

Interoperability Guide

Novell® GroupWise®

7

June 25, 2008

www.novell.com



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 \(http://www.novell.com/info/exports/\)](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.

Copyright © 2003-2008 Novell, Inc. All rights reserved. No part of this publication may be reproduced, photocopied, stored on a retrieval system, or transmitted without the express written consent of the publisher.

Novell, Inc. has intellectual property rights relating to technology embodied in the product that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more of the U.S. patents listed on the [Novell Legal Patents Web page \(http://www.novell.com/company/legal/patents/\)](http://www.novell.com/company/legal/patents/) and one or more additional patents or pending patent applications in the U.S. and in other countries.

Novell, Inc.
404 Wyman Street, Suite 500
Waltham, MA 02451
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 \(http://www.novell.com/documentation\)](http://www.novell.com/documentation).

Novell Trademarks

For Novell trademarks, see the [Novell Trademark and Service Mark list \(http://www.novell.com/company/legal/trademarks/tmlist.html\)](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	17
Part I Novell Cluster Services on NetWare	19
1 Introduction to GroupWise 7 and Novell Cluster Services on NetWare	21
2 Planning GroupWise in a NetWare Cluster	23
2.1 Meeting Software Version Requirements	24
2.2 Installing Novell Cluster Services	24
2.3 Planning a New Clustered Domain	25
2.4 Planning a New Clustered Post Office	26
2.5 Planning a New Library for a Clustered Post Office	27
2.6 Deciding Whether to Cluster-Enable the Shared Volumes Used by GroupWise	27
2.7 Ensuring Successful Name Resolution for GroupWise Volumes	29
2.8 Deciding How to Install and Configure the Agents in a Cluster	31
2.8.1 Planning Secondary IP Addresses and Cluster-Unique Port Numbers for Agents in the Cluster	31
2.8.2 Determining Appropriate Failover Paths for the Agents	33
2.8.3 Deciding Where to Install the Agent Software	34
2.8.4 Deciding Whether to Run the Agents in Protected Memory	36
2.8.5 Planning the NetWare Agent Installation	37
2.9 GroupWise Clustering Worksheets	37
2.9.1 System Clustering Worksheet	37
2.9.2 IP Address Worksheet	40
2.9.3 Agent Clustering Worksheet	41
3 Setting Up a Domain and Post Office in a NetWare Cluster	43
3.1 Preparing the Cluster for GroupWise	43
3.1.1 Ensuring Required Software Versions	43
3.1.2 Cluster-Enabling Shared Volumes for Use with GroupWise	43
3.1.3 Configuring Short Name Resolution	44
3.2 Setting Up a New GroupWise System in a Cluster	46
3.3 Creating a New Secondary Domain in a Cluster	48
3.4 Creating a New Post Office in a Cluster	49
3.5 Installing and Configuring the MTA and the POA in a Cluster	50
3.5.1 Installing the Agent Software in a Cluster	51
3.5.2 Editing Clustered Agent Startup Files	52
3.5.3 Configuring the GroupWise Volume Resource to Load and Unload the Agents	53
3.5.4 Setting Up New Instances of the Agents without Installing the Agent Software	57
3.6 Testing Your Clustered GroupWise System	59
3.7 Managing Your Clustered GroupWise System	60
3.7.1 Updating GroupWise Objects with Cluster-Specific Descriptions	60
3.7.2 Using Novell Remote Manager on NetWare 6.x	61
3.7.3 Knowing What to Expect in MTA and POA Failover Situations	64
3.8 What's Next	64
3.9 Clustering Quick Checklists	65
3.9.1 GroupWise System Quick Checklist	65

3.9.2	Domain Quick Checklist	66
3.9.3	Post Office Quick Checklist	67

4 Implementing the Internet Agent in a NetWare Cluster 69

4.1	Planning the Internet Agent in a Cluster	69
4.1.1	Planning a Domain for the Internet Agent	70
4.1.2	Deciding Whether to Cluster-Enable the Internet Agent Volume	70
4.1.3	Determining an Appropriate Failover Path for the Internet Agent Volume	70
4.1.4	Planning a Secondary IP Address and Cluster-Unique Port Numbers for the Internet Agent and Its MTA	71
4.1.5	Preparing Your Firewall for the Internet Agent	71
4.1.6	Deciding Where to Install the Internet Agent and Its MTA	72
4.1.7	Deciding Whether to Run the Internet Agent and Its MTA in Protected Memory	72
4.1.8	Planning the MTA Installation	73
4.1.9	Planning the Internet Agent Installation	73
4.2	Setting Up the Internet Agent in a Cluster	73
4.2.1	Cluster-Enabling a Shared Volume for Use with the Internet Agent	73
4.2.2	Creating a Domain for the Internet Agent	74
4.2.3	Installing the MTA for the Internet Agent Domain	74
4.2.4	Installing and Configuring the Internet Agent in a Cluster	74
4.2.5	Testing the Clustered Internet Agent	82
4.3	Managing the Internet Agent in a Cluster	83
4.3.1	Updating GroupWise Objects with Cluster-Specific Descriptions	83
4.3.2	Knowing What to Expect in an Internet Agent Failover Situation	84
4.4	Internet Agent Clustering Worksheet	84
4.5	Internet Agent Quick Checklist	86

5 Implementing WebAccess in a NetWare Cluster 89

5.1	Understanding the WebAccess Components	89
5.2	Planning WebAccess in a Cluster	89
5.2.1	Planning a New Domain for the WebAccess Agent	90
5.2.2	Deciding Whether to Cluster-Enable the WebAccess Agent Volume	90
5.2.3	Determining an Appropriate Failover Path for the WebAccess Agent Volume	91
5.2.4	Planning a Secondary IP Address and Cluster-Unique Port Numbers for the WebAccess Agent and Its MTA	91
5.2.5	Deciding Where to Install the WebAccess Agent and Its MTA	91
5.2.6	Deciding Whether to Run the WebAccess Agent and Its MTA in Protected Memory	92
5.2.7	Planning the MTA Installation	92
5.2.8	Planning the WebAccess Installation	92
5.3	Setting Up WebAccess in a Cluster	93
5.3.1	Cluster-Enabling a Shared Volume for Use with the WebAccess Agent	93
5.3.2	Creating a Domain for the WebAccess Agent	94
5.3.3	Installing the MTA for the WebAccess Agent Domain	94
5.3.4	Installing and Configuring the WebAccess Agent in a Cluster	94
5.3.5	Testing Your Clustered WebAccess Installation	100
5.4	Managing WebAccess in a Cluster	101
5.4.1	Updating GroupWise Objects with Cluster-Specific Descriptions	101
5.4.2	Knowing What to Expect in WebAccess Failover Situations	102
5.4.3	Updating the WebAccess Agent Configuration File (commgr.cfg)	102
5.5	WebAccess Clustering Worksheet	103
5.6	WebAccess Quick Checklist	104

6	Implementing GroupWise Gateways in a Novell Cluster	107
7	Monitoring a GroupWise System in a NetWare Cluster	109
8	Backing Up a GroupWise System in a NetWare Cluster	111
8.1	Using GW TSA in a NetWare 5.1 Cluster	111
8.2	Using TSAFSGW in a NetWare 6.x Cluster	112
9	Updating a GroupWise System in a NetWare Cluster	115
10	Moving an Existing GroupWise 7 System into a NetWare Cluster	117
11	Implementing Messenger in a NetWare Cluster	119
11.1	Planning Your Messenger System in a Cluster	119
11.1.1	Understanding Your Cluster	119
11.1.2	Planning Messenger Administration	119
11.1.3	Deciding Where to Install the Messenger Agent Software	120
11.1.4	Planning the Messenger Agent Installation	122
11.2	Setting Up Your Messenger System in a Cluster	122
11.2.1	Installing to Each Node in the Cluster	123
11.2.2	Installing to a Messenger Volume	123
11.3	Messenger Clustering Worksheet	127
Part II	Novell Cluster Services on Linux	129
12	Introduction to GroupWise 7 and Novell Cluster Services on Linux	131
13	Planning GroupWise in a Linux Cluster	133
13.1	Installing Novell Cluster Services on Linux	134
13.2	Planning a Clustered Software Distribution Directory	135
13.3	Planning a New Clustered Domain	136
13.4	Planning a New Clustered Post Office	137
13.5	Planning a New Library for a Clustered Post Office	137
13.6	Deciding How to Install and Configure the Linux Agents in a Cluster	138
13.6.1	Recording Secondary IP Addresses for the Agents	138
13.6.2	Determining Appropriate Failover Lists for the Agents	138
13.6.3	Determining Cluster Resource Information for the Linux Agents	139
13.6.4	Planning the Linux Agent Installation	139
13.7	GroupWise Clustering Worksheets	139
13.7.1	System Clustering Worksheet	140
13.7.2	Agent Clustering Worksheet	141
14	Setting Up a Domain and a Post Office in a Linux Cluster	143
14.1	Setting Up a New GroupWise System in a Linux Cluster	143
14.2	Creating a New Secondary Domain in a Linux Cluster	144
14.3	Creating a New Post Office in a Linux Cluster	145
14.4	Installing and Configuring the MTA and the POA in a Cluster	146
14.4.1	Installing and Setting Up the Agents in Your Cluster	146

14.4.2	Changing Agent Paths to Locations on GroupWise Partitions	150
14.4.3	Configuring GroupWise Cluster Resources to Load and Unload the Agents	152
14.4.4	Setting Up New Instances of the Agents without Installing the Agent Software	158
14.5	Testing Your Clustered GroupWise System.	160
14.6	Managing Your Clustered GroupWise System.	161
14.6.1	Updating GroupWise Objects with Cluster-Specific Descriptions	161
14.6.2	Knowing What to Expect in MTA and POA Failover Situations	163
14.7	What's Next	163
14.8	Clustering Quick Checklists	164
14.8.1	GroupWise System Quick Checklist	164
14.8.2	Domain Quick Checklist.	164
14.8.3	Post Office Quick Checklist	165

15 Implementing the Internet Agent in a Linux Cluster 167

15.1	Planning the Internet Agent in a Linux Cluster.	167
15.1.1	Planning a Domain for the Internet Agent	168
15.1.2	Selecting the Internet Agent Partition and Secondary IP Address	168
15.1.3	Determining an Appropriate Failover List for the Internet Agent	169
15.1.4	Determining Cluster Resource Information for the Internet Agent.	169
15.1.5	Preparing DNS for the Clustered Internet Agent	169
15.1.6	Preparing Your Firewall for the Clustered Internet Agent	169
15.1.7	Planning the MTA Installation	170
15.1.8	Planning the Internet Agent Installation	170
15.2	Setting Up the Internet Agent in a Linux Cluster	171
15.2.1	Creating a Domain for the Internet Agent	171
15.2.2	Installing the MTA for the Internet Agent Domain	171
15.2.3	Installing and Configuring the Internet Agent in a Cluster	171
15.3	Testing the Internet Agent in a Linux Cluster.	182
15.4	Managing the Internet Agent in a Linux Cluster.	183
15.4.1	Updating GroupWise Objects with Cluster-Specific Descriptions	183
15.4.2	Knowing What to Expect in an Internet Agent Failover Situation	184
15.5	Internet Agent Clustering Worksheet	184
15.6	Internet Agent Quick Checklist.	185

16 Implementing WebAccess in a Linux Cluster 187

16.1	Understanding the WebAccess Components	187
16.2	Planning the WebAccess Agent in a Linux Cluster	187
16.2.1	Planning a Domain for the WebAccess Agent.	188
16.2.2	Selecting the WebAccess Agent Partition and Secondary IP Address	189
16.2.3	Determining an Appropriate Failover List for the WebAccess Agent.	189
16.2.4	Determining Cluster Resource Information for the WebAccess Agent	189
16.2.5	Planning the MTA Installation	190
16.2.6	Planning the WebAccess Agent Installation	190
16.2.7	Planning the WebAccess Application Installation	190
16.3	Setting Up the WebAccess Agent in a Linux Cluster	190
16.3.1	Creating a Domain for the WebAccess Agent	191
16.3.2	Installing the MTA for the WebAccess Agent Domain.	191
16.3.3	Installing and Configuring the WebAccess Agent in a Cluster.	191
16.3.4	Installing and Configuring the WebAccess Application in a Cluster	201
16.4	Testing the WebAccess Agent in a Linux Cluster	201
16.5	Managing the WebAccess Agent in a Linux Cluster	202
16.5.1	Updating GroupWise Objects with Cluster-Specific Descriptions	203
16.5.2	Knowing What to Expect in a WebAccess Agent Failover Situation	204
16.6	WebAccess Agent Clustering Worksheet	204

16.7	WebAccess Agent Quick Checklist	205
17	Implementing GroupWise Monitor in a Linux Cluster	207
17.1	Understanding the Monitor Components	207
17.2	Planning GroupWise Monitor in a Linux Cluster.	208
17.2.1	Selecting a Domain for Access during Monitor Agent Installation	208
17.2.2	Selecting an MTA for the Monitor Agent to Access after Installation	208
17.2.3	Selecting the Monitor Agent Partition and Secondary IP Address.	209
17.2.4	Determining an Appropriate Failover List for the Monitor Agent	209
17.2.5	Determining Cluster Resource Information for the Monitor Agent	209
17.2.6	Planning the Monitor Agent Installation	210
17.3	Setting Up GroupWise Monitor in a Linux Cluster	210
17.3.1	Installing and Configuring the Monitor Agent on Each Node in Your Cluster	210
17.3.2	Configuring the Monitor Agent Cluster Resource to Load and Unload the Monitor Agent	213
17.4	Testing the Monitor Agent in a Linux Cluster	217
17.5	Managing the Monitor Agent in a Linux Cluster	218
17.6	Monitor Agent Clustering Worksheet	218
17.7	Monitor Agent Quick Checklist	219
18	Backing Up a GroupWise System in a Linux Cluster	221
19	Updating a GroupWise System in a Linux Cluster	223
20	Moving an Existing Linux GroupWise 7 System into a Linux Cluster	225
21	Moving a Clustered GroupWise 7 System from NetWare to Linux	227
22	Implementing Messenger in a Linux Cluster	229
22.1	Planning Your Messenger System in a Linux Cluster	229
22.1.1	Understanding Your Cluster	229
22.1.2	Selecting the Messenger Partition and Secondary IP Address	230
22.1.3	Determining an Appropriate Failover List for the Messenger Agents	230
22.1.4	Determining Cluster Resource Information for the Messenger Agents	230
22.1.5	Planning the Messenger Agent Installation	231
22.2	Setting Up Your Messenger System in a Linux Cluster	231
22.2.1	Creating Your Messenger System and Installing the Messenger Agents	231
22.2.2	Changing Messenger Paths to Locations on the Messenger Partition	233
22.2.3	Configuring the Messenger Cluster Resource to Load and Unload the Messenger Agents	235
22.3	Testing Your Clustered Messenger System.	239
22.4	Managing Your Clustered Messenger System.	240
22.5	Messenger Clustering Worksheet	240
22.6	Messenger Clustering Quick Checklist.	241
Part III	Novell Teaming and Conferencing	243
23	Using GroupWise with Novell Teaming	245
23.1	Understanding How Novell Teaming Interacts with eDirectory and GroupWise	245

23.2	Meeting Novell Teaming System Requirements on OES 2 or SLES 10 SP1 Servers	245
23.2.1	Linux Open File Limit	246
23.2.2	Java Development Kit Requirement	246
23.2.3	MySQL Setup	247
23.2.4	Potential Web Server Port Conflicts	247
23.3	Installing Novell Teaming with GroupWise	248
23.3.1	Preparing for the Novell Teaming Installation	248
23.3.2	Performing the Novell Teaming Installation	248
23.4	Configuring LDAP Authentication	249
23.5	Configuring GroupWise to Support Your Teaming Installation	249
23.5.1	Configuring Outgoing E-Mail from Teaming to GroupWise	249
23.5.2	Configuring Incoming E-Mail from GroupWise Users to the Novell Teaming Site	251
23.5.3	Testing GroupWise Integration with Novell Teaming	257
23.6	Receiving Notification of Teaming Site Activity in Your GroupWise Client	259
23.7	Receiving Calendar Postings from the Teaming Site	259
23.8	Adding Portlets to Novell Teaming	259
23.8.1	Adding Liferay Portlets	259
23.8.2	Adding GroupWise Portlets	262
 24 Using GroupWise with Conferencing		263
24.1	Configuring GroupWise as the Conferencing E-Mail System	263
24.1.1	Preparing for GroupWise Integration with Conferencing	263
24.1.2	Integrating GroupWise with Novell Conferencing	263
24.1.3	Testing GroupWise as the Novell Conferencing E-Mail System	264
 25 Streamlining Authentication to Teaming or Conferencing		265
25.1	Using iChain for Authenticating to Teaming or Conferencing	265
25.1.1	Meeting iChain Requirements	265
25.1.2	Setting Up an iChain Web Server Accelerator for Teaming or Conferencing	265
25.1.3	Adding the New Web Server Accelerator to the iChain Server Object in ConsoleOne	266
25.1.4	Using iChain for Authentication	267
25.2	Using Novell Access Manager for Authenticating to Teaming or Conferencing	267
 Part IV Other Novell Products		269
 26 GroupWise DirXML Driver for Novell Identity Manager		271
26.1	Identity Manager Warnings in ConsoleOne	271
26.1.1	Recovering a Deleted GroupWise Account	271
26.1.2	Grafting Users	272
26.1.3	Converting an External Entity to a User	272
26.1.4	Converting a User to an External Entity	272
26.1.5	Associating a GroupWise Object with an eDirectory Object	272
26.1.6	Disassociating a GroupWise Object's Attributes from an eDirectory Object	272
26.1.7	Resolving an Invalid Association	273
26.1.8	Disabling the DirXML Warnings	273
26.1.9	Enabling the DirXML Warnings	273

27 GroupWise Customization Tools	275
28 Novell exteNd	277
Part V Heartbeat on Linux	279
29 Introduction to GroupWise 7 and Heartbeat on Linux	281
29.1 Understanding Heartbeat	281
29.2 Meeting System Requirements	281
29.3 Installing Heartbeat Packages	282
29.4 Configuring Heartbeat	282
29.5 Starting Heartbeat	283
29.6 Managing Heartbeat.	283
29.7 Setting Up Shared Storage	283
29.8 Setting Up Node Fencing or Resource Fencing.	284
29.9 What's Next	284
30 Planning GroupWise in a Heartbeat Cluster	285
30.1 Planning Your Heartbeat Cluster for GroupWise	285
30.1.1 Recording Heartbeat Cluster Information	286
30.1.2 Planning GroupWise Cluster Organization	286
30.1.3 Planning Cluster Resources and Resource Groups	287
30.1.4 Planning Secondary IP Addresses	287
30.2 Planning a Software Distribution Directory for a Cluster	288
30.3 Planning a New Clustered Domain	288
30.4 Planning a New Clustered Post Office.	289
30.5 Planning a New Library for a Clustered Post Office.	290
30.6 Considering How to Install and Configure the Linux Agents in a Cluster.	290
30.6.1 Recording Secondary IP Addresses for the Agents	290
30.6.2 Determining Appropriate Heartbeat Constraints for GroupWise Cluster Resource Groups.	291
30.6.3 Planning the Linux Agent Installation.	293
30.7 GroupWise Clustering Worksheets	293
30.7.1 Heartbeat Clustering Worksheet	293
30.7.2 GroupWise Clustering Worksheet	294
31 Setting Up a Heartbeat Cluster	297
31.1 Starting the HA Management Client.	297
31.2 Creating a Cluster Resource Group	298
31.3 Creating Native Heartbeat Resources	299
31.3.1 EVMS Container Resources	299
31.3.2 File System Resources	300
31.3.3 IP Address Resources.	301
31.4 Creating Heartbeat Constraints	302
31.4.1 Place Constraints.	302
31.4.2 Order Constraints	303
31.4.3 Colocation Constraints.	304
31.5 Starting the Native Heartbeat Resources.	305

32	Setting Up a Domain and a Post Office in a Heartbeat Cluster	307
32.1	Setting Up a New GroupWise System in a Heartbeat Cluster	307
32.2	Creating a New Secondary Domain in a Heartbeat Cluster	308
32.3	Creating a New Post Office in a Heartbeat Cluster	309
32.4	Installing and Configuring the MTA and the POA in a Heartbeat Cluster	310
32.4.1	Running the GroupWise Installation Program on the Initial Node	311
32.4.2	Running the GroupWise Installation Program on Subsequent Nodes	313
32.4.3	Creating and Configuring GroupWise Cluster Resources	315
32.4.4	Testing Your Agent Installation on Each Node	317
32.4.5	Changing Agent Paths to Locations on the Shared Storage Partition	318
32.4.6	Setting Up New Instances of the Agents without Installing the Agent Software	319
32.5	Testing Your Clustered GroupWise System	320
32.6	Managing Your Clustered GroupWise System	321
32.6.1	Updating GroupWise Objects with Cluster-Specific Descriptions	321
32.6.2	Knowing What to Expect in MTA and POA Failover Situations	322
32.7	What's Next	323
33	Implementing the Internet Agent in a Heartbeat Cluster	325
33.1	Planning the Internet Agent in a Heartbeat Cluster	325
33.1.1	Planning a Cluster Resource Group for the Internet Agent	326
33.1.2	Planning a Domain for the Internet Agent	326
33.1.3	Recording the Internet Agent Secondary IP Address	327
33.1.4	Determining Appropriate Heartbeat Constraints for the Internet Agent	327
33.1.5	Preparing DNS for the Clustered Internet Agent	327
33.1.6	Preparing Your Firewall for the Clustered Internet Agent	327
33.1.7	Planning the MTA Installation	328
33.1.8	Planning the Internet Agent Installation	328
33.2	Setting Up the Internet Agent in a Heartbeat Cluster	329
33.2.1	Setting Up Native Heartbeat Cluster Resources for the Internet Agent	329
33.2.2	Creating a Domain for the Internet Agent	329
33.2.3	Installing the MTA for the Internet Agent Domain	330
33.2.4	Installing and Configuring the Internet Agent in a Cluster	330
33.3	Testing the Internet Agent in a Heartbeat Cluster	338
33.4	Managing the Internet Agent in a Heartbeat Cluster	339
33.4.1	Updating GroupWise Objects with Cluster-Specific Descriptions	339
33.4.2	Knowing What to Expect in an Internet Agent Failover Situation	340
33.5	Internet Agent Clustering Worksheet	340
34	Implementing WebAccess in a Heartbeat Cluster	343
34.1	Understanding the WebAccess Components	343
34.2	Planning the WebAccess Agent in a Heartbeat Cluster	343
34.2.1	Planning a Cluster Resource Group for the WebAccess Agent	344
34.2.2	Planning a Domain for the WebAccess Agent	344
34.2.3	Recording the WebAccess Secondary IP Address	345
34.2.4	Determining Appropriate Heartbeat Constraints for the WebAccess Agent	345
34.2.5	Planning the MTA Installation	346
34.2.6	Planning the WebAccess Agent Installation	346
34.2.7	Planning the WebAccess Application Installation	346
34.3	Setting Up the WebAccess Agent in a Heartbeat Cluster	346
34.3.1	Setting Up Native Heartbeat Resources for the WebAccess Agent	347
34.3.2	Creating a Domain for the WebAccess Agent	347
34.3.3	Installing the MTA for the WebAccess Agent Domain	347
34.3.4	Installing and Configuring the WebAccess Agent in a Heartbeat Cluster	348

34.3.5	Installing and Configuring the WebAccess Application on Your Web Server	355
34.4	Testing the WebAccess Agent in a Heartbeat Cluster	356
34.5	Managing the WebAccess Agent in a Heartbeat Cluster	356
34.5.1	Updating GroupWise Objects with Cluster-Specific Descriptions	357
34.5.2	Knowing What to Expect in a WebAccess Agent Failover Situation	357
34.6	WebAccess Agent Clustering Worksheet	358
35	Implementing GroupWise Monitor in a Heartbeat Cluster	361
36	Backing Up a GroupWise System in a Heartbeat Cluster	363
37	Updating a GroupWise System in a Heartbeat Cluster	365
38	Moving an Existing Linux GroupWise 7 System into a Heartbeat Cluster	367
39	Implementing Messenger in a Heartbeat Cluster	369
39.1	Planning Your Messenger System in a Heartbeat Cluster	369
39.1.1	Planning a Cluster Resource Group for the Messenger Agents	369
39.1.2	Recording the Messenger Agents' Secondary IP Address	370
39.1.3	Determining Appropriate Constraints for the Messenger Agents	370
39.1.4	Planning the Messenger Agent Installation	371
39.2	Setting Up Your Messenger System in a Heartbeat Cluster	371
39.2.1	Setting Up Native Heartbeat Resources for the Messenger Agents	371
39.2.2	Creating Your Messenger System and Installing the Messenger Agents	372
39.2.3	Changing Messenger Paths to Locations on the Messenger Partition	376
39.3	Testing Your Clustered Messenger System	378
39.4	Managing Your Clustered Messenger System	378
39.5	Messenger Clustering Worksheet	379
Part VI	Microsoft Clustering Services on Windows	381
40	Introduction to GroupWise 7 and Microsoft Clusters	383
41	Planning GroupWise in a Microsoft Cluster	385
41.1	Setting Up Your Microsoft Cluster	386
41.2	Planning a New Clustered Domain	386
41.3	Planning a New Clustered Post Office	387
41.4	Planning a New Library for a Clustered Post Office	388
41.5	Planning GroupWise Resource Groups	388
41.6	Planning Shared Administrative Resources	389
41.7	Ensuring Successful Name Resolution for GroupWise Resource Groups	389
41.8	Deciding How to Install and Configure the Agents in a Cluster	391
41.8.1	Planning Cluster-Unique Port Numbers for Agents in the Cluster	391
41.8.2	Deciding Where to Install the Agent Software	393
41.8.3	Planning the Agent Services	395
41.8.4	Planning the Windows Agent Installation	395
41.9	GroupWise Clustering Worksheets	396
41.9.1	System Clustering Worksheet	396
41.9.2	Network Address Worksheet	398
41.9.3	Agent Clustering Worksheet	399

42	Setting Up a Domain and Post Office in a Microsoft Cluster	401
42.1	Preparing the Cluster for GroupWise	401
42.1.1	Creating GroupWise Resource Groups	401
42.1.2	Creating Agent Service Resources	401
42.1.3	Configuring Short Name Resolution	402
42.2	Setting Up a New GroupWise System in a Cluster	403
42.3	Creating a New Secondary Domain in a Cluster	404
42.4	Creating a New Post Office in a Cluster.	405
42.5	Installing and Configuring the MTA and the POA in a Cluster	406
42.5.1	Installing the Agent Software in a Cluster	407
42.5.2	Editing Clustered Agent Startup Files	408
42.5.3	Setting Up New Instances of the Agents without Installing the Agent Software	408
42.6	Testing Your Clustered GroupWise System.	409
42.7	Managing Your Clustered GroupWise System.	409
42.7.1	Updating GroupWise Objects with Cluster-Specific Descriptions	409
42.7.2	Knowing What to Expect in MTA and POA Failover Situations	410
42.8	What's Next	411
43	Implementing the Internet Agent in a Microsoft Cluster	413
43.1	Planning the Internet Agent in a Cluster	413
43.1.1	Planning a Domain for the Internet Agent	413
43.1.2	Planning the Internet Agent Resource Group	414
43.1.3	Planning Cluster-Unique Port Numbers for the Internet Agent and Its MTA	414
43.1.4	Preparing Your Firewall for the Internet Agent	415
43.1.5	Deciding Where to Install the Internet Agent and Its MTA.	415
43.1.6	Planning the MTA Installation	416
43.1.7	Planning the Internet Agent Installation	416
43.2	Setting Up the Internet Agent in a Cluster	416
43.2.1	Setting Up the Internet Agent Resource Group	417
43.2.2	Creating a Domain for the Internet Agent	417
43.2.3	Installing the MTA for the Internet Agent Domain	417
43.2.4	Installing and Configuring the Internet Agent in a Cluster	417
43.2.5	Testing the Clustered Internet Agent.	420
43.3	Managing the Internet Agent in a Cluster.	420
43.3.1	Updating GroupWise Objects with Cluster-Specific Descriptions	420
43.3.2	Knowing What to Expect in an Internet Agent Failover Situation	421
43.4	Internet Agent Clustering Worksheet	422
44	Implementing WebAccess in a Microsoft Cluster	423
44.1	Understanding the WebAccess Components	423
44.2	Planning WebAccess in a Cluster	423
44.2.1	Setting Up Your Web Server in the Microsoft Cluster	424
44.2.2	Planning a New Domain for the WebAccess Agent.	424
44.2.3	Planning the WebAccess Resource Group	425
44.2.4	Planning Cluster-Unique Port Numbers for the WebAccess Agent and Its MTA.	425
44.2.5	Deciding Where to Install the WebAccess Agent and Its MTA	426
44.2.6	Planning the MTA Installation	426
44.2.7	Planning the WebAccess Installation.	426
44.3	Setting Up WebAccess in a Cluster	427
44.3.1	Setting Up the WebAccess Resource Group.	427
44.3.2	Creating a Domain for the WebAccess Agent	428
44.3.3	Installing the MTA for the WebAccess Agent Domain.	428
44.3.4	Installing the WebAccess Agent in a Cluster.	428

44.3.5	Installing and Configuring the WebAccess Application in a Cluster	429
44.3.6	Testing Your Clustered WebAccess Installation	430
44.4	Managing WebAccess in a Cluster	430
44.4.1	Updating GroupWise Objects with Cluster-Specific Descriptions	431
44.4.2	Knowing What to Expect in WebAccess Failover Situations	432
44.4.3	Updating the WebAccess Agent Configuration File (commgr.cfg)	432
44.5	WebAccess Clustering Worksheet	433
45	Implementing GroupWise Gateways in a Microsoft Cluster	435
46	Monitoring a GroupWise System in a Microsoft Cluster	437
47	Backing Up a GroupWise System in a Microsoft Cluster	439
48	Moving an Existing GroupWise 7 System into a Microsoft Cluster	441
49	Implementing Messenger in a Microsoft Cluster	443
49.1	Planning Your Messenger System in a Cluster	443
49.1.1	Understanding Your Cluster	443
49.1.2	Planning Messenger Administration	443
49.1.3	Deciding Where to Install the Messenger Agent Software	444
49.1.4	Planning the Messenger Agent Installation	445
49.2	Setting Up Your Messenger System in a Cluster	446
49.2.1	Installing the Messenger Agents to Each Node in the Cluster	446
49.2.2	Installing the Messenger Agents to a Shared Disk	446
49.3	Messenger Clustering Worksheet	447
Part VII	Non-GroupWise Clients	449
50	Outlook Express	451
51	Microsoft Outlook	453
52	Evolution	455
52.1	GroupWise Features Available in Evolution	455
52.2	Configuring Evolution	456
Part VIII	Mobile Devices	459
53	GroupWise Mobile Server, Powered by Intellisync	461
54	BlackBerry Enterprise Server	463
55	GroupWise PDA Connect	465
55.1	Prerequisites	465
55.2	Downloading the Software	465
55.3	Third-Party Partners	465

Part IX Documentation Updates	467
A June 25, 2008	469
B March 14, 2008 (GroupWise 7 SP3)	471
C October 24, 2007	473
D April 16, 2007 (GroupWise 7 SP 2)	475
E September 29, 2006	477
F August 2, 2006	479
G June 15, 2006 (GroupWise 7 SP1)	481
H November 30, 2005	485

About This Guide

This Novell® *GroupWise*® 7 *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 NetWare,” on page 19
- ◆ Part II, “Novell Cluster Services on Linux,” on page 129
- ◆ Part III, “Novell Teaming and Conferencing,” on page 243
- ◆ Part IV, “Other Novell Products,” on page 269

Third-Party Products

- ◆ Part V, “Heartbeat on Linux,” on page 279
 - ◆ Part VI, “Microsoft Clustering Services on Windows,” on page 381
 - ◆ Part VII, “Non-GroupWise Clients,” on page 449
-

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

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, or go to www.novell.com/documentation/feedback.html and enter your comments there.

Documentation Updates

For the most recent version of the *GroupWise 7 Interoperability Guide*, visit the [Novell GroupWise 7 documentation Web site \(http://www.novell.com/documentation/gw7\)](http://www.novell.com/documentation/gw7).

Additional Documentation

For additional GroupWise documentation, see the following guides at the [Novell GroupWise 7 documentation Web site \(http://www.novell.com/documentation/gw7\)](http://www.novell.com/documentation/gw7):

- ◆ *Installation Guide*
- ◆ *Administration Guide*
- ◆ *Multi-System Administration Guide*
- ◆ *Troubleshooting Guides*
- ◆ *GroupWise Client User Guides*
- ◆ *GroupWise Client Frequently Asked Questions (FAQ)*

Documentation Conventions

In Novell documentation, a greater-than symbol (>) is used to separate actions within a step and items within a cross-reference path.

A trademark symbol (™, ®, etc.) denotes a Novell trademark. An asterisk denotes a third-party trademark.

When a single pathname can be written with a backslash for some platforms or a forward slash for other platforms, the pathname is presented with a backslash. Users of platforms that require a forward slash, such as Linux*, should use forward slashes as required by your software.

When a startup switch can be written with a forward slash for some platforms or a double hyphen for other platforms, the startup switch is presented with a forward slash. Users of platforms that require a double hyphen, such as Linux, should use double hyphens as required by your software.

Novell Cluster Services on NetWare

- ♦ Chapter 1, “Introduction to GroupWise 7 and Novell Cluster Services on NetWare,” on page 21
- ♦ Chapter 2, “Planning GroupWise in a NetWare Cluster,” on page 23
- ♦ Chapter 3, “Setting Up a Domain and Post Office in a NetWare Cluster,” on page 43
- ♦ Chapter 4, “Implementing the Internet Agent in a NetWare Cluster,” on page 69
- ♦ Chapter 5, “Implementing WebAccess in a NetWare Cluster,” on page 89
- ♦ Chapter 6, “Implementing GroupWise Gateways in a Novell Cluster,” on page 107
- ♦ Chapter 7, “Monitoring a GroupWise System in a NetWare Cluster,” on page 109
- ♦ Chapter 8, “Backing Up a GroupWise System in a NetWare Cluster,” on page 111
- ♦ Chapter 9, “Updating a GroupWise System in a NetWare Cluster,” on page 115
- ♦ Chapter 10, “Moving an Existing GroupWise 7 System into a NetWare Cluster,” on page 117
- ♦ Chapter 11, “Implementing Messenger in a NetWare Cluster,” on page 119

Introduction to GroupWise 7 and Novell Cluster Services on NetWare

1

Before implementing GroupWise[®] 7 with Novell[®] Cluster Services[™], make sure you have a solid understanding of Novell Cluster Services by reviewing the following information resources:

- ♦ **AppNote:** *An Introduction to Novell Cluster Services* (http://developer.novell.com/research/appnotes/1999/may/01/a990501_.pdf)
- ♦ **Novell Open Enterprise Server (OES) Product Documentation:** *OES Novell Cluster Services 1.8 Administration Guide for NetWare* (http://www.novell.com/documentation/oes/cluster_admin/data/h4hgu4hs.html#bktitle)
- ♦ **NetWare 6.5 Product Documentation:** *Novell Cluster Services* (<http://www.novell.com/documentation/ncs65>)
- ♦ **NetWare 6 Product Documentation:** *Novell Cluster Services* (<http://www.novell.com/documentation/ncs6p>)
- ♦ **NetWare 5.1 Product Documentation:** *Novell Cluster Services* (<http://www.novell.com/documentation/ncs>)

When you review the information 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 2 to 32 NetWare[®] 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 NetWare server that is part of a cluster.

resource: An IP address, volume, application, service, and so on, that can function successfully anywhere in the cluster. The volumes where domains and post offices reside are a specific type of cluster resources termed “volume resources.” In this section, the terms “cluster resource” and “volume resource” are used instead of “resource” to avoid confusion with GroupWise resources (such as conference rooms and projectors).

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 setting up cluster resources, you need to consider what components need to fail over together in order to continue functioning.

fan-out-failover: The configuration where cluster resources from a failed node fail over to 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 WebAccess, 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.

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

shared volume: A volume in a shared disk system that can be accessed from any cluster node that needs the data stored on it.

cluster-enabled shared volume: A shared volume for which a Volume Resource object has been created in Novell eDirectory™. The properties of the Volume Resource object provide load and unload scripts for programs installed on the volume, failover/failback/migration policies for the volume, and the failover path for the volume. Cluster-enabling is highly recommended for GroupWise.

GroupWise volume: As used in this section, a cluster-enabled shared volume that is used for GroupWise, such as for storing a domain, post office, software distribution directory, and so on. This section also uses the terms Internet Agent volume, WebAccess Agent volume, Messenger volume, and gateway volume in a similar manner.

storage area network (SAN): The cluster nodes together with their shared disk system and shared volumes.

virtual server: A logical server, rather than a physical server, to which cluster-enabled shared volumes are tied.

active/active mode: The configuration of a clustered application where the application runs simultaneously on multiple nodes in the cluster. Active/active mode is recommended when the GroupWise MTA, POA, Internet Agent, and WebAccess Agent run in protected memory because protected memory isolates them from each other, even if they are running on the same node.

active/passive mode: The configuration of a clustered application where the application runs on only one node at a time in the cluster. The GroupWise MTA, POA, Internet Agent, and WebAccess Agent must run in active/passive mode if they are not running in protected memory because only one instance of each agent/database combination can be running at the same time in the cluster.

Planning GroupWise in a NetWare Cluster

2

The majority of this part of the *GroupWise 7 Interoperability Guide* (Chapter 2, “Planning GroupWise in a NetWare Cluster,” on page 23 through Chapter 8, “Backing Up a GroupWise System in a NetWare Cluster,” on page 111) 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™. If you already have an existing GroupWise 7 system and need to configure it to work in a newly installed cluster, see Chapter 10, “Moving an Existing GroupWise 7 System into a NetWare Cluster,” on page 117.

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 7 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 (MTA and POA)

During the planning process, component configuration alternatives are explained. For example, you might want the domain and post office together on the same shared volume or on different shared volumes. You might want to install the agents to standard `sys:\system` directories or to manually created `vol:\system` directories on shared volumes where domains and post offices reside. You might or might not need to run the agents in protected memory.

The “System Clustering Worksheet” on page 37 lists all the information you need as you set up GroupWise in a clustering environment. You should print the worksheet and fill it out as you complete the tasks listed below:

- ♦ Section 2.1, “Meeting Software Version Requirements,” on page 24
- ♦ Section 2.2, “Installing Novell Cluster Services,” on page 24
- ♦ Section 2.3, “Planning a New Clustered Domain,” on page 25
- ♦ Section 2.4, “Planning a New Clustered Post Office,” on page 26
- ♦ Section 2.5, “Planning a New Library for a Clustered Post Office,” on page 27
- ♦ Section 2.6, “Deciding Whether to Cluster-Enable the Shared Volumes Used by GroupWise,” on page 27
- ♦ Section 2.7, “Ensuring Successful Name Resolution for GroupWise Volumes,” on page 29
- ♦ Section 2.8, “Deciding How to Install and Configure the Agents in a Cluster,” on page 31
- ♦ Section 2.9, “GroupWise Clustering Worksheets,” on page 37

After you have completed the tasks and filled out “System Clustering Worksheet” on page 37, you are ready to continue with Chapter 3, “Setting Up a Domain and Post Office in a NetWare Cluster,” on page 43.

2.1 Meeting Software Version Requirements

GroupWise 7 can be clustered on a system that meets the following requirements:

- GroupWise 7
- A supported version of NetWare® with the latest Support Pack
 - ◆ OES NetWare
 - ◆ NetWare 6.5
 - ◆ NetWare 6
 - ◆ NetWare 5.1

IMPORTANT: Novell Cluster Services does not support mixed NetWare versions within a cluster.

SYSTEM CLUSTERING WORKSHEET

Under **Item 1: Software Version Updates for Cluster**, mark any updates required for nodes in the cluster to ensure that all nodes in the cluster are running the same version of NetWare.

2.2 Installing Novell Cluster Services

Install Novell Cluster Services by following the instructions provided in the documentation for your version of NetWare, as listed in **Chapter 1, “Introduction to GroupWise 7 and Novell Cluster Services on NetWare,”** on page 21.

The installation process includes:

- ◆ Meeting hardware and software requirements
- ◆ Setting up a shared disk system
- ◆ Creating a new NetWare 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
- ◆ Mounting the shared volumes where you will set up GroupWise domains and post offices and install the GroupWise agents

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

SYSTEM CLUSTERING WORKSHEET

Under **Item 2: eDirectory Tree for Cluster**, record the name of the eDirectory tree where the new NetWare Cluster object has been created.

Under **Item 3: Cluster Name**, record the name of the NetWare Cluster object that you created for your GroupWise system.

Under **Item 4: Cluster Context**, record the full context of the NetWare Cluster object.

Under **Item 5: Nodes in Cluster**, list the nodes that you have added to the cluster.

The number of nodes and shared volumes that are available 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 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.
- ◆ In a like manner, WebAccess provides user access to GroupWise mailboxes across the Internet through users' Web browsers. It is another good candidate for 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 e-mail 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 would be 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 Internet Agent 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 clustering environment. It all depends on the specific needs of your particular GroupWise system and its users.

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 will be creating the primary domain as you complete the tasks in this section. In preparation, review “[Planning Your Basic GroupWise System](#)”, then print and fill out the “[Basic GroupWise System Worksheet](#)” in “[Installing a Basic GroupWise System](#)” in the *GroupWise 7 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 will be creating a new secondary domain. In preparation, review “[Planning a New Domain](#)”, then print and fill out the “[Domain Worksheet](#)” in “[Domains](#)” in the *GroupWise 7 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 volume 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 [item 17](#). If you are filling out the Domain Worksheet, stop with [item 10](#).

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 10: Domain Name](#), transfer the domain name and database directory to the System Clustering Worksheet.

Under [Item 7: Shared Volume for Domain](#), transfer the domain location to the System Clustering Worksheet. You will fill out the rest of the information under item 7 later.

IMPORTANT: Do not create the new domain until you are instructed to do so in [Chapter 3](#), “[Setting Up a Domain and Post Office in a NetWare Cluster](#),” on page 43.

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 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 “[Post Office Worksheet](#)” in “[Post Offices](#)” in the *GroupWise 7 Administration Guide*. When you specify the locations for the post office directories, include the shared volumes 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 11: Post Office Name](#), transfer the post office name and database location to the System Clustering Worksheet.

If you will create the post office on a different shared volume from where the domain is located, under [Item 8: Shared Volume for Post Office](#), transfer the post office location to the System Clustering Worksheet. You will fill out the rest of the information under item 8 later.

IMPORTANT: Do not create the new post office until you are instructed to do so in [Chapter 3](#), “[Setting Up a Domain and Post Office in a NetWare Cluster](#),” on page 43.

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. You can plan a library for a clustered post office by following the standard instructions provided in “[Creating and Managing Libraries](#)” in the *GroupWise 7 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 14: Library 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, specify a user name and password that the POA can use to access the volume where the document storage area will reside.

IMPORTANT: Do not create the new library until you are instructed to do so in [Chapter 3, “Setting Up a Domain and Post Office in a NetWare Cluster,”](#) on page 43.

2.6 Deciding Whether to Cluster-Enable the Shared Volumes Used by GroupWise

Cluster-enabling the shared volumes where domains and post offices reside greatly simplifies GroupWise administration. If you are creating a new GroupWise system, you might also want to cluster-enable shared volumes for the GroupWise administration snap-ins to ConsoleOne[®] and for the GroupWise software distribution directory so that these locations are always available within the cluster. To review the concept of cluster-enabled shared volumes, see the applicable section of clustering documentation for your version of NetWare, as listed in [Chapter 1, “Introduction to GroupWise 7 and Novell Cluster Services on NetWare,”](#) on page 21.

The advantages of cluster-enabling GroupWise volumes include:

- ◆ Drive mappings always occur through the virtual server associated with the cluster-enabled volume, rather than through a physical server. This guarantees that you can always map a drive to the domain or post office database no matter which node it is currently located on.
- ◆ The GroupWise snap-ins to ConsoleOne always work no matter which node is running ConsoleOne.
- ◆ Cluster-enabling the domain volume and installing the GroupWise agents to this volume guarantees that the GroupWise snap-ins to ConsoleOne can always find the configuration files that they need to access.
- ◆ When you rebuild a domain database or a post office database, you do not need to determine which node the database is currently located on.
- ◆ Help desk personnel do not need to be trained to determine where GroupWise is running before they connect to a domain to create a new GroupWise user.

When you cluster-enable a volume, additional eDirectory objects are created:

Table 2-1 eDirectory Objects Used in a Cluster

eDirectory Object	Object Name and Description
	<p>clustername_volumename (default object name) A new Volume object represents the cluster-enabled volume. It is created by renaming the original Volume object that was tied to a physical server and associating it with a virtual server instead.</p> <p>For example, if your cluster name is GWCLUSTER and your original volume name is gwvol1, the new Volume object representing the cluster-enabled volume is named gwcluster_gwvol1.</p>
	<p>clustername_volumename_SERVER (default object name) A new Server object represents the virtual server to which the new cluster-enabled volume is tied.</p> <p>Continuing with the above example, the new Server object representing the virtual server is named GWCLUSTER_GWVOL1_SERVER.</p>
	<p>volumename_SERVER.clustername (default object name) A new Volume Resource object stores property information for the cluster-enabled volume, such as start, failover, and failback mode information and load/unload scripts. These modes and scripts enable the cluster-enabled volume to function much like an independent server; hence, the SERVER portion of its name. The Volume Resource object is created in the Cluster container object.</p> <p>Continuing with the above example, the new Volume Resource object is named GWVOL1_SERVER.GWCLUSTER.</p>

IMPORTANT: Notice that the default object names include the underscore () character. Some DNS name servers cannot resolve object names that include underscore characters. If you have met the system requirements described in [Section 2.1, “Meeting Software Version Requirements,”](#) on [page 24](#), you can rename these objects as needed when you cluster enable the volume.

Cluster-enabling the shared volumes used by GroupWise is highly recommended. Throughout the rest of this document, the term “GroupWise volume” means “a cluster-enabled shared volume used by GroupWise.”

SYSTEM CLUSTERING WORKSHEET

Under [Item 6: Shared Volumes for GroupWise Administration](#), list any shared volumes you want to use for GroupWise administration purposes. For example, you might have a shared pub: volume with a public directory where you install the GroupWise snap-ins to ConsoleOne instead of installing them on multiple administrator workstations. You might have a shared apps: volume where you create the GroupWise software distribution directory. Mark whether or not you want to cluster-enable the GroupWise administration volumes.

Under [Item 7: Shared Volume for Domain](#), specify the name of the shared volume where you will create the domain. Mark whether or not you want to cluster-enable the domain volume. Also mark whether you will place the post office on the same volume with the domain.

If you want the post office on a different volume from where the domain is located, under [Item 8: Shared Volume for Post Office](#), specify the name of the shared volume where you will create the post office. Mark whether or not you want to cluster-enable the post office volume.

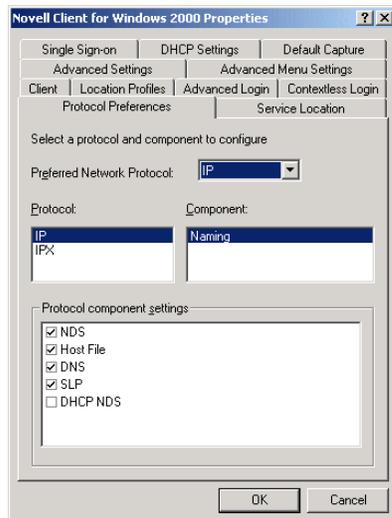
IMPORTANT: Because cluster-enabling the volumes where domains and post offices reside is so strongly recommended, this documentation does not include the steps for setting up domains and post offices on non-cluster-enabled volumes. If you decide not to cluster-enable GroupWise volumes, you should adjust the steps presented in this documentation for your system's specialized needs. Novell Cluster Services does provide a GroupWise Mail Server template for use when creating GroupWise Cluster Resource objects instead of cluster-enabled Volume Resource objects.

2.7 Ensuring Successful Name Resolution for GroupWise Volumes

Because you are using cluster-enabled volumes for GroupWise domains and post offices, 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 address of the current, physical location of that domain within your cluster. It is through short name resolution that all GroupWise cluster resources (such as domain and post office volumes) 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:

Figure 2-1 Novell Client Preferences Property Page



Short name resolution methods that pertain to clustering your GroupWise system are discussed below:

Table 2-2 *Short Name Resolution Methods*

Short Name Resolution Method	Description
eDirectory	<p>You can use eDirectory to resolve short names into specific network addresses. However, when using 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.</p>
Hosts File	<p>Windows uses the following files when performing short name resolution at the workstation:</p> <ul style="list-style-type: none">◆ Windows 2000/XP/2003: <code>\winnt\system32\drivers\etc\hosts</code> <p>Using these files at the Windows workstation is not a preferred method for TCP/IP name resolution (except perhaps for the administrator's workstation).</p> <p>However, whenever you cluster-enable a volume, you should add its virtual server to the <code>sys:\etc\hosts</code> file of all nodes in the cluster.</p>
DNS	<p>Perhaps the most common short name resolution option is Domain Name Service (DNS). As with the <code>hosts</code> file, it is good practice to place all of your virtual servers into DNS.</p> <p>For short name resolution to work using DNS, the client workstation must either belong to the same DNS zone (such as <code>provo.novell.com</code>) as the cluster resource, or the cluster resource zone must be configured in the client's DNS suffix search path under TCP/IP settings for the workstation.</p> <p>The underscore (<code>_</code>) character is part of default cluster-related object names. Because it is not supported by the DNS RFC, some DNS name servers cannot resolve default cluster-related object names. If you are using such a DNS name server on NetWare 5.1, make sure you have installed the latest Novell Cluster Services snap-in to ConsoleOne, so that you can change cluster-related object names as needed to remove the underscore characters.</p>
SLP	<p>NetWare 6.x and NetWare 5.1 both use Service Location Protocol (SLP) to advertise service information across TCP/IP-based networks, which provides short name resolution of TCP/IP-based cluster resources within the network. On NetWare 6.x, Novell Cluster Services propagates virtual server information into SLP by default.</p> <p>On NetWare 5.1, Novell Cluster Services does not propagate virtual server information into SLP by default. If you want to propagate virtual server information to SLP on NetWare 5.1, you can run the (unsupported) CVSBIND utility, which gives you reliable short name resolution within your cluster regardless of shortcomings you might encounter with other name resolution methods.</p>

Specific setup instructions for each of these short name resolution methods will be provided in [Chapter 3, "Setting Up a Domain and Post Office in a NetWare Cluster,"](#) on page 43.

2.8 Deciding How to Install and Configure the Agents in a Cluster

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

- ♦ [Section 2.8.1, “Planning Secondary IP Addresses and Cluster-Unique Port Numbers for Agents in the Cluster,”](#) on page 31
- ♦ [Section 2.8.2, “Determining Appropriate Failover Paths for the Agents,”](#) on page 33
- ♦ [Section 2.8.3, “Deciding Where to Install the Agent Software,”](#) on page 34
- ♦ [Section 2.8.4, “Deciding Whether to Run the Agents in Protected Memory,”](#) on page 36
- ♦ [Section 2.8.5, “Planning the NetWare Agent Installation,”](#) on page 37

2.8.1 Planning Secondary IP Addresses and Cluster-Unique Port Numbers for Agents in the Cluster

The GroupWise agents listen on all IP addresses, both primary and secondary, that are bound to the server on their specified port numbers. 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 a cluster-unique port number, even though each agent is using a unique secondary IP address. 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 [“IP Address Worksheet”](#) on page 40 to help you plan secondary IP addresses and cluster-unique port numbers for all GroupWise agents.

The following filled-out version of the IP 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
Development	(same as MTA)	1677	7101	7181
Manufacturing	172.16.5.82	1678	7102	7182

Internet Agent Information

Internet Agent	GWIA IP Address	MTA MTP Port	MTA Live Remote Port	MTA HTTP Port	GWIA HTTP Port
GWIA Domain MTA	172.16.5.83	7110	7677	7183	N/A
Internet Agent (GWIA)	(same as MTA)	N/A	N/A	N/A	9850

WebAccess Information

WebAccess Agent	WebAccess IP Address	MTA MTP Port	MTA HTTP Port	WebAccess Agent Port	WebAccess HTTP Port
WebAccess Domain MTA	172.16.5.84	7120	7184	N/A	N/A
WebAccess Agent (GWINTER)	(same as MTA)	N/A	N/A	7205	7205 (same as agent)

This example places the Development post office on the same node and on the same GroupWise volume with the Provo1 domain; therefore, the Provo1 MTA and the Development POA can use the same secondary IP address. The Manufacturing post office is placed on a different node on a different GroupWise volume, so that the Manufacturing post office has a different secondary IP address.

The example also illustrates that the MTA, the POA, and the Internet Agent use different port numbers for agent ports and HTTP ports. In contrast, the WebAccess Agent uses the same port number for the agent port and the HTTP port.

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.

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. For more information, see [“Simplifying Client/Server Access with a GroupWise Name Server”](#) in [“Post Office Agent”](#) in the *GroupWise 7 Administration Guide*.

IP ADDRESS WORKSHEET

Fill out the [“IP Address Worksheet”](#) on page 40 to help you plan secondary IP addresses and cluster-unique port numbers for all GroupWise agents in the cluster. (MTA, POA, Internet Agent, WebAccess Agent).

After you have filled out the IP Address Worksheet, transfer the secondary IP addresses and cluster-unique port numbers from the IP Address Worksheet to the System Clustering Worksheet and the Agent Clustering Worksheet so that they are available in the sequence in which you will need them as you set up GroupWise in a cluster.

SYSTEM CLUSTERING WORKSHEET

If you are setting up a new GroupWise system, under **Item 6: Shared Volumes for GroupWise Administration**, specify secondary IP addresses for your GroupWise administration volumes.

Under **Item 7: Shared Volume for Domain**, use the domain MTA secondary IP address from the IP Address Worksheet as the domain volume IP address.

If you are planning the post office on a different volume from the domain, under **Item 8: Shared Volume for Post Office**, use the post office POA secondary IP address from the IP Address Worksheet as the post office volume IP address.

AGENT CLUSTERING WORKSHEET

Under **Item 4: MTA Network Information**, transfer the secondary IP address and cluster-unique port numbers for the MTA from the IP Address Worksheet to the Agent Clustering Worksheet.

Under **Item 7: POA Network Information**, transfer the secondary IP address and cluster-unique port numbers for the POA from the IP Address Worksheet to the Agent Clustering Worksheet.

2.8.2 Determining Appropriate Failover Paths for the Agents

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

When a node fails, its volumes should not all fail over together to the same secondary node. Instead, the volumes 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 volumes 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 3: Domain Failover Path**, list the nodes that you want to have in the domain volume failover path. The MTA might need to run on any node that the domain volume fails over to.

If you are planning the post office on a different GroupWise volume from where the domain is located, under **Item 6: Post Office Failover Path**, list the nodes that you want to have in the post office volume failover path. The POA might need to run on any node that the post office volume fails over to.

2.8.3 Deciding Where to Install the Agent Software

When you install the NetWare MTA and POA in a clustering environment, you can choose between two different installation locations:

Table 2-3 *Agent Software Installation Locations*

Location	Description
<code>sys:\system</code> on each node in the cluster	<p>This is the default location provided by the Agent Installation program. Because the agents must be installed on each node where they might need to run during a failover situation, you need to do one of the following if you select this alternative:</p> <ul style="list-style-type: none">◆ Run the Agent Installation program multiple times in order to install the agent software and to create the agent startup files on each node that is on a GroupWise volume failover path.◆ Run the Agent Installation program, then copy the agent software and startup files to each node that is on a GroupWise volume failover path.
<code>system</code> directory on each GroupWise volume	<p>If you create a <code>vol:\system</code> directory on a GroupWise volume, the agent software and startup files fail over and back with the domains and post offices that the agents service.</p> <p>Unless you have a very small GroupWise system with all domains and post offices on a single GroupWise volume, you still need to install the agent software multiple times, once to each GroupWise volume.</p>

A simple way to look at the agent location alternatives would be that if you have fewer nodes on failover paths than you have GroupWise volumes for domains and post offices, then it would be most efficient to install the agent software to the nodes. Conversely, if you have fewer GroupWise volumes than you have nodes on failover paths, then it would be most efficient to install the agent software to the GroupWise volumes. However, there are issues to consider that extend beyond efficiency during installation.

The following sections can help you choose which installation location would be best for your clustered GroupWise system:

- ◆ [“Advantages of a `vol:\system` Directory on Each GroupWise Volume” on page 34](#)
- ◆ [“Disadvantages of a `vol:\system` Directory on Each GroupWise Volume” on page 35](#)
- ◆ [“Recommendation” on page 35](#)

Advantages of a `vol:\system` Directory on Each GroupWise Volume

Using a `vol:\system` directory on each GroupWise volume has several advantages:

- ◆ If you change information in the agent startup files, you only need to change it in one place, not on every node on any GroupWise volume failover path.
- ◆ Having the agent startup files on the same GroupWise volume as the domain or post office makes them easy to find.

- ◆ 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 on a GroupWise volume failover path. This prevents the potential problem of having a domain or post office fail over to a location where a different version of the agent software is installed.
- ◆ If you ever need to add or replace a physical server in the cluster, you only need to install NetWare and Novell Cluster Services to the new server, then add that node to the appropriate failover paths. No extra GroupWise configuration is necessary because there are no `sys:\system` dependencies for the GroupWise agents.
- ◆ If you want to back up the GroupWise software, you do not have to include the `sys:\system` directory in the backup.

Disadvantages of a `vol:\system` Directory on Each GroupWise Volume

Installing the agents on a GroupWise volume does have some disadvantages:

- ◆ GroupWise administrators who are used to the GroupWise agents being installed in `sys:\system` 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. Accidentally installing a GroupWise Support Pack to the default location of `sys:\system` would not have the desired results if the original agent software was installed to the `vol:\system` directory on a GroupWise volume.

Recommendation

Whichever method you choose, be consistent throughout the entire cluster. Either put all the GroupWise agents on the GroupWise volumes with the domains and post offices they service, or put them all in `sys:\system` directories. If you put them on GroupWise volumes, make sure there are no agent files in `sys:\system` directories to confuse the issue at a later time.

Even if you choose to install the agents to multiple `sys:\system` directories, you can still store the agent startup files on the GroupWise volumes. The significant advantage of this approach is that you only have 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 a `vol:\system` directory on a GroupWise volume or to `sys:\system` on 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 and location from the System Clustering Worksheet to the Agent Clustering Worksheet.

Under **Item 5: Post Office Name**, transfer the post office name and location from the System Clustering Worksheet to the Agent Clustering Worksheet.

2.8.4 Deciding Whether to Run the Agents in Protected Memory

On a NetWare server, using protected memory allows you to create isolated memory spaces where NLM™ programs can run without affecting other NLM programs running on the same node. This contributes to the high availability of the cluster. Using protected memory has the following advantages:

- ♦ When using protected memory, the node can restart a specific memory space if any NLM program within that memory space abends. This allows for recovery without failing the entire node, which enhances both up time and database integrity.
- ♦ Using protected memory gives you the ability to unload a single instance of an agent, rather than all instances.
- ♦ If you use protected memory, you can run the agents in active/active mode, rather than active/passive mode.

If you have any possibility of the same type of GroupWise agent loading multiple times on any node in the cluster, you must use protected memory so that you can unload agents individually. Check your failover paths ([Agent Clustering Worksheet items 3 and 6](#)) for failover combinations where multiple instances of the same type of agent might need to run on the same node.

Protected memory does result in higher memory utilization (about 5% to 10%) and a slight performance penalty. Make sure your nodes have sufficient memory to handle the number of memory spaces that might reside on them. Keep in mind that if you load the MTA and the POA in different memory spaces, the agent engine (`gwenn5.nlm`) will load twice on the node. Remember to provide memory for any GroupWise volumes that could fail over to a node, in addition to that node's regular processing load.

IMPORTANT: For optimum stability, we strongly recommend that you run the agents in protected memory, with one agent per memory space.

AGENT CLUSTERING WORKSHEET

Under [Item 8: Load Agents in Protected Memory?](#), mark whether or not you need to run the GroupWise agents in protected memory.

If you will use protected memory, provide one or two unique protected memory space names. If you will create the domain and post office on the same GroupWise volume, the MTA and POA can use the same memory space, although this is not recommended. If you will create the domain and post office on different GroupWise volumes, the MTA and POA must use different memory spaces.

If you will use protected memory, a user name and password for the POA to use to access its post office volume might be required, depending on the version of NetWare you are using.

Provide a user name and password if you are using the following versions of NetWare:

- ♦ NetWare 5.1 Support Pack 2 or earlier
- ♦ Initial release of NetWare 6

A user name and password are no longer needed on the following later versions of NetWare.

- ♦ NetWare 5.1 Support Pack 3 or later
 - ♦ NetWare 6 Support Pack 1 or later
-

2.8.5 Planning the NetWare Agent Installation

Aside from the cluster-specific issues discussed in the preceding sections, the considerations involved in planning to install the GroupWise NetWare 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 7 Installation Guide* for each location where you will install the NetWare MTA and/or POA.

Fill out the NetWare Agent Worksheet, taking into account the following cluster-specific issues:

GROUPWISE AGENT INSTALLATION WORKSHEET

Under [Item 2: Agents and Locations](#), mark POA Local to Post Office and MTA Local to Domain. In a clustering environment, a domain or post office and its agent must reside on the same GroupWise volume in order to fail over together.

Under [Item 3: Installation Path](#), take into account your decision based on “[Deciding Where to Install the Agent Software](#)” on page 34.

Under [Item 4: Configure GroupWise Agents for Clustering](#), mark Yes. This will cause the Agent Installation program to customize the agent startup files for clustering.

Under [Item 6: Domains](#) and [Item 7: Post Offices](#), refer to the Domain and Post Office Worksheets you filled out in [Section 2.3, “Planning a New Clustered Domain,”](#) on page 25 and [Section 2.4, “Planning a New Clustered Post Office,”](#) on page 26, and to the IP Address Worksheet you completed during “[Planning Secondary IP Addresses and Cluster-Unique Port Numbers for Agents in the Cluster](#)” on page 31.

Under [Item 10: Launch GroupWise Agents Now](#), mark No. You will configure the agents to run in protected mode later.

IMPORTANT: Do not install the NetWare agent software until you are instructed to do so in [Chapter 3, “Setting Up a Domain and Post Office in a NetWare Cluster,”](#) on page 43.

Skip to [Chapter 3, “Setting Up a Domain and Post Office in a NetWare Cluster,”](#) on page 43.

2.9 GroupWise Clustering Worksheets

- ♦ [Section 2.9.1, “System Clustering Worksheet,”](#) on page 37
- ♦ [Section 2.9.2, “IP Address Worksheet,”](#) on page 40
- ♦ [Section 2.9.3, “Agent Clustering Worksheet,”](#) on page 41

2.9.1 System Clustering Worksheet

Item	Explanation
1) Software Version Updates for Cluster:	List any servers that need to be updated so that all nodes in the cluster are running the same version of NetWare. For more information, see Section 2.1, “Meeting Software Version Requirements,” on page 24.

Item	Explanation
2) eDirectory Tree for Cluster:	<p>Record the eDirectory tree where you created the new Novell Cluster object when you installed Novell Cluster Services.</p> <p>For more information, see Section 2.2, “Installing Novell Cluster Services,” on page 24</p>
3) Cluster Name: Cluster IP Address:	<p>Record the name of the new NetWare 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.</p> <p>For more information, see Section 2.2, “Installing Novell Cluster Services,” on page 24.</p>
4) Cluster Context:	<p>Record the full context where you created the new NetWare Cluster object.</p> <p>For more information, see Section 2.2, “Installing Novell Cluster Services,” on page 24.</p>
5) Nodes in Cluster	<p>List the nodes that are part of the cluster that you set up for your GroupWise system.</p> <p>For more information, see Section 2.2, “Installing Novell Cluster Services,” on page 24.</p>
6) Shared Volumes for GroupWise Administration: Cluster Enabled?	<p>Specify the names (<i>cluster_volume</i>) of the shared volumes where the GroupWise administration snap-ins to ConsoleOne and the GroupWise software distribution directory will reside.</p> <p>For cluster-enabling, specify the IP addresses of the virtual servers (<i>volume_SERVER.cluster</i>) to which the cluster-enabled volumes are tied.</p>
<ul style="list-style-type: none"> ◆ Yes (highly recommended) Cluster volume IP addresses 	
<ul style="list-style-type: none"> ◆ No 	<p>For more information, see Section 2.6, “Deciding Whether to Cluster-Enable the Shared Volumes Used by GroupWise,” on page 27.</p>
GroupWise Administration Snap-ins to ConsoleOne	
<ul style="list-style-type: none"> ◆ public directory ◆ Other 	
GroupWise Software Distribution Directory	
<ul style="list-style-type: none"> ◆ <code>\grpwise\software</code> directory ◆ Other 	

Item	Explanation
7) Shared Volume for Domain: Cluster Enabled?	Specify the name (<i>cluster_volume</i>) of the shared volume where the GroupWise domain will reside.
<ul style="list-style-type: none"> ◆ Yes (highly recommended) <p>Cluster volume IP address</p>	For cluster-enabling, specify the IP address of the virtual server (<i>volume_SERVER.cluster</i>) to which the cluster-enabled volume is tied.
<ul style="list-style-type: none"> ◆ No <p>Post Office on Same Volume as Domain?</p>	For more information, see Section 2.4, "Planning a New Clustered Post Office," on page 26 and Section 2.6, "Deciding Whether to Cluster-Enable the Shared Volumes Used by GroupWise," on page 27 .
<ul style="list-style-type: none"> ◆ Yes ◆ No 	
8) Shared Volume for Post Office: Cluster Enabled?	Specify the name (<i>cluster_volume</i>) of the shared volume where the GroupWise post office will reside.
<ul style="list-style-type: none"> ◆ Yes (highly recommended) <p>Cluster volume IP address</p>	For cluster-enabling, specify the IP address of the virtual server (<i>volume_SERVER.cluster</i>) to which the cluster-enabled volume is tied.
<ul style="list-style-type: none"> ◆ No 	For more information, see Section 2.4, "Planning a New Clustered Post Office," on page 26 and Section 2.6, "Deciding Whether to Cluster-Enable the Shared Volumes Used by GroupWise," on page 27 .
9) IP Address Resolution Methods:	Mark the short name address resolution methods you want to implement to ensure that the UNC paths stored in ConsoleOne can be successfully resolved into physical network addresses.
<ul style="list-style-type: none"> ◆ eDirectory ◆ hosts file ◆ DNS ◆ SLP (highly recommended) 	For more information, see Section 2.7, "Ensuring Successful Name Resolution for GroupWise Volumes," on page 29
10) Domain Name: Domain Database Location:	Specify a unique name for the domain. Specify the directory on the GroupWise volume where you want to create the new domain.
	For more information, see Section 2.3, "Planning a New Clustered Domain," on page 25 .
11) Post Office Name: Post Office Database Location:	Specify a unique name for the post office. Specify the directory on the GroupWise volume where you want to create the post office.
	For more information, see Section 2.4, "Planning a New Clustered Post Office," on page 26 .

Item	Explanation
12) Document Storage Area Location:	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.
<ul style="list-style-type: none"> ◆ At the post office ◆ Outside the post office ◆ Separate post office 	For more information, see Section 2.5, “Planning a New Library for a Clustered Post Office,” on page 27.
Document Storage Area Access	
<ul style="list-style-type: none"> ◆ POA <code>/user</code> startup switch setting ◆ POA <code>/password</code> startup switch setting 	

2.9.2 IP Address Worksheet

- ◆ [“Domain Information”](#) on page 40
- ◆ [“Post Office Information”](#) on page 40
- ◆ [“Internet Agent Information”](#) on page 40
- ◆ [“WebAccess Information”](#) on page 41

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

Internet Agent Information

Internet Agent	GWIA IP Address	MTA MTP Port	MTA Live Remote Port	MTA HTTP Port	GWIA HTTP Port
GWIA Domain MTA					N/A
Internet Agent (GWIA)	(same)	N/A	N/A	N/A	

WebAccess Information

WebAccess Agent	WebAccess IP Address	MTA MTP Port	MTA HTTP Port	WebAccess Agent Port	WebAccess HTTP Port
WebAccess Domain MTA				N/A	N/A
WebAccess Agent (GWINTER)	(same)	N/A	N/A		

2.9.3 Agent Clustering Worksheet

Item	Explanation
<p>1) agent installation location:</p> <ul style="list-style-type: none"> ◆ <code>vol:\system</code> on the GroupWise volume ◆ <code>sys:\system</code> on each node <p>Consolidate multiple startup files on GroupWise volume?</p>	<p>Mark the location where you will install the agent software.</p> <p>If necessary, specify the location where you will consolidate multiple agent startup files on a GroupWise volume.</p> <p>For more information, see “Deciding Where to Install the Agent Software” on page 34.</p>
<p>2) Domain Name:</p> <p>Domain Location:</p>	<p>Transfer this information from the System Clustering Worksheet (item 10).</p>
<p>3) Domain Failover Path:</p>	<p>List other nodes in the cluster where the GroupWise domain and its MTA could fail over.</p> <p>For more information, see “Determining Appropriate Failover Paths for the Agents” on page 33.</p>
<p>4) MTA Network Information:</p> <ul style="list-style-type: none"> ◆ MTA IP address ◆ MTA message transfer port ◆ MTA HTTP port 	<p>Gather the MTA network address information from the “IP Address Worksheet” on page 40.</p> <p>For more information, see “Planning Secondary IP Addresses and Cluster-Unique Port Numbers for Agents in the Cluster” on page 31.</p>
<p>5) Post Office Name:</p> <p>Post Office Location:</p>	<p>Transfer this information from the System Clustering Worksheet (item 11).</p>
<p>6) Post Office Failover Path:</p>	<p>List other nodes in the cluster where the GroupWise post office and its POA could fail over.</p> <p>For more information, see “Determining Appropriate Failover Paths for the Agents” on page 33.</p>

Item	Explanation
7) POA Network Information:	Gather the POA network address information from the "IP Address Worksheet" on page 40.
<ul style="list-style-type: none"> ◆ POA IP address ◆ POA client/server port ◆ POA message transfer port ◆ POA HTTP port 	For more information, see "Planning Secondary IP Addresses and Cluster-Unique Port Numbers for Agents in the Cluster" on page 31.
8) Load Agents in Protected Memory?	Mark whether you need to run the agents in protected memory. If so, specify a unique address space for each agent. For the POA, specify a user name and password if required by your version of NetWare.
<ul style="list-style-type: none"> ◆ No ◆ Yes 	For more information, see "Deciding Whether to Run the Agents in Protected Memory" on page 36.
MTA protected address space POA protected address space POA /user startup switch setting POA /password startup switch setting	

Setting Up a Domain and Post Office in a NetWare Cluster

3

You should have already reviewed “[Planning GroupWise in a NetWare Cluster](#)” on page 23 and filled out the “[System Clustering Worksheet](#)” on page 37, the “[IP Address Worksheet](#)” on page 40, and the “[Agent Clustering Worksheet](#)” on page 41. You are now ready to complete the following tasks to set up GroupWise® in a clustering environment:

- ♦ [Section 3.1, “Preparing the Cluster for GroupWise,”](#) on page 43
- ♦ [Section 3.2, “Setting Up a New GroupWise System in a Cluster,”](#) on page 46
- ♦ [Section 3.3, “Creating a New Secondary Domain in a Cluster,”](#) on page 48
- ♦ [Section 3.4, “Creating a New Post Office in a Cluster,”](#) on page 49
- ♦ [Section 3.5, “Installing and Configuring the MTA and the POA in a Cluster,”](#) on page 50
- ♦ [Section 3.6, “Testing Your Clustered GroupWise System,”](#) on page 59
- ♦ [Section 3.7, “Managing Your Clustered GroupWise System,”](#) on page 60
- ♦ [Section 3.8, “What’s Next,”](#) on page 64
- ♦ [Section 3.9, “Clustering Quick Checklists,”](#) on page 65

3.1 Preparing the Cluster for GroupWise

After you have installed Novell® Cluster Services™, as described in *Novell Cluster Services Overview and Installation*, complete the following tasks to prepare the cluster for your GroupWise system:

- ♦ [Section 3.1.1, “Ensuring Required Software Versions,”](#) on page 43
- ♦ [Section 3.1.2, “Cluster-Enabling Shared Volumes for Use with GroupWise,”](#) on page 43
- ♦ [Section 3.1.3, “Configuring Short Name Resolution,”](#) on page 44

3.1.1 Ensuring Required Software Versions

Double-check each node in the cluster to make sure it meets the requirements described in [Section 2.1, “Meeting Software Version Requirements,”](#) on page 24.

3.1.2 Cluster-Enabling Shared Volumes for Use with GroupWise

To cluster-enable a shared volume for use with GroupWise:

- 1 Select a System Clustering Worksheet item (6, 7, or 8) where you marked Yes under Cluster Enabled?.
- 2 Complete the steps in the cluster-enabling section of the cluster documentation for your version of NetWare®, as listed in [Chapter 1, “Introduction to GroupWise 7 and Novell Cluster Services on NetWare,”](#) on page 21.

The System Clustering Worksheet provides the volume to cluster-enable for use the GroupWise, the cluster-enabled volume IP address, and the failover path for the GroupWise volume.

For a review of the new Novell eDirectory™ objects that are created when you cluster-enable a shared volume, see [Section 2.6, “Deciding Whether to Cluster-Enable the Shared Volumes Used by GroupWise,”](#) on page 27.

If you have installed the latest version of ConsoleOne® and the Novell Cluster Services snap-in, you can rename the cluster-related objects in case your DNS name server cannot resolve object names that include the underscore (`_`) character.

- 3 Repeat [Step 1](#) and [Step 2](#) above for the other shared volumes on your System Clustering Worksheet that need to be cluster-enabled.
- 4 Continue with [Configuring Short Name Resolution](#).

3.1.3 Configuring Short Name Resolution

To ensure that GroupWise volumes are always locatable, configure the short name resolution methods that you want to rely on for GroupWise ([System Clustering Worksheet item 9](#)):

- ♦ [“eDirectory”](#) on page 44
- ♦ [“Hosts Files”](#) on page 45
- ♦ [“DNS”](#) on page 45
- ♦ [“SLP”](#) on page 46

After configuring your selected short name resolution methods, continue with the task you need to perform:

- ♦ [“eDirectory”](#) on page 44
- ♦ [“Hosts Files”](#) on page 45
- ♦ [“DNS”](#) on page 45
- ♦ [“SLP”](#) on page 46

eDirectory

Most commonly, you will use eDirectory to resolve the UNC path of a volume into its network address. For example, on the workstation where you run ConsoleOne, you need to map a drive to the location of a domain directory so that ConsoleOne can access the domain database. You could use a `map` command as shown in the example below:

Syntax:

```
map drive: = .cluster_volume.context
```

Example:

```
map m: = .GWCLUSTER_GWVOL1.GWServers
```

When specifying the `map` commands, use [System Clustering Worksheet item 3](#) for *cluster*. Use [System Clustering Worksheet item 7](#) or [8](#) for each *volume* where a domain or post office resides. Use [System Clustering Worksheet item 4](#) for *context*.

Hosts Files

Because each GroupWise volume where you plan to create a domain or post office has been associated with a virtual server, you should add lines for the new virtual servers to one or more of the following files as needed:

- ◆ **NetWare:**
sys:\etc\hosts
(on all nodes in the cluster; recommended)
- ◆ **Windows 2000/2003:**
\winnt\system32\drivers\etc\hosts
(on the administrator's workstation only; optional)

The lines you add to a hosts file could look similar to the following example:

Syntax:

```
IP_address cluster_volume_SERVER.context cluster_volume_SERVER
```

Remember that *cluster_volume_SERVER* represents the name of the virtual server created when you cluster-enabled the volume.

Example:

```
172.16.5.81 gwcluster_gwvol1_SERVER.gwcluster.com  
                                           gwcluster_gwvol1_SERVER
```

When specifying the lines in the hosts files, use [System Clustering Worksheet item 7 or 8](#) for each *IP_address* and *volume* where a domain or post office resides. Use [System Clustering Worksheet item 3](#) for *cluster*. Use [System Clustering Worksheet item 4](#) for *context*.

DNS

Because each GroupWise volume where you plan to create a domain or post office has been associated with a virtual server, you should add all your new virtual servers to DNS. Then you could use a `map` command as shown in the example below (all on one line, of course):

Syntax:

```
map drive: = \\volume_SERVER.cluster.com\volume
```

Remember that *volume_SERVER* represents the name of the Volume Resource object created when you cluster-enabled the volume. A cluster-enabled volume can function like a server, as these commands illustrate.

Example:

```
map m: = \\gwvol1_SERVER.gwcluster.com\gwvol1
```

Or, if the ConsoleOne workstation is in the same DNS domain as the GroupWise volume:

Syntax:

```
map drive: = \\volume_SERVER\volume
```

Example

```
map m: = \\gwvol1_SERVER\gwvol1
```

When specifying the map commands you will need, use [System Clustering Worksheet item 7 or 8](#) for each *volume* where a domain or post office resides. Use [System Clustering Worksheet item 3](#) for *cluster*.

SLP

On NetWare 6.x, Novell Cluster Services automatically propagates virtual server information into SLP and provides the most reliable name resolution.

On NetWare 5.1, Novell Cluster Services does not propagate virtual server information into SLP by default. If you want to use SLP for name resolution on NetWare 5.1, you must download the (unsupported) CVSBIND utility from the Technical Information Document [NWCS Bindery Tool](#) (<http://www.novell.com/support/search.do?cmd=displayKC&docType=kc&externalId=10062624&sliceId=&dialogID=9016836&stateId>). Install CVSBIND according to the instructions included with the download, then determine the server-specific commands you will need to use with CVSBIND.

Syntax:

```
cvsbind add cluster_volume_SERVER ip_address
cvsbind del cluster_volume_SERVER ip_address
```

Remember that *cluster_volume_SERVER* represents the name of the virtual server created when you cluster-enabled the volume.

Example:

```
cvsbind add gwcluster_gwvol1_SERVER 172.16.5.81
cvsbind del gwcluster_gwvol1_SERVER 172.16.5.81
```

Later, in [“Modifying the Volume Resource Load Script for the Agents” on page 53](#) and [“Modifying the Volume Resource Unload Script for the Agents” on page 54](#), you need to add the `cvsbind` commands to the load and unload scripts for GroupWise volume resources.

3.2 Setting Up a New GroupWise System in a 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 on the same GroupWise volume or on two different volumes. 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 as needed.

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

- 1 If necessary, map a drive to each GroupWise administration volume ([System Clustering Worksheet item 6](#)).

- 2 Map a drive to the GroupWise volume for the domain ([System Clustering Worksheet item 7](#)) and, if needed, to the GroupWise volume 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.

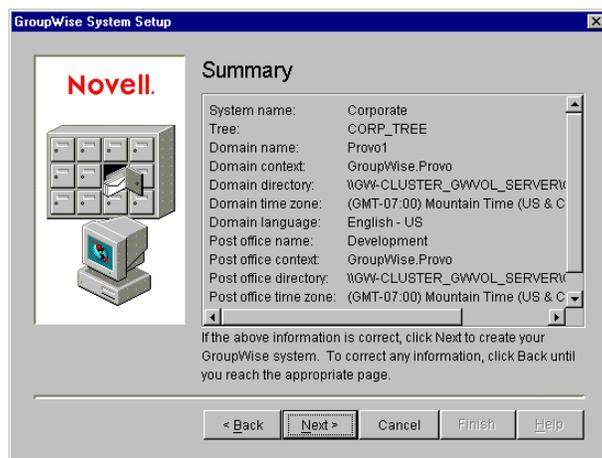
The GroupWise volume name will be *cluster_volume*. For assistance with mapping a drive to a cluster-enabled volume, see [“Configuring Short Name Resolution” on page 44](#).

- 3 Manually create the domain directory ([System Clustering Worksheet item 10](#)) and the post office directory ([System Clustering Worksheet item 12](#)).

This step is not required, but in a clustered environment, the following step is easier if the domain directory already exists.

- 4 Run the GroupWise Installation Advisor to set up your initial GroupWise system, following the steps provided in [“NetWare and Windows: Setting Up a Basic GroupWise System”](#) in [“Installing a Basic GroupWise System”](#) in the *GroupWise 7 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 GroupWise volume accessed in [Step 1](#) above.
- ◆ When you specify the domain directory and post office directory, be sure to browse through the GroupWise volume accessed in [Step 2](#) to select the directory created in [Step 3](#) above.
- ◆ For the post office link type, select *TCP/IP Link*.
- ◆ When providing the MTA and POA network address information, use the Agent Clustering Worksheet that you filled out in [Section 2.8, “Deciding How to Install and Configure the Agents in a Cluster,” on page 31](#). 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 virtual server name with the GroupWise volume.



- 5 When you have finished creating the primary domain and the initial post office, continue with installing the GroupWise Agents, starting with [Step 4 on page 51](#) in [Section 3.5, “Installing and Configuring the MTA and the POA in a Cluster,” on page 50](#).

3.3 Creating a New Secondary Domain in a Cluster

After you have set up the primary domain and initial post office, as described in [Section 3.2](#), “[Setting Up a New GroupWise System in a Cluster](#),” on page 46, you can create additional secondary domains as needed.

To create a new secondary domain in a clustering environment:

- 1 Map a drive to the GroupWise volume for the domain ([System Clustering Worksheet item 7](#)) where the new secondary domain will be created.

The GroupWise volume name will be *cluster_volume*. For assistance with mapping a drive to a cluster-enabled volume, see “[Configuring Short Name Resolution](#)” on page 44.
- 2 Manually create the domain directory ([System Clustering Worksheet item 10](#)).

This step is not required, but in a clustered environment, [Step 5](#) is easier if the domain directory already exists.
- 3 If you selected `vol:\system` on GroupWise volume as the agent installation location (under [Agent Clustering Worksheet item 1](#)), create the `vol:\system` directory on the GroupWise volume accessed in [Step 1](#).

or

If you selected `sys:\system` on each node, decide which node you will install the agents to first.
- 4 In ConsoleOne, connect to the primary domain in your GroupWise system, as described in “[Connecting to a Domain](#)” in “[Domains](#)” in the *GroupWise 7 Administration Guide*.
- 5 Create the new domain, following the steps provided in “[Creating the New Domain](#)” in “[Domains](#)” in the *GroupWise 7 Administration Guide*. Keep in mind the following cluster-specific details:
 - ♦ Use the Domain Worksheet you filled out in [Section 2.3](#), “[Planning a New Clustered Domain](#),” on page 25 to fill in the fields in the Create GroupWise Domain dialog box.
 - ♦ In the *Domain Database Location* field, be sure to browse through the drive you mapped in [Step 1](#) to the domain directory you created in [Step 2](#) 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 Secondary IP Addresses and Cluster-Unique Port Numbers for Agents in the Cluster](#)” on page 31 for the secondary IP address and cluster-unique port numbers that you need to specify in order to configure the link.
- 6 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 7 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.
- 7 Make sure you are still connected to the primary domain.

- 8 Rebuild the domain database for the new domain, following the steps provided in “[Rebuilding Domain or Post Office Databases](#)” in “[Databases](#)” in the *GroupWise 7 Administration Guide*. Be sure to browse to the database location (under [System Clustering Worksheet item 10](#)) through the virtual server that was created when you cluster-enabled the GroupWise volume.
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.
- 9 Continue with [Creating a New Post Office in a Cluster](#).

3.4 Creating a New Post Office in a Cluster

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

To create a new post office in a clustering environment:

- 1 If you marked Yes for Post Office on Same Volume as Domain? (under [System Clustering Worksheet item 7](#)), map a drive to the GroupWise volume for the domain ([System Clustering Worksheet item 7](#)).
or
Map a drive to the GroupWise volume for the post office ([System Clustering Worksheet item 8](#)).
The GroupWise volume name will be *cluster_volume*. For assistance with mapping a drive to a cluster-enabled volume, see “[Configuring Short Name Resolution](#)” on page 44.
- 2 Manually create the post office directory ([System Clustering Worksheet item 12](#)).
This step is not required, but in a clustered environment, [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 7 Administration Guide*.
- 4 Create the new post office, following the steps provided in “[Creating the New Post Office](#)” in “[Post Offices](#)” in the *GroupWise 7 Administration Guide*. Keep in mind the following cluster-specific details:
 - ♦ Use the Post Office Worksheet you filled out in [Section 2.4, “Planning a New Clustered Post Office,”](#) on page 26 to fill in the fields in the Create GroupWise Post Office dialog box.
 - ♦ In the *Post Office Database Location* field, be sure to browse through the drive you mapped 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 14](#)), 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.

- ♦ 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 Secondary IP Addresses and Cluster-Unique Port Numbers for Agents in the Cluster](#)” on [page 31](#) for the secondary IP address and cluster-unique port numbers that you need to specify in order to configure the link.

5 Right-click the new Post Office object, then click *Properties*.

6 Click *GroupWise > Post Office Settings*; 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
- ♦ Enable 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 7 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 7 Administration Guide*. Be sure to browse to the database location ([under System Clustering Worksheet item 11](#)) through the virtual server that was created when you cluster-enabled the GroupWise volume.

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 If you want to create a library with its document storage area outside the post office directory, ([System Clustering Worksheet item 14](#)), follow the steps in “[Setting Up a Basic Library](#)” or “[Setting Up a Full-Service Library](#)” in “[Libraries and Documents](#)” in the *GroupWise 7 Administration Guide*, after you have completely finished setting up the clustered post office.

11 Continue with [Installing and Configuring the MTA and the POA in a Cluster](#).

3.5 Installing and Configuring the MTA and the POA in a 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 volume where you want to install the agent software:

- ♦ [Section 3.5.1, “Installing the Agent Software in a Cluster,” on page 51](#)
- ♦ [Section 3.5.2, “Editing Clustered Agent Startup Files,” on page 52](#)

- ♦ [Section 3.5.3, “Configuring the GroupWise Volume Resource to Load and Unload the Agents,” on page 53](#)
- ♦ [Section 3.5.4, “Setting Up New Instances of the Agents without Installing the Agent Software,” on page 57](#)

Under some circumstances, 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. For example:

- ♦ You have created a new domain and/or post office on a GroupWise volume where the agent software is already installed in the `vol:\system` directory of the GroupWise volume.
- ♦ In your GroupWise system, the agent software is already installed to multiple `sys:\system` directories.

In these circumstances, follow the instructions in [“Setting Up New Instances of the Agents without Installing the Agent Software” on page 57](#) instead of completing the tasks above.

3.5.1 Installing the Agent Software in a Cluster

To install the MTA and the POA:

- 1 Map a drive to the GroupWise volume for the domain ([Agent Clustering Worksheet item 2](#)) or the post office ([Agent Clustering Worksheet item 5](#)).

The GroupWise volume name will be `cluster_volume`. For assistance with mapping a drive to a cluster-enabled volume, see [“Configuring Short Name Resolution” on page 44](#).

- 2 If you selected `vol:\system` on GroupWise volume as the agent installation location (under [Agent Clustering Worksheet item 1](#)), create the `vol:\system` directory on the GroupWise volume accessed in [Step 1](#).

or

If you selected `sys:\system` on each node, decide which node you will install the agents to first.

- 3 Start the Agent Installation program, following the steps provided in [“Installing the NetWare Agent Software”](#) in [“Installing GroupWise Agents”](#) in the *GroupWise 7 Installation Guide*.
- 4 Install the NetWare agents, keeping in mind the following cluster-specific details:
 - ♦ Use the NetWare Agent Clustering Worksheet that you filled out in [“Planning the NetWare Agent Installation” on page 37](#) to fill in the fields during the agent installation process.
 - ♦ In the Installation Path dialog box, be sure to browse through the drive you mapped in [Step 1](#) to the location you chose in [Step 2](#) above. Also select *Configure GroupWise Agents for Clustering*.
 - ♦ In the Domains / Post Offices dialog box, 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. In the *HTTP Port* field, specify the cluster-unique HTTP port planned for each agent (under [Agent Clustering Worksheet items 4 and 7](#)).
 - ♦ In the Installation Complete dialog box, do not select *Launch GroupWise Agents Now*. You will configure the agents to launch in protected mode later.

- 5 If you need to install the agents to `sys:\system` on multiple nodes in the cluster:
 - 5a Repeat [Step 4](#), mapping new drives as needed.
 - 5b If you marked Yes for Consolidate Multiple Startup Files on GroupWise Volume? (under [Agent Clustering Worksheet item 1](#)), copy one complete set of agent startup files and the `grpwise.ncf` file to the planned location, then delete all agent startup files, as well as the `grpwise.ncf` file, from the `sys:\system` directories to avoid future confusion.

The `grpwise.ncf` file includes a `load` command for each instance of each agent. You will use this information later when you create the load and unload scripts for the volume resources.
- 6 Continue with [Editing Clustered Agent Startup Files](#).

3.5.2 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.

Because you selected Configure GroupWise Agents for Clustering during installation, the Agent Installation program set the MTA `/home` startup switch and the POA `/home` startup switch using the format:

```
volume:\directory
```

so that the startup files are valid no matter which node the agents are currently running on.

The Agent Installation program also adds a `/cluster` startup switch to POA startup files to ensure that GroupWise clients detect the clustering environment and try more persistently to reconnect in a failover, fallback, or migration situation.

One additional manual modification of POA startup files is required for robust functionality in a clustering environment. Uncomment the `/ip` startup switch and provide the secondary IP address of the GroupWise volume where the post office is located ([Agent Clustering Worksheet item 7](#)). This information is available to the POA in its eDirectory object properties. However, in some failover situations, reconnection to the MTA is improved when the information is immediately available to the POA in its startup file.

If you are running the POA in protected memory and your version of NetWare requires it, add the `/user` and `/password` startup switches (under [Agent Clustering Worksheet item 8](#)) in order to provide a user name and password that the POA can use to access its post office volume.

If the POA needs to access a remote document storage area, add the `/user` and `/password` startup switches (under [System Clustering Worksheet item 12](#)) in order to provide a user name and password that the POA can use to access the volume 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.

If you have connection problems between the MTA and the POA, you can use the `/user` and `/password` startup switches in the MTA startup file as well.

3.5.3 Configuring the GroupWise Volume Resource to Load and Unload the Agents

The properties of the Volume Resource object define how the GroupWise volume functions within the cluster, how NLM™ programs are loaded and unloaded, and how failover and failback situations are handled. At this point, you might have one volume resource with a domain and post office on it, or you might have two volume resources, one for the domain and one for the post office. Complete the following tasks for each volume resource:

- ♦ “Modifying the Volume Resource Load Script for the Agents” on page 53
- ♦ “Modifying the Volume Resource Unload Script for the Agents” on page 54
- ♦ “Setting the Failover Path and Policies for the Agents” on page 56

Modifying the Volume Resource Load Script for the Agents

The volume resource load script executes whenever the GroupWise volume comes online.

To set up the load script:

- 1 In ConsoleOne, browse to and select the Cluster object.
If necessary, click *View > Console View* to display its contents.
- 2 Right-click the Volume Resource object (*volume_SERVER*), then click *Properties > Load* to display the default volume resource load script for the GroupWise volume.
- 3 Make the following changes to the default load script:
 - ♦ Remove the `trustmig` command. It is not necessary to migrate trustees for a GroupWise volume. Removing this line helps the load script to execute faster.
 - ♦ On NetWare 5.1, if you selected SLP as a short name resolution method, add the `cvsbinding add` command for the GroupWise volume to the load script.

```
cvsbinding add cluster_volume_SERVER IP_address
```

- ♦ If you selected `vol:\system` on GroupWise volume as the agent installation location ([Agent Clustering Worksheet item 1](#)), add a `search add` command to add the new `vol:\system` directory to the server search path.

```
search add volume:\system
```

- ♦ If you selected `sys:\system` on each node as the installation location ([Agent Clustering Worksheet item 1](#)) but you are storing the agent startup files on the GroupWise volume, add that location to the server search path.
- ♦ If you marked No under Load Agents in Protected Memory? ([Agent Clustering Worksheet item 8](#)), add the following `abend` recovery options:

```
set auto restart after abend = 2
set auto restart after abend delay time = 0
set auto restart down timeout = 60
set developer option = off
```

These settings provide the best possible handling of GroupWise databases in the event that an `abend` should occur within the cluster when the agents are not running in protected memory.

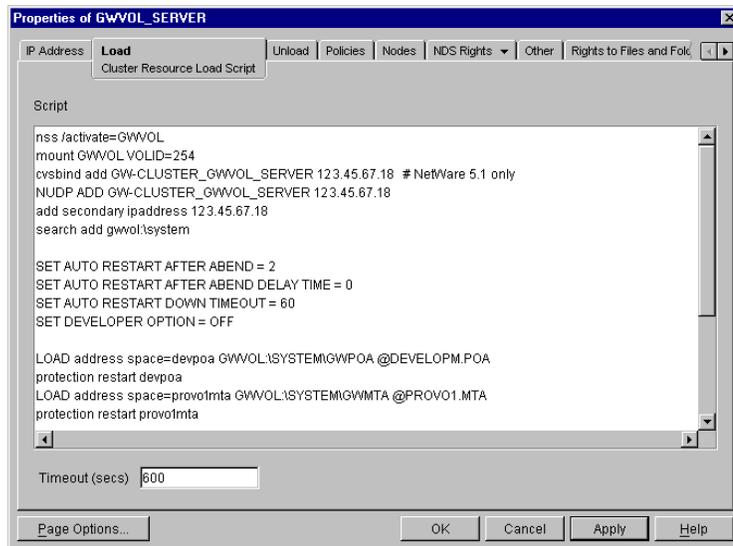
- ◆ Transfer the agent load commands from the `grpwise.ncf` file into the load script. Use `Ctrl+C` to copy and `Ctrl+V` to paste text into the load script page. Then delete or rename the `grpwise.ncf` file to avoid future confusion.

```
load volume:\system\agent.nlm @startup_file
```

- ◆ If you marked Yes under Load Agents in Protected Memory? ([Agent Clustering Worksheet item 8](#)), add the address space parameter to the agent load commands to specify the protected address space for each agent. Add a `protection restart` command for each address space name.

```
load address space=addr_space_name
                        volume:\system\agent.nlm @startup_file
protection restart name
```

The result would look similar to the following example:



NOTE: The `set` commands are needed in the load script only when the agents are *not* running in protected memory. The address space parameters are needed in the `load` commands only when the agents *are* running in protected memory.

- 4 Click *Apply* to save the load script.
- 5 If necessary, click *OK* to confirm that you must offline and then online the volume resource in order for the changes to take effect.
- 6 Continue with [Modifying the Volume Resource Unload Script for the Agents](#).

Modifying the Volume Resource Unload Script for the Agents

The volume resource unload script executes whenever the GroupWise volume 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.

To set up the unload script:

- 1 In ConsoleOne, in the properties pages for the Volume Resource object (*volume_SERVER*), click *Unload* to display the default volume resource unload script.
- 2 Make the following changes to the default unload script:
 - ♦ If you marked Yes under Load Agents in Protected Memory? ([Agent Clustering Worksheet item 8](#)), add an unload address space command for each address space. Add an unload kill address space command to ensure that the address space is completely cleaned up.

```
unload address space=addr_space_name
unload kill address space=addr_space_name
```

If your system seems to be trying to kill the address space before the GroupWise agents have been completely unloaded, resulting in the agents hanging in the unloading state, load the *delay.nlm* program and set a delay of several seconds before issuing the unload kill address space command to allow the GroupWise agents adequate time to unload completely. The length of the delay varies from system to system; ten seconds is a good starting place.

```
unload address space=addr_space_name
load delay.nlm
delay 10
unload kill address space=addr_space_name
```

- ♦ If you marked No under Load Agents in Protected Memory? ([Agent Clustering Worksheet item 8](#)), create an unload command parallel to each load command that you placed in the load script.

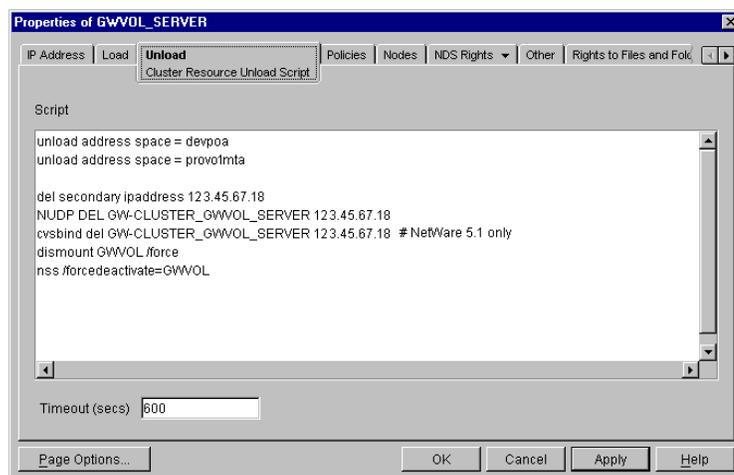
```
unload agent.nlm
```

- ♦ On NetWare 5.1, if you selected SLP as a short name resolution method, add the *cvsbinding del* command for the GroupWise volume to the unload script.

```
cvsbinding del cluster_volume_SERVER IP_address
```

- ♦ Remove the *trustmig* command just like you did in the load script.

The result would look similar to the following example:

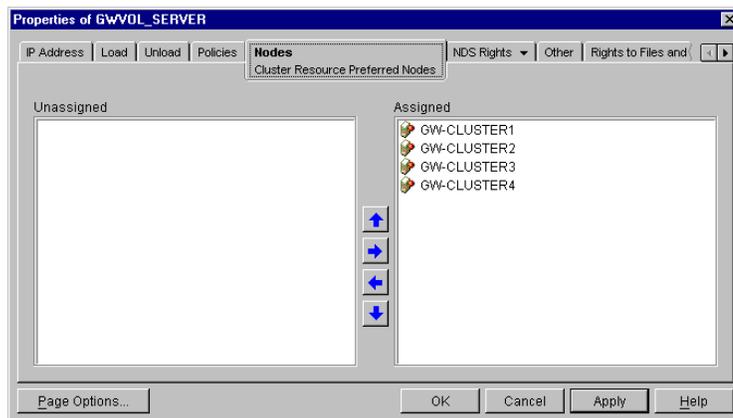


- 3 Click *Apply* to save the unload script.
- 4 If necessary, click *OK* to confirm that you must offline and then online the volume resource in order for the changes to take effect.
- 5 Continue with [Setting the Failover Path and Policies for the Agents](#).

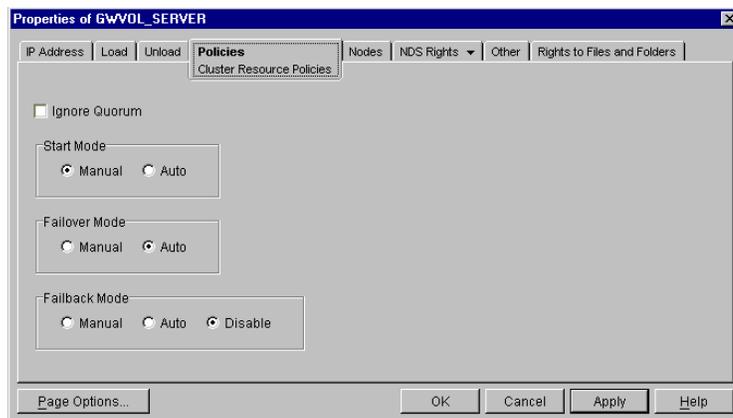
Setting the Failover Path and Policies for the Agents

To modify the failover path and policies for a GroupWise volume resource:

- 1 In ConsoleOne, in the properties pages for the Volume Resource object (*volume_SERVER*), click *Nodes* to display the default failover path for the GroupWise volume resource.



- 2 Arrange the nodes in the cluster into the desired failover path for the domain or post office volume (under [Agent Clustering Worksheet items 3 or 6](#)).
- 3 Click *Apply* to save the failover path.
- 4 Click *Policies* to display the default start, failover, and failback policies.



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

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

If you are considering changing these defaults, see the section about failover and failback modes in the cluster documentation for your version of NetWare, as listed in [Chapter 1, “Introduction to GroupWise 7 and Novell Cluster Services on NetWare,”](#) on page 21.

- 5 Click *OK* when you are finished editing the GroupWise volume resource properties.
- 6 Skip to [Section 3.6, “Testing Your Clustered GroupWise System,”](#) on page 59.

3.5.4 Setting Up New Instances of the 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 57
- ♦ [“Modifying Existing Load and Unload Scripts”](#) on page 57

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. If the existing agent software is located in the `vol:\system` directory of a GroupWise volume, the startup files are there as well. If the existing agent software is located in multiple `sys:\system` directories, the startup files might be located there as well, or they might be in a directory on a GroupWise volume.

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 startup file looking for other active (not commented out) startup switches, then modify them as needed for the new agent.

For example, you might find that the `/httpport` switch is active and needs to be changed to a cluster-unique port number for the new agent.
- 4 Save the new startup file.
- 5 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 volume resource load scripts.

If you created the new domain and/or post office on a new GroupWise volume, skip back to [“Configuring the GroupWise Volume Resource to Load and Unload the Agents”](#) on page 53 for instructions to create new load and unload scripts.

If you created the new domain and/or post office on an existing GroupWise volume, most of the configuration of the volume resource has already been done. You just need to add new `load` and `unload` commands to the existing scripts. Continue with the steps below:

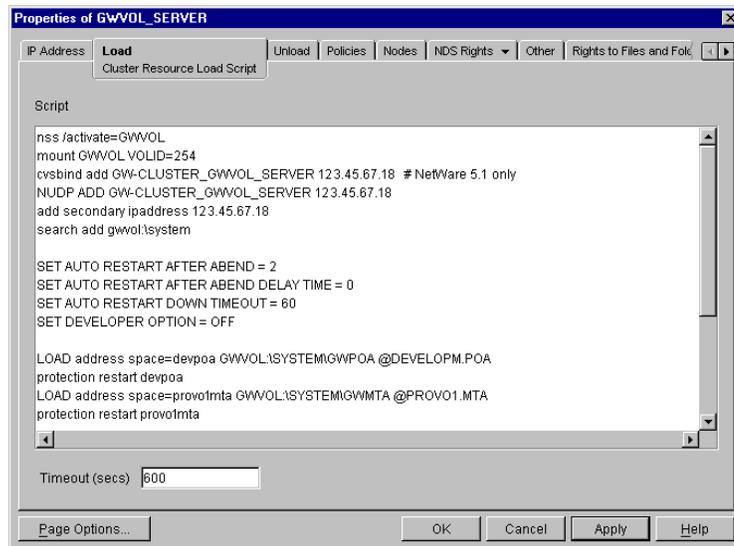
To modify existing load and unload scripts:

- 1 In ConsoleOne, browse to and select the Cluster object.

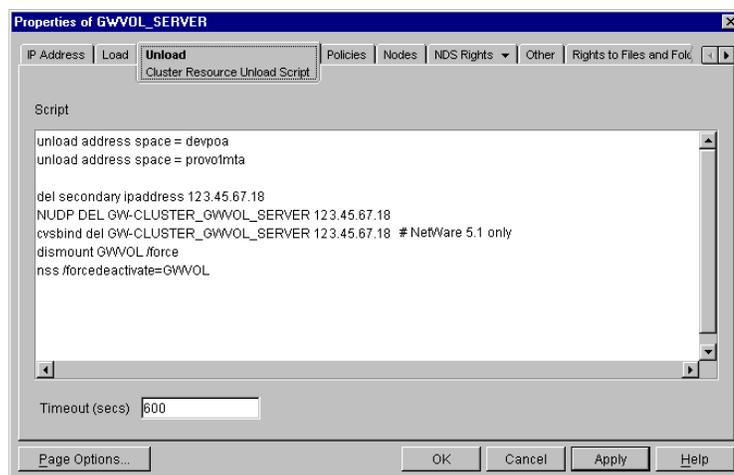
If necessary, click *View > Console View* to display its contents.

- 2 Right-click the Volume Resource object (*volume_SERVER*), then click *Properties > Load* to display the volume resource load script for the GroupWise volume.
- 3 Following the pattern of the existing `load` commands, add `load` 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.

The results would be similar to the following example:



- 4 Click *Apply* to save the modified load script.
- 5 Click *Unload*
- 6 Add corresponding `unload` commands for the new instances of the agents.



- 7 Click *Apply* to save the modified unload script.

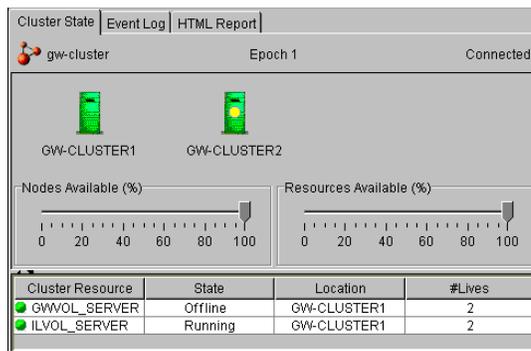
You might want to review other properties of the Volume Resource object, such as the failover path on the Nodes page and the failover/failback/migration procedures on the Policies page, in light of the fact that an additional domain and/or post office now resides on the GroupWise volume.

- 8 Change other *Volume Resource* properties as needed.
- 9 Click *OK* to save the modified properties.
- 10 In the Cluster State View, take the GroupWise volume 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 [Section 3.6, “Testing Your Clustered GroupWise System,” on page 59](#).

3.6 Testing Your Clustered GroupWise System

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

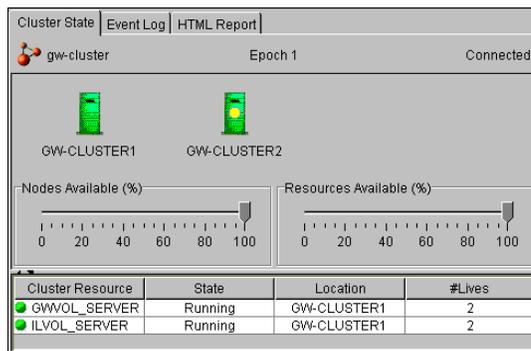
- 1 In ConsoleOne, select the Cluster object, then click *View > Cluster State*.



The new GroupWise volume resource shows *Offline* in the *State* column.

- 2 Click the new GroupWise volume resource, then click *Online*.

The *State* column for the volume resource now displays *Running*.



- 3 Observe the server console where the MTA and/or POA are loading to see that they start and run correctly.

If you are using protected memory, you can use the `protection` command at the server console prompt to list all the address spaces on the node and what NLM programs are running in each one.

- 4 Click the new GroupWise volume resource, then click *Offline*.
The *State* column for the volume resource returns to Offline.
- 5 Observe the server console where the MTA and/or POA are unloading to see that they shut down correctly.
If you are using protected memory, you can use the `protection` command again to make sure that the address spaces used by the GroupWise agents are no longer present.
- 6 Repeat **Step 2** whenever you are ready to bring the new GroupWise volume resource online permanently.
On NetWare 6.x, these actions can also be performed from your Web browser. See [“Using Novell Remote Manager on NetWare 6.x” on page 61](#).
- 7 Continue with [Managing Your Clustered GroupWise System](#).

3.7 Managing Your Clustered GroupWise System

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

- ♦ [Section 3.7.1, “Updating GroupWise Objects with Cluster-Specific Descriptions,” on page 60](#)
- ♦ [Section 3.7.2, “Using Novell Remote Manager on NetWare 6.x,” on page 61](#)
- ♦ [Section 3.7.3, “Knowing What to Expect in MTA and POA Failover Situations,” on page 64](#)

3.7.1 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 that 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” on page 60](#)
- ♦ [“Recording Cluster-Specific Information for a Post Office and Its POA” on page 61](#)
- ♦ [“Recording Cluster-Specific Information for a Software Distribution Directory” on page 61](#)

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 cluster-enabled volume 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 secondary IP address of the cluster-enabled domain volume 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 secondary IP address of its cluster-enabled volume 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 secondary IP address of the cluster-enabled post office volume 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 61.

or

Skip to [“Knowing What to Expect in MTA and POA Failover Situations”](#) on page 64.

Recording Cluster-Specific Information for a Software Distribution Directory

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

- 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 cluster-enabled volume where the software distribution directory resides.
- 4 Click *OK*, then click *Close* to save the software distribution directory description.
- 5 Skip to [“Knowing What to Expect in MTA and POA Failover Situations”](#) on page 64.

3.7.2 Using Novell Remote Manager on NetWare 6.x

On NetWare 6.x, you can use Novell Remote Manager to manage many aspects of your GroupWise cluster from your Web browser. For instructions on setting up and accessing this useful network administration utility, see the NetWare 6.x *Novell Remote Manager Administration Guide* at the [Novell Documentation Web site \(http://www.novell.com/documentation\)](http://www.novell.com/documentation). Cluster management features are automatically added to Novell Remote Manager when you install Novell Cluster Services.

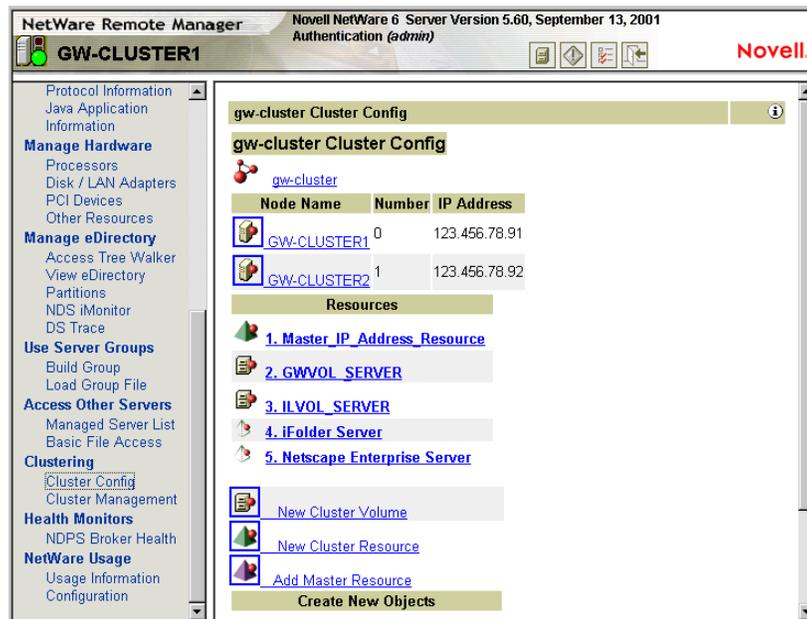
After you have accessed Novell Remote Manager, you might find that many GroupWise cluster management tasks are easier to perform with Novell Remote Manager than with ConsoleOne. The following sections help you configure and manage the cluster using Novell Remote Manager:

- ♦ “Configuring Your GroupWise Cluster” on page 62
- ♦ “Managing Your GroupWise Cluster” on page 63

Configuring Your GroupWise Cluster

On the main Novell Remote Manager page:

- 1 In the left frame, scroll down to the *Clustering* section, then click *Cluster Config*.



The Cluster Configuration page displays the cluster name, the nodes in the cluster, and the resources in the cluster. It also enables you to create new GroupWise Volume Resource objects (termed Cluster Volumes in the Novell Remote Manager interface).

- 2 Click the cluster name to display the Cluster object properties:

Cluster Configuration Fields: gw-cluster

Revision	261
IP Address	123.456.78.99
Port Number	7023
Quorum	
Membership	2
Timeout	60
Protocol	
Tolerance	8
Heartbeat	1
Master Watchdog	1
Slave Watchdog	8
Max Retransmit	30
GIPC Stream	
Resource Priorities	
Email Reporting	

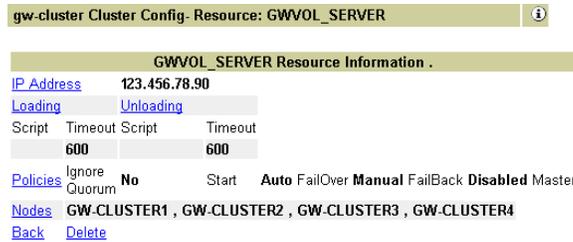
Click a linked item to edit the Cluster object properties. Click your browser’s *Back* button to return to the Cluster Configuration page.

- 3 On the Cluster Configuration page, click a server to display the Server object properties:



Click a linked item to edit the Server object properties. Click *Back* or *Delete* to perform the specified action.

- 4 On the Cluster Configuration page, click a GroupWise volume to display the Volume Resource object properties:



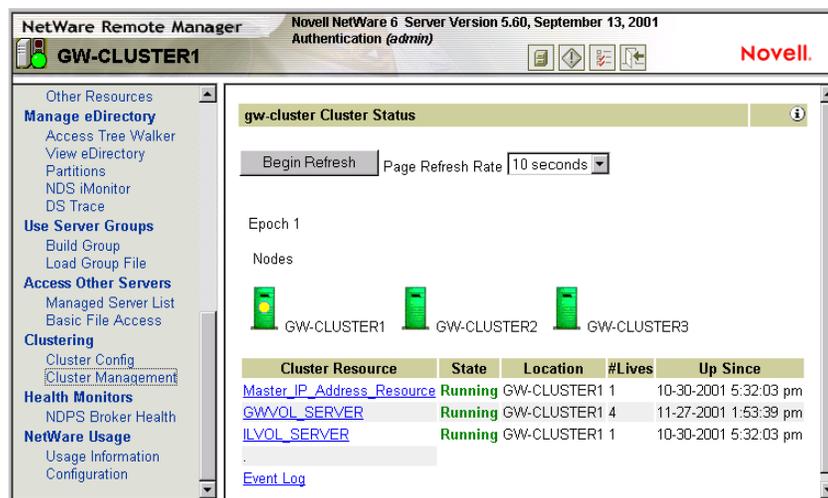
Click a linked item to edit Volume Resource object properties. Click *Back* or *Delete* to perform the specified action.

- 5 On the Cluster Configuration page, click *New Cluster Volume* to create a new GroupWise Volume Resource object, then follow the instructions to provide the information needed to create the new Cluster Volume object.

Managing Your GroupWise Cluster

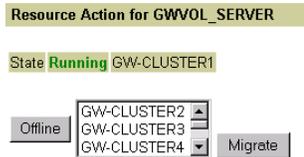
On the main Novell Remote Manager page:

- 1 In the left frame, scroll down to the Clustering section, then click *Cluster Management*.



The Cluster Status page displays the nodes and volume resources in the cluster. The master node in the cluster is marked with a yellow ball. Status information is listed for each volume resource. You can set the refresh rate for the status information at the top of the Cluster Status page.

- 2 Select a page refresh rate, then click *Begin Refresh* so that the page automatically refreshes at the selected rate.
- 3 Click a cluster resource to bring it online, take it offline, or migrate it to another node.



To take the currently running volume resource offline, click *Offline*. To migrate the volume resource, select a node from the drop-down list, then click *Migrate*.

- 4 On the Cluster Resource page, click *Event Log* to view a list of cluster events.
The event log can help you resolve problems with cluster functioning.

3.7.3 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 would need to restart the GroupWise client would be 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 and then 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.

In comparison to failover, 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](#).

3.8 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 7 Administration Guide*.
- ♦ Setting up the GroupWise client software and helping users to get started using it. See “[Client](#)” in the *GroupWise 7 Administration Guide*. Also see the *GroupWise 7 Windows Client User Guide*.
- ♦ Connecting your clustered GroupWise system to the Internet. See [Chapter 4, “Implementing the Internet Agent in a NetWare Cluster,”](#) on page 69.

- ♦ Providing access to users' GroupWise mailboxes from their Web browsers. See [Chapter 5, "Implementing WebAccess in a NetWare Cluster,"](#) on page 89.
- ♦ Connecting your clustered GroupWise system to other e-mail systems through GroupWise gateways. See [Chapter 6, "Implementing GroupWise Gateways in a Novell Cluster,"](#) on page 107.
- ♦ Monitoring the status of your clustered GroupWise system from your Web browser. See [Chapter 7, "Monitoring a GroupWise System in a NetWare Cluster,"](#) on page 109.
- ♦ Backing up your clustered GroupWise system. See [Chapter 8, "Backing Up a GroupWise System in a NetWare Cluster,"](#) on page 111.

3.9 Clustering Quick Checklists

- ♦ [Section 3.9.1, "GroupWise System Quick Checklist,"](#) on page 65
- ♦ [Section 3.9.2, "Domain Quick Checklist,"](#) on page 66
- ♦ [Section 3.9.3, "Post Office Quick Checklist,"](#) on page 67

3.9.1 GroupWise System Quick Checklist

- Plan your new clustered GroupWise system.
See [Chapter 2, "Planning GroupWise in a NetWare Cluster,"](#) on page 23.
- Cluster-enable the volumes where GroupWise domains and post offices will reside.
See [Section 3.1.2, "Cluster-Enabling Shared Volumes for Use with GroupWise,"](#) on page 43.
- Make sure that short name resolution works throughout your network.
See [Section 3.1.3, "Configuring Short Name Resolution,"](#) on page 44.
- Create the primary domain and initial post office in your new clustered GroupWise system.
See [Section 3.2, "Setting Up a New GroupWise System in a Cluster,"](#) on page 46.
- Set up the agents for the primary domain and the initial post office.
See [Section 3.5, "Installing and Configuring the MTA and the POA in a Cluster,"](#) on page 50.
- Modify the volume resource load script(s):
 - ♦ Remove the `trustmig` command
 - ♦ Add the `cvsbind add` command (NetWare 5.1 only; optional)
 - ♦ Add the `search add` command (optional)
 - ♦ If you will *not* run the agents in protected memory, add the `set auto restart` commands and the `set developer option = off` command
 - ♦ Add the agent load command(s)
 - ♦ If you *will* run the agents in protected memory, add the `address space` parameter to the load command(s) and add a corresponding `protection restart` command for each address space
 See ["Modifying the Volume Resource Load Script for the Agents"](#) on page 53.
- Modify the volume resource unload script(s):
 - ♦ Add the agent or address space `unload` command(s)

- ◆ Add the `cvsbinding del` command if you used the `cvsbinding add` command in the load script (NetWare 5.1 only; optional)
- ◆ Remove the `trustmig` command

See “[Modifying the Volume Resource Unload Script for the Agents](#)” on page 54.

- ❑ Set up the volume failover path(s) and policies.

See “[Setting the Failover Path and Policies for the Agents](#)” on page 56.

- ❑ Test your new clustered GroupWise system.

See [Section 3.6, “Testing Your Clustered GroupWise System,”](#) on page 59.

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

See [Section 3.7, “Managing Your Clustered GroupWise System,”](#) on page 60.

3.9.2 Domain Quick Checklist

- ❑ Plan your new clustered domain.

See [Section 2.3, “Planning a New Clustered Domain,”](#) on page 25.

- ❑ Cluster-enable the volume where the domain will reside.

See [Section 3.1.2, “Cluster-Enabling Shared Volumes for Use with GroupWise,”](#) on page 43.

- ❑ Make sure that short name resolution for the new domain volume works throughout your network.

See [Section 3.1.3, “Configuring Short Name Resolution,”](#) on page 44.

- ❑ Create the new domain.

See [Section 3.3, “Creating a New Secondary Domain in a Cluster,”](#) on page 48.

- ❑ Set up the MTA for the new domain.

See [Section 3.5, “Installing and Configuring the MTA and the POA in a Cluster,”](#) on page 50.

- ❑ Modify the domain volume resource load script:

- ◆ Remove the `trustmig` command
- ◆ Add the `cvsbinding add` command (NetWare 5.1 only; optional)
- ◆ Add the `search add` command (optional)
- ◆ If you will *not* run the MTA in protected memory, add the `set auto restart` commands and the `set developer option = off` command
- ◆ Add the MTA `load` command
- ◆ If you *will* run the MTA in protected memory, add the `address space` parameter to the MTA `load` command and add a corresponding `protection restart` command for the address space

See “[Modifying the Volume Resource Load Script for the Agents](#)” on page 53.

- ❑ Modify the domain volume resource unload script:

- ◆ Add the MTA or address space `unload` command

- ◆ Add the `cvswind del` command if you used the `cvswind add` command in the load script (NetWare 5.1 only; optional)
- ◆ Remove the `trustmig` command

See “[Modifying the Volume Resource Unload Script for the Agents](#)” on page 54.

- ❑ Set up the domain volume failover path and policies.

See “[Setting the Failover Path and Policies for the Agents](#)” on page 56.

- ❑ Test your new clustered domain.

See [Section 3.6, “Testing Your Clustered GroupWise System,”](#) on page 59.

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

See [Section 3.7, “Managing Your Clustered GroupWise System,”](#) on page 60.

3.9.3 Post Office Quick Checklist

- ❑ Plan your new clustered post office.

See [Section 2.4, “Planning a New Clustered Post Office,”](#) on page 26.

- ❑ Cluster-enable the volume where the post office will reside.

See [Section 3.1.2, “Cluster-Enabling Shared Volumes for Use with GroupWise,”](#) on page 43.

- ❑ Make sure that short name resolution for the new post office volume works throughout your network.

See [Section 3.1.3, “Configuring Short Name Resolution,”](#) on page 44.

- ❑ Create the new post office.

See [Section 3.4, “Creating a New Post Office in a Cluster,”](#) on page 49.

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

See [Section 3.5, “Installing and Configuring the MTA and the POA in a Cluster,”](#) on page 50.

- ❑ Add the `/ip` startup switch to the POA startup file in order to provide the secondary IP address of the post office volume. If you are running the POA in protected memory, add the `/user` and `/password` startup switches so the POA can access the volume.

See “[Editing Clustered Agent Startup Files](#)” on page 52.

- ❑ Modify the post office volume resource load script:

- ◆ Remove the `trustmig` command
- ◆ Add the `cvswind add` command (NetWare 5.1 only; optional)
- ◆ Add the `search add` command (optional)
- ◆ If you will *not* run the POA in protected memory, add the `set auto restart` commands and the `set developer option = off` command
- ◆ Add the POA `load` command
- ◆ If you *will* run the POA in protected memory, add the `address space` parameter to the POA `load` command and add a corresponding `protection restart` command for the address space

See “[Modifying the Volume Resource Load Script for the Agents](#)” on page 53.

- ❑ Modify the post office volume resource unload script:
 - ◆ Add the POA or address space `unload` command
 - ◆ Add the `cvsbind del` command if you used the `cvsbind add` command in the load script (NetWare 5.1 only; optional)
 - ◆ Remove the `trustmig` command

See [“Modifying the Volume Resource Unload Script for the Agents”](#) on page 54.

- ❑ Set up the post office volume failover path and policies.

See [“Setting the Failover Path and Policies for the Agents”](#) on page 56.

- ❑ Test your new clustered post office.

See [Section 3.6, “Testing Your Clustered GroupWise System,”](#) on page 59.

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

See [Section 3.7, “Managing Your Clustered GroupWise System,”](#) on page 60.

Implementing the Internet Agent in a NetWare Cluster

4

You should already have set up at least a basic GroupWise[®] system, as described in [Chapter 2, “Planning GroupWise in a NetWare Cluster,” on page 23](#) and [Chapter 3, “Setting Up a Domain and Post Office in a NetWare Cluster,” on page 43](#). As part of this process, the [“System Clustering Worksheet” on page 37](#) and the [“IP Address Worksheet” on page 40](#) 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 in a cluster.

- ◆ [Section 4.1, “Planning the Internet Agent in a Cluster,” on page 69](#)
- ◆ [Section 4.2, “Setting Up the Internet Agent in a Cluster,” on page 73](#)
- ◆ [Section 4.3, “Managing the Internet Agent in a Cluster,” on page 83](#)
- ◆ [Section 4.4, “Internet Agent Clustering Worksheet,” on page 84](#)
- ◆ [Section 4.5, “Internet Agent Quick Checklist,” on page 86](#)

4.1 Planning the Internet Agent in a 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 each GroupWise gateway, including the Internet Agent.

The [Section 4.4, “Internet Agent Clustering Worksheet,” on page 84](#) lists all the information you need as you set up the Internet Agent in a clustering environment. You should print the worksheet and fill it out as you complete the tasks listed below:

- ◆ [Section 4.1.1, “Planning a Domain for the Internet Agent,” on page 70](#)
- ◆ [Section 4.1.2, “Deciding Whether to Cluster-Enable the Internet Agent Volume,” on page 70](#)
- ◆ [Section 4.1.3, “Determining an Appropriate Failover Path for the Internet Agent Volume,” on page 70](#)
- ◆ [Section 4.1.4, “Planning a Secondary IP Address and Cluster-Unique Port Numbers for the Internet Agent and Its MTA,” on page 71](#)
- ◆ [Section 4.1.5, “Preparing Your Firewall for the Internet Agent,” on page 71](#)
- ◆ [Section 4.1.6, “Deciding Where to Install the Internet Agent and Its MTA,” on page 72](#)
- ◆ [Section 4.1.7, “Deciding Whether to Run the Internet Agent and Its MTA in Protected Memory,” on page 72](#)
- ◆ [Section 4.1.8, “Planning the MTA Installation,” on page 73](#)
- ◆ [Section 4.1.9, “Planning the Internet Agent Installation,” on page 73](#)

4.1.1 Planning a Domain for the Internet Agent

The considerations involved in planning a domain for the Internet Agent are much the same as planning any other domain. In preparation, review “[Planning a New Domain](#)”, then print and fill out the “[Domain Worksheet](#)” in “[Domains](#)” in the *GroupWise 7 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 volume where you want the domain directory to reside.
- ◆ 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: Shared Volume for Internet Agent](#), transfer the domain location to the Internet Agent Clustering Worksheet.

Under [Item 2: Internet Agent Domain Name](#), transfer the domain name and database directory to the Internet Agent Clustering Worksheet.

4.1.2 Deciding Whether to Cluster-Enable the Internet Agent Volume

You should plan to cluster-enable the shared volume where the Internet Agent domain will reside. For a review of the benefits of cluster-enabling volumes, see [Section 2.6, “Deciding Whether to Cluster-Enable the Shared Volumes Used by GroupWise,”](#) on [page 27](#), which describes the issues in the context of planning MTA and POA installations.

INTERNET AGENT CLUSTERING WORKSHEET

Under [Item 1: Shared Volume for Internet Agent](#), mark Yes under Cluster Enabled?.

Cluster-enabling relies on successful short name resolution throughout your system. Review [Section 2.7, “Ensuring Successful Name Resolution for GroupWise Volumes,”](#) on [page 29](#), which describes the issues in the context of planning MTA and POA installations.

4.1.3 Determining an Appropriate Failover Path for the Internet Agent Volume

As with the MTA and the POA, you need to decide which nodes in the cluster would be appropriate locations for the Internet Agent volume to fail over to. For a review of failover paths, see [“Determining Appropriate Failover Paths for the Agents”](#) on [page 33](#), which describes the issues in the context of planning MTA and POA installations.

INTERNET AGENT CLUSTERING WORKSHEET

Under **Item 3: Internet Agent Failover Path**, list the nodes that you want to have in the Internet Agent volume failover path.

4.1.4 Planning a Secondary IP Address and Cluster-Unique Port Numbers for the Internet Agent and Its MTA

As with the MTA and the POA, the Internet Agent needs a secondary IP address and cluster-unique port numbers. As part of planning to install the MTA and POA, you should already have determined the secondary IP address and cluster-unique port numbers for the Internet Agent and its MTA as you filled out the **“IP Address Worksheet” on page 40**. 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 MTA secondary IP address and cluster-unique port numbers from the Internet Agent section of the IP Address Worksheet to the Internet Agent Clustering Worksheet.

Under **Item 1: Shared Volume for Internet Agent**, copy the MTA secondary IP address under Cluster Volume IP Address as well, because they are the same.

Under **Item 7: Internet Agent Network Information**, transfer the Internet Agent secondary IP address (the same as for its MTA) and the cluster-unique Internet Agent port number from the Internet Agent section of the IP Address Worksheet to the Internet Agent Clustering Worksheet.

4.1.5 Preparing Your Firewall for the Internet Agent

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

Table 4-1 *Standard Ports*

Protocol	Standard Port
IMAP4	143
LDAP	389
POP3	110
SMTP	25

By default, the Internet Agent 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 Internet Agent volume failover path.

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

As another alternative, you can configure the Internet Agent to use its secondary IP address for sending as well as receiving messages. Setup instructions for this configuration are provided in [“Forcing Use of the Internet Agent Secondary IP Address” on page 81](#), which you can complete after installing the Internet Agent.

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

4.1.6 Deciding Where to Install the Internet Agent and Its MTA

As with the MTA and the POA, you can choose to install the Internet Agent and its MTA to the `sys:\system` directory of each node or to a `vol:\system` directory on the Internet Agent volume. For a discussion of these alternatives, see [“Deciding Where to Install the Agent Software” on page 34](#), which describes the issues in the context of planning MTA and POA installations. If you only have one Internet Agent for your GroupWise system with several servers in its failover path, it is an easy choice: Install it once to a `vol:\system` directory on the Internet Agent volume.

INTERNET AGENT CLUSTERING WORKSHEET

Under [Item 4: MTA Installation Location](#) and [Item 6: Internet Agent Installation Location](#), mark whether you will install the Internet Agent and its MTA to a `vol:\system` directory on the Internet Agent volume or to `sys:\system` on each node in the cluster. If necessary, specify where the MTA startup file and the Internet Agent configuration file will be stored.

4.1.7 Deciding Whether to Run the Internet Agent and Its MTA in Protected Memory

As with the MTA and the POA, you can choose whether to run the Internet Agent in protected memory. For a review of the benefits of protected memory, see [“Deciding Whether to Run the Agents in Protected Memory” on page 36](#), which describes the issues in the context of planning MTA and POA installations.

You might think that protected memory is not necessary if you have only one Internet Agent for your GroupWise system because it could never fail over to a node where another Internet Agent was running. However, because the Internet Agent in a cluster is installed into its own domain with its own MTA, this MTA could fail over to a node where another MTA was already running. Therefore, it is safest to load the MTA into protected memory. Loading the Internet Agent into protected memory is also recommended. Load each agent into its own memory space and mark each memory space as restartable.

INTERNET AGENT CLUSTERING WORKSHEET

Under [Item 8: Load Internet Agent and Its MTA in Protected Memory?](#), mark whether you need to run the Internet Agent and its MTA in protected memory. If you do, provide a protected memory address space name for each agent.

4.1.8 Planning the MTA Installation

Follow the instructions in [“Planning the NetWare Agent Installation” on page 37](#), then return to this point. After you follow the instructions, you have a filled-out NetWare Agent Worksheet to use when you install the MTA.

IMPORTANT: Do not install the NetWare MTA until you are instructed to do so in [Section 4.2, “Setting Up the Internet Agent in a Cluster,” on page 73](#).

4.1.9 Planning the Internet Agent Installation

Aside from the cluster-specific issues discussed in the preceding sections, the considerations involved in planning to install the Internet Agent are the same in a clustering environment as for any other environment. Review the installation instructions in [“NetWare and Windows: Installing the Internet Agent Software”](#) in [“Installing the GroupWise Internet Agent”](#) in the *GroupWise 7 Installation Guide*. You might want to print this section and write down the types of planning information you have provided on worksheets in other sections. You will need this information as you install the Internet Agent in your cluster.

IMPORTANT: Do not install the Internet Agent software until you are instructed to do so in [Section 4.2, “Setting Up the Internet Agent in a Cluster,” on page 73](#).

4.2 Setting Up the Internet Agent in a Cluster

You should already have reviewed [Section 4.1, “Planning the Internet Agent in a Cluster,” on page 69](#) and filled out the [Section 4.4, “Internet Agent Clustering Worksheet,” on page 84](#). You are now ready to complete the following tasks to set up the Internet Agent in a clustering environment:

- ♦ [Section 4.2.1, “Cluster-Enabling a Shared Volume for Use with the Internet Agent,” on page 73](#)
- ♦ [Section 4.2.2, “Creating a Domain for the Internet Agent,” on page 74](#)
- ♦ [Section 4.2.3, “Installing the MTA for the Internet Agent Domain,” on page 74](#)
- ♦ [Section 4.2.4, “Installing and Configuring the Internet Agent in a Cluster,” on page 74](#)
- ♦ [Section 4.2.5, “Testing the Clustered Internet Agent,” on page 82](#)

4.2.1 Cluster-Enabling a Shared Volume for Use with the Internet Agent

To cluster-enable the Internet Agent shared volume:

- 1 Complete the cluster-enabling steps in the applicable section of the cluster documentation for your version of NetWare, as listed in [Chapter 1, “Introduction to GroupWise 7 and Novell Cluster Services on NetWare,” on page 21](#).

The Internet Agent Clustering Worksheet provides the volume to cluster-enable, the cluster-enabled volume IP address, and the failover path for the Internet Agent volume.

For a review of the new Novell® eDirectory™ objects that are created when you cluster-enable a shared volume, see [Section 2.6, “Deciding Whether to Cluster-Enable the Shared Volumes Used by GroupWise,” on page 27](#).

If you have installed the latest version of ConsoleOne® and the Novell Cluster Services snap-in, you can rename the cluster-related objects in case your DNS name server cannot resolve object names that include the underscore (_) character.

- 2 To ensure successful short name resolution, add entries for the Internet Agent virtual server to support your preferred methods of short name resolution, as described in [“Configuring Short Name Resolution” on page 44](#).
- 3 To ensure that the Internet Agent has incoming and outgoing access to the Internet, make sure your firewall is properly configured, as described in [“Preparing Your Firewall for the Internet Agent” on page 71](#).
- 4 Continue with [Creating a Domain for the Internet Agent](#).

4.2.2 Creating a Domain for the Internet Agent

The Internet Agent domain will be a secondary domain. To create it, follow the instructions in [Section 3.3, “Creating a New Secondary Domain in a Cluster,” on page 48](#), 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 Internet Agent domain.

Continue with [Installing the MTA for the Internet Agent Domain](#).

4.2.3 Installing the MTA for the Internet Agent Domain

The MTA for the Internet Agent domain can be installed just like any other MTA in your clustered GroupWise system. Follow the instructions in [“Installing the Agent Software in a Cluster” on page 51](#), then return to this point.

You do not need to edit the MTA startup file. You do not need to modify the Volume Resource properties until after you have installed the Internet Agent.

Continue with [Installing and Configuring the Internet Agent in a Cluster](#).

4.2.4 Installing and Configuring the Internet Agent in a Cluster

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

- ♦ [“Installing the Internet Agent Software in a Cluster” on page 74](#)
- ♦ [“Configuring the Internet Agent Volume Resource to Load and Unload the Internet Agent and Its MTA” on page 75](#)
- ♦ [“Enabling Internet Addressing for Your Clustered GroupWise System” on page 80](#)
- ♦ [“Verifying Internet Agent Object Properties” on page 80](#)

Installing the Internet Agent Software in a Cluster

- 1 Map a drive to the Internet Agent volume ([Internet Agent Clustering Worksheet item 1](#)).

The Internet Agent volume name will be *cluster_volume*. For assistance with mapping a drive to a cluster-enabled volume, see [“Configuring Short Name Resolution” on page 44](#).

2 If you selected `vol:\system` on Internet Agent Volume as the Internet Agent installation location ([Internet Agent Clustering Worksheet item 6](#)), create the `vol:\system` directory on the Internet Agent volume accessed in [Step 1](#).

or

If you selected `sys:\system` on Each Node, decide which node you will install the Internet Agent to first, then map a drive to `sys:\system` on that node.

3 Start the Internet Agent Installation program and install the NetWare® Internet Agent, following the steps provided in “[NetWare and Windows: Installing the Internet Agent Software](#)” in “[Installing the GroupWise Internet Agent](#)” in the *GroupWise 7 Installation Guide*. Keep in mind the following cluster-specific details:

- ♦ Use the notes you made during “[Planning the Internet Agent Installation](#)” on [page 73](#) to fill in the fields during the Internet Agent installation process.
- ♦ In the Installation Path dialog box, be sure to browse through the drive you mapped to the location you chose in [Step 2](#) above. Deselect *Update AUTOEXEC File* and select *Configure GroupWise Agents for Clustering*.
- ♦ In the GroupWise Domain dialog box, be sure to browse through the drive you mapped in [Step 1](#) to the domain directory ([Internet Agent Clustering Worksheet item 2](#)).
- ♦ The Internet Agent Installation program creates the `gwia.ncf` file, which includes the `load` command for the Internet Agent. You need this information later when you create the load script for the Volume Resource object.

4 If you need to install the Internet Agent to `sys:\system` to each node in the cluster:

4a Repeat [Step 3](#), mapping new drives as needed.

4b If you marked Yes for Consolidate Multiple Configuration Files on Internet Agent Volume? (under [Internet Agent Clustering Worksheet item 6](#)), copy the `gwia.cfg` file to the planned location, then delete it from the `sys:\system` directories to avoid future confusion.

5 Make sure you have completed all the tasks described in “[Installing the GroupWise Internet Agent](#)” in the *GroupWise 7 Installation Guide*.

The Internet Agent starts automatically immediately after Installation.

6 Stop each Internet Agent you have installed before configuring it for clustering.

7 Continue with [Configuring the Internet Agent Volume Resource to Load and Unload the Internet Agent and Its MTA](#).

Configuring the Internet Agent Volume Resource to Load and Unload the Internet Agent and Its MTA

The properties of the Volume Resource object define how the Internet Agent volume functions within the cluster, how NLM™ programs are loaded and unloaded, and how failover and failback situations are handled. Complete the following tasks for the Internet Agent volume:

- ♦ “[Modifying the Volume Resource Load Script for the Internet Agent](#)” on [page 75](#)
- ♦ “[Modifying the Volume Resource Unload Script for the Internet Agent](#)” on [page 77](#)
- ♦ “[Setting the Failover Path and Policies for the Internet Agent](#)” on [page 79](#)

Modifying the Volume Resource Load Script for the Internet Agent

The volume resource load script executes whenever the Internet Agent volume comes online.

To set up the load script:

- 1 In ConsoleOne, browse to and select the Cluster object.

If necessary, click *View > Console View* to display its contents.

- 2 Right-click the Volume Resource object (*volume_SERVER*), then click *Properties > Load* to display the default volume resource load script for the Internet Agent volume.

The next step assumes that this is the first time you have edited this load script. If other GroupWise agents are already running from this volume, some of the modifications have already been made.

- 3 Make the following changes to the default load script:

- ♦ Remove the `trustmig` command. It is not necessary to migrate trustees for the Internet Agent volume. Removing this line helps the load script to execute faster.
- ♦ On NetWare 5.1, if you are using SLP as a short name resolution method, as described in “[Configuring Short Name Resolution](#)” on page 44, add the `cvsbinding` add command for the Internet Agent volume to the load script.

```
cvsbinding add cluster_volume_SERVER IP_address
```

- ♦ If you selected `vol:\system` on Internet Agent volume as the installation location ([Internet Agent Clustering Worksheet items 4 and 6](#)), add a `search` add command to add the new `vol:\system` directory to the server search path.

```
search add volume:\system
```

- ♦ If you selected `sys:\system` on each node as the installation location ([Internet Agent Clustering Worksheet items 4 and 6](#)) but you are storing the MTA startup file and the Internet Agent configuration file on the Internet Agent volume, add that location to the server search path.
- ♦ If you marked No under Load Internet Agent and Its MTA in Protected Memory? ([Internet Agent Clustering Worksheet item 8](#)), add the following `abend` recovery options:

```
set auto restart after abend = 2
set auto restart after abend delay time = 0
set auto restart down timeout = 60
set developer option = off
```

These settings provide the best possible handling of GroupWise databases in the event that an `abend` should occur within the cluster.

- ♦ Transfer the MTA `load` command from the `grpwise.ncf` file located in the `vol:\system` directory into the load script. Use Ctrl+C and Ctrl+V to copy and paste text into the load script page. Then delete or rename the `grpwise.ncf` file to avoid future confusion.

```
load volume:\system\gwmnta.nlm @domain.mta
```

- ♦ Add a delay so that the MTA is fully loaded before the Internet Agent starts to load:

```
load delay.nlm
delay 10
```

The length of the delay varies from system to system; ten seconds is a good starting place.

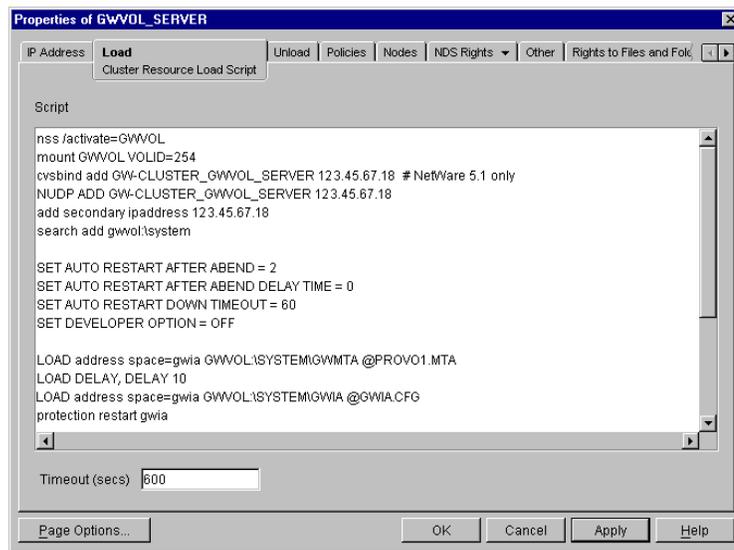
- ◆ Transfer the Internet Agent `load` command from the `gwia.ncf` file located in the `vol:\system` directory into the load script. Use `Ctrl+C` and `Ctrl+V` to copy and paste text into the load script page. Then delete or rename the `gwia.ncf` file to avoid future confusion.

```
load volume:\system\gwia.nlm @gwia.cfg
```

- ◆ If you marked Yes under Load Internet Agent and Its MTA in Protected Memory? ([Internet Agent Clustering Worksheet item 8](#)), add the `address space` parameter to the `load` commands to specify the protected address space where the Internet Agent and its MTA will run. Add a `protection restart` command for the address space name.

```
load address space=addr_space_name
                        volume:\system\gwmta.nlm @domain.mta
load address space=addr_space_name
                        volume:\system\gwia.nlm @gwia.cfg
protection restart addr_space_name
```

The result would look similar to the following example:



NOTE: The `set` commands are needed in the load script only when the MTA and the Internet Agent are *not* running in protected memory. The address space parameters are needed in the `load` commands only when the MTA and the Internet Agent *are* running in protected memory.

- 4 Click *Apply* to save the load script.
- 5 If necessary, click *OK* to confirm that you must offline and then online the volume resource in order for the changes to take effect.
- 6 Continue with [Modifying the Volume Resource Unload Script for the Internet Agent](#).

Modifying the Volume Resource Unload Script for the Internet Agent

The volume resource unload script executes whenever the Internet Agent volume 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.

To set up the unload script:

- 1 In ConsoleOne, in the properties pages for the Volume Resource object (*volume_SERVER*), click *Unload* to display the default volume resource unload script.

The next step assumes that this is the first time you have edited this unload script. If other GroupWise agents are already running from this volume, some of the modifications have already been made.

- 2 Make the following changes to the default unload script:

- ♦ If you marked Yes under Load Internet Agent and Its MTA in Protected Memory? (**Internet Agent Clustering Worksheet item 8**), add an unload address space command and an unload kill address space command to ensure that the address space is completely cleaned up.

```
unload address space=addr_space_name
unload kill address space=addr_space_name
```

If your system seems to be trying to kill the address space before the Internet Agent and its MTA have been completely unloaded, resulting in the agents hanging in the unloading state, set a delay of several seconds before issuing the unload kill address space command to allow the Internet Agent and its MTA adequate time to unload completely. The length of the delay varies from system to system; ten seconds is a good starting place.

```
unload address space=addr_space_name
delay 10
unload kill address space=addr_space_name
```

- ♦ If you marked No under Load Internet Agent and Its MTA in Protected Memory? (**Internet Agent Clustering Worksheet items 8**), create an unload command parallel to each load command that you placed in the load script.

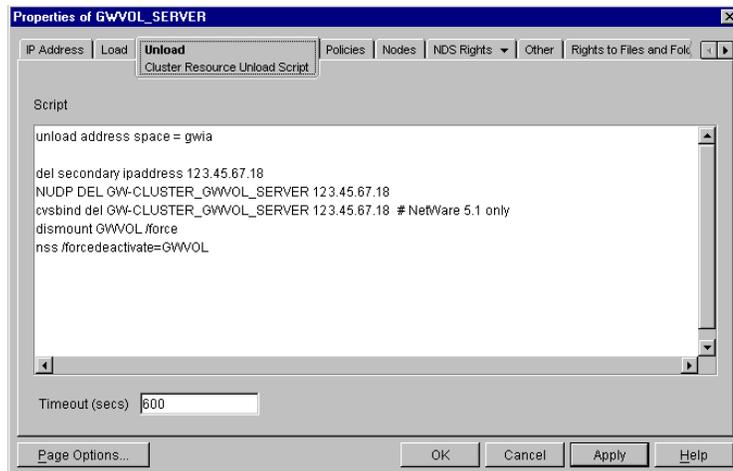
```
unload gwia.nlm
unload gwmta.nlm
```

- ♦ On NetWare 5.1, if you are using SLP as a short name resolution method, add the cvsbind del command for the Internet Agent volume to the unload script.

```
cvsbind del cluster_volume_SERVER IP_address
```

- ♦ Remove the trustmig command just like you did in the load script.

The result would look similar to the following example:

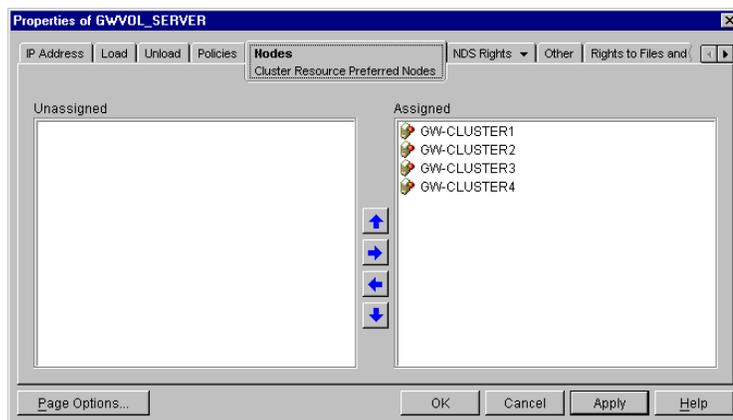


- 3 Click *Apply* to save the unload script.
- 4 If necessary, click *OK* to confirm that you must offline and then online the volume resource in order for the changes to take effect.
- 5 Continue with **Setting the Failover Path and Policies for the Internet Agent**.

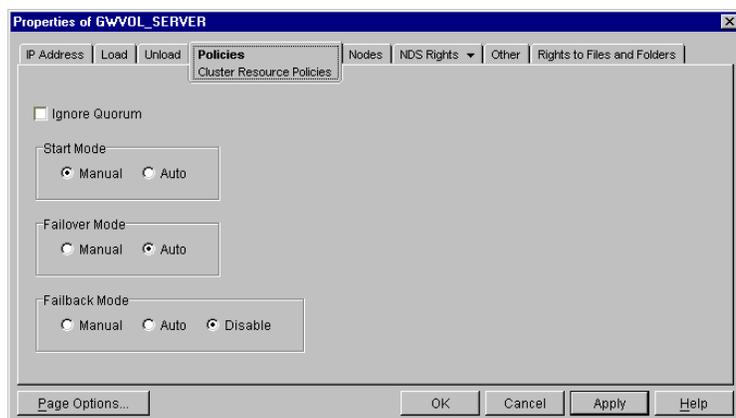
Setting the Failover Path and Policies for the Internet Agent

To modify the failover path and policies for the Internet Agent volume resource:

- 1 In ConsoleOne, in the properties pages for the Volume Resource object (*volume_SERVER*), click *Nodes* to display the default failover path for the Internet Agent volume resource.



- 2 Arrange the nodes in the cluster into the desired failover path for the Internet Agent volume (**Internet Agent Clustering Worksheet item 3**).
- 3 Click *Apply* to save the failover path.
- 4 Click *Policies* to display the default start, failover, and failback policies.



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

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

If you are considering changing these defaults, see the applicable section about failover and failback modes in the cluster documentation for your version of NetWare, as listed in [Chapter 1, “Introduction to GroupWise 7 and Novell Cluster Services on NetWare,” on page 21.](#)

- 5 Click *OK* when you are finished editing the Internet Agent volume resource properties.
- 6 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 Internet Agent is no different from setting it up for an Internet Agent in a any other environment. Follow the instructions in [“Enabling Internet Addressing”](#) in [“System”](#) in the *GroupWise 7 Administration Guide*, then return to this point.

Verifying Internet Agent Object Properties

During installation of the Internet Agent, the Internet Agent 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 Internet Agent.

- ◆ [“Accessing Internet Agent Object Properties” on page 80](#)
- ◆ [“Verifying the Reference to the Volume Resource” on page 81](#)
- ◆ [“Verifying the Reference to the Virtual Server” on page 81](#)
- ◆ [“Verifying Post Office Links” on page 81](#)
- ◆ [“Forcing Use of the Internet Agent Secondary IP Address” on page 81](#)

Accessing Internet Agent Object Properties

- 1 In ConsoleOne, browse to and select the Internet Agent domain in order to display its contents.

- 2 Right-click the Internet Agent object, then click *Properties*.
- 3 Continue with **Verifying the Reference to the Volume Resource**.

Verifying the Reference to the Volume Resource

In the Internet Agent object properties pages:

- 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 Volume Resource object name (*volume_SERVER*) of the Internet Agent volume, not the name of a physical server in the cluster.
- 3 Make changes if necessary.
- 4 Continue with **Verifying the Reference to the Virtual Server**.

Verifying the Reference to the Virtual Server

In the Internet Agent object properties pages:

- 1 Click *Server Directories*.
- 2 Verify that the displayed directories match the virtual server name (*cluster_volume_SERVER*) associated with the Volume Resource object, not the name of a physical server in the cluster.
- 3 Make changes if necessary.
- 4 Continue with **Verifying Post Office Links**.

Verifying Post Office Links

In the Internet Agent object properties pages:

- 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 Internet Agent.
- 3 Verify that the Links column displays the secondary IP addresses of the GroupWise volumes where post offices reside, not the IP addresses of any physical servers in the cluster.
- 4 Make changes if necessary.
- 5 Continue with **Forcing Use of the Internet Agent Secondary IP Address**.

Forcing Use of the Internet Agent Secondary IP Address

If you want the Internet Agent to send outgoing messages on its secondary IP address, rather than using the default of its primary IP address:

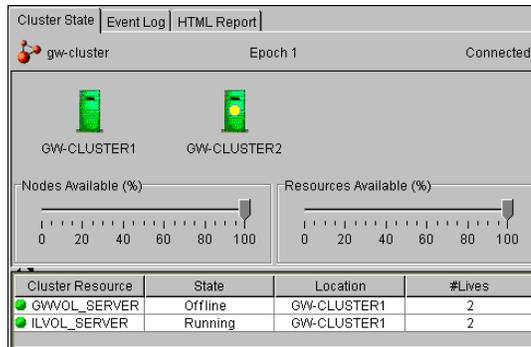
- 1 Click *GroupWise > Network Address*.
- 2 In the *TCP/IP Address* field, provide the secondary IP address (under **Internet Agent Clustering Worksheet item 1**) for the Internet Agent to use for sending outgoing messages.
- 3 Click *SMTP/MIME*, then click *Settings*.
- 4 Select *Bind to TCP/IP Address at Connection Time*.

- 5 Click *OK*.
- 6 Continue with **Testing the Clustered Internet Agent**.

4.2.5 Testing the Clustered Internet Agent

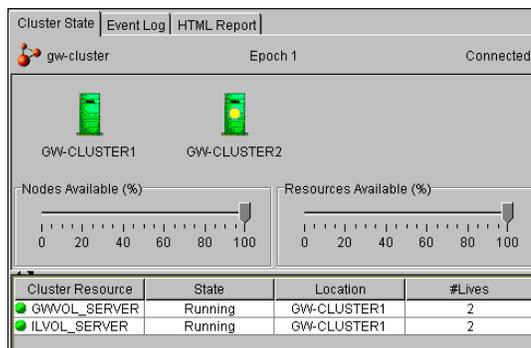
After you have configured the Internet Agent volume resource, you can test the load and unload scripts by bringing the Internet Agent volume online and taking it offline again.

- 1 In ConsoleOne, select the Cluster object, then click *View > Cluster State*.



The new Internet Agent volume resource shows *Offline* in the *State* column.

- 2 Click the new Internet Agent volume resource, then click *Online*.



The *State* column for the volume resource now displays *Running*.

- 3 Observe the server console where the Internet Agent and its MTA are loading to see that they start and run correctly.

If you are using protected memory, you can use the `protection` command at the server console prompt to list all the address spaces on the node and what NLM programs are running in each one.

- 4 Click the new Internet Agent volume resource, then click *Offline*.

The *State* column for the volume resource returns to *Offline*.

- 5 Observe the server console where the Internet Agent and its MTA are unloading to see that they shut down correctly.

If you are using protected memory, you can use the `protection` command again to make sure that the address space used by the Internet Agent and its MTA is no longer present.

- 6 Repeat [Step 2](#) whenever you are ready to bring the new Internet Agent volume resource online permanently.
On NetWare 6.x, these actions can also be performed from your Web browser. See [“Using Novell Remote Manager on NetWare 6.x”](#) on page 61.
- 7 Continue with [Managing the Internet Agent in a Cluster](#).

4.3 Managing the Internet Agent in a Cluster

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

- ♦ [Section 4.3.1, “Updating GroupWise Objects with Cluster-Specific Descriptions,”](#) on page 83
- ♦ [Section 4.3.2, “Knowing What to Expect in an Internet Agent Failover Situation,”](#) on page 84

4.3.1 Updating GroupWise Objects with Cluster-Specific Descriptions

After installing the Internet Agent 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 Internet Agent Domain and Its MTA”](#) on page 83
- ♦ [“Recording Cluster-Specific Information about the Internet Agent”](#) on page 84

Recording Cluster-Specific Information about the Internet Agent Domain and Its MTA

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

- 1 In ConsoleOne, browse to and right-click the Domain object, then click *Properties*.
- 2 In the *Description* field of the Internet Agent domain Identification page, provide a cluster-specific description of the Internet Agent domain, including the secondary IP address of its cluster-enabled volume and the cluster-unique port numbers used by its MTA.
- 3 Click *OK* to save the Internet Agent domain description.
- 4 Select the Internet Agent 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 cluster-enabled Internet Agent domain volume 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 Internet Agent](#).

Recording Cluster-Specific Information about the Internet Agent

With the contents of the Internet Agent domain still displayed:

- 1 Right-click the Internet Agent object, then click *Properties*.
- 2 Click *GroupWise*, then click *Identification*.
- 3 In the *Description* field, record the secondary IP address of the cluster-enabled Internet Agent domain volume and the cluster-unique port numbers used by the Internet Agent.
This information appears on the Internet Agent console, no matter which node in the cluster it is currently running on.
- 4 Click *OK* to save the Internet Agent information.
- 5 Continue with [Knowing What to Expect in an Internet Agent Failover Situation](#).

4.3.2 Knowing What to Expect in an Internet Agent Failover Situation

The failover behavior of the MTA for the Internet Agent 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 64.

Failover of the Internet Agent 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 Internet Agent restarts quickly in its failover location so users can reconnect quickly.

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

4.4 Internet Agent Clustering Worksheet

Item	Explanation
1) Shared Volume for Internet Agent:	Specify the name (<i>cluster_volume</i>) of the shared volume where the Internet Agent domain will be created.
Cluster Enabled?	For cluster-enabling, specify the IP addresses of the virtual server (<i>volume_SERVER.cluster</i>) to which the cluster-enabled volume is tied.
♦ Yes (highly recommended)	
Cluster volume IP address	For more information, see “Deciding Whether to Cluster-Enable the Internet Agent Volume” on page 70.
♦ No	
2) Internet Agent Domain Name:	Specify a unique name for the Internet Agent domain. Specify the directory on the Internet Agent volume where you want to create the new domain.
Domain Database Location:	For more information, see “Planning a Domain for the Internet Agent” on page 70.

Item	Explanation
3) Internet Agent Failover Path:	<p>List other nodes in the cluster where the Internet Agent and its MTA could fail over.</p> <p>For more information, see “Determining an Appropriate Failover Path for the Internet Agent Volume” on page 70.</p>
4) MTA Installation Location:	<p>Mark the location where you will install the MTA software. If necessary, specify the location where you will consolidate multiple MTA startup files on an Internet Agent volume.</p> <p>For more information, see “Deciding Where to Install the Internet Agent and Its MTA” on page 72.</p>
<ul style="list-style-type: none"> ◆ <code>vol:\system</code> on Internet Agent volume ◆ <code>sys:\system</code> on each node <p>Consolidate multiple MTA startup files on Internet Agent volume?</p>	
5) MTA Network Information:	<p>Gather the MTA network address information from the Internet Agent section of the “IP Address Worksheet” on page 40.</p> <p>For more information, see “Planning a Secondary IP Address and Cluster-Unique Port Numbers for the Internet Agent and Its MTA” on page 71.</p>
<ul style="list-style-type: none"> ◆ MTA IP address ◆ MTA message transfer port ◆ MTA live remote port ◆ MTA HTTP port 	
6) Internet Agent Installation Location:	<p>Mark the location where you will install the Internet Agent software.</p> <p>If necessary, specify the location where you will consolidate multiple Internet Agent configuration files (<code>gwia.cfg</code>) on an Internet Agent volume.</p> <p>For more information, see “Deciding Where to Install the Internet Agent and Its MTA” on page 72.</p>
<ul style="list-style-type: none"> ◆ <code>vol:\system</code> on the Internet Agent volume ◆ <code>sys:\system</code> on each node <p>Consolidate multiple Internet Agent configuration files on Internet Agent volume?</p>	
7) Internet Agent Network Information:	<p>Gather the Internet Agent network address information from the Internet Agent section of the “IP Address Worksheet” on page 40.</p> <p>For more information, see “Planning a Secondary IP Address and Cluster-Unique Port Numbers for the Internet Agent and Its MTA” on page 71.</p>
<ul style="list-style-type: none"> ◆ Internet Agent IP address ◆ Internet Agent HTTP port 	
8) Load Internet Agent and Its MTA in Protected Memory?	<p>Mark whether you need to run the Internet Agent and its MTA in protected memory. If so, specify a unique address space for each agent.</p> <p>For more information, see “Deciding Whether to Run the Internet Agent and Its MTA in Protected Memory” on page 72.</p>
<ul style="list-style-type: none"> ◆ No ◆ Yes 	
Protected address space names:	
<ul style="list-style-type: none"> ◆ MTA: ◆ Internet Agent: 	

4.5 Internet Agent Quick Checklist

- ❑ Plan the new clustered Internet Agent, including the new domain required to house the Internet Agent in a clustering environment.
See [Section 4.1, “Planning the Internet Agent in a Cluster,” on page 69.](#)
- ❑ Make sure your firewall is configured to accommodate the Internet Agent.
See [Section 4.1.5, “Preparing Your Firewall for the Internet Agent,” on page 71.](#)
- ❑ Cluster-enable the volume where the Internet Agent domain will reside.
See [Section 5.3.1, “Cluster-Enabling a Shared Volume for Use with the WebAccess Agent,” on page 93.](#)
- ❑ Create the new Internet Agent domain.
See [Section 4.2.2, “Creating a Domain for the Internet Agent,” on page 74.](#)
- ❑ Set up the MTA for the new Internet Agent domain.
See [Section 4.2.3, “Installing the MTA for the Internet Agent Domain,” on page 74.](#)
- ❑ Install the Internet Agent.
See [“Installing the Internet Agent Software in a Cluster” on page 74.](#)
- ❑ Modify the Internet Agent volume resource load script:
 - ◆ Remove the `trustmig` command
 - ◆ Add the `cvshint add` command (NetWare 5.1 only; optional)
 - ◆ Add the `search add` command (optional)
 - ◆ If you will *not* run the MTA and the Internet Agent in protected memory, add the `set auto restart` commands and the `set developer option = off` command
 - ◆ Add the `load` commands for the MTA and the Internet Agent, separating them with a `delay` command
 - ◆ If you *will* run the MTA and the Internet Agent in protected memory, add the address space parameter to the `load` commands and add a `protection restart` command for the address spaceSee [“Modifying the Volume Resource Load Script for the Internet Agent” on page 75.](#)
- ❑ Modify the Internet Agent volume resource unload script:
 - ◆ Add the MTA and Internet Agent or address space `unload` command(s)
 - ◆ Add the `cvshint del` command if you used the `cvshint add` command in the load script (NetWare 5.1 only; optional)
 - ◆ Remove the `trustmig` commandSee [“Modifying the Volume Resource Unload Script for the Internet Agent” on page 77.](#)
- ❑ Set up the Internet Agent volume failover path and policies.
See [“Setting the Failover Path and Policies for the Internet Agent” on page 79.](#)
- ❑ Enable Internet Addressing for the clustered Internet Agent.
See [“Enabling Internet Addressing for Your Clustered GroupWise System” on page 80.](#)
- ❑ Double-check the cluster-specific Internet Agent object properties.
See [“Verifying Internet Agent Object Properties” on page 80.](#)

- ❑ Test the clustered Internet Agent.
See [Section 4.2.5, “Testing the Clustered Internet Agent,”](#) on page 82.
- ❑ Record cluster-specific information in the properties pages of the GroupWise objects associated with the Internet Agent.
See [Section 4.3.1, “Updating GroupWise Objects with Cluster-Specific Descriptions,”](#) on page 83.

Implementing WebAccess in a NetWare Cluster

5

You should already have set up at least a basic GroupWise® system, as described in [Chapter 2, “Planning GroupWise in a NetWare Cluster,”](#) on page 23 and [Chapter 3, “Setting Up a Domain and Post Office in a NetWare Cluster,”](#) on page 43. As part of this process, the [“System Clustering Worksheet”](#) on page 37 and the [“IP Address Worksheet”](#) on page 40 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 WebAccess in a cluster.

- ♦ [Section 5.1, “Understanding the WebAccess Components,”](#) on page 89
- ♦ [Section 5.2, “Planning WebAccess in a Cluster,”](#) on page 89
- ♦ [Section 5.3, “Setting Up WebAccess in a Cluster,”](#) on page 93
- ♦ [Section 5.4, “Managing WebAccess in a Cluster,”](#) on page 101
- ♦ [Section 5.5, “WebAccess Clustering Worksheet,”](#) on page 103
- ♦ [Section 5.6, “WebAccess Quick Checklist,”](#) on page 104

5.1 Understanding the WebAccess Components

If you are not familiar with GroupWise WebAccess, review [“GroupWise WebAccess Overview”](#) in [“Installing GroupWise WebAccess”](#) in the *GroupWise 7 Installation Guide*.

As you plan WebAccess in a clustering environment, you must keep in mind that WebAccess consists of two components:

- ♦ WebAccess Agent (`gwinter.nlm`) that will be associated with a GroupWise WebAccess domain in the cluster
- ♦ WebAccess Application (a Java* servlet) that will be added to your Web server. You can install the WebAccess Application on a clustered Web server or a non-clustered Web server. How to set up and manage a clustered Web server is beyond the scope of the GroupWise documentation. If you have not clustered your Web server, you can install the WebAccess Application on a server that is outside of the cluster where the WebAccess Agent is installed.

5.2 Planning WebAccess in a 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 each GroupWise gateway, including the WebAccess Agent.

[Section 5.5, “WebAccess Clustering Worksheet,”](#) on page 103 lists all the information you need as you set up the WebAccess Agent in a clustering environment. You should print the worksheet and fill it out as you complete the tasks listed below:

- ♦ [Section 5.2.1, “Planning a New Domain for the WebAccess Agent,”](#) on page 90

- ◆ Section 5.2.2, “Deciding Whether to Cluster-Enable the WebAccess Agent Volume,” on page 90
- ◆ Section 5.2.3, “Determining an Appropriate Failover Path for the WebAccess Agent Volume,” on page 91
- ◆ Section 5.2.4, “Planning a Secondary IP Address and Cluster-Unique Port Numbers for the WebAccess Agent and Its MTA,” on page 91
- ◆ Section 5.2.5, “Deciding Where to Install the WebAccess Agent and Its MTA,” on page 91
- ◆ Section 5.2.6, “Deciding Whether to Run the WebAccess Agent and Its MTA in Protected Memory,” on page 92
- ◆ Section 5.2.7, “Planning the MTA Installation,” on page 92
- ◆ Section 5.2.8, “Planning the WebAccess Installation,” on page 92

5.2.1 Planning a New Domain for the WebAccess Agent

The considerations involved in planning a domain for the WebAccess Agent are much the same as planning any other domain. In preparation, review “[Planning a New Domain](#)”, then print and fill out the “[Domain Worksheet](#)” in “[Domains](#)” in the *GroupWise 7 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 volume where you want the domain directory to reside.
- ◆ 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 WebAccess Clustering Worksheet.

WEBACCESS CLUSTERING WORKSHEET

Under [Item 1: Shared Volume for WebAccess Agent](#), transfer the domain location from the Domain Worksheet to the WebAccess Clustering Worksheet.

Under [Item 2: WebAccess Agent Domain Name](#), transfer the domain name and database directory from the Domain Worksheet to the WebAccess Clustering Worksheet.

5.2.2 Deciding Whether to Cluster-Enable the WebAccess Agent Volume

You should plan to cluster-enable the shared volume where the WebAccess Agent domain will reside. For a review of the benefits of cluster-enabling volumes, see [Section 2.6, “Deciding Whether to Cluster-Enable the Shared Volumes Used by GroupWise,” on page 27](#), which describes the issues in the context of planning MTA and POA installations.

WEBACCESS CLUSTERING WORKSHEET

Under [Item 1: Shared Volume for WebAccess Agent](#), mark Yes under Cluster Enabled?.

Cluster-enabling relies on successful short name resolution throughout your system. Review [Section 2.7, “Ensuring Successful Name Resolution for GroupWise Volumes,” on page 29](#), which describes the issues in the context of planning MTA and POA installations.

5.2.3 Determining an Appropriate Failover Path for the WebAccess Agent Volume

As with the MTA and the POA, you need to decide which nodes in the cluster are appropriate locations where the WebAccess Agent volume could fail over. For a review of failover paths, see [“Determining Appropriate Failover Paths for the Agents” on page 33](#), which describes the issues in the context of planning MTA and POA installations.

WEBACCESS CLUSTERING WORKSHEET

Under [Item 4: WebAccess Agent Failover Path](#), list the nodes that you want to have in the WebAccess Agent volume failover path.

5.2.4 Planning a Secondary IP Address and Cluster-Unique Port Numbers for the WebAccess Agent and Its MTA

As with the MTA and the POA, the WebAccess Agent needs a secondary IP address and cluster-unique port numbers. As part of planning to install the MTA and POA, you should already have determined the secondary IP address and cluster-unique port numbers for the WebAccess Agent and its MTA as you filled out the [“IP Address Worksheet” on page 40](#). If you do not have a filled-out copy of this worksheet for your system, print it now and fill in current system information.

WEBACCESS CLUSTERING WORKSHEET

Under [Item 6: MTA Network Information](#), transfer the MTA secondary IP address and cluster-unique port numbers from the WebAccess section the IP Address Worksheet to the WebAccess Clustering Worksheet.

Under [Item 1: Shared Volume for WebAccess Agent](#), copy the MTA secondary IP address under Cluster Volume IP Address as well, because they are the same.

Under [Item 8: WebAccess Agent Network Information](#), transfer the WebAccess Agent secondary IP address (the same as for its MTA) and the cluster-unique WebAccess Agent port number from the WebAccess section of the IP Address Worksheet to the WebAccess Clustering Worksheet.

5.2.5 Deciding Where to Install the WebAccess Agent and Its MTA

As with the MTA and the POA, you can choose to install the WebAccess Agent and its MTA to the `sys:\system` directory of each node or to a `vol:\system` directory on the WebAccess Agent volume. For a discussion of these alternatives, see [“Deciding Where to Install the Agent Software” on page 34](#), which describes the issues in the context of planning MTA and POA installations. If you only have one WebAccess Agent for your GroupWise system with several nodes in its failover path, it is an easy choice.

WEBACCESS CLUSTERING WORKSHEET

Under [Item 5: MTA Installation Location](#) and [Item 7: WebAccess Agent Installation Location](#), mark whether you will install the WebAccess Agent and its MTA to `sys:\system` on each node in the cluster or to a `vol:\system` directory on the WebAccess Agent volume. Also specify where the MTA startup file will be stored.

5.2.6 Deciding Whether to Run the WebAccess Agent and Its MTA in Protected Memory

As with the MTA and the POA, you can choose whether to run the WebAccess Agent in protected memory. For a review of the benefits of protected memory, see [“Deciding Whether to Run the Agents in Protected Memory” on page 36](#), which describes the issues in the context of planning MTA and POA installations.

You might think that protected memory is not necessary if you have only one WebAccess Agent for your GroupWise system because it could never fail over to a node where another WebAccess Agent was running. However, because the WebAccess Agent in a cluster is installed into its own domain with its own MTA, this MTA could fail over to a node where another MTA was already running. Therefore, it is safest to load the WebAccess Agent and its MTA into protected memory. Load each agent into its own memory space and mark each memory space as restartable.

WEBACCESS CLUSTERING WORKSHEET

Under [Item 8: Load WebAccess Agent and Its MTA in Protected Memory?](#), mark whether you need to run the WebAccess Agent and its MTA in protected memory. If you do, provide a protected memory address space name for each agent.

5.2.7 Planning the MTA Installation

Follow the instructions in [“Planning the NetWare Agent Installation” on page 37](#), then return to this point. After you follow the instructions, you will have a filled-out NetWare Agent Worksheet to use when you install the MTA.

IMPORTANT: Do not install the NetWare MTA until you are instructed to do so in [Section 5.3, “Setting Up WebAccess in a Cluster,” on page 93](#).

5.2.8 Planning the WebAccess Installation

Aside from the cluster-specific issues discussed in the preceding sections, the considerations involved in planning to install WebAccess are the same in a clustering environment as for any other environment. Review [“Planning GroupWise WebAccess”](#), then print and fill out the [“GroupWise WebAccess Installation Worksheet”](#) in [“Installing GroupWise WebAccess”](#) in the *GroupWise 7 Installation Guide*. When you set up WebAccess in a cluster, you install the WebAccess Agent and the WebAccess Application in two separate steps:

- ◆ [“Planning the WebAccess Agent Installation” on page 93](#)
- ◆ [“Planning the WebAccess Application Installation” on page 93](#)

IMPORTANT: Do not install the WebAccess software until you are instructed to do so in [Section 5.3, “Setting Up WebAccess in a Cluster,”](#) on page 93.

Planning the WebAccess Agent Installation

For the WebAccess Agent, fill out items 2 through 12 on the GroupWise WebAccess Installation Worksheet, taking into account the following cluster-specific issues:

WEBACCESS INSTALLATION WORKSHEET

Under [Item 2: Installation Directory](#), take into account your decision recorded on the WebAccess Clustering Worksheet ([Item 7: WebAccess Agent Installation Location](#)).

Under [Item 3: Server Address](#), transfer the IP address and port number from the WebAccess Clustering Worksheet ([Item 8: WebAccess Agent Network Information](#)) filled out in [“Planning a Secondary IP Address and Cluster-Unique Port Numbers for the WebAccess Agent and Its MTA”](#) on page 91.

Under [Item 4: Enable Clustering Support?](#), mark Yes. This causes the WebAccess Installation program to customize the WebAccess files for clustering.

Under [Item 5: Domain Directory Path](#), transfer the domain directory from the Domain Worksheet you filled out in [“Planning a New Domain for the WebAccess Agent”](#) on page 90.

Planning the WebAccess Application Installation

The WebAccess Application can be installed on a clustered or non-clustered Web server. How to set up and manage a clustered Web server is beyond the scope of the GroupWise documentation. For WebAccess Application planning information, see [“Determining the WebAccess and WebPublisher Applications’ Configuration”](#) in [“Installing GroupWise WebAccess”](#) in the *GroupWise 7 Installation Guide*.

5.3 Setting Up WebAccess in a Cluster

You should already have reviewed [Section 5.2, “Planning WebAccess in a Cluster,”](#) on page 89 and filled out the [Section 5.5, “WebAccess Clustering Worksheet,”](#) on page 103. You are now ready to complete the following tasks to set up the WebAccess Agent in a clustering environment:

- ◆ [Section 5.3.1, “Cluster-Enabling a Shared Volume for Use with the WebAccess Agent,”](#) on page 93
- ◆ [Section 5.3.2, “Creating a Domain for the WebAccess Agent,”](#) on page 94
- ◆ [Section 5.3.3, “Installing the MTA for the WebAccess Agent Domain,”](#) on page 94
- ◆ [Section 5.3.4, “Installing and Configuring the WebAccess Agent in a Cluster,”](#) on page 94
- ◆ [Section 5.3.5, “Testing Your Clustered WebAccess Installation,”](#) on page 100

5.3.1 Cluster-Enabling a Shared Volume for Use with the WebAccess Agent

- 1 Complete the cluster-enabling steps in the applicable section of cluster documentation for your version of NetWare, as listed in [Chapter 1, “Introduction to GroupWise 7 and Novell Cluster Services on NetWare,”](#) on page 21.

The WebAccess Clustering Worksheet provides the volume to cluster-enable, the cluster-enabled volume IP address, and the failover path for the WebAccess volume.

For a review of the new Novell eDirectory™ objects that are created when you cluster-enable a shared volume, see [Section 2.6, “Deciding Whether to Cluster-Enable the Shared Volumes Used by GroupWise,” on page 27.](#)

If you have installed the latest version of ConsoleOne® and the Novell Cluster Services snap-in, you can rename the cluster-related objects in case your DNS name server cannot resolve object names that include the underscore (_) character.

- 2 To ensure successful short name resolution, add entries for the WebAccess Agent virtual server to support your preferred methods of short name resolution, as described in [“Configuring Short Name Resolution” on page 44.](#)
- 3 Continue with [Creating a Domain for the WebAccess Agent.](#)

5.3.2 Creating a Domain for the WebAccess Agent

The WebAccess Agent domain will be a secondary domain. To create it, follow the instructions in [Section 3.3, “Creating a New Secondary Domain in a Cluster,” on page 48,](#) taking your information from the WebAccess Clustering Worksheet, rather than the System Clustering Worksheet, then return to this point.

Do not create any post offices in the WebAccess Agent domain.

Continue with [Installing the MTA for the WebAccess Agent Domain.](#)

5.3.3 Installing the MTA for the WebAccess Agent Domain

The MTA for the WebAccess Agent domain can be installed just like any other MTA in your clustered GroupWise system. Follow the instructions in [“Installing the Agent Software in a Cluster” on page 51,](#) then return to this point.

You do not need to edit the MTA startup file. You do not need to modify the Volume Resource properties until after you have installed the WebAccess Agent.

Continue with [Installing and Configuring the WebAccess Agent in a Cluster.](#)

5.3.4 Installing and Configuring the WebAccess Agent in a Cluster

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

- ♦ [“Installing the WebAccess Agent Software in a Cluster” on page 95](#)
- ♦ [“Configuring the WebAccess Agent Volume Resource to Load and Unload the WebAccess Agent and Its MTA” on page 96](#)

Installing the WebAccess Agent Software in a Cluster

The WebAccess Agent is the component of your WebAccess installation that accesses post offices and libraries to retrieve information for WebAccess client users.

- 1 Map a drive to the WebAccess Agent volume ([WebAccess Clustering Worksheet item 1](#)) where the WebAccess domain is located.

The WebAccess Agent volume name will be *cluster_volume*. For assistance with mapping a drive to a cluster-enabled volume, see [“Configuring Short Name Resolution” on page 44](#).

- 2 If you selected *vol:\system* on WebAccess Agent volume as the WebAccess Agent installation location ([WebAccess Clustering Worksheet item 7](#)), create the *vol:\system* directory on the WebAccess Agent volume accessed in [Step 1](#).

or

if you selected *sys:\system* on each node, decide which node you will install the WebAccess agent to first, then map a drive to its *sys:\system* directory.

- 3 Start the WebAccess Installation program and install the NetWare WebAccess Agent, following Step 1 through Step 5 provided in [“Installing the WebAccess Agent”](#) in [“Installing GroupWise WebAccess”](#) in the [GroupWise 7 Installation Guide](#). Keep in mind the following cluster-specific details:

- ♦ In the Components dialog box, select only *GroupWise WebAccess Agent*.
Do not install the WebAccess Application at this time.
- ♦ Use items 2 through 12 on the GroupWise WebAccess Installation Worksheet that you filled out in [“Planning the WebAccess Installation” on page 92](#) to fill in the fields during the WebAccess Agent installation process.
- ♦ In the Network Address dialog box, select *Configure GroupWise Agents for Clustering*.
- ♦ In the Installation Path dialog box, be sure to browse through the drive you mapped to the location you chose in [Step 2](#) above.
- ♦ In the Gateway Directory dialog box, be sure to browse to the domain directory through the drive you mapped in [Step 1](#) above.
- ♦ In the Start Applications dialog box, deselect *Start the GroupWise WebAccess Agent*.
- ♦ The WebAccess Installation program creates the *strtweb.ncf* and *stopweb.ncf* files, which include the *load* and *unload* commands for the WebAccess Agent. You use this information later when you create the load and unload scripts for the WebAccess Volume Resource object.

- 4 If you need to install the WebAccess Agent to *sys:\system* on multiple nodes in the cluster, repeat [Step 4](#), mapping new drives as needed.
- 5 Make sure you have completed all the WebAccess Agent tasks described in [“NetWare and Windows: Setting Up GroupWise WebAccess”](#) in [“Installing GroupWise WebAccess”](#) in the [GroupWise 7 Installation Guide](#), but do not start the WebAccess Agent at this time.
- 6 Continue with [Configuring the WebAccess Agent Volume Resource to Load and Unload the WebAccess Agent and Its MTA](#).

Configuring the WebAccess Agent Volume Resource to Load and Unload the WebAccess Agent and Its MTA

The properties of the Volume Resource object define how the WebAccess Agent volume functions within the cluster, how NLM™ programs are loaded and unloaded, and how failover and fallback situations are handled. Complete the following tasks for the WebAccess Agent volume:

- ♦ “Modifying the Volume Resource Load Script for the WebAccess Agent” on page 96
- ♦ “Modifying the Volume Resource Unload Script for the WebAccess Agent” on page 98
- ♦ “Setting the Failover Path and Policies for the WebAccess Agent” on page 99
- ♦ “Installing and Configuring the WebAccess Application in a Cluster” on page 100

Modifying the Volume Resource Load Script for the WebAccess Agent

The volume resource load script executes whenever the WebAccess Agent volume comes online.

To set up the load script:

- 1 In ConsoleOne®, browse to and select the Cluster object.
If necessary, click *View > Console View* to display its contents.
- 2 Right-click the Volume Resource object (*volume_SERVER*), then click *Properties > Load* to display the default volume resource load script for the WebAccess Agent volume.
The next step assumes that this is the first time you have edited the load script. If other GroupWise agents are already running from this volume, some of the modifications have already been made.
- 3 Make the following changes to the default load script:
 - ♦ Remove the `trustmig` command. It is not necessary to migrate trustees for the WebAccess Agent volume. Removing this line helps the load script to execute faster.
 - ♦ On NetWare 5.1, if you are using SLP as a short name resolution method, as described in “Configuring Short Name Resolution” on page 44, add the `cvsbind add` command for the WebAccess Agent volume to the load script.

```
cvsbind add cluster_volume_SERVER IP_address
```

- ♦ If you selected `vol:\system` on WebAccess Agent volume as the installation location (WebAccess Clustering Worksheet items 5 and 7), add a `search add` command to add the new `vol:\system` directory to the server search path.

```
search add volume:\system
```

- ♦ If you selected `sys:\system` on each node as the installation location (WebAccess Clustering Worksheet items 5 and 7) but you are storing the MTA startup file on the WebAccess Agent volume, add that location to the server search path.
- ♦ If you marked No under Load WebAccess Agent and Its MTA in Protected Memory? (WebAccess Clustering Worksheet item 9), add the following `abend` recovery options:

```
set auto restart after abend = 2
set auto restart after abend delay time = 0
set auto restart down timeout = 60
set developer option = off
```

These settings provide the best possible handling of GroupWise databases in the event that an `abend` should occur within the cluster.

- ◆ Transfer the MTA load command from the `grpwise.ncf` file located in the `vol:\system` directory into the load script. Use Ctrl+C and Ctrl+V to copy and paste text into the load script page. Then delete or rename the `grpwise.ncf` file to avoid future confusion.

```
load volume:\system\gwmta.nlm @domain.mta
```

- ◆ Add a delay so that the MTA is fully loaded before the WebAccess Agent starts to load:

```
load delay
delay 10
```

The length of the delay varies from system to system; ten seconds is a good starting place.

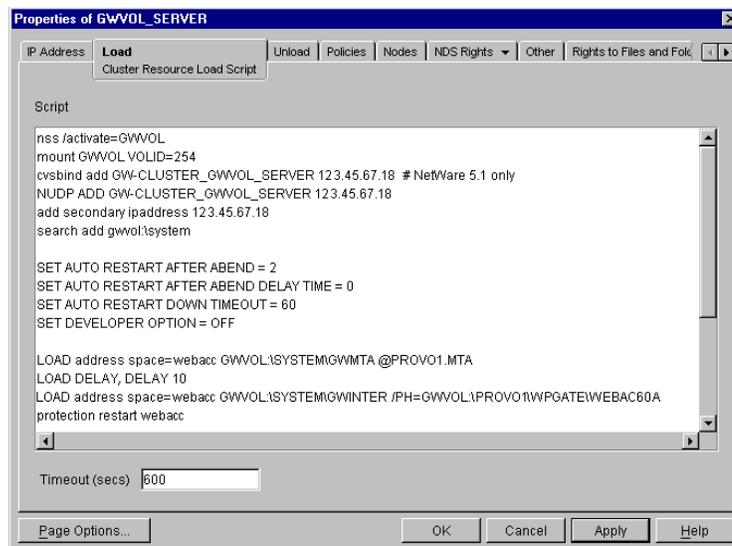
- ◆ Transfer the WebAccess Agent load command from the `strtweb.ncf` file located in the `vol:\system` directory into the load script. Use Ctrl+C and Ctrl+V to copy and paste text into the load script page.

```
load volume:\system\gwinter.nlm
  /ph=volume:\domain\wpgate\webac70a
  /user=username /PASSWORD=password
```

- ◆ If you marked Yes under Load WebAccess Agent and Its MTA in Protected Memory? ([WebAccess Clustering Worksheet item 9](#)), add the `address space` parameter to the load commands to specify the protected address space where the WebAccess Agent and its MTA will run. Add a `protection restart` command for the address space name.

```
load address space=addr_space_name
  volume:\system\gwmta.nlm @domain.mta
load address space=addr_space_name
  volume:\system\gwinter.nlm
  /ph=volume:\domain\wpgate\webac70a
  /user=username /password=password
protection restart addr_space_name
```

The result would look similar to the following example:



NOTE: The `set` commands are needed only when the MTA and the WebAccess Agent are not running in protected memory. The address space parameters are needed in the `load` commands only when the MTA and the WebAccess Agent are running in protected memory.

- 4 Click *Apply* to save the load script.
- 5 If necessary, click *OK* to confirm that you must offline and then online the volume resource in order for the changes to take effect.
- 6 Continue with **Modifying the Volume Resource Unload Script for the WebAccess Agent**.

Modifying the Volume Resource Unload Script for the WebAccess Agent

The volume resource unload script executes whenever the WebAccess Agent volume 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.

To set up the unload script:

- 1 In ConsoleOne, in the properties pages for the Volume Resource object (*volume_SERVER*), click *Unload* to display the default volume resource unload script.

The next step assumes that this is the first time you have edited the unload script. If other GroupWise agents are already running from this volume, some of the modifications have already been made.

- 2 Make the following changes to the default unload script:

- ♦ If you marked Yes under Load WebAccess Agent and Its MTA in Protected Memory? (**WebAccess Clustering Worksheet item 9**), add an `unload address space` command and an `unload kill address space` command to ensure that the address space is completely cleaned up.

```
unload address space=addr_space_name
unload kill address space=addr_space_name
```

If your system seems to be trying to kill the address space before the WebAccess Agent and its MTA have been completely unloaded, resulting in the agents hanging in the unloading state, set a delay of several seconds before issuing the `unload kill address space` command to allow the WebAccess Agent and its MTA adequate time to unload completely. The length of the delay varies from system to system; ten seconds is a good starting place.

```
unload address space=addr_space_name
delay 10
unload kill address space=addr_space_name
```

- ♦ If you marked No under Load WebAccess Agent and Its MTA in Protected Memory? (**WebAccess Clustering Worksheet items 9**), create an `unload` command parallel to each `load` command that you placed in the load script.

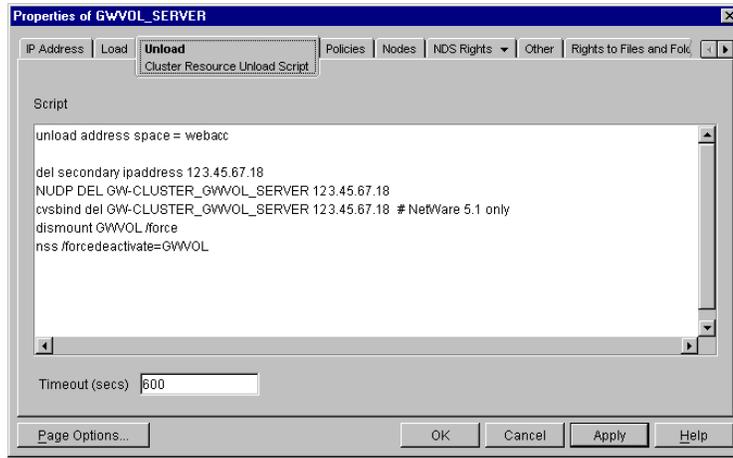
```
unload gwinter.nlm
unload gwmta.nlm
```

- ♦ On NetWare 5.1, if you are using SLP as a short name resolution method, add the `cvsbinding del` command for the WebAccess Agent volume to the unload script.

```
cvsbind del cluster_volume_SERVER ip_address
```

- ◆ Remove the `trustmig` command just like you did in the load script.

The result would look similar to the following example:

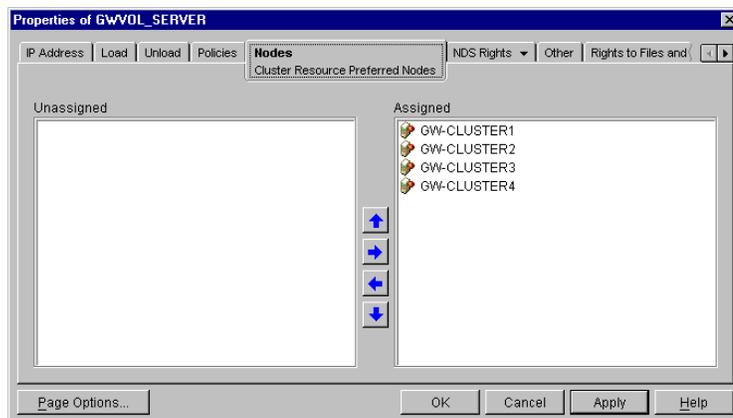


- 3 Click *Apply* to save the unload script.
- 4 If necessary, click *OK* to confirm that you must offline and then online the volume resource in order for the changes to take effect.
- 5 Continue with [Setting the Failover Path and Policies for the WebAccess Agent](#).

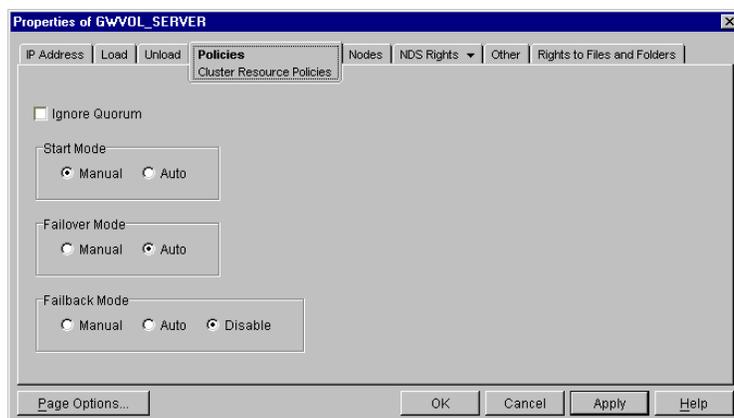
Setting the Failover Path and Policies for the WebAccess Agent

To modify the failover path and policies for the WebAccess Agent volume resource:

- 1 In ConsoleOne, in the properties pages for the Volume Resource object (*volume_SERVER*), click *Nodes* to display the default failover path for the WebAccess Agent volume resource.



- 2 Arrange the nodes in the cluster into the desired failover path for the WebAccess Agent volume ([WebAccess Clustering Worksheet item 4](#)).
- 3 Click *Apply* to save the failover path.
- 4 Click *Policies* to display the default start, failover, and failback policies.



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

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

If you are considering changing these defaults, see the section about failover and failback mode modes in the clustering documentation for your version of NetWare, as listed in [Chapter 1, “Introduction to GroupWise 7 and Novell Cluster Services on NetWare,”](#) on page 21.

- 5 Click *OK* when you are finished editing the WebAccess Agent volume resource properties.
- 6 Continue with [“Installing and Configuring the WebAccess Agent in a Cluster”](#) on page 94.

Installing and Configuring the WebAccess Application in a Cluster

If you have clustered your Web server, you must install the WebAccess Application on each node where the Web server is installed, as described in [“Installing the WebAccess Application and WebPublisher Application”](#) in [“Installing GroupWise WebAccess”](#) in the *GroupWise 7 Installation Guide*.

5.3.5 Testing Your Clustered WebAccess Installation

Remember that the WebAccess Agent volume resource and the WebAccess Application on your Web server are separate and could fail over to different nodes at different times.

To thoroughly test your WebAccess installation:

- 1 Make sure the initial combination of WebAccess Agent volume resource and the WebAccess Application installed on your Web server is functioning properly.
- 2 Migrate the WebAccess Agent volume resource to each node on its failover path, making sure it functions with your Web server.
- 3 If your Web server is clustered, migrate the Web server cluster resource to a different node, migrate the WebAccess Agent volume resource to each node in the Web server cluster resource failover path, then make sure each combination works.
- 4 Repeat [Step 3](#) for each Web server cluster resource.

5.4 Managing WebAccess in a Cluster

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

- ♦ [Section 5.4.1, “Updating GroupWise Objects with Cluster-Specific Descriptions,” on page 101](#)
- ♦ [Section 5.4.2, “Knowing What to Expect in WebAccess Failover Situations,” on page 102](#)
- ♦ [Section 5.4.3, “Updating the WebAccess Agent Configuration File \(commgr.cfg\),” on page 102](#)

5.4.1 Updating GroupWise Objects with Cluster-Specific Descriptions

After installing WebAccess 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 WebAccess Agent Domain and Its MTA” on page 101](#)
- ♦ [“Recording Cluster-Specific Information about the WebAccess Agent” on page 102](#)

Recording Cluster-Specific Information about the WebAccess Agent Domain and Its MTA

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

- 1 In ConsoleOne, browse to and right-click the Domain object, then click *Properties*.
- 2 In the *Description* field of the WebAccess Agent domain Identification page, provide a cluster-specific description of the WebAccess Agent domain, including the secondary IP address of its cluster-enabled volume and the cluster-unique port numbers used by its MTA.

You might also want to include cluster-specific information about the WebAccess Application, such as the secondary IP address of the Web server cluster resource where the WebAccess Application is installed.

- 3 Click *OK* to save the WebAccess Agent domain description.
- 4 Select the WebAccess Agent 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 cluster-enabled WebAccess Agent domain volume 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 Internet Agent](#).

Recording Cluster-Specific Information about the WebAccess Agent

With the contents of the WebAccess Agent domain still displayed:

- 1 Right-click the WEBAC70A object, then click *Properties*.
- 2 Click *GroupWise*, then click *Identification*.
- 3 In the *Description* field, record the secondary IP address of the cluster-enabled WebAccess Agent domain volume and the cluster-unique port numbers used by the WebAccess Agent.
This information appears on the WebAccess Agent console, no matter which node in the cluster it is currently running on.
- 4 Click *OK* to save the WebAccess Agent information.
- 5 Continue with [Knowing What to Expect in MTA and POA Failover Situations](#).

5.4.2 Knowing What to Expect in WebAccess Failover Situations

The failover behavior of the MTA for the WebAccess Agent 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 64](#).

The WebAccess Application caches users’ credentials on the node where it is running. Therefore, if that node fails, or if the WebAccess Application migrates to a different node, the cached credentials are lost. Consequently, the user needs to restart the WebAccess client in order to re-authenticate and re-establish the credentials.

If the WebAccess Agent fails over or migrates, the user receives an error message that the WebAccess Agent is no longer available. However, after the WebAccess Agent starts in its new location, the WebAccess Application passes the cached user credentials to the WebAccess Agent and the user reconnects automatically without having to re-authenticate.

As with the MTA and the POA, migration of the WebAccess Agent takes longer than failover. However, the WebAccess Agent restarts quickly so that users are able to reconnect quickly.

5.4.3 Updating the WebAccess Agent Configuration File (commgr.cfg)

As part of installing WebAccess, the WebAccess Agent configuration file (`commgr.cfg`) is created in the following subdirectory:

```
domain\wpgate\webac70a
```

It is also automatically copied to the following Web server subdirectory:

```
sys:\novell\webaccess
```

If you change WebAccess agent configuration information (for example, if you change its ip address), the information is changed in the following file:

```
domain\wpgate\webac70a\commgr.cfg
```

because the domain is on a cluster-enabled volume, and it is changed in the following file:

```
sys:\novell\webaccess\commgr.cfg
```

on the node where the WebAccess Application is currently running. However, the other nodes on the WebAccess Application failover path are not currently available for update. therefore, you must manually copy the updated `commgr.cfg` file to the `sys:\novell\webaccess` subdirectory on each node in the WebAccess Application failover path.

5.5 WebAccess Clustering Worksheet

Item	Explanation
1) Shared Volume for WebAccess Agent:	Specify the name (<i>cluster_volume</i>) of the shared volume where the WebAccess Agent domain will be created.
Cluster Enabled?	For cluster-enabling, specify the IP addresses of the virtual server (<i>volume_cluster</i>) to which the cluster-enabled volume is tied.
<ul style="list-style-type: none"> ◆ Yes (highly recommended) 	For more information, see “Deciding Whether to Cluster-Enable the WebAccess Agent Volume” on page 90.
Cluster volume IP address	
<ul style="list-style-type: none"> ◆ No 	
2) WebAccess Agent Domain Name:	Specify a unique name for the WebAccess Agent domain. Specify the directory on the WebAccess Agent volume where you want to create the new domain.
Domain Database Location:	For more information, see “Planning a New Domain for the WebAccess Agent” on page 90.
3) WebAccess Agent Failover Path:	List other nodes in the cluster where the WebAccess Agent and its MTA could fail over.
	For more information, see “Determining an Appropriate Failover Path for the WebAccess Agent Volume” on page 91.
4) MTA Installation Location:	Mark the location where you will install the MTA software.
<ul style="list-style-type: none"> ◆ <code>vol:\system</code> on WebAccess Agent volume 	If necessary, specify the location where you will consolidate multiple MTA startup files on a WebAccess Agent volume.
<ul style="list-style-type: none"> ◆ <code>sys:\system</code> on each node 	For more information, see “Deciding Where to Install the WebAccess Agent and Its MTA” on page 91.
Consolidate multiple MTA startup files on WebAccess Agent volume?	
5) MTA Network Information:	Gather the MTA network address information from the WebAccess section of the “IP Address Worksheet” on page 40.
<ul style="list-style-type: none"> ◆ MTA IP address ◆ MTA message transfer port ◆ MTA live remote port ◆ MTA HTTP port 	For more information, see “Planning a Secondary IP Address and Cluster-Unique Port Numbers for the WebAccess Agent and Its MTA” on page 91.
6) WebAccess Agent Installation Location:	Mark the location where you will install the WebAccess Agent software.
<ul style="list-style-type: none"> ◆ <code>vol:\system</code> on the WebAccess Agent volume 	For more information, see “Deciding Where to Install the WebAccess Agent and Its MTA” on page 91.
<ul style="list-style-type: none"> ◆ <code>sys:\system</code> on each node 	

Item	Explanation
7) WebAccess Agent Network Information:	Gather the WebAccess Agent network address information from the WebAccess section of the “IP Address Worksheet” on page 40.
<ul style="list-style-type: none"> ◆ WebAccess Agent IP address ◆ WebAccess Agent HTTP port 	For more information, see “Planning a Secondary IP Address and Cluster-Unique Port Numbers for the WebAccess Agent and Its MTA” on page 91.
8) Load WebAccess Agent and Its MTA in Protected Memory?	Mark whether you need to run the WebAccess Agent and its MTA in protected memory. If so, specify a unique address space for each agent.
<ul style="list-style-type: none"> ◆ No ◆ Yes 	For more information, see “Deciding Whether to Run the WebAccess Agent and Its MTA in Protected Memory” on page 92.
Protected address space names:	
<ul style="list-style-type: none"> ◆ MTA: ◆ WebAccess: 	

5.6 WebAccess Quick Checklist

- Plan the new clustered WebAccess installation, including the new domain required to house the WebAccess Agent in a clustering environment.
See [Section 5.2, “Planning WebAccess in a Cluster,” on page 89.](#)
 - Cluster-enable the volume where the WebAccess Agent domain will reside.
See [Section 5.3.1, “Cluster-Enabling a Shared Volume for Use with the WebAccess Agent,” on page 93.](#)
 - Create the new WebAccess Agent domain.
See [Section 5.3.2, “Creating a Domain for the WebAccess Agent,” on page 94.](#)
 - Set up the MTA for the new WebAccess Agent domain.
See [Section 5.3.3, “Installing the MTA for the WebAccess Agent Domain,” on page 94.](#)
 - Install the WebAccess Agent.
See [Section 5.3.4, “Installing and Configuring the WebAccess Agent in a Cluster,” on page 94.](#)
 - Modify the WebAccess Agent volume resource load script:
 - ◆ Remove the `trustmig` command
 - ◆ Add the `cvswind add` command (NetWare 5.1 only; optional)
 - ◆ Add the `search add` command (optional)
 - ◆ If you will *not* run the MTA and the WebAccess Agent in protected memory, add the `set auto restart` commands and the `set developer option = off` command
 - ◆ Add the `load` commands for the MTA and the WebAccess Agent, separating them with a `delay` command
 - ◆ If you *will* run the MTA and the WebAccess Agent in protected memory, add the `address space` parameter to the `load` commands and add a `protection restart` command for the address space
- See [“Modifying the Volume Resource Load Script for the WebAccess Agent” on page 96.](#)

- ❑ Modify the WebAccess Agent volume resource unload script:
 - ◆ Add the MTA and WebAccess Agent or address space unload command(s)
 - ◆ Add the `cvsvbind del` command if you used the `cvsvbind add` command in the load script (NetWare 5.1 only; optional)
 - ◆ Remove the `trustmig` command

See “[Modifying the Volume Resource Unload Script for the WebAccess Agent](#)” on page 98.

- ❑ Set up the WebAccess Agent volume failover path and policies.

See “[Setting the Failover Path and Policies for the WebAccess Agent](#)” on page 99.

- ❑ Test the clustered WebAccess Agent.

See [Section 5.3.5, “Testing Your Clustered WebAccess Installation,”](#) on page 100.

- ❑ Record cluster-specific information in the properties pages of the GroupWise objects associated with the WebAccess Agent.

See [Section 5.4.1, “Updating GroupWise Objects with Cluster-Specific Descriptions,”](#) on page 101.

Implementing GroupWise Gateways in a Novell Cluster

6

A significant 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 each GroupWise gateway. The gateway domain should be created on a cluster-enabled volume. This enables the gateway to fail over independently from other GroupWise components.

If you have set up the Internet Agent or WebAccess in your clustered GroupWise system, you should already have the skills necessary to set up a GroupWise gateway as well.

GroupWise gateways that have not received recent development have not been thoroughly tested in a clustering environment. If you are currently using such GroupWise gateways, you might want to leave them outside of your cluster.

Monitoring a GroupWise System in a NetWare Cluster

7

Because the GroupWise® 7 Monitor currently runs on Windows and Linux, rather than NetWare®, you cannot run GroupWise Monitor in a NetWare cluster. However, GroupWise Monitor can easily monitor a clustered GroupWise system from a vantage point outside the NetWare cluster.

When you first install Monitor, it gathers information about agents to monitor from a domain database (`wpdomain.db`). This provides the secondary IP address of each agent (the IP address associated with the cluster-enabled volume where the agent's domain or post office resides). 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 Monitor must use secondary IP addresses to monitor the agents in a clustered GroupWise system, the Discover Machine and Discover Network options do not work in a cluster. Secondary 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 secondary IP address.

For instructions on setting up GroupWise Monitor, see “[Installing GroupWise Monitor](#)” in the *GroupWise 7 Installation Guide*.

Backing Up a GroupWise System in a NetWare Cluster

8

The GroupWise[®] Target Service Agent (GWTSA) is a GroupWise-specific API that works with compatible backup software to provide reliable backups of a running GroupWise system on NetWare[®] 5.1. On NetWare 6.x, improved functionality is available with the GroupWise Target Service Agent for File Systems (TSAFSGW). These TSAs can be used in a clustered GroupWise system with appropriate preparation and understanding of how the TSAs work. For background information about the GroupWise TSAs, see “[Target Service Agents](#)” in “[Databases](#)” in the *GroupWise 7 Administration Guide*.

- ♦ [Section 8.1, “Using GWTSA in a NetWare 5.1 Cluster,” on page 111](#)
- ♦ [Section 8.2, “Using TSAFSGW in a NetWare 6.x Cluster,” on page 112](#)

8.1 Using GWTSA in a NetWare 5.1 Cluster

In a clustering environment, GWTSA must be installed and loaded on each node from which your backup software backs up any portion of your GroupWise system.

You should not install GWTSA on the cluster-enabled volumes where domains and post offices are located. GWTSA cannot currently run in protected memory, nor can it run re-entrantly. If multiple GroupWise volumes failed over to the same node, and if that node attempted to start an instance of GWTSA for each GroupWise volume, only the first instance would run successfully. Subsequent instances of GWTSA would fail and the domains and post offices specified for backup by those instances would not be available to your backup software.

To accommodate the variable locations of data to back up from a clustered GroupWise system, the `.ncf` file for GWTSA on each node should be edited to include a `/home` startup switch for every domain and post office on every cluster-enabled volume that might ever be mounted on that node. Set the `/home` startup switch to the physical volume name (without the server name), rather than the cluster-enabled volume name, when specifying the location of each domain and post office. For example:

Correct:

```
/home-gwvol1:\gwdom
```

Incorrect:

```
/home-gwcluster_gwvol1:\gwdom
```

Before your backup software starts backing up your GroupWise system, it calls GWTSA to verify the accessibility of each physical volume specified by the `/home` startup switch.

Your backup software then backs up those domains and post offices that are currently accessible to the GroupWise TSA and skips those that are not accessible. The backup runs successfully, even if some locations are not currently mounted on the node where the GroupWise TSA is running. Locations skipped when backups are run from one node will be caught when backups are run from another node where the GroupWise TSA is also running.

If the node where GW TSA is running goes down, the backup needs to be rerun when the node is up again, unless the nodes where the domain and post office volumes failed over are also configured to back them up. Using GW TSA, there is currently no way to restart a backup from a checkpoint position. By configuring redundant backup jobs on all nodes where domains and post offices could fail over, you can ensure that domains and post offices are always backed up, regardless of the node they are mounted on when the backup takes place.

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.

8.2 Using TSAFSGW in a NetWare 6.x Cluster

In a clustering environment, TSAFSGW must be installed and loaded on each node from which your backup software backs up any portion of your GroupWise system. To accommodate the variable locations of data to back up from a clustered GroupWise system, the `.ncf` file for TSAFSGW on each node should be edited to include a `/home` startup switch for every domain and post office on every cluster-enabled volume that might ever be mounted on that node.

If you are using TSAFSGW, the `/vserver` switch enables you to specify the name of a virtual server in your NetWare cluster. Then you can use the `/home` switch to specify shared volumes and paths rather than physical volumes and paths. For example:

```
/vserver-clustername_volumename_SERVER  
/home-volumename:\domaindirectory  
/home-volumename:\postofficedirectory
```

For example, if you have a domain named NewYork and a post office named Engineering, with libraries named Design and Training, the switches for TSAFSGW would be:

```
/vserver-CORPORATE_GROUPWISE_SERVER  
/home-gw:\gwsystem\newyork  
/home-gw:\gwsystem\engineering
```

You can use this same configuration file on every node in the cluster. TSAFSGW can identify the libraries based on information provided by the post office. Your backup software would then list the following locations to back up, based on the previous example:

```
GroupWise Cluster System  
[DOM] newyork  
[PO] engineering  
[DMS] lib0001  
[BLB] design_store  
[DMS] lib0002  
[BLB] training_store
```

From the list provided in your backup software, you can select *GroupWise Cluster System* to back up all the objects listed, or you can select individual objects for backup.

Because TSAFSGW must be installed and loaded on all nodes in the cluster, it can connect to the virtual server from all nodes. If the node where TSAFSGW is performing a backup goes down, that node is dismounted and the next node is mounted to the virtual server. However, because TSAFSGW creates its list of resources when it loads, it is not aware of the newly mounted node. It must be unloaded and reloaded in order to notify your backup software of the change. Your backup

software acknowledges the disruption and attempts to reconnect to the next node. When the next node is fully up and responding, the backup recommences, starting with the resource that was being backed up when the disruption occurred.

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.

Updating a GroupWise System in a NetWare Cluster

9

In a NetWare[®] cluster, you have the option of installing the GroupWise[®] software on each node in the cluster or on a GroupWise volume along with a domain or post office, as described in [Section 2.8.3, “Deciding Where to Install the Agent Software,” on page 34](#). Before you run the GroupWise Installation program to install updated software, make sure you understand where in your cluster the GroupWise software is already installed.

If your existing GroupWise software is installed on each cluster node, it is very important to update all nodes on the failover path of each domain and post office at once, 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 path 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.

If your existing GroupWise software is installed on a GroupWise volume along with a domain or post office, updating the agent software on the GroupWise volume needs to be done only once, because the agents fail over along with the domain or post office they service.

Keep in mind these cluster-specific details as you follow the instructions in “[Update](#)” in the [GroupWise 7 Installation Guide](#) to update your GroupWise system in a NetWare cluster.

Moving an Existing GroupWise 7 System into a NetWare Cluster

10

If you are adding the high availability benefits of Novell® Cluster Services™ to a GroupWise® 7 system that is already up and running, the first step is to install Novell Cluster Services following the instructions in the clustering documentation for your version of NetWare®, as listed in [Chapter 1, “Introduction to GroupWise 7 and Novell Cluster Services on NetWare,” on page 21](#). You should also review [Chapter 1, “Introduction to GroupWise 7 and Novell Cluster Services on NetWare,” on page 21](#) 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 could transfer individual post offices where the needs for high availability are greatest. You could 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 volumes in your shared disk system you will use for GroupWise administration (ConsoleOne® and the software distribution directory).
- ❑ Decide which shared volumes in your shared disk system you will use for GroupWise domains and post offices.
- ❑ Decide which nodes in your storage area network you will have on failover paths for the GroupWise agents.
- ❑ Review [Chapter 2, “Planning GroupWise in a NetWare Cluster,” on page 23](#). Fill out the [“System Clustering Worksheet” on page 37](#) to help you decide which domains and post offices you will move to which shared volumes.
- ❑ Review [“Planning Secondary IP Addresses and Cluster-Unique Port Numbers for Agents in the Cluster” on page 31](#) and fill out the [“IP Address Worksheet” on page 40](#) to select secondary IP addresses for cluster-enabled volumes and to specify cluster-specific port numbers for all of your GroupWise agents.
- ❑ Select the first shared volume that will be part of your clustered GroupWise system and cluster-enable it, following the instructions in [“Cluster-Enabling Shared Volumes for Use with GroupWise” on page 43](#) and [“Configuring Short Name Resolution” on page 44](#).
- ❑ Move a domain and/or post office onto the cluster-enabled volume, following the instructions in [“Moving a Domain” in “Domains”](#) or [“Moving a Post Office” in “Post Offices”](#) in the [GroupWise 7 Administration Guide](#).
- ❑ Review [Section 2.8, “Deciding How to Install and Configure the Agents in a Cluster,” on page 31](#), fill out the [“Agent Clustering Worksheet” on page 41](#), and install the agents as needed for the first clustered domain and/or post office, following the instructions in [Section 3.5, “Installing and Configuring the MTA and the POA in a Cluster,” on page 50](#). This includes setting up the load and unload scripts for the cluster-enabled volume.
- ❑ Test the first component of your clustered GroupWise system following the instructions in [Section 3.6, “Testing Your Clustered GroupWise System,” on page 59](#).

- ❑ Take care of the cluster management details described in [Section 3.7, “Managing Your Clustered GroupWise System,”](#) on page 60.
- ❑ 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 27.
- ❑ Move GroupWise administration into the cluster as needed.
- ❑ Add other components to your clustered GroupWise system as needed, following the instructions in:
 - ◆ [Chapter 4, “Implementing the Internet Agent in a NetWare Cluster,”](#) on page 69
 - ◆ [Chapter 5, “Implementing WebAccess in a NetWare Cluster,”](#) on page 89.
 - ◆ [Chapter 6, “Implementing GroupWise Gateways in a Novell Cluster,”](#) on page 107
 - ◆ [Chapter 7, “Monitoring a GroupWise System in a NetWare Cluster,”](#) on page 109
 - ◆ [Chapter 8, “Backing Up a GroupWise System in a NetWare Cluster,”](#) on page 111

Implementing Messenger in a NetWare Cluster

11

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 NetWare Cluster,” on page 23](#) and [Chapter 3, “Setting Up a Domain and Post Office in a NetWare Cluster,” on page 43](#). As part of the process of setting up GroupWise in the cluster, you filled out the “[System Clustering Worksheet](#)” on page 37. Some of the information from that worksheet is helpful as you implement Messenger in your cluster.

- ♦ [Section 11.1, “Planning Your Messenger System in a Cluster,” on page 119](#)
- ♦ [Section 11.2, “Setting Up Your Messenger System in a Cluster,” on page 122](#)
- ♦ [Section 11.3, “Messenger Clustering Worksheet,” on page 127](#)

11.1 Planning Your Messenger System in a 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 11.3, “Messenger Clustering Worksheet,” on page 127](#) 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:

- ♦ [“Understanding Your Cluster” on page 119](#)
- ♦ [“Planning Messenger Administration” on page 119](#)
- ♦ [“Deciding Where to Install the Messenger Agent Software” on page 120](#)
- ♦ [“Planning the Messenger Agent Installation” on page 122](#)

11.1.1 Understanding Your Cluster

Fill out items 1 through 5 on the [Section 11.3, “Messenger Clustering Worksheet,” on page 127](#) with information about your cluster. This information corresponds to items 1-5 on the “[System Clustering Worksheet](#)” on page 37. For background information, see [Section 2.1, “Meeting Software Version Requirements,” on page 24](#) and [Section 2.2, “Installing Novell Cluster Services,” on page 24](#).

11.1.2 Planning Messenger Administration

If you have set up a cluster-enabled shared volume for GroupWise administration, as described in [Section 2.6, “Deciding Whether to Cluster-Enable the Shared Volumes Used by GroupWise,” on page 27](#), you can use the same cluster-enabled shared volume for the Messenger administration files. For example, you might have a shared `pub:` volume with a `\public` directory where you install the Messenger snap-in to ConsoleOne®. Or you can install the Messenger snap-in on one or more administrator workstations.

MESSENGER CLUSTERING WORKSHEET

Under [Item 6: Installation Location for Messenger Administration](#), mark where you want to install the Messenger snap-in to ConsoleOne.

If you plan to install the Messenger snap-in to ConsoleOne on a cluster-enabled shared volume, under [Item 7: Shared Volume for Messenger Administration](#), list the IP address of the shared volume and the directory where you want to install the Messenger snap-in.

IMPORTANT: Cluster-enabling relies on successful short name resolution throughout your system. Review [Section 2.7, “Ensuring Successful Name Resolution for GroupWise Volumes,” on page 29](#), which describes the issues in the context of planning MTA and POA installations for GroupWise.

11.1.3 Deciding Where to Install the Messenger Agent Software

When you install the NetWare[®] Messenger agents in a clustering environment, you can choose between two different installation locations:

Table 11-1 *Messenger Agent Software Installation Locations*

Location	Description
Each node in the cluster	The <code>sys:\system</code> directory is the default location provided by the Messenger Installation program.
Shared volume	If you create a <code>vol:\system</code> directory on a cluster-enabled shared volume, the agent software and startup files fail over and back along with supporting files such as the Messenger archive.

IMPORTANT: Cluster-enabling relies on successful short name resolution throughout your system. Review [Section 2.7, “Ensuring Successful Name Resolution for GroupWise Volumes,” on page 29](#), which describes the issues in the context of planning MTA and POA installations for GroupWise.

You must install to a cluster-enabled shared volume 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 volume, you should plan to install the Archive Agent separately outside the cluster.

MESSENGER CLUSTERING WORKSHEET

Under [Item 8: Installation Location for Messenger Agents](#), mark where you want to install the Messenger agent software.

Continue with the planning instructions for the installation location you want to use:

- ◆ [“Each Node in the Cluster” on page 121](#)
- ◆ [“Shared Volume” on page 121](#)

Each Node in the Cluster

Make sure you have filled out [item 5](#) on the Messenger Clustering Worksheet with a complete list of nodes in the cluster. Skip to [“Planning the Messenger Agent Installation” on page 122](#).

Shared Volume

For convenience throughout the rest of this section, the term “Messenger volume” means “a cluster-enabled shared volume where the Messenger agents are installed.” Complete the following planning tasks for the Messenger volume:

- ♦ [“Selecting the Messenger Volume” on page 121](#)
- ♦ [“Determining an Appropriate Failover Path for the Messenger Volume” on page 121](#)
- ♦ [“Selecting IP Address Resolution Methods for the Messenger Volume” on page 121](#)

Selecting the Messenger Volume

If you are not planning to enable archiving, or if you are not anticipating a large Messenger archive, you can use one Messenger volume. 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 volume for the Archive Agent and the archive database. The steps in this section cover setting up the Messenger agents on a single Messenger volume.

MESSENGER CLUSTERING WORKSHEET

Under [Item 9: Shared Volume for Messenger Agents](#), record the name and IP address of the Messenger volume.

Determining an Appropriate Failover Path for the Messenger Volume

By default, a Messenger volume is configured to have all nodes in the cluster in its failover path, organized in ascending alphanumeric order. Only one node at a time can have the Messenger volume mounted and active. If a Messenger volume’s preferred node fails, the volume fails over to the next node in the failover path.

MESSENGER CLUSTERING WORKSHEET

Under [Item 10: Failover Path for Messenger Volume](#), list the nodes that you want to have in the Messenger volume failover path. The Messenger agents might need to run on any node that the Messenger volume fails over to.

Selecting IP Address Resolution Methods for the Messenger Volume

Because you are using a cluster-enabled volume for the Messenger agents, you must ensure that short name resolution is always successful. For background information, see [Section 2.7, “Ensuring Successful Name Resolution for GroupWise Volumes,” on page 29](#).

MESSENGER CLUSTERING WORKSHEET

Under [Item 11: IP Address Resolution Methods](#), mark the short name address resolution methods you want to implement to ensure that the UNC paths stored in ConsoleOne can be successfully resolved into the physical network address of the Messenger volume.

11.1.4 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 *Messenger 2.0 Installation Guide*. Transfer the following information from the Messenger Clustering Worksheet to the Messenger System Worksheet:

- ◆ For [Item 3: Installation Path](#) on the Messenger System Worksheet:
 - ◆ If you are installing the Messenger agents to each node in the cluster, use `sys:\system`.
 - ◆ If you are installing the Messenger agents to a Messenger volume, use `volume:\system`, where *volume* is the name of the Messenger volume from [Item 9: Shared Volume for Messenger Administration](#) on the Messenger Clustering Worksheet.
- ◆ Under [Item 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 3: Cluster Identification](#) on the Messenger Clustering Worksheet.
 - ◆ If you are installing the Messenger agents to a Messenger volume, specify the Messenger volume IP address from [Item 9: Shared Volume for Messenger Agents](#) on the Messenger Clustering Worksheet.
- ◆ Under [Item 13: Configure Agents for Clustering?](#) on the Messenger System Worksheet, mark *Yes*. This adds the `/cluster` switch to the agent startup files. The `/cluster` switch tells the Messenger agents to use the virtual server name of the cluster or the Messenger volume rather than the specific server name in pathnames obtained from agent object properties in Novell eDirectory™ or from startup switches. This enables the Messenger agents to access the location no matter which node it is currently running on. This applies to the agents’ working directory, queue directory, log file directory, and so on.
- ◆ 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 volume, use `volume:\directory`, where *volume* is the name of the Messenger administration volume from [Item 7: Shared Volume for Messenger Administration](#) on the Messenger Clustering Worksheet and *directory* is typically `\public`.

Continue with [Setting Up Your Messenger System in a Cluster](#).

11.2 Setting Up Your Messenger System in a Cluster

You should have already reviewed [Section 11.1, “Planning Your Messenger System in a Cluster,”](#) on [page 119](#) and filled out the [Section 11.3, “Messenger Clustering Worksheet,”](#) on [page 127](#) and the “[Novell Messenger Worksheet](#)” in the “[Messenger 2.0 Installation Guide](#)”. Follow the instructions for the installation location you have chosen:

- ◆ [Section 11.2.1, “Installing to Each Node in the Cluster,”](#) on [page 123](#)
- ◆ [Section 11.2.2, “Installing to a Messenger Volume,”](#) on [page 123](#)

11.2.1 Installing to Each Node in the Cluster

There are two methods of installing the Messenger agents to each node in the cluster:

- ♦ Run the Messenger Installation program multiple times in order to install the agent software and to create the agent startup files on each node in the cluster.
- ♦ Run the Messenger Installation program, then copy the Messenger agent software and startup files to each node in the cluster.

Use whichever method you prefer, following the steps provided in “[Starting the Messenger Installation Program](#)” and “[Creating Your Messenger System](#)” in “[Installing a Novell Messenger System](#)” in the *Messenger 2.0 Installation Guide*. Make each node in the cluster active to make sure that the Messenger agents start successfully.

11.2.2 Installing to a Messenger Volume

Complete the following tasks to set up your Messenger system on a Messenger volume:

- ♦ “[Preparing the Cluster for Messenger](#)” on page 123
- ♦ “[Running the Messenger Installation Program](#)” on page 123
- ♦ “[Configuring the Messenger Volume Resource to Load and Unload the Messenger Agents](#)” on page 124
- ♦ “[Copying LDAP and QuickFinder Files to the Messenger Volume](#)” on page 125
- ♦ “[Testing Your Clustered Messenger System](#)” on page 125

Preparing the Cluster for Messenger

Cluster preparation for Messenger is the same as cluster preparation for GroupWise. Review [Section 3.1, “Preparing the Cluster for GroupWise,”](#) on page 43 before running the Messenger installation program.

Running the Messenger Installation Program

The Messenger Installation program walks you through setting up your Messenger system and installing the Messenger agents.

- 1 If necessary, map a drive to the Messenger administration volume ([Messenger Clustering Worksheet item 7](#)).
- 2 Map a drive to the Messenger volume ([Messenger Clustering Worksheet item 9](#)).

The Messenger volume name will be *cluster_volume*. For assistance with mapping a drive to a cluster-enabled volume, see “[Configuring Short Name Resolution](#)” on page 44.

- 3 Run the Messenger Installation program at an administrator workstation to set up your Messenger system, following the steps provided in “[Starting the Messenger Installation Program](#)” and “[Creating Your Messenger System](#)” in “[Installing a Novell Messenger System](#)” in the *Messenger 2.0 Installation Guide*. Keep in mind the following cluster-specific details:
 - ♦ When you specify the Messenger installation directory, be sure to browse to the location through the Messenger volume accessed in [Step 2](#) above.
 - ♦ When you specify the ConsoleOne directory, be sure to browse to the location through the Messenger administration volume accessed in [Step 1](#) above.
 - ♦ On the Start Copying Files page, the server object name should be the virtual server name, not a physical server name.
- 4 When you have finished creating your Messenger system, continue with [Configuring the Messenger Volume Resource to Load and Unload the Messenger Agents](#).

Configuring the Messenger Volume Resource to Load and Unload the Messenger Agents

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

- 1 In ConsoleOne, browse to and select the Cluster object.
If necessary, click *View > Console View* to display its contents.
- 2 Right-click the Volume Resource object (*volume_SERVER*), then click *Properties > Load* to display the default volume resource load script for the Messenger volume.
The volume resource load script executes whenever the Messenger volume comes online.
- 3 Add the following lines to the load script:

```
load volume:\novell\nm\ma\nmma.nlm @volume:novell\nm\ma\strdup.ma
load volume:\novell\nm\aa\nmaa.nlm @volume:novell\nm\aa\strdup.aa
```

where *volume* is the name of the Messenger volume ([Messenger Clustering Worksheet item 9](#)).

For example:

```
load msgr:\novell\nm\ma\nmma.nlm @msgr:novell\nm\ma\strdup.ma
load msgr:\novell\nm\aa\nmaa.nlm @msgr:novell\nm\aa\strdup.aa
```

As an alternative, you can start both Messenger agents with a single command:

```
volume:\novell\nm\nmstart.ncf
```

- 4 Click *Apply* to save the load script.
- 5 Click *Unload*.
- 6 Add the following lines to the unload script:


```
unload nmma.nlm
unload nmaa.nlm
```
- 7 Click *Apply* to save the unload script.
- 8 Click *Nodes* to display the default failover path for the Messenger volume.
- 9 Arrange the nodes in the cluster into the desired failover path for the Messenger volume ([Messenger Clustering Worksheet item 10](#)).

- 10 Click *Apply* to save the failover path.
- 11 Click *Policies* to display the default start, failover, and failback policies.
By default, a volume resource:
 - ♦ Fails over automatically if the node it is running on fails
 - ♦ Starts automatically on the net node in its failover path
 - ♦ Continues running at its failover location even after its most preferred node is again available
- 12 Change the policies if necessary, then click *OK*.
- 13 Continue with **Copying LDAP and QuickFinder Files to the Messenger Volume**.

Copying LDAP and QuickFinder Files to the Messenger Volume

During installation of the Messenger agents, some files were copied to `sys:\system` of the node where the Messenger volume was mounted. These files must be copied one time to the cluster volume that is hosting the Messenger agents.

From the node where the Messenger volume was mounted during installation, copy the following files to the cluster volume:

Table 11-2 *LDAP and QuickFinder Files to Copy*

Copy From:	Copy To:
<code>sys:\system\ldapsdk.nlm</code>	<code>clustervol:\system\ma</code>
<code>sys:\system\ldapssl.nlm</code>	<code>clustervol:\system\ma</code>
<code>sys:\system\ldapx.nlm</code>	<code>clustervol:\system\ma</code>
<code>sys:\system\qfind217.nlm</code>	<code>clustervol:\system\aa</code>

If you are running in a language other than English, copy the following files from the appropriate numbered language subdirectory on the NetWare server to the cluster volume:

Table 11-3 *Language-Specific Files to Copy*

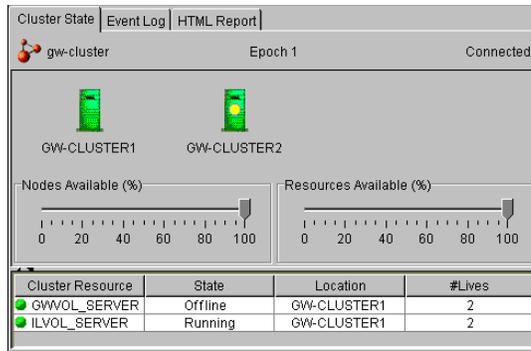
Copy From:	Copy To:
<code>sys:\system\nls\language_code\nmma.msg</code>	<code>clustervol:\system\ma</code>
<code>sys:\system\nls\language_code\nmaa.msg</code>	<code>clustervol:\system\aa</code>

Continue with **Testing Your Clustered Messenger System**.

Testing Your Clustered Messenger System

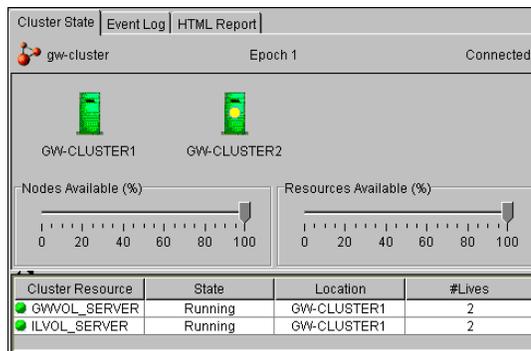
After you have configured the Messenger volume resource, you can test the load and unload scripts by bringing the Messenger volume online and taking it offline again.

- 1 In ConsoleOne, select the Cluster object, then click *View > Cluster State*.



The new Messenger volume resource shows *Offline* in the *State* column.

- Click the new Messenger volume resource, then click *Online*.



The *State* column for the volume resource now displays *Running*.

- Observe the server console where the Messenger agents are loading to see that they start and run correctly.
- Click the new Messenger volume resource, then click *Offline*.
The *State* column for the volume resource returns to *Offline*.
- Observe the server console where the Messenger agents are unloading to see that they shut down correctly.
- Repeat **Step 2** whenever you are ready to bring the new Messenger volume resource online permanently.

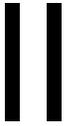
On NetWare 6.x, these actions can also be performed from your Web browser. See [“Using Novell Remote Manager on NetWare 6.x”](#) on page 61.

11.3 Messenger Clustering Worksheet

Item	Explanation
1) Software Version Updates for Cluster: <ul style="list-style-type: none">◆ Support Pack 3 or higher for NetWare 5.1◆ Latest ConsoleOne Snap-In for Novell Cluster Services™	<p>Mark any updates that the nodes in your cluster need in order to meet the system requirements for Messenger system in a cluster.</p> <p>To review the background information provided for GroupWise clustering, see Section 2.1, "Meeting Software Version Requirements," on page 24.</p>
2) eDirectory Tree for Cluster:	<p>Record the eDirectory tree where you created the Novell Cluster object when you installed Novell Cluster Services.</p> <p>To review the background information provided for GroupWise clustering, see Section 2.2, "Installing Novell Cluster Services," on page 24.</p>
3) Cluster Identification: Cluster Name: Cluster IP Address:	<p>Record the name of the name of the NetWare Cluster object where your Messenger system will be located. Also record the virtual IP address of the cluster that will remain constant regardless of which node is currently active.</p> <p>To review the background information provided for GroupWise clustering, see Section 2.2, "Installing Novell Cluster Services," on page 24.</p>
4) Cluster Context:	<p>Record the full context where you created the NetWare Cluster object.</p> <p>To review the background information provided for GroupWise clustering, see Section 2.2, "Installing Novell Cluster Services," on page 24.</p>
5) Nodes in Cluster	<p>List the nodes that are part of the cluster.</p> <p>To review the background information provided for GroupWise clustering, see Section 2.2, "Installing Novell Cluster Services," on page 24.</p>
6) Installation Location for Messenger Administration: <ul style="list-style-type: none">◆ Administrator workstation(s)◆ Shared volume	<p>Mark the location where you will install the Messenger snap-in to ConsoleOne. For more information, see "Planning Messenger Administration" on page 119.</p>

Item	Explanation
7) Shared Volume for Messenger Administration:	If you plan to install the Messenger snap-in to ConsoleOne on a shared volume, specify the name (<i>cluster_volume</i>) of the shared volume where you will install it.
Cluster Volume IP Address:	Specify the IP addresses of the virtual server (<i>volume_SERVER.cluster</i>) to which the shared volume is tied.
Installation Location for Messenger Snap-In to ConsoleOne:	Specify the directory where you will install the Messenger snap-in to ConsoleOne on the shared volume.
<ul style="list-style-type: none"> ◆ /public directory ◆ Other directory 	<p>To review the background information about cluster-enabled volumes provided for GroupWise, see Section 2.6, “Deciding Whether to Cluster-Enable the Shared Volumes Used by GroupWise,” on page 27.</p>
8) Installation Location for Messenger Agents:	Mark the location where you will install the Messaging Agent software.
<ul style="list-style-type: none"> ◆ Each node in the cluster ◆ Shared volume 	For more information, see “Deciding Where to Install the Messenger Agent Software” on page 120.
9) Shared Volume for Messenger Agents:	If you plan to install the Messenger agents on a shared volume, specify the name (<i>cluster_volume</i>) of the shared volume.
Cluster volume IP address:	Specify the IP address of the virtual server (<i>volume_SERVER.cluster</i>) to which the cluster-enabled volume is tied.
	<p>To review the background information about cluster-enabled volumes provided for GroupWise, see Section 2.6, “Deciding Whether to Cluster-Enable the Shared Volumes Used by GroupWise,” on page 27.</p>
	For more information, see “Selecting the Messenger Volume” on page 121.
10) Failover Path for Messenger Shared Volume:	If you plan to install the Messenger agents on a shared volume, list other nodes in the cluster where the Messenger agents could fail over.
	For more information, see “Determining an Appropriate Failover Path for the Messenger Volume” on page 121.
11) IP Address Resolution Methods:	Mark the short name address resolution methods you want to implement to ensure that the UNC paths stored in ConsoleOne can be successfully resolved into physical network addresses.
<ul style="list-style-type: none"> ◆ eDirectory ◆ hosts file ◆ DNS ◆ SLP (highly recommended) 	<p>To review the background information provided for GroupWise, see Section 2.7, “Ensuring Successful Name Resolution for GroupWise Volumes,” on page 29.</p>

Novell Cluster Services on Linux



- ♦ Chapter 12, “Introduction to GroupWise 7 and Novell Cluster Services on Linux,” on page 131
- ♦ Chapter 13, “Planning GroupWise in a Linux Cluster,” on page 133
- ♦ Chapter 14, “Setting Up a Domain and a Post Office in a Linux Cluster,” on page 143
- ♦ Chapter 15, “Implementing the Internet Agent in a Linux Cluster,” on page 167
- ♦ Chapter 16, “Implementing WebAccess in a Linux Cluster,” on page 187
- ♦ Chapter 17, “Implementing GroupWise Monitor in a Linux Cluster,” on page 207
- ♦ Chapter 18, “Backing Up a GroupWise System in a Linux Cluster,” on page 221
- ♦ Chapter 19, “Updating a GroupWise System in a Linux Cluster,” on page 223
- ♦ Chapter 20, “Moving an Existing Linux GroupWise 7 System into a Linux Cluster,” on page 225
- ♦ Chapter 21, “Moving a Clustered GroupWise 7 System from NetWare to Linux,” on page 227
- ♦ Chapter 22, “Implementing Messenger in a Linux Cluster,” on page 229

Introduction to GroupWise 7 and Novell Cluster Services on Linux

12

Before implementing GroupWise® 7 with Novell® Cluster Services™ on Linux, make sure you have a solid understanding of Novell Cluster Services by reviewing the *OES Novell Cluster Services 1.8.2 Administration Guide for Linux*. 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.

NOTE: Although a cluster can include both Linux and NetWare® servers, GroupWise components on Linux servers can fail over only to other Linux servers.

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 “[Cluster Enabling Traditional Linux Volumes on Shared Disks](#)”, “[Creating NSS Volumes](#)”, and “[Creating NSS Shared Disk Partitions and Pools](#)” and in “[Installation and Setup](#)” in the *OES Novell Cluster Services 1.8.2 Administration Guide for Linux*.

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 “[Cluster Enabling Traditional Linux Volumes on Shared Disks](#)” and “[Cluster Enabling NSS Pools and Volumes](#)” in “[Installation and Setup](#)” in the *OES Novell Cluster Services 1.8.2 Administration Guide for Linux*.

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, etc.

cluster resource: A shared partition, secondary IP address, application, service, Web server, etc., 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 setting 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 Internet Agent, each cluster resource could 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, etc.

Planning GroupWise in a Linux Cluster

13

The majority of this part of the *GroupWise 7 Interoperability Guide* (Chapter 13, “Planning GroupWise in a Linux Cluster,” on page 133 through Chapter 18, “Backing Up a GroupWise System in a Linux Cluster,” on page 221) 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 7 system on OES Linux and need to configure it to work in a newly installed cluster, see Chapter 20, “Moving an Existing Linux GroupWise 7 System into a Linux Cluster,” on page 225. If you already have an existing clustered GroupWise 7 system on OES NetWare, see Chapter 21, “Moving a Clustered GroupWise 7 System from NetWare to Linux,” on page 227.

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 7 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) and Post Office Agent (POA)

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.

Section 13.7.1, “System Clustering Worksheet,” on page 140 lists the information you need as you set up GroupWise in a clustering environment. You should print the worksheet and fill it out as you complete the tasks listed below:

- ♦ Section 13.1, “Installing Novell Cluster Services on Linux,” on page 134
- ♦ Section 13.2, “Planning a Clustered Software Distribution Directory,” on page 135
- ♦ Section 13.3, “Planning a New Clustered Domain,” on page 136
- ♦ Section 13.4, “Planning a New Clustered Post Office,” on page 137
- ♦ Section 13.5, “Planning a New Library for a Clustered Post Office,” on page 137
- ♦ Section 13.6, “Deciding How to Install and Configure the Linux Agents in a Cluster,” on page 138
- ♦ Section 13.7, “GroupWise Clustering Worksheets,” on page 139

After you have completed the tasks and filled out Section 13.7.1, “System Clustering Worksheet,” on page 140 and Section 13.7.2, “Agent Clustering Worksheet,” on page 141, you are ready to continue with Chapter 14, “Setting Up a Domain and a Post Office in a Linux Cluster,” on page 143.

13.1 Installing Novell Cluster Services on Linux

Install Novell Cluster Services on OES Linux by following the instructions provided in the *OES Novell Cluster Services 1.8.2 Administration Guide for Linux*.

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 “[Cluster Enabling Traditional Linux Volumes on Shared Disks](#)”, “[Creating NSS Volumes](#)”, and “[Creating NSS Shared Disk Partitions and Pools](#)” in “[Installation and Setup](#)” in the *OES Novell Cluster Services 1.8.2 Administration Guide for Linux*.

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 “[Cluster Enabling Traditional Linux Volumes on Shared Disks](#)” and “[Cluster Enabling NSS Pools and Volumes](#)” in “[Installation and Setup](#)” in the *OES Novell Cluster Services 1.8.2 Administration Guide for Linux*.

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 (*reiserfs*, *ext3*, etc.), device name (*sda2*, *hda1*, etc.), and mount point directory (*/mnt*, */mail*, etc.). 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 shared partitions that are available for use in your GroupWise system.

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 e-mail 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 would be 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 Internet Agent 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 High Availability Service for the Linux GroupWise Agents](#)” in “[Installing GroupWise Agents](#)” in the *GroupWise 7 Installation Guide*. The 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.

13.2 Planning a Clustered Software Distribution Directory

During creation of a GroupWise system, 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 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 7 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.

13.3 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 Your Basic GroupWise System**”, then print and fill out the “**Basic GroupWise System Worksheet**” in “**Installing a Basic GroupWise System**” in the *GroupWise 7 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 “**Domain Worksheet**” in “**Domains**” in the *GroupWise 7 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 **item 17**. If you are filling out the Domain Worksheet, stop with **item 10**.

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 14, “Setting Up a Domain and a Post Office in a Linux Cluster,”** on page 143.

13.4 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 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 “[Post Office Worksheet](#)” in “[Post Offices](#)” in the *GroupWise 7 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 14, “Setting Up a Domain and a Post Office in a Linux Cluster,”](#) on page 143.

13.5 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 7 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 14, “Setting Up a Domain and a Post Office in a Linux Cluster,”](#) on page 143.

13.6 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 and POA in your clustered GroupWise system:

- ♦ [Section 13.6.1, “Recording Secondary IP Addresses for the Agents,” on page 138](#)
- ♦ [Section 13.6.2, “Determining Appropriate Failover Lists for the Agents,” on page 138](#)
- ♦ [Section 13.6.3, “Determining Cluster Resource Information for the Linux Agents,” on page 139](#)
- ♦ [Section 13.6.4, “Planning the Linux Agent Installation,” on page 139](#)

13.6.1 Recording Secondary IP Addresses for the 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 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. For more information, see “[Simplifying Client/Server Access with a GroupWise Name Server](#)” in “[Post Office Agent](#)” in the *GroupWise 7 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.

13.6.2 Determining Appropriate Failover Lists for the 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 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.

13.6.3 Determining Cluster Resource Information for the Linux Agents

A cluster resource is a **shared partition**, secondary IP address, application, service, Web server, etc., that can function successfully anywhere in the cluster. Cluster resources include the GroupWise agents and the Messenger agents. When using 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/gwssystem` 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.

AGENT CLUSTERING WORKSHEET

Under **Item 7: Cluster Resource Information**, list the mount point and secondary IP address for the GroupWise partition where the domain and post office will be located.

13.6.4 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 7 Installation Guide* for each location where you will install the Linux MTA and/or POA.

IMPORTANT: Do not install the Linux agent software until you are instructed to do so in **Chapter 14, “Setting Up a Domain and a Post Office in a Linux Cluster,”** on page 143.

Skip to **Chapter 14, “Setting Up a Domain and a Post Office in a Linux Cluster,”** on page 143.

13.7 GroupWise Clustering Worksheets

- ◆ **Section 13.7.1, “System Clustering Worksheet,”** on page 140
- ◆ **Section 13.7.2, “Agent Clustering Worksheet,”** on page 141

13.7.1 System Clustering Worksheet

Item	Explanation
1) eDirectory Tree for Cluster:	<p>Record the eDirectory tree where you created the new Novell Cluster object when you installed Novell Cluster Services.</p> <p>For more information, see Section 13.1, "Installing Novell Cluster Services on Linux," on page 134</p>
2) Cluster Name: Master IP Address:	<p>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.</p> <p>For more information, see Section 13.1, "Installing Novell Cluster Services on Linux," on page 134.</p>
3) Cluster Context:	<p>Record the full context where you created the new Cluster object.</p> <p>For more information, see Section 13.1, "Installing Novell Cluster Services on Linux," on page 134.</p>
4) Nodes in Cluster: <ul style="list-style-type: none">◆ File system type◆ Device name◆ Mount point directory	<p>List the nodes that are part of the cluster that you set up for your GroupWise system. Also list the file system type (<i>reiserfs</i>, <i>ext3</i>, etc.), device name (<i>sda2</i>, <i>hda1</i>, etc.), and mount point directory (<i>/mnt</i>, <i>/mail</i>, etc.) for each. You need this information as you create load and unload scripts for GroupWise agents.</p> <p>For more information, see Section 13.1, "Installing Novell Cluster Services on Linux," on page 134.</p>
5) Shared Partitions in Cluster:	<p>List the shared partitions that are available for use in your GroupWise system.</p> <p>For more information, see Section 13.1, "Installing Novell Cluster Services on Linux," on page 134.</p>
6) GroupWise Partition for Software Distribution Directory: Secondary IP Address: Directory:	<p>If desired, specify the name of the shared partition where the GroupWise software distribution directory will reside and the full path to its location.</p> <p>For more information, see Section 13.2, "Planning a Clustered Software Distribution Directory," on page 135.</p>
7) GroupWise Partition for Domain: Secondary IP Address: Post Office on Same Partition as Domain? <ul style="list-style-type: none">◆ Yes◆ No	<p>Specify the name of the shared partition where the GroupWise domain will reside and its secondary IP address.</p> <p>For more information, see Section 13.3, "Planning a New Clustered Domain," on page 136.</p>

Item	Explanation
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 13.4, “Planning a New Clustered Post Office,” on page 137.
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 13.3, “Planning a New Clustered Domain,” on page 136.
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 13.4, “Planning a New Clustered Post Office,” on page 137.
11) Document Storage Area Location: <ul style="list-style-type: none"> ◆ 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 13.5, “Planning a New Library for a Clustered Post Office,” on page 137.

13.7.2 Agent Clustering Worksheet

Item	Explanation
1) Domain Name: Domain Location:	Transfer this information from the System Clustering Worksheet (item 9).
2) Domain Failover List:	List other nodes in the cluster where the GroupWise domain and its MTA could fail over. For more information, see Section 13.6.2, “Determining Appropriate Failover Lists for the Agents,” on page 138.
3) MTA Network Information: <ul style="list-style-type: none"> ◆ 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 13.6.1, “Recording Secondary IP Addresses for the Agents,” on page 138.
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 could fail over. For more information, see Section 13.6.2, “Determining Appropriate Failover Lists for the Agents,” on page 138.

Item	Explanation
6) POA Network Information:	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.
<ul style="list-style-type: none"> ◆ POA IP address ◆ POA client/server port ◆ POA message transfer port ◆ POA HTTP port 	See Section 13.6.1, "Recording Secondary IP Addresses for the Agents," on page 138.
7) Cluster Resource Information	List the cluster resource information for the GroupWise partition where the domain and post office serviced by the agents are located.
<ul style="list-style-type: none"> ◆ Path to the cluster resource mount point ◆ IP address of the cluster resource 	For more information, see Section 13.6.3, "Determining Cluster Resource Information for the Linux Agents," on page 139.

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

14

You should have already reviewed [Chapter 13, “Planning GroupWise in a Linux Cluster,”](#) on page 133 and filled out the [Section 13.7.1, “System Clustering Worksheet,”](#) on page 140 and the [Section 13.7.2, “Agent Clustering Worksheet,”](#) on page 141. You are now ready to complete the following tasks to set up GroupWise® in a clustering environment on Linux:

- ♦ [Section 14.1, “Setting Up a New GroupWise System in a Linux Cluster,”](#) on page 143
- ♦ [Section 14.2, “Creating a New Secondary Domain in a Linux Cluster,”](#) on page 144
- ♦ [Section 14.3, “Creating a New Post Office in a Linux Cluster,”](#) on page 145
- ♦ [Section 14.4, “Installing and Configuring the MTA and the POA in a Cluster,”](#) on page 146
- ♦ [Section 14.5, “Testing Your Clustered GroupWise System,”](#) on page 160
- ♦ [Section 14.6, “Managing Your Clustered GroupWise System,”](#) on page 161
- ♦ [Section 14.7, “What’s Next,”](#) on page 163
- ♦ [Section 14.8, “Clustering Quick Checklists,”](#) on page 164

14.1 Setting Up a New GroupWise System in a Linux Cluster

The 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.
If necessary, you can download ConsoleOne for Linux from the [Novell Product Downloads site](http://download.novell.com) (<http://download.novell.com>). ConsoleOne is always installed in `/usr/Consoleone`.
- 3 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 From the *GroupWise 7 Administrator for Linux* CD, run the GroupWise Installation program, as described in [“Starting the GroupWise Installation Program on Linux”](#) in [“Installing a Basic GroupWise System”](#) in the *GroupWise 7 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 7 Installation Guide*.
- 8 When providing the MTA and POA network address information, use the Agent Clustering Worksheet that you filled out in [Section 13.6, “Deciding How to Install and Configure the Linux Agents in a Cluster,” on page 138](#). 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 from the CD.
- 10 Skip to “[Running the GroupWise Installation Program on the Preferred Node](#)” on page 147.

14.2 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 14.1, “Setting Up a New GroupWise System in a Linux Cluster,” on page 143](#), 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 7 Administration Guide*.
- 3 Create the new domain, following the steps provided in “[Creating the New Domain](#)” in “[Domains](#)” in the *GroupWise 7 Administration Guide*.
Use the Domain Worksheet you filled out in [Section 13.3, “Planning a New Clustered Domain,” on page 136](#) 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 7 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 7 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 MTA and the POA in a Cluster](#) to install the MTA software for the new domain.

14.3 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 located.
- 2 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 7 Administration Guide*.
- 4 Create the new post office, following the steps provided in “[Creating the New Post Office](#)” in “[Post Offices](#)” in the *GroupWise 7 Administration Guide*.

Use the Post Office Worksheet you filled out in [Section 13.4, “Planning a New Clustered Post Office,” on page 137](#) to fill in the fields in the Create GroupWise Post Office dialog box.

- 5 Refer to the Agent Clustering Worksheet that you filled in during [Section 13.6, “Deciding How to Install and Configure the Linux Agents in a Cluster,” on page 138](#) for the secondary IP address and port numbers that you need to specify in order to configure the link.
- 6 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 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
- ♦ Enable 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 7 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 7 Administration Guide*. Be sure to browse to the database location ([under System Clustering Worksheet item 11](#)) 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 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 7 Administration Guide*, after you have completely finished setting up the clustered post office.
- 11 Continue with [Installing and Configuring the MTA and the POA in a Cluster](#) to install the POA software for the new post office.

14.4 Installing and Configuring the MTA and the POA in a Cluster

By following the instructions in [Section 14.1, “Setting Up a New GroupWise System in a Linux Cluster,”](#) on page 143, you installed the MTA and the POA on the first node in your cluster, as well as created the initial domain and post office in your GroupWise system. You are now ready to install and configure the agents on additional nodes and set up the agent software for use in your cluster.

- ♦ [Section 14.4.1, “Installing and Setting Up the Agents in Your Cluster,”](#) on page 146
- ♦ [Section 14.4.2, “Changing Agent Paths to Locations on GroupWise Partitions,”](#) on page 150
- ♦ [Section 14.4.3, “Configuring GroupWise Cluster Resources to Load and Unload the Agents,”](#) on page 152
- ♦ [Section 14.4.4, “Setting Up New Instances of the Agents without Installing the Agent Software,”](#) on page 158

If you have added a new secondary domain or a new post office to an existing GroupWise system, 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 14.4.4, “Setting Up New Instances of the Agents without Installing the Agent Software,”](#) on page 158 instead of completing the tasks above.

14.4.1 Installing and Setting Up the 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 GroupWise Installation Program on the Preferred Node”](#) on page 147
- ♦ [“Running the GroupWise Installation Program on Subsequent Nodes”](#) on page 148
- ♦ [“Testing Your Agent Installation on Each Node”](#) on page 150

After you have installed and tested the agents on each node in the cluster, continue with [Section 14.4.2, “Changing Agent Paths to Locations on GroupWise Partitions,”](#) on page 150.

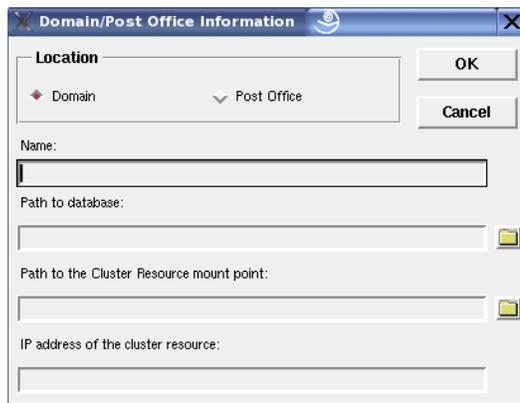
Running the GroupWise Installation Program on the Preferred Node

- 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 14.1, “Setting Up a New GroupWise System in a Linux Cluster,” on page 143](#), start the 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 GroupWise Agents on Linux”](#) in [“Installing GroupWise Agents”](#) in the *GroupWise 7 Installation Guide*.
- 4 Configure the Linux agents according to the [Section 13.7.2, “Agent Clustering Worksheet,” on page 141](#) that you filled out in [Section 13.6, “Deciding How to Install and Configure the Linux Agents in a Cluster,” on page 138](#), 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:

- ♦ 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` 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 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 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 Configure the agents to run as a non-`root` user, as described in the applicable section of the *GroupWise 7 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”
- 6 Continue with *Running the GroupWise Installation Program on Subsequent Nodes*.

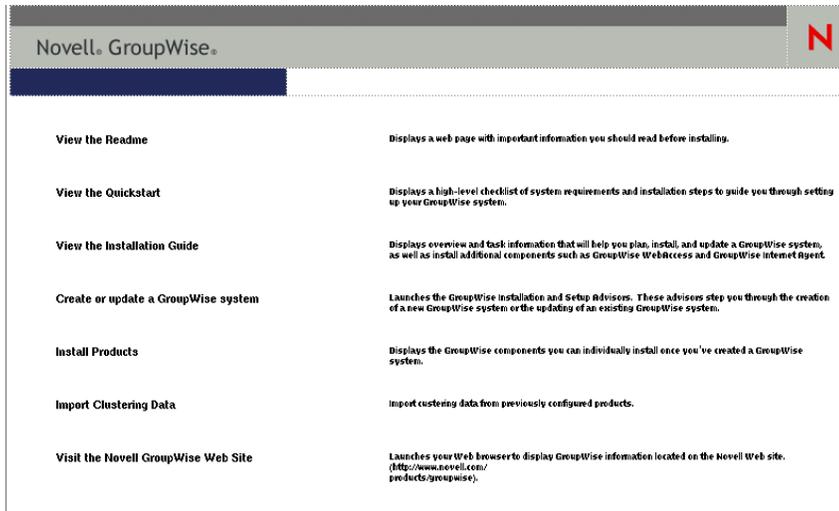
Running the GroupWise Installation Program on Subsequent Nodes

- 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 14.1, “Setting Up a New GroupWise System in a Linux Cluster,”* on page 143, start the GroupWise Installation program.

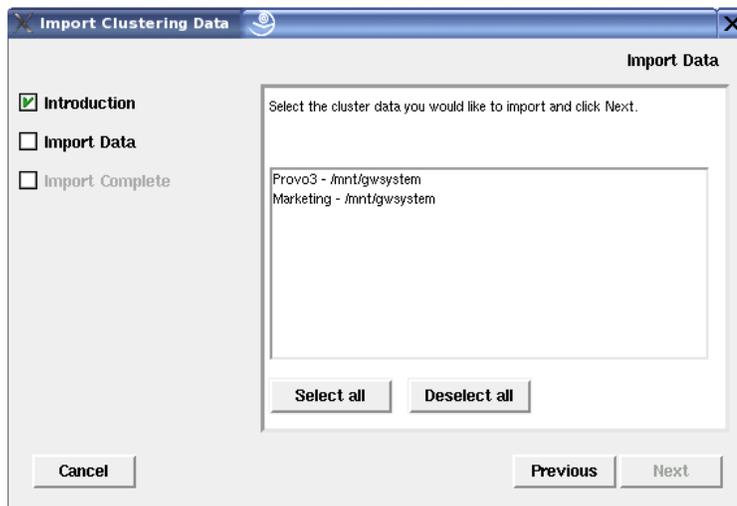
IMPORTANT: You should select the *Configure GroupWise for Clustering* option again.



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 7 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 Agent Installation on Each Node**.

Testing Your Agent Installation on Each Node

1 Test the agents 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 7 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 After you can see that the agents start successfully, test them by starting them as daemons, as described in “Starting the Linux GroupWise Agents as Daemons” in “Installing GroupWise Agents” in the *GroupWise 7 Administration Guide*.

```
/etc/init.d/grpwise start  
/etc/init.d/grpwise status
```

4 Stop the agents.

```
/etc/init.d/grpwise stop  
/etc/init.d/grpwise status
```

5 Return to “Running the GroupWise Installation Program on the Preferred Node” on page 147 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 Agent Paths to Locations on GroupWise Partitions**.

14.4.2 Changing 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 MTA Message Log File Path” on page 151
- ♦ “Setting the MTA Certificate and Key File Path” on page 151
- ♦ “Setting the POA Certificate and Key File Path” on page 151

Setting the 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 7 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 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 7 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 7 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 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 7 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 7 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*.

14.4.3 Configuring GroupWise Cluster Resources to Load and Unload the 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:

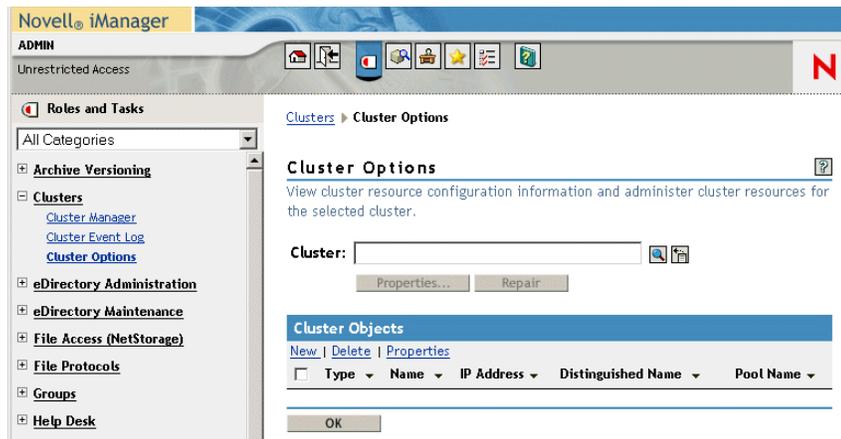
- ♦ “Modifying the Cluster Resource Load Script for the Agents” on page 152
- ♦ “Modifying the Cluster Resource Unload Script for the Agents” on page 155
- ♦ “Setting the Failover List and Policies for the Agents” on page 157

Modifying the Cluster Resource Load Script for the 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:

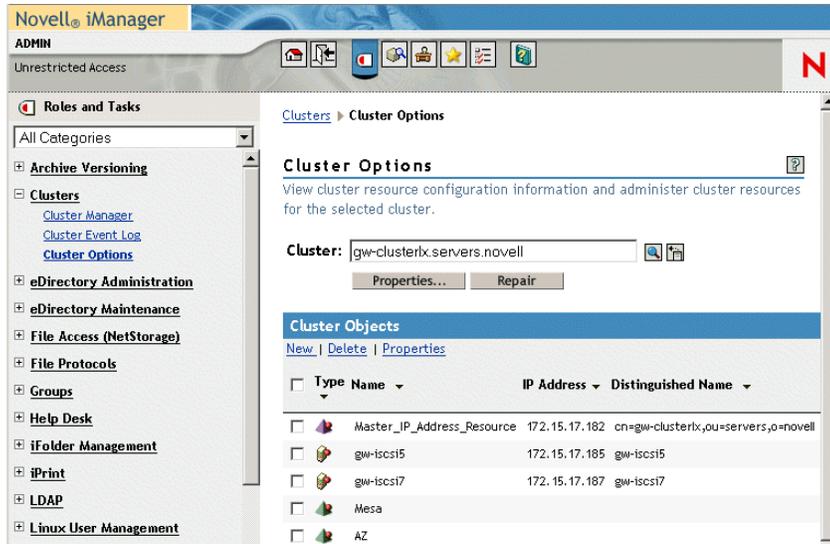
- 1 Expand Clusters, then click *Cluster Options*.



- 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 *Properties*.

The default load script from a generic IP service template appears as follows:

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

# mount the file system
exit_on_error mount -t reiserfs /dev/evms/vol /mnt/generic

# add the IP address
exit_on_error add_secondary_ipaddress a.b.c.d

# start the service
exit_on_error /etc/init.d/myservice start

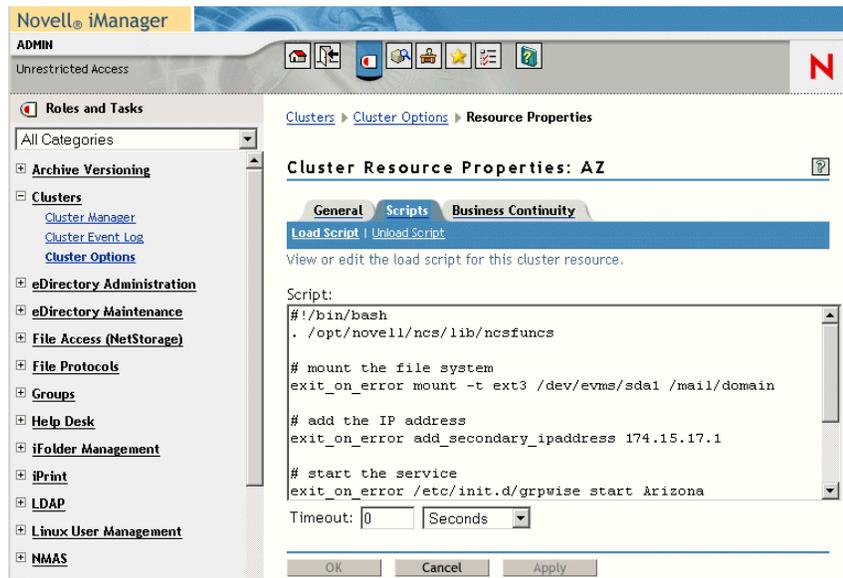
# return status
exit 0
```

- 5 If this is an NSS volume or a shared pool, make the following changes to set up the load script for the agents:
 - 5a As needed, in the `mount` command, change `reiserfs` to whatever file system type is in use on nodes in the cluster ([System Clustering Worksheet item 5](#)).
 - 5b In the `mount` command, change `vol` to the actual device name in use on nodes in the cluster ([System Clustering Worksheet item 5](#)).
 - 5c In the `mount` command, change `/mnt/generic` to the mount point directory in use on nodes in the cluster ([System Clustering Worksheet item 5](#)).
 - 5d In the `add_secondary_ipaddress` command, change `a.b.c.d` to the secondary IP address of the GroupWise partition ([System Clustering Worksheet item 7](#) or [System Clustering Worksheet item 8](#)).
 - 5e In the `start service` command, change `myservice start` to the command to start the GroupWise agent or agents that you want to run on this GroupWise partition.

If you created a domain on the partition, you need to run the MTA. If you created a post office on the partition, you need to run the POA. If you created both a domain and a post office, you need to run both the MTA and the POA. Use the following commands as needed for your GroupWise partition:

```
/etc/init.d/grpwise start domain
/etc/init.d/grpwise start post_office.domain
```

The result would be similar to the following:



6 If this is a traditional Linux volume, use the following load script:

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

# define the IP address
RESOURCE_IP=123.123.1.

# define the file system type
MOUNT_FS=reiserf

# define the device
exit_on_error evms -f /var/opt/novell/ncs/ContainerActivate -rl
                                                    Share `uname -n`
MOUNT_DEV=/dev/evms/Share/dat

# define the mount point
MOUNT_POINT=/mnt/mount_point

# mount the file system
exit_on_error mount -t $MOUNT_FS $MOUNT_DEV $MOUNT_POINT

# add the IP address
exit_on_error add_secondary_ipaddress $RESOURCE_IP
```

```

/etc/init.d/grpwise start domain
/etc/init.d/grpwise start post_office.domain

exit 0

```

Make the following changes to set up the load script for the agents:

- 6a** On the RESOURCE_IP line, change 123.123.1.1 to the secondary IP address of the GroupWise partition ([System Clustering Worksheet item 7](#) or [System Clustering Worksheet item 8](#))
- 6b** As needed on the MOUNT_FS line, change reiserfs to whatever file system type is in use on nodes in the cluster ([System Clustering Worksheet item 5](#)).
- 6c** On the MOUNT_DEV line, change /dev/evms/Share/dat to the actual device name in use on nodes in the cluster ([System Clustering Worksheet item 5](#)).
- 6d** On the MOUNT_POINT line, change /mnt/mount_point to the mount point directory in use on nodes in the cluster ([System Clustering Worksheet item 5](#)).
- 6e** If you created a domain on the partition, use the following command to start the MTA:

```

/etc/init.d/grpwise start domain

```

- 6f** If you created a post office on the partition, use the following command to start the POA:

```

/etc/init.d/grpwise start post_office.domain

```

- 7** Click *OK* to save the load script.

Modifying the Cluster Resource Unload Script for the 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 Cluster Resource Properties page of the GroupWise cluster resource, click *Scripts > Unload Script*.

The default unload script appears as follows:

```

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

# request service stop
ignore_error /etc/init.d/myservice stop

# stop service otherwise
sleep 8
ignore_error fuser -k /mnt/generic

# del the IP address
ignore_error del_secondary_ipaddress a.b.c.d

# umount the file system
exit_on_error umount /mnt/generic

# return status
exit 0

```

2 If this is an NSS volume or a shared pool, make the following changes:

2a In the service stop command, change `myservice stop` to the command to stop the GroupWise agent or agents that you want to run on this GroupWise partition.

If you created a domain on the partition, you need to stop the MTA. If you created a post office on the partition, you need to stop the POA. If you created both a domain and a post office, you need to stop both the MTA and the POA. Use the following commands as needed for your GroupWise partition:

```
/etc/init.d/grpwise stop domain
/etc/init.d/grpwise stop post_office.domain
```

2b Adjust the `sleep` command as needed so that the agents can shut down normally on your system without being inadvertently killed by the kill service command.

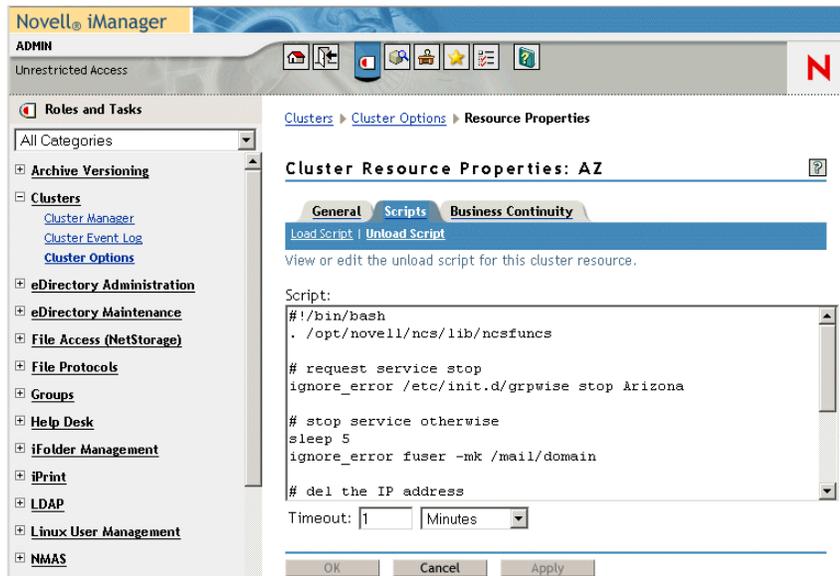
2c In the kill service command (used if the agents do not stop normally), change `-k` to `-mk`. The `-m` parameter obtains the PID numbers of the processes to kill.

2d In the kill service command, change `/mnt/generic` to the mount point directory used in the load script.

2e In the `del_secondary_ipaddress` command, change `a.b.c.d` to the secondary IP address used in the load script.

2f In the `umount` command, change `/mnt/generic` to the mount point directory used in the load script.

The result would be similar to the following:



3 If this is a traditional Linux volume, use the following unload script:

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

/etc/init.d/grpwise stop domain
/etc/init.d/grpwise stop post_office.domain

# define the IP address
```

```

RESOURCE_IP=172.16.5.18

# define the mount point
MOUNT_POINT=/mnt/mount_point

sleep 8

ignore_error fuser -k $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

```

Make the following changes to set up the unload script for the agents.

3a Use the following command to stop the MTA:

```
/etc/init.d/grpwise stop domain
```

3b Use the following command to stop the POA:

```
/etc/init.d/grpwise stop post_office.domain
```

3c On the `RESOURCE_IP` line, change `172.16.5.18` to the secondary IP address used in the load script.

3d On the `MOUNT_POINT` line, change `/mnt/mount_point` to the mount point directory used in the load script.

3e Adjust the `sleep` command as needed so that the agents can shut down normally on your system without being inadvertently killed by the `fuser -k` command that follows.

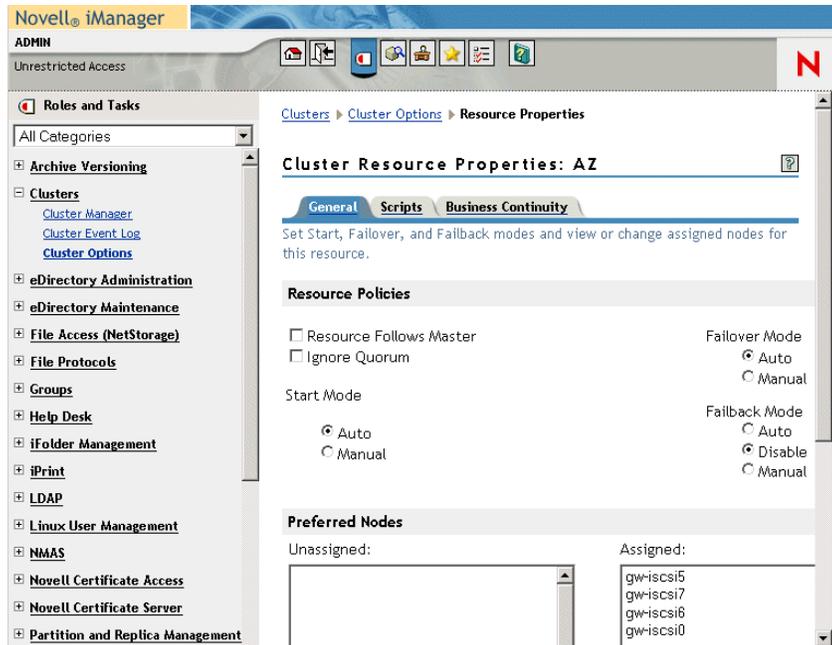
3f In the `fuser -k` command (used if the agents do not stop normally), change `-k` to `-mk`.

The `-m` parameter obtains the PID numbers of the processes to kill.

4 Click *OK* to save the unload script.

Setting the Failover List and Policies for the Agents

1 On the 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 “[Setting Start, Failover, and Failback Modes](#)” in “[Installation and Setup](#)” in the *OES Novell Cluster Services 1.8.2 Administration Guide for Linux*.

- 2 Under *Preferred Nodes*, 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*.

14.4.4 Setting Up New Instances of the 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 158
- ◆ “[Modifying Existing Load and Unload Scripts](#)” on page 159

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 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 looking for 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 High Availability Service for the Linux GroupWise Agents](#)”.
- 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 14.4.3, “Configuring GroupWise Cluster Resources to Load and Unload the Agents,”](#) on [page 152](#) 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 has already been 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.

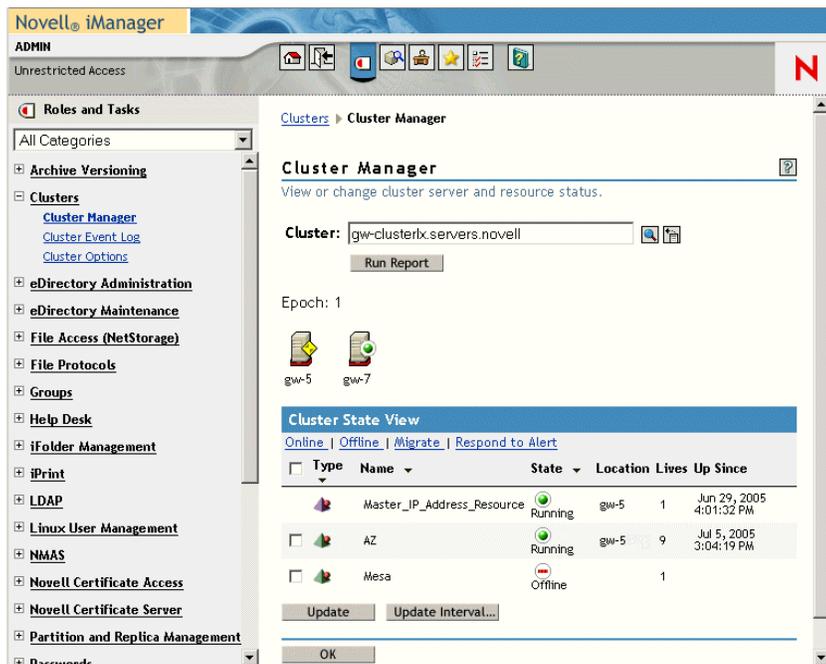
You might want to review other properties of the Cluster Resource object, such as the failover list and the failover/failback/migration procedures on the General page, in light of the fact that an additional domain and/or post office now resides on the GroupWise partition.

- 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](#).

14.5 Testing Your Clustered GroupWise System

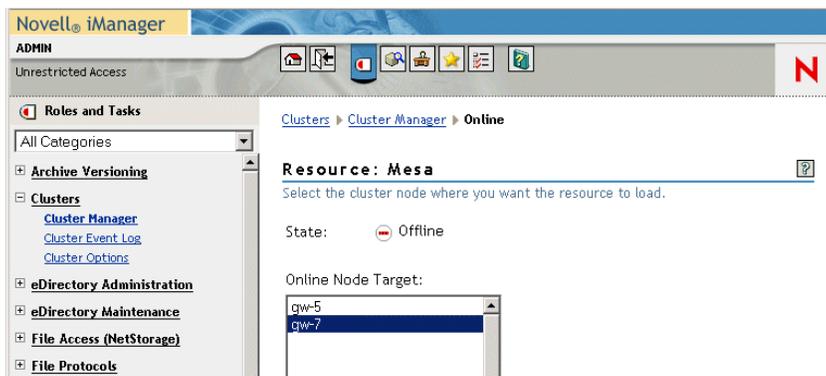
After you have configured the GroupWise cluster resources, you can test the load and unload scripts by bringing the GroupWise cluster resource online and taking it offline again.

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



The new GroupWise cluster resource shows *Offline* in the *State* column.

- 2 Click the new GroupWise cluster resource, then click *Online*.



- 3 Select the cluster node where you want to online the GroupWise cluster resource, then click *OK*.

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:

```
/etc/init.d/grpwise status domain
/etc/init.d/grpwise status post_office.domain
```

- 5 Select the new GroupWise cluster resource, then click *Offline*.
The *State* column for the GroupWise cluster resource returns to *Offline*.
- 6 Use the same command you used in [Step 4](#) to verify if they have stopped.
- 7 Repeat [Step 2](#) whenever you are ready to bring the new GroupWise cluster resource online permanently.
- 8 Continue with [Managing Your Clustered GroupWise System](#).

14.6 Managing Your Clustered GroupWise System

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

- ♦ [Section 14.6.1, “Updating GroupWise Objects with Cluster-Specific Descriptions,” on page 161](#)
- ♦ [Section 14.6.2, “Knowing What to Expect in MTA and POA Failover Situations,” on page 163](#)

14.6.1 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” on page 161](#)
- ♦ [“Recording Cluster-Specific Information for a Post Office and Its POA” on page 162](#)
- ♦ [“Recording Cluster-Specific Information for a Software Distribution Directory” on page 163](#)

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

- 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 If applicable, continue with [Recording Cluster-Specific Information for a Software Distribution Directory](#).

or

Skip to [Section 14.6.2, "Knowing What to Expect in MTA and POA Failover Situations,"](#) on page 163.

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 and POA Failover Situations](#).

14.6.2 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 and then 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.

In comparison to failover, 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](#).

14.7 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 7 Administration Guide*.
- ♦ Setting up the GroupWise client software and helping users to get started using it. See “[Client](#)” in the *GroupWise 7 Administration Guide*. Also see the *GroupWise 7 Windows Client User Guide*.
- ♦ Connecting your clustered GroupWise system to the Internet. See [Chapter 15, “Implementing the Internet Agent in a Linux Cluster,”](#) on page 167.
- ♦ Providing access to users’ GroupWise mailboxes from their Web browsers. See [Chapter 16, “Implementing WebAccess in a Linux Cluster,”](#) on page 187.

- ♦ Monitoring the status of your clustered GroupWise system from your Web browser. See [Chapter 17, “Implementing GroupWise Monitor in a Linux Cluster,”](#) on page 207.
- ♦ Backing up your clustered GroupWise system. See [Chapter 18, “Backing Up a GroupWise System in a Linux Cluster,”](#) on page 221.

14.8 Clustering Quick Checklists

- ♦ [Section 14.8.1, “GroupWise System Quick Checklist,”](#) on page 164
- ♦ [Section 14.8.2, “Domain Quick Checklist,”](#) on page 164
- ♦ [Section 14.8.3, “Post Office Quick Checklist,”](#) on page 165

14.8.1 GroupWise System Quick Checklist

- Plan your new clustered GroupWise system.
See [Chapter 13, “Planning GroupWise in a Linux Cluster,”](#) on page 133.
- Create the primary domain and initial post office in your new clustered GroupWise system.
See [Section 14.1, “Setting Up a New GroupWise System in a Linux Cluster,”](#) on page 143.
- Set up the agents for the primary domain and the initial post office.
See [Section 14.4, “Installing and Configuring the MTA and the POA in a Cluster,”](#) on page 146.
- Modify the cluster resource load scripts:
See [“Modifying the Cluster Resource Load Script for the Agents”](#) on page 152.
- Modify the cluster resource unload scripts:
See [“Modifying the Cluster Resource Unload Script for the Agents”](#) on page 155.
- Set up the cluster failover lists and policies.
See [“Setting the Failover List and Policies for the Agents”](#) on page 157.
- Test your new clustered GroupWise system.
See [Section 14.5, “Testing Your Clustered GroupWise System,”](#) on page 160.
- Record cluster-specific information in the properties pages of the GroupWise objects that the information pertains to.
See [Section 14.6, “Managing Your Clustered GroupWise System,”](#) on page 161.

14.8.2 Domain Quick Checklist

- Plan your new clustered domain.
See [Section 13.3, “Planning a New Clustered Domain,”](#) on page 136.
- Create the new domain.
See [Section 14.2, “Creating a New Secondary Domain in a Linux Cluster,”](#) on page 144.
- Set up the MTA for the new domain.
See [Section 14.4, “Installing and Configuring the MTA and the POA in a Cluster,”](#) on page 146.

- ❑ Modify the domain cluster resource load script.
See [“Modifying the Cluster Resource Load Script for the Agents” on page 152.](#)
- ❑ Modify the domain cluster resource unload script.
See [“Modifying the Cluster Resource Unload Script for the Agents” on page 155.](#)
- ❑ Set up the domain failover list and policies.
See [“Setting the Failover List and Policies for the Agents” on page 157.](#)
- ❑ Test your new clustered domain.
See [Section 14.5, “Testing Your Clustered GroupWise System,” on page 160.](#)
- ❑ Record cluster-specific information in the properties pages of the GroupWise objects that the information pertains to.
See [Section 14.6, “Managing Your Clustered GroupWise System,” on page 161.](#)

14.8.3 Post Office Quick Checklist

- ❑ Plan your new clustered post office.
See [Section 13.4, “Planning a New Clustered Post Office,” on page 137](#) and [Section 13.5, “Planning a New Library for a Clustered Post Office,” on page 137.](#)
- ❑ Create the new post office.
See [Section 14.3, “Creating a New Post Office in a Linux Cluster,” on page 145.](#)
- ❑ Set up the POA for the new post office.
See [Section 14.4, “Installing and Configuring the MTA and the POA in a Cluster,” on page 146.](#)
- ❑ Modify the post office cluster resource load script:
See [“Modifying the Cluster Resource Load Script for the Agents” on page 152.](#)
- ❑ Modify the post office cluster resource unload script:
See [“Modifying the Cluster Resource Unload Script for the Agents” on page 155.](#)
- ❑ Set up the post office failover list and policies.
See [“Setting the Failover List and Policies for the Agents” on page 157.](#)
- ❑ Test your new clustered post office.
See [Section 14.5, “Testing Your Clustered GroupWise System,” on page 160.](#)
- ❑ Record cluster-specific information in the properties pages of the GroupWise objects that the information pertains to.
See [Section 14.6, “Managing Your Clustered GroupWise System,” on page 161.](#)

Implementing the Internet Agent in a Linux Cluster

15

You should already have set up at least a basic GroupWise® system, as described in [Chapter 13, “Planning GroupWise in a Linux Cluster,” on page 133](#) and [Chapter 14, “Setting Up a Domain and a Post Office in a Linux Cluster,” on page 143](#). As part of this process, you filled you the “[System Clustering Worksheet](#)” on page 140. 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 in a cluster.

- ♦ [Section 15.1, “Planning the Internet Agent in a Linux Cluster,” on page 167](#)
- ♦ [Section 15.2, “Setting Up the Internet Agent in a Linux Cluster,” on page 171](#)
- ♦ [Section 15.3, “Testing the Internet Agent in a Linux Cluster,” on page 182](#)
- ♦ [Section 15.4, “Managing the Internet Agent in a Linux Cluster,” on page 183](#)
- ♦ [Section 15.5, “Internet Agent Clustering Worksheet,” on page 184](#)
- ♦ [Section 15.6, “Internet Agent Quick Checklist,” on page 185](#)

15.1 Planning the Internet Agent in a Linux Cluster

A major system configuration difference between the Internet Agent in a clustering environment and the Internet Agent in a regular environment is that you need to create a separate domain to house each GroupWise gateway, including the Internet Agent.

The [Section 15.5, “Internet Agent Clustering Worksheet,” on page 184](#) lists the information you need as you set up the Internet Agent in a clustering environment. You should print the worksheet and fill it out as you complete the tasks listed below:

- ♦ [Section 15.1.1, “Planning a Domain for the Internet Agent,” on page 168](#)
- ♦ [Section 15.1.2, “Selecting the Internet Agent Partition and Secondary IP Address,” on page 168](#)
- ♦ [Section 15.1.3, “Determining an Appropriate Failover List for the Internet Agent,” on page 169](#)
- ♦ [Section 15.1.4, “Determining Cluster Resource Information for the Internet Agent,” on page 169](#)
- ♦ [Section 15.1.5, “Preparing DNS for the Clustered Internet Agent,” on page 169](#)
- ♦ [Section 15.1.6, “Preparing Your Firewall for the Clustered Internet Agent,” on page 169](#)
- ♦ [Section 15.1.7, “Planning the MTA Installation,” on page 170](#)
- ♦ [Section 15.1.8, “Planning the Internet Agent Installation,” on page 170](#)

15.1.1 Planning a Domain for the Internet Agent

The considerations involved in planning a domain for the Internet Agent are much the same as planning any other domain. In preparation, review “[Planning a New Domain](#)”, then print and fill out the “[Domain Worksheet](#)” in “[Domains](#)” in the *GroupWise 7 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 Internet Agent partition because it is where the Internet Agent 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 Internet Agent**, transfer the domain location to the Internet Agent Clustering Worksheet.

Under **Item 2: Internet Agent 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 15.2.1, “Creating a Domain for the Internet Agent,”](#) on page 171.

15.1.2 Selecting the Internet Agent Partition and Secondary IP Address

As with the MTA and the POA, the Internet Agent needs a secondary IP address that remains the same no matter which node in the cluster it is running on. You can place the Internet Agent and its domain on a **GroupWise partition** where a domain or post office already reside, which means that the Internet Agent 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 Internet Agent 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 Internet Agent**, specify the secondary IP address for the Internet Agent partition.

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

Under **Item 6: Internet Agent Network Information**, copy the same secondary IP address.

15.1.3 Determining an Appropriate Failover List for the Internet 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 MTA and the POA, you need to decide which nodes in the cluster are appropriate locations for the **Internet Agent partition** to fail over to. You must install the Internet Agent software on all of the nodes where you want the Internet Agent to be able to fail over. For a review of failover lists, see **Section 13.6.2, "Determining Appropriate Failover Lists for the Agents," on page 138**, which describes the issues in the context of planning MTA and POA installations.

INTERNET AGENT CLUSTERING WORKSHEET

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

15.1.4 Determining Cluster Resource Information for the Internet Agent

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

INTERNET AGENT CLUSTERING WORKSHEET

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

15.1.5 Preparing DNS for the Clustered Internet Agent

In order for the Internet Agent 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 Internet Agent partition. A DNS A record associates the secondary IP address with the hostname.

15.1.6 Preparing Your Firewall for the Clustered Internet Agent

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

Table 15-1 Standard Ports

Protocol	Standard Port
IMAP4	143
LDAP	389
POP3	110
SMTP	25

By default, the Internet Agent 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 Internet Agent partition failover list.

If the Internet Agent has a large number of nodes on its failover list, you could configure the Internet Agent 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 Internet Agent was running. This reduces the amount of modification to your firewall required to set up the Internet Agent. However, if the relay host goes down, outgoing messages are delayed.

As another alternative, you can configure the Internet Agent to use its secondary IP address for sending as well as receiving messages. Setup instructions for this configuration are provided in [“Forcing Use of the Internet Agent Secondary IP Address” on page 181](#), which you can complete after installing the Internet Agent.

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

15.1.7 Planning the MTA Installation

Follow the instructions in [Section 13.6.4, “Planning the Linux Agent Installation,” on page 139](#), 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 15.2, “Setting Up the Internet Agent in a Linux Cluster,” on page 171](#).

15.1.8 Planning the Internet Agent Installation

Aside from the cluster-specific issues discussed in the preceding sections, the considerations involved in planning to install the Internet Agent are the same in a clustering environment as for any other environment. Review the installation instructions in [“Linux: Installing the Internet Agent Software”](#) in [“Installing the GroupWise Internet Agent”](#) in the *GroupWise 7 Installation Guide*. Use the [“GroupWise Internet Agent Installation Worksheet”](#) to record the planning information you will need as you install the Internet Agent in your cluster.

IMPORTANT: Do not install the Internet Agent software until you are instructed to do so in [Setting Up the Internet Agent in a Linux Cluster](#).

15.2 Setting Up the Internet Agent in a Linux Cluster

You should already have reviewed [Section 15.1, “Planning the Internet Agent in a Linux Cluster,” on page 167](#) and filled out [Section 15.5, “Internet Agent Clustering Worksheet,” on page 184](#). You are now ready to complete the following tasks to set up the Internet Agent in a clustering environment:

- ♦ [Section 15.2.1, “Creating a Domain for the Internet Agent,” on page 171](#)
- ♦ [Section 15.2.2, “Installing the MTA for the Internet Agent Domain,” on page 171](#)
- ♦ [Section 15.2.3, “Installing and Configuring the Internet Agent in a Cluster,” on page 171](#)

15.2.1 Creating a Domain for the Internet Agent

The Internet Agent domain will be a secondary domain. To create it, follow the instructions in [Section 14.2, “Creating a New Secondary Domain in a Linux Cluster,” on page 144](#), 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 Internet Agent domain.

After you have created the domain, continue with [Installing the MTA for the Internet Agent Domain](#).

15.2.2 Installing the MTA for the Internet Agent Domain

The MTA for the Internet Agent domain can be installed just like any other MTA in your clustered GroupWise system. Follow the instructions in [Section 14.4.1, “Installing and Setting Up the Agents in Your Cluster,” on page 146](#), 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 Internet Agent.

After you have installed the MTA, continue with [Installing and Configuring the Internet Agent in a Cluster](#).

15.2.3 Installing and Configuring the Internet Agent in a Cluster

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

- ♦ [“Installing and Setting Up the Internet Agent Software in Your Cluster” on page 172](#)
- ♦ [“Configuring the Clustered Internet Agent for SSL” on page 176](#)
- ♦ [“Configuring the Internet Agent Cluster Resource to Load and Unload the Internet Agent and Its MTA” on page 176](#)
- ♦ [“Enabling Internet Addressing for Your Clustered GroupWise System” on page 180](#)
- ♦ [“Verifying Internet Agent Object Properties” on page 180](#)

Installing and Setting Up the Internet Agent Software in Your Cluster

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

- ♦ “Running the Internet Agent Installation Program on the Preferred Node” on page 172
- ♦ “Running the Internet Agent Installation Program on Subsequent Nodes” on page 173
- ♦ “Testing Your Internet Agent Installation on Each Node” on page 175

Running the Internet Agent Installation Program on the Preferred Node

- 1 Make sure that the Internet Agent software is available in the software distribution directory you created in [Step 6](#) in [Section 14.1](#), “Setting Up a New GroupWise System in a Linux Cluster,” on page 143.
- 2 Mount the Internet Agent partition ([Internet Agent Clustering Worksheet item 1](#)) where the Internet Agent message queues are located.
- 3 From the software distribution directory, start the GroupWise Installation program and select *Configure GroupWise for Clustering*.



- 4 Install the Internet Agent software, following the steps provided in “[Linux: Installing the Internet Agent Software](#)” in “[Installing the GroupWise Internet Agent](#)” in the *GroupWise 7 Installation Guide*.
- 5 Configure the Internet Agent according to the “[GroupWise Internet Agent Installation Worksheet](#)” that you filled out in [Section 15.1](#), “Planning the Internet Agent in a Linux Cluster,” on page 167, 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:

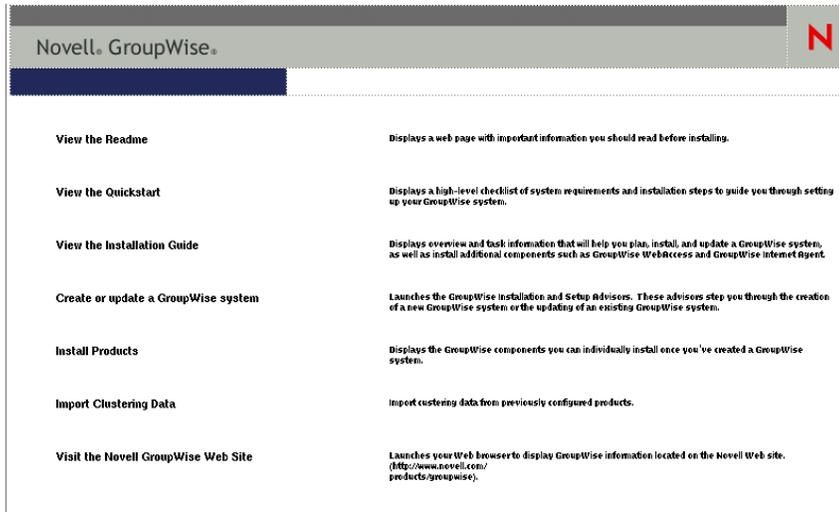
- ◆ The Internet Agent startup file (`gwia.cfg`) is created in `mount_point/groupwise/agents/share` on the shared resource so that the Internet Agent uses the same startup file regardless of which cluster node it is running on. The `--home` switch includes the mount point and the path to the database so that the startup file is valid when mounted to each cluster node.
 - ◆ The `--cluster` switch is added to the Internet Agent startup file to inform the Internet Agent 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 Internet Agent 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 Internet Agent 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 Internet Agent.
 - ◆ 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 Internet Agent on the preferred node is available when you install the Internet Agent on subsequent nodes.
 - ◆ The Internet Agent is not configured to start automatically on system startup. In a cluster, you do not want the Internet Agent to start automatically whenever a node restarts.
- 6** Configure the Internet Agent to run as a non-`root` user, as described in the applicable section of the *GroupWise 7 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 Internet Agent Installation Program on Subsequent Nodes](#).

Running the Internet Agent Installation Program on Subsequent Nodes

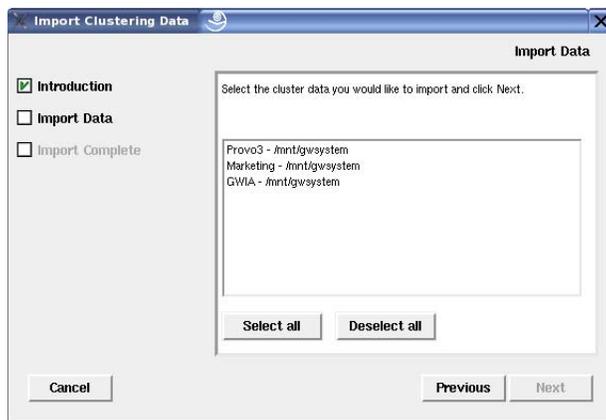
- 1** On the next node in the Internet Agent failover list, mount the Internet Agent partition ([Internet Agent Clustering Worksheet item 1](#)) where the Internet Agent message queues are located.
- 2** From the software distribution directory you created in [Step 6](#) in [Section 14.1](#), “[Setting Up a New GroupWise System in a Linux Cluster](#),” on [page 143](#), 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 Internet 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 Internet Agent to run as a non-`root` user, as described in the applicable section of the *GroupWise 7 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 Internet Agent failover list.
After you install and configure the Internet Agent on each node in its failover list, the cluster node is ready for the Internet Agent to fail over to it.
- 8 Continue with [Testing Your Internet Agent Installation on Each Node](#).

Testing Your Internet Agent Installation on Each Node

- 1 Test the Internet Agent by starting it with a user interface, as described in “[Linux: Starting the Internet Agent](#)” in “[Installing the GroupWise Internet Agent](#)” in the *GroupWise 7 Installation Guide*.

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

- 2 Stop the Internet Agent by clicking *File > Exit*.
- 3 After you can see that the Internet Agent stopped successfully, test it by starting it as a daemon, as described in “[Starting the Linux GroupWise Agents as Daemons](#)” in “[Installing GroupWise Agents](#)” in the *GroupWise 7 Installation Guide*.

```
/etc/inet.d/grpwise start domain.gwia  
/etc/inet.d/grpwise status domain.gwia
```

- 4 Stop the Internet Agent.

```
/etc/inet.d/grpwise stop domain.gwia  
/etc/inet.d/grpwise status domain.gwia
```

- 5 Make sure you have completed all the tasks described in “[Installing the GroupWise Internet Agent](#)” in the *GroupWise 7 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 Internet Agent Installation Program on the Preferred Node](#)” on [page 172](#) for each node in the Internet Agent failover list.

When you have installed the Internet Agent on all of the nodes in the Internet Agent failover list, continue with [Configuring the Clustered Internet Agent for SSL](#).

Configuring the Clustered Internet Agent for SSL

If you plan to enable SSL, as described in “[Securing Internet Agent Connections with SSL](#)” in “[Internet Agent](#)” in the *GroupWise 7 Administration Guide*, you must make the SSL certificate file and key file available to the Internet Agent 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 Internet Agent 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 7 Administration Guide*.
- 3 In ConsoleOne, browse to and select the Domain object.
- 4 Right-click the Internet Agent 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 Internet Agent Cluster Resource to Load and Unload the Internet Agent and Its MTA](#).

Configuring the Internet Agent Cluster Resource to Load and Unload the Internet Agent and Its MTA

The properties of the Cluster Resource object define how the [Internet Agent partition](#) functions within the cluster, how the Internet Agent is loaded and unloaded, and how failover and failback situations are handled. Complete the following tasks for the Internet Agent cluster resource:

- ♦ [“Modifying the Cluster Resource Load Script for the Internet Agent and Its MTA”](#) on page 176
- ♦ [“Modifying the Cluster Resource Unload Script for the Internet Agent and Its MTA”](#) on page 178
- ♦ [“Setting the Failover List and Policies for the Internet Agent and Its MTA”](#) on page 180

Modifying the Cluster Resource Load Script for the Internet Agent and Its MTA

The cluster resource load script executes whenever the Internet 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 Internet Agent cluster resource is located.
- 3 Click the Cluster object to display the cluster resources that belong to the cluster.
- 4 Select the Internet Agent cluster resource that you created when you set up the Internet Agent partition, then click *Properties*.

The default load script from a generic IP service template appears as follows:

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

# mount the file system
exit_on_error mount -t reiserfs /dev/evms/vol /mnt/generic

# add the IP address
exit_on_error add_secondary_ipaddress a.b.c.d

# start the service
exit_on_error /etc/init.d/myservice start

# return status
exit 0
```

5 If this is an NSS volume or a shared pool, make the following changes to set up the Internet Agent load script:

- 5a** As needed, in the `mount` command, change `reiserfs` to whatever file system type is in use on nodes in the cluster.
- 5b** In the `mount` command, change `vol` to the actual device name in use on nodes in the cluster.
- 5c** In the `mount` command, change `/mnt/generic` to the mount point directory in use on nodes in the cluster.
- 5d** In the `add_secondary_ipaddress` command, change `a.b.c.d` to the secondary IP address of the Internet Agent partition (**Internet Agent Clustering Worksheet item 1**)
- 5e** In the start service command, change `myservice start` to the command to start the Internet Agent and its MTA.

```
/etc/init.d/grpwise start domain
/etc/init.d/grpwise start domain.gwia
```

6 If this is a traditional Linux volume, use the following load script:

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

# define the IP address
RESOURCE_IP=123.123.1.

# define the file system type
MOUNT_FS=reiserf

# define the device
exit_on_error evms -f /var/opt/novell/ncs/ContainerActivate -rl
                                                    Share `uname -n`

MOUNT_DEV=/dev/evms/Share/dat

# define the mount point
MOUNT_POINT=/mnt/mount_point

# mount the file system
exit_on_error mount -t $MOUNT_FS $MOUNT_DEV $MOUNT_POINT

# add the IP address
```

```

exit_on_error add_secondary_ipaddress $RESOURCE_IP

/etc/init.d/grpwise start domain
/etc/init.d/grpwise start domain.gwia

exit 0

```

Make the following changes to set up the load script for the Internet Agent:

- 6a** On the `RESOURCE_IP` line, change `123.123.1.1` to the secondary IP address of the GroupWise partition (**Internet Agent Clustering Worksheet item 1**).
- 6b** As needed on the `MOUNT_FS` line, change `reiserfs` to whatever file system type is in use on nodes in the cluster.
- 6c** On the `MOUNT_DEV` line, change `/dev/evms/Share/dat` to the actual device name in use on nodes in the cluster.
- 6d** On the `MOUNT_POINT` line, change `/mnt/mount_point` to the mount point directory in use on nodes in the cluster.
- 6e** If you created a domain on the partition, use the following command to start the MTA:

```
/etc/init.d/grpwise start domain
```

- 6f** Use the following commands to start the Internet Agent and its MTA:

```

/etc/init.d/grpwise start domain
/etc/init.d/grpwise start domain.gwia

```

- 7** Click *OK* to save the load script.

Modifying the Cluster Resource Unload Script for the Internet Agent and Its MTA

The cluster resource unload script executes whenever the Internet Agent 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 Cluster Resource Properties page of the Internet Agent cluster resource, click *Scripts > Unload Script*.

The default unload script appears as follows:

```

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

# request service stop
ignore_error /etc/init.d/myservice stop

# stop service otherwise
sleep 8
ignore_error fuser -k /mnt/generic

# del the IP address
ignore_error del_secondary_ipaddress a.b.c.d

# umount the file system
exit_on_error umount /mnt/generic

# return status
exit 0

```

2 If this is an NSS volume or a shared pool, make the following changes:

2a In the stop service command, change `myservice stop` to the command to stop the Internet Agent and its MTA.

```
/etc/init.d/grpwise stop domain.gwia  
/etc/init.d/grpwise stop domain
```

2b Adjust the `sleep` command as needed so that the Internet Agent and its MTA can shut down normally on your system without being inadvertently killed by the kill service command.

2c In the kill service command (used if the Internet Agent and its MTA do not stop normally), change `-k` to `-mk`.

The `-m` parameter obtains the PID numbers of the processes to kill.

2d In the kill service command, change `/mnt/generic` to the mount point directory used in the load script.

2e In the `del_secondary_ipaddress` command, change `a.b.c.d` to the secondary IP address used in the load script.

2f In the `umount` command, change `/mnt/generic` to the mount point directory used in the load script.

3 If this is a traditional Linux volume, use the following unload script:

```
#!/bin/bash  
. /opt/novell/ncs/lib/ncsfuncs  
  
/etc/init.d/grpwise stop domain.gwia  
/etc/init.d/grpwise stop domain  
  
# define the IP address  
RESOURCE_IP=172.16.5.18  
  
# define the mount point  
MOUNT_POINT=/mnt/mount_point  
  
sleep 8  
  
ignore_error fuser -k $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
```

Make the following changes to set up the unload script for the Internet Agent.

3a Use the following command to stop the Internet Agent:

```
/etc/init.d/grpwise stop domain.gwia
```

3b Use the following command to stop its MTA:

```
/etc/init.d/grpwise stop domain
```

- 3c** On the `RESOURCE_IP` line, change `172.16.5.18` to the secondary IP address used in the load script.
 - 3d** On the `MOUNT_POINT` line, change `/mnt/mount_point` to the mount point directory used in the load script.
 - 3e** Adjust the `sleep` command as needed so that the agents can shut down normally on your system without being inadvertently killed by the `fuser -k` command that follows.
 - 3f** In the `fuser -k` command (used if the agents do not stop normally), change `-k` to `-mk`. The `-m` parameter obtains the PID numbers of the processes to kill.
- 4** Click *OK* to save the unload script.

Setting the Failover List and Policies for the Internet Agent and Its MTA

- 1** On the Cluster Resource Properties page of the Internet 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 availableIf you are considering changing these defaults, see “[Setting Start, Failover, and Failback Modes](#)” in “[Installation and Setup](#)” in the *OES Novell Cluster Services 1.8.2 Administration Guide for Linux*.
- 2** Under *Preferred Nodes*, arrange the nodes in the cluster into the desired failover list for the Internet Agent ([Internet Agent Clustering Worksheet item 3](#)).
- 3** Click *OK*.

Enabling Internet Addressing for Your Clustered GroupWise System

Setting up Internet addressing for a clustered Internet Agent is no different from setting it up for an Internet Agent in any other environment. Follow the instructions in “[Configuring Internet Addressing](#)” in “[Internet Agent](#)” in the *GroupWise 7 Administration Guide*, then return to this point.

Verifying Internet Agent Object Properties

During installation of the Internet Agent, the Internet Agent 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 Internet Agent.

- ♦ “[Accessing Internet Agent Object Properties](#)” on page 180
- ♦ “[Verifying the Reference to the Internet Agent Cluster Resource](#)” on page 181
- ♦ “[Verifying the Reference to the Mount Point Directory](#)” on page 181
- ♦ “[Verifying Post Office Links](#)” on page 181
- ♦ “[Forcing Use of the Internet Agent Secondary IP Address](#)” on page 181

Accessing Internet Agent Object Properties

- 1** In ConsoleOne, browse to and select the Internet Agent domain in order to display its contents.

- 2 Right-click the Internet Agent object, then click *Properties*.
- 3 Continue with [Verifying the Reference to the Internet Agent Cluster Resource](#).

Verifying the Reference to the Internet Agent Cluster Resource

In the Internet Agent object properties pages:

- 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 display the hostname that corresponds to the secondary IP address of the Internet Agent cluster resource. For more information, see [Section 15.1.5, "Preparing DNS for the Clustered Internet Agent," on page 169](#).
- 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 Internet Agent object properties pages:

- 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 Internet Agent object properties pages:

- 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 Internet Agent.
- 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 Internet Agent Secondary IP Address](#).

Forcing Use of the Internet Agent Secondary IP Address

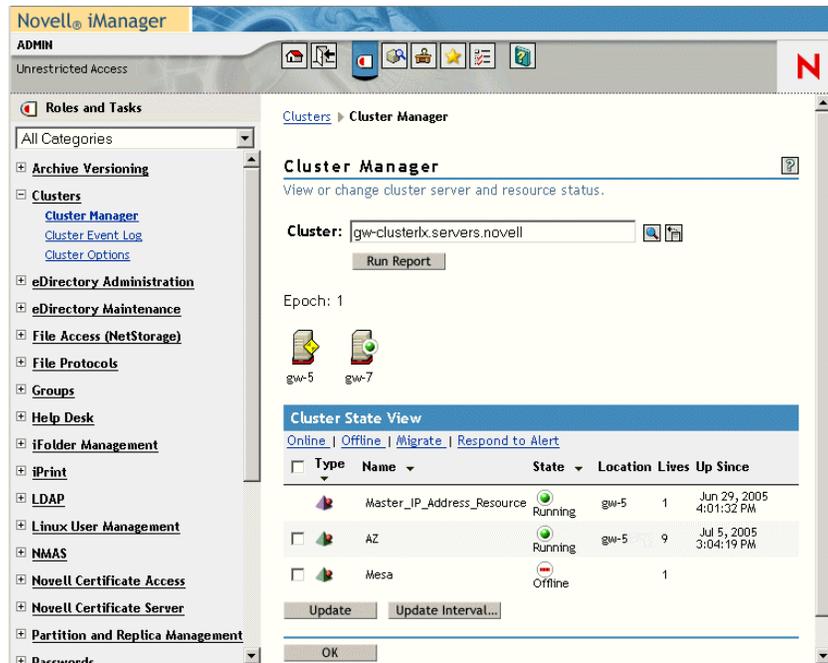
If you want the Internet Agent to send outgoing messages on its secondary IP address, rather than using the default of its primary IP address:

- 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 Internet Agent to use for sending outgoing messages.
- 3 Select *Bind Exclusively to TCP/IP Address*.
- 4 Click *OK*.
- 5 Continue with [Testing the Internet Agent in a Linux Cluster](#).

15.3 Testing the Internet Agent in a Linux Cluster

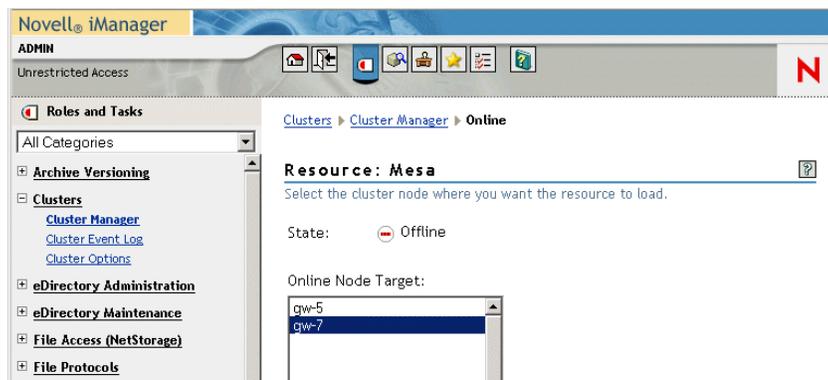
After you have configured the Internet Agent cluster resource, you can test the load and unload scripts by bringing the cluster resource online and taking it offline again.

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



The new Internet Agent cluster resource shows *Offline* in the *State* column.

- 2 Click the new Internet Agent cluster resource, then click *Online*.



- 3 Select the cluster node where you want to online the Internet Agent cluster resource, then click *OK*.

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

- 4 At the server where the Internet Agent is starting, use the following command to see that the Internet Agent has started:

```
/etc/init.d/grpwise status domain.gwia
```

- 5 Select the new Internet Agent cluster resource, then click *Offline*.
The *State* column for the Internet Agent cluster resource returns to *Offline*.
- 6 Use the same command that you used in [Step 4](#) to verify that the Internet Agent has stopped.
- 7 Repeat [Step 2](#) whenever you are ready to bring the new Internet Agent cluster resource online permanently.
- 8 Continue with [Managing Your Clustered GroupWise System](#).

15.4 Managing the Internet Agent in a Linux Cluster

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

- ♦ [Section 15.4.1, “Updating GroupWise Objects with Cluster-Specific Descriptions,” on page 183](#)
- ♦ [Section 15.4.2, “Knowing What to Expect in an Internet Agent Failover Situation,” on page 184](#)

15.4.1 Updating GroupWise Objects with Cluster-Specific Descriptions

After installing the Internet Agent 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 Internet Agent Domain and Its MTA” on page 183](#)
- ♦ [“Recording Cluster-Specific Information about the Internet Agent” on page 184](#)

Recording Cluster-Specific Information about the Internet Agent Domain and Its MTA

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

- 1 In ConsoleOne, browse to and right-click the Domain object, then click *Properties*.
- 2 In the *Description* field of the Internet Agent domain Identification page, provide a cluster-specific description of the Internet Agent domain, including the secondary IP address of its [GroupWise partition](#).
- 3 Click *OK* to save the Internet Agent domain description.
- 4 Select the Internet Agent 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 Internet Agent](#).

Recording Cluster-Specific Information about the Internet Agent

With the contents of the Internet Agent domain still displayed:

- 1 Right-click the Internet Agent 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 Internet Agent domain is located.

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

- 4 Click *OK* to save the Internet Agent information.
- 5 Continue with [Knowing What to Expect in an Internet Agent Failover Situation](#).

15.4.2 Knowing What to Expect in an Internet Agent Failover Situation

The failover behavior of the MTA for the Internet Agent domain is the same as for an MTA in a regular domain. See [Section 14.6.2, “Knowing What to Expect in MTA and POA Failover Situations,”](#) on page 163.

Failover of the Internet Agent 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 Internet Agent restarts quickly in its failover location so users can reconnect quickly.

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

15.5 Internet Agent Clustering Worksheet

Item	Explanation
1) GroupWise Partition for the Internet Agent:	Specify the GroupWise partition where the Internet Agent domain will be created, along with its secondary IP address.
Secondary IP Address:	For more information, see Section 15.1.2, “Selecting the Internet Agent Partition and Secondary IP Address,” on page 168.
2) Internet Agent Domain Name:	Specify a unique name for the Internet Agent domain. Specify the directory on the GroupWise partition where you want to create the new domain.
Domain Database Location:	For more information, see Section 15.1.1, “Planning a Domain for the Internet Agent,” on page 168.

Item	Explanation
3) Internet Agent Failover List:	List other nodes in the cluster where the Internet Agent and its MTA could fail over. For more information, see Section 15.1.3, “Determining an Appropriate Failover List for the Internet Agent,” on page 169.
4) Cluster Resource Mount Point:	Specify the mount point directory where the Internet Agent domain will be mounted. For more information, see Section 15.1.4, “Determining Cluster Resource Information for the Internet Agent,” on page 169.
5) MTA Network Information:	Record the MTA network address information that you will need as you install the MTA.
<ul style="list-style-type: none"> ◆ MTA IP address ◆ MTA message transfer port ◆ MTA live remote port ◆ MTA HTTP port 	For more information, see Section 15.1.7, “Planning the MTA Installation,” on page 170.
6) Internet Agent Network Information:	Record the Internet Agent network address information that you will need to install the Internet Agent.
<ul style="list-style-type: none"> ◆ Internet Agent IP address ◆ Internet Agent HTTP port 	For more information, see Section 15.1.8, “Planning the Internet Agent Installation,” on page 170.

15.6 Internet Agent Quick Checklist

- Plan the new clustered Internet Agent, including the new domain required to house the Internet Agent in a clustering environment.
See [Section 15.1, “Planning the Internet Agent in a Linux Cluster,”](#) on page 167.
- Make sure DNS includes the secondary IP address of the Internet Agent partition
See [Section 15.1.5, “Preparing DNS for the Clustered Internet Agent,”](#) on page 169
- Make sure your firewall is configured to accommodate the Internet Agent.
See [Section 15.1.6, “Preparing Your Firewall for the Clustered Internet Agent,”](#) on page 169.
- Create the new Internet Agent domain.
See [Section 15.2.1, “Creating a Domain for the Internet Agent,”](#) on page 171.
- Set up the MTA for the new Internet Agent domain.
See [Section 15.2.2, “Installing the MTA for the Internet Agent Domain,”](#) on page 171.
- Install the Internet Agent.
See [“Installing and Setting Up the Internet Agent Software in Your Cluster”](#) on page 172.
- Modify the Internet Agent cluster resource load script.
See [“Modifying the Cluster Resource Load Script for the Internet Agent and Its MTA”](#) on page 176.
- Modify the Internet Agent cluster resource unload script.
See [“Modifying the Cluster Resource Unload Script for the Internet Agent and Its MTA”](#) on page 178.

- ❑ Set up the Internet Agent failover list and policies.
See [“Setting the Failover List and Policies for the Internet Agent and Its MTA”](#) on page 180.
- ❑ Enable Internet Addressing for the clustered Internet Agent.
See [“Enabling Internet Addressing for Your Clustered GroupWise System”](#) on page 180.
- ❑ Double-check the cluster-specific Internet Agent object properties.
See [“Verifying Internet Agent Object Properties”](#) on page 180.
- ❑ Test the clustered Internet Agent.
See [Section 15.3, “Testing the Internet Agent in a Linux Cluster,”](#) on page 182.
- ❑ Record cluster-specific information in the properties pages of the GroupWise objects associated with the Internet Agent.
See [Section 15.4.1, “Updating GroupWise Objects with Cluster-Specific Descriptions,”](#) on page 183.

Implementing WebAccess in a Linux Cluster

16

You should already have set up at least a basic GroupWise® system, as described in [Chapter 13, “Planning GroupWise in a Linux Cluster,”](#) on page 133 and [Chapter 14, “Setting Up a Domain and a Post Office in a Linux Cluster,”](#) on page 143. As part of this process, you filled out the “[System Clustering Worksheet](#)” on page 140. 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 WebAccess Agent in a cluster.

- ♦ [Section 16.1, “Understanding the WebAccess Components,”](#) on page 187
- ♦ [Section 16.2, “Planning the WebAccess Agent in a Linux Cluster,”](#) on page 187
- ♦ [Section 16.3, “Setting Up the WebAccess Agent in a Linux Cluster,”](#) on page 190
- ♦ [Section 16.4, “Testing the WebAccess Agent in a Linux Cluster,”](#) on page 201
- ♦ [Section 16.5, “Managing the WebAccess Agent in a Linux Cluster,”](#) on page 202
- ♦ [Section 16.6, “WebAccess Agent Clustering Worksheet,”](#) on page 204
- ♦ [Section 16.7, “WebAccess Agent Quick Checklist,”](#) on page 205

16.1 Understanding the WebAccess Components

If you are not familiar with GroupWise WebAccess, review “[GroupWise WebAccess Overview](#)” in “[Installing GroupWise WebAccess](#)” in the *GroupWise 7 Installation Guide*.

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

- ♦ WebAccess Agent ([gwinter](#)) that will be associated with a domain in your GroupWise system
- ♦ WebAccess Application (a Java servlet) that will be added to your Web server (Apache). The WebAccess Application is currently not supported on a cluster resource.

You install the WebAccess Agent on each node in the cluster. You install the WebAccess Application to your Web server, which must not be clustered. This means that the WebAccess client login to the Web server at the following URL:

```
http://secondary_IP_address/gw/webacc
```

is available only when the Web server is running.

16.2 Planning the WebAccess Agent in a Linux Cluster

In a cluster, you need to create a separate domain to house each GroupWise gateway, including the WebAccess Agent.

Section 16.6, “WebAccess Agent Clustering Worksheet,” on page 204 lists the information you need as you set up the WebAccess Agent in a clustering environment. You should print the worksheet and fill it out as you complete the tasks listed below:

- ♦ Section 16.2.1, “Planning a Domain for the WebAccess Agent,” on page 188
- ♦ Section 16.2.2, “Selecting the WebAccess Agent Partition and Secondary IP Address,” on page 189
- ♦ Section 16.2.3, “Determining an Appropriate Failover List for the WebAccess Agent,” on page 189
- ♦ Section 16.2.4, “Determining Cluster Resource Information for the WebAccess Agent,” on page 189
- ♦ Section 16.2.5, “Planning the MTA Installation,” on page 190
- ♦ Section 16.2.6, “Planning the WebAccess Agent Installation,” on page 190
- ♦ Section 16.2.7, “Planning the WebAccess Application Installation,” on page 190

16.2.1 Planning a Domain for the WebAccess Agent

The considerations involved in planning a domain for the WebAccess Agent are much the same as planning any other domain. In preparation, review “[Planning a New Domain](#)”, then print and fill out the “[Domain Worksheet](#)” in “[Domains](#)” in the *GroupWise 7 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 WebAccess Agent partition because it is where the WebAccess Agent 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 WebAccess Agent Clustering Worksheet.

WEBACCESS AGENT CLUSTERING WORKSHEET

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

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

IMPORTANT: Do not create the new domain until you are instructed to do so in [Section 16.3.1, “Creating a Domain for the WebAccess Agent,”](#) on page 191.

16.2.2 Selecting the WebAccess Agent Partition and Secondary IP Address

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

WEBACCESS AGENT CLUSTERING WORKSHEET

Under **Item 1: GroupWise Partition for WebAccess Agent**, specify the secondary IP address for the WebAccess Agent partition.

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

Under **Item 6: WebAccess Agent Network Information**, copy the same secondary IP address.

16.2.3 Determining an Appropriate Failover List for the WebAccess 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 MTA and the POA, you need to decide which nodes in the cluster are appropriate locations for the **WebAccess Agent partition** to fail over to. You must install the WebAccess Agent software on all of the nodes where you want the WebAccess Agent to be able to fail over. For a review of failover lists, see **Section 13.6.2, "Determining Appropriate Failover Lists for the Agents," on page 138**, which describes the issues in the context of planning MTA and POA installations.

WEBACCESS AGENT CLUSTERING WORKSHEET

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

16.2.4 Determining Cluster Resource Information for the WebAccess Agent

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

WEBACCESS AGENT CLUSTERING WORKSHEET

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

16.2.5 Planning the MTA Installation

Follow the instructions in **Section 13.6.4, “Planning the Linux Agent Installation,”** on page 139, then return to this point. After you follow the instructions, you 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 16.3, “Setting Up the WebAccess Agent in a Linux Cluster,”** on page 190.

16.2.6 Planning the WebAccess Agent Installation

Aside from the cluster-specific issues discussed in the preceding sections, the considerations involved in planning to install the WebAccess Agent are the same in a clustering environment as for any other environment. Review the installation instructions in “**Installing the Linux WebAccess Agent**” in “**Installing GroupWise WebAccess**” in the *GroupWise 7 Installation Guide*. Use the “**GroupWise WebAccess Installation Worksheet**” to record the planning information you will need as you install the WebAccess Agent in your cluster.

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

16.2.7 Planning the WebAccess Application Installation

The WebAccess Application must be installed on a non-clustered Web server. For WebAccess Application planning information, see “**Determining the WebAccess and WebPublisher Applications’ Configuration**” in “**Installing GroupWise WebAccess**” in the *GroupWise 7 Installation Guide*.

16.3 Setting Up the WebAccess Agent in a Linux Cluster

You should already have reviewed **Section 16.2, “Planning the WebAccess Agent in a Linux Cluster,”** on page 187 and filled out **Section 16.6, “WebAccess Agent Clustering Worksheet,”** on page 204. You are now ready to complete the following tasks to set up the WebAccess Agent in a clustering environment:

- ♦ **Section 16.3.1, “Creating a Domain for the WebAccess Agent,”** on page 191
- ♦ **Section 16.3.2, “Installing the MTA for the WebAccess Agent Domain,”** on page 191
- ♦ **Section 16.3.3, “Installing and Configuring the WebAccess Agent in a Cluster,”** on page 191
- ♦ **Section 16.3.4, “Installing and Configuring the WebAccess Application in a Cluster,”** on page 201

16.3.1 Creating a Domain for the WebAccess Agent

The WebAccess Agent domain will be a secondary domain. To create it, follow the instructions in [Section 14.2, “Creating a New Secondary Domain in a Linux Cluster,” on page 144](#), taking your information from the WebAccess Agent Clustering Worksheet, rather than the System Clustering Worksheet, then return to this point.

Do not create any post offices in the WebAccess Agent domain.

After you have created the domain, continue with [Installing the MTA for the WebAccess Agent Domain](#).

16.3.2 Installing the MTA for the WebAccess Agent Domain

The MTA for the WebAccess Agent domain can be installed just like any other MTA in your clustered GroupWise system. Follow the instructions in [Section 14.4.1, “Installing and Setting Up the Agents in Your Cluster,” on page 146](#), 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 WebAccess Agent.

After you have installed the MTA, continue with [Installing and Configuring the WebAccess Agent in a Cluster](#).

16.3.3 Installing and Configuring the WebAccess Agent in a Cluster

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

- ♦ [“Installing and Setting Up the WebAccess Agent Software in Your Cluster” on page 191](#)
- ♦ [“Configuring Clustered Logging” on page 195](#)
- ♦ [“Configuring the Clustered WebAccess Agent for SSL” on page 195](#)
- ♦ [“Configuring the WebAccess Agent Cluster Resource to Load and Unload the WebAccess Agent and Its MTA” on page 196](#)
- ♦ [“Verifying WebAccess Agent Object Properties” on page 200](#)

Installing and Setting Up the WebAccess Agent Software in Your Cluster

The WebAccess Agent must be installed on each node in its failover list ([WebAccess Agent Clustering Worksheet item 3](#)).

- ♦ [“Running the WebAccess Agent Installation Program on the Preferred Node” on page 191](#)
- ♦ [“Running the WebAccess Agent Installation Program on Subsequent Nodes” on page 193](#)
- ♦ [“Testing Your WebAccess Agent Installation on Each Node” on page 194](#)

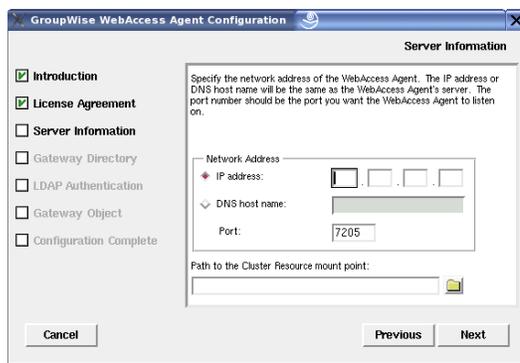
Running the WebAccess Agent Installation Program on the Preferred Node

- 1 Make sure that the WebAccess Agent software is available in the software distribution directory you created in [Step 6 in Section 14.1, “Setting Up a New GroupWise System in a Linux Cluster,” on page 143](#).

- 2 Mount the WebAccess Agent partition ([WebAccess Agent Clustering Worksheet item 1](#)) where the WebAccess Agent message queues are located.
- 3 From the software distribution directory, start the GroupWise Installation program and select *Configure GroupWise for Clustering*.



- 4 Install the WebAccess Agent software, following the steps provided in “[Installing the Linux WebAccess Agent](#)” in “[Installing GroupWise WebAccess](#)” in the *GroupWise 7 Installation Guide*.
- 5 Configure the WebAccess Agent according to the “[GroupWise WebAccess Installation Worksheet](#)” that you filled out in [Section 16.2, “Planning the WebAccess Agent in a Linux Cluster,”](#) on page 187, 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 WebAccess Agent startup file (*webac70a.waa*) is created in *mount_point/groupwise/agents/share* on the shared resource so that the WebAccess Agent 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.
- ♦ 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 WebAccess Agent.

- ◆ 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 WebAccess Agent on the preferred node is available when you install the WebAccess Agent on subsequent nodes.
- ◆ The WebAccess Agent is not configured to start automatically on system startup. In a cluster, you do not want the WebAccess Agent to start automatically whenever a node restarts.

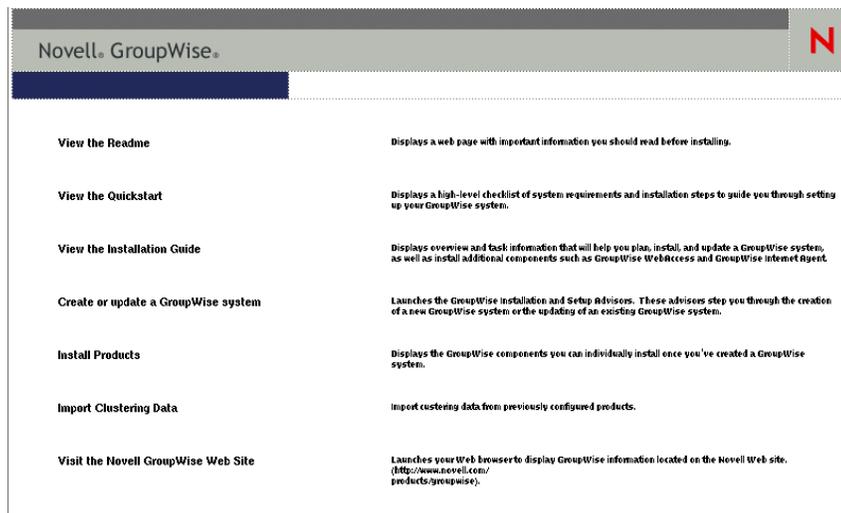
6 Continue with [Running the WebAccess Agent Installation Program on Subsequent Nodes](#).

Running the WebAccess Agent Installation Program on Subsequent Nodes

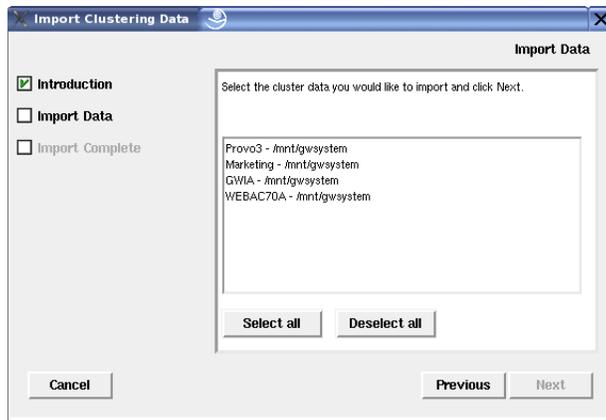
- 1 On the next node in the WebAccess Agent failover list, mount the WebAccess Agent partition ([WebAccess Agent Clustering Worksheet item 1](#)) where the WebAccess Agent message queues are located.
- 2 From the software distribution directory you created in [Step 6 in Section 14.1, “Setting Up a New GroupWise System in a Linux Cluster,” on page 143](#), 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 WebAccess 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 Repeat [Step 1](#) through [Step 5](#) for each cluster node in the WebAccess Agent failover list.

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

- 7 Continue with [Testing Your WebAccess Agent Installation on Each Node](#).

Testing Your WebAccess Agent Installation on Each Node

- 1 Test the WebAccess Agent by starting it with a user interface, as described in [“Starting the Linux WebAccess Agent”](#) in [“Installing GroupWise WebAccess”](#) in the *GroupWise 7 Installation Guide*.

```
/opt/novell/groupwise/agents/bin/gwinter --show @webac70a.waa &
```

- 2 Stop the WebAccess Agent by closing the window that it is running in.
- 3 After you can see that the WebAccess Agent stopped successfully, test it by starting it as a daemon, as described in [“Starting the Linux GroupWise Agents as Daemons”](#) in [“Installing GroupWise Agents”](#) in the *GroupWise 7 Installation Guide*.

```
/etc/inet.d/grpwise start webac70a
/etc/inet.d/grpwise status webac70a
```

- 4 Stop the WebAccess Agent.

```
/etc/inet.d/grpwise stop webac70a
/etc/inet.d/grpwise status webac70a
```

- 5 Repeat **Step 1** through **Step 4** for each node in the WebAccess Agent failover list.
- 6 Continue with **“Configuring Clustered Logging” on page 195**

Configuring Clustered Logging

The default location for the WebAccess Agent log files is on the cluster nodes along with the WebAccess Agent software, rather than being located with the domain on the WebAccess Agent partition. To avoid having multiple copies of log files in multiple locations, you should set the location in ConsoleOne.

- 1 On the WebAccess Agent partition where the domain is located, create the directory where you want to store WebAccess Agent log files.
- 2 In ConsoleOne, browse to and select the Domain object.
- 3 Right-click the WebAccess Agent object, then click *Properties*.
- 4 Click *GroupWise > Log Settings*.
- 5 In the *Log File Path* field, browse to and select the directory you created in **Step 1**, then click *OK*.
- 6 If you want the WebAccess Agent to use SSL connections, continue with **Configuring the Clustered WebAccess Agent for SSL**
or
Skip to **“Configuring the WebAccess Agent Cluster Resource to Load and Unload the WebAccess Agent and Its MTA” on page 196**.

Configuring the Clustered WebAccess Agent for SSL

If you plan to enable SSL, as described in **“Securing WebAccess Agent Connections with SSL”** in **“WebAccess”** in the *GroupWise 7 Administration Guide*, you need to make the SSL certification file and key file available to the WebAccess Agent in the cluster. The default locations for the SSL certificate file and key file are on the cluster nodes along with the WebAccess Agent software, rather than being located with the domain on a GroupWise partition. To avoid having multiple copies of these files in multiple locations, you should set the locations in ConsoleOne®.

- 1 On the WebAccess Agent 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 7 Administration Guide*.
- 3 In ConsoleOne, browse to and select the Domain object.
- 4 Right-click the WebAccess Agent 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 After you have set these locations to the WebAccess Agent partition, continue with [Configuring the WebAccess Agent Cluster Resource to Load and Unload the WebAccess Agent and Its MTA](#).

Configuring the WebAccess Agent Cluster Resource to Load and Unload the WebAccess Agent and Its MTA

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

- ♦ [“Modifying the Cluster Resource Load Script for the WebAccess Agent and Its MTA” on page 196](#)
- ♦ [“Modifying the Cluster Resource Unload Script for the WebAccess Agent and Its MTA” on page 198](#)
- ♦ [“Setting the Failover List and Policies for the WebAccess Agent and Its MTA” on page 199](#)

Modifying the Cluster Resource Load Script for the WebAccess Agent and Its MTA

The cluster resource load script executes whenever the WebAccess 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 WebAccess Agent cluster resource is located.
- 3 Click the Cluster object to display the cluster resources that belong to the cluster.
- 4 Select the WebAccess Agent cluster resource that you created when you set up the WebAccess Agent partition, then click *Properties*.

The default load script from a generic IP service template appears as follows:

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

# mount the file system
exit_on_error mount -t reiserfs /dev/evms/vol /mnt/generic

# add the IP address
exit_on_error add_secondary_ipaddress a.b.c.d

# start the service
exit_on_error /etc/init.d/myservice start

# return status
exit 0
```

- 5 If this is an NSS volume or a shared pool, make the following changes to set up the WebAccess Agent load script:
 - 5a As needed, in the `mount` command, change `reiserfs` to whatever file system type is in use on nodes in the cluster.

- 5b** In the mount command, change `vol` to the actual device name in use on nodes in the cluster.
- 5c** In the mount command, change `/mnt/generic` to the mount point directory in use on nodes in the cluster.
- 5d** In the `add_secondary_ipaddress` command, change `a.b.c.d` to the secondary IP address of the WebAccess Agent partition ([WebAccess Agent Clustering Worksheet item 1](#))
- 5e** In the start service command, change `myservice start` to the command to start the WebAccess Agent and its MTA.

```
/etc/init.d/grpwise start domain
/etc/init.d/grpwise start webac70a
```

6 If this is a traditional Linux volume, use the following load script:

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

# define the IP address
RESOURCE_IP=123.123.1.

# define the file system type
MOUNT_FS=reiserf

# define the device
exit_on_error evms -f /var/opt/novell/ncs/ContainerActivate -rl
                                                    Share `uname -n`
MOUNT_DEV=/dev/evms/Share/dat

# define the mount point
MOUNT_POINT=/mnt/mount_point

# mount the file system
exit_on_error mount -t $MOUNT_FS $MOUNT_DEV $MOUNT_POINT

# add the IP address
exit_on_error add_secondary_ipaddress $RESOURCE_IP

/etc/init.d/grpwise start domain
/etc/init.d/grpwise start domain.webac70a

exit 0
```

Make the following changes to set up the load script for the agents:

- 6a** On the `RESOURCE_IP` line, change `123.123.1.1` to the secondary IP address of the GroupWise partition ([WebAccess Agent Clustering Worksheet item 1](#))
- 6b** As needed on the `MOUNT_FS` line, change `reiserfs` to whatever file system type is in use on nodes in the cluster.
- 6c** On the `MOUNT_DEV` line, change `/dev/evms/Share/dat` to the actual device name in use on nodes in the cluster.
- 6d** On the `MOUNT_POINT` line, change `/mnt/mount_point` to the mount point directory in use on nodes in the cluster.
- 6e** Use the following command to start the MTA:

```
/etc/init.d/grpwise start domain
```

6f Use the following command to start the WebAccess Agent:

```
/etc/init.d/grpwise start domain.webac70a.
```

7 Click *OK* to save the load script.

Modifying the Cluster Resource Unload Script for the WebAccess Agent and Its MTA

The cluster resource unload script executes whenever the WebAccess Agent 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 Cluster Resource Properties page of the WebAccess Agent cluster resource, click *Scripts > Unload Script*.

The default unload script appears as follows:

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

# request service stop
ignore_error /etc/init.d/myervice stop

# stop service otherwise
sleep 8
ignore_error fuser -k /mnt/generic

# del the IP address
ignore_error del_secondary_ipaddress a.b.c.d

# umount the file system
exit_on_error umount /mnt/generic

# return status
exit 0
```

2 If this is an NSS volume or a shared pool, make the following changes:

2a In the stop service command, change `myervice stop` to the command to stop the WebAccess Agent and its MTA.

```
/etc/init.d/grpwise stop webac70a
/etc/init.d/grpwise stop domain
```

2b Adjust the `sleep` command as needed so that the WebAccess Agent and its MTA can shut down normally on your system without being inadvertently killed by the `kill` command.

2c In the `kill service` command (used if the WebAccess Agent and its MTA do not stop normally), change `-k` to `-mk`.

The `-m` parameter obtains the PID numbers of the processes to kill.

2d In the `kill service` command, change `/mnt/generic` to the mount point directory used in the load script.

- 2e** In the `del_secondary_ipaddress` command, change `a.b.c.d` to the secondary IP address used in the load script.
 - 2f** In the `umount` command, change `/mnt/generic` to the mount point directory used in the load script.
- 3** If this is a traditional Linux volume, use the following unload script:

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

/etc/init.d/grpwise stop domain.webac70a
/etc/init.d/grpwise stop domain

# define the IP address
RESOURCE_IP=172.16.5.18

# define the mount point
MOUNT_POINT=/mnt/mount_point

sleep 8

ignore_error fuser -k $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
```

Make the following changes to set up the unload script for the agents.

- 3a** Use the following command to stop the WebAccess Agent:

```
/etc/init.d/grpwise stop domain.webac70a
```

- 3b** Use the following command to stop the MTA:

```
/etc/init.d/grpwise stop domain
```

- 3c** On the `RESOURCE_IP` line, change `172.16.5.18` to the secondary IP address used in the load script.
- 3d** On the `MOUNT_POINT` line, change `/mnt/mount_point` to the mount point directory used in the load script.
- 3e** Adjust the `sleep` command as needed so that the agents can shut down normally on your system without being inadvertently killed. by the `fuser -k` command that follows.
- 3f** In the `fuser -k` command (used if the agents do not stop normally), change `-k` to `-mk`. The `-m` parameter obtains the PID numbers of the processes to kill.

- 4** Click *OK* to save the unload script.

Setting the Failover List and Policies for the WebAccess Agent and Its MTA

- 1** On the Cluster Resource Properties page of the WebAccess 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 “[Setting Start, Failover, and Failback Modes](#)” in “[Installation and Setup](#)” in the *OES Novell Cluster Services 1.8.2 Administration Guide for Linux*.

- 2 Under *Preferred Nodes*, arrange the nodes in the cluster into the desired failover list for the WebAccess Agent ([WebAccess Agent Clustering Worksheet item 3](#)).
- 3 Click *OK*.

Verifying WebAccess Agent Object Properties

During installation of the WebAccess Agent, the WebAccess Agent 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 WebAccess Agent.

- ♦ “[Accessing WebAccess Agent Object Properties](#)” on page 200
- ♦ “[Verifying Post Office Links](#)” on page 200
- ♦ “[Forcing Use of the Web Access Agent Secondary IP Address](#)” on page 200

Accessing WebAccess Agent Object Properties

- 1 In ConsoleOne, browse to and select the WebAccess Agent domain in order to display its contents.
- 2 Right-click the WebAccess Agent object, then click *Properties*.
- 3 Continue with [Verifying Post Office Links](#).

Verifying Post Office Links

In the WebAccess object properties pages:

- 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 WebAccess Agent.
- 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 Web Access Agent Secondary IP Address](#).

Forcing Use of the Web Access Agent Secondary IP Address

If you want the WebAccess Agent to always use its secondary IP address, rather than using the primary IP address:

- 1 Click *GroupWise > Network Address*.

- 2 In the *TCP/IP Address* field, provide the secondary IP address ([WebAccess Agent Clustering Worksheet item 1](#)) for the Internet Agent.
- 3 Select *Bind Exclusively to TCP/IP Address*.
- 4 Click *OK*.
- 5 Continue with [Testing the WebAccess Agent in a Linux Cluster](#).

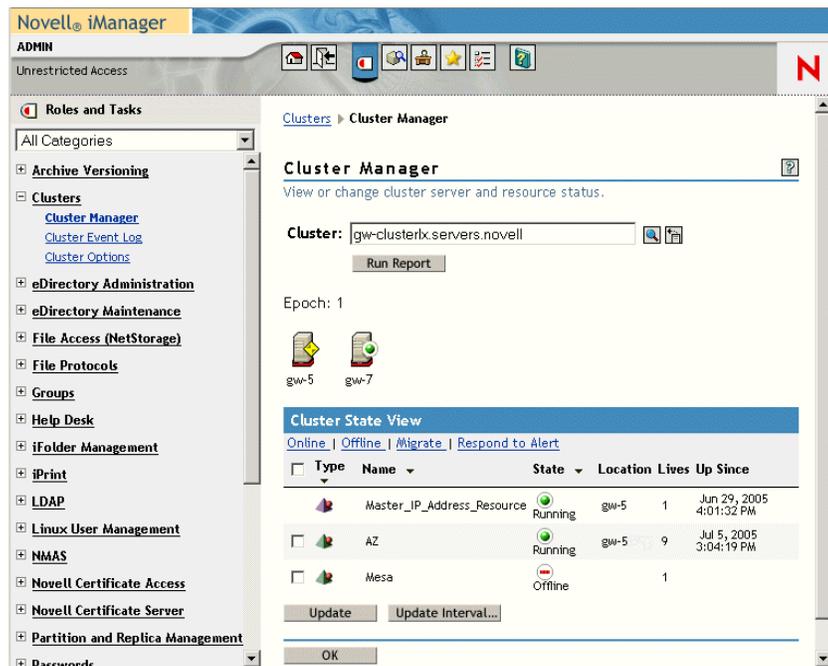
16.3.4 Installing and Configuring the WebAccess Application in a Cluster

If you have clustered your Web server, you must install the WebAccess Application on each node where the Web server is installed, as described in “[Installing the WebAccess Application and WebPublisher Application](#)” in “[Installing GroupWise WebAccess](#)” in the *GroupWise 7 Installation Guide*.

16.4 Testing the WebAccess Agent in a Linux Cluster

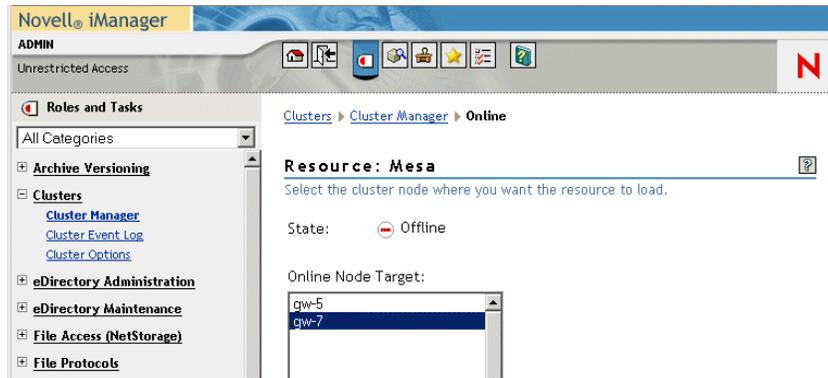
After you have configured the WebAccess Agent cluster resource, you can test the load and unload scripts by bringing the cluster resource online and taking it offline again.

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



The new WebAccess Agent cluster resource shows *Offline* in the *State* column.

- 2 Click the new WebAccess Agent cluster resource, then click *Online*.



- 3 Select the cluster node where you want to online the WebAccess Agent cluster resource, then click *OK*.

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

- 4 At the server where the WebAccess Agent is starting, use the following command to see that the Internet Agent has started:

```
/etc/init.d/grpwise status webac70a
```

- 5 Select the new WebAccess Agent cluster resource, then click *Offline*.

The *State* column for the WebAccess Agent cluster resource returns to *Offline*.

- 6 Use the same command that you used in [Step 4](#) to verify that the WebAccess Agent has stopped.
- 7 Repeat [Step 2](#) whenever you are ready to bring the new Internet Agent cluster resource online permanently.
- 8 Continue with [Managing the WebAccess Agent in a Linux Cluster](#).

16.5 Managing the WebAccess Agent in a Linux Cluster

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

- ♦ [Section 16.5.1, “Updating GroupWise Objects with Cluster-Specific Descriptions,”](#) on page 203
- ♦ [Section 16.5.2, “Knowing What to Expect in a WebAccess Agent Failover Situation,”](#) on page 204

16.5.1 Updating GroupWise Objects with Cluster-Specific Descriptions

After installing the WebAccess Agent 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 WebAccess Agent Domain and Its MTA” on page 203](#)
- ♦ [“Recording Cluster-Specific Information about the WebAccess Agent” on page 203](#)

Recording Cluster-Specific Information about the WebAccess Agent Domain and Its MTA

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

- 1 In ConsoleOne, browse to and right-click the Domain object, then click *Properties*.
- 2 In the *Description* field of the WebAccess Agent domain Identification page, provide a cluster-specific description of the WebAccess Agent domain, including the secondary IP address of its **GroupWise partition**.
- 3 Click *OK* to save the WebAccess Agent domain description.
- 4 Select the Internet Agent 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 WebAccess Agent](#).

Recording Cluster-Specific Information about the WebAccess Agent

With the contents of the WebAccess Agent domain still displayed:

- 1 Right-click the WebAccess Agent 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 WebAccess Agent domain is located.
This information appears on the WebAccess Agent console, no matter which node in the cluster it is currently running on.
- 4 Click *OK* to save the WebAccess Agent information.
- 5 Continue with [Knowing What to Expect in an Internet Agent Failover Situation](#).

16.5.2 Knowing What to Expect in a WebAccess Agent Failover Situation

The failover behavior of the MTA for the WebAccess Agent domain is the same as for an MTA in a regular domain. See [Section 14.6.2, “Knowing What to Expect in MTA and POA Failover Situations,”](#) on page 163.

When the WebAccess Agent fails over, the WebAccess client user sees the following message:

```
Unable to communicate with the GroupWise WebAccess Agent
```

The user just needs to be patient. When the WebAccess Agent comes up on the next node, the user can continue working without logging in again.

If the POA for the user’s post office fails over, the WebAccess client becomes unresponsive, but there is no error message. Again, the user should be patient until the POA comes up on the next node. Then the user can continue working without logging in again.

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

16.6 WebAccess Agent Clustering Worksheet

Item	Explanation
1) GroupWise Partition for the WebAccess Agent:	Specify the GroupWise partition where the WebAccess Agent domain will be created, along with its secondary IP address.
Secondary IP Address:	For more information, see Section 16.2.2, “Selecting the WebAccess Agent Partition and Secondary IP Address,” on page 189.
2) WebAccess Agent Domain Name:	Specify a unique name for the Internet Agent domain. Specify the directory on the GroupWise partition where you want to create the new domain.
Domain Database Location:	For more information, see Section 16.2.1, “Planning a Domain for the WebAccess Agent,” on page 188.
3) WebAccess Agent Failover List:	List other nodes in the cluster where the WebAccess Agent and its MTA could fail over. For more information, see Section 16.2.3, “Determining an Appropriate Failover List for the WebAccess Agent,” on page 189.
4) Cluster Resource Mount Point:	Specify the mount point directory where the WebAccess Agent domain will be mounted. For more information, see Section 16.2.4, “Determining Cluster Resource Information for the WebAccess Agent,” on page 189.

Item	Explanation
5) MTA Network Information:	Record the MTA network address information that you will need as you install the MTA.
<ul style="list-style-type: none"> ◆ MTA IP address ◆ MTA message transfer port ◆ MTA live remote port ◆ MTA HTTP port 	For more information, see Section 16.2.5, “Planning the MTA Installation,” on page 190.
6) WebAccess Agent Network Information:	Record the WebAccess Agent network address information that you will need to install the WebAccess Agent.
<ul style="list-style-type: none"> ◆ WebAccess Agent IP address ◆ WebAccess Agent HTTP port 	For more information, see Section 16.2.6, “Planning the WebAccess Agent Installation,” on page 190.

16.7 WebAccess Agent Quick Checklist

- Plan the new clustered WebAccess Agent, including the new domain required to house the WebAccess Agent in a clustering environment.
See [Section 16.2, “Planning the WebAccess Agent in a Linux Cluster,”](#) on page 187.
- Create the new WebAccess Agent domain.
See [Section 16.3.1, “Creating a Domain for the WebAccess Agent,”](#) on page 191.
- Set up the MTA for the new WebAccess Agent domain.
See [Section 16.3.2, “Installing the MTA for the WebAccess Agent Domain,”](#) on page 191.
- Install the WebAccess Agent.
See [Section 16.3.3, “Installing and Configuring the WebAccess Agent in a Cluster,”](#) on page 191.
- Modify the WebAccess Agent cluster resource load script.
See [“Modifying the Cluster Resource Load Script for the WebAccess Agent and Its MTA”](#) on page 196.
- Modify the WebAccess Agent cluster resource unload script.
See [“Modifying the Cluster Resource Unload Script for the WebAccess Agent and Its MTA”](#) on page 198.
- Set up the WebAccess Agent failover list and policies.
See [“Setting the Failover List and Policies for the WebAccess Agent and Its MTA”](#) on page 199.
- Double-check the cluster-specific WebAccess Agent object properties.
See [“Verifying WebAccess Agent Object Properties”](#) on page 200.
- Test the clustered WebAccess Agent.
See [Section 16.4, “Testing the WebAccess Agent in a Linux Cluster,”](#) on page 201.
- Record cluster-specific information in the properties pages of the GroupWise objects associated with the WebAccess Agent.

See [Section 16.5.1, “Updating GroupWise Objects with Cluster-Specific Descriptions,”](#) on [page 203](#).

Implementing GroupWise Monitor in a Linux Cluster

17

You should already have set up at least a basic GroupWise® system, as described in [Chapter 13](#), “[Planning GroupWise in a Linux Cluster](#),” on page 133 and [Chapter 14](#), “[Setting Up a Domain and a Post Office in a Linux Cluster](#),” on page 143. As part of this process, [Section 13.7.1](#), “[System Clustering Worksheet](#),” on page 140 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.

- ◆ [Section 17.1](#), “[Understanding the Monitor Components](#),” on page 207
- ◆ [Section 17.2](#), “[Planning GroupWise Monitor in a Linux Cluster](#),” on page 208
- ◆ [Section 17.3](#), “[Setting Up GroupWise Monitor in a Linux Cluster](#),” on page 210
- ◆ [Section 17.4](#), “[Testing the Monitor Agent in a Linux Cluster](#),” on page 217
- ◆ [Section 17.5](#), “[Managing the Monitor Agent in a Linux Cluster](#),” on page 218
- ◆ [Section 17.6](#), “[Monitor Agent Clustering Worksheet](#),” on page 218
- ◆ [Section 17.7](#), “[Monitor Agent Quick Checklist](#),” on page 219

17.1 Understanding the Monitor Components

If you are not familiar with GroupWise Monitor, review “[GroupWise Monitor Overview](#)” in “[Installing GroupWise Monitor](#)” in the *GroupWise 7 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 URLs are not available if the Web server is down:

Table 17-1 Monitor Web Console URLs

NetWare or Windows Web Server:	<code>http://Web_server_address/gw/gwmonitor</code>
--------------------------------------	---

Linux Web Server:	<code>http://Web_server_address/gwmon/gwmonitor</code>
----------------------	--

17.2 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.

Section 17.6, “Monitor Agent Clustering Worksheet,” on page 218 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 17.2.1, “Selecting a Domain for Access during Monitor Agent Installation,” on page 208
- ♦ Section 17.2.2, “Selecting an MTA for the Monitor Agent to Access after Installation,” on page 208
- ♦ Section 17.2.3, “Selecting the Monitor Agent Partition and Secondary IP Address,” on page 209
- ♦ Section 17.2.4, “Determining an Appropriate Failover List for the Monitor Agent,” on page 209
- ♦ Section 17.2.5, “Determining Cluster Resource Information for the Monitor Agent,” on page 209
- ♦ Section 17.2.6, “Planning the Monitor Agent Installation,” on page 210

17.2.1 Selecting a Domain for Access during 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 Internet Agent, as described in Section 15.2.1, “Creating a Domain for the Internet Agent,” on page 171, 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.

17.2.2 Selecting an MTA for the 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.

17.2.3 Selecting the Monitor Agent Partition and Secondary IP Address

As with the MTA and the POA, 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.

17.2.4 Determining an Appropriate Failover List for the 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 would be 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 13.6.2, "Determining Appropriate Failover Lists for the Agents," on page 138**, which describes the issues in the context of planning MTA and POA installations.

MONITOR CLUSTERING WORKSHEET

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

17.2.5 Determining Cluster Resource Information for the Monitor Agent

A cluster resource is a **shared partition**, secondary IP address, application, service, Web server, etc., that can function successfully anywhere in the cluster. Cluster resources include the GroupWise agents and the Messenger agents. When using 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.

17.2.6 Planning the 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 the Linux Monitor Agent](#)” in “[Installing GroupWise Monitor](#)” in the *GroupWise 7 Installation Guide*. Use the “[GroupWise Monitor Installation Worksheet](#)” 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 17.3, “Setting Up GroupWise Monitor in a Linux Cluster,”](#) on page 210.

17.3 Setting Up GroupWise Monitor in a Linux Cluster

GroupWise Monitor depends on a Web server, Apache in particular 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.

- ♦ [Section 17.3.1, “Installing and Configuring the Monitor Agent on Each Node in Your Cluster,”](#) on page 210
- ♦ [Section 17.3.2, “Configuring the Monitor Agent Cluster Resource to Load and Unload the Monitor Agent,”](#) on page 213

17.3.1 Installing and Configuring the 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.

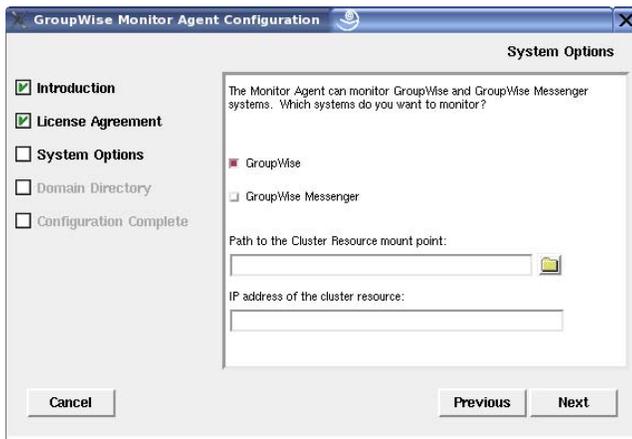
- ♦ [“Running the Monitor Installation Program on the Preferred Node”](#) on page 210
- ♦ [“Running the Monitor Agent Installation Program on Subsequent Nodes”](#) on page 211
- ♦ [“Configuring the Monitor Agent Web Console for SSL”](#) on page 212
- ♦ [“Testing the Monitor Agent Installation on Each Node”](#) on page 212

Running the 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 14.1, “Setting Up a New GroupWise System in a Linux Cluster,”](#) on page 143.
- 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 the Monitor Agent software, following the steps provided in “[Installing the Linux Monitor Agent](#)” in “[Installing GroupWise Monitor](#)” in the *GroupWise 7 Installation Guide*.
- 5 Configure the Monitor Agent, following the steps provided in “[Configuring the Linux Monitor Agent](#)” in “[Installing GroupWise Monitor](#)” in the *GroupWise 7 Installation Guide*, paying 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.
- 6 Continue with [Running the Monitor Agent Installation Program on Subsequent Nodes](#)

Running the 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 14.1](#), “[Setting Up a New GroupWise System in a Linux Cluster](#),” on page 143, 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 Monitor Agent Web Console for SSL](#).

Configuring the 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 Monitor Agent Installation on Each Node](#).

Testing the 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 7 Installation Guide*.

```
/etc/initd/grpwise-ma start
/etc/initd/grpwise-ma status
```

- 2 Then stop the Monitor Agent.

```
/etc/initd/grpwise-ma stop
/etc/initd/grpwise-ma status
```

- 3 Return to “[Running the Monitor Installation Program on the Preferred Node](#)” on page 210 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 fail over list, continue with [Configuring the Monitor Agent Cluster Resource to Load and Unload the Monitor Agent](#).

17.3.2 Configuring the Monitor Agent Cluster Resource to Load and Unload the 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 Monitor Agent](#)” on page 213
- ♦ “[Modifying the Cluster Resource Unload Script for the Monitor Agent](#)” on page 215
- ♦ “[Setting the Failover List and Policies for the Monitor Agent](#)” on page 216

Modifying the Cluster Resource Load Script for the 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 *Properties*.

The default load script from a generic IP service template appears as follows:

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

# mount the file system
exit_on_error mount -t reiserfs /dev/evms/vol /mnt/generic

# add the IP address
exit_on_error add_secondary_ipaddress a.b.c.d

# start the service
exit_on_error /etc/init.d/myservice start

# return status
exit 0
```

- 5 If this is an NSS volume or a shared pool, make the following changes to set up the Monitor Agent load script:
 - 5a As needed, in the `mount` command, change `reiserfs` to whatever file system type is in use on nodes in the cluster.
 - 5b In the `mount` command, change `vol` to the actual device name of the device in use on nodes in the cluster.
 - 5c In the `mount` command, change `/mnt/generic` to the mount point directory in use on nodes in the cluster.

- 5d** In the `add_secondary_ipaddress` command, change `a.b.c.d` to the secondary IP address of the Monitor Agent cluster resource (**Monitor Agent Clustering Worksheet item 1**)
- 5e** In the start service command, change `myservice start` to the command to start the Monitor Agent.

```
/etc/init.d/grpwise-ma start
```

- 6** If this is a traditional Linux volume, use the following load script:

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

# define the IP address
RESOURCE_IP=123.123.1.

# define the file system type
MOUNT_FS=reiserf

# define the device
exit_on_error evms -f /var/opt/novell/ncs/ContainerActivate -rl
                                                    Share `uname -n`
MOUNT_DEV=/dev/evms/Share/dat

# define the mount point
MOUNT_POINT=/mnt/mount_point

# mount the file system
exit_on_error mount -t $MOUNT_FS $MOUNT_DEV $MOUNT_POINT

# add the IP address
exit_on_error add_secondary_ipaddress $RESOURCE_IP

/etc/init.d/grpwise-ma start

exit 0
```

Make the following changes to set up the load script for the Monitor Agent:

- 6a** On the `RESOURCE_IP` line, change `123.123.1.1` to the secondary IP address of the GroupWise partition (**Monitor Agent Clustering Worksheet item 1**).
- 6b** As needed, on the `MOUNT_FS` line, change `reiserfs` to whatever file system type is in use on nodes in the cluster.
- 6c** On the `MOUNT_DEV` line, change `/dev/evms/Share/dat` to the actual device name in use on nodes in the cluster.
- 6d** On the `MOUNT_POINT` line, change `/mnt/mount_point` to the mount point directory in use on nodes in the cluster.
- 6e** Use the following command to start the Monitor Agent:

```
/etc/init.d/grpwise-ma start
```

- 7** Click *OK* to save the load script.

Modifying the Cluster Resource Unload Script for the Monitor Agent

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

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

The default unload script appears as follows:

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

# request service stop
ignore_error /etc/init.d/myservice stop

# stop service otherwise
sleep 8
ignore_error fuser -k /mnt/generic

# del the IP address
ignore_error del_secondary_ipaddress a.b.c.d

# umount the file system
exit_on_error umount /mnt/generic

# return status
exit 0
```

- 2 If this is an NSS volume or a shared pool, make the following changes:

- 2a In the stop service command, change `myservice stop` to the command to stop the Monitor Agent.

```
/etc/init.d/grpwise-ma stop
```

- 2b Adjust the `sleep` command as needed so that the Monitor Agent can shut down normally on your system without being inadvertently killed by the `fuser -k` command that follows.

- 2c In the kill service command (used if the Monitor Agent does not stop normally), change `-k` to `-mk`.

The `-m` parameter obtains the PID number of the process to kill.

- 2d In the kill service command, change `/mnt/generic` to the mount point directory used in the load script.

- 2e In the `del_secondary_ipaddress` command, change `a.b.c.d` to the secondary IP address used in the load script.

- 2f In the `umount` command, change `/mnt/generic` to the mount point directory used in the load script.

- 3 If this is a traditional Linux volume, use the following unload script:

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

/etc/init.d/grpwise-ma stop

# define the IP address
```

```

RESOURCE_IP=172.16.5.18

# define the mount point
MOUNT_POINT=/mnt/mount_point

sleep 8

ignore_error fuser -k $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

```

Make the following changes to set up the unload script for the Monitor Agent.

3a Use the following command to stop the Monitor Agent:

```
/etc/init.d/grpwise-ma stop
```

3b On the `RESOURCE_IP` line, change `172.16.5.18` to the secondary IP address used in the load script.

3c On the `MOUNT_POINT` line, change `/mnt/mount_point` to the mount point directory used in the load script.

3d Adjust the `sleep` command as needed so that the Monitor Agent can shut down normally on your system without being inadvertently killed by the `fuser -k` command that follows.

3e In the `fuser -k` command (used if the Monitor Agent does not stop normally), change `-k` to `-mk`.

The `-m` parameter obtains the PID numbers of the processes to kill.

4 Click *OK* to save the unload script.

Setting the Failover List and Policies for the Monitor Agent

1 On the 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 “[Setting Start, Failover, and Failback Modes](#)” in “[Installation and Setup](#)” in the *OES Novell Cluster Services 1.8.2 Administration Guide for Linux*.

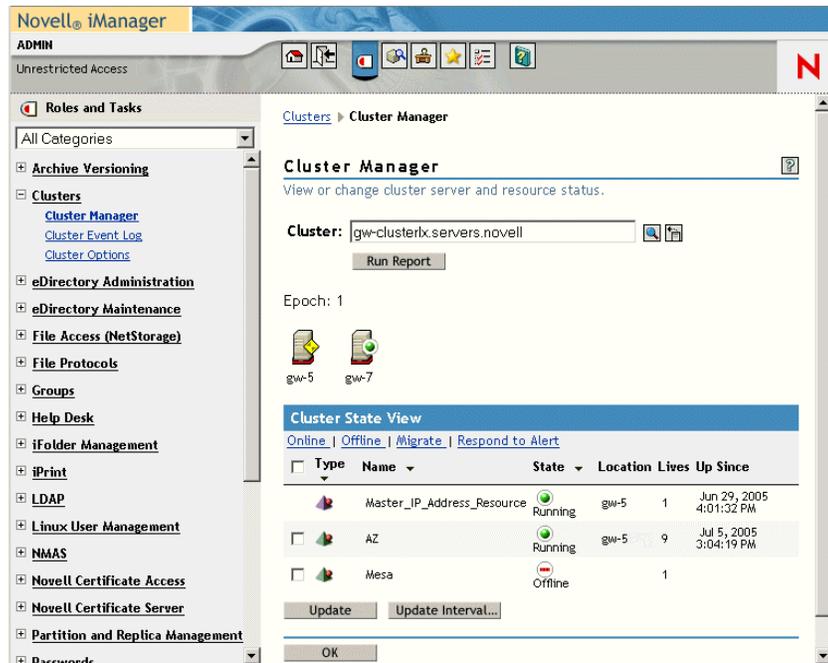
2 Under *Preferred Nodes*, arrange the nodes in the cluster into the desired failover list for the Monitor Agent ([Monitor Agent Clustering Worksheet item 4](#)).

3 Click *OK*.

17.4 Testing the Monitor Agent in a Linux Cluster

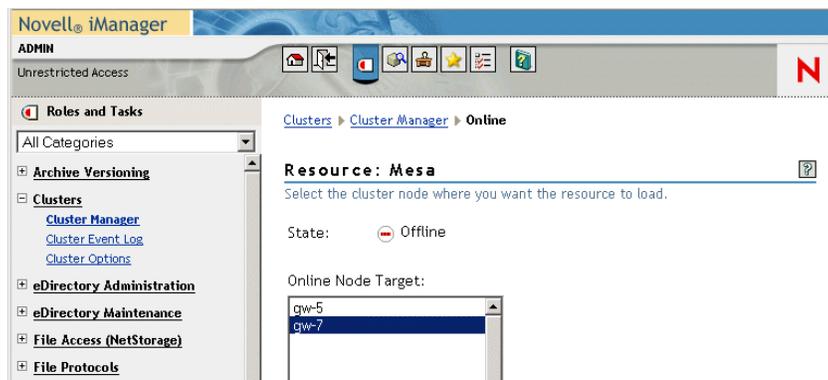
After you have configured the Monitor Agent cluster resource, you can test the load and unload scripts by bringing the Monitor Agent cluster resource online and taking it offline again.

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



The new Monitor Agent cluster resource shows *Offline* in the *State* column.

- 2 Click the new Monitor Agent cluster resource, then click *Online*.



- 3 Select the cluster node where you want to online the Monitor Agent cluster resource, then click *OK*.

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 returns to *Offline*.
- 6 Use the same command that you used in [Step 4](#) to verify that the Monitor Agent has stopped.
- 7 Repeat [Step 2](#) 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](#).

17.5 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 may 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.

17.6 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 17.2.3, "Selecting the Monitor Agent Partition and Secondary IP Address," on page 209.
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 17.2.1, "Selecting a Domain for Access during Monitor Agent Installation," on page 208.
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 17.2.2, "Selecting an MTA for the Monitor Agent to Access after Installation," on page 208
4) Monitor Agent Failover List:	List the nodes in the cluster where the Monitor Agent could fail over. For more information, see Section 17.2.4, "Determining an Appropriate Failover List for the Monitor Agent," on page 209.
5) Cluster Resource Information <ul style="list-style-type: none"> ◆ 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 17.2.5, "Determining Cluster Resource Information for the Monitor Agent," on page 209.

17.7 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 17.2, “Planning GroupWise Monitor in a Linux Cluster,”](#) on page 208.
- ❑ Install the Monitor Agent on all nodes in the Monitor Agent’s failover list.

See [Section 17.3.1, “Installing and Configuring the Monitor Agent on Each Node in Your Cluster,”](#) on page 210.
- ❑ Modify the Monitor Agent cluster resource load script.

See [“Modifying the Cluster Resource Load Script for the Monitor Agent”](#) on page 213.
- ❑ Modify the Monitor Agent cluster resource unload script.

See [“Modifying the Cluster Resource Unload Script for the Monitor Agent”](#) on page 215.
- ❑ Set up the Monitor Agent failover list and policies.

See [“Setting the Failover List and Policies for the Monitor Agent”](#) on page 216.
- ❑ Test the clustered Monitored Agent.

See [Section 17.4, “Testing the Monitor Agent in a Linux Cluster,”](#) on page 217.

Backing Up a GroupWise System in a Linux Cluster

18

To back up GroupWise® data in a Linux cluster, 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 7 Administration Guide*.

Updating a GroupWise System in a Linux Cluster

19

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 7 Installation Guide* to update your GroupWise system in a NetWare cluster.

Moving an Existing Linux GroupWise 7 System into a Linux Cluster

20

If you are adding the high availability benefits of Novell® Cluster Services™ to a GroupWise® 7 system that is already up and running, the first step is to install Novell Cluster Services following the instructions in *OES Novell Cluster Services 1.8.2 Administration Guide for Linux*. You should also review [Chapter 12, “Introduction to GroupWise 7 and Novell Cluster Services on Linux,” on page 131](#) 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 could transfer individual post offices where the needs for high availability are greatest. You could 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 13, “Planning GroupWise in a Linux Cluster,” on page 133](#). Fill out the [Section 13.7.1, “System Clustering Worksheet,” on page 140](#) 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 7 Administration Guide*.
- Review [Section 13.6, “Deciding How to Install and Configure the Linux Agents in a Cluster,” on page 138](#), fill out the [Section 13.7.2, “Agent Clustering Worksheet,” on page 141](#), and install the agents as needed for the first clustered domain and/or post office, following the instructions in [Section 14.4, “Installing and Configuring the MTA and the POA in a Cluster,” on page 146](#). 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 14.5, “Testing Your Clustered GroupWise System,” on page 160](#).
- Take care of the cluster management details described in [Section 14.6, “Managing Your Clustered GroupWise System,” on page 161](#).
- Move more domains and post offices into the cluster as needed. If you have GroupWise libraries, see [Section 13.5, “Planning a New Library for a Clustered Post Office,” on page 137](#).
- Add other components to your clustered GroupWise system as needed, following the instructions in:
 - ♦ [Chapter 15, “Implementing the Internet Agent in a Linux Cluster,” on page 167](#)
 - ♦ [Chapter 16, “Implementing WebAccess in a Linux Cluster,” on page 187](#).

- ♦ Chapter 17, “Implementing GroupWise Monitor in a Linux Cluster,” on page 207
- ♦ Chapter 18, “Backing Up a GroupWise System in a Linux Cluster,” on page 221

Moving a Clustered GroupWise 7 System from NetWare to Linux

21

For general information about moving from a NetWare® cluster to a Linux cluster, see “[Converting a NetWare Cluster to Linux](#)” in “[Installation and Setup](#)” in the *OES Novell Cluster Services 1.8.2 Administration Guide for Linux*. It is possible to have a cluster that includes both NetWare and Linux servers. Therefore, you can move your GroupWise® 7 system from NetWare servers to Linux servers one component at a time. However, all of the servers on the failover list of each GroupWise component must be of the same platform. For example, a domain and MTA on a NetWare server cannot fail over to a Linux server; they must fail over to another NetWare server.

Ideally, you should administer a GroupWise system on Linux using the Linux version of ConsoleOne®, which is available from [Novell Product Downloads site \(http://download.novell.com\)](http://download.novell.com).

The *GroupWise 7 Installation Guide* provides the following sections to help you move components of your GroupWise system to Linux:

- ♦ “[Transitioning GroupWise Administration to Linux](#)”
- ♦ “[Manually Migrating a Post Office and Its POA to Linux](#)”
- ♦ “[Manually Migrating a Domain and Its MTA to Linux](#)”
- ♦ “[Manually Migrating the Internet Agent to Linux](#)”
- ♦ “[Manually Migrating WebAccess to Linux](#)”
- ♦ “[Manually Migrating Monitor to Linux](#)”

GroupWise 6.5 cannot run in a cluster on Linux. Therefore, if you have a clustered GroupWise 6.5 system, you must update it to GroupWise 7 before you can move it into a Linux cluster.

Implementing Messenger in a Linux Cluster

22

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 13, “Planning GroupWise in a Linux Cluster,” on page 133](#) and [Chapter 14, “Setting Up a Domain and a Post Office in a Linux Cluster,” on page 143](#). As part of the process of setting up GroupWise in the cluster, you filled out the [“System Clustering Worksheet” on page 140](#). Some of the information from that worksheet is helpful as you implement Messenger in your cluster.

- ♦ [Section 22.1, “Planning Your Messenger System in a Linux Cluster,” on page 229](#)
- ♦ [Section 22.2, “Setting Up Your Messenger System in a Linux Cluster,” on page 231](#)
- ♦ [Section 22.3, “Testing Your Clustered Messenger System,” on page 239](#)
- ♦ [Section 22.4, “Managing Your Clustered Messenger System,” on page 240](#)
- ♦ [Section 22.5, “Messenger Clustering Worksheet,” on page 240](#)
- ♦ [Section 22.6, “Messenger Clustering Quick Checklist,” on page 241](#)

22.1 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 22.5, “Messenger Clustering Worksheet,” on page 240](#) 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:

- ♦ [Section 22.1.1, “Understanding Your Cluster,” on page 229](#)
- ♦ [Section 22.1.2, “Selecting the Messenger Partition and Secondary IP Address,” on page 230](#)
- ♦ [Section 22.1.3, “Determining an Appropriate Failover List for the Messenger Agents,” on page 230](#)
- ♦ [Section 22.1.4, “Determining Cluster Resource Information for the Messenger Agents,” on page 230](#)
- ♦ [Section 22.1.5, “Planning the Messenger Agent Installation,” on page 231](#)

22.1.1 Understanding Your Cluster

As described in [Section 13.1, “Installing Novell Cluster Services on Linux,” on page 134](#), 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 [Section 13.7.1, “System Clustering Worksheet,” on page 140](#).

22.1.2 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.

22.1.3 Determining an Appropriate Failover List for the 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.

22.1.4 Determining Cluster Resource Information for the Messenger Agents

A cluster resource is a **shared partition**, secondary IP address, application, service, Web server, etc., 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.

22.1.5 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 *Messenger 2.0 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](#).

22.2 Setting Up Your Messenger System in a Linux Cluster

You should have already reviewed [Section 22.1, “Planning Your Messenger System in a Linux Cluster,”](#) on page 229 and filled out the [Section 22.5, “Messenger Clustering Worksheet,”](#) on page 240 and the “[Novell Messenger Worksheet](#)” in the *Messenger 2.0 Installation Guide*.

- ♦ [Section 22.2.1, “Creating Your Messenger System and Installing the Messenger Agents,”](#) on page 231
- ♦ [Section 22.2.2, “Changing Messenger Paths to Locations on the Messenger Partition,”](#) on page 233
- ♦ [Section 22.2.3, “Configuring the Messenger Cluster Resource to Load and Unload the Messenger Agents,”](#) on page 235

22.2.1 Creating Your Messenger System and Installing the 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 Messenger Installation Program on the Preferred Node”](#) on page 231
- ♦ [“Running the Messenger Installation Program on Subsequent Nodes”](#) on page 232
- ♦ [“Setting Up Non-root Access on NSS Volumes on Each Node”](#) on page 232
- ♦ [“Testing Your Messenger Agent Installation on Each Node”](#) on page 232

Running the 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 From the *Messenger 2 Administrator for Linux* CD, 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 *Messenger 2.0 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 *Messenger 2.0 Installation Guide*.
- 7 Continue with [Running the Messenger Installation Program on Subsequent Nodes](#).

Running the 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 From the *Messenger 2 Administrator for Linux* CD, 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 “[Creating NSS Volumes](#)” in “[Installation and Setup](#)” in the *OES Novell Cluster Services 1.8.2 Administration Guide for Linux*, 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 *Messenger 2.0 Installation Guide*. Then continue with [Testing Your Messenger Agent Installation on Each Node](#).

Testing Your 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 *Messenger 2.0 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 Messenger Installation Program on the Preferred Node” on page 231 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](#).

22.2.2 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:

- ♦ “Setting the Store Path” on page 233
- ♦ “Setting the Messaging Agent Queue Path” on page 233
- ♦ “Setting the Archive Agent Queue Path” on page 234
- ♦ “Setting the Messaging Agent Log Path” on page 234
- ♦ “Setting the Archive Agent Log Path” on page 234

After settings these directories, continue with [Section 22.2.3, “Configuring the Messenger Cluster Resource to Load and Unload the Messenger Agents,”](#) on page 235.

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 *Messaging 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*.

22.2.3 Configuring the Messenger Cluster Resource to Load and Unload the 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 fallback situations are handled.

- ♦ “[Modifying the Cluster Load Script for the Messenger Agents](#)” on page 235
- ♦ “[Modifying the Cluster Resource Unload Script for the Messenger Agents](#)” on page 237
- ♦ “[Setting the Failover List and Policies for the Messenger Agents](#)” on page 238

Modifying the Cluster Load Script for the 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 *Properties*.

The default load script from a generic IP service template appears as follows:

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

# mount the file system
exit_on_error mount -t reiserfs /dev/evms/vol /mnt/generic

# add the IP address
exit_on_error add_secondary_ipaddress a.b.c.d

# start the service
exit_on_error /etc/init.d/myservice start

# return status
exit 0
```

- 5 If this is an NSS volume or a shared pool, make the following changes to set up the Messenger load script:
 - 5a As needed, in the `mount` command, change `reiserfs` to whatever file system type is in use on nodes in the cluster.
 - 5b In the `mount` command, change `vol` to the actual device name in use on nodes in the cluster.

- 5c** In the mount command, change `/mnt/generic` to the mount point directory in use on nodes in the cluster.
- 5d** In the `add_secondary_ipaddress` command, change `a.b.c.d` to the secondary IP address of the Messenger partition.
- 5e** In the start service command, change `myservice start` to the command to start the Messenger agents.

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

- 6** If this is a traditional Linux volume, use the following load script:

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

# define the IP address
RESOURCE_IP=123.123.1.

# define the file system type
MOUNT_FS=reiserf

# define the device
exit_on_error evms -f /var/opt/novell/ncs/ContainerActivate -rl
                                                    Share `uname -n`
MOUNT_DEV=/dev/evms/Share/dat

# define the mount point
MOUNT_POINT=/mnt/mount_point

# mount the file system
exit_on_error mount -t $MOUNT_FS $MOUNT_DEV $MOUNT_POINT

# add the IP address
exit_on_error add_secondary_ipaddress $RESOURCE_IP
/etc/init.d/novell-nmma start
/etc/init.d/novell-nmaa start

exit 0
```

Make the following changes to set up the load script for the Messenger agents:

- 6a** On the `RESOURCE_IP` line, change `123.123.1.1` to the secondary IP address of the GroupWise partition ([Messenger Clustering Worksheet item 6](#) and [Messenger Clustering Worksheet item 7](#))
- 6b** As needed on the `MOUNT_FS` line, change `reiserfs` to whatever file system type is in use on nodes in the cluster.
- 6c** On the `MOUNT_DEV` line, change `/dev/evms/Share/dat` to the actual device name in use on nodes in the cluster.
- 6d** On the `MOUNT_POINT` line, change `/mnt/mount_point` to the mount point directory in use on nodes in the cluster.
- 6e** Use the following command to start the Messaging Agent:

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

- 6f** Use the following command to start the Archive Agent:

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

- 7 Click *Apply* to save the load script.

Modifying the Cluster Resource Unload Script for the Messenger Agents

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

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

The default unload script appears as follows:

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

# request service stop
ignore_error /etc/init.d/myservice stop

# stop service otherwise
sleep 8
ignore_error fuser -k /mnt/generic

# del the IP address
ignore_error del_secondary_ipaddress a.b.c.d

# umount the file system
exit_on_error umount /mnt/generic

# return status
exit 0
```

- 2 If this is an NSS volume or a shared pool, make the following changes to the Messenger unload script:

- 2a In the stop service command, change `myservice stop` to the command to stop the Messenger agents.

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

- 2b Adjust the `sleep` command as needed so that the Messenger agents can shut down normally on your system without being inadvertently killed by the `fuser -k` command that follows.

- 2c In the kill service command, used if the Messenger agents do not stop normally, change `-k` to `-mk`.

The `-m` parameter obtains the PID numbers of the processes to kill.

- 2d In the kill service command, change `/mnt/generic` to the mount point directory used in the load script.

- 2e In the `del_secondary_ipaddress` command, change `a.b.c.d` to the secondary IP address used in the load script.

- 2f In the `umount` command, change `/mnt/generic` to the mount point directory used in the load script.

- 3 If this is a traditional Linux volume, use the following unload script:

```

#!/bin/bash
. /opt/novell/ncs/lib/ncsfncs

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

# define the IP address
RESOURCE_IP=172.16.5.18

# define the mount point
MOUNT_POINT=/mnt/mount_point

sleep 8

ignore_error fuser -k $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

```

Make the following changes to set up the unload script for the Messenger agents.

3a Use the following command to stop the Messaging Agent:

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

3b Use the following command to stop the Archive Agent:

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

3c On the `RESOURCE_IP` line, change `172.16.5.18` to the secondary IP address used in the load script.

3d On the `MOUNT_POINT` line, change `/mnt/mount_point` to the mount point directory used in the load script.

3e Adjust the `sleep` command as needed so that the Messenger agents can shut down normally on your system without being inadvertently killed by the `fuser -k` command that follows.

3f In the `fuser -k` command (used if the agents do not stop normally), change `-k` to `-mk`. The `-m` parameter obtains the PID numbers of the processes to kill.

4 Click *Apply* to save the unload script.

Setting the Failover List and Policies for the Messenger Agents

1 On the 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

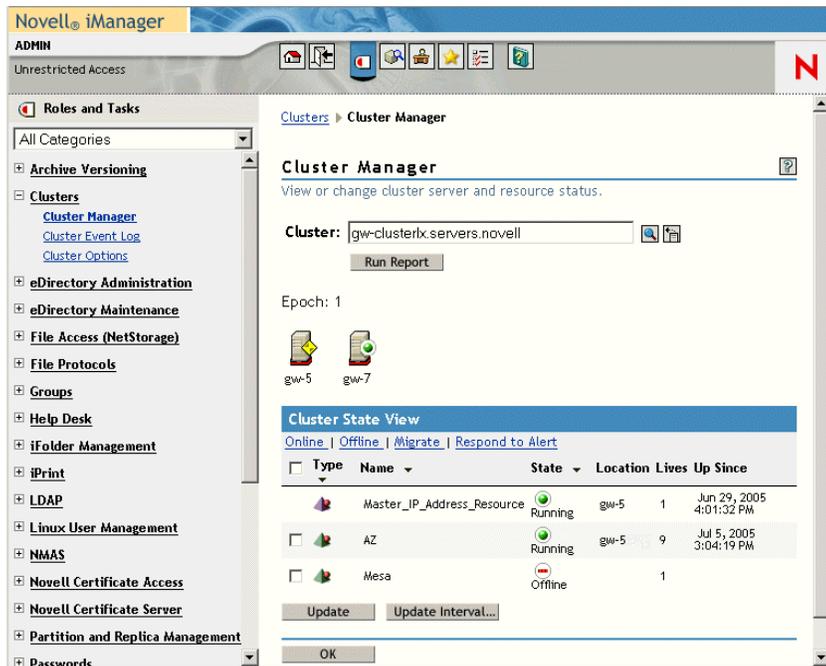
If you are considering changing these defaults, see “[Setting Start, Failover, and Failback Modes](#)” in “[Installation and Setup](#)” in the *OES Novell Cluster Services 1.8.2 Administration Guide for Linux*.

- 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*.

22.3 Testing Your Clustered Messenger System

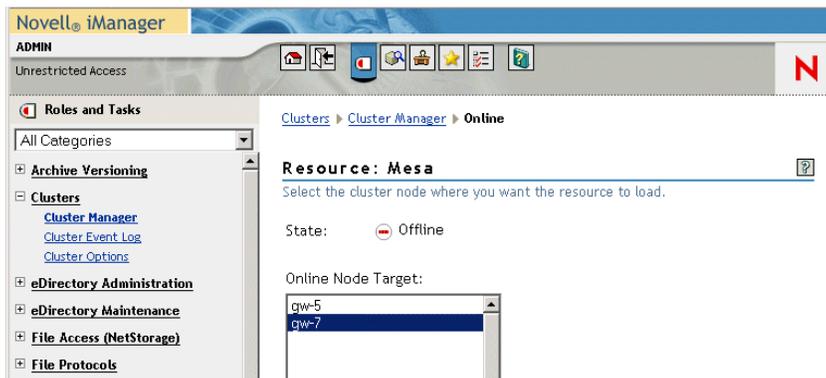
After you have configured the Messenger cluster resource, you can test the load and unload scripts by bringing the Messenger cluster resource online and taking it offline again.

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



The new Messenger cluster resource shows *Offline* in the *State* column.

- 2 Click the new Messenger cluster resource, then click *Online*.



- 3 Select the cluster node where you want to online the Messenger cluster resource, then click *OK*.
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 2](#) whenever you are ready to bring the new Messenger cluster resource online permanently.
- 8 Continue with [Managing Your Clustered Messenger System](#).

22.4 Managing Your Clustered Messenger System

If the node where your Messenger system is running goes down, it fails over to the next node in its fail over 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.

22.5 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 13.1, "Installing Novell Cluster Services on Linux," on page 134.
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 2.2, "Installing Novell Cluster Services," on page 24.
3) Cluster Context:	Record the full context where you created the Cluster object. For more information, see Section 13.1, "Installing Novell Cluster Services on Linux," on page 134.

Item	Explanation
4) Nodes in Cluster <ul style="list-style-type: none"> ◆ File system type ◆ Device name ◆ Mount point 	<p>List the nodes that are part of the cluster that will include Messenger. Also list technical information, including file system type (<code>reiserfs</code>, <code>ext3</code>, etc.), device name (<code>sda2</code>, <code>hda1</code>, etc.), and mount point directory (<code>/mnt</code>, <code>/mail</code>, etc.) in use on the nodes the cluster. You need this information as you create load and unload scripts for the Messenger agents.</p> <p>For more information, see Section 13.1, “Installing Novell Cluster Services on Linux,” on page 134.</p>
5) Shared Partitions in Cluster:	<p>List the shared partitions that are available for use in your Messenger system.</p> <p>For more information, see Section 13.1, “Installing Novell Cluster Services on Linux,” on page 134.</p>
6) Messenger Partition for Messaging Agent: Secondary IP address: Use Same Partition for Archive Agent? <ul style="list-style-type: none"> ◆ Yes ◆ No 	<p>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.</p> <p>For more information, see “Selecting the Messenger Volume” on page 121.</p>
7) Messenger Partition for Archive Agent? Secondary IP Address:	<p>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 could choose to use a separate shared partition for the Archive Agent.</p> <p>For more information, see “Selecting the Messenger Volume” on page 121.</p>
8) Failover List for Messenger Agents:	<p>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.</p> <p>For more information, see “Determining an Appropriate Failover Path for the Messenger Volume” on page 121.</p>
9) Mount Point for Shared Storage:	<p>Specify the mount point directory where the shared resource is mounted to the cluster node where the Messenger Agents run.</p> <p>For more information, see Section 22.1.4, “Determining Cluster Resource Information for the Messenger Agents,” on page 230.</p>

22.6 Messenger Clustering Quick Checklist

- Plan your clustered Messenger system.

See [Section 22.1, “Planning Your Messenger System in a Linux Cluster,” on page 229.](#)

- Create your Messenger system and install the Messenger agents.

See [Section 22.2.1, “Creating Your Messenger System and Installing the Messenger Agents,”](#) on page 231.

- ❑ 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 232.

- ❑ In ConsoleOne, change the locations of various Messenger files from their default locations on local nodes to the Messenger partition that is always available not matter what node the Messenger agents are running on.

See [Section 22.2.2, “Changing Messenger Paths to Locations on the Messenger Partition,”](#) on page 233.

- ❑ Modify the Messenger agents cluster resource load script.

See [“Modifying the Cluster Load Script for the Messenger Agents”](#) on page 235.

- ❑ Modify the Messenger agents cluster resource unload script.

See [“Modifying the Cluster Resource Unload Script for the Messenger Agents”](#) on page 237.

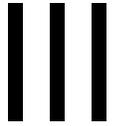
- ❑ Set up the Messenger agents failover list and policies.

See [“Setting the Failover List and Policies for the Messenger Agents”](#) on page 238.

- ❑ Test your clustered Messenger system.

See [Section 22.3, “Testing Your Clustered Messenger System,”](#) on page 239.

Novell Teaming and Conferencing



Before installing Novell® Teaming or the Conferencing component of Novell Teaming + Conferencing, you should thoroughly review the documentation provided at the [Novell Teaming and Conferencing documentation Web site \(http://www.novell.com/documentation/team_plus_conf\)](http://www.novell.com/documentation/team_plus_conf). These guides provide detailed product installation and configuration instructions, but they do not include specific instructions for integrating Novell Teaming or Conferencing with eDirectory™ or GroupWise®. This section of the *GroupWise 8 Interoperability Guide* supplies these product-specific instructions.

- ♦ [Chapter 23, “Using GroupWise with Novell Teaming,” on page 245](#)
- ♦ [Chapter 24, “Using GroupWise with Conferencing,” on page 263](#)
- ♦ [Chapter 25, “Streamlining Authentication to Teaming or Conferencing,” on page 265](#)

For background information about Novell Teaming and Novell Teaming + Conferencing, see the [Collaboration Tools and Organizational Success white paper \(http://www.novell.com/rc/docrepository/public/37/basedocument.2007-09-18.8712884885/4622072%20-%20FINAL_en.pdf\)](http://www.novell.com/rc/docrepository/public/37/basedocument.2007-09-18.8712884885/4622072%20-%20FINAL_en.pdf).

Using GroupWise with Novell Teaming

23

When you install Novell® Teaming together with eDirectory™ and GroupWise®, some configuration steps are required to integrate the applications.

- ♦ [Section 23.1, “Understanding How Novell Teaming Interacts with eDirectory and GroupWise,” on page 245](#)
- ♦ [Section 23.2, “Meeting Novell Teaming System Requirements on OES 2 or SLES 10 SP1 Servers,” on page 245](#)
- ♦ [Section 23.3, “Installing Novell Teaming with GroupWise,” on page 248](#)
- ♦ [Section 23.4, “Configuring LDAP Authentication,” on page 249](#)
- ♦ [Section 23.5, “Configuring GroupWise to Support Your Teaming Installation,” on page 249](#)
- ♦ [Section 23.6, “Receiving Notification of Teaming Site Activity in Your GroupWise Client,” on page 259](#)
- ♦ [Section 23.7, “Receiving Calendar Postings from the Teaming Site,” on page 259](#)
- ♦ [Section 23.8, “Adding Portlets to Novell Teaming,” on page 259](#)

23.1 Understanding How Novell Teaming Interacts with eDirectory and GroupWise

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

- ♦ Teaming can use eDirectory for LDAP authentication of Teaming users. This means that you do not need to create Teaming users manually. Teaming can create its user accounts based on the users that already exist in eDirectory.
- ♦ Teaming can use GroupWise as its integrated e-mail system. This means that e-mail messages sent from the Teaming site are delivered to GroupWise mailboxes. It also means that GroupWise users can post items to Teaming folders by sending e-mail messages to Teaming folders.
- ♦ Teaming information can be displayed in the GroupWise Windows client.

23.2 Meeting Novell Teaming System Requirements on OES 2 or SLES 10 SP1 Servers

“Prerequisites” in the *Novell Teaming 1.0 Installation Guide* lists the standard system requirements for Novell Teaming. If you are installing Teaming on Novell Open Enterprise (OES) 2 or SUSE Linux Enterprise Server (SLES) 10 SP1, a default operating system installation does not meet these system requirements. The following sections explain how to update an OES 2 or SLES 10 SP1 server so that it meets the Teaming system requirements:

- ♦ [Section 23.2.1, “Linux Open File Limit,” on page 246](#)

- ◆ Section 23.2.2, “Java Development Kit Requirement,” on page 246
- ◆ Section 23.2.3, “MySQL Setup,” on page 247
- ◆ Section 23.2.4, “Potential Web Server Port Conflicts,” on page 247

23.2.1 Linux Open File Limit

The default limit on the number of open files on SLES 10 is not sufficient for Novell Teaming.

To configure SLES 10 with an appropriate open file limit:

- 1 Edit the `/etc/security/limits.conf` file in an ASCII text editor.
- 2 Add the following lines to the bottom of the list, following the format of the example lines:

```
* hard nofile 65535
* soft nofile 4096
```

- 3 Save the file, then exit the text editor.

23.2.2 Java Development Kit Requirement

In “Prerequisites” in the *Novell Teaming 1.0 Installation Guide* states that Teaming requires Sun Java Development Kit (JDK) 1.5.0_11 or IBM JDK 1.5. A JDK is not provided with OES 2 or SLES 10 SP1. You must install a JDK on your Linux server before you install Teaming.

IMPORTANT: Make sure you download the JDK, not the Java Runtime Environment (JRE).

- 1 Go to the following URL:
[Java SE Downloads - Previous Release \(http://java.sun.com/javase/downloads/index_jdk5.jsp\)](http://java.sun.com/javase/downloads/index_jdk5.jsp)
The update you need is listed as “JDK 5.0 Update *nn*.”
- 2 Click *Download* next to this update.
- 3 In the *Platform* field, select Linux.
- 4 Select *I agree...*, then click *Continue* to accept Sun’s License Agreement.
- 5 Select *Linux RPM in Self-Extracting File*, then save the file to an empty temporary directory on your Linux server.
- 6 As the `root` user, change to that temporary directory, then use the following command to make sure that the download arrived safely:

```
ls -l
```

You should see a file named `jdk-1_5_0_nn-linux-i586-rpm.bin`.

- 7 Change the permissions on the file to include execute permissions:

```
chmod +x jdk-1_5_0_nn-linux-i586.bin
```

- 8 Run the self-extracting file:

```
./jdk-1_5_0_nn-linux-i586-rpm.bin
```

This creates a file named `jdk-1_5_0_nn-linux-i586.rpm`, and a directory named `/usr/java/jdk1.5.0_nn`.

The Sun JDK is now installed and configured on your Linux server. The Teaming installer asks for the JDK directory.

If you would prefer to use the IBM JDK, refer to the [IBM developerWorks \(http://www.ibm.com/developerworks/java/jdk/linux\)](http://www.ibm.com/developerworks/java/jdk/linux) Web site.

23.2.3 MySQL Setup

Although MySQL is installed by default on OES 2 and SLES 10 SP1, it is not configured with a password, nor is it configured to start automatically. Follow the steps below to set up MySQL for use with Novell Teaming:

- 1 In YaST, click *System > System Services*.
- 2 Select the MySQL service, then click *Enable*.
- 3 Click *Finish*, then exit YaST.
- 4 In a terminal window, become the `root` user.
- 5 To verify that MySQL has started, use the following command:

```
ps -eaf | grep mysql
```

You should see MySQL processes running.

- 6 Set the MySQL root password:

```
mysqladmin -root password root
```

This sets the password for the MySQL root user to `root`, which matches the default MySQL configuration in the Teaming installer.

The following are some basic and useful MySQL commands:

Action	Command
Stop MySQL	<code>/etc/initd/mysql stop</code>
Start MySQL	<code>/etc/initd/mysql start</code>
Show MySQL status	<code>mysqladmin -r root -p extended-status</code>

If you would like to administer MySQL using a GUI interface, you can download tools from:

[MySQL GUI Tools Downloads \(http://dev.mysql.com/downloads/gui-tools/5.0.html\)](http://dev.mysql.com/downloads/gui-tools/5.0.html)

For more information about MySQL, see:

[MySQL Documentation \(http://dev.mysql.com/doc\)](http://dev.mysql.com/doc)

23.2.4 Potential Web Server Port Conflicts

Novell Teaming includes its own Web server. If the Linux server where you plan to install Novell Teaming already has a Web server on it, make sure that you specify different, unique port numbers when you install Novell Teaming or remove the existing Web server from the Teaming server. Teaming does not start successfully if another Web server is already running on the Teaming server using the same port numbers.

23.3 Installing Novell Teaming with GroupWise

Complete instructions for installing Novell Teaming are provided in the *Novell Teaming 1.0 Installation Guide*. The sections below detail the GroupWise-specific information you need as you install Novell Teaming.

- ♦ [Section 23.3.1, “Preparing for the Novell Teaming Installation,” on page 248](#)
- ♦ [Section 23.3.2, “Performing the Novell Teaming Installation,” on page 248](#)

23.3.1 Preparing for the Novell Teaming Installation

Before you install Novell Teaming with eDirectory and GroupWise, you need to gather the following system and GroupWise information:

- ♦ The DNS hostname of the server where you plan to install Teaming.
- ♦ The DNS hostname of the server where the GroupWise Internet Agent is running. The Internet Agent functions as an SMTP host to exchange e-mail messages between Teaming site users and GroupWise mailboxes.
- ♦ The Internet domain name of your company or organization.
- ♦ A GroupWise username and password for a GroupWise account that represents the Teaming Site Manager account (by default, username `admin`, password `admin`) associated with the Teaming site. For example, you might plan to create a GroupWise user named `TeamingAdmin` for this purpose. This user does not need any administrative rights. It is used as the administrative e-mail address from which the Teaming site sends messages to GroupWise mailboxes. It is also used to own the GroupWise resources associated with Teaming folder where you want Teaming users to be able to post items from GroupWise

NOTE: When you enable LDAP authentication and synchronization, if your eDirectory Admin User object is within the search scope for Teaming users, the Teaming `admin` user becomes synchronized to your eDirectory Admin user. Therefore, it simplifies the situation to set the password for the Teaming `admin` user to your eDirectory Admin password.

- ♦ A GroupWise resource for each Teaming folder where GroupWise users need to be able to post items by sending e-mail messages. For example, if a sales team has a Discussion folder, you could create a resource representing the Discussion folder, such as a resource named `salesdisc`, for example. Then, when sales people want to post items to the Discussion folder, they can mail a posting to an e-mail address similar to `salesdisc@corporate.net`. The message is delivered to the resource mailbox, where it is picked up by the Teaming SMTP server for posting in the designated folder on the Teaming site.

23.3.2 Performing the Novell Teaming Installation

As you follow the instructions in “[Installing Novell Teaming](#)” in the *Novell Teaming 1.0 Installation Guide*, use the system and GroupWise information you gathered in [Section 23.3.1, “Preparing for the Novell Teaming Installation,” on page 248](#).

23.4 Configuring LDAP Authentication

As you follow the instructions in “[Configuring Liferay for LDAP Authentication](#)” in “[Configuring Novell Teaming](#)” in the *Novell Teaming 1.0 Administration Guide*, use `cn=@screen_name@` as the search filter for compatibility with eDirectory.

23.5 Configuring GroupWise to Support Your Teaming Installation

Novell Teaming requires an integrated e-mail system in order for the Teaming site and the GroupWise system to exchange messages.

- ♦ [Section 23.5.1, “Configuring Outgoing E-Mail from Teaming to GroupWise,”](#) on page 249
- ♦ [Section 23.5.2, “Configuring Incoming E-Mail from GroupWise Users to the Novell Teaming Site,”](#) on page 251
- ♦ [Section 23.5.3, “Testing GroupWise Integration with Novell Teaming,”](#) on page 257

23.5.1 Configuring Outgoing E-Mail from Teaming to GroupWise

During installation, you set up outgoing E-Mail in the [Outbound E-Mail Configuration](#) window of the Teaming installer. If you need to change the outbound e-mail configuration after installation, you need to use the Liferay Enterprise Admin portlet and the Teaming Administration portlet on the Teaming site. Also after installation, you need to create the GroupWise user that you configured to send outgoing mail from the Teaming site.

- ♦ [“Adding Portlets for Outgoing E-Mail Setup”](#) on page 249
- ♦ [“Configuring Outgoing E-Mail from Novell Teaming to GroupWise Users”](#) on page 250

Adding Portlets for Outgoing E-Mail Setup

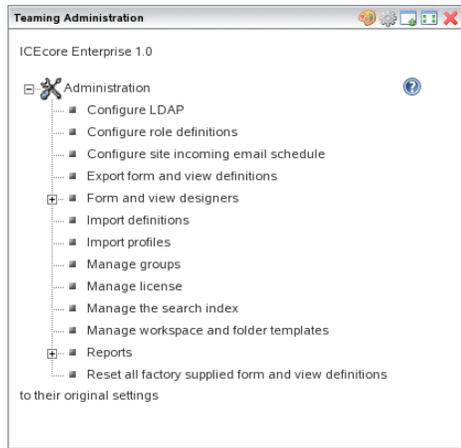
You configure LDAP using the Liferay Enterprise Admin portlet and the Teaming Administration portlet. If you have not already added these portlets to your Teaming home page, do so now.

- 1 Click *Add Content*.
- 2 Click *Admin*, then click *Add* beside Enterprise Admin to add the Liferay Enterprise Admin portlet to your Teaming home page.



The screenshot shows the 'Enterprise Admin' portlet interface. At the top, there are tabs for 'Users', 'Organizations', 'Locations', 'User Groups', and 'Roles'. The 'Users' tab is selected. Below the tabs, there are input fields for 'First Name', 'Middle Name', and 'Last Name'. Below these are fields for 'Screen Name' and 'Email Address'. To the right of the 'Email Address' field is an 'Active' dropdown menu with 'Yes' selected. At the bottom left, there is an 'And' dropdown menu and a 'Search' button.

- 3 Click *Teaming*, then click *Add* beside Teaming Administration to add the Novell Teaming Administration portlet to your Teaming home page.



4 If necessary, close the Add Content window.

When you add the Novell Teaming Administration portlet, the window closes automatically.

When you add the Liferay Enterprise Admin portlet, the window remains open so you can add more Liferay portlets.

Configuring Outgoing E-Mail from Novell Teaming to GroupWise Users

Outgoing e-mail from your Teaming site consists of automated notifications of changes in your Teaming site's content and messages from Teaming site users sent to GroupWise mailboxes.

Your Novell Teaming system needs an administrator e-mail account that is responsible for sending notifications from Teaming folders and e-mail messages from Teaming site users to GroupWise mailboxes. During installation, you specified this account in the **Outbound E-Mail Configuration** window in the Teaming installer. You use ConsoleOne to create the Teaming administrator account in GroupWise.

1 In ConsoleOne, browse to and right-click the container where you want to create the Teaming administrator User object, then click *New > User*.

- 2 In the *Name* field, specify the GroupWise username for the Teaming administrator user (for example, TeamingAdmin).
- 3 Provide a surname for the new user.
The *Unique ID* field is filled in to match the *Name* field by default.
- 4 Assign a password to the eDirectory account.
- 5 Add the Teaming administrator user to a GroupWise post office.
- 6 Select *Define Additional Properties*, if desired, then click *OK* to create the Teaming administrator user.

23.5.2 Configuring Incoming E-Mail from GroupWise Users to the Novell Teaming Site

Incoming e-mail to your Teaming site consists of postings that GroupWise users make to the Teaming site by sending e-mail messages rather than by logging into the Teaming site. Postings can consist of contributions to discussion groups, wikis, or blogs, files for sharing with the team, calendar items, and so on. You should already have created at least one workspace by following the instructions in “[Create Your Initial Workspaces](#)” in “[Configuring Novell Teaming](#)” in the *Novell Teaming 1.0 Administration Guide*.

- ◆ “[Enabling Incoming E-Mail](#)” on page 251
- ◆ “[Creating GroupWise Resources for Teaming Folders](#)” on page 252
- ◆ “[Creating a Distribution List for GroupWise Resources](#)” on page 252
- ◆ “[Setting Passwords for GroupWise Resources](#)” on page 253
- ◆ “[Associating a Resource with a Novell Teaming Folder](#)” on page 254
- ◆ “[Configuring the GroupWise Internet Agent for POP or IMAP](#)” on page 256

Enabling Incoming E-Mail

Incoming e-mail is enabled in the [Inbound E-Mail Configuration](#) windows in the Teaming installer. This window is available when you perform an Advanced installation.

To enable incoming e-mail after installation:

- 1 As the Teaming Site Manager user (username `admin`), display the Teaming site Home page.
- 2 In the Novell Teaming Administration portlet, click *Configure Site Incoming Email Schedule*.

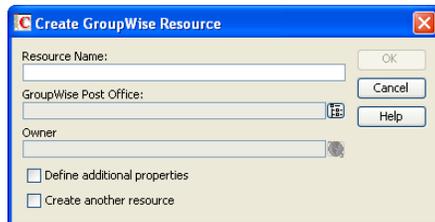
The screenshot shows the Novell Teaming Administration interface. At the top, there is a navigation bar with "Novell Teaming" on the left and "Welcome Mary Admin!" on the right. Below this is a secondary navigation bar with "Home A1" and "Add Page" buttons. The main content area is titled "Teaming Administration" and contains a "Portal" link. The "Incoming mail schedule" section is active, showing a table for "Incoming addresses" with columns for "Delete", "Address", and "Associated folder". Below the table, there are configuration options for the incoming mail schedule, including a radio button for "Every day", a radio button for "Weekly (on the days selected below)", a radio button for "At time", and a radio button for "Repeat every". The "At time" option is selected, with the time set to 12:15 GMT and a repeat interval of 0.25 hours. There are "Apply" and "Close" buttons at the bottom of the configuration window.

- 3 Select *Enable Incoming Mail*.
- 4 Select *Every Day*.
or
Select *Weekly*, then select the days of the week when you want the Teaming SMTP server to check for incoming mail.
- 5 Select *Repeat Every*.
- 6 Leave the default of 0.25 hours.
This configures the Teaming SMTP server to check the GroupWise Teaming mailbox for incoming messages every 15 minutes. You can adjust this schedule as necessary to meet the needs of your Teaming site.
- 7 Click *Apply* to save your incoming e-mail settings.

Creating GroupWise Resources for Teaming Folders

In order for GroupWise users to post items to Teaming folders, each folder needs a unique e-mail address. This is accomplished by creating a GroupWise resource that corresponds to each folder where e-mail posting is needed. For example, if a sales team has a Discussion folder, you could create a resource named salesdisc. Then, when sales people want to post items to the Discussion folder, they can mail a posting to an e-mail address similar to salesdisc@corporate.net. The message is delivered to the resource mailbox, where it is picked up by the Teaming SMTP server for posting in the designated folder on the Teaming site.

- 1 In ConsoleOne, browse to and right-click the container object where the Teaming administrator User object is located, then click *New > Resource*.

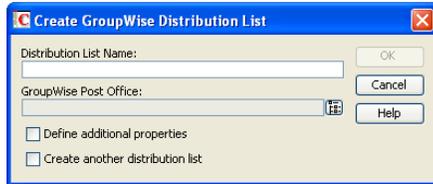


- 2 Specify the name for the resource.
- 3 Select the post office that the Teaming administrator user belongs to.
- 4 Select the Teaming administrator user as the resource owner.
- 5 Click *OK*.
- 6 Repeat **Step 1** through **Step 5** as needed to create resources for all Teaming folders where postings from GroupWise are needed.

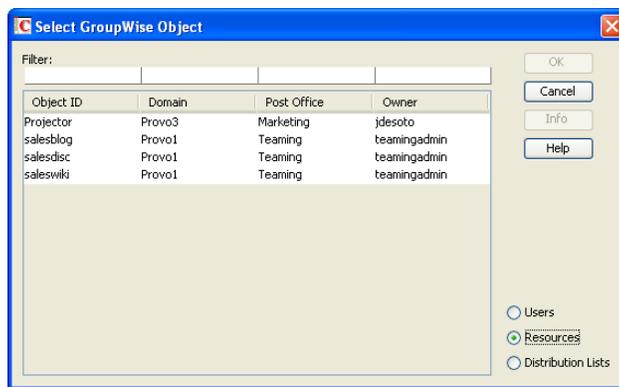
Creating a Distribution List for GroupWise Resources

For convenience in creating an Internet Agent class of service specifically for Teaming resources, you need to create a distribution list that includes all the resources that correspond to Teaming folders.

- 1 In ConsoleOne, browse to and right-click the container object where the Teaming administrator User object is located, then click *New > Distribution List*.



- 2 Specify the name for the distribution list (for example, TeamingResources).
- 3 Select the post office that the Teaming administrator user belongs to.
- 4 Select *Define Additional Properties*, then click OK.
- 5 On the Distribution List object Identification page, click *GroupWise > Membership*, then click *Add*.
- 6 Select *Resources* to filter the list by resources,



- 7 Select the resources you have created for use with Teaming, then click *OK*.
- 8 Click *OK* again to add the listed resources to the distribution list.

Setting Passwords for GroupWise Resources

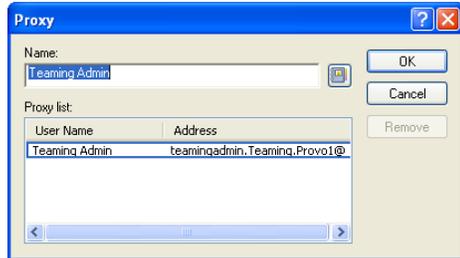
Because the Teaming administrator user owns all the resources that correspond to Teaming folders, the Teaming administrator can proxy into the resource mailboxes to set passwords.

- 1 Start the GroupWise Windows client as the Teaming administrator user.
 - 1a Double-click the GroupWise icon on your desktop.
 - 1b When prompted for your personal GroupWise password, click *Cancel* to display the Novell GroupWise Startup dialog box.



- 1c Replace your username with the username of the Teaming administrator user.

- 1d Specify the administrator user's password.
 - 1e Select *Online*, then specify the IP address and TCP port of the administrator user's POA.
 - 1f Click *OK* to log in to the administrator user's mailbox.
- 2 In the upper left corner of the GroupWise client window, click *Online > Proxy*.



- 3 Click the *Address Book* icon, then select the Novell GroupWise Address Book.
- 4 In the *Look For* field, start typing the name of the resource until it is highlighted in the Address List, then click *OK* to proxy into the resource mailbox.
This also adds the resource to your proxy list, so that the next time you need to proxy into this resource, you can simply click *Proxy > resource_name*.
- 5 In the resource mailbox, click *Tools > Options > Security*.
Use the same password for the resource mailbox that you used for the Teaming administrator user's mailbox.
- 6 Type the password twice for verification, then click *OK*.
- 7 In the upper left corner of the GroupWise client window, click *Proxy*, then select your username to return to your own mailbox.
- 8 Repeat **Step 2** through **Step 7** for each resource you have created.

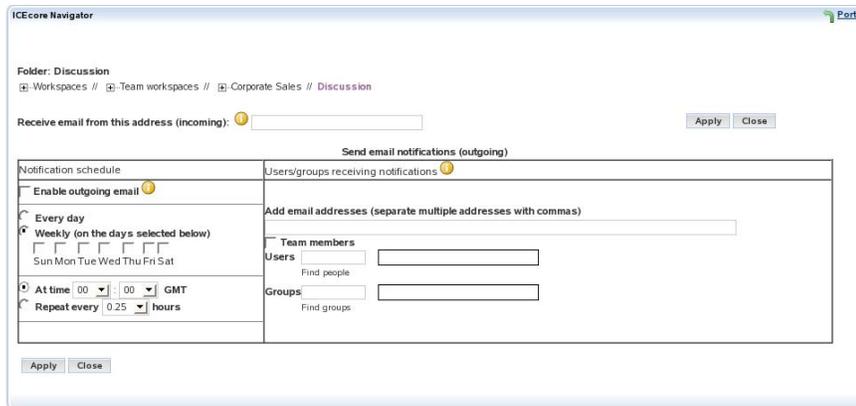
Associating a Resource with a Novell Teaming Folder

Each Teaming folder that can receive postings by way of e-mail must be configured with its resource e-mail address.

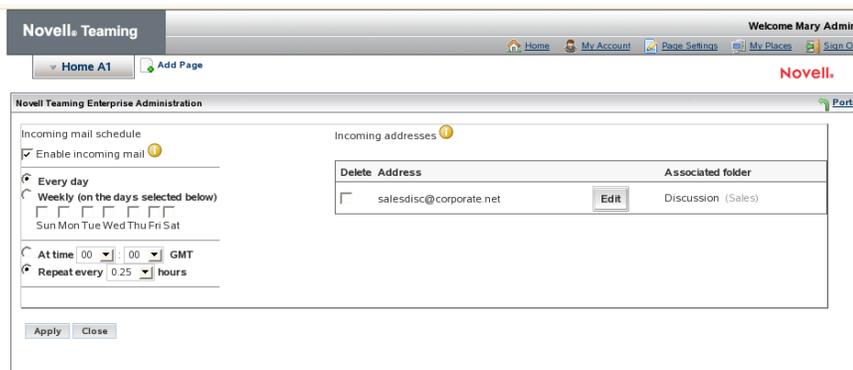
- 1 Log in to your Teaming site as the Teaming Site Manager user (by default, username `admin`, password `admin`).
- 2 Browse to and open the folder that you want to associate with a resource.



- 3 Click *Manage*, then click *Email Settings*.



- 4 In the *Receive Email from This Address (Incoming)* field:
 - 4a Specify the e-mail address of the resource (for example, salesdisc@corporate.net).
 - 4b In the *Password* field, specify the resource password.
 - 4c Click *Apply*.
- 5 Select *Enable Outgoing Email* if you want users to be able to subscribe to notifications for this folder, then click *Apply*.
If desired, you can set up a unique outgoing notification schedule for this folder as well.
- 6 Repeat **Step 2** through **Step 5** for each Teaming folder that you want GroupWise users to be able to post to by way of e-mail.
- 7 When you are finished adding resource e-mail addresses to folders, return to the Teaming home page.
- 8 In the Teaming Administration portlet, click *Configure Site Incoming Email Schedule*.
The e-mail addresses for the resources that you have provided on folders have been added to the Incoming Addresses list.



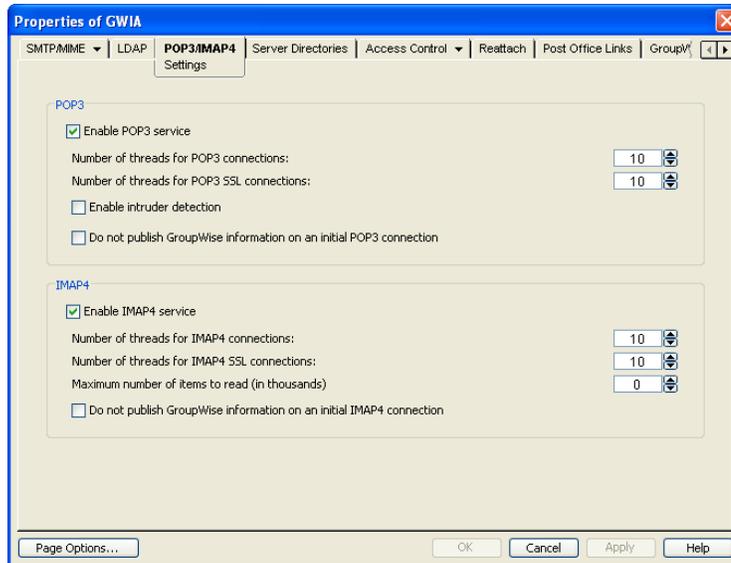
- 9 For each e-mail address, click *Edit*, specify the password for the e-mail account, then click *OK*.
- 10 When all the e-mail addresses have passwords associated with them, return to the Teaming home page.

Configuring the GroupWise Internet Agent for POP or IMAP

By default, the GroupWise Internet Agent is configured for both POP and IMAP. Whether your Teaming site uses POP or IMAP is determined by your selection in the [Inbound E-Mail Configuration](#) window when you ran the Teaming installer.

To verify that the Internet Agent is properly configured:

- 1 In ConsoleOne, browse to and select the Domain object where the Internet Agent object is located, right-click the Internet Agent object, then click *Properties*.
- 2 Click *GroupWise > POP3/IMAP4*.

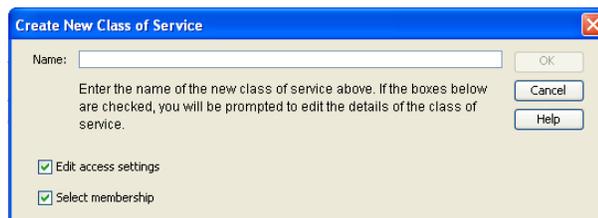


- 3 Make sure that either POP or IMAP is enabled, to match the configuration you provided in the [Inbound E-Mail Configuration](#) window when you ran the Teaming installer.
- 4 If necessary, adjust the POP or IMAP settings as needed for your Internet Agent.

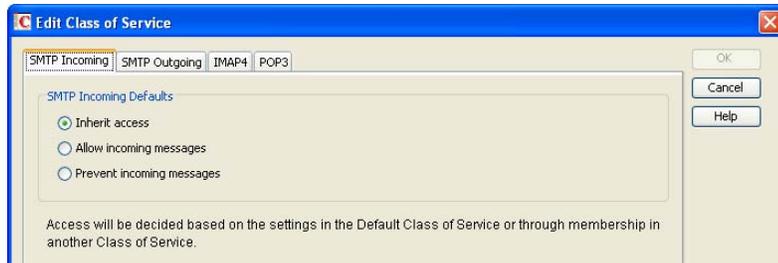
For more information, see “[Configuring POP3/IMAP4 Services](#)” in “[Internet Agent](#)” in the [GroupWise 7 Administration Guide](#).

- 5 If you are using POP, create a new class of service specifically for interacting with the Teaming site:

- 5a Click *Access Control*, then click *Create*.



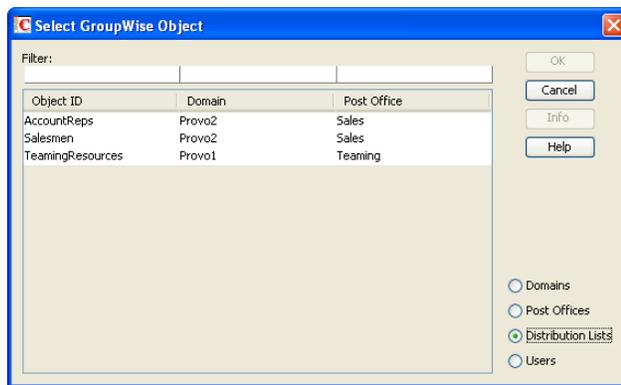
- 5b Specify a name for the new class of service, such as Teaming Site, then click *OK*.



5c On the *SMTP Incoming* tab, select *Allow Incoming Messages*.

5d On the *SMTP Outgoing* tab, select *Allow Outgoing Messages*.

5e On the *POP3* tab, select *Allow Access*, select *Delete Messages from GroupWise Mailbox after Download*, then click *OK*.



5f Select the distribution list of Teaming resources, then click *OK* to add the new class of service.

6 Click *OK* to save the settings.

23.5.3 Testing GroupWise Integration with Novell Teaming

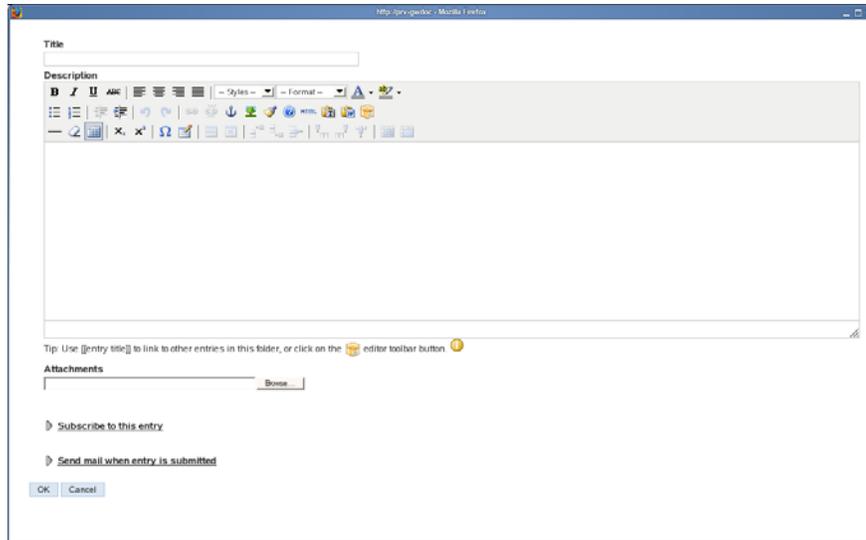
After you have configured outgoing and incoming e-mail, you are ready to test your e-mail integration.

- ♦ “Testing Outgoing E-Mail from Your Teaming Site” on page 257
- ♦ “Testing an Incoming Posting to Your Teaming Site” on page 258

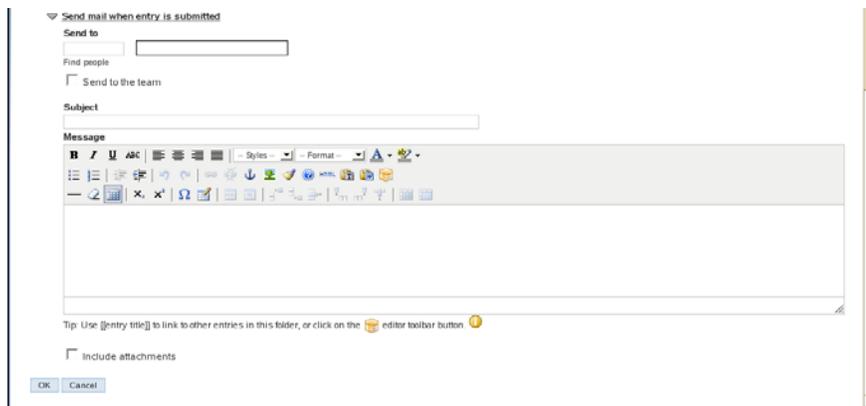
Testing Outgoing E-Mail from Your Teaming Site

To verify that e-mail is successfully delivered from the Teaming site to the GroupWise system:

- 1** Log in as yourself to the Teaming site.
- 2** Browse to and open a Teaming folder where you are a member of the team.
- 3** Click *New*, then add a new entry to the folder.



- 4 Fill in the *Title* and *Description* fields to create a test posting, then click *Send Mail When Entry Is Submitted*.



- 5 Select *Send to the Team*.
- 6 Fill in the *Subject* and *Message* fields to create a test notification message, then click *OK*.
- 7 Check your GroupWise mailbox to see if you received the notification of the new posting.
- 8 Check with other members of the team, especially remote team members, to verify that they also received the notification.
- 9 Open the message, then click the link to verify that it opens the page of the Teaming site from which the notification was sent.

Testing an Incoming Posting to Your Teaming Site

To verify that e-mail is successfully delivered from the GroupWise system to the Teaming site:

- 1 In the GroupWise client, send a message to one of the nickname addresses you created in [“Associating a Resource with a Novell Teaming Folder” on page 254](#).
- 2 Open that folder on the Teaming site and see if your message has been posted.

23.6 Receiving Notification of Teaming Site Activity in Your GroupWise Client

You can configure any Teaming folder to send you a notification when information in that folder changes.

- 1 Log in to your Teaming account.
- 2 Browse to and open a Teaming folder that you want to be able to monitor from GroupWise.
- 3 Click *Subscriptions > Email Notification*.



- 4 Select the type of e-mail notification you want to receive for this folder, then click *OK*.

From now on, you can monitor changes in that Teaming folder from GroupWise. The notification message you receive provides a link to the folder, so that you can quickly log into the Teaming site and check the new information.

23.7 Receiving Calendar Postings from the Teaming Site

Appointments created on a Teaming calendar arrive just like appointments created on a GroupWise calendar. When you open the Teaming appointment, you can accept it or decline it. If you accept it, the Teaming appointment is placed on your GroupWise calendar.

23.8 Adding Portlets to Novell Teaming

Portlets are optional user interface components that you can add to your Novell teaming site to increase its functionality.

- ♦ [Section 23.8.1, “Adding Liferay Portlets,” on page 259](#)
- ♦ [Section 23.8.2, “Adding GroupWise Portlets,” on page 262](#)

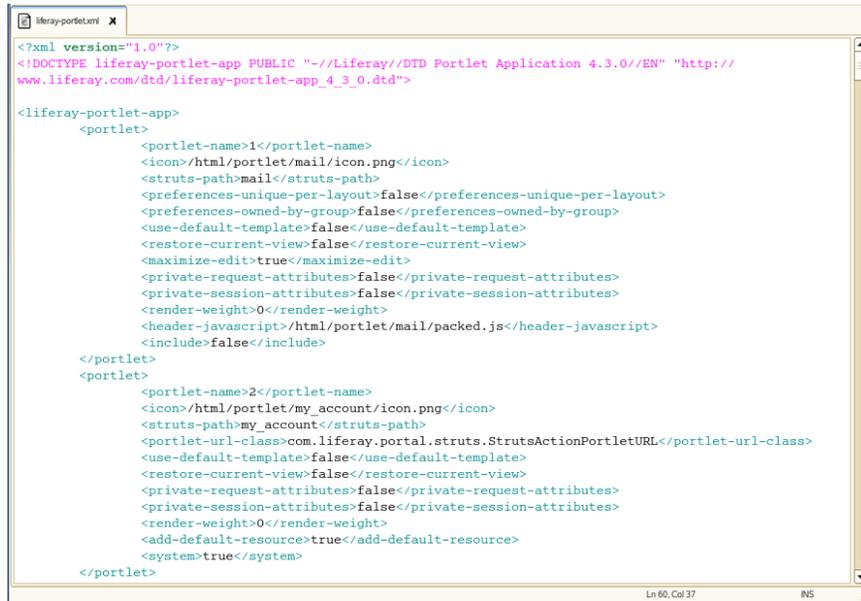
23.8.1 Adding Liferay Portlets

Novell Teaming includes the [Liferay 4.3 portal](http://www.liferay.com/web/guest/home). (<http://www.liferay.com/web/guest/home>) and makes some Liferay portlets available under the *Add Content* link on the Teaming Home page, specifically those that are useful in a Teaming environment. If you want, you can enable any of the Liferay portlets so that they appear under the *Add Content* link.

The file that lists all available Liferay portlets is named `liferay-portlet.xml` and is located in the following directory:

NOTE: The part of the path before the webapps directory might be different on your system, depending on where you installed Novell Teaming.

- 1 Edit the liferay-portlet.xml file.

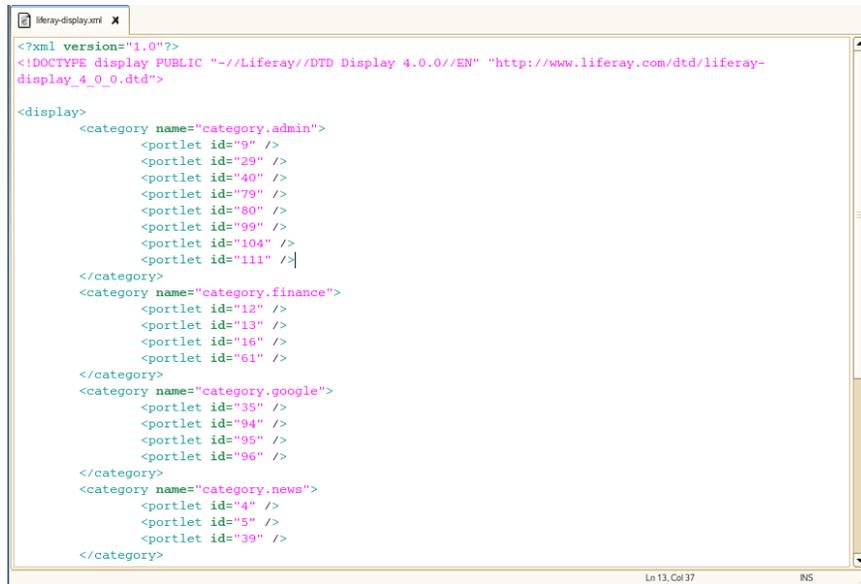


```
<?xml version="1.0"?>
<!DOCTYPE liferay-portlet-app PUBLIC "-//Liferay//DTD Portlet Application 4.3.0//EN" "http://
www.liferay.com/dtd/liferay-portlet-app_4_3_0.dtd">

<liferay-portlet-app>
  <portlet>
    <portlet-name>1</portlet-name>
    <icon>/html/portlet/mail/icon.png</icon>
    <struts-path>mail</struts-path>
    <preferences-unique-per-layout>>false</preferences-unique-per-layout>
    <preferences-owned-by-group>>false</preferences-owned-by-group>
    <use-default-template>>false</use-default-template>
    <restore-current-view>>false</restore-current-view>
    <maximize-edit>>true</maximize-edit>
    <private-request-attributes>>false</private-request-attributes>
    <private-session-attributes>>false</private-session-attributes>
    <render-weight>0</render-weight>
    <header-javascript>/html/portlet/mail/packed.js</header-javascript>
    <include>>false</include>
  </portlet>
  <portlet>
    <portlet-name>2</portlet-name>
    <icon>/html/portlet/my_account/icon.png</icon>
    <struts-path>my_account</struts-path>
    <portlet-url-class>com.liferay.portal.struts.StrutsActionPortletURL</portlet-url-class>
    <use-default-template>>false</use-default-template>
    <restore-current-view>>false</restore-current-view>
    <private-request-attributes>>false</private-request-attributes>
    <private-session-attributes>>false</private-session-attributes>
    <render-weight>0</render-weight>
    <add-default-resource>>true</add-default-resource>
    <system>true</system>
  </portlet>
</liferay-portlet-app>
```

Each portlet is defined between `<portlet>` and `</portlet>` tags. Each portlet has a name which is a number from 1 to 113.

- 2 To turn on a portlet, edit the edit the value of the `<include>` tag from `false` to `true`.
- 3 Make a note of the portlet function and its corresponding number.
- 4 Repeat **Step 2** and **Step 3** for each Liferay portlet that you want to enable.
- 5 Save and close the `liferay-portlet.xml` file.
- 6 Edit the `liferay-display.xml` file.

A screenshot of a text editor window titled 'liferay-display.xml'. The code is XML and defines a 'display' container with five categories: 'category.admin', 'category.finance', 'category.google', 'category.news', and 'category.tools'. Each category contains a list of portlets with specific IDs. The 'category.admin' category contains portlets with IDs 9, 29, 40, 79, 80, 99, 104, and 111. The 'category.finance' category contains portlets with IDs 12, 13, 16, and 61. The 'category.google' category contains portlets with IDs 35, 94, 95, and 96. The 'category.news' category contains portlets with IDs 4, 5, and 39. The 'category.tools' category is empty. The XML code is as follows:

```
<?xml version="1.0"?>
<!DOCTYPE display PUBLIC "-//Liferay//DTD Display 4.0.0/EN" "http://www.liferay.com/dtd/liferay-
display_4_0_0.dtd">

<display>
  <category name="category.admin">
    <portlet id="9" />
    <portlet id="29" />
    <portlet id="40" />
    <portlet id="79" />
    <portlet id="80" />
    <portlet id="99" />
    <portlet id="104" />
    <portlet id="111" />
  </category>
  <category name="category.finance">
    <portlet id="12" />
    <portlet id="13" />
    <portlet id="16" />
    <portlet id="61" />
  </category>
  <category name="category.google">
    <portlet id="35" />
    <portlet id="94" />
    <portlet id="95" />
    <portlet id="96" />
  </category>
  <category name="category.news">
    <portlet id="4" />
    <portlet id="5" />
    <portlet id="39" />
  </category>
  <category name="category.tools">
  </category>
</display>
```

Each category is defined between `<category . . . >` and `</category>` tags. Each portlet has a name that is defined within the `<category . . . >` tag. The default set of tags in the `liferay-portlet.xml` file are:

admin
finance
google
news
tools

Additional Liferay-defined categories include:

cms
journal
collaboration
community
entertainment
category.polls
category.religion
category.christianity
category.sample
category.shopping
category.wiki
category.workflow
category.wsrp

- 7 To add the portlet to an existing category on the portlet menu, add a line similar to the following between the `<category . . . >` and `</category>` tags.

```
<portlet id="nn" />
```

where *nn* is the numeric name of the portlet.

or

To use a Liferay-defined category for the portlet, add a set of lines similar to the following:

```
<category name="category.category_name">
    <portlet id="nn" />
</category>
```

where *category_name* is one of the pre-defined Liferay categories (for example, entertainment) and *nn* is the numeric name of one of the Liferay portlets (for example, 14 for games).

If you want to define your own unique categories, you must edit the `/Language-ext.properties` file located in the following directory:

```
/opt/icecore/liferay-portal-tomcat-5.5-jdk5-4.3.0/webapps/ROOT/
                                WEB-INF/classes/content
```

NOTE: The part of the path before the `webapps` directory might be different on your system, depending on where you installed Novell Teaming.

- 8** Repeat **Step 7** for each portlet you defined in **Step 2** and **Step 3**.
- 9** Save and close the `liferay-display.xml` file.
- 10** Restart Novell Teaming to make the additional portlets available to Teaming users.

23.8.2 Adding GroupWise Portlets

GroupWise Mail and Calendar portlets are available for adding GroupWise functionality into your Teaming site. For instructions, see [Novell Collaboration Portlets \(http://developer.novell.com/wiki/index.php/Novell_Collaboration_Portlets\)](http://developer.novell.com/wiki/index.php/Novell_Collaboration_Portlets).

Using GroupWise with Conferencing

24

When you install the Conferencing component of Novell® Teaming + Conferencing together with Novell eDirectory™ and GroupWise®, some configuration steps are required to integrate the applications.

- ♦ [Section 24.1, “Configuring GroupWise as the Conferencing E-Mail System,” on page 263](#)

24.1 Configuring GroupWise as the Conferencing E-Mail System

By default, Conferencing is configured to send meeting invitation messages from the administrator account of the SMTP server on the host where Conferencing is installed. You do not need to use the GroupWise Internet Agent for this purpose, but you can.

- ♦ [Section 24.1.1, “Preparing for GroupWise Integration with Conferencing,” on page 263](#)
- ♦ [Section 24.1.2, “Integrating GroupWise with Novell Conferencing,” on page 263](#)
- ♦ [Section 24.1.3, “Testing GroupWise as the Novell Conferencing E-Mail System,” on page 264](#)

24.1.1 Preparing for GroupWise Integration with Conferencing

Before you install Conferencing with GroupWise, you need to gather the following GroupWise information:

- ♦ A GroupWise user who has been designated as a Postmaster for the Internet Agent. This user receives messages about problems with the Conferencing server.
- ♦ The IP address or DNS hostname of a server where the GroupWise Internet Agent is running.
- ♦ A GroupWise username and password for the Conferencing mailer to use when connecting to the GroupWise Internet Agent to transfer meeting invitation messages into the GroupWise system (optional).

24.1.2 Integrating GroupWise with Novell Conferencing

GroupWise integration is established when you install Conferencing.

During the installation, you are prompted:

```
Do you want to change the system administrator e-mail configuration? (y/n/?)
```

- 1 Type `y`, then press Enter.
- 2 Type the GroupWise Postmaster e-mail address, then press Enter.
You can enter multiple e-mail addresses if needed.
- 3 Press Enter again to end the list of addresses.
The e-mail address you entered is displayed for verification.

4 Type `y`, then press Enter.

Next, you are prompted:

```
Do you want to change the mailer configuration? (y/n/?)
```

5 Type `y`, then press Enter.

Next, you are prompted:

```
What SMTP server will the Conferencing mailer use?
```

6 Specify the IP address or DNS hostname of the server where the GroupWise Internet Agent is running, then press Enter.

Next, you are prompted:

```
Use user and password authentication for SMTP_server? (y/n/?)
```

7 If you want the Conferencing mailer to authenticate to the GroupWise Internet Agent, type `y`, then press Enter.

7a Specify a valid GroupWise username, then press Enter.

7b Specify the password for the username, then press Enter.

8 Continue as usual with the Conferencing installation.

24.1.3 Testing GroupWise as the Novell Conferencing E-Mail System

- 1** Make sure that users receive the Conferencing Welcome e-mail, as described in “[Installing the Conferencing Client](#)” in the *Conferencing 1.0 User Guide*
- 2** When you create meetings, as described in “[Working with Meetings](#)” in the *Conferencing 1.0 User Guide*, make sure that users are being notified of meetings.

Streamlining Authentication to Teaming or Conferencing

25

You can implement single sign-on for use with Teaming or Conferencing.

- [Section 25.1, “Using iChain for Authenticating to Teaming or Conferencing,” on page 265](#)
- [Section 25.2, “Using Novell Access Manager for Authenticating to Teaming or Conferencing,” on page 267](#)

25.1 Using iChain for Authenticating to Teaming or Conferencing

You can use Novell® iChain® to eliminate a dual user login into your network and into Teaming or Conferencing. The instructions in this section assume that you have an understanding of iChain, as described on the [Novell iChain 2.3 Documentation Web site \(http://www.novell.com/documentation/ichain23\)](http://www.novell.com/documentation/ichain23) and that you have iChain set up and running on your system.

There are many ways to configure iChain. This section illustrates one possible way to configure iChain to support Teaming or Conferencing. Before following the steps in this section, you must have Teaming or Conferencing, as well as iChain, installed, configured, and running.

- [Section 25.1.1, “Meeting iChain Requirements,” on page 265](#)
- [Section 25.1.2, “Setting Up an iChain Web Server Accelerator for Teaming or Conferencing,” on page 265](#)
- [Section 25.1.3, “Adding the New Web Server Accelerator to the iChain Server Object in ConsoleOne,” on page 266](#)
- [Section 25.1.4, “Using iChain for Authentication,” on page 267](#)

25.1.1 Meeting iChain Requirements

In order to get the best performance and reliability from iChain with Teaming and Conferencing, you must install iChain 2.3 Support Pack 5 Release 4 version 2.3.410. This software is available on the iChain *Patches* tab on the [Novell Downloads Web site \(http://download.novell.com\)](http://download.novell.com). Follow the installation instructions that are provided with the patch.

25.1.2 Setting Up an iChain Web Server Accelerator for Teaming or Conferencing

- 1 Access the iChain Proxy Administration Tool at the following URL:

`http://proxy_server_address:port/appliance/config.html`

- 2 Click *Configure*, then click *Insert* to create a new Web server accelerator for Teaming or Conferencing

The new accelerator is enabled by default.

- 3 In the *Name* field, provide a unique and descriptive name for the new accelerator.

For example, you might want to call it Teaming or Conferencing, as appropriate.

- 4 Select *Allow Pages to Be Cached at the Browser*.
- 5 Select *Enable Multi-Homing*.
 - 5a In the Multi-Homing Options dialog box, select *Domain-Based Multi-Homing* to configure the Teaming or Conferencing URL as a DNS name prepended to your Internet domain name, for example:
`http://teaming.corporate.net`
The A record for the DNS name must already exist. The Proxy Administration Tool does not create it for you.
 - 5b In the *DNS Name* field, specify the DNS A record.
 - 5c Click *OK* to save your multi-homing settings.
- 6 If you have created a custom login page for your Teaming or Conferencing Web site, specify it in the *Custom Login Page Location* field.
The default location for custom login pages is `sys:\etc\proxy\data`. The custom login page must be an HTML file with a `.htm` extension. If it is located in a directory other than the default, specify the full pathname for the file.
- 7 Select *Enable Secure Exchange*.
 - 7a In the *Port* field on the right, specify the port number that the iChain proxy server should use to communicate with the Web server where Teaming or Conferencing is installed.
 - 7b If desired, select *Enable Secure Access between the iChain Proxy and the Origin Web Server*.
 - 7c Click *OK* to save your secure exchange options.
- 8 Under the Web Server Addresses box, click *Insert*.
 - 8a Specify the IP address or DNS hostname of the Web server where you have installed Teaming or Conferencing.
 - 8b Click *OK* to add the Web server to the list in the Web Server Accelerator dialog box.
- 9 Click *OK* to save the new Web server accelerator.

25.1.3 Adding the New Web Server Accelerator to the iChain Server Object in ConsoleOne

- 1 Start ConsoleOne in a location where the iChain snap-ins are installed.
- 2 Browse to and right-click the iChain Server object, then click *Properties*.
- 3 Click *Protected Resource* to display a list of protected resources.
- 4 Click the *Plus* icon to add a new protected resource.
 - 4a In the *Resource Name* field, provide a unique and descriptive name for the new protected resource, which is the Web server accelerator.
 - 4b In the *URL Prefix* field, specify the part of the URL that precedes the application-specific part of the URL; for example:
`teaming.corporate.net/*`

- 4c** Select the type of access you want to provide for users to view the URL: *Secure*, *Restricted*, or *Public*.
- 4d** Click *OK* to save the new protected resource.
- 5** Select the new protected resource, then click the *Parameters* icon to display the OLAC Parameters dialog box.
 - 5a** In the *Name* column, specify *Authorization*.
 - 5b** In the *Data Source* column, specify *ldap*.
 - 5c** In the *Value* column, specify *cn*.

These settings add an extended HTTP request header called X-Authorization that stores each user's cn (common name). The cn is retrieved from the LDAP server by the iChain OLAC process so that users can log in automatically.
 - 5d** Click *OK* to save the OLAC parameters.
- 6** When prompted, click *Yes* to refresh the iChain proxy configuration with the new changes.
- 7** Provide the password to the proxy server, then click *OK* to perform the refresh operation immediately.

25.1.4 Using iChain for Authentication

Now that you have created an iChain Web server accelerator for Teaming or Conferencing and have configured the iChain Server object for the new Web server accelerator, users should be able to authenticate to Teaming or Conferencing in a single step, using their eDirectory or LDAP passwords.

25.2 Using Novell Access Manager for Authenticating to Teaming or Conferencing

Novell Access Manager is not currently supported for use with Novell Teaming or Conferencing.

Other Novell Products

IV

- ♦ Chapter 26, “GroupWise DirXML Driver for Novell Identity Manager,” on page 271
- ♦ Chapter 27, “GroupWise Customization Tools,” on page 275
- ♦ Chapter 28, “Novell exteNd,” on page 277

GroupWise DirXML Driver for Novell Identity Manager

26

The GroupWise® DirXML® driver for use with Novell® Identity Manager provides data integration between users in Novell eDirectory™ with GroupWise accounts in your GroupWise system. For example, the driver can create e-mail accounts automatically when employees are hired. The driver can also disable an e-mail 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 gives information about certain administrative actions in ConsoleOne® that require you to stop the GroupWise DirXML driver or disable a user's association:

- ♦ [Section 26.1, “Identity Manager Warnings in ConsoleOne,” on page 271](#)

For additional information, see:

- ♦ [Novell Identity Manager \(http://www.novell.com/documentation/dirxml20\)](http://www.novell.com/documentation/dirxml20)
- ♦ [Identity Manager Drivers \(http://www.novell.com/documentation/dirxmldrivers\)](http://www.novell.com/documentation/dirxmldrivers)

26.1 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 and usually restart the GroupWise DirXML driver or re-enable the user's association when you have completed the action. By default, these activities generate a warning message in ConsoleOne:

- ♦ [Section 26.1.1, “Recovering a Deleted GroupWise Account,” on page 271](#)
- ♦ [Section 26.1.2, “Grafting Users,” on page 272](#)
- ♦ [Section 26.1.3, “Converting an External Entity to a User,” on page 272](#)
- ♦ [Section 26.1.4, “Converting a User to an External Entity,” on page 272](#)
- ♦ [Section 26.1.5, “Associating a GroupWise Object with an eDirectory Object,” on page 272](#)
- ♦ [Section 26.1.6, “Disassociating a GroupWise Object's Attributes from an eDirectory Object,” on page 272](#)
- ♦ [Section 26.1.7, “Resolving an Invalid Association,” on page 273](#)
- ♦ [Section 26.1.8, “Disabling the DirXML Warnings,” on page 273](#)
- ♦ [Section 26.1.9, “Enabling the DirXML Warnings,” on page 273](#)

26.1.1 Recovering a Deleted GroupWise Account

- 1 Using the DirXML Management role in Novell iManager, stop the GroupWise DirXML driver.
- 2 Recover the deleted account, as described in “[Recovering Deleted GroupWise Accounts](#)” in “[Databases](#)” in the *GroupWise 7 Administration Guide*.
- 3 Using the DirXML Management role, restart the GroupWise DirXML driver.

26.1.2 Grafting Users

- 1 If you are grafting the users into a different eDirectory tree, on the *DirXML* tab of each User object in Novell iManager, disable the association with the GroupWise DirXML driver.
- 2 Using the DirXML Management role in Novell iManager, 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 7 Administration Guide*.
- 4 If you grafted the users into a different eDirectory tree, on the *DirXML* tab of each User object, enable the association with the GroupWise DirXML driver in the new tree.
- 5 Using the DirXML Management role, restart the GroupWise DirXML driver for the tree into which you grafted the users.

26.1.3 Converting an External Entity to a User

- 1 Using the DirXML Management role in Novell iManager, stop the GroupWise DirXML driver.
- 2 Convert the external entity, as described in “**Convert External Entity to User**” in “**System**” in the *GroupWise 7 Administration Guide*.
- 3 Using the DirXML Management role, restart the GroupWise DirXML driver.

26.1.4 Converting a User to an External Entity

- 1 On the *DirXML* tab of the User object in Novell iManager, 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 7 Administration Guide*.

26.1.5 Associating a GroupWise Object with an eDirectory Object

- 1 Using the DirXML Management role in Novell iManager, stop the GroupWise DirXML driver.
- 2 Establish the association, as described in “**Associate Objects**” in “**System**” in the *GroupWise 7 Administration Guide*.
- 3 Using the DirXML Management role, restart the GroupWise DirXML driver.

26.1.6 Disassociating a GroupWise Object’s Attributes from an eDirectory Object

- 1 On the *DirXML* tab of the User object in Novell iManager, disable the association with the GroupWise DirXML driver.
- 2 Disassociate the objects, as described in “**Disassociate GroupWise Attributes**” in “**System**” in the *GroupWise 7 Administration Guide*.
- 3 On the *DirXML* tab of the User object, enable the association with the GroupWise DirXML driver.

26.1.7 Resolving an Invalid Association

- 1 On the *DirXML* tab of the User object in Novell iManager, disable the association with the GroupWise DirXML driver.
- 2 Resolve the invalid association, as described in “Invalid Associations” in “System” in the *GroupWise 7 Administration Guide*.

26.1.8 Disabling the DirXML Warnings

- 1 In ConsoleOne, deselect *Display DirXML Warnings* in any DirXML warning dialog box.

26.1.9 Enabling the DirXML Warnings

- 1 In ConsoleOne, click *Tools > GroupWise System Operations > System Preferences*.
- 2 On the *Admin Preferences* tab, select *Display DirXML Warnings*.
- 3 Click *OK*.

The GroupWise® Software Developer Kit provides tools for customizing GroupWise to the specific needs of your organization. It includes the following components:

- ♦ **WebAccess Customization:** Lets you modify the WebAccess client HTML source files to include your own graphics or company information. You can also enhance the WebAccess client by creating additional calendar views. For more information, see [GroupWise WebAccess Customization \(http://developer.novell.com/wiki/index.php/GroupWise_WebAccess_Customization\)](http://developer.novell.com/wiki/index.php/GroupWise_WebAccess_Customization).
- ♦ **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 \(http://developer.novell.com/wiki/index.php/GroupWise_Object_API\)](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 \(http://developer.novell.com/wiki/index.php/GroupWise_Administrative_Object_API\)](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 \(http://developer.novell.com/wiki/index.php/GroupWise_C3PO\)](http://developer.novell.com/wiki/index.php/GroupWise_C3PO).
- ♦ **GroupWise Tokens:** Let 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 \(http://developer.novell.com/wiki/index.php/GroupWise_Tokens\)](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 e-mail auditing. For more information, see [GroupWise Trusted Application API \(http://developer.novell.com/wiki/index.php/GroupWise_Trusted_Application_API\)](http://developer.novell.com/wiki/index.php/GroupWise_Trusted_Application_API).
- ♦ **GroupWise SOAP:** Provides access to GroupWise data, through defined standards, 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) (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) (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) (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) (http://developer.novell.com/wiki/index.php/GroupWise_Controls_for_ActiveX).

In earlier versions of GroupWise[®], WebAccess stored its default user interface files in the following directory on the Web server:

```
tomcat_root/webapps/gw/WEB-INF/classes/com/novell/webaccess/  
templates/frames
```

In GroupWise 7, the user interface files are stored in the `css` subdirectory of templates rather than the frames subdirectory.

If you have exteNd™ portlets configured to use WebAccess, you must copy some of the exteNd template files from the frames directory into the css directory. Refer to the *Identity Manager Accessory Portlet Reference Guide* (http://developer.novell.com/wiki/index.php/GroupWise_Controls_for_ActiveX) to determine which exteNd template files you should copy to the css directory. Do not copy the entire contents of the frames directory into the css directory; this would damage the new GroupWise 7 WebAccess user interface.

In addition, modify the WebAccess URL in the exteNd Portal Preferences from:

```
http://Web_server_address/servlet/webacc
```

to:

```
http://Web_server_address/gw/webacc
```


Heartbeat on Linux



- ♦ Chapter 29, “Introduction to GroupWise 7 and Heartbeat on Linux,” on page 281
- ♦ Chapter 30, “Planning GroupWise in a Heartbeat Cluster,” on page 285
- ♦ Chapter 31, “Setting Up a Heartbeat Cluster,” on page 297
- ♦ Chapter 32, “Setting Up a Domain and a Post Office in a Heartbeat Cluster,” on page 307
- ♦ Chapter 33, “Implementing the Internet Agent in a Heartbeat Cluster,” on page 325
- ♦ Chapter 34, “Implementing WebAccess in a Heartbeat Cluster,” on page 343
- ♦ Chapter 35, “Implementing GroupWise Monitor in a Heartbeat Cluster,” on page 361
- ♦ Chapter 36, “Backing Up a GroupWise System in a Heartbeat Cluster,” on page 363
- ♦ Chapter 37, “Updating a GroupWise System in a Heartbeat Cluster,” on page 365
- ♦ Chapter 38, “Moving an Existing Linux GroupWise 7 System into a Heartbeat Cluster,” on page 367
- ♦ Chapter 39, “Implementing Messenger in a Heartbeat Cluster,” on page 369

Introduction to GroupWise 7 and Heartbeat on Linux

29

Before implementing GroupWise® with Heartbeat on SUSE® Linux Enterprise Server (SLES), make sure that your shared area network (SAN) is properly set up and that all nodes in the cluster can access it.

- ♦ [Section 29.1, “Understanding Heartbeat,” on page 281](#)
- ♦ [Section 29.2, “Meeting System Requirements,” on page 281](#)
- ♦ [Section 29.3, “Installing Heartbeat Packages,” on page 282](#)
- ♦ [Section 29.4, “Configuring Heartbeat,” on page 282](#)
- ♦ [Section 29.5, “Starting Heartbeat,” on page 283](#)
- ♦ [Section 29.6, “Managing Heartbeat,” on page 283](#)
- ♦ [Section 29.7, “Setting Up Shared Storage,” on page 283](#)
- ♦ [Section 29.8, “Setting Up Node Fencing or Resource Fencing,” on page 284](#)
- ♦ [Section 29.9, “What’s Next,” on page 284](#)

29.1 Understanding Heartbeat

Make sure you have a solid understanding of the environment into which you are installing GroupWise. Heartbeat is an open source project that provides clustering services on Linux. The following Web sites provide helpful information:

- ♦ [The High-Availability Linux Project \(http://www.linux-ha.org\)](http://www.linux-ha.org)
- ♦ [Heartbeat \(http://www.linux-ha.org/Heartbeat\)](http://www.linux-ha.org/Heartbeat)
- ♦ [High Availability Management Client \(http://www.linux-ha.org/GuiGuide\)](http://www.linux-ha.org/GuiGuide)
- ♦ [High Availability Admin Tools \(http://www.linux-ha.org/v2/AdminTools\)](http://www.linux-ha.org/v2/AdminTools)
- ♦ [Open Cluster Framework \(OCF\) Resource Agent \(http://www.linux-ha.org/OCFResourceAgent\)](http://www.linux-ha.org/OCFResourceAgent)

29.2 Meeting System Requirements

The following software is required to set up a GroupWise system in a Heartbeat cluster:

- ♦ SUSE Linux Enterprise Server 10, plus the latest update
- ♦ Heartbeat 2, plus the latest update
- ♦ GroupWise 7, plus Support Pack 2 or later
- ♦ Novell® eDirectory™ 8.7 or later, plus the latest Support Pack
- ♦ ConsoleOne® 1.3.6 or later

For a complete list of GroupWise requirements, see “[GroupWise System Requirements](#)” in “[Installation](#)” in the *GroupWise 7 Installation Guide*.

29.3 Installing Heartbeat Packages

The following Heartbeat 3 packages are required for use with GroupWise:

- ♦ `heartbeat-2.0.7-1.2` (<http://www.novell.com/products/linuxpackages/suselinux/heartbeat.html>)
- ♦ `heartbeat-pils-2.0.7-1.2` (<http://www.novell.com/products/linuxpackages/suselinux/heartbeat-pils.html>)
- ♦ `heartbeat-stonith-2.0.7-1.2` (<http://www.novell.com/products/linuxpackages/suselinux/heartbeat-stonith.html>)

These packages are not included in a default SLES 10 installation. You must select them manually, then the SLES 10 installation process installs them for you.

29.4 Configuring Heartbeat

Set up Heartbeat 2 on each node in the cluster by following the instructions provided in *Installing a Heartbeat 2 Cluster Using YaST in the SUSE Linux Enterprise Server 10 Installation and Administration Guide* (http://www.novell.com/documentation/sles10/pdfdoc/sles_admin/sles_admin.pdf). The setup process creates basic Heartbeat configuration files:

- ♦ `/etc/ha.d/ha.cf` (<http://www.linux-ha.org/ha.cf>)

This is the main Heartbeat configuration file. The default file provides the names of the nodes in the cluster, sets up communication among them, and performs other basic cluster setup. For convenience and for best performance with GroupWise, add the following lines to the `ha.cf` file:

```
use_logd yes
keepalive 2
deadtime 30
warntime 10
initdead 120
```

To configure the `logd` daemon referenced in the first line above, copy the `/usr/share/doc/packages/heartbeat/ha_logd.cf` file to `/etc/ha.d`. Remove the comment character (`#`) from the `debugfile` and `logfile` lines. If desired, change the directories where log files are stored.

In order for the nodes to communicate with each other, you must configure each node so that the Heartbeat UDP port of 694 is open. You must copy the customized `ha.cf` to each node in the cluster.

- ♦ `/etc/ha.d/authkeys` (<http://www.linux-ha.org/authkeys>)

This file provides an authentication string and must not be readable by any user other than `root`. After providing the authentication string as `root`, set the mode to 0400 or 0600. You must copy the `authkeys` file to each node in the cluster.

- ♦ `/var/lib/heartbeat/crm/cib.xml` (<http://www.linux-ha.org/ClusterResourceManager/DTD1.0/Annotated>)

This cluster resource manager (CRM) file is created automatically by Heartbeat when you use the HA Management Client to create cluster resources, resource groups, constraints, etc., as described in [Chapter 31, “Setting Up a Heartbeat Cluster,” on page 297](#). Do not manually modify the `cib.xml` file after starting Heartbeat.

29.5 Starting Heartbeat

The Heartbeat packages must be installed and the configuration files set up on each node in the cluster. Then use the following command to start Heartbeat on each node in the cluster:

```
rheartbeat start
```

To make sure that Heartbeat starts each time the node restarts, use the following command:

```
chkconfig heartbeat on
```

Two log files provide status information about the functioning of the Heartbeat cluster and can be viewed by using the following commands:

```
tail -f /var/log/ha-log  
tail -f /var/log/ha-debug
```

29.6 Managing Heartbeat

Two types of programs are available for managing your Heartbeat cluster:

- ♦ **Command line utilities** (<http://www.linux-ha.org/v2/AdminTools>)
 - ♦ `crm_mon` provides an overview of the cluster
 - ♦ `crm_resource` views, modifies, and migrates resources (such as GroupWise agents)
 - ♦ `crm_standby` migrates all resources from a node, such as when maintenance is required
 - ♦ `cibadmin` lets you modify the configuration information in the `cib.xml` file
- ♦ **GUI HA Management Client** (`/usr/lib/heartbeat/haclient.py`) (<http://www.linux-ha.org/GuiGuide>)

When you run this program, you are prompted to log in to the cluster as the `hacluster` user or you can be a member of the `haclient` group in order to log in. By default, the `hacluster` user does not have a password. You must give it one in order to run the HA Management Client.

29.7 Setting Up Shared Storage

After your Heartbeat cluster is installed, configured, and running smoothly, you are ready to set up partitions on your shared storage device for use with GroupWise. You can use the Enterprise Volume Management System (EVMS) to create partitions for file systems where GroupWise domains and post offices will be located. EVMS is an open source project that provides volume management services on Linux. These Web sites provide helpful information:

- ♦ **Enterprise Volume Management System** (<http://evms.sourceforge.net>)
- ♦ **EVMS User Guide** (http://evms.sourceforge.net/user_guide)

Use the following command to start the EVMS Administration Utility:

```
evmsgui
```

Do the following to prepare partitions for use by GroupWise:

- ◆ Delete any existing partitions. They are not the correct type for use with GroupWise.
- ◆ Create private cluster container partitions. These partitions are cluster aware and prevent two nodes from being active on the same partition at the same time.

29.8 Setting Up Node Fencing or Resource Fencing

Although private cluster container partitions protect against two nodes being active on the same partition, there are circumstances where the protection they offer is insufficient. For example, if each of two nodes each think that the other node is dead, they might both write to the same partition and cause damage to data on the shared partition.

To protect against the data damage that can be caused by two nodes writing to the same partition at the same time, you should implement node fencing or resource fencing in your cluster. Node fencing locks resources away from a node whose status is uncertain. Resource fencing ensures exclusive access to a given resource and provides finer granularity, so that a node with uncertain status can still be safely used under certain circumstances. In addition to node fencing and resource fencing, some hardware provides self-fencing resources that take care of the problem for you.

There are numerous ways to implement fencing. The following Web sites offer additional information about fencing:

- ◆ [Split-Brain \(http://linux-ha.org/SplitBrain\)](http://linux-ha.org/SplitBrain)
- ◆ [Fencing \(http://linux-ha.org/fencing\)](http://linux-ha.org/fencing)
- ◆ [Node Fencing \(STONITH\) \(http://linux-ha.org/STONITH\)](http://linux-ha.org/STONITH)
- ◆ [Resource Fencing \(http://linux-ha.org/ResourceFencing\)](http://linux-ha.org/ResourceFencing)
- ◆ [Self-Fencing Resources \(http://linux-ha.org/SelfFencingResource\)](http://linux-ha.org/SelfFencingResource)

29.9 What's Next

Now that you have set up your Heartbeat cluster, you are ready to start installing your GroupWise system.

- ◆ [Chapter 32, “Setting Up a Domain and a Post Office in a Heartbeat Cluster,” on page 307](#)
- ◆ [Chapter 33, “Implementing the Internet Agent in a Heartbeat Cluster,” on page 325](#)
- ◆ [Chapter 34, “Implementing WebAccess in a Heartbeat Cluster,” on page 343](#)

Planning GroupWise in a Heartbeat Cluster

30

The majority of this part of the *GroupWise 7 Interoperability Guide* (Chapter 30, “Planning GroupWise in a Heartbeat Cluster,” on page 285 through Chapter 36, “Backing Up a GroupWise System in a Heartbeat Cluster,” on page 363) is designed for those who are creating a new GroupWise® system, or at least new domains and post offices, in a Heartbeat cluster on Linux.

If you already have an existing GroupWise 7 system on Linux and need to configure it to work in a newly installed Heartbeat cluster, see Chapter 38, “Moving an Existing Linux GroupWise 7 System into a Heartbeat Cluster,” on page 367.

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 7 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 Heartbeat cluster:

- ♦ A new GroupWise system consisting of the primary domain and the initial post office
- ♦ A new secondary domain
- ♦ A new post office
- ♦ GroupWise Message Transfer Agent (MTA) and Post Office Agent (POA)

The “Heartbeat Clustering Worksheet” on page 293 and the “System Clustering Worksheet” on page 140 list the information you need as you set up GroupWise in a Heartbeat cluster. You should print the worksheets and fill them out as you complete the tasks listed below:

- ♦ Section 30.1, “Planning Your Heartbeat Cluster for GroupWise,” on page 285
- ♦ Section 30.2, “Planning a Software Distribution Directory for a Cluster,” on page 288
- ♦ Section 30.3, “Planning a New Clustered Domain,” on page 288
- ♦ Section 30.4, “Planning a New Clustered Post Office,” on page 289
- ♦ Section 30.5, “Planning a New Library for a Clustered Post Office,” on page 290
- ♦ Section 30.6, “Considering How to Install and Configure the Linux Agents in a Cluster,” on page 290
- ♦ Section 30.7, “GroupWise Clustering Worksheets,” on page 293

After you have completed the tasks and filled out the “Heartbeat Clustering Worksheet” on page 293 and the “System Clustering Worksheet” on page 140, you are ready to continue with Chapter 31, “Setting Up a Heartbeat Cluster,” on page 297.

30.1 Planning Your Heartbeat Cluster for GroupWise

- ♦ Section 30.1.1, “Recording Heartbeat Cluster Information,” on page 286

- ◆ Section 30.1.2, “Planning GroupWise Cluster Organization,” on page 286
- ◆ Section 30.1.3, “Planning Cluster Resources and Resource Groups,” on page 287
- ◆ Section 30.1.4, “Planning Secondary IP Addresses,” on page 287

30.1.1 Recording Heartbeat Cluster Information

Now that you have installed Heartbeat on Linux, you need to record key information about the cluster. In a Heartbeat cluster, most items that you add to the cluster are referred to as cluster resources.

HEARTBEAT CLUSTERING WORKSHEET

Under **Item 1: HA Management Client Password**, record the password you want to use to provide access to the GUI HA Management Client for administering your Heartbeat cluster.

Under **Item 2: Nodes in the Cluster**, list the nodes that you want to use for your GroupWise system. Identify each node by its DNS hostname and the IP address that it is physically configured with.

Under **Item 3: EVMS Container Resources**, list the EVMS shared storage partitions that you want to use for GroupWise domains and post offices. An EVMS container resource needs a resource name and the name of an EMVS storage container.

Under **Item 4: File System Resources**, for each EVMS shared storage partition, provide the necessary information about the file system on that partition. List the device name (for example, `/dev/evms/gwssystem/gwvol1`), mount directory (for example, `/mnt/gwssystem`), and file system type (for example, `reiserfs` or `ext3`) for the file system. In addition to the device, the directory, and the file system type, a file system resource also needs a resource name.

NOTE: The `reiserfs` file system type is recommended for storing GroupWise data because it most efficiently handles the large numbers of small files that are stored in domains and post offices.

30.1.2 Planning GroupWise Cluster Organization

The number of nodes that are available in the cluster and the amount of shared storage in your SAN strongly influence where you can place GroupWise domains, post offices, and agents. 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 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 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 e-mail 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 would be 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 Internet Agent 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 POP or IMAP clients by GroupWise users.
- ♦ The WebAccess Agent might or might not require high availability in your GroupWise system, depending on the importance of mailbox access from a Web browser.

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.

30.1.3 Planning Cluster Resources and Resource Groups

In a GroupWise Heartbeat cluster, the required cluster resources for each resource group are an EVMS container, a file system, a secondary IP address, and one or more GroupWise agents. All of these types of resources, as members of a cluster resource group, fail over from one node to another while providing uninterrupted service to GroupWise users. So, for example, the resource group for a GroupWise domain would include the EVMS container and file system where the domain directory will be located and the secondary IP address that will be used by the MTA that services the domain. The domain directory and secondary IP address must be available before the MTA can start.

Typically, Heartbeat resources and groups are named with the object type first, followed by an underscore (`_`), followed by a unique name. For example, if you are creating a new GroupWise system and you want to put the primary domain in the same cluster resource group as the initial post office, you could name the resource group `group_gwprimary`. You could name the domain resource `resource_gwmta` and the post office resource `resource_gwpoa`. In this example, the domain and the post office would fail over together.

You might want the domain and post office together in the same cluster resource group or in different cluster resource groups. If you want the domain and the post office to fail over separately, you need two separate cluster resource groups. For the domain, you could create a resource group named `group_domain1` and a resource named `resource_mta1`. For the post office, you could create a resource group named `group_post1` and a resource named `resource_poal`.

HEARTBEAT CLUSTERING WORKSHEET

Under **Item 6: Cluster Resource Groups for GroupWise Components**, list the cluster resource groups that you want to create for GroupWise components. Initially, you plan one or two resource groups for the primary domain and initial post office. Later, you plan additional resource groups for additional GroupWise components. A GroupWise resource needs a resource name and the name of the domain or post office equivalent to its eDirectory™ object name.

30.1.4 Planning Secondary IP Addresses

Each cluster resource group needs a unique secondary IP address in the cluster. The secondary IP address remains constant regardless of which node the cluster resource group is running on. It is independent of all physical IP addresses in your network.

HEARTBEAT CLUSTERING WORKSHEET

Under **Item 5: IP Address Resources**, list the secondary IP addresses that you want to use for GroupWise agents. The IP address resource also needs a unique name. Initially, you plan one or two secondary IP addresses for the MTA and the POA for the primary domain and initial post office. Later, you plan additional secondary IP addresses for additional resource groups.

30.2 Planning a Software Distribution Directory for a Cluster

During creation of a GroupWise system, 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 the shared storage partition so that you install it only once but it is still always available.

IMPORTANT: You must place 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 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 GroupWise Installation program.

For background information about software distribution directories, see “**Software Directory Management**” in “**System**” in the *GroupWise 7 Administration Guide*.

GROUPWISE CLUSTERING WORKSHEET

Under **Item 1: Software Distribution Directory**, list the full path for the software distribution directory.

30.3 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 cluster, you are creating the primary domain as you complete the tasks in this section. To prepare, review “**Planning Your Basic GroupWise System**”, then print and fill out the “**Basic GroupWise System Worksheet**” in “**Installing a Basic GroupWise System**” in the *GroupWise 7 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 in the cluster. To prepare, review “**Planning a New Domain**”, then print and fill out the “**Domain Worksheet**” in “**Domains**” in the *GroupWise 7 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 shared storage, not on the node where you will be running the GroupWise Installation program. Therefore, each node will mount the domain directory or post office directory when its resource group is active on that node.
- ♦ 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 [item 17](#). If you are filling out the Domain Worksheet, stop with [item 10](#).

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

GROUPWISE CLUSTERING WORKSHEET

Under [Item 2: Domain Name](#), transfer the domain name and domain directory to the GroupWise Clustering Worksheet. Specify the mount point where the domain directory will be mounted to each node when it becomes active.

HEARTBEAT CLUSTERING WORKSHEET

Under [Item 7: Cluster Resources for Domain](#), specify the name of the cluster resource group that will represent the domain (for example, `group_dom1`), the name of the resource for the MTA (for example, `resource_mta2`, and the name of the domain (for example, `Waltham1`).

IMPORTANT: Do not create the new domain until you are instructed to do so in [Chapter 32](#), “[Setting Up a Domain and a Post Office in a Heartbeat Cluster](#),” on page 307.

30.4 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 is planned on the “[Basic GroupWise System Worksheet](#)” in the *GroupWise 7 Installation Guide*. To plan additional new post offices, review “[Planning a New Post Office](#)”, then print and fill out the “[Post Office Worksheet](#)” in “[Post Offices](#)” in the *GroupWise 7 Administration Guide*. When you specify the location for the post office directory, remember that it will be on shared storage, not on the node where you will be running the GroupWise Installation program. Therefore, each node will mount the post office directory when its resource group is active on that node.

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

GROUPWISE CLUSTERING WORKSHEET

Under [Item 4: Post Office Name](#), transfer the post office name and database location to the GroupWise Clustering Worksheet. Specify the mount point where the post office directory will be mounted to each node when it becomes active.

HEARTBEAT CLUSTERING WORKSHEET

Under **Item 8: Cluster Resource for Post Office**, specify the name of the cluster resource group that will represent the post office (for example, `group_po2`), the name of the resource for the POA (for example, `resource_poa2`), and the name of the post office, including its domain (for example, `Waltham1.Sales`).

IMPORTANT: Do not create the new post office until you are instructed to do so in **Chapter 32, “Setting Up a Domain and a Post Office in a Heartbeat Cluster,”** on page 307.

30.5 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 server. 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 directory. Ideally, you should create a dedicated post office just to house the library. This provides easy access for all users.

You can plan a library for a clustered post office by following the standard instructions provided in “**Creating and Managing Libraries**” in the *GroupWise 7 Administration Guide* and filling out the “**Basic Library Worksheet**” or the “**Full-Service Library Worksheet**”. Then provide the library information on the GroupWise Clustering Worksheet.

GROUPWISE CLUSTERING WORKSHEET

Under **Item 6: 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 32, “Setting Up a Domain and a Post Office in a Heartbeat Cluster,”** on page 307.

30.6 Considering 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 and POA in your clustered GroupWise system:

- ♦ **Section 30.6.1, “Recording Secondary IP Addresses for the Agents,”** on page 290
- ♦ **Section 30.6.2, “Determining Appropriate Heartbeat Constraints for GroupWise Cluster Resource Groups,”** on page 291
- ♦ **Section 30.6.3, “Planning the Linux Agent Installation,”** on page 293

30.6.1 Recording Secondary IP Addresses for the Agents

By default, the GroupWise agents listen on all IP addresses that are bound to the server. This means that any time there is a possibility of two of the same type of agent starting on the same node, it is important that each agent use the secondary IP address associated with its cluster resource group, not

the physical IP address of the node. Secondary IP addresses are created by setting up IP address resources, as described in [Section 30.1.4, “Planning Secondary IP Addresses,” on page 287](#). The IP address resource 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 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. For more information, see [“Simplifying Client/Server Access with a GroupWise Name Server”](#) in [“Post Office Agent”](#) in the *GroupWise 7 Administration Guide*.

GROUPWISE CLUSTERING WORKSHEET

Under [Item 3: MTA Network Information](#), select an IP address resource from the Heartbeat Clustering Worksheet ([item 5](#)) and record the secondary IP address associated with the IP address resource on the GroupWise Clustering Worksheet.

Under [Item 5: POA Network Information](#), select an IP address resource from the Heartbeat Clustering Worksheet ([item 5](#)) and record the secondary IP address associated with the IP address resource on the GroupWise Clustering Worksheet.

30.6.2 Determining Appropriate Heartbeat Constraints for GroupWise Cluster Resource Groups

By default, a cluster resource group has all nodes in the cluster available for failover. Only one node at a time can have a particular cluster resource group mounted and active. If a cluster resource group’s initial node fails, the resource group fails over to another node in the cluster.

You should customize the Heartbeat constraints for each GroupWise cluster resource group based on the fan-out-failover principle. When a node fails, its cluster resource groups should not all fail over together to the same node. Instead, the resource groups 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 cluster resource groups 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.

When you create a Heartbeat constraint, you give it a unique ID, associate it with one or more cluster resource groups, and provide information specific to the constraint type. Heartbeat offers three types of constraints that control the failover behavior of cluster resource groups in the cluster. A cluster resource group can use multiple types of constraints to assure the desired failover behavior.

- ◆ [“Place Constraint” on page 292](#)
- ◆ [“Order Constraint” on page 292](#)
- ◆ [“Colocation Constraint” on page 292](#)

Place Constraint

The place constraint controls where nodes in a cluster resource group can fail over. A place constraint includes a score. A score of 100 (the default) means that the resource group to which the constraint is applied should run on the node assigned to the place constraint whenever possible. A score of 50 means that the resource group to which the constraint is applied can fail over to the assigned node when the initial node fails. A score of 0 means that the resource group never runs on that node, perhaps because that node is used for some other program. A score of INFINITY means that the resource group can only run on that node.

HEARTBEAT CLUSTERING WORKSHEET

Under **Item 9: Cluster Resource Group Constraints for Domain**, list the nodes that might need to mount the domain. The MTA might need to run on any node that the domain constraints allow. Therefore, you will install the agent software on all of the nodes where the domain could possibly fail over. List the score for each node.

If you are planning the post office in a different cluster resource group from where the domain is located, under **Item 10: Cluster Resource Group Constraints for Post Office**, list the nodes that might need to mount the post office. The POA might need to run on any node that the post office constraints allow. Therefore, you will install the agent software on all of the nodes where the post office could possibly fail over. List the score for each node.

Order Constraint

The order constraint lets you place a cluster resource group either before or after another resource group. For example, if you have the domain and post office in separate resource groups and you want the domain to start before the post office, you would use an order constraint.

HEARTBEAT CLUSTERING WORKSHEET

Under **Item 9: Cluster Resource Group Constraints for Domain**, list any order constraints regarding the domain.

If you are planning the post office in a different cluster resource group from where the domain is located, under **Item 10: Cluster Resource Group Constraints for Post Office**, list any order constraints regarding the post office.

Colocation Constraint

The colocation constraint lets you designate resource groups that must run together on the same node. For example, you might want the MTA and the Internet Agent for a domain to always run together whenever the domain fails over. Or you might want the MTA and the WebAccess Agent to always run together on the same node. To guarantee that two resource groups always run together, give them a score of INFINITY in the colocation constraint.

HEARTBEAT CLUSTERING WORKSHEET

Under **Item 9: Cluster Resource Group Constraints for Domain**, list any resource groups that you want to run along with the domain.

If you are planning the post office in a different cluster resource group from where the domain is located, under **Item 10: Cluster Resource Group Constraints for Post Office**, list any resource groups that you want to run along with the post office.

30.6.3 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 7 Installation Guide* for each domain and post office for which you will install the Linux MTA or POA.

IMPORTANT: Do not install the Linux agent software until you are instructed to do so in [Chapter 32, “Setting Up a Domain and a Post Office in a Heartbeat Cluster,”](#) on page 307.

Skip to [Chapter 32, “Setting Up a Domain and a Post Office in a Heartbeat Cluster,”](#) on page 307.

30.7 GroupWise Clustering Worksheets

- ◆ [Section 30.7.1, “Heartbeat Clustering Worksheet,”](#) on page 293
- ◆ [Section 30.7.2, “GroupWise Clustering Worksheet,”](#) on page 294

30.7.1 Heartbeat Clustering Worksheet

Item	Explanation
1) HA Management Client Password:	<p>Specify the password you use to access the GUI HA Management Client.</p> <p>For more information, see Section 30.1.1, “Recording Heartbeat Cluster Information,” on page 286.</p>
2) Nodes in Cluster:	<p>List the nodes in your Heartbeat cluster that you will use in your GroupWise system. Identify each node by its hostname and its physical IP address on your network.</p> <p>For more information, see Section 30.1.1, “Recording Heartbeat Cluster Information,” on page 286.</p>
3) EVMS Container Resources:	<p>Specify the names of the EVMS failover container resources that you want to represent partitions of shared storage provided by the SAN in your cluster.</p> <p>For more information, see Section 30.1.1, “Recording Heartbeat Cluster Information,” on page 286.</p>
4) File System Resources:	<p>List the EVMS device name (for example, <code>/dev/evms/gwssystem/gwvol1</code>), the mount point directory (<code>/mnt</code>, <code>/gwssystem</code>, etc.), and file system type (<code>reiserfs</code>, <code>ext3</code>, etc.), for each file system.</p> <p>For more information, see Section 30.1.1, “Recording Heartbeat Cluster Information,” on page 286.</p>

Item	Explanation
5) IP Address Resources:	<p>List the IP address resources and secondary IP addresses that will be assigned to GroupWise components.</p> <p>For more information, see Section 30.1.4, “Planning Secondary IP Addresses,” on page 287.</p>
6) Cluster Resource Groups for GroupWise Components	<p>If you want your domain and post office to fail over together, you need one cluster resource group. If you want them to fail over separately, you need two cluster resource groups.</p> <p>For more information, see Section 30.1.3, “Planning Cluster Resources and Resource Groups,” on page 287.</p>
7) Cluster Resource for Domain:	Specify the name of the cluster resource group where the GroupWise domain will reside, along with the name of the domain.
Domain Name:	
Post Office in Same Resource Group as Domain?	For more information, see Section 13.3, “Planning a New Clustered Domain,” on page 136.
<ul style="list-style-type: none"> ◆ Yes ◆ No 	
8) Cluster Resource for Post Office:	Specify the name of the cluster resource group where the GroupWise post office will reside, along with the name of the post office.
Post Office Name:	
	For more information, see Section 13.4, “Planning a New Clustered Post Office,” on page 137.
9) Cluster Resource Group Constraints for Domain:	Specify the constraints you want to put on the domain resource group to control what nodes it can run on.
<ul style="list-style-type: none"> ◆ Place ◆ Order ◆ Colocation 	For more information, see Section 13.6.2, “Determining Appropriate Failover Lists for the Agents,” on page 138.
10) Cluster Resource Constraints for Post Office:	Specify the constraints you want to put on the post office resource group to control what nodes it can run on.
<ul style="list-style-type: none"> ◆ Place ◆ Order ◆ Colocation 	For more information, see Section 13.6.2, “Determining Appropriate Failover Lists for the Agents,” on page 138.

30.7.2 GroupWise Clustering Worksheet

Item	Explanation
1) Software Distribution Directory:	<p>Specify the full path of the directory where the GroupWise software distribution directory will reside.</p> <p>For more information, see Section 13.2, “Planning a Clustered Software Distribution Directory,” on page 135.</p>

Item	Explanation
2) Domain Name: Domain Directory: Mount Point:	Specify the name for the domain. Specify the directory on the resource group partition where you want to create the new domain. Specify the mount point where each node will mount the domain directory. For more information, see Section 13.3, “Planning a New Clustered Domain,” on page 136.
3) MTA Network Information: <ul style="list-style-type: none"> ◆ Secondary IP address ◆ Message transfer port: 7100 ◆ HTTP port: 7180 	Select a secondary IP address and its cluster resource from Item 5 on the “Heartbeat Clustering Worksheet” on page 293. You can use the default port numbers if they are not already used on your system.
4) Post Office Name: Post Office Directory: Mount Point:	Specify the name for the post office. Specify the directory on the resource group partition where you want to create the post office. Specify the mount point where each node will mount the post office directory. For more information, see Section 13.4, “Planning a New Clustered Post Office,” on page 137.
5) POA Network Information: <ul style="list-style-type: none"> ◆ Secondary IP address ◆ Client/server port: 1677 ◆ Message transfer port: 7101 ◆ HTTP port: 7181 	Select an IP address resource and its secondary IP address from Item 5 on the “Heartbeat Clustering Worksheet” on page 293. You can use the default port numbers if they are not already used on your system.
6) Document Storage Area Location: <ul style="list-style-type: none"> ◆ 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 13.5, “Planning a New Library for a Clustered Post Office,” on page 137.

Setting Up a Heartbeat Cluster

31

You use the HA Management Client program to configure your Heartbeat cluster.

- ♦ Section 31.1, “Starting the HA Management Client,” on page 297
- ♦ Section 31.2, “Creating a Cluster Resource Group,” on page 298
- ♦ Section 31.3, “Creating Native Heartbeat Resources,” on page 299
- ♦ Section 31.4, “Creating Heartbeat Constraints,” on page 302
- ♦ Section 31.5, “Starting the Native Heartbeat Resources,” on page 305

31.1 Starting the HA Management Client

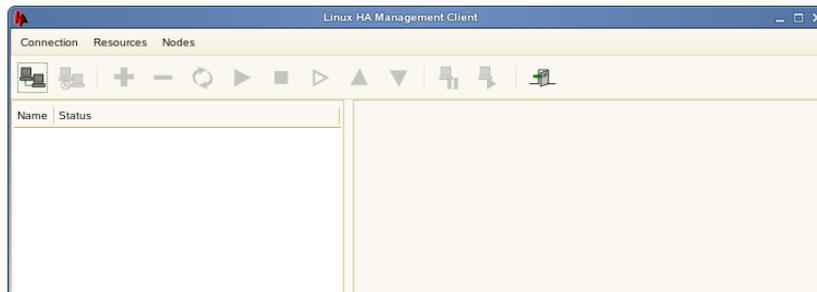
- 1 On a node in the cluster, become the `hacluster` user in a terminal window.
- 2 If you have not yet set a password for the `hacluster` user, you need to set one ([Heartbeat Clustering Worksheet item 1](#)):
 - 2a Enter the following command to set a password:

```
passwd hacluster
```

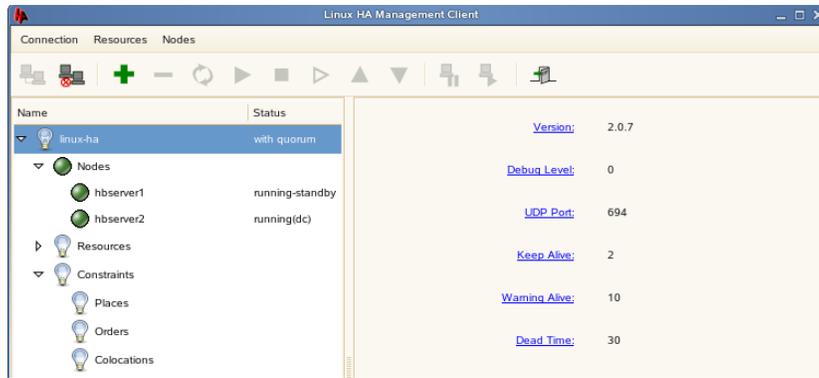
- 2b Enter the desired password.

- 3 Enter the following command to start the HA Management Client:

```
/usr/lib/heartbeat/haclient.py
```



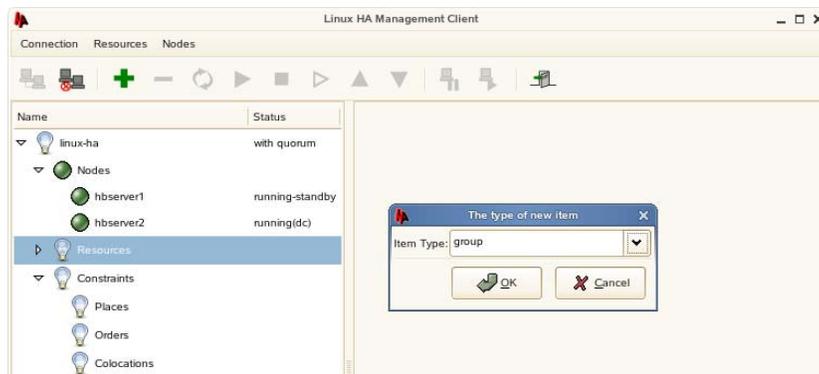
- 4 Click *Login to Cluster*, type the password you established in [Step 2](#), then click *OK* to display the main HA Management Client window.



- 5 Make sure that all nodes ([Heartbeat Clustering Worksheet item 2](#)) are listed.
- 6 Continue with [Creating a Cluster Resource Group](#).

31.2 Creating a Cluster Resource Group

- 1 In the [HA Management Client](#), select *Resources*, then click *Add New Item*.



- 2 In the *Item Type* drop-down list, select *group*, then click *OK*.



- 3 In the *ID* field, specify the name of the GroupWise resource group ([Heartbeat Clustering Worksheet item 6](#)).
- 4 Repeat [Step 1](#) through [Step 3](#) for the cluster resource groups you listed in [Section 30.1.3, “Planning Cluster Resources and Resource Groups,”](#) on page 287.
- 5 Continue with [Creating Native Heartbeat Resources](#).

31.3 Creating Native Heartbeat Resources

When you create a cluster resource group for use with GroupWise, the group must always include an EVMS container resource, a file system resource, and an IP address resource. These are considered native Heartbeat resources, because all Heartbeat clusters have them.

Internally, GroupWise domains and post offices are directory structures containing files and databases that must be housed in a file system on shared storage. An EVMS container resource provides the mount point where each node mounts the shared storage where the domain and post office are located. Each GroupWise agent requires a unique IP address in order to run.

As opposed to native resources, GroupWise resources for domains and post offices must be created after you install the GroupWise software, so that the `GroupWise` script is available to enhance the capabilities of the HA Management Client. By creating native resources before you install the GroupWise software, you can use them to conveniently fail over the nodes as you install the software on multiple nodes.

- ◆ [Section 31.3.1, “EVMS Container Resources,” on page 299](#)
- ◆ [Section 31.3.2, “File System Resources,” on page 300](#)
- ◆ [Section 31.3.3, “IP Address Resources,” on page 301](#)

For more information about Heartbeat resources, consult the Heartbeat documentation that is available on the Web.

31.3.1 EVMS Container Resources

In the [HA Management Client](#):

- 1 Select a resource group, then click *Add New Item*.
- 2 In the *Item Type* field, select *native*, then click *OK*.

Name	Class/Provider	Description
aaeventd	lsb	AppArmor Notification and Reporting
acpid	lsb	Listen and dispatch ACPI events from the kernel
alsasound	lsb	alsasound
apache	ocf/heartbeat	Apache web server
apache	heartbeat	apache

Name	Value	Description
target_role	stopped	press "Default" or "Start" button in toolbar/menu to start the resource

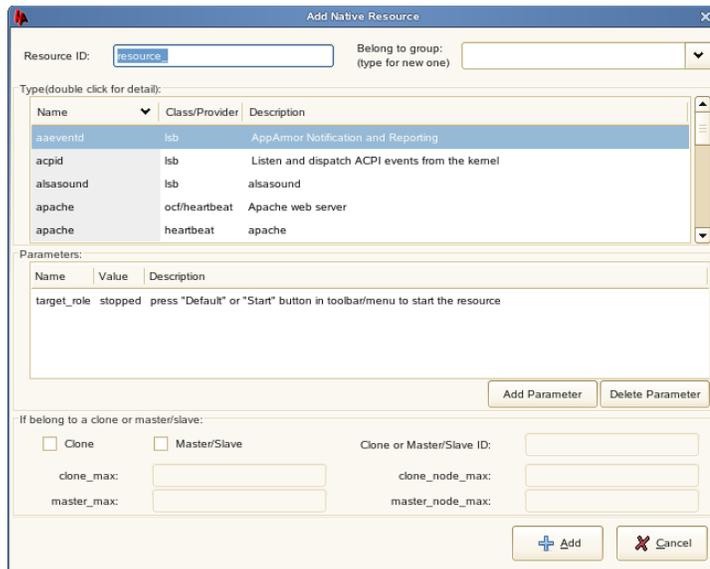
- 3 In the *Resource ID* field, specify a unique name for the EVMS container resource ([Heartbeat Clustering Worksheet item 3](#)).

- 4 In the *Belong to Group* field, select the cluster resource group where you want to create the EVMS contain resource.
- 5 In the *Type* list, select `evms_failover`.
- 6 In the *Parameters* list, specify the EVMS container name for the shared storage:
 - 6a Click *Add Parameter*.
 - 6b In the *Name* drop-down list, select `I`.
 - 6c In the *Value* field, specify the EVMS container name ([Heartbeat Clustering Worksheet item 3](#)), then click *OK*
- 7 Click *Add* to create the EVMS container resource in the resource group.
- 8 Repeat [Step 1](#) through [Step 7](#) to create an EVMS container resource in each cluster resource group.
- 9 Continue with [File System Resources](#).

31.3.2 File System Resources

In the [HA Management Client](#):

- 1 Select a resource group, then click *Add New Item*.
- 2 In the *Item Type* field, select `native`, then click *OK*.



- 3 In the *Resource ID* field, specify a unique name for the file system resource ([Heartbeat Clustering Worksheet item 4](#)).
- 4 In the *Belong to Group* field, select the resource group where you want to create the file system resource.
- 5 In the *Type* list, select *Filesystem (ocf/heartbeat)*.
- 6 In the *Parameters* list:
 - 6a Set the device parameter to the shared storage device you have selected for the GroupWise component ([Heartbeat Clustering Worksheet item 4](#)).

- 6b Set the directory parameter to the mount point ([Heartbeat Clustering Worksheet item 4](#)).
- 6c Set the fstype parameter to the type of file system on the device ([Heartbeat Clustering Worksheet item 4](#)).
- 7 Click *Add* to create the file system resource.
- 8 Repeat [Step 1](#) through [Step 7](#) to create a file system resources in each cluster resource group.

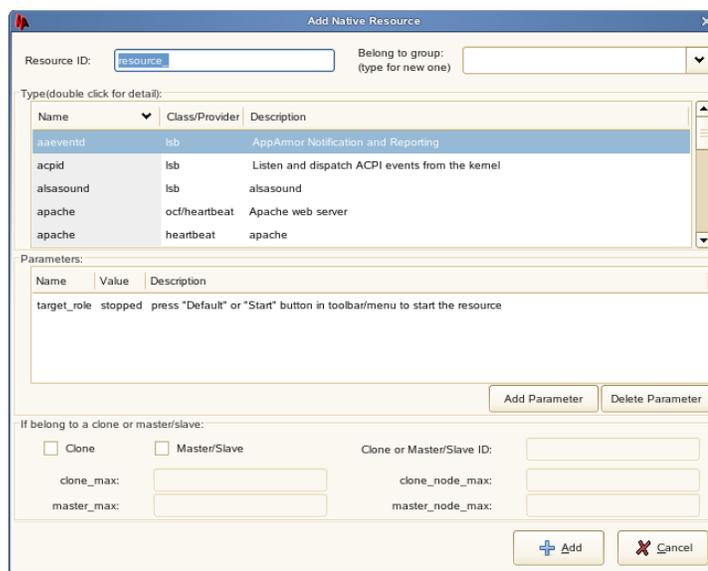
IMPORTANT: To protect the data on the file systems from damage caused by two nodes mistakenly writing to the same file system at the same time, see [Section 29.8, “Setting Up Node Fencing or Resource Fencing,”](#) on page 284.

- 9 Continue with [IP Address Resources](#).

31.3.3 IP Address Resources

In the [HA Management Client](#):

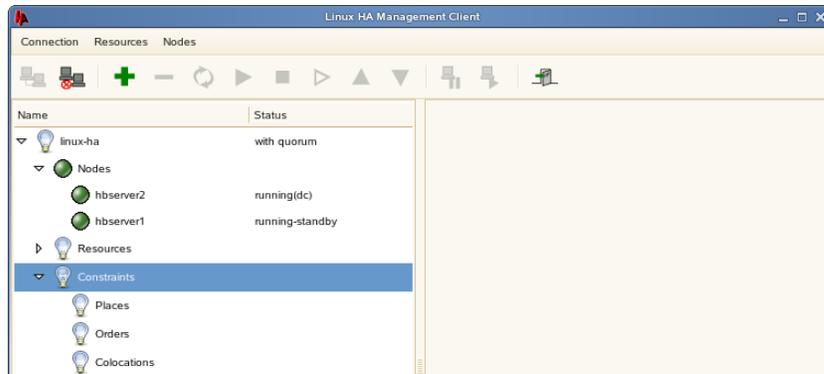
- 1 Select a new cluster resource group, then click *Add New Item*.
- 2 In the *Item Type* field, select *native*, then click *OK*.



- 3 In the *Resource ID* field, specify a unique name for the IP address resource ([Heartbeat Clustering Worksheet item 5](#)).
- 4 In the *Belong to Group* field, select the resource group where you want to create the IP address resource.
- 5 In the *Type* list, select *IPaddr (ocf/heartbeat)*.
- 6 In the *Parameters* list, set the *ip* parameter to the secondary IP address you have selected for the GroupWise agent ([GroupWise Clustering Worksheet item 3](#) for the MTA or [item 5](#) for the POA).
- 7 Click *Add* to create the IP address resource in the cluster resource group.
- 8 Repeat [Step 1](#) through [Step 7](#) to create an IP address resource in each cluster resource group.
- 9 Continue with [Creating Heartbeat Constraints](#).

31.4 Creating Heartbeat Constraints

Heartbeat constraints control which nodes in the cluster each resource group can run on and when. In the HA Management Client, constraints are located below nodes and resources.



The following sections provide some basics to help you get your GroupWise system set up in a Heartbeat cluster.

- ◆ [Section 31.4.1, “Place Constraints,” on page 302](#)
- ◆ [Section 31.4.2, “Order Constraints,” on page 303](#)
- ◆ [Section 31.4.3, “Colocation Constraints,” on page 304](#)

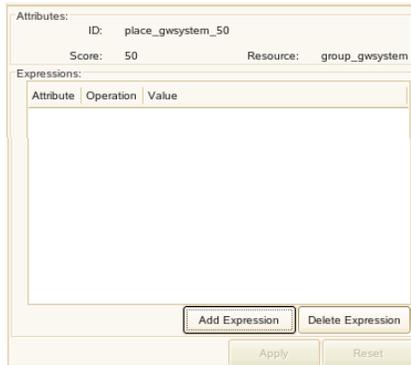
For more information about Heartbeat constraints, consult the Heartbeat documentation that is available on the Web.

31.4.1 Place Constraints

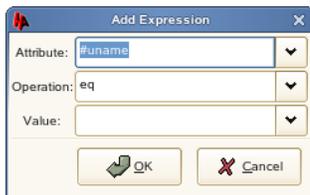
- 1 In the **HA Management Client**, select *Places*, then click *Add New Item*.
- 2 In the *Item Type* drop-down list, select *place*, then click *OK*.



- 3 Specify the unique ID for the place constraint, select the resource group that you want it to apply to, set the score ([Heartbeat Clustering Worksheet item 9](#) for a domain and [item 10](#) for a post office), then click *OK* to add the place constraint.



- 4 Click *Add Expression* to add attributes to the place constraint.



- 5 Accept the defaults of *#uname* (name of the node) and *eq* (equals), in the *Value* field, select the hostname of the node that you want associated with this place constraint ([Heartbeat Clustering Worksheet item 9](#) for a domain and [item 10](#) for a post office), then click *OK*.

For example, if the place constraint has a score of 100, you would select the node where you want the cluster resource group to run on a regular basis. If the place constraint has a score of 50, you would select the node where you want the cluster resource group to fail over if its initial node fails.

- 6 Repeat [Step 1](#) through [Step 5](#) to create additional place constraints as needed.
- 7 Continue with [Order Constraints](#).

31.4.2 Order Constraints

- 1 In the [HA Management Client](#), select *Orders*, then click *Add New Item*.
- 2 In the *Item Type* drop-down list, select *order*, then click *OK*.



- 3 Specify the unique ID for the order constraint, select a node, select a type of order (*before* or *after*), select another node ([Heartbeat Clustering Worksheet item 9](#) for a domain and [item 10](#) for a post office), then click *OK*.

ID: order_dompo

From: group_gwdomain

Type: before

To: group_gwpost

Apply Reset

You can change the attributes of the order constraint at any time if the needs of your cluster change.

- 4 Repeat **Step 1** through **Step 3** for each order constraint that you need in your GroupWise system.
- 5 Continue with **Colocation Constraints**.

31.4.3 Colocation Constraints

- 1 In the **HA Management Client**, select *Colocations*, then click *Add New Item*.
- 2 In the *Item Type* drop-down list, select *colocation*, then click *OK*.

Add Colocation Constraint

ID: colocation

From: group_gwssystem

To: group_gwssystem

Score: INFINITY

OK Cancel

- 3 Specify the unique ID for the colocation constraint, select a resource group, select another resource group that you always want to run on the same node as the first resource group (**Heartbeat Clustering Worksheet item 9** for a domain and **item 10** for a post office), select *INFINITY* to indicate that they must always run together, then click *OK*.

ID: colocation_syspridom

From: group_gwssystem

To: group_gwdomain

Score: INFINITY

Apply Reset

You can change the attributes of the colocation constraint at any time if the needs of your cluster change.

- 4 Repeat **Step 1** through **Step 3** for each colocation constraint that you need in your GroupWise system
- 5 Continue with **Starting the Native Heartbeat Resources**.

31.5 Starting the Native Heartbeat Resources

In the **HA Management Client**:

- 1 In each cluster resource group, select the EVMS container resource, then click *Start Resource*.
- 2 Select the file system resource, then click *Start Resource*.
- 3 Select the IP address resource, then click *Start Resource*.

When a resource starts successfully, the Status column indicates that it is running.

- 4 When all the resources you have created so far are running, continue with **Setting Up a Domain and a Post Office in a Heartbeat Cluster**.

Setting Up a Domain and a Post Office in a Heartbeat Cluster

32

You should have already reviewed [Chapter 30, “Planning GroupWise in a Heartbeat Cluster,”](#) on page 285 and filled out the [“Heartbeat Clustering Worksheet”](#) on page 293 and the [“System Clustering Worksheet”](#) on page 140. You should also have set up a basic Heartbeat cluster, as described in [Chapter 31, “Setting Up a Heartbeat Cluster,”](#) on page 297. You are now ready to complete the following tasks to set up GroupWise® in your Heartbeat cluster:

- ♦ [Section 32.1, “Setting Up a New GroupWise System in a Heartbeat Cluster,”](#) on page 307
- ♦ [Section 32.2, “Creating a New Secondary Domain in a Heartbeat Cluster,”](#) on page 308
- ♦ [Section 32.3, “Creating a New Post Office in a Heartbeat Cluster,”](#) on page 309
- ♦ [Section 32.4, “Installing and Configuring the MTA and the POA in a Heartbeat Cluster,”](#) on page 310
- ♦ [Section 32.5, “Testing Your Clustered GroupWise System,”](#) on page 320
- ♦ [Section 32.6, “Managing Your Clustered GroupWise System,”](#) on page 321
- ♦ [Section 32.7, “What’s Next,”](#) on page 323

32.1 Setting Up a New GroupWise System in a Heartbeat Cluster

The GroupWise Installation program walks you through setting up the primary domain and an initial post office owned by the primary domain. You might be creating your primary domain and initial post office in the same cluster resource group or in two different cluster resource groups. After you have created the primary domain and the 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.

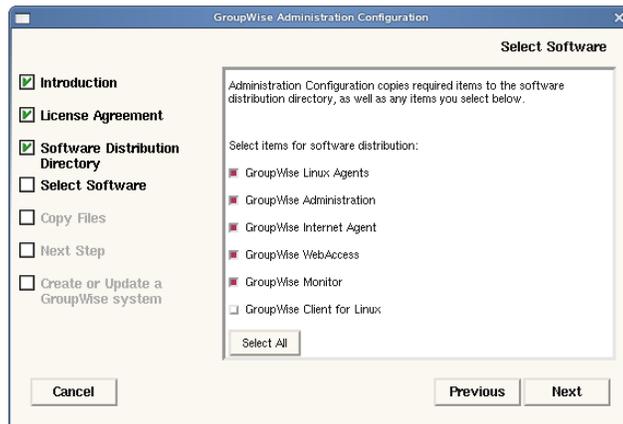
To set up the primary domain and initial post office for a new GroupWise system in a Heartbeat cluster:

- 1 Start on the initial node where you want the domain directory to be mounted ([Heartbeat Clustering Worksheet item 9](#)).
- 2 Make sure that ConsoleOne® is installed on the node.
If necessary, you can download ConsoleOne for Linux from the [Novell Product Downloads site \(http://download.novell.com\)](http://download.novell.com). ConsoleOne is always installed in `/usr/ConsoleOne`.
- 3 If necessary, mount the partition where you want to create the GroupWise software distribution directory ([GroupWise Clustering Worksheet item 1](#)).
- 4 In the **HA Management Client**, start the EVMS container resource and the file system resource for the domain cluster resource group in order to mount the shared storage partition for the domain ([GroupWise Clustering Worksheet item 2](#)) and, if needed, mount the shared storage partition for the post office ([GroupWise Clustering Worksheet item 4](#)), where the primary domain and the initial post office for your new GroupWise system will be created.

- From the *GroupWise 7 Administrator for Linux* CD, run the GroupWise Installation program (install), as described in “Starting the GroupWise Installation Program on Linux” in “Installing a Basic GroupWise System” in the *GroupWise 7 Installation Guide*.

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

- When you set up the software distribution directory, copy all of 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 primary domain and the initial post office.

- From the GroupWise 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 7 Installation Guide*.

- When providing the MTA and POA network address information, use their secondary IP addresses ([GroupWise Clustering Worksheet item 3](#) for the MTA and [item 5](#) for the POA).

The information you provide is used to configure the MTA and POA eDirectory™ 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.

- 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 from the *GroupWise 7 Administrator for Linux* CD.

IMPORTANT: You will install the agent software from the software distribution directory that you just created, not from the CD.

- Skip to “Running the GroupWise Installation Program on the Preferred Node” on page 147.

32.2 Creating a New Secondary Domain in a Heartbeat Cluster

After you have set up the primary domain and initial post office, as described in [Chapter 32, “Setting Up a Domain and a Post Office in a Heartbeat Cluster,”](#) on page 307, you can create additional secondary domains in your GroupWise system as needed.

To create a new secondary domain in a clustering environment:

- 1 In the **HA Management Client**, start the EVMS container resource and the file system resource in order to mount the shared storage 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 7 Administration Guide*.
- 3 Create the new domain, following the steps provided in “**Creating the New Domain**” in “**Domains**” in the *GroupWise 7 Administration Guide*.

Use the Domain Worksheet you filled out in **Section 13.3, “Planning a New Clustered Domain,”** on page 136 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 7 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 7 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 **Section 14.4, “Installing and Configuring the MTA and the POA in a Cluster,”** on page 146 to install the MTA software for the new domain.

32.3 Creating a New Post Office in a Heartbeat Cluster

You can create a new post office in the same cluster resource group where its domain resides or in a separate cluster resource group. If the post office and its domain are in the same cluster resource group, they fail over together. If they are in separate cluster resource groups, they fail over separately.

To create a new post office in a clustering environment:

- 1 In the **HA Management Client**, start the EVMS container resource and the file system resource in order to mount the shared storage partition where the domain that will own the new post office is located.
- 2 If necessary, start the EVMS container resource and the file system resource 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 7 Administration Guide*.

- 4 Create the new post office, following the steps provided in “[Creating the New Post Office](#)” in “[Post Offices](#)” in the *GroupWise 7 Administration Guide*.

Use the Post Office Worksheet you filled out in [Section 13.4, “Planning a New Clustered Post Office,”](#) on page 137 to fill in the fields in the Create GroupWise Post Office dialog box.

- 5 Refer to the GroupWise Clustering Worksheet that you filled in during [Section 13.6, “Deciding How to Install and Configure the Linux Agents in a Cluster,”](#) on page 138 for the secondary IP address and port numbers that you need to specify in order to configure the link.
- 6 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 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
- ♦ Enable 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 7 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 7 Administration Guide*.

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 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 7 Administration Guide*, after you have completely finished setting up the clustered post office.
- 11 Continue with [Installing and Configuring the MTA and the POA in a Cluster](#) to install the POA software for the new post office.

32.4 Installing and Configuring the MTA and the POA in a Heartbeat Cluster

By following the instructions in [Section 14.1, “Setting Up a New GroupWise System in a Linux Cluster,”](#) on page 143, you created the primary domain and initial post office in your GroupWise system. You are now ready to install and configure the agents, first on the initial node, and then on additional nodes.

The agents must be installed on each node where the domain might be mounted ([Heartbeat Clustering Worksheet item 9](#)) and where the post office might be mounted ([Heartbeat Clustering Worksheet item 10](#)).

- ♦ [Section 32.4.1, “Running the GroupWise Installation Program on the Initial Node,”](#) on page 311
- ♦ [Section 32.4.2, “Running the GroupWise Installation Program on Subsequent Nodes,”](#) on page 313
- ♦ [Section 32.4.3, “Creating and Configuring GroupWise Cluster Resources,”](#) on page 315
- ♦ [Section 32.4.4, “Testing Your Agent Installation on Each Node,”](#) on page 317
- ♦ [Section 32.4.5, “Changing Agent Paths to Locations on the Shared Storage Partition,”](#) on page 318
- ♦ [Section 32.4.6, “Setting Up New Instances of the Agents without Installing the Agent Software,”](#) on page 319

If you have added a new secondary domain or a new post office to an existing GroupWise cluster resource group, 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 32.4.6, “Setting Up New Instances of the Agents without Installing the Agent Software,”](#) on page 319 instead of running the GroupWise Installation program.

32.4.1 Running the GroupWise Installation Program on the Initial Node

- 1 In the [HA Management Client](#), select a node where you want to run the agents ([Heartbeat Clustering Worksheet item 9](#) for a domain and [item 10](#) for a post office), then click *Make the Node Active*.

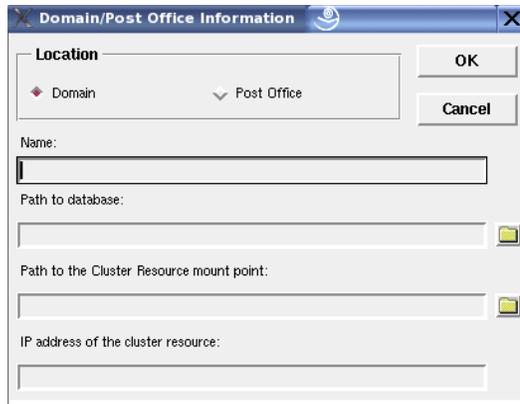
This applies the secondary IP address to the node and mounts the file system where the domain and post office are located, thus setting up the proper environment for an agent installation.

- 2 From the software distribution directory you created in [Step 6](#) in [Section 14.1, “Setting Up a New GroupWise System in a Linux Cluster,”](#) on page 143, start the GroupWise Installation program (`install`).

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



- 3 Install the agent software, following the steps provided in [“Installing the GroupWise Agents on Linux”](#) in [“Installing GroupWise Agents”](#) in the *GroupWise 7 Installation Guide*.
- 4 Configure the Linux agents, paying special attention to the cluster resource information on the Domain/Post Office Information page.



The path to the cluster resource mount point is the mount point for the domain ([GroupWise Clustering Worksheet item 2](#)) or the post office ([GroupWise Clustering Worksheet item 4](#)). The IP address of the cluster resource is the secondary IP address for the MTA ([GroupWise Clustering Worksheet item 3](#)) or the POA ([GroupWise Clustering Worksheet item 5](#)).

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

- ◆ The agent startup files are created in `/mount_point/groupwise/agents/share` on the shared storage partition so that the agents use the same startup files regardless of which node they run on. The `--home` switch includes the mount point and the path to the domain database or the post office database so that the startup file is valid when mounted to each 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 for the cluster resource group 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 node the agents are running on.
- ◆ The `--log` startup switch is set to a location on the shared storage partition (`/mount_point/groupwise/agents/log`) so that agent logging information is written to the same log file regardless of which node the agents run on.
- ◆ The GroupWise High Availability service (GWhA) is automatically configured on the current node and its configuration file (`gwha.conf`) is created in the `/etc/opt/novell/groupwise` directory of the node.
- ◆ The OCF Resource Agent script is installed. The OCF Resource Agent provides support for the GroupWise agents (POA, MTA, Internet Agent, and WebAccess Agent) and the Messenger agents (Messaging Agent and Archive Agent) in the HA Management Client. The script is installed as part of the `novell-groupwise-gwha` package and depends on the High Availability service configuration file (`gwha.conf`) to identify the agents that are being managed by the HA Management Client. The script is named `GroupWise` and is located in the `/usr/lib/ocf/resource.d/novell` directory on the current node.

- ◆ A `clusterimport.conf` file is created in the `gwinst` subdirectory of the software distribution directory from which you ran the GroupWise Installation program, so that the cluster-specific information collected when you configured the agents on the initial 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 For optimum security, configure the agents to run as a non-root user, as described in “Running the Linux GroupWise Agents as a Non-root User” in “Installing GroupWise Agents” in the *GroupWise 7 Installation Guide*.
 - 6 Continue with [Running the GroupWise Installation Program on Subsequent Nodes](#).

32.4.2 Running the GroupWise Installation Program on Subsequent Nodes

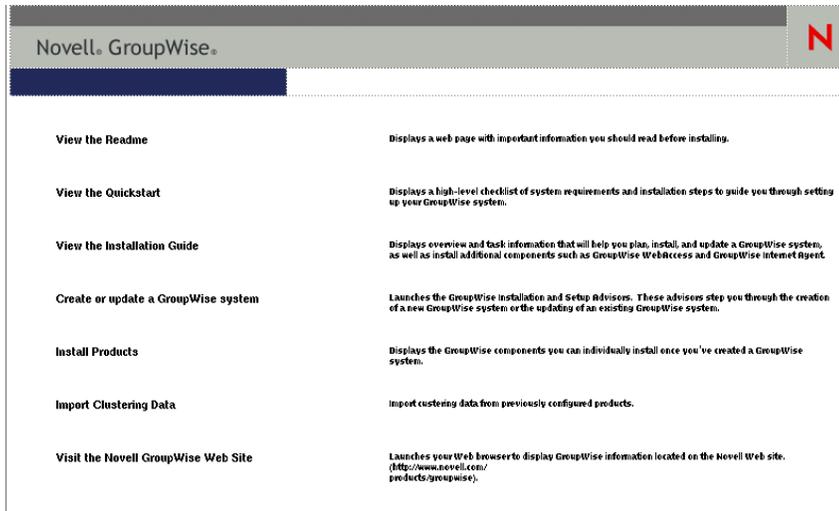
- 1 In the [HA Management Client](#), select a different node where you want to run the agents ([Heartbeat Clustering Worksheet item 9](#) for a domain and [item 10](#) for a post office), then click *Make the Node Active*.

This applies the secondary IP address to this node and mounts the file system where the domain and post office are located, thus setting up the proper environment for an agent installation on the next node.
- 2 From the software distribution directory you created in [Step 6 in Section 14.1, “Setting Up a New GroupWise System in a Linux Cluster,” on page 143](#), start the GroupWise Installation program.

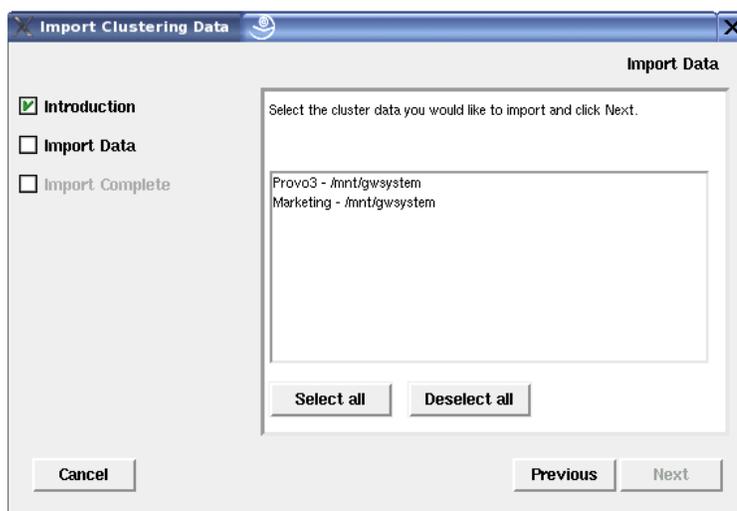
IMPORTANT: You should select the *Configure GroupWise for Clustering* option again.



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 node as usual, but do not use the *Configure* option. You will use the *Import Clustering Data* option instead of the *Configure* option.
- 4 On the main page of the GroupWise 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 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 node. It is configured specifically for the agents you just selected.

- ♦ The GroupWise High Availability service (GWAHA) is automatically configured on the current node and its configuration file (`gwha.conf`) is created in the `/etc/opt/novell/groupwise` directory of the node. It is configured specifically for the agents you just selected.
- ♦ The OCF Resource Agent script (`GroupWise`), required by the HA Management Client to support GroupWise resources, is installed in the `/usr/lib/ocf/resource.d/novell` directory on the current node.

Because the agent startup files and log files are located on the shared storage partition with the domain or post office, they do not need to be customized for each node.

- 6 For optimum security, configure the agents to run as a non-root user, as described in “Running the Linux GroupWise Agents as a Non-root User” in “Installing GroupWise Agents” in the *GroupWise 7 Installation Guide*.
- 7 Repeat [Step 1](#) through [Step 6](#) for each node where you want the agents to be able to fail over ([Heartbeat Clustering worksheet item 9](#) for a domain and [item 10](#) for a post office).
- 8 When the agent software has been installed on all nodes specified by the Heartbeat constraints ([Heartbeat Clustering Worksheet item 9](#) for a domain and [item 10](#) for a post office), continue with [Creating and Configuring GroupWise Cluster Resources](#).

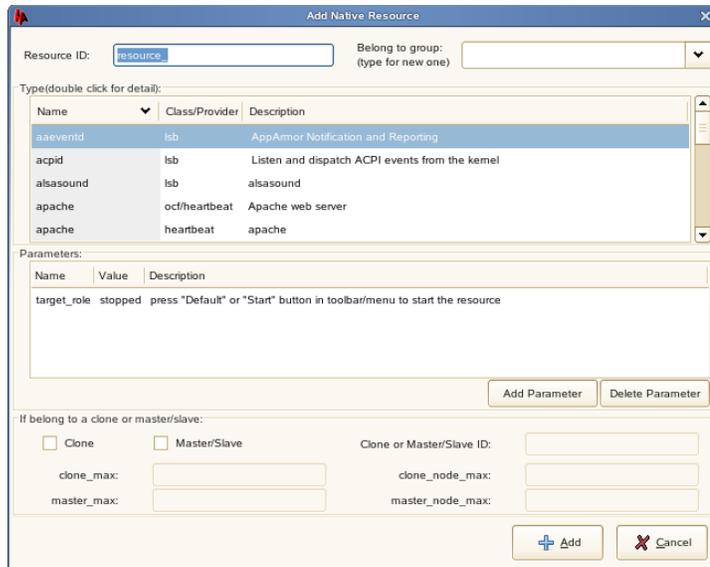
32.4.3 Creating and Configuring GroupWise Cluster Resources

Now that you have installed the GroupWise software, you can create and configure GroupWise resources in the cluster resource groups you have already created in [Chapter 31, “Setting Up a Heartbeat Cluster,”](#) on page 297:

- ♦ “Creating GroupWise Resources” on page 315
- ♦ “Enabling Automatic GroupWise Agent Restart” on page 316
- ♦ “Preventing GroupWise Agent Failback” on page 317

Creating GroupWise Resources

- 1 In the [HA Management Client](#), select a cluster resource group, then click *Add New Item*.
- 2 In the *Item Type* field, select *native*, then click *OK*.



- 3 In the *Resource ID* field, specify a unique name for the GroupWise resource ([Heartbeat Clustering Worksheet item 7](#) for a domain or [item 8](#) for a post office).
- 4 In the *Belong to Group* field, select the cluster resource group where you want to create the GroupWise resource.
- 5 In the *Type* list, select *GroupWise (ocf/novell)*.
The OCF Resource Agent script (*GroupWise*) added this item to the list.
- 6 In the *Parameters* list, set the value of the *object_name* parameter to the identifier for the agent that services the domain or post office ([GroupWise Clustering Worksheet item 2](#) or [item 4](#)).
For a domain, the identifier is the domain name. For a post office, the identifier is the name of the post office, followed by the name of the domain it belongs to (*PostOffice.Domain*).
- 7 Click *Add* to create the GroupWise resource.
You need to add one GroupWise resource for the domain and one for the post office, regardless of whether they are in the same resource group or different resource groups.
- 8 Continue with [Enabling Automatic GroupWise Agent Restart](#).

Enabling Automatic GroupWise Agent Restart

If you want Heartbeat to automatically restart a GroupWise agent that has stopped:

- 1 In the [HA Management Client](#), select each GroupWise resource, then click *Operations*.



- 2 Click *Add Operation*, then select *monitor*.

- 3 Select the time interval for how often you want the GroupWise agent to be polled to see if it is still running.

The GroupWise script causes the agent to restart, using the GroupWise High Availability service. (GWHA)

- 4 Click *OK* to save the monitoring setting.

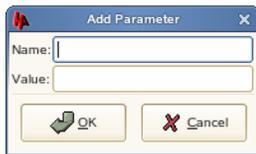
Because Heartbeat can perform such agent monitoring, you do not need to set up GroupWise Monitor to perform this task, as described in “[Configuring the Monitor Agent to Communicate with the High Availability Service](#)” in “[Installing GroupWise Monitor](#)” in the *GroupWise 7 Installation Guide*. Heartbeat and the GroupWise High Availability Service can work together to keep the GroupWise agents running on Linux.

- 5 Continue with [Preventing GroupWise Agent Failback](#).

Preventing GroupWise Agent Failback

If you do not want a resource group to automatically fail back when the initial node is online again:

- 1 In the **HA Management Client**, select each GroupWise resource group, click *Parameters*, then click *Add Parameter*.



- 2 In the *Name* field, specify `resource_stickiness`.
- 3 In the *Value* field, specify `INFINITY`.
- 4 Click *Apply* to save the setting.
- 5 Continue with [Testing Your Agent Installation on Each Node](#).

32.4.4 Testing Your Agent Installation on Each Node

- 1 On the node where the shared storage partition is currently mounted, test the agents 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 7 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 After you can see that the agents start successfully, test them by starting them as daemons, as described in “[Starting the Linux GroupWise Agents as Daemons](#)” in “[Installing GroupWise Agents](#)” in the *GroupWise 7 Administration Guide*.

```
/etc/init.d/grpwise start  
/etc/init.d/grpwise status
```

- 4 Stop the agents.

```
/etc/init.d/grpwise stop  
/etc/init.d/grpwise status
```

- 5 In the **HA Management Client**, activate a different node where you have installed the agent software.
- 6 Repeat **Step 1** through **Step 5** for each node where you have installed the GroupWise agents.
- 7 After you have tested the agents on all nodes manually, start and stop their GroupWise resources in the HA Management Client to make sure that their Heartbeat resources are properly configured.

If you receive the error `Unmanaged on node`, then you have failed to test the agents on one or more nodes. On each node, view the `gwha.conf` file in the `/etc/opt/novell/groupwise` directory. Each section should be preceded by the name of a domain for an MTA or a post office and domain for a POA. If any section is preceded by a startup file variable (for example, `@startup.mta`), then that agent has not yet been initialized by being run manually. To resolve the problem, you can run the agent manually on that node, or you can replace the startup file variable with the appropriate domain and/or post office information, being sure to follow the same format as the other section headings in the `gwha.conf` file.

- 8 When you have installed and tested the agents on all of the nodes where the domain and post office can fail over, continue with **Changing Agent Paths to Locations on GroupWise Partitions**.

32.4.5 Changing Agent Paths to Locations on the Shared Storage Partition

The default locations for some agent files are on the nodes along with the agent software, rather than being located with the domain and post office on a shared storage partition. To avoid having multiple copies of these files in multiple locations, you should set the locations in ConsoleOne.

- ♦ “Setting the MTA Message Log File Path” on page 318
- ♦ “Setting the MTA Certificate and Key File Path” on page 319
- ♦ “Setting the POA Certificate and Key File Path” on page 319

When you have changed the paths, skip to **Section 14.5, “Testing Your Clustered GroupWise System,”** on page 160.

Setting the 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 7 Administration Guide*:

- 1 On the shared storage partition of the cluster resource group 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 MTA Certificate and Key File Path

If you plan to enable SSL for the MTA, as described in “[Securing the Domain with SSL Connections to the MTA](#)” in “[Message Transfer Agent](#)” in the *GroupWise 7 Administration Guide*:

- 1 On the shared storage partition of the cluster resource group where the domain is located, create the directory where you want to store the certificate file 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 7 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 POA Certificate and Key File Path

If you plan to enable SSL for the POA, as described in “[Securing the Post Office with SSL Connections to the POA](#)” in “[Post Office Agent](#)” in the *GroupWise 7 Administration Guide*:

- 1 On the shared storage partition of the cluster resource group where the post office is located, create the directory where you want to store the certificate file 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 7 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*.

NOTE: When the MTA and the POA run in the same cluster resource group, they can use the same certificate file and key file.

- 9 Skip to [Section 14.5, “Testing Your Clustered GroupWise System,”](#) on page 160.

32.4.6 Setting Up New Instances of the Agents without Installing the Agent Software

If you want to run a second instance of an agent that is already installed on a node, you do not need to run the GroupWise Installation program a second time. You can simply create a new startup file for the second instance of the agent.

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 GroupWise Installation program creates agent startup files in `/mount_point/groupwise/agents/share` on the shared storage partition.

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 looking for 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 High Availability Service for the Linux GroupWise Agents](#)” in “[Installing GroupWise Agents](#)” in the *GroupWise 7 Installation Guide*.

- 8 As if you had run the GroupWise Installation program, skip back to [Section 32.4.3, “Creating and Configuring GroupWise Cluster Resources,”](#) on page 315.

32.5 Testing Your Clustered GroupWise System

After you have created and configured the GroupWise cluster resources, you can use the following methods to test whether the GroupWise agents fail over as expected:

- ♦ In the [HA Management Client](#), select a node where a GroupWise agent is running, then click *Make the Node Standby*. The agent should fail over to another node and start running again. This can also be done from the command line using the [crm_standby utility](http://www.linux-ha.org/v2/AdminTools/crm_standby) (http://www.linux-ha.org/v2/AdminTools/crm_standby).
- ♦ Using the [crm_resource utility](http://www.linux-ha.org/v2/AdminTools/crm_resource) (http://www.linux-ha.org/v2/AdminTools/crm_resource), manually migrate a node where a GroupWise agent is running.

If you receive the error `Unmanaged on node`, see [Step 7](#) in [Section 32.4.4, “Testing Your Agent Installation on Each Node,”](#) on page 317 for troubleshooting assistance.

If you have set up your Heartbeat cluster in a test situation first, you can also use the following methods for testing failover:

- ♦ Power off the node where the GroupWise agent is running. The agent should fail over to another node and start running again.
- ♦ Simulate the agent stopping unexpectedly (`kill -9 agent_PID`). If you have configured the GroupWise resource for automatic restart, as described in “[Enabling Automatic GroupWise Agent Restart](#)” on page 316, the agent should start running again on the same node.

WARNING: These last two tests should not be performed in a production system because they can result in damage to GroupWise databases.

32.6 Managing Your Clustered GroupWise System

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

- ♦ [Section 32.6.1, “Updating GroupWise Objects with Cluster-Specific Descriptions,” on page 321](#)
- ♦ [Section 32.6.2, “Knowing What to Expect in MTA and POA Failover Situations,” on page 322](#)

32.6.1 Updating GroupWise Objects with Cluster-Specific Descriptions

After setting up your GroupWise system in a Heartbeat cluster, 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 cluster changes.

- ♦ [“Recording Cluster-Specific Information for a Domain and Its MTA” on page 321](#)
- ♦ [“Recording Cluster-Specific Information for a Post Office and Its POA” on page 322](#)
- ♦ [“Recording Cluster-Specific Information for a Software Distribution Directory” on page 322](#)

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 cluster resource group and shared storage partition where the domain is located.
- 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 associated with the domain’s cluster resource group.

When the Linux MTA runs with a user interface, this information appears on the MTA server console.

- 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, provide a cluster-specific description of the post office, including the cluster resource group and shared storage partition where the post office is located.
- 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 associated with the post office's cluster resource group.

When the Linux POA runs with a user interface, this information appears on the POA server console.

- 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 Continue with [Recording Cluster-Specific Information for a Software Distribution Directory](#).

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 shared storage 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 and POA Failover Situations](#).

32.6.2 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 and then 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.

In comparison to failover, 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.

32.7 What's Next

Now that you have at least one GroupWise domain and post office up and running in a Heartbeat 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 7 Administration Guide*.
- ◆ Setting up the GroupWise client software and helping users to get started using it. See “Client” in the *GroupWise 7 Administration Guide*. Also see the *GroupWise 7 Windows Client User Guide* and the *GroupWise 7 Cross-Platform Client User Guide*.
- ◆ Connecting your clustered GroupWise system to the Internet. See Chapter 33, “Implementing the Internet Agent in a Heartbeat Cluster,” on page 325.
- ◆ Providing access to users’ GroupWise mailboxes from their Web browsers. See Chapter 34, “Implementing WebAccess in a Heartbeat Cluster,” on page 343 and the *GroupWise 7 WebAccess Client User Guide*.
- ◆ Monitoring the status of your clustered GroupWise system from your Web browser. See Chapter 35, “Implementing GroupWise Monitor in a Heartbeat Cluster,” on page 361.
- ◆ Backing up your clustered GroupWise system. See Chapter 36, “Backing Up a GroupWise System in a Heartbeat Cluster,” on page 363.

Implementing the Internet Agent in a Heartbeat Cluster

33

You should already have set up at least a basic GroupWise® system in your Heartbeat cluster, as described in [Chapter 30, “Planning GroupWise in a Heartbeat Cluster,” on page 285](#), [Chapter 31, “Setting Up a Heartbeat Cluster,” on page 297](#), and [Chapter 32, “Setting Up a Domain and a Post Office in a Heartbeat Cluster,” on page 307](#). As part of this process, you filled out the [Heartbeat Clustering Worksheet](#). If you do not have access to the filled-out worksheet, print the worksheet now and fill out the Heartbeat cluster information as it currently exists on your system. You need this information as you implement the Internet Agent in a Heartbeat cluster.

- ◆ [Section 33.1, “Planning the Internet Agent in a Heartbeat Cluster,” on page 325](#)
- ◆ [Section 33.2, “Setting Up the Internet Agent in a Heartbeat Cluster,” on page 329](#)
- ◆ [Section 33.3, “Testing the Internet Agent in a Heartbeat Cluster,” on page 338](#)
- ◆ [Section 33.4, “Managing the Internet Agent in a Heartbeat Cluster,” on page 339](#)
- ◆ [Section 33.5, “Internet Agent Clustering Worksheet,” on page 340](#)

33.1 Planning the Internet Agent in a Heartbeat Cluster

A major system configuration difference between the Internet Agent in a clustering environment and the Internet Agent in a regular environment is that you need to create a separate domain to house each GroupWise gateway, including the Internet Agent.

The [“Internet Agent Clustering Worksheet” on page 184](#) lists the information you need to set up the Internet Agent in a Heartbeat cluster. You should print the worksheet and fill it out as you complete the tasks listed below:

- ◆ [Section 33.1.1, “Planning a Cluster Resource Group for the Internet Agent,” on page 326](#)
- ◆ [Section 33.1.2, “Planning a Domain for the Internet Agent,” on page 326](#)
- ◆ [Section 33.1.3, “Recording the Internet Agent Secondary IP Address,” on page 327](#)
- ◆ [Section 33.1.4, “Determining Appropriate Heartbeat Constraints for the Internet Agent,” on page 327](#)
- ◆ [Section 33.1.5, “Preparing DNS for the Clustered Internet Agent,” on page 327](#)
- ◆ [Section 33.1.6, “Preparing Your Firewall for the Clustered Internet Agent,” on page 327](#)
- ◆ [Section 33.1.7, “Planning the MTA Installation,” on page 328](#)
- ◆ [Section 33.1.8, “Planning the Internet Agent Installation,” on page 328](#)

33.1.1 Planning a Cluster Resource Group for the Internet Agent

As with the MTA and the POA, the Internet Agent needs cluster nodes, native Heartbeat resources, GroupWise resources, and a cluster resource group in order to function in your Heartbeat cluster. For a review of these cluster components, see [Section 13.1, “Installing Novell Cluster Services on Linux,”](#) on page 134.

INTERNET AGENT CLUSTERING WORKSHEET

Under **Item 1: Internet Agent Nodes**, list the nodes that you want to use for the Internet Agent. Identify each node by its DNS hostname and the IP address that it is physically configured with.

Under **Item 2: Internet Agent Cluster Resource Group**, specify the name of the cluster resource group to hold the native Heartbeat and GroupWise resources associated with the Internet Agent.

Under **Item 3: Native Heartbeat Resources**, list the resource names and parameters for the EVMS shared storage partition, the file system, and the secondary IP address that you want to use for the Internet Agent and its domain.

Under **Item 4: GroupWise Resources**, list the resource names and parameters for the Internet Agent and the MTA.

33.1.2 Planning a Domain for the Internet Agent

The considerations involved in planning a domain for the Internet Agent are much the same as planning any other domain. In preparation, review “[Planning a New Domain](#)”, then print and fill out the “[Domain Worksheet](#)” in “[Domains](#)” in the *GroupWise 7 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 the shared storage partition, not on the node where you run the GroupWise Installation program.
- ◆ 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 6: Internet Agent Domain Name**, transfer the domain name and domain directory to the Internet Agent Clustering Worksheet. Also specify the mount point directory where each node will mount the Internet Agent domain.

IMPORTANT: Do not create the new domain until you are instructed to do so in [Section 15.2.1, “Creating a Domain for the Internet Agent,”](#) on page 171.

33.1.3 Recording the Internet Agent Secondary IP Address

As with the MTA and the POA, the Internet Agent needs a secondary IP address that remains the same no matter which node in the cluster it is running on. You can place the Internet Agent and its domain in a cluster resource group where a domain or post office already reside, which means that the Internet Agent 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 Internet Agent and its domain in its own cluster resource group, which means that it has its own secondary IP address and fails over independently.

INTERNET AGENT CLUSTERING WORKSHEET

Under **Item 7: MTA Network Information**, select an IP address resource from the Heartbeat Clustering Worksheet (**item 5**) and record the secondary IP address associated with the resource on the Internet Agent Internet Agent Clustering Worksheet.

Under **Item 8: Internet Agent Network Information**, copy the same secondary IP address.

33.1.4 Determining Appropriate Heartbeat Constraints for the Internet Agent

By default, a Heartbeat cluster is configured to have all nodes in the cluster available for failover. Only one node at a time can have a particular cluster resource group mounted and active. If a cluster resource group's initial node fails, the resource group fails over to another node in the cluster.

You should customize the Heartbeat constraints for each GroupWise cluster resource group based on the fan-out-failover principle. As with the MTA and the POA, you need to select the nodes in the cluster where the Internet Agent cluster resource group can fail over. You must install the Internet Agent software on all of the nodes where you want the Internet Agent to be able to fail over. For a review of the fan-out-failover principle, see **Section 13.6.2, "Determining Appropriate Failover Lists for the Agents,"** on page 138, which describes the issues in the context of planning MTA and POA installations.

INTERNET AGENT CLUSTERING WORKSHEET

Under **Item 3: Internet Agent Cluster Resource Group Constraints**, list the nodes where you want the Internet Agent to be able to fail over.

33.1.5 Preparing DNS for the Clustered Internet Agent

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

33.1.6 Preparing Your Firewall for the Clustered Internet Agent

The Internet Agent receives incoming messages on its secondary IP address, not the physical IP address of the node where it is running. Your firewall configuration must be modified to allow inbound TCP/IP traffic from the Internet to the Internet Agent secondary IP address on the following standard ports:

Table 33-1 Standard Ports

Protocol	Standard Port
IMAP4	143
LDAP	389
POP3	110
SMTP	25

By default, the Internet Agent sends outgoing messages on the *primary 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 where the Internet Agent could fail over.

If the Internet Agent has a large number of nodes available, you could configure the Internet Agent 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 Internet Agent was running. This reduces the amount of modification to your firewall required to set up the Internet Agent. However, if the relay host goes down, outgoing messages are delayed.

As another alternative, you can configure the Internet Agent to use its secondary IP address for sending as well as receiving messages. Setup instructions for this configuration are provided in [“Forcing Use of the Internet Agent Secondary IP Address” on page 181](#), which you can complete after installing the Internet Agent software.

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

33.1.7 Planning the MTA Installation

Follow the instructions in [Section 13.6.4, “Planning the Linux Agent Installation,” on page 139](#), then return to this point. After you follow the instructions, you will have a filled-out [System 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 15.2, “Setting Up the Internet Agent in a Linux Cluster,” on page 171](#).

33.1.8 Planning the Internet Agent Installation

Aside from the cluster-specific issues discussed in the preceding sections, the considerations involved in planning to install the Internet Agent are the same in a clustering environment as for any other environment. Review the installation instructions in [“Linux: Installing the Internet Agent Software”](#) in [“Installing the GroupWise Internet Agent”](#) in the *GroupWise 7 Installation Guide*. Use the [“GroupWise Internet Agent Installation Worksheet”](#) to record the planning information you will need as you install the Internet Agent in your cluster.

IMPORTANT: Do not install the Internet Agent software until you are instructed to do so in [Setting Up the Internet Agent in a Linux Cluster](#).

33.2 Setting Up the Internet Agent in a Heartbeat Cluster

You should already have reviewed [Section 15.1, “Planning the Internet Agent in a Linux Cluster,” on page 167](#) and filled out [Section 15.5, “Internet Agent Clustering Worksheet,” on page 184](#). You are now ready to complete the following tasks to set up the Internet Agent in a clustering environment:

- ♦ [Section 33.2.1, “Setting Up Native Heartbeat Cluster Resources for the Internet Agent,” on page 329](#)
- ♦ [Section 33.2.2, “Creating a Domain for the Internet Agent,” on page 329](#)
- ♦ [Section 33.2.3, “Installing the MTA for the Internet Agent Domain,” on page 330](#)
- ♦ [Section 33.2.4, “Installing and Configuring the Internet Agent in a Cluster,” on page 330](#)

33.2.1 Setting Up Native Heartbeat Cluster Resources for the Internet Agent

- 1 Start the HA Management Client, as described in [Section 31.1, “Starting the HA Management Client,” on page 297](#).
- 2 Create a cluster resource group for the Internet agent and its domain ([Internet Agent Clustering Worksheet item 2](#)), as described in [Section 31.2, “Creating a Cluster Resource Group,” on page 298](#).
- 3 Create the EVMS container resource, the file system resource, and the IP address resource for the cluster resource group ([Internet Agent Clustering Worksheet item 3](#)), as described in [Section 31.3, “Creating Native Heartbeat Resources,” on page 299](#).

As opposed to native Heartbeat resources, GroupWise resources for the Internet Agent and its MTA must be created after you install the Internet Agent software, so that the `GroupWise` script is available to enhance the capabilities of the HA Management Client. By creating native resources before you install the Internet Agent software, you can use them to conveniently fail over the nodes as you install the Internet Agent software on multiple nodes.

- 4 Set up constraints on the nodes where the Internet Agent and its domain can fail over ([Internet Agent Clustering Worksheet item 5](#)), as described in [Section 31.4, “Creating Heartbeat Constraints,” on page 302](#)
- 5 Continue with [Creating a Domain for the Internet Agent](#).

33.2.2 Creating a Domain for the Internet Agent

The Internet Agent domain will be a secondary domain. To create it, follow the instructions in [Section 14.2, “Creating a New Secondary Domain in a Linux Cluster,” on page 144](#), taking your information from the Internet Agent Clustering Worksheet rather than the GroupWise Clustering Worksheet, then return to this point.

Do not create any post offices in the Internet Agent domain.

After you have created the domain, continue with [Installing the MTA for the Internet Agent Domain](#).

33.2.3 Installing the MTA for the Internet Agent Domain

The MTA for the Internet Agent domain can be installed just like any other MTA in your clustered GroupWise system. Follow the instructions in [Section 14.4, “Installing and Configuring the MTA and the POA in a Cluster,”](#) on page 146, then return to this point.

After you have installed the MTA, continue with [Installing and Configuring the Internet Agent in a Cluster](#).

33.2.4 Installing and Configuring the Internet Agent in a Cluster

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

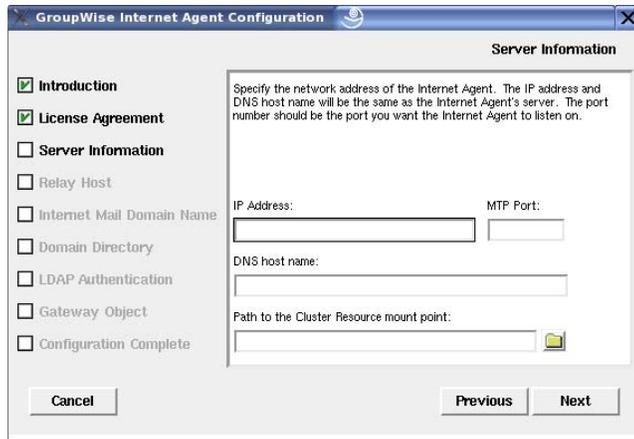
- ♦ [“Running the Internet Agent Installation Program on the Initial Node”](#) on page 330
- ♦ [“Running the Internet Agent Installation Program on Subsequent Nodes”](#) on page 332
- ♦ [“Creating and Configuring GroupWise Cluster Resources for the Internet Agent”](#) on page 333
- ♦ [“Testing Your Internet Agent Installation on Each Node”](#) on page 335
- ♦ [“Configuring the Clustered Internet Agent for SSL”](#) on page 336
- ♦ [“Enabling Internet Addressing for Your Clustered GroupWise System”](#) on page 337
- ♦ [“Verifying Internet Agent Object Properties”](#) on page 337

Running the Internet Agent Installation Program on the Initial Node

- 1 Make sure that the Internet Agent software is available in the software distribution directory you created in [Step 6](#) in [Section 14.1, “Setting Up a New GroupWise System in a Linux Cluster,”](#) on page 143.
- 2 Start the resources in the Internet Agent cluster resource group ([Internet Agent Clustering Worksheet item 2](#)) in order to mount the domain directory where the Internet Agent message queues will be located.
- 3 From the software distribution directory, start the GroupWise Installation program and select *Configure GroupWise for Clustering*.



- 4 Install the Internet Agent software, following the steps provided in [“Linux: Installing the Internet Agent Software”](#) in [“Installing the GroupWise Internet Agent”](#) in the *GroupWise 7 Installation Guide*.
- 5 Configure the Internet Agent according to the [“GroupWise Internet Agent Installation Worksheet”](#) that you filled out in [Section 15.1, “Planning the Internet Agent in a Linux Cluster,”](#) on page 167, paying special attention to the cluster resource information on the Server Information page.



The IP address is the secondary IP address associated with the Internet Agent cluster resource group ([Internet Agent Clustering Worksheet item 8](#)). The path to the cluster resource mount point is the mount point for the Internet Agent domain ([Internet Agent Clustering Worksheet item 6](#)).

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

- ◆ The Internet Agent startup file (`gwia.cfg`) is created in `/mount_point/groupwise/agents/share` on the shared storage partition so that the Internet Agent uses the same startup file regardless of which node it runs on. The `--home` switch includes the mount point and the path to the domain database so that the startup file is valid when mounted on each node.
- ◆ The `--cluster` switch is added to the Internet Agent startup file to inform the Internet Agent that it is running in a cluster.
- ◆ The `--ip` startup switch is set to the secondary IP address of the cluster resource group where the domain is located. This ensures that the Internet Agent runs with the same IP address regardless of which node it runs on.
- ◆ The `--log` startup switch is set to a location on the shared storage partition (`mount_point/groupwise/agents/log`) so that Internet Agent logging information is written to the same log file regardless of which node the Internet Agent runs on.
- ◆ If this is the first GroupWise agent installed on this node, the GroupWise High Availability service (GWHA) 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 node, the `gwha.conf` file is updated to include the Internet Agent.
- ◆ If it is not already present, the OCF Resource Agent script (`GroupWise`) is installed in the `/usr/lib/ocf/resource.d/novell` directory on the node for use by the HA Management Client.
- ◆ 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 configure the Internet Agent on the initial node is available when you install the Internet Agent on subsequent nodes.
- ◆ The Internet Agent is not configured to start automatically on system startup. In a cluster, you do not want the Internet Agent to start automatically whenever a node restarts.

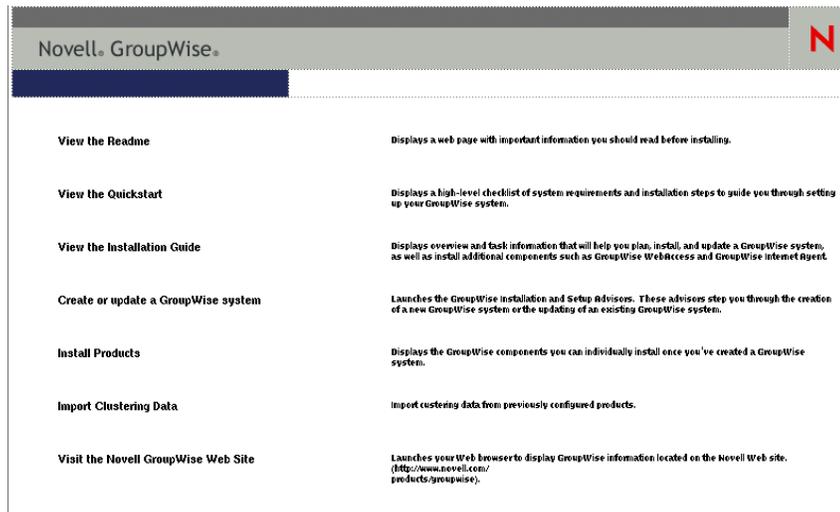
- 6 For optimum security, configure the Internet Agent to run as a non-root user, as described in “Running the Linux GroupWise Agents as a Non-root User” in “Installing GroupWise Agents” in the *GroupWise 7 Installation Guide*.
- 7 Continue with [Running the Internet Agent Installation Program on Subsequent Nodes](#).

Running the Internet Agent Installation Program on Subsequent Nodes

- 1 In the **HA Management Client**, select a different node where you want to run the Internet Agent (**Internet Agent Clustering Worksheet item 3**), then click *Make the Node Active*.
- 2 From the software distribution directory you created in **Step 6 in Section 14.1, “Setting Up a New GroupWise System in a Linux Cluster,” on page 143**, 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 Internet Agent software on the node as usual, but do not use the *Configure* option. You will use the *Import Clustering Data* option instead of 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 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 node. It is configured specifically for the agents you just selected.
- ♦ The GroupWise High Availability service (GWHA) is automatically configured on the current 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.
- ♦ If it is not already present, the OCF Resource Agent script (`GroupWise`) is installed in the `/usr/lib/ocf/resource.d/novell` directory on the node for use by the HA Management Client.

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

- 6 For optimum security, configure the Internet Agent to run as a non-`root` user, as described in “[Running the Linux GroupWise Agents as a Non-root User](#)” in “[Installing GroupWise Agents](#)” in the *GroupWise 7 Installation Guide*.
- 7 Repeat [Step 1](#) through [Step 6](#) for each node where you want the Internet Agent to be able to fail over ([Internet Agent Clustering Worksheet item 3](#)).
- 8 Continue with [Creating and Configuring GroupWise Cluster Resources for the Internet Agent](#).

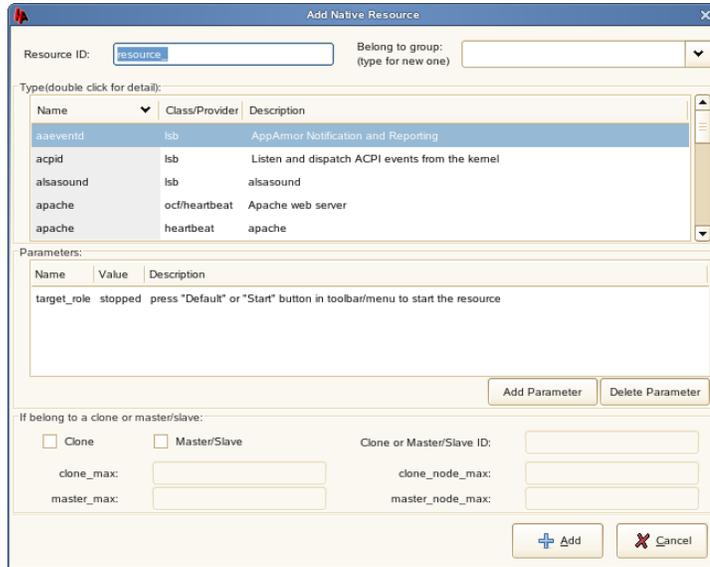
Creating and Configuring GroupWise Cluster Resources for the Internet Agent

Now that you have installed the Internet Agent software, you can create and configure GroupWise resources in the cluster resource group you created in [Section 33.2.1, “Setting Up Native Heartbeat Cluster Resources for the Internet Agent,”](#) on page 329:

- ♦ “[Creating GroupWise Resources for the Internet Agent](#)” on page 334
- ♦ “[Enabling Automatic Internet Agent Restart](#)” on page 334
- ♦ “[Preventing Internet Agent Failback](#)” on page 335

Creating GroupWise Resources for the Internet Agent

- 1 In the **HA Management Client**, select a cluster resource group, then click *Add New Item*.
- 2 In the *Item Type* field, select *native*, then click *OK*.



- 3 In the *Resource ID* field, specify a unique name for the GroupWise resource (**Internet Agent Clustering Worksheet item 4**).
- 4 In the *Belong to Group* field, select the cluster resource group where you want to create the GroupWise resource for the Internet Agent
- 5 In the *Type* list, select *GroupWise (ocf/novell)*.
The OCF Resource Agent script (*GroupWise*) added this item to the list.
- 6 In the *Parameters* list, set the value of the *object_name* parameter to the identifier for the Internet Agent (**Internet Agent Clustering Worksheet item 4**).
For the Internet Agent, the identifier is the domain name, followed by the eDirectory™ object name of the Internet Agent (*Domain.GWIA*).
- 7 Click *Add* to create the GroupWise resource.
- 8 Create another GroupWise resource in the cluster resource group for the Internet Agent's MTA.
For a domain, the identifier is the domain name.
- 9 Continue with **Enabling Automatic Internet Agent Restart**.

Enabling Automatic Internet Agent Restart

If you want Heartbeat to automatically restart an Internet Agent that has stopped:

- 1 In the **HA Management Client**, select the GroupWise resource for the Internet Agent, then click *Operations*.



- 2 Click *Add Operation*, then select *monitor*.
- 3 Select the time interval for how often you want the Internet Agent to be polled to see if it is still running.

The GroupWise script causes the Internet Agent to restart, using the GroupWise High Availability service. (GWHA)

- 4 Click *OK* to save the monitoring setting.
- 5 Repeat the process for the GroupWise resource for the MTA in the Internet Agent's cluster resource group.

Because Heartbeat can perform such agent monitoring, you do not need to set up GroupWise Monitor to perform this task, as described in “[Configuring the Monitor Agent to Communicate with the High Availability Service](#)” in “[Installing GroupWise Monitor](#)” in the *GroupWise 7 Installation Guide*. Heartbeat and the GroupWise High Availability Service can work together to keep the Internet Agent running on Linux.

- 6 Continue with [Preventing Internet Agent Failback](#).

Preventing Internet Agent Failback

If you do not want the Internet Agent cluster resource group to automatically fail back when the initial node is on-line again:

- 1 In the [HA Management Client](#), select the GroupWise cluster resource group for the Internet Agent, click *Parameters*, then click *Add Parameter*.



- 2 In the *Name* field, specify `resource_stickiness`.
- 3 In the *Value* field, specify `INFINITY`.
- 4 Click *Apply* to save the setting.
- 5 Continue with [Testing Your Internet Agent Installation on Each Node](#).

Testing Your Internet Agent Installation on Each Node

- 1 Test the Internet Agent by starting it with a user interface on the currently active node, as described in “[Linux: Starting the Internet Agent](#)” in “[Installing the GroupWise Internet Agent](#)” in the *GroupWise 7 Installation Guide*.

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

- 2 Stop the Internet Agent by clicking *File > Exit*.

- 3 After you can see that the Internet Agent stopped successfully, test it by starting it as a daemon, as described in “Starting the Linux GroupWise Agents as Daemons” in “Installing GroupWise Agents” in the *GroupWise 7 Installation Guide*.

```
/etc/inet.d/grpwise start domain.gwia  
/etc/inet.d/grpwise status domain.gwia
```

- 4 Stop the Internet Agent.

```
/etc/inet.d/grpwise stop domain.gwia  
/etc/inet.d/grpwise status domain.gwia
```

- 5 Also test the MTA, as described in Section , “Testing Your Agent Installation on Each Node,” on page 150.
- 6 Repeat Step 1 through Step 4 on each node where you have installed the Internet Agent.
- 7 Make sure you have completed all the tasks described in “Installing the GroupWise Internet Agent” in the *GroupWise 7 Installation Guide*.

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

- 8 After you have tested the Internet Agent on all nodes manually, start and stop the Internet Agent resource in the HA Management Client to make sure that it is properly configured.

If you receive the error `Unmanaged on node`, then you have failed to test the Internet Agent on one or more nodes. On each node, view the `gwha.conf` file in the `/etc/opt/novell/groupwise` directory. Each Internet Agent section should be preceded by the name of the domain followed by the eDirectory object name of the Internet Agent (GWIA by default). If any section is preceded by a startup file name (for example, `@gwia.conf`), then that Internet Agent has not yet been initialized by being run manually. To resolve the problem, you can run the Internet Agent manually on that node, or you can replace the `@gwia.cnf` with the appropriate domain and Internet Agent object information, being sure to follow the same format as the other Internet Agent section headings in the `gwha.conf` file.

- 9 Continue with [Configuring the Clustered Internet Agent for SSL](#).

Configuring the Clustered Internet Agent for SSL

If you plan to enable SSL, as described in “Securing Internet Agent Connections with SSL” in “Internet Agent” in the *GroupWise 7 Administration Guide*, you must make the SSL certificate file and key file available to the Internet Agent in the cluster. The default locations for the SSL certificate file and key file are on the nodes along with the Internet Agent software, rather than being located with the domain on the shared storage partition. To avoid having multiple copies of these files in multiple locations, you should set the locations in ConsoleOne®.

- 1 On the shared storage partition where the Internet Agent domain is located, create the directory where you want to store the certificate file 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 7 Administration Guide*.
- 3 In ConsoleOne, browse to and select the Domain object.
- 4 Right-click the Internet Agent 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 [Enabling Internet Addressing for Your Clustered GroupWise System](#).

Enabling Internet Addressing for Your Clustered GroupWise System

Setting up Internet addressing for a clustered Internet Agent is no different from setting it up for an Internet Agent in any other environment. Follow the instructions in “[Configuring Internet Addressing](#)” in “[Internet Agent](#)” in the *GroupWise 7 Administration Guide*, then return to this point.

Verifying Internet Agent Object Properties

During installation of the Internet Agent, the Internet Agent 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 Internet Agent.

- ♦ “[Accessing Internet Agent Object Properties](#)” on page 337
- ♦ “[Verifying the Reference to the Internet Agent Secondary IP Address](#)” on page 337
- ♦ “[Verifying the Reference to the Mount Point Directory](#)” on page 337
- ♦ “[Verifying Post Office Links](#)” on page 338
- ♦ “[Forcing Use of the Internet Agent Secondary IP Address](#)” on page 338

Accessing Internet Agent Object Properties

- 1 In ConsoleOne, browse to and select the Internet Agent domain in order to display its contents.
- 2 Right-click the Internet Agent object, then click *Properties*.
- 3 Continue with [Verifying Internet Agent Object Properties](#).

Verifying the Reference to the Internet Agent Secondary IP Address

In the Internet Agent object property page:

- 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 display the hostname that corresponds to the secondary IP address of the Internet Agent cluster resource group. For more information, see [Section 15.1.5, “Preparing DNS for the Clustered Internet Agent,”](#) on page 169.
- 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 Internet Agent object property page:

- 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 Internet Agent object property page:

- 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 Internet Agent.
- 3 Verify that the *Links* column displays the secondary IP addresses of the cluster resource groups that the post offices belong to, not the IP addresses of any nodes in the cluster.
- 4 Make changes if necessary.
- 5 Continue with **Forcing Use of the Internet Agent Secondary IP Address**.

Forcing Use of the Internet Agent Secondary IP Address

If you want the Internet Agent to send outgoing messages on its secondary IP address, rather than using the default of its primary IP address:

- 1 In ConsoleOne, click *GroupWise > Network Address*.
- 2 In the *TCP/IP Address* field, provide the secondary IP address (**Internet Agent Clustering Worksheet item 6**) for the Internet Agent to use for sending outgoing messages.
- 3 Select *Bind Exclusively to TCP/IP Address*.
- 4 Click *OK*.
- 5 Continue with **Testing the Internet Agent in a Linux Cluster**.

33.3 Testing the Internet Agent in a Heartbeat Cluster

After you have installed and configured the Internet Agent in your Heartbeat cluster, you can use the following methods to test whether the Internet Agent fails over as expected:

- ♦ In the **HA Management Client**, select a node where the Internet Agent is running, then click *Make the Node Standby*. The Internet Agent should fail over to another node and start running again. This can also be done from the command line using the `crm_standby` utility (http://www.linux-ha.org/v2/AdminTools/crm_standby).
- ♦ Using the `crm_resource` utility (http://www.linux-ha.org/v2/AdminTools/crm_resource), manually migrate a node where the Internet Agent is running.

If you receive the error `Unmanaged on node`, see **Step 8** in “**Testing Your Internet Agent Installation on Each Node**” on page 335 for troubleshooting assistance.

If you have set up your Heartbeat cluster in a test situation first, you can also use the following methods for testing Internet Agent failover:

- ♦ Power off the node where the Internet Agent is running. The Internet Agent should fail over to another node and start running again.
- ♦ Simulate the Internet Agent stopping unexpectedly (`kill -9 agent_PID`). If you have configured the Internet Agent for automatic restart, as described in “**Enabling Automatic GroupWise Agent Restart**” on page 316, the agent should start running again on the same node.

WARNING: These last two tests should not be performed on a production system because they can result in damage to GroupWise databases.

33.4 Managing the Internet Agent in a Heartbeat Cluster

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

- ♦ [Section 33.4.1, “Updating GroupWise Objects with Cluster-Specific Descriptions,” on page 339](#)
- ♦ [Section 33.4.2, “Knowing What to Expect in an Internet Agent Failover Situation,” on page 340](#)

33.4.1 Updating GroupWise Objects with Cluster-Specific Descriptions

After installing the Internet Agent 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 Internet Agent Domain and Its MTA” on page 339](#)
- ♦ [“Recording Cluster-Specific Information about the Internet Agent” on page 339](#)

Recording Cluster-Specific Information about the Internet Agent Domain and Its MTA

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

- 1 In ConsoleOne, browse to and right-click the Domain object, then click *Properties*.
- 2 In the *Description* field of the Internet Agent domain Identification page, provide a cluster-specific description of the Internet Agent domain, including the cluster resource group and shared storage partition where the domain is located.
- 3 Click *OK* to save the Internet Agent domain description.
- 4 Select the Internet Agent 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 associated with the Internet Agent domain’s cluster resource group.
- 7 Click *OK* to save the MTA description.
- 8 Continue with [Recording Cluster-Specific Information about the Internet Agent](#).

Recording Cluster-Specific Information about the Internet Agent

With the contents of the Internet Agent domain still displayed:

- 1 Right-click the Internet Agent object, then click *Properties*.

- 2 Click *GroupWise > Identification*.
- 3 In the *Description* field, record the secondary IP address associated with the Internet Agent domain's cluster resource group.
- 4 Click *OK* to save the Internet Agent information.
- 5 Continue with [Knowing What to Expect in an Internet Agent Failover Situation](#).

33.4.2 Knowing What to Expect in an Internet Agent Failover Situation

The failover behavior of the MTA for the Internet Agent domain is the same as for an MTA in a regular domain. See [Section 14.6.2, "Knowing What to Expect in MTA and POA Failover Situations,"](#) on page 163.

Failover of the Internet Agent 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 Internet Agent restarts quickly in its failover location so users can reconnect quickly.

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

33.5 Internet Agent Clustering Worksheet

Item	Explanation
1) Internet Agent Nodes:	<p>List the nodes where you want the Internet Agent to be able to fail over. Identify each node by its DNS hostname and the IP address that it is physically configured with.</p> <p>For more information, see Section 33.1.1, "Planning a Cluster Resource Group for the Internet Agent," on page 326.</p>
2) Internet Agent Cluster Resource Group:	<p>Specify the name of the cluster resource group to hold the resources required for the Internet Agent.</p> <p>For more information, see Section 15.1.1, "Planning a Domain for the Internet Agent," on page 168.</p>
3) Native Heartbeat Resources:	<p>List names for the native Heartbeat resources and the values for their parameters to be used in the Internet Agent cluster resource group.</p> <ul style="list-style-type: none"> ◆ EVMS container ◆ File system ◆ Secondary IP address <p>For more information, see Section 33.1.1, "Planning a Cluster Resource Group for the Internet Agent," on page 326.</p>
4) GroupWise Resources:	<p>List the names for the GroupWise resources and the values for their parameters to be used in the Internet Agent cluster resource group.</p> <ul style="list-style-type: none"> ◆ MTA ◆ Internet Agent <p>For more information, see Section 33.1.1, "Planning a Cluster Resource Group for the Internet Agent," on page 326</p>

Item	Explanation
5) Internet Agent Constraints: <ul style="list-style-type: none"> ◆ Place ◆ Order ◆ Colocation 	Specify the Heartbeat constraints that you want to put on the Internet Agent cluster resource group to control what nodes it can run on. For more information, see Section 15.1.3, “Determining an Appropriate Failover List for the Internet Agent,” on page 169
6) Internet Agent Domain Name: Domain Directory: Mount Point:	Specify a unique name for the Internet Agent domain. Specify the directory on the shared storage partition where you want to create the new domain. Specify the mount point on each node where you plan to mount the domain. For more information, see Section 15.1.1, “Planning a Domain for the Internet Agent,” on page 168 .
7) MTA Network Information: <ul style="list-style-type: none"> ◆ Secondary IP address ◆ Message transfer port ◆ HTTP port 	Record the MTA network address information that you will need as you install the MTA. For more information, see Section 15.1.7, “Planning the MTA Installation,” on page 170 .
8) Internet Agent Network Information: <ul style="list-style-type: none"> ◆ Secondary IP address ◆ HTTP port 	Record the Internet Agent network address information that you will need to install the Internet Agent. For more information, see Section 15.1.8, “Planning the Internet Agent Installation,” on page 170 .

Implementing WebAccess in a Heartbeat Cluster

34

You should already have set up at least a basic GroupWise® system, as described in [Chapter 30](#), “[Planning GroupWise in a Heartbeat Cluster](#),” on page 285, [Chapter 31](#), “[Setting Up a Heartbeat Cluster](#),” on page 297, and [Chapter 32](#), “[Setting Up a Domain and a Post Office in a Heartbeat Cluster](#),” on page 307. As part of this process, you filled out the “[Heartbeat Clustering Worksheet](#)” on page 293. 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 WebAccess Agent in a cluster.

- ♦ [Section 34.1](#), “[Understanding the WebAccess Components](#),” on page 343
- ♦ [Section 34.2](#), “[Planning the WebAccess Agent in a Heartbeat Cluster](#),” on page 343
- ♦ [Section 34.3](#), “[Setting Up the WebAccess Agent in a Heartbeat Cluster](#),” on page 346
- ♦ [Section 34.4](#), “[Testing the WebAccess Agent in a Heartbeat Cluster](#),” on page 356
- ♦ [Section 34.5](#), “[Managing the WebAccess Agent in a Heartbeat Cluster](#),” on page 356
- ♦ [Section 34.6](#), “[WebAccess Agent Clustering Worksheet](#),” on page 358

34.1 Understanding the WebAccess Components

If you are not familiar with GroupWise WebAccess, review “[GroupWise WebAccess Overview](#)” in “[Installing GroupWise WebAccess](#)” in the *GroupWise 7 Installation Guide*.

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

- ♦ WebAccess Agent ([gwinter](#)) that will be associated with a domain in your GroupWise system
- ♦ WebAccess Application (a [Java servlet](#)) that will be added to your Web server (Apache). The WebAccess Application is currently not supported on a Heartbeat cluster.

You install the WebAccess Agent on each node in the cluster. You install the WebAccess Application to your Web server, which must not be clustered. This means that the WebAccess client login to the Web server at http://secondary_IP_address/gw/webacc is available only when the Web server is running.

34.2 Planning the WebAccess Agent in a Heartbeat Cluster

In a cluster, you need to create a separate domain to house each GroupWise gateway, including the WebAccess Agent.

The “[WebAccess Agent Clustering Worksheet](#)” on page 204 lists the information you need to set up the WebAccess Agent in a clustering environment. You should print the worksheet and fill it out as you complete the tasks listed below:

- ◆ [Section 34.2.1, “Planning a Cluster Resource Group for the WebAccess Agent,”](#) on page 344
- ◆ [Section 34.2.2, “Planning a Domain for the WebAccess Agent,”](#) on page 344
- ◆ [Section 34.2.3, “Recording the WebAccess Secondary IP Address,”](#) on page 345
- ◆ [Section 34.2.4, “Determining Appropriate Heartbeat Constraints for the WebAccess Agent,”](#) on page 345
- ◆ [Section 34.2.5, “Planning the MTA Installation,”](#) on page 346
- ◆ [Section 34.2.6, “Planning the WebAccess Agent Installation,”](#) on page 346
- ◆ [Section 34.2.7, “Planning the WebAccess Application Installation,”](#) on page 346

34.2.1 Planning a Cluster Resource Group for the WebAccess Agent

As with the MTA and the POA, the WebAccess Agent needs cluster nodes, native Heartbeat resources, GroupWise resources, and a cluster resource group in order to function in your Heartbeat cluster. For a review of these cluster components, see [Section 13.1, “Installing Novell Cluster Services on Linux,”](#) on page 134.

WEBACCESS AGENT CLUSTERING WORKSHEET

Under [Item 1: WebAccess Agent Nodes](#), list the nodes that you want to use for the WebAccess Agent. Identify each node by its DNS hostname and the IP address that it is physically configured with.

Under [Item 2: WebAccess Agent Cluster Resource Group](#), specify the name of the cluster resource group to hold the native Heartbeat and GroupWise resources associated with the WebAccess Agent.

Under [Item 3: Native Heartbeat Resources](#), list the native resource names and parameter values for the EVMS shared storage partition, the file system, and the secondary IP address that you want to use for the WebAccess Agent and its domain.

Under [Item 4: GroupWise Resources](#), list the resource names and parameter values for the WebAccess Agent and the MTA.

34.2.2 Planning a Domain for the WebAccess Agent

The considerations involved in planning a domain for the WebAccess Agent are much the same as planning any other domain. In preparation, review “[Planning a New Domain](#)”, then print and fill out the “[Domain Worksheet](#)” in “[Domains](#)” in the *GroupWise 7 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 shared storage partition, not on the node where you run the GroupWise Installation program.
- ◆ 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 WebAccess Agent Clustering Worksheet.

WEBACCESS AGENT CLUSTERING WORKSHEET

Under **Item 6: WebAccess Agent Domain Name**, transfer the domain name and domain directory to the WebAccess Agent Clustering Worksheet.

IMPORTANT: Do not create the new domain until you are instructed to do so in **Section 16.3.1, “Creating a Domain for the WebAccess Agent,”** on page 191.

34.2.3 Recording the WebAccess Secondary IP Address

As with the MTA and the POA, the WebAccess Agent needs a secondary IP address that remains the same no matter which node in the cluster it runs on. You can place the WebAccess Agent and its domain in a cluster resource group where a domain or post office already reside, which means that the WebAccess Agent 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 WebAccess Agent and its domain in its own cluster resource group, which means that it has its own secondary IP address and fails over independently.

WEBACCESS AGENT CLUSTERING WORKSHEET

Under **Item 7: MTA Network Information**, select an IP address resource from the Heartbeat Clustering Worksheet (**item 5**) and record the secondary IP address associated with the resource on the WebAccess Agent Clustering Worksheet.

Under **Item 8: WebAccess Agent Network Information**, copy the same secondary IP address.

34.2.4 Determining Appropriate Heartbeat Constraints for the WebAccess Agent

By default, a Heartbeat cluster is configured to have all nodes in the cluster available for failover. Only one node at a time can have a particular cluster resource group mounted and active. If a cluster resource group’s initial node fails, the resource group fails over to another node in the cluster.

You should customize the Heartbeat constraints for each GroupWise cluster resource group based on the fan-out-failover principle. As with the MTA and the POA, you need to decide which nodes in the cluster are appropriate locations where the WebAccess Agent cluster resource group can fail over. You must install the WebAccess Agent software on all of the nodes where you want the WebAccess Agent to be able to fail over. For a review of the fan-out-failover principle, see **Section 13.6.2, “Determining Appropriate Failover Lists for the Agents,”** on page 138, which describes the issues in the context of planning MTA and POA installations.

WEBACCESS AGENT CLUSTERING WORKSHEET

Under **Item 5: WebAccess Agent Resource Group Constraints**, list the nodes where you want the WebAccess Agent to fail over.

34.2.5 Planning the MTA Installation

Follow the instructions in [Section 13.6.4, “Planning the Linux Agent Installation,”](#) on page 139, then return to this point. After you follow the instructions, you have a filled-out [System 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 16.3, “Setting Up the WebAccess Agent in a Linux Cluster,”](#) on page 190.

34.2.6 Planning the WebAccess Agent Installation

Aside from the cluster-specific issues discussed in the preceding sections, the considerations involved in planning to install the WebAccess Agent are the same in a clustering environment as for any other environment. Review the installation instructions in “[Installing the Linux WebAccess Agent](#)” in “[Installing GroupWise WebAccess](#)” in the *GroupWise 7 Installation Guide*. Use the “[GroupWise WebAccess Installation Worksheet](#)” to record the planning information you will need as you install the WebAccess Agent in your cluster.

IMPORTANT: Do not install the WebAccess Agent software until you are instructed to do so in [Section 16.3, “Setting Up the WebAccess Agent in a Linux Cluster,”](#) on page 190.

34.2.7 Planning the WebAccess Application Installation

The WebAccess Application must be installed on a non-clustered Web server. For WebAccess Application planning information, see “[Determining the WebAccess and WebPublisher Applications’ Configuration](#)” in “[Installing GroupWise WebAccess](#)” in the *GroupWise 7 Installation Guide*.

34.3 Setting Up the WebAccess Agent in a Heartbeat Cluster

You should already have reviewed [Section 16.2, “Planning the WebAccess Agent in a Linux Cluster,”](#) on page 187 and filled out the “[WebAccess Agent Clustering Worksheet](#)” on page 204. You are now ready to complete the following tasks to set up the WebAccess Agent in a clustering environment:

- ◆ [Section 34.3.1, “Setting Up Native Heartbeat Resources for the WebAccess Agent,”](#) on page 347
- ◆ [Section 34.3.2, “Creating a Domain for the WebAccess Agent,”](#) on page 347
- ◆ [Section 34.3.3, “Installing the MTA for the WebAccess Agent Domain,”](#) on page 347
- ◆ [Section 34.3.4, “Installing and Configuring the WebAccess Agent in a Heartbeat Cluster,”](#) on page 348
- ◆ [Section 34.3.5, “Installing and Configuring the WebAccess Application on Your Web Server,”](#) on page 355

34.3.1 Setting Up Native Heartbeat Resources for the WebAccess Agent

- 1 Start the HA Management Client, as described in [Section 31.1, “Starting the HA Management Client,”](#) on page 297.
- 2 Create a cluster resource group for the WebAccess agent and its domain ([WebAccess Agent Clustering Worksheet item 2](#)), as described in [Section 31.2, “Creating a Cluster Resource Group,”](#) on page 298.
- 3 Create the EVMS container resource, the file system resource, and the IP address resource for the cluster resource group ([WebAccess Agent Clustering Worksheet item 3](#)), as described in [Section 31.3, “Creating Native Heartbeat Resources,”](#) on page 299.

As opposed to native Heartbeat resources, GroupWise resources for the WebAccess Agent and its MTA must be created after you install the WebAccess Agent software, so that the GroupWise script is available to enhance the capabilities of the HA Management Client. By creating native resources before you install the WebAccess Agent software, you can use them to conveniently fail over the nodes as you install the WebAccess Agent software on multiple nodes.

- 4 Set up constraints on the nodes where the WebAccess Agent and its domain can fail over ([WebAccess Agent Clustering Worksheet item 5](#)), as described in [Section 31.4, “Creating Heartbeat Constraints,”](#) on page 302
- 5 Refine the configuration of the WebAccess Agent and MTA resources, as described in [“Enabling Automatic GroupWise Agent Restart”](#) on page 316 and [“Preventing GroupWise Agent Failback”](#) on page 317.
- 6 Continue with [Creating a Domain for the WebAccess Agent](#).

34.3.2 Creating a Domain for the WebAccess Agent

The WebAccess Agent domain will be a secondary domain. To create it, follow the instructions in [Section 14.2, “Creating a New Secondary Domain in a Linux Cluster,”](#) on page 144, taking your information from the WebAccess Agent Clustering Worksheet, rather than the GroupWise Clustering Worksheet, then return to this point.

Do not create any post offices in the WebAccess Agent domain.

After you have created the domain, continue with [Installing the MTA for the WebAccess Agent Domain](#).

34.3.3 Installing the MTA for the WebAccess Agent Domain

The MTA for the WebAccess Agent domain can be installed just like any other MTA in your clustered GroupWise system. Follow the instructions in [Section 14.4, “Installing and Configuring the MTA and the POA in a Cluster,”](#) on page 146, then return to this point.

After you have installed the MTA, continue with [Installing and Configuring the WebAccess Agent in a Cluster](#).

34.3.4 Installing and Configuring the WebAccess Agent in a Heartbeat Cluster

After you have created a domain for the WebAccess Agent and installed the MTA for that domain, you are ready to install and configure the WebAccess Agent. The WebAccess Agent must be installed on each node where it is allowed to fail over ([WebAccess Agent Clustering Worksheet item 5](#)).

- ♦ “Running the WebAccess Agent Installation Program on the Initial Node” on page 348
- ♦ “Running the WebAccess Agent Installation Program on Subsequent Nodes” on page 349
- ♦ “Creating and Configuring GroupWise Cluster Resources for the WebAccess Agent” on page 351
- ♦ “Testing Your WebAccess Agent Installation on Each Node” on page 353
- ♦ “Configuring Clustered Logging” on page 354
- ♦ “Configuring the Clustered WebAccess Agent for SSL” on page 354
- ♦ “Verifying WebAccess Agent Object Properties” on page 355

Running the WebAccess Agent Installation Program on the Initial Node

- 1 Make sure that the WebAccess Agent software is available in the software distribution directory you created in [Step 6 in Section 14.1, “Setting Up a New GroupWise System in a Linux Cluster,”](#) on page 143.
- 2 Start the resources in the WebAccess Agent resource group ([WebAccess Agent Clustering Worksheet item 2](#)) in order to mount the domain directory where the WebAccess Agent message queues are located.
- 3 From the software distribution directory, start the GroupWise Installation program and select *Configure GroupWise for Clustering*.



- 4 Install the WebAccess Agent software, following the steps provided in “[Installing the Linux WebAccess Agent](#)” in “[Installing GroupWise WebAccess](#)” in the *GroupWise 7 Installation Guide*.
- 5 Configure the WebAccess Agent according to the “[GroupWise WebAccess Installation Worksheet](#)” that you filled out in [Section 16.2, “Planning the WebAccess Agent in a Linux Cluster,”](#) on page 187, paying special attention to the cluster resource information on the Server Information page.



The IP address is the secondary IP address associated with the IP address resource in the WebAccess Agent resource group ([WebAccess Agent Clustering Worksheet item 7](#)). The path to the cluster resource mount point is the mount point for the WebAccess Agent domain ([WebAccess Agent Clustering Worksheet item 6](#)).

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

- ◆ The WebAccess Agent startup file (`webac70a.waa`) is created in `/mount_point/groupwise/agents/share` on the shared storage partition so that the WebAccess Agent uses the same startup file regardless of which node it runs 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 node.
- ◆ If this is the first GroupWise agent installed on this node, the GroupWise High Availability service (GWHA) 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 node, the `gwha.conf` file is updated to include the WebAccess Agent.
- ◆ If it is not already present, the OCF Resource Agent script (`GroupWise`) is installed in the `/usr/lib/ocf/resource.d/novell` directory on the node.
- ◆ 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 WebAccess Agent on the initial node is available when you install the WebAccess Agent on subsequent nodes.
- ◆ The WebAccess Agent is not configured to start automatically on system startup. In a cluster, you do not want the WebAccess Agent to start automatically whenever a node restarts.

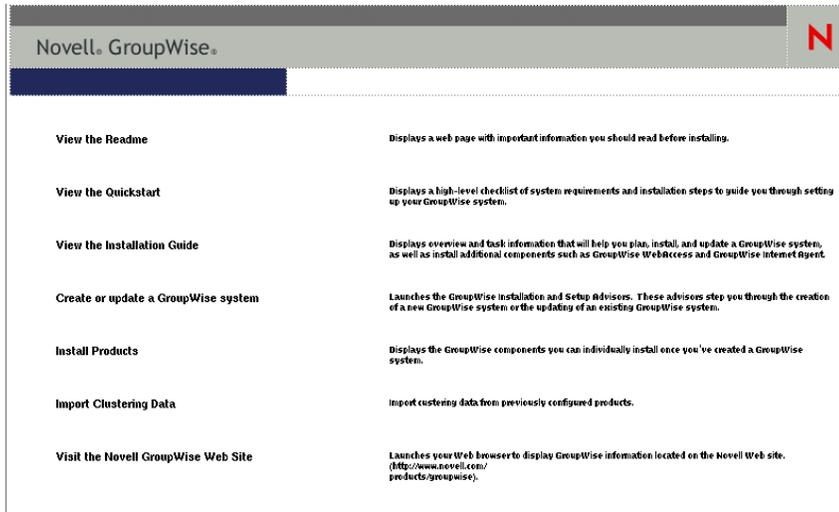
6 Continue with [Running the WebAccess Agent Installation Program on Subsequent Nodes](#).

Running the WebAccess Agent Installation Program on Subsequent Nodes

- 1** In the [HA Management Client](#), select a different node where you want to run the WebAccess Agent ([WebAccess Agent Clustering Worksheet item 5](#)), then click *Make the Node Active*.
- 2** From the software distribution directory you created in [Step 6 in Section 14.1, “Setting Up a New GroupWise System in a Linux Cluster,” on page 143](#), start the GroupWise Installation program and select *Configure GroupWise for Clustering*.



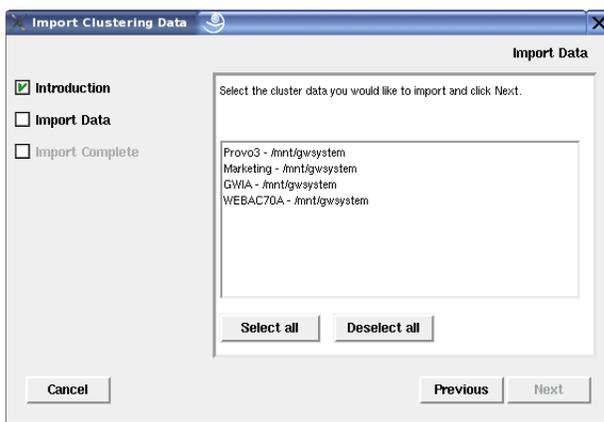
Because of the existence of the `clusterimport.conf` file in the `gwinst` subdirectory of the software distribution directory, a new installation option, *Import Clustering Data*, is now available on the main GroupWise Installation program page.



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

You will use the *Import Clustering Data* option instead of 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 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 node. It is configured specifically for the agents you just selected.
- ♦ The GroupWise High Availability service is automatically configured on the current 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.
- ♦ If it is not already present, the OCF Resource Agent script (`GroupWise`) is installed in the `/usr/lib/ocf/resource.d/novell` directory on the node to increase the capabilities of the HA Management Client.

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

- 6 Continue with [Creating and Configuring GroupWise Cluster Resources for the WebAccess Agent](#).

Creating and Configuring GroupWise Cluster Resources for the WebAccess Agent

Now that you have installed the WebAccess Agent software, you can create and configure GroupWise resources in the cluster resource group you created in [Section 34.3.1, “Setting Up Native Heartbeat Resources for the WebAccess Agent,”](#) on page 347:

- ♦ [“Creating GroupWise Resources for the WebAccess Agent”](#) on page 351
- ♦ [“Enabling Automatic WebAccess Agent Restart”](#) on page 352
- ♦ [“Preventing WebAccess Agent Failback”](#) on page 353

Creating GroupWise Resources for the WebAccess Agent

- 1 In the [HA Management Client](#), select a cluster resource group, then click *Add New Item*.
- 2 In the *Item Type* field, select *native*, then click *OK*.

The screenshot shows the 'Add Native Resource' dialog box. At the top, there is a 'Resource ID' field containing 'resource' and a 'Belong to group:' dropdown menu. Below this is a table with columns 'Name', 'Class/Provider', and 'Description'. The table lists several resources, with 'apache' (ocf/heartbeat) selected. Below the table is a 'Parameters:' section with a table containing one row: 'target_role' with value 'stopped' and description 'press "Default" or "Start" button in toolbar/menu to start the resource'. At the bottom, there are checkboxes for 'Clone' and 'Master/Slave', and input fields for 'clone_max', 'master_max', 'clone_node_max', and 'master_node_max'. The dialog ends with 'Add' and 'Cancel' buttons.

Name	Class/Provider	Description
aaeventd	lsb	AppArmor Notification and Reporting
acpid	lsb	Listen and dispatch ACPI events from the kernel
alsasound	lsb	alsasound
apache	ocf/heartbeat	Apache web server
apache	heartbeat	apache

Name	Value	Description
target_role	stopped	press "Default" or "Start" button in toolbar/menu to start the resource

- 3 In the *Resource ID* field, specify a unique name for the GroupWise resource ([WebAccess Agent Clustering Worksheet item 4](#)).
- 4 In the *Belong to Group* field, select the cluster resource group where you want to create the GroupWise resource for the WebAccess Agent
- 5 In the *Type* list, select *GroupWise (ocf/novell)*.
The OCF Resource Agent script (*GroupWise*) added this item to the list.
- 6 In the *Parameters* list, set the value of the *object_name* parameter to the identifier for the WebAccess Agent ([WebAccess Agent Clustering Worksheet item 4](#)).
For the WebAccess Agent, the identifier is the domain name, followed by the eDirectory object name of the WebAccess Agent (*Domain.webac70a*).
- 7 Click *Add* to create the GroupWise resource.
- 8 Create another GroupWise resource in the cluster resource group for the WebAccess Agent's MTA.
For a domain, the identifier is the domain name.
- 9 Continue with [Enabling Automatic WebAccess Agent Restart](#).

Enabling Automatic WebAccess Agent Restart

If you want Heartbeat to automatically restart a WebAccess Agent that has stopped:

- 1 In the [HA Management Client](#), select the GroupWise resource for the WebAccess Agent, then click *Operations*.

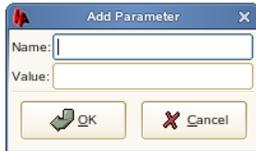


- 2 Click *Add Operation*, then select *monitor*.
- 3 Select the time interval for how often you want the WebAccess Agent to be polled to see if it is still running.
The *GroupWise* script causes the WebAccess Agent to restart, using the GroupWise High Availability service. (GWHA)
- 4 Click *OK* to save the monitoring setting.
- 5 Repeat the process for the GroupWise resource for the MTA in the WebAccess Agent's cluster resource group.
Because Heartbeat can perform such agent monitoring, you do not need to set up GroupWise Monitor to perform this task, as described in "[Configuring the Monitor Agent to Communicate with the High Availability Service](#)" in "[Installing GroupWise Monitor](#)" in the [GroupWise 7 Installation Guide](#). Heartbeat and the GroupWise High Availability Service can work together to keep the WebAccess Agent running on Linux.
- 6 Continue with [Preventing WebAccess Agent Failback](#).

Preventing WebAccess Agent Failback

If you do not want the WebAccess Agent cluster resource group to automatically fail back when the initial node is on-line again:

- 1 In the **HA Management Client**, select the GroupWise cluster resource group for the WebAccess Agent, click *Parameters*, then click *Add Parameter*.



- 2 In the *Name* field, specify `resource_stickiness`.
- 3 In the *Value* field, specify `INFINITY`.
- 4 Click *Apply* to save the setting.
- 5 Continue with **Testing Your WebAccess Agent Installation on Each Node**.

Testing Your WebAccess Agent Installation on Each Node

- 1 Test the WebAccess Agent by starting it with a user interface, as described in “**Starting the Linux WebAccess Agent**” in “**Installing GroupWise WebAccess**” in the *GroupWise 7 Installation Guide*.

```
/opt/novell/groupwise/agents/bin/gwinter --show @webac70a.waa &
```

- 2 Stop the WebAccess Agent by closing the window that it is running in.
- 3 After you can see that the WebAccess Agent stopped successfully, test it by starting it as a daemon, as described in “**Starting the Linux GroupWise Agents as Daemons**” in “**Installing GroupWise Agents**” in the *GroupWise 7 Installation Guide*.

```
/etc/inet.d/grpwise start webac70a  
/etc/inet.d/grpwise status webac70a
```

- 4 Stop the WebAccess Agent.

```
/etc/inet.d/grpwise stop webac70a  
/etc/inet.d/grpwise status webac70a
```

- 5 Also test the MTA on each node, as described in **Section** , “**Testing Your Agent Installation on Each Node,**” on page 150.
- 6 Repeat **Step 1** through **Step 5** for each node where you installed the WebAccess Agent.
- 7 After you have tested the WebAccess Agent on all nodes manually, start and stop the WebAccess Agent resource in the HA Management Client to make sure that it is properly configured.

If you receive the error `Unmanaged` on node, then you have failed to test the WebAccess Agent on one or more nodes. On each node, view the `gwha.conf` file in the `/etc/opt/novell/groupwise` directory. Each WebAccess Agent section should be preceded by `WEBAC70A`. If any section is preceded by a startup file name (for example, `@webac70a.waa`), then that WebAccess Agent has not yet been initialized by being run

manually. To resolve the problem, you can run the WebAccess Agent manually on that node, or you can replace the @webac70a.waa with WEBAC70A, being sure to follow the same format as the other WebAccess Agent section headings in the gwha.conf file.

- 8 When you have installed the WebAccess Agent on all of the nodes where you want it to fail over, continue with [Configuring Clustered Logging](#).

Configuring Clustered Logging

The default location for the WebAccess Agent log files is on the nodes along with the WebAccess Agent software, rather than being located with the domain on the WebAccess Agent shared storage partition. To avoid having multiple copies of log files in multiple locations, you should set the location in ConsoleOne®.

- 1 On the WebAccess Agent shared storage partition where the domain is located, create the directory where you want to store WebAccess Agent log files.
- 2 In ConsoleOne, browse to and select the Domain object.
- 3 Right-click the WebAccess Agent object, then click *Properties*.
- 4 Click *GroupWise > Log Settings*.
- 5 In the *Log File Path* field, browse to and select the directory you created in [Step 1](#), then click *OK*.
- 6 Continue with [Configuring the Clustered WebAccess Agent for SSL](#)

Configuring the Clustered WebAccess Agent for SSL

If you plan to enable SSL, as described in “[Securing WebAccess Agent Connections with SSL](#)” in “[WebAccess](#)” in the *GroupWise 7 Administration Guide*, you need to make the SSL certification file and key file available to the WebAccess Agent in the cluster. The default locations for the SSL certificate file and key file are on the nodes along with the WebAccess Agent software, rather than being located with the domain on a shared storage partition. To avoid having multiple copies of these files in multiple locations, you should set the locations in ConsoleOne.

- 1 On the WebAccess Agent shared storage partition, create the directory where you want to store the certificate file 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 7 Administration Guide*.
- 3 In ConsoleOne, browse to and select the Domain object.
- 4 Right-click the WebAccess Agent 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 After you have set these locations to the WebAccess Agent partition, continue with [Verifying WebAccess Agent Object Properties](#).

Verifying WebAccess Agent Object Properties

During installation of the WebAccess Agent, the WebAccess Agent 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 WebAccess Agent.

- ♦ “[Accessing WebAccess Agent Object Properties](#)” on page 355
- ♦ “[Verifying Post Office Links](#)” on page 355
- ♦ “[Forcing Use of the Web Access Agent Secondary IP Address](#)” on page 355

Accessing WebAccess Agent Object Properties

- 1 In ConsoleOne, browse to and select the WebAccess Agent domain in order to display its contents.
- 2 Right-click the WebAccess Agent object, then click *Properties*.
- 3 Continue with [Verifying Post Office Links](#).

Verifying Post Office Links

In the WebAccess object property page:

- 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 WebAccess Agent.
- 3 Verify that the *Links* column displays the secondary IP addresses of the cluster resource group 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 Web Access Agent Secondary IP Address](#).

Forcing Use of the Web Access Agent Secondary IP Address

If you want the WebAccess Agent to always use its secondary IP address, rather than using the primary IP address:

- 1 In ConsoleOne, click *GroupWise > Network Address*.
- 2 In the *TCP/IP Address* field, provide the secondary IP address ([WebAccess Agent Clustering Worksheet item 8](#)) for the WebAccess Agent.
- 3 Select *Bind Exclusively to TCP/IP Address*.
- 4 Click *OK*.
- 5 Continue with [Installing and Configuring the WebAccess Application on Your Web Server](#).

34.3.5 Installing and Configuring the WebAccess Application on Your Web Server

The WebAccess Application cannot be installed in a cluster. Follow the standard installation and configuration instructions in “[Installing and Configuring the WebAccess Application and WebPublisher Application](#)” in “[Installing GroupWise WebAccess](#)” in the *GroupWise 7 Installation Guide*.

34.4 Testing the WebAccess Agent in a Heartbeat Cluster

After you have installed and configured the WebAccess Agent in your Heartbeat cluster, and the WebAccess Application with your Web server, you can use the following methods to test whether the WebAccess Agent fails over as expected:

- ♦ In the **HA Management Client**, select a node where the WebAccess Agent is running, then click *Make the Node Standby*. The WebAccess Agent should fail over to another node and start running again. This can also be done from the command line using the **crm_standby utility** (http://www.linux-ha.org/v2/AdminTools/crm_standby).
- ♦ Using the **crm_resource utility** (http://www.linux-ha.org/v2/AdminTools/crm_resource), manually migrate a node where the WebAccess Agent is running.

If you receive the error `Unmanaged on node`, see **Step 7** in “**Testing Your WebAccess Agent Installation on Each Node**” on **page 353** for troubleshooting assistance.

If you have set up your Heartbeat cluster in a test situation first, you can also use the following methods for testing WebAccess Agent failover:

- ♦ Power off the node where the WebAccess Agent is running. The WebAccess Agent should fail over to another node and start running again.
- ♦ Simulate the WebAccess Agent stopping unexpectedly (`kill -9 agent_PID`). If you have configured the WebAccess Agent for automatic restart, as described in “**Enabling Automatic GroupWise Agent Restart**” on **page 316**, the agent should start running again on the same node.

WARNING: These last two tests should not be performed on a production system because they can result in damage to GroupWise databases.

34.5 Managing the WebAccess Agent in a Heartbeat Cluster

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

- ♦ **Section 34.5.1, “Updating GroupWise Objects with Cluster-Specific Descriptions,”** on **page 357**
- ♦ **Section 34.5.2, “Knowing What to Expect in a WebAccess Agent Failover Situation,”** on **page 357**

34.5.1 Updating GroupWise Objects with Cluster-Specific Descriptions

After installing the WebAccess Agent 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 WebAccess Agent Domain and Its MTA” on page 357](#)
- ♦ [“Recording Cluster-Specific Information about the WebAccess Agent” on page 357](#)

Recording Cluster-Specific Information about the WebAccess Agent Domain and Its MTA

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

- 1 In ConsoleOne, browse to and right-click the Domain object, then click *Properties*.
- 2 In the *Description* field of the WebAccess Agent domain Identification page, provide a cluster-specific description of the WebAccess Agent domain, including the cluster resource group and shared storage partition where the domain is located.
- 3 Click *OK* to save the WebAccess Agent domain description.
- 4 Select the WebAccess Agent 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 cluster resource group.
When the MTA runs with a user interface, this information appears on the MTA server console.
- 7 Click *OK* to save the MTA description.
- 8 Continue with [Recording Cluster-Specific Information about the WebAccess Agent](#).

Recording Cluster-Specific Information about the WebAccess Agent

With the contents of the WebAccess Agent domain still displayed:

- 1 Right-click the WebAccess Agent object, then click *Properties*.
- 2 Click *GroupWise > Identification*.
- 3 In the *Description* field, record the secondary IP address of the cluster resource group.
- 4 Click *OK* to save the WebAccess Agent information.
- 5 Continue with [Knowing What to Expect in a WebAccess Agent Failover Situation](#).

34.5.2 Knowing What to Expect in a WebAccess Agent Failover Situation

The failover behavior of the MTA for the WebAccess Agent domain is the same as for an MTA in a regular domain. See [Section 14.6.2, “Knowing What to Expect in MTA and POA Failover Situations,” on page 163](#).

When the WebAccess Agent fails over, the WebAccess client user sees the following message:

```
Unable to communicate with the GroupWise WebAccess Agent
```

The user just needs to be patient. When the WebAccess Agent comes up on the next node, the user can continue working without logging in again and no data is lost.

If the POA for the user's post office fails over, the WebAccess client becomes unresponsive, but there is no error message. Again, the user should be patient until the POA comes up on the next node. Then the user can continue working without logging in again and no data is lost.

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

34.6 WebAccess Agent Clustering Worksheet

Item	Explanation
1) WebAccess Agent Nodes:	<p>List the nodes where you want the WebAccess Agent to fail over. Identify each node by its DNS hostname and the IP address that it is physically configured with.</p> <p>For more information, see Section 34.2.1, "Planning a Cluster Resource Group for the WebAccess Agent," on page 344.</p>
2) WebAccess Agent Cluster Resource Group:	<p>Specify the name of the cluster resource group to hold the cluster resources required for the WebAccess Agent.</p> <p>For more information, see Section 34.2.1, "Planning a Cluster Resource Group for the WebAccess Agent," on page 344.</p>
3) Native Heartbeat Resources:	<p>List names for the native resources and the values for their parameters to be used in the WebAccess Agent cluster resource group.</p> <ul style="list-style-type: none">◆ EVMS container◆ File system◆ Secondary IP address <p>For more information, see Section 34.2.1, "Planning a Cluster Resource Group for the WebAccess Agent," on page 344.</p>
4) GroupWise Resources:	<p>List the names for the GroupWise resources and the values for their parameters to be used in the WebAccess Agent cluster resource group.</p> <ul style="list-style-type: none">◆ MTA◆ WebAccess Agent <p>For more information, see Section 34.2.1, "Planning a Cluster Resource Group for the WebAccess Agent," on page 344</p>
5) WebAccess Agent Constraints:	<p>Specify the Heartbeat constraints you want to put on the WebAccess Agent resource group to control the nodes where it can fail over.</p> <ul style="list-style-type: none">◆ Place◆ Order◆ Colocation <p>For more information, see Section 34.2.1, "Planning a Cluster Resource Group for the WebAccess Agent," on page 344</p>

Item	Explanation
6) WebAccess Agent Domain Name: Domain Directory: Mount Point:	Specify a unique name for the WebAccess Agent domain. Specify the directory on the shared storage partition where you want to create the new domain. Specify the mount point on each node where you plan to mount the shared storage partition. For more information, see Section 16.2.1, “Planning a Domain for the WebAccess Agent,” on page 188.
7) MTA Network Information: <ul style="list-style-type: none"> ◆ Secondary IP address ◆ Message transfer port (default 7100) ◆ HTTP port (default 7180) 	Record the MTA network address information that you will need as you install the MTA. For more information, see Section 16.2.5, “Planning the MTA Installation,” on page 190.
8) WebAccess Agent Network Information: <ul style="list-style-type: none"> ◆ Secondary IP address ◆ HTTP port (default 7211) 	Record the WebAccess Agent network address information that you will need to install the WebAccess Agent. For more information, see Section 16.2.6, “Planning the WebAccess Agent Installation,” on page 190.

Implementing GroupWise Monitor in a Heartbeat Cluster

35

Monitor is not currently supported in a Heartbeat cluster.

Backing Up a GroupWise System in a Heartbeat Cluster

36

To back up GroupWise® data in a Heartbeat cluster, 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 7 Administration Guide*.

Updating a GroupWise System in a Heartbeat Cluster

37

In a Heartbeat 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 nodes where the GroupWise software is already installed.

It is very important to update, at the same time, all nodes where each domain and post office can fail over because each domain and post office should be serviced by only one version of the agent software. If you do not update all nodes 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 7 Installation Guide* to update your GroupWise system in a Linux cluster.

Moving an Existing Linux GroupWise 7 System into a Heartbeat Cluster

If you are adding the high availability benefits of Heartbeat to a Linux GroupWise® 7 system that is already up and running, the first step is to install Heartbeat, as described in [Chapter 29, “Introduction to GroupWise 7 and Heartbeat on Linux,”](#) on page 281.

You do not need to transfer your entire GroupWise system into the cluster all at once. You could transfer individual post offices where the needs for high availability are greatest. You could transfer a domain and all of its post offices at the same time. You might decide that you do not 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 storage partitions in your cluster you want to use for GroupWise domains and post offices.
- Determine the nodes in the cluster where you want GroupWise agents to fail over.

If any nodes in your cluster are servers where GroupWise agents have been running in the past and have been starting automatically, use the following command to make sure that they do not start automatically in the future:

```
chkconfig grpwise off
```

In a cluster, the GroupWise agents should not start automatically when a node restarts.

- Review [Chapter 30, “Planning GroupWise in a Heartbeat Cluster,”](#) on page 285 and [Chapter 31, “Setting Up a Heartbeat Cluster,”](#) on page 297. Fill out the “Heartbeat Clustering Worksheet” on page 293 and the “System Clustering Worksheet” on page 140 to help you decide which domains and post offices you want move to which shared storage partitions.
- Move a domain and/or post office onto the shared storage partition, following the instructions in “Moving a Domain” in “Domains” or “Moving a Post Office” in “Post Offices” in the *GroupWise 7 Administration Guide*.
- Using the information on the “System Clustering Worksheet” on page 140, install the agents as needed for the first clustered domain and/or post office, following the instructions in [Section 14.4, “Installing and Configuring the MTA and the POA in a Cluster,”](#) on page 146.
- Test the first component of your clustered GroupWise system, following the instructions in [Section 14.5, “Testing Your Clustered GroupWise System,”](#) on page 160.
- Take care of the cluster management details described in [Section 14.6, “Managing Your Clustered GroupWise System,”](#) on page 161.

- ❑ Move more domains and post offices into the cluster as needed. If you have GroupWise libraries, see [Section 13.5, “Planning a New Library for a Clustered Post Office,”](#) on page 137.
- ❑ Add other components to your clustered GroupWise system as needed, following the instructions in:
 - ◆ [Chapter 33, “Implementing the Internet Agent in a Heartbeat Cluster,”](#) on page 325
 - ◆ [Chapter 34, “Implementing WebAccess in a Heartbeat Cluster,”](#) on page 343.

Implementing Messenger in a Heartbeat Cluster

39

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 30, “Planning GroupWise in a Heartbeat Cluster,” on page 285](#), [Chapter 31, “Setting Up a Heartbeat Cluster,” on page 297](#), and [Chapter 32, “Setting Up a Domain and a Post Office in a Heartbeat Cluster,” on page 307](#). As part of the process of setting up GroupWise in the cluster, you filled out the [“Heartbeat Clustering Worksheet” on page 293](#). Some of the information from that worksheet is helpful as you implement Messenger in your cluster.

- ♦ [Section 39.1, “Planning Your Messenger System in a Heartbeat Cluster,” on page 369](#)
- ♦ [Section 39.2, “Setting Up Your Messenger System in a Heartbeat Cluster,” on page 371](#)
- ♦ [Section 39.3, “Testing Your Clustered Messenger System,” on page 378](#)
- ♦ [Section 39.4, “Managing Your Clustered Messenger System,” on page 378](#)
- ♦ [Section 39.5, “Messenger Clustering Worksheet,” on page 379](#)

39.1 Planning Your Messenger System in a Heartbeat 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. The [“Messenger Clustering Worksheet” on page 240](#) 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:

- ♦ [Section 39.1.1, “Planning a Cluster Resource Group for the Messenger Agents,” on page 369](#)
- ♦ [Section 39.1.2, “Recording the Messenger Agents’ Secondary IP Address,” on page 370](#)
- ♦ [Section 39.1.3, “Determining Appropriate Constraints for the Messenger Agents,” on page 370](#)
- ♦ [Section 39.1.4, “Planning the Messenger Agent Installation,” on page 371](#)

39.1.1 Planning a Cluster Resource Group for the Messenger Agents

As with the GroupWise agents, the Messenger agents need cluster nodes, native Heartbeat resources, Messenger resources, and a cluster resource group in order to function in your Heartbeat cluster. If you put the Messenger agents in the same cluster resource group, they fail over together. If you put them in separate cluster resource groups, then they fail over separately.

For a review of these cluster components, see [Section 13.1, “Installing Novell Cluster Services on Linux,” on page 134](#).

MESSENGER CLUSTERING WORKSHEET

Under **Item 1: Messenger Agent Nodes**, list the nodes that you want to use for the Messenger agents. Identify each node by its DNS hostname and the IP address that it is physically configured with.

Under **Item 2: Messaging Agent Cluster Resource Group**, specify the name of the cluster resource group to hold the native and Messenger resources associated with the Messaging Agent.

Under **Item 3: Archive Agent Cluster Resource Group**, specify the name of the cluster resource group to hold the native and Messenger resources associated with the Archive Agent.

Under **Item 4: Native Heartbeat Resources**, list the cluster resource names and parameters for the EVMS shared storage partition, the file system, and the secondary IP address that you want to use for the Messenger agents.

Under **Item 5: Messenger Resources**, list the cluster resource names and parameters for the Messaging Agent and the Archive Agent. The identifiers for the Messenger agents are their full eDirectory™ object names. You can look them up in the `/etc/opt/novell/groupwise/gwha.conf` file, which is the configuration file for the GroupWise High Availability service. The required format is `.messagingagent.server_nameserver.messengerservice.organization`.

39.1.2 Recording the Messenger Agents' Secondary IP Address

If you are not anticipating a large Messenger archive, you can use one cluster resource group 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 cluster resource group for the Archive Agent and its archive database. The steps in this section focus on setting up the Messenger agents in a single cluster resource group.

MESSENGER CLUSTERING WORKSHEET

Under **Item 7: Messaging Agent Network Information**, record the secondary IP address of the Messaging Agent cluster resource group.

If you want a separate cluster resource group for archiving, under **Item 8: Archiving Agent Network Information**, record the secondary IP address of the Archiving Agent cluster resource group.

39.1.3 Determining Appropriate Constraints for the Messenger Agents

By default, a Heartbeat cluster is configured to have all nodes in the cluster available for failover. Only one node at a time can have a particular cluster resource group mounted and active. If a cluster resource group's initial node fails, the resource group fails over to another node in the cluster. As with the GroupWise agents, you need to select the nodes in the cluster where you want the Messenger agents to fail over. You must install the Messenger software on all of the nodes where you want the Messenger agents to be able to fail over.

MESSENGER CLUSTERING WORKSHEET

Under **Item 6: Messenger Agent Constraints**, list the nodes where you want the Messenger agents to fail over.

39.1.4 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 *Messenger 2.0 Installation Guide*. Messenger must be installed on each node where you want the Messenger agents to be able to fail over ([Messenger Clustering Worksheet item 6](#))

Continue with [Setting Up Your Messenger System in a Linux Cluster](#).

39.2 Setting Up Your Messenger System in a Heartbeat Cluster

You should have already reviewed [Section 22.1, “Planning Your Messenger System in a Linux Cluster,”](#) on page 229 and filled out the “[Messenger Clustering Worksheet](#)” on page 240 and the “[Novell Messenger Worksheet](#)” in the *Messenger 2.0 Installation Guide*.

- ♦ [Section 39.2.1, “Setting Up Native Heartbeat Resources for the Messenger Agents,”](#) on page 371
- ♦ [Section 39.2.2, “Creating Your Messenger System and Installing the Messenger Agents,”](#) on page 372
- ♦ [Section 39.2.3, “Changing Messenger Paths to Locations on the Messenger Partition,”](#) on page 376

39.2.1 Setting Up Native Heartbeat Resources for the Messenger Agents

- 1 Start the HA Management Client, as described in [Section 31.1, “Starting the HA Management Client,”](#) on page 297.
- 2 Create a one or two cluster resource groups for the Messenger agents ([Messenger Clustering Worksheet item 2](#) and [item 3](#)), as described in [Section 31.2, “Creating a Cluster Resource Group,”](#) on page 298.
- 3 Create the EVMS container resource, the file system resource, and the IP address resource for the cluster resource group ([Messenger Clustering Worksheet item 4](#)), as described in [Section 31.3, “Creating Native Heartbeat Resources,”](#) on page 299.

As opposed to native Heartbeat resources, Messenger resources for the Messenger agents must be created after you install the Messenger software, so that the `GroupWise` script is available to enhance the capabilities of the HA Management Client. By creating native resources before you install the Messenger software, you can use them to conveniently fail over the nodes as you install the Messenger agent software on multiple nodes.

- 4 Set up constraints on the nodes where the Messenger agents can fail over ([Messenger Clustering Worksheet item 5](#)), as described in [Section 31.4, “Creating Heartbeat Constraints,”](#) on page 302
- 5 Continue with [Creating Your Messenger System and Installing the Messenger Agents](#).

39.2.2 Creating Your Messenger System and Installing the 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 where you want the Messenger agents to be able to fail over, in order to install the Messenger software on each node, but you do not create any more objects in eDirectory.

- ♦ “Running the Messenger Installation Program on the Initial Node” on page 372
- ♦ “Running the Messenger Installation Program on Subsequent Nodes” on page 372
- ♦ “Creating and Configuring Messenger Cluster Resources for the Messenger Agents” on page 373
- ♦ “Testing Your Messenger Agent Installation on Each Node” on page 375

Running the Messenger Installation Program on the Initial Node

- 1 In the **HA Management Client**, start the resources in the Messenger resource group (**Messenger Clustering Worksheet item 2**) in order to mount the shared storage partition.
- 2 From the *Messenger 2 Administrator for Linux* CD, 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 *Messenger 2.0 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 partition (**Messenger Clustering Worksheet item 4**).
- 6 When asked if you want to start the agents, enter *n* for *No*.
- 7 Set up your Messenger system, following the steps provided in “Configuring Your Messenger System on Linux” in “Installing a Novell Messenger System” in the *Messenger 2.0 Installation Guide*.
- 8 Continue with **Running the Messenger Installation Program on Subsequent Nodes**.

Running the Messenger Installation Program on Subsequent Nodes

- 1 In the **HA Management Client**, select a different node where you want to run the Messenger agents (**Messenger Clustering Worksheet item 6**), then click *Make the Node Active*.
- 2 From the *Messenger 2 Administrator for Linux* CD, 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 partition (**Messenger Clustering Worksheet item 4**).

The Messenger Installation program then accesses the Messenger files that were created on the shared storage partition when the Messenger agents were installed on the initial 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 where you want the Messenger agents to fail over.
- 9 Continue with **Creating and Configuring Messenger Cluster Resources for the Messenger Agents**.

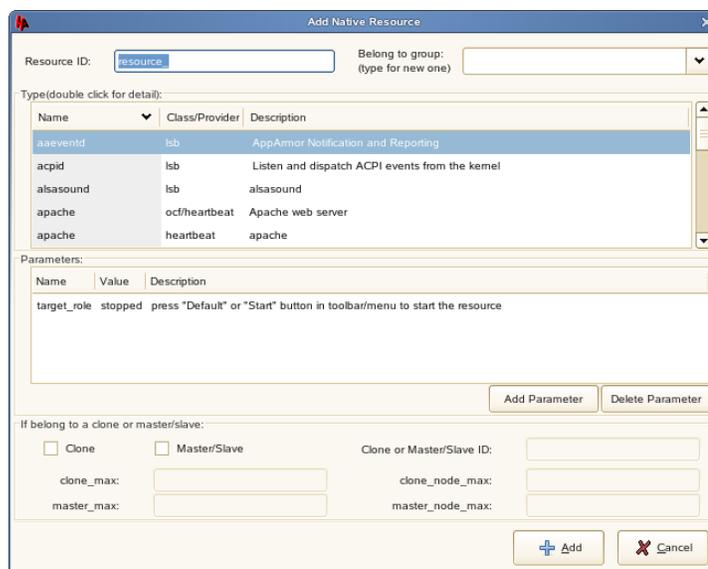
Creating and Configuring Messenger Cluster Resources for the Messenger Agents

Now that you have installed the Messenger agent software, you can create and configure Messenger resources in the cluster resource group you created in **Section 39.2.1, “Setting Up Native Heartbeat Resources for the Messenger Agents,”** on page 371:

- ♦ “Creating Messenger Resources for the Messenger Agents” on page 373
- ♦ “Enabling Automatic Messaging Agent Restart” on page 374
- ♦ “Preventing Messaging Agent Failback” on page 375

Creating Messenger Resources for the Messenger Agents

- 1 In the **HA Management Client**, select the Messaging Agent cluster resource group (**Messenger Clustering Worksheet item 2**}, then click *Add New Item*.
- 2 In the *Item Type* field, select *native*, then click *OK*.



- 3 In the *Resource ID* field, specify a unique name for the Messenger resource (**Messenger Clustering Worksheet item 2**).

- 4 In the *Belong to Group* field, select the Messenger cluster resource group.
- 5 In the *Type* list, select *GroupWise (ocf/novell)*.
The OCF Resource Agent script (*GroupWise*) adds this item to the list. If no *GroupWise* agents are installed on this node, you can copy the *GroupWise* script from the `/usr/lib/ocf/resource.d/novell` directory from another node to this node so that the *GroupWise (ocf/novell)* resource type appears in the list.
- 6 In the *Parameters* list, set the value of the `object_name` parameter to the identifier for the Messaging Agent (**Messenger Clustering Worksheet item 4**).
For the Messaging Agent, the identifier is the full eDirectory object name of the Messaging Agent.
- 7 Click *Add* to create the Messenger resource.
- 8 If desired, create another Messenger resource in the cluster resource group for the Archive Agent.
- 9 Continue with **Enabling Automatic Messaging Agent Restart**.

Enabling Automatic Messaging Agent Restart

If you want Heartbeat to automatically restart a Messaging Agent that has stopped:

- 1 In the **HA Management Client**, select the Messenger resource for the Messaging Agent, then click *Operations*.

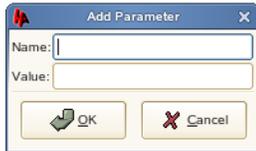


- 2 Click *Add Operation*, then select *monitor*.
- 3 Select the time interval for how often you want the Messaging Agent to be polled to see if it is still running.
The *GroupWise* script causes the Messaging Agent to restart, using the *GroupWise High Availability service*. (GWHA)
- 4 Click *OK* to save the monitoring setting.
- 5 If necessary, repeat the process for the Messenger resource for the Archive Agent.
Because Heartbeat can perform such agent monitoring, you do not need to set up *GroupWise Monitor* to perform this task, as described in “**Configuring the Monitor Agent to Communicate with the High Availability Service**” in “**Installing GroupWise Monitor**” in the *GroupWise 7 Installation Guide*. Heartbeat and the *GroupWise High Availability Service* can work together to keep the Messenger agents running on Linux.
- 6 Continue with **Preventing Messaging Agent Failback**.

Preventing Messaging Agent Failback

If you do not want the Messaging Agent cluster resource group to automatically fail back when the initial node is on-line again:

- 1 In the **HA Management Client**, select the Messenger cluster resource group for the Messaging Agent, click *Parameters*, then click *Add Parameter*.



- 2 In the *Name* field, specify `resource_stickiness`.
- 3 In the *Value* field, specify `INFINITY`.
- 4 Click *Apply* to save the setting.
- 5 If necessary, repeat the process for the Archive Agent.
- 6 Continue with **Testing Your Messenger Agent Installation on Each Node**.

Testing Your 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 *Messenger 2.0 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 Repeat **Step 1** and **Step 2** for each node where you installed the Messenger software.
- 4 After you have tested the Messenger agents on all nodes manually, start and stop the each Messenger agent resource in the HA Management Client to make sure that it is properly configured.

If you receive the error `Unmanaged` on node, then you have failed to test the Messenger agents on one or more nodes. On each node, view the `gwha.conf` file in the `/etc/opt/novell/groupwise` directory. Each Messaging Agent section should be preceded by:

```
.messagingagent.server_nameserver.messengerservice.organization
```

Each Archive Agent section should be preceded by:

```
.archiveagent.server_nameserver.messengerservice.organization.
```

If any section is preceded by a startup file name (for example, `@startup.ma`), then that Messenger agent has not yet been initialized by being run manually. To resolve the problem, you can run the Messenger agent manually on that node, or you can replace the `@startup.ma` or `@startup.aa` with:

```
.messagingagent.server_nameserver.messengerservice.organization  
.archiveagent.server_nameserver.messengerservice.organization
```

being sure to follow the same format as the other Messenger agent section headings in the `gwha.conf` file.

- 5 When you have installed the Messenger agents on all of the nodes where you want them to fail over, continue with [Changing Messenger Paths to Locations on the Messenger Partition](#).

39.2.3 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 shared storage 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:

- ♦ [“Setting the Store Path” on page 376](#)
- ♦ [“Setting the Messaging Agent Queue Path” on page 376](#)
- ♦ [“Setting the Archive Agent Queue Path” on page 377](#)
- ♦ [“Setting the Messaging Agent Log Path” on page 377](#)
- ♦ [“Setting the Archive Agent Log Path” on page 378](#)

After setting these locations, continue with [Section 22.3, “Testing Your Clustered Messenger System,” on page 239](#).

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 shared storage partition.

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

- 1 Choose a directory for the Messaging Agent queue and create that directory on the shared storage 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 shared storage partition.

- 1 Choose a directory for the Archive Agent queue and create that directory on the shared storage 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 *Messaging 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 shared storage partition.

- 1 Choose a directory for the Messaging Agent log files and create that directory on the shared storage 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 shared storage partition.

- 1 Choose a directory for the Archive Agent queue and create that directory on the shared storage 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*.

39.3 Testing Your Clustered Messenger System

After you have created and configured the Messenger cluster resources, you can use the following methods to test whether the Messenger agents fail over as expected:

- ♦ In the **HA Management Client**, select a node where a Messenger agent is running, then click *Make the Node Standby*. The agent should fail over to another node and start running again. This can also be done from the command line using the `crm_standby` utility (http://www.linux-ha.org/v2/AdminTools/crm_standby).
- ♦ Using the `crm_resource` utility (http://www.linux-ha.org/v2/AdminTools/crm_resource), manually migrate a node where a Messenger agent is running.

If you receive the error `Unmanaged on node`, see **Step 4** in “**Testing Your Messenger Agent Installation on Each Node**” on **page 375** for troubleshooting assistance.

If you have set up your Heartbeat cluster in a test situation first, you can also use the following methods for testing failover:

- ♦ Power off the node where a Messenger agent is running. The agent should fail over to another node and start running again.
- ♦ Simulate the Messenger agent stopping unexpectedly (`kill -9 agent_PID`). If you have configured the Messenger agent resource for automatic restart, as described in “**Enabling Automatic GroupWise Agent Restart**” on **page 316**, the agent should start running again on the same node.

WARNING: These last two tests should not be performed on a production system because they can result in damage to Messenger databases.

39.4 Managing Your Clustered Messenger System

If the node where your Messenger system is running goes down, it fails over to another node in the cluster. 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.

39.5 Messenger Clustering Worksheet

Item	Explanation
1) Messenger Agent Nodes:	<p>List the nodes where you want the Messenger agents to fail over. Identify each node by its DNS hostname and the IP address that it is physically configured with.</p> <p>For more information, see Section 22.1.1, “Understanding Your Cluster,” on page 229.</p>
2) Messaging Agent Cluster Resource Group:	<p>Specify the name of the cluster resource group to hold the resources required for the Messaging Agent.</p>
Use Same Resource Group for Archive Agent?	<p>For more information, see “Understanding Your Cluster” on page 229.</p> <ul style="list-style-type: none"> ◆ Yes ◆ No
3) Archive Agent Cluster Resource Group:	<p>Specify the name of the cluster resource group to hold the resources required for the Archive Agent.</p> <p>For more information, see “Understanding Your Cluster” on page 229.</p>
4) Native Heartbeat Resources:	<p>List names for the native Heartbeat resources and the values for their parameters to be used in the Messenger agent cluster resource group.</p> <ul style="list-style-type: none"> ◆ EVMA container ◆ File system ◆ Secondary IP address <p>For more information, see “Understanding Your Cluster” on page 229.</p>
5) Messenger Resources:	<p>List the names for the Messenger resources and the values for their parameters to be used in the Messenger agent cluster resource group.</p> <ul style="list-style-type: none"> ◆ Messaging Agent ◆ Archive Agent <p>For more information, see “Understanding Your Cluster” on page 229.</p>
6) Messenger Agent Constraints:	<p>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.</p> <p>For more information, see “Determining an Appropriate Failover List for the Messenger Agents” on page 230.</p>
7) Messaging Agent Network Information	<p>Record the Messaging Agent network address information that you will need to install the Messaging Agent.</p> <ul style="list-style-type: none"> ◆ Secondary IP address <p>For more information, see Section 22.1.2, “Selecting the Messenger Partition and Secondary IP Address,” on page 230.</p>

Item	Explanation
8) Archive Agent Network Information <ul style="list-style-type: none"><li data-bbox="280 310 545 338">◆ Secondary IP address	Record the Archive Agent network address information that you will need to install the Messaging Agent. For more information, see Section 22.1.2, “Selecting the Messenger Partition and Secondary IP Address,” on page 230.

Microsoft Clustering Services on Windows

VI

- ◆ Chapter 40, “Introduction to GroupWise 7 and Microsoft Clusters,” on page 383
- ◆ Chapter 41, “Planning GroupWise in a Microsoft Cluster,” on page 385
- ◆ Chapter 42, “Setting Up a Domain and Post Office in a Microsoft Cluster,” on page 401
- ◆ Chapter 43, “Implementing the Internet Agent in a Microsoft Cluster,” on page 413
- ◆ Chapter 44, “Implementing WebAccess in a Microsoft Cluster,” on page 423
- ◆ Chapter 45, “Implementing GroupWise Gateways in a Microsoft Cluster,” on page 435
- ◆ Chapter 46, “Monitoring a GroupWise System in a Microsoft Cluster,” on page 437
- ◆ Chapter 47, “Backing Up a GroupWise System in a Microsoft Cluster,” on page 439
- ◆ Chapter 48, “Moving an Existing GroupWise 7 System into a Microsoft Cluster,” on page 441
- ◆ Chapter 49, “Implementing Messenger in a Microsoft Cluster,” on page 443

Introduction to GroupWise 7 and Microsoft Clusters

40

Before implementing GroupWise[®] 7 in a Microsoft* cluster, make sure you have a solid understanding of Microsoft clustering technologies by reviewing the following information resources:

- ♦ [Cluster Technologies Community Center \(http://www.microsoft.com/windowsserver2003/community/centers/clustering/more_resources.asp\)](http://www.microsoft.com/windowsserver2003/community/centers/clustering/more_resources.asp)
- ♦ [Windows 2003 Clustering Services \(http://www.microsoft.com/windowsserver2003/technologies/clustering/default.mspix\)](http://www.microsoft.com/windowsserver2003/technologies/clustering/default.mspix)
- ♦ [Windows 2000 Clustering Technologies \(http://www.microsoft.com/windows2000/technologies/clustering/default.asp\)](http://www.microsoft.com/windows2000/technologies/clustering/default.asp)

When you review the information 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 would fail over to another node so that users could 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 WebAccess, each resource group could 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.

Planning GroupWise in a Microsoft Cluster

41

The majority of this part of the guide ([Chapter 41, “Planning GroupWise in a Microsoft Cluster,” on page 385](#) through [Chapter 47, “Backing Up a GroupWise System in a Microsoft Cluster,” on page 439](#)) is designed for those who are creating a new GroupWise® system, or at least new domains and post offices, in a Microsoft cluster. If you already have an existing GroupWise 6.x system and need to configure it to work in a newly installed cluster, see [Chapter 48, “Moving an Existing GroupWise 7 System into a Microsoft Cluster,” on page 441](#).

When you implement a new GroupWise system or a new domain or post office in a Microsoft cluster, overall GroupWise system design does not need to change substantially. For a review, see [“Installing a Basic GroupWise System”](#) in the *GroupWise 7 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 Microsoft 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 and POA)

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:\grpwise` directory on each node or to manually create a `drive:\grpwise` 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” on page 396](#) lists all the information you need as you set up GroupWise in a Microsoft cluster. You should print the worksheet and fill it out as you complete the tasks listed below:

- ♦ [Section 41.1, “Setting Up Your Microsoft Cluster,” on page 386](#)
- ♦ [Section 41.2, “Planning a New Clustered Domain,” on page 386](#)
- ♦ [Section 41.3, “Planning a New Clustered Post Office,” on page 387](#)
- ♦ [Section 41.4, “Planning a New Library for a Clustered Post Office,” on page 388](#)
- ♦ [Section 41.5, “Planning GroupWise Resource Groups,” on page 388](#)
- ♦ [Section 41.6, “Planning Shared Administrative Resources,” on page 389](#)
- ♦ [Section 41.7, “Ensuring Successful Name Resolution for GroupWise Resource Groups,” on page 389](#)
- ♦ [Section 41.8, “Deciding How to Install and Configure the Agents in a Cluster,” on page 391](#)
- ♦ [Section 41.9, “GroupWise Clustering Worksheets,” on page 396](#)

After you have completed the tasks and filled out the [“System Clustering Worksheet” on page 396](#), you are ready to continue with [Chapter 42, “Setting Up a Domain and Post Office in a Microsoft Cluster,” on page 401](#).

41.1 Setting Up Your Microsoft Cluster

As you set up your Microsoft 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 Microsoft 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.
- ◆ In a like manner, WebAccess provides user access to GroupWise mailboxes across the Internet through users' Web browsers. It is another good candidate for the 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 e-mail 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 Internet Agent 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.

41.2 Planning a New Clustered Domain

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

- ◆ **Primary Domain:** If you are setting up a new GroupWise system in a Microsoft cluster, you will be creating the primary domain as you complete the tasks in this section. In preparation, review “[Planning Your Basic GroupWise System](#)”, then print and fill out the “[Basic GroupWise](#)

System Worksheet” in “*Installing a Basic GroupWise System*” in the *GroupWise 7 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 will be creating a new secondary domain. In preparation, review “*Planning a New Domain*”, then print and fill out the “*Domain Worksheet*” in “*Domains*” in the *GroupWise 7 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 [item 17](#). If you are filling out the Domain Worksheet, stop with [item 10](#).

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 3](#), “*Setting Up a Domain and Post Office in a NetWare Cluster*,” on page 43.

41.3 Planning a New Clustered Post Office

The considerations involved in planning a new post office in a Microsoft 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 “*Post Office Worksheet*” in “*Post Offices*” in the *GroupWise 7 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 3](#), “*Setting Up a Domain and Post Office in a NetWare Cluster*,” on page 43.

41.4 Planning a New Library for a Clustered Post Office

The considerations involved in planning a library in a Microsoft 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 7 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 3, “Setting Up a Domain and Post Office in a NetWare Cluster,”](#) on page 43.

41.5 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 could have a single GroupWise resource group so that your whole GroupWise system would fail over together. More typically, multiple domains and post offices are located throughout your organization, so you would 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 could include a domain, a post office, or both, is termed a “GroupWise resource group.”

Each GroupWise resource group has associated with it a list of possible owners. The possible owners are the nodes to which the resource group could 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 Internet Agent or WebAccess Agent 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 names for the Internet Agent or the WebAccess Agent. 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.

41.6 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 could 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.

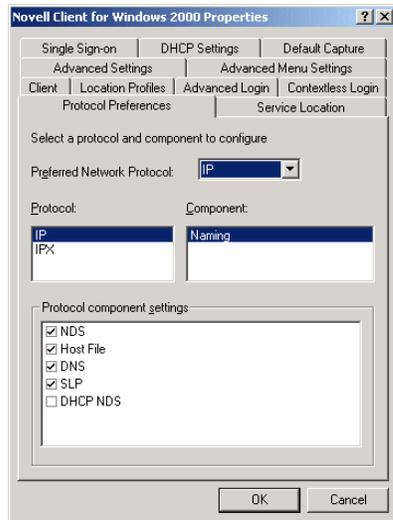
41.7 Ensuring Successful Name Resolution for GroupWise Resource Groups

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:

Figure 41-1 *Novell Client Preferences Property Page*



Short name resolution methods that pertain to your clustered GroupWise system are discussed below:

Table 41-1 *Short Name Resolution Methods*

Short Name Resolution Method	Description
eDirectory	You can use Novell eDirectory™ to resolve short names into specific network addresses. However, when using 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 2000/2003 uses the <code>\winnt\system32\drivers\etc\hosts</code> 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).

Short Name Resolution Method	Description
DNS	<p>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.</p> <p>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.</p>

Specific setup instructions for each of these short name resolution methods are provided in [Chapter 42, "Setting Up a Domain and Post Office in a Microsoft Cluster,"](#) on page 401.

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.

41.8 Deciding How to Install and Configure the Agents in a Cluster

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

- [Section 41.8.1, "Planning Cluster-Unique Port Numbers for Agents in the Cluster,"](#) on page 391
- [Section 41.8.2, "Deciding Where to Install the Agent Software,"](#) on page 393
- [Section 41.8.3, "Planning the Agent Services,"](#) on page 395
- [Section 41.8.4, "Planning the Windows Agent Installation,"](#) on page 395

41.8.1 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.

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 398](#) 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
Development	(same as MTA)	1677	7101	7181
Manufacturing	172.16.5.82	1678	7102	7182

Internet Agent Information

Internet Agent	GWIA IP Address	MTA MTP Port	MTA Live Remote Port	MTA HTTP Port	GWIA HTTP Port
GWIA Domain MTA	172.16.5.83	7110	7677	7183	N/A
Internet Agent (GWIA)	(same as MTA)	N/A	N/A	N/A	9850

WebAccess Information

WebAccess Agent	WebAccess IP Address	MTA MTP Port	MTA HTTP Port	WebAccess Agent Port	WebAccess HTTP Port
WebAccess Domain MTA	172.16.5.84	7120	7184	N/A	N/A
WebAccess Agent (GWINTER)	(same as MTA)	N/A	N/A	7205	7205 (same as agent)

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. The Internet Agent and the WebAccess Agent each have their own domains and separate resource groups.

The example also illustrates that the MTA, the POA, and the Internet Agent use different port numbers for agent ports and HTTP ports. In contrast, the WebAccess Agent uses the same port number for the agent port and the HTTP port.

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.

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 7 Administration Guide*.

NETWORK ADDRESS WORKSHEET

Fill out the “[Network Address Worksheet](#)” on page 398 to help you determine resource group IP addresses and cluster-unique port numbers for all GroupWise agents in the cluster. (MTA, POA, Internet Agent, WebAccess Agent). 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.

41.8.2 Deciding Where to Install the Agent Software

In a Microsoft cluster, the agents must run as Windows services. When you install the Windows MTA and POA, you can choose between two different installation locations:

Table 41-2 *Agent Software Installation Locations*

Location	Description
Each node in the cluster	The <code>c:\grpwise</code> directory is the default location provided by the Agent Installation program.
Shared disk	If you create a <code>drive:\grpwise</code> 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 Microsoft 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 would be best for your clustered GroupWise system:

- ♦ “Advantages of Installing to a Shared Disk” on page 394
- ♦ “Disadvantages of Installing to a Shared Disk” on page 394
- ♦ “Recommendation” on page 394

Advantages of Installing to a Shared Disk

Using a *drive:\grpwise* 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:\grpwise* 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. Accidentally installing a GroupWise Support Pack to the default location of *c:\grpwise* 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:\grpwise* 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:\grpwise* directories on nodes to confuse the issue at a later time.

Even if you choose to install the agents to the `c:\grpwise` 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 only have 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.

41.8.3 Planning the Agent Services

In a Microsoft cluster, the MTA and POA must be set up as service resources. 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.

41.8.4 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 Microsoft 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 7 Installation Guide* for each location where you will install the Windows MTA and/or POA.

Fill out the Windows Agent Worksheet.

GROUPWISE AGENT INSTALLATION WORKSHEET

Under **Item 2: Agents and Locations**, mark *POA Local to Post Office* and *MTA Local to Domain*. In a Microsoft cluster, a domain or post office and its agent must be located on the same node in order to fail over together.

Under **Item 3: Installation Path**, take into account your decision based on **“Deciding Where to Install the Agent Software”** on page 393.

Under **Item 8: Installation Options**, mark *Install as Windows Services*.

Under **Item 6: Domains** and **Item 7: Post Offices**, refer to the Domain and Post Office Worksheets you filled out in **Section 41.2, “Planning a New Clustered Domain,”** on page 386 and **Section 41.3, “Planning a New Clustered Post Office,”** on page 387, and to the Network Address Worksheet you completed during **“Planning Cluster-Unique Port Numbers for Agents in the Cluster”** on page 391.

IMPORTANT: Do not install the Windows agent software until you are instructed to do so in **Chapter 42, “Setting Up a Domain and Post Office in a Microsoft Cluster,”** on page 401.

Skip to **Chapter 42, “Setting Up a Domain and Post Office in a Microsoft Cluster,”** on page 401.

41.9 GroupWise Clustering Worksheets

- ♦ **Section 41.9.1, “System Clustering Worksheet,”** on page 396
- ♦ **Section 41.9.2, “Network Address Worksheet,”** on page 398
- ♦ **Section 41.9.3, “Agent Clustering Worksheet,”** on page 399

41.9.1 System Clustering Worksheet

Item	Explanation
1) Cluster Name:	Record the name of the name of your Microsoft cluster. For more information, see Section 41.1, “Setting Up Your Microsoft Cluster,” on page 386.
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 41.1, “Setting Up Your Microsoft Cluster,” on page 386.

Item	Explanation
3) Resources for GroupWise Administration:	List any shared locations that you want to set up for ConsoleOne or the software distribution directory.
ConsoleOne:	For more information, see Section 41.6, "Planning Shared Administrative Resources," on page 389.
Shared disk:	
Possible owners:	
Software Distribution Directory:	
Shared disk:	
Possible owners:	
4) Resource Group for Domain:	Specify the information for the domain resource group.
Network name:	For more information, see Section 41.2, "Planning a New Clustered Domain," on page 386.
IP address:	
Physical disk:	
File share:	
MTA service:	
Possible owners:	
Post Office in Same Resource Group as Domain?	
♦ Yes	
♦ No	
5) Resource Group for Post Office:	Specify the information for the post office resource group.
Network name:	For more information, see Section 41.3, "Planning a New Clustered Post Office," on page 387.
IP address:	
Physical disk:	
File share:	
POA service:	
Possible owners:	
6) IP Address Resolution Methods:	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.
♦ eDirectory	
♦ hosts file	
♦ DNS	For more information, see Section 41.7, "Ensuring Successful Name Resolution for GroupWise Resource Groups," on page 389
7) Domain Name:	Specify a unique name for the domain. Specify the directory where you want to create the new domain.
Domain Directory:	For more information, see Section 41.2, "Planning a New Clustered Domain," on page 386.
8) Post Office Name:	Specify a unique name for the post office. Specify the directory where you want to create the post office.
Post Office Directory:	For more information, see Section 41.3, "Planning a New Clustered Post Office," on page 387.

Item	Explanation
9) Document Storage Area Location:	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.
<ul style="list-style-type: none"> ◆ At the post office ◆ Outside the post office ◆ Separate post office 	For more information, see Section 41.4, “Planning a New Library for a Clustered Post Office,” on page 388.
Document Storage Area Access	
<ul style="list-style-type: none"> ◆ POA <code>/user</code> startup switch setting ◆ POA <code>/password</code> startup switch setting 	

41.9.2 Network Address Worksheet

- ◆ [“Domain Information”](#) on page 398
- ◆ [“Post Office Information”](#) on page 398
- ◆ [“Internet Agent Information”](#) on page 398
- ◆ [“WebAccess Information”](#) on page 399

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

Internet Agent Information

Internet Agent	GWIA IP Address	MTA MTP Port	MTA Live Remote Port	MTA HTTP Port	GWIA HTTP Port
GWIA Domain MTA					N/A
Internet Agent (GWIA)	(same)	N/A	N/A	N/A	

WebAccess Information

WebAccess Agent	WebAccess IP Address	MTA MTP Port	MTA HTTP Port	WebAccess Agent Port	WebAccess HTTP Port
WebAccess Domain MTA				N/A	N/A
WebAccess Agent (GWINTER)	(same)	N/A	N/A		

41.9.3 Agent Clustering Worksheet

Item	Explanation
1) Agent installation location:	Mark the location where you will install the agent software.
<ul style="list-style-type: none"> ◆ Shared disk with domain or post office ◆ Each node in the cluster 	<p>If necessary, specify the location where you will store agent startup files on the same shared disk with the domain or post office.</p> <p>For more information, see “Deciding Where to Install the Agent Software” on page 393.</p>
Consolidate startup files?	
2) Domain Name:	Transfer this information from the System Clustering Worksheet (item 6).
Domain Directory:	
3) MTA Service Resource:	List other nodes in the cluster where the domain resource group could fail over and any resources that must be online before the MTA can start.
Service name:	
Possible owners:	
Resource dependencies:	For more information, see “Planning the Agent Services” on page 395 .
4) MTA Network Information:	Gather the MTA network address information from the “Network Address Worksheet” on page 398 .
MTA IP address:	
MTA message transfer port:	For more information, see “Planning Cluster-Unique Port Numbers for Agents in the Cluster” on page 391 .
MTA HTTP port:	
5) Post Office Name:	Transfer this information from the System Clustering Worksheet (item 7).
Post Office Directory:	
6) POA Service Resource:	List other nodes in the cluster where post office resource group could fail over and any resources that must be online before the POA can start.
Service name:	
Possible owners:	
Resource dependencies:	For more information, see “Planning the Agent Services” on page 395 .

Item	Explanation
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 398 . For more information, see "Planning Cluster-Unique Port Numbers for Agents in the Cluster" on page 391 .

Setting Up a Domain and Post Office in a Microsoft Cluster

42

You should have already reviewed “[Planning GroupWise in a Microsoft Cluster](#)” on page 385 and filled out the “[System Clustering Worksheet](#)” on page 396, the “[Network Address Worksheet](#)” on page 398, and the “[Agent Clustering Worksheet](#)” on page 399. You are now ready to complete the following tasks to set up GroupWise® in your Microsoft cluster:

- ◆ [Section 42.1, “Preparing the Cluster for GroupWise,”](#) on page 401
- ◆ [Section 42.2, “Setting Up a New GroupWise System in a Cluster,”](#) on page 403
- ◆ [Section 42.3, “Creating a New Secondary Domain in a Cluster,”](#) on page 404
- ◆ [Section 42.4, “Creating a New Post Office in a Cluster,”](#) on page 405
- ◆ [Section 42.5, “Installing and Configuring the MTA and the POA in a Cluster,”](#) on page 406
- ◆ [Section 42.6, “Testing Your Clustered GroupWise System,”](#) on page 409
- ◆ [Section 42.7, “Managing Your Clustered GroupWise System,”](#) on page 409
- ◆ [Section 42.8, “What’s Next,”](#) on page 411

42.1 Preparing the Cluster for GroupWise

After you have set up your Microsoft cluster and become familiar with its functioning, as described in [Chapter 40, “Introduction to GroupWise 7 and Microsoft Clusters,”](#) on page 383, complete the following tasks to prepare the cluster for your GroupWise system:

- ◆ [Section 42.1.1, “Creating GroupWise Resource Groups,”](#) on page 401
- ◆ [Section 42.1.2, “Creating Agent Service Resources,”](#) on page 401
- ◆ [Section 42.1.3, “Configuring Short Name Resolution,”](#) on page 402

42.1.1 Creating GroupWise Resource Groups

Create the needed domain and post office resource groups in your Microsoft cluster ([System Clustering Worksheet items 3 and 4](#)), as planned in [Section 41.2, “Planning a New Clustered Domain,”](#) on page 386 and [Section 41.3, “Planning a New Clustered Post Office,”](#) on page 387.

42.1.2 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 395.

42.1.3 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 41.7, “Ensuring Successful Name Resolution for GroupWise Resource Groups,”](#) on page 389.

- ♦ “eDirectory” on page 402
- ♦ “Hosts Files” on page 402
- ♦ “DNS” on page 402

After configuring your selected short name resolution methods, continue with the task you need to perform:

- ♦ [Section 42.2, “Setting Up a New GroupWise System in a Cluster,”](#) on page 403
- ♦ [Section 42.3, “Creating a New Secondary Domain in a Cluster,”](#) on page 404
- ♦ [Section 42.4, “Creating a New Post Office in a Cluster,”](#) on page 405

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 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 `\winnt\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 could look similar to the following example (all on one line, of course):

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 files, 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 3](#) for *cluster*. Use [System Clustering Worksheet item 4](#) 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.

42.2 Setting Up a New GroupWise System in a 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 Microsoft cluster:

- 1 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 in a Microsoft cluster, 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 “[NetWare and Windows: Setting Up a Basic GroupWise System](#)” in “[Installing a Basic GroupWise System](#)” in the *GroupWise 7 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 providing the MTA and POA network address information, use the Agent Clustering Worksheet that you filled out in [Section 41.8, “Deciding How to Install and Configure the Agents in a Cluster,” on page 391](#). The information you provide will be 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 Cluster](#)” on page 407.

The GroupWise Installation Advisor starts the Agent Installation program for you.

42.3 Creating a New Secondary Domain in a Cluster

After you have set up the primary domain and initial post office, as described in [Section 42.2, “Setting Up a New GroupWise System in a Cluster,” on page 403](#), you can create additional secondary domains as needed.

To create a new secondary domain in a Microsoft cluster:

- 1 Create a domain resource group for the new domain, as described in [“Creating GroupWise Resource Groups” on page 401](#).
- 2 Create an MTA service resource for the domain’s MTA, as described in [“Creating Agent Service Resources” on page 401](#).
- 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 in a clustered environment, [Step 7](#) is easier if the domain directory already exists.
- 5 If you selected the same shared disk with the domain as the agent installation location ([Agent Clustering Worksheet item 1](#)), create the `drive:\grpwise` directory on the drive accessed in [Step 3](#).
or
If you selected `c:\grpwise` 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 7 Administration Guide*](#).
- 7 Create the new domain, following the steps provided in [“Creating the New Domain” in “Domains” in the *GroupWise 7 Administration Guide*](#). Keep in mind the following cluster-specific details:
 - ♦ Use the Domain Worksheet you filled out in [Section 41.2, “Planning a New Clustered Domain,” on page 386](#) 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 391](#) 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 7 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 7 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 Cluster](#).

42.4 Creating a New Post Office in a 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 Microsoft cluster:

- 1 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 in a Microsoft cluster, [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 7 Administration Guide*.
- 4 Create the new post office, following the steps provided in “[Creating the New Post Office](#)” in “[Post Offices](#)” in the *GroupWise 7 Administration Guide*. Keep in mind the following cluster-specific details:
 - ♦ Use the Post Office Worksheet you filled out in [Section 41.3, “Planning a New Clustered Post Office,”](#) on page 387 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 391 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
- ♦ Enable 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 7 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 7 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 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 7 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 Cluster](#).

42.5 Installing and Configuring the MTA and the POA in a 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 42.5.1, “Installing the Agent Software in a Cluster,” on page 407](#)
- ♦ [Section 42.5.2, “Editing Clustered Agent Startup Files,” on page 408](#)
- ♦ [Section 42.5.3, “Setting Up New Instances of the Agents without Installing the Agent Software,” on page 408](#)

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:\grpwise` directory for the resource group.
- ◆ In your GroupWise system, the agent software is already installed to the `c:\grpwise` 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 408 instead of completing the tasks listed above.

42.5.1 Installing the Agent Software in a Cluster

To install the MTA and the POA:

- 1 Map a drive to the shared disk of the domain resource group ([Agent Clustering Worksheet item 2](#)) or the post office resource group ([Agent Clustering Worksheet item 5](#)).
- 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 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:\grpwise` 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:\grpwise` 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 7 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 395 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 If you need to install the agent software to `c:\grpwise` 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 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:\grpwise` directories on the nodes to avoid future confusion.
- 8 Continue with **Editing Clustered Agent Startup Files**.

42.5.2 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.

Because you mapped a drive to the shared disk of the GroupWise resource group 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.

One manual modification of POA startup files is required for robust functionality in a Microsoft 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 42.6, “Testing Your Clustered GroupWise System,”** on page 409.

42.5.3 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.

If the existing agent software is located in the `drive:\grpwise` directory of a shared disk with a domain or post office, the startup files are located there as well. If the existing agent software is located in the `c:\grpwise` 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](#).

42.6 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](#).

42.7 Managing Your Clustered GroupWise System

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

- ♦ [Section 42.7.1, “Updating GroupWise Objects with Cluster-Specific Descriptions,”](#) on page 409
- ♦ [Section 42.7.2, “Knowing What to Expect in MTA and POA Failover Situations,”](#) on page 410

42.7.1 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 that 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”](#) on page 409
- ♦ [“Recording Cluster-Specific Information for a Post Office and Its POA”](#) on page 410
- ♦ [“Recording Cluster-Specific Information for a Software Distribution Directory”](#) on page 410

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 410.

or

Skip to [“Knowing What to Expect in MTA and POA Failover Situations”](#) on page 410.

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](#).

42.7.2 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 would need to restart the GroupWise client would be 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 and then 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.

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](#).

42.8 What's Next

Now that you have at least one GroupWise domain and post office up and running in your Microsoft 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 7 Administration Guide*.
- ◆ Setting up the GroupWise client software and helping users to get started using it. See “[Client](#)” in the *GroupWise 7 Administration Guide*. Also see the *GroupWise 7 Windows Client User Guide*.
- ◆ Connecting your clustered GroupWise system to the Internet. See [Chapter 4, “Implementing the Internet Agent in a NetWare Cluster,”](#) on page 69.
- ◆ Providing access to users’ GroupWise mailboxes from their Web browsers. See [Chapter 5, “Implementing WebAccess in a NetWare Cluster,”](#) on page 89.
- ◆ Connecting your clustered GroupWise system to other e-mail systems through GroupWise gateways. See [Chapter 6, “Implementing GroupWise Gateways in a Novell Cluster,”](#) on page 107.
- ◆ Monitoring the status of your clustered GroupWise system from your Web browser. See [Chapter 7, “Monitoring a GroupWise System in a NetWare Cluster,”](#) on page 109.
- ◆ Backing up your clustered GroupWise system. See [Chapter 8, “Backing Up a GroupWise System in a NetWare Cluster,”](#) on page 111.

Implementing the Internet Agent in a Microsoft Cluster

43

You should already have set up at least a basic GroupWise® system, as described in [Chapter 41](#), “[Planning GroupWise in a Microsoft Cluster](#),” on page 385 and [Chapter 42](#), “[Setting Up a Domain and Post Office in a Microsoft Cluster](#),” on page 401. As part of this process, the “[System Clustering Worksheet](#)” on page 396 and the “[Network Address Worksheet](#)” on page 398 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 in a cluster.

- ◆ [Section 43.1](#), “[Planning the Internet Agent in a Cluster](#),” on page 413
- ◆ [Section 43.2](#), “[Setting Up the Internet Agent in a Cluster](#),” on page 416
- ◆ [Section 43.3](#), “[Managing the Internet Agent in a Cluster](#),” on page 420
- ◆ [Section 43.4](#), “[Internet Agent Clustering Worksheet](#),” on page 422

43.1 Planning the Internet Agent in a 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 each GroupWise gateway, including the Internet Agent. The Internet Agent is faster and more stable when it runs on the same server with its domain. In a cluster, creating a separate domain for the Internet Agent ensures that the Internet Agent and its domain always fail over together.

[Section 43.4](#), “[Internet Agent Clustering Worksheet](#),” on page 422 lists all the information you need as you set up the Internet Agent in a Microsoft cluster. You should print the worksheet and fill it out as you complete the tasks listed below:

- ◆ [Section 43.1.1](#), “[Planning a Domain for the Internet Agent](#),” on page 413
- ◆ [Section 43.1.2](#), “[Planning the Internet Agent Resource Group](#),” on page 414
- ◆ [Section 43.1.3](#), “[Planning Cluster-Unique Port Numbers for the Internet Agent and Its MTA](#),” on page 414
- ◆ [Section 43.1.4](#), “[Preparing Your Firewall for the Internet Agent](#),” on page 415
- ◆ [Section 43.1.5](#), “[Deciding Where to Install the Internet Agent and Its MTA](#),” on page 415
- ◆ [Section 43.1.6](#), “[Planning the MTA Installation](#),” on page 416
- ◆ [Section 43.1.7](#), “[Planning the Internet Agent Installation](#),” on page 416

43.1.1 Planning a Domain for the Internet Agent

The considerations involved in planning a domain for the Internet Agent are much the same as planning any other domain. In preparation, review “[Planning a New Domain](#)”, then print and fill out the “[Domain Worksheet](#)” in “[Domains](#)” in the *GroupWise 7 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 Internet Agent](#), transfer the disk drive to the Internet Agent Clustering Worksheet.

Under [Item 2: Internet Agent Domain Name](#), transfer the domain name and directory to the Internet Agent Clustering Worksheet.

43.1.2 Planning the Internet Agent Resource Group

The Internet Agent resource group is similar to the GroupWise resource groups you have already set up, as described in [Section 41.5, “Planning GroupWise Resource Groups,” on page 388](#) and [“Creating GroupWise Resource Groups” on page 401](#). The Internet Agent resource group contains a domain whose only role is to connect the Internet Agent into your clustered GroupWise system. It also contains two agent service resources, one for the MTA that services the domain and one for the Internet Agent. If you plan to have more than one Internet Agent in the cluster, they must never fail over to the same node.

INTERNET AGENT CLUSTERING WORKSHEET

Under [Item 1: Resource Group for Internet Agent](#), specify the network name and other required information for the Internet Agent resource group.

To ensure successful short name resolution, add entries for the Internet Agent network name to support your preferred methods of short name resolution, as described in [“Configuring Short Name Resolution” on page 402](#).

43.1.3 Planning Cluster-Unique Port Numbers for the Internet Agent and Its MTA

As with the MTA and the POA, the Internet Agent 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 Internet Agent and its MTA as you filled out the [“Network Address Worksheet” on page 398](#). 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 Internet Agent section of the Network Address Worksheet to the Internet Agent Clustering Worksheet.

Under **Item 7: Internet Agent Network Information**, transfer the resource group IP address (the same as for its MTA) and the cluster-unique Internet Agent port number from the Internet Agent section of the Network Address Worksheet to the Internet Agent Clustering Worksheet.

43.1.4 Preparing Your Firewall for the Internet Agent

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

Table 43-1 Standard Ports

Protocol	Standard Port
IMAP4	143
LDAP	389
POP3	110
SMTP	25

By default, the Internet Agent 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 Internet Agent resource group's possible owners list.

If the Internet Agent has a large number of nodes in its possible owners list, you could configure the Internet Agent to send outgoing messages to a relay host, which would then send them out through the firewall using its own IP address rather than the IP address of the particular node where the Internet Agent is running. This reduces the amount of modification to your firewall required to set up the Internet Agent. However, if the relay host goes down, all outgoing messages are delayed.

As another alternative, you can configure the Internet Agent 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 Internet Agent Secondary IP Address" on page 81**, which you can complete after installing the Internet Agent.

In preparation for installing the Internet Agent, configure your firewall as needed to handle the Internet Agent's use of node and resource group IP addresses when sending and receiving messages.

43.1.5 Deciding Where to Install the Internet Agent and Its MTA

The default Internet Agent installation directory is `c:\grpwise\gwia`. As with the MTA and the POA, you can choose to install the Internet Agent and its MTA to each node in the cluster or to the shared disk of the Internet Agent resource group. For a discussion of these alternatives, see

“[Deciding Where to Install the Agent Software](#)” on page 393, which describes the issues in the context of planning MTA and POA installations. As with the MTA and POA, the Internet Agent and its MTA must be installed as Windows services.

INTERNET AGENT CLUSTERING WORKSHEET

Under [Item 4: MTA Installation Location](#) and [Item 6: Internet Agent Installation Location](#), mark whether you will install the Internet Agent and its MTA to the shared disk of the Internet Agent resource group or to each node in the cluster. If necessary, specify where the MTA startup file and the Internet Agent configuration file (`gwia.cfg`) will be stored.

43.1.6 Planning the MTA Installation

Follow the instructions in “[Planning the Windows Agent Installation](#)” on page 395 to plan the MTA installation for the Internet Agent 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 43.2](#), “[Setting Up the Internet Agent in a Cluster](#),” on page 416.

43.1.7 Planning the Internet Agent Installation

Aside from the cluster-specific issues discussed in the preceding sections, the considerations involved in planning to install the Internet Agent are the same in a Microsoft cluster as for any other environment. Review “[NetWare and Windows: Installing the Internet Agent Software](#)”, then print and fill out the “[GroupWise Internet Agent Installation Worksheet](#)” in “[Installing the GroupWise Internet Agent](#)” in the *GroupWise 7 Installation Guide*. You need this information as you install the Internet Agent in your cluster.

IMPORTANT: Do not install the Internet Agent software until you are instructed to do so in [Section 43.2](#), “[Setting Up the Internet Agent in a Cluster](#),” on page 416.

43.2 Setting Up the Internet Agent in a Cluster

You should already have reviewed [Section 43.1](#), “[Planning the Internet Agent in a Cluster](#),” on page 413 and filled out [Section 43.4](#), “[Internet Agent Clustering Worksheet](#),” on page 422. You are now ready to complete the following tasks to set up the Internet Agent in a Microsoft cluster:

- ◆ [Section 43.2.1](#), “[Setting Up the Internet Agent Resource Group](#),” on page 417
- ◆ [Section 43.2.2](#), “[Creating a Domain for the Internet Agent](#),” on page 417
- ◆ [Section 43.2.3](#), “[Installing the MTA for the Internet Agent Domain](#),” on page 417
- ◆ [Section 43.2.4](#), “[Installing and Configuring the Internet Agent in a Cluster](#),” on page 417
- ◆ [Section 43.2.5](#), “[Testing the Clustered Internet Agent](#),” on page 420

43.2.1 Setting Up the Internet Agent Resource Group

- 1 Create the Internet Agent resource group and agent services resources ([Internet Agent Clustering Worksheet item 1](#)), as planned in “[Planning the Internet Agent Resource Group](#)” on page 414.
- 2 To ensure successful short name resolution, add entries for the Internet Agent network name to support your preferred methods of short name resolution, as described in “[Configuring Short Name Resolution](#)” on page 402.
- 3 To ensure that the Internet Agent has incoming and outgoing access to the Internet, make sure your firewall is properly configured, as described in “[Preparing Your Firewall for the Internet Agent](#)” on page 415.
- 4 Continue with [Creating a Domain for the Internet Agent](#).

43.2.2 Creating a Domain for the Internet Agent

The Internet Agent domain will be a secondary domain. To create it, follow the instructions in [Section 42.3, “Creating a New Secondary Domain in a Cluster,”](#) on page 404, 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 Internet Agent domain.

Continue with [Installing the MTA for the Internet Agent Domain](#).

43.2.3 Installing the MTA for the Internet Agent Domain

The MTA for the Internet Agent domain can be installed just like any other MTA in your clustered GroupWise system. Follow the instructions in “[Installing the Agent Software in a Cluster](#)” on page 407, then return to this point.

You do not need to edit the MTA startup file.

Continue with [Installing and Configuring the Internet Agent in a Cluster](#).

43.2.4 Installing and Configuring the Internet Agent in a Cluster

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

- ♦ “[Installing the Internet Agent Software in a Cluster](#)” on page 417
- ♦ “[Enabling Internet Addressing for Your Clustered GroupWise System](#)” on page 418
- ♦ “[Verifying Internet Agent Object Properties](#)” on page 418

Installing the Internet Agent Software in a Cluster

- 1 Map a drive to the shared disk of the Internet Agent 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 Internet Agent as a Windows service ([System Clustering Worksheet item 2](#)).

- 3 If you plan to install the Internet Agent software to the shared disk of the Internet Agent resource group ([Internet Agent Clustering Worksheet item 6](#)), create the `drive:\grpwise\gwia` directory on the shared disk accessed in [Step 1](#).
or
If you plan to install the Internet Agent software to each node in the cluster, create the `c:\grpwise\gwia` directory on the drive accessed in [Step 2](#)).
- 4 Start the Internet Agent Installation program, following the steps provided in “[NetWare and Windows: Installing the Internet Agent Software](#)” in “[Installing the GroupWise Internet Agent](#)” in the *GroupWise 7 Installation Guide*.
- 5 Install the Windows Internet Agent, keeping in mind the following cluster-specific details:
 - ♦ Use the Windows Internet Agent Clustering Worksheet that you filled out in “[Planning the Internet Agent Installation](#)” on page 73 to fill in the fields during the Internet 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 *Run WebAccess 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 Internet Agent 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 Internet Agent software to a shared disk, you need to repeat the installation process for each node so that the Internet Agent gets set up as a Windows service on each node.
- 7 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:\grpwise\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 7 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 Internet Agent is no different from setting it up for an Internet Agent in a any other environment. Follow the instructions in “[Enabling Internet Addressing](#)” in “[System](#)” in the *GroupWise 7 Administration Guide*, then continue with [Verifying Internet Agent Object Properties](#).

Verifying Internet Agent Object Properties

During installation of the Internet Agent, the Internet Agent 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 Internet Agent.

- ♦ “[Accessing Internet Agent Object Properties](#)” on page 419
- ♦ “[Verifying the Reference to the Network Name for Use by DNS](#)” on page 419

- ♦ “Verifying the Reference to the Network Name in Directory Paths” on page 419
- ♦ “Verifying Post Office Links” on page 419
- ♦ “Forcing Use of the Internet Agent Resource Group IP Address” on page 420

Accessing Internet Agent Object Properties

- 1 In ConsoleOne[®], browse to and select the Internet Agent domain in order to display its contents.
- 2 Right-click the Internet Agent object, then click *Properties*.
- 3 Continue with [Verifying the Reference to the Volume Resource](#).

Verifying the Reference to the Network Name for Use by DNS

In the Internet Agent 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 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 Internet Agent 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 Make changes if necessary.
- 4 Continue with [Verifying Post Office Links](#).

Verifying Post Office Links

In the Internet Agent 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 Internet Agent.
- 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 Make changes if necessary.
- 5 Continue with [Forcing Use of the Internet Agent Resource Group IP Address](#).

Forcing Use of the Internet Agent Resource Group IP Address

If you want the Internet Agent 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 ([Internet Agent Clustering Worksheet item 1](#)) for the Internet Agent 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 Internet Agent](#).

43.2.5 Testing the Clustered Internet Agent

After you have set up the Internet Agent resource group, you can test it by manually bringing it online and taking it offline again.

Continue with [Managing the Internet Agent in a Cluster](#).

43.3 Managing the Internet Agent in a Cluster

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

- ♦ [Section 43.3.1, “Updating GroupWise Objects with Cluster-Specific Descriptions,” on page 420](#)
- ♦ [Section 43.3.2, “Knowing What to Expect in an Internet Agent Failover Situation,” on page 421](#)

43.3.1 Updating GroupWise Objects with Cluster-Specific Descriptions

After installing the Internet Agent 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 Internet Agent Domain and Its MTA” on page 420](#)
- ♦ [“Recording Cluster-Specific Information about the Internet Agent” on page 421](#)

Recording Cluster-Specific Information about the Internet Agent Domain and Its MTA

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

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

- 2 In the *Description* field of the Internet Agent domain Identification page, provide a cluster-specific description of the Internet Agent domain, including its resource group IP address and the cluster-unique port numbers used by its MTA.
- 3 Click *OK* to save the Internet Agent domain description.
- 4 Select the Internet Agent 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 Internet Agent**.

Recording Cluster-Specific Information about the Internet Agent

With the contents of the Internet Agent domain still displayed:

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

43.3.2 Knowing What to Expect in an Internet Agent Failover Situation

The failover behavior of the MTA for the Internet Agent 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 410**.

Failover of the Internet Agent itself is more complex. The various e-mail 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 e-mail client and restart it to reestablish the connection after the failover process is complete. Fortunately, the Internet Agent restarts quickly in its failover location so users can reconnect quickly.

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

43.4 Internet Agent Clustering Worksheet

Item	Explanation
1) Resource Group for Internet Agent:	Specify the information for the Internet Agent resource group.
Network name:	For more information, see “Planning the Internet Agent Resource Group” on page 414.
IP address:	
Physical disk:	
File share:	
MTA service resource:	
Internet Agent service resource:	
Possible owners:	
2) Internet Agent Domain Name:	Specify a unique name for the Internet Agent domain. Specify the directory on the physical disk that belongs to the Internet Agent resource group where you want to create the new domain.
Domain Directory:	For more information, see “Planning a Domain for the Internet Agent” on page 413.
4) MTA Installation Location:	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 Internet Agent is installed.
<ul style="list-style-type: none"> ◆ Shared disk of the Internet Agent resource group ◆ Each node in the cluster 	For more information, see “Deciding Where to Install the Internet Agent and Its MTA” on page 415.
Consolidate MTA startup files?	
5) MTA Network Information:	Gather the MTA network address information from the Internet Agent section of the “Network Address Worksheet” on page 398.
MTA IP address:	
MTA message transfer port:	For more information, see “Planning Cluster-Unique Port Numbers for the Internet Agent and Its MTA” on page 414.
MTA live remote port:	
MTA HTTP port	
6) Internet Agent Installation Location:	Mark the location where you will install the Internet Agent software.
<ul style="list-style-type: none"> ◆ Shared disk in the Internet Agent resource group ◆ Each node in the cluster 	If necessary, specify the location on the shared disk of the Internet Agent resource group where you will consolidate the Internet Agent configuration files (<code>gwia.cfg</code>) from the various nodes where it is installed.
Consolidate configuration files?	For more information, see “Deciding Where to Install the Internet Agent and Its MTA” on page 415.
7) Internet Agent Network Information:	Gather the Internet Agent network address information from the Internet Agent section of the “Network Address Worksheet” on page 398.
Internet Agent IP address:	
Internet Agent HTTP port:	For more information, see “Planning Cluster-Unique Port Numbers for the Internet Agent and Its MTA” on page 414.

Implementing WebAccess in a Microsoft Cluster

44

You should already have set up at least a basic GroupWise[®] system, as described in [Chapter 41, “Planning GroupWise in a Microsoft Cluster,” on page 385](#) and [Chapter 42, “Setting Up a Domain and Post Office in a Microsoft Cluster,” on page 401](#). As part of this process, the “[System Clustering Worksheet](#)” on page 396 and the “[Network Address Worksheet](#)” on page 398 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 will need this information as you implement WebAccess in a cluster.

- ♦ [Section 44.1, “Understanding the WebAccess Components,” on page 423](#)
- ♦ [Section 44.2, “Planning WebAccess in a Cluster,” on page 423](#)
- ♦ [Section 44.3, “Setting Up WebAccess in a Cluster,” on page 427](#)
- ♦ [Section 44.4, “Managing WebAccess in a Cluster,” on page 430](#)
- ♦ [Section 44.5, “WebAccess Clustering Worksheet,” on page 433](#)

44.1 Understanding the WebAccess Components

If you are not familiar with GroupWise WebAccess, review “[GroupWise WebAccess Overview](#)” in “[Installing GroupWise WebAccess](#)” in the *GroupWise 7 Installation Guide*.

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

- ♦ WebAccess Agent (`gwinter.exe`) that will be associated with a GroupWise WebAccess domain
- ♦ WebAccess Application (a Java servlet) that will be added to your Web server

44.2 Planning WebAccess in a 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 each GroupWise gateway, including the WebAccess Agent. The WebAccess Agent is faster and more stable when it runs on the same server with its domain. In a cluster, creating a separate domain for the WebAccess Agent ensures that the WebAccess Agent and its domain always fail over together.

[Section 44.5, “WebAccess Clustering Worksheet,” on page 433](#) lists all the information you need as you set up the WebAccess Agent and the WebAccess Application in a clustering environment. You should print the worksheet and fill it out as you complete the tasks listed below:

- ♦ [Section 44.2.1, “Setting Up Your Web Server in the Microsoft Cluster,” on page 424](#)
- ♦ [Section 44.2.2, “Planning a New Domain for the WebAccess Agent,” on page 424](#)
- ♦ [Section 44.2.3, “Planning the WebAccess Resource Group,” on page 425](#)

- ◆ Section 44.2.4, “Planning Cluster-Unique Port Numbers for the WebAccess Agent and Its MTA,” on page 425
- ◆ Section 44.2.5, “Deciding Where to Install the WebAccess Agent and Its MTA,” on page 426
- ◆ Section 44.2.6, “Planning the MTA Installation,” on page 426
- ◆ Section 44.2.7, “Planning the WebAccess Installation,” on page 426

44.2.1 Setting Up Your Web Server in the Microsoft Cluster

Before you install WebAccess, your Web server must already be set up and running in the cluster. Make sure that it can fail over and fail back successfully.

As you set up your Web server, record the following key configuration information on the WebAccess Clustering Worksheet:

WEBACCESS CLUSTERING WORKSHEET

Under **Item 7: Physical Web Servers**, list the nodes in the cluster where you are installing the Web server software.

Under **Item 8: Web Server IP Address**, record the secondary IP address of the Web server resource that you create.

Under **Item 9: Hardware Virtual Server Information**, record the dedicated IP address for the Web site and the document root directory.

Because the WebAccess Application is installed to a subdirectory of the Web server installation directory (*directory\com\novell\webaccess*), the WebAccess Application cannot be installed on a shared disk. Instead, you will install it to each node in the cluster where the Web server has been installed.

44.2.2 Planning a New Domain for the WebAccess Agent

The considerations involved in planning a domain for the WebAccess Agent are much the same as planning any other domain. In preparation, review “**Planning a New Domain**”, then print and fill out the “**Domain Worksheet**” in “**Domains**” in the *GroupWise 7 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 physical disk in your shared disk system 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 WebAccess Clustering Worksheet.

WEBACCESS CLUSTERING WORKSHEET

Under **Item 1: Resource Group for WebAccess Agent**, transfer the shared disk from the Domain Worksheet to the WebAccess Clustering Worksheet.

Under **Item 2: WebAccess Agent Domain Name**, transfer the domain name and directory from the Domain Worksheet to the WebAccess Clustering Worksheet.

44.2.3 Planning the WebAccess Resource Group

The WebAccess resource group is similar to the domain and post office resource groups you have already set up, as described in [Section 41.5, “Planning GroupWise Resource Groups,” on page 388](#) and [“Creating GroupWise Resource Groups” on page 401](#). The WebAccess resource group contains a domain whose only role is to connect the WebAccess Agent into your clustered GroupWise system. It also contains two agent service resources, one for the MTA that services the domain and one for the WebAccess Agent. If you plan to have more than one WebAccess Agent in the cluster, they must never fail over to the same node.

WEBACCESS CLUSTERING WORKSHEET

Under **Item 1: Resource Group for WebAccess**, specify the network name and other required information for the WebAccess resource group.

To ensure successful short name resolution, add entries for the WebAccess network name to support your preferred methods of short name resolution, as described in [“Configuring Short Name Resolution” on page 402](#).

44.2.4 Planning Cluster-Unique Port Numbers for the WebAccess Agent and Its MTA

As with the MTA and the POA, the WebAccess Agent needs cluster-unique port numbers. As part of planning to install the MTA and POA, you should already have determined the IP address and cluster-unique port numbers for the WebAccess Agent and its MTA as you filled out the [“Network Address Worksheet” on page 398](#). If you do not have a filled-out copy of this worksheet for your system, print it now and fill in current system information.

WEBACCESS CLUSTERING WORKSHEET

Under **Item 1: Resource Group for WebAccess**, transfer the WebAccess resource group IP address.

Under **Item 4: MTA Network Information**, transfer the cluster-unique MTA port numbers from the WebAccess section of the Network Address Worksheet to the WebAccess Clustering Worksheet.

Under **Item 6: WebAccess Agent Network Information**, transfer the cluster-unique WebAccess Agent port number from the WebAccess section of the Network Address Worksheet to the WebAccess Clustering Worksheet.

44.2.5 Deciding Where to Install the WebAccess Agent and Its MTA

As with the MTA and the POA, you can choose to install the WebAccess Agent and its MTA to each node in the cluster or to the shared disk of the WebAccess resource group. For a discussion of these alternatives, see [“Deciding Where to Install the Agent Software” on page 393](#), which describes the issues in the context of planning MTA and POA installations.

WEBACCESS CLUSTERING WORKSHEET

Under [Item 3: MTA Installation Location](#) and [Item 5: WebAccess Agent Installation Location](#), mark whether you will install the WebAccess Agent and its MTA to each node in the cluster or to the shared disk of the WebAccess resource group. Also specify where the MTA startup file will be stored.

44.2.6 Planning the MTA Installation

Follow the instructions in [“Planning the MTA Installation” on page 426](#), 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 44.3, “Setting Up WebAccess in a Cluster,” on page 427](#).

44.2.7 Planning the WebAccess Installation

Aside from the cluster-specific issues discussed in the preceding sections, the considerations involved in planning to install WebAccess are the same in a clustering environment as for any other environment. Review [“Planning GroupWise WebAccess”](#), then print and fill out the [“GroupWise WebAccess Installation Worksheet”](#) in [“Installing GroupWise WebAccess”](#) in the *GroupWise 7 Installation Guide*. When you set up WebAccess in a cluster, you will install the WebAccess Agent and the WebAccess Application in two separate steps:

- ◆ [“Planning the WebAccess Agent Installation” on page 426](#)
- ◆ [“Planning the WebAccess Application Installation” on page 427](#)

IMPORTANT: Do not install the WebAccess software until you are instructed to do so in [Section 44.3, “Setting Up WebAccess in a Cluster,” on page 427](#).

Planning the WebAccess Agent Installation

For the WebAccess Agent, fill out items 2 through 12 on the GroupWise WebAccess Installation Worksheet, taking into account the following cluster-specific issues:

WEBACCESS INSTALLATION WORKSHEET

Under **Item 2: Installation Directory**, take into account your decision recorded on the WebAccess Clustering Worksheet (**Item 5: WebAccess Agent Installation Location**).

Under **Item 3: Server Address**, transfer the IP address and port number from the WebAccess Clustering Worksheet (**Item 6: WebAccess Agent Network Information**) filled out in “**Planning Cluster-Unique Port Numbers for the WebAccess Agent and Its MTA**” on page 425.

Under **Item 5: Domain Directory Path**, transfer the domain directory from the Domain Worksheet you filled out in “**Planning a New Domain for the WebAccess Agent**” on page 424.

Planning the WebAccess Application Installation

For the WebAccess Application, fill out items 13 through 19 on the GroupWise WebAccess Installation Worksheet, taking into account the following cluster-specific issues:

WEBACCESS INSTALLATION WORKSHEET

Under **Item 13: Web Server Type and Root Directory**, mark the Web server you have installed in your cluster and specify the Web server root directory.

Under **Item 16: Novell Root Directory**, specify a directory on the Web server where you want to install the WebAccess Agent configuration file. the default is `c:\novell`

44.3 Setting Up WebAccess in a Cluster

You should already have reviewed “**Planning GroupWise in a Microsoft Cluster**” on page 385 and filled out the **Section 44.5, “WebAccess Clustering Worksheet,”** on page 433. You are now ready to complete the following tasks to set up the WebAccess Agent in a clustering environment:

- ◆ **Section 44.3.1, “Setting Up the WebAccess Resource Group,”** on page 427
- ◆ **Section 44.3.2, “Creating a Domain for the WebAccess Agent,”** on page 428
- ◆ **Section 44.3.3, “Installing the MTA for the WebAccess Agent Domain,”** on page 428
- ◆ **Section 44.3.4, “Installing the WebAccess Agent in a Cluster,”** on page 428
- ◆ **Section 44.3.5, “Installing and Configuring the WebAccess Application in a Cluster,”** on page 429
- ◆ **Section 44.3.6, “Testing Your Clustered WebAccess Installation,”** on page 430

44.3.1 Setting Up the WebAccess Resource Group

- 1** Create the WebAccess resource group and agent services resources (**WebAccess Clustering Worksheet item 1**), as planned in “**Planning the WebAccess Resource Group**” on page 425.
- 2** To ensure successful short name resolution, add entries for the WebAccess Agent network name to support your preferred methods of short name resolution, as described in “**Configuring Short Name Resolution**” on page 402.
- 3** Continue with **Creating a Domain for the WebAccess Agent**.

44.3.2 Creating a Domain for the WebAccess Agent

The WebAccess Agent domain will be a secondary domain. To create it, follow the instructions in [Section 42.3, “Creating a New Secondary Domain in a Cluster,” on page 404](#), taking your information from the WebAccess Clustering Worksheet, rather than the System Clustering Worksheet, then return to this point.

Do not create any post offices in the WebAccess Agent domain.

Continue with [Installing the MTA for the WebAccess Agent Domain](#).

44.3.3 Installing the MTA for the WebAccess Agent Domain

The MTA for the WebAccess Agent domain can be installed just like any other MTA in your clustered GroupWise system. Follow the instructions in [“Installing the Agent Software in a Cluster” on page 407](#), then return to this point.

You do not need to edit the MTA startup file.

Continue with [Installing the WebAccess Agent in a Cluster](#).

44.3.4 Installing the WebAccess Agent in a Cluster

After you have created a domain for the WebAccess Agent and installed the MTA for that domain, you are ready to install and configure the WebAccess Agent. The WebAccess Agent is the component of your WebAccess installation that accesses post offices and libraries to retrieve information for WebAccess client users.

- 1 Map a drive to the shared disk of the WebAccess resource group ([WebAccess Clustering Worksheet item 1](#)).
- 2 Map a drive to `c:\` on the first node in the cluster where you will set up the WebAccess Agent as a Windows service ([System Clustering Worksheet item 2](#)).
- 3 If you plan to install the WebAccess Agent software to the shared disk of the WebAccess resource group ([WebAccess Clustering Worksheet item 5](#)), create the `drive:\grpwise\webacc` directory on the WebAccess shared disk accessed in [Step 1](#).
or
If you plan to install the WebAccess Agent software to each node in the cluster, create the `c:\grpwise\webacc` directory on the drive accessed in [Step 2](#).
- 4 Start the WebAccess Installation program, following the steps provided in [“Installing the WebAccess Agent”](#) in [“Installing GroupWise WebAccess”](#) in the [GroupWise 7 Installation Guide](#).
- 5 Install the Windows WebAccess Agent, keeping in mind the following cluster-specific details:
 - ♦ On the Components page select only *GroupWise WebAccess Agent*.
Do not install the WebAccess Application at this time.
 - ♦ Use items 2 through 12 on the GroupWise WebAccess Installation Worksheet that you filled out in [“Planning the WebAccess Installation” on page 426](#) to fill in the fields during the WebAccess Agent installation process.
 - ♦ On the Installation Path page, be sure to browse through the mapped drive to the installation directory you created in [Step 3](#) above.

- ♦ On the Gateway Directory page, be sure to browse to the domain directory through the drive you mapped in **Step 1** above.
 - ♦ On the Execution Options page, be sure that *Run WebAccess Agent as a Windows Service* is selected.
 - ♦ On the Start Applications page, deselect *Start the GroupWise WebAccess Agent*.
- 6** Repeat **Step 4** and **Step 5**, mapping a drive to each node in the cluster.
- Even if you installed the WebAccess Agent software to a shared disk, you need to repeat the installation process for each node so that the Internet Agent gets set up as a Windows service on each node.
- 7** Make sure you have completed all the WebAccess Agent tasks described in “**NetWare and Windows: Setting Up GroupWise WebAccess**” in “**Installing GroupWise WebAccess**” in the *GroupWise 7 Installation Guide*, but do not start the WebAccess Agent at this time.
- 8** Continue with **Installing and Configuring the WebAccess Application in a Cluster**.

44.3.5 Installing and Configuring the WebAccess Application in a Cluster

Recall that the WebAccess Agent is the component of your WebAccess installation that accesses post offices and libraries to retrieve information for WebAccess client users. The WebAccess Application provides the link between the WebAccess Agent and the WebAccess clients’ Web browsers.

To install the WebAccess Application:

- 1** Map a drive to the shared disk of the WebAccess resource group (**WebAccess Clustering Worksheet item 1**) where the WebAccess domain is located.
- 2** Map a drive to the first Web server node where you want to install the WebAccess Application (**WebAccess Clustering Worksheet item 7**).
- 3** If the Web server node where you are going to install the WebAccess Application is currently running any applications that rely on Java or on the Web server, migrate those applications to another node in the cluster. If any GroupWise agents are running on the node, migrate the agents.
- 4** Stop the Web server.
- 5** Start the WebAccess Installation program as you did when you installed the WebAccess Agent (**Step 5 on page 428**). Keep in mind the following cluster-specific details:
 - ♦ On the Components page, select only *GroupWise WebAccess Application*.
 - ♦ Use items 13 through 19 on the GroupWise WebAccess Installation Worksheet that you filled out in “**Planning the WebAccess Installation**” on **page 426** to fill in the fields during the WebAccess Application installation process.
 - ♦ On the Gateway Directory page, be sure to browse to the WebAccess gateway directory (*domain\wpgate\webac70a*) through the drive you mapped in **Step 1** above.
 - ♦ On the Web Server Information page be sure to browse to the Web server root directory through the drive you mapped in **Step 2** above.
 - ♦ On the Start Applications page, deselect *Restart Web Server*.

- 6 Make sure you have completed all the WebAccess Application tasks described in “[NetWare and Windows: Setting Up GroupWise WebAccess](#)” in “[Installing GroupWise WebAccess](#)” in the *GroupWise 7 Installation Guide*.
- 7 Copy the `directory\docs\com` directory from the server where you just installed the WebAccess Application to the document root directory of the Web server.
- 8 Restart the Web server.
- 9 Offline and then online the Web server to reestablish its resource group IP address.
- 10 Repeat [Step 2](#) through [Step 9](#) for each Web server node in the Web server resource group possible owners list ([WebAccess Clustering Worksheet item 1](#)).
- 11 Continue with [Testing Your Clustered WebAccess Installation](#).

44.3.6 Testing Your Clustered WebAccess Installation

Remember that the WebAccess resource group and the Web server resource group are separate resource groups that could fail over to different nodes at different times.

To thoroughly test your WebAccess installation:

- 1 Make sure the initial combination of WebAccess resource group and Web server resource group is functioning properly.
- 2 Migrate the WebAccess resource group to each node in its possible owners list, making sure it functions with the initial Web server node.
- 3 Migrate the Web server to a different node, migrate the WebAccess resource group to each node in its possible owners list, then make sure each combination works.
- 4 Repeat [Step 3](#) for each Web server resource group.
- 5 Continue with [Managing WebAccess in a Cluster](#).

44.4 Managing WebAccess in a Cluster

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

- ♦ [Section 44.4.1, “Updating GroupWise Objects with Cluster-Specific Descriptions,”](#) on page 431
- ♦ [Section 44.4.2, “Knowing What to Expect in WebAccess Failover Situations,”](#) on page 432
- ♦ [Section 44.4.3, “Updating the WebAccess Agent Configuration File \(commgr.cfg\),”](#) on page 432

44.4.1 Updating GroupWise Objects with Cluster-Specific Descriptions

After installing WebAccess 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 WebAccess Agent Domain and Its MTA” on page 431](#)
- ♦ [“Recording Cluster-Specific Information about the WebAccess Agent” on page 431](#)

Recording Cluster-Specific Information about the WebAccess Agent Domain and Its MTA

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

- 1 In ConsoleOne, browse to and right-click the Domain object, then click *Properties*.
- 2 In the *Description* field of the WebAccess Agent domain Identification page, provide a cluster-specific description of the WebAccess Agent domain, including the resource group IP address and the cluster-unique port numbers used by its MTA.

You might also want to include cluster-specific information about the WebAccess Application, such as the resource group IP address of the Web server where the WebAccess Application is installed.

- 3 Click *OK* to save the WebAccess domain description.
- 4 Select the WebAccess 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 WebAccess 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 WebAccess Agent](#).

Recording Cluster-Specific Information about the WebAccess Agent

With the contents of the WebAccess domain still displayed:

- 1 Right-click the WEBAC70A object, then click *Properties*.
- 2 Click *GroupWise > Identification*.
- 3 In the *Description* field, record the WebAccess resource group IP address and the cluster-unique port numbers used by the WebAccess Agent.

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

- 4 Click *OK* to save the WebAccess Agent information.
- 5 Continue with [Knowing What to Expect in WebAccess Failover Situations](#).

44.4.2 Knowing What to Expect in WebAccess Failover Situations

The failover behavior of the MTA for the WebAccess 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 410.

The WebAccess Application caches users’ credentials on the node where it is running. Therefore, if that node fails, or if the WebAccess Application migrates to a different node, the cached credentials are lost. Consequently, the user needs to restart the WebAccess client in order to re-authenticate and re-establish the credentials.

If the WebAccess Agent fails over or migrates, the user receives an error message that the WebAccess Agent is no longer available. However, after the WebAccess Agent starts in its new location, the WebAccess Application passes the cached user credentials to the WebAccess Agent and the user reconnects automatically without having to re-authenticate.

As with the MTA and the POA, migration of the WebAccess Agent takes longer than failover. However, the WebAccess Agent restarts quickly so that users are able to reconnect quickly.

44.4.3 Updating the WebAccess Agent Configuration File (commgr.cfg)

As part of installing WebAccess, the WebAccess Agent configuration file (`commgr.cfg`) is created in the following subdirectory:

```
domain\wpgate\webac70a
```

It is also automatically copied to the following Web server subdirectory:

```
drive:\novell\webaccess
```

If you change WebAccess agent configuration information (for example, if you change its IP address), the information is changed in the following file:

```
domain\wpgate\webac70a\commgr.cfg
```

because the domain is on the shared disk of a resource group, and it is changed in the following file:

```
drive:\novell\webaccess\commgr.cfg
```

on the node where the WebAccess Application is currently running. However, the other nodes in the Web server possible owners list are not currently available for update. Therefore, you must manually copy the updated `commgr.cfg` file to the `drive:\novell\webaccess` subdirectory on each node in the Web serve possible owners list.

44.5 WebAccess Clustering Worksheet

Item	Explanation
1) Resource Group for WebAccess Agent:	Specify the information for the WebAccess resource group.
Group name:	For more information, see “Planning the WebAccess Resource Group” on page 425.
Network name:	
IP address:	
Shared disk:	
Share name:	
MTA service resource:	
WebAccess Agent service resource:	
Possible owners	
2) WebAccess Agent Domain Name:	Specify a unique name for the WebAccess Agent domain. Specify the directory on the WebAccess Agent resource group disk where you want to create the new domain.
Domain Directory:	For more information, see “Planning a New Domain for the WebAccess Agent” on page 424.
3) MTA Installation Location:	Mark the location where you will install the MTA software.
<ul style="list-style-type: none"> ◆ Shared disk of WebAccess resource group ◆ Each node in the cluster 	<p>If necessary, specify the location where you will consolidate the MTA startup files on the shared disk of the WebAccess resource group.</p> <p>For more information, see “Deciding Where to Install the WebAccess Agent and Its MTA” on page 426.</p>
4) MTA Network Information:	Gather the MTA network address information from the WebAccess section of the “Network Address Worksheet” on page 398.
MTA IP address:	
MTA message transfer port:	For more information, see “Planning Cluster-Unique Port Numbers for the WebAccess Agent and Its MTA” on page 425.
MTA live remote port:	
MTA HTTP port:	
5) WebAccess Agent Installation Location:	Mark the location where you will install the WebAccess Agent software.
<ul style="list-style-type: none"> ◆ Shared disk of WebAccess resource group ◆ Each node in the cluster 	<p>For more information, see “Deciding Where to Install the WebAccess Agent and Its MTA” on page 426.</p>
6) WebAccess Agent Network Information:	Gather the WebAccess Agent network address information from the WebAccess section of the “Network Address Worksheet” on page 398.
WebAccess Agent IP address:	
WebAccess Agent HTTP port:	For more information, see “Planning Cluster-Unique Port Numbers for the WebAccess Agent and Its MTA” on page 425.

Item	Explanation
7) Physical Web Servers:	List the servers in the cluster where you are installing the Web server for use with WebAccess. For more information, see “Setting Up Your Web Server in the Microsoft Cluster” on page 424.
8) Web Server IP Address:	Record the secondary IP address for the Web server in the cluster. For more information, see “Setting Up Your Web Server in the Microsoft Cluster” on page 424.
9) Hardware Virtual Server Information:	Record the hardware virtual server information for your shared disk system.
♦ Dedicated IP address: ♦ Document root	For more information, see “Setting Up Your Web Server in the Microsoft Cluster” on page 424.

Implementing GroupWise Gateways in a Microsoft Cluster

45

A significant 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 each GroupWise gateway. The gateway domain should be created in its own resource group. This enables the gateway to fail over independently from other GroupWise components.

If you have set up the Internet Agent or WebAccess in your clustered GroupWise system, you should already have the skills necessary to set up a GroupWise gateway as well.

GroupWise gateways that have not received recent development have not been thoroughly tested in a clustering environment. If you are currently using such GroupWise gateways, you might want to leave them outside of your cluster.

Monitoring a GroupWise System in a Microsoft Cluster

46

GroupWise® Monitor is similar to WebAccess in that it relies on a Web server for communication with administrators' Web browsers. Consequently, the setup procedure for GroupWise Monitor in a Microsoft cluster is similar to the setup procedure for WebAccess. If you have set up WebAccess in your clustered GroupWise system, you should already have the skills necessary to set up GroupWise Monitor as well.

When you first install Monitor, 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 Monitor 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 7 Installation Guide*.

Backing Up a GroupWise System in a Microsoft Cluster

47

The issues involved in backing up a GroupWise® system in a Microsoft 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/partnerguid>).

Moving an Existing GroupWise 7 System into a Microsoft Cluster

48

If you are adding the high availability benefits of a Microsoft cluster to a GroupWise® 7 system that is already up and running, the first step is to set up the cluster and review [Chapter 40, “Introduction to GroupWise 7 and Microsoft Clusters,” on page 383](#) 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 could transfer individual post offices where the needs for high availability are greatest. You could 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 41, “Planning GroupWise in a Microsoft Cluster,” on page 385](#). Fill out the “[System Clustering Worksheet](#)” on page 396 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 391 and fill out the “[Network Address Worksheet](#)” on page 398 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 401 and “[Configuring Short Name Resolution](#)” on page 402.
- 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 7 Administration Guide*.
- Review [Section 41.8, “Deciding How to Install and Configure the Agents in a Cluster,” on page 391](#), fill out the “[Agent Clustering Worksheet](#)” on page 399, and install the agents as needed for the first clustered domain and/or post office, following the instructions in [Section 42.5, “Installing and Configuring the MTA and the POA in a Cluster,” on page 406](#).
- Test the first component of your clustered GroupWise system following the instructions in [Section 42.6, “Testing Your Clustered GroupWise System,” on page 409](#).
- Take care of the cluster management details described in [Section 42.7, “Managing Your Clustered GroupWise System,” on page 409](#).
- Move more domains and post offices into the cluster as needed. If you have GroupWise libraries, see [Section 41.4, “Planning a New Library for a Clustered Post Office,” on page 388](#).

- ❑ Move GroupWise administration into the cluster as needed.
- ❑ Add other components to your clustered GroupWise system as needed, following the instructions in:
 - ◆ [Chapter 43, “Implementing the Internet Agent in a Microsoft Cluster,” on page 413](#)
 - ◆ [Chapter 44, “Implementing WebAccess in a Microsoft Cluster,” on page 423.](#)
 - ◆ [Chapter 45, “Implementing GroupWise Gateways in a Microsoft Cluster,” on page 435](#)
 - ◆ [Chapter 46, “Monitoring a GroupWise System in a Microsoft Cluster,” on page 437](#)
 - ◆ [Chapter 47, “Backing Up a GroupWise System in a Microsoft Cluster,” on page 439](#)

Implementing Messenger in a Microsoft Cluster

49

Novell® Messenger does not require the existence of a GroupWise® system in your Microsoft cluster, but presumably one has already been set up as described in [Chapter 41, “Planning GroupWise in a Microsoft Cluster,” on page 385](#) and [Chapter 42, “Setting Up a Domain and Post Office in a Microsoft Cluster,” on page 401](#). As part of the process of setting up GroupWise in your cluster, you filled out the [“System Clustering Worksheet” on page 396](#). Some of the information from this worksheet will be helpful as you implement Messenger in your cluster.

- ♦ [Section 49.1, “Planning Your Messenger System in a Cluster,” on page 443](#)
- ♦ [Section 49.2, “Setting Up Your Messenger System in a Cluster,” on page 446](#)
- ♦ [Section 49.3, “Messenger Clustering Worksheet,” on page 447](#)

49.1 Planning Your Messenger System in a 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 49.3, “Messenger Clustering Worksheet,” on page 447](#) 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:

- ♦ [Section 49.1.1, “Understanding Your Cluster,” on page 443](#)
- ♦ [Section 49.1.2, “Planning Messenger Administration,” on page 443](#)
- ♦ [Section 49.1.3, “Deciding Where to Install the Messenger Agent Software,” on page 444](#)
- ♦ [Section 49.1.4, “Planning the Messenger Agent Installation,” on page 445](#)

49.1.1 Understanding Your Cluster

Fill out items 1 and 2 in [Section 49.3, “Messenger Clustering Worksheet,” on page 447](#) with information about your cluster. This information corresponds to items 1 and 2 on the [“System Clustering Worksheet” on page 396](#) that you filled out for GroupWise. For background information, see [Section 41.1, “Setting Up Your Microsoft Cluster,” on page 386](#).

49.1.2 Planning Messenger Administration

If you have set up a shared disk for GroupWise administration, as described in [Section 41.6, “Planning Shared Administrative Resources,” on page 389](#), 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 could fail over to.

49.1.3 Deciding Where to Install the Messenger Agent Software

In a Microsoft cluster, the Messenger agents must run as Windows services. When you install the Windows Messenger Agents, you can choose between two different installation locations:

Table 49-1 Messenger Agent Software Installation Locations

Location	Description
Each node in the cluster	The <code>c:\novell\nm</code> directory is the default installation location provided by the Messenger Installation program.
Shared disk	If you create a <code>drive:\novell\nm</code> 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 Microsoft 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 444](#)
- ◆ [“Planning the Messenger Agents on a Shared Disk” on page 445](#)

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 445](#).

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**.

49.1.4 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 *Messenger 2.0 Installation Guide*. Transfer the following information from the Messenger Clustering Worksheet to the Messenger System Worksheet:

- ◆ For **Item 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 **Item 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 Microsoft 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 Cluster**.

49.2 Setting Up Your Messenger System in a Cluster

You should have already reviewed [Section 49.1, “Planning Your Messenger System in a Cluster,”](#) on page 443 and filled out [Section 49.3, “Messenger Clustering Worksheet,”](#) on page 447 and the “[Novell Messenger Worksheet](#)” in the *Messenger 2.0 Installation Guide*. Follow the instructions for the installation location you have chosen:

- ♦ [Section 49.2.1, “Installing the Messenger Agents to Each Node in the Cluster,”](#) on page 446
- ♦ [Section 49.2.2, “Installing the Messenger Agents to a Shared Disk,”](#) on page 446

49.2.1 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 *Messenger 2.0 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 *Messenger 2.0 Installation Guide*.

49.2.2 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 446
- ♦ [“Running the Messenger Installation Program”](#) on page 447
- ♦ [“Testing the Clustered Messenger Agents”](#) on page 447

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 444.
- 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 402.
- 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 “[Setting Up Your Novell Messenger System](#)” in “[Installing a Novell Messenger System](#)” in the *Messenger 2.0 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 “[Planning the Messenger Agent Installation](#)” on [page 122](#) 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 Microsoft 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 *Messenger 2.0 Installation Guide*

49.3 Messenger Clustering Worksheet

Item	Explanation
1) Cluster Identification:	Record the name and IP address of your Microsoft cluster.
Cluster name:	For more information, see Section 41.1, “Setting Up Your Microsoft Cluster,” on page 386 .
Cluster IP address	
2) Nodes in Cluster:	List the servers that are included in your Microsoft cluster.
	For more information, see Section 41.1, “Setting Up Your Microsoft Cluster,” on page 386 .

Item	Explanation
<p>3) Installation Location for Messenger Agents:</p> <ul style="list-style-type: none"> ◆ Each node in the cluster Consolidate startup files? ◆ Shared disk 	<p>Mark the location where you will install the Messenger agent software.</p> <p>For more information, see “Deciding Where to Install the Messenger Agent Software” on page 444.</p>
<p>4) Resource Group for Messenger Agents</p> <p>Network name: IP address: Physical disk: File share: Messaging Agent service: Archive Agent service: Possible owners</p>	<p>If you plan to install the Messenger agent software to a shared disk, provide the information about the shared disk you want to use.</p> <p>For more information, see “Planning the Messenger Agents on a Shared Disk” on page 445.</p>
<p>5) Installation Location for Messenger Administration:</p> <ul style="list-style-type: none"> ◆ Administrator workstation(s) ◆ Shared disk 	<p>Mark the location where you want to install the Messenger snap-in to ConsoleOne.</p> <p>For more information, see “Planning Messenger Administration” on page 443.</p>
<p>6) Resource for Messenger Administration:</p> <p>Network name: IP address: Physical disk: File share: Possible owners</p>	<p>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.</p> <p>For more information, see Section 41.6, “Planning Shared Administrative Resources,” on page 389.</p>
<p>7) IP Address Resolution Methods:</p> <ul style="list-style-type: none"> ◆ eDirectory ◆ hosts file ◆ DNS 	<p>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.</p> <p>For more information, see Section 41.7, “Ensuring Successful Name Resolution for GroupWise Resource Groups,” on page 389.</p>

Non-GroupWise Clients

VII

If your users already have a common POP, IMAP, or SOAP e-mail client that comes with Linux or Windows, they can continue to use it to access their GroupWise® mailboxes. Users of non-GroupWise® e-mail clients retain the feature sets of their familiar e-mail clients, but many GroupWise features are not available to such users because they are not offered in POP, IMAP, and SOAP e-mail clients. For example, calendaring is available only if the POP, IMAP, or SOAP client supports iCal.

- ♦ [Chapter 50, “Outlook Express,” on page 451](#)
- ♦ [Chapter 51, “Microsoft Outlook,” on page 453](#)
- ♦ [Chapter 52, “Evolution,” on page 455](#)

Outlook Express

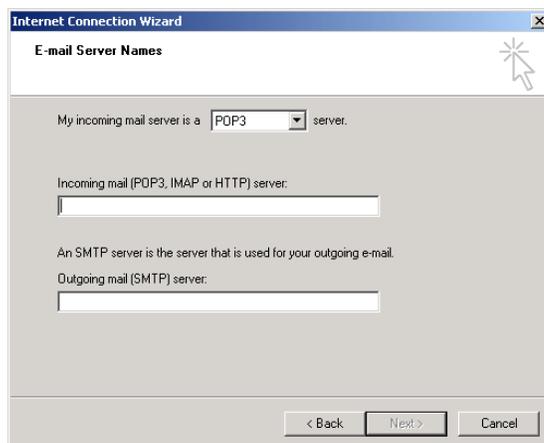
50

The GroupWise[®] Internet Agent is required in order for users to access their mailboxes using non-GroupWise clients. If you have not already installed the Internet Agent, follow the instructions in the *GroupWise 7 Installation Guide*.

In order for users to access their GroupWise mailboxes from a third-party e-mail client, they must configure their e-mail 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 e-mail 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 username and password, then click *OK*.

Microsoft Outlook

51

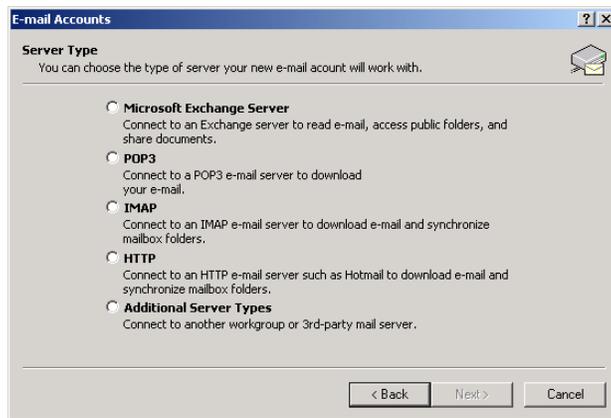
The GroupWise® Internet Agent is required in order for users to access their mailboxes using non-GroupWise clients. If you have not already installed the Internet Agent, follow the instructions in the *GroupWise 7 Installation Guide*, available on the [GroupWise 7 Documentation Web site \(http://www.novell.com/documentation/gw7\)](http://www.novell.com/documentation/gw7).

If your users have been using the Microsoft Outlook e-mail 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, then click *OK*.
- 4 Select *Add a New E-Mail Account*, then click *Next*.



- 5 Select *POP3* or *IMAP* as your incoming mail server type, then click *Next*.

- 6** Provide the e-mail account settings for the type of server you selected.
- 7** Click *Test Account Settings* to make sure that you have provided the information correctly.
- 8** Click *Next*, then click *Finish*.

You can now access your GroupWise mailbox using Microsoft Outlook by selecting the profile you just created.

For additional functionality, you can install the GroupWise Connector for Microsoft Outlook, available along with the GroupWise client from the [Novell Product Downloads site \(http://download.novell.com\)](http://download.novell.com). For more information, see the *GroupWise Connector for Microsoft Outlook Quick Start* on the [GroupWise 7 Documentation Web site \(http://www.novell.com/documentation/gw7\)](http://www.novell.com/documentation/gw7)

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, an integral part of the Internet-connected desktop.

Evolution can help you work in a group by handling e-mail, 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 e-mail, or to send e-mail to a contact or appointment. People who get lots of e-mail will appreciate advanced features like vFolders, which let you save searches as though they were ordinary e-mail folders.

If you have Evolution 2.4 installed, you can access accounts on Novell® GroupWise® 7.

- ♦ [Section 52.1, “GroupWise Features Available in Evolution,” on page 455](#)
- ♦ [Section 52.2, “Configuring Evolution,” on page 456](#)

52.1 GroupWise Features Available in Evolution

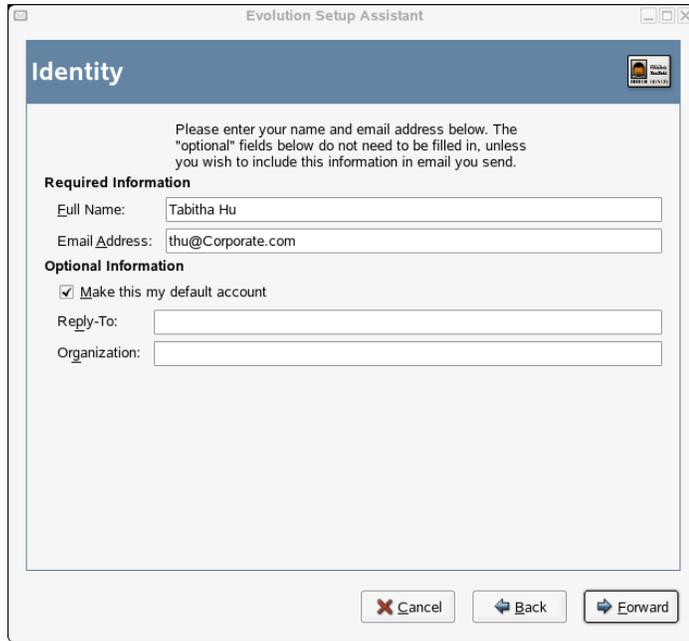
Evolution connecting to GroupWise supports the following basic GroupWise features:

- ♦ 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 is supported 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 e-mail messages with a single click.
- ♦ Proxy
 - ♦ Assign Proxy access to other users.
 - ♦ View other users' accounts through Proxy access.

52.2 Configuring Evolution

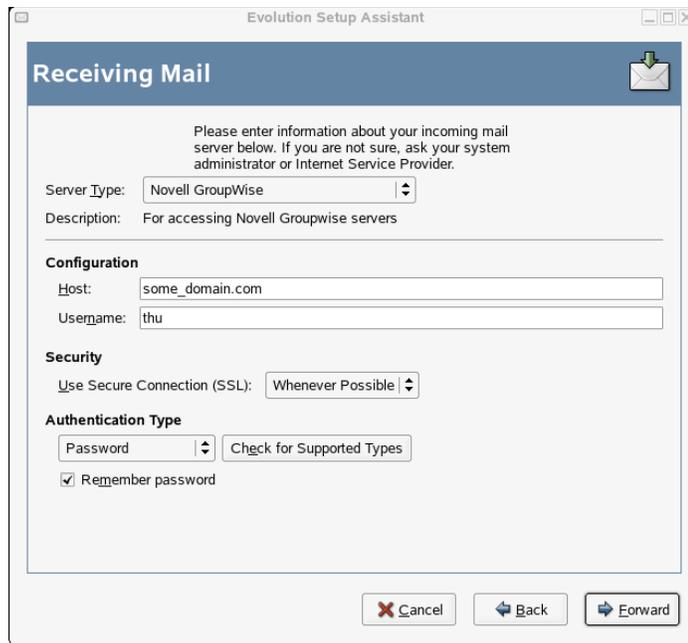
In order for users to access their GroupWise mailboxes from Evolution, they must configure Evolution to access their GroupWise accounts.

- 1 Click *Edit > Preferences*, then click *Mail Accounts*.
- 2 Click *Add*.
- 3 On the Identity page, type your e-mail address, then click *Forward*.



The screenshot shows the 'Evolution Setup Assistant' window with the 'Identity' tab selected. The window title is 'Evolution Setup Assistant'. The 'Identity' section has a blue header. Below the header, there is a message: 'Please enter your name and email address below. The "optional" fields below do not need to be filled in, unless you wish to include this information in email you send.' The form is divided into two sections: 'Required Information' and 'Optional Information'. Under 'Required Information', there are two text input fields: 'Full Name' with the value 'Tabitha Hu' and 'Email Address' with the value 'thu@Corporate.com'. Under 'Optional Information', there is a checked checkbox labeled 'Make this my default account', and two more text input fields: 'Reply-To' and 'Organization', both of which are currently empty. At the bottom of the window, there are three buttons: 'Cancel' (with a red X icon), 'Back' (with a left arrow icon), and 'Forward' (with a right arrow icon).

- 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*.

Mobile Devices

VIII

If you own a mobile device, you can synchronize it with GroupWise®. GroupWise has provided GroupWise Mobile Server for synchronizing several of the most common device. In addition, GroupWise has teamed up with BlackBerry* for synchronizing of BlackBerry devices.

- ♦ [Chapter 53, “GroupWise Mobile Server, Powered by Intellisync,” on page 461](#)
- ♦ [Chapter 54, “BlackBerry Enterprise Server,” on page 463](#)
- ♦ [Chapter 55, “GroupWise PDA Connect,” on page 465](#)

GroupWise Mobile Server, Powered by Intellisync

53

Using GroupWise® Mobile Server, you can synchronize Personal Information Manager (PIM) and e-mail data from Novell® GroupWise to Windows* CE, Windows Mobile*-based Smartphones, Symbian* OS, Palm OS* handheld devices, and SyncML* devices.

The GroupWise Mobile Server includes the following modules from Intellisync*:

- ♦ E-mail Accelerator (excluding POP3, IMAP, Exchange Connector, Lotus Notes Connector, Workgroup, and PC Monitor)
- ♦ GroupWise Connector
- ♦ Mobile device synchronization

You can download the GroupWise Mobile Server software from [Novell Downloads \(http://download.novell.com/index.jsp\)](http://download.novell.com/index.jsp).

For additional information documentation about GroupWise Mobile Server, visit the [GroupWise Mobility Web site \(http://www.novell.com/documentation/gwmobility/index.html\)](http://www.novell.com/documentation/gwmobility/index.html).

BlackBerry Enterprise Server

54

Novell® and Research In Motion (RIM) 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 their Novell GroupWise messaging application and other enterprise tools to mobile professionals. The BlackBerry Enterprise Solution provides users with mobile access to e-mail, 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 GroupWise Messenger support, enhanced support for PowerPoint and Web Doc attachments, group- and role-based administration, localized data pass-through and SMS/PIN/call log auditing.

For product information about BlackBerry Enterprise Server for Novell GroupWise, see the [BlackBerry GroupWise Product Web site \(http://www.blackberry.com/products/software/server/groupwise/quickstart.shtml\)](http://www.blackberry.com/products/software/server/groupwise/quickstart.shtml). To learn about the latest BlackBerry offers for GroupWise users, see the [GroupWise Product Web site \(http://www.novell.com/products/groupwise/rim.html\)](http://www.novell.com/products/groupwise/rim.html).

For GroupWise-specific BlackBerry documentation, see the [BlackBerry Technical Knowledge Center \(http://www.blackberry.com/knowledgecenterpublic/livelink.exe/fetch/2000/8067/645045/7963/customview.html?func=ll&objId=7963&objAction=browse&sort=name44\)](http://www.blackberry.com/knowledgecenterpublic/livelink.exe/fetch/2000/8067/645045/7963/customview.html?func=ll&objId=7963&objAction=browse&sort=name44).

For GroupWise-specific BlackBerry Knowledgebase articles, look up "GroupWise" in the [BlackBerry Technical Solutions Center \(http://www.blackberry.com/btsc/search.do?languages=&rwTarget=%2Fbtsc%2FrfPlayerWidget.do&searchMode=GuidedSearch&searchString=groupwise&cmd=search&productFamily=&contextType=gs\)](http://www.blackberry.com/btsc/search.do?languages=&rwTarget=%2Fbtsc%2FrfPlayerWidget.do&searchMode=GuidedSearch&searchString=groupwise&cmd=search&productFamily=&contextType=gs)

GroupWise PDA Connect

55

If you own a mobile device (Palm OS or Pocket PC), you can synchronize it with GroupWise® using GroupWise PDA Connect. GroupWise PDA Connect is designed to work on a Windows computer to synchronize data between GroupWise and a PDA device. It is comprised of a synchronization engine and translators, which are used for seamless integration with the data source's features and data.

- ♦ [Section 55.1, “Prerequisites,” on page 465](#)
- ♦ [Section 55.2, “Downloading the Software,” on page 465](#)
- ♦ [Section 55.3, “Third-Party Partners,” on page 465](#)

55.1 Prerequisites

Before installing GroupWise PDA Connect:

- ♦ Ensure that Microsoft ActiveSync* is installed for the Pocket PC or Palm HotSync* Manager is installed for the Palm OS.
- ♦ Ensure that you are not synchronizing anything in Outlook (such as e-mail, calendar, tasks, or notes) in either ActiveSync or PocketMirror*.

55.2 Downloading the Software

GroupWise PDA Connect 1.0 is available for download at the [Novell Product Downloads Web site \(http://download.novell.com\)](http://download.novell.com).

The download is an executable (.exe) file.

- 1 Select *GroupWise* as the product, then click *Submit Search*.
- 2 Click *GroupWise PDA Connect 1.0 SP1 Multi Lingual*.
- 3 Click the filename (*setupgwpda701.exe*), then follow the instructions to download the file.
- 4 Click *Start > Run > Browse*.
- 5 Select the *setupgwpda701.exe* file on the local or network drive.
- 6 Click *OK* to run the GroupWise PDA Connect Installation program.
- 7 Follow the on-screen instructions provided in the GroupWise PDA Connect Installation to install the software.

55.3 Third-Party Partners

Novell partners with several third-party companies for synchronization of mobile devices. A complete list of the third-party partners can be found in the [Novell Partner Product Guide \(http://www.novell.com/partnerguid\)](http://www.novell.com/partnerguid).

Documentation Updates



This section lists updates to the *Interoperability Guide* that have been made since the initial release of Novell® GroupWise® 7. 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 *Interoperability Guide* was republished. Within each dated section, the updates are listed by the section title.

The *GroupWise 7 Interoperability Guide* has been updated on the following dates:

- ♦ [Appendix A, “June 25, 2008,” on page 469](#)
- ♦ [Appendix B, “March 14, 2008 \(GroupWise 7 SP3\),” on page 471](#)
- ♦ [Appendix C, “October 24, 2007,” on page 473](#)
- ♦ [Appendix D, “April 16, 2007 \(GroupWise 7 SP 2\),” on page 475](#)
- ♦ [Appendix E, “September 29, 2006,” on page 477](#)
- ♦ [Appendix F, “August 2, 2006,” on page 479](#)
- ♦ [Appendix G, “June 15, 2006 \(GroupWise 7 SP1\),” on page 481](#)
- ♦ [Appendix H, “November 30, 2005,” on page 485](#)

June 25, 2008

A

Location	Change
Novell Teaming and Conferencing	
"Novell Teaming and Conferencing" on page 243	Updated the section to correspond with Teaming and Conferencing 1.0 Support Pack 3.
Section 25.1.3, "Adding the New Web Server Accelerator to the iChain Server Object in ConsoleOne," on page 266	Changed X-Authentication to X-Authorization.

Location	Change
Novell Cluster Services on NetWare	
Section 5.1, "Understanding the WebAccess Components," on page 89	Emphasized that if you have not clustered your Web server, you can install the WebAccess Application on a Web server that is outside the cluster where the WebAccess Agent is installed.
"Using TSAFSGW in a NetWare 6.x Cluster" on page 112	Corrected the description of the functionality of TSAFSGW in a cluster.
Microsoft Clustering Services	
"Planning GroupWise Resource Groups" on page 388	Explained that multiple Internet Agents or multiple WebAccess Agents cannot run on the same node at the same time.
PolyServe Matrix Server	PolyServe Matrix Server is no longer supported.
Mobile Devices	
"BlackBerry Enterprise Server" on page 463	Added links to BlackBerry documentation and support sites.

October 24, 2007

C

Location	Change
Novell Teaming and Conferencing	
"Novell Teaming and Conferencing" on page 243	Added a new section about integrating eDirectory™ and GroupWise® with the Novell® Teaming and Novell Teaming + Conferencing products.

April 16, 2007 (GroupWise 7 SP 2)

D

Location	Change
Novell Cluster Services on NetWare	
“Copying LDAP and QuickFinder Files to the Messenger Volume” on page 125	Improved the instructions for handling LDAP and QuickFinder files after installing the Messenger agents in a cluster.
Heartbeat on Linux	
Part II, “Novell Cluster Services on Linux,” on page 129	Added a new section about setting up a GroupWise system in a Heartbeat cluster on Linux.

September 29, 2006

E

Location	Change
Novell Cluster Services on NetWare	
Section 3.5.2, "Editing Clustered Agent Startup Files," on page 52	Mentioned use of the /user and /password startup switches for the MTA.

Location	Change
Novell Cluster Services on Linux	
Section 14.4.3, “Configuring GroupWise Cluster Resources to Load and Unload the Agents,” on page 152	Added new examples of load and unload scripts for use with Linux traditional file system disk partitions.
“Configuring the Internet Agent Cluster Resource to Load and Unload the Internet Agent and Its MTA” on page 176	Added new examples of load and unload scripts for use with Linux traditional file system disk partitions.
“Configuring the WebAccess Agent Cluster Resource to Load and Unload the WebAccess Agent and Its MTA” on page 196	Added new examples of load and unload scripts for use with Linux traditional file system disk partitions.
Section 17.3.2, “Configuring the Monitor Agent Cluster Resource to Load and Unload the Monitor Agent,” on page 213	Added new examples of load and unload scripts for use with Linux traditional file system disk partitions.
Section 22.2.3, “Configuring the Messenger Cluster Resource to Load and Unload the Messenger Agents,” on page 235	Added new examples of load and unload scripts for use with Linux traditional file system disk partitions.

June 15, 2006 (GroupWise 7 SP1)



Location	Change
Novell Cluster Services on NetWare	
Section 2.8.4, "Deciding Whether to Run the Agents in Protected Memory," on page 36	Strongly recommended running the agents in protected memory, with each agent in a separate memory space.
Section 4.1.7, "Deciding Whether to Run the Internet Agent and Its MTA in Protected Memory," on page 72	Strongly recommended running the Internet Agent as well as its MTA in protected memory, with each agent in a separate memory space.
"Implementing WebAccess in a NetWare Cluster" on page 89	Removed reference to Netscape Enterprise Server. No instructions for clustering any Web server are provided in the GroupWise documentation.
Chapter 8, "Backing Up a GroupWise System in a NetWare Cluster," on page 111	Described the new TSAFSGW functionality associated with the /vserver switch.
Chapter 9, "Updating a GroupWise System in a NetWare Cluster," on page 115	Added a new section about updating GroupWise software in a NetWare cluster.
"Configuring the Messenger Volume Resource to Load and Unload the Messenger Agents" on page 124	Added a command to start both Messenger agents with a single command.
Novell Cluster Services on Linux	
Section 13.2, "Planning a Clustered Software Distribution Directory," on page 135	Explained the need for a complete software distribution directory in order to use the new <i>Configure GroupWise for Clustering</i> option of the GroupWise Installation program.
Section 13.6.1, "Recording Secondary IP Addresses for the Agents," on page 138	Explained that the <i>Configure GroupWise for Clustering</i> option of the GroupWise Installation program sets the --ip switch for you automatically.
Section 13.6.3, "Determining Cluster Resource Information for the Linux Agents," on page 139	Explained the new fields on the Domains / Post Offices page of the <i>Configuration</i> option in the GroupWise Installation program.

Location	Change
Section 14.1, "Setting Up a New GroupWise System in a Linux Cluster," on page 143	Explained that you should copy all agent software into the software distribution directory, and that you should exit the GroupWise Installation program without installing the agent software from the CD.
"Running the GroupWise Installation Program on the Preferred Node" on page 147	Explained that you install the agent software from the software distribution directory, not the CD; explained the new cluster resource options on the Domains / Post Offices page.
"Running the GroupWise Installation Program on Subsequent Nodes" on page 148	Explained the new <i>Import Clustering Data</i> option.
"Creating New Startup Files" on page 158	Explained that, in a cluster, startup files are located in <code>mount_point/groupwise/agents/share</code> on the shared resource, rather than in the standard location on each cluster node.
Section 15.1.4, "Determining Cluster Resource Information for the Internet Agent," on page 169	Explained the new field on the Server Information page of the <i>Configuration</i> option in the GroupWise Installation program.
"Running the Internet Agent Installation Program on the Preferred Node" on page 172	Explained that you install the Internet Agent software from the software distribution directory, not the CD; explained the new cluster resource options on the Server Information page.
"Running the Internet Agent Installation Program on Subsequent Nodes" on page 173	Explained the new <i>Import Clustering Data</i> option.
Chapter 16, "Implementing WebAccess in a Linux Cluster," on page 187	Added a new section for WebAccess, because the WebAccess Agent can now be installed in a cluster.
Section 17.2.5, "Determining Cluster Resource Information for the Monitor Agent," on page 209	Explained the new field on the System Options page of the <i>Configuration</i> option in the GroupWise Installation program.
"Running the Monitor Installation Program on the Preferred Node" on page 210	Explained that you install the Monitor Agent software from the software distribution directory, not the CD; explained the new cluster resource options on the System Options page.
"Running the Monitor Agent Installation Program on Subsequent Nodes" on page 211	Explained that you do not need to use the <i>Import Clustering Data</i> option to configure the Monitor Agent in a cluster; you just need to use the <i>Install</i> option on each subsequent node.
Chapter 19, "Updating a GroupWise System in a Linux Cluster," on page 223	Added a new section about updating GroupWise software in a Linux cluster.

Location	Change
Section 22.1.4, "Determining Cluster Resource Information for the Messenger Agents," on page 230	Explained the mount point for shared storage.
"Running the Messenger Installation Program on the Preferred Node" on page 231	Explained the new mount point prompt when installing the Messenger agents in a cluster.
"Running the Messenger Installation Program on Subsequent Nodes" on page 232	Explained that the Messenger Installation program gathers agent configuration information from the shared resource to create defaults for configuring the Messenger agents on subsequent nodes.
PolyServe Matrix Server	PolyServe Matrix Server is now supported.
Non-GroupWise Clients	
"Non-GroupWise Clients" on page 449	Added an example of a possible limitation of a POP, IMAP, or SOAP e-mail client.

November 30, 2005

H

Location	Change
Novell exteNd	
Chapter 28, "Novell exteNd," on page 277	Added information about Novell® exteNd™ compatibility with GroupWise® WebAccess and provided a link to the exteNd documentation that explains how to configure exteNd for use with GroupWise.
Entire Guide	Page design reformatted to comply with revised Novell documentation standards.
