Interoperability Guide

GroupWise 2012

April 16, 2013

Novell

Legal Notices

Novell, Inc. makes no representations or warranties with respect to the contents or use of this documentation, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, Novell, Inc. reserves the right to revise this publication and to make changes to its content, at any time, without obligation to notify any person or entity of such revisions or changes.

Further, Novell, Inc. makes no representations or warranties with respect to any software, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, Novell, Inc. reserves the right to make changes to any and all parts of Novell software, at any time, without any obligation to notify any person or entity of such changes.

Any products or technical information provided under this Agreement may be subject to U.S. export controls and the trade laws of other countries. You agree to comply with all export control regulations and to obtain any required licenses or classification to export, re-export, or import deliverables. You agree not to export or re-export to entities on the current U.S. export exclusion lists or to any embargoed or terrorist countries as specified in the U.S. export laws. You agree to not use deliverables for prohibited nuclear, missile, or chemical biological weaponry end uses. See the Novell International Trade Services Web page (http://www.novell.com/info/exports/) for more information on exporting Novell software. Novell assumes no responsibility for your failure to obtain any necessary export approvals.

Novell, Inc.
1800 South Novell Place
Provo, UT 84606
U.S.A.
www.novell.com

Online Documentation: To access the online documentation for this and other Novell products, and to get updates, see the Novell Documentation Web site (http://www.novell.com/documentation).

Novell Trademarks

For Novell trademarks, see the Novell Trademark and Service Mark list (http://www.novell.com/company/legal/trademarks/tmlist.html).

Third-Party Materials

All third-party trademarks are the property of their respective owners.

Contents

About This Guide

This Novell *GroupWise 2012 Interoperability Guide* helps you use GroupWise in the context of other software products. The guide provides assistance with Novell products and third-party products:

Novell Products	◆ Part I, “Novell Cluster Services on Linux,” on page 11 ◆ Part II, “Novell Vibe,” on page 107 ◆ Part III, “Novell ZENworks,” on page 117 ◆ Part IV, “Other Novell Products,” on page 147
Third-Party Products	◆ Part V, “Microsoft Clustering on Windows,” on page 155 ◆ Part VI, “Non-GroupWise Email Clients,” on page 211 ◆ Part VII, “Mobile Devices,” on page 221

For information about additional GroupWise-related software from GroupWise partners, see the Novell Partner Product Guide (http://www.novell.com/partnerguide).

For troubleshooting assistance, see:

- *GroupWise 2012 Troubleshooting 1: Error Messages*
- *GroupWise 2012 Troubleshooting 2: Solutions to Common Problems*
- *GroupWise 2012 Troubleshooting 3: Message Flow and Directory Structure*
- Novell Support and Knowledgebase (http://www.novell.com/support)

  To search the GroupWise documentation from the Novell Support Web site, click *Advanced Search*, select *Documentation* in the *Search In* drop-down list, select *GroupWise* in the *Products* drop-down list, type the search string, then click *Search*.
- GroupWise Support Forums (http://forums.novell.com/forumdisplay.php?&f=356)
- GroupWise Support Community (http://www.novell.com/support/products/groupwise)
- GroupWise Cool Solutions (http://www.novell.com/coolsolutions/gwmag/index.html)

Audience

This guide is intended for network administrators who install and administer GroupWise.

Feedback

We want to hear your comments and suggestions about this manual and the other documentation included with this product. Please use the User Comment feature at the bottom of each page of the online documentation.

Additional Documentation

For additional GroupWise documentation, see the following guides at the Novell GroupWise 2012 documentation Web site (http://www.novell.com/documentation/groupwise2012):

- *Installation Guide*
- *Server Migration Guide*
- *Administration Guide*
- *Multi-System Administration Guide*
- *Interoperability Guide*
- *Troubleshooting Guides*
- *GroupWise User Frequently Asked Questions (FAQ)*
- *GroupWise User Guides*
- *GroupWise User Quick Starts*

Novell Cluster Services on Linux

1 Introduction to GroupWise 2012 and Novell Cluster Services on Linux

Before implementing GroupWise 2012 with Novell Cluster Services on Linux, make sure you have a solid understanding of Novell Cluster Services by reviewing the OES Linux Clustering documentation for your version of OES Linux (http://www.novell.com/documentation/oes.html). When you review this information, you discover that clustering employs very specialized terminology. The following brief glossary provides basic definitions of clustering terms and relates them to your GroupWise system:

cluster: A grouping of from two to 32 servers configured using Novell Cluster Services so that data storage locations and applications can transfer from one server to another without interrupting their availability to users.

node: A clustered server; in other words, a single server that is part of a cluster.

shared disk system: The hardware housing the physical disks that are shared among the cluster nodes.

shared partition: A disk partition in a shared disk system that can be accessed from any cluster node that needs the data stored on it. On Linux, Novell Cluster Services supports shared partitions (Linux traditional file system disk partitions), shared NSS volumes (Novell Storage Services volumes), and shared pools (virtual servers).

NOTE: For simplicity, this section uses the term "shared partition" to represent any of these three storage configuration alternatives. For more information, see the *Novell Cluster Services Administration Guide for Linux* for your version of OES Linux (http://www.novell.com/documentation/oes.html).

---

cluster-enabled shared partition: A shared partition for which a Cluster Resource object has been created in Novell eDirectory. The properties of the Cluster Resource object provide load and unload scripts for applications and services installed on the partition, failover/failback/migration policies for the applications and services, and the failover list for the partition.

IMPORTANT: Cluster-enabling is required for GroupWise. For more information, see the *Novell Cluster Services Administration Guide for Linux* for your version of OES Linux (http://www.novell.com/documentation/oes.html).

GroupWise partition: As used in this section, a cluster-enabled shared partition that is used for GroupWise, such as for housing a domain, a post office, or a software distribution directory.

Messenger partition: As used in this section, a cluster-enabled shared partition that is used for Messenger, such as for storing conversation files, log files, temporary files, queue directories, and so on.

cluster resource: A shared partition, secondary IP address, application, service, Web server, and so on, that can function successfully anywhere in the cluster. Cluster resources include the GroupWise agents and the Messenger agents.

failover: The process of moving cluster resources from a failed node to a functional node so that availability to users is uninterrupted. For example, if the node where the POA is running goes down, the POA and its post office fail over to a secondary node so that users can continue to use GroupWise. When you set up cluster resources, you must consider what components need to fail over together in order to continue functioning.

fan-out-failover: The configuration where cluster resources from a single failed node fail over to several different nodes in order to distribute the load from the failed node across multiple nodes. For example, if a node runs a cluster resource consisting of a domain and its MTA, another cluster resource consisting of a post office and its POA, and a third cluster resource for the GWIA, each cluster resource can be configured to fail over separately to different secondary nodes.

failback: The process of returning cluster resources to their preferred node after the situation causing the failover has been resolved. For example, if a POA and its post office fail over to a secondary node, that cluster resource can be configured to fail back to its preferred node when the problem is resolved.

migration: The process of manually moving a cluster resource from its preferred node to a secondary node for the purpose of performing maintenance on the preferred node, temporarily lightening the load on the preferred node, and so on.

2 Planning GroupWise in a Linux Cluster

The majority of this part of the *GroupWise 2012 Interoperability Guide* (Chapter 2, "Planning GroupWise in a Linux Cluster," on page 15 through Chapter 8, "Backing Up a GroupWise System in a Linux Cluster," on page 85) is designed for those who are creating a new GroupWise system, or at least new domains and post offices, in the context of Novell Cluster Services on Linux.

If you already have an existing GroupWise 2012 system on OES Linux and need to configure it to work in a newly installed cluster, see Chapter 10, "Moving an Existing Linux GroupWise 2012 System into a Linux Cluster," on page 89.

When you implement a new GroupWise system or a new domain or post office in a clustering environment, overall GroupWise system design does not need to change substantially. For a review, see "Installing a Basic GroupWise System" in the *GroupWise 2012 Installation Guide*. However, the configuration of individual components of your GroupWise system will be significantly different. This section helps you plan the following GroupWise components in a cluster:

- A new GroupWise system consisting of the primary domain and the initial post office
- A new secondary domain
- A new post office
- The GroupWise agents: Message Transfer Agent (MTA), Post Office Agent (POA), and Document Viewer Agent

During the planning process, component configuration alternatives are explained. For example, you might want the domain and post office together on the same shared partition or on different shared partitions.

The System Clustering Worksheet and the Agent Clustering Worksheet list the information you need as you set up GroupWise in a clustering environment. You should print these worksheets and fill them out as you complete the tasks listed below:

- Section 2.1, "Installing Novell Cluster Services on Linux," on page 16
- Section 2.2, "Planning a Clustered Software Distribution Directory," on page 17
- Section 2.3, "Planning a New Clustered Domain," on page 18
- Section 2.4, "Planning a New Clustered Post Office," on page 19
- Section 2.5, "Planning a New Library for a Clustered Post Office," on page 19
- Section 2.6, "Deciding How to Install and Configure the Linux Agents in a Cluster," on page 20
- Section 2.7, "GroupWise Clustering Worksheets," on page 22

After you have completed the tasks and filled out the "System Clustering Worksheet" on page 22 and the "Agent Clustering Worksheet" on page 23, you are ready to continue with Chapter 3, "Setting Up a Domain and a Post Office in a Linux Cluster," on page 25.

2.1 Installing Novell Cluster Services on Linux

Install Novell Cluster Services on OES Linux by following the instructions provided in the *Novell Cluster Services Administration Guide for Linux* for your version of OES Linux (http://www.novell.com/documentation/oes.html).

The cluster installation process includes:

- Meeting hardware and software requirements
- Setting up a shared disk system
- Creating a new Cluster object to represent the cluster in Novell eDirectory
- Adding nodes to the cluster
- Installing the Novell Cluster Services software on all nodes in the cluster
- Creating shared partitions, shared NSS volumes, or shared pools as needed for your cluster, as described in the *Novell Cluster Services Administration Guide for Linux* for your version of OES Linux (http://www.novell.com/documentation/oes.html).

  NOTE: For simplicity in this section, the term "shared partition" is intended to include any of these shared storage alternatives.

- Cluster-enabling any of these shared storage alternatives, as described in the *Novell Cluster Services Administration Guide for Linux* for your version of OES Linux (http://www.novell.com/documentation/oes.html).

  IMPORTANT: Cluster-enabling is required for GroupWise.

- Mounting the shared partitions where you want to set up GroupWise domains and post offices.

As you install Novell Cluster Services on Linux, record key information about the cluster on the System Clustering Worksheet:

SYSTEM CLUSTERING WORKSHEET

Under Item 1: eDirectory Tree for Cluster, record the name of the eDirectory tree where the new Cluster object has been created.

Under Item 2: Cluster Name, record the name of the Cluster object that you created for your GroupWise system.

Under Item 3: Cluster Context, record the full context of the Cluster object.

Under Item 4: Nodes in Cluster, list the nodes that you have added to the cluster. Include the file system information about each partition, including file system type (`nss`, `reiserfs`, `ext3`, and so on), device name (`sda2`, `hda1`, and so on), and mount point directory (`/media/nss`, `/mnt`, `/mail`, and so on). You need this information when you set up the load and unload scripts for the GroupWise cluster resources.

Under Item 5: Shared Partitions, list the volume names and volume IDs for the shared partitions that are available for use in your GroupWise system. You need this information when you set up the load and unload scripts for the GroupWise cluster resources.

The number of nodes and shared partitions that are available in the cluster strongly influences where you can place GroupWise domains and post offices. You have several alternatives:

- Your whole GroupWise system can run in a single cluster.

- Parts of your GroupWise system can run in one cluster while other parts of it run in one or more other clusters.
- Parts of your GroupWise system can run in a cluster while other parts run outside of the cluster, on non-clustered servers.

If you do not have the system resources to run all of your GroupWise system in a clustering environment, you must decide which parts have the most urgent need for the high availability provided by clustering. Here are some suggestions:

- Post offices and their POAs must be available in order for users to access their GroupWise mailboxes. Therefore, post offices and their POAs are excellent candidates for the high availability provided by clustering.
- Domains and their MTAs are less noticeable to users when they are unavailable (unless users in different post offices happen to be actively engaged in an email discussion when the MTA goes down). On the other hand, domains and their MTAs are critical to GroupWise administrators, although administrators might be more tolerant of a down server than end users are. Critical domains in your system are the primary domain and, if you have one, a hub or routing domain. These domains should be in the cluster, even if other domains are not.
- The GWIA might or might not require high availability in your GroupWise system, depending on the importance of immediate messaging across the Internet and the use of POP3 or IMAP4 clients by GroupWise users.
- The Monitor Agent is a vital partner with the GroupWise High Availability service, described in "Enabling the GroupWise High Availability Service for the Linux GroupWise Agents" in "Installing GroupWise Agents" in the *GroupWise 2012 Installation Guide*. The GroupWise High Availability service automatically restarts agents that go down under circumstances that do not cause the entire server to go down. If you want this protection for your GroupWise agents, you can run the Monitor Agent in your cluster.

There is no right or wrong way to implement GroupWise in a clustering environment. It all depends on the specific needs of your particular GroupWise system and its users.

2.2 Planning a Clustered Software Distribution Directory

During creation of a GroupWise system on Linux, you are prompted to create a software distribution directory. You can create the software distribution directory on each node where you install the GroupWise software or you can create it on a GroupWise partition so that you install it only once but it is still always available.

IMPORTANT: You must create the software distribution directory in a location that is available to all nodes in the cluster if you want to take advantage of the *Configure GroupWise for Clustering* option of the Linux GroupWise Installation program. This option simplifies the process of installing the agent software to multiple nodes in the cluster. It eliminates the need to provide the same agent configuration information multiple times. The installation instructions in this section are based on using the *Configure GroupWise for Clustering* option of the Installation program.

For background information about software distribution directories, see "Software Directory Management" in "System" in the *GroupWise 2012 Administration Guide*.

SYSTEM CLUSTERING WORKSHEET

If you want to have your GroupWise software distribution directory as part of your cluster, under Item 6: GroupWise Partition for Software Distribution Directory, list the GroupWise partition and associated secondary IP address for the software distribution directory. List the full path for the software distribution directory, regardless of whether it is located on a GroupWise partition or on each node in the cluster.

2.3 Planning a New Clustered Domain

The considerations involved in planning a new domain in a clustering environment are essentially the same as for any other environment.

- **Primary Domain:** If you are setting up a new GroupWise system in a clustering environment, you are creating the primary domain as you complete the tasks in this section. To prepare, review "Planning a Basic GroupWise System", then print and fill out the "Basic GroupWise System Worksheet" in "Installing a Basic GroupWise System" in the *GroupWise 2012 Installation Guide*. This covers planning the primary domain and an initial post office in the primary domain.
- **Secondary Domain:** If your GroupWise system already exists, you are creating a new secondary domain. To prepare, review "Planning a New Domain", then print and fill out the "New Domain Summary Sheet" in "Domains" in the *GroupWise 2012 Administration Guide*.

Regardless of the type of domain you are creating, keep in mind the following cluster-specific details as you fill out the worksheet you need:

- When you specify the location for the domain directory (and for a new GroupWise system, the post office directory) on the worksheet, remember that it will be on a GroupWise partition, not on the node where you will be running the GroupWise Installation program.
- Do not concern yourself with the GroupWise agent information on the worksheet. You will plan the agent installation later. If you are filling out the Basic GroupWise System Worksheet, stop with Post Office Settings. If you are filling out the Domain Worksheet, stop with Domain Administrator.

When you have completed the worksheet, transfer the key information from the Basic GroupWise System Worksheet or the Domain Worksheet to the System Clustering Worksheet.

SYSTEM CLUSTERING WORKSHEET

Under Item 9: Domain Name, transfer the domain name and database directory to the System Clustering Worksheet.

Under Item 7: GroupWise Partition for Domain, transfer the domain location to the System Clustering Worksheet. Also specify the secondary IP address of the shared partition where you plan to create the domain.

IMPORTANT: Do not create the new domain until you are instructed to do so in Chapter 3, "Setting Up a Domain and a Post Office in a Linux Cluster," on page 25.

2.4 Planning a New Clustered Post Office

The considerations involved in planning a new post office in a clustering environment are essentially the same as for any other environment. The initial post office in a new GroupWise system on Linux is planned on the Basic GroupWise System Worksheet. To plan additional new post offices, review "Planning a New Post Office ", then print and fill out the "New Post Office Summary Sheet" in "Post Offices" in the *GroupWise 2012 Administration Guide*. When you specify the location for the post office directory, remember that it will be on a GroupWise partition, not on the node where you will be running the GroupWise Installation program.

When you have completed the worksheet, transfer key information from the Basic GroupWise System Worksheet or the Post Office Worksheet to the System Clustering Worksheet.

SYSTEM CLUSTERING WORKSHEET

Under Item 10: Post Office Name, transfer the post office name and database location to the System Clustering Worksheet. Also specify the secondary IP address of the shared partition where you plan to create the domain.

If you will create the post office on a different GroupWise partition from where the domain is located, under Item 8: Shared Partition for Post Office, transfer the post office location to the System Clustering Worksheet. Also specify the secondary IP address of the shared partition where you plan to create the post office.

IMPORTANT: Do not create the new post office until you are instructed to do so in Chapter 3, "Setting Up a Domain and a Post Office in a Linux Cluster," on page 25.

2.5 Planning a New Library for a Clustered Post Office

The considerations involved in planning a new library in a clustering environment are essentially the same as for any other environment. However, in a Linux cluster, you should not plan to locate a document storage area on a remote storage area. If you choose to place it outside the post office directory structure, it should still be located on the same server with the post office.

You can plan a library for a clustered post office by following the standard instructions provided in "Creating and Managing Libraries" in the *GroupWise 2012 Administration Guide* and filling out the "Basic Library Worksheet" or the "Full-Service Library Worksheet". Then provide the library information on the System Clustering Worksheet.

SYSTEM CLUSTERING WORKSHEET

Under Item 11: Document Storage Area Location, mark where you want to create the library's document storage area.

IMPORTANT: Do not create the new library until you are instructed to do so in Chapter 3, "Setting Up a Domain and a Post Office in a Linux Cluster," on page 25.

2.6 Deciding How to Install and Configure the Linux Agents in a Cluster

There are several cluster-specific issues to consider as you plan to install the Linux MTA, POA, and DVA in your clustered GroupWise system:

- Section 2.6.1, "Recording Secondary IP Addresses for the Linux Agents," on page 20
- Section 2.6.2, "Determining Appropriate Failover Lists for the Linux Agents," on page 21
- Section 2.6.3, "Determining Cluster Resource Information for the Linux Agents," on page 21
- Section 2.6.4, "Planning the Linux Agent Installation," on page 21

2.6.1 Recording Secondary IP Addresses for the Linux Agents

By default, the GroupWise agents listen on all IP addresses, both primary and secondary, that are bound to the server. This means that any time there is a possibility of two of the same type of agent loading on the same node, it is important that each agent use the appropriate secondary IP address of the GroupWise partition. The secondary IP address moves with each agent when it fails over, so that, in the case of the POA, GroupWise clients do not lose their connections to the POA. When you use the *Configure GroupWise for Clustering* option, the GroupWise Installation program sets the --ip switch in each agent startup file to its unique secondary IP address.

If you are planning to set up a GroupWise name server to help GroupWise clients locate their post offices, make sure that the default POA port number of 1677 is used somewhere in the cluster. For more information, see "Simplifying Client/Server Access with a GroupWise Name Server" in "Post Office Agent" in the *GroupWise 2012 Administration Guide*.

If you want to install the DVA along with the POA, it can participate in the cluster by failing over along with the POA and its post office. However, clustering is not necessary in order to provide high availability for the DVA. As an alternative to clustering, you can install the DVA on multiple servers outside the cluster, as described in "Scaling Your DVA Installation" in "Document Viewer Agent" in the *GroupWise 2012 Administration Guide*.

AGENT CLUSTERING WORKSHEET

Under Item 3: MTA Network Information, transfer the domain secondary IP address from the System Clustering Worksheet to the Agent Clustering Worksheet.

Under Item 6: POA Network Information, transfer the post office secondary IP address from the System Clustering Worksheet to the Agent Clustering Worksheet.

Under Item 7: DVA Network Information, transfer the post office secondary IP address from the System Clustering Worksheet to the Agent Clustering Worksheet.

2.6.2 Determining Appropriate Failover Lists for the Linux Agents

By default, a GroupWise partition is configured to have all nodes in the cluster in its failover list, organized in ascending alphanumeric order. Only one node at a time can have a particular GroupWise partition mounted and active. If a GroupWise partition's preferred node fails, the partition fails over to the next node in the failover list. You should customize the failover list for each GroupWise partition based on the fan-out-failover principle.

When a node fails, its partitions should not all fail over together to the same secondary node. Instead, the partitions should be distributed across multiple nodes in the cluster. This prevents any one node from shouldering the entire processing load typically carried by another node. In addition, some partitions should never have the potential of being mounted on the same node during a failover situation. For example, a post office and POA that service a large number of very active GroupWise client users should never fail over to a node where another very large post office and heavily loaded POA reside. If they did, users on both post offices would notice a decrease in responsiveness of the GroupWise client.

AGENT CLUSTERING WORKSHEET

Under Item 2: Domain Failover List, list the nodes that you want to have in the domain partition failover list. The MTA might need to run on any node that the domain partition fails over to. Therefore, you will install the agent software on all of the nodes in the domain failover list.

If you are planning the post office on a different GroupWise partition from where the domain is located, under Item 5: Post Office Failover List, list the nodes that you want to have in the post office partition failover list. The POA (and the DVA if you installed along with the POA) might need to run on any node that the post office partition fails over to. Therefore, you will install the agent software on all of the nodes in the post office failover list.

2.6.3 Determining Cluster Resource Information for the Linux Agents

A cluster resource is a shared partition, secondary IP address, application, service, Web server, and so on, that can function successfully anywhere in the cluster. Cluster resources include the GroupWise agents and the Messenger agents. When you use the *Configure GroupWise for Clustering* option, the GroupWise Installation program needs to know the mount point for the GroupWise partition where it will create the domain and post office. For example, you might create a `/mnt/gwsystem` mount point, or you might create `/mnt/dom1` and `/mnt/po1` mount points. The Installation program also needs to know the secondary IP address of the GroupWise partition. In the Installation program, you provide these pieces of information in the *Cluster Resource Mount Point* and *IP Address of the Cluster Resource* fields.

AGENT CLUSTERING WORKSHEET

Under Item 8: Cluster Resource Information, list the mount point and secondary IP address for the GroupWise partition where the domain and post office will be located.

2.6.4 Planning the Linux Agent Installation

Aside from the cluster-specific issues discussed in the preceding sections, the considerations involved in planning to install the GroupWise Linux agents are the same in a clustering environment as for any other environment. Review "Planning the GroupWise Agents", then print and fill out the "GroupWise Agent Installation Worksheet" in "Installing GroupWise Agents" in the *GroupWise 2012 Installation Guide* for each location where you will install the Linux MTA, POA, and/or DVA.

IMPORTANT: Do not install the Linux agent software until you are instructed to do so in Chapter 3, "Setting Up a Domain and a Post Office in a Linux Cluster," on page 25.

2.7 GroupWise Clustering Worksheets

2.7.1 System Clustering Worksheet

Item	Explanation
1) eDirectory Tree for Cluster:	Record the eDirectory tree where you created the new Novell Cluster object when you installed Novell Cluster Services for Linux. For more information, see Section 2.1, "Installing Novell Cluster Services on Linux," on page 16
2) Cluster Name: Master IP Address:	Record the name of the new Cluster object that you created for your GroupWise system. Also record the virtual IP address of the cluster that will remain constant regardless of which node is currently active. For more information, see Section 2.1, "Installing Novell Cluster Services on Linux," on page 16.
3) Cluster Context:	Record the full context where you created the new Cluster object. For more information, see Section 2.1, "Installing Novell Cluster Services on Linux," on page 16.
4) Nodes in Cluster: ◆ File system type ◆ Device name ◆ Mount point directory	List the nodes that are part of the cluster that you set up for your GroupWise system. Also list the file system type (`reiserfs`, `ext3`, and so on), device name (`sda2`, `hda1`, and so on), and mount point directory (`/mnt`, `/mail`, and so on) for each. You need this information as you create load and unload scripts for GroupWise agents. For more information, see Section 2.1, "Installing Novell Cluster Services on Linux," on page 16.
5) Shared Partitions in Cluster: ◆ Volume name ◆ Volume ID	List the volume name and volume ID for each shared partition that is available for use in your GroupWise system. For more information, see Section 2.1, "Installing Novell Cluster Services on Linux," on page 16.

Item	Explanation
6) GroupWise Partition for Software Distribution Directory: Secondary IP Address: Directory:	If desired, specify the name of the shared partition where the GroupWise software distribution directory will reside and the full path to its location. For more information, see Section 2.2, “Planning a Clustered Software Distribution Directory,” on page 17.
7) GroupWise Partition for Domain: Secondary IP Address: Post Office on Same Partition as Domain? ◆ Yes ◆ No	Specify the name of the shared partition where the GroupWise domain will reside and its secondary IP address. For more information, see Section 2.3, “Planning a New Clustered Domain,” on page 18.
8) GroupWise Partition for Post Office: Secondary IP Address:	Specify the name of the shared partition where the GroupWise post office will reside and its secondary IP address. For more information, see Section 2.4, “Planning a New Clustered Post Office,” on page 19.
9) Domain Name: Domain Directory:	Specify a unique name for the domain. Specify the directory on the GroupWise partition where you want to create the new domain. For more information, see Section 2.3, “Planning a New Clustered Domain,” on page 18.
10) Post Office Name: Post Office Directory:	Specify a unique name for the post office. Specify the directory on the GroupWise partition where you want to create the post office. For more information, see Section 2.4, “Planning a New Clustered Post Office,” on page 19.
11) Document Storage Area Location: ◆ At the post office ◆ Outside the post office ◆ Separate post office	If you need a library for a clustered post office, mark where you want to create its document storage area and provide a directory if necessary. For more information, see Section 2.5, “Planning a New Library for a Clustered Post Office,” on page 19.

2.7.2 Agent Clustering Worksheet

Item	Explanation
1) Domain Name: Domain Location:	Transfer this information from the System Clustering Worksheet (item 9).

Item	Explanation
2) Domain Failover List:	List other nodes in the cluster where the GroupWise domain and its MTA can fail over. For more information, see Section 2.6.2, “Determining Appropriate Failover Lists for the Linux Agents,” on page 21.
3) MTA Network Information: ◆ MTA IP address ◆ MTA message transfer port ◆ MTA HTTP port	Record the MTA network address information for the server where the MTA will run. The MTA IP address is the same as the domain secondary IP address in the cluster. See Section 2.6.1, “Recording Secondary IP Addresses for the Linux Agents,” on page 20.
4) Post Office Name: Post Office Location:	Transfer this information from the System Clustering Worksheet (item 10).
5) Post Office Failover List:	List other nodes in the cluster where the GroupWise post office and its POA can fail over. For more information, see Section 2.6.2, “Determining Appropriate Failover Lists for the Linux Agents,” on page 21.
6) POA Network Information: ◆ POA IP address ◆ POA client/server port ◆ POA message transfer port ◆ POA HTTP port	Record the POA network address information for the server where the POA will run. The POA IP address is the same as the post office secondary IP address in the cluster. See Section 2.6.1, “Recording Secondary IP Addresses for the Linux Agents,” on page 20.
7) DVA Network Information: ◆ DVA IP address ◆ DVA HTTP port	If you are installing the DVA along with the POA so that it can participate in the cluster, record the DVA network address information for the server where the DVA will run. When you install the DVA along with the POA, the DVA IP address is the same as the post office secondary IP address in the cluster. See Section 2.6.1, “Recording Secondary IP Addresses for the Linux Agents,” on page 20.
8) Cluster Resource Information ◆ Path to the cluster resource mount point ◆ IP address of the cluster resource	List the cluster resource information for the GroupWise partition where the domain and post office serviced by the agents are located. For more information, see Section 2.6.3, “Determining Cluster Resource Information for the Linux Agents,” on page 21.

3 Setting Up a Domain and a Post Office in a Linux Cluster

You should have already reviewed Chapter 2, "Planning GroupWise in a Linux Cluster," on page 15 and filled out the "System Clustering Worksheet" on page 22 and the "Agent Clustering Worksheet" on page 23. You are now ready to complete the following tasks to set up GroupWise in a clustering environment on Linux:

3.1 Setting Up a New GroupWise System in a Linux Cluster

The Linux GroupWise Installation program walks you through setting up the primary domain and an initial post office in the primary domain. You might be creating your primary domain and initial post office on the same GroupWise partition or on two different partitions. After you have created the primary domain and initial post office and then installed the GroupWise agents on multiple nodes in the cluster, you can create additional secondary domains and post offices as needed.

To set up the primary domain and initial post office for a new GroupWise system in a clustering environment:

1. Start with the first node on the domain failover list (Agent Clustering Worksheet item 2).
2. Make sure that ConsoleOne is installed on the node.

   You must use the version of ConsoleOne that is included with GroupWise 2012. For installation instructions, see "Installing Linux ConsoleOne" in "System" in the *GroupWise 2012 Administration Guide*.
3. (Conditional) If necessary, mount the GroupWise partition where you want to create the GroupWise software distribution directory (System Clustering Worksheet item 6).
4. Mount the GroupWise partition for the domain (System Clustering Worksheet item 7) and, if needed, the GroupWise partition for the post office (System Clustering Worksheet item 8), where the primary domain and the initial post office for your new GroupWise system will be created.

5 Run the Linux GroupWise Installation program, as described in "Starting the Linux GroupWise Installation Program" in "Installing a Basic GroupWise System" in the *GroupWise 2012 Installation Guide*.

IMPORTANT: Do not select the *Configure GroupWise for Clustering* option at this time.

6 When you set up the software distribution directory, install all the agent software.

Although this is not required when creating your initial domain and post office, it makes installation of the other GroupWise agents easier after you have created the initial domain and post office.

7 From the Installation program, run ConsoleOne to set up your initial GroupWise system, as described in "Using ConsoleOne to Create Your Basic GroupWise System" in "Installing a Basic GroupWise System" in the *GroupWise 2012 Installation Guide*.

8 When providing the MTA and POA network address information, use the Agent Clustering Worksheet that you filled out in Section 2.6, "Deciding How to Install and Configure the Linux Agents in a Cluster," on page 20. The information you provide is used to configure the MTA and POA objects in the domain and post office even though you have not yet installed the agent software on any nodes in the cluster.

Do not create users in the post office at this time.

9 When you have finished creating the primary domain and the initial post office, click *Finish* to exit the GroupWise Installation program without installing the agent software.

10 Skip to "Installing and Configuring the Agents in a Linux Cluster" on page 28.

3.2 Creating a New Secondary Domain in a Linux Cluster

After you have set up the primary domain and initial post office, as described in Section 3.1, "Setting Up a New GroupWise System in a Linux Cluster," on page 25, you can create additional secondary domains in your GroupWise system as needed.

To create a new secondary domain in a clustering environment:

1 Mount the GroupWise partition where the new secondary domain will be created.

2 In ConsoleOne, connect to the primary domain in your GroupWise system, as described in "Connecting to a Domain" in "Domains" in the *GroupWise 2012 Administration Guide*.

3 Create the new domain, following the steps provided in "Creating the New Domain" in "Domains" in the *GroupWise 2012 Administration Guide*.

Use the Domain Worksheet you filled out in Section 2.3, "Planning a New Clustered Domain," on page 18 to fill in the fields in the Create GroupWise Domain dialog box.

4 In the *Link to Domain* field, link the new domain to the primary domain of your GroupWise system.

5 Use the Link Configuration tool to change the links from the new domain to all other domains in the cluster to direct TCP/IP links, following the steps provided in "Changing the Link Protocol between Domains to TCP/IP" in "Message Transfer Agent" in the *GroupWise 2012 Administration Guide*.

Although a complete mesh link configuration is the most efficient, it might not be feasible in all situations. Set up as many direct TCP/IP links as possible for best MTA performance in the cluster.

6 Make sure you are still connected to the primary domain.

7 Rebuild the domain database for the new domain, following the steps provided in "Rebuilding Domain or Post Office Databases" in "Databases" in the *GroupWise 2012 Administration Guide*.

The database rebuild is necessary in order to transfer the MTA configuration information and the domain link information into the secondary domain database, because the MTA for the new domain is not yet running.

8 Skip to Installing and Configuring the Agents in a Linux Cluster to install the MTA software for the new domain.

3.3 Creating a New Post Office in a Linux Cluster

You can create a new post office on the same GroupWise partition where its domain resides or on a separate GroupWise partition. If the post office and its domain are on the same GroupWise partition, they fail over together. If they are on separate GroupWise partitions, they fail over separately.

To create a new post office in a clustering environment on Linux:

1 Mount the GroupWise partition where the domain that will own the new post office is located.

2 (Conditional) If necessary, mount the GroupWise partition for the new post office

3 In ConsoleOne, connect to the GroupWise domain where you want to create the new post office, as described in "Connecting to a Domain" in "Domains" in the *GroupWise 2012 Administration Guide*.

4 Create the new post office, following the steps provided in "Creating the New Post Office" in "Post Offices" in the *GroupWise 2012 Administration Guide*.

Use the Post Office Worksheet you filled out in Section 2.4, "Planning a New Clustered Post Office," on page 19 to fill in the fields in the Create GroupWise Post Office dialog box.

5 Refer to the Agent Clustering Worksheet that you filled out during Section 2.6, "Deciding How to Install and Configure the Linux Agents in a Cluster," on page 20 for the secondary IP address and port numbers that you need to specify in order to configure the link.

6 (Optional) If you want to create a library at the post office, select *Create Library*.

This option creates the document storage area for the library under the post office directory and is not recommended for large libraries.

7 (Optional) Right-click the new POA object, then click *Properties*.

On the POA Agent Settings and Scheduled Events pages, you might want to specify unique times for the following POA activities to prevent multiple POAs from performing the same activities on the same node at the same time during a failover situation:

- Start User Upkeep
- Generate Address Book for Remote
- Start QuickFinder Indexing
- Mailbox/Library Maintenance Event

For more information about these repetitive POA activities, see "Performing Nightly User Upkeep", "Regulating Indexing", and "Scheduling Database Maintenance" in "Post Office Agent" in the *GroupWise 2012 Administration Guide*.

8 Make sure you are still connected to the domain that owns the new post office.

9 Rebuild the post office database for the new post office, following the steps provided in "Rebuilding Domain or Post Office Databases" in "Databases" in the *GroupWise 2012 Administration Guide*. Be sure to browse to the database location under (System Clustering Worksheet item 8) through the GroupWise partition.

The database rebuild is necessary in order to transfer the POA configuration information and the post office link information into the post office database, because the POA for the new post office is not yet running.

10 (Optional) If you want to create a library with its document storage area outside the post office directory, follow the steps in "Setting Up a Basic Library" or "Setting Up a Full-Service Library" in "Libraries and Documents" in the *GroupWise 2012 Administration Guide*, after you have completely finished setting up the clustered post office.

11 Continue with Installing and Configuring the Agents in a Linux Cluster to install the POA software for the new post office.

3.4 Installing and Configuring the Agents in a Linux Cluster

By following the instructions in Section 3.1, "Setting Up a New GroupWise System in a Linux Cluster," on page 25, you created the initial domain and post office in your GroupWise system. You are now ready to install and configure the MTA, POA, and DVA, and set up the agent software for use in your cluster.

- Section 3.4.1, "Installing and Setting Up the Linux Agents in Your Cluster," on page 28
- Section 3.4.2, "Changing Linux Agent Paths to Locations on GroupWise Partitions," on page 33
- Section 3.4.3, "Configuring GroupWise Cluster Resources to Load and Unload the Linux Agents," on page 34
- Section 3.4.4, "Setting Up New Instances of the Linux Agents without Installing the Agent Software," on page 39

IMPORTANT: If you have added a new secondary domain or a new post office to an existing GroupWise partition, the agent software has already been installed and you simply need to create a new startup file specific to the new domain or post office. In these circumstances, follow the instructions in Section 3.4.4, "Setting Up New Instances of the Linux Agents without Installing the Agent Software," on page 39 instead of completing the first three tasks above.

3.4.1 Installing and Setting Up the Linux Agents in Your Cluster

The agents must be installed on each node in domain failover list (Agent Clustering Worksheet item 2) and the post office failover list (Agent Clustering Worksheet item 5).

- "Running the Linux GroupWise Installation Program on the Preferred Node" on page 28
- "Running the Linux GroupWise Installation Program on Subsequent Nodes" on page 30
- "Testing Your Linux Agent Installation on Each Node" on page 32

After you have installed and tested the agents on each node in the cluster, continue with Section 3.4.2, "Changing Linux Agent Paths to Locations on GroupWise Partitions," on page 33.

Running the Linux GroupWise Installation Program on the Preferred Node

1 Mount the GroupWise partition for the domain (System Clustering Worksheet item 7) or the post office (System Clustering Worksheet item 8).

2 From the software distribution directory you created in Step 6 in Section 3.1, "Setting Up a New GroupWise System in a Linux Cluster," on page 25, start the Linux GroupWise Installation program.

IMPORTANT: This time, you should select the *Configure GroupWise for Clustering* option.

3 Install the agent software, following the steps provided in "Installing the Linux GroupWise Agents" in "Installing GroupWise Agents" in the *GroupWise 2012 Installation Guide*.

4 Configure the Linux agents according to the "Agent Clustering Worksheet" on page 23 that you filled out in Section 2.6, "Deciding How to Install and Configure the Linux Agents in a Cluster," on page 20, paying special attention to the cluster resource information on the Domains / Post Offices page.

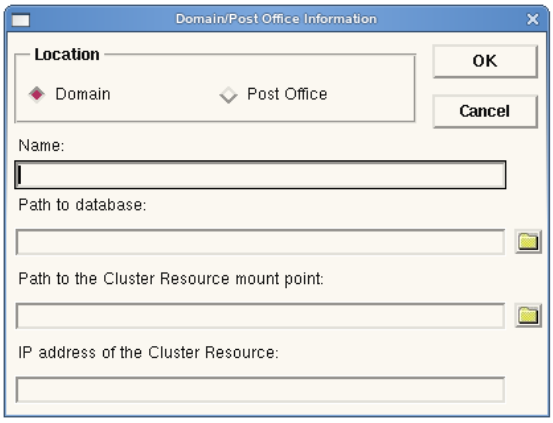

As a result of selecting *Configure GroupWise for Clustering* on the preferred node, the following cluster-specific configuration actions are performed for the MTA and the POA:

- The agent startup files are created in *mount_point*`/groupwise/agents/share` on the shared resource so that the agents use the same startup files regardless of which cluster node the agents are running on. The --home startup switch includes the mount point and the path to the database so that the startup file is valid when mounted to each cluster node.
- The --cluster startup switch is added to the agent startup files to inform the agents that they are running in a cluster.
- The --ip startup switch is set to the secondary IP address of the shared resource where the domain and post office are located. This ensures that the MTA and the POA run with the same IP address regardless of which cluster node the agents are running on.
- The --log startup switch is set to a location on the shared resource (*mount_point*`/groupwise/agents/log`) so that agent logging information is written to the same log file regardless of which cluster node the agents are running on.
- The GroupWise High Availability service is automatically configured on the current cluster node and its configuration file (`gwha.conf`) is created in the `/etc/opt/novell/groupwise` directory.
- A `clusterimport.conf` file is created in the `gwinst` subdirectory of the software distribution directory from which you ran the Linux GroupWise Installation program, so that the cluster-specific information collected when you configured the agents on the preferred node is available when you install the agents on subsequent nodes.
- The agents are not configured to start automatically on system startup. In a cluster, you do not want the agents to start automatically whenever a node restarts.

5 (Conditional) If you enable the DVA for clustering, configure it according to the Agent Clustering Worksheet, paying special attention to the cluster resource information:

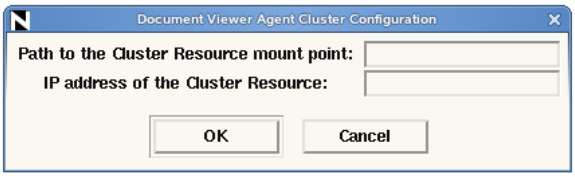

As a result of selecting *Enable the Document Viewer Agent for Clustering* on the preferred node, the following cluster-specific configuration actions are performed for the DVA:

- The DVA startup file is created in *mount_point*`/groupwise/agents/share` on the shared resource so that the DVA uses the same startup file regardless of which cluster node it is running on. The DVA --home startup switch includes the mount point so that the startup file is valid when mounted to each cluster node.
- The DVA --ip startup switch is set to the secondary IP address of the shared resource. This ensures that the DVA runs with the same IP address regardless of which cluster node it is running on.
- The DVA --log startup switch is set to a location on the shared resource (*mount_point*`/groupwise/agents/log`) so that DVA logging information is written to the same log file regardless of which cluster node it is running on.
- The GroupWise High Availability service configuration file (`gwha.conf`) is updated with the configuration information for the DVA.

6 Configure the agents to run as a non-`root` user, as described in the applicable section of the *GroupWise 2012 Installation Guide*:

- "Running the Linux GroupWise Agents as a Non-root User"
- "Setting Up Non-root Access on an NSS Volume on Novell Open Enterprise Server Linux"

7 Continue with Running the Linux GroupWise Installation Program on Subsequent Nodes.

Running the Linux GroupWise Installation Program on Subsequent Nodes

1 On the next node in the GroupWise agent failover list, mount the GroupWise partition for the domain (System Clustering Worksheet item 7) or the post office (System Clustering Worksheet item 8).

2 From the software distribution directory you created in Step 6 in Section 3.1, "Setting Up a New GroupWise System in a Linux Cluster," on page 25, start the Linux GroupWise Installation program.

IMPORTANT: You should select the *Configure GroupWise for Clustering* option again.

Because of the existence of the `clusterimport.conf` file in the `gwinst` subdirectory of the software distribution directory, a new installation option, *Import Clustering Data*, is now available on the main GroupWise Installation program page.

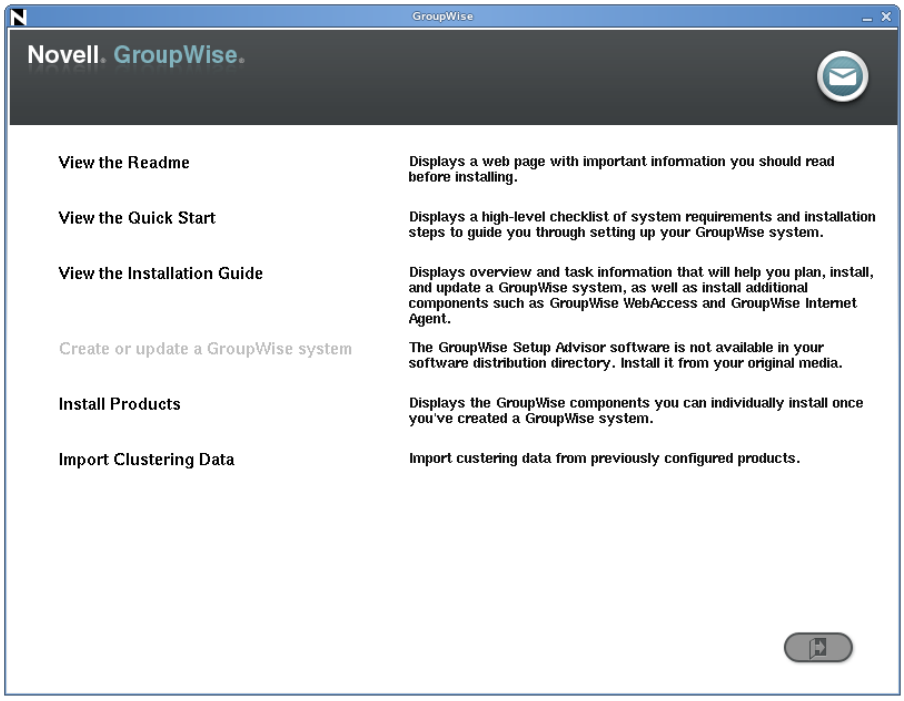

3 Install the agent software on the cluster node as usual, but do not use the *Configure* option.

4 On the main page of the Installation program, click *Import Clustering Data*, then click *Next*.

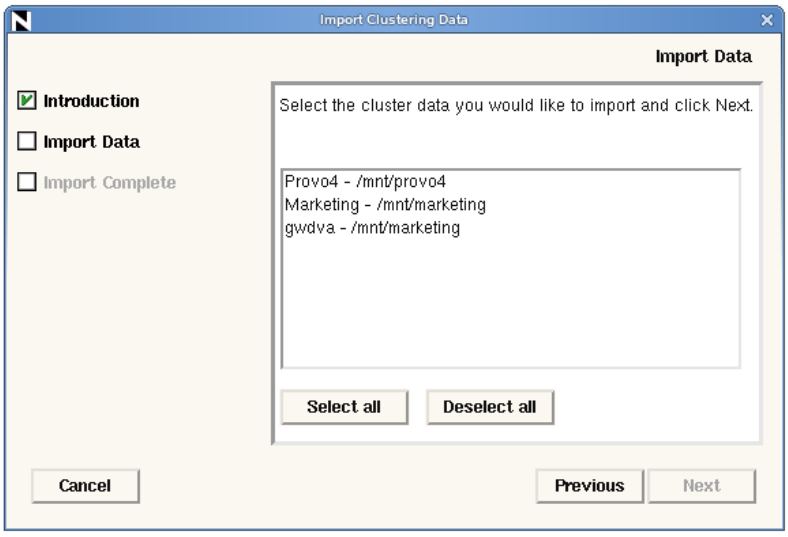

All GroupWise agents that you have installed from the software distribution directory are listed, based on the information stored in the `clusterimport.conf` file.

5 Select the GroupWise agents that you want to configure on the current cluster node, then click *OK*.

The *Import Clustering Data* option performs the following configuration actions for each subsequent node where you run it:

- The `grpwise` script is created in the `/etc/init.d` directory on the current cluster node. It is configured specifically for the agents you just selected.
- The GroupWise High Availability service is automatically configured on the current cluster node and its configuration file (`gwha.conf`) is created in the `/etc/opt/novell/groupwise` directory. It is configured specifically for the agents you just selected.

Because the agent startup files and log files are stored on the shared resource, they do not need to be customized for each cluster node.

6 Configure the agents to run as a non-`root` user, as described in the applicable section of the *GroupWise 2012 Installation Guide*:

- "Running the Linux GroupWise Agents as a Non-root User"
- "Setting Up Non-root Access on an NSS Volume on Novell Open Enterprise Server Linux"

7 Repeat Step 1 through Step 6 for each cluster node in the GroupWise agent failover list.

After you install and configure the agents on each node in each agent's failover list, the cluster node is ready for the GroupWise agent to fail over to it.

8 Continue with Testing Your Linux Agent Installation on Each Node.

Testing Your Linux Agent Installation on Each Node

1 Test the MTA and the POA by starting them with a user interface, as described in "Starting the Linux Agents with a User Interface" in "Installing GroupWise Agents" in the *GroupWise 2012 Administration Guide*.

```
/opt/novell/groupwise/agents/bin/gwmta --show @domain.mta &
/opt/novell/groupwise/agents/bin/gwpoa --show @post_office.poa &
```

2 Stop the agents by clicking *File* > *Exit*.

3 (Conditional) If you installed the DVA, test it by starting and stopping it:

```
rcgrpwise start gwdva
rcgrpwise stop gwdva
```

4 After you can see that the agents start successfully, test them by starting them as daemons, as described in "Starting the Linux Agents as Daemons" in "Installing GroupWise Agents" in the *GroupWise 2012 Administration Guide*.

```
rcgrpwise start
rcgrpwise status
```

5 Stop the agents.

```
rcgrpwise stop
rcgrpwise status
```

6 Return to "Running the Linux GroupWise Installation Program on the Preferred Node" on page 28 for each node in the domain or post office failover list.

When you have installed the agents on all of the nodes in the domain and post office failover lists, continue with Changing Linux Agent Paths to Locations on GroupWise Partitions.

3.4.2 Changing Linux Agent Paths to Locations on GroupWise Partitions

The default locations for some agent files are on the cluster nodes along with the agent software, rather than being located with the domain and post office on one or more GroupWise partitions. To avoid having multiple copies of these files in multiple locations, you should set the locations in ConsoleOne.

Setting the Linux MTA Message Log File Path

If you plan to enable message logging, as described in "Enabling MTA Message Logging" in "Message Transfer Agent" in the *GroupWise 2012 Administration Guide*:

1. On the GroupWise partition where the domain is located, create the directory where you want to store MTA message log files.
2. In ConsoleOne, browse to and select the Domain object.
3. Right-click the MTA object, then click *Properties*.
4. Click *GroupWise > Message Log Settings*.
5. In the *Message Log File Path* field, browse to and select the directory you created in Step 1, then click *OK*.

Setting the Linux MTA Certificate and Key File Path

If you plan to enable SSL, as described in "Securing the Domain with SSL Connections to the MTA" in "Message Transfer Agent" in the *GroupWise 2012 Administration Guide*:

1. On the GroupWise partition where the domain is located, create the directory where you want to store the certificate and key file required for SSL.
2. Copy the certificate file and key file into the new directory.

   If you need assistance obtaining these files, see "Server Certificates and SSL Encryption" in "Security Administration" in the *GroupWise 2012 Administration Guide*.
3. In ConsoleOne, browse to and select the Domain object.
4. Right-click the MTA object, then click *Properties*.
5. Click *GroupWise > SSL Settings*.
6. In the *Certificate File* field, browse to and select the certificate file.
7. In the *SSL Key File* field, browse to and select the key file.
8. Click *OK*.

Setting the Linux POA Certificate and Key File Path

If you plan to enable SSL, as described in "Securing the Post Office with SSL Connections to the POA" in "Post Office Agent" in the *GroupWise 2012 Administration Guide*:

1. On the GroupWise partition where the post office is located, create the directory where you want to store the certificate and key file required for SSL.
2. Copy the certificate file and key file into the new directory.

If you need assistance obtaining these files, see "Server Certificates and SSL Encryption" in "Security Administration" in the *GroupWise 2012 Administration Guide*.

3 In ConsoleOne, browse to and select the Post Office object.

4 Right-click the POA object, then click *Properties*.

5 Click *GroupWise* > *SSL Settings*.

6 In the *Certificate File* field, browse to and select the certificate file.

7 In the *SSL Key File* field, browse to and select the key file.

8 Click *OK*.

3.4.3 Configuring GroupWise Cluster Resources to Load and Unload the Linux Agents

The properties of the Cluster Resource object associated with the GroupWise partition define how partitions function within the cluster, how agents are loaded and unloaded, and how failover and failback situations are handled. At this point, you might have one cluster resource for a GroupWise partition with a domain and post office on it, or you might have two cluster resources for two GroupWise partitions, one for the domain and one for the post office. Complete the following tasks for each cluster resource:

- "Modifying the Cluster Resource Load Script for the Linux Agents" on page 34
- "Modifying the Cluster Resource Unload Script for the Linux Agents" on page 37
- "Setting the Failover List and Policies for the Linux Agents" on page 39

Modifying the Cluster Resource Load Script for the Linux Agents

The cluster resource load script executes whenever the GroupWise partition comes online. On OES Linux, all cluster management activities are performed in Novell iManager.

To set up the load script in iManager:

1 Expand *Clusters*, then click *Cluster Options*.

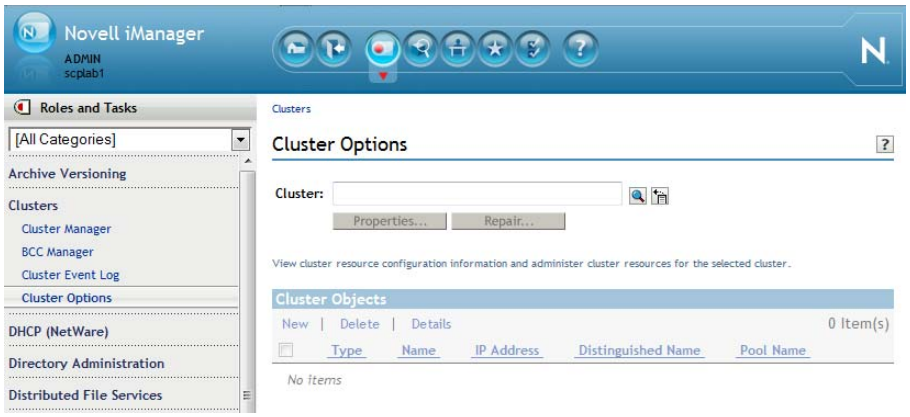

2 In the *Cluster* field, browse to the Cluster object where the GroupWise cluster resource is located.

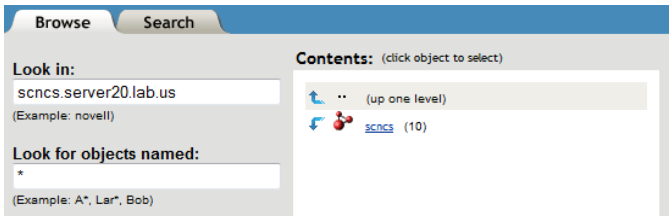

3 Click the Cluster object to display the cluster resources that belong to the cluster.

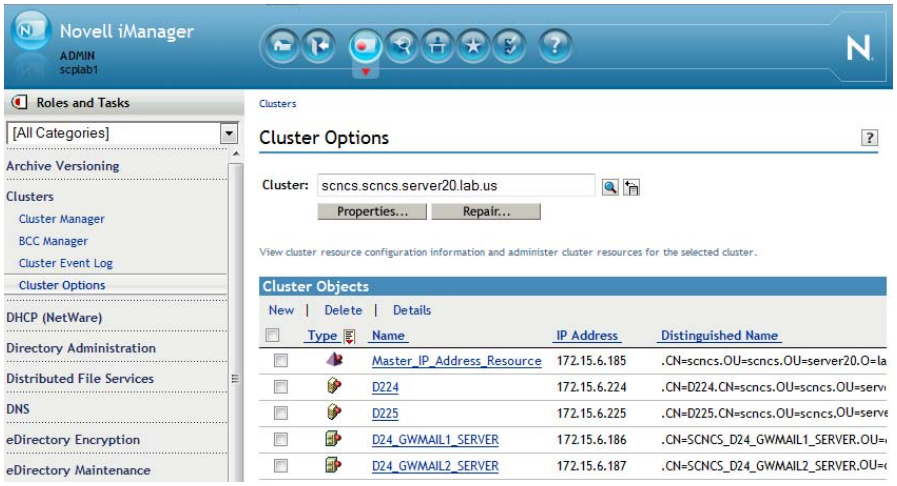

4 Select the GroupWise Cluster Resource object that you created when you set up the GroupWise partition, then click *Details*.

5 Click *Scripts* > *Load Script*.

6 (Conditional) If this is an NSS volume or a shared pool, use a load script similar to the following example, depending on the configuration of your cluster and nodes:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

# mount filesystem
exit_on_error ncpcon mount /opt="noatime,nodiratime" volume_name=volume_ID

# add IP address
exit_on_error add_secondary_ipaddress gw_partition_ip_address

# start service
exit_on_error /etc/init.d/grpwise start domain
exit_on_error /etc/init.d/grpwise start post_office.domain
exit_on_error /etc/init.d/grpwise start gwdva
```

6a In the `mount filesystem` section, specify the volume name and volume ID of the GroupWise partition that you are clustering (System Clustering Worksheet item 5).

6b In the `add IP address` section, specify the secondary IP address of the GroupWise partition (System Clustering Worksheet item 7 or System Clustering Worksheet item 8).

6c In the `start service` section, use the commands to start the specific GroupWise agents that you want to run on this GroupWise partition.

If you created a domain on the partition, you need to start the MTA. If you created a post office on the partition, you need to start the POA. If you created both a domain and a post office, you need to start both the MTA and the POA. If you installed the DVA, you need to start it as well.

7 (Conditional) If this is a traditional Linux volume, use a load script similar to the following example, depending on the configuration of your cluster and nodes:

```
#! /bin/bash
. /opt/novell/ncs/lib/ncsfunc

# define IP address
RESOURCE_IP=gw_partition_ip_address

# define filesystem type
MOUNT_FS=filesystem

# define device (if using EVMS)
exit_on_error evms -f /var/opt/novell/ncs/ContainerActivate -rl
                                                    Share `uname -n`
MOUNT_DEV=/dev/evms/Share/dat

# define mount point
MOUNT_POINT=/mnt/mount_point_directory

# mount filesystem
exit_on_error mount -t $MOUNT_FS $MOUNT_DEV $MOUNT_POINT -o noatime,nodiratime

# add IP address
exit_on_error add_secondary_ipaddress $RESOURCE_IP

# start service
exit_on_error /etc/init.d/grpwise start domain
exit_on_error /etc/init.d/grpwise start post_office.domain
exit_on_error /etc/init.d/grpwise start gwdva

exit 0
```

7a In the `define IP address` section, specify the secondary IP address of the GroupWise partition (System Clustering Worksheet item 7 or System Clustering Worksheet item 8)

7b In the `define filesystem type` section, specify the filesystem type that is in use on the nodes in the cluster (System Clustering Worksheet item 5).

7c In the `define mount point` section, specify the mount point directory in use for the nodes in the cluster (System Clustering Worksheet item 5).

7d In the `start service` section, use the commands to start the specific GroupWise agents that you want to run on this GroupWise partition.

If you created a domain on the partition, you need to start the MTA. If you created a post office on the partition, you need to start the POA. If you created both a domain and a post office, you need to start both the MTA and the POA. If you installed the DVA, you need to start it as well.

8 Click *OK* to save the load script.

Modifying the Cluster Resource Unload Script for the Linux Agents

The cluster resource unload script executes whenever the GroupWise partition goes offline. Programs should be unloaded in the reverse order of how they were loaded. This ensures that supporting programs are not unloaded before programs that rely on them in order to function properly.

1. On the iManager Cluster Resource Properties page of the GroupWise cluster resource, click *Scripts* > *Unload Script*.
2. (Conditional) If this is an NSS volume or a shared pool, use an unload script similar to the following example, depending on the configuration of your cluster and nodes:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

# request service stop
ignore_error /etc/init.d/grpwise stop domain
ignore_error /etc/init.d/grpwise stop post_office.domain
ignore_error /etc/init.d/grpwise stop gwdva

# stop service otherwise
sleep 8
ignore_error pkill -fx "'/opt/novell/groupwise/agents/bin/gwmta
          @/media/nss/volume_name/groupwise/agents/share/domain_name.mta'"
ignore_error pkill -fx "'/opt/novell/groupwise/agents/bin/gwpoa
          @/media/nss/volume_name/groupwise/agents/share/
                                                  post_office_name.poa'"
ignore_error pkill -fx "'/opt/novell/groupwise/agents/bin/gwdva
          @/media/nss/volume_name/groupwise/agents/share/gwdva.dva'"

# delete IP address
ignore_error del_secondary_ipaddress gw_partition_ip_address

# unmount filesystem
exit_on_error umount /mnt/mount_point_directory

# return status
exit 0
```

   2a In the `request service stop` section, use the commands to stop the specific GroupWise agents that are running on this GroupWise partition.

   2b In the `stop service otherwise` section, adjust the `sleep` command as needed so that the agents can shut down normally on your system without being inadvertently killed by the `pkill` command that follows.

   2c In the `delete IP address` section, specify the secondary IP address of the GroupWise partition.

   2d In the `unmount filesystem` section, specify the mount point directory in use for the nodes in the cluster.

   2e (Conditional) If you are running the GroupWise High Availability service (gwha), stop it before the script stops the agents, then start it again at the end of the unload script.

   This prevents the GroupWise High Availability service from trying to restart the agents while the script is trying to stop them.

   Add the following section before the `request service stop` section:

```
# Temporarily disable the gwha service under xinetd
ignore_error /sbin/chkconfig -s gwha off
ignore_error kill -HUP `pidof xinetd`
```

Add the following section before the `return status` section:

```
# Restart the gwha service under xinetd
ignore_error /sbin/chkconfig -s gwha on
ignore_error kill -HUP `pidof xinetd`
```

3 (Conditional) If this is a traditional Linux volume, use an unload script similar to the following example, depending on the configuration of your cluster and nodes:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

# request service stop
ignore_error /etc/init.d/grpwise stop domain
ignore_error /etc/init.d/grpwise stop post_office.domain
ignore_error /etc/init.d/grpwise stop gwdva

# stop service otherwise
sleep 8
ignore_error pkill -fx "'/opt/novell/groupwise/agents/bin/gwmta
          @/media/nss/volume_name/groupwise/agents/share/domain_name.mta'"
ignore_error pkill -fx "'/opt/novell/groupwise/agents/bin/gwpoa
          @/media/nss/volume_name/groupwise/agents/share/
                                                  post_office_name.poa'"

# define IP address
RESOURCE_IP=gw_partition_ip_address

# define mount point
MOUNT_POINT=/mnt/mount_point_directory

# delete IP address
ignore_error del_secondary_ipaddress $RESOURCE_IP

# unmount filesystem
exit_on_error umount $MOUNT_POINT

# return status
exit 0
```

3a In the `request service stop` section, use the commands to stop the specific GroupWise agents that are running on this GroupWise partition.

3b In the `stop service otherwise` section, adjust the `sleep` command as needed so that the agents can shut down normally on your system without being inadvertently killed by the `pkill` command that follows.

3c In the `define IP address` section, specify the secondary IP address of the GroupWise partition.

3d In the `define mount point` section, specify the mount point directory in use for the nodes in the cluster.

3e (Conditional) If you are running the GroupWise High Availability service (gwha), stop it before the script stops the agents, then start it again at the end of the unload script.

This prevents the GroupWise High Availability service from trying to restart the agents while the script is trying to stop them.

Add the following section before the `request service stop` section:

```
# Temporarily disable the gwha service under xinetd
ignore_error /sbin/chkconfig -s gwha off
ignore_error kill -HUP `pidof xinetd`
```

Add the following section before the `return status` section:

```
# Restart the gwha service under xinetd
ignore_error /sbin/chkconfig -s gwha on
ignore_error kill -HUP `pidof xinetd`
```

4 Click *OK* to save the unload script.

Setting the Failover List and Policies for the Linux Agents

1 On the iManager Cluster Resource Properties page of the GroupWise cluster resource, click *General*.

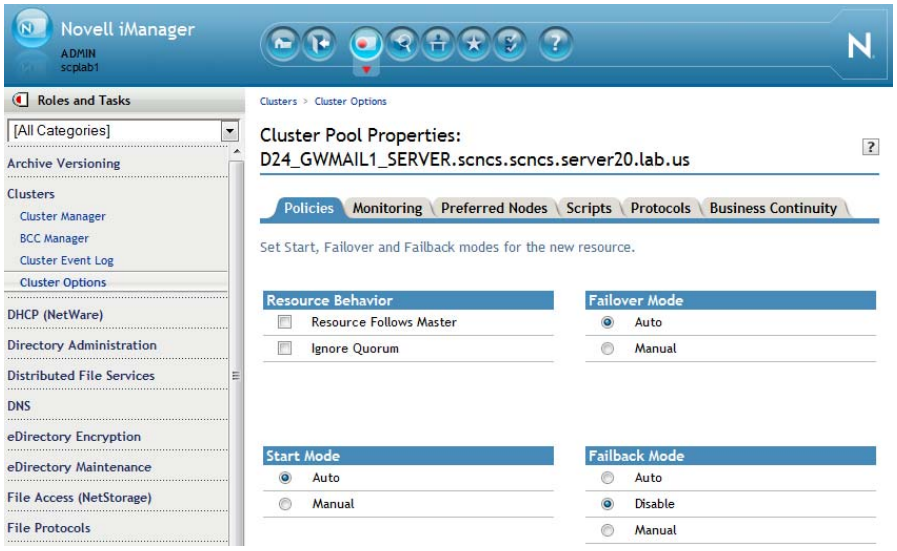

The default policy settings are often appropriate. By default, a cluster resource:

- Fails over automatically if the node it is running on fails
- Starts automatically on the next node in its failover list
- Continues running at its failover location, even after its most preferred node is again available

  If you are considering changing these defaults, see the *Novell Cluster Services Administration Guide for Linux* for your version of OES Linux (http://www.novell.com/documentation/oes.html).

2 Click *Preferred Nodes*, then arrange the nodes in the cluster into the desired failover list for the domain or post office (under Agent Clustering Worksheet items 3 or 6).

3 Click *OK*.

3.4.4 Setting Up New Instances of the Linux Agents without Installing the Agent Software

There are two steps to setting up new instances of the agents without installing the agent software:

- "Creating New Startup Files" on page 40
- "Modifying Existing Load and Unload Scripts" on page 40

Creating New Startup Files

Each MTA startup file is named after the domain it services, with a `.mta` extension. Each POA startup file is named after the post office it services, with a `.poa` extension. When you select the *Configure GroupWise for Clustering* option, the Linux GroupWise Installation program creates agent startup files in *`mount_point`*`/groupwise/agents/share` on the shared resource.

To create a new startup file without installing the agent software:

1. Make a copy of an existing startup file and name it after the domain or post office that will be serviced by the agent.
2. Edit the setting of the --home startup switch to point to the location of the new domain directory or post office directory. Be careful to maintain the syntax of the original line.
3. Scroll down through the new startup file to find other active (not commented out) startup switches, then modify them as needed for the new agent.
4. Save the new startup file.
5. Edit the GroupWise High Availability service configuration file (`/etc/opt/novell/groupwise/gwha.conf`).
6. Make a copy of the section for an existing domain and its MTA or post office and its POA, then modify the information for the new domain or post office and its accompanying agent.
7. Save the `gwha.conf` file.

   For more information about the High Availability service, see "Enabling the GroupWise High Availability Service for the Linux GroupWise Agents" in "Installation" in the *GroupWise 2012 Installation Guide*.
8. Continue with Modifying Existing Load and Unload Scripts.

Modifying Existing Load and Unload Scripts

The agent startup file names are part of the `load` commands found in GroupWise cluster resource load scripts.

If you created the new domain and/or post office on a new GroupWise partition, skip back to Section 3.4.3, "Configuring GroupWise Cluster Resources to Load and Unload the Linux Agents," on page 34 for instructions to create new load and unload scripts.

If you created the new domain and/or post office on an existing GroupWise partition, most of the configuration of the cluster resource is already done. You just need to add new service start and stop commands to the existing scripts. Continue with the steps below:

To modify existing load and unload scripts:

1. In iManager, expand Cluster, then click *Cluster Options*.
2. In the *Cluster* field, browse to and click the Cluster object where the GroupWise cluster resource is located.
3. Select a cluster resource that contains a GroupWise partition, then click *Properties* > *Scripts*.
4. Following the pattern of the existing service start commands, add start commands for the new instances of the agents you are setting up. Use Ctrl+C to copy and Ctrl+V to paste lines in the load script page.
5. Click *Apply* to save the modified load script.
6. Click *Unload Script*.
7. Add corresponding service stop commands for the new instances of the agents.

8 Click *Apply* to save the modified unload script.

You might want to review other properties of the Cluster Resource object, such as the failover list and the failover/failback/migration procedures on the General page, because an additional domain and/or post office now resides on the GroupWise partition.

9 Change other Cluster Resource properties as needed.

10 Click *OK* to save the modified properties.

11 In the Cluster Manager, take the GroupWise partition offline and then bring it online again to test the new startup files and the modified load and unload scripts. If you need assistance with these tasks, see Testing Your Clustered GroupWise System on Linux.

3.5 Testing Your Clustered GroupWise System on Linux

After you have configured the GroupWise cluster resources, you can test the load and unload scripts by bringing the GroupWise cluster resource online and taking it offline again.

1 In iManager, expand *Clusters*, then click *Cluster Manager*.

2 Browse to the Cluster object to display the current cluster state.

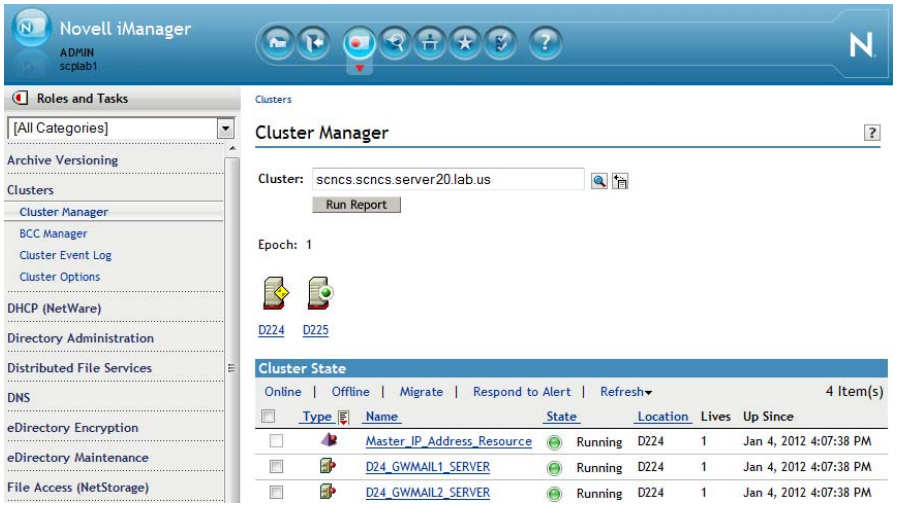

3 (Conditional) If the new GroupWise cluster resource shows *Offline* in the *State* column, select the new GroupWise cluster resource, then click *Online*.

After a moment, the GroupWise cluster resource displays *Running* in the *State* column.

4 At the server where the MTA and/or POA are starting, use the following command to see if they are running:

```
rcgrpwise status domain
rcgrpwise status post_office.domain
```

5 Select the new GroupWise cluster resource, then click *Offline*.

The *State* column for the GroupWise cluster resource displays *Offline*.

6 Use the same command you used in Step 4 to verify if they have stopped.

7 Repeat Step 3 whenever you are ready to bring the new GroupWise cluster resource online permanently.

8 Continue with Managing Your Clustered GroupWise System on Linux.

3.6 Managing Your Clustered GroupWise System on Linux

After you have set up a basic clustered GroupWise system, you should consider some long-term management issues.

3.6.1 Updating GroupWise Objects with Cluster-Specific Descriptions

After setting up your clustered GroupWise system, while the cluster-specific information is fresh in your mind, you should record the cluster-specific information as part of the GroupWise objects in ConsoleOne so that you can easily refer to it later. Be sure to update the information in the GroupWise objects if the configuration of your system changes.

Recording Cluster-Specific Information for a Domain and Its MTA

To permanently record important cluster-specific information for the domain:

1. In ConsoleOne, browse to and right-click the Domain object, then click *Properties*.
2. In the *Description* field of the domain Identification page, provide a cluster-specific description of the domain, including the secondary IP address of its GroupWise partition.
3. Click *OK* to save the domain description.
4. Select the Domain object to display its contents.
5. Right-click the MTA object, then click *Properties*.
6. In the *Description* field of the MTA Identification page, record the secondary IP address of the domain's GroupWise partition.

   This information appears on the MTA server console, no matter which node in the cluster it is currently running on.
7. Click *Apply* to save the description.
8. Click *Network Address*.
9. In the *TCP/IP Address* field, provide the secondary IP address that you provided in the GroupWise Installation program for use with the --ip switch in the MTA startup file.
10. Select *Bind Exclusively to TCP/IP Address*.

    This records this vital information in eDirectory as well as in the MTA startup file.
11. Click *OK* to save the MTA description and secondary IP address.
12. Continue with Recording Cluster-Specific Information for a Post Office and Its POA.

Recording Cluster-Specific Information for a Post Office and Its POA

To permanently record important cluster-specific information for a post office:

1. In ConsoleOne, browse to and right-click the Post Office object, then click *Properties*.

2 In the *Description* field of the post office Identification page:

   2a Provide a cluster-specific description of the post office, including the secondary IP address of its GroupWise partition.

   2b (Conditional) If you installed and clustered the DVA along with the POA, make a note of that configuration choice.

3 Click *OK* to save the post office description.

4 Select the Post Office object to display its contents.

5 Right-click the POA object, then click *Properties*.

6 In the *Description* field of the POA Identification page, record the secondary IP address of the post office's GroupWise partition.

   This information appears on the POA server console, no matter which node in the cluster it is currently running on.

7 Click *Apply* to save the description.

8 Click *Network Address*.

9 In the *TCP/IP Address* field, provide the secondary IP address that you provided in the GroupWise Installation program for use with the --ip switch in the POA startup file.

10 Select *Bind Exclusively to TCP/IP Address*.

   This records this vital information in eDirectory as well as in the POA startup file.

11 Click *OK* to save the POA description and secondary IP address.

12 (Conditional) If applicable, continue with Recording Cluster-Specific Information for a Software Distribution Directory.

   or

   Skip to Section 3.6.2, "Knowing What to Expect in MTA, POA, and DVA Failover Situations," on page 43.

Recording Cluster-Specific Information for a Software Distribution Directory

To permanently record important cluster-specific information about a software distribution directory located on a GroupWise partition:

1 In ConsoleOne, click *Tools* > *System Operations* > *Software Directory Management*.

2 Select the software distribution directory, then click *Edit*.

3 In the description field, record the secondary IP address of the GroupWise partition where the software distribution directory resides.

4 Click *OK*, then click *Close* to save the software distribution directory description.

5 Continue with Knowing What to Expect in MTA, POA, and DVA Failover Situations.

3.6.2 Knowing What to Expect in MTA, POA, and DVA Failover Situations

In a failover situation, the MTA and the POA might need to perform some database repair as they start on the new node. The time required depends on the size of the databases involved.

Typically, the POA returns to full functionality faster than the MTA. This benefits GroupWise client users, who can reconnect to their mailboxes very quickly and probably do not notice if messages to users in other post offices are not delivered immediately. The only time a user needs to restart the

GroupWise client is if he or she was actually in the process of sending a message when the POA went down. Notify can continue running even if the connection to the POA becomes unavailable because it reconnects automatically when the POA is again available.

The MTA typically takes some time reestablishing the links to its post offices, other domains, and gateways, but this situation usually resolves itself in a few minutes without administrator intervention. If it does not, you can manually restart the MTA to speed up the process.

The DVA must reestablish its HTTP connections with one or more POAs and WebAccess Applications. Typically, this occurs quite quickly.

In comparison to failover, migration typically takes longer because the POA and the MTA methodically terminate their threads and close their databases as part of their normal shutdown procedure. However, as a result, no database repair is required when these agents start up again in their new location.

Continue with What's Next.

3.7 What's Next

Now that you have at least one GroupWise domain and post office up and running in a clustering environment, you are ready to proceed with the rest of your GroupWise system setup by:

- Adding users to post offices.

  See "Users" in the *GroupWise 2012 Administration Guide*.
- Setting up the GroupWise Windows client software and helping users to get started using it.

  See "Client" in the *GroupWise 2012 Administration Guide*. Also see the *GroupWise 2012 Windows Client User Guide*.
- Connecting your clustered GroupWise system to the Internet.

  See Chapter 5, "Implementing the Internet Agent in a Linux Cluster," on page 49.
- Monitoring the status of your clustered GroupWise system from your Web browser.

  See Chapter 6, "Implementing GroupWise Monitor in a Linux Cluster," on page 69.
- Backing up your clustered GroupWise system.

  See Chapter 8, "Backing Up a GroupWise System in a Linux Cluster," on page 85.

3.8 Clustering Quick Checklists

3.8.1 GroupWise System Quick Checklist

❒ Plan your new clustered GroupWise system.

See Chapter 2, "Planning GroupWise in a Linux Cluster," on page 15.

❒ Create the primary domain and initial post office in your new clustered GroupWise system.

See Section 3.1, "Setting Up a New GroupWise System in a Linux Cluster," on page 25.

- ❒ Set up the agents for the primary domain and the initial post office.

  See Section 3.4, “Installing and Configuring the Agents in a Linux Cluster,” on page 28.
- ❒ Modify the cluster resource load scripts:

  See “Modifying the Cluster Resource Load Script for the Linux Agents” on page 34.
- ❒ Modify the cluster resource unload scripts:

  See “Modifying the Cluster Resource Unload Script for the Linux Agents” on page 37.
- ❒ Set up the cluster failover lists and policies.

  See “Setting the Failover List and Policies for the Linux Agents” on page 39.
- ❒ Test your new clustered GroupWise system.

  See Section 3.5, “Testing Your Clustered GroupWise System on Linux,” on page 41.
- ❒ Record cluster-specific information in the properties pages of the GroupWise objects that the information pertains to.

  See Section 3.6, “Managing Your Clustered GroupWise System on Linux,” on page 42.

3.8.2 Domain Quick Checklist

- ❒ Plan your new clustered domain.

  See Section 2.3, “Planning a New Clustered Domain,” on page 18.
- ❒ Create the new domain.

  See Section 3.2, “Creating a New Secondary Domain in a Linux Cluster,” on page 26.
- ❒ Set up the MTA for the new domain.

  See Section 3.4, “Installing and Configuring the Agents in a Linux Cluster,” on page 28.
- ❒ Modify the domain cluster resource load script.

  See “Modifying the Cluster Resource Load Script for the Linux Agents” on page 34.
- ❒ Modify the domain cluster resource unload script.

  See “Modifying the Cluster Resource Unload Script for the Linux Agents” on page 37.
- ❒ Set up the domain failover list and policies.

  See “Setting the Failover List and Policies for the Linux Agents” on page 39.
- ❒ Test your new clustered domain.

  See Section 3.5, “Testing Your Clustered GroupWise System on Linux,” on page 41.
- ❒ Record cluster-specific information in the properties pages of the GroupWise objects that the information pertains to.

  See Section 3.6, “Managing Your Clustered GroupWise System on Linux,” on page 42.

3.8.3 Post Office Quick Checklist

- ❒ Plan your new clustered post office.

  See Section 2.4, “Planning a New Clustered Post Office,” on page 19 and Section 2.5, “Planning a New Library for a Clustered Post Office,” on page 19.
- ❒ Create the new post office.

See Section 3.3, “Creating a New Post Office in a Linux Cluster,” on page 27.

- ❒ Set up the POA for the new post office.

  See Section 3.4, “Installing and Configuring the Agents in a Linux Cluster,” on page 28.

- ❒ Modify the post office cluster resource load script:

  See “Modifying the Cluster Resource Load Script for the Linux Agents” on page 34.

- ❒ Modify the post office cluster resource unload script:

  See “Modifying the Cluster Resource Unload Script for the Linux Agents” on page 37.

- ❒ Set up the post office failover list and policies.

  See “Setting the Failover List and Policies for the Linux Agents” on page 39.

- ❒ Test your new clustered post office.

  See Section 3.5, “Testing Your Clustered GroupWise System on Linux,” on page 41.

- ❒ Record cluster-specific information in the properties pages of the GroupWise objects that the information pertains to.

  See Section 3.6, “Managing Your Clustered GroupWise System on Linux,” on page 42.

4 Implementing the Document Viewer Agent in a Linux Cluster

Clustering is not currently supported for the Document Viewer Agent (DVA).

However, clustering is not necessary in order to provide high availability for the DVA, because the DVA does not require the database access that is required by the MTA and POA. As an alternative to clustering the DVA, you can install the DVA on multiple servers outside the cluster, as described in "Scaling Your DVA Installation" in "Document Viewer Agent" in the *GroupWise 2012 Administration Guide*.

5 Implementing the Internet Agent in a Linux Cluster

You should already have set up at least a basic GroupWise system, as described in Chapter 2, "Planning GroupWise in a Linux Cluster," on page 15 and Chapter 3, "Setting Up a Domain and a Post Office in a Linux Cluster," on page 25. As part of this process, you filled out the "System Clustering Worksheet" on page 22. If you do not have access to the filled-out worksheet, print the worksheet now and fill in the clustering information as it currently exists on your system. You need this information as you implement the Internet Agent (GWIA) in a cluster.

5.1 Planning the Internet Agent in a Linux Cluster

A major system configuration difference between the GWIA in a clustering environment and the GWIA in a regular environment is that you need to create a separate domain to house the GWIA in the clustering environment.

The "Internet Agent Clustering Worksheet" on page 66 lists the information you need as you set up the GWIA in a clustering environment. You should print the worksheet and fill it out as you complete the tasks listed below:

5.1.1 Planning a Domain for the GWIA

The considerations involved in planning a domain for the GWIA are much the same as planning any other domain. In preparation, review "Planning a New Domain", then print and fill out the "New Domain Summary Sheet" in "Domains" in the *GroupWise 2012 Administration Guide*.

Keep in mind the following cluster-specific details:

- When you specify the location for the domain directory on the Domain Worksheet, remember that it is on a GroupWise partition, not on the node where you running the GroupWise Installation program. This location is referred to as the GWIA partition because it is where the GWIA message queues are located.
- Do not concern yourself with the GroupWise agent information on the Domain Worksheet. You can stop with item 10. You will plan the MTA installation later.

When you have completed the Domain Worksheet, transfer the key information from the Domain Worksheet to the Internet Agent Clustering Worksheet.

INTERNET AGENT CLUSTERING WORKSHEET

Under Item 1: GroupWise Partition for the GWIA, transfer the domain location to the Internet Agent Clustering Worksheet.

Under Item 2: GWIA Domain Name, transfer the domain name and database directory to the Internet Agent Clustering Worksheet.

IMPORTANT: Do not create the new domain until you are instructed to do so in Section 5.2.1, "Creating a Domain for the GWIA," on page 53.

5.1.2 Selecting the GWIA Partition and Secondary IP Address

As with the MTA, the POA, and the DVA, the GWIA needs a secondary IP address that remains the same no matter which node in the cluster it is running on. You can place the GWIA and its domain on a GroupWise partition where a domain or post office already reside, which means that the GWIA shares the same secondary IP address as that domain or post office and fails over along with that domain or post office. Or you can place the GWIA and its domain on its own GroupWise partition, which means that it has its own secondary IP address and fails over independently.

INTERNET AGENT CLUSTERING WORKSHEET

Under Item 1: GroupWise Partition for GWIA, specify the secondary IP address for the GWIA partition.

Under Item 5: MTA Network Information, copy the same secondary IP address.

Under Item 6: GWIA Network Information, copy the same secondary IP address.

IMPORTANT: You must configure the GWIA to bind exclusively to the secondary IP address. Novell Cluster Services uses Postfix to send cluster email alerts using the primary IP address. Postfix and the GWIA both default to port 25. A conflict results unless you configure the GWIA so that it does not use the primary IP address. Instructions are provided in "Forcing Use of the GWIA Secondary IP Address" on page 62.

5.1.3 Determining an Appropriate Failover List for the Linux GWIA

By default, a GroupWise partition is configured to have all nodes in the cluster in its failover list, organized in ascending alphanumeric order. Only one node at a time can have a particular GroupWise partition mounted and active. If a GroupWise partition's preferred node fails, the partition fails over to the next node in the failover list. You should customize the failover list for each GroupWise partition based on the fan-out-failover principle.

As with the MTA, the POA, and the DVA, you need to decide which nodes in the cluster are appropriate locations for the GWIA partition to fail over to. You must install the GWIA software on all of the nodes where you want the GWIA to be able to fail over. For a review of failover lists, see Section 2.6.2, "Determining Appropriate Failover Lists for the Linux Agents," on page 21, which describes the issues in the context of planning installations for the other GroupWise agents.

INTERNET AGENT CLUSTERING WORKSHEET

Under Item 3: GWIA Failover List, list the nodes that you want in the GWIA partition failover list.

5.1.4 Determining Cluster Resource Information for the Linux GWIA

A cluster resource is a shared partition, secondary IP address, application, service, Web server, and so on, that can function successfully anywhere in the cluster. Cluster resources include the GroupWise agents and the Messenger agents. When you use the *Configure GroupWise for Clustering* option, the GroupWise Installation program needs to know the mount point for the GroupWise partition where the GWIA domain is located.

INTERNET AGENT CLUSTERING WORKSHEET

Under Item 4: Cluster Resource Mount Point, list the mount point for the GroupWise partition where the GWIA domain is located.

5.1.5 Preparing DNS for the Clustered Linux GWIA

In order for the GWIA partition to be recognized on your network, DNS must have an MX record that includes the hostname corresponding to the secondary IP address of the GWIA partition. A DNS A record associates the secondary IP address with the hostname.

5.1.6 Preparing Your Firewall for the Clustered Linux GWIA

The GWIA receives incoming messages on the secondary IP address of the GWIA partition. Your firewall configuration must be modified to allow inbound TCP/IP traffic from the Internet to the GWIA secondary IP address on the following standard ports:

Protocol	Standard Port
IMAP4	143
LDAP	389
POP3	110
SMTP	25

By default, the GWIA sends outgoing messages on the primary IP address of the server where it is running. If you decide to use this default configuration, your firewall must be configured to allow outbound TCP/IP traffic from all nodes in the GWIA partition failover list. However, because Postfix uses the primary IP address for sending cluster email alerts, the default GWIA configuration is not recommended. To avoid a port conflict, configure the GWIA to bind to the secondary IP address, as described in "Forcing Use of the GWIA Secondary IP Address" on page 62.

If the GWIA has a large number of nodes on its failover list, you can configure the GWIA to send outgoing messages to a relay host, which then sends them out through the firewall using its own IP address rather than the address of the particular node where the GWIA was running. This reduces the amount of modification to your firewall required to set up the GWIA. However, if the relay host goes down, outgoing messages are delayed.

In preparation for installing the GWIA, configure your firewall as needed to handle the GWIA's use of primary and secondary IP addresses when sending and receiving messages.

5.1.7 Planning the Linux MTA Installation

Follow the instructions in Section 2.6.4, "Planning the Linux Agent Installation," on page 21, then return to this point. After you follow the instructions, you will have a filled-out Agent Clustering Worksheet to use when you install the MTA.

IMPORTANT: Do not install the Linux MTA until you are instructed to do so in Section 5.2, "Setting Up the GWIA in a Linux Cluster," on page 52.

5.1.8 Planning the Linux GWIA Installation

Aside from the cluster-specific issues discussed in the preceding sections, the considerations involved in planning to install the GWIA are the same in a clustering environment as for any other environment. Review the installation instructions in "Linux: Installing the GWIA" in "Installing the GroupWise Internet Agent" in the *GroupWise 2012 Installation Guide*. Use the "GroupWise Internet Agent Installation Worksheet" to record the planning information you will need as you install the GWIA in your cluster.

IMPORTANT: Do not install the GWIA software until you are instructed to do so in Setting Up the GWIA in a Linux Cluster.

5.2 Setting Up the GWIA in a Linux Cluster

You should already have reviewed Section 5.1, "Planning the Internet Agent in a Linux Cluster," on page 49 and filled out the "Internet Agent Clustering Worksheet" on page 66. You are now ready to complete the following tasks to set up the GWIA in a clustering environment:

- Section 5.2.1, "Creating a Domain for the GWIA," on page 53
- Section 5.2.2, "Installing the MTA for the GWIA Domain," on page 53
- Section 5.2.3, "Installing and Configuring the Linux GWIA in a Cluster," on page 53

5.2.1 Creating a Domain for the GWIA

The GWIA domain will be a secondary domain. To create it, follow the instructions in Section 3.2, "Creating a New Secondary Domain in a Linux Cluster," on page 26, taking your information from the GWIA Clustering Worksheet, rather than the System Clustering Worksheet, then return to this point.

Do not create any post offices in the GWIA domain.

After you have created the domain, continue with Installing the MTA for the GWIA Domain.

5.2.2 Installing the MTA for the GWIA Domain

The MTA for the GWIA domain can be installed just like any other MTA in your clustered GroupWise system. Follow the instructions in Section 3.4.1, "Installing and Setting Up the Linux Agents in Your Cluster," on page 28, then return to this point.

You do not need to edit the MTA startup file. You do not need to modify the Cluster Resource object properties until after you have installed the GWIA.

After you have installed the MTA, continue with Installing and Configuring the Linux GWIA in a Cluster.

5.2.3 Installing and Configuring the Linux GWIA in a Cluster

After you have created a domain for the GWIA and installed the MTA for that domain, you are ready to install and configure the GWIA.

Installing and Setting Up the Linux GWIA Software in Your Cluster

The GWIA must be installed on each node in its failover list (Internet Agent Clustering Worksheet item 3).

Running the Linux GWIA Installation Program on the Preferred Node

1. Make sure that the GWIA software is available in the software distribution directory you created in Step 6 in Section 3.1, "Setting Up a New GroupWise System in a Linux Cluster," on page 25.
2. Mount the GWIA partition (Internet Agent Clustering Worksheet item 1) where the GWIA message queues are located.

3 From the software distribution directory, start the GroupWise Installation program and select *Configure GroupWise for Clustering*.

4 Install the GWIA software, following the steps provided in "Linux: Installing the GWIA" in "Installing the GroupWise Internet Agent" in the *GroupWise 2012 Installation Guide*.

5 Configure the GWIA according to the "GroupWise Internet Agent Installation Worksheet" that you filled out in Section 5.1, "Planning the Internet Agent in a Linux Cluster," on page 49, paying special attention to the cluster resource information on the Server Information page.

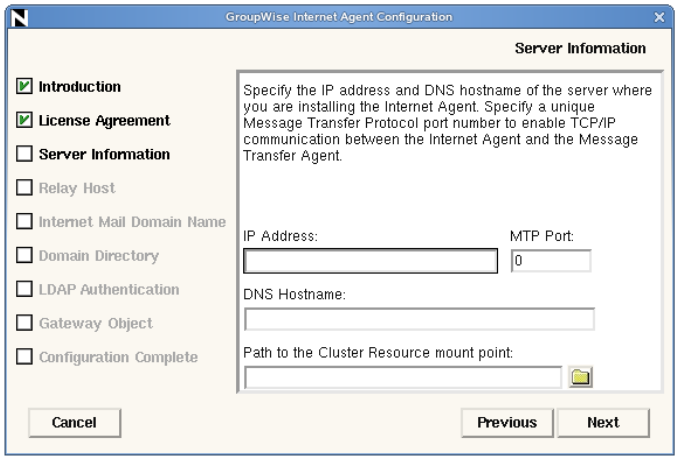

As a result of selecting *Configure GroupWise for Clustering* on the preferred node, the following cluster-specific configuration actions are performed:

- The GWIA startup file (`gwia.cfg`) is created in *`mount_point`*`/groupwise/agents/share` on the shared resource so that the GWIA uses the same startup file regardless of which cluster node it is running on. The --home switch includes the mount point and the path to the database so that the startup file is valid when mounted to each cluster node.
- The --cluster switch is added to the GWIA startup file to inform the GWIA that it is running in a cluster.
- The --ip startup switch is set to the secondary IP address of the shared resource where the domain is located. This ensures that the GWIA runs with the same IP address regardless of which cluster node it is running on.
- The --log startup switch is set to a location on the shared resource (*`mount_point`*`/groupwise/agents/log`) so that GWIA logging information is written to the same log file regardless of which cluster node it is running on.
- If this is the first GroupWise agent installed on this cluster node, the GroupWise High Availability service is automatically configured and its configuration file (`gwha.conf`) is created in the `/etc/opt/novell/groupwise` directory. If another GroupWise agent has already been installed on this cluster node, the `gwha.conf` file is updated to include the GWIA.

- The `clusterimport.conf` file in the `gwinst` subdirectory of the software distribution directory from which you ran the GroupWise Installation program is updated, so that the cluster-specific information collected when you configured the GWIA on the preferred node is available when you install the GWIA on subsequent nodes.
- The GWIA is not configured to start automatically on system startup. In a cluster, you do not want the GWIA to start automatically whenever a node restarts.

6 Configure the GWIA to run as a non-`root` user, as described in the applicable section of the *GroupWise 2012 Installation Guide*:

- "Running the Linux GroupWise Agents as a Non-root User"
- "Setting Up Non-root Access on an NSS Volume on Novell Open Enterprise Server Linux"

7 Continue with Running the Linux GWIA Installation Program on Subsequent Nodes.

Running the Linux GWIA Installation Program on Subsequent Nodes

1 On the next node in the GWIA failover list, mount the GWIA partition (Internet Agent Clustering Worksheet item 1) where the GWIA message queues are located.

2 From the software distribution directory you created in Step 6 in Section 3.1, "Setting Up a New GroupWise System in a Linux Cluster," on page 25, start the GroupWise Installation program and select *Configure GroupWise for Clustering*.

Because of the existence of the `clusterimport.conf` file in the `gwinst` subdirectory of the software distribution directory, a new installation option, *Import Clustering Data*, is now available on the main GroupWise Installation program page.

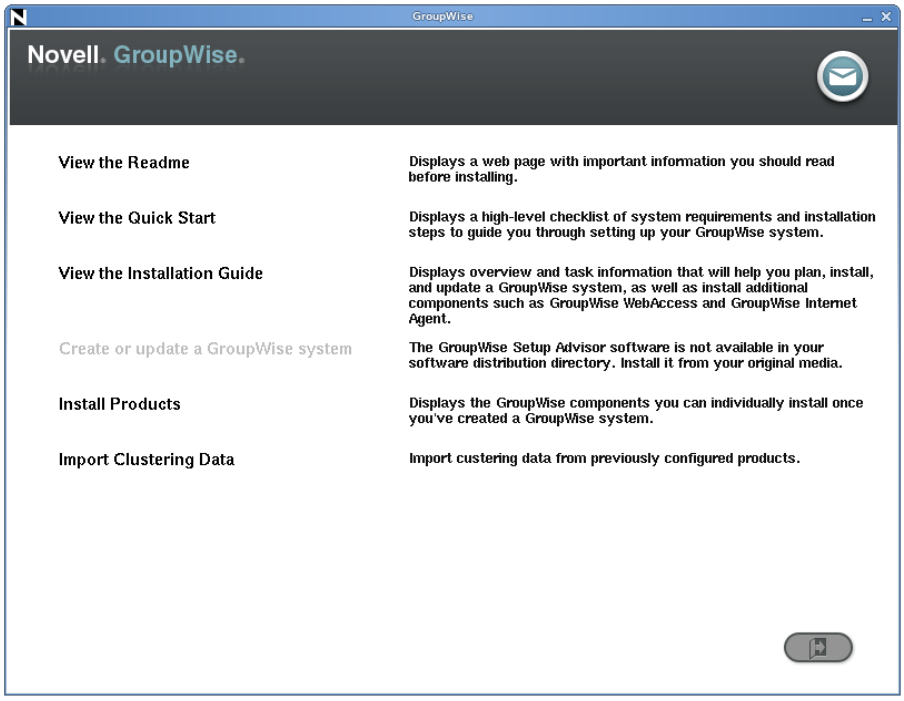

3 Install the GWIA software on the cluster node as usual, but do not use the *Configure* option.

4 On the main page of the Installation program, click *Import Clustering Data*, then click *Next*.

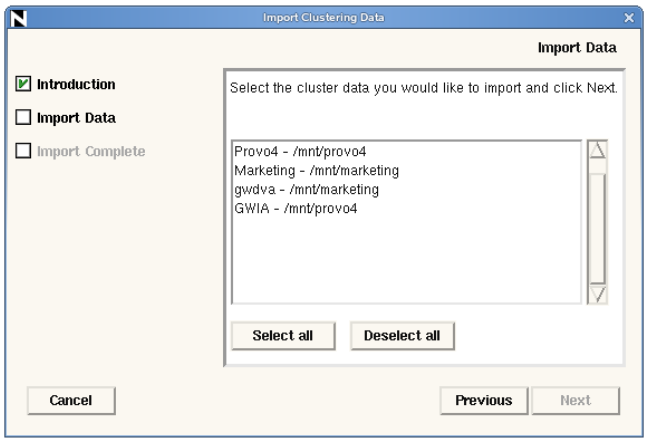

All GroupWise agents that you have installed from the software distribution directory are listed, based on the information stored in the `clusterimport.conf` file.

5 Select the GWIA, then click *OK*.

The *Import Clustering Data* option performs the following configuration actions for each subsequent node where you run it:

- The `grpwise` script is created in the `/etc/init.d` directory on the current cluster node.
- The GroupWise High Availability service is automatically configured on the current cluster node and its configuration file (`gwha.conf`) is created in the `/etc/opt/novell/groupwise` directory. It is configured specifically for the GWIA.

Because the GWIA configuration file (`gwia.cfg`) and log files are stored on the shared resource, they do not need to be customized for each cluster node.

6 Configure the GWIA to run as a non-`root` user, as described in the applicable section of the *GroupWise 2012 Installation Guide*:

- "Running the Linux GroupWise Agents as a Non-root User"
- "Setting Up Non-root Access on an NSS Volume on Novell Open Enterprise Server Linux"

7 Repeat Step 1 through Step 6 for each cluster node in the GWIA failover list.

After you install and configure the GWIA on each node in its failover list, the cluster node is ready for the GWIA to fail over to it.

8 Continue with Testing Your Linux GWIA Installation on Each Node.

Testing Your Linux GWIA Installation on Each Node

1 Test the GWIA by starting it with a user interface, as described in "Linux: Starting the GWIA" in "Installing the GroupWise Internet Agent" in the *GroupWise 2012 Installation Guide*.

```
/opt/novell/groupwise/agents/bin/gwia --show @gwia.cfg &
```

2 Stop the GWIA by clicking *File* > *Exit*.

3 After you can see that the GWIA stopped successfully, test it by starting it as a daemon, as described in "Starting the Linux Agents as Daemons" in "Installing GroupWise Agents" in the *GroupWise 2012 Installation Guide*.

```
rcgrpwise start domain.gwia
rcgrpwise status domain.gwia
```

4 Stop the GWIA.

```
rcgrpwise stop domain.gwia
rcgrpwise status domain.gwia
```

5 Make sure you have completed all the tasks described in "Installing the GroupWise Internet Agent" in the *GroupWise 2012 Installation Guide*.

A few tasks, such as assigning a postmaster, are not dealt with in this cluster-oriented section.

6 Repeat the steps in "Running the Linux GWIA Installation Program on the Preferred Node" on page 53 for each node in the GWIA failover list.

When you have installed the GWIA on all of the nodes in the GWIA failover list, continue with Configuring the Clustered Linux GWIA for SSL.

Configuring the Clustered Linux GWIA for SSL

If you plan to enable SSL, as described in "Securing GWIA Connections with SSL" in "Internet Agent" in the *GroupWise 2012 Administration Guide*, you must make the SSL certificate file and key file available to the GWIA in the cluster. The default locations for the SSL certificate file and key file are

on the cluster nodes along with the GroupWise software, rather than being located with the domain and post office on one or more GroupWise partitions. To avoid having multiple copies of these files in multiple locations, you should set the locations in ConsoleOne.

1. On the GWIA partition, create the directory where you want to store the certificate and key file required for SSL.
2. Copy the certificate file and key file into the new directory.

   If you need assistance obtaining these files, see "Server Certificates and SSL Encryption" in "Security Administration" in the *GroupWise 2012 Administration Guide*.
3. In ConsoleOne, browse to and select the Domain object.
4. Right-click the GWIA object, then click *Properties*.
5. Click *GroupWise* > *SSL Settings*.
6. In the *Certificate File* field, browse to and select the certificate file.
7. In the *SSL Key File* field, browse to and select the key file.
8. Click *OK*.
9. Continue with Configuring the Linux GWIA Cluster Resource to Load and Unload the GWIA and Its MTA.

Configuring the Linux GWIA Cluster Resource to Load and Unload the GWIA and Its MTA

The properties of the Cluster Resource object define how the GWIA partition functions within the cluster, how the GWIA is loaded and unloaded, and how failover and failback situations are handled. Complete the following tasks for the GWIA cluster resource:

- "Modifying the Cluster Resource Load Script for the Linux GWIA and Its MTA" on page 58
- "Modifying the Cluster Resource Unload Script for the Linux GWIA and Its MTA" on page 60
- "Setting the Failover List and Policies for the Linux GWIA and Its MTA" on page 62

Modifying the Cluster Resource Load Script for the Linux GWIA and Its MTA

The cluster resource load script executes whenever the GWIA cluster resource comes online.

To set up the load script in iManager:

1. Expand *Clusters*, then click *Cluster Options*.
2. In the *Cluster* field, browse to the Cluster object where the GWIA cluster resource is located.
3. Click the Cluster object to display the cluster resources that belong to the cluster.
4. Select the GWIA cluster resource that you created when you set up the GWIA partition, then click *Details*.
5. Click *Scripts* > *Load Script*.

6 (Conditional) If this is an NSS volume or a shared pool, use a load script similar to the following example, depending on the configuration of your cluster and nodes:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

# mount filesystem
exit_on_error ncpcon mount /opt="noatime,nodiratime" volume_name=volume_ID

# add IP address
exit_on_error add_secondary_ipaddress gwia_partition_ip_address

# start service
exit_on_error /etc/init.d/grpwise start domain
exit_on_error /etc/init.d/grpwise start domain.gwia
```

6a In the `mount filesystem` section, specify the volume name and volume ID of the GWIA partition that you are clustering (System Clustering Worksheet item 5):

6b In the `add IP address` section, specify the secondary IP address of the GWIA partition (System Clustering Worksheet item 7)

6c In the `start service` section, provide the commands to start the MTA first, following by the GWIA.

7 (Conditional) If this is a traditional Linux volume, use a load script similar to the following example, depending on the configuration of your cluster and nodes:

```
#! /bin/bash
. /opt/novell/ncs/lib/ncsfunc

# define IP address
RESOURCE_IP=gwia_partition_ip_address

# define filesystem type
MOUNT_FS=filesystem

# define device (if using EVMS)
exit_on_error evms -f /var/opt/novell/ncs/ContainerActivate -rl
                                                    Share `uname -n`
MOUNT_DEV=/dev/evms/Share/dat

# define mount point
MOUNT_POINT=/mnt/mount_point_directory

# mount filesystem
exit_on_error mount -t $MOUNT_FS $MOUNT_DEV $MOUNT_POINT -o noatime,nodiratime

# add IP address
exit_on_error add_secondary_ipaddress $RESOURCE_IP

# start service
exit_on_error /etc/init.d/grpwise start domain
exit_on_error /etc/init.d/grpwise start domain.gwia

exit 0
```

7a In the `define IP address` section, specify the secondary IP address of the GWIA partition (Internet Agent Clustering Worksheet item 1).

7b In the `define filesystem type` section, specify the filesystem type that is in use in use on the nodes in the cluster (System Clustering Worksheet item 5).

7c In the `define mount point` section, specify the mount point directory in use for the nodes in the cluster (System Clustering Worksheet item 5).

7d In the `start service` section, provide the commands to start the MTA first, following by the GWIA.

8 Click *OK* to save the load script.

Modifying the Cluster Resource Unload Script for the Linux GWIA and Its MTA

The cluster resource unload script executes whenever the GWIA cluster resource goes offline. Programs should be unloaded in the reverse order of how they were loaded. This ensures that supporting programs are not unloaded before programs that rely on them in order to function properly.

1 On the iManager Cluster Resource Properties page of the GWIA cluster resource, click *Scripts* > *Unload Script*.

2 (Conditional) If this is an NSS volume or a shared pool, use an unload script similar to the following example, depending on the configuration of your cluster and nodes:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

# request service stop
ignore_error /etc/init.d/grpwise stop domain.gwia
ignore_error /etc/init.d/grpwise stop domain

# stop service otherwise
sleep 8
ignore_error pkill -fx "'/opt/novell/groupwise/agents/bin/gwia
     @/media/nss/volume_name/groupwise/agents/share/gwia.cfg'"
ignore_error pkill -fx "'/opt/novell/groupwise/agents/bin/gwmta
     @/media/nss/volume_name/groupwise/agents/share/domain_name.mta'"

# delete IP address
ignore_error del_secondary_ipaddress gwia_partition_ip_address

# unmount filesystem
exit_on_error umount /mnt/mount_point_directory

# return status
exit 0
```

2a In the `request service stop` section, provide the commands to stop the GWIA first, followed by the MTA.

2b In the `stop service otherwise` section, adjust the `sleep` command as needed so that the agents can shut down normally on your system without being inadvertently killed by the `pkill` command the follows.

2c In the `delete IP address` section, specify the secondary IP address of the GWIA partition.

2d In the `unmount filesystem` section, specify the mount point directory in use for the nodes in the cluster.

2e (Conditional) If you are running the GroupWise High Availability service (gwha), stop it before the script stops the agents, then start it again at the end of the unload script.

This prevents the GroupWise High Availability service from trying to restart the agents while the script is trying to stop them.

Add the following section before the commands to stop the agents:

```
# Temporarily disable the gwha service under xinetd
ignore_error /sbin/chkconfig -s gwha off
ignore_error kill -HUP `pidof xinetd`
```

Add the following section before the `return status` section:

```
# Restart the gwha service under xinetd
ignore_error /sbin/chkconfig -s gwha on
ignore_error kill -HUP `pidof xinetd`
```

3 (Conditional) If this is a traditional Linux volume, use an unload script similar to the following example, depending on the configuration of your cluster and nodes:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

# request service stop
ignore_error /etc/init.d/grpwise stop domain.gwia
ignore_error /etc/init.d/grpwise stop domain

# stop service otherwise
sleep 8
ignore_error pkill -fx "'/opt/novell/groupwise/agents/bin/gwia
     @/media/nss/volume_name/groupwise/agents/share/gwia.cfg'"
ignore_error pkill -fx "'/opt/novell/groupwise/agents/bin/gwmta
     @/media/nss/volume_name/groupwise/agents/share/domain_name.mta'"

# define IP address
RESOURCE_IP=gwia_partition_ip_address

# define mount point
MOUNT_POINT=/mnt/mount_point_directory

# delete IP address
ignore_error del_secondary_ipaddress $RESOURCE_IP

# unmount filesystem
exit_on_error umount $MOUNT_POINT

# return status
exit 0
```

3a In the `request service stop` section, provide the commands to stop the GWIA first, followed by the MTA.

3b In the `stop service otherwise` section, adjust the `sleep` command as needed so that the agents can shut down normally on your system without being inadvertently killed by the `pkill` command the follows.

3c In the `define IP address` section, specify the secondary IP address of the GWIA partition.

3d In the `define mount point` section, specify the mount point directory in use for the nodes in the cluster.

3e (Conditional) If you are running the GroupWise High Availability (gwha) service, stop it before the script stops the agents, then start it again at the end of the unload script.

This prevents the GroupWise High Availability service from trying to restart the agents while the script is trying to stop them.

Add the following section before the `request service stop` section:

```
# Temporarily disable the gwha service under xinetd
ignore_error /sbin/chkconfig -s gwha off
ignore_error kill -HUP `pidof xinetd`
```

Add the following section before the `return status` section:

```
# Restart the gwha service under xinetd
ignore_error /sbin/chkconfig -s gwha on
ignore_error kill -HUP `pidof xinetd`
```

4 Click *OK* to save the unload script.

Setting the Failover List and Policies for the Linux GWIA and Its MTA

1. On the iManager Cluster Resource Properties page of the GWIA cluster resource, click *General*.

   The default policy settings are often appropriate. By default, a cluster resource:

   - Fails over automatically if the node it is running on fails
   - Starts automatically on the next node in its failover list
   - Continues running at its failover location, even after its most preferred node is again available

     If you are considering changing these defaults, see the *Novell Cluster Services Administration Guide for Linux* for your version of OES Linux (http://www.novell.com/documentation/oes.html).

2. Under *Preferred Nodes*, arrange the nodes in the cluster into the desired failover list for the GWIA (Internet Agent Clustering Worksheet item 3).
3. Click *OK*.

Enabling Internet Addressing for Your Clustered GroupWise System

Setting up Internet addressing for a clustered GWIA is no different from setting it up for an GWIA in any other environment. Follow the instructions in "Configuring Internet Addressing" in "Internet Agent" in the *GroupWise 2012 Administration Guide*, then return to this point.

Forcing Use of the GWIA Secondary IP Address

Novell Cluster Services uses Postfix to send cluster email alerts using the primary IP address. Postfix and the GWIA both default to using port 25. You must configure the GWIA to bind exclusively to the secondary IP address in order to avoid a port conflict between Postfix and the GWIA.

1. Click *GroupWise > Network Address*.
2. In the *TCP/IP Address* field, provide the secondary IP address (Internet Agent Clustering Worksheet item 1) for the GWIA to use for sending outgoing messages.
3. Select *Bind Exclusively to TCP/IP Address*.
4. Click *OK*.
5. Continue with Verifying GWIA Object Properties.

Verifying GWIA Object Properties

During installation of the GWIA, the GWIA object should have been configured correctly. However, it can be helpful to verify certain cluster-specific information in order to familiarize yourself with the configuration of a clustered GWIA.

- "Accessing GWIA Object Properties" on page 63
- "Verifying the Reference to the GWIA Cluster Resource" on page 63
- "Verifying the Reference to the Mount Point Directory" on page 63
- "Verifying Post Office Links" on page 63

Accessing GWIA Object Properties

1. In ConsoleOne, browse to and select the GWIA domain in order to display its contents.
2. Right-click the GWIA object, then click *Properties*.
3. Continue with Verifying the Reference to the GWIA Cluster Resource.

Verifying the Reference to the GWIA Cluster Resource

In the GWIA object properties pages in ConsoleOne:

1. Click *SMTP/MIME > Settings*.
2. Verify the contents of the *Hostname/DNS "A Record" Name* field.

   This field displays the hostname as currently configured in DNS. It should display the hostname that corresponds to the secondary IP address of the GWIA cluster resource. For more information, see Section 5.1.5, "Preparing DNS for the Clustered Linux GWIA," on page 51.
3. Make changes if necessary.
4. Continue with Verifying the Reference to the Mount Point Directory.

Verifying the Reference to the Mount Point Directory

In the GWIA object properties pages in ConsoleOne:

1. Click *Server Directories*.
2. Verify that the displayed directories match the mount point directory and the domain directory.
3. Make changes if necessary.
4. Continue with Verifying Post Office Links.

Verifying Post Office Links

In the GWIA object properties pages in ConsoleOne:

1. Click *Post Office Links*.
2. Verify that the *Access Mode* column displays *C/S* (for client/server mode) for all post offices serviced by the GWIA.
3. Verify that the *Links* column displays the secondary IP addresses of the GroupWise partitions where post offices reside, not the IP addresses of any nodes in the cluster.
4. Make changes if necessary.
5. Continue with Forcing Use of the GWIA Secondary IP Address.

5.3 Testing the Internet Agent in a Linux Cluster

After you have configured the GWIA cluster resource, you can test the load and unload scripts by bringing the cluster resource online and taking it offline again.

1. In iManager, expand *Clusters*, then click *Cluster Manager*.
2. Browse to the Cluster object to display the current cluster state.

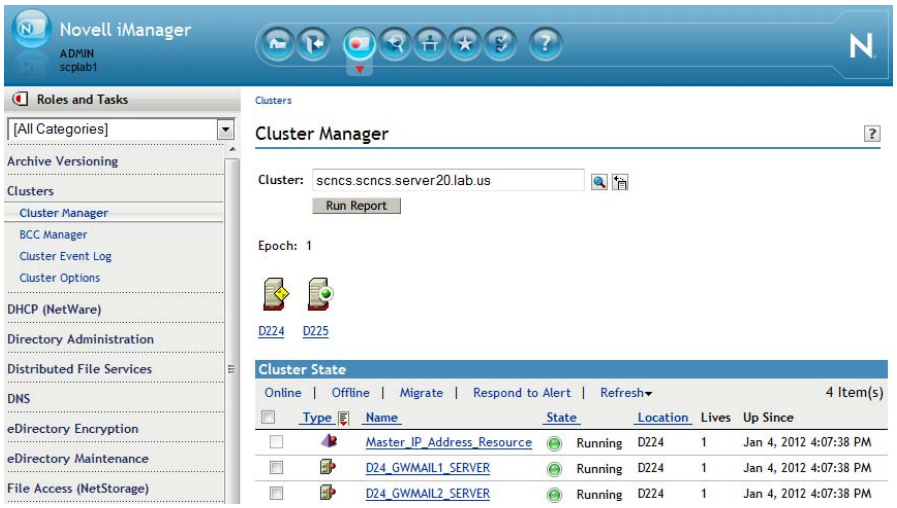

3. (Conditional) If the new GWIA cluster resource shows *Offline* in the *State* column, click the new GWIA cluster resource, then click *Online*.

   After a moment, the GWIA cluster resource displays *Running* in the *State* column.
4. At the server where the GWIA is starting, use the following command to see that the GWIA has started:

   ```
   /etc/init.d/grpwise status domain.gwia
   ```
5. Select the new GWIA cluster resource, then click *Offline*.

   The *State* column for the GWIA cluster resource returns to *Offline*.
6. Use the same command that you used in Step 4 to verify that the GWIA has stopped.
7. Repeat Step 3 whenever you are ready to bring the new GWIA cluster resource online permanently.
8. Continue with Managing Your Clustered GroupWise System on Linux.

5.4 Managing the Internet Agent in a Linux Cluster

After you have installed the GWIA in a cluster, you should consider some long-term management issues.

- Section 5.4.1, “Updating GroupWise Objects with Cluster-Specific Descriptions,” on page 65
- Section 5.4.2, “Knowing What to Expect in an GWIA Failover Situation,” on page 66

5.4.1 Updating GroupWise Objects with Cluster-Specific Descriptions

After installing the GWIA in your clustered GroupWise system, while the cluster-specific information is fresh in your mind, you should record the cluster-specific information as part of the GroupWise objects in ConsoleOne so that you can easily refer to it later. Be sure to update the information in the GroupWise objects if the configuration of your system changes.

- "Recording Cluster-Specific Information about the GWIA Domain and Its MTA" on page 65
- "Recording Cluster-Specific Information about the GWIA" on page 65

Recording Cluster-Specific Information about the GWIA Domain and Its MTA

To permanently record important cluster-specific information for the GWIA domain:

1. In ConsoleOne, browse to and right-click the Domain object, then click *Properties*.
2. In the *Description* field of the GWIA domain Identification page, provide a cluster-specific description of the GWIA domain, including the secondary IP address of its GroupWise partition.
3. Click *OK* to save the GWIA domain description.
4. Select the GWIA Domain object to display its contents.
5. Right-click the MTA object, then click *Properties*.
6. In the *Description* field of the MTA Identification page, record the secondary IP address of the GroupWise partition.

   This information appears on the MTA console, no matter which node in the cluster it is currently running on.
7. Click *OK* to save the MTA description.
8. Continue with Recording Cluster-Specific Information about the GWIA.

Recording Cluster-Specific Information about the GWIA

With the contents of the GWIA domain still displayed in ConsoleOne:

1. Right-click the GWIA object, then click *Properties*.
2. Click *GroupWise*, then click *Identification*.
3. In the *Description* field, record the secondary IP address of the GroupWise partition where the GWIA domain is located.

   This information appears on the GWIA console, no matter which node in the cluster it is currently running on.
4. Click *OK* to save the GWIA information.
5. Continue with Knowing What to Expect in an GWIA Failover Situation.

5.4.2 Knowing What to Expect in an GWIA Failover Situation

The failover behavior of the MTA for the GWIA domain is the same as for an MTA in a regular domain. See Section 3.6.2, “Knowing What to Expect in MTA, POA, and DVA Failover Situations,” on page 43.

Failover of the GWIA itself is more complex. The various clients (POP3, IMAP4, and LDAP) receive an error message that the node is not available. Most of the clients do not attempt to reconnect automatically, so the user must exit the client and restart it to reestablish the connection after the failover process is complete. Fortunately, the GWIA restarts quickly in its failover location so users can reconnect quickly.

As with the MTA, the POA, and the DVA, migration of the GWIA takes longer than failover. In fact, the GWIA can seem especially slow to shut down properly as it finishes its normal processing and stops its threads. For a busy GWIA, you might need to wait several minutes for it to shut down properly.

5.5 Internet Agent Clustering Worksheet

Item	Explanation
1) GroupWise Partition for the GWIA: Secondary IP Address:	Specify the GroupWise partition where the GWIA domain will be created, along with its secondary IP address. For more information, see Section 5.1.2, “Selecting the GWIA Partition and Secondary IP Address,” on page 50.
2) GWIA Domain Name: Domain Database Location:	Specify a unique name for the GWIA domain. Specify the directory on the GroupWise partition where you want to create the new domain. For more information, see Section 5.1.1, “Planning a Domain for the GWIA,” on page 50.
3) GWIA Failover List:	List other nodes in the cluster where the GWIA and its MTA can fail over. For more information, see Section 5.1.3, “Determining an Appropriate Failover List for the Linux GWIA,” on page 51.
4) Cluster Resource Mount Point:	Specify the mount point directory where the GWIA domain will be mounted. For more information, see Section 5.1.4, “Determining Cluster Resource Information for the Linux GWIA,” on page 51.
5) MTA Network Information: ◆ MTA IP address ◆ MTA message transfer port ◆ MTA live remote port ◆ MTA HTTP port	Record the MTA network address information that you will need as you install the MTA. For more information, see Section 5.1.7, “Planning the Linux MTA Installation,” on page 52.

Item	Explanation
6) GWIA Network Information: • GWIA IP address • GWIA HTTP port	Record the GWIA network address information that you will need to install the GWIA. For more information, see Section 5.1.8, “Planning the Linux GWIA Installation,” on page 52.

5.6 Internet Agent Quick Checklist

- ❒ Plan the new clustered GWIA, including the new domain required to house the GWIA in a clustering environment.

  See Section 5.1, “Planning the Internet Agent in a Linux Cluster,” on page 49.
- ❒ Make sure DNS includes the secondary IP address of the GWIA partition

  See Section 5.1.5, “Preparing DNS for the Clustered Linux GWIA,” on page 51
- ❒ Make sure your firewall is configured to accommodate the GWIA.

  See Section 5.1.6, “Preparing Your Firewall for the Clustered Linux GWIA,” on page 51.
- ❒ Create the new GWIA domain.

  See Section 5.2.1, “Creating a Domain for the GWIA,” on page 53.
- ❒ Set up the MTA for the new GWIA domain.

  See Section 5.2.2, “Installing the MTA for the GWIA Domain,” on page 53.
- ❒ Install the GWIA.

  See “Installing and Setting Up the Linux GWIA Software in Your Cluster” on page 53.
- ❒ Modify the GWIA cluster resource load script.

  See “Modifying the Cluster Resource Load Script for the Linux GWIA and Its MTA” on page 58.
- ❒ Modify the GWIA cluster resource unload script.

  See “Modifying the Cluster Resource Unload Script for the Linux GWIA and Its MTA” on page 60.
- ❒ Set up the GWIA failover list and policies.

  See “Setting the Failover List and Policies for the Linux GWIA and Its MTA” on page 62.
- ❒ Enable Internet Addressing for the clustered GWIA.

  See “Enabling Internet Addressing for Your Clustered GroupWise System” on page 62.
- ❒ Double-check the cluster-specific GWIA object properties.

  See “Verifying GWIA Object Properties” on page 62.
- ❒ Test the clustered GWIA.

  See Section 5.3, “Testing the Internet Agent in a Linux Cluster,” on page 64.
- ❒ Record cluster-specific information in the properties pages of the GroupWise objects associated with the GWIA.

  See Section 5.4.1, “Updating GroupWise Objects with Cluster-Specific Descriptions,” on page 65.

6 Implementing GroupWise Monitor in a Linux Cluster

You should already have set up at least a basic GroupWise system, as described in Chapter 2, "Planning GroupWise in a Linux Cluster," on page 15 and Chapter 3, "Setting Up a Domain and a Post Office in a Linux Cluster," on page 25. As part of this process, the "System Clustering Worksheet" on page 22 was filled out. If you do not have access to the filled-out worksheet, print the worksheet now and fill in the clustering and network address information as it currently exists on your system. You need this information as you implement Monitor in a cluster.

6.1 Understanding the Monitor Components

If you are not familiar with GroupWise Monitor, review "GroupWise Monitor Overview" in "Installing GroupWise Monitor" in the *GroupWise 2012 Installation Guide*.

As you plan Monitor in a clustering environment, you must keep in mind that you will plan and set up two Monitor components:

- Monitor Agent (gwmon) that will be associated with a domain in your GroupWise system
- Monitor Application (a Java servlet) that will be added to your Web server (Apache). You must install the Monitor Application on a non-clustered Web server.

You install the Monitor Agent on each node in the cluster. You install the Monitor Application to your Web server, which must not be clustered. This means that the Monitor Agent Web console at the following URL is always available, because it is part of the cluster:

`http://secondary_IP_address:8200`

However, the Monitor Web console at the following URL is not available if the Web server is down:

`http://Web_server_address/gwmon/gwmonitor`

6.2 Planning GroupWise Monitor in a Linux Cluster

A major system configuration difference between the Monitor Agent and other GroupWise agents is that the Monitor Agent needs access to a domain during installation but does not need permanent access to a domain thereafter.

The "Monitor Agent Clustering Worksheet" on page 80 lists information you need as you set up Monitor in a clustering environment. You should print the worksheet and fill it out as you complete the planning tasks listed below:

- Section 6.2.1, "Selecting a Domain for Access during Linux Monitor Agent Installation," on page 70
- Section 6.2.2, "Selecting an MTA for the Linux Monitor Agent to Access after Installation," on page 70
- Section 6.2.3, "Selecting the Monitor Agent Partition and Secondary IP Address," on page 71
- Section 6.2.4, "Determining an Appropriate Failover List for the Linux Monitor Agent," on page 71
- Section 6.2.5, "Determining Cluster Resource Information for the Linux Monitor Agent," on page 71
- Section 6.2.6, "Planning the Linux Monitor Agent Installation," on page 71

6.2.1 Selecting a Domain for Access during Linux Monitor Agent Installation

During installation, the Monitor Agent Installation program needs access to a domain database (`wpdomain.db`) in order to obtain information about agents to monitor. You might want to use the domain you created for use with the GWIA, as described in Section 5.2.1, "Creating a Domain for the GWIA," on page 53, although you can use any domain in your GroupWise system.

MONITOR CLUSTERING WORKSHEET

Under Item 2: Domain Name, specify the domain and domain directory that the Monitor Agent Installation program can use to obtain information about your GroupWIse system.

6.2.2 Selecting an MTA for the Linux Monitor Agent to Access after Installation

After installation, you can configure the Monitor Agent to be independent of a domain database. To do this, you configure the Monitor Agent to communicate with an MTA by way of TCP/IP.

MONITOR CLUSTERING WORKSHEET

Under Item 3: MTA IP Address, specify the MTA IP address and message transfer port that the Monitor Agent can use after installation to communicate with an MTA to obtain agent information.

6.2.3 Selecting the Monitor Agent Partition and Secondary IP Address

As with the MTA, the POA, and the DVA, the Monitor Agent needs a secondary IP address that remains the same no matter which node in the cluster it is running on. You can associate the Monitor Agent with the domain that was accessed during installation or with any other domain, so that they fail over together, or you can associate the Monitor Agent with its own shared partition, so that it fails over independently of any domain.

MONITOR CLUSTERING WORKSHEET

Under Item 1: GroupWise Partition for Monitor Agent, specify the secondary IP address for the Monitor Agent.

6.2.4 Determining an Appropriate Failover List for the Linux Monitor Agent

By default, a GroupWise partition is configured to have all nodes in the cluster in its failover list, organized in ascending alphanumeric order. Only one node at a time can have a particular GroupWise partition mounted and active. If a GroupWise partition's preferred node fails, the partition fails over to the next node in the failover list. You should customize the failover list for each GroupWise partition based on the fan-out-failover principle.

As with the other agents, you need to decide which nodes in the cluster are appropriate locations for the Monitor Agent to fail over to. You must install the Monitor Agent software on all of the nodes where you want the Monitor Agent to be able to fail over. For a review of failover lists, see Section 2.6.2, "Determining Appropriate Failover Lists for the Linux Agents," on page 21, which describes the issues in the context of planning MTA, POA, and DVA installations.

MONITOR CLUSTERING WORKSHEET

Under Item 4: Monitor Agent Failover List, list the nodes that you want in the Monitor Agent failover list.

6.2.5 Determining Cluster Resource Information for the Linux Monitor Agent

A cluster resource is a shared partition, secondary IP address, application, service, Web server, and so on, that can function successfully anywhere in the cluster. Cluster resources include the GroupWise agents and the Messenger agents. When you use the *Configure GroupWise for Clustering* option, the GroupWise Installation program needs to know the mount point for the GroupWise partition where it can access a domain database in order to gather information about agents to monitor. The Installation program also needs to know the secondary IP address of the GroupWise partition.

MONITOR AGENT CLUSTERING WORKSHEET

Under Item 5: Cluster Resource Information, list the mount point and secondary IP address for the GroupWise partition where the domain and post office will be located.

6.2.6 Planning the Linux Monitor Agent Installation

Aside from the cluster-specific issues discussed in the preceding sections, the considerations involved in planning to install the Monitor Agent are the same in a clustering environment as for any other environment. Review the installation instructions in "Installing and Configuring the Linux Monitor Agent" in "Installing GroupWise Monitor" in the *GroupWise 2012 Installation Guide*. Use the "GroupWise Monitor Installation Worksheets" to record the types of planning information you need as you install the Monitor Agent in your cluster.

IMPORTANT: Do not install the Monitor Agent software until you are instructed to do so in Section 6.3, "Setting Up GroupWise Monitor in a Linux Cluster," on page 72.

6.3 Setting Up GroupWise Monitor in a Linux Cluster

GroupWise Monitor depends on the Apache Web server on Linux. However, Apache is not typically installed in a cluster and the Monitor Application is not supported in a cluster. Therefore, these instructions do not include that task.

6.3.1 Installing and Configuring the Linux Monitor Agent on Each Node in Your Cluster

The Monitor Agent must be installed on each node in the Monitor Agent failover list (Monitor Agent Clustering Worksheet item 4). The Monitor Application is installed to your Web server and is therefore not installed on nodes in the cluster.

Running the Linux Monitor Installation Program on the Preferred Node

1. Make sure that the Monitor Agent software is available in the software distribution directory you created in Step 6 in Section 3.1, "Setting Up a New GroupWise System in a Linux Cluster," on page 25.
2. Mount the GroupWise partition (Monitor Agent Clustering Worksheet item 2) where the Monitor Agent Installation program can access a domain database.
3. From the software distribution directory, start the Installation program and select *Configure GroupWise for Clustering*.

4. Install and configure the Monitor Agent software, following the steps provided in "Installing and Configuring the Linux Monitor Agent" in "Installing GroupWise Monitor" in the *GroupWise 2012 Installation Guide*.

Pay special attention to the cluster resource information on the System Options page.

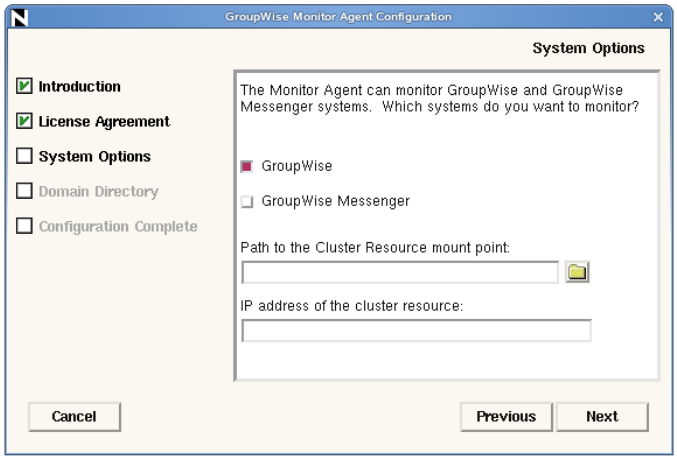

As a result of selecting *Configure GroupWise for Clustering* on the preferred node, the following cluster-specific configuration actions are performed:

- The Monitor Agent configuration file (`monitor.xml`) is created in *`mount_point`*`/groupwise/agents/share` on the shared resource so that the Monitor Agent uses the same configuration file regardless of which cluster node it is running on. The HOME_PATH option includes the mount point and the path to the database so that the configuration file is valid when mounted to each cluster node.
- The --log startup switch in the `grpwise-ma` script is set to a location on the shared resource (*`mount_point`*`/groupwise/agents/log`) so that Monitor Agent logging information is written to the same log file regardless of which cluster node it is running on. Gateway accounting files that you can use to generate reports are stored in the `acct` subdirectory of this location.
- The Monitor Agent is not configured to start automatically on system startup. In a cluster, you do not want the Monitor Agent to start automatically whenever a node restarts.

5 Continue with Running the Linux Monitor Agent Installation Program on Subsequent Nodes

Running the Linux Monitor Agent Installation Program on Subsequent Nodes

1 On the next node in the Monitor Agent failover list, mount the GroupWise partition (Monitor Agent Clustering Worksheet item 2) where the Monitor Agent Installation program can access a domain database.

2 From the software distribution directory you created in Step 6 in Section 3.1, "Setting Up a New GroupWise System in a Linux Cluster," on page 25, start the GroupWise Installation program and select *Configure GroupWise for Clustering*.

3 Install the Monitor Agent software on the cluster node as usual, but do not use the *Configure* option.

For the Monitor Agent, you do not need to import clustering data on subsequent nodes as you do for the other GroupWise agents.

4 Repeat Step 1 through Step 3 for each cluster node in the Monitor Agent failover list.

After you install the Monitor Agent on each node in its failover list, the cluster node is ready for the Monitor Agent to fail over to it.

5 Continue with Configuring the Linux Monitor Agent Web Console for SSL.

Configuring the Linux Monitor Agent Web Console for SSL

If you plan to secure the Monitor Web console using SSL, you need to provide an SSL certificate file. You can place the file on the Monitor Agent partition, rather than each node.

1 Create a directory on the Monitor Agent partition where you want to store the certificate file.

2 In the `grpwise-ma` script, use the --httpcertfile switch to specify the full path to the directory you created in Step 1.

Continue with Testing the Linux Monitor Agent Installation on Each Node.

Testing the Linux Monitor Agent Installation on Each Node

1 Test the Monitor by starting it as a daemon, as described in "Starting the Linux Monitor Agent as a Daemon" in "Installing GroupWise Monitor" in the *GroupWise 2012 Installation Guide*.

```
rcgrpwise-ma start
rcgrpwise-ma status
```

2 Then stop the Monitor Agent.

```
rcgrpwise-ma stop
rcgrpwise-ma status
```

3 Return to "Running the Linux Monitor Installation Program on the Preferred Node" on page 72 for each node in the Monitor Agent failover list (Monitor Agent Clustering Worksheet item 4)

When you have installed the Monitor Agent on all of the nodes in the Monitor Agent failover list, continue with Configuring the Monitor Agent Cluster Resource to Load and Unload the Linux Monitor Agent.

6.3.2 Configuring the Monitor Agent Cluster Resource to Load and Unload the Linux Monitor Agent

The properties of the Monitor Agent Cluster Resource object define how the Monitor Agent functions within the cluster, how the Monitor Agent is loaded and unloaded, and how failover and failback situations are handled. Complete the following tasks for the Monitor Agent cluster resource:

- "Modifying the Cluster Resource Load Script for the Linux Monitor Agent" on page 74
- "Modifying the Cluster Resource Unload Script for the Linux Monitor Agent" on page 77
- "Setting the Failover List and Policies for the Linux Monitor Agent" on page 78

Modifying the Cluster Resource Load Script for the Linux Monitor Agent

The cluster resource load script executes whenever the Monitor Agent cluster resource comes online.

To set up the load script in iManager:

1. Expand Clusters, then click *Cluster Options*.
2. In the *Cluster* field, browse to the Cluster object where the Monitor Agent cluster resource is located.
3. Click the Cluster object to display the cluster resources that belong to the cluster.
4. Select the Monitor Agent cluster resource that you created when you set up the Monitor Agent partition, then click *Details*.
5. Click *Scripts* > *Load Script*.
6. (Conditional) If this is an NSS volume or a shared pool, use a load script similar to the following example, depending on the configuration of your cluster and nodes:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

# mount filesystem
exit_on_error ncpcon mount /opt="noatime,nodiratime" volume_name=volume_ID

# add IP address
exit_on_error add_secondary_ipaddress monitor_partition_ip_address

# start service
exit_on_error /etc/init.d/grpwise-ma start

# return status
exit 0
```

   6a In the `mount filesystem` section, specify the volume name and volume ID of the GroupWise partition that you are clustering (System Clustering Worksheet item 5).

   6b In the `add IP address` section, specify the secondary IP address of the GroupWise partition (System Clustering Worksheet item 7).

7 (Conditional) If this is a traditional Linux volume, use a load script similar to the following example, depending on the configuration of your cluster and nodes:

```
#! /bin/bash
. /opt/novell/ncs/lib/ncsfunc

# define IP address
RESOURCE_IP=gw_partition_ip_address

# define filesystem type
MOUNT_FS=filesystem

# define device (if using EVMS)
exit_on_error evms -f /var/opt/novell/ncs/ContainerActivate -rl
                                                    Share `uname -n`
MOUNT_DEV=/dev/evms/Share/dat

# define mount point
MOUNT_POINT=/mnt/mount_point_directory

# mount filesystem
exit_on_error mount -t $MOUNT_FS $MOUNT_DEV $MOUNT_POINT -o noatime,nodiratime

# add IP address
exit_on_error add_secondary_ipaddress $RESOURCE_IP

# start service
exit_on_error /etc/init.d/grpwise-ma start

exit 0
```

7a In the `define IP address` section, specify the secondary IP address of the GroupWise partition (Monitor Agent Clustering Worksheet item 1).

7b In the `define filesystem type` section, specify the file system type that is in use on the nodes in the cluster (System Clustering Worksheet item 5).

7c In the `define mount point` section, specify the mount point directory in use for the nodes in the cluster (System Clustering Worksheet item 5).

8 Click *OK* to save the load script.

Modifying the Cluster Resource Unload Script for the Linux Monitor Agent

The cluster resource unload script executes whenever the Monitor Agent cluster resource goes offline.

1. On the iManager Cluster Resource Properties page of the Monitor Agent cluster resource, click *Scripts* > *Unload Script*.
2. (Conditional) If this is an NSS volume or a shared pool, use an unload script similar to the following example, depending on the configuration of your cluster and nodes:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

# request service stop
ignore_error /etc/init.d/grpwise-ma stop

# stop service otherwise
sleep 8
ignore_error pkill -fx "'/opt/novell/groupwise/agents/bin/gwmon
                            --home /domain_directory'"

# delete IP address
ignore_error del_secondary_ipaddress gw_partition_ip_address

# unmount filesystem
exit_on_error umount /mnt/mount_point_directory

# return status
exit 0
```

   2a In the `stop service otherwise` section, adjust the `sleep` command as needed so that the Monitor Agent can shut down normally on your system without being inadvertently killed by the `pkill` command that follows.

   2b In the `delete IP address section`, specify the secondary IP address of the GroupWise partition.

   2c In the `unmount file system` section, specify the mount point directory in use for the nodes in the cluster.

3 (Conditional) If this is a traditional Linux volume, use an unload script similar to the following example, depending on the configuration of your cluster and nodes:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

# request service stop
/etc/init.d/grpwise-ma stop

# stop service otherwise
sleep 8
ignore_error pkill -fx "'/opt/novell/groupwise/agents/bin/gwmon
                                                --home /domain_directory'"

# define IP address
RESOURCE_IP=172.16.5.18

# define mount point
MOUNT_POINT=/mnt/mount_point

# del the IP address
ignore_error del_secondary_ipaddress $RESOURCE_IP

# umount the file system
exit_on_error umount $MOUNT_POINT

# return status
exit 0
```

3a In the `stop service otherwise` section, adjust the `sleep` command as needed so that the Monitor Agent can shut down normally on your system without being inadvertently killed by the `pkill` command the follows.

3b In the `define IP address` section, specify the secondary IP address of the GroupWise partition.

3c In the `define mount point` section, specify the mount point directory used in use for the nodes in the cluster.

4 Click *OK* to save the unload script.

Setting the Failover List and Policies for the Linux Monitor Agent

1 On the iManager Cluster Resource Properties page of the Monitor Agent cluster resource, click *General*.

The default policy settings are often appropriate. By default, a cluster resource:

- Fails over automatically if the node it is running on fails
- Starts automatically on the next node in its failover list
- Continues running at its failover location, even after its most preferred node is again available

If you are considering changing these defaults, see the *Novell Cluster Services Administration Guide for Linux* for your version of OES Linux (http://www.novell.com/documentation/oes.html).

2 Under *Preferred Nodes*, arrange the nodes in the cluster into the desired failover list for the Monitor Agent (Monitor Agent Clustering Worksheet item 4).

3 Click *OK*.

6.4 Testing the Monitor Agent in a Linux Cluster

After you have configured the Monitor Agent cluster resource, you can test the load and unload scripts by bringing the Monitor Agent cluster resource online and taking it offline again.

1. In iManager, expand Clusters, then click *Cluster Manager*.
2. Browse to the Cluster object to display the current cluster state.

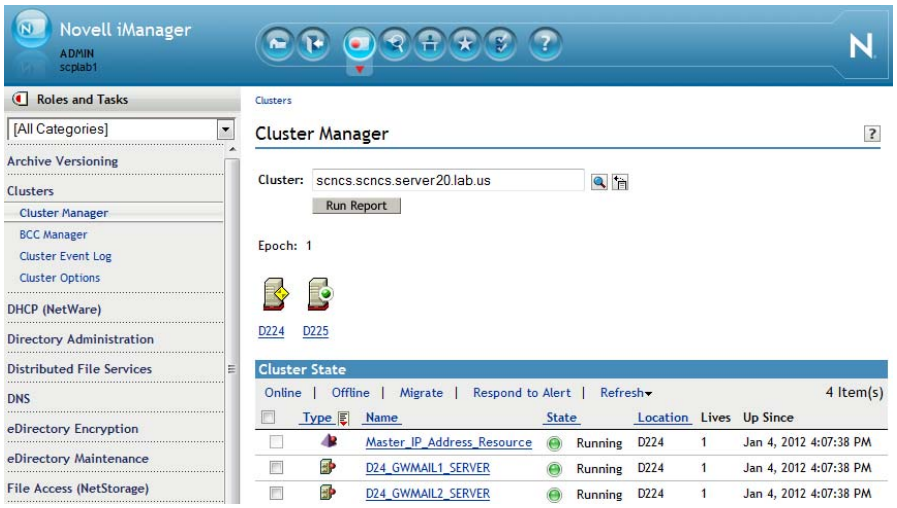

3. (Conditional) If the new Monitor Agent cluster resource shows *Offline* in the *State* column, click the new Monitor Agent cluster resource, then click *Online*.

   After a moment, the Monitor Agent cluster resource displays *Running* in the *State* column.
4. At the server where the Monitor Agent is starting, use the following command to see that the Monitor Agent has started:

   ```
   /etc/init.d/grpwise-ma status
   ```
5. Select the new Monitor Agent cluster resource, then click *Offline*.

   The *State* column for the Monitor Agent cluster resource displays *Offline*.
6. Use the same command that you used in Step 4 to verify that the Monitor Agent has stopped.
7. Repeat Step 3 whenever you are ready to bring the new Monitor Agent cluster resource online permanently.
8. Continue with Managing the Monitor Agent in a Linux Cluster.

6.5 Managing the Monitor Agent in a Linux Cluster

When the Monitor Agent fails over, it must repoll all the monitored agents to ascertain their current status. This might take a few moments, depending on the number of agents being monitored. However, no action is necessary on your part as the Monitor Agent starts on the next node in the cluster.

6.6 Monitor Agent Clustering Worksheet

Item	Explanation
1) GroupWise Partition for Monitor Agent: Secondary IP Address:	Specify the name of the Cluster Resource object for the Monitor Agent, along with its secondary IP address. For more information, see Section 6.2.3, "Selecting the Monitor Agent Partition and Secondary IP Address," on page 71.
2) GroupWise Partition for Domain Domain Name: Domain Directory:	Specify a GroupWise partition where there is a domain database from which the Monitor Agent can gather information about agents to monitor. Also provide the domain name and directory. For more information, see Section 6.2.1, "Selecting a Domain for Access during Linux Monitor Agent Installation," on page 70.
3) MTA IP Address: MTA MTP Port Number:	If you want the Monitor Agent to be able to fail over independently, specify the IP address and message transfer port number of an MTA with which the Monitor Agent can communicate, as an alternative to accessing a domain database. For more information, see Section 6.2.2, "Selecting an MTA for the Linux Monitor Agent to Access after Installation," on page 70
4) Monitor Agent Failover List:	List the nodes in the cluster where the Monitor Agent can fail over. For more information, see Section 6.2.4, "Determining an Appropriate Failover List for the Linux Monitor Agent," on page 71.
5) Cluster Resource Information ◆ Path to the cluster resource mount point ◆ IP address of the cluster resource	List the cluster resource information for the GroupWise partition where the domain is located so that the Monitor Agent can access its domain database for information about agents to monitor. For more information see, Section 6.2.5, "Determining Cluster Resource Information for the Linux Monitor Agent," on page 71.

6.7 Monitor Agent Quick Checklist

- ❒ Plan the new clustered Monitor Agent, including a domain to access during installation to gather information about agents to monitor

  See Section 6.2, "Planning GroupWise Monitor in a Linux Cluster," on page 70.

- ❒ Install the Monitor Agent on all nodes in the Monitor Agent's failover list.

  See Section 6.3.1, "Installing and Configuring the Linux Monitor Agent on Each Node in Your Cluster," on page 72.

- ❒ Modify the Monitor Agent cluster resource load script.

See “Modifying the Cluster Resource Load Script for the Linux Monitor Agent” on page 74.

- ❒ Modify the Monitor Agent cluster resource unload script.

  See “Modifying the Cluster Resource Unload Script for the Linux Monitor Agent” on page 77.

- ❒ Set up the Monitor Agent failover list and policies.

  See “Setting the Failover List and Policies for the Linux Monitor Agent” on page 78.

- ❒ Test the clustered Monitored Agent.

  See Section 6.4, “Testing the Monitor Agent in a Linux Cluster,” on page 79.

7 Implementing the GroupWise Web Applications in a Cluster

Clustering the GroupWise Web applications (WebAccess, Calendar Publishing Host, and Monitor) is not currently supported.

Because the GroupWise Web applications are installed into your Web server, clustering the GroupWise Web applications would entail clustering your entire Web server. If you are already successfully clustering your Web server, you should be able to install the GroupWise Web applications into that clustered Web server in the same way that you have installed other Web applications into your clustered Web server. However, you cannot receive support to set up this configuration of the GroupWise Web applications. For guidance, refer to the "WebAccess Application (Web Server) (http://www.novell.com/communities/node/12389/oes2-rolling-cluster-upgrade-netware-part-2-migrating-groupwise#WebAccess%20Application%20(Web%20Server))" section in a helpful GroupWise 8 Cool Solutions article. The principles are the same for GroupWise 8 and GroupWise 2012.

8 Backing Up a GroupWise System in a Linux Cluster

To back up GroupWise data in a Linux cluster, you can use the GroupWise Database Copy (DBCopy) utility to copy the data from the live GroupWise system to a static location for backup. For more information, see "Backing Up GroupWise Databases" and "GroupWise Database Copy Utility" in "Databases" in the *GroupWise 2012 Administration Guide*.

To restore data in a clustering environment, you must run your backup/restore software on the node where the location to restore is currently mounted.

9 Updating a GroupWise System in a Linux Cluster

In a Linux cluster, you must install the GroupWise software on each node in the cluster. Before you run the GroupWise Installation program to install updated software, make sure you know all the cluster nodes where the GroupWise software is already installed.

It is very important to update all nodes on the failover list of each domain and post office at the same time because each domain and post office should be serviced by only one version of the agent software. If you do not update all nodes on the failover list at once, there is the potential for a domain or post office to be serviced by a different version of the agent software during a failover situation. This can cause database problems.

Keep in mind these cluster-specific details as you follow the instructions in "Update" in the *GroupWise 2012 Installation Guide* to update your GroupWise system in a Linux cluster.

10 Moving an Existing Linux GroupWise 2012 System into a Linux Cluster

If you are adding the high availability benefits of Novell Cluster Services to a GroupWise 2012 system that is already up and running, the first step is to install Novell Cluster Services, following the instructions in *Novell Cluster Services Administration Guide for Linux* for your version of OES Linux (http://www.novell.com/documentation/oes.html). You should also review Chapter 1, "Introduction to GroupWise 2012 and Novell Cluster Services on Linux," on page 13 to help you apply clustering principles and practices to your GroupWise system.

You do not need to transfer your entire GroupWise system into the cluster all at once. You can transfer individual post offices where the needs for high availability are greatest. You can transfer a domain and all of its post offices at the same time. You might decide that you don't need to have all of your GroupWise system running in the cluster.

This section provides a checklist to help you get started with moving your GroupWise system into a clustering environment:

- ❒ Decide which shared partitions in your cluster you want to use for GroupWise domains and post offices.
- ❒ Decide which nodes in your storage area network you want have on failover lists for the GroupWise agents.
- ❒ Review Chapter 2, "Planning GroupWise in a Linux Cluster," on page 15. Fill out the "System Clustering Worksheet" on page 22 to help you decide which domains and post offices you want move to which shared partitions.
- ❒ Move a domain and/or post office onto the GroupWise partition, following the instructions in "Moving a Domain" in "Domains" or "Moving a Post Office" in "Post Offices" in the *GroupWise 2012 Administration Guide*.
- ❒ Review Section 2.6, "Deciding How to Install and Configure the Linux Agents in a Cluster," on page 20, fill out the "Agent Clustering Worksheet" on page 23, and install the agents as needed for the first clustered domain and/or post office, following the instructions in Section 3.4, "Installing and Configuring the Agents in a Linux Cluster," on page 28. This includes setting up the load and unload scripts for the GroupWise partition.
- ❒ Test the first component of your clustered GroupWise system, following the instructions in Section 3.5, "Testing Your Clustered GroupWise System on Linux," on page 41.
- ❒ Take care of the cluster management details described in Section 3.6, "Managing Your Clustered GroupWise System on Linux," on page 42.
- ❒ Move more domains and post offices into the cluster as needed. If you have GroupWise libraries, see Section 2.5, "Planning a New Library for a Clustered Post Office," on page 19.
- ❒ Add other components to your clustered GroupWise system as needed, following the instructions in:
  - Chapter 5, "Implementing the Internet Agent in a Linux Cluster," on page 49

- Chapter 6, “Implementing GroupWise Monitor in a Linux Cluster,” on page 69
- Chapter 8, “Backing Up a GroupWise System in a Linux Cluster,” on page 85

11 Implementing Messenger in a Linux Cluster

Novell Messenger does not require the existence of a GroupWise system in the cluster, but presumably one has already been set up as described in Chapter 2, "Planning GroupWise in a Linux Cluster," on page 15 and Chapter 3, "Setting Up a Domain and a Post Office in a Linux Cluster," on page 25. As part of the process of setting up GroupWise in the cluster, you filled out the "System Clustering Worksheet" on page 22. Some of the information from that worksheet is helpful as you implement Messenger in your cluster.

11.1 Planning Your Messenger System in a Linux Cluster

Because the Messenger agents are not associated with GroupWise domains or post offices, the Messenger agents are easier to implement in a cluster than are the GroupWise agents. Section 11.5, "Messenger Clustering Worksheet," on page 103 lists the information you need as you set up the Messenger agents in a clustering environment. You should print the worksheet and fill it out as you complete the tasks listed below:

11.1.1 Understanding Your Cluster

As described in Section 2.1, "Installing Novell Cluster Services on Linux," on page 16, you set up your cluster with a certain number of shared partitions and cluster resources.

MESSENGER CLUSTERING WORKSHEET

Under Items 1-5, record information about your cluster. This information corresponds to items 1-5 on the "System Clustering Worksheet" on page 22.

11.1.2 Selecting the Messenger Partition and Secondary IP Address

If you are not planning to enable archiving, or if you are not anticipating a large Messenger archive, you can use one Messenger partition for both the Messaging Agent and the Archive Agent. If you anticipate archiving a large number of messages so that the Messenger archive grows very large, you might want to have a separate Messenger partition for the Archive Agent and its archive database. The steps in this section focus on setting up the Messenger agents on a single Messenger partition.

MESSENGER CLUSTERING WORKSHEET

Under Item 6: Shared Partition for Messenger, record the name and secondary IP address of the Messenger partition in your cluster.

If you want a separate Messenger partition for archiving, under Item 7: Shared Partition for Archiving, record the name and secondary IP address of the archiving partition in your cluster.

11.1.3 Determining an Appropriate Failover List for the Linux Messenger Agents

By default, a Messenger partition is configured to have all nodes in the cluster in its failover list, organized in ascending alphanumeric order. Only one node at a time can have the Messenger partition mounted and active and the Messenger agents running. If a Messenger partition's preferred node fails, the partition fails over to the next node in the failover list. The Messenger agents might need to run on any node that the Messenger partition fails over to.

MESSENGER CLUSTERING WORKSHEET

Under Item 8: Failover List for Messenger Agents, list the nodes that you want to have in the Messenger agents' failover list.

11.1.4 Determining Cluster Resource Information for the Linux Messenger Agents

A cluster resource is a shared partition, secondary IP address, application, service, Web server, and so on, that can function successfully anywhere in the cluster. Cluster resources include the GroupWise agents and the Messenger agents. When you are installing the Messenger agents in a cluster, the Messenger Installation program needs to know the mount point for the Messenger partition where it can store agent startup files, log files, SSL certificate files, and the `uid.conf` file that enables the Messenger agents to run as a non-`root` user. By storing these files on a shared partition, the Messenger agents can access the files regardless of which node in the cluster the agents are currently running on.

MESSENGER AGENT CLUSTERING WORKSHEET
Under Item 9: Mount Point for Shared Storage, list the mount point directory for the Messenger partition where the Messenger startup and other files will be located.

11.1.5 Planning the Linux Messenger Agent Installation

Aside from the cluster-specific issues discussed in the preceding sections, the considerations involved in planning to install the Messenger agents are the same in a clustering environment as for any other environment. Review "Planning Your Novell Messenger System", then print and fill out the "Novell Messenger Worksheet" in "Installing a Novell Messenger System" in the *Novell Messenger 2.2 Installation Guide*. Messenger must be installed on each node in the failover list (Messenger Clustering Worksheet item 8)

Continue with Setting Up Your Messenger System in a Linux Cluster.

11.2 Setting Up Your Messenger System in a Linux Cluster

You should have already reviewed Section 11.1, "Planning Your Messenger System in a Linux Cluster," on page 91 and filled out the "Messenger Clustering Worksheet" on page 103 and the "Novell Messenger Worksheet" in the *Novell Messenger 2.2 Installation Guide*.

11.2.1 Creating Your Messenger System and Installing the Linux Messenger Agents

The Messenger Installation program walks you through setting up your Messenger system and installing the Messenger agents. The first time you run the Messenger Installation program, you create your Messenger system, which includes creating various Messenger objects in eDirectory and installing the Messenger software on the node where you run the Messenger Installation program. After that, you run the Messenger Installation program on each node in the Messenger failover list to install the Messenger software on each node, but you do not create any more objects in eDirectory.

Running the Linux Messenger Installation Program on the Preferred Node

1. Mount the Messenger partition (Messenger Clustering Worksheet item 6) on the mount point for shared storage (Messenger Clustering Worksheet item 9).
2. Run the Messenger Installation program, following the steps provided in "Starting the Messenger Installation Program on Linux" in "Installing a Novell Messenger System" in the *Novell Messenger 2.2 Installation Guide*.
3. When asked if you are installing to a cluster, enter `y` for *Yes*.
4. From the options list, enter `1` for *Create a new system*.
5. Specify the mount point for the shared storage.
6. Set up your Messenger system, following the steps provided in "Configuring Your Messenger System on Linux" in "Installing a Novell Messenger System" in the *Novell Messenger 2.2 Installation Guide*.
7. Continue with Running the Linux Messenger Installation Program on Subsequent Nodes.

Running the Linux Messenger Installation Program on Subsequent Nodes

1. On the next node in the Messenger failover list (Messenger Cluster Worksheet item 8), mount the Messenger partition on the mount point for shared storage.
2. Run the Messenger Installation program.
3. When asked if you are installing to a cluster, enter `y` for *Yes*.
4. From the options list, enter `2` for *Install a new server to an existing system*.
5. Specify the mount point for the shared storage.

   The Messenger Installation program then accesses the Messenger files that were created on the shared storage when the Messenger agents were installed on the preferred node. From these files, the Messenger Installation program lists the probable configuration for the Messenger agents you are installing on the current node.
6. Enter `1` for *Proceed with these settings*.

   or

   Enter `2` for *Change the settings*, then modify the configuration for the Messenger agents as needed.
7. When asked if you want to start the agents, enter `n` for *No*.
8. Repeat Step 1 through Step 7 for each node on the Messenger failover list.
9. Continue with Setting Up Non-root Access on NSS Volumes on Each Node.

Setting Up Non-root Access on NSS Volumes on Each Node

If your cluster uses NSS volumes, as described in the *Novell Cluster Services Administration Guide for Linux* for your version of OES Linux (http://www.novell.com/documentation/oes.html), you must set up non-`root` access to those NSS volumes, as described in "Setting Up Non-root Access on an NSS Volume on Novell Open Enterprise Server Linux" in "Installing a Novell Messenger System" in the *Novell Messenger 2.2 Installation Guide*. Then continue with Testing Your Linux Messenger Agent Installation on Each Node.

Testing Your Linux Messenger Agent Installation on Each Node

1 Test the Messenger agents by starting them as daemons, as described in "Starting the Linux Messenger Agents" in "Installing a Novell Messenger System" in the *Novell Messenger 2.2 Installation Guide*.

```
/etc/init.d/novell-nmma start
/etc/init.d/novell-nmaa start
/etc/init.d/novell-nmma status
/etc/init.d/novell-nmaa status
```

2 Stop the Messenger agents.

```
/etc/init.d/novell-nmma stop
/etc/init.d/novell-nmaa stop
/etc/init.d/novell-nmma status
/etc/init.d/novell-nmaa status
```

3 Return to "Running the Linux Messenger Installation Program on the Preferred Node" on page 94 for each node in the Messenger failover list (Messenger Clustering Worksheet item 8).

When you have installed the Messenger agents on all of the nodes in the Messenger failover list, continue with Changing Messenger Paths to Locations on the Messenger Partition.

11.2.2 Changing Messenger Paths to Locations on the Messenger Partition

During installation, various Messenger paths are set to locations on the node where the software is installed. After installation, you need to set these paths to locations on the Messenger partition, so that the files stored at these locations are available to the Messenger agents regardless of which node in the cluster the agents are running on:

After settings these directories, continue with Section 11.2.3, "Configuring the Messenger Cluster Resource to Load and Unload the Linux Messenger Agents," on page 97.

Setting the Store Path

The store path is the location where you want the archive created. During installation, the default store path is created in `/var/opt/novell/messenger/aa/store` on each node, but you need the archive to be stored on the Messenger partition.

1 Choose a directory where you want to store the archive and create that directory on the Messenger partition.

2 In ConsoleOne, browse to and select the Novell Messenger Service object (MessengerService), then click *Messenger Server* > *Archive Agent*.

3 Right-click the File Module object, then click *Properties*.

4 In the *Store Path* field, specify your archive store path, then click *OK*.

Setting the Messaging Agent Queue Path

When archiving is enabled, the Messaging Agent passes conversations to the Archive Agent when the conversations are completed. If the Messaging Agent cannot communicate with the Archive Agent when it has a conversation to archive, it saves the conversation in its holding directory (queue) until it can communicate with the Archive Agent again. During installation, the default Messaging Agent queue path is created in `/var/opt/novell/messenger/ma/queue`, but you need the queue directory to be located on the Messenger partition.

1. Choose a directory for the Messaging Agent queue and create that directory on the Messenger partition.
2. In ConsoleOne, browse to and select the Novell Messenger Service object (MessengerService), then click *Messenger Server*.
3. Right-click the Messaging Agent object, then click *Properties*.
4. Click *Agent* > *Messaging*.
5. In the *Messaging Queue Path* field, specify the Messaging Agent queue path, then click *OK*.

Setting the Archive Agent Queue Path

When the Archive Agent receives a conversation to archive, if it is already busy processing other conversations, it temporarily stores the conversation in its holding directory (queue). During installation, the default Archive Agent queue path is created in `/var/opt/novell/messenger/aa/queue`, but you need the queue directory to be located on the Messenger partition.

1. Choose a directory for the Archive Agent queue and create that directory on the Messenger partition.
2. In ConsoleOne, browse to and select the Novell Messenger Service object (MessengerService), then click *Messenger Server*.
3. Right-click the Archive Agent object, then click *Properties*.
4. Click *Agent* > *Messaging*.
5. In the *Archive Queue Path* field, specify the Archive Agent queue path, then click *OK*.

Setting the Messaging Agent Log Path

During installation, the default Messaging Agent log path is created in `/var/opt/novell/log/messenger/ma`, but you need the log file directory to be located on the Messenger partition.

1. Choose a directory for the Messaging Agent log files and create that directory on the Messenger partition.
2. In ConsoleOne, browse to and select the Novell Messenger Service object (MessengerService), then click *Messenger Server*.
3. Right-click the Messaging Agent object, then click *Properties*.
4. Click *Agent* > *Log Settings*.
5. In the *Log Files Path* field, specify the Messaging Agent log path, then click *OK*.

Setting the Archive Agent Log Path

During installation, the default Archive Agent log path is created in `/var/opt/novell/log/messenger/aa`, but you need the log file directory to be located on the Messenger partition.

1. Choose a directory for the Archive Agent queue and create that directory on the Messenger partition.
2. In ConsoleOne, browse to and select the Novell Messenger Service object (MessengerService), then click *Messenger Server*.
3. Right-click the Archive Agent object, then click *Properties*.
4. Click *Agent* > *Log Settings*.
5. In the *Log Files Path* field, specify the Archive Agent log path, then click *OK*.

11.2.3 Configuring the Messenger Cluster Resource to Load and Unload the Linux Messenger Agents

The properties of the Messenger Cluster Resource object define how the Messenger partition functions within the cluster, how the Messenger agents are loaded and unloaded, and how failover and failback situations are handled.

- "Modifying the Cluster Load Script for the Linux Messenger Agents" on page 97
- "Modifying the Cluster Resource Unload Script for the Linux Messenger Agents" on page 99
- "Setting the Failover List and Policies for the Linux Messenger Agents" on page 101

Modifying the Cluster Load Script for the Linux Messenger Agents

To set up the load script in iManager:

1. Expand *Clusters*, then click *Cluster Options*.
2. In the *Cluster* field, browse to the Cluster object where the Messenger cluster resource is located.
3. Click the Cluster object to display the cluster resources that belong to the cluster.
4. Select the Messenger cluster resource that you created when you set up the Messenger partition, then click *Details*.
5. Click *Scripts* > *Load Script*.

6 (Conditional) If this is an NSS volume or a shared pool, use a load script similar to the following example, depending on the configuration of your cluster and nodes:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

# mount filesystem
exit_on_error ncpcon mount /opt="noatime,nodiratime" volume_name=volume_ID

# add IP address
exit_on_error add_secondary_ipaddress messenger_partition_ip_address

# start service
exit_on_error /etc/init.d/novell-nmma start
exit_on_error /etc/init.d/novell-nmaa start

# return status
exit 0
```

6a In the `mount filesystem` section, specify the volume name and volume ID of the Messenger partition that you are clustering (System Clustering Worksheet item 5).

6b In the `add ip address` section, specify the secondary IP address of the Messenger partition (Messenger Clustering Worksheet item 6 or Messenger Clustering Worksheet item 7).

6c In the `start service` section, provide the commands to start the Messaging Agent first, following by the Archive Agent.

7 (Conditional) If this is a traditional Linux volume, use an unload script similar to the following example, depending on the configuration of your cluster and nodes:

```
#! /bin/bash
. /opt/novell/ncs/lib/ncsfunc

# define IP address
RESOURCE_IP=messenger_partition_ip_address

# define filesystem type
MOUNT_FS=filesystem

# define device (if using EVMS)
exit_on_error evms -f /var/opt/novell/ncs/ContainerActivate -rl
                                                    Share `uname -n`
MOUNT_DEV=/dev/evms/Share/dat
```

```
# define mount point
MOUNT_POINT=/mnt/mount_point_directory

# mount file system
exit_on_error mount -t $MOUNT_FS $MOUNT_DEV $MOUNT_POINT -o noatime,nodiratime

# add IP address
exit_on_error add_secondary_ipaddress $RESOURCE_IP

# start service
exit_on_error /etc/init.d/novell-nmma start
exit_on_error /etc/init.d/novell-nmaa start

exit 0
```

7a In the `define IP address` section, specify the secondary IP address of the Messenger partition (Messenger Clustering Worksheet item 6 or Messenger Clustering Worksheet item 7)

7b In the `define filesystem type` section, specify the filesystem type that is in use on nodes in the cluster (System Clustering Worksheet item 5).

7c In the `define mount point` section, specify the mount point directory in use for the nodes in the cluster (System Clustering Worksheet item 5).

8 Click *Apply* to save the load script.

Modifying the Cluster Resource Unload Script for the Linux Messenger Agents

The cluster resource unload script executes whenever the Messenger cluster resource goes offline.

1 On the iManager Cluster Resource Properties page of the Monitor Agent cluster resource, click *Scripts* > *Unload Script*.

2 (Conditional) If this is an NSS volume or a shared pool, use an unload script similar to the following example, depending on the configuration of your cluster and nodes:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

# request service stop
ignore_error /etc/init.d/novell-nmma stop
ignore_error /etc/init.d/novell-nmaa stop

# stop service otherwise
sleep 8
ignore_error pkill -fx "'/etc/init.d/novell-nmma stop'"
ignore_error pkill -fx "'/etc/init.d/novell-nmaa stop'"

# delete IP address
ignore_error del_secondary_ipaddress messenger_partition_ip_address

# unmount filesystem
exit_on_error umount /mnt/mount_point_directory

# return status
exit 0
```

2a In the `stop service otherwise` section, adjust the `sleep` command as needed so that the Messenger agents can shut down normally on your system without being inadvertently killed by the `pkill` command that follows.

2b In the `delete IP address` section, specify the secondary IP address of the Messenger partition.

2c In the `unmount filesystem` section, specify the mount point directory in use for the nodes in the cluster.

2d (Conditional) If you are running the GroupWise High Availability service (gwha), stop it before the script stops the Messenger agents, then start it again at the end of the unload script.

This prevents the GroupWise High Availability service from trying to restart the Messenger agents while the script is trying to stop them.

Add the following section before the `request service stop` section:

```
# Temporarily disable the gwha service under xinetd
ignore_error /sbin/chkconfig -s gwha off
ignore_error kill -HUP `pidof xinetd`
```

Add the following section before the `return status` section:

```
# Restart the gwha service under xinetd
ignore_error /sbin/chkconfig -s gwha on
ignore_error kill -HUP `pidof xinetd`
```

3 (Conditional) If this is a traditional Linux volume, use an unload script similar to the following example, depending on the configuration of your cluster and nodes:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

# request service stop
ignore_error /etc/init.d/novell-nmma stop
ignore_error /etc/init.d/novell-nmaa stop

# stop service otherwise
sleep 8
ignore_error pkill -fx "'/etc/init.d/novell-nmma stop'"
ignore_error pkill -fx "'/etc/init.d/novell-nmaa stop'"

# define IP address
RESOURCE_IP=messenger_partition_ip_address

# define mount point
MOUNT_POINT=/mnt/mount_point_directory

# delete IP address
ignore_error del_secondary_ipaddress $RESOURCE_IP

# umount filesystem
exit_on_error umount $MOUNT_POINT

# return status
exit 0
```

3a In the `stop service otherwise` section, adjust the `sleep` command as needed so that the Messenger agents can shut down normally on your system without being inadvertently killed by the `pkill` command that follows.

3b In the `define IP address` section, specify the secondary IP address of the Messenger partition.

3c In the `define mount point` section, specify the mount point directory in use for the nodes in the cluster.

3d (Conditional) If you are running the GroupWise High Availability service (gwha), stop it before the script stops the Messenger agents, then start it again at the end of the unload script.

This prevents the GroupWise High Availability service from trying to restart the Messenger agents while the script is trying to stop them.

Add the following section before the `request service stop` section:

```
# Temporarily disable the gwha service under xinetd
ignore_error /sbin/chkconfig -s gwha off
ignore_error kill -HUP `pidof xinetd`
```

Add the following section before the `return status` section:

```
# Restart the gwha service under xinetd
ignore_error /sbin/chkconfig -s gwha on
ignore_error kill -HUP `pidof xinetd`
```

4 Click *Apply* to save the unload script.

Setting the Failover List and Policies for the Linux Messenger Agents

1 On the iManager Cluster Resource Properties page of the Messenger cluster resource, click *General*.

The default policy settings are often appropriate. By default, a cluster resource:

- Fails over automatically if the node it is running on fails
- Starts automatically on the next node in its failover list
- Continues running at its failover location, even after its most preferred node is again available

If you are considering changing these defaults, see the *Novell Cluster Services Administration Guide for Linux* for your version of OES Linux (http://www.novell.com/documentation/oes.html).

2 Under *Preferred Nodes*, arrange the nodes in the cluster into the desired failover list for the Messenger agents (under Messenger Clustering Worksheet item 3).

3 Click *OK*.

11.3 Testing Your Messenger System in a Linux Cluster

After you have configured the Messenger cluster resource, you can test the load and unload scripts by bringing the Messenger cluster resource online and taking it offline again.

1. In iManager, expand *Clusters*, then click *Cluster Manager*.
2. Browse to the Cluster object to display the current cluster state.

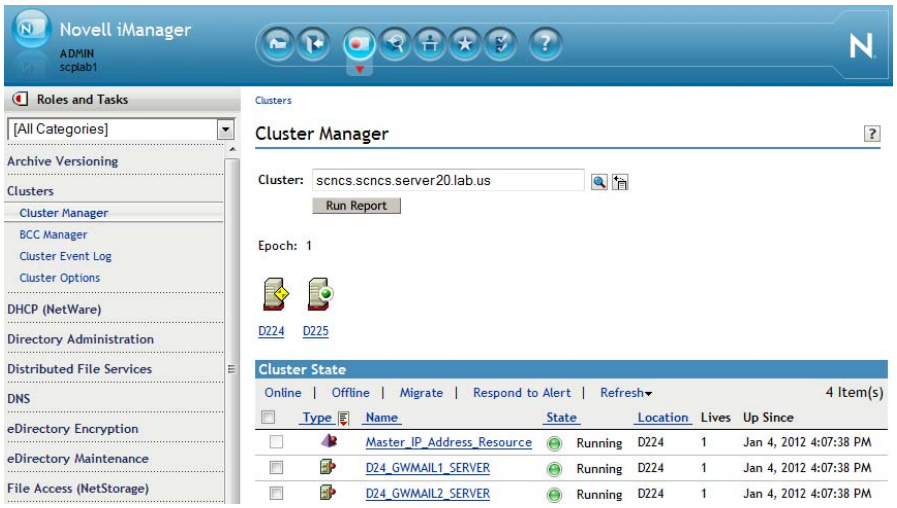

3. (Conditional) If the new Messenger cluster resource shows *Offline* in the *State* column, click the new Messenger cluster resource, then click *Online*.

   After a moment, the Messenger cluster resource displays *Running* in the *State* column.
4. At the server where the Messenger agents are starting, use the following commands to see that the Messenger agents have started:

   ```
   /etc/init.d/novell-nmma status
   /etc/init.d/novell-nmaa status
   ```
5. Select the new Messenger cluster resource, then click *Offline*.

   The *State* column for the Messenger cluster resource returns to *Offline*.
6. Use the same command that you used in Step 4 to verify that the Messenger agents have stopped.
7. Repeat Step 3 whenever you are ready to bring the new Messenger cluster resource online permanently.
8. Continue with Managing Your Messenger System in a Linux Cluster.

11.4 Managing Your Messenger System in a Linux Cluster

If the node where your Messenger system is running goes down, it fails over to the next node in its failover list. Messenger clients reconnect automatically as soon as the Messaging Agent restarts on the next node. Users who are actively carrying on conversations notice the interruption, but do not need to do anything to reestablish their conversation when the Messaging Agent is up and running again.

In comparison to failover, migration typically takes longer because the Messaging Agent methodically terminates its thread as part of its normal shutdown procedure.

11.5 Messenger Clustering Worksheet

Item	Explanation
1) eDirectory Tree for Cluster:	Record the eDirectory tree where you created the Novell Cluster object when you installed Novell Cluster Services. For more information, see Section 2.1, “Installing Novell Cluster Services on Linux,” on page 16.
2) Cluster Name: Master IP Address:	Record the name of the name of the Cluster object where your Messenger system will be located. Also record the master IP address of the cluster. For more information, see Section 2.1, “Installing Novell Cluster Services on Linux,” on page 16.
3) Cluster Context:	Record the full context where you created the Cluster object. For more information, see Section 2.1, “Installing Novell Cluster Services on Linux,” on page 16.
4) Nodes in Cluster ◆ FIle system type ◆ Device name ◆ Mount point	List the nodes that are part of the cluster that will include Messenger. Also list technical information, including file system type (`reiserfs`, `ext3`, and so on), device name (`sda2`, `hda1`, and so on), and mount point directory (`/mnt`, `/mail`, and so on) in use on the nodes the cluster. You need this information as you create load and unload scripts for the Messenger agents. For more information, see Section 2.1, “Installing Novell Cluster Services on Linux,” on page 16.
5) Shared Partitions in Cluster:	List the shared partitions that are available for use in your Messenger system. For more information, see Section 2.1, “Installing Novell Cluster Services on Linux,” on page 16.
6) Messenger Partition for Messaging Agent: Secondary IP address: Use Same Partition for Archive Agent? ◆ Yes ◆ No	The Messaging Agent software is installed on each node in its failover list but it does use a shared partition to store its log files, temporary files, and queue directories. Specify the name of the shared partition in the your cluster that the Messaging Agent can use for these purposes. also specify the secondary IP address of that shared partition. For more information, see Section 11.1.2, “Selecting the Messenger Partition and Secondary IP Address,” on page 92.

Item	Explanation
7) Messenger Partition for Archive Agent? Secondary IP Address:	In addition to the data storage needs of the Messaging Agent, the Archive Agent can be configured to store all instant message conversations. It is possible that this might consume a large quantity of disk space. If so, you can choose to use a separate shared partition for the Archive Agent. For more information, see Section 11.1.5, "Planning the Linux Messenger Agent Installation," on page 93.
8) Failover List for Messenger Agents:	List other nodes in the cluster where the Messenger agents can fail over. You might want to have the same nodes on the both agents' lists or have separate lists for each agent. It depends on the loads that each agent will be carrying. For more information, see Section 11.1.3, "Determining an Appropriate Failover List for the Linux Messenger Agents," on page 92.
9) Mount Point for Shared Storage:	Specify the mount point directory where the shared resource is mounted to the cluster node where the Messenger Agents run. For more information, see Section 11.1.4, "Determining Cluster Resource Information for the Linux Messenger Agents," on page 92.

11.6 Messenger Clustering Quick Checklist

- ❒ Plan your clustered Messenger system.

  See Section 11.1, "Planning Your Messenger System in a Linux Cluster," on page 91.

- ❒ Create your Messenger system and install the Messenger agents.

  See Section 11.2.1, "Creating Your Messenger System and Installing the Linux Messenger Agents," on page 93.

- ❒ If you use NSS volumes in your cluster, configure the Messenger agents so that they run as a non-root user.

  See "Setting Up Non-root Access on NSS Volumes on Each Node" on page 94.

- ❒ In ConsoleOne, change the locations of various Messenger files from their default locations on local nodes to the Messenger partition that is always available no matter what node the Messenger agents are running on.

  See Section 11.2.2, "Changing Messenger Paths to Locations on the Messenger Partition," on page 95.

- ❒ Modify the Messenger agents cluster resource load script.

  See "Modifying the Cluster Load Script for the Linux Messenger Agents" on page 97.

- ❒ Modify the Messenger agents cluster resource unload script.

  See "Modifying the Cluster Resource Unload Script for the Linux Messenger Agents" on page 99.

- ❒ Set up the Messenger agents failover list and policies.

  See "Setting the Failover List and Policies for the Linux Messenger Agents" on page 101.
- ❒ Test your clustered Messenger system.

  See Section 11.3, "Testing Your Messenger System in a Linux Cluster," on page 102.

II Novell Vibe

Before installing Novell Vibe 3, you should thoroughly review the documentation provided at the Novell Vibe 3 documentation Web site (http://www.novell.com/documentation/vibe32). These guides provide detailed product installation and configuration instructions, but they do not include specific instructions for integrating Novell Vibe with eDirectory or GroupWise. This section of the *GroupWise 2012 Interoperability Guide* supplies these product-specific instructions.

12 Configuring GroupWise for Use with Novell Vibe

When you install Novell Vibe with eDirectory and GroupWise, some configuration steps are required to integrate the applications.

12.1 Understanding How Novell Vibe Interacts with eDirectory and GroupWise

When you install Novell Vibe in an environment where eDirectory and GroupWise are already set up, the products interact in the following ways:

- Vibe can use eDirectory for LDAP authentication of Vibe users. This means that you do not need to create Vibe users manually. Vibe can create its user accounts based on the users that already exist in eDirectory.
- Vibe can use GroupWise as its integrated email system. This means that email messages sent from the Vibe site are delivered to GroupWise mailboxes. It also means that GroupWise users can post items to Vibe folders by sending email messages to Vibe folders.
- Vibe information can be displayed in the GroupWise Windows client. Starting in GroupWise 8.0.2, you can drag and drop GroupWise items into Vibe folders in the GroupWise Folder List to post items to the corresponding folders in your Vibe site. You can also use the GroupWise Find feature to search your Vibe site.

12.2 Using eDirectory as the Vibe LDAP Directory

For instructions, see the following sections of the *Novell Vibe 3.2 Installation Guide*:

- "Gathering Directory Services Information" in "Planning a Basic Vibe Installation"
- "Adding Users to Your Vibe Site" in "Basic Installation"

12.3 Using GroupWise as the Vibe Email System

For setup instructions, see the following sections of the *Novell Vibe 3.2 Installation Guide*:

- "Gathering Outbound E-Mail Information" in "Planning a Basic Vibe Installation"
- "Enabling Inbound E-Mail" in "Planning a Basic Vibe Installation"

See also the following section of the *Novell Vibe 3.2 Administration Guide*:

- "Configuring E-Mail Integration" in "Site Setup"

For basic email usage instructions, see the following sections of the *Novell Vibe 3.2 User Guide*:

- "Sending E-Mail to Team Members and Announcing the Workspace after Its Creation" in "Managing and Using Workspaces"
- "Subscribing to E-Mail Notifications from a Folder", "Setting Up a Folder to Receive Entries Via E-Mail" and "Adding Entries to a Folder Via E-Mail" in "Managing and Using Folders"
- "Sending E-Mail from within Vibe" in "Connecting With Your Co-Workers"

See also the following sections of the *Novell Vibe 3.2 Advanced User Guide*:

- "Enabling Folders to Receive Entries through E-Mail" and "Configuring Folders to Send E-Mail Notifications to Other Users" in "Managing Folders"
- "Sending E-Mail Notifications" in "Creating and Managing Workflows"
- "E-Mailing Files and Attachments to the Vibe Site When You Are Over Your Quota" in "Managing Your Data Quota"
- "Sending E-Mail" in "Using Vibe on Your Mobile Phone"

12.4 Enabling GroupWise/Vibe Integration for GroupWise Windows Client Users

Before you can integrate GroupWise and Vibe, your Vibe site must be set up, as described in the *Novell Vibe 3.2 Installation Guide*.

1. In ConsoleOne, browse to and select a Domain object, Post Office object, or User object where you want to make Vibe available to GroupWise Windows client users.
2. Click *Tools* > *GroupWise Utilities* > *Client Options*.
3. Click *Environment* > *Vibe*.

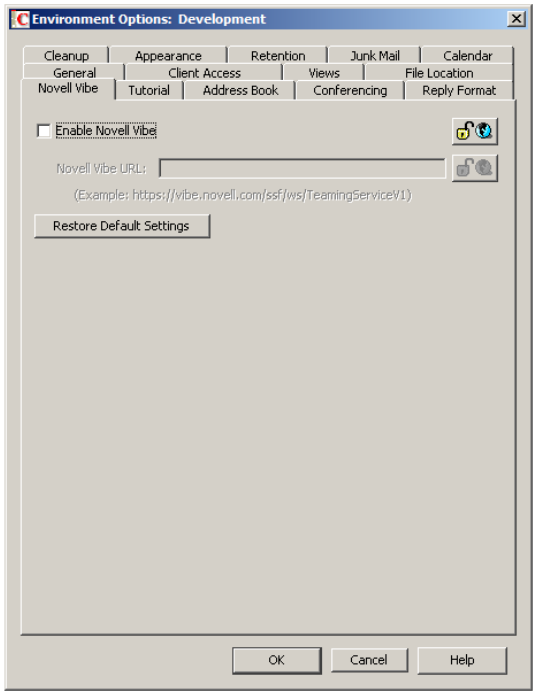

4 Select *Enable Novell Vibe*.

5 Provide the Vibe URL:

5a Specify the fully qualified IP address or DNS hostname of the Vibe server:

`vibe_server.domain`

For example:

`vibe.yourcompanyname.com`

ConsoleOne provides the rest of the default Vibe URL, which uses a secure HTTPS connection, assumes the default port number, and includes the default location for the Vibe Web service that communicates with other applications:

`https://vibe_server.domain/ssf/ws/TeamingServiceV1`

5b (Conditional) If you want to use HTTP instead of HTTPS, include it in the *Novell Vibe URL* field, for example:

`http://vibe.yourcompanyname.com`

5c (Conditional) If Vibe is not configured with the default HTTPS port, include the port number after the hostname, for example:

`vibe.yourcompanyname.com:444`

5d (Conditional) If Vibe is not installed in the default location, include the path to `TeamingServiceV1`, for example:

`vibe.yourcompanyname.com/Web/Vibe/TeamingServiceV1`

6 Click *OK*.

IMPORTANT: In order for GroupWise users to take advantage of GroupWise/Vibe integration, they must provide their GroupWise email address in their Vibe profile.

---

13 Accessing Your Vibe Site from the GroupWise Windows Client

Before you can access the Vibe site from the GroupWise Windows client, you must add your GroupWise email address to your Vibe profile, as described in "Modifying Your Profile" in "Getting Started" in the *Novell Vibe 3.2 User Guide*.

The *GroupWise and Vibe Quick Start* (http://www.novell.com/documentation/groupwise2012/pdfdoc/gw2012_qs_vibe/gw2012_qs_vibe.pdf) describes the Vibe functionality that becomes available in the GroupWise Windows client as a result of the GroupWise/Vibe integration.

14 Streamlining Authentication to Vibe

You can implement single sign-on for use with Novell Vibe, so that users do not need to log in separately to GroupWise and Vibe. See the following sections in "Planning an Advanced Vibe Installation" in the *Novell Vibe 3.2 Installation Guide*:

- "Configuring Single Sign-On with Novell Access Manager"
- "Configuring Single Sign-On with Internet Information Services for Windows"

III Novell ZENworks

15 Using ZENworks Configuration Management to Distribute the GroupWise Windows Client

You can use the Configuration Management functionality in Novell ZENworks 10 to distribute the GroupWise Windows client to workstations.

IMPORTANT: This information assumes that you are familiar with ZENworks 10.3 Configuration Management. For background information, or for help completing the ZENworks tasks outlined in the steps below, see the ZENworks Configuration Management documentation at the Novell ZENworks Documentation Web site (http://www.novell.com/documentation/zcm10).

For a video demonstration of this feature, see:

- Novell GroupWise 2012 Deployed Using ZCM 11 SP1 Part 1 (http://www.youtube.com/watch?v=V9Wa2ETf7ag&list=UUZHLAuzLykBK09uEJesc2og&index=2&feature=plcp)
- Novell GroupWise 2012 Deployed Using ZCM 11 SP1 Part 2 (http://www.youtube.com/watch?v=0XWzqxOJpos&list=UUZHLAuzLykBK09uEJesc2og&index=1&feature=plcp)

15.1 Creating a GroupWise MST Transform File

An MST (Media Stream Transforms) file enables you to customize the functioning of an MSI (Microsoft Installer) file that has been provided with a software program. The GWTuner utility allows you to customize your GroupWise MSI installation. GWTuner creates an MST transform file named `groupwise.mst`, which you can use when performing an MSI install with ZENworks. You must have write access to the software distribution directory to use the GroupWise 2012 Tuner application.

1. (Conditional) If you have not already done so, create a GroupWise software distribution directory on the Windows server from which you want to distribute the GroupWise Windows client software.

   For instructions, see "Creating a Software Distribution Directory" in "System" in the *GroupWise 2012 Administration Guide*. The software distribution directory must include at least the client and administration components of GroupWise.
2. On the Windows server, ensure that you have write access to the software distribution directory.
3. Run the following program:

software_distribution_directory\admin\utility\tools\gwtuner.exe

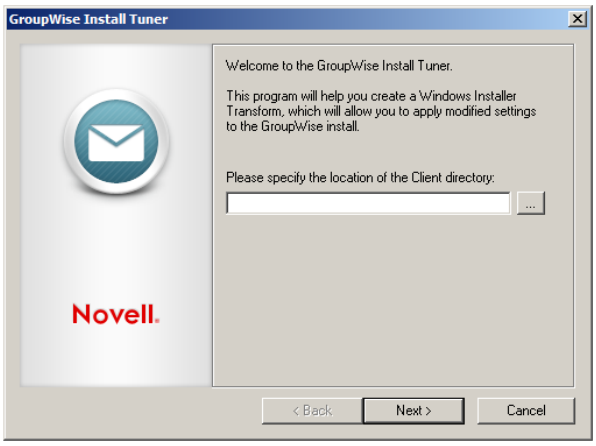

4 Browse to and select the following directory, then click *Next*.

software_distribution_directory\client

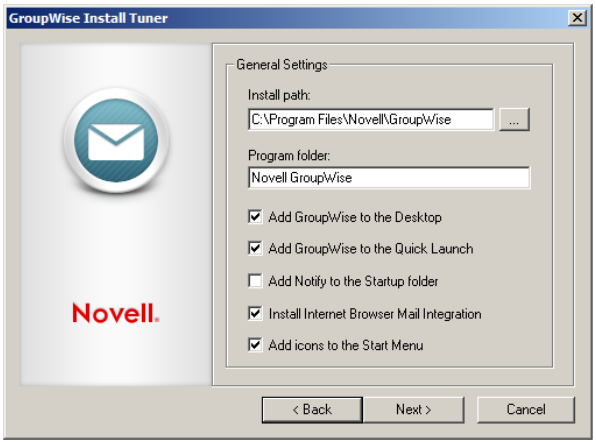

5 In the *Install path* field, specify where you want to install the Windows client software on users' workstations.

The typical location varies depending on the architecture of the workstations:

32-bit architecture: `C:\Program Files\Novell\GroupWise`

64-bit architecture: `C:\Program Files (x86)\Novell\GroupWise`

If all of the workstations in your environment have the same architecture, specify the appropriate location for that architecture. If you have both 32-bit and 64-bit workstations in your environment, you can specify either path in GWTuner. In ZENworks Configuration Management, you will create a separate launch action for each architecture.

6 Select Windows client installation options as needed.

For information about the Windows client installation options, see "Using GroupWise Windows Client Custom Installation Options" in "System" in the *GroupWise 2012 Administration Guide*.

IMPORTANT: When you use ZENworks Configuration Management to install the Windows client, the following selections are recommended in GWTuner:

6a Deselect *Add GroupWise to the Desktop*.

6b Deselect *Add GroupWise to the Quick Launch*.

6c Deselect *Add Icons to the Start Menu*.

It is preferable to set these options later in ZENworks Configuration Management, rather than in GWTuner.

6d Click *Next*.

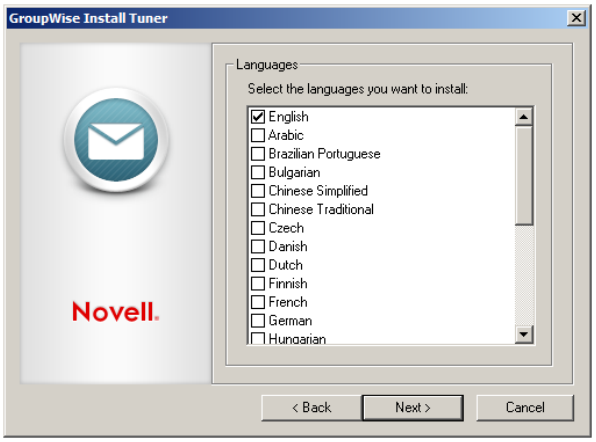

7 Select all the languages that you want to install on users' workstations. then click *Next*.

8 Select the default Windows client startup language, then click *Next*.

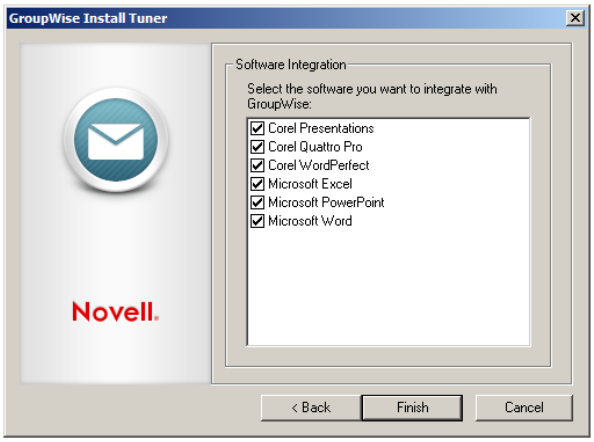

9 Unless you are actively using GroupWise Document Management Services (DMS), deselect all of the integrations, then click *Finish*.

10 Click *OK* to exit GWTuner.

The following file is created:

`software_distribution_directory\client\win32\groupwise.mst`

11 Continue with Creating a Bundle for the GroupWise Windows Client Software.

15.2 Creating a Bundle for the GroupWise Windows Client Software

ZENworks Configuration Management deploys software as bundles that include all the files and instructions required to successfully install software on users' workstations.

- Section 15.2.1, "Creating a New GroupWise Bundle," on page 122
- Section 15.2.2, "Adding Installation Actions to the GroupWise Bundle," on page 125
- Section 15.2.3, "Adding a Launch Action to the GroupWise Bundle," on page 135
- Section 15.2.4, "Displaying GroupWise Bundle Activity on User Workstations," on page 137
- Section 15.2.5, "Publishing the GroupWise Bundle," on page 138

15.2.1 Creating a New GroupWise Bundle

To create a ZENworks Configuration Management bundle for the GroupWise Windows client software:

1 In ZENworks Control Center, click *Bundles*, then click *New* > *Bundle*.

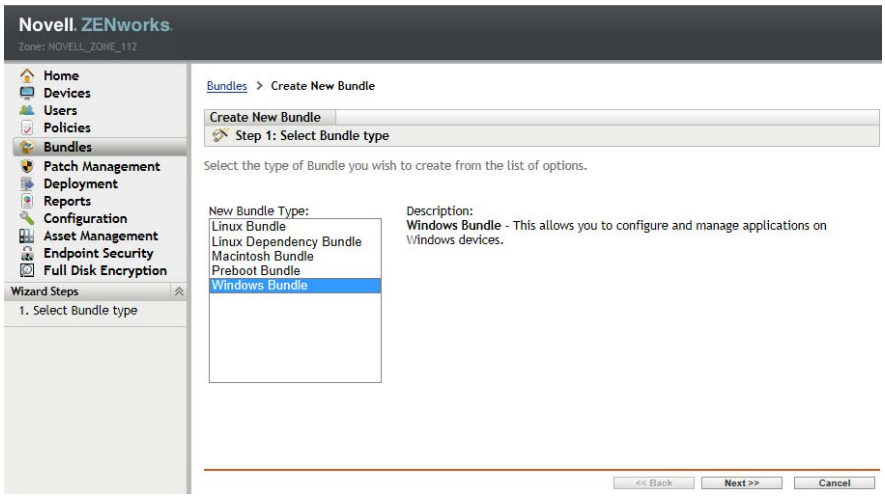

2 Select *Windows Bundle,* then click *Next*.

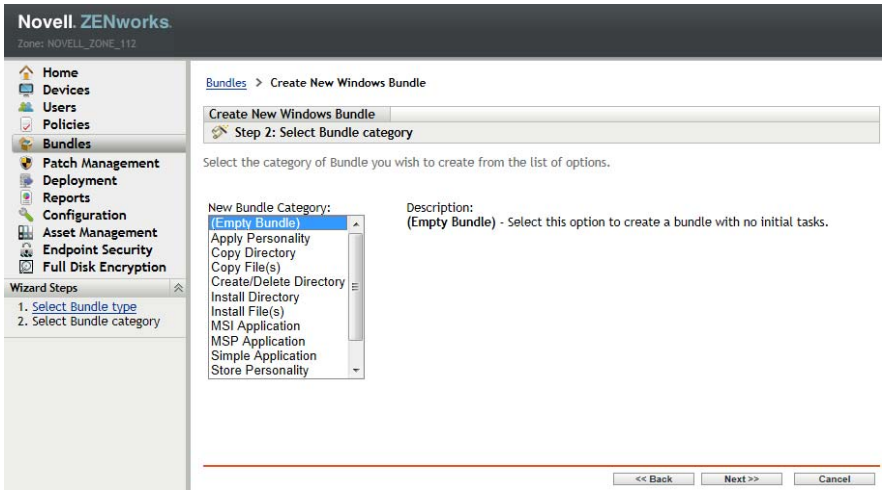

3 Click *Next* to accept the default of *(Empty Bundle)*.

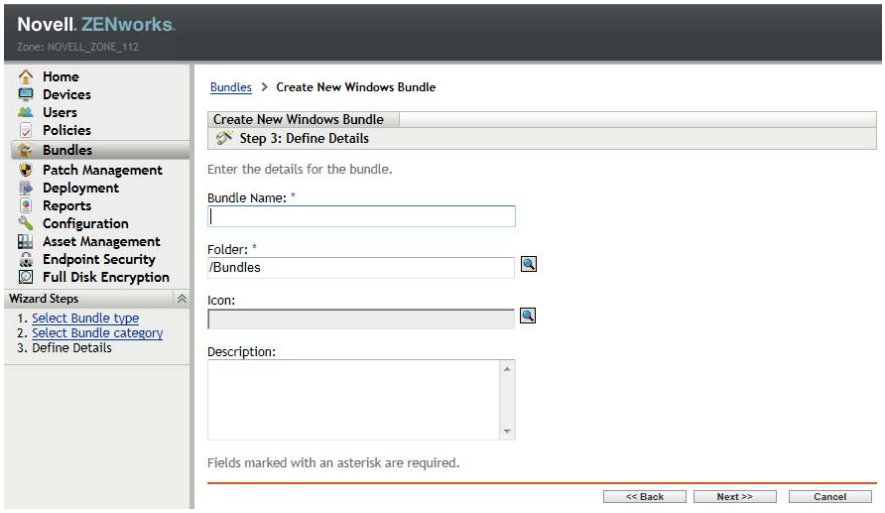

4 Define the new bundle:

4a In the *Bundle Name* field, specify a unique name for the bundle, such as GW2012.

4b In the *Icon* field, browse to and select the following file:

`software_distribution_directory\client\win32\grpwise12.ico`

4c Click *Next* to display the bundle summary.

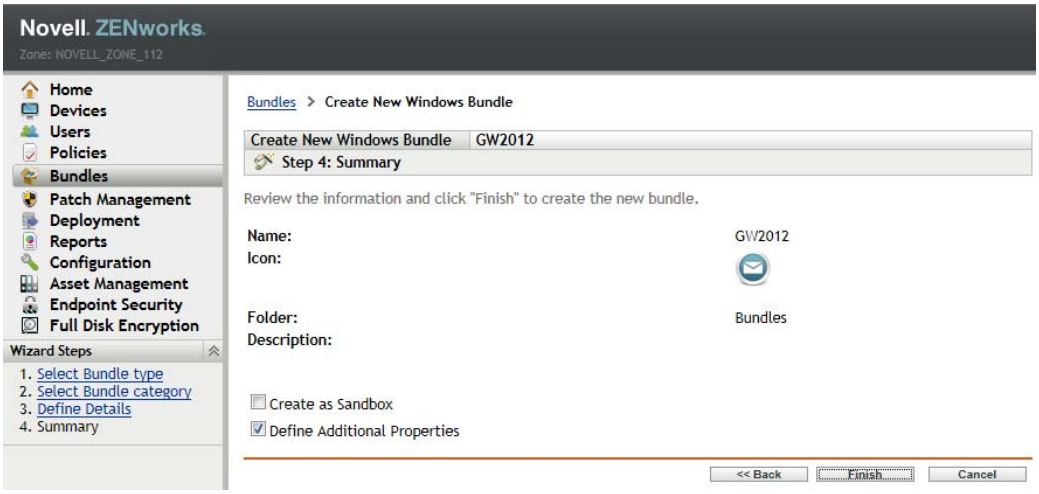

4d Select *Create as Sandbox* so that you can set up all the components of the bundle before publishing the bundle, then click *Finish*.

Define Additional Properties is selected by default, which enables you to conveniently set up all the bundle components.

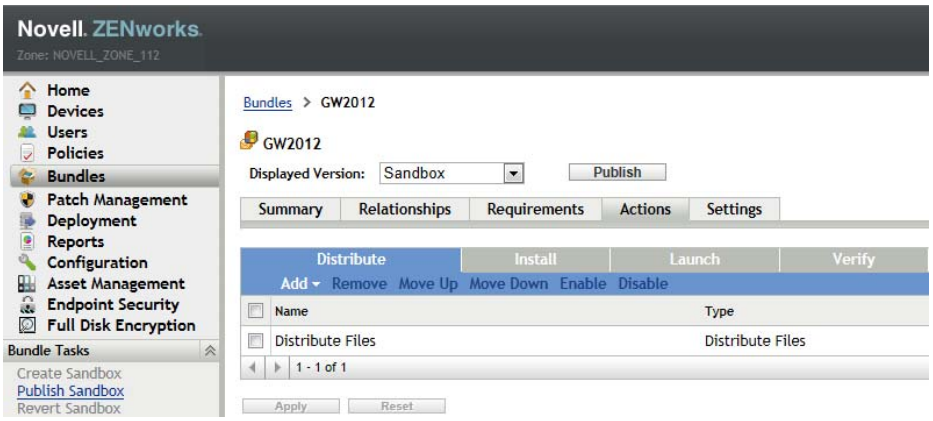

5 Click *Install* to create installation actions for the GroupWise bundle:

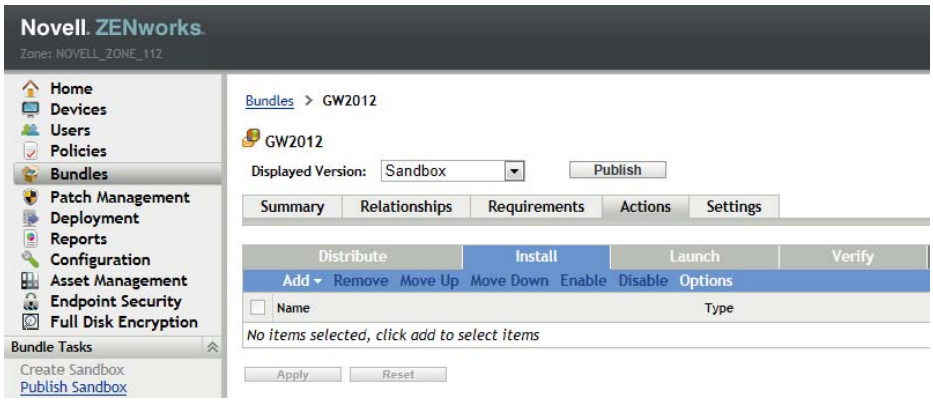

6 Continue with Adding Installation Actions to the GroupWise Bundle.

15.2.2 Adding Installation Actions to the GroupWise Bundle

Several installation actions are required in order to install the GroupWise Windows client software.

- "Creating an Install Files Action to Copy the GroupWise Windows Client Software to Temporary Location" on page 125
- "Creating a Run Script Action" on page 127
- "Creating an Install Files Action to Add the GroupWise MST Transform File to the GroupWise Windows Client Installation" on page 129
- "Creating a Launch Executable Action" on page 132

Creating an Install Files Action to Copy the GroupWise Windows Client Software to Temporary Location

1 On the *Install* tab, click *Add* > *Install File(s)* to add an Install File(s) action.

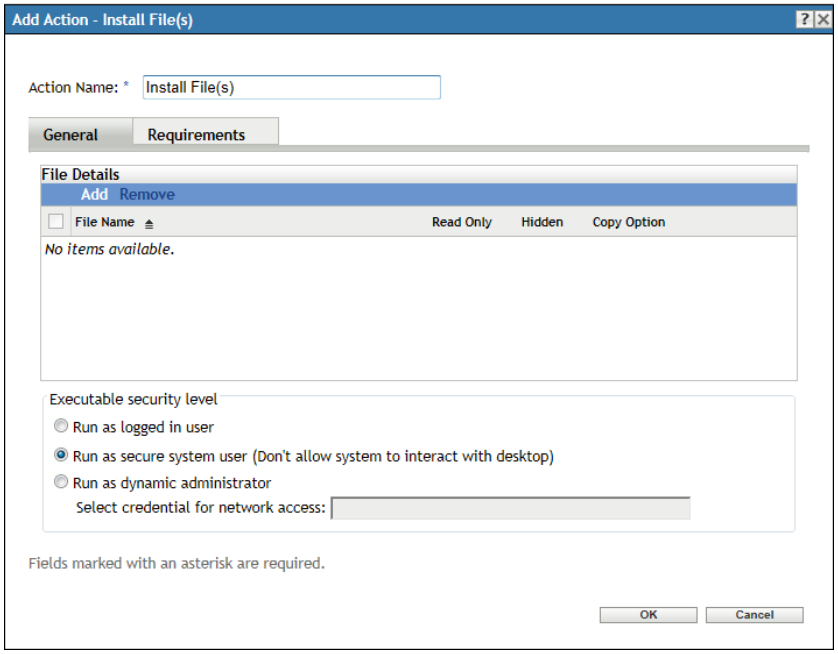

The *Install File(s)* option copies the files that you select to a location that you specify. The first action you create copies the GroupWise Windows software files to a temporary location on your local server.

2 In the *Action Name* field, customize the name to clearly identify the action, for example:

```
Install Files to Temporary Directory
```

3 Click *Add* to open the Select Files dialog box.

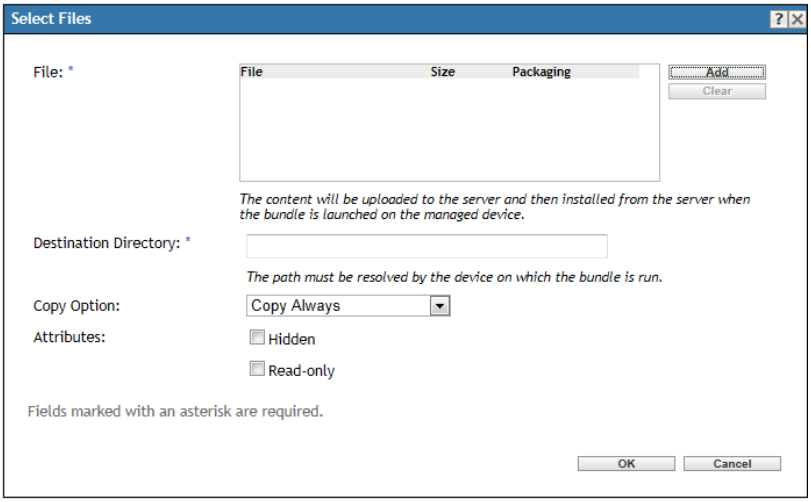

4 Click *Add*.

5 Browse to the following directory:

```
software_distribution_directory\admin\utility\setupip
```

6 Select the following files:

- `extract.bat`
- `setupip.fil`
- `setupip.xx` (where *xx* is a two-letter language code)

7 (Conditional) If your version of ZENworks Configuration Management provides the *Do not compress or encrypt uploaded content*, select this option.

The largest file, `setupip.fil`, is already compressed. The other files are small and would not benefit substantially from compression.

The uploaded files display in the *File* list.

8 In the *Destination Directory* field, specify the temporary directory where you want to create a copy of the GroupWise Windows client software, for example:

```
c:\gwclient
```

If the directory does not already exist, it is automatically created.

9 Click *OK* to close the Select Files dialog box and return to the Install File(s) dialog box with the selected files listed.

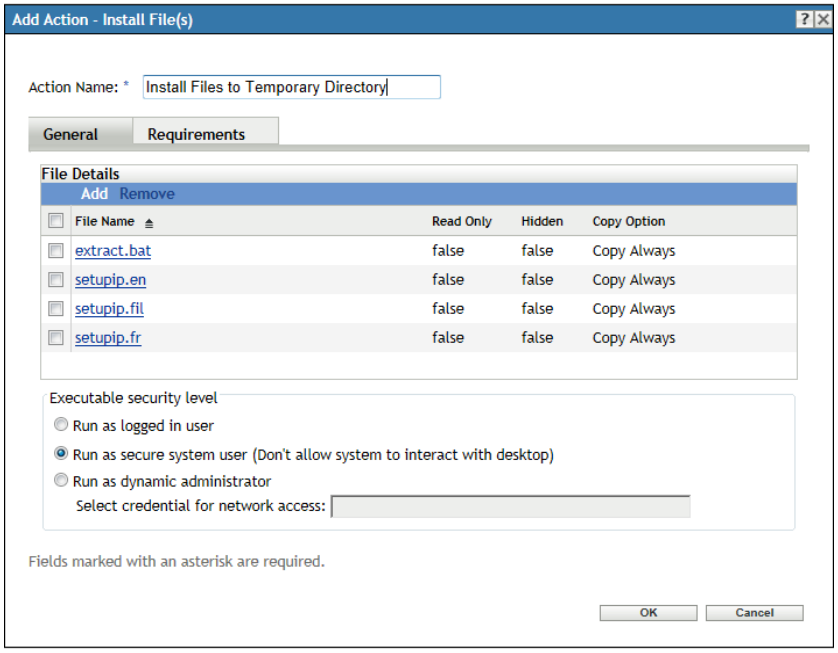

10 (Optional) Select *Run as dynamic administrator*.

The default of *Run as secure system user* is also appropriate for most systems.

11 Click *OK* to close the Install File(s) dialog box and return to the GroupWise Bundles page with the first installation task listed.

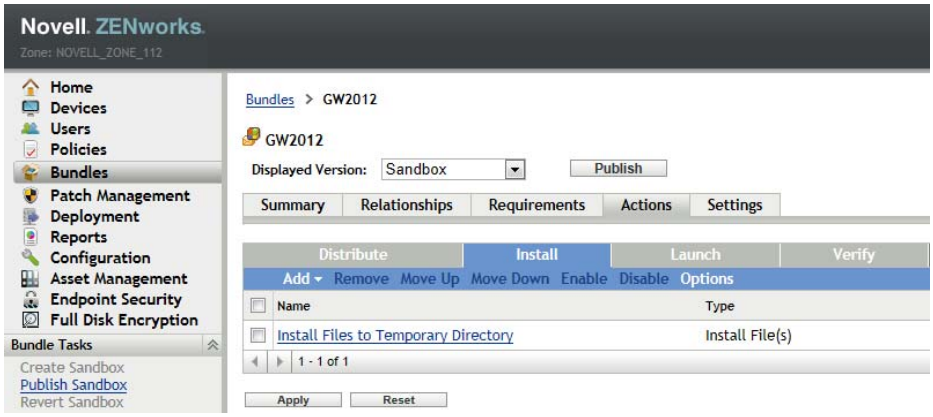

12 Click *Apply* to save the first installation action.

13 Continue with Creating a Run Script Action.

Creating a Run Script Action

1 On the *Install* tab, click *Add* > *Run Script* to add a Run Script action.

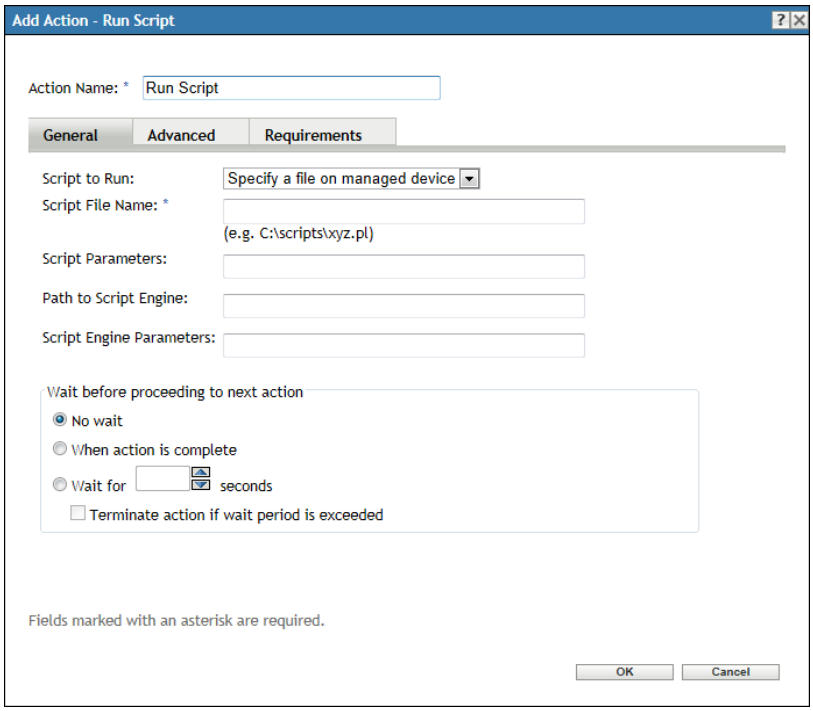

The *Run Script* option runs the script that you specify to extract all the Windows client software files from the `setupip.fil` file into a `win32` subdirectory under `gwclient`.

2 In the *Action Name* field, customize the name to clearly identify the action, for example:

```
Run Script to Extract GroupWise Client Software
```

3 In the *Script File Name* field, specify the full path to the batch file for extracting the GroupWise Windows client software, for example:

```
c:\gwclient\extract.bat
```

4 In the *Wait before proceeding to next action* box, select *When action is complete*.

5 Click *Advanced*.

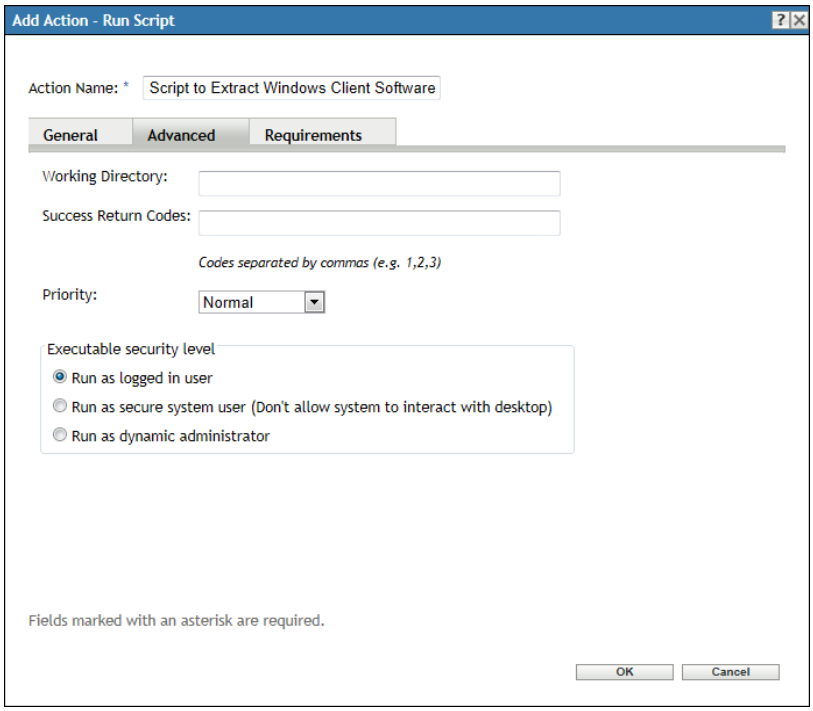

6 In the *Executable security level* box:

Select *Run as secure system user*.

or

Select *Run as dynamic administrator*.

7 Click *OK* to add this installation action to the list.

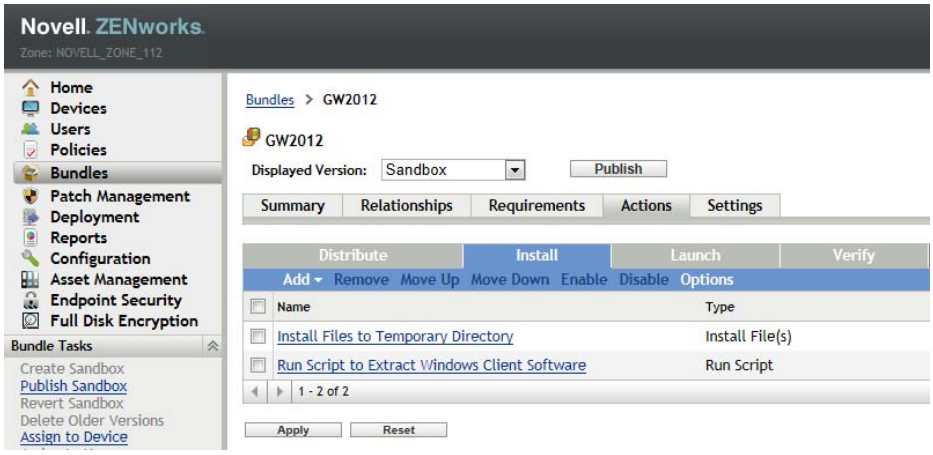

8 Click *Apply* to save this installation action.

9 Continue with Creating an Install Files Action to Add the GroupWise MST Transform File to the GroupWise Windows Client Installation.

Creating an Install Files Action to Add the GroupWise MST Transform File to the GroupWise Windows Client Installation

1 On the *Install* tab, click *Add* > *Install File(s)* to add an Install File(s) action.

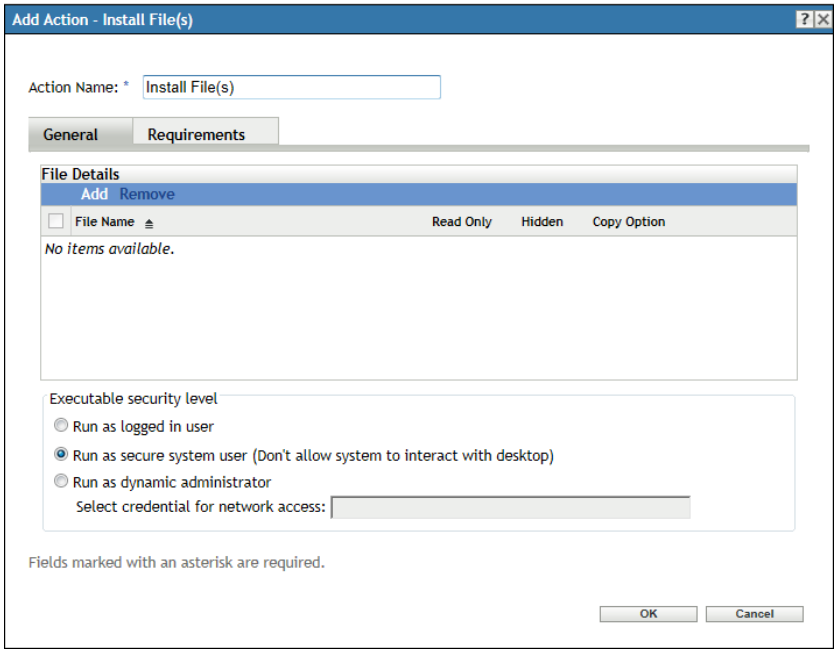

The *Install File(s)* option copies the GroupWise MST transform file to the temporary location on your workstation.

2 In the *Action Name* field, customize the name to clearly identify the action, for example:

```
Install File groupwise.mst
```

3 Click *Add* to open the Select Files dialog box.

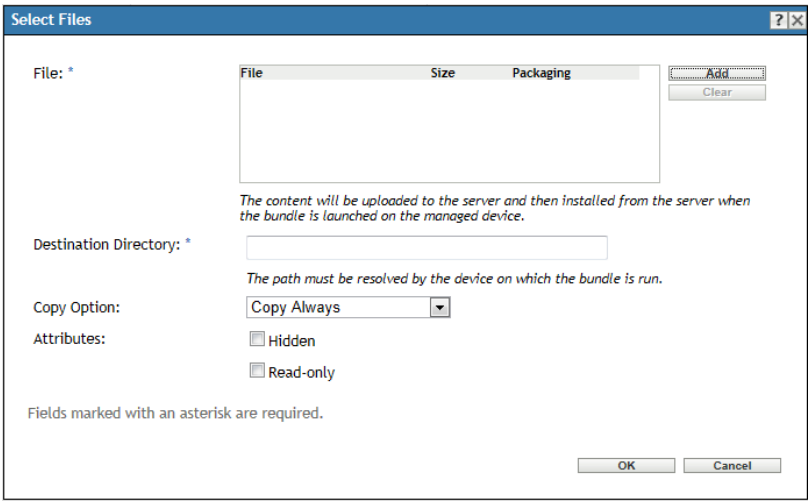

4 Click *Add*.

5 Browse to the following directory:

```
software_distribution_directory\client\win32
```

6 Select the `groupwise.mst` file.

7 (Conditional) If your version of ZENworks Configuration Management provides the *Do not compress or encrypt uploaded content*, select this option.

The `groupwise.mst` file is very small and does not benefit from compression.

The uploaded files display in the *File* list.

8 In the *Destination Directory* field, specify the `win32` subdirectory of the temporary directory where you created the copy of the GroupWise Windows client software, for example:

```
c:\gwclient\win32
```

9 Click *OK* to close the Select Files dialog box and return to the Install File(s) dialog box with the `groupwise.mst` file listed.

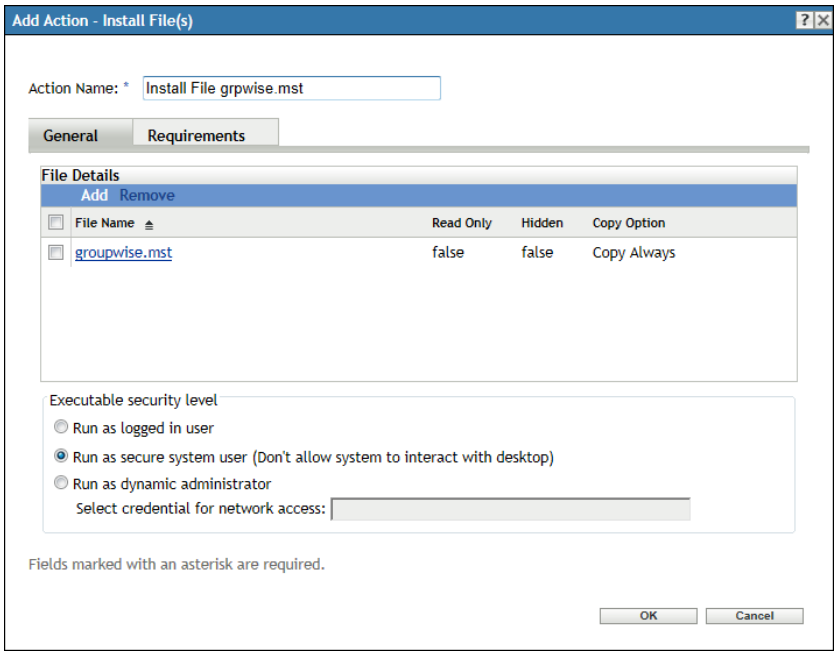

10 (Optional) Select *Run as dynamic administrator*.

The default of *Run as secure system user* is also appropriate for most systems.

11 Click *OK* to close the Install File(s) dialog box and return to the GroupWise Bundles page with this installation task listed.

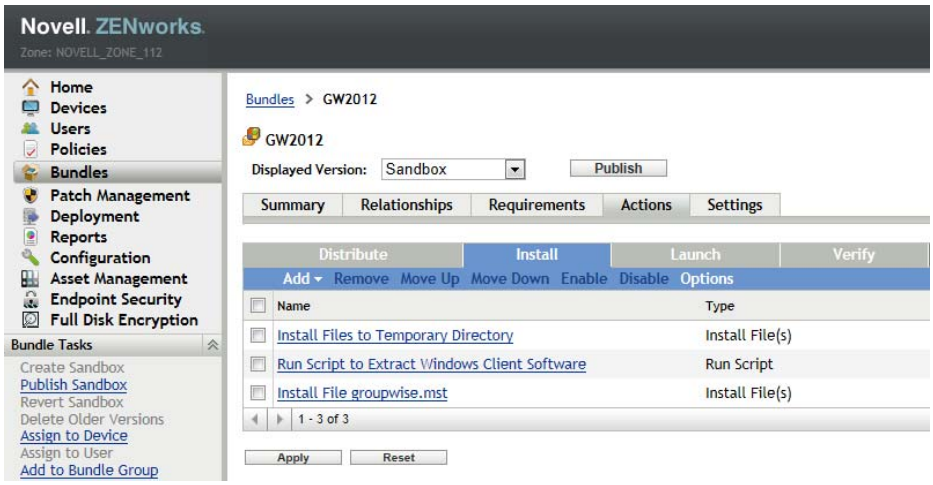

12 Click *Apply* to save this installation action.

13 Continue with Creating a Launch Executable Action.

Creating a Launch Executable Action

1 On the *Install* tab, click *Add* > *Launch Executable* to add a Launch Executable action.

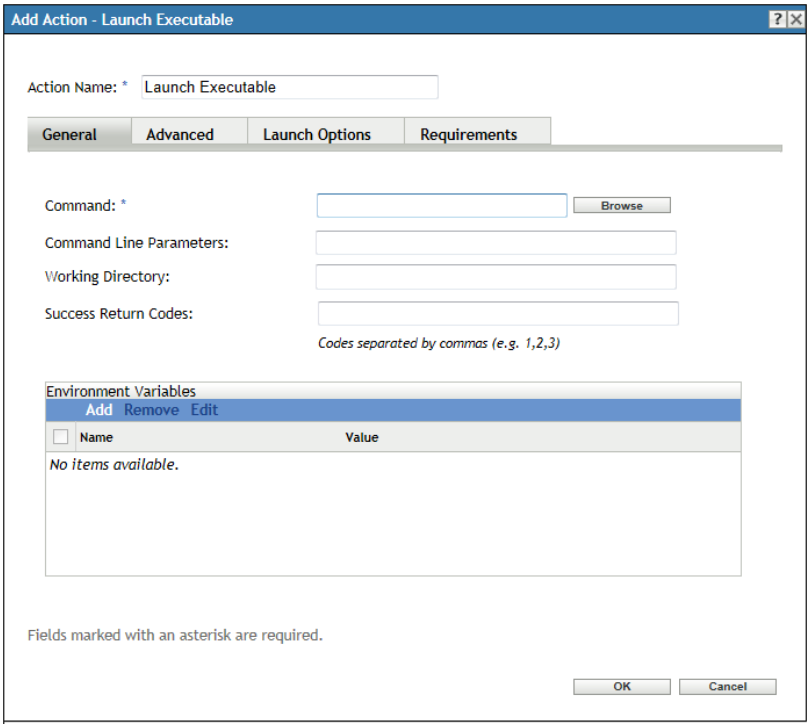

The *Launch Executable* option launches the `install.bat` file to install the GroupWise Windows client software on users' workstations.

2 In the *Action Name* field, customize the name to clearly identify the action, for example:

```
Launch Executable install.bat
```

3 In the *Command* field, specify the full path to the executable file that was extracted from the `setupip.fil` file, for example:

```
c:\gwclient\win32\install.bat
```

4 In the *Command Line Parameters* field, specify one of the following parameters:

- **/unattended:** Users see the window displayed by the `groupwise.msi` file while the GroupWise client software is being installed. The users do not need to do anything during the process, but they are informed that the installation is taking place.
- **/silent:** Users see nothing happening while the GroupWise client software is being installed, so they do not know that it has occurred.

The instructions that follow assume that `/silent` has been specified.

5 In the *Environment Variables* box:

5a Click *Add*.

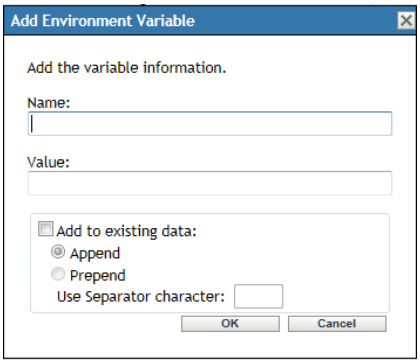

Two environment variables must be set in order to customize the interaction of the `groupwise.msi` file with the `groupwise.mst` file.

5b In the *Name* field, specify `GW_INST_TRANSFORM_FILE`.

This environment variable provides the name of the GroupWise MST transform file. The default location of the file is the same as the location of the MSI file, so you do not need to specify a full path for the file.

5c In the *Value* field, specify `groupwise.mst`, then click *OK* to return to the Launch Executable dialog box.

5d Click *Add* again to add the second environment variable.

5e In the *Name* field, specify `GW_INST_REMOVE_MSI`.

This environment variable allows for updates of previous versions of the GroupWise Windows client that were installed using the `groupwise.msi` file. For example, it allows the automatic removal of obsolete GroupWise icons.

5f In the *Value* field, specify `True`, then click *OK* to return to the Launch Executable dialog box.

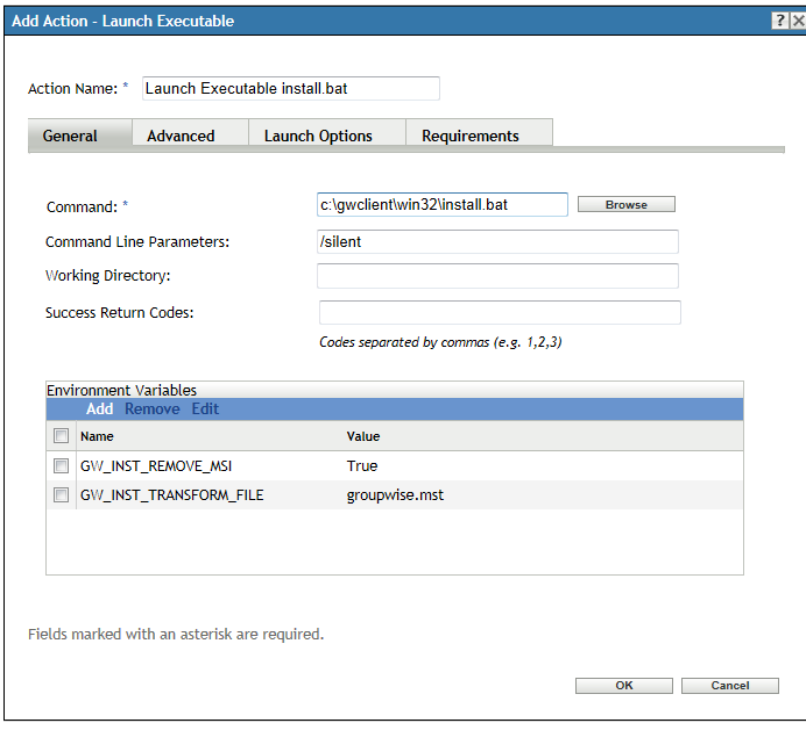

6 Click *Advanced*.

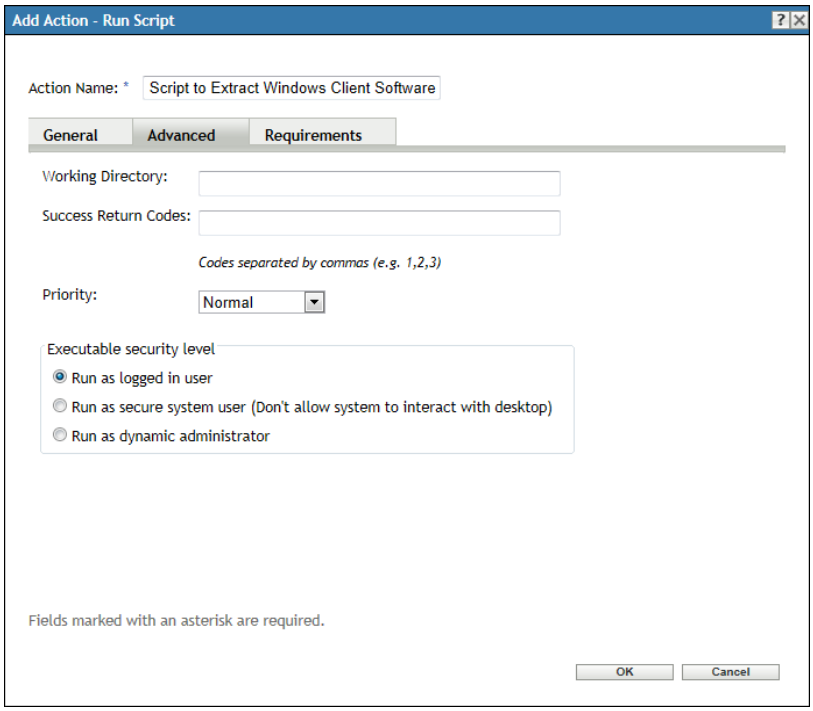

7 In the *Wait before proceeding to next action* box, select *When action is complete*.

8 In the *Executable security level* box:

Select *Run as secure system user*.

or

Select *Run as dynamic administrator*.

9 Click *OK* to close the Launch Executable dialog box and return to the GroupWise Bundles page with this installation task listed.

10 Click *Apply* to save this installation action.

11 (Optional) Add additional installation actions based on your preferences for how you want to install the GroupWise Windows client software.

12 Continue with Adding a Launch Action to the GroupWise Bundle.

15.2.3 Adding a Launch Action to the GroupWise Bundle

The launch action is the command to start the GroupWise Windows client on users workstations. It is useful for testing your GroupWise bundle. You might not want to use it after you have tested your GroupWise bundle.

1 On the GroupWise Bundle page, click *Launch*.

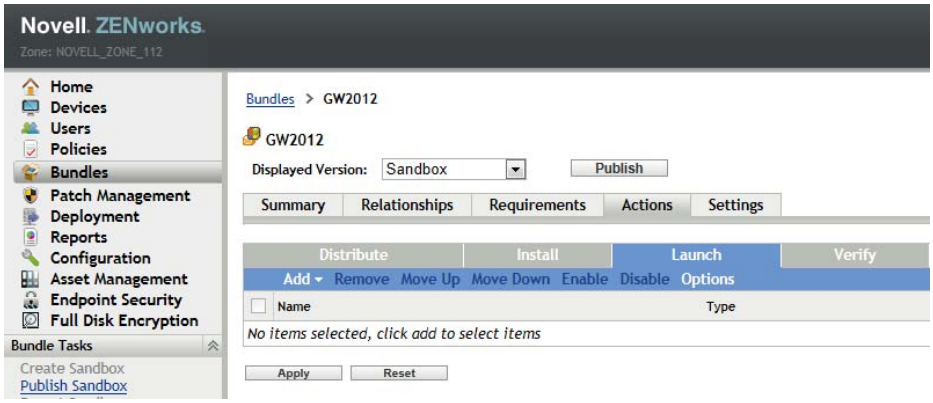

2 Click *Add* > *Launch Executable* to open the Launch Executable dialog box.

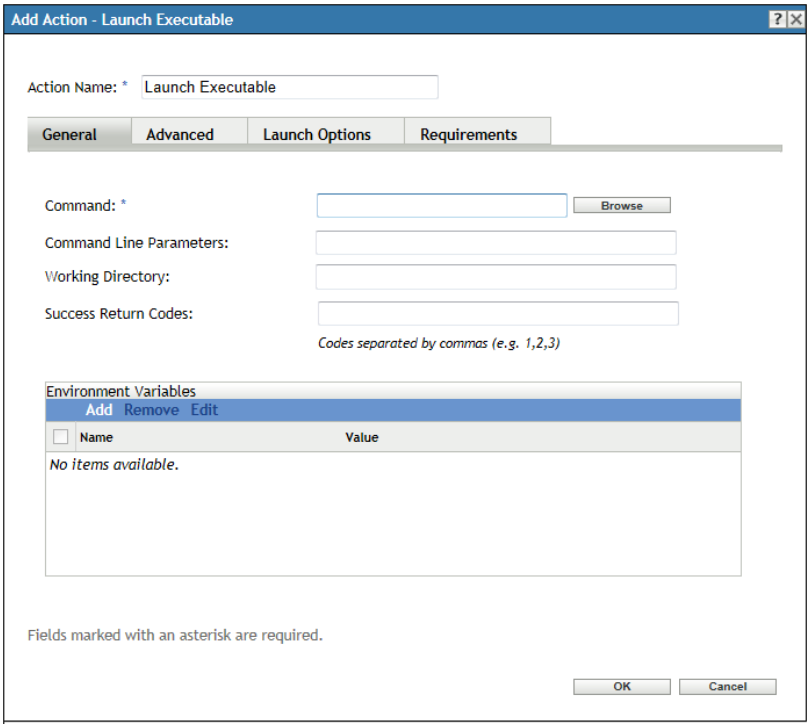

3 In the *Action Name* field, customize the name to clearly identify the action, for example:

```
Launch Executable grpwise.exe
```

4 In the *Command* field, specify the command to start the GroupWise Windows client.

The command depends on the architecture of the Windows workstation where the GroupWise Windows client is being installed:

32-bit architecture: `C:\Program Files\Novell\GroupWise\grpwise.exe`

64-bit architecture: `C:\Program Files (x86)\Novell\GroupWise\grpwise.exe`

For testing purposes, specify the command for the architecture of workstation where you plan to test the bundle.

5 (Conditional) If your environment includes both 32-bit and 64-bit workstations, define the system requirements for this launch action:

5a Click *Requirements* > *Add Filter*.

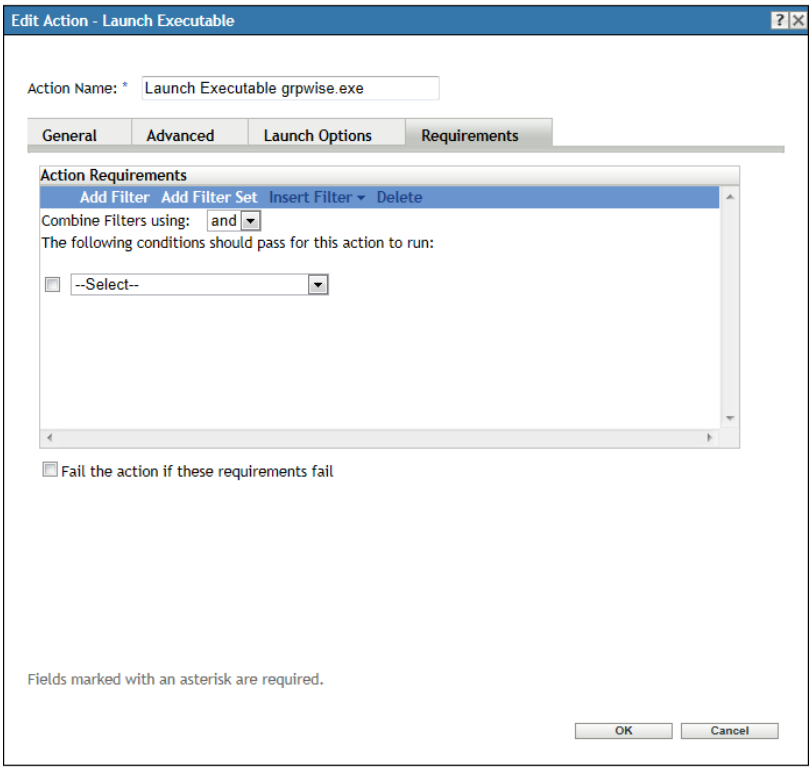

5b In the drop-down list, select *Architecture*, then select *32* or *64*.

5c Click *General* to return to the *General* tab.

5d In the *Action Name* field, customize the name with the architecture you selected for the launch action.

6 (Optional) In the *Command Line Parameters* field, add any desired client startup options, as listed in "Startup Options for the GroupWise Windows Client" in "Client" in the *GroupWise 2012 Administration Guide*.

For this installation action, no *Advanced* settings are required, because the default settings of *No wait* and *Run as logged in user* are appropriate.

7 Click *OK* to return to the Launch Executable dialog box.

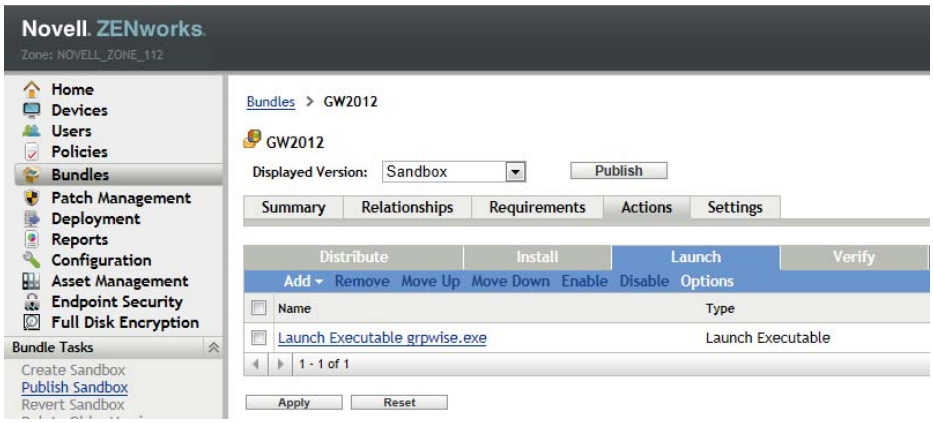

8 Click *Apply* to save the launch action.

9 (Conditional) If you selected an architecture requirement in Step 5, repeat Step 1 through Step 8 to create a launch action for the other workstation architecture in your environment.

10 Continue with Displaying GroupWise Bundle Activity on User Workstations.

15.2.4 Displaying GroupWise Bundle Activity on User Workstations

By default, ZENworks Configuration Management does not display any user interface when it is installing a software bundle. You can configure the GroupWise bundle to provide feedback to the user about how the installation is progressing.

1 On the GroupWise Bundle page, click *Summary*:

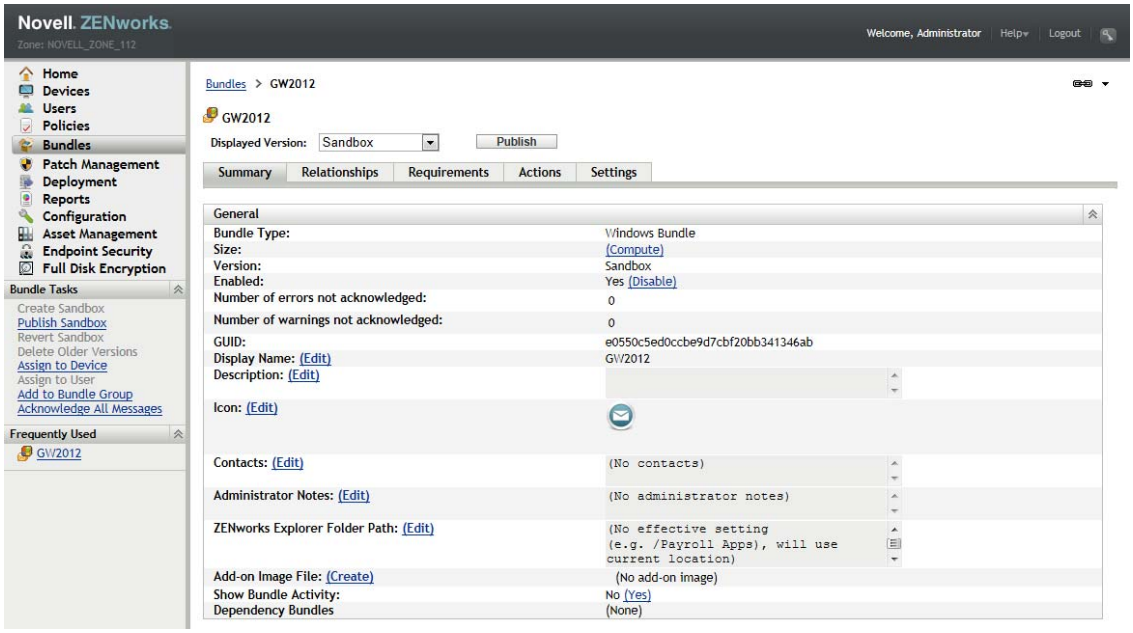

2 In the *Show Bundle Activity* field, click *Yes*.

With this setting enabled, users receive messages showing status information as the GroupWise Windows software is downloaded and installed on their workstations.

3 Continue with Section 15.2.5, "Publishing the GroupWise Bundle," on page 138.

15.2.5 Publishing the GroupWise Bundle

1 On the GroupWise Bundle page, click *Publish*.

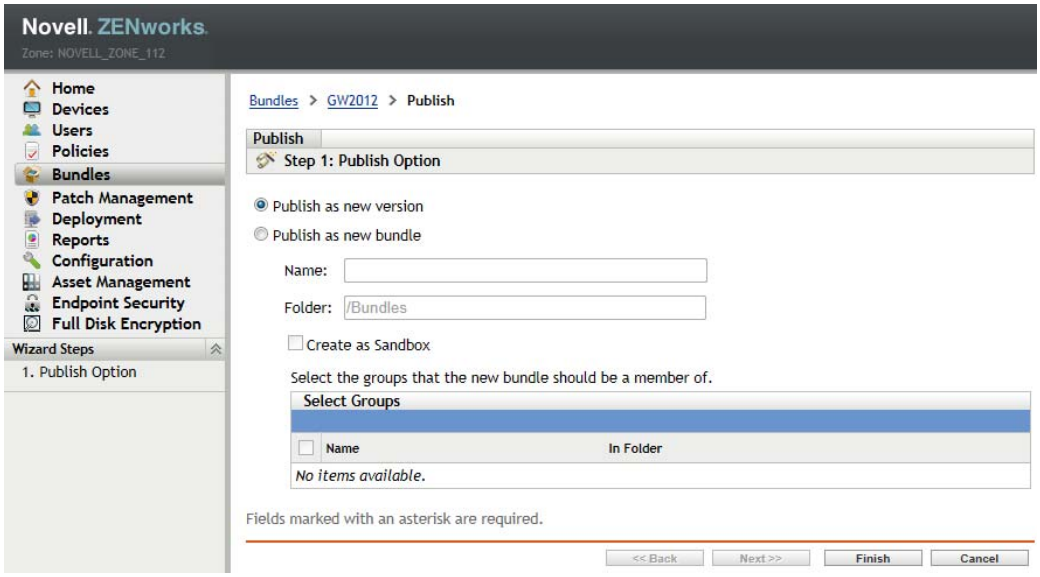

The default setting of *Publish as new version* is appropriate for the GroupWise bundle.

2 Click *Finish* to publish the GroupWise bundle.

This changes the *Displayed Version* field from *Sandbox* to *Published*.

3 Continue with Section 15.3, "Associating the GroupWise Bundle with Devices on Your Network," on page 138.

15.3 Associating the GroupWise Bundle with Devices on Your Network

This section assumes that you already have all the workstations where you want to install the GroupWise Windows client software defined as devices in ZENworks Configuration Management.

1 On the GroupWise Bundle page, click *Relationships*.

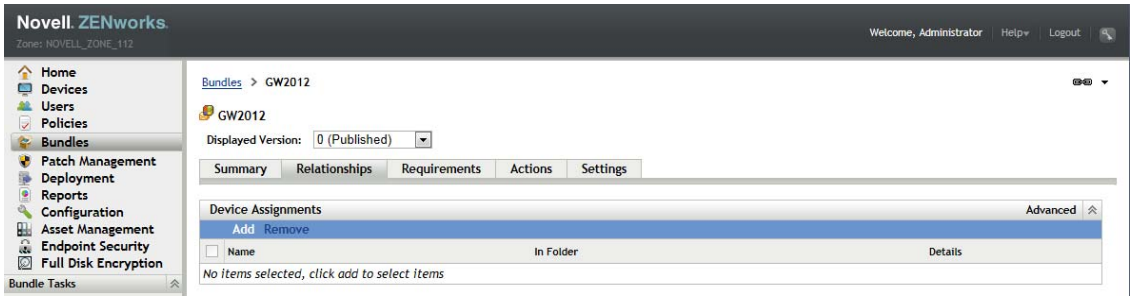

2 In the *Device Assignments* box, click *Add*.

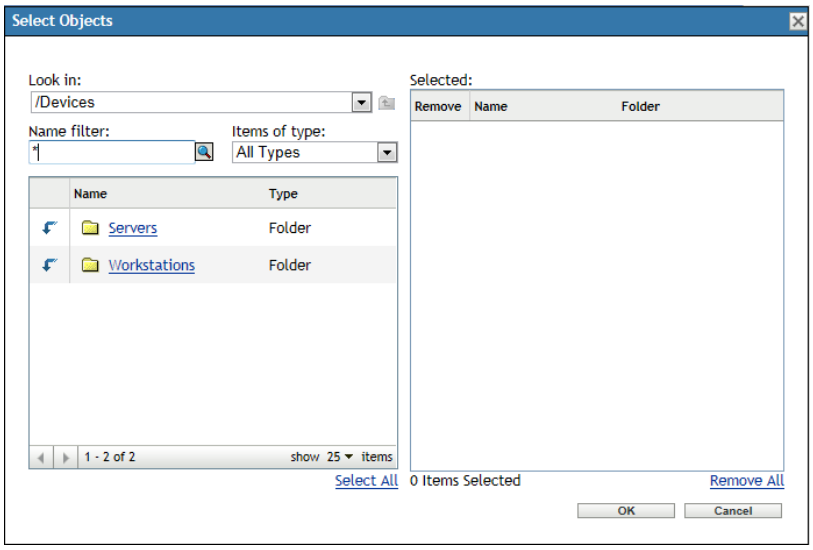

3 Display your available workstations and dynamic workstation groups.

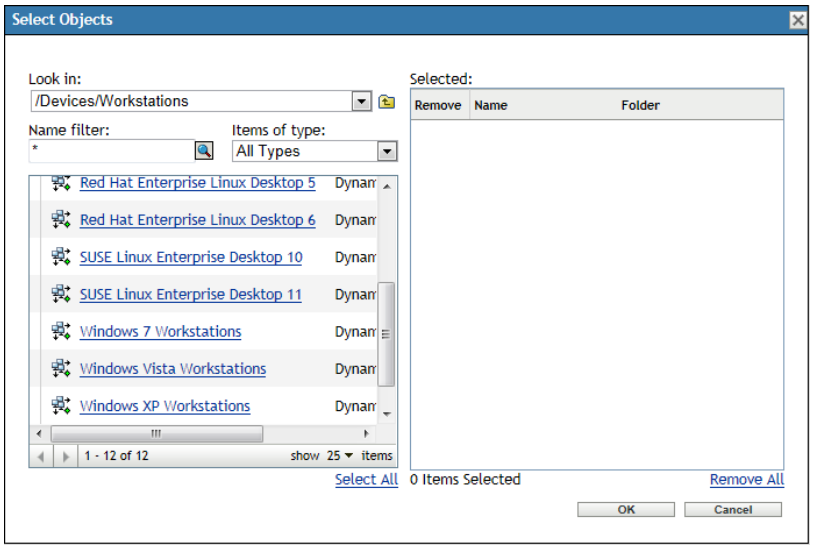

4 Select the workstation where you want to test the GroupWise bundle, then click *OK*.

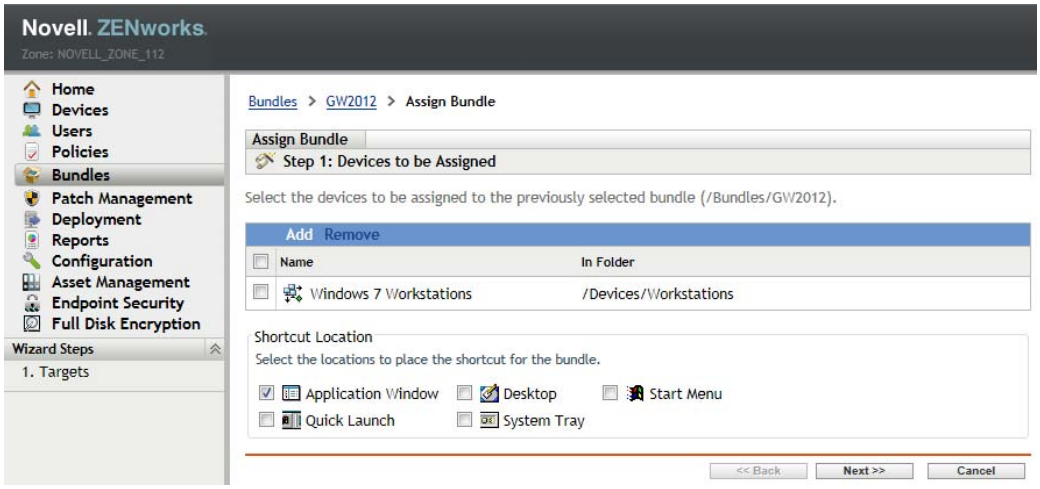

5 Select the locations where you want the GroupWise icon to display, then click *Next*.

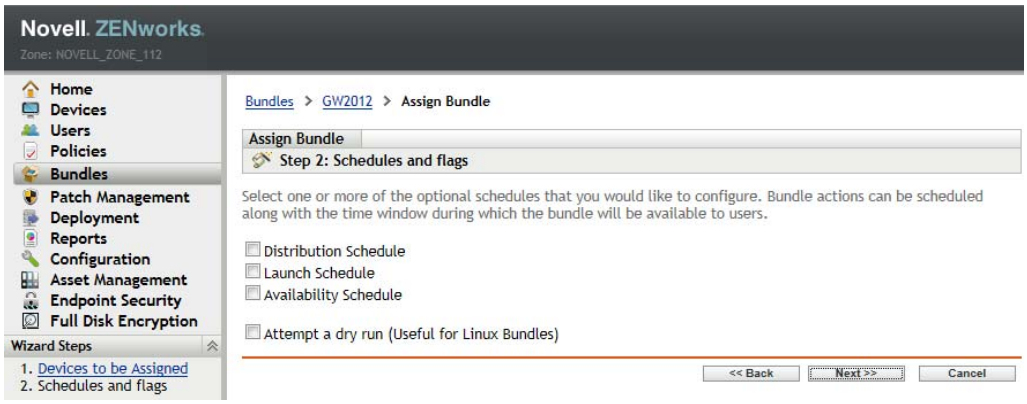

6 Select *Distribution Schedule*, then click *Next*.

7 For testing purposes, select *Now* in the *Schedule Type* field.

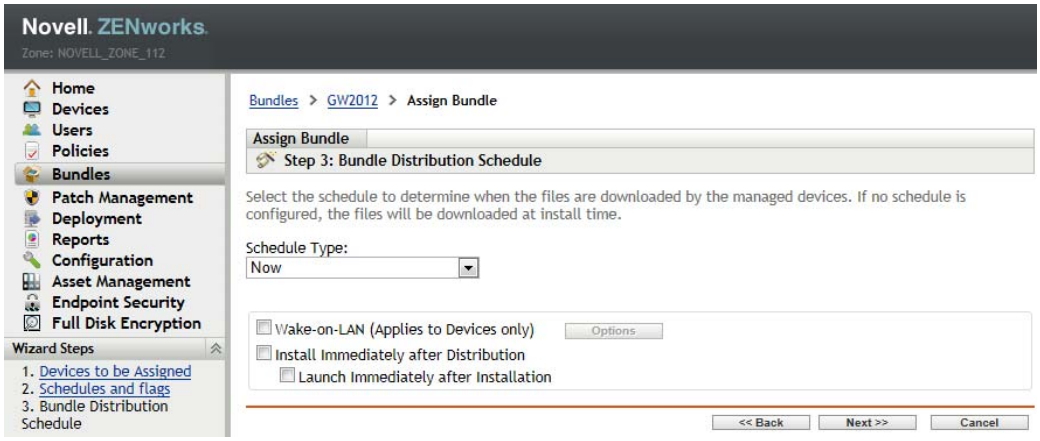

8 Select *Install immediately after Distribution*.

9 Select *Launch immediately after Installation*.

10 Click *Next*, review the bundle distribution summary, then click *Finish* to display the GroupWise bundle relationships with the selected workstation listed.

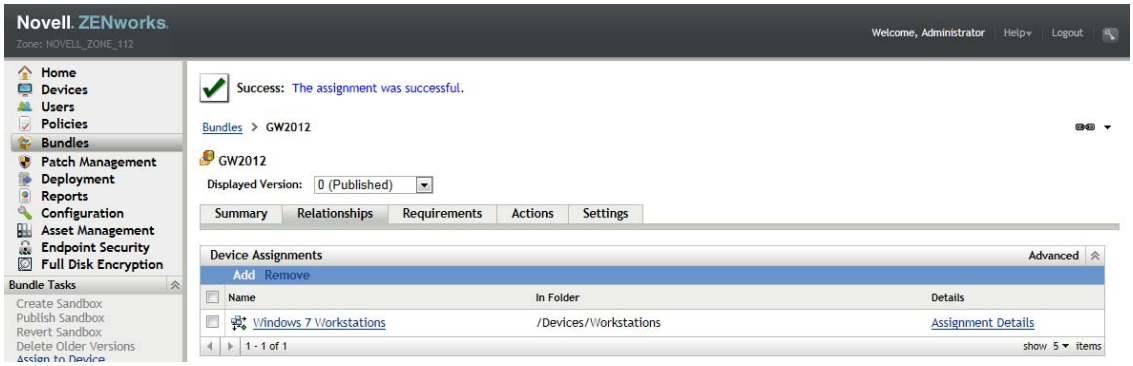

11 Continue with Testing the GroupWise Bundle on a Device.

15.4 Testing the GroupWise Bundle on a Device

Because you set the Distribution Schedule to *Now*, you should be able to go to the device you selected in Step 4 in Section 15.3, "Associating the GroupWise Bundle with Devices on Your Network," on page 138 and see the GroupWise Windows client software distribution take place. A progress window appears while the software is being downloaded.

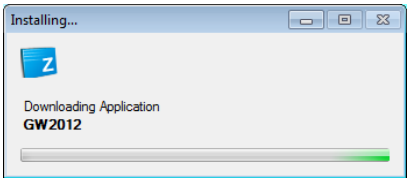

Then the GroupWise icon appears on the desktop with a down-arrow indicator showing that the GroupWise client software is being downloaded

The download process takes several minutes. After the download has been verified, the installation begins.

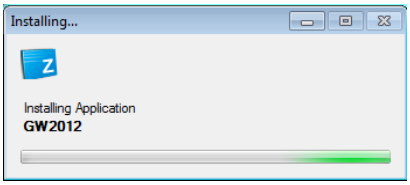

The indicator on the GroupWise icon changes to show that the GroupWise Windows client is being installed.

The installation process takes several minutes. When the installation is complete, the GroupWise icon simply indicates that the GroupWise software is being managed by ZENworks Configuration Management.

If an earlier version of GroupWise was already installed on the workstation, you also see that the old GroupWise icon is replaced by the new GroupWise icon.

Because of the way you configured the GroupWise bundle for testing purposes in Section 15.2.3, "Adding a Launch Action to the GroupWise Bundle," on page 135, the GroupWise Windows client should start immediately on the Windows workstation.

After the GroupWise bundle successfully deploys the GroupWise Windows client software on a test workstation, you can proceed with the rollout of the new GroupWise version.

16 Using ZENworks Linux Management to Distribute the GroupWise Linux Client

You can install the GroupWise Linux client and agents using Novell ZENworks Linux Management. Refer to the Novell ZENworks Linux Management Product Web site (http://www.novell.com/products/zenworks/linuxmanagement/) and the Novell ZENworks Linux Management Documentation Web site (http://www.novell.com/documentation/zlm73) for more information.

17 Using ZENworks Application Virtualization for GroupWise and Messenger

Novell ZENworks Application Virtualization lets you convert applications that run on Microsoft Windows into self-contained virtual applications. After being virtualized, an application becomes a single, isolated file that runs instantly from anywhere, including a thumb drive or other removable media. Unlike traditional installation methods, the single virtual application file does not require a separate setup process, and does not rely on external components and runtimes, reboots, or administrative privileges. After virtualization, the application is isolated from other system applications, preventing DLL conflicts and other deployment nightmares, yet the experience for the application's user is unchanged.

For instructions on virtualizing the GroupWise Windows client, see "Preparing GroupWise and GroupWise Notify for Virtualization" in the *ZENworks Integration and Streaming Guide* (http://www.novell.com/documentation/zav80).

For instructions on virtualizing the Messenger Windows client, see "Preparing Novell Messenger for Virtualization" in the *ZENworks Integration and Streaming Guide* (http://www.novell.com/documentation/zav80).

IV Other Novell Products

18 GroupWise DirXML Driver for Novell Identity Manager

The GroupWise DirXML driver for use with Novell Identity Manager provides data integration between users in Novell eDirectory and GroupWise accounts in your GroupWise system. For example, the driver can create email accounts automatically when employees are hired. The driver can also disable an email account when a user is no longer active. This configurable solution gives you the ability to increase productivity and streamline business processes by integrating GroupWise and eDirectory.

This guide provides information about certain administrative actions in ConsoleOne that require you to stop the GroupWise DirXML driver or disable a user's association:

- Section 18.1, "Identity Manager Warnings in ConsoleOne," on page 149

For additional information, see:

- Novell Identity Manager (http://www.novell.com/documentation/idm401)
- Identity Manager Drivers (http://www.novell.com/documentation/idm401drivers)

18.1 Identity Manager Warnings in ConsoleOne

Some GroupWise administrative actions in ConsoleOne require that you stop the GroupWise DirXML driver or disable a user's association with it before you perform the action. Most GroupWise administrative actions in ConsoleOne require that you manually restart the GroupWise DirXML driver or re-enable the user's association when you have completed the action, but a few do not. By default, these activities generate a warning message in ConsoleOne:

- Section 18.1.1, "Recovering a Deleted GroupWise Account," on page 150
- Section 18.1.2, "Grafting Users," on page 150
- Section 18.1.3, "Converting an External Entity to a User," on page 150
- Section 18.1.4, "Converting a User to an External Entity," on page 150
- Section 18.1.5, "Associating a GroupWise Object with an eDirectory Object," on page 150
- Section 18.1.6, "Disassociating a GroupWise Object's Attributes from an eDirectory Object," on page 151
- Section 18.1.7, "Resolving an Invalid Association," on page 151
- Section 18.1.8, "Disabling the DirXML Warnings," on page 151
- Section 18.1.9, "Enabling the DirXML Warnings," on page 151

18.1.1 Recovering a Deleted GroupWise Account

Using the DirXML Management role in Novell iManager:

1. Stop the GroupWise DirXML driver.
2. Recover the deleted account, as described in "Recovering Deleted GroupWise Accounts" in "Databases" in the *GroupWise 2012 Administration Guide*.
3. Restart the GroupWise DirXML driver.

18.1.2 Grafting Users

Using the DirXML Management role in Novell iManager:

1. (Conditional) If you are grafting the users into a different eDirectory tree, go to the *DirXML* tab of each User object, then disable the association with the GroupWise DirXML driver.
2. Stop the GroupWise DirXML driver for the tree into which you are grafting the users.
3. Graft the users, as described in "Graft GroupWise Objects" in "Databases" in the *GroupWise 2012 Administration Guide*.
4. (Conditional) If you grafted the users into a different eDirectory tree, go to the *DirXML* tab of each User object, then enable the association with the GroupWise DirXML driver in the new tree.
5. Restart the GroupWise DirXML driver for the tree into which you grafted the users.

18.1.3 Converting an External Entity to a User

Using the DirXML Management role in Novell iManager:

1. Stop the GroupWise DirXML driver.
2. Convert the external entity, as described in "Convert External Entity to User" in "System" in the *GroupWise 2012 Administration Guide*.
3. Restart the GroupWise DirXML driver.

18.1.4 Converting a User to an External Entity

In Novell iManager:

1. Go to the *DirXML* tab of the User object in Novell iManager, then disable the association with the GroupWise DirXML driver.
2. Convert the user, as described in "Convert User to External Entity" in "System" in the *GroupWise 2012 Administration Guide*.

18.1.5 Associating a GroupWise Object with an eDirectory Object

Using the DirXML Management role in Novell iManager:

1. Stop the GroupWise DirXML driver.
2. Establish the association, as described in "Associate Objects" in "System" in the *GroupWise 2012 Administration Guide*.
3. Restart the GroupWise DirXML driver.

18.1.6 Disassociating a GroupWise Object's Attributes from an eDirectory Object

In Novell iManager:

1. Go to the *DirXML* tab of the User object, then disable the association with the GroupWise DirXML driver.
2. Disassociate the objects, as described in "Disassociate GroupWise Attributes" in "System" in the *GroupWise 2012 Administration Guide*.
3. Go to the *DirXML* tab of the User object, then enable the association with the GroupWise DirXML driver.

18.1.7 Resolving an Invalid Association

In Novell iManager:

1. Go to the *DirXML* tab of the User object, then disable the association with the GroupWise DirXML driver.
2. Resolve the invalid association, as described in "Invalid Associations" in "System" in the *GroupWise 2012 Administration Guide*.

18.1.8 Disabling the DirXML Warnings

In ConsoleOne:

1. Deselect *Display DirXML Warnings* in any DirXML warning dialog box.

18.1.9 Enabling the DirXML Warnings

In ConsoleOne:

1. Click *Tools* > *GroupWise System Operations* > *System Preferences*.
2. On the *Admin Preferences* tab, select *Display DirXML Warnings*.
3. Click *OK*.

19 GroupWise Customization Tools for Developers

The GroupWise Software Developer Kit provides tools for customizing GroupWise to the specific needs of your organization. It includes the following components:

- **GroupWise Object API:** Lets you create your own client application. It provides access to the Address Book, along with documents, mail messages, appointments, tasks, notes, phone messages, and workflow items. The GroupWise Object API supports COM Automation, which is an industry standard for interfacing applications and is simple to use with languages such as Delphi, Visual Basic, and C++. For more information, see GroupWise Object API (http://developer.novell.com/wiki/index.php/GroupWise_Object_API).
- **GroupWise Administrative Object API:** Lets you see, use, and manipulate GroupWise administration information from outside GroupWise. You can use the GroupWise Administrative Object API through COM languages, such as Visual Basic, Delphi, and object-oriented languages such as C++. It also supports COM Automation, which is an industry standard for interfacing applications. For more information, see GroupWise Administrative Object API (http://developer.novell.com/wiki/index.php/GroupWise_Administrative_Object_API).
- **GroupWise C3PO (Custom 3rd-Party Object):** Works with C++, Delphi, or Visual Basic to let you add menu and toolbar items to trigger applications. For example, you can modify the GroupWise client toolbar or define new record types in the GroupWise information store. For more information, see GroupWise C3PO (http://developer.novell.com/wiki/index.php/GroupWise_C3PO).
- **GroupWise Tokens:** Lets you manipulate the GroupWise client interface by subscribing to internal token events or by publishing new tokens to the client. It names low-level events, such as “save a file” or “send mail,” which allows you to extend GroupWise functionality. A C3PO lets you extend GroupWise objects and the Object API lets you see and manipulate the GroupWise information store from outside GroupWise. In addition, tokens let your solution command the GroupWise client from DLLs and DDE scripts, using the Third-Party Handler. You can also use tokens to create Visual Basic executables that users can run from the client interface. For more information, see GroupWise Tokens (http://developer.novell.com/wiki/index.php/GroupWise_Tokens).
- **GroupWise Trusted Applications:** Enables you to develop applications that can log in to any user’s mailbox without supplying the user’s password and perform various tasks such as virus scanning, content filtering, or email auditing. For more information, see GroupWise Trusted Application API (http://developer.novell.com/wiki/index.php/GroupWise_Trusted_Application_API).
- **GroupWise SOAP:** Uses defined standards to provide access to GroupWise data directly from the GroupWise post office. The standards used include HTTP, SOAP, XML, XML schemas, and WSDL. HTTP, SOAP, and XML are used to transport data from between computers. XML schemas define the structure and the types of GroupWise data that is transported. The

GroupWise WSDL (Service Descriptive Language) combines everything into a GroupWise Web service. For more information, see GroupWise Web Service (SOAP) (http://developer.novell.com/wiki/index.php/GroupWise_Web_Service_%28SOAP%29).

- **GroupWise MAPI:** Uses a set of object-oriented functions that provide messaging capabilities. The Messaging Application Programming Interface (MAPI) is used by mail-enabled applications to create, transfer, and store messages, as well as to handle complex addressing information. MAPI objects are data structures that support a set of properties and that comply with the component object model, which requires that objects support one or more interfaces or sets of functions. For more information, see GroupWise MAPI (http://developer.novell.com/wiki/index.php/GroupWise_MAPI).
- **GroupWise Events:** Provides event notification to registered third-party applications, and is responsive to queries, while not significantly degrading the overall performance of the GroupWise Post Office Agent (POA). This functionality is included in the GroupWise Object API (http://developer.novell.com/wiki/index.php/GroupWise_Object_API).
- **GroupWise Controls for ActiveX:** Lets you embed an Address Book or Name Completion COM Control (OCX) in your Visual Basic, Delphi, and C++ solutions. OCX properties let you customize user access to Address Book contents and control return information for your solution to use. For more information, see GroupWise Controls for ActiveX (http://developer.novell.com/wiki/index.php/GroupWise_Controls_for_ActiveX).

V Microsoft Clustering on Windows

20 Introduction to GroupWise 2012 and Clustering on Windows

Before implementing GroupWise 2012 in a cluster on Windows, make sure you have a solid understanding of Microsoft clustering technologies by reviewing the following resources:

- Windows Server 2008 – Failover Clustering (http://www.microsoft.com/en-us/server-cloud/windows-server/failover-clustering-network-load-balancing.aspx)
- Windows Server 2003 – Microsoft Cluster Service (MSCS) (http://msdn.microsoft.com/en-us/library/ms952401.aspx)

When you review the resources recommended above, you discover that clustering employs very specialized terminology. The following brief glossary provides basic definitions of clustering terms and relates them to your GroupWise system:

cluster: A grouping of from two to eight Windows servers configured so that data storage locations and applications can transfer from one server to another without interrupting their availability to users.

node: A clustered server; in other words, a single Windows server that is part of a cluster.

active node: A node in the cluster that is actively running programs. An active node makes its resources available in the cluster.

passive node: A node in the cluster that is not currently running programs, but is waiting for an active node to fail. A passive node does not make its resources available in the cluster until an active node fails over to it.

resource: A data storage location or application. For example, a domain directory and the MTA for the domain are resources. A post office directory and the POA for the post office are resources.

resource group: Two or more resources that must fail over together in order to remain functional. For example, for a domain to be functional, the domain directory and its MTA must fail over together. For a post office to be functional, the post office directory and its POA must fail over together

physical disk: The physical location where resources are created or installed. For example, a domain or post office directory is created on a physical disk. The agent software is installed on a physical disk.

shared disk: A physical disk that can be made active on any node in the cluster.

failover: The process of moving resources and resource groups on a shared disk from a failed node to a functional node so that availability to users is uninterrupted. For example, if the node where the POA is running goes down, the post office resource group fails over to another node so that users can continue to use GroupWise.

fan-out-failover: The configuration where resources and resource groups from a failed node fail over to different nodes in order to distribute the load from the failed node across multiple nodes in the cluster. For example, if a node runs a resource group consisting of a domain and its MTA, another resource group consisting of a post office and its POA, and a third resource group for the GWIA, each resource group can be configured to fail over separately to different nodes in the cluster.

failback: The process of returning resources and resource groups to their original node after the situation causing the failover has been resolved. For example, if a POA and its post office fail over to another node in the cluster, that resource group can be configured to fail back to its original node when the problem is resolved.

shared disk system: The hardware housing the physical disks that are shared among the nodes in the cluster. The C: drives in the clustered nodes are not part of the shared disk system. Each C: drive belongs to its own server.

storage area network (SAN): The clustered nodes together with their shared disk system and shared physical disks.

21 Planning GroupWise in a Windows Cluster

The majority of this part of the guide (Chapter 21, "Planning GroupWise in a Windows Cluster," on page 159 through Chapter 26, "Backing Up a GroupWise System in a Windows Cluster," on page 201) is designed for those who are creating a new GroupWise system, or at least new domains and post offices, in a Windows cluster. If you already have an existing GroupWise 2012 system and need to configure it to work in a newly installed cluster, see Chapter 27, "Moving an Existing GroupWise 2012 System into a Windows Cluster," on page 203.

When you implement a new GroupWise system or a new domain or post office in a Windows cluster, your overall GroupWise system design does not need to change substantially. For a review, see "Installing a Basic GroupWise System" in the *GroupWise 2012 Installation Guide*. However, the configuration of individual components of your GroupWise system is significantly different. This section helps you plan the following GroupWise components in a Windows cluster:

- A new GroupWise system consisting of the primary domain and the initial post office
- A new secondary domain
- A new post office
- The GroupWise agents (MTA, POA, and DVA)

During the planning process, component configuration alternatives are explained. For example, you might want the domain and post office together in the same resource group or in separate resource groups. You might want to install the agents to the standard `c:\Program Files\Novell\GroupWise Server\Agents` directory on each node or to manually create a *`drive`*`:\Program Files\Novell\GroupWise Server\Agents` directory for each shared disk for domains and post offices so that the agents fail over with the domains and post offices they service.

The System Clustering Worksheet lists all the information you need as you set up GroupWise in a Windows cluster. You should print the worksheet and fill it out as you complete the tasks listed below:

After you have completed the tasks and filled out the "System Clustering Worksheet" on page 170, you are ready to continue with Chapter 22, "Setting Up a Domain and Post Office in a Windows Cluster," on page 175.

21.1 Setting Up Your Windows Cluster

As you set up your Windows cluster, record key information about the cluster on the System Clustering Worksheet:

SYSTEM CLUSTERING WORKSHEET

Under Item 1: Cluster Name, record the name of your Windows cluster.

Under Item 2: Nodes in Cluster, list the servers that you have added to the cluster.

The number of nodes in the cluster strongly influences where you place GroupWise domains and post offices. You have several alternatives:

- Your whole GroupWise system can run in a single cluster.
- Parts of your GroupWise system can run in one cluster while other parts of it run in one or more other clusters.
- Parts of your GroupWise system can run in a cluster while other parts run outside of the cluster, on non-clustered servers.

If you do not have the system resources to run all of your GroupWise system in the cluster, you must decide which parts have the most urgent need for the high availability provided by clustering. Here are some suggestions:

- Post offices and their POAs must be available in order for users to access their GroupWise mailboxes. Therefore, post offices and their POAs are excellent candidates for the high availability provided in a cluster.
- Domains and their MTAs are less noticeable to GroupWise client users when they are unavailable (unless users in different post offices happen to be actively engaged in an email discussion when the MTA goes down). On the other hand, domains and their MTAs are critical to GroupWise administrators, although administrators might be more tolerant of a down server than client users are. Critical domains in your GroupWise system are the primary domain and, if you have one, a hub or routing domain. These domains should be in the cluster, even if other domains are not.
- The GWIA might or might not require high availability in your GroupWise system, depending on the importance of immediate messaging across the Internet and the use of POP3 or IMAP4 clients by GroupWise users.

There is no right or wrong way to implement GroupWise in a cluster. It all depends on the specific needs of your particular GroupWise system and its users.

21.2 Planning a New Clustered Domain

The considerations involved in planning a new domain in a Windows cluster are essentially the same as for any other environment.

- **Primary Domain:** If you are setting up a new GroupWise system in a Windows cluster, you create the primary domain as you complete the tasks in this section. In preparation, review "Planning a Basic GroupWise System", then print and fill out the "Basic GroupWise System Worksheet" in "Installing a Basic GroupWise System" in the *GroupWise 2012 Installation Guide*. This covers planning the primary domain and an initial post office in the primary domain.
- **Secondary Domain:** If your GroupWise system already exists, you create a new secondary domain. In preparation, review "Planning a New Domain", then print and fill out the "New Domain Summary Sheet" in "Domains" in the *GroupWise 2012 Administration Guide*.

Regardless of the type of domain you are creating, keep in mind the following cluster-specific details as you fill out the worksheet you need:

- When you specify the location for the domain directory (and for a new GroupWise system, the post office directory) on the worksheet, include the shared disk where you want the directory to reside.
- Do not concern yourself with the GroupWise agent information on the worksheet. You will plan the agent installation later. If you are filling out the Basic GroupWise System Worksheet, stop with Post Office Settings. If you are filling out the Domain Worksheet, stop with Domain Administrator.

When you have completed the worksheet, transfer the key information from the Basic GroupWise System Worksheet or the Domain Worksheet to the System Clustering Worksheet.

SYSTEM CLUSTERING WORKSHEET

Under Item 7: Domain Name, transfer the domain name and directory to the System Clustering Worksheet.

IMPORTANT: Do not create the new domain until you are instructed to do so in Chapter 22, "Setting Up a Domain and Post Office in a Windows Cluster," on page 175.

21.3 Planning a New Clustered Post Office

The considerations involved in planning a new post office in a Windows cluster are essentially the same as for any other environment. The initial post office in a new GroupWise system is planned on the Basic GroupWise System Worksheet. To plan additional new post offices, review "Planning a New Post Office ", then print and fill out the "New Post Office Summary Sheet" in "Post Offices" in the *GroupWise 2012 Administration Guide*. When you specify the locations for the post office directories, include the shared disks where you want the post office directories to reside.

When you have completed the worksheet, transfer key information from the Basic GroupWise System Worksheet or the Post Office Worksheet to the System Clustering Worksheet.

SYSTEM CLUSTERING WORKSHEET

Under Item 8: Post Office Name, transfer the post office name and directory to the System Clustering Worksheet.

IMPORTANT: Do not create the new post office until you are instructed to do so in Chapter 22, "Setting Up a Domain and Post Office in a Windows Cluster," on page 175.

21.4 Planning a New Library for a Clustered Post Office

The considerations involved in planning a library in a Windows cluster are essentially the same as for any other environment. You can plan a library for a new clustered post office by following the standard instructions provided in "Creating and Managing Libraries" in the *GroupWise 2012 Administration Guide* and filling out the "Basic Library Worksheet" or the "Full-Service Library Worksheet". Then provide the library information on the System Clustering Worksheet.

SYSTEM CLUSTERING WORKSHEET

Under Item 9: Document Storage Area Location, mark where you want to create the library's document storage area.

If the document storage area will be located outside the post office directory structure and outside the cluster, specify a user name and password that the POA can use to access the server where the document storage area will reside.

IMPORTANT: Do not create the new library until you are instructed to do so in Chapter 22, "Setting Up a Domain and Post Office in a Windows Cluster," on page 175.

21.5 Planning GroupWise Resource Groups

Resource groups ensure that resources that depend on each other fail over together. If your GroupWise system is very small (for example, one domain and one post office), you can have a single GroupWise resource group so that your whole GroupWise system fails over together. More typically, multiple domains and post offices are located throughout your organization, so you must set up a resource group for each domain and post office.

A resource group for a domain or post office must include the following types of resources:

- **Network Name:** The virtual name by which the domain or post office resource group will be known on the network, regardless of which node it is active on.
- **IP Address:** The virtual IP address that will be associated with the network name, regardless of which node the domain or post office resource group is active on.
- **Physical Disk:** The drive letter where the domain or post office directory will be located, used when mapping a drive to the physical disk.
- **File Share:** The name of the physical disk, used when mapping a drive to the physical disk.
- **Generic Service:** The GroupWise agent, running as a Windows service, that will service the domain or post office.

For convenience, you might want to name each resource group after the domain or post office it represents. In this documentation, a resource group that can include a domain, a post office, or both, is termed a "GroupWise resource group."

Each GroupWise resource group has a list of possible owners associated with it. The possible owners are the nodes to which the resource group can fail over. By default, a resource group is configured to have all nodes in the cluster in its possible owners list, organized in ascending alphanumeric order. Only one node at a time can have a particular GroupWise resource group active. If a resource group's

current owner node fails, the resource group fails over to the next node in the possible owners list. You should customize the owners list for each GroupWise resource group based on the fan-out-failover principle.

When a node fails, its resource groups should not all fail over together to the same node in the cluster. Instead, the resource groups should be distributed across multiple nodes throughout the cluster. This prevents any one node from shouldering the entire processing load typically carried by another node. In addition, some GroupWise resource groups should never have the potential of failing over to the same node. For example, a post office and POA that service a large number of very active GroupWise client users should never fail over to a node where another very large post office and heavily loaded POA reside. If they did, users on both post offices would notice a decrease in responsiveness of the GroupWise client.

IMPORTANT: If you are planning more than one DVA or GWIA in the cluster, you must ensure that they can never fail over to the same node at the same time. You cannot customize the Windows service name for the DVA or the GWIA. Therefore, only one of each can run on a server. The Windows service names for POAs and MTAs include the name of the post office or domain that they service, so this limitation does not apply to POAs and MTAs.

SYSTEM CLUSTERING WORKSHEET

Under Item 4: Resource Group for Domain, specify the network name and other required information for the domain resource group. Mark whether you will place the post office in the same resource group with the domain.

If you want the post office in a different resource group from where the domain is located, under Item 5: Resource Group for Post Office, specify the network name and other required information for the post office resource group.

21.6 Planning Shared Administrative Resources

Depending on your administrative needs, you might or might not want to set up shared administrative resources. For example, you might want to have a shared disk where you install the GroupWise snap-ins to ConsoleOne instead of installing them on multiple administrator workstations. You might also have a shared disk where you create the GroupWise software distribution directory. These shared disks can be configured to fail over as part of your clustered environment.

SYSTEM CLUSTERING WORKSHEET

Under Item 3: Resources for GroupWise Administration, list any shared disks you want to use for GroupWise administration purposes.

21.7 Ensuring Successful Name Resolution for GroupWise Resource Groups

When you establish GroupWise resource groups, you establish network names for the locations of domains and post offices. The network names remain constant no matter which node in the cluster the domain or post office is currently active on. Because you are using virtual network names, not physical locations, you must ensure that short name resolution is always successful. For example, in ConsoleOne, if you right-click a Domain object in the GroupWise View and then click *Connect*,

ConsoleOne must be able to resolve the domain database location, as provided in the *UNC Path* field, to the network name of that domain within your cluster. It is through short name resolution that all GroupWise resource groups are accessed and managed in ConsoleOne.

A client program (such as ConsoleOne) that runs on a Windows workstation, can be configured to use several different short name resolution methods. To see which methods are in use at a particular workstation, view the protocol preferences for the Novell Client that is installed on the Windows workstation:

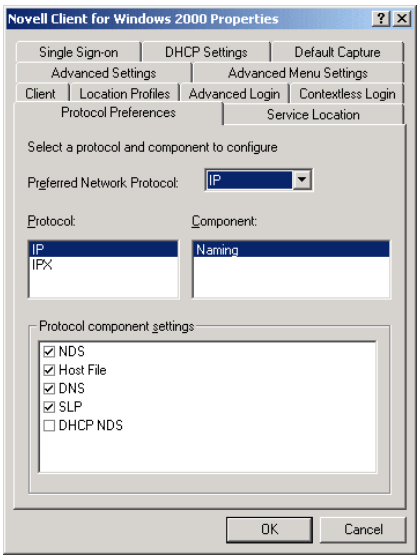

Short name resolution methods that pertain to your clustered GroupWise system are discussed below:

Short Name Resolution Method	Description
eDirectory	You can use Novell eDirectory to resolve short names into specific network addresses. However, when you use eDirectory for short name resolution, you must remember to consider current context in the name resolution process. eDirectory short name resolution works only if your current context is the same as the context of the eDirectory object you need to access.
Hosts File	Windows XP/Vista/7 uses the `\Windows\System32\drivers\etc\hosts` file when performing short name resolution at the workstation. Using this file at the Windows workstation is not a preferred method for short name resolution (except perhaps for the administrator's workstation).
DNS	Perhaps the most common short name resolution option is Domain Name Service (DNS). As with the hosts file, it is good practice to place all the network names of your GroupWise resource groups into DNS. For short name resolution to work using DNS, the client workstation must either belong to the same DNS zone (such as `support.novell.com`) as the resource group, or the cluster resource zone must be configured in the client's DNS suffix search path under TCP/IP settings for the workstation.

Specific setup instructions for each of these short name resolution methods are provided in Chapter 22, "Setting Up a Domain and Post Office in a Windows Cluster," on page 175.

SYSTEM CLUSTERING WORKSHEET

Under Item 6: IP Address Resolution Methods, mark which methods you want to implement in order to resolve the locations stored as UNC paths in ConsoleOne into the network names of the GroupWise resource groups.

21.8 Deciding How to Install and Configure the Agents in a Windows Cluster

There are several cluster-specific issues to consider as you plan to install the Windows MTA, POA, and DVA in your clustered GroupWise system:

- Section 21.8.1, "Planning Cluster-Unique Port Numbers for Agents in the Cluster," on page 165
- Section 21.8.2, "Deciding Where to Install the Agent Software," on page 167
- Section 21.8.3, "Planning the Agent Services," on page 169
- Section 21.8.4, "Planning the Windows Agent Installation," on page 169

21.8.1 Planning Cluster-Unique Port Numbers for Agents in the Cluster

By default, the GroupWise agents listen on all IP addresses, both primary and secondary, that are bound to the server on their specified port numbers. The primary IP address is the IP address of the physical node. Secondary IP addresses are the IP addresses associated with GroupWise resource groups. The secondary IP address moves with each agent when it fails over, so that, in the case of the POA, GroupWise clients do not lose their connections to the POA.

If you are planning to set up a GroupWise name server to help GroupWise clients locate their post offices, make sure that the default POA port number of 1677 is used somewhere in the cluster. For more information, see "Simplifying Client/Server Access with a GroupWise Name Server" in "Post Office Agent" in the *GroupWise 2012 Administration Guide*.

If you want to install the DVA along with the POA, it can participate in the cluster by failing over along with the POA and its post office. However, clustering is not necessary in order to provide high availability for the DVA. As an alternative to clustering, you can install the DVA on multiple servers outside the cluster, as described in "Scaling Your DVA Installation" in "Document Viewer Agent" in the *GroupWise 2012 Administration Guide*.

Any time there is a possibility of two of the same type of agent running on the same node, it is important that each agent use a cluster-unique port number, even though each agent is using the unique IP address established for each GroupWise resource group. The best way for you to avoid port conflicts is to plan your cluster so that each agent in the cluster runs on a cluster-unique port. Print out a copy of the "Network Address Worksheet" on page 171 to help you plan cluster-unique port numbers for all GroupWise agents in your GroupWise system.

The following filled-out version of the Network Address Worksheet illustrates one way this can be done:

Domain Information

Domain	MTA IP Address	MTA MTP Port	MTA HTTP Port
Provo1	172.16.5.81	7100	7180

Post Office Information

Post Office	POA IP Address	POA C/S Port	POA MTP Port	POA HTTP Port	DVA HTTP Port
Development	172.16.5.81 (same as MTA)	1677	7101	7181	8301
Manufacturing	172.16.5.82	1678	7102	7182	8311

GWIA Information

GWIA	GWIA IP Address	MTA MTP Port	MTA HTTP Port	GWIA HTTP Port
GWIA Domain MTA	172.16.5.83	7110	7183	N/A
GWIA	172.16.5.83 (same as MTA)	N/A	N/A	9850

This example places the Development post office in the same resource group with the Provo1 domain; therefore, the Provo1 MTA and the Development POA use the same IP address. The Manufacturing post office is placed in a different resource group, so the Manufacturing post office has a different IP address. A DVA accompanies the POA in each post office resource group. The GWIA has its own domain and separate resource group.

The example also illustrates that the MTA, the POA, and the GWIA use different port numbers for agent ports and HTTP ports.

The example uses default port numbers where possible. For example, the default MTA message transfer port is 7100 and the default POA client/server port is 1677. Incrementing port numbers are used in the example when multiple components have the same type of ports. For example, port numbers 1677 and 1678 are both POA client/server ports and port numbers 7180 through 7184 are all HTTP ports. Incrementing from the default port numbers generates unique, though related, port numbers.

NOTE: During normal operation, the DVA uses a range of port numbers. Therefore, the base port numbers for DVAs should be spaced out. For example, using the default port number of 8301 for the first DVA and using 8311 for the second port number allows the first DVA to use ports 8302 through 8310 during its normal processing without conflicting with the second DVA in the cluster, which can use ports 8311 and above.

If you are going to set up a GroupWise name server to help GroupWise clients locate their post offices, make sure that the default POA port number of 1677 is used somewhere in the cluster and specify the IP address of the post office resource group, not the IP address of a specific node. For more information, see "Simplifying Client/Server Access with a GroupWise Name Server" in "Post Office Agent" in the *GroupWise 2012 Administration Guide*.

NETWORK ADDRESS WORKSHEET

Fill out the "Network Address Worksheet" on page 171 to help you determine resource group IP addresses and cluster-unique port numbers for all GroupWise agents in the cluster (MTA, POA, DVA, and GWIA). Refer to the IP addresses you planned for the domain and post office resource groups under items 4 and 5 on the System Clustering Worksheet.

After you have filled out the Network Address Worksheet, transfer the IP addresses and the cluster-unique port numbers from the Network Address Worksheet to the Agent Clustering Worksheet so that they will be available in the sequence in which you will need them as you set up the GroupWise agents in the cluster.

AGENT CLUSTERING WORKSHEET

Under Item 4: MTA Network Information, transfer the resource group IP address and cluster-unique port numbers for the MTA from the Network Address Worksheet to the Agent Clustering Worksheet.

Under Item 7: POA Network Information, transfer the resource group IP address and cluster-unique port numbers for the POA from the Network Address Worksheet to the Agent Clustering Worksheet.

Under Item 8: DVA Network Information, transfer the resource group IP address and cluster-unique port numbers for the DVA from the Network Address Worksheet to the Agent Clustering Worksheet.

21.8.2 Deciding Where to Install the Agent Software

In a Windows cluster, the agents must run as Windows services. When you install the Windows MTA, POA, and DVA, you can choose between two different installation locations:

Location	Description
Each node in the cluster	The `c:\Program Files\Novell\GroupWise Server\Agents` directory is the default location provided by the Agent Installation program.
Shared disk	If you create a *`drive`*`:\program Files\Novell\GroupWise Server\Agents` directory on the same shared disk with the domain or post office directory, the agent software and startup files fail over and back with the domains and post offices that the agents service.

Because the agents must be installed as Windows services in a Windows cluster, you must initially run the Agent Installation program for each node in the cluster so that the Windows services for the agents get created, regardless of where you are planning to run the agents from. However, for updates, you need to run the Agent Installation program only once if you are running the agents from a shared disk.

The following sections can help you choose which installation location is best for your clustered GroupWise system:

- "Advantages of Installing to a Shared Disk" on page 168
- "Disadvantages of Installing to a Shared Disk" on page 168
- "Recommendation" on page 168

Advantages of Installing to a Shared Disk

Using a *drive*:\Program Files\Novell\GroupWise Server\Agents directory for each GroupWise shared disk has several advantages:

- When you update the agent software, you only need to update it in one place for a particular domain or post office, not on every node in the resource group's possible owners list. This prevents the potential problem of having a domain or post office fail over to a node where a different version of the agent software is installed.
- Having the agent startup files on the same node as the domain or post office makes them easy to find.
- If you change information in the agent startup files, you only need to change it in one place, not on every node in the resource group's possible owners list.
- If you want to back up the GroupWise data, you can back up the domain and/or post office directories and the agent software from the same shared disk.

Disadvantages of Installing to a Shared Disk

Installing the agents on the same shared disk with a domain or post office does have some disadvantages:

- You must install the agent software each time you create a new domain or post office on a new shared disk.
- GroupWise administrators who are used to the GroupWise agents being installed in c:\Program Files\Novell\GroupWise Server\Agents might be confused by not finding them there in the clustered GroupWise system.
- You must remember where you installed the GroupWise agents when you update the agent software. Accidently installing a GroupWise Support Pack to the default location of c:\Program Files\Novell\GroupWise Server\Agents on the active node would not have the desired results if the original agent software was installed to a shared disk.

Recommendation

Whichever method you choose, be consistent throughout the entire cluster. Either put all the GroupWise agents on the shared disks with the domains and post offices they service, or put them all in c:\Program Files\Novell\GroupWise Server\Agents directories on all nodes. If you put them on shared disks with domains and post offices, make sure there are no agent files in c:\Program Files\Novell\GroupWise Server\Agents directories on nodes to create future confusion.

Even if you choose to install the agents to the c:\Program Files\Novell\GroupWise Server\Agents directory of multiple nodes, you can still store the agent startup files on shared disks with the domains and post offices. The significant advantage of this approach is that you have only one startup file to modify per agent.

AGENT CLUSTERING WORKSHEET

Under Item 1: Agent Installation Location, mark whether you will install the agent software to the shared disk with a domain or post office, or to each node in the cluster. If necessary, specify where the agent startup files will be stored.

Under Item 2: Domain Name, transfer the domain name, shared disk, and directory from the System Clustering Worksheet to the Agent Clustering Worksheet.

Under Item 5: Post Office Name, transfer the post office name, shared disk, and directory from the System Clustering Worksheet to the Agent Clustering Worksheet.

21.8.3 Planning the Agent Services

In a Windows cluster, the MTA and POA must be set up as service resources. If you decide to cluster the DVA, it runs along with the POA in the same service resource. A service resource for a GroupWise agent must include the following information:

- **Name:** The name by which the agent service will be listed in the resource group (for example, MTA Service or POA Service).
- **Possible Owners:** The list of nodes in the cluster to which the GroupWise agent can fail over (the same as the possible owners of the resource group to which the agent service belongs).
- **Resource Dependencies:** Other resources in the resource group that must be online before the GroupWise agent can start on a new node (for example, the Group IP Address resource and the Physical Disk resource where the domain or post office directory is located).

AGENT CLUSTERING WORKSHEET

Under Item 3: MTA Service Resource, specify the MTA service resource name and list any possible resource dependencies.

Under Item 6: POA Service Resource, specify the POA service resource name and list any possible resource dependencies.

21.8.4 Planning the Windows Agent Installation

Aside from the cluster-specific issues discussed in the preceding sections, the considerations involved in planning to install the GroupWise Windows agents are the same in a Windows cluster as in any other environment. Review "Planning the GroupWise Agents", then print and fill out the "GroupWise Agent Installation Worksheet" in "Installing GroupWise Agents" in the *GroupWise 2012 Installation Guide* for each location where you will install the Windows MTA and/or POA and DVA.

GROUPWISE AGENT INSTALLATION WORKSHEET

Under Installation Path, take into account your decision based on "Deciding Where to Install the Agent Software" on page 167.

Under Windows Installation Options, mark *Install as Windows Services*.

Under Domain Information and Post Office Information, refer to the Domain and Post Office Worksheets you filled out in Section 21.2, "Planning a New Clustered Domain," on page 161 and Section 21.3, "Planning a New Clustered Post Office," on page 161, and to the Network Address Worksheet you completed during "Planning Cluster-Unique Port Numbers for Agents in the Cluster" on page 165.

IMPORTANT: Do not install the Windows agent software until you are instructed to do so in Chapter 22, “Setting Up a Domain and Post Office in a Windows Cluster,” on page 175.

Skip to Chapter 22, “Setting Up a Domain and Post Office in a Windows Cluster,” on page 175.

21.9 GroupWise Clustering Worksheets

21.9.1 System Clustering Worksheet

Item	Explanation
1) Cluster Name:	Record the name of the name of your Windows cluster. For more information, see Section 21.1, “Setting Up Your Windows Cluster,” on page 160.
2) Nodes in Cluster:	List the servers that are part of the cluster that you set up for your GroupWise system. For more information, see Section 21.1, “Setting Up Your Windows Cluster,” on page 160.
3) Resources for GroupWise Administration: ConsoleOne: Shared disk: Possible owners: Software Distribution Directory: Shared disk: Possible owners:	List any shared locations that you want to set up for ConsoleOne or the software distribution directory. For more information, see Section 21.6, “Planning Shared Administrative Resources,” on page 163.
4) Resource Group for Domain: Network name: IP address: Physical disk: File share: MTA service: Possible owners Post Office in Same Resource Group as Domain? ◆ Yes ◆ No	Specify the information for the domain resource group. For more information, see Section 21.2, “Planning a New Clustered Domain,” on page 161.

Item	Explanation
5) Resource Group for Post Office: Network name: IP address: Physical disk: File share: POA service: Possible owners:	Specify the information for the post office resource group. For more information, see Section 21.3, “Planning a New Clustered Post Office,” on page 161.
6) IP Address Resolution Methods: ◆ eDirectory ◆ hosts file ◆ DNS	Mark the short name address resolution methods you want to implement to ensure that the UNC paths stored in ConsoleOne with network names can be successfully resolved into physical network addresses. For more information, see Section 21.7, “Ensuring Successful Name Resolution for GroupWise Resource Groups,” on page 163.
7) Domain Name: Domain Directory:	Specify a unique name for the domain. Specify the directory where you want to create the new domain. For more information, see Section 21.2, “Planning a New Clustered Domain,” on page 161.
8) Post Office Name: Post Office Directory:	Specify a unique name for the post office. Specify the directory where you want to create the post office. For more information, see Section 21.3, “Planning a New Clustered Post Office,” on page 161.
9) Document Storage Area Location: ◆ At the post office ◆ Outside the post office ◆ Separate post office Document Storage Area Access ◆ POA /user startup switch setting ◆ POA /password startup switch setting	If you need a library for a clustered post office, mark where you want to create its document storage area and provide a directory if necessary. For more information, see Section 21.4, “Planning a New Library for a Clustered Post Office,” on page 162.

21.9.2 Network Address Worksheet

Domain Information

Domain	MTA IP Address	MTA MTP Port	MTA HTTP Port

Post Office Information

Post Office	POA IP Address	POA C/S Port	POA MTP Port	POA HTTP Port	DVA HTTP Port

GWIA Information

GWIA	GWIA IP Address	MTA MTP Port	MTA HTTP Port	GWIA HTTP Port
GWIA Domain MTA				N/A
GWIA	(same)	N/A	N/A	

21.9.3 Agent Clustering Worksheet

Item	Explanation
1) Agent installation location: ◆ Shared disk with domain or post office ◆ Each node in the cluster Consolidate startup files?	Mark the location where you will install the agent software. If necessary, specify the location where you will store agent startup files on the same shared disk with the domain or post office. For more information, see “Deciding Where to Install the Agent Software” on page 167.
2) Domain Name: Domain Directory:	Transfer this information from the System Clustering Worksheet (item 6).
3) MTA Service Resource: Service name: Possible owners: Resource dependencies:	List other nodes in the cluster where the domain resource group can fail over and any resources that must be online before the MTA can start. For more information, see “Planning the Agent Services” on page 169.

Item	Explanation
4) MTA Network Information: MTA IP address: MTA message transfer port: MTA HTTP port:	Gather the MTA network address information from the “Network Address Worksheet” on page 171. For more information, see “Planning Cluster-Unique Port Numbers for Agents in the Cluster” on page 165.
5) Post Office Name: Post Office Directory:	Transfer this information from the System Clustering Worksheet (item 7).
6) POA Service Resource: Service name: Possible owners: Resource dependencies:	List other nodes in the cluster where post office resource group can fail over and any resources that must be online before the POA can start. For more information, see “Planning the Agent Services” on page 169.
7) POA Network Information: POA IP address POA client/server port POA message transfer port POA HTTP port	Gather the POA network address information from the “Network Address Worksheet” on page 171. For more information, see “Planning Cluster-Unique Port Numbers for Agents in the Cluster” on page 165.
8) DVA Network Information: DVA IP address DVA HTTP port	Gather the DVA network address information from the “Network Address Worksheet” on page 171. For more information, see “Planning Cluster-Unique Port Numbers for Agents in the Cluster” on page 165.

22 Setting Up a Domain and Post Office in a Windows Cluster

You should have already reviewed "Planning GroupWise in a Windows Cluster" on page 159 and filled out the "System Clustering Worksheet" on page 170, the "Network Address Worksheet" on page 171, and the "Agent Clustering Worksheet" on page 172. You are now ready to complete the following tasks to set up GroupWise in your Windows cluster:

22.1 Preparing the Cluster for GroupWise

After you have set up your Windows cluster and become familiar with its functioning, as described in Chapter 20, "Introduction to GroupWise 2012 and Clustering on Windows," on page 157, complete the following tasks to prepare the cluster for your GroupWise system:

22.1.1 Creating GroupWise Resource Groups

Create the needed domain and post office resource groups in your Windows cluster (System Clustering Worksheet items 3 and 4), as planned in Section 21.2, "Planning a New Clustered Domain," on page 161 and Section 21.3, "Planning a New Clustered Post Office," on page 161.

22.1.2 Creating Agent Service Resources

Within each GroupWise resource group, create the MTA or POA service resource (Agent Clustering Worksheet items 3 and 6), as planned in "Planning the Agent Services" on page 169.

22.1.3 Configuring Short Name Resolution

To ensure that GroupWise resource groups are always locatable on the network, configure the short name resolution methods that you want to rely on for your clustered GroupWise system (System Clustering Worksheet item 9), as planned in Section 21.7, "Ensuring Successful Name Resolution for GroupWise Resource Groups," on page 163.

- "eDirectory" on page 176
- "Hosts Files" on page 176
- "DNS" on page 176

After configuring your selected short name resolution methods, continue with the following tasks as needed:

- Section 22.2, "Setting Up a New GroupWise System in a Windows Cluster," on page 177
- Section 22.3, "Creating a New Secondary Domain in a Windows Cluster," on page 177
- Section 22.4, "Creating a New Post Office in a Windows Cluster," on page 179

eDirectory

ConsoleOne uses Novell eDirectory to resolve the UNC path of a domain or post office directory into its network name in the cluster. For example, on the workstation where you run ConsoleOne, you need to map a drive to the location of a domain directory by using the network name of the domain resource group so that ConsoleOne can access the domain database no matter which node in the cluster it is active on.

Hosts Files

Because each GroupWise resource group has been associated with a network name, you should add lines for the new network names to the `c:\Windows\System32\drivers\etc\hosts` file as needed This should only be done on the administrator's workstation.

The lines you add to a hosts file might look similar to the following example:

Syntax:

```
IP_address network_name.context
```

Remember that *network_name* represents the name of the virtual server, which remains unchanged regardless of which node is currently active.

Example:

```
172.16.5.81   gwcluster.novell.com
```

When specifying the lines in the hosts file, use System Clustering Worksheet item 7 or 8 for each *IP_address* and *network_name* where a domain or post office resides. Use System Clustering Worksheet item 1 for *network_name*. Use your Internet domain name for *context*.

DNS

Because each GroupWise resource group has been associated with a virtual network name, you should add all your new network names to DNS.

22.2 Setting Up a New GroupWise System in a Windows Cluster

The GroupWise Installation Advisor walks you through setting up the primary domain and an initial post office in the primary domain. You might be creating your primary domain and initial post office in the same resource group or in two different resource groups. After you have created the primary domain and initial post office and installed the GroupWise agents, you can create additional secondary domains and post offices in the cluster as needed.

To set up the primary domain and initial post office for a new GroupWise system in a Windows cluster:

1. (Conditional) If necessary, map a drive to each GroupWise administration shared disk (System Clustering Worksheet item 3).
2. Map a drive to the shared disk of the domain resource group (System Clustering Worksheet item 6) and, if needed, to the shared disk of the post office resource group (System Clustering Worksheet item 7), where the primary domain and the initial post office for your new GroupWise system will be created.
3. Manually create the domain directory (System Clustering Worksheet item 6) and the post office directory (System Clustering Worksheet item 7).

   This step is not required, but the following step is easier if the directory already exists.
4. Run the GroupWise Installation Advisor to set up your initial GroupWise system, following the steps provided in "Windows: Setting Up a Basic GroupWise System" in "Installing a Basic GroupWise System" in the *GroupWise 2012 Installation Guide*. Keep in mind the following cluster-specific details:
   - When you specify the ConsoleOne directory and the software distribution directory, be sure to browse to each location through the shared disk accessed in Step 1 above.
   - When you specify the domain directory and post office directory, be sure to browse through the shared disk accessed in Step 2 to select the directory created in Step 3 above.
   - For the post office link type, select *TCP/IP Link*.
   - When you provide the MTA and POA network address information, use the Agent Clustering Worksheet that you filled out in Section 21.8, "Deciding How to Install and Configure the Agents in a Windows Cluster," on page 165. The information you provide is used to configure the MTA and POA objects in the domain and post office even though you have not yet installed the agent software.
   - Do not create users in the post office at this time.
   - In the Summary dialog box, the domain directory and post office directory that you browsed to should display as UNC paths using the network name of the GroupWise resource group, not the name of a specific node in the cluster.
5. When you have finished creating the primary domain and the initial post office, continue with installing the GroupWise Agents, starting with Step 5 in "Installing the Agent Software in a Windows Cluster" on page 180.

   The GroupWise Installation Advisor starts the Agent Installation program for you.

22.3 Creating a New Secondary Domain in a Windows Cluster

After you have set up the primary domain and initial post office, as described in Section 22.2, "Setting Up a New GroupWise System in a Windows Cluster," on page 177, you can create additional secondary domains as needed.

To create a new secondary domain in a Windows cluster:

1. Create a domain resource group for the new domain, as described in "Creating GroupWise Resource Groups" on page 175.
2. Create an MTA service resource for the domain's MTA, as described in "Creating Agent Service Resources" on page 175.
3. Map a drive to the shared disk of the domain resource group (System Clustering Worksheet item 7) where the new secondary domain will be created.
4. Manually create the domain directory (System Clustering Worksheet item 7).

   This step is not required, but Step 7 is easier if the domain directory already exists.
5. (Conditional) If you selected the same shared disk with the domain as the agent installation location (Agent Clustering Worksheet item 1), create the `drive:\Program Files\Novell\GroupWise Server\Agents` directory on the drive accessed in Step 3.

   or

   If you selected `c:\Program Files\Novell\GroupWise Server\Agents` on each node in the cluster, decide which node to install the agents to first.
6. In ConsoleOne, connect to the primary domain in your GroupWise system, as described in "Connecting to a Domain" in "Domains" in the *GroupWise 2012 Administration Guide*.
7. Create the new domain, following the steps provided in "Creating the New Domain" in "Domains" in the *GroupWise 2012 Administration Guide*. Keep in mind the following cluster-specific details:
   - Use the Domain Worksheet you filled out in Section 21.2, "Planning a New Clustered Domain," on page 161 to fill in the fields on the Create GroupWise Domain page.
   - In the *Domain Database Location* field, be sure to browse through the drive you accessed in Step 3 to the domain directory you created in Step 4 above.
   - In the *Link to Domain* field, link the new domain to the primary domain of your GroupWise system.
   - The *Configure Link* option is selected by default. Select *TCP/IP Link to the Other Domain*. Refer to the Agent Clustering Worksheet that you filled out in "Planning Cluster-Unique Port Numbers for Agents in the Cluster" on page 165 for the resource group IP address and cluster-unique port numbers that you need to specify in order to configure the link.
8. Use the Link Configuration tool to change the links from the new domain to all other domains in the cluster to direct TCP/IP links, following the steps provided in "Changing the Link Protocol between Domains to TCP/IP" in "Message Transfer Agent" in the *GroupWise 2012 Administration Guide*.

   Although a complete mesh link configuration is the most efficient, it might not be feasible in all situations. Set up as many direct TCP/IP links as possible for best MTA performance in the cluster.
9. Make sure you are still connected to the primary domain.
10. Rebuild the domain database for the new domain, following the steps provided in "Rebuilding Domain or Post Office Databases" in "Databases" in the *GroupWise 2012 Administration Guide*. Be sure to browse to the database location (System Clustering Worksheet item 7) through the shared disk you accessed in Step 3 to the domain directory you created in Step 4 above.

    The database rebuild is necessary in order to transfer the MTA configuration information and the domain link information into the secondary domain database, because the MTA for the new secondary domain is not yet running.
11. Continue with Creating a New Post Office in a Windows Cluster.

22.4 Creating a New Post Office in a Windows Cluster

You can create a new post office in the same resource group where its domain is located or in a separate resource group. If the post office and its domain are in the same resource group, they fail over together. If they are in separate resource groups, they fail over separately.

To create a new post office in a Windows cluster:

1. (Conditional) If you selected *Yes* for *Post Office in Same Resource Group as Domain?* (under System Clustering Worksheet item 4), map a drive to the shared disk of the domain resource group.

   or

   Map a drive to the shared disk of the post office resource group (System Clustering Worksheet item 5).

2. Manually create the post office directory (System Clustering Worksheet item 8).

   This step is not required, but Step 4 is easier if the post office directory already exists.

3. In ConsoleOne, connect to the GroupWise domain where you want to create the new post office, as described in "Connecting to a Domain" in "Domains" in the *GroupWise 2012 Administration Guide*.

4. Create the new post office, following the steps provided in "Creating the New Post Office" in "Post Offices" in the *GroupWise 2012 Administration Guide*. Keep in mind the following cluster-specific details:

   - Use the Post Office Worksheet you filled out in Section 21.3, "Planning a New Clustered Post Office," on page 161 to fill in the fields on the Create GroupWise Post Office page.
   - In the *Post Office Database Location* field, be sure to browse through the shared disk you accessed in Step 1 to the post office directory you created in Step 2 above.
   - If you want to create a library at the post office (System Clustering Worksheet item 9), select *Create Library*.
   - The *Configure Link* option is selected by default. Select *TCP/IP Link from Domain to New Post Office*. Refer to the Agent Clustering Worksheet that you filled in during "Planning Cluster-Unique Port Numbers for Agents in the Cluster" on page 165 for the resource group IP address and cluster-unique port numbers that you need to specify in order to configure the link.

5. In ConsoleOne, right-click the new Post Office object, then click *Properties*.

6. Click *GroupWise* > *Post Office Settings*, then in the *Access Mode* field, select *Client/Server Only*.

7. Right-click the new POA object, then click *Properties*.

   On the POA Agent Settings and Scheduled Events pages, you might want to specify unique times for the following POA activities to prevent multiple POAs from performing the same activities on the same node at the same time during a failover situation:

   - Start User Upkeep
   - Generate Address Book for Remote
   - Start QuickFinder Indexing
   - Mailbox/Library Maintenance Event

   For more information about these repetitive POA activities, see "Performing Nightly User Upkeep", "Regulating Indexing", and "Scheduling Database Maintenance" in "Post Office Agent" in the *GroupWise 2012 Administration Guide*.

8. Make sure you are still connected to the domain that owns the new post office.

9 Rebuild the post office database for the new post office, following the steps provided in "Rebuilding Domain or Post Office Databases" in "Databases" in the *GroupWise 2012 Administration Guide*.

   Be sure to browse to the database location (System Clustering Worksheet item 7) through the shared disk you accessed in Step 1 to the post office directory you created in Step 2 above.

   The database rebuild is necessary in order to transfer the POA configuration information and the post office link information into the post office database, because the POA for the new post office is not yet running.

10 (Conditional) If you want to create a library (System Clustering Worksheet item 9) for the new clustered post office, follow the steps in "Setting Up a Basic Library" or "Setting Up a Full-Service Library" in "Libraries and Documents" in the *GroupWise 2012 Administration Guide*, after you have completely finished setting up the new clustered post office.

11 Continue with Installing and Configuring the MTA and the POA in a Windows Cluster.

22.5 Installing and Configuring the MTA and the POA in a Windows Cluster

After you have created a new domain and/or post office, you are ready to install and configure the GroupWise agents. Complete all the tasks below if you are setting up a new GroupWise system or if you have created a new GroupWise resource group where you want to install the agent software:

- Section 22.5.1, "Installing the Agent Software in a Windows Cluster," on page 180
- Section 22.5.2, "Editing Clustered Agent Startup Files," on page 181
- Section 22.5.3, "Setting Up New Instances of the Agents without Installing the Agent Software," on page 182

Under some circumstances, the agent software has already been installed in the cluster and you simply need to create a new startup file specific to the new domain or post office. For example:

- You have created a new domain and/or post office in a GroupWise resource group where the agent software is already installed in the *drive*:\Program Files\Novell\GroupWise Server\Agents directory for the resource group.
- In your GroupWise system, the agent software is already installed to the c:\Program Files\Novell\GroupWise Server\Agents directory on each node in the cluster.

In these circumstances, follow the instructions in "Setting Up New Instances of the Agents without Installing the Agent Software" on page 182 instead of completing the tasks listed above.

22.5.1 Installing the Agent Software in a Windows Cluster

To install the MTA, POA, and DVA:

1 Map a drive to the shared disk of the domain resource group (Agent Clustering Worksheet item 3) or the post office resource group (Agent Clustering Worksheet item 6).

2 Map a drive to c:\ on the first node in the cluster where you will set up the agents as Windows services (System Clustering Worksheet item 2).

3 (Conditional) If you plan to install the agent software to the shared disk of the domain or post office resource group (under Agent Clustering Worksheet item 1), create the *drive*:\Program Files\Novell\GroupWise Server\Agents directory on the shared disk accessed in Step 1.

   or

If you plan to install the agent software to each node in the cluster, create the `c:\Program Files\Novell\GroupWise Server\Agents` directory on the drive accessed in Step 2.

4 Start the Agent Installation program, following the steps provided in "Installing the Windows Agent Software" in "Installing GroupWise Agents" in the *GroupWise 2012 Installation Guide*.

5 Install the Windows agents, keeping in mind the following cluster-specific details:

- Use the Windows Agent Clustering Worksheet that you filled out in "Planning the Windows Agent Installation" on page 169 to fill in the fields during the agent installation process.
- On the Installation Path page, be sure to browse through the mapped drive to the directory you created in Step 3 above. Be sure that *Install as Windows Services* is selected.
- On the Domains / Post Offices page, click *Add* for each domain and post office that the agents will service. In the *Path to Database* field, be sure to browse through the drive you mapped in Step 1 above to the domain directory or the post office directory on the shared disk.
- On the Installation Complete page, do not select *Launch GroupWise Agents Now*.

6 (Conditional) If you need to install the agent software to `c:\Program Files\Novell\GroupWise Server\Agents` on each node in the cluster, repeat Step 4 and Step 5, mapping a drive to each node in the cluster.

or

If you installed the agent software to a shared disk and need only to set up the agents as Windows services on each node, repeat Step 4 and Step 5, mapping drives to new nodes as needed. On the Installation Options page, select only the *Install as Windows Services* option to speed up the installation process for each node.

7 (Conditional) If you installed the agent software to each node and you selected *Yes* for *Consolidate Startup Files?* (under Agent Clustering Worksheet item 1), copy one complete set of agent startup files to the planned location on the shared disk, then delete all agent startup files from the `c:\Program Files\Novell\GroupWise Server\Agents` directories on the nodes to avoid future confusion.

8 Continue with Editing Clustered Agent Startup Files.

22.5.2 Editing Clustered Agent Startup Files

By default, the Agent Installation program creates agent startup files in the agent installation directory. Each MTA startup file is named after the domain it services, with a `.mta` extension. Each POA startup file is named after the post office it services, with a `.poa` extension. The DVA startup file is named `gwdva.dva`.

Because you mapped a drive to the shared disk of the GroupWise resource group by using the physical disk and file share information from the resource group, the setting for the MTA /home startup switch and the POA /home startup switch are always correct, no matter which node in the cluster the domain and post office are currently active on. The same holds true for the DVA /home startup switch.

One manual modification of POA startup files is required for robust functionality in a Windows cluster. Uncomment the /ip startup switch and provide the IP address of the post office resource group (Agent Clustering Worksheet item 7). This information is available to the POA in its eDirectory object properties. However, in some failover situations, the POA reconnects to the MTA more quickly when the information is immediately available to the POA in its startup file.

If the POA needs to access a remote document storage area that is outside the cluster, add the /user and /password startup switches (under System Clustering Worksheet item 9) in order to provide a user name and password that the POA can use to access the server where the document storage area resides. As an alternative to startup switches, you can assign the POA object all rights except Supervisor and Access control, as long as the remote document storage area is located in the same tree with the post office.

Skip to Section 22.6, "Testing Your Clustered GroupWise System," on page 182.

22.5.3 Setting Up New Instances of the Agents without Installing the Agent Software

To set up new instances of the agents without installing the agent software, you simply create new startup files. Each MTA startup file is named after the domain it services, with a `.mta` extension. Each POA startup file is named after the post office it services, with a `.poa` extension.

NOTE: You cannot set up multiple instances of the DVA in the same post office resource, because the name of the DVA service is always GWDVA. Therefore, there can be only one DVA in a post office resource.

If the existing agent software is located in the *`drive`*`:\Program Files\Novell\GroupWise Server\Agents` directory of a shared disk with a domain or post office, the startup files are also located there. If the existing agent software is located in the `c:\Program Files\Novell\GroupWise Server\Agents` directory on each node in the cluster, the startup files might be located there, or they might be located on the shared disk with the domain or post office.

To create a new startup file without installing the agent software:

1. Make a copy of an existing startup file and name it after the domain or post office that will be serviced by the new instance of the agent.
2. Edit the setting of the /home startup switch to point to the location of the new domain directory or post office directory. Be careful to maintain the syntax of the original line, using the physical disk and file share provided in the GroupWise resource group.
3. Scroll down through the startup file looking for other active (not commented out) startup switches, then modify them as needed for the new instance of the agent.
4. Save the new startup file.
5. Continue with Testing Your Clustered GroupWise System.

22.6 Testing Your Clustered GroupWise System

After you have configured the GroupWise resource group, you can test the failover and failback functionality by bringing the GroupWise resource group online and taking it offline again.

Continue with Managing Your Clustered GroupWise System.

22.7 Managing Your Clustered GroupWise System

After you have set up a basic clustered GroupWise system, you should consider some long-term management issues.

22.7.1 Updating GroupWise Objects with Cluster-Specific Descriptions

After setting up your clustered GroupWise system, while the cluster-specific information is fresh in your mind, you should record the cluster-specific information as part of the GroupWise objects in ConsoleOne so that you can easily refer to it later. Be sure to keep the information recorded in the GroupWise objects up-to-date if the configuration of your system changes.

Recording Cluster-Specific Information for a Domain and Its MTA

To permanently record important cluster-specific information for the domain:

1. In ConsoleOne, browse to and right-click the Domain object, then click *Properties*.
2. In the *Description* field of the domain Identification page, provide a cluster-specific description of the domain, including the resource group IP address and the cluster-unique port numbers used by its MTA.
3. Click *OK* to save the domain description.
4. Select the Domain object to display its contents.
5. Right-click the MTA object, then click *Properties*.
6. In the *Description* field of the MTA Identification page, record the domain resource group IP address and the cluster-unique port numbers used by the MTA.

   This information appears on the MTA console, no matter which node in the cluster it is currently running on.
7. Click *OK* to save the MTA description.
8. Continue with Recording Cluster-Specific Information for a Post Office and Its POA.

Recording Cluster-Specific Information for a Post Office and Its POA

To permanently record important cluster-specific information for a post office:

1. In ConsoleOne, browse to and right-click the Post Office object, then click *Properties*.
2. In the *Description* field of the post office Identification page, provide a cluster-specific description of the post office, including the resource group IP address and the cluster-unique port numbers used by its POA.
3. Click *OK* to save the post office description.
4. Select the Post Office object to display its contents.
5. Right-click the POA object, then click *Properties*.

6 In the *Description* field of the POA Identification page, record the post office resource group IP address and the cluster-unique port numbers used by the POA.

   This information appears on the POA console, no matter which node in the cluster it is currently running on.

7 Click *OK* to save the POA description.

8 If necessary, continue with "Recording Cluster-Specific Information for a Software Distribution Directory" on page 184.

   or

   Skip to "Knowing What to Expect in MTA and POA Failover Situations" on page 184.

Recording Cluster-Specific Information for a Software Distribution Directory

To permanently record important cluster-specific information about a software distribution directory located on a shared disk:

1 In ConsoleOne, click *Tools* > *System Operations* > *Software Directory Management*.

2 Select the software distribution directory, then click *Edit*.

3 In the *Description* field, record the IP address of the cluster resource where the software distribution directory resides.

4 Click *OK*, then click *Close* to save the software distribution directory description.

5 Continue with Knowing What to Expect in MTA and POA Failover Situations.

22.7.2 Knowing What to Expect in MTA and POA Failover Situations

In a failover situation, the agents might need to perform some database repair as they start on the new node. The time required depends on the size of the databases involved.

Typically, the POA returns to full functionality faster than the MTA. This benefits GroupWise client users, who can reconnect to their mailboxes very quickly and probably do not notice if messages to users in other post offices are not delivered immediately. The only time a user needs to restart the GroupWise client is if he or she was actually in the process of sending a message when the POA went down. Notify can continue running even if the connection to the POA becomes unavailable because it reconnects automatically when the POA is again available.

The MTA typically takes some time reestablishing the links to its post offices, other domains, and gateways, but this situation usually resolves itself in a few minutes without administrator intervention. If it does not, you can manually restart the MTA to speed up the process.

The DVA must reestablish its HTTP connections with one or more POAs and WebAccess Applications. Typically, this occurs quite quickly.

In comparison to failover, manual migration typically takes longer because the agents methodically terminate their threads and close their databases as part of their normal shutdown procedure. However, as a result, no database repair is required when the agents start up again in their new location.

Continue with What's Next.

22.8 What's Next

Now that you have at least one GroupWise domain and post office up and running in your Windows cluster, you are ready to proceed with the rest of your GroupWise system setup by:

- Adding users to post offices.

  See "Users" in the *GroupWise 2012 Administration Guide*.

- Setting up the GroupWise client software and helping users to get started using it.

  See "Client" in the *GroupWise 2012 Administration Guide*. Also see the *GroupWise 2012 Windows Client User Guide*.

- Connecting your clustered GroupWise system to the Internet.

  See Chapter 23, "Implementing the Internet Agent in a Windows Cluster," on page 187.

- Monitoring the status of your clustered GroupWise system from your Web browser.

  See Chapter 25, "Monitoring a GroupWise System in a Windows Cluster," on page 199.

- Backing up your clustered GroupWise system.

  See Chapter 26, "Backing Up a GroupWise System in a Windows Cluster," on page 201.

23 Implementing the Internet Agent in a Windows Cluster

You should already have set up at least a basic GroupWise system, as described in Chapter 21, "Planning GroupWise in a Windows Cluster," on page 159 and Chapter 22, "Setting Up a Domain and Post Office in a Windows Cluster," on page 175. As part of this process, the "System Clustering Worksheet" on page 170 and the "Network Address Worksheet" on page 171 were filled out. If you do not have access to the filled-out worksheets, print the worksheets now and fill in the clustering and network address information as it currently exists on your system. You need this information as you implement the Internet Agent (GWIA) in a cluster.

23.1 Planning the Internet Agent in a Windows Cluster

A main system configuration difference between a GroupWise system in a clustering environment and a GroupWise system in a regular environment is that you need to create a separate domain to house the GWIA. The GWIA is faster and more stable when it runs on the same server with its domain. In a cluster, creating a separate domain for the GWIA ensures that the GWIA and its domain always fail over together.

Internet Agent Clustering Worksheet lists all the information you need as you set up the GWIA in a Windows cluster. You should print the worksheet and fill it out as you complete the tasks listed below:

23.1.1 Planning a Domain for the GWIA

The considerations involved in planning a domain for the GWIA are much the same as planning any other domain. In preparation, review "Planning a New Domain", then print and fill out the "New Domain Summary Sheet" in "Domains" in the *GroupWise 2012 Administration Guide*.

Keep in mind the following cluster-specific details:

- When you specify the location for the domain directory on the Domain Worksheet, include the shared disk where you want the domain directory to be located.
- Do not concern yourself with the GroupWise agent information on the Domain Worksheet. You can stop with item 10. You will plan the MTA installation later.

When you have completed the Domain Worksheet, transfer the key information from the Domain Worksheet to the Internet Agent Clustering Worksheet.

INTERNET AGENT CLUSTERING WORKSHEET

Under Item 1: Resource Group for the GWIA, transfer the disk drive to the Internet Agent Clustering Worksheet.

Under Item 2: GWIA Domain Name, transfer the domain name and directory to the Internet Agent Clustering Worksheet.

23.1.2 Planning the GWIA Resource Group

The GWIA resource group is similar to the GroupWise resource groups you have already set up, as described in "Planning GroupWise Resource Groups" on page 162 and "Creating GroupWise Resource Groups" on page 175. The GWIA resource group contains a domain whose only role is to connect the GWIA into your clustered GroupWise system. It also contains two agent service resources, one for the MTA that services the domain and one for the GWIA.

INTERNET AGENT CLUSTERING WORKSHEET

Under Item 1: Resource Group for GWIA, specify the network name and other required information for the GWIA resource group.

To ensure successful short name resolution, add entries for the GWIA network name to support your preferred methods of short name resolution, as described in "Configuring Short Name Resolution" on page 176.

23.1.3 Planning Cluster-Unique Port Numbers for the GWIA and Its MTA

As with the MTA, POA, and DVA, the GWIA needs cluster-unique port numbers. As part of planning to install the MTA and POA, you should already have determined the resource group IP address and cluster-unique port numbers for the GWIA and its MTA as you filled out the "Network Address Worksheet" on page 171. If you do not have a filled-out copy of this worksheet for your system, print it now and fill in current system information.

INTERNET AGENT CLUSTERING WORKSHEET

Under Item 5: MTA Network Information, transfer the resource group IP address and cluster-unique port numbers from the GWIA section of the Network Address Worksheet to the Internet Agent Clustering Worksheet.

Under Item 7: GWIA Network Information, transfer the resource group IP address (the same as for its MTA) and the cluster-unique GWIA port number from the GWIA section of the Network Address Worksheet to the Internet Agent Clustering Worksheet.

23.1.4 Preparing Your Firewall for the GWIA

The GWIA receives incoming messages on the IP address of the GWIA resource group. Your firewall configuration must be modified to allow inbound TCP/IP traffic from the Internet to the GWIA IP address on the following standard ports:

Protocol	Standard Port
IMAP4	143
LDAP	389
POP3	110
SMTP	25

By default, the GWIA sends outgoing messages on the IP address of the node where it is running. If you decide to use this default configuration, your firewall must be configured to allow outbound TCP/IP traffic from all nodes on the GWIA resource group's possible owners list.

If the GWIA has a large number of nodes in its possible owners list, you can configure the GWIA to send outgoing messages to a relay host, which then sends them out through the firewall using its own IP address rather than the IP address of the particular node where the GWIA is running. This reduces the amount of modification to your firewall required to set up the GWIA. However, if the relay host goes down, all outgoing messages are delayed.

As another alternative, you can configure the GWIA to use its resource group IP address for sending as well as receiving messages. Setup instructions for this configuration are provided in "Forcing Use of the GWIA Resource Group IP Address" on page 193, which you can complete after installing the GWIA.

In preparation for installing the GWIA, configure your firewall as needed to handle the GWIA's use of node and resource group IP addresses when sending and receiving messages.

23.1.5 Deciding Where to Install the GWIA and Its MTA

The default GWIA installation directory is `c:\Program Files\Novell\GroupWise Server\GWIA`. As with the MTA, the POA, and the DVA, you can choose to install the GWIA and its MTA to each node in the cluster or to the shared disk of the GWIA resource group. For a discussion of these alternatives, see "Deciding Where to Install the Agent Software" on page 167, which describes the issues in the context of planning MTA and POA installations. As with the MTA, POA, and DVA, the GWIA and its MTA must be installed as Windows services.

INTERNET AGENT CLUSTERING WORKSHEET

Under Item 4: MTA Installation Location and Item 6: GWIA Installation Location, mark whether you will install the GWIA and its MTA to the shared disk of the GWIA resource group or to each node in the cluster. If necessary, specify where the MTA startup file and the GWIA configuration file (`gwia.cfg`) will be stored.

23.1.6 Planning the MTA Installation

Follow the instructions in "Planning the Windows Agent Installation" on page 169 to plan the MTA installation for the GWIA domain, then return to this point. After you follow the instructions, you will have a filled-out Windows Agent Worksheet to use when you install the MTA.

IMPORTANT: Do not install the Windows MTA until you are instructed to do so in Section 23.2, "Setting Up the Internet Agent in a Windows Cluster," on page 190.

23.1.7 Planning the GWIA Installation

Aside from the cluster-specific issues discussed in the preceding sections, the considerations involved in planning to install the GWIA are the same in a Windows cluster as for any other environment. Review "Windows: Installing the GWIA Software", then print and fill out the "GroupWise Internet Agent Installation Worksheet" in "Installing the GroupWise Internet Agent" in the *GroupWise 2012 Installation Guide*. You need this information as you install the GWIA in your cluster.

IMPORTANT: Do not install the GWIA software until you are instructed to do so in Section 23.2, "Setting Up the Internet Agent in a Windows Cluster," on page 190.

23.2 Setting Up the Internet Agent in a Windows Cluster

You should already have reviewed Section 23.1, "Planning the Internet Agent in a Windows Cluster," on page 187 and filled out Section 23.4, "Internet Agent Clustering Worksheet," on page 195. You are now ready to complete the following tasks to set up the GWIA in a Windows cluster:

- Section 23.2.1, "Setting Up the GWIA Resource Group," on page 190
- Section 23.2.2, "Creating a Domain for the GWIA," on page 191
- Section 23.2.3, "Installing the MTA for the GWIA Domain," on page 191
- Section 23.2.4, "Installing and Configuring the GWIA in a Windows Cluster," on page 191
- Section 23.2.5, "Testing the Clustered GWIA," on page 194

23.2.1 Setting Up the GWIA Resource Group

1. Create the GWIA resource group and agent services resources (Internet Agent Clustering Worksheet item 1), as planned in "Planning the GWIA Resource Group" on page 188.
2. To ensure successful short name resolution, add entries for the GWIA network name to support your preferred methods of short name resolution, as described in "Configuring Short Name Resolution" on page 176.

3 To ensure that the GWIA has incoming and outgoing access to the Internet, make sure your firewall is properly configured, as described in "Preparing Your Firewall for the GWIA" on page 189.

4 Continue with Creating a Domain for the GWIA.

23.2.2 Creating a Domain for the GWIA

The GWIA domain will be a secondary domain. To create it, follow the instructions in Section 22.3, "Creating a New Secondary Domain in a Windows Cluster," on page 177, taking your information from the Internet Agent Clustering Worksheet, rather than the System Clustering Worksheet, then return to this point.

Do not create any post offices in the GWIA domain.

Continue with Installing the MTA for the GWIA Domain.

23.2.3 Installing the MTA for the GWIA Domain

The MTA for the GWIA domain can be installed just like any other MTA in your clustered GroupWise system. Follow the instructions in "Installing the Agent Software in a Windows Cluster" on page 180, then return to this point.

You do not need to edit the MTA startup file.

Continue with Installing and Configuring the GWIA in a Windows Cluster.

23.2.4 Installing and Configuring the GWIA in a Windows Cluster

After you have created a domain for the GWIA and installed the MTA for that domain, you are ready to install and configure the GWIA.

Installing the GWIA Software in a Windows Cluster

1 Map a drive to the shared disk of the GWIA resource group (Internet Agent Clustering Worksheet item 1).

2 Map a drive to `c:\` on the first node in the cluster where you will set up the GWIA as a Windows service (System Clustering Worksheet item 2).

3 (Conditional) If you plan to install the GWIA software to the shared disk of the GWIA resource group (Internet Agent Clustering Worksheet item 6), create the *`drive`*`:\Progarm Files\Novell\Groupwise Server\GWIA` directory on the shared disk accessed in Step 1.

or

If you plan to install the GWIA software to each node in the cluster, create the `c:\Program Files\Novell\GroupWise Server\GWIA` directory on the drive accessed in Step 2).

4 Start the GWIA Installation program, following the steps provided in "Windows: Installing the GWIA Software" in "Installing the GroupWise Internet Agent" in the *GroupWise 2012 Installation Guide*.

5 Install the Windows GWIA, keeping in mind the following cluster-specific details:

- Use the Windows Internet Agent Clustering Worksheet that you filled out in Section 23.1, "Planning the Internet Agent in a Windows Cluster," on page 187 to fill in the fields during the GWIA installation process.
- On the Installation Path page, be sure to browse through the mapped drive to the directory you created in Step 3 above. Be sure that *Run Internet Agent as a Windows Service* is selected.
- On the GroupWise Domain page, be sure to browse through the drive you mapped in Step 1 to the domain directory on the shared disk.
- On the Post Installation Task List page, deselect *Launch Internet Agent Now* so that the Installation program does not start the GWIA after installation is complete.

6 Repeat Step 4 and Step 5, mapping a drive to each node in the cluster.

Even if you installed the GWIA software to a shared disk, you need to repeat the installation process for each node so that the GWIA is set up as a Windows service on each node.

7 (Conditional) If you installed the software to each node in the cluster and you selected *Yes* for *Consolidate Configuration Files?* (under Internet Agent Clustering Worksheet item 6), copy the `gwia.cfg` file to the planned location on the shared disk, then delete it from the `c:\Program Files\Novell\GroupWise Server\GWIA` directory on each node to avoid future confusion.

8 Make sure you have completed all the tasks described in "Installing the GroupWise Internet Agent" in the *GroupWise 2012 Installation Guide*.

9 Continue with Enabling Internet Addressing for Your Clustered GroupWise System.

Enabling Internet Addressing for Your Clustered GroupWise System

Setting up Internet addressing for a clustered GWIA is no different from setting it up for a GWIA in a any other environment. Follow the instructions in "Enabling Internet Addressing" in "System" in the *GroupWise 2012 Administration Guide*, then continue with Verifying GWIA Object Properties.

Verifying GWIA Object Properties

During installation of the GWIA, the GWIA object should have been configured correctly. However, it can be helpful to verify certain cluster-specific information in order to familiarize yourself with the configuration of a clustered GWIA.

Accessing GWIA Object Properties

1 In ConsoleOne, browse to and select the GWIA domain in order to display its contents.

2 Right-click the GWIA object, then click *Properties*.

3 Continue with Verifying the Reference to the Network Name for Use by DNS.

Verifying the Reference to the Network Name for Use by DNS

In the GWIA object properties page tabs:

1. Click *SMTP/MIME > Settings*.
2. Verify the contents of the *Hostname/DNS "A Record" Name* field.

   It displays the hostname as currently configured in DNS. It should match the network name of the domain resource group, not the name of a node in the cluster.
3. (Conditional) Make changes if necessary.
4. Continue with Verifying the Reference to the Network Name in Directory Paths.

Verifying the Reference to the Network Name in Directory Paths

In the GWIA object properties page tabs:

1. Click *Server Directories*.
2. Verify that the displayed directories match the network name of the domain resource group, not the name of a node in the cluster.
3. (Conditional) Make changes if necessary.
4. Continue with Verifying Post Office Links.

Verifying Post Office Links

In the GWIA object properties page tabs:

1. Click *Post Office Links*.
2. Verify that the *Access Mode* column displays *C/S* (for client/server mode) for all post offices serviced by the clustered GWIA.
3. Verify that the *Links* column displays the IP addresses of the post office resource groups, not the IP addresses of any nodes in the cluster.
4. (Conditional) Make changes if necessary.
5. Continue with Forcing Use of the GWIA Resource Group IP Address.

Forcing Use of the GWIA Resource Group IP Address

If you want the GWIA to send outgoing messages on its resource group IP address, rather than using the default the node IP address:

1. Click *GroupWise > Network Address*.
2. In the *TCP/IP Address* field, provide the resource group IP address (GWIA Clustering Worksheet item 1) for the GWIA to use for sending outgoing messages.
3. Click *SMTP/MIME*, then click *Settings*.
4. Select *Bind to TCP/IP Address at Connection Time*.
5. Click *OK*.
6. Continue with Testing the Clustered GWIA.

23.2.5 Testing the Clustered GWIA

After you have set up the GWIA resource group, you can test it by manually bringing it online and taking it offline again.

Continue with Managing the Internet Agent in a Windows Cluster.

23.3 Managing the Internet Agent in a Windows Cluster

After you have installed the GWIA in a cluster, you should consider some long-term management issues.

23.3.1 Updating GroupWise Objects with Cluster-Specific Descriptions

After installing the GWIA in your clustered GroupWise system, while the cluster-specific information is fresh in your mind, you should record that cluster-specific information as part of the GroupWise objects in ConsoleOne so that you can easily refer to it later. Be sure to update the information recorded in the GroupWise objects if the configuration of your system changes.

Recording Cluster-Specific Information about the GWIA Domain and Its MTA

To permanently record important cluster-specific information for the GWIA domain:

1. In ConsoleOne, browse to and right-click the Domain object, then click *Properties*.
2. In the *Description* field of the GWIA domain Identification page, provide a cluster-specific description of the GWIA domain, including its resource group IP address and the cluster-unique port numbers used by its MTA.
3. Click *OK* to save the GWIA domain description.
4. Select the GWIA Domain object to display its contents.
5. Right-click the MTA object, then click *Properties*.
6. In the *Description* field of the MTA Identification page, record the domain resource group IP address and the cluster-unique port numbers used by the MTA.

   This information appears on the MTA console, no matter which node in the cluster it is currently running on.
7. Click *OK* to save the MTA description.
8. Continue with Recording Cluster-Specific Information about the GWIA.

Recording Cluster-Specific Information about the GWIA

With the contents of the GWIA domain still displayed:

1. Right-click the GWIA object, then click *Properties*.
2. Click *GroupWise*, then click *Identification*.

3 In the *Description* field, record the resource group IP address and the cluster-unique port numbers used by the GWIA.

   This information appears on the GWIA console, no matter which node in the cluster it is currently running on.

4 Click *OK* to save the GWIA information.

5 Continue with Knowing What to Expect in a GWIA Failover Situation.

23.3.2 Knowing What to Expect in a GWIA Failover Situation

The failover behavior of the MTA for the GWIA domain is the same as for an MTA in a regular domain. See "Knowing What to Expect in MTA and POA Failover Situations" on page 184.

Failover of the GWIA itself is more complex. The various email clients (POP3, IMAP4, and LDAP) receive an error message when the server they were connected to becomes unavailable. Most of the clients do not attempt to reconnect automatically, so the user must exit the email client and restart it to reestablish the connection after the failover process is complete. Fortunately, the GWIA restarts quickly in its failover location so users can reconnect quickly.

As with the MTA and the POA, manual migration of the GWIA takes longer than failover. In fact, the GWIA can seem especially slow to shut down properly, as it finishes its normal processing and stops its threads. For a busy GWIA, you might need to wait several minutes for it to shut down properly when you are manually migrating it.

23.4 Internet Agent Clustering Worksheet

Item	Explanation
1) Resource Group for GWIA: Network name: IP address: Physical disk: File share: MTA service resource: GWIA service resource: Possible owners:	Specify the information for the GWIA resource group. For more information, see "Planning the GWIA Resource Group" on page 188.
2) GWIA Domain Name: Domain Directory:	Specify a unique name for the GWIA domain. Specify the directory on the physical disk that belongs to the GWIA resource group where you want to create the new domain. For more information, see "Planning a Domain for the GWIA" on page 188.

Item	Explanation
4) MTA Installation Location: ◆ Shared disk of the GWIA resource group ◆ Each node in the cluster Consolidate MTA startup files?	Mark the location where you will install the MTA software. If necessary, specify the location where you will consolidate the MTA startup files from the various nodes where the GWIA is installed. For more information, see “Deciding Where to Install the GWIA and Its MTA” on page 189.
5) MTA Network Information: MTA IP address: MTA message transfer port: MTA HTTP port	Gather the MTA network address information from the GWIA section of the “Network Address Worksheet” on page 171. For more information, see “Planning Cluster-Unique Port Numbers for the GWIA and Its MTA” on page 188.
6) GWIA Installation Location: ◆ Shared disk in the GWIA resource group ◆ Each node in the cluster Consolidate configuration files?	Mark the location where you will install the GWIA software. If necessary, specify the location on the shared disk of the GWIA resource group where you will consolidate the GWIA configuration files (`gwia.cfg`) from the various nodes where it is installed. For more information, see “Deciding Where to Install the GWIA and Its MTA” on page 189.
7) GWIA Network Information: GWIA IP address: GWIA HTTP port:	Gather the GWIA network address information from the GWIA section of the “Network Address Worksheet” on page 171. For more information, see “Planning Cluster-Unique Port Numbers for the GWIA and Its MTA” on page 188.

24 Implementing the Document Viewer Agent in a Windows Cluster

If you want to install the Document Viewer Agent (DVA) along with the POA, it can participate in the cluster by failing over along with the POA and its post office. Follow the instructions in Section 22.4, "Creating a New Post Office in a Windows Cluster," on page 179.

However, clustering is not necessary in order to provide high availability for the DVA, because the DVA does not require the database access that is required by the MTA and POA. As an alternative to clustering the DVA, you can install the DVA on multiple servers outside the cluster, as described in "Scaling Your DVA Installation" in "Document Viewer Agent" in the *GroupWise 2012 Administration Guide*.

25 Monitoring a GroupWise System in a Windows Cluster

The setup procedure for the Monitor Agent in a Windows cluster is similar to the setup procedure for the POA and the MTA. Therefore, you should already have the skills necessary to set up the Monitor Agent as well.

When you first install the Monitor Agent, it gathers information about agents to monitor from a domain database (`wpdomain.db`). This provides the resource group IP address of each agent. When an agent fails over or migrates to a different node, its status in Monitor displays as *Not Listening* until it is up and running again, at which time its status returns to *Normal*.

Because the Monitor Agent must use resource group IP addresses to monitor the agents in a clustered GroupWise system, the *Discover Machine* and *Discover Network* options do not work in a cluster. Resource group IP addresses cannot be obtained by examining the network itself. If you need to add agents to monitor, use the *Add Agent* option and provide the agent's resource group IP address.

For instructions on setting up GroupWise Monitor, see "Installing GroupWise Monitor" in the *GroupWise 2012 Installation Guide*.

26 Backing Up a GroupWise System in a Windows Cluster

The issues involved in backing up a GroupWise system in a Windows cluster are the same as in backing up any GroupWise system that is running on Windows. If you want to back up your GroupWise system while it is running, you must use backup software that can back up open files. If your backup software cannot back up open files, then you must stop all GroupWise agents before running the backup and start them again when the backup is finished. This means that GroupWise users cannot be logged into their mailboxes while backups are running.

To find backup software that is compatible with GroupWise, see the *Novell Partner Product Guide* (http://www.novell.com/partnerguide).

27 Moving an Existing GroupWise 2012 System into a Windows Cluster

If you are adding the high availability benefits of a Windows cluster to a GroupWise 2012 system that is already up and running, the first step is to set up the cluster and review Chapter 20, "Introduction to GroupWise 2012 and Clustering on Windows," on page 157 to help you apply clustering principles and practices to your GroupWise system.

You do not need to transfer your entire GroupWise system into the cluster all at once. You can transfer individual post offices where the needs for high availability are greatest. You can transfer a domain and all of its post offices at the same time. You might decide that you don't need to have all of your GroupWise system running in the cluster.

This section provides a checklist to help you get started with moving your GroupWise system into a Microsoft cluster:

- ❒ Decide which shared disks you will use for GroupWise administration (ConsoleOne and the software distribution directory).
- ❒ Decide which shared disks you will use for GroupWise domains and post offices.
- ❒ Plan the resource groups for domains and post offices.
- ❒ Review Chapter 21, "Planning GroupWise in a Windows Cluster," on page 159. Fill out the "System Clustering Worksheet" on page 170 to help you decide which domains and post offices you will move to which shared disks.
- ❒ Review "Planning Cluster-Unique Port Numbers for Agents in the Cluster" on page 165 and fill out the "Network Address Worksheet" on page 171 to record resource group IP addresses and to specify cluster-specific port numbers for all of your GroupWise agents.
- ❒ Select the first shared disk that will be part of your clustered GroupWise system and set up the resource group for it, following the instructions in "Creating GroupWise Resource Groups" on page 175 and "Configuring Short Name Resolution" on page 176.
- ❒ Move a domain and/or post office onto the shared disk, following the instructions in "Moving a Domain" in "Domains" or "Moving a Post Office" in "Post Offices" in the *GroupWise 2012 Administration Guide*.
- ❒ Review Section 21.8, "Deciding How to Install and Configure the Agents in a Windows Cluster," on page 165, fill out the "Agent Clustering Worksheet" on page 172, and install the agents as needed for the first clustered domain and/or post office, following the instructions in Section 22.5, "Installing and Configuring the MTA and the POA in a Windows Cluster," on page 180.
- ❒ Test the first component of your clustered GroupWise system following the instructions in Section 22.6, "Testing Your Clustered GroupWise System," on page 182.
- ❒ Take care of the cluster management details described in Section 22.7, "Managing Your Clustered GroupWise System," on page 183.

- ❒ Move more domains and post offices into the cluster as needed. If you have GroupWise libraries, see Section 21.4, “Planning a New Library for a Clustered Post Office,” on page 162.
- ❒ Move GroupWise administration into the cluster as needed.
- ❒ Add other components to your clustered GroupWise system as needed, as described in the following sections:
  - Chapter 23, “Implementing the Internet Agent in a Windows Cluster,” on page 187
  - Chapter 25, “Monitoring a GroupWise System in a Windows Cluster,” on page 199
  - Chapter 26, “Backing Up a GroupWise System in a Windows Cluster,” on page 201

28 Implementing Messenger in a Windows Cluster

Novell Messenger does not require the existence of a GroupWise system in your Novell cluster, but presumably one has already been set up as described in Chapter 21, "Planning GroupWise in a Windows Cluster," on page 159 and Chapter 22, "Setting Up a Domain and Post Office in a Windows Cluster," on page 175. As part of the process of setting up GroupWise in your cluster, you filled out the "System Clustering Worksheet" on page 170. Some of the information from this worksheet will be helpful as you implement Messenger in your cluster.

28.1 Planning Your Messenger System in a Windows Cluster

Because the Messenger agents are not associated with GroupWise domains or post offices, the Messenger agents are easier to implement in a cluster than the GroupWise agents. Messenger Clustering Worksheet lists all the information you need as you set up the Messenger agents in a clustering environment. You should print the worksheet and fill it out as you complete the tasks listed below:

28.1.1 Understanding Your Windows Cluster

Fill out items 1 and 2 in Section 28.3, "Messenger Clustering Worksheet," on page 209 with information about your cluster. This information corresponds to items 1 and 2 on the "System Clustering Worksheet" on page 170 that you filled out for GroupWise. For background information, see Section 21.1, "Setting Up Your Windows Cluster," on page 160.

28.1.2 Planning Messenger Administration

If you have set up a shared disk for GroupWise administration, as described in Section 21.6, "Planning Shared Administrative Resources," on page 163, you can use the same shared disk for the Messenger administration files. For example, you might want to have a shared disk where you install the Messenger snap-in to ConsoleOne instead of installing it to multiple administrator workstations.

MESSENGER CLUSTERING WORKSHEET

Under Item 5: Installation Location for Messenger Administration, mark whether you want to install the Messenger snap-in to ConsoleOne to administrator workstations or to a shared disk.

If you plan to install the Messenger snap-in to ConsoleOne to a shared disk, under Item 6: Resource for Messenger Administration, list the network name and IP address of the shared disk, the physical disk name and file share for mapping to it, and the nodes in the cluster that it can fail over to.

28.1.3 Deciding Where to Install the Messenger Agent Software

In a Windows cluster, the Messenger agents must run as Windows services. When you install the Windows Messenger Agents, you can choose between two different installation locations:

Location	Description
Each node in the cluster	The `c:\novell\nm` directory is the default installation location provided by the Messenger Installation program.
Shared disk	If you create a *`drive`*`:\novell\nm` directory on a shared disk, the Messenger agent software and startup files fail over and fail back along with supporting files such as the Messenger archive. **IMPORTANT:** You must install to a shared disk if you do not want a separate Messenger archive to be created on each node where the Archive Agent runs. If you do not want to use a shared disk, you should plan to install the Archive Agent separately outside the cluster.

Because the Messenger agents must be installed as Windows services in a Windows cluster, you must initially run the Messenger Installation program for each node in the cluster so that the Windows services for the agents get created, regardless of where you are planning to run the Messenger agents from. However, for updates, you need to run the Messenger Installation program only once if you are running the Messenger agents from a shared disk.

MESSENGER CLUSTERING WORKSHEET

Under Item 3: Installation Location for Messenger Agents, mark whether you want to install the Messenger agent software to each node in the cluster or to a shared disk.

Continue with the planning instructions for the installation location you want to use:

- "Planning the Messenger Agents on Each Node in the Cluster" on page 206
- "Planning the Messenger Agents on a Shared Disk" on page 207

Planning the Messenger Agents on Each Node in the Cluster

Make sure you have filled out item 2 on the Messenger Clustering Worksheet with a complete list of nodes in the cluster where you need to install the Messenger agents. Skip to "Planning the Messenger Agent Installation" on page 207.

Planning the Messenger Agents on a Shared Disk

If you do not anticipate a large Messenger archive, you can use one Messenger shared disk. If you anticipate archiving a large number of messages so that the Messenger archive grows very large, you might want to have a separate Messenger shared disk for the Archive Agent and the archive database. The steps in this section cover setting up the Messenger agents on a single shared disk.

MESSENGER CLUSTERING WORKSHEET

Under Item 4: Resource Group for Messenger Agents, plan the network name and IP address of the resource group, the physical disk and share name for mapping to it, the agent service names, and the nodes in the cluster where the Messenger resource group can fail over.

Continue with Planning the Messenger Agent Installation.

28.1.4 Planning the Messenger Agent Installation

Aside from the cluster-specific issues discussed in the preceding sections, the considerations involved in planning to install the Messenger agents are the same in a clustering environment as for any other environment. Review "Planning Your Novell Messenger System", then print and fill out the "Novell Messenger Worksheet" in "Installing a Novell Messenger System" in the *Novell Messenger 2.2 Installation Guide*. Transfer the following information from the Messenger Clustering Worksheet to the Messenger System Worksheet:

- For "3) Installation Path"on the Messenger System Worksheet:
  - If you are installing the Messenger agents to each node in the cluster, use `c:\novell\nm`.
  - If you are installing the Messenger agents to a shared disk, use *`drive`*`:\novell\nm` where *drive* is the shared disk from Item 4: Resource Group for Messenger Agents on the Messenger Clustering Worksheet.
- Under "12) Server Address" on the Messenger System Worksheet:
  - If you are installing the Messenger agents to each node in the cluster, use the cluster IP address from Item 1: Cluster Identification on the Messenger Clustering Worksheet.
  - If you are installing the Messenger agents to a shared disk, specify the Messenger resource group IP address from Item 4: Resource Group for Messenger Agents on the Messenger Clustering Worksheet.
- Under Item 13: Configure Agents for Clustering? on the Messenger System Worksheet, mark *No*. This applies to the Messenger Agents running with Novell Cluster Services, not in a Windows cluster.
- Under Item 14: Admin Configuration on the Messenger System Worksheet:
  - If you are installing the Messenger snap-in to ConsoleOne to an administrator workstation, use the location where ConsoleOne is already installed (typically `c:\novell\consoleone\`*`version_number`*).
  - If you are installing the Messenger snap-in to ConsoleOne to a shared disk, use *`drive`*`:\`*`directory`*, where *drive* is the shared disk from Item 6: Resource for Messenger Administration on the Messenger Clustering Worksheet and *directory* is typically `c:\novell\consoleone\`*`version_number`*.

Continue with Setting Up Your Messenger System in a Windows Cluster.

28.2 Setting Up Your Messenger System in a Windows Cluster

You should have already reviewed Section 28.1, "Planning Your Messenger System in a Windows Cluster," on page 205 and filled out Section 28.3, "Messenger Clustering Worksheet," on page 209 and the "Novell Messenger Worksheet" in the *Novell Messenger 2.2 Installation Guide*. Follow the instructions for the installation location you have chosen:

- Section 28.2.1, "Installing the Messenger Agents to Each Node in the Cluster," on page 208
- Section 28.2.2, "Installing the Messenger Agents to a Shared Disk," on page 208

28.2.1 Installing the Messenger Agents to Each Node in the Cluster

1. Follow the steps provided in "Starting the Messenger Installation Program" and "Creating Your Messenger System" in "Installing a Novell Messenger System" in the *Novell Messenger 2.2 Installation Guide* for each node in the cluster.
2. After you have installed the software to each node in the cluster, if you selected *Yes* for *Consolidate Startup Files?* (under Messenger Clustering Worksheet item 3), copy the Messenger agent startup files to the planned location on the shared disk, then delete them from the `c:\novell\nm\ma` and `c:\novell\nm\aa` directories on each node to avoid future confusion.
3. Make each node in the cluster active to make sure that the Messenger agents start successfully on each node.
4. Continue setting up your Messenger system following the instructions in "What's Next" in "Installing a Novell Messenger System" in the *Novell Messenger 2.2 Installation Guide*.

28.2.2 Installing the Messenger Agents to a Shared Disk

Complete the following tasks to set up your Messenger system on a shared disk:

- "Setting Up the Messenger Resource Group" on page 208
- "Running the Messenger Installation Program" on page 208
- "Testing the Clustered Messenger Agents" on page 209

Setting Up the Messenger Resource Group

1. Create the Messenger resource group and agent services resources (Messenger Clustering Worksheet item 4), as planned in "Planning the Messenger Agents on Each Node in the Cluster" on page 206.
2. To ensure successful short name resolution, add entries for the Messenger network name to support your preferred methods of short name resolution, as described in "Configuring Short Name Resolution" on page 176.
3. Continue with Running the Messenger Installation Program.

Running the Messenger Installation Program

1. If necessary, map a drive to the shared disk for Messenger administration (Messenger Clustering worksheet item 6) where you will install the Messenger snap-ins to ConsoleOne.
2. Map a drive to the shared disk of the Messenger resource group (Messenger Clustering Worksheet item 4) where you will install the Messenger agent software.

3 Map a drive to `c:\` on the first node in the cluster (Messenger Clustering Worksheet item 2) where you will set up the Messenger agents as a Windows services.

4 Start the Messenger Installation program, following the steps provided in "Starting the Messenger Installation Program" in "Installing a Novell Messenger System" in the *Novell Messenger 2.2 Installation Guide*.

5 Install the Windows Messenger agents, keeping in mind the following cluster-specific details:

- Use the Novell Messenger System Worksheet that you filled out in Section 28.1.4, "Planning the Messenger Agent Installation," on page 207 to fill in the fields during the Messenger installation process.
- When you specify the Messenger installation directory, be sure to browse to the location through the drive mapped in Step 2 above.
- When you specify the ConsoleOne directory, be sure to browse to the location through the drive mapped in Step 1 above.
- On the Setup Complete page, do not select *Launch Agents Now*.

6 Repeat Step 4 and Step 5, mapping a drive to each node in the cluster.

Initially, you need to repeat the installation process for each node so that the Messenger agents are set up as Windows services on each node. For updates, you need to install only once to the shared disk.

7 Continue with Testing the Clustered Messenger Agents.

Testing the Clustered Messenger Agents

After you have set up the Messenger agents on a shared disk in your Windows cluster, you can test them by manually bringing the Messenger resource group online and taking it offline again.

Continue setting up your Messenger system following the instructions in "What's Next" in "Installing a Novell Messenger System" in the *Novell Messenger 2.2 Installation Guide*.

28.3 Messenger Clustering Worksheet

Item	Explanation
1) Cluster Identification: Cluster name: Cluster IP address	Record the name and IP address of your Windows cluster. For more information, see Section 21.1, "Setting Up Your Windows Cluster," on page 160.
2) Nodes in Cluster:	List the servers that are included in your Windows cluster. For more information, see Section 21.1, "Setting Up Your Windows Cluster," on page 160.
3) Installation Location for Messenger Agents: ◆ Each node in the cluster Consolidate startup files? ◆ Shared disk	Mark the location where you will install the Messenger agent software. For more information, see "Deciding Where to Install the Messenger Agent Software" on page 206.

Item	Explanation
4) Resource Group for Messenger Agents Network name: IP address: Physical disk: File share: Messaging Agent service: Archive Agent service: Possible owners	If you plan to install the Messenger agent software to a shared disk, provide the information about the shared disk you want to use. For more information, see “Planning the Messenger Agents on a Shared Disk” on page 207.
5) Installation Location for Messenger Administration: ◆ Administrator workstation(s) ◆ Shared disk	Mark the location where you want to install the Messenger snap-in to ConsoleOne. For more information, see “Planning Messenger Administration” on page 205.
6) Resource for Messenger Administration: Network name: IP address: Physical disk: File share: Possible owners	If you want to install the Messenger snap-in to ConsoleOne to a shared disk, provide the required information about the shared disk you want to use. For more information, see Section 21.6, “Planning Shared Administrative Resources,” on page 163.
7) IP Address Resolution Methods: ◆ eDirectory ◆ hosts file ◆ DNS	Mark the short name address resolution methods you want to implement to ensure that the UNC paths stored in ConsoleOne with network names can be successfully resolved into physical network addresses. For more information, see Section 21.7, “Ensuring Successful Name Resolution for GroupWise Resource Groups,” on page 163.

VI Non-GroupWise Email Clients

If your users already have a common POP, IMAP, or SOAP email client that comes with Linux or Windows, they can continue to use it to access their GroupWise mailboxes. Users of non-GroupWise email clients retain the feature sets of their familiar email clients, but many GroupWise features are not available to such users because they are not offered in POP, IMAP, and SOAP email clients. For example, calendaring is available only if the POP, IMAP, or SOAP client supports iCal.

29 Outlook Express

The GroupWise Internet Agent (GWIA) is required in order for users to access their mailboxes using non-GroupWise clients. If you have not already installed the GWIA, follow the instructions in the *GroupWise 2012 Installation Guide*.

In order for users to access their GroupWise mailboxes from a third-party email client, they must configure their email clients to access their GroupWise accounts. For example, Outlook Express users would follow steps similar to the following:

NOTE: Steps might vary depending on the versions of Windows and Outlook Express installed on the workstation.

---

1. In Outlook Express, click *Tools* > *Accounts* > *Add* > *Mail*.
2. Follow the prompts and provide personal information until you are prompted for the email server information.

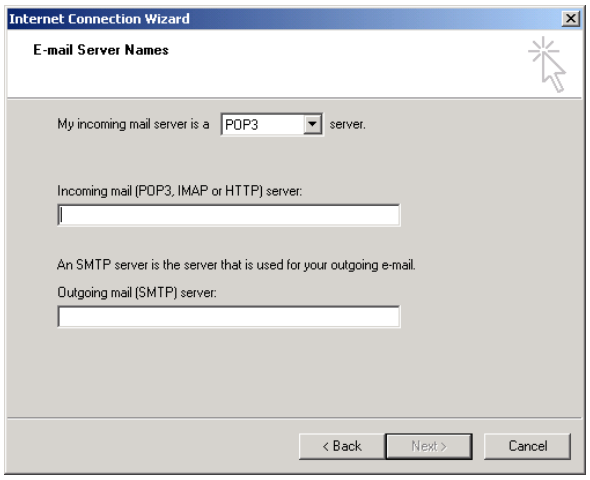

3. Select *POP3* or *IMAP* as your incoming mail server type.
4. In the *Incoming* and *Outgoing Mail* fields, specify the IP address or hostname of your outgoing mail server, then click *Next*.
5. Continue following the prompts and providing personal information until the new account has been set up in Outlook Express.
6. Click *Tools* > *Accounts*.

7 Select the new account you just created, then click *Properties* > *Servers*.

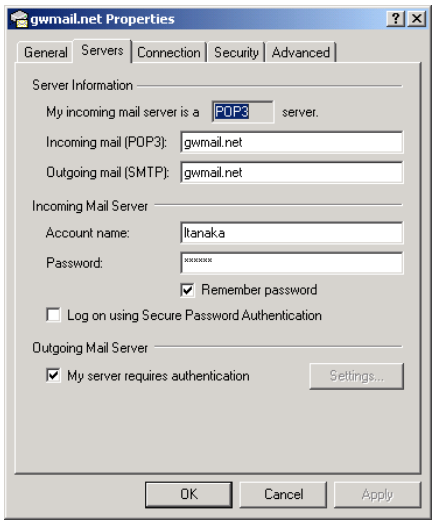

8 Select *My Server Requires Authentication*, then click *OK*.

The default setting for server authentication is *Use Same Settings as My Incoming Mail Server*, so you do not need to change any settings.

9 To access your GroupWise mailbox in Outlook Express, click *Tools* > *Send and Receive*.

10 Click the IP address or hostname of your mail server.

11 Provide your user name and password, then click *OK*.

30 Microsoft Outlook

The GroupWise Internet Agent (GWIA) is required in order for users to access their mailboxes using non-GroupWise clients. If you have not already installed the GWIA, follow the instructions in the *GroupWise 2012 Installation Guide*.

If your users have been using the Microsoft Outlook email client that comes with Microsoft Office, they can continue to use POP or IMAP in it to access their GroupWise mailboxes.

In order for users to access their GroupWise mailboxes from Outlook, they must configure Windows to access their GroupWise accounts. For example, Outlook users would follow steps similar to the following.

NOTE: Steps might vary depending on the versions of Windows and Outlook installed on the workstation.

1. In the Windows Control Panel, double-click *Mail*.
2. Click *Show Profiles* > *Add* to add a new profile for your GroupWise account.
3. Type a name for the new profile, the click *OK*.

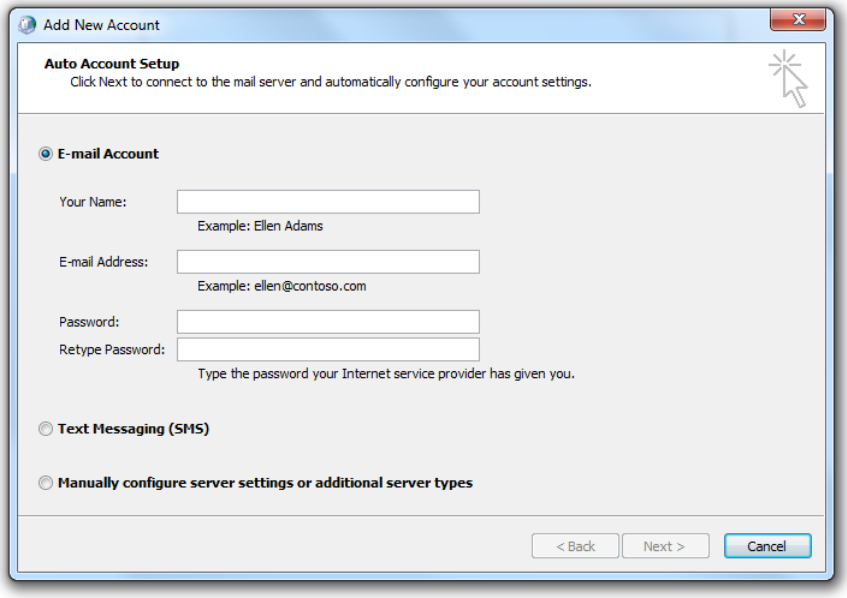

4. Select *Manually configure server settings or additional server types*, then click *Next*.

5 Click *Next* to accept *Internet E-mail* as the mail service.

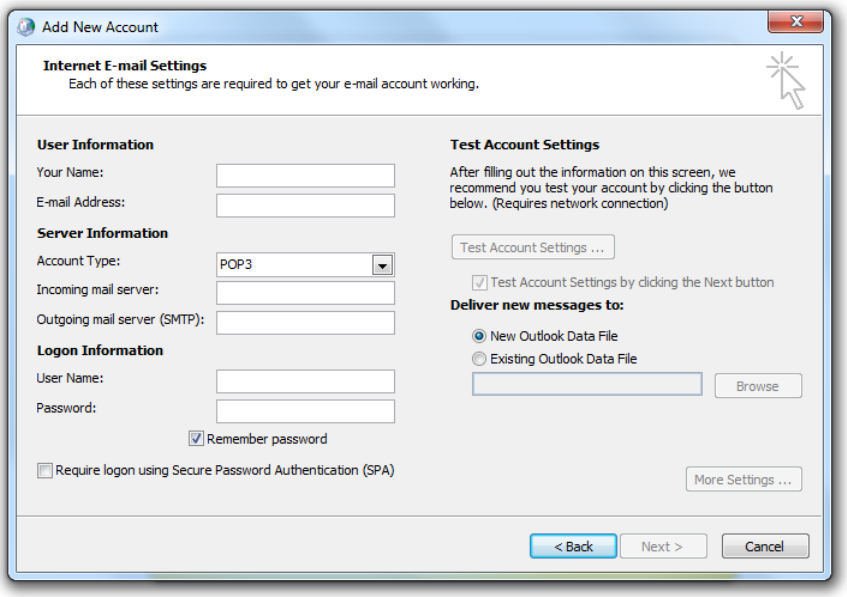

6 Provide the email account settings for your GroupWise system.

7 Click *Test Account Settings* to make sure that you have provided the information correctly.

8 Click *Next*, then click *Finish*.

You can now use Microsoft Outlook to access your GroupWise mailbox by selecting the profile you just created.

31 Evolution

Evolution makes the tasks of storing, organizing, and retrieving your personal information easy, so you can work and communicate more effectively with others. It's a highly evolved groupware program that can be an integral part of the Internet-connected desktop.

Evolution can help you work in a group by handling email, address, and other contact information, and one or more calendars. It can do that on one or more computers, connected directly or over a network, for one person or for large groups.

With Evolution, you can accomplish your most common daily tasks quickly. For example, it takes only one or two clicks to enter appointment or contact information sent to you by email, or to send email to a contact or appointment. People who get a lot of email will appreciate advanced features like vFolders, which let you save searches as though they were ordinary email folders.

If you have Evolution 2.4 or later installed, you can access accounts in Novell GroupWise 8 or later.

31.1 GroupWise Features Available in Evolution

Evolution connecting to GroupWise supports the following basic GroupWise features:

- Mail
  - View mail and folders stored on the GroupWise system.
  - Send mail from you GroupWise account.
  - Convert mail to a task or meeting.
- Calendar
  - Send and receive appointment and meeting requests. Allows Evolution users to schedule meetings and view attendee availability for other users on GroupWise.
  - Receive an iCalendar meeting request and add it to your calendar. It is saved to your GroupWise calendar.
- Contacts
  - Address Completion for your GroupWise address books, including the corporate address book, the Frequent Contacts address book, and your personal address book.
  - Adding vCards to the Address Book. If you receive a vCard attachment and click *Save in Address Book*, it is saved to your Personal address book. New Address Book entries can be added to your Personal address book from received email messages with a single click.
- Proxy
  - Assign Proxy access to other users.
  - View other users' accounts through Proxy access.

31.2 Configuring Evolution

In order for users to access their GroupWise mailboxes from Evolution, they must configure Evolution to access their GroupWise accounts.

1. In Evolution, click *Edit* > *Preferences*, then click *Mail Accounts*.
2. Click *Add*.
3. On the Identity page, type your email address, then click *Forward*.

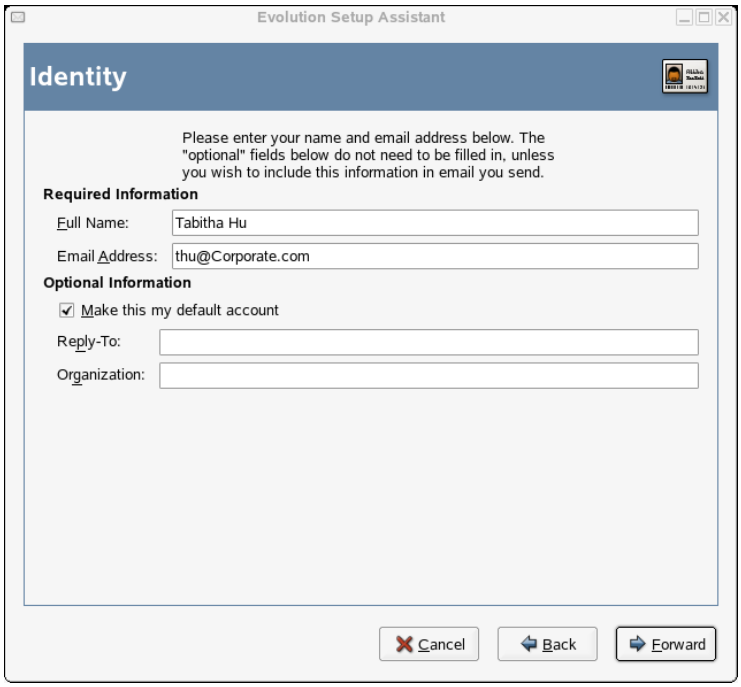

4. On the Receiving Mail page, select *Novell GroupWise* as your server type.
5. Type the name of your mail server, your user name, and select whether to use SSL.

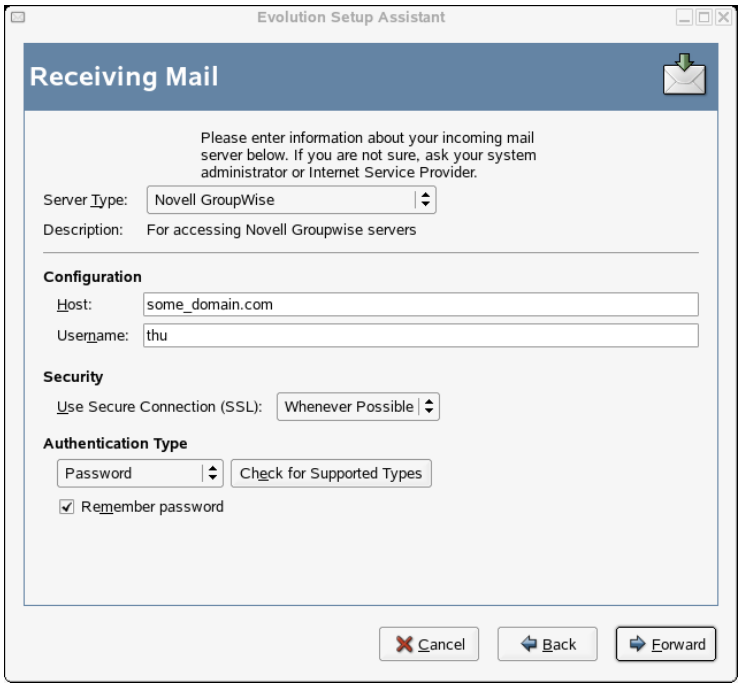

6 Click *Forward*.

7 On the Receive Options page, select if you want Evolution to automatically check for new mail.

If you select this option, you need to specify how often Evolution should check for new messages.

8 Select if you want to check for new messages in all folders.

9 Select if you want to apply filters to new messages in the Inbox on the server.

10 Select if you want to check new messages for junk content.

11 Select if you want to only check for junk messages in the Inbox folder.

12 Select if you want to automatically synchronize remote mail locally.

13 Type your Post Office Agent SOAP port number in the *Post Office Agent SOAP Port* field, then click *Forward*.

If you are unsure of what your Post Office Agent SOAP port number is, contact your system administrator.

14 On the Account Management page, type the name for the account, then click *Forward*.

15 Click *Apply*.

VII Mobile Devices

If you own a mobile device, you can synchronize it with GroupWise. GroupWise uses the Novell Data Synchronizer Mobility Pack for synchronization with many of the most common mobile devices. In addition, GroupWise has teamed up with BlackBerry for synchronization with BlackBerry devices.

32 Novell Data Synchronizer Mobility Pack

You can use Novell Data Synchronizer to synchronize email and other Personal Information Manager (PIM) data from Novell GroupWise to mobile devices. The Mobility Pack includes Data Synchronizer, the GroupWise Connector, and the Mobility Connector. Additional connectors can be added to a Synchronizer system to synchronizer GroupWise data to other supported applications.

For more information, see:

- Novell Data Synchronizer documentation Web page (http://www.novell.com/documentation/datasynchronizer1)
- Novell Data Synchronizer Connectors documentation Web page (http://www.novell.com/documentation/datasync_connectors1)

33 BlackBerry Enterprise Server

Novell and Research In Motion collaborate to deliver stellar support to the thousands of customers accessing GroupWise on BlackBerry devices. This partnership has resulted in strong solutions for end users and administrators alike.

The BlackBerry Enterprise Solution provides a complete wireless platform that allows organizations to extend the their Novell GroupWise messaging application and other enterprise tools to mobile professionals. The BlackBerry Enterprise Solution provides users with mobile access to email, instant messaging (IM), calendar, personal information management (PIM) and applications, all from a single wireless device. In addition, with BlackBerry push technology, these users are automatically sent up-to-date information while they're on the go.

BlackBerry Enterprise Server software is an important element of the BlackBerry Enterprise Solution. It is designed to provide IT departments with simplified management and centralized control of wireless devices in a secure, scalable and flexible architecture. BlackBerry Enterprise Server v.4.1 for Novell GroupWise includes several new features to enhance end user productivity and back-end administration. These features include Novell Messenger support, enhanced support for PowerPoint and Web Doc attachments, group-based and role-based administration, localized data pass-through, and SMS/PIN/call log auditing.

For more information about BlackBerry Enterprise Server for Novell GroupWise, see the BlackBerry Enterprise Server for Novell GroupWise product Web site (http://na.blackberry.com/eng/services/business/server/full/).

For documentation, see the BlackBerry Enterprise Server for Novell GroupWise product documentation Web site (http://docs.blackberry.com/en/admin/subcategories/?userType=2&category=BlackBerry+Enterprise+Server+for+Novell+GroupWise).

For GroupWise-specific BlackBerry articles, look up "GroupWise" in the BlackBerry Technical Solution Center. (http://na.blackberry.com/eng/support)

For support information, look up "GroupWise" in the BlackBerry Technical Knowledge Center (http://www.blackberry.com/knowledgecenterpublic/livelink.exe).

VIII Documentation Updates

This section lists updates to the *GroupWise 2012 Interoperability Guide* that have been made since the initial release of GroupWise 2012. The information helps you to keep current on documentation updates and, in some cases, software updates (such as a Support Pack release).

The information is grouped according to the date when the *GroupWise 2012 Interoperability Guide* was republished. Within each dated section, the updates are listed by the names of the main table of contents sections.

The *GroupWise 2012 Interoperability Guide* has been updated on the following dates:

A April 16, 2013 (GroupWise 2012 SP2)

Location	Change
Novell Cluster Services on Linux	
"Modifying the Cluster Resource Unload Script for the Linux Agents" on page 37	Corrected the syntax of quotation marks for the `pkill` command in unload scripts.
"Modifying the Cluster Resource Unload Script for the Linux GWIA and Its MTA" on page 60	
"Modifying the Cluster Resource Unload Script for the Linux Monitor Agent" on page 77	
"Modifying the Cluster Resource Unload Script for the Linux Messenger Agents" on page 99	

B September 20, 2012 (GroupWise 2012 SP1)

Location	Change
Non-GroupWise Email Clients	
Chapter 30, “Microsoft Outlook,” on page 215	Updated for Outlook 2010.