once the external database is synchronized with the ADM database, you can create custom queries to collect data and generate custom reports. The *VMware vCenter Application Discovery Manager Repository Reference Guide* provides information about setting up and configuring the ADM external database. To help you with the custom reports, it also provides examples of reports that can be generated from the ADM external database.

Solver 9

This chapter describes the **Solver** tab that allows you to generate useful reports and to help you solve important business initiatives. Topics include:

Overview

ADM includes a series of best practice solutions to help you with your strategic business initiatives like license contract auditing; application, software, and server consolidation; disaster recovery planning; campus relocations; mergers and acquisitions; compliance with the Sarbanes-Oxley Act of 2002 section 404 (SOX); and many others. These analytics and best-practice solutions provide you with information to help you optimize your business applications and their resources.

Reports in the Solver Tab

Table 9-1 lists all the reports you can generate from this tab.

Table 9-1. Reports in the Solver Tab

Business Initiative	Definition	Reports
Application upgrade	Before upgrading your applications, use this solution to identify applications, hosts, and services that are heavily utilized and are excellent candidates for upgrading. This helps you improve your application performance and service levels.	■ Create a short list of hosts that are upgrade candidates. ■ Create a short list of services that are upgrade candidates. ■ List all hosts that are heavily used and could be upgraded. ■ List all services that are heavily used and could be upgraded. ■ Create a short list of the most active service clients for services that are upgrade candidates. ■ Create a short list of the most active application clients for applications that are upgrade candidates. ■ Determine the impact of hosts that are upgrade candidates on other hosts and services. ■ Graph the demand placed on hosts that are upgrade candidates over a specified time. ■ Graph the demand placed on the services that are upgrade candidates over a specified time. ■ Determine the impact of applications that are upgrade candidates on other hosts and services.

Table 9-1. Reports in the Solver Tab

Business Initiative	Definition	Reports
Application consolidation	Before running an application consolidation initiative, use this solution to identify applications, hosts and services that are under utilized and are excellent candidates for consolidation.	■ Create a short list of hosts that are retirement candidates. ■ Create a short list of services that are retirement candidates. ■ List all hosts that are not heavily used and could be retired. ■ List all services that are not heavily used and could be retired. ■ Determine the impact of the hosts that are retirement candidates on other hosts and services. ■ Determine the impact of business application retirement on other hosts and services. ■ Create a short list of hosts that have minimal dependency on other hosts and services.
Application migration	Before migrating your applications to new vendors, versions, or systems, use this solution to get a clear picture of your application architecture, their dependencies, and the demand placed on each application, host, and service.	■ List all hosts that support your application infrastructure. ■ List all services that support your application infrastructure. ■ Show which other hosts and services are dependent on the hosts that support this application. ■ Graph the demand placed on hosts that is migrated over a specified time. ■ Graph the demand placed on services that is migrated over a specified time. ■ Determine the impact of applications that are upgrade candidates on other hosts and services.
Mergers and acquisitions	As you acquire or divest business operations, use this solution to minimize the impact of acquiring or selling assets and ensure your business applications continue functioning without interruption.	■ Create a short list of hosts that are being acquired or sold. ■ Create a short list of services that are being acquired or sold. ■ Show which other hosts and services are dependent on hosts that are being acquired or sold. ■ Graph the demand placed on hosts that are being acquired or sold over a specified time. ■ Graph the demand placed on services that are being acquired or sold over a specified time. ■ Determine the impact of applications that are upgrade candidates on other hosts and services.
Disaster recovery planning	Create and automatically maintain accurate and up-to-date documentation of your disaster recovery plans. You can also use this solution to audit your disaster recovery plans, ensuring your business will continue without interruption.	■ List all hosts that support your application infrastructure. ■ List all services that support your application infrastructure. ■ Show which other hosts and services are dependent on hosts in your application infrastructure. ■ Determine the impact of applications that are upgrade candidates on other hosts and services.
SOX compliancy audit	Section 404 of the Sarbanes-Oxley Act requires you to document your key financial applications, amongst other things. This solution provides you with the necessary information required for Sarbanes-Oxley compliance readiness.	■ Show the list of hosts that support your critical financial applications. ■ Show the list of services that support your critical financial applications. ■ Show the list of hosts that depend on other hosts that support your critical financial applications. ■ Graph the demand placed on hosts that support your critical financial applications. ■ Graph the demand placed on services that support your critical financial applications.

Index