RAYMANAGESOFT i

A New Generation of
Software Deployment

RAYVENTORY
is part of RayManageSofti

RayManageSofti is part of RaySuite.

Configuration RayManageSofti 11.4

rayNET

Configuration
for release 11.4 *infinity*

Raynet GmbH Germany
See our website for locations.
www.raynet.de

Table of Contents

Introduction

RMS Configuration is a reference on different areas of RayManageSofti that benefit from configuration for your specific enterprise needs.

The main areas of configuration covered in this manual are:

- *Configuring your administration server* provides information about how you can change the appearance and behavior of the administration server.
- *Selecting distribution servers* describes how to configure managed devices to use the appropriate set of algorithms for selecting distribution locations from which to download content. Note that this chapter deals only with the configuration of the managed device. It does not deal with extending base product functionality with additional algorithms tailored to your enterprise. That facility is available, and is documented in *RMS Reference: Customization and Extension*.
- *Policy merging and distribution* discusses ways to configure your Deployment Manager environment to optimize server-side policy merging and distribution operations
- *Configuring environments with multiple domains* describes how to configure environments with multiple domains.
- *Configuring byte-level differentiation* examines the use of byte-level differentiation, and provides guidelines for testing the feature.
- *Configuring inventory* describes the different types of inventory data that can be collected on managed devices, and describes how to configure the inventory agent on managed devices. It also describes options about importing the inventory data into the RayManageSofti database.
- *Listening or polling for distribution jobs* describes how to set or modify the way your distribution servers receive distribution jobs from their parent servers.
- *Configuring reporting* describes how to configure Deployment Manager reporting.
- *Using remote control software with Deployment Manager* describes how to integrate your remote control software with Deployment Manager to enable RayManageSofti administrators to remotely connect to other managed devices.
- *Security overview* offers an overview of issues and solutions in the security aspect of software management
- *RayManageSofti and personal firewalls* describes how Deployment Manager interacts with personal firewall software such as the Windows Internet Connection Firewall (ICF).
- *Digital signing* examines the use of digital signatures and the protection these offer to packages (and therefore to the system overall)
- *Trusted locations* introduces trusted locations, and includes the configuration options for using trusted locations and digital signature together.
- *User groups* summarizes the user account types needed in a fully operational system.

- *Configuring administration server rights* describes security configuration on administration servers.
- *Configuring distribution server rights* describes security configuration on distribution servers.
- *Configuring managed device rights* describes security configuration on managed devices.
- *HTTPS security configuration* describes how and why to install a web server certificate on each IIS server that will use HTTPS.
- *Role-based security* describes how to set user access to RayManageSofti console tasks.
- *Configuring web proxy servers* describes how to configure your system to download data through an HTTP proxy server.
- *Moving the database* describes the processes for moving your RayManageSofti database from one server to another.
- *Bandwidth settings* describes how to specify your own bandwidth settings, used to configure volume of traffic between administration servers and distribution servers and distribution locations.
- *Data and log file processing* describes how to configure the behavior of the IIS web application and Deployment Manager data importers used to load data into the RayManageSofti database.

Before You Proceed

This manual is intended for experienced system administrators. The manual also assumes:

- A working implementation of Deployment Manager, the structure of which you understand well
- A working installation of Active Directory, and you are familiar with the SDOU structure implemented in your enterprise

Be aware:
If you intend to use Deployment Manager for ongoing software management, it is strongly recommended that you use Active Directory. Deployment Manager interacts with Active Directory to extend Windows Group Policy, providing sophisticated and efficient software management.

As an alternative for smaller enterprises, Deployment Manager does provide a simple native policy mechanism. This alternative is less flexible than using Active Directory.

- A general picture of how Deployment Manager works, at least to the depth of the *System overview* chapter of the *RMS Software Deployment*
- Skills and experience in administering Windows-based systems in particular

Additional Resources

The flexibility of the Deployment Manager system means that there are many different ways that you can configure it. As well, its interaction with many other parts of your infrastructure means that a successful implementation requires careful planning and a broad understanding of all the elements.
For these reasons, we strongly recommend that you seek assistance from trained Raynet consultants in preparing your implementation.
In addition to the expertise of Raynet consultants and the extensive product documentation, there are several other resources available to you.

- The *RMS Getting Started* provides an overview of documentation contents. Go there first to explore the available guides and references delivered along with RayManageSofti.
- The *RMS Release Notes*, included on the installation media, cover compatibility issues and late-breaking news about this release.
- *Application Help* is installed on the administration server and by default on managed devices:
  - Help on the administration server is intended for qualified personnel operating various aspects of the RayManageSofti system: administrators, package creators, release managers, and the like. Much of the content of the application help is replicated in this manual, but each contains some additional material.
  - Help on managed devices is intended for end-users, and provides a quick introduction to the self-service capabilities that RayManageSofti provides for optional installations.
- The support section of the product website *www.RayManageSoft.com* includes lists of what's new, known problems, and problems repaired for each release of RayManageSofti. There is also an extensive, searchable knowledge base of specialized articles and notes about many aspects of the product. Your enterprise should already have an account name and password for the support area. If not, contact your Raynet representative.
- Raynet and its partners offer a range of training courses that can also be customized to meet your requirements. For more information on these courses, speak with your Raynet consultant.

Feature Availability

With RayManageSofti 11.4 *infinity*, Raynet introduces a re-invented license control system. By implementing three individually settable options per RayManageSofti feature, customers can easily configure the most suitable RMS toolbox for their corporate needs. Each feature can be enabled, disabled, or set into demo mode. The demo mode allows for preview feature-testing, whilst preventing full feature usage. Features in demo-mode are marked by a flag to indicate the limited state. Some features, such as the AD policy editor, base on the availability of others (e. g. the policy snap-in).
This document describes RayManageSofti in its full stage of expansion. Therefore, if an individual license does not cover all features, readers might miss snap-ins, workflows, wizards etc. as within their actual user interface.

RMS consultants and sales representatives are the right contact persons for a list of available combinations and for advice on the preparation of custom-made RayManageSofti 11.4 *infinity* licenses.

Documentation Requests

We welcome your suggestions and input on the various documentation resources available with RayManageSofti and its components. Your comments and requests can be forwarded through your Raynet support representative.

Manual Conventions

The following typesetting conventions are used in this manual:

- Cross references to headings or chapters in this manual, or to other manuals, are shown in italics:
  "See *RMS Reference: System Reference* for..."
- Quotations from your computer screen (titles, prompts, and so on) are shown in bold:
  "The **Receive Packages Wizard** appears."
- Code syntax, file samples, directory paths, entries that you may type on screen, and the like are shown in a monospaced font:
  "The default directory is `C:\ManageSoft...`"
- Italics may also be used for emphasis: "This manual is *not* intended..."
- Bold may also be used for inline headings: "**Target**: Indicates a target frame..."

Two note formats are used in RayManageSofti documentation

This is the basic format for giving additional information to the current topic.

It can come with four different headings:

Be aware:
This note format contains important information related to your current activity. You should not skip over this text.

Note:
This format is used for items of interest that relate to the current discussion.

Best practice:
If there is a best practice approach to the current topic you can decide if you want to follow it, or stick to your own plan.

Tip:
Tips are designed to help you find the easiest and quickest way to work with RayManageSofti.

The second format is for very serious alerts.

WARNING
The information here may save you from data loss. Pay particular attention.

Registry

In this manual, registry keys are preceded by the text `[Registry]`. This text represents the location of all RayManageSofti registry entries in the registry:

- On Windows administration servers, distribution servers and managed devices, RayManageSofti registry entries are usually stored under the key:

  ```
  HKEY_LOCAL_MACHINE\SOFTWARE\ManageSoft Corp\
  ```

- On non-Windows managed devices, registry entries are stored in the `/var/opt/managesoft/etc/`

`config.ini` file. Within this file, registry keys are shown in square brackets.

For example: `[ManageSoft\Usage Agent\CurrentVersion]`.
The lines below each key show the registry entries set under that key.
For more information about RayManageSofti preferences and the registry, see *RMS Reference: Preferences* for Managed Devices.

Example Usage

When the manual refers to the registry entry `[Registry]\Launcher\CurrentVersion\ForceReboot`, this represents:

- In the Windows registry, `HKEY_LOCAL_MACHINE\SOFTWARE\ManageSoftCorp\ManageSoft\Launcher\CurrentVersion\ForceReboot`
- On non-Windows devices, the `ForceReboot=True|False` line in the `config.ini` file, located under the section heading [`ManageSoft\Launcher\CurrentVersion`].

Configuring Your Administration Server

During installation of Deployment Manager, your administration server is automatically configured, based on the choices you made during the installation and standard installation defaults.

Many of these configuration items are controlled through the Windows registry on the core server. However, the most commonly changed configuration items can be set through an MMC snap-in.

What Can You Configure?

A large number of configuration items are available. As examples, the configuration items available include:

- **Administration Server**: Reports, repository, backward-compatibility
- **Discovery and adoption**: Remote execution commands and ports, various paths, OU for discovered devices
- **Distribution**: Replication permissions, batching, job queue location, timing and job retry
- **Importers**: Connection strings, domains, and SQL error codes that are regarded as permanent failures
- **Packaging**: Compression, time stamp format
- **Policy**: Containers, security groups, domain controllers, policy distribution
- **Security**: Folders, URLs, command line parameters, authenticode
- **Software**: Default package states
- **Tracing**: On/off configuration for many trace classes
- **Wake on LAN**: Folders and timing

Editing Your Administration Server Configuration

The most commonly changed administration server configuration items can be set from the Deployment Manager Configuration Console, accessible by

- navigating through your Windows **Start** menu**: All programs** > **Deployment Manager** > **Configuration** > **Deployment Manager Configuration**
  or
- selecting the **Deployment Manager Configuration tile** from your **start screen**

Note:
The **Product Activation Wizard** may display once you open the console, if you are using a temporary time out license or if the number of devices you are managing approaches or exceeds the number you are licensed to manage. If it does appear, select **Start** to go ahead.

To change administration server configuration settings:

1. Access the **Deployment Manager configuration** node.

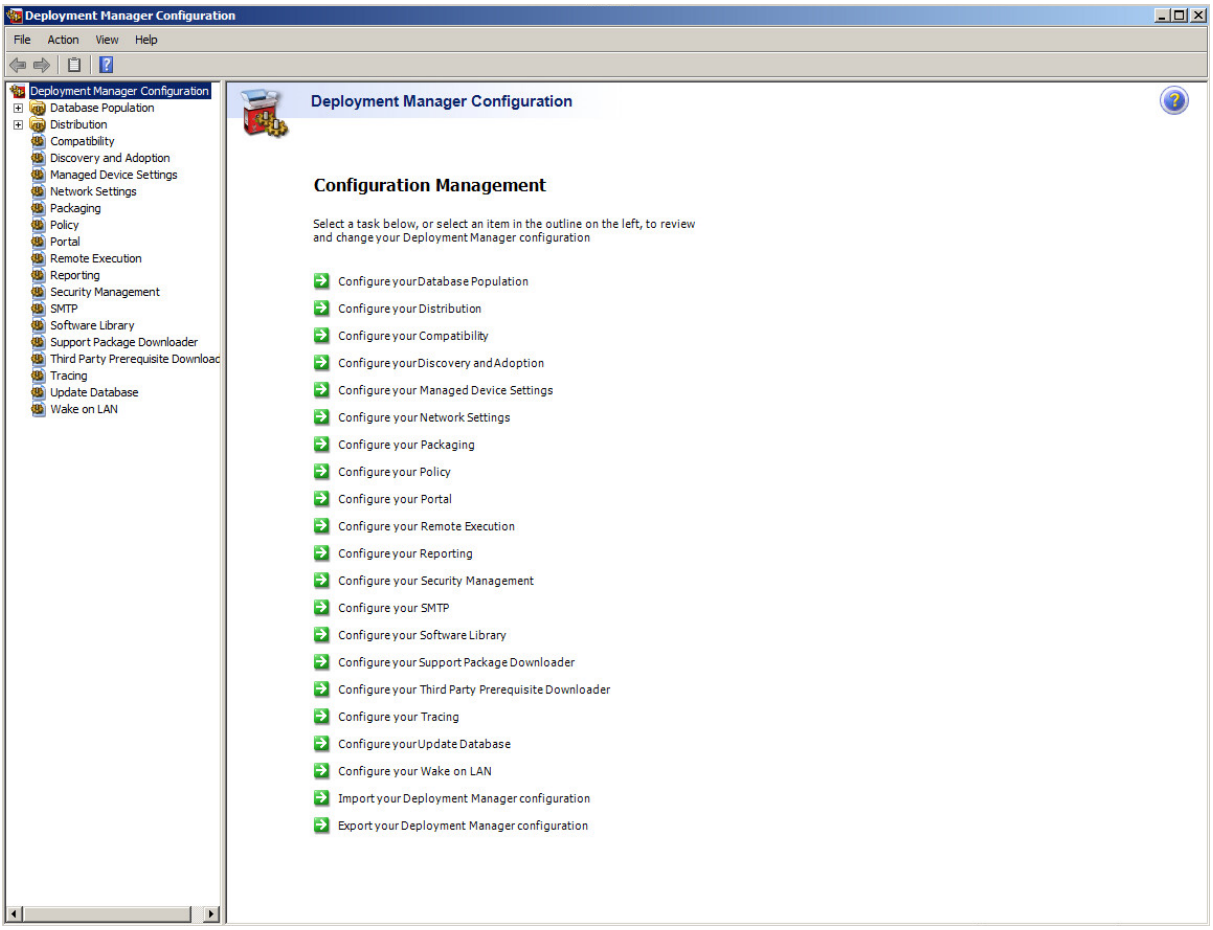

2. Select configuration settings from the details pane or browse through the nodes in the console tree to display the appropriate configuration page. A sample page is shown below.

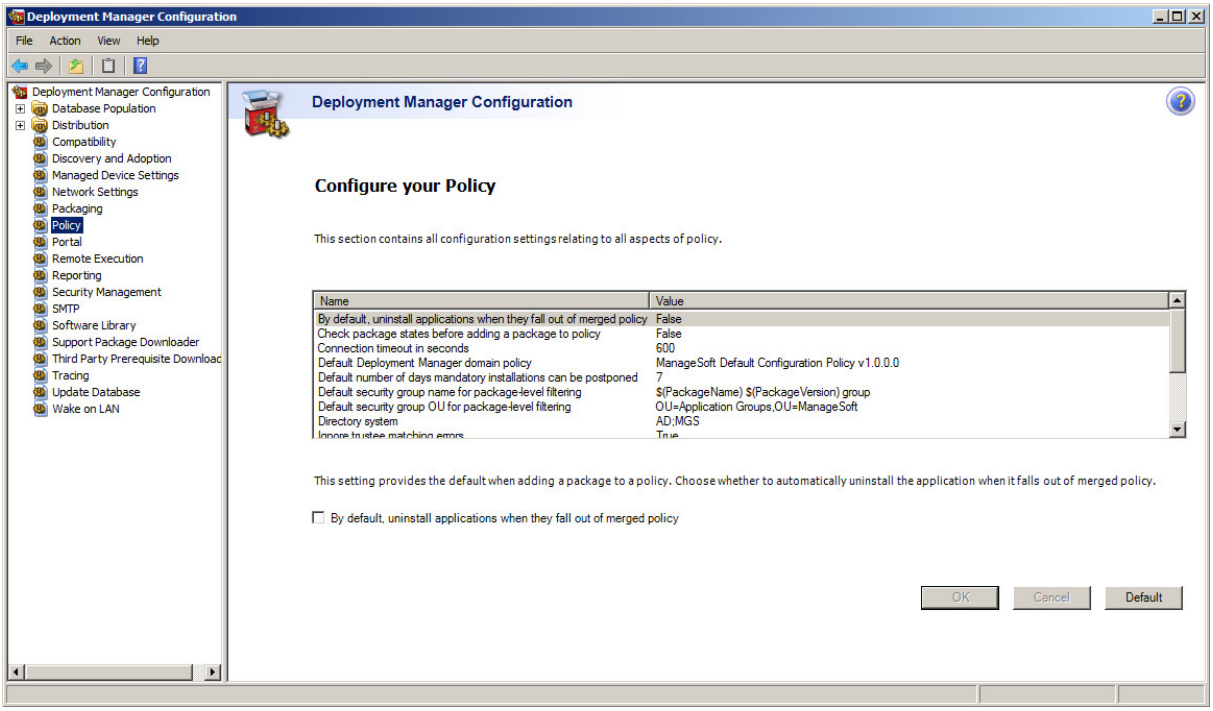

3. To change the value of a configuration setting, select the name from the list.

   A field to modify the setting is displayed below the configuration settings table. Depending on the type of value, you may be able to select from drop-down options, or type in the required number or string. In some cases, you can select a checkbox to set an option.

   To reset the value of a setting to the Deployment Manager default, click **Default**.

4. Click **OK** to confirm your changes.

Importing and Exporting Configuration

Facilities are provided for you to export your configuration settings to a data file and import your configuration settings from a data file.

These facilities are provided to assist you when troubleshooting.

To Export Your Configuration

1. Start the **Deployment Manager Configuration console** as described above.

2. Right-click the Deployment Manager Configuration node and select **Export**.

   A dialog is displayed for you to specify the export file.

3. Click **Save** to confirm the export.

4. An information dialogue is presented.

Click **OK** to end the export routine.

To Import Your Configuration

1. Start the **Deployment Manager Configuration console** as described above.
2. Right-click the Deployment Manager Configuration node and select **Import**.

   A dialog is displayed for you to specify the import file.
3. Click **Open** to confirm the import.
4. An information dialogue is presented.

   Click **OK** to end the export routine.

Configuring Remote Administration Consoles

Remote administration consoles obtain some configuration settings from their own registry, and other settings from the registry of the administration server:

- For operating preferences managed through the **Deployment Manager Configuration** console, remote consoles always retrieve these settings from the administration server. If a system administrator changes a setting using the Configuration feature from a remote console, the changes are saved on the administration server.
- For preferences set in other dialogs on the RayManageSofti console (often property dialogs), the settings are stored on the administration server if they relate to the general operation of Deployment Manager (for example, distribution settings), or are stored on the remote console if they relate to the user's operating choices (for example, settings associated with the operating preferences for a wizard).
- A number of preferences are always read from the local registry on the remote console. See *Settings on the local registry* below.

Accessing Administration Server Settings

Remote consoles can only access administration server settings within a local area network and only if the account used to log in to Deployment Manager on the remote console has permission to read (and in some cases, write) Deployment Manager registry keys. By default, members of the **MGS Administrators** group have read and write permissions to the `HKEY_LOCAL_MACHINE\SOFTWARE\ManageSoft Corp\ManageSoft` hive. In addition, the remote console cannot access administration server settings through a firewall.

If the remote console is unable to read the administration server registry settings (either because it is not within a LAN or because the account does not have appropriate permissions), the remote console will use any equivalent settings in the local registry on the remote console.

Settings on the Local Registry

The following registry settings are always read from the local registry on a RayManageSofti remote console:

Stored under [Registry]\ManageSoft\...	Local registry setting
\	InstallDir ETAPInstallDir ETAPVersion
\ Administrator \ CurrentVersion	AccessMode LockDLL
\ Administrator \ CurrentVersion \ AvailableFeatures	Configuration Licensing RPMPackaging OSXPackaging SolarisPackaging Console SoftwareLibrary Packaging SnapshotWizard WindowsInstallerWizard ClassicWizard MergeWizard Distribution Reports Scheduling ADPolicy DeploymentPolicyEditing Reporting ThirdPartyInstaller Wizard WakeOnLAN
\ Distributor \ CurrentVersion	SMDIndexFolder SMDIndexName
\ EventLog \ CurrentVersion	EventLogPolicy
\ Framework \ CurrentVersion \ Applications \ PackageBrowser	AppClass AppDSO AppIcon AppName
\ Framework \ CurrentVersion \ Applications \ Packer	AppClass AppDSO AppIcon AppName
\ Framework \ CurrentVersion \ Applications \ Reports	AppClass AppDSO AppIcon AppName
\ Framework \ CurrentVersion \ Applications \ Scheduler	AppClass AppDSO AppIcon AppName

Stored under [Registry]\ManageSoft\...	Local registry setting
\ Framework \ CurrentVersion \ Extensions \ DPE	ExtDSO ExtClass ExtName
\ Packer \ CurrentVersion	signcode signcode_orig AutoCreateLocation AutoCreateINI
\ Packer \ CurrentVersion \ Installation Snapshot Wizard	NDPSaveDir
\ Profiler \ CurrentVersion \ Mapper	MetaPackageStageDir StageDir
\ Replication Agent \ CurrentVersion	ueueLocation JobQueueRemote JobQueueProtocol PluginDirectory
\ Reporter \ CurrentVersion	ReportsRefreshRate PortalHomePage
\ Repository \ CurrentVersion	RepositoryLocation
\ Scheduler \ CurrentVersion \ UploaderRules	Inventory Log PolicyComplianceLog ReplicateLog
\ Warehouse \ CurrentVersion	AdministratorsGroup DistributorsGroup ReportUsersGroup TaskpadURL BestPractiseDocLocation
\ Warehouse \ CurrentVersion \ SoftwareLibrary	ReceiveNDPSourceDir ReceiveOSDSourceDir
\ WarehouseConsole \ CurrentVersion \ Arbitrary CommandWizard	CommandFile
\ WarehouseConsole \ CurrentVersion \ RemoteExecutionWizard	UserMustHaveWriteToAttribute
\ SecurityPatch \ CurrentVersion	OnlineHelpFile OnlineHelpTopics

Configuring Security Groups for Package-level Filtering

You can also set the name and location of security groups, used in package-level filtering. Package-level filtering allows you to minimize the number of Group Policy Objects in Active Directory by bundling packages that are

not part of a standard application set into a Group Policy Object.

For more information on package-level filtering, see *Using security groups* in the *Deployment policies* chapter of the *RMS Software Deployment*.

To set the name and location of security groups used in package-level filtering:

1. In the backwards compatibility console tree*, right-click the root node and select **Properties** from the context menu.

   The **Package Filtering tab** is displayed.

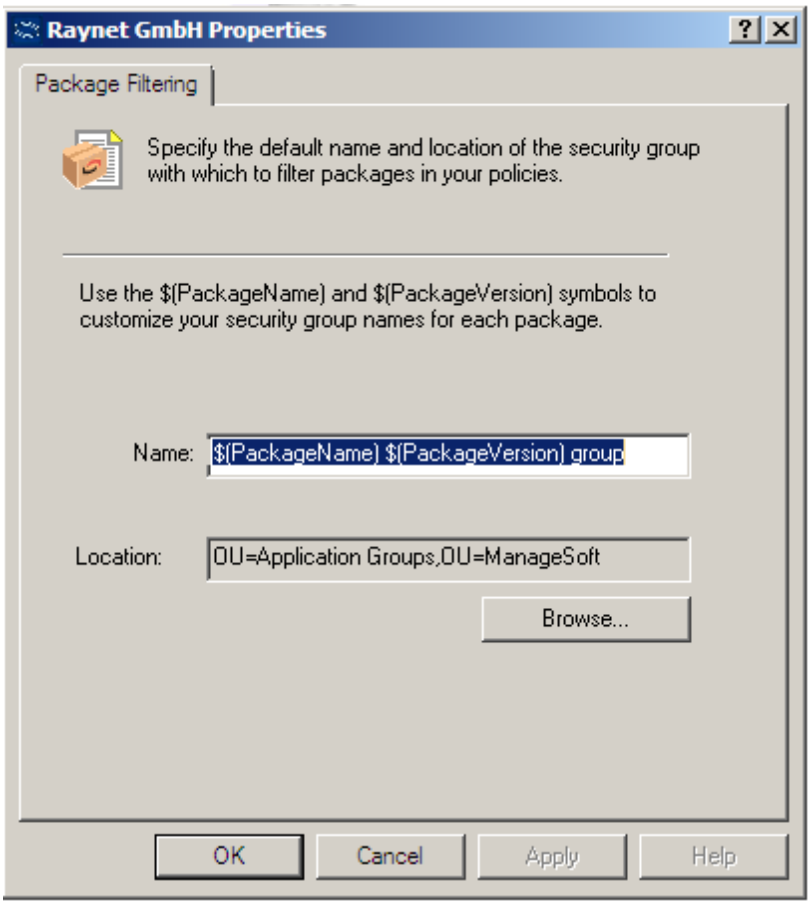

2. Enter the naming convention you wish your security groups to have.

   These are symbols used to vary the name of the security group to match the package name and version number. For example: **$(PackageName) $(PackageVersion)** group.

3. Browse to the location you wish your security groups to be created in.

Note:
This location must already exist.

This information is used in the **Add package to Policy wizard**, where you specify security groups for a

particular package.

* Some RayManageSofti functionalities are exclusively available via the backwards compatibility console. To open this console navigate to the Program Files folder on your Deployment Manager administration server and execute the `eta.msc` file within the ManageSoft directory.

Selecting Distribution Servers

This chapter describes how managed devices determine which distribution server to use for a download or upload operation, and how you can configure these selections.

Configure here means choosing between a combination from a number of supplied options. It is also possible to extend those options by adding algorithms of your own. The latter extension is not covered in this manual. Instead, see *RMS Reference: Customization and Extension*.

Overview

RayManageSofti distribution servers are used as staging posts between the administration server and managed devices. Managed devices retrieve data from the distribution locations on distribution servers (or other computers). They also upload data to reporting locations on distribution servers.

Details of all distribution locations and reporting locations in your RayManageSofti distribution hierarchy are stored on each managed device. You can update the record of locations by distributing a *Failover locations package* from the RayManageSofti console. For more information, refer to the *Distribution system* chapter of *the RMS Software Deployment*.

A list of distribution servers from which a managed device may collect packages and other updates is established from settings on the managed device. As some distribution servers are better suited to a given managed device than others, the managed device can prioritize the list of distribution servers. (You may also override this capability by setting a fixed priority.)

To determine which distribution server to use for upload and download activities, managed devices use a set of rules, called an *algorithm*, to assign priorities to each distribution server. The distribution server with the highest priority (lowest numeric value) is used. This is commonly referred to as determining the closest distribution server, although the server may not be physically closest to the managed device.

In the event that the highest priority distribution server cannot be used (for example, because of network problems), the distribution server with the next highest priority is used, and so on until the download or upload activity can be performed. It is possible to configure the number of, and period between, attempts to connect to each server.

The following section describes how Windows managed devices select distribution servers. For information on UNIX and Macintosh managed devices, refer to *The following section describes how Windows managed devices select distribution servers. For information on UNIX and Macintosh managed devices, refer to*.

On Windows Managed Devices

Details of all distribution locations and reporting locations in your RayManageSofti distribution hierarchy are stored in the Windows registry. You can configure the registry settings to assign priorities to distribution servers.

Download and Upload Locations

The details about RayManageSofti download and upload locations stored in the Windows registry include folder and host names, and port numbers. This information is stored in the **DownloadSettings** and **UploadSettings**

registry keys which are located under:

Download settings	`[Registry]\ManageSoft\Common\DownloadSettings\[GUID]` where GUID is the GUID of the distribution location
Upload settings	`[Registry]\ManageSoft\Common\UploadSettings\[GUID]` where GUID is the GUID of the reporting location

The following Registry Editor screen shows sample download settings.

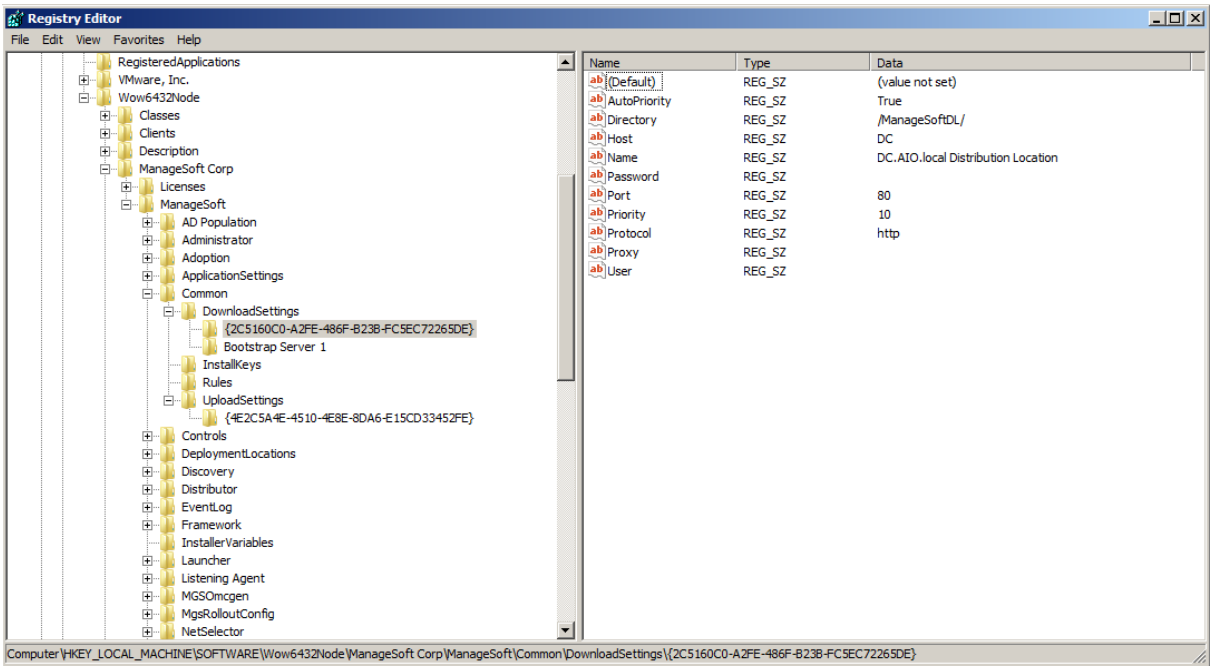

Two of the Windows registry keys stored for each location are associated with assigning priorities to distribution servers: **AutoPriority** and **Priority**.

AutoPriority determines whether a distribution server has a priority calculated at the time of download or upload, or whether it is fixed at the value declared in the **Priority** registry key.

By default, the distribution server used in adoption of the managed device is assigned a low priority (100) so that it is naturally overridden by servers identified in managed device settings.

When **AutoPriority** is `True`, each managed device uses its selected RayManageSofti (or your custom) algorithm to determine the priority of this location at the time of upload or download. This is the recommended behavior, as it optimizes system performance over time despite variations in network configuration.

When **AutoPriority** is `False`, the value of the **Priority** registry key declares the distribution server's fixed priority. For further information on these settings, refer to *RMS Preferences for Managed Devices*.

If you need to override the intended behavior of the **SelectorAlgorithm** setting and assign a fixed priority to a distribution (or reporting) location for specific managed devices, you may:

- Set the **AutoPriority** and **Priority** registry keys manually, where only one or two managed devices are affected (for example, debugging)

- Create a small custom package on the administration server to distribute the appropriate registry settings to managed devices within a security group (see the *RMS Packaging* for more information)

- Use Active Directory templates to distribute the registry settings (see your Active Directory documentation for more information).

Be aware:
Whenever Deployment Manager updates the prioritized list of distribution locations, the same prioritized list is also written to the Microsoft Source List. This ensures that updates triggered through the Microsoft Windows Installer from sources, such as Add/Remove Programs, use the same prioritized list of distribution locations as updates triggered by Deployment Manager.

Prioritizing Distribution Servers

When determining the closest distribution server, a managed device needs to know which algorithm to use to prioritize distribution servers. It also needs to know the range of priority values to assign to those servers. This information is stored in three registry keys on each managed device: **HighestPriority**, **LowestPriority**, and **SelectorAlgorithm**. These keys are located under:

```
[Registry]\ManageSoft\NetSelector\CurrentVersion
```
or
```
HKEY_CURRENT_USER\Software\ManageSoft Corp\ManageSoft\NetSelector\CurrentVersion
```

For further information, see *RMS Preferences for Managed Devices*.

The following Registry Editor screen shows sample **NetSelector** settings:

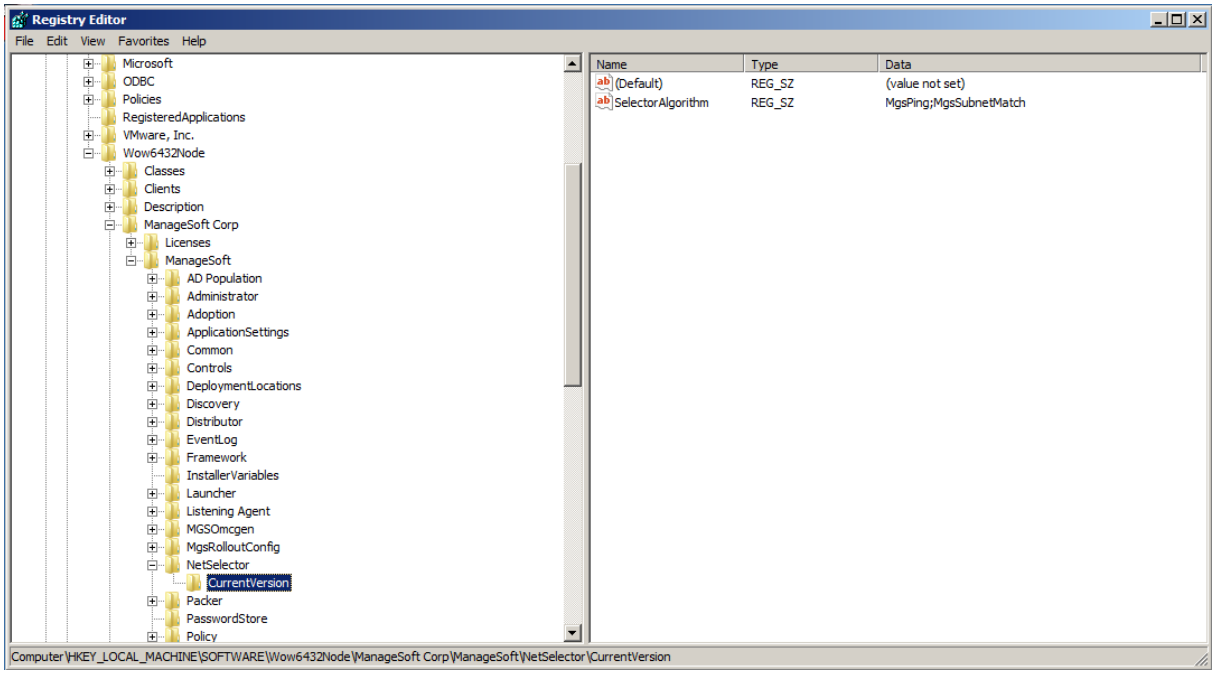

If there is no **NetSelector** registry entry, you may need to add the appropriate registry key.

Supplied Algorithms

Deployment Manager includes the following algorithms for assigning priorities to distribution servers (you can write and use your own DLL to create others):

- *MgsADSiteMatch*: Match to Active Directory site
- *MgsBandwidth*: Bandwidth priorities
- *MgsDHCP*: Retrieve location list from DHCP server options
- *MgsDomainMatch*: Match to domain name
- *MgsIPMatch:* Match to IP address
- *MgsNameMatch*: Match prefixes of computer names
- *MgsPing*: Fastest responding server
- *MgsRandom*: Random priorities
- *MgsServersFromAD*: Retrieve location list from Active Directory
- *MgsSubnetMatch*: Match to subnet

How you use these algorithms depends on the structure of your distribution hierarchy and how you want to spread the load of file uploads and downloads across that hierarchy.

For example, random priorities may be useful if you want to ensure a good load balance across distribution servers, while *domain matching* may tend to favor particular servers. Alternatively, random priorities might not be suitable if there are particular distribution servers that you want to target.

You can choose a combination of fixed and assigned dynamic priorities, where you determine the priorities of some servers, but allow the priorities of other servers to be assigned by the Deployment Manager algorithms. For example, there may be a number of servers to which you want to specifically assign low priorities, but let Deployment Manager randomly assign priorities amongst the servers you want to use more often. In this scenario, for each location that requires a fixed priority, you would set the **AutoPriority** registry key to `False` and manually assign a value for the **Priority** registry key.

For mobile users, whose upload and download speeds can be severely affected by the decision of which distribution server to use, the Deployment Manager *IP matching* algorithm is most often favored. When a mobile user connects to the network and is allocated an IP address, Deployment Manager uses the distribution server with the closest match to that IP address, matching components of the IP address from left to right.

For example, when a mobile user who travels to different geographic regions is in Germany, the IP matching algorithm assigns higher priority to the company's European distribution servers. When the same mobile user travels to New York, the company's US distribution servers get highest priority.

Using a combination of IP matching and random algorithms would help balance the load between servers with similar IP addresses.

Algorithms like *site matching* and *name matching* can be useful in large enterprises. Among other effects, you can use them to limit the range of *fastest response servers*, and thereby control network impacts.

You can use more than one algorithm. When you do this, Deployment Manager assigns priorities according to the first algorithm, then reassigns priorities according to subsequent ones. This can be used to refine priorities, or to ensure a shared load across servers, as the following example shows.

Example: `SelectorAlgorithm="MgsDomainMatch;MgsRandom(3)"`

This means that Deployment Manager should sort all servers using a domain match, then randomize the top three servers.

An empty string means that no algorithm is used to determine the download or upload location. In this case, the last priority values assigned to locations are used.

By combining algorithms, you can also limit the servers that other algorithms will test, since they will all ignore any server with a priority set to `invalid` (or indeed, to any string that does not represent an integer). For example, a combination such as

```
MgsADSiteMatch(, true);MgsSubnetMatch(, true);MgsPing(3)
```

will prioritize the fastest three servers within the managed device's site and subnet.

Remember that the algorithms only assign values to locations that have the **AutoPriority** registry key set to `True`.

Be aware:
When unable to differentiate between locations, the Deployment Manager algorithms assign priorities according to the order in which locations are listed in the registry settings.

MgsADSiteMatch: Match to Active Directory Site

Be aware:
This algorithm is only available if the managed device is joined to an Active Directory domain, and **Active Directory Sites and Services** is correctly configured.

This algorithm is particularly useful when the managed device preferences list a large number of servers (perhaps all the servers in an enterprise), and you want to restrict transfers to those within the site.

MgsADSiteMatch takes the following steps:

- Identifies the Active Directory site in which this managed device is a member.
- Collects the list of selected servers from the preference settings (normally in the machine hive of the registry). If the algorithm is in use by the installation agent, it collects the server list from the preference **DownloadSettings**. If the algorithm is in use by the upload agent, it collects the server list from the preference **UploadSettings**.
- For each server in the list, it looks up the Active Directory site in which the server is a member. You can restrict the domain controller used for these server site queries using the `localSiteDCOnly` parameter (see below).
- If the server's site is the same as the site for the managed device, the algorithm assigns a high priority. Where there are multiple such servers, their relative priority is left unchanged.
- If the server is in a different site than the managed device, the behavior depends on the `discardForeign` parameter (see below).

Syntax

```
MgsADSiteMatch( int limit, boolean discardForeign, boolean localSiteDCOnly )
```

where:

- `Limit` is an optional integer setting the maximum number of servers to which priorities will be assigned by this algorithm.
- `discardForeign` is an optional boolean represented by the case-insensitive strings `true` or `false`. The default is `false`.
  - If `false` or omitted, the algorithm assigns a low priority to servers that are not in the same Active Directory site as the managed device. Where there are multiple such servers, their relative priority is maintained.
  - If `true`, the algorithm sets the priorities of any servers that are not in the same Active Directory site as the managed device to the string literal `invalid`. The installation agent and the upload agent will not use the servers for transfers.
- `localSiteDCOnly` is an optional boolean represented by the case-insensitive strings `true` or `false`. The default is `false`.
  - If `false` or omitted, the algorithm does not require the domain controller used for server site queries to be in the same Active Directory site as the managed device.
  - If `true`, the algorithm restricts the server site queries to a domain controller that is in the same Active Directory site as the managed device. If there is no domain controller in the same site as the managed device, the algorithm treats all servers as being in a different Active Directory site as the managed device.

Example of MgsADSiteMatch Algorithm Results

For managed device in Active Directory site: Boston

Location	AD Site of server	Auto-Priority	Discard-Foreign	Priority before algorithm	Priority after algorithm
A	Melbourne	True	False True	1 1	2 Invalid
B	Boston	True	False True	3 3	1 1
C	Boston	False	Don't care	4	4
D	Frankfurt	True	False True	2 2	3 invalid

In the example shown, the algorithm determines priorities in the following way:

- Location C has **AutoPriority** set to `false`, so that its priority value of 4 is excluded from the allocation.
- If the algorithm parameter `discardForeign` is `false` or missing, Locations A and D will be given lower priorities (with their relative order maintained). However, if `discardForeign` is `true`, both Locations A and D will be marked invalid because they are not in the same site as the managed device.

MgsBandwidth: Bandwidth Priorities

This algorithm prioritizes servers based on the end-to-end bandwidth available to the server. It uses an average of ping requests where packets of different sizes are sent as part of the calculation. Unlike **MgsPing**, there is no parallelism in querying servers, therefore this algorithm should only be used in scenarios that do not require parallel pinging.

MgsBandwidth estimates the total bandwidth available between the local machine and each server; not the amount of currently unused bandwidth. The estimate is more accurate when there is less traffic on the network.

Syntax

```
MgsBandwidth( int limit )
```

where:

- `limit` is an optional integer setting the maximum number of servers to which priorities will be assigned by this algorithm.

Example of MgsBandwidth Algorithm Results

Location	Bandwidth	Priority before algorithm	Priority after algorithm
A	LAN (100Mbps)	Blank	1
B	LAN (10Mbps)	Blank	2
C	WAN (56K Modem)	Blank	4
D	WAN (ADSL)	Blank	3

MgsDHCP: Retrieve Location List from DHCP Server Options

This algorithm prioritizes according to lists of servers specified in a DHCP property. Whenever the installation or upload agents need to prioritize distribution or reporting locations, NetSelector sends a DHCP broadcast to request the value of the DHCP property corresponding to the option number specified in the algorithm. The value contains a list of one or more servers that are to be prioritized. You should configure this value on each DHCP scope in your organization. The servers listed in the value should be the servers that you wish managed devices within the DHCP scope to access. For instructions on how to configure the DHCP server option, see the next page.

Syntax

```
MgsDHCP( int limit, boolean discardForeign, int option )
```

where:

- `limit` is the maximum number of servers to prioritize using the algorithm. If limit is not specified (empty), all servers listed in the DHCP server option will be prioritized.

- `discardForeign` is a boolean flag indicating whether to discard servers not selected by the algorithm from

the failover list (`true`), or whether to keep the servers in the failover list but prioritize them lower than all servers selected by the algorithm (`false`). Defaults to `false` if not specified.

- `option` is an optional integer setting of the option configured on the DHCP server that contains the failover server information. Valid values are 0 to 255. The default is 123

For example: `MgsDHCP(,true,123)`

Example of MgsDHCP Algorithm Results

For DHCP server option set to `ds-prs-01.tmnis.org,ds-prs-02.tmnis.org`.

In the example shown, the fact that **AutoPriority** has been set to `False` for `ds-prs-01` prevents it from being given the highest priority, despite its pre-eminent position in the DHCP server option listing.

Location	Auto-Priority	Discard-Foreign	Priority before algorithm	Priority after algorithm
ds-cls-01	True	False True	1 1	2 Invalid
ds-prs-01	False	False True	4 4	4 4
ds-cls-02	True	False True	2 2	3 Invalid
ds-prs-02	True	False True	3 3	1 1

To Configure the DHCP Server Option

The DHCP option must be a string type and have an option number matching the one used by the NetSelector algorithm. You will typically need to configure this on each DHCP scope in your organization on which there are managed devices.

The syntax of the value is:

```
serverlist[,...n]
```

where `serverlist = [servername | random(servername[,...n])]`.

Some examples are:

- `srv1,srv2,srv3`
  Prioritizes srv1 first, followed by srv2 and srv3
- `random(srv1,srv2,srv3)`
  Prioritizes srv1, srv2, and srv3 in random order
- `random(srv1,srv2,srv3),srv4`
  Prioritizes srv1, srv2, srv3 in random order, followed by srv4

- `random(srv1,srv2,srv3),random(srv4,srv5,srv6)`
  Prioritizes srv1, srv2, and srv3 in random order, then srv4, srv5, and srv6 in random order

- `srv0,random(srv2,srv2,srv3),random(srv4,srv5,srv6)`
  Prioritizes srv0 first, followed by srv1, srv2, and srv3 in random order, then srv4, srv5, and srv6 in random order

- `srv0,random(srv1,srv2,srv3),srv4`
  Prioritizes srv0 first, followed by srv1, srv2, and srv3 in random order, followed lastly by srv4

The details about configuring a DHCP option will depend on the DHCP server that you are using. These instructions describe how to configure the DHCP option on a Microsoft DHCP Server:

1. Create a custom DHCP scope option;

   a. Start the DHCP MMC snap-in (**Start** > **Programs** > **Administrative Tools** > **DHCP**).
   b. Right-click the server name and select **Set Predefined Options...**
      The **Predefined Options and Values** dialog is displayed.
   c. Click **Add...**
   d. In the **Name** field, enter a descriptive name for the option.
   e. In the **Code** field, enter an option number such as 123. This is the option parameter that is passed to **MgsDHCP**.
   f. Set the **Data type** to string.
   g. Click **OK**.

2. Enable and configure the DHCP option:

   a. Start the DHCP MMC snap-in (**Start** > **Programs** > **Administrative Tools** > **DHCP**).
   b. Right-click **<Scope Name> Scope Options** and select **Configure Options**.
   c. Select the option and enter the server list in the **String** field.
   d. Click **OK**.

3. Repeat the previous step for every DHCP scope.

MgsDomainMatch: Match to Domain Name

This algorithm assigns priorities based on the domain name. The name closest to that of the managed device is given the highest priority.

MgsDomainMatch checks each domain component, from right to left.

Syntax

```
MgsDomainMatch( int limit )
```

where:

- `limit` is an optional integer setting the maximum number of servers to which priorities will be assigned by this algorithm.

Example of MgsDomainMatch Algorithm Results

For managed device domain: `abc.com.au`

Location	Domain	Auto-Priority	Priority before algorithm	Priority after algorithm	Normalized priority after algorithm
A	abc.com.au	True	Blank	6	1
B	abb.com.au	True	Blank	42	2
C	abc.com.de	True	Blank	83	3
D	abb.com.de	True	Blank	104	4
E	abb.com.de	False	5	5	5

In the example shown, the algorithm determines priorities in the following way:

- Location A matches the domain name of the managed device so it is assigned a very high priority (for example, 5). This priority is then indexed by the order of locations in the registry. It is first, so the priority assigned is 5+1 (6).

Be aware:
Because multiple locations may be assigned the same priority, Deployment Manager indexes priorities according to the order in which locations are listed in the registry.

- Location B matches the third and second domain components, but not the first. It is given a medium priority (for example, 40). After indexing, the priority is 40 + 2 (42).
- Location C does not match the third domain component, so it is given a lower priority (for example, 80). After indexing, the priority is 80 + 3 (83).
- Location D only matches the middle domain component and is given a very low priority (for example, 100). After indexing, the priority is 100 + 4 (104).
- Location E has a fixed priority, so priority 5 is excluded from the algorithm.
- Deployment Manager normalizes the priorities of locations A to D to fit within the range specified by the server selection settings. In this example, the range is 1-5.

MgsIPMatch: Match to IP Address

This algorithm assigns priorities based on similarities in the IP address. Address components are converted to binary numbers and compared, left to right. Priority is assigned according to the longest common (matching) bit in the binary number.

By comparing the IP address of the managed device against each server, any servers within a subnet will be given higher priority because the network address portion of the IP will be the same. If two servers are within the same subnet, then the value of the local host ID determines the priorities of the two servers.

Syntax

```
MgsIPMatch( int limit )
```

where:

- `limit` is an optional integer setting the maximum number of servers to which priorities will be assigned by this algorithm.

Example of MgsIPMatch Algorithm Results

For managed device IP address: `123.3.45.44`

Location	IP address	Auto-Priority	Priority before algorithm	Priority after algorithm	Normalized priority after algorithm
A	123.3.45.56	True	Blank	21	2
B	123.3.44.46	True	Blank	42	3
C	123.5.45.56	True	Blank	63	4
D	123.3.45.32	True	Blank	14	1
E	123.3.45.42	False	5	5	5

In the example shown, the algorithm determines priorities in the following way:

- Location A matches the first three domain components, so there is little binary difference from the managed device's IP address. The algorithm assigns a high priority (for example 20). It is then indexed by the order of location entries in the registry. Location A is first in the list, so the priority is recalculated as 20+1 (21)

Be aware:
Because multiple locations may be assigned the same priority, Deployment Manager indexes priorities according to the order in which locations are listed in the registry.

- Location B matches on the first and second address components and has a greater binary difference than location A. It is given a medium high priority (for example, 40). This is indexed and the priority is recalculated as 40 + 2 (42).

- Location C matches the managed device on the first domain component and the third component, but not on the last address component. It has a high binary difference and is given a low priority (for example, 60). This is indexed and the priority is recalculated as 60 + 3 (63).

- Location D is similar to location A because it doesn't match the last domain component of the managed device, but it does match on the first three domain components. It actually has a smaller binary difference, and is assigned a very high priority (for example, 10). It is then indexed by the order of location entries in the registry; it is fourth in the list, so the priority is recalculated as 10 + 4 (14).

- Location E has a fixed priority, so priority 5 is excluded from the algorithm.

- Deployment Manager normalizes the priorities to fit within the range specified by the server selection settings (described below). In this example, the range is 1-5.

MgsNameMatch: Match Prefixes of Computer Names

This algorithm compares the first characters of the host name of each server with the first characters of the name of the managed device on which the algorithm is running. Servers whose names match according to this check are prioritized above servers whose names do not match, while retaining the relative order of existing priorities.

The name of the managed device on which the algorithm is running is determined by the **MachineName** preference. For more information, see *RMS Preferences for Managed Devices*.

MgsNameMatch may be useful in situations where computers' names have a prefix determined according to their location or site. It is common to use this algorithm in conjunction with another algorithm, such as `MgsNameMatch(5,,true);MgsRandom`. This combination selects all servers that have at least five matching characters in their names, and then randomizes access (to achieve load balancing).

Syntax

`MgsNameMatch``( int matchLength, int limit, boolean discardForeign )`

where:

- `matchLength` is the number of characters to compare in the managed device and server host names
- `limit` is an optional integer setting the maximum number of servers to which priorities will be assigned by this algorithm (if this setting is not specified, no limit is imposed)
- `discardForeign` is an optional boolean represented by the case-insensitive strings `True` or `False`. The default is `False`.
  - If `false`, the algorithm sets the priorities of any servers which are not matched by the algorithm to a lower priority than all servers that are matched by the algorithm.
  - If `True`, the algorithm sets the priorities of non-matching servers to the literal string invalid. The installation agent and upload agent will ignore this server.

Example of MgsNameMatch Algorithm Results

For managed device **MachineName**: `Bost0014`

Server host name	match-Length	Discard-Foreign	Priority before algorithm	Priority after algorithm
BostonS01	5	True	5	Invalid
BostonS02	4	Don't care	44	2
ChicagS03	4	False	2	4
BostonS06	4	Don't care	1	1

Server host name	match-Length	Discard-Foreign	Priority before algorithm	Priority after algorithm
BostonS04	4	Don't care	4	3

In the example shown, the algorithm determines priorities in the following way:

- In the first row, the `matchLength` is set so high that the name matching fails on the fifth character. Since `discardForeign` is `True`, the priority after a failure must be set to invalid, and this server will be ignored. (In the remaining rows, the `matchLength` is corrected to 4.)
- All the remaining Boston servers pass the matching test, so their priorities are maintained in the same relative order.
- The Chicago server fails the match. Since `discardForeign` is `False`, it remains in the list, but its priority is set lower than all matching servers.

MgsPing: Fastest Response Server

This algorithm pings each server three times with a timeout of three seconds. An unlimited number of servers can be prioritized, however a maximum of 50 are 'pinged' in parallel at the same time. The general rule is that for 50 actively responding servers, the prioritization will be complete in fewer than five seconds.

The servers with the shortest response time are assigned the highest priority.

WARNING
Carefully consider the multiplier effect of using this algorithm on a large number of managed devices addressing a large number of servers. It could have a negative impact on network performance, particularly at tightly-scheduled policy update times. The `maxHops` parameter is useful for limiting the scope of pings.

Where pinging is desirable, consider also using the facility for random spread of update times, or combine it with another algorithm to limit the set of servers tested.

Syntax

```
MgsPing( int limit, boolean discardForeign, int maxHops )
```

where:

- `limit` is an optional integer setting the maximum number of servers to which priorities will be assigned by this algorithm.
- `discardForeign` is an optional boolean represented by the case-insensitive strings `True` or `False`. The default is `False`.
  - If `False`, the algorithm sets the priorities of any servers which do not respond to a lower priority than all servers that do respond.
  - If `True`, the algorithm sets the priorities of non-matching servers to the literal string invalid. The installation agent and upload agent will ignore this server.

- `maxHops` is an optional integer setting the maximum number of network segments that the ping will traverse. A value of 1 means that there must be no routers between the two computers. Valid values are 1 to 255. The default is 127.

Example of MgsPing Algorithm Results

Location	Average round trip time	Priority before algorithm	Priority after algorithm
A	100 ms	Blank	2
B	50 ms	Blank	1
C	200 ms	Blank	3
D	400 ms	Blank	4

MgsRandom: Random Priorities

This algorithm assigns random priorities to the download/upload servers, within the range specified by the **HighestPriority** and **LowestPriority** registry keys. For more information about these registry keys, see *Prioritizing distribution servers.*

This ensures that all managed devices configured to use the same list of locations do not use the same location for download and/or upload, spreading the load between servers.

Syntax

```
MgsRandom( int limit )
```

where:

- `limit` is an optional integer setting the maximum number of servers to which priorities will be assigned by this algorithm.

Example of MgsRandom Algorithm Results

For:
Managed device domain: `abc.com.au`
Managed device IP address: `123.3.45.44`

Location	Domain	IP address	Auto Priority	Priority before algorithm	Priority after algorithm
A	abc.com.au	123.3.45.56	True	Blank	3
B	abc.com.de	123.3.44.46	True	Blank	1
C	abb.com.au	123.5.45.56	True	Blank	2
D	abbb.com. de	123.3.45.32	True	Blank	4

Location	Domain	IP address	Auto Priority	Priority before algorithm	Priority after algorithm
E	abc.com.au	123.3.45.42	False	5	5

In the example shown, the algorithm determines priorities in the following way:

- Location E has a fixed priority, so priority 5 is excluded from the algorithm
- Priorities 1-4 are assigned randomly to all other locations

MgsServersFromAD: Retrieve Location List from AD

This algorithm prioritizes according to lists of servers specified in Active Directory. Consequently, it is not useful for:

- Managed devices that are not known to Active Directory
- Managed devices that are running legacy Windows operating systems, and that are not running the Active Directory client
- "Black box" implementations of Active Directory at the central administration server that do not include rolling out Active Directory throughout the enterprise.

Note:
To configure this algorithm, you need to use **ADSI Edit** (`adsiedit.msc`), a GUI tool that acts as a low-level editor for Active Directory.

MgsServersFromAD can be used with either client-side or server-side policy merging.

Syntax

```
MgsServersFromAD( int limit, boolean discardForeign, string dnPrefix )
```

where:

- `limit` is the maximum number of servers to prioritize using the algorithm. If this setting is not specified (empty), all servers listed in Active Directory will be prioritized.
- `discardForeign` is a boolean flag indicating whether to discard servers not selected by the algorithm from the failover list (`true`), or whether to keep the servers in the failover list but prioritize them lower than all servers selected by the algorithm (`false`). Defaults to `false` if not specified.
- `dnPrefix` is the prefix (quoted with double quotes) to add to a computer's distinguished name (DN), its subnet DN and its site DN in order to find an object in Active Directory with failover information. Failover information is obtained from the Description attribute in the first such object found in AD. Defaults to "CN=ManageSoft" if not specified.

As usual for distinguished names, any of the following special characters must be escaped with a backslash (\) character wherever they appear in distinguished name components (see following example):

```
, = + < > # ;
```

Additionally, any double quote (") characters in `dnPrefix` must be similarly escaped.

For example:
`MgsServersFromAD(,true,"CN=MGS Servers\, For Failover")`

Including this specification in the **SelectorAlgorithm** preference:

- Prioritizes all servers specified in the **Description** attribute of the first of the following objects that are found in AD:
  - `CN=MGS Servers\, For Failover,CN=<computer name>,<computer's OU/ container DN>`
  - `CN=MGS Servers\, For Failover,CN=<subnet>,CN=Subnets,CN=Sites,CN=Configuration,<domain DN>`
  - `CN=MGS Servers\, For Failover,CN=<site name>,CN=Sites,CN=Configuration,<domain DN>`
- Discards any servers not specified in Active Directory from the failover list.

When the managed device uses this algorithm, it looks up each of these three locations (computer name, subnet, and site name) in AD based on the configuration of that managed device. Each managed device calculates its own set of these three AD queries.

For:
Managed device MachineName: `mypc`
Managed device OU: `Desktops`
Managed device domain: `abc.com`
Managed device subnet: `172.16.34.0`
Managed device site: `melbourne`

Including this specification in the **SelectorAlgorithm** preference, the managed device looks for objects in AD in the following order:

- `CN=MGS Servers\, For Failover,CN=mypc,OU=Desktops,DC=abc,DC=com`
- `CN=MGS Servers\, For Failover,CN=172.16.34.0,CN=Subnets,CN=Sites,CN=Configuration,DC=abc,DC=com`
- `CN=MGS Servers\, For Failover,CN=melbourne,CN=Sites,CN=Configuration,DC=abc,DC=com`

If one of these objects is found, the managed device checks the **Description** attribute of this object and extracts the server information from this list. The failover server information must be a comma separated list of server hostnames (as specified in the RayManageSofti distribution hierarchy), using the following syntax:

`serverlist[,...n]`

where `serverlist = [servername | random(servername[,...n])]`.

Some examples are:

- `srv1,srv2,srv3`
  Prioritizes srv1 first, followed by srv2 and srv3

- `random(srv1,srv2,srv3)`
  Prioritizes srv1, srv2, and srv3 in random order
- `random(srv1,srv2,srv3),srv4`
  Prioritizes srv1, srv2, srv3 in random order, followed by srv4
- `random(srv1,srv2,srv3),random(srv4,srv5,srv6)`
  Prioritizes srv1, srv2, and srv3 in random order, then srv4, srv5, and srv6 in random order
- `srv0,random(srv2,srv2,srv3),random(srv4,srv5,srv6)`
  Prioritizes srv0 first, followed by srv1, srv2, and srv3 in random order, then srv4, srv5, and srv6 in random order
- `srv0,random(srv1,srv2,srv3),srv4`
  Prioritizes srv0 first, followed by srv1, srv2, and srv3 in random order, followed lastly by srv4

Example of MgsServersFromAD Algorithm Results

For:
Description attribute set to `ds-prs-01.tmnis.org,ds-prs-02.tmnis.org`.

Location	Auto-Priority	discard-Foreign	Priority before algorithm	Priority after algorithm
ds-cls-01	True	False True	1 1	2 Invalid
ds-prs-01	False	False True	4 4	4 4
ds-cls-02	True	False True	2 2	3 Invalid
ds-prs-02	True	False True	3 3	1 1

In the example shown, the fact that **AutoPriority** has been set to `False` for ds-prs-01 prevents it from being given the highest priority, despite its preeminent position in the Active Directory listing.

Taking Care Not to Orphan Managed Devices

When you set the `discardForeign` flag to `True`, any servers not found in Active Directory are discarded. There is a possibility that managed devices may become orphaned from Deployment Manager if the managed device's list of download locations does not contain any of the servers listed in Active Directory. If this occurs, the managed device will not attempt to download any packages, including any updated failover settings packages.

To avoid this situation, and still use the `discardForeign` option, it is recommended that you modify the **Update Failover Locations** scheduled event for managed devices to include an alternative algorithm for identifying the distribution server to use. In the parameters for this event, include an option similar to

```
-o ScheduleAlgorithm=MgsServersFromAD
```

For details about modifying the Update Failover Locations event, see the *Scheduling* chapter of the *RMS Software Deployment*.

If a value is not specified for `discardForeign`, it defaults to `False`. Whenever the managed device runs the scheduled event to update its failover locations (download and upload locations), the managed device will be able to use download locations that are not servers listed in Active Directory, if it cannot use those in Active Directory.

MgsSubnetMatch: Match to Subnet

This algorithm scans through all servers in the current subnet and moves them to the front of the priority list, while retaining the relative order of existing priorities.

This is particularly useful if combining with another algorithms, such as **MgsPing;MgsSubnetMatch**. The combination provides a resultant set of priorities based on the fastest responding servers in the current subnet, followed by the fastest responding servers outside the subnet.

Syntax

```
MgsSubnetMatch( int limit, boolean discardForeign )
```

where:

- `limit` is an optional integer setting the maximum number of servers to which priorities will be assigned by this algorithm.
- `discardForeign` is an optional boolean represented by the case-insensitive strings `true` or `false`.
    - When `true`, **MgsSubnetMatch** sets the priorities of any servers that are not in the same subnet as the managed device to the string literal invalid. The installation agent and upload agent will ignore this server.

Example of MgsSubnetMatch Algorithm Results

For:
Managed device with IP address: `172.16.34.53`
Managed device subnet mask: `255.255.248.0`
Managed device subnet: `172.16.34.0`

Location	Domain	IP address	Priority before algorithm	Priority after algorithm
A	LAN (100Mbps)	172.16.45.40	1	3
B	LAN (10Mbps)	169.15.13.12	2	4
C	WAN (56K Modem)	172.16.34.40	4	2
D	WAN (ADSL)	172.16.34.60	3	1

In the example shown, the algorithm determines priorities in the following way:

- Location C and D are in the same subnet as the client, thus have higher priority than A and B.

On Macintosh, Linux, and UNIX Managed Devices

Information about how to prioritize distribution servers for Macintosh, Linux, and UNIX managed devices is located in the `managesoft.xconf` file. This file is installed in `/etc/managesoft` by default. It is an XML file that you can edit using a text editor of your choice.

The following sections describe how you can configure managed devices for download or upload operations.

Download and Upload Locations

At installation time, you specify download and reporting locations for a managed device. (For details about the installation process, see the *Installing Deployment Manager on managed devices* chapter of the *RMS Implementation*.) These details are stored in XML format files:

- Download location is stored in
  `$(CommonAppDataFolder)/etc/download.xconf`
  (by default: `/var/opt/managesoft/etc/download.xconf`)

- Reporting location is stored in
  `$(CommonAppDataFolder)/etc/upload.xconf`
  (by default: `/var/opt/managesoft/etc/upload.xconf`)

To change any of these details, you can edit the files directly, or run `/opt/managesoft/bin/managesoft-configure`. The values you enter when you run this script overwrite any previous values.

Configuring the Download Locator

The download locator provides RayManageSofti with locations from which to download files. There are a number of algorithms available that govern the choice of download location:

- **ipmatchlocator** provides download locations in order based on IP address. The IP address of the distribution server closest to that of the device under management is returned first. **This is the default setting**.

- **nativepinglocator**, an alternative implementation of **pinglocator**, which uses ICMP sockets.

- **orderedlocator** provides download locations in the order in which they are specified in the `download.xconf` configuration file. This file is written to when managed devices are bootstrapped or managed device settings are deployed, so the order specified in this file is the same order used by the managed device to retrieve data from remote sources.

- **pinglocator** provides download locations in the order in which they respond to a ping (shortest response time first). This implementation uses UDP sockets, and requires UDP echo to be enabled. For more information, see *Troubleshooting* in the *Installing Deployment Manager on managed devices* chapter of the *RMS Implementation*.

  This is the recommended setting for low-bandwidth sites.

- **randomlocator** (the default algorithm) provides a download location by random selection from the available locations. It is not recommended for low-bandwidth sites (instead, use **pinglocator**).

The download locator is defined in `managesoft.xconf` in the following lines:

```
<!--+
| Download Locator
|
| The Download Locator's task is to provide valid download
| locations for the RayManageSofti client to use when retrieving files.
|
| The following locator types are available:
|
| * ipmatchlocator - returns locations in closest IP match order
| * randomlocator - randomizes the list of available locations
| * orderedlocator - returns locations in configuration order
| * pinglocator      - orders locations according to ping time
|      Note, this class require UDP ECHO to be enabled on the remote hosts
| * nativepinglocator - orders location according to ping time
|      Note, this class requires root privileges
+-->
<randomlocator id="download" default="true" logger="locator.download">
<file location="/var/opt/managesoft/etc/download.xconf"/>
<last-good location="/var/opt/managesoft/etc/lastdownload.xconf"/>
</randomlocator>
```

To use an algorithm other than randomlocator, replace the `<randomlocator ...>` section with one of `<ipmatchlocator>`, `<orderedlocator ...>`, `<pinglocator...>`, or `<nativepinglocator ...>`.

In all cases:

- `<file location="/var/opt/managesoft/etc/download.xconf"/>` specifies the location on the managed device in which to store the download location returned by the selection algorithm
- `<last-good location="/var/opt/managesoft/etc/lastdownload.xconf"/>` specifies the location on the managed device that stores the location of the last download location from which files were successfully retrieved

Configuring the Upload Locator

The upload locator provides RayManageSofti with locations to which to upload files. Its behavior is similar to that of the download locator, and its configuration is identical. There are five algorithms available that govern the choice of upload location:

- **ipmatchlocator** provides upload locations in order based on IP address. The IP address of the distribution location closest to that of the device under management is returned first. **This is the default setting**.
- **nativepinglocator**, an alternative implementation of **pinglocator**, which uses ICMP sockets.
- **orderedlocator** provides upload locations in the order in which they are specified in the `upload.xconf` configuration file. This file is written to when managed devices are bootstrapped or managed device settings are deployed, so the order specified in this file is the same order used by the managed device to retrieve data from remote sources.
- **pinglocator** provides upload locations in the order in which they respond to a ping (shortest response time first). This implementation uses UDP sockets, and requires UDP echo to be enabled. For more information, see *Troubleshooting* in the *Installing Deployment Manager on managed devices* chapter of the *RMS Implementation*.

This is the recommended setting for low-bandwidth sites.

- **randomlocator** (the default algorithm) provides an upload location by random selection from the available locations. This setting is not recommended for low-bandwidth sites (instead, use **pinglocator**).

The upload locator is defined in `managesoft.xconf` in the following lines:

```
<!--+
| Upload Locator
|
| The Upload Locator's task is to provide valid upload
| locations for the RayManageSofti client to use when uploading files.
|
| The following locator types are available:
|
| * ipmatchlocator - returns locations in closest IP match order
| * randomlocator - randomizes the list of available locations
| * orderedlocator - returns locations in configuration order
| * pinglocator      - orders locations according to ping time
|      Note, this class require UDP ECHO to be enabled on the remote hosts
| * nativepinglocator - orders location according to ping time
|      Note, this class requires root privileges
+-->
<randomlocator id="upload" logger="locator.upload">
<file location="/var/opt/managesoft/etc/upload.xconf"/>
<last-good location="/var/opt/managesoft/etc/lastupload.xconf"/>
</randomlocator>
```

To use an algorithm other than randomlocator, replace the `<randomlocator ...>` section with one of `<ipmatchlocator ...>`, `<orderedlocator ...>`, `<pinglocator ...>`, or `<nativepinglocator ...>`.

In all cases:

- `<file location="/var/opt/managesoft/etc/upload.xconf"/>` specifies the location on the managed device in which to store the upload location returned by the selection algorithm
- `<last-good location="/var/opt/managesoft/etc/lastupload.xconf"/>` specifies the location on the managed device that stores the location of the last upload location from which files were successfully uploaded.

Policy Merging and Distribution

This chapter describes the various ways of configuring your Deployment Manager environment to optimize server-side policy merging and distribution operations. In it, you will:

- Read about "resultant set of policy (RSoP) groups" and how they can speed up server-side policy merging
- Determine whether or not your organizational and Active Directory structures offer you a way to streamline policy distribution.

Note that this chapter is not relevant if you are exclusively using client-side policy merging.

RSoP Groups

Deployment Manager provides policy merging options which can dramatically reduce the time taken for server-side policy merging. For appropriately-structured environments, Deployment Manager can use "resultant set of policy (RSoP) groups" (calculated internally) to reduce computation.

Whether or not you can benefit from these policy merging enhancements depends on the range of managed devices you must support. Deployment Manager offers several compatibility options, and you must configure your administration server appropriately for your environment. Instructions for this configuration are provided in:

- *Configure support for policy merging enhancements* in both the *Installing a core server* and *Installing a combined server* chapters of the *RMS Implementation*
- *Configure support for policy merging enhancements* in the *Upgrading administration servers* chapter of the *RMS Upgrade*

Using those instructions, you can reconfigure your RSoP grouping support at any time (not only during installation or upgrade). For example, if you upgrade older managed devices, you might immediately change your RSoP grouping support setting to take advantage of policy merging enhancements previously unavailable to you.

The RSoP grouping support setting also affects the format and location of merged policy (`.npl`) files, as noted in the *File formats* chapter of the *RMS System Reference*.

Policy Distribution

In general operation, after server-side policy merging has been performed, all merged policy files are distributed to all distribution servers. However, in environments where the RayManageSofti distribution hierarchy closely matches the Active Directory OU structure, with a distribution server at each Active Directory site, it is possible to streamline the distribution of merged policy files. This is because merged policy file names contain information about the managed devices and users to which they are targeted, including their domains and organizational units. This makes it possible to determine which merged policy files will be required on which distribution servers.

You can configure distribution servers so that they only receive merged policy files required by managed devices and users in the site served by the distribution server.

To do so, set the following registry key value on each distribution server:

```
[Registry]\ManageSoft\Policy\CurrentVersion\DistributedPolicyArchives
```

Set **DistributedPolicyArchives** to a string consisting of a pipe (|) separated set of values. Distribution servers will then download only merged policy files whose names contain any of the specified values, ignoring merged policy files whose names do not match.

Example

```
DistributedPolicyArchives= OU=Office1,OU=Offices,tmnis.org_domain|tmnis.org_domain_rsop
```

The distribution server will download any archive files containing either of the two strings (`OU=Office1,OU=Offices,tmnis.org_domain`, or `tmnis.org_domain_rsop`). This means that the distribution server will download all merged policy files:

- For the OU `Office1` and any OUs under `Office1`
- Containing `tmnis.org_domain_rsop` in the filename. This will pick up all RSoP merged policy files for the `tmnis.org` domain. These are files calculated internally by Deployment Manager to reduce repetitive computation during policy merging. They may be referenced from any policy in the domain, and therefore it's a good idea to add `_rsop` to the **PolicyArchives** string, to make sure all RSoP files are available on every distribution server.

Automatic Policy Merge

Certain actions carried out within RayManageSofti require a policy merge to proceed further. These actions include starting a Rollout Plan, staring a wave within a Rollout Plan, and adding package(s) or bundle(s) to a policy. By default, RayManageSofti carries out this task by automatically executing a policy merge in the background.

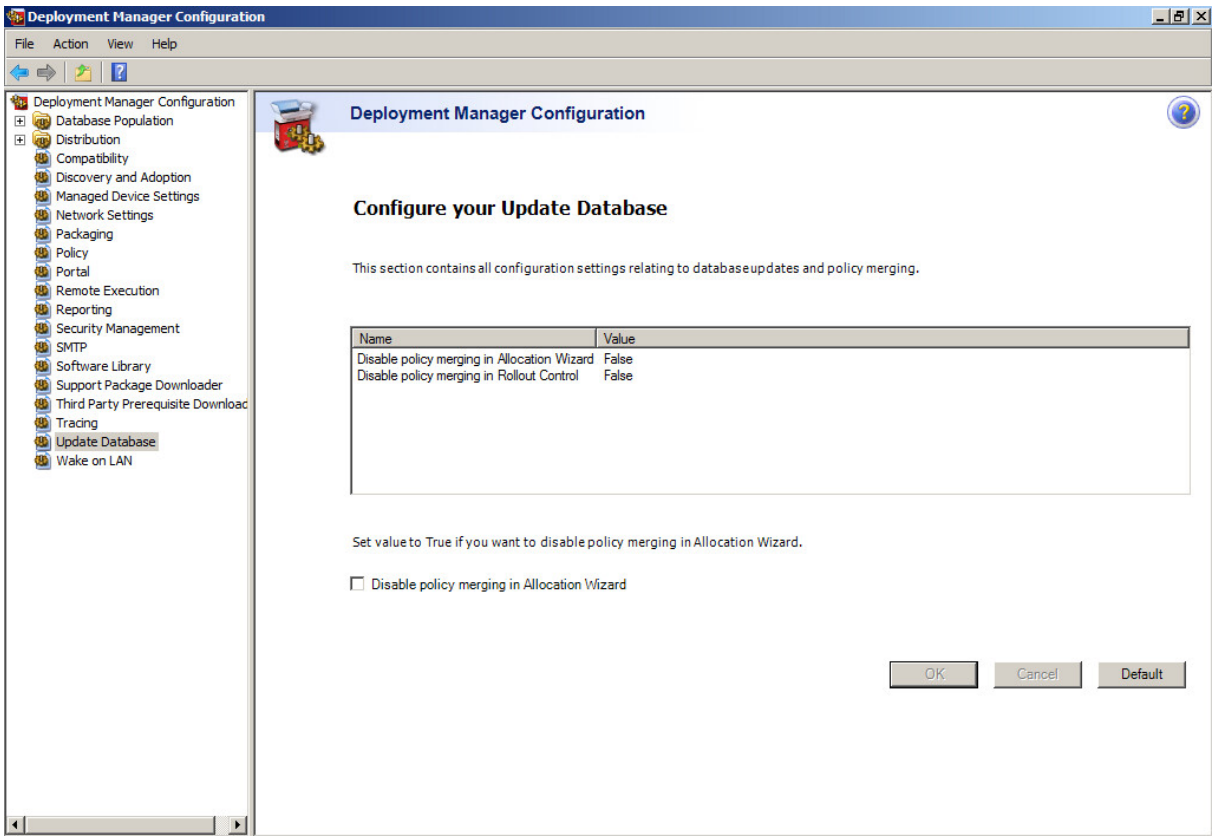

In certain scenarios within an enterprise environment, it may be desired to have an option, where this automatic policy merge could be disabled. This can be done in the Deployment Manager Configuration by checking the **Disable policy merging in Allocation Wizard** checkbox in the **Update Database** node.

Be aware:
Please use this option wisely as it may not be possible to proceed further with certain actions that require a policy merge.

Configuring Environments with Multiple Domains

If you manage software across multiple Active Directory domains, some domain data is stored in the RayManageSofti SQL database. Domain data is obtained and stored in the database:

- During installation
- During upgrade from a previous release of Deployment Manager
- During Active Directory reconciliation

Configuration Options

You can configure how Deployment Manager handles domains. You may:

- Correct the automatically-detected domain type
- Set whether to search for alternate domain controllers
- Set whether to include each domain in reconciliations with the RayManageSofti database
- Specify whether or not to include each domain when Deployment Manager performs policy merges
- Within each domain, specify whether merged policy should be generated for users, computers, or both

To Configure Domains

1. From the **Settings** node in Deployment Manager, select the **advanced** tab and click on the **Open multiple domain configuration tool** button.

   When this application starts, it queries the Active Directory global catalog and populates the RayManageSofti database with details about all domains in the global catalog. It does not update the details of domains that already have details recorded in the database.

   The **Deployment Manager Domain Configuration** dialog is displayed.

   The **Domain name** field shows the name of the domain currently selected in the list of **Domains** on the left.

   If any managed devices reference the domain using the NT 4 version of its name, that version of the name is displayed in the **NT 4 domain name** field. Note that this information is retrieved from managed device inventory records, and this field displays the first such mapping discovered in the inventory.

2. In the **Domain type** field, choose the type of this domain (*Active Directory* or *Unspecified*).

   The **Domain controller** field shows the last domain controller accessed for details about this domain. The field is blank if RayManageSofti could not identify the domain controller from inventory records.

3. Use the **When this domain controller is unavailable, do not fail-over to alternative servers** check box to specify whether or not to access an alternative domain controller for details about this domain if the identified domain controller is not available.

4. Use the **Computers** and **Users** check boxes to specify whether policy merging and reconciliation should

apply to users, computers, or both.

These check boxes do not configure any reconciliation or policy merging activities: you must specify what actions should occur using the **Reconcile** and **Policy** tabs (see below).

See *How user and computer selections affect reconciling and merging* for more details about your Computers and Users selections on the **Identity** tab affect reconciliation and policy merging.

Be aware:
The policy merge also updates RayManageSofti access rights with any changes to Active Directory group memberships. However, if you clear the **Users** check box, this step is removed from the policy merge.

For any domain that contains users who may wish to access to the RayManageSofti console or reports, make sure you select the **Users** check box. If you clear this check box, the users currently in the group may not match the users with access rights.

5. When you have finished making changes, select **File** > **Save**.

   Your domain configuration is saved.

6. If you do not want to make further changes to this or any other listed domain, select **File** > **Exit**.

To Set Reconciliation Options

1. On the **Deployment Manager Domain Configuration** dialog, click the **Reconcile** tab.

   The current settings are shown for Active Directory reconciliation for the domain selected in the list of Domains on the left.

2. Use the **Reconcile the Deployment Manager database with Active Directory for this domain** check box to set whether or not this domain is included in the RayManageSofti automatic reconciliation between the RayManageSofti SQL database and Active Directory.

   Normally leave this check box set (on), as this reconciliation provides accurate reports in the RayManageSofti console. You might clear the check box (off) if you are not managing the selected domain with RayManageSofti.

   If you selected **Computers** on the **Identity** tab, and you select **Reconcile the Deployment Manager database with Active Directory for this domain**, Active Directory data about all computers in this domain will be reconciled with the RayManageSofti database.

   If you selected **Users** on the **Identity** tab, and you select **Reconcile the Deployment Manager database with Active Directory for this domain**, Active Directory data about all users in this domain will be reconciled with the RayManageSofti database.

   See *How user and computer selections affect reconciling and merging* for more details about your Computers and Users selections on the **Identity** tab affect reconciliation and policy merging.

3. If merged policy calculations should use groups and memberships from the root domain, you need to reconcile data from the root Active Directory domain with the RayManageSofti database. If you do not want merged policy calculations to use groups and memberships from the root domain, clear the **Include groups and memberships from the root domain** check box so that you do not attempt to reconcile this data.

You might choose to clear this check box if:

- RayManageSofti has limited access to Active Directory. If RayManageSofti cannot access the root directory, and this check box is selected, reconciliation will fail.
- You reconcile policy data for all domains at the same time. In this instance, it is time-consuming and unnecessary for each domain to reload root domain data, so you could clear this check box for all domains other than the root domain.

(The Merge Deployment Manager policies scheduled task, discussed in the *Scheduled tasks* chapter of *RMS System Reference*, reconciles data, and generates merged policies, for all domains configured for reconciliation and policy merging.)

4. Discovered devices can be automatically assigned to sites based on IP addresses. If you will assign discovered devices to sites defined in Active Directory, select the **Reconcile the Deployment Manager database with sites from Active Directory for this domain** check box.
5. The **Results from the last reconciliation panel** shows:
   - The time and date that this domain's data in the RayManageSofti database was last reconciled with Active Directory
   - Whether or not the last Active Directory reconciliation succeeded
6. When you have finished making changes, select **File** > **Save**.

   Your domain configuration is saved.
7. If you do not want to make further changes to this or any other listed domain, select **File** > **Exit**.

To Set Policy Merging Options

1. On the **Deployment Manager Domain Configuration** dialog, click the **Policy** tab.

   The current settings are shown for the domain selected in the list of domains on the left.
2. Use the **Merge Deployment Manager policies for this domain** check box to specify whether to generate merged policies for this domain.

   You could clear this check box if every managed device in the selected domain is using client-side policy merging, or if you are not managing this domain with RayManageSofti. If any of the managed devices in this domain (including, for example, mobile computers) are set for server-side policy merging, be sure that this check box is selected.

   If you selected **Computers** on the **Identity** tab, and you select **Merge Deployment Manager policies for this domain**, merged policies will be generated for all managed devices in this domain.

   If you selected **Users** on the **Identity** tab, and you select **Merge Deployment Manager policies for this domain**, merged policies will be generated for all users in this domain.
3. The **Results from the last policy merge** panel shows:
   - The time and date of the last merged policy generation for this domain

- Whether or not the last merged policy generation succeeded

4. When you have finished making changes, select **File** > **Save**.

   Your domain configuration is saved.

5. If you do not want to make further changes to any other listed domain, select **File** > **Exit**.

How User and Computer Selections Affect Reconciling and Merging

It is helpful to work through an example of how the User and Computer selections you make on the **Identity** tab affect reconciliation and policy merging operations. The example relates to user data. The outcomes for computer data are parallel.

The example refers to the generation of merged policies, which assumes that server-side policy merging is in operation. If client-side policy merging is being used, much of the same data is processed, as Deployment Manager reporting relies on knowledge of the resultant set of policy (RSOP) for each user and managed device.

Imagine that the RayManageSofti database currently contains records about users Alexandra, Fernando, and Denis.

Since the RayManageSofti database was last reconciled with Active Directory, user Fernando has been removed from Active Directory.

The following table outlines the outcomes when each combination of options is selected.

Users check box (Identity tab)	Reconcile... check box (Reconcile tab)	Merge... check box (Policy tab)	Result
Selected	Selected	Selected	• User Fernando removed from RayManageSofti database • Merged policies generated for users Alexandra and Denis
Selected	Cleared	Selected	• Merged policies generated for users Alexandra, Fernando, and Denis.
Selected	Selected	Cleared	• User Fernando removed from RayManageSofti database • No merged policies are generated. Existing merged policies for users Alexandra, Fernando, and Denis remain in effect
Cleared	Selected	Selected	• No changes to user data in the RayManageSofti database • No updates to Deployment Manager user access rights • Merged policies generated for users Alexandra, Fernando, and Denis

Users check box (Identity tab)	Reconcile... check box (Reconcile tab)	Merge... check box (Policy tab)	Result
Cleared	Cleared	Selected	• No changes to user data in the RayManageSofti database • No updates to Deployment Manager user access rights • Merged policies generated for users Alexandra, Fernando, and Denis
Cleared	Selected	Cleared	• No changes to user data in the RayManageSofti database • No updates to Deployment Manager user access rights • No merged policies are generated. Existing merged policies for users Alexandra, Fernando, and Denis remain in effect

Configuring Byte-level Differentiation

Byte-level differentiation is a way of minimizing network transfers. Using patent-protected algorithms, it compares the contents of an existing package with its desired state, and determines the smallest possible transfers to bring the local copy to that state. The "desired state" may mean upgrading to a new version of a package to meet changed policy, or it may mean repairing an existing package (self-heal).

There is an extensive overview of byte-level differentiation in the *RMS Software Deployment*, in the *System overview* chapter (look for the *Download differentiation* section in the discussion of the distribution system). This chapter assumes that you are familiar with that overview.

In this Chapter

This chapter focuses on the configuration of the byte-level differentiation feature. It provides sufficient guidance for you to understand when byte-level differentiation should operate, and to test the savings in network transmissions when it does so.

Since byte-level differentiation imposes some small computational overhead, it is designed to be easily configurable to meet your requirements. Out of the box, it is configured to operate only when the balance between its computational overhead and network transmission costs is favorable. In other words, by default Deployment Manager always takes the fastest option - any time that it would take longer to compute differences than to transmit a whole file, Deployment Manager transfers the file. Only if byte-level differentiation will save time it is invoked.

However, there may be circumstances where you want to minimize network transmissions at the expense of slightly longer elapsed times per downloading computer. Deployment Manager allows you to configure byte-level differentiation to meet such needs.

The first part of this chapter focuses on the conditions necessary for byte-level differentiation to run at all. Included throughout this section are details about modifying the configuration. Once you have these conditions in place, so that byte-level differentiation can operate, you can use the later part of this chapter to guide your testing of byte-level differentiation, and subsequent trouble-shooting.

Preconditions for Operation

There are several prerequisites for the operation of byte-level differentiation. These are summarized in the following flow diagram, and discussed in this section.

Note:
The flow diagram is a diagnostic chart, and does not represent code operations.

Summary File

As discussed in the overview of byte-level differentiation in the *RMS Software Deployment*, the Deployment Manager algorithm does not require that differencing files between any two package versions are maintained. Instead, for each file for which byte-level differentiation is required, a small summary file is included as part of the package.

In general, you should never need to interact with summary files, as their generation and handling is automatic. They are mentioned here because they will help your understanding of the following sections.

Each summary file is generated automatically on the administration server at packaging time. Attributes giving its location (`updatehref`) and size (`updatesize`) are added to the implementation archive (`.ndc` file) at packaging time. Because it is listed in the implementation archive, it is distributed as expected with all other elements of the package, and sits in the same directory as the file it summarizes. (Neither compression nor byte-level differencing is ever applied to the summary file itself.) It can be identified by the filename extension `.bls` added to the filename of the file it summarizes. For example, a compressed DLL named `Agm.dll.gz` will have an accompanying summary file `Agm.dll.gz.bls`.

It is a reasonable debugging step to check that `.bls` files are present in distribution locations for files that you are testing. Do not allow manual cleanups of distribution servers and distribution locations to remove `.bls` files.

Be aware:
Summary files are handled slightly differently on distribution servers and on managed devices. A distribution server requires that the summary file of the old version is valid and present with the old version of the package in its local repository of packages. It does not recalculate summary files, making the remaining calculations for byte-level differentiation somewhat faster than they would otherwise be.

In contrast, managed devices allow for a greater risk of file corruption, and recalculate the summary files for installed versions.

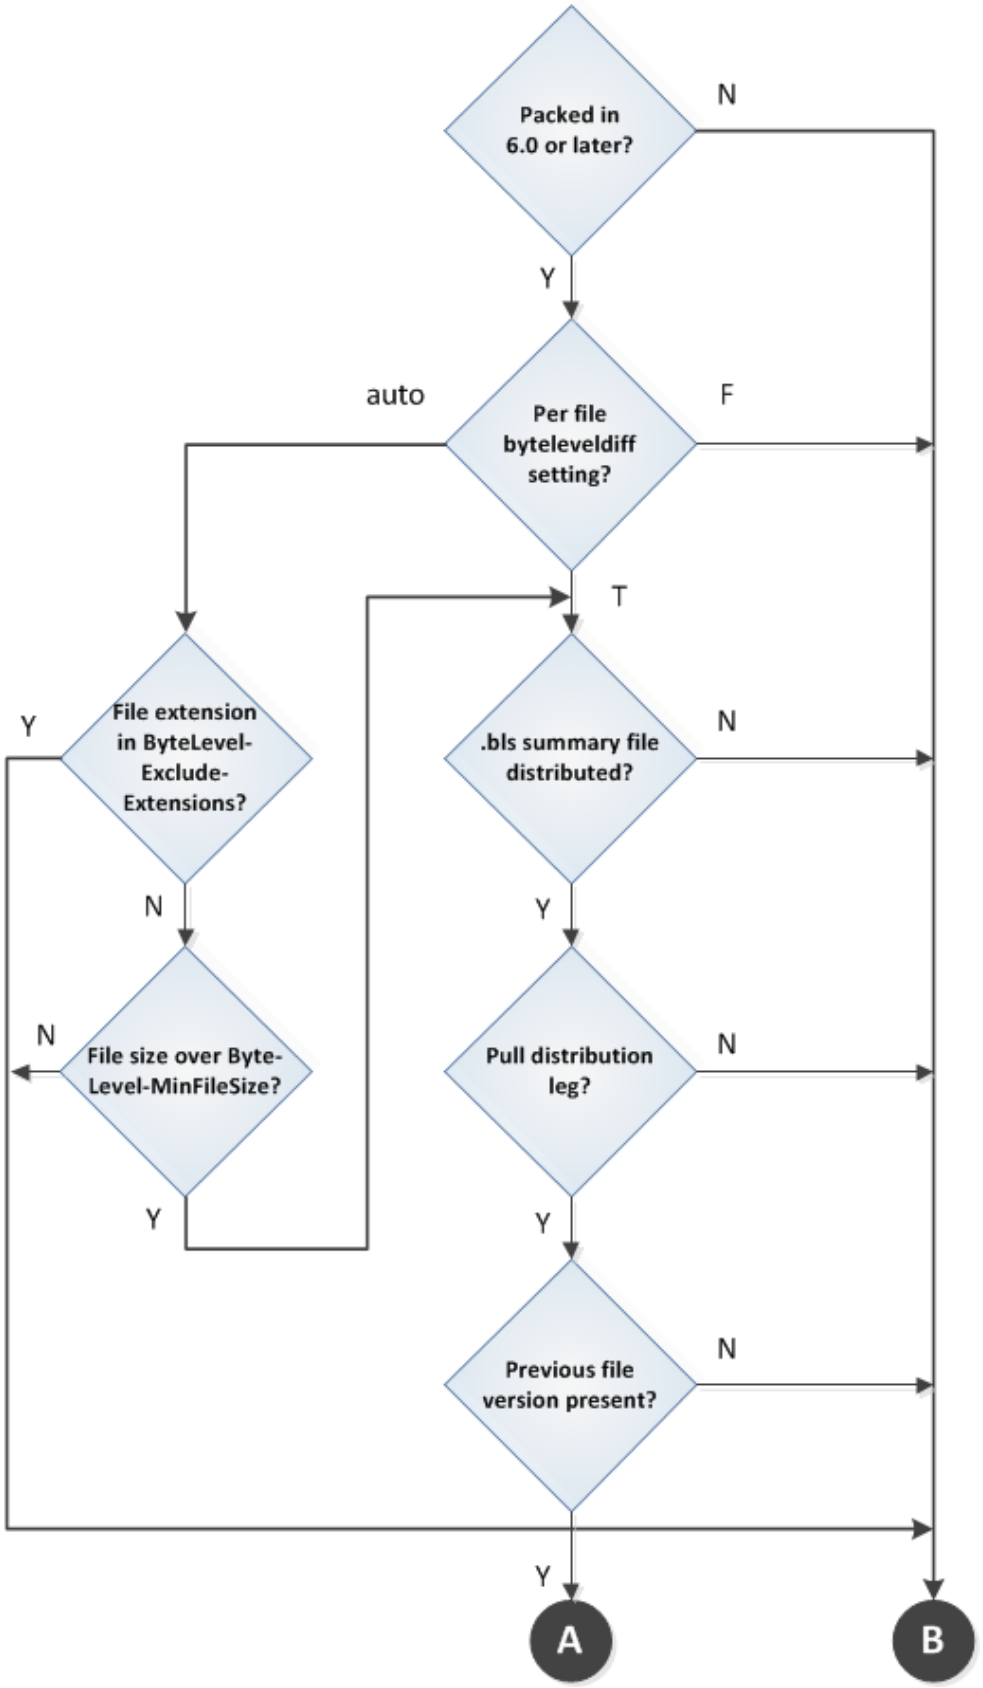
Packed in 6.0 or later?
N
Y
auto
Per file byteleveldiff setting?
F
T
File extension in ByteLevel-Exclude-Extensions?
Y
N
.bls summary file distributed?
N
Y
File size over Byte-Level-MinFileSize?
N
Y
Pull distribution leg?
N
Y
Previous file version present?
N
Y
A
B

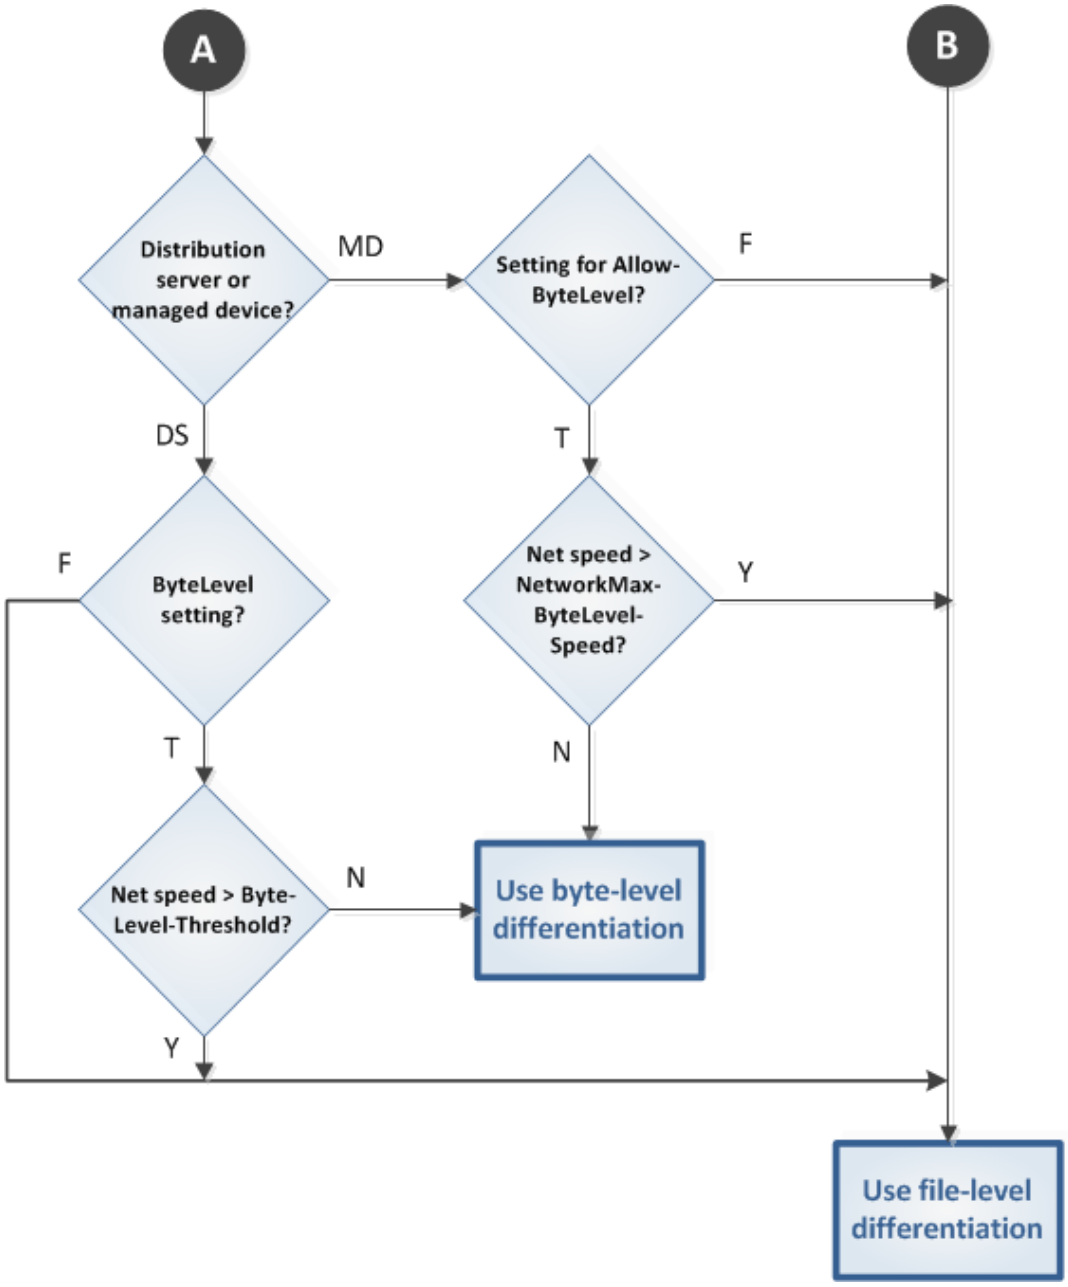

Product Version

The first requirement for operation is that byte-level differentiation is at all possible.

- Any package on which byte-level differentiation is required must be packed in a console of Deployment Manager 6.0 or later. Support for byte-level differentiation was added to the packaging agent at release 6.0.

Be aware:
RayManageSofti requires adequate virtual memory on the downloading computer to perform byte-level differentiation. As a guideline, the virtual memory requirement is something less than 2% of the per-file download size. For example, 4Mb of available virtual memory is adequate for byte-level differentiation of a downloaded file size of about 256Mb. If downloading computers do not have

adequate virtual memory, you should disable byte-level differentiation on those computers. For more information on disabling byte-level differentiation, see *Distribution leg configured* and *Managed device configured*.

Updating Earlier Package Formats

You may have in production packages prepared in NETDEPLOY GLOBAL 5.5 or earlier versions. (Prior to release 6.0, Deployment Manager was known as NETDEPLOY GLOBAL.) There are two possible paths for upgrading these packages to take advantage of byte-level differentiation at the *next* upgrade. The simpler path allows for byte-level differentiation only over the final leg, downloading to the managed device; while a slightly more complex process provides for byte-level differentiation to distribution servers as well. (Both paths provide for full byte-level differentiation of subsequent versions. Only the immediately forthcoming update is affected by the upgrade path you choose.)

Managed Device Byte-level Differentiation Only

In this simple path, you:

- Make no change to existing packages.
- Prepare the new version of each package in RayManageSofti 11.4 *infinity*.
- Distribute the new version in your usual manner.

Deployment Manager distributes the new packages to distribution servers and distribution locations using file-level differentiation. This means that only changed files are transmitted, but for any change in a file, the entire file is transmitted.

Managed devices (using RayManageSofti 11.4 *infinity*) will use byte-level differentiation to download the updated package versions. (For more information on the differences between downloads to managed devices and distribution servers, see *Summary file*.)

All future updates to these packages will use byte-level differentiation system-wide.

System-wide Byte-level Differentiation

If you need to take immediate advantage of byte-level differentiation in downloads to distribution servers for your next upgrade, follow this procedure:

1. In the RayManageSofti console, open the project for any preexisting package you need to upgrade in this way. (This is the *old* version of each package, not the upgrade.)

   Deployment Manager automatically adds to the project file and implementation archive the extra file attributes required for byte-level differentiation, and sets them to practical defaults.

2. Making *no other changes* to the projects, repack and redistribute.

   Since the application files are unchanged, they are neither redistributed nor updated. However, the summary files used for byte-level differentiation are new in each package, and are distributed. For more information, see the section on *Summary file*.

3. Now prepare, pack, and distribute the new versions of any packages as usual.

   Since the older versions have already been upgraded appropriately, byte-level differentiation will now be applied system-wide, on distribution servers as well as on managed devices.

All future updates to these packages will use byte-level differentiation system-wide.

File Type and File Size

Deployment Manager applies byte-level differentiation on a file-by-file basis, within any package. Within the package project (`.ndp` file), each file has an attribute `byteleveldiff` that may have one of three values:

- `True`
  byte-level differentiation is applied to this file
- `False`
  byte-level differentiation is not applied to this file
- `Auto`
  Deployment Manager determines whether to use byte-level differentiation or not.

This attribute is available on the administration server through the **Resources** dialog for any file within a package of the **software library** under the **Software** node. Add or edit a file and navigate to the **general** dialogue. Within the user interface, the three values of the attribute are represented as **Yes**, **No**, and **Auto** within the **Distribution options** section.

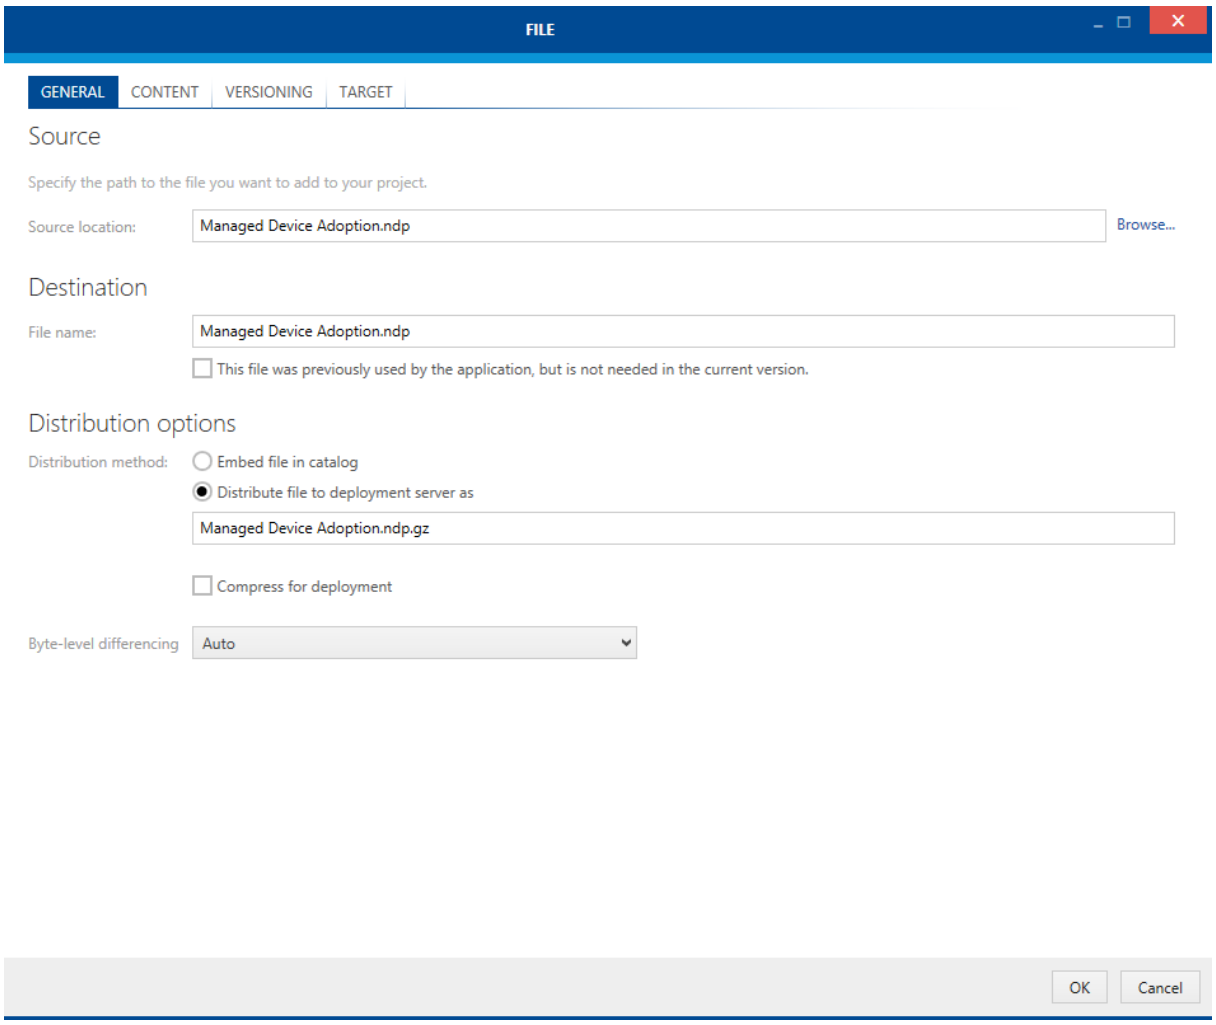

For any file for which the setting is **Auto**, Deployment Manager determines at packaging time whether to use byte-level differentiation according to two characteristics of the file:

- If the file name extension is included in a list of file types for which byte-level differentiation is prohibited, the summary file will not be calculated.

  The default list of file types is `jpg;gif;png`. These files are normally excluded because they are image files liable to radical change at each revision, making byte-level differentiation inefficient.

- If the file size is under the minimum value for byte-level differentiation, the summary file will not be calculated.

  The default minimum size for byte-level differentiation is 32k. This minimum has been determined as a likely point below which the calculations for byte-level differentiation may take longer than downloading the whole file.

To summarize the positive case, if the file is a permitted type and the size is above the minimum specified, a summary (`.bls`) file is generated, and referenced in the implementation archive. For more information, see *Summary file*.

Downstream from the administration server, distribution servers and managed devices use the presence of the summary file to determine whether to use byte-level differentiation when downloading this file.

Pull Distribution Leg

File transfer in the RayManageSofti distribution hierarchy may be configured for pull (target fetch) or push (source-driven) over any leg of the distribution hierarchy.

Only pull distribution legs support byte-level differentiation. This means that byte-level differentiation will not operate over any of the following:

- Connections to older HTTP or HTTPS servers that do not support byte ranges as part of their fetch requests
- The connection from a distribution server to a distribution location (except, of course, for a linked distribution location that is physically located on the active distribution server)
- Any distribution legs configured for push distribution.

Be aware:
While byte-level differentiation is available for FTP connections, it can offer less benefit than other file transfer methods. This is because FTP does not support the ability to transfer a specific block of a file. It does allow file transfer from an identified start location in a file, but must transfer from that start location through to the end of the file. For changes near the end of a file, the benefits of byte-level differentiation can still be realized. But when differences occur at the beginning of a file, the entire file must be transferred.

Naturally, pull-distribution legs over which byte-level differentiation may operate include the last leg from the distribution location to the managed device.

All other things being equal, the availability of byte-level differentiation is a very strong argument in favor of configuring for pull distribution, since it can greatly reduce network traffic.

File Update

To state the obvious, byte-level differentiation can operate only where a package is updating an earlier version and the earlier version exists on the downloading distribution server or managed device. The whole principle of byte-level differentiation is to download only the differences needed to bring the earlier version to its new desired state, so that the existence of the earlier version is a prerequisite.

You identify different versions of packages in the software library during the packaging process. Within the distribution hierarchy, Deployment Manager supports both:

- Replacement of the old version with a new version in the same directory
- Parallel distribution of two versions using different directories (this allows for progressive rollout of a new version to pilot production groups while the old version is still in use on other managed devices)

In both cases, the new version is constructed by combining the common parts of the two versions (from files already available locally) with the downloaded differences.

Where a previous version of a file does not exist on a particular distribution server or managed device, RayManageSofti downloads the entire file.

Platform Differences in File Updates

On Linux, UNIX, and Macintosh platforms, the name of a file may change when a package is upgraded. For example, RPM packages contain a version name as part of the file name.

In this situation, RayManageSofti must determine which software file is being upgraded. To do this, it uses fuzzy logic to check for matches between an upgraded package file name and original files names. This ensures that RayManageSofti correctly upgrades software even when there are differences in the package file names.

On these platforms, byte-level differencing only occurs when a package version is incremented. It does not apply if a package is updated and repacked but the version number of the package is unchanged. Also see *RMS Preferences for Managed Devices* for details about configuring byte-level differencing on these platforms.

Distribution Leg Configured

Each leg of your distribution hierarchy that is configured for pull distribution is normally configured for byte-level differentiation. This is because the calculation of differences is faster on distribution servers than on managed devices, so that the overheads are reduced (for more information, refer back to *Summary file*). This means that byte-level differentiation is cost-effective for distribution hierarchies even across fast LAN links.

It is possible to turn off byte-level differentiation if required. There are two attributes for each server listed in the `hierarchy.cfg` file, which determine this behavior:

- `BYTELEVEL` is a boolean set by default to `True`. If for some reason you wish to permanently turn off byte-level differentiation for a particular server, set this to `False`.
- `BYTELEVELTHRESHOLD` is the limit (in kilobits per second) of network speed above which byte-level differentiation is to be temporarily disabled for this server. A special-case value of zero means that byte-level differentiation is always enabled.

  `BYTELEVELTHRESHOLD` may be set to the speed above which the delays for differencing computations are in

the same order as network latency, and it becomes as efficient to transfer the whole file. On distribution servers, RayManageSofti compares this threshold with the network bandwidth the distribution server is currently allowed to use over the appropriate network link. This allowance may vary over time, depending on the bandwidth optimization settings you made on the administration server. If the current allowance is less than the `BYTELEVELTHRESHOLD` value, byte-level differentiation may operate.

However, byte-level differencing calculations on distribution servers are faster than on managed devices (for more information, see *Summary file*). This means that for distribution servers, byte-level differentiation is almost always more efficient than file-level differentiation. For this reason, the default setting on distribution servers is zero, so that byte-level differentiation is always enabled.

Disabling byte-level differentiation for a particular distribution server does not affect its use further down the distribution hierarchy. The necessary package attributes and summary files are transferred in any case, and are available for use by other distribution servers or managed devices.

If for some reason you wish to disable byte-level differentiation for one or more servers, use the following procedure:

1. Ensure that the servers have been defined on the administration server as usual.
2. Close the RayManageSofti console.
3. Locate the `hierarchy.cfg` file on your core administration server. The default location is `C:\ManageSoft\Repository\DeploymentLocations\Common\hierarchy.cfg`.
4. Edit the file in a text editor, searching for the friendly name of the server(s) in question to locate each entry within the file.
5. To unconditionally disable byte-level differentiation for each server, changes its boolean attribute to `BYTELEVEL="False"`.
6. To disable byte-level differentiation above some level of network performance, set the bandwidth parameter to your chosen number of kilobits (not kilobytes).
   For example: `BYTELEVELTHRESHOLD="10000.0"` turns off byte-level differentiation for an (unthrottled) 10Mbit/second network, but leaves it enabled for lower-speed networks.
7. Reopen the RayManageSofti console, and refresh the distribution hierarchy.

To repeat: the recommended practice is to leave byte-level differentiation always enabled for distribution servers.

Managed Device Configured

Just as there are two settings for configuring byte-level differentiation to distribution servers, so there are parallel settings for the managed device. These settings are exposed as preferences on the managed device (for more information, see *RMS Preferences for Managed Devices*):

- `AllowByteLevel` is a boolean set by default to `True`. If for some reason you wish to permanently turn off byte-level differentiation for a particular managed device, set this to `False`.
- `NetworkMaxByteLevelSpeed` is the limit (in bytes per second) of estimated network speed above which byte-level differentiation is to be temporarily disabled for this managed device.

The estimated network speed is established using ping testing (once for each new distribution location from

which this device attempts to download), and takes into account bandwidth optimization settings. It is the speed above which the delays for differencing computations are in the same order as network latency, and it becomes as efficient to transfer the whole file. The default setting is 262144 bytes (about the speed of a 2Mb/s WAN). This means that byte-level differentiation does not normally operate on lightly-loaded LAN networks. It comes into effect when a network is heavily loaded, or across slow links. Once the decision is made and file transfer commences, network performance is not re-evaluated during this download. (If there is a server failure, and the managed device fails over to a new distribution location, the network performance is re-evaluated for the new location.)

For more information on updating these (and other) preferences for managed devices, see *RMS Preferences for Managed Devices*.

Note:
It is possible for Raynet consultants to disable the download staging area on managed devices. Doing this is deprecated for production environments. For completeness, we note here that turning off the staging area also disables byte-level differentiation. Staging area settings should not be an issue in any live environment.

Managed Device Operation

When RayManageSofti determines that a file is to be downloaded to a managed device using byte-level differentiation, it recalculates the summary file(s) for the installed (old) version of the software. This allows for the possibility of corruption in any files while the software has been installed on the managed device. The summary calculated for the old, installed version of each file is then compared with the summary file for the new version, downloaded first from the distribution location. From these summary files, RayManageSofti calculates the bytes required to bring the current state of the application on the managed device to its desired state (new version).

The RayManageSofti technology transfers the smallest practical block of bytes needed to create the new file version, regardless of where the source file (new version) differs from the reference file (old version). For example, if the original file is represented like this:

```
0123456789
```

then a file changed in this manner:

```
x123456789
```

will transfer about the same number of bytes as a file changed in this manner:

```
012345678x
```

or, indeed, this:

```
01234x56789
```

The position of changes in the file does not determine the size of the transfer. In particular, notice in the last example that bytes moved down the file by an earlier insertion are not downloaded.

The patented download algorithms operate equally well on uncompressed files and compressed files (without requiring decompression and recompression). Download features like automatic resumption after an interruption and automatic failover in the event of server problems continue to work with byte-level differentiation.

Testing Byte-level Differentiation

The descriptions in the previous sections of this chapter make clear that there are many elements that must come together before byte-level differentiation becomes operational.

You may wish to test byte-level differentiation and measure the differences in network transfer sizes. To do this, you must ensure that byte-level differentiation operates both for the test file in general and in particular over the individual distribution leg under test.

The easiest download to manage is over the final leg to the managed device, as here the command-line tools are readily available with appropriate command-line options to facilitate testing. For information about command-line tools, see *RMS System Reference*. For more detail about individual preference options, see *RMS Preferences for Managed Devices*.

Offline Testing

You may conduct preliminary testing of two versions of a package to determine what the downloads will be using byte-level differentiation. This can help determine whether you should set byte-level differentiation for a particular file or not.

A utility for this purpose is included on the RayManageSofti Deployment Manager product DVD. It is available at `<...>\ManageSoft\Tools\BLD comparison\mgsbdcmp.exe`. More extensive documentation for the tool is also supplied on the DVD.

This utility is run on the command line, passing the names (in order) of two file versions to be tested. The files must be raw (uncompressed). If they will eventually be compressed in the distributed package, use the `-z` option, and the utility will calculate the differences as if the files were compressed.

The complete command line is therefore: `mgsbdcmp.exe [-z] fileversion1 fileversion2`

The output from the utility takes the following form:

```
C:\> mgsbdcmp.exe firstversion.dat secondversion.dat
Summarizing file firstversion.dat: 100%
Summarizing file secondversion.dat: 100%
Results of comparing firstversion.dat and secondversion.dat 96%
match: 473 of 489 blocks match 96%
saving: 65339 bytes of BLD data required compared to file size of 1885184
```

The first line above is the command line for two uncompressed files. The next two lines show progress as the utility performs the calculations for byte-level differentiation. Finally the results are presented. In this instance, RayManageSofti chose to divide the files into 489 blocks (this number is variable across files), 96% of which were unchanged. Transferring the changes and the summary file required 65kb. This is less than 4% of the size of the complete second version file, which was over 1.8Mb.

Testing through Distribution

The key ingredients of a more complete test framework include the following:

- Use a large enough file, and one which is not of a type for which byte-level differentiation is prohibited. For

example, a 60k text file might be convenient, since small changes in two versions of the file are conveniently compared by eye (as well as with various tools).

For testing, it may be more convenient to temporarily disable compression when packaging test files. This makes it a little easier to 'read' what is happening to the files.

- Package the initial version of the test files in a RayManageSofti console. Distribute this version of the package to your test distribution location and download and install it on the test managed device.

- Make appropriately small changes in the original file, and repackage it as a new version of the original package. Distribute this new version to the test distribution location, but do not automatically update the managed device (for example, do not modify policy for the test managed device).

- Set up your monitoring tools to trace the network traffic to the test managed device. (Alternatively, you could rely on the log file as described below.)

- Use the command-line installation tool to download and install the new version of the package. Remember that byte-level differentiation is generally configured not to operate across high-speed LAN connections, so you may need to vary the network speed parameter to enable byte-level differentiation while keeping testing times short.

  For instance, you could use an option like this:

  ```
  ndlaunch
  -o NetworkMaxByteLevelSpeed=12500000
  -o LogFile=myTest.log myTestFile.osd
  ```

  This speed setting is high enough to allow byte-level differentiation on a 100Mbit/s LAN connection (recall that the option is specified in bytes).

- After the download, examine `myTest.log` for information about how much data was downloaded.

- Your network analysis should first show the download of the summary (`.bis`) file for the test file, followed by a relatively small number of bytes to download changes. The summary file should be less than 5% of the size of the test file. The smallest download of the revised test file is likely to be in the order of 512 bytes, and the mean block size in practical tests is likely to be about 4k bytes.

Troubleshooting Byte-level Differentiation

If, during testing, you do not get the results you expect from byte-level differentiation, this section of the chapter summarizes the settings to check.

1. Package uses byte-level differentiation

Check the implementation archive (`.ndc` file) of the package for the `updatehref` and `updatesize` attributes. If these settings are missing, check the following:

- Is the project packaged in Deployment Manager 6.0 or later? If not, repackage in a RayManageSofti console. See *Product version*.

- Is the file of a type for which byte-level differentiation is prohibited? If so, use a different file type for testing. See *File type and file size*.

- Is the file too small for byte-level differentiation? If so, use a larger file for testing. See *File type and file size*.

- Within the Software node of the RayManageSofti console, are the properties of this file set to allow byte-level differentiation? If not, set for automatic use. See *File type and file size*.

2. **Summary file is in place**

   Check the source directory from which the test file is being fetched, and ensure that the `.bls` summary file is present. If not (and test 1 above is valid), check distribution reports on the administration server for possible distribution problems. Also ensure that the attributes for summary file location (`updatehref`) and size (`updatesize`) are correctly populated in the package archive (`.ndc` file). For more information, see *Summary file*.

3. **Previous version of file available locally**

   Ensure that the previous version of the test file is in place on the target computer. See *File update*.

4. **Pull distribution and byte-level differentiation are configured for this leg**

   Ensure that pull distribution is configured for the test download (see *Pull distribution leg*).

   - If the test computer is a distribution server, check its attributes in the `hierarchy.cfg` file, and ensure that its `ByteLevel` attribute is set to `True`, and that the `ByteLevelThreshold` is set to zero (for more information, see *Distribution leg configured*).

   - For a managed device, check the preferences in the registry, and ensure that `AllowByteLevel` is `True`, and `NetworkMaxByteLevelSpeed` is set to a value higher than the achieved network speeds for this download leg.

   - For a non-Windows managed device (one running UNIX, Linux, or OS X), check the content of the `managesoft.xconf` file. In the managesoft-configuration element, ensure that both the cache-files and byte-level-update-enabled attributes are set to `True`.

     ```
     <managesoft-configuration ...
     cache-files="true"
     byte-level-update-enabled="true"
     ...>
     ```

     For more information, see the chapter on *Preferences on UNIX and Macintosh managed devices* in the *RMS Preferences for Managed Devices* manual.

For more information and assistance, contact your Raynet consultant.

Summary

Building on the overview of byte-level differentiation provided in the *RMS Software Deployment*, this chapter has:

- Examined all the factors that determine when and where byte-level differentiation will operate
- Provided guidelines for testing byte-level differentiation
- Summarized the key items to check when trouble-shooting byte-level differentiation.

Configuring Inventory

In this chapter, you will:

- Discover the different ways in which RayManageSofti can collect inventory data
- Read about the different sets of inventory information that can be collected on managed devices
- Find references to relevant other parts of the documentation set

About Inventory Data

RayManageSofti can obtain inventory data from two sources:

- Pre-implementation inventory uses a range of techniques to gather a baseline hardware and software inventory of enterprise computers before RayManageSofti is implemented. In this case there are no RayManageSofti agents installed on client computers. These techniques are described in the *RMS Discovery*.

- Post-implementation inventory relies on the activities of RayManageSofti agents installed throughout the system. The processes of collecting and collating post-implementation inventory data are described in detail in the *Inventory system flow* chapter of *RMS System Reference*.

Reporting on Inventory Data

Both pre- and post-implementation data is available to the Deployment Manager inventory reports, described in the *Reporting* chapter of the *RMS Software Deployment*.

Types of Inventory Data

On Windows, the inventory agent can collect different sets of inventory information:

- **Full inventory** is exactly what the name suggests: a complete listing of all applications (or other software, depending on agent configuration) found on the managed device at the time of the audit. This kind of inventory has no prerequisite.

- **Differential inventory** transmits a smaller set of information: The differences between the list of installed applications at audit time and the list generated at the last full inventory. It's important to be clear that each differential inventory reports differences from the last full inventory.

- **Incremental inventory** reports (potentially) a still smaller set: The differences between the list of applications at audit time and the list generated by the last inventory, of whatever kind.

Configuring the Inventory Agent

On managed devices, the inventory agent is responsible for gathering data about installed hardware and software.

Full, Differential, or Incremental Inventories

RayManageSofti can be configured to conduct either differential or incremental inventories, and to automatically insert a full inventory on every `n`th trigger. For details, see *Options to control differential inventory* in *RMS Preferences for Managed Devices*.

On Macintosh and UNIX managed devices, the inventory agent can collect only one set of inventory information, as specified in your `inventory.xconf` file. Refer to *RMS Preferences for Managed Devices* for details about configuring the inventory agent on Macintosh and UNIX managed devices.

Hardware and Software Classes

You can configure the hardware and software classes tracked by the inventory agent on managed devices. For example, you can configure it to perform complete searches of the disk drives on the managed device and to report classes of files, or even every file found. See the *RMS Preferences for Managed Devices* volume for details about configuring the inventory agent.

Configuring Inventory Data Import

Inventory information received from managed devices is processed into the RayManageSofti database by either:

- An ASP.NET web application, which imports data using IIS and following the process described at *Importing data using IIS*
- A scheduled task, **Import Deployment Manager inventories**, described in the *Scheduled tasks* chapter of *RMS System Reference*. It runs the Deployment Manager data importer, `mgsimport.exe`, which is described in the *Command line tools* chapter of *RMS System Reference*

Name Matching During Installation Data Import

During the import of inventory records, the Deployment Manager data importer performs name matching on installation evidence from both Deployment Manager and Windows Installer, so that where there is installation evidence both from the Deployment Manager and Windows Installer caches, only one installation evidence record is created. Installation evidence is matched against application names.

Listening or Polling for Distribution Jobs

In this chapter, you will:

- Become familiar with the processes of polling and listening for distribution jobs
- Identify the settings on each distribution server that control these processes
- Find out how distribution server settings are set in the distribution hierarchy file
- Learn how to change the way distribution jobs are processed
- Understand configuration requirements on parent distribution servers to allow polling to occur

How Distribution Jobs are Processed

Once distribution has been initiated, distribution jobs are processed periodically by the **Process Deployment Manager distribution jobs** scheduled task.

For distribution jobs that are intended for another distribution server or distribution location further down the distribution hierarchy, a distribution server uses information in the distribution hierarchy file (`hierarchy.cfg`) to determine which is the next distribution server in the hierarchy, and whether that distribution server is configured to listen for jobs or poll for jobs.

If the Child Distribution Server is Configured for Listening

The originating distribution server contacts the next distribution server in the distribution hierarchy.

The child distribution server then passes a request to the parent distribution agent, and the two distributions are connected. The distribution job is then distributed to the child distribution server.

If the Child Distribution Server is Configured for Polling

The originating distribution server writes the job in an encoded form into a specific job queue for the child distribution server.

At predetermined times, the child distribution server polls its parent distribution server to determine whether there are any jobs waiting to be distributed to the child distribution server.

What about Push / Pull Distribution?

Regardless of whether a distribution server is configured for polling or listening, data is always transferred between distribution servers in the following way:

- Files are **pulled** by active distribution servers from their parent distribution servers
- Files are **pushed** by the administration server or active distribution servers to distribution locations on passive distribution servers

Choosing to Poll or Listen

For Windows distribution servers, you may choose whether each distribution server will poll or listen for jobs. Listening is often the most appropriate option. For this reason, listening is the default action for distribution servers.

However, polling is appropriate in the following circumstances:

- There is a firewall between the distribution server and its parent that only allows connections to be initiated by the child distribution server.
- There is a reason to control or limit network connections. For example, some enterprises may incur a cost each time a connection is made. In this circumstance, it is possible to configure polling to occur at predetermined intervals. (When a distribution server is configured for listening, connections cannot be controlled. They occur every time a job needs to be passed down from a parent server.)

Network traffic is minimal for both polling and listening. However, polling does generate slightly more traffic, as polling occurs at predetermined times, regardless of whether there are any jobs to be processed. When listening, traffic is only generated when a job is available to be received.

Configuring a Distribution Server for Polling

Each distribution server's initial settings for processing distribution jobs are set during installation of **Deployment Manager for distribution servers** on the server.

These settings can be recorded:

- In the `mgssetup.ini` file used for bootstrapping the distribution server
- In the Install wizard that can be run from the **RayManageSofti Deployment Manager** product DVD (or the installation resources provided for you)

Mgssetup.ini Settings

The options that you can set in the `mgssetup.ini` file to determine how a distribution server handles distribution jobs are:

- `DSJOBQLOCATION`
  Name of the parent distribution server used to poll for jobs.
- `DSJOBQPORT`
  The port used to poll for jobs.
- `DSJOBQPROTOCOL`
  Polling protocol. Set to `http` or `https`.
- `DSJOBQREMOTE`
  Method of processing distribution job (set to `1` for polling, or `0` for listening).
- `DSLISTENPORT`

The port used to listen for jobs.

Install Wizard Settings

If you are installing a distribution server using the Install wizard, set the method of polling or listening for jobs on the **Job Retrieval Method** page.

See the *RMS Implementation* for details.

Changing Distribution Job Configuration

Configuration settings that determine whether a distribution server listens or polls for distribution jobs are stored in two places:

- Registry settings on each distribution server
- The distribution hierarchy file, which is defined on the RayManageSofti administration server and distributed to all distribution servers as part of the standard distribution process

To alter distribution job processing for a distribution server, the registry settings on the distribution server and the hierarchy file on the administration server must both be changed.

Changing Registry Settings

The registry settings that control how distribution jobs are processed on each distribution server are described in the following table:

To change...	Update this registry key...
Whether this server listens or polls the parent distribution server for distribution jobs	`[Registry]\ManageSoft\Replication Agent\CurrentVersion\JobQueueRemote` To enable this distribution server's connection agent to listen for connection attempts from the parent distribution server, set this key to `False`. To enable this distribution server's distribution agent to periodically poll the parent distribution server for distribution jobs, set this key to `True`. This option is only supported if a web server is installed on the parent distribution server.
The listening port	`[Registry]\ManageSoft\Listening Agent\CurrentVersion\Port` By default, this is set to `7010`. Change the port number as required.
The polling port	`[Registry]\ManageSoft\Replication Agent\CurrentVersion\JobQueuePort` By default, this is set to `80` for HTTP protocol, or `443` for HTTPS protocol. Change the port number as required.
The polling protocol	`[Registry]\ManageSoft\Replication Agent\CurrentVersion\`

To change...	Update this registry key...
	`JobQueueProtocol` By default, this is set to `http`. You can change the protocol to `https` if required.

The parent distribution server's job queue location (for polling)	`[Registry]\ManageSoft\Replication Agent\CurrentVersion\JobQueueLocation`
	
	During installation, this key is calculated by the protocol, server name, and port number to:
	`[Protocol]://[Server Name]:[Port]/ManageSoftJQ/mgsjobsrv.exe`
	
	You should not have to change this value manually, as it is updated whenever the distribution hierarchy is verified on the administration server. The server name is set according to the server name defined in the distribution hierarchy. It can be changed using the Distribution node on the RayManageSofti console.
	
	You may need to configure port details using IIS (Internet Information Services Manager) on the parent distribution server to ensure that the location is valid.

Changing the Distribution Hierarchy File

The distribution hierarchy file (`hierarchy.cfg`) is created on the administration server, and is distributed to all distribution servers. The *Distribution system flow* chapter in the *RMS System Reference* describes how this distribution occurs.

You can change an individual distribution server's port details on the RayManageSofti console. To do this, update the **Distribution Jobs** tab in the **Distribution Server properties**. The changes you make are saved to the distribution hierarchy file (`hierarchy.cfg`). See the *Distribution system* chapter of the *RMS Software Deployment* for details.

To change the port number for multiple distribution servers at once, you may find it more convenient to change the `hierarchy.cfg` file directly. To do this:

1. On the RayManageSofti administration server, open your preferred text editor.
2. Open `hierarchy.cfg`.

   This is typically located in: `C:\ManageSoft\Repository\DeploymentLocations\Common\`
3. Find the details you wish to change and replace them with the details (port number, for example) of your choice.
4. Save and close the file.

If a Distribution Server Cannot Identify Its Parent...

In some circumstances active distribution servers that are configured to poll for jobs are not able to identify their parent servers in the distribution server hierarchy. This occurs if:

- The distribution server's name in the RayManageSofti distribution hierarchy is different than the fully qualified DNS hostname that the distribution server determines for itself

- The distribution server's parent in the RayManageSofti distribution hierarchy has changed

If any of these conditions is `true`, your distributions will fail. You must update the distribution server's details so that it knows its parent in the RayManageSofti distribution hierarchy. To do so:

1. On the distribution server that is failing to receive jobs, open the **replication agent log file** (located by default in `$(TempDirectory)\ManageSoft\replag.log`) for viewing.

   The log file should show an attempt every minute to poll the distribution server's parent for jobs.

2. Search for the message *Could not find DS UID for hostname <host>*.

   If the message is not present in the log file, proceed to step 3.

   If this message is present in the log file, it indicates that the specified hostname can't be found in the RayManageSofti distribution hierarchy, and you should complete the following steps:

   - Change the `hostname` in the distribution server's properties on the RayManageSofti console to match the name in the log message (see the *Distribution system* chapter in the *RMS Software Deployment* for details about changing the `hostname`).

   - Wait to see if this fixes the problem. If it does, you do not need to perform any more steps in this process.

3. Search for the most recent log file message containing the text *Polling http://<parent>/ManageSoftJQ/mgsjobsrv.exe?<args>*.

   If the `<parent>` hostname specified in this message is not the current parent of the distribution server, this suggests that the parent of this distribution server has changed.

   - On the distribution server that is failing to receive jobs, use `regedit` or `regedt32` to set the value of `[Registry]\ManageSoft\Replication Agent\CurrentVersion\JobQueueLocation` to the hostname of the correct parent.

   - Wait to see if this fixes the problem. If it does, you do not need to perform any more steps in this process.

4. Copy the file that contains details about the current distribution hierarchy (`C:\ManageSoft\Repository\DeploymentLocations\Common\hierarchy.cfg`) from your administration server to the same location on the distribution servers.

   Wait to see if this fixes the problem. If it does, you do not need to perform any more steps in this process

5. On the distribution server that is failing to receive jobs, set the value of `[Registry]\ManageSoft\Replication Agent\CurrentVersion\LocalAddress` by completing the following steps:

   - On the administration server, open `C:\ManageSoft\Repository\DeploymentLocations\Common\hierarchy.cfg` in the text editor of your choice.

   - Locate the `<Replicator>` element for the distribution server in the hierarchy. It will look something like:
     ```
     <Replicator
     SERVER="{F4D153B2-C6A0-4C6D-AC3D-EE5E10CB7CE5}"
     HOST="florida-ds.tmnis.com" PORT="none">
     ```

   - On the distribution server that is failing to receive jobs, use `regedit` or `regedt32` to set the value of `[Registry]\ManageSoft\Replication Agent\CurrentVersion\LocalAddress` to the value of the `SERVER` attribute you identified at the previous step. Include the curly braces `{}`.

This `LocalAddress` value is then used by the distribution server to identify itself in the distribution hierarchy.

Configuring a Parent Distribution Server to Allow Polling

The job server program on a polling distribution server accesses the **ManageSoftJQ** share on the parent distribution server to check for distribution jobs.

This share is automatically configured on distribution servers.

Configuring Reporting

In this chapter, you will read how to configure the appearance and behavior of reports. For instructions about generating and viewing reports, refer to the *Reporting* section of the *Deployment Manager application help*, or to the *RMS Software Deployment*.

There are a number of ways that you can customize reporting in RayManageSofti. See the following sections:

- *To register reports with Reporting Services*
- *To change the default graph refresh rate*
- *To alter report pagination*
- *To change the default date range for application usage monitoring*
- *To change the default graph size for application usage*
- *To change the default color for application usage chart columns*
- *To add custom report pages*
- *To change a report's definition*
- *To build your own reports*
- *Managing the asset reports catalog*

To Register Reports with Reporting Services

Reports can be exported to a variety of formats, including PDF (`.pdf`, suitable for printing), Microsoft Excel (`.xls`), a text file of comma-separated values (`.csv`), and XML (`.xml`). This functionality relies on reports being registered with Microsoft Reporting Services, which must be running on the server on which the RayManageSofti reports server is installed.

If Microsoft Reporting Services is installed at the time you install or upgrade to Deployment Manager, you can automatically register all reports with Reporting Services, so that you can print and export data from them. For details, refer to the *RMS Implementation* and the *RMS Upgrade*.

If Microsoft Reporting Services is not installed at the time you install or upgrade to Deployment Manager, and you later decide to install it so that you can print and export data from reports, you must register all reports after installing Reporting Services. Consult the Microsoft Reporting Services documentation for installation instructions. Scripts are provided with your Deployment Manager installation to register all reports with Reporting Services. Instructions for using them are provided below.

If you add custom reports, you must register them with Reporting Services before you can print or export data from them:

- You can register **asset reports** through the asset reports catalog. See *Managing the asset reports catalog* for details.

- You can register **operational reports** using scripts provided with your Deployment Manager installation. Read on for details.

Registering Deployment Manager Reports with Reporting Services

To register Deployment Manager reports with Reporting Services:

1. Make sure the location of the `rs.exe` executable installed with Reporting Services is in your `PATH` environment variable. By default, `rs.exe` exists in `C:\Program Files\Microsoft SQL Server\80\Tools\Binn`.

2. Open a command prompt (**Start** > **Run** > **cmd**).

3. Change directory (`cd`) to `<RayManageSofti>\Reporter\Examples\PrintableReports`, where `<RayManageSofti>` is the location of your Deployment Manager installation (by default, `C:\Program Files\ManageSoft`).

4. Run this command, substituting the values appropriate for your environment as described below:

```
rs -i DeployManageSoftReports.rss -s localhost/ReportServer
-v domainName="<domain>"
-v installDir="<installDir>"
-v mgsDataServerName="<DataServer>"
-v mgsDataBaseName="<DatabaseName>"
-v mgsDataAuthType="<DataAuthenticationType>"
-v mgsCredRetrieval="”<CredentialType>”
-v mgsDataUserName="<DOMAIN\Username>"
-v mgsDataPassword="<password>"
-v mgsAdministratorGroup="MGS Administrators"
-v mgsReportUsers="MGS Report Users"
```

   where:

   - `<domain>` is the NT domain name to be used during authentication if WindowsNT is specified as the `mgsDataAuthType` value.
     For example `-v domainName="tmnis.org"`.

   - `<installDir>` is the root directory of your Deployment Manager installation.
     For example `-v installDir="C:\Program Files\ManageSoft\"`.

Be aware:
You must escape the trailing slash (`\`) with a second slash to prevent it from escaping the closing `"`. Without escaping the trailing slash, this command will not run from a DOS (command) prompt.

   - `<DataServer>` is the DNS name of the RayManageSofti data server.
     For example `-v mgsDataServerName="localhost"`.

   - `<DatabaseName>` is the name of the RayManageSofti database, typically *ManageSoft*.
     For example, `-v mgsDataBaseName="ManageSoft"`.

   - `<DataAuthenticationType>` is the type of authentication to use when connecting to the RayManageSofti database. Permitted values are `WindowsNT` and `SQLServer`.
     For example, `-v MgsDataAuthType="SQLServer"`.

   - `<CredentialType>` is the method used to retrieve credentials. Permitted values are `Integrated`, `Store`, `Prompt`, and `None`.

For example, `-v mgsCredRetrieval="Integrated"`.
For more information about managing credentials with Reporting Services, see the product documentation.

- `<domain\Username>` is the username (qualified with domain name, if required) to use when connecting to the RayManageSofti database.
  For example, `-v mgsDataUserName="tmnis.org\reports"`.

- `<password>` is the password for the username provided.
  For example, `-v mgsDataPassword="j7ndfldk9a"`.

- Mgs Administrators is the default security group for users with administration privileges, but you can specify a different security group.
  For example, `-v mgsAdministratorGroup="MySecurityGroup"`.

- MGS Report Users is the default security group for users with privileges for running Deployment Manager reports, but you can specify a different security group.
  For example, `-v mgsReportUsers="MySecurityGroup"`.

Be aware:
You must specify each parameter, even if you provide an empty value (`""`).

This registers all operational reports with Reporting Services.

5. To register asset reports, repeat the procedure, first changing directory to `<ManageSoft>\Reporter\Web\MGSDefns\En`, and also substituting `DeployAssetReports.rss` in place of `DeployManageSoftReports.rss`.

6. If you also use Security Manager and want to be able to print and export its report data, repeat the procedure again, this time changing directory to the location where `DeploySecMngrReports.rss` is installed, and substituting this script name with `DeployManageSoftReports.rss`.

7. Open your text editor of choice.

8. Browse to your Deployment Manager reporting installation, and open the configuration file `Web.config`.

   The default location is `C:\Program Files\ManageSoft\Reporter\Web\Web.config`.

9. Check for the following entries within the `<appSettings>` section of this file.

Entry	Description

add
`key="ManageSoft.Reports.MSReportingServer"`
`value="http://<machineName>/ReportServer"/>`
	This entry enables export features in reporting. Make sure that there is only one occurrence of this key in your `Web.config` file. Multiple entries will prevent reports from being exported.

add
`key="ManageSoft.Reports.AdminGroup"`
`value="MGS Administrators"/>`
	This entry is used to configure the name of the security group that contains users who are able to modify the list of my assets reports.
	(This task is available from the **Reports** node on the RayManageSofti console. Select **my assets**, and then select **Manage reports catalog**.)
	Normally, this entry should be set to **MGS**

Entry	Description
	Administrators.

add
`key="ManageSoft.Reports.SecurityDomain"`
`value="DEVOCEAN"/>`
	This entry is used to configure the domain of the security group identified by the `ManageSoft.Reports.AdminGroup` key described above.
	Normally, this entry should be set to identify the domain in which RayManageSofti for administration servers is installed. Enter the domain name in NT format.

If reports were not registered when you installed Deployment Manager, these entries may not be present.

Search to find the entries. If you find the entries, edit them to suit your environment.

If these entries do not exist, add them to the `<appSettings>` section.

10. Save and close the `web.config` file.

11. To verify that your configuration is complete, export a Deployment Manager report to each of the available formats: `.pdf`, `.xls`, `.csv`, and `.xml`.

Removing Report Registrations

Scripts are also provided for you to run if you ever need to remove report registrations:

- `RemoveManageSoftReports.rss` removes both operational and asset report registrations
- `RemoveSecMngrReports.rss` removes Security Manager report registrations

In each case, it is best to change directory (`cd`) to the location of the `.rss` script before executing it.

To Change the Default Graph Refresh Rate

The default graph rate is used by all graphs in the top level reports pages (**my organization**, **my assets**, and any custom pages you have added) and detailed report pages, and is also the default for graphs in the Deployment Manager dashboard. For individual graphs in the dashboard, you can over-ride this default setting (see the *Reporting* chapter of the *RMS Software Deployment*).

To change the default graph rate:

1. Close any RayManageSofti console or web browser using Deployment Manager reporting.

2. Open your text editor of choice.

3. Browse to your Deployment Manager reporting installation, and open the configuration file `Web.config`.

   The default location is `C:\Program Files\ManageSoft\Reporter\Web\Web.config`.

4. Look for the following lines:

```
<!-- PUBLIC : Default value of the refresh rate used for graphs.
Value is in seconds -->
<add key="ManageSoft.Reports.RefreshRate" value="60" />
```

The initial default value is `60` (seconds).

5. Change the content of value to the desired refresh interval, in seconds.
   Example: `<add key="ManageSoft.Reports.RefreshRate" value="120" />`

Be aware:
All graphs use your setting as the refresh rate (except dashboard graphs for which you have set a special refresh rate). Make sure that you choose a refresh rate that both keeps your information current, and can be supported by your server and bandwidth availability.

6. Save and close the file.

7. Restart Deployment Manager.

   Any items added to your dashboard from this point will use the new refresh rate by default. Pre-existing dashboard items will continue to use their previous settings.

To Alter Report Pagination

To alter the pagination of reports for on-screen and print display:

1. Close any RayManageSofti console or web browser using Deployment Manager reporting.

2. Open your text editor of choice.

3. Browse to your Deployment Manager reporting installation, and open the configuration file `Web.config`.

   The default location is `C:\Program Files\ManageSoft\Reporter\Web\Web.config`.

4. In the section `<!-- REPORTS CONFIG -->`, look for the following lines:

```
<!-- PUBLIC : Default value for enabling pagination -->
<add key="ManageSoft.Reports.Pagination_Default" value="yes" />
<!-- PUBLIC : Default value for default page size if pagination enabled -->
<add key="ManageSoft.Reports.PageSize_Default" value="20" />
```

5. To turn off report pagination (so that all reports present as continuous long lists), change the first of these keys to have the value "`no`":
   `<add key="ManageSoft.Reports.Pagination_Default" value="no" />`

Be aware:
If you have large amounts of data, leaving pagination turned on is strongly recommended.

6. To change the number of items per page of your reports, ensure that the first of these keys retains its default value of "`yes`", and edit the second to have the required numerical value.

   For example, to increase to 30 items per report page:

```
<add key="ManageSoft.Reports.PageSize_Default" value="30" />
```

Determining the correct number for your environment is a compromise between the amount of scrolling per page and the amount of switching between pages.

Note:
Remember that you can use filtering to limit the length of reports.

To Change the Default Date Range for Application Usage Monitoring

The default date range for filtering application usage reports is 30 days. To change this range, do the following:

1. Close any RayManageSofti console or web browser using Deployment Manager reporting.
2. Open your text editor of choice.
3. Browse to your Deployment Manager reporting installation, and open the configuration file `Web.config`.

   The default location is: `C:\Program Files\ManageSoft\Reporter\Web\Web.config`
4. Look for the following line:
   ```
   <add key="ManageSoft.Reports.AppUsageReportDefaultDateRange" value="30"/>
   ```
5. Change the number `30` to the number of days you want the application usage reports to use as a default date range.

Be aware:
It is recommended that you copy the commented examples provided in the `Web.config` file, paste them outside the commented text and edit them to suit your needs.

6. Save and close.

To Change the Default Graph Size for Application Usage

The View graph for application usage has a default size of 740x440. This size fits in the standard layout of the reports when viewed through the RayManageSofti console at a screen resolution of 1024x768 pixels.

To change this size, do the following:

1. Close any RayManageSofti console or web browser using Deployment Manager reporting.
2. Open your text editor of choice.
3. Browse to your Deployment Manager reporting installation, and open the configuration file `Web.config`.

   The default location is: `C:\Program Files\ManageSoft\Reporter\Web\Web.config`

4. Look for the following line:

```
<add key="ManageSoft.Reports.AppUsageGraphHeight" value="440"/
<add key="ManageSoft.Reports.AppUsageGraphWidth" value="740"/>
```

5. Change the height and width to the desired number of pixels.

Be aware:
It is recommended that you copy the commented examples provided in the `Web.config` file, paste them outside the commented text and edit them to suit your needs.

6. Save and close.

To Change the Default Color for Application Usage Chart Columns

The columns on an application usage chart are light green by default. To change this color, do the following:

1. Close any RayManageSofti console or web browser using Deployment Manager reporting.
2. Open your text editor of choice.
3. Browse to your Deployment Manager reporting installation, and open the configuration file `Web.config`.

   The default location is: `C:\Program Files\ManageSoft\Reporter\Web\Web.config`
4. Look for the following line:

```
<add key="ManageSoft.Reports.SingleChartColor" value="#00CC33" />
```

5. Change the color to your desired color.

   The color is specified in RGB hexadecimal format. For example "`#FF0000`" is pure red.
6. Save and close.

To Configure Operating System Legends and Labels for the OS Summary

When you first open the **Reports** node in the RayManageSofti administration console, the **my organization** page is displayed. At the right-hand side, it shows a graphical summary of operating systems. You can configure the legends and labels for the operating systems that appear on this graph through a `Web.config` file.

To modify operating system legends and labels:

1. In a text editor of your choice, open `C:\Program Files\ManageSoft\Reporter\Web\Web.config` for editing.

2. Locate the strings
   `<add key="ManageSoft.Reports.OSFullNames" …>`
   and
   `<add key="ManageSoft.Reports.OSShortNames" …>`

3. Add, change, or delete operating system details as required.

Be aware:
Each operating system must have entries in both the `OSFullNames` and in `OSShortNames` value strings, and the entries must occur in the same location in each value string, or legends and labels will be mismatched.

When adding an operating system:

- At the end of the `ManageSoft.Reports.OSFullNames` value string, add the new OS name (for example Windows Server Longhorn), followed by a semicolon (`;`).
- At the end of the `ManageSoft.Reports.OSShortNames` value string, add the abbreviation for the new OS name (for example WSL), followed by a semicolon (`;`).

4. Save and close the file.

To Change a Report's Definition

Source files containing the definitions of all Deployment Manager reports are installed with the product.

Under the Deployment Manager reporting installation location (by default, `C:\Program Files\ManageSoft\Reporter`):

- Operational reports definitions are stored under the Web folder
- Asset reports definitions are stored under `Web\MGSDefns`. Under this folder are folders named for the native language of the reports. For example, the `EN` folder contains the English language report definitions.

You can modify these report definitions to customize the reports they produce. Report definitions are editable in a text editor, or using Microsoft Visual Studio .NET 2003 (you must have the Report Designer from Microsoft SQL Server Reporting Services installed).

To Build Your Own Reports

You can create custom reports in Report Definition Language (RDL) format, and add them to your asset reports catalog (see *To add a report*).

RDL is an XML-based industry standard used to define reports. You can create custom reports in any application that can generate RDL files.

Microsoft SQL Server Reporting Services includes Report Designer, a report authoring tool hosted in the Microsoft Visual Studio Environment.

Managing the Asset Reports Catalog

Deployment Manager provides a standard set of categories and asset reports, to which you can add your own categories and custom reports as required.

To add custom reports to locations other than the my assets page, refer to *To add custom report pages* and *Adding custom reports to a report page*. For details about creating your own custom reports, see *To build your own reports*.

To Add an Asset Reports Category

To add an asset reports category:

1. Make sure you are positioned on the **my assets** page:

    a. On the RayManageSofti console, select the **Reports** node.
    b. The **my organization** page is displayed.
    c. Click **my assets**.

2. Click the **Manage reports catalog** button. 

    The Manage Asset Reports page is displayed.

3. Select the asset reports category below which you want your new category to appear.

4. Click the **Add a category beneath the item selected below** button.

    The **Add Category** page is displayed.

5. In the **Category Name** field, enter the name of your new category.

6. Click the **Add the category to the report catalog** button.

The **Manage Asset Reports screen** is re-displayed, with the new category. The plus (+) sign on the new folder icon shows that this is an added category, not one of the standard set supplied with Deployment Manager.

To Remove an Asset Reports Category

You can remove report categories that you have added to the asset reports catalog. You cannot remove the standard categories provided by Deployment Manager.

To remove an asset reports category:

1. Make sure you are positioned on the **my assets** page:

    a. On the RayManageSofti console, select the **Reports** node.
    b. The **my organization** page is displayed.
    c. Click **my assets**.

2. Click the **Manage reports catalog** button. 

The **Manage Asset Reports** page is displayed.

3. Select the asset reports category you want to delete.
4. Click the **Remove the category selected below** button.
5. On the resulting dialog, click **OK** to confirm that you want to delete the category.

The category is deleted, and the **Manage Asset Reports** screen is re-displayed.

To Add a Report

If you have created custom reports, you can add them to your Deployment Manager asset reports catalog. You can add:

- Web-based reports that use technologies including HTML, ASP, ASP.NET and JavaScript.
- Reports in Report Definition Language (RDL) format. (RDL is an XML-based industry standard used to define reports. See *To build your own reports* for more details.)
- Crystal reports

Be aware:
To add Crystal reports, the Crystal embedded DLLs must be installed on the RayManageSofti reports server.

(You can also add custom web-based reports to the my organization page, or custom pages. See *Adding custom reports to a report page* for details.)

To add a report to the catalog:

1. Make sure you are positioned on the **my assets** page:
   a. On the RayManageSofti console, select the **Reports** node.
   b. The **my organization** page is displayed.
   c. Click **my assets**.
2. Click the **Manage reports catalog** button.

   The **Manage Asset Reports** page is displayed.
3. Select the asset reports category or report below which you want the new report to appear.
4. Click the **Add a new report beneath the item selected below** button.

   The **Add Report** page is displayed.
5. From the **Report Type** list, select the type of report you are adding to the asset reports catalog.

   The fields on the screen may be updated according to the type of report you select.

6. Complete the remaining fields on the screen:
   - If the **Filename** field appears, enter (or browse for) the location of the report to be added to the catalog.
   - If the **Title** field appears, enter a description of the report being added to the catalog.
   - If the **URL** field appears, enter the web address from which the report is available.
7. Click the **Add report** button.

   The report is added, and the **Add Report** screen is refreshed.
8. Continue to add reports, or click **Back** to return to the **Manage Asset Reports** page.

To Remove an Asset Report

You can remove reports that you have added to the asset reports catalog. You cannot remove the standard reports provided by Deployment Manager.

To remove an asset reports category:

1. Make sure you are positioned on the **my assets** page:
   a. On the RayManageSofti console, select the **Reports** node.
   b. The **my organization** page is displayed.
   c. Click **my assets**.
2. Click the **Manage reports catalog** button. 

   The **Manage Asset Reports** page is displayed.
3. Select the asset report you want to delete.
4. Click the **Remove the report selected below** button.
5. On the resulting dialog, click **OK** to confirm that you want to delete the report.

The report is deleted, and the **Manage Asset Reports** screen is re-displayed.

To Register an Asset Report with Windows Reporting Services

If you want to be able to print and/or export a report to supported formats (`.xls`, `.csv`, and `.pdf`), the report must be registered with Windows Reporting Services. All asset reports supplied with Deployment Manager are registered with Windows Reporting Services, so typically you will require this option only when you have added a new report.

To register a report with Windows Reporting Services:

1. Make sure you are positioned on the **my assets** page:
   a. On the RayManageSofti console, select the **Reports** node.

b. The **my organization** page is displayed.
c. Click **my assets**.

2. Click the **Manage reports catalog** button. 

   The **Manage Asset Reports** page is displayed.

3. Select the report you want to register with Windows Reporting Services.

4. Click the **Register the selected report with Windows Reporting Services** button.

The report is registered with Windows Reporting Services, and the **Windows Reporting Service registration** page is displayed with the results of the registration.

To Register All Asset Reports with Windows Reporting Services

If you want to be able to print and/or export reports to supported formats (`.xls`, `.csv`, and `.pdf`), each report must be registered with Windows Reporting Services.

All asset reports supplied with Deployment Manager are registered with Windows Reporting Services, so typically you will require this option only when you have added a number of reports.

To register all reports with Windows Reporting Services:

1. Make sure you are positioned on the **my assets** page:

   a. On the RayManageSofti console, select the **Reports** node.
   b. The **my organization** page is displayed.
   c. Click **my assets**.

2. Click the **Manage reports catalog** button. 

   The **Manage Asset Reports** page is displayed.

3. Click the **Register all reports with Windows Reporting Services** button.

The reports are registered with Windows Reporting Services, and the **Windows Reporting Service registration** page is displayed with the results of the registrations.

To Add Custom Report Pages

You can add one or more custom report pages at the same level as the **my organization** and **my assets** pages. To do this:

1. Close any RayManageSofti console or web browser using Deployment Manager reporting.

2. Open your text editor of choice.

3. Browse to your Deployment Manager reporting installation, and open the configuration file `Web.config`.

The default location is: `C:\Program Files\ManageSoft\Reporter\Web\Web.config`

4. Look for the following line:
   `<!-- CUSTOM TOP-LEVEL REPORTS MENU LIST -->`

Be aware:
It is recommended that you copy the commented examples provided in the `Web.config` file, paste them outside the commented text and edit them to suit your needs.

5. Locate the line
   `<add key="ManageSoft.Reports.TopMenuList" value="1;2"/>` and add a semi-colon (`;`) and page number to it. For example, `;3`.

6. Add an additional `<Add key />` section with the details of your new report page. It's easiest to copy and paste an existing section, and edit it. Example:

```
<add key="ManageSoft.Reports.Menu.3.DisplayLabel" value="TMNIS custom"/>
<!-- Optional URL params for organizationToday.aspx-->
<add key="ManageSoft.Reports.Menu.3.Href"
value="Main/organizationToday.aspx?browrep=EMSHierarchyBrowser
&amp;browdetrep=In formation"/>
<!-- Optional -->
<add key="ManageSoft.Reports.Menu.3.Target"
value="CenterFrame"/>
<!-- Optional -->
<add key="ManageSoft.Reports.Menu.3.Tooltip"
value="Click to go to TMNIS custom reports page"/>
<!-- Optional -->
<add key="ManageSoft.Reports.Menu.3.DefaultStartup" value="true"/>
<add key="ManageSoft.Reports.RefreshRate" value="0"/>
```

7. Save and close the file.

8. Restart Deployment Manager.

To add reports to this page, follow the instructions below.

Adding Custom Reports to a Report Page

You can add web-page based custom reports to your Deployment Manager reporting environment. Your custom report web pages may use `HTML`, `ASP`, `ASP.NET` and `JavaScript`. This section explains how to link those pages into Deployment Manager. (You can add reports to the my assets page through the reports catalog. Refer to *Managing the asset reports catalog* for details.)

1. Close any RayManageSofti console or web browser using Deployment Manager reporting.

2. Open your text editor of choice.

3. Browse to your Deployment Manager reporting installation, and open the configuration file `Web.config`.

   The default location is: `C:\Program Files\ManageSoft\Reporter\Web\Web.config`

4. Look for the following line:
   `<!-- CUSTOM REPORTS URLS -->`

The custom report URL list is a semicolon-separated (";") list used to store the IDs for any URLs which you want to display.

Be aware:
It is recommended that you copy the commented examples provided in the `Web.config` file, paste them outside the commented text and edit them to suit your needs.

5. Add as many reports as you require.

   The report URL lines consist of the following details:

   - The first line is a list of the IDs of the custom report URLs to be displayed.

     In this example, three reports are being defined:
     ```
     <add key="ManageSoft.Reports.CustomUrlList" value="Rep1;Rep2;Rep3" />
     ```

   - The next lines are the details of your reports.

   ### Mandatory Details

   - `ReportTitle`
   - `LinkLabel`
   - `HREF`

   Rep1, for example, could look like this:

   ```
   <add key="ManageSoft.Reports.CustomUrlDetails.Rep1.ReportTitle"
   value="Custom Title 01" />
   <add key="ManageSoft.Reports.CustomUrlDetails.Rep1.LinkLabel"
   value="Label 1" />
   <add key="ManageSoft.Reports.CustomUrlDetails.Rep1.Href"
   value="Report1_table.asp" />
   ```

   ### Optional Details

   - `Target`
     Indicates a target frame or page in which the URL is to be displayed

   - `PassParams`
     Indicates that options given to the Deployment Manager reports are to be added to the URL.

   #### Example

   ```
   <add key="ManageSoft.Reports.CustomUrlDetails.Rep2.ReportTitle"
   value="Custom Title 02" />
   <add key="ManageSoft.Reports.CustomUrlDetails.Rep2.LinkLabel"
   value="Label 2" />
   <add key="ManageSoft.Reports.CustomUrlDetails.Rep2.Href"
   value="http://externalserver.com/"/>
   <add key="ManageSoft.Reports.CustomUrlDetails.Rep2.Target"
   value="_new"     />
   <add key="ManageSoft.Reports.CustomUrlDetails.Rep3.ReportTitle"
   value="Custom Title 03" />
   <add key="ManageSoft.Reports.CustomUrlDetails.Rep3.LinkLabel"
   value="Label 3" />
   <add key="ManageSoft.Reports.CustomUrlDetails.Rep3.Href"
   ```

```
value="page3.asp"/>
<add key="ManageSoft.Reports.CustomUrlDetails.Rep3.PassParams"
value="yes" />
```

6. Save and close the file.

7. Restart Deployment Manager.

Using Remote Control Software with Deployment Manager

Deployment Manager can integrate with software that can remotely control computers. The ability to remotely control user computers is useful for support staff that can use it to demonstrate tasks, or to help analyze and debug unresolved and possibly ill-defined problems.

Deployment Manager provides the widely accepted and freely redistributable TightVNC as a RayManageSofti package to be deployed to Windows managed devices. Alternatively, Deployment Manager integrates with your existing remote control solution.

This chapter provides information about the installation, configuration, and use of remote control software with Deployment Manager.

After reading this chapter, you will:

- Understand how to configure Deployment Manager to work with remote control software
- Know how to deploy the TightVNC remote control software to Windows managed devices across your enterprise
- Be able to configure Deployment Manager to use remote control software other than TightVNC
- Understand some of the security issues involved in the use of remote control software

Components of the Deployment Manager Remote Control Solution

There are three components of the Deployment Manager remote control solution:

- Managed devices from which operators will initiate remote control of other end-user computers.
- Managed devices that need to be controlled remotely.
- The RayManageSofti reports server. Managed devices that initiate remote control of other computers can access the reports server to view Remote Control and Diagnostics reports about other managed devices and initiate remote control access from these reports. See *Using remote control software* for more details.

The following diagram illustrates the way in which Deployment Manager integrates with remote control software.

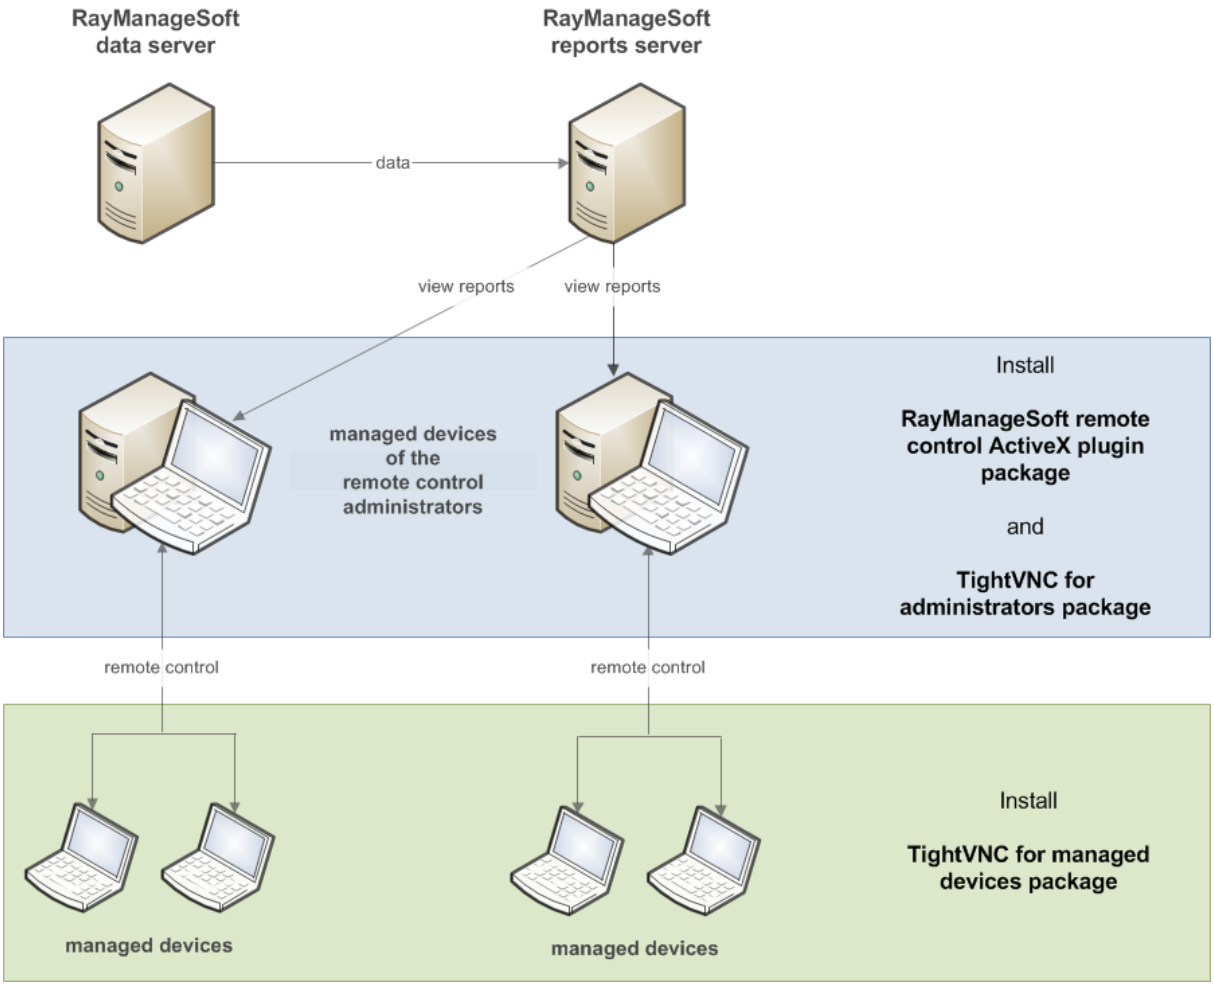

To enable Deployment Manager remote control, components must be installed on:

- Managed devices from which operators will initiate remote control of other end-user computers
- Managed devices that need to be controlled remotely

On Computers that will Initiate Remote Control

On every computer that will initiate remote control, you must install:

- Remote control software
- The Remote Control ActiveX object

Remote Control Software

This can be existing remote control software, or TightVNC.

While Deployment Manager supplies a default remote control package (TightVNC), it is not installed during a standard installation of the Deployment Manager product. This allows for enterprises which already have remote control software in place, or wish to use an alternate product.

To install TightVNC, you can use the RayManageSofti package, `TightVNC for Administrators.ndp` distributed with Deployment Manager.

The Remote Control ActiveX Object

The Remote Control ActiveX object enables operators to run the remote control software directly from the Deployment Manager report page. The report page searches for the ActiveX object on the controlling computer.

The ActiveX object is *not* automatically installed with Deployment Manager on the reports server, as this would allow every reports user to have access to remote control. Instead, the ActiveX object has to be installed as a package under a standard policy that dictates who should have remote control capability.

The package `Remote Control Plugin.ndp` is automatically available from your software library. To install the ActiveX object on managed devices, deploy this package using standard Deployment Manager software deployment routines.

When the ActiveX object runs the remote control software, it expects the following registry keys to be set on the operators' computer under `HKEY_CURRENT_USER or [Registry]\ManageSoft\RemoteControl\CurrentVersion`:

- `ConnectorCommand=<Full path to remote control executable>`
- `ConnectorOptions=$(dnsMachineID)`

These registry keys are set automatically when the `TightVNC for Administrators.ndp` package is installed (see *Remote control software* above). However, in some circumstances, you may need to manually set these keys. For more information, see *Setting registry keys on computers that initiate control*.

On Managed Devices that will Be Remotely Controlled

On every managed device that you want to control from a remote computer, you must install TightVNC (or already have other remote control software installed).

To do this on Windows managed devices, you can use a RayManageSofti package, `TightVNC for Managed Devices.ndp`.

Configuring Computers that will Initiate Remote Control

This section describes how to install Deployment Manager components on computers that will remotely control other managed devices. There are three steps in the installation:

- Install TightVNC for Administrators
- Install the Deployment Manager ActiveX object
- Set registry keys, if necessary

Installing TightVNC for Administrators

To enable computers to initiate remote control using TightVNC:

1. Receive the package `TightVNC for Administrators.ndp` into the RayManageSofti software library. (See the *Software* section of the *RMS Software Deployment* for details about receiving packages.)

   This package is used to install the TightVNC viewer. It will also set registry entries that allow the ActiveX object to initiate the TightVNC viewer.

2. Assign this package to policies targeting operators who are eligible to initiate remote control, such as System administrators, Help desk operators and Training Officers.

3. Distribute the package through the distribution hierarchy.

   After distribution through the hierarchy, the package will be available for managed devices to install.

Installing the Remote Control ActiveX Package

To install the ActiveX package on managed devices that will remotely control other managed devices:

1. Receive the package `Remote Control ActiveX Plugin.ndp` into the software library. (See the *Software* section of the *RMS Software Deployment* for details about receiving packages.)

2. Assign this package to policies targeting operators who are eligible to initiate remote control, such as System administrators, Help desk operators and Training Officers.

3. Distribute the package through the distribution hierarchy.

   After distribution through the hierarchy, the package will be available for managed devices to install.

Setting Registry Keys on Computers that Initiate Control

When the `TightVNC for Administrators.ndp` package is installed on an operator's computer, the **ConnectorCommand** and **ConnectorOptions** registry keys are automatically set under:

- `HKEY_CURRENT_USER\Software\ManageSoft Corp\ManageSoft\RemoteControl\CurrentVersion`
  or
- `[Registry]\ManageSoft\RemoteControl\CurrentVersion`

If these keys do not exist under `HKEY_CURRENT_USER`, then the equivalent key under `HKEY_LOCAL_MACHINE` is used.

ConnectorCommand

This key contains the full path to the remote control executable. For TightVNC, its value is `<Full path>\vncviewer.exe`.

ConnectorOptions

This key contains the values of any command line options applicable to the remote control executable.

For TightVNC, the command line is `vncviewer.exe <machine name>`, so **ConnectorOptions** should be set to a single argument specifying the hostname or IP address of the machine to which TightVNC will connect.

The following special values can be used in the value of **ConnectorOptions** to identify the computer to which a connection is being made:

- **$(dnsMachineID)**
  the fully qualified domain name of the computer
- **$(machineID)**
  the machine name (not qualified by any domain)
- **$(fullADMachineID)**
  the machine name, qualified with the computer's Active Directory domain name, in DNS format
- **$(ipAddress)**
  the IP address of the computer

Depending on your network setup, any of the above special values may be suitable. The default configured by the TightVNC for Administrators package (distributed with Deployment Manager) is **$(dnsMachineID)**.

Setting these Keys Manually

You must manually configure these registry keys if:

- You are using existing remote control software, or
- You installed TightVNC directly (not using the package distributed with Deployment Manager)

This can be done by adding the registry keys to the package being used to rollout the remote control software.

Setting these Keys for other Remote Control Software

For remote control software other than TightVNC, the registry keys will contain different values to the examples shown earlier. The values will depend on the command line used to initiate this remote control software.

For example: If another product used the command line `remotecontrolSoftware.exe -machine <machine name>` the registry keys could be set as follows:

ConnectorCommand: `<Full path>\remotecontrolSoftware.exe`
ConnectorOptions: `-machine $(dnsMachineID)`

Configuring Computers to be Remotely Controlled

This section describes how to install TightVNC on Windows managed devices that will be controlled by a remote computer.

Installing TightVNC on Windows Computers to be Remotely Controlled

To install TightVNC on Windows managed devices that can be remotely controlled by other managed devices:

1. Receive the package `TightVNC for Managed Devices.ndp` into the software library. (See the *Software* section of the *RMS Software Deployment* for details about receiving packages.)

2. If necessary, edit the package and change the default project password. See *Changing the password for TightVNC* for instructions.

   Why change the default password?

   The default project for TightVNC includes the password `managesoft`. This password is used by TightVNC to protect against unauthorized use. This password is configured by TightVNC in the registry in an encrypted format. You should check your security procedures to see if this default should be changed.

   If you want to have different passwords for different groups of operators, create a separate customized TightVNC for Managed Devices package for each group.

3. If necessary, change registry settings to determine whether end-users are given the opportunity to accept or reject requests to initiate a remote control session. For details about TightVNC registry settings, refer to TightVNC product documentation.

   More about configuring remote control behavior

   In the TightVNC for Managed Devices package, TightVNC is configured to present the end-user with a dialog that will allow them to accept or reject a request to initiate a remote control session. The request will automatically be accepted if there is no response from the user within 30 seconds. You can change this behavior by configuring the TightVNC registry setting that governs this behavior in the TightVNC for Managed Devices package before distributing it.

4. Assign this package to policies targeting managed devices or end-users whose devices are to be controlled.

After distribution through the hierarchy, the package will be available for managed devices to install.

Changing the Password for TightVNC

The default password is configured in TightVNC for Managed Devices, encrypted, and stored in the registry. To change the password:

1. Copy `C:\Program Files\TightVNC\WinVNC.exe` to a test computer.

Be aware:
If you are using a computer which is already running the WinVNC service, stop the service.

2. Run the executable.

3. Double-click the icon in your system tray.

   The **TightVNC properties** dialog is displayed.

4. Change the password.

Be aware:
You may want to change the password for the *view only* mode at the same time.

5. Close the dialog.

6. Select **Start** > **Run**...

7. Type **regedit** and click **OK**.

   The **Registry Editor** is opened.

8. Expand `HKEY_CURRENT_USER` to show `HKEY_CURRENT_USER\SOFTWARE\ORL\WinVNC3`.

9. Right-click WinVNC3 and select **Export** to export the registry key to a `.reg` file.

10. Open the `.reg` file in a text editor.

11. Locate the registry setting **Password**.

    Its value will be set to a hex value. For example, `hex:eb,33,52,ad,94,02,bc,76`.

12. Remove `hex:` and all commas from the string, leaving a string in the format `eb3352ad9402bc76`. Make a note of this string.

13. Save and close the file.

14. Using the Registry Editor, delete the registry key `HKEY_CURRENT_USER\SOFTWARE\ORL`.

15. Right-click the TightVNC icon in the system tray and select **Close VNC**.

Be aware:
If you stopped the WinVNC service earlier, restart it now.

16. Transfer the file containing the password value to your administration server.

17. From the **software library** within the **Software** node, open the **TightVNC for Managed Devices project** for editing.

18. Under the **resources** tab, open the **registry** section.

19. Expand `HKEY_LOCAL_MACHINE` until `HKEY_LOCAL_MACHINE\SOFTWARE\ORL\WinVNC3\Default` is visible.

20. Double-click **Password** to open the edit dialogue.

21. Replace the encrypted string within value with the one copied from TightVNC.

Be aware:
You will need to repeat these steps to change the view only password for TightVNC. The registry setting for the view only password is **PasswordViewOnly**.

22. Save, pack, and distribute the package to the relevant distribution locations.

Using Remote Control Software

Once the remote control packages have been installed you can start remote control sessions from the **Remote control and diagnostics** option on the reporting interface.

You can report by user to view the managed devices visited by each user, or report by managed device to view inventory data or initiate remote control.

By User

The search will report a list of users and the managed devices visited by each user. For user based searching, Deployment Manager must be configured to gather and resolve user based inventory.

From this report, you can:

- Click a user to view a list of packages, their installation policy, and any installation details
- Click **Software** to list software on each computer

By Managed Device

This report lists details of the last hardware and software inventory reported.

From this report, you can:

- Click a managed device to list packages, their installation policy and other installation details
- Click **Hardware** to list hardware found on a managed device
- Click **Software** to list software found on a managed device
- Select the Remote Control icon  to automatically start a remote control session

Security Overview

This brief chapter introduces security aspects of the RayManageSofti products. It is intended as a lead-in to any of the subsequent chapters on different elements of the security offering: rights and permissions controls, digital signature, and trusted locations.

Computer security is a diverse and sophisticated area of expertise. It includes at least five aspects of major concern, summarized with the acronym **RAPID**:

Rights

Rights management means setting privilege levels, access control lists, and other file system permissions so that the only people able to take specified actions are the ones duly authorized to do so. In RayManageSofti products, rights are usually assigned to security groups. These rights are then inherited by members of the groups. Information about rights and access for RayManageSofti is contained in the chapters on rights configuration.

Authentication

Authentication means validating the credentials of the user (or process) account. Identity is usually established by having each user provide a valid (and unique) account name and valid password before access is granted. Once authenticated, each account can be looked up in Access Control Lists to determine its authorized actions and access. Because there are often several accounts that need access to any given part of the RayManageSofti system, authorization is often handled through security group memberships. In summary, a user or process logs on to an account using an authenticated account name and password; then the group memberships of that account provide authorization for access.

Privacy

Privacy deals with whether your intellectual property is secure from unauthorized access or (in the extreme) from theft. In part, privacy is established by the careful control of rights, and for use inside a trusted corporation, this may be sufficient. However, when transmissions may travel in part over the Internet, you may need to protect against greater levels of malice, such as packet sniffing. In this area, RayManageSofti provides support for HTTPS as the industry-standard method of securing wide-area transmissions.

Integrity

Integrity concerns center around whether a transmission has been corrupted, either by network failures or by malicious intent. A related concern is whether a deliverable is truly what it purports to be (a package may be uncorrupted and technically valid and still contain a malicious payload). The combined area of package integrity is critical in RayManageSofti, since its essence is delivering packages to computers over a network. Therefore RayManageSofti provides an array of techniques to assure package integrity:

- MD5 cryptographic digests ensure that the software delivered to the managed device is a good copy of the original stored on the administration server. This guards against deliberate or accidental corruption during transmission. Where transmissions are also encrypted with HTTPS (such as across Internet links), this provides an additional check on package integrity.

- The optional use of digital signing validates that a package is unchanged since it was signed (that is, that it has not been attacked prior to transmission from the administration server). There's more information in the chapter *Digital signing*.

- Another option that provides additional assurance for the package catalog is the use of trusted locations. Managed devices may be configured to collect catalogs only from trusted locations, further reducing the risk

of malicious packages being inserted (provided that the trusted servers are well managed). Trusted locations are often combined with digital signing. There is more information in the chapter on *Trusted locations*.

Detection

Detecting when an incursion occurs is an operational matter. Good practice includes a range of measures, starting with email or cellphone alerts when incursion is detected, careful management and protection of log files, and so on.

Of the RAPID set, the most important in a RayManageSofti implementation are rights management (authorization), authentication, and integrity.

Because RayManageSofti is a system for delivering packages to computers over a network, there is a risk that it can be used maliciously to deliver hostile content.

Preventing such an attack from outside a company is primarily a function of corporate firewalls. However, some threats are internal, and it is important that system administrators and end-users can be assured that packages installed on their computers (particularly when those installations may be silent) are not malicious.

The Problem

Computer security is a complex issue, and it can be argued that no computer on a network can ever be made absolutely secure. RayManageSofti is not a security tool, and it cannot be used to prevent end-users (or others) from doing things that they are otherwise authorized to do.

The goal, however, is to prevent the use of RayManageSofti offering additional opportunities for threat and malice. In principle, malice is possible because RayManageSofti may require elevated privileges to perform a package installation. For example, when a non-administrator user is logged on while a machine policy calls for an all-user update, RayManageSofti elevates privileges for its installation agent to allow updates to proceed. Therefore, if a malicious person can arrange for a package to be installed with elevated privileges, there could be a significant security breach.

In summary, there need to be adequate safeguards (both technological and procedural) to prevent unprivileged users from gaining privileged access. Since the most likely form of attack is for a malicious package to be installed under elevated privileges, the critical element of defense is to be able to trust packages.

The Solution

There are several parts to a complete solution for trust of RayManageSofti packages. Fortunately, the solution is not onerous. The elements include:

- Normal business processes to ensure that only trustworthy individuals are authorized to access package preparation facilities in the central administration server. If you use separate packaging computers, access must be similarly controlled there as well. This is a matter of rights management, for which see the chapters on rights configuration.

- The use of trusted locations for the distribution of unsigned packages that are validated by the enterprise, and known to be good. The key security issue then shifts to securing the trusted locations against interference. For more on trusted locations, see the chapter on *Trusted locations*.

- The use of digital signatures to certify packages distributed from untrusted locations. Untrusted locations may

include servers that cannot be duly secured, external hard drives, or CDs used for distributing large packages. (Since CD drives can load any digital content from a CD, the drive should never be regarded as a trusted location.) Of course, it is possible to use digital signatures with all packages for additional peace of mind. For more on digital signatures, see the chapter on *Digital signing*.

Rights management will always be required in a security-conscious enterprise. RayManageSofti allows the flexibility to combine the other two approaches, trusted locations and digital signing, for the best balance of corporate security with maximum flexibility and minimal overhead.

When considering rights management, the following sections may be useful:

- RayManageSofti security mode
  An introduction to the matrix of the main segments of RayManageSofti and the types of rights controls applicable in each.

- General access considerations
  An overview that introduces some concepts and terminology, and provides a summary of general controls on access to the system.

- Configuring RayManageSofti rights
  A list of the elements of the system, with notes to help you understand the documentation approach to each of them.

- User groups
  A summary of all the account types needed in a fully operational system.

- Configuring administration server rights

- Configuring distribution server rights

- Configuring managed device rights

Each of the last three sections provides considerable detail, including a section overview, details of installation settings, a process and data flow diagram, and discussion of all features and their system impacts.

The Benefits

The combination of rights management, trusted locations, and digital signatures provides a powerful solution to prevent malicious use of deployed packages. Even though the only completely secure computer is powered down and disconnected from all networks, the approach gives a high level of assurance against attack through the manipulation of packages.

What's more, the solution provides this assurance level for a very small investment. As explained in the following chapters, the overheads to implement trusted locations and package signing to operate in this more secure manner are very small. The functionality in RayManageSofti means that very little has to change to provide the improved security.

Furthermore, the overheads remain low even if the corporation is switching from an existing, working RayManageSofti environment that has been using unsigned packages, to the use of signed packages. With due care, the changeover process is neither risky nor onerous.

The Limits

The RayManageSofti security offerings suffer the same limitations as similar industry-standard schemes. For example, with digital signing, it cannot be made absolutely impossible for certificates to be spoofed; and using trusted locations does not defend against DNS spoofing or IP spoofing. However, these technologies are currently the leading, industry-standard approaches to content authentication.

In the end, then, can these solutions prevent end-users setting up a malicious package for installation on their own computers, gaining elevated privileges, and then using this as a base for wider attack?

We noted at the beginning that RayManageSofti is not a security tool as such, and it cannot prevent activities which are currently permitted by the system. But by exactly the same token, on a secure (locked down) system, the end-user would need elevated administrator privileges to launch that scenario. In other words, "building your own launching pad" is restricted to those administrator-privileged users who should already be in the trusted inner circle. Non-administrator users could not play out that scenario.

In short, the combined use of trusted locations and digital signatures is a good solution for the goal: to prevent RayManageSofti being used to increase the risks of abuse by providing wrongly-elevated privileges to unauthorized personnel.

Assuring Integrity for the whole Application

In Security overview, we saw that package integrity has two closely-related aspects:

- **Protection against interference during storage in the software library**. This is provided through digital signatures applied at the completion of packaging and testing. The digital signature certifies that the package legitimately has corporate approval. Trusted locations can also assist in protecting packages during storage on distribution locations.

- **Protection against a "Trojan Horse" attack**, where an approved package is subsequently corrupted to contain malicious content, perhaps even during transmission. This aspect warrants further attention.

Digital signatures also partially address downstream file integrity - that is, a digital signature also testifies that the one file to which the signature is attached has not been tampered with since it was signed. However, in the RayManageSofti system where there are a number of separate files (which makes for network efficiencies), what prevents *other* elements of a duly certificated package from being corrupted during distribution, perhaps with deliberate and malicious intent?

One prevention would be to require the signing of every separate file. In the standard Microsoft distribution scenario, that is the requirement for multi-file structures. But the architecture of RayManageSofti allows for much greater efficiency, applying a digital signature once at the package (or application) level, so that there is not the potential for repeated interruptions during installation. This is possible because RayManageSofti already provides assurance of package integrity during transmission.

Quite independently of digital signing, all RayManageSofti packages are protected by MD5 cryptographic digests. Conceptually, an MD5 digest is like a super-checksum, applied to an entire file at a time. It is not possible in practice to hack a package and fudge its MD5 digests to remain correct. The MD5 digests ensure that the package which arrives at the managed device is an exact copy of the package stored on the server. This is true even when the package is distributed across unsecured or poor-quality connections. In other words, if you can protect the package *catalog*, you are assured of the integrity of the entire package. For example, when using trusted locations, it is a sufficient assurance of package integrity to collect only the catalog from a trusted location. Other package elements can be collected from other servers, and the MD5 digests provide assurance

of file integrity. (That said, RayManageSofti still provides for those who prefer to insist that *all* files come from trusted locations.)

The use of MD5 digests also means that, even without digital signing, the standard facilities of RayManageSofti provide protection again an uneducated attack (or, more likely, corruption due to network errors). It is not sufficient just to hack application files or implementation catalogs: such corruption will be detected and rejected in the normal course of operations. A better educated attacker might also try recreating new MD5 digests in all the right places to conceal the changes. This variation is stopped by the use of digital signatures, since these tie the overall package integrity back to the moment of signing by the publisher.

Thus it is standard operation for RayManageSofti to assure the integrity of every element of delivered packages. When this is combined with digital signatures to validate the identity of the publisher, you have good assurance of both: Legitimacy (or authenticity) and integrity.

The Question of Encryption

We noted at the beginning that secrecy or privacy are the separate provinces of encryption. The purpose of encryption is to prevent the interception of content during distribution. Encryption renders the content unreadable to third parties.

There is no direct connection between certifying the publisher (with digital signing) and keeping the content private (with encryption). Nor (perhaps more obviously) is there any connection between the use of trusted locations and encryption. For encryption, RayManageSofti supports the use of standard HTTPS to make secure connections between managed devices and distribution servers, and from the administration server to the distribution hierarchy.

Next

With the above overview in mind, you can explore in the next chapters the key elements of the RayManageSofti security offering: interaction with firewall software, digital signing, trusted locations, and rights configuration.

RayManageSofti and Personal Firewalls

This chapter provides details about how RayManageSofti works in conjunction with personal firewall software.

What is Personal Firewall Software?

A personal firewall is software on a computer designed to protect the computer from attack from the Internet, and possibly also prevent harmful content resident on the computer from being released to the Internet.

Just as dedicated firewall servers provide an interface between a private computer network and the broader, public Internet, a personal firewall is the interface between a computer and any network. All incoming content must pass through the firewall software, and (depending on the personal firewall software being used and its configuration) the firewall software may also inspect outgoing traffic. Incoming or outgoing traffic that does not meet the security criteria specified by the firewall's rules is blocked.

There are typically three classes of firewall rule:

- *Allow* rules specify how to recognize traffic as "safe". Safe traffic passes through the firewall unimpeded.
- *Deny* rules specify how to recognize dangerous traffic. Dangerous traffic is blocked.
- *Ask* rules specify that when traffic is received from specific IP addresses and ports, it should ask the end-user whether or not the content should be allowed access to the computer.

This chapter explains in general terms how RayManageSofti can work in conjunction with personal firewall software so that enterprise security is not compromised. It cannot address every configuration option of every personal firewall solution, so you should use the information in this chapter in conjunction with the documentation for the particular personal firewall solution in use, and (if necessary) your Raynet consultant.

RayManageSofti and Windows Firewalls

The Windows XP Internet Connection Firewall (ICF) (on Windows XP) and Windows Firewall (on Windows XP SP2 and Windows Server 2003 SP1) primarily block inbound traffic to the computer unless the traffic matches up with an outbound network request. With ICF or Windows Firewall running, RayManageSofti actions which are initiated from the managed device (such as requests for policy update, packages, and so on) work without additional configuration of the Windows firewall.

Configuration of Windows Firewall is required to enable RayManageSofti to function if:

- The File protocol is to be used to distribute packages from or upload status data to Windows XP SP2 or Windows Server 2003 SP1 computers
- Any of the HTTP, HTTPS, or FTP protocols are to be used for distributing packages from or uploading status data to Windows XP SP2 or Windows Server 2003 SP1 computers
- Windows XP SP2 or Windows Server 2003 SP1 computers are to be discovered and adopted under management, or execute remote execution tasks

- Windows XP SP2 or Windows Server 2003 SP1 computers are to be used for peer-to-peer file sharing

ICF is not enabled by default in Windows XP installations. ICF settings are configured separately for each connection.

Windows XP SP2 and Windows Server 2003 SP1 installations enable Windows Firewall on all connections (wired and wireless LAN, dial-up, and virtual private network (VPN)). Windows Firewall is also enabled by default on new connections. Windows Firewall allows specification of global settings to apply to each connection on a computer, and local settings, which override any global settings for the connection(s) to which they apply.

Configuring ICF

If you use ICF and ADM templates, you can set **Administrative Templates** > **Network** > **Network Connections** > **Prohibit use of Internet Connection Firewall** to `true` to disable ICF for all managed devices running Windows XP that connect to the domain.

Another approach is to configure ICF to open a limited number of ports required for RayManageSofti actions that are not initiated on the managed device or distribution server you are configuring (see *Port numbers used by RayManageSofti*). ICF must be configured locally on each computer. Consult the ICF documentation for details.

Configuring Windows Firewall

You can choose either to:

- Completely disable Windows Firewall (not recommended unless other firewall software is operating)
- Selectively configure Windows Firewall to permit RayManageSofti to function

If you choose to configure Windows Firewall, you will need to configure it on any Windows XP SP2 or Windows Server 2003 SP1 computers:

- Operating as distribution servers or administration servers
- Operating as managed devices
- That are not currently running RayManageSofti, but on which you want to perform a zero-touch inventory, or that will be targeted for adoption into management

Further details about the configuration required on each of these classes of computers are provided below.

You can configure Windows Firewall:

- As part of the initial rollout of Windows XP SP2 or Windows Server 2003 SP1, using a `netfw.inf` file
- On distribution servers, during installation of, or upgrade to, Deployment Manager 11.4 *infinity*
- Using Group Policy (recommended if you have Active Directory deployed throughout your enterprise)
- Using ADM and ADMX templates or logon scripts

- By deploying a RayManageSofti package containing a script to execute (useful if all Windows XP SP2 or Windows Server 2003 SP1 computers in your enterprise are already under RayManageSofti management)

- Locally on each computer (not recommended in a managed environment, for obvious reasons).

Further details about the recommended methods of configuring Windows Firewall are provided below.

A Summary of Windows Firewall Configuration

Windows Firewall requires different configuration depending on the role (distribution server, administration server, managed device, computer not yet under management) of the computer being configured.

Distribution Servers

When you are installing or upgrading RayManageSofti on distribution servers, the installation/upgrade script provides the option to configure the Windows Firewall exceptions necessary to allow RayManageSofti to function. Refer to the *RMS Implementation* and/or the *RMS Upgrade* for details.

If you chose not to configure Windows Firewall automatically during RayManageSofti installation, you must configure Windows Firewall using one of these methods:

- *Configuring Windows Firewall during rollout*
- *Configuring Windows Firewall using Group Policy*
- *Configuring Windows Firewall using ADM templates*
- *Configuring Windows Firewall using a RayManageSofti package*
- Alternatively, develop your own script to configure Windows Firewall according to your requirements.

Job Server

Raynet recommends that distribution servers should be configured to poll for jobs from parent distribution servers rather than listen for jobs. Distribution servers that use the listening agent require additional configuration to allow inbound TCP connections on the listening port (typically port `7010`). Set one of the following exceptions:

- For **Define program exceptions**, type this definition string:

```
%ProgramFiles%\ManageSoft\Replicator\ndlisten.exe:*:enabled:ManageSoft Connection Agent
```

- For **Define port exceptions**, type this definition string:

```
7010:TCP:*:Enabled:ManageSoft Connection Agent
```

In both these examples, `*` is specified as the subnet, meaning that incoming TCP connections on this port are accepted from computers anywhere on the network. Raynet recommends the use of localsubnet when appropriate.

Connections from Child Distribution Servers and Managed Devices

Distribution servers receive inbound connection requests from child distribution servers and managed devices on ports `139` and `445`. Enable the **Allow file and print sharing exception** to allow these connections.

Distribution servers also receive ICMP echo "ping" requests if managed devices are configured to ping before attempting to establish a connection. (This behavior is governed by the **NetworkSense** preference.) Enabling the **Allow file and print sharing exception** permits these (as does enabling **Allow remote administration exception** or the **Allow ICMP exceptions** with **Allow inbound echo request** selected).

In their capacity as Web/FTP servers, distribution servers receive inbound connection attempts on ports `21` and `80`. Configure IIS to allow communication on these ports. (Consult the IIS documentation for instructions.)

Administration Servers

If you are running a core server on a Windows XP SP2 or Windows Server 2003 SP1 computer, the Windows Firewall configuration required is the same as that required for distribution servers, with the exception that the listening agent is not used on the administration server. That is, you must configure IIS to allow inbound connections on ports `21`, `80`, and `443`. Consult the IIS documentation for instructions.

If you are running a reports server on a Windows XP SP2 or Windows Server 2003 SP1 computer, you must configure IIS to allow inbound connections on port `80`.

If you are running a data server on a separate physical server from your core and/or reports servers, you will need to configure Windows Firewall exceptions to allow SQL Server to listen to the network and receive TCP/IP connections. Refer to the Microsoft website (in particular, http://support.microsoft.com/kb/841249/en-us) for details about configuring Windows Firewall for use with SQL Server.

Other Windows XP SP2 or Windows Server 2003 SP1 Computers

The following sections outline the Windows Firewall configuration required on computers other than distribution servers.

Discovery, Adoption, Remote Execution, and Zero-touch Inventory

In order to successfully use RayManageSofti discovery (including determine remote execution credentials), adoption, remote execution, and zero-touch inventory functionality, you must configure Windows Firewall to accept inbound connections from the relevant distribution server on port `445`. You can achieve this in any of these ways:

- By enabling the **Allow file and print sharing exception**
- By enabling the **Allow remote administration exception** (recommended)
- By configuring **Define port exceptions** with port `445` and transport TCP

Peer-to-peer File Sharing

If you will configure managed devices to enable peer-to-peer file sharing, you will need to use the Windows Firewall **Define port exceptions** setting to selectively open the UDP and TCP ports used for file sharing. By default, port `6087` is used for both, but you can configure this using the **PeerSearchPort** preference for UDP and the **PeerPullPort** preference for TCP.

Remote Control of Managed Devices

If you use Deployment Manager in conjunction with remote control software such as TightVNC, you must perform some additional configuration of Windows Firewall on managed devices running Windows XP SP2 or Windows Server 2003 SP1.

If your remote control application uses any fixed ports, use the Windows Firewall Define port exceptions setting to selectively open the ports used by the remote control application.

If your remote control application uses dynamically-assigned ports for incoming connections, use the Windows Firewall Define program exceptions setting to specify the filename of the remote control application.

Be aware:
Use this same process to configure Windows Firewall to operate with other software, such as Symantec Ghost, if required.

Configuring Windows Firewall During Rollout

Note:
If computers in your enterprise already have Windows XP SP2 or Windows Server 2003 SP1 installed, this section is not relevant. It is only appropriate for configuring Windows Firewall during Windows XP SP2 or Windows Server 2003 SP1 rollout.

You can configure Windows Firewall during installation of Windows XP SP2 or Windows Server 2003 SP1.

If you have not already done so, download *Using the Windows Firewall INF File in Microsoft Windows XP Service Pack 2* from the Microsoft website (*http://www.microsoft.com/en-us/download/details.aspx?id=18996*), and familiarize yourself with its contents.

You can download a sample `Netfw.inf` file from the RayManageSofti knowledge base. You can replace the default `Netfw.inf` file with this sample file, or copy and paste the relevant lines from the sample file to the `Netfw.inf` file you will use with your Windows XP SP2 or Windows Server 2003 SP1 rollout.

To edit or replace the default `Netfw.inf` file with this sample file:

1. If you are using a CD image of Windows, copy it to a local file system so that you can edit `Netfw.inf`. (If your Windows image is already on a file system, proceed to the next step.)

2. Use `expand.exe` to decompress a copy of `Netfw.in_` from the `ic` or `ip` directory (the copies of the file are the same).

Be aware:
Using the Windows Firewall INF File in Microsoft Windows XP Service Pack 2 incorrectly lists the location of `Netfw.in_` as `Cd_drive:\I386\Netfw.in_`.

   The expanded `Netfw.in_` file is renamed `Netfw.inf`.

3. Replace `Netfw.inf` with the sample file from the RayManageSofti knowledge base, or edit this file to suit your environment.

   If you are using the sample file, you need to uncomment the appropriate lines to perform the configuration in your environment, according to whether or not the computers being configured are connected to an Active Directory domain. Uncomment lines by removing their leading semicolons (`;`).

```
[...]
; Uncomment these two lines to enable Remote Administration when
; connected to the domain
;HKLM,"SYSTEM\CurrentControlSet\Services\SharedAccess\Parameters\
FirewallPolicy\DomainProfile\RemoteAdminSettings","Enabled",0x00010001, 1
;HKLM,"SYSTEM\CurrentControlSet\Services\SharedAccess\Parameters\
FirewallPolicy\DomainProfile\RemoteAdminSettings","RemoteAddresses",0x00000000,"*"
[...]
; Uncomment these two lines to enable Remote Administration when not
connected to the domain
;HKLM,"SYSTEM\CurrentControlSet\Services\SharedAccess\Parameters\
FirewallPolicy\StandardProfile\RemoteAdminSettings","Enabled",0x00010001,1
```

```
;HKLM,"SYSTEM\CurrentControlSet\Services\SharedAccess\Parameters\
FirewallPolicy\StandardProfile\RemoteAdminSettings","RemoteAddresses",0x00000000,"*"
```

4. Use `makecab.exe` to recompress `Netfw.inf` and rename it to `Netfw.in_`.

5. Replace the copies of `Netfw.in_` in both the `ip` and `ic` directories of your Windows image with your updated version.

Configuring Windows Firewall Using Group Policy

If you have Active Directory deployed throughout your enterprise, you can configure Windows Firewall using Group Policy.

If you have not already done so, download the document *Deploying Windows Firewall Settings for Microsoft® Windows® XP with Service Pack 2* from the Microsoft website (*http://www.microsoft.com/en-us/download/details.aspx?id=7405*).

Follow the instructions in *Deploying Windows Firewall Settings for Microsoft® Windows® XP with Service Pack 2* to update your Group Policy objects with the new Windows Firewall settings.

Then, for each of the Windows Firewall settings discussed above that is relevant to your enterprise, follow the instructions in *Deploying Windows Firewall Settings for Microsoft® Windows® XP with Service Pack 2* to specify the Windows Firewall settings for appropriate Group Policy objects.

There are two profiles that use the Windows Firewall Group Policy settings:

- The **domain** profile is used when a Windows XP SP2 or Windows Server 2003 SP1 computer connects to the network that contains the organization's domain controllers
- The **standard** profile is used in other cases (for example when a computer connects to the Internet using a public ISP rather than through enterprise networks)

Raynet recommends that the firewall exceptions should be configured under the domain profile only, since the RayManageSofti functions allowed by these exceptions are typically only required when the computer is operating in the enterprise network environment.

Configuring Windows Firewall Using ADM Templates

If you do not have Active Directory implemented across your enterprise, you cannot configure Windows Firewall through Group Policy. An alternative in NT domains is to use ADM templates (administrative templates), which allow configuration of both the User (`HKEY_CURRENT_USER`) and Local Machine (`HKEY_LOCAL_MACHINE`) sections of the registry database. (You could also use logon scripts, but they only work for users with administrative privileges.)

You can use this file with the System Policy Editor (`poledit.exe`) to create one or more policy (`.POL`) files. These policy files are stored on each domain controller to which Windows XP SP2 or Windows Server 2003 SP1 computers may connect.

Be aware:
Current versions of Windows do not cover firewall configuration with ADM templates. Please refer to the Microsoft knowledge base to get information on ADM and its successor technology ADMX:

- *http://support.microsoft.com/kb/816662/en*
- *http://msdn.microsoft.com/en-us/library/windows/desktop/bb851669.aspx*

- *http://msdn.microsoft.com/en-us/library/bb530196.aspx*

Configuring Windows Firewall Using a RayManageSofti Package

If all Windows XP SP2 and Windows Server 2003 SP1 computers in your enterprise are already under RayManageSofti management, you can create one or more packages containing scripts to configure Windows Firewall.

You might choose to create two packages:

- One containing a script to configure Windows Firewall on distribution servers and administration servers
- One containing a script to configure Windows Firewall on managed devices

In each case, make sure that the script runs as `SYSTEM`, not the user currently logged in.

Consult your Raynet support representative for help if necessary.

RayManageSofti and other Personal Firewall Products

Most personal firewall products can be configured to inspect both incoming and outgoing traffic.

These firewalls may notify end-users when an outbound network connection is being initiated, and allow the end-user to block the outbound request. If you are using personal firewall software that behaves this way, and you want it to continue to behave this way, you will need to educate end-users in your organization about what requests should be accepted.

Alternatively, consult the firewall software's documentation to determine how to configure it to allow RayManageSofti to operate quietly in that environment. In general, you will need to create rules that specify which programs are allowed to make outbound requests, and define parameters for them including port numbers and network locations.

Port Numbers Used by RayManageSofti

The port numbers used by managed devices, distribution servers, and administration servers are detailed in a series of tables (below). The tables detail:

- The port number used - you may need to configure your firewall software to allow connections on this port.
- The protocol used on this port.
- Whether the initial request is incoming (to the device) or outgoing (initiated by the device). Most firewall software automatically allows outgoing connections, and accepts response traffic. If your firewall software does not operate this way, special configuration may be required.
- What system software or RayManageSofti feature uses this port for communication.

In general, RayManageSofti uses standard ports configured on your computers for LDAP requests, DNS queries, HTTP/HTTPS/FTP requests, and so on.

Port Numbers Used on Administration Servers

On administration servers, the ports used by RayManageSofti are:

Port number	Protocol	Incoming / Outgoing	Used for
21 or your configured port	TCP	Outgoing	Uploads or downloads configured to use FTP. This is the control connection, used for commands.
53	TCP or UDP	Outgoing	DNS queries
80, or your configured port	TCP	Incoming / Outgoing	HTTP, if used for uploads or downloads
88	TCP	Outgoing	Kerberos (Authentication on devices running Windows 2000 or later releases)
135 ports above 1024	TCP or UDP	Outgoing	RPC, for automatic update of policy, and for adding packages to policy
137 138 139	TCP	Outgoing	NetBIOS, if configured for operation over TCP/IP
138	UDP	Outgoing	WINS
139 and/or 445	TCP	Incoming	SMB file shares
389	TCP or UDP	Outgoing	LDAP queries for policy generation
443, or your configured port	TCP	Incoming / Outgoing	HTTPS, if used for uploads or downloads
Port numbers above 1024 (dynamic)	TCP	Incoming (for active mode servers) / Outgoing (for passive mode servers)	Uploads or downloads configured to use FTP. Your firewall software must be configured to open these ports when needed—most firewall software automatically allows outgoing traffic, so no special configuration may be required.
1443	TCP	Incoming	SQL Server on the data server
1443	TCP	Outgoing	SQL Server on the core and reports servers
7010, or your configured port for the RayManageSofti listening agent	TCP	Incoming	Listens for RayManageSofti instructions such as requests to distribute packages

Port Numbers Used on Distribution Servers

Distribution servers that are running personal firewall software will typically need to be configured to allow RayManageSofti managed devices to request data from the network.

Depending on your personal firewall software, you may need to set up a rule to allow access by RayManageSofti managed devices, or simply configure the software to allow unsolicited inbound network traffic on the required ports.

On distribution servers, the ports used by RayManageSofti are:

Port number	Protocol	Incoming / Outgoing	Used for
21 or your configured port	TCP	Outgoing	Uploads or downloads configured to use FTP. This is the control connection, used for commands.
80, or your configured port	TCP	Incoming / Outgoing	HTTP, if used for uploads or downloads
88	TCP	Outgoing	Kerberos (Authentication on devices running Windows 2000 or later releases)
135 ports above 1024	TCP or UDP	Outgoing	RPC, for automatic update of policy, and for adding packages to policy
137 138 139	TCP	Outgoing	NetBIOS, if configured for operation over TCP/IP
139 and/or 445	TCP	Incoming (for distribution servers with file shares)/ Outgoing (for distribution servers that use files shares)	SMB file shares
Port numbers above 1024 (dynamic)	TCP	Incoming (for active mode servers) / Outgoing (for passive mode servers)	Uploads or downloads configured to use FTP. This dynamically assigned port number is the connection used for data transfer. Your firewall software must be configured to open these ports when needed - most firewall software automatically allows outgoing traffic, so no special configuration may be required.
7010, or your configured port for the RayManageSofti listening agent	TCP	Incoming	Listens for RayManageSofti instructions such as requests to distribute packages
443, or your configured port	TCP	Incoming / Outgoing	HTTPS, if used for uploads or downloads

Port Numbers Used on Managed Devices

On managed devices, RayManageSofti listens on ports for activity only if:

- Remote execution tasks are in progress
- Peer-to-peer file sharing is enabled

Otherwise, the managed device is event-driven, with events occurring at times specified by schedules. The actions that the managed device then performs may require that some ports are not blocked.

On managed devices, the ports used by RayManageSofti are:

Port number	Protocol	Incoming / Outgoing	Used for
21 or your configured port	TCP	Outgoing	Uploads or downloads configured to use FTP. This is the control connection, used for commands.
Port numbers above 1024 (dynamic)	TCP	Incoming (for active mode servers) / Outgoing (for passive mode servers)	Uploads or downloads configured to use FTP. This dynamically assigned port number is the connection used for data transfer. Your firewall software must be configured to open this port when needed.
53	TCP or UDP	Outgoing	DNS queries
80, or your configured port	TCP	Incoming / Outgoing	HTTP, if used for uploads or downloads
88	TCP	Outgoing	Kerberos (Authentication on devices running Windows 2000 or later releases)
135 ports above 1024	TCP or UDP	Outgoing	RPC, for automatic update of policy, and for adding packages to policy. Inter-process communication between computers. Blocking this port largely stops one computer from communicating with another remotely.
139	UDP	Ingoing / Outgoing	NetBIOS name resolution. NetBIOS allows discovery of computers in a network. Blocking this port prevents NetBIOS from seeing certain hosts.
139 and/or 445	TCP	Outgoing	SMB file shares
161	UDP	Incoming / Outgoing	Used by the SNMP service, if it is running. This port may be scanned during discovery, if the discovery process is configured to do so. Refer to the Discovery chapter of the *RMS Discovery* for details.
445 port numbers above 1024 (dynamic)	UDP	Incoming / Outgoing	Port 445 supports AD / DNS name resolution. DNS allows discovery of computers in a network. Blocking this port prevents NetBIOS from seeing certain hosts.
389	TCP or UDP	Outgoing	LDAP queries for client-side policy generation
443, or your configured port	TCP	Incoming / Outgoing	HTTPS, if used for uploads or downloads
1080, or your configured port	TCP	Incoming/ Outgoing	SOCKS, if used for uploads or downloads

Port number	Protocol	Incoming / Outgoing	Used for
1230 (UDP)	UDP	Incoming	Wake on LAN
6087, or your configured ports (UDP and TCP)	TCP or UDP	Incoming/ Outgoing	Peer-to-peer file sharing
7020, or your configured port (see the *Discovery* chapter of the *RMS Discovery*)	TCP	Incoming	The TCP-based remote execution listening agent, if used
16992 and 16993, AMT ports (see the *Working with AMT-capable devices* chapter of the *RMS Discovery*)	TCP	Incoming	This port is used when collecting data from an AMT-capable managed device at the beginning of a remote execution task

Port Numbers Used for the Remote Console

If there is a firewall between a machine running the remote console and the AS, the ports listed in the table below will need to be opened for the listed operations to occur.

Port number	Protocol	Incoming / Outgoing	Used for
TCP: 139 and 445 UDP: 137 and 138	TCP or UDP	Incoming / Outgoing	File and Printer Sharing
443 (if used)	TCP	Incoming / Outgoing	Secure World Wide Web Services (HTPS)
1443	TCP	Incoming / Outgoing	SQL Server Windows NT
80 (if used)	TCP	Incoming / Outgoing	World Wide Web Service (HTTP)
7010	TCP	Incoming / Outgoing	MGS Listening Agent

Digital Signing

This chapter provides both an overview of digital signing and details about setting up its use in Deployment Manager. For a preliminary overview of where digital signing fits within the wider scope of the Deployment Manager security offering, see the chapter *Security overview*. For unfamiliar terms, check the *RMS Glossary*.

The Basics of Digital Signing

Digital signing is a system of delegated trust. In this respect it is similar to the issue of passports by national governments. In a passport, the government (as the trusted authority) is certifying that you are who you say you are. With digital signatures, an authority is saying that the software publisher is who it says it is.

Trusted Authority

The trust chain therefore begins with the initial trusted authority. In the digital world, these are called certifying authorities. They are commercial enterprises established to bear witness that they have validated the identity of another enterprise. The best known of these certifying authorities is VeriSign. The Raynet recommendation is that you make use of the combination of VeriSign as the certification authority with standard Microsoft technology for validating the certificates and signatures (as discussed later). Therefore in this discussion, we will focus entirely on VeriSign as the certification authority.

A certification authority is recognized through its own certificates of authority. The VeriSign certificates are installed as a standard part of the Windows operating system. By using the standard digital signing approach, therefore, you need take no action to “educate” managed devices about the certification authority. That step is completed already.

It remains for the certification authority to issue the software publisher with the analog of a passport - a digital authority declaring that the publisher is accurately represented.

Certifying Technology

Publisher certification depends on the use of public key cryptography. In public key cryptography systems, every entity has two complementary keys: a public key and a private key. As the names imply, public keys are widely distributed to users, while private keys are kept safe and only used by their owners. A public key is derived from the private key, and can be used to successfully decrypt something encrypted by the private key - and indeed, only things encrypted with that private key. For more information on public keys and private keys, please see *Introduction to Public Key Cryptography* at *http://www.verisign.com.au/repository/tutorial/cryptography/intro1.shtml*.

In summary then, if you know that a public key is from your friend, then you can be certain that anything you decrypt with that public key is also from your friend.

In the digital world, Microsoft has set up its operating systems with the public keys of approved certification authorities. It also provides mechanisms for checking encryptions against those public keys. The operating system therefore tells us when something has been encrypted by the certifying authority, such as VeriSign.

What VeriSign encrypts is called a software publishing credential. This identifies a publisher, and includes the publisher's public key. It is like a passport, or a letter of introduction. In effect, it says "VeriSign has investigated

this publisher and declares it to be trustworthy. Duly and unmistakably signed." Trust is thereby delegated from VeriSign's public key to the publisher's public key, distributed in the software publishing credential.

Only a file that has been digitally signed with the publisher's private key can be successfully verified using the same publisher's public key. That public key is distributed along with the software publishing credential, so that anyone receiving the credential is also able to verify files signed with the publisher's private key. From the end-user's viewpoint, a file successfully verified using the publisher's public key can only have been digitally signed using the publisher's private key. Since the publisher holds this private key very securely, this authenticates the source of the code.

There is a further benefit. For speed, the signing and checking algorithms work with cryptographic file digests. This has the beneficial side-effect that a file successfully verified is also certified to be a good copy of the original signed file, and has not been tampered with.

Be aware:
Do not confuse this with the similar use of MD5 cryptographic digests by Deployment Manager. Although a similar technology, the latter applies to all files in a RayManageSofti package, and is completely independent of digital signing technology.

Be clear that we have been discussing cryptography only in the context of the digital signature itself. Digital signing does not encrypt the file it protects.

So far we can see that the technology ensures that what the publisher signs can be validated as being from the publisher. But this technology is (like justice) blind: anyone can issue a public key/private key pair. Thus we come back to the certification authority - an authority able to say that the private key is held by the company it claims to represent.

Getting Certification

Your enterprise applies to VeriSign for recognition in an online process. As part of that process you determine your private key, and from that moment on, you keep it absolutely secure. (VeriSign recommends having it on a floppy disk kept under lock and key. A secured CD would serve the same purpose.) This private key is never sent to VeriSign or anyone else, so if you lose the private key, you lose the ability to sign digital content. It is worth having a backup copy in a very secure location.

At application time you also create a password to protect your use of the private key. That password must be entered every time the private key is used to sign content. The password itself should also be recorded in an appropriate backup location and access-restricted. (Remember to make provision for succession and emergency use.)

VeriSign will then take 3-5 days to verify the identity of the applying enterprise. When the enterprise is approved, you receive a PIN with instructions on how to collect your software publishing credential (`.spc`) file. This is your certificate that applies VeriSign's authority to your public key/private key pair. It is equivalent to your passport as an identity document. VeriSign's service mark for this is Digital ID.

Signing Packages

With the combination of your private key (kept secure) and your public key contained in your software publishing credential, you can sign files using standard Microsoft technology called Authenticode. With Deployment Manager, it is even more convenient: with a single setting (described in *Procedures*), you can validate entire applications, also using the Authenticode technology.

The Microsoft technology (with standard settings) advises end-users of the certificate status, and asks whether to

proceed (even if the certificate is not validated). Having this question pop up many times for different application elements could be very annoying. In contrast, the Deployment Manager approach manages to combine file separation (for smaller downloads and easier updates) with a single validating signature.

There is another loosely-related aspect that package-level signing brings up: the implicit *assumption* of end-user approval. Under management with Deployment Manager, it is no longer a requirement that you allow end-users to have personal approval of the publisher's certificate. You *may* provide that level of feedback to end-users if you wish to do so; but the principle is that the RayManageSofti installation agent is checking the certificates, and providing the approval to install the package.

Timestamping

Because public / private key pairs are based on mathematical relationships which can theoretically be "cracked" with a great deal of time and effort, it is a well-established security principle that digital certificates should expire. Your VeriSign Digital ID will expire one year after it is issued.

However, most software is intended to have a lifetime of longer than one year. To avoid having to resign software every time your certificate expires, VeriSign and Microsoft introduced a timestamping service. Now, when you sign packages, attaching a VeriSign URL allows the signature to be timestamped. As a result, when your code is downloaded, end-users (and indeed the RayManageSofti installation agent) will be able to distinguish between two kinds of signature:

- Files signed with a certificate that was invalid or expired at the time of signing, which should *not* be trusted.
- Files signed with a certificate which was valid at the time the file was signed, but which has subsequently expired. Such files *should* be trusted.

This means that you will not need to worry about re-signing packages when your Digital ID expires. VeriSign is currently the only certification authority offering the timestamping service. It is a free service to VeriSign customers.

On the Managed Device

From the above discussion, it should be clear that digital-signature checking occurs on the managed device. Each managed device needs three things to achieve this checking:

- The client-side Authenticode software, installed on any Microsoft operating system using Microsoft Internet Explorer 4.0 or later, or the Internet Client SDK (all devices under management with RayManageSofti should comply)
- The certificate of the certification authority (which, in the case of VeriSign, is present on all the same Microsoft operating systems)
- The public key, signature, and software publishing credential of the software publisher (in this case, your enterprise), all supplied automatically with the signed file as part of the signing process.

One thing that the managed device does not need is access to the Internet.

The process of verification with Deployment Manager is explained in detail later in this chapter.

Worldwide Use

Cryptographic technologies are regarded by many governments as weapons of war, and are therefore subject to export controls (particularly for very secure encryptions).

Public and private key cryptography today makes use of either 40-bit or 128-bit encryption. In general 40-bit encryption is considered adequate for many business purposes, but 128-bit encryption is required where there is a real threat of significant hostile decrypting resources being applied.

VeriSign is not restricted by export restrictions. Their service, consisting essentially of an authentication service for public keys but not the keys themselves, is unaffected by U.S. export regulations, and therefore does not have any export restrictions. For more information, see: *https://www.symantec.com/ssl-certificates*.

Hostile Publishers

One consideration about the trust chain is that it has to allow reasonable accreditation for applicants, not all of whom have the public visibility of a Microsoft. This exposes two related risks:

- VeriSign may wrongly authorize an impostor intent on misusing the trust chain to do harm. Clearly it is in VeriSign's best commercial interests to ensure that this does not happen, as this would undermine the trust that is the essence of their business offering.
- A rogue employee of another duly authorized software publisher may misuse that employer's certification to do harm.

Hopefully, all enterprises realize the importance of policing proper use of digital signatures, and put in place fail-safe controls to prevent such abuse by individuals.

Industry Standard or Custom

It should be noted that the Microsoft Authenticode technologies allow for enterprises to serve as their own certification authorities, creating custom certificates. This is a useful technology for testing purposes, especially while waiting for the issue of an industry-standard `.spc` file.

In production use, Raynet recommends against custom certificates and in favor of using the combination of VeriSign with Microsoft technologies, for the following reasons:

- There is less infrastructure required
- There is less set-up of managed devices required (the CA certification is already provided)
- Using standard technologies makes it easier to widen the applicable scope of the SPC certification, for example to send authorized packages to business partners
- Only VeriSign offers the timestamping service

How RayManageSofti Uses Digital Signatures

This section outlines the differences between signed and unsigned packages in RayManageSofti, how the processing varies between the two, and how these processes help protect against attack.

Signed and Unsigned Packages

There are only two technical differences between signed and unsigned packages in Deployment Manager. They are:

- An addition to the content of the package's implementation archive (`.ndc` file)
- An additional small file conveying the digital signature

Digitally signed packages are only available for Windows managed devices.

Additions to the Implementation Archive

The implementation archive is the core file in the package for describing the application files, user settings, and other information required to control the installation of the software.

For a signed package, there is one extra conditional statement in the implementation archive. Its purpose is to declare the presence of the digital signature file, and it takes the following form:

```
# if ($(OS)=="Win95") || ($(OS)=="WinNT")
File: authcode.cab;
href = packagename.ndc.cab;
action = Authenticode;
...
# endif
```

Be aware:
The operating system identifier "Win95" identifies Windows 95, 98, and Me, and the identifier "WinNT" identifies all versions based on NT.

This declaration is used in the process of certificate checking, described shortly. For information about attacks that remove this section from the implementation archive, see *Summary: How secure is secure?*.

The Digital Signature File

A normal RayManageSofti package contains a catalog (`.osd` file), an implementation archive (`.ndc` file), and compressed application files.

Signed RayManageSofti packages also include a digital signature file (`.ndc.cab` file). This takes the form of a CAB file, signed with Microsoft Authenticode technology that contains a single item: the MD5 digest calculated at packaging time for the package's implementation archive (`.ndc`).

Therefore this file is used both to carry the digital signature and to provide further validation of the package, as described in the process flow below.

Process Flow: Checking the Certificate

Refer to this figure to better follow the upcoming discussion.

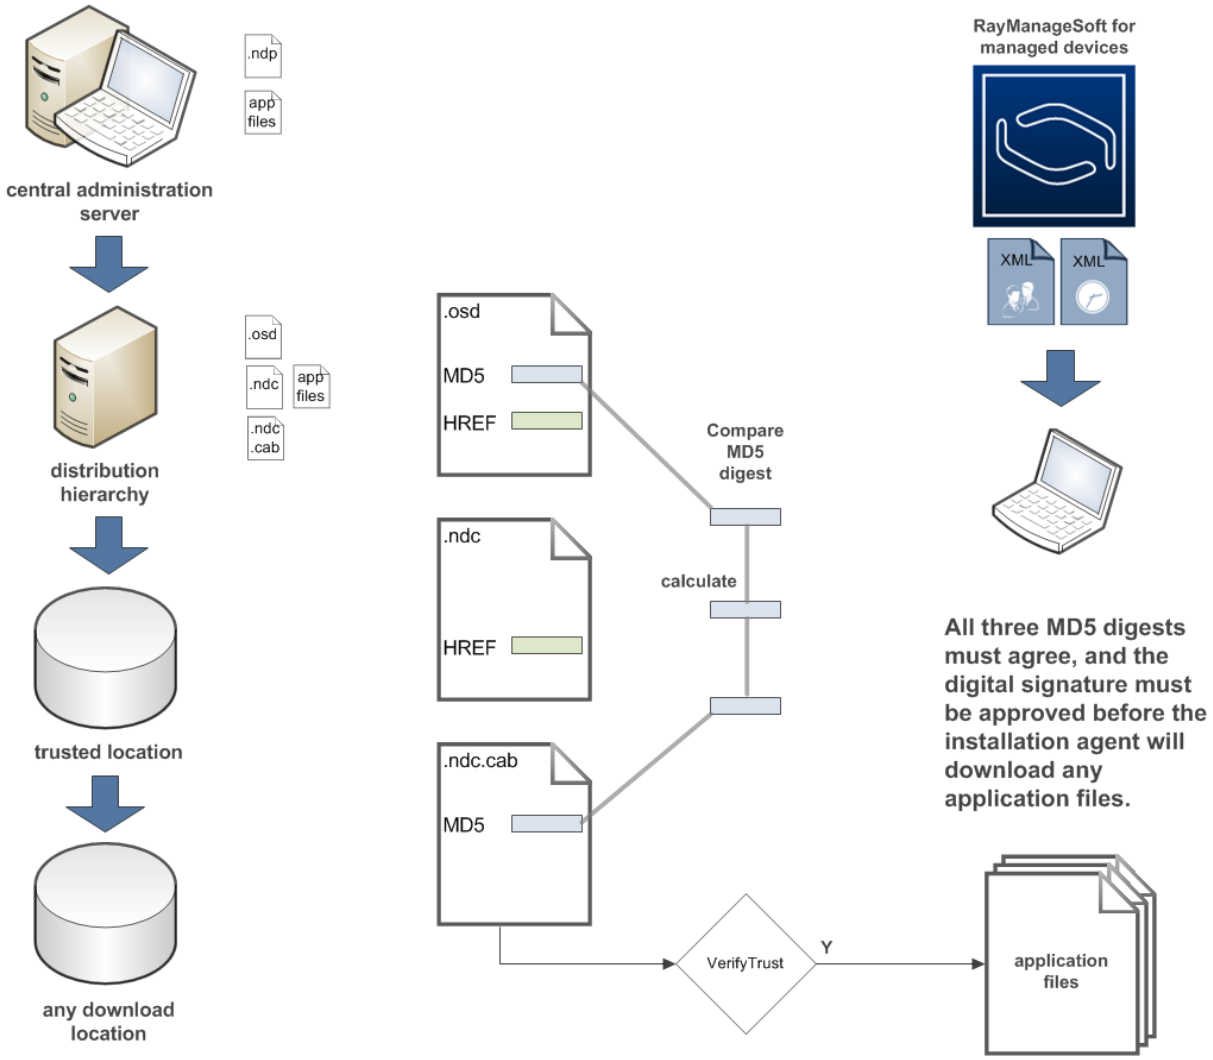

Before the process of publisher and package validation begins, several obvious prerequisites occur:

- On the administration server, a package creator creates a project file (`.ndp`) that describes the files, sourced from the software library, to include in the package, and provides details about the installation of the application on managed devices. For more information, see *Security tab-fields* in the *RMS Software Deployment* or the *application help*. All that is required is selecting the check box for digital signing, and supplying some standard file locations (discussed further below).

- Given these details, the Deployment Manager packaging agent compresses the application source files for distribution, and prepares three meta-files to complete the package: the application catalog (`.osd` file), the implementation archive (`.ndc` file), and the digital signature file (`.ndc.cab` file).

- All these files move from the administration server staging area through the distribution hierarchy in the normal distribution process. There are no changes to any of the files throughout distribution.

Everything is now in place for the verification process to begin. In due course, the package is presented to the RayManageSofti installation agent on the managed device. Note that it makes no difference to the process how this process is triggered, whether through policy, through a scheduled event, through end-user action in the selection agent, or any other method. The process remains identical in every case:

1. The installation agent checks the current preference settings to determine whether a digital signature check is required. There may be two reasons why one is not required:

   - Preference settings may not support digital signatures at all.
   - Preference settings may support the joint use of digital signatures with trusted locations. In this case, if the package comes from a trusted location, the signature checking is skipped, and the installation follows normal processes.

   Preference settings may be set in the registry or by command line options. See *Summary: How secure is secure?* for more about protection of these settings. The relevant preference settings are listed in the details section below. The use of preference settings to combine digital signing with trusted locations is detailed in *Mixing digital signatures with trusted locations*. For details on the permissible values for preference settings, see *RMS Preferences for Managed Devices*.

2. Since a digital signature check is required, the process continues with the download of the application catalog (`.osd` file). This contains both the URL of the implementation archive (`.ndc` file) and the MD5 digest of the implementation archive.

3. The implementation archive is downloaded from the URL specified in the catalog. (If the package is being delivered on a CD, the implementation archive is loaded from the CD.)

4. The installation agent calculates a fresh MD5 digest for the implementation archive just received.

5. Assuming that the preference setting **CheckCatalogDigest** is `True`, the installation agent validates that the calculated value for the MD5 digest matches the value calculated on the administration server. If these two digests are not identical, then the download enters the normal retry cycle (to protect against accidental disruption of a given transfer), and if the difference persists, the transfer is failed.

6. With the integrity of the implementation archive (`.ndc` file) now verified, the installation agent looks up the URL for the digital signature file (see *Additions to the implementation archive*). This is where the process is specialized when digital signatures are in use.

   If the package has not been digitally signed, the installation is terminated immediately, and (depending on preferred levels of feedback) an alert may be displayed to the end-user.

   End-users may next see the following message:

   

   If, on the other hand, the package has been digitally signed, the process continues.

7. For a signed package, the digital signature file is now downloaded. This file, which is less than 8KB in size, is the only additional download ever needed in the use of digital signatures compared with standard operations.

8. The installation agent now extracts from the digital signature (`.ndc.cab`) file the copy of the MD5 digest for

the implementation archive (`.ndc` file), as recorded on the administration server. This copy of the MD5 digest must also match the one calculated in step 4. If not, the installation is terminated immediately, and an error message may be displayed for the end-user (depending on the preference settings for level of feedback).

This means that there is potentially a three-way validation of the implementation archive, comparing its present calculated value with values separately stored in the application catalog and the digital signature file.

9. With all files now cross-checked, the Microsoft API for WinVerifyTrust is called on the digital signature file. This returns either acceptance or rejection of the signature on the CAB file, based on the root certificate authority already available on the managed device. Note that this step does not require any Internet (or indeed network) connection.

   Note that the Microsoft WinVerifyTrust technology tries to provide the end-user with feedback on the certification outcomes, and to ask for Yes/No authorization to proceed. Whether or not these messages are shown to the end-user under RayManageSofti depends on the interaction of preference settings with outcomes:

   - If **UserInteractionLevel** is set to any setting less than `Full`, all Microsoft messages are suppressed, and the installation will proceed if, and only if, the signature is approved. This is the recommended setting to use with RayManageSofti and digital signatures, since it ensures that end-users cannot over-ride policy or security.

   - If the **UserInteractionLevel** for RayManageSofti is set to `Full`, the following choices are applied:

     - If the signature is **accepted**, the Microsoft dialogs will be displayed to the end-user. Note that this allows the end-user the option of saying "No" to an installation, even though the digital signature is approved. The Microsoft dialog does not take account of policy settings that may declare installation to be mandatory.

     - If the signature is **rejected**, RayManageSofti tests the setting of **ForceValidSignature**.

       - If **ForceValidSignature** is `True`, the end-user is notified of the failure, and the installation stops without giving the end-user the Microsoft choice of proceeding anyway. We strongly recommend setting `ForceValidSignature = True` whenever you use digital signatures.

       - If **ForceValidSignature** is `False`, the Microsoft dialog will be displayed, giving the end-user the choice of proceeding with the installation even though the signature has not been verified.

WARNING
This last option allows the end-user the choice of breaching security by permitting installation of an unsigned or wrongly-signed application.
An installation that fails because of a digital signature failure is shown with other installation failures in Deployment Manager reporting, and the installation logs will report the reason for failure.

Summary: How Secure is Secure?

How does the process described above help to secure RayManageSofti against abuse? There are a number of lines of attack against which a secure package is a valuable defensive contribution.

Attacks against the Package

The major risk to security with Deployment Manager is that an attacker will fabricate a package that contains malicious code, and will abuse the elevated privileges that permit RayManageSofti to operate.

The above detailed discussion makes clear that the use of digital signing (particularly in combination with trusted locations) effectively closes off that risk.

- The implementation archive is strongly protected. It is impossible to hack an archive without detection against the MD5 digest recorded at packing time. Recall that the digest value recorded in the application catalog (`.osd`) and in the digital signature file (`.ndc.cab`) must exactly match the recalculated MD5 digest of the delivered archive.
- Even if an attacker has the means to create valid catalogs and implementation archives, you remain protected by the security of your software publishing credential.

Attacks against the Application Catalog

An attacker could hack an application catalog (`.osd` file), since this is not protected by its own MD5 digest. However, damage is unlikely for these reasons:

- The catalog (`.osd`) contains little of value for an attack (the location of the implementation archive and the application name)
- The implementation archive (`.ndc`) remains central to the installation process, and remains protected even if the catalog is compromised
- Any attack on the catalog does not affect the software publishing credential (the publisher's certificate that authorizes approved applications)
- If further assurance is desirable, structure your distribution hierarchy such that the application catalogs (`.osd` files) are collected only from trusted locations

Attacks against the Implementation Archive

A possible attack is to hack the implementation archive to remove the statement that identifies the digital signature file. This requires that the attacker also has the ability to create new, valid MD5 digests, and can also hack the application catalog and the digital signature file to replace the MD5 digest copies in each one. Assuming that all this is achieved, the package would then be an unsigned package, and will be treated as such when signed packages are in use. In other words, if it comes from a trusted location it is a risk. If it does not come from a trusted location, it will be rejected as an unsigned package.

Attacks by Turning off Digital Signing

Another potential loophole might appear to be the use of preference settings to turn on (and off) the use of digital signing (see step 1 in the process listing above). What prevents a malicious user from turning off digital signing on his or her own managed device, and then using that as a platform for attack?

The answer is that the settings for the use of digital signatures reside in the machine hive of the registry, not the individual user hive; and those registry settings are locked. This means that in order to change the settings on a locked-down managed device, the user must already have the elevated privileges that are the 'benefit' to be gained by hacking RayManageSofti. (Note too that access to registry settings can be further restricted to particular users using the Windows access control list.) In short, an enterprise must manage the trust of those with administrator privileges. Without those privileges, it is not possible (on a locked-down computer) to modify the use of digital signatures with RayManageSofti.

Attacks through Policy or Schedules

None of the triggers for the installation agent (such as a scheduled event, a policy check, a call from the selection agent, or a command line) have any capacity to change privileges. Only the installation agent can do that, and only in response to the package settings. The process discussion and diagram above make clear that the protection of the package is not affected by the way that the package was presented to the installation agent. The package validation always occurs.

Therefore, protecting against attacks on (or through) the package also protects against all these other paths. Without the package, an attacker can do nothing.

The Next Line of Defense

Protecting the package is critical, and the use of digital signing (particularly with trusted locations) provides, as we have seen, a strong defense. It is therefore reasonable to say that, once digital signing and trusted locations are in use, the most likely risk is from a 'loose cannon' amongst approved employees. Only those with access to your packaging facilities and your publishing credential (including your private key and its protective password) will be able to emulate your approved packages sufficiently well to get past the protection that RayManageSofti provides on the managed device. And only those with pre-existing administrator privileges on a managed device could turn that into a platform for wider attack, with or without RayManageSofti.

Since this is the case, your security policy must pay particular attention to the following critical issues:

- Restricted access to packaging facilities
- Security of the password protecting the private key, and procedures for timely password renewal
- Protection of your software publishing credential (preventing it being copied to a hostile location)
- Restricting the issue of administrator privileges on managed devices
- Logging of all packaging and other activity (even by approved personnel) on core servers

Impact of Implementing Security

We have now seen the methods that Deployment Manager uses to protect against delivery of fraudulent packages, and can see the benefits of that assurance. It is now appropriate to consider the overheads that the additional security invokes. A reasonable question to ask is whether the potential risks justify the overall cost package. The short answer is that, where risk moderation is important, the additional costs are negligible in relation to the benefits. This section provides a quick summary, and there is more detail about setup requirements and detailed procedures in *Procedures*.

Impacts can be categorized as those on deployment operation, those on the network, and other business impacts.

Impacts on Deployment Operations

As the following overview shows, there is very little impact on deployment operations.

Preparing New Packages

Once the infrastructure is in place (as detailed later), the only change in packaging is to select one extra check

box and ensure that the pointers to the required files are in place. The rest is automatic.

Distributing New Packages

There are no changes to the distribution process. For each application there is one additional small file that carries the digital signature. This is distributed automatically with the rest of the package files.

Verifying Existing Signed and Unsigned Packages

Verification is the process on the managed device of checking applications against metadata to ensure there is no corruption. If the check fails, most often repair attempts follow immediately.

What impact does the distribution of signed packages have on verification and repair? Once digital signing is in use, you can choose whether to also distribute some unsigned packages from trusted locations. This means it remains possible to have a mix of signed and unsigned packages into the future.

Except for the proper refusal to update anything with an unsigned package that does not come from a trusted location, all verification and repair processes proceed as usual. This table summarizes all the cases:

Update and package	With trusted locations	Without trusted locations
Signed update to a signed package	Normal process	Normal process
Unsigned update to an unsigned or signed package	From a trusted location, normal process	Not from a trusted location, the update will properly be refused
Signed update to an unsigned package	Normal process	Normal process

Distributing Updates to Packages

There is no change to the preparation and distribution of updates. Naturally, it would be normal practice to continue to sign updates for packages that were once distributed as signed.

Distributing Packages with Dependencies

There is no change required for packages with dependencies, other than the obvious one: dependencies of signed packages must also be either signed or delivered from trusted locations. (It is generally more straight-forward just to sign them.)

Reporting

There are no changes to reporting. Checking the certificates installed on the end-user's computer becomes one extra step in resolving any installation failures.

User Changes

In the normal case, switching to signed packages need have no impact on end-users. You can set preferences in such a way that even the expected failure case (delivering an unsigned package from an untrusted location) does not disturb the end-user. For more about user feedback levels, see *Preparing a settings update for managed devices*.

Impact on Network and other Resources

Network impacts from switching to signed packages are small. There is one additional digital signature file (`.ndc.cab`) per application, under 8K in size that is transferred and checked before any other application files are transferred. With normal structures in place for the spreading of policy checking and update loads, the impact generally remains negligible. The computational load of validating the signatures is also very small.

Impact of Certificate Renewals

Potentially the most significant impact to the use of digital signatures can come at the annual certificate renewal. Sadly, the renewal does not extend the existing certificate, but provides a new one. (This is arguably better security practice, as it reduces the risk of compromise of the certificate.)

The downside is that package installations may fail, or worse, that end-users may see dialogs that allow them to install packages that have invalid signatures. There are two ways to ensure that this does not happen.

- The traditional way (and still the required way with any certification authority other than VeriSign) is for you to be very punctual in obtaining your new certificate. Then all applications must have their security settings (the file pointers within application details) updated. Then they all have to be repacked and redistributed with the new software publishing credential.

- The easier way is to use the timestamping service (see *Timestamping*) every time you sign a package. Provided that you have timestamped all your digital signatures, you never need to worry about certificates expiring for those packages. Timestamping saves you from re-signing previously packed applications. With timestamps in use, you need sign with the new certificate only those packages you are packing or repacking after the expiry of the old certificate.

The comparison of effort at renewal time is a strong argument in favor of using the timestamp with every digital signing of an application package.

Setting Up Infrastructure for Digital Signing

This section provides an overview of the changes in infrastructure needed to support the use of digital signatures with Deployment Manager. Although there are, naturally, a number of things to do, the effort involved is small relative to the benefits to be gained.

This section assumes a working Deployment Manager implementation, and notes only the infrastructure *changes* required. For the related issues of migrating existing packages, see *Migrating existing packages*.

The Ordered Process

Here is a suggested ordering of the overall process. Details of different steps will be found below.

1. It's important to first determine whether you will be supporting the use of digital signatures with the operation of trusted locations (as recommended). If so, it is best to set up the trusted locations first, as their operation can smooth the changeover to the use of digital signatures. For information on trusted locations, see the chapter *Trusted locations*.

2. Next, determine which packages need to be digitally signed.

    - If you are not using trusted locations, all packages will need to be signed.

- If you are mixing digital signatures and trusted locations, you may choose not to sign (or perhaps not to sign yet) packages that will be delivered only through trusted locations.

On the other hand, it is probably just as easy to standardize the packaging process to sign all packages, and there are no significant penalties for doing so.

3. Make the necessary changes to the administration server infrastructure (detailed at *Administration server infrastructure*).

4. With steps 2 and 3 completed, package creators can start on applying the digital signatures to all required packages. See *Migrating existing packages*.

5. When all required packages have been signed, tested, and redistributed, you can switch on the digital signature functionality for managed devices with a settings update package. See *Preparing a settings update for managed devices*.

The following sections provide more detail on these and related steps.

Administration Server Infrastructure

The following steps are needed to change the administration server implementation:

- The Authenticode signing executable needs to be installed. This file is freely available. Detailed instructions for installation are included in *Procedures*.

- The private key must be established as part of gaining certification from VeriSign. Details are available from VeriSign. Procedures must be set in place for securing this private key between uses, and for duly authorizing its use.

- The software publishing credential obtained from VeriSign must be stored appropriately where accessible from the administration server.

- The administration server must have Internet access to be able to timestamp signed packages.

- You may want to consider whether to separate the distribution locations used for different elements of application packages. For example, decide whether you want to ensure that application catalogs are collected only from trusted locations.

Distribution Hierarchy Changes

No changes are needed to the distribution hierarchy itself for the use of digital signatures.

Managed Device Changes

For digital signatures, there are three changes to the machine hive of the registry, and three matching locks to prevent these changes being overwritten by other methods of setting preferences. Note that this should be the last change implemented, and all packages needing signatures should be repacked and tested before you do this. Details about setting these for new computers and updating them in existing managed devices are included in *Procedures*.

Software Costs

Digital certificates that allow for signing packages are available from suppliers like VeriSign. For VeriSign certificates, signatures can be either 40-bit, for adequate security for most purposes, or 128-bit, when the validity and security of the signature is critical. Certificates are renewed annually for a few hundred (U.S.) dollars (check the Verisign website for current prices). Your enterprise receives two enabling files: a software publishing credential (`.spc` file) and a private key for the enterprise digital signature (`.pvk` file, created during the VeriSign registration process but never shared with VeriSign or any other enterprise). For more information, see *The basics of digital signing*.

The annual billing also means that certificates expire each 12 months and must be renewed. As noted earlier, a timestamp facility means that a digital signature applied under certification can be recognized as valid even after its authorizing certificate has expired. There is no charge for the timestamp service. This is considered very worthwhile. Without it, annual renewal of the certificates also means annual repackaging of all signed applications with the new digital signature, and redistribution.

The enterprise also requires a Microsoft digital signature wizard, an executable (`signcode.exe`) that applies the signature to an appropriate file. There is no charge for this software. It is freely available by download from the Microsoft website (see detailed instructions in *Procedures*) and is also distributed in the .NET framework SDK.

In summary, the direct costs of software and supporting infrastructure are negligible in comparison with the increased security that can result.

RayManageSoft Upgrade Changes

Once digital signing is in use, there are two minor changes needed in the process of rolling out upgrades to the RayManageSofti product itself.

- The upgrade package for RayManageSofti must itself be signed with the approved, valid digital signature.
- The `mgssetup.ini` configuration file must be modified to reflect the registry settings needed to continue use of digital signatures (the factory-supplied copy is not configured for digital signatures). Details of the registry settings are included among the *Procedures*.

With these two small changes, self-updating will proceed as usual.

Migrating Existing Packages

In migrating a working implementation of RayManageSofti from unsigned to signed distribution, the other area of concern is the question of existing packages. These are discussed in more detail here.

Repackaging Requirements

Once the administration server infrastructure is in place (see *Administration server infrastructure*), the requirement is to open each application's project file, select one check box and provide a set of command line parameters, and save the changes. If signing is the only change, you should not change the application name, the package name, or the version number of the package.

The effort involved is literally seconds per package. If there are a large number of packages to switch over, your Raynet consultant can arrange to make the changes to all of your projects at one time. Signing details are covered in *Procedures*.

Redistribution Overheads

After testing, each modified package must be redistributed from the administration server, using the normal distribution process (see the *RMS Software Deployment* for details). If you have your distribution hierarchy set up to use pull distribution, there will be minimum impact. Assuming that there are not concurrent changes to application files, each application will only need three small files distributed: the catalog, the implementation archive, and the digital signature. Switching to signed packages does not cause compressed application files to be redistributed.

Managed Device Overheads

It is not required to specially update all the existing applications and packages installed on managed devices. You can leave existing packages in place. They will be upgraded to signed versions of the packages naturally, as the next upgrade requirement arrives.

Notice that it will not matter whether the upgrade requirement occurs between the distribution of signed packages and turning on managed devices to use signed packages. When the registry settings are not yet turned on for signed packages, any signature in packages is ignored.

For more information on mixing signed and unsigned packages, refer back to *Verifying existing signed and unsigned packages*.

We have seen that the infrastructure changes are not too onerous and the switch-over itself is straight-forward. With that overview, we can now examine individual procedures in specific detail.

Procedures

In this section are step-by-step procedures to set up and use digital signing with RayManageSofti.

Setting Up the Certification Files

The prerequisite for operations with digital signatures is that necessary files have been installed on the same administration server on which packaging is performed. On that server:

1. Download the digital signing utility `codesigningx86.exe` from *http://msdn.microsoft.com/en-us/library/8s9b9yaz.aspx* into a temporary location such as `C:\Temp`.

2. Run the executable. The file is a self-extracting executable that by default extracts to `\inetsdk\bin`.

Preparing a Test Certificate

Digital signatures are validated by reference to a certificate from a certification authority. If you already have a publisher's certificate issued by a supplier like VeriSign, skip this process.

However, if you are still waiting for your official certificate, you can use this method to create a temporary test certificate that allows you to evaluate and debug your processes while you wait. Do not release packages into production based on the test certificate. Subsequently, you can switch to the externally-verified certificate before moving into production. Here is the process for setting up your test certificate.

1. From a command prompt (**Start** > **Run** > **cmd**), enter the following commands.

   ```
   cd inetsdk\bin
   ```

2. Optional: If you would like to check available options:

   ```
   MAKECERT /!
   ```

3. `MAKECERT -r -cy authority -n "CN=Test Root" -sv root.pvk root.cer`

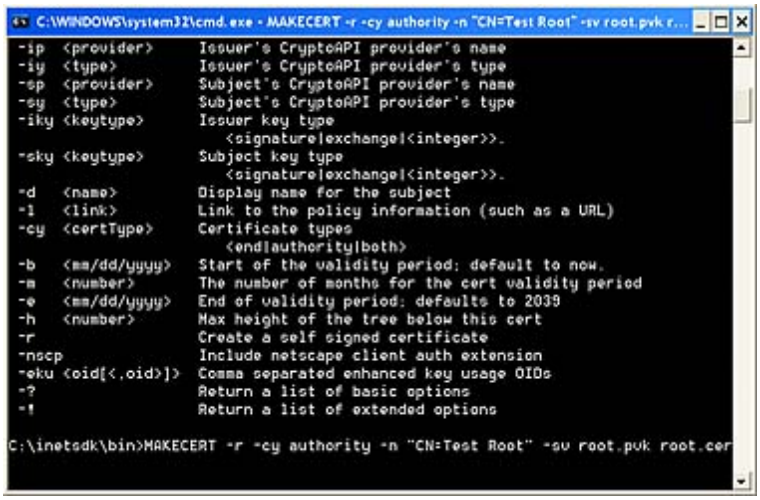

4. You are prompted to create and confirm a password for your private key. Type your password identically in both fields.

5. You must now make use of that password to create the private key that is protected by the password. Enter the same password in the next dialog that is displayed.

6. You have created a root certificate and a root private key. For testing purposes, these will act as the top level certification authority (CA) - a temporary substitute for a supplier such as VeriSign.

   Now we will repeat a very similar process to make a publisher's certificate and key.

7. `MAKECERT -iv root.pvk -ic root.cer -n "CN=Your Name" -sv user.pvk user.cer`

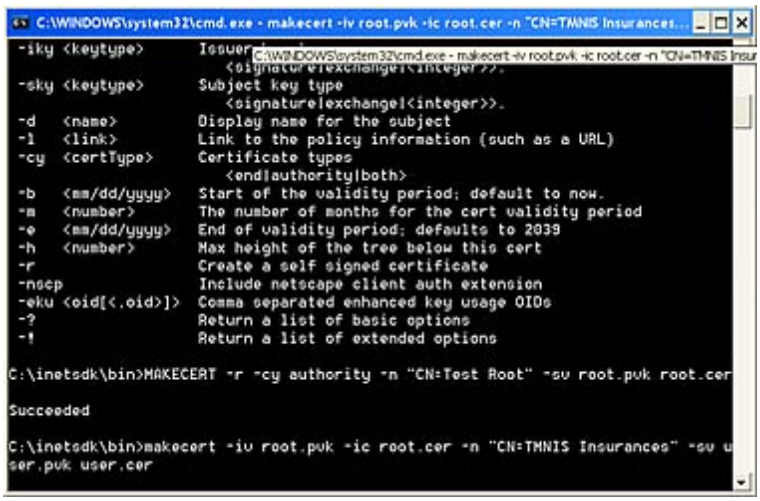

8. Again, you are prompted to create and verify the private key passwords for this publisher (`user.pvk`). Of course, make these different from the CA passwords.

9. Now use the same password to create the publisher's private key.

10. You are now acting as the certification authority, issuing a new publisher's certificate and private key. To authorize these, you must provide the CA signature by providing your original CA password.

The command window reflects your success. You now have two certificates and private keys: one for a CA, and one for a publisher.

11. The final certification step is to create a software publishing credential (`.spc` file) that records the authorization by combining the root and user certificates. (These certificates do not need to be timestamped, as they have a long life span.)

    Enter: `CERT2SPC root.cer user.cer user.spc`

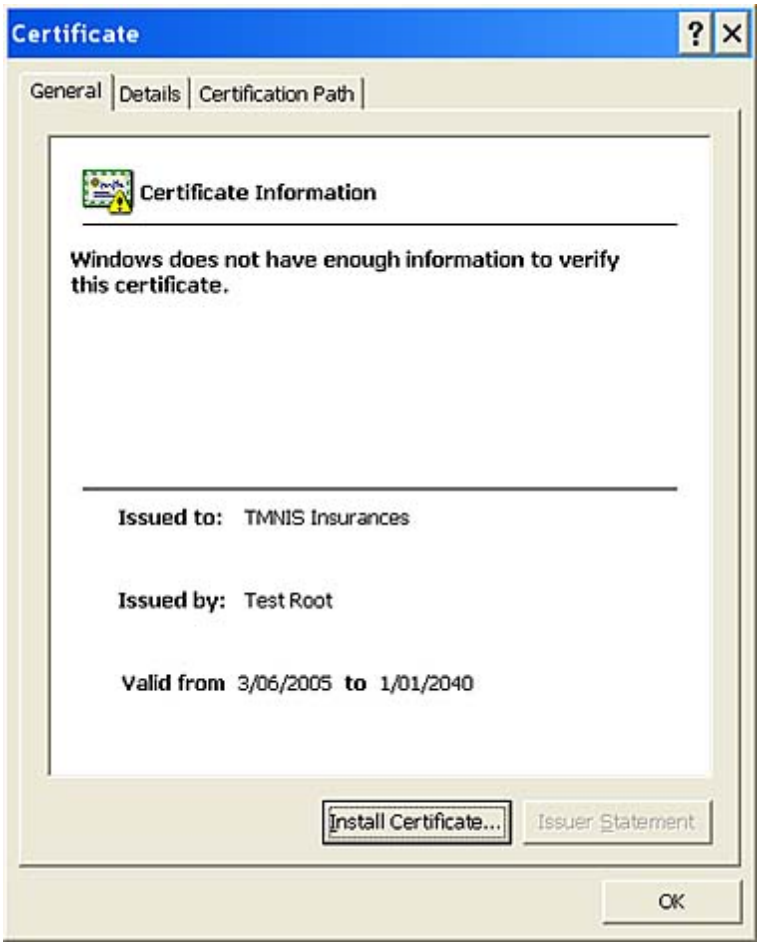

12.The test root certificate (`root.cer`) must now be installed on test managed devices.

For a small pilot group, it is easiest to do this manually:

a. Copy `certmgr.exe` from the same directory as `codesigningx86.exe` (by default, `inetsdk\bin`) onto the target managed device.

b. On the managed device, run `certmgr -add root.cer -s ROOT`

When the time comes to use real certificates, the CA root certificates for VeriSign and Microsoft are already installed on all recent Windows platforms.

Testing Whether a Signature Will Pass Inspection

You can use a command line tool to validate whether the Microsoft technology will return success or failure on any given package. This is an accurate predictor of whether the RayManageSofti installation agent will attempt to install the signed package.

In a command window, enter:

```
chktrust.exe -q <URL of the *.ndc.cab>
```

The message *CertMgr Succeeded* shows success.

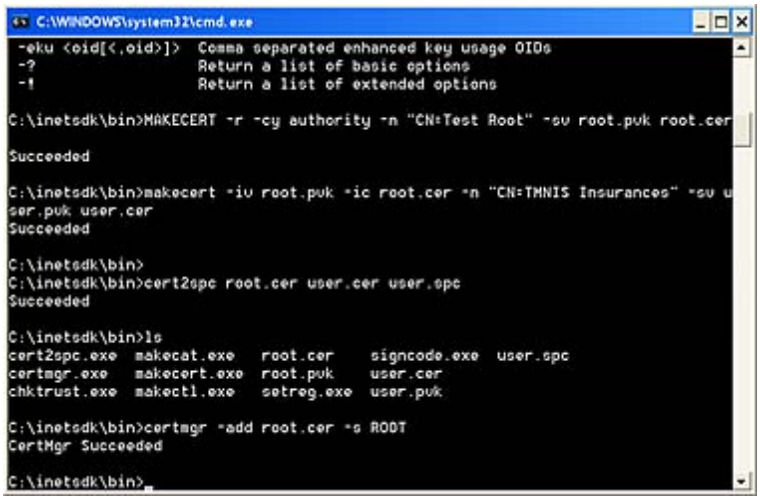

Signing a Package

Be aware:
This process is only supported on Windows managed devices.

To sign a package for distribution, locate the package within the software library, and proceed as follows (see also *Security tab-fields* in the *RMS Software Deployment*):

1. Double-click the **package** within the **software library** list.

2. In the opening edit dialog, click on the **security** section.

3. Scroll down and select the **Digitally sign my Deployment Manager package** check box.

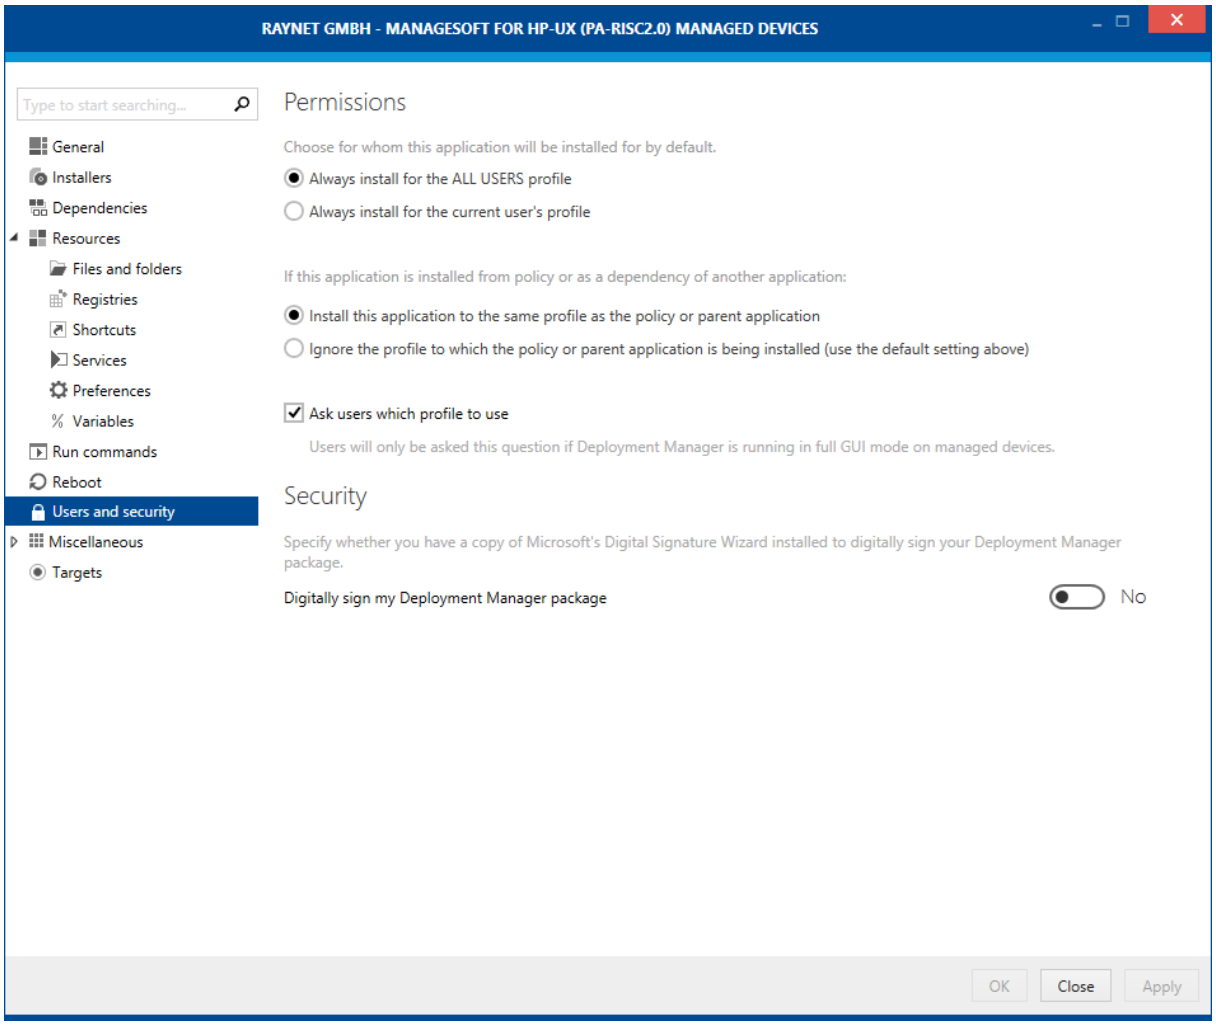

This allows you to specify the path to the digital signature wizard (the `codesigningx86.exe` executable installed earlier).

4. Enter the command line for the digital signature wizard in the form:

```
\full_path_to_executable\codesigningx86.exe
-spc <location of software publishing credential>
-v <location of enterprise private key>
-t <timestamp URL>
```

For example:

```
\inetsdk\bin\codesigningx86.exe
-spc C:\verisign\mycredentials.spc
-v C:\verisign\myprivatekey.pvk
-t http://timestamp.verisign.com/scripts/timstamp.dll
```

Be aware:
"`timstamp.dll`" does not contain the letter 'e'.

5. If all other details for the package are already correct, distribute this package in the usual way.

Preparing a Settings Update for Managed Devices

The use of digital signatures ultimately relies on settings on the managed device. Without the appropriate settings, digital signatures on packages will simply be ignored.

There are three occasions when you may need to set the appropriate preferences on managed devices:

- When initializing a new computer
- When updating the RayManageSofti product on the managed device
- When switching a managed device from using non-signed packages to using signed packages

Be aware:
For details about the preference settings discussed below, see *RMS Preferences for Managed Devices*.

Modifying the .ini File

For the first two of the above cases (for new computers or for an update of RayManageSofti itself), you need to make edits to the factory-supplied `mgssetup.ini` configuration file. Make the edits in the registry settings section, adding the six settings 1-6, as shown here:

```
; Registry settings to be created under
; HKLM\Software\ManageSoft Corp\ManageSoft\Launcher\CurrentVersion [Launcher]
desc0 = LogFile
val0 = C:\launcher.log
; ... etc.
desc1 = ForceValidSignature val1 = True
; ... etc.
desc2 = VerifyCatalogSigned val2 = True
; ... etc.
desc3 = CheckCatalogDigest val3 = True
; ... etc.
desc4 = ForceValidSignatureFxd val4 = True
; ... etc.
desc5 = VerifyCatalogSignedFxd val5 = True
; ... etc.
desc6 = CheckCatalogDigestFxd val6 = True
; ... etc.
```

Be aware:
These settings apply three registry settings to the machine registry hive, and three locks that prevent the same registry settings being over-ridden by the user's registry hive, by command line options, by `.ini` file settings, by project variables, or by any other method.

Updating Registry Settings

If you are switching managed devices from unsigned to signed packages, prepare a new package that contains only the required registry settings, and target it through policy to all managed devices that are making the change. For details, see the *Managed device settings* chapter of the *RMS Software Deployment*.

There are six settings to make. For each one, the location in the registry is `[Registry]\ManageSoft\Launcher\CurrentVersion`

If you have created a new package for these settings, you will first need to cycle through this list of keys to create the registry keys at each level.

Be aware:
"Launcher" is an older name for the installation agent.

Creating Registry Keys

Your goal is to create, under your chosen project, the hierarchy shown below:

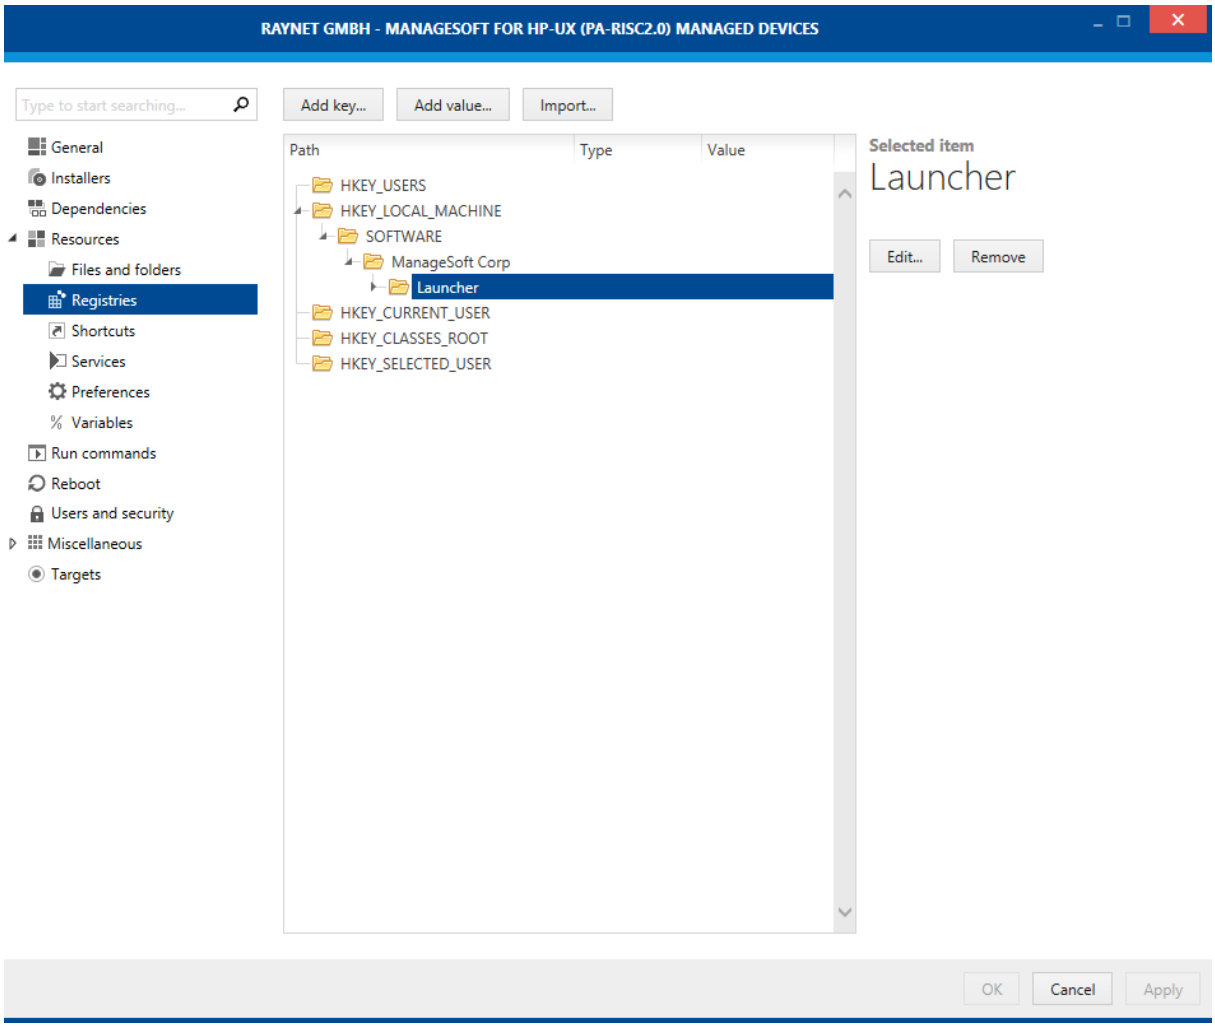

1. Double-click the **package** within the **software library** list.

2. In the opening edit dialog, click on the **resource** section.
3. Open the **registry** dialogue.
4. Navigate to the existing key you want to extend and focus it by clicking on it.
5. Click the **New Key** button above the registry structure area.

   The **New registry key properties** dialog is displayed.

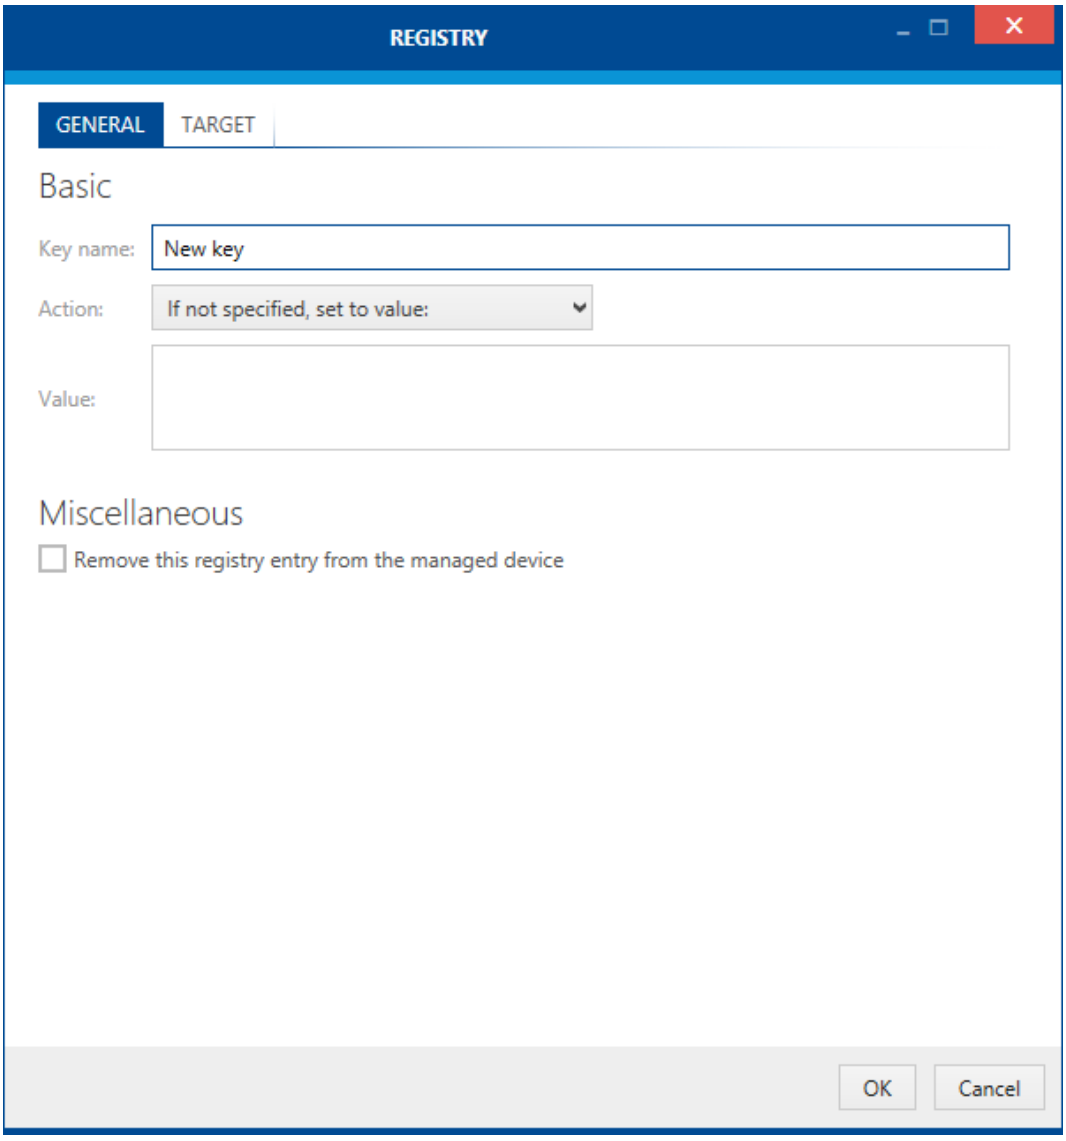

6. In the **Key name** field, accurately type the name of one key (the next one) in the hierarchy. In order, the keys to be entered are:

```
a.SOFTWARE
b.ManageSoft Corp
c.ManageSoft
d.Launcher
e.CurrentVersion
```

Be aware:
"`ManageSoft Corp`" has a space - "`CurrentVersion`" does not.

7. Select the check box for **Remove Registry key from the managed device only if it is empty**. (Accept the remaining settings unchanged.)

8. Click **OK**.

9. Right-click the registry key name that you have most recently created. Repeat steps 3-8 until you have completed setting all the registry keys.

   The final pass should look like the screen capture shown.

Be aware:
Double-check your typing. There can be no validation of the names you create here. A typing error will mean your settings have no effect, as the managed device will be looking elsewhere.

Creating the Registry Entries

Now that your package knows about the relevant registry keys from the managed device, you can enter the six name/value pair entries.

Be aware:
If you are also creating settings for the use of trusted locations, see the chapter on *Trusted locations*, as you may wish to set other relevant registry settings at the same time.

For each entry, proceed as follows:

1. Click the `CurrentVersion` key you just created.

2. Click on the **New Value** button above the registry structure area.

   The **New registry entry properties** dialog is displayed.

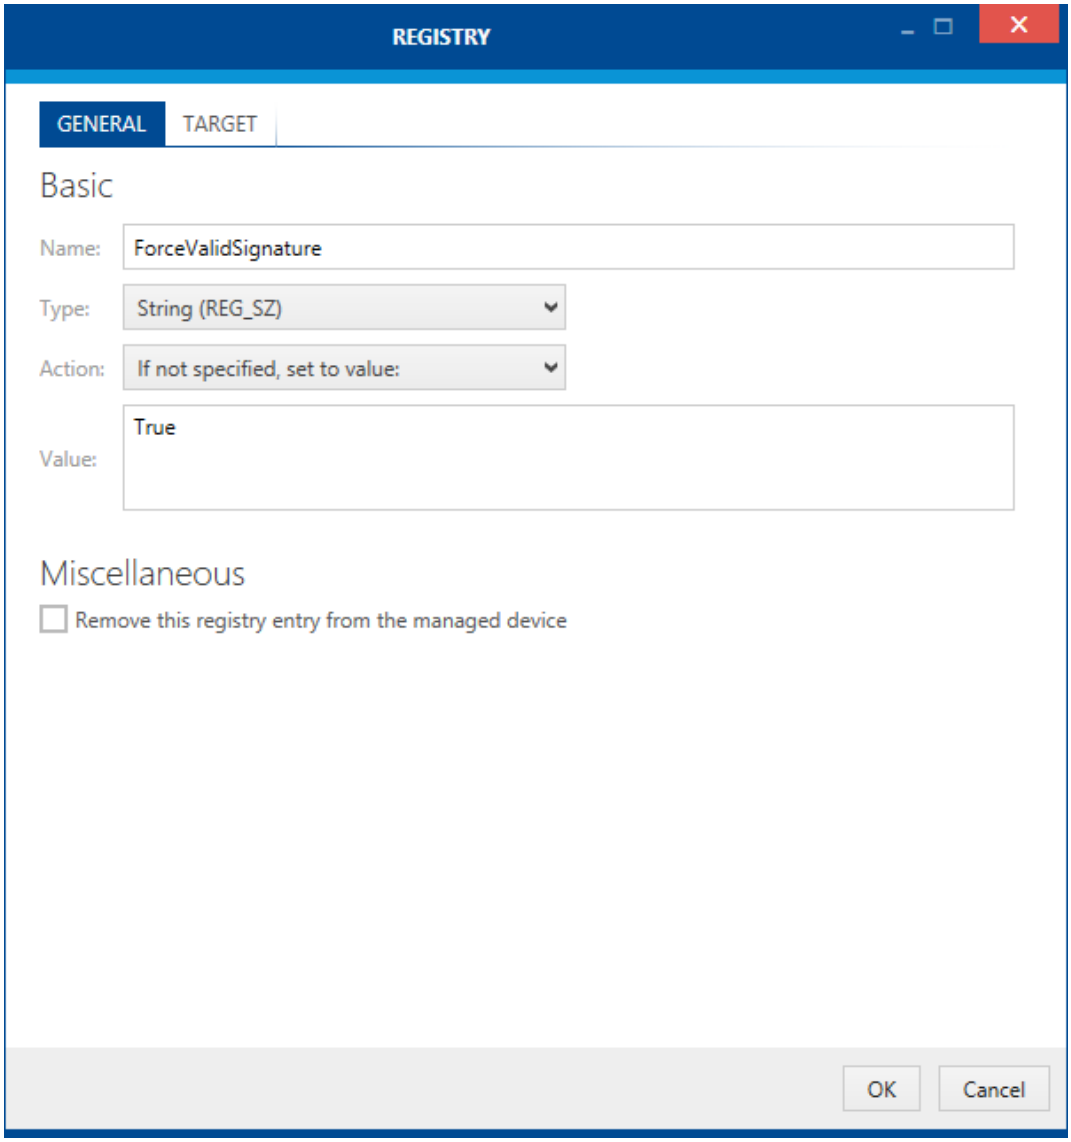

3. Type the name of the next entry into the Name field.

   The six names to be entered sequentially are the same ones used for the `.ini` file, three registry settings and three locks to prevent those settings being over-written:

   ```
   a.ForceValidSignature
   b.VerifyCatalogSigned
   c.CheckCatalogDigest
   d.ForceValidSignatureFxd
   e.VerifyCatalogSignedFxd
   f.CheckCatalogDigestFxd
   ```

   Enter one of these names at a time.

4. The **Type** field must be set to String (`REG_SZ`) each time.

5. The **Action** field must be **Always set to value:** each time.

6. The **Value** field must read **True** each time.

7. Ensure that the check box for **Remove this Registry entry from the managed device** is *not* selected.

8. Click **OK**.

9. Repeat steps 1-8 until you have entered all the registry entries.

Complete the Package

Validate and pack the package of settings in the usual way.

You may use a digital signature with this package. If it happens that the package arrives at any managed device already set to use digital signatures, the package will be applied without incident. For those managed devices not yet set to use digital signing, the signature will be ignored.

When your specification is complete, distribute the package as usual.

Trusted Locations

The term *trusted locations* refers to a second system of trust delegation available within RayManageSofti. (Digital signatures are the first such system of trust delegation.) Specifically, you determine that have strongly secured a distribution location against unauthorized write access. In short, you trust it to contain valid files. You then declare that trust to an appropriate set of managed devices. In effect, you are saying "If a file comes from location X through port Y with protocol Z, you can trust it."

Conversely, you may also declare other download locations to be untrustworthy, so that managed devices do not collect package catalogs from them. While these locations are called "excluded locations", we often use the generic term "trusted locations" to refer to the whole mechanism of identifying both trustworthy and untrustworthy locations.

The purpose of the trusted locations mechanism is to identify those download locations that are within the corporate infrastructure, and efficiently exclude other locations (such as CD-ROM drives) that might otherwise be used to load suspect packages. You can choose whether to turn on this facility for all your managed devices, or for some more limited subset of managed devices that face increased risk of attack.

In this chapter, keep clearly in mind the distinction in RayManageSofti between distribution *servers* and distribution *locations*. A server is a computer, and a location is a file address (like a URL). One server can host many locations. Trust is applied to distribution locations, not to servers as a whole.

Why Use Trusted Locations?

If your enterprise is [believed to be] highly secure and all your end-users are well trusted, you may have no need of trusted locations.

When you need some additional assurance about packages being delivered to your managed devices, trusted locations are a relatively low-cost way to get a known outcome. Trusted locations have zero direct costs (contrasted with the small annual costs of the digital signing solution). They are an efficient way to declare the 'valid' members of your distribution hierarchy and thereby exclude other locations from which end-users may try (wittingly or not) to install packages.

Performance Impacts

As noted in the *Security overview* chapter, the standard use of MD5 digests within RayManageSofti means that if you can guarantee the package catalog, Deployment Manager will guarantee the entire package contents. For that reason, RayManageSofti only checks the trust settings for a server when deciding whether or not to collect the catalog (`.osd` file). In this way, the use of trusted locations does not present a significant performance impact.

Trusted or Excluded Locations

The mechanism for trusted locations includes ways to specify locations that are trusted, and those that are untrusted. When you declare only untrusted locations, all other unlisted locations are trusted. In most enterprises, it is easier over time to maintain a [shorter] list of trusted locations than to maintain a [longer] list of excluded locations. Raynet recommends in general that you use the trusted locations option. This is also the safer option, in that unknown locations are by default untrusted until you declare them to be safe.

Identifying Trusted Locations

It is evident that managed devices have to be informed about trusted and excluded locations. The information is stored in a database (in the sense of structured data records, but not an SQL database) that can be delivered in one of two ways. We will first examine the data structure, and then the delivery methods.

Data Structure for a Trusted Location

Whichever delivery method is used, the data structure for describing a trusted (or excluded) location is the same.

- The record is identified by a unique key that identifies the location. You should use descriptive names that will be meaningful during later maintenance. A corporate naming convention can help to ensure that the key remains both unique and meaningful.

- Each record then includes the following attributes. All attributes are mandatory (the values cannot be an empty string). Note that the wildcards `*` (asterisk, replacing any number of characters) and `?` (question mark, replacing a single character) are permitted within most attribute values:

Protocol	One of the case-sensitive string literals
	
	`http`
	`https`
	`ftp`
	`file`
	`*`
	
	This defines the protocol for data transfer to the managed device. In this case the wildcard `*` means “any protocol”, and wildcards are not otherwise permitted in the string literals.

Host	The name or IP address of the distribution server.
	
	Examples:
	`*.yourhost.com`
	`123.45.67.*`

Port	The port number for trusted transfers. Typical port numbers are `80` for HTTP, `443` for HTTPS, and `21` for FTP. You may also choose to specify special ports, depending on the structure of your server and its security. If you wish to specify “any port”, use the wildcard `*` in place of a port number.
Directory	The full directory path to the trusted location, relative to the root of the server. • Use forward slashes (not backslashes) as the separators. • Include the leading / for the server root. • To restrict trust to exactly this directory, leave the last slash off. Example: `/corporate/trusted` • To include all the subdirectories of this trusted location in the trust declaration, add a trailing wildcard. Example: `/corporate/trusted/*` • This path is a string literal, and you cannot use project variables or registry key values in it.

Data Storage for Trusted Locations

There are two delivery methods to get the trusted location dataset to Windows managed devices: the registry; and a global configuration file.

UNIX and Macintosh managed devices read trusted locations from their `$(CommonAppDataFolder)/etc/download.xconf` (by default `/var/opt/managesoft/etc/download.xconf`) configuration files. Refer to the *RMS Preferences for Managed Devices* for details.

In Registry Keys

Data that define trusted locations can be stored under either the current user hive or the machine hive of the Windows registry. Unless you have specific reasons for setting trusted locations for particular users, Raynet recommends using the machine hive for the following reasons:

- Recalling the evaluation order of RayManageSofti preference settings, a user preference will override any machine settings. This is useful for those times when you do want a specific override for a small set of end-users, provided that you have used the machine hive for normal corporate settings. (For more on the evaluation order of preference settings, see *RMS Preferences for Managed Devices*.)

- Machine settings require administrator privileges to change. User settings can (with a few exceptions under unusual conditions) be altered by the user, and it's rare that you want users to adjust trusted location settings.

For this reason, this chapter documents the settings for the machine hive.

Be aware:
The ability to store the trusted locations *dataset* in the registry is quite separate from the set of controls for trusted locations that are also available as registry settings. The *control* settings are documented in *Controlling the use of trusted locations*.

Trusted Locations

To identify a trusted location, create a key for the location under the following registry key:

```
[Registry]\ManageSoft\Launcher\CurrentVersion\TrustedLocations
```

`TrustedLocations` will not exist by default, and you may have to declare that key first. An example of a completed key setting with a descriptive name is:

```
[Registry]\ManageSoft\Launcher\CurrentVersion\TrustedLocations\CentralTrust01
```

For each excluded location, continue with *Completing the data record* below.

Excluded Locations

To identify an excluded location, create a key for the location under the following registry key:

```
[Registry]\ManageSoft\Launcher\CurrentVersion\ExcludedLocations
```

ExcludedLocations will not exist by default, and you may have to declare that key first. An example of a completed key setting using a descriptive name is:

```
[Registry]\ManageSoft\Launcher\CurrentVersion\ExcludedLocations\InfectedServer01
```

For each excluded location, continue with the next subsection to complete the data record.

Completing the Data Record

Once you have established the appropriate key for the trusted (or excluded) location, complete the data record by adding the registry value names/strings for the attributes described in *Data structure for a trusted location*. For example, for the Host entry create a string-typed value named `Host` with a value of the appropriate form, such as `123.45.67.*`.

Repeat the process for all required trusted or excluded locations.

If you are making these settings on a test computer, it is adequate to make the settings manually in the registry. If you are preparing settings for application to a number of managed devices, there are two ways that you can deliver the settings:

- Provided that you have Active Directory implemented throughout your enterprise, you can use the Group Policy to propagate registry changes through ADM templates to your managed devices. See the Active Directory documentation for more information.

- Whether or not you have Active Directory implemented throughout your enterprise, you can prepare a RayManageSofti package of registry settings. The process is documented in the *RMS Packaging*, particularly the section on registry entries, and is repeated in this manual in Updating registry settings. This package can then be distributed to appropriate targets through policy in the usual way.

We have examined the first method of delivering the dataset of trusted or excluded locations, through registry settings. We now return to the second method, the global configuration file.

In the Global Configuration File

The global configuration file is another way to distribute RayManageSofti preference settings.

Choosing whether to Use a Global Configuration File

A global configuration file can be used separately, or you can use it in conjunction with registry settings. In the latter case, the order in which preference settings are applied means that you can use the global configuration file to set the baseline set of preferences, and then override specific settings for a subset of computers with registry settings. For more background, see the *RMS Preferences for Managed Devices*.

You should also consider the question of scalability of a global configuration file contrasted with the use of registry settings. Once registry settings are delivered to the individual managed device, there is no network access needed to read the settings. In contrast, a global configuration file will be accessed by a large number of managed devices at frequent intervals. This can lead to scalability problems. You can partially address these by using DFS aliasing for shares. This would allow you to use a system-wide alias for the configuration file that pointed to (say) regional copies that managed devices could access at local LAN speeds.

The Format of a Global Configuration File

The RayManageSofti global configuration file is in the standard Windows `.ini` file format:

- All lines set flush left with no leading spaces
- Comments starting with a semicolon
- Section titles enclosed in square brackets
- Key entries in the format `keyname = keyvalue` and repeated as many times as necessary for all the keys within a section.

Within this format, you need to insert one section for each trusted location, and one section for each excluded location.

Trusted Locations

For each trusted location, start a separate section and include the unique key for the location as part of the section name. The section label is **TrustLocation**. For example, a completed section heading for the example used earlier (in registry settings) would be:

```
[TrustLocation CentralTrust01]
```

Excluded Locations

For each excluded location, start a separate section and include the unique key for the location as part of the section name. The section label is **ExcludeLocation**. For example, a completed section heading for the example used earlier (in registry settings) would be:

```
[ExcludeLocation InfectedServer01]
```

Completing the Data Record

Once you have established the appropriate section heading for the trusted (or excluded) location, complete the data record by adding the key names/value pairs for the attributes described in *Data structure for a trusted location*. For example, the completed value entry for the Host value might look like this: `Host=123.45.67.*`

Repeat the process for all required trusted or excluded locations.

Example

An extract from a global configuration file with the settings for trusted FTP locations might therefore be as follows:

```
; RayManageSofti trusted locations configuration settings
;
; This file lists trusted and untrusted distribution locations
; as a baseline for all managed devices. Local changes may be
; effected through additional registry settings.
;
[TrustLocation CorporateFTP01] Protocol=ftp
Host=123.45.67.*
Port=21 Directory=/ftp/corporate/trusted/*

[Launcher]
...
```

Be aware:
The settings under the `Launcher` section are the controls for trusted locations. These are documented in *Controlling the use of trusted locations*. "Launcher" is an old name for the RayManageSofti installation agent.

Location of the Global Configuration File

Once completed, a configuration file can be located wherever it suits your network infrastructure, with the following points taken into account:

- Give due thought to securing the copy of any configuration file that lists trusted locations. Exposed and unprotected, it presents a security weakness.
- Consider the scaling implications and whether to use aliases. For more information refer back to *Choosing whether to use a global configuration file.*
- Managed devices need to know the location of your global configuration file. You declare the entire path and filename in the preference setting **GlobalConfigSource** (normally set in the machine hive of the registry). See the *RMS Preferences for Managed Devices* for more details.

Resolving Trusted and Excluded Locations

RayManageSofti copes with incomplete sets of data on trusted (and excluded) locations in the following ways:

- If trusted locations are turned off (the **UseTrustDatabase** preference is set to `False`, as discussed in *Controlling the use of trusted locations*), the dataset is entirely ignored.
- With trusted locations turned on and no locations identified under the **TrustedLocations** key, all locations are trusted except those under the **ExcludedLocations** key.
- With trusted locations turned on and one or more locations identified under the **TrustedLocations** key, only locations listed under the **TrustedLocations** key are trusted, unless they also appear under the **ExcludedLocations** key.
- Locations that appear under both the **TrustedLocations** key and the **ExcludedLocations** key are not trusted.

Raynet recommends the use of (at least some) trusted locations in preference to relying exclusively on excluded locations. It is easier to declare your list of corporate locations (thus disallowing everything else) than it is to think of all the possible locations you might want to exclude.

We have now completed the data structure, the dataset storage, the data delivery methods, and data resolution rules for recording trusted locations. With the data in place, we can turn to controlling their use.

Controlling the Use of Trusted Locations

There is a set of several preference settings that determine the use of trusted locations and whether (and how) they interact with the use of digital signatures. Each of the preference settings is documented in detail in the *RMS Preferences for Managed Devices*. In this section, we will group related preference settings together to explain their interactions and operation.

As noted before, there are several places that you could make the following control settings. For the reasons listed in *Data storage for trusted locations*, the Raynet recommendation is that you make these control settings in `[Registry]\ManageSoft\Launcher\CurrentVersion`.

If you decide to make control settings in a global configuration file, the section is called `[Launcher]`.

Turning on Trusted Locations

A managed device will use trusted locations and avoid excluded locations only if this preference setting is `true`.

Value name	Setting	Comment
`UseTrustDatabase`	`True`	The fundamental on/off switch for the use of trusted locations.

Securing Trusted Location Settings

Recall that the order of resolving preference settings from various sources means that:

- A setting in a global configuration file will be overridden by any registry setting
- A machine registry setting will be overridden by a user registry setting
- All user settings may be overridden by (and some cases overwritten by) a command line option

For more information about this, and about the standard locking mechanisms that can prevent this cycle, see the *RMS Preferences for Managed Devices*.

In the specific case of trusted locations, there are two locking mechanisms:

Value name	Setting	Comment
`UseTrustDatabaseFxd`	`True`	Prevents overriding the basic on/off switch for trusted locations use.
`TrustDatabaseFxd`	`True`	Requires that a user account needs administrator privileges before changing *any* dataset entries for trusted or excluded locations.

Reading from the Global Configuration File

A managed device will first read a global configuration file for baseline trusted (and excluded) location settings if the following preferences are set:

Value name	Setting	Comment
`UseTrustDatabaseFxd`	`True`	Prevents overriding the basic on/off switch for trusted locations use.
`GlobalConfigSource`	`<loc>/<file>`	Setting the path and filename to the global configuration file turns on its use.

Be aware:
Settings in the global configuration file will be extended or overridden by any settings in higher-priority preferences declarations, including registry settings.

Mixing Digital Signatures with Trusted Locations

There are a number of settings used in various combinations to control the interaction of trusted locations with digital signatures. For more on digital signatures, see the chapter *Digital signing*. The Raynet recommendation is the fourth of the following options.

To Use Trusted Locations Only

To use trusted locations only, with checking of digital signatures turned off, ensure that the following registry entries (listed alphabetically) are set:

Value name	Setting	Comment
CheckCatalogDigest	True	See documentation on digital signing
CheckCatalogDigestFxd	True	Lock on previous item
ForceValidSignature	N/A	Not checked
ForceValidSignatureFxd	N/A	Don't care
UseTrustDatabase	True	Turns on use of trusted locations
UseTrustDatabaseFxd	True	Lock on previous item
VerifyCatalogSigned	False	Turns off digital signing
VerifyCatalogSignedFxd	True	Lock on previous item
VerifyTrustOrSign	False	Disallows the logical OR of digital signing or trusted locations
VerifyTrustOrSignFxd	True	Lock on previous item

To Use Digital Signatures Only

To use digital signatures only, without checking whether files are delivered from a trusted location, ensure that the following registry entries are set:

Value name	Setting	Comment
CheckCatalogDigest	True	See documentation on digital signing
CheckCatalogDigestFxd	True	Lock on previous item
ForceValidSignature	True	See documentation on digital signing
ForceValidSignatureFxd	True	Lock on previous item
UseTrustDatabase	False	Turns on use of trusted locations
UseTrustDatabaseFxd	True	Lock on previous item

Value name	Setting	Comment
VerifyCatalogSigned	True	Turns on digital signing
VerifyCatalogSignedFxd	True	Lock on previous item
VerifyTrustOrSign	False	Disallows the logical OR of digital signing or trusted locations
VerifyTrustOrSignFxd	True	Lock on previous item

To Use Digital Signatures and Trusted Locations

To combine the use of digital signatures and trusted locations such that every file must be validated by digital signature and delivered from a trusted location (logical AND), ensure that the following registry entries are set:

Value name	Setting	Comment
CheckCatalogDigest	True	See documentation on digital signing
CheckCatalogDigestFxd	True	Lock on previous item
ForceValidSignature	True	See documentation on digital signing
ForceValidSignatureFxd	True	Lock on previous item
UseTrustDatabase	True	Turns on use of trusted locations
UseTrustDatabaseFxd	True	Lock on previous item
VerifyCatalogSigned	True	Turns on digital signing
VerifyCatalogSignedFxd	True	Lock on previous item
VerifyTrustOrSign	False	Disallows the logical OR of digital signing or trusted locations
VerifyTrustOrSignFxd	True	Lock on previous item

To Use Digital Signatures or Trusted Locations

To combine the use of digital signatures and trusted locations (as recommended in *Digital signing*) such that

- Packages with valid digital signatures will be accepted from any location
- Packages without approved digital signatures will be accepted from trusted locations

(that is, a logical OR), ensure that the following registry entries are set:

Value name	Setting	Comment
CheckCatalogDigest	True	See documentation on digital signing
CheckCatalogDigestFxd	True	Lock on previous item
ForceValidSignature	True	See documentation on digital signing
ForceValidSignatureFxd	True	Lock on previous item

Value name	Setting	Comment
UseTrustDatabase	N/A	Not checked
UseTrustDatabaseFxd	N/A	Don't care
VerifyCatalogSigned	N/A	Not checked
VerifyCatalogSignedFxd	N/A	Don't care
VerifyTrustOrSign	True	Allows use of either digital signatures or trusted locations on package, and causes the UseTrustDatabase and VerifyCatalogSigned switches to be ignored
VerifyTrustOrSignFxd	True	Lock on previous item

User Groups

This chapter contains a summary of all the user account types needed in a fully operational system.

After reading this chapter, you should also refer to specific rights configuration information for:

- *Configuring administration server rights*
- *Configuring distribution server rights*
- *Configuring managed device rights*

User Account Categories

To understand the default RayManageSofti rights configuration and what configurations are possible, it is necessary to understand the basic categories of users that interact with various features of RayManageSofti.

Users can be members of many of these categories at the same time. Several of the roles outlined will commonly be performed by the same person, but for clarity it is necessary to elaborate all the possible roles by which RayManageSofti can be used.

These user categories are used to detail the permissions and/or users that apply to certain elements of RayManageSofti. Standard system user names (for example **Everyone** from Microsoft Windows) are also used where appropriate. In this case they will have the same meaning as in the system they originate from. There are also some users who should in general have access to everything. For example the `SYSTEM` user on a Windows NT system should have access to everything by default. These users are not specified throughout this document on each specific directory.

User roles are typically performed by people. However, there are also several types of users within the RayManageSofti infrastructure that only exist as accounts for running services, scheduled tasks, and other background tasks requiring no user interaction.

Security Policies

There are also standard security policies applied to the `Program Files` and `C:\Users` (on Vista and later) and `Documents and Settings` (on earlier Windows platforms) folders which are not affected by the default RayManageSofti rights configuration. These are not specified explicitly in this document, but need to be remembered when applying settings.

Console User Roles

Access to the RayManageSofti console is, by default, denied to all users except those in the **MGS Administrators** and **MGS Reporting Users** security groups.

You can modify this access by assigning (or denying) rights to other security groups. Details are provided later in the chapter *Role-based security*.

Accounts Used by People

Users fall broadly into the following categories:

Deployment Manager Administrator

Users who are responsible for maintaining the RayManageSofti infrastructure. Typical tasks they perform include:

- Defining and maintaining the distribution hierarchy
- Reviewing reports for errors
- Applying desired security configurations

Users in this category should be members of the **MGS Administrators** security group, which is created during installation of Deployment Manager on the administration server. Typically, members of other user categories that access the administration server or distribution servers should also be members of this security group. Membership of this group can be edited using the Active Directory Users and Computers tool. Users in this security group would normally also be members of **Domain Admins** in order to access Active Directory.

Deployment Manager Reports Viewer

Users who need to view the contents of the Deployment Manager reports. Typical tasks they perform include monitoring the state of the RayManageSofti infrastructure for any problems that occur. Users in this category should be members of the **MGS Reporting Users** security group, which is created during installation of Deployment Manager. Membership of this group can be edited using the Active Directory Users and Computers tool.

Deployment Manager Managed Device User

Users who are under the control of RayManageSofti for software management. Typical tasks they perform include installation and removal of packages (automatically through the installation agent, or manually using the selection agent). This will typically be an organization's **Domain Users**, however in certain situations it may be a subset of **Domain Users**. For example, you may wish a particular distribution location to only be available to some Domain Users.

Accounts Used by RayManageSofti

These types of users only exist as accounts for:

- Running services
- Running scheduled tasks
- Running IIS web applications

These accounts require no user interaction, and can be broken down into the categories listed below.

Be aware:
The installation of Deployment Manager on administration servers establishes a single username for all scheduled tasks (see *User name for scheduled tasks on the administration server*). If you wish to

provide separate user accounts for different scheduled tasks, you may edit the username within the definition of each task.

Administration Server Security Groups

During installation of RayManageSofti for administration servers, a number of Active Directory user groups are created (and in some cases given default memberships), and security configuration controls are set.

The user groups are

- MGS Administrators
- MGS Data Modifiers
- MGS Data Readers
- MGS Distributors
- MGS Field Technicians Tool
- MGS Report Users

See *Default user groups* for further details.

Deployment Manager Application Usage Importer Task Users

User accounts under which the application usage importer scheduled tasks run on administration servers.

Deployment Manager Connection Agent Service Users

User accounts under which the connection agent service runs on distribution servers.

Deployment Manager Discovery Importer Task Users

User accounts under which the data importer scheduled tasks run to process discovery data on administration servers.

Deployment Manager Distribution Server Status Importer Task Users

User accounts on administration servers under which scheduled tasks to process distribution server status data run.

Deployment Manager Distribution Task Users

User accounts under which scheduled tasks for the distribution agent run on the administration server and distribution servers. Users in this category should be members of the **MGS Distributors** security group which is created during installation of Deployment Manager. The membership of this group can be edited using the Active Directory Users and Computers tool.

Be aware:
Users need to be manually added to the **MGS Distributors** security group. This is not done automatically during installation, as it requires knowledge of the distribution hierarchy, which is not known at this time.

Also be aware:
After adding a computer account to the **MGS Distributors** security group, you must reboot that computer to enable it to recognize the changes to security group membership.

Also be aware:
Typically, users from parent domains in a multi-domain environment do not have access to the job queue in other domains. If this access is required (for example, you want the distribution agent scheduled task to be run by a user in the parent domain), you must manually configure this.

Deployment Manager IIS User

User account under which IIS web applications run on administration servers.

Deployment Manager Inventory Processor Task Users

User accounts under which the inventory agent scheduled tasks run on administration servers.

Deployment Manager Managed Device Event Importer Task Users

User accounts on administration servers under which scheduled tasks for the managed device event importer run.

Deployment Manager managed device task users

User accounts under which the scheduled tasks run on a managed device.

Deployment Manager Merged Policy Generator Task Users

User accounts under which the merged policy generation is run on the administration server. (Merged policy is always run on the administration server to populate the RayManageSofti database, and may also be used in server-side policy merging.) Users in this category must have domain administrator rights in order to access the security descriptors in Active Directory. Note that these rights must be conferred through membership in the Domain Admins group, and not through the Domain Administrators group (a Microsoft built-in group that appears to exist in only some installations of Active Directory).

For more details about rights configuration for merged policy generation, refer to *Application usage importer.*

Deployment Manager Remote Execution Task Users

User accounts under which the remote execution server agent scheduled tasks run on administration servers and distribution servers.

Deployment Manager Remote Execution Task Importer Task Users

User accounts under which the remote execution task importer scheduled tasks run on administration servers and distribution servers.

Deployment Manager Uploader Scheduled Task Users

User accounts under which the upload agent scheduled tasks run on distribution servers and managed devices.

Deployment Manager Wake on LAN Task Users

There are two areas where user accounts are required for Wake on LAN:

1. **Generate Deployment Manager Wake on LAN jobs** (`mgswolauto.exe`) is the scheduled task that runs on the RayManageSofti administration server.

   Because this task requires access to the RayManageSofti database, users in this category need to be members of the **MGS Administrators** security group which is created during installation of Deployment Manager.

2. The **Wake On LAN wizard** requires access to Active Directory.

   Users in this category must have domain administrator rights in order to access the security descriptors in Active Directory. Note that these rights must be conferred through membership in the **Domain Admins** group, and not through the **Domain Administrators** group (a Microsoft built-in group that appears to exist in only some installations of Active Directory).

Configuring Administration Server Rights

This chapter describes the default rights and permissions configuration for administration servers, as well as the minimum permissions settings that you can implement for RayManageSofti to function as a whole.

It also describes interactions between various Deployment Manager features and agents, and the security permissions required to enable these interactions.

Rights and permissions settings include username/password logon authentication (for example, for database access or access to a FTP server) and file system permissions controlling access to files, directories, and other file system elements such as shared drives.

The layout of this information is described in *Configuring RayManageSofti rights*.

Be aware:
The same name may be used for security settings on the administration server, distribution servers, and managed devices. Where this occurs, the abbreviation AS, DS, or MD appears after the name of the setting.

RayManageSofti Security Model

This section describes the overall security model implemented by RayManageSofti. Read this section before attempting to analyze the data flow of the system, presented in the next section.

The diagram below depicts the two security types applied across the RayManageSofti system.

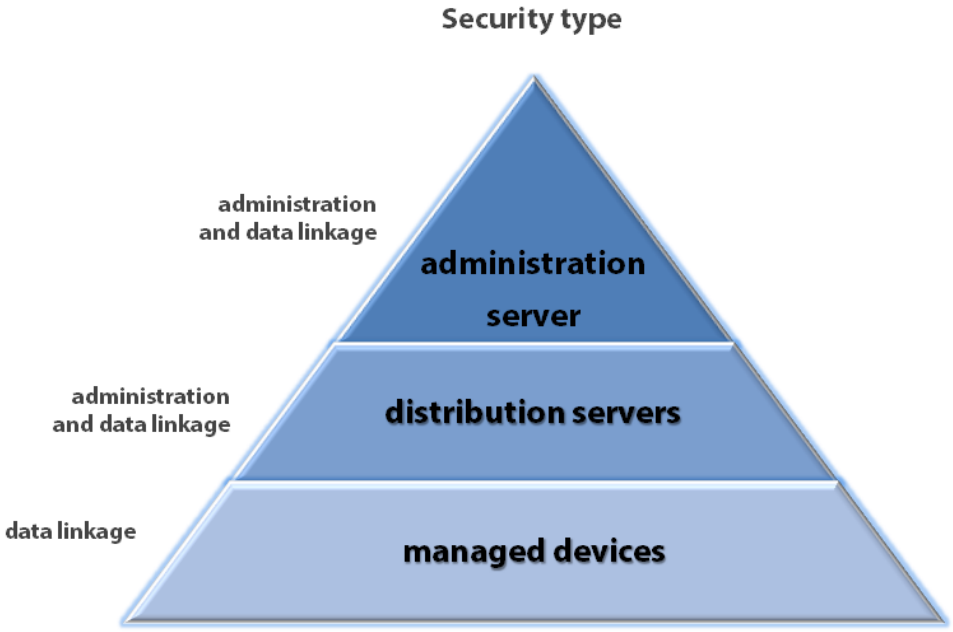

Types of security applied to the three main segments in the RayManageSofti system.

Administration security relates to configuration, data access, and processing of data that is essential to maintaining the smooth running of the system. This includes security imposed when:

- Processing logs
- Modifying the distribution hierarchy
- Distributing jobs and packages
- Modifying package assignments in policies
- Creating schedules, packages, policies, and client settings
- Reconciling the RayManageSofti database with Active Directory
- Updating the RayManageSofti database with new policy assignment data
- Generating reports
- Maintaining license allocations

Data linkage security relates to the distribution and reporting chain, from the administration server to the managed device and back again. This includes:

- Permissions applied to distribution locations
- Permissions applied to reporting locations
- Permissions applied to the package stage

General Access Considerations

This chapter covers all points of entry to and exit from RayManageSofti on administration servers. A point of entry or exit is anywhere data goes out of the direct control of a RayManageSofti agent, including data written to a permanent or temporary store or transferred between RayManageSofti agents (for example, using a network connection). You can apply security settings to each of these points of entry or exit to control access and modification rights.

All the share names described are used by the current release of RayManageSofti. However the permissions settings discussed also apply to any backwards-compatible share names in use, for example **nDUpdate** instead of **ManageSoftDL**.

Default Rights Configuration

The default RayManageSofti rights configuration (resulting from a standard installation) aims to provide a fully functioning system with adequate security, as easily as possible.

Organizations may want to further restrict access of this sort. As security settings and permissions are organization-specific, automatic configuration of all of these items is not possible at installation time. They are also typically very dynamic, so maintenance of these settings is also very important.

Rights settings do not just include permissions on files and valid user accounts. Consideration should also be given to who has access to the RayManageSofti administration server, distribution servers, and distribution locations. This includes physical access to the computers (on the console) and also remote access. All of these elements need to be included together in any successful security configuration.

Default Security in Windows Environments

To provide a fully functioning system that can operate without major rights modifications, read access is set at a low level for most elements. For example, most data files/locations in RayManageSofti are readable by everyone. Full control, however, is more tightly restricted on initial installation. For example, the RayManageSofti reporting features are only made available to users in the **MGS Administrators** and **MGS Reporting Users** security groups.

Considerations for Your Rights Configuration

Important areas to consider when implementing the rights configuration for an organization's RayManageSofti infrastructure include the following:

Who Has Access to the Administration Server?

Any user who has access to the administration server can perform actions such as applying applications to policies, which can in turn have large impacts on managed devices in an enterprise. Therefore access to the administration server should be restricted as tightly as possible.

For security measures you can take to restrict access to reports, see the *Role-based security* chapter.

Who Has Access to the RayManageSofti Database?

The Deployment Manager database is the central store for most of the information concerning an organization's Deployment Manager infrastructure. The integrity of this information is essential to the successful functioning of Deployment Manager. Therefore access to the database should be restricted as tightly as possible.

Who Has Access to Critical RayManageSofti Configuration / Task Data?

Both the hierarchy files and the job queue are key elements to the security of the distribution system of RayManageSofti. Therefore access to these files should be restricted as tightly as possible.

Who Has Access to Distributed Applications, Policies, Schedules, and Managed Device Settings?

Care should be taken over which users can access specific applications, policies, schedules, and managed device settings. For example, an enterprise may not want its main financial program accessible to every user. Permissions should be set appropriately on distribution servers and distribution locations to restrict user access to the packages on them. Access to registry settings should be carefully controlled, and many preferences set in the registry should be locked (for more information, see *RMS Preferences for Managed Devices*).

NTFS Setting Considerations

NTFS file permissions on the underlying directories for NTLM shares also apply through NTLM share access. File server settings should therefore be considered in conjunction with the appropriate NTFS settings.

Also, NTFS file permissions on the underlying directories for the IIS virtual directories also apply through FTP, HTTP, and HTTPS access. FTP, HTTP, and HTTPS server settings should therefore be considered in conjunction

with the appropriate NTFS settings.

This means that for "anonymous" FTP or HTTP access, the user configured in IIS as the anonymous user must also be granted access through the NTFS permissions.

Configuring RayManageSofti Rights

This chapter and later chapters describe security considerations for:

- *Configuring administration server rights*
- *Configuring distribution server rights*
- *Configuring managed device rights*

When configuring rights, it may be convenient to consider security issues within a number of broad topics:

- RayManageSofti features
- Associated data files/locations
- Communications
- Databases
- Servers

Each of the rights configuration chapters describes these topics using a common layout, as described below.

About Features

A feature is essentially something that enables a user to perform a task. It does not necessarily have a direct mapping to an executable or program on its own. It may be a collection of programs or a part of one program. A feature will make use of various data files/locations, communications, databases, and servers to perform its task.

Features are described in this guide in the following table format:

Aspect	Description
Description	A brief description of what the feature is and does
Program	The element that implements this feature. This may be an executable (or group of executables) or an element of the user interface
Users	List of users that can run this feature
Data files/locations	List of data files/locations this feature interacts with
Communications	List of any communications this feature interacts with
Databases	List of any databases this feature interacts with
Servers	List of any servers this feature interacts with

About Data Files / Locations

Data files / locations are essentially storage points for data that reside on a file system. There may be collections of files or one specific file. Files may also be permanently present or only present while in use.

Each data file / location is described using the following table format:

Aspect	Description
Description	A brief description of the data file/location, including its purpose and contents
Default location	The default location for this data file/location
Settings location	The location where the permission settings for this data file/location can be controlled
Default permissions	The default permissions for this data file/location. Access rights are specified using the standard Microsoft Windows NTFS permission terms (Full control, Modify, Read & execute, List folder contents, Read and Write)
Minimum permissions	The tightest set of permissions that can be applied to this data file/location that still allow the product to function. Access rights are specified using the standard Microsoft Windows NTFS permission terms (Full control, Modify, Read & execute, List folder contents, Read and Write)

About Communications

Communications are any network connections that may be established between two features or a feature and a server.

After a brief description of the communication, including its purpose, details are provided the following table format:

Aspect	Description
Initiator	The feature that initiates the communication
Acceptor	The feature that accepts the communication
Protocol	The protocol used by the communication, including details of any authentication performed

Ports

To assist with firewall and network routing configuration, this section documents the TCP and UDP ports used in each of the network communications required by the product features. Ports are described for each computer, including both incoming and outgoing communications. This means that any given communication will appear twice in this listing, once for the sending computer and once for the receiving computer. The details are arranged in the following 'card' format:

Aspect	Description
Source/ Destination	Identifies the partner (remote) device or computer in this communication

Aspect	Description
Port	The port number, which is a local port for incoming communications, and the port on the receiving device for outgoing communications
App Protocol	The application protocol (or sometimes, the standard use) of this port number. For example, port 80 is the default for Hypertext Transfer Protocol (HTTP)
Config'able?	Whether or not the port number can be configured. If this column shows N, you must use only the specified port number. If this column shows Y, you may reconfigure the port number for this communication
N/w Protocol	The network protocol used for the communication
Purpose	A summary of the reason for the communication
Notes	Optionally, additional comments

About Databases

Databases are any storage of data used by the product which is not directly stored in the file system. They may be true databases (for example the RayManageSofti database) or pseudo-databases (for example Active Directory).

Databases are described using the following table format:

Aspect	Description
Description	A brief description of the database, including its purpose and contents
Default location	The default location for this database (if relevant)
Users	List of the users and their permissions on the database

About Servers

Servers are any data source/repository that the product uses. Examples include web servers, FTP servers, and file servers.

Servers are described using the following table format:

Aspect	Description
Description	A brief description of the server, including its purpose and what it supplies and/or stores
Default location	The default location for this server
Settings location	The location where the permission settings for this server can be controlled
Default permissions	The default permissions for this server
Minimum permissions	The tightest set of permissions that can be applied to this server while still allowing the product to function

When determining authorization for a location, server permissions are considered. If the user meets the

permissions set for the server, then the permissions of the underlying NTFS folders are then assessed.

Administration Server Overview

The administration server is the heart of the RayManageSofti system.

It is the source of all administration data and also serves as the root node distribution server. Therefore, there are some key elements of the administration server which deserve special attention in your permissions configuration.

In particular, the RayManageSofti database must be secure, as it is the starting point of most actions and contains the central definition of the state of the RayManageSofti infrastructure.

You will also need to restrict access to certain administration processes, not all of which expose the same levels of risk.

Consider the figure below. There are two main groups of administration processes:

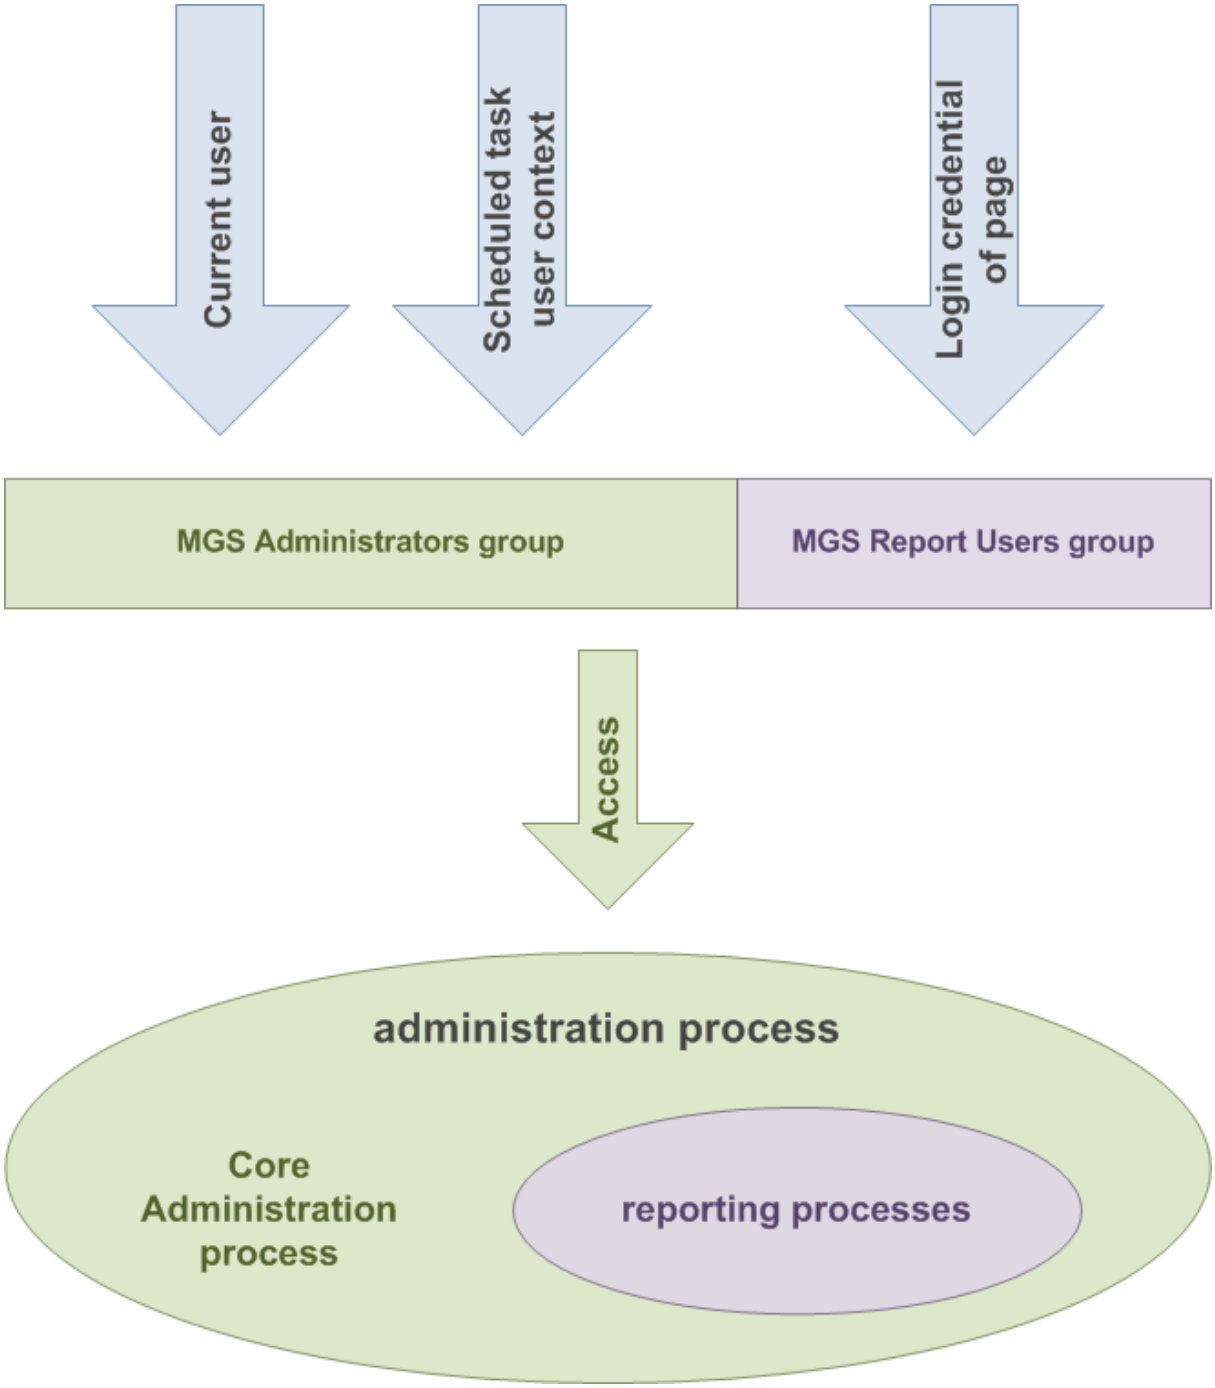

The security model for Administration

- **Core administration processes**, such as processing of inventory logs or software distribution. (Note that some of these processes, such as data and log file processing into the RayManageSofti database, can also be performed directly from distribution servers.)
- **Reporting processes**, such as generating a package installation report. This area needs to be controlled, since potentially sensitive information concerning your RayManageSofti infrastructure can be viewed here.

Both are treated as the one entity called administration process. There are many entry points to the

administration process, including the RayManageSofti console, executables in the common directory, and Internet Information Server.

The diagram shows that, wherever the entry point may be, access to the administration process should always go through two main rights groups:

- **MGS Administrators** group
- **MGS Report Users** group

These are among the groups created during installation of RayManageSofti for administration servers, and discussed in more detail under *Default user groups*.

Security for Administration Server Data Links

With respect to security for data linkages, the administration server has at least one distribution location, one reporting location, and one package stage. These locations are used by the administration server, distribution servers, and managed devices as communication links. As a consequence, these locations need permissions appropriate for all managed devices and child distribution servers. Since these locations are only directories, it is simply a matter of setting up directory shares or virtual directories in IIS. Specific details are included under each of the share names later in this section.

Administration Server Installation

During installation of Deployment Manager for administration servers, a number of Active Directory user groups are created (and in some cases given default memberships), and there are three security configuration controls that are set. The user groups are discussed first.

Default User Groups

MGS Administrators Group

Members of the **MGS Administrators** group have full access to the RayManageSofti system, including reports. During installation of Deployment Manager, additions are made to this group automatically. Remembering that your Deployment Manager installation may be spread across three kinds of servers (core, data, and reports servers), the permissions depend slightly on which of the servers is being installed.

- In all cases, the current (installing) user is added to the **MGS Administrators** group.
- For the core server only, the current computer is also added to the **MGS Administrators** group. This allows processes to run in the system context.

By default, all automated processes such as processing of logs and Active Directory reconciliation are initiated through a scheduled task, in the specified user context. The user account specified in the scheduled task must be a member of the **MGS Administrators** group.

Manual processes are run in the context of the currently logged in user. Again, the current user must belong to the **MGS Administrators** group.

In an Active Directory environment, members of this group are automatically replicated as members of the user

group **MGS Data Modifiers** (see next).

MGS Data Modifiers Group

The **MGS Data Modifiers** group is a local user group created on the Deployment Manager data server. It enables members to modify data within the RayManageSofti database.

Note:
As well as providing a mechanism for the separation of roles, there are advantages in using Microsoft SQL Server with local security groups rather than AD groups. These include the facility to access the database if domains become unavailable.

In an Active Directory environment, AD replication ensures that all members of the **MGS Administrators** group are also defined as members of the **MGS Data Modifiers** group.

In a non-AD environment, any members of the **MGS Administrators** group are added at installation time to the **MGS Data Modifiers** group, but any changes made to the **MGS Administrators** group after that time are not synchronized automatically. You must manually apply changes to the **MGS Data Modifiers** group to replicate any changes you have made to membership of the **MGS Administrators** group.

Raynet recommends that you ensure that the members of the **MGS Administrators** group and **MGS Data Modifiers** group are identical.

MGS Data Readers Group

The **MGS Data Readers** group is a local user group created on the Deployment Manager data server. It enables members to read data within the RayManageSofti database.

Note:
As well as providing a mechanism for the separation of roles, there are advantages in using Microsoft SQL Server with local security groups rather than AD groups. These include the facility to access the database if domains become unavailable.

In an Active Directory environment, AD replication ensures that all members of the **MGS Report Users** group are also defined as members of the **MGS Data Readers** group.

In a non-AD environment, any members of the **MGS Report Users** group are added at installation time to the **MGS Data Readers** group, but any changes made to the **MGS Report Users** group after that time are not synchronized automatically. You must manually apply changes to the **MGS Data Readers** group to replicate any changes you have made to membership of the **MGS Report** Users group.

Raynet recommends that you ensure that the members of the **MGS Report Users** group and **MGS Data Readers** group are identical.

MGS Distributors Group

The **MGS Distributors** group is used by the administration server to authenticate itself to child distribution servers. For more information, see *Configuring distribution server rights*.

MGS Field Technicians Group

The **MGS Field Technicians** group is created as an empty group when you install the administration server.

Members of this security group may override security policies on managed devices.

MGS Report Users Group

Membership in the **MGS Report Users** group confers access privileges only to the Deployment Manager report pages. Unlike the RayManageSofti console (which does not require you to log in), the reports area requires authentication at the start of the session. The type of authentication is determined by an installation setting (see *Authentication methods for reporting access*). If, during installation of the reports server, you selected Windows basic authentication, no automatic additions were made to this group. If you selected a Windows domain account instead (anonymous access), that account was automatically added to the **MGS Report Users** group during installation of the reports server.

In an Active Directory environment, members of this group are automatically replicated as members of the user group **MGS Data Readers** (see earlier).

SQL Server Database Roles

Two database roles are created in SQL Server to manage access to the RayManageSofti database. A number of default security groups are created and added as members of these database roles. When users are assigned to one of these security groups, they are given database permissions associated with the relevant database roles:

This database role...	contains these members...	and has this level of access...
`mgs_writer`	**MGS Data Modifiers**	Read and write access to all data, based on the standard database roles **db_owner**.
`mgs_reader`	**MGS Data Readers** **MGS Data Modifiers** **mgs_writer**	Read access to all data, based on the standard database role **db_datareader**.

Assigning "Log on Locally" Permissions to User Groups

For Web Services on Remote Administration Consoles

Members of the **MGS Administrators** security group require the right to log on locally to the (core) administration server to access web services. This is particularly relevant if your administration server is installed on a domain controller, since non-privileged users do not typically have the right to log on locally to such servers by default.

If these rights are not set, the RayManageSofti console will open on the administration console computer without any details in the console tree.

These rights are normally controlled by applying log on locally policy settings on the administration server, using the **Domain Controller Security Policy console**. (Expand the console tree as follows to locate the log on locally settings: **Windows Settings** > **Security Settings** > **Local Policies** > **User Rights Assignment**.)

You should ensure that all members of the **MGS Administrators** security group have this right.

For the Reports Server

Members of the **MGS Report Users** and **MGS Administrators** security groups require the right to log on locally to the reports server to access the Deployment Manager reports subsystems. This is particularly relevant if your reports server is installed on a domain controller, since non-privileged users do not typically have the right to log on locally to such servers by default.

These rights are normally controlled by applying log on locally policy settings on the report server, using the **Domain Controller Security Policy** console. (Expand the console tree as follows to locate the log on locally settings: **Windows Settings** > **Security Settings** > **Local Policies** > **User Rights Assignment**.)

You should ensure that all members of the **MGS Report Users** and **MGS Administrators** security groups have this right.

Refer to your Microsoft documentation about Group Policy and Windows security settings for further information.

Other Controls

Role-based Security

For security measures you can take to restrict access to tasks on the RayManageSofti console, see the *Role-based security* chapter.

Database User Name

The user name to access the RayManageSofti database through SQL Server is set on the **Configure Administration Servers** page when installing the data server.

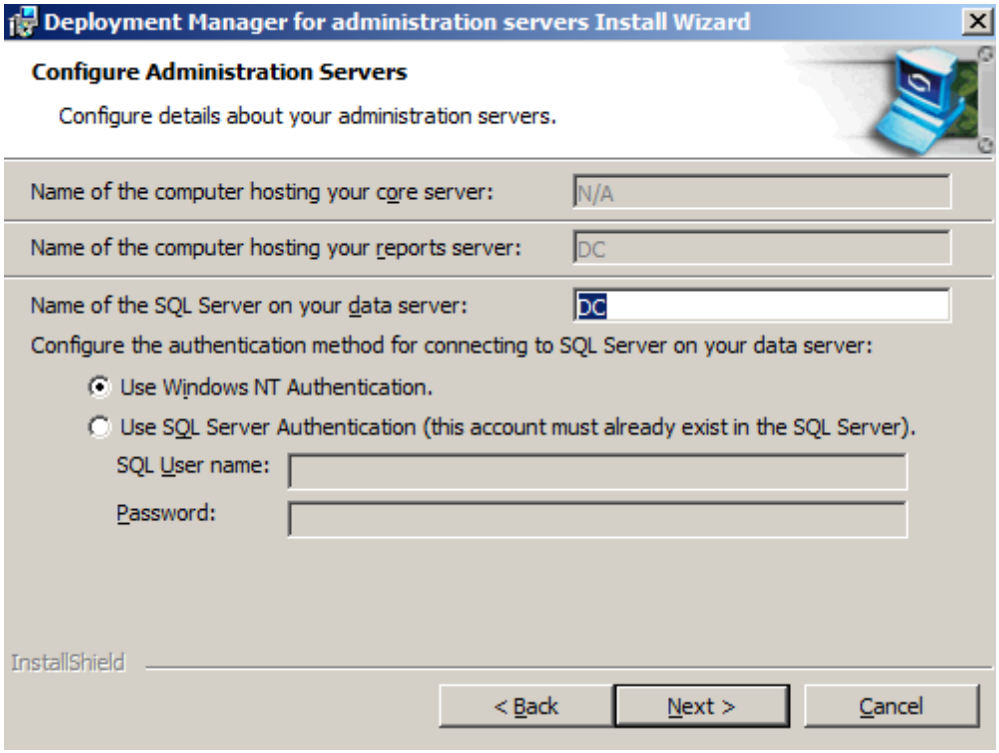

Be aware:
If you install the data and reports servers on separate computers, the same information should be specified for both installations, including the same authentication method and the individual details entered.

If you select **Use Windows NT Authentication**, any users that require access to the entire RayManageSofti database (not just the Reporting section) must be added to the **MGS Administrators** security group. The user performing the administration server installation will automatically be added to the security group. SQL Server must have been installed in **Integrated** authentication mode to enable this mode to be used.

If you select **Use SQL Server Authentication**, you must then specify a user that already exists in the SQL Server. The administration server installation will grant the specified user the necessary permissions on the RayManageSofti database tables and stored procedures. SQL Server must have been installed in **Mixed** authentication mode to enable this mode to be used.

Authentication Methods for Reporting Access

You can control the authentication method used for access to reporting on the reports server. This is set on the **Configure Reports Server Authentication** dialog when installing the reports server.

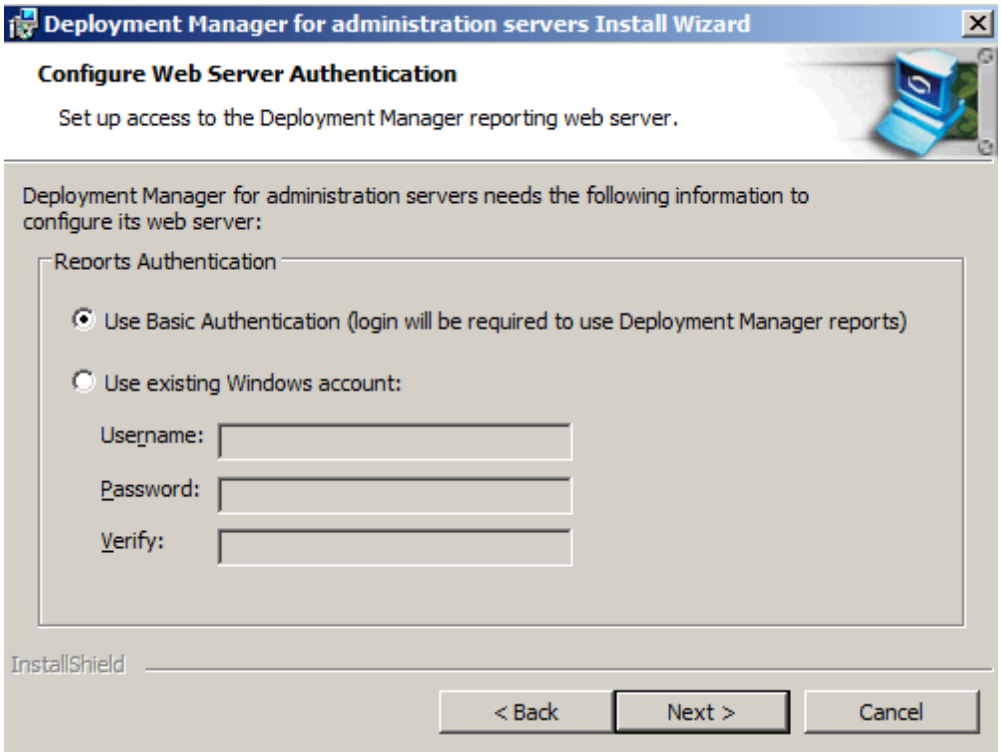

If you select **Use Basic Authentication**, any users that require access to the reporting feature must be added to the **MGS Reporting Users** security group. The users will be prompted to enter their username and password the first time they access the feature.

If you select **Use existing Windows domain account**, you are using what is known to Microsoft Internet Information Server as "anonymous access". In this mode, the specified account is added to the **MGS Reporting Users** security group. This account is then automatically used by IIS when accessing the RayManageSofti database from the reporting web page.

User Name for Scheduled Tasks on the Administration Server

Another important control is the user name under which the scheduled tasks on the administration server should be run. This is set in the credentials section of the **Configure Scheduled Tasks** page when installing the core server.

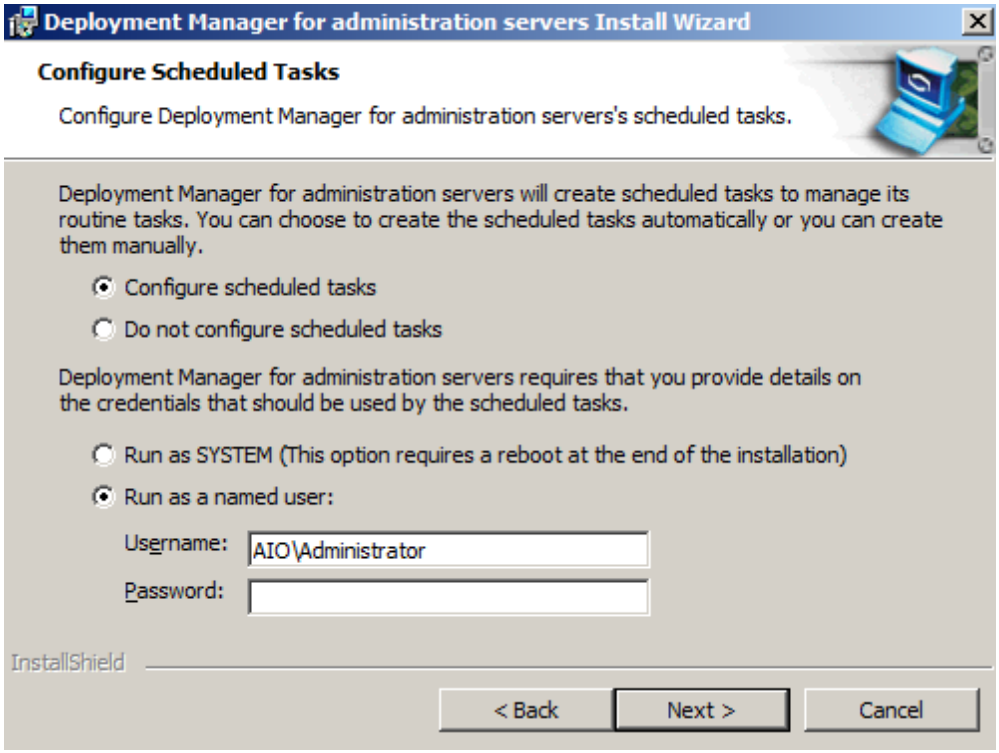

This setting becomes the standard username applied to all scheduled tasks. You may choose:

- To leave this username for all scheduled tasks. In this case you must provide adequate permissions for the most demanding scheduled tasks. For more information see *Deployment Manager distribution task users* and *Deployment Manager merged policy generator task users*.
- To edit the tasks so that they have separate usernames with different permissions as appropriate for the different tasks.

For a full description of the tasks established on the core server when Deployment Manager is installed, see the *Scheduled tasks* chapter of *RMS System Reference*.

User Name for IIS Web Applications on the Administration Server

During installation, you need to identify the user name to be used by IIS web applications on the administration server. This is set in the credentials section of the **Setup IIS Configuration** page when installing the core server.

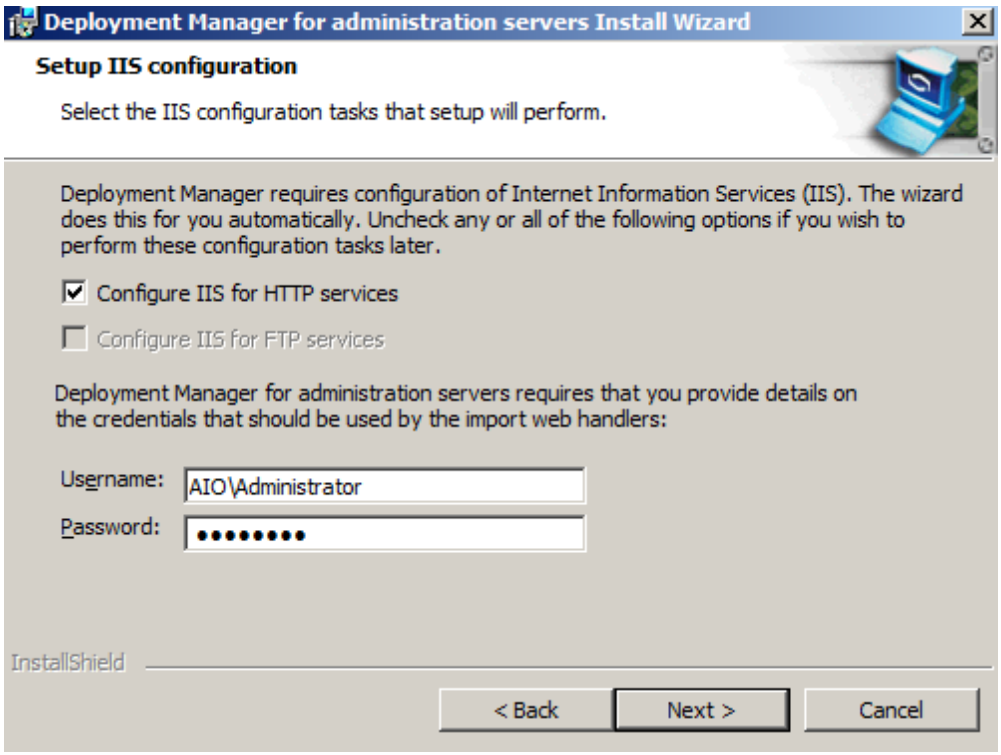

This setting becomes the standard username used by IIS web applications. See the *Preparing the Environment* chapter of the *RMS Implementation* for details about choosing an appropriate user account.

User Permissions for Remote Administration Consoles

A login dialog displays when you start Deployment Manager on the remote console. The account name defaults to the account under which you logged on to the remote console computer. Change this account name if necessary (for example, if your system is configured to require a different account name on the central administration server). However, typically you would use the same account name to log in to Deployment Manager as you used to log in on this computer initially, for the following reasons:

- Access to web services on the RayManageSofti console (**Managed Device Settings** and **Devices** nodes) is determined by the permissions of the account you use to log in to Deployment Manager
- Access to other nodes in the RayManageSofti console is determined by the permissions of the account under which you are logged on to the remote console computer.

The user account that you use to access the remote console must have access rights to the:

- RayManageSofti repository (including the software library folder structure and distribution hierarchy records) on the administration server.
- `Program Files\ManageSoft` directory (access rights to this directory are granted to members of the **MGS Administrators** group by default).
- Registry on the administration server. By default, members of the **MGS Administrators** group have full control rights to the `[Registry]\ManageSoft` key (and sub keys) and read-only access to the `HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\SecurePipeServers\WinReg` key and all

subkeys.

Without these permissions, the remote console will fail to run.

Be aware:
After adjusting user rights and group memberships, a policy merge must be performed to apply the new settings. To trigger a policy merge manually, use the `polmerge.exe`. Details regarding this agent are available within the *Merged policy generator* section of this reference. Also see the *Client-side policy merging* chapter, especially the section on how to *Add the policy information to the SQL database*, of the *RMS System Reference*.

Administration Server Features

In this section we examine a number of administration server features in detail. (For a definition of what is meant by features in this context, see *About features*) These include:

- *Administration console*
- *Application usage importer*
- *Data importer*
- *Deployment policy editor*
- *Discovery agent (AS)*
- *Discovery wizard*
- *Distribution agent (AS)*
- *Distribution agent importer*
- *Distribution hierarchy editor*
- *Distribution wizard*
- *Job server program (AS)*
- *Merged policy generator*
- *Package editor*
- *Package receiver*
- *Remote execution server agent (AS)*
- *Remote execution task importer*
- *Reporting*
- *Schedule editor*
- *Wake on LAN auto generator*
- *Wake on LAN wizard*

Administration Console

A RayManageSofti administration console enables users to perform tasks on the administration server without requiring the user to log on to the physical computer on which the administration server is installed. Users may remotely access and manipulate data stored on the RayManageSofti administration server using the standard RayManageSofti console.

Because the remote administration console uses the data and files stored on the central RayManageSofti administration server, any user of an administration console requires the same level of access to the data and files as any user of the central RayManageSofti administration server.

All data repositories (database, data files, and directories) accessed by an administration console are listed in the

table below.

Aspect	Description
Program	RayManageSofti console
Users	Users in the **MGS Administrators** group
Data files / locations	Job queue, RayManageSofti data
Communications	Web services
Databases	RayManageSofti database
Servers	RayManageSofti services

Application Usage Importer

The application usage importer (`mgsimport.exe`) is one of the many programs that make up the implementation of the Process RayManageSofti logs and inventory process in the data flow. This program is invoked through a (configurable period) scheduled task.

The `mgsimport.exe` program parses all application usage data logs located in **UsageData**, a subdirectory of **Incoming**. Logs that are successfully parsed and processed are removed from the directory. Logs that cannot be processed are moved to a **Backup** subdirectory.

The RayManageSofti database is updated with information from these logs. As such, the user credentials (set in the scheduled task) must have read and write access to the **Incoming** directory. User credentials must also have full access to the RayManageSofti database, and related stored procedures.

Aspect	Description
Program	`mgsimport.exe`
Users	Deployment Manager application usage importer task users
Data files / locations	Incoming
Communications	None
Databases	RayManageSofti database
Servers	None

Data Importer

The data importer (`mgsimport.exe`) is one of the many programs that make up the implementation of the **Process RayManageSofti logs and inventory** process in the data flow. This program is invoked through a series of (configurable period) scheduled tasks, each of which initiates resolving of a particular type of data into the RayManageSofti database.

If no arguments are specified, the `mgsimport.exe` program parses all data logs located in subdirectories of **Incoming**. Logs that are successfully parsed and processed are removed from the directory. Logs that cannot be processed are moved to a **BadLogs** subdirectory.

The RayManageSofti database is updated with information from these logs. As such, the user credentials (set in the scheduled task) must have read and write access to the **Incoming** directory. User credentials must also have full access to the RayManageSofti database, and related stored procedures.

Aspect	Description
Program	`mgsimport.exe`
Users	• Deployment Manager discovery importer task users • Deployment Manager distribution server status importer task users • Deployment Manager inventory processor task users • Deployment Manager managed device event importer task users • Deployment Manager security compliance importer task users (if Security Manager is installed)
Data files / locations	Incoming
Communications	None
Databases	RayManageSofti database
Servers	None

Deployment Policy Editor

The deployment policy editor is the user interface to the **Assign packages** and **Generate policy files** processes in the data flow. The **Assign packages** process involves attaching details of the chosen package, and optionally an associated schedule, to a policy in Active Directory. The editor must therefore have read access to **Software** and the **Schedules** directory in the repository. In addition, the user context in which the editor is run (the currently logged-in user) must also have read/write access to Active Directory and RayManageSofti database.

The **Generate policy files** process examines the policies of the specified Organizational Unit in Active Directory and generates RayManageSofti policy files. The files are placed in the **package stage** (see *Package stage (DS)*).

Aspect	Description
Program	RayManageSofti console, `Mgspoled.exe`
Users	User in the **MGS Administrators** group
Data files / locations	• Software • Schedules • Package stage
Communications	None
Databases	• RayManageSofti database • Active Directory
Servers	None

Discovery Agent (AS)

The discovery agent program, `mgsdisco.exe`, is the implementation of the **Discover devices** process in the administration server data flow. It is triggered by scheduled tasks that were previously set up by the Discovery wizard. Refer to the *Discovery* chapter of the *RMS Discovery* for more details.

Aspect	Description
Program	`mgsdisco.exe`
Users	Deployment Manager Discovery Agent Service users
Data files / locations	• JobQueue • RemoteExecution
Communications	None
Databases	None
Servers	None

Discovery Wizard

The discovery wizard is the implementation of the **Discover devices** process in the data flow. It is the users' interface to the discovery system. The wizard is a means to kick start a discovery process but does not perform the actual discovery. See the *Discovery agent (AS)* for more information on the discovery process.

All data repositories (database, data files and directories) accessed by this process are listed in the table below.

Aspect	Description
Program	RayManageSofti console
Users	Users in the **MGS Administrators** group
Data files / locations	• Incoming • JobQueue • RemoteExecution
Communications	None
Databases	None
Servers	None

Distribution Agent (AS)

The **Replicate data** process is implemented through `ndrspawn.exe` and `ndrplag.exe`. The `ndrspawn.exe` executable is a long running process that polls the job queue for new jobs. It is initially started by a scheduled task in the user context specified during installation of the administration server.

When new jobs are added to the queue, `ndrspawn.exe` invokes `ndrplag.exe` to process the jobs. After a job is processed, it is removed from the queue by `ndrplag.exe`. This implies that the user credentials, set in the scheduled task, must have read and write access to the job queue.

The tasks of replicating packages to distribution locations, and propagating jobs to distribution servers are performed by `ndrplag.exe`. A summary of the replication process is as follows:

1. Retrieve the distribution configuration to determine the current system state.

2. Check the **smart distribution database** to determine if the packages need to be distributed to local distribution locations.

3. Push the packages out (from the package stage to distribution locations) if necessary and update the smart distribution database.

4. If the distribution server listens for jobs, the distribution agent on the parent administration server contacts the **connection agent** on the distribution server to initiate the distribution agent on the distribution server. The administration server creates distribution jobs and propagates them to all child servers.

   If the distribution server polls for jobs, the administration server stores jobs for that distribution server in a separate job queue for the server. Polling is initiated by the distribution agent on the distribution server, and the **job server program** on the administration server sends the jobs in the job queue to the distribution server.

5. All errors and successes are logged to the **Incoming** directory.

For further information about the connection between the administration server and distribution server, see the distribution system process flow in *RMS System Reference*.

All data repositories (database, data files and directories) accessed by this process are listed in the table below.

Aspect	Description
Program	`ndrspawn.exe, ndrplag.exe`
Users	Deployment Manager distribution task users
Data files / locations	• Package stage • Distribution Location • Incoming • Job Queue • Distribution configuration
Communications	• Distribution Job Propagate • HTTP / HTTPS Distribution transfer • FTP Distribution transfer • NTLM Distribution transfer
Databases	Smart Distribution Database
Servers	• Package Stage Web Server on the administration server and distribution servers • Package Stage FTP Server on the administration server and distribution servers • Package Stage File Server on the administration server and distribution servers • Distribution Location Web Server on the administration server and distribution servers • Distribution Location FTP Server on the administration server and distribution servers

Aspect	Description
	• Distribution Location File Server on the administration server and distribution servers

Distribution Agent Importer

`ndrap.exe` is one of the many programs that make up the implementation of the **Process Deployment Manager logs and inventory** process in the data flow. This program is invoked through a regular (configurable period) scheduled task. `ndrap.exe` parses all distribution event logs located in **Distributor**, a subdirectory of **Incoming**. Logs that are successfully parsed and processed are removed from the directory. Logs that cannot be processed are moved to a **BadLogs** subdirectory.

The RayManageSofti database is updated with information from these logs. As such, the user credentials (set in the scheduled task) must have read and write access to the **Incoming** directory (see *Incoming (DS)*). In addition, the user credentials must also have full access to the RayManageSofti database, and related stored procedures.

Aspect	Description
Program	`ndrap.exe`
Users	Deployment Manager managed device event importer task users
Data files / locations	Incoming
Communications	None
Databases	RayManageSofti database
Servers	None

RayManageSofti logs distribution event data to the Windows event log according to the settings in `[Registry]\ManageSoft\EventLog\CurrentVersion\EventLogPolicy`.

If the value of this registry setting is...	RayManageSofti logs...
All	All distribution events
Failures	Only failed distribution events
Significant	The final result of the distribution event (success or failure), but not intermediate progress messages

Distribution Hierarchy Editor

This distribution hierarchy editor is the user interface to the **Configure hierarchy details** process in the data flow. It consists of wizards and property pages which are used to create logical distribution servers and locations. Data from the **Distribution Configuration** (see *Distribution configuration (AS)*) are loaded into the system when the RayManageSofti console initially loads up, and are represented graphically under the distribution node. As servers and locations are created or deleted, the distribution configuration and the RayManageSofti database are updated. Aside from the wizards, three key actions can be performed on the distribution hierarchy:

1. Synchronize

Synchronize the database records with the actual contents of distribution locations in the hierarchy. This is done by placing a job in the job queue, which is propagated down the hierarchy. In due course, logs (which detail the contents of each distribution server) are returned to the administration server. The information is used to update the database's image of each server.

2. **Verify**
   Sends out configuration jobs to all distribution servers in the hierarchy. In due course logs from each server are returned to the administration server, indicating whether the server's configuration failed or succeeded.

3. **Rebuild**
   Based on information stored in the RayManageSofti database (an image of the desired content of each distribution server), automatically distribute packages to the appropriate distribution servers.

These actions are also considered to be a part of the configuration process and are represented as such in the data flow.

Aspect	Description
Program	RayManageSofti console
Users	**MGS Administrators**
Data files / locations	• Job Queue • Distribution Configuration
Communications	None
Databases	RayManageSofti database
Servers	None

Distribution Wizard

The distribution wizard is the implementation of the **Distribute packages** process in the data flow. It is the users' interface to the distribution system. The wizard is a means to kick starting a distribution process but does not perform the actual distribution. See the *Distribution agent (AS)* for more information on the actual distribution process. The wizard reads data from the following directories:

- Distribution configuration
- Software library
- Schedules
- Policies

Information gathered is made available to the user. Once the user has chosen what to distribute and where to distribute it to the wizard proceeds to check if any packages (schedules and policies are also considered packages when readied for distribution) needs to be repacked. This is done by comparing the packages' MD5 digest to the values stored in the smart distribution database. The pack process places the final package in the package stage.

Finally, jobs are created for the distribution and placed in the job queue. Successes and failures are logged to the **Incoming** directory. The **Replicate Data** process takes over from that point.

All data repositories (database, data files and directories) accessed by this process are listed in the table below.

Aspect	Description
Program	RayManageSofti console
Users	User in the **MGS Administrators** group
Data files / locations	• Package Stage • Incoming • Job Queue • Distribution Configuration • Software Library • Schedules • Policies
Communications	None
Databases	Smart Distribution Database
Servers	None

Job Server Program (AS)

The job server program (`mgsjobsrv.exe`) responds to job requests from child distribution servers that are configured to poll for distribution jobs. It is called through an HTTP/HTTPS request that contains the UID of the child distribution server and MD5 details for components of the distribution hierarchy that the child distribution server already knows about.

In response to the call from the child distribution server, `mgsjobsrv.exe` looks for a job queue file for the appropriate UID. It then reads out each job. If the MD5 associated with the job has not been received in the HTTP/HTTPS request, `mgsjobsrv.exe` creates a new job to pass the relevant hierarchy details to the child distribution server.

Aspect	Description
Program	`mgsjobsrv.exe`
Users	Deployment Manager distribution task users
Data files / locations	Job Queue
Communications	Distribution Job Poll
Databases	None
Servers	None

Merged Policy Generator

The `polmerge.exe` program merges policy for computers and users from a set of Active Directory domains and populates the Resultant Set of Policy (RSoP) information in the RayManageSofti database.

This program is invoked through the scheduled task **Merge Deployment Manager policies** that recurs on a configurable cycle. It can also be invoked through the scheduled task **Reconcile Deployment Manager directory tables with AD** to reconcile Active Directory data without merging policy. This is useful for domains which are not under management but contain users to whom rights may be granted.

In an environment that has more than one Active Directory domain, the domain of the user that runs the scheduled task must be trusted by each domain whose policy data is to be merged. Policy merging only requires read access to the domain being merged; write access is not required. Write access is only required to edit policy in a domain.

Minimum Active Directory Rights

The account used to execute `polmerge.exe` must have read access to the following objects in each domain that is being merged:

- The **CN=System** container
- The full contents of all policies (under the **CN=Policies,CN=System** container) that contain any RayManageSofti packages
- All users and computers which are to be managed using Deployment Manager
- All security groups which are used to:
  - Filter policies and packages
  - Restrict and control access to Deployment Manager administration functions
- Organizational units and other container objects containing any of the users, computers, or groups mentioned above

In addition, rights are required to access the root domain of the forest containing the domain being merged and read all groups which are used to filter policies and packages in the domain being merged.

You can use the `-x` command line option to restrict the policy merging process so that it does not require access to the root domain. This is useful if you know that the root domain does not contain security groups that are used to filter policies and packages. For further information, see *Generating merged policy* in the *Command line tools* chapter of *RMS System Reference*.

Default Deployment Manager Privileges

The default configuration when Deployment Manager is installed is that:

- Scheduled events run either under the `SYSTEM` account on the local computer if the local computer is not a domain controller (if the local computer is a domain controller, the `SYSTEM` user has Domain Administration rights and another user is nominated to run scheduled events), or under a local user account of your choice
- Deployment Manager will attempt to merge all configured domains
- No cross-domain trusts are configured
- No administrator rights are granted
- Only the domain of the installing user is configured to be merged
- No check is made that the scheduled task user has the required rights

WARNING
The default configuration described above is suitable only for merging policies for a single domain. If you want to merge policies for multiple domains, you must modify accounts and permissions to meet the requirements documented in this section. For further information, see *Configuring environments with multiple domains*.

Symptoms and Performance

If read access is not available to any required object in the domain, the policy merging process may either:

- Terminate. (If you are running in console mode, an error message is displayed.)
- Finish successfully but leave an unexpected result in the RayManageSofti database.

For example, if a computer account that is being managed by Deployment Manager was not readable, the computer would be deleted from the RayManageSofti database by the policy merging process.

Possible Reasons Why the Required Rights are not Available Include:

- You have not changed the default configuration and are using multiple domains
- You have configured an additional domain to be merged without granting the required rights

Be aware:
The `polmerge` algorithm traverses and loads much of the content of your Active Directory. The domain controller outputs are efficient, but their creation is the predominant contributor to the `polmerge` runtime.

Security in a Multi-domain Environment

The simplest way to ensure that the **Merge Deployment Manager policies** task has sufficient permissions is to grant Enterprise Administrator rights to the account that will run the scheduled task, and operate the task with the `-a` option to merge all domains. As this is not the most secure configuration, an alternative is to configure multiple scheduled tasks, each running as a different account and each running a merge of just one or a few domains (using the `-d` option).

Aspect	Description
Program	`polmerge.exe`
Users	Deployment Manager merged policy generator task users
Data files / locations	None
Communications	None
Databases	• RayManageSofti database • Active Directory
Servers	None

Package Editor

The package editor is the user interface to the **Create packages** and **Modify packages** processes in the data flow. The editor is accessible through the RayManageSofti console (**software** node) on the administration server. The editor requires read and write access to Software when loading and saving packages. In addition, when a package is saved (new or modified), an entry is written to the RayManageSofti database. The user context in which the console is run must also have full access to the RayManageSofti data base.

Aspect	Description
Program	RayManageSofti console (Software node)
Users	User in the **MGS Administrators** group
Data files / locations	Software Library
Communications	None
Databases	RayManageSofti database
Servers	None

Package Receiver

The package receiver is the implementation of the **Import packages** process in the data flow. The process, though automated, is identical to the **Create packages** process. As such, the data stores accessed are the same, as are the permissions required.

Aspect	Description
Program	RayManageSofti console (Software node)
Users	User in the **MGS Administrators** group
Data files / locations	Software Library
Communications	None
Databases	RayManageSofti database
Servers	None

Remote Execution Server Agent (AS)

The remote execution server agent, `mgsresa.exe`, is the implementation of the **Remotely execute tasks on managed devices** process in the administration server data flow. It is triggered by scheduled tasks that were previously set up by the Discovery or Remote Command wizards. Refer to the *Discovery* chapter of the *RMS Discovery* for more details.

Aspect	Description
Program	`mgsresa.exe`
Users	RayManageSofti Remote Execution Server Agent Users

Aspect	Description
Data files / locations	• JobQueue • RemoteExecution
Communications	None
Databases	None
Servers	None

Remote Execution Task Importer

The remote execution task importer (`mgstaskgat.exe`) is one of the many programs that make up the implementation of the **Process ManageSoft logs and inventory** process in the data flow. This program is invoked through a (configurable period) scheduled task.

The `mgstaskgat.exe` program parses all remote execution task data logs located in **ActionStatus**, a subdirectory of **Incoming**. Logs that are successfully parsed and processed are removed from the directory. Logs that cannot be processed are moved to a **BadLogs** subdirectory.

The RayManageSofti database is updated with information from these logs. As such, the user credentials (set in the scheduled task) must have read and write access to the **Incoming** directory. User credentials must also have full access to the RayManageSofti database, and related stored procedures.

Aspect	Description
Program	`mgstaskgat.exe`
Users	Deployment Manager remote execution task importer task users
Data files / locations	Incoming
Communications	None
Databases	RayManageSofti database
Servers	None

Reporting

This feature provides users with functionality to query and view reports in a friendly format. It covers the **Generate reports** and **Update query database** processes in the data flow. Despite the naming of the latter process, this feature will only ever query the database. No data are modified in the tables.

The Reporting feature has two main sections:

- Active Directory browser
- The reports

The Active Directory browser reads data from Active Directory, which is used to filter reports.

The reports feature queries data from the database through the following chain of authentication:

1. User logs in to the reports page.

2. The credentials are verified by Microsoft Internet Information Server.

3. The backend server code is executed to query the RayManageSofti database. For this to succeed, the credentials passed in must have access rights (including execute permissions) to the directory where the code is located (usually `C:\Program Files\ManageSoft\Reporter\Web`). The credentials must also have access rights set in RayManageSofti role-based security. For more details about role-based security, see the *Role-based security* chapter.

4. The same credentials are passed from Internet Information Server to SQL Server when the RayManageSofti database is queried. SQL Server will grant rights to read data from the RayManageSofti database to members of either the **MGS Data Readers** group, or the **MGS Data Modifiers** group, in Active Directory. Therefore the credentials passed must also be a member of either of those groups.

Be aware:
Access to the reporting feature requires a user to be a member of the **MGS Report Users** or **MGS Administrators** group. Access to the database, however, requires a user to be a member of the **MGS Data Readers** or **MGS Data Modifiers** group. Typically, members of the **MGS Report Users** group and **MGS Data Readers** group are identical, and members of the **MGS Administrators** group and **MGS Data Modifiers** group are identical.

The same credentials are passed on from Internet Information Server to Active Directory as required. The credentials passed must have read access to all relevant organizational units within Active Directory.

Aspect	Description
Program	Microsoft Internet Information Server
Users	Deployment Manager reports viewer, a member of either **MGS Report Users** or **MGS Administrators** groups
Data files / locations	Reporter\Web
Communications	None
Databases	• RayManageSofti database • Active Directory
Servers	• Reporting web server • ASP.NET web server • Web controls web server

Schedule Editor

The schedule editor is the user interface to the **Create new schedules** and **Modify existing schedules** processes in the data flow. The editor is accessible through the RayManageSofti console (**scheduling** node) on the administration server. When the scheduling node is opened, all schedules in the **Schedules** directory are loaded into the system. Any modifications or new schedules are written to the directory. When a new schedule is created, the resultant package name is also written to the RayManageSofti database.

Aspect	Description
Program	RayManageSofti console (Scheduling)
Users	User in the **MGS Administrators** group
Data files / locations	Schedules
Communications	None
Databases	RayManageSofti database
Servers	None

Wake on LAN Auto Generator

The Wake on LAN auto generator gets its initial policy information from the RayManageSofti database which requires read access for Wake on LAN.

It later uses the Wake on LAN wizard interface to generate the **Wake on LAN** tasks. In this case, the same security rules for the Wake on LAN wizard apply.

Aspect	Description
Program	`mgswolauto.exe` (run via scheduled task)
Users	User in the **MGS Administrators** group
Data files / locations	None
Communications	None
Databases	RayManageSofti database
Servers	None

Wake on LAN Wizard

The wizard retrieves data from Active Directory, the RayManageSofti database and distribution hierarchy file. For this to occur, Wake on LAN must have read access to these three locations.

It also stores Wake on LAN information in the RayManageSofti database for reporting purposes. For this to occur, Wake on LAN must have write access to the RayManageSofti database.

The Wake on LAN wizard creates jobs for the distribution system which it needs to place in the job queue. For this to occur, read and write access must be given to Wake on LAN.

Aspect	Description
Program	Active Directory Users and Computers (from the All Tasks section of the context menu on domain and OU nodes), Deployment Policy Editor from the context menu of the computer settings for RayManageSofti
Users	User in the **MGS Administrators** group
Data files /	None

Aspect	Description
locations	
Communications	None
Databases	• RayManageSofti Database • Active Directory • Distribution Hierarchy
Servers	None

Administration Server Data Files / Locations

This section describes the major data storage locations on the administration server:

- *Distribution configuration (AS)*
- *Distribution location (AS)*
- *Incoming (AS)*
- *Job queue (AS)*
- *ManageSoft Data (ManageSoft$)*
- *ManageSoftJQ (AS)*
- *Package stage (AS)*
- *RemoteExecution\Actions (AS)*
- *RemoteExecution\Public (AS)*
- *Reporter\Web*
- *Schedules*
- *Software*
- *Tools*

In addition to the files and locations described below, `C:\ManageSoft\Repository\History\` stores backup copies of projects, schedules, the distribution hierarchy, and so on. These files can be used to revert to an earlier version of a file if required (discuss this with your Raynet consultant). You can safely remove these files if you no longer require backups.

Distribution Configuration (AS)

The distribution configuration is the storage location for all currently active distribution hierarchies. This includes distribution servers, distribution locations, and distribution groups. The following components require access to this location:

- **Distribution wizard**
  Loads distribution configuration data, which are displayed to the user

- **Distribution hierarchy editor**
  Loads distribution configuration data to populate the Distribution node in the RayManageSofti console, and writes data back to this location when the hierarchy is modified

- **Distribution agent**
  Reads data from this location when data and packages are replicated to child servers and distribution locations.

Default location	`C:\ManageSoft\Repository\DeploymentLocations\Common`	
Settings location	Directory properties, Security	
Default permissions	**User**	**Access rights**
	MGS Administrators	Full control
	Administrators	Full control
	System	Full control
	Note: Default permissions are inherited from the parent folder `C:\ManageSoft`	
Minimum permissions	**User**	**Access rights**
	MGS Administrators	Full control
	Deployment Manager distribution task users	Read, List folder contents

Distribution Location (AS)

The storage location for all applications, policies, schedules and managed device settings that have been made available to managed devices.

Default location	• `C:\ManageSoft\LocalDeployment (if not linked)` • `C:\ManageSoft\Staging\Common (if linked)` shared as ManageSoftDL (FTP, HTTP, HTTPS) and ManageSoftDL$ (file)	
Settings location	Directory properties, Security	
Default permissions	**User**	**Access rights**
	Everyone	Full control
	MGS Administrators	Full control
	Administrators	Full control
	Anonymous logon	Full control
	Network	Full control
	System	Full control
Minimum permissions	**User**	**Access rights**
	MGS Administrators	Full control
	Deployment Manager distribution task users	Full control
	Deployment Manager Managed Device User	Read, List folder contents
	Domain Computers	Read, List folder contents
	Domain Controllers	Read, List folder contents

The web server, FTP server, and file server settings should be considered in conjunction with these NTFS file permissions.

Incoming (AS)

This location is used by administration server components, child distribution servers, and managed devices to report data back to the system. Managed device users and child distribution servers must have write access to this location. In addition, the user context in which the various administration server scheduled tasks are run must have full access to this location.

Default location	For FTP and file uploads: `C:\ManageSoft\Incoming`, shared as ManageSoftRL (FTP) and ManageSoftRL$ (file) For HTTP and HTTPS uploads: `C:\Program Files\ManageSoft\DotNet`	
Settings location	Directory properties, Security	
Default permissions	**User**	**Access rights**
	Everyone	Full control
	MGS Administrators	Full control
	Administrators	Full control
	Anonymous logon	Full control
	Network	Full control
	System	Full control
Minimum permissions	**User**	**Access rights**
	Domain Computers	Full control
	Domain Controllers	Full control
	Deployment Manager discovery importer task user	Full control
	Deployment Manager distribution server status importer task users	Full control
	Deployment Manager distribution task users	Full control
	Deployment Manager inventory processor task users	Full control
	Deployment Manager managed device event importer task users	Full control
	Deployment Manager Managed Device User	Full control
	Deployment Manager remote execution task importer task users	Full control
	Deployment Manager uploader scheduled task users	Full control

The web server, FTP server, and file server settings should be considered in conjunction with these NTFS file permissions.

Job Queue (AS)

The Job Queue is the storage location for all distribution jobs that are yet to be processed. It is accessed by:

- The Distribution Wizard, which adds distribution jobs to the queue
- Distribution Hierarchy Editor, which adds configuration jobs to the queue
- Distribution Agent, which processes and removes jobs from the queue

Full control must be granted to two types of users:

- MGS Administrators
- The user context in which Deployment Manager scheduled tasks are run

Default location	`C:\Documents and Settings\All Users\Application Data\ManageSoft Corp\ManageSoft\Replication Agent\JobQueue` shared as ManageSoftJQ$	
Settings location	Directory properties, Security	
Default permissions	**User**	**Access rights**
	Administrators	Full control
	MGS Administrators	Full control
	System	Full control
	MGS Report users	Write only
	Note: Default permissions are inherited from the parent folder `C:\Documents and Settings\All Users\Application Data\ManageSoft Corp`.	
Minimum permissions	**User**	**Access rights**
	MGS Administrators	Full control
	Deployment Manager distribution task users	Full control

ManageSofti Data (ManageSoft$)

The RayManageSofti Data folder is the storage location for all RayManageSofti data, including the following items:

- Packages
- Schedules
- Managed device settings

- Failover locations
- RayManageSofti database
- Distribution hierarchy

It is also the parent folder for the following locations:

- Distribution configuration
- Distribution location
- Incoming
- Package stage
- Remote execution actions
- Schedules
- Software library

Full control must be granted to two types of users:

- MGS Administrators
- The user context in which Deployment Manager scheduled tasks are run

Default location	`C:\ManageSoft shared as ManageSoft$`	
Settings location	Directory properties, Security	
Default permissions	**User**	**Access rights**
	Administrators	Full control
	MGS Administrators	Full control
	System	Full control
Minimum permissions	**User**	**Access rights**
	MGS Administrators	Full control
	Deployment Manager Scheduled Task Users	Full control

ManageSoftJQ (AS)

This virtual folder is created at installation time.

The following components require access to this location:

- **Job server program**
  Polls parent distribution server for distribution jobs.

Default location	`C:\Program Files\ManageSoft\Replicator`
Settings location	Internet Services Manager

Default permissions	User	Access rights
	Anonymous	Read, Execute permission on scripts and executables
Minimum permissions	**User**	**Access rights**
	Anonymous	Read, Execute permission on scripts and executables

Package Stage (AS)

The package stage is intended as an intermediate area. The directory typically contains packages, schedules, policy files, and managed device settings packages. It is used by many RayManageSofti components, all of which are accessed through the RayManageSofti console. Users of the RayManageSofti console must have write access to this area.

In addition, child distribution servers must also have read access to this location, as data are pulled from here to their local directories.

Default location	`C:\ManageSoft\Staging\Common`, shared as ManageSoftDS (FTP, HTTP, HTTPS) and ManageSoftDS$ (file)	
Settings location	Directory properties, Security	
Default permissions	**User**	**Access rights**
	Everyone	Full control
	MGS Administrators	Full control
	Administrators	Full control
	Anonymous logon	Full control
	Network	Full control
	System	Full control
Minimum permissions	**User**	**Access rights**
	MGS Administrators	Full control
	Deployment Manager distribution task users	Full control
	MGS Distributors	Read, List folder contents

The web server, FTP server, and file server settings should be considered in conjunction with these NTFS file permissions.

RemoteExecution \ Actions (AS)

This is the location where remote execution actions are stored, prior to their retrieval by managed devices.

The physical folder to which this maps is accessed by the remote execution scheduler, which writes remote execution jobs to this location.

For more information about the remote execution subsystem, see the *Discovery* chapter of the *RMS Discovery*.

Default location	`C:\ManageSoft\RemoteExecution\Actions\Public`	
Settings location	Directory properties, Security	
Default permissions	**User**	**Access rights**
	Administrators	Full control
	MGS Administrators	Full control
	System	Full control
Minimum permissions	**User**	**Access rights**
	Administrators	Full control
	MGS Administrators	Full control
	System	Full control

The file server settings should be considered in conjunction with these NTFS file permissions.

RemoteExecution \ Public (AS)

This is the location where remote execution tools are stored, including the managed device agent, adoption executables, and remote inventory executables.

The physical folder to which this maps is accessed by the remote execution subsystem on managed devices:

- Managed devices run `mgsreca.exe` directly from the **ClientAgent** subfolder
- Managed devices run the adoption agent from the **Adoption** subfolder to install Deployment Manager for managed devices
- Managed devices run zero-touch inventory from the **Inventory** subfolder

For more information about the remote execution subsystem, see the *Discovery* chapter of the *RMS Discovery*.

Default location	`C:\Program Files\ManageSoft\RemoteExecution\Public`	
Settings location	Directory properties, Security	
Default permissions	**User**	**Access rights**
	Administrators	Full control
	Anonymous logon	Read and execute, list file contents, read
	Everyone	Read and execute, list file contents, read
	Network	Read and execute, list file contents, read
	System	Full control

Minimum permissions	User	Access rights
	Administrators	Full control
	Anonymous logon	Read and execute, list file contents, read
	Everyone	Read and execute, list file contents, read
	Network	Read and execute, list file contents, read
	System	Full control

The file server settings should be considered in conjunction with these NTFS file permissions.

Reporter \ Web

This directory contains all backend code used for generating report pages. If the level of access is changed for any group, check that it does not affect any report users, either directly or implicitly through group membership.

Be aware:
A defect in ASP.NET (documented at *http://support.microsoft.com/kb/317955*) means that the minimum permissions on this directory must also be treated as minimum permissions for each ancestor directory of Reporter\Web, up to and including the root. In particular, **MGS Administrators** and **MGS Report Users** must have Read & Execute permissions on `Reporter\Web` and all its ancestor directories. (You also need to extend the same permissions downward to subdirectories to enable execution of other files in the application.)

Also be aware:
When ASPX applications are first run or re-configured, .NET compiles pages into a temporary area (the current default on Windows 2000 is `C:\WINNT\Microsoft.NET\Framework\v1.0.3705\Temporary ASP.NET Files`). Below this are subdirectories for each .NET application, including `managesoftrp`. By default, administrator users have read/write permissions on this temporary directory structure, but other authorized users only have read permissions. This means that, for non-administrator users, the .NET write will fail. So that reports can operate, you need to also give authorized users read/write permissions on the appropriate temporary directories.

Default location	`C:\Program Files\ManageSoft\Reporter\Web` (and all ancestor and child directories - see notes above), and `C:\WINNT\Microsoft.NET\Framework\v1.0.3705\Temporary ASP.NET Files\managesoftrp`	
Settings location	Directory properties, Security	
Default permissions	Default permissions are inherited from the parent folder. Permissions are only specified directly from the `Program Files` folder, and these defaults vary for different operating systems.	
Minimum permissions	**User**	**Access rights**
	Administrators	Full control
	MGS Administrators	Read & Execute
	MGS Report Users	Read & Execute
	Server Operators	Modify

	SYSTEM	Full control
	CREATOR OWNER	Full control

Schedules

The Schedules directory stores details of all schedules that are currently on the administration server, which may or may not have been distributed. It does not include details of old schedules that have been deleted from the administration server, but still reside on distribution servers and locations.

The following components require access to this location:

- **Schedule editor**
  Loads schedules into the RayManageSofti Console and writes new / modified schedules to the directory
- **Deployment policy editor**
  Loads schedules details into a list, which allows the user to choose that schedule to assign to a policy
- **Distribution wizard**
  Loads into memory details of the schedule, which are used to generate distribution jobs. In addition, the wizard will pack the schedule into a package, which is saved to the package stage

Default location	`C:\ManageSoft\Repository\Schedules`	
Settings location	Directory properties, Security	
Default permissions	**User**	**Access rights**
	MGS Administrators	Full control
	Administrators	Full control
	System	Full control
	Note: Default permissions are inherited from the parent folder `C:\ManageSoft`	
Minimum permissions	**User**	**Access rights**
	MGS Administrators	Full control

Software

The software library within the Software node is the location where all projects (of packages) are stored. The physical directory to which this maps is accessed by the following components:

- **Package editor**
  Loads existing package details into the Software node in the RayManageSofti console. New projects and modified package details are written out to this directory.
- **Distribution wizard**
  Loads the list of packages (available for distribution) into a list where the operator selects the packages for

distribution.

- **Package receiver**
  Writes into the directory new project details for the package that has been received.

- **Deployment policy editor**
  Loads the list of available packages into a read-only list, where the operator selects the packages to assign to a policy.

Default location	`C:\ManageSoft\Repository\Packages`	
Settings location	Directory properties, Security	
Default permissions	**User**	**Access rights**
	MGS Administrators	Full control
	Administrators	Full control
	System	Full control
	Note: Default permissions are inherited from the parent folder `C:\ManageSoft`	
Minimum permissions	**User**	**Access rights**
	MGS Administrators	Full control

Tools

This directory contains common tools that the RayManageSofti administrator may want access to from other computers. It is provided for convenient access. RayManageSofti does not use the share for any purpose. Since this location contains core RayManageSofti system files, it is advisable to limit access to administrators only.

Default location	`C:\Program Files\ManageSoft\Common`, shared as ManageSoftTL$ (file)	
Settings location	Directory properties, Security	
Default permissions	Default permissions are inherited from the parent folder. Permissions are only specified directly from the `Program Files` folder, and these defaults vary for different operating systems.	
Minimum permissions	**User**	**Access rights**
	MGS Administrators	Read, Execute
	Deployment Manager distribution task users	Read, Execute
	Administrators (Local computer security group)	Full control
	System	Full control

The corresponding file share permissions are usually more restrictive, denying all write access via the share by

default.

Administration Server Communications

This section describes network connections that may be established between two RayManageSofti features, either on the administration server, or spread across an administration server and another server:

- *Distribution job propagate (AS)*
- *Distribution job poll (AS)*
- *Distribution transfer FTP*
- *Distribution transfer HTTP/HTTPS*
- *Distribution transfer NTLM*
- *Web services*

Distribution Job Propagate (AS)

Performs the propagation of a distribution job from the administration server to a child distribution server.

Aspect	Description
Initiator	Distribution Agent
Acceptor	Connection Agent
Protocol	Proprietary protocol. If parent server authentication is enabled, then the user running the Distribution Agent must be a member of the MGS Distributors security group

Distribution Job Poll (AS)

Polls for distribution jobs from the administration server distribution server to the software library.

Aspect	Description
Initiator	Distribution Agent
Acceptor	Job server program
Protocol	HTTP or HTTPS, Windows integrated authentication (anonymous access allowed by default)

Distribution Transfer FTP

Performs the transfer of data (files, directory listings) over the FTP protocol.

Aspect	Description
Initiator	Distribution Agent
Acceptor	• Package stage FTP server on the administration server or distribution servers • Distribution location FTP server on the administration server or distribution servers

Aspect	Description
Protocol	FTP, username and password authentication (“anonymous” is the default user)

Distribution Transfer HTTP / HTTPS

Performs the transfer of data (files, directory listings) over the HTTP or HTTPS protocols. See *HTTPS security configuration* for further information about HTTPS security configuration

Aspect	Description
Initiator	Distribution Agent
Acceptor	• Package Stage Web Server on the administration server or distribution servers • Distribution Location Web Server on the administration server or distribution servers
Protocol	HTTP or HTTPS, username/password authentication or Windows integrated authentication.

Distribution Transfer NTLM

Performs the transfer of data (files, directory listings) over the NTLM protocol.

Aspect	Description
Initiator	Distribution Agent
Acceptor	• Package stage file server on the administration server or distribution servers • Distribution location file server on the administration server or distribution servers
Protocol	NTLM, Windows integrated authentication

Web Services

Administration consoles access web services exposed by the central RayManageSofti administration server to perform some of their tasks.

Aspect	Description
Initiator	Administration console
Acceptor	RayManageSofti services
Protocol	SOAP / XML via HTTP or HTTPS, Windows integrated authentication

Administration Server Databases

The RayManageSofti databases include the core *RayManageSofti database* and two other stores of data:

- *Active Directory*

- *Smart distribution database (AS)*

RayManageSofti Database

The RayManageSofti database is the core of the system. It is referenced by all major components (nodes in the RayManageSofti console), as well as RayManageSofti scheduled tasks. The database contains data relating to:

- Users
- Managed devices
- Distribution servers and locations
- The Active Directory hierarchy
- Packages, and all other management data
- Inventory data
- Stored procedures used by the reporting and core system

Manage access to this database very carefully, as changes to permissions, in the stored procedures, tables or logins, could have catastrophic consequences on the RayManageSofti system. On installation of the product, two groups are added to the database login:

- **MGS Data Modifiers** - Has full access to the RayManageSofti system. Typically, this local group contains all members of the **MGS Administrators** group.

- **MGS Data Readers** - Has access to the reports data. Typically, this local group contains all members of the **MGS Report Users** group.

These logins should never be removed or renamed, as the RayManageSofti database tables grants permissions to users based on these groups. When users are added to these groups in Active Directory, they will automatically have access to the appropriate data in the database.

Aspect	Description
Default location	`C:\ManageSoft\Repository\Database`
Users	• MGS Data Modifiers have access to all tables and stored procedures • Deployment Manager Data Readers have access to the tables and stored procedures relevant to the reporting feature

Active Directory

Active Directory is a database storing information about the structure of your enterprise. Software management data relevant to Deployment Manager is also stored here. Various Deployment Manager components extract data from here, which are then used to update the RayManageSofti database. The following components require access to Active Directory:

- **Deployment policy editor**
  Writes new package assignment data to Active Directory

- **Merged policy generator**
  Reads package assignment data from Active Directory, which are used to update the RayManageSofti database

- **Reports feature**
  Reads the organizational structure from Active Directory, used by users to filter reports information

- **AD Reconcile Agent**
  Reads data from Active Directory to synchronize the RayManageSofti database

Aspect	Description
Default location	Active Directory Server
Users	Domain administrators

Smart Distribution Database (AS)

The smart distribution database is a caching mechanism used by the system to avoid unnecessary distribution or packing of packages (this includes schedules, policies, and client settings data). It is accessed by the following components:

- **Distribution wizard**
  The wizard checks this database to determine if a package needs to be repacked before being distributed. If a package does need to be repacked, the smart distribution database is updated with the package's new details.

- **Distribution agent**
  The agent always checks the smart distribution database to determine if a package needs to be distributed/redistributed. When a package is distributed to a server or location for the first time, it is recorded in the database. If a package was redistributed to a server or location, the database is updated with this information.

Aspect	Description
Default location	`C:\ManageSoft\Repository\DeploymentLocations\distributeLog.smd`
Users	Users in the MGS Administrators group

Additional Servers

Additional servers are any data source/repository that RayManageSofti uses. On the administration server, these include:

- *ASP.NET web server*
- *Distribution location file server (AS)*
- *Distribution location FTP server (AS)*
- *Distribution location web server (AS)*
- *Incoming file server (AS)*
- *Incoming FTP server (AS)*
- *Incoming web server (AS)*
- *Package stage file server (AS)*
- *Package stage FTP server (AS)*
- *Package stage web server (AS)*
- *RayManageSofti data file server*
- *RayManageSofti job queue server*

- *RayManageSoft services*
- *Remote execution task file server (AS)*
- *Remote execution task FTP server (AS)*
- *Remote execution task web server (AS)*
- *Remote execution tools server (AS)*
- *Reporting web server*
- *Tools file server (AS)*
- *Web controls web server*

ASP.NET Web Server

This is the web server hosting the ASP.NET controls, used to draw graphs in the reporting pages. The most immediate sign of permission problems in this area is when no summary graphs are drawn.

Be aware:
NTFS file permissions on the underlying directories for the IIS virtual directories also apply through HTTP and HTTPS access. The settings here should therefore be considered in conjunction with the appropriate NTFS settings. Access checks for both the share permissions and the NTFS permissions must succeed before a user is granted access.

Default location	`http://{hostname}/aspnet_client`
Settings location	Internet Services Manager
Default configuration	Basic authentication, anonymous access, or integrated Windows authentication as configured by .NET installation.
Minimum configuration	If the reporting web server uses anonymous access, then the ASP.NET web server also requires anonymous access. Otherwise the minimum requirement is basic authentication.

Distribution Location File Server (AS)

This is the file sharing service hosting the ManageSoftDL$ file share for package distribution.

Be aware:
NTFS file permissions on the underlying directories for the NTLM shares also apply through NTLM share access. Access checks for both the share permissions and the NTFS permissions must succeed before a user is granted access. The settings here should therefore be considered in conjunction with the appropriate NTFS settings (see *Distribution location (AS)*).

Access must be granted to the following types of users:

- **Administration server scheduled task user**
  Read/write access is required, as the Distribution Agent (which writes data to this location) is usually invoked in that context

- **Managed device users**
  This is the location from which managed device users download packages, and read access is required.

Default location	`file:////{hostname}/ManageSoftDL$`	
Settings location	Directory properties, Sharing, Permissions	
Default permissions	**User**	**Access rights**
	Everyone	Full control
Minimum permissions	**User**	**Access rights**
	Everyone	Full control

Distribution Location FTP Server (AS)

This is the FTP server hosting the ManageSoftDL FTP share for package distribution.

Be aware:
NTFS file permissions on the underlying directories for the IIS virtual directories also apply through FTP access. Access checks for both the share permissions and the NTFS permissions must succeed before a user is granted access. The settings here should therefore be considered in conjunction with the appropriate NTFS settings (see *Distribution location (AS)*).
This also means that for the "anonymous" FTP user to work, the user configured in IIS as the anonymous user must also be granted access through the NTFS permissions.

Access must be granted to the following types of users:

- **Administration server scheduled task user**
  Read/write access is required, as the Distribution Agent (which writes data to this location) is usually invoked in that context

- **Managed device users**
  This is the location from which managed device users download packages, and read access is required.

Default location	`ftp://{hostname}/ManageSoftDL`
Settings location	Internet Services Manager
Default permissions	Read, write
Minimum permissions	Read, write

Distribution Location Web Server (AS)

This is the web server hosting the ManageSoftDL web share for package distribution.

Be aware:
NTFS file permissions on the underlying directories for the IIS virtual directories also apply through HTTP and HTTPS access. Access checks for both the share permissions and the NTFS permissions must succeed before a user is granted access. The settings here should therefore be considered in conjunction with the appropriate NTFS settings (see *Distribution location (AS)*).

Access must be granted to the following types of users:

- **Administration server scheduled task user**
  Read/write access is required, as the Distribution Agent (which writes data to this location) is usually invoked in that context

- **Managed device users**
  This is the location from which managed device users download packages, and read access is required

Default location	`http://{hostname}/ManageSoftDL`
Settings location	Internet Services Manager
Default permissions	Read, write
Minimum permissions	Read, write

Incoming File Server (AS)

File sharing service hosting the ManageSoftRL$ file share for distribution and managed device logs for uploading.

Be aware:
NTFS file permissions on the underlying directories for the NTLM shares also apply through NTLM share access. Access checks for both the share permissions and the NTFS permissions must succeed before a user is granted access. The settings here should therefore be considered in conjunction with the appropriate NTFS settings (see *Incoming (AS)*).

The following components require access to this location:

- **Managed devices users**
  Upload logs and inventories to this location

- **Child distribution servers**
  Upload logs and inventories to this location

- **All administration server event processors**
  Update the RayManageSofti database with data from this location. When the logs/inventories are successfully processed, the files are removed from this directory

- **Distribution wizard**
  Writes distribution logs to this location

- **Distribution agent**
  Writes distribution logs to this location

Default location	`file:////{hostname}/ManageSoftRL$`	
Settings location	Directory properties, Sharing, Permissions	
Default permissions	**User**	**Access rights**

	Everyone	Full control
Minimum permissions	**User**	**Access rights**
	Everyone	Full control

Incoming FTP Server (AS)

This is the FTP server hosting the ManageSoftRL FTP share for distribution and managed device logs for uploading.

Be aware:
NTFS file permissions on the underlying directories for the IIS virtual directories also apply through FTP access. Access checks for both the share permissions and the NTFS permissions must succeed before a user is granted access. The settings here should therefore be considered in conjunction with the appropriate NTFS settings (see *Incoming (AS)*).

The following components require access to this location:

- **Managed devices users**
  Upload logs and inventories to this location
- **Child distribution servers**
  Upload logs and inventories to this location
- **All administration server event processors**
  Update the RayManageSofti database with data from this location. When the logs/inventories are successfully processed, the files are removed from this directory
- **Distribution wizard**
  Writes distribution logs to this location
- **Distribution agent**
  Writes distribution logs to this location

Default location	`ftp://{hostname}/ManageSoftRL`
Settings location	Internet Services Manager
Default permissions	Read, write
Minimum permissions	Read, write

Incoming Web Server (AS)

This is the web server hosting the ManageSoftRL HTTP or HTTPS share for distribution and managed device logs for uploading.

Be aware:
NTFS file permissions on the underlying directories for the IIS virtual directories also apply through HTTP or HTTPS access. Access checks for both the share permissions and the NTFS permissions must

succeed before a user is granted access. The settings here should therefore be considered in conjunction with the appropriate NTFS settings (*Incoming (AS)*).

The following components require access to this location:

- **Managed devices users**
  Upload logs and inventories to this location.

- **Child distribution servers**
  Upload logs and inventories to this location.

- **All administration server event processors, including IIS handlers**
  Update the RayManageSofti database with data from files uploaded to this location.

  Files handled by IIS handlers are only stored in this directory if they encounter transient processing failures. Re-processing is attempted by a scheduled task running the data importer (`mgsimport.exe`).

  Files processed by other data importers are stored in this directory until processed, and removed after successful processing.

- **Distribution wizard**
  Writes distribution logs to this location.

- **Distribution agent**
  Writes distribution logs to this location.

Default location	`http://{hostname}/ManageSoftRL`
Settings location	Internet Services Manager
Default permissions	Read, write
Minimum permissions	Read, write

RayManageSofti Data File Server

This is the file server that allows remote RayManageSofti administration consoles to access data and files from the central RayManageSofti administration server.

Be aware:
NTFS file permissions on the underlying directories for the NTLM shares also apply through NTLM share access. Access checks for both the share permissions and the NTFS permissions must succeed before a user is granted access. The settings here should therefore be considered in conjunction with the appropriate NTFS settings (see *RayManageSofti Data (ManageSoft$)*).

Default location	`file:////(hostname)/ManageSoft$`	
Settings location	Directory properties, Sharing, Permissions	
Default permissions	**User**	**Access rights**
	Everyone	Full control

Minimum permissions	User	Access rights
	Everyone	Read

RayManageSofti Job Queue Server

This is the file server that allows remote RayManageSofti administration consoles to access the job queue of the central RayManageSofti administration server.

Be aware:
NTFS file permissions on the underlying directories for the NTLM shares also apply through NTLM share access. Access checks for both the share permissions and the NTFS permissions must succeed before a user is granted access. The settings here should therefore be considered in conjunction with the appropriate NTFS settings (see *Job queue (AS)*).

Default location	file:////(hostname)/ManageSoftJQ$	
Settings location	Directory properties, Sharing, Permissions	
Default permissions	User	Access rights
	Everyone	Full control
Minimum permissions	User	Access rights
	Everyone	Read

RayManageSofti Services

This is the web server that hosts Microsoft .NET Web Services. These services provide remote RayManageSofti administration consoles with the business logic functionality from the central RayManageSofti administration server.

Be aware:
The settings here should be considered in conjunction with the NTFS permissions required to run these services. Checks for NTFS permissions must succeed before a user is able to run these services.

Default location	http://(hostname)/ManageSoftServices
Settings location	Internet Services Manager
Default permissions	Integrated Windows authentication
Minimum permissions	Integrated Windows authentication

Package Stage File Server (AS)

This is the file sharing service hosting the ManageSoftDS$ file share for package distribution.

Be aware:
NTFS file permissions on the underlying directories for the NTLM shares also apply through NTLM share access. Access checks for both the share permissions and the NTFS permissions must succeed before a user is granted access. The settings here should therefore be considered in conjunction with the appropriate NTFS settings (see *Package stage (AS)*).

This location is where child distribution servers download data from the administration server.

Default location	`file:////{hostname}/ManageSoftDS$`	
Settings location	Directory properties, Sharing, Permissions	
Default permissions	**User**	**Access rights**
	Everyone	Full control
Minimum permissions	**User**	**Access rights**
	Everyone	Read

Package Stage FTP Server (AS)

This is the FTP server hosting the ManageSoftDS FTP share for package distribution.

Be aware:
NTFS file permissions on the underlying directories for the IIS virtual directories also apply through FTP access. Access checks for both the share permissions and the NTFS permissions must succeed before a user is granted access. The settings here should therefore be considered in conjunction with the appropriate NTFS settings (see *Package stage (AS)*).

This also means that for the "anonymous" FTP user to work, the user configured in IIS as the anonymous user must also be granted access through the NTFS permissions.

This location is where child distribution servers download data from the administration server.

Default location	`ftp://{hostname}/ManageSoftDS`
Settings location	Internet Services Manager
Default permissions	Read, write
Minimum permissions	Read

Package Stage Web Server (AS)

This is the web server hosting the ManageSoftDS web share for package distribution.

Be aware:
NTFS file permissions on the underlying directories for the IIS virtual directories also apply through HTTP and HTTPS access. Access checks for both the share permissions and the NTFS permissions must succeed before a user is granted access. The settings here should therefore be considered in

conjunction with the appropriate NTFS settings (see *Package stage (AS)*).

This location is where child distribution servers download data from the administration server.

Default location	`http://{hostname}/ManageSoftDS`
Settings location	Internet Services Manager
Default permissions	Read, write
Minimum permissions	Read

Remote Execution Task File Server (AS)

This is the file sharing service hosting the ManageSoftREA$ file share for storing remote execution actions prior to their retrieval by managed devices.

Be aware:
NTFS file permissions on the underlying directories for the NTLM shares also apply through NTLM share access. Access checks for both the share permissions and the NTFS permissions must succeed before a user is granted access. The settings here should therefore be considered in conjunction with the appropriate NTFS settings (see *RemoteExecution\Actions (AS)*).

The remote execution scheduler agent requires access to this location.

Default location	`file:////{hostname}/ManageSoftREA$` `file:////{hostname}/mgsREA$`	
Settings location	Directory properties, Sharing, Permissions	
Default permissions	**User**	**Access rights**
	Everyone	Read
Minimum permissions	**User**	**Access rights**
	Everyone	Read

Remote Execution Task FTP Server (AS)

This is the FTP server hosting the ManageSoftREA FTP share for storing remote execution actions prior to their retrieval by managed devices.

Be aware:
NTFS file permissions on the underlying directories for the IIS virtual directories also apply through FTP access. Access checks for both the share permissions and the NTFS permissions must succeed before a user is granted access. The settings here should therefore be considered in conjunction with the appropriate NTFS settings (see *RemoteExecution\Actions (AS)*).

The remote execution scheduler agent requires access to this location.

Default location	`ftp://{hostname}/ManageSoftREA`
Settings location	Internet Services Manager
Default permissions	Read
Minimum permissions	Read

Remote Execution Task Web Server (AS)

This is the web server hosting the ManageSoftREA web share for storing remote execution actions prior to their retrieval by managed devices.

Be aware:
NTFS file permissions on the underlying directories for the NTLM shares also apply through HTTP and HTTPS access. Access checks for both the share permissions and the NTFS permissions must succeed before a user is granted access. The settings here should therefore be considered in conjunction with the appropriate NTFS settings (see *RemoteExecution\Actions (AS)*).

The remote execution scheduler agent requires access to this location.

Default location	`http://{hostname}/ManageSoftREA`
Settings location	Internet Services Manager
Default permissions	Read
Minimum permissions	Read

Remote Execution Tools Server (AS)

This is the file sharing service hosting the ManageSoftRET$ file share. This share is used to store the remote execution client agent, and the programs that run managed device adoptions and inventories.

Be aware:
NTFS file permissions on the underlying directories for the NTLM shares also apply through NTLM share access. Access checks for both the share permissions and the NTFS permissions must succeed before a user is granted access. The settings here should therefore be considered in conjunction with the appropriate NTFS settings (see *RemoteExecution\Public (AS)*).

The remote execution subsystem requires access to this location.

Default location	`file:////{hostname}/ManageSoftRET$` `file:////{hostname}/mgsRET$`	
Settings location	Directory properties, Sharing, Permissions	
Default permissions	**User**	**Access rights**
	Everyone	Read
	Anonymous logon	Read

	Network	Read
Minimum permissions	**User**	**Access rights**
	Everyone	Read
	Anonymous logon	Read
	Network	Read

Reporting Web Server

RayManageSofti reporting is an application hosted on a Microsoft Internet Information Server (IIS) website. The site is referenced through the address `http://{hostname}/ManageSoftRP`, which by default points to the physical location `C:\Program Files\ManageSoft\Reporter\Web`. This directory contains the scripts and other elements that create the reporting framework.

Be aware:
Data to populate the reports is drawn from the RayManageSofti database.

- Using basic authentication: When the reporting code queries the RayManageSofti database, IIS passes to SQL Server the credentials that the user provided when logging in to the site.
- Using anonymous access: When the reporting code queries the RayManageSofti database, IIS passes to SQL Server the username and password you registered during reporting setup (see *Authentication methods for reporting access*).

By default, during reporting setup the same users will be granted permissions for database reads. For more information on database access, see *RayManageSofti database*.

Also be aware:
Database queries are filtered by information about Organizational Units retrieved from Active Directory. Therefore the same users also require read access to Active Directory. This is set through normal Active Directory processes (see your Active Directory documentation for more information).

For further clarification, see the Generate reports process in the Administration server data flow.

Default location	`http://{hostname}/ManageSoftRP` (points to `C:\Program Files\ManageSoft\Reporter\Web`)
Settings location	Internet Services Manager
Default configuration	Basic authentication or anonymous access, as selected during reporting setup (see *Authentication methods for reporting access*). **Note**: Other authentication methods supported by IIS (such as integrated Windows authentication) are not supported by Deployment Manager reporting.

Tools File Server (AS)

This is the file sharing service hosting the ManageSoftTL$ file share for common tool access.

Be aware:
NTFS file permissions on the underlying directories for the NTLM shares also apply through NTLM share access. Access checks for both the share permissions and the NTFS permissions must succeed before a user is granted access. The settings here should therefore be considered in conjunction with the appropriate NTFS settings (see *Tools*).

This directory contains common tools that the RayManageSofti administrator may want to access from other computers. It is provided for convenient access. RayManageSofti does not use the share for any purpose. Since this location contains core RayManageSofti system files, it is advisable to limit access to administrators only.

Default location	`file:////{hostname}/ManageSoftTL$`	
Settings location	Directory properties, Sharing, Permissions	
Default permissions	**User**	**Access rights**
	Everyone	Read
Minimum permissions	**User**	**Access rights**
	Everyone	Read

Web Controls Web Server

Web server hosting the Microsoft Office Web Controls, used by the AD Browser in Reporting. The users that access reports must have access to this component. However, be aware that Microsoft security updates for web controls may alter existing permissions, which may cause the reports to stop working.

Be aware:
NTFS file permissions on the underlying directories for the IIS virtual directories also apply through HTTP and HTTPS access. Access checks for both the share permissions and the NTFS permissions must succeed before a user is granted access. The settings here should therefore be considered in conjunction with the appropriate NTFS settings.

Default location	`http://{hostname}/webctrl_client`
Settings location	Internet Services Manager
Default configuration	Basic authentication, anonymous access, or integrated Windows authentication as configured by web controls installation.
Minimum configuration	If the reporting web server uses anonymous access, then the web controls web server also requires anonymous access. Otherwise the minimum requirement is basic authentication.

Configuring Distribution Server Rights

The Deployment Manager for distribution servers software is a key element in any security configuration involving RayManageSofti. The data stored here is directly accessed by all of the managed device users. Therefore access to the data needs to be carefully controlled in order to maintain the data's integrity.

This chapter describes the default rights and permissions configuration for RayManageSofti distribution servers, as well as the minimum permissions settings that you can implement for RayManageSofti to function as a whole.

It also describes interactions between various RayManageSofti features and agents, and the security permissions required to enable these interactions.

Rights and permissions settings include username/password logon authentication (for example, for database access or access to a FTP server) and file system permissions controlling access to files, directories, and other file system elements such as shared drives.

The layout of this information is described in *Configuring RayManageSofti rights*.

Be aware:
The same name may be used for security settings on the administration server, distribution servers, and managed devices. Where this occurs, the abbreviation AS, DS, or MD appears after the name of the setting.

Distribution Server Overview

The distribution server is an extension of the administration server. It runs the same replication process (used in distribution of data) as the administration server. As such, the same **data linkage** security applies, with some additions:

- The distribution server has an upload process which uploads logs to the parent distribution server or administration server, and runs as a scheduled task
- The distribution server creates Wake on LAN tasks for distribution as necessary
- RayManageSofti administrators can specify usernames and passwords for use when executing tasks remotely on managed devices

Other than the user interface for specifying usernames and passwords, the distribution server has no user interface. The purpose of distribution is to automatically **replicate data** from the administration server to distribution locations (leaf nodes of the distribution hierarchy) where it can be accessed by managed devices.

The replicate data process is started by a scheduled task. The user context of the scheduled task must have access to the RayManageSofti files on the distribution server. The only entry points into the system are from parent servers (distribution or administration) when new job requests arrive or from the distribution server itself when it polls the parent server for new jobs. The entry point is determined by the job retrieval method selected for the distribution server.

On installation, the distribution server can be configured to authenticate the originator of the requested job. If the option is selected, the **replicate data** process will only accept jobs from users whom are in the **MGS Distribution** group.

If the child server of this server uses the same authentication, the user account of the scheduled task that starts the replicate data process must be a member of the **MGS Distributors** group.

In large implementations, the distribution server can also **resolve data** into the RayManageSofti database, in order to help minimize load on the administration server. To better understand this process, read the Chapter *Configuring administration server rights*, especially the details about the **MGS Data Modifiers** group and permissions required to access the RayManageSofti database.

Default Security in Windows Environments

To provide a fully functioning system that can operate without major rights modifications, read access is set at a low level for most elements. For example, most data files/locations in RayManageSofti are readable by everyone. Full control, however, is more tightly restricted on initial installation.

For example, the two shares used for the distribution system (**ManageSoftDS** and **ManageSoftDL**) give read access to everyone by default, but full control is restricted to `SYSTEM` and members of the **Administrators** and **MGS Administrators** groups. This ensures that anybody who can access these shares can read the packages that have been distributed, but only restricted personnel have the ability to remove or otherwise modify or damage the packages that managed device users will download. Further protection is offered by digital signing and MD5 cryptographic digests.

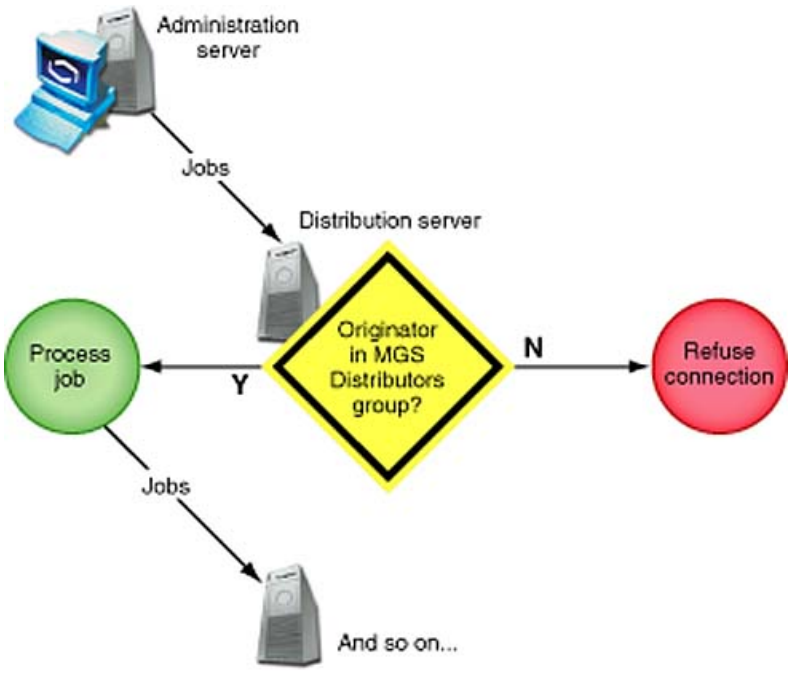

Distribution Server Installation

During the installation of the `RayManageSofti` for distribution servers software, there are a number of security configuration controls to be set.

On Windows distribution servers, there are three controls:

- *Authenticating parent connections*

- *User name for scheduled tasks*
- *Job retrieval method*

You set these controls using the install wizard.

Authenticating Parent Connections

The first control on Windows distribution servers is to choose whether connections from parent distribution servers (or the administration server) to the connection agent should be authenticated or not. You can only set this if you select a **Custom** installation.

It is set on the **Authentication** page.

If you choose to authenticate the parent server, authentication is performed by the connection agent when a parent distribution server (or the administration server) attempts to establish a connection.

The connection agent requests the user name which owns the process that is attempting the connection. It then checks membership of the **MGS Distributors** security group in the current domain (or the domain of the parent server if this is specified) to ensure that the connecting user is a member of the security group. If this is true, then the connection is allowed. Otherwise the connection is rejected.

The authentication stops invalid users from accessing the connection agent and potentially inserting malicious or invalid distribution jobs into the distribution hierarchy.

For more information about parent authentication, see the discussion under *Connection agent*.

User Name for Scheduled Tasks

The second security control is the name under which the scheduled tasks on the distribution server (distribution agent and upload agent) should be run.

This can be set on the **Configure Scheduled Tasks** page, for both **Complete** and **Custom** installations.

The user that is selected must have all the permissions required for the **Deployment Manager Distribution Scheduled Task Users** on this distribution server, the parent distribution server, and all direct child distribution servers and distribution locations.

If you choose to run scheduled tasks as `SYSTEM` user, you must reboot the distribution server after installation to allow the distribution server to recognize its security group membership.

Job Retrieval Method

The third control is the job retrieval method. This control determines how distribution is initiated on the distribution server:

- It can **listen** for connection attempts from the parent distribution server
- It can use the **job server program** to periodically **poll** the parent distribution server for distribution jobs, if a web server is installed on the parent distribution server.

The selection you make affects the way that the distribution agent operates, as described in *Distribution agent (DS)*.

For Windows distribution servers, the job retrieval method is set on the **Job Retrieval Method** page.

Be aware:
The job retrieval method, protocol, parent distribution server, and ports can also be set in the `mgssetup.ini` file used to bootstrap a managed distribution server.

Distribution Server Features

In this section we will examine a number of distribution server features in detail.
These include:

- *Connection agent*
- *Discovery agent (DS)*
- *Distribution agent (DS)*
- *Job server program (DS)*
- *Password store manager*
- *Remote execution server agent (DS)*
- *Upload agent (DS)*
- *Wake on LAN packet generator program*

Connection Agent

This component implements part of the **Replicate data** process in the distribution server data flow (see next page).

Deployment Manager for distribution servers uses a service (the **Deployment Manager Connection Agent**) to accept distribution jobs from parent distribution servers.

In this context, the administration server is also a distribution server.

The connection agent accepts connections on a specific network port (by default, this port is `7010`) and places the received jobs into the job queue for processing.

By default, this connection is not secure. This means any computer and user can pretend to be a distribution server and submit potentially malicious jobs into the job queue. On Windows distribution servers, enabling parent server authentication secures this connection so only known parent distribution servers are allowed to connect.

It is enabled during installation (through the **Custom** installation option) on the following dialog in the **Deployment Manager for distribution servers - Install Wizard**.

The following distribution hierarchy diagram will be used as a reference to illustrate various points in the description below. In the diagram, the user specified for each server represents the user under which the RayManageSofti distribution agent scheduled task runs.

This discussion assumes that parent server authentication is enabled.

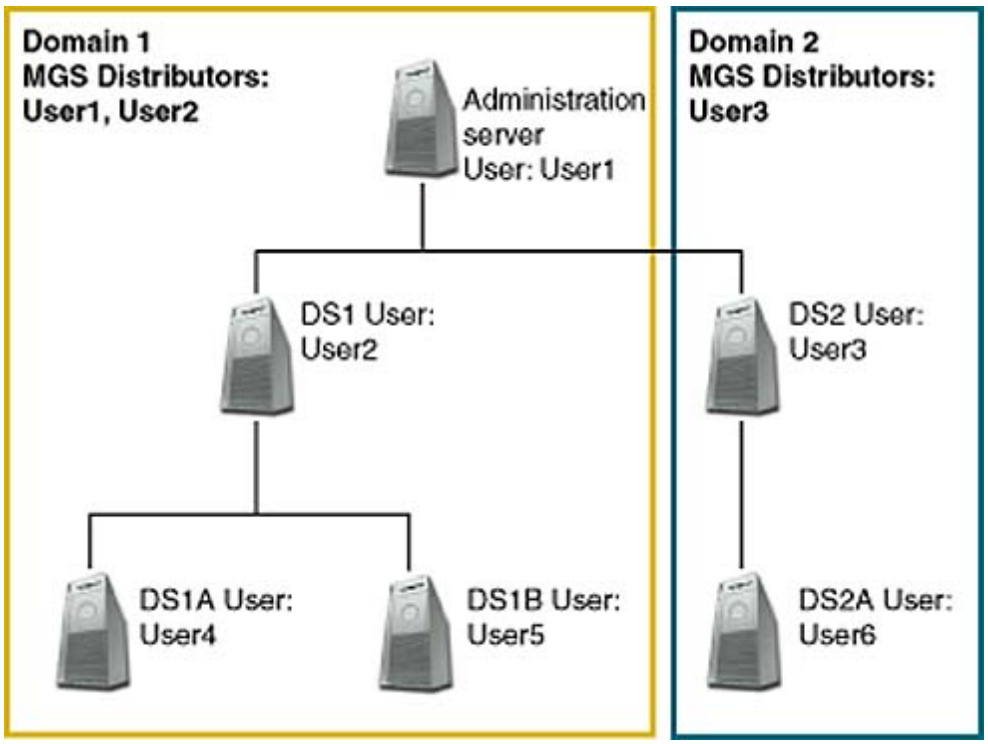

How Distribution Servers Enable Connection

Whenever the connection agent receives a connection, a check is performed to determine if this connection should be allowed. The check involves querying the user of the process at the other end of the connection and determining if the user is a member of the **MGS Distributors** security group.

This means that if the distribution server is in a different Windows domain (for example DS2 in the diagram is in Domain 2), the security group has to be manually created using (for example) **Active Directory Users and Computers**.

The scheduled task user will also need to be added to this group.

The scheduled task users for distribution servers that are leaf nodes in the hierarchy (for example DS1A, DS1B and DS2A in the diagram), do not need to be added to the security group. This is because these distribution servers do not attempt connections to any child distribution servers.

The user being checked is the user under which the RayManageSofti distribution agent scheduled task is running on the parent distribution server, not the local distribution server. (The parent can also be the administration server.)

The check is performed on the **MGS Distributors** security group in the current Windows domain. If the parent distribution server is in another domain (for example the administration server relative to DS2), the check needs to be performed on the **MGS Distributors** security group in that domain (Domain 1 in this example).

In this case this domain name should be specified in the appropriate field on the **Deployment Manager for distribution servers - Install Wizard** page. Appropriate trust relationships will also need to be established between the domains so that the security group membership check can be performed.

The **Deployment Manager for distribution servers - Install Wizard** will automatically create the **MGS Distributors** security group in the Windows domain where the administration server belongs (Domain 1 in the

diagram).

The Install Wizard will add the user under which the Deployment Manager distribution agent scheduled task runs on the administration server to the security group (User1 in the diagram).

The **Deployment Manager for distribution servers - Install Wizard** will not automatically create the security group or add the scheduled task user to the security group.

Be aware:
Users need to be manually added to the **MGS Distributors** security group. This is not done automatically during installation, as it requires knowledge of the distribution hierarchy, which is not known at this time.

Also be aware:
After adding a user to the **MGS Distributors** security group, you must reboot the distribution server to enable it to recognize the changes to security group membership.

On Windows distribution servers, the settings for parent server authentication can be altered after installation through the following registry keys, located under:

```
[Registry]\ManageSoft\Replication Agent\CurrentVersion
```

Key	Value	Purpose
AuthenticationEnabled	True or False	Determines whether parent server authentication is enabled or not
DistributorsGroup	MGS Distributors	The security group name to check against
ParentServerDomain	Windows Domain Name	The domain name for the parent distribution server (typically the current domain name)

Once changed, the new settings will take effect the next time a connection is attempted to this distribution server. Any distribution jobs currently in progress are unaffected by the changes because the connection has already been established. No scheduled tasks or services need to be stopped or restarted.

To verify the settings have been successfully applied, a verification of the distribution hierarchy should be performed.

To perform this:

1. Select the **Infrastructure** node in the RayManageSofti console and click on the **Open infrastructure editor** button from within the details pane.

2. Right-click the **distribution servers** node within the infrastructure editor tree.

3. Select **Verify Distribution Servers** configuration.

4. From the **Reports** node, view the **Distribution Servers report** to confirm a successful configuration. All distribution servers should be marked as configured once the verification is complete.

How Parent Authentication Works

To illustrate how parent server authentication works, consider the distribution of a package from the administration server to all of the distribution servers in the hierarchy from the diagram above.

First of all a distribution job will be added to the job queue on the administration server. The Deployment Manager distribution agent scheduled task will process the distribution job.

It will do this by contacting DS1 first of all and connecting to the connection agent. The connection agent on DS1 will query the user making the connection (User1 in this case). In Windows environments, it will check whether the user is a member of the **MGS Distributors** security group in Domain 1. Upon a successful check, the connection will be accepted and the job will be added to the job queue on DS1.

The same process will be followed to pass the job on to DS1A and DS1B, however the check will be performed upon User2.

The administration server then attempts to contact DS2, again connecting to the connection agent. In this case the connection agent on DS2 will query the user making the connection (User1 again).

A check whether this user is a member of the **MGS Distributors** security group in Domain 1 will then be performed.

As this is a different domain to the domain of DS2, the **ParentServerDomain** Registry key on DS2 will need to specify Domain 1. Upon a successful check, the connection will be accepted and the distribution job added to the job queue on DS2.

The same process will be followed to pass the job onto DS2A, however the check will be performed upon User3 in the **MGS Distributors** security group in Domain 2.

Parent Authentication Failures

When a parent distribution server cannot be authenticated by one of its child distribution servers, then the current job (whether distribute, remove, synchronize or configure) will be marked as failed in the reports with one of the following reasons given:

Reason	Explanation
Remote user `<username>` is not a member of security group `<groupname>` and is therefore not authorized to connect.	(Windows distribution servers only) The user under which the scheduled task is running on the parent distribution server, is not a member of the **MGS Distributors** (or other specified name) security group. Therefore the connection is denied. To correct this error, add the user to the security group. **Note**: This error may also occur if you have not rebooted a distribution server after adding a user to the **MGS Distributors** security group. Reboot the distribution server to enable it to recognize the change in security group membership.
Failed to negotiate authentication on connection from server `<hostname>`.	This is a general message to indicate that the connection from the given hostname failed. A more specific error message will also be provided.

Reason	Explanation
Authentication is required. A message containing no authentication details was received from the server `<hostname>`, running an earlier version of RMS. A server software upgrade is necessary to support authentication.	This distribution server (version 11.4 *infinity* or later), with parent server authentication enabled, received a connection attempt from a parent distribution server of an earlier version. To enable authentication, the software on the parent distribution server will need to be upgraded.
Failed to look up security group `<groupname>`: error `<errorcode>`.	(Windows distribution servers only) The **MGS Distributors** (or other specified security group name) could not be found. The extended Windows error code is also specified. To correct this error, check the definition of the security group.
Group `<groupname>` was found in domain `<domainname>`, not in domain `<domainname>`	The security group specified was found in a different domain to that specified as the parent server domain. To correct this error, check the definition of the security group.
User `<username>` is not a member of group `<groupname>`	The user under which the scheduled task is running on the parent distribution server is not a member of the **MGS Distributors** (or other specified name) security group. Therefore, the connection is denied. To correct this error, add the user to the security group.

Aspect	Explanation
Program	`ndlisten.exe`
Users	Deployment Manager Connection Agent Service Users (Windows)
Data files / locations	None
Communications	Distribution Job Propagate
Databases	None
Servers	None

Discovery Agent (DS)

The discovery agent program, `mgsdisco.exe`, is the implementation of the **Discover devices** process in the distribution server data flow. It is triggered by scheduled tasks that were previously set up by the **Discovery wizard**. Refer to the *Discovery* chapter of the *RMS Discovery* for more details.

Distribution Agent (DS)

The distribution agent works in much the same way as the administration server distribution agent (see *Distribution agent (AS)*), although the data flow is slightly different because it is not the root node in the distribution hierarchy.

The `ndrspawn.exe` and `ndrplag.exe` executables are the implementations of the Replicate data process in the distribution server data flow. `ndrspawn.exe` is a long running process that polls the job queue for new jobs.

It is initially started by a scheduled task in the user context specified during installation of the administration server (see *User name for scheduled tasks*). When new jobs are added to the queue, `ndrspawn.exe` invokes `ndrplag.exe` to process the jobs.

After a job is processed, it is removed from the queue by `ndrplag.exe`. This implies that the user credentials (set in the scheduled task) must have read and write access to the job queue.

Be aware:
When using Wake on LAN on a Windows distribution server, the account the replication agent is run as must be able to create scheduled tasks as `LocalSystem`. For more information, see *Wake on LAN packet generator program*.

The task of replicating packages to distribution locations, and propagating jobs to distribution servers are performed by `ndrplag.exe`. The replication process is as follows:

1. Receive distribution jobs from the parent server.

2. Check the distributed objects database (see *Distributed objects database*) for a version of each package that may already exist on this server. If a suitable one exists, only differences between the package on the parent server and the package on this server will be downloaded.

3. Download, from the parent server, all packages as specified in the distribution job.

4. Retrieve the distribution configuration to determine the current system state.

5. Check the smart distribution database to determine if the packages need to be distributed to local distribution locations.

6. Push the packages out (from the package stage to distribution locations) if necessary and update the smart distribution database.

7. If the child distribution server listens for jobs, the distribution agent on the parent server contacts the connection agent on the child server to initiate the distribution agent on the child server. The parent distribution server creates distribution jobs and propagates them to all child servers.

   If the child distribution server polls for jobs, the parent distribution server stores jobs for the child server in a separate job queue for the server. Polling is initiated by the distribution agent on the child distribution server, and the job server program on the parent distribution server sends the jobs in the job queue to the child server.

8. All errors and successes are logged to the **Incoming** directory.

   All data repositories (database, data files and directories) accessed by this process are listed in the table below.

Aspect	Explanation
Program	`ndrspawn.exe, ndrplag.exe` (Windows)
Users	Deployment Manager Distribution Scheduled Task Users (Windows)
Data files / locations	• Package Stage • Distribution Location • Incoming • Job Queue

Aspect	Explanation
	• Wake on LAN • Distribution Configuration • Distribution Cache
Communications	• Distribution Job Propagate • HTTP or HTTPS Distribution Transfer • FTP Distribution Transfer • NTLM Distribution Transfer
Databases	• Smart Distribution Database • Distributed Objects Database
Servers	• Package Stage Web Server • Package Stage FTP Server • Package Stage File Server • Distribution Location Web Server • Distribution Location FTP Server • Distribution Location File Server

Job Server Program (DS)

The job server program `mgsjobsrv.exe` responds to job requests from child distribution servers that are configured to poll for distribution jobs. It is called through an HTTP/HTTPS request that contains the UID of the child distribution server and MD5 details for components of the distribution hierarchy that the child distribution server already knows about.

In response to the call from the child distribution server, `mgsjobsrv.exe` looks for a job queue file for the appropriate UID. It then reads out each job. If the MD5 associated with the job has not been received in the HTTP/HTTPS request, `mgsjobsrv.exe` creates a new job to pass the relevant hierarchy details to the child distribution server.

Aspect	Explanation
Program	`mgsjobsrv.exe`
Users	Deployment Manager Distribution Scheduled Task Users (Windows)
Data files / locations	Job Queue
Communications	Distribution Job Poll
Databases	None
Servers	None

Password Store Manager

The password store manager program (`mgspswdw.exe`) is executed by RayManageSofti administrators to create a store of usernames and passwords for use when remotely executing tasks on managed devices. The password store is described more fully in the *Discovery* chapter of the *RMS Discovery*.

Aspect	Explanation
Program	`mgspswdw.exe`
Users	Deployment Manager administrator (Windows)
Data files / locations	On Windows distribution servers, passwords are encrypted using 52-bit DES encryption, and stored in the Windows registry. The key used for encryption is generated the first time any password is added. It is stored as a private data object in the registry. The key is accessible by any user with Administrator privileges.
Communications	None
Databases	None
Servers	None

Remote Execution Server Agent (DS)

The remote execution server agent, `mgsresa.exe`, is the implementation of the **Remotely execute tasks on managed devices** process in the distribution server data flow. It is triggered by scheduled tasks that were previously set up by the Discovery or Remote Command wizards. Refer to the *Discovery* chapter of the *RMS Discovery* for more details.

Aspect	Explanation
Program	`mgsresa.exe` (Windows)
Users	Deployment Manager Remote Execution Server Agent Users (Windows)
Data files / locations	• JobQueue • RemoteExecution
Communications	None
Databases	None
Servers	None

Upload Agent (DS)

This component is the implementation of the **Upload to parent server** process in the distribution server data flow. It passes distribution server and managed device logs and inventories, stored in Incoming to the parent distribution server. Once the logs are successfully uploaded, the files are removed from the Incoming directory.

This component is configured from the administration server, using client data settings. These settings tell child distribution servers and managed devices how and where logs should be propagated back up the distribution chain.

Aspect	Explanation
Program	`ndupload.exe` (Windows)
Users	Deployment Manager Uploader Scheduled Task Users (Windows)

Aspect	Explanation
Data files / locations	Incoming
Communications	• FTP Upload Transfer • HTTP or HTTPS Upload Transfer • NTLM Upload Transfer
Databases	None
Servers	• Incoming FTP Server on the administration server and distribution servers • Incoming Web Server on the administration server and distribution servers (using the HTTP/HTTPS protocol's efficient PUT method to receive files instead of the more common POST method) • Incoming File Server on the administration server and distribution servers

Wake on LAN Packet Generator Program

The Wake on LAN packet generator program (`mgswolgen.exe`) is executed by the scheduled task created to perform a Wake on LAN task.

Alternatively, if that scheduled task is scheduled to occur as soon as possible or was scheduled in the past, it is immediately executed by the distribution agent when new Wake on LAN data is received.

The packet generator program:

- Reads details of the Wake on LAN task from the file stored in the Wake on LAN tasks data location.
- Attempts to wake all the targeted managed devices.
- Writes a log of results into the **Incoming** directory. This log is uploaded for processing on the administration server.

Aspect	Explanation
Program	`mgswolgen.exe` (Windows)
Users	Deployment Manager Distribution Scheduled Task Users
Data files / locations	• Wake on LAN tasks • Incoming
Communications	None
Databases	None
Servers	None

Distribution Server Data Files / Locations

The key data storage locations on distribution servers are:

- *Distribution cache (DS)*
- *Distribution configuration (DS)*

- *Distribution location (DS)*
- *Incoming (DS)*
- *Job queue (DS)*
- *ManageSoftJQ (DS)*
- *Package stage (DS)*
- *Remote execution actions (DS)*
- *Remote execution tools (DS)*
- *Wake on LAN tasks (DS)*

Distribution Cache (DS)

The storage location for files during download before they are moved to the package stage. Files are copied from here to the package stage. Rather than inheriting the permissions of the package stage, the copied files retain their distribution cache permissions (NTFS permissions).

Aspect	Description
Default location	`C:\ManageSoft\Distribution`
Settings location	• Directory properties • Security
Default Windows permissions	As for the distribution location and package stage (the combined settings of both)
Minimum Windows permissions	As for the distribution location and package stage (the combined settings of both)

Distribution Configuration (DS)

The distribution configuration is the storage location for all currently active distribution hierarchies. This includes information on distribution servers and locations. The distribution agent writes new distribution hierarchy data to this directory and reads data from here when packages and jobs are replicated to child servers.

Default location	`C:\ManageSoft\Repository\DeploymentLocations\Common`	
Settings location	• Directory properties • Security	
Default Windows permissions	Default permissions are inherited from the root folder (`C:\` for Windows). Permissions set directly in this folder vary for different operating systems. They are not set by RayManageSofti.	
Minimum Windows permissions	**User**	**Access rights**
	MGS Administrators	Full control
	Deployment Manager Distribution Scheduled Task Users	Full control

Distribution Location (DS)

The storage location for all applications, policies, schedules and managed device settings that have been made available to managed devices.

Be aware:
If this is a linked distribution location, this will be the same physical location as the package stage location and therefore the permissions need to be a combination of the permissions for a package stage location and distribution location.

Access is required by the following components:

- **Managed device users**
  This is where they download packages from.
- **Distribution agent**
  Packages are delivered here by the agent.

Default location	`C:\ManageSoft\LocalDeployment` (if not linked) `C:\ManageSoft\Staging\Common` (if linked) shared as ManageSoftDL (FTP, HTTP, HTTPS) and ManageSoftDL$ (file)	
Settings location	• Directory properties • Security	
Default Windows permissions	**User**	**Access rights**
	Everyone	Full control
Minimum Windows permissions	**User**	**Access rights**
	Deployment Manager Distribution Scheduled Task Users	Full control
	Deployment Manager Managed Device User	Read, List folder contents
	Domain Computers	Read, List folder contents
	Domain Controllers	Read, List folder contents

The web server, FTP server, and file server settings should be considered in conjunction with these NTFS file permissions.

Incoming (DS)

This location is used by the distribution agent, child distribution servers, and managed devices to report data back to the system, and to store logs and inventories. Managed device users and child distribution servers must have write access to this location. In addition, the user context in which the various scheduled tasks are run must have full access to this location.

Default location	`C:\ManageSoft\Incoming`, shared as ManageSoftRL (FTP, HTTP, HTTPS) and

	ManageSoftRL$ (file)	
Settings location	• Directory properties • Security	
Default Windows permissions	**User**	**Access rights**
	Everyone	Full control
	Administrators	Full control
	Anonymous Logon	Full control
	Network	Full control
	System	Full control
Minimum Windows permissions	**User**	**Access rights**
	Deployment Manager Uploader Scheduled Task Users	Full control
	Deployment Manager Distribution Scheduled Task Users	Full control
	Domain Computers	Full control
	Domain Controllers	Full control

The web, FTP, and file server settings should be considered in conjunction with these NTFS file permissions.

Job Queue (DS)

The job queue is the storage location for all distribution jobs that are yet to be processed.

The distribution agent writes new jobs to the queue when they are received from the parent server. After a job has been processed, the agent removes it from the queue.

Default location	`C:\Documents and Settings\All Users\Application Data\ManageSoft Corp\ManageSoft\Replication Agent\JobQueue`	
Settings location	• Directory properties • Security	
Default Windows permissions	Default permissions are inherited from the parent folder `C:\Documents and Settings\All Users\Application Data`. Permissions set directly in this folder vary for different operating systems. They are not set by RayManageSofti.	
Minimum Windows permissions	**User**	**Access rights**
	MGS Administrators	Full control
	Deployment Manager Distribution Scheduled Task Users	Full control
	Deployment Manager Connection Agent Service Users	Full control

ManageSoftJQ (DS)

This virtual folder is created at installation time.

The following components require access to this location:

- **Job server program**
  Polls parent distribution server for distribution jobs.

Default location	Windows `C:\Program Files\ManageSoft\Replicator`	
Settings location	Internet Services Manager	
Default Windows permissions	**User**	**Access rights**
	Anonymous	Read, Execute permission on scripts and executables
Minimum Windows permissions	**User**	**Access rights**
	Anonymous	Read, Execute permission on scripts and executables

Package Stage (DS)

The storage location for all applications, policies, schedules and managed device settings that have been distributed by the administration server. Access is required by the following components:

- **Child distribution servers**
  Requires access for pull distribution (when the child servers downloads packages from the parent server).

- **Distribution agent**
  Writes incoming packages to this location, and where appropriate (if instructed to by the distribution job) copies the files to local distribution locations.

Default location	`C:\ManageSoft\Staging\Common`, shared as ManageSoftDS (FTP, HTTP, HTTPS) and ManageSoftDS$ (file)	
Settings location	• Directory properties • Security	
Default Windows permissions	**User**	**Access rights**
	Everyone	Full control
Minimum Windows permissions	**User**	**Access rights**
	Deployment Manager Distribution Scheduled Task Users	Full control
	MGS Distributors	Read, List folder contents

The web server, FTP server, and file server settings should be considered in conjunction with these NTFS file

permissions.

Remote Execution Actions (DS)

This is the location where remote execution actions are stored, prior to their retrieval by managed devices.

The physical folder to which this maps is accessed by the remote execution scheduler agent, which writes remote execution jobs to this location.

For more information about the remote execution subsystem, see the *Discovery* chapter of the *RMS Discovery*.

Default location	`C:\Program Files\ManageSoft\RemoteExecution\Actions\Public`	
Settings location	• Directory properties • Security	
Default Windows permissions	**User**	**Access rights**
	Administrators	Full control
	MGS Administrators	Full control
	System	Full control
Minimum Windows permissions	**User**	**Access rights**
	Administrators	Full control
	MGS Administrators	Full control
	System	Full control

The file server settings should be considered in conjunction with these NTFS file permissions.

Remote Execution Tools (DS)

This is the location where remote execution tools are stored, including the client agent, adoption executables, and remote inventory executables.

The physical folder to which this maps is accessed by the remote execution subsystem on managed devices:

- Managed devices run `mgsreca.exe` directly from the **ClientAgent** subfolder
- Managed devices run the adoption agent from the **Adoption** subfolder to install Deployment Manager for managed devices
- Managed devices run zero-touch inventory from the **Inventory** subfolder

For more information about the remote execution subsystem, see the *Discovery* chapter of the *RMS Discovery*.

Default location	`C:\Program Files\ManageSoft\RemoteExecution\Public`
Settings location	• Directory properties • Security

Default Windows permissions	**User**	**Access rights**
	Administrators	Full control
	Anonymous logon	Read and execute, list file contents, read
	Everyone	Read and execute, list file contents, read
	Network	Read and execute, list file contents, read
	System	Full control
Minimum Windows permissions	**User**	**Access rights**
	Administrators	Full control
	Anonymous logon	Read and execute, list file contents, read
	Everyone	Read and execute, list file contents, read
	Network	Read and execute, list file contents, read
	System	Full control

The file server settings should be considered in conjunction with these NTFS file permissions.

Wake on LAN Tasks (DS)

This location stores Wake on LAN task files. These files contain the list of devices to be woken.

Distribution Scheduled Task Users must have full access to this location.

Default location	Windows `C:\Documents and Settings\All Users\Application Data\ManageSoft Corp\ManageSoft\WakeOnLAN`	
Settings location	• Directory properties • Security	
Default Windows permissions	Default permissions are inherited from the parent folder `C:\Documents and Settings\All Users\Application Data`. Permissions set directly in this folder vary for different operating systems. They are not set by RayManageSofti.	
Minimum Windows permissions	**User**	**Access rights**
	Deployment Manager Distribution Scheduled Task Users	Full control

Distribution Server Communications

This section describes network connections that may be established between two distribution server features or a distribution server feature and a server:

- *Distribution job propagate (DS)*
- *Distribution job poll (DS)*
- *Distribution transfer HTTP/HTTPS (DS)*
- *Distribution transfer FTP (DS)*
- *Distribution transfer NTLM (DS)*
- *Upload transfer FTP (DS)*
- *Upload transfer HTTP/HTTPS (DS)*
- *Upload transfer NTLM (DS)*

Distribution Job Propagate (DS)

Performs the propagation of a distribution job from this distribution server to a child distribution server.

Aspect	Description
Initiator	Distribution agent
Acceptor	Connection agent
Protocol	Proprietary protocol. If parent server authentication is enabled, then the user running the distribution agent must be a member of the MGS Distributors security group.

Distribution Job Poll (DS)

Polls for distribution jobs from this distribution server to a parent distribution server.

Aspect	Description
Initiator	Distribution agent
Acceptor	Job server program
Protocol	HTTP or HTTPS, Windows integrated authentication (anonymous access allowed by default)

Distribution Transfer HTTP / HTTPS (DS)

Performs the transfer of data (files, directory listings) over the HTTP/HTTPS protocol. See *HTTPS security configuration* for further information about HTTPS security configuration

Aspect	Description
Initiator	Distribution agent

Aspect	Description
Acceptor	• Package stage web server • Distribution location web server
Protocol	HTTP/HTTPS, username/password authentication or Windows integrated authentication

Distribution Transfer FTP (DS)

Performs the transfer of data (files, directory listings) over the FTP protocol.

Aspect	Description
Initiator	Distribution agent
Acceptor	• Package stage FTP server • Distribution location FTP server
Protocol	FTP, username and password authentication ("anonymous" is the default user)

Distribution Transfer NTLM (DS)

Performs the transfer of data (files, directory listings) over the NTLM protocol.

Aspect	Description
Initiator	Distribution agent
Acceptor	• Package stage file server • Distribution location file server
Protocol	NTLM, Windows integrated authentication

Upload Transfer FTP (DS)

Performs the transfer of logs over the FTP protocol.

Aspect	Description
Initiator	Distribution agent
Acceptor	Incoming FTP server on the administration server or parent distribution servers
Protocol	FTP, username and password authentication (“anonymous” is the default user)

Upload Transfer HTTP / HTTPS (DS)

Performs the transfer of logs over the HTTP or HTTPS protocol. See *HTTPS security configuration* for further information about HTTPS security configuration

Aspect	Description
Initiator	Distribution agent
Acceptor	Incoming Web server on the administration server or parent distribution servers
Protocol	HTTP/HTTPS, username/password authentication or Windows integrated authentication

Upload Transfer NTLM (DS)

Performs the transfer of logs over the NTLM protocol.

Aspect	Description
Initiator	Distribution agent
Acceptor	Incoming File Server on the administration server or parent distribution servers
Protocol	NTLM, Windows integrated authentication

Distribution Server Databases

The distribution server data stores include the distributed objects database and smart distribution database.

Distributed Objects Database

This database is used for Windows distribution servers. It contains the details of distributed packages names and versions for use with byte-level differencing. The distribution agent always checks this database before downloading any packages. If the required package already exists on this server, only the difference between the version of the package on the parent server and version of the package on this server is downloaded. The database is updated each time a package is downloaded.

Aspect	Description
Default location	`C:\ManageSoft\Repository\Database\DistributedObjectsDB`
Users	Deployment Manager Distribution Scheduled Task Users

Smart Distribution Database (DS)

The smart distribution database is a caching mechanism used by the system to avoid unnecessary distribution of packages (this includes schedules, policies, client settings data). It is accessed by the following components:

- **Distribution agent**
  The agent always checks the smart distribution database to determine if a package needs to be distributed/ redistributed. When a package is distributed to a server or location for the first time, it is recorded in the database. If a package was redistributed to a server or location, the database is updated with this information.

Aspect	Description
Default location	`C:\ManageSoft\Repository\DeploymentLocations\distributeLog.smd`
Users	Deployment Manager Distribution Scheduled Task Users

Server Types for Distribution Servers

Servers are any data source/repository that RayManageSofti uses. On distribution servers, these include:

- *Distribution location file server (DS)*
- *Distribution location FTP server (DS)*
- *Distribution location web server (DS)*
- *Incoming file server (DS)*
- *Incoming FTP server (DS)*
- *Incoming web server (DS)*
- *Package stage file server (DS)*
- *Package stage FTP server (DS)*
- *Package stage web server (DS)*
- *Remote execution tasks server (DS)*
- *Remote execution tools server (DS)*

Distribution Location File Server (DS)

File sharing service hosting the ManageSoftDL$ file share for package distribution.

Be aware:
File permissions (NTFS) on the underlying directories for the shares also apply through share access. Access checks for both the share permissions and the file permissions must succeed before a user is granted access. The settings here should therefore be considered in conjunction with the appropriate file settings (see *Distribution location (DS)*).

Access must be granted to the following types of users:

- **Administration server scheduled task user**
  Read/write access is required, as the Distribution Agent (which writes data to this location) is usually invoked in that context

- **Managed device users**
  This is the location from which managed device users download packages, and read access is required.

Default location	`file:////{hostname}/ManageSoftDL$`	
Settings location	• Directory properties • Sharing • Permissions	
Default permissions	**User**	**Access rights**
	Everyone	Full control
Minimum permissions	**User**	**Access rights**
	Everyone	Full control

Distribution Location FTP Server (DS)

FTP server hosting the ManageSoftDL FTP share for package distribution.

Be aware:
File permissions (NTFS) on the underlying directories also apply through FTP access.

Access checks for both the share permissions and the file permissions must succeed before a user is granted access. The settings here should therefore be considered in conjunction with the appropriate file settings (see *Distribution location (DS)*).

On Windows distribution servers, this also means that for the "anonymous" FTP user to work, the user configured in IIS as the anonymous user must also be granted access through the file permissions.

Access must be granted to the following types of users:

- **Administration server scheduled task user**
  Read/write access is required, as the Distribution Agent (which writes data to this location) is usually invoked in that context

- **Managed device users**
  This is the location from which managed device users download packages, and read access is required.

Default location	`ftp://{hostname}/ManageSoftDL`
Settings location	Internet Services Manager
Default permissions	Read, write
Minimum permissions	Read, write

Distribution Location Web Server (DS)

Web server hosting the ManageSoftDL web share for package distribution.

Be aware:
File permissions (NTFS) on the underlying directories also apply through HTTP and HTTPS access.

Access checks for both the share permissions and the file permissions must succeed before a user is granted access. The settings here should therefore be considered in conjunction with the appropriate file settings (see *Distribution location (DS)*).

Access must be granted to the following types of users:

- **Administration server scheduled task user**
  Read/write access is required, as the Distribution Agent (which writes data to this location) is usually invoked in that context
- **Managed device users**
  This is the location from which managed device users download packages, and read access is required

Default location	`http://{hostname}/ManageSoftDL`
Settings location	Internet Services Manager
Default permissions	Read, write
Minimum permissions	Read, write

Incoming File Server (DS)

File sharing service hosting the ManageSoftRL$ file share for distribution and managed device logs for uploading.

Be aware:
File permissions (NTFS) on the underlying directories for the shares also apply through share access.

Access checks for both the share permissions and the file permissions must succeed before a user is granted access. The settings here should therefore be considered in conjunction with the appropriate file settings (see *Incoming (DS)*).

Access must be granted to the following types of users:

- **Managed devices users**
  This is the location to which managed device users upload logs and inventories.
- **Child distribution servers**
  This is the location to which child distribution servers upload logs and inventories.
- **All administration server event processors**
  This is the location from which administration server event processors update the RayManageSofti database with data. When the logs/inventories are successfully processed, the files are removed from this directory.

- **Distribution wizard**
  This is the location to which the distribution wizard writes distribution logs.

- **Distribution agent**
  This is the location to which the distribution agent writes distribution logs.

Default location	`file:////{hostname}/ManageSoftRL$`	
Settings location	• Directory properties • Sharing • Permissions	
Default permissions	**User**	**Access rights**
	Everyone	Full control
Minimum permissions	**User**	**Access rights**
	Everyone	Full control

Incoming FTP Server (DS)

FTP server hosting the ManageSoftRL FTP share for distribution and managed device logs for uploading.

Be aware:
File permissions (NTFS) on the underlying directories also apply through FTP access.

Access checks for both the share permissions and the file permissions must succeed before a user is granted access. The settings here should therefore be considered in conjunction with the appropriate file settings (see *Incoming (DS)*).

Access must be granted to the following types of users:

- **Managed devices users**
  This is the location to which managed device users upload logs and inventories.

- **Child distribution servers**
  This is the location to which child distribution servers upload logs and inventories.

- **All administration server event processors**
  This is the location from which administration server event processors update the RayManageSofti database with data. When the logs/inventories are successfully processed, the files are removed from this directory.

- **Distribution wizard**
  This is the location to which the distribution wizard writes distribution logs.

- **Distribution agent**
  This is the location to which the distribution agent writes distribution logs.

Default location	`ftp://{hostname}/ManageSoftRL`
Settings location	Internet Services Manager

Default permissions	Read, write
Minimum permissions	Read, write

Incoming Web Server (DS)

Web server hosting the ManageSoftRL HTTP share for distribution and managed device logs for uploading.

Be aware:
File permissions (NTFS) on the underlying directories also apply through HTTP and HTTPS access.

Access checks for both the share permissions and the file permissions must succeed before a user is granted access. The settings here should therefore be considered in conjunction with the appropriate file settings (see *Incoming (DS)*).

Access must be granted to the following types of users:

- **Managed devices users**
  This is the location to which managed device users upload logs and inventories.
- **Child distribution servers**
  This is the location to which child distribution servers upload logs and inventories.
- **All administration server event processors**
  This is the location from which administration server event processors update the RayManageSofti database with data. When the logs/inventories are successfully processed, the files are removed from this directory.
- **Distribution wizard**
  This is the location to which the distribution wizard writes distribution logs.
- **Distribution agent**
  This is the location to which the distribution agent writes distribution logs.

Default location	`http://{hostname}/ManageSoftRL`
Settings location	Internet Services Manager
Default permissions	Read, write
Minimum permissions	Read, write

Package Stage File Server (DS)

File sharing service hosting the ManageSoftDS$ file share for package distribution.

Be aware:
File permissions (NTFS) on the underlying directories for the shares also apply through share access.

Access checks for both the share permissions and the file permissions must succeed before a user is granted access. The settings here should therefore be considered in conjunction with the appropriate file settings (see *Package stage (DS)*).

This location is where child distribution servers download data from the administration server.

Default location	`file:////{hostname}/ManageSoftDS$`	
Settings location	• Directory properties • Sharing • Permissions	
Default permissions	**User**	**Access rights**
	Everyone	Full control
Minimum permissions	**User**	**Access rights**
	Everyone	Read

Package Stage FTP Server (DS)

FTP server hosting the ManageSoftDS FTP share for package distribution.

Be aware:
File permissions (NTFS) on the underlying directories also apply through FTP access.

Also be aware:
Access checks for both the share permissions and the file permissions must succeed before a user is granted access. The settings here should therefore be considered in conjunction with the appropriate file settings (see *Package stage (DS)*).

For Windows distribution servers, this also means that for the “anonymous” FTP user to work, the user configured in IIS as the anonymous user must also be granted access through the file permissions.

This location is where child distribution servers download data from the administration server.

Default location	`ftp://{hostname}/ManageSoftDS`
Settings location	Internet Services Manager
Default permissions	Read, write
Minimum permissions	Read

Package Stage Web Server (DS)

Web server hosting the ManageSoftDS web share for package distribution.

Be aware:
File permissions (NTFS) on the underlying directories also apply through HTTP and HTTPS access.

Access checks for both the share permissions and the file permissions must succeed before a user is granted access. The settings here should therefore be considered in conjunction with the appropriate

file settings (see *Package stage (DS)*).

This location is where child distribution servers download data from the administration server.

Default location	`http://{hostname}/ManageSoftDS`
Settings location	Internet Services Manager
Default permissions	Read, write
Minimum permissions	Read

Remote Execution Tasks Server (DS)

This is the file sharing service hosting the ManageSoftREA$ file share for storing remote execution actions prior to their retrieval by managed devices.

Be aware:
NTFS file permissions on the underlying directories for the NTLM shares also apply through NTLM share access. Access checks for both the share permissions and the NTFS permissions must succeed before a user is granted access. The settings here should therefore be considered in conjunction with the appropriate NTFS settings (see *Remote execution actions (DS)*).

The remote execution scheduler agent requires access to this location.

Default location	`file:////{hostname}/ManageSoftREA$`	
Settings location	• Directory properties • Sharing • Permissions	
Default permissions	**User**	**Access rights**
	Everyone	Read
Minimum permissions	**User**	**Access rights**
	Everyone	Read

Remote Execution Tools Server (DS)

This is the file sharing service hosting the ManageSoftRET$ file share. This share is used to store the remote execution client agent, and the programs that run managed device adoptions and inventories.

Be aware:
File permissions (NTFS) on the underlying directories for the shares also apply through share access.

Be aware:
Access checks for both the share permissions and the file permissions must succeed before a user is granted access. The settings here should therefore be considered in conjunction with the appropriate file settings (see *Remote execution tools (DS)*).

The remote execution subsystem requires access to this location.

Default location	`file:////{hostname}/ManageSoftRET$`	
Settings location	• Directory properties • Sharing • Permissions	
Default permissions	**User**	**Access rights**
	Everyone	Read
	Anonymous logon	Read
	Network	Read
Minimum permissions	**User**	**Access rights**
	Everyone	Read
	Anonymous logon	Read
	Network	Read

Configuring Managed Device Rights

The Deployment Manager for managed devices software provides the fewest opportunities for tightening up the default security configuration. However, data stored by a managed device is accessed only by peer managed devices (if managed devices are configured for peer-to-peer file sharing), and not by any other part of the RayManageSofti infrastructure, and therefore the need for tight security is considerably lessened.

This chapter describes the default rights and permissions configuration for managed devices, as well as the minimum permissions settings that you can implement for RayManageSofti to function as a whole.

It also describes interactions between various RayManageSofti features and agents, and the security permissions required to enable these interactions.

Rights and permissions settings include username/password logon authentication (for example, for database access or access to a FTP server) and file system permissions controlling access to files, directories, and other file system elements such as shared drives.

The layout of this information is described in *Configuring RayManageSofti rights*.

Be aware:
The same name may be used for security settings on the administration server, distribution servers, and managed devices. Where this occurs, the abbreviation AS, DS, or MD appears after the name of the setting.

Managed Device Overview

In the context of the RayManageSofti system, the managed device need only deal with **data linkage** type security (see *RayManageSofti security model*).

Typical data exchanges for managed devices are downloading packages from the distribution server (or administration server) to the managed device, and uploading logs back to the servers. Permissions on the download and upload locations are set on the servers. The managed device has no control over that aspect of the system.

Managed devices do not host data locations that are accessed by other parts of the RayManageSofti system. The one exception to this is the peer cache, which may be shared with peer managed devices. The peer cache need only offer read access to peer managed devices, since they do not write data to it.

In terms of multi-user security, RayManageSofti relies, to a large extent, on the operating system. In normal operation, RayManageSofti data are generated by a number of scheduled tasks which update policy, upload logs, and so on.

Storage Locations for RayManageSofti Data

Where the data is stored depends on which user runs the scheduled task:

- Data generated by scheduled tasks running as a particular user are placed in:

- On Vista and later, under `C:\Users\{User Name}`
- On earlier Windows platforms, the current user's folder under `Documents and Settings`

- Data generated by scheduled tasks that run in system context are saved to the **All Users** folder. Protection of that data is based on the mechanisms provided by the underlying operating system to prevent users accessing other users' data without proper authorization.

Managed Device Installation

The installation wizard for the Deployment Manager for managed devices software does not offer any options for altering the default configuration of security settings.

Managed Device Features

In this section we will examine a number of managed device features in detail. These include:

- *Application usage agent*
- *Inventory agent*
- *Installation agent*
- *Policy client*
- *Schedule agent*
- *Selection agent*
- *Upload agent (MD)*

Application Usage Agent

The application usage agent is the implementation of the **Gather application usage data** process in the managed device data flow. The agent collects summaries of application usage on a managed device.

By default, this agent is installed but disabled when Deployment Manager for managed devices is installed. The *RMS Implementation* discusses how to enable it during installation. The *Adoption* chapter of the *RMS Discovery* contains details about configuring this agent when automatically bringing devices under management.

This agent runs inside the security service on Windows devices, and as a daemon on non-Windows devices.

The resulting logs will be written to the **All Users** folder on Windows devices, and to the **$(CommonAppDataFolder)/log** folder (by default, `/var/opt/managesoft/log`) on non-Windows devices.

Aspect	Description
Program	`mgssecsvc.exe` (Windows devices) `mgsusageag` (non-Windows devices)
Users	Local system (Windows devices) root (non-Windows devices)
Data files/locations	• Software cache • Installation agent package cache (Windows devices)

Aspect	Description
	• Application usage data • Application usage data uploads • Package database (non-Windows devices)
Communications	None
Databases	None
Servers	None

Inventory Agent

The inventory agent is the implementation of the **Gather inventory** process in the managed device data flow. The agent collects an inventory of hardware and software installed on a managed device.

Typically, this agent is invoked by a regular scheduled task. Depending on the nature of the task (system inventory or user inventory), the resulting logs are written to either the `All Users` folder or the current user's folder. The upload agent requires **Full** access to these directories, as it must delete logs after a successful upload. Synchronizing permissions for the upload agent and inventory agent on the inventory folders is trivial, as long as schedule tasks are created for both processes in the same user context.

Aspect	Description
Program	`ndtrack.exe` (Windows devices) `ndtrack` (non-Windows devices)
Users	Deployment Manager Managed Device Scheduled Task Users
Data files/locations	• Software cache • Installation agent package cache • Inventories • Inventory uploads • Package database (non-Windows devices only)
Communications	None
Databases	None
Servers	None

Installation Agent

The installation agent is the implementation of the **Install on local computer** process in the managed device data flow.

On Windows devices, it consists of `ndlaunch.exe` (the interface) and `ndserv.exe` (handles the actual installation request). `ndserv.exe` runs under the Local system account on the managed device for XP.

On non-Windows devices, it consists of `ndlaunch`. The installation process is as follows:

1. Check the cache for evidence of the requested packages. If the package exists, the agent may only need to download updates.

2. On Windows managed devices, the location of the cache depends on whether or not the managed device is configured for peer-to-peer file sharing.

3. Download the packages (policies, schedules, managed device settings) to a local directory. Packages are downloaded from the nearest distribution location or from a peer managed device (for Windows managed devices configured for peer-to-peer file sharing).

4. Install the package (and policies...) onto the system. Make a backup of the package contents file in the software cache, and a backup of the `.osd` catalog file in the package cache. The two caches are used to for self-healing and uninstall purposes.

5. Write out a log to the installation agent logs.

The installation agent is typically invoked by the policy client or schedule agent, running on a scheduled task. Either the local machine's computer account, or the user context in which the scheduled tasks are run must have access to the distribution location on the distribution server.

Aspect	Description
Program	`ndlaunch.exe`, `ndserv.exe` (Windows devices) `ndlaunch` (non-Windows devices)
Users	• Deployment Manager managed device user • Deployment Manager managed device scheduled task users
Data files/locations	• Software Cache • Installation Agent Package Cache • Managed Device Schedules • Installation Agent Logs • Policies
Communications	• HTTP/HTTPS Installation Transfer • FTP Installation Transfer • NTLM Installation Transfer • TCP Installation Transfer (by managed devices configured for file sharing, transferring files from peers) • UDP broadcast (by managed devices configured for file sharing, requesting files from peers)
Databases	None
Servers	• Distribution location web server on the administration server and distribution servers • Distribution location FTP server on the administration server and distribution servers • Distribution location file server on the administration server and distribution servers

Policy Client

The policy client is the implementation of the Generate policy file for this user / managed device process in the managed device data flow. When client-side policy merging is enabled, the policy client queries policy data from Active Directory. Policies, throughout the Organizational Units, which apply to this user (or computer, depending on whether you are dealing with machine/user policy) are merged and written out into the policy file, in the Policies directory (see *Policies (MD)*). The program is typically executed as a scheduled task.

On Windows devices, the user context in which the task is run must have access to Active Directory and the local Policies folder.

Aspect	Description
Program	`mgspolicy.exe` (Windows devices) `mgspolicy` (non-Windows devices)
Users	Deployment Manager Managed Device Scheduled Task Users
Data files/locations	Policies
Communications	None
Databases	Active Directory
Servers	None

Schedule Agent

The schedule agent handles all scheduled tasks on the managed device. It is the user interface (through scheduled tasks) between the major processes and the user. When this agent is run, it must have access to the managed device schedules location (see *Managed device schedules (MD)*), as well as execute permission on the following executables:

- `ndlaunch.exe` (Windows) or `ndlaunch` (non-Windows)
- `ndtrack.exe` (Windows) or `ndtrack` (non-Windows)
- `mgspolicy.exe` (Windows) or `mgspolicy` (non-Windows)
- `ndupload.exe` (Windows) or `ndupload` (non-Windows)

Aspect	Description
Program	`ndschedag.exe` (Windows devices) `ndschedag` (non-Windows devices)
Users	Deployment Manager Managed Device Scheduled Task Users
Data files/locations	Managed Device Schedules
Communications	None
Databases	None
Servers	None

Selection Agent

The selection agent is the implementation of the **Retrieve installed software details for this user** process in the managed device data flow. It is a user interface on the managed device to manage packages currently installed and those which are available to be installed. It retrieves information from the software cache for the current user which is stored in a separate directory for each user.

Aspect	Description
Program	`ndselect.exe` (Windows devices) `selector` (non-Windows devices)
Users	Deployment Manager Managed Device User
Data files/locations	• SoftwareCache • Installation Agent Package Cache
Communications	None
Databases	None
Servers	None

Upload Agent (MD)

The upload agent is the implementation of the **Upload files** process in managed device data flow. The upload agent is typically invoked by a scheduled task, running in local system context for all users on Windows devices, and running as root on non-Windows devices.

The **Incoming** directory on the distribution server or administration server must allow write access to all managed devices. If user scheduled events are enabled, the directory must also allow write access for all managed device users.

Aspect	Description
Program	`ndupload.exe` (Windows devices) `upload` (non-Windows devices)
Users	Deployment Manager Uploader Scheduled Task Users
Data files/locations	• Installation agent logs • Inventory uploads
Communications	• FTP upload transfer • HTTP/HTTPS upload transfer • NTLM upload transfer
Databases	None
Servers	• Incoming FTP server on the administration server and distribution servers • Incoming Web Server on the administration server and distribution servers (using the HTTP/HTTPS protocol's efficient PUT method to receive files instead of the more common POST method) • Incoming file server on the administration server and distribution servers

Managed Device Data Files / Locations

The key data store locations on managed devices are described in this section. They include:

- *Application usage data*

- *Application usage data uploads*
- *Installation agent logs*
- *Installation agent package cache*
- *Inventories*
- *Inventory uploads*
- *Managed device schedules (MD)*
- *Peer cache*
- *Policies (MD)*
- *Software cache*

Application Usage Data

This is the storage location for application usage data logs that have been generated prior to upload. It also stores important state information for the application usage agent. The following components require access to this location:

- **Application usage agent**
  Generates application usage data logs to this location.

Default location	For Vista and later:
	`C:\Users\{User Name}\Application Data\ManageSoft Corp\ManageSoft\Common\Uploads\UsageData`
	
	For earlier Windows platforms:
	`C:\Documents and Setting\{User Name`
	
	Non-Windows devices:
	`$(CommonAppDataFolder)/ManageSoft Corp/ManageSoft/Usage Agent/UsageData`
	
	(by default `/var/opt/managesoft/ManageSoft Corp/ManageSoft/Usage Agent/UsageData`)	

Settings location	• Directory properties • Security	
Default permissions	Default permissions are inherited from the parent folder: • `C:\Users\{User Name}` (for Vista and later) • `C:\Documents and Settings\{User Name}` (for XP) Permissions vary for different operating systems and are not set by RayManageSofti.	
Minimum permissions	**Individual users**	**Access rights**
	User	Full control
	Local System	Full control
	All users	**Access rights**
	Administrators	Full control

	Local System	Full control

Application Usage Data Uploads

This is the storage location for application usage data logs that have been generated prior to upload. The following components require access to this location:

- **Application usage agent**
  Generates application usage data logs to this location.

- **Upload agent**
  Uploads logs to the distribution server, removing the files after they are successfully uploaded.

Default location	For Vista and later:
	`C:\Users\{User Name}\Application Data\ManageSoft Corp\ManageSoft\Common\Uploads\UsageData`
	
	For earlier Windows platforms:
	`C:\Documents and Setting\{User Name}\Application Data\ManageSoft\ManageSoft\Common\Uploads\UsageData`
	
	For non-Windows platforms:
	`$(CommonAppDataFolder)/uploads/UsageData`
	(by default, `/var/opt/managesoft/uploads/UsageData`)	

Settings location	• Directory properties • Security	
Default permissions	Default permissions are inherited from the parent folder: • `C:\Users\{User Name}` (for Vista and later) • `C:\Documents and Settings\{User Name}` (for XP) Permissions vary for different operating systems and are not set by RayManageSofti.	
Minimum permissions	**Individual users**	**Access rights**
	User	Full control
	Local System	Full control
	All users	**Access rights**
	Administrators	Full control
	Local System	Full control

Installation Agent Logs

This is the storage location for all event logs that have been generated on the managed device and are awaiting upload to a reporting location. In general, there are two types of installations: For the current user only and for all users.

On installation of packages for all users, installation logs are saved to the `All Users` folder. Components accessing this location must run in system context. On installation of packages for the current user, logs are saved to the current user's folder in `C:\Users` (for Vista and later) or `Documents and Settings` (for earlier Windows platforms). Components accessing this location must run in the current users' context. The following components require access to this location:

- **Installation agent**
  Writes installation logs to this location.

- **Upload agent**
  Uploads installation logs to the distribution server, removing the files after it is successfully uploaded.

RayManageSofti logging behavior on Macintosh and UNIX managed devices is configured in `managesoft.xlog`. For further information, see the *Tracing and logging* chapter of *RMS System Reference*.

<table>
<tr><td>Default location</td><td colspan="2">For Vista and later:
• <code>C:\Users\{User Name}\Application Data\ManageSoft Corp\ManageSoft\Common\Uploads\Logs</code> (for user inventories)
• <code>C:\Users\All Users\Application Data\ManageSoft Corp\ManageSoft\Common\Uploads\Logs</code> (for computer inventories)

For earlier Windows platforms:
• <code>C:\Documents and Setting\{User Name}\Application Data\ManageSoft Corp\ManageSoft\Common\Uploads\Logs</code> (for user inventories)
• <code>C:\Documents and Setting\All Users\Application Data\ManageSoft Corp\ManageSoft\Common\Uploads\Logs</code> (for computer inventories)

For non-Windows platforms:
<code>$(CommonAppDataFolder)/uploads/logs</code>
(by default, <code>/var/opt/managesoft/uploads/logs</code>)</td></tr>
<tr><td>Settings location</td><td colspan="2">• Directory properties
• Security</td></tr>
<tr><td>Default permissions</td><td colspan="2">Default permissions are inherited from the parent folder:
• <code>C:\Users\{User Name} (for Vista and later)</code>
• <code>C:\Documents and Settings\{User Name} (for XP)</code>

Permissions vary for different operating systems and are not set by RayManageSofti.</td></tr>
<tr><td rowspan="4">Minimum permissions</td><td>Individual users</td><td>Access rights</td></tr>
<tr><td>User</td><td>Full control</td></tr>
<tr><td>Local System</td><td>Full control</td></tr>
<tr><td>All users</td><td>Access rights</td></tr>
<tr><td></td><td>Administrators</td><td>Full control</td></tr>
</table>

	Local System	Full control

Installation Agent Package Cache

The installation agent package cache is the storage location for all application packages that have been downloaded to the managed device. The cache is used by the following components:

- **Installation agent**
  The cache is accessed when the agent installs/uninstalls a software package.
- **Selection agent**
  Lists packages in the cache.
- **Application usage agent**
- **Inventory agent**

If peer-to-peer file sharing is not enabled, files are downloaded from the nearest distribution location. If peer-to-peer file sharing is enabled, files are downloaded from either the local peer cache or a peer cache on another managed device.

Description	The storage location for all application packages that have been downloaded to the managed device.	

Default location	Windows devices:
	`C:\Program Files\ManageSoft\Launcher\PkgCache`
	
	Non-Windows devices:
	`$(CommonAppDataFolder)/cache`
	(by default, `/var/opt/managesoft/cache`)	

Settings location	• Directory properties • Security	
Default permissions	Default permissions are inherited from the parent folder `C:\Program Files`. Permissions on this folder vary for different operating systems and are not set by RayManageSofti.	
Minimum permissions	**User**	**Access rights**
	Deployment Manager Managed Device User	Read, Execute
	Deployment Manager Managed Device Scheduled Task Users	Read, Execute
	Local System	Full control

Inventories

This is the storage location for all inventories that have been generated on the managed device. In general, there are two types of inventories: User and computer.

For Windows devices, computer inventories are saved to the **All Users** folder. Components accessing this location must run in system context. User inventories are saved to the current user's folder in `Users` (for Vista and later) or `Documents and Settings` (for earlier Windows platforms). Components accessing this location must run in the current user's context.

For non-Windows devices, only computer inventories are generated. The following components require access to this location:

- **Inventory agent**
  Generates inventories to this location.

Default location	For Vista and later: • `C:\Users\{User Name}\Application Data\ManageSoft Corp\ManageSoft\Tracker\Inventories` (for user inventories) • `C:\Users\All Users\Application Data\ManageSoft Corp\ManageSoft\Tracker\Inventories` (for computer inventories) For earlier Windows platforms: • `C:\Documents and Setting\{User Name}\Application Data\ManageSoft Corp\ManageSoft\Tracker\Inventories` (for user inventories) • `C:\Documents and Setting\All Users\Application Data\ManageSoft Corp\ManageSoft\Tracker\Inventories` (for computer inventories) For non-Windows devices: `$(CommonAppDataFolder)/uploads/Inventories` (by default, `/var/opt/managesoft/uploads/Inventories`)	
Settings location	• Directory properties • Security	
Default permissions	Default permissions are inherited from the parent folder: • `C:\Users\{User Name}` (for Vista and later) • `C:\Documents and Settings\{User Name}` (for XP) Permissions vary for different operating systems and are not set by RayManageSofti.	
Minimum permissions	**Individual users**	**Access rights**
	User	Full control
	Local System	Full control
	All users	**Access rights**
	Administrators	Full control
	Local System	Full control

Inventory Uploads

This is the storage location for all inventories that have been generated on the managed device and are awaiting upload to a reporting location. In general, there are two types of inventories: User and computer.

Computer inventories are saved to the `All Users` folder. Components accessing this location must run in system context. User inventories are saved to the current user's folder in `Users` (for Vista and later) or `Documents and Settings` (for earlier Windows platforms). Components accessing this location must run in the current users' context.

The following components require access to this location:

- **Inventory agent**
  Copies inventories to this location when they are ready to be uploaded to the distribution server.

- **Upload agent**
  Uploads inventories to the distribution server, removing the files after it is successfully uploaded.

Default location	For Vista and later: • `C:\Users\{User Name}\Application Data\ManageSoft Corp\ManageSoft\ Common\Uploads\Inventories` (for user inventories) • `C:\Users\All Users\Application Data\ManageSoft Corp\ManageSoft\ Common\Uploads\Inventories` (for computer inventories) For earlier Windows platforms: • `C:\Documents and Setting\{User Name}\Application Data\ManageSoft Corp\ManageSoft\ Common\Uploads\Inventories` (for user inventories) • `C:\Documents and Setting\All Users\Application Data\ManageSoft Corp\ManageSoft\ Common\Uploads\Inventories` (for computer inventories) For non-Windows devices: `$(CommonAppDataFolder)/uploads/Inventories` (by default, `/var/opt/managesoft/uploads/Inventories`)	
Settings location	• Directory properties • Security	
Default permissions	Default permissions are inherited from the parent folder: • `C:\Users\{User Name} (for Vista and later)` • `C:\Documents and Settings\{User Name} (for XP)` Permissions vary for different operating systems and are not set by RayManageSofti.	
Minimum	**Individual users**	**Access rights**

	User	Full control
	Local System	Full control
	All users	**Access rights**
	Administrators	Full control
	Local System	Full control

Managed Device Schedules (MD)

This is the storage location for all schedules that have been installed on the managed device. In general, there are two types of schedules: User and computer.

On Windows devices, computer schedules are saved to the **All Users** folder. Components accessing this location must run in system context. User schedules are saved to the current user's folder in `Users` (for Vista and later) or `Documents and Settings` (for earlier Windows platforms). Components accessing this location must run in the current users' context.

On non-Windows devices, only computer schedules are available. The default location is described in the table below.

The following components require access to this location:

- **Schedule agent**
  Runs schedules stored in this location. For non-Windows devices, the schedule agent also installs the scheduled tasks to this location.

- **Installation agent**
  Installs the scheduled tasks to this location (Windows devices only).

Default location	For Vista and later: • `C:\Users\{User Name}\Application Data\ManageSoft Corp\ManageSoft\Schedule Agent\Schedules` (for user schedules) • `C:\Users\All Users\Application Data\ManageSoft Corp\ManageSoft\Schedule Agent\Schedules` (for computer schedules) For earlier Windows platforms: • `C:\Documents and Setting\{User Name}\Application Data\ManageSoft Corp\ManageSoft\Schedule Agent\Schedules` (for user schedules) • `C:\Documents and Setting\All Users\Application Data\ManageSoft Corp\ManageSoft\Schedule Agent\Schedules` (for computer schedules) For non-Windows devices: `$(CommonAppDataFolder)/scheduler/schedules` (by default, `/var/opt/managesoft/scheduler/schedules`)

Settings location	• Directory properties • Security	
Default permissions	Default permissions are inherited from the parent folder: • `C:\Users\{User Name} (for Vista and later)` • `C:\Documents and Settings\{User Name} (for XP)` Permissions vary for different operating systems and are not set by RayManageSofti.	
Minimum permissions	**Individual users**	**Access rights**
	User	Full control
	Local System	Full control
	All users	**Access rights**
	Administrators	Full control
	Local System	Full control

Peer Cache

This section is only applicable to Windows managed devices.

When peer-to-peer file-sharing is enabled, a storage location on managed devices operates a local distribution location. When the installation agent downloads packages for installation, it downloads them from the peer cache on the same managed device or on a peer managed device.

The installation agent downloads files from peer caches to the local installation agent package cache.

The cache is used by the installation agent. The cache is accessed as a network download location when the installation agent attempts to download a software package.

Description	The storage location for all application packages that have been downloaded to the managed device when peer-to-peer file sharing is enabled.	
Default location	`$(CommonAppDataFolder)\ManageSoft Corp\ManageSoft\PeerCache`	
Settings location	• Directory properties • Security	
Default permissions	Default permissions are inherited from the parent folder `C:\Program Files`. Permissions on this folder vary for different operating systems and are not set by RayManageSofti.	
Minimum permissions	**User**	**Access rights**
	Deployment Manager Managed Device User	Read, Execute
	Deployment Manager Managed Device Scheduled Task Users	Read, Execute

	Local System	Full control

Policies (MD)

This is the storage location for all policies that have been installed on the managed device. In general, there are two types of policies: User and machine.

On Windows devices, machine policies are saved to the `All Users` folder. Components accessing this location must run in system context. User policies are saved to the current user's folder in `Users` (for Vista and later) or `Documents and Settings` (for earlier Windows platforms). Components accessing this location must run in the current users' context.

On non-Windows devices, only machine policies are available. The following components require access to this location:

- **Policy client**
  Policy files (`.npl`), generated from Active Directory data, are saved to this location.
- **Installation agent**
  Reads data from this location when the policy files are being installed.
- **Selection agent**
  Displays the contents of this location.

Default location	For Vista and later: • `C:\Users\{User Name}\Application Data\ManageSoft Corp\ManageSoft\Policy Client\Policies` (for user policies) • `C:\Users\All Users\Application Data\ManageSoft Corp\ManageSoft\Policy Client\Policies` (for machine policies) For earlier Windows platforms: • `C:\Documents and Setting\{User Name}\Application Data\ManageSoft Corp\ManageSoft\Policy Client\Policies` (for user policies) • `C:\Documents and Setting\All Users\Application Data\ManageSoft Corp\ManageSoft\Policy Client\Policies` (for machine policies) For non-Windows devices: `$(CommonAppDataFolder)/cache/package/Policies/Merged` (by default, `/var/opt/managesoft/cache/package/Policies/Merged`)
Settings location	Directory properties, Security

Default permissions	Default permissions are inherited from the parent folder: • `C:\Users\{User Name} (for Vista and later)` • `C:\Documents and Settings\{User Name} (for XP)` Permissions vary for different operating systems and are not set by RayManageSofti.	
Minimum permissions	**Individual users**	**Access rights**
	User	Full control
	Local System	Full control
	All users	**Access rights**
	Administrators	Full control
	Local System	Full control

Software Cache

The software cache is the storage location for all package content files of applications that have been installed on the managed device. This location is accessed by the installation agent. On non-Windows devices, this location is also accessed by the inventory agent.

Default location	On Windows devices:
	`C:\Program Files\ManageSoft\Launcher\Cache`
	
	On non-Windows devices:
	`$(CommonAppDataFolder)/live`
	(by default, `/var/opt/managesoft/live`)	

Settings location	• Directory properties • Security	
Default permissions	Default permissions are inherited from the parent folder `C:\Program Files`. Permissions on this folder vary for different operating systems and are not set by RayManageSofti.	
Minimum permissions	**User**	**Access rights**
	Deployment Manager Managed Device User	Read, Execute
	Deployment Manager Managed Device Scheduled Task Users	Read, Execute
	Local System	Full control

Managed Device Communications

This section describes network connections that may be established between two managed device features or a managed device feature and a server:

- *Installation transfer FTP*
- *Installation transfer HTTP/HTTPS*
- *Installation transfer NTLM*
- *Installation transfer TCP*
- *Upload transfer FTP (MD)*
- *Upload transfer HTTP/HTTPS (MD)*
- *Upload transfer NTLM (MD)*

Installation Transfer FTP

Performs the transfer of data (files, directory listings) over the File Transfer Protocol.

Aspect	Description
Initiator	Installation agent
Acceptor	Distribution location FTP server
Protocol	FTP, username and password authentication (“anonymous” is the default user)

Installation Transfer HTTP / HTTPS

Performs the transfer of data (files, directory listings) over the HTTP protocol. See *HTTPS security configuration* for further information about HTTPS security configuration

Aspect	Description
Initiator	Installation agent
Acceptor	Distribution location web server
Protocol	HTTP/HTTPS, username/password and Windows integrated authentication

Installation Transfer NTLM

Performs the transfer of data (files, directory listings) over the NTLM protocol.

Aspect	Description
Initiator	Installation agent
Acceptor	Distribution location file server
Protocol	NTLM, Windows integrated authentication

Installation Transfer TCP

Performs the transfer of data (files, directory listings) using TCP.

Aspect	Description
Initiator	Installation agent
Acceptor	Peer managed device's peer download agent
Protocol	TCP

Upload Transfer FTP (MD)

Performs the transfer of logs over the FTP protocol.

Aspect	Description
Initiator	Upload agent
Acceptor	Incoming FTP server on administration server or parent distribution servers
Protocol	FTP, username and password authentication ("anonymous" is the default user)

Upload Transfer HTTP / HTTPS (MD)

Performs the transfer of logs over the HTTP protocol. See *HTTPS security configuration* for further information about HTTPS security configuration

Aspect	Description
Initiator	Distribution agent
Acceptor	Incoming Web server on the administration server or parent distribution servers
Protocol	HTTP/HTTPS, username/password and Windows integrated authentication

Upload Transfer NTLM (MD)

Performs the transfer of logs over the NTLM protocol.

Aspect	Description
Initiator	Upload agent
Acceptor	Incoming file server on administration server or parent distribution servers
Protocol	NTLM, Windows integrated authentication

Managed Device Databases

Managed devices access the Active Directory store of data.

Managed devices running operating systems other than Windows access a package database.

Active Directory (MD)

The Active Directory contains, among other things, software management information for Windows managed devices and users in the organization. If the managed device is configured for client-side policy merging, the policy client will directly query Active Directory for software assignment/entitlement data. Data retrieved from Active Directory is merged into a RayManageSofti policy file (`.npl`), which gets installed on the managed device.

Aspect	Description
Default location	Active Directory server
Users	Everyone

Package Database (Non-Windows Devices)

The package database is an index file allowing the inventory agent and application usage agent on non-Windows devices fast access to details of the native packages installed on managed devices.

Aspect	Description
Default location	`$(CommonAppDataFolder)/cache/packagedb.cdb`
Users	root

HTTPS Security Configuration

The HTTPS protocol provides SSL (Secure Sockets Layer) encryption to secure data communications. To use SSL, you must configure your system to use web server certificates.

A web server certificate must be installed on each IIS server in your distribution hierarchy that will use HTTPS. This web server certificate is sent whenever data is transferred from the IIS server.

You must also configure managed devices to accept the web server certificates sent from web servers.

This chapter describes how to:

- *Create the web server certificate*
- *Configure IIS servers in the distribution hierarchy*
- *Configure managed devices*

Creating a Web Server Certificate

You can choose to use a certificate from a certification authority (discussed in some detail in *The basics of digital signing*) or you can use a local Microsoft Certificate Server. The instructions in this chapter describe how to create and use a local Microsoft Certificate Server.

If you choose to use a local Microsoft Certificate Server, you will need to install one of the following:

- **Microsoft Certificate Services**

- **Certificate Server**
  (available on Windows Server 2003 in **Add/Remove Windows Components**; and in Windows Server 2008 from **Start** > **Administrative Tools** > **Server Manager**, select the **Roles** node and select **Add Roles**)

To prepare for and create a web server certificate:

1. Install one of these products on a server that will become your certificate server. You may find it most convenient to install it on your administration server.

   When installing this software, the simplest option is to select the installation option for a stand-alone root CA, although you can choose alternative options.

2. If the server is running Windows 2000 service pack 3 or earlier, you must also either install Microsoft hotfix Q323172, or upgrade to service pack 4. If hotfix Q323172 has been installed prior to the certificate software, refer to Hotfix Q328595 for additional information.

3. Use the certificate software to create a web server certificate. Refer to the relevant product documentation for details.

   The certificate you create will become the key that allows downloads and uploads to occur on IIS servers using the HTTPS protocol.

Configuring IIS Servers

Use these procedures for any IIS distribution servers or distribution locations that will use the HTTPS protocol. (This protocol is set in the properties of the distribution server or distribution location on the administration console.)

To install a certificate on an IIS server, you must:

- Create a certificate file on the IIS server
- Request the issue of certificate details from the certificate server
- Issue the certificate from the certificate server
- Install the certificate details into the certificate file when they are issued from the certificate server
- Complete the installation in Internet Services Manager

To Create a Certificate File

On the IIS server:

1. Start **Internet Services Manager**.

2. Expand the node for the server's computer name. On Windows 2003, also expand the **Web sites** node.

3. Right-click **Default Web site** and select **Properties** from the menu.
   The **Default Web Site Properties** dialog is displayed.

4. Select the **Directory Security** tab.

5. Click **Server Certificate**.
   The **Web Server Certificate Wizard** is displayed.

6. Use the wizard to create a new certificate, taking note of the following settings:

   - On the **Delayed or Immediate Request** page, select the **Prepare the request now, but send it later** radio button.
   - On the Y**our Site's Common Name** page, type the fully qualified domain name of this server as it is required to appear in any https: URLs (for example, `www.raynet.de`).

The certificate file that you create is a place-holder for you to install the certificate information from the certificate created on the certificate server.

To Request the Issue of a Certificate

On the IIS server:

1. Use your internet browser to browse to `http://<certificate server>/certsrv`.
   where certificate server is the name of the server on which you created the certificate.

2. Use the web page options to request a certificate.

3. Select the option to submit a web server certificate request using a bas64 encoded `PKCS #10` file.

4. Browse to the file you created in the previous section.

   In some cases you may not have permission to upload the file. In this scenario, copy the text from the file into the **Saved Request** text box.

5. Click **Submit** to submit the request.

The certificate request is submitted to the certificate server (where Microsoft Certificate Services or Certificate Server is installed).

To Issue the Certificate from the Certificate Server

On the certificate server:

1. From the **Start** menu, select **All Programs** > **Administrative Tools** > **Certification Authority**.

2. Browse to find the certificate request you submitted for the IIS server in the list of pending certificate requests.

3. Right-click on the certificate request and select **All Tasks** > **Issue** from the menu.

This issues the certificate to the IIS server. The next step is to install the issued certificate on the IIS server.

To Install the Certificate on the IIS Server

On the IIS server:

1. Use your internet browser to browse to
   `http://<certificate server>/certsrv`
   where certificate server is the name of the certificate server on which you issued the certificate.

2. Select the option to view the status of outstanding requests (Windows 2003).

3. Select the certificate.

4. Select the **DER encoded** option.

5. Select the option to download the certificate.

6. A dialog is displayed for you to specify the download location. Save the certificate to a file on the IIS server.

The web certificate is now installed on the IIS server.

To Complete the Certificate Installation in Internet Services Manager

On the IIS server:

1. Start **Internet Services Manager**.

2. Expand the console list for the server's computer name.

Right-click **Default Web site** and select **Properties** from the menu.

The **Default Web Site Properties** dialog is displayed.

3. Select the **Directory Security** tab.
4. Click **Server Certificate**.

   The **Web Server Certificate Wizard** is displayed.

5. Complete the wizard, taking note of the following settings:

   - Select **Enter response from CA request**.
   - Browse to the certificate file you copied from the certificate server.

6. Use your internet browser to browse to
   `http://<web server>/certsrv`
   where web server is the name of your IIS server on which you are installing the certificate.
7. Double-click the padlock at the bottom of the window to validate the server certificate.

If You Need to Copy a Certificate to a Different File Type

If the pending certificate is not listed when you attempt to install it on the IIS server, the certificate request you processed must be copied to a different file type.

On the certificate server:

1. From the **Start** menu, select **All Programs** > **Administrative Tools** > **Certification Authority**.
2. Select **Issued Certificates** and find the certificate.
3. Right-click the certificate and select **Open** from the menu.
4. Select the **Details** tab.
5. Select **Copy to file**.
   The **Certificate Export Wizard** starts.
6. Click **Next**.
   The **Export File Format** page is displayed.
7. Select **DER encoded binary X.509**.
8. Click **Next**.
   The **File to Export** page is displayed.
9. Type the name of a file to save the web server certificate.
10. Click **Next**.
    A **summary** page is displayed.

11.Click **Finish**.

12.Copy the file to the IIS server.

Configuring Managed Devices

On managed devices that will connect to an IIS server using HTTPS protocol, you need to configure the managed device to accept the web server certificates sent from web servers.

Configuring Managed Devices to Accept Web Server Certificates

When an IIS server connects to a managed device, it issues a web server certificate with data that is transferred. The managed device will check whether the web server certificate will be accepted by users on the managed device.

This step identifies the certification authority of the certificate server in the Local Computer Trusted Root CA store. This then allows any web server certificate issued by that CA to be accepted on the managed device.

To identify the certification authority on the managed device:

1. Use your internet browser to browse to
   `http://<certificate server>/certsrv`
   where certificate server is the server on which Microsoft Certificate Services or Certificate Server is installed.

2. Select the option to retrieve or download a CA certificate.

3. Select the CA certificate.

4. Click **Download CA certificate**.

5. Save the CA certificate to a file.

6. Select the option to install this CA certificate path/chain.

   This adds the certificate to the trusted root certification authorities store.
   A number of warnings may display. Select the appropriate options to continue processing.

The certification authority is now identified on this managed device.

Managed Devices and Certificate Revocation Lists

A web server certificate is applied to the data being downloaded or uploaded from an HTTPS webserver. When receiving web server certificates from servers, RayManageSofti checks the CA (certification authority) server to ensure that the certificates are not on the CRL (certificate revocation list).

You can use the `CheckCertificateRevocation` preference to prevent RayManageSofti from performing this check. See the *RMS Preferences for Managed Devices* for further details about this preference.

Role-based Security

The RayManageSofti console provides users with powerful editing and reporting facilities that can be used to perform software management across an enterprise and view detailed information about the enterprise.

Some enterprises may wish to restrict access to these facilities to particular users, or to limit the organizational units upon which users may perform actions.

System administrators can use the RayManageSofti role-based security system to implement role-based security measures relating to:

- Setting security rights
- Report access
- Adding packages to policy for Deployment Manager and Security Manager
- Allocating users and computers to deployment policies for Deployment Manager and Security Manager

Access to the RayManageSofti Console

Access to the RayManageSofti console is, by default, denied to all users except those in the **MGS Administrators** and **MGS Reporting Users** security groups. (For more information about RayManageSofti security groups, see *Default user groups*.)

You can modify this access by assigning (or denying) rights to other security groups.

What Rights Can You Assign?

Role-based security defines a number of resources, which represent an area of data, such as software distribution or hardware assets.

For each resource, there are two actions that may be permitted:

- **Read** access allows the user to view the data
- **Modify** access allows the user to view the data and add, change or delete the data

If a user attempts to perform an action for which they do not have access, a message similar to the following is displayed:

Access Denied. The user `<name>` is not permitted to view report `<ReportName>`

Rights can be granted to Active Directory groups to perform a specified action for a specified resource.

These rights are represented by rules that are created, modified, and deleted using the **Assign Rights wizard**. A rule can apply across all organizational units, or can be filtered to include only specified organizational units.

Changes to Active Directory Groups

When you use the Assign Rights wizard to assign rights to an Active Directory group, the rights apply immediately to all members of the group.

However, over time, group memberships can change. Access rights are automatically updated each time a policy merge is run for users. See the *Deployment Policies* chapter of the *RMS Software Deployment* for details about merging policy.

Be aware:
In a multi-domain environment, domains can be configured so that policy merge is only run for computers. If you choose not to run policy merges for users, changes to group membership are not reflected in the RayManageSofti database.

For domains with users that require rights to the RayManageSofti console and reports, it is important to ensure that the domain is configured to run user policy merges. See *Configuring environments with multiple domains* for details.

To Assign or Deny Rights to a Group

The Assign Rights wizard allows you to:

- Assign new rights to a resource
- Deny rights to a resource
- Edit existing rights
- Limit rights to specific organizational units (where applicable)

To use the Assign Rights wizard:

1. From the RayManageSofti console, open the **Settings** node, and click on the **Open Assign Rights wizard** button which is available from the **permissions** tab.

Be aware:
Access to this wizard is restricted, by default, to members of the **MGS Administrators** security group. If you do not have access, a message similar to the following is displayed: `The current user does not have the required access to assign rights.`

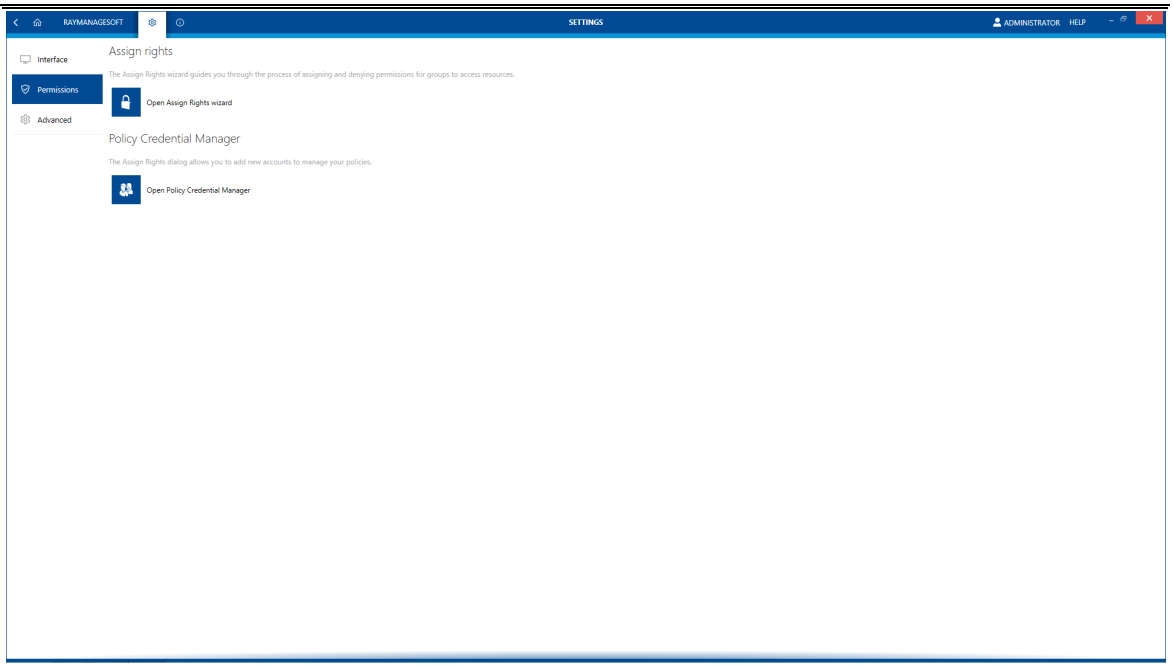

If you have appropriate rights, the **Assign Rights wizard** is displayed.

2. Click **Next** to select the group whose rights you will modify.

   The **Group Selection** page is displayed.

3. Click **Browse**.

   The **Browse for Groups** page is displayed.

4. Navigate to find the group whose rights you want to define.

5. Select the check box next to the group name.

6. Click **OK** to return to the **Group Selection** page.

   Details of the group you selected are displayed.

7. Click **Next** to select the rights you want to assign or deny.

   The **Assign Rights for Selected Group** page is displayed.

   This page lists actions that have already been assigned or denied to the selected group.

   Use this page to create, change or remove rules that assign or deny access to the selected resource.

If you want to...	Take this action...
Create a new right	Click **Add**. The Add a right page is displayed. See *Adding a new right* for further details.
Edit a right	Select the rule you want to change.

If you want to...	Take this action...
	Click **Edit**. The **Edit a right** page is displayed. See *Editing a right* for further details.
Revert to original rights	Click **Revert**. Any new rules or changes to existing rules made in this editing session are reversed. Any rules deleted in this editing session are reinstated.
Remove a right	Select the rule you want to remove. Click **Remove**. The rule is removed.

8. When you have finished working with the rights of the selected group, click **Next** to view a summary of changes to rights that you have made.

   The **Summary** page is displayed.

   This page identifies any rules you have created, modified, or removed.

9. Check that the details you have specified are correct. Click **Back** to return to earlier pages and update details if necessary.

10. Click **Next** to complete the changes.

    The wizard makes the selected changes to group rights, and displays a status bar to illustrate the progress of these changes.

    When this task is completed, the **Rights Assignment Succeeded** page is displayed.

    This page indicates that the rights have been successfully changed.

11. Click **Finish** to complete this wizard.

If Rights Could Not Be Assigned

If your changes cannot be saved, the **Rights Assignment Failed** page is displayed. This page indicates that the rights have not been changed.
Possible reasons for the failure include:

- The Active Directory data in the RayManageSofti database may not be up-to-date.

  Perform an AD reconciliation to update the database. See the *Command line tools* chapter of *RMS System Reference*.

- Problems accessing the RayManageSofti database.

  For example, there may be no connection established between the administration server and the database server. The network connection may be disrupted, or the computer hosting the database server may be down, or switched off.

- Security data files are unavailable or have been removed.

You can enable tracing for this wizard using a `.config` file. For details, see the *Tracing and Logging* chapter of *RMS System Reference*.

Adding a New Right

When you select to add a new right, the **Enter Details of Access Right** page is displayed.

1. Select a resource to allow or deny to the selected group. To do this:

   - Click **Browse**.

     The **Browse for Resources** page is displayed.

     This page displays a hierarchical representation of the resources that are available to RayManageSofti console users.

   - Browse the resources to find the one that you want to assign or deny to the selected group.
   - Select a resource.
   - Click **OK** to return to the **Enter Details of Access Right** page.

     The selected resource name is displayed in the **Resource** field.

2. In the **Type** field, select **Allow** if you are creating a rule to allow access, or select **Deny** to deny access to the selected resource.
3. In the **Right** field, select the action that you want to allow or deny.

   For reporting resources, the only right available for selection is the right to view data.

4. Where applicable, you can choose to limit the rule to a selected organizational unit.

   To do this:

   - Click **Browse...** to display a Browse page
   - Navigate to select the organizational unit to which this rule will apply
   - Click **OK** to return to the **Enter Details of Access Right** page

5. Click **OK** to add the rule and return to the **Assign Rights for Selected Group** page.

Editing a Right

When you select to edit a rule, the **Edit Details of Access Right** page is displayed.

Use this page to modify the details of this rule as required.

1. In the **Type** field, select **Allow** if you are creating a rule to allow access, or select **Deny** to deny access to the selected resource.
2. In the **Right** field, select the action that you want to allow or deny.

   For reporting resources, the only right available for selection is the right to view data.

3. Where applicable, you can choose to limit the rule to a selected organizational unit.

   To do this:

   - Click **Browse...** to display a **Browse** page
   - Navigate to select the organizational unit to which this rule will apply
   - Click **OK** to return to the **Edit Details of Access Right** page

4. Click **OK** to return to the **Assign Rights for Selected Group** page.

Configuring Web Proxy Servers

In this chapter, you will learn how to configure your administration server to download data from the Internet through a web proxy server.

This configuration setting may be required if your administration server connects to the Internet using a web proxy server and you want to download the following types of data from the Internet:

- Deployment Manager support packages
- Security Manager third party prerequisite packages
- Security Manager security patches and errata files

To Configure Web Proxy Settings

To change the web proxy settings, you can use the RayManageSofti Configuration tool:

1. The **Deployment Manager configuration console** is accessible from **Start** > **All programs** > **Deployment Manager Tools** > **Configuration** > **Deployment Manager Configuration**.

Note:
The **Product Activation Wizard** may display once you open the console, if you are using a temporary time out license or if the number of devices you are managing approaches or exceeds the number you are licensed to manage. If it does appear, select **Start** to go ahead..

2. In the console tree, select **Network settings**.

   The **Configure your Network Settings** page is displayed, with the Web proxy server account name selected for editing.

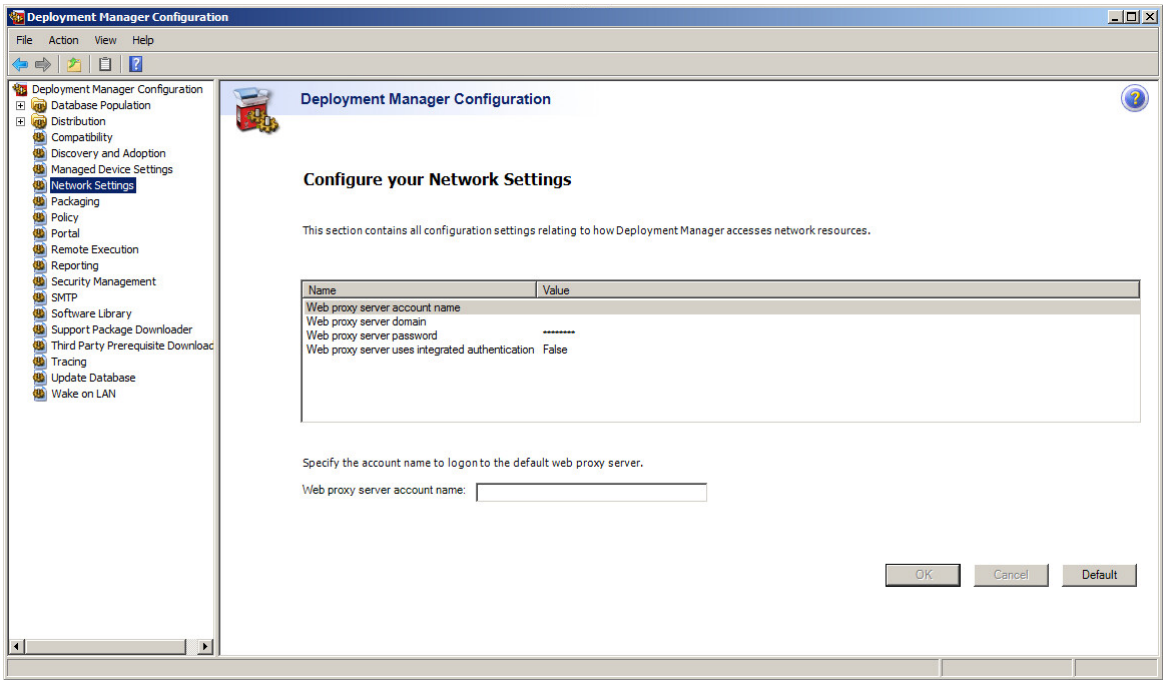

By default, RayManageSofti uses the Internet HTTP proxy server settings used by the operating system.

3. To use integrated authentication credentials when downloading data from the internet:

   - Select the **Web proxy server uses integrated authentication** setting.

     Additional fields display below the list of settings.

   - Select the **Web proxy users integrated authentication** check box.

4. Alternatively, to use credentials of a specific user account when downloading data from the internet, identify the user account in the following settings on this page.

   - **Web proxy server account name**
   - **Web proxy server domain**
   - **Web proxy server password**
     (the password shows an asterisk in this field, even if the password is blank)

5. Click **OK**.

6. Close the **Deployment Manager Configuration** window.

7. Restart Deployment Manager so that your changed settings will take effect.

In the Registry

Web proxy settings are stored on the administration server in the registry under:

```
[Registry]\ManageSoft\SecurityPatch\CurrentVersion\Credentials
```

When you change settings on the **Network Settings** page, the following registry keys are changed:

Setting	Registry key
Web proxy server account name	`username_1`
Web proxy server domain	`domain_1`
Web proxy server password	`password_1`
Web proxy users integrated authentication	`IntegratedAuth`

Moving the Database

This chapter describes the process to follow if you want to move your RayManageSofti database from one server to another.

Throughout this chapter:

- **Data server** refers to the physical server on which the RayManageSofti database and the data server component of Deployment Manager for administration servers is installed
- **Core server** refers to the physical server on which the core server component of Deployment Manager for administration servers is installed

Choose which of the following statements most closely matches your requirements, and follow its instructions below to move your RayManageSofti database:

- **Scenario 1: Combined data and core server: moving to a separate data server**

  I have an existing physical server which is a combined core and data server. I want to move the data server to a separate physical server.
  Follow the instructions in *Scenario 1: Combined data and core server: moving to a separate data server.*

- **Scenario 2: Separate data and core servers: moving to a new data server**

  I have existing separate data and core servers. I want to move the data server to a new physical server.
  Follow the instructions for *Scenario 2: Separate data and core servers: moving to a new data server.*

- **Scenario 3: Combined data and core server: moving to a new physical server**

  I have an existing physical server which is a combined core and data server. I want to move this installation and data to a new physical server.
  Follow the instructions for *Scenario 3: Combined data and core server: moving to a new physical server.*

- **Scenario 4: Separate data and core servers: moving to a new core server**

  I have existing separate core and data servers. I want to move the core server to a new physical server.
  Follow the instructions for *Scenario 4: Separate data and core servers: moving to a new core server.*

Scenario 1: Combined Data and Core Server: Moving to a Separate Data Server

I have an existing physical server which is a combined core and data server. I want to move the data server to a separate physical server.

Complete these steps:

1. On the server that currently hosts the RayManageSofti database ("old server"), check the collation of the RayManageSofti database. See *About database collations* for more details.

2. On the server to which you will move the data server ("new data server"), make sure all required system software is installed, according to the specifications in the *RMS Implementation*.

Be aware:
If you are installing SQL Server, perform a custom install, and set the default database collation order to be the same as the collation order of the current RayManageSofti database (see *About database collations* for more details). Refer to SQL Server installation instructions for details about setting the default database collation order.

If SQL Server is already installed, check the collations of the `tempdb` and `master` databases, and change the collation of the RayManageSofti database on your old server to match.

3. Transfer data from your old server to your new server. See *Transferring data between servers* for details.

4. On your new server, install the data server component of Deployment Manager for administration servers, identifying the new database.

5. On your new server, update the new database with RayManageSofti data copied from your old server. See *Updating RayManageSofti settings* for details.

6. On your old server, update your RayManageSofti settings to point to the new data server. See *Updating RayManageSofti settings* for detailed instructions.

Scenario 2: Separate Data and Core Servers: Moving to a New Data Server

I have existing separate data and core servers. I want to move the data server to a new physical server.

Complete these steps:

1. On the server that currently hosts the RayManageSofti database ("old data server"), check the collation of the RayManageSofti database. Refer to *About database collations* for more details about database collations.

2. On the server to which you will move the RayManageSofti database ("new data server"), install all required system software, according to the specifications in the *RMS Implementation*.

Be aware:
If you are installing SQL Server, perform a custom install, and set the default database collation order to be the same as the collation order of the current RayManageSofti database (see *About database collations* for more details). Refer to SQL Server installation instructions for details about setting the default database collation order.

If SQL Server is already installed, check the collations of the `tempdb` and `master` databases, and change the collation of the RayManageSofti database on your old server to match.

3. Transfer data from your old data server to your new data server. See *Transferring data between servers* for details.

4. On your new data server, install the data server component of Deployment Manager for administration servers, identifying the new database.

5. On your core server, update RayManageSofti settings to point to the new data server. See *Updating RayManageSofti settings* for detailed instructions.

Scenario 3: Combined Data and Core Server: Moving to a New Physical Server

I have an existing physical server which is a combined core and data server. I want to move this installation to a new physical server.

1. On the server that currently hosts RayManageSofti ("old server"), check the collation of the RayManageSofti database. Refer to *About database collations* for more details about database collations.

2. On the new server, install all required system software, according to the specifications in the *RMS Implementation*.

Be aware:
If you are installing SQL Server, perform a custom install, and set the default database collation order to be the same as the collation order of the current RayManageSofti database (see *About database collations* for more details). Refer to SQL Server installation instructions for details about setting the default database collation order.

If SQL Server is already installed, check the collations of the `tempdb` and `master` databases, and change the collation of the RayManageSofti database on your old server to match.

3. Disable RayManageSofti processing on the old server:

   a. Wait for any running distribution jobs to complete.
   (You can check if any distribution jobs are currently running from **Start** > **Settings** > **Control Panel** > **Scheduled Tasks**.)

   b. Remove write permission for the Everyone group on `C:\ManageSoft\Incoming`.
   This effectively disables **ManageSoftRL** (**ManageSoftRL$**), preventing any more log or inventory files being uploaded to it.

   c. Wait until RayManageSofti scheduled tasks have processed any log or inventory files from `C:\ManageSoft\Incoming` into the database before proceeding to the next step. (Files are removed once they have been successfully processed, so wait until the folder is empty.)

4. Transfer data from your old server to your new server. See *Transferring data between servers* for details.

5. Install RayManageSofti on your new server, identifying the new database.

6. Enable RayManageSofti processing on the new server:

a. On any other administration servers, update the distribution server properties to reflect the new location from which packages are pulled.

   Refer to *To change the administration server distribution server's properties* in the *Distribution system* chapter of the *RMS Software Deployment* for further details if required.

b. On this administration server, update any local distribution and reporting locations.

Refer to *To update a distribution location's properties* and *To update a reporting location's properties* in the *Distribution system* chapter of the *RMS Software Deployment* for further details if required.

c. Verify the distribution hierarchy.

Refer to *Managing your distribution hierarchy* in the *Distribution system* chapter of the *RMS Software Deployment* for further details if required.

d. Redistribute any managed device settings packages that include the administration server's distribution or reporting locations.

Refer to *Distribution and installation* in the *Software* chapter of the *RMS Software Deployment* for further details if required.

Scenario 4: Separate Data and Core Servers: Moving to a New Core Server

I have existing separate core and data servers. I want to move the core server to a new physical server.

Complete these steps:

1. On the server that currently hosts RayManageSofti ("old server"), check the collation of the RayManageSofti database. Refer to *About database collations* for more details about database collations.

2. On the new server, install all required system software, according to the specifications in the *RMS Implementation*.

Be aware:
If you are installing SQL Server, perform a custom install, and set the default database collation order to be the same as the collation order of the current RayManageSofti database (see *About database collations* for more details). Refer to SQL Server installation instructions for details about setting the default database collation order.

If SQL Server is already installed, check the collations of the `tempdb` and `master` databases, and change the collation of the RayManageSofti database on your old server to match.

3. Disable RayManageSofti processing on the old server. To do so:

   a. Wait for any running distribution jobs to complete.
   (You can check if any distribution jobs are currently running from **Start** > **Settings** > **Control Panel** **Scheduled Tasks**.)

   b. Remove write permission for the Everyone group on `C:\ManageSoft\Incoming`.

   This effectively disables **ManageSoftRL** (**ManageSoftRL$**), preventing any more log or inventory files being uploaded to it.

   c. Wait until RayManageSofti scheduled tasks have processed any log or inventory files from `C:\ManageSoft\Incoming` into the database before proceeding to the next step. (Files are removed once they have been successfully processed, so wait until the folder is empty.)

4. Transfer data (except for the database) from your old server to your new server. See *Transferring data between servers* for details.

5. Install RayManageSofti on your new server.

6. Enable RayManageSofti processing on the new server by completing these steps (refer to the *Distribution system* chapter of the *RMS Software Deployment* for further details if required):

   a. On any other administration servers, update the location from which packages are pulled.

   b. On this administration server, update any local distribution and reporting locations.

   c. Verify the distribution hierarchy.

   d. Redistribute any managed device settings packages that include the administration server's distribution or reporting locations.

About Database Collations

Microsoft SQL Server supports several collations. A collation encodes the rules governing the proper use of characters for either a language, such as `Greek` or `Polish`, or an alphabet, such as `Latin1_General` (the Latin alphabet used by Western European languages). Collations govern the physical storage of character strings in SQL Server.

You can read more about collations and SQL Server on the Microsoft website. When SQL Server is installed, a default collation is specified. Each database created uses this default collation if no collation is specified for the new database. Collations can also be specified for character columns. If no collation is specified, character columns use the default collation for the database.

When you are moving a RayManageSofti database to a new server, it is important that the collations of the `master`, `tempdb`, and `ManageSoft` databases are identical, or you will encounter collation compatibility issues, including errors such as:

```
-- This statement gets a code page conversion not allowed error
-- because the temporary table is created in tempdb, which has a
-- different default collation than ManageSoft.
```

Checking Database Collations

To check the collation for a database:

1. Open SQL Server Enterprise Manager (**Start** > **Programs** > **Microsoft SQL Server Enterprise Manager**).

2. In the console tree, expand the hierarchy until the **Databases** node is visible.

   Expand the **Databases** node.

3. Right-click the name of the database whose collation you want to check, and select **Properties**.

4. The collation name is displayed at the bottom of the dialog.

Transferring Data Between Servers

You need to transfer several sets of data from your old server to your new server:

- The RayManageSofti SQL database (and other databases, if you are using add-on products). See *Transferring databases between servers*.
- Other RayManageSofti data, including packages, schedules, and distribution hierarchy data. By default, this data is stored under `C:\ManageSoft`. See *Transferring RayManageSofti data between servers*.
- Configuration data. See *Transferring RayManageSofti configuration settings*.

Transferring Databases Between Servers

The process for transferring the RayManageSofti database between servers depends on whether the collation order of the database to be transferred is the same as the default collation order of SQL databases on the new server. Make sure you have read and understood the information about collations at *About database collations*.

If the default collation order and the collation order of the `tempdb` and `master` databases on the new server matches the collation order of the RayManageSofti database on your old server, you can simply back up the database on the old server, and restore it on the new server. See *If the collation order is the same* for details.

If the default collation order for databases on the new server, or the collations of the `tempdb` and `master` databases differ from the collation order of the RayManageSofti database on your old server, you must export data from the old database and import it to the new database using DTS procedures. See *If the collation order is different* for details.

Before you transfer data by either of these methods, you should stop scheduled tasks and services that operate

on the database.

Stopping Scheduled Tasks

Before backing up or exporting data from the RayManageSofti database, you should stop scheduled tasks that may affect it. Scheduled tasks must be stopped on both the core server and data server, so you may need to perform this process on more than one physical server. (Most RayManageSofti scheduled tasks run on the core server.)

To stop scheduled tasks:

1. Start Control Panel (**Start** > **Settings** > **Control Panel**).

   The **Control Panel** dialog is displayed.

2. Double-click **Scheduled Tasks**.

   The **Scheduled Tasks** dialog is displayed.

3. For each task whose name contains "`ManageSoft`", do the following:

   a. Right-click the task name.

   b. If the task is running, the **End** option is available. Select it to stop the task.

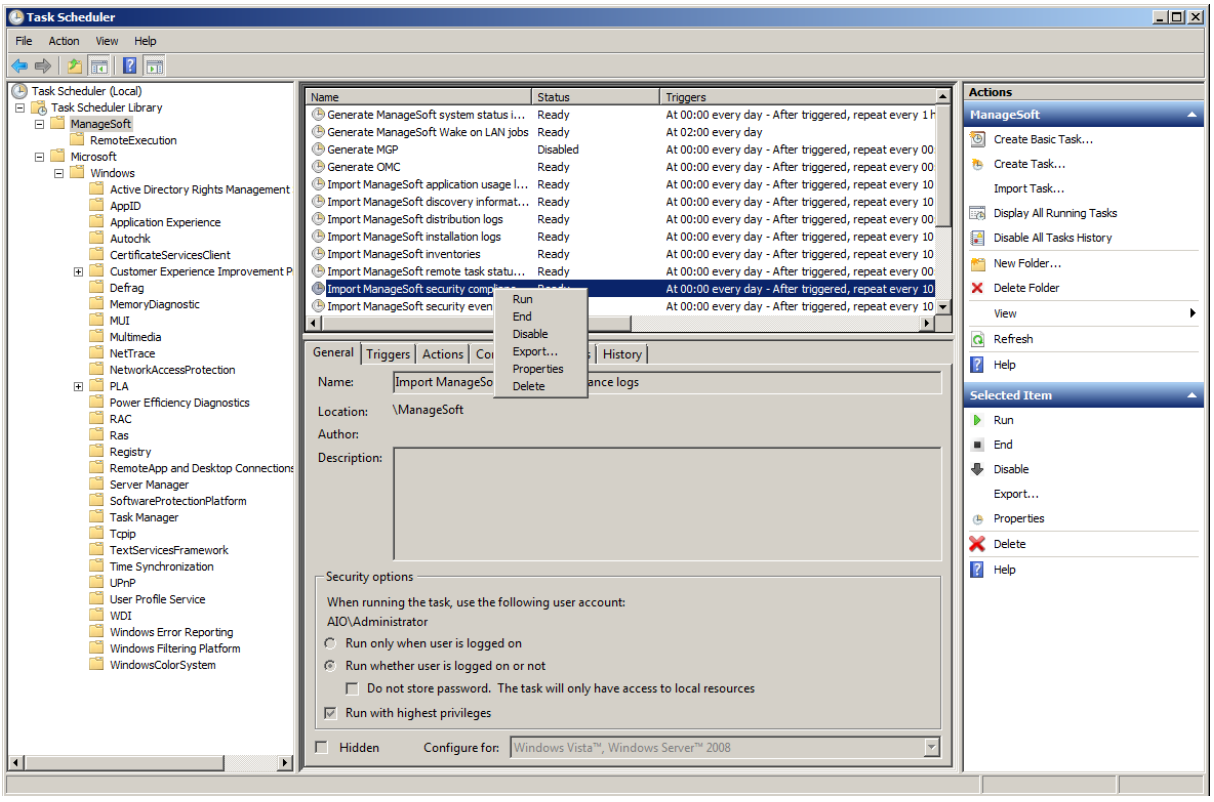

4. Disable each `ManageSoft` scheduled task.

   To do so, for each task whose name contains "`ManageSoft`", do the following:

- Right-click the task name and select **Disable**.

Stopping SQL Jobs

Before backing up or exporting data from the RayManageSofti database, you should stop SQL jobs that may affect the RayManageSofti database.

1. Start Enterprise Manager (**Start** > **Programs** > **Microsoft SQL Server** > **Enterprise Manager**).
2. In the console tree, drill down through **SQL Server Group** > **Local** > **Management** > **SQL Server Agent** and click **Jobs**.
3. For each job in the right hand pane whose name contains "`ManageSoft`", do the following:
   a. Right-click the job name and select **Disable job**.
   b. Right-click the job name and select **Stop job**.
4. Click **OK**.

If the Collation Order is the Same

If the collation order of the RayManageSofti database "`ManageSoft`" on your old server matches that of the `tempdb` and `master` databases on your new server, you can simply back up the database on the old server, and restore it on the new server. To do so:

1. Back up the RayManageSofti database:
   a. Start Enterprise Manager (**Start** > **Programs** > **Microsoft SQL Server** > **Enterprise Manager**).
   b. In the console tree, expand the hierarchy until the **Databases** node is visible.
   c. Expand the **Databases** node so that the RayManageSofti database "`ManageSoft`" is visible.
   d. Right-click `ManageSoft`, and select **All Tasks** > **Backup Database...**
   The SQL Server Backup - ManageSoft dialog is displayed.
   e. Make sure the **Database - complete** radio button is selected.
   f. In the **Destination** panel, specify the location to which to back up the database.
   g. Select the **Overwrite existing media** radio button.
   h. Click the **Options** tab.
   i. Check the **Verify backup upon completion** box.
   j. Click **OK** to start the backup.
   k. Click **OK** to close the dialog confirming that the backup has completed.
   l. Close Enterprise Manager.
2. Transfer the backup file to your new server.
3. Restore the RayManageSofti database:
   a. Start Enterprise Manager (**Start** > **Programs** > **Microsoft SQL Server** > **Enterprise Manager**).
   b. In the console tree, expand the hierarchy until the **Databases** node is visible.
   c. Right-click **Databases**, and select **All Tasks** > **Restore Database...**
   The **Restore database** dialog is displayed.
   d. In the Restore as database field, enter "`ManageSoft`".
   e. Make sure the **Database** radio button is selected.
   f. In the Parameters panel, select `ManageSoft` from the **Show backups of database** drop-down list, and

in the lower part of the panel select the backup file you transferred from your old server.

g. Select the **Options** tab and enter paths for the restored database and transaction log files.
h. Click **OK**.
The **Restore progress** dialog is displayed while the restore occurs.
i. Click **OK** to close the dialog that confirms that the restore operation completed successfully.
j. Close Enterprise Manager.

If the Collation Order is Different

If the collation order of the RayManageSofti database "`ManageSoft`" on your old server is different than that of the `tempdb` and `master` databases on your new server, you must export the data from the old database, and import it to the new one. To do so:

1. On your old server, back up the RayManageSofti database "`ManageSoft`":

   a. Start Enterprise Manager (**Start** > **Programs** > **Microsoft SQL Server** > **Enterprise Manager**).
   b. In the console tree, expand the hierarchy until the **Databases** node is visible.
   c. Expand the Databases node so that the "`ManageSoft`" database is visible.
   d. Right-click ManageSoft, and select **All Tasks** > **Backup Database...**
   The **SQL Server Backup - ManageSoft** dialog is displayed.
   e. Make sure the **Database - complete** radio button is selected.
   f. In the **Destination** panel, specify the location to which to back up the database.
   g. Select the **Overwrite existing media** radio button.
   h. Click the **Options** tab.
   i. Check the **Verify backup upon completion** box.
   j. Click **OK** to start the backup.
   k. Click **OK** to close the dialog confirming that the backup has completed.
   l. Close Enterprise Manager.

2. Use the DTS Import/Export wizard to transfer data from the RayManageSofti database on your old server into the new RayManageSofti database. On your old server:

   a. Right-click the `ManageSoft` database, and do one of the following:
      - Select **Tasks** > **Export Data...** (SQL Server 2005) The **SQL Server Import and Export** wizard starts.
      - Select **All Tasks** > **Export Data** (SQL Server 2000). The **Data Transformation Services Import/ Export** Wizard starts.
   b. Click **Next**.
   The **Choose a Data Source** page is displayed.
   c. In the **Server** pull-down list, choose the current server name or (local).
   d. From the **Database** pull-down list, choose `ManageSoft`.
   e. Click **Next**.
   The **Choose a destination** page is displayed.
   f. In the **Server** pull-down list, choose the name of the new server.
   g. From the Database pull-down list, choose `ManageSoft`.
   h. Click **Next**.
   The **Specify Table Copy or Query** page is displayed.
   i. Select **Copy objects and data between SQL Server databases**.
   j. Click **Next**.
   The **Select Objects to Copy** page is displayed.
   k. Check the **Use Collation** box.
   l. Uncheck the **Use default options** box.
   m. Click **Options...**
   The **Advanced Copy Options dialog** is displayed.

n. Check the **Copy SQL Server logins** (Windows and SQL Server logins) box.
Copying logins ensures that SQL Server logins such as **MGS Administrators**, which are required by the database, are created on the new server at the same time as the database. Otherwise, these logins are not created until RayManageSofti for administration servers is installed.
o. Click **OK**.
The **Select Objects to Copy** page is redisplayed.
p. Click **Next**.
The **Save, schedule, and replicate package** page is displayed.
q. Click **Next**.
The **Save DTS Package** page is displayed.
r. Click **Next**.
The **Complete the Wizard page** is displayed (SQL Server 2005) or the **Completing the DTS Import/ Export Wizard** page is displayed (SQL Server 2000).
s. Click **Finish**.
The RayManageSofti database "*ManageSoft*" and its data is transferred from your old server to your new server. This may take some time, depending on the size of your database.
t. When the database transfer has completed, click **OK** on the dialog that appears.
u. Click **Done**.

Starting Scheduled Tasks

On your old server, if you stopped the task scheduler earlier (see *Stopping scheduled tasks*), you should restart it now:

1. Start Control Panel (**Start** > **Settings** > **Control Panel**).
The **Control Panel** dialog is displayed.

2. Double-click **Scheduled Tasks**.
The **Scheduled Tasks** dialog is displayed.

3. From the **Advanced** menu, select **Start Using Task Scheduler**.
The task scheduler restarts, and tasks will run when next scheduled.

Enabling SQL Jobs

On your new data server, you must ensure that the SQL jobs that you disabled and stopped earlier (see *Stopping SQL jobs*) are enabled. To do so:

1. Start Enterprise Manager (**Start** > **Programs** > **Microsoft SQL Server** > **Enterprise Manager**).

2. In the console tree, drill down through **SQL Server Group** > **Local** > **Management** > **SQL Server Agent** and click **Jobs**.

3. For each job in the right hand pane whose name contains "`ManageSoft`", right-click the **job name** and select **Enable job**.

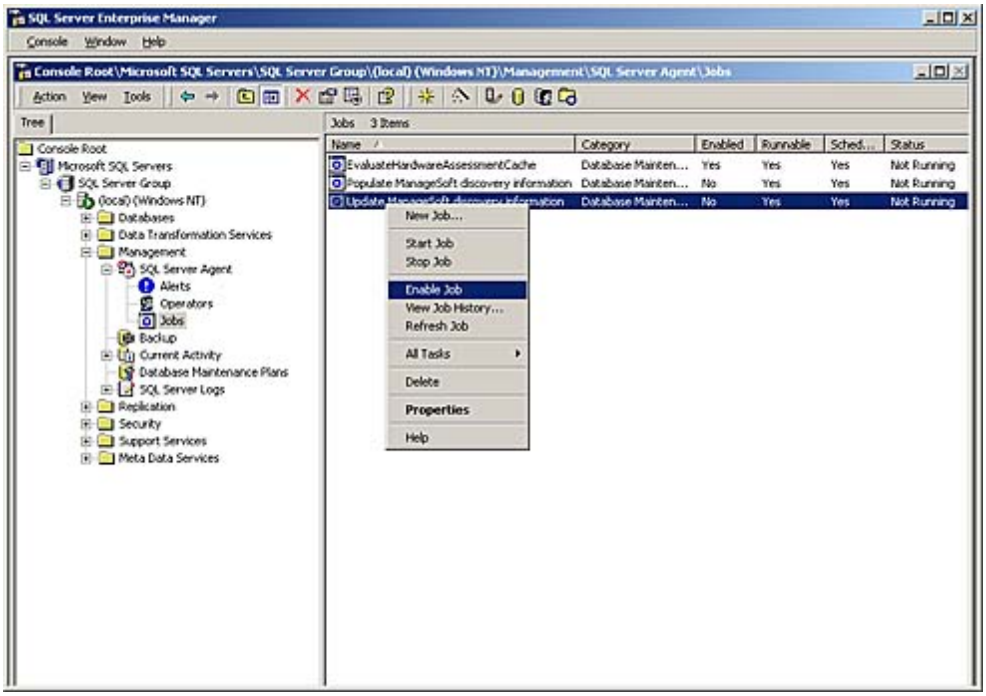

4. Click **OK**.

Transferring RayManageSofti Data Between Servers

RayManageSofti data is stored under `C:\ManageSoft`. Data such as packages, schedules, and distribution hierarchy data is stored under `C:\ManageSoft\Repository`. `C:\ManageSoft\Staging` contains files created when projects are packed.

In general, you will copy all data under `C:\ManageSoft` on your old server to your new server. The following exceptions may apply:

- If your SQL database exists under this location, do not copy it or its transaction logs, as you will copy the database separately.
- You may choose to copy only `<RayManageSofti>\Repository`. If you do so, packages will be repacked at the time they are distributed from your new server (because you have not copied the Staging area of your old server to your new server). The checksum of the new package may differ from the old (because of environmental differences on the new server such as locale setting, or daylight savings time), which may trigger unnecessary package updates.

To copy RayManageSofti data from your old server to your new server:

1. Make a copy of your RayManageSofti data from your existing installation.
2. Transfer these files to the new server.

Transferring RayManageSofti Configuration Settings

RayManageSofti configuration settings are stored in a series of `.config` files:

- Reporting configuration settings are stored in `C:\Program Files\ManageSoft\Reporter\Web\Web.config`
- Other configuration settings are accessible through the **Deployment Manager Configuration Console**

If you want to recreate your current configuration settings on your new server, after you have installed RayManageSofti on your new server, copy the appropriate `.config` files from the old server over their equivalent files on the new server.

You can export settings from the Deployment Manager Configuration Console on your old server, and import them on your new server. To do so:

1. The Deployment Manager configuration console is accessible from **Start** > **All programs** > **Deployment Manager Tools** > **Configuration** > **Deployment Manager Configuration**.

Note:
The **Product Activation Wizard** may display once you open the console, if you are using a temporary time out license or if the number of devices you are managing approaches or exceeds the number you are licensed to manage. If it does appear, select **Start** to go ahead.

2. Click on the root entry **Deployment Manager Configuration**.
3. Select **Export your Deployment Manager configuration** from the action list in the details pane on the right.
   The **Export to** dialog is displayed.
4. Complete the **File name** field by entering, or browsing for, a location and file name at which to store the configuration data.
5. Click **Save**.
6. The configuration data is exported.
7. On the dialog that reports that the configuration data has been successfully exported, click **OK**.
8. Transfer the file of configuration data to your new server.
9. On your new server, start the **Deployment Manager Configuration Console**.
10. Click on the **root** entry Deployment Manager Configuration and select **Import your Deployment Manager configuration** from the action list in the details pane on the right.
    The **Import from** dialog is displayed.
11. Complete the **File name** field by entering, or browsing for, the location and file name at which the configuration data is stored.
12. Click **Open**.
    The configuration data is imported.
13. On the dialog that reports that the configuration data has been successfully exported, click **OK**.

Updating RayManageSofti Settings

If you are continuing to use the RayManageSofti installation on your old server, you must update it to point to the location of the database on your new server. To do so, complete these steps on your core server:

1. The **Deployment Manager Configuration Console** is accessible from **Start** > **All programs** > **Deployment Manager Tools** > **Configuration** > **Deployment Manager Configuration**.

Note:
The **Product Activation Wizard** may display once you open the console, if you are using a temporary time out license or if the number of devices you are managing approaches or exceeds the number you are licensed to manage. If it does appear, select **Start** to go ahead.

2. In the console tree, click **Database Population**.

   The **Configure your database population settings** page is displayed, with the Database connection string selected for editing.

3. Edit the contents of the **Database connection string** field to record the name of the server on which you have created the new database.

   Example old string:
   ```
   Driver={SQL Server};
   Server=AS-BPO-01;
   Database=ManageSoft;
   Trusted_Connection=yes;
   ```

   Example new string:
   ```
   Driver={SQL Server};
   Server=AS-CLS-02;
   Database=ManageSoft;
   Trusted_Connection=yes;
   ```

4. Click **OK**.

5. Close the **Deployment Manager Configuration** window.

6. Restart **Deployment Manager** so that your changed settings will take effect.

Changing Server and Database Connection Values for Reports

You will also need to change server and database connection values in the reports `Web.config` file to point to the new database. To do so:

1. Browse to your RayManageSofti reporting installation, and open the configuration file `Web.config` in the text editor of your choice.
   (The default location is `C:\Program Files\ManageSoft\Reporter\Web\Web.config`.)

2. In the `DATA ACCESS CONFIG` section, locate the setting for the server connection value, for example `<add key="ManageSoft.DataAccess.server" value="AS- BPO-01"/>`.
   Set the value to the new server name.

3. Make sure the values of `ManageSoft.DataAccess.database` and `ManageSoft.DataAccess.trusted_connection` are appropriate for the new server and database.

4. Save and close the file.

Migration Checklist

To confirm that your RayManageSofti data and SQL database(s) have been successfully migrated to your new server(s), check that:

- You can view reports
- From the RayManageSofti console, you can view details about the distribution hierarchy
- You can distribute packages throughout your distribution hierarchy
- Distribution results are appearing on the administration server
- Log and inventory files are being uploaded to the administration server
- Data from log and inventory files is being resolved into the database

Bandwidth Settings

When you define and configure distribution servers and distribution locations (see the *Distribution* chapter of the *RMS Software Deployment* for details), you can specify the amount of bandwidth that can be used per second for data transfer between the administration server and the distribution server/distribution location.

Specifying Your Own Bandwidth Settings

RayManageSofti supplies a number of default bandwidth settings, but you can specify your own, additional, bandwidth settings using this process:

1. On your administration server, start **regedit**.

2. Create the registry key `[Registry]\ManageSoft\Distributor\CurrentVersion\NetworkBandwidthSettings`.

   Under the registry key you created at the previous step, create a string value for each bandwidth setting that you want to create. Use the name `Bandwidth<N>` and set the value to `<number> <unit>`, where:

   - `<N>` is an integer
   - `<number>` is the value of the bandwidth to be used. Entries are rounded to one decimal place for display purposes.
     For example, if you enter `128.765, 128.8` is displayed.

   The unit of bandwidth is not changed where rounding occurs, but values are always written in kilobits to the configuration file. Rounding will occur when converting values to kilobits for storage in the configuration file.

   In the configuration file, `128.765` Kbits will be *recorded* as `128.8` kilobits
   If you enter `0.01` Mbits, it will be *displayed* as `0.0` Mbits, but *behave* as `10` Kbits. (`10` Kbits is the value that will be written to the configuration file.)

   - A space follows `<number>`.
   - `<unit>` is one of (units are presented in size order):

     - Kilobits (or kbit, kbits, kilobit, kilobits, kbps)
     - Megabits (or mbit, mbits, megabit, megabits, mbps)
     - Gigabits (or gbit, gbits, gigabit, gigabits, gbps)
     - Terabits (or tbit, tbits, terabit, terabits, tbps)

     If no unit is specified, Kilobits is assumed.

     Examples:
     Bandwidth1 128.8 kbps, Bandwidth2 64 terabits, Bandwidth3 500.

     The settings you enter are immediately available in the **Available bandwidth** pull-down list on the B**andwidth Management** dialog for affected distribution servers and distribution locations.

3. When you have finished entering bandwidth settings, close **regedit**.

Data and Log File Processing

When operations such as event processing, discovery, and installation occur on distribution servers or managed devices, data about the results of the operations are uploaded through the distribution hierarchy and added to the RayManageSofti database. (Depending how your environment is configured, data may be uploaded all the way to the administration server, and processed into the database from there, or processed on distribution servers that can remotely update the database on the administration server.)

The types of data added to the database include:

- Data about devices discovered on your network
- Logging data about distribution activities throughout the RayManageSofti hierarchy
- Logging data about tasks performed on managed devices, such as inventory gathering, policy application, package installation, remote task execution, and others

This chapter describes the methods by which data are imported into the RayManageSofti database, and discusses the ways in which you can configure this behavior. There are some general configuration items, and also specialized configuration you can perform to:

- Help optimize the performance of your administration server or distribution server
- Facilitate troubleshooting

How Files are Processed

Depending on their type, files are processed in one of two ways:

- By an ASP.NET web application
- By a scheduled task that runs a RayManageSofti importer

Importing Data Using IIS

When Deployment Manager is installed, IIS on administration servers (and, optionally, distribution servers) is configured so that the ManageSoftRL virtual directory is associated with a web application. This web application intercepts all files uploaded to this location, and - depending on their destination URL - either handles them directly, or passes them on to be handled by some other process.

How individual files are handled is determined by the location under the **ManageSoftRL** virtual directory to which each file is uploaded. For example:

- Files uploaded to `ManageSoftRL\discovery` are processed as files of data about discovered devices
- Files uploaded to `ManageSoftRL\logs` are processed as files of data about package installation on managed devices

The file extension is of secondary importance, and is mostly used to apply security directives, such as preventing the upload of files with `.exe` extensions, as they may be malicious. A file with a `.disco` extension uploaded to `ManageSoftRL\logs` will be treated as if it were a file of package installation data, and processing will

therefore fail, as its data will not be in the expected format. The default configuration of IIS during Deployment Manager installation ensures that upload locations for all file types are properly configured - for example, files with `.disco` extensions are uploaded to `ManageSoftRL\discovery`.

By default, if file processing fails, the file will be saved. The location to which it is saved, and whether (and when) reprocessing is attempted, depends on the type of error and configuration options. See *Retrying failed items* for more detail.

File types that are not handled by the IIS web application - for example, files of application usage data - are saved to the file system, and are later processed by scheduled tasks.

Importing Data Using RayManageSofti Importers

Data from files that are not handled by IIS, or for which IIS's processing attempt failed, are added into the RayManageSofti database by RayManageSofti data importers. There are different importers for different data file types, as summarized in the *Command line tools* chapter of *RMS System Reference*.

Although importers are typically run by scheduled tasks (see the *Scheduled tasks* chapter of *RMS System Reference*), you can also run them from the command line. Details of how to do this are provided in the *Command line tools* chapter of *RMS System Reference*.

Configuring Data File Processing

You can perform some configuration of how data and log files are processed using the `web.config` file located in `C:\Program Files\ManageSoft\DotNet`.

It is an XML file that can be edited in the text editor of your choice.

Some of the configuration that can be performed through the `web.config` file is described here. You will notice other content in the file that is not discussed here - in general, that is an indication that you should not be changing the default settings without guidance from your Raynet consultant.

When you make changes to the `web.config` file, ASP.NET restarts the web application, causing `web.config` to be reloaded automatically. This means that changes take effect immediately. You do not need to restart IIS. When the web application restarts, statistics are reset.

File Types

File types accepted for resolution into the database

By default, files with the following extensions will be accepted and processed into the database:

- `.disco`
  Files containing data about discovered devices
- `.gz`
  Compressed files, typically application usage (`.mmi`), inventory (`.ndi`), or security compliance (`.msa`) data files
- `.log, .txt, .xml`

Files containing status data

- `.mmi`
  Files containing application usage data
- `.msa, .msel`
  Files containing security compliance data
- `.mss`
  Files containing distribution server status data
- `.ndi`
  Files containing inventory data
- `.rea, .red, .rer, .rev, .taskstatus`
  Files of remote task execution data

You can read more about the structure and contents of each of these file types in the *File formats* chapter of *RMS System Reference*.

If for some reason you do not want to import data from one or more of these file types into the RayManageSofti database, you can remove the file extension from the list of accepted file types. To do so:

1. Open `web.config` for editing.
2. Locate this line: `<add key="AllowedExtensions" value="gz, txt, xml, log, ndi, disco, msa, mmi, rea, red, rer, rev, taskstatus, mss, msel" />`
3. Remove the file extensions that you do not want processed.
4. Save and close the file.

The administration server or distribution server will refuse files with any extensions not listed here, generating an `HTTP 403 FRAMEWORK_REJECTED_BY_SECURITY` error.

File Types to Process Using an IIS Web Application

Of the file types accepted for processing into the database, a subset is handled by an IIS web application. Whether files are processed by the IIS web application or by some other mechanism depends on the location under the **ManageSoftRL** virtual directory to which each file is uploaded.

By default, files of the types listed in the table below are processed by an IIS web application.

Data type	URL to which data is uploaded (see note)	File extension
Application usage	`ManageSoftRL/usagedata`	`.mmi`
Log of attempted use of prohibited hardware devices	`ManageSoftRL/SecurityEventLogs`	`.msel`
Discovery	`ManageSoftRL/discovery`	`.disco`
Distribution server system status	`ManageSoftRL/systemstatus`	`.mss`

Data type	URL to which data is uploaded (see note)	File extension
Installation logs	`ManageSoftRL/logs`	`.log`
Inventories	`ManageSoftRL/inventories`	`.ndi and .ndi.gz`
Security compliance (generated by security analysis)	`ManageSoftRL/securityanalysis`	`.msa`

Note:
This location is the internally-recognized location under the reporting location (**ManageSoftRL**) to which files are directed through the distribution hierarchy.

Data may not ever be stored as physical files. Data that are successfully processed into the RayManageSofti SQL database by a web application is not ever stored as a physical file on the administration server or distribution server that performs the processing. Data that fail processing, or that are to be processed by a different mechanism than a web application, may be stored in a file in a directory under `C:\ManageSoft\Incoming`, in subfolders including:

- `Discovery`
- `Inventories`
- `Logs`
- `SecurityAnalysis`
- `SecurityEventLogs`
- `SystemStatus`
- `UsageData`

File extensions are checked against a list of prohibited extensions so that, for example, RayManageSofti administrators can prevent the upload of executable files, which may present a security risk.

You can edit the `web.config` file to change the types of files to be processed using IIS. This can be useful, for example, if you are troubleshooting and want to view files before they are processed.

To edit the list of file extensions which are automatically handled by an IIS web application:

1. Open `web.config` for editing.

2. Locate this line:

```
<add key="ImportTypesProcessedDirectly" value="logs, inventories, discovery,
securityanalysis, systemstatus, usagedata, securityevent" />
```

3. Remove the file types that you do not want processed by the web application.

4. Save and close the file.

Files of types that are not handled by IIS are saved on the administration or distribution server, into a subdirectory of `C:\ManageSoft\Incoming`, and processed by scheduled tasks running `mgsimport.exe` or other data importers (as described in the *Scheduled tasks* chapter of *RMS System Reference*).

If you do not want RayManageSofti data importers to process the files, for example so that they remain on disk for you to examine while troubleshooting, you must also disable the scheduled tasks that process files from `C:\ManageSoft\Incoming` and its subdirectories:

If you stop IIS handling files of this type...	Disable this scheduled task...
discovery	Import ManageSoft discovery Information
inventories	Import ManageSoft inventories
logs	Import ManageSoft installation logs
securityanalysis	Import ManageSoft security compliance logs
securityevent	Import ManageSoft security event information
systemstatus	Import ManageSoft system status information
usagedata	Import ManageSoft application usage logs

General Configuration Options

There are some general configuration tasks you can perform, including:

- Configuring RayManageSofti so that file processing requests are rejected if there are known problems with the running environment that will prevent files from being processed
- Specifying when file processing status should be returned to distribution servers and managed devices further down the distribution hierarchy
- Dictating what should happen if file processing fails
- Specifying whether to perform file processing synchronously or asynchronously

These options are discussed in turn in this section.

Later sections cover:

- Configuration for performance
- Configuration for troubleshooting

Rejecting Requests if the Environment is Not Ready

You can specify that file processing requests should be rejected when the database is not available, as this environmental problem will prevent successful file processing.

To do this, you nominate the number of requests that must fail because of environmental problems, and the frequency with which the environment should be polled for readiness:

1. Open `web.config` for editing.
2. Locate this line:

   ```
   <add key="EnvironmentGoverningThreshold" value="10" />
   ```

   After this number of requests have failed due to environmental factors, RayManageSofti will automatically reject further file processing requests until an environmental poll indicates that file processing is likely to succeed again.

3. Edit the value if you wish.

4. Locate this line:

   ```
   <add key="EnvironmentCheckingInterval" value="60" />
   ```

   This specifies the number of seconds between polls of the environment. This setting is only used when the value of `EnvironmentGoverningThreshold` is greater than 0.

5. Edit this number if you wish.

6. Save and close the file.

File Downloads and Status Reporting

Deployment Manager can report the status of file processing:

- When it has received a complete file from further down the distribution hierarchy (this is the default)

  If success is reported when a server has received a complete file, distribution servers and managed devices do not need to keep copies of the file any more, or perform any further processing of them. (Even if processing of the file fails after it has reported that it has successfully received the file, the file will not be uploaded again through the distribution hierarchy. The administration or distribution server is responsible for dealing with the failed file processing. See *Retrying failed items*.)

- When it has completed file processing.

  If success is only reported when the administration or distribution server has successfully processed the file, processing failure will result in a failure code being returned to distribution servers and managed devices further down the distribution hierarchy. This failure prompts them to resubmit the file for upload through the distribution hierarchy and processing on the administration or distribution server.

In its default configuration, the administration server or distribution server returns a success code when it has received a full copy of the file. You can change this behavior.

1. Open `web.config` for editing.

2. Locate the line that contains:

   ```
   <add key="CommitContentToServerSideWhenFullyReceived" ... />
   ```

   To return a success code after the file has been received, ensure its value is set to `true`:

   ```
   <add key="CommitContentToServerSideWhenFullyReceived" value="true" />
   ```

   To return a success code only after the file has been processed, ensure its value is set to `false`:

   ```
   <add key="CommitContentToServerSideWhenFullyReceived" value="false" />
   ```

3. Locate the line that contains:

   ```
   <add key="ReceiveFullContentInAdvance" ... />
   ```

   If you want a success code to be returned when a file has been received, ensure this value is set to `true`:

```
<add key="ReceiveFullContentInAdvance" value="true" />
```

This ensures that the administration server or distribution server does not start processing the file until it has received it entirely.

If this value is false, processing will commence while the file is being received by the administration or distribution server, which is incompatible with returning a success code before processing.

4. If you set `ReceiveFullContentInAdvance=true`, locate the line that contains:

```
<add key="MaxSizeLimitOnReceiveFullContentInAdvance" ... />
```

Only files smaller than the size specified in `MaxSizeLimitOnReceiveFullContentInAdvance` will be received in advance. (This value is in KB. So if `10240` is specified, only files smaller than 10MB will be received completely in advance of processing.)

Processing of files greater than this size will commence while the rest of the file is being transferred. To always receive files in advance before processing, regardless of size, set the value of `MaxSizeLimitOnReceiveFullContentInAdvance` to `0`:

```
<add key="MaxSizeLimitOnReceiveFullContentInAdvance" value="0" />
```

5. Save and close the file.

Retrying Failed Items

Using a Web Application

When the web application fails to process incoming data, it attempts to determine if the failure is transient or permanent. Transient errors can be caused by failures to connect to the database, a database concurrency or locking problem, and so on. Permanent failures are caused by problems like invalid formats.

You can configure a list of SQL errors to be treated as permanent errors. See *Blacklisting SQL errors*.

If the administration or distribution server is configured with `CommitContentToServerSideWhenFullyReceived=true`, data whose processing failed due to a transient error are saved on the server. The location to which data files are saved depends on the setting of `AutoRetryCommittedItems`:

- If `AutoRetryCommittedItems` is `true`, files are saved into a subdirectory of `C:\ManageSoft\Incoming`, according to the type of the file. (For example, discovery data whose processing fails are saved to `C:\ManageSoft\Incoming\Discovery`.) Files in these directories will be reprocessed by scheduled tasks that run the data importer `mgsimport.exe`. You can read about `mgsimport.exe` in the *Command line tools* chapter of *RMS System Reference*, and about scheduled tasks in the *Scheduled tasks* chapter of the same volume.

  Files whose processing failed due to a permanent error are saved into a `C:\ManageSoft\Incoming\*\BadLogs` directory, where `*` varies according to the type of the file.

  Within `BadLogs`, files are saved to one of four subdirectories according to the type of error encountered during processing:

  - `invalidcompression`, for data that could not be successfully decompressed
  - `invalidformat`, for data that is not in the expected format

- `invaliddata`, for data that is unusable for some reason other than formatting errors (for example, data that contains references to distribution servers that do not exist)
- `failure`, for failures that do not fit into any of the preceding categories

- If `AutoRetryCommittedItems` is `false`, files are saved into `C:\ManageSoft\*\ReTry`, where `*` varies according to the type of the file.

  Files in these subdirectories will not automatically be reprocessed, but you can manually attempt to reprocess them as part of your troubleshooting operations. To do so, run the appropriate data importer from the command line, or move the files to the Incoming directory above the `ReTry` directory and trigger the scheduled task that processes the files.

To set the value of `AutoRetryCommittedItems`:

1. Open `web.config` for editing.
2. Locate this line: `<add key="AutoRetryCommittedItems" value="true" />`
3. Set the value to `true` or `false`, as required.
4. Save and close the file.

Using Mgsimport.exe

For files processed by `mgsimport.exe` that generate errors, `mgsimport.exe` attempts to determine if the failure is transient or permanent.

Files whose processing failed due to a transient error are retained on the processing server in the same location. Processing will be attempted on those files next time the scheduled task runs.

Files whose processing failed due to a permanent error are moved to a `BadLogs` subdirectory below the faulty file's current location.

Using Deployment Manager importers (other than Mgsimport.exe)

Files that are processed by other importers and that do not process successfully are moved to a `BadLogs` subdirectory below the file's current location.

Blacklisting SQL Errors

You can specify a list of SQL errors that are to be regarded as permanent. When a permanent SQL error is encountered, no further attempt is made to process the data file that caused the error. For transient SQL errors, further attempts may be made to process the data file - see *Retrying failed items*.

To specify which SQL errors are permanent:

1. Start Deployment Manager.
2. Use the mouse to hover over the **Configuration** icon, and select **Configure Administrator** from the context menu that appears.

   The **Configuration Management** page is displayed.

3. Select **Configure your Database Population**.

   The **Configure your Database Population** page is displayed.

4. Select **Fail import on SQL exception codes**.

   A field is displayed in the bottom half of the page for you to modify the setting.

5. Enter the SQL exception codes that should be treated as permanent failures. Separate exception codes with commas.
   For example, `4060,18456,547,2627,2601`.

Be aware:
The list of SQL exception codes is stored at `[Registry]\ManageSoft\SQLExceptionBlackList`.

6. Click **OK**.

7. The list of exception codes is stored. Data files that prompt these errors will not be re-processed.

Synchronous or Asynchronous Processing

You can choose to process files that are uploaded to the administration or distribution server synchronously or asynchronously. You are advised to leave the default - asynchronous - configuration in place unless advised by your Raynet consultant to change it.

In other parts of the `web.config` file, you will see settings that take effect only when asynchronous mode is in operation, or only when synchronous mode is in operation. Again, leave the defaults unchanged unless advised otherwise by your Raynet consultant.

Configuring for Performance

This section covers configuration that affects performance of your administration server. It includes details about obtaining statistics about current performance, and discusses items that can be configured specifically for your environment.

File Processing Statistics

The optimum values to specify when configuring file processing for your organization depend on the specifics of your environment. Deployment Manager can produce statistics about file processing, throughput, and load that can help you to configure processing appropriately for your environment.

There is a "pipeline" for file processing: files enter the pipeline when they are accepted for processing, they are opened and read, their data is imported to the database, and files then exit the pipeline. Statistics are available for the overall time that files are in the pipeline, and also for the time that files are actively being processed (their data is being loaded into the database), and time lost to overhead such as opening files, verifying user identity and permissions, queueing files, and so on.

Available Statistics

The following statistics are available:

- `TrackProcessingPerformanceStats` turns on calculations of how long the overall processing of files takes from the moment files enter the pipeline to the moment they exit.

- `TrackProcessingLoadStats` turns on calculations that determine how many files are in the pipeline simultaneously.

- `TrackProcessingThroughputStats` turns on monitoring of throughput (the rate at which files exit the pipeline) in number of files per hour.

- `TrackExecutionPerformanceStats` calculates how long the actual data import phase of processing takes. This is different than `TrackProcessingPerformanceStats` because:

  - It only includes the time spent actively importing data, not other time spent in the pipeline
  - Some files never reach the stage of having their data imported (if they are rejected because of invalid formatting, for example)

- `TrackExecutionLoadStats` calculates how many files are simultaneously being processed in the data import phase (how many threads are concurrently executing import routines)

- `TrackExecutionThroughputStats` monitors throughput for the data import phase (the rate at which files exit the active data import phase) in number of files per hour.

- `TrackOverheadAndDelaysStats` calculates how much time is lost to overheads such as:

  - Files waiting in a queue for processing
  - Accessing the file system
  - Waiting for operating system services
  - Multi-threading switching, queueing, and locking
  - User identity switching

  Time lost to these overheads is calculated as `Time spent in pipeline - Time spent in active data import phase`.

- `TrackAdvanceContentReceivingPerfStats` calculates how long is spent receving content from files. This can only be measured when `ReceiveFullContentInAdvance` is `true`. Otherwise, its value is `0`.

- `TrackAdvanceContentDecompressingPerfStats` calculates how much time is spent receiving and decompressing content from files. This can only be measured when `ReceiveFullContentInAdvance` is `true`. Otherwise, its value is `0`.

Statistics are reset each time the web application restarts (for example, after changes are made to `web.config`.)

Specifying what Statistics to Monitor

In `web.config`, ensure that the values of the appropriate keys are set to true for the statistics you want to monitor:

```
<add key="TrackProcessingPerformanceStats" value="true" />
<add key="TrackProcessingLoadStats" value="true" />
<add key="TrackProcessingThroughputStats" value="true" />
<add key="TrackExecutionPerformanceStats" value="true" />
<add key="TrackExecutionLoadStats" value="true" />
<add key="TrackExecutionThroughputStats" value="true" />
```

```
<add key="TrackOverheadAndDelaysStats" value="true" />
<add key="TrackAdvanceContentReceivingPerfStats" value="true" />
<add key="TrackAdvanceContentDecompressingPerfStats" value="true" />
```

Specifying How Frequently to Print or Publish Statistics

If configured to do so, Deployment Manager gathers statistics about file processing. These statistics can be:

- Written to a trace file, with the frequency specified by `StatsTracingInterval`
- Published to Windows performance counters, immediately, or with the frequency specified by `StatsPublishingInterval`

You can view statistics published to Windows performance counters through **Performance Logs and Alerts** (**Start** > **Settings** > **Control Panel** > **Administrative Tools** > **Performance** > **System Monitor**).

To specify when statistics should be written to the trace file or published to Windows performance counters:

1. Open `web.config` for editing.
2. Locate this line: `<add key="StatsTracingInterval" value="0" />`
3. Edit the value to reflect the number of file processing requests after which statistics should be written to the trace file. A value of `0` means that statistics will never be written to the trace file.
4. Locate this line: `<add key="StatsPublishingInterval" value="0" />`
5. Edit the value to reflect the number of file processing requests after which statistics should be published to Windows performance counters. A value of `0` means that statistics will never be published.
6. Locate this line: `<add key="PublishAveragesDirectly" value="true" />`

   Windows performance counters automatically calculate and display averages for all statistics being published.

   For throughput-based statistics, these may not be accurate, since they are averages of the reported values.

   The following scenario assumes that Windows performance counter statistics are published every 100 file requests.

   Imagine that 100 files are processed in one minute. In order to publish hourly statistics, this figure is extrapolated to an hourly throughput of 6000 files. Imagine that the next 100 files take 59 minutes to process. This is extrapolated to a rounded hourly throughput figure of 100 files. The average of the two sets of published statistics shows an average hourly throughput of 3050 files.

   By contrast, Deployment Manager calculates the average number of files processed per hour based on the total number of files processed and the elapsed time. It will show an average hourly throughput of 200 files.

   To calculate and print/publish Deployment Manager averages with other statistics, make sure `PublishAveragesDirectly` is set to `true`. A value of `false` means that Deployment Manager-calculated averages will not be calculated and printed or published.
7. Save and close the file.

Limiting the Numbers of Files to Process

You can restrict the number of files accepted by the administration or distribution server for processing at one time. This is useful for managing the performance of the server, preventing file processing requests from causing system overload.

When the maximum number of files has been accepted for processing, incoming files from further down the distribution hierarchy are rejected. Distribution servers and managed devices will automatically retry failed uploads, so the files will be accepted and processed later.

The numbers and types of files that can be processed simultaneously without adversely affecting performance of your administration or distribution server will depend on its specifications, the configuration of your Deployment Manager implementation, and your network configuration. A Raynet consultant can help you to perform benchmark tests and determine suitable values for your environment. It is also useful to gather statistics in your environment. See *File processing statistics* for details about settings in `web.config` that govern statistics gathering.

WorkloadMaxLimit

By default, the maximum number of files that will be accepted for processing by the IIS web application is 100. This is specified by:

```
<add key="WorkloadMaxLimit" value="100" />
```

You can edit this number if required.

WorkloadMaxLimit Settings

Alternatively, you can limit the numbers of files to be processed by specifying maximum numbers for particular types of log files. To do so:

1. Open `web.config` for editing.
2. Make sure that the value of `WorkloadMaxLimitAppliesPerImportType` is `true`:

   ```
   <add key="WorkloadMaxLimitAppliesPerImportType" value="true" />
   ```

3. Edit the following lines to specify the maximum number of each type of file processing request that can be in the processing pipeline on the administration or distribution server at one time:

   ```
   <add key="WorkloadMaxLimit-logs" value="100" />
   <add key="WorkloadMaxLimit-inventories" value="100" />
   <add key="WorkloadMaxLimit-discovery" value="100" />
   <add key="WorkloadMaxLimit-securityanalysis" value="100" />
   <add key="WorkloadMaxLimit-systemstatus" value="100" />
   ```

4. Save and close the file.

Limiting the Numbers of Files Processed Simultaneously

Data imports to the RayManageSofti database are multi-threaded, meaning that data from more than one file can be imported at one time. You can limit the number of simultaneous data import operations, to manage performance of your administration or distribution server.

ConcurrencyMaxLimit

To set an overall limit on the number of files to be processed simultaneously, set value of `ConcurrencyMaxLimit`:

```
<add key="ConcurrencyMaxLimit" value="10" />
```

ConcurrencyMaxLimit Settings

Rather than limiting the overall number of files that can be simultaneously processed, you can limit the number of a particular type of file to be processed simultaneously:

1. Open `web.config` for editing.
2. Make sure that the value of `ConcurrencyMaxLimitAppliesPerImportType` is `true`:

   ```
   <add key="ConcurrencyMaxLimitAppliesPerImportType" value="true" />
   ```

3. Edit the following lines to specify the maximum number of each type of file that can be present on the administration or distribution server at one time:

   ```
   <add key="ConcurrencyMaxLimit-logs" value="10" />
   <add key="ConcurrencyMaxLimit-inventories" value="10" />
   <add key="ConcurrencyMaxLimit-discovery" value="10" />
   <add key="ConcurrencyMaxLimit-securityanalysis" value="10" />
   <add key="ConcurrencyMaxLimit-systemstatus" value="10" />
   ```

4. Save and close the file.

Configuring the Age of Cached Data

Some data - such as data extracted from Active Directory, and security compliance data - does not change frequently. To optimize performance, the IIS web application and Deployment Manager importers use a cached copy of this data when required. You can configure the interval at which the cached copies of data are refreshed:

1. Open `web.config` for editing.
2. Locate this line: `<add key="DirectoryBinder_CacheRefreshInterval" value="120" />`
3. If you wish, change the frequency with which the cache of data extracted from Active Directory is updated. The value is in seconds.
4. Locate this line: `<add key="SecurityAnalysisImport_CacheRefreshInterval" value="1200"/>`
5. If you wish, change the frequency with which the security compliance data cache is updated. The value is in seconds.
6. Save and close the file.

Configuring for Troubleshooting

This section discusses configuration changes that you might make temporarily during troubleshooting operations.

General configuration options, and performance-specific configuration options, are discussed earlier in this chapter.

Tracing Data File Processing

If you want to conduct tracing on file processing operations, you must turn on the appropriate level of tracing in the trace profile (`.trace`) file, as discussed in the *Tracing and logging* chapter of *RMS System Reference*.

In addition, to ensure that tracing is reinitialized after you make changes to the trace profile file, ensure that `DynamicTracingReInitialization` is `true`:

1. Open `web.config` for editing.
2. Locate this line: `<add key="DynamicTracingReInitialization" value="true" />`
3. If the value is currently `false`, set it to `true`.
4. Save and close the file.
5. If `DynamicTracingReInitialization=false`, you must restart IIS after making changes to `etap.trace`, before the changes will take effect.

RayManageSofti is part of the RaySuite

More information online
www.raynet.de

Raynet GmbH
Technologiepark 20
33100 Paderborn, Germany
T +49 5251 54009-0
F +49 5251 54009-29
info@raynet.de

www.raynet.de